Small Samples No More:
Probing the Evolution of Massive Stars

Trevor Z. Dorn-Wallenstein

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2021

Reading Committee:
Emily M. Levesque, Chair
Julianne Dalcanton
James Davenport

Program Authorized to Offer Degree:
Astronomy

University of Washington

Abstract

Small Samples No More:
Probing the Evolution of Massive Stars

Trevor Z. Dorn-Wallenstein

Chair of the Supervisory Committee:
Professor Emily M. Levesque
Astronomy

Evolved massive stars — the post-main sequence descendants of stars with initial masses higher than roughly 8 solar masses — are rare yet critically important objects. As residents of their host galaxies, they inject radiation and matter into their surroundings on short timescales before exploding as supernovae. Individually, they are fascinating astrophysical laboratories in which many of the unknowns of stellar evolution coalesce. Due to their rarity, these multitudinous unknowns remain under-constrained. In this work, I attempt to understand evolved massive stars using a variety of techniques that have only recently begun to be applied to these interesting objects. My studies of populations of young stars reveal that the massive star binary fraction can be inferred using only simple demographic statistics. However, these methods can only be used given large numbers of well-classified stars, and I show that even using advanced machine learning techniques, existing data are insufficient to classify these stars. Finally, I demonstrate the immense potential of using asteroseismology to probe the interiors of evolved massive stars, and discover a new class of pulsating supergiant.

TABLE OF CONTENTS

Page

LIST OF FIGURES

LIST OF TABLES

Table Number Page

ACKNOWLEDGMENTS

In this thesis, I'm going to use the word "I" a lot to describe various ways in which various experiments and projects were conducted in order to learn things about the Universe. This is, of course, the preferred style of this dissertation, but it paints a very incorrect picture of the work that was done. So often we picture science as one person pondering the universe, a person that is often represented in our media as white, and male. But science is *not* done by individuals. Some of the biggest discoveries that took place during my time as a graduate student at the University of Washington — including the first gravitational wave observations, the first ever image of the event horizon of a black hole, and countless exciting discoveries in planetary science — were published in papers with dozens, hundreds, or even thousands of coauthors. We stand on the shoulders of giants, yes, but we also stand shoulder to shoulder with our peers and colleagues, leaning and being leaned on.

And so whenever you read "I" in this work, know that it is not just me. It is also my coauthors and collaborators, my support network, my friends, my family. Now is the time when I have the great honor of naming them, and thanking them. Professor Emily Levesque, Professor James Davenport, Dr. Kathryn Neugent, Dr. Daniela Huppenkothen, Dr. Brett Morris, Dr. John Ruan, and Keyan Gootkin have all appeared next to me on papers I have published, and their ideas, advice, inspiration, and sheer genius have been invaluable to me. I've also had the privilege of collaborating with Professor Philip — or is it Fred? — Massey, and Ishan Coutinho. Along the way I have been helped along by my thesis committee: Professor Emily Levesque, Professor James Davenport, Professor Julianne Dalcanton, and my Grad School Representative, Professor Cecilia Aragon. I thank them not only for their wisdom and guidance, but for giving me what may be the longest General Exam topic ever: Physical Mechanisms and Uncertainties for Nucleosynthetic

Yields from Massive Stars as a Function of Mass, Metallicity, and Evolutionary Phase. What a mouthful! And of course, before I knew any of these people, my fledgling research career was aided by my mentors at Wesleyan University, Professor Roy Kilgard and Professor Ed Moran, and my REU supervisor at Colgate University, Professor Jeff Bary.

In addition to my collaborators and mentors, I have been supported by a group of truly wonderful humans collectively known as the Astronomy graduate students at the University of Washington. Though many of them have long since graduated, each grad has left a mark on me, whether through trauma bonding while studying for the qualifying exam, collaboration on our coursework, troubleshooting research problems, quiet conversations held after a long w(h)ine time, marriages, celebrations, disasters, dinners, walks, hikes, adventures, or just a short reprieve of companionship in one of our offices. Thank you all. And naturally, I need to shout out Seattle's best (and only) all-astronomer rock band, Night Lunch. What started as a way to blow off steam and practice our instruments evolved into a full-fledged *thing*, and we were awesome. We played out, we recorded a freakin' album (no seriously, go on Spotify and search for Am I Falling), and had the time of our lives. Nicole, Brett, Spencer: y'all are amazing. Simple as that.

Of course, I need to mention my family. I grew up in quite a chaotic home, but one full of love and support, a fact for which I continue to grow increasingly grateful. Michael, Eileen, Megan, Aidan, thank you all. More recently, I have been building a family of my own. For the past 4-ish years, I have had the honor of sharing my life with my partner — now fiancée! — soon-to-be-Dr. Natalie Nicole Sanchez: rock goddess, fashionista, and galaxy-simulator extraordinaire. No leave, only stay? And lastly, our dog, Poderosa "I Can Sleep Anywhere I Want, Just Watch Me" "Sock Thief" Don-Wallanchez: the most disrespectful, grumpy, cheese-obsessed long haired miniature dachshund ever to grace this earth. Who knew one tiny dog could demand so many snuggles?

Most of observational astronomy done in the United States relies on facilities, Universities, and observatories that are situated on land formerly and currently occupied by Indigenous peoples, and are the legacy of violent settler colonialism. For example, the University of Washington and the areas of Seattle in which I have lived and worked are on the traditional, and unceded land of the first people of Seattle, the Duwamish People (Dkhw Duw'Absh) past and present. The Duwamish are currently not recognized by the United States government, and promises made to the Duwamish 150 years ago in the Treaty of Point Elliot have not been kept. I honor with gratitude the land itself and the Duwamish Tribe. For more information, please see duwamishtribe.org, and realrentduwamish.org Additionally, throughout my time at UW I have had the privilege of using the Apache Point Observatory 3.5-meter telescope, which is on the stolen land of the Mescalero Apache tribe (mescaleroapachetribe.com), as well as the Gemini North telescope at Mauna Kea Observatory which is on the stolen land of the Kānaka ʻŌiwi people (see Kahanamoku et al. [129]). I stand in solidarity with the Pu'uhonua o Pu'uhuluhulu Maunakea in their effort to preserve this sacred space for native Hawaiʻians.

Finally, thanks to you for reading this! It's quite a long work, so I understand if you need to take your time, or only read it in bits and pieces. I won't be offended. I hope you enjoy!

— Trevor Z. Dorn-Wallenstein, 6/4/2021

DEDICATION

to Noodle and Pod

Chapter 1

INTRODUCTION

1.1 What is a star?

No, really, what is a star? Throughout the text that follows, I'm going to be devoting quite a lot of attention to stars of varying temperatures, sizes, compositions, masses, even shapes! So we should take some time here and now to make sure we're on the same page.[1] So... what is a star?

Well, it's complicated. To start with, let's turn to pop culture, in particular, the band They Might Be Giants. In addition to contemplating the nomenclature of ancient cities, They Might Be Giants devoted some attention to stellar astrophysics in their cover of Tom Glazer's song *Why Does the Sun Shine*. In it, they claim that

> *[a star] is a mass of incandescent gas, a gigantic nuclear furnace.*

This seems as good a place to start as any. Going off of this line, we can conclude:

- Stars are massive: we measure the masses of stars relative to the mass of the sun, which we define as $1M_\odot = 1.989 \times 10^{30}$ kg — about a nonillion times the amount of ice cream I would be comfortable eating in one sitting.

- Stars are incandescent; they glow. Specifically, the sun glows with an intensity of $1L_\odot = 3.828 \times 10^{26}$ W — about ten septillion times the brightness of a typical lightbulb that illuminates an ice cream store.[2]

---

[1] For reference, we should both be on page one at the moment.

[2] You can tell where my mind is at as I write this.

- Without being too pedantic about the difference between gas and plasma, stars are indeed made of gas. As Dr. Cecilia H. Payne demonstrated in her Ph.D. thesis, that gas is predominantly hydrogen [228]. We measure the composition of stars as fractions by mass of their components. For example, at the surface of the sun today, the mass fraction of Hydrogen is $X_{surf,\odot} = 0.7438 \pm 0.0054$, the mass fraction of Helium is $Y_{surf,\odot} = 0.2423 \pm 0.0054$, and the mass fraction of... well, everything else (including ice cream) that astronomers insist on lumping together and calling "metals" is $Z_{surf,\odot} = 0.0139 \pm 0.0006$,[3] according to our most recent measurements [12].

- Stars are gigantic, with the radius of the sun being $1R_\odot = 6.9634 \times 10^5$ km — far longer than I would be willing to walk for ice cream.

- Stars are nuclear furnaces; that is, their incandescence is powered by nuclear fusion in their cores. And while a cubic meter of the solar core produces approximately 35 times less energy per second than I would gain by eating ice cream at a rate of one standard 68-gram scoop per minute [272], the sheer volume of the sun (remember, it's gigantic!) means that this process is capable of keeping the sun going for billions of years! Cool, huh?

This definition of a star definitely works, and describes everything you really need to know... until it breaks down. For example, a supernova explosion satisfies all of these properties. Is a supernova a star? I would argue that it isn't! Meanwhile, there are stars (or stellar-like objects) that don't obey all of these laws. Neither white dwarfs nor neutron stars are supported by nuclear fusion, nor are they entirely gaseous.[4] Finally, by assigning a single measurement of a star's size (its radius) we are implicitly assuming that stars are spherically symmetric, when in fact this

---

[3] Z is also called the "metallicity," a term which can also refer to the logarithm of the ratio of the number of iron atoms to the number of hydrogen atoms measured at the stellar surface, scaled to the sun ($[Fe/H] = \log(N(\mathrm{Fe})/N(\mathrm{H})) - \log(N(\mathrm{Fe})/N(\mathrm{H}))_\odot$) *or* the logarithm of the ratio of the number of oxygen atoms to the number of Hydrogen atoms measured in the gas of hot HII regions plus twelve for some reason ($\log(\mathrm{O/H}) + 12$). It's complicated.

[4] White dwarfs crystallize as they cool, while neutron stars have a solid crust and an interior that is best described as "weird."

couldn't be further from the truth! Stars rotate, and centrifugal force distorts the shape of the stellar surface [316]. Stars are also frequently born into binary systems, and tidal forces can affect the shape of the star as well [120]. Many stars can also pulsate non-radially [5] or shed mass asymmetrically [e.g. 261, 96], resulting in further deviations from this simple, one-dimensional picture of a star as a sphere.

For now, let's go with this definition: a star is an object whose behavior is determined by a number of effects, including, but not limited to gravity, nuclear physics, thermodynamics, hydrodynamics, quantum mechanics, and relativity. For the most part, we can approximate their structure using a set of four equations, collectively called the *equations of stellar structure*. These equations are by no means complete, and assume spherical symmetry, but they give us a good idea of the physics involved. Most importantly, they can be modified to incorporate all of the effects that aren't included at first glance. First, we require the star to be in *hydrostatic equillibrium*: for a given layer of the star at a radius r, we require that the downward force of gravity is matched by the difference in pressure just below and just above r. In symbols:

$$\frac{dP}{dr} = -\frac{GM_r\rho}{r^2} \tag{1.1}$$

where P is the pressure, M_r is the mass contained within a sphere of radius r, ρ is the density, and G is the gravitational constant. We now require *mass continuity*: we can write how M_r changes across a small layer of the star given the density and radius:

$$\frac{dM_r}{dr} = 4\pi r^2\rho \tag{1.2}$$

Both of these formulae are in *Eulerian* coordinates, where we are interested in the structure as a function of a spatial coordinate (in this case, r). While we are perhaps more used to working with spatial coordinates (whether or not they are more intuitive), it can be more useful to track the structure of the star as a function of the mass coordinate, M_r (*Lagrangian* coordinates). Given the wide range of densities found in stars, physical quantities like the pressure can change rapidly over small spatial distances, or very little over the majority of the stellar radius. Recasting our equations into Lagrangian coordinates can help avoid these high gradients, allowing us to study the structure using a variable that scales nicely with the physics. Figure 1.1 shows two quantities computed from a stellar model:[5] the (log of the square of the) Brunt-Väisälä and Lamb frequencies,[6] plotted as a function of the radial coordinate (Eulerian coordinates, top), and as a function of the mass coordinate (Lagrangian coordinates, bottom). The model itself is quite interesting: from the top panel we can see the star's radius is approximately 400 $R_\odot$, while from the bottom, we can see its mass is about 16 $M_\odot$: a very different star than our sun! For the majority of this star's *spatial* extent, however, very little in the way of anything interesting is happening in these two quantities. But by examining the bottom panel, we see some neat behavior through the majority of the star's *mass*. Nevermind what the neat behavior actually is; this example simply serves to illustrate that unless the math is made significantly more complex, Lagrangian coordinates are the way to go.

Using Equation (1.2), we can now rewrite the equation of hydrostatic equilibrium in Lagrangian coordinates:

$$\frac{dP}{dM_r} = -\frac{GM_r}{4\pi r^4} \tag{1.3}$$

Continuing onward, we can track the generation of energy through the star using the *energy equation*:

$$\frac{dL_r}{dM_r} = \epsilon - \epsilon_\nu \tag{1.4}$$

where L_r can be thought of as the luminosity measured at a radius r, ϵ is the rate of energy generation per unit mass, and ϵ_ν is the rate of energy loss to neutrinos per unit mass. Outside of the stellar core, the right hand side of (1.4) goes to 0, and the luminosity is constant through the stellar structure. Finally, we can describe how energy moves through the star with the *thermal transport equation*. In *radiative equilibrium*, heat is predominantly carried through the star by radiation.

---

[5] I'll talk about how this model was created in Chapter 6. For now, just know it's a spherically symmetric model of a nonrotating 23 $M_\odot$ star just after it has lost its *envelope*, or the inert — a.k.a. not doing nuclear fusion — layers of material that surround the *core* of the star where nuclear fusion happens.

[6] Don't worry about what these mean right now. At the moment, just know that they're interesting when the lines are wiggly or spiky or nonexistent.

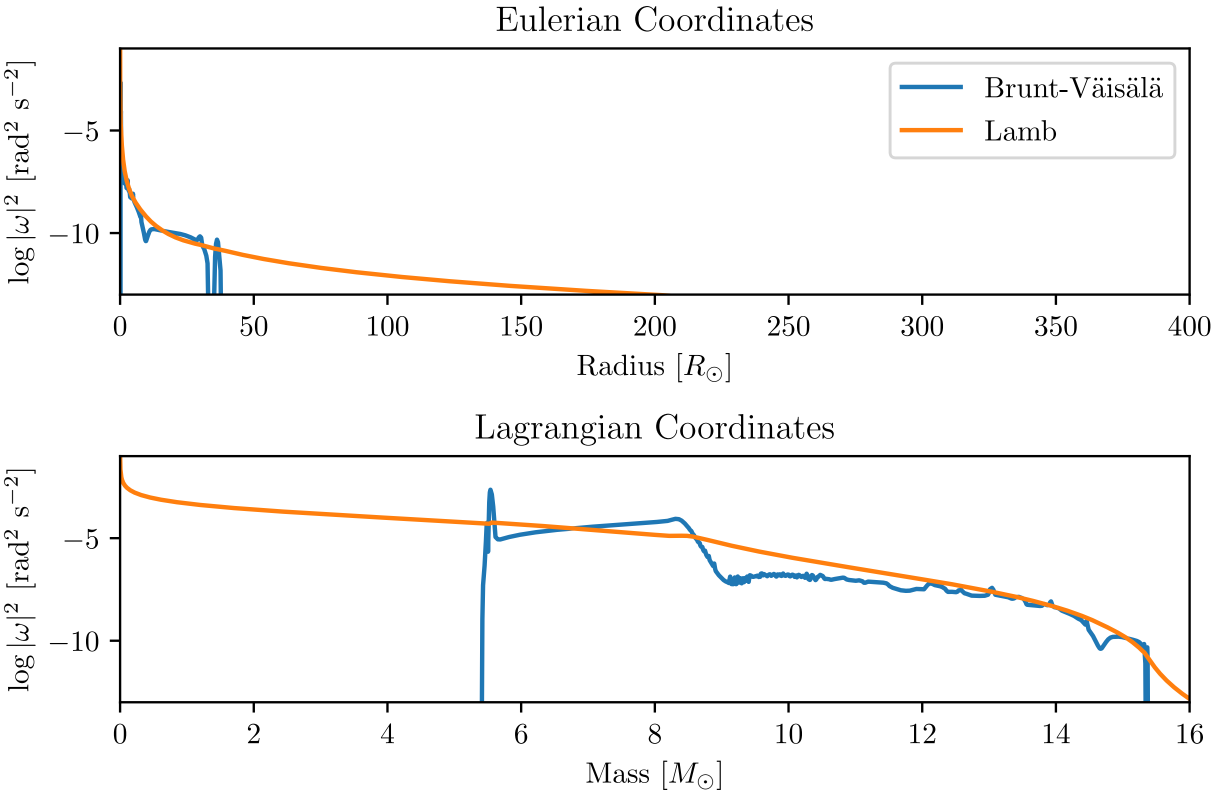

Figure 1.1 Stellar model showing the log of the square of the Brunt-Väisälä and Lamb frequencies in Eulerian coordinates (top), and Lagrangian coordinates (bottom). While the same exact model is plotted, the Lagrangian coordinates allow us to see much of the interesting behavior in this star.

This leads to a temperature gradient given by

$$\frac{dT}{dM_r} = -\frac{3\kappa L_r}{64\pi^2 acr^4 T^4} \tag{1.5}$$

where T is the temperature, κ is the opacity of the layer, a is the "radiative constant", and c is the speed of light. In regions where luminosity is high like the cores of massive stars, or where opacity is high and temperature is low as in the envelopes of cool stars, radiative energy transport is inefficient and *convective energy transport* can occur instead. While convection is a hydrodynamical process requiring careful treatment, approximations like the commonly-used "mixing-length theory" [25] can help us approximate the temperature gradient in these regions.

And that's it! Take your favorite differential equation solver, adopt some appropriate boundary conditions, and have at it! Then watch as your model star... does nothing. Absolutely nothing. No variable in the above equations is explicitly dependent on time, so the model is entirely static. That's good, isn't it? After all, the sun is more or less the same from day-to-day.[7] But stars — and really everything else in the Universe — aren't particularly interesting if they don't change. Thankfully, we know that stars *do* change. Most spectacularly, some stars explode, which is quite a change indeed! But equations (1.2)-(1.5) aren't explicitly time-dependent. So what gives? Why do stars change?

1.2 Why do stars evolve?

Before I answer this question, I should probably be clear about what I mean. There are many stellar phenomena that change on human timescales. For example, the equator of the sun rotates once every 24.47 days (the poles rotate somewhat slower, once every ∼38 days). Sunspots follow the solar rotation, appearing and disappearing from view. The sunspots themselves are formed and disappear on timescales of a few days to a few months. The frequency at which sunspots appear on the sun also changes, with a period of about 11 years, over which time the solar magnetic field

---

[7] Unless you happen to live in Seattle, where occasionally the sun disappears entirely, sometimes for months at a time. Philosophers and astrophysicists have long pondered this phenomenon, and wondered why anyone would choose to live in a place where it happens. They remain baffled, but here we are.

drastically realigns itself. But none of these phenomena are reflective of fundamental changes in the sun; it remains a 1 $R_\odot$, 1 $M_\odot$, 1 $L_\odot$ star. The kind of changes that I am interested in reflect changes to the stellar structure itself; what I will call *stellar evolution*.

1.2.1 Evolutionary Timescales

Stellar evolution can be caused by a number of different processes which operate on multiple timescales, some nearly instantaneous, and some longer than the age of the Universe. By looking at the physics of stars however, we can identify a few characteristic timescales. The first is what we call the *dynamical timescale* — also called the free-fall timescale. Imagine we took a star, and turned off everything but gravity. With no pressure, no thermal effects, no nuclear fusion, the star would start to collapse under its own gravity in a time proportional to its mean density, $\bar{\rho}$:

$$\tau_{\mathrm{Dyn}} \propto (G\bar{\rho})^{-1/2} \tag{1.6}$$

The dynamical timescale is usually quite fast (only about 30 minutes in the sun), and is only relevant for a small number of brief evolutionary phases, including the very beginnings (star formation) and utter ends (core collapse) of massive stars!

Turning on the rest of physics except nuclear fusion (for now), if a star contracts under its gravity, it loses gravitational potential energy, of which approximately half is converted into heat, which increases the pressure to stave off further collapse. According to the virial theorem, the remaining half is radiated away as light. The timescale on which a star's gravtational energy can be completely converted to radiation is the *thermal timescale*:

$$\tau_{\mathrm{Therm}} = \frac{GM^2}{RL} \tag{1.7}$$

where M is the star's mass, R is its radius, and L is its luminosity. Before the realization that nuclear fusion powered stars, it was thought that this was the mechanism through which the sun shined. However, while the thermal timescale is significantly longer than the dynamical timescale, it's still astrophysically quite short; the sun's thermal timescale is about 30 Myr, while our observations indicate that the Earth is older by a few billion years! Clearly, the sun does not evolve on the thermal timescale. That said, there are a few instances where the thermal timescale is relevant, mostly for protostars, and during the subgiant phase.[8]

The final timescale is the *nuclear timescale*. We now know that gravitational collapse in stars is held at bay not just by thermal processes, but by the generation of energy in their cores driven by nuclear fusion. Thus the nuclear timescale is the timescale on which all of the material available for nuclear fusion can be burned at the current rate. Neglecting a few factors of order unity, this makes the nuclear timescale roughly equal to

$$\tau_{Nuc} = \frac{Mc^2}{L} \tag{1.8}$$

which, for the sun, is approximately 10 Gyr. Now that's more like it! Stars like our sun evolve on Gyr timescales, and analyzing these approximate timescales was the necessary step to realizing why: nuclear fusion!

1.2.2 How do stars evolve?

How is nuclear fusion responsible for stellar evolution? Looking at the equations of stellar structure (1.2)-(1.5), there's only one variable in there that explicitly relates to nuclear fusion, and that's the energy generation, ϵ (and, to a lesser extent, the energy lost to neutrinos, ϵ_ν, but I'm going to ignore neutrinos for the most part). Sun-like stars fuse hydrogen into helium via the *proton-proton chain* (or "p-p chain"), which looks something like this:

$$4p + 2e^- \rightarrow^4 He + 2\nu_e + \gamma \tag{1.9}$$

---

[8] Things also get interesting when the thermal timescale gets to be comparable to the dynamical timescale (usually when the ratio L/M is approximately 10^3), but we'll save that for Chapter 6.

In words: four protons (p, a.k.a., ionized hydrogen) and two electrons (e^-) are converted into a nucleus of Helium-4 (^4He), producing a couple of electron neutrinos (ν_e, that immediately zip away, leeching some of the energy away from this reaction), and γ-rays.[9] In addition to creating energy in the form of γ-ray photons that are immediately absorbed into their surroundings, this reaction subtly changes the overall composition of the part of the star where this reaction took place. Now, there's slightly less hydrogen, and a little bit more helium. This makes it ever so slightly harder for the first reaction in the p-p chain (the formation of deuterium from two hydrogen ions) to occur. To compensate, the star's core contracts in order to increase the density, making fusion reactions more likely. In response, the entire star expands slightly, and increases its overall luminosity; nuclear fusion drives stellar evolution!

Of course there's another mechanism through which fusion leaves its mark on the equations of stellar structure. Analyzing (1.5), we see that the temperature gradient depends on the opacity, κ. Opacity calculations are difficult, requiring techniques well-beyond the scale of this introduction. However, suffice it to say that the opacity is strongly dependent on the chemical composition. In the core of the sun, electrons are the key sources of opacity. In a medium of pure hydrogen, there's roughly one electron per atomic nucleus floating around. However, once all of the hydrogen has been converted into helium, there are now four electrons per atomic nucleus. This significantly affects the opacity, changing the temperature gradient, and altering the stellar structure. While I've focused for now on the p-p chain, the other forms of nuclear fusion further impact the composition of the star and thus the opacity. From here, phenomena like convection are capable of mixing the material that has been processed by nuclear fusion throughout the stellar structure. These compositional changes impact how energy is transported through the star, and affect the stellar evolution. Returning now to our imaginary stellar model that we created at the end of the previous section, we now know how to make it evolve: we take our static model, estimate all of the compositional changes due to nuclear burning, convection, and other mixing processes, recalculate the stellar model after implementing these changes, rinse, repeat, and watch the star evolve.

---

[9] In reality, the p-p chain is shorthand for multiple branching reaction pathways, with some limiting steps and low-probability quantum mechanical funtimes. (1.9) is just the overall result of these reactions.

Doing this by hand is a total pain, but an enormous number of stellar evolution codes exist [259], allowing us to generate theoretical models of how stars evolve. In general, a stellar astrophysicist takes (or makes) stellar evolution models that incorporate different physics, or tunes the knobs of the physics that are already in the models, and compares the models to the data in an attempt to make the models and data agree. Sometimes[10] they don't, implying that we either need better models, better data, or both (usually both). For a plethora of reasons, the most disagreement is found when we study massive stars — stars with initial masses of about $8M_\odot$ or above.[11] Coincidentally, these stars are extremely important, both astrophysically speaking, and on a very human level. While I've focused on the sun thus far to keep things as relevant to us humans as possible, it's now time to turn to these stellar behemoths.

1.3 Massive Stars & Their Attendant Uncertainties

In all honesty, we (the massive stars community) overuse our metaphors. Massive stars are magnifying glasses. Massive stars are lighthouses. Go to any massive stars conference[12], and you'll hear these phrases thrown around with reckless abandon. And for good reason: these metaphors tend to be quite apt. Perhaps the best overused metaphor to use here is one that adorns the first paragraphs of many a scientific paper. ***Massive stars are cosmic engines.*** It's a particularly good metaphor because it calls to mind a few mental images: heat and pressure, physical reactions (including explosions), and power!

Massive stars certainly do output power, in the (astro)physical sense that they give of a tremendous amount of energy every second in the form of light! On the main sequence, we observe massive stars as O and early B type stars, which emit immense amounts of ultraviolet radiation. This radiation is energetic enough to heat and ionize nearby gas. Incidentally, this is bad news for any nearby star systems that are attempting to form planets, as UV radiation is capable of disinte-

---

[10] Most of the time

[11] I'll explain where this threshold comes from in the next section, but for now, you'll have to take it on faith.

[12] or, as in my case, the 2015 Annie Jump Cannon prize talk at the AAS meeting given by a certain member of my thesis committee who has had a tremendous impact on my career

grating protoplanetary disks. As a result, any attempt to ponder which locations in our Universe are best suited for life must take into account the existence of massive stars.

Like engines, massive stars are also the sites of critical transformative processes, in this case nucleosynthesis. Unlike the sun's p-p chain, massive stars (in fact all stars more massive than about $1.3M_{\odot}$) have hot enough cores that the far-more-efficient CNO cycle takes priority. Just like the p-p chain, the CNO cycle has multiple branches, each consisting of multiple stages that are more- or less-favorable from a quantum mechanical perspective. Regardless, the net process results in either

$$4p + 2e^{-} \rightarrow^{4} He + 2e^{+} + 2e^{-} + 2\nu_{e} + 3\gamma \tag{1.10}$$

or

$$4p + 2e^{-} \rightarrow^{4} He + 2\nu_{e} + 3\gamma \tag{1.11}$$

The CNO cycle is still hydrogen fusion; four hydrogen nuclei still go in, one helium nucleus still comes out, and we get similar biproducts plus or minus a lepton or two. But along the way are various intermediary steps that use carbon (C), nitrogen (N), and oxygen (O) [13] as catalysts, transforming one into the other in a neverending cycle that relentlessly churns out helium and energy! One of the rate-limiting steps is one that uses up N-14 and produces O-15. It turns out that this particular reaction requires a very low-probability event to occur, resulting in a buildup of nitrogen. So while we wouldn't call this process "nitrogen fusion," hydrogen fusion in massive stars does produce a large amount of nitrogen[14]. Over the course of this evolution, the star begins as a main sequence O or B star, evolving from the "zero age main sequence" (ZAMS, defined as when hydrogen burning first kicks off) to the "terminal age main sequence" (TAMS, defined as when the core runs out of hydrogen).

After a star fuses all of its hydrogen, it keeps going! First, it fuses helium into carbon by the triple-α process (so named due to its use of α particles, a.k.a. helium nuclei). The first step actually

---

[13] hence the CNO cycle. Fluorine (F) plays a role as well, but does not take a part in the dominant reaction pathways that occur.

[14] which, by the way, makes up about 78.09% of our atmosphere, no big deal.

absorbs energy in order to fuse helium into unstable beryllium:

$$^{4}He +^{4} He \rightarrow^{8} Be \tag{1.12}$$

and in the infinitesimal fraction of a second before the beryllium decays back into helium,[15] a third helium nucleus enters the fray and makes carbon:

$$^{4}He +^{8} Be \rightarrow^{12} C + 2\gamma \tag{1.13}$$

Now for a low mass star, this is generally where the story ends.[16] But as we go to higher and higher masses, stellar cores are capable of getting hotter and hotter, and more and more advanced stages of nuclear burning can be reached: carbon into oxygen, oxygen into neon, neon into magnesium, magnesium into silicon, and, via an elaborate chain of reactions, all the way into nickel-56. It turns out that 56 is the highest number of nucleons you can put into a nucleus and still get energy out via nuclear fusion (or fission, incidentally). A very small amount of energy still comes from the radioactive decay of nickel into cobalt and then iron, but at the end of the day, massive stars wind up with an inert core of Fe-56, surrounded by an onion-like structure of shells, each burning a different kind of fuel. Now, this is a catastrophic situation for a star, but we're going to pause the story here. Note that massive stars fuse many of the elements that are necessary for life to exist here on earth. As they do, their cores shrink, and their envelopes expand and cool, resulting in the star evolving from the TAMS to become a red supergiant (RSG).

Somehow this freshly nucleosynthesized matter needs to make its way from the cores of massive stars into the interstellar medium (ISM) from whence we came. After all, we are carbon-based lifeforms relying on oxygen in nitrogen-rich atmosphere, so how did that stuff get here? Through some mechanism, the star needs to shed matter. Now a low-mass star like our sun loses mass via a fast but very low-density coronal wind, resulting in an incredibly small mass loss rate of

---

[15] The half life of beryllium-8 is only 8.17×10^{-17} seconds.

[16] Some of the most massive low mass stars are capable of making oxygen or neon, but such objects are relatively rare.

$\sim 10^{-14} M_{\odot}$ yr^{-1}. By contrast, the typical mass loss rates seen in massive stars are orders of magnitude higher, and significantly contribute to the evolution of massive stars. Remember our equations of stellar structure: they all depend explicitly on the mass of the star. By changing the mass of the star, the boundary conditions we use to solve these equations are altered, and the evolutionary trajectory from hot main sequence star to cool red supergiant can be interrupted by the star losing a significant amount of its envelope. The resulting exposed stellar core then evolves to become a hot and blue object with a strong stellar wind called a Wolf-Rayet star or WR.

How does this mass loss happen for massive stars? Main sequence OB stars as well as Wolf-Rayet stars emit enormous amounts of ultraviolet radiation, imparting a significant amount of radiation pressure onto metals in their atmospheres via resonant UV absorption lines, resulting in mass loss rates from 10^{-8} to $10^{-5} M_{\odot}$ yr^{-1}. Cool supergiants are also capable of driving strong stellar winds via radiation pressure on circumstellar dust grains, with mass loss rates ranging from 10^{-6} to $10^{-4} M_{\odot}$ yr^{-1}. Finally, massive stars can experience various outbursts during their lifetimes. These events, driven by either intense radiation pressure or pulsations, can even result in the star shedding up to a few solar masses per year during so-called "giant eruptions." All of this material has been enriched by nucleosynthetic products thanks to various internal processes in the star,[17] and so massive stars pollute the ISM over the majority of their lifetimes from the main sequence, through supergiant phases, all the way up until the development of an iron core.

It is here, once the star's nuclear fuel is utterly spent and the core is an inert lump of iron, that we arrive at the end of a massive star's life. With no nuclear fusion, the core is only supported by electron degeneracy pressure that comes as a result of the Pauli exclusion principle. And degeneracy pressure isn't all that powerful. Eventually the star has a massive enough iron core that electron degeneracy can't save it. The core collapses under its own weight and the rest of the star falls in after it. Once the core has shrunk to roughly the density of an atom, various nuclear forces as well as neutron degeneracy pressure rapidly halt its collapse. The star's envelope collides with the now-incredibly-hard proto-neutron star core, and bounces, exploding outward in a violent explo-

---

[17] More on such mixing processes to come...

sion called a supernova.[18] Any matter not blown off by earlier stellar winds is now ejected into the ISM in spectacular fashion. All that's left behind is a neutron star, an incredibly dense[19] city-size hunk of quantum mechanical weirdness. Finally, if the remnant is more massive than about 3 $M_\odot$, not even neutron degeneracy pressure can save it, and it collapses into a black hole.

And that's the entire life cycle of a massive star! Nuclear fusion and mass loss combined cause the star to evolve until it no longer can, at which point the cosmic engine sputters and grinds to a halt: the core collapses, the star explodes, a compact object gets left behind, and the next generation of stars can be formed from the stardust that remains. It turns out we have a name for this scenario that describes how a massive star evolves: it's called the "Conti Scenario" [49], and divides massive star evolution into three pathways. The "blue-red" pathway describes stars with initial masses between about 8 and 25 $M_\odot$, which start their lifes as OB main sequence stars, become RSGs, then explode as type II-P supernovae. Stars with initial masses between about 25 and 50 $M_\odot$ do a "blue-red-blue" path, evolving from O stars to RSGs, then losing enough mass to become Wolf-Rayet stars, before exploding as type Ib or Ic supernovae. Finally, the evolutionary branch for stars more massive than 50 $M_\odot$ is just called "blue." After (or even during) the main sequence, such stars experience one of those outbursting phases I referenced earlier (this one called the "Luminous Blue Variable" or LBV stage), loses its envelope, and becomes a WR before exploding as a Ib or Ic supernova.

If that sounds a little too qualitative or hand-wavy to you, that's because it is! Our models of how massive stars evolve suffer from a number of uncertainties that prevent us from making more quantitative predictions. Most of this uncertainty is driven by the fact that massive stars are *incredibly* rare. When a cold cloud of gas and dust collapses under its gravity to begin forming stars, it does not form stars equally: instead, less massive stars that are easier to make are formed much more often. The function that dictates the number of stars being formed in an infinitesimal bin of mass between M and $M + dM$ is called the initial mass function (IMF), and we usually

---

[18] It turns out neutrinos play an incredibly important role here as well, but this simple mental image will suffice for now.

[19] A teaspoon of neutron star crust weighs about ten tons! And yes, I said crust. Neutron star surfaces are solid.

characterize it as a power-law:[20]

$$\Phi(M) = \frac{dN}{dM} \propto M^{-\alpha} \tag{1.14}$$

and it seems like $\alpha = 2.35$ (a.k.a. the "Salpeter" IMF [251]) almost everywhere we look [113]. In other words, for every 10 $M_\odot$ star in a stellar population, there are about 200 1 $M_\odot$ stars. And for every 40,000 1 $M_\odot$ star, you only get a single 100 $M_\odot$ star. This means that observational studies of massive stars have been traditionally inhibited by small sample sizes, resulting in a number of uncertainties in key physical processes that govern how massive stars evolve.

1.3.1 Mass loss

Throughout their lifetimes, massive stars rapidly lose mass via their stellar winds. This was first realized in the 60s and 70s, where ultraviolet spectroscopic observations revealed blueshifted absorption lines that arose in the winds of OB stars [203, 269]. Early stellar evolution models that incorporated stellar winds demonstrated that mass loss significantly altered the evolution of massive stars, broadening the main sequence, causing stars to evolve to higher luminosity, and leading to the creation of Wolf-Rayet stars [47, 41]. However, a number of uncertainties remain in our understanding of mass loss in massive stars. While I will only focus on a only few key mechanisms for the remainder of this work, it is important to highlight some broader areas of uncertainty.

Observationally inferring mass lots rates requires an estimate of the mass density of the wind at a location where the wind is moving at its terminal velocity v_∞. In hot stars (mostly O and B dwarfs and supergiants, but also WRs), many observational diagnostics (e.g., radio observations and emission lines like Hα) are sensitive not to ρ but to ρ^2. Because winds are not smooth but clumpy, this results in an overestimate of mass loss rates compared to those derived by hard-to-observe diagnostics sensitive to ρ (like observations of ultraviolet P Cygni absorption) [84]. This implies that common recipes for mass loss (e.g. [315]) that are derived from theoretical models

---

[20] Some authors (e.g. [36]) have adopted more complicated functional forms for the low mass IMF, but in general, the high mass end of the IMF is a power law.

of a smooth, radiatively driven wind[21] should be reduced by some unknown “clumping factor” f_{cl} that can take on values anywhere from 3 to 10. However, even this reduction overpredicts the observed ρ-dependent mass loss rates of low luminosity O and B stars [235]. While this *weak wind problem* continues to stymie us, recent theoretical work has shown that shocks driven by the “line-deshadowing instability” [219] may heat the winds of such stars above their effective temperatures, resulting in an illusory reduction of mass loss rates derived from density-sensitive diagnostics. If these results are confirmed, this would potentially solve the weak-wind problem, bringing ρ- and clumping-corrected ρ^2-based measurements of mass loss in low luminosity OB stars back into alignment with each other.

Even after a star has evolved into a cool supergiant, mass loss remains important, and uncertain. Mass loss rates in RSGs are often inferred from mid-infrared observations, which reveal dusty circumstellar envelopes [323]. RSGs do not emit a significant amount of the UV radiation that is thought to needed to drive a wind radiatively. Instead, either via pulsations or turbulence [132], matter is gently wafted off the surface of the star until it is sufficiently cool enough for dust to form ($\sim$1000-1500 K), at which point radiation pressure on the dust can drive a slow, yet quite dense wind. In the most luminous RSGs, this wind is capable of stripping the envelope entirely, producing a WR. Varying prescriptions for RSG mass loss exist (e.g. [62, 216, 308]), which contain significant uncertainties. Depending on which recipe is adopted, the mass at which a star can shed its envelope via mass loss can change quite drastically from 20 to 50 $M_{\odot}$ [76, 77].

Finally, luminous stars can exhibit outbursts and eruptions. Most notable of these eruptive stars are the luminous blue variables (LBVs), descendants of $\gtrsim 50 M_{\odot}$ stars which are mostly luminous, sometimes blue, and occasionally variable. Notable members of this class have been cataloged since the early 17th century [23], and include η Car, S Doradus, and P Cygni, all of which have undergone strong outbursts in between long periods of quiescence. During these events, some unknown instability[22] triggers a powerful episode of mass loss during which the star’s optical

---

[21] The specific model is called the “CAK” model after the last initials of the authors that developed it [34]

[22] This instability is thought to be tied to these stars’ proximity to the Eddington limit, the point at which radiation pressure overcomes gravity, see [266] for a slightly out-of-date review

brightness dramatically increases,[23] and the star sheds anywhere from 0.1 to 10 $M_\odot$ over a very small amount of time (the entire LBV phase, including quiescent periods, is thought to last only $\sim 10^{4-5}$ yr, or $\sim$1-10% of the main sequence lifetime). Due to their short lifetimes, an incredibly small number of LBVs exist in the Galaxy, and significant debate has arisen regarding which stars are "true" LBVs [241], and whether LBVs even are massive stars to begin with![24] In addition to LBVs, we also observe two other important classes of outbursting star about which we know even less: B[e] supergiants which are dusty objects closely related to LBVs [149]; and certain high luminosity yellow supergiants[25] that occupy a region of the Hertzprung-Russell diagram (HRD) known as the *yellow void* [216, 282], where the atmosphere is thought to become dynamically unstable. Despite the tremendous amount of mass lost during these phases, stellar evolution models rarely include them due to our uncertainties in which stars experience them, how long they last, and what physical mechanism(s) drive them.

1.3.2 Rotation

Rotation is an ubiquitous property of stars, and its signal is observed across the entire mass spectrum. Low mass stars like our sun exhibit star spots, and as these spots rotate in and out of view, the brightnesses of spotted stars vary by a few percent [202]. While massive main sequence stars don't possess dark spots, they do display *bright* spots, which imprint themselves in high-cadence time-series photometry [24, 32, 18, 17, 125, 237, 229]. Spectroscopic measurements of suitable resolution can also yield rotational velocities by measuring the rotational broadening of absorption or emission lines. This method has a key caveat that only the component of the rotational velocity that is projected into the line-of-sight[26] can be measured, and in general, the inclination angle i is unknown [e.g., 114]. Regardless of the method, observations reveal that massive stars appear to be

---

[23] Even to the point where the eruption can be mistaken for a supernova. Interestingly enough, such supernova imposters can then actually explode shortly thereafter, e.g. [191]

[24] See a long and contentious series of publications [268, 118, 1].

[25] sometimes inexplicably called "hypergiants."

[26] A.k.a. $v_{rot} \sin i$, where i is the angle of the rotational axis relative to the line-of-sight, such that if we view the star pole-on, $i = 0$, and no rotational broadening will be observed

rapid rotators, with some stars rotating quite close to the *critical velocity*, v_{crit}, the point at which centrifugal force overcomes gravity and starts to rip the star apart.[27]

Rotation has a significant effect on the physical properties of massive stars. Most notably as far as the observational properties are concerned, rotation alters the very shapes of stars [316], causing the stellar equator to bulge outwards. Because the flux from a given surface area is proportional to the *effective gravity* which is lessened at the equators by centrifugal force, this leads to *gravity darkening/brightening*, a process where the poles of the star are hotter and brighter than the equator. As a result, an effective temperature measured assuming the star does not rotate can differ from the temperature of the poles/equator by several thousand K in the case of B stars [82]. In an extreme case, when stars rotate close to Ω_{crit}, they can form *decretion disks* out of material ejected from the equator. These disks are illuminated and heated by the central star, causing the entire system to display emission lines whose shapes are dependent on the inclination angle [243]. So-called *Be stars* (where the "e" stands for emission) make up $\sim$15-20% of all B-type stars [243].

In addition to its effects on the observed surface properties of stars, rotation can also impact their interiors. Rotationally induced mixing via diffusion and meridional circulation [329] can change the chemical structures of stars — which, as I've mentioned, impacts their evolution — and alters the surface element abundances. Finally, mass loss is enhanced in rotating stars. This is both due to radiative effects caused by gravity darkening [174], and mechanical mass loss due to the reduced equatorial gravity [93]. All-in-all, the initial rotation rate is a critical quantity in addition to the initial mass and metallicity that influences the evolution of massive stars. However, rotation rates are difficult to measure directly, and the implementation of rotational effects in stellar evolution codes is nontrivial.

1.3.3 Binary Interactions

Our star is a lonely one. At a distance of about 4.25 light years [86], Proxima Centauri is our closest stellar neighbor. Proxima, on the other hand, is itself only 0.2 light years away from *its*

---

[27] Both v_{rot} and v_{crit} can also be replaced with the *rotation frequency*, typically expressed as an angular frequency Ω_{rot} or Ω_{crit}.

closest neighbor, α Centauri AB, a distance to small that it is gravitationally bound, orbiting α Centauri AB once every 550,000 years. α Centauri AB itself is in fact a star system composed of two stars that orbit each other once every 80 years. *Multiple systems* like α Centauri come in various flavors and architectures determined by the number, masses, and kinematics of the stars in the system, but in general, they are extremely common. About $\sim$44% of all sun-like stars are born into a multiple system of one kind or another, with the companion fraction (the average number of companions per star) being around 0.6 [75].

Stars born in close enough binaries experience interactions, which occur primarily via tides and mass-transfer [120], the latter of which can occur via both stellar winds and Roche Lobe Overflow. Both mechanisms can change the angular momentum of the system, affecting either the orbital angular momentum, the rotational angular momentum of each star in the system, or the net angular momentum which is removed by gas ejected from the system entirely [65], which in turn can lead to further interactions, including those where the stars come into contact with each other. In the most extreme cases, the system enters a brief common envelope phase, which may be followed by a merger, depending on the orbital energy of the system [220, 67]. The effects of these interactions on the evolution of both stars in the system is heavily dependent on the evolutionary state of each star [151], which depends on the initial period and mass ratio[28] of the system. This leaves a large parameter space that must be fully explored in order to sample the entire range of binary effects; initial orbital periods can run from over 10,000 days (in which case the two components of the system are effectively single stars) to under a day [198], while mass ratios can be incredibly small, or $\sim$1. Furthermore, while observations are biased towards finding more massive companions, secondaries with low mass ratios are still capable of significantly altering the evolution of their primary at short orbital periods [77].

When it comes to massive stars, multiplicity is critical. In a radial velocity survey of 71 Galactic O stars in nearby open clusters, Sana et al. [253] directly searched for massive binary systems, and used their results to infer the intrinsic binary fraction f_{bin}[29] as well the distributions of binary

---

[28] Defined as the ratio of the mass of the less massive secondary star to the mass of the primary, $q = M_2/M_2$.

[29] Simply defined as the fraction of stars with a binary companion. While triplets like α Centauri or even higher-

parameters. They reported a tendency for binaries to favor close systems with mass ratios drawn from a uniform distribution. They also found a high binary fraction $f_{bin} > 70\%$ when including longer-period systems, and subsequent work [see e.g., 75, 254, 198] has also found evidence that the evolution of massive stars is dominated by binary interactions in close systems. However, other observations place the short-period binary fraction for O-type stars at ∼30-35% [e.g., 89, 252]. For post-main-sequence massive stars, the observed binary fraction for Wolf-Rayet stars is ∼30% [207], while the binary fraction of yellow and red supergiants appears to vary between ∼15-40% depending on the metallicity of the host galaxy [212, 210].

While I will discuss in detail the effects of binary interactions on both individual massive stars as well as stellar populations in Chapters 2 and 3, suffice it to say that they are important, impacting everything from the explodability of individual massive stars [153] to the radiation that plays a key part in reionizing the universe [277]. Critically, they are poorly understood. For decades, binarity was seen as a secondary effect relative to rotation,[30] and if it was included in stellar evolution models at all, it was coarsely approximated by, e.g., artificially removing a star's entire hydrogen envelope once it reached a given stage in its evolution [77, and references therein]. It is only recently that detailed models of stars in binaries have been produced. However, these models include many free parameters and uncertainties, including ones as fundamental as: what is the binary fraction, and how does it — as well as the distribution of initial periods and mass ratios — depend on initial mass and metallicity?

1.3.4 Mixing

At various times through this introduction, I've referred to "mixing" as an important, yet vague process that nonetheless is important for the evolution of stars. To be more specific, consider

---

order multiple systems exist, I largely neglect them in this work. This assumption is mostly valid, as the majority of such systems are either unstable, or form hierarchical systems consisting of tight binaries orbiting each other at large separations. For a neat visualization, see @ThreeBodyBot, an account on Twitter that posts videos of simulated three-body systems that clearly demonstrate this fact.

[30] This is largely based on the fact that both effects lead to higher mass loss, and while not all massive stars are in binaries (we think...), all stars rotate.

some species (e.g., hydrogen). For a spherically symmetric star, mixing is usually modelled as a time-dependent diffusive process, wherein at some location r the mass fraction of this species X_i changes at a rate given by

$$\frac{\partial X_i}{\partial t} = \mathcal{R}_i + \frac{1}{\rho r^2}\frac{\partial}{\partial r}\left(D\rho r^2 \frac{\partial X_i}{\partial r}\right) \tag{1.15}$$

where $\mathcal{R}_i$ covers changes due to nuclear fusion, and D is the *mixing coefficient*, which we usually take as the sum over a number of *mixing processes*.[31] Each of these processes can be described by some *mixing profile*, $D(r)$.

The dominant mixing process in stars is convection, which occurs when radiation is inefficient at transporting energy through the star by Eq. (1.5), and instead, the energy must be directly carried by matter moving within the star. In such cases, without convection, the star would generate a steep temperature gradient, which we compute relative to the pressure to make the math somewhat easier:[32]

$$\nabla \equiv \frac{\partial \ln T}{\partial \ln P} \tag{1.16}$$

With radiative energy transport,

$$\nabla_{\text{rad}} = \frac{3}{16\pi acG}\frac{\kappa P}{T^4}\frac{L}{M_r} \tag{1.17}$$

which you may recognize as a reformulated version of Eq. (1.5), and so long as the temperature gradient is shallow enough, convection won't occur. However, once ∇_{rad} is steeper than a temperature gradient that could be accomplished *adiabatically* (∇_{ad}), that layer becomes unstable, and convection occurs (i.e., the *Schwarzschild criterion*: stability is achieved when $\nabla_{\text{rad}} < \nabla_{\text{ad}}$). Alternatively, gradients in the chemical composition — which can result from mixing! — can increase the stability of a region to convection. In these cases, the *Ledoux criterion* for stability is

---

[31] Note that I'm expressing this equation in Eulerian coordinates. Conversion to Lagrangian coordinates is simple enough, though the equation is less easy to parse.

[32] In general, the pressure monotonically decreases from the stellar core to the surface as a result of (1.3). So high values of ∇ mean that a small increase in pressure (corresponding to moving inward in the star) results in a strong increase in temperature. In the same scenario, the gradient of the temperature relative to r or M_r would be *negative*, with a high absolute value. In general, I like to stay positive, hence taking derivatives relative to the pressure.

adopted:

$$\nabla_{\rm rad} < \nabla_L \tag{1.18}$$

where

$$\nabla_L = \nabla_{\rm ad} + \frac{\varphi}{\delta}\nabla_\mu \tag{1.19}$$

μ is the mean molecular weight, which, for fully ionized gas is sum over every species i:

$$\mu = \left(\sum_i \frac{X_i(1+Z_i)}{\mu_i}\right)^{-1} \tag{1.20}$$

with Z_i being the charge. The corresponding value of μ for neutral gas can be derived by setting all $Z_i = 0$, and for partially ionized gas, can be determined by the Saha ionization equation [245, 246]. The remaining quantities are defined as

$$\varphi = \left(\frac{\partial \ln \rho}{\partial \ln \mu}\right)_{P,T} \tag{1.21}$$

$$\delta = \left(\frac{\partial \ln \rho}{\partial \ln T}\right)_{P,\mu} \tag{1.22}$$

$$\nabla_\mu = \frac{\partial \ln \mu}{\partial \ln P} \tag{1.23}$$

While modeling the mixing of chemical species by convection really does require complex, three-dimensional simulations, the *mixing-length theory* developed by Erika Böhm-Vitense [25] seems to be the best approximation for convection in stars that we have. In mixing-length theory, convection results in a mixing coefficient

$$D_{\rm conv} = \frac{1}{3}\alpha_{\rm mlt} H_p v_{\rm conv} \tag{1.24}$$

where H_p is the pressure scale height (the physical distance over which the pressure decreases by a factor of e), $v_{\rm conv}$ is the typical velocity of gas blobs that are convecting, and $\alpha_{\rm mlt}$ is a free parameter that appears to vary between 1.5 and 2.2 based on both observations [217, 197, 5] and hydrodynamic simulations [314, 298, 124] (though even larger or smaller values are sometimes used, and are, in fact, equally physically plausible). Even though this is a relatively small range in values for $\alpha_{\rm mlt}$, small tweaks in this parameter can seriously impact the evolution of the star, resulting in, e.g., significant changes in the coolest temperatures that stars can reach as red supergiants [43]. Even worse, we have no clue whether $\alpha_{\rm mlt}$ has a universal value, or whether it is dependent on the stellar mass and metallicity, whether it changes through a star's lifetime, or even whether it should be constant throughout the entire stellar structure at a given time!

Beyond mixing length theory, additional mixing processes are also important, including three that are intricately linked to convection. The first is *convective overshooting*. At the interface between convective and radiative zones, the upwelling/downwelling material still has momentum, and some of the gas can penetrate into the convective zone, where it diffuses into its surroundings. A commonly used overshooting model is *exponential overshooting* [111]: given the convective mixing coefficient, $D_{\rm conv,0}$ measured some small fraction of the pressure scale height $f_{\rm ov,D}$ into the convective zone, the mixing coefficient due to overshooting is given by

$$D_{\rm ov} = D_{\rm conv,0} e^{-\frac{2z}{f_{\rm ov} H_p}} \tag{1.25}$$

where z is the distance into the radiative zone from the edge of the convective zone, and $f_{\rm ov}$ is yet another free parameter in addition to $f_{\rm ov,D}$. Both of these variables are usually assigned values of around 10^{-3} or so [128], and, like $\alpha_{\rm mlt}$ are horribly unconstrained. The second convective-adjacent process is *semiconvection*, which occurs in regions that are Schwarzschild-unstable, but stable under the Ledoux criterion. Though we do not model it like this in practice, semiconvection also occurs within the Earth, where it forms a "staircase" of thin layers that alternate between convective and radiative energy transport ([274] and references therein). The semiconvective mixing coefficient is given by

$$D_{\rm sc} = \alpha_{\rm sc} \left(\frac{K}{6 C_P \rho} \right) \frac{\nabla - \nabla_{\rm ad}}{\nabla_L - \nabla} \tag{1.26}$$

where K is the radiative conductivity, and C_P is the heat capacity at constant pressure, and $\alpha_{\rm sc}$ is (you guessed it) yet another free parameter that is of-order 10^{-2} [78]. Finally, *thermohaline mixing* occurs in layers that are Ledoux-stable, yet develop an inversion in the chemical gradient, causing the heavier material to sink and exchange heat along long, thin "fingers" [33], which we observe on the Earth's oceans when evaporation causes the ocean surface to develop a high concentration of salt. The mixing coefficient [136, 300] is given by

$$D_{\rm th} = \alpha_{\rm th}\left(\frac{3K}{2C_P\rho}\right)\frac{\frac{\varphi}{\delta}\nabla_\mu}{\nabla - \nabla_{\rm ad}} \tag{1.27}$$

with the associated free parameter $\alpha \mathrm{th}$, which is of-order unity.

As if this weren't enough uncertainty, physical processes beyond convection can result in additional mixing, especially in radiative zones that wouldn't otherwise mix. These processes can be macroscopic or microscopic [233, 250]. Though beyond the scope of this work, large scale turbulent flows [193, 279] and magnetic fields (via their effects on the motions of gas and the rotational evolution of the star [250]) are additional macroscopic sources of mixing. More relevant here are mixing due to rotation, and to wave dynamics. Though covered in detail in Chapters 3 and 6, in brief: rotation adds to the mixing coefficient via both a shear term, and a term from *meridional circulation* [291], while internal gravity waves launched at the interface between convective and radiative layers can directly transport material, and redistribute angular momentum in the star [290]. Amongst microscopic mixing processes, atomic diffusion plays a key role. Temperature and pressure gradients cause heavier elements to sink in the star, while chemical gradients (e.g. from the chemically enriched core to the relatively metal-poor envelope) have the opposite effect. Finally, radiative forces can levitate atoms in ways that depend critically on both local physical conditions, and the atomic structure. These four processes can be combined [296] to derive the diffusive velocities of every species v_i, and Eq. (1.15) then becomes

$$\frac{\partial X_i}{\partial t} = \mathcal{R}_i - \frac{1}{\rho r^2}\frac{\partial}{\partial r}(\rho r^2 X_i v_i) + \frac{1}{\rho r^2}\frac{\partial}{\partial r}\left(D\rho r^2\frac{\partial X_i}{\partial r}\right) \tag{1.28}$$

where the second term on the right hand side accounts for diffusion. In practice, implementing diffusion is devilishly complicated; thankfully, diffusion isn't critically important for the evolution of massive stars, but is a factor to consider.

All of these mixing phenomena are wildly unconstrained. While some observables (e.g., the location of the TAMS, or the temperature of the coldest RSGs in a population) can be used to constrain one of the free parameters mentioned above, the derived mixing profiles $D(M)$ are highly degenerate: many internal structures can yield the same set of observables. Recent breakthroughs in asteroseismology [3] have demonstrated the power of this technique for inferring mixing profiles. However, both our theoretical and observational capabilities to apply asteroseismology to massive stars are sorely lacking, as discussed in Chapter 6.

1.4 Why Evolved Massive Stars?

With all of these uncertainties, why should we (as I do in this thesis) focus on massive stars that have left the main sequence and dived into the glorious and peculiar menagerie of evolutionary states that I will collectively call "evolved massive stars"? After all, all massive stars are quite rare already by dint the IMF (Equation (1.14)). And because they spend so much of their lifetimes ($\sim$90%) on the main sequence, evolved massive stars are that much more difficult to find! Shouldn't we just focus on O and B dwarfs?

As a matter of fact, it turns out that things go quite well when we do. In general, evolutionary models agree fairly well on the behavior of massive stars on the main sequence. But once the models leave the main sequence, things go off the rails. Martins & Palacios [181] compared six different evolutionary codes, each of which predicted a drastically different behavior for the post main sequence evolution of a model with identical initial conditions. Some models evolve straight across the HRD to become RSGs. Others dip in luminosity before becoming RSGs. Still more become RSGs, then exhibit a so-called "blue loop" in the HRD as core helium fusion ignites. And all six models adopt reasonable, yet slightly different implementations of the physics discussed here! As a result, a measurement of a star's surface properties — its effective temperature, surface gravity, chemical abundances, rotation rates, and status as a binary star — cannot be used to uniquely infer its mass, age, or evolutionary status.

This behavior lead Kippenhahn & Weigert in their 1990 textbook [137] to describe evolved massive stars as

> *...a sort of magnifying glass, revealing relentlessly the faults of calculations of earlier phases.*

Now they aren't wrong. But I like to turn this way of thinking on its head. I think that if you have a magnifying glass, you should use it! If evolved massive stars are the objects in which the effects of mass loss, rotation, binary interactions, mixing, and more are the most pronounced, we should use them as precision probes of these processes. Evolved massive stars aren't the problem; they are the *solution*, our key to unlocking many mysteries in the evolution of massive stars.

Of course, if it were that easy, someone would have done it already. Unfortunately, until recently, our magnifying glass as been a little foggy. Observationally, small samples sizes coupled with the difficulty in obtaining measurements with high enough precision have impeded traditional studies of massive stars. Theoretically, many of the effects I've listed have been approximated due to a lack of computational investment/resources. But now in the third decade of the era[33] we call "the two thousands", things are changing; our once-foggy magnifying glass is clearing. New space-based observatories will either launch soon[34], or have already begun observing, and will revolutionize/are revolutionizing the ways in which we think about studying evolved massive stars, while cutting edge ground-based instrumentation is allowing us to follow-up on the most interesting objects. Simultaneously, theorists have begun grappling with nitty gritty physics, providing us with detailed implementations of the key effects that impact massive star evolution. This means that we can now begin to take the techniques and methods that have long been used to study low mass stars, and apply them in new ways to study evolved massive stars!

And so at last we come to the point of this thesis: to use evolved massive stars as magnifying glasses. I do this in a few different ways. In Chapters 2 and 3, I use cutting edge stellar evolution models to study the impact of binary interactions and rotation on the massive star demographics of resolved stellar populations, and demonstrate that both effects manifest themselves in unique and non-degenerate ways. The models I present require large catalogs of evolved massive stars

---

[33] I hesitate to call it a millenium, as the last 18 months have felt like an eon.

[34] Assuming launch windows do not continue to be pushed back indefinitely.

that have been reliably classified in order to make inferences, and in Chapter 4, I explore the possibilities (and shortcomings) of using machine learning to classify thousands of stars using only broadband photometry. In Chapters 5 and 6, I explore another avenue of probing the *interiors* of evolved massive stars using asteroseismology. Finally, in Chapter 7, I will briefly summarize all of my results, and discuss how the study of evolved massive stars will change in the 2020s and beyond.

With me so far?

Ok, let's begin.

Chapter 2

STELLAR POPULATION DIAGNOSTICS OF THE MASSIVE STAR BINARY FRACTION

2.1 Introduction

Comparing theoretical and observed populations of massive stars can be an incredibly powerful tool for understanding stellar evolution. Massive stars are luminous, and can be easily seen in the Local Group; photometric catalogs are readily available [e.g., the Local Group Galaxy Survey, LGGS; 186, 185], from which massive stars can be selected after filtering for foreground contaminants [e.g., 188]. The relative abundance of various subtypes of massive stars can then be used as a probe of stellar physics. Reproducing the observed data reflects our understanding of the relative lifetimes of these evolutionary phases, and therefore our ability to predict the impact of massive stars on their surroundings — e.g., chemical yields, ionizing radiation, and mechanical feedback. However, most stellar evolution models assume that stars evolve in relative isolation, without any influence from a binary companion.

Massive stars in interacting binary systems face drastically differing evolutionary pathways than their single-star cousins, causing ensembles of binary stars to appear notably different as a function of age, metallicity, and the underlying statistical distributions of the binary parameters of the systems. Measuring these properties directly is possible; however, such surveys are time intensive, and require detailed understandings of the completeness of the survey and the sensitivity of the observational method to systems of varying periods, inclinations, and mass ratios. Correcting for these effects in small samples can be difficult, making it hard to generalize the results to the entire population of massive stars. Thus any inferences made about young stellar populations are inherently polluted by unresolved binaries that have not been accounted for [67].

The predicted number of almost every subtype of massive star depends upon binarity. Perhaps the most notable effect is an increase in the expected number of stripped-envelope stars (i.e., Wolf-Rayet stars, WRs) with f_{bin}, which occurs due to RLOF onto the secondary star. Other, more subtle, effects can alter the number of massive stellar subtypes observed through time, e.g., red supergiants (RSGs), yellow supergiants (YSGs), blue supergiants (BSG) and the various WR subtypes (WC, WN, etc.). Indeed, recent work has argued that entire subclasses of massive stars may exclusively be the product of binary evolution (e.g., Luminous Blue Variables, LBVs; see [268, 118]). One may conclude then that using stellar count diagnostics as a probe of stellar physics is hopeless in the presence of binary stars with unknown properties. More optimistically, we seek to understand the effect of binaries on star count diagnostics to determine if they can be used to disentangle binary effects from single-star evolution.

This chapter is laid out as follows. In §2.2, I describe the Binary Population and Spectral Synthesis (BPASS) code and the population synthesis method we employ to generate theoretical predictions for the abundance of various subtypes. In §2.3 we describe various ratios of these subtypes that are sensitive to age, metallicity, and binary fraction. §2.4 describes how these ratios can be applied to real data, while §2.5 includes our presciption for handling incomplete samples of massive stars, and the effect of incompleteness on the inferred results. We apply these ratios to populations with complicated star-formation histories in §2.6 before concluding in §2.7. This work originally appeared in Dorn-Wallenstein et al. (2018) [68].

2.2 Creating Theoretical Populations With a Physical Treatment of Binaries

2.2.1 Binary Evolution and BPASS

The Binary Population and Spectral Synthesis code [BPASS, 77, 276] incorporates many of these effects in a custom stellar evolution code that is evaluated for single and binary stars on a dense grid[1] of initial primary and secondary masses (M_1 and M_2), initial periods P, and mass ratios ($q \equiv M_2/M1$) at 12 metallicities. We express the metallicity as a mass fraction Z, and BPASS

---

[1] Details on the grids of parameter values and more can be found in the BPASS v2.2 User Manual, currently hosted online at `https://drive.google.com/file/d/1IYCYf5Bxt1WmqPuFTYLQ7kpN-hKY2SAp/view?usp=sharing`

adopts metallicities in the range $10^{-5} \leq Z \leq 0.04$. Note that we assume solar metallicity $Z_\odot = 0.014$ [13]. Figure 2.1 shows solar metallicity evolutionary tracks from BPASS v2.2 for primary stars of initial mass M_1 between 15 and 50 $M_\odot$, companions with mass ratio $q = 0.9$, and initial orbital periods between 10 and 1000 days, as well as the corresponding single-star evolution track ($P \to \infty$). At the widest orbital separations, the primary stars evolve more or less identically to their single counterparts until the very end of their lives, where they fill their Roche lobes as yellow or red supergiants. This is most noticeable for the 30 $M_\odot$ tracks, where the primary of the $P = 10^3$ day binary only fills its Roche Lobe when it is close to the Hayashi limit. The subsequent mass transfer reduces the luminosity of the primary (see Figure 13.1 of [150], with data from [64]), and causes the primary to end its life as a lower luminosity Wolf-Rayet star.

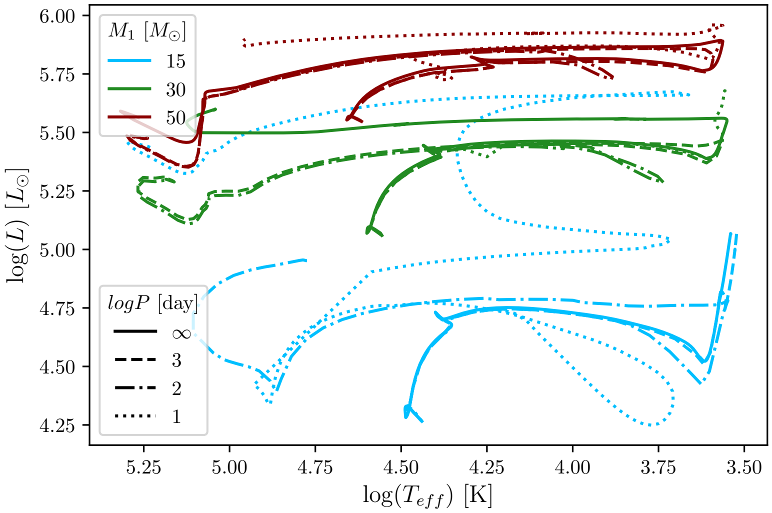

Figure 2.1 BPASS solar metallicity stellar evolution tracks for 15 (blue), 30 (green), and 50 (red) $M_\odot$ primary stars, assuming a mass ratio $q = 0.9$, and initial orbital periods of 1000 (dashed), 100 (dash-dotted), and 10 days (dotted) respectively. The corresponding single-star evolution track is shown with the solid line. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [68].

At progressively shorter periods, the effects of mass transfer become increasingly extreme. This is especially drastic for the 10 $M_\odot$ models, where mass transfer begins to occur earlier and earlier in the star's post-main sequence life, drastically altering its evolution. In the most extreme case presented ($P = 10$ days), the primary experiences multiple episodes of mass transfer both to and from the secondary, ultimately becoming an incredibly luminous Wolf-Rayet star instead of reaching the Hayashi limit and ending its life as a red supergiant as would be expected for an isolated 10 $M_\odot$ star. We note that these are extreme examples chosen to illustrate the range of behavior that occurs in close binaries for specific combinations of system parameters; each individual model is ultimately assigned a very small weight in the ensuing population synthesis due to the large number of models (see §2.2.2).

If massive stars truly favor high binary fractions and short orbital periods then very few stars will evolve completely free from the influence of a companion. This has a drastic effect on the relative numbers of stars of a given subtype in a population. For example, the cluster Westerlund 1 is a single-age ($\sim$ 5 Myr, [145]) massive cluster known to contain both red supergiants and Wolf-Rayet stars [44]. Single-star evolution predicts that these stages are evolved from stars in two almost entirely disjoint sets of initial masses, implying that single-aged clusters containing both RSGs and WRs should only exist for an incredibly narrow window of time after an initial starburst. However, allowing for the formation of Wolf-Rayet stars via RLOF-induced channel increases the overlap in the initial masses of RSG and WR progenitors.

2.2.2 Population Synthesis

With the complete set of single and binary stellar evolution tracks, we assembled synthetic populations by weighting each model according to the likelihood that it would be formed in an instantaneous burst of star formation. When considering only single stars, the weighting is calculated according to an initial mass function (IMF), the probability $\Phi(M)$ of a star being formed at a given mass. Φ is typically parametrized as a power law or broken power law. BPASS allows for population synthesis assuming one of nine IMFs, including IMFs with low-mass exponential cutoffs [36], and the classical Salpeter IMF with a slope of -2.35 [251]. We adopt the BPASS default, which is a broken power law with slope -1.3 below 0.5 $M_\odot$, and a slope of -2.35 for higher masses, with a maximum mass of 300 $M_\odot$. Because we are mostly considering massive stars, the shape of the low-mass IMF should have little effect on our results.

When adding binary stars, individual models must also be weighted according to the distributions of the fundamental natal parameters P and q of the binary system. BPASS v2.2 adopts the distribution parameters from [198], who found that these distrubitions are interrelated with, for example, the power law slope and twin ($q = 1$) fraction of the mass ratio distribution depending on the initial mass and period[2].

Finally the weightings are normalized to ensure that the entire population forms $10^6 M_\odot$ of stars. For each metallicity, we create two synthetic populations: one composed entirely of single stars using the input files provided in the BPASS v2.2 data release, and one composed entirely of binary stars using custom input files provided by the BPASS team (J. J. Eldridge 2018, private communication). While no $f_{bin} = 0$ or 1 populations have been observed, creating these populations allows us to generate results with tailored intermediate values of f_{bin} by mixing both populations accordingly. Note that the binary input files provided in the BPASS v2.2 data release assume the binary fractions found by [198], which would enforce an implicit maximum f_{bin} on our results; we instead use our custom $f_{bin} = 1$ population to avoid biasing our results.

2.2.3 Number Counts vs. Time

We now examine each stellar evolution track to determine the evolutionary phases that it goes through by assigning a subtype to all timesteps in the model, using the model parameters listed in Table 2.1; we largely adapt the classification scheme from [77]. First a check is done on $\log(T_{eff})$ and X to determine if the star is a Wolf-Rayet star — i.e., $\log(T_{eff}) \geq 4.45$ and $X \leq 0.4$. If it is and $X > 10^{-3}$, it is a WNH star; otherwise, it is classified as a WN or WC star depending on the ratio of $C + O$ to Y. If a star is not a WR, a $\log(T_{eff})$ and $\log(g)$ check is performed to determine if the star is a Of star — i.e., an evolved O star with a particularly strong stellar wind. These rare stars are included as a separate class because they are particularly strong sources of He II emission lines [31]. If the star is neither a WR nor a Of star, it is then assigned a classical MK spectral type

---

[2] The exact distribution parameters can be found in Table 13 of [198]

Table 2.1. BPASS model parameters used to label timesteps with an evolutionary phase.

Parameter	Description	Unit
$\log(L)$	Logarithm of the luminosity	$L_\odot$
$\log(T_{eff})$	Logarithm of the effective temperature	K
$\log(g)$	Logarithm of the surface gravity	cm s^{-2}
X	Hydrogen Surface Mass Fraction	-
Y	Helium Surface Mass Fraction	-
C	Carbon Surface Mass Fraction	-
O	Oxygen Surface Mass Fraction	-

based on its effective temperature. The exact numerical criteria used for our classification are listed in Table 2.2. Figure 2.2 serves as an illustration of the various temperature criteria used, compared to evolutionary tracks for single stars with initial masses between 5 and 50 $M_\odot$.

Once a timestep is assigned a label, it is then assigned to at least one of 51 time bins that are logarithmically-spaced between 10^6 and 10^{11} years in 0.1 dex increments. The weight of that model from the input file is then adjusted by the size of the model timestep (accounting for the fact that some timesteps cross the boundaries of the logarithmic time bin, and thus contribute to two time bins in differing amounts), and its final weight added to an entry in an array corresponding to its assigned label, time bin, and luminosity. We then create output arrays for each subtype by summing the array over the luminosity axis. We also create outputs assuming minimum luminosities for each subtype in 0.1 dex steps between $\log(L) = 3$ and 6. All of the arrays have been compiled into a single file, which we make available online.

Finally, as a post-processing step for this chapter, we label stars above $\log(L) = 4.9$ as supergiants, classifying O, Of, B, and A stars as BSGs, F and G stars as YSGs, and K and M stars as RSGs. As discussed by [77], this is based on the luminosity criterion used by [184] to ensure that lower-mass AGB stars were not included in their sample of RSGs, and applied to the rest of the supergiants for consistency.

For most of our analyses, we apply the same luminosity cutoff when considering the number of Wolf-Rayet stars in a population. As discussed previously, binary interactions are capable of stripping low mass stars that would be otherwise incapable of losing that much mass through stellar winds or instabilities alone. This results in a large number of "Wolf-Rayet" stars at ages well older than when the last WR stars are expected to disappear. These stars may appear as both binary systems and single stars (in the case of stripped secondaries), and exhibit a range of spectra from classical WR spectra to hot subdwarfs [100]. Very few such systems have been found — e.g. ϕ Persei [94], FY CMa [231], 59 Cyg [232], 60 Cyg [320], HD 45166 [281, 105]. This may be due to detectability issues (as the "Wolf-Rayet" primary can be far less luminous than the mass-gaining secondary), or a lack of atmospheric models for these stars. The luminosity cutoff for WRs attempts to mitigate this issue; however, BPASS still predicts the existence of luminous yet low-mass WRs well after ages of 10 Myr. No WRs have been found in intermediate-age clusters, which is in tension with a high binary fraction for massive stars.

Figure 2.3 shows the predicted number counts for times between 10^6 and $10^{7.5}$ years at three example metallicities for subtypes O, B, BSG, YSG, RSG, WR, WN, and WC. We then find N_{max}, the maximum expected number of stars of each subtype at both values of f_{bin}, and scale by the appropriate N_{max} so the maximum number of each subtype at each f_{bin} is 1 for clarity. The $f_{bin} = 1$ and $f_{bin} = 0$ populations are indicated with solid and dashed lines respectively. While neither is representative of a physical sample of stars, the comparison between the two is useful for understanding the effects of binarity on a population.

As expected, the Wolf-Rayet stages are most heavily affected: WRs at $f_{bin} = 1$ appear slightly earlier, while the age at which the most WRs are predicted is significantly later. There is also a converse effect on the RSGs, which is most noticeable at low metallicity, because the stellar-wind channel of WR creation is diminished, allowing the the effects of binarity on the WR population to dominate. Because of our $\log(L) \geq 4.9$ cutoff, the WRs created by binary interactions are massive ($\geq 8\ M_\odot$), and would thus turn into RSGs if they were single stars. In Figure 2.4, we show the total number of RSGs (red) and WRs (purple) per $10^6\ M_\odot$ of stellar material created in both the single (dashed) and binary (solid) populations at $Z = 0.002$. In the $f_{bin} = 0$ population, we see the

Table 2.2. Criteria used to label regions of the HR diagram to classify evolution tracks. Adapted from Table 3 of [77]. The luminosity cutoff for WR stars is not always applied; see §2.5 for details.

Label	Criteria
WNH	$\log(T_{eff}) \geq 4.45$ $X \leq 0.4$
WN	$\log(T_{eff}) \geq 4.45$ $X \leq 10^{-3}$ $(C+O))/Y \leq 0.03$
WC	$\log(T_{eff}) \geq 4.45$ $X \leq 10^{-3}$ $(C+O))/Y > 0.03$
O	$\log(T_{eff}) \geq 4.48$
Of	$\log(T_{eff}) \geq 4.519$ $\log(g) > 3.676\log(T_{eff}) + 13.253$
B	$4.041 \leq \log(T_{eff}) < 4.48$
A	$3.9 \leq \log(T_{eff}) < 4.041$
F/G	$3.66 \leq \log(T_{eff}) < 3.9$
K	$3.55 \leq \log(T_{eff}) < 3.66$
M	$\log(T_{eff}) < 3.55$
BSG	O + Of + B + A $\log(L) \geq 4.9$
YSG	F/G $\log(L) \geq 4.9$
RSG	K + M $\log(L) \geq 4.9$
WR	WNH + WN + WC $\log(L) \geq 4.9$

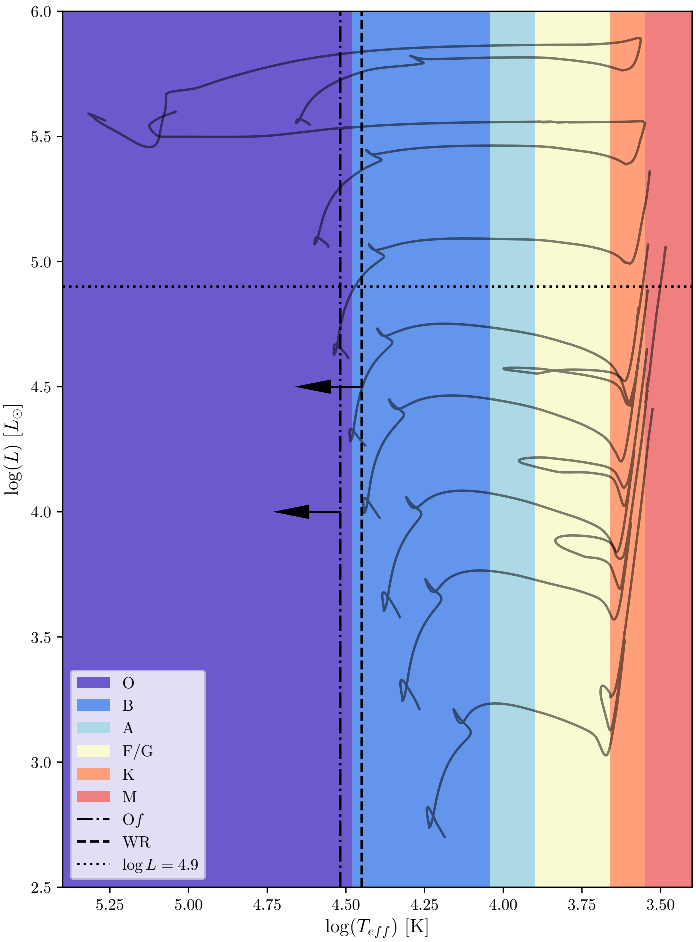

Figure 2.2 Visualization of the criteria used to count massive stellar subtypes from [77]. Spectral types are indicated by the colored patches. The minimum temperatures for Of and WR stars are shown by the dash-dotted and dashed lines respectively; additional criteria, and the criteria for various WR subtypes are in Table 2.2. For comparison, single-star solar metallicity BPASS tracks from 5 to 50 $M_\odot$ are shown in gray. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [68].

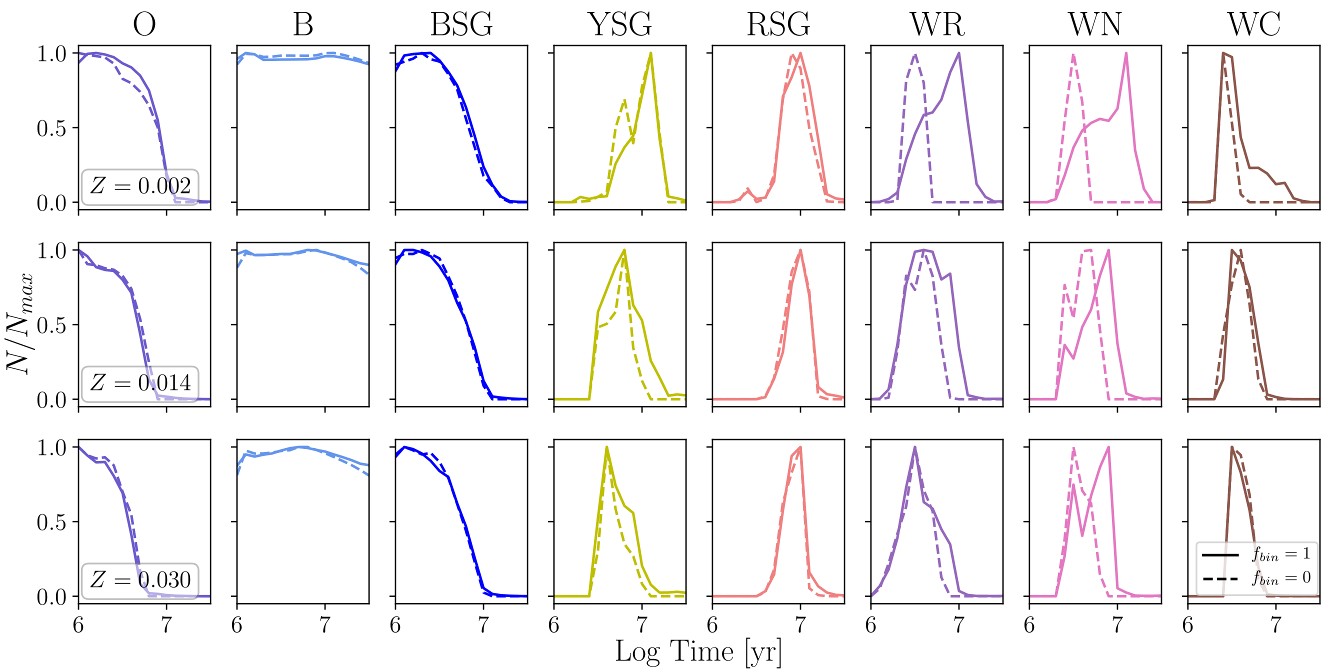

Figure 2.3 Number of various stellar subtypes in the $Z = 0.002$ (top), $Z = 0.014$ (middle), and $Z = 0.03$ (bottom) populations between 10^6 and $10^{7.5}$ years, scaled so that the maximum number of any subtype is 1. The solid and dashed lines indicate the $f_{bin} = 1$ and $f_{bin} = 0$ populations respectively.

expected behavior: there are far fewer WRs created, and they coexist with RSGs for a very narrow window of time. However, in the $f_{bin} = 1$ population, the number of RSGs is suppressed by a factor of ~ 4, and we see WRs at far later times.

