Improving synthetic aperture radar measurements of surface movement and snow depth in mountain environments

George Brencher

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2025

Reading Committee:

David Shean, Chair

Scott Henderson

Alexander Handwerger

Program Authorized to Offer Degree:

Civil & Environmental Engineering

University of Washington

Abstract

Improving synthetic aperture radar measurements of
surface movement and snow depth in mountain environments

George Brencher

Chair of the Supervisory Committee:
David Shean
Civil & Environmental Engineering

Mountainous regions store critical water resources and produce devastating natural hazards. As climate change disproportionately impacts mountainous regions, accurate and timely observations are needed for adaptive resource and hazard management and to understand changing cryospheric and geomorphic processes. Synthetic aperture radar (SAR) can provide these observations, but SAR-based measurements are subject to noise and errors that reduce their reliability. Relying on the extensive SAR archive, this dissertation develops workflows that integrate emerging data science approaches, including deep learning, with established geophysical methods to improve SAR-based measurements of surface movement and snow depth in mountainous terrain.

In Chapter 1, I used a convolutional neural network (CNN) to remove atmospheric noise from interferometric synthetic aperture radar (InSAR) interferograms. The CNN was trained using thousands of Sentinel-1 interferograms and exploits differences in the spatial and topographic structure of atmospheric noise and deformation signals, without relying on external atmospheric data. This approach outperforms commonly used atmospheric correction methods and reveals

previously obscured centimeter-scale deformation of rock glaciers and landslides in the Rocky Mountains. These improvements enable more reliable interpretation of subtle surface kinematics in high-relief terrain.

In Chapter 2, I developed a fused InSAR and SAR feature tracking approach to quantify surface displacement of moraines damming glacial lakes. Combining InSAR and feature tracking results in improved displacement time series that are more accurate than those produced using either method alone. Application to the Imja Lake moraine dam in Nepal reveals decimeter-scale cumulative subsidence over a seven-year period and widespread buried ice. I validated these results using very-high-resolution satellite stereo digital elevation models. The observed displacement patterns are consistent with year-round ice flow and warm-season ice melt. These results provide new constraints on the processes contributing to moraine dam degradation and have direct implications for glacial lake outburst flood (GLOF) hazard assessment.

In Chapter 3, I extended this approach to the 23 moraine-dammed glacial lakes in Nepal which are the highest priority for monitoring. I used seasonal change in InSAR coherence as a proxy for buried ice presence. I found that most moraine dams contain buried ice that produces surface displacement of centimeters to tens of centimeters per year. Analysis of displacement components indicates that the observed deformation reflects a combination of ice melt and ice flow, with the relative contribution of each process varying between sites. I found evidence for extensive buried ice in several moraine dams previously classified as ice-free, which substantially changes the conclusions of prior hazard assessments. These results can be used to improve GLOF hazard assessments and modelling studies.

In Chapter 4, I developed a deep-learning approach for regional snow depth prediction across the Western United States. I trained a U-Net CNN using a large archive of airborne lidar snow

depth measurements and multi-modal inputs including SAR backscatter, optical imagery, topography, and coarse-resolution physical model outputs. The final CNN substantially outperforms existing approaches for near real-time prediction of Western U.S. snow depth in accuracy, precision, and resolution. It can be applied to create spatially continuous maps of snow depth over the entire Western U.S. and dense snow depth time series over the past decade. This work establishes a new benchmark for regional snow depth prediction performance, with implications for future operational forecasting.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

When I reflect on how I ended up here, writing the acknowledgements section of a (mostly) finished PhD dissertation, I first think of my family. Other than threatening to call the police on me if I continued to pick up snakes, my parents always encouraged and enabled my love of nature. I am grateful to them for that. Precious to me are the moments when we shared nature together: me and Dad out on the boat, me and Mom walking in the woods, and me, Mom, and Dad snowshoeing and sledding.

As grateful as I am to my parents for fostering my love of nature, I am perhaps even more grateful to them for having me do all kinds of things that did not come naturally to me, for challenging me, and for doing their best to make sure I experienced a healthy dose of failure. These experiences prepared me for the challenging moments (months?) when grinding out a PhD felt more like being trapped in a torture prison of my own stupidity than frolicking in a meadow of learning. And of course, I am grateful for how much thought, time, energy, and money you put into my education.

Tait and Reed, thank you for all the time we spent outside playing in the yard and elsewhere. I have not yet unraveled the ways our dynamic has shaped my psyche, but I am sure it explains a lot about me. Thanks for supporting me at my best and worst. Deming, thank you for joining our crazy family, and thanks for keeping it real.

Micaela Roy is the love of my life. If she were only that, and not much else, I would just need to thank her for existing. But I'm lucky, because I have a tremendous amount more to thank her for. My gratefulness to her makes the work "thank" feel inadequate. I will spare readers who've made it here an exhaustive list of all the ways Micaela has supported me during my PhD. It is at the point that it rises from a series of actions to a feeling that I carry with me, of being generally and unconditionally supported. I'll tell you this, it's a pretty good feeling.

Of course, I owe a huge thanks to the broader group of family members and friends who supported me at the very beginning of this path: Patrick, Patricia, Grandpa T and Grandma Linda,

Grandma Lorna and Grandpa John, Grandma Linda and Grandpa Jim, Nana, Uncle Phil, Aunt Johanna and Uncle Tony, Uncle Kurt and Uncle Dan, Roger, and everyone else. Your encouragement and belief in me continue to make me want to do my best work.

I must acknowledge the Choate Science Research Program, Deron Chang, and particularly Rachel Gritzer for the truly excellent mentorship that propelled me toward a career in science. In particular, I still frequently think back to your expertise on giving an effective talk. I thank Michael Singer for taking me on in his lab group when I was a high schooler who did not know anything. I still think of you every time I cross paths with a caterpillar. I also want to thank all the other members of the Singer Lab for being generous with their time and welcoming to me. And of course, I need to recognize the other members of the SRP biology group. You set a high bar for a first research community, and it was a privilege to learn from you and joke around with you. I was proud to know you all in high school, and I am proud to see what you have accomplished since then.

I have other high school classmates to acknowledge. Will and Max, thank you for joining me on what is still the best adventure I have had in my lifetime. You may not know it, but seeing the West with you was a transformative event in my life. If it were not for you, I might be a biologist. Will, thank you for being such a great friend to me, even during tough moments. I am very proud to know you.

I was privileged to have excellent mentors in high school, and somehow I won the mentor lottery again in college. Kirsten Coe played an important role in attracting me to science by showing me that scientists could be cool and athletic, while also being articulate and brilliant. Kirsten took me to my first academic conference, Botany 2019 in Tucson, and made it an incredibly fun trip. Kirsten, thank you for being such a kind and committed mentor. In the biology department, I also want to acknowledge Tom Root. I count his invertebrate biology class among the best and most fun classes I have ever taken.

There are many people I must acknowledge from the Middlebury Geology Department. First, I want to acknowledge Tom Manley. His oceanography class was the first earth sciences class I was ever able to take. I enrolled when I was more than halfway through my biology degree, which I thought was what people who studied nature did. I will admit that I took it because I wanted to go on the boat. But that class inspired me to explore earth science. Once I started taking more

geology classes, I had the opportunity to learn from the rest of Middlebury's excellent geology faculty. I had great classes with Dave West and Peter Ryan. Dave did more to teach me how to write a lab report than anyone else. Pete inspired me to think about the implications of earth science research for people. I also want to thank the other Middlebury geology majors I worked alongside, particularly the 2020 thesis class. Field trips would not have been nearly as fun or interesting without that community.

I likely would not have done a PhD if not for Kristina Walowski. I was, and still am, inspired by her approach to teaching, science, and life. It is special to spend time with someone so smart, talented, and hard-working, who is also thoughtful and kind. Kristina's Bedrock Geology of Vermont class was genuinely mind-expanding for me. It feels like every third time I meet with Kristina she says something that changes my perspective on science and academia. I feel lucky that I have been able to stay in touch with her over the years since we both left Middlebury.

Jeff Munroe, who co-advised my undergraduate thesis, was a great mentor. On some level, I am still following the research directions he pointed me toward. Tromping around the Uintas with him and soaking up as much knowledge as I could is a core science memory for me. I admire Jeff's boldness as a researcher. He cannot be contained by the bounds of a single discipline, country, or language. He is constantly asking new questions, solving new problems, and trying new things. To me, he remains the blueprint for being a productive researcher at a liberal arts college, which I may attempt to follow someday.

I probably owe more to Al Handwerger's guidance and example than to anyone else. Even though it was not a job responsibility for him, he invested a significant amount of time teaching me InSAR as an undergraduate, setting me up for success in graduate school and beyond. He completely broke with what I expected a remote sensing scientist to be. In all the ways I expected someone in his role to be anxious and nerdy, he instead appeared relaxed and cool. As an undergraduate, Al invited me to Los Angeles to visit JPL. We toured the campus, talked to some InSAR OGs, went surfing, hung out at Ballast Point Brewing, and walked around on a landslide. I went to Trader Joe's for the first time. I felt like I had found my people and made it to where I wanted to go. It gave me a vision for a possible way of living, which I have been pursuing in my own way ever since. I am not sure whether Al has noticed that I am following him around the

country and doing many of the things he did. I feel incredibly grateful for all the time Al has invested in mentoring me, and just hanging out with me, over the past seven (!) years.

I want to thank everyone I worked with at GSI Service Group. Even though I decided to return to school, I learned a lot during my time at GSI that served me later on. In particular, my communication and collaboration skills improved substantially. I also saw what applied science looks like in practice, and I frequently think back on this experience during my research.

I want to acknowledge some people in the InSAR community. First, I want to thank everyone involved in the UNAVCO InSAR short course. I especially want to thank Franz Meyer and the broader ASF SAR community, who are behind many of the InSAR resources I use regularly, as well as Zhang Yunjun, whose MintPy software has been extremely useful in my research.

Once at UW, I had the privilege of learning from and working with many extraordinary people. I first want to acknowledge the CEE faculty I took classes with or sought advice from. Faisal Hossain impressed me with his work connecting science with stakeholders. He was a great addition to my Qualifying Exam committee, helping to solidify my understanding of remote sensing. Jessica Lundquist is an excellent question asker and problem solver. I learned more than I ever expected to know about snow from her, but I also learned a great deal about thinking through complicated issues. I admire that she never seems afraid to ask questions or express skepticism, while also being an impressively supportive mentor. I want to thank Bart Nijssen for being an engaged and insightful member of my Final Exam committee while also serving as department Chair. Conversations with Bart during my closed session gave me useful new perspectives on my research.

I would also like to acknowledge some ESS faculty. I owe most of what I think I know about glaciers to Knut Christianson and Gerard Roe, both enthusiastic and engaging lecturers, and both patient with me when I struggled to understand something. I profoundly enjoyed Alison Duvall's hillslope geomorphology class. Not only did it give me tools I needed to understand the signals I was observing in my research, but it was also fun and got me outside. I also want to thank Alison for serving on my general and final exam committees as the graduate school representative. Her perspective was the perfect addition, and she brought a lot of enthusiasm, joy, and kindness. For a number of reasons, I must acknowledge David Schmidt. First, I want to thank him for advising Al, which allowed Al to advise me. His InSAR knowledge has also been a resource to me throughout

my PhD, both while serving on my Qualifying Exam committee and through his InSAR class. In particular, that class was hugely helpful in pulling together the many stray concepts I had picked up through my research into a systematic understanding of SAR and InSAR. He has always been very generous with his time and experience. I want to thank Caily Condit for offering a really useful and fun "Preparing Future Faculty" class, probably the most immediately applicable class I took at UW. I would also like to acknowledge Ian Joughin, another grand-advisor, having advised David Shean. When I was applying to graduate schools, he directed me to David, and his research has inspired me throughout my PhD.

Outside of ESS and CEE, I want to acknowledge Eli Shlizerman in Applied Mathematics, whose excellent Inferring Structure of Complex Systems course was my first exposure to machine learning and gave me the skills to begin applying it in my research. I am sure that a version of this course taught carelessly would have gone completely over my head, but everything was so clearly explained that most concepts were easy to grasp. This course had an outsized impact on my research, and I still refer back to the material.

The eScience Institute is responsible for a number of programs and events that made a big impact on my research. I first want to acknowledge everyone involved in the eScience Hackweeks, and in particular Anthony Arendt. I participated in three Hackweeks during my time as a PhD student at UW. From a community-building standpoint, they were extremely successful. I met many people at Hackweeks who went on to become close collaborators and good friends. At the Hackweeks where I was a project lead, I also had opportunities to develop leadership, teaching, and collaboration skills that I expect to apply going forward. I want to thank everyone who participated in the projects I led and lent their skills and enthusiasm to working on problems I care about. Another eScience program I benefited from was the Data Science Incubator Program. Being able to discuss data science problems with people from other fields, who brought different perspectives, was interesting, rewarding, and useful. I also want to thank everyone at eScience involved in the Software Carpentries workshops. These workshops taught me the basic skills I needed to get started with Python, without which my research would have been impossible. Lastly, I want to thank Valentina Staneva, with whom I collaborated on two workshops teaching GitHub Actions for scientific data workflows. I really enjoyed working with Valentina, and the exploration

of GitHub Actions we did while preparing for those workshops has made it a valuable tool in much of my research since then.

I am tremendously grateful to Scott Henderson for his role in my academic journey through graduate school. He has been both an excellent collaborator and an excellent mentor. I feel that the most productive moments of my time in grad school were often the moments when I was working most closely with Scott. While serving as my mentor in the eScience Incubator Program, Scott built a software package that directly facilitated the type of data processing I needed to do for my project, including creatively working around having to pay for compute. That effort probably saved me six months of work. One type of interaction with Scott that I always enjoy is when I Google an issue I have encountered and find myself on a GitHub page where Scott fixed the issue eight months earlier. I once did the math: Scott has made improvements to open-source software packages at a rate of 3.3 per day over the past five years. I assume many people primarily interact with Scott through his seeming omnipresence online, building, fixing, and improving software. They are missing out. In person, Scott is kind, funny, and humble. He is also principled, thoughtful, and candid. Working with Scott has been one of the brightest parts of the PhD process for me. A huge thank you to Scott.

There are many people in my research community, at UW and beyond, whom I owe thanks to. I want to acknowledge the spicy-snow team (and H. P. Marshall, who supported our work). I learned a lot from working with you all, and had a lot of fun. Maybe the real Sentinel-1 snow depth signal was the friends we made along the way.

At UW, I want to acknowledge the CEE graduate students in the fishbowl and treehouse. I really appreciated how everyone supported one another with both research and life problems. It was a lot of fun hanging out at lunch over the years. When I finally participated in intramural sports, it was a fun and a welcome distraction. Shoutout to the hydroballers.

I also want to acknowledge the graduate students in ESS, who welcomed me into their community. Not only have I been inspired by their work, but I have also made some lifelong friendships. I first need to thank the brilliant Anna Ledeczi, a great friend who kept inviting me to ESS events until I eventually came to one. I want to thank Ruben De La Calle for being a ridiculously good friend who supported me through the ups and downs of my PhD. I want to thank Daniel Otto, a kindred spirit, for being one of the few people I have met who always seems to be

on the same wavelength as me. I feel like I could talk with Daniel about research forever and not get bored. I want to thank the gifted Trent Thomas for all the awesome research conversations over the years. Not only does Trent inspire me, he gets me fired up. I want to thank Peter Lindquist for always sticking to his guns, being funny, and making pizza. I want to thank Jensen DeGrande for always showing up for me. I want to thank Autum Downey for all the laughs. An excellent commiserator.

This brings me to the members of the TACO lab group, without whom this dissertation certainly would not have been possible. Over the years, the postdocs in our group were a great source of knowledge, advice, and perspective for me. We also shared many adventures and fun moments that had nothing to do with science. I want to thank Romain Hugonnet for generously sharing his boundless knowledge of uncertainty (a paradox?), and for being kind and humble. I want to thank Grégoire Guillet for taking the time to understand and discuss my research with me, and for inspiring me to give Bayesian statistics a closer look. I want to thank Karthik Venkataramani for the time he spent sharing his deep knowledge of machine learning, his experiences, and his advice with me. Lastly, I want to thank Ben Purinton for helping me with ASP and for sharing his many interesting life experiences.

When I arrived at UW in 2021, there were three senior graduate students working toward graduation. They happen to be some of the best people I have ever met. First, I want to acknowledge Michelle Hu. I have always been a little intimidated by Michelle because she is so cool and put together. I admire her confidence, organizational skills, and patience. Michelle is the kind of person I would want to be marooned on an island with after a plane crash. For the same reasons, she was also a great person to know in graduate school, which at times has similarities to a disaster scenario. Second, I want to acknowledge Friedrich Knuth. Friedrich sat next to me, and we quickly formed a close connection. If you wanted to argue that in-person work matters, Friedrich and my daily interactions would be good evidence. We were constantly showing each other new things, helping one another work through issues, and sharing ups and downs. I admire Friedrich's perseverance, software skills, and approach to life. I am an oldest child, but Friedrich has felt like something of an older brother to me. Third, I want to acknowledge Shashank Bhushan. Shashank supported me in many ways throughout my PhD, giving excellent feedback, taking time to explain things, and even cooking dinner for the lab group. Shashank does excellent work. When

I need to model a talk, poster, or document after an example, I usually ask Shashank. He is so caring, funny, and hardworking that it is impossible not to like him.

Among the later TACO group graduate students, I first want to acknowledge Seth Vanderwilt. A Seattle native, Seth did more to acquaint me with the outdoors in Washington than anyone else. Seth is one of those rare people with exceptional outdoor experience and exceptional computer skills, and he shared both with me throughout my PhD. He was especially helpful as I began adding machine learning to my research. Second, I want to acknowledge Ayush Gupta. Though we have not known each other long, I have quickly come to admire Ayush's impressive machine learning skills, willingness to try new things, and positive outlook. I have really enjoyed discussing research with Ayush, and I am excited to work with him more in the future. Third, I want to acknowledge Jack Hayes. I love having Jack in the lab group. He is smart, funny, kind, hardworking, and generous. In the relatively short time we have known each other, he has become a close friend who supported me and my research in many different ways.

For the sake of time, I have to resist the urge to write pages of praise for Eric Gagliano, though in some ways that would be easier than distilling my appreciation for him into a few sentences. Eric and I spent four years in the TACO lab together, worked on numerous projects of many different types, and are graduating in the same quarter. I cannot begin to enumerate the many ways Eric supported me, both academically and otherwise, through all kinds of wonderful and terrible events. There were times when I felt I was putting too much baggage on Eric, and all he ever did was ask for more. Having experienced graduate school together, there are things I truly feel only Eric and I can understand. In some ways, I think we have more differences than similarities, though we have always shared a sense of humor. Through one of life's many quirks, we're bonded now in a way that I believe will last a long time, and I'm very grateful for that.

Last among the TACO group, I need to acknowledge my advisor, David Shean. I am not sure people outside academia fully appreciate the amount of time and energy that goes into advising a PhD student. What David put into my development as a scientist goes beyond professional duty, and beyond what most friends would do for one another. It is probably not quite the amount of work required to raise a child. It is probably closer to the amount of work involved in, patiently, day after day, week after week, year after year, trying to teach a gorilla sign language. It is not always easy for the teacher or the gorilla, though at times it can be glorious for both. It is a profound

challenge and a profound act of service, and I think it should be considered a profound achievement. It is something I will never forget or be able to repay. That we find it in ourselves to do this for each other, I honestly believe, is a testament to humanity. In this case, it is a testament to David's humanity.

Finally, I want to acknowledge the funding sources that made this work possible: the Burges Endowed Presidential Fellowship, the National Science Foundation Graduate Research Fellowship, and various NASA grants.

DEDICATION

To my parents

Chapter 1

INTRODUCTION

When moist air is forced over mountains, it must ascend and cool, causing condensation and precipitation. The result is that mountainous regions often receive more rainfall and snowfall than surrounding lowlands. Some of this water is naturally stored in seasonal snow, glaciers, ice-containing permafrost, and in surface and groundwater. Consequently, mountains serve both as a supply of water resources and as a source of water-related hazards. Seasonal snow, glaciers, and rock glaciers collectively provide runoff used by billions of people (Barnett et al., 2005; Mankin et al., 2015, Jones et al., 2019; Gascoin, 2024). Meanwhile, landslides and glacial lake outburst floods (GLOFs) collectively cause thousands of fatalities and billions of dollars in damage each year (Petley et al., 2005; Carrivick and Tweed, 2016).

As climate change disproportionately impacts mountainous regions (Beniston, 2003; Mountain Research Initiative EDW Working Group, 2015; Huss et al., 2017), accurate and timely observations are needed to quantify changing water resources and identify, monitor, and understand changing natural hazards (Dozier et al., 2016; Bormann et al., 2018; National Academies of Science, 2018; Huggel et al., 2010; Scaioni et al. 2014). Analysis of these observations can provide insight into a wide range of processes in mountain environments. In this dissertation, I primarily focus on measurements of the two following quantities, which are immediately useful for resource and hazard management and can be used to answer a range of important science questions: surface movement and snow depth.

Surface movement measurements can be used to detect and monitor active landslides (e.g. Rosi et al., 2018) and quantify subsidence of dams impounding glacial lakes. They can further be used to understand how precipitation, temperature, and insolation control the movement of the diverse assemblage of dynamic features found in mountainous regions (e.g. Handwerger et al., 2013; Brencher et al., 2021). Surface movement measurements can be used to invert for the thickness of glaciers (e.g. Millan et al., 2022) and rock glaciers, providing information about the quantity of water they store.

Snow depth is the primary source of variability in snow-water equivalent (SWE), the defining hydrologic variable of the seasonal snowpack (Sturm et al., 2010). Measurements of snow depth

can be used to understand the processes that control the variability of SWE across spatial and temporal scales, and for operational applications such as runoff forecasting.

Unfortunately, in situ data collection in mountainous regions is a substantial logistical challenge. Remote sensing offers an alternative way to make these observations. Synthetic aperture radar (SAR) is an active remote sensing technique that involves sending a pulse of microwaves toward a target and measuring the signal returned from the target to the sensor. SAR data can provide information about both surface movement and snow depth. However, these data require extensive processing to produce analysis-ready products and are subject to noise and errors that can decrease their utility. Using the large data quantities stored in SAR data archives, we can apply developing data science approaches alongside established geophysical approaches to mitigate these issues, unlocking the vast quantity of information available in SAR data. Specifically, machine learning-based approaches have significant potential to improve SAR processing and analysis.

1.1 Objectives and dissertation outline

This dissertation focuses on improving SAR-based measurements of two key variables for water resource and hazard management in mountainous regions: surface movement and snow depth.

To measure surface movement, we use Interferometric SAR (InSAR) and SAR feature tracking. InSAR relates the phase offset from successive SAR acquisitions to surface displacement (Bürgmann et al., 2000; Rosen et al., 2000). SAR feature tracking involves 2D cross-correlation of backscatter amplitude values in successive radar images to measure offsets in the range (across-track) and azimuth (along-track) directions.

To measure snow depth, we use the change in the ratio of vertically polarized to horizontally polarized backscatter. This change is caused by depolarization of the radar signal as it interacts with the snowpack, which is related to the depth of the snowpack (Lievens et al., 2019, 2022).

Using data science approaches to improve measurements made with these techniques, we are able to answer a range of important science questions about changing water resources and hazards in the mountainous regions.

In Chapter 2, we train a convolutional neural network to remove a primary source of error in InSAR-based measurements of surface movement: atmospheric noise. Our correction substantially outperforms existing correction approaches and improves the accuracy of surface movement measurements of rock glaciers and landslides in the Rocky Mountains, USA.

In Chapter 3, we develop a SAR-based approach to quantify degradation of moraines that dam dangerous glacial lakes. By combining InSAR and feature tracking measurements, we use the strengths of both methods to provide a more complete record and understanding of moraine dam kinematics. We apply our approach to the Imja Lake moraine dam in Nepal.

In Chapter 4, we scale our approach to measure subsidence and map the extent of buried ice for 23 potentially dangerous glacial lakes in Nepal. We find that most of these dams contain internal ice and exhibit measurable subsidence. We discuss implications for the structural properties of moraine dams and integrate our results with existing risk assessment frameworks to provide pertinent information to regional stakeholders.

In Chapter 5, we work to improve SAR-based snow depth retrievals over the Western United States. We use a large archive of lidar snow depth data to train a convolutional neural network to predict snow depth from remote sensing data. Our approach outperforms existing empirical algorithms and sets new benchmarks for accuracy and resolution in large scale remote sensing-based snow depth retrievals over mountainous regions.

Chapter 2

REMOVING ATMOSPHERIC NOISE FROM INSAR INTERFEROGRAMS IN MOUNTAINOUS REGIONS WITH A CONVOLUTIONAL NEURAL NETWORK

This chapter has been published as: **Brencher, G**., Henderson, S. T., & Shean, D. E. (2025). ***Removing atmospheric noise from InSAR interferograms in mountainous regions with a convolutional neural network***. *Computers & Geosciences*, *194*, 105771.

2.1 Abstract

Atmospheric noise in interferometric synthetic aperture radar (InSAR)-derived estimates of surface deformation often obscures real displacement signals, especially in mountainous regions. As climate change disproportionately impacts the mountain cryosphere, a reliable technique for atmospheric correction in high-relief terrain is increasingly important. We developed and implemented a statistical machine learning atmospheric correction approach that relies on the differing spatial and topographic characteristics of slow-moving periglacial features and atmospheric noise. Our correction is applied at the native spatial and temporal resolution of the InSAR data, does not require external atmospheric reanalysis data, and can correct both stratified and turbulent atmospheric noise.

Using Sentinel-1 data from 2017–2022, we trained a convolutional neural network (CNN) on observed atmospheric noise from 330 short-baseline interferograms and observed displacement signals from time series inversion of 1322 interferograms. We applied our trained CNN to correct 251 additional interferograms over an out-of-region application area, which were inverted to create displacement time series. We used the Rocky Mountains in New Mexico, Colorado, and Wyoming as our training, validation, testing, and application areas. When applied to our testing dataset, our correction offered performance improvements of 131%, 208%, and 68% in structural similarity index measure over corrections using atmospheric reanalysis data, phase correlation with topography, and high-pass filtering, respectively. The CNN-corrected time series reveals previously obscured kinematic behavior of rock glaciers and other features in the application

dataset. Our flexible, robust approach can be used to correct arbitrary InSAR data to analyze subtle surface deformation signals for a range of science and engineering applications.

2.2 Introduction

Interferometric synthetic aperture radar (InSAR) is a remote sensing technique that can measure millimeter-scale deformation of the Earth's surface from space (Bürgmann et al., 2000). This sensitivity makes InSAR a powerful tool for deformation monitoring with an array of applications in Earth science and engineering disciplines (Massonnet and Feigl, 1998; Rosen et al., 2000). However, noise caused by interaction of the radar signal with the atmosphere limits the accuracy of InSAR-derived surface displacement maps (Zebker et al., 1997). Atmospheric noise is especially detrimental in mountainous regions, where high-relief topography causes atmospheric turbulence and variable relative humidity over short spatial and temporal scales. This situation is exacerbated by the limited availability of validation data, short snow-free observation periods, and poor options for reliable stable reference points (Delacourt et al., 1998; Loibl et al., 2019; Remy et al., 2015). Along with issues caused by steep slopes (layover and radar shadow (Eineder, 2004; Colesanti and Wasowski, 2006; Grebby et al., 2021)), atmospheric noise makes effective processing and interpretation of InSAR data challenging in mountain environments.

As climate change disproportionately impacts the mountain cryosphere (Huss et al., 2017), a generalized atmospheric correction approach for InSAR data is increasingly necessary to monitor and understand changing glaciers (Dehecq et al., 2019), permafrost (Delaloye et al., 2010), ice-cored moraines (Ravanel et al., 2018), and landslides (Stoffel and Huggel, 2012). Fortunately, the recent availability of large SAR data archives, on-demand InSAR processing services, user-friendly machine learning tools, and cloud-computing resources provide opportunities to develop and test new methods to correct atmospheric noise.

Anisotropy in atmospheric temperature, pressure, and relative humidity introduces radar propagation velocity variations that cause phase delays in two-way radar returns (Zebker et al., 1997). At C-band wavelengths (~5.5 cm), most of this delay occurs in the lower 5 km of the troposphere, where water vapor partial pressure is highest (Ding et al., 2008; Bekaert et al., 2015b). Tropospheric conditions change significantly on sub-hourly time-scales, causing differing two-way phase delays during successive radar acquisitions and introducing an atmospheric component into the observed interferometric phase offset (Emardson et al., 2003; Ding et al., 2008). This offset manifests as apparent line-of-sight displacement signals of up to 15–20 cm that are unrelated to

deformation of the ground surface, with characteristic spatial patterns ranging from short (km) to longer (tens of km) scale (Lohman and Simons, 2005; Li et al., 2007; Bekaert et al., 2015b; Murray et al., 2019). Atmospheric noise is typically described as "stratified" or "turbulent," depending on its physical origin. Stratified atmospheric noise is caused by coherent changes in vertical gradients of pressure, temperature, and water vapor; it typically manifests as elevation-dependent noise (Ding et al., 2008; Bekaert et al., 2015b). Turbulent atmospheric noise is caused by turbulent mixing of water vapor in the lower troposphere, and typically manifests as km-scale noise (Ding et al., 2008; Bekaert et al., 2015b). Both sources of noise are prominent in mountainous regions. Li et al., (2007) found that a spatial resolution of 0.3 km was required to characterize and remove 90% of atmospheric noise in C-band interferograms over the Shanghai region in China, which is mostly flat. Finer spatial resolution (<0.3 km) is required where rugged topography shortens noise length scales.

2.2.1 Atmospheric noise correction: previous work

Atmospheric noise corrections can be separated into two categories – methods that rely on external atmospheric data and phase-based methods that characterize noise empirically. Typical sources of external atmospheric data used for interferogram correction include ground-based meteorological observations (Hanssen and Feijt, 1997; Delacourt et al., 1998), observed path delays at fixed global navigation satellite systems (GNSS) reference stations (Janssen et al., 2004; Z. Li et al., 2006a; Yu et al., 2018; Kinoshita, 2022), global meteorological reanalysis data (Doin et al., 2009; Jolivet et al., 2011, 2014), regional atmospheric models (Foster et al., 2006; Shen et al., 2019), and near-infrared water vapor products from spaceborne spectrometers (e.g., Moderate Resolution Imaging Spectroradiometer (MODIS) (Z. Li et al., 2006), Medium Resolution Imaging Spectrometer (MERIS) (Z. Li et al., 2006b)).

Corrections based on external atmospheric data suffer from some common challenges in mountainous environments. Limited availability of in-situ data over large areas can degrade the accuracy of corrections based on GNSS and atmospheric model data (Gao et al., 2018; Jain et al., 2019; Hamm et al., 2020). Corrections based on spatially and temporally coarse or incomplete external data may require substantial interpolation or resampling in both space and time to match the resolution and acquisition timestamps of the SAR data (Gong et al., 2010). In mountainous environments, where atmospheric conditions change over short spatial (hundreds of meters) and

temporal (minutes) scales, even sophisticated interpolation and downscaling likely fails to capture the full complexity of the true tropospheric delays.

Empirical methods to correct atmospheric noise use relationships between surface elevation and phase to remove elevation-dependent noise (Beauducel et al., 2000; Bekaert et al., 2015a, Liang et al., 2018; Murray et al., 2021; Zebker, 2021) or temporal operations like stacking (Zebker et al., 1997; Ferretti et al., 1999) or smoothing over time. These temporal operations are not a viable correction option for a limited number of interferograms or where surface deformation occurs on short timescales relative to the observation interval. Corrections exploiting relationships between surface elevation and phase are not suitable for correcting turbulent atmospheric noise, which may not be correlated with topography. Despite more than two decades of research, mitigation of atmospheric noise remains a significant challenge for satellite InSAR, with no universal approach for mountainous regions (Murray et al., 2019).

2.2.2 Atmospheric noise correction with machine learning

For most applications, surface displacement signals over features of interest and atmospheric noise have distinct spatial and/or temporal characteristics. This situation is ideal for computer vision machine learning approaches including convolutional neural networks (CNNs), which are frequently employed for image denoising applications (Jain and Seung, 2008; Gondara, 2016; Ilesanmi and Ilesanmi, 2021). In the field of InSAR processing, CNNs have been used to filter wrapped interferograms (Mukherjee et al., 2018; Vitale et al., 2022; Yang et al., 2022), to unwrap interferograms (Wu et al., 2020; Sica et al., 2022; Sorkhabi et al., 2022), and to detect deformation related to ground subsidence (Schwegmann et al., 2017; Anantrasirichai et al., 2021), volcanic activity (Anantrasirichai et al., 2018, 2019), mining (Sica et al., 2022; Wu et al., 2020), landslides (Liu et al., 2022), and earthquakes (Brengman and Barnhart, 2021).

Neural networks have also increasingly been used for InSAR atmospheric corrections. Chen et al., (2020) designed a UNet CNN architecture to remove atmospheric noise from TerraSAR-X data over Hong Kong, providing correction results that outperformed a GNSS-based method. Rouet-Leduc et al., (2021) used a CNN to extract cumulative displacement related to fault activity from short interferogram time series. Zhao et al., (2021) developed a recurrent neural network to remove turbulent atmospheric noise from InSAR time series, providing improved results compared to simple gaussian filtering. Chen et al., (2022) implemented a multi-layer perceptron neural network to remove elevation-dependent atmospheric noise from interferograms. Zhou et al.,

(2023) designed a deep fully connected neural network to remove atmospheric noise from Sentinel-1 interferograms, also finding improved results compared to a GNSS-based method. Zhou et al., (2024) explored the effect of a deep neural network correction on an InSAR time series. These studies relied on simulated atmospheric noise and displacement signals to train their neural networks. While effective, if the simulated data are too simple relative to real interferograms, the resulting models may struggle to generalize, impacting correction performance.

Here, we present a CNN atmospheric correction based solely on the differing spatial and topographic characteristics of atmospheric noise and displacement signals of interest. We demonstrate a robust method for training dataset preparation using real InSAR data rather than simulated data. Our correction is applied at the native spatial and temporal resolution of the original InSAR data. It does not rely on external atmospheric data or synthetic training data, and it can remove both stratified and turbulent atmospheric noise in mountainous terrain. We developed and tested this reproducible correction approach using publicly available SAR data, on-demand InSAR processing services, cloud-based machine learning resources, and open-source software. Section 2 describes the data and study sites. Section 3 details our methods for training dataset preparation, CNN architecture and hyperparameters, alternative correction approaches implemented for comparison, evaluation of results, and CNN correction application. Section 4 discusses the results of training and validation, performance of different correction approaches, testing performance, CNN correction characteristics, and CNN correction application. Section 5 includes a summary and conclusions.

2.3 Data and study site

2.3.1 Data

We used Copernicus Sentinel-1 C-band single-look complex (SLC) radar data collected between June 1 and October 15 during the period from 2017 to 2022. This summer date range was selected to provide the maximum number of mostly snow-free interferograms. Data were acquired in interferometric wide (IW) swath mode (250-km swath width) with vertical co-polarization (VV) along ascending (satellite moving north and looking east) relative orbits 151 (AT151) and 49 (AT49) and descending (satellite moving south and looking west) relative orbit 56 (DT56). No data were available for relative orbit 49 during 2022. The Sentinel-1 SLCs have a spatial resolution

of 22 m in azimuth (i.e., along-track direction) and 2.7–3.5 m in range (i.e., across-track direction) depending on the incidence angle. The minimum revisit time is generally 12 days.

Interferograms were processed with the Alaska Satellite Facility Hybrid Pluggable Processing Pipeline (ASF HyP3) (Hogenson et al., 2023). We performed multi-looking (10 looks in range and 2 looks in azimuth) resulting in interferograms with 40 m pixel spacing. HyP3 uses the 2021 release of the 30-m Copernicus "GLO-30" digital elevation model (DEM) to remove the topographic component of the phase and geocode the interferograms (European Space Agency, 2021). The GLO-30 DEM product has an absolute vertical accuracy of <4 m (LE90) and an absolute horizontal accuracy of <6 m (CE90) (NSDDA, 1998; European Space Agency, 2022). We used the same DEM products during CNN training. The DEM tiles were corrected to provide height above the WGS84 ellipsoid, resampled to match the pixel spacing of the interferograms, and reprojected to the local Universal Transverse Mercator (UTM) zone (Hogenson et al., 2023).

2.3.2 Study area and application area

To train, validate, and test our model, we defined a ~187,000 km^2 study area in the Rocky Mountains of New Mexico, Colorado, and Wyoming, USA (Figure 2.1). This region has a large area above treeline (Figure 2.1 inset) and high density of slow-moving periglacial features (Janke, 2005; Johnson et al., 2021; Routt, 2022). The study area spans several distinct mountain ranges, offering a broad sample of terrain characteristics over an elevation range from ~1300–4400 m. The study area contains 3894 mapped rock glaciers with variable orientation covering a total of ~360 km^2 (Johnson et al., 2021). The mean (± 1 standard deviation) individual rock glacier area is 0.093 ± 0.10 km^2 with mean major and minor axis length of 500 ± 310 m and 240 ± 130 m, respectively (Figure A.1). While there are no published deformation rates for these features, the mean LOS velocity of rock glaciers in the nearby Uinta Mountains ranged from ~0–3 cm/yr during the period from 2016 to 2019 (Brencher et al., 2021). The study area likely contains other slow-moving, coherent features undergoing deformation due to permafrost creep, solifluction, and landslide motion (Mirus et al., 2020).

In addition to the main study area, we also defined a ~45000 km^2 out-of-region "application area" in Wyoming (Figure 2.1B) to demonstrate that our CNN correction can be applied to an independent interferogram time series. This application area spans an elevation range of ~1200–4200 m and contains 411 mapped rock glaciers covering a total of ~55 km^2 (Johnson et al., 2021).

The mean individual rock glacier area is 0.14 ± 0.21 km^2 with mean major and minor axis length of 600 ± 360 m and 300 ± 140 m, respectively (Figure A.1).

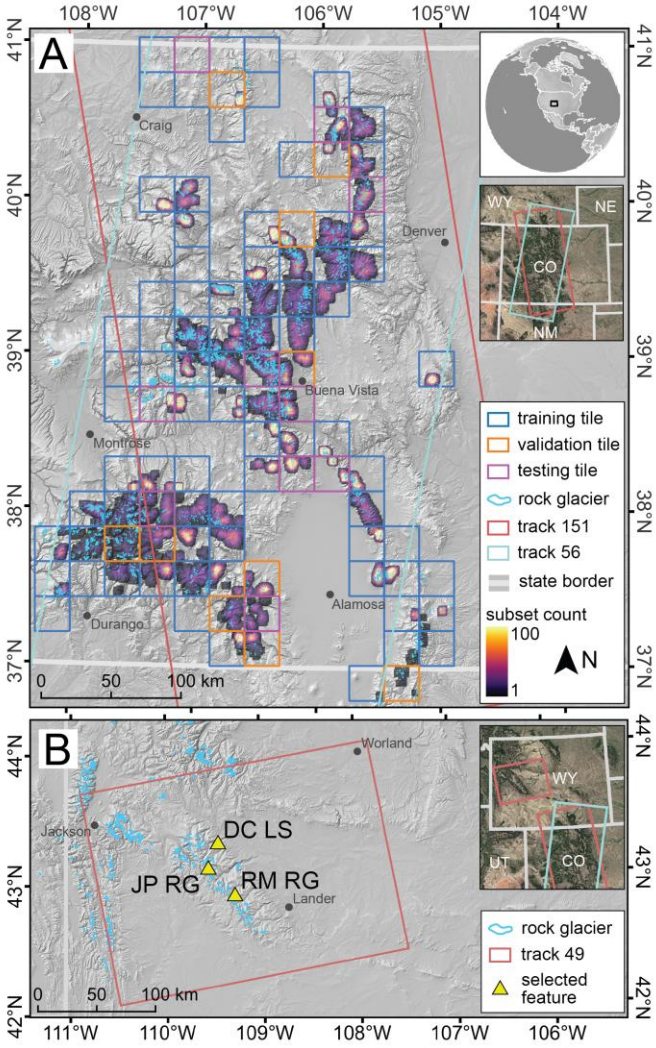

Figure 2.1: Maps showing study area and application area. A) Heatmap of the study area showing the spatial distribution and density of 128 by 128-pixel subsets used for training, validation, and testing. Rock glacier inventory is from Johnson et al., (2021). B) Detail of the application area. Yellow triangles show locations of selected features of interest: the Dinwoody Creek landslide (DC LS), Jackson Peak rock glacier (JP RG), and Roberts Mountain rock glacier (RM RG).

We selected three representative features in the application area for detailed analysis (Figure 2.1B,

Table 2.1). The Jackson Peak rock glacier is on the northeast side of a moraine from the Bull Lake Glacier. The Dinywood Creek landslide is part of a large translational landslide complex

moving northeast along a bedding plane with a clear northwest-trending scarp. The Roberts Mountain rock glacier flows northeast from a cirque north of Roberts Mountain and terminates in a small lake.

Table 2.1: Characteristics and results for selected representative features in the application area. The mean ± standard deviation displacement values were computed from spatially averaged displacement values for each feature over five summer periods between 2017 and 2022. The CNN correction significantly reduces apparent interannual variability due to atmospheric noise. See Figure 2.7 for additional details.

feature name	location	area (km^2)	median elevation (m)	uncorrected mean summer displacement (cm)	CNN-corrected mean summer displacement (cm)
Jackson Peak rock glacier	43.12582°, -109.59104°	0.038	3687	0.19 ± 3.06	0.77 ± 0.23
Dinywood Creek landslide	43.31791°, -109.47372°	0.307	2718	0.65 ± 1.81	0.72 ± 0.39
Roberts Mountain rock glacier	42.92269°, -109.29148°	0.170	3405	0.37 ± 3.05	0.71 ± 0.62

2.4 Methods

2.4.1 Training dataset preparation

We prepared training, validation, and testing datasets composed of 1) characteristic atmospheric noise maps and 2) characteristic displacement signal maps (Figure 2.2). These maps were combined to form "composite interferograms." The characteristic atmospheric noise maps and the corresponding characteristic displacement signal maps were processed further to create the "target atmospheric noise" and the "target displacement signal" datasets used during training. The following sections describe these steps in detail.

Extending the logic outlined in (Hanssen et al., 1999), we assume that the magnitude of atmospheric noise in short-baseline interferograms greatly exceeds the line-of-sight magnitude of coherent periglacial displacement signals, and that nearly all apparent displacement is atmospheric noise. This is supported by the ~2.3 cm mean apparent LOS displacement over the study area in short-baseline interferograms (12 days), nearly 70 times greater than the expected ~0.03 cm mean LOS displacement for rock glaciers over a 12-day period (Brencher et al., 2021).

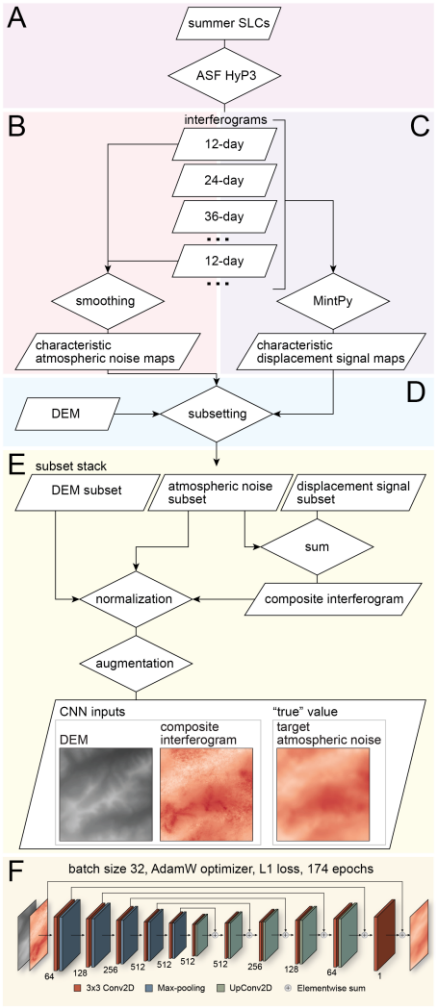

Figure 2.2: Data preparation and training flowchart. A) Interferograms were prepared from summer Sentinel-1 single-look complex (SLC) synthetic aperture radar data. B) Characteristic atmospheric noise maps were prepared by smoothing short temporal-baseline interferograms. C) Characteristic displacement signal maps were prepared by inverting densely connected interferogram networks and calculating low-noise 6-year mean displacement rate maps. D) subsets were prepared from 128 by 128-pixel stacks of characteristic atmospheric noise maps, characteristic displacement signal maps, and the digital elevation model (DEM). E) Composite interferograms were formed by adding scaled subsets of characteristic displacement signal maps and characteristic atmospheric noise maps. Composite interferograms, DEMs, and target atmospheric noise subsets were then normalized and augmented before training. F) A U-Net convolutional neural network (CNN) was trained to predict atmospheric noise, given input interferograms and accompanying DEMs. Architecture illustration modified from (Stucker and Schindler, 2022).

2.4.1.1 Characteristic atmospheric noise maps

To prepare characteristic atmospheric noise maps, we processed all possible short temporal baseline (12-day) Sentinel-1 interferograms over the study area, resulting in 330 interferograms from 3 ascending and 3 descending frames (Table 2.2; Figure A.2). To eliminate any lingering apparent short-spatial-scale surface displacement signals, which are mostly found in low coherence areas (cliffs, vegetation, surface water, and snow), we applied a 2-D gaussian filter to smooth each short-baseline interferogram. As a data augmentation measure, the parameters for the gaussian filter (kernel size between 13x13 and 21x21 pixels and corresponding standard deviation between 3 and 5) were randomly selected throughout training. This range of parameters largely eliminates short-spatial-scale surface displacement signals while preserving realistic atmospheric noise. The resulting smoothed short-baseline interferograms serve as our characteristic atmospheric noise maps.

Table 2.2: Summary of interferogram datasets.

	product	orbit	interferogram count	mean connections
training, validation, testing	target signal	AT151	591	6.68
	target signal	DT56	731	7.17
	target noise	AT151	147	1.72
	target noise	DT56	183	1.79
application	application interferograms	AT49	251	9.47

2.4.1.2 Characteristic displacement signal maps

Characteristic displacement signal maps were derived using a time series processing approach to generate low-noise mean displacement rate maps (e.g. Brencher et al., 2021; Munroe and Handwerger, 2024). We assume that the magnitude of summer LOS displacement signals for active periglacial features far exceeds the magnitude of any remaining atmospheric noise in these mean displacement rate maps. These maps can then be used to approximate representative displacement rate signals for the full summer study period each year.

We used the Miami INsar Time-series software in PYthon (MintPy) package (Yunjun et al., 2019) to perform time series inversion and produce mean LOS displacement rate maps for each summer period. We created densely connected interferogram networks for each relative orbit (Table 2.2; Figure A.2) by processing Sentinel-1 interferogram pairs (temporal baselines of 6–60 days) for each summer period (591 total ascending interferograms, 731 total descending interferograms, mean of 6.94 interferograms created per SLC). Before inversion, atmospheric noise was removed from each interferogram using the empirical method described by (Zebker, 2021). Following inversion, linear deramping and DEM error corrections were performed with MintPy. To produce the final characteristic ascending and descending displacement signal maps, we computed the per-pixel mean of all six summer LOS displacement signal maps for each flight direction. Our characteristic displacement signal maps represent the mean surface displacement rate over the six-summer study period.

2.4.1.3 Subset stack preparation

To define the training, validation, and testing areas, we created a grid of 25 by 25 km tiles (n=100) that overlapped the rock glacier inventory (Johnson et al., 2021) in the study area. Following standard approaches, 10% (n=10) of these tiles were randomly selected for validation and 10% (n=10) were randomly selected for testing, leaving 80% (n=80) for training (Figure 2.1A).

Within each tile, we randomly selected 128 by 128-pixel windows ("subsets") with median elevation above 3300 m (approximately the lower bound for most rock glaciers in the study area) to prepare "subset stacks" containing a characteristic atmospheric noise map, a characteristic displacement signal map, and the resampled Copernicus DEM. This random subsetting process was repeated for all characteristic atmospheric noise maps and all 25-km tiles to generate 14,896 training subset stacks, 2,217 validation subset stacks, and 2,765 testing subset stacks over the study area.

2.4.1.4 Composite interferogram preparation

Prior to forming composite interferograms for each subset stack, we performed another iteration of the (Zebker, 2021) empirical correction to remove any remaining local elevation-dependent atmospheric noise in the characteristic displacement signal map. To sample the full range of realistic unwrapped phase displacement values (in radians) expected for real seasonal to interannual displacement variability, we multiplied the characteristic displacement signal maps

(output from MintPy in m/yr) in each subset stack by a random scalar drawn from a lognormal distribution with a mean of 2 and a standard deviation of 0.4. This scaled displacement map serves as the "target displacement signal" for each composite interferogram. The corresponding "target atmospheric noise" for each composite interferogram is the smoothed characteristic atmospheric noise map in the subset stack.

We added the scaled target displacement signals to the corresponding target atmospheric noise to create composite interferograms with a range of effective signal-to-noise ratios (SNR of 0.01–4.32, median of 0.17). For an observed atmospheric noise magnitude of ~2.3 cm (see Section 2.3.1), a SNR of 0.17 corresponds to a target displacement signal magnitude of ~0.4 cm over the composite interferogram temporal baseline. Our final composite interferograms contain both atmospheric noise from the characteristic atmospheric noise maps and displacement signals from the characteristic displacement signal maps.

2.4.1.5 Data normalization and augmentation

The composite interferograms and target displacement signals were normalized between -1 and 1 using the global minimum and maximum values for all composite interferograms and target displacement signals across the full training, validation, and testing datasets. Each corresponding DEM subset was normalized between -1 and 1 using the global minimum and maximum elevation values across the entire training, validation, and testing datasets. Normalization was performed in order to place inputs on the same scale, stabilizing gradient descent and improving model convergence (LeCun et al., 2002).

The stacks of normalized training data (composite interferogram, target displacement signal, and DEM) were randomly mirrored horizontally and vertically and rotated by a random angle throughout training. In addition, each composite interferogram and target displacement signal was randomly inverted (multiplied by -1) during training to expose the CNN to an equal number of signals with positive and negative range change.

2.4.2 Convolutional neural network

We tested several CNN architectures, including the architecture developed by Rouet-Leduc et al. (2021). We observed the best performance with ResDepth, a CNN with a U-Net architecture developed by Stucker and Schindler (2022). In addition to the skip connections typical of U-Nets, which add the outputs of encoder layers to the outputs of decoder layers at the same level,

ResDepth includes an additional "outer skip connection" where the input image is added to the model output. Our implementation using PyTorch (Paszke et al., 2019) had five downsampling and upsampling layers and two input channels: an interferogram and a corresponding DEM. It had 64 initial filters, increasing to a maximum of 512 filters, all of size 3 by 3 pixels. Rectified linear unit (ReLu) activation was performed after each convolutional layer except the final one. Our training sought to minimize the L1 loss between the model output and the target atmospheric noise. We used the AdamW optimizer (Kingma and Ba, 2017; Loshchilov and Hutter, 2017) with an initial learning rate of 0.0003, a weight decay of 0.1 and a batch size of 32. The input composite interferogram and DEM chip size was 128 by 128 pixels such that overall input dimensions during training was 32 by 2 by 128 by 128. The learning rate was decreased by a factor of 10 when training loss failed to improve after 5 epochs. Early stopping was triggered when validation loss failed to improve after 10 epochs. We trained for 174 epochs on an Nvidia T4 GPU.

2.4.3 Alternative correction approaches

We implemented three widely-used atmospheric noise corrections for comparison with the CNN results: 1) a correction based on European Centre for Medium-Range Weather Forecasts Reanalysis 5th Generation (ERA5) data (Jolivet et al., 2011), 2) an empirical correction based on the relationship between phase and elevation that uses multiple reference points (hereafter, "MuRP") (Zebker, 2021), and 3) a high-pass filter (HPF) correction. We used the pyaps3 package to predict atmospheric noise for the ERA5 correction. These ERA5 and MuRP atmospheric noise prediction maps were subset alongside our characteristic atmospheric noise maps such that each 128 by 128 pixel composite interferogram had an accompanying ERA5 and MuRP noise prediction. The ERA5 and MuRP corrections were applied by subtracting the noise prediction from the composite interferogram. The pixel location with the lowest magnitude in the target displacement signal map was selected as the local reference point for the ERA5 correction. For the HPF correction, we applied a 2-D gaussian filter with a kernel size of 17x17 pixels and a standard deviation of 4 pixels to the composite interferogram, and then subtracted the smoothed composite interferogram from the original composite interferogram.

2.4.4 Evaluation

The output from the trained CNN is a predicted atmospheric noise map, which is subtracted from the input composite interferogram yielding a corrected displacement signal map. We used the

structural similarity index measure (SSIM; Wang et al., 2004) implemented in the scikit-image package (Walt et al., 2014) to compare the corrected displacement signal map to the corresponding target displacement signal map in the testing dataset. The SSIM measures the perceptual similarity of two 2D arrays on a scale between -1 and 1, where higher values correspond to a higher degree of similarity. By comparing luminance, intensity, and contrast in a moving window, SSIM accounts for structural differences between images and is therefore a better metric of perceptual similarity than traditional metrics such as root mean squared error (RMSE; Wang et al., 2004). We computed SSIM values for each of the 2765 subset stacks in the testing dataset to obtain a "SSIM distribution" for each correction approach, and then compared the SSIM distribution median of the three approaches. A Wilcoxon signed-rank test, which is nonparametric and paired, was used to establish the significance of differences between the SSIM distributions for the four correction methods. A Kolmogorov–Smirnov test was used to determine whether the SSIMs were normally distributed. We also computed the coefficient of determination (R^2) and RMSE for the corrected vs. target displacement signal values of all pixels in the testing dataset composite interferograms.

2.4.5 Application

2.4.5.1 Application dataset preparation

Our out-of-region application dataset was composed of densely connected interferogram networks created by processing all possible AT49 Sentinel-1 interferograms that intersected the application area during each summer period (251 total interferograms; Table 2.2, Figure A.2). We refer to these interferograms as "interferogram observations" to distinguish them from the composite interferograms prepared for training, validation, and testing. As with the training, validation, and testing dataset preparation, the interferogram observations and associated Copernicus DEM products were normalized between -1 and 1 for the corresponding global minimum and maximum values of the application dataset.