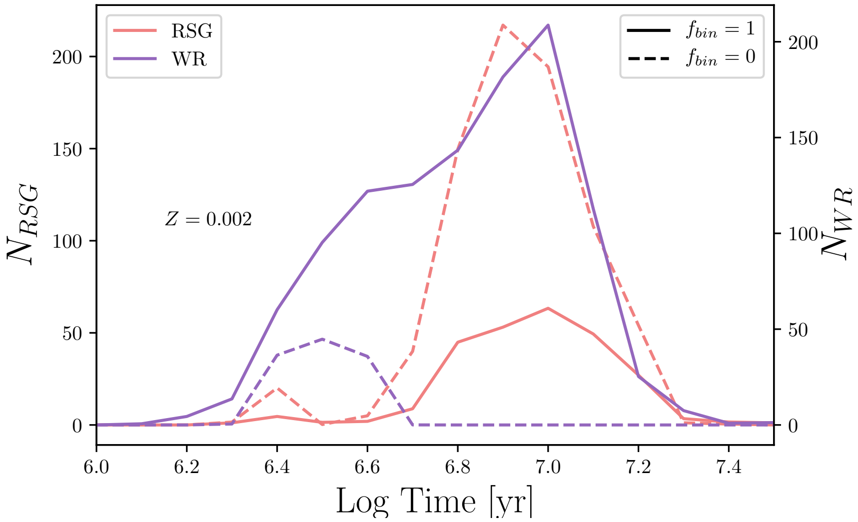

Figure 2.4 Absolute number of RSGs and WRs per $10^6\ M_\odot$ of stellar mass created for both binaries (solid) and singles (dashed) at $Z = 0.002$. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [68].

We note that we do not include Luminous Blue Variables (LBVs) in our classification scheme. Due to their eruptive outbursts, LBVs are certainly important to stellar populations, and we highlight their interesting properties in Chapter 5. However, the term has been applied to a set of objects with a wide variety of photometric and spectroscopic behavior [48], such that the exact definition of what is and isn't an LBV often varies from source to source, and only tens of *confirmed* LBVs (i.e., those which have been observed in a S Dor-type outburst) exist in the entire Local Group [305, 183, 241]. Additionally, distances to a set of LBVs and LBV candidates from the second data release of the Gaia survey [85] were derived by [267]. In many cases, the updated distances are smaller than those previously reported, implying that LBVs may occupy a different region of the HR diagram than is often assumed. Due to the present uncertainty in the evolutionary status of

LBVs, and the lack of a clear consensus in how to observationally classify a statistically significant number of them without long-term monitoring — a fact that we will address in Chapter 4 — we choose to not consider the LBV evolutionary phase in our analysis.

We can now construct the expected number counts for realistic populations with a given f_{bin} by mixing the two populations in proportion. However, these values are all relative to the total mass formed in a population, M_*. Estimating M_* can be a difficult exercise, as it can depend heavily on the lower-mass IMF, as well as the measured age and distance to the population. As we demonstrate, measuring the exact age of a population can be complicated by the presence of stellar binaries. Thus instead of comparing our number count predictions directly to populations, we can construct diagnostic ratios using massive stars. These ratios are independent of both the total mass and shape of the low-mass IMF.

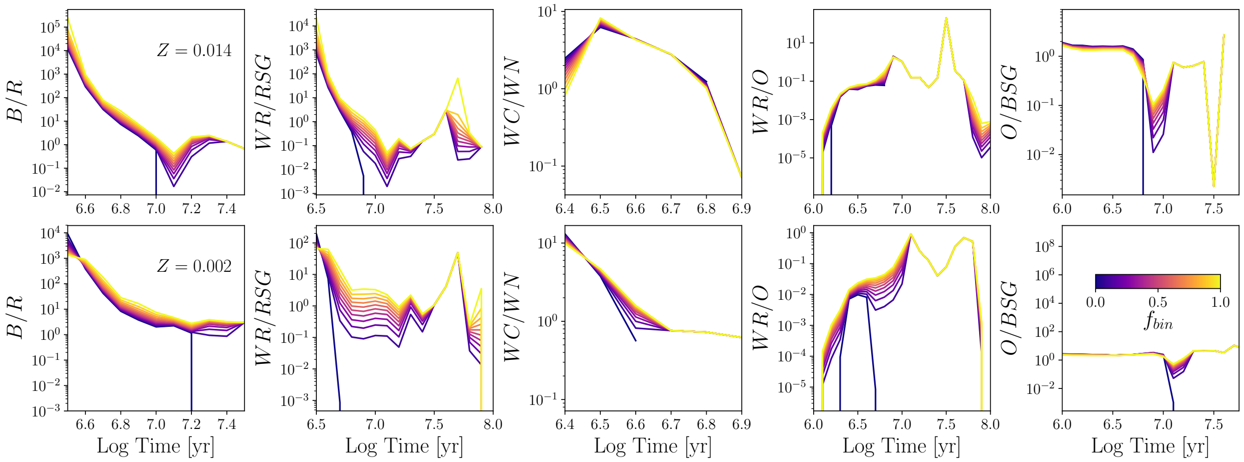

Figure 2.5 From left to right: B/R, WR/RSG, WC/WN, WR/O, and O/BSG vs. log time at $Z = 0.014$ (top row) and $Z = 0.002$ bottom row, calculated for f_{bin} between 0 and 1, as indicated by the colorbar on the top right. Axis limits on the abscissa are chosen to highlight the timescales on which these ratios are most dependent on varying f_{bin}. For all subtypes except O stars, a minimum luminosity of $\log(L) = 4.9$ is enforced. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [68].

2.3 Diagnostic Ratios

We first construct the predicted number counts for subtypes in a population with a given f_{bin}. We calculate the abundance of a subtype S at time t and metallicity Z as

$$S(t, f_{bin}, Z) = f_{bin} S_b(t, Z) + (1 - f_{bin}) S_s(t, Z) \tag{2.1}$$

where S_b and S_s are the abundances in the $f_{bin} = 1$ and $f_{bin} = 0$ populations respectively. We begin by looking only at simple stellar populations (SSPs, i.e., instantaneous bursts of star formation) to determine the effect of age, metallicity, and varying f_{bin}, before examining more complicated populations. Figure 2.5 shows the values of five different ratios vs. time for SSPs with solar metallicity (top row) and $Z = 0.002$ (approximately the metallicity of the SMC, bottom row), and binary fractions between 0 (purple) and 1 (yellow) as indicated by the colorbar. The bounds of the time axis have been chosen to highlight the time range during which each ratio is most dependent on f_{bin}.

2.3.1 B/R

One of the most frequently used diagnostics is the ratio of the number of blue supergiants to red supergiants (B/R). Its earliest uses were to corroborate the then-putative metallicity gradient in M33 [318, 301]. While the trend of B/R increasing with increasing metallicity stymied theoretical models [152], it has still been used as a metallicity diagnostic by multiple studies [182]. B/R is mostly sensitive to the physics governing a star's rightward evolution in the HR diagram after the main sequence, and thus is dependent on rotation and convection for single stars. Note that subsequent leftward movement in the HR diagram occurs during the final stages of a star's life and is quite rapid.

Because massive stars evolve at approximately constant luminosity, and reach their coolest temperatures (i.e., largest radii) during their first crossing of the HR diagram, the first instance of RLOF for a binary must occur during this initial rightward movement. This interrups the star's normal evolution, and causes it to evolve blueward on the HR diagram. Therefore, it would make sense that binary interactions reduce the number of red supergiants, increasing B/R. The first column of Figure 2.5 shows the predicted B/R values in our SSPs. We find that binarity does increase B/R at most times by factors of 2-10. Considering that errors on number count ratios in star clusters can be an order of magnitude or more in all but the most massive clusters with IMFs that are well populated out to tens of $M_{\odot}$, measuring this effect requires exquisite statistics. However, B/R varies by many orders of magnitude as a function of time, implying that it is a much better age diagnostic.

2.3.2 Wolf-Rayet Ratios

After B/R, perhaps the most-used number count ratios involve Wolf-Rayet stars. In the single-star paradigm (the "Conti scenario", [49]), they evolve from the most massive progenitors, and the full sequence from WN to WC/WO stars are thought to be a progression of increasingly stripped stellar envelopes. As subsequent layers are revealed, the products of more and more advanced nuclear fusion stages that have been mixed to those layers are revealed. Thus WRs are useful probes of extremely rapid mass loss. WRs have an observed binary fraction of $\sim$30% [207], to say nothing of the intrinsic binary fraction or WRs that originated as secondary stars of systems that have since been disrupted by supernovae. It is thus important to discuss WR-based diagnostic ratios in the context of stellar binaries.

WR/RSG

WRs and RSGs are thought to evolve from progenitors with two mostly disjoint sets of initial masses, so their coexistance in a star cluster only occurs for an incredibly narrow window in time (e.g., the dashed lines in Figure 2.4). Thus, with the notable exception of Westerlund 1, which we discuss later in this section, WR/RSG has most often been used in the literature as a metallicity diagnostic in galaxies: [171] note that WR/RSG changed by factors of up to 90 in the Milky Way as a function of galactocentric distance between 7 and 13 kpc. Moreover, they proposed that

WR/RSG (or more accurately, its inverse) is an even more sensitive metallicity diagnostic than B/R. This is because the relative abundance of both subtypes is highly sensitive to the exact mass ranges of their progenitors, which in turn is affected by metallicity-dependent mass-loss. However, quite recently (and after the original publication of this work), updated data showed that RSG/WR was in-fact roughly independent of metallicity, bringing the data into alignment with the models presented here [189].

As discussed in §3.2.3, binary interactions have an incredibly drastic effect on the relative numbers of both subtypes, especially at low metallicity. Thus it is unsurprising that the behavior of WR/RSG is *incredibly* dependent on f_{bin}. The second column of Figure 2.5 shows WR/RSG for our SSPs. As expected, at $f_{bin} = 0$, $WR/RSG \rightarrow 0$ by ~ 5 Myr. However, once binaries are included, more WRs are produced, so WR/RSG has defined values well after this time. Indeed, WR/RSG takes on values spanning multiple orders of magnitude as a function of both age and f_{bin}. Issues of "missing" old WRs notwithstanding, we note for now that if significant numbers of these WRs produced through binary evolution channels are found in populations with ages of a few 10s of Myr, WR/RSG can be a powerful diagnostic of both f_{bin} and age in SSPs.

WC/WN

A second often-used diagnostic, WC/WN, uses only the relative abundance of WR subtypes. Compared to the rest of the ratios discussed, WC and WN stars arise from a mostly-overlapping set of initial masses (at least from the single-star perspective). Most interestingly, it is sensitive only to the lifetimes of WR phases, and should be mostly independent of both the IMF and which channel produces WRs. Thus, as proposed by [311] and [110], WC/WN is solely a function of the metallicity and temperature dependence of Wolf-Rayet winds. The third column of Figure 2.5 shows WC/WN vs. time at solar and subsolar metallicity for varying binary fraction. As expected, there is minimal dependence on f_{bin} at almost all times, except in the lower metallicity population for a brief window around $\log t = 6.6$. Thus, for most metallicities/ages, WC/WN should indeed be a useful diagnostic, free from the influence of unresolved binaries.

WR/O

A third diagnostic, WR/O, is a probe of a large swath of the mass spectrum of massive stars. Both Galactic WR catalogs [e.g., 303] and surveys of the Local Group [186, 185] have made data available in environments spanning a wide range of stellar masses, metallicities, and star formation histories. However, like WR/RSG, WR/O is especially succeptible to contamination by binaries [170]. The fourth column of Figure 2.5 shows WR/O vs. time for different metallicities and binary fractions. While difficult to see due to the large y-axis scale of the plot, WR/O is indeed affected by including binaries at both early and late times.

2.3.3 O/BSG

Finally, we introduce a ratio that is rarely discussed in the literature: O/BSG. This ratio is largely sensitive to the spectral type of the main sequence turnoff, and thus main sequence lifetimes. The final column of Figure 2.5 shows O/BSG vs. time. For most ages, O/BSG is insensitive to f_{bin}, and generally declines from early to late times as the turnoff moves to later spectral types. However, for a small window around $\log t \approx 6.75/7$ ($Z = 0.014$/$Z = 0.002$ respectively), O/BSG exhibits f_{bin}-dependent behavior. This is likely due to stars that experience moderate amounts of RLOF, and evolve blueward, but haven't lost enough to of their H envelopes to become WR stars. These stars are then classified as O, increasing O/BSG.

2.4 Comparisons with Real Data

Given a complete sample of massive stars in a population, it is possible to calculate stellar count ratios to compare to predictions. However, because the different subtypes require various and typically time-intensive methods for discovery and classification, it is often the case that the data for individual types of stars must be assembled from a variety of inhomogeneous sources. Thus only a few subtypes may have been cataloged, from which only a few ratios can be calculated. Therefore it is critically important to choose ratios that are best suited to the population under consideration — i.e., suitable diagnostics of age, metallicity, or f_{bin}.

Massive stars are rare, and the abundance of evolved massive stellar subtypes is subject to Poisson noise in star clusters with $M_* \sim 10^5\ M_\odot$ at most. Thus, great care must be made when comparing the theoretical to observed values. With a SSP we can calculate the value of arbitrary ratios at infinite signal to noise on a grid of ages and binary fractions. To estimate the error on the real data, we follow [244], and assume the error on each number count measurement N is $\sqrt{N}$. Thus, for subtypes X and Y with observed number counts X and Y, the error[3] of the measured ratio X/Y is

$$\sigma_{X/Y} = \frac{X}{Y}\sqrt{\frac{1}{X} + \frac{1}{Y}} \tag{2.2}$$

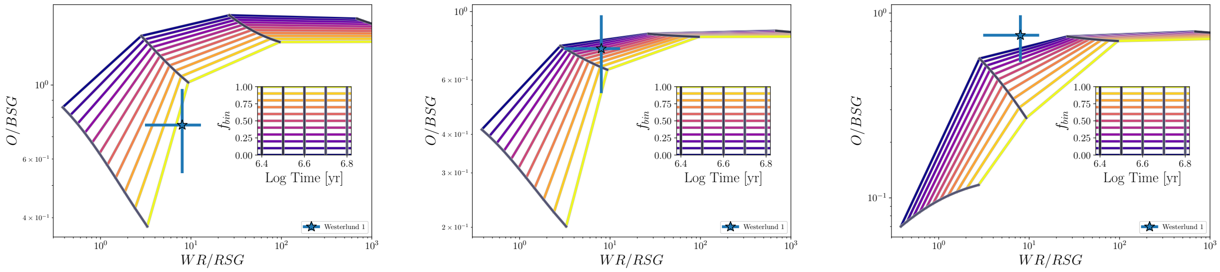

Figure 2.6 *Left*: Diagnostic two-ratio plot, applied to the young super star cluster Wd1, which has a cohort of WR, RSG, BSG, and O stars. The inset grid is for reference when interpreting the figure, indicating f_{bin} increases from purple to yellow, and (log) time increases from dark to light grey. *Center*: Identical, but assuming completeness limits consistent with the data: WR stars are complete down to $\log(L) = 5.1$, all O stars are supergiants ($\log(L) \geq 4.9$), and the RSG sample is complete ($\log(L) \geq 4.9$). ***Right***: Identical, but assuming an overly conservative completeness limit of $\log(L) = 5.2$ for all species. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [68].

The following examples illustrate a few possibilities. We invite the reader to make use of our publicly available code to explore the entire parameter space to develop diagnostic plots suitable

---

[3] It is important to note that is a very naive assumption; X and Y are discrete Poisson variables, and their ratio X/Y is not normally distributed. Thus assigning an equal-tailed confidence interval to X/Y via $\sigma_{X/Y}$ is incorrect. As we only wish to compare the rough scales of the spread in model predictions and the typical uncertainty of measured ratios, using Bayesian inference or other methods to construct more accurate confidence intervals is beyond the scope of this chapter. However, authors wishing to make quantitive inferences absolutely should make robust error estimates, as we do in Chapter 3.

to their dataset. We note that, while the assumption that real star clusters are true SSPs is suspect [99], there are a variety of open clusters with an age spread less than the time resolution of our SSPs.

We first consider a young (< 10 Myr) solar metallicity population, which has a wealth of both main sequence and evolved massive stars. Many Galactic super star clusters are this age, and are massive enough to have well-populated IMFs out to tens of $M_{\odot}$. Westerlund 1 (Wd1) in particular is notable for being a $M_{*} \approx 5 \times 10^{4}\ M_{\odot}$ cluster [6], with a well studied cohort of evolved massive stars [44, 51]. Notably, [51] used the diagnostic ratio $WR/(RSG+YHG)$ (where $YHG \equiv$yellow hypergiant) to estimate the age of Wd1 as ∼$4.5 - 5$ Myr. We can now directly compare the BPASS models with the observed number count data to determine Wd1's age while allowing for a variable f_{bin}.

To obtain constraints on both quanitities, we need to use two star count ratios. When the grids of f_{bin} and age are projected into the ratio space, the ensuing topology can be complicated, making inferrence difficult. Thus it is critical to choose ratios such that the grid of parameters remains somewhat orthogonal (or at the very least, non-degenerate). The left panel of Figure 2.6 shows the predicted values for O/BSG vs. WR/RSG at solar metallicity for $\log t$ between 6.4 and 6.8 (2.5 and 6.3 Myr). The inset shows the grid of parameter values, with f_{bin} increasing from purple to yellow, and time increasing from dark to light grey. At the earliest times, the model grid is highly degenerate. However, in the latest time bins, different values of age and f_{bin} yield a large spread of possible ratio values. Using this diagram, one can plot the observed values of the two ratios, find the colors of lines that the point intersects, and use the inset to directly determine a binary fraction and age.

The data from Wd1 are indicated with the blue cross, which assumes $\sqrt{N}$ errors. From [44] and [51], we count 22 O stars, 29 BSGs, 24 WR stars, and 3 RSGs. We assume here that the sample in [44] and [51] is complete down to $\log(L) = 4.9$ for all subtypes but O stars, and that all O stars have been found. Note that the O star sample as reported consists mostly of supergiants. At the approximate age of Wd1, the main sequence turnoff is ∼30-40 $M_{\odot}$, implying there are still undetected main sequence O stars; we discuss the implications of an incomplete O star sample in the next section. It may be the case that some WRs and RSGs are obscured by dust, and thus that the sample of evolved stars is also incomplete. However, radio studies of Wd1 have failed to produce previously unknown dust-enshrouded members [e.g., 8, 74].

With no further assumptions made about the completeness of the data, we can infer that Wd1 has a high binary fraction of $f_{bin} \gtrsim 0.7$, and an age of 5-6.3 Myr. The age is consistent with previous studies, while the measured binary fraction is consistent with the results from [253]. This example highlights the importance of constructing ratios that are appropriate for the population under consideration. For a younger (∼3 Myr) cluster, the errors on the measurements of WR/RSG and O/BSG would have made any inferrence impossible. However, the grid covers a much larger area in ratio space at the latest time shown, implying WR/RSG vs. O/BSG is an even more sensitive metric at older ages.

We next consider the same ratios in a cluster with an age of approximately 10-20 Myr ($\log t \approx 7 - 7.3$) in Figure 2.7. We calculate WR/RSG and O/BSG for $\log t$ between 6.9 and 7.4 (8 and 25 Myr), and binary fractions between 0 and 1. A reference grid to aid in the interpretation of the plot is shown in the inset. Once again these ratios separate well in this space for all but the latest times shown, and span multiple orders of magnitude, implying that this is an incredibly useful diagnostic plot. Unfortunately, all of the WR stars that contribute to WR/RSG at this age are the heretofore mostly undiscovered products of binary evolution. As discussed previously, stripped "old" WR stars that evolve from lower-mass progenitors are incredibly rare and most may not be as spectroscopically obvious as their more massive single-star cousins. In a spectroscopic survey of $h + \chi$ Persei, a 13-14 Myr double cluster, [264] reached the main sequence down to a spectral type of A1, and found no low-mass WR candidates; however, not all of the bright, blue stars were observed. Future deep observing campaigns and theoretical work may yet reveal these stars; unfortunately, this plot remains unusable until then.

While these are only two examples using the same set of ratios, any combination of stellar types (and associated completeness limits, which we discuss further in §2.5), metallicites (within the set of metallicities modelled with BPASS), and f_{bin} can be used. We make all code that we wrote to generate these plots available online and provide additional examples at `https://github.com/tzdwi/Dia`

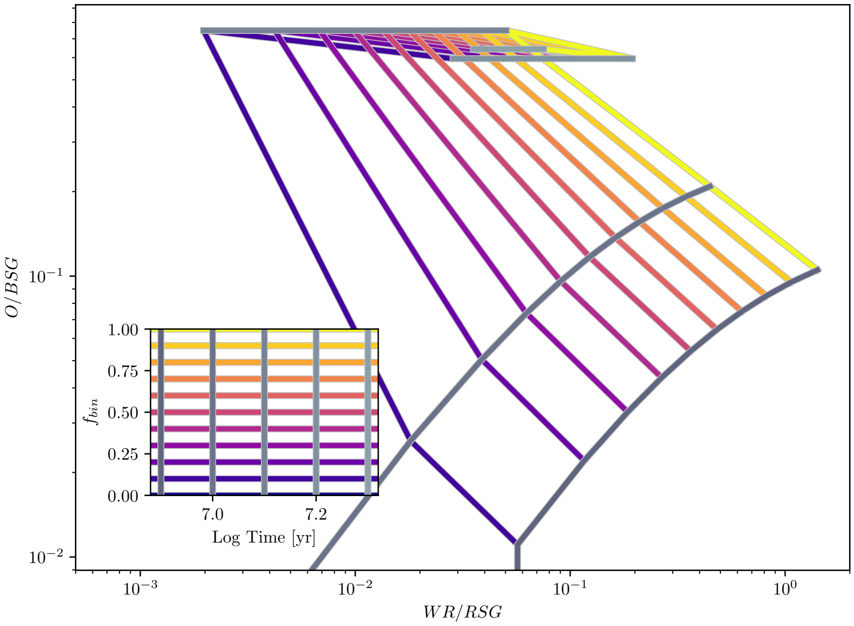

Figure 2.7 Similar diagnostic plot to Figure 2.6 for a theoretical ∼20 Myr cluster. Note: at this age, no single-star WRs are left, so the $f_{bin} = 0$ portion of the model isn't visible. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [68].

2.5 Accounting for Observational Completeness in Real Samples

Unfortunately, real surveys of Galactic and extragalactic populations are hindered by issues such as source confusion/crowding, inconsistent source classification, and incompleteness, the last of which we discuss here. Throughout this work, we have assumed that WR stars and supergiants were limited to $\log(L) \geq 4.9$, while stars on the main sequence could be found and classified with infinite precision.

Correcting for incompleteness in observed samples of massive stars is a difficult task that can introduce additional uncertainty into an already highly uncertain measurement. We can instead account for incompleteness in our synthetic populations by increasing the lower luminosity criteria for individual subtypes. Consider, for example, a spectroscopic survey that is only complete down

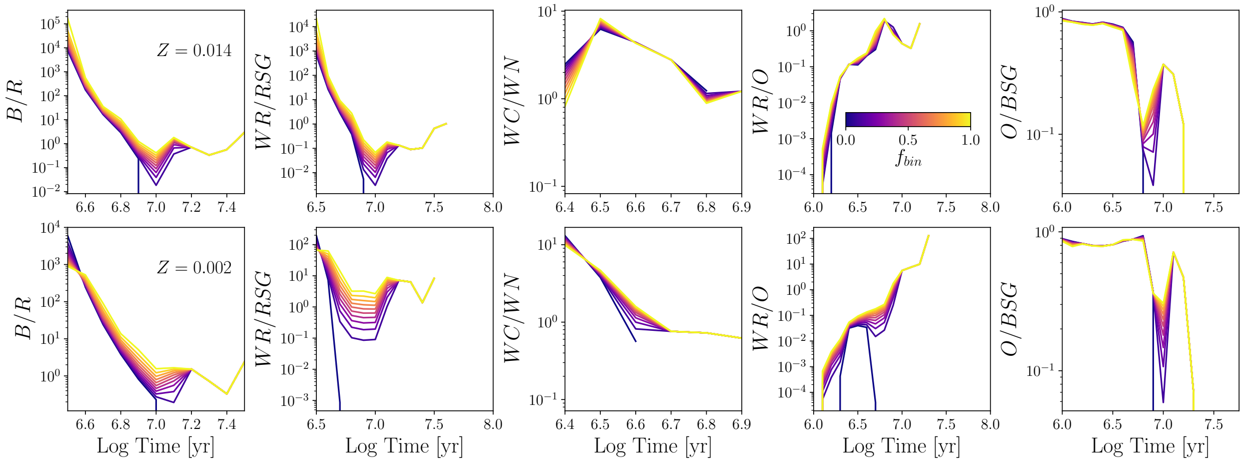

Figure 2.8 Identical to Figure 2.5, except assuming a minimum luminosity of $\log(L) = 5.2$ for all subtypes. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [68].

to $\log(L) = 5.2$, below which no stars are classified. Figure 2.8 shows identical ratios to Figure 2.5 after accounting for this completeness limit. Compared to Figure 2.5, almost all of the ratios become far less dependent on f_{bin}, especially at later times. This means that results obtained using ratios applied to incomplete samples will be less affected by binaries that weren't accounted for in the analysis. Conversely, incompleteness makes the task of simultaneously measuring age and f_{bin} much more difficult.

In practical terms, survey completeness is often expressed as a limiting magnitude in an optical bandpass, which can be transformed into a limiting absolute magnitude in that band. Thanks to blackbody physics, a magnitude limit like this corresponds to different bolometric luminosity limits for different spectral types. Therefore the above example is a toy model. In actuality, a survey will be more sensitive to bluer stars, and thus the exact luminosity limits should be carefully chosen to match the limits of the data.

As an example, we return again to Wd1. The catalog of [44] mostly does not include main sequence O stars; instead, most of the OBA stars they found belong to the cluster's supergiant cohort ($\log(L) \geq 4.9$). Using a combination of narrow- and broad-band imaging and spectroscopy to hunt for WR stars, [51] only find stars brighter than $M_{bol} = -8.2$ ($\log(L) \approx 5.18$). In the center panel of Figure 2.6, we again plot O/BSG vs. WR/RSG for our theoretical populations and for Wd1. However, in this figure, we impose a minimum WR luminosity of 5.1, and a minimum O luminosity of 4.9. These changes yield a model grid that predicts smaller values for both O/BSG and WR/RSG. The implied age of Wd1 is now slightly younger at 4-5 Myr, while the uncertainty in the observed ratios makes it impossible to measure a value of f_{bin}. We note that [51] made no such correction for completeness when using $WR/(RSG+YHG)$ to estimate the age of Wd1.

The right panel of Figure 2.6 is identical to the left and center panels, but assumes an overly conservative completeness limit of $\log(L) = 5.2$ for all species. The model grid predicts smaller values for both ratios, an age consistent with [51], and an upper limit for the binary fraction of $f_{bin} \lesssim 0.4$, which is inconsistent with current measurements of f_{bin}. This stresses the importance of choosing completeness limits that are consistent with the data, rather than relying on conservative assumptions.

Ultimately, completeness limits combined with uncertain values of f_{bin} can affect the theoretical values of star count ratios, hindering measurements of age or f_{bin}. Thus, great care must be taken to ensure that the Poisson noise of the measurement is smaller than the anticipated effect of binaries or incompleteness. For example, if precision of ~ 0.1 dex in a star count ratio $R = X/Y$ is required to measure a value of f_{bin} to within 0.1, and the expected value of the R is $R \approx 1$, the observer should find approximately 100 of each type of star in order to obtain the necessary precision, assuming they use $\sqrt{N}$ errors. Of course, the exact number of stars required changes with the age and metallicity of the population, but, as a rough rule of thumb, sample sizes of ~ 100s are necessary to make precision measurements. While only the most massive of galactic star clusters have the requisite number of stars, entire galaxies do have enough massive stars. We now turn our attention to calculating star count ratios in galaxies with complex star formation histories.

2.6 Ratios in Complex Star Formation Histories

While our discussion thusfar has been limited to SSPs, star count ratios are most frequently applied to entire galaxies, where they are used both as a benchmark of the physics implemented in a stellar evolution code and a test of our understanding of these physics. However, galaxies have complex

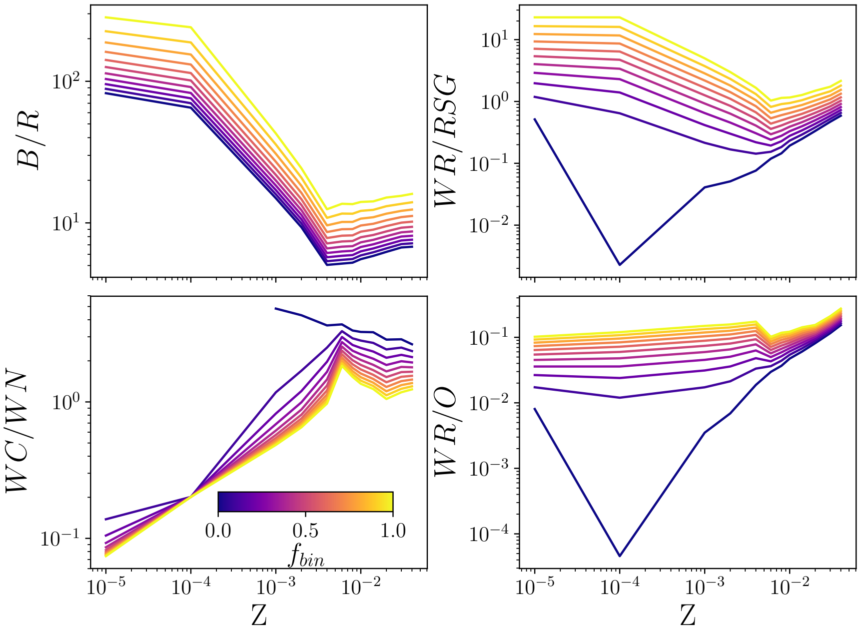

Figure 2.9 B/R (top left), WR/RSG (top right), WC/WN (bottom left), and WR/O for galaxies with constant star formation, and values of f_{bin} between 0 and 1; the color mapping for f_{bin} is identical to previous plots. A minimum luminosity of $\log(L) = 4.9$ is assumed for all subtypes. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [68].

star formation histories (SFHs); the galaxy as we see it today can be seen as the integrated set of countless populations formed between the onset of star formation and today, weighted by the star formation history, Ψ. We implement complex star formation histories in our code as follows. The number of a subtype S at time t for a SSP is $S(t)$. The total number seen in a population with star formation history $\Psi(t)$ is

$$S_{tot}(t_{i_{max}}) = \sum_{i=0}^{i_{max}} \Psi(t_i) S(t_i) \Delta t_i \tag{2.3}$$

where t_i is the i^{th} BPASS time bin with width Δt_i, and i_{max} corresponds to the total age of the population. Because BPASS uses 51 logarithmically spaced time bins from 1 Myr to 100 Gyr, Δt_i

is thus

$$\Delta t_i = \begin{cases} 10^{6.05} & i = 0 \\ 10^{6.15+0.1i} - 10^{6.05+0.1i} & 1 \leq i \leq 51 \end{cases} \tag{2.4}$$

We note that the definition of Ψ is such that $\Psi(t_0)$ is the star formation rate 1 Myr *ago* (i.e., the youngest BPASS time bin), not the star formation rate when the population is 1 Myr old. In the following, we adopt a constant star formation rate of $\Psi(t) = 1\ M_\odot$ yr^{-1} following [77]; however, our code allows for arbitrary SFHs, discretized to the default 51 age bins in BPASS. Note that, though we compute populations assuming constant SFR for all BPASS age bins, populations of massive stars are only sensitive to (at most) the previous 50-100 Myr of star formation, after which point the numbers of massive stars reach an equilibrium.[4] Because of this fact, the implicit assumption of constant metallicity for the entire population, while still unrealistic, is slightly more tenable.

Figure 2.9 shows WR/O, B/R, WR/RSG, and WC/WN as a function of metallicity in galaxies with constant SFR after allowing the massive star populations to reach equilibrium. We also implement a minimum luminosity of $\log(L) = 4.9$ through the remainder of this section, for consistency with extragalactic samples which are typically incomplete below this luminosity. For all four ratios plotted, introducing binary stars adds more than an order of magnitude spread in the predicted values for these ratios at most metallicities. This implies that inferences of metallicity, or star formation rate from number count ratios assuming only single star models are incorrect. However, the order-of-magnitude effect of binaries along with the improvement in statistics afforded by studying galaxies instead of star clusters implies that we can analyze extragalactic populations within the self-consistent framework of BPASS to make qualitative statements about binary populations.

Figure 2.10 shows B/R vs. WR/RSG and WR/O vs. WC/WN, with lines of constant f_{bin} and Z for populations with constant star formation, along with inset reference grids in a similar fashion to Figure 2.6. We also plot real data from the SMC (blue, $Z = 0.002$); LMC (orange,

---

[4] Indeed, the ionizing spectra of young populations in starburst galaxies reaches an equilibrium far earlier, at ~ 5 Myr [134].

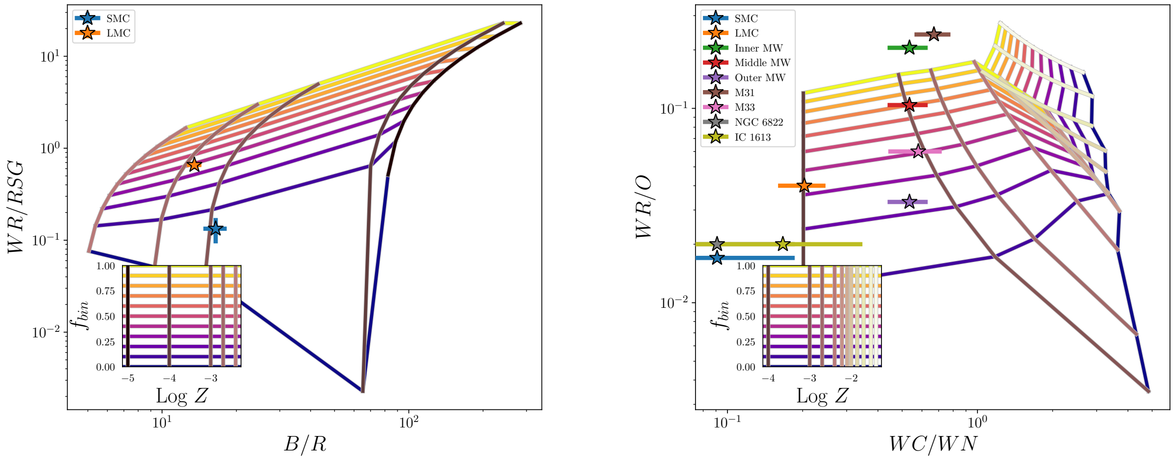

Figure 2.10 B/R vs. WR/RSG (left) and WR/O vs. WC/WN (right) for constant star formation populations with varying metallicity and f_{bin} (metallicity in the right plot is limited to $10^{-5} \leq Z \leq 0.004$). The insets are similar references grids to the plots in Figure 2.6. Observed values of these ratios (including errors where provided by the original authors) for various populations around the local group are plotted — note that WR/O is equal in NGC 6822 and IC 1613). B/R data are from [184]; WR/RSG data are from [184] for RSGs, [206] for SMC WRs, and [209] for LMC WRs; WR/O data are from [172]; and WC/WN data are from [244] for the Galactic WRs; [11] in NGC 6822 and IC 1613; and [206] and [208, 209] in the Magellanic Clouds, M31, and M33. All subtypes are assumed to have a minimum luminosity $\log(L) = 4.9$, including O stars, to account for incompleteness in extragalactic studies. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [68].

Table 2.3. Raw counts of massive stars used to calculate ratios and errors, as well as the value of WR/O from [172].

Galaxy	BSG	RSG	WR	WC	WN	WR/O
SMC	1484	90	12	1	11	0.017
LMC	3164	234	154	26	128	0.04
Inner MW	—	—	—	—	—	0.205
Middle MW	—	—	—	46[a]	86[a]	0.104
Outer MW	—	—	—	—	—	0.033
M31	—	—	—	62	92	0.24
M33	—	—	—	26	45	0.06
NGC 6822	—	—	—	1	11	0.02
IC 1613	—	—	—	1	6	0.02

Note. —

$Z = 0.006$); solar neighborhood stars with galactocentric distances of $6 \le R_{GC} < 7.5$ kpc (green, supersolar metallicity), $7.5 \le R_{GC} < 9$ kpc (red, $Z \approx Z_\odot$), and $9 \le R_{GC} < 11$ kpc (purple, subsolar metallicity); M31 (brown, $Z \approx 2Z_\odot$); the inner region of M33 (pink, $Z \approx Z_\odot$); NGC 6822 (yellow, $Z = 0.005$); and IC 1613 (grey, $Z = 0.002$). B/R data are from [184] and WR/RSG data are compiled from [184], [206], and [208]. WR/O data are all from [172], who do not report raw numbers to allow us to estimate an error. WC/WN data are from [244] in the Milky Way (where we only use their data from the solar circle to compare with the smaller annuli from [172]); [11] in NGC 6822 and IC 1613; and [206] and [208, 209] in the Magellanic Clouds, M31, and M33. [208] report both $\sqrt{N}$ and rigorous asymmetric errors based on their completeness limits. For consistency between varying sources of data we only use the $\sqrt{N}$ errors, but note that quantitative analyses should adopt rigorous error calculations. We tabulate the raw numbers used to calculate the ratios and errors, as well as the value of WR/O from [172] in Table 2.3.

The first thing that we see is that f_{bin} and Z projected into the ratio spaces are well separated across multiple orders of magnitude, and are roughly orthogonal; i.e., WR/RSG and WR/O are good tracers of f_{bin}, while B/R and WC/WN trace metallicity. Thus, changes in the physics in BPASS that only affect the lifetime of one evolutionary phase (e.g., an implementation of meridional circulation, which would increase the BSG lifetime, [77]), will predominantly change the absolute inferred f_{bin} of Z; the relative values inferred are still valid.

When compared to data, we see that the metallicity that one might infer based on these grids is systematically lower than the true metallicity of the galaxies (e.g., on the left-hand plot, one might assume that $Z_{SMC} \approx 0.001$ and $Z_{LMC} \approx 0.002$; the right-hand plot shows an even worse correspondence). This may be due to at least one of a few possibilities:

1. The data are subject to inconsistent classification and difficult-to-quantify completeness-limits.

2. The completeness of these samples extends to lower luminosities than are assumed in the model grids. For example, a lower luminosity cutoff in the models would add more WNs than WCs (because lower luminosity WRs will be lower mass, and thus more likely to be WNs), and more Os than WRs (due to the relative lifetimes and abundances of both subtypes), shifting the grid down and to the left in the right hand plot.

3. The mass loss through both single- and binary- star channels predicted by BPASS is incorrect. Mass loss will govern the age at which massive stars transition between various evolutionary stages, and WC/WN and B/R are both sensitive to these lifetimes.

4. The effects of rotation included in BPASS are incorrect. BPASS only uses approximate physics to simulate rotation. Including e.g., meridional circulation, would increase the number of BSGs, which would then increase the metallicity inferred by B/R. Future work will examine the effects of rotation using stellar evolution codes that adopt more detailed treatments of rotational phenomena.

Assuming that, while the absolute values of metallicity or f_{bin} or Z inferred from these plots will change depending on the exact physical prescriptions, the relative values will remain largely consistent, we see that f_{bin} in all of these galaxies appears to increase with actual metallicity. If WR/O is indeed a good tracer of f_{bin}, this should also be apparent when plotting WR/O vs. Z. Figure 2.11 shows the values of WR/O and Z as listed in Table 6 of [172], along with a linear fit to $\log{(WR/O)}$ vs. $Z_{\rm actual}$. Note that we use an asterisk in the x-axis label because the metallicities reported in [172] assume $Z_{\odot} = 0.02$, instead of 0.014 [13]. The fit clearly shows the trend that WR/O, and thus f_{bin}, increases with Z (under the assumptions above).

While more work is necessary both in the BPASS models and in observational studies in order to quanitify the exact relationship between f_{bin} and Z, this is still a very intriguing finding. A correlation between f_{bin} and Z for massive stars has not previously been reported. [236] considered low-mass stars, and found that metal-poor main sequence stars with $0.625 \leq B - V \leq 1.0$ were more likely to have stellar companions at a confidence of 2.4σ. If the binary fraction instead increases with metallicity for more massive stars, this would point to a fundamental difference in how stellar binaries are formed in different mass regimes. Indeed, this result makes intuitive sense; increased metal line cooling in a molecular cloud would make the cloud cooler and denser, while carrying away none of the angular momentum and making the formation of binaries more likely. We note that this argument applies only to the formation of binary systems, and not their *evolution*. Indeed, at lower metallicity, binaries are expected to be *found* at closer orbital separations due to weaker stellar winds being less efficient at losing angular momentum, thus increasing orbital separations [66, 165]. BPASS does include these physics, but this result is merely a statement about the natal, intrinsic value of f_{bin}.

It is important to note that there are caveats to this result: if the binary fraction is dependent on metallicity, then it is also likely that the period and mass ratio distributions are as well, while the BPASS inputs implicitly assume only the binary parameter distributions from [198]. Future work will include varying the period and mass ratio distributions and considering any potential dependence on metallicity. We also make no attempt at quantifying this putative relationship between f_{bin} and Z, as such a result is not the focus of this work, and will require dedicated observational and theoretical study.

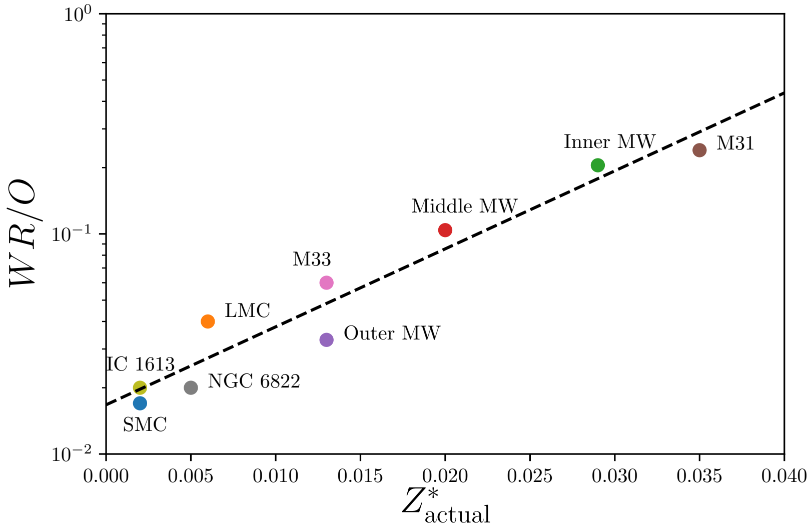

Figure 2.11 WR/O vs. actual metallicity for galaxies in the local group. All data are from [172]; those authors assumed $Z_\odot = 0.02$. For self-consistency, we only report Z as listed in Table 6 of [172]; thus we put an asterisk in the x label. The dashed black line is a linear fit to $\log{(WR/O)}$ vs. $Z^*_{\rm actual}$. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [68].

2.7 Summary and Conclusion

We summarize our results as follows:

- When applied to SSPs, accounting for binary effects when using star count ratios is incredibly important. If the binary fraction isn't known, an order-of-magnitude spread can be introduced into the theoretical prediction.

- We find including binaries and imposing a completeness limit can both introduce $\gtrsim 0.1$ dex changes in the inferred $\log(\mathrm{age}/\mathrm{Myr})$ of star clusters.

- Similar incompleteness and binary effects can manifest themselves in more complex systems of stars. However, because star count ratios can be subject to fewer systematics and small-number statistics, proper treatment of stellar binaries can yield interesting results; for example, the existing measurements of the relative numbers of WRs and RSGs around the Local Group are now, as of this year, of sufficient quality to infer the relative importance of binary interactions — subject to modelling uncertainties as well as poorly constrained recent star formation histories — over more than an order of magnitude in metallicity [189].

- Combinations of star count ratios can be used to indirectly measure the massive star binary fraction, provided the data have well understood completeness limits and assumed errors. This method works in large or distant populations where direct measurement of f_{bin} via spectroscopic studies of individual stars is otherwise impossible. Where direct measurements are also possible, this method can also illuminate where the BPASS stellar evolution physics can be improved.

Chapter 3

A COMPARISON OF ROTATING AND BINARY STELLAR EVOLUTION MODELS: EFFECTS ON MASSIVE STAR POPULATIONS

3.1 Introduction

As discussed in Chapter 1, rotation is a ubiquitous property of stars, affecting their evolution through a variety of means, mostly via enhanced mixing and mass loss [329, 93]. The Geneva stellar evolution code [76, 90, 104] is the current state of the art implementation of these effects in massive stars. Simultaneously, many massive stars are born into binary and higher-order systems [253, 75, 254, 198], and binary systems with dwarf, supergiant, Wolf-Rayet, and compact components are commonly observed [252, 207, 205, 211]. Interactions in binary systems through both tides [120] and mass transfer can completely disrupt the evolution of both stars, creating otherwise-impossible evolutionary states, affecting populations of massive stars, especially at low metallicity [277]. The Binary Population and Spectral Synthesis code (BPASS, [77, 276]) is the current state of the art rigorous implementation of these effects.

Despite the importance of binary interactions and rotation, the theory of single, nonrotating stars has thusfar been successful at predicting the evolution of stars of varying compositions and initial masses (the “Conti scenario”, [49]). The addition of the effects of rotation has further refined single-star evolution to reproduce observed surface abundance enhancements and mass loss rates [e.g., 195]. These successes are often used as evidence that binary interactions are a secondary effect, or can be approximated by simple methods [e.g., by instantaneously removing a model star’s H envelope after it leaves the main sequence; for a review of recent stellar evolution models featuring simple implementations of binary interactions, see 77, and citations therein]. Indeed, at first glance some of the effects predicted by binary stellar evolution models (such as enhanced surface abundances, stronger mass loss, broader mass and age ranges for Wolf-Rayet progenitors, and harder ionizing spectra produced by stellar populations) are identical to the predictions from rotating stellar models (e.g. [160]). Regardless, massive binary systems *do* exist, necessitating detailed modeling of their evolution as well as observations of the frequency of binary systems to understand their impact on stellar populations.

Finding binary systems via radial velocity variability (e.g., [253]; [252]) requires long time baselines, high spectral resolution, and long integration times, and is currently infeasible for stars beyond the Magellanic Clouds (MCs), or for some stars such as red supergiant binaries [211]. For more-distant systems, we are left studying (semi-)resolved stellar populations. Simultaneously, measuring rotation periods requires high-cadence time-series photometry [24, 32, 18, 17, 125, 237, 229]. Spectroscopic measurements yield projected rotational velocities, with the caveat that the inclination of the rotational axis to the line of sight is unknown [e.g., 114]. In either case, both methods again require high quality observations of nearby stars. Fortunately, massive stars are luminous, and can be resolved in galaxies around the Local Group and beyond. Previously (Chapter 2), we used the BPASS models to construct grids of synthetic populations with varying metallicity (Z), star formation history (SFH), and the natal binary fraction (f_{bin}), and predicted the frequency of various evolutionary phases. Here, we incorporate stellar evolution models that include rotation [76, 90, 104], and perform a detailed comparison between the two model sets. This work originally appeared in Dorn-Wallenstein et al. (2020a) [70].

In §3.2, we compare the two evolutionary codes, from the BPASS and Geneva groups, used to generate binary and rotating stellar tracks respectively. We also detail the population synthesis method we employ to generate theoretical rotating and binary stellar populations and subsequent predictions for the frequency of various spectral types in these populations. In §3.3 we describe the observables we generate from the two synthetic populations, and introduce a novel Bayesian framework of estimating these observables from data. We compare our synthetic rotating and binary populations to each other and to observations of massive star populations with both simple and complex star-formation histories before considering the results of our comparison and the implications for using rotating and binary stellar evolution models to interpret future observations

Table 3.1. Model parameters used to label model timesteps with an evolutionary phase, and other variables introduced in the text.

Parameter	Description	Unit
$\log(L)$	Logarithm of the luminosity	$L_\odot$
$\log(T_{eff})$	Logarithm of the effective temperature	K
$\log(g)$	Logarithm of the surface gravity	cm s^{-2}
X	Hydrogen Surface Mass Fraction	-
Y	Helium Surface Mass Fraction	-
C	Carbon Surface Mass Fraction (Sum of ^{12}C and ^{13}C for Geneva tracks)	-
N	Nitrogen Surface Mass Fraction	-
O	Oxygen Surface Mass Fraction	-
$\log t$	Logarithm of time	yr
Z	Mass fraction metals, $Z_\odot = 0.014$	-
f_{bin}	Binary fraction	-
f_{rot}	Rotating fraction	-
f	Generic term for either f_{bin} or f_{rot}	-
n_S	Observed frequency of an arbitrary spectral type S	-
R_{S_1/S_2}	Observed ratio of the frequency of two spectral types, S$_1$ and S$_2$	-
$\hat{n}_S$	Intrinsic frequency of an arbitrary spectral type S	-
$\hat{R}_{S_1/S_2}$	Intrinsic ratio of the frequency of two spectral types, defined as $\hat{R}_{S_1/S_2} \equiv \hat{n}_{S_1}/\hat{n}_{S_2}$	-

of massive star populations (§3.4).

3.2 Creating Theoretical Populations

3.2.1 The Models

In Chapter 2, we created synthetic populations using BPASS version 2.2.1, which incorporates the effects of both tides and mass transfer to predict the evolution of single and binary stars on a dense grid of initial primary and secondary masses (M_1 and M_2), initial periods P, and mass ratios ($q \equiv M_2/M1$) at 13 metallicities. We express the metallicity as a mass fraction Z, and BPASS adopts metallicities in the range $10^{-5} \leq Z \leq 0.04$. Note that as in Chapter 2, we assume solar metallicity $Z_\odot = 0.014$ [13]. Binary interactions are modeled as enhanced mass loss/gain from/onto its model stars via RLOF. The orbital energy is tracked throughout, allowing BPASS to model a broad range of simulated evolutionary scenarios.

Here, we also incorporate evolutionary tracks generated by the Geneva code which is the current state of the art implementation of rotating stellar evolutionary models, and has been used to study the impact of rotation on the radiative output [160], chemical yields [112], and final fates [194] of massive stars. Stellar evolution is modeled at two different initial rotation rates (nonrotating, with $v_{ini}/v_{crit} = 0$, and rotating, with $v_{ini}/v_{crit} = 0.4$) and at three different metallicities: $Z = 0.014$ [76], $Z = 0.002$ [90], and $Z = 0.0004$ [104]. Horizontal diffusion coefficients in the rotating models are calculated following [329]. Meridional circulation is calculated as described by [176], and the two effects are combined following [35]. Angular momentum transport is included, and angular momentum is conserved following [93]. Finally, rotation-enhanced radiative [174] and mechanical [93] mass loss are also implemented; for a detailed disucssion see [76].

3.2.2 Population Synthesis

Evolutionary tracks in hand, we synthesize a population by weighing each track with the initial mass function, $\Phi(M)$. We adopt the default form of Φ in BPASS v2.2.1, which is a broken power law with slope -1.3 below 0.5 $M_\odot$, and a slope of -2.35 for higher masses, with a minimum mass of 0.1 $M_\odot$ and a maximum mass of 300 $M_\odot$, normalized so the total stellar mass is $10^6 M_\odot$. Binary models are also weighted according to the distributions of the fundamental natal period P and mass ratio q from [198]. Because the Geneva tracks are sampled on a much coarser grid of initial mass than the BPASS single star tracks, we linearly interpolate the available Geneva tracks onto the BPASS single star initial mass grid following [92], and adopt the IMF weighting from the BPASS v2.2.1 inputs[1]. Additionally, both rotating and nonrotating tracks at all three metallicities are only available between 1.7 and 120 $M_\odot$. We choose not to introduce any correction factors to the IMF weights — e.g., boosting the weight of the 120 $M_\odot$ model to represent all stars with initial masses between 120 and 300 $M_\odot$ — to ensure that tracks with identical masses are weighted identically. Because no single star below the 1.7 $M_\odot$ threshold would become any of the stellar types considered here, only the exclusion of these very massive stars would affect our results. However, these stars are so rare, and their lifetimes so short, that their impact on our synthetic populations is minimal.

We create four sets of synthetic populations: one composed entirely of single, nonrotating stars using the input files provided in the BPASS v2.2 data release ($f_{bin} = 0$), one composed entirely of binary stars ($f_{bin} = 1$) using custom input files provided by the BPASS team (J. J. Eldridge 2018, private communication), one composed of single, nonrotating stars from the Geneva models (hereafter referred to as the population with "rotating fraction" $f_{rot} = 0$), and one composed entirely of rotating stars ($f_{rot} = 1$). We note that there are no single stars in the custom $f_{bin} = 1$ population, though the distribution of periods and mass ratios is identical to the default BPASS v2.2 binary population, which is drawn from [198].

3.2.3 Number Counts vs. Time

Photometric surveys of nearby massive stars [e.g., the Local Group Galaxy Survey, LGGS; 186, 185], can yield fairly complete catalogs after filtering for foreground contaminants [e.g., 188], and follow-up narrow-band surveys can be used to find evolved emission line stars (e.g., [206,

---

[1] Details on the mass grids, parameter values and more can be found in the BPASS v2.2 User Manual, currently hosted online at `bpass.auckland.ac.nz`

Table 3.2. Criteria used to classify evolution tracks, similar to Table 2.2. Adapted from Table 3 of [77]. We specify where the Geneva tracks are classified with different criteria than the BPASS tracks.

Label	BPASS Criteria	Geneva Criteria
WNH	$\log(T_{eff}) \geq 4.45$ $X \leq 0.4$	$\log(T_{eff}) \geq 4.0$ $X \leq 0.3$
WN	$\log(T_{eff}) \geq 4.45$ $X \leq 10^{-3}$ $(C+O))/Y \leq 0.03$	$\log(T_{eff}) \geq 4.0$ $X \leq 10^{-5}$ $N > C$
WC	$\log(T_{eff}) \geq 4.45$ $X \leq 10^{-3}$ $(C+O))/Y > 0.03$	$\log(T_{eff}) \geq 4.0$ $X \leq 10^{-5}$ $N \leq C$
O	$\log(T_{eff}) \geq 4.48$	$\log(T_{eff}) \geq 4.5$
Of	$\log(T_{eff}) \geq 4.519$ $\log(g) > 3.676 \log(T_{eff}) + 13.253$	
B	$4.041 \leq \log(T_{eff}) < 4.48$	$4.041 \leq \log(T_{eff}) < 4.5$
A	$3.9 \leq \log(T_{eff}) < 4.041$	$3.8 \leq \log(T_{eff}) < 4.041$
F/G	$3.66 \leq \log(T_{eff}) < 3.9$	$3.66 \leq \log(T_{eff}) < 3.8$
K	$3.55 \leq \log(T_{eff}) < 3.66$	
M	$\log(T_{eff}) < 3.55$	
BSG	O + Of + B + A $\log(L) \geq 4.9$	
YSG	F/G $\log(L) \geq 4.9$	
RSG	K + M $\log(L) \geq 4.9$	
WR	WNH + WN + WC $\log(L) \geq 4.9$	

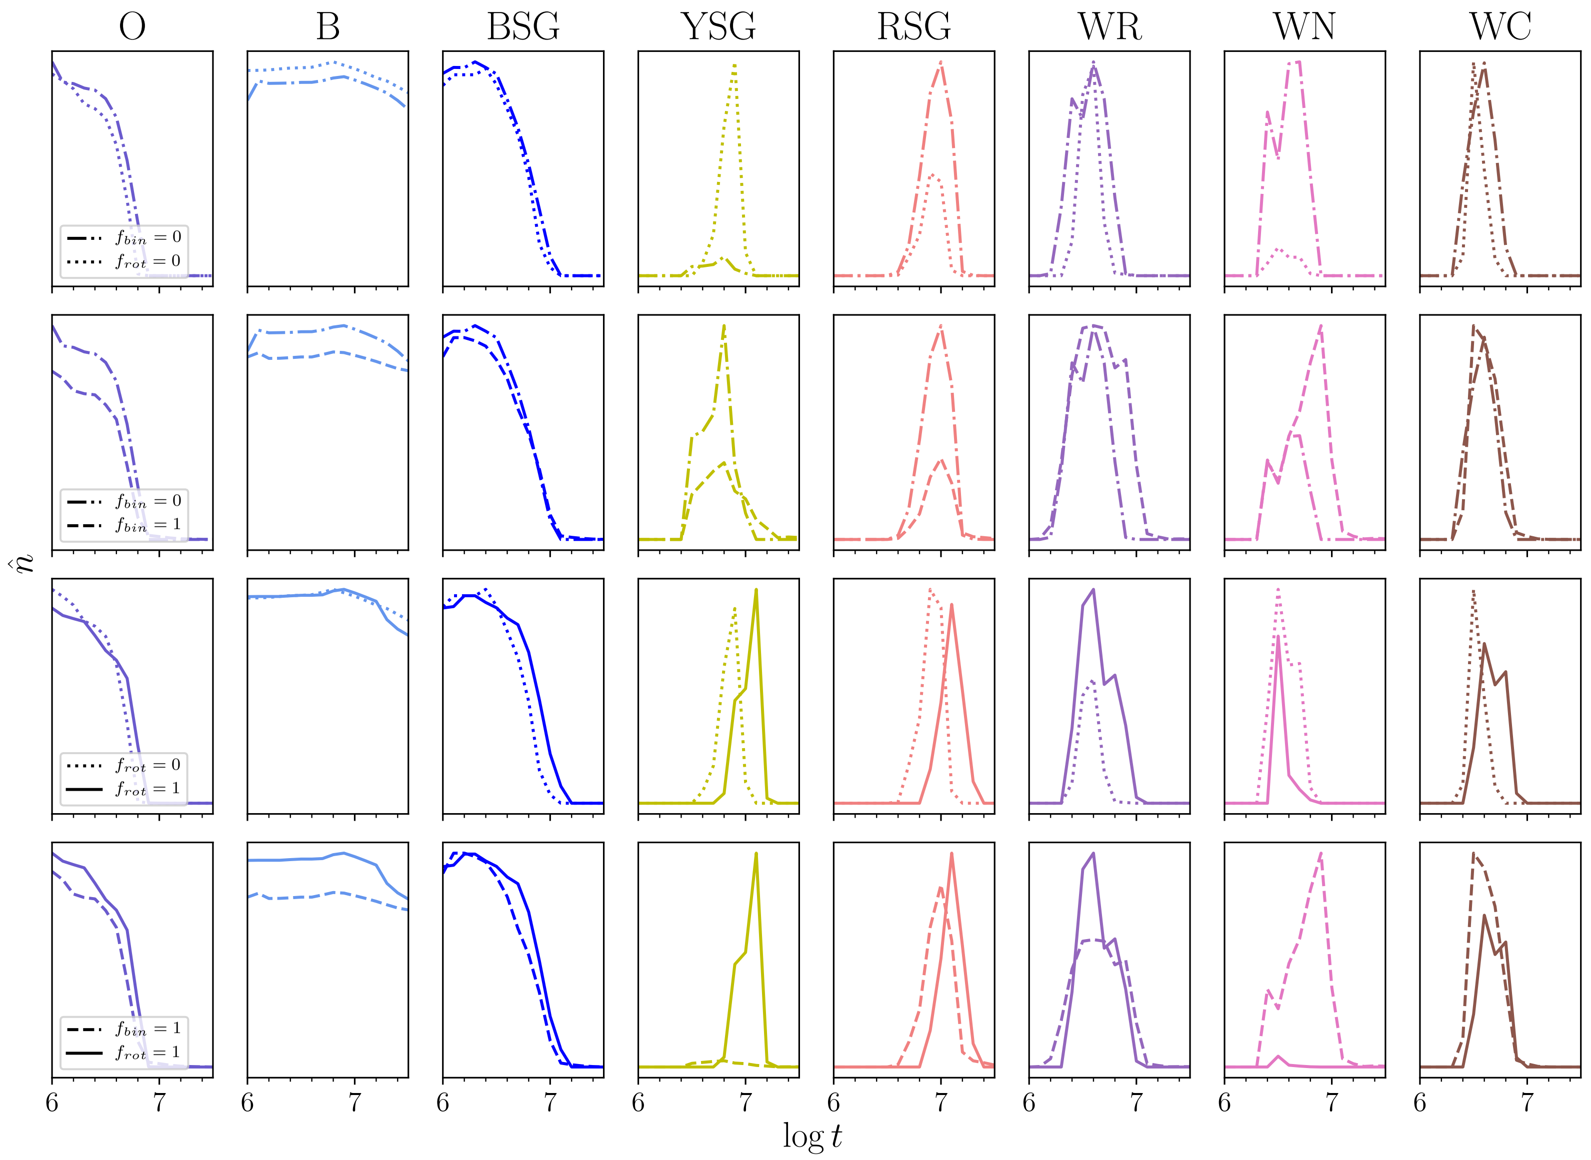

Figure 3.1 Number of various stellar subtypes at $Z = 0.014$ between 10^6 and $10^{7.5}$ years. The top row compares the $f_{bin} = 0$ population (dash-dotted) with the $f_{rot} = 0$ population (dotted), the second row compares the $f_{bin} = 1$ (dashed) and $f_{bin} = 0$ (dash-dotted) populations, the third row compares the $f_{rot} = 1$ (solid) and $f_{rot} = 0$ (dotted) populations, and the bottom row compares the $f_{bin} = 1$ (dashed) population with the $f_{rot} = 1$ (solid) population. Both lines in a given panel are plotted on the same linear y-scale to allow for comparison.

208, 209]). Even without follow-up spectroscopy, photometric measurements can then be used to categorize stars into broad spectral types. Thus it is useful to classify all timesteps of a given evolutionary track into one of a number of coarse spectral types, using the position of the model timestep on the HR diagram, as well as its surface composition, based on the model parameters listed in Table 3.1; we largely adapt the classification scheme from [77], shown in Table 3.2. [90] use a similar classification scheme, with slightly different temperature or composition thresholds. All labels refer explicitly to spectral types. WNH corresponds to Hydrogen-rich Wolf-Rayet stars; Of stars are O stars with particularly strong winds and He II emission [31]. The BSG, YSG, RSG, and WR numbers are computed by summing the numbers of the indicated species, and applying the relevant luminosity threshold, as described below.

Where they are different, we adopt the Geneva criteria to classify the Geneva tracks, and use the BPASS criteria in cases where the BPASS classification is more specific than the Geneva classification (e.g., BPASS distinguishes between K and M stars). This choice is motivated by two facts. Firstly, in some cases, these criteria are used within the individual evolutionary tracks to distinguish between different prescriptions for, e.g., mass-loss. Secondly, the coupling between the outermost layer of a model star and a model stellar atmosphere in order to produce a synthetic spectrum is non-trivial. Indeed, the criteria for classifying WR stars have nothing to do with the mass-loss rate, which might be observed from such a spectrum. In lieu of synthesizing a spectrum for each timestep of each model, we defer to the creators of each code in how to best classify their models. However, it should be noted that the exact choice of the values presented in Table 3.2 do not drastically affect the results [196]. When we attempted to classify the Geneva tracks using the BPASS criteria, we found little substantive changes, except the rotating Geneva models produce more WN stars (still less than half the WNs produced by BPASS).

We then use the classification, age, and weight of each track to find the frequency of each spectral type in Table 3.2 as a function of time, binned to 51 time bins that are logarithmically-spaced between 10^6 and 10^{11} years in 0.1 dex increments as described in Chapter 2. We also assign each model to one of 31 luminosity bins with 0.1 dex width between $\log(L) = 3$ and 6. Thus we can apply coarse luminosity thresholds to mimic observational completeness limits. This allows us to account for unresolved binaries by making the assumption that all stars below a given luminosity, including secondaries, are not detected, while all stars above this threshold are. While a somewhat simplistic assumption, the evolved stages of massive star evolution are so short-lived that most evolved massive stars have main sequence companions [e.g., 205, 211]. Of course, binaries with two evolved components do exist (e.g., WR+WR binaries), but are usually detectable via their wind interactions. In the case that spectral observations are of insufficient SNR or resolution to classify the components of such a binary, the WC/WN ratio would be unreliable. We encourage observational efforts dedicated to making a census of the massive component of stellar populations to discuss their insensitivity to unresolved binaries.

Each column in Figure 3.1 shows the number, $\hat{n}$, of the spectral types in Table 3.2. For clarity, values on the y-axis are not shown, but both lines in a given panel are plotted on the same, linear scale. Line styles are used to indicate the four different populations, with dash-dotted lines for $f_{bin} = 0$, dashed for $f_{bin} = 1$, dotted for $f_{rot} = 0$, and solid for $f_{rot} = 1$. The top row corresponds to the single, nonrotating populations from both evolutionary codes, the second row shows the $f_{bin} = 0$ and $f_{bin} = 1$ populations, the third row shows the $f_{rot} = 0$ and $f_{rot} = 1$ populations, and the bottom row shows the $f_{bin} = 1$ and $f_{rot} = 1$ populations. Note that, as in Chapter 2, we do not include Luminous Blue Variables in our analysis. This is due to the present uncertainty in the evolutionary status of LBVs [268, 118, 1], and the lack of a clear consensus in how to observationally classify a statistically significant number of them without long-term monitoring.

3.2.4 Diagnostic Ratios

Because the calculated number counts are the frequency of each subtype per $10^6 M_\odot$ of stars formed, direct application to observed populations requires an estimate of the stellar mass of a population, M_*. Such measurements are often model dependent, and are based on inferences of the sometimes-undetected low-mass end of the population. Instead, ratios of the frequency of these types (hereafter "number count ratios") are independent of the stellar mass, while remaining sensitive to both rotation and binary interactions.

We first construct the predicted number counts for subtypes in a population with a given f_{bin} or

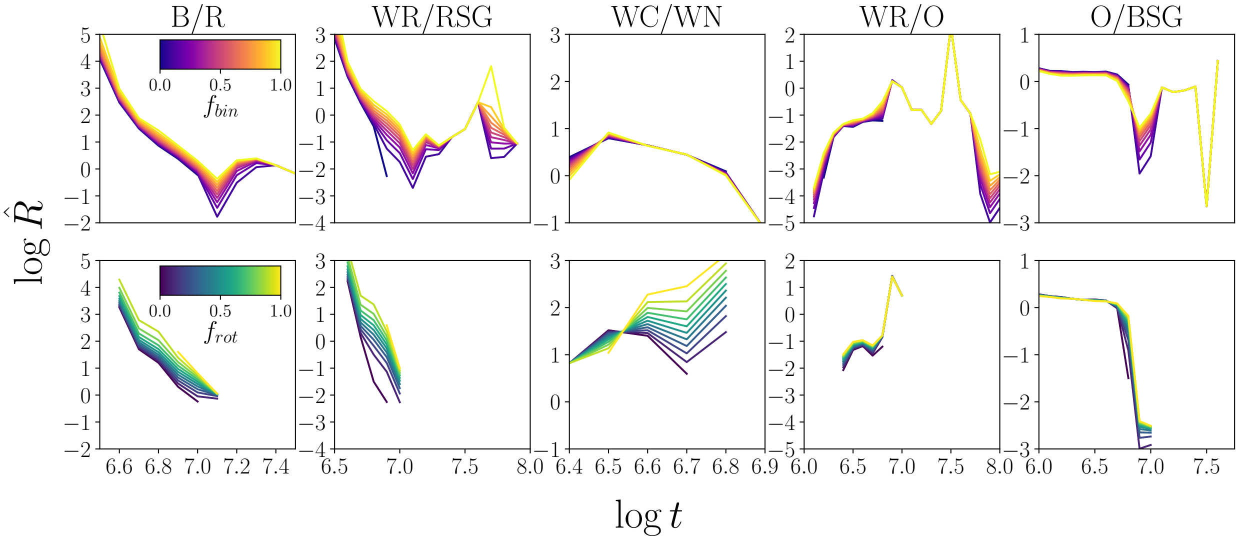

Figure 3.2 Predicted values of five diagnostic ratios (where B/R is shorthand for the ratio of blue to red supergiants), at solar metallicity, for the BPASS (top) and Geneva (bottom) populations as a function of time, for values of f between 0 and 1, as indicated by the colorbar. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [70].

f_{rot} (generically f hereafter). All notation used is summarized in Table 3.1. We calculate $\hat{n}_S$, the frequency of a subtype S at time t and metallicity Z as

$$\hat{n}_S(t, f, Z) = f\hat{n}_{S_{f=1}}(t, Z) + (1 - f)\hat{n}_{S_{f=0}}(t, Z) \tag{3.1}$$

where $\hat{n}_{S_{f=0}}$ and $\hat{n}_{S_{f=1}}$ are the frequencies in the $f = 0$ and $f = 1$ populations respectively. We note that, for the BPASS populations, the naive interpretation of f_{bin} is rather straightforward: f_{bin} is the fraction of binary stars in the population[2]. However, for the Geneva populations, f_{rot} would then correspond to the fraction of stars in the population born with $v_{ini}/v_{crit} = 0.4$ (i.e., rapid rotators), while realistic stars rotate with a range of velocities, from nonrotating to almost critically rotating; another interpretation might then be for f_{rot} to correspond to the average initial rotation rate (i.e., $\langle v_{ini}/v_{crit}\rangle = 0.4f_{rot}$, see, for example, [160]). Ultimately, in both cases, f is simply

---

[2] There is a secondary factor here, in that f_{bin} is mass-dependent [75]. However, all of the evolutionary phases here are mostly descended from O and early B stars (at least in the single-star paradigm), for which observed samples are too small to determine any mass dependence.

a factor used to linearly combine the *output number counts* from each population, and does not necessarily correspond to the fraction of binary/rotating star models that are used, a fact that the reader should be aware of when interpreting our results.

We then calculate number count ratios using

$$\hat{R}_{S_1/S_2}(t, f, Z) = \hat{n}_{S_1}(t, f, Z)/\hat{n}_{S_2}(t, f, z) \tag{3.2}$$

for two subtypes S_1 and S_2.

In Chapter 2 we described four number count ratios frequently found in the literature, as well as a novel number count ratio, O/BSG. Here we briefly describe these ratios, and the physical effects that they probe. We note that, in theory, our model populations would allow us to perform a search for the ratios and completeness limits that would best differentiate between different channels of stellar evolution. However the existing data tends to focus only on individual species, making the available space of ratios that can be measured quite small. In future work, we plan to perform this search in order to guide observers to evolutionary species for which an accurate census is most useful for constraining stellar evolution.

- *B/R*: The ratio of the number of BSGs to RSGs (B/R) is among the most frequently used ratios, and has a long history as a metallicity diagnostic [318, 301, 152, 182]. It is sensitive to the physics governing the timescale of rightward evolution of stars on the HR diagram — rotational/convective mixing — as well as the interruption of a star's expansion by RLOF.

- *WR/RSG*: In the single star paradigm, the WR/RSG ratio probes the boundary between stars that experience only redward evolution and stars that lose enough mass to evolve blueward at the end of their lives [49], and was therefore thought to be sensitive to metallicity [171] until quite recently [189]. Mass loss via binary channels serves to artificially boost this ratio by decreasing the number of RSGs, and commensurately increasing the number of WRs.

- *WC/WN*: The WC/WN ratio probes the evolution of stars that have already lost enough mass to become WRs. Thus it is mostly insensitive to the binary fraction, and is a very sensitive diagnostic of radiative mass loss in the WR phase [311, 110].

- *WR/O*: The WR/O ratio probes the largest swatch of the mass spectrum considered here. As the only ratio in the literature with main sequence components, it is the most subject to contamination by unresolved O+O binaries [170] that is difficult to address via our simple implementation of completeness limits.

- *O/BSG*: In Chapter 2, we introduced the O/BSG ratio. Both species are recovered by photometric censuses of bright blue stars in stellar populations, and don't require narrow band imaging or spectroscopic follow-up to detect. While, in theory, it is mostly sensitive to the main sequence lifetime (and thus rotational and convective mixing), some main sequence O stars are luminous enough to be classified as BSGs in our scheme. As we demonstrated, O/BSG is mostly insensitive to the binary fraction, except in a narrow window around $\log t = 7$. This is due to lower mass stars losing enough of their envelopes via RLOF to evolve blueward, without losing enough Hydrogen to be labelled as WR stars — observationally these might be classified as sdB/O stars, which are not counted separately from O and B stars in our classification scheme (see Section 3.2 and Figure 10 of [77]). This effect boosts the O/BSG ratio relative to the single star population until all of the O stars have evolved. Thereafter, only small numbers of models in the binary population reside within the Luminosity-Temperature-Composition boundary of O stars or BSGs at different times, causing rapid changes to O/BSG (for a particularly drastic example of a star rapidly entering and leaving the O and BSG regimes, see Figure 1 of [68]).

Figure 3.2 shows the values of five different ratios vs. time for solar metallicity BPASS (top row) and Geneva (bottom row) populations with $0 \leq f \leq 1$ as indicated by the colorbar. Panels in the same column have identical bounds on the vertical axis for comparison. The bounds of the time axis have been chosen to highlight the time range during which each ratio is most dependent on f.

3.3 Results

3.3.1 Comparing Rotating and Binary Model Populations

The differences between the single/nonrotating O and B stars in the top row of Figure 3.1 are minimal, where the dotted line shows the nonrotating Geneva population, and the dash-dotted line shows the BPASS single-star population. We warn the reader that, in general, the BPASS user manual cautions against using the $f_{bin} = 0$ population in isolation; only the $f_{bin} > 0$ populations can be considered reliable. However, we show the $f_{bin} = 0$ results for completeness, and discuss the differences between the two $f = 0$ populations in detail here and throughout the text.

Significant differences arise in the yellow supergiant phase, where the Geneva models predict the existence of far more YSGs. However, theoretical uncertainty in this very short-lived phase has long stymied our understanding of massive star evolution [137]; for this reasons we do not use the YSG phase in our subsequent ratio diagnostics and caution against using it as a diagnostic of stellar population properties until it is better understood. The single BPASS models produce approximately twice as many RSGs as the nonrotating Geneva models. This is largely due to the fact that the BPASS models cross the HR diagram quicker (reflected in the significantly smaller number of YSGs compared to the Geneva models), increasing the amount of time the stars spend as RSGs before ending their lives. However, the two model sets also adopt slightly different mass loss prescriptions during the RSG phase: the Geneva tracks use mass loss rates from [239, 240] for the models less massive than 12 $M_{\odot}$ during the RSG phase, and a combination of mass loss rates from [62], [287], and [309] for more massive models, while BPASS only uses rates from [62] that are higher on average. This serves to modulate the increased numbers of RSGs seen in the single BPASS models. This difference in mass loss rates carries over into the WR phase, where BPASS produces far more H-deficient WN stars, while the Geneva models form more H-rich WNH stars (not shown). However, the overall numbers of WR stars (and WC stars) are similar between the two codes.

The second and third rows illustrate the various effects of binary interactions and rotation. The effects of both rotation and binarity can be seen increasingly clearly from the least to most evolved phases. In particular, RLOF decreases the number of YSGs and RSGs. This causes a boost in the number of WR stars in the $f_{bin} = 1$ population at late times. Rotation prolongs the length of the early evolutionary phases, serving to delay the onset of the RSG and YSG phases. Rotating models also produce higher mass loss rates thanks to luminosity-dependent mass loss prescriptions and the higher luminosities of rotating stars, beginning at the terminal age main sequence and persisting through their post-main-sequence evolution [173, 76], as well as contributions from mechanical mass loss (e.g., mass loss via the stellar equator from matter rotating above the critical velocity). We also see an increase in the number of WR stars formed in the rotating Geneva populations; this boost primarily manifests as an increase of WNH stars, with slight *decreases* in the number of WN and WC stars. This is the result of a longer lifetime for the WNH phase and subsequent shorter lifetimes for the WN and WC phase (a consequence of rotational mixing), as well as efficient mass loss producing lower-mass WNH stars (for more discussion of rotation effects of WR subtypes see [91]). Finally, the bottom panel serves as a comparison between the binary BPASS population (dashed line) and the rotating Geneva population (solid line) as shown in the above rows. In summary: rotation causes a delay in the appearance of evolved supergiants and an increase in the number of WNH stars, while binary interactions effectively trade evolved supergiants for WR stars. However, the exact timing and degree of these effects as a function of initial mass/luminosity is more complicated, and this effect manifests in the observed stellar populations as we will demonstrate.

The predictions of the number count ratios in Figure 3.2 are fairly similar between the Geneva and BPASS populations at $\log t \lesssim 6.5 - 7$. Both binary interactions and rotation introduce similar effects, especially in B/R and WR/RSG. The Geneva models predict a higher overall value of WC/WN before $\log t \sim 6.75$ (due to the increase/decrease of WNH/WN stars respectively in the Geneva models), and a higher local maximum of WR/O at $\log t \sim 6.8$. At this time, the WR component of the Geneva populations becomes increasingly dominated by WC and WNH stars at increasing f_{rot}, while the BPASS WRs are largely WN type, which is reflected in the multiple orders of magnitude difference in the prediction for WC/WN. At increasingly later times, more stars in binaries become stripped. Models that lose enough Hydrogen increase the number of WRs well after $\log t \sim 6.8$, when the Geneva populations predict the last WRs have died, while models with only moderate mass transfer/loss become O stars/BSGs, depending on their luminosity — this is reflected in the rapid changes in the O/BSG ratio at late times. Overall, depending on the ratio chosen and the approximate age of the population being analyzed, different ratios are most sensitive to age, f_{bin}, f_{rot} or all three. For example, WC/WN is incredibly sensitive to f_{rot} in moderately evolved populations, while B/R is a relatively powerful age indicator in the earliest populations.

Below solar metallicity, the dominant differences are that all WR stars in the Geneva populations are H-rich (WNH in our labelling scheme), and binary interactions become increasingly important for producing WRs in the BPASS populations. The former is consistent with [90] and [104], and is a known feature of the Geneva models [155]. While this might indicate that all H-deficient WRs at low metallicity are formed by binary interactions, other possibilities and evolutionary pathways exist, which we defer to work focused more specifically on WR populations.

3.3.2 Comparisons with Real Data

Ensuring Self-Consistency

Two important effects must be considered before directly comparing some observed number count ratio to a theoretical prediction of this ratio, to ensure that the quantities being compared are identical:

- The value (and corresponding uncertainty) reported by an observer must be an estimate of the *underlying* number count ratio, $\hat{R}$ (an intrinsic characteristic of the stellar population belonging to the set of real numbers), rather than the raw observed ratio, R (which is a characteristic of the data belonging to the set of rational numbers due to the integer nature of the measurement).