2.4.5.2 CNN Application

Each interferogram observation and DEM in the application dataset was supplied to the trained CNN to generate a prediction of atmospheric noise. The output atmospheric noise prediction was subtracted from the input interferogram observation, which was then denormalized. Time series inversion of the CNN-corrected interferogram network was performed with MintPy, as described above, and the same inversion process was repeated for the uncorrected interferogram network to

prepare a control dataset for comparison. To produce a single mean LOS velocity map, we computed the per-pixel mean of all five summer LOS velocity maps.

To evaluate the CNN correction, we examined the cumulative LOS displacement time series for each of the three representative features in the application area. To quantify residual atmospheric noise for each feature, we identified a relatively flat and exposed "stable area" with high coherence. The distance between each feature and the corresponding stable area varied due to availability of appropriate terrain. To generate feature displacement time series and uncertainty estimates, we computed the median displacement of all 40 m pixels within the feature polygon and the corresponding stable area at each time step.

2.5 Results and discussion

2.5.1 Training and validation

Following successful training (Figure A.3), the CNN generated atmospheric noise predictions for the training and validation datasets that matched the "target" atmospheric noise used to prepare the input composite interferograms (Figure 2.3). When the predicted atmospheric noise was subtracted from the composite interferogram for each subset in the training and validation datasets, the remaining signals included: 1) high-frequency noise over low-coherence areas, and 2) displacement signals with distinct edges over known periglacial features and landslides (Figure 2.3).

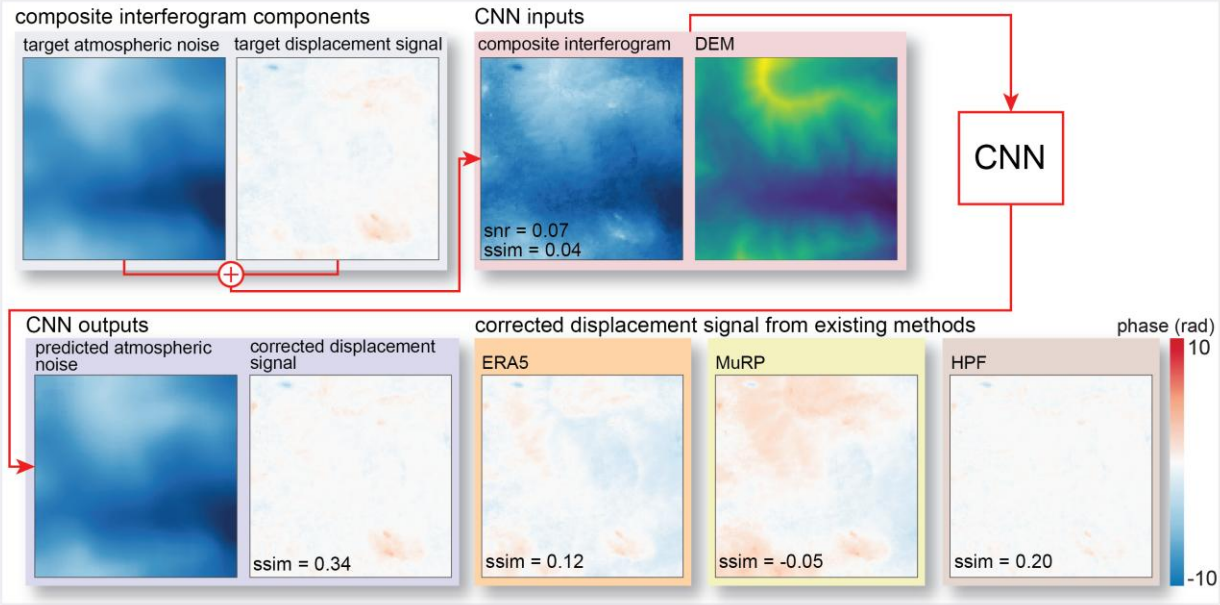

Figure 2.3: Sample results for one validation subset stack. First row shows 128 by 128 pixel subset inputs for the CNN. Second row shows the CNN-predicted atmospheric noise map, the resulting corrected displacement signal map, and the corresponding corrected displacement signal maps from the ERA5, MuRP, and HPF correction methods for comparison. The corrected displacement signal from the CNN most closely resembles the target displacement signal.

2.5.2 Correction performance comparison

The ERA5, MuRP, and HPF corrections did not perform as well as the CNN. The ERA5 correction performance improved when atmospheric noise was correlated with topography. However, the ERA5 correction approach sometimes produced atmospheric noise predictions that did not resemble the target atmospheric noise (Figure A.4). This result is expected, as the 31-km ERA5 grid cells are too large to resolve small-scale turbulent noise. Qualitatively, the ERA5 correction success was also sensitive to the choice of reference point location. The ERA5 correction results could likely be improved by manually picking a custom reference point for each composite interferogram, though this is impractical for large datasets with thousands of interferograms.

The MuRP correction performance appeared to improve when atmospheric noise was correlated with topography or uniform across the composite interferogram. The MuRP correction could likely be improved by applying it to smaller sections of the full interferogram to capture local relationships between phase and topography that vary spatially.

The HPF correction was more effective for composite interferograms with lower SNR values, with a tendency to overpredict atmospheric noise for higher SNR values. A HPF correction could likely be improved by manually tuning the relevant parameters (e.g., kernel size) for each composite interferogram, but again, this is impractical for large datasets.

2.5.3 Testing

For all correction approaches, the testing dataset SSIM distribution was not normal (p=0.00, n=2765) and we compare the median of each SSIM distribution (Figure 2.4B). The CNN-corrected displacement signal had the highest SSIM distribution median, significantly larger (p=0.00, n=2765) than the SSIM distribution median for the ERA5, MuRP and HPF corrections (Figure 2.4B). The CNN correction offered a 131% improvement in SSIM distribution median over the ERA5 correction, a 208% improvement over the MuRP correction, and a 68% improvement over the HPF correction for the testing dataset (Figure 2.4B). The CNN-corrected composite interferograms also had lower RMSE and R^2 than composite interferograms corrected with the alternative methods (Figure 2.4A).

2.5.4 CNN correction characteristics

Unlike the alternative correction approaches, our CNN correction approach is adaptive, and when applied to a large dataset, does not require parameter tuning to make realistic corrections for composite interferograms with differing atmospheric noise and displacement signal characteristics. The ERA5 correction is based on external meteorological reanalysis data and sometimes makes inaccurate atmospheric noise predictions compared to the observed atmospheric noise in composite interferograms (Figure A.4), explaining its broad SSIM distribution (Figure 2.4B). The CNN, MuRP and HPF corrections are derived from the observed data, and consequently tend to consistently reduce atmospheric noise by deriving corrections from the data itself. Unlike the HPF correction, the CNN removes entire displacement signals in the output atmospheric noise predictions rather than only distinct edges, resulting in improved performance over high-pass filtering.

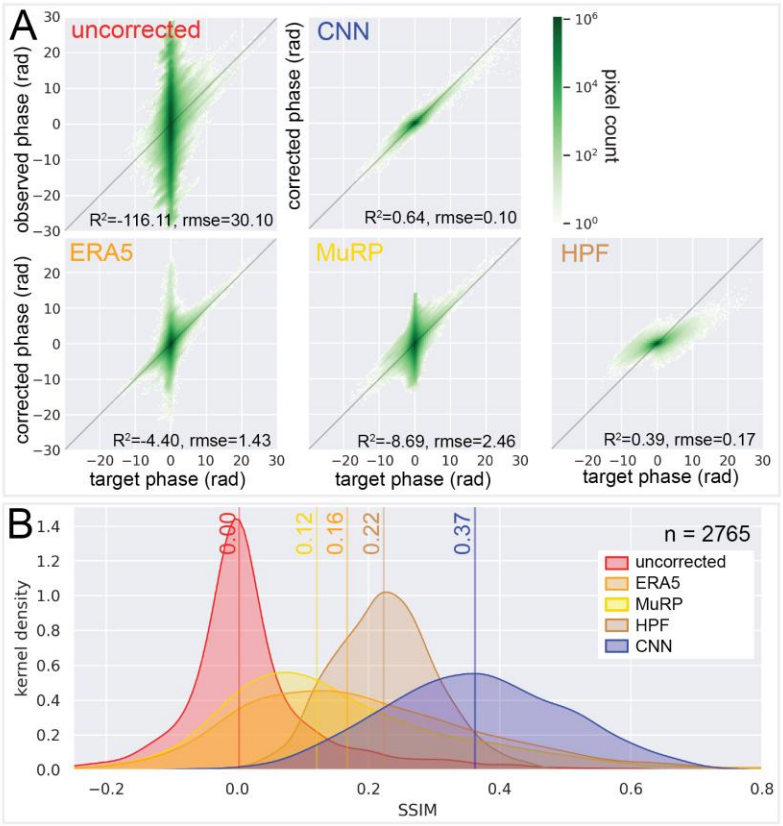

Figure 2.4: A) Scatterplots showing corrected (y axis) vs. target (x axis) displacement signal values for all pixels in the testing dataset. The CNN-corrected displacement signal has the largest R^2 and smallest RMSE values (see annotations in each scatterplot). B) Structural similarity index measure (SSIM) distributions and median values (vertical lines) for each correction approach. Higher SSIM values indicate greater similarity between the corrected and target displacement signal for each subset stack. The distributions show SSIM values for all subset stacks (n=2765) in the testing dataset, with highest SSIM distribution median for the CNN correction.

2.5.5 Application area

The CNN atmospheric noise predictions were similar to the real atmospheric noise in the interferogram observations over the out-of-region application area, including both elevation-dependent and turbulent noise (Figure 2.5). Apparent displacement signals spanning large areas (>10 km^2) were generally interpreted as atmospheric noise by the CNN, even where they had well-defined edges (e.g., unwrapping artifacts, Figure A.5).

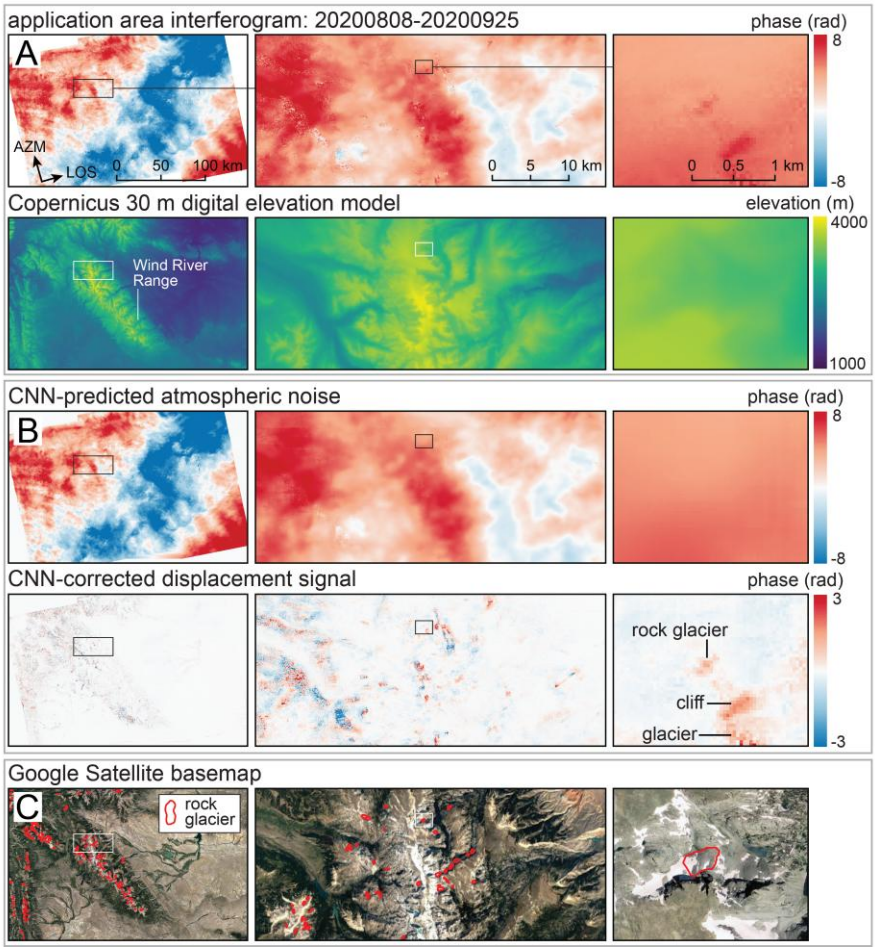

Figure 2.5: Sample results for an August-September 2020 interferogram observation in the application area. Positive (red) signals indicate apparent line-of-sight (LOS) displacement away from the satellite. A) Example CNN inputs (interferogram observation and DEM). The center and right columns show detail around a rock glacier with downslope displacement signal roughly aligned with the LOS. B) Example CNN-predicted atmospheric noise and corrected displacement signal from the out-of-region application dataset. Note that turbulent and stratified atmospheric noise is removed in the corrected displacement signal, while small-scale displacement signals associated with moving features (rock glacier, glacier) are preserved. C) Google Satellite basemap for context.

Following time-series processing, the uncorrected displacement time series shows apparent landscape-scale positive and negative LOS displacements (Figure 2.6A, C, E) due to atmospheric noise, while the CNN-corrected time series shows small-scale displacement signals over mountains, where we expect deformation of dynamic rock glacier features (Figure 2.6B, D, F).

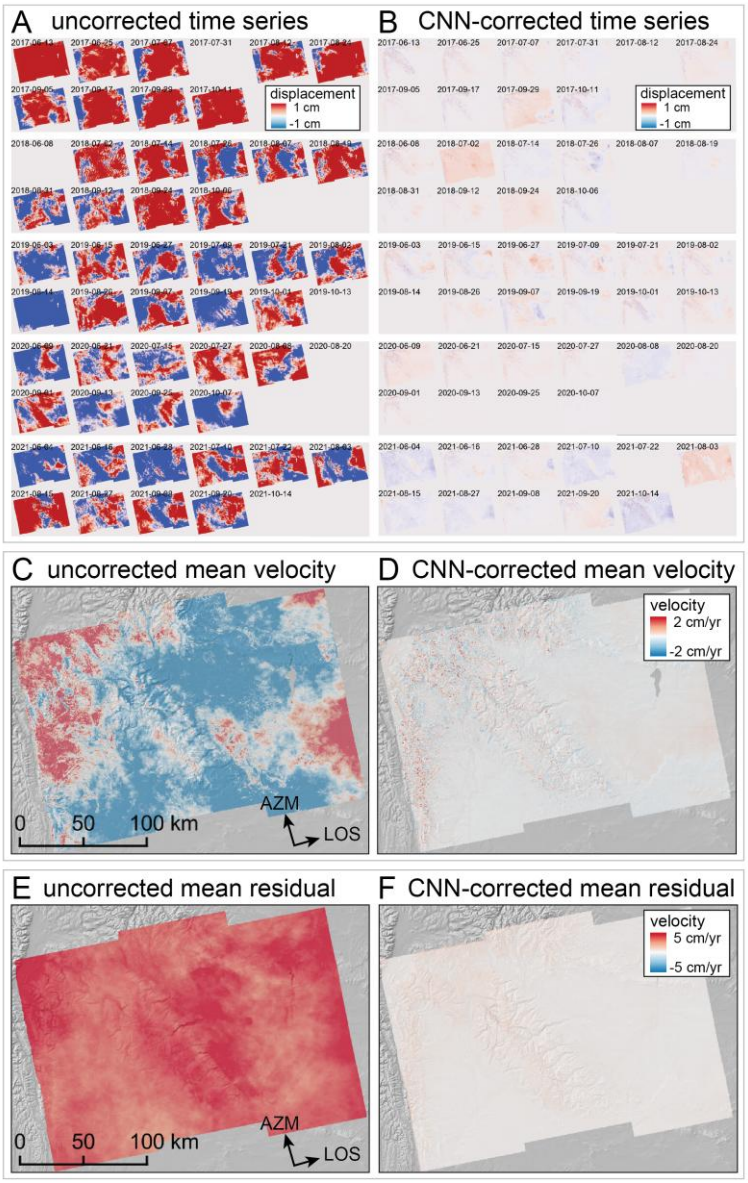

Figure 2.6: Cumulative displacement time series for the application area from inversion of the summer interferogram observation networks for the period spanning 2017 to 2021 (no data were collected for relative orbit 49 in 2022). Positive (red) signals indicate apparent displacement away from the satellite. A) Uncorrected and (B) CNN-corrected time series. Five-year mean velocity maps from the (C) uncorrected and (D) corrected time series. Mean of residuals for all time steps in the (E) uncorrected and (F) corrected time series inversion. Note persistent elevation-dependent atmospheric noise in the uncorrected cumulative displacement time series, mean velocity map, and mean velocity residual map. The atmospheric noise is largely removed from the CNN-corrected products.

2.5.6 Rock glacier and landslide kinematics

We observe distinct positive displacement signals for all three selected features in the CNN-corrected mean velocity map (Figure 2.7B), which is consistent with expected displacement in the downslope direction. The uncorrected summer displacement time series show large interannual variability and physically unrealistic changes in displacement direction and magnitude over both the feature of interest (Figure 2.7C) and stable area (Figure 2.7D) for each site. We attribute these apparent displacements to local atmospheric noise.

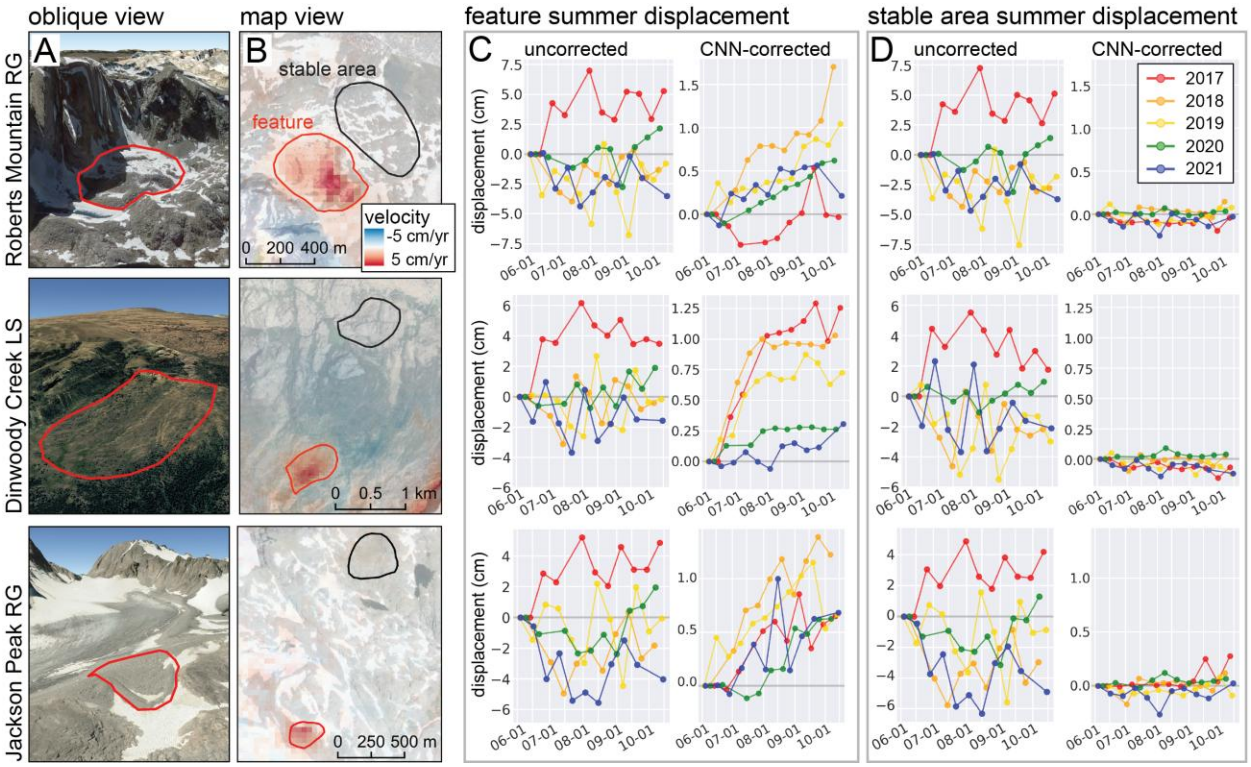

Figure 2.7: Observed summer LOS displacement from June 1 to October 15 for selected features and stable areas in the uncorrected and CNN-corrected time series during the 5 years between 2017 and 2021. Columns show (A) Oblique view of Google Satellite basemap for each feature. B) Map of the 5-year mean of CNN-corrected LOS velocity (Figure 2.6D). C) Uncorrected and corrected summer displacement time series of representative features. D) Uncorrected and corrected summer displacement time series of corresponding stable areas. The uncorrected time series show greater variability, and the apparent stable area displacement history is similar to the corresponding feature displacement history. In the CNN-corrected time series, all features moved away from the satellite over time (as expected for downslope motion), with limited apparent stable area displacement.

By comparison, the CNN-corrected time series for these features generally show positive cumulative LOS displacement of ~0.3 to 1.3 cm each summer (Figure 2.7C, D; Table 1). The displacement direction was generally downslope, with some apparent changes in direction potentially due to spurious signals over low-coherence pixels within the feature polygon. Following CNN correction, the cumulative LOS displacement over stable surfaces was near-zero. Standard deviation of seasonal stable area displacement was on average 1.58 cm for the uncorrected time series and 0.054 cm for the CNN-corrected time series, indicating that the CNN correction effectively removed most atmospheric noise.

These results demonstrate the ability of the CNN correction to capture the kinematics of coherent, slow-moving surface features with high temporal resolution, even where there is significant atmospheric noise and limited stable reference areas in the immediate vicinity of the feature of interest. For example, the Dinywood Creek landslide (Figure 2.7) displayed notable intra-annual variability in the CNN-corrected time series, with faster velocities in early summer and slower velocities in late summer during most years, potentially related to timing of precipitation and meltwater infiltration.

2.5.7 Generalization

We found that training loss diverged from validation loss after only a few epochs (Figure A.3), suggesting that generalization represents a challenge for the current CNN correction approach. This finding is consistent with minor residual atmospheric noise visible at some time steps in the application area time series, especially over non-mountainous terrain (Figure 2.6B). Generalization could likely be improved by 1) limiting to in-region training data; for example, training a network on 2015–2019 data and applying it to 2020–2023 data for the same area, 2) increasing the coverage and diversity of terrain represented in the training data, and/or 3) further data augmentation and network architecture refinement. Generalization of CNNs in remote sensing is a widespread challenge and an area of ongoing research (e.g., Mnih, 2013; Maggiori et al., 2017; Ahmad et al., 2021; Gong et al., 2021), and our correction approach will likely benefit from future work in these research areas.

2.5.8 Limitations and future work

Our CNN correction is based on the differing spatial and topographic characteristics of atmospheric noise and slow-moving periglacial features. This approach involves two major assumptions.

The first assumption is that the true rock glacier displacement contribution to observed LOS phase offset is negligible in 6- and 12-day interferograms. This may not hold where InSAR is applied to detect faster displacements, like those associated with surging glaciers or earthquakes, which display significant phase offsets in short temporal baseline interferograms. However, atmospheric correction becomes less important in these high SNR cases, as standard time series approaches applied to densely connected interferogram networks will effectively capture the kinematic behavior of fast-moving features.

The second assumption is that apparent displacement signals in the characteristic displacement signal maps derived from time series processing have fundamentally different spatial and topographic characteristics than the atmospheric noise. We expect the performance of our CNN correction to degrade as the spatial and topographic characteristics of the displacement signal of interest approach those of the local atmospheric noise. In these cases, the CNN atmospheric noise prediction could include some true displacement signals or fail to include some atmospheric noise.

As such, our current CNN correction approach is well-suited for alpine features like smaller slow-moving landslides, stagnant debris-covered glaciers, rock glaciers, and ice-cored moraines. It is likely less appropriate for features spanning several kilometers with limited displacement signals, indistinct boundaries, and limited adjacent stable terrain (e.g., ice sheet interior). For example, the elevation-dependent displacement signals related to inflation or deflation of large volcanoes may appear visually similar to stratified atmospheric noise.

We recommend that the dominant spatial wavelength of the atmospheric noise should be at least 2x the dominant spatial wavelength of the feature of interest for best results. Future efforts should systematically evaluate the effect of different pixel size, radar wavelength, acquisition geometry, and geographic location on CNN correction performance.

Our CNN correction improved the accuracy of cumulative InSAR displacement measurements, removing apparent displacement in stable areas and correcting apparent rapid fluctuations in rock glacier displacement direction along the satellite LOS. Unfortunately, we did

not have prior knowledge of the true 3D surface displacement vectors for any of the features of interest, which would provide independent validation for our displacement time series results. In the future, we hope to train and evaluate a similar CNN correction for landslides and rock glaciers with multiple years of in situ surface displacement measurements.

2.6 Conclusions

We developed and implemented a statistical machine learning approach for correcting atmospheric noise in satellite interferometric synthetic aperture radar (InSAR) data over mountainous regions. We used on-demand InSAR processing and cloud computing resources to train, validate and test a denoising convolutional neural network (CNN) using characteristic atmospheric noise and displacement signals derived from Sentinel-1 interferograms over the Rocky Mountains in New Mexico, Colorado, and Wyoming, where slow-moving periglacial features above treeline are abundant. Our correction approach offers several advantages over existing techniques as it does not require external atmospheric data, is applied at the native spatial and temporal resolution of individual InSAR datasets, and can correct both stratified and turbulent atmospheric noise.

We found that our CNN correction significantly outperformed other existing corrections involving ERA5 reanalysis data, empirical phase, and a high-pass filter approaches. The CNN correction offered a 131%, 208% and 68% improvement, respectively, in structural similarity index measure (SSIM) metric that objectively quantifies similarity to the target displacement signal, as well as improvements in R^2 and RMSE. We found that the CNN successfully removed atmospheric noise from displacement time series for an out-of-region application area, revealing consistent downslope summer displacement of three selected features (two rock glaciers and a landslide).

Our CNN correction approach is suitable for applications involving 1) slow-moving features with negligible true displacement in short-baseline (6- and 12-day) interferograms and 2) distinct spatial scales for features of interest and atmospheric noise. In these circumstances, our CNN correction represents a step toward realizing the full potential of InSAR to monitor any coherent feature displacement with high spatial and temporal resolution and mm-scale precision. The improved CNN-corrected displacement measurements can reveal the seasonal to interannual kinematics of many other alpine glaciers and periglacial features without in-situ instrumentation, improving our understanding of their dynamics and ongoing response to climate change.

Chapter 3

QUANTIFYING DEGRADATION OF THE IMJA LAKE MORAINE DAM WITH FUSED INSAR AND SAR FEATURE TRACKING TIME SERIES

This chapter has been accepted by *The Cryosphere* as: **Brencher, G.**, Henderson, S. T., & Shean, D. E. ***Quantifying degradation of the Imja Lake moraine dam with fused InSAR and SAR feature tracking time series***.

3.1 Abstract

Glacial lake outburst flood (GLOF) hazards are often tied to the structural properties of the moraines that dam glacial lakes. Traditional investigations of moraine dam structure and degradation involve costly and logistically challenging in situ geophysical and repeat topographic surveys, which can only be performed for a small number of sites. We developed a scalable satellite remote sensing approach using interferometric synthetic aperture radar (InSAR), InSAR coherence, and SAR feature tracking to precisely measure moraine dam surface displacement and map the extent of buried ice. We combined time series from ascending and descending Sentinel-1 orbits to investigate vertical and horizontal surface displacement for the period 2017–2024 with ~12-day temporal sampling.

We applied our approach to quantify degradation of the Imja Lake moraine dam in the Everest Region of Nepal. We validated our SAR-based displacement measurements using 3D displacement measurements from very-high-resolution satellite stereo digital elevation models that span the study period. We found that a 0.3 km^2 area of the moraine dam has cumulatively subsided ~90 cm over the 7-year study period. Seasonal change in InSAR coherence provides evidence for buried ice throughout the moraine dam. We observed consistent downward and eastward displacement throughout the colder months, which we attribute to ice flow. The magnitude of downward vertical surface velocity increases in the warmer months, likely due to melting of buried ice. Our observations provide new insights into the timing and magnitude of the processes that control moraine dam evolution, with broader implications for regional GLOF hazard assessment and mitigation.

3.2 Introduction

Glaciers in High Mountain Asia (HMA), Earth's largest glacierized region outside the poles, are expected to lose between ~29% and 67% of their total mass by the end of the century (Hock et al., 2019; Rounce et al., 2020). Worldwide, glacier thinning and retreat are associated with an increase in the formation of glacial lakes (Shugar et al., 2020). These lakes constitute a significant hazard, as they can drain in sudden, catastrophic glacial lake outburst flood (GLOF) events that flood downstream valleys and can result in loss of life and damage to infrastructure (Mool et al., 2011; Riaz et al., 2014; Carrivick and Tweed, 2016). For example, the September 2023 GLOF event in Sikkim, India, destroyed the Teesta III hydroelectric dam, washed away 15 bridges, stranded 3,000 tourists, and left at least 74 dead with more missing (Ali Badal, 2023; Choudhury and Hussain, 2023; Kumar and Travelli, 2023; Sebastian, 2023).

Effective management of GLOF hazards involves 1) identification and monitoring of glacial lakes that pose a significant hazard, and 2) mitigation of this hazard before GLOFs occur, potentially through engineering solutions. The former has been primarily accomplished through hazard assessments ranging from a single lake (e.g., Rana et al., 2000; Budhathoki et al., 2010; Wang et al., 2018; Sattar et al., 2021) to regional scale (e.g., Mool et al., 2011; Wang et al., 2012; Fujita et al., 2013; Allen et al., 2016; Rounce et al., 2017). Glacial lakes can be dammed by moraines, bedrock, and glacier ice (Rick et al., 2022). Where glacial lakes are dammed by moraines, hazard assessments frequently consider moraine dam stability, the presence of buried ice within moraine dams, potential GLOF triggering events, and downstream impacts (Rounce et al., 2016). Where moraine dam instability is not identified as a primary GLOF trigger mechanism, melting of buried ice can increase lake area, reduce dam width and height, and provide potential pathways for seepage and piping (Richardson and Reynolds, 2000; Emmer and Cochachin, 2013; Medeu et al., 2022). As such, the evolution of moraine dams may substantially increase GLOF likelihood.

As glacial lakes, moraine dams, and the surrounding landscape change over time (Huggel et al., 2010; Shugar and Clague, 2011; Kellerer-Pirklbauer et al., 2012; Ravanel et al., 2018), hazard assessment and mitigation tasks require repeated observations (Fujita et al., 2009; Rounce et al., 2017; Emmer et al., 2018). Satellite remote sensing has been used to create glacial lake inventories, track glacial lake development (e.g., Fujita et al., 2009; Nie et al., 2018; Shugar et al., 2020), and recently, to monitor glacial lake dam and bank evolution (Haritashya et al., 2018; Scapozza et al.,

2019; Wangchuk et al., 2022; Yang et al., 2022; Jiang et al., 2023; Yang et al., 2023; Yu et al., 2024). These remote measurements of surface displacement can be used to assess moraine dam stability, infer moraine dam structure, and quantify changes in moraine dam topography, providing critical information for GLOF hazard assessments and potentially guiding prioritization for in situ surveys and hazard mitigation strategies.

3.2.1 Moraine dam evolution

Ice-cored moraines are found globally in modern glacier forelands (Østrem, 1959). These features, which range in size from the decameter to kilometer scale, are typically formed at glacier margins when ice overlain by insulating debris cover is disconnected from the active glacier through differential ablation (Østrem, 1959; Kjær & Krüger, 2001; Lukas, 2011). Ice-cored moraines often impound meltwater to form proglacial glacial lakes, becoming moraine dams.

Once cut off from the active upstream glacier, ice-cored moraine dam degradation proceeds through multiple related processes, which change the topographic and structural characteristics of moraine dams over time. Understanding these mechanisms is important for both methods development and interpretation of surface displacement measurements from remote sensing observations. Several processes contribute to degradation of ice-cored moraines, including 1) downwasting, 2) thermokarst, 3) backwasting, and 4) debris reworking.

Downwasting, or vertical subsidence of ice-cored features, occurs due to melting and sublimation of buried ice (Lukas, 2011). Downwasting rates have been observed to vary between 0.003 and 4.8 m/yr (Table B.1), and spatiotemporal variability is correlated with mean annual temperature and debris thickness (Irvine-Fynn et al., 2011). The downwasting rate of ice-cored moraine dams may have significant implications for GLOF hazards, as it relates directly to changes in the freeboard height of the dam and the shape of the shoreline, leading to lake expansion.

Thermokarst occurs where melting of glacial ice or ground ice beneath unconsolidated sediments creates pseudokarst landforms, including kettles and sinkholes (Healy, 1975). Thermokarst landforms may contribute to rapid surface lowering, as melting of buried ice is accelerated by the presence of liquid water (Richardson & Reynolds, 2000b). Remnant glacial structures including water- and debris-filled crevasses likely exert a primary control over thermokarst landform location (Richardson & Reynolds, 2000b).

Backwasting, or lateral retreat of near-vertical ice walls and ice-cored slopes, is the dominant ice loss mechanism for most ice-cored moraines (Krüger & Kjær, 2000; Lukas, 2011). Unlike

downwasting, backwasting rates are primarily controlled by local factors, such as ice cliff exposure direction, ice cliff albedo, wind conditions, cloudiness, and proximity to water bodies, rather than debris thickness or mean annual temperature (Johnson, 1971; Schomacker, 2008). Backwasting rates can vary from 0.65 to 24 m/yr (Schomacker, 2008).

Lakes act as heat and sediment reservoirs, and ice-cored moraine dam degradation can be accelerated by interaction with the glacial lake and surface ponds (Johnson, 1971; Driscoll, 1980; Pickard, 1983; Watanabe et al., 1995). Removal and redistribution of moraine dam debris via debris flows and fluvial transport exposes ice-cored slopes to backwasting, changes the spatial pattern of debris insulation, and removes debris from the system, which can affect structural stability (Lukas et al., 2005; Schomacker & Kjær, 2008). Debris flows cause backwasting of the shoreline away from the lake, decreasing moraine width and increasing lake area (Watanabe et al., 1995). In addition to debris flows, other types of landslides have been observed on moraine dams (Richardson & Reynolds, 2000b). Debris and ice mass movements can also indirectly contribute to GLOF risk by changing the shape and structure of moraine dams, and/or directly trigger GLOFs by creating waves that overtop the moraine dam (e.g. Zhang et al., 2025).

3.2.2 Remote sensing of moraine dam degradation

C-band SAR can operate during both day and night, penetrate clouds, and penetrate thin, dry snow (Bürgmann et al., 2000; Sun et al., 2015). Both InSAR and SAR feature tracking (in original range–azimuth coordinates) can provide measurements of surface motion in the radar line-of-sight (LOS) direction.

InSAR relates the phase offset from successive SAR acquisitions to surface displacement, potentially with mm-level accuracy (Bürgmann et al., 2000; Rosen et al., 2000). InSAR has frequently been applied to measure displacement of ice-rich features, including glaciers and ice sheets (Massonnet and Feigl, 1998; Rosen et al., 2000), rock glaciers (Bertone et al., 2022), and permafrost (Zhang et al., 2022). Recently, InSAR has also been applied to quantify surface movement of glacial lake dams (Scapozza et al., 2019; Yang et al., 2022; Jiang et al., 2023; Yang et al., 2023; Yu et al., 2024).

While InSAR offers unparalleled precision, several factors can reduce measurement accuracy, including atmospheric noise, layover, and radar shadow. InSAR requires 1) coherent surface change with similar scatterer characteristics between radar acquisitions, and 2) sufficiently small phase-change gradients between adjacent pixels to allow reliable phase unwrapping (Itoh, 1982;

Handwerger et al., 2015). While errors caused by atmospheric conditions and acquisition geometry can be avoided or corrected, errors caused by rapid surface change may cause underestimation of true LOS surface velocity. InSAR coherence, a unitless measure of the "sameness" of surface scatterers between acquisitions, can be used to discriminate between reliable and unreliable InSAR measurements (e.g., Schmidt and Bürgmann, 2003). InSAR coherence can also be used to identify significant change in surface characteristics, and low coherence has been used to map the extent of desert erosion (e.g., Cabré et al., 2020), landslides (e.g., Ohki et al., 2020; Jacquemart and Tiampo, 2021), flooding (e.g., Chini et al., 2019), and debris-covered glaciers (e.g., Atwood et al., 2010; Frey et al., 2012; Lippl et al., 2018).

Feature tracking (pixel offset) techniques offer an alternative method to retrieve surface displacement measurements from SAR data when InSAR is not possible due to loss of coherence or large displacement gradients between adjacent pixels. Feature tracking involves 2-D cross-correlation of backscatter amplitude values in successive radar images to measure offsets in the range (across-track) and azimuth (along-track) directions. This technique has been used to measure glacier surface velocity (e.g., Fahnestock et al., 1993; Strozzi et al., 2002, 2008; Fahnestock et al., 2016) and fast-moving features surrounding glacial lakes (Scapozza et al., 2019; Yang et al., 2022). While feature-tracking measurements offer lower resolution and precision than InSAR measurements, and measurement quality can be degraded by large surface changes, they are more robust to large surface displacements and rapid surface change (Joughin, 2002; Zheng et al., 2023).

Here, we present a method to create a single cumulative displacement time series from InSAR and feature tracking measurements using a variation of the small baseline subset approach (SBAS; Berardino et al., 2002; Schmidt and Bürgmann, 2003). The SBAS approach uses a network of displacement observations to solve for the displacement of a given pixel between successive acquisitions. It was developed for InSAR interferograms (Berardino et al., 2002) and has subsequently been applied to SAR feature tracking (e.g., Casu et al., 2011; Euillades et al., 2016; Guo et al., 2020; Samsonov et al., 2021) and optical feature tracking offsets (e.g., Bontemps et al., 2018; Altena et al., 2019; Lacroix et al., 2019). The SBAS approach, however, has issues during periods with poor coherence between successive acquisitions, which can introduce errors in the resulting displacement time series. This limitation has previously been addressed by constraining the time series inversion using geophysical constraints parameterized in time (López-Quiroz et al., 2009) and spatial covariance (Jolivet and Simons, 2018). We address this issue by constraining the

per-pixel InSAR time series inversion with additional observations from SAR feature tracking, which is more robust to decorrelation than InSAR. Our approach does not require prior assumptions about the displacement patterns of pixels over time.

Previous studies combined InSAR and feature tracking to produce self-consistent surface displacement maps over large ice sheets and ice caps, where surface velocity is often too fast for InSAR alone (e.g., Joughin, 2002; Gray et al., 2005; Liu et al., 2007; Sánchez-Gámez and Navarro, 2017; Joughin et al., 2018). To our knowledge, no studies have combined these techniques to measure year-round surface displacement of ice-cored moraine features, where we expect seasonal variability in surface melt rates and associated debris reworking.

3.3 Imja Lake Study Site

Imja Lake is a ~2.8 by 0.6 km glacial lake covering ~1.6 km^2 in the Everest region of Nepal (Figure 3.1). Imja Lake is impounded on its north and south sides by lateral moraines, above which rise the steep slopes of adjacent mountain peaks: Imja Tse to the north and Ombigaichen to the south. To the east, the ~860 m wide calving front of the Lhotse Shar and Imja Glaciers terminates in the lake. Imja Lake formed in the early 1970s as supraglacial lakes coalesced on the stagnant, debris-covered tongue of Imja Glacier. It expanded at a roughly linear rate of 0.02 km^2/yr until 1997, then at a rate of 0.03 km^2/yr from 1997 to 2020 (Watanabe et al., 2009; Gupta et al., 2023).

To the west, Imja Lake is impounded by an ice-cored terminal moraine covering 0.62 km^2. The Imja Lake moraine dam (hereafter, the moraine dam) has a low-relief, hummocky surface with ridges and furrows, ponds, and ice cliffs, with a mean elevation of 4,982 m (orthometric height above the EGM2008 geoid). The lake drains across the moraine via a series of linked ponds. An artificial drainage channel was constructed downstream of the westmost pond in 2016 (UNDP, 2012; Khadra, 2016). The moraine dam has an enclosed concave-up surface and broadly slopes down toward the lake and surface ponds (Figure B.1). Immediately northwest of the moraine dam is a lobe of the debris-covered Lhotse Glacier.

Geophysical studies using ground-penetrating radar (GPR) and electrical resistivity tomography (ERT) have documented the extent and thickness of buried ice in the moraine dam, with continuous buried ice tens of meters thick to the east transitioning to discontinuous blocks of buried ice to the west (Hambrey et al., 2008; Somos-Valenzuela et al., 2012; Dahal et al., 2018).

Two studies quantified Imja Lake moraine dam degradation rates using in situ measurements obtained with traditional and Global Positioning System (GPS) surveying techniques (Watanabe

et al., 1995; Fujita et al., 2009), and one study measured moraine dam displacement by differencing the 30-m Shuttle Radar Topographic Mission Global 1 (SRTM-GL1) DEM acquired in February 2000 with a 2-m WorldView stereo DEM acquired in February 2016 (Haritashya et al., 2018). These studies document meter- to sub-meter-scale subsidence and horizontal motion over the surface of the moraine dam, with the fastest motion occurring in the northeastern area.

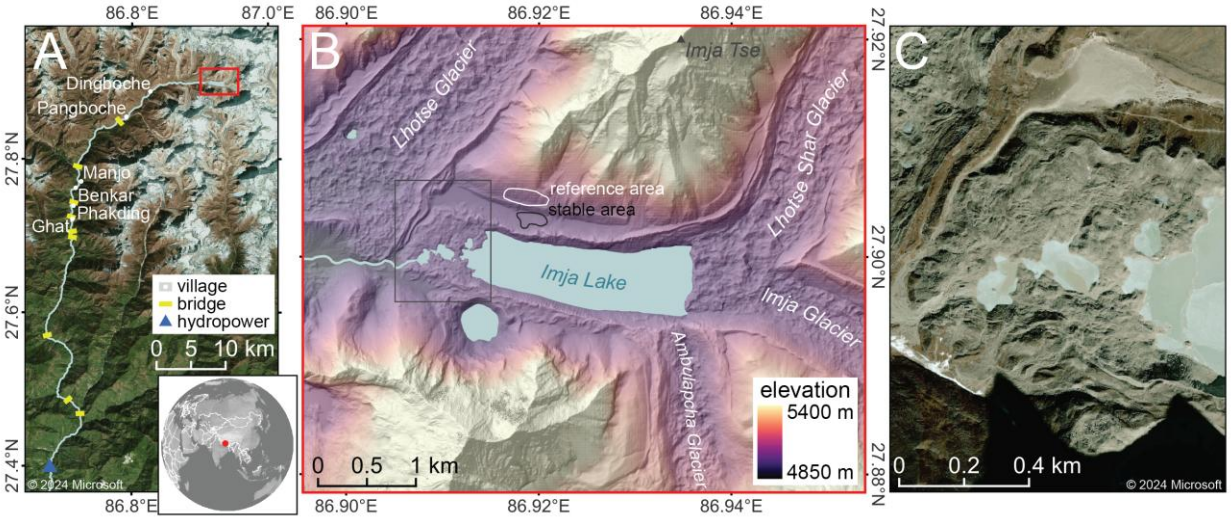

Figure 3.1: Context maps for the Imja Lake moraine dam. A) Microsoft satellite basemap showing the area downstream of Imja Lake, including settlements and infrastructure along the Dudh Koshi river. B) Color-shaded relief map from median composite of EarthDEM strips acquired during the 2014-2019 period (Porter et al., 2022). C) © Microsoft satellite basemap detail of the Imja Lake moraine dam. Microsoft product screen shots reprinted with permission from Microsoft Corporation.

3.4 Data

We downloaded all available Copernicus Sentinel-1 C-band single-look complex (SLC) images collected over the study site from January 1, 2017 to March 1, 2024. The revisit time for this period was generally 12 days; we did not include data before 2017 due to longer gaps between acquisitions. All images were acquired in interferometric wide (IW) swath mode with vertical co-polarization (VV) along ascending (satellite moving north and looking east) relative orbit 12 and descending (satellite moving south and looking west) relative orbit 121. To minimize download time and storage, we limited processing to SLC bursts covering our study site, with a total of 214 ascending bursts (Burst ID 012_023790_IW1) and 227 descending bursts (Burst ID 121_258661_IW2). Burst dimensions are ~20 km in azimuth (along-track direction) by ~85 km in

range (across-track direction), with pixel spacing of 14.1 m in azimuth and 2.3 m in range (Figure B.2).

We used the 2022_1 release of the 30-m Copernicus GLO-30 DEM (European Space Agency, 2021) to remove the topographic component of phase and geocode the SAR data products. The GLO-30 DEM product has an absolute vertical accuracy (LE90) of <4 m and an absolute horizontal accuracy (CE90) of <6 m (NSSDA, 1998; European Space Agency, 2022).

We downloaded daily 2-m air temperature and total precipitation data from the Pyramid Weather Station (UCAR/NCAR-Earth Observing Laboratory et al., 2011), which is located ~11 km northwest of the Imja Lake moraine dam at an elevation of 4951 m, about 30 m lower than the elevation of the moraine dam. We aggregated available daily measurements between October 1, 2002 and December 31, 2009 to derive monthly climatology (mean, minimum, maximum and standard deviation) for the ~7-year period.

3.5 Methods

By combining InSAR and feature tracking measurements, we leverage the strengths of both methods to provide a more complete record and improved understanding of moraine dam kinematics. This approach captures large seasonal variations, including slow cold-season displacement with InSAR and rapid warm-season surface change with feature tracking. However, feature tracking measures the displacement of discrete kernels of neighboring pixels, so it provides inherently lower spatial resolution than per-pixel InSAR displacements (Joughin, 2002). By combining InSAR and feature tracking measurements, we effectively degrade the spatial resolution of the InSAR measurements in exchange for improved displacement accuracy. This compromise is essential for applications involving debris-covered ice, as large surface change during the warmer months reduces InSAR coherence between Sentinel-1 acquisitions. It is also important to note that extreme, rapid and/or highly localized surface change from processes like iceberg calving, backwasting, and fast-moving debris flows will still present challenges for the combined approach, potentially resulting in underestimation of surface deformation. Nonetheless, combining InSAR and feature tracking offers improved displacement time series that are more accurate than those produced independently using either method alone.

3.5.1 SAR data processing

To extract surface displacement information from the Sentinel-1 SLC products, we generated both interferograms and feature tracking offsets using the Hybrid Pluggable Processing Pipeline (HyP3) ISCE2 Plugin (Hogenson et al., 2023), which enables batch processing with the Jet Propulsion Laboratory InSAR Scientific Computing Environment, ISCE version 2.6.3 (Rosen et al., 2012; Figure 3.2).

We used the insar_tops_burst workflow to process interferograms with five looks (effectively, a spatial averaging factor of five) in the range direction and one look in the azimuth direction, geocoded to 20 m by 20 m pixels (Figure B.2; Figure B.3). To form redundant networks of ascending and descending interferograms, we created interferograms for each burst acquisition using the three nearest burst acquisitions in time. In total, we processed 636 ascending burst interferograms and 675 descending burst interferograms (Table 3.1).

We selected a local "reference area" and "stable area" (Figure 3.1) as close to the moraine dam as possible, with high mean InSAR coherence for the entire study period. The reference area (~0.06 km^2, ~144 pixels) was used to remove atmospheric noise, while the stable area (~0.03 km^2, ~87 pixels) was used to assess derived measurement accuracy and precision.

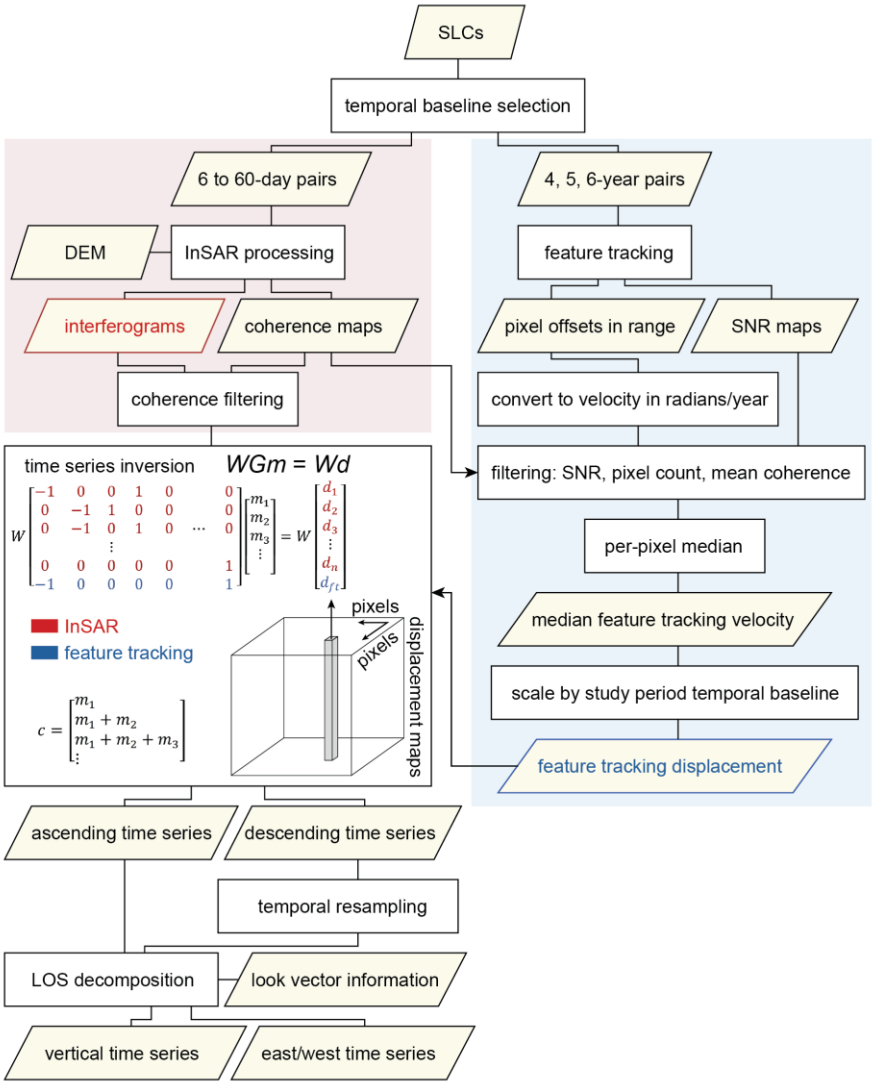

Figure 3.2: Flowchart illustrating our combined InSAR and feature-tracking approach. The time series inversion schematic is adapted from Schmidt and Bürgmann (2003). The equation relates the observed phase change (d) from each of n interferograms and a single cumulative feature tracking displacement product (d_{ft}) to the change in phase between acquisitions (m) at a given pixel. The design matrix G describes the temporal baseline of each pair, with -1 corresponding to the primary acquisition date and 1 corresponding to the secondary acquisition date. The diagonal weight matrix W includes coherence values. The cumulative phase change time series (c) is prepared from the inversion results for each pixel.

We co-registered all SLCs and performed feature tracking on the SLC products in radar coordinates at the native resolution using a modified version of the insar_tops_burst workflow and

the ‘dense ampcor’ ISCE2 feature tracking routine (Figure 3.2). We selected kernel sizes and skip sizes based on the characteristics of the moraine dam and an empirical sensitivity analysis. The moraine dam is roughly square (~800 by 900 m) in ground range. To maximize the number of unique feature-tracking measurements over the moraine dam that do not extend beyond its edges, we chose to use kernel sizes and skip sizes that were roughly square in ground range. Kernels that span the moraine dam and the surrounding area are likely to produce poor matches due to intra-kernel variability in the magnitude and direction of surface displacement.

Smaller kernel sizes will result in poor matching, while larger kernel sizes will reduce the spatial resolution of the output feature tracking measurements. We found that a kernel size of 10 pixels in azimuth and 60 pixels in range (141 by 138 m in ground range) was the smallest possible square kernel size that provided high-quality matches. We used a skip size of 4 pixels in azimuth and 20 pixels in range, or 56 by 46 m in ground range, to minimize interpolation without increasing processing time. Smaller skip sizes substantially increased processing time with negligible improvement, while larger skip sizes required more interpolation between sparse measurements. The search window size was set to 10 pixels in azimuth and 30 pixels in range, or 141 m by 69 m, which is much larger than the expected cumulative surface displacement of the moraine dam over the study period.

To estimate the accuracy of our feature tracking measurements, we first calculated the ground-range resolution of the slant-range SLCs. Following Notti et al. (2010) we computed the ratio of slant-range resolution to ground-range resolution (R-index) for the SLCs, considering terrain slope, terrain aspect, and the radar LOS vectors. We used this index to calculate the SLC resolution in the ground-range direction (Figure B.4). Over the moraine dam, the median ground-range resolution was 4.38 m for the ascending SLCs and 3.82 m for the descending SLCs. Assuming perfect SLC co-registration and a conservative feature-tracking precision of 0.2 pixels (Fialko et al., 2001; Strozzi et al., 2002), the theoretical accuracy of our feature tracking measurements is 2.82 m in azimuth, 0.88 m in range for the ascending SLCs, and 0.76 m in range for the descending SLCs. Given the expected mean surface velocity of <1 m per year for the moraine dam (Fujita et al., 2009), we generated feature tracking offsets with multi-year temporal baselines to maximize displacement between acquisitions. As longer temporal baselines could result in degraded quality due to significant changes in surface scatterer characteristics, we generated offsets for all possible feature tracking image pairs with temporal baselines of 4, 5, and 6 years (Table 3.1). In total, we

generated 192 ascending and 210 descending feature tracking offset maps with corresponding signal-to-noise ratio (SNR) maps. The expected precision in the range direction for the 4-, 5-, and 6-year feature tracking offset maps is 0.22 m/yr, 0.18 m/yr, and 0.15 m/yr, respectively.

Table 3.1: Summary of InSAR and feature tracking datasets.

Orbit	Acquisitions	Dataset	Pair count	Mean temporal baseline	Temporal baseline range
ascending track 12, Burst ID 012_023790_IW1	214	InSAR	636	24.3 days	12-48 days
		feature tracking	192	4.7 years	4-6 years
descending track 121, Burst ID 121_258661_IW2	227	InSAR	675	23.0 days	6-60 days
		feature tracking	210	4.7 years	4-6 years

3.5.2 Time series inversion

Since both InSAR interferograms and the component of feature-tracking motion in the slant-range direction are measurements of displacement along the radar LOS (Figure B.5), these measurements can be combined to produce more accurate estimates of LOS surface displacement over time. We used the Miami INsar Time-series software in PYthon (MintPy) package (Yunjun et al., 2019) to prepare cumulative LOS displacement time series from all InSAR and feature tracking observations for each relative orbit.

To avoid cumulative displacement errors caused by poor feature tracking matches, we aggregated the available feature tracking offset products over time and prepared a single median velocity product for each relative orbit. These median velocity products were then used to constrain the cumulative displacement solution from the network of InSAR interferograms from the same relative orbit (Figure 3.2).

Several preprocessing steps were required to produce the median feature tracking velocity products for each relative orbit. We first converted all feature tracking offsets in range from pixels to C-band phase in radians, the same unit as our interferograms. We then masked all pixels with signal-to-noise ratio lower than 8 to remove unreliable measurements (Casu et al., 2011). This threshold was set to minimize the resulting magnitude of noise in the stable area. Next, we calculated the median surface displacement in range observed in the reference area and subtracted it from each offset map to mitigate potential co-registration bias. At this point we divided each feature tracking offset map in range by its temporal baseline and computed a per-pixel median rate for the entire time series stack, resulting in a single median velocity product for each relative orbit.

We masked any pixels with less than 10 feature-tracking observations or a mean InSAR coherence greater than 0.85 during late summer (day of year 220–280, the lowest coherence period), as the InSAR measurements alone should provide accurate displacement time series for those pixels. We also computed the per-pixel standard deviation of the feature tracking offset maps to assess variability due to real changes and measurement error.

Prior to time series inversion, we removed interferograms with mean coherence less than 0.6 over the moraine dam, unless they were required to form a continuous network. This filter effectively removed interferograms impacted by intermittent snow accumulation, larger (>200 m) perpendicular baselines, or large surface change over the temporal baseline. To convert our per-pixel median feature tracking velocity (displacement rate) products into a cumulative displacement product analogous to an interferogram for the full 2017 to 2024 study period, we multiplied by the appropriate temporal baseline for the ascending (7.29 years) and descending (7.36 years) time series. The cumulative feature tracking displacement product was assigned a uniform coherence of 0.6, which is the minimum allowable mean coherence for the moraine dam area. The network of per-pixel displacement measurements, including the interferograms and the ~7-year cumulative feature tracking displacement product, was then inverted using phase variance, a measure of expected phase noise calculated using coherence (Tough et al., 1995; Hanssen, 2001), as a weight (Yunjun et al., 2019). As coherence is used to calculate phase variance, the cumulative feature tracking displacement product is weighted lower than interferograms where coherence is greater than 0.6.

Following the time series inversion, we computed and removed the median apparent displacement over the reference area from the LOS displacement estimate at each time step to mitigate atmospheric noise. The standard deviation of the median apparent displacement values over the reference area was 3.1 cm for the ascending data (n=214) and 0.5 cm for the descending data (n=227). The moraine dam is approximately 900 by 800 m and the reference area is only 250 m northwest of the moraine dam. Given the expected correlation length scales of several km for atmospheric noise in mountains (e.g., Bekaert et al., 2015), the apparent displacement in the reference area should be similar to the atmospheric noise over the moraine dam. To evaluate remaining measurement uncertainty, we first computed the mean apparent displacement in the stable area at each time step, then computed the standard deviation of this spatial mean over all time steps to characterize the magnitude of remaining atmospheric noise. The standard deviation

was 0.2 cm for the ascending orbit and 0.4 cm for the descending orbit. Finally, we delineated a "moving area" polygon over the moraine dam including pixels with mean velocity of ~1 cm/yr above the background noise level in the ascending and descending cumulative displacement time series.

3.5.3 Surface displacement from LOS decomposition

Given two cumulative displacement measurements over the moraine dam, each with a different LOS from different relative orbits, we can solve for two components of the 3D surface displacement vector (Wright et al., 2004; Fuhrmann & Garthwaite, 2019). Since the mean LOS incidence angle at the site is 35.8° for the ascending burst and 37.8° for the descending burst, and the mean LOS azimuth (defined clockwise from north) is 79.3° for the ascending burst and 280.5° for the descending burst, we chose to solve for the vertical (up/down) and horizontal (east/west) surface displacement components. We assume that any north/south displacement would be negligible when projected into the available LOS vectors. The LOS unit vector ($\hat{l}$) of each pixel can be defined as:

$$\hat{l} = \cos(\gamma)\sin(\theta)\,\hat{n} + \sin(\gamma)\sin(\theta)\,\hat{e} - \cos(\theta)\,\hat{z}, \tag{3.1}$$

for ascending ($\hat{l}_{asc}$) and descending ($\hat{l}_{des}$) geometry, where γ is the LOS azimuth angle, θ is the LOS incidence angle defined relative to surface-normal direction, and $\hat{n}$, $\hat{e}$ and $\hat{z}$ represent the north, east, and up axes, respectively, which define the up/down ($\hat{l}_{ud} = 0\hat{n} + 0\hat{e} + 1\hat{z}$) and east/west ($\hat{l}_{ew} = 0\hat{n} + 1\hat{e} + 0\hat{z}$) unit vectors.

If we assume that the observed cumulative ascending ($d_{asc}(t)$) and descending ($d_{des}(t)$) displacement at time t is only sensitive to the up/down ($d_{ud}(t)$) and east/west (($d_{ew}(t)$) surface displacement components, we obtain the following equations:

$$d_{asc(t)} = \hat{l}_{asc} \cdot \hat{l}_{ud} * d_{ud}(t) + \hat{l}_{asc} \cdot \hat{l}_{ew} * d_{ew}(t), \tag{3.2}$$

$$d_{des(t)} = \hat{l}_{des} \cdot \hat{l}_{ud} * d_{ud}(t) + \hat{l}_{des} \cdot \hat{l}_{ew} * d_{ew}(t), \tag{3.3}$$

We can then rearrange these equations to solve for the up/down and east/west displacement components:

$$d_{ud}(t) = \frac{\hat{l}_{asc}\cdot\hat{l}_{ew} * d_{des}(t) - \hat{l}_{des}\cdot\hat{l}_{ew} * d_{asc}(t)}{\hat{l}_{asc}\cdot\hat{l}_{ew} * \hat{l}_{des}\cdot\hat{l}_{ud} - \hat{l}_{des}\cdot\hat{l}_{ew} * \hat{l}_{asc}\cdot\hat{l}_{ud}}, \tag{3.4}$$

$$d_{ew}(t) = \frac{\hat{l}_{asc} \cdot \hat{l}_{ud} * d_{ud}(t) - d_{asc}(t)}{\hat{l}_{asc} \cdot \hat{l}_{ew}}, \quad (3.5)$$

To accomplish this, we interpolated the descending orbit cumulative displacement time series to match the acquisition dates of the ascending orbit cumulative displacement time series, and then solved for the vertical and east/west displacement at every pixel at each ascending acquisition time.

To compute the vertical and east/west surface velocity, we divided the observed cumulative displacement by the time between acquisition dates. We aggregated these velocity component time series to compute the per-pixel monthly median values over the entire study period.

To identify and map areas of buried ice, we calculated a winter coherence map using the per-pixel median coherence for all 12-day temporal baseline observations in winter to early spring (day of year 0–100), and a summer coherence map using the per-pixel median coherence for observations in late summer (day of year 220–280). We attribute large observed seasonal decreases in InSAR coherence to surface changes caused by melting of buried ice during warm months.

3.5.4 Preparation of a high-resolution 3D displacement validation dataset from stereo DEMs

We identified and processed two Maxar WorldView-3 in-track stereo image pairs spanning the 2017–2024 study period, including the February 11, 2016 WorldView-3 stereo pair (catalog IDs 104001001854B000, 10400100175C2D00) used by Haritashya et al. (2018), and a new in-track stereo pair acquired on January 30, 2025 (10400100A0D45D00, 10400100A193EE00). We used the latest version (3.5.0-alpha) of the Ames Stereo Pipeline (Shean et al., 2016; Beyer et al., 2018) and the processing settings outlined by Bhushan and Shean (2021) and Bhushan et al. (2024) to prepare the corresponding DEMs with 1-m posting.

We co-registered the January 30, 2025 DEM to the February 11, 2016 DEM using the demcoreg (Shean et al., 2023) implementation of the Nuth and Kääb (2011) algorithm (Figure B.6). We then corrected residual in-track "jitter" and cross-track "CCD array geometry" artifacts (see Shean et al., 2016) by fitting a Savitzky-Golay filter (window of 101 px, 2nd-order polynomial basis function) to the median of elevation difference residuals over static surfaces for each row and column (Figure B.7).

We prepared final cumulative elevation change and elevation change rate maps for the ~9-year period (Figure 3.3C). The final elevation difference residuals over all exposed static control surfaces used during co-registration had a median of 0.00 m and normalized median absolute deviation (NMAD) of 0.14 m (<0.5 px) for the ~9-year period (Figure B.8), corresponding to a 1-

sigma uncertainty of ~1.6 cm/yr. The elevation difference residuals over the stable area used for calibration of the InSAR and feature-tracking results (see Figure 3.1 for context) had a mean bias of -0.54 cm and a standard deviation of 7.77 cm, corresponding to apparent vertical displacement rates of -0.06 cm/yr and 0.87 cm/yr, respectively.

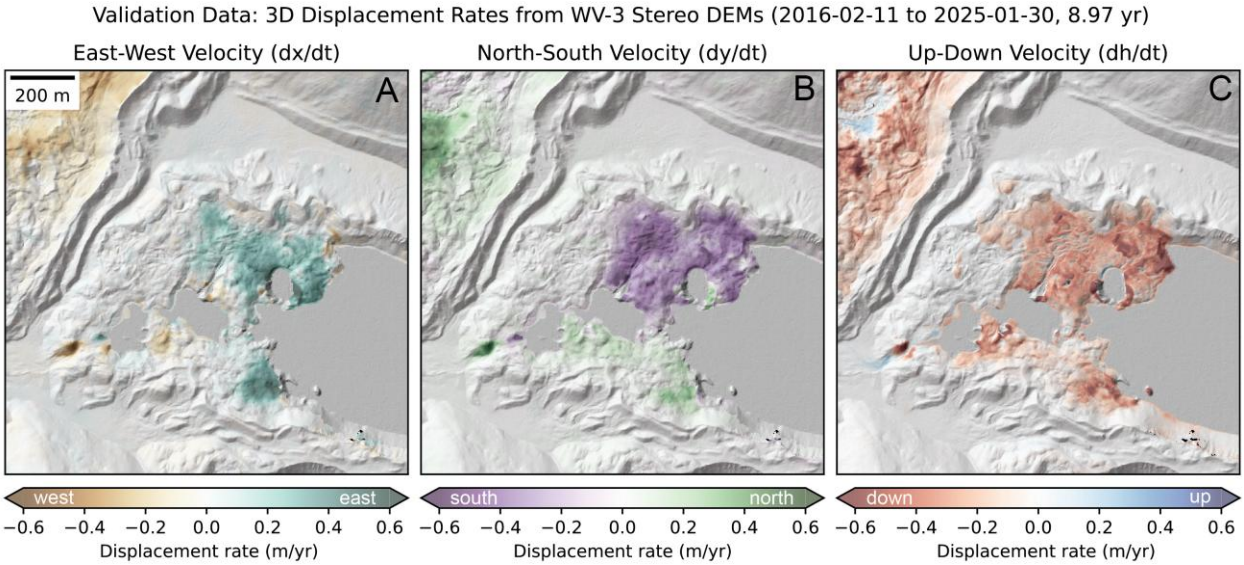

Figure 3.3: Long-term (~9-year) east/west, north/south, and vertical surface velocity validation data, calculated from feature tracking of co-registered 1-m DEM hillshade products and a vertical DEM difference product. These products document the spatial extent and magnitude of moraine dam surface motion during the study period with enough detail to capture signals associated with ice cliff retreat, downslope flow, and anthropogenic cut/fill and sediment deposition associated with the October 2016 lowering project (UNDP, 2012; Khadra, 2016). These products serve as validation for the combined InSAR and SAR feature-tracking time series, which offer detailed temporal evolution of these processes (Figure 3.4, Figure 3.7), with reduced spatial resolution.

Following co-registration, we created shaded relief maps for the two DEMs using the "combined" hillshade option in the gdaldem utility (Rouault et al., 2024). We digitized the lake and pond shoreline using the January 30, 2025 orthoimages and shaded relief map, and masked the hillshade products over water. We then used the Ames Stereo Pipeline correlator to prepare dense horizontal sub-pixel displacement maps tracking surface motion between the hillshade products (disp_mgm_corr.py utility of Bhushan and Shean [2025]), following the methods outlined in Bhushan et al. (2024). The final horizontal velocity (m/yr) components (Figure 3.3A–B) were calculated based on the 1-m DEM pixel size and the ~9-year time period. To assess the uncertainty

of the validation data, we computed the mean and standard deviation of the observed east/west (-1.5 cm/yr, 2.3 cm/yr, respectively) and north/south (-2.0 cm/yr, 1.9 cm/yr, respectively) velocity components over the stable area.

3.5.5 Validation of combined InSAR and SAR feature-tracking products

To enable direct comparison of our time series results with the DEM-derived validation data, we first masked pixels with a cold-season coherence of <0.7 (Figure 3.4). These areas contain rapid winter displacement caused by backwasting of ice cliffs, which have large signals in the DEM difference products that we do not expect to capture with Sentinel-1 SAR observations. We next resampled the 1-m DEM-derived products to match the 20-m grid of the Sentinel-1 SAR observations using median resampling, then smoothed the DEM-derived products to match the SAR feature tracking resolution using a 7 by 7 px (~140x140 m) Gaussian kernel. We computed zonal statistics (mean, standard deviation) for each velocity component (up/down, east/west) over the moraine dam. We also computed R^2 values between the InSAR/SAR velocity component and the corresponding DEM-derived validation velocity component. Finally, we used the north/south component of DEM-derived velocity to assess expected bias caused by our up/down and east/west decomposition (see Section B.2).

3.5.6 Comparison with InSAR-only and SAR-feature-tracking-only SBAS time series

We prepared two additional time series from the same Sentinel-1 SLC products: a standard InSAR-only SBAS time series and a SAR-feature-tracking-only SBAS time series. The InSAR-only SBAS time series used the same processing parameters as described for our combined approach, but did not include the scaled median velocity product from feature tracking for the inversion (blue row in Figure 3.2 system of equations). The SAR-feature-tracking-only SBAS time series included the same set of 4-, 5-, and 6-year temporal baseline pairs from our combined approach (192 ascending and 210 descending), as well as all possible image pairs with a 3-year temporal baseline (182 ascending and 132 descending), to ensure that the system was overdetermined (total ascending n=374, descending n=342). As with the DEM-derived validation, we computed zonal statistics (mean, standard deviation) for each velocity component (up/down, east/west) for the InSAR-only and SAR-feature-tracking-only SBAS time series over the moraine dam.

3.6 Results

3.6.1 InSAR coherence

We examined the mean InSAR coherence from all pairs on both the ascending and descending orbits over the Imja Lake moraine dam during the full 2017–2024 study period (Figure 3.4). The pair temporal baseline strongly controlled InSAR coherence over the moraine dam moving area (Figure 3.4C, G), with shorter temporal baselines (i.e., 12 days) offering higher mean coherence than longer temporal baselines (i.e., 36–60 days). Over the stable area, we observed no obvious relationship between temporal baseline and coherence (Figure 3.4D, H). Descending orbit coherence was on average slightly higher (mean 0.80) than ascending orbit coherence (0.74) over the moraine dam moving area for 12-day temporal baseline interferograms.

We observed a systematic seasonal cycle in InSAR coherence values over the moraine dam moving area. Coherence remained largely stable from the beginning of the year until early June. For 12-day temporal baseline interferograms, mean coherence decreased by approximately 0.3 from June through mid-September. Coherence then increased from mid-September through mid-November. This seasonal decrease in coherence lags about 60 days behind the timing of positive mean daily maximum temperatures at the nearby Pyramid Station (Figure 3.4C, G). No significant seasonal coherence change was observed over the stable area (Figure 3.4D, H).

The decrease in coherence is spatially constrained. Visual inspection of mean coherence maps from the annual high-coherence period (December–May) and low-coherence period (June–November) revealed that coherence decreased over almost all of the moraine dam moving area (Figure 3.4A-D). Several areas of low coherence in December–May expanded and displayed a further decrease in coherence during June–November. Areas of high-coherence in December–May also displayed a decrease in coherence during June–November. Low-coherence areas were mostly consistent between ascending and descending orbits, with minor discrepancies potentially explained by differences in the LOS and acquisition time (18:00 local for ascending, 06:00 local for descending) of the two relative orbits.

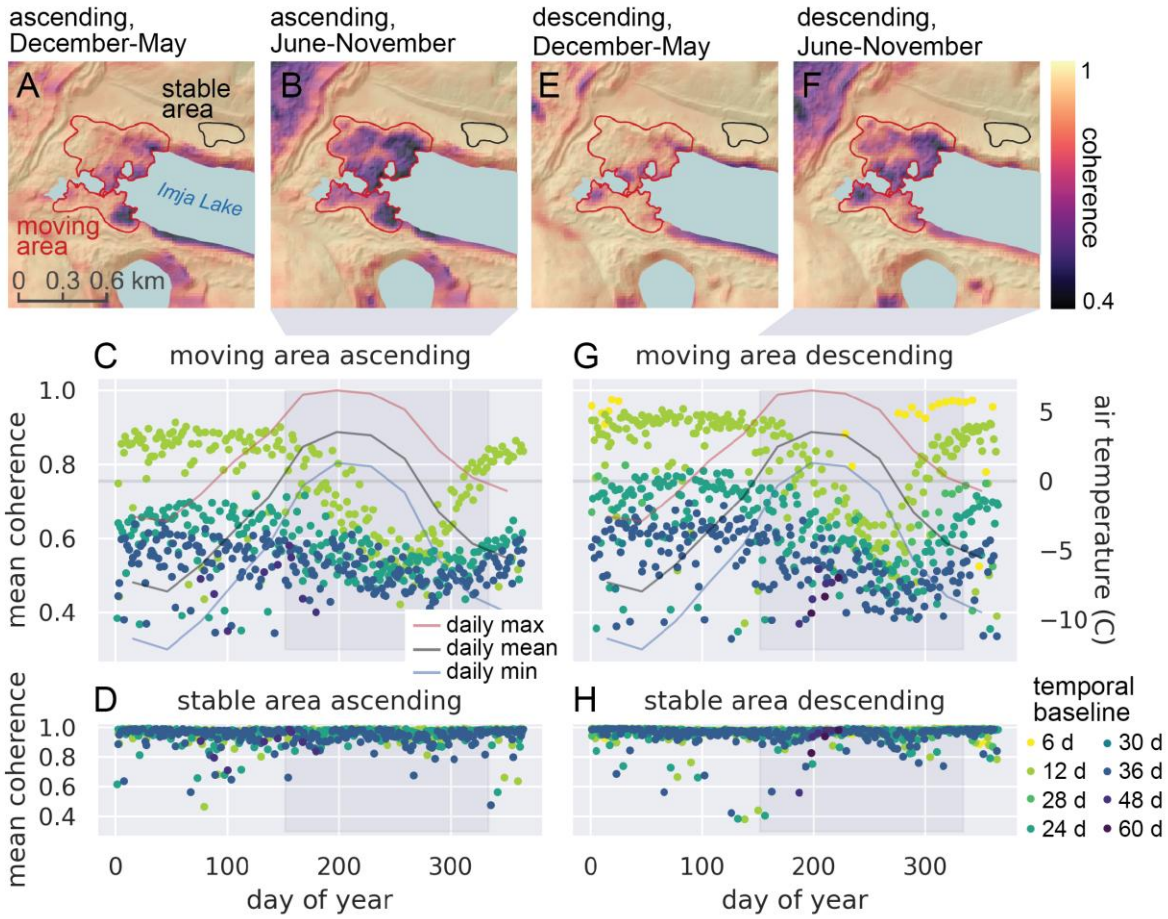

Figure 3.4: Imja Lake moraine dam InSAR coherence and air temperature. Mean coherence maps for all 439 12-day pairs of ascending Sentinel-1 SAR images during the colder months (December-May, A) and warmer months (June-November, B). C) Mean coherence of the moraine dam moving area in all ascending interferograms. Lines show mean daily maximum, mean, and minimum air temperature for each month from the nearby Pyramid Weather Station (UCAR/NCAR-Earth Observing Laboratory et al., 2011). The gray line corresponds to 0˚ C. D) Mean coherence of the stable area for all ascending interferograms. Panels E-H show corresponding plots for descending pairs.

Applying our technique to identify and map areas of buried ice, we observed a substantial decrease (0.34) in median coherence over the moraine dam moving area between the winter to early spring (0.92) and late summer (0.58) (Figure 3.5). There was no measurable decrease over the stable area.

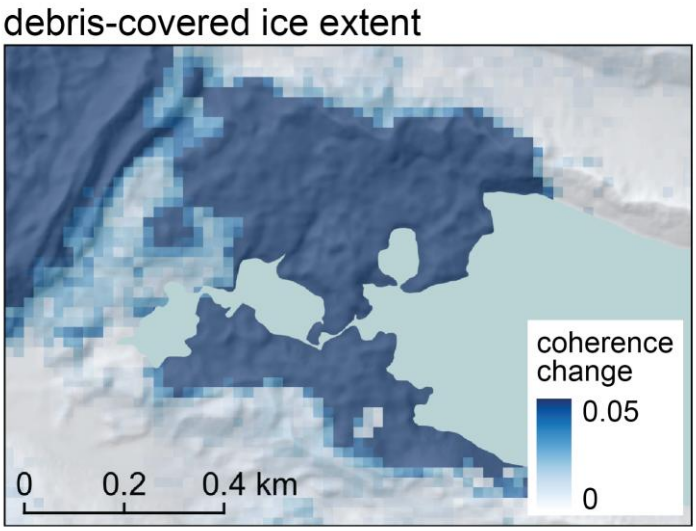

Figure 3.5: Seasonal change in median coherence offers a proxy to map the spatial distribution of buried ice.

3.6.2 Feature-tracking offset maps

The feature-tracking offset maps showed coherent motion of the moraine dam surface. The spatial mean of the per-pixel median LOS velocity over the moving area was +13.6 cm/yr for the ascending orbit and +7.8 cm/yr for the descending orbit, with a maximum per-pixel median LOS velocity of +46.1 cm/yr for the ascending orbit and +29.3 cm/yr for the descending orbit (Figure 3.6A, B). For both orbits, higher velocity was observed on the eastern side of the moving area near the lake shoreline. The mean SNR of the feature tracking products over the moraine dam moving area was 11.4 for the ascending orbit and 14.0 for the descending orbit (Figure 3.6C, D). The spatial mean of the per-pixel standard deviation of LOS velocity was 13.7 cm/yr for the ascending orbit and 25.9 cm/yr for the descending orbit (Figure 3.6E, F). The spatial mean of per-pixel median LOS velocity over the stable area was 0.7 cm/yr for the ascending track and -0.3 cm/yr for the descending track, though the feature tracking velocity values were ultimately masked in these areas with high InSAR coherence prior to inversion (Section 3.5.2). Positive LOS velocities are also observed over the Lhotse Glacier northwest of the moraine dam (Figure 3.6A, B).

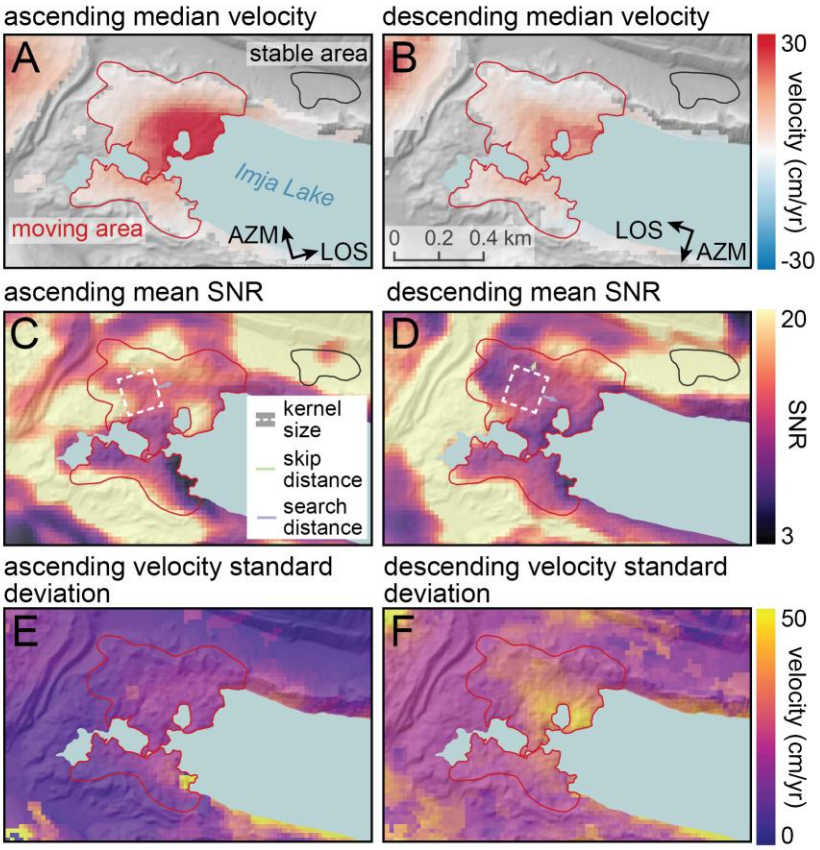

Figure 3.6: Sentinel-1 SAR image feature tracking results for the full 2017-2024 time series, including all 192 pairs on ascending (left) and 210 pairs on descending (right) orbits. A, B) Per-pixel median feature tracking velocity in the slant-range direction (along the line-of-sight). Results are masked in areas with high InSAR coherence (Figure 3.4), where feature tracking products are not used during inversion. C, D) Per-pixel median signal-to-noise (SNR) ratio. E, F) Per-pixel standard deviation of the feature tracking velocity in range direction. Higher standard deviation values are observed for locations with 1) greater temporal variability, 2) fewer reliable displacement measurements, and/or 3) larger measurement error.

3.6.3 Cumulative displacement time series from combined InSAR and feature tracking inversion

Most pixels in the moving area showed positive range change (i.e., movement away from the satellite) over the course of the study period for both ascending and descending orbits (Figure 3.7A, B). The mean total LOS displacement of the moraine dam moving area over the 7-year study period was 97.7 cm for the ascending orbit and 50.5 cm for the descending orbit, with a maximum LOS displacement of 338.6 cm for the ascending orbit and 207.4 cm for the descending orbit. The largest displacements were observed on the east side of the moving area near the lake shore, while

limited displacement was observed on the northern, southern, and western margins of the moving area. While many pixels displayed constant LOS velocities over time, others showed apparent seasonal fluctuations (Figure 3.7E, F). To assess error caused by atmospheric noise, we first computed the mean apparent displacement in the stable area at each time step, then computed the standard deviation over all time steps. These values were 0.2 cm/yr for the ascending orbit and 0.3 cm/yr for the descending orbit, indicating that our approach for atmospheric noise removal was effective.

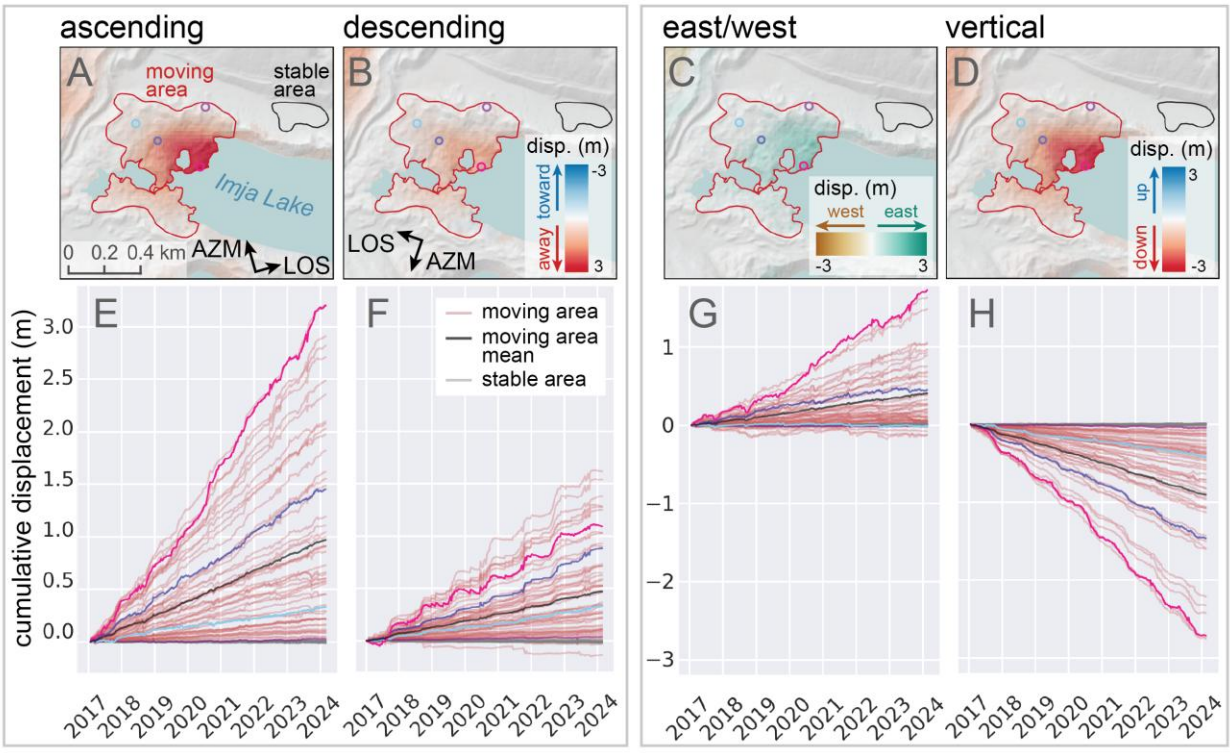

Figure 3.7: Cumulative displacement time series from the combined InSAR and feature tracking inversion. A–D) Maps showing total cumulative line-of-sight (LOS) displacement of the moraine dam and stable area over the study period for the ascending (A) and descending (B) orbits, and the decomposed east/west (C) and vertical (D) directions. Colored circles show pixel locations for the 1st, 25th, 75th, and 99th percentile of cumulative ascending displacement values in the moving area. E–H). Cumulative displacement time series for 50 randomly selected pixels in the moving area (red lines) and 50 randomly selected pixels in the stable area (grey lines). The black lines show the mean of all pixels in the moving area. The colored lines show the cumulative displacement time series for the pixels marked with corresponding colored circles in the maps.

After decomposing the LOS displacements into vertical and east/west displacements, we found that the observed LOS displacement of the moraine dam is consistent with downward and eastward motion (Figure 3.7, Figure 3.8). The mean total vertical displacement of the moraine dam moving area over the study period was -93.4 cm, with a maximum total vertical displacement of -337.0 cm. The mean total eastward displacement was 38.3 cm, with a maximum displacement of 174.1 cm. Over the stable area, the mean total vertical and horizontal displacements were 0.7 cm and 0.8 cm, respectively. While some pixels showed constant velocity over time, the pixels with the highest velocity displayed more seasonal variability, with increased velocity during the warmer months (Figure 3.7G, H).

The monthly median velocity maps show year-round movement of the moraine dam with a clear seasonal cycle. Vertical velocity was lowest (between -8.6 cm/yr and -10.1 cm/yr) between January and April (Figure 3.8A, B). The moraine dam moving area vertical velocity increased from May to October, when the mean vertical velocity reached -21.7 cm/yr (Figure 3.8A, B). The vertical velocity over the moraine dam moving area decreased in November and December (Figure 3.8A).From December-May, the mean horizontal velocity of the moraine dam moving area was 6.0 to 8.3 cm/yr to the east, with limited pixel-to-pixel variability (Figure 3.8A, C). Between June and September, some pixels began to move westward, and pixel-to-pixel variability increased (Figure 3.8A, C). Mean east/west velocity in the moving area decreased from June through September (Figure 3.8A, C).

From December to May, the mean horizontal velocity of the moraine dam moving area was 6.0 to 8.3 cm/yr to the east, with limited pixel-to-pixel variability (Figure 3.8A, C). Between June and September, some pixels began to move westward, and pixel-to-pixel variability increased (Figure 3.8A, C). The mean east/west velocity in the moving area decreased from June through September (Figure 3.8A, C).

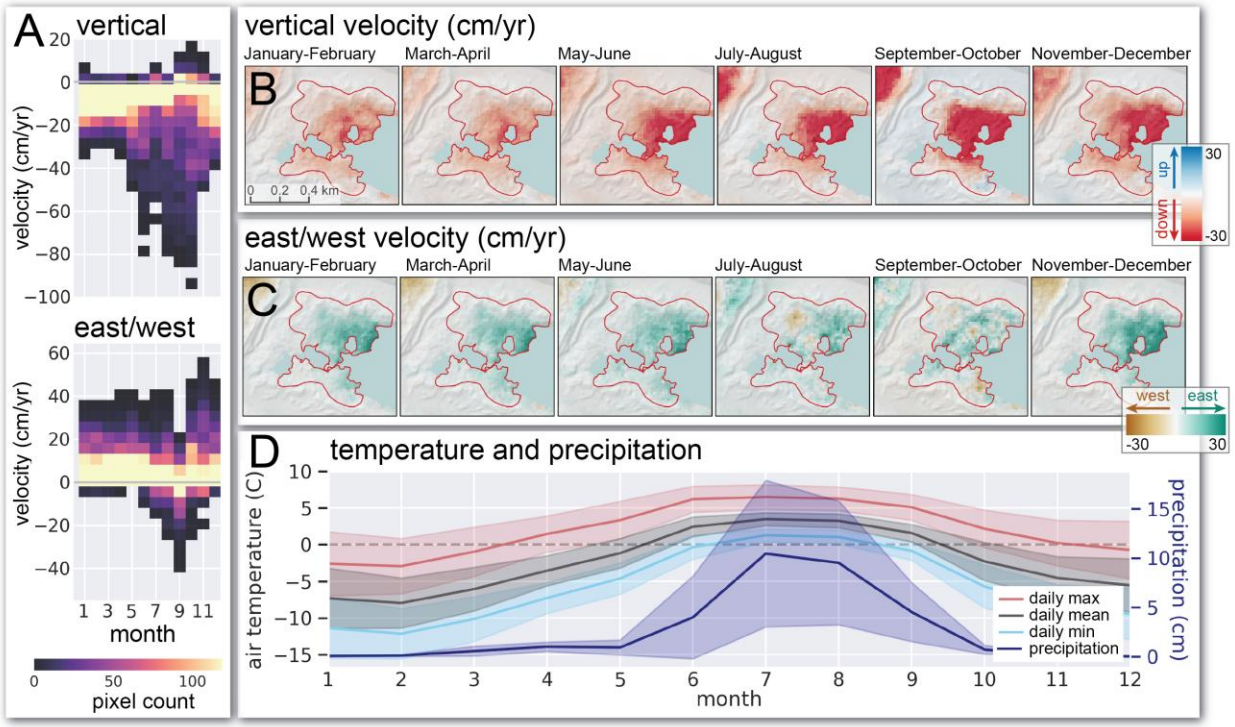

Figure 3.8: Aggregated monthly surface velocity of the Imja Lake moraine dam from the combined InSAR and feature tracking inversion for the 2017 to 2024 study period. A) Histograms showing distribution of median monthly surface velocity values in the moving area. B, C) Maps showing vertical and east/west median surface velocity during bimonthly periods. D) Mean daily air temperature and precipitation for each month at the nearby Pyramid Weather Station (UCAR/NCAR-Earth Observing Laboratory et al., 2011). Shaded areas show one standard deviation.

3.6.4 Validation

Our analysis shows that both the magnitude and direction of the combined InSAR and feature tracking velocity time series are consistent with the corresponding DEM-derived validation components (Figure 3.9; Table 3.2). The mean vertical velocity in the moving area from DEM differencing was -15.0 cm/yr. The mean vertical velocity from our combined InSAR and feature-tracking time series approach was -13.0 cm/yr. The InSAR-only SBAS underestimates both vertical and horizontal velocity (Figure 3.9). The feature-tracking-only SBAS captures velocity magnitude, but fails to capture short-term variability and contains large errors.

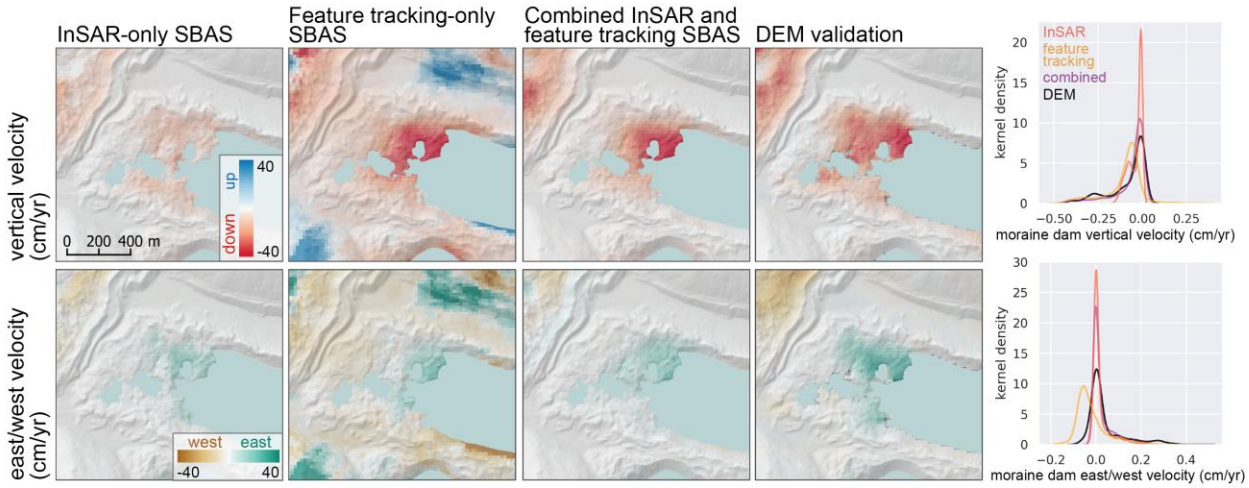

Figure 3.9: Mean vertical (top) and east/west (bottom) velocity products from InSAR-only SBAS (left), feature-tracking-only SBAS (left center), our combined InSAR and SAR feature tracking SBAS approach (right center), and the downsampled DEM-derived validation data (right), with corresponding kernel density plots for the moraine dam moving area. Our combined approach provides the best agreement with the validation data.

Table 3.2: Moraine dam velocity components (up/down and east/west) from InSAR-only SBAS, feature-tracking-only SBAS, our combined InSAR and feature tracking SBAS approach, and the DEM-derived validation data. The R^2 values were computed from per-pixel differences of each approach with the corresponding DEM-derived validation data.

	Mean (U/D)	Mean (E/W)	Std (U/D)	Std (E/W)	R^2 (U/D)	R^2 (E/W)
InSAR-only SBAS	-3.7	1.8	3.9	3.9	0.63	0.82
Feature-tracking-only SBAS	-10.2	2.0	10.6	6.9	0.67	0.35
Combined SBAS	-6.7	2.4	9.9	4.6	0.78	0.58
DEM-derived validation	-8.0	4.0	10.6	8.1	–	–

3.7 Discussion

Our surface displacement and coherence change results can be used to both infer the primary processes responsible for moraine dam degradation and to constrain the timing and magnitude of those processes.

3.7.1 Ice melt

We interpret the observed seasonal decrease in coherence to be caused by melting of buried ice and associated debris reworking. The lag between the onset of air temperatures above freezing and melt onset (Figure 3.4) is likely related to debris cover insulation. Both the melting of buried ice and debris reworking will change the moraine surface geometry, altering the scattering characteristics of the surface and decreasing coherence.

Where the magnitude of the seasonal change in coherence is large (Figure 3.5), we assume that buried ice is present. We find evidence for buried ice over most of the moraine dam, including areas displaying near-zero surface displacement. This is consistent with previous geophysical survey results. Hambrey et al. (2008) used both GPR and ERT to document buried glacier ice tens of meters thick (over 40 meters in places) overlain by up to 20 m of debris along the eastern portion of the moraine. Somos-Valenzuela et al. (2012) found consistent evidence for buried ice throughout 13 GPR transects across the moraine dam, ranging in bottom depth from 0–65 m. The deepest ice bottom depths (40–60 m) were documented on the northeast moraine dam, while shallower ice bottom depths (20–40 m) were documented around the southwest margin of the moraine dam, near the outlet. Dahal et al. (2018) conducted a dense ERT survey of the area surrounding the moraine dam outlet to the southwest, finding evidence for discontinuous blocks of glacier ice at depths varying from 0–20 m. Note that Hambrey et al. (2008) and Dahal et al. (2018) estimated debris thickness but were not able to locate the bottom of the debris-covered ice, and thus presented only a lower bound on ice thickness. Meanwhile Somos-Valenzuela et al. (2012) presented an ice-bottom depth, but did not estimate debris thickness, resulting in only an upper bound on ice thickness. We observe a smaller seasonal change in coherence over the southwest corner of the dam near the outlet (Figure 3.5), which is consistent with the isolated ice deposits mapped by Dahal et al. (2018). The magnitude of observed seasonal coherence change is likely also related to overlying debris thickness, as thicker debris 1) provides more insulation, suppressing melt, and 2) experiences less surface change per volume of ice melt at depth.

While other processes could also contribute to decreased coherence during warmer months, we did not observe loss of coherence over the stable area during the warmer months (Figure 3.4D, H). This suggests that seasonal melt of massive, buried ice is primarily responsible for the observed coherence loss over the moraine dam.

The observed spatial variability in seasonal coherence change may be explained by debris thickness and proximity to the lake and surface ponds, which can enhance downwasting and backwasting (Johnson, 1971; Driscoll, 1980; Watanabe et al., 1995). Low coherence during colder months was observed in some locations, especially bordering the lake and surface ponds (Figure 3.4A, C). These areas tend to coincide with areas experiencing rapid subsidence in our DEM-derived validation data (Figure 3.3). When air temperatures are below freezing, these locations potentially experience additional ice degradation and associated debris reworking due to melt caused by contact with water and by calving.

Overall, our measured vertical velocity values (Figure 3.7, Figure 3.8; Table 3.2) are lower than those documented by earlier studies. Watanabe et al. (1995) measured the position of five painted boulders in November 1989 and October 1994 and documented cumulative vertical displacements of -0.3 to -13.5 m with a mean of -4.8 m over the 5-year period (corresponding to vertical velocity range of -0.06 to -2.70 m/yr and mean of -1.0 m/yr). Fujita et al. (2009) performed topographic surveys using a theodolite in November 2001 and GPS receivers in April 2002 and October 2007. They created digital elevation models (DEMs) with 1 m posting for each survey and differenced to estimate cumulative vertical surface displacement. They report cumulative vertical displacement values for the "left bank" and "right bank" of -1.63 ± 1.71 m and -1.97 ± 1.67 m (mean ± 1-sigma), respectively, with a total range of vertical velocity values over the moraine dam of -0.06 to -1.03 m/yr for the 6-year period from 2001 to 2007. Haritashya et al. (2018) measured mean vertical velocity over the moraine dam between February 2000 and February 2016 of less than -0.5 m/yr, with a maximum of -1.9 m/yr, and uncertainty of 0.7 m/yr. The disagreement with our measured vertical velocity values likely results from some combination of inaccurate and/or incomplete sampling of the moraine dam surface in previous work, and/or a real decrease in the moraine dam vertical velocity over time.

During the warmer months, we observed a clear increase in downward vertical velocity over much of the moraine dam (Figure 3.8A, B), likely associated with melting of buried ice. The highest downward velocity over the moving area was observed in October with mean velocity of -21 cm/yr. Seasonal increases in velocity were first observed over the northern margin of the lake and surface ponds, suggesting earlier seasonal melt in these locations. Small areas of apparent upward displacement at the southern margin of the lake (Figure 3.8B) are likely caused by consistent unwrapping errors due to low coherence (Figure 3.4B, D). East-west displacement

during the warmer months was variable, with some areas mostly moving east and some areas mostly moving west (Figure 3.8C). This variability in horizontal displacement direction could be caused by the variable topography of the moraine dam, as the apparent direction of displacement caused by melt of sloped ice surfaces depends on the surface aspect.

If we extrapolate the observed mean subsidence rates for the 2017 to 2024 period forward in time and consider the current moraine dam topography (Figure 3.1, Figure B.1), we find that the northeast area of the moraine dam bordering the lake edge is most vulnerable to subsidence below the current lake level. This topographic evolution would result in westward expansion of Imja Lake, decreasing the width of the moraine dam.

3.7.2 Ice flow

During the colder months, we observed widespread downward and eastward displacement of the moraine dam surface, with a mean velocity within the moving area of -10.5 and 7.3 cm/yr, respectively (Figure 3.8). The high coherence values over the moraine dam during colder months (Figure 3.4) provide high confidence in these measurements. Assuming that air temperatures during the colder months were mostly below freezing (Figure 3.4C, G), we do not attribute this displacement to melt of buried ice. While debris settling and frost creep could potentially explain some component of downward and horizontal motion, they are not expected to cause displacement of this magnitude on these time scales (Colman and Pierce, 1986; Bursik, 1991; Hallet and Putkonen, 1994) and we do not observe similar motion over debris outside the moraine dam. We instead attribute this motion primarily to flow of the thick, massive ice within the moraine dam.

Downward and eastward movement is consistent with along-slope flow for the broad moraine dam surface slopes (Figure B.1). The mean surface slope of the moving area is 9.5˚. Somos-Valenzuela et al. (2012) found evidence for ice bottom depths between 30 and 60 m over much of the moraine dam. Assuming debris thickness of less than 1–20 m (Hambrey et al., 2008), buried ice in the moraine dam should have thickness of 10 to 60 m. We generally observe higher east/west and vertical surface velocity during colder months over areas with deeper ice bottom depths documented by earlier studies. Near-zero velocity was observed over the southwest portion of the moraine dam, where Dahal et al. (2018) observed discontinuous blocks of buried ice. The observed mean winter velocity magnitude (combining the vertical and east/west components) of 12.8 cm/yr over the moving area of the moraine dam is consistent with expected surface velocity values for a simple ice flow model and the range of measured ice thickness (see Section B.1, Figure B.9).

Previous surface displacement measurements also provide evidence for flow of buried ice within the moraine dam. Watanabe et al. (1995) described westward movement of several boulders along the outlet channel of -0.96 m/yr. Fujita et al. (2009) described horizontal flow rates of less than 0.17 m/yr in an unspecified direction. These results potentially suggest that at some point before our study period, the flow direction of the moraine dam evolved from original westward, downvalley flow (the original flow direction of Imja Glacier) to the eastward, upvalley flow (toward Imja Lake) observed during our 2017–2024 study period.

The observed change in flow direction is likely related to evolving thickness of lower Imja Glacier through differential ablation both above and below the lake surface (Kjær and Krüger, 2001), changing longitudinal stresses associated with decoupling from the active flowing glacier during retreat, and interaction with the expanding lake. We observed a concave-upward transverse profile (bottom-right panel of Figure B.10) for the Imja Lake moraine dam during our study period. A similar configuration, where the glacier surface slopes downward from the lateral and terminal moraines toward the centerline, is observed for many debris-covered glaciers in the Everest region (e.g., Nuptse Glacier, Figure B.11). This scenario introduces driving stress toward the glacier centerline (downslope), which is consistent with the observed direction of motion in our horizontal surface velocity maps. Our detailed displacement time series can potentially be used to constrain ice and debris evolution models, which are needed to better understand the evolution of the Imja Lake moraine dam and implications for changing structural stability.

3.7.3 Limitations

While our InSAR and feature tracking results improve our understanding of the processes responsible for moraine dam degradation, it is important to discuss some limitations. Underestimation of rapid, localized surface displacement can partly explain why our maximum vertical velocity over the moraine dam surface was only -0.47 m/yr compared to -1.9 m/yr from Haritashya et al. (2018). We also see evidence of localized backwasting in our DEM-derived validation dataset (Figure 3.3) that is not captured in our InSAR and feature-tracking displacement time series (Figure 3.7, Figure 3.8).

We also acknowledge limitations of the Sentinel-1 acquisition geometry, with only ~10% of any real north/south displacement mapped into LOS measurements for this location. We combined data from ascending and descending passes to calculate vertical and east/west displacement components, assuming that north/south displacement does not substantially contribute to the

displacement observations. In fact, some small contribution from north/south displacement is expected, which potentially introduces a source of error in our measurements of surface velocity (Brouwer and Hanssen, 2021). Based on the north/south displacement from the DEM-derived validation dataset, we expect that our LOS decomposition could introduce a bias of less than ~4% in mean vertical velocity over the moraine dam moving area and less than ~3% in mean east/west velocity over the moraine dam moving area (Section B.2, Figure B.12, Figure B.13).

3.8 Conclusions

We demonstrated that Sentinel-1 InSAR and SAR feature tracking can be used to measure the displacement of an ice-cored moraine dam bounding Imja Lake in Nepal. We validated our measurements using a high-resolution 3D displacement dataset derived from stereo DEMs. By combining InSAR and feature tracking measurements, we created a surface displacement time series that is more accurate than those created using InSAR or feature tracking alone. We found that a 0.3 km^2 area of the moraine dam subsided an average of ~90 cm over the 7-year study period, and the northeastern area of the moraine dam bordering the lake edge is most vulnerable to subsidence below the current lake level.

We observed a systematic seasonal cycle in InSAR coherence values over the moraine dam surface, with a substantial decrease in coherence occurring about 60 days after mean daily maximum air temperature increased above 0°C. We showed that seasonal changes in InSAR coherence can be used to map buried ice within the moraine dam, with an observed spatial distribution that is consistent with previous geophysical surveys.

Our results suggest that the primary processes contributing to moraine dam degradation are buried ice melt and ice flow. Buried ice melt occurs in the warmer months (summer and fall), causing a seasonal increase in downward vertical velocity and variable horizontal velocity over the surface of the moraine dam. Ice flow occurs year-round and results mostly in eastward movement. Our surface velocity measurements, alongside historical measurements, provide evidence that the Imja Lake moraine dam is undergoing kinematic changes, which we attribute to changing ice thickness, past decoupling from the Imja Glacier, and interaction with Imja Lake.

Collectively, our findings demonstrate the potential for satellite SAR remote sensing to quantify the evolution of moraine dam degradation and to provide new information about moraine dam internal structure. This information may be valuable when assessing GLOF hazards, performing in situ investigations, and planning hazard remediation activities. We are extending

this analysis to other high-priority moraine dams in the region and hope to integrate additional remote sensing measurements of surface change (e.g., NISAR, optical feature tracking) in the future.

Chapter 4

MAPPING BURIED ICE, QUANTIFYING SURFACE MOTION, AND IMPROVING GLOF HAZARD ASSESSMENTS FOR GLACIAL LAKE MORAINE DAMS IN NEPAL WITH SENTINEL-1 INSAR

This chapter is in preparation as: **Brencher, G.**, Henderson, S. T., & Shean, D. E. ***Mapping buried ice, quantifying surface motion, and improving GLOF hazard assessments for glacial lake moraine dams in Nepal with Sentinel-1 InSAR***

4.1 Abstract

Moraine-dammed glacial lakes can produce destructive glacial lake outburst floods (GLOFs). Buried ice within moraine dams can impact both GLOF risk and the characteristics of potential GLOFs. However, hazard assessments and GLOF modelling studies frequently lack information about the spatiotemporal distribution of buried ice in moraine dams, potentially decreasing their accuracy. In this work, we use Sentinel-1 InSAR to quantify surface movement and map buried ice for moraines damming 23 high-priority glacial lakes in Nepal. We find that most moraine dams contain buried ice, and that buried ice typically results in surface movement of centimeters to tens of centimeters per year. We find that while the spatial distribution of buried ice is related to moraine dam morphology, these characteristics alone are not sufficient to identify buried ice presence. We infer the presence of extensive buried ice in eight moraine dams previously thought to be ice-free, substantially changing the results of previous hazard assessments. By analyzing the vertical and east/west components of surface movement, we find evidence that the observed surface deformation is likely caused by both ice melt and ice flow, which vary in magnitude between moraine dams. The fastest-moving moraine dams are Imja, Barun, and Thugali, suggesting that these locations have relatively thick (tens of m) buried ice layers and relatively thin debris cover. Overall, our results can be used to improve the accuracy and efficiency of hazard assessments and GLOF modelling efforts, and to better understand moraine dam evolution and the role of buried ice in GLOF events.

4.2 Introduction

Glacial lake outburst floods (GLOFs) are a threat to millions of people and to critical infrastructure, including hydropower facilities that dominate economic production in some high-mountain regions (Tshering and Tamang, 2004; Carrivick & Tweed, 2016; Taylor et al., 2023). For example, a July 8, 2025 GLOF damaged four hydropower plants in Nepal and is estimated to have rendered 8% of Nepal's power production capacity inoperative (Nepali Times, 2025). Some of the most destructive GLOFs in recent years have been produced by moraine-dammed glacial lakes (e.g., South Lhonak Lake, Sattar et al., 2025). For these lakes, both GLOF likelihood and characteristics are controlled in part by moraine dam morphology and structure (Huggel et al., 2004; McKillop and Clague, 2007; Westoby et al., 2014). One important control is the 3D distribution of buried ice within the moraine dam.

Buried ice is frequently preserved in Little Ice Age moraine dams long after the glacier has retreated (Richardson & Reynolds 2000). Buried ice can influence GLOF hazard in several ways. As it degrades over time, it can weaken the moraine dam and create conduits for lake water, potentially leading to piping and "self-destructive" dam failure (Richardson & Reynolds 2000; Clague and Evans, 2000). Loss of buried ice can change dam height, width, and slope, factors which control the likelihood of moraine dam overtopping and breaching (McKillop & Clague, 2007; Fujita et al., 2013). The presence of buried ice can also affect how the dam breaching process unfolds (Worni et al., 2012; Westoby et al., 2014).