- The theoretical population must approximate the observed population, and reflect the completeness of the catalog of massive stars (which may vary with spectral type).

We first consider how to estimate number count ratios and confidence intervals from observed data. Here we present a novel framework for making point-estimates (with corresponding uncertainties) for an *intrinsic* number count ratio. Much like photons in the low-count regime, the frequency of finding a given spectral type is determined by Poisson statistics (this is especially true for counting massive stars, where "shot noise" is the dominant source of uncertainty). In particular, the measurement uncertainty of the frequency of two subtypes, n_1 and n_2, can be approximated as $\sigma_{n_1} = \sqrt{n_1}$ and $\sigma_{n_2} = \sqrt{n_2}$ respectively. Past studies that calculate the number count ratio [e.g., 184, 208], apply traditional propagation of uncertainties, and report the *observed* ratio $R = n_1/n_2$, with corresponding uncertainty $\sigma_R = R\sqrt{n_1^{-1} + n_2^{-1}}$. As discussed by [208], this approach is problematic, and more sophisticated corrections can be made. However, an additional problem exists in that, with few enough stars (e.g., the Wolf-Rayet population of the SMC, where 1 WC star is known; [209]), or where the true underlying ratio is large, yet finite, $n_2 = 0$ is well-within a "3σ" error bar, and a measurement of $R = \infty$ could have been made.

This is a well-studied problem in the X-ray astronomy community, with a tractable solution within a Bayesian framework. [222] derive the posterior probability distribution of colors and hardness ratios for X-ray sources, in the limit of few (or no) photon counts. The problem here corresponds to a special case where the background is guaranteed to be 0, which simplifies the calculations somewhat; unknown sample contamination by foreground stars can be accounted for with minimal added complexity.

Say we measure a ratio $R = n_1/n_2$. Both n_1 and n_2 are assumed to be Poisson variables with (unknown) expectation values $\hat{n}_1$ and $\hat{n}_2$. What we wish to report is an estimate of the *true* ratio, $\hat{R} \equiv \hat{n}_1/\hat{n}_1$, which is a property of the underlying stellar population — indeed, it is the exact quantity plotted in Figure 3.2. From Bayes' theorem, the probability distribution for $\hat{n}_1$ given a measurement of n_1 is given by

$$p(\hat{n}_1|n_1) \propto p(\hat{n}_1)p(n_1|\hat{n}_1) \tag{3.3}$$

and similarly for $p(\hat{n}_2|n_2)$, where $p(\hat{n}_1)$ reflects our prior knowledge on the value of $\hat{n}_1$, and

$p(n_1|\hat{n}_1)$ is the likelihood of drawing n_1 from a Poisson distribution with expectation value $\hat{n}_1$. For a prior, we adopt $p(\hat{n}_1) \propto \hat{n}_1^{\phi-1}$. As discussed in [222] and [304], this is a special case of a γ-prior:

$$p(\hat{n}_1, \alpha, \beta) = \frac{1}{\Gamma(\alpha)} \beta^\alpha \hat{n}_1^{\alpha-1} e^{-\beta \hat{n}_1} \tag{3.4}$$

with $\alpha = \phi$ and $\beta \to 0$. This choice of prior ensures that the posterior probability function takes the same parametric form as the likelihood function. [222] found that, in Monte Carlo simulations, the choice of ϕ only has a moderate impact on the coverage (the percentage of simulations where the ground truth value of $\hat{R}$ is within a 95% confidence interval). We choose $\phi = 1/2$, which generally provides the best coverage for the observed number counts reported in typical extragalactic surveys.

Assuming $\hat{n}_1$ and $\hat{n}_2$ are independent (i.e., no stars of type 1 would also be counted as type 2[3]), the joint posterior distribution is $p(\hat{n}_1, \hat{n}_2|n_1, n_2) = p(\hat{n}_1|n_1)p(\hat{n}_2|n_2)$. Transforming $\hat{n}_1 = \hat{R}\hat{n}_2$ and marginalizing over $\hat{n}_2$,

$$p(\hat{R}|n_1, n_2) d\hat{R} = d\hat{R} \int_{\hat{n}_2} d\hat{n}_2 \hat{n}_2 p(\hat{R}\hat{n}_2, \hat{n}_2|n_1, n_2) \tag{3.5}$$

Utilizing Eq. (3.3), and substituting in the prior and likelihood functions,

$$p(\hat{R}|n_1, n_2) \propto \int_{\hat{n}_2} d\hat{n}_2 \hat{R}^{\phi-1} \hat{n}_2^{2\phi-1} \frac{\hat{R}^{n_1} \hat{n}_2^{(n_1+n_2)} e^{-\hat{n}_2(\hat{R}+1)}}{n_1! n_2!} \tag{3.6}$$

We use `emcee` [79], a Markov Chain Monte Carlo package, to sample the joint posterior probability distribution for $\hat{R}$ and $\hat{n}_2$ with 100 walkers initialized around the observed value R, 500 burn-in steps that are discarded, and an additional 3000 steps to explore the stationary distribution of walkers. In cases with $R \to \infty$ or $R \to 0$, we force R to be in the range $[10^{-10}, 10^5]$ when initializing the walkers. We estimate the value of both $\hat{R}$ and $\hat{n}_2$, as well as a 68% (1σ) confidence

---

[3] In one example below, n_1 is subset of n_2. There, a simple transformation can be made, but handling more complex situations is nontrivial.

interval, using the 16th-, 50th-, and 84th-percentile values of the samples.[4] The key advantages of this method are that the estimated quantity can be directly compared to the model predictions, that the reported errorbars correspond to the actual posterior probability distribution (and can be assymetric), and that the estimate of $\hat{R}$ is meaningful even if n_1 or n_2 are 0.

The challenge of accounting for complete samples is discussed in depth in Chapter 2. Here we reiterate that accounting for incompleteness in the observed samples — here defined as the lowest luminosity to which all stars of a given subtype have been found — is critical, and incorrectly handling or ignoring this effect can result in biases of ~ 0.1 dex in the estimated population age, and lead to incorrect results. In order to account for this effect, we include a luminosity cutoff $L_{\rm cut}$, that can be tuned for each subtype under consideration, and does not include models with $L < L_{\rm cut}$. Thus, even if the sample is assembled spectroscopically, or has a limiting magnitude in some photometric band that corresponds to different luminosity thresholds depending on the effective temperatures of the different subtypes, the model populations can be adjusted accordingly.

We note that assuming an observed sample is 100% complete above L_{cut}, and no stars are detected below it is a somewhat simplistic assumption. Below we compare our models to actual observed samples. In the two star clusters that we focus on, the data mostly come from focused studies that are designed to detect a given species. These stars are the brightest objects in a given part of the color magnitude diagram, have been followed up spectroscopically to remove contaminants, and the sample should be complete above the lowest luminosity star. In the Magellanic Clouds, the current existing samples of WRs and RSGs claim to be mostly-complete censuses of both species down to quite stringent magnitude limits [209, 208]. Some confusion arises in the counting of BSGs; however, the topic of bright blue stars in the Magellanic Clouds is currently being debated. We defer here to authors with more expertise [1]. Finally, observational bias, particularly in spectroscopic searches for WR stars, is likely to lead to missed weak-lined WR stars or WR stars with less evolved companions when not carefully accounted for, as discussed by [209].

---

[4] Software for performing these calculations, as well as reproducing all of the results in this work, is available online at `https://github.com/tzdwi/Diagnostics/`

Starburst Comparisons

We now wish to test our models in an environment where both sets of populations produce roughly identical predictions in a simple stellar population. From Figure 3.2, the best examples are young (< 10 Myr) star clusters, where we can assume that all of the stars belong to a single burst of star formation [see caveats in 99]. At these young ages, most WRs formed by binary interactions are evolved from progenitors that were massive enough to become WRs anyway. There are very few such clusters with enough confirmed members to adequately sample the IMF. With a mass of $M_* \approx 5 \times 10^4\ M_\odot$ cluster [6], and a well-studied cohort of evolved massive stars [44, 51], including a large number of BSGs, and an appreciable amount of WRs and RSGs, Westerlund 1 (Wd 1) is perhaps the best Galactic test bench for our model populations. In Chapter 2, we demonstrated that, when including binary effects and accounting for completeness, we can use two number count ratios, to estimate an age consistent with [51], who use a single diagnostic ratio and did not account for binarity or completeness. We can now apply our updated proscription for estimating number count ratios, as well as the rotating populations.

We first apply the Monte Carlo method described above to estimate the value of two ratios, O/BSG and WR/RSG. Using data from [44] and [51], we count $n_O = 22$, $n_{BSG} = 29$, $n_{WR} = 24$, and $n_{RSG} = 3$. As an illustration, the samples of the joint posterior distribution for $\hat{R}_{WR/RSG}$ and $\hat{n}_{RSG}$ are shown in the bottom-left panel of Figure 3.3, along with the marginalized posterior distributions for each parameter and accompanying point estimates (solid blue vertical line) and 68% confidence intervals (dashed black vertical lines). Note that because n_{RSG} is so low, the distribution of $\hat{R}_{WR/RSG}$ is very skewed, and the 97.5th-percentile upper limit is much higher than the reported 84th-percentile. All of the O stars in our sample are also blue supergiants, and so our assumption of independent variables no longer holds. Instead, we estimate $\hat{R}_{O/(BSG-O)}$ where BSG-O refers to all BSGs that are not O stars. We then transform each posterior sample of O/(BSG-O) into a sample of O/BSG estimates. With this method we measure intrinsic values of $\hat{R}_W R/RSG = 7.619^{+6.748}_{-3.194}$ and $\hat{R}_{O/BSG} = 0.755^{+0.072}_{-0.084}$.

We now wish to estimate age and f from the data. Figure 3.4 shows the predictions for

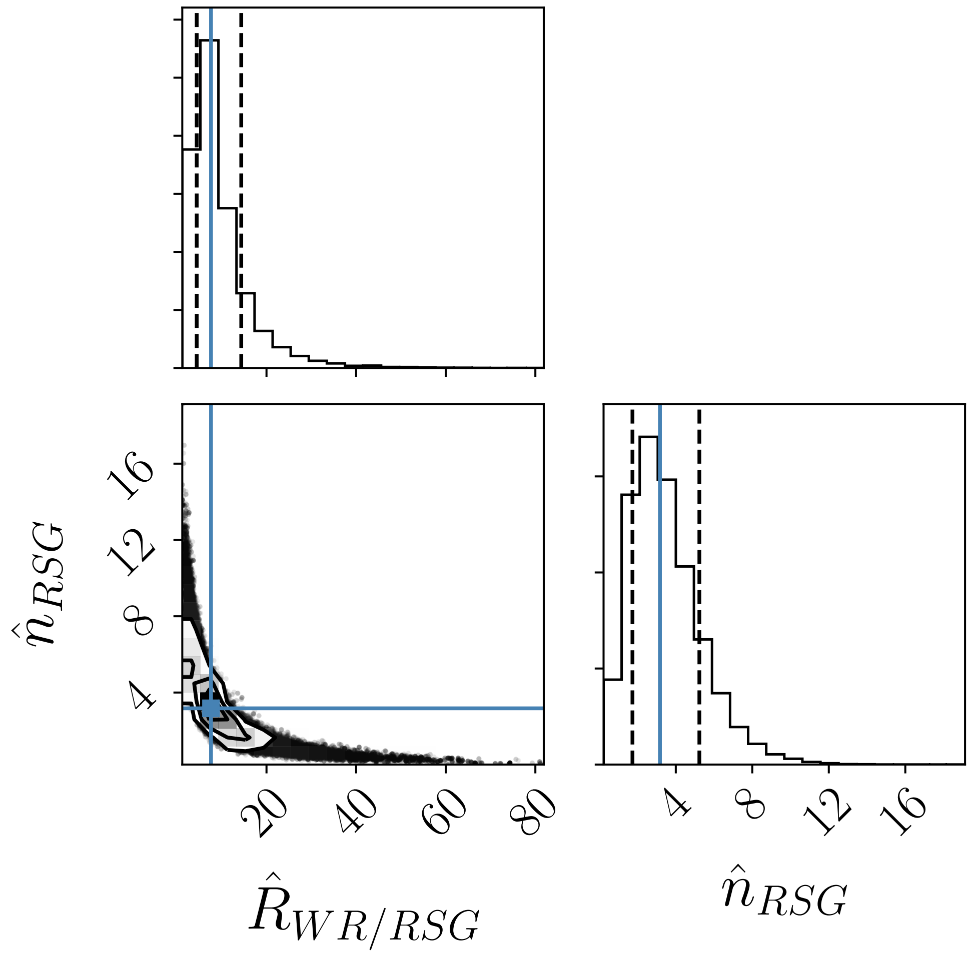

Figure 3.3 Posterior distribution samples of the estimate of WR/RSG for Wd 1. The 1-D histograms show the marginalized posterior distribution for the true ratio $\hat{R}$, and the true number of RSGs, $\hat{n}_{RSG}$, as well as the point estimates (in blue vertical lines) and 68% confidence intervals (in dashed vertical lines). Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [70].

WR/RSG vs. O/BSG, calculated on a grid of Geneva (left) and BPASS (center) populations with varying age (with $6.4 \leq \log t \leq 6.9$) and f. Lines of constant age and f are shown; the inset plot can be used to translate from the ratio-space into f and log age. The models incorporate completeness limits consistent with the lowest luminosities of each subtype reported by [44] and [51] (specificially, $L_{\text{cut}} = 4.9, 4.9, 5.1, 4.9$ for O stars, BSGs, WRs, and RSGs respectively). The right panel shows both grids overlain on top of each other. Combined, the two sets of models predict more-or-less identical values of both ratios as a function of age, especially for the $f = 0$ populations (shown in purple in all three panels). The data for Wd 1, plotted in blue, are consistent with an age of $\log t = 6.7$ (or approximately 5 Myr) using both model sets. This example is useful for validating our models, in that we demonstrate that the data are consistent with both model sets, as expected. Interestingly, the lines of constant age in the two grids are nearly orthogonal, implying that, with higher signal to noise data, these ratios can yield a measurement of both f_{bin} and f_{rot}.

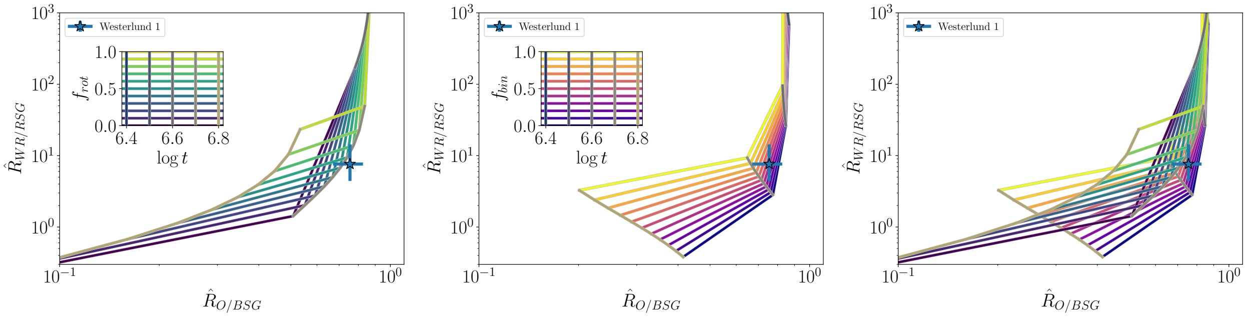

Figure 3.4 $\hat{R}_{WR/RSG}$ vs $\hat{R}_{O/BSG}$ for the Geneva (left) and BPASS (center) populations, calculated on a grid of $6.4 \leq \log t \leq 6.9$, and $0 \leq f \leq 1$, as shown by the inset panels. The rightmost panel shows both grids overlain on top of each other, with identical color coding. An estimate for $\hat{R}_{WR/RSG}$ and $\hat{R}_{O/BSG}$ in Westerlund 1, as well as corresponding 68% confidence intervals, is calculated using data from [44] and [51] and shown in blue. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [70].

In the example above focused on Wd 1, most of the stars in the sample were OB dwarfs or supergiants, where the differences between binary and rotating scenarios are small. For intermediate age clusters ($\log t \gtrsim 7$), more stars are expected to be found in increasingly evolved states, and the model grids no longer overlap. Figure 3.5 is similar to Figure 3.4, but calculated for SSPs with log-ages between 6.9 and 7.4 (approximately 8 and 25 Myr), with a luminosity threshold of $\log L = 4.9$ for all subtypes but O stars (comparable with the sample in [54], to which we compare below). The two model grids are separated by orders of magnitude in O/BSG, and showcase quite different behavior as a function of age in both ratios. This is due to the WR stars in the Geneva models dying, while increasing amounts of primary stars in the BPASS models are being stripped by binary interactions. These models boost the value of both ratios, depending on their luminosity and whether they lose enough Hydrogen to be classifed as WRs; if not, they tend to instead be classified as O stars. Note that we cannot visually compare the $f = 0$ case from both model sets, where $\hat{R}_{WR/RSG} \rightarrow 0$. Both grids reflect this, as they asymptote off the bottom left of the plot.

There exists only one stellar population that is massive enough to test our synthetic populations in this age regime: $h + \chi$ Persei. Photometric and spectroscopic studies of the members of $h + \chi$ Per have determined ages of 13-14 Myr ($\log(t) \approx 7.1$) for both clusters [264, 54]. While these studies have revealed a population of O stars, BSGs, and RSGs, no obvious WRs have been found, and thus we were unable to compare the data with the models in Chapter 2. Now, we can use the data from Table 3 in [54] to count 1 O star (HD 14434), 29 BSGs, and 7 RSGs, and estimate $\hat{R}_WR/RSG = 0.032^{+0.115}_{0.029}$ and $\hat{R}_O/BSG = 0.041^{+0.050}_{0.026}$, assuming 0 observed WRs. These values are plotted in blue in Figure 3.5. Note that the samples O stars and BSGs are *independent*, and we do not need to perform the same transformation as above. Because [54] study the main sequence down to G0 dwarfs, the luminosity cutoffs applied here are consistent with the data. The point estimates for both ratios are consistent with the BPASS models, though the Geneva models are not completely excluded. Interestingly, this difference is driven primarily by our measurement of O/BSG, and not WR/RSG (which we estimate without observing any WR stars). Using the BPASS model grid, the data corresponds to an age of ~ 10 Myr, consistent with previous age estimates.

Constant Star Formation Histories

While these simple stellar populations in our Galaxy are useful, the 68% confidence intervals on our number count ratio estimates are too wide to accurately determine age and f_{bin}/f_{rot}. To obtain a larger sample of evolved massive stars, we turn to galaxies in the Local Group. Not only do these galaxies contain more massive stars, they also sample a broad range of metallicities. However, these populations have complex star formation histories (SFHs), which we describe in Chapter 2.

Here, we only consider populations that are constantly forming stars, following [77]. Despite including stars from all age bins, all of the evolved types of massive stars considered here are

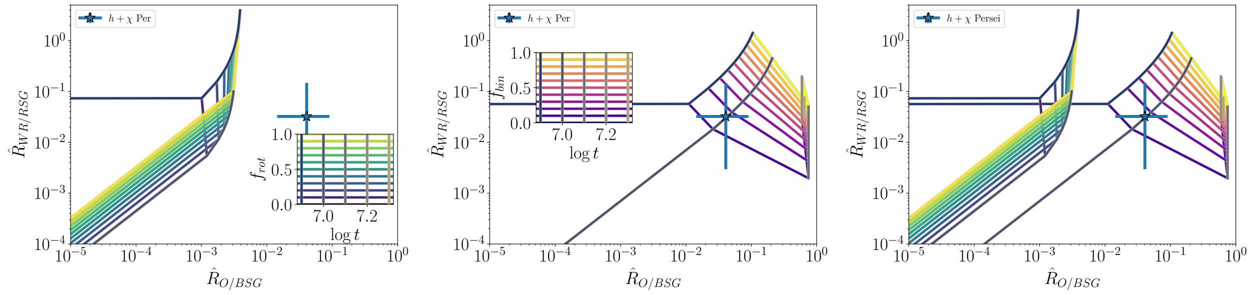

Figure 3.5 Similar to Figure 3.4: $\hat{R}_{WR/RSG}$ vs $\hat{R}_{O/BSG}$ for the Geneva (left) and BPASS (center) populations, now calculated on a grid of $6.9 \leq \log t \leq 7.4$. An estimate for $\hat{R}_{WR/RSG}$ and $\hat{R}_{O/BSG}$ in $h + \chi$ Persei, as well as corresponding 68% confidence intervals, is calculated using data from [54] and shown in blue. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [70].

only sensitive to, at most, 50-100 Myr of star formation, after which the populations of massive stars reach an equilibrium. Figure 3.6 shows WR/O, B/R, WR/RSG, and WC/WN as a function of metallicity and f_{rot} for galaxies constantly forming stars. A minimum luminosity of $\log(L) = 4.9$ is applied to be consistent with typical extragalactic studies (i.e., where only the supergiant population is complete). As discussed above, the Geneva models don't produce WC or WN stars below solar metallicity; the predictions at $Z_\odot$ are presented instead as individual points. Consistent with past models and existing observations [302, 117, 171], B/R is the least sensitive ratio to f_{rot}, and is a good indicator of metallicity, while WR/RSG is largely independent of metallicity [a fact that has long been in tension with observations, see section 6.4 of 157, and citations therein].

Recent measurements of the frequency of all three subtypes are only available for the Magellanic Clouds (MCs). Figure 3.7 shows $\hat{R}_{WR/RSG}$ vs. $\hat{R}_{B/R}$ using the Geneva (left), and BPASS models (center), with both grids overlaid (right), with lines of constant f_{rot}, f_{bin} and Z, as indicated with the inset grids. We note that the $f = 0$ case from both grids do not agree; BPASS predicts approximately an order of magnitude smaller values in both ratios. This is largely due to the significant difference in the number of RSGs predicted (see the top row of Figure 3.1), and thus, the difference in the mass loss prescription used by each code. However, both grids reside in largely the same part of this ratio space. Furthermore, consistent with Figure 3.6, B/R remains a useful

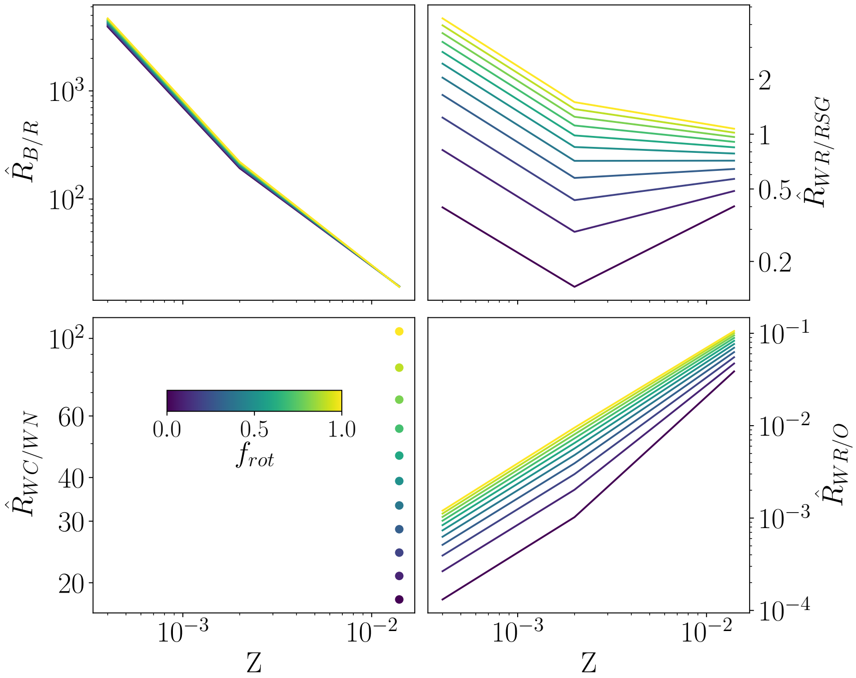

Figure 3.6 Predicted values of B/R (top left), WR/RSG (top right), WC/WN (bottom left), and WR/O (bottom right) for Geneva populations with constant star formation, and values of f_{rot} between 0 and 1, as indicated by the color bar. A minimum luminosity of $\log(L) = 4.9$ is assumed for all subtypes. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [70].

proxy for metallicity, and WR/RSG is a useful diagnostic of either f_{bin} or f_{rot}.

For both the Large Magellanic Cloud (LMC, plotted in orange) and Small Magellanic Cloud (SMC, plotted in blue), we estimate $\hat{R}_{B/R}$ using data from [184] and $\hat{R}_{WR/RSG}$ from [184], [206], and [208], and find $\hat{R}_{B/R} = 13.521^{+0.937}_{0.865}$, $\hat{R}_{WR/RSG} = 0.659^{+0.071}_{0.064}$ in the LMC, and $\hat{R}_{B/R} = 16.449^{+1.898}_{1.658}$, $\hat{R}_{WR/RSG} = 0.134^{+0.046}_{0.036}$ in the SMC. The Geneva models fail to reproduce the observations, while the BPASS models underpredict the metallicity of both clouds ($Z \approx 0.001/0.002$ for the SMC/LMC respectively). In both cases, the data indicate that the initial f_{bin} or f_{rot} is larger in the higher metallicity LMC than in the lower metallicity SMC, consistent with the results in [68]. We stress here that these results refer to the *natal* values of these parameters. Indeed, the exact opposite trend is predicted and observed in the frequency of X-ray binaries

[22]; however, this is a reflection the metallicity-dependent angular momentum evolution of binary systems. Measurements of the frequency of O star binaries (which in theory have not undergone mass transfer interactions in our model populations, and should serve as a decent proxy for the initial value of f_{bin}) reveal that f_{bin} does indeed increase from the LMC ($f_{bin} \sim 0.5$, [252]) to the Galaxy ($f_{bin} \sim 0.7$, [253]). Meanwhile, the observed rotation rates of metal-poor stars are expected to be higher (i.e., f_{rot} decreases with metallicity), due to stars being more compact at low Z [39]. However, the agreement between the overall trends in the data and those predicted by both model sets imply that number count ratios may be the best way to infer f_{bin} or f_{rot} in galaxies where direct measurements are currently infeasible.

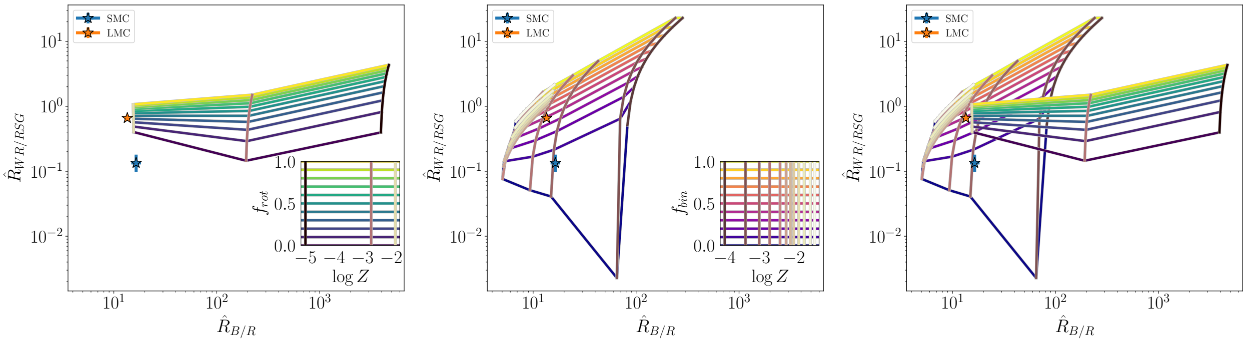

Figure 3.7 Similar to Figure 3.4: $\hat{R}_{WR/RSG}$ vs $\hat{R}_{B/R}$ for the Geneva (left) and BPASS (center) populations, assuming constant star formation, and calculated on a grid of $10^{-5} \leq Z \leq 0.04$. Estimates for $\hat{R}_{WR/RSG}$ and $\hat{R}_{B/R}$ in the LMC and SMC, as well as corresponding 68% confidence intervals, are shown in orange and blue, respectively. Data are from [184], [206], and [208]. Copyright AAS. Reproduced with permission from Dorn-Wallenstein & Levesque [70].

3.4 Discussion and Conclusion

Before discussing our results, we note two important caveats. Our results here assume that both stellar evolution codes accurately describe the evolution of the modelled stars. If this were true, then the $f = 0$ cases from each population should be identical, as both populations contain only single, nonrotating stars. Of course, both stellar evolution codes make different assumptions; while we refer the reader to the papers describing both codes for more details, we note that the different mass loss prescriptions in particular can result in the differing numbers of predicted cool supergiants.

We are furthermore assuming that both codes are completely accurate descriptions of rotation or binary interactions. Indeed, both codes are now seen as the industry standard for modelling their respective effects in massive stars. However, both codes do exhibit shortcomings. For example, the Geneva group only provides models that rotate at $0.4v_{crit}$. More rapidly rotating models are available, but only in a limited mass window below 15 $M_{\odot}$. With rapidly rotating models at higher mass incorporated into our model populations, we would be able to probe stars with significantly enhanced rotational mixing and enhanced mass loss — perhaps increasing the number of WNs while decreasing the number of WNHs.

Meanwhile, the BPASS team is the only group that explicitly models RLOF in binary systems and makes their results publicly available. However, as noted in [77], BPASS makes a number of simplifying assumptions in its treatment of circular orbits, rotation, and common envelope evolution. Furthermore, BPASS does not model systems with initial orbital periods shorter than one day. Regardless, BPASS is still successful at producing all classes of observed binary systems; in theory, some merger products are not modelled, but those products are either the result of very short-period systems — which merge very early, and evolve as single stars (J. Eldridge 2018, private communication), which may make result in an effectively top-heavier IMF — or are incredibly rare objects (e.g., Thorne-Żytkow Objects, [294, 295]).

These caveats aside, it is still important to examine how these two widely-used codes compare to each other and to observations when predicting the evolution and populations of massive stars, and to consider this comparison when interpreting the use and application of these models in future work. Our main results are as follows:

1. While rotation and binary interactions predict qualitatively similar effects on stellar populations, the predictions for the detailed makeup of simulated simple and complex stellar populations show dramatic differences between the two evolutionary scenarios. While some of these can be attributed to fundamental differences in the Geneva and BPASS codes (as

discussed in Section 3.2.1), it is also clear that rotation and binarity have quantitatively different effects on the evolution of massive star populations (in particular, the diagnostic number count ratios discussed in Section 3.3.1, and shown in Figure 3.2).

2. We introduced a novel Bayesian method of estimating both the value and error of diagnostic ratios in the low-number count regime typical of samples of massive stars. This method, combined with our implementation of completeness limits in the model populations, make our comparisons between models and data more accurate.

3. The data from observed Galactic populations agree with the grids of simulated SSPs, but suffer from poor signal-to-noise. The increased sample sizes in Local Group galaxies allows us to make higher-precision estimates of diagnostic ratios, and show that measurements of the natal binary fraction or rotation rate of stellar populations beyond the Magellanic Clouds may be possible with currently obtainable data. In the coming decades, JWST and the *Nancy Grace Roman Space Telescope* are scheduled to launch, giving us immediate access to red/optical-mid IR photometry, as well as sparsely sampled lightcurves for massive stars well-beyond the Local Group. Using comparable existing data in the Galaxy and Magellanic Clouds, it is possible to accurately fit for coarse spectral types akin to those used here (Dorn-Wallenstein et al. in prep).

4. Figures 3.4, 3.5, and 3.7 show a comparison between intrinsic number count ratios derived from grids of model stellar populations and inferred from observed data. For the starburst populations, we show that the ages consistent with the observed data are reasonably close to already published age estimates after correcting for completeness. However, the true power of this technique lies in its ability to derive the values and uncertainties of unknown parameters that are critically important for stellar evolution — e.g. f_{bin} or f_{rot} — in environments where traditional means of measuring these quantities are expensive or impossible given current technology. In this case, deriving the likelihood of obtaining a given set of star count ratios given a set of values for these parameters is nontrivial. Future work will focus on

applying Approximate Bayesian Computation (ABC, [286]) to derive constraints on f_{bin} and compare them to existing values found in the literature [253, 254].

5. Finally, we stress that rotating stars are found in binary (and higher order) systems. A complete model of stellar evolution should not designate one or the other effect as secondary. Rather, as we demonstrate here, binary interactions and rotation and produce both similar and contradictory effects in stellar populations, and a rigorous simultaneous treatment of both is necessary.

Chapter 4

THE PHOTOMETRIC CLASSIFICATION OF EVOLVED MASSIVE STARS: PREPARING FOR THE ERA OF *WEBB* AND *ROMAN* WITH MACHINE LEARNING

Evolved massive stars are observed in a menagerie of exotic evolutionary phases. While the challenge of connecting these states with a self-consistent theory of stellar evolution has seen rapid advancement since the original introduction of the "Conti Scenario" [49], the effects of rotation, magnetic fields, internal mixing processes, and binary interactions on the evolution of massive stars are still the subject of much theoretical effort [e.g., 76, 77]. While individual massive stars can be used as precision probes of these processes, *ensembles* of evolved massive stars can also significantly constrain stellar evolution. This can be done by comparing the integrated spectra of massive stars [e.g. 160], or by studying a detailed census of resolved populations of massive stars as done in Chapters 2 and 3, as well as [278].

Using the demographics of stellar populations to constrain stellar evolution requires large and accurately-classified samples of evolved massive stars. Such samples will be achievable in the coming years with the launch of the James Webb Space Telescope (*Webb*) and the Nancy Grace Roman Space Telescope (*Roman*). Among the instrumentation on *Webb* and the proposed instrumentation for *Roman* are photometers equipped with filters spanning a broad wavelength baseline from 0.5 to 28 μm. The resolution of *Webb* will allow us to identify and study in detail individual luminous stars throughout the Local Volume to distances of $\sim$4 Mpc [e.g. 127], while the impressive 0.218 deg^2 field of view of *Roman* will allow us to efficiently survey nearby galaxies in a small number of pointings [273]. Combined, observations from both missions will give astronomers access to precise infrared measurements of vast numbers of evolved massive stars. But without sophisticated methods of identifying and classifying these stars, the science return afforded by such a large increase in expected sample sizes will be significantly reduced.

Classification of massive stars from broadband photometry is often done by adopting simple linear cuts in color-magnitude space [e.g., 186, 188], and — most critically — do not include rare emission line objects, whose classification requires dedicated narrow-band surveys (sometimes with custom-designed filters, e.g. [209]), and ideally follow-up spectroscopy, both of which require extensive telescope time. Stars with line emission are often the post-main sequence evolved states of massive stars, in which the effects of rotation, binary interactions, and chemical mixing are the most pronounced; as such, the stars that place the most valuable constraints on unknown stellar physics are also the hardest to detect via traditional means. It is therefore worthwhile to determine whether there are alternative ways to classify massive stars that avoid using traditional and expensive methods.

Machine learning classification based on broad-band photometry has a rich history in the literature. With the advent of large surveys like the Sloan Digital Sky Survey (SDSS, [328]), optical data could be coupled with space-based MIR data to find the stellar locus in a 10-dimensional color-space [59]. Recent efforts to separate stars from quasars, or perform a regression on effective temperature with machine learning on photometric data have been successful [179, 15]; however, these studies are often focused on main sequence, low mass stars. This is an understandable choice given the rarity of evolved, high mass stars, the absence of reliable distances to calculate luminosities from which to select putative massive stars, and the fact that follow-up spectroscopy is necessary in order to confirm a star's membership in many important classes. However, with the advent of data from the *Gaia* mission [85], luminosities can be easily determined, and putative massive stars can be identified, allowing us to finally explore the application of machine learning to massive star classification.

At present, we can mimic the observing capabilities of *Webb* and *Roman* by combining data from *Gaia* (which has a red-optical bandpass), the Two Micron All Sky Survey (2MASS, [263], near-infrared), and the Wide Field Infrared Survey Explorer (WISE, [324], mid-infrared). WISE provides the additional benefit of having scanned the sky approximately every six months, yielding lightcurves spanning a $\sim$ 7-year baseline from which we can extract variability metrics for most stars observed. While *Roman* and *Webb* will not be observing the entire sky in this fashion, determining whether variability can aid in the classification of evolved massive stars will determine whether observers should seek repeated observations of a stellar population.

We wish to determine whether we can

1. Assemble a sample of evolved massive stars with available classifications as a training data set,
2. Construct a machine learning classifier that can reject low mass red contaminants and identify likely emission line objects in order to optimise available telescope time on the most promising targets,
3. Determine whether variability metrics estimated from WISE lightcurves can aid in these tasks, and
4. Determine which photometric bandpasses and variability metrics contribute the most to making accurate classifications.

Here we utilize a support vector machine classifier (SVC) trained only on broad-band photometry and simple metrics derived from WISE lightcurves to classify a large sample of evolved massive stars. We describe our sample selection and labeling method in §4.1. §4.2.1 details the calculation of the simple metrics derived from the WISE lightcurves, and describes the overall behavior of the stars in our sample. We explain our classification algorithm, discuss its successes and shortcomings in §4.3, and apply it to a training sample of 2500 stars before presenting our recommendations and concluding in §4.4. This work originally appeared in Dorn-Wallenstein et al. (2021) [69].

4.1 Sample Selection & Labeling

For any machine learning algorithm, a high-quality training set with accurate labels is necessary. The second data release (DR2) of the *Gaia* mission [85] contains precise photometry in three bands

(G, G_{BP}, and G_{RP}) and geometric parallaxes (ϖ) for 1.3 billion stars in the Milky Way (MW) and Magellanic Clouds. Because the parallax measurements suffer from some systematics [164], and many objects have high fractional errors (σ_{ϖ}/ϖ) or negative measured parallax, Bailer-Jones et al. (2018) [16] calculated Bayesian distance estimates for the majority of stars in *Gaia* DR2, using a prior based on the spatial distribution of stars in the MW. Figure 4.1 shows the difference between the distance inferred by Bailer-Jones et al., r_{est}, and a naive distance derived by inverting the reported *Gaia* measurements of ϖ for $\sim 10,000$ putative massive stars (as described below). The dashed line indicates where $r_{est} = 1/\varpi$. While the two distance estimates are roughly consistent for nearby stars, more distant stars are much closer than inferred from biased naive distance estimates.

We first perform a cross-match between the Bailer-Jones et al. catalog [16] and the existing cross-match between *Gaia* DR2 and the ALLWISE data release. ALLWISE [56] contains photometry in four mid-infrared (MIR) bands — $W1$ (3.4 μm), $W2$ (4.6 μm), $W3$ (12 μm), and $W4$ (22 μm) — derived from co-added images obtained during the original WISE mission, as well as $W1$ and $W2$ images obtained in the post-cryogenic NEOWISE mission [177]. We select all stars with successful distance estimates (i.e., where `result_flag` $= 1$ if the distance estimate is the mode of the posterior distribution, 2 if it is the median, and 0 for a failed estimate, see Bailer-Jones et al. [16] for more details) that satisfy

$$\begin{aligned} &M_G = G - 5\log r_{est} + 5 \leq -1.5, \\ &W1 < 14. \end{aligned} \tag{4.1}$$

Since Bailer-Jones et al. used a Galactic prior, stars in the Large and Small Magellanic Clouds (LMC/SMC respectively) have distances that are considerably underestimated. Thus, we also match the catalog in Gaia Collaboration et al. (2018b) [87] to the ALLWISE/*Gaia* cross-match, and select stars with $W1 < 14$ and $M_G \leq -1.5$, assuming distance moduli of 19.05/18.52 for the SMC/LMC respectively [143, 142] and combining the two cross-matches while dropping duplicate stars. This results in a total of 452,283 stars. The left panel of Figure 4.2 shows the logarithmic density on the sky of all stars selected from the *Gaia* DR2 database; the Galactic plane and Magellanic Clouds are clearly visible.

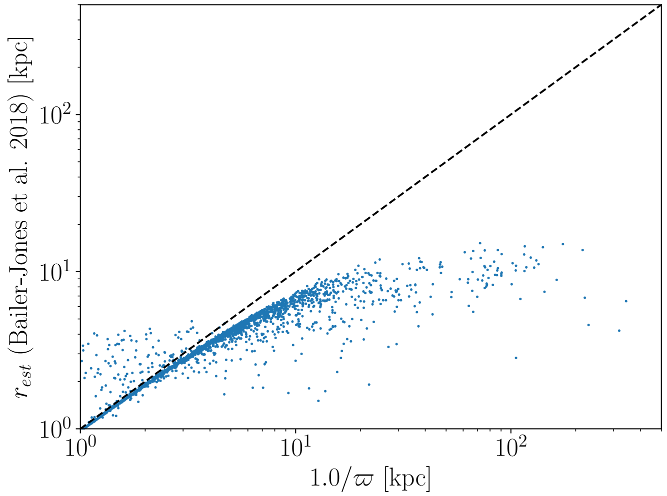

Figure 4.1 Distance from Bailer-Jones et al. [16] versus distance inferred via inverting the reported ϖ from *Gaia* DR2 for $\sim$ 10000 putative massive stars. The dashed line shows where $r_{est} = 1/\varpi$. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

We then estimate the reddening in the *Gaia* bandpasses using the published estimate for A_G from *Gaia* DR2, and coefficients from [180] to calculate $E(G_{BP} - G_{RP})$. For Galactic stars without A_G estimates, we assume $A_G = 0$, and for stars in the Magellanic Clouds, we assume the average value of A_G and $E(G_{BP} - G_{RP})$ using R_V measurements from [97] and $E(B - V)$ from [185]. Using these quanitites, we calculate the intrinsic $G_{BP} - G_{RP}$ and M_G for all stars.

We next construct color-magnitude diagrams (CMDs) in the *Gaia* filters, which we can use to select massive stars — i.e., stars with initial mass $M_i \geq 8M_\odot$ (right panels of Figure 4.2). To find the limiting magnitude expected for massive stars at a given *Gaia* color, we use the MESA Isochrones & Stellar Tracks (MIST, [73, 42, 223, 224, 225]) isochrones with metallicity $[Fe/H] = 0, -0.5, -1$ for the Galaxy, LMC, and SMC respectively, and rotation speed relative to critical of

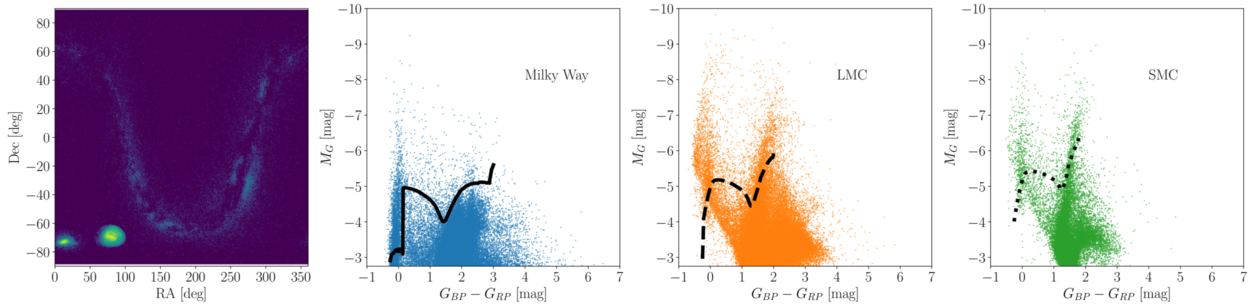

Figure 4.2 Left panel shows the density of stars selected from the *Gaia* database on the sky. Intensity of the colormap corresponds to the logarithm of the number of stars in each bin. The right three panels show the *Gaia* CMD for stars selected from the *Gaia* DR2 database brighter than $M_G = -1.5$ (though we only plot stars brighter than $M_G = -2.75$ to highlight the likely massive stars). Galactic stars are in blue (center left panel), LMC stars are in orange (center right panel), and SMC stars are in green (right panel). The solid, dashed, and dotted lines represent our minimum-luminosity criteria to select massive stars in the Galaxy, LMC, and SMC respectively. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

$v/v_{crit} = 0.4$. We then selected the faintest isochrone point of any age with $M_i \geq 8$ M$_\odot$ in 100 equally-spaced bins of color in the range $-0.25 \leq G_{BP} - G_{RP} \leq 3$. We note that the oldest MIST time bin is $10^{10.3}$ yr (older than the age of the Universe), but by selecting points with $M_i \geq 8$ M$_\odot$, none of the selected points are older than ~40 Myr. These isochrone points form a boundary in the *Gaia* CMD that represents the faintest luminosities reached by any massive star at any point during its evolution, and no fainter massive stars are expected to be found. However, many stars with $M_i < 8$ M$_\odot$ also lie above this boundary — mostly RGB and AGB stars — so while our sample maximized the number of massive stars selected, it is not constructed to be free of contamination. Note that the thresholds accurately capture the slope of the main sequence for all three galaxies, as well as the $G_{BP} - G_{RP}$ color corresponding to the Hayashi limit.

We select all stars brighter than the corresponding luminosity threshold for their host galaxy, resulting in 9784 stars. From this sample, we select all stars fainter than the saturation limit in $W1$ (8) and $W2$ (7) with valid measurements listed in the ALLWISE catalog for the three bluest WISE bands (excluding $W4$, where the signal-to-noise is often poor). We also convert the $W1$

Table 4.1. Common names, coordinates, host galaxies, and *Gaia* measurements of 6,484 putative massive stars, ordered by Right Ascension. r_{est} from [16] is given for Galactic stars. Listed values of G and $G_{BP} - G_{RP}$ are uncorrected for extinction.

Common Name	R.A. [deg]	Dec [deg]	Host Galaxy	r_{est} [kpc]	G [mag]	A_G [mag]	$G_{BP} - G_{RP}$ [mag]
HD 236270	0.17442287	55.72245665	MW	2.162	9.07	0.94	0.24
LS I +64 10	0.38838658	64.51219232	MW	5.305	11.55	1.33	0.57
LS I +60 69	0.55506135	60.43828347	MW	5.866	11.85	1.23	0.61
BD+62 2353	0.59458742	62.90087875	MW	5.243	9.81	0.48	0.37
HD 73	1.40408512	43.40139506	MW	1.869	8.19	0.12	−0.15
HD 240496	1.42175475	58.49541068	MW	2.499	9.70	1.55	0.68
WISE J000559.28-790653.3	1.49713706	−79.11483482	SMC	-	13.94	0.21	1.05
LS I +59 30	1.70503555	59.85955733	MW	4.006	10.86	1.19	0.50
BD+57 2870	1.82960982	58.33785301	MW	3.893	9.84	1.42	0.82
BD+62 1	1.88805102	63.08030731	MW	2.893	10.29	1.30	0.53

Note. — This table is published in its entirety in the machine-readable format. A portion is shown here for guidance regarding its form and content.

and $W2$ magnitudes (and uncertainties) to fluxes, and filter for stars with signal-to-noise ratio greater than 3. This results in a final sample of 6484 objects. We discuss the makeup of this sample in §4.1.1. Table 4.1 lists the names, coordinates, host galaxies, distances from Bailer-Jones et al. [16], and *Gaia* photometry for these stars. We query `Vizier` [218] using `astroquery` to download JHK_s photometry from the 2-micron All Sky Survey (2MASS, [263]) for all stars. We also query SIMBAD [322] and download the common name (`MAIN_ID`), spectral type (contained in the `MK_Spectral_Type` and `SP_Type` fields), and object type (`OType`) for each star, the latter two of which we use to assign labels.

4.1.1 Label Assignment

For our final sample of ∼6500 stars, we wish to assign the best available estimate of its evolutionary state. These labels can be used to compare to the predictions of stellar population models. Note that these evolutionary states (which are theoretical concepts) are mostly tied to spectral appearance (which is an observable quantity). Therefore we are assuming that, e.g., all stars with Wolf-Rayet spectra are in the same evolutionary state (namely, the descendants of massive stars with high luminosities that have lost their envelopes via strong winds), and that all stars in that evolutionary state are observed as Wolf-Rayet stars. We know that at least the former isn't true, as some stars lose their envelopes due to interactions with a binary companion [77], and the latter is also questionable since such stars may or may not appear similar to classical Wolf-Rayets [100]. Nevertheless, we assume that the assigned labels are a reasonable approximation for a star's evolutionary state with this caveat in mind.

At present, a database of homogeneously classified massive stars does not exist. While all-sky spectroscopic surveys have observed many massive stars, the machine learning pipelines that produce the effective temperatures, surface gravities, and chemical compositions that would allow us to accurately classify our sample do not cover the parameter regime in which massive stars reside [e.g. 88]. As a result, we must use the heterogeneous classification data available on SIMBAD. For each star, we apply a decision tree that results in the star receiving a single label. Figure 4.3 shows a flowchart that summarizes our labeling scheme. Note that this process highly tailored to this dataset, and some branches in the decision tree serve only to accurately label very small numbers of stars with unique spectral types (e.g., spectroscopically peculiar stars or X-ray binaries). Deriving labels for known massive stars using existing sources is not trivial, and our labelling scheme would be entirely different if a large sample of massive stars with well-measured temperatures, surface gravities, and chemical abundances were available.

We first use the common name and *Gaia* `source_id` of the star to determine if the star belongs to the catalog of confirmed Luminous Blue Variables (LBVs) presented in [241][1]. Non-LBVs are classified as WR stars if "W" is in the spectral type, or the SIMBAD `OType` is "*WR". Non-WRs with "K" or "M" in their spectral type are classified as either Red Supergiants (RSGs) or "C/S/Giant" if their SIMBAD `SP_Type` contains "III" — we keep all such low-mass contaminants in our sample as distinguishing between RSGs and luminous low-mass giants is still a difficult problem [188, 326, 213]. The resulting sample of RSGs is pure; of the five RSGs that don't have

---

[1] Note that we chose to ignore LBVs in Chapter 2 and 3 due to the inhomogeneous nature of the class as well as the difficulty of identifying LBVs observationally. However, the method that we develop is capable of classifying LBVs (see §4.3), and so we choose to include them in the classifier.

luminosity class I, all of them are luminosity class Ia-II, Ib-II, or Iab-II, which are consistent with bona fide RSGs [162]. This is a good test of the *Gaia* DR2 parallaxes and [16] distances, as cool subgiants and dwarfs with luminosity class IV or V would have been erroneously classified as RSGs using our criteria had they been included in our sample due to inaccurate distance and M_G measurements.

Non-RSGs with “F” or “G” in their spectral type are classified as Yellow Supergiants (YSGs). However, this includes a number of blue stars. Further inspection of these stars reveals a number of objects whose `MK_Spectral_Type` field contradictorily indicates these are hot stars, with spectral type O, B, or A. As these stars have *Gaia* photometry consistent with hot stars, we classify them as such (see below). Eight low-mass yellow stars are also included in our sample. Stars with “III” or “V” in their spectral types are classified as Yellow Dwarfs. While luminosity class III formally denotes giant stars, only one yellow giant is in the sample, and so we assign it the Yellow Dwarf label. We note that for extragalactic samples, foreground dwarfs can usually be filtered based on proper motions, while dwarfs belonging to the stellar population under study can be excluded just based on their apparent magnitudes. Nonetheless, we retain this class to avoid confusion between these stars and true YSGs. All YSGs that aren’t hot stars or dwarfs keep their YSG label.

Of the objects that have not yet been classified, stars with “[e]” in their spectral type are classified as OB[e] stars, while non-OB[e] stars with spectral types including an “e” (without brackets, and without “pec” in their spectral type) are classified as OBAe stars. If the star is not yet classified and O, B, or A are in the spectral type with no additional information, they are classified as generic OBA stars. OBA stars with “III” or “IV” in their spectral type are classified as evolved OBA stars, stars with “V” in their spectral type are classified as OBA main sequence stars, and finally, stars with “I” in their spectral type are labeled as OBA supergiants. All stars that have not been assigned a label at this stage are either Carbon stars or S stars (which are assigned the C/S/Giant class), stars labeled only as Variables in SIMBAD (e.g., LPVs, semi-regular variables, or just variables, without other spectral information) which are assigned the Miscellaneous Variable classification, or stars with no identifying information/no confirmed designation (e.g., the SIMBAD `OType` is “Star” or contains "Candidate") which are classified as Unknown/Candidate.

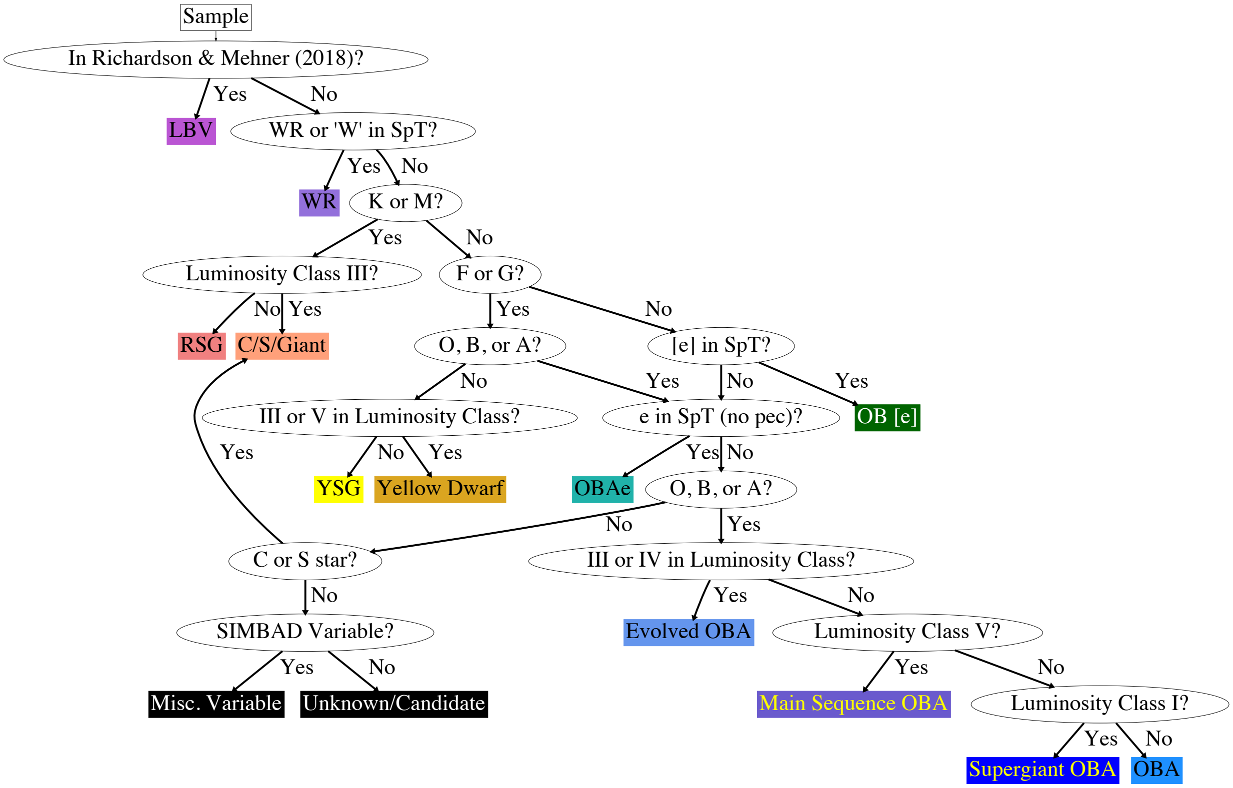

Figure 4.3 Flowchart illustrating the process by which stars are assigned labels, as described in text. Each star begins in the top left and is assigned a label by following a series of binary decisions. This process is complex, and demonstrates the difficulty in deriving useful labels for massive stars. For example, some stars with F or G in their spectral types are actually hot OBA stars (as described in text), and require special handling. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

Finally, we include an "Is Binary" flag for all stars, which is 1 for stars classified as Eclipsing or Spectroscopic Binaries, High Mass X-ray Binaries, or Ellipsoidal Variables, or if they have a compound spectral type (e.g., WN8 + O6V),[2] and 0 otherwise; 102 stars are flagged as binaries. This flag is separate from the labeling process shown in in Figure 4.3. Because photometry of binary systems can be misleading [205], and binary systems exhibit a broad range of variability

---

[2] This does not include stars with the "OB+" spectral type, which is an outdated class that describes OB stars with weaker absorption lines that would now be classified as OB supergiants.

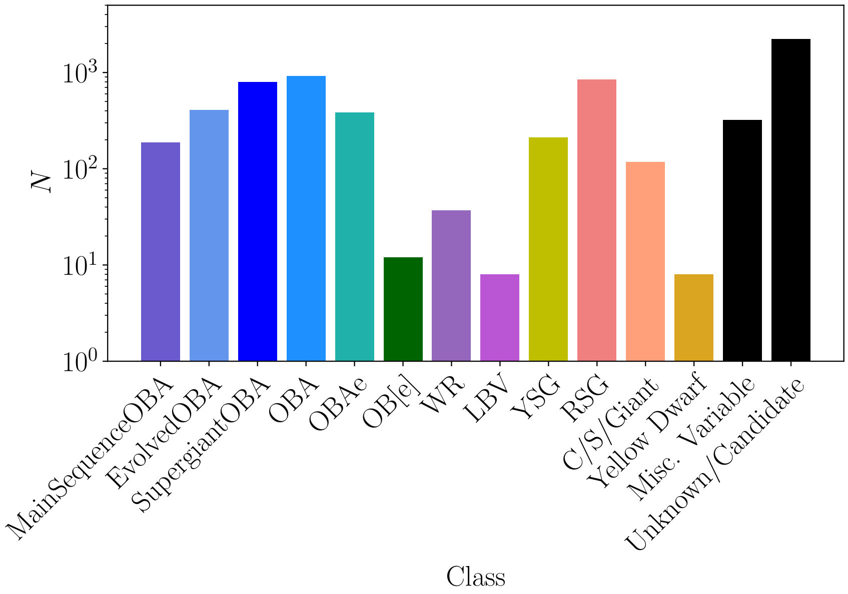

Figure 4.4 The makeup of our sample of massive stars. Note that the sample is dominated by OBA stars and cool supergiants. Non-OBAe emission line stars — OB[e] stars, WRs, and LBVs — are the rarest massive stars in our sample, despite being stars of great scientific interest. For readability, we have used a logarithmic y-axis to display our sample statistics. Note that in practice, differences in the number of stars per class are much larger than they might appear here. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

that is not intrinsic to the individual components, we exclude these stars from our classifier.

Figure 4.4 shows the makeup of our sample. Approximately 30% of our sample (2550 stars) belong to the Miscellaneous Variable and Unknown/Candidate classes, which we do not use to train our classifier; instead, we use the classifier to assign tentative classifications to these stars in §4.3. The rest of the sample is dominated by luminous OBA stars and cool supergiants, with very few LBVs, OB[e] stars, and WR. This is unsurprising given the incredibly short expected lifetimes of the evolutionary phases in which massive stars exhibit line emission relative to the time they spend as main sequence OBA stars or helium-burning RSGs [76]. However, despite the fact that massive stars spend 90% of their lifetimes on the main sequence, only 2309 stars ($\sim$60% of labeled stars in our sample) have OBA spectral types, implying that cool supergiants are *overrepresented*

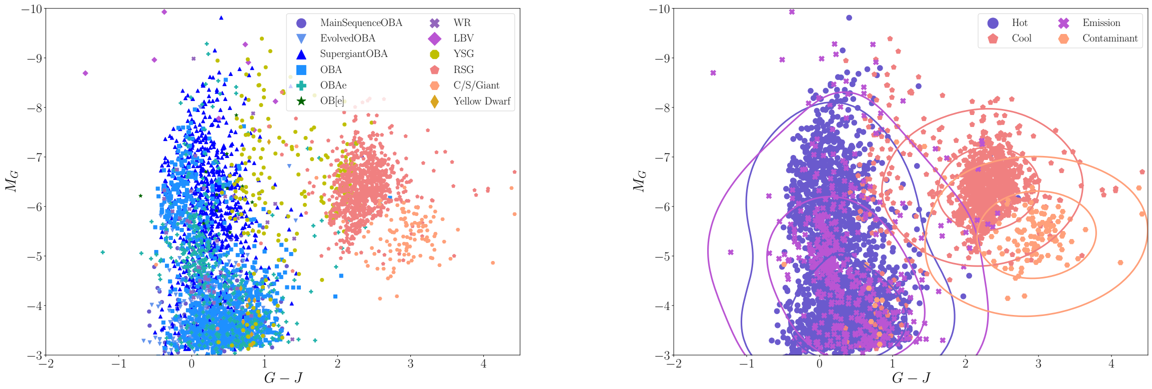

Figure 4.5 M_G vs. $G-J$ for putative massive stars. *Left*: Stars are colored by their label. While we also use shapes to distinguish between stars in different classes, this illustrates a key difficulty faced by our classifier: the classes have significant overlap with each other in the CMD. For example, the coolest/warmest YSGs have identical optical photometry to RSGs/OBA supergiants respectively, while the different classes of hot stars are impossible to distinguish from one another by eye. *Right*: Stars are colored by their coarse label. Contours for each coarse class correspond to 0.5 and 0.1 times the maximum value of a kernel density estimate of the distribution of each class in the CMD. Even in the coarse labels, the contours for hot stars and emission line stars are nearly identical. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

in our sample. This is unsurprising given that MIR photometry is ideal for observing these stars. Regardless, the availability of training data is imbalanced across different classes, especially in the rare classes, which will impact the performance of the classifier if not properly addressed. [37]. We discuss this issue for this particular sample in §4.3.1.

The left panel of Figure 4.5 shows the M_G vs. $G-J$ CMD for all stars in our sample that aren't labeled as miscellaneous variables or unknown/candidate, colored by their label. $G-J$ correlates reasonably well with effective temperature in main sequence stars [58], and in this case is especially useful for distinguishing from the near vertical main sequence/blue supergiants and the significantly cooler yellow and red supergiants. From this plot, it is clear that many stars are misclassified in SIMBAD (the worst example is one particular red star classified as an OBA star), reducing the effectiveness of any machine learning algorithm, and propagating biases into the results. When we examined our sample in detail, we found that yellow supergiants are especially prone to this problem: 81/212 YSGs in the sample (38%) have $G - J < 1$, consistent with the optical colors of much hotter stars. Indeed, this problem would have been worse had we not corrected for the presence of OBA stars in the initial sample of YSGs. This issue may originate from bad distance estimates for individual stars (which explains why our sample includes F and G dwarfs), bad estimates of reddening (given our usage of monochromatic extinction coefficients), previously unidentified variability, or the fact that many of these spectral types were determined via stellar spectra taken on photographic plates [for example, many spectral types for stars in the LMC come from 1972 9]. We expect this issue to propagate into our results, increasing the confusion between YSGs and hot stars.[3]

Because we expect objects in some classes — especially those with an evolutionary link such as main sequence, evolved, and supergiant OBA stars — to appear similar in the training data set, we also assign all stars a coarse label: all classes of OBA stars excluding OB[e] and OBAe are labeled “Hot”; RSGs and YSGs are labeled “Cool”; WRs, LBVs, and both OB[e] and OBAe stars are labeled “Emission” (EM for short); C/S/Giant stars and Yellow Dwarfs are labeled “Contaminant”; and miscellaneous variables and unknown/candidates are labeled “Unknown/Candidate.” The results of this labeling scheme are summarized in Table 4.2 which shows the number of stars with a given refined class that are assigned a particular coarse label. The right panel of Figure 4.5 shows the same CMD, with points colored by their coarse label. This leads to some improvement: each coarse class lies in the approximate region of the CMD that one would expect. Regardless, it is evident that selecting any one of these classes solely from this optical photometry would be diffi-

---

[3] It would certainly be possible to tailor our dataset by removing the “worst” stars, thus cleaning up the boundaries between classes. However, by doing this we would be making several assumptions about where different classes of stars reside in our feature space, with no way of knowing whether these enforced boundaries actually divide stars in physically different evolutionary states. Indeed, we *expect* the boundaries between classes to be fuzzy, because our discrete labels are an approximation of a continuum of evolutionary states — a fact that we have otherwise swept under the rug. That said, a significant amount of the overlap between classes *is* due to the poor quality of existing labels. Instead of trying to guess which stars are poorly labelled, and which ones truly reside in the overlap between classes, we instead wish to see how the quality of existing labels impacts the performance of our classifier. With better labels (or even effective temperatures and surface gravities) based on spectroscopy from large ground based surveys, this problem would significantly improve. However, as we discuss in §4.4, while this data exists, attempts to develop spectroscopic labeling for massive stars have been nonexistant until extremely recently.

cult. The “cool” class has significant overlap with the “hot” class — largely driven by the YSGs — emission line stars can be found at a range of $G_{BP} - G_{RP}$ colors, and there is significant overlap between low-mass contaminants that will end their lives as white dwarfs and true massive stars that will end their lives in supernovae explosions. This point is emphasized by the contours, which correspond to 0.5 and 0.1 times the maximum value of a kernel density estimate of the distribution of each class, which replaces each point with a kernel function (in this case a two-dimensional Gaussian centered on the point), and sums the kernels to estimate the underlying distribution. We do this using the `KernelDensity` estimator from `sklearn`, and use a similar cross-validation scheme described below to find a suitable bandwidth for the kernel (i.e., the width parameter of the Gaussian).

We use these coarse labels to train a second classifier. While these coarse labels lose some specificity, each coarse class contains more stars, hopefully increasing the performance of a classifier trained on these labels. Furthermore, they still retain physical information while increasing the number of stars in each class: the “cool” label contains stars with convective envelopes, while “hot” stars contain radiative envelopes. Meanwhile, emission line stars are notable for their variability. It is our hope that this second classifier will still address two of our stated goals: to identify emission line stars, and to reject contaminating low mass stars.

4.2 WISE Lightcurves

Variability in evolved massive stars has been well-characterized at timescales from minutes to decades [e.g. 46, 71, 270]. In a study of massive stars in the Whirlpool Galaxy (M51), [46] found that almost half of the stars brighter than $M_I = -7$ were variable, with red stars nearing a variability fraction of 1. Both red and extremely luminous blue stars exhibited quite high amplitude ($\Delta I \geq 0.3$) variability. For spectral energy distributions (SEDs) dominated by purely stellar light, mid-infrared (MIR) flux measurements (and thus variability) is sensitive to (variations in) the bolometric luminosity. However, for stars with significant circumstellar dust components in their SEDs, MIR variability is correlated with both intrinsic bolometric variability, and with dust creation/destruction processes in the circumstellar medium [for example, in RSGs where it is

Table 4.2. Number of stars in a class that are assigned a given coarse label, not including the Miscellanous Variable or Unknown/Candidate labels.

	Coarse Label			
Refined Label	Hot	Emission	Cool	Contaminant
Main Sequence OBA	187			
Evolved OBA	409			
Supergiant OBA	798			
OBA	915			
OBAe		383		
OB[e]		12		
WR		37		
LBV		8		
YSG			212	
RSG			847	
C/S/Giant				118
Yellow Dwarf				8
Total	2309	440	1059	126

correlated with the mass loss rate, e.g. 327].

The WISE mission provides lightcurves from stars in all parts of the sky, observed over a $\sim$7 year baseline. Due to the scanning law adopted by WISE [324], most stars not on the ecliptic poles are visited approximately every $\sim$ 180 days. All stars have a $\sim$ 3-year data gap from when WISE was placed in hibernation in February 2011 and when it was reactivated in December 2013. WISE initially observed simultaneously in four filters during its primary mission: $W1$ (3.4 μm), $W2$ (4.6 μm), $W3$ (12 μm), and $W4$ (22 μm). However, it was reduced to using only the two bluest bands in its post-cryogenic survey mode called "NEOWISE". The time, duration, and number of individual observations during each $\sim$ 180 day visit depends on spatial geometry of the WISE scanning program, i.e. stars closer to the ecliptic poles have longer duration visits (often exceeding a week) with many epochs per visit, while star near the equator have very short visits (typically a couple days) with only a few epochs per visit. Because WISE lightcurves possess such non-uniform cadence, extracting detailed physics for most individual stars is difficult. However, the WISE lightcurves place *fantastic* constraints on MIR variability amplitudes on longer timescales, especially for evolved massive stars whose highest-amplitude variability occurs over $\sim$year timescales. Such amplitude and timescale estimates are related to the physical parameters of the star, potentially aiding in classification.

For every star selected in §4.1 we queried the Single-Exposure ("L1b") source databases for all phases of the WISE mission, including the original 4-band, partial cryogenic 3-band, and post-cryogenic 2-band NEOWISE tables. We used `astroquery` to pull data in the region within 3 arcsec of the known source location. To ensure high quality data for all recovered epochs, we require the photometric quality flag to be `PH_QUAL=A`, the contamination flag to be `CC_FLAGS=00`, the number of deblended sources flag to be `NB=1`, and the PSF photometry fit quality (defined as the reduced χ^2) in $W1$ to be `w1rchi2<5`.

Only two of our stars did not have usable data from WISE: WISE J074911.48-102000.2 (HD 63554) has no lightcurve available online, and WISE J050128.62-701120.2, which does not have any corresponding object nearby on SIMBAD. When calculating each of the variability metrics below, we instead record a value of `NaN` (i.e., missing data). For the remaining stars, we ignore the

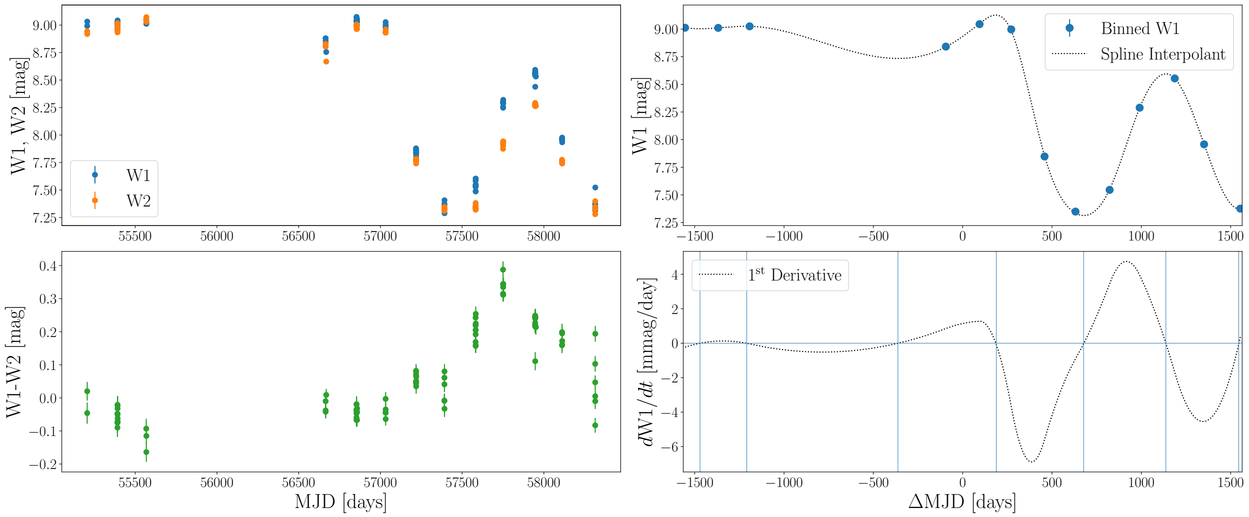

Figure 4.6 Example lightcurve for WISE J000536.97+432405.0. *Top left*: Raw lightcurve, with $W1$ points plotted as blue errorbars, and $W2$ points plotted in orange. *Bottom left*: Variability in $W1 - W2$ plotted as green errorbars. *Top right*: Binned $W1$ lightcurve. Blue points are binned data (errorbars are smaller than the points). Black dotted line is the B-spline interpolation. Time has been adjusted so the lightcurve is centered on $t = 0$. *Bottom right*: First derivative of the interpolant. Vertical blue lines show the times where the derivative crosses zero, indicated by the horizontal blue line. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

$W3$ and $W4$ data here due to the lower signal to noise and significantly shorter observing baselines due to the loss of cryogenic observations after the original WISE mission. As WISE observes simultaneously in all bands, we can construct $W1 - W2$ lightcurves without any interpolation, and simply subtract the $W2$ data from the $W1$ data to obtain the $W1 - W2$ color curve. The left panels of Figure 4.6 show an example set of lightcurves for WISE J000536.97+432405.0 (=HD 73), a B1.5IV star that illustrates the typical observing cadence and variability of a bright star in our sample.

4.2.1 Variability Metrics

Amplitude

For each of the three lightcurves of each object, we wish to extract simple metrics that describe the amplitude and timescale of variability. We choose χ^2 about the median defined as

$$\chi^2 = \sum \left(\frac{M_i - \tilde{M}}{\sigma_i}\right)^2 \tag{4.2}$$

and the reduced-χ^2

$$\chi^2_{red} = \chi^2/(N-1) \tag{4.3}$$

Where M_i is a magnitude measurement, $\tilde{M}$ is the median of the lightcurve, σ_i is the corresponding error on the data point, and N is the number of points in the lightcurve. We also calculate the Median Absolute Deviation (MAD), and Error-Weighted MAD (EWM):

$$MAD = \mathrm{Median}(|M_i - \tilde{M}|) \tag{4.4}$$

$$EWM = \mathrm{Median}(|M_i - \tilde{M}|/\sigma_i) \tag{4.5}$$

If the filtered and cleaned lightcurve only contains one good measurement (or no good measurements), we automatically give it $\chi^2 = \chi^2_{red} = \mathrm{MAD} = \mathrm{EWM} =$ `NaN`. We describe our method for treating missing data below.

The top panels of Figure 4.7 show the distributions of $\log \chi^2_{red}$ and $\log$EWM derived for our sample. Values from $W1$ lightcurves are in blue, $W2$ in orange, and $W1-W2$ in green. While the distribution of amplitudes are similar in both $W1$ and $W2$, stars tend to have significantly smaller amplitude color variability. The bottom left panel shows a scatter plot of χ^2_{red} vs. EWM for all three lightcurves. While the two measures correlate reasonably well with each other, there is a branch of stars whose lightcurves have high χ^2_{red} and low EWM in $W1$, $W2$, and $W1-W2$; because the EWM is robust to outliers, χ^2_{red} is an effective probe of lightcurves with sudden brightening/fading events, while EWM is an effective selector for lightcurves that display consistent variability. This

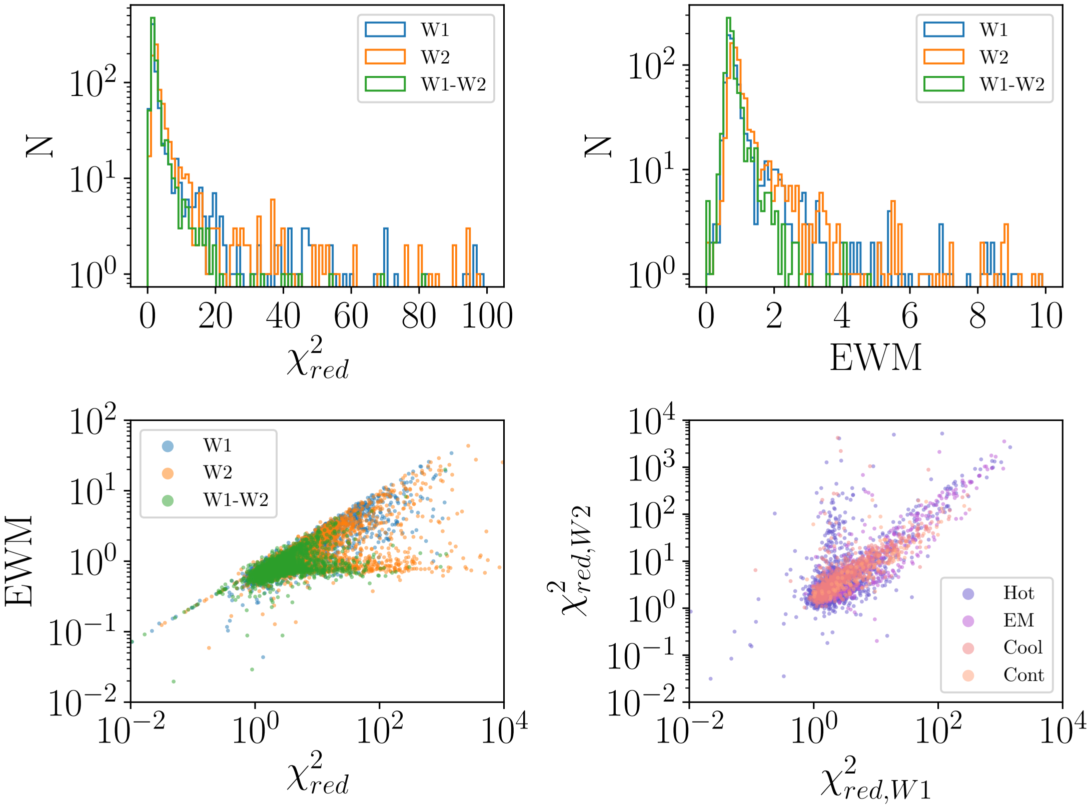

Figure 4.7 *Top row*: Distribution of derived $\log \chi^2_{red}$ (left) and $\log EWM$ (right) values in the $W1$ (blue), $W2$ (orange), and $W1 - W2$ (green) lightcurves. *Bottom row*: Scatter plots comparing different amplitude metrics. The left panel shows EWM vs. χ^2_{red} in the $W1$, $W2$, and $W1 - W2$ lightcurves (using the same color-coding). The right panel shows χ^2_{red} in $W2$ vs. χ^2_{red} in $W1$, with each point colored by its coarse label. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

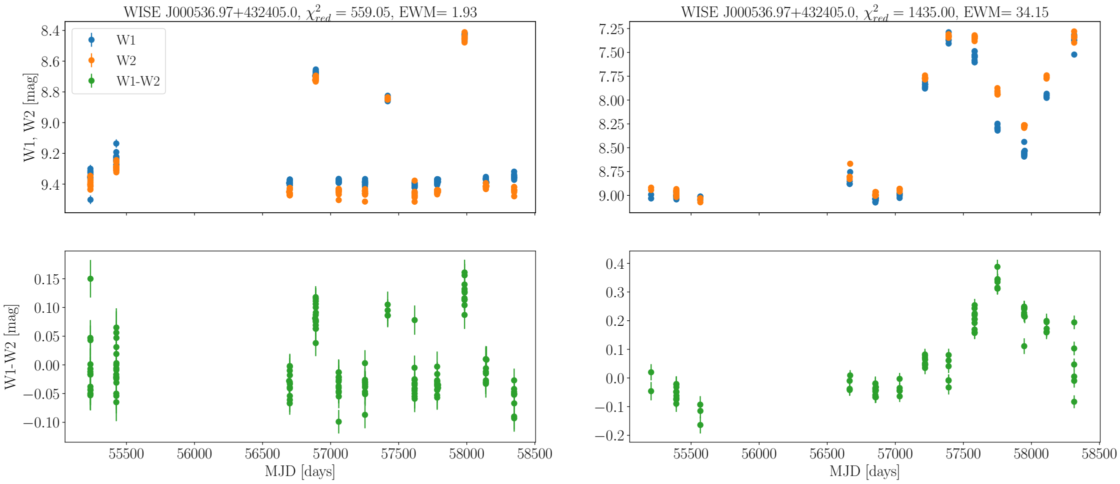

Figure 4.8 *Left*: Lightcurves in $W1$ and $W2$ (top) and $W1 - W2$ bottom for an example star with high χ^2_{red} and low EWM in $W1$. Three flaring events are clearly visible, during which the star becomes slightly redder; otherwise, the star only displays modest variability. *Left*: Similar, but for a star with high values of both χ^2_{red} and EWM in $W1$. By contrast, this star shows consistent, smooth variability.

can be seen in Figure 4.8, which shows two example lightcurves, one with high χ^2_{red} and low EWM (left panels), and one with both high χ^2_{red} and high EWM (right panels); the first star shows three short-lived flaring events during which the star becomes slightly redder, superimposed over low-level variability, while the second shows consistent smooth variability in both WISE bands and the color lightcurve.

The bottom right panel of Figure 4.7 shows a scatter plot of χ^2_{red} in $W1$ vs. in $W2$, with each point colored by its coarse label. A distinct branch of stars that are much more variable in $W2$ than $W1$ is clearly evident; oddly, the distributions of classes, EWM, and broadband colors in this branch are consistent with the whole sample, and no similar branch exists in the measured EWM values. Visual inspection of the lightcurves of stars with $\chi^2_{red} < 10$ in $W1$ and $\chi^2_{red} > 100$ in $W2$ shows that these stars appear to have higher signal to noise $W2$ measurements than $W1$, and have one observation during which the star apparently becomes considerably redder, achieving $W1 - W2$ values as high as ~ 4; an example of this behavior is shown in Figure 4.9. Examining

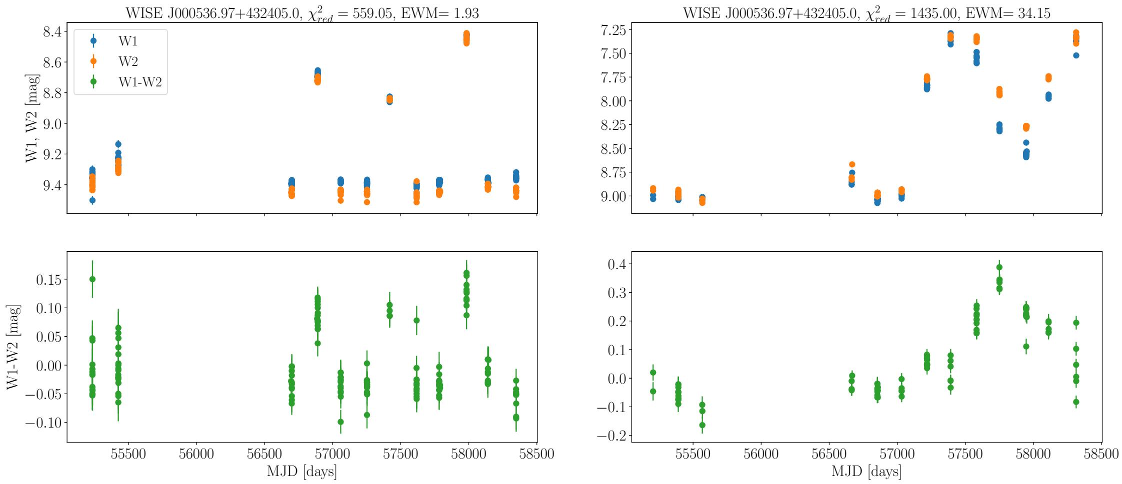

Figure 4.9 Lightcurve in $W1$ and $W2$ (top) and $W1 - W2$ bottom for an example star with high χ^2_{red} in $W2$, but low χ^2_{red} in $W1$. The star is mostly non-variable except for a single observation during which it becomes ∼0.2 mag brighter in both bands, and ∼1.2 mag redder.

the times at which these extreme reddening events occur shows a preference for times during the cryogenic WISE survey, implying that this behavior is likely instrumental in origin, despite our filtering using the provided quality flags. Nonetheless, we include χ^2_{red} as it does not map perfectly onto EWM, and only 98 stars fall in this regime. We do not yet know whether χ^2_{red}, EWM, both metrics, or neither are useful features for classification, so we keep both with the intent of exploring their importance below.

Timescale

Many methods exist for estimating dominant timescales in lightcurves. [46] use the Lomb-Scargle Periodogram [166, 256] to search for periodic variables. However, this approach suffers from numerous, well-known issues (including accurate period recovery at low signal to noise), and false peaks can easily be mistaken for real timescales, especially in highly-irregularly sampled data, as is the case for the WISE lightcurves. [270] use a Gaussian process interpolation scheme coupled with a wavelet analysis to estimate timescales in massive stars in M31 observed by the Palomar Transient Factory (PTF). However, with so few data points, we found it difficult to obtain a reliable fit with a Gaussian process, and even when the fit was successful, the resulting interpolant had a large standard deviation in between WISE visits. The resulting measurements of the characteristic timescale was more reflective of the kernel used.

Instead, we turn to a a spline-based interpolation method, which is analogous to certain Gaussian process methods [135]. We first subtract half of the sum of the times of the first and last available observations, so that the lightcurve is centered at $t = 0$. We then bin the observations in each visit. Visits are defined as sets of points separated in time by less than a defined threshold. Due to the WISE scanning law, some stars near the ecliptic poles have visits separated by less than the typical ~ 180 days. Therefore, we adopt 50 days as the threshold for visits. Two stars in our sample are close enough to the ecliptic pole to be observed nearly continuously such that we erroneously record two “visits”: one each during the cryogenic and post-cryogenic surveys. However, neither star is strongly variable and thus this small edge case does not substantially impact our subsequent analyses.

For all observations in a given visit, we calculate the mean time and $W1$/$W2$/$W1 - W2$ measurement. We use `scipy.interpolate.splrep` in Python to find the 3$^{\mathrm{rd}}$-order basis spline (a.k.a. B-spline, which performs a spline fit using spline basis functions, [61]) representation of the binned lightcurves, adopting a smoothing factor `s=10`. This returns the knots, B-spline coefficients, and degree of the spline. By definition, 3$^{\mathrm{rd}}$-order splines are differentiable, so we use `scipy.interpolate.splev` to evaluate the first derivative of the spline interpolant, and find the times when the derivative changes sign — i.e., when the lightcurve reaches a maximum or minimum. As metrics of the characteristic timescale of the lightcurve, we calculate the frequency of zero-crossings of the first derivative of the spline interpolant, ν_0 — calculated as the number of times the derivative passes through zero, divided by the time baseline of the lightcurve — $\langle \Delta t \rangle$, the mean of the differences between successive zero-crossings, and the standard deviation of the differences between successive zero-crossings, $\sigma_{\Delta t}$. For stars with fewer than four visits, we automatically assign $\nu_0 = \langle \Delta t \rangle = \sigma_{\Delta t} = \texttt{NaN}$. The right panels of Figure 4.6 show this process on the $W1$ lightcurve plotted in the left panels. The blue points are the binned $W1$ measurements (the errors are smaller than the size of the points), and the dotted black lines are the spline interpolant (top right) and corresponding derivative (bottom right).

While this is a simple method that yields multiple estimates of variability timescale, it is important to note that it is dependent on both the variability amplitude and the sampling. For example, a non-variable object whose lightcurve is poorly sampled may appear to be variable due to measurement noise (which does not have a characteristic timescale), and the derived timescale from this method will thus be more reflective of the sampling than anything else. Thankfully, in many cases such a variable would have a low EWM value. However, it is possible that a star may enter a period of low-amplitude variability resulting in false zero-crossings of the first derivative of the spline interpolant (e.g., the first three WISE visits in lightcurve in Figure 4.6; it is possible that the first few zero crossings in the bottom right panel may not be real.) These systematics are difficult to work around in sparsely sampled lightcurves, and are an important caveat to keep in mind.

4.3 Machine Learning

4.3.1 Classifier Selection

We wish to train an algorithm that takes as input the broadband photometry and variability metrics derived for our sample, and outputs spectral type classifications. Many machine learning classifiers exist; of these, we wish to choose a flexible model with well-understood mathematics, while avoiding techniques like neural networks that can be difficult to interpret. Of the classifiers available in the `sklearn` package, we decided to test a Random Forest (RF) classifier [30] — which consists of a collection of decision trees trained on random subsets of samples and features — a Support Vector Machine (SVM) classifier [50] — which identifies hyperplanes in the feature space that separate different classes — and a Gaussian process (GP) classifier — which models the function determining the probability of a star being a given class at a location in the feature space as a multidimensional Gaussian distribution whose properties are determined entirely by a covariance function (a.k.a. a kernel function), coupled with a linking function (usually the logit function) to make discrete class predictions [238]. We refer the reader to these publications, as well as to the

Table 4.3. List of features passed to our machine learning classifiers, as well as clarifying definitions where relevant. WISE photometry used to calculate colors and magnitudes is from the ALLWISE data release [56]. All variability metrics are calculated from the WISE $W1$, $W2$, and $W1 - W2$ lightcurves.

Feature	Definition
Colors & Magnitudes	
M_G	Absolute magnitude in *Gaia* G band.
$G - J$	From *Gaia* and 2MASS photometry.
$J - H$	From 2MASS photometry.
$H - K_s$	From 2MASS photometry.
$K_s - W1$	From 2MASS and WISE photometry.
$W1 - W2$	From WISE photometry.
$W2 - W3$	From WISE photometry.
$W3 - W4$	From WISE photometry.
M_{W1}	Absolute magnitude in WISE $W1$ band.
Variability Metrics	
$\log \chi^2_{red}$	Log of the reduced χ^2.
log EWM	Log of the error-weighted Median Absolute Deviation.
ν_0	Frequency of zero-crossings of the first derivative of the spline interpolant.
$\log\langle \Delta t \rangle$	Log of the average time between zero-crossings.
$\log \sigma_{\Delta t}$	Log of the standard deviation of zero-crossing times.

`sklearn` documentation[4] for the mathematics and implementation details of each classifier. In the multi-class case, a collection of classifiers are trained on each possible pair of classes (one-versus-one or "ovo"), generating a total of $N_{classes}(N_{classes} - 1)/2$ classifiers where $N_{classes}$ is the number of classes. Labels are assigned to test samples by allowing each classifier to vote, and the label with the most votes is chosen [139].

Each type of classifier has a number of hyperparameters that affect the performance of the classifier. For the RF classifier, `n_estimators` specifies the number of trees in the forest, `max_depth` specifies how many branches each decision tree in the forest can have, and `max_features`

---

[4] https://scikit-learn.org/stable/index.html

specifies the maximum number of features each tree is trained on. We also set `class_weight=balanced`, which weighs samples when fitting to account for the different frequencies of each class in the data.

For the SVM classifier (SVC), C is a regularization parameter that governs the tradeoff between maximizing the margin and misclassifications in the training set. Higher values of C will force the SVC to correctly classify every point, resulting in poor generalization (i.e. overfitting). The SVC requires that the distance between two points in the feature space is defined as the inner product of two vectors in the feature space, $\langle \vec{X}_i, \vec{X}_j \rangle$. Because the boundaries between classes in our sample are not guaranteed to be linear, one can project the samples into a much higher dimension space via a mapping function, Φ, where distances between two vectors in this space are calculated as $\langle \vec{Z}_i, \vec{Z}_j \rangle = \langle \Phi(\vec{X}_i), \Phi(\vec{X}_j) \rangle$. In reality, the transformed feature space can be incredibly high dimensional, and explicitly mapping the data into this high-dimensional space is computationally inefficient. Instead, we can adopt a kernel function, K that defines distances in the higher dimensional space, e.g., $K(\vec{X}_i, \vec{X}_j) = \langle \vec{Z}_i, \vec{Z}_j \rangle$. Because the kernel only takes the measured features as input and outputs a number, using a kernel function implicitly maps the input feature space into a high dimensional space without specifying Φ.

Common choices of the kernel function include a linear kernel (i.e., the Euclidean distance between individual samples in the feature space) and the "radial basis function" (RBF) kernel:

$$K(\vec{X}_i, \vec{X}_j) = e^{-\gamma ||\vec{X}_i - \vec{X}_j||^2} \tag{4.6}$$

where γ governs the influence of the kernel function; lower values result in increasingly linear boundaries, while high values result in the decision function being entirely depended on individual points, creating small islands of a given class centered on each training point. The advantage of non-linear models like a SVM with a RBF kernel over linear methods is that the decision boundaries can be much more flexible; the tradeoff is that the contribution of individual features to the classifier cannot be easily calculated without the unknown function Φ (see §4.3.3). The "optimum" kernel and hyperparameters are chosen via a cross-validation strategy described below. Finally, we also set `class_weight=balanced` for the SVC, which automatically sets the value C for

Table 4.4. Feature values and assigned labels for all stars in our sample, ordered by Right Ascension. Missing numbers are indicated with "-".

Common Name	M_G [mag]	$G-J$ [mag]	$W1-W2$ [mag]	$\log\chi^2_{red,W1}$	$\log\langle\Delta t\rangle_{W1}$ [d]	Label	Coarse Label
HD 236270	-3.54	0.35	0.14	0.389	-	RSG	Cool
LS I +64 10	-3.41	0.72	0.00	-	-	OBA	Hot
LS I +60 69	-3.22	0.88	-0.03	0.061	-	OBAe	EM
BD+62 2353	-4.27	0.47	-0.04	0.110	2.816	SupergiantOBA	Hot
HD 73	-3.29	-0.66	-0.05	3.157	2.701	EvolvedOBA	Hot
HD 240496	-3.84	1.01	0.01	-0.064	2.764	OBA	Hot
LS I +59 30	-3.34	0.67	-0.04	0.234	-	OBA	Hot
BD+57 2870	-4.53	1.17	0.00	0.046	3.135	SupergiantOBA	Hot
BD+62 1	-3.31	0.90	0.27	1.622	2.813	EvolvedOBA	Hot

Note. — This table is published in its entirety in the machine-readable format. A portion is shown here for guidance regarding its form and content. We note that only a subset of the features listed in Table 4.3 are shown here. All features are listed in the machine-readable version.

class i to $CN_{\text{samples}}/(N_{\text{classes}}N_i)$ where N_{samples} is the size of the sample, N_{classes} is the number of classes, and N_i is the number of objects in the sample belonging to the class. This serves to weight rarer classes more heavily. The GP classifier's only hyperparameter is a choice of kernel, which defines the covariance function of the Gaussian process.

To fit our classifiers, we first remove all miscellaneous variables and unknown/candidates, as well as known binaries, and one star with bad J photometry, the Be star HD 53032. For features, we use the intrinsic calculated value of M_G, as well as (uncorrected for extinction) $G-J$, $J-H$, $H-K_s$, K_s-W1, $W1-W2$, $W3-W4$, and χ^2_{red}, EWM, ν_0, $\langle\Delta t\rangle$, and $\sigma_{\Delta t}$ in all three lightcurves — we indicate the WISE band or color that each variability metric corresponds to with a subscript hereafter. The input features and a brief description where relevant are listed in Table 4.3. Because χ^2_{red}, EWM, $\langle\Delta t\rangle$, and $\sigma_{\Delta t}$ have significant dynamic range, we use the base-10 logarithm of these features. Features values as well as labels for each star in our sample are given in Table 4.4. We note that only a subset of the features from Table 4.3 are listed here due to the number of features. The table in its entirety will be made available in a machine-readable format.

We now randomly split our sample into a training set with 70% of the samples, and a test set with the remaining 30%, using a stratification strategy to ensure the proportions of the classes in both sets are equal. The test set is withheld until we are ready to assess the performance of the chosen classifier. We then use `sklearn.preprocessing.StandardScaler` in Python to scale the training data such that each feature has 0 mean and unit variance. Because the data have missing values, we then use `sklearn.impute.IterativeImputer` which uses a Bayesian ridge regression to predict and replace missing values.

To test the accuracy of the imputer, we select only the rows from the training set with no missing data. For each feature with missing data ($\log \chi^2_{red,W1-W2}$, $\log \mathrm{EWM}_{W1-W2}$, $\log \sigma_{\Delta t,W1}$, $\log \sigma_{\Delta t,W2}$, and $\log \sigma_{\Delta t,W1-W2}$) we randomly choose 200 objects, replace the value of the feature with `NaN` for only these objects, transform the data using the scaler and imputer, and calculate the fractional error between the true value and the imputed value. The returned fractional errors for each feature centered around 0, and had a low scatter with the exception of $\log \mathrm{EWM}_{W1-W2}$. However, this has little impact on the classifier, as only two objects in the actual training set have missing values for $\log \mathrm{EWM}_{W1-W2}$. We also repeated this procedure for coarse labels: we select 200 random objects with a given coarse label, replace a random feature from the list of features with missing data with `NaN` for each object, and again calculate the fractional error between the true and imputed values. We find that the imputer performs poorly on Cool and Contaminant stars. Given that all of the features with missing data are linked to variability, a significant fraction of red supergiants display high amplitude variability [46], the MIR variability of AGB stars (which make up the bulk of the Contaminants) is higher than RSGs in a given magnitude range [327], and our sample of Cool and Contaminant stars contains objects in quite different evolutionary states that nevertheless have similar colors and magnitudes, it is unsurprising that the imputer is unable to predict the variability properties of these stars. In §4.3.3, we will discuss the impact of these features on the overall performance of our classifier.

For each classifier, we then initialize a corresponding `sklearn` classifier object (e.g., `sklearn.svm.SVC`). To settle on the best values for the hyperparameters we use `sklearn.model_selection.GridSearchCV` to perform a cross-validation search on a grid of hyperparameters, using a stratified K-fold strategy with $k = 5$ to ensure that each fold has a representative distribution of classes. For the RF, we search for `n_estimators` between 10 and 150 in steps of 10, `max_depth` between 10 and 100 in steps of 10, and allow `max_features` to be either `sqrt`, `log2`, or `None` (where the maximum number of features individual trees are trained on is the square root of, base-2 logarithm of, or equal to the number of features, see the documentation for details). We also allow `max_depth` to take on the default value (`None`), such that individual trees can be grown until each leaf only contains one sample.

For the SVC, we search for values of C on a logarithmic grid with 1 dex spacing between 0.01 and 100, and for the RBF kernel, we search for values of γ on a similar grid between 0.01 and 10. Additionally, we allow γ to be the default values of `1/n_features` (where `n_features` is the number of features). For the GP classifier, we only vary the kernel, as `sklearn` automatically optimizes the kernel hyperparameters. We let the kernel be either linear, RBF, or the default (a special case of the RBF kernel with the length scale equal to 1).

Each classifier object has a default method to score each set of hyperparameters, e.g. the accuracy of predicted labels compared to true labels. However, the classes in our training set are unbalanced (e.g. Figure 4.4), so inaccurately classifying every single LBV, for example, would have little impact on the overall accuracy of the classifier. To account for this, we instead use the *balanced accuracy* [204, 107], which weighs each sample by the frequency of that sample's class in the training set. Other options for scoring criteria exist, including some that help maximize the classifier's precision such as the weighted F_1 score and Cohen's kappa [45]. We experimented with using these scores, and found that using the balanced accuracy minimizes misclassifications across all classes (reflected in the diagonal in the left panels of Figures 4.11 and 4.14). Note that this choice implicitly selects a classifier that performs well across all classes, and is not optimized for specific classes. Future work will explore the possibility of tuning a classifier to find specific classes of rare stars.

Finally, we explore three variations of a voting classifier. Such a classifier consists of an ensemble of individual classifiers, each which "votes" by assigning a class to a given sample. The final assigned class can either be chosen with a "hard" (the class with the most votes wins) or "soft" (class assignments are weighted by the probabilities output by each classifier) strategy. We construct two voting classifiers that each use a different voting strategy, using RF, SVC, and GP classifiers as the individual components. We refer to these as the Voting (Hard) and Voting (Soft) classifiers. We also make a third voting classifier that also uses a soft voting strategy, but the votes from each component classifier weighted by the balanced accuracy determined via cross-validation. We refer to this as the Voting (Weighted) classifier. We score each voting classifier by averaging the balanced accuracy taken from five stratified folds of the data. Figure 4.10 shows the balanced accuracy for the three optimized classifiers, as well as the three voting classifiers; the SVC performs "best," though all classifiers return similarly low balanced accuracies between $\sim$0.4 and 0.55. Both the Voting (Soft) and Voting (Weighted) classifiers perform comparably with the worst classifier, the GP. This is due to the fact that, while the SVC often selects one individual class with high probability, both the RF and GP tend to select multiple classes with high probability (with the GP sometimes selecting all classes with roughly equal probability, usually slightly favoring the classes selected by the RF). This can result in both the RF and GP voting for the wrong class with higher probability than the correct vote from the SVC, leading to the poor observed performance.

4.3.2 SVC Performance

The procedure above results in values for the SVC hyperparameters of `kernel = linear`, and `C = 0.01`. With these hyperparameters, we fit the SVC to the training set, use the scaler and imputer that were previously fit to the training set to transform the test set, and use the SVC to predict the labels of the test set. The left-hand panel of Figure 4.11 shows the raw number of stars in the test set with the true label given on the y-axis, and the predicted label given on the x-axis. The center and right-hand panels show the *confusion* and *efficiency* matrices, derived by taking the raw matrix and normalizing each row/column by the total number of stars in that row/column, respectively. Thus the i, j entry in the confusion matrix (center panel) corresponds to the fraction of objects in the test set belonging to class i (shown on the y-axis) that are assigned class j (shown on the x-axis). Entries along the diagonal are the completeness (also called the recall in some contexts), i.e., the percentage of a given class that is accurately recovered by the classifier. The i, j entry in the efficiency matrix (right-hand panel) is the fraction of objects in the test set classified as j, that belong to class i. Entries along the diagonal are equivalent to the precision (equivalent

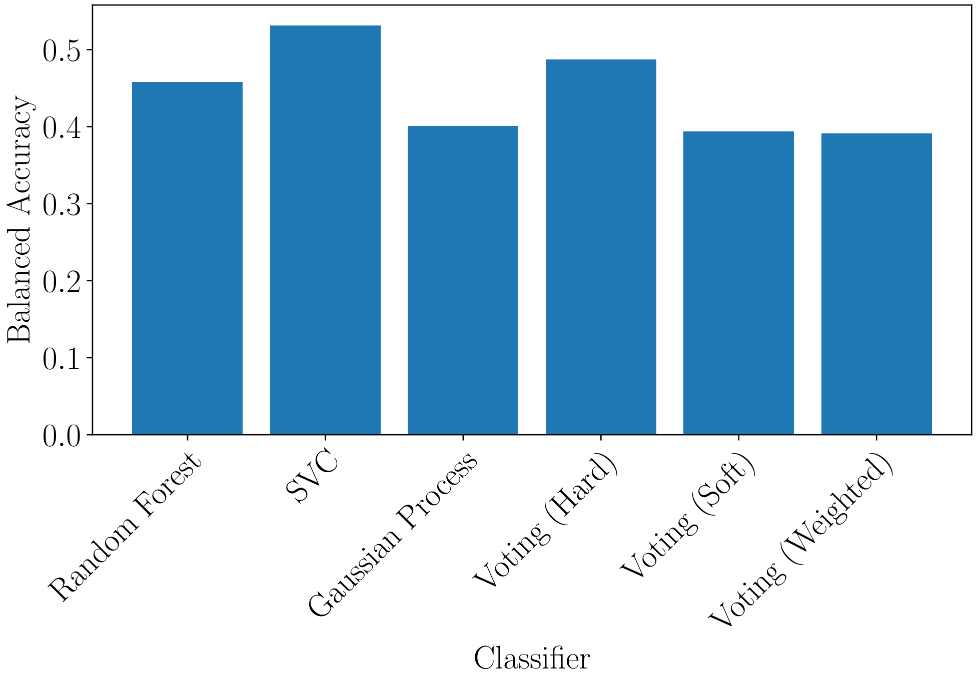

Figure 4.10 Balanced accuracy for each optimized classifier, averaged over five foldings of the data. The SVC is the best overall, with a balanced accuracy of 0.53. Among the three voting classifiers, the Voting (Hard) classifier performs best with a balanced accuracy of 0.49, still below the SVC. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

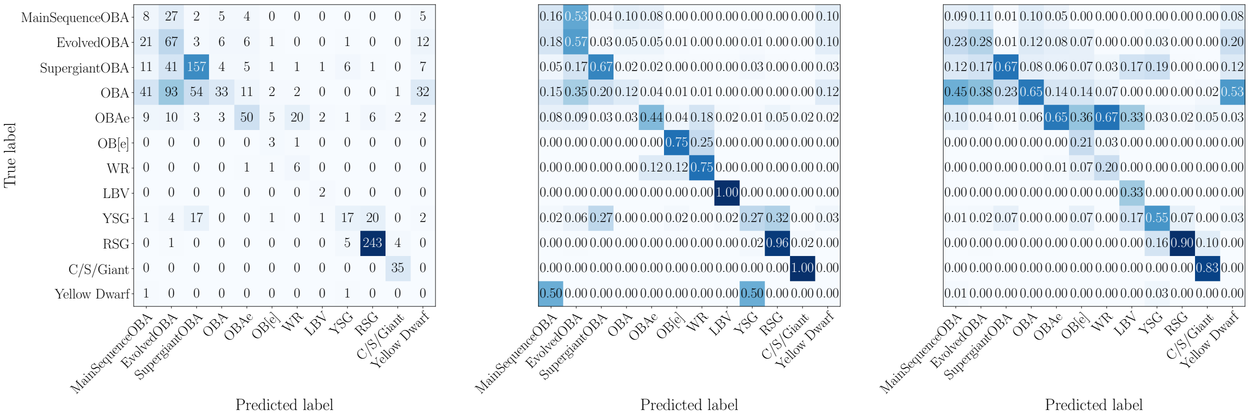

	MainSequenceOBA	EvolvedOBA	SupergiantOBA	OBA	OBAe	OB[e]	WR	LBV	YSG	RSG	C/S/Giant	Yellow Dwarf
MainSequenceOBA	8	27	2	5	4	0	0	0	0	0	0	5
EvolvedOBA	21	67	3	6	6	1	0	0	1	0	0	12
SupergiantOBA	11	41	157	4	5	1	1	1	6	1	0	7
OBA	41	93	54	33	11	2	2	0	0	0	1	32
OBAe	9	10	3	3	50	5	20	2	1	6	2	2
OB[e]	0	0	0	0	0	3	1	0	0	0	0	0
WR	0	0	0	0	1	1	6	0	0	0	0	0
LBV	0	0	0	0	0	0	0	2	0	0	0	0
YSG	1	4	17	0	0	1	0	1	17	20	0	2
RSG	0	1	0	0	0	0	0	0	5	243	4	0
C/S/Giant	0	0	0	0	0	0	0	0	0	0	35	0
Yellow Dwarf	1	0	0	0	0	0	0	0	1	0	0	0

	MainSequenceOBA	EvolvedOBA	SupergiantOBA	OBA	OBAe	OB[e]	WR	LBV	YSG	RSG	C/S/Giant	Yellow Dwarf
MainSequenceOBA	0.16	0.53	0.04	0.10	0.08	0.00	0.00	0.00	0.00	0.00	0.00	0.10
EvolvedOBA	0.18	0.57	0.03	0.05	0.05	0.01	0.00	0.00	0.01	0.00	0.00	0.10
SupergiantOBA	0.05	0.17	0.67	0.02	0.02	0.00	0.00	0.00	0.03	0.00	0.00	0.03
OBA	0.15	0.35	0.20	0.12	0.04	0.01	0.01	0.00	0.00	0.00	0.00	0.12
OBAe	0.08	0.09	0.03	0.03	0.44	0.04	0.18	0.02	0.01	0.05	0.02	0.02
OB[e]	0.00	0.00	0.00	0.00	0.00	0.75	0.25	0.00	0.00	0.00	0.00	0.00
WR	0.00	0.00	0.00	0.00	0.12	0.12	0.75	0.00	0.00	0.00	0.00	0.00
LBV	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00
YSG	0.02	0.06	0.27	0.00	0.00	0.02	0.00	0.02	0.27	0.32	0.00	0.03
RSG	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.02	0.96	0.02	0.00
C/S/Giant	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	0.00
Yellow Dwarf	0.50	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.50	0.00	0.00	0.00

	MainSequenceOBA	EvolvedOBA	SupergiantOBA	OBA	OBAe	OB[e]	WR	LBV	YSG	RSG	C/S/Giant	Yellow Dwarf
MainSequenceOBA	0.09	0.11	0.01	0.10	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.08
EvolvedOBA	0.23	0.28	0.01	0.12	0.08	0.07	0.00	0.00	0.03	0.00	0.00	0.20
SupergiantOBA	0.12	0.17	0.67	0.08	0.06	0.07	0.03	0.17	0.19	0.00	0.00	0.12
OBA	0.45	0.38	0.23	0.65	0.14	0.14	0.07	0.00	0.00	0.00	0.02	0.53
OBAe	0.10	0.04	0.01	0.06	0.65	0.36	0.67	0.33	0.03	0.02	0.05	0.03
OB[e]	0.00	0.00	0.00	0.00	0.00	0.21	0.03	0.00	0.00	0.00	0.00	0.00
WR	0.00	0.00	0.00	0.00	0.01	0.07	0.20	0.00	0.00	0.00	0.00	0.00
LBV	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.33	0.00	0.00	0.00	0.00
YSG	0.01	0.02	0.07	0.00	0.00	0.07	0.00	0.17	0.55	0.07	0.00	0.03
RSG	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.16	0.90	0.10	0.00
C/S/Giant	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.83	0.00
Yellow Dwarf	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.03	0.00	0.00	0.00

Figure 4.11 *Left*: Matrix showing the number of stars in the test set with true label indicated on the y-axis that are assigned the label on the x-axis. *Center*: Confusion matrix for the SVC, calculated by normalizing each row of the left panel by the total number of stars in that row. Values correspond to the *fraction* of samples in the test set with true label indicated on the y-axis that are assigned the label on the x-axis, such that the values along the diagonal are the fraction of each class that is correctly classified. *Right*: Efficiency matrix for the SVC, calculated by normalizing each column of the left panel by the total number of stars in that column. Values in each box correspond to the fraction of samples in the test set assigned the label on the x-axis that belong to the class on the y-axis, such that the values along the diagonal correspond to the precision (one minus the contamination). Darker colors in all panels correspond to more/a higher fraction of stars. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

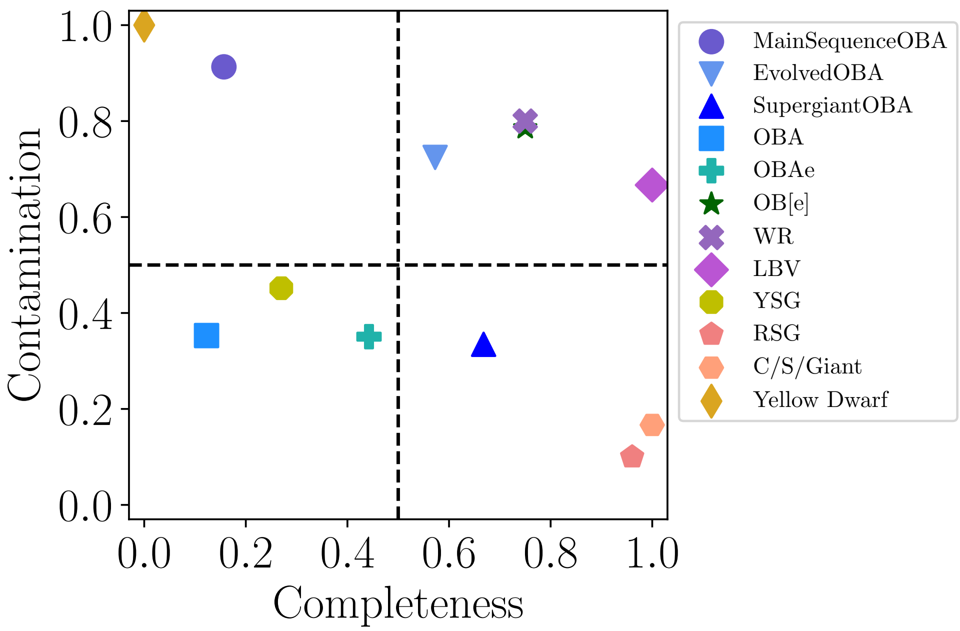

Figure 4.12 Completeness versus contamination of each class in the test set, as classified by the SVC. A high completeness value implies members of that class are accurately classified, while a low contamination value implies that an object classified as such is likely to belong to that class. The figure is roughly divided into four quadrants; stars with classes in the bottom-right quadrant can be considered to be well-classified, in the sense that they have high completeness and low contamination. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

to one minus the contamination), i.e., the percentage of an observed class that is made up of true members of that class. Figure 4.12 shows the completeness versus the contamination for each class. Completeness is just the diagonal of the corresponding row/column of the confusion matrix, and contamination is one minus the diagonal of the corresponding row in the efficiency matrix.[5]

The SVC performs poorly on non-supergiant OBA stars. This is perhaps unsurprising given that both the observed colors and interior structures of OBA stars as they evolve from the zero-age main

---

[5] We note that a variety of terms are used in the classification problem, some of which (i.e., completeness and contamination) are familiar to astronomy, which we briefly summarize here. The completeness (or recall) is also referred to as the true positive rate in the binary classification case. The accuracy refers to the the sum of the diagonal in the left panel of Figure 4.11 divided by the number of objects in the test set. The contamination is also called the false positive rate in the binary classification case; the precision refers to one minus the contamination.

sequence (ZAMS) to the terminal-age main sequence (TAMS) do not change drastically compared to the much more evolved states that we also consider. The classifier classifies OBAe stars with somewhat lower contamination compared to main sequence and evolved OBA stars, though with comparably low completeness. True OBAe stars are either misclassified as other types of OBA star or as WRs, while stars falsely labeled as OBAe are mostly true OBA stars, with the exception of one true WR. 75% of WR stars are recovered, but only 6/30 stars identified as WRs in the test set are true WRs; given the importance of WRs for both the physics of mass loss and studying evolved massive stellar populations [68, 70], future work will focus on developing a classifier specifically for identifying WR stars.

All LBVs in the test set are recovered; while such high accuracy is often seen as a sign of overfitting, we choose not to focus on this subclass, given both the disputed evolutionary status of LBVs [268, 118, 1] and the fact that only two LBVs exist in the training set. Yellow supergiants are only classified with 27% accuracy. As discussed in §4.1 and shown in Figure 4.5, the yellow supergiant label is assigned to stars with optical colors consistent with hot stars as well as RSGs. This is reflected in the types of stars that YSGs are mistaken for, as well as the stars that are mistaken for YSGs. Overall, the classifier performs best on the coolest stars in the sample. RSGs are classified with 96% accuracy, and only 10% contamination. Meanwhile the classifier performs exceptionally well at identifying low mass contaminants, at the cost of misclassifying four RSGs, two OBAe stars, and one OBA star.

Overall, an SVC trained on these refined labels appears to have little use. With the exception of RSGs and low-mass giants, the remaining classes have low accuracy, high contamination, or both. We nonetheless use the SVC to predict labels for the 2550 stars initially labeled as “Miscellaneous Variable” or “Unknown/Candidate.” We identify 79 candidate RSGs and 36 candidate C/S/Giant stars, of which we expect $\sim$71 and 30 to be genuine, respectively, given the efficiency matrix. We list the candidate RSGs in Table 4.5. A small spectroscopic observing campaign would easily confirm the ability of this classifier to correctly identify RSGs and low-mass giants.

Table 4.5. Common names and coordinates of stars predicted to be RSGs by the SVC trained on refined labels.

Common Name	R.A. [deg]	Dec [deg]
SP77 48-11	81.07900941	−70.43417562
WISE J185608.58-163255.1	284.03575762	−16.54867009
W61 19-14	83.07777354	−67.52941938
OGLE BRIGHT-LMC-MISC-169	72.94711333	−69.32348227
WISE J194127.64+385155.3	295.36520609	38.86536427
NGC 2004 BBBC 431	82.69161818	−67.29036242
[KWV2015] J045626.51-692350.6	74.11062177	−69.39740804
W61 6-54	85.54018822	−69.21978048
WISE J064232.30-715243.3	100.63462144	−71.87871874
W61 6-34	85.51621031	−69.21870016

Note. — This table is published in its entirety in the machine-readable format. A portion is shown here for guidance regarding its form and content.

4.3.3 Performance on Coarse Labelling

Examining Figures 4.11 and 4.12, we see that, while the classifier is not especially accurate except for the coolest stars, the classifier is *roughly* useful for sorting the test set into broad categories: different types of OBA stars are mostly (mis)classified as other classes of OBA stars; the same is true for emission line stars (OBAe, OB[e], WR, and LBV) and cool stars (YSG, RSG, C/S/Giant).

For this reason, we also utilize the coarse labels introduced in §4.1. We repeat the entire process described above, beginning with the selection of the classifier. Figure 4.13 shows the balanced accuracy for each of the classifiers discussed above, trained on the coarse labels, using a five-fold cross-validation to optimize the hyperparameters of each classifier. We find that once again, the SVC yields the highest balanced accuracy (0.876).

We keep the same scaled and imputed training and testing sets, and perform a five-fold cross-validation as before to find the optimal hyperparameters, which are `kernel=rbf`, `C = 1`, and `gamma = 1/n_features`. With these hyperparameters, we fit the SVC to the training set before predicting labels for the test set. Figure 4.14 shows the confusion and efficiency matrices similar to Figure 4.11, while Figure 4.15 shows the completeness versus the contamination similar

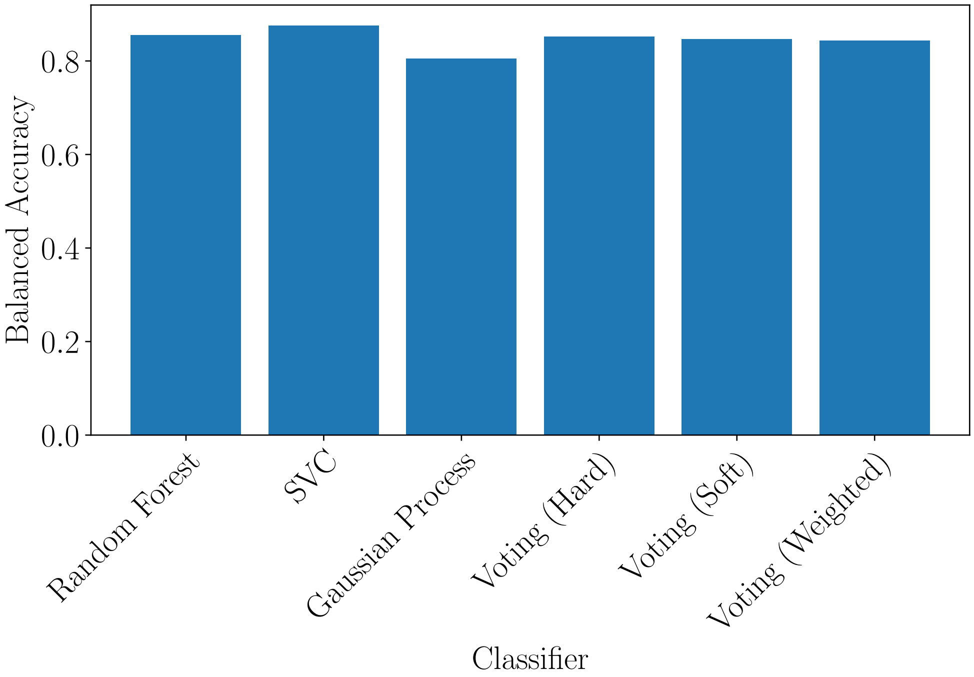

Figure 4.13 Similar to Figure 4.10 for classifiers trained on the coarse labels. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

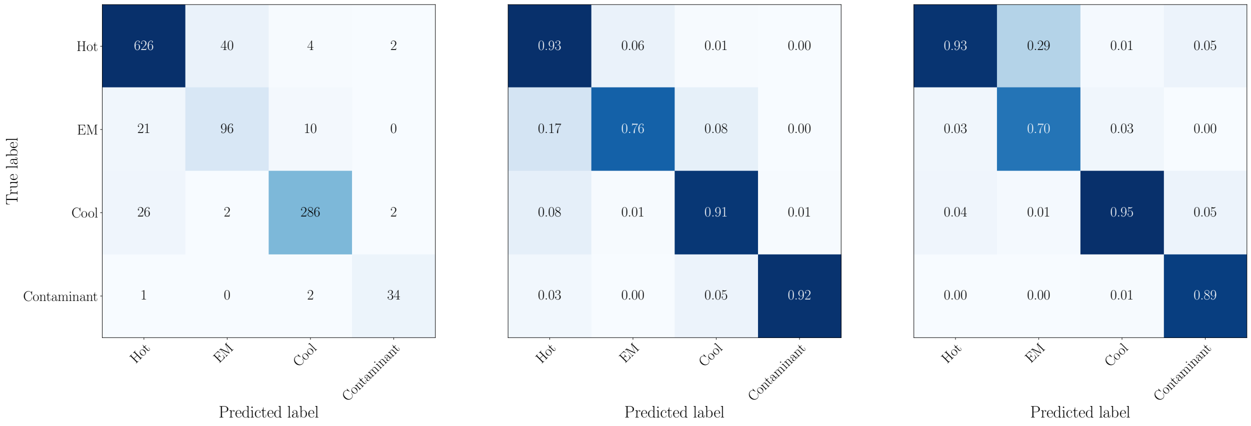

Figure 4.14 Similar to Figure 4.11 for the SVC trained using the coarse labels. Note that significantly more stars fall along the diagonal of each plot, reflecting the improved performance of the SVC on the coarse labels. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

to Figure 4.12. All told, the SVC performs significantly better compared to the classifier trained on the refined labels, recovering all classes with $\geq 75\%$ completeness and $\leq 30\%$ contamination.

Of the emission line stars that are correctly identified, 83 are OBAe stars, three are OB[e] stars, eight are WRs, and two are LBVs. This is 73%, 75%, 100%, and 100%, respectively of these stars that are in the test set, implying that the performance of the SVC on emission line stars is not dominated entirely by OBAe stars (which comprise the majority of emission line stars in the test set). Of the stars mislabeled as contaminants, two are OBA stars and two are RSGs. One true C/S/Giant star, and two Yellow Dwarfs are misclassified.

We then use the SVC to predict the coarse labels for the same 2550 stars as above. Figure 4.16 shows the distribution of these predicted labels. The majority (2472 stars) are labeled as Hot. 63 stars are labeled as Cool, three of which are already identified in SIMBAD as candidate AGBs or RGBs. 14 of these stars are labeled as emission line stars, of which 9-10 are likely to actually be emission line stars, assuming 30% contamination. We list all 2550 stars' common names, coordinates, and predicted coarse label in Table A.1.

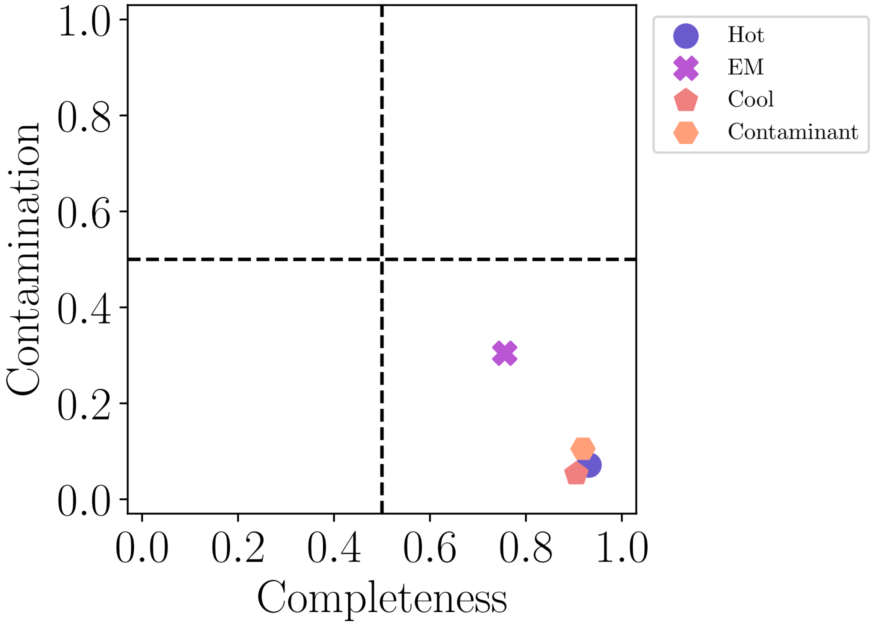

Figure 4.15 Similar to Figure 4.12 for the SVC trained using the coarse labels. All coarse classes have high completeness and low contamination. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

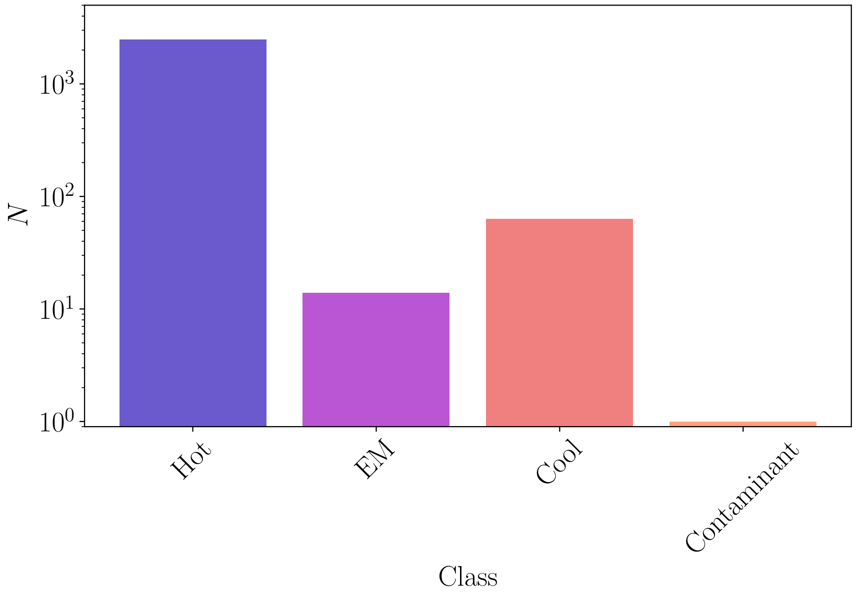

Figure 4.16 Distribution of coarse labels assigned to 2550 stars with no previously known class. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

Feature Importance

We can also identify which features contribute most to the overall performance of the classifier on the coarse labels. To do this, we initialize a new `SVC` object with the same hyperparameters, and perform a "greedy search" over features, defined as follows: For each feature in the scaled and imputed training set, we train the SVC on just this feature across five stratified folds of the training set, and record the average and standard deviation of the balanced accuracy. We select the feature that yields the highest average balanced accuracy. This has the advantage of ensuring the contribution of each feature to the balanced accuracy is stable across subsets of the data. We then train the SVC on all combinations of this feature and the remaining features, selecting which combination again yields the highest average balanced accuracy. This process is repeated until all features are used.

Figure 4.17 illustrates this process. The x-axis shows the feature that is selected at each stage of the greedy search. The y-axis shows the mean balanced accuracy of the SVC at that stage.

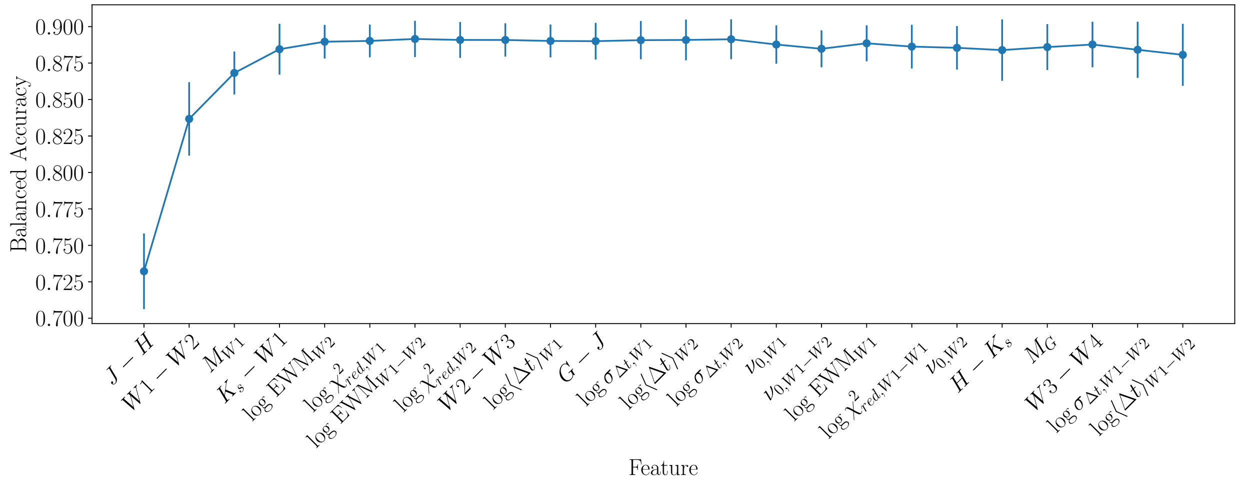

Figure 4.17 Mean balanced accuracy of the SVC for coarse labels trained on successively added features, calculated from five stratified folds of the data. The balanced accuracy reaches a maximum after the first seven features. Errorbars indicate the standard deviation of the balanced accuracy across folds. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

Errorbars show the standard deviation of the balanced accuracy across the five folds of the training set. The balanced accuracy reaches a maximum after the first seven features: $J-H$, $W1-W2$, M_{W1}, K_s-W1, $\log EWM_{W2}$, $\log \chi^2_{red,W2}$, and $\log EWM_{W1-W2}$. However, the contribution to the balanced accuracy from all but the first four features is small. This suggests that variability amplitude is a useful, though not critical metric to obtain, while variability timescales are not necessary. Finally this suggests that photometry bluer than J-band is also unnecessary.

We can also examine the importance of each feature for classifying individual classes. We perform the same greedy search over the features, instead calculating a performance metric that focuses on the performance on a specific class. One option is the F_β measure:

$$F_\beta = (1+\beta^2)\frac{\text{completeness}\cdot\text{precision}}{\beta^2\cdot\text{completeness}+\text{precision}} \tag{4.7}$$

where β is a free parameter that sets the relative importance of completeness compared to precision.

Common choices are $\beta = 1$ (i.e., F_1, a harmonic mean of completeness and precision), $\beta = 0.5$, and $\beta = 2$ [40]. We adopt F_2 (i.e., $\beta = 2$), because we prioritize generating complete samples of rare massive stars.

The left panels of Figure 4.18 show the F_2 measure as a function of successively added features, calculated specifically for hot stars (top), emission line stars (second panel), cool stars (third panel), and contaminants (bottom). The results for both hot and cool stars are mostly similar to the results for the overall classifier in Figure 4.17, in the sense that the best performance is reached after including a mix of near- and mid-infrared colors and magnitudes. The main difference is that M_G is the fourth most important feature for classifying hot stars.

For emission line stars, a maximum in the mean F_2 is reached after 11 features: $W1 - W2$, M_{W1}, G_J, $\log\langle\Delta t\rangle_{W1}$, $\log\langle\Delta t\rangle_{W2}$, $\log\sigma_{\Delta t,W2}$, $W3 - W4$, $W2 - W3$, $K_s - W1$, $\nu_{0,W2}$, and $J - H$. However, given the errorbars, only the first three features contribute meaningfully, with the remaining features consistent with a constant value of F_2. Interestingly, compared to its contribution to the overall balanced accuracy of the classifier, bluer photometry (signified by the presence of $G - J$ in the above list) is much more important for identifying emission line stars. While variability metrics are included in the above list, they do not significantly contribute to the F_2 score.

For contaminants, a total of 15 features are required in order to maximize F_2: $J - H$, $\log\sigma_{\Delta t,W1-W2}$, $\log\sigma_{\Delta t,W2}$, $\log\langle\Delta t\rangle_{W1}$, $\log\langle\Delta t\rangle_{W1-W2}$, $\nu_{0,W2}$, $\log\sigma_{\Delta t,W1}$, $\nu_{0,W1-W2}$, M_{W1}, $H - K_s$, $K_s - W1$, $\log \text{EWM}_{W2}$, $\log\chi^2_{red,W1-W1}$, $\log \text{EWM}_{W1-W2}$, and $\log\chi^2_{red,W1}$. Notably, the F_2 measure first *decreases* as features are added, before increasing to the maximum after M_{W1}. This trend is unintuitive compared to the other panels in the Figure. It may be a result of the fact that increased features improve the precision of the classifier at the cost of completeness, resulting in a decrease in F_2 due to the increased weighting of completeness.

To demonstrate the capabilities of the classifier using a limited set of features, we plot the scaled and imputed test set — which was not used in the greedy search algorithm — in the right panels of Figure 4.18, using only the two most important features in each row. Stars belonging to the corresponding coarse class are plotted as larger, colored points, with gray points in the background corresponding to stars in test set with different coarse labels.. In all cases, most members of the

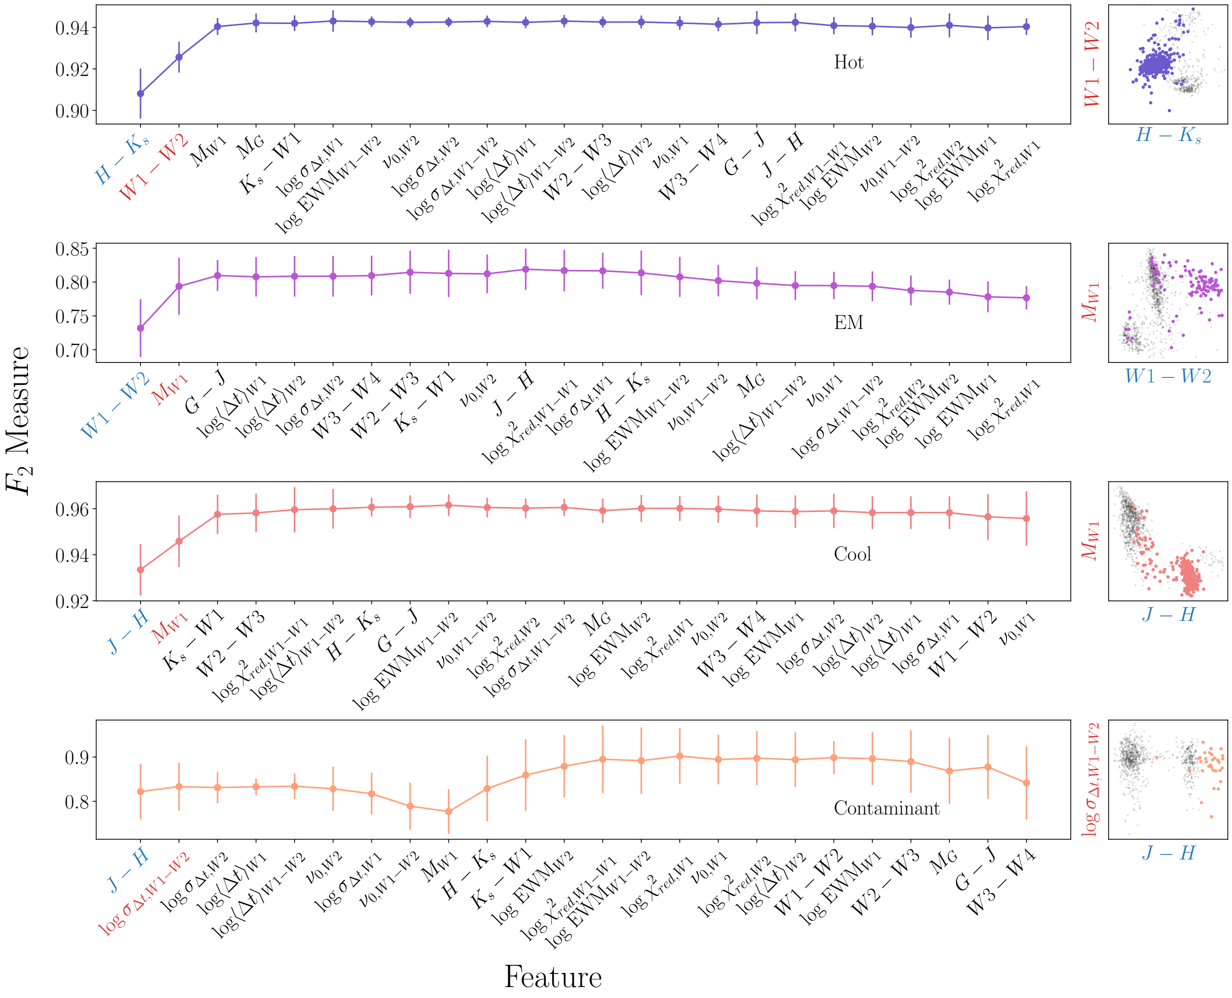

Figure 4.18 Left panels are similar to Figure 4.17, except using the F_2 measure calculated for hot stars (top), emission line stars (second panel), cool stars (third panel), and contaminants (bottom). The right panel in each row shows a scatter plot of only first and second most important features (indicated with blue and red text respectively) drawn from the test set, with stars belonging to the corresponding class in each row highlighted. Note that the features plotted are the scaled and imputed values, *not* the original values listed in Table 4.4. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [69].

test set with that label are well-separated from the other stars. As expected from the left panels, most of the separation is along the x-axis which corresponds to the most important feature, with the second-most important feature plotted on the y-axis providing some additional differentiation, especially for hot stars. In the case of contaminants, the second-most important feature provides little-to-no additional information, consistent with the lack of change in F_2 with increased features.

We conclude that while small numbers of features can be used to classify hot, cool, and (remarkably) emission line stars with high accuracy and precision ($F_2 \gtrsim 0.9$ for hot and cool stars, and $F_2 \gtrsim 0.8$ for emission line stars), a large number of features is necessary in order to maximize the number of accurately identified contaminants. This includes time domain features, where the drastically different structures of old AGB and RGB stars compared to massive cool supergiants may be imprinted. Already extragalactic massive star samples are contaminated by foreground giants in the Milky Way halo; distant stars that can be resolved by *Webb* and *Roman* will have comparable brightnesses to cool dwarfs that are too faint to be filtered out using astrometry from *Gaia* . Developing the infrastructure to reliably remove these contaminating objects from massive star samples will be that much more critical.

4.4 Discussion & Conclusion

In the coming decades, space-based infrared observatories like *Webb* and *Roman* will give us access to unprecedentedly large samples of evolved massive stars. Therefore, we need to be prepared to leverage these data to search for stars in the most interesting evolutionary states. Obtaining spectroscopy of individual stars does not scale well at the size of the expected samples, while linear cuts in color-magnitude space are too simplistic and ignore emission line objects. Here we have demonstrated the promising performance of a support vector machine trained on $\sim 0.5 - 22$ μm photometry and simple variability metrics. However, with currently available labels, we are not able to construct a classifier that performs well at the level of granularity needed for many science cases.

Our main results are summarized as follows:

- We have assembled a large sample of evolved massive stars using distances from *Gaia* DR2 and [16], with high precision infrared photometry from *Gaia* , 2MASS, and WISE.

- Using SIMBAD, we assign labels to all stars, and find that the sample contains a number of low mass contaminants.

- We find that, of the classification methods we applied, a support vector machine classification (SVC) algorithm is best at accurately labeling evolved massive stars. The SVC is fast, and has the added benefit that the underlying mathematics are well-understood.

- The SVC trained on refined labels is capable of identifying low mass red giant contaminants with high accuracy. However, the overall performance of this classifier is quite poor, and we do not recommend its use at present.

- The SVC trained using coarse spectral types performs better, as measured with the balanced accuracy score. We find higher completeness and lower contamination (Figures 4.14 and 4.15) compared to the SVC trained on the refined labels (Figures 4.11 and 4.15). With this classifier, we identify 14 candidate emission line stars from a sample of ~ 2500 unlabeled stars. We plan to obtain spectroscopy of these stars to confirm our results.

- We find that the SVC performs equally as well with only a small subset of features. These features are mostly infrared colors and absolute magnitudes — i.e., those least affected by reddening — with small contributions from infrared variability metrics. However, if we change our performance metric to one that focuses on emission line stars, optimal performance of the classifier requires some red-optical photometry. We find that the added benefit of using variability metrics may not be worth the investment in telescope time in order to measure them. Of course, this is only the case for the sparsely sampled lightcurves in our sample; with the advent of the Legacy Survey of Space and Time (LSST) conducted at the Vera Rubin Observatory, multi-color variability metrics can be estimated from well-sampled optical lightcurves for a significantly larger sample of evolved massive stars, and this claim can be reevaluated.

Ultimately, the performance of the SVC trained on the refined labels is poor. All stars in the sample are bright ($W1 < 14$), and the input features we use are easily measured, implying the classifier is not limited by the quality of the data. However, the labelling itself is not of sufficient accuracy, as can be seen in Figure 4.5. Labels are derived inhomogeneously, and many are from spectroscopy that is now more than 50 years old. Unfortunately, these are the best labels available for this sample. At present, though curated lists of different subclasses of massive stars exist [e.g. 241], no unified catalog of massive stars in our Galaxy or the Magellanic Clouds exists.

Modern all-sky surveys such as SDSS [328] (as well the APOGEE and APOGEE-2 sub-surveys [178]) and LAMOST [53] have already given us access to precision photometric and spectroscopic measurements of unprecedented numbers of stars. Massive stars are bright, and so the existing data is of suitable signal to noise to perform spectroscopic classification. However, they are often excluded from analyses that provide value-added measurements like effective temperatures, surface gravities, compositions, radial velocities, and more that can be used to accurately classify massive stars, with the exception of some incredibly recent preliminary work focused exclusively on main sequence massive stars [e.g. 325] In order to prepare ourselves for the era of *Webb* and *Roman*, we must develop pipelines specifically tuned for evolved massive stars. This is especially true for the classes that are underrepresented in our dataset, i.e., rare emission line stars.

Along with better labels, more data will become available via future data releases of the *Gaia* mission. The recent early third *Gaia* release contains modest improvements in precision and sample size that are unlikely to affect our results given the high quality of the photometry in our sample. However, the full *Gaia* DR3 will contain low-resolution spectra as well as epoch photometry for a limited number of sources, which have the potential to significantly improve the performance of a machine learning classifier. On the horizon, the Legacy Survey of Space and Time (LSST) conducted at the Vera Rubin Observatory will measure the multi-color variability of massive stars at higher cadence, while its large telescope aperture will help define a much larger sample. As we demonstrate with Figure 4.18, it is possible to select features that maximize the performance of the SVC for specific classes. With a larger sample, we may be able to optimize the SVC to search for specific classes of evolved massive stars.

Chapter 5

SHORT TERM VARIABILITY OF EVOLVED MASSIVE STARS WITH *TESS*

5.1 Introduction

The environments in and around evolved massive stars are complex and unique astrophysical laboratories. Much of the information about the physics of these stars is encoded within their variability. However, due to their rarity, the behavior of massive stars in the time domain is still poorly studied by high-precision space-based instruments. Thus, the critical physical ingredients that inform our models of evolved massive stars (e.g., the distribution of rotation rates, asteroseismically determined masses and radii, short-timescale wind-driven variability and more) are still poorly constrained by observations.

Main sequence massive stars are observed as O and B dwarfs earlier than spectral type ∼B3 [108]. Because massive stars spend approximately 90% of their lifetimes on the main sequence [76], unevolved OB dwarfs and some supergiants have been observed extensively using space-based photometry, revealing rotational modulation, wind-driven variability, pulsations, and more [see, e.g., 248, 24, 4, 32, 18, 17, 125, 237, 57, 144].

Beyond these early phases in their evolution, massive stars are poorly understood at short (<1 week) timescales. [46] demonstrated that these stars display a rich variety of variability on day-to-decade-long timescales. However, no massive stars evolved beyond the blue supergiant phase were observed at higher cadence by *Kepler* or K2 before the *Kepler* spacecraft ran out of fuel, and only small samples of evolved stars at specific evolutionary phases have been observed with targeted campaigns using MOST, or the BRITE constellation [e.g. 199, 237, 200]. In other words, there have been no systematic studies of high frequency variability in evolved massive stars, despite the fact that many of the effects observable in main sequence massive stars become strongly amplified in their evolved descendants. In particular, stars in transitional states with lifetimes of only a few 10^4 years (e.g., Yellow Supergiants, luminous blue variables, "slash" stars, etc.) have gone completely unobserved by space-based photometry missions, are quite poorly understood theoretically, and, as we demonstrate here, exhibit multiple modes of variability on ∼day timescales.

The Transiting Exoplanet Survey Satellite (TESS, [242]) is a nearly all-sky photometry mission targeting ∼20,000 bright stars per year at a two-minute cadence (with full-frame images for ∼20 million stars every thirty minutes), yielding approximately 27 days of continuous photometry for stars close to the ecliptic plane, with longer light curves for stars observed by multiple spacecraft pointings. Large numbers of evolved massive stars in the Galaxy and the Magellanic Clouds are bright enough to be observed by TESS over the course of its nominal two-year mission, allowing us to finally explore variability in evolved massive stars at high frequency with fantastic precision. In particular, the Large Magellanic Cloud is located in the TESS Southern Continuous Viewing Zone, providing us with year-long continuous lightcurves; this dataset presents an extremely exciting opportunity to study evolved massive stars in entirely new ways. Here we present analysis of the first evolved massive star light curves to become available from TESS sectors 1 and 2. In §5.2, we discuss our sample selection. Results for each star are presented in §5.3. We discuss the relevance of our findings for stellar evolution theory, and the prospects of a dedicated TESS campaign to observe evolved massive stars in §5.4, before concluding in §5.5. This work originally appeared in Dorn-Wallenstein et al. (2019) [71].

5.2 *Sample Selection and Data Processing*

We first identified candidate evolved massive stars in the Galaxy using the accurate astrometry published in the second data release (DR2) of the *Gaia* [85], following a similar procedure as described in Chapter 4. While many of the details are identical, we briefly describe them here. We acquired the TESS Sectors 1 and 2 target lists available online, uploaded them to the ESA *Gaia* archive[1], and searched for objects in *Gaia* DR2 and the TESS target lists that were separated

---

[1] https://gea.esac.esa.int/archive/

Table 5.1. TESS two-minute cadence targets. TIC # and T magnitude are from the TIC; J magnitude is from 2MASS [55], and 3.6, 4.5, 5.8, and 8.0 μm magnitudes are from the Spitzer SAGE LMC survey [26].

Common Name	Evolutionary Stage	TIC #	T [mag]	J [mag]	$J-[3.6]$ [mag]	$J-[4.5]$ [mag]	$J-[5.8]$ [mag]	$J-[8.0]$ [mag]
S Dor	LBV	179305185	9.16	8.683	0.923	1.036	1.164	1.373
HD 269953	YSG	404850274	9.22	8.588	1.311	1.89	2.363	3.427
HD 269582	Ofpe/WN9 → LBV	279957111	9.33	12.041	2.814	3.188	3.255	3.612
HD 270046	YSG	389437365	9.45	8.713	0.712	0.748	0.933	0.894
HD 270111	YSG	389565293	9.63	9.073	0.535	0.578	0.599	0.480
HD 269331	YSG	179206253	9.82	9.594	0.373	0.457	0.525	0.580
HD 269110	YSG	40404470	10.01	9.320	0.560	0.643	0.695	0.891
HD 268687	YSG	29984014	10.21	9.693	0.397	0.523	0.608	0.706

by less than 1”. We then derived absolute magnitudes using either the Bayesian distance estimates from Bailer-Jones [16] or the available catalogs of *Gaia* targets in the Magellanic Clouds [87], and estimated reddening using the published extinction in the *Gaia* band, A_G and coefficients from [180]. We make the cross-matched data, as well as the estimated M_G and $G_{BP}-G_{RP}$ publicly available online.[2] For stars without an estimated A_G, we also estimate a lower limit to the absolute magnitude

$$M_G \leq G - 5\log_{10} r_{est} + 5 \tag{5.1}$$

and an upper limit to the intrinsic $G_{BP}-G_{RP}$ by assuming $E(G_{BP}-G_{RP}) = 0$. For stars in the Magellanic Clouds, we adopt distance moduli of 19.05/18.52 for the SMC/LMC respectively [143, 142], and assume the values of R_V from [97] and $E(B-V)$ from [185] to calculate the average A_G and $E(G_{BP}-G_{RP})$ towards both MCs.

We then construct dust- and distance-corrected color-magnitude diagrams (CMDs), which we can use to select massive stars from targets observed by TESS using a similar procedure to that described in §4.1, using MIST isochrones. In particular we chose isochrones with metallicity $[Fe/H] = 0.25$ for the Galaxy, and rotation speed relative to critical of $v/v_{crit} = 0.4$ — the $[Fe/H] = 0.25$ isochrones were chosen to follow the general distribution of main sequence stars,

---

[2] https://github.com/tzdwi/TESS-Gaia

which we wish to avoid in our sample [e.g., 58]. For the LMC (SMC) we choose the corresponding $[Fe/H] = -0.5$ (-1) isochrones. We used these isochrones to create a boundary in color-magnitude space that represents the faintest luminosities reached by any massive star at any point during its evolution. We show the boundary for each metallicity isochrone set, as well as *Gaia* colors and absolute magnitudes in Figure 5.1. Points in blue are stars for which our estimate of M_G and $G_{BP} - G_{RP}$ include the extinction, and stars in orange are those without estimates of A_G in *Gaia* DR2. Stars in green (purple) are stars belonging to the LMC (SMC), as identified by [87]. The black/green/purple lines denotes our luminosity cutoff for selecting massive stars in the Galaxy/LMC/SMC.

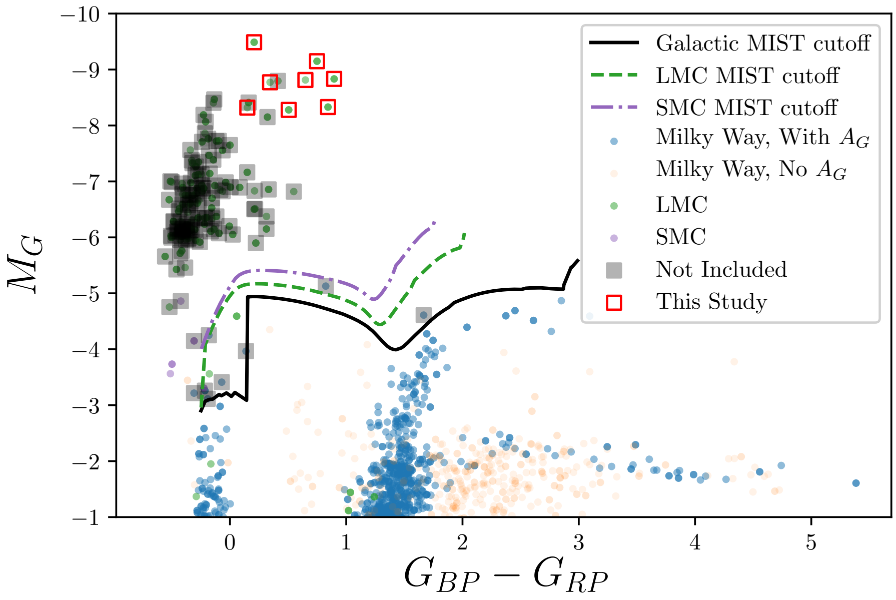

Figure 5.1 *Gaia* CMD for TESS Sector 1 and 2 targets. Galactic stars with an estimate of A_G are in blue, while stars without an A_G estimate are in orange; for these stars, colors are upper limits, and magnitudes are lower limits. Points in green (purple) are in the LMC (SMC) The black, green, and purple lines represent our minimum-luminosity criteria to select massive stars in the Galaxy, LMC, and SMC respectively. Stars in grey boxes are either low mass stars, or relatively unevolved O and B stars. The red boxes indicate the eight evolved massive stars we select for this study. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [71].

Of these stars, many are Galactic long period variables (LPVs), and a number are main sequence or giant OB stars, as well as some Be stars. These hot stars were all observed for specific OB or Be asteroseismology programs [229]. As we are interested in evolved massive stars, we remove these objects. We also remove several targets that are not spectroscopically confirmed as massive stars, located in extremely crowded fields (with multiple bright targets located in a single 21" TESS pixel), or that are members of multiple systems (which will be studied in future work). This leaves us with a total of eight evolved massive stars in Sectors 1 and 2 observed at two-minute cadence with TESS. We list the evolutionary stage, T magnitude and ID number from the TESS Input Catalog [TIC, 280], 2MASS J magnitude [55], and 2MASS/IRAC colors using data from [26] in Table 5.1. We indicate the locations of these stars with red boxes in Figure 5.1. Points outlined in grey boxes are either low mass AGB LPVs or relatively unevolved O and B stars that we ignore for this study. All remaining targets belong to the LMC. While it would potentially be exciting to examine the dependence of any observed variability on metallicity, all eight of these stars will eventually be observed for a year during the TESS prime mission due to their location in the Southern Continuous Viewing Zone.

5.2.1 Data Cleaning

The TESS team released raw light curves and full-frame images from Sectors 1 and 2 on 6 December 2018. We downloaded all light curves for stars observed at two-minute cadence from The Mikulski Archive for Space Telescopes (MAST), and selected the light curves associated with the TIC numbers of our targets. Because these light curves are processed by the first iteration of the TESS pipeline, we err on the side of caution, assuming that the raw light curves contain numerous instrumental effects. Thus we select the `PDCSAP_FLUX` lightcurves, which have been corrected for some systematics. We then normalize by dividing the data by the median flux. For targets observed in both Sectors 1 and 2, we choose to median-divide each Sector individually before concatenating the light curves. While this helps to eliminate Sector-to-Sector offsets and systematics, it can also erase variations at timescales longer than $\sim$1 month, making us less sensitive on those timescales. Finally, HD 268687 displays a gradual brightening at the beginning of Sector 1, which we attribute to telescope systematics. We fit this light curve with a 7^{th}-order polynomial, and normalize the data by the fit in order to remove this increase in flux that would mask otherwise interesting behavior. We plot all of the normalized light curves, along with a rolling 128-point median in orange, in Figure 5.2. The resulting data have pseudo-Nyquist frequencies f_{Ny} (calculated from the average time-difference between data points) between 320 and 330 d^{-1}. Due to the ∼30-day observing window per TESS sector, the expected width of peaks in the periodograms presented (the Rayleigh resolution, defined as the inverse of the observing baseline T) is 0.036 day^{-1} for HD 269582, HD 270111, and HD 269331, and 0.018 day^{-1} for the remaining stars.

5.2.2 *Iterative Prewhitening*

We now wish to describe the variability of each star in terms of sinusoids, in order to identify any significant periodicities in the data. For this purpose, a procedure known as *prewhitening*[3] can be applied, which identifies a significant frequency in the data and subtracts a sinusoid from the lightcurve, proceding iteratively until some noise floor is reached. We adapt the prewhitening method described in Blomme et al. [24] as follows: we first subtract the mean value of the flux. We then use Astropy [14, 292] to calculate the Lomb-Scargle Periodogram [166, 256] on the unsmoothed data for frequencies between 1/30 day^{-1} and f_{Ny}, adopting the default Astropy grid-spacing heuristic of 5 times the Rayleigh resolution. The lower frequency limit is set to avoid overinterpreting low-frequency systematics that may exist between TESS sectors. Using the `psd` normalization, the resulting periodograms are in units of power.

At the j^{th} stage of prewhitening, we calculate the periodogram, and select the frequency of the peak with the highest power, f_{max}. We then use the `curve_fit` routine in SciPy [126] to fit the current prewhitened flux at each time t_i with a sin function $A_j \sin(2\pi f_j t_i + \phi_j)$, where A_j is the semi-amplitude, f_j is the frequency, and ϕ_j is the phase. A_j and ϕ_j are allowed to vary freely, and f_j is bound to the range $f_{max} \pm 1/T$. To calculate the errors on each parameter, we use the

---

[3] This is somewhat of a misnomer, as "prewhitening" implies that we are applying this procedure in order to prepare for some subsequent analysis, when in fact it is the frequencies identified through prewhitening that we are interested in.

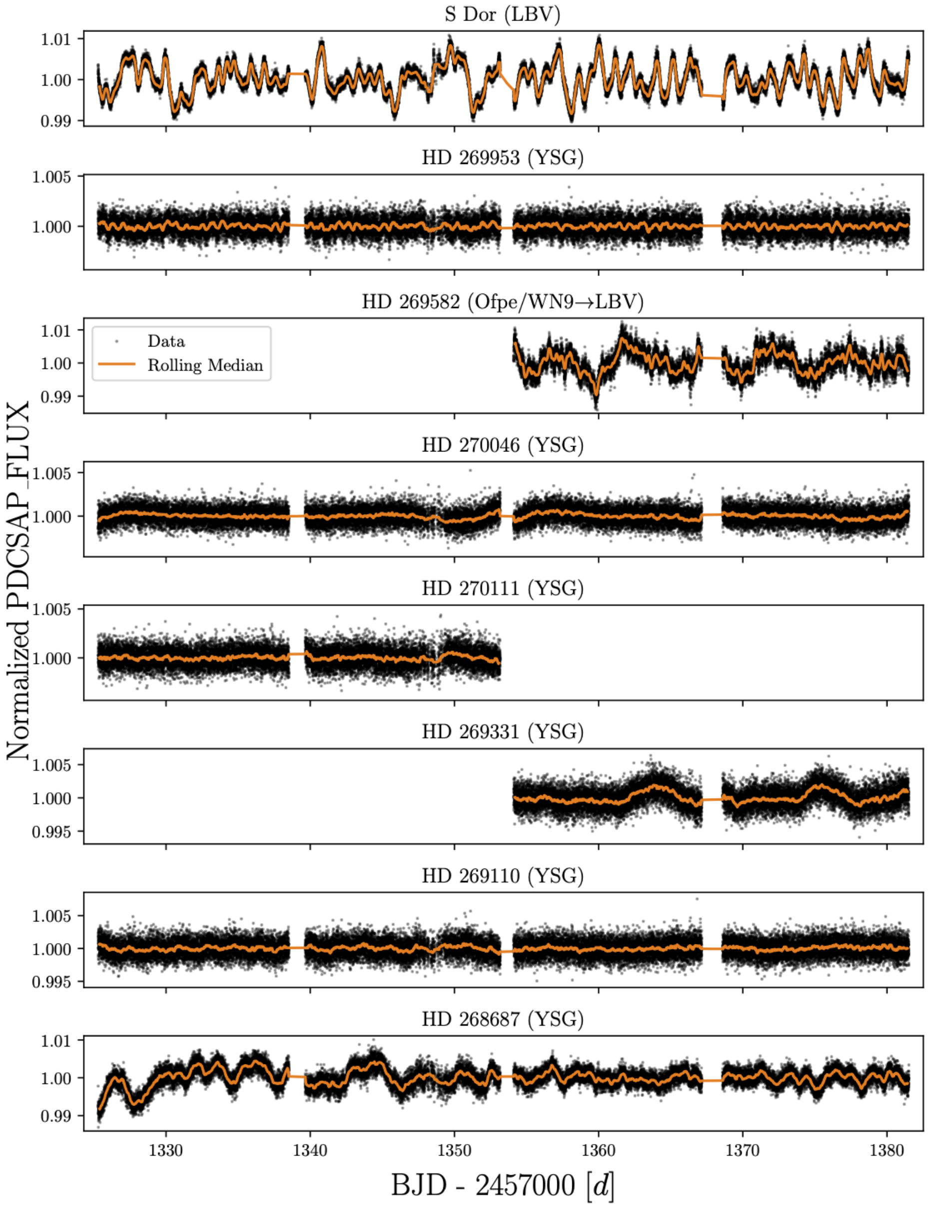

Figure 5.2 Normalized TESS light curves for the eight target stars. Data are in black points, and a rolling 128-point median is plotted in orange. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [71].

formulae given in [169] and [201]:

$$\epsilon(f_j) = \sqrt{\frac{6}{N}}\frac{1}{\pi T}\frac{\sigma_j}{A_j} \tag{5.2}$$

$$\epsilon(A_j) = \sqrt{\frac{2}{N}}\sigma_j \tag{5.3}$$

$$\epsilon(\phi_j) = \sqrt{\frac{2}{N}}\frac{\sigma_j}{A} \tag{5.4}$$

where σ_j is the standard deviation of the flux at the j^{th} prewhitening stage.