Despite the complex role of buried ice in shaping GLOF hazard, some hazard assessments fail to account for buried ice within moraine dams, and many GLOF hazard assessments treat moraine dams as either "ice-cored" or not (e.g. McKillop & Clague, 2007; Wang et al., 2012; Rounce et al., 2016; Dubey and Goyal, 2020, Zhang et al., 2022). Remote sensing approaches designed to infer the presence of an ice core are largely qualitative, and have not been rigorously validated (Rounce et al., 2016). For example, McKillop and Clague (2007) and Wang et al. (2012) classified rounded moraines with minor superimposed ridges as ice-cored, and narrow, sharp-crested moraines with angular cross-sections as ice-free. Other studies have used surface features including ice cliffs, karst topography, and outlet channels that change over time to infer the presence of buried ice (Rounce et al., 2016). Misclassification of moraine dams based on these approaches can have substantial impacts on the conclusions of hazard assessments.

In reality, the distribution of ice within moraine dams is variable over both space and time, and cannot be captured by a simple binary framework. Failure to account for these details can decrease the accuracy of hazard assessments. The continued use of binary "ice-cored" or "not ice-cored" classification approach likely reflects both limited data on moraine dam composition and the difficulty of incorporating more complex ice distribution data into existing hazard assessment frameworks. While advances in remote sensing capabilities have improved topographic characterization of moraine dams and surrounding terrain (e.g. Haritashya et al., 2018), mapping of buried ice still typically relies on in situ geophysical methods. These field-based approaches cannot be applied to thousands of potentially hazardous glacial lakes, and the high heterogeneity within moraine dams means that even field-based mapping efforts are prone to undersampling moraine dam composition (Westoby et al., 2014).

In addition to creating a challenge for hazard assessments, buried ice also complicates numerical modeling of GLOFs (Westoby et al. 2014). The material properties of ice and glacial till are markedly different, and the presence, extent, and characteristics of buried ice can strongly influence breach development (Worni et al., 2012). These properties can alter breach initiation and growth, shaping the GLOF outflow hydrograph and controlling peak discharge, which are important for estimating downstream impacts. Uncertainty in modeling the breach process contributes to errors in predicted GLOF impacts. Unfortunately, buried ice is typically not accounted for in numerical models of moraine dam breaching, as 1) most numerical models used to simulate dam breaching are not designed to represent ice (Westoby et al. 2014), and 2) the spatial distribution of ice within moraine dams is often unknown. Numerical modelling of moraine dam failure has generally been performed for moraine dams assumed to have no buried ice (e.g. Somos-Valenzuela et al., 2016; Mergili et al., 2018), or in locations where ice is thought to be buried deep enough to not be involved in realistic moraine dams breach scenarios (e.g. Lala et al., 2017).

Remote sensing can serve as a key source of information about buried ice in moraine dams, which can be used to improve both GLOF hazard assessments and GLOF modelling efforts. Satellite interferometric synthetic aperture radar (InSAR) can be used to measure subtle changes in the moraine dam surface caused by melt and deformation of buried ice, quantifying moraine dam degradation and mapping the distribution of buried ice. Previous efforts have used digital elevation model (DEM) differencing to quantify surface change associated with buried ice (e.g.

Fujita et al., 2008; Haritashya et al., 2018). Sentinel-1 InSAR offers near-global coverage with 12-day temporal resolution, penetration through clouds, sensitivity to sub-centimeter surface displacement, and a decade-long data archive. In this work, we adapt the Sentinel-1 InSAR methods developed by Brencher et al. (2025) to study the moraine dams at 23 glacial lakes in Nepal identified as "potentially dangerous" and a high priority for monitoring by Mool et al. (2001) and ICIMOD (2011).

4.3 Methods

We generally followed the methods of Brencher et al. (2025) to quantify surface displacement and map internal ice in moraine dams, with a few modifications. We employ an interferogram stacking approach (Zebker et al., 1997) to capture peak seasonal velocity in the late summer, rather than the time-series approach of Brencher et al. (2025) to examine seasonal change in moraine dam velocity. This approach mitigates the impact of data gaps caused by seasonal snow and long periods between Sentinel-1 acquisitions at the cost of detailed information on the temporal evolution of the moraine dams.

4.3.1 Data

We downloaded all ascending and descending Sentinel-1 burst single-look complex (SLC) data between January 1, 2017 and May 15, 2024 over each lake (Table C.1). Data collected prior to the Sentinel-1B commissioning in 2017 were not used due to long gaps between acquisitions. Where multiple ascending and descending bursts covered a lake, we chose the burst with a smaller local incidence angle to maximize the vertical component of the line-of-sight (LOS) displacement vector, improving sensitivity to moraine dam subsidence. Data were acquired in interferometric wide (IW) swath mode with vertical co-polarization (VV) and with a resolution of 14.1 m in azimuth and 2.3 m in range. For each burst SLC, we also used the corresponding layover and shadow masks distributed with the Observational Products for End-Users from Remote Sensing Analysis (OPERA) Radiometric Terrain Corrected (RTC) Sentinel-1 products (OPERA, 2023). We used the 2021 release of the 30-m Copernicus "GLO-30" (European Space Agency, 2021) digital elevation model (DEM) for topographic analysis and InSAR processing.

4.3.2 Interferometric processing

Interferograms and interferometric coherence maps were generated using the HyP3 ISCE2 Plugin (version 0.9.3; Hogenson et al., 2020), which enables batch processing with the Jet Propulsion Laboratory InSAR Scientific Computing Environment (ISCE) version 2.6.3 software (Rosen et al., 2012). We processed interferograms using a given burst SLC and the three subsequently acquired burst SLCs, typically resulting in temporal baselines of 12, 24, and 36 days (with occasional 6-, 18-, and 30-day temporal baselines) using the "insar_tops_burst" workflow. Interferograms were processed with five looks in range and one look in azimuth and geocoded to 20-m square pixels.

4.3.3 InSAR coherence change and buried ice mapping

Following Brencher et al. (2025), we use seasonal change in InSAR coherence to infer the location of buried ice within moraine dams. We assume that warm-season melt of buried ice results in surface changes that reduce InSAR coherence (Brencher et al., 2025). To delineate the extent of buried ice, we used interferograms with short 12-day temporal baselines, which show the largest seasonal coherence change over debris-covered ice (Brencher et al., 2025). We then examined the median coherence of all interferograms (with variable temporal baselines of 6-36 days) over the feature moving area (Section 4.3.5) to define high-coherence (cold, snow-free) and low-coherence (warm, snow-free) time periods. The mean duration of the low-coherence period spanned 60 days, from DOY 228-288 (August 16-October 14), while the high-coherence period spanned ~DOY 0-100 for 7 sites with minimal snow or ~DOY 120-180 for 16 sites with significant cold-season snow (Table C.2). Next, we calculated the seasonal coherence change, which is the difference between the median coherence in the high-coherence period and the low-coherence period. We used median rather than mean coherence to mitigate the impact of anomalous events (i.e., snow accumulation during storms) that temporarily decrease coherence, but do not reflect surface change related to seasonal degradation of buried ice.

To identify buried-ice pixels, we used Gaussian mixture modeling (GMM) to estimate a threshold in seasonal coherence change. Gaussian mixture models represent an observed distribution as the weighted sum of two or more Gaussian components, and are frequently used in remote sensing to estimate the probability that a pixel belongs to a given class for change detection (e.g., Bruzzone & Prieto, 2002; Çelik, 2011) and classification (e.g. Ju et al., 2003; Berge & Solberg, 2006; Wan et al., 2019) tasks. For each moraine dam, we modeled the distribution of

seasonal-coherence-change values to determine whether a subpopulation of pixels with reduced warm-season coherence was present. We expect that pixels without buried ice will follow an approximately Gaussian distribution centered near zero, while pixels containing buried ice will form a distinct group of positive values corresponding to a warm-season decrease in coherence.

To assess whether this secondary distribution was appropriate for a given moraine dam, we fit both a single-Gaussian model and a two-Gaussian model using maximum-likelihood estimation with the GMM implementation in scikit-learn (Pedregosa et al., 2011). We compared the two models using the Bayesian Information Criterion (BIC; Schwarz, 1978), a measure which balances goodness of fit and model complexity and is commonly applied to select the appropriate number of Gaussian components in GMM (e.g., Steele & Raftery, 2010; McLachlan & Rathnayake 2014; Wan et al., 2019). Lower BIC values indicate a better-supported model. For moraine dams where the two-Gaussian model yielded a lower BIC than the single-Gaussian model, we defined the buried-ice threshold as the smallest seasonal coherence-change value where the probability of belonging to the buried-ice subpopulation exceeded that of the non-buried ice subpopulation. We then calculated the median threshold across all moraine dams and applied this median threshold to classify buried-ice pixels. We applied a single threshold across all moraine dams, rather than a unique threshold at each moraine dam, to allow for the identification of smaller areas of buried ice (<5 pixels) which may be insufficient for the two-Gaussian model. To evaluate our threshold, we performed a sensitivity analysis by systematically varying the threshold and evaluating the effect on the total buried ice area within each moraine dam.

4.3.4 Surface velocity

To minimize effects of snow and capture maximum activity associated with melt of buried ice, we characterized surface velocity during the low-snow warm period of the year. Warm-season ascending and descending LOS velocity were calculated by first selecting interferograms with temporal baseline <= 12 days to mitigate temporal decorrelation, then masking areas of layover and radar shadow. We then included only interferograms formed from acquisitions between day-of-year 200 and 300 (July 19 - October 27), and subtracted the mean displacement in the reference area (Section 4.3.5) to remove local atmospheric noise. We next filtered each interferogram to include only pixels with coherence greater than 0.6, which mitigates the effect of potential phase unwrapping errors over the moraine dam. We then calculated LOS velocity for these pixels in each interferogram by dividing the total displacement by the temporal baseline. We then computed the

per-pixel median summer LOS velocity for any pixels with 5 or more valid LOS velocity measurements. Using the azimuth and incidence angles for each burst SLC, we then combined the ascending and descending median LOS velocity maps to solve for the east/west and vertical velocity components, following the approach described and validated in Brencher et al. (2025). We evaluated agreement between these velocity components and the Brencher et al. (2025) DEM velocity products for the Imja lake moraine dam using the coefficient of determination (R^2).

We summarize the deformation observed in the median summer velocity maps using spatial statistics. For the LOS and east/west components, we report the spatial mean of the absolute velocity values over each moraine dam to characterize overall movement magnitude and to avoid cancellation between positive and negative values that arise from viewing geometry and complex deformation patterns. For the vertical component, we report the spatial mean of the signed velocity values because subsidence is the expected signal and the sign carries physical meaning. To characterize the fastest-moving parts of each moraine dam, we compute spatial percentile metrics from the same median summer velocity products. Specifically, we report the 95th percentile of the absolute LOS velocities and the 5th percentile of the vertical velocities (fastest subsiding pixels). These percentile-based metrics reduce sensitivity to potential outliers. After computing these statistics separately for each moraine dam, we take the mean across all sites with moving areas to represent the typical motion magnitude and the typical fastest rates of deformation observed throughout the region.

4.3.5 Delineation of moraine dam, reference, stable and moving areas

To delineate moraine dam features, we mapped all non-bedrock material separating each glacial lake from the downstream valley using an ensemble of satellite basemaps (Microsoft Bing; Google Satellite; ESRI World Imagery; all accessed May-June 2024) and the GLO-30 DEM (Figure C.1). In other words, we delineated material that, if removed, would result in the impounded lake entering a downstream area. This approach was intended to reduce the subjectivity of dam delineation in the absence of detailed subsurface information and to identify the areas where moraine dam subsidence would contribute most significantly to GLOF risk, rather than to include all material that plays some role in impounding the lake.

Reference areas were used to set a local reference point for interferograms to mitigate the effects of atmospheric noise and potential unwrapping errors, which may introduce an apparent phase shift over the entire area of interest. Stable areas were used to assess potential error of the

surface displacement measurements. Reference and stable areas were selected over relatively flat surfaces with high year-round coherence and near-zero LOS velocity, as close as possible to the moraine dam (Figure C.2). The distance between the moraine dam centroids and the reference and stable area centroids ranged from 193 to 2460 m, with a mean distance of 811 m. Moving areas of the moraine dam were defined where coherent LOS surface velocity was greater than or equal to 1 cm/yr.

4.3.6 Uncertainty

To assess uncertainty in buried ice extent, we calculated the apparent area of pixels classified as buried ice (Section 4.3.3) in the stable areas, which should be ice-free. To assess uncertainty in surface velocity, we calculated the mean and standard deviation of apparent ascending, descending, vertical, and east/west velocity in the stable areas. Assuming no motion in the stable area, the mean apparent velocity serves as a measure of bias, while the standard deviation of apparent velocity serves as a measure of dispersion.

4.4 Results

4.4.1 Moraine dam characteristics

The mean moraine dam area for our 23 sites was 0.19 ± 0.15 km^2 with a mean elevation of 4941 ± 498 m (± corresponds to 1 standard deviation; Table C.3). Of the 23 dams, 12 were affected by layover or radar shadow for at least one Sentinel-1 orbit. Of dams affected, layover or radar shadow covered on average 10 ± 14% of the moraine dam area. Ten of the moraine dams are best classified as narrow and sharp-crested, and 13 moraine dams are best classified as wide and rounded (Table C.4).

We observed coherent moving areas for 19 of the 23 moraine dams (Figure 4.1, Table C.3). The moving area size ranged from 0.007 to 0.29 km^2, with a mean of 0.085 km^2, corresponding to 43 ± 25% of the total moraine dam area (Table C.3).

4.4.2 Buried ice mapping

Applying GMM to model the distribution of seasonal coherence change values (Figure C.3) yielded thresholds for buried ice identification ranging from 0.01 (Barun) to 0.15 (gakal_gl_0008). The median seasonal coherence change threshold, which we applied to identify buried ice pixels at all moraine dams, was 0.04.

Of the 23 moraine dams, 18 contained more than three pixels classified as buried ice (Figure 4.1). Across these 18 dams, buried ice area ranged from 0.003 to 0.3 km², with a mean of 0.068 km². Between 3% and 62% of the moraine dam area was classified as buried ice, with an average of 28%; 87% of moraine dam buried ice area was located within a moving area (Figure 4.1).

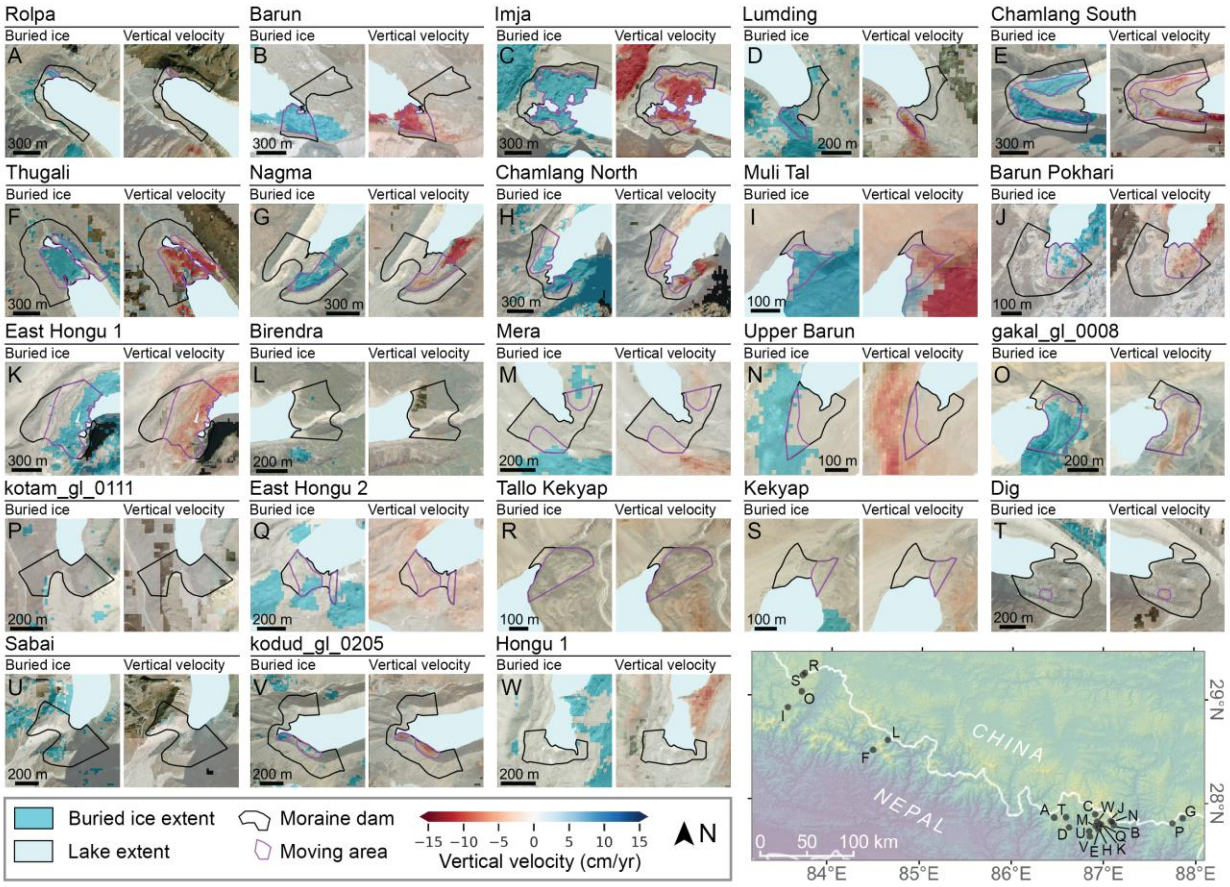

Figure 4.1: Buried ice extent and median vertical velocity for 23 high-priority moraine dams in Nepal during the period from 2017-2024. Buried ice extent inferred from seasonal InSAR coherence change. Most moraine dams contain subsiding areas and areas of buried ice. Panel order corresponds to the ranking of priority for investigation from ICIMOD (2011). Basemap is Microsoft satellite basemap. Microsoft product screen shots reprinted with permission from Microsoft Corporation. See Figure C.5 for corresponding maps of East/West velocity.

Total buried-ice area and the proportion of moraine-dam area classified as buried ice were more sensitive to decreases in the seasonal coherence-change threshold than to increases (Figure C.4). Across thresholds from 0.02 to 0.06, total buried-ice area varied by ~45%, and the mean proportion of dam area classified as buried ice varied by ~55%, indicating relatively high sensitivity to

threshold choice. The number of moraine dams containing buried-ice pixels ranged from 21 to 18 across this range, indicating that buried ice presence is moderately sensitive to threshold choice.

4.4.3 Surface velocity

Across the three highest-priority moraine dams (Rolpa, Barun, and Imja; Mool et al., 2001; ICIMOD, 2011), the spatial pattern and magnitude of surface velocity vary widely. At the Rolpa moraine dam, surface movement is restricted to a small patch on the northwest flank (0.01 km², 6 percent of the moraine dam area). This area is affected by layover in the ascending orbit, but the descending orbit shows a mean absolute LOS velocity of 5.1 cm/yr (Figure 4.1, Figure C.6A). At the Barun moraine dam, deformation is concentrated in the southern portion of the moraine dam, which has a mean vertical velocity of −13.9 cm/yr, while the northern portion remains largely stable (Figure 4.1). At the Imja moraine dam, more than half of the moraine dam surface is deforming (58 percent; Table C.3), with a mean subsidence rate of −13.2 cm/yr. These examples illustrate the diversity of deformation patterns present in moraine dams across the region.

Across the 19 moraine dams, the mean absolute LOS velocity (Section 4.3.4) of the moving areas was 4.7 ± 0.5 cm/yr for the ascending orbit and 4.8 ± 0.5 cm/yr for the descending orbit (± 1-sigma uncertainty computed over stable area; Section 4.3.6, Figure C.6). The mean of 95th percentile absolute LOS velocity (Section 4.3.4), representing the fastest-moving pixels for each site, was 11.4 cm/yr for the ascending orbit and 10.8 cm/yr for the descending orbit. The mean vertical velocity for all of the moving areas was -5.1 ± 0.4 cm/yr, while mean 5th percentile vertical velocity (Section 4.3.4), representing the fastest-subsiding pixels for each site, was -11.1 cm/yr (Figure 4.2, Table C.5). The Barun, Imja, and Thugali moraine dams had the largest magnitude 5th percentile vertical velocity values of -29.6, -28.7, and -28.9 cm/yr, respectively (Table C.5). The mean absolute value east/west velocity across all moving areas was 4.1 ± 0.6 cm/yr, similar in magnitude to the vertical velocities (Figure C.5, Figure C.7).

In the stable areas, the mean LOS velocity was 0.4 ± 0.5 cm/yr for the ascending orbit and 0.4 ± 0.5 cm/yr for the descending orbit. The mean vertical velocity was -0.5 ± 0.4 cm/yr, while mean east/west velocity was -0.1 ± 0.6 cm/yr.

We observe good agreement between our InSAR velocity products and the 3D velocity validation dataset from Brencher et al. (2025) for the Imja Lake moraine dam, with an R^2 of 0.73 for the vertical velocity and 0.65 for the East/West velocity (Figure C.8).

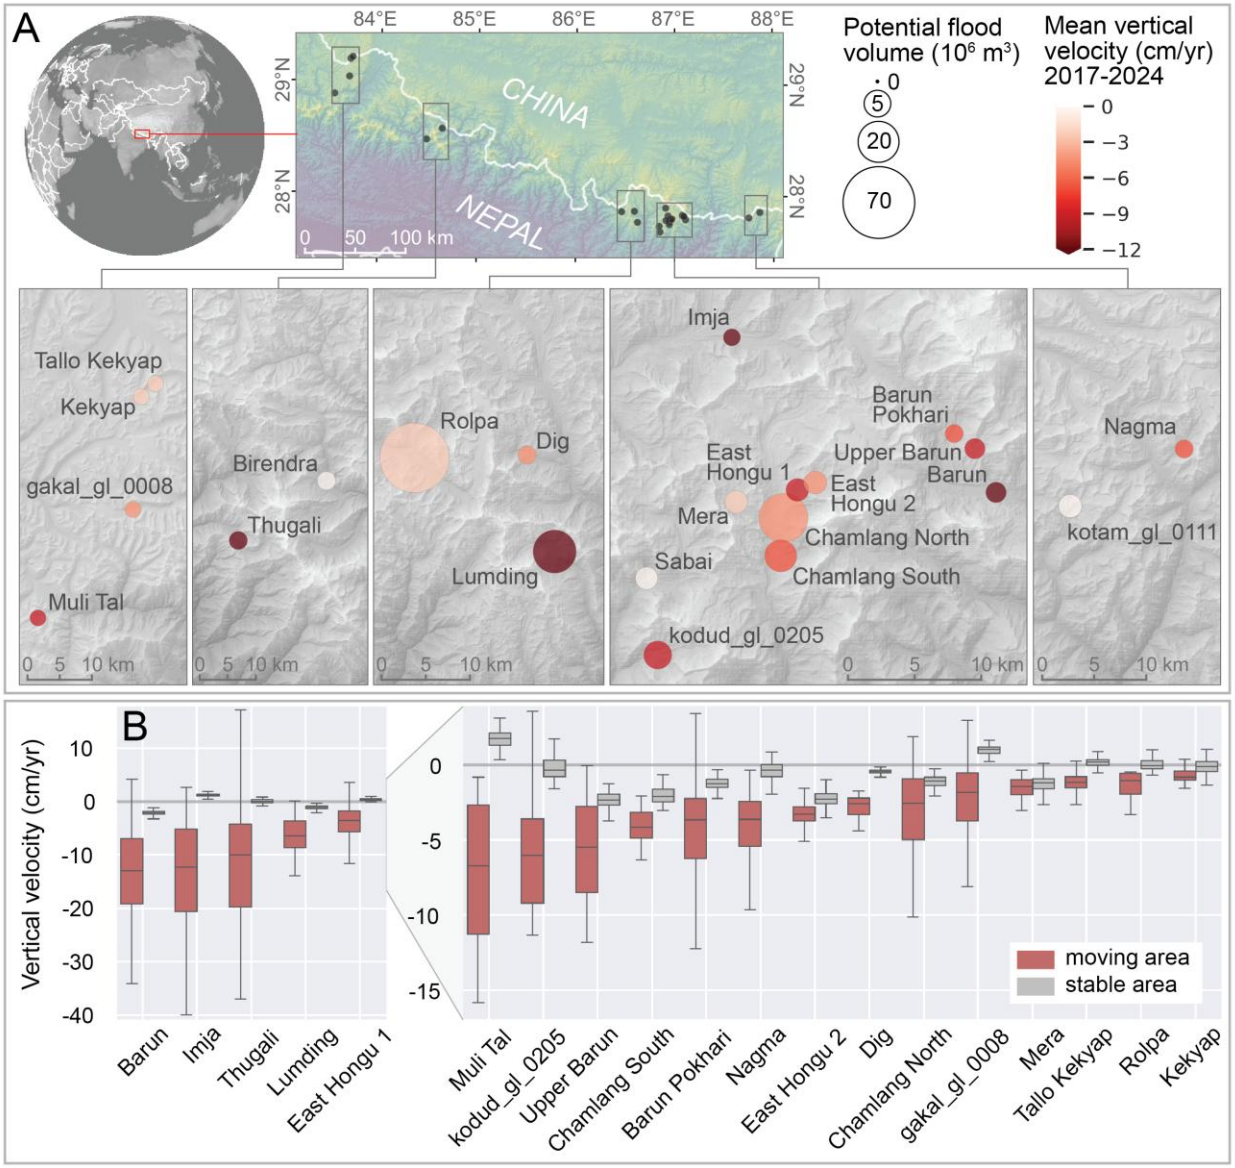

Figure 4.2: Vertical subsidence of moving areas for 23 high-priority glacial lake moraine dams in Nepal for the 7.37-year period between January 1, 2017 and May 15, 2024. A) Local maps of hazardous lake clusters showing mean vertical velocity of the moving area (color) and potential flood volume (marker area, Rounce et al. (2017)). B) Vertical velocity distributions for the moving areas and corresponding stable areas of each moraine dam. Sites are sorted on the x-axis based on median vertical velocity, with the 5 fastest-subsiding moraine dams on separate axes in left panels. Figure C.5 shows East/West velocity distributions.

4.5 Discussion

4.5.1 Buried ice in moraine dams

Our results show that seasonal change in InSAR coherence can be used to detect the presence and map the extent of buried ice in moraine dams and surrounding areas. We documented evidence for buried ice in 18 of the 23 high-priority moraine dams in Nepal. A similar coherence change was detected over nearby debris-covered glaciers and moraines that do not impound glacial lakes (Figure C.9). Minimum coherence for pixels in the moraine dam moving areas occurs in the early fall, suggesting that this is the timing of peak melt for buried ice within the moraine dams.

We found that the spatial distribution of buried ice was related to moraine dam geometry, which in turn reflects how the glacial lake formed. Buried ice was identified across most of the surface for wider, more rounded moraine dams such as those at Imja, Thugali, East Hongu 1, and gakal_gl_008 (Figure 4.1, Table C.4). These dams likely formed from remnant lower debris-covered glacier ice, as supraglacial lakes coalesced and expanded both up- and down-glacier (e.g., Imja Lake, Watanabe et al., 2009). Extensive buried glacier ice was left behind as the lake expanded and the active portion of the glacier retreated up-valley (Richardson and Reynolds, 2000).

Buried ice was frequently identified around the interior edge of the moraine dam, below the dam crest of narrow, sharp-crested moraine dams, such as Rolpa, Chamlang South, Nagma, kodud_gl_0205 (Figure 4.1). We hypothesize that at these sites, following glacier retreat, only buried ice adjacent to the narrow terminal and lateral moraines was preserved, rather than a wide portion of the lower glacier. Buried ice within lateral moraines frequently extended far up-valley of the moraine dam area. This ice was likely insulated by debris from the moraine and may also receive reduced insolation due to shading from steep moraine slopes and adjacent valley walls.

Our methods provide an objective approach to identify buried ice, and to evaluate previous, more qualitative approaches. We found strong evidence of buried ice in both narrow, sharp-crested dams and broad, rounded dams, suggesting that topographic form alone is not a reliable predictor of the presence of buried ice in this region (Table C.4). In many cases, the location of buried ice did correspond to surface features including ice cliffs, karst topography, surface water, and minor arcuate ridges. However, we also identified buried ice in locations where characteristic surface features were absent, for example, at the Upper Barun moraine dam, which has a relatively smooth surface with no ice cliffs, karst features, or surface water (Figure 4.1).

The binary presence of an ice core is an important factor in many hazard assessment frameworks, and misclassification of ice-cored moraine dams can substantially change resulting hazard and risk. Rounce et al. (2017) used the presence of an ice core to determine if a moraine dam is susceptible to "self-destructive failure." We find evidence for buried ice covering more than 10% of the moraine dam area in eight of the 13 moraine dams that Rounce et al. classified as "ice-free" (Table C.4). The presence of ice in these moraine dams was not sensitive to threshold choice, indicating high confidence. Applying their criteria, this reclassification shifts Nagma, Upper Barun, and gakal_gl_0008 from moderate to high hazard, and Muli Tal from high to very-high hazard.

Conversely, we did not find evidence for buried ice in all locations that display surface features (e.g., minor superimposed ridges, kettles, humocky textures) indicative of ice presence. For example, we found no evidence of internal ice in Hongu 1, one of ten dams Rounce et al. classified as ice-cored. Based on this finding, the hazard category for Hongu 1 would move from high to moderate. These results highlight the value of our approach to identify ice-cored moraine dams and improve the accuracy of hazard assessments.

We also emphasize the need for hazard frameworks to incorporate more detailed information about the spatial distribution, observed subsidence, and horizontal motion of buried ice, rather than relying solely on binary presence. For example, the kodud_gl_0205 moraine dam has small, discontinuous deposits of buried ice (Figure 4.1). While this dam might be classified as ice-cored, degradation of these buried ice deposits is unlikely to significantly affect its self-destructive potential or topographic characteristics such as freeboard height or height-to-width ratio. Furthermore, in all narrow, sharp-crested moraine dams with an existing breach, we found no evidence of buried ice along the breach itself, except for stray pixels near the top of the kotam_gl_0111 moraine dam breach (Figure 4.1E, G, L, P, T, U, V). This suggests that buried ice might play a limited role in GLOF initiation and evolution for these dams, assuming flow through the existing outflow channel, and not through a new breach. This also suggests that using indicators like change in the outflow channel to infer whether a moraine dam contains buried ice may not be effective, as buried ice is frequently absent in the outflow channel but present in other parts of the moraine dam. In contrast, for wide, rounded moraine dams (Table C.4), ice was frequently distributed throughout, including along existing outflow channels. At these locations, buried ice may play an important role in any potential dam breach event.

In some cases, we detected east/west and vertical movement in areas of moraine dams where coherence was high throughout the warm season, for example, at the Kekyap and Tallo Kekyap moraine dams (Figure 4.1). At these moraine dams, there may be no seasonal change in coherence (and associated buried ice detection) because the primary cause of surface movement is flow of massive ice, which occurs year-round, rather than melt of near-surface ice and incoherent surface debris layer motion. This interpretation is supported by the presence of ridges and furrows on the moraine dam and the strong downslope component of surface movement (Figure C.5R, S). Debris may be thick enough that warm-season melt of internal ice is minimal and does not result in surface reorganization sufficient to decrease coherence. Based on surface features and the observed deformation, we still interpret these moraine dams as ice-cored, but seasonal coherence change does not capture the presence of buried ice. This reflects a limitation of our buried ice mapping approach where buried ice experiences minimal melt in the warm season, potentially as a result of topographic shading, low temperatures, and thick and insulating debris cover. At these locations, measurements of surface deformation can still be used to help infer the presence of buried ice, rather than relying on surface morphology alone. Conversely, our buried ice mapping approach is also limited where surface change results in low coherence throughout the year (e.g., over a fast-moving glacier surface).

Our buried ice mapping approach can provide detailed information on the spatial distribution of internal ice. In the future, we hope to extend our approach to provide information on the depth of buried ice, i.e., the debris thickness. Debris thickness is an important additional factor in whether buried ice needs to be considered when modelling a moraine dam breach (Lala et al., 2017). If debris cover is very thick, overtopping waves may result in erosion sufficient to produce a GLOF without involving the underlying buried ice. We hypothesize that the magnitude of coherence change is related to the rate of ice core downwasting, which is in turn related to debris thickness, thermal conductivity, and temperature. By developing an empirical relationship between magnitude of coherence change, temperature, debris properties, and debris thickness, we hope to be able to estimate the thickness of debris overlying buried ice in moraine dams.

4.5.2 Moraine dam movement

We detected coherent surface motion in 19 of the 23 high-priority moraine dams in Nepal, indicating that movement on the order of centimeters to tens of centimeters per year is common in moraine dams that impound large glacial lakes in this region (Figure 4.1). Moving areas typically coincided with areas where seasonal coherence change indicated that buried ice is present, supporting the interpretation that surface deformation is caused by ice degradation or ice flow within the dam. In general, we expect that melt of buried ice will result primarily in subsidence, and occur mostly during the warm season, while flow of buried ice will result in both subsidence and horizontal motion, likely along the large-scale surface slope direction (Brencher et al., 2025). We found that most moraine dams with measurable vertical subsidence also exhibited measurable east/west motion (Figure C.5). Moraine dams with higher rates of subsidence tended to have higher rates of east/west motion (Figure C.7). This suggests that for most moraine dams, the internal deformation and flow of buried ice contribute significantly to overall moraine dam displacement. Buried ice layers 10–50 m thick with surface slopes below 15° are expected to produce surface-parallel movement on the order of centimeters to tens of centimeters per year (Figure B.9). Where debris is thick and insulating and/or air temperatures remain sufficiently low, the flow of buried ice may be the primary process explaining surface deformation. This is likely the case at Kekyap and Tallo Kekyap, where there is clear, extensive surface movement but no warm-season coherence loss observed at other moraine dams.

Melt rates are likely controlled by thickness and thermal conductivity of the debris cover, temperature, topographic shading, and hydrology. Where debris is relatively thin and/or temperatures are high, warm-season ice melt likely plays a larger role in moraine dam surface movement. The highest downward and east/west velocity (>10 cm/yr) were observed at the Imja, Barun, and Thugali moraine dams (Table C.5). Given that these moraine dams have average surface slopes, the relatively high east/west velocities are likely the result of relatively thicker massive buried ice. The relatively high observed vertical velocity likely indicates thinner debris cover and more melt, compared to other moraine dams. Overall, our results suggest that moraine dams with buried ice can continue to actively deform decades after glacier retreat. Extrapolating observed average subsidence rates of centimeters per year and assuming buried ice layers with thickness of ~10-50 m (based on the magnitude of horizontal motionREF), some buried ice could potentially remain in these 19 moraine dams for hundreds of years.

The rate of moraine dam subsidence has important implications for hazard assessment. Very slow summer subsidence rates (less than 1 cm/yr) likely indicate minimal melt of buried ice and a more gradual degradation of the ice core. This may reduce the likelihood of rapid internal channel development that leads to self-destructive failure. In contrast, where summer subsidence rates are higher, buried ice is likely melting more rapidly, and the ice core is degrading more quickly, which may reduce moraine dam stability or create channels for piping.

Observations of moraine dam surface velocity can also inform how frequently hazard assessments should be repeated. Rounce et al. (2017) recommend reassessments based on the rate of lake expansion. Glacial lakes with moraines dams that have surface velocity of tens of centimeters per year (e.g. Imja, Barun, Thugali) should also potentially be subject to repeated hazard assessments, even in the absence of lake growth. Over decadal timescales, the relevant topographic characteristics of a given moraine dam could change substantially. For moraine dams experiencing rapid deformation, GLOF modelling studies should also use updated high-resolution DEMs.

Future work should seek to isolate observed surface motion due to ablation of near-surface buried ice and surface motion related to flow of massive ice. One approach to disentangle these sources of movement is to consider only cold-season surface movement, when we expect minimal ablation, and can attribute observed surface motion to ice flow/deformation (e.g., Brencher et al., 2025). This approach should be possible where snow does not significantly decrease cold-season InSAR coherence. Cold-season surface movement due to ice flow, alongside topographic information and some parameters for the material properties of ice, can potentially be used to estimate the thickness of the buried ice. By subtracting the vertical component of cold-season displacement caused by flow of massive ice from the vertical component of warm-season displacement, we can potentially isolate and estimate the vertical displacement caused by ablation. Constraints on ice thickness, ablation rate, and debris thickness would provide useful additional information for GLOF hazard assessment and modelling, and for understanding the future evolution of ice-cored moraine dams.

4.6 Summary and conclusions

We used seven years of Sentinel-1 InSAR observations to measure surface deformation and map the distribution of buried ice for 23 high-priority moraine dams in Nepal between 2017 and 2024.

We found that most moraine dams (19) contain buried ice and have moving areas, which deform centimeters to tens of centimeters per year. We observed considerable variability in the size of moraine dam moving areas, the rates of moraine dam movement, and the spatial distribution of buried ice within moraine dams. We found that while the distribution of buried ice is related to moraine dam topographic characteristics and surface morphology, these characteristics alone are inadequate to accurately infer the presence of buried ice.

Our results suggest that the binary ice-cored classification frequently used in GLOF hazard assessments is insufficient to fully characterize the role of buried ice for GLOF hazards and characteristics. We identified evidence for extensive buried ice in eight moraine dams previously reported to be ice-free, and our method can provide more detailed information about the spatial distribution of buried ice in moraine dams over time. Future hazard assessments should incorporate this information to better consider the role of buried ice for moraine dam stability. Future numerical modelling studies that seek to simulate moraine dam breach evolution should account for the distribution of buried ice.

We found evidence for both melt and flow of buried ice within moraine dams. The relative contribution of these processes to overall surface movement is variable across moraine dams, and likely depends on debris thickness and thermal conductivity, ice thickness, and moraine geometry. Our results show that buried ice remains in most moraine dams and continues to deform for decades following glacier retreat. Observed moraine dam surface velocity should be considered when determining appropriate intervals and allocating resources for detailed hazard assessments for these glacial lakes.

Our results can inform more advanced geophysical modeling efforts, including estimating the debris thickness overlying buried ice, parsing the relative contributions to surface motion from ablation and ice flow, and estimating the thickness of buried ice. Overall, our results show that publicly available satellite SAR data can provide essential observations to improve both GLOF hazard assessments and modelling studies. Future integration of NISAR, commercial SAR constellations, and high-resolution DEMs will improve our ability to capture moraine dam surface change and deepen our understanding of their evolution.

Chapter 5

ACCURATE SNOW DEPTH PREDICTIONS ACROSS THE WESTERN U.S. USING A DEEP LEARNING MODEL TRAINED ON 7 YEARS OF AIRBORNE LIDAR SNOW DEPTH MEASUREMENTS

This chapter is in preparation as: **Brencher, G.**, Shean, D. E., Gagliano, E., and Henderson, S. T. ***Accurate snow depth predictions across the Western U.S. using a deep learning model trained on 7 years of airborne lidar snow depth measurements***

5.1 Abstract

Seasonal snow provides water for billions of people, but current regional observations of snow depth lack the spatial and temporal resolution needed for effective water resource management. Meanwhile, the publicly available archive of high-quality snow depth products has grown substantially, providing an opportunity to train multi-modal machine learning models to create improved snow depth products.

We compiled, corrected, and processed 233 airborne lidar snow depth products acquired between 2016 and 2023, containing ~271,000 km^2 of snow-covered area, similar to the area of Colorado. We used this dataset to train a U-Net convolutional neural network (CNN) to predict snow depth at 50-m resolution from publicly available input data sources including Sentinel-1 SAR backscatter images, Sentinel-2 multi-spectral images, the COP30 DEM, and coarse SNODAS physical model predictions. We performed extensive feature selection analyses to identify the best-performing input features for the final model.

Our final CNN substantially outperforms existing approaches for near real-time prediction of Western U.S. snow depth in accuracy, precision, and resolution, with an R^2 of 0.87 and a mean squared error of 0.31 m. Precision was best for shallow (<0.5 m) snow over low surface slopes, with decline in performance for extreme depth >5 m. We note excellent performance for both low and high fractional forest cover, overcoming a major challenge for snow depth prediction approaches in closed-canopy forests.

We developed an open-source Python package ("deep-snow") that can be used to predict snow depth for anywhere in the Western U.S. on any date over the past decade. We used this package to

create spatially continuous maps of snow depth over the entire Western U.S. and time series of snow depth for select test dataset locations with a 12-day interval over an 8-year period from 2016 to 2024. These predictions were validated against in-situ snow depth measurements from Snow Telemetry (SNOTEL) and California Cooperative Snow Surveys (CCSS) snow monitoring stations.

This work establishes a new benchmark for regional snow depth prediction performance, and demonstrates the potential of data-driven deep learning to better understand snowpack evolution, with implications for future operational forecasting.

5.2 Introduction

Seasonal snow is a critical water resource for global populations and ecosystems, and an important component of the global climate system. In the Western U.S., seasonal snow is responsible for 53% of total river discharge and 71% of discharge in mountainous areas (Li et al., 2017), providing essential water for agriculture, hydropower, municipal use, and outdoor recreation (Burakowski & Magnusson, 2012; Sturm et al., 2017). The amount of water stored in seasonal snow can be quantified as snow-water equivalent (SWE, with units of meters w.e.), which depends on snow depth and snow density. There is considerable spatiotemporal variability in SWE across the Western U.S.; the magnitude of annual April 1 SWE anomalies are typically 25-60% of the long-term (>50 year) average at any given site (Cayan, 1996), and can be much larger during years of snow droughts and snow deluges (e.g. Marshall et al., 2024). Given this interannual variability, effective management of water resources requires coordinated annual measurements to characterize the distribution of SWE across the Western U.S. Measurements of snow depth are more important than measurements of snow density for accurate SWE estimation, as variability in snow depth causes most of the spatial and temporal variability in SWE (Sturm et al., 2010; Painter et al., 2016), and uncertainty in bulk snow density is limited if snow depth is known (Avanzi et al., 2015; Gascoin et al., 2024), especially near peak SWE.

Each year, a suite of in-situ and remote sensing techniques are employed to measure snow depth around the world; these measurements are complemented by physical and/or statistical model predictions of snow depth. Each approach has strengths and limitations (see Hoppinen et al., 2024) and none alone can characterize the spatial distribution of snow depth at <100 m resolution with regional coverage, which is required to achieve many resource management and science goals (National Academies of Sciences, 2018). In particular, existing approaches tend to

perform poorly in mountainous environments due to complex terrain, deeper (>2-3 m) snowpack, and sparse in-situ measurements (Bormann et al., 2018).

Over the past decade, accurate, high-resolution observations of snow depth have been prepared using airborne lidar (e.g., Painter et al., 2016), satellite lidar (e.g., Besso et al., 2024), airborne photogrammetry (e.g., Meyer et al., 2022), and satellite photogrammetry (e.g., Hu et al., 2023). Airborne lidar consistently offers the most accurate remote sensing snow depth measurements, with validation efforts reporting uncertainty of centimeters to tens of centimeters (Harpold et al., 2014; Grünewald and Lehning, 2015; Painter et al., 2016; Currier et al., 2019; McGrath et al., 2019, Hu et al., 2023). These products have primarily been used for operational basin-scale applications such as runoff forecasting (e.g. Lahmers et al., 2022) and for validation of alternative approaches to measure snow depth and SWE (e.g. Bair et al., 2016; Deschamps-Berger et al., 2020; Yang et al., 2023). The public archive of airborne lidar snow depth measurements now includes hundreds of observations distributed over basins throughout the Western U.S. (Painter, 2018).

While airborne lidar observations have limited coverage in time and space, they can be used to train machine learning models that fuse data from different sources to produce distributed snow depth predictions. Once a machine learning model is trained, it can be efficiently applied to predict snow depth for dates and locations where airborne and/or in-situ snow depth measurements are not available. Previous efforts have employed machine learning models to make distributed snow depth predictions from data sources including spaceborne synthetic aperture radar (e.g., Dunmire et al., 2024; Broxton et al. 2024), meteorological station measurements (e.g., Herbert et al. 2025), passive microwave observations (e.g., Hu et al., 2021), airborne InSAR (e.g., Alabi et al., 2025), climate reanalysis data (e.g., Saranathan et al., 2025), meteorological data (e.g., Pflug et al., 2025), and snow cover (e.g., Betato et al., 2025). Many of these efforts used tree-based approaches (e.g., random forest, XGBoost), which are computationally efficient to train, require relatively small amounts of training data, and are more interpretable than deep-learning approaches. However, tree-based models consider each grid cell independently and cannot make use of valuable spatial information from neighboring grid cells. Given that snow tends to be highly spatially autocorrelated over short (<100 m) length scales (Trujillo et al., 2009), a machine learning approach that combines information from multiple grid cells to make predictions can likely offer improved performance.

Here, we present an approach that uses a convolutional neural network (CNN) to predict snow depth in the Western U.S. While CNNs are more complex and less easily interpretable than tree-based machine learning models, they naturally include multi-scale spatial context in their predictions, making them an appealing option for a wide variety of image processing tasks. Our CNN is trained using hundreds of airborne lidar snow depth products, including the entire Airborne Snow Observatories (ASO) archive from 2016-2023. We explore the performance of CNN snow depth predictions across the Western U.S. based on contemporaneous Sentinel-1 SAR, Sentinel-2 multi-spectral, and SNOw Data Assimilation System (SNODAS) model data. Our CNN can be applied over locations where no airborne lidar snow depth data or snow monitoring stations are available, helping to fill regional mountain snow observation gaps in both space and time.

5.3 Data

5.3.1 Target data: airborne lidar snow depth products

Snow depth products derived from differential airborne lidar serve as the target dataset that we use to train and evaluate our CNN. We downloaded all available ASO 50-m lidar snow depth products acquired during the 2016 to 2023 period from the NSIDC (Painter, 2018) and the ASO, Inc. archive. We also downloaded all nine 0.5-m Quantum Spatial, Inc (QSI) lidar snow depth products from the 2020 and 2021 SnowEx campaigns (Abedisi et al., 2022). In total, we considered 255 airborne lidar snow depth products. All products were downsampled by spatial averaging to match the 50-m resolution of the ASO snow depth products.

5.3.1.1 Quality control

We manually reviewed all airborne lidar snow depth products to identify artifacts and other quality issues (Figure D.1). We observed a border of anomalous zero snow depth values near the edge of valid snow depth pixels (Figure D.1B) in 25 products. We manually clipped these snow depth products to the extent of the valid snow depth pixels. We also identified clusters of physically-unrealistic, extreme snow depth values in 21 snow depth products, including some with zero values around the border. These artifacts were also frequently observed around the edge of the valid snow depth pixels (Figure D.1A, C, D). In the absence of additional data to validate these measurements, we chose to manually identify and remove these artifacts from the snow depth products, as slightly decreasing the dataset size is a worthwhile tradeoff to avoid training a machine learning model on spurious signals. We excluded three snow depth products with extensive artifacts. We also

removed 19 airborne lidar snow depth products where there was no corresponding Sentinel-1 acquisition spatially overlapping the snow depth product within 2 weeks of the lidar acquisition date. Our cleaned airborne lidar dataset consisted of 233 snow depth products.

5.3.1.2 Airborne lidar snow depth uncertainty

The reported mean absolute error (MAE) of the 50-m resolution ASO snow depth products is ~0.02 m based on 80 in-situ measurements near Tioga Pass, California (Painter, 2016). While uncertainty has not been assessed for the QSI snow depth products at 50-m resolution, the MAE of the original 0.5-m products is 0.06 m based on average snow depth within a 3-m buffer around six SNOTEL sites (Hoppinen et al., 2024). Independent assessments of the 3-m ASO products using in-situ and ground-penetrating radar snow depth measurements at Grand Mesa in Colorado reported root mean squared error (RMSE) between 0.14 and 0.46 m (McGrath et al., 2019; Hu et al., 2023). Spatial averaging of the 3-m ASO snow depth product to 50-m resolution is expected to decrease measurement error by a factor of ~4 (Painter et al., 2016). Applying this scaling factor, the estimated corresponding RMSE for the 50-m products should be between 0.03 and 0.11 m. Accuracy is likely to be reduced over steep and forested terrain (Deems et al., 2013). As such, we conservatively assume the uncertainty of our 50-m resolution airborne lidar snow depth products is 0.10 m. Overall, we expect that the uncertainty of the airborne lidar snow depth dataset is low enough that it can serve as effective training data and be used to evaluate our CNN, though we acknowledge that some component of the observed difference between model predictions and lidar snow depth measurements will be due to measurement error.

5.3.2 Input data

To prepare our input dataset for the CNN, we first selected data sources that provide either direct or contextual information about snow depth on a given date. We sought to include established and validated data products with broad spatial and temporal coverage. All data were reprojected to the appropriate local UTM zone, resampled, and cropped to match the 50-m grid spacing and bounds of the corresponding airborne lidar snow depth product. Where upsampling was required to match the airborne lidar snow depth grid, we used bilinear resampling, and where downsampling was required, we used spatial averaging.

5.3.2.1 Sentinel-1 SAR backscatter images

Sentinel-1 is a constellation of C-band synthetic aperture radar (SAR) satellites in sun-synchronous, near-polar orbits. In interferometric wide-swath mode, Sentinel-1 has a swath width of 250 km and a resolution of 5 m in range and 20 m in azimuth. For most of the Western U.S., the Sentinel-1 revisit interval is less than 12 days, with shorter revisit intervals (~3-6 days) for locations covered by swaths associated with multiple orbits. The first Sentinel-1A satellite was launched on April 3, 2014 by the European Space Agency Copernicus Programme, while the second Sentinel-1B was launched on April 25, 2016, improving temporal coverage.

Previous work has demonstrated that changes in Sentinel-1 backscatter can provide information about changes in snow depth where snow is dry (Lievens et al., 2019, Lievens et al., 2022, Brangers et al., 2024; Hoppinen et al., 2024). We used Sentinel-1 radiometrically terrain corrected (RTC) image products available on Microsoft Planetary Computer (Small, 2011; Microsoft Open Source, 2022), which have a pixel spacing of 10 m. We identified data for two periods for each airborne lidar snow depth acquisition: snow-on data collected within a four-week window (+/- 2 weeks) around the snow-on airborne lidar acquisition date, and snow-off data collected within a four-week window around September 10 of the previous year. We limited our search to data collected along the descending pass (~6:00 AM local acquisition time) to minimize the impacts of wet snow (more likely for ascending pass with ~6:00 PM local acquisition time), which attenuates the C-band radar signal and reduces the usefulness of the returns for measuring snow depth. We further limited our query to Sentinel-1 acquisitions with both VV and VH polarizations. For cases with multiple Sentinel-1 acquisitions within the four-week window that fully contained the entire airborne lidar snow depth product extent, we chose the acquisition closest in time to the airborne lidar acquisition for the snow-on data, and closest to September 10 for the snow-off acquisition. Otherwise, we selected the Sentinel-1 acquisition with the most complete coverage of the airborne lidar snow depth product within the four-week period. The mean difference in Sentinel-1 acquisition date and the airborne lidar acquisition date was -1.1 days, with a standard deviation of 4.0 days. For both the snow-on and snow-off Sentinel-1 data, we converted backscatter amplitude to units of decibels and calculated the polarization ratio, as described in Lievens et al. (2022) and Hoppinen et al. (2024). In addition, we calculated the change in polarization ratio between the annual snow-off observation and each corresponding snow-on observation.

5.3.2.2 Sentinel-2 multispectral images

Sentinel-2 is a constellation of high-resolution, multi-spectral optical imaging satellites in near-polar, sun-synchronous orbits. The first Sentinel-2 satellite was launched by the European Space Agency Copernicus Programme on June 23, 2015. Over the Western U.S., the maximum Sentinel-2 revisit interval is five days, with shorter revisit interval (~2-3 days) for locations covered by swaths from multiple orbits. Sentinel-2 data can provide important contextual information related to snow depth, including snow cover extent, snow reflectance, and vegetation and surface water extent. We used the Sentinel-2 Level-2A data products (Collection 1; European Space Agency, 2021a) hosted on Microsoft Planetary Computer (Microsoft Open Source, 2022). We identified Sentinel-2 acquisitions within a four-week window around each lidar snow depth acquisition with the greatest number of useful pixels over the airborne lidar snow depth product extent (i.e., pixels not identified as cloud cover, missing data, etc. in the Scene Classification Map). The mean difference in Sentinel-2 acquisition date and the airborne lidar acquisition date was -1.1 days, with a standard deviation of 3.8 days. We calculated normalized difference indices for snow (NDSI), water (NDWI), and vegetation (NDVI) using standard bands (Table 5.1).

5.3.2.3 SNOw Data Assimilation System (SNODAS) model output

The SNOw Data Assimilation System (SNODAS) products (NOHRS, 2004) include daily snow properties for the contiguous U.S. at 1-km resolution. SNODAS uses downscaled numerical weather model predictions to drive a physically-based, spatially-distributed energy- and mass-balance snow model, and assimilates satellite-derived, airborne, and in situ observations of snow-covered area and SWE. We extracted and upsampled the SNODAS snow depth product for the same date as each airborne lidar snow depth product. While the SNODAS snow depth products contain known issues (e.g., Bair et al., 2016; Yang et al., 2023), they offer a valuable initial coarse-resolution estimate of daily snow depth for the entire Western U.S. domain.

5.3.2.4 Fractional forest cover (FCF)

Forests influence the snowpack through multiple processes including canopy interception, shading, and modification of local wind and air temperature (e.g., Varhola et al., 2010). Forests also influence radar backscatter in snow-covered areas, potentially altering relationships between Sentinel-1 polarization ratios and snow depth (Lievens et al. 2022). We prepared fractional forest

cover (FCF, values from 0-100%) products from the PROBA-V global land cover dataset for the year 2019 (Buchhorn et al., 2020).

Table 5.1: Details for all data in the training dataset, including original data product resolution and citations. All input data (Section 5.3.2) were resampled to match the 50-m grid of the airborne lidar snow depth products (Section 5.3.1).

Name	**Description**	**Original Product Resolution (m)**	**Citation**
Airborne lidar snow depth			
lidar_sd	Airborne Snow Observatories (ASO) L4 lidar snow depth	50	Painter (2018)
lidar_sd	SnowEx lidar snow depth	0.5	Adebisi et al. (2022)
Sentinel-1 RTC SAR			
snowon_vv	Snow-on co-polarized (VV) return in dB	22 (range) by 20 (azimuth)	Small (2011); Microsoft Open Source (2022)
snowon_vh	Snow-on cross-polarized (VH) return in dB	"	"
snowoff_vv	Snow-off co-polarized return in dB	"	"
snowoff_vh	Snow-off cross-polarized return	"	"
snowon_cr	Snow-on cross ratio	"	"
snowoff_cr	Snow-off cross ratio	"	"
delta_cr	Change in cross ratio	"	"
Snow-on Sentinel-2 Level-2A			
blue	Blue band (center wavelength 0.49 μm)	10	European Space Agency (2021a); Microsoft Open Source (2022)
green	Green band (0.56 μm)	10	"
red	Red band (0.67 μm)	10	"
nir	Near-infrared band (0.84 μm)	10	"
swir1	Short-wave infrared band 16 (1.61 μm)	20	"
swir2	Short-wave infrared band 22 (2.19 μm)	20	"
ndvi	Normalized difference vegetation index (nir - red)/(nir + red)	10	--
ndwi	Normalized difference water index (green - nir)/(green + nir)	10	--
ndsi	Normalized difference snow index (green - swir1)/(green + swir1)	10	--
SNOw Data Assimilation System (SNODAS)			
snodas_sd	Snow depth	1000	NOHRSC (2004)
Fractional forest cover (FCF)			
fcf	PROBA-V global land cover, fractional forest cover from 0 to 100%	100	Buchhorn et al. (2020)
Copernicus GLO-30 digital elevation model (DEM) products			
elevation	Orthometric height above EGM2008 geoid	30	European Space Agency (2021b)
slope	Terrain slope	30	European Space Agency (2021b); Horn (1981)
aspect	Terrain aspect	30	"

northness	Sine of terrain aspect		“
eastness	Cosine of terrain aspect		“
curvature	Terrain curvature	30	European Space Agency (2021b); Zevenbergen and Thorne (1987)
tpi	Topographic position index	30	European Space Agency (2021b); Weiss (2001)
tri	Terrain ruggedness index	30	European Space Agency (2021b); Reily et al. (1999)
Other			
latitude	Geodetic latitude of each pixel	50	--
longitude	Longitude of each pixel	50	--
dowy	Day of water year for airborne lidar snow depth acquisition	50	--

5.3.2.5 Terrain elevation

Terrain exerts a primary control over the distribution of snow on the landscape. We prepared subsets of the Copernicus GLO-30 Digital Elevation Model (European Space Agency, 2021b) to match each 50-m airborne lidar snow depth product. The GLO-30 DEM product contains orthometric elevation above the EGM2008 geoid, with an absolute vertical accuracy (LE90) of <4 m and an absolute horizontal accuracy (CE90) of <6 m (NSDDA, 1998; European Space Agency, 2022). We also prepared derived terrain products (slope, aspect, northness, eastness, curvature, topographic position index, and terrain roughness index) from the resampled GLO-30 DEM using the xDEM package (xDEM contributors, 2024).

5.3.2.6 Auxiliary spatiotemporal information

We prepared auxiliary metadata raster data to capture essential spatiotemporal information for each airborne lidar snow depth product. We created raster coordinate maps with pixel latitude and longitude values. We also created a day of water year (DOWY) raster map using the acquisition date of each airborne lidar snow depth product, with a constant decimal day value for valid pixels with lidar snow depth coverage.

5.3.3 In situ validation data

To validate our snow depth predictions, we identified and prepared in situ snow depth measurements from all available Snow Telemetry (SNOTEL) and California Cooperative Snow Surveys (CCSS) stations (USDA NRCS, 2025) using the “easysnowdata” Python package (Gagliano, 2025).

5.4 Methods

5.4.1 Dataset preparation

5.4.1.1 Training, validation, and testing data

We defined training, validation, and testing areas using grids of tiles encompassing the airborne lidar products in each UTM zone. A tile size of 32 by 32 km was chosen because it is small enough to roughly capture the extent of most airborne lidar snow depth products (Figure 5.1), but large enough to provide observations of statistically independent (uncorrelated) snow depth values. We randomly assigned 80% of the tiles to training, 10% to validation, and 10% to testing (Figure 5.1).

For each 32-km tile, we divided the airborne lidar snow depth products and accompanying input data into 25 unique, evenly spaced 128 by 128-px (6.4 by 6.4 km) subsets. We chose an even spatial sampling scheme rather than a random spatial sampling scheme to ensure that all available data were included in the training, testing, and validation datasets while minimizing the size of those datasets in storage. We created corresponding binary missing data maps to record gaps in the airborne lidar, Sentinel-1, or Sentinel-2 data for a given subset. During training, we did not evaluate loss over locations with any missing input data.

The subsetting process resulted in 12,380 unique training data subsets, 2,160 unique validation data subsets, and 1,516 unique testing data subsets. The total area of the airborne lidar snow depth products included in the dataset is ~422,000 km^2, similar to the total area of California. This includes ~71,000 km^2 of unique coverage, with multi-temporal coverage over ~55,000 km^2. Of the ~422,000 km^2, approximately ~271,000 km^2 was snow-covered, for a total of ~1.08 x 108 50 by 50 m grid cells, similar to the area of Colorado. The test dataset, used to evaluate model performance, contained ~28,000 km^2 and ~1.12 x 10^7 grid cells. The total size of the compressed training, validation and testing dataset is 54.5 gigabytes.

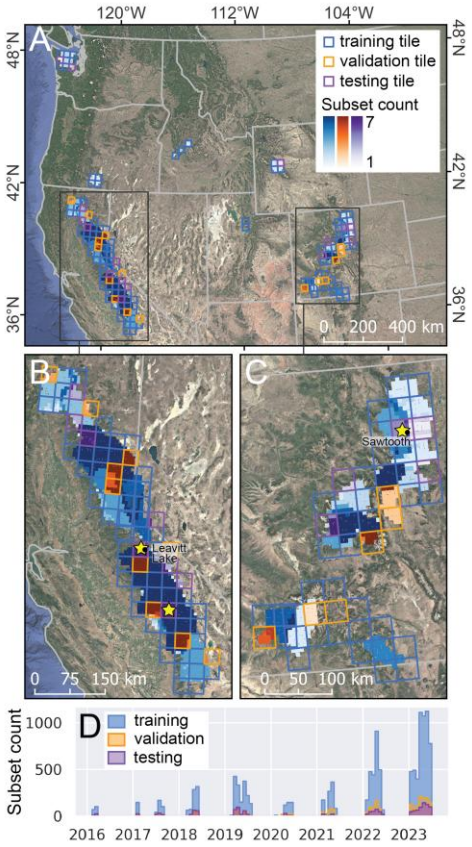

Figure 5.1: A) Spatial and temporal distribution of data in the training, validation, and testing datasets over the full Western U.S. domain, overlaid on the Google satellite basemap. Most Airborne Snow Observatories data were collected over the Sierra Nevada in California (B) and the Rocky Mountains in Colorado (C). Stars indicate location of three test tiles selected for additional analysis. Dots and labels show the locations of the predicted snow depth time series (Figure 8) for snow monitoring stations within the test dataset. D) Monthly time series showing the total count of 128 by 128-pixel (6.4 by 6.4 km) subsets used for training, validation, and testing.

5.4.1.2 Dataset normalization and augmentation

Before training, input features were normalized and clamped between 0 and 1 to speed up and stabilize convergence (LeCun et al., 2002). Input features were normalized using either their 0.1 and 99.9th percentile values across the full training dataset or their known global range (e.g., latitude was normalized using -90 as the minimum and 90 as the maximum). Normalization using percentile values involves a tradeoff, as very high and very low values that may be valid will saturate at 0 and 1 respectively. However, for input features with highly tailed distributions of values (e.g., Sentinel-2 surface reflectance), this strategy ensures that values occupy a similar range

of values to other input features rather than only occupying a portion of the range between 0 and 1. Meanwhile, normalization based on a known global range may force values to occupy a smaller portion of the range between 0 and 1, but ensures that the trained CNN can be applied outside the training area.

Data augmentation consisted of randomly rotating (0, 90, 180 or 270 degrees) and randomly flipping (horizontal and vertical mirroring) all input features throughout training. While this process could potentially disrupt the CNN's ability to apply consistent directional information to predict snow depth (e.g., ignoring regions expected to be affected by radar layover for typical descending pass Sentinel-1 acquisition geometry), we found that data augmentation improved training outcomes for our dataset size.

5.4.2 CNN architecture

We used a UNet CNN implemented with PyTorch (Paszke et al., 2019) with five downsampling and upsampling layers, and 64 initial filters increasing to a maximum of 512 filters, all of size 3 by 3 grid cells. Rectified linear unit (ReLu) activation was performed after each convolutional layer except the final one. Our training minimized the mean squared error (MSE) loss between the CNN prediction and the normalized target airborne lidar snow depth.

5.4.3 Feature selection

Feature selection is an essential step of model training, as 1) including features that are not helpful to the CNN can impair the ability of the CNN to generalize, 2) features that are highly correlated with one another can impair effective training and provide redundant information (LeCun et al., 2002), 3) a smaller number of features decreases CNN complexity and simplifies CNN interpretability, and 4) reliance on many external data sources increases the chances that a data availability issue will disrupt CNN application in practice. To identify the optimal group of features we performed a series of experiments, including 1) Monte Carlo feature importance, 2) Principal Component Analysis, 3) an addition experiment, which are described in detail in the following sections.

5.4.3.1 Monte Carlo feature importance

Monte Carlo experiments for feature selection involve training models on many randomly-selected groups of features in order to understand how inclusion of specific features contributes to model performance (e.g., Dramiński et al., 2008). We implemented a Monte Carlo-style experiment to

determine the relative importance of each potential input feature. We randomly selected groups of five features and trained the CNN for 100 epochs on a subsample of 512 randomly selected training samples with 256 randomly selected validation samples. Each sample has dimensions of 5 features by 128 by 128 pixels. Batch size (the number of samples processed before updating the model) was set to 64 to minimize step-to-step variability in validation loss (Godbole et al., 2023). The learning rate was set to 0.0004. We performed 217 trials of this experiment such that each feature was included an average of 37.4 times, and recorded the minimum validation loss achieved during each trial. We then compared the distribution of minimum validation loss for trials where a given feature was included, and trials where that feature was not included. Features that are more useful to the CNN are expected to result in better performance when they are included compared to when they are not included. A Welch's t-test, which assumes a normal distribution of values but does not assume equal variance or sample sizes, was used to establish the significance of differences between groups.

5.4.3.2 Principal Component Analysis

Principal Component Analysis (PCA) is commonly used to select a group of variables that contribute the most to the variation in a dataset (Jolliffe, 2002). We assessed the uniqueness of each of our features by performing a PCA on the validation dataset using scikit-learn (Pedregosa et al., 2011). Each two-dimensional feature was flattened into a one-dimensional array and any pixel with data gaps for any feature was excluded. As all features were normalized during data loading, no additional scaling was applied. Feature loadings on the first two principal components were used to evaluate each feature's contribution to overall variance. Features with distinct loadings were interpreted as more unique, while features with similar loadings were considered potentially redundant.

Based on the Monte Carlo experiment and PCA results, we selected a group of "core" variables which are important relative to other features and non-redundant. This group represents an initial feature configuration for model training that yields reasonable performance and serves as a foundation for further tuning of the feature group (Godbole et al., 2023).

5.4.3.3 Addition experiment

Addition experiments for feature selection involve systematically adding features to evaluate their effect on model performance. These experiments can be computationally expensive, as they

involve training a model on many different groups of features. They are effectively the reverse of commonly performed ablation experiments (e.g. Meyes et al., 2019) which involve systematically removing features. While ablation experiments have been used for recent snow depth machine learning model feature selection efforts (Betato et al., 2025), we chose to perform an addition experiment, as it is less computationally expensive than an ablation experiments when the final feature group is expected to be small relative to the total number of features being evaluated. Here, we further reduce computational expense by starting with a core group of features which we expect (based on the previous experiments) to be included in the optimal feature group, reducing the number of combinations of features for evaluation. We then determine which combination of the non-core "ancillary" features results in optimal CNN performance when included alongside the core features (See Figure D.2 for a schematic of the addition experiment). The experiment proceeds through a number of different rounds, during which we evaluate increasingly large combinations of ancillary features. During each round, we performed multiple trials, where the CNN was trained on a given combination of features.

First, we performed a control run, where we trained the CNN using only the core features to establish a baseline performance metric for the validation dataset. We performed 20 trials where we trained the CNN for 150 epochs on a subsample of 512 randomly selected training samples with 256 randomly selected validation samples. The batch size was set to 64 and the learning rate was set to 0.0004. We recorded the minimum validation loss achieved during each trial.

In Round 1 of the addition experiment, for each of the 21 ancillary features, we trained the CNN using a feature group including that ancillary features and the core features. As in the control run, we performed 20 trials for each group and retained the minimum validation loss achieved in each trial. If mean validation loss for the group improved relative to the mean baseline validation loss (achieved using the core features alone) and the probability that improvement was due to chance was less than 50% (based on Welch's t-tests) the ancillary feature advanced to the next round. The group with the best mean validation loss was used as the baseline for the next round.

In Round 2, we evaluated each possible combination of two features from the set of ancillary features that advanced from the previous round. Each of these combinations of ancillary features was included alongside the core features and used to train the CNN for 20 trials, as in the prior rounds.

After Round 2, only a single combination of features resulted in improved performance relative to the baseline with a less than 50% chance the improvement was due to chance. We selected these ancillary features to augment our core features in the final feature group for hyperparameter tuning and CNN training.

5.4.4 Hyperparameter tuning

We followed standard practices for CNN hyperparameter tuning (e.g., Godbole et al., 2023), using the AdamW optimizer (Loshchilov and Hutter, 2017), an MSE loss function, and a CNN depth of 5 layers based on exploratory experiments on a portion of the training data. With these hyperparameters fixed, we then tuned the learning rate and weight decay following a standard two-stage random search approach (Godbole et al., 2023). We first performed a coarse random search over a large parameter space. We performed 20 trials of up to 500 epochs on the full dataset, using learning rate decay and early stopping to prevent overfitting. Based on the best-performing trials from this stage, we then performed a fine random search over a smaller parameter space. We selected the hyperparameters from the trial achieving the minimum validation loss as our final set of hyperparameters for CNN training. To train the final CNN, we folded the validation dataset into the training dataset to maximize the data available for fitting.

5.4.5 CNN evaluation

5.4.5.1 Error metrics

We trained our "final CNN" using the final feature group and final set of hyperparameters. To evaluate our final CNN, we performed error analysis using our test dataset to quantify systematic and random errors in the resulting snow depth predictions. We first used our final CNN to generate snow depth predictions for the entire test dataset, and calculated the prediction error (residuals) by subtracting the de-normalized predicted snow depth from the target airborne lidar snow depth measurements for each grid cell. For each grid cell i we compute the snow depth residual δSD_i:

$$\Delta SD_i = \widehat{SD}_i - SD_{lidar,i} \tag{5.1}$$

where $\widehat{SD}_i$ is predicted snow depth and $SD_{lidar,i}$ is the measured snow depth. For snow-covered grid cells (airborne lidar snow depth > 0), we calculated commonly used performance metrics including Pearson Correlation Coefficient (R), Coefficient of Determination (R^2), relative error (percent error relative to the airborne lidar value), MSE, and MAE, in addition to the mean, standard

deviation, median, and normalized median absolute deviation (NMAD) of the residuals. The median and NMAD metrics are robust to outliers and should be more representative of true bias and dispersion than the mean and standard deviation. We also calculated MSE and MAE for snow-free grid cells (airborne lidar snow depth = 0).

While the 50-m airborne lidar snow depth measurements are treated as the reference in this analysis, they have a measurement uncertainty (1-sigma) of approximately 0.10 m (Section 5.3.1.2). These lidar errors are expected to vary across terrain and land cover types and may be correlated with the final CNN error. Because the spatial structure and correlation of lidar uncertainty are not known, we cannot fully separate model error from measurement error or perform a complete uncertainty propagation. As a result, our reported error metrics represent the combined effects of model error and lidar error and may slightly overestimate the model error relative to the true snow depth.

5.4.5.2 CNN performance with different input data sources

To explore how different input data sources affect CNN performance, we performed an ablation study, retraining the CNN after removing one of the input data sources (e.g., SNODAS, Sentinel-2) from the input dataset. We otherwise used the same features from the final feature group and the final hyperparameter set during training. We then evaluated test dataset performance using the metrics described in Section 5.4.5.1 to better understand how the CNN makes use of data from different sources.

5.4.5.3 Bias and dispersion analyses

To evaluate the final CNN performance under varying conditions, we examined how both bias and dispersion of residuals were associated with potential covariates, including static spatial variables for terrain elevation, terrain slope, terrain aspect, forest cover fraction (FCF), mountain range (Snethlage et al., 2022), snow class (Sturm and Liston, 2021), and temporal variables such as year and month of year. Because the variance of the residuals may change as a function of these covariates (heteroscedasticity), we divided continuous variables into ten quantiles, and computed bias (mean of signed residuals) and dispersion (standard deviation of residuals) within each quantile bin. We used quantiles rather than fixed bin intervals to ensure an adequate sample count ($n \approx 1.1 \times 10^6$) in each bin.

To model the expected bias and dispersion as a function of input conditions, we identified a subset of variables most strongly correlated with the bias and dispersion of the prediction error. We then performed multidimensional linear interpolation across the binned variable space to estimate expected bias and dispersion values for each bin combination.

To assess whether the modeled bias and dispersion explained the systematic components of the error, we first corrected the predicted snow depth using the modeled bias, then recomputed the prediction error using the bias-corrected snow depth. We then standardized, dividing by the modeled dispersion (the standard deviation of the residuals). To assess the extent to which the distribution of the resulting standardized residuals matched a normal distribution, we evaluated skewness and kurtosis and examined deviations from the 1:1 line in quantile-quantile plots.

5.4.6 CNN ensemble

While we initially evaluated the performance of the single "final CNN", we also evaluated the performance of a CNN ensemble, including five different CNNs trained with combinations of hyperparameters (Table D.1; Section 5.4.4). The ensemble snow depth prediction was calculated as the median of snow depth predictions from the five individual CNNs.

5.4.7 CNN inference

We developed an open-source Python package called deep-snow (Zenodo citation here) that prepares all input datasets and applies our final CNN to generate snow depth predictions for a user-specified area of interest and date range. The earliest date for which valid predictions can be made is June 27, 2015 with the start of Sentinel-2 data availability. The geographic extent for which valid predictions can be made depends on the user-selected CNN. For CNNs trained with SNODAS data (Section 5.3.2.3), the valid extent is the Contiguous United States. For CNNs trained without SNODAS data, the valid extent is near-global, with limited gaps due to available Sentinel-1 and Sentinel-2 coverage. However, considering that the training dataset was limited to the Western U.S., we recommend limiting inference to the Western U.S. with the current models.

We now describe how the deep-snow software predicts snow depth for a given area of interest on a single date of interest ("target date"). Regardless of the size of the area of interest, the software implements the following strategy to prepare valid input datasets.

Rather than using a single best Sentinel-1 and Sentinel-2 image as input data, as we do when constructing the training dataset, we first construct Sentinel-1 and Sentinel-2 composites by

calculating the pixel-wise median of respective data products overlapping the area of interest within a 12-day window around the target date. By creating composites rather than using a single image, we mitigate issues caused by incomplete coverage of the area of interest. This is important where the area of interest is too large to be covered by a single Sentinel-1 or Sentinel-2 scene. For the Sentinel-2 mosaics, we mask cloudy pixels in individual images using the Scene Classification Map prior to composite generation.

To identify valid data for the composite generation, we search for images within a 12-day window around the date of interest. Where there are no available Sentinel-1 or Sentinel-2 overpasses within the initial time window, we increase the window length by four days until an overpass is included. In practice, these date range parameters can also be adjusted by the user using known Sentinel-1 and Sentinel-2 intervals for their area of interest.

For inference/prediction, all input features are divided into 1024 by 1024 pixel tiles with 100 pixel overlap at tile boundaries. The outer 50 pixels of the predicted snow depth tiles are removed to mitigate potential poor UNet performance near image edges (Innamorati et al., 2018), and the cropped tiles are mosaicked into a single continuous snow depth prediction for the user-defined area of interest.

To generate snow depth time series, we apply the CNN as described above to generate a snow depth map for each 12-day period within the longer period of interest. Using a 12-day interval ensures that unique Sentinel-1 and Sentinel-2 data are available for each prediction. In practice, unique predictions can be generated for each day during the period of interest, as SNODAS snow depth products have a daily interval, though it is possible that some of the Sentinel-1 and Sentinel-2 input data will be shared across multiple days.

5.4.8 Sample predictions for evaluation in this study

To explore the performance of our CNN over broad spatial extents, we generated a single continuous snow depth prediction for the entire Western U.S. for a single day. We selected the date of maximum average snow depth measured at the in-situ monitoring stations in 2024, which was March 13.

To explore the performance of our CNN over long time periods, we generated a continuous snow-depth time series with 12-day interval from October 10, 2017 to May 1, 2025 (230 predictions) for two test tiles including portions of the upper Tuolumne River basin in California and the upper Colorado River basin in Colorado.

5.4.9 In situ validation

In addition to the model validation described above (Section 5.4.5), we validated our snow depth predictions using in situ snow depth measurements (Section 5.3.3). For the March 13, 2024 Western U.S. snow depth prediction, we first calculated prediction error at each station using the nearest 50-m grid cell on that date (in situ snow depth minus predicted snow depth), then calculated the mean and standard deviation of the error and the median relative error (percent error relative to in situ snow depth measurement). For evaluation of the snow depth time series predictions, we calculated Pearson's correlation coefficient (R) between predicted and measured snow depth for a representative station location with a continuous, high-quality in situ measurement record within each test tile. These stations included the Leavitt Lake, California SNOTEL station (site number 574) and the Sawtooth, Colorado SNOTEL station (site number 1251).

5.5 Results

We first describe the results of our feature selection experiments, which we used to select the feature group used to train the final CNN model. Next, we describe the final CNN model predictions and report on model performance. Lastly, we describe the results of applying the model and validating the model predictions.

5.5.1 Feature selection

The Monte Carlo experiment indicated that the SNODAS snow depth was the most valuable feature, followed by Sentinel-2 bands and band indices (Figure 5.2A). The MSE for the validation dataset was significantly lower when SNODAS snow depth, Sentinel-2 blue, red, near-infrared, and green bands, and Sentinel-2-derived NDSI and NDWI (distributions with asterisks on left side of Figure 5.2A) were included among the features used to train the CNN vs. when they were not included. When the SNODAS snow depth was included in a randomly selected feature group, the most valuable additional features included various terrain characteristics and day of water year (Figure D.3).

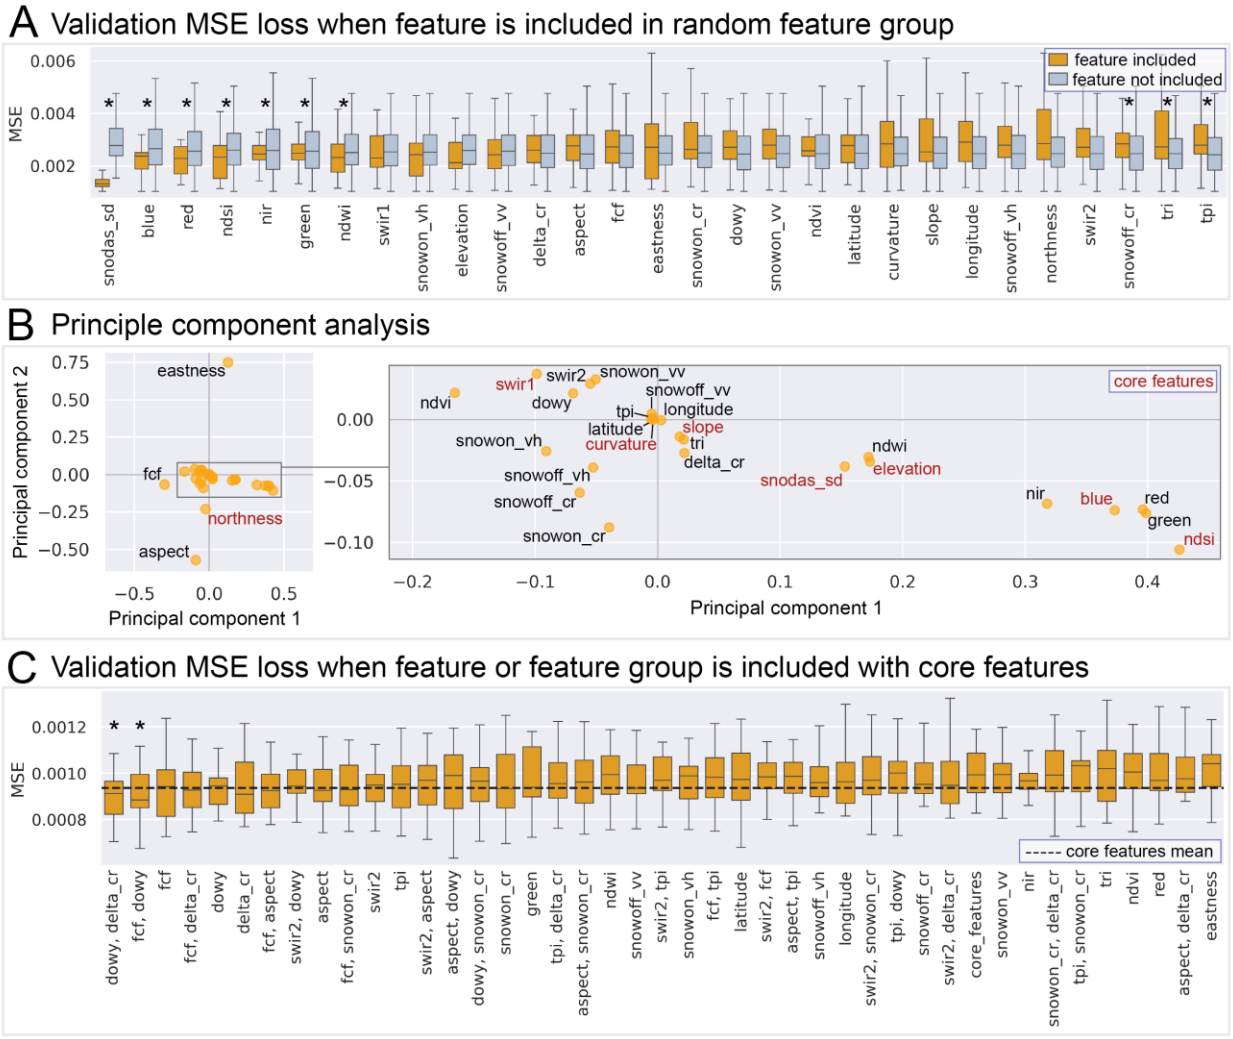

Figure 5.2: Summarized results of feature selection experiments. A) Results of Monte Carlo experiment (Section 5.4.3.1). Asterisks indicate a significant difference in MSE when a given feature is included in the feature group used to train the CNN vs. when the feature is not included. Features are arranged on the x-axis by their probability of improving MSE loss when included. The 7 features with asterisks on the left side of the plot show significant decrease in MSE (improvement) when included, whereas the 3 features on the right side show significant increase in MSE. B) Results of principal component analysis for all features (Section 5.4.3.2). The positions of features indicate their loadings on the first two principal components. The final set of "core" features are labeled in red, for reference. C) Results of addition experiments (Section 5.4.3.3). Asterisks indicate a significant difference in minimum validation loss for a feature group when compared to the core features only. Feature groups are ordered by mean MSE loss.

Following the PCA experiment, we found that the first principal component accounted for 29% of the validation dataset variance. This component was positively associated with features expected to have high values over snow (e.g., NDSI, Sentinel-2 visible bands) and negatively associated with features expected to have low values over snow (e.g., FCF, NDVI; Figure 5.2). The second principal component explained 22% of the variance and was positively associated with features linked to east-facing slopes (e.g., eastness) and negatively associated with features characteristic of north-facing slopes (e.g., northness, NDSI). Based on these two experiments, we selected a core feature group including SNODAS snow depth, Sentinel-2 NDSI, the Sentinel-2 blue and shortwave infrared 1 bands, elevation, slope, northness, and curvature.

During the addition experiment, we found that many "non-core" ancillary features improved validation loss when used to train the CNN alongside the core features (Figure 5.2C). The best-performing features groups tended to include day of water year, fractional forest cover, and Sentinel-1 delta cross ratio.

The final feature group (both "core" and "non-core ancillary" features) used to train the "final CNN" included SNODAS snow depth, Sentinel-2 NDSI, the Sentinel-2 blue and shortwave infrared 1 bands, elevation, slope, northness, curvature, day of water year, fractional forest cover, and Sentinel-1 delta cross ratio.

5.5.2 Final CNN snow depth predictions

The final CNN snow depth predictions for the test dataset (Section 5.4.5) largely capture the measured distribution of snow over the landscape with different terrain characteristics and during different months of the year (Figure 5.3). The CNN snow depth predictions generally captured small-scale (<200 m) structures in snow depth (Figure 5.4). Prediction error was largest in areas where snow depth is controlled by local wind redistribution and avalanching, i.e., steep slopes above the treeline (Figure 5.5). Qualitatively, the predicted snow depth maps in the test dataset show slightly less detail than the original measured airborne lidar snow depth products, especially in alpine environments where snow depth varied on shorter length scales (Figure 5.5).

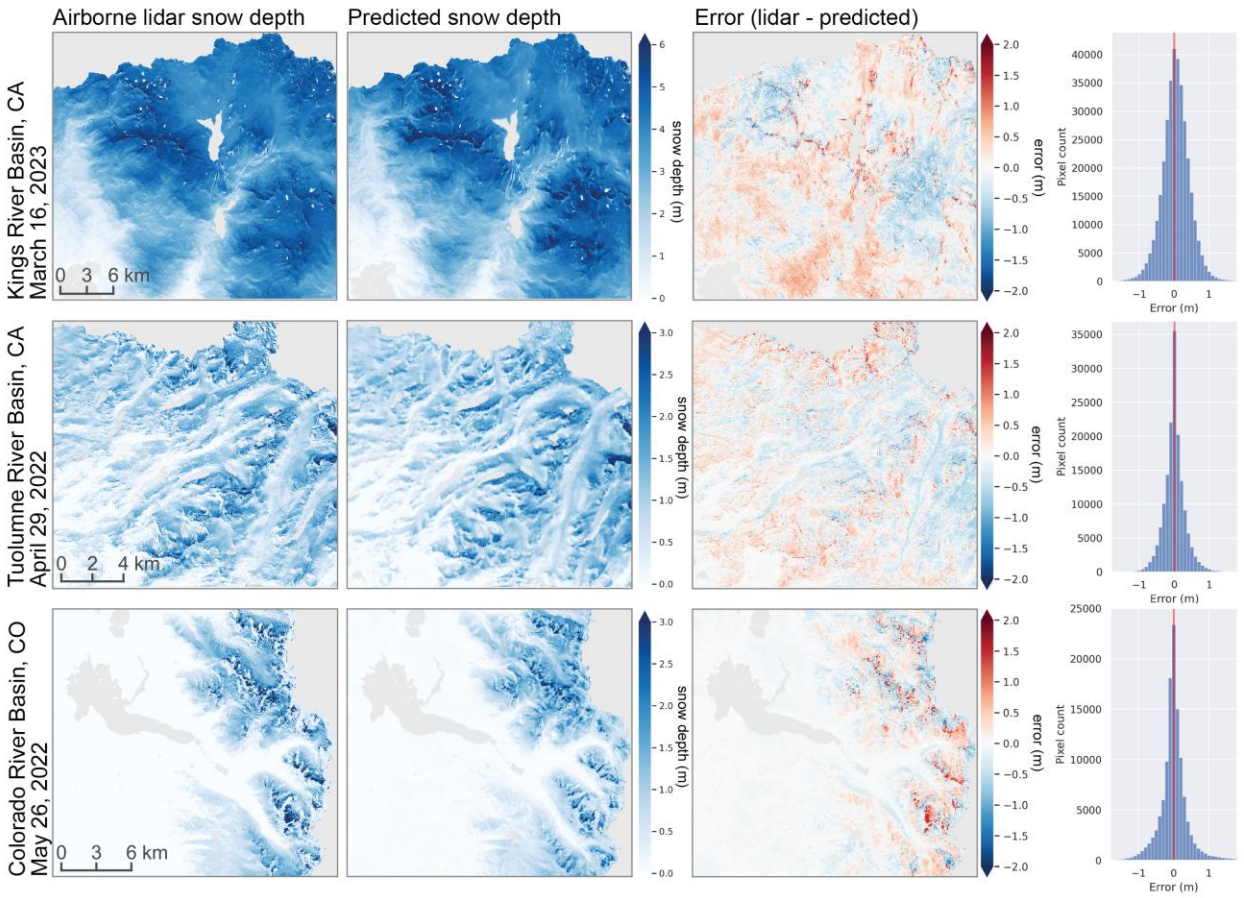

Figure 5.3: Example CNN snow depth predictions on the corresponding airborne lidar acquisition date for three unseen test tiles (stars in Figure 5.1). Columns show the ground truth airborne lidar snow depth, predicted CNN snow depth, the per-pixel difference map and histograms. Only snow-covered grid cells (airborne lidar snow depth > 0) are included in the histograms, as CNN performance is excellent (zero difference) over nearly all snow-free grid cells (Section 5.5.3.1). The black outlines in the middle row show the location of Figure 5.4.

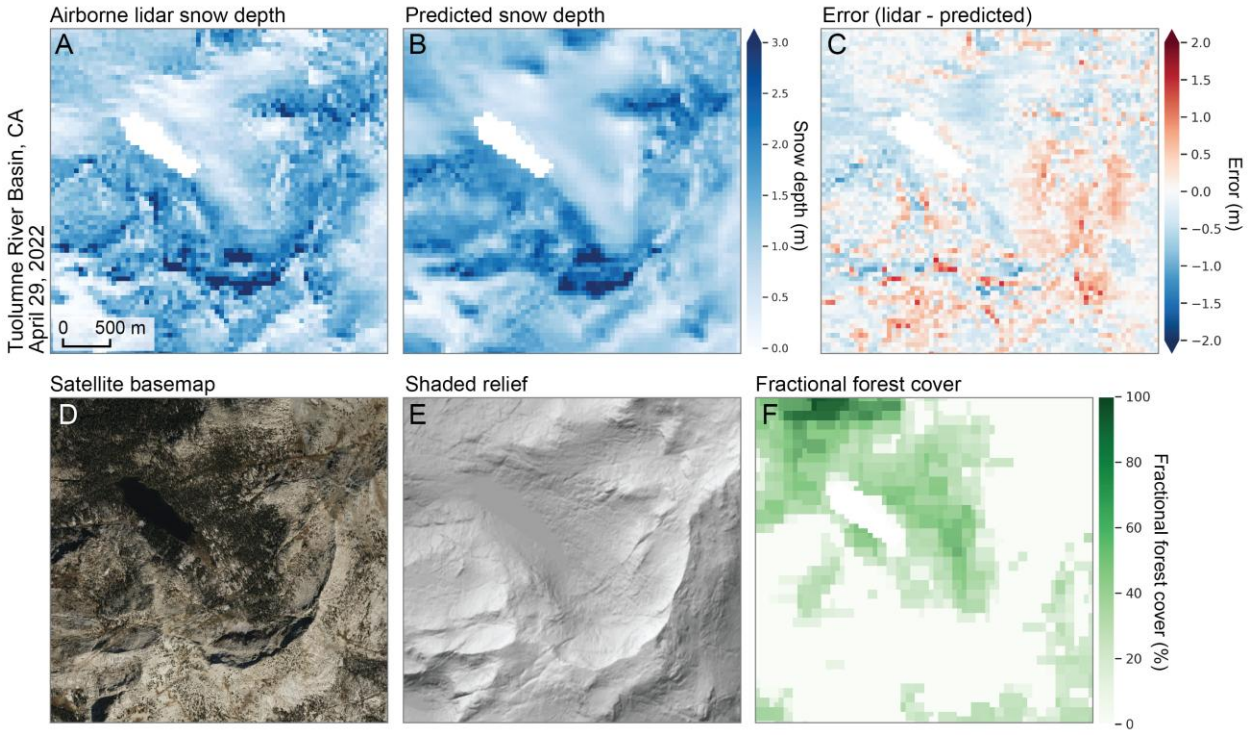

Figure 5.4: Example showing detail of CNN snow depth predictions for a portion of an unseen test tile in the upper Tuolumne River Basin, CA (see Figure 5.1 for context). A) Ground truth airborne lidar snow depth, B) predicted final CNN snow depth, C) the per-pixel difference, D) Microsoft Bing Satellite Basemap, E) Shaded relief of 10-m U.S. Geological Survey 3D Elevation Program digital terrain model (USGS, 2025), F) Fractional forest cover (Buchhorn et al., 2020). See Figure 5.3 for final CNN snow depth predictions for the entire test tile.

5.5.3 CNN performance

5.5.3.1 Overall error metrics

Over all grid cells in the testing dataset (Figure 5.1) with measured airborne lidar snow depth greater than zero, the final CNN had a Pearson correlation coefficient R of 0.93 and MSE of 0.31 m (Table 5.2; Figure 5.5). The median of the prediction error was -0.01 m, indicating low overall bias, with an NMAD of 0.35 m, indicating moderate precision.

Over grid cells with no snow, the final CNN had a MSE of 0.02 m and MAE of 0.04 m. Where airborne lidar snow depth was greater than zero, the CNN predicted zero snow depth for only 1.6% of grid cells. These false negatives typically occurred over grid cells with very shallow snow; the median snow depth resulting in a false negative prediction was 0.02 m, which is within the expected uncertainty of the lidar snow depth products.

The CNN ensemble (Section 5.4.6) had improved performance metrics compared to the individual CNNs (Table 5.2, Figure D.4).

Table 5.2: Performance of different models for snow-covered grid cells in the test dataset.

Model	R	MSE (m)	MAE (m)	Mean of residuals (m)	Standard deviation of residuals (m)	Median of residuals (m)	NMAD of residuals (m)	Median relative error (%)	NMAD relative error (%)
Final CNN	0.93	0.31	0.36	0.01	0.55	-0.01	0.35	20.20	21.46
Final CNN (bias-corrected)	0.94	0.29	0.34	0.02	0.53	0.00	0.33	19.47	21.22
CNN ensemble	0.94	0.28	0.33	0.02	0.53	-0.01	0.33	19.17	20.52
SNODAS only	0.85	0.74	0.60	-0.25	0.82	-0.22	0.59	31.17	36.74

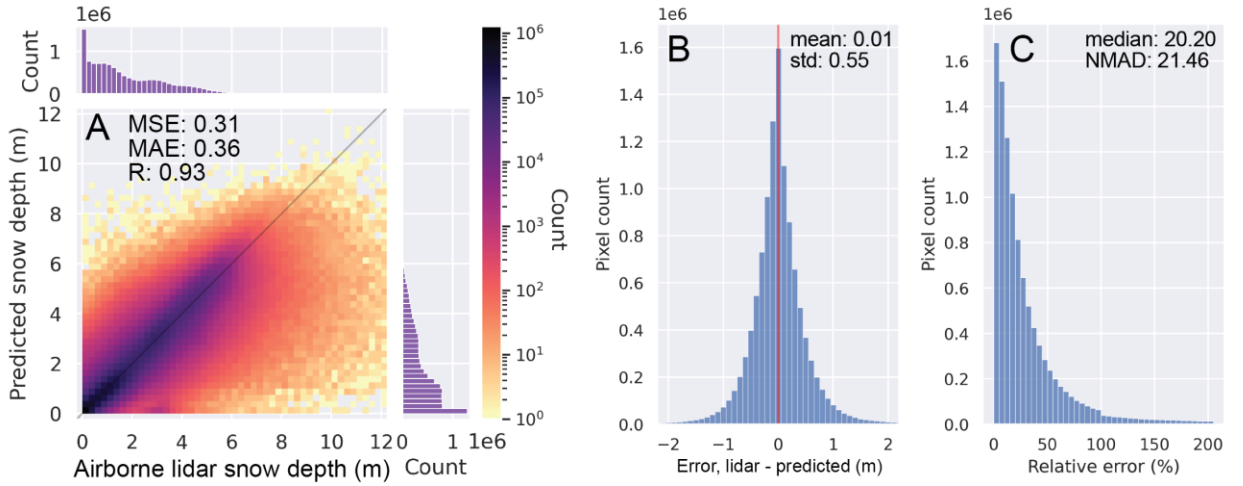

Figure 5.5: Final CNN performance for snow-covered grid cells in the test dataset (n≈11.2M grid cells covering ~28,000 km^2). A) 2D histogram showing airborne lidar snow depth and predicted snow depth. Grid cells with an airborne lidar snow depth of zero, where CNN performance is excellent, are not included in the plots or the reported metrics. B) Histogram of the residuals calculated by subtracting the predicted snow depth from the airborne lidar snow depth. The difference is attributed to CNN prediction error. Overall, CNN prediction bias is low. C) Histogram of relative error. We report the median and normalized median absolution deviation (a robust measure of dispersion, equivalent to the standard deviation for a normal distribution), to limit the influence of outliers related to negligible airborne lidar snow depth (<0.01 m), which may result in large relative error even where absolute error is small. Note that false negatives, where predicted snow depth is zero and airborne lidar snow depth is not zero, always result in a relative error of 100% and are not included in this plot.

5.5.3.2 Performance with different data sources

Excluding SNODAS snow depth from the training data (Section 5.4.5.2) substantially degraded CNN performance on the test dataset across all evaluation metrics (Table D.2). For example, the R^2 decreased from 0.87 to 0.79 and the MSE increased from 0.31 to 0.48. The CNN trained without SNODAS snow depth also exhibited a mean underprediction of 0.25 m on the test dataset. Performance also declined when Sentinel-2 and topographic features were excluded from training, though to a smaller degree. Mean bias on the test dataset increased from 0.01 m to 0.03 m and 0.06 m when Sentinel-2 and COP30-derived features were excluded from the training data. In contrast, removing Sentinel-1 features, including change in polarization cross ratio, fractional forest cover, and day of water year, from the training dataset led to only minor reductions in performance.

5.5.3.3 Bias with covariates

To explore potential bias in the final CNN snow depth predictions, we examined the relationship between prediction bias on the test dataset and potential covariates. We found bias associated with certain covariate categories defined by quantile binning (Section 5.4.5.3). Aspect showed the strongest association with bias. The final CNN overpredicted snow depth on north- and south-facing slopes (aspect range of 324–36° and 144–216°, respectively) by 0.09 m, and underpredicted snow depth on east- and west-facing slopes (aspect range of 72–108° and 252–288°, respectively) by 0.15 m (Figure 5.6D). The final CNN tended to underpredict snow depth greater than 4 m (~10% of snow-covered grid cells) by an average of 0.33 m and snow depth greater than 6 m (<1% of snow-covered grid cells) by an average of 1.78 m. When predicted snow depth exceeded 3.1 m, the final CNN overpredicted snow depth by an average of 0.08 m (Figure 5.6A). It also overpredicted snow depth by 0.06 m on flat slopes (< 6°) and underpredicted snow depth by 0.10 m on steep slopes (> 27°, Figure 5.6C). In areas with high fractional forest cover (> 0.9), the final CNN overpredicted snow depth by an average of 0.07 m (Figure 5.6B).

We observed no clear bias related to mountain range, snow class, or month of the year (Figure 5.6E-G). We observe variable bias during different years, with a mean bias magnitude of 0.02 m (Figure 5.6H). After applying the bias correction to the final CNN (Section 5.4.5.3), most error metrics improved (Table 5.2). The interannual bias decreased slightly, with the mean magnitude dropping from 0.073 to 0.066 m, suggesting that the observed interannual bias is not fully explained by predicted snow depth, slope, aspect, or forest cover. After bias correction, we observed a new pattern in monthly bias (Figure D.7). The bias-corrected final CNN underpredicted

snow depth in March by an average of 0.10 m and overpredicted snow depth in May, June, and July by an average of 0.04 m. This indicates that prior to bias correction, the observed bias related to month of year was masked by the larger bias related to spatial characteristics.

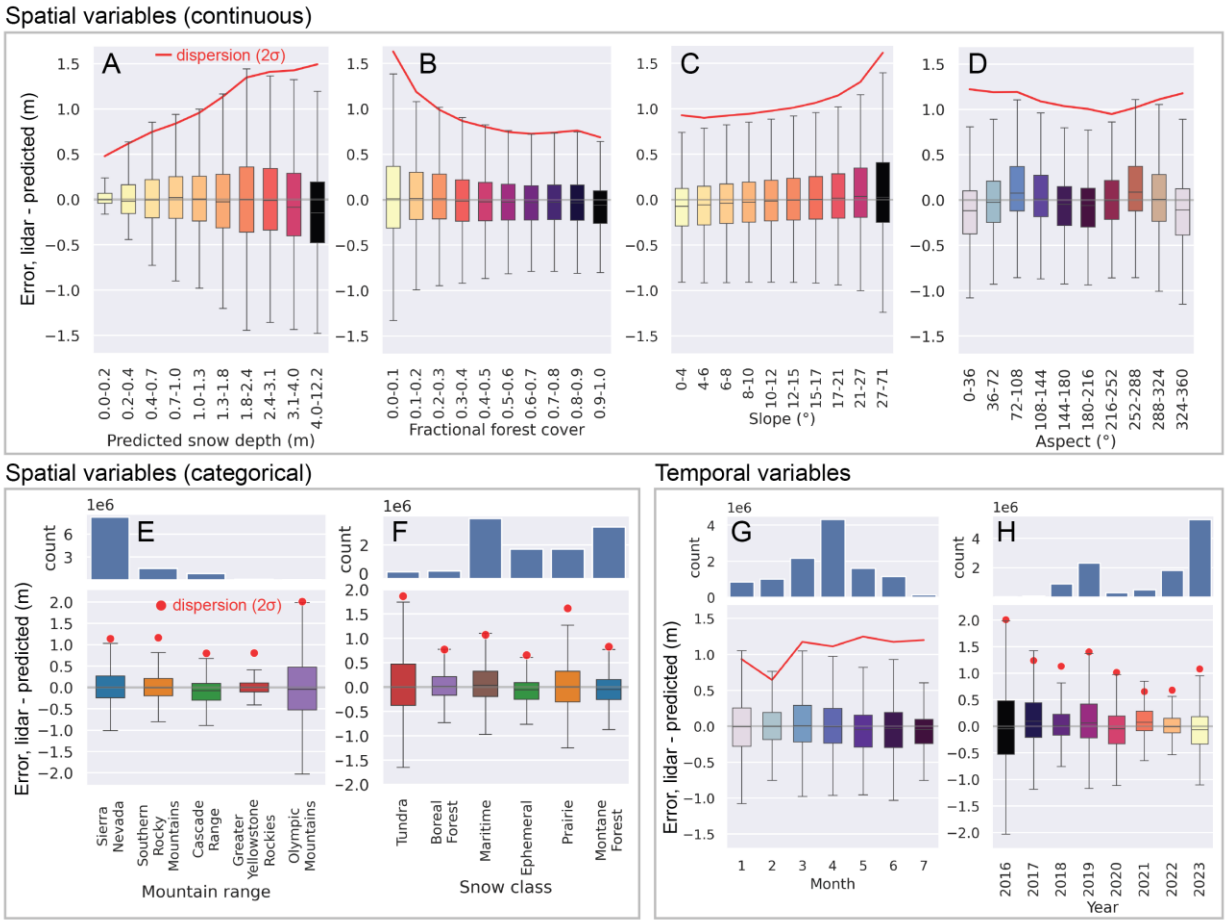

Figure 5.6: Observed relationships between final CNN prediction error and key spatial and temporal variables. Red lines and red markers show 2-sigma dispersion. Grid cells with airborne lidar snow depth of zero are not included. Prediction error vs. continuous spatial variables including A) predicted snow depth, B) fractional forest cover, C) terrain slope, and D) terrain aspect. Note that bin ranges are not necessarily uniform across the 10 quantiles, so a linear color ramp corresponding to the full range of the x-axis variable is used for box color. Prediction error for E) different mountain ranges, and F) snow classes, with accompanying histograms showing the number of grid cells in each category. Prediction error for temporal variables, including G) month of the year, and H) individual years.

5.5.3.4 Dispersion with covariates

We observed substantial heteroscedasticity in snow depth prediction error (airborne lidar – predicted) across the test dataset, with increased error dispersion for greater predicted snow depth and surface slope, and decreased error dispersion for greater fractional forest cover (Figure 5.6A–C). Dispersion of error was also lower on southwest-facing slopes and higher on north-facing slopes (Figure 5.6D). Dispersion of error was largest on steep (> 27°), north-facing slopes with deep snow (> 1.8 m) and no forest cover. Conversely, dispersion was smallest on relatively shallow (< 8°), southwest-facing slopes with shallow snow (< 0.2 m) and substantial forest cover (> 50%).

To account for this variability, we standardized the prediction error using four continuous spatial variables: predicted snow depth, slope, aspect, and fractional forest cover (Figure D.6). After standardization, the error distribution of the test dataset approximated normality (Figure D.7). The observed variability in error dispersion across mountain ranges, snow classes, months of the year, and years (Figure 5.6E–H), was removed by the standardization (Figure D.7). This suggests that snow depth, terrain, and forest characteristics largely explain heteroscedasticity of the prediction error.

Notable exceptions included persistently high dispersion in the Olympic Mountains across all years (Figure D.7E), high dispersion across the entire domain in 2016 (Figure D.7H), as well as a slight increase in dispersion from February to May (Figure D.7G). Apart from these cases, the standardized error dispersion was similar across regions, snow classes, and years. We observe nominal dispersion and bias in 2023, which can be considered an outlier year in some areas of the Western U.S. due to extreme snowfall (Marshall et al., 2024).

5.5.4 Final CNN inference and in situ validation

We applied our final CNN (Section 5.4.4) to predict snow depth across the contiguous Western U.S. at 50 m resolution on March 13, 2024, the date of peak snow depth in water year 2024 across the full domain (Figure 5.7). This computation required about one hour using the deep-snow batch computing workflow.

Qualitatively, the final CNN snow depth prediction is realistic, reproducing expected spatial distribution and magnitudes across the Western U.S. There are some artifacts at Sentinel-2 swath boundaries caused by 1) differing snow cover on different Sentinel-2 acquisition dates within the 12-day window around the target date and 2) no available cloud-free imagery. When compared to

the in-situ snow depth from the SNOTEL and CCSS stations (n=863), the mean prediction error (in situ minus prediction) was 0.18 m (underprediction) with a standard deviation of 0.49 m (Figure D.8). This result is consistent with the monthly bias observed in March for the test dataset, after accounting for biases due to snow depth, terrain slope and aspect, and forest cover fraction (Section 5.5.3.3). The median relative error was 26.97%.

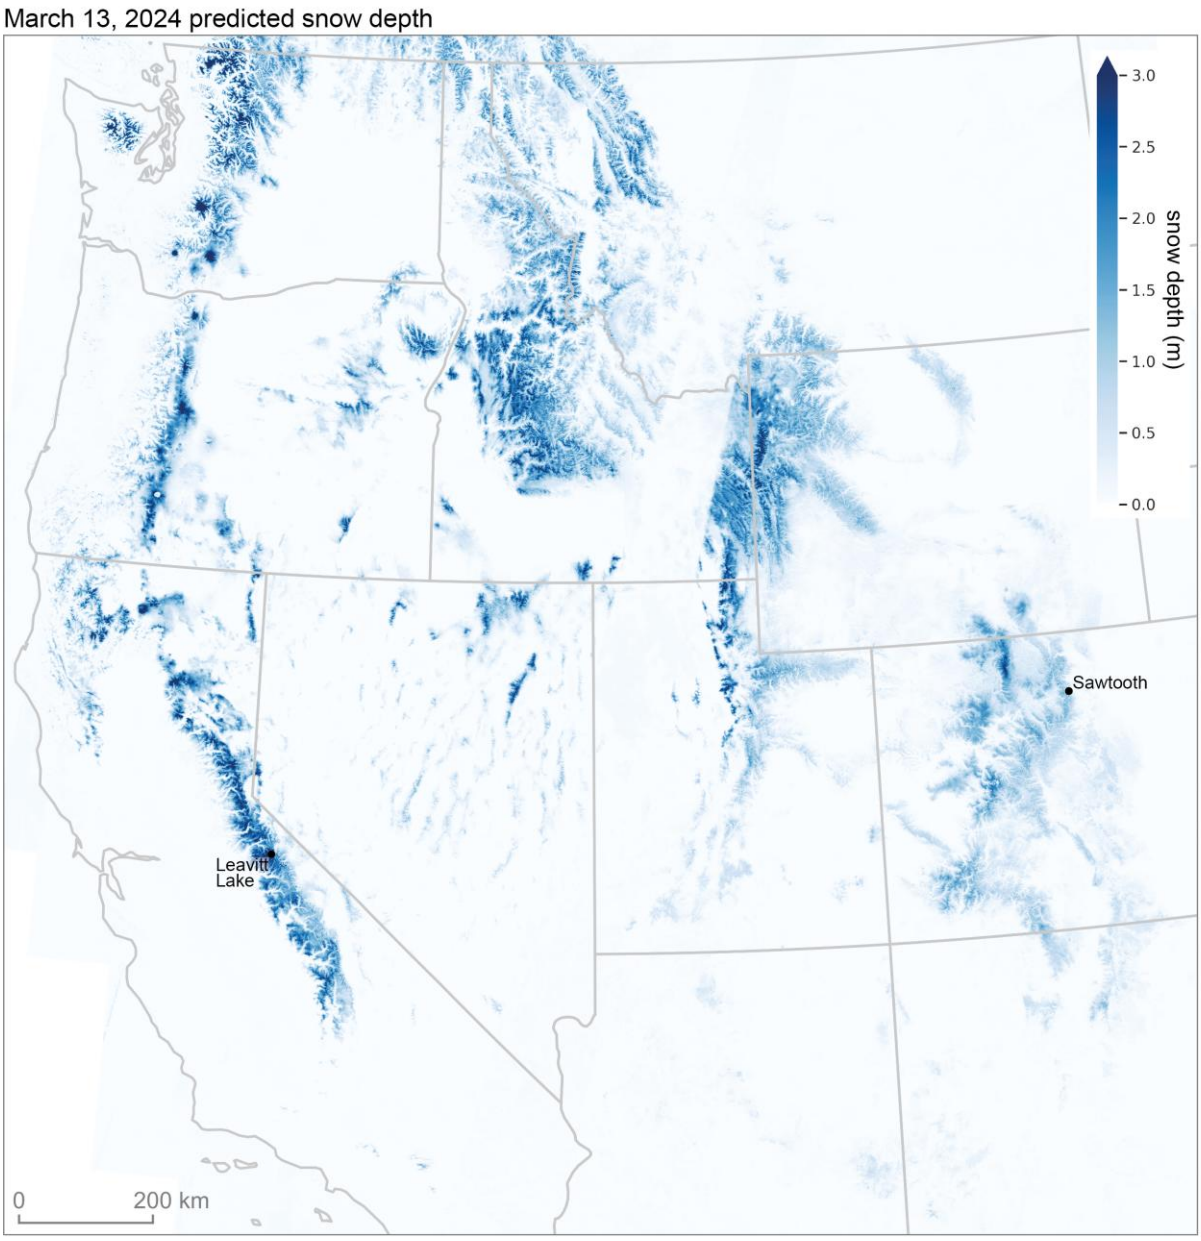

Figure 5.7: CNN snow depth prediction for March 13, 2024, the date of maximum snow depth for in-situ monitoring stations over the Western U.S. Prediction spatial resolution is 50 m. Marked locations correspond to the snow monitoring stations featured in Figure 5.8.