We then subtract the fit from the flux and begin the next phase. As a stopping criterion, we adopt the Bayesian Information Content (BIC, [257])

$$BIC = -2\ln(\mathcal{L}) + m\ln(N) \tag{5.6}$$

where m is the total number of terms in the fit ($\equiv 3j$), N is the number of points in the lightcurve, and $\mathcal{L}$ is the likelihood, defined such that

$$-2\ln(\mathcal{L}) = \sum_{i=1}^{N} \frac{(y_i - F(t_i, \Theta_m))^2}{\sigma_i^2} \tag{5.7}$$

to within a constant, where y_i are the original normalized fluxes, $F(t_i, \Theta_m)$ is the sum of all of the fit sinusoids in the current and preceding prewhitening stages evaluated at times t_i and fit parameters Θ_m, and σ_i are the normalized errors in the original light curve [234, Sect. 15.1]. To determine when we have reached the noise level of the light curve, we continue prewhitening until we reach a minimum in the fit's BIC.

This procedure results in a list of frequencies, amplitudes, and phases that can, in principle, describe the variability of each star. However, a number of the derived frequencies are quite similar to each other (i.e., the difference in frequencies is within the Rayleigh resolution). These similar and spurious frequencies can arise due to the short length of the observing baseline [167]. Therefore, we filter the list of frequencies by imposing a requirement that frequencies must be separated by more than $1.5/T$. In cases where such pairs of similar frequencies are found, we discard the frequency found at the later stage of prewhitening. We list N, f_{Ny}, the Rayleigh resolution $1/T$, the number of frequencies found via prewhitening, and the number of unique frequencies in Table 5.2. The unique frequencies, amplitudes, and phases (as well as corresponding formal errors) found for each star are listed in §A.2.

We note that, while this prewhitening procedure is capable of accurately describing the TESS light curves, significant frequency-dependent astrophysical and instrumental noise exists in these data, and produces spurious coherent structures that may be mistaken for periodicity. Determining the significance of a detected periodic signal under the null hypothesis of such noise is a non-trivial task. While analytic methods exist to estimate significance in the case of white or power-law noise [e.g., 20, 312], significance tests in the case of a more complex noise model are inconsistent in various subfields in the literature. No suitable physical model for the astrophysical noise, represented in Equation (5.12), has been proposed; therefore, any estimate of the significance of the frequencies found here would be model dependent. However, we do test to determine if the value of the periodogram peaks closest to the recovered frequencies exceed a power corresponding to a false-alarm level of 1% under the null hypothesis of white noise, as described by [20], and note the cases where frequencies discussed do not, both in text and in §A.2.

In the case of the stochastic variability — a.k.a. red noise, manifested as a low frequency power excess in the periodogram — discussed below, [24] present a method of evaluating the significance of frequencies under noise of the form discussed here. After fitting the periodogram with Equation (5.12), we rescale the noise function so that the integral over the frequency range considered is equal to the variance of the flux times f_{Ny}. For each frequency, we evaluate:

$$P(z > Z) = 1 - e^{-e^{-0.93Z+\ln{(0.8N)}}} \tag{5.8}$$

which gives the probability that a stochastic noise process would result in an amplitude z that exceeds a threshold

$$Z(f) = \frac{A_j^2 N}{4\sigma_{red}^2} \tag{5.9}$$

where A_j is the amplitude found by prewhitening, N is the number of points in the lightcurve, and σ_{red} is the value of the rescaled noise function at the frequency under consideration. A frequency is considered significant under red noise if $P < 0.01$. Finally, we estimate the signal-to-noise ratio (SNR) for each frequency, calculated as the ratio of the amplitude to the standard deviation of the lightcurve after prewhitening.

Suspiciously, all frequencies found are significant under this criterion, while no frequency has $SNR > 2$, despite all frequencies included quantitatively improving the overall fit to the lightcurve

Table 5.2. Summary of light curve characteristics.

Common Name	N	f_{Ny} [day^{-1}]	$1/T$ [day^{-1}]	Number of frequencies	Unique frequencies
S Dor	36296	322.8060468379108	0.0177879072510214	71	41
HD 269953	36418	323.89171072083246	0.01778794028727421	14	14
HD 269582	18226	332.52345927965945	0.036490914598590884	33	25
HD 270046	36291	322.76203971547	0.017787932748165886	14	10
HD 270111	18101	324.59880106819276	0.03586727083626439	7	6
HD 269331	18279	333.49073111685124	0.03649094333262405	10	6
HD 269110	36403	323.7577400015701	0.017787909455610686	11	7
HD 268687	36412	323.83746567083887	0.017787891882718898	64	33

according to the adopted BIC criterion (Equation (5.6)). As we mention above, this result is likely model dependent; it is possible that many of these frequencies (if not all in some cases) are entirely attributable to the noise process. Alternatively, the red noise is much higher amplitude than the instrumental noise at low frequencies, and is not removed by our prewhitening procedure; as a result, the SNR we are measuring is mostly reflective of the ratio in power between two astrophysical processes (the identified coherent frequency and the stochastic red noise background), rather than our confidence in the detection of each signal. Finally, we do not have a generative model that would allow us to fit and remove the red noise from the data. Ultimately, future TESS sectors will provide insight into the low frequency regime, and allow us to conduct, for example, time-frequency analyses over a much longer baseline to determine the persistence of these frequencies, as we show in Chapter 6.

5.3 Results

5.3.1 Yellow Supergiants

A 25 $M_{\odot}$ solar-metallicity will begin its life as an O star; which, in the Geneva evolutionary tracks [76] appears as an O6 dwarf on the zero-age main sequence. After 7 Myr, it evolves into a B0 supergiant, at which point it crosses the HR diagram in under a Myr to become a RSG, at which point it experiences strong mass loss via dust-driven winds. Due to this mass loss, Approximately

500 kyr after this, it evolves bluewards once more due to having lost a significant amount of its envelope through stellar winds, after which it becomes a Wolf-Rayet star [190]. During both rightward and leftward crossings of the HR diagram, the star undergoes an incredibly brief ($\sim 10^4$ yr) yellow supergiant (YSG) phase.[4] Thus, while the luminosities and effective temperatures of two given YSGs may be identical, their ages, masses, surface abundances, and interior structures may be radically different depending on whether they are underoing a blueward or redward evolution in the HR diagram. Signatures of these differences may be imprinted in the TESS lightcurves. While it is likely that a cool YSG that has previously undergone a RSG phase will be accompanied by a dusty envelope, the dust may be photodissociated as the star's effective temperature increases, which is consistent with the decreasing abundance of circumstellar dust species around increasingly hot evolved massive stars in Table 1 of [321]. Therefore, it is possible that variability may be the best or most unambiguous means of distinguishing between rightward- and leftward-moving YSGs. Indeed, this logic has been applied to *blue* supergiants, where variability has been used to identify candidate post-red supergiant objects [249]. Finding leftward-moving YSGs places a valuable upper limit on the initial masses of stars that explode as RSGs before they can become YSGs (a.k.a., "the red supergiant problem", e.g., [265]).

HD 269953

HD 269953 is a G0 YSG, assigned a luminosity class of 0 by [133], which is in agreement with its high luminosity $\log(L/L_{\odot}) = 5.437$ from [215]. Coupled with its temperature $T_{eff} = 4920$ K, this implies a radius of 566 $R_{\odot}$ from the Stephan-Boltzmann law. While the light curve presented in Figure 5.2 initially appears to be dominated by noise, the light curve smoothed by a 128-point rolling median appears to show coherent oscillations. We re-plot the smoothed, mean-subtracted light curve in the top panel of Figure 5.3. The periodogram, plotted in the lower panel of Figure 5.3 shows multiple strong peaks at frequencies above 1 day^{-1}. The prewhitening procedure described

---

[4] Note that in Chapters 2, 3, and 4, we restricted our definition of YSGs to F- and G-type supergiants — i.e., stars with $T_{eff} < 8000$ K. Here, and in Chapter 6, we expand this definition to include A supergiants, as many of these stars have left the main sequence and are in the midst of crossing the HR diagram.

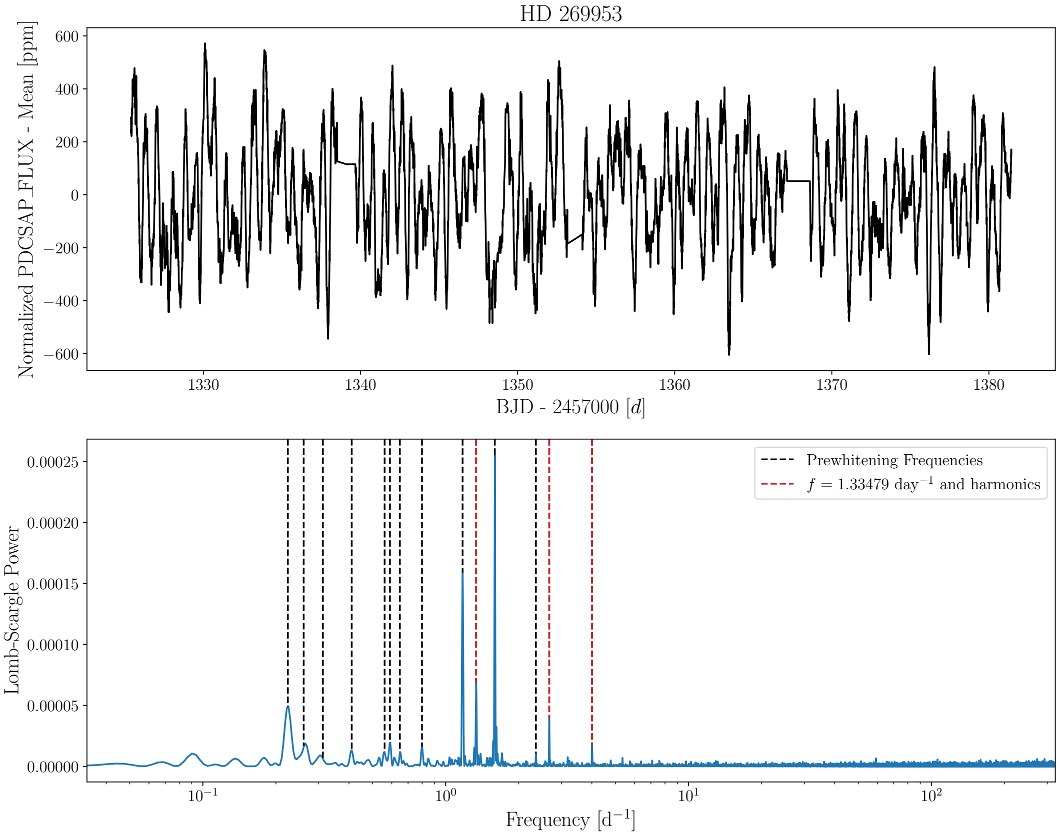

Figure 5.3 *Top*: Light curve for HD 269953, after smoothing with a 128-point rolling median, showing coherent variability at approximately 600 ppm. *Bottom*: Lomb-Scargle Periodogram. The black lines indicate frequencies found via prewhitening. The $f = 1.33479$ day^{-1} peak and it's two detected harmonics are indicated with red vertical lines. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [71].

in §5.2.2 reveals the presence of 14 unique frequencies, indicated by the vertical black lines. We search for harmonics of the form f_1/f_0 that satisfy

$$nf_0 - f_1 \leq \sqrt{\big(n\epsilon(f_0)\big)^2 + \big(\epsilon(f_1)\big)^2} \tag{5.10}$$

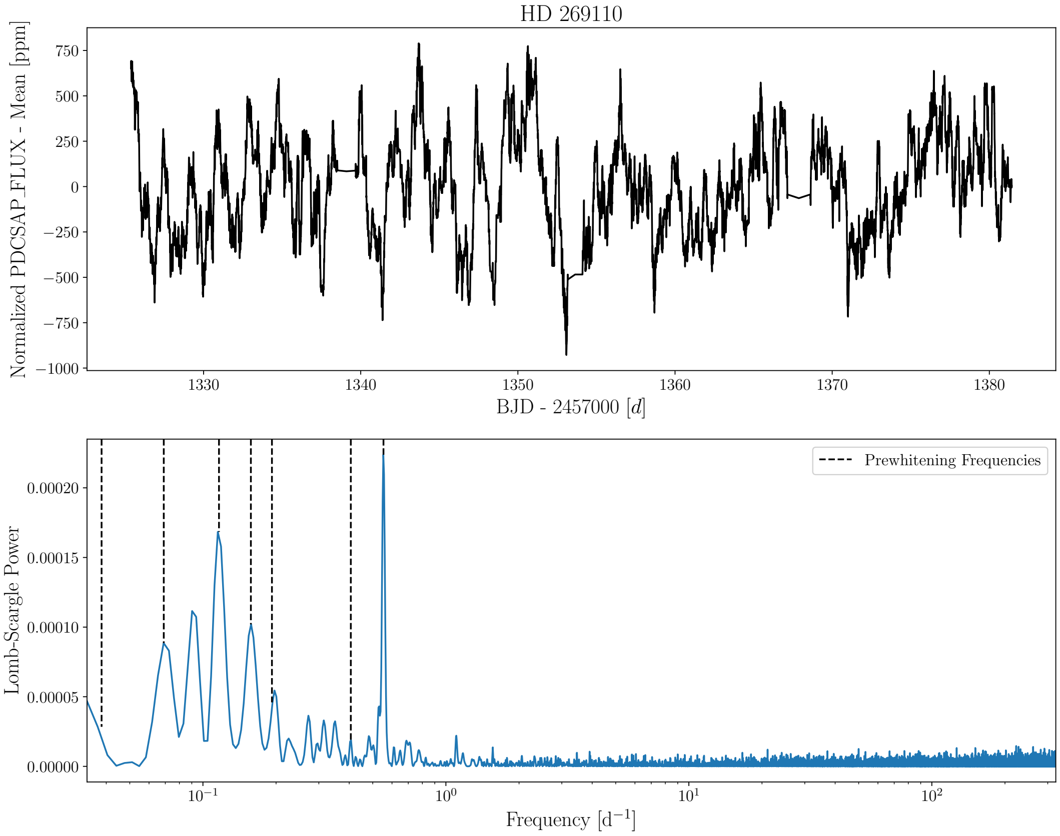

Figure 5.4 Similar to Figure 5.3, for HD 269110. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [71].

where n is an integer greater than 1. In the light curve of HD 269593, we detect the first ($n = 2$) and second ($n = 3$) harmonics of $f = 1.3347924$ day^{-1}, which we indicate with the vertical red lines. While we search for harmonics up to $n = 10$ for all sources, the low amplitudes and frequency resolution make it unlikely that any high harmonics found are real, and we caution against over-interpreting high harmonics that may not exist. For example, this search also revealed the ninth harmonic of $f = 0.26019548$ day^{-1}, but the corresponding peak in the periodogram does not surpass the 1% false alarm level. Given the additional null detection of any of the preceding

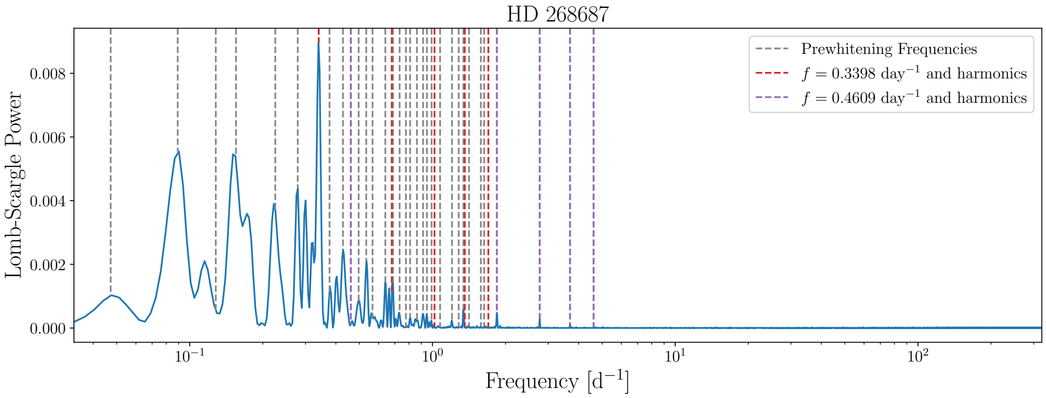

Figure 5.5 Similar to the bottom panel of Figure 5.3 for HD 268687. We indicate the strongest peak and its first four harmonics with vertical red lines; these harmonics were not detected via the procedure described in §5.2.2. We also find harmonics of a ~ 0.4609 day^{-1} signal, indicated in purple. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [71].

harmonics, we deem this to be a chance coincidence.

We also searched for frequency combinations in the form $f_0 + f_1 = f_2$, such that

$$f_0 + f_1 - f_2 \leq \sqrt{\left(\epsilon(f_0)\right)^2 + \left(\epsilon(f_1)\right)^2 + \left(\epsilon(f_2)\right)^2} \tag{5.11}$$

We find that the strongest peak in the periodogram at 1.59360677 day^{-1} can be represented as the sum of two peaks at 0.26019548 and 1.3347924 day^{-1}. If variability at these frequencies is driven by pulsations, this may indicate that some of the modes are interacting with each other.

With high amounts of IR reddening, HD 269953 is the dustiest YSG in our sample, making it quite likely that it is in a post-RSG phase, and thus evolving blueward. Therefore, a direct measurement of its mass via asteroseismology, using the multiple sets of frequencies presented here would be an incredibly valuable constraint on stellar evolution. Unfortunately, it is located in the star-forming LMC cluster NGC 2085, and thus is subject to a high degree of crowding in TESS's 21" pixels. While HD 269953 is the brightest star by far in the field, we reserve further analysis of the light curve until more advanced tools are developed to extract light curves from crowded regions in TESS.

HD 269110 & HD 268687

HD 269110 is a lower luminosity YSG with $\log(L/L_{\odot}) = 5.251$, $T_{eff} = 5624$ K [215], and thus a radius of 445 $R_{\odot}$. It has a spectral type of G0I from [9]. Similar to HD 269953, the light curve initially appears to be just noise, while the light curve smoothed by a 128-point rolling median shows coherent variability at approximately 750 ppm. Figure 5.4 shows the smoothed, mean-subtracted light curve in the top panel, and the periodogram in the bottom. Prewhitening reveals a set of 7 unique frequencies in the lightcurve. On average, these frequencies indicated slower pulsations than those found in the other YSGs. With physical properties and photometry consistent with all other YSGs other than HD 269953, it is not clear why might be the case; it is possible that, as the second-faintest star in the sample, we simply don't detect higher frequency signals. The first harmonic of the strongest peak at 0.55 day^{-1} (a period of 1.81 days) can be seen by eye in the periodogram; however, it was not recovered by our prewhitening procedure. [24] found that the BIC was the most conservative stopping criterion in their prewhitening procedure, resulting in the fewest detected frequencies (which may also explain why we do not recover one of the low frequency peaks seen in the periodogram). Alternatively, this harmonic may not be real. Data from future TESS sectors will allow us to more firmly establish or disprove the existence of this harmonic. A group of peaks centered at $\nu_{max} = 0.115$ day^{-1} with an average frequency spacing $\Delta\nu = 0.032$ day^{-1} is also visible. Similar structures have been found in the periodograms of many less massive pulsating stars, for which asteroseismic models have yielded precise measurements of their masses and deep insight into their core structures. Comparable measurements of evolved massive stars would provide a previously inaccessible constraint on a poorly understood phase of massive stellar evolution. However, suitable models of YSG pulsations at these timescales do not yet exist.

HD 268687 is classified as a F6Ia supergiant by [9], and has a luminosity $\log(L/L_{\odot}) = 5.169$ and effective temperature $T_{eff} = 6081$ K from [215], implying a radius of 346 $R_{\odot}$. The periodogram, shown in Figure 5.5, with prewhitening frequencies shown in grey, displays a clear peak at 0.3398 day^{-1}, corresponding to a period of 2.95 days. We indicate this frequency and its first four harmonics with vertical red lines; none of these exact harmonics are found by prewhitening, however, many of the frequencies recovered are close to the harmonics (though the latter three do not surpass the 1% false alarm level). Additionally, the series of peaks at ~ 0.3 day^{-1} appears to repeat with smaller amplitudes at a spacing of ~ 1.2 day^{-1} until ~ 0.8 day^{-1}, though the exact shape of the peaks changes. The autocorrelation function calculated from the periodogram has a series of peaks with Δf between 0.06 and 0.25 day^{-1}. A total of 33 unique frequencies are revealed by prewhitening, making its lightcurve the most complex of the YSGs studied. We do recover the third ($n = 4$), fifth, seventh, and ninth harmonics of $f = 0.46094861$ day^{-1} (the last of which does not surpass the 1% false alarm level), which we indicate with vertical purple lines. While the fundamental corresponds to a minimum in the periodogram, the harmonics correspond to the only peaks above ~ 2 day^{-1}. Our search for combinations of frequencies yields two frequencies (0.86314365 and 1.07507415 day^{-1}) that mix with the dominant 0.3398 day^{-1} signal, along with 8 other combinations. Finally, a broad bump of peaks accompany the dominant peak at lower frequencies.

All told, the similarities in their periodograms suggest both YSGs are in similar physical sates, and both show clear peaks in their periodograms on timescales of 2-3 days, in addition to structure at higher frequencies, and a series of peaks at lower frequencies. We can rule out some possible sources of the dominant signal in both light curves. If both stars are approximately 25 M$_\odot$ and the variability is due to binary interactions with a companion, the companion would have to be approximately 64,000 M$_\odot$ to be in a 2.95-day Keplerian orbit outside of the stellar surface of HD 268687, and 360,000 M$_\odot$ to be in a 1.81-day Keplerian orbit around HD 269110. We determine both scenarios to be highly implausible.

Rather than interpreting the frequency structures as being due to asteroseismically interesting stellar pulsations, it is possible that the brightness modulations are instead due to one or more spots on the surface of the stars, causing the apparent luminosity to change as the star rotates. If the typical spot latitude were at the stellar equator, then the star would be rotating at approximately 6,000 km s^{-1} for HD 268687, and 12,000 km s^{-1} for HD 269110, well beyond the critical velocity for both stars. However, we cannot rule out a nearly-polar spot, which may explain the change in the shape of the variability in HD 268687 with time, but would require invoking severe surface differential rotation, as well as extremely fast spot decay times to explain the change of the variability in HD 269110.

We therefore favor the interpretation that we are observing coherent pulsational variability in both stars, which is also consistent with the change in the shape of the variability, in addition to the apparent "frequency comb" seen in the periodogram of HD 269110. YSGs have been observed to vary with periods of many tens of days [10] caused by He ionization-driven radial pulsations. Perhaps we are observing a very high-order harmonic of a radial mode. Alternately, oscillations in a non-radial mode may be causing this variability. All of our targets are in the LMC and will be observed almost continuously for a year. Given that the Rayleigh resolution scales as the inverse of the observing baseline, if the variability we've observed is caused by p- or g-mode pulsations, some of the peaks in the periodogram may resolve into additional frequency combs characteristic of these pulsations after a full year of TESS data. Another option is that the variability is caused by Rossby waves (or r-mode oscillations, [221]), which appear as "hump and spike" shapes in the periodogram [247], which have been observed in main sequence F and G stars. While the fundamental mode is located at a slightly lower frequency than the rotational frequency (and hence we run into the same problems as above), higher-azimuthal-order frequencies can arise. Unfortunately the amplitude of the oscillations declines sharply at the higher orders, implying that the rotation speeds would only be a factor of a few slower, which is still physically implausible.

HD 270046, HD 269331, & HD 270111

Two of the remaining three YSGs, HD 270046 and HD 270111, have been poorly studied thus far, and display minimal variability in existing data before TESS. The 5th data release of the RAdial Velocity Experiment (RAVE, [146]) contains spectra for HD 270046. While the multiple optical spectra obtained by RAVE yield a wide range of atmospheric parameters (and distances inconsistent with HD 270046's membership in the LMC), measurements of T_{eff} inferred from infrared flux are all between 6200 and 6360 K, with a mean of 6275 K. We were unable to find atmospheric parameters for HD 270111. However, with spectral types of F8Ia and G5I, and TESS

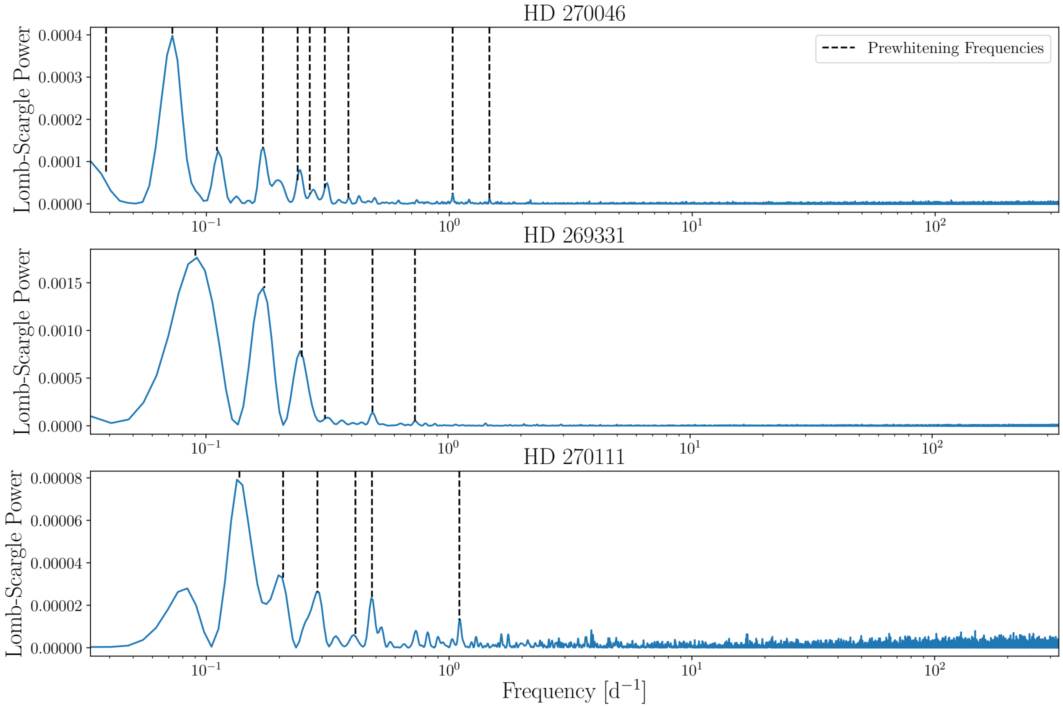

Figure 5.6 Similar to Figure 5.5 for HD 270046 and HD 270111. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [71].

magnitudes/infrared colors consistent with HD 269110 and HD 268687, it is fair to assume both stars are fairly typical yellow supergiants. The final YSG, HD 269331, has atmospheric parameters from Neugent et al. (2012) [215]. While its temperature is consistent with the other YSGs in their sample, Ardeberg (1972) [9] assign a spectral type of A5Ia to HD 269331. Its lightcurve (Figure 5.2) displays two prominent bumps. However, HD 269331 was not observed by TESS in Sector 1, so we are unable to see if these bumps are repeating patterns. We present the periodograms of all three stars in Figure 5.6.

Prewhitening reveals the presence of 10, 6, and 6 unique frequencies in the lightcurves of HD 270046, HD 269331, and HD 270111 respectively, none of which have a normalized semi-amplitude larger than 0.0002, and no peaks surpass the 1% false alarm level. Some higher order

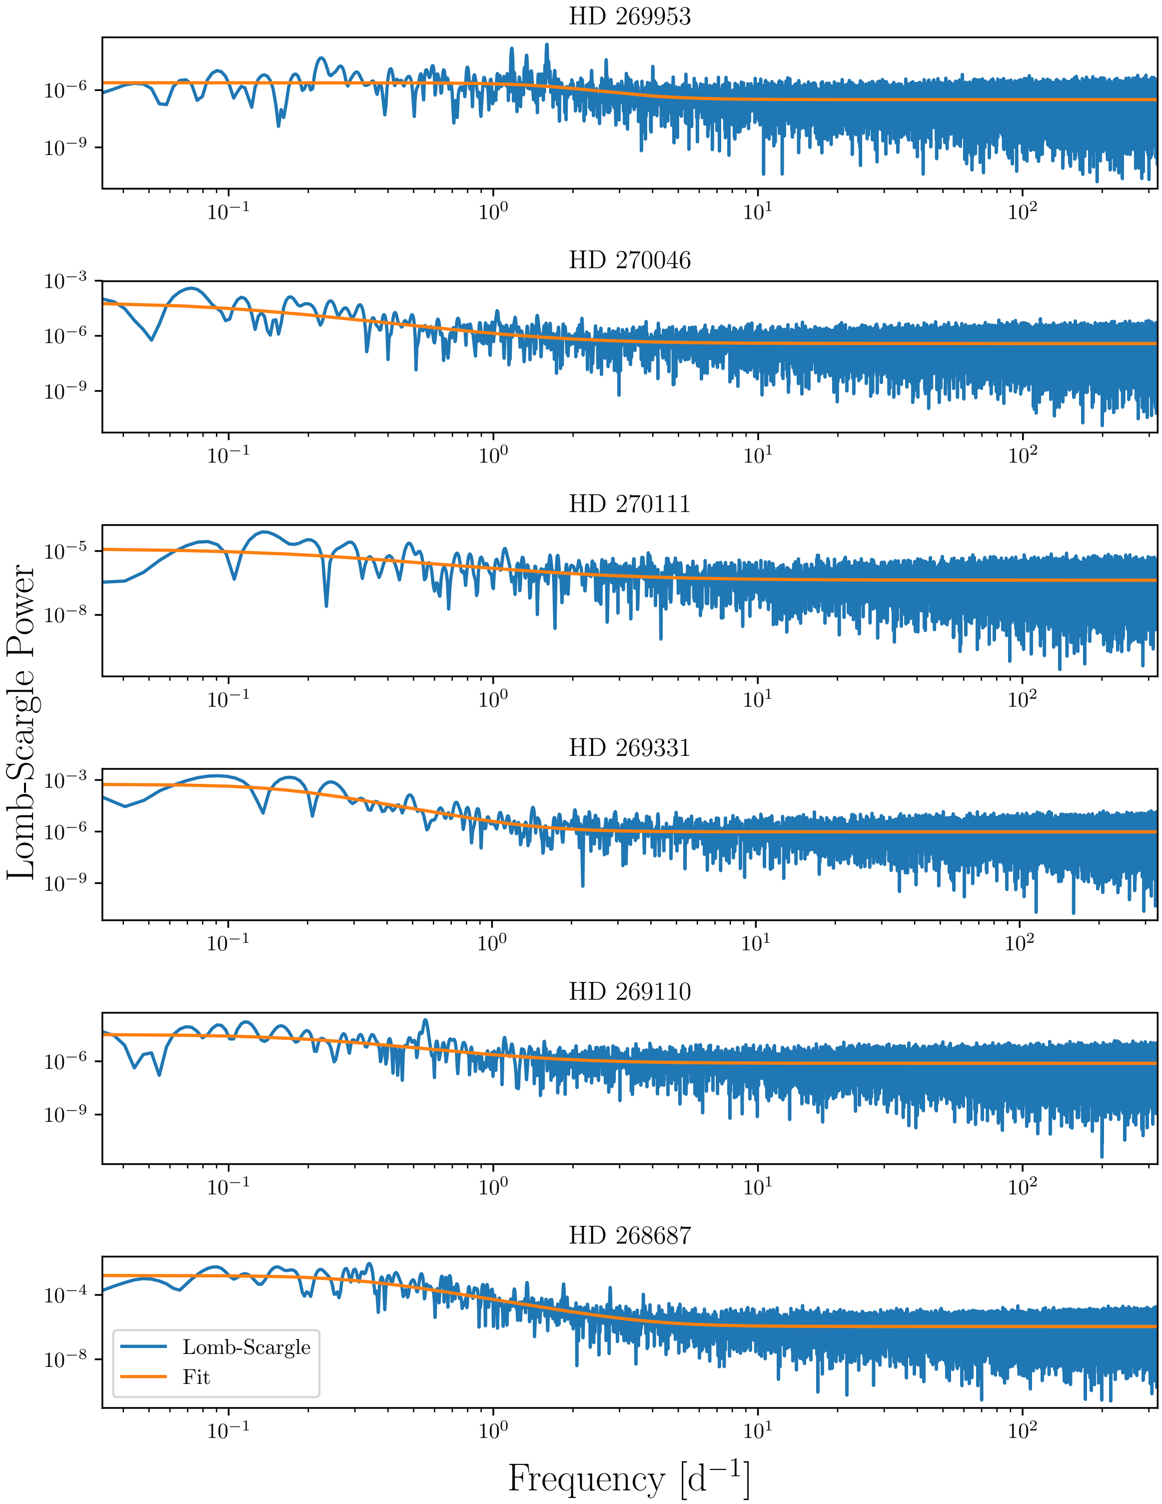

Figure 5.7 Lomb-Scargle periodograms calculated between 1/30 day^{-1} and the pseudo-Nyquist frequency for all six YSGs. Fits using equation (5.12) are in orange. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [71].

($n \geq 7$) harmonics of various frequencies are found in HD 270046, which we dismiss as coincidental. We also see a combination of the highest peak at $f = 0.07243$ day^{-1} and the lowest frequency detected at 0.03865 day^{-1}, seen at 0.11055 day^{-1}. However, it appears as if the low frequency detected by prewhitening corresponds to a peak in the periodogram that actually lies below our minimum frequency cutoff. In HD 270111, we detect the second and seventh harmonics of the strongest peak at $f = 0.136908$ day^{-1}; the second harmonic is also the first harmonic of the small peak at 0.207207 day^{-1}. With such small amplitudes relative to the noise in the light curve, we refrain from discussing these patterns until higher signal to noise periodograms are obtainable with future TESS sectors.

Noise Properties of YSG Light Curves

In addition to the peaks in periodograms of the three YSGs discussed above, background noise exists for all six YSGs. Is this noise instrumental or astrophysical? Astrophysical red noise is seemingly ubiquitous in the light curves of hot massive stars as discussed in §5.1, and thus it would be unsurprising for it to manifest in these cooler stars. When plotted in log-scale, some of the periodograms in Figure 5.7 appear to display red noise (especially HD 268687). To model the background, we follow [24], and use `curve_fit` to fit the Lomb-Scargle periodogram with the function

$$\alpha(f) = \frac{\alpha_0}{1 + (2\pi\tau f)^\gamma} + \alpha_w \tag{5.12}$$

from [275], where f is the frequency, α_0 is the power as $f \to 0$, τ is a characteristic timescale, and α_w is an additional parameter we add in to model the white noise floor at the highest frequencies (ostensibly equal to the instrumental noise). We perform this fit after calculating the base-10 logarithm of both the Lomb-Scargle power and the fitting function to avoid artificial weighting of real peaks at high frequencies.

The periodograms and fits for all six YSGs are shown in Figure 5.7. The parameter values and 1σ error estimates are compiled in Table 5.3, and compared to the physical properties of the stars when available in [215]. It is immediately clear that the noise characteristics of all of the light curves differ, indicating that the source of the noise is likely astrophysical. HD 269953 is quantitatively different from the other YSGs in all parameters but the white noise component of the fit. Notably the red noise power-law component of the fit is only readily apparent over a narrow range of frequencies $\sim 2 - 4$ day^{-1}, but the power law slope is approximately twice as steep as all YSGs but HD 268687. Combined with its status as the dustiest YSG in the sample, it is clear that this object warrants further follow-up in the short-timescale regime. Finally, both of the F supergiants have significantly higher power at the lowest frequencies (α_0) especially HD 268687. With a larger sample of YSGs, comparisons between physical quantities and noise parameters will help constrain the origin of this noise, which has not been detected until now.

In summary, we find a wide range of behavior in the YSGs in our sample, from coherent variability that we attribute to stellar pulsations, to astrophysical red noise. This variability is found on timescales from hours to days, far faster than any previously-observed YSG variability. Comparing the properties of the YSGs, it is readily apparent that the periodogram of HD 269953 is different from the periodograms of the other YSGs discussed above. The strongest peaks in its periodogram are located at almost an order of magnitude higher frequency, and it displays very little of the low frequency peaks seen in the periodograms of the other YSGs. Additionally, the slope of the red noise is much flatter. Coupled with its apparently more-evolved state, this suggests a distinct difference in variability between pre- and post-RSG yellow supergiants.

5.3.2 Luminous Blue Variables

Luminous blue variables (LBVs) are a phenomenological class consisting of extremely luminous stars that show signs of dramatic variability. They are arguably one of the least understood stellar evolutionary phases, and are perhaps best characterized by their giant eruptions (such as those famously associated with η Carina and P Cygni), bright enough to be mistaken as supernovae. In some cases these ”impostor” events are followed by true supernovae on timescales of a few years, as in the case of SN 2009ip, which underwent two outbursts in 2009 and 2010 before undergoing what appears to be a terminal explosion in 2012, e.g. [191, 81]. However, LBVs also experience large episodic variations in their effective temperatures on timescales of months to years, known

Table 5.3. Summary of the fit results to the periodograms of the six YSGs in our sample, along with their physical properties when available from [215]. The T_{eff} for HD 270046 is the mean value inferred from the infrared flux in [146]. .

Common Name	Lit. Spectral Type	$\log(L/L_\odot)$	T_{eff}/K	$\alpha_0/10^{-4}$	$\tau/10^{-2}\mathrm{d}$	γ	$\alpha_w/10^{-5}$
HD 269953	G0 0 [133]	5.437	4920	0.02 ± 0.002	8.79 ± 0.75	3.06 ± 0.28	0.03 ± 0.0001
HD 270046	F8Ia [9]	—	6275	0.66 ± 0.306	172.14 ± 63.31	1.76 ± 0.09	0.04 ± 0.0002
HD 270111	G5I [255]	—	—	0.13 ± 0.056	84.90 ± 42.58	1.40 ± 0.13	0.04 ± 0.0003
HD 269331	A5Ia [9]	5.307	6457	5.56 ± 2.208	100.80 ± 22.10	2.83 ± 0.19	0.10 ± 0.0006
HD 269110	G0I [9]	5.251	5624	0.32 ± 0.093	77.23 ± 20.29	1.87 ± 0.13	0.08 ± 0.0003
HD 268687	F6Ia [9]	5.169	6081	16.03 ± 2.940	54.06 ± 5.25	2.79 ± 0.07	0.11 ± 0.0005

as ”S Dor variations”. With their bolometric luminosities remaining almost constant, these S Dor variations manifest as horizontal evolution on the Hertzsprung-Russell diagram between their hot and cool states that differ by 10-20 kK. In addition to these large scale variations, LBVs also exhibit ∼0.1 mag irregular microvariability on timescales of weeks to months [2].

The evolutionary state of LBVs, their status as single or binary stars, and the physical mechanisms driving the S Dor variations are all topics of current debate (see [268, 118, 1], Levesque & Lamers 2019). One possibility is that pulsations may be important for driving mass loss for S Doradus variability [168], and may therefore be observable. Indeed, a simple estimate of the dynamical/free-fall timescale for a typical LBV from [2] yields

$$t_{dyn} \approx 0.6 \left(\frac{R_*}{10^{12}\mathrm{cm}}\right)^{3/2} \left(\frac{M_*}{100\mathrm{M}_\odot}\right)^{-1/2} \mathrm{days}, \tag{5.13}$$

and variability on this timescale should be easily observable by TESS. However, LBVs tend to be surrounded by a complex and sometimes dusty CSM that can attenuate and modulate these pulsations. Indeed, both LBVs studied here have incredibly red colors in Table 5.1, indicating that any intrinsic variability detected by TESS has been processed by circumstellar material. All told, understanding the short timescale variability of LBVs can offer incredibly valuable insight into the physical state of LBVs and their immediate environments.

HD 269582

HD (sometimes HDE) 269582 was observed as a H-rich Ofpe/WN9 or WN10h Wolf-Rayet star as recently as the mid 1990s [52]. However, since 2003, it has entered an outbursting LBV state, rapidly brightening in V-band as it cooled to a late-B/early-A spectral type, accompanied by drastic changes in various line profiles [317]. Because HD 269582 appears to be newly entering the LBV phase, studying its variability can be quite instructive. Indeed, a link between light curve structure and outbursts has been proposed for Be stars [115, 147]; such a link for LBVs may even be testable with an entire year of observations.

The light curve for HD 269582 presented in the third panel of Figure 5.2 shows coherent $\sim$1%-level variability on timescales of a few days. The periodogram shown in Figure 5.8 shows a strong peak, detected with prewhitening at 0.20327588 days^{-1} (corresponding to a 4.919-day period) with small peaks to either side. Though TESS only observed HD 269582 for 5 full cycles of this measured period, the shape of the light curve from cycle to cycle changes noticeably. This can be seen in the dynamic plot in Figure 5.9, showing the flux as a function of phase from cycle to cycle. The phase of maximum luminosity appears to shift from cycle to cycle, while the amplitude of modulation decreases. Prewhitening reveals the presence of a total of 25 unique frequencies, most of which are small amplitude peaks above ~ 1 day^{-1}. Among those frequencies, we find no convincing harmonics. Interestingly, two frequencies, $f = 1.6023258$ and 3.1987555 day^{-1}, are found twice each in our search for sums of frequencies.

Similar dominant periods and changes in the light curve shape were observed in WR 110 by Chenne et al. (2011) [38]. The 30-day light curve presented there appears remarkably similar to the TESS light curve of HD 269582. [38] attributed the behavior of WR 110 to a CIR in the wind, implying that we are measuring the rotational frequency. It is also possible that this frequency and the surrounding peaks in the periodogram, or the higher frequencies found by prewhitening are nonradial pulsations. Longer monitoring by TESS will enable us to resolve these peaks further, and build a more physical model with well-sampled parameter distributions, and spectroscopic monitoring would allow us to confirm a CIR scenario.

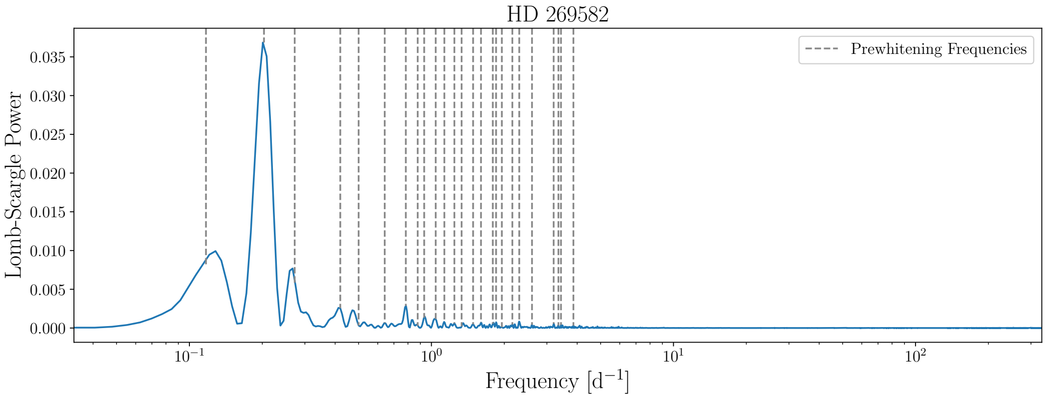

Figure 5.8 Periodogram for HD 269582, showing a clear peak corresponding to a period of 4.919 days. All frequencies found via prewhitening are indicated with grey vertical lines. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [71].

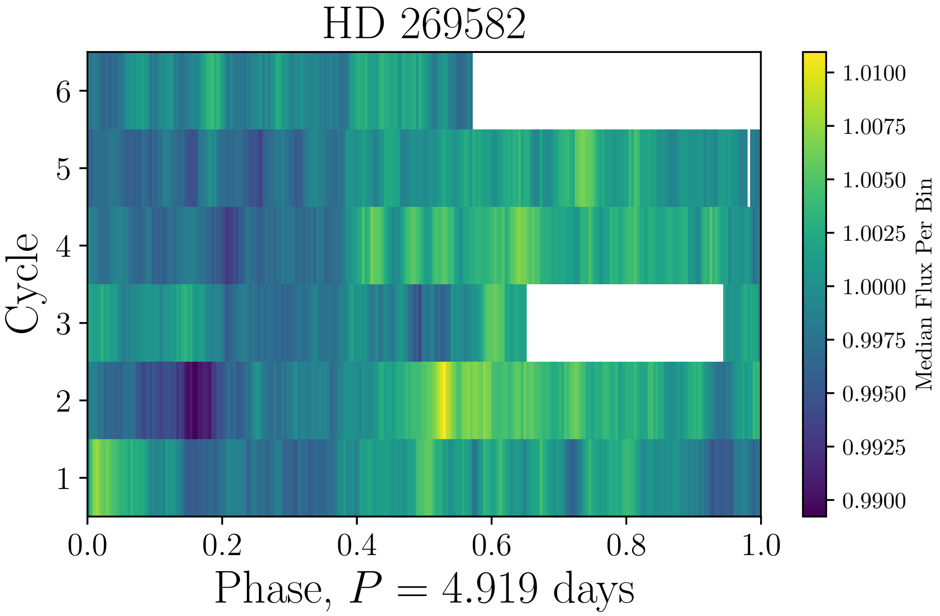

Figure 5.9 Dynamic plot, phased to a 4.919-day period for HD 269582, showing the variability from cycle to cycle. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [71].

S Doradus

S Doradus is the prototypical S Dor variable, with a long history of photometric and spectroscopic observations. [307] detected a ∼7 year period in S Dor's light curve, which [2] argued is more likely to be a timescale associated with the duration of individual flaring events.

The light curve presented in Figure 5.2 shows strong ∼1% variations on sub-day timescales. From the periodogram (Figure 5.10), it is clear that the variability displayed by S Dor is quite complicated. Prewhitening reveals a total of 41 unique frequencies, the most of any star in the sample. However, we find no harmonics in this list of frequencies. Similar to HD 269582, we find two frequencies (0.66258409 and 1.6460147 day^{-1}) that are each the sum of two different pairs of frequencies. Due to the lack of any single dominant signal, the complexity in the periodogram, and the current theoretical debate on the physical origin of S Dor outbursts, we reserve further modelling until a longer baseline TESS light curve is available, in the hopes of measuring lower frequencies, and resolving the periodogram peaks better.

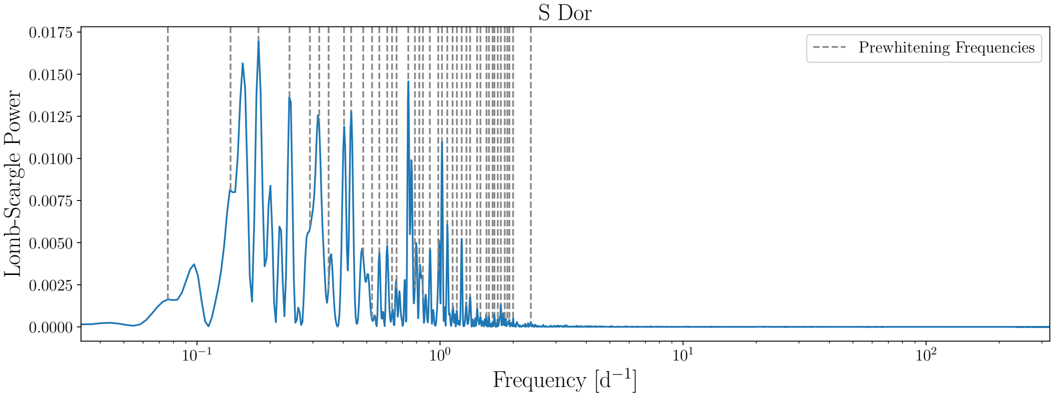

Figure 5.10 Similar to Figure 5.8 for S Dor. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [71].

LBV Noise Properties

In addition to our search for coherent variability in the two LBVs, we also analyze the noise properties of their light curves, using Equation 5.12 to fit the (log of the) Lomb-Scargle periodograms between 1/30 day^{-1} and the pseudo-Nyquist frequency. The resulting fit parameters are presented in Table 5.4, and the fits themselves are shown in Figure 5.11.

The two LBVs have fairly distinct fit properties, as may be expected given their different tem-

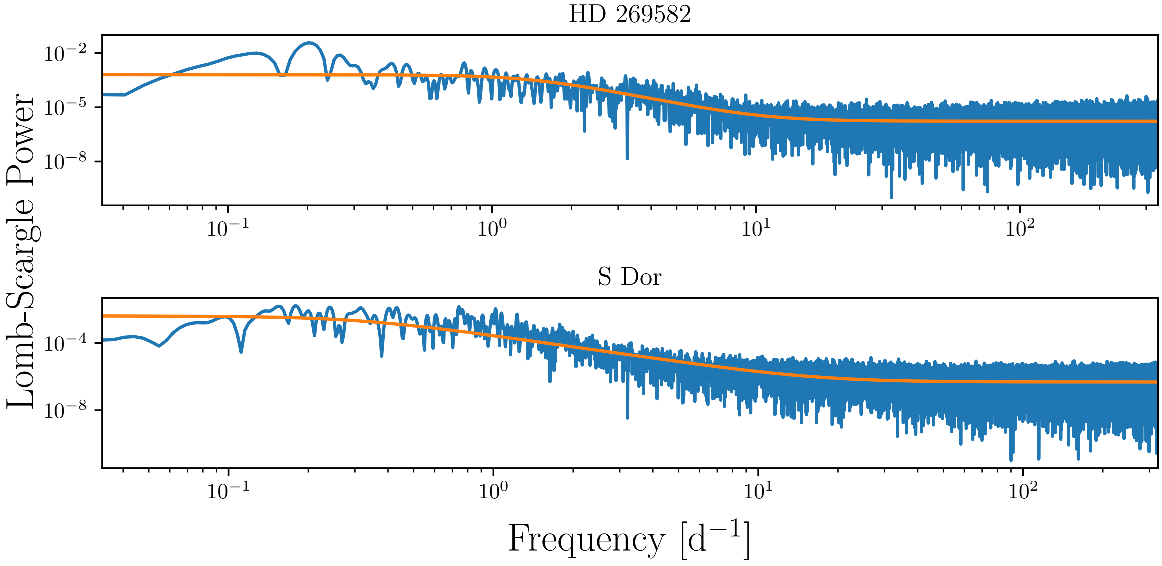

Figure 5.11 Lomb-Scargle periodograms calculated between 1/30 day^{-1} and the pseudo-Nyquist frequency for both LBVs. Fits using equation (5.12) are in orange. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [71].

peratures, and the recent evolution of HD 269582 into an LBV state. The noise in the S Dor light curve has a higher α_0, which may be expected given that the strength of the dominant period in HD 269582. Additionally, we find $\tau \approx 0.5$ day and $\gamma = 2.29$, while HD 2696582 has $\tau \approx 0.1$ and $\gamma = 2.96$. In longer-cadence AAVSO data, [2] fit the power spectra with a pure power law model, and found slopes closer to 2 for strongly flaring objects, and flatter slopes for LBVs in quiescence. While the TESS data don't probe the low-frequency regime measured by [2], they do indicate that, in HD 268582, the slope of the stochastic noise is steeper than expected, while in S Dor, the region of the power spectrum where the power-law behavior dominates extends over a wide range of frequencies. This suggests that the variability on sub-day timescales in LBVs may be generated by a mixture of physical processes.

Table 5.4. Summary of the fit results to the periodograms of the two LBVs in our sample..

Common Name	$\alpha_0/10^{-4}$	$\tau/10^{-2}$d	γ	$\alpha_w/10^{-5}$
HD 269582	6.33 ± 0.74	11.12 ± 0.76	2.96 ± 0.08	0.17 ± 0.0011
S Dor	41.23 ± 7.10	50.55 ± 4.68	2.29 ± 0.03	0.05 ± 0.000

5.4 Discussion

From this small sample of stars it is impossible to make many sweeping inferences. However, the broad range of light curve characteristics, unexpected characteristic time scales, and the structured noise properties displayed by almost every star in this sample make it clear that rare, evolved massive stars are prime candidates for study with TESS and subsequent missions.

Of the light curves that display clear periodicity, only one (the LBV HD 269582) appears to be on a timescale that could be consistent with a rotational period. This is unsurprising given the typical inflated radii of yellow supergiants (∼hundreds of $R_\odot$), and the incredible complexity of the lightcurve of S Doradus. However, the detection of possible rotational modulation is extremely exciting. Rotation is a deeply important parameter for massive stars which can have drastic effects on their evolution [76]. Current samples of measured rotation periods in massive stars are insufficient to statistically measure the distribution of rotation rates, leaving us with spectroscopic measurements [e.g., 114] which are hindered by the unknown inclination of the star relative to the line of sight. *Kepler* has revolutionized the study of stellar rotation for low-mass (FGKM) stars [e.g. 310, 58]. It is our hope that by extending our measurements of massive stellar rotation periods beyond the main sequence, TESS will provide similar advancements.

Stars with periodicities inconsistent with rotation (or binary interactions) possess variability on timescales consistent with pulsations. Pulsational modes can give us deep insight into fundamental stellar properties like mass and radius. However, models of both radial and nonradial pulsations in evolved massive stars have only recently been made available [e.g., 123]. Developing suitable models will allow us to constrain the interior structures of massive stars, and understand energy transport at an unprecedented level. The impact of wave-energy deposition in the last century of a massive star's life can have important impacts on its pre-supernova evolution [83], and measuring the pulsational properties of the most massive stars will give us valuable constraints on the masses of supernova progenitors.

Finally, red noise is a ubiquitous property in all of the light curves. Whether this noise arises from decoherent pulsations, surface granulations in the cooler stars, wind instabilities in the hot stars, or some other process entirely, measuring the noise characteristics of a large sample of massive stars will allow us to search for trends as a function of evolutionary stage, which can give us some insight into the physical processes involved. All told, studying evolved massive stars at short timescales can help us answer many unsolved problems in massive star evolution.

5.5 Summary & Conclusion

Our main results are summarized as follows:

- We study eight evolved massive stars. We find distinct periodicity in five stars, including two luminous blue variables, and three yellow supergiants. We are unable to constrain the source of the variability in any case.

- The light curve of one YSG, HD 269953, displays unique properties not shared by its fellow YSGs. We suggest that it is in a post-RSG evolutionary phase.

- All of the YSGs display red noise in their light curves, suggesting that the noise is likely astrophysical in origin.

- The LBV HD 269582 displays 1% variability at $\sim$5 day timescales. While the shape of the variability changes, it is possibly due to a rotation period that is imprinting itself into the wind of HD 269582 via a co-rotating interaction region.

- S Doradus exhibits incredible complexity in its peridogram at frequencies below ~ 1.5 day^{-1}, with a total of 41 unique frequencies found via prewhitening.
- Both LBVs display red noise. The noise in S Dor is stronger (as parametrized by α_0), less steep, and has a longer characteristic timescale τ compared to HD 269582.

We wish to emphasize that evolved massive stars have never been studied before with high cadence space based photometry. As the observed baseline increases for stars in the TESS sourthern CVZ, the periodogram peaks will grow sharper, and allow us to probe lower frequencies for comparison with previous studies. However, our tentative results presented here highlight the potential for studying massive stars in this domain. It is clear that new models are required to explain the observed variability, which will allow these data to give us an incredibly deep insight into the physics of evolved massive stars.

Chapter 6

A NEW CLASS OF COOL, PULSATING SUPERGIANTS

6.1 Introduction

The exact evolutionary pathway a star takes in its late phases and its final fate as a function of initial mass and composition are both incredibly sensitive to the physics of mass loss. Recent advances have explored the important roles that interior mixing, pulsations, rotation, binary interactions, and magnetic fields can play on the evolution of the most massive stars [159]. However, a discrepancy still exists between the highest mass observed red supergiants ($M \lesssim 25M_{\odot}$) and the highest mass observed supernova II-P progenitors ($M \lesssim 17-20M_{\odot}$, see Smartt et al. 2009 [265]). Many solutions to this *red supergiant problem* have been proposed by observers and theorists alike, including considering extinction effects to account for underestimated progenitor luminosities (e.g. [319], [141]); reexamining bolometric corrections that are extremely sensitive to effective temperature to accurately estimate supernova progenitor luminosities(e.g. [162, 158]); attemping to quantify biases in progenitor mass estimates (e.g. [60], [140]); intricately mapping the landscape of explodability in stellar models [283, 284, 285]; and incorporating well-tested prescriptions for RSG mass loss showing that the highest-mass ($\geq$20$M_{\odot}$) RSGs may simply evolve back to the blue side of the H-R diagram before explosion [e.g. 76, 213].

A more direct solution is to find yellow or blue stars that have likely already experienced a RSG phase [98], explicitly determining the mass of stars that can no longer end their lives as RSGs. Such post-RSGs allow us to place critical observational constraints on which stars *do not* simply evolve redward from the main sequence and then explode. The literature contains many examples of methods for finding these post-RSGs. For example, surface abundance enhancements of CNO-cycle elements are indicative of both envelope loss and convective mixing that extends from the envelope to the core during the RSG phase. Alternately, stars with evidence of past strong mass loss are likely candidate post-RSGs, and can be detected either via infrared excesses caused by warm circumstellar dust, or by direct detection of ejected mass [e.g. 119, 261], assuming that the CSM is detectable. Finally, though stars' first crossing of the HR diagram proceeds relatively unimpeded, a small number of very luminous yellow supergiants are observed to encounter the "yellow void" where their envelopes become dynamically unstable, resulting in outbursts [216, 282]. Such outbursting stars have been proposed to be post-RSGs.

A possibility that has only recently been explored is to identify post-RSGs by searching for stable pulsations in evolved massive stars. While massive stars crossing the HR diagram are not expected to pulsate, so-called strange modes can arise in stars with high ratios of luminosity to mass (L/M). Saio et al. (2013) [249] proposed this as the mechanism driving variability in variable B and A supergiants (α Cygni variables). While α Cygni variables are promising candidate post-RSG objects, their observed surface abundances are discrepant with predictions from stellar evolution models [249]. Therefore, a better route may be to explore the Hertzprung gap between blue and red supergiants, and search for pulsating *yellow* supergiants (YSGs) that may pulsate for the same reason as α Cygni variables, and whose spectra are more amenable to precision abundance analyses. The discovery of a group of pulsators that are cooler than α Cygni variables, and separate from both the "yellow void" and the brightest Cepheids, would thus be of great use.

Unfortunately, concrete predictions for whether such a group of pulsators should exist are currently out of reach given that theoretical modeling of envelope stability in this regime of the HR diagram still encounters convergence difficulties [122]. Thankfully, where theory falls short, observations may provide a path forward. The *Transiting Exoplanet Survey Satellite* (*TESS* , [242]) is collecting lightcurves of the brightest stars across 85% of the sky. While its primary mission has been to search for exoplanets, many of the brightest massive stars have already been observed at two-minute cadence for approximately 27 days at minimum, with observations of stars in the northern and southern continuous viewing zones (CVZs) lasting an entire year. *TESS* has thus allowed us to measure microvariability in an unprecedentedly large sample of cool supergiants. In Chapter 5, we examined a small sample of evolved massive stars that had been observed in *TESS* Sectors 1 and 2, and found evidence for fast pulsations in three YSGs. Here, we utilize the first 22 Sectors of *TESS* data, and report the discovery of a group of yellow supergiants that exhibit rapid (< 1 d) multiperiodic variability. These stars are more luminous and warmer than the classical Cepheid instability strip, fainter than outbursting yellow "hypergiants," and notably cooler than the coolest α Cygni variables. We describe our sample and methodology in §6.2, then characterize and discuss the stochastic variability that is ubiquitous in the sample. We present a new class of fast yellow pulsating supergiants in §6.3, discuss the importance of this new class in §6.4 before concluding in §6.5. This work originally appeared in Dorn-Wallenstein et al. (2020b) [72].

6.2 Methodology

6.2.1 Sample Selection

We first created a sample of cool supergiants with well-measured effective temperature ($T_{\rm eff}$) and luminosity (L), drawing from a large sample of YSGs and RSGs in the Large Magellanic Cloud (LMC) from Neugent et al. (2012) [215], who used radial velocities to confirm the membership of stars in the LMC, along with updated formulae derived from Kurucz [148] and MARCS [106] to obtain $T_{\rm eff}$ and $\log L/L_\odot$ from $J-K$ photometry. Because we wished to provide a full picture of the HR diagram including RSGs, and none of the RSGs published in [215] have been observed by *TESS* (as described below), we also include the Galactic RSGs from [116], [162, 158, 161], and the unique RSG WOH G64 [163]. Finally, we discard stars with $\log L/L_\odot \leq 4$ to avoid contamination by lower-mass evolved stars (see, for example, [157]). The positions of the YSGs in the HR diagram have a typical error of 0.015 dex in $\log T_{\rm eff}$ and 0.10 dex in $\log L/L_\odot$ respectively. For RSGs observed by Levesque et al. [162, 158, 161, 163], we used their published estimates of $T_{\rm eff}$ and $\log L/L_\odot$ (derived from M_{bol}) where available. The $T_{\rm eff}$ measurements have typical uncertainties of ±25 K for M stars ($T_{\rm eff} \lesssim 3810$ K, $\log T_{\rm eff} \lesssim 3.581$), and ±100 K for K stars. The $\log L/L_\odot$ measurements have typical uncertainties of ~ 0.1 dex. One RSG not studied by Levesque et al., V772 Cen (= HD 101712), is a known RSG+B star binary. A spectrum of V772 Cen is published in [121], and available on the *Vizier* online service [218]. We obtained this spectrum, and used it to estimate $T_{\rm eff}$ and $\log L/L_\odot$ for the RSG member of the binary following

Levesque et al. [162], with comparable uncertainties. Because the method we use relies on the depths of the TiO bands, we do not expect the B star companion to significantly impact our T_{eff} and $\log L/L_{\odot}$ estimates. No suitable archival spectrophotometry exists from which we could estimate T_{eff} and $\log L/L_{\odot}$ for fourteen of the RSGs.

6.2.2 *TESS Observations*

We crossmatched our sample of cool supergiants to the latest version of the *TESS* Input Catalog (TIC, [280]) available on the Mikulski Archive for Space Telescopes (MAST). From this sample, we selected all stars with a magnitude in the *TESS* bandpass fainter than $T = 4$ (where *TESS* begins to saturate), and brighter than $T = 12$ to obtain sufficient signal-to-noise ratios (SNR) in the lightcurves to detect sub-ppt-level variability (see Chapter 5). We also omitted any stars with calculated `contratio` values above 0.1 to mitigate contamination by nearby stars[1]. We then downloaded target lists for *TESS* Sectors 1-22[2], and select the cool supergiants that have been observed at two-minute cadence. This results in a total of 28 YSGs and 48 RSGs.

We can then use TIC to estimate parameters for the fourteen RSGs without temperatures and luminosities. We use the T_{eff} estimate published in the TIC, as well as the radius measurement to estimate $\log L/L_{\odot}$. For stars in both the TIC and [162], the parameters from the TIC show generally good agreement with the results in [162] to within the errors. However, we do expect the errors on both parameters (especially luminosity) to be significant as the relations used to compute stellar parameters are only validated on dwarfs and giants [see §2.2 of 280]. Table 6.1 shows the name, TIC number, coordinates, proper motions, *TESS* magnitude, $\log T_{\text{eff}}$, and $\log L/L_{\odot}$ of each star, as well as the source used to determine their position in the HR diagram, whether the star is a RSG or YSG, and if the star is an α Cygni variable or belongs to the newly identified class of pulsating yellow supergiants (see below).

---

[1] See [280] for the exact definition of `contratio`

[2] *TESS* target lists are available online at `https://tess.mit.edu/observations/target-lists/`

Table 6.1. Names, TIC numbers, coordinates, proper motions, *TESS* magnitudes and positions in the HR diagram of the cool supergiants observed by *TESS,* ordered by effective temperature from coolest to warmest. The source of the $T_{\rm eff}$ and $\log L/L_\odot$ measurements is indicated, where N corresponds to [215], L to [162, 158, 161, 163], I to [121], and T to the TIC [280]. Typical uncertainties in $\log T_{\rm eff}$ and $\log L/L_\odot$ are 0.015 dex and 0.10 dex respectively in [215]. M stars from [162] have uncertainties of 25 K and 0.1 dex respectively, while the uncertainties in $T_{\rm eff}$ in K stars are somewhat larger (100 K). Quantities for RSGs derived from the TIC show good agreement with the values published by [162] where overlap exists. We also indicate whether the star is a RSG or YSG (indicated with "R" or "Y" respectively), and whether the star is a candidate α Cygni variable or belongs to the newly identified class of pulsators.

Common Name	TIC Number	R.A. [deg]	Dec [deg]	μ_α [mas/yr]	μ_δ [mas/yr]	T [mag]	$\log T_{\rm eff}$ [K]	$\log L/L_\odot$ $L_\odot$	Source	RSG/YSG?	Var. Type
V1092 Cen	290678703	174.10924243	-61.31944611	-6.709	0.744	5.290	3.534	4.448	T	R	-
HS Cas	52782147	17.08300080	63.58652909	-2.450	-0.357	5.887	3.535	4.560	T	R	-
HD 143183	423407817	240.40092730	-54.14322405	-2.301	-3.620	4.245	3.537	5.222	T	R	-
BD+35 4077	136034302	305.30862034	35.62126593	-2.846	-4.499	5.684	3.556	4.768	L	R	-
AD Per	348314378	35.12084468	56.99312317	-0.066	-1.423	5.357	3.543	4.587	T	R	-
KY Cyg	15065085	306.49184826	38.35213201	-3.574	-6.279	4.898	3.544	5.432	L	R	-
TYC 8626-2180-1	459005094	161.46107152	-59.48870180	-7.080	1.750	4.571	3.547	4.936	L	R	-
V589 Cas	399355842	26.52283837	60.99352149	-0.952	-0.488	5.843	3.547	4.716	L	R	-
RS Per	348607532	35.60122973	57.10947226	-0.371	-0.931	5.084	3.550	5.156	L	R	-
V602 Car	467450857	168.37488668	-60.09134769	-5.425	2.183	4.945	3.550	5.020	L	R	-
W Per	251118305	42.65788594	56.98341594	0.243	-1.991	5.625	3.550	4.732	L	R	-
V396 Cen	443405175	199.35433424	-61.58398415	-4.770	-1.758	4.580	3.550	5.212	L	R	-
BI Cyg	13249363	305.34119647	36.93214587	-2.929	-5.223	4.738	3.553	5.352	L	R	-
BC Cyg	13325866	305.41061705	37.53303272	-3.856	-5.835	5.094	3.553	5.280	L	R	-
SU Per	348528265	35.52872734	56.60413801	-0.617	-1.490	4.650	3.553	4.952	L	R	-
PZ Cas	272324954	356.01366443	61.78949643	-3.110	-1.808	4.972	3.556	5.324	L	R	-
ST Cep	63963820	337.54474090	57.00085201	-3.517	-2.837	5.150	3.556	4.088	L	R	-
RW Cyg	15888421	307.21079278	39.98178278	-3.255	-5.511	4.596	3.556	5.156	L	R	-
TZ Cas	378292562	358.23432055	61.00233067	-3.220	-2.075	5.562	3.556	4.988	L	R	-
BU Per	264731552	34.72204574	57.42132329	-0.526	-1.106	5.898	3.556	4.764	L	R	-
V349 Car	457427613	157.39738942	-57.96638247	-7.191	3.632	5.250	3.559	4.808	L	R	-
V774 Cas	399433806	26.75004525	60.37232574	-1.068	-0.601	5.853	3.559	4.616	L	R	-
HD 95687	466289471	165.39899669	-61.04883831	-6.746	1.084	4.647	3.559	4.948	L	R	-
V441 Per	445664243	36.34108308	57.43726049	-0.254	-1.559	5.283	3.559	4.820	L	R	-
HD 303250	458834083	161.08350153	-58.06484800	-6.875	2.935	5.585	3.559	4.936	L	R	-
RT Car	458861722	161.19645089	-59.41336782	-7.450	2.914	6.417	3.559	5.260	L	R	-
V772 Cen	321656644	175.45585098	-63.41457099	-5.508	1.089	5.258	3.560	4.630	I	R	-
HD 101007	319508664	174.23716722	-61.18277794	-6.647	0.928	4.885	3.562	4.368	T	R	-
V648 Cas	450147792	42.76645187	57.85553435	-0.184	-1.252	5.912	3.562	4.900	L	R	-
IX Car	465185147	162.60957843	-59.98238045	-6.054	2.311	4.902	3.562	5.128	L	R	-
W Cep	65034243	339.11484739	58.42609816	-3.329	-2.132	5.207	3.566	5.466	T	R	-
V910 Cen	290681168	173.93730736	-61.57806090	-6.699	0.937	5.304	3.568	4.516	L	R	-
V528 Car	466325776	165.77563786	-60.91072867	-7.130	1.875	4.335	3.568	4.912	L	R	-
YZ Per	245588987	39.60591607	57.04616613	-0.119	-1.391	5.160	3.568	4.684	L	R	-
V362 Aur	285640583	81.79257440	29.92105466	-0.678	-2.892	4.886	3.568	4.620	L	R	-
PR Per	348442493	35.42670692	57.86281915	-0.788	-1.328	5.425	3.570	4.440	T	R	-
FZ Per	348314886	35.24852231	57.15832387	-0.696	-1.223	5.649	3.571	4.468	T	R	-
V809 Cas	265186608	349.84905043	62.73977569	-2.257	-2.004	4.117	3.574	4.472	L	R	-
V439 Per	348671468	35.79610521	57.19943969	-0.308	-0.920	5.803	3.580	4.420	L	R	-
V605 Cas	348436054	35.09359712	59.67136417	-0.711	-0.953	5.710	3.585	4.920	T	R	-
41 Gem	337334476	105.06593015	16.07900049	-2.088	-4.853	4.230	3.597	4.341	T	R	-
RW Cep	422108142	335.77923003	55.96322672	-3.616	-2.349	4.370	3.597	5.470	T	R	-
HD 155603	188405014	258.61523030	-39.76665102	-0.900	-1.087	4.138	3.601	4.870	T	R	-
NR Vul	435670188	297.54969991	24.92338263	-2.320	-5.807	5.421	3.602	5.348	L	R	-
QY Pup	334352580	116.91052662	-15.99068889	-2.162	3.511	4.898	3.608	4.756	T	R	-
HD 17958	390806332	44.10270614	64.33244354	-3.739	0.017	4.124	3.623	4.548	L	R	-
HD 33299	367172191	77.64572922	30.79754031	-0.015	-3.142	5.274	3.633	4.044	L	R	-

Table 6.1 (cont'd)

Common Name	TIC Number	R.A. [deg]	Dec [deg]	μ_α [mas/yr]	μ_δ [mas/yr]	T [mag]	$\log T_{\rm eff}$ [K]	$\log L/L_\odot$ $L_\odot$	Source	RSG/YSG?	Var. Type
AZ Cas	444831689	25.56880634	61.42120644	-2.198	-0.263	7.053	3.656	4.550	T	R	-
SK -67 57	40603917	77.96699736	-67.16603943	1.552	0.179	11.736	3.656	4.519	N	Y	-
HV 883	30526897	75.03151958	-68.45001791	1.785	-0.043	11.196	3.680	4.841	N	Y	-
HD 269953	404850274	85.05069622	-69.66801469	1.718	0.692	9.267	3.692	5.437	N	Y	FYPS
HD 269110	40404470	77.29420213	-69.60339017	2.081	0.252	10.038	3.750	5.251	N	Y	FYPS
HD 268687	29984014	72.73273606	-69.43125133	1.833	-0.114	10.465	3.784	5.169	N	Y	FYPS
HD 269840	277108449	84.04200662	-68.92812902	1.487	0.682	10.132	3.791	5.335	N	Y	FYPS
HD 269902	277300045	84.53992899	-69.10592146	1.707	0.628	9.790	3.793	5.352	N	Y	FYPS
HD 269331	179206253	79.50763757	-69.56049032	1.772	0.291	10.114	3.810	5.307	N	Y	-
RMC 137	404768745	84.65400862	-69.08552151	1.886	0.908	11.878	3.847	4.543	N	Y	-
CPD-69 430	277172433	84.23672142	-69.27176831	1.778	0.620	11.922	3.857	4.581	N	Y	-
W61 27-27	277025859	84.01580306	-69.02503035	1.526	0.556	10.770	3.861	4.493	N	Y	-
HD 269392	179376451	79.96588265	-69.88570140	1.979	0.257	11.961	3.865	4.605	N	Y	-
HD 269128	40518041	77.59495623	-68.77328288	1.862	0.284	9.189	3.872	5.134	N	Y	-
HD 269700	425081475	82.96784116	-68.54412683	1.602	0.401	8.808	3.882	5.069	N	Y	-
HD 270151	389749856	87.26183616	-70.04170277	1.756	0.847	10.561	3.897	4.635	N	Y	-
CPD-69 491	404852071	85.20335494	-69.28089004	1.801	0.533	10.298	3.914	4.665	N	Y	-
HD 270754	294872353	71.76854552	-67.11475533	0.754	0.830	11.191	3.915	4.927	N	Y	-
HD 269655	391810734	82.65769726	-68.41088876	1.495	0.972	11.123	3.924	4.497	N	Y	-
HD 269997	404933493	85.33497338	-69.08535421	1.658	0.890	9.313	3.927	4.970	N	Y	-
W61 6-77	389363675	85.56405367	-69.22241953	1.658	0.424	11.069	3.969	4.502	N	Y	-
HD 269777	276864600	83.57692409	-67.30380846	1.373	0.795	11.153	3.976	5.067	N	Y	-
CPD-69 394	276936320	83.65036911	-69.76013942	1.642	0.341	10.745	3.984	4.571	N	Y	-
HD 269992	404967301	85.36531401	-69.80104224	1.969	0.735	9.256	3.990	5.096	N	Y	-
HD 269786	277022505	83.76501015	-69.75056435	1.872	0.508	9.655	4.000	5.116	N	Y	-
HD 269101	40343782	77.43830593	-68.76940998	1.760	-0.125	10.577	4.027	4.799	N	Y	α Cyg
SK -69 68	40515514	77.49501756	-69.11716271	2.311	0.180	11.526	4.029	4.611	N	Y	α Cyg
HD 268798	30317301	74.28356450	-68.42008185	1.942	-0.074	10.103	4.033	5.071	N	Y	α Cyg
HD 269769	276936458	83.62856915	-69.78104781	1.783	0.526	10.700	4.037	4.714	N	Y	α Cyg

Using the Python package `astroquery`, we queried MAST and downloaded all available two-minute cadence lightcurves for each target. The data are provided by the *TESS* Science Processing Operations Center (SPOC), and include two flux measurements as a function of time: a simple aperture photometry measurement (`SAP_FLUX`) and flux measurements that have been corrected for systematic trends in the data (`PDCSAP_FLUX`). The time at each cadence is the photon arrival time at the solar system barycenter, correcting for the position and movement of the *TESS* spacecraft. For the following, we used the `PDCSAP_FLUX` lightcurves. To stitch together lightcurves from different *TESS* sectors, we divided each sector's lightcurve by its median flux.

6.2.3 *Stochastic Low Frequency Variability Across the Upper HR Diagram*

The top panel of Figure 6.1 shows the lightcurve of the YSG HD 269953, a star that we concluded is a likely post-RSG in Chapter 5, based on its pulsations and infrared excess measured by *Spitzer*

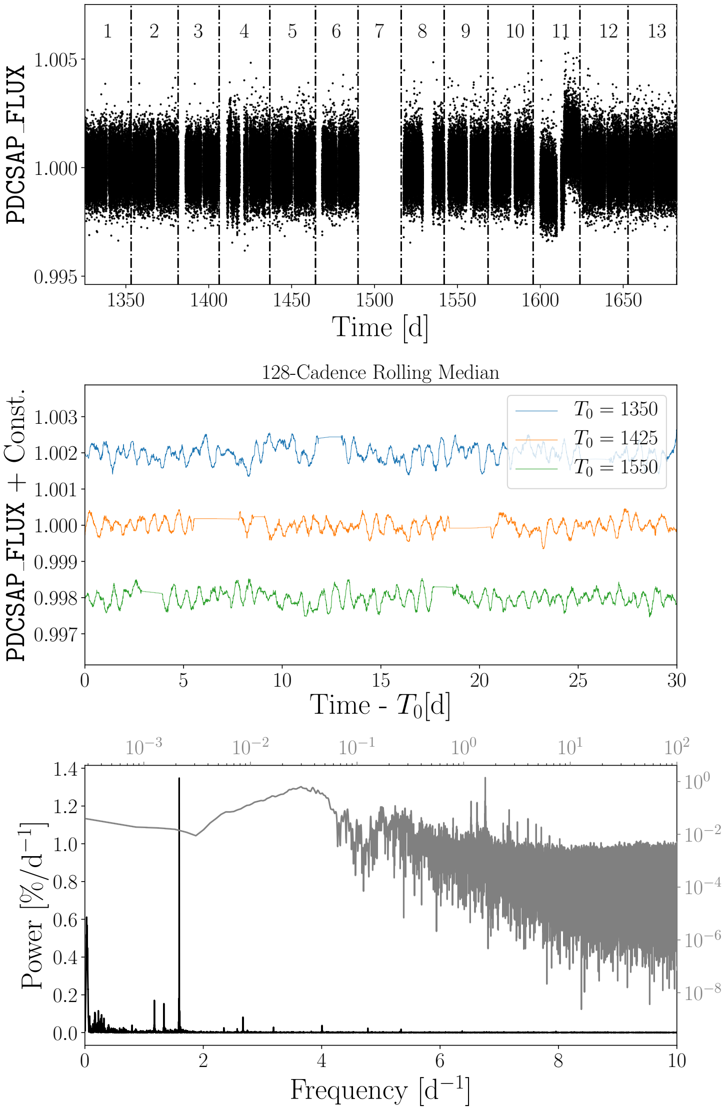

Figure 6.1 (*Top*): `PDCSAP_Flux` lightcurve of the YSG HD 269953. Sector numbers are indicated, with the boundaries between *TESS* sectors marked as dash-dotted black lines. (*Middle*): Same as above, zooming in to three 30-day windows, each beginning at the epoch given in the legend, and smoothing with a 128-cadence rolling median to highlight coherent variability. (*Bottom*): Periodogram of the entire unsmoothed lightcurve. Power is multiplied by 100. Linear (logarithmic) scaling for both frequency is shown for the black (grey) line, with corresponding black (grey) axis labels. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [72].

[26]. Individual `PDCSAP_FLUX` measurements are shown as black points[3]. Dash-dotted vertical

---

[3] A rapid dimming/brightening event is visible in the Sector 11 lightcurve that arises due to a combination of systematics in the detrending and a discontinuity after the mid-sector downlink. However, as revealed in the wavelet

lines show the boundaries between *TESS* sectors, and sector numbers are indicated. The middle panel shows a zoom-in to three different thirty-day portions of the lightcurve, now plotting the data after smoothing with a 128-cadence rolling median. Oscillations can clearly be seen. We analyze the frequencies of these oscillations using Lomb-Scargle periodogram [166, 256] of the unsmoothed data calculated with the `astropy` package (bottom panel). We use the `psd` normalization option, and divide the power by the number of points in the lightcurve. The resulting quantity is equivalent to the absolute value of the power spectral density, $|PSD|$, in units of normalized flux squared. We use the default `astropy` heuristics for choosing the frequency grid; the maximum frequency is set by the pseudo-Nyquist frequency, $f_{Ny} = 1/(2\langle\Delta t\rangle)$ (where $\langle\Delta t\rangle$ is the mean difference in time between two consecutive observations) and the frequency spacing is five times smaller than the Rayleigh frequency, $f_R = 1/T$ where T is the time baseline of the entire lightcurve. The black periodogram is plotted with a linear scaling (corresponding to the black axes labels), and the grey periodogram is plotted with logarithmic scaling on both the frequency and power axes (corresponding to the grey axes labels). Both are scaled to be in units of $\%/d^{-1}$.