Our final CNN snow depth time series predictions showed general agreement with the timing and magnitude of in-situ snow depth measurements at the Leavitt Lake, CA and Sawtooth, CO SNOTEL stations, with R values of 0.96 and 0.83, respectively (Figure 5.8). At both locations, the final CNN overpredicted snow depth between October and January in some years (e.g., 2022, 2024). At the Sawtooth SNOTEL station, the final CNN largely captured the temporal pattern of snow depth fluctuation during the winter, but tended to underpredict snow depth by about 0.5 m between January and May.

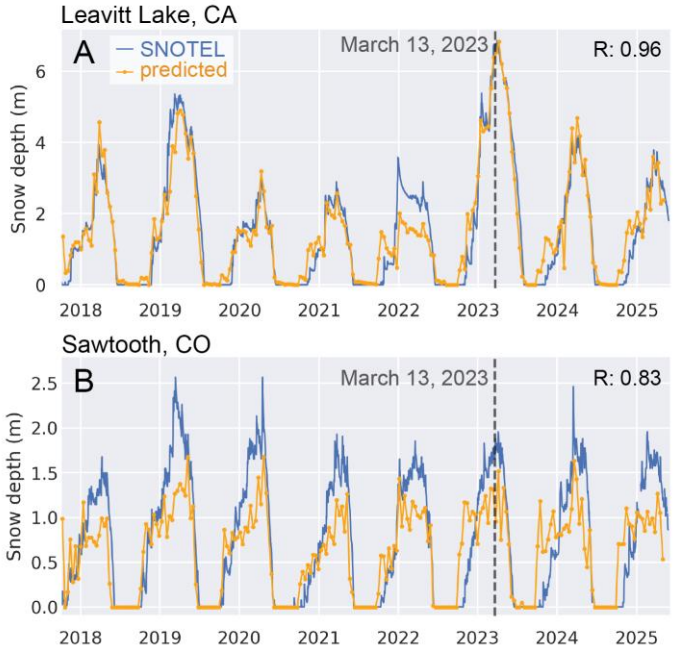

Figure 5.8: Snow depth time series for the final CNN predictions (orange) and observed in situ measurements (blue) at the Leavitt Lake, CA and Sawtooth SNOTEL stations, both located within unseen test tiles (Figure 5.1, Figure 5.3). Vertical dashed lines correspond to the March 13, 2024 date for the CONUS snow depth map shown in Figure 5.7.

5.6 Discussion

5.6.1 Feature importance

5.6.1.1 Sentinel-1

None of the input features derived from Sentinel-1 data (including snow-on and snow-off VV and VH RTC backscatter, snow-on and snow-off backscatter polarization ratio, and change in polarization ratio) were found to be important relative to other features in the Monte Carlo

experiment. However, in the addition experiment, including the Sentinel-1 change in polarization ratio alongside the core features consistently led to a slight improvement in CNN performance. Removing this feature from the training data resulted in a small decline in test performance (Table D.2). This finding is consistent with Dunmire et al. (2024), who reported that change in Sentinel-1 cross ratio was less important than variables such as snow cover fraction, cumulative snow cover fraction, elevation, snow class, and day of snow year. Mirza et al. (2025) found that including Sentinel-1 snow depth in a data assimilation framework alongside snow disappearance date resulted in limited or no improvement in performance. The Sentinel-1 polarization cross ratio as calculated here is likely not strongly correlated with snow depth across much of the test data. Its utility to the CNN might improve if the change in cross ratio was integrated over time and scaled to approximate snow depth (e.g., Lievens et al., 2022; Hoppinen et al., 2024), potentially providing a noisy initial estimate. Broxton et al. (2023) found that snow depth products derived from Sentinel-1 (rather than change in cross ratio) had the highest importance from a set of features that included snow density, tree cover, temperature, elevation, month, latitude, or longitude, though this input dataset did not include any other features directly related to snow depth.

5.6.1.2 SNODAS input

The relative importance of the SNODAS snow depth as an input feature, despite its coarse resolution and poor standalone performance on the test dataset (Figure D.5), suggests that the CNNs benefit from incorporating prior snow depth estimates that can be refined, rather than relying solely on predictors of snow depth. This finding is consistent with previous work showing that a random forest model using SNOTEL snow depth can effectively predict the difference between SNOTEL values and snow depth at nearby grid cells (Herbert et al., 2025). When SNODAS snow depth was included among the input features during our Monte Carlo experiment, the most informative additional variables were primarily terrain-related, suggesting that topographic characteristics are important for refining SNODAS snow depth estimates (Figure D.3).

Among all input data sources, removing SNODAS snow depth from the set of features used to train the CNN resulted in the greatest decrease in CNN performance across evaluation metrics. However, the CNN trained without SNODAS (Section 5.4.5.2; Section 5.5.3.2) still substantially outperformed the upsampled SNODAS product on the test dataset. This suggests that the CNN is not simply downscaling the coarse SNODAS snow depth information; it is combining

complementary information from SNODAS and other input features. This interpretation is further supported by a comparison of observed biases. The upsampled SNODAS snow depth product has a mean bias of -0.21 m, whereas the CNN snow depth predictions have a bias of 0.01 m (Table 5.2). The corresponding shift in the central tendency of the distribution of input SNODAS snow depth to the distribution of CNN snow depth predictions is atypical of standard statistical downscaling methods, which generally preserve the central tendency, and suggests that additional sources of snow depth information are being used.

5.6.2 CNN performance

5.6.2.1 Comparison with other snow depth models

To our knowledge, our CNN provides the most accurate and precise near-real-time snow depth maps produced to date for the western U.S. at resolutions less than 100 m. One alternative approach is the Lievens et al. (2022) empirical model (Section 5.3.2.1). Hoppinen et al. (2025) found that the 90 m output from this model had an RMSE of 0.92 m and an R value of 0.46 compared to airborne lidar snow depth observations over sites across the Rocky Mountains. Our CNN represents a large improvement over these results, with an RMSE of 0.55 m and an R value of 0.93 at 50 m resolution.

We also compare the performance of our final CNN and other recent machine learning approaches for predicting snow depth on a single date. Dunmire et al. (2024) reported a mean absolute error of 0.33 m and a Pearson R value of 0.78 for non-zero snow depth using an XGBoost model trained at 100 m resolution over the European Alps. Resampling our CNN snow depth predictions on the test dataset to 100-m grid cell spacing improved our MAE to 0.31 m and improved our Pearson R to 0.95. Herbert et al. (2025) reported a median RMSE of ~0.45 m when applying their random forest model, trained in Colorado at 50 m resolution, for basin-scale snow depth predictions. They evaluate performance over snow-free grid cells and in snow-covered areas with snow depth < 5 m, and rely on snow depth measurements from nearby SNOTEL stations as input features. When we restrict our evaluation to test dataset grid cells with snow depth < 5 m and include snow-free grid cells, we obtain a comparable RMSE of 0.44 m. Broxton et al. (2023) reported R^2 and relative mean absolute difference (RMAD) values of 0.76 and 0.40 m over the Sierra Nevada and 0.70 and 0.38 m, respectively, over the Colorado Rockies for their random forest bias-corrected and gap-filled Sentinel-1-derived snow depth predictions at 500 m resolution.

Our test dataset RMAD is 0.18 m, which represents a substantial improvement. In summary, our approach matches or exceeds the performance and resolution of previous efforts while offering improved spatial resolution and broader applicability across the western U.S., including in areas without in-situ monitoring stations.

5.6.2.2 Underprediction of deep snow

The final CNN exhibits strong overall performance on the test dataset, but interpretation of bias and precision under certain conditions reveals where the CNN is less effective and can potentially be improved. One notable limitation is a tendency to underpredict snow depth greater than 4 m (Figure 5.5A). Underprediction of extreme values is a common issue in machine learning regression tasks, including snow depth prediction (e.g., Dunmire et al., 2024; Herbert et al., 2025). This underprediction likely stems from the distribution of snow depth in the training data, where moderate snow depth are far more common than extreme snow depth. This problem may be further exacerbated by the use of loss functions such as mean squared error (MSE), which disproportionately penalize large errors. As snow depth predictions diverge from the central tendency of the training distribution, the risk of large penalties increases, discouraging the CNN from predicting extreme snow depth values. To address this, some prior studies have used fine-tuning strategies that reweight the loss based on sample frequency, improving performance on rare, high-magnitude cases at the potential cost of overall CNN accuracy (e.g., Lang et al., 2023).

5.6.2.3 Aspect bias

The CNN tends to overpredict snow depth on north- and south-facing slopes and underpredict on east- and west-facing slopes. When terrain features (i.e., elevation, slope, curvature, northness) were excluded from the training dataset, a different and stronger bias emerged: the CNN overpredicted snow on south-facing slopes and underpredicted on north-facing slopes. This pattern is consistent with the general tendency for north-facing slopes in the Northern Hemisphere to accumulate deeper snow due to reduced insolation, causing CNN underprediction. Supplying the CNN with topographic information, including northness, mitigates but does not eliminate aspect-related bias, suggesting that terrain features help the CNN adjust for aspect-driven snow depth variation, but additional features, training data, or model architectures may be needed to fully resolve this issue.

5.6.2.4 Monthly bias

When we applied the bias correction to our CNN snow depth predictions (Section 5.4.5.3), a bias related to month of the year became apparent (Figure D.7). This seasonal bias may arise from the available temporal distribution of the training data, which spans January to July, but is concentrated in April compared to other months (Figure 5.6G). The April bias in the test dataset is negligible, but when the CNN is applied to dates in February and March, it tends to overpredict snow depth; when applied to dates in May through July, it tends to underpredict. While these opposing biases may largely cancel for the full January through July period, they can be non-negligible on specific dates. For example, on March 13, 2024, the final CNN model had a 0.18 m negative bias relative to in-situ measurements (Figure D.8).

These results may indicate that the CNN does not effectively make use of the day of water year feature in its current form. Future exploration of alternative approaches to encode this temporal information (e.g., encoding positionally within an image) could improve performance. Another option would be to adjust the loss function to increase the weight for snow depth during months that are less frequent in the training dataset, which could reduce temporal bias, but may affect overall performance. We don't expect these options to improve snow depth predictions during times of the year that are not included in the training or testing dataset, namely the early accumulation season, where the CNN tends to overpredict snow depth. The ideal way to address this issue is to expand the training dataset to include more data from different times of the year, though this depends on the availability of high-quality snow depth products. An alternative is to apply post-hoc bias correction using snow monitoring station data or the test dataset.

5.6.2.5 Interannual bias

While we observed minor interannual biases that were not explained by differences in predicted snow depth, aspect and slope, or fractional forest cover, these biases do not appear to reflect an inability to capture interannual variability in snow depth. If the CNN were failing in this regard, we would expect predictions close to the mean or median snow depth of the training dataset for a given day of year and location, regardless of the year. In this way, a CNN that does not capture interannual variability could minimize error, though it would result in overprediction in low-snow years and underprediction in high-snow years. We observe the opposite: in 2023, a high accumulation year across much of the Western U.S. (Marshall et al., 2024), the final CNN slightly overpredicts snow depth, while in 2021, a low accumulation year, it slightly underpredicts (Figure

5.6). This pattern suggests that the CNN can capture interannual variability but may still have small annual residual biases.

5.6.2.6 CNN precision

We found that CNN precision was highest in areas with relatively shallow snow, low slopes, and dense forest cover, and lowest in areas with deep snow, steep slopes, and sparse or no forest cover. These patterns likely reflect 1) the dominant physical processes controlling snow accumulation and redistribution, and 2) the spatial length scales over which snow depth varies. In steep terrain above the treeline, snow depth is strongly influenced by complex processes such as wind redistribution and avalanching, which produce highly heterogeneous snow depth distributions. These fine-scale patterns are not captured by the coarse-resolution SNODAS snow depth product and must instead be inferred from topographic information, snow extent, texture, and color, and potentially changes in Sentinel-1 cross-polarization ratio. Snow depth in such areas can vary over short distances (tens of meters, e.g., between a wind-scoured ridge and an adjacent leeward slope), limiting the usefulness of nearby grid cells for inferring snow depth at a given location. While the CNN captures many of these patterns remarkably well in unseen terrain (Figure 5.3), larger prediction errors in these settings are expected. In contrast, in relatively flat, forested areas, snow depth is less influenced by wind or avalanching, and spatial patterns of snow depth tend to vary over longer length scales (hundreds of meters). The range of likely snow depth is narrower, and neighboring grid cells are more likely to have similar values. In these environments, the CNN has access to more consistent contextual information and is operating within a more constrained prediction range, both of which contribute to improved precision.

5.6.3 Considerations for operational use

Our CNN can be applied to predict snow depth on a given date at the basin, mountain range, and regional scale with a temporal resolution of about 12 days. The temporal resolution can also be improved where Sentinel-1 revisit times are shorter or when Sentinel-1 data are not included in the version of the CNN used for prediction, which results in slightly degraded test dataset performance (Table D.2).

When combined with snow density estimates (i.e., empirical density models; Hill et al. (2019), Sturm et al. (2021)), our CNN snow depth predictions can be used to estimate SWE for a range of

applications. Figure D.9 demonstrates this capability for a distributed SWE estimate derived from our CNN snow depth output for an arbitrary location in the Western U.S. domain.

Our CNN snow depth predictions (and corresponding SWE estimates) may be especially valuable in ungauged basins where limited or no snow depth information is available and late in the water year when the snowline advances above the elevation of most snow monitoring stations. In these situations, our CNN likely substantially outperforms alternatives for real-time snow depth estimation, such as SNODAS and the Lievens et al. (2022) algorithm.

The CNN is computationally efficient enough for near-real-time inference on a standard workstation or cloud environment, making it suitable for operational use in daily snow monitoring workflows. To support operational use, we provide a batch-computing workflow that runs on currently free cloud infrastructure and requires no user-provided compute (deep-snow repo). This workflow can generate a snow depth map for the entire Western U.S. in approximately one hour, or an 8-year daily snow depth time series for a 1.5° × 1.5° tile in about one hour. All predicted snow depth output is provided at 50 m resolution. We achieved the best performance using the median snow depth prediction from an ensemble of five CNNs, and with bias-correction applied based on the test dataset. A bias-corrected CNN ensemble is likely to produce the most accurate and precise results in an operational setting. Prediction uncertainty can potentially be approximated from the ensemble spread, which may offer a useful proxy for confidence in individual grid-cell estimates in the absence of in-situ data.

Although overall bias on the test dataset was very low, we found non-negligible intra- and interannual biases that may impact the utility of CNN predictions (Figure 5.6, Figure 5.8, Figure D.6). Where available, in-situ monitoring stations should be used to assess CNN bias and precision and potentially to perform bias correction. This is especially important for snow depth predictions in the early accumulation season, which is not represented in the training dataset. In this period, comparisons to in-situ measurements show consistent overprediction of snow depth. Periodic retraining with additional years of data may improve CNN performance, particularly for early- or late-season snow conditions not well represented in the current training dataset.

Some issues that may degrade performance in operational settings can be mitigated by expanding the capabilities of the deep-snow software to automatically select an appropriate CNN variant. One potential challenge is pervasive cloud cover, which eliminates Sentinel-2 as a useful input feature. If no cloud-free Sentinel-2 data are available near the target date for a given location,

the deep-snow software could use a CNN variant trained without Sentinel-2 input features. While we expect that a CNN variant without Sentinel-2 input will exhibit reduced performance (Table D.2), it likely outperforms most alternatives and may still be useful for many applications. Similarly, outside of the Western U.S., the deep-snow software could use a CNN variant that does not include SNODAS snow depth input, since all other input data products have near-global coverage. In addition to degraded performance without SNODAS input (Table D.2), we expect less reliable predictions outside of the Western U.S. using the current generation of CNNs due to challenges with generalization. A promising strategy to address this would be to fine-tune the CNNs using high-quality snow depth products from other regions. The fine-tuned CNNs may outperform models trained solely on smaller, region-specific datasets.

5.6.4 Future work

In this work, we explored the potential of a standard UNet, an established CNN architecture in image processing, to predict snow depth over the Western U.S.. This architecture offers the advantage of leveraging spatial autocorrelation in snow depth by incorporating information from a local neighborhood of grid cells when making predictions. Previous work has demonstrated the potential of model architectures that exploit temporal autocorrelation in snow depth and SWE, using information from earlier in the water year to inform predictions on a given date (e.g., snow depth: Liu et al., 2024; SWE: Cui et al., 2023; Pflug et al., 2025; Saranathan et al., 2025). These temporally-aware models are well-suited for incorporating meteorological input features, enabling them to model snow depth evolution over time based on physical processes. A model architecture that propagates information across both space and time to generate distributed snow depth predictions will likely outperform models that rely solely on information across space or time.

While snow depth predictions from machine learning models can be improved by refining model architecture, an equally important opportunity for improvement lies in expanding the spatial and temporal coverage of the training data. Most training data used in this study were collected over the California Sierra Nevada and the Colorado Rocky Mountains, with limited airborne lidar acquisitions available from other regions of the Western U.S.. There are currently few airborne lidar snow depth measurements from early in the water year (October to December) or late in the water year (August to September), which contributes to poor CNN performance during the early season. Increasing the number of airborne lidar acquisitions outside of California and Colorado, and in October, November, and December, would likely lead to substantial improvements in CNN

performance across the full year and throughout the Western U.S. Although these acquisitions would be less useful for operational water resource management needs, they would provide essential coverage for a representative training dataset.

Finally, we demonstrated the capability of machine learning methods to fuse an arbitrary number of input features to make snow depth predictions. Future work should explore performance improvements from including additional data sources, including for example meteorological data, snow reanalysis data, in-situ data, and additional satellite data (e.g., ICESat-2). As new techniques to measure and predict snowpack properties are developed (e.g., those based on L-band InSAR from the NISAR mission), the resulting products can be used as input features for updated machine learning models. This adaptability makes machine learning an attractive option for combining evolving data sources to produce improved predictions of snowpack properties.

5.7 Summary

We trained a UNet CNN to predict snow depth over the Western U.S. using input data from SNODAS, Sentinel-2, Sentinel-1, and the COP30 DEM. The CNN was trained on 233 quality-controlled airborne lidar snow depth products, covering ~422,000 km^2, similar to the total area of California, and containing ~271,000 km^2 of snow-covered area, similar to the area of Colorado. We performed extensive feature importance analyses to inform our choice of input features and found that including SNODAS snow depth strongly improved CNN performance, while including Sentinel-1 change in cross ratio led to a small improvement in performance. CNN evaluation on our test dataset showed that our approach substantially outperforms previous methods for generating snow depth maps over the Western U.S. at <100 m resolution. Performance was similar to or better than that of comparable machine learning approaches, and unlike many other approaches, our CNN does not rely on in-situ snow depth from nearby monitoring stations.

We developed an open-source Python package allowing users to apply our CNN to generate snow depth predictions. Using this package, we produced a snow depth map for the entire Western U.S. on March 13, 2024, and created eight-year snow depth time series at two representative snow monitoring stations in the test dataset.

CNN snow depth predictions may be particularly valuable in ungauged basins and late in the water year, when the snowline rises above the elevation of most monitoring stations. In practice, we recommend applying the CNN as an ensemble and bias-correcting snow depth predictions using either the test dataset or in-situ observations. This work demonstrates the potential of

machine learning to produce high-resolution snow depth maps over large areas like the Western U.S. Future work should explore model architectures that leverage both spatial and temporal autocorrelation to improve snow depth predictions. Expanding the archive of high-quality distributed snow depth datasets to cover periods early and late in the water year and regions beyond the Sierra Nevada and Colorado Rockies would also help create a more representative training dataset and improve CNN performance over the Western U.S.

BIBLIOGRAPHY

Adebisi, N., Marshall, H., Vuyovich, C. M., Elder, K., Hiemstra, C., & Durand, M. (2022). *SnowEx20–21 QSI lidar snow depth 0.5 m UTM grid (SNEX20_QSI_SD, Version 1)* [Data set]. NASA National Snow and Ice Data Center Distributed Active Archive Center. https://doi.org/10.5067/VBUN16K365DG

Ahmad, M., Mazzara, M., & Distefano, S. (2021). Regularized CNN feature hierarchy for hyperspectral image classification. *Remote Sensing, 13*, 2275. https://doi.org/10.3390/rs13122275

Alabi, I. O., Marshall, H. P., Mead, J., & Trujillo, E. (2025). Advancing terrestrial snow depth monitoring with machine learning and L-band InSAR data: A case study using NASA's SnowEx 2017 data. *Frontiers in Remote Sensing, 5*, 1481848. https://doi.org/10.3389/frsen.2024.1481848

Ali Badal, L. (2023, October 6). Indian soldier's body found in Teesta handed over to BSF. *Dhaka Tribune.* https://www.dhakatribune.com/bangladesh/nation/327225/indian-soldier%E2%80%99s-body-found-in-teesta-handed-over

Allen, S. K., Linsbauer, A., Randhawa, S. S., Huggel, C., Rana, P., & Kumari, A. (2016). Glacial lake outburst flood risk in Himachal Pradesh, India: An integrative and anticipatory approach considering current and future threats. *Natural Hazards, 84*(3), 1741–1763. https://doi.org/10.1007/s11069-016-2511-x

Altena, B., Scambos, T., Fahnestock, M., & Kääb, A. (2019). Extracting recent short-term glacier velocity evolution over southern Alaska and the Yukon from a large collection of Landsat data. *The Cryosphere, 13*(3), 795–814. https://doi.org/10.5194/tc-13-795-2019

Anantrasirichai, N., Biggs, J., Albino, F., & Bull, D. (2019). Application of convolutional neural networks to detect slow, sustained deformation in InSAR time series. *Geophysical Research Letters, 46*, 11850–11858. https://doi.org/10.1029/2019GL084993

Anantrasirichai, N., Biggs, J., Albino, F., Hill, P., & Bull, D. (2018). Application of machine learning to classification of volcanic deformation in routinely generated InSAR data. *Journal of Geophysical Research: Solid Earth, 123*, 6592–6606. https://doi.org/10.1029/2018JB015911

Anantrasirichai, N., Biggs, J., Kelevitz, K., Sadeghi, Z., Wright, T., Thompson, J., Achim, A. M., & Bull, D. (2021). Detecting ground deformation in the built environment using sparse satellite InSAR data with a convolutional neural network. *IEEE Transactions on Geoscience and Remote Sensing, 59*, 2940–2950. https://doi.org/10.1109/TGRS.2020.3018315

Anderson, R. S., & Anderson, S. P. (2025). Lingering beneath crumbling walls: The origin of Holocene rock glaciers. *Geology, 53*(8), 663–667. https://doi.org/10.1130/G53180.1

Avanzi, F., De Michele, C., & Ghezzi, A. (2015). On the performances of empirical regressions for the estimation of bulk snow density. *Geografia Fisica e Dinamica Quaternaria, 38*(2), 105–112. https://doi.org/10.4461/GFDQ.2015.38.10

Bair, E. H., Rittger, K., Davis, R. E., Painter, T. H., & Dozier, J. (2016). Validating reconstruction of snow water equivalent in California's Sierra Nevada using measurements from the NASA Airborne Snow Observatory. *Water Resources Research, 52*(11), 8437–8460. https://doi.org/10.1002/2016WR018704

Bales, R. C., Molotch, N. P., Painter, T. H., Dettinger, M. D., Rice, R., & Dozier, J. (2006). Mountain hydrology of the western United States. *Water Resources Research, 42*(8). https://doi.org/10.1029/2005WR004387

Barnett, T. P., Adam, J. C., & Lettenmaier, D. P. (2005). Potential impacts of a warming climate on water availability in snow-dominated regions. *Nature, 438*, 303–309. https://doi.org/10.1038/nature04141

Bartholomew, I., Nienow, P., Mair, D., Hubbard, A., King, M. A., & Sole, A. (2010). Seasonal evolution of subglacial drainage and acceleration in a Greenland outlet glacier. *Nature Geoscience, 3*(6), 408–411. https://doi.org/10.1038/ngeo863

Bayer, B., Simoni, A., Mulas, M., Corsini, A., & Schmidt, D. (2018). Deformation responses of slow-moving landslides to seasonal rainfall in the Northern Apennines, measured by InSAR. *Geomorphology, 308*, 293–306. https://doi.org/10.1016/j.geomorph.2018.02.020

Beauducel, F., Briole, P., & Froger, J.-L. (2000). Volcano-wide fringes in ERS synthetic aperture radar interferograms of Etna (1992–1998): Deformation or tropospheric effect? *Journal of Geophysical Research: Solid Earth, 105*, 16391–16402. https://doi.org/10.1029/2000JB900095

Bekaert, D. P. S., Hooper, A., & Wright, T. J. (2015a). A spatially variable power-law tropospheric correction technique for InSAR data. *Journal of Geophysical Research: Solid Earth, 120*, 1345–1356. https://doi.org/10.1002/2014JB011558

Bekaert, D. P. S., Walters, R. J., Wright, T. J., Hooper, A. J., & Parker, D. J. (2015b). Statistical comparison of InSAR tropospheric correction techniques. *Remote Sensing of Environment, 170*, 40–47. https://doi.org/10.1016/j.rse.2015.08.035

Beniston, M. (2003). Climatic change in mountain regions: A review of possible impacts. *Climatic Change, 59*(1), 5–31. https://doi.org/10.1023/A:1024458411589

Beniston, M., Stoffel, M., Clarvis, M. H., & Quevauviller, P. (2014). Assessing climate change impacts on the quantity of water in Alpine regions: Foreword to the adaptation and policy implications of the EU/FP7 ACQWA project. *Environmental Science & Policy, 43*, 1–4. https://doi.org/10.1016/j.envsci.2014.01.009

Berardino, P., Fornaro, G., Lanari, R., & Sansosti, E. (2002). A new algorithm for surface deformation monitoring based on small baseline differential SAR interferograms. *IEEE Transactions on Geoscience and Remote Sensing, 40*(11), 2375–2383. https://doi.org/10.1109/TGRS.2002.803792

Bertone, A., Barboux, C., Bodin, X., Bolch, T., Brardinoni, F., Caduff, R., Christiansen, H. H., Darrow, M. M., Delaloye, R., Etzelmüller, B., Humlum, O., Lambiel, C., Lilleøren, K. S., Mair, V., Pellegrinon, G., Rouyet, L., Ruiz, L., & Strozzi, T. (2022). Incorporating InSAR kinematics into rock glacier inventories: Insights from 11 regions worldwide. *The Cryosphere, 16*(7), 2769–2792. https://doi.org/10.5194/tc-16-2769-2022

Betato, A., Rodríguez, H. D., French, N., James, T., & Remeseiro, B. (2025). MAPunet: High-resolution snow depth mapping through U-Net pixel-wise regression. *Remote Sensing Applications: Society and Environment, 101477.* https://doi.org/10.1016/j.rsase.2025.101477

Beyer, R. A., Alexandrov, O., & McMichael, S. (2018). The Ames Stereo Pipeline: NASA's open-source software for deriving and processing terrain data. *Earth and Space Science, 5*(9), 537–548. https://doi.org/10.1029/2018EA000409

Bhushan, S., & Shean, D. (2025). *uw-cryo/debris_cover_smb (Version 0.4.1)* [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.14816790

Bhushan, S., & Shean, D. (2021). *Chamoli disaster pre-event 2-m DEM composite: September 2015 (Version 1.0)* [Data set]. Zenodo. https://doi.org/10.5281/zenodo.4554647

Bhushan, S., Shean, D., Hu, J. Y. M., Guillet, G., & Rounce, D. R. (2024). Deriving seasonal and annual surface mass balance for debris-covered glaciers from flow-corrected satellite stereo DEM time series. *Journal of Glaciology, 70*, e6. https://doi.org/10.1017/jog.2024.57

Bolch, T., Peters, J., Yegorov, A., Pradhan, B., Buchroithner, M., & Blagoveshchensky, V. (2012). Identification of potentially dangerous glacial lakes in the northern Tian Shan. In *Terrigenous mass movements: Detection, modelling, early warning and mitigation using geoinformation technology* (pp. 369–398). https://doi.org/10.1007/978-3-642-25495-6_12

Bontemps, N., Lacroix, P., & Doin, M.-P. (2018). Inversion of deformation-field time series from optical images and application to the long-term kinematics of slow-moving landslides in Peru. *Remote Sensing of Environment, 210*, 144–158. https://doi.org/10.1016/j.rse.2018.02.023

Bormann, K. J., Brown, R. D., Derksen, C., & Painter, T. H. (2018). Estimating snow-cover trends from space. *Nature Climate Change, 8*(11), 924–928. https://doi.org/10.1038/s41558-018-0318-3

Brangers, I., Marshall, H. P., De Lannoy, G., Dunmire, D., Mätzler, C., & Lievens, H. (2024). Tower-based C-band radar measurements of an alpine snowpack. *The Cryosphere, 18*(7), 3177–3193. https://doi.org/10.5194/tc-18-3177-2024

Brencher, G., Handwerger, A. L., & Munroe, J. S. (2021). InSAR-based characterization of rock glacier movement in the Uinta Mountains, Utah, USA. *The Cryosphere, 15*(10), 4823–4844. https://doi.org/10.5194/tc-15-4823-2021

Brengman, C. M. J., & Barnhart, W. D. (2021). Identification of surface deformation in InSAR using machine learning. *Geochemistry, Geophysics, Geosystems, 22*, e2020GC009204. https://doi.org/10.1029/2020GC009204

Brouwer, W. S., & Hanssen, R. F. (2021). An analysis of InSAR displacement vector decomposition fallacies and the strap-down solution. In *2021 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)* (pp. 2927–2930). https://doi.org/10.1109/IGARSS47720.2021.9554216

Broxton, P., Ehsani, M. R., & Behrangi, A. (2024). Improving mountain snowpack estimation using machine learning with Sentinel-1, the Airborne Snow Observatory, and University of Arizona snowpack data. *Earth and Space Science, 11*(3), e2023EA002964. https://doi.org/10.1029/2023EA002964

Buchhorn, M., Lesiv, M., Tsendbazar, N. E., Herold, M., Bertels, L., & Smets, B. (2020). Copernicus global land cover layers—Collection 2. *Remote Sensing, 12*(6), 1044. https://doi.org/10.3390/rs12061044

Budhathoki, K. P., Bajracharya, O., & Pokharel, B. (2010). Assessment of Imja Glacier lake outburst flood risk in Dudh Koshi River basin using remote sensing techniques. *Journal of Hydrology and Meteorology, 7*(1), 75–91. https://doi.org/10.3126/jhm.v7i1.5618

Burakowski, E., & Magnusson, M. (2012). *Climate impacts on the winter tourism economy in the United States.* Natural Resources Defense Council. http://www.nrdc.org/globalwarming/climate-impacts-winter-tourism.asp

Bürgmann, R., Rosen, P. A., & Fielding, E. J. (2000). Synthetic aperture radar interferometry to measure Earth's surface topography and its deformation. *Annual Review of Earth and Planetary Sciences, 28*(1), 169–209. https://doi.org/10.1146/annurev.earth.28.1.169

Bursik, M. (1991). Relative dating of moraines based on landform degradation, Lee Vining Canyon, California. *Quaternary Research, 35*(3), 451–455. https://doi.org/10.1016/0033-5894(91)90057-C

Cabré, A., Remy, D., Aguilar, G., Carretier, S., & Riquelme, R. (2020). Mapping rainstorm erosion associated with an individual storm from InSAR coherence loss validated by field evidence for the Atacama Desert. *Earth Surface Processes and Landforms, 45*(9), 2091–2106. https://doi.org/10.1002/esp.4868

Carrivick, J. L., & Tweed, F. S. (2016). A global assessment of the societal impacts of glacier outburst floods. *Global and Planetary Change, 144*, 1–16. https://doi.org/10.1016/j.gloplacha.2016.07.001

Casu, F., Manconi, A., Pepe, A., & Lanari, R. (2011). Deformation time-series generation in areas characterized by large displacement dynamics: The SAR amplitude pixel-offset SBAS technique. *IEEE Transactions on Geoscience and Remote Sensing, 49*(7), 2752–2763. https://doi.org/10.1109/TGRS.2010.2104325

Cayan, D. R. (1996). Interannual climate variability and snowpack in the western United States. *Journal of Climate, 9*(5), 928–948. https://doi.org/10.1175/1520-0442(1996)009%3C0928:ICVASI%3E2.0.CO;2

Chao, W. A., Wu, T. R., Ma, K. F., Kuo, Y. T., Wu, Y. M., Zhao, L., … & Tsai, Y. L. (2018). The large Greenland landslide of 2017: Was a tsunami warning possible? *Seismological Research Letters, 89*(4), 1335–1344. https://doi.org/10.1785/0220170160

Chen, C., Dai, K., Tang, X., Cheng, J., Pirasteh, S., Wu, M., Shi, X., Zhou, H., & Li, Z. (2022). Removing InSAR topography-dependent atmospheric effects based on deep learning. *Remote Sensing, 14*, 4171. https://doi.org/10.3390/rs14174171

Chen, Y., Bruzzone, L., Jiang, L., & Sun, Q. (2020). ARU-Net: Reduction of atmospheric phase screen in SAR interferometry using an attention-based deep residual U-Net. *IEEE Transactions on Geoscience and Remote Sensing, 59*, 5780–5793. https://doi.org/10.1109/TGRS.2020.3021765

Chini, M., Pelich, R., Pulvirenti, L., Pierdicca, N., Hostache, R., & Matgen, P. (2019). Sentinel-1 InSAR coherence to detect floodwater in urban areas: Houston and Hurricane Harvey as a test case. *Remote Sensing, 11*(2), 107. https://doi.org/10.3390/rs11020107

Choudhury, S. N., & Hussain, Z. (2023, October 5). Floods in India's Sikkim state leave 10 dead, 82 missing. *Reuters.* https://www.reuters.com/world/india/nearly-two-dozen-indian-troops-missing-after-flash-flood-ani-2023-10-04/

Cicoira, A., Beutel, J., Faillettaz, J., & Vieli, A. (2019). Water controls the seasonal rhythm of rock glacier flow. *Earth and Planetary Science Letters, 528*, 115844. https://doi.org/10.1016/j.epsl.2019.115844

Clague, J. J., & Evans, S. G. (2000). A review of catastrophic drainage of moraine-dammed lakes in British Columbia. *Quaternary Science Reviews, 19*(17–18), 1763–1783. https://doi.org/10.1016/S0277-3791(00)00090-1

Colesanti, C., & Wasowski, J. (2006). Investigating landslides with space-borne synthetic aperture radar interferometry. *Engineering Geology, 88*, 173–199. https://doi.org/10.1016/j.enggeo.2006.09.013

Colman, S. M., & Pierce, K. L. (1986). Glacial sequence near McCall, Idaho: Weathering rinds, soil development, morphology, and other relative-age criteria. *Quaternary Research, 25*(1), 25–42. https://doi.org/10.1016/0033-5894(86)90041-4

Costa, J. E., & Schuster, R. L. (1988). The formation and failure of natural dams. *GSA Bulletin, 100*(7), 1054–1068. https://doi.org/10.1130/0016-7606(1988)100<1054:TFAFON>2.3.CO;2

Cuffey, K. M., & Paterson, W. S. B. (2010). *The physics of glaciers* (4th ed.). Butterworth-Heinemann.

Cui, G., Anderson, M., & Bales, R. (2023). Mapping of snow water equivalent by a deep-learning model assimilating snow observations. Journal of Hydrology, 616, 128835. https://doi.org/10.1016/j.jhydrol.2022.128835

Currier, W. R., Pflug, J., Mazzotti, G., Jonas, T., Deems, J. S., Bormann, K. J., … & Lundquist, J. D. (2019). Comparing aerial lidar observations with terrestrial lidar and snow-probe transects from NASA's 2017 SnowEx campaign. Water Resources Research, 55(7), 6285–6294. https://doi.org/10.1029/2018WR024533

Dahal, P., Paudyal, K., & Rajaure, S. (2018). Geophysical study on the moraine dam of Imja Glacial Lake in eastern Nepal using the electrical resistivity tomography method. *Journal of Nepal Geological Society, 55*, 15–22. https://doi.org/10.3126/jngs.v55i1.22784

Deems, J. S., Painter, T. H., & Finnegan, D. C. (2013). Lidar measurement of snow depth: A review. Journal of Glaciology, 59(215), 467–479. https://doi.org/10.3189/2013JoG12J154

Dehecq, A., Gourmelen, N., Gardner, A. S., Brun, F., Goldberg, D., Nienow, P. W., Berthier, E., Vincent, C., Wagnon, P., & Trouvé, E. (2019). Twenty-first century glacier slowdown driven by mass loss in High Mountain Asia. *Nature Geoscience, 12*, 22–27. https://doi.org/10.1038/s41561-018-0271-9

Delacourt, C., Briole, P., & Achache, J. A. (1998). Tropospheric corrections of SAR interferograms with strong topography: Application to Etna. *Geophysical Research Letters, 25*, 2849–2852. https://doi.org/10.1029/98GL02112

Delaloye, R., Lambiel, C., & Gärtner-Roer, I. (2010). Overview of rock glacier kinematics research in the Swiss Alps. *Geographica Helvetica, 65*, 135–145. https://doi.org/10.5194/gh-65-135-2010

Deschamps-Berger, C., Gascoin, S., Berthier, E., Deems, J., Gutmann, E., Dehecq, A., … & Dumont, M. (2020). Snow depth mapping from stereo satellite imagery in mountainous terrain: Evaluation using airborne lidar data. The Cryosphere, 14, 2925–2940. https://doi.org/10.5194/tc-14-2925-2020

Ding, X., Li, Z., Zhu, J., Feng, G., & Long, J. (2008). Atmospheric effects on InSAR measurements and their mitigation. *Sensors, 8*, 5426–5446. https://doi.org/10.3390/s8095426

Doin, M.-P., Lasserre, C., Peltzer, G., Cavalié, O., & Doubre, C. (2009). Corrections of stratified tropospheric delays in SAR interferometry: Validation with global atmospheric models. *Journal of Applied Geophysics, 69*, 35–50. https://doi.org/10.1016/j.jappgeo.2009.03.010

Dozier, J., Bair, E. H., & Davis, R. E. (2016). Estimating the spatial distribution of snow water equivalent in the world's mountains. Wiley Interdisciplinary Reviews: Water, 3(3), 461–474. https://doi.org/10.1002/wat2.1140

Dramiński, M., Rada-Iglesias, A., Enroth, S., Wadelius, C., Koronacki, J., & Komorowski, J. (2008). Monte Carlo feature selection for supervised classification. Bioinformatics, 24(1), 110–117. https://doi.org/10.1093/bioinformatics/btm486

Driscoll, F. G. (1980). Wastage of the Klutlan ice-cored moraines, Yukon Territory, Canada. *Quaternary Research, 14*(1), 31–49. https://doi.org/10.1016/0033-5894(80)90005-8

Dubey, S., & Goyal, M. K. (2020). Glacial lake outburst flood hazard, downstream impact, and risk over the Indian Himalayas. Water Resources Research, 56(4), e2019WR026533. https://doi.org/10.1029/2019WR026533

Dunmire, D., Lievens, H., Boeykens, L., & De Lannoy, G. J. (2024). A machine learning approach for estimating snow depth across the European Alps from Sentinel-1 imagery. Remote Sensing of Environment, 314, 114369. https://doi.org/10.1016/j.rse.2024.114369

Eineder, M. (2004). Problems and solutions for InSAR digital elevation model generation of mountainous terrain. In *Proceedings of the Fringe 2003 Workshop* (Vol. 550, p. 18). ESA.

Emardson, T. R., Simons, M., & Webb, F. H. (2003). Neutral atmospheric delay in interferometric synthetic aperture radar applications: Statistical description and mitigation. *Journal of Geophysical Research: Solid Earth, 108*. https://doi.org/10.1029/2002JB001781

Emmer, A., & Cochachin, A. (2013). The causes and mechanisms of moraine-dammed lake failures in the Cordillera Blanca, North American Cordillera, and Himalayas. *AUC Geographica, 48*(2), 5–15. https://doi.org/10.14712/23361980.2014.23

Emmer, A., Vilímek, V., & Zapata, M. L. (2018). Hazard mitigation of glacial lake outburst floods in the Cordillera Blanca (Peru): The effectiveness of remedial works. *Journal of Flood Risk Management, 11*(S1), S489–S501. https://doi.org/10.1111/jfr3.12241

Etzelmüller, B. (2000). Quantification of thermo-erosion in pro-glacial areas: Examples from Svalbard. *Zeitschrift für Geomorphologie, 44*, 343–361. https://doi.org/10.1127/zfg/44/2000/343

Euillades, L. D., Euillades, P. A., Riveros, N. C., Masiokas, M. H., Ruiz, L., Pitte, P., … & Balbarani, S. (2016). Detection of glacier displacement time series using SAR. *Remote Sensing of Environment, 184*, 188–198. https://doi.org/10.1016/j.rse.2016.07.003

European Space Agency. (2021a). *Copernicus Sentinel-2 (processed by ESA), MSI Level-2A BOA Reflectance Product, Collection 1*. https://doi.org/10.5270/S2_-znk9xsj

European Space Agency. (2021b). *Copernicus Digital Elevation Model (DEM)* [Computer software]. https://registry.opendata.aws/copernicus-dem/

European Space Agency. (2022). *Copernicus DEM product handbook (Version 4.0).* https://spacedata.copernicus.eu/documents/20123/121239/GEO1988-CopernicusDEM-SPE-002_ProductHandbook_I4.0.pdf

Fahnestock, M., Bindschadler, R., Kwok, R., & Jezek, K. (1993). Greenland Ice Sheet surface properties and ice dynamics from ERS-1 SAR imagery. *Science, 262*(5139), 1530–1534. https://doi.org/10.1126/science.262.5139.1530

Fahnestock, M., Scambos, T., Moon, T., Gardner, A., Haran, T., & Klinger, M. (2016). Rapid large-area mapping of ice flow using Landsat 8. *Remote Sensing of Environment, 185*, 84–94. https://doi.org/10.1016/j.rse.2015.11.023

Fang, Y., Liu, Y., & Margulis, S. A. (2022). A western United States snow reanalysis dataset over the Landsat era from water years 1985 to 2021. Scientific Data, 9, 677. https://doi.org/10.1038/s41597-022-01768-7

Ferretti, A., Prati, C., & Rocca, F. (1999). Multibaseline InSAR DEM reconstruction: The wavelet approach. *IEEE Transactions on Geoscience and Remote Sensing, 37*, 705–715. https://doi.org/10.1109/36.752187

Foster, J., Brooks, B., Cherubini, T., Shacat, C., Businger, S., & Werner, C. L. (2006). Mitigating atmospheric noise for InSAR using a high-resolution weather model. *Geophysical Research Letters, 33*. https://doi.org/10.1029/2006GL026781

Fuhrmann, T., & Garthwaite, M. C. (2019). Resolving three-dimensional surface motion with InSAR: Constraints from multi-geometry data fusion. *Remote Sensing, 11*(3), Article 241. https://doi.org/10.3390/rs11030241

Fujita, K., Sakai, A., Nuimura, T., Yamaguchi, S., & Sharma, R. R. (2009). Recent changes in Imja Glacial Lake and its damming moraine in the Nepal Himalaya revealed by in situ surveys and multi-temporal ASTER imagery. *Environmental Research Letters, 4*(4), 045205. https://doi.org/10.1088/1748-9326/4/4/045205

Fujita, K., Sakai, A., Takenaka, S., Nuimura, T., Surazakov, A. B., Sawagaki, T., & Yamanokuchi, T. (2013). Potential flood volume of Himalayan glacial lakes. *Natural Hazards and Earth System Sciences, 13*(7), 1827–1839. https://doi.org/10.5194/nhess-13-1827-2013

Fujita, K., Suzuki, R., Nuimura, T., & Sakai, A. (2008). Performance of ASTER and SRTM DEMs and their potential for assessing glacial lakes in the Lunana region, Bhutan Himalaya. Journal of Glaciology, 54(185), 220–228. https://doi.org/10.3189/002214308784886179

Gagliano, E. (2025). *egagli/easysnowdata (Version 0.0.15)* [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.14741502

Gao, Y., Xiao, L., Chen, D., Xu, J., & Zhang, H. (2018). Comparison between past and future extreme precipitations simulated by global and regional climate models over the Tibetan Plateau. *International Journal of Climatology, 38*, 1285–1297. https://doi.org/10.1002/joc.5243

Gascoin, S. (2024a). A call for an accurate presentation of glaciers as water resources. Wiley Interdisciplinary Reviews: Water, 11(2), e1705. https://doi.org/10.1002/wat2.1705

Gascoin, S., Luojus, K., Nagler, T., Lievens, H., Masiokas, M., Jonas, T., … & De Rosnay, P. (2024b). Remote sensing of mountain snow from space: Status and recommendations. Frontiers in Earth Science, 12, 1381323. https://doi.org/10.3389/feart.2024.1381323

Godbole, V., Dahl, G. E., Gilmer, J., Shallue, C. J., & Nado, Z. (2023). *Deep learning tuning playbook* (Version 1) [Technical report]. https://github.com/google-research/tuning_playbook

Gondara, L. (2016). Medical image denoising using convolutional denoising autoencoders. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)* (pp. 241–246). https://doi.org/10.1109/ICDMW.2016.0041

Gong, T., Zheng, X., & Lu, X. (2021). Cross-domain scene classification by integrating multiple incomplete sources. *IEEE Transactions on Geoscience and Remote Sensing, 59*, 10035–10046. https://doi.org/10.1109/TGRS.2020.3034344

Gong, W., Meyer, F., Webley, P. W., Morton, D., & Liu, S. (2010). Performance analysis of atmospheric correction in InSAR data based on the Weather Research and Forecasting model (WRF). In *2010 IEEE International Geoscience and Remote Sensing Symposium* (pp. 2900–2903). https://doi.org/10.1109/IGARSS.2010.5652267

Gray, L., Joughin, I., Tulaczyk, S., Spikes, V. B., Bindschadler, R., & Jezek, K. (2005). Evidence for subglacial water transport in the West Antarctic Ice Sheet through three-dimensional satellite radar interferometry. *Geophysical Research Letters, 32*(3). https://doi.org/10.1029/2004GL021387

Grebby, S., Sowter, A., Gee, D., Athab, A., Barreda-Bautista, B. D. L., Girindran, R., & Marsh, S. (2021). Remote monitoring of ground motion hazards in high mountain terrain using InSAR: A case study of the Lake Sarez area, Tajikistan. *Applied Sciences, 11*, 8738. https://doi.org/10.3390/app11188738

Grünewald, T., & Lehning, M. (2015). Are flat-field snow depth measurements representative? A comparison of selected index sites with areal snow depth measurements at the small-catchment scale. Hydrological Processes, 29(7), 1717–1728. https://doi.org/10.1002/hyp.10295

Guo, L., Li, J., Li, Z. W., Wu, L. X., Li, X., Hu, J., … & Li, Z. Q. (2020). The surge of the Hispar Glacier, central Karakoram: SAR 3-D flow velocity time series and thickness changes. *Journal of Geophysical Research: Solid Earth, 125*(7), e2019JB018945. https://doi.org/10.1029/2019JB018945

Gupta, V., Rakkasagi, S., Rajpoot, S., Imanni, H. S. E., & Singh, S. (2023). Spatiotemporal analysis of Imja Lake to estimate the downstream flood hazard using the SHIVEK approach. *Acta Geophysica, 71*(5), 2233–2244. https://doi.org/10.1007/s11600-023-01124-2

Hallet, B., & Putkonen, J. (1994). Surface dating of dynamic landforms: Young boulders on aging moraines. *Science, 265*(5174), 937–940. https://doi.org/10.1126/science.265.5174.937

Hambrey, M. J., Quincey, D. J., Glasser, N. F., Reynolds, J. M., Richardson, S. J., & Clemmens, S. (2008). Sedimentological, geomorphological and dynamic context of debris-mantled glaciers, Mount Everest (Sagarmatha) region, Nepal. *Quaternary Science Reviews, 27*(25), 2361–2389. https://doi.org/10.1016/j.quascirev.2008.08.01

Hamm, A., Arndt, A., Kolbe, C., Wang, X., Thies, B., Boyko, O., Reggiani, P., Scherer, D., Bendix, J., & Schneider, C. (2020). Intercomparison of gridded precipitation datasets over a sub-region of the central Himalaya and the southwestern Tibetan Plateau. *Water, 12*, 3271. https://doi.org/10.3390/w12113271

Handwerger, A. L., Fielding, E. J., Huang, M. H., Bennett, G. L., Liang, C., & Schulz, W. H. (2019). Widespread initiation, reactivation, and acceleration of landslides in the Northern California Coast Ranges due to extreme rainfall. *Journal of Geophysical Research: Earth Surface, 124*(7), 1782–1797. https://doi.org/10.1029/2019JF005035

Handwerger, A. L., Roering, J. J., & Schmidt, D. A. (2013). Controls on the seasonal deformation of slow-moving landslides. *Earth and Planetary Science Letters, 377*, 239–247.

Handwerger, A. L., Roering, J. J., Schmidt, D. A., & Rempel, A. W. (2015). Kinematics of earthflows in the Northern California Coast Ranges using satellite interferometry. *Geomorphology, 246*, 321–333. https://doi.org/10.1016/j.geomorph.2015.06.003

Hanisch, J., Ambarish, P., Grabs, W., Dixit, A., & Reynolds, J. (2000). GLOF mitigation strategies—Lessons learned from studying the Thulagi Glacier Lake, Nepal. *Journal of Nepal Geological Society, 22*. https://doi.org/10.3126/jngs.v22i0.32375

Hanssen, R. F. (2001). *Radar interferometry: Data interpretation and error analysis.* Kluwer Academic Publishers.

Hanssen, R. F., & Feijt, A. (1997). A first quantitative evaluation of atmospheric effects on SAR interferometry. In *ERS SAR interferometry: Proceedings of the Fringe 96 Workshop, Zurich, Switzerland, 30 September–2 October 1996* (Vol. 406, p. 277).

Hanssen, R. F., Weckwerth, T. M., Zebker, H. A., & Klees, R. (1999). High-resolution water vapor mapping from interferometric radar measurements. *Science, 283*(5406), 1297–1299. https://doi.org/10.1126/science.283.5406.1297

Haritashya, U. K., Kargel, J. S., Shugar, D. H., Leonard, G. J., Strattman, K., Watson, C. S., … & Regmi, D. (2018). Evolution and controls of large glacial lakes in the Nepal Himalaya. *Remote Sensing, 10*(5), 798.

Harpold, A. A., Guo, Q., Molotch, N., Brooks, P. D., Bales, R., Fernandez-Diaz, J. C., … & Lucas, R. (2014). LiDAR-derived snowpack data sets from mixed conifer forests across the western United States. Water Resources Research, 50(3), 2749–2755. https://doi.org/10.1002/2013WR013935

Healy, T. R. (1975). Thermokarst—A mechanism of de-icing ice-cored moraines. *Boreas, 4*(1), 19–23. https://doi.org/10.1111/j.1502-3885.1975.tb00676.x

Herbert, J., Raleigh, M. S., & Small, E. E. (2025). Using a random forest model to combine airborne lidar and SNOTEL data for daily estimates of snow depth across mountain drainage basins of Colorado. *Authorea Preprints.*

Hill, D. F., Burakowski, E. A., Crumley, R. L., Keon, J., Hu, J. M., Arendt, A. A., … & Wolken, G. J. (2019). Converting snow depth to snow water equivalent using climatological variables. The Cryosphere, 13(7), 1767–1784. https://doi.org/10.5194/tc-13-1767-2019

Hock, R., Bliss, A., Marzeion, B., Giesen, R. H., Hirabayashi, Y., Huss, M., Radić, V., & Slangen, A. B. A. (2019). GlacierMIP – A model intercomparison of global-scale glacier mass-balance models and projections. *Journal of Glaciology, 65*(251), 453–467. https://doi.org/10.1017/jog.2019.22

Hogenson, K., Kristenson, H., Kennedy, J., Johnston, A., Rine, J., Logan, T., Zhu, J., Williams, F., Herrmann, J., Smale, J., & Meyer, F. (2023). *Hybrid Pluggable Processing Pipeline (HyP3): A cloud-native infrastructure for generic processing of SAR data* [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.7644424

Höhle, J., & Höhle, M. (2009). Accuracy assessment of digital elevation models by means of robust statistical methods. ISPRS Journal of Photogrammetry and Remote Sensing, 64(4), 398–406. https://doi.org/10.1016/j.isprsjprs.2009.02.003

Hoppinen, Z., Palomaki, R. T., Brencher, G., Dunmire, D., Gagliano, E., Marziliano, A., … & Marshall, H. P. (2024). Evaluating snow depth retrievals from Sentinel-1 volume scattering over NASA SnowEx sites. *The Cryosphere, 18*(11), 5407–5430.

Horn, B. K. (1981). Hill shading and the reflectance map. *Proceedings of the IEEE, 69*(1), 14–47.