As found previously in Chapter 5, the periodogram displays prominent peaks, superimposed upon a frequency-dependent background. The background shows rising power at low frequency that levels off at the lowest frequencies — i.e., red noise or *stochastic low frequency variability* – that is clearly visible in log scaling. Examining the periodograms of the entire sample shows that SLFV is present throughout the region of the HR diagram occupied by our sample. Adding on the fact that SLFV has been identified in hot O and B stars [24, 29, 27, 28], its presence throughout this sample of A-M supergiants suggests that it is in fact a ubiquitous feature of massive stars. Figure 6.2 shows the periodograms of four stars. The power is normalized to have a maximum value of 1, and an arbitrary offset constant is added for clarity. The top two periodograms are calculated for two “normal” supergiants that are representative of the overall sample: the red supergiant BD+35 4077, and the yellow supergiant HD 270754. Their periodograms are dominated by SLFV, and they display no strong peaks at frequencies above 1 d^{-1}. The bottom two periodograms belong to

---

analysis below as well as a by-eye inspection of the periodogram computed only on the light curve before this event, this discontinuity only manifests itself as a low frequency transient and has no effect on the recovered frequencies.

two yellow supergiants: HD 269101 and HD 268687. While both stars' variability are dominated by SLFV, they also show visible peaks in their periodograms at high frequencies.

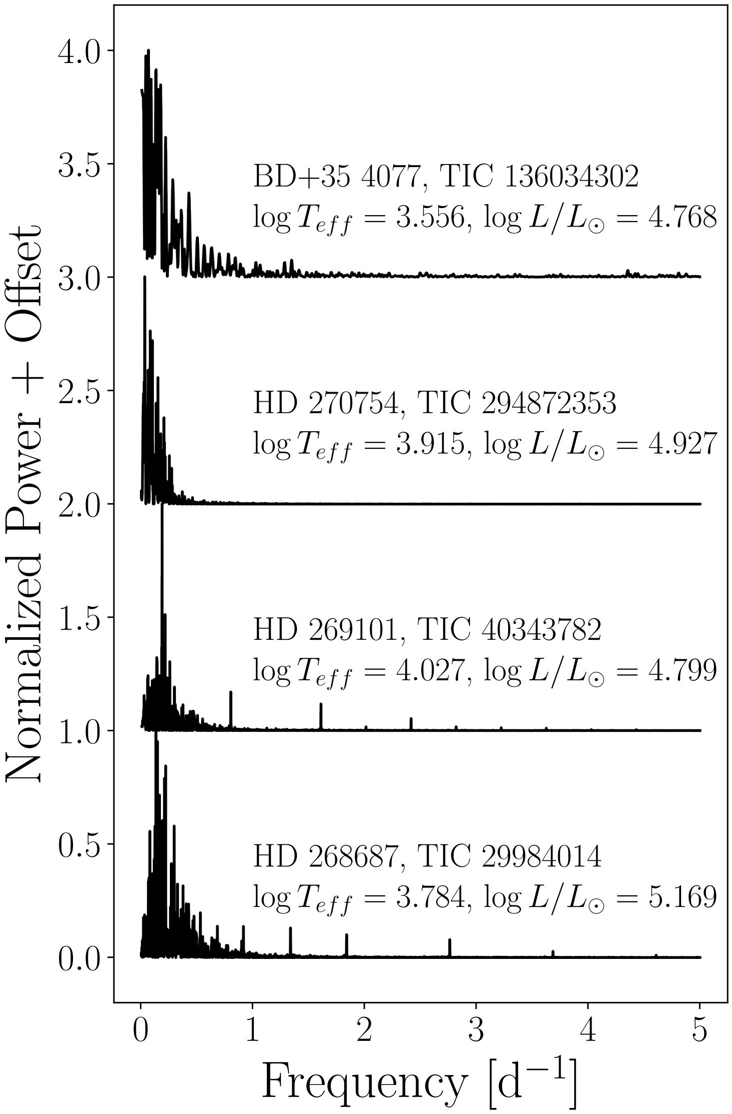

Figure 6.2 Periodograms of four stars that are representative of the entire sample: the RSG BD+35 4077, the YSG HD 270754 (neither of which appear to pulsate), the candidate α Cygni variable HD 269101, and HD 268687, which belongs to the newly identified class of pulsating YSGs. The power has been normalized by the maximum power, and an arbitrary offset has been applied for clarity. Stochastic Low Frequency Variability dominates the power at low frequencies. However, it is possible to see real peaks superimposed on the background of the bottom two periodograms. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [72].

Figure 6.2 suggests a potentially interesting dichotomy among the stars in our sample, distinguished by the presence or absence of high frequency pulsations. Quantifying the presence or absence of these pulsations requires that we characterize the underlying broad spectrum of noise in the periodogram (i.e., the SLFV). We proceed here by first characterizing all star's SLFV, allowing us to look for correlations between the SLFV and other stellar properties, before removing the SLFV from the periodogram and identifying the narrow features in the periodogram. We present an automated search for these features and then analyze their properties as a function of stellar subtype, expanding on the examples shown in Figure 6.2.

To characterize the SLFV, we follow [24] and [29, 27, 28] and fit the amplitude spectrum ($\alpha(f)$, obtained by taking the square root of the PSD) with a phenomenological model. As in Chapter 5, we adopt the function

$$\alpha(f) = \frac{\alpha_0}{1 + (2\pi\tau f)^\gamma} + \alpha_w \tag{6.1}$$

from [275], where f is the frequency, α_0 is the amplitude as $f \rightarrow 0$ in units of normalized flux, τ is a characteristic timescale in days on which the noise is correlated, γ sets the slope of the red noise, and α_w is an additional parameter we add in to model the white noise floor at the highest frequencies, also in units of normalized flux. We note that Equation (6.1) is equivalent to the function adopted by [29, 27, 28], and the characteristic frequency in those works is equivalent to $\nu_{char} = (2\pi\tau)^{-1}$. Figure 6.3 shows the square of Equation (6.1), $|\alpha(f)|^2$, with $\alpha_0 = 1$, and $\alpha_w = 10^{-3}$. The top (bottom) panel shows the effects of varying τ (γ) at constant $\gamma = 2$ ($\tau = 1$). The characteristic timescale τ determines the timescale/frequency on which the power-law behavior dominates over the uncorrelated noise at the lowest and highest frequencies, while the slope γ sets the range of frequencies where this occurs, with higher slopes resulting in a narrower frequency range.

We use the `curve_fit` routine within the `scipy` package to fit the base-10 logarithm of the amplitude spectrum to avoid the artificial weighting of real peaks at high frequencies. Note that we do not first remove any peaks from the periodogram before performing this fit. However, compared to the periodograms calculated from, e.g., the CoRoT lightcurves of hot stars studied by

Table 6.2. Names, TIC numbers, fit parameters from Equation (6.1), and corresponding errors for all stars in our sample.

Common Name	TIC Number	α_0 ppt	τ d	γ	α_w ppt
V1092 Cen	290678703	194.9581 ± 17.3575	0.3464 ± 0.0318	1.755 ± 0.049	7.3583 ± 0.0077
HS Cas	52782147	414.3034 ± 33.9049	0.1181 ± 0.0095	1.784 ± 0.042	11.6105 ± 0.0196
HD 143183	423407817	114.8980 ± 12.2582	0.2139 ± 0.0267	1.601 ± 0.061	7.9003 ± 0.0111
BD+35 4077	136034302	198.9679 ± 25.9046	0.6764 ± 0.0985	1.544 ± 0.056	8.5699 ± 0.0107
AD Per	348314378	173.5579 ± 10.9315	0.0597 ± 0.0039	2.268 ± 0.092	18.0802 ± 0.0291
KY Cyg	15065085	286.2191 ± 37.0462	0.6387 ± 0.0998	1.359 ± 0.042	13.1549 ± 0.0172
TYC 8626-2180-1	459005094	546.2201 ± 69.9013	0.3636 ± 0.0479	1.456 ± 0.032	9.0229 ± 0.0141
V589 Cas	399355842	92.4473 ± 7.3775	0.0915 ± 0.0080	2.142 ± 0.113	12.1910 ± 0.0188
RS Per	348607532	383.0987 ± 28.7313	0.0859 ± 0.0069	1.654 ± 0.037	14.4608 ± 0.0264
V602 Car	467450857	163.9396 ± 10.5384	0.1408 ± 0.0098	2.165 ± 0.093	21.4533 ± 0.0220

Note. — This table is published in its entirety in the machine-readable format. A portion is shown here for guidance regarding its form and content.

[27], the power of the low frequency excess seen in this sample is far stronger than the power of the observed peaks, with the exception of HD 269953, and so we don't expect the fit parameters to be significantly affected. In the case of HD 269953, fitting the logarithm of the amplitude spectrum mitigates any significant effect.

Figure 6.4 shows histograms with the distribution of the best-fit SLFV parameters for YSGs in yellow, with RSGs stacked on each bin in red. While the distributions for α_0 and γ appear to be similar between RSGs and YSGs, we can see that the YSGs have characteristic timescales clustered around $\log \tau/\mathrm{day} \approx 0.5$, while the distribution of $\log \tau$ for RSGs has a long tail extending to quite short timescales ($\log \tau/\mathrm{day} \approx -1.5$). We also see a clear difference in the instrumental noise, α_w, which is systematically lower for the YSGs. All YSGs are in the LMC, and thus were observed nearly continuously during the first year of *TESS* observations, with the exception of small gaps when e.g., individual stars passed through the gaps between CCDs in *TESS* camera 4. Therefore, we expect the signal in the periodogram (relative to the intrumental noise) to be much higher for these stars. In contrast, all the RSGs in our sample are Galactic, and most were only observed for a single *TESS* sector which reduces the lightcurves' sensitivity for RSGs.

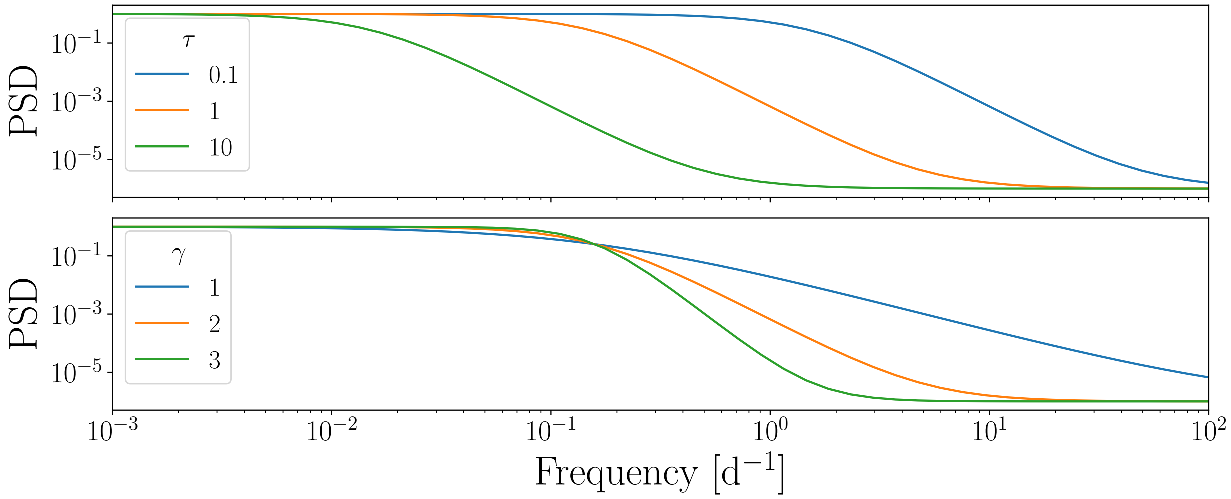

Figure 6.3 Plotting Eq. (6.1), squared so the units are comparable to the Lomb-Scargle periodogram, with $\alpha_0 = 1$, $\alpha_w = 10^{-3}$ and varying τ (top), and γ (bottom). Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [72].

To examine the behavior of these parameters further, we also plot them as a function of $\log T_{\text{eff}}$ in Figure 6.5. Errorbars are calculated from the covariance matrix returned by `curve_fit`, and the color corresponds to increasing $\log L/L_\odot$ from darker to lighter colors. The amplitude of the SLFV α_0 (top left) is suppressed at temperatures around $\log T_{\text{eff}} \sim 3.7 - 3.8$. The characteristic timescale τ (top right) slightly increases with increasing temperature. While we don't place any bounds on τ when performing the fits, there is an implicit upper limit to τ that can be seen in the data, as the SPOC processing pipeline and our sector-combining procedure effectively erases correlations in the lightcurves on timescales longer than a few days. The slope of the noise γ (bottom left) clusters around $\gamma = 2$ for the cooler stars, while favoring slightly smaller values for YSGs with $\log T_{\text{eff}} \gtrsim 3.7$. Finally, we again see that the instrumental noise α_w (bottom right) is systematically lower for the YSGs.

Convection or Something More?

As demonstrated above, the characteristic timescale of the background noise, τ, appears to increase with increasing temperatures. Given that τ roughly corresponds to the characteristic timescale over

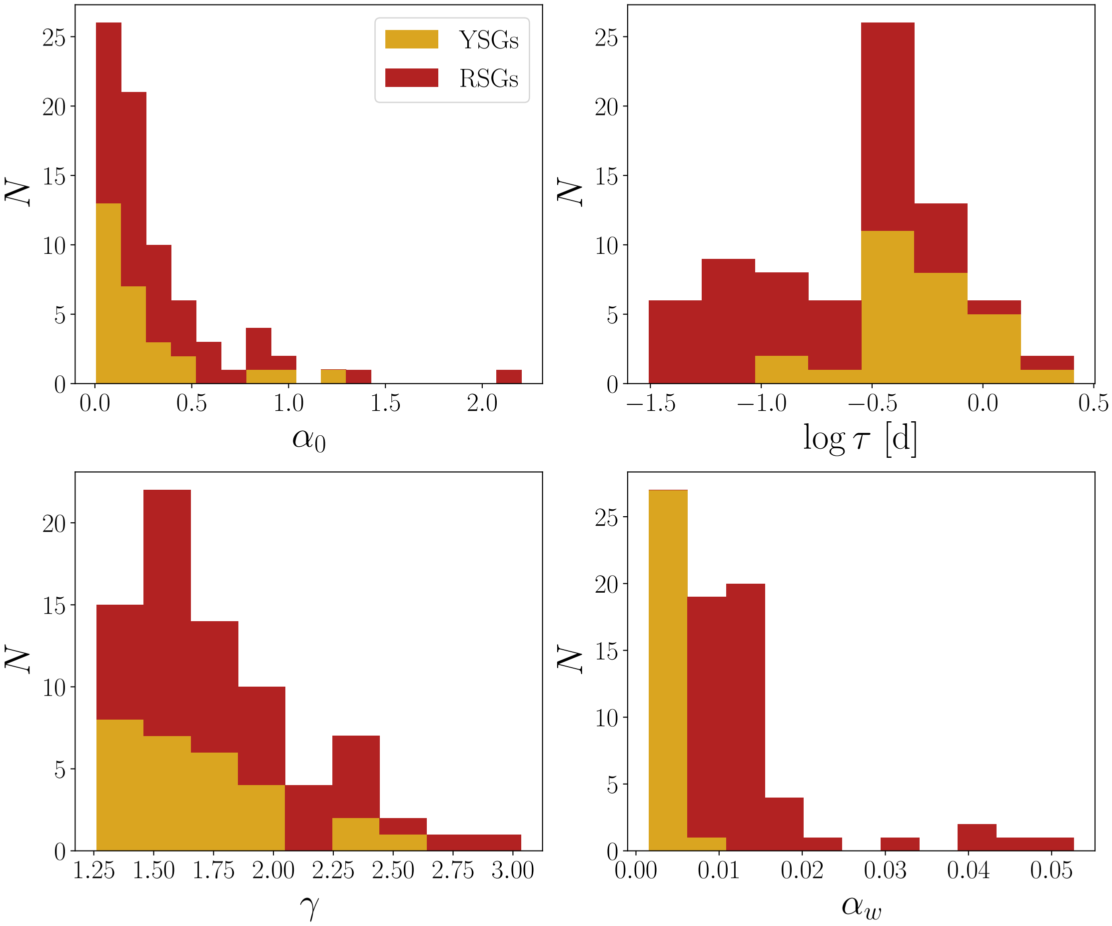

Figure 6.4 Histogram of the best-fit values for amplitude α_0 (top-left), characteristic timescale τ (top-right), power-law slope γ (lower-left), and instrumental noise α_w (lower-right). Bins are split into their contribution from YSGs in yellow and from RSGs in red. The YSGs have both higher characteristic timescales and lower instrumental noise values than the RSGs.

which the stochastic variability is correlated, the observed correlation gives us some clues as to the origin of the noise.

One possibility is that the noise is due to granulation in these stars' outer layers. Much lower mass yellow stars like our Sun show similar low-frequency power excesses due to granulation, however, noise due to convective granulation should leave the opposite scaling as what we observe. Kallinger et al. (2014) [130] demonstrate that, from first principles, the characteristic convective

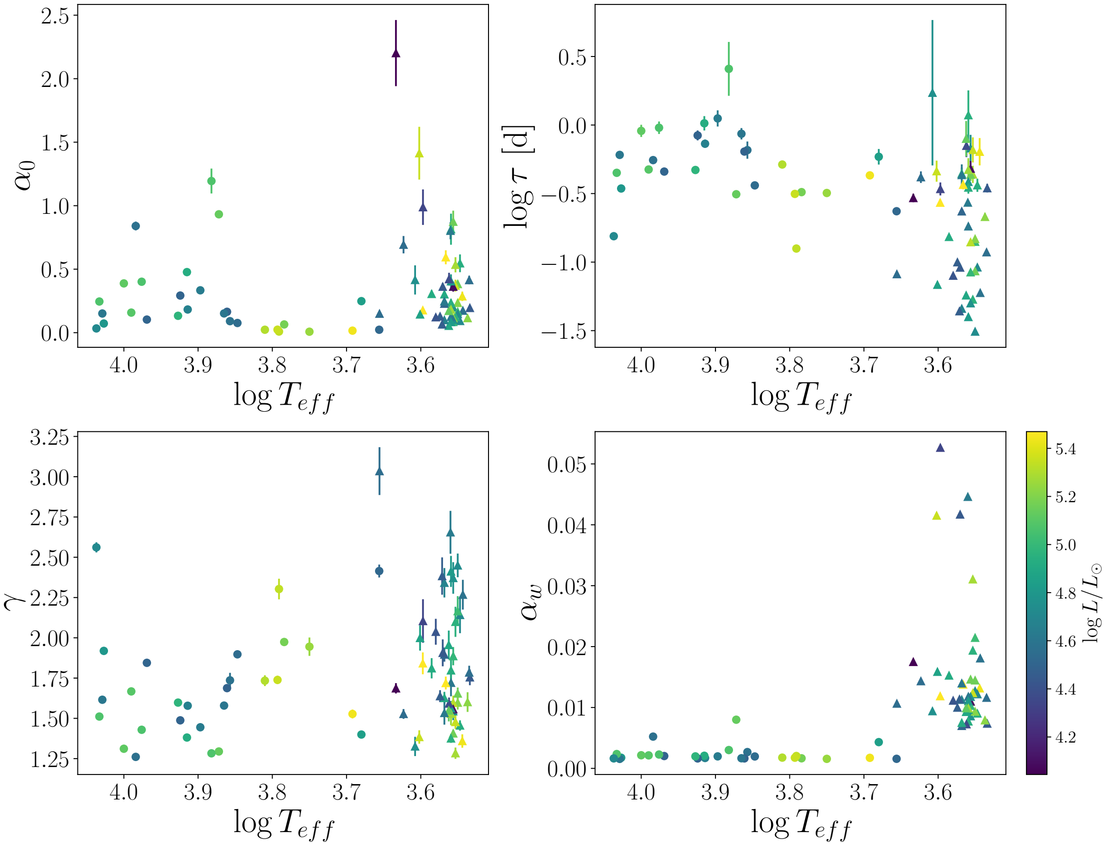

Figure 6.5 Best-fit values for amplitude α_0 (top-left), characteristic timescale τ (top-right), power-law slope γ (lower-left), and instrumental noise α_w (lower-right) as a function of $\log T_{\text{eff}}$. Stars are color-coded by luminosity, with YSGs plotted as circles, and RSGs as triangles. Errorbars are calculated from the covariance matrix returned by `curve_fit`.

timescale, τ_{conv} (denoted τ_{eff} in Kallinger et al.) scales with the surface gravity, g, and effective temperature as $g^{-0.85}T_{\text{eff}}^{-0.4}$ [130]. Because $g \propto MR^{-2}$ and $R^{-2} \propto T_{\text{eff}}^4 L^{-1}$, this implies $\tau_{conv} \propto T_{\text{eff}}^{-3.4}(L/M)^{0.85}$. Thus as a massive star evolves rightward at essentially constant L in the HR diagram, the convective timescale increases strongly as the star's surface cools (with a small boost as the star's L/M increases as the star loses mass). This is the exact opposite of the trend we observe, implying that the low-frequency variability that we see in our sample is not (solely) a result of surface convection, at least in the warmer stars.

Another possible explanation for the behavior of τ in Figure 6.5 comes from studies of hotter

($T_{\text{eff}} > 40$ kK) stars, where correlated stochastic variability has been linked with sub-surface convection zones [24], that may interact with pulsations [230], or possibly through internal gravity waves (IGW) arising from the boundary of internal convective and radiative layers ([29], though see [154] for possible caveats). Perhaps the stochastic variability in this sample is connected with that seen in hot stars? Recently, [28] characterized stochastic variability in a sample of 70 OBA stars spanning a range of temperatures above $\sim 10^4$ K and masses between ~ 5 and 80 $M_\odot$ (see Figure 2 in that work). Unfortunately, their sample has very few post-main sequence stars, especially in the mass range of our sample, and so we are unable to construct a complete sequence of α_0, γ, or $\tau \approx 1/\nu_{char}$ as stars evolve across the HR diagram. If stochastic variability proves to be attributable to IGWs, such an evolutionary sequence would be an incredibly powerful means of applying asteroseismology to massive stars as they near the ends of their lives.

It also remains possible, however, that sub-surface processes that *aren't* IGWs are causing the low frequency stochastic variability. Therefore, we cannot uniquely identify the stochastic variability with any particular source, and reserve such identification for a study of low frequency stochastic variability in massive stars across the entire HR diagram. That said, the RSGs display power law slopes clustered around $\gamma = 2$, consistent with what was found by Kiss et al. (2006) [138] in AAVSO data of RSGs — though in a significantly higher frequency range, and with lower amplitude than the lightcurves studied by Kiss et al. — and which was attributed by those authors to convective processes. Indeed, though the observed scaling of τ with T_{eff} is inconsistent with surface convection, the timescales of simulated turbulent convection in stellar *interiors* (albeit in lower mass stars) are not expected to obey the simple scaling presented above [101].

Indeed, Kiss et al. (2006) found that there is no point at which the power spectra of RSGs turn over; as observing time increases, more power at low frequencies is recovered. We note that while our analysis suggests longer timescales in warmer massive stars, it is not clear exactly how strong the correlation is. Because the *TESS* observations of RSGs only span a few sectors at most, and correlations on long timescales are smoothed out by the detrending performed by the SPOC, we might expect to see values of τ clustered at the maximum value possible given the detrending (a few tens of days). However, RSGs display the *smallest* observed values of τ

seen in our sample, suggesting that, at the precision of *TESS* , RSG microvariability is not entirely consumed by convective noise.

One possible way to distinguish between subsurface convection and IGWs from core convection may be to measure the macroturbulent velocity of these stars, and compare these measurements with the stars' locations in the HR diagram and observed values of α_0, τ, and γ (see both [102] and [28]). It is also entirely possible that both scenarios are at play, and contributing to the observed stochastic variability in this sample. Regardless of the physical origin, stochastic variability is ubiquitous across the upper HR diagram, from hot stars [e.g. 28] to cool supergiants (this work and [138]). Characterizing this variability and determining its origin has the potential to offer critical insight into the evolution of massive star interiors from birth to death.

6.3 Fast Yellow Pulsating Supergiants

After fitting the amplitude spectrum of each star to obtain the best-fit model, $\hat{\alpha}(f)$, we divide its signal out of the power spectrum by computing $|PSD|/\hat{\alpha}(f)^2$. The resulting quantity has no formal definition, but is incredibly useful at showing the power of peaks relative to the background. Hereafter, we refer to the background-normalized power spectrum as the *residual power spectrum* (RPS). Figure 6.6 shows the HR diagram, where each star in the sample is replaced by its RPS between 0 and 5 d^{-1}, normalized by its maximum value and scaled to fit in the plot. Note that because of this scaling, the relative heights of peaks in two different RPS have no relation, but the relative heights of two peaks in the same RPS are meaningful. In particular, plotting the RPS in this way allows us to simultaneously assess the approximate signal to noise in the periodogram as a function of each star's location in the HR diagram. A subset of the nonrotating, $Z = 0.006$ evolutionary tracks calculated with MESA and described below are plotted as thin black lines with the initial masses.

For the majority of stars in this sample, their lightcurves appear to be entirely composed of SLFV, and their RPS show either no peaks, or small peaks with low signal to noise. There are, however, two groupings of stars with high signal-to-noise peaks in their RPS. The first group, comprised of a vertical strip of four stars with $\log T_{\mathrm{eff}} \approx 4.0$ lies in the region of the HR diagram

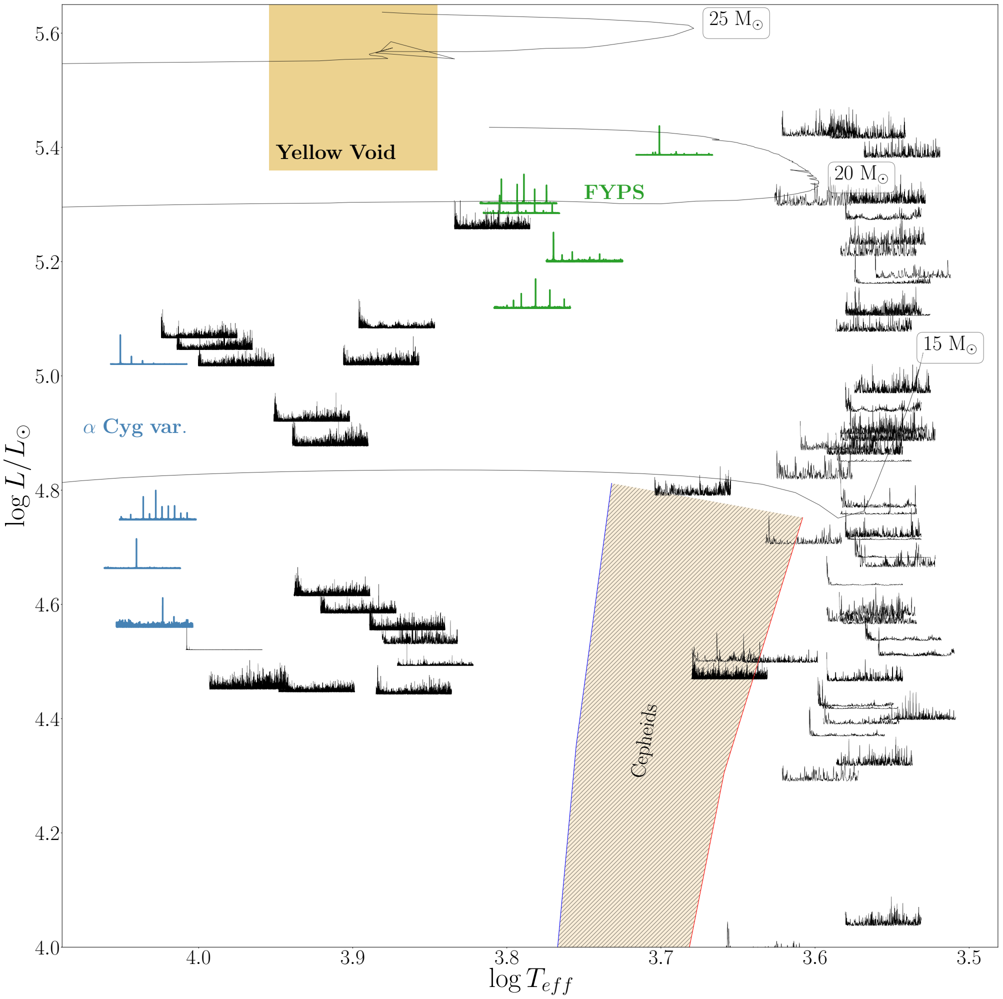

Figure 6.6 HR diagram showing the residual power spectra of each star in our sample between 0 and 5 d^{-1}, obtained after dividing out the SLFV in each periodogram, and normalized by the maximum value. Each RPS is centered on the T_{eff} and L of the star it corresponds to. The rough boundaries of the yellow void is shown as a goldenrod rectangle. The Cepheid instability strip derived from nonrotating, LMC-metallicity ($Z = 0.006$) stellar models on their first crossing of the HR diagram from [7] are shown by the yellow crosshatched region. Four α Cygni variables are highlighted in blue. Separate from all three regions of instability, we find five stars, highlighted in green, that display prominent, high signal-to-noise peaks in their RPS, which we name Fast Yellow Pulsating Supergiants (FYPS). For reference, we plot a subset of the nonrotating, $Z = 0.006$ evolutionary tracks calculated with MESA that are described in the text as solid black lines, with their initial masses indicated by the corresponding box. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [72].

where α Cyg variables are expected to be found[4], which we highlight in blue. We searched for any past work that has studied their variability, as listed in SIMBAD [322], and find the following:

- HD 268798 was previously identified as an eclipsing binary by [19], and as a rotational variable with ellipsoidal variations by [229]. It has not previously been identified as an α Cyg variable.
- HD 269101 was identified as a candidate α Cyg variable by [19], and incorrectly identified as a Slowly Pulsating B-star by [229] (likely due to its entry in SIMBAD erroneously listing it as an early-B supergiant).
- HD 269769 has not previously been studied in the time domain, and thus has not previously been identified as an α Cyg variable.
- Sk-69° 68 has not previously been studied in the time domain, and thus has not previously been identified as an α Cyg variable.

Apart from these candidate α Cyg variables, we also find a cluster of five stars, all with $5.1 \leq \log L/L\odot \leq 5.5$ and $3.69 \leq \log T_{\text{eff}} \leq 3.8$ with high signal to noise peaks in their RPS. We highlight these stars in green in Figure 6. This region of the HR diagram contains no other stars[5]. We list the common names (found on SIMBAD), TIC numbers, coordinates, temperatures, and luminosities for the five stars in Table 6.3. They are well-separated in the HR diagram from the lower-right edge of the yellow void — a region of the HR diagram occupied by a very small

---

[4] The frequencies found in the lightcurves of these stars are somewhat higher than in other α Cyg variables [249]. However, most observations of α Cyg variables have been taken from the ground, where detecting frequencies around ~ 1 d^{-1} is difficult. While not the focus of our work, we note these candidate α Cyg variables for use by other authors.

[5] One star, HD 269331, has a similar luminosity and a $\log T_{\text{eff}}$ that is 0.02 dex higher than the warmest identified pulsator. Its RPS is low signal-to-noise and shows no significant peaks. In Chapter 5 we identified two prominent bumps in the first two sectors of *TESS* data. Examination of the remaining sectors shows that these bumps are present throughout the lightcurve, and that the star is displays variability with an amplitude of ~ 1 ppt. With no significant RPS peaks, we exclude HD 269331 from our subsequent analyses; however it is possible that HD 269331 is a genuine member of this novel class of supergiant pulsator.

Table 6.3. Names, TIC numbers, coordinates, temperatures, and luminosities of the five pulsating YSGs.

Common Name	TIC #	R.A. °	Dec. °	$\log T_{\rm eff}$ K	$\log L/L_{\odot}$ $L_{\odot}$
HD 269953	404850274	85.050696	-69.668015	3.692	5.437
HD 269110	40404470	77.294202	-69.603390	3.750	5.251
HD 268687	29984014	72.732736	-69.431251	3.784	5.169
HD 269840	277108449	84.042007	-68.928129	3.791	5.335
HD 269902	277300045	84.539929	-69.105921	3.793	5.352

number of stars that exhibit extreme variability and mass loss due to dynamical instabilities in their atmospheres [63] — which is shown in goldenrod, and the upper-left edge of the Cepheid instability strip (shown in cross-hatched orange, derived from $Z = 0.006$, nonrotating stellar models on their first crossing of the HR diagram; see [7]). This group includes the three pulsating YSGs previously identified in Chapter 5, as well as two newly identified stars. Many of the frequencies found in their lightcurves (see below) are on timescales shorter than 1 day, and, as discussed in Chapter 5, are hard to explain with rotational or orbital effects given the large radii of YSGs. Furthermore, while it is possible that these frequencies may arise in the winds of these stars, we deem it unlikely that only YSGs in this region of the HR diagram would show coherent modulations in their winds.[6] Finally, spectra of all five stars from [215] indicate that they are all fairly typical YSGs, though HD 269902 has a slightly weaker Ca II triplet. Therefore, we adopt the name “Fast Yellow Pulsating Supergiants” (FYPS)[7] for these stars, and discuss them below.

---

[6] Regardless, with the notable exception of a small number of incredibly luminous YSGs that have undergone outbursts, YSG winds remain poorly understood.

[7] The authors acknowledge the poor adjective ordering in this acronym. However, we believe that FYPS is easier to pronounce than FPYS.

6.3.1 Chance or New Class?

All five of the FYPS are located in the LMC. Each *TESS* pixel is 21" on a side ($\sim$ 17 ly at the distance of the LMC). Furthermore, YSGs are found in crowded regions with many hot young stars, making it highly unlikely that the starlight in the optimal aperture defined by the SPOC is coming only from these stars. This effect is somewhat mitigated by *TESS* 's relatively red passband (centered at 7865 Å); while in bluer passbands, the flux in the aperture may contain significant flux from nearby O and B stars, cool, evolved evolutionary phases of massive stars that dominate the flux in the aperture are significantly rarer due to their shorter lifetimes. Additionally, the binary fraction of massive stars is high [253, 252, 198]; even the most evolved red supergiants that are the most likely to have interacted and merged with a companion have a binary fraction of $\sim 20\%$ [212]. Therefore it is possible that we are recovering five stars with pulsating companions that happen to be located in a small region of the HR diagram by chance. Assuming the initial massive star binary fraction (as well as the initial period and mass ratio distributions) is roughly constant for the masses of stars in our sample, it is effectively equally likely for any star in our sample to have a pulsating companion (with decreasing likelihood for the largest stars in our sample). Finally, detrending of time series photometry can generate spurious low frequency peaks, which propagate to higher frequencies via harmonics and combinations with real peaks [e.g. 297]. Potentially, it could appear as if these five stars, which may be ordinary YSGs, are pulsating when they are in fact not.

Therefore, we need to assess the likelihood that the five stars identified above are otherwise normal YSGs, and their lightcurves are all contaminated by starlight from actual pulsating stars (whether from nearby stars in the aperture or a binary companion), or contain spurious periodic signals introduced by detrending. If five stars were randomly selected as "pulsators" from our sample due to contamination, we would not expect them to be found in such a small region of the HR diagram. To determine the extent to which crowding may have influenced this detection, we can ask the question: assuming any of the stars in our sample could have randomly been contaminated by pulsators, how likely is it to draw five stars from our sample and have them form a grouping in the HR diagram with equal or lesser size. We can answer this question directly with a simple bootstrap analysis. For each of the 18,474,840 unique subsamples of 5 stars, we calculate the dimensions, $\Delta \log T_{\rm eff}$ and $\Delta \log L/L_{\odot}$, of the smallest box in the HR diagram that contains the subsample. Note that this analysis does not account for the decreasing likelihood of finding a binary companion around a large supergiant, and so in just the case of contamination by a binary companion, the derived likelihood is an uppper limit. Figure 6.7 shows the two-dimensional histogram of $\Delta \log T_{\rm eff}$ and $\Delta \log L/L_{\odot}$. The actual values of $\Delta \log T_{\rm eff}$ and $\Delta \log L/L_{\odot}$ calculated from the five FYPS is indicated by the red star. We find that only 8.8% (1.2%) of all possible subsamples have equal or lesser ranges of $\Delta \log T_{\rm eff}$ ($\Delta \log L/L_{\odot}$). All told, only 0.07% of all possible subsamples are bounded by a smaller region in the HR diagram. If we repeat this calculation with only the LMC YSGs, this number decreases to 0.03%. Thus, we deem it exceedingly unlikely that the lightcurves of these five stars (and only these five stars) happened to have randomly been contaminated by a nearby pulsator or pulsating companion, and conclude that we have discovered a genuine new class of pulsating star.

6.3.2 Variability and Pulsation Frequencies of FYPS

Prewhitening

The extraction and measurement of individual pulsation frequencies from lightcurves with frequency-independent noise is a fairly well-defined procedure: the strongest peak can be identified based on its amplitude or FAP, and a prewhitening procedure can be applied to iteratively fit and subtract sinusoids from the lightcurve corresponding to the extracted frequencies until some noise threshold is reached [e.g., 24]. The resulting lists of frequencies, amplitudes, and phases in principle, completely describe the coherent variability found in the lightcurve. However, in the case of frequency dependent noise (in this case, stochastic low frequency variability of astrophysical origin), spurious peaks that are random fluctuations superimposed on the noise are extracted, while true peaks that lie on top of the much lower-amplitude white noise at higher frequencies can be ignored. To account for this effect, we adopt the procedure used by [24], with the following modifications and

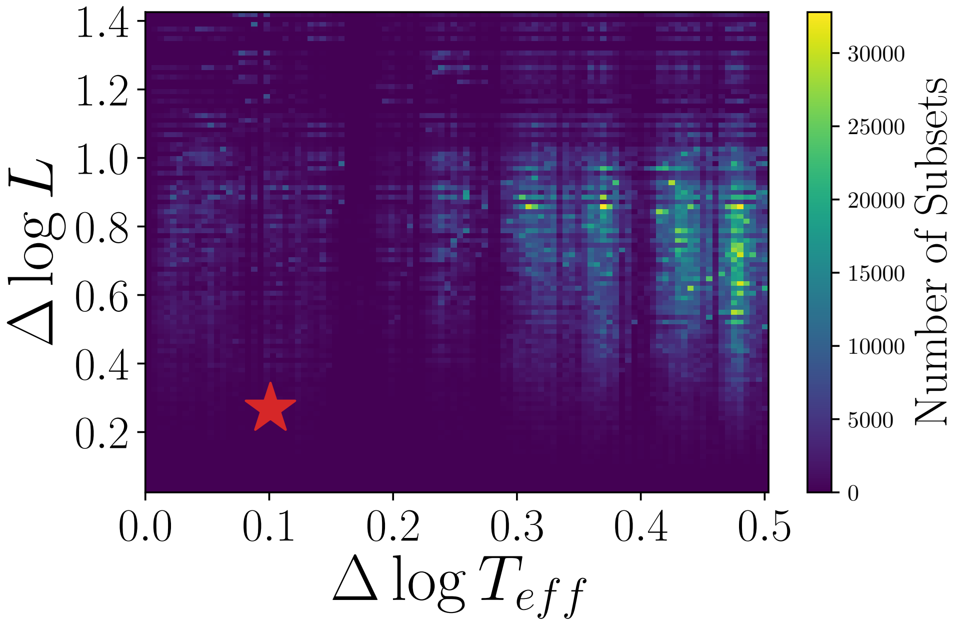

Figure 6.7 Two-dimensional histogram of $\Delta \log T_{\rm eff}$ and $\Delta \log L/L_\odot$ for all simulated five-star subsets. The actual $\Delta \log T_{\rm eff}$ and $\Delta \log L/L_\odot$ of the real sample of FYPS is shown with the red star. We conclude that the association of FYPS in the HR diagram is unlikely to arise by chance. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [72].

stopping criterion:

1. At each stage of prewhitening, we fit the (amplitude) spectrum as described above, and obtain the RPS before selecting a frequency to prewhiten. We note that this is the *opposite* procedure adopted by [29, 27, 28], who prewhiten coherent frequencies before fitting the amplitude spectra to characterize the stochastic background. This is because the power of the stochastic variability is much higher than that of the peaks in the periodogram, especially at low frequency. However, because we are not removing the stochastic variability from the lightcurve itself, this only helps us locate the peaks in the RPS. The frequencies, amplitudes, and phases we obtain are otherwise identical to what we would obtain following [29, 27, 28].

2. Chapter 5 identified multiple harmonics of some recovered frequencies. To properly treat potential harmonics, at each stage of prewhitening, we fit both the selected frequency and

the amplitude and phase of its first two harmonics. There are some instances where we select a frequency that is itself a harmonic of another lower amplitude frequency, and so the fundamental is not removed by the harmonic fit. We note all instances when this occurs below.

3. In addition to saving the best-fit parameters of each sinusoid, we calculate the associated errors on each parameter, using the formulae given in [169] and [201]:

$$\epsilon(f_j) = \sqrt{\frac{6}{N}}\frac{1}{\pi T}\frac{\sigma_j}{A_j} \tag{6.2}$$

$$\epsilon(A_j) = \sqrt{\frac{2}{N}}\sigma_j \tag{6.3}$$

$$\epsilon(\phi_j) = \sqrt{\frac{2}{N}}\frac{\sigma_j}{A_j} \tag{6.4}$$

$$\tag{6.5}$$

where N is the number of points in the lightcurve, T is the time baseline of the lightcurve, f_j, A_j, and ϕ_j are the frequency, amplitude, and phase extracted at the j^{th} prewhitening stage, and σ_j is the standard deviation of the flux at the same stage. We also record the value of the RPS at the selected frequency, and the SNR, calculated as the peak height divided by the standard deviation of the RPS in a narrow window between $f_{max} \pm 2f_R$ and $f_{max} \pm 7f_R$.

4. As a stopping criterion, we proceed until we reach a minimum in the Bayesian Information Content (BIC, [257]) of the fit:

$$BIC = -2\ln(\mathcal{L}) + m\ln(N) \tag{6.6}$$

$$-2\ln(\mathcal{L}) = \sum_{i=1}^{N}\frac{(y_i - F(t_i, \Theta_m))^2}{\sigma_i^2} \tag{6.7}$$

where $\mathcal{L}$ is the likelihood (to within a constant), m is the number of free parameters in the fit ($7j + 7$ at the j^{th} stage of prewhitening, beginning with $j = 0$), y_i are the original fluxes observed at times t_i, $F(t_i, \Theta_m)$ is the sum of all of the fit sinusoids evaluated at fit parameters Θ_m, and σ_i are the normalized errors in the original lightcurve [234, §15.1].

As a postprocessing step, we discard frequencies that are quite similar to each other (i.e., the difference in frequencies is within $1.5f_R$, keeping the earliest frequency found). These similar and spurious frequencies can arise due to the short length of the observing baseline [167]. The unique frequencies, amplitudes, and phases, corresponding formal errors, RPS peak heights, and RPS SNRs found for each star are listed in Appendix A.3. The frequencies extracted from the FYPS are all between ~ 0.5 and ~ 4.6 d^{-1} with semi-amplitudes ranging between ~ 40 and ~ 280 ppm.

From the final list of extracted frequencies, we search for harmonics of the form $f_j/f_i = n$, where n is an integer greater than 1, that satisfy

$$nf_i - f_j \leq \sqrt{\left(n\epsilon(f_i)\right)^2 + \left(\epsilon(f_j)\right)^2} \tag{6.8}$$

i.e., f_j is an exact integer multiple of f_i to within the errors, and the k^{th} harmonic corresponds to $n = k + 1$ (e.g., the first harmonic is $n = 2$). We also search for frequency combinations in the form $f_i + f_j = f_k$, such that

$$f_1 + f_j - f_k \leq \sqrt{\left(\epsilon(f_i)\right)^2 + \left(\epsilon(f_j)\right)^2 + \left(\epsilon(f_k)\right)^2}. \tag{6.9}$$

Wavelet Analysis

In addition to the frequencies extracted from the entire lightcurve, we also attempt to determine whether the frequencies and amplitudes are stable. To that end, we employ a time-frequency analysis to search for variability in the dominant frequencies. We calculate the Weighted Wavelet Z-transform [WWZ, 80], an extension of wavelet analysis with the Morlet wavelet:

$$\psi(t, \tau, \omega) = e^{i\omega(t-\tau) - c\omega^2(t-\tau)^2} \tag{6.10}$$

where here τ is the time of the center of the wavelet, $\omega = 2\pi f$ is the angular frequency, and c sets the width of the Gaussian envelope and is chosen to be sufficiently small so that the wavelet decays appreciably over the course of one cycle. Here we adopt $c = 0.0125$, following [80]; in principle, smaller or larger values of c can be chosen to alter the time and frequency resolution, which is frequency dependent. The discrete wavelet transform can be converted into a projection onto the continuous basis functions

$$\Phi_1(t,\omega,\tau) = 1 \tag{6.11}$$

$$\Phi_2(t,\omega,\tau) = \cos\omega(t-\tau) \tag{6.12}$$

$$\Phi_3(t,\omega,\tau) = \sin\omega(t-\tau) \tag{6.13}$$

and we now fold the Gaussian envelope into a weighting function for each data point at time t_i that depends on the frequency and time center of the wavelet:

$$w_i(\tau,\omega) = e^{-c\omega^2(t_i-\tau)^2} \tag{6.14}$$

This change ensures that a wavelet centered on a gap in the data won't pick up small amplitude random fluctuations on either side of the gap, thereby suppressing false power that can often arise in wavelet transformations of unevenly sampled data [e.g. 289, 288]. At each τ and ω, we calculate the number of *effective* data points,

$$N_{eff}(\tau,\omega) = \frac{\left(\sum_i w_i(\tau,\omega)\right)^2}{\left(\sum_i w_i^2(\tau,\omega)\right)} \tag{6.15}$$

and the weighted variance of the flux (which we denote x):

$$V_x(\tau,\omega) = \langle x|x\rangle - \langle \Phi_1|x\rangle^2 \tag{6.16}$$

where the inner product of two functions, $\langle f|g\rangle$ is defined as

$$\langle f|g\rangle = \frac{\sum_i w_i(\tau,\omega) f(t_i) g(t_i)}{\sum_j w_j(\tau,\omega)} \tag{6.17}$$

We also calculate the weighted variance of a sinusoidal fit to the model:

$$V_y(\tau,\omega) = \langle y|y\rangle - \langle \Phi_1|y\rangle^2 \tag{6.18}$$

defining

$$y = \vec{y}_\Phi \cdot \vec{\Phi} \tag{6.19}$$

where $\vec{\Phi}$ is a vector containing the basis functions, Φ. $\vec{y}_\Phi$ is a vector containing the coefficients of the projection onto the basis functions,

$$\vec{y}_\Phi = S^{-1}\vec{y}_b \tag{6.20}$$

with the entries of the matrix S equal to $S_{ab} = \langle \Phi_a|\Phi_b\rangle$ and entries in the vector $\vec{y}_b$ are $y_b = \langle \Phi_b|x\rangle$. With these ingredients, we can calculate the WWZ at each τ and ω as

$$WWZ = \frac{(N_{eff} - 3)V_y}{2(V_x - V_y)} \tag{6.21}$$

Finally, we set $WWZ = 0$ if $min(t_i - \tau) > 2\pi/\omega$ (i.e., the nearest data point is more than one cycle away from the center of the wavelet), to reduce computation cost.

HD 269953

HD 269953 is the brightest FYPS discovered with $\log L/L_\odot = 5.437$. We initially analyzed it in Chapter 5, and suggested that it was a post-RSG, largely due to its infrared excess hinting at past mass loss. It was previously studied by [306], who noted its variability. Figure 6.8 shows the RPS of HD 269953, with the four unique frequencies recovered by prewhitening in grey. After search-

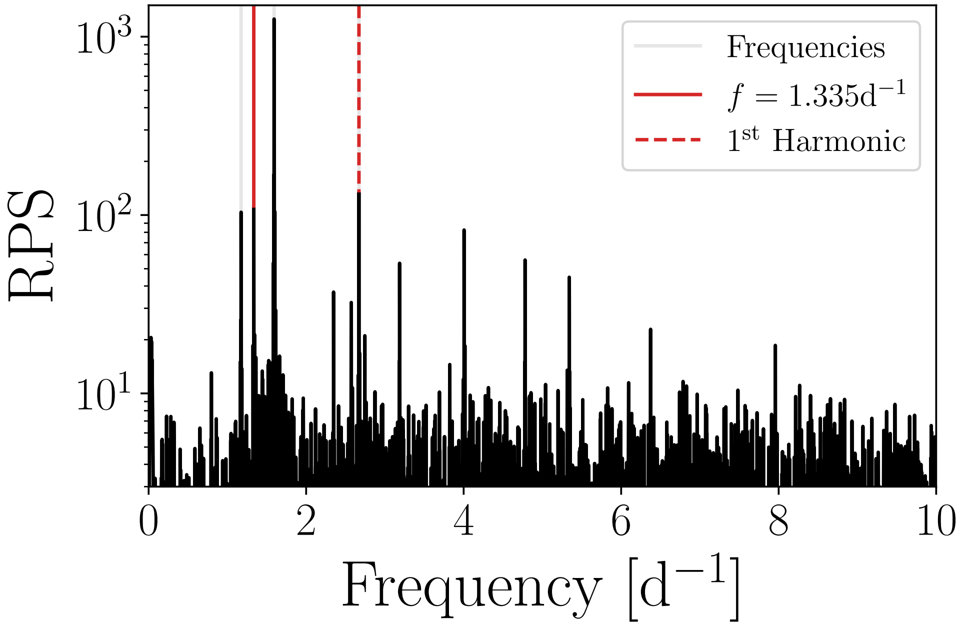

Figure 6.8 RPS for HD 269953, with frequencies extracted by prewhitening indicated by vertical grey lines. Red lines correspond to $f = 1.335$ d^{-1} and its harmonic. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [72].

ing for harmonics, we find one instance where the fundamental at 1.335 d^{-1} has a lower amplitude than the first harmonic at 2.671 d^{-1} and so both frequencies are recovered. No combination frequencies are recovered, implying the presence of three independent frequencies in the lightcurve of HD 269953. The presence of these non-aliased frequencies may indicate the presence of multiple oscillation modes in HD 269953. Alternately, the aperture may contain two different pulsating stars. Without a better model of YSG pulsations, we cannot conclusively determine which scenario is more likely; however, if the pulsation mechanism of FYPS is identical to that of α Cygni variables, the presence of multiple modes may not be surprising [131].

Figure 6.9 shows the lightcuve (top, black), rolling median of the flux (top, green), RPS (right), and WWZ (center) calculated on 500 linearly spaced time points, and 1000 frequency points on a $\log_2$ grid between $\log_2(2\pi f) = -1$ ($f \approx 0.08$ d^{-1}) and 5.5 ($f \approx 7.2$ d^{-1}). Frequencies extracted via prewhitening are shown as horizontal white (WWZ) or grey (RPS) lines. At frequencies below

1 d^{-1}, the WWZ shows transient events that are associated with times where systematics in the data appear to be present — e.g. the discontinuity after the mid-sector downlink at $\mathrm{Time} \approx 1620$ days. However, the wavelet map demonstrates that these transient events have no effect on the highest amplitude frequencies in the RPS. At higher frequencies, the frequency of maximum power in the WWZ rapidly changes as a function of time, appearing to oscillate between the three lowest extracted frequencies. There even appear to be times when the peak in the WWZ almost disappears (e.g., around $\mathrm{Time} = 1400$; note the gap at $\mathrm{Time} = 1500$ is due to a gap in the data). This indicates that the detected pulsations are not stable, with amplitudes changing on timescales of days. Perhaps the modes are stochastically excited and damped on these timescales. YSG interiors are complex, with multiple boundaries between convective and radiative zones, which may be responsible for driving the pulsations. However, with no reliable models of YSG pulsations, we can only speculate at this time. Of the seven higher frequency peaks detected, five are harmonics of lower-frequency signals as discussed. Some brief, low amplitude transient events are associated with the peaks in the RPS corresponding to harmonics of lower frequency signals. Unfortunately, a drawback of the WWZ (and most time-frequency analyses in general) is that potential interesting high-frequency features are smeared out in exchange for increased time resolution.

HD 269110

HD 269110 was also discussed in Chapter 5, and had the lowest frequency signals detected there. It has not been analyzed by any other modern variability studies. With our updated prewhitening scheme, we extract four unique frequencies, including three incredibly close to each other: a main peak with the highest SNR at $f = 0.553$ d^{-1}, and two small peaks, each separated from the main peak by $\Delta f = \pm 0.011$ d^{-1} (~ 10 times the resolution of the periodogram). The peak in the RPS corresponding to the first harmonic of this frequency shows similar structure, though the higher frequency subpeak is very low signal to noise. The lightcurve, RPS, and WWZ are shown in Figure 6.10. Only the three closely spaced peaks can be seen in the WWZ, and those peaks appear to fade and reappear semi-regularly. While the modulation in the WWZ is not sinusoidal, extracting the WWZ in a 0.2 d^{-1} band around the closest frequency to the main peak and calculating the power

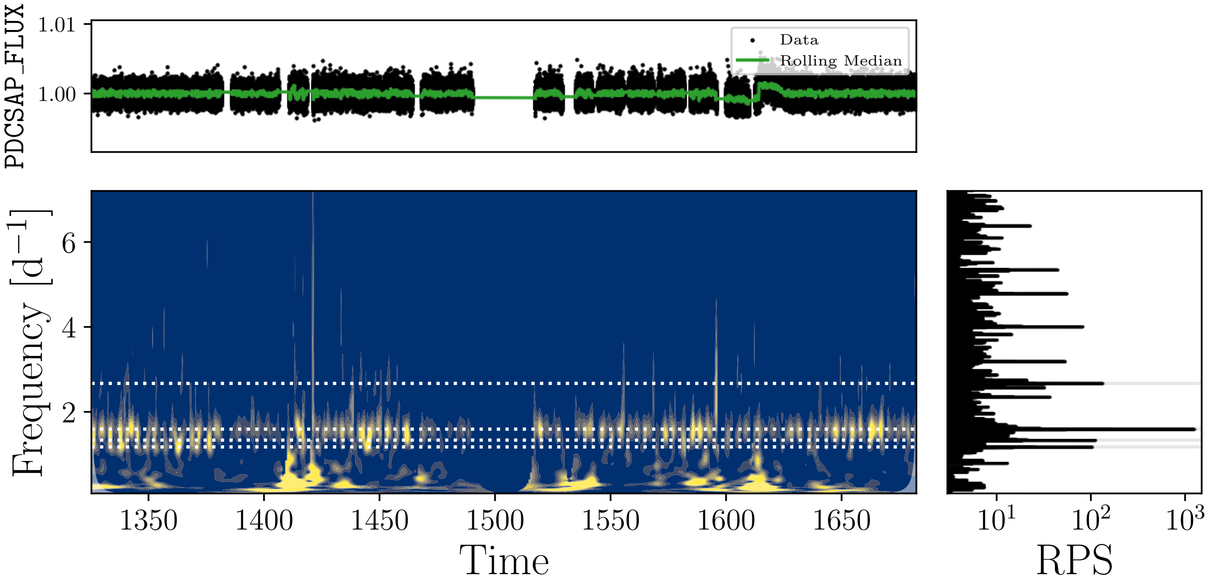

Figure 6.9 (*Top*): *TESS* lightcurve of HD 269953 in black points, with green indicating the 128-cadence rolling median. (*Right*): RPS, with frequencies identified by prewhitening in grey. (*Center*): Weighted Wavelet Z-transform (WWZ, as defined in text) of the lightcurve as a function of time and frequency. Higher values of the WWZ are shown in yellow, and lower values in blue. Identified frequencies are shown as horizontal white lines, and region of the WWZ where the center of the wavelet is within one cycle ($1/f$) of the beginning and end of the data are shaded in white. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [72].

spectrum reveals a strong peak at $f = 0.010$ d^{-1}, quite similar to the frequency difference between the three peaks recovered by prewhitening. One possibility is that HD 269110 is a binary system. Pulsating stars in binaries exhibit frequency modulation similar to what we see [e.g. 262], and eccentric close binaries can induce modulations in pulsation amplitudes on the orbital timescale [see 293, for a lower-mass example], also in line with what we have detected. An alternate hypothesis is that the $f = 0.553$ d^{-1} is split by rotational effects [§II.B.3 3], and 0.011 d^{-1} is the rotational frequency.

If HD 269110 is a binary system with a 20 $M_\odot$ YSG primary and a $1/0.011 \approx 91$ d orbital period, a companion star would have to be $\sim 100 M_\odot$ in order for the semimajor axis of the orbit to be larger than the stellar radius derived from the parameters listed in Table 6.3. Such an object would have to be a black hole in order to not be significantly brighter than the primary, in which case it would still be more massive than any known stellar mass black hole, a scenario we deem to be incredibly unlikely. Alternately, rotational modulation of the WWZ on a ~ 90 d timescale requires invoking a misalignment of the pulsational and rotational axes. Furthermore, a 91 day rotation period is either quite close to, or exceeds the critical rotation periods of YSGs in the Geneva models, depending on the mass. Regardless, the similarity between Δf and the characteristic timescale extracted from the WWZ is intriguing, and warrants follow-up observations. We discuss the implications on the evolutionary status of FYPS below.

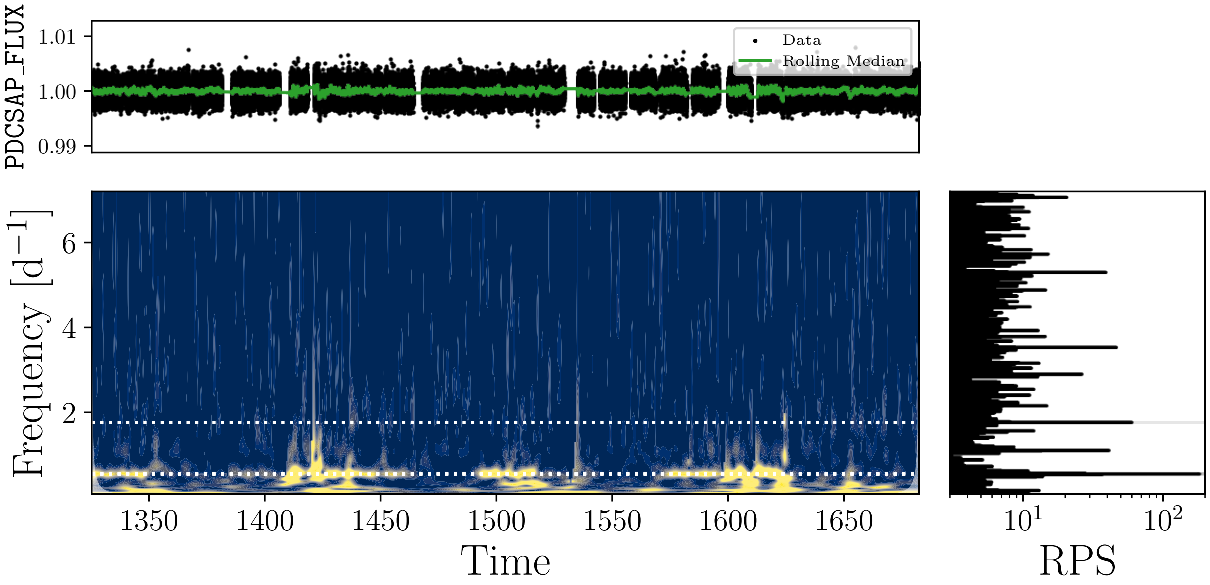

Figure 6.10 Similar to Figure 6.9 for HD 269110. Note that the tallest peak in the RPS is the triplet of split frequencies at $f = 0.553$ d^{-1}. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [72].

HD 268687, HD 269840, & HD 269902

HD 268687 was the third candidate pulsating YSG found in Chapter 5. However, the most prominent frequencies we had found previously were superimposed upon the part of the power spectrum dominated by low frequency stochastic variability. Apart from Chapter 5, it was identified as a

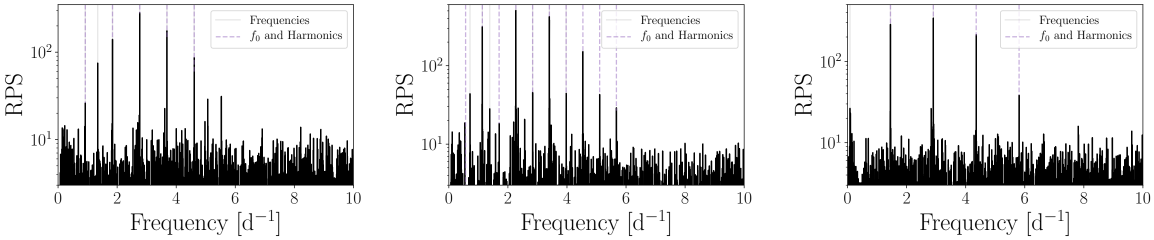

Figure 6.11 RPS for HD 268687 (left), HD 269840 (center), and HD 269902 (right). Light grey lines show frequencies extracted by prewhitening, and light purple lines show the inferred fundamental, f_0, and harmonics, as discussed in text. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [72].

variable star in OGLE photometry [299]. With prewhitening, we extract six unique frequencies. The RPS and extracted frequencies are shown in the left panel of Figure 6.11. The RPS displays a broad comb of peaks, all but one of which correspond to periods faster than one day. None of the low-frequency peaks identified in Chapter 5 are recovered; however, the peaks previously identified as harmonics of the (now nondetected) dominant peak *are* recovered. From this, we conclude that the dominant peak identified in early *TESS* data without correcting for the SLFV may have belonged to the overall pattern of peaks, even though it is no longer detected. Searching for harmonics in the recovered frequencies reveals one frequency at 1.844 d^{-1} with a detected first harmonics at 3.687 d^{-1}. A search for combination frequencies shows that the sums of the lowest frequency peak and both the fundamental and first harmonic are also recovered frequencies. Interestingly, the four recovered frequencies in ascending order starting with 1.844 d^{-1} are roughly equally spaced, with a frequency difference of $\Delta f = 0.922$ d^{-1}.

Pulsations in both HD 269840 and HD 269902 are newly discovered; the former was identified as a variable in OGLE photometry [299]. Similar to HD 268687, the RPS of both stars show a broad frequency comb (center and right panels of Figure 6.11) with a total of seven and three unique frequencies extracted from HD 269840 and HD 269902 respectively. We do not find any remaining harmonics or combination frequencies amongst the frequencies extracted from both stars. As in HD 268687, the four frequencies recovered in HD 269840 above 2.270 d^{-1} are regularly spaced with a spacing $\Delta f \approx 0.567$ d^{-1} (though the spacing between the highest two frequencies is 0.001

d^{-1} higher). No such regular spacing is found in HD 269902.

Motivated by the apparent regular spacing of peaks in the RPS of all three stars, we searched for evidence that each of the three FYPS exhibits a harmonic chain of peaks. In all three panels of Figure 6.11, we assume that one of the extracted frequencies is a fundamental frequency and plot its first few harmonics. In the case of HD 269840, we instead assume that the frequency of the tallest peak at $f = 2.270$ d^{-1} is four times the frequency of the fundamental, which lines up with a peak in the RPS that is not selected by our prewhitening procedure. In all three cases, we find an exceptional match between the assumed harmonic chain and most of the peaks in the RPS. However, the actual observed frequencies are inconsistent with regular harmonic patterns to within the errors

This fact could be due to one of two causes. The offsets from an even spacing pattern could be caused by, e.g., structural glitches [§IV.B 3]. Such glitches can be used to assess sharp features in the stellar structure that would otherwise be inaccessible by other means. Alternately, we may have underestimated the errors for the extracted frequency. Instrumental correlations exist in data taken by the *Kepler* and CoRoT missions, and can add to the uncertainty in extracted frequencies [258]. As the theory of pulsations in YSGs is still nascent, we have no asteroseismic model for these stars, and thus cannot determine whether the observed offset from an even spacing pattern is astrophysical. However, we can determine the extent to which the uncertainty in the extracted frequencies may be underestimated, assuming that the frequencies *should* be evenly spaced.

For each star, we calculate the difference between the observed frequencies and the closest frequency in an evenly-spaced grid extrapolated from the fundamental frequencies assumed above, Δf, ignoring the observed frequency if the difference between it and the closest predicted frequency is more than 0.005 d^{-1}. We also calculated the associated uncertainties, $\sigma_{\Delta f}$ by adding the errors of the observed frequencies and the assumed fundamental frequency in quadrature. Finally, for the entire collection of Δf measurements for all three stars, we calculate the reduced chi-squared,

$$\chi^2_{red} = \frac{1}{N} \sum \frac{\Delta f}{D\sigma_{\Delta f}} \tag{6.22}$$

where D is a correction factor to account for underestimated errors, and N is the total number of frequencies in the lightcurves of all three stars. If any scatter in the values of Δf around 0 is driven by measurement error, we can find the value of D for which $\chi^2_{red} \approx 1$ — i.e., the extent to which our errors are underestimated. Doing so yields a value of $D \approx 3$, consistent with typical values of D for similar space-based photometric observations [typically ∼2-10, 258]; of course, this requires that we have correctly identified the right fundamental frequency, that the frequency spacing is indeed a regular pattern, and that these offsets are not *actually* due to astrophysical effects.

Ultimately, until systematic correlations in *TESS* data are better quantified and we are able to generate accurate asteroseismic models of FYPS, we will not be able to make a concrete determination of the cause of the offset between the observed frequencies and a regularly spaced harmonic series. The wavelet analysis used above may be able to help diagnose the behavior of the observed frequencies. Unfortunately, HD 268687 has a measured SLFV amplitude of $\alpha = 0.064 \pm 0.002$, the highest of all of the discovered FYPS, and the WWZ of the lightcurve is entirely dominated by low-frequency transient features associated with this stochastic variability. The values of α are smaller for HD 269840 and HD 269902. We plot the WWZs of both stars in Figure 6.12.

In HD 269902, most of the power in the WWZ is contained in the lower frequencies where the stochastic low frequency background dominates — unsurprisingly given that the star has the second-largest value of α out of the five FYPS (0.024 ± 0.001). Only by altering the scaling of the image of the WWZ are we able to see the ridge associated with the lowest frequency peak, and it is only detected with low signal-to-noise. In both stars, the the lowest frequency peak in the RPS is associated with a broad band of power in the WWZ, similar to the WWZ of HD 269953, with a secondary band appearing at $\sim 2-3$ d^{-1} in HD 269840, at the approximate location of the highest peak in the RPS.

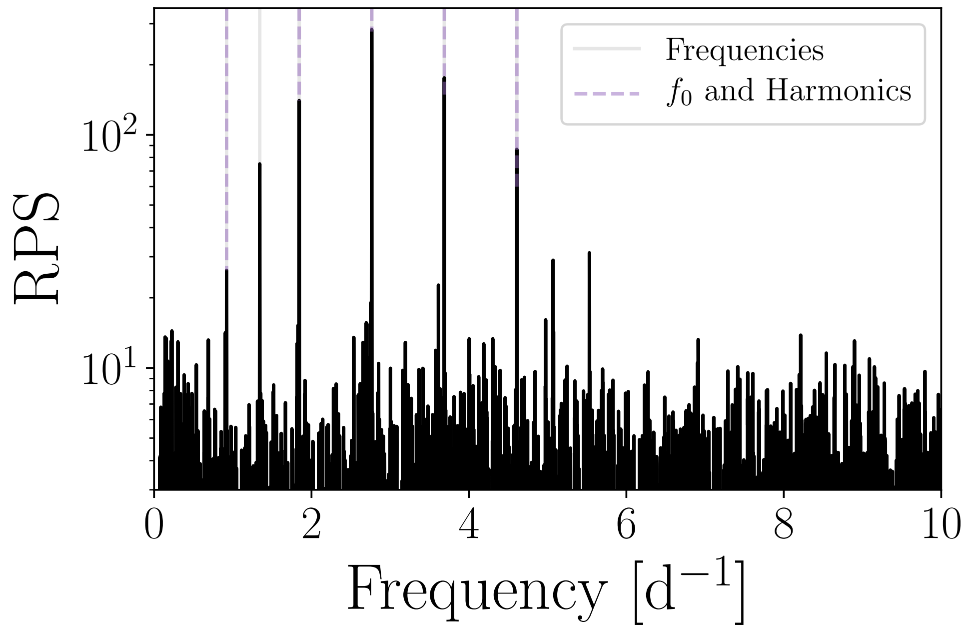

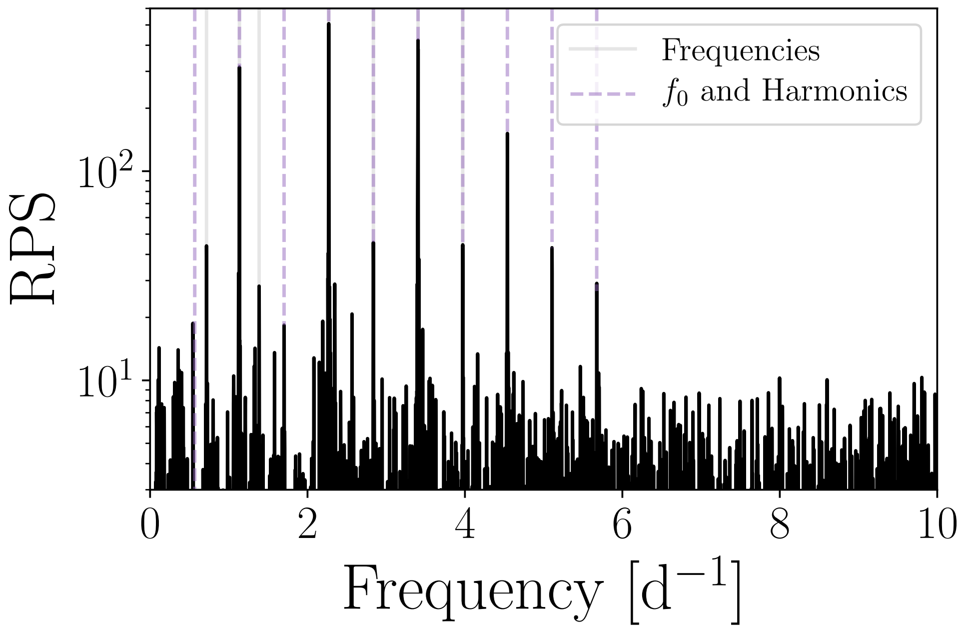

Figure 6.12 Similar to Figure 6.8 for HD 269840 (top) and HD 269902 (bottom). Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [72].

6.4 Discussion

6.4.1 Asteroseismic Modeling

The discovery of fast pulsations in YSGs has very interesting implications for the study of the interiors of evolved massive stars. The 25 $M_\odot$ solar-metallicity Geneva model [76] has a main sequence lifetime of approximately 7 Myr. It then crosses the HR diagram in under 1 Myr, and half a Myr later, has evolved bluewards once more to become a Wolf-Rayet star [190]. Due to the incredibly short lifetime of YSGs on both crossings of the HR diagram, theoretical uncertainty has long stymied our understanding of massive star evolution [137]. Pulsation frequencies extracted from long-baseline lightcurves assembled from space-based observations are perhaps the most precise measurements we can make; typical values of $f/\epsilon(f)$ for frequencies listed in Tables A.10-A.14 are $\sim 10^{-4} - 10^{-5}$. In better-studied stars, such precision allows for the diagnosis of incredibly complicated physics, and is truly the benchmark of testing stellar evolution theory [3]. No asteroseismic models for YSGs that reliably converge exist [122], and so the era of precision YSG asteroseismology is not yet upon us. However, we can compare the frequencies observed in the FYPS with the characteristic Lamb and Brunt-Väisälä angular frequencies in a model YSG:

$$S_\ell^2 = \frac{\ell(\ell+1)c_s^2}{r^2} \tag{6.23}$$

$$N^2 = \frac{g}{H_p}[\delta(\nabla_{ad} - \nabla) + \phi\nabla_\mu] \tag{6.24}$$

where ℓ is the harmonic degree, c_s is the sound speed, r is the radius of the local shell, g is the gravity, and the remaining variables are defined as in [5]. Modes with angular frequencies $|\omega| > |N|$ and $|\omega| > |S_\ell|$ are mostly restored by pressure (*p*-modes), and modes with $|\omega| < |N|$ and $|\omega| < |S_\ell|$ are mostly restored by buoyancy (*g*-modes). While rudimentary, such a comparison would illustrate the approximate regions of the stellar structure the observed pulsations will allow us to probe.

To do this, we use version 12778 of Modules for Experiments in Stellar Astrophysics (MESA, [223, 224, 225, 226, 227]) to evolve a grid of nonrotating stellar models with initial masses between

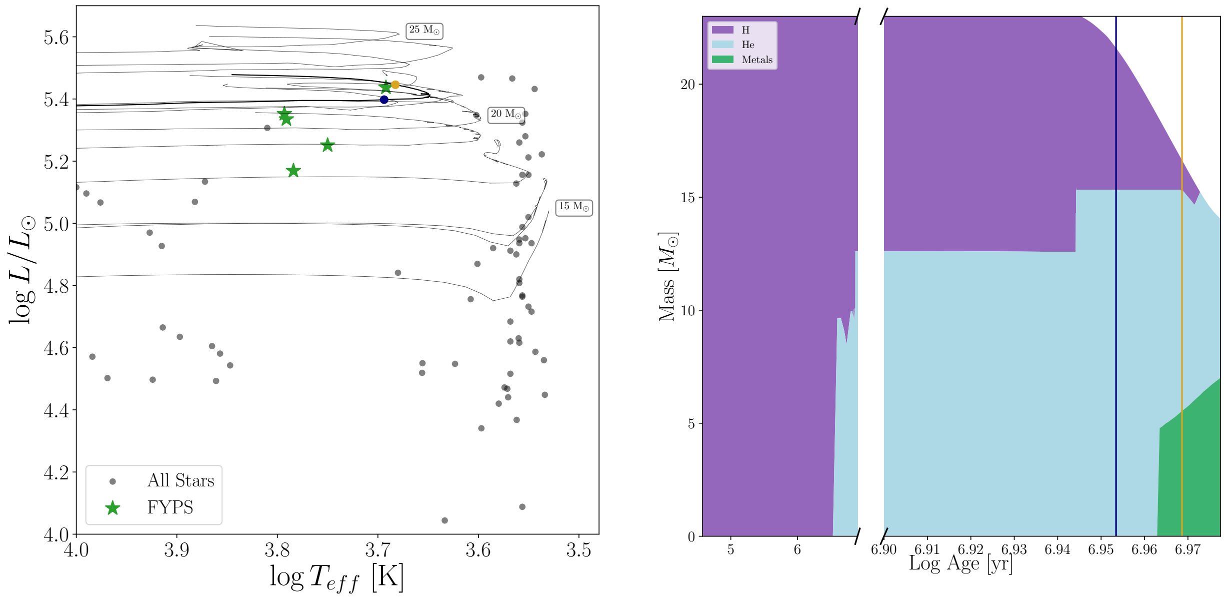

Figure 6.13 *Left*: HR diagram with non-FYPS plotted as black points, and FYPS plotted as green stars. $Z = 0.006$ evolutionary tracks calculated with MESA are shown as solid grey lines with their initial mass indicated. The 23 $M_\odot$ model is shown as a solid black line. The navy point is the pre-RSG timestep in this model closest to the position of HD 269953 in the HR diagram. The goldenrod point is the closest post-RSG timestep in the same evolutionary track. *Right*: Evolution of the structure of the 23 $M_\odot$ MESA model as a function of time. The purple/light blue/green regions are the parts of the star (in mass coordinates) dominated by H/He/metals, respectively. The time axis is broken to illustrate the rapid post-main sequence evolution, including the strong mass loss in the star's final stages. The vertical navy line shows the age of the closest pre-RSG timestep to HD 269953 in the HR diagram. The vertical goldenrod shows the age of the closest post-RSG timestep. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [72].

15 and 30 $M_\odot$, with a spacing of 1 $M_\odot$, initial metallicity set to $Z = 0.006$, and initial helium mass fraction set to $Y = 0.25 + 1.5Z = 0.259$. We first evolve the models through the entire pre-main sequence stage until the entirety of the star's luminosity comes from nuclear burning then introduce more elaborate physics. Convective mixing follows the `mlt++` prescription from [224], with overshoot following [78]. For mass loss, we use the "Dutch" [95] and "Vink" [315] cool/hot wind schemes respectively, and adopt an efficiency of $\eta = 0.8$ for the former [e.g. 175]. Because the luminosity in the outer layers of red supergiants can exceed the Eddington luminosity due to high opacity in these layers, we follow [76], and increase the mass-loss rate in by a factor of 3 when the luminosity exceeds five times the Eddington luminosity. Finally, we use Type 2 opacities for when the star has extra C/O during and after He burning. With these controls, all but the 28 $M_\odot$ model successfully ran; as we focus the remainder of our analysis on a single model, we chose not to introduce additional controls for this one model. Our inlist files, including our timestep and spatial resolution controls, are available online at `https://github.com/tzdwi/TESS`.

The post-ZAMS evolutionary tracks are plotted in the left panel of Figure 6.13, with the 15, 20, 25, and 30 $M_\odot$ tracks labeled. We note that evolutionary modeling of post main-sequence massive stars is fraught with uncertainty; different prescriptions for overshooting, or mass loss [e.g. 181], or different choices of input physics like binary interactions or rotation [e.g. 70] can radically alter the evolutionary pathway a given stellar model might take. Nonetheless, we can use these evolutionary tracks to find a model that may approximate the structure and evolution of the FYPS. Interestingly, the stars with initial mass $M \geq 18 M_\odot$ lose enough mass to begin turning around on the HR diagram. Models more massive than 19 $M_\odot$ become luminous and warm enough to encounter the yellow void, at which point the models begin to exhibit rapid changes in their luminosities and effective temperatures on ~month to year timescales — for clarity, we do not show the post-RSG portion of the tracks after they reach an effective temperature hotter than 7000 K. The stars in our sample are plotted as grey points, and the FYPS are plotted as green stars. The dark line shows the model that passes the closest to the position of HD 269953 in the HR diagram, which has an initial mass of 23 $M_\odot$. We show the evolution of the interior structure of this model throughout its lifetime in the right panel of Figure 6.13. Each colored region shows the part of the star (in mass coordinates) dominated by H (purple), He (light blue), and metals (green) respectively as a function of the age of the star.

The navy point and navy line in both panels of Figure 6.13 indicate the pre-RSG timestep whose temperature and luminosity best match the observed values for HD 269953. At this point, the model is ~ 9.0 Myr old, has a current mass of $21.6 M_{\odot}$, and has $\log T_{\rm eff} = 3.694$, and $\log L/L_{\odot} = 5.398$. By this time, the star has begun core He fusion, has created $\sim 0.8 M_{\odot}$ of C, and is losing significant mass from its envelope. Similarly, the goldenrod point and vertical line in Figure 6.13 correspond to the post-RSG timestep in the same model that is closest to HD 269953. This model is 9.3 Myr old, has undergone extensive mass loss as an RSG, and has $\log T_{\rm eff} = 3.683$ and $\log L/L_{\odot} = 5.446$. It is still fusing He in its core, but has built up significant C and O mass. It's current mass is only 16.6 $M_{\odot}$

We show the interior structures of both models as a function of mass coordinate in the top panels of Figure 6.14; the pre-RSG model is on the left, the post-RSG model is on the right. The density profiles, normalized by the central density ρ_c, are shown in dark yellow, while the composition is shown as the profiles of X, Y, and Z with identical colors as the right panel of Figure 6.13. We calculate the Lamb and Brunt-Väisälä frequencies in each model. The bottom panels of Figure 6.14 show the logarithm of N^2 (blue) and S_1^2 (orange). The frequencies detected in the *TESS* lightcurve of HD 269953 are shown as horizontal black dashed lines. Two regions exist within the pre-RSG model where *g*-modes are able to propagate (shaded in blue), while *p*-modes are able to propagate throughout the envelope of the star (shaded in orange). The inner *g*-mode cavity corresponds to the region outside of the core that experienced previous H burning, while the outer cavity cavity corresponds to a chemically-stratified outer envelope. In the post-RSG model, the innermost *g*-mode cavity has moved outward, and combined with the outer cavity. Sharp features in the stellar structure can be seen in both characteristic frequencies.[8]

We stress that this is *not* an asteroseismic analysis of the star; we have made no attempt at

---

[8] We note that the sharp steps in the composition profiles of the envelopes of both models (and the corresponding spiky behavior of N^2 in this region) are due to small discontinuities in the resolution of the MESA model, and are not real features.

predicting the excited frequencies in either model, and are by no means identifying the observed frequencies with p- or g-modes, let alone more complicated phenomena such as strange modes. Furthermore, the exact treatment of mixing (including semiconvection and thermohaline mixing, which we do not include in our simple models) can have an incredibly strong infuence on the pre-supernova structure of the star [e.g. 78]. Indeed, with no reliable interior models of YSGs, we cannot identify modes in order to conduct a full asteroseismic analysis, which may allow us to constrain these physics as well as the evolutionary status of FYPS. However, this rudimentary comparison does illuminate the regions of the stellar structure that the observed pulsations might probe, as well as the drastically different interior structures and pulsational properties seen in pre- and post-RSG models that reside in quite similar regions of the HR diagram. One possible step towards mode identification is to see whether the highest amplitude pulsation frequencies scale with the observed parameters of the stars to ascertain whether the pulsations may be in the acoustic, gravity, or gravito-intertial regime. We searched for correlations between the strongest observed frequencies and $\log T_{eff}$, $\log L/L_{\odot}$, and $\log R^{-2}$ (as a proxy for $\log g$), but a sample of only five FYPS is insufficient to find any obvious trends. Futhermore, the available spectra are of insufficient resolution to measure $v \sin i$, which are typically less than 10 km s $^{-1}$ in YSGs [21]. With more FYPS, and higher resolution spectroscopy we may be able to conclusively determine the origin of these pulsations. Again, we emphasize that this demonstration serves only to motivate future work on these stars, and show the potential of FYPS for asteroseismology of YSGs.

6.4.2 Evolutionary Status: Leftward vs. Rightward Evolution, and the Red Supergiant Problem

Examining Figure 6.13, the FYPS appear to be descended from stars with initial masses of $M_{FYPS} \gtrsim 20 M_{\odot}$. As discussed above, the maximum observed luminosity of Supernovae II-P progenitors is significantly lower than the maximum luminosity of field RSGs (the red supergiant problem), implying that stars with an initial mass above $M_{max} \approx 20 M_{\odot}$ don't end their lives as red supergiants that explode [140]. The coincidence between M_{max} and M_{FYPS} is consistent with perhaps the most natural solution to the red supergiant problem: high mass RSGs aren't found as supernova progenitors because they do not explode as RSGs. They instead evolve bluewards on the HR di-

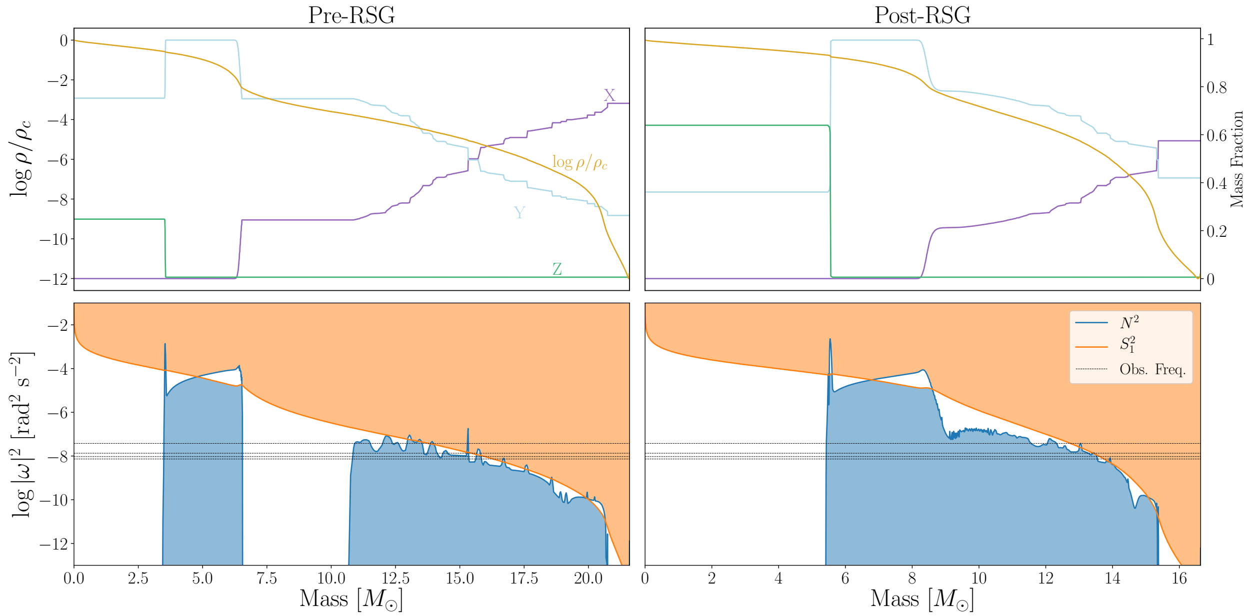

Figure 6.14 *Top*: Interior structure of the closest pre-RSG (left) and post-RSG (right) 28 $M_\odot$ MESA models to HD 269953. The dark yellow line shows the density profile, while the purple, light blue, and green lines correspond to the profiles of X, Y, and Z, respectively, as in the right panel of Figure 6.13. *Bottom*: Calculated values of the log of the square of the Brunt-Väisälä (blue line) and Lamb (orange line) frequencies as a function of the mass coordinate within both MESA models. Here we only show S_1 for simplicity. Regions where $\omega < |N|$ and $|\omega| < |S_1|$ (i.e., where *g*-modes of a given frequency can propagate) are shaded in blue, and regions where $\omega > |N|$ and $|\omega| > |S_1|$ (i.e., where *p*-modes of a given frequency can propagate) are shaded in orange. Frequencies extracted from the *TESS* lightcurve of HD 269953 are shown as horizontal black dashed lines. Copyright AAS. Reproduced with permission from Dorn-Wallenstein et al. [72].

agram after losing significant mass as RSGs, attaining high enough values of L/M to excite the rich spectrum of observed pulsation modes as predicted in the post-RSG models of [249]; FYPS are thus post-RSG objects.[9]

While this hypothesis offers a tantalizing solution, putting it on more solid footing would re-

---

[9] Note that the MESA models more massive than 18 $M_\odot$ lose enough mass as RSGs to evolve bluewards in the HR diagram. This is not seen in the low-metallicity Geneva models [91, 155]. If FYPS are genuine post-RSGs, this would require an increase in the mass loss rates used in standard stellar models. RSG mass loss rates have only been directly measured in very small numbers of stars in the Galaxy [e.g. 192], and in even fewer LMC RSGs [e.g. 308]. RSG mass loss is also episodic, and while constant mass loss prescriptions can reproduce the makeup of RSG populations [213], they are still only an approximation to the actual mass loss histories of individual objects.

quire additional evidence that FYPS are indeed in a post-RSG phase. One route is to search for evidence of strong past mass loss. The TIC contains photometry from the Wide Field Infrared Exlorer (WISE, [324]), in four mid-infrared bands: $W1$ (3.4 μm), $W2$ (4.6 μm), $W3$ (12 μm), and $W4$ (22 μm). For most stars, the WISE bands are in the Rayleigh-Jeans tail of their spectral energy distribution (SED), and they therefore have colors ~ 0 in Vega passbands [59]. Deviations from 0 may be attributable to infrared molecular bands, especially in the spectra of red supergiants (Dicenzo & Levesque, in prep). However, they may also be indicative of circumstellar dust.

We inspected the WISE colors of all stars in our sample. Only the brightest and coolest FYPS, HD 269953, has significantly larger values in each WISE color than nearby stars on the HR diagram. The remaining FYPS are mostly indistinguishable from their fellow stars, though HD 269840 also has a slightly larger $W3 - W4$ value than other nearby stars. The three FYPS discussed in §6.3.2 have somewhat *smaller* values of $W3 - W4$, more consistent with the RSGs than the YSGs. While suggestive that the two most luminous FYPS may have some warm dust contributing to their SEDs, this is not concrete evidence that all five FYPS are in a post-RSG phase. Of course, these mid-infrared measurements are only sensitive to warm dust located relatively close to the star. Perhaps HD 269953 is a less-evolved FYPS with a warm and close CSM, while the CSM of the remaining FYPS have cooled and are undetectable. Probing cooler dust further from the stellar surface would require higher resolution imaging at longer wavelengths [e.g. 261]. If the mechanism powering the pulsations in FYPS also operates at higher metallicity, it would be possible to find a more nearby FYPS in our Galaxy in order to perform such observations.

It is also important to note that the circumstellar dust produced by RSGs is known to be larger grain than that found in the interstellar medium, and that incorrectly accounting for "gray" extinction from this dust can lead to underestimates of L in RSGs (e.g. [187, 260, 109, 156]). If the same is true for the circumstellar dust found around post-RSG YSGs (produced during the RSG phase), it is possible that this could lead to underestimating the L of these stars, which would in turn lead to underestimating stellar mass when comparing the stars to evolutionary tracks on the HR diagram. Quantifying this requires careful observations of these stars' circumstellar environments and dust properties.

One concrete counterargument to the post-RSG hypothesis is found in HD 269110. If HD 269110 is a binary, and the pulsation amplitudes are orbitally modulated, such a binary system would have to be relatively close and eccentric, implying that the two stars have not previously exchanged significant mass. If the primary had previously experienced strong mass-loss as an RSG, the orbit would most likely have circularized instead. However, as discussed above, a companion star with a 91 day orbit would have to be $\sim 100 M_\odot$ in order for the semimajor axis of the orbit to be larger than the stellar radius. As rotation is an equally unlikely culprit due to the low surface gravity/critical rotation speeds typical of YSGs, a more thorough characterization of HD 269110 is required.

Ultimately, it is hard to draw any conclusions from such a small sample of pulsators. A total of 341 YSGs identified from [116], [271], [313], [103] — all in the Milky Way — [214] — SMC — and [215] have entries in the TIC. Assuming our sample of YSGs observed by *TESS* is representative of the entire population, $\sim 5/27 = 19\%$ of YSGs are FYPS, implying that *TESS* can detect ~ 63 FYPS. Even if *TESS* observes half as many FYPS, we would be able to better characterize the boundaries of the region in the HR diagram in which they reside, and whether the presence/behavior of FYPS is dependent on metallicity.

6.5 Conclusions

Our main results are summarized as follows:

- We study the *TESS* lightcurves of 76 cool supergiants with accurate temperatures and luminosities. For YSGs located in the Large Magellanic Cloud, these lightcurves span a time baseline of a year. We discover that low-frequency stochastic variability is ubiquitous in these stars, and rule out surface convection as the underlying cause for all but the red supergiants. This implies that this variability, also observed in main sequence O stars, is a constant feature of massive stars throughout their lifetimes.

- After removing the contribution of this background variability from the periodograms of the stars in our sample, we find two regions in the HR diagram with pulsating stars. Four of these stars are candidate α Cygni variables, of which three are newly identified as such. The remaining five pulsating stars are clustered in a region of the HR diagram not previously identified as a region of instability.

- We rule out binarity, spurious signals, and chance alignment with nearby stars causing us to mistakenly find these pulsators in the same part of the HR diagram by chance, and conclude that these five stars comprise a real class, that we dub Fast Yellow Pulsating Supergiants (FYPS).
- We extract pulsation frequencies from the FYPS lightcurve using a new procedure to account for the stochastic background, search for harmonics and frequency combinations from the extracted frequencies, and calculate the Weighted Wavelet Z-transform (WWZ) of the lightcurves to study the time-dependent behavior of these frequencies
- In HD 269953 and HD 269840 (and perhaps HD 269902), the lowest frequency strong periodogram peak extracted is coincident with a broad ridge of power in the WWZ that rapidly switches frequencies and amplitudes on $\sim$day timescales.
- HD 268687, HD 269840, and HD 269902 show a broad comb of frequencies, with most extracted peaks being close to but not exactly in a harmonic chain.
- One of the frequencies found in HD 269110 is split into a triplet, and the WWZ of the lightcurve at that frequency is modulated on the same timescale as the difference between peaks in the triplet. We are unable to determine whether or not this is due to rotation, binary effects, or some other cause.
- We introduce the possibility that FYPS are post-RSG objects that have lost enough mass as RSGs to attain high luminosity-to-mass ratios, and excite pulsations [as in 249]. This possibility is bolstered by the coincidence between the lowest estimated mass of the FYPS and the highest mass RSG progenitors. While we are unable to concretely determine the exact

evolutionary status of the FYPS, future work to determine whether or not they are genuine post-RSG objects is of extreme importance, both theoretically by modeling their pulsation frequencies, and observationally by finding more FYPS. Regardless of the evolutionary status of FYPS, their pulsational properties will be of great use in the study of the interiors of evolved massive stars.

Chapter 7

CONCLUSION

Having reached the end,[1] it is worthwhile to now take a look back, see what we've learned, and maybe even think a bit about what's next. After all, much of this work has knocked on doors that haven't yet been opened for massive stars. If you'll allow me, I'd like to indulge in some speculation of what the future may bring.

I view this work as being neatly separable into three discrete chunks: Chapters 2 and 3 concern themselves with modeling massive stars using approaches based in stellar population synthesis; Chapter 4 is focused on applying ideas of machine learning to a field where "big data" has, until recently, meant a few hundred stars; and Chapters 5 and 6 showcase the power of studying the high frequency variability of massive stars. I'll now discuss each of these topics in turn, before ending everything with what I hope to be a satisfying summary.

7.1 Studying Populations of Massive Stars

To be a stellar theorist must be quite difficult. Crafting stellar models requires not only the acknowledgement but the embracing of countless unknowns. This includes both known unknowns — for example, the mixing length efficiency or the mass loss rates — and the unknown unknowns that won't reveal themselves until observers like me come along and ruin the collective days of theorists everywhere.

Many of these quantities can be constrained by careful observations of massive stars. Asteroseismology has tremendous potential as a tool for probing the interior structures of stars, and constraining evolutionary models. But as I've said once or twice already, massive stars are rare. In order to gather large enough samples, we must turn to the massive star content of entire galax-

---

[1] Congratulations, by the way! As much fun as this has been to write, it must have been quite something to read.

ies. Beyond individual bright objects in the Local Group, however, the careful observations that guide us in our Milky Way aren't feasible, and so we must study the ensemble properties of stellar populations.

In Chapter 2, I introduced a method of estimating an unknown quantity, the massive star binary fraction f_{bin}, using the demographics of a stellar population, while accounting for the fact that our samples of massive stars are typically incomplete. In Chapter 3, I extended this analysis to stellar models that incorporate the effects of rotation, and developed a statistical framework that ensures the measurements we are making are the same quantities that are predicted by theoretical models. This work demonstrated that the models are remarkably consistent with the data, leading to the intriguing possiblity that we could actually measure the massive star binary fraction in distant galaxies. I also found tentative evidence that the massive star binary fraction increases with metallicity, a fact that has tremendous implications on both the formation and evolution of massive stars. Indeed, recent work [210] has found direct evidence of this f_{bin}-Z relationship, albeit only in red supergiants. Current work that uses a much improved sample of massive stars [189] has also used the population synthesis method I have developed to find that this result is robust to a variable, and unknown star formation history.

Of course, this work only implements a somewhat crude, *post-hoc* method of varying the binary fraction in the models. This is largely because the population synthesis itself — calculating weights and creating the synthetic populations — is computationally expensive. However, in theory, much of the computational cost can be done before weighing each model, with the only remaining step being the final summation over evolutionary tracks. In this case, provided weights could be generated quickly, the stellar populations could be synthesized "on the fly" using simple linear algebra in order to forward model the observed data. Such methods would allow for not only the comparison of models to data, but the actual *inference* of the binary fraction using Bayesian techniques. And why stop at f_{bin}? Throughout the analysis, I assumed the distributions of the initial periods and mass ratios from [198], but if the binary fraction changes with metallicity, why shouldn't these quantities as well? With the high quality data of massive stars that will be observed by *Webb*, *Roman*, and LSST at the Vera Rubin observatory, the time to solve these computational issues is now.

7.2 *Massive Stars in the Era of Big Data*

For quite some time, the era of "big data" seemed to be a distant future for most of the massive star community. Our samples are just too small, our methods of observation too inhomogeneous, our error bars too large. And so many of our studies relied upon intense observational campaigns focused on individual stars, or on "surveys" of dozens of poorly-characterized objects. But the times, they are a-changin'.

In Chapter 4, I demonstrated that data from the *Gaia* mission constrains the absolute magnitudes and colors of stars in our Galaxy and the Magellanic Clouds with fantastic precision, allowing us to construct huge samples of massive stars, with no need for narrowband surveys, or spectroscopic follow up. So are we ready to take out our fancy machine learning techniques, turn the crank, and go? Much like most rhetorical questions in science, the answer is no. We are severely limited by the existing quality of our labels: identifications of a given star as a red supergiant, or luminous blue variable, or whatever class of massive star you happen to care about, are just too unreliable to give us the results we want. This is an urgent problem given the imminent launch of the *James Webb Space Telescope*[2] and the planned 2025 launch of the *Nancy Grace Roman Space Telescope*, which will observe luminous stars 10s of Megaparsecs from our Galaxy. With limited ability to obtain follow-up observations of these stars from the ground, we're going to *need* machine learning techniques.

Thankfully, the data necessary to classify the stars that we can use in future training samples *already exists*: ground-based spectroscopic surveys like the APO Galaxy Evolution Experiment (APOGEE) and the Large Sky Area Multi-Object Fibre Spectroscopic Telescope survey (LAM-

---

[2] It is important to note that before and during his tenure as NASA administrator, James Webb oversaw an era of persecution, harassment, and purging of queer people in government service known as the "Lavender Scare;" evidence indicates that Webb was personally involved in the creation and enforcement of the policies leading to this period. If the stated purpose of the telescope named after Webb is to "reveal the unknown for the benefit of all humankind," we must ask ourselves how *Webb* can inspire us — all of us — in such a way, while the legacy of Webb himself is one of persecution? At the time of this writing, NASA is currently investigating Webb's record. I believe this is not enough, and I urge my readers to urge NASA to rename the telescope after a more worthy figure.

OST) are measuring temperatures, surface gravities, and compositions for millions of stars. Unfortunately, the focus on their spectroscopic pipelines — which themselves use machine learning — is on low mass dwarfs and giants. However, with additional information available on Galactic massive stars from *Gaia* , it is time to invest our resources into machine learning and massive stars.

7.3 Towards an Asteroseismic Future

If you're anything like me from a few years ago, then maybe "asteroseismology" sounds like magic, a whispered incantation to summon measurements of seemingly impossible quantities; I mean come on, measuring how fast a star's core is rotating using just single-band photometry? You've gotta be joking. But think about it like this: the frequencies listed in §A.3 are measured with a precision $f/\epsilon(f)$ of about 10^5. That's like measuring the temperature of Betelgeuse to the nearest tenth of a Kelvin, when in reality, our measurements of stellar temperatures have errors on the order of a few tens to a couple hundred Kelvin! And these frequencies are incredibly sensitive to the conditions inside the star, in a way that effective temperatures, surface gravities, and even surface abundances just *aren't*. And so it's no surprise that missions like *Kepler* and *TESS* are changing the game when it comes to stellar astrophysics.

Unsurprisingly, many of the advancements enabled by asteroseismology have been for low mass stars. Part of this was an observational problem: massive stars — especially evolved massive stars — are rare, and their average density on the sky is low. This is made even worse by the fact that the massive stars observable by small telescopes (like *Kepler* and *TESS*) cluster together, either in the Galactic plane or in the Milky Way's satellite galaxies. This means that missions with small fields of view who won't be named[3] just *didn't* observe any evolved massive stars — at least not during "Schmepler's" original 4-year mission.

But *TESS* , with its large field of view and all-sky observing strategy, is fixing this problem. In Chapter 5 I showed that even with single sectors of observations, TESS revealed both stochastic and coherent variability in yellow supergiants and luminous blue variables. In Chapter 6, I showed

---

[3] Ok, one of their names rhymes with "Schmepler."

that with an entire year of data for stars in the continuous viewing zones, we can measure the frequencies of coherent modes with incredible precision, while constraining stochastic low frequency variability. Indeed, this SLFV seems tied to a star's location in the HR diagram (Figure 6.4), and therefore its mass and evolutionary status.

The remaining stumbling block for us is now one of theory. As discussed in depth in the introduction, low mass stellar models are generally in agreement, but models for the post-main sequence evolution of massive stars diverge rapidly from one another. Even worse, recent efforts to model the stability of evolved massive stars to pulsations have found difficulties getting models to converge in the part of the HR diagram where asteroseismic analysis could have the greatest impact (i.e., the region occupied by fast yellow pulsating supergiants). While the existence of FYPS has tremendous potential for unravelling the mysteries of massive star evolution, this will only be possible if stellar theorists turn their attention to the pulsations of massive stars crossing the HR diagram.

7.4 In Summary...

...massive stars are cool![4] They are tremendously important objects given their impact on their surroundings, as well as their status as complex astrophysical laboratories in which our theories of stellar evolution can truly be put to the test. In order to truly understand these fantastic objects, we must take a broad scope, using every tool in our collective toolbox. Some unknown quantities (like the binary fraction, or the distribution of natal rotation rates, etc.) require the observation and modelling of large populations of massive stars. In particular, simple demographic information of young stellar populations can be incredibly insightful (Chapters 2 and 3). However, these methods need large numbers of stars with well-measured properties (e.g., their rough evolutionary stage, or even better, temperatures and gravities). Making these measurements at scale is only going to be possible with machine learning techniques, a prospect for which we are currently unprepared (Chapter 4). Finally, some of our studies of massive stars are going to have to stay close to home,

---

[4] Or hot. Whatever.

bringing the full might of asteroseismology to bear, and providing us with truly unprecedented glimpses into the interiors of massive stars (Chapter 5 and 6). But this can only be done with models that don't currently exist.

If there is one guiding theme to this work, it is the underlying push and pull that exists between new technologies — techniques or telescopes or models — that have allowed us to take enormous strides in the last 5-10 years, and the resulting mess when these steps have taken us into new lands where our current knowledge and abilities are no longer sufficient. It is my sincerest hope that this thesis has outlined this tension and described a few necessary steps that we as a community must take to prepare to study massive stars in the 2020s and beyond. I also hope you've enjoyed my work! It is the result of six years of intense thought and research; twenty-four years of formal education; and countless late nights, lines of code, cups of coffee, housing disasters, personal drama, joy, sadness, and everything in between. And after all that, I can't wait to see what's next. Thank you for reading.

BIBLIOGRAPHY

[1] E. Aadland, P. Massey, K. F. Neugent, and M. R. Drout. Shedding Light on the Isolation of Luminous Blue Variables. *AJ*, 156:294, December 2018.

[2] P. Abolmasov. Stochastic variability of luminous blue variables. *NewA*, 16:421–429, November 2011.

[3] C. Aerts. Probing the interior physics of stars through asteroseismology. *Reviews of Modern Physics*, 93(1):015001, January 2021.

[4] C. Aerts, S. Simón-Díaz, C. Catala, C. Neiner, M. Briquet, N. Castro, V. S. Schmid, M. Scardia, M. Rainer, E. Poretti, P. I. Pápics, P. Degroote, S. Bloemen, R. H. Østensen, M. Auvergne, A. Baglin, F. Baudin, E. Michel, and R. Samadi. Low-amplitude rotational modulation rather than pulsations in the CoRoT B-type supergiant HD 46769. *A&A*, 557:A114, September 2013.

[5] Conny Aerts, Jørgen Christensen-Dalsgaard, and Donald W. Kurtz. *Asteroseismology*. 2010.

[6] M. Andersen, M. Gennaro, W. Brandner, A. Stolte, G. de Marchi, M. R. Meyer, and H. Zinnecker. Very low-mass stellar content of the young supermassive Galactic star cluster Westerlund 1. *A&A*, 602:A22, June 2017.

[7] R. I. Anderson, H. Saio, S. Ekström, C. Georgy, and G. Meynet. On the effect of rotation on populations of classical Cepheids. II. Pulsation analysis for metallicities 0.014, 0.006, and 0.002. *A&A*, 591:A8, June 2016.

[8] H. Andrews, D. Fenech, R. K. Prinja, J. S. Clark, and L. Hindson. Asymmetric ejecta of cool supergiants and hypergiants in the massive cluster Westerlund 1. *MNRAS*, 477:L55–L59, June 2018.

[9] A. Ardeberg, J. P. Brunet, E. Maurice, and L. Prevot. Spectrographic and photometric observations of supergiants and foreground stars, in the direction of the Large Magellanic Cloud. *Astronomy and Astrophysics Supplement Series*, 6:249–309, July 1972.

[10] A. Arellano Ferro. Periodicity and pulsational mode of five bright yellow supergiants. *MNRAS*, 216:571–587, October 1985.

[11] T. E. Armandroff and P. Massey. Wolf-Rayet stars in NGC 6822 and IC 1613. *ApJ*, 291:685–692, April 1985.

[12] M. Asplund, A. M. Amarsi, and N. Grevesse. The chemical make-up of the Sun: A 2020 vision. *arXiv e-prints*, page arXiv:2105.01661, May 2021.

[13] M. Asplund, N. Grevesse, A. J. Sauval, and P. Scott. The Chemical Composition of the Sun. *ARA&A*, 47:481–522, September 2009.

[14] Astropy Collaboration, T. P. Robitaille, E. J. Tollerud, P. Greenfield, M. Droettboom, E. Bray, T. Aldcroft, M. Davis, A. Ginsburg, A. M. Price-Whelan, W. E. Kerzendorf, A. Conley, N. Crighton, K. Barbary, D. Muna, H. Ferguson, F. Grollier, M. M. Parikh, P. H. Nair, H. M. Unther, C. Deil, J. Woillez, S. Conseil, R. Kramer, J. E. H. Turner, L. Singer, R. Fox, B. A. Weaver, V. Zabalza, Z. I. Edwards, K. Azalee Bostroem, D. J. Burke, A. R. Casey, S. M. Crawford, N. Dencheva, J. Ely, T. Jenness, K. Labrie, P. L. Lim, F. Pierfederici, A. Pontzen, A. Ptak, B. Refsdal, M. Servillat, and O. Streicher. Astropy: A community Python package for astronomy. *A&A*, 558:A33, October 2013.

[15] Yu Bai, JiFeng Liu, Song Wang, and Fan Yang. Machine Learning Applied to Star-Galaxy-QSO Classification and Stellar Effective Temperature Regression. *AJ*, 157(1):9, Jan 2019.

[16] C. A. L. Bailer-Jones, J. Rybizki, M. Fouesneau, G. Mantelet, and R. Andrae. Estimating Distance from Parallaxes. IV. Distances to 1.33 Billion Stars in Gaia Data Release 2. *AJ*, 156:58, August 2018.

[17] L. A. Balona. Rotational modulation in B stars observed by the Kepler K2 mission. *MNRAS*, 457:3724–3731, April 2016.

[18] L. A. Balona, A. S. Baran, J. Daszyńska-Daszkiewicz, and P. De Cat. Analysis of Kepler B stars: rotational modulation and Maia variables. *MNRAS*, 451:1445–1459, August 2015.

[19] L. A. Balona, G. Handler, S. Chowdhury, D. Ozuyar, C. A. Engelbrecht, G. M. Mirouh, G. A. Wade, A. David-Uraz, and M. Cantiello. Rotational modulation in TESS B stars. *MNRAS*, 485(3):3457–3469, May 2019.

[20] R. V. Baluev. Assessing the statistical significance of periodogram peaks. *MNRAS*, 385:1279–1285, April 2008.

[21] B. Barbuy, J. R. de Medeiros, and A. Maeder. Carbon, nitrogen and oxygen abundances in yellow supergiants. *A&A*, 305:911, January 1996.

[22] Krzysztof Belczynski, Aleksander Sadowski, and Frederic A. Rasio. A Comprehensive Study of Young Black Hole Populations. *ApJ*, 611(2):1068–1079, Aug 2004.

[23] Willem Janszoon Blaeu. Himmelsglob. *Skokloster Castle*, 1602.

[24] R. Blomme, L. Mahy, C. Catala, J. Cuypers, E. Gosset, M. Godart, J. Montalban, P. Ventura, G. Rauw, T. Morel, P. Degroote, C. Aerts, A. Noels, E. Michel, F. Baudin, A. Baglin, M. Auvergne, and R. Samadi. Variability in the CoRoT photometry of three hot O-type stars. HD 46223, HD 46150, and HD 46966. *A&A*, 533:A4, September 2011.

[25] E. Böhm-Vitense. Über die Wasserstoffkonvektionszone in Sternen verschiedener Effektivtemperaturen und Leuchtkräfte. Mit 5 Textabbildungen. *ZAP*, 46:108, January 1958.

[26] A. Z. Bonanos, D. L. Massa, M. Sewilo, D. J. Lennon, N. Panagia, L. J. Smith, M. Meixner, B. L. Babler, S. Bracker, M. R. Meade, K. D. Gordon, J. L. Hora, R. Indebetouw, and B. A. Whitney. Spitzer SAGE Infrared Photometry of Massive Stars in the Large Magellanic Cloud. *AJ*, 138:1003–1021, October 2009.

[27] D. M. Bowman, C. Aerts, C. Johnston, M. G. Pedersen, T. M. Rogers, P. V. F. Edelmann, S. Simón-Díaz, T. Van Reeth, B. Buysschaert, A. Tkachenko, and S. A. Triana. Photometric detection of internal gravity waves in upper main-sequence stars. I. Methodology and application to CoRoT targets. *A&A*, 621:A135, January 2019.

[28] D. M. Bowman, S. Burssens, S. Simón-Díaz, P. V. F. Edelmann, T. M. Rogers, L. Horst, F. K. Roepke, and C. Aerts. Photometric detection of internal gravity waves in upper main-sequence stars. II. Combined TESS photometry and high-resolution spectroscopy. *arXiv e-prints*, page arXiv:2006.03012, June 2020.

[29] Dominic M. Bowman, Siemen Burssens, May G. Pedersen, Cole Johnston, Conny Aerts, Bram Buysschaert, Mathias Michielsen, Andrew Tkachenko, Tamara M. Rogers, Philipp V. F. Edelmann, Rathish P. Ratnasingam, Sergio Simón-Díaz, Norberto Castro, Ehsan Moravveji, Benjamin J. S. Pope, Timothy R. White, and Peter De Cat. Low-frequency gravity waves in blue supergiants revealed by high-precision space photometry. *Nature Astronomy*, 3:760–765, May 2019.

[30] Leo Breiman. Random forests. *Mach. Learn.*, 45(1):5–32, October 2001.

[31] J. Brinchmann, M. Pettini, and S. Charlot. New insights into the stellar content and physical conditions of star-forming galaxies at z = 2-3 from spectral modelling. *MNRAS*, 385:769–782, April 2008.

[32] B. Buysschaert, C. Aerts, S. Bloemen, J. Debosscher, C. Neiner, M. Briquet, J. Vos, P. I. Pápics, R. Manick, V. S. Schmid, H. Van Winckel, and A. Tkachenko. Kepler's first view of O-star variability: K2 data of five O stars in Campaign 0 as a proof of concept for O-star asteroseismology. *MNRAS*, 453:89–100, October 2015.

[33] M. Cantiello and N. Langer. Thermohaline mixing in evolved low-mass stars. *A&A*, 521:A9, October 2010.

[34] J. I. Castor, D. C. Abbott, and R. I. Klein. Radiation-driven winds in Of stars. *ApJ*, 195:157–174, January 1975.

[35] B. Chaboyer and J. P. Zahn. Effect of horizontal turbulent diffusion on transport by meridional circulation. *A&A*, 253:173–177, January 1992.

[36] G. Chabrier. Galactic Stellar and Substellar Initial Mass Function. *PASP*, 115:763–795, July 2003.

[37] Nitesh V. Chawla. *Data Mining for Imbalanced Datasets: An Overview*, pages 875–886. Springer US, Boston, MA, 2010.

[38] A. N. Chené, A. F. J. Moffat, C. Cameron, R. Fahed, R. C. Gamen, L. Lefèvre, J. F. Rowe, N. St-louis, V. Muntean, A. De La Chevrotière, D. B. Guenther, R. Kuschnig, J. M. Matthews, S. M. Rucinski, D. Sasselov, and W. W. Weiss. WR 110: A Single Wolf-Rayet Star with Corotating Interaction Regions in its Wind? *ApJ*, 735:34, July 2011.

[39] C. Chiappini, R. Hirschi, G. Meynet, S. Ekström, A. Maeder, and F. Matteucci. A strong case for fast stellar rotation at very low metallicities. *A&A*, 449(2):L27–L30, Apr 2006.

[40] Nancy Chinchor and Ph D. Muc-4 evaluation metrics. In *In Proceedings of the Fourth Message Understanding Conference*, pages 22–29, 1992.

[41] Cesare Chiosi and Andre Maeder. The evolution of massive stars with mass loss. *ARA&A*, 24:329–375, January 1986.

[42] J. Choi, A. Dotter, C. Conroy, M. Cantiello, B. Paxton, and B. D. Johnson. Mesa Isochrones and Stellar Tracks (MIST). I. Solar-scaled Models. *ApJ*, 823:102, June 2016.

[43] Sang-Hyun Chun, Sung-Chul Yoon, Moo-Keon Jung, Dong Uk Kim, and Jihoon Kim. Evolutionary Models of Red Supergiants: Evidence for A Metallicity-dependent Mixing Length and Implications for Type IIP Supernova Progenitors. *ApJ*, 853(1):79, January 2018.

[44] J. S. Clark, I. Negueruela, P. A. Crowther, and S. P. Goodwin. On the massive stellar population of the super star cluster ¡ASTROBJ¿Westerlund 1¡/ASTROBJ¿. *A&A*, 434:949–969, May 2005.

[45] Jacob Cohen. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46, 1960.

[46] C. Conroy, J. Strader, P. van Dokkum, A. E. Dolphin, D. R. Weisz, J. W. Murphy, A. Dotter, B. D. Johnson, and P. Cargile. A Complete Census of Luminous Stellar Variability on Day to Decade Timescales. *ApJ*, 864:111, September 2018.

[47] P. S. Conti. On the relationship between Of and WR stars. *Memoires of the Societe Royale des Sciences de Liege*, 9:193–212, January 1975.

[48] P. S. Conti. Basic Observational Constraints on the Evolution of Massive Stars. In A. Maeder and A. Renzini, editors, *Observational Tests of the Stellar Evolution Theory*, volume 105 of *IAU Symposium*, page 233, 1984.

[49] P. S. Conti, C. D. Garmany, C. De Loore, and D. Vanbeveren. The evolution of massive stars - The numbers and distribution of O stars and Wolf-Rayet stars. *ApJ*, 274:302–312, November 1983.

[50] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, Sep 1995.

[51] P. A. Crowther, L. J. Hadfield, J. S. Clark, I. Negueruela, and W. D. Vacca. A census of the Wolf-Rayet content in Westerlund 1 from near-infrared imaging and spectroscopy. *MNRAS*, 372:1407–1424, November 2006.

[52] P. A. Crowther and L. J. Smith. Fundamental parameters of Wolf-Rayet stars. VI. Large Magellanic Cloud WNL stars. *A&A*, 320:500–524, April 1997.

[53] Xiang-Qun Cui, Yong-Heng Zhao, Yao-Quan Chu, Guo-Ping Li, Qi Li, Li-Ping Zhang, Hong-Jun Su, Zheng-Qiu Yao, Ya-Nan Wang, Xiao-Zheng Xing, Xin-Nan Li, Yong-Tian Zhu, Gang Wang, Bo-Zhong Gu, A. Li Luo, Xin-Qi Xu, Zhen-Chao Zhang, Gen-Rong Liu, Hao-Tong Zhang, De-Hua Yang, Shu-Yun Cao, Hai-Yuan Chen, Jian-Jun Chen, Kun-Xin Chen, Ying Chen, Jia-Ru Chu, Lei Feng, Xue-Fei Gong, Yong-Hui Hou, Hong-Zhuan Hu, Ning-Sheng Hu, Zhong-Wen Hu, Lei Jia, Fang-Hua Jiang, Xiang Jiang, Zi-Bo Jiang, Ge Jin, Ai-Hua Li, Yan Li, Ye-Ping Li, Guan-Qun Liu, Zhi-Gang Liu, Wen-Zhi Lu, Yin-Dun Mao, Li Men, Yong-Jun Qi, Zhao-Xiang Qi, Huo-Ming Shi, Zheng-Hong Tang, Qing-Sheng Tao, Da-Qi Wang, Dan Wang, Guo-Min Wang, Hai Wang, Jia-Ning Wang, Jian Wang, Jian-Ling Wang, Jian-Ping Wang, Lei Wang, Shu-Qing Wang, You Wang, Yue-Fei Wang, Ling-Zhe Xu, Yan Xu, Shi-Hai Yang, Yong Yu, Hui Yuan, Xiang-Yan Yuan, Chao Zhai, Jing Zhang, Yan-Xia Zhang, Yong Zhang, Ming Zhao, Fang Zhou, Guo-Hua Zhou, Jie Zhu, and Si-Cheng Zou. The Large Sky Area Multi-Object Fiber Spectroscopic Telescope (LAMOST). *Research in Astronomy and Astrophysics*, 12(9):1197–1242, September 2012.

[54] Thayne Currie, Jesus Hernandez, Jonathan Irwin, Scott J. Kenyon, Susan Tokarz, Zoltan Balog, Ann Bragg, Perry Berlind, and Mike Calkins. The Stellar Population of h and χ Persei: Cluster Properties, Membership, and the Intrinsic Colors and Temperatures of Stars. *ApJS*, 186(2):191–221, Feb 2010.

[55] R. M. Cutri, M. F. Skrutskie, S. van Dyk, C. A. Beichman, J. M. Carpenter, T. Chester, L. Cambresy, T. Evans, J. Fowler, J. Gizis, E. Howard, J. Huchra, T. Jarrett, E. L. Kopan, J. D. Kirkpatrick, R. M. Light, K. A. Marsh, H. McCallon, S. Schneider, R. Stiening, M. Sykes, M. Weinberg, W. A. Wheaton, S. Wheelock, and N. Zacarias. VizieR Online Data Catalog: 2MASS All-Sky Catalog of Point Sources (Cutri+ 2003). *VizieR Online Data Catalog*, 2246, June 2003.

[56] R. M. Cutri, E. L. Wright, T. Conrow, J. W. Fowler, P. R. M. Eisenhardt, C. Grillmair, J. D. Kirkpatrick, F. Masci, H. L. McCallon, S. L. Wheelock, S. Fajardo-Acosta, L. Yan, D. Benford, M. Harbut, T. Jarrett, S. Lake, D. Leisawitz, M. E. Ressler, S. A. Stanford, C. W.

Tsai, F. Liu, G. Helou, A. Mainzer, D. Gettings, A. Gonzalez, D. Hoffman, K. A. Marsh, D. Padgett, M. F. Skrutskie, R. P. Beck, M. Papin, and M. Wittman. Explanatory Supplement to the AllWISE Data Release Products. Explanatory Supplement to the AllWISE Data Release Products, November 2013.

[57] J. Daszynska-Daszkiewicz, P. Walczak, W. Szewczuk, and A. Pamyatnykh. What have we learnt from pulsations of B-type stars? *ArXiv e-prints*, November 2018.

[58] J. R. A. Davenport and K. R. Covey. Rotating Stars from Kepler Observed with Gaia DR2. *ApJ*, 868:151, December 2018.

[59] James R. A. Davenport, Željko Ivezić, Andrew C. Becker, John J. Ruan, Nicholas M. Hunt-Walker, Kevin R. Covey, Alexia R. Lewis, Yusra AlSayyad, and Lauren M. Anderson. The SDSS-2MASS-WISE 10-dimensional stellar colour locus. *MNRAS*, 440(4):3430–3438, Jun 2014.

[60] Ben Davies and Emma R. Beasor. The 'Red Supergiant Problem': the upper luminosity boundary of type-II supernova progenitors. *MNRAS*, January 2020.

[61] Carl. De Boor. *A practical guide to splines / Carl de Boor*. Springer-Verlag New York, 1978.

[62] C. de Jager, H. Nieuwenhuijzen, and K. A. van der Hucht. Mass loss rates in the Hertzsprung-Russell diagram. *A&AS*, 72:259–289, Feb 1988.

[63] Cornelis de Jager. The yellow hypergiants. *A&A Rv*, 8(3):145–180, January 1998.

[64] C. De Loore, J. P. de Grève, and D. Vanbeveren. Parameters of Massive Stars of Main Sequence Evolution without and with Stellar Wind. *A&AS*, 34:363, December 1978.

[65] S. E. de Mink, N. Langer, R. G. Izzard, H. Sana, and A. de Koter. The Rotation Rates of Massive Stars: The Role of Binary Interaction through Tides, Mass Transfer, and Mergers. *ApJ*, 764:166, February 2013.

[66] S. E. de Mink, O. R. Pols, and S.-C. Yoon. Binaries at Low Metallicity: Ranges For Case A, B and C Mass Transfer. In B. W. O'Shea and A. Heger, editors, *First Stars III*, volume 990 of *American Institute of Physics Conference Series*, pages 230–232, March 2008.

[67] S. E. de Mink, H. Sana, N. Langer, R. G. Izzard, and F. R. N. Schneider. The Incidence of Stellar Mergers and Mass Gainers among Massive Stars. *ApJ*, 782:7, February 2014.

[68] T. Z. Dorn-Wallenstein and E. M. Levesque. Stellar Population Diagnostics of the Massive Star Binary Fraction. *ArXiv e-prints*, October 2018.

[69] Trevor Z. Dorn-Wallenstein, James R. A. Davenport, Daniela Huppenkothen, and Emily M. Levesque. Photometric Classifications of Evolved Massive Stars: Preparing for the Era of Webb and Roman with Machine Learning. *arXiv e-prints*, page arXiv:2102.02829, February 2021.

[70] Trevor Z. Dorn-Wallenstein and Emily M. Levesque. A Comparison of Rotating and Binary Stellar Evolution Models: Effects on Massive Star Populations. *arXiv e-prints*, page arXiv:2004.13040, April 2020.

[71] Trevor Z. Dorn-Wallenstein, Emily M. Levesque, and James R. A. Davenport. Short-term Variability of Evolved Massive Stars with TESS. *ApJ*, 878(2):155, Jun 2019.

[72] Trevor Z. Dorn-Wallenstein, Emily M. Levesque, Kathryn F. Neugent, James R. A. Davenport, Brett M. Morris, and Keyan Gootkin. Short-term Variability of Evolved Massive Stars with TESS. II. A New Class of Cool, Pulsating Supergiants. *ApJ*, 902(1):24, October 2020.

[73] A. Dotter. MESA Isochrones and Stellar Tracks (MIST) 0: Methods for the Construction of Stellar Isochrones. *ApJS*, 222:8, January 2016.

[74] S. M. Dougherty and J. S. Clark. Radio observations of the massive stellar cluster Westerlund 1. *ArXiv e-prints*, May 2007.

[75] G. Duchêne and A. Kraus. Stellar Multiplicity. *ARA&A*, 51:269–310, August 2013.

[76] S. Ekström, C. Georgy, P. Eggenberger, G. Meynet, N. Mowlavi, A. Wyttenbach, A. Granada, T. Decressin, R. Hirschi, U. Frischknecht, C. Charbonnel, and A. Maeder. Grids of stellar models with rotation. I. Models from 0.8 to 120 $M_\odot$ at solar metallicity ($Z = 0.014$). *A&A*, 537:A146, January 2012.

[77] J. J. Eldridge, E. R. Stanway, L. Xiao, L. A. S. McClelland, G. Taylor, M. Ng, S. M. L. Greis, and J. C. Bray. Binary Population and Spectral Synthesis Version 2.1: Construction, Observational Verification, and New Results. *PASA*, 34:e058, November 2017.

[78] R. Farmer, C. E. Fields, I. Petermann, Luc Dessart, M. Cantiello, B. Paxton, and F. X. Timmes. On Variations Of Pre-supernova Model Properties. *ApJS*, 227(2):22, December 2016.

[79] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee: The MCMC Hammer. *PASP*, 125:306, March 2013.

[80] Grant Foster. Wavelets for period analysis of unevenly sampled time series. *AJ*, 112:1709, Oct 1996.

[81] M. Fraser, R. Kotak, A. Pastorello, A. Jerkstrand, S. J. Smartt, T.-W. Chen, M. Childress, G. Gilmore, C. Inserra, E. Kankare, S. Margheim, S. Mattila, S. Valenti, C. Ashall, S. Benetti, M. T. Botticella, F. E. Bauer, H. Campbell, N. Elias-Rosa, M. Fleury, A. Gal-Yam, S. Hachinger, D. A. Howell, L. Le Guillou, P.-F. Léget, A. Morales-Garoffolo, J. Polshaw, S. Spiro, M. Sullivan, S. Taubenberger, M. Turatto, E. S. Walker, D. R. Young, and B. Zhang. SN 2009ip at late times - an interacting transient at +2 years. *MNRAS*, 453:3886–3905, November 2015.

[82] Y. Frémat, J. Zorec, A. M. Hubert, and M. Floquet. Effects of gravitational darkening on the determination of fundamental parameters in fast-rotating B-type stars. *A&A*, 440(1):305–320, September 2005.

[83] J. Fuller. Pre-supernova outbursts via wave heating in massive stars - I. Red supergiants. *MNRAS*, 470:1642–1656, September 2017.

[84] A. W. Fullerton, D. L. Massa, and R. K. Prinja. The Discordance of Mass-Loss Estimates for Galactic O-Type Stars. *ApJ*, 637(2):1025–1039, February 2006.

[85] Gaia Collaboration, A. G. A. Brown, A. Vallenari, T. Prusti, J. H. J. de Bruijne, C. Babusiaux, C. A. L. Bailer-Jones, M. Biermann, D. W. Evans, L. Eyer, and et al. Gaia Data Release 2. Summary of the contents and survey properties. *A&A*, 616:A1, August 2018.

[86] Gaia Collaboration, A. G. A. Brown, A. Vallenari, T. Prusti, J. H. J. de Bruijne, C. Babusiaux, M. Biermann, O. L. Creevey, D. W. Evans, L. Eyer, A. Hutton, F. Jansen, C. Jordi, S. A. Klioner, U. Lammers, L. Lindegren, X. Luri, F. Mignard, C. Panem, D. Pourbaix, S. Randich, P. Sartoretti, C. Soubiran, N. A. Walton, F. Arenou, C. A. L. Bailer-Jones, U. Bastian, M. Cropper, R. Drimmel, D. Katz, M. G. Lattanzi, F. van Leeuwen, J. Bakker, C. Cacciari, J. Castañeda, F. De Angeli, C. Ducourant, C. Fabricius, M. Fouesneau, Y. Frémat, R. Guerra, A. Guerrier, J. Guiraud, A. Jean-Antoine Piccolo, E. Masana, R. Messineo, N. Mowlavi, C. Nicolas, K. Nienartowicz, F. Pailler, P. Panuzzo, F. Riclet, W. Roux, G. M. Seabroke, R. Sordo, P. Tanga, F. Thévenin, G. Gracia-Abril, J. Portell, D. Teyssier, M. Altmann, R. Andrae, I. Bellas-Velidis, K. Benson, J. Berthier, R. Blomme, E. Brugaletta, P. W. Burgess, G. Busso, B. Carry, A. Cellino, N. Cheek, G. Clementini, Y. Damerdji, M. Davidson, L. Delchambre, A. Dell'Oro, J. Fernández-Hernández, L. Galluccio, P. García-Lario, M. Garcia-Reinaldos, J. González-Núñez, E. Gosset, R. Haigron, J. L. Halbwachs, N. C. Hambly, D. L. Harrison, D. Hatzidimitriou, U. Heiter, J. Hernández, D. Hestroffer, S. T. Hodgkin, B. Holl, K. Janßen, G. Jevardat de Fombelle, S. Jordan, A. Krone-Martins, A. C. Lanzafame, W. Löffler, A. Lorca, M. Manteiga, O. Marchal, P. M. Marrese, A. Moitinho, A. Mora, K. Muinonen, P. Osborne, E. Pancino, T. Pauwels, J. M. Petit, A. Recio-Blanco, P. J. Richards, M. Riello, L. Rimoldini, A. C. Robin, T. Roegiers, J. Rybizki, L. M. Sarro, C. Siopis, M. Smith, A. Sozzetti, A. Ulla, E. Utrilla, M. van Leeuwen, W. van Reeven, U. Abbas, A. Abreu Aramburu, S. Accart, C. Aerts, J. J. Aguado, M. Ajaj,

G. Altavilla, M. A. Álvarez, J. Álvarez Cid-Fuentes, J. Alves, R. I. Anderson, E. Anglada Varela, T. Antoja, M. Audard, D. Baines, S. G. Baker, L. Balaguer-Núñez, E. Balbinot, Z. Balog, C. Barache, D. Barbato, M. Barros, M. A. Barstow, S. Bartolomé, J. L. Bassilana, N. Bauchet, A. Baudesson-Stella, U. Becciani, M. Bellazzini, M. Bernet, S. Bertone, L. Bianchi, S. Blanco-Cuaresma, T. Boch, A. Bombrun, D. Bossini, S. Bouquillon, A. Bragaglia, L. Bramante, E. Breedt, A. Bressan, N. Brouillet, B. Bucciarelli, A. Burlacu, D. Busonero, A. G. Butkevich, R. Buzzi, E. Caffau, R. Cancelliere, H. Cánovas, T. Cantat-Gaudin, R. Carballo, T. Carlucci, M. I. Carnerero, J. M. Carrasco, L. Casamiquela, M. Castellani, A. Castro-Ginard, P. Castro Sampol, L. Chaoul, P. Charlot, L. Chemin, A. Chiavassa, M. R. L. Cioni, G. Comoretto, W. J. Cooper, T. Cornez, S. Cowell, F. Crifo, M. Crosta, C. Crowley, C. Dafonte, A. Dapergolas, M. David, P. David, P. de Laverny, F. De Luise, R. De March, J. De Ridder, R. de Souza, P. de Teodoro, A. de Torres, E. F. del Peloso, E. del Pozo, M. Delbo, A. Delgado, H. E. Delgado, J. B. Delisle, P. Di Matteo, S. Diakite, C. Diener, E. Distefano, C. Dolding, D. Eappachen, B. Edvardsson, H. Enke, P. Esquej, C. Fabre, M. Fabrizio, S. Faigler, G. Fedorets, P. Fernique, A. Fienga, F. Figueras, C. Fouron, F. Fragkoudi, E. Fraile, F. Franke, M. Gai, D. Garabato, A. Garcia-Gutierrez, M. García-Torres, A. Garofalo, P. Gavras, E. Gerlach, R. Geyer, P. Giacobbe, G. Gilmore, S. Girona, G. Giuffrida, R. Gomel, A. Gomez, I. Gonzalez-Santamaria, J. J. González-Vidal, M. Granvik, R. Gutiérrez-Sánchez, L. P. Guy, M. Hauser, M. Haywood, A. Helmi, S. L. Hidalgo, T. Hilger, N. Hładczuk, D. Hobbs, G. Holland, H. E. Huckle, G. Jasniewicz, P. G. Jonker, J. Juaristi Campillo, F. Julbe, L. Karbevska, P. Kervella, S. Khanna, A. Kochoska, M. Kontizas, G. Kordopatis, A. J. Korn, Z. Kostrzewa-Rutkowska, K. Kruszyńska, S. Lambert, A. F. Lanza, Y. Lasne, J. F. Le Campion, Y. Le Fustec, Y. Lebreton, T. Lebzelter, S. Leccia, N. Leclerc, I. Lecoeur-Taibi, S. Liao, E. Licata, E. P. Lindstrøm, T. A. Lister, E. Livanou, A. Lobel, P. Madrero Pardo, S. Managau, R. G. Mann, J. M. Marchant, M. Marconi, M. M. S. Marcos Santos, S. Marinoni, F. Marocco, D. J. Marshall, L. Martin Polo, J. M. Martín-Fleitas, A. Masip, D. Massari, A. Mastrobuono-Battisti, T. Mazeh, P. J. McMillan, S. Messina, D. Michalik, N. R. Millar, A. Mints, D. Molina, R. Molinaro, L. Molnár,

P. Montegriffo, R. Mor, R. Morbidelli, T. Morel, D. Morris, A. F. Mulone, D. Munoz, T. Muraveva, C. P. Murphy, I. Musella, L. Noval, C. Ordénovic, G. Orrù, J. Osinde, C. Pagani, I. Pagano, L. Palaversa, P. A. Palicio, A. Panahi, M. Pawlak, X. Peñalosa Esteller, A. Penttilä, A. M. Piersimoni, F. X. Pineau, E. Plachy, G. Plum, E. Poggio, E. Poretti, E. Poujoulet, A. Prša, L. Pulone, E. Racero, S. Ragaini, M. Rainer, C. M. Raiteri, N. Rambaux, P. Ramos, M. Ramos-Lerate, P. Re Fiorentin, S. Regibo, C. Reylé, V. Ripepi, A. Riva, G. Rixon, N. Robichon, C. Robin, M. Roelens, L. Rohrbasser, M. Romero-Gómez, N. Rowell, F. Royer, K. A. Rybicki, G. Sadowski, A. Sagristà Sellés, J. Sahlmann, J. Salgado, E. Salguero, N. Samaras, V. Sanchez Gimenez, N. Sanna, R. Santoveña, M. Sarasso, M. Schultheis, E. Sciacca, M. Segol, J. C. Segovia, D. Ségransan, D. Semeux, S. Shahaf, H. I. Siddiqui, A. Siebert, L. Siltala, E. Slezak, R. L. Smart, E. Solano, F. Solitro, D. Souami, J. Souchay, A. Spagna, F. Spoto, I. A. Steele, H. Steidelmüller, C. A. Stephenson, M. Süveges, L. Szabados, E. Szegedi-Elek, F. Taris, G. Tauran, M. B. Taylor, R. Teixeira, W. Thuillot, N. Tonello, F. Torra, J. Torra, C. Turon, N. Unger, M. Vaillant, E. van Dillen, O. Vanel, A. Vecchiato, Y. Viala, D. Vicente, S. Voutsinas, M. Weiler, T. Wevers, Ł. Wyrzykowski, A. Yoldas, P. Yvard, H. Zhao, J. Zorec, S. Zucker, C. Zurbach, and T. Zwitter. Gaia Early Data Release 3. Summary of the contents and survey properties. *A&A*, 649:A1, May 2021.

[87] Gaia Collaboration, A. Helmi, F. van Leeuwen, P. J. McMillan, D. Massari, T. Antoja, A. C. Robin, L. Lindegren, U. Bastian, F. Arenou, C. Babusiaux, M. Biermann, M. A. Breddels, D. Hobbs, C. Jordi, E. Pancino, C. Reylé, J. Veljanoski, A. G. A. Brown, A. Vallenari, T. Prusti, J. H. J. de Bruijne, C. A. L. Bailer-Jones, D. W. Evans, L. Eyer, F. Jansen, S. A. Klioner, U. Lammers, X. Luri, F. Mignard, C. Panem, D. Pourbaix, S. Randich, P. Sartoretti, H. I. Siddiqui, C. Soubiran, N. A. Walton, M. Cropper, R. Drimmel, D. Katz, M. G. Lattanzi, J. Bakker, C. Cacciari, J. Castañeda, L. Chaoul, N. Cheek, F. De Angeli, C. Fabricius, R. Guerra, B. Holl, E. Masana, R. Messineo, N. Mowlavi, K. Nienartowicz, P. Panuzzo, J. Portell, M. Riello, G. M. Seabroke, P. Tanga, F. Thévenin, G. Gracia-Abril, G. Comoretto, M. Garcia-Reinaldos, D. Teyssier, M. Altmann, R. Andrae, M. Audard, I. Bellas-Velidis,

K. Benson, J. Berthier, R. Blomme, P. Burgess, G. Busso, B. Carry, A. Cellino, G. Clementini, M. Clotet, O. Creevey, M. Davidson, J. De Ridder, L. Delchambre, A. Dell'Oro, C. Ducourant, J. Fernández-Hernández, M. Fouesneau, Y. Frémat, L. Galluccio, M. García-Torres, J. González-Núñez, J. J. González-Vidal, E. Gosset, L. P. Guy, J. L. Halbwachs, N. C. Hambly, D. L. Harrison, J. Hernández, D. Hestroffer, S. T. Hodgkin, A. Hutton, G. Jasniewicz, A. Jean-Antoine-Piccolo, S. Jordan, A. J. Korn, A. Krone-Martins, A. C. Lanzafame, T. Lebzelter, W. Löffler, M. Manteiga, P. M. Marrese, J. M. Martín-Fleitas, A. Moitinho, A. Mora, K. Muinonen, J. Osinde, T. Pauwels, J. M. Petit, A. Recio-Blanco, P. J. Richards, L. Rimoldini, L. M. Sarro, C. Siopis, M. Smith, A. Sozzetti, M. Süveges, J. Torra, W. van Reeven, U. Abbas, A. Abreu Aramburu, S. Accart, C. Aerts, G. Altavilla, M. A. Álvarez, R. Alvarez, J. Alves, R. I. Anderson, A. H. Andrei, E. Anglada Varela, E. Antiche, B. Arcay, T. L. Astraatmadja, N. Bach, S. G. Baker, L. Balaguer-Núñez, P. Balm, C. Barache, C. Barata, D. Barbato, F. Barblan, P. S. Barklem, D. Barrado, M. Barros, M. A. Barstow, S. Bartholomé Muñoz, J. L. Bassilana, U. Becciani, M. Bellazzini, A. Berihuete, S. Bertone, L. Bianchi, O. Bienaymé, S. Blanco-Cuaresma, T. Boch, C. Boeche, A. Bombrun, R. Borrachero, D. Bossini, S. Bouquillon, G. Bourda, A. Bragaglia, L. Bramante, A. Bressan, N. Brouillet, T. Brüsemeister, E. Brugaletta, B. Bucciarelli, A. Burlacu, D. Busonero, A. G. Butkevich, R. Buzzi, E. Caffau, R. Cancelliere, G. Cannizzaro, T. Cantat-Gaudin, R. Carballo, T. Carlucci, J. M. Carrasco, L. Casamiquela, M. Castellani, A. Castro-Ginard, P. Charlot, L. Chemin, A. Chiavassa, G. Cocozza, G. Costigan, S. Cowell, F. Crifo, M. Crosta, C. Crowley, J. Cuypers, C. Dafonte, Y. Damerdji, A. Dapergolas, P. David, M. David, P. de Laverny, F. De Luise, R. De March, D. de Martino, R. de Souza, A. de Torres, J. Debosscher, E. del Pozo, M. Delbo, A. Delgado, H. E. Delgado, P. Di Matteo, S. Diakite, C. Diener, E. Distefano, C. Dolding, P. Drazinos, J. Durán, B. Edvardsson, H. Enke, K. Eriksson, P. Esquej, G. Eynard Bontemps, C. Fabre, M. Fabrizio, S. Faigler, A. J. Falcão, M. Farràs Casas, L. Federici, G. Fedorets, P. Fernique, F. Figueras, F. Filippi, K. Findeisen, A. Fonti, E. Fraile, M. Fraser, B. Frézouls, M. Gai, S. Galleti, D. Garabato, F. García-Sedano, A. Garofalo, N. Garralda, A. Gavel, P. Gavras, J. Gerssen, R. Geyer, P. Gi-

acobbe, G. Gilmore, S. Girona, G. Giuffrida, F. Glass, M. Gomes, M. Granvik, A. Gueguen, A. Guerrier, J. Guiraud, R. Gutiérrez-Sánchez, W. Hofmann, G. Holland, H. E. Huckle, A. Hypki, V. Icardi, K. Janßen, G. Jevardat de Fombelle, P. G. Jonker, Á. L. Juhász, F. Julbe, A. Karampelas, A. Kewley, J. Klar, A. Kochoska, R. Kohley, K. Kolenberg, M. Kontizas, E. Kontizas, S. E. Koposov, G. Kordopatis, Z. Kostrzewa-Rutkowska, P. Koubsky, S. Lambert, A. F. Lanza, Y. Lasne, J. B. Lavigne, Y. Le Fustec, C. Le Poncin-Lafitte, Y. Lebreton, S. Leccia, N. Leclerc, I. Lecoeur-Taibi, H. Lenhardt, F. Leroux, S. Liao, E. Licata, H. E. P. Lindstrøm, T. A. Lister, E. Livanou, A. Lobel, M. López, S. Managau, R. G. Mann, G. Mantelet, O. Marchal, J. M. Marchant, M. Marconi, S. Marinoni, G. Marschalkó, D. J. Marshall, M. Martino, G. Marton, N. Mary, G. Matijevič, T. Mazeh, S. Messina, D. Michalik, N. R. Millar, D. Molina, R. Molinaro, L. Molnár, P. Montegriffo, R. Mor, R. Morbidelli, T. Morel, D. Morris, A. F. Mulone, T. Muraveva, I. Musella, G. Nelemans, L. Nicastro, L. Noval, W. O'Mullane, C. Ordénovic, D. Ordóñez-Blanco, P. Osborne, C. Pagani, I. Pagano, F. Pailler, H. Palacin, L. Palaversa, A. Panahi, M. Pawlak, A. M. Piersimoni, F. X. Pineau, E. Plachy, G. Plum, E. Poggio, E. Poujoulet, A. Prša, L. Pulone, E. Racero, S. Ragaini, N. Rambaux, M. Ramos-Lerate, S. Regibo, F. Riclet, V. Ripepi, A. Riva, A. Rivard, G. Rixon, T. Roegiers, M. Roelens, M. Romero-Gómez, N. Rowell, F. Royer, L. Ruiz-Dern, G. Sadowski, T. Sagristà Sellés, J. Sahlmann, J. Salgado, E. Salguero, N. Sanna, T. Santana-Ros, M. Sarasso, H. Savietto, M. Schultheis, E. Sciacca, M. Segol, J. C. Segovia, D. Ségransan, I. C. Shih, L. Siltala, A. F. Silva, R. L. Smart, K. W. Smith, E. Solano, F. Solitro, R. Sordo, S. Soria Nieto, J. Souchay, A. Spagna, F. Spoto, U. Stampa, I. A. Steele, H. Steidelmüller, C. A. Stephenson, H. Stoev, F. F. Suess, J. Surdej, L. Szabados, E. Szegedi-Elek, D. Tapiador, F. Taris, G. Tauran, M. B. Taylor, R. Teixeira, D. Terrett, P. Teyssand ier, W. Thuillot, A. Titarenko, F. Torra Clotet, C. Turon, A. Ulla, E. Utrilla, S. Uzzi, M. Vaillant, G. Valentini, V. Valette, A. van Elteren, E. Van Hemelryck, M. van Leeuwen, M. Vaschetto, A. Vecchiato, Y. Viala, D. Vicente, S. Vogt, C. von Essen, H. Voss, V. Votruba, S. Voutsinas, G. Walmsley, M. Weiler, O. Wertz, T. Wevems, Ł. Wyrzykowski, A. Yoldas, M. Žerjal, H. Ziaeepour, J. Zorec, S. Zschocke, S. Zucker, C. Zurbach, and

T. Zwitter. Gaia Data Release 2. Kinematics of globular clusters and dwarf galaxies around the Milky Way. *A&A*, 616:A12, Aug 2018.

[88] Ana E. García Pérez, Carlos Allende Prieto, Jon A. Holtzman, Matthew Shetrone, Szabolcs Mészáros, Dmitry Bizyaev, Ricardo Carrera, Katia Cunha, D. A. García-Hernández, Jennifer A. Johnson, Steven R. Majewski, David L. Nidever, Ricardo P. Schiavon, Neville Shane, Verne V. Smith, Jennifer Sobeck, Nicholas Troup, Olga Zamora, David H. Weinberg, Jo Bovy, Daniel J. Eisenstein, Diane Feuillet, Peter M. Frinchaboy, Michael R. Hayden, Fred R. Hearty, Duy C. Nguyen, Robert W. O'Connell, Marc H. Pinsonneault, John C. Wilson, and Gail Zasowski. ASPCAP: The APOGEE Stellar Parameter and Chemical Abundances Pipeline. *AJ*, 151(6):144, June 2016.

[89] C. D. Garmany, P. S. Conti, and P. Massey. Spectroscopic studies of O type stars. IX - Binary frequency. *ApJ*, 242:1063–1076, December 1980.

[90] C. Georgy, S. Ekström, P. Eggenberger, G. Meynet, L. Haemmerlé, A. Maeder, A. Granada, J. H. Groh, R. Hirschi, N. Mowlavi, N. Yusof, C. Charbonnel, T. Decressin, and F. Barblan. Grids of stellar models with rotation. III. Models from 0.8 to 120 $M_\odot$ at a metallicity $Z = 0.002$. *A&A*, 558:A103, October 2013.

[91] C. Georgy, S. Ekström, G. Meynet, P. Massey, E. M. Levesque, R. Hirschi, P. Eggenberger, and A. Maeder. Grids of stellar models with rotation. II. WR populations and supernovae/GRB progenitors at $Z = 0.014$. *A&A*, 542:A29, Jun 2012.

[92] C. Georgy, A. Granada, S. Ekström, G. Meynet, R. I. Anderson, A. Wyttenbach, P. Eggenberger, and A. Maeder. Populations of rotating stars. III. SYCLIST, the new Geneva population synthesis code. *A&A*, 566:A21, Jun 2014.

[93] Cyril Georgy. *Anisotropic mass loss and stellar evolution: from Be Stars to Gamma Ray Bursts*. PhD thesis, Geneva Observatory, Geneva University, September 2010.

[94] D. R. Gies, W. G. Bagnuolo, Jr., E. C. Ferrara, A. B. Kaye, M. L. Thaller, L. R. Penny, and G. J. Peters. Hubble Space Telescope Goddard High Resolution Spectrograph Observations of the Be + sdO Binary ϕ Persei. *ApJ*, 493:440–450, January 1998.

[95] E. Glebbeek, E. Gaburov, S. E. de Mink, O. R. Pols, and S. F. Portegies Zwart. The evolution of runaway stellar collision products. *A&A*, 497(1):255–264, April 2009.

[96] Keyan Gootkin, Trevor Dorn-Wallenstein, Jamie R. Lomax, Gwendolyn Eadie, Emily M. Levesque, Brian Babler, Jennifer L. Hoffman, Marilyn R. Meade, Kenneth Nordsieck, and John P. Wisniewski. 13 yr of P Cygni Spectropolarimetry: Investigating Mass Loss through Hα, Periodicity, and Ellipticity. *ApJ*, 900(2):162, September 2020.

[97] K. D. Gordon, G. C. Clayton, K. A. Misselt, A. U. Landolt, and M. J. Wolff. A Quantitative Comparison of the Small Magellanic Cloud, Large Magellanic Cloud, and Milky Way Ultraviolet to Near-Infrared Extinction Curves. *ApJ*, 594:279–293, September 2003.

[98] Michael S. Gordon and Roberta M. Humphreys. Red Supergiants, Yellow Hypergiants, and Post-RSG Evolution. *Galaxies*, 7(4):92, December 2019.

[99] S. Gossage, C. Conroy, A. Dotter, J. Choi, P. Rosenfield, P. Cargile, and A. Dolphin. Age Determinations of the Hyades, Praesepe, and Pleiades via MESA Models with Rotation. *ApJ*, 863:67, August 2018.

[100] Y. Götberg, S. E. de Mink, J. H. Groh, T. Kupfer, P. A. Crowther, E. Zapartas, and M. Renzo. Spectral models for binary products: Unifying subdwarfs and Wolf-Rayet stars as a sequence of stripped-envelope stars. *A&A*, 615:A78, July 2018.

[101] L. Grassitelli, L. Fossati, N. Langer, A. Miglio, A. G. Istrate, and D. Sanyal. Relating turbulent pressure and macroturbulence across the HR diagram with a possible link to γ Doradus stars. *A&A*, 584:L2, December 2015.

[102] L. Grassitelli, L. Fossati, N. Langer, S. Simón-Díaz, N. Castro, and D. Sanyal. Metallicity dependence of turbulent pressure and macroturbulence in stellar envelopes. *A&A*, 593:A14, August 2016.

[103] R. O. Gray, M. G. Napier, and L. I. Winkler. The Physical Basis of Luminosity Classification in the Late A-, F-, and Early G-Type Stars. I. Precise Spectral Types for 372 Stars. *AJ*, 121:2148–2158, April 2001.

[104] J. H. Groh, S. Ekström, C. Georgy, G. Meynet, A. Choplin, P. Eggenberger, R. Hirschi, A. Maeder, L. J. Murphy, I. Boian, and E. J. Farrell. Grids of stellar models with rotation. IV. Models from 1.7 to 120 $M_\odot$ at a metallicity $Z = 0.0004$. *A&A*, 627:A24, July 2019.

[105] J. H. Groh, A. S. Oliveira, and J. E. Steiner. The qWR star HD 45166. II. Fundamental stellar parameters and evidence of a latitude-dependent wind. *A&A*, 485:245–256, July 2008.

[106] B. Gustafsson, B. Edvardsson, K. Eriksson, U. G. Jørgensen, Å. Nordlund, and B. Plez. A grid of MARCS model atmospheres for late-type stars. I. Methods and general properties. *A&A*, 486(3):951–970, August 2008.

[107] I. Guyon, K. Bennett, G. Cawley, H. J. Escalante, S. Escalera, Tin Kam Ho, N. Macià, B. Ray, M. Saeed, A. Statnikov, and E. Viegas. Design of the 2015 chalearn automl challenge. In *2015 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2015.

[108] G. M. H. J. Habets and J. R. W. Heintze. Empirical bolometric corrections for the main-sequence. *A&AS*, 46:193–237, November 1981.

[109] X. Haubois, B. Norris, P. G. Tuthill, C. Pinte, P. Kervella, J. H. Girard, N. M. Kostogryz, S. V. Berdyugina, G. Perrin, S. Lacour, A. Chiavassa, and S. T. Ridgway. The inner dust shell of Betelgeuse detected by polarimetric aperture-masking interferometry. *A&A*, 628:A101, August 2019.

[110] P. Hellings and D. Vanbeveren. Evolution of massive stars with low metal abundance holding for the Magellanic Clouds. *A&A*, 95:14–17, February 1981.

[111] F. Herwig, T. Bloecker, D. Schoenberner, and M. El Eid. Stellar evolution of low and intermediate-mass stars. IV. Hydrodynamically-based overshoot and nucleosynthesis in AGB stars. *A&A*, 324:L81–L84, August 1997.

[112] R. Hirschi, G. Meynet, and A. Maeder. Yields of rotating stars at solar metallicity. *A&A*, 433(3):1013–1022, April 2005.

[113] A. M. Hopkins. The Dawes Review 8: Measuring the Stellar Initial Mass Function. *PASA*, 35:e039, November 2018.

[114] W. Huang, D. R. Gies, and M. V. McSwain. A Stellar Rotation Census of B Stars: From ZAMS to TAMS. *ApJ*, 722:605–619, October 2010.

[115] A.-L. Huat, A.-M. Hubert, F. Baudin, M. Floquet, C. Neiner, Y. Frémat, J. Gutiérrez-Soto, L. Andrade, B. de Batz, P. D. Diago, M. Emilio, F. Espinosa Lara, J. Fabregat, E. Janot-Pacheco, B. Leroy, C. Martayan, T. Semaan, J. Suso, M. Auvergne, C. Catala, E. Michel, and R. Samadi. The B0.5IVe CoRoT target HD 49330. I. Photometric analysis from CoRoT data. *A&A*, 506:95–101, October 2009.

[116] R. M. Humphreys. Studies of luminous stars in nearby galaxies. I. Supergiants and O stars in the Milky Way. *ApJS*, 38:309–350, December 1978.

[117] R. M. Humphreys and K. Davidson. Studies of luminous stars in nearby galaxies. III. Comments on the evolution of the most massive stars in the Milky Way and the Large Magellanic Cloud. *ApJ*, 232:409–420, Sep 1979.

[118] R. M. Humphreys, K. Weis, K. Davidson, and M. S. Gordon. On the Social Traits of Luminous Blue Variables. *ApJ*, 825:64, July 2016.

[119] Roberta M. Humphreys, Kris Davidson, and Nathan Smith. Crossing the Yellow Void: Spatially Resolved Spectroscopy of the Post-Red Supergiant IRC +10420 and Its Circumstellar Ejecta. *AJ*, 124(2):1026–1044, August 2002.

[120] J. R. Hurley, C. A. Tout, and O. R. Pols. Evolution of binary stars and the effect of tides on binary populations. *MNRAS*, 329:897–928, February 2002.

[121] Valentin D. Ivanov, Lodovico Coccato, Mark J. Neeser, Fernando Selman, Alessandro Pizzella, Elena Dalla Bontà, Enrico M. Corsini, and Lorenzo Morelli. MUSE library of stellar spectra. *A&A*, 629:A100, September 2019.

[122] C. S. Jeffery and H. Saio. Radial pulsation as a function of hydrogen abundance. *MNRAS*, 458(2):1352–1373, May 2016.

[123] Y.-F. Jiang, M. Cantiello, L. Bildsten, E. Quataert, O. Blaes, and J. Stone. Outbursts of luminous blue variable stars from variations in the helium opacity. *Nature*, 561:498–501, September 2018.

[124] Yan-Fei Jiang, Matteo Cantiello, Lars Bildsten, Eliot Quataert, and Omer Blaes. Local Radiation Hydrodynamic Simulations of Massive Star Envelopes at the Iron Opacity Peak. *ApJ*, 813(1):74, November 2015.

[125] C. Johnston, B. Buysschaert, A. Tkachenko, C. Aerts, and C. Neiner. Detection of intrinsic variability in the eclipsing massive main-sequence O+B binary HD 165246. *MNRAS*, 469:L118–L122, July 2017.

[126] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. [Online; accessed ¡today¿].

[127] Olivia C. Jones, Margaret Meixner, Kay Justtanont, and Alistair Glasse. Probing the Dusty Stellar Populations of the Local Volume Galaxies with JWST/MIRI. *ApJ*, 841(1):15, May 2017.

[128] S. Jones, C. Ritter, F. Herwig, C. Fryer, M. Pignatari, M. G. Bertolli, and B. Paxton. H ingestion into He-burning convection zones in super-AGB stellar models as a potential site for intermediate neutron-density nucleosynthesis. *MNRAS*, 455(4):3848–3863, February 2016.

[129] Sara Kahanamoku, Rosie 'Anolani Alegado, Aurora Kagawa-Viviani, Katie Leimomi Kamelamela, Brittany Kamai, Lucianne M Walkowicz, Chanda Prescod-Weinstein, Mithi Alexa de los Reyes, and Hilding Neilson. A Native Hawaiian-led summary of the current impact of constructing the Thirty Meter Telescope on Maunakea. *arXiv e-prints*, page arXiv:2001.00970, January 2020.

[130] T. Kallinger, J. De Ridder, S. Hekker, S. Mathur, B. Mosser, M. Gruberbauer, R. A. García, C. Karoff, and J. Ballot. The connection between stellar granulation and oscillation as seen by the Kepler mission. *A&A*, 570:A41, October 2014.

[131] A. Kaufer, O. Stahl, B. Wolf, A. W. Fullerton, T. Gaeng, C. A. Gummersbach, I. Jankovics, J. Kovacs, H. Mandel, J. Peitz, T. Rivinius, and T. Szeifert. Long-term spectroscopic monitoring of BA-type supergiants. III. Variability of photospheric lines. *A&A*, 320:273–286, April 1997.

[132] N. D. Kee, J. O. Sundqvist, L. Decin, A. de Koter, and H. Sana. Analytic, dust-independent mass-loss rates for red supergiant winds initiated by turbulent pressure. *A&A*, 646:A180, February 2021.

[133] Philip C. Keenan and Raymond C. McNeil. The Perkins Catalog of Revised MK Types for the Cooler Stars. *The Astrophysical Journal Supplement Series*, 71:245, October 1989.

[134] L. J. Kewley, M. A. Dopita, R. S. Sutherland, C. A. Heisler, and J. Trevena. Theoretical Modeling of Starburst Galaxies. *ApJ*, 556:121–140, July 2001.

[135] George S. Kimeldorf and Grace Wahba. A Correspondence Between Bayesian Estimation

on Stochastic Processes and Smoothing by Splines. *The Annals of Mathematical Statistics*, 41(2):495 – 502, 1970.

[136] R. Kippenhahn, G. Ruschenplatt, and H. C. Thomas. The time scale of thermohaline mixing in stars. *A&A*, 91(1-2):175–180, November 1980.

[137] Rudolf Kippenhahn and Alfred Weigert. *Stellar Structure and Evolution*. 1990.

[138] L. L. Kiss, G. M. Szabó, and T. R. Bedding. Variability in red supergiant stars: pulsations, long secondary periods and convection noise. *MNRAS*, 372:1721–1734, November 2006.

[139] S. Knerr, L. Personnaz, and Gérard Dreyfus. Single-layer learning revisited: A stepwise procedure for building and training a neural network. *Neurocomputing: Algorithms, Architectures and applications*, 1, 01 1990.

[140] C. S. Kochanek. On the red supergiant problem. *MNRAS*, 493(4):4945–4949, April 2020.

[141] C. S. Kochanek, R. Khan, and X. Dai. On Absorption by Circumstellar Dust, with the Progenitor of SN 2012aw as a Case Study. *ApJ*, 759(1):20, November 2012.

[142] G. Kovács. The distance modulus of the Large Magellanic Cloud based on double-mode RR Lyrae stars. *A&A*, 363:L1–L4, November 2000.

[143] G. Kovács. The distance modulus of the Small Magellanic Cloud based on double-mode Cepheids. *A&A*, 360:L1–L4, Aug 2000.

[144] J. Krtička and A. Feldmeier. Light variations due to the line-driven wind instability and wind blanketing in O stars. *A&A*, 617:A121, September 2018.

[145] N. Kudryavtseva, W. Brandner, M. Gennaro, B. Rochau, A. Stolte, M. Andersen, N. Da Rio, T. Henning, E. Tognelli, D. Hogg, S. Clark, and R. Waters. Instantaneous Starburst of the Massive Clusters Westerlund 1 and NGC 3603 YC. *ApJL*, 750:L44, May 2012.

[146] Andrea Kunder, Georges Kordopatis, Matthias Steinmetz, Tomaž Zwitter, Paul J. McMillan, Luca Casagrande, Harry Enke, Jennifer Wojno, Marica Valentini, Cristina Chiappini, Gal Matijevič, Alessand ro Siviero, Patrick de Laverny, Alejandra Recio-Blanco, Albert Bijaoui, Rosemary F. G. Wyse, James Binney, E. K. Grebel, Amina Helmi, Paula Jofre, Teresa Antoja, Gerard Gilmore, Arnaud Siebert, Benoit Famaey, Olivier Bienaymé, Brad K. Gibson, Kenneth C. Freeman, Julio F. Navarro, Ulisse Munari, George Seabroke, Borja Anguiano, Maruša Žerjal, Ivan Minchev, Warren Reid, Joss Bland-Hawthorn, Janez Kos, Sanjib Sharma, Fred Watson, Quentin A. Parker, Ralf-Dieter Scholz, Donna Burton, Paul Cass, Malcolm Hartley, Kristin Fiegert, Milorad Stupar, Andreas Ritter, Keith Hawkins, Ortwin Gerhard, W. J. Chaplin, G. R. Davies, Y. P. Elsworth, M. N. Lund, A. Miglio, and B. Mosser. The Radial Velocity Experiment (RAVE): Fifth Data Release. *AJ*, 153:75, Feb 2017.