Hu, J. M., Shean, D., & Bhushan, S. (2023). Six consecutive seasons of high-resolution mountain snow depth maps from satellite stereo imagery. *Geophysical Research Letters, 50*(24), e2023GL104871.

Hu, Y., Che, T., Dai, L., & Xiao, L. (2021). Snow depth fusion based on machine learning methods for the Northern Hemisphere. Remote Sensing, 13(7), 1250. https://doi.org/10.3390/rs13071250

Huggel, C., Cochachin, A., Drenkhan, F., Fluixá-Sanmartín, J., Frey, H., García Hernández, J., Jurt, C., Muñoz, R., Price, K., & Vicuña, L. (2020). Glacier Lake 513, Peru: Lessons for early warning service development. *WMO Bulletin, 69(1),* 45–52. https://doi.org/10.5167/UZH-186748

Huggel, C., Haeberli, W., Kääb, A., Bieri, D., & Richardson, S. (2004). An assessment procedure for glacial hazards in the Swiss Alps. Canadian Geotechnical Journal, 41(6), 1068–1083. https://doi.org/10.1139/t04-053

Huggel, C., Salzmann, N., Allen, S., Caplan-Auerbach, J., Fischer, L., Haeberli, W., … & Wessels, R. (2010). Recent and future warm extreme events and high-mountain slope stability. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 368(1919), 2435–2459. https://doi.org/10.1098/rsta.2010.0078

Hugonnet, R., Brun, F., Berthier, E., Dehecq, A., Mannerfelt, E. S., Eckert, N., & Farinotti, D. (2022). Uncertainty analysis of digital elevation models by spatial inference from stable terrain. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 15*, 6456–6472.

Huss, M., Bookhagen, B., Huggel, C., Jacobsen, D., Bradley, R. S., Clague, J. J., Vuille, M., Buytaert, W., Cayan, D. R., Greenwood, G., Mark, B. G., Milner, A. M., Weingartner, R., & Winder, M. (2017). Toward mountains without permanent snow and ice. *Earth's Future, 5*, 418–435. https://doi.org/10.1002/2016EF000514

Ikeda, A., Matsuoka, N., & Kääb, A. (2008). Fast deformation of perennially frozen debris in a warm rock glacier in the Swiss Alps: An effect of liquid water. *Journal of Geophysical Research: Earth Surface, 113*(F1). https://doi.org/10.1029/2007JF000859

Ilesanmi, A. E., & Ilesanmi, T. O. (2021). Methods for image denoising using convolutional neural network: A review. *Complex & Intelligent Systems, 7*, 2179–2198. https://doi.org/10.1007/s40747-021-00428-4

Innamorati, C., Ritschel, T., Weyrich, T., & Mitra, N. J. (2018). Learning on the edge: Explicit boundary handling in CNNs. *arXiv*. https://arxiv.org/abs/1805.03106

International Centre for Integrated Mountain Development. (2011). *Glacial lakes and glacial lake outburst floods in Nepal.* ICIMOD.

Irvine-Fynn, T. D. L., Barrand, N. E., Porter, P. R., Hodson, A. J., & Murray, T. (2011). Recent High-Arctic glacial sediment redistribution: A process perspective using airborne lidar. *Geomorphology, 125*(1), 27–39.

Itoh, K. (1982). Analysis of the phase unwrapping algorithm. *Applied Optics, 21*(14), 2470–2470. https://doi.org/10.1364/AO.21.002470

Iverson, N. R. (2010). Shear resistance and continuity of subglacial till: Hydrology rules. *Journal of Glaciology, 56*(200), 1104–1114. https://doi.org/10.3189/002214311796406220

Jacquemart, M., & Tiampo, K. (2021). Leveraging time series analysis of radar coherence and normalized difference vegetation index ratios to characterize pre-failure activity of the Mud Creek landslide, California. *Natural Hazards and Earth System Sciences, 21*(2), 629–642. https://doi.org/10.5194/nhess-21-629-2021

Jain, S., Mishra, S. K., Salunke, P., & Sahany, S. (2019). Importance of the resolution of surface topography vis-à-vis atmospheric and surface processes in the simulation of the climate of Himalaya–Tibet highland. *Climate Dynamics, 52*, 4735–4748. https://doi.org/10.1007/s00382-018-4411-0

Jain, V., & Seung, H. S. (2008). Natural image denoising with convolutional networks. In *Advances in neural information processing systems*. Curran Associates.

Janke, J. R. (2005). Long-term flow measurements (1961–2002) of the Arapaho, Taylor, and Fair Rock Glaciers, Front Range, Colorado. *Physical Geography, 26*, 313–336. https://doi.org/10.2747/0272-3646.26.4.313

Janssen, V., Ge, L., & Rizos, C. (2004). Tropospheric corrections to SAR interferometry from GPS observations. *GPS Solutions, 8*, 140–151. https://doi.org/10.1007/s10291-004-0099-1

Jiang, L., Fan, X., Deng, Y., Zou, C., Feng, Z., Djukem, D. L. W., … & Xu, Q. (2023). Combining geophysics, remote sensing and numerical simulation to assess GLOFs: Case study of the Namulacuo Lake in the southeastern Tibetan Plateau. *Science of the Total Environment, 880*, 163262.

Johnson, G., Chang, H., & Fountain, A. (2021). Active rock glaciers of the contiguous United States: Geographic information system inventory and spatial distribution patterns. *Earth System Science Data, 13*(8), 3979–3994. https://doi.org/10.5194/essd-13-3979-2021

Johnson, P. G. (1971). Ice-cored moraine formation and degradation, Donjek Glacier, Yukon Territory, Canada. *Geografiska Annaler: Series A, Physical Geography, 53*(3–4), 198–202. https://doi.org/10.1080/04353676.1971.11879845

Johnson, P. G. (1992). Stagnant glacier ice, St. Elias Mountains, Yukon. *Geografiska Annaler: Series A, Physical Geography, 74*(1), 13–19. https://doi.org/10.1080/04353676.1992.11880345

Jolivet, R., Agram, P. S., Lin, N. Y., Simons, M., Doin, M.-P., Peltzer, G., & Li, Z. (2014). Improving InSAR geodesy using global atmospheric models. *Journal of Geophysical Research: Solid Earth, 119*, 2324–2341. https://doi.org/10.1002/2013JB010588

Jolivet, R., & Simons, M. (2018). A multipixel time series analysis method accounting for ground motion, atmospheric noise, and orbital errors. *Geophysical Research Letters, 45*(4), 1814–1824. https://doi.org/10.1002/2017GL076533

Jolivet, R., Grandin, R., Lasserre, C., Doin, M.-P., & Peltzer, G. (2011). Systematic InSAR tropospheric phase delay corrections from global meteorological reanalysis data. *Geophysical Research Letters, 38*. https://doi.org/10.1029/2011GL048757

Jolliffe, I. T. (2002). Choosing a subset of principal components or variables. In *Principal component analysis* (pp. 92–114). Springer.

Jones, D. B., Harrison, S., Anderson, K., & Whalley, W. B. (2019). Rock glaciers and mountain hydrology: A review. Earth-Science Reviews, 193, 66–90. https://doi.org/10.1016/j.earscirev.2019.04.001

Joughin, I., Smith, B. E., & Howat, I. M. (2018). A complete map of Greenland ice velocity derived from satellite data collected over 20 years. *Journal of Glaciology, 64*(243), 1–11. https://doi.org/10.1017/jog.2017.73

Joughin, I. (2002). Ice-sheet velocity mapping: A combined interferometric and speckle-tracking approach. *Annals of Glaciology, 34*, 195–201. https://doi.org/10.3189/172756402781817978

Kellerer-Pirklbauer, A., Lieb, G. K., Avian, M., & Carrivick, J. (2012). Climate change and rock-fall events in high mountain areas: Numerous and extensive rock falls in 2007 at Mittlerer Burgstall, central Austria. *Geografiska Annaler: Series A, Physical Geography, 94*(1), 59–78. https://doi.org/10.1111/j.1468-0459.2011.00449.x

Kenner, R., Phillips, M., Beutel, J., Hiller, M., Limpach, P., Pointner, E., & Volken, M. (2017). Factors controlling velocity variations at short-term, seasonal and multiyear time scales, Ritigraben rock

glacier, western Swiss Alps. *Permafrost and Periglacial Processes, 28*(4), 675–684. https://doi.org/10.1002/ppp.1953

Khadra, N. S. (2016, October 31). *Nepal drains dangerous Everest lake.* BBC News. https://www.bbc.com/news/world-asia-37797559

King, L., & Volk, M. (1994). Glaziologie und Glazialmorphologie des Liefde- und Bockfjordgebietes, NW-Spitzbergen. *Zeitschrift für Geomorphologie, 97*, 145–159.

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv*. https://doi.org/10.48550/arXiv.1412.6980

Kinoshita, Y. (2022). Development of InSAR neutral atmospheric delay correction model by use of GNSS ZTD and its horizontal gradient. *IEEE Transactions on Geoscience and Remote Sensing, 60*, 1–14. https://doi.org/10.1109/TGRS.2022.3188988

Kjær, K. H., & Krüger, J. (2001). The final phase of dead-ice moraine development: Processes and sediment architecture, Kötlujökull, Iceland. *Sedimentology, 48*(5), 935–952. https://doi.org/10.1046/j.1365-3091.2001.00402.x

Krüger, J., & Kjær, K. H. (2000). De-icing progression of ice-cored moraines in a humid, subpolar climate, Kötlujökull, Iceland. *The Holocene, 10*(6), 737–747. https://doi.org/10.1191/09596830094980

Kumar, H., & Travelli, A. (2023, October 6). *A calamitous flood shows the dangers lurking in melting glaciers. The New York Times.* https://www.nytimes.com/2023/10/06/world/asia/india-flood-sikkim-climate-change.html

Lacroix, P., Araujo, G., Hollingsworth, J., & Taipe, E. (2019). Self-entrainment motion of a slow-moving landslide inferred from Landsat-8 time series. *Journal of Geophysical Research: Earth Surface, 124*(5), 1201–1216. https://doi.org/10.1029/2018JF004920

Lahmers, T. M., Kumar, S. V., Rosen, D., Dugger, A., Gochis, D. J., Santanello, J. A., … Dunlap, R. (2022). Assimilation of NASA's Airborne Snow Observatory snow measurements for improved hydrological modeling: A case study enabled by the coupled LIS/WRF-Hydro system. *Water Resources Research, 58*(3), e2021WR029867.

Lala, J. M., Rounce, D. R., & McKinney, D. C. (2018). Modeling the glacial lake outburst flood process chain in the Nepal Himalaya: Reassessing Imja Tsho's hazard. *Hydrology and Earth System Sciences, 22*, 3721–3737. https://doi.org/10.5194/hess-22-3721-2018

Lang, N., Jetz, W., Schindler, K., & Wegner, J. D. (2023). A high-resolution canopy height model of the Earth. Nature Ecology & Evolution, 7(11), 1778–1789. https://doi.org/10.1038/s41559-023-02206-6

LeCun, Y., Bottou, L., Orr, G. B., & Müller, K. R. (2002). Efficient backprop. In G. B. Orr & K. R. Müller (Eds.), *Neural networks: Tricks of the trade* (pp. 9–50). Springer. https://doi.org/10.1007/3-540-49430-8_2

Li, D., Wrzesien, M. L., Durand, M., Adam, J., & Lettenmaier, D. P. (2017). How much runoff originates as snow in the western United States, and how will that change in the future? Geophysical Research Letters, 44(12), 6163–6172. https://doi.org/10.1002/2017GL073551

Li, Z. W., Ding, X. L., Huang, C., Zou, Z. R., & Chen, Y. L. (2007). Atmospheric effects on repeat-pass InSAR measurements over Shanghai region. *Journal of Atmospheric and Solar-Terrestrial Physics, 69*, 1344–1356. https://doi.org/10.1016/j.jastp.2007.04.007

Li, Z., Fielding, E. J., Cross, P., & Muller, J.-P. (2006a). Interferometric synthetic aperture radar atmospheric correction: Medium Resolution Imaging Spectrometer and Advanced Synthetic Aperture Radar integration. *Geophysical Research Letters, 33*. https://doi.org/10.1029/2005GL025299

Li, Z., Fielding, E. J., Cross, P., & Muller, J.-P. (2006b). Interferometric synthetic aperture radar atmospheric correction: GPS topography-dependent turbulence model. *Journal of Geophysical Research: Solid Earth, 111*. https://doi.org/10.1029/2005JB003711

Li, Z., Muller, J. P., Cross, P., Albert, P., Fischer, J., & Bennartz, R. (2006c). Assessment of MERIS near-infrared water vapor products to correct ASAR interferometric measurements. *International Journal of Remote Sensing, 27*, 349–365. https://doi.org/10.1080/01431160500307342

Liang, H., Zhang, L., Ding, X., Lu, Z., & Li, X. (2018). Toward mitigating stratified tropospheric delays in multitemporal InSAR: A quadtree-aided joint model. *IEEE Transactions on Geoscience and Remote Sensing, 57*, 291–303. https://doi.org/10.1109/TGRS.2018.2853706

Lievens, H., Brangers, I., Marshall, H. P., Jonas, T., Olefs, M., & De Lannoy, G. (2022). Sentinel-1 snow depth retrieval at sub-kilometer resolution over the European Alps. The Cryosphere, 16(1), 159–177. https://doi.org/10.5194/tc-16-159-2022

Lievens, H., Demuzere, M., Marshall, H. P., Reichle, R. H., Brucker, L., Brangers, I., … De Lannoy, G. J. (2019). Snow depth variability in the Northern Hemisphere mountains observed from space. Nature Communications, 10, 4629. https://doi.org/10.1038/s41467-019-12566-y

Lippl, S., Vijay, S., & Braun, M. (2018). Automatic delineation of debris-covered glaciers using InSAR coherence derived from X-, C- and L-band radar data: A case study of Yazgyl Glacier. *Journal of Glaciology, 64*(247), 811–821. https://doi.org/10.1017/jog.2018.70

Liu, H., Zhao, Z., & Jezek, K. C. (2007). Synergistic fusion of interferometric and speckle-tracking methods for deriving surface velocity from interferometric SAR data. *IEEE Geoscience and Remote Sensing Letters, 4*(1), 102–106. https://doi.org/10.1109/LGRS.2006.885885

Liu, W., Lin, Z., Hu, Y., Tian, A., Yuan, X., & Wickert, J. (2024). Snow depth retrieval using detrended SNR from GNSS-R with bidirectional GRU. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing.*

Liu, Y., Yao, X., Gu, Z., Zhou, Z., Liu, X., Chen, X., & Wei, S. (2022). Automatic recognition of landslides using InSAR images and an improved Mask R-CNN model in the Eastern Tibet Plateau. *Remote Sensing, 14*, 3362. https://doi.org/10.3390/rs14143362

Lohman, R. B., & Simons, M. (2005). Some thoughts on the use of InSAR data to constrain models of surface deformation: Noise structure and data downsampling. *Geochemistry, Geophysics, Geosystems, 6*. https://doi.org/10.1029/2004GC000841

Loibl, D., Bookhagen, B., Valade, S., & Schneider, C. (2019). OSARIS: The "Open Source SAR Investigation System" for automatized parallel InSAR processing of Sentinel-1 time series with emphasis on cryosphere applications. *Frontiers in Earth Science, 7*.

López-Quiroz, P., Doin, M.-P., Tupin, F., Briole, P., & Nicolas, J.-M. (2009). Time series analysis of Mexico City subsidence constrained by radar interferometry. *Journal of Applied Geophysics, 69*(1), 1–15. https://doi.org/10.1016/j.jappgeo.2009.02.006

Loshchilov, I., & Hutter, F. (2017). Decoupled weight decay regularization. *arXiv Preprint*, arXiv:1711.05101. https://doi.org/10.48550/arXiv.1711.05101

Lu, Z., & Meyer, D. J. (2002). Study of high SAR backscattering caused by an increase of soil moisture over a sparsely vegetated area: Implications for characteristics of backscattering. *International Journal of Remote Sensing, 23*(6), 1063–1074. https://doi.org/10.1080/01431160110040035

Lukas, S., Nicholson, L. I., Ross, F. H., & Humlum, O. (2005). Formation, melt-out processes and landscape alteration of high-Arctic ice-cored moraines: Examples from Nordenskiöld Land, central Spitsbergen. *Polar Geography, 29*(3), 157–187. https://doi.org/10.1080/789610198

Lukas, S. (2011). Ice-cored moraines. In V. P. Singh, P. Singh, & U. K. Haritashya (Eds.), *Encyclopedia of snow, ice and glaciers* (pp. 616–619). Springer. https://doi.org/10.1007/978-90-481-2642-2

Maggiori, E., Tarabalka, Y., Charpiat, G., & Alliez, P. (2017). Can semantic labeling methods generalize to any city? The INRIA aerial image labeling benchmark. In *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)* (pp. 3226–3229). https://doi.org/10.1109/IGARSS.2017.8127684

Mankin, J. S., Viviroli, D., Singh, D., Hoekstra, A. Y., & Diffenbaugh, N. S. (2015). The potential for snow to supply human water demand in the present and future. *Environmental Research Letters, 10*(11), 114016. https://doi.org/10.1088/1748-9326/10/11/114016

Marshall, A. M., Abatzoglou, J. T., Rahimi, S., Lettenmaier, D. P., & Hall, A. (2024). California's 2023 snow deluge: Contextualizing an extreme snow year against future climate change. *Proceedings of the National Academy of Sciences, 121*(20), e2320600121.

Massonnet, D., & Feigl, K. L. (1998). Radar interferometry and its application to changes in the Earth's surface. *Reviews of Geophysics, 36*(4), 441–500. https://doi.org/10.1029/97RG03139

Mattson, L. E., & Gardner, J. S. (1991). Mass wasting on valley-side ice-cored moraines, Boundary Glacier, Alberta, Canada. *Geografiska Annaler: Series A, Physical Geography, 73*(3–4), 123–128. https://doi.org/10.1080/04353676.1991.11880337

Mattson, L. E. (1993). Ablation on debris-covered glaciers: An example from the Rakhiot Glacier, Punjab, Himalaya. *IAHS Publication, 218*, 289–296.

McGrath, D., Webb, R., Shean, D., Bonnell, R., Marshall, H. P., Painter, T. H., … Brucker, L. (2019). Spatially extensive ground-penetrating radar snow depth observations during NASA's 2017 SnowEx campaign: Comparison with in situ, airborne, and satellite observations. Water Resources Research, 55(11), 10026–10036. https://doi.org/10.1029/2019WR024907

McKenzie, G. D. (1969). Observations on a collapsing kame terrace in Glacier Bay National Monument, south-eastern Alaska. *Journal of Glaciology, 8*(54), 413–425. https://doi.org/10.3189/S0022143000027003

McKillop, R. J., & Clague, J. J. (2007). Statistical, remote-sensing–based approach for estimating the probability of catastrophic drainage from moraine-dammed lakes in southwestern British Columbia. *Global and Planetary Change, 56*(1–2), 153–171.

Mergili, M., Emmer, A., Juřicová, A., Cochachin, A., Fischer, J. T., Huggel, C., & Pudasaini, S. P. (2018). How well can we simulate complex hydro-geomorphic process chains? The 2012 multi-lake outburst flood in the Santa Cruz Valley (Cordillera Blanca, Perú). *Earth Surface Processes and Landforms, 43*(7), 1373–1389.

Meyes, R., Lu, M., De Puiseau, C. W., & Meisen, T. (2019). Ablation studies in artificial neural networks. *arXiv Preprint*, arXiv:1901.08644.

Microsoft Open Source, McFarland, M., Emanuele, R., Morris, D., & Augspurger, T. (2022). *microsoft/PlanetaryComputer: October 2022 (2022.10.28)* [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.7261897

Millan, R., Mouginot, J., Rabatel, A., & Morlighem, M. (2022). Ice velocity and thickness of the world's glaciers. *Nature Geoscience, 15*(2), 124–129.

Minchew, B. M., & Meyer, C. R. (2020). Dilation of subglacial sediment governs incipient surge motion in glaciers with deformable beds. *Proceedings of the Royal Society A, 476*(2238), 20200033. https://doi.org/10.1098/rspa.2020.0033

Mirus, B. B., Jones, E. S., Baum, R. L., Godt, J. W., Slaughter, S., Crawford, M. M., … McCoy, K. M. (2020). Landslides across the USA: Occurrence, susceptibility, and data limitations. *Landslides, 17*, 2271–2285. https://doi.org/10.1007/s10346-020-01424-4

Mirza, B. N., Small, E. E., & Raleigh, M. S. (2025). Evaluating the utility of Sentinel-1 in a data assimilation system for estimating snow depth in a mountainous basin. *EGUsphere*, 1–31.

Mnih, V. (2013). *Machine learning for aerial image labeling* (Doctoral dissertation, University of Toronto).

Mool, P. K., Maskey, P. R., Koirala, A., Joshi, S. P., Lizong, W., Shrestha, A. B., … Shrestha, R. B. (2011). *Glacial lakes and glacial lake outburst floods in Nepal.* ICIMOD. https://policycommons.net/artifacts/1516291/glacial-lakes-and-glacial-lake-outburst-floods-in-nepal/2192662/

Mountain Research Initiative EDW Working Group. (2015). Elevation-dependent warming in mountain regions of the world. Nature Climate Change, 5, 424–430. https://doi.org/10.1038/nclimate2563

Mukherjee, S., Zimmer, A., Kottayil, N. K., Sun, X., Ghuman, P., & Cheng, I. (2018). CNN-based InSAR denoising and coherence metric. In *2018 IEEE Sensors* (pp. 1–4). https://doi.org/10.1109/ICSENS.2018.8589920

Munroe, J. S., & Handwerger, A. L. (2024). Contemporary movement of rock glaciers in the La Sal and Uinta Mountains, Utah, USA. *Quaternary Science Advances, 14*, 100188. https://doi.org/10.1016/j.qsa.2024.100188

Murray, K. D., Lohman, R. B., & Bekaert, D. P. (2021). Cluster-based empirical tropospheric corrections applied to InSAR time series analysis. *IEEE Transactions on Geoscience and Remote Sensing, 59*, 2204–2212. https://doi.org/10.1109/TGRS.2020.3003271

Murray, K. D., Bekaert, D. P. S., & Lohman, R. B. (2019). Tropospheric corrections for InSAR: Statistical assessments and applications to the Central United States and Mexico. *Remote Sensing of Environment, 232*, 111326. https://doi.org/10.1016/j.rse.2019.111326

National Academies of Sciences, Engineering, and Medicine. (2018). *Thriving on our changing planet: A decadal strategy for Earth observation from space.* National Academies Press. https://doi.org/10.17226/24938

National Operational Hydrologic Remote Sensing Center. (2004). *Snow Data Assimilation System (SNODAS) data products at NSIDC* (G02158, Version 1) [Data set]. National Snow and Ice Data Center. https://doi.org/10.7265/N5TB14TC

Nepali Times. (2025, July 8). *Rasuwa flood likely a GLOF: Deadly Nepal–China border flood could have been glacial lake(s) that burst in Tibet.* https://nepalitimes.com/news/rasuwa-flood-likely-a-glof

Neupane, R., Chen, H., & Cao, C. (2019). Review of moraine dam failure mechanisms. *Geomatics, Natural Hazards and Risk, 10*(1), 1948–1966. https://doi.org/10.1080/19475705.2019.1652210

Nie, Y., Liu, Q., Wang, J., Zhang, Y., Sheng, Y., & Liu, S. (2018). An inventory of historical glacial lake outburst floods in the Himalayas based on remote sensing observations and geomorphological analysis. *Geomorphology, 308*, 91–106. https://doi.org/10.1016/j.geomorph.2018.02.002

Niggli, L., Allen, S., Frey, H., Huggel, C., Petrakov, D., Raimbekova, Z., Reynolds, J., & Wang, W. (2024). GLOF risk management experiences and options: A global overview. In *Oxford Research Encyclopedia of Natural Hazard Science*. Oxford University Press. https://doi.org/10.1093/acrefore/9780199389407.013.540

Nolan, M., Fatland, D. R., & Hinzman, L. (2003). DInSAR measurement of soil moisture. *IEEE Transactions on Geoscience and Remote Sensing, 41*(12), 2802–2813. https://doi.org/10.1109/TGRS.2003.817211

Notti, D., Davalillo, J. C., Herrera, G., & Mora, O. (2010). Assessment of the performance of X-band satellite radar data for landslide mapping and monitoring: Upper Tena Valley case study. *Natural Hazards and Earth System Sciences, 10*(9), 1865–1875. https://doi.org/10.5194/nhess-10-1865-2010

NSDDA. (1998). *Geospatial positioning accuracy standards part 3: National standard for spatial data accuracy* (FGDC-STD-007.3-1998).

Nuth, C., & Kääb, A. (2011). Co-registration and bias corrections of satellite elevation data sets for quantifying glacier thickness change. *The Cryosphere, 5*(1), 271–290. https://doi.org/10.5194/tc-5-271-2011

Ohki, M., Abe, T., Tadono, T., & Shimada, M. (2020). Landslide detection in mountainous forest areas using polarimetry and interferometric coherence. *Earth, Planets and Space, 72*(1), 67. https://doi.org/10.1186/s40623-020-01191-5

OPERA. (2023). *OPERA Level-2 Coregistered Single-Look Complex from Sentinel-1 (OPERA_L2_CSLC-S1, Version 1)* [Data set]. NASA Alaska Satellite Facility Distributed Active Archive Center. https://doi.org/10.5067/SNWG/OPERA_L2_CSLC-S1_V1

Østrem, G. (1959). Ice melting under a thin layer of moraine and the existence of ice cores in moraine ridges. *Geografiska Annaler, 41*(4), 228–230. https://doi.org/10.1080/20014422.1959.11907953

Østrem, G. (1965). Problems of dating ice-cored moraines. *Geografiska Annaler: Series A, Physical Geography, 47*(1), 1–38. https://doi.org/10.1080/04353676.1965.11879710

Painter, T. H. (2018). *ASO L4 Lidar Snow Depth 50 m UTM Grid (ASO_50M_SD, Version 1)* [Data set]. NASA National Snow and Ice Data Center Distributed Active Archive Center. https://doi.org/10.5067/STOT5I0U1WVI

Painter, T. H., Berisford, D. F., Boardman, J. W., Bormann, K. J., Deems, J. S., Gehrke, F., … Winstral, A. (2016). The Airborne Snow Observatory: Fusion of scanning lidar, imaging spectrometer, and physically based modeling for mapping snow water equivalent and snow albedo. *Remote Sensing of Environment, 184*, 139–152. https://doi.org/10.1016/j.rse.2016.06.013

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., … Chintala, S. (2019). PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems* (Vol. 32). Curran Associates.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., … Duchesnay, É. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research, 12*, 2825–2830.

Petley, D. N., Dunning, S. A., & Rosser, N. J. (2005). The analysis of global landslide risk through the creation of a database of worldwide landslide fatalities. In O. Hungr, R. Fell, R. Couture, & E. Eberhardt (Eds.), *Landslide risk management* (pp. 377–384). CRC Press.

Pflug, J. M., Kumar, S., Hall, D. K., Riggs, G. A., Konapala, G., Whitney, K., … Arsenault, K. R. (2025). Lightweight and regionally transferrable snow water equivalent estimation using a long short-term memory network. *Authorea Preprints*.

Porter, C., Howat, I., Husby, E., Noh, M. J., Khuvis, S., Danish, E., … Morin, P. (2022). *EarthDEM – Strips (Version 1)* [Data set]. Harvard Dataverse. https://doi.org/10.7910/DVN/LHE9O7

Rana, B., Shrestha, A. B., Reynolds, J. M., & Aryal, R. (2000). Hazard assessment of the Tsho Rolpa Glacier Lake and ongoing remediation measures. *Journal of Nepal Geological Society, 22*, 563–570.

Ravanel, L., Duvillard, P.-A., Jaboyedoff, M., & Lambiel, C. (2018). Recent evolution of an ice-cored moraine at the Gentianes Pass, Valais Alps, Switzerland. *Land Degradation & Development, 29*(10), 3693–3708. https://doi.org/10.1002/ldr.3088

Riley, S. J., DeGloria, S. D., & Elliot, R. A. (1999). A terrain ruggedness index that quantifies topographic heterogeneity. *Intermountain Journal of Sciences, 5*(1–4), 23–27.

Rémy, D., Chen, Y., Froger, J.-L., Bonvalot, S., Cordoba, L., & Fustos, J. (2015). Revised interpretation of recent InSAR signals observed at Llaima volcano (Chile). *Geophysical Research Letters, 42*, 3870–3879. https://doi.org/10.1002/2015GL063872

Riaz, S., Ali, A., & Baig, M. N. (2014). Increasing risk of glacial lake outburst floods as a consequence of climate change in the Himalayan region. *Jàmbá: Journal of Disaster Risk Studies, 6*(1), Article 110. https://doi.org/10.4102/jamba.v6i1.110

Richardson, S. D., & Reynolds, J. M. (2000a). An overview of glacial hazards in the Himalayas. *Quaternary International, 65–66*, 31–47. https://doi.org/10.1016/S1040-6182(99)00035-X

Richardson, S. D., & Reynolds, J. M. (2000b). Degradation of ice-cored moraine dams: Implications for hazard development. *IAHS Publication, 264*.

Rick, B., McGrath, D., Armstrong, W., & McCoy, S. W. (2022). Dam type and lake location characterize ice-marginal lake area change in Alaska and NW Canada between 1984 and 2019. *The Cryosphere, 16*(1), 297–314. https://doi.org/10.5194/tc-16-297-2022

Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention – MICCAI 2015* (pp. 234–241). Springer. https://doi.org/10.1007/978-3-319-24574-4_28

Rosen, P. A., Gurrola, E., Sacco, G. F., & Zebker, H. (2012). The InSAR scientific computing environment. In *EUSAR 2012: 9th European Conference on Synthetic Aperture Radar* (pp. 730–733).

Rosen, P. A., Hensley, S., Joughin, I. R., Li, F. K., Madsen, S. N., Rodríguez, E., & Goldstein, R. M. (2000). Synthetic aperture radar interferometry. *Proceedings of the IEEE, 88*(3), 333–382. https://doi.org/10.1109/5.838084

Rosi, A., Tofani, V., Tanteri, L., Tacconi Stefanelli, C., Agostini, A., Catani, F., & Casagli, N. (2018). The new landslide inventory of Tuscany (Italy) updated with PS-InSAR: Geomorphological features and landslide distribution. *Landslides, 15*, 5–19.

Ross, A. B. (1976). *A form and process response study of a terminal ice-cored ablation moraine* (Doctoral dissertation). University of Ottawa, Canada.

Rouet-Leduc, B., Jolivet, R., Dalaison, M., Johnson, P. A., & Hulbert, C. (2021). Autonomous extraction of millimeter-scale deformation in InSAR time series using deep learning. *Nature Communications, 12*, 6480. https://doi.org/10.1038/s41467-021-26254-3

Rounce, D. R., Hock, R., & Shean, D. E. (2020). Glacier mass change in High Mountain Asia through 2100 using the open-source Python Glacier Evolution Model (PyGEM). *Frontiers in Earth Science, 7*, 331.

Rounce, D. R., McKinney, D. C., Lala, J. M., Byers, A. C., & Watson, C. S. (2016). A new remote hazard and risk assessment framework for glacial lakes in the Nepal Himalaya. *Hydrology and Earth System Sciences, 20*(9), 3455–3475. https://doi.org/10.5194/hess-20-3455-2016

Rounce, D. R., Watson, C. S., & McKinney, D. C. (2017). Identification of hazard and risk for glacial lakes in the Nepal Himalaya using satellite imagery from 2000–2015. *Remote Sensing, 9*(7), 654. https://doi.org/10.3390/rs9070654

Routt, A. (2022). *An investigation of intra-annual flow rates of a rock glacier in the San Juan Mountains of Southwest Colorado* (Master's thesis). University of Colorado Colorado Springs.

Samsonov, S., Tiampo, K., & Cassotto, R. (2021). Measuring the state and temporal evolution of glaciers in Alaska and Yukon using SAR-derived 3D time series of glacier surface flow. *The Cryosphere, 15*(9), 4221–4243. https://doi.org/10.5194/tc-15-4221-2021

Sánchez-Gámez, P., & Navarro, F. J. (2017). Glacier surface velocity retrieval using D-InSAR and offset-tracking techniques applied to ascending and descending passes of Sentinel-1 data for southern Ellesmere Ice Caps, Canadian Arctic. *Remote Sensing, 9*(5), 442. https://doi.org/10.3390/rs9050442

Saranathan, A. M., Saeedimoghaddam, M., Smith, B., Raghunandan, D., Nearing, G., & Pelissier, C. (2025). A physically driven long short-term memory model for estimating snow water equivalent over the continental United States. *arXiv Preprint*, arXiv:2504.20129.

Sattar, A., Cook, K. L., Rai, S. K., Berthier, E., Allen, S., Rinzin, S., … Bhat, S. Y. (2025). The Sikkim flood of October 2023: Drivers, causes, and impacts of a multihazard cascade. *Science, 387*(6740), eads2659.

Sattar, A., Goswami, A., Kulkarni, A. V., Emmer, A., Haritashya, U. K., Allen, S., Frey, H., & Huggel, C. (2021). Future glacial lake outburst flood (GLOF) hazard of the South Lhonak Lake, Sikkim Himalaya. *Geomorphology, 388*, 107783. https://doi.org/10.1016/j.geomorph.2021.107783

Sawagaki, T., Lamsal, D., Byers, A. C., & Watanabe, T. (2012). Changes in surface morphology and glacial lake development of Chamlang South Glacier in the eastern Nepal Himalaya since 1964. *Global Environmental Research, 16*, 83–94.

Scaioni, M., Longoni, L., Melillo, V., & Papini, M. (2014). Remote sensing for landslide investigations: An overview of recent achievements and perspectives. *Remote Sensing, 6*(10), 9600–9652.

Scapozza, C., Ambrosi, C., Cannata, M., & Strozzi, T. (2019). Glacial lake outburst flood hazard assessment by satellite Earth observation in the Himalayas (Chomolhari area, Bhutan). *Geographica Helvetica, 74*(1), 125–139.

Schmidt, D. A., & Bürgmann, R. (2003). Time-dependent land uplift and subsidence in the Santa Clara Valley, California, from a large InSAR data set. *Journal of Geophysical Research: Solid Earth, 108*(B9). https://doi.org/10.1029/2002JB002267

Schomacker, A. (2008). What controls dead-ice melting under different climate conditions? A discussion. Earth-Science Reviews, 90(3–4), 103–113. https://doi.org/10.1016/j.earscirev.2008.08.003

Schomacker, A., & Kjær, K. H. (2008). Quantification of dead-ice melting in ice-cored moraines at the high-Arctic glacier Holmströmbreen, Svalbard. *Boreas, 37*(2), 211–225. https://doi.org/10.1111/j.1502-3885.2007.00014.x

Schwegmann, C. P., Kleynhans, W., Engelbrecht, J., Mdakane, L. W., & Meyer, R. G. V. (2017). Subsidence feature discrimination using deep convolutional neural networks in synthetic aperture radar imagery. In *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)* (pp. 4626–4629). https://doi.org/10.1109/IGARSS.2017.8128031

Sebastian, M. (2023, October 5). *Sikkim: Race against time to save 102 missing in India floods.* BBC News. https://www.bbc.com/news/business-66997995

Shean, D. E., Alexandrov, O., Moratto, Z. M., Smith, B. E., Joughin, I. R., Porter, C., & Morin, P. (2016). An automated, open-source pipeline for mass production of digital elevation models from very-high-resolution commercial stereo satellite imagery. *ISPRS Journal of Photogrammetry and Remote Sensing, 116*, 101–117. https://doi.org/10.1016/j.isprsjprs.2016.03.012

Shean, D., Bhushan, S., Lilien, D., Knuth, F., Schwat, E., Meyer, J., Sharp, M., & Hu, M. (2023). *dshean/demcoreg: v1.1.1 – Compatibility and documentation improvements (Version 1.1.1)* [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.7730376

Shen, L., Hooper, A., & Elliott, J. (2019). A spatially varying scaling method for InSAR tropospheric corrections using a high-resolution weather model. *Journal of Geophysical Research: Solid Earth, 124*, 4051–4068. https://doi.org/10.1029/2018JB016189

Shugar, D. H., & Clague, J. J. (2011). The sedimentology and geomorphology of rock avalanche deposits on glaciers. *Sedimentology, 58*(7), 1762–1783. https://doi.org/10.1111/j.1365-3091.2011.01238.x

Shugar, D. H., Burr, A., Haritashya, U. K., Kargel, J. S., Watson, C. S., Kennedy, M. C., … Strattman, K. (2020). Rapid worldwide growth of glacial lakes since 1990. *Nature Climate Change, 10*(10), 939–945. https://doi.org/10.1038/s41558-020-0855-4

Sica, F., Calvanese, F., Scarpa, G., & Rizzoli, P. (2022). A CNN-based coherence-driven approach for InSAR phase unwrapping. *IEEE Geoscience and Remote Sensing Letters, 19*, 1–5. https://doi.org/10.1109/LGRS.2020.3029565

Small, D. (2011). Flattening gamma: Radiometric terrain correction for SAR imagery. *IEEE Transactions on Geoscience and Remote Sensing, 49*(8), 3081–3093.

Snethlage, M. A., Geschke, J., Ranipeta, A., Jetz, W., Yoccoz, N. G., Körner, C., … Urbach, D. (2022). A hierarchical inventory of the world's mountains for global comparative mountain science. *Scientific Data, 9*, 149.

Somos-Valenzuela, M. A., Chisolm, R. E., Rivas, D. S., Portocarrero, C., & McKinney, D. C. (2016). Modeling a glacial lake outburst flood process chain: The case of Lake Palcacocha and Huaraz, Peru. *Hydrology and Earth System Sciences, 20*(6), 2519–2543.

Somos-Valenzuela, M., McKinney, D. C., Byers, A. C., Rounce, D. R., & Portocarrero, C. (2013). *Modeling mitigation strategies for risk reduction at Imja Lake, Nepal.* Center for Research in Water Resources, University of Texas at Austin. http://hdl.handle.net/2152/27736

Somos-Valenzuela, M., McKinney, D. C., Byers, A. C., Voss, K., Moss, J., & McKinney, J. C. (2012). *Ground penetrating radar survey for risk reduction at Imja Lake, Nepal.* Center for Research in Water Resources, University of Texas at Austin. http://hdl.handle.net/2152/19751

Sorkhabi, O. M., Nejad, A. S., & Khajehzadeh, M. (2022). Evaluation of Isfahan City subsidence rate using InSAR and artificial intelligence. *KSCE Journal of Civil Engineering, 26*, 2901–2908. https://doi.org/10.1007/s12205-022-2013-1

Stoffel, M., & Huggel, C. (2012). Effects of climate change on mass movements in mountain environments. *Progress in Physical Geography: Earth and Environment, 36*, 421–439. https://doi.org/10.1177/0309133312441010

Strozzi, T., Kouraev, A., Wiesmann, A., Wegmüller, U., Sharov, A., & Werner, C. (2008). Estimation of Arctic glacier motion with satellite L-band SAR data. *Remote Sensing of Environment, 112*(3), 636–645. https://doi.org/10.1016/j.rse.2007.06.007

Strozzi, T., Luckman, A., Murray, T., Wegmüller, U., & Werner, C. L. (2002). Glacier motion estimation using SAR offset-tracking procedures. *IEEE Transactions on Geoscience and Remote Sensing, 40*(11), 2384–2391. https://doi.org/10.1109/TGRS.2002.805079

Stucker, C., & Schindler, K. (2022). ResDepth: A deep residual prior for 3D reconstruction from high-resolution satellite images. *ISPRS Journal of Photogrammetry and Remote Sensing, 183*, 560–580. https://doi.org/10.1016/j.isprsjprs.2021.11.009

Sturm, M., & Liston, G. E. (2021). Revisiting the global seasonal snow classification: An updated dataset for Earth system applications. *Journal of Hydrometeorology, 22*(11), 2917–2938.

Sturm, M., Goldstein, M. A., & Parr, C. (2017). Water and life from snow: A trillion-dollar science question. *Water Resources Research, 53*(5), 3534–3544.

Sturm, M., Taras, B., Liston, G. E., Derksen, C., Jonas, T., & Lea, J. (2010). Estimating snow water equivalent using snow depth data and climate classes. *Journal of Hydrometeorology, 11*(6), 1380–1394.

Sun, S., Che, T., Wang, J., Li, H., Hao, X., Wang, Z., & Wang, J. (2015). Estimation and analysis of snow water equivalents based on C-band SAR data and field measurements. *Arctic, Antarctic, and Alpine Research, 47*(2), 313–326. https://doi.org/10.1657/AAAR00C-13-135

Taylor, C., Robinson, T. R., Dunning, S., Carr, J. R., & Westoby, M. (2023). Glacial lake outburst floods threaten millions globally. *Nature Communications, 14*(1), 487.

Tonkin, T. N., Midgley, N. G., Cook, S. J., & Graham, D. J. (2016). Ice-cored moraine degradation mapped and quantified using an unmanned aerial vehicle: A case study from a polythermal glacier in Svalbard. *Geomorphology, 258*, 1–10. https://doi.org/10.1016/j.geomorph.2015.12.019

Tough, J. A., Blacknell, D., & Quegan, S. (1995). A statistical description of polarimetric and interferometric synthetic aperture radar data. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences, 449*(1937), 567–589. https://doi.org/10.1098/rspa.1995.0059

Trujillo, E., Ramirez, J. A., & Elder, K. J. (2009). Scaling properties and spatial organization of snow depth fields in sub-alpine forest and alpine tundra. Hydrological Processes, 23(11), 1575–1590. https://doi.org/10.1002/hyp.7270

Tshering, S., & Tamang, B. (2004, October). *Hydropower—Key to sustainable, socio-economic development of Bhutan*. In *United Nations Symposium on Hydropower and Sustainable Development* (pp. 27–29).

UCAR/NCAR Earth Observing Laboratory. (2011). *CEOP_AP: Himalayas Surface Meteorology and Radiation Data Set* (Version 1.0) [Data set]. UCAR/NCAR Earth Observing Laboratory. https://doi.org/10.26023/ZZHE-8SAB-040R

UNDP. (2012). *Community Based Flood and Glacial Lake Outburst Risk Reduction Project* [Project document]. https://www.undp.org/sites/g/files/zskgke326/files/migration/np/aa21b4ccde4230b2b26dda751438572149d185892d482a1cf417b3a1737ce05a.pdf

USDA NRCS. (2025). *SNOwpack TELemetry Network (SNOTEL)* [Data set]. U.S. Department of Agriculture, Natural Resources Conservation Service. https://catalog.data.gov/dataset/snowpack-telemetry-network-snotel

USGS. (2025). *USGS 1/3 Arc Second n39w120 20250514* [Data set]. U.S. Geological Survey. https://www.sciencebase.gov/catalog/item/68269e8cd4be024c53ba6019

Varhola, A., Coops, N. C., Weiler, M., & Moore, R. D. (2010). Forest canopy effects on snow accumulation and ablation: An integrative review of empirical results. Journal of Hydrology, 392(3–4), 219–233. https://doi.org/10.1016/j.jhydrol.2010.08.009

Vitale, S., Ferraioli, G., & Pascazio, V. (2022). A CNN-based solution for InSAR phase denoising. In *IGARSS 2022 – IEEE International Geoscience and Remote Sensing Symposium* (pp. 931–934). https://doi.org/10.1109/IGARSS46834.2022.9883956

Viviroli, D., & Weingartner, R. (2004). The hydrological significance of mountains: From regional to global scale. Hydrology and Earth System Sciences, 8(6), 1017–1030. https://doi.org/10.5194/hess-8-1017-2004

van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Gouillart, E., & Yu, T. (2014). scikit-image: Image processing in Python. *PeerJ, 2*, e453. https://doi.org/10.7717/peerj.453

Wang, W., Gao, Y., Iribarren Anacona, P., Lei, Y., Xiang, Y., Zhang, G., Li, S., & Lu, A. (2018). Integrated hazard assessment of Cirenmaco glacial lake in Zhangzangbo Valley, Central Himalayas. *Geomorphology, 306*, 292–305. https://doi.org/10.1016/j.geomorph.2015.08.013

Wang, X., Liu, S., Ding, Y., Guo, W., Jiang, Z., Lin, J., & Han, Y. (2012). An approach for estimating the breach probabilities of moraine-dammed lakes in the Chinese Himalayas using remote-sensing data. *Natural Hazards and Earth System Sciences, 12*(10), 3109–3122. https://doi.org/10.5194/nhess-12-3109-2012

Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing, 13*(4), 600–612. https://doi.org/10.1109/TIP.2003.819861

Watanabe, T., Kameyama, S., & Sato, T. (1995). Imja Glacier dead-ice melt rates and changes in a supra-glacial lake, 1989–1994, Khumbu Himal, Nepal: Danger of lake drainage. *Mountain Research and Development, 15*(4), 293–300. https://doi.org/10.2307/3673805

Watanabe, T., Lamsal, D., & Ives, J. D. (2009). Evaluating the growth characteristics of a glacial lake and its degree of danger of outburst flooding: Imja Glacier, Khumbu Himal, Nepal. *Norsk Geografisk Tidsskrift–Norwegian Journal of Geography, 63*(4), 255–267. https://doi.org/10.1080/00291950903368367

Weiss, A. (2001, July). *Topographic position and landforms analysis*. Poster presented at the ESRI User Conference, San Diego, CA.

Westoby, M. J., Glasser, N. F., Brasington, J., Hambrey, M. J., Quincey, D. J., & Reynolds, J. M. (2014). Modelling outburst floods from moraine-dammed glacial lakes. Earth-Science Reviews, 134, 137–159. https://doi.org/10.1016/j.earscirev.2014.03.009

Wilk, M. B., & Gnanadesikan, R. (1968). Probability plotting methods for the analysis of data. Biometrika, 55(1), 1–17. https://doi.org/10.1093/biomet/55.1.1

Worni, R., Stoffel, M., Huggel, C., Volz, C., Casteller, A., & Luckman, B. (2012). Analysis and dynamic modeling of a moraine failure and glacier lake outburst flood at Ventisquero Negro, Patagonian Andes (Argentina). Journal of Hydrology, 444, 134–145. https://doi.org/10.1016/j.jhydrol.2012.04.013

Wright, T. J., Parsons, B. E., & Lu, Z. (2004). Toward mapping surface deformation in three dimensions using InSAR. *Geophysical Research Letters, 31*(1). https://doi.org/10.1029/2003GL018827

Wu, Z., Zhang, H., Wang, Y., Wang, T., & Wang, R. (2020). A deep learning based method for local subsidence detection and InSAR phase unwrapping: Application to mining deformation monitoring. In *IGARSS 2020 – IEEE International Geoscience and Remote Sensing Symposium* (pp. 20–23). https://doi.org/10.1109/IGARSS39084.2020.9323342

xDEM contributors. (2024). *xDEM* (Version 0.1.0) [Computer software]. Zenodo. https://doi.org/10.5281/zenodo.11492983

Yang, K., Rittger, K., Musselman, K. N., Bair, E. H., Dozier, J., Margulis, S. A., & Molotch, N. P. (2023). Intercomparison of snow water equivalent products in the Sierra Nevada, California, using airborne snow observatory data and ground observations. *Frontiers in Earth Science, 11*, 1106621.

Yang, L., Lu, Z., Ouyang, C., Zhao, C., Hu, X., & Zhang, Q. (2023). Glacial lake outburst flood monitoring and modeling through integrating multiple remote sensing methods and HEC-RAS. *Remote Sensing, 15*(22), 5327.

Yang, L., Lu, Z., Zhao, C., Kim, J., Yang, C., Wang, B., & Wang, Z. (2022). Analyzing the triggering factors of glacial lake outburst floods with SAR and optical images: A case study in Jinweng Co, Tibet, China. *Landslides, 19*(4), 855–864.

Yang, W., He, Y., Yao, S., Zhang, L., Cao, S., & Wen, Z. (2022). An InSAR interferogram filtering method based on multi-level feature fusion CNN. *Sensors, 22*, 5956. https://doi.org/10.3390/s22165956

Yu, C., Li, Z., & Penna, N. T. (2018). Interferometric synthetic aperture radar atmospheric correction using a GPS-based iterative tropospheric decomposition model. *Remote Sensing of Environment, 204*, 109–121. https://doi.org/10.1016/j.rse.2017.10.038

Yu, Y., Li, B., Li, Y., & Jiang, W. (2024). Retrospective analysis of glacial lake outburst flood (GLOF) using AI Earth InSAR and optical images: A case study of South Lhonak Lake, Sikkim. *Remote Sensing, 16*(13), 2307.

Yunjun, Z., Fattahi, H., & Amelung, F. (2019). Small baseline InSAR time series analysis: Unwrapping error correction and noise reduction. *Computers & Geosciences, 133*, 104331. https://doi.org/10.1016/j.cageo.2019.104331

Zebker, H. (2021). Accuracy of a model-free algorithm for temporal InSAR tropospheric correction. *Remote Sensing, 13*(3), 409.

Zebker, H. A., Rosen, P. A., & Hensley, S. (1997). Atmospheric effects in interferometric synthetic aperture radar surface deformation and topographic maps. *Journal of Geophysical Research: Solid Earth, 102*(B4), 7547–7563.

Zevenbergen, L. W., & Thorne, C. R. (1987). Quantitative analysis of land surface topography. *Earth Surface Processes and Landforms, 12*(1), 47–56.

Zhang, T., Wang, W., & An, B. (2025). A massive lateral moraine collapse triggered the 2023 South Lhonak Lake outburst flood, Sikkim Himalayas. *Landslides, 22*(2), 299–311. https://doi.org/10.1007/s10346-024-02358-x

Zhang, T., Wang, W., Gao, T., An, B., & Yao, T. (2022). An integrative method for identifying potentially dangerous glacial lakes in the Himalayas. *Science of the Total Environment, 806*, 150442.

Zhang, Z., Lin, H., Wang, M., Liu, X., Chen, Q., Wang, C., & Zhang, H. (2022). A review of satellite synthetic aperture radar interferometry applications in permafrost regions: Current status, challenges, and trends. *IEEE Geoscience and Remote Sensing Magazine, 10*(3), 93–114. https://doi.org/10.1109/MGRS.2022.3170350

Zhao, Z., Wu, Z., Zheng, Y., & Ma, P. (2021). Recurrent neural networks for atmospheric noise removal from InSAR time series with missing values. *ISPRS Journal of Photogrammetry and Remote Sensing, 180*, 227–237. https://doi.org/10.1016/j.isprsjprs.2021.08.009

Zheng, W., Bhushan, S., Van Wyk de Vries, M., Kochtitzky, W., Shean, D., Copland, L., Dow, C., Jones-Ivey, R., & Pérez, F. (2023). GLAcier Feature Tracking testkit (GLAFT): A statistically and physically based framework for evaluating glacier velocity products derived from optical satellite image feature tracking. *The Cryosphere, 17*(9), 4063–4078. https://doi.org/10.5194/tc-17-4063-2023

Zhou, H., Dai, K., Pirasteh, S., Li, R., Xiang, J., & Li, Z. (2023). InSAR spatial-heterogeneity tropospheric delay correction in steep mountainous areas based on deep learning for landslide monitoring. *IEEE Transactions on Geoscience and Remote Sensing, 61*, 1–14. https://doi.org/10.1109/TGRS.2023.3307477

Zhou, H., Dai, K., Wen, N., Xiang, J., & Li, Z. (2024). The influence of topography-dependent atmospheric delay for InSAR time series results and the deep neural network correction. *IEEE Geoscience and Remote Sensing Letters, 21*, 1–5. https://doi.org/10.1109/LGRS.2024.3379982

Appendix A

SUPPLEMENTARY MATERIAL FOR CHAPTER 1: REMOVING ATMOSPHERIC NOISE FROM INSAR INTERFEROGRAMS IN MOUNTAINOUS REGIONS WITH A CONVOLUTIONAL NEURAL NETWORK

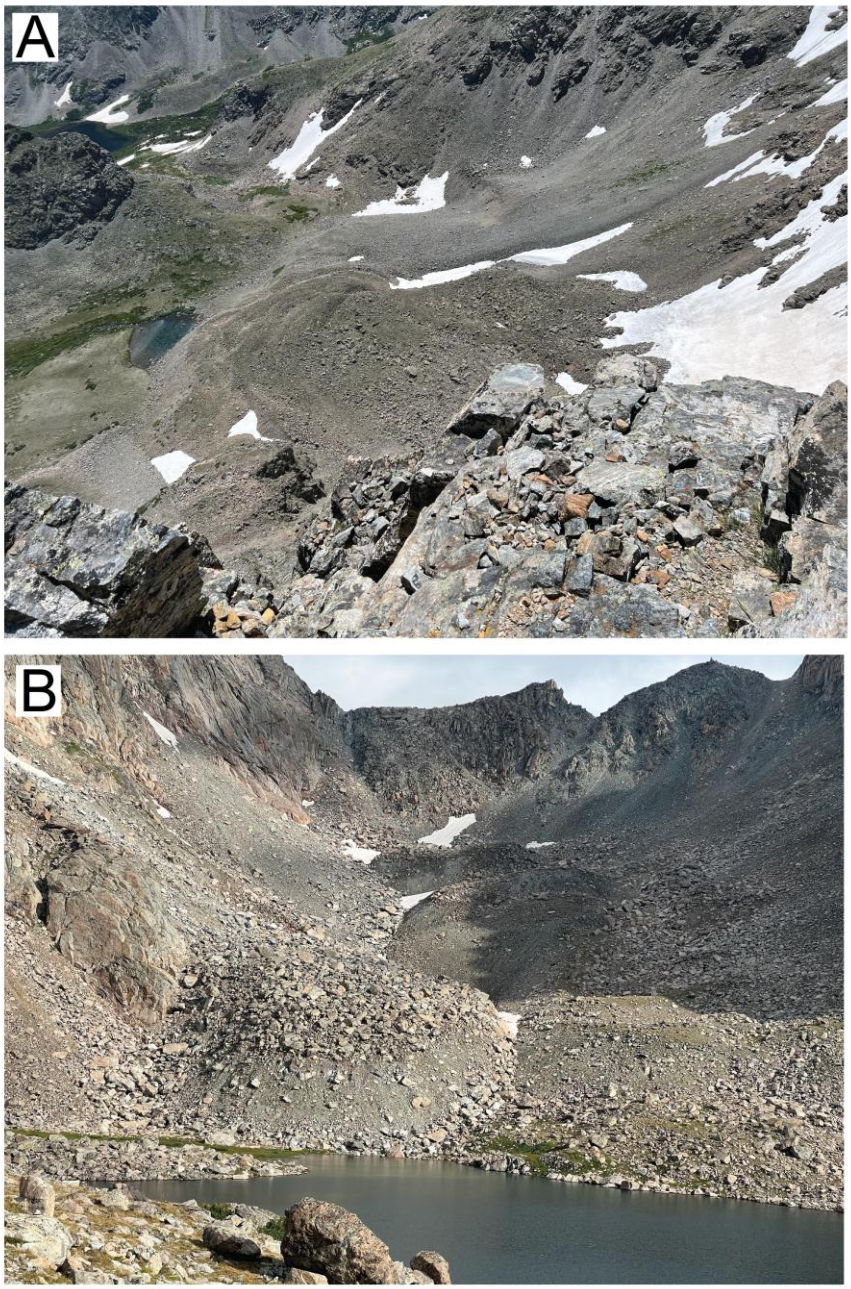

Figure A.1: Examples of rock glaciers in the study and application areas. (a) A study area rock glacier in the Front Range, Colorado, USA. Rock glacier motion is well-represented in the characteristic displacement signal maps used to train the CNN. (b) An application area rock glacier in the Wind River Range, Wyoming, USA.