[147] Donald W. Kurtz, Hiromoto Shibahashi, Simon J. Murphy, Timothy R. Bedding, and Dominic M. Bowman. A unifying explanation of complex frequency spectra of γ Dor, SPB and Be stars: combination frequencies and highly non-sinusoidal light curves. *MNRAS*, 450:3015–3029, July 2015.

[148] R. L. Kurucz. Model Atmospheres for Population Synthesis. In Beatriz Barbuy and Alvio Renzini, editors, *The Stellar Populations of Galaxies*, volume 149 of *IAU Symposium*, page 225, January 1992.

[149] H. J. G. L. M. Lamers. Stars with the B[e] Phenomenon: Summary and Personal Conclusions. In Michaela Kraus and Anatoly S. Miroshnichenko, editors, *Stars with the B[e] Phenomenon*, volume 355 of *Astronomical Society of the Pacific Conference Series*, page 371, December 2006.

[150] H. J. G. L. M. Lamers and J. P. Cassinelli. *Introduction to Stellar Winds*. June 1999.

[151] N. Langer. Presupernova Evolution of Massive Single and Binary Stars. *ARA&A*, 50:107–164, September 2012.

[152] N. Langer and A. Maeder. The problem of the blue-to-red supergiant ratio in galaxies. *A&A*, 295:685, March 1995.

[153] E. Laplace, S. Justham, M. Renzo, Y. Götberg, R. Farmer, D. Vartanyan, and S. E. de Mink. Different to the core: the pre-supernova structures of massive single and binary-stripped stars. *arXiv e-prints*, page arXiv:2102.05036, February 2021.

[154] Daniel Lecoanet, Matteo Cantiello, Eliot Quataert, Louis-Alexandre Couston, Keaton J. Burns, Benjamin J. S. Pope, Adam S. Jermyn, Benjamin Favier, and Michael Le Bars. Low-frequency Variability in Massive Stars: Core Generation or Surface Phenomenon? *ApJL*, 886(1):L15, Nov 2019.

[155] Claus Leitherer, Sylvia Ekström, Georges Meynet, Daniel Schaerer, Katerina B. Agienko, and Emily M. Levesque. The Effects of Stellar Rotation. II. A Comprehensive Set of Starburst99 Models. *ApJS*, 212(1):14, May 2014.

[156] E. Levesque, H. Lamers, and A. de Koter. Rotation and Mass Loss in Luminous Blue Variables. In *American Astronomical Society Meeting Abstracts*, volume 52 of *American Astronomical Society Meeting Abstracts*, page 131.07, Jan 2020.

[157] E. M. Levesque. *Astrophysics of Red Supergiants*. IOP Publishing Ltd., December 2017.

[158] E. M. Levesque, P. Massey, K. A. G. Olsen, B. Plez, G. Meynet, and A. Maeder. The Effective Temperatures and Physical Properties of Magellanic Cloud Red Supergiants: The Effects of Metallicity. *ApJ*, 645:1102–1117, July 2006.

[159] Emily M. Levesque. Symposium Summary. In J. J. Eldridge, J. C. Bray, L. A. S. McClelland, and L. Xiao, editors, *The Lives and Death-Throes of Massive Stars*, volume 329 of *IAU Symposium*, pages 339–345, November 2017.

[160] Emily M. Levesque, Claus Leitherer, Sylvia Ekstrom, Georges Meynet, and Daniel Schaerer. The Effects of Stellar Rotation. I. Impact on the Ionizing Spectra and Integrated Properties of Stellar Populations. *ApJ*, 751(1):67, May 2012.

[161] Emily M. Levesque, Philip Massey, K. A. G. Olsen, and Bertrand Plez. The Coolest Stars in the Clouds: Unusual Red Supergiants in the Magellanic Clouds. *arXiv e-prints*, page arXiv:0708.2914, Aug 2007.

[162] Emily M. Levesque, Philip Massey, K. A. G. Olsen, Bertrand Plez, Eric Josselin, Andre Maeder, and Georges Meynet. The Effective Temperature Scale of Galactic Red Supergiants: Cool, but Not As Cool As We Thought. *ApJ*, 628(2):973–985, August 2005.

[163] Emily M. Levesque, Philip Massey, Bertrand Plez, and Knut A. G. Olsen. The Physical Properties of the Red Supergiant WOH G64: The Largest Star Known? *AJ*, 137(6):4744–4752, June 2009.

[164] L. Lindegren, J. Hernández, A. Bombrun, S. Klioner, U. Bastian, M. Ramos-Lerate, A. de Torres, H. Steidelmüller, C. Stephenson, D. Hobbs, U. Lammers, M. Biermann, R. Geyer, T. Hilger, D. Michalik, U. Stampa, P. J. McMillan, J. Castañeda, M. Clotet, G. Comoretto, M. Davidson, C. Fabricius, G. Gracia, N. C. Hambly, A. Hutton, A. Mora, J. Portell, F. van Leeuwen, U. Abbas, A. Abreu, M. Altmann, A. Andrei, E. Anglada, L. Balaguer-Núñez, C. Barache, U. Becciani, S. Bertone, L. Bianchi, S. Bouquillon, G. Bourda, T. Brüsemeister, B. Bucciarelli, D. Busonero, R. Buzzi, R. Cancelliere, T. Carlucci, P. Charlot, N. Cheek, M. Crosta, C. Crowley, J. de Bruijne, F. de Felice, R. Drimmel, P. Esquej, A. Fienga, E. Fraile, M. Gai, N. Garralda, J. J. González-Vidal, R. Guerra, M. Hauser, W. Hofmann, B. Holl, S. Jordan, M. G. Lattanzi, H. Lenhardt, S. Liao, E. Licata, T. Lister, W. Löffler, J. Marchant, J.-M. Martin-Fleitas, R. Messineo, F. Mignard, R. Morbidelli, E. Poggio, A. Riva, N. Rowell, E. Salguero, M. Sarasso, E. Sciacca, H. Siddiqui, R. L. Smart, A. Spagna, I. Steele, F. Taris, J. Torra, A. van Elteren, W. van Reeven, and A. Vecchiato. Gaia Data Release 2. The astrometric solution. *A&A*, 616:A2, August 2018.

[165] T. Linden, V. Kalogera, J. F. Sepinsky, A. Prestwich, A. Zezas, and J. S. Gallagher. The Effect of Starburst Metallicity on Bright X-ray Binary Formation Pathways. *ApJ*, 725:1984–1994, December 2010.

[166] N. R. Lomb. Least-squares frequency analysis of unequally spaced data. *Ap&SS*, 39:447–462, February 1976.

[167] G. L. Loumos and T. J. Deeming. Spurious results from Fourier analysis of data with closely spaced frequencies. *Ap&SS*, 56:285–291, July 1978.

[168] C. C. Lovekin and J. A. Guzik. Pulsations as a driver for LBV variability. *MNRAS*, 445:1766–1773, December 2014.

[169] L. B. Lucy and M. A. Sweeney. Spectroscopic binaries with circular orbits. *AJ*, 76:544–556, August 1971.

[170] A. Maeder. The relative numbers of Wolf-Rayet stars in galaxies with active star formation - The metallicity effect. *A&A*, 242:93–111, February 1991.

[171] A. Maeder, J. Lequeux, and M. Azzopardi. The numbers of red supergiants and WR stars in galaxies - an extremely sensitive indicator of chemical composition. *A&A*, 90:L17–L20, October 1980.

[172] A. Maeder and G. Meynet. New models of Wolf-Rayet stars and comparison with data in galaxies. *A&A*, 287:803–816, July 1994.

[173] A. Maeder and G. Meynet. Stellar evolution with rotation. VI. The Eddington and Omega -limits, the rotational mass loss for OB and LBV stars. *A&A*, 361:159–166, Sep 2000.

[174] A. Maeder and G. Meynet. The Evolution of Rotating Stars. *ARA&A*, 38:143–190, 2000.

[175] A. Maeder and G. Meynet. Stellar evolution with rotation. VII. . Low metallicity models and the blue to red supergiant ratio in the SMC. *A&A*, 373:555–571, July 2001.

[176] Andre Maeder and Jean-Paul Zahn. Stellar evolution with rotation. III. Meridional circulation with MU -gradients and non-stationarity. *A&A*, 334:1000–1006, June 1998.

[177] A. Mainzer, J. Bauer, T. Grav, J. Masiero, R. M. Cutri, J. Dailey, P. Eisenhardt, R. S. McMillan, E. Wright, R. Walker, R. Jedicke, T. Spahr, D. Tholen, R. Alles, R. Beck, H. Brandenburg, T. Conrow, T. Evans, J. Fowler, T. Jarrett, K. Marsh, F. Masci, H. McCallon, S. Wheelock, M. Wittman, P. Wyatt, E. DeBaun, G. Elliott, D. Elsbury, T. Gautier, IV, S. Gomillion, D. Leisawitz, C. Maleszewski, M. Micheli, and A. Wilkins. Preliminary Results from NEOWISE: An Enhancement to the Wide-field Infrared Survey Explorer for Solar System Science. *ApJ*, 731:53, April 2011.

[178] Steven R. Majewski, Ricardo P. Schiavon, Peter M. Frinchaboy, Carlos Allende Prieto, Robert Barkhouser, Dmitry Bizyaev, Basil Blank, Sophia Brunner, Adam Burton, Ricardo Carrera, S. Drew Chojnowski, Kátia Cunha, Courtney Epstein, Greg Fitzgerald, Ana E. García Pérez, Fred R. Hearty, Chuck Henderson, Jon A. Holtzman, Jennifer A. Johnson, Charles R. Lam, James E. Lawler, Paul Maseman, Szabolcs Mészáros, Matthew Nelson, Duy Coung Nguyen, David L. Nidever, Marc Pinsonneault, Matthew Shetrone, Stephen Smee, Verne V. Smith, Todd Stolberg, Michael F. Skrutskie, Eric Walker, John C. Wilson, Gail Zasowski, Friedrich Anders, Sarbani Basu, Stephane Beland, Michael R. Blanton, Jo Bovy, Joel R. Brownstein, Joleen Carlberg, William Chaplin, Cristina Chiappini, Daniel J. Eisenstein, Yvonne Elsworth, Diane Feuillet, Scott W. Fleming, Jessica Galbraith-Frew, Rafael A. García, D. Aníbal García-Hernández, Bruce A. Gillespie, Léo Girardi, James E. Gunn, Sten Hasselquist, Michael R. Hayden, Saskia Hekker, Inese Ivans, Karen Kinemuchi, Mark Klaene, Suvrath Mahadevan, Savita Mathur, Benoît Mosser, Demitri Muna, Jeffrey A. Munn, Robert C. Nichol, Robert W. O'Connell, John K. Parejko, A. C. Robin, Helio Rocha-Pinto, Matthias Schultheis, Aldo M. Serenelli, Neville Shane, Victor Silva Aguirre, Jennifer S. Sobeck, Benjamin Thompson, Nicholas W. Troup, David H. Weinberg, and Olga Zamora. The Apache Point Observatory Galactic Evolution Experiment (APOGEE). *AJ*, 154(3):94, September 2017.

[179] S. Makhija, S. Saha, S. Basak, and M. Das. Separating stars from quasars: Machine learning investigation using photometric data. *Astronomy and Computing*, 29:100313, Oct 2019.

[180] K. Malhan, R. A. Ibata, and N. F. Martin. Ghostly tributaries to the Milky Way: charting the halo's stellar streams with the Gaia DR2 catalogue. *MNRAS*, 481:3442–3455, December 2018.

[181] F. Martins and A. Palacios. A comparison of evolutionary tracks for single Galactic massive stars. *A&A*, 560:A16, December 2013.

[182] P. Massey. MASSIVE STARS IN THE LOCAL GROUP: Implications for Stellar Evolution and Star Formation. *ARA&A*, 41:15–56, 2003.

[183] P. Massey, R. T. McNeill, K. A. G. Olsen, P. W. Hodge, C. Blaha, G. H. Jacoby, R. C. Smith, and S. B. Strong. A Survey of Local Group Galaxies Currently Forming Stars. III. A Search for Luminous Blue Variables and Other Hα Emission-Line Stars. *AJ*, 134:2474–2503, December 2007.

[184] P. Massey and K. A. G. Olsen. The Evolution of Massive Stars. I. Red Supergiants in the Magellanic Clouds. *AJ*, 126:2867–2886, December 2003.

[185] P. Massey, K. A. G. Olsen, P. W. Hodge, G. H. Jacoby, R. T. McNeill, R. C. Smith, and S. B. Strong. A Survey of Local Group Galaxies Currently Forming Stars. II. UBVRI Photometry of Stars in Seven Dwarfs and a Comparison of the Entire Sample. *AJ*, 133:2393–2417, May 2007.

[186] P. Massey, K. A. G. Olsen, P. W. Hodge, S. B. Strong, G. H. Jacoby, W. Schlingman, and R. C. Smith. A Survey of Local Group Galaxies Currently Forming Stars. I. UBVRI Photometry of Stars in M31 and M33. *AJ*, 131:2478–2496, May 2006.

[187] P. Massey, B. Plez, E. M. Levesque, K. A. G. Olsen, G. C. Clayton, and E. Josselin. The Reddening of Red Supergiants: When Smoke Gets in Your Eyes. *ApJ*, 634:1286–1292, December 2005.

[188] P. Massey, D. R. Silva, E. M. Levesque, B. Plez, K. A. G. Olsen, G. C. Clayton, G. Meynet, and A. Maeder. Red Supergiants in the Andromeda Galaxy (M31). *ApJ*, 703:420–440, September 2009.

[189] Philip Massey, Kathryn F. Neugent, Trevor Z. Dorn-Wallenstein, J. J. Eldridge, E. R. Stanway, and Emily M. Levesque. Testing Evolutionary Models with Red Supergiant and Wolf-Rayet Populations. *arXiv e-prints*, page arXiv:2107.08304, July 2021.

[190] Philip Massey, Kathryn F. Neugent, and Emily M. Levesque. The evolution of massive stars: bridging the gap in the Local Group. *Philosophical Transactions of the Royal Society of London Series A*, 375:20160267, September 2017.

[191] J. C. Mauerhan, N. Smith, A. V. Filippenko, K. B. Blanchard, P. K. Blanchard, C. F. E. Casper, S. B. Cenko, K. I. Clubb, D. P. Cohen, K. L. Fuller, G. Z. Li, and J. M. Silverman. The unprecedented 2012 outburst of SN 2009ip: a luminous blue variable star becomes a true supernova. *MNRAS*, 430:1801–1810, April 2013.

[192] N. Mauron and E. Josselin. The mass-loss rates of red supergiants and the de Jager prescription. *A&A*, 526:A156, February 2011.

[193] Casey A. Meakin and David Arnett. Turbulent Convection in Stellar Interiors. I. Hydrodynamic Simulation. *ApJ*, 667(1):448–475, September 2007.

[194] G. Meynet, V. Chomienne, S. Ekström, C. Georgy, A. Granada, J. Groh, A. Maeder, P. Eggenberger, E. Levesque, and P. Massey. Impact of mass-loss on the evolution and pre-supernova properties of red supergiants. *A&A*, 575:A60, Mar 2015.

[195] G. Meynet and A. Maeder. Stellar evolution with rotation. V. Changes in all the outputs of massive star models. *A&A*, 361:101–120, September 2000.

[196] G. Meynet and A. Maeder. Stellar evolution with rotation. X. Wolf-Rayet star populations at solar metallicity. *A&A*, 404:975–990, Jun 2003.

[197] A. Miglio and J. Montalbán. Constraining fundamental stellar parameters using seismology. Application to α Centauri AB. *A&A*, 441(2):615–629, October 2005.

[198] M. Moe and R. Di Stefano. Mind Your Ps and Qs: The Interrelation between Period (P) and Mass-ratio (Q) Distributions of Binary Stars. *ApJS*, 230:15, June 2017.

[199] A. F. J. Moffat, S. V. Marchenko, L. Lefèvre, A. N. Chené, N. St-Louis, B. E. Zhilyaev, C. Aerts, H. Saio, G. A. H. Walker, J. M. Matthews, R. Kuschnig, C. Cameron, J. F. Rowe, D. B. Guenther, S. M. Rucinski, D. Sasselov, and W. W. Weiss. Pulsations Beneath the Winds: Unique Precise Photometry from MOST. In A. de Koter, L. J. Smith, and Laurens B. F. M. Waters, editors, *Mass Loss from Stars and the Evolution of Stellar Clusters*, volume 388 of *Astronomical Society of the Pacific Conference Series*, page 29, June 2008.

[200] A. F. J. Moffat, N. St-Louis, D. Carlos-Leblanc, N. D. Richardson, H. Pablo, and T. Ramiaramanantsoa. Wolf-Rayet Stars with BRITE. In G. A. Wade, D. Baade, J. A. Guzik, and R. Smolec, editors, *3rd BRITE Science Conference*, volume 8, pages 37–42, August 2018.

[201] M. H. Montgomery and D. Odonoghue. A derivation of the errors for least squares fitting to time series data. *Delta Scuti Star Newsletter*, 13:28, July 1999.

[202] Brett Michael Morris. *The Effects of Stellar Magnetic Activity and Variability on Observations of Exoplanets*. PhD thesis, University of Washington, January 2019.

[203] Donald C. Morton. Mass Loss from Three OB Supergiants in Orion. *ApJ*, 150:535, November 1967.

[204] Lawrence Mosley. In *A balanced approach to the multi-class imbalance problem*, 2013.

[205] K. F. Neugent, E. M. Levesque, and P. Massey. Binary Red Supergiants: A New Method for Detecting B-type Companions. *AJ*, 156:225, November 2018.

[206] K. F. Neugent and P. Massey. The Wolf-Rayet Content of M33. *ApJ*, 733:123, June 2011.

[207] K. F. Neugent and P. Massey. The Close Binary Frequency of Wolf-Rayet Stars as a Function of Metallicity in M31 and M33. *ApJ*, 789:10, July 2014.

[208] K. F. Neugent, P. Massey, and C. Georgy. The Wolf-Rayet Content of M31. *ApJ*, 759:11, November 2012.

[209] K. F. Neugent, P. Massey, and N. Morrell. A Modern Search for Wolf-Rayet Stars in the Magellanic Clouds. IV. A Final Census. *ArXiv e-prints*, July 2018.

[210] Kathryn F. Neugent. The Red Supergiant Binary Fraction as a Function of Metallicity in M31 and M33. *ApJ*, 908(1):87, February 2021.

[211] Kathryn F. Neugent, Emily M. Levesque, Philip Massey, and Nidia I. Morrell. Binary Red Supergiants. II. Discovering and Characterizing B-type Companions. *ApJ*, 875(2):124, Apr 2019.

[212] Kathryn F. Neugent, Emily M. Levesque, Philip Massey, Nidia I. Morrell, and Maria R. Drout. The Red Supergiant Binary Fraction of the Large Magellanic Cloud. *arXiv e-prints*, page arXiv:2007.15852, July 2020.

[213] Kathryn F. Neugent, Philip Massey, Cyril Georgy, Maria R. Drout, Michael Mommert, Emily M. Levesque, Georges Meynet, and Sylvia Ekström. The Luminosity Function of Red Supergiants in M31. *ApJ*, 889(1):44, January 2020.

[214] Kathryn F. Neugent, Philip Massey, Brian Skiff, Maria R. Drout, Georges Meynet, and Knut A. G. Olsen. Yellow Supergiants in the Small Magellanic Cloud: Putting Current Evolutionary Theory to the Test. *ApJ*, 719(2):1784–1795, August 2010.

[215] Kathryn F. Neugent, Philip Massey, Brian Skiff, and Georges Meynet. Yellow and Red Supergiants in the Large Magellanic Cloud. *ApJ*, 749:177, April 2012.

[216] H. Nieuwenhuijzen and C. de Jager. Atmospheric accelerations and the stability of dynamic supergiant atmospheres. *A&A*, 302:811, October 1995.

[217] A. Noels, N. Grevesse, P. Magain, C. Neuforge, A. Baglin, and Y. Lebreton. Calibration of the alpha Centauri system : metallicity and age. *A&A*, 247:247, July 1991.

[218] F. Ochsenbein, P. Bauer, and J. Marcout. The VizieR database of astronomical catalogues. *A&AS*, 143:23–32, Apr 2000.

[219] Stanley P. Owocki. *Instabilities in the Envelopes and Winds of Very Massive Stars*, volume 412, page 113. 2015.

[220] B. Paczynski. Common Envelope Binaries. In P. Eggleton, S. Mitton, and J. Whelan, editors, *Structure and Evolution of Close Binary Systems*, volume 73 of *IAU Symposium*, page 75, 1976.

[221] J. Papaloizou and J. E. Pringle. Non-radial oscillations of rotating stars and their relevance to the short-period oscillations of cataclysmic variables. *Monthly Notices of the Royal Astronomical Society*, 182(3):423–442, 1978.

[222] Taeyoung Park, Vinay L. Kashyap, Aneta Siemiginowska, David A. van Dyk, Andreas Zezas, Craig Heinke, and Bradford J. Wargelin. Bayesian Estimation of Hardness Ratios: Modeling and Computations. *ApJ*, 652(1):610–628, Nov 2006.

[223] B. Paxton, L. Bildsten, A. Dotter, F. Herwig, P. Lesaffre, and F. Timmes. Modules for Experiments in Stellar Astrophysics (MESA). *ApJS*, 192:3, January 2011.

[224] B. Paxton, M. Cantiello, P. Arras, L. Bildsten, E. F. Brown, A. Dotter, C. Mankovich, M. H. Montgomery, D. Stello, F. X. Timmes, and R. Townsend. Modules for Experiments in Stellar Astrophysics (MESA): Planets, Oscillations, Rotation, and Massive Stars. *ApJS*, 208:4, September 2013.

[225] B. Paxton, P. Marchant, J. Schwab, E. B. Bauer, L. Bildsten, M. Cantiello, L. Dessart, R. Farmer, H. Hu, N. Langer, R. H. D. Townsend, D. M. Townsley, and F. X. Timmes. Modules for Experiments in Stellar Astrophysics (MESA): Binaries, Pulsations, and Explosions. *ApJS*, 220:15, September 2015.

[226] Bill Paxton, Josiah Schwab, Evan B. Bauer, Lars Bildsten, Sergei Blinnikov, Paul Duffell, R. Farmer, Jared A. Goldberg, Pablo Marchant, Elena Sorokina, Anne Thoul, Richard H. D. Townsend, and F. X. Timmes. Modules for Experiments in Stellar Astrophysics (MESA): Convective Boundaries, Element Diffusion, and Massive Star Explosions. *ApJS*, 234(2):34, February 2018.

[227] Bill Paxton, R. Smolec, Josiah Schwab, A. Gautschy, Lars Bildsten, Matteo Cantiello, Aaron Dotter, R. Farmer, Jared A. Goldberg, Adam S. Jermyn, S. M. Kanbur, Pablo Marchant, Anne Thoul, Richard H. D. Townsend, William M. Wolf, Michael Zhang, and F. X. Timmes. Modules for Experiments in Stellar Astrophysics (MESA): Pulsating Variable Stars, Rotation, Convective Boundaries, and Energy Conservation. *ApJS*, 243(1):10, July 2019.

[228] Cecilia Helena Payne. *Stellar Atmospheres; a Contribution to the Observational Study of High Temperature in the Reversing Layers of Stars.* PhD thesis, RADCLIFFE COLLEGE., January 1925.

[229] May G. Pedersen, Sowgata Chowdhury, Cole Johnston, Dominic Bowman, Conny Aerts, Gerald Handler, Peter De Cat, Coralie Neiner, Alexandre David-Uraz, Derek Buzasi, Andrew Tkachenko, Sergio Simon-Diaz, Ehsan Moravveji, James Sikora, Giovanni M. Mirouh, Catherine C. Lovekin, Matteo Cantiello, Jadwiga Daszynska-Daszkiewicz, Andrzej Pigulski, Roland K. Vanderspek, and George R. Ricker. Diverse Variability of O and B Stats Revealed from 2-minute Cadence Light Curves in Sectors 1 and 3 of the TESS Mission: Selection of an Asteroseismic Sample. *arXiv e-prints*, page arXiv:1901.07576, January 2019.

[230] J. Perdang. CA Simulations of 2D Stellar Atmosphere Pulsations. In M. Goupil, Z. Koláth, N. Nardetto, and P. Kervella, editors, *EAS Publications Series*, volume 38 of *EAS Publications Series*, pages 43–62, 2009.

[231] G. J. Peters, D. R. Gies, E. D. Grundstrom, and M. V. McSwain. Detection of a Hot Subdwarf Companion to the Be Star FY Canis Majoris. *ApJ*, 686:1280–1291, October 2008.

[232] G. J. Peters, T. D. Pewett, D. R. Gies, Y. N. Touhami, and E. D. Grundstrom. Far-ultraviolet Detection of the Suspected Subdwarf Companion to the Be Star 59 Cygni. *ApJ*, 765:2, March 2013.

[233] M. Pinsonneault. Mixing in Stars. *ARA&A*, 35:557–605, January 1997.

[234] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical recipes in FORTRAN. The art of scientific computing*. Cambridge University Press, 1992.

[235] J. Puls, J. S. Vink, and F. Najarro. Mass loss from hot massive stars. *A&A Rv*, 16:209–325, December 2008.

[236] D. Raghavan, H. A. McAlister, T. J. Henry, D. W. Latham, G. W. Marcy, B. D. Mason, D. R. Gies, R. J. White, and T. A. ten Brummelaar. A Survey of Stellar Families: Multiplicity of Solar-type Stars. *ApJS*, 190:1–42, September 2010.

[237] T. Ramiaramanantsoa, A. F. J. Moffat, R. Harmon, R. Ignace, N. St-Louis, D. Vanbeveren, T. Shenar, H. Pablo, N. D. Richardson, I. D. Howarth, I. R. Stevens, C. Piaulet, L. St-Jean, T. Eversberg, A. Pigulski, A. Popowicz, R. Kuschnig, E. Zocłońska, B. Buysschaert, G. Handler, W. W. Weiss, G. A. Wade, S. M. Rucinski, K. Zwintz, P. Luckas, B. Heathcote, P. Cacella, J. Powles, M. Locke, T. Bohlsen, A.-N. Chené, B. Miszalski, W. L. Waldron, M. M. Kotze, E. J. Kotze, and T. Böhm. BRITE-Constellation high-precision time-dependent photometry of the early O-type supergiant ζ Puppis unveils the photospheric drivers of its small- and large-scale wind structures. *MNRAS*, 473:5532–5569, February 2018.

[238] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. 2006.

[239] D. Reimers. Circumstellar absorption lines and mass loss from red giants. *Memoires of the Societe Royale des Sciences de Liege*, 8:369–382, Jan 1975.

[240] D. Reimers. On the absolute scale of mass-loss in red giants. I. Circumstellar absorption lines in the spectrum of the visual companion of alpha ^{1}Her. *A&A*, 61:217–224, Oct 1977.

[241] N. D. Richardson and A. Mehner. The 2018 Census of Luminous Blue Variables in the Local Group. *Research Notes of the American Astronomical Society*, 2(3):121, July 2018.

[242] G. R. Ricker, J. N. Winn, R. Vanderspek, D. W. Latham, G. Á. Bakos, J. L. Bean, Z. K. Berta-Thompson, T. M. Brown, L. Buchhave, N. R. Butler, R. P. Butler, W. J. Chaplin, D. Charbonneau, J. Christensen-Dalsgaard, M. Clampin, D. Deming, J. Doty, N. De Lee, C. Dressing, E. W. Dunham, M. Endl, F. Fressin, J. Ge, T. Henning, M. J. Holman, A. W. Howard, S. Ida, J. M. Jenkins, G. Jernigan, J. A. Johnson, L. Kaltenegger, N. Kawai, H. Kjeldsen, G. Laughlin, A. M. Levine, D. Lin, J. J. Lissauer, P. MacQueen, G. Marcy, P. R. McCullough, T. D. Morton, N. Narita, M. Paegert, E. Palle, F. Pepe, J. Pepper, A. Quirrenbach, S. A. Rinehart, D. Sasselov, B. Sato, S. Seager, A. Sozzetti, K. G. Stassun, P. Sullivan, A. Szentgyorgyi, G. Torres, S. Udry, and J. Villasenor. Transiting Exoplanet Survey Satellite (TESS). *Journal of Astronomical Telescopes, Instruments, and Systems*, 1(1):014003, January 2015.

[243] Thomas Rivinius, Alex C. Carciofi, and Christophe Martayan. Classical Be stars. Rapidly rotating B stars with viscous Keplerian decretion disks. *A&A Rv*, 21:69, October 2013.

[244] C. K. Rosslowe and P. A. Crowther. Spatial distribution of Galactic Wolf-Rayet stars and implications for the global population. *MNRAS*, 447:2322–2347, March 2015.

[245] M. N. Saha. Ionisation in the Solar Chromosphere. *Nature*, 105(2634):232–233, April 1920.

[246] M. N. Saha. On a Physical Theory of Stellar Spectra. *Proceedings of the Royal Society of London Series A*, 99(697):135–153, May 1921.

[247] H. Saio, D. W. Kurtz, S. J. Murphy, V. L. Antoci, and U. Lee. Theory and evidence of global Rossby waves in upper main-sequence stars: r-mode oscillations in many Kepler stars. *MNRAS*, 474:2774–2786, February 2018.

[248] H. Saio, R. Kuschnig, A. Gautschy, C. Cameron, G. A. H. Walker, J. M. Matthews, D. B. Guenther, A. F. J. Moffat, S. M. Rucinski, D. Sasselov, and W. W. Weiss. MOST Detects

g- and p-Modes in the B Supergiant HD 163899 (B2 Ib/II). *ApJ*, 650:1111–1118, October 2006.

[249] Hideyuki Saio, Cyril Georgy, and Georges Meynet. Evolution of blue supergiants and α Cygni variables: puzzling CNO surface abundances. *MNRAS*, 433(2):1246–1257, August 2013.

[250] Maurizio Salaris and Santi Cassisi. Chemical element transport in stellar evolution models. *Royal Society Open Science*, 4(8):170192, August 2017.

[251] E. E. Salpeter. The Luminosity Function and Stellar Evolution. *ApJ*, 121:161, January 1955.

[252] H. Sana, A. de Koter, S. E. de Mink, P. R. Dunstall, C. J. Evans, V. Hénault-Brunet, J. Maíz Apellániz, O. H. Ramírez-Agudelo, W. D. Taylor, N. R. Walborn, J. S. Clark, P. A. Crowther, A. Herrero, M. Gieles, N. Langer, D. J. Lennon, and J. S. Vink. The VLT-FLAMES Tarantula Survey. VIII. Multiplicity properties of the O-type star population. *A&A*, 550:A107, February 2013.

[253] H. Sana, S. E. de Mink, A. de Koter, N. Langer, C. J. Evans, M. Gieles, E. Gosset, R. G. Izzard, J.-B. Le Bouquin, and F. R. N. Schneider. Binary Interaction Dominates the Evolution of Massive Stars. *Science*, 337:444, July 2012.

[254] H. Sana, J.-B. Le Bouquin, S. Lacour, J.-P. Berger, G. Duvert, L. Gauchet, B. Norris, J. Olofsson, D. Pickel, G. Zins, O. Absil, A. de Koter, K. Kratter, O. Schnurr, and H. Zinnecker. Southern Massive Stars at High Angular Resolution: Observational Campaign and Companion Detection. *ApJS*, 215:15, November 2014.

[255] N. Sanduleak. A deep objective-prism survey for Large Magellanic Cloud members. *Contributions from the Cerro Tololo Inter-American Observatory*, 89, January 1970.

[256] J. D. Scargle. Studies in astronomical time series analysis. II - Statistical aspects of spectral analysis of unevenly spaced data. *ApJ*, 263:835–853, December 1982.

[257] Gideon Schwarz. Estimating the dimension of a model. *Ann. Statist.*, 6(2):461–464, 03 1978.

[258] A. Schwarzenberg-Czerny. *An astronomer's guide to period searching*, volume 292 of *Astronomical Society of the Pacific Conference Series*, page 383. 2003.

[259] Martin Schwarzschild. *Structure and evolution of the stars.* 1958.

[260] P. Scicluna, R. Siebenmorgen, R. Wesson, J. A. D. L. Blommaert, M. Kasper, N. V. Voshchinnikov, and S. Wolf. Large dust grains in the wind of VY Canis Majoris. *A&A*, 584:L10, December 2015.

[261] Dinesh Shenoy, Roberta M. Humphreys, Terry J. Jones, Massimo Marengo, Robert D. Gehrz, L. Andrew Helton, William F. Hoffmann, Andrew J. Skemer, and Philip M. Hinz. Searching for Cool Dust in the Mid-to-far Infrared: The Mass-loss Histories of the Hypergiants μ Cep, VY CMa, IRC+10420, and ρ Cas. *AJ*, 151(3):51, March 2016.

[262] Hiromoto Shibahashi and Simon J. Murphy. Asteroseismology as a new window to statistics of binaries. *Bulletin de la Societe Royale des Sciences de Liege*, 88:147–162, December 2019.

[263] M. F. Skrutskie, R. M. Cutri, R. Stiening, M. D. Weinberg, S. Schneider, J. M. Carpenter, C. Beichman, R. Capps, T. Chester, J. Elias, J. Huchra, J. Liebert, C. Lonsdale, D. G. Monet, S. Price, P. Seitzer, T. Jarrett, J. D. Kirkpatrick, J. E. Gizis, E. Howard, T. Evans, J. Fowler, L. Fullmer, R. Hurt, R. Light, E. L. Kopan, K. A. Marsh, H. L. McCallon, R. Tam, S. Van Dyk, and S. Wheelock. The Two Micron All Sky Survey (2MASS). *AJ*, 131(2):1163–1183, Feb 2006.

[264] C. L. Slesnick, L. A. Hillenbrand, and P. Massey. The Star Formation History and Mass Function of the Double Cluster h and χ Persei. *ApJ*, 576:880–893, September 2002.

[265] S. J. Smartt, J. J. Eldridge, R. M. Crockett, and J. R. Maund. The death of massive stars - I. Observational constraints on the progenitors of Type II-P supernovae. *MNRAS*, 395:1409–1437, May 2009.

[266] N. Smith. Mass Loss: Its Effect on the Evolution and Fate of High-Mass Stars. *ARA&A*, 52:487–528, August 2014.

[267] N. Smith, M. Aghakhanloo, J. W. Murphy, K. G. Stassun, M. R. Drout, and J. H. Groh. On the Gaia DR2 distances for Galactic Luminous Blue Variables. *ArXiv e-prints*, May 2018.

[268] N. Smith and R. Tombleson. Luminous blue variables are antisocial: their isolation implies that they are kicked mass gainers in binary evolution. *MNRAS*, 447:598–617, February 2015.

[269] Jr. Snow, T. P. and E. B. Jenkins. A catalog of 0.2 Å resolution far-ultraviolet stellar spectra measured with Copernicus. *ApJS*, 33:269–360, March 1977.

[270] Monika D. Soraisam, Lars Bildsten, Maria R. Drout, Thomas A. Prince, Thomas Kupfer, Frank Masci, Russ R. Laher, and Shrinivas R. Kulkarni. Variability of Massive Stars in M31 from the Palomar Transient Factory. *ApJ*, 893(1):11, April 2020.

[271] J. R. Sowell. A survey of Balmer-line profiles and IRAS fluxes in forty yellow supergiants. *AJ*, 100:834–848, September 1990.

[272] S. Specter, F. Bellisle, S. Hémery-Véron, P. Fiquet, F. Bornet, and G. Slama. Reducing ice cream energy density does not condition decreased acceptance or engender compensation following repeated exposure. *European Journal of Clinical Nutrition*, 52:703–710, 1998.

[273] D. Spergel, N. Gehrels, J. Breckinridge, M. Donahue, A. Dressler, B. S. Gaudi, T. Greene, O. Guyon, C. Hirata, J. Kalirai, N. J. Kasdin, W. Moos, S. Perlmutter, M. Postman, B. Rauscher, J. Rhodes, Y. Wang, D. Weinberg, J. Centrella, W. Traub, C. Baltay, J. Colbert, D. Bennett, A. Kiessling, B. Macintosh, J. Merten, M. Mortonson, M. Penny, E. Rozo,

D. Savransky, K. Stapelfeldt, Y. Zu, C. Baker, E. Cheng, D. Content, J. Dooley, M. Foote, R. Goullioud, K. Grady, C. Jackson, J. Kruk, M. Levine, M. Melton, C. Peddie, J. Ruffa, and S. Shaklan. WFIRST-2.4: What Every Astronomer Should Know. *arXiv e-prints*, page arXiv:1305.5425, May 2013.

[274] H. C. Spruit. Semiconvection: theory. *A&A*, 552:A76, April 2013.

[275] V. Stanishev, Z. Kraicheva, H. M. J. Boffin, and V. Genkov. PX Andromedae: Superhumps and variable eclipse depth. *A&A*, 394:625–632, November 2002.

[276] E. R. Stanway and J. J. Eldridge. Re-evaluating old stellar populations. *MNRAS*, 479:75–93, September 2018.

[277] E. R. Stanway, J. J. Eldridge, and G. D. Becker. Stellar population effects on the inferred photon density at reionization. *MNRAS*, 456:485–499, February 2016.

[278] E. R. Stanway, J. J. Eldridge, and A. A. Chrimes. Binary fraction indicators in resolved stellar populations and supernova-type ratios. *MNRAS*, 497(2):2201–2212, July 2020.

[279] E. I. Staritsin. Turbulent entrainment at the boundaries of the convective cores of main-sequence stars. *Astronomy Reports*, 57(5):380–390, May 2013.

[280] K. G. Stassun, R. J. Oelkers, J. Pepper, M. Paegert, N. De Lee, G. Torres, D. W. Latham, S. Charpinet, C. D. Dressing, D. Huber, S. R. Kane, S. Lépine, A. Mann, P. S. Muirhead, B. Rojas-Ayala, R. Silvotti, S. W. Fleming, A. Levine, and P. Plavchan. The TESS Input Catalog and Candidate Target List. *AJ*, 156:102, September 2018.

[281] J. E. Steiner and A. S. Oliveira. The qWR star HD 45166. I. Observations and system parameters. *A&A*, 444:895–904, December 2005.

[282] Richard B. Stothers and Chao-wen Chin. Yellow Hypergiants as Dynamically Unstable Post-Red Supergiant Stars. *ApJ*, 560(2):934–936, October 2001.

[283] T. Sukhbold, T. Ertl, S. E. Woosley, J. M. Brown, and H.-T. Janka. Core-collapse Supernovae from 9 to 120 Solar Masses Based on Neutrino-powered Explosions. *ApJ*, 821:38, April 2016.

[284] T. Sukhbold, S. E. Woosley, and A. Heger. A High-resolution Study of Presupernova Core Structure. *ApJ*, 860:93, June 2018.

[285] Tuguldur Sukhbold and Scott Adams. Missing red supergiants and carbon burning. *MNRAS*, 492(2):2578–2587, February 2020.

[286] Mikael Sunnåker, Alberto Giovanni Busetto, Elina Numminen, Jukka Corander, Matthieu Foll, and Christophe Dessimoz. Approximate Bayesian Computation. *PLoS Computational Biology*, 9(1):e1002803, January 2013.

[287] R. J. Sylvester, C. J. Skinner, and M. J. Barlow. Silicate and hydrocarbon emission from Galactic M supergiants. *MNRAS*, 301(4):1083–1094, Dec 1998.

[288] K. Szatmary, J. Vinko, and J. Gal. Application of wavelet analysis in variable star research. I. Properties of the wavelet map of simulated variable star light curves. *A&AS*, 108:377–394, December 1994.

[289] Karoly Szatmary and Jozsef Vinko. Periodicities of the light curve of the semiregular variable star Y Lyncis. *MNRAS*, 256:321–328, May 1992.

[290] S. Talon. Transport Processes in Stars: Diffusion, Rotation, Magnetic fields and Internal Waves. In C. Charbonnel and J. P. Zahn, editors, *EAS Publications Series*, volume 32 of *EAS Publications Series*, pages 81–130, November 2008.

[291] S. Talon, J. P. Zahn, A. Maeder, and G. Meynet. Rotational mixing in early-type stars: the main-sequence evolution of a $9M_{sun-}$ star. *A&A*, 322:209–217, June 1997.

[292] The Astropy Collaboration, A. M. Price-Whelan, B. M. Sipőcz, H. M. Günther, P. L. Lim, S. M. Crawford, S. Conseil, D. L. Shupe, M. W. Craig, N. Dencheva, A. Ginsburg, J. T. Van-

derPlas, L. D. Bradley, D. Pérez-Suárez, M. de Val-Borro, T. L. Aldcroft, K. L. Cruz, T. P. Robitaille, E. J. Tollerud, C. Ardelean, T. Babej, M. Bachetti, A. V. Bakanov, S. P. Bamford, G. Barentsen, P. Barmby, A. Baumbach, K. L. Berry, F. Biscani, M. Boquien, K. A. Bostroem, L. G. Bouma, G. B. Brammer, E. M. Bray, H. Breytenbach, H. Buddelmeijer, D. J. Burke, G. Calderone, J. L. Cano Rodríguez, M. Cara, J. V. M. Cardoso, S. Cheedella, Y. Copin, D. Crichton, D. DÁvella, C. Deil, É. Depagne, J. P. Dietrich, A. Donath, M. Droettboom, N. Earl, T. Erben, S. Fabbro, L. A. Ferreira, T. Finethy, R. T. Fox, L. H. Garrison, S. L. J. Gibbons, D. A. Goldstein, R. Gommers, J. P. Greco, P. Greenfield, A. M. Groener, F. Grollier, A. Hagen, P. Hirst, D. Homeier, A. J. Horton, G. Hosseinzadeh, L. Hu, J. S. Hunkeler, Ž. Ivezić, A. Jain, T. Jenness, G. Kanarek, S. Kendrew, N. S. Kern, W. E. Kerzendorf, A. Khvalko, J. King, D. Kirkby, A. M. Kulkarni, A. Kumar, A. Lee, D. Lenz, S. P. Littlefair, Z. Ma, D. M. Macleod, M. Mastropietro, C. McCully, S. Montagnac, B. M. Morris, M. Mueller, S. J. Mumford, D. Muna, N. A. Murphy, S. Nelson, G. H. Nguyen, J. P. Ninan, M. Nöthe, S. Ogaz, S. Oh, J. K. Parejko, N. Parley, S. Pascual, R. Patil, A. A. Patil, A. L. Plunkett, J. X. Prochaska, T. Rastogi, V. Reddy Janga, J. Sabater, P. Sakurikar, M. Seifert, L. E. Sherbert, H. Sherwood-Taylor, A. Y. Shih, J. Sick, M. T. Silbiger, S. Singanamalla, L. P. Singer, P. H. Sladen, K. A. Sooley, S. Sornarajah, O. Streicher, P. Teuben, S. W. Thomas, G. R. Tremblay, J. E. H. Turner, V. Terrón, M. H. van Kerkwijk, A. de la Vega, L. L. Watkins, B. A. Weaver, J. B. Whitmore, J. Woillez, and V. Zabalza. The Astropy Project: Building an inclusive, open-science project and status of the v2.0 core package. *ArXiv e-prints*, January 2018.

[293] Susan E. Thompson, Mark Everett, Fergal Mullally, Thomas Barclay, Steve B. Howell, Martin Still, Jason Rowe, Jessie L. Christiansen, Donald W. Kurtz, Kelly Hambleton, Joseph D. Twicken, Khadeejah A. Ibrahim, and Bruce D. Clarke. A Class of Eccentric Binaries with Dynamic Tidal Distortions Discovered with Kepler. *ApJ*, 753(1):86, July 2012.

[294] K. S. Thorne and A. N. Zytkow. Red giants and supergiants with degenerate neutron cores. *ApJL*, 199:L19–L24, July 1975.

[295] K. S. Thorne and A. N. Zytkow. Stars with degenerate neutron cores. I. Structure of equilibrium models. *ApJ*, 212:832–858, Mar 1977.

[296] Anne A. Thoul, John N. Bahcall, and Abraham Loeb. Element Diffusion in the Solar Interior. *ApJ*, 421:828, February 1994.

[297] A. Tkachenko, C. Aerts, A. Yakushechkin, J. Debosscher, P. Degroote, S. Bloemen, P. I. Pápics, B. L. de Vries, R. Lombaert, M. Hrudkova, Y. Frémat, G. Raskin, and H. Van Winckel. Detection of a large sample of γ Doradus stars from Kepler space photometry and high-resolution ground-based spectroscopy. *A&A*, 556:A52, August 2013.

[298] Regner Trampedach, Robert F. Stein, Jørgen Christensen-Dalsgaard, Åke Nordlund, and Martin Asplund. Improvements to stellar structure models, based on a grid of 3D convection simulations - II. Calibrating the mixing-length formulation. *MNRAS*, 445(4):4366–4384, December 2014.

[299] K. Ulaczyk, M. K. Szymański, A. Udalski, M. Kubiak, G. Pietrzyński, I. Soszyński, Ł. Wyrzykowski, R. Poleski, W. Gieren, A. R. Walker, and A. Garcia-Varela. Variable Stars from the OGLE-III Shallow Survey in the Large Magellanic Cloud. *AcA*, 63(2):159–179, June 2013.

[300] Roger K. Ulrich. Thermohaline Convection in Stellar Interiors. *ApJ*, 172:165, February 1972.

[301] S. van den Bergh. The Galaxies of the Local Group. *JRASC*, 62:219, October 1968.

[302] Sidney van den Bergh. The Brightest Red and Blue Stars in Galaxies. *ApJL*, 183:L123, Aug 1973.

[303] K. A. van der Hucht. The VIIth catalogue of galactic Wolf-Rayet stars. *NewAR*, 45:135–232, February 2001.

[304] D. A. van Dyk, A. Connors, V. L. Kashyap, and A. Siemiginowska. Analysis of Energy Spectra with Low Photon Counts via Bayesian Posterior Simulation. *ApJ*, 548:224–243, February 2001.

[305] A. M. van Genderen. S Doradus variables in the Galaxy and the Magellanic Clouds. *A&A*, 366:508–531, February 2001.

[306] A. M. van Genderen, A. Jones, and C. Sterken. Light variations of alpha Cygni variables in the Magellanic Clouds. *Journal of Astronomical Data*, 12:4, December 2006.

[307] A. M. van Genderen, C. Sterken, and M. de Groot. New discoveries on the S DOR phenomenon based on an investigation of the photometric history of the variables AG Car, S Dor and η Car. *A&A*, 318:81–98, February 1997.

[308] J. Th. van Loon, M. R. L. Cioni, A. A. Zijlstra, and C. Loup. An empirical formula for the mass-loss rates of dust-enshrouded red supergiants and oxygen-rich Asymptotic Giant Branch stars. *A&A*, 438(1):273–289, July 2005.

[309] J. Th. van Loon, M. A. T. Groenewegen, A. de Koter, N. R. Trams, L. B. F. M. Waters, A. A. Zijlstra, P. A. Whitelock, and C. Loup. Mass-loss rates and luminosity functions of dust-enshrouded AGB stars and red supergiants in the LMC. *A&A*, 351:559–572, Nov 1999.

[310] J. L. van Saders, T. Ceillier, T. S. Metcalfe, V. Silva Aguirre, M. H. Pinsonneault, R. A. García, S. Mathur, and G. R. Davies. Weakened magnetic braking as the origin of anomalously rapid rotation in old field stars. *Nature*, 529:181–184, January 2016.

[311] D. Vanbeveren and P. S. Conti. On the binary frequency distribution and evolution of Wolf-Rayet stars. *A&A*, 88:230–239, August 1980.

[312] S. Vaughan. A simple test for periodic signals in red noise. *A&A*, 431:391–403, February 2005.

[313] K. A. Venn. CNO Abundances and the Evolutionary Status of Galactic, A-Type Supergiants. *ApJ*, 449:839, August 1995.

[314] Maxime Viallet, Casey Meakin, David Arnett, and Miroslav Mocák. Turbulent Convection in Stellar Interiors. III. Mean-field Analysis and Stratification Effects. *ApJ*, 769(1):1, May 2013.

[315] J. S. Vink, A. de Koter, and H. J. G. L. M. Lamers. Mass-loss predictions for O and B stars as a function of metallicity. *A&A*, 369:574–588, April 2001.

[316] H. von Zeipel. The radiative equilibrium of a slightly oblate rotating star. *MNRAS*, 84:684–701, June 1924.

[317] N. R. Walborn, R. C. Gamen, N. I. Morrell, R. H. Barbá, E. Fernández Lajús, and R. Angeloni. Active Luminous Blue Variables in the Large Magellanic Cloud. *AJ*, 154:15, July 2017.

[318] M. F. Walker. Two-color composite photographs of M 33. *AJ*, 69:744, November 1964.

[319] Joseph J. Walmswell and John J. Eldridge. Circumstellar dust as a solution to the red supergiant supernova progenitor problem. *MNRAS*, 419(3):2054–2062, January 2012.

[320] L. Wang, D. R. Gies, and G. J. Peters. Detection of the Ultraviolet Spectrum of the Hot Subdwarf Companion of 60 Cygni (B1 Ve) from a Survey of IUE Spectra of Be Stars. *ApJ*, 843:60, July 2017.

[321] L. B. F. M. Waters. Circumstellar Dust in Massive Stars. In C. Leitherer, P. D. Bennett, P. W. Morris, and J. T. Van Loon, editors, *Hot and Cool: Bridging Gaps in Massive Star Evolution*, volume 425 of *Astronomical Society of the Pacific Conference Series*, page 267, June 2010.

[322] M. Wenger, F. Ochsenbein, D. Egret, P. Dubois, F. Bonnarel, S. Borde, F. Genova, G. Jasniewicz, S. Laloë, S. Lesteven, and R. Monier. The SIMBAD astronomical database. The CDS reference database for astronomical objects. *A&AS*, 143:9–22, Apr 2000.

[323] L. A. Willson. Mass loss from cool stars. In Mario Livio, editor, *Unsolved Problems in Stellar Evolution*, volume 12, page 227, January 2000.

[324] E. L. Wright, P. R. M. Eisenhardt, A. K. Mainzer, M. E. Ressler, R. M. Cutri, T. Jarrett, J. D. Kirkpatrick, D. Padgett, R. S. McMillan, M. Skrutskie, S. A. Stanford, M. Cohen, R. G. Walker, J. C. Mather, D. Leisawitz, T. N. Gautier, III, I. McLean, D. Benford, C. J. Lonsdale, A. Blain, B. Mendez, W. R. Irace, V. Duval, F. Liu, D. Royer, I. Heinrichsen, J. Howard, M. Shannon, M. Kendall, A. L. Walsh, M. Larsen, J. G. Cardon, S. Schick, M. Schwalm, M. Abid, B. Fabinsky, L. Naes, and C.-W. Tsai. The Wide-field Infrared Survey Explorer (WISE): Mission Description and Initial On-orbit Performance. *AJ*, 140:1868–1881, December 2010.

[325] Maosheng Xiang, Hans-Walter Rix, Yuan-Sen Ting, Rolf-Peter Kudritzki, Charlie Conroy, Eleonora Zari, Jian-Rong Shi, Norbert Przybilla, Maria Ramirez-Tannus, Andrew Tkachenko, Sarah Gebruers, and Xiao-Wei Liu. Stellar labels for hot stars from low-resolution spectra - I. the HotPayne method and results for 330,000 stars from LAMOST DR6. *arXiv e-prints*, page arXiv:2108.02878, August 2021.

[326] Ming Yang, Alceste Z. Bonanos, Bi-Wei Jiang, Jian Gao, Panagiotis Gavras, Grigoris Maravelias, Yi Ren, Shu Wang, Meng-Yao Xue, Frank Tramper, Zoi T. Spetsieri, and Ektoras Pouliasis. Evolved massive stars at low-metallicity. I. A source catalog for the Small Magellanic Cloud. *A&A*, 629:A91, September 2019.

[327] Ming Yang, Alceste Z. Bonanos, Bi-Wei Jiang, Jian Gao, Meng-Yao Xue, Shu Wang, Man I. Lam, Zoi T. Spetsieri, Yi Ren, and Panagiotis Gavras. Red supergiant stars in the Large Magellanic Cloud. II. Infrared properties and mid-infrared variability. *A&A*, 616:A175, Sep 2018.

[328] D. G. York, J. Adelman, J. E. Anderson, Jr., S. F. Anderson, J. Annis, N. A. Bahcall, J. A. Bakken, R. Barkhouser, S. Bastian, E. Berman, W. N. Boroski, S. Bracker, C. Briegel, J. W. Briggs, J. Brinkmann, R. Brunner, S. Burles, L. Carey, M. A. Carr, F. J. Castander, B. Chen, P. L. Colestock, A. J. Connolly, J. H. Crocker, I. Csabai, P. C. Czarapata, J. E. Davis, M. Doi, T. Dombeck, D. Eisenstein, N. Ellman, B. R. Elms, M. L. Evans, X. Fan, G. R. Federwitz, L. Fiscelli, S. Friedman, J. A. Frieman, M. Fukugita, B. Gillespie, J. E. Gunn, V. K. Gurbani, E. de Haas, M. Haldeman, F. H. Harris, J. Hayes, T. M. Heckman, G. S. Hennessy, R. B. Hindsley, S. Holm, D. J. Holmgren, C.-h. Huang, C. Hull, D. Husby, S.-I. Ichikawa, T. Ichikawa, Ž. Ivezić, S. Kent, R. S. J. Kim, E. Kinney, M. Klaene, A. N. Kleinman, S. Kleinman, G. R. Knapp, J. Korienek, R. G. Kron, P. Z. Kunszt, D. Q. Lamb, B. Lee, R. F. Leger, S. Limmongkol, C. Lindenmeyer, D. C. Long, C. Loomis, J. Loveday, R. Lucinio, R. H. Lupton, B. MacKinnon, E. J. Mannery, P. M. Mantsch, B. Margon, P. McGehee, T. A. McKay, A. Meiksin, A. Merelli, D. G. Monet, J. A. Munn, V. K. Narayanan, T. Nash, E. Neilsen, R. Neswold, H. J. Newberg, R. C. Nichol, T. Nicinski, M. Nonino, N. Okada, S. Okamura, J. P. Ostriker, R. Owen, A. G. Pauls, J. Peoples, R. L. Peterson, D. Petravick, J. R. Pier, A. Pope, R. Pordes, A. Prosapio, R. Rechenmacher, T. R. Quinn, G. T. Richards, M. W. Richmond, C. H. Rivetta, C. M. Rockosi, K. Ruthmansdorfer, D. Sandford, D. J. Schlegel, D. P. Schneider, M. Sekiguchi, G. Sergey, K. Shimasaku, W. A. Siegmund, S. Smee, J. A. Smith, S. Snedden, R. Stone, C. Stoughton, M. A. Strauss, C. Stubbs, M. SubbaRao, A. S. Szalay, I. Szapudi, G. P. Szokoly, A. R. Thakar, C. Tremonti, D. L. Tucker, A. Uomoto, D. Vanden Berk, M. S. Vogeley, P. Waddell, S.-i. Wang, M. Watanabe, D. H. Weinberg, B. Yanny, N. Yasuda, and SDSS Collaboration. The Sloan Digital Sky Survey: Technical Summary. *AJ*, 120:1579–1587, September 2000.

[329] J. P. Zahn. Circulation and turbulence in rotating stars. *A&A*, 265:115–132, November 1992.

Appendix A

SUPPLEMENTARY DATA

A.1 Coarse labels for 2550 stars

Table A.1 lists all 2550 stars with no known label, and predicted labels generated by the SVC trained on coarse labels.

A.2 Frequencies found via prewhitening

The following tables contain the list of unique frequencies (separated by $1.5/T$) found by the prewhitening procedure described in §5.2.2, not accounting for the red noise that we later identify as stochastic low frequency variability.

Table A.1. Common names, coordinates, and predicted labels of 2550 stars input to the SVC trained on coarse labels.

Common Name	R.A. [deg]	Dec [deg]	Predicted Coarse Label
WISE J000559.28-790653.3	1.49713706	−79.11483482	Hot
TYC 4500-1480-1	2.86210879	79.08686958	Hot
BD+61 45	5.25504760	62.77064970	Hot
NGC 104 LEE 2520	5.41170226	−72.21106679	Hot
WISE J002203.44-693554.7	5.51434821	−69.59851087	Hot
WISE J002207.43-742212.1	5.53102165	−74.37003199	Hot
WISE J002318.05-742326.4	5.82523611	−74.39068759	Hot
WISE J002340.20-750446.9	5.91756693	−75.07972556	Hot
WISE J002758.92-764527.2	6.99552600	−76.75757402	Hot
WISE J002759.32-742119.8	6.99734043	−74.35552728	Hot

Note. — This table is published in its entirety in the machine-readable format. A portion is shown here for guidance regarding its form and content.

Table A.2. Unique frequencies, amplitudes, phases, and formal errors for S Dor found via prewhitening. For each frequency, we specify the signal to noise as defined in text, and whether the corresponding signal is significant under the assumption of white noise.

f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	White Noise Significant?
0.17914669	0.00018699	1.37355329	0.02618908	2.1115	0.0191	1.3475	Y
0.74026962	0.00019105	1.29287459	0.02518611	−2.1163	0.0195	1.2684	Y
0.40287251	0.00019943	1.14940943	0.02337377	3.0257	0.0203	1.1276	Y
0.43090953	0.00019300	1.14783830	0.02258911	2.2956	0.0197	1.1261	Y
1.01866464	0.00020589	1.03719972	0.02177483	−2.7873	0.0210	1.0176	Y
0.31848154	0.00019611	1.05415657	0.02108043	1.5970	0.0200	1.0342	Y
0.29162660	0.00017781	1.12253453	0.02035210	−2.9533	0.0181	1.1013	Y
0.24027468	0.00021217	0.90005164	0.01947179	0.9284	0.0216	0.8830	Y
0.13744498	0.00021613	0.83184153	0.01833204	−1.7119	0.0220	0.8161	Y
1.07115176	0.00022219	0.78722343	0.01783562	1.8799	0.0227	0.7723	Y
0.78858490	0.00022709	0.74934953	0.01735188	−1.3781	0.0232	0.7352	Y
0.90828327	0.00021351	0.77605388	0.01689571	1.7982	0.0218	0.7614	Y
0.48307755	0.00020698	0.77700877	0.01639875	−2.3994	0.0211	0.7623	Y
0.34807760	0.00022384	0.69613397	0.01588865	2.3071	0.0228	0.6830	Y
0.85068961	0.00022519	0.67309428	0.01545560	−1.3659	0.0230	0.6603	Y
0.56272556	0.00022740	0.64883341	0.01504492	−2.4447	0.0232	0.6365	Y
1.22788836	0.00022903	0.62763187	0.01465752	−0.5198	0.0234	0.6157	Y
0.98355185	0.00025223	0.54339970	0.01397567	−2.0356	0.0257	0.5331	Y
0.07590602	0.00025525	0.52566523	0.01368190	−1.4919	0.0260	0.5157	Y
0.60660139	0.00025350	0.51849540	0.01340258	0.4052	0.0258	0.5087	Y
1.33015001	0.00024735	0.47928094	0.01208818	1.8572	0.0252	0.4702	Y
0.52554769	0.00029181	0.39080118	0.01162824	2.6079	0.0298	0.3834	Y
0.63478994	0.00028666	0.38523180	0.01126049	0.3253	0.0292	0.3779	Y
1.12676389	0.00031689	0.33764956	0.01091027	−2.3219	0.0323	0.3313	Y
1.77927887	0.00032095	0.32894596	0.01076545	1.6639	0.0327	0.3227	Y
1.46672799	0.00031496	0.33088598	0.01062659	−2.3816	0.0321	0.3246	Y

Table A.2 (cont'd)

f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	White Noise Significant?
1.42147429	0.00035459	0.27557523	0.00996403	−1.6050	0.0362	0.2704	Y
1.28479399	0.00035017	0.27609702	0.00985824	−2.3815	0.0357	0.2709	Y
0.66258409	0.00034800	0.27480157	0.00975139	−0.1970	0.0355	0.2696	Y
1.67496645	0.00035948	0.26019401	0.00953745	2.8436	0.0367	0.2553	Y
1.59025206	0.00034506	0.26826890	0.00943897	−1.8810	0.0352	0.2632	Y
1.72867338	0.00034824	0.26282574	0.00933267	0.6700	0.0355	0.2578	Y
1.89236098	0.00034930	0.25913313	0.00922956	2.8694	0.0356	0.2542	Y
1.17357337	0.00038477	0.22163510	0.00869573	1.7063	0.0392	0.2174	Y
0.82216092	0.00040298	0.20429334	0.00839472	−1.6606	0.0411	0.2004	Y
1.84679016	0.00041217	0.19482554	0.00818812	0.8427	0.0420	0.1911	Y
1.99892588	0.00039723	0.20056820	0.00812403	0.4310	0.0405	0.1968	Y
1.64601470	0.00039939	0.19780379	0.00805566	−0.7468	0.0407	0.1941	Y
1.92956387	0.00041821	0.18733961	0.00798889	−0.3576	0.0426	0.1838	Y
1.55320350	0.00045138	0.16847434	0.00775430	−0.8069	0.0460	0.1653	Y
2.36401361	0.00048031	0.15626593	0.00765330	1.8377	0.0490	0.1533	Y

Table A.3. Unique frequencies, amplitudes, phases, and formal errors for HD 269953 found via prewhitening. For each frequency, we specify the signal to noise as defined in text, and whether the corresponding signal is significant under the assumption of white noise.

f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	White Noise Significant?
1.59360677	0.00033360	0.16933302	0.00576003	−2.4637	0.0340	0.2248	Y
1.17415786	0.00045342	0.12308820	0.00569092	−0.7787	0.0462	0.1634	Y
1.33479243	0.00067704	0.08190471	0.00565440	−1.2113	0.0690	0.1087	Y
0.22397835	0.00077286	0.07154318	0.00563809	2.3700	0.0788	0.0950	Y
2.67017901	0.00084583	0.06522685	0.00562564	0.3010	0.0862	0.0866	Y
0.58964787	0.00123743	0.04450296	0.00561529	−1.1783	0.1262	0.0591	Y
4.00557547	0.00130137	0.04228006	0.00561047	1.8289	0.1327	0.0561	Y
0.79888166	0.00134452	0.04089120	0.00560608	2.3207	0.1371	0.0543	Y
2.35305773	0.00136719	0.04018379	0.00560199	−2.6328	0.1394	0.0533	Y
0.56055207	0.00141498	0.03879914	0.00559803	0.5170	0.1443	0.0515	Y
0.26019548	0.00145392	0.03773483	0.00559432	1.0507	0.1483	0.0501	Y
0.41029670	0.00166929	0.03284581	0.00559081	2.5195	0.1702	0.0436	Y
0.64779043	0.00171080	0.03203336	0.00558813	−3.1395	0.1744	0.0425	Y
0.31217690	0.00180051	0.03042365	0.00558561	0.1242	0.1836	0.0404	N

Table A.4. Unique frequencies, amplitudes, phases, and formal errors for HD 269582 found via prewhitening. For each frequency, we specify the signal to noise as defined in text, and whether the corresponding signal is significant under the assumption of white noise.

f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	White Noise Significant?
0.20327588	0.00026040	2.82300142	0.03653911	1.9457	0.0129	1.4285	Y
0.11705849	0.00038216	1.56893390	0.02980298	−1.8281	0.0190	0.7939	Y
0.27226816	0.00053917	1.02731918	0.02753193	−2.3371	0.0268	0.5199	Y
0.78368008	0.00069788	0.76229473	0.02644290	0.0658	0.0347	0.3857	Y
0.42011270	0.00076838	0.67630930	0.02583021	1.8420	0.0382	0.3422	Y
0.93338470	0.00082913	0.59356109	0.02446216	2.5364	0.0412	0.3004	Y
1.04145053	0.00111576	0.42869614	0.02377533	−1.8655	0.0555	0.2169	Y
1.60232584	0.00110823	0.42783042	0.02356714	−2.4842	0.0551	0.2165	Y
2.30579937	0.00112019	0.41940449	0.02335219	1.7965	0.0557	0.2122	Y
1.84981762	0.00115772	0.39499834	0.02273024	1.7261	0.0575	0.1999	Y
3.19875548	0.00124431	0.36446275	0.02254161	−0.4055	0.0618	0.1844	Y
1.24393495	0.00124528	0.36155380	0.02237916	−2.3360	0.0619	0.1830	Y
0.64113112	0.00136209	0.32818766	0.02221935	−0.2110	0.0677	0.1661	Y
2.60327615	0.00142004	0.31293170	0.02208795	−2.7382	0.0706	0.1584	Y
2.15944146	0.00143368	0.30823824	0.02196565	−1.2146	0.0713	0.1560	Y
1.48653861	0.00144310	0.30457385	0.02184704	3.0061	0.0717	0.1541	Y
0.87749148	0.00150520	0.28886982	0.02161222	−2.1677	0.0748	0.1462	Y
3.42963659	0.00152746	0.28324793	0.02150505	−1.3241	0.0759	0.1433	Y
1.79273050	0.00169152	0.25455605	0.02140248	−0.6712	0.0841	0.1288	Y
1.95391049	0.00168754	0.25415661	0.02131865	−1.6620	0.0839	0.1286	N
1.33240989	0.00234358	0.18158923	0.02115312	−0.4892	0.1165	0.0919	N
1.13080040	0.00161657	0.26272929	0.02111096	−2.4513	0.0804	0.1329	Y
0.50022915	0.00173592	0.24363786	0.02102219	−0.9353	0.0863	0.1233	Y
3.86015782	0.00167586	0.25044910	0.02086229	−1.1760	0.0833	0.1267	Y
3.34708610	0.00170259	0.24554829	0.02078023	−2.5729	0.0846	0.1243	Y

Table A.5. Unique frequencies, amplitudes, phases, and formal errors for HD 270046 found via prewhitening. For each frequency, we specify the signal to noise as defined in text, and whether the corresponding signal is significant under the assumption of white noise.

f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	White Noise Significant?
0.07243476	0.00030333	0.20507787	0.00634302	−0.3777	0.0309	0.2496	Y
0.17088628	0.00053075	0.11542925	0.00624698	1.5664	0.0541	0.1405	Y
0.11054852	0.00062375	0.09772997	0.00621592	2.9262	0.0636	0.1189	Y
0.03864776	0.00059915	0.10138306	0.00619388	−1.7175	0.0611	0.1234	Y
0.23757177	0.00081819	0.07373290	0.00615146	−0.3369	0.0834	0.0897	Y
0.30727002	0.00088671	0.06790027	0.00613926	−2.1785	0.0904	0.0826	Y
0.26646463	0.00108406	0.05544526	0.00612890	−1.8481	0.1105	0.0675	Y
1.03232899	0.00124699	0.04814869	0.00612224	−2.4317	0.1272	0.0586	Y
0.38422638	0.00153431	0.03905225	0.00610975	2.7074	0.1565	0.0475	Y
1.46213112	0.00165260	0.03621682	0.00610298	−1.3304	0.1685	0.0441	Y

Table A.6. Unique frequencies, amplitudes, phases, and formal errors for HD 270111 found via prewhitening. For each frequency, we specify the signal to noise as defined in text, and whether the corresponding signal is significant under the assumption of white noise.

f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	White Noise Significant?
0.13690821	0.00134854	0.13655040	0.00931212	2.5988	0.0682	0.1563	Y
0.20720738	0.00228811	0.08001971	0.00925903	−2.9088	0.1157	0.0916	Y
0.48123388	0.00267880	0.06819859	0.00923862	−2.4779	0.1355	0.0781	Y
0.28689396	0.00253835	0.07186560	0.00922492	−1.0078	0.1284	0.0823	Y
0.41136823	0.00327074	0.05561540	0.00919881	0.6688	0.1654	0.0637	N
1.10220750	0.00358087	0.05074735	0.00918952	2.2761	0.1811	0.0581	N

Table A.7. Unique frequencies, amplitudes, phases, and formal errors for HD 269331 found via prewhitening. For each frequency, we specify the signal to noise as defined in text, and whether the corresponding signal is significant under the assumption of white noise.

f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	White Noise Significant?
0.09042104	0.00050776	0.61699402	0.01557183	−0.4965	0.0252	0.4745	Y
0.17425772	0.00049282	0.60733804	0.01487734	−2.6545	0.0245	0.4670	Y
0.24876181	0.00072428	0.39309944	0.01415176	−0.3416	0.0360	0.3023	Y
0.48741510	0.00153999	0.18011694	0.01378726	0.6008	0.0765	0.1385	Y
0.31028888	0.00155887	0.17711953	0.01372399	−2.5111	0.0775	0.1362	Y
0.72997431	0.00315091	0.08710461	0.01364212	−0.4636	0.1566	0.0670	Y

Table A.8. Unique frequencies, amplitudes, phases, and formal errors for HD 269110 found via prewhitening. For each frequency, we specify the signal to noise as defined in text, and whether the corresponding signal is significant under the assumption of white noise.

f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	White Noise Significant?
0.55347685	0.00054988	0.15765766	0.00883984	−3.0872	0.0561	0.1345	Y
0.11644097	0.00063650	0.13560965	0.00880137	0.0247	0.0649	0.1156	Y
0.15753531	0.00084280	0.10207525	0.00877217	−2.7811	0.0859	0.0871	Y
0.06899622	0.00092764	0.09256332	0.00875550	0.3983	0.0946	0.0789	Y
0.19233848	0.00121647	0.07031901	0.00872248	−2.6869	0.1240	0.0600	Y
0.03827322	0.00123650	0.06911817	0.00871466	−2.0575	0.1261	0.0589	Y
0.40617680	0.00160856	0.05304164	0.00870001	−0.8097	0.1640	0.0452	N

Table A.9. Unique frequencies, amplitudes, phases, and formal errors for HD 268687 found via prewhitening. For each frequency, we specify the signal to noise as defined in text, and whether the corresponding signal is significant under the assumption of white noise.

f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	White Noise Significant?
0.33981559	0.00016013	0.99396295	0.01623009	2.0925	0.0163	0.7147	Y
0.08920241	0.00018662	0.80772984	0.01537070	1.5642	0.0190	0.5808	Y
0.15510687	0.00020617	0.70324644	0.01478409	-1.6240	0.0210	0.5056	Y
0.27854537	0.00024130	0.58182295	0.01431571	-0.7820	0.0246	0.4183	Y
0.42764474	0.00023359	0.58723924	0.01398757	-0.7239	0.0238	0.4222	Y
0.12802564	0.00026270	0.49788111	0.01333659	1.4333	0.0268	0.3580	Y
0.63934778	0.00026317	0.48760456	0.01308489	-3.0826	0.0268	0.3506	Y
0.04723986	0.00032221	0.37892065	0.01244933	2.3780	0.0329	0.2725	Y
0.22510722	0.00031700	0.38006651	0.01228505	-2.1114	0.0323	0.2733	Y
0.53353809	0.00032444	0.36148395	0.01195896	-2.5173	0.0331	0.2599	Y
0.49649588	0.00039730	0.29148694	0.01180883	-2.3401	0.0405	0.2096	Y
0.68791268	0.00041399	0.27292858	0.01152133	-1.9944	0.0422	0.1962	Y
1.34086866	0.00045333	0.24730899	0.01143192	-2.5120	0.0462	0.1778	Y
0.37663919	0.00048990	0.22604532	0.01129199	-2.2111	0.0500	0.1625	Y
0.91512039	0.00049874	0.21952667	0.01116421	2.4091	0.0509	0.1578	Y
1.84286938	0.00049901	0.21823149	0.01110428	-1.1779	0.0509	0.1569	Y
0.73238778	0.00054246	0.19967911	0.01104494	2.9946	0.0553	0.1436	Y
0.80885741	0.00057409	0.18694716	0.01094361	-1.4790	0.0585	0.1344	Y
0.56637230	0.00058752	0.18116978	0.01085355	-0.1779	0.0599	0.1303	Y
0.94887342	0.00059490	0.17821977	0.01081098	0.5464	0.0607	0.1281	Y
2.76492751	0.00062369	0.16936638	0.01077111	0.0375	0.0636	0.1218	Y
0.86314365	0.00069636	0.15073177	0.01070295	-0.4226	0.0710	0.1084	Y
0.77723411	0.00076226	0.13697958	0.01064685	-1.6051	0.0777	0.0985	N
1.20371478	0.00083481	0.12400270	0.01055564	0.6327	0.0851	0.0892	Y
3.68650176	0.00097373	0.10559870	0.01048488	-0.9334	0.0993	0.0759	Y
0.99151843	0.00101754	0.10078308	0.01045694	2.3637	0.1038	0.0725	Y

Table A.9 (cont'd)

f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	White Noise Significant?
0.46094861	0.00124317	0.08212864	0.01041091	−1.0581	0.1268	0.0591	Y
1.28252486	0.00126920	0.08030607	0.01039306	2.7736	0.1294	0.0577	N
1.41451720	0.00145098	0.07008822	0.01036985	0.1645	0.1480	0.0504	N
1.58281137	0.00152920	0.06630589	0.01033905	−1.7477	0.1559	0.0477	N
1.63275628	0.00163583	0.06188241	0.01032218	−1.3461	0.1668	0.0445	N
1.07507415	0.00166747	0.06067859	0.01031709	2.8570	0.1700	0.0436	N
4.60899785	0.00171384	0.05900802	0.01031210	−3.1032	0.1748	0.0424	N

A.3 Frequencies found via prewhitening

The following tables contain the list of unique frequencies (separated by $1.5/T$) found in the lightcurves of the FYPS by the prewhitening procedure described in §6.3.2, now taking into account the stochastic low frequency variability. Frequencies that are exact harmonics of other frequencies are indicated. Combination frequencies recovered are noted in the comments of each table. However, these are only the "exact" harmonics and combination frequencies recovered to within the precision of the observed frequencies. As discussed in §6.3.2, frequencies with spacings that are close to but not exact harmonics and combinations are recovered in HD 268687, HD 269840, and HD 269902; in these stars, only f_4, f_4 and f_6, and f_2 respectively don't belong to these sequences of near-harmonics.

Table A.10. Unique frequencies, amplitudes, phases, and formal errors for HD 269953 found via prewhitening. For each frequency, we specify the SNR as defined in text, and the height of the associated peak in the RPS at that stage of prewhitening.

Frequency	f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	RPS Peak Height
f_0	1.59347960	0.00002335	0.17835463	0.00269747	2.2998	0.0151	53.8321	1187.4096
f_1^*	2.67052158	0.00009247	0.04452312	0.00266624	0.2701	0.0599	32.9642	150.6893
f_2	1.33523329	0.00006790	0.06058007	0.00266395	-0.5273	0.0440	20.1564	133.2001
f_3	1.17424693	0.00006774	0.06062141	0.00265962	0.2279	0.0439	29.5097	118.7172

Note. — $*$: harmonics of f_2.

Table A.11. Unique frequencies, amplitudes, phases, and formal errors for HD 269110 found via prewhitening. For each frequency, we specify the SNR as defined in text, and the height of the associated peak in the RPS at that stage of prewhitening.

Frequency	f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	RPS Peak Height
f_0	0.55280981	0.00003493	0.16642430	0.00376526	-1.7283	0.0226	19.2635	173.6190
f_1	1.76353979	0.00013526	0.04277167	0.00374665	1.0248	0.0876	28.0123	59.3401
f_2	0.54185885	0.00007463	0.07747331	0.00374458	-1.2620	0.0483	14.7085	40.2324
f_3	0.56377424	0.00009201	0.06277635	0.00374066	-2.2209	0.0596	7.5121	29.0397

Table A.12. Unique frequencies, amplitudes, phases, and formal errors for HD 268687 found via prewhitening. For each frequency, we specify the SNR as defined in text, and the height of the associated peak in the RPS at that stage of prewhitening.

Frequency	f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	RPS Peak Height
f_0	2.76530509	0.00006476	0.18020143	0.00755737	−2.9889	0.0419	63.1039	268.6123
f_1^*	3.68693149	0.00010893	0.10696509	0.00754615	2.7915	0.0705	86.6110	171.2117
f_2	1.84352713	0.00005631	0.20680066	0.00754233	−2.5874	0.0365	44.4964	148.7754
f_3	4.60860763	0.00018083	0.06428032	0.00752813	2.2392	0.1171	57.2225	88.0887
f_4	1.34234726	0.00004934	0.23553062	0.00752674	3.0867	0.0320	50.1173	87.3329
f_5	0.92187126	0.00004733	0.24496247	0.00750823	−2.2787	0.0307	8.2893	33.9434

Note. — $*$: harmonics of f_2. $f_3 = f_1 + f_5$. $f_0 = f_2 + f_5$.

Table A.13. Unique frequencies, amplitudes, phases, and formal errors for HD 269840 found via prewhitening. For each frequency, we specify the SNR as defined in text, and the height of the associated peak in the RPS at that stage of prewhitening.

Frequency	f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	RPS Peak Height
f_0	2.26965893	0.00005279	0.19821747	0.00463827	−1.0031	0.0234	60.4289	490.7758
f_1	3.40445242	0.00007777	0.13350616	0.00460197	2.2648	0.0345	45.3612	460.6608
f_2	1.13468853	0.00003708	0.27915368	0.00458731	1.8206	0.0164	41.2142	332.8405
f_3	2.83703781	0.00020544	0.04966845	0.00452268	2.2047	0.0911	33.1261	61.2621
f_4	0.71676916	0.00007334	0.13905708	0.00451996	0.5576	0.0325	18.0708	57.1411
f_5	3.97206611	0.00024415	0.04160807	0.00450260	−0.9063	0.1082	32.0854	55.6234
f_6	1.38512974	0.00013130	0.07734552	0.00450113	2.4916	0.0582	24.1347	52.1294

Table A.14. Unique frequencies, amplitudes, phases, and formal errors for HD 269902 found via prewhitening. For each frequency, we specify the SNR as defined in text, and the height of the associated peak in the RPS at that stage of prewhitening.

Frequency	f_j [day^{-1}]	$\epsilon(f_j)$ [day^{-1}]	A_j [ppt]	$\epsilon(A_j)$ [ppt]	ϕ_j [radians]	$\epsilon(\phi_j)$ [radians]	SNR	RPS Peak Height
f_0	2.90337860	0.00006011	0.12717671	0.00455148	-1.2740	0.0358	62.4969	333.8202
f_1	1.45175166	0.00003744	0.20368400	0.00454088	-1.4444	0.0223	53.9124	307.3930
f_2	0.08163766	0.00002411	0.31421154	0.00451126	-2.6455	0.0144	5.5540	21.4397

VITA

Trevor Dorn-Wallenstein is a Ph.D. Candidate in the Department of Astronomy at the University of Washington, where he studies massive stars using a variety of techniques, from astroseismology to population synthesis. He earned his B.A. in 2015 at Wesleyan University in Middletown, Connecticut, where he double-majored in Physics and Astronomy. His undergraduate thesis, which examined the X-ray source population of the Whirlpool Galaxy, M51, earned high honors. After matriculating at UW the following Fall, he passed the last written Qualifying Exam administered by the department, and earned his M.S. in Spring of 2016. He passed his General Exam — Physical Mechanisms and Uncertainties for Nucleosynthetic Yields from Massive Stars as a Function of Mass, Metallicity, and Evolutionary Phase — in November of 2018. After defending his thesis in August of 2021, Trevor plans to continue working as a postdoctoral scholar at UW for the next year. He currently lives in Seattle with his finacée, Nicole Sanchez, and their longhaired miniature dachshund, Poderosa Don-Wallanchez.