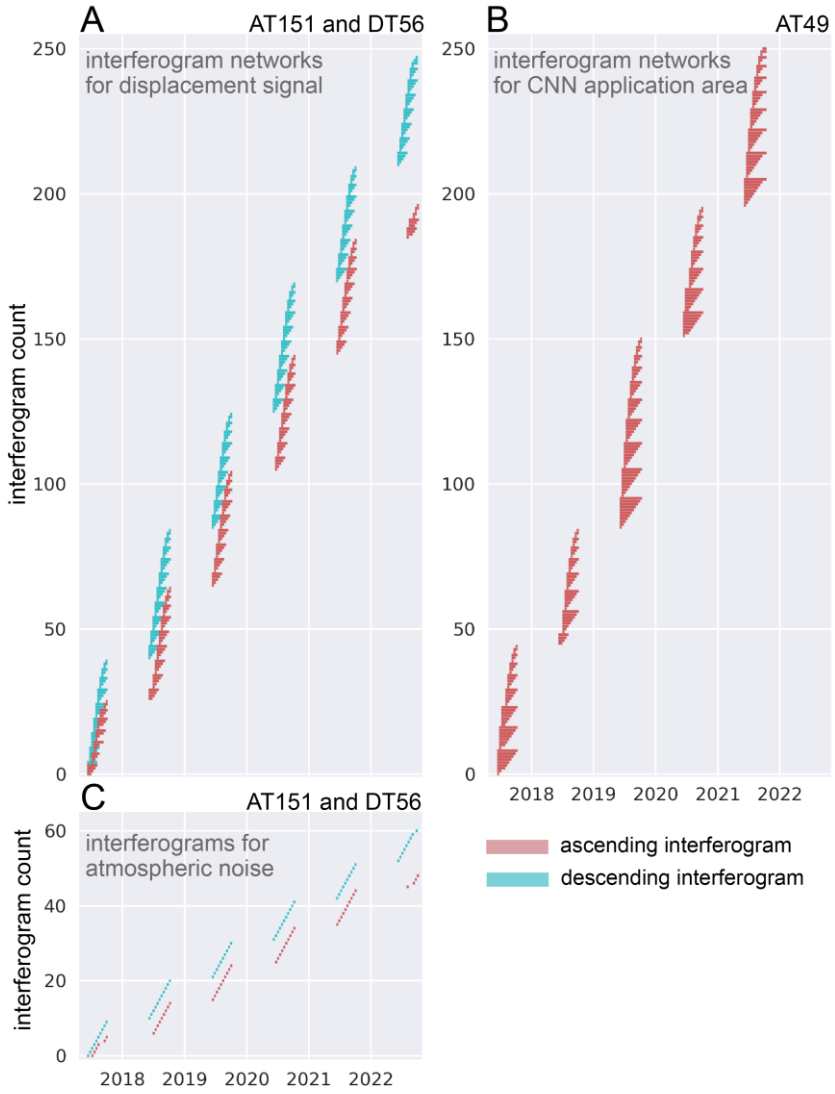

Figure A.2: Temporal baselines of all interferograms used for training, validation, testing, and application. Each bar spans the temporal baseline of a single interferogram, respectively beginning and ending at the primary and secondary acquisition dates. (A, C) Interferograms for the study area (Figure 2.1). For a given acquisition date, three interferograms were generated, as three Sentinel-1 frames are required to cover the study area. Here, we only show one bar per acquisition date. (A) Interferogram networks from ascending track 151 (AT151, red) and descending track 56 (DT56, cyan) used to prepare characteristic displacement signal maps for training, validation, and testing in the study area. (C) Short-baseline interferograms used to prepare characteristic atmospheric noise maps. (B) Interferogram networks from ascending track 49 (AT49, red) used to generate the application dataset.

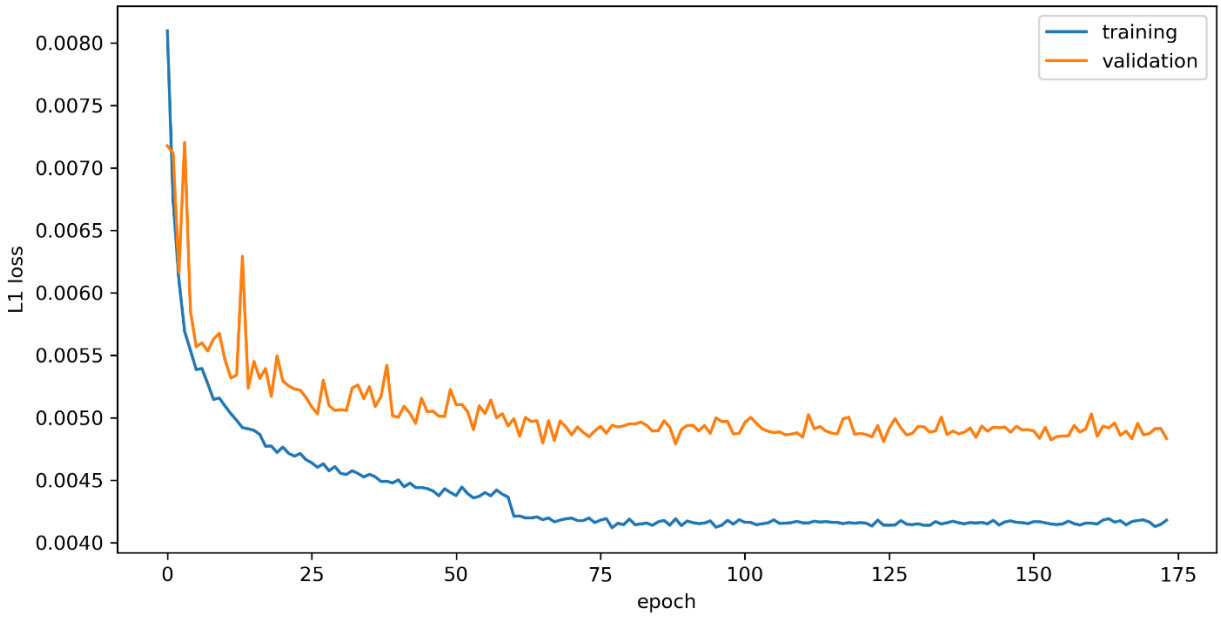

Figure A.3: Training curve of the convolutional neural network. Mean L1 loss (mean absolute error between the target displacement signal and the CNN-corrected displacement signal) is plotted for training (blue, n = 14,896 interferograms) and validation data (orange, n = 2,217 interferograms) for each epoch. Early stopping was triggered at epoch 174 to limit overfitting.

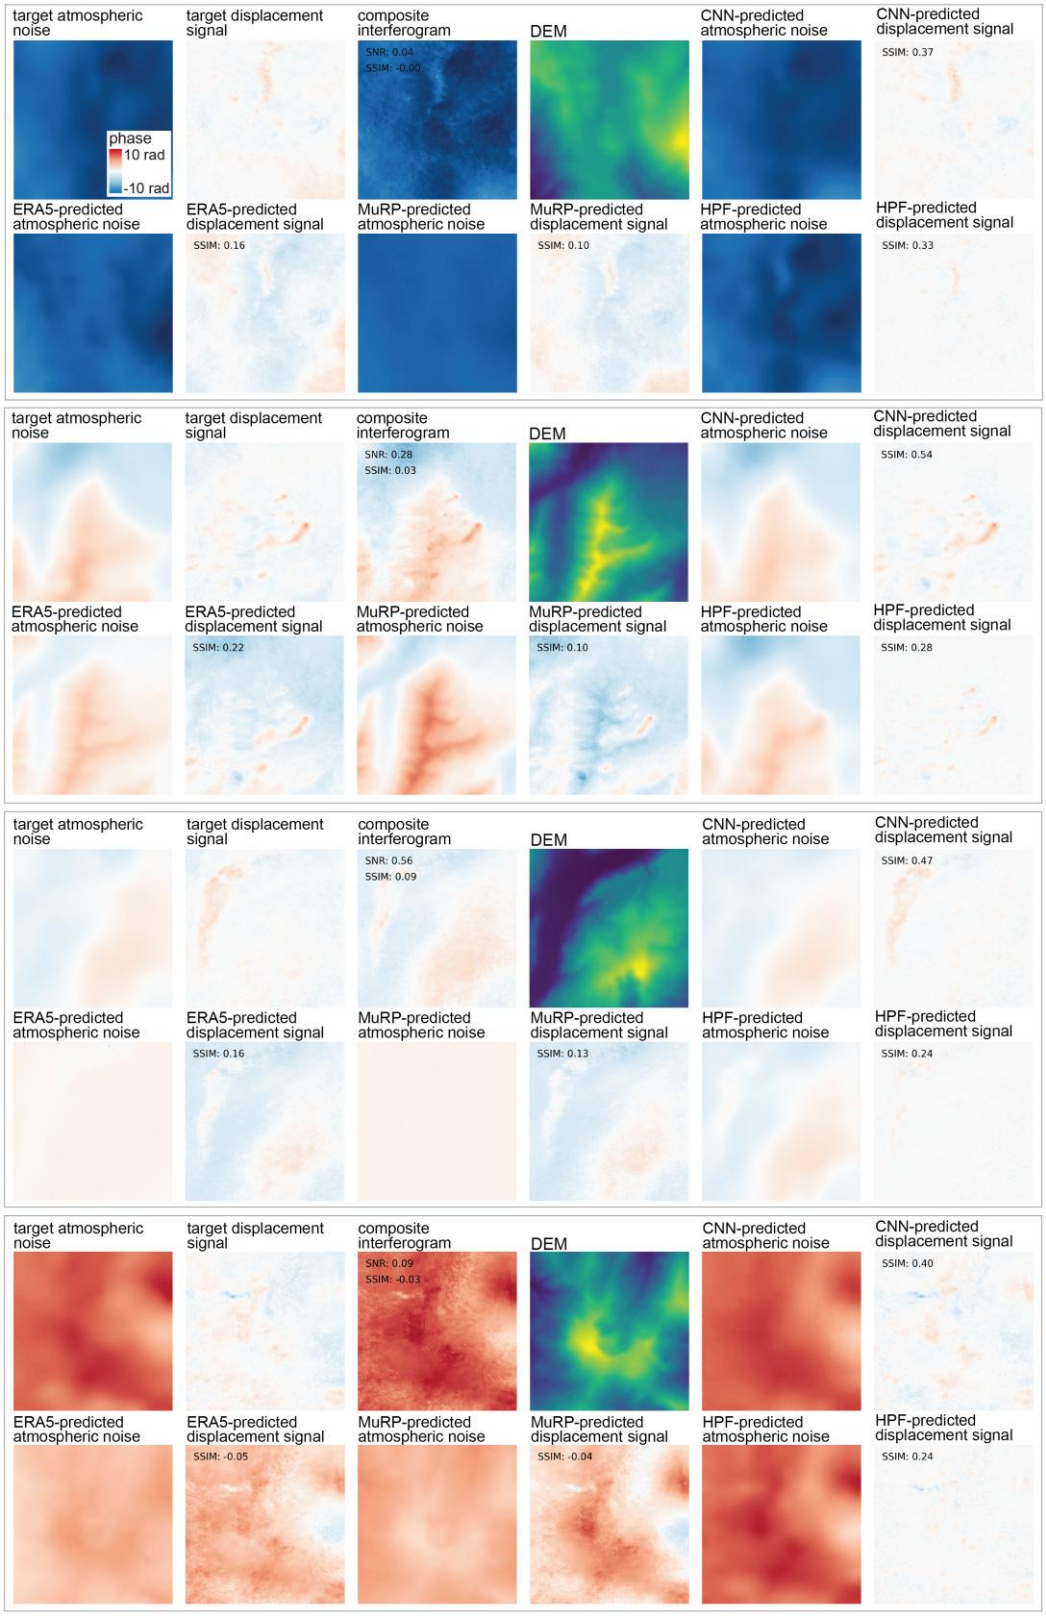

Figure A.4: Comparison between all correction approaches for four random 128 by 128-pixel validation subset stacks.

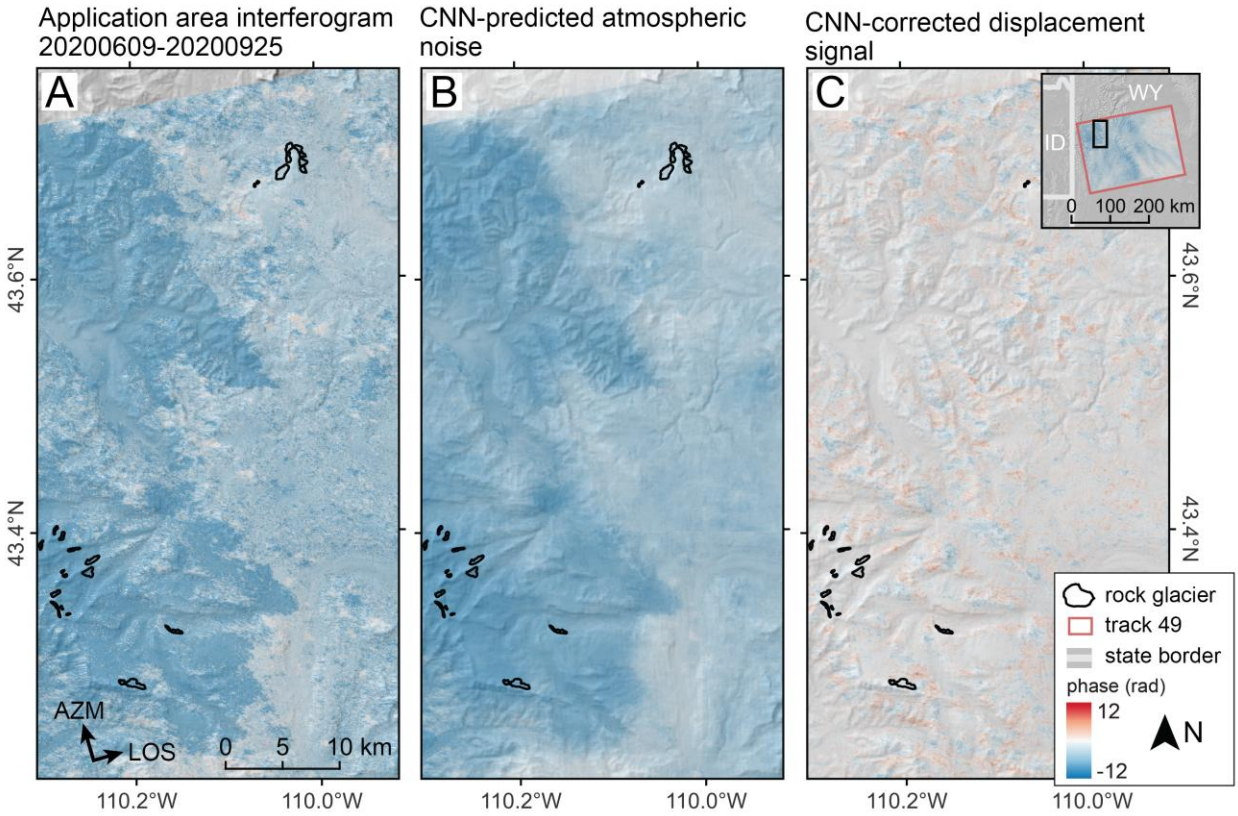

Figure A.5: CNN atmospheric noise correction results over unwrapping error artifacts. A) Application area interferogram from 20200609-20200925 with sharp phase discontinuity due to unwrapping error. B) CNN-predicted atmospheric noise for the application area interferogram. The CNN interprets large-scale apparent displacement signals as atmospheric noise, even where they have distinct edges. C) The CNN-corrected displacement signal map preserves small-scale apparent displacement signals.

Appendix B

SUPPLEMENTARY MATERIAL FOR CHAPTER 2: QUANTIFYING DEGRADATION OF THE IMJA LAKE MORAINE DAM WITH FUSED INSAR AND SAR FEATURE TRACKING TIME SERIES

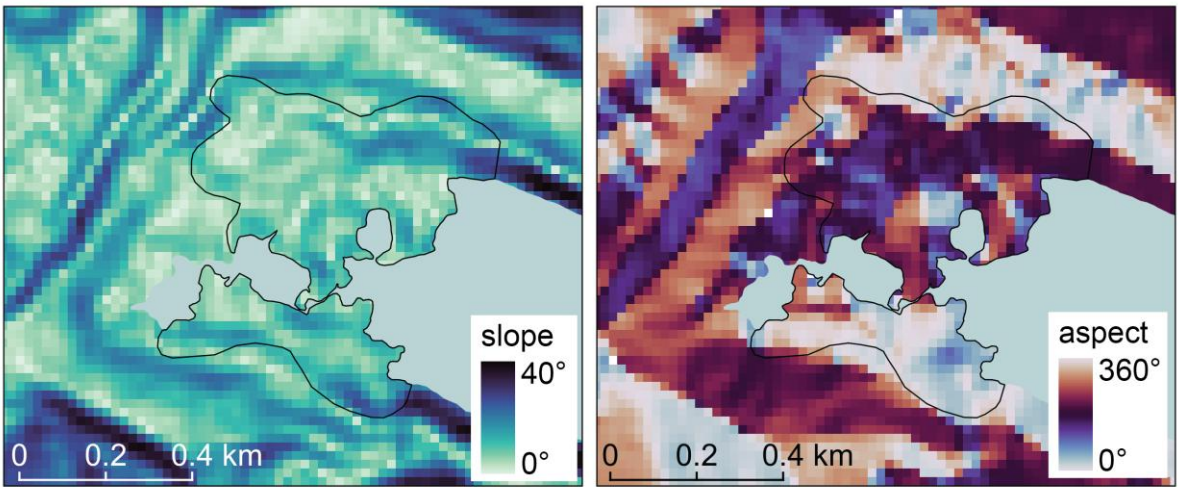

Figure B.1: Surface slope (left) and aspect (right) of the Imja Lake moraine dam from the COP30 composite digital elevation model prepared from TSX/TDX scenes acquired between 2011 and 2015 (European Space Agency, 2022).

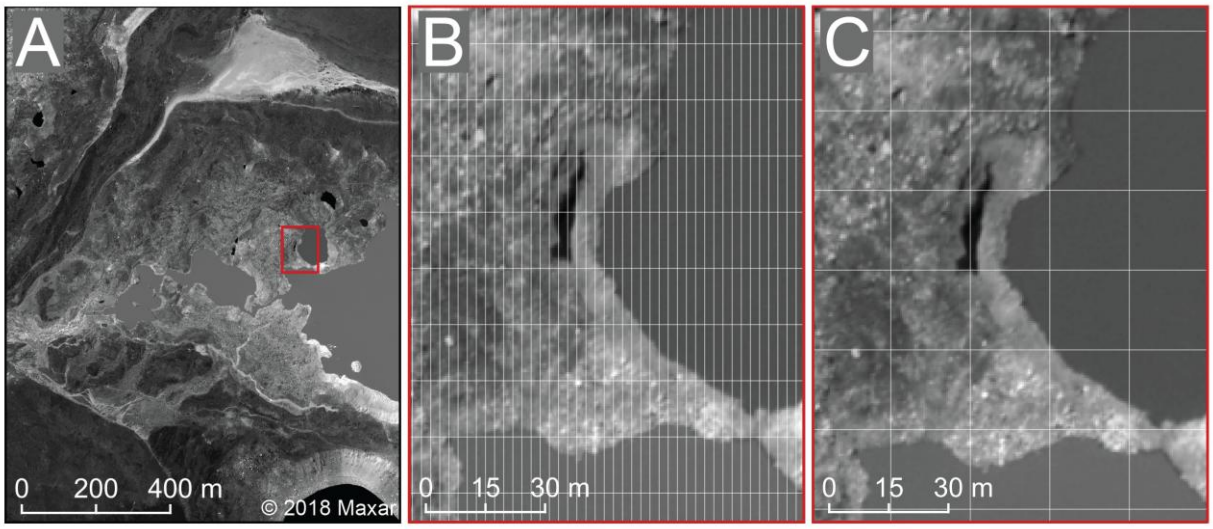

Figure B.2: Illustration of Sentinel-1 product resolution over the Imja Lake moraine dam. (A) Maxar WorldView-1 panchromatic orthoimage from June 19, 2018 (Catalog ID 10200100758B9800) with 0.53 m ground sample distance. (B) Detail of pond and lake margin denoted by red box in A, with 14.1 by 2.3 m grid overlaid showing native Sentinel-1 single-look complex (SLC) azimuth and range resolution, respectively. (C) Same as B, with 20 by 20 m grid showing the pixel spacing of our combined interferometric synthetic aperture radar (InSAR) and feature tracking products.

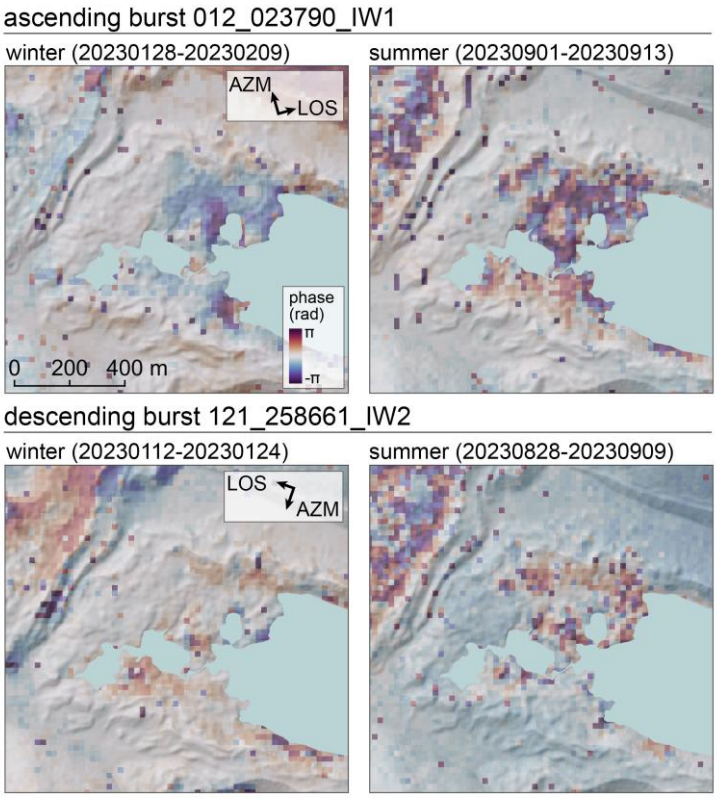

Figure B.3: Examples of wrapped interferograms. Note that no local reference point has been set for these interferograms, and they still contain some atmospheric noise over the moraine dam.

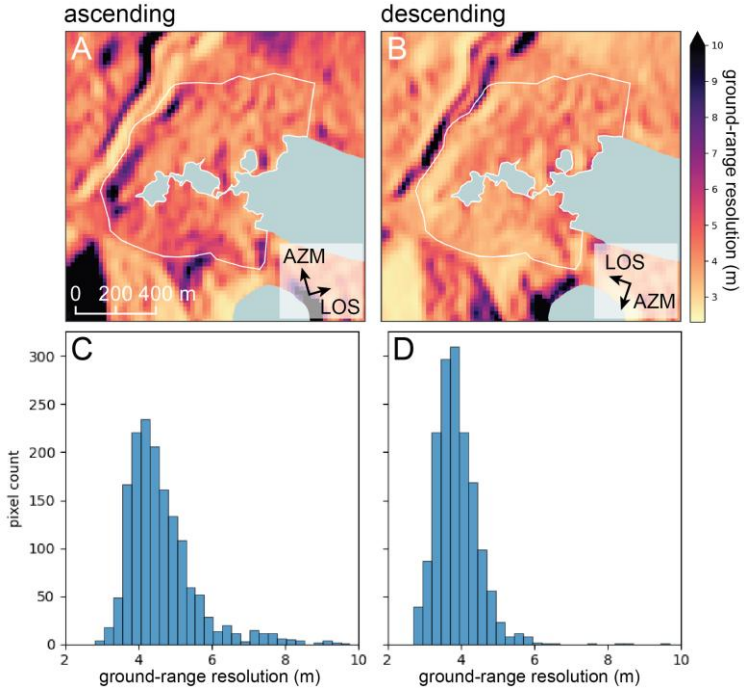

Figure B.4: Ground-range resolution of ascending (A) and descending (B) SLCs. C and D show the distribution of ground-range resolution values for pixels over the moraine dam (white outline).

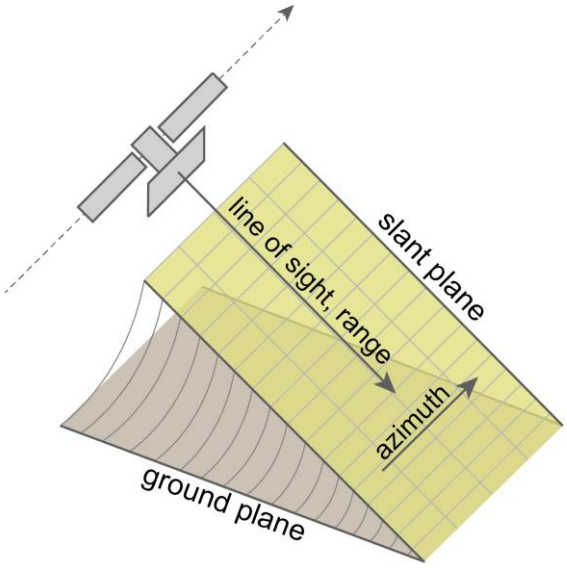

Figure B.5: SAR acquisition geometry.

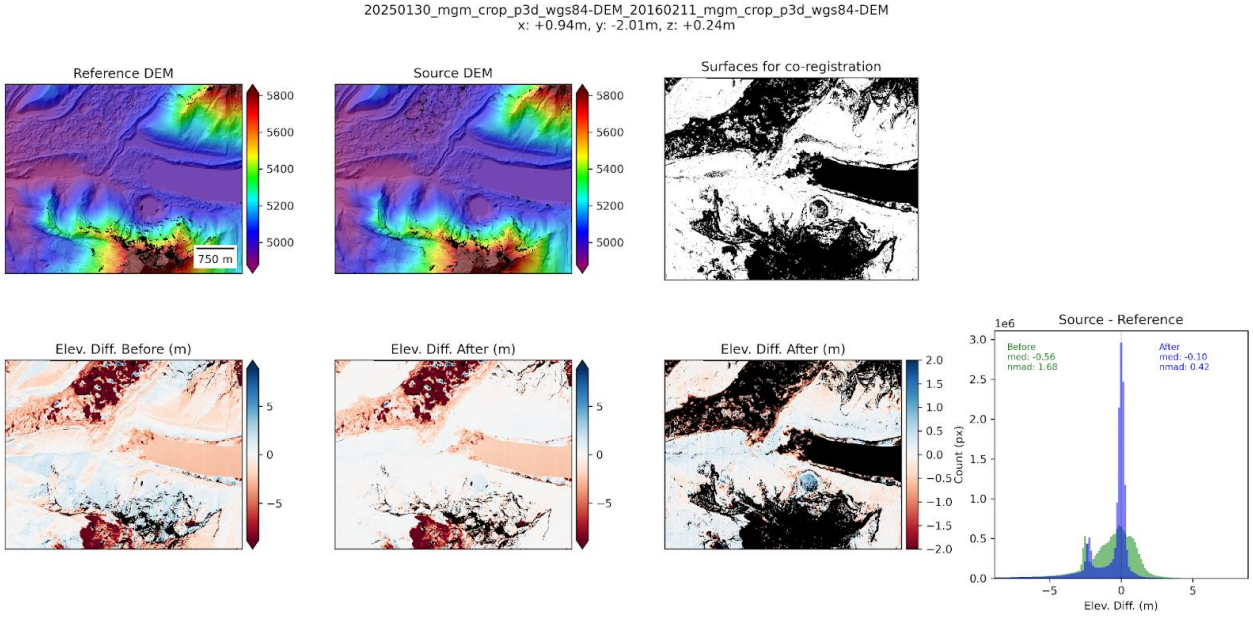

Figure B.6: Co-registration results for the February 11, 2016 DEM (reference) and January 30, 2025 DEM (source). The latter was shifted (+0.94 m, -2.01 m, +0.24 m, see title) to minimize residuals over unmasked surfaces (white in the top right figure). Bottom row shows elevation difference maps before co-registration (left), and after co-registration (center). Bottom right map shows enhanced color stretch and histogram shows unmasked values before and after co-registration. Note that some large negative values (-5 to -2.5 m) observed for unmasked pixels near glacier margins are included in the histogram, but these outliers did not affect the robust co-registration.

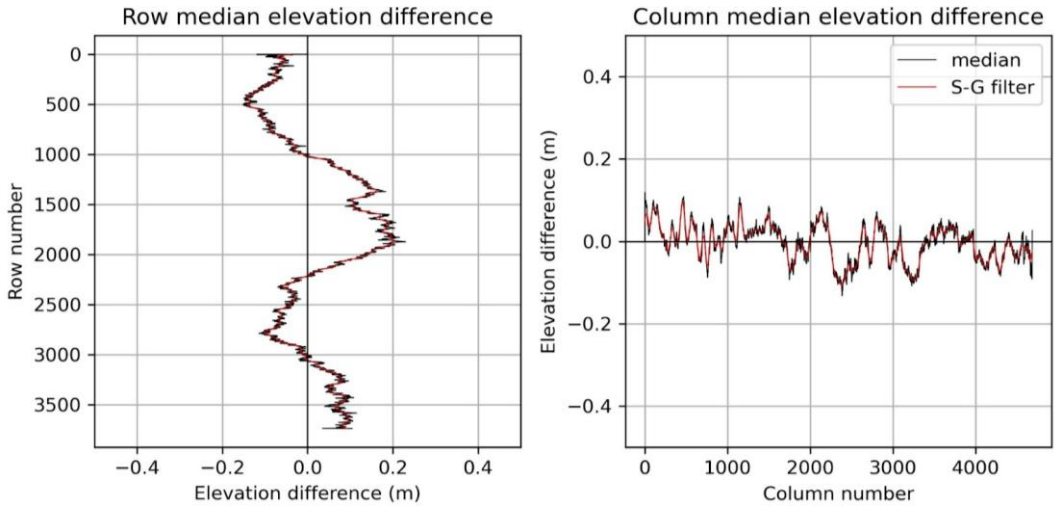

Figure B.7: Row (left panel) and column (right panel) median elevation difference values (black) over static control surfaces after DEM co-registration (see Figure S6), showing residual artifacts due unmodeled attitude error ("jitter", left) and detector array geometry calibration (right). See Shean et al. (2016) and the "ct_at_correction_wrapper" function of demcoreg/coreglib.py for more details. Red line shows the smoothed model from a Savitzky-Golay filter, which was used to correct all rows and columns in the unmasked elevation difference product.

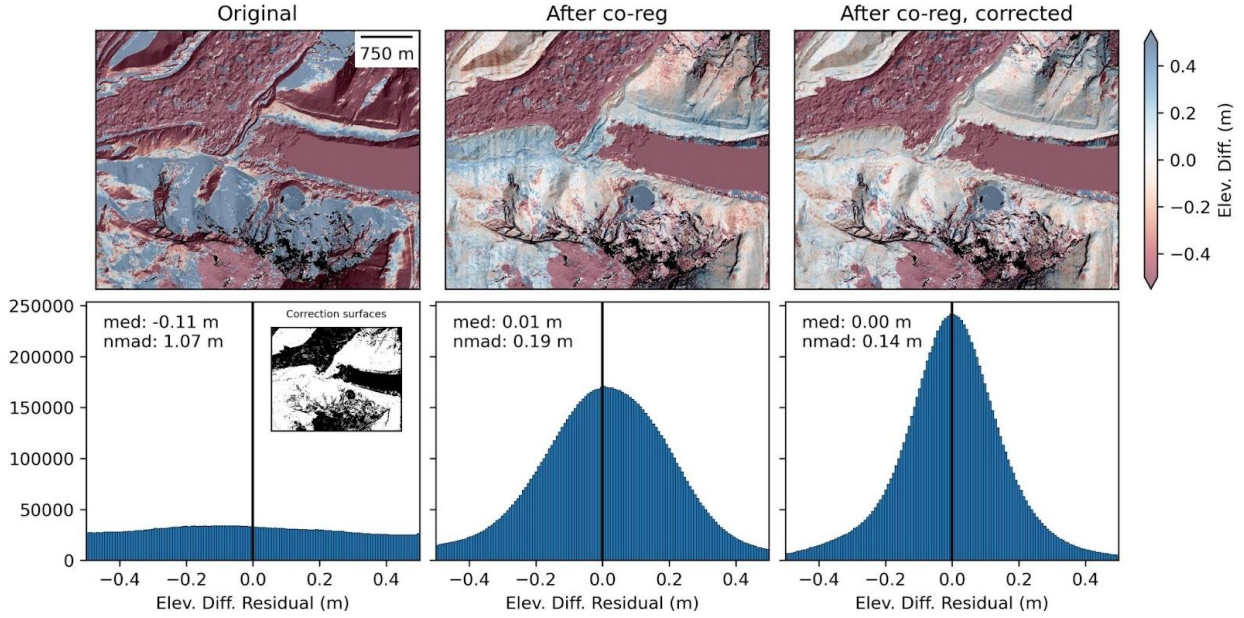

Figure B.8: Elevation difference maps (top row) and histograms (bottom row) of residuals over exposed surfaces (white areas in inset axes of bottom left panel) assumed to have no elevation change during the study period. Left shows original difference map, center shows difference map after co-registration, right shows final difference map after co-registration and correction of row/column median values.

B.1 Expected ice flow

We computed the expected surface-parallel velocity (u_s) due to internal deformation of buried ice using a simple 1-D model with no basal sliding (Eq. 8.35 in Cuffey & Paterson, 2010):

$$u_s = \frac{2A}{n+1}\tau_b{}^n H \tag{B.1}$$

where A is a flow rate factor (typically $2.4*10^{-24}$ $s^{-1}Pa^{-3}$ for temperate glacier ice), n is the flow law exponent (3), H is the ice thickness, and τ_b is the basal shear stress, defined as:

$$\tau_b = \rho g H \sin(\theta) \tag{B.2}$$

where ρ is ice density (917 kg/m^3), g is acceleration due to gravity (9.81 m/s^2), H is ice thickness, and θ is surface slope. We computed expected surface-parallel velocity due to ice flow for a range of ice thicknesses (0-70 m) and surface slope values (0-20°) for the Imja Lake moraine dam.

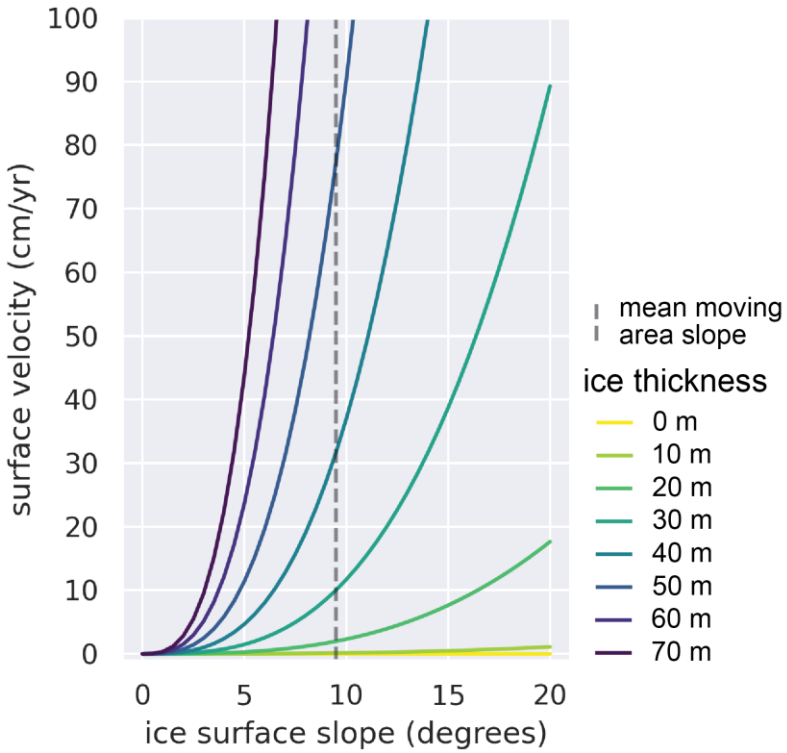

Figure B.9: Expected along-slope surface velocity due to internal deformation (no sliding) of temperate glacier ice (Eq 8.35 in Cuffey & Paterson, 2010) for a range of ice thicknesses and surface slope for the Imja Lake moraine dam.

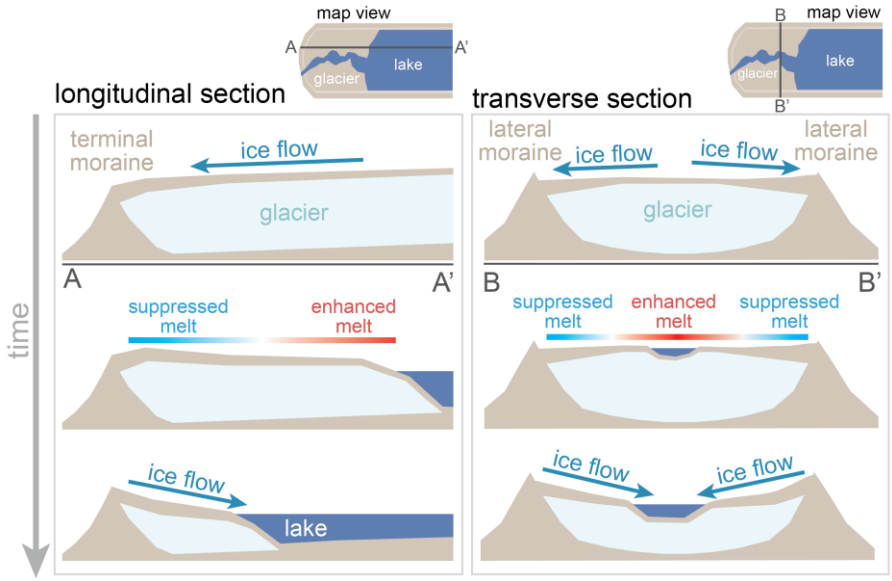

Figure B.10: Schematic showing topographic inversion at the Imja Lake moraine dam. Melt is enhanced at the lake edges through interaction with liquid water and ice cliff retreat. Melt is suppressed near the lateral moraines due to insulation from thicker debris and the effects of local shading.

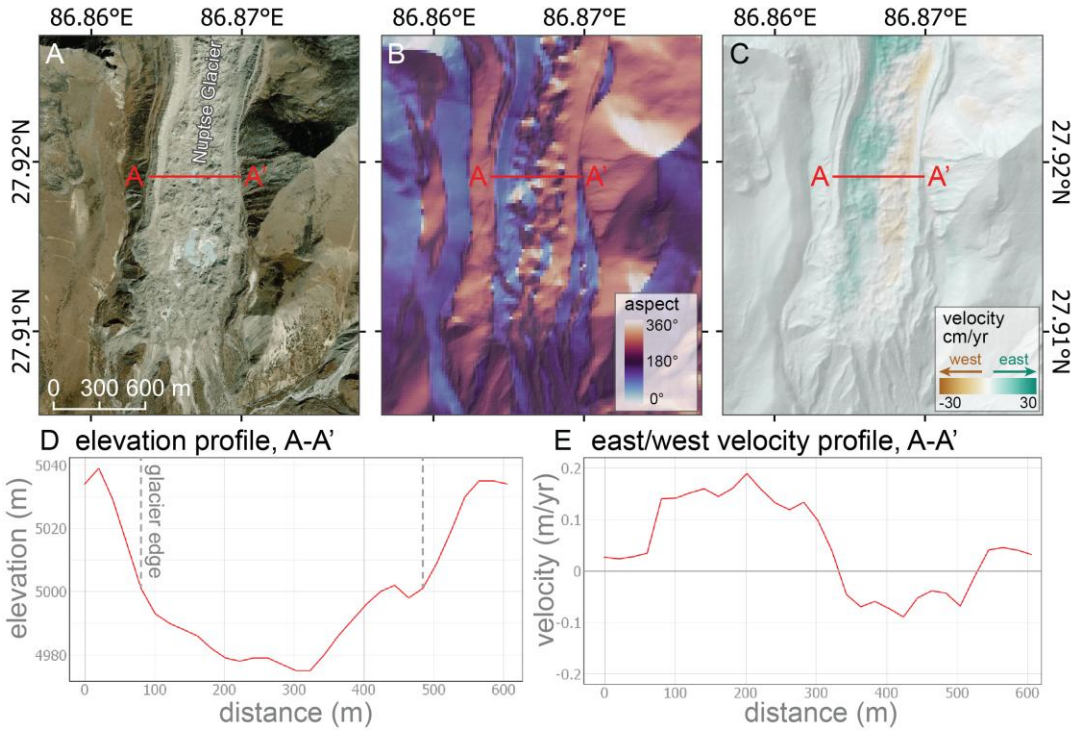

Figure B.11: The Lower Nuptse Glacier, located to the west-northwest of Imja Lake. Note the general surface aspect (B), concave-up transverse profile (D) and pattern of surface motion toward the glacier centerline (C/E).

Table B.1: Summary of published downwasting rates for ice-cored moraines. Expanded from Irvine-Fynn et al. (2011).

Downwasting rate (m/yr)	Location	Reference
<4.8	Jotunheimen, Norway and Kebnekaise, Sweden	Østrem (1965)
0.048-0.08	Southeastern Alaska, USA	Mckenzie (1969)
0.1	Yukon, Canada	Ross (1976)
0.003-0.3	Yukon, Canada	Driscoll (1980)
<5	Alberta, Canada	Mattson et al. (1991)
1.0	Yukon, Canada	Johnson (1992)
0.4-4.8	Punjab Himalaya, Pakistan	Mattson (1993)
0.05	Svalbard	King & Volk (1994)
0.1-0.4	Svalbard	Etzelmüller (2000)
0.3-2.5	Kötlujökull, Iceland	Krüger & Kjær (2000)
0.4-3.5	Svalbard	Lukas et al. (2005)
0.9	Svalbard	Schomacker & Kjær (2008)
0.7	Svalbard	Irvine-Fynn et al. (2011)
0.16	Svalbard	Tonkin et al. (2016)
0.5	South Tyrol, Italy	Kunz et al. (2021)

B.2 LOS decomposition bias

To investigate potential bias caused by the assumptions made for our LOS decomposition (Section 3.5.3), we used the observed north/south velocity measurements from the DEM-derived validation data (Section 3.5.4, Figure 3.3) to estimate the expected bias in the east/west and vertical velocity components from our decomposition. To do this, we solved the following equations for vertical and east/west velocity:

$$v_{asc} = \hat{l}_{asc} \cdot \hat{l}_{ud} * v_{ud} + \hat{l}_{asc} \cdot \hat{l}_{ew} * v_{ew} + \hat{l}_{asc} \cdot \hat{l}_{ns} * v_{ns} \quad \text{(B.3)}$$

$$v_{des} = \hat{l}_{des} \cdot \hat{l}_{ud} * v_{ud} + \hat{l}_{des} \cdot \hat{l}_{ew} * v_{ew} + \hat{l}_{asc} \cdot \hat{l}_{ns} * v_{ns} \quad \text{(B.4)}$$

Where $\hat{l}$ represents a unit vector, v is the mean velocity over the study period, and the subscripts (*asc*, *des*, *ud*, *ew*, and *ns*) correspond to the ascending, descending, up/down, east/west, and north/south components, respectively. The v_{asc} and v_{des} values are the mean ascending and descending LOS velocity magnitude from our combined InSAR and feature tracking time series

approach. For this analysis we substitute the north/south velocity from the DEM-derived validation data as the v_{ns} value in each equation. Thus, rather than assuming that the v_{ns} north/south contribution to the LOS velocity is 0, as in Section 3.5.3, here we include the observed north/south velocity magnitude when solving for the v_{ud} and v_{ew} components (Figure B.12).

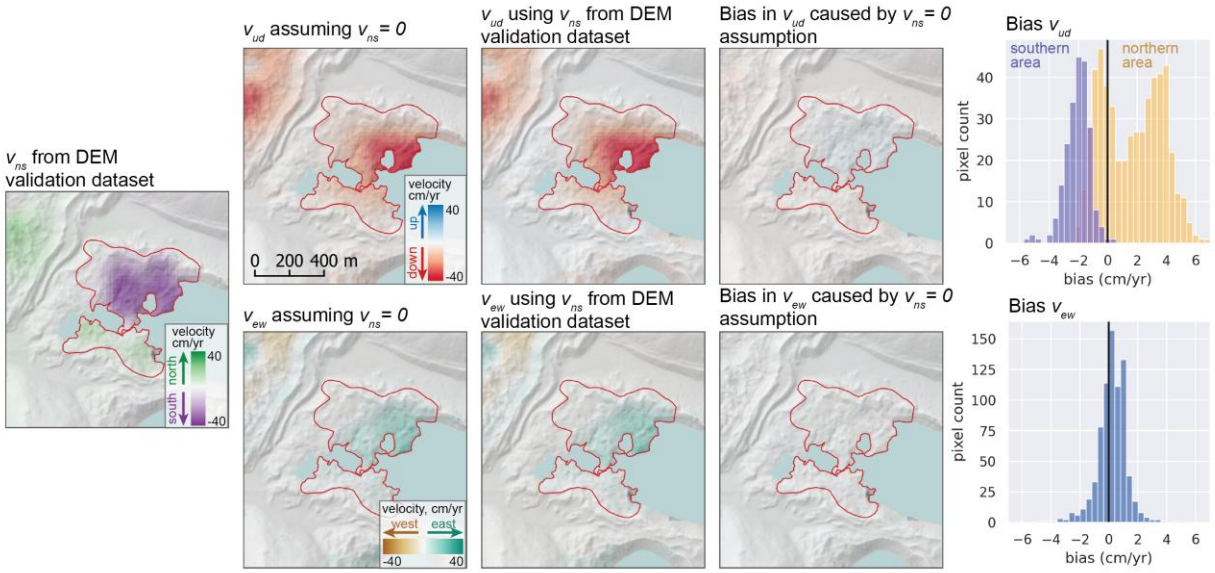

Figure B.12: Estimated bias in up/down and east/west velocity components caused by ignoring the north/south velocity in the LOS decomposition. *v* is the mean velocity over the study period, and the subscripts (*ud*, *ew*, and *ns*) correspond to the up/down, east/west, and north/south components, respectively. (left) The north/south velocity components from the DEM-derived validation dataset. (Second column) The up/down and east/west velocity components, assuming that the north/south velocity component is 0. (Center column) The up/down and east/west velocity components calculated using the north/south velocity component from the DEM-derived validation dataset. (Fourth column and right column) Difference maps and histograms showing estimated bias caused by neglecting the north/south velocity component. Expected bias over the moving area is small (~3-4%) compared to signal magnitude.

The DEM-derived validation data show that the northern portion of the moraine dam was moving southward, with a mean velocity of -19 cm/yr, while the southern portion of the moraine dam was moving northward, with a mean velocity of 7 cm/yr. Including these displacements during the LOS decomposition leads to a slight underestimation of the vertical velocity magnitude over the northern portion (mean bias of 1.6 cm/yr or ~10%) and a slight overestimation of the vertical

velocity magnitude over the southern portion (mean bias of -2.1 cm/yr or ~35%). When averaged over the entire moving area, these biases largely cancel, with a mean total bias of +0.6 cm/yr, or ~4% of the observed mean vertical velocity from the DEM-derived validation data. In the east/west direction, including the true north/south displacements results in a slight overestimation of eastward velocity (mean bias of 0.2 cm or 3% of the observed mean east/west velocity from the DEM-derived validation data.

We next address the question of whether the "downward signal" observed during the warm season could be caused by north/south displacement. To do this, we calculated the change in north/south velocity that would be required to produce the observed magnitude of seasonal change in vertical velocity (-8.0 cm/yr). We first subtracted the January/February vertical velocity (Figure 3.8) from our September/October vertical velocity to quantify the observed seasonal change in vertical velocity. We then projected this observed vertical seasonal change into the ascending LOS, and calculated the theoretical magnitude of north/south velocity change that would be required to achieve the same change in ascending LOS velocity.

The seasonal change in velocity over the moving area can either be explained by a -8.0 cm/yr change in the mean vertical velocity or a +62 cm/yr change in the mean north/south velocity (Figure B.13). A much larger change in the north/south velocity is needed because Sentinel-1 LOS measurements are largely insensitive to north/south motion. Within the moving area, some pixels would require a northward change in velocity of more than 3 m/yr. The -8.0 cm/yr mean vertical change is much more realistic, and can be explained by the well-documented physical process of seasonal ice melt in these areas (Irvine-Fynn et al., 2011). The alternative explanation would require a large seasonal transition to *rapid*, *upslope* motion over most of the moraine dam moving area, which has no physical basis and is inconsistent with the DEM-derived validation data.

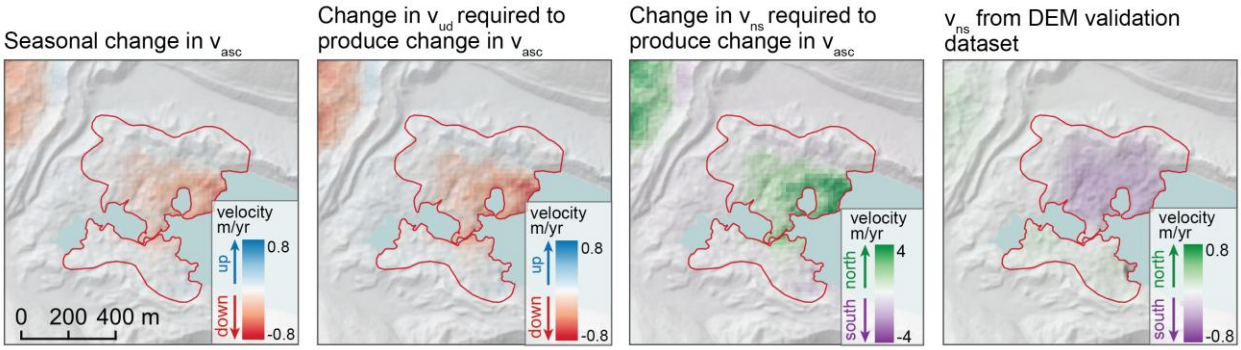

Figure B.13: Observed seasonal change in ascending LOS velocity (left) and the corresponding change in vertical (second left) or north/south velocity (second right) required to produce the same LOS seasonal change. *v* is velocity, and the subscripts (*asc*, *ud*, and *ns*) correspond to the ascending LOS, up/down, and north/south directions, respectively. Note the expanded color bar range for the north/south velocity plot. The required change in the vertical direction has a much smaller magnitude than the required change in velocity in the north/south direction, as the LOS vector has a small northward component. The required change in velocity in the north/south direction is inconsistent with the observed north/south velocity from the DEM-derived validation data (right).

Appendix C

SUPPLEMENTARY MATERIAL FOR CHAPTER 3: MAPPING BURIED ICE, QUANTIFYING SURFACE MOTION, AND IMPROVING GLOF HAZARD ASSESSMENTS FOR GLACIAL LAKE MORAINE DAMS IN NEPAL WITH SENTINEL-1 INSAR

Table C.1: SAR datasets for each glacial like identified as a high hazard by Mool et al. (2001) and ICIMOD (2011).

ICIMOD ID	common name	ascending		descending	
		burst ID	acquisitions, interferograms	burst ID	acquisitions, interferograms
kotak_gl_0009	Rolpa	012_023790_IW1	220, 654	121_258662_IW2	233, 690
koaru_gl_0009	Barun	012_023789_IW1	222, 660	121_258662_IW1	233, 693
kodud_gl_0184	Imja	012_023790_IW1	220, 654	121_258661_IW2	234, 688
kodud_gl_0036	Lumding	012_023789_IW1	222, 660	121_258662_IW2	233, 690
kodud_gl_0242	Chamlang South	012_023789_IW1	222, 660	121_258662_IW2	233, 690
gamar_gl_0018	Thugali	085_180592_IW1	221, 657	019_039569_IW2	230, 684
kotam_gl_0135	Nagma	012_023788_IW2	221, 657	048_101863_IW3	209, 621
kodud_gl_0241	Chamlang North	012_023789_IW1	222, 660	121_258662_IW2	233, 690
gakal_gl_0004	Muli Tal	158_337392_IW2	221, 657	092_196365_IW1	222, 661
kodud_gl_0012	Barun Pokhari	012_023789_IW1	222, 660	121_258662_IW1	233, 693
kodud_gl_0238	East Hongu 1	012_023789_IW1	222, 660	121_258662_IW2	233, 690
gabud_gl_0009	Birendra	085_180593_IW1	221, 657	019_039569_IW2	230, 684
kodud_gl_0220	Mera	012_023789_IW1	222, 660	121_258662_IW2	233, 690
koaru_gl_0016	Upper Barun	012_023789_IW1	222, 660	121_258662_IW1	233, 693
gakal_gl_0008	--	158_337392_IW3	221, 657	092_196364_IW1	222, 660
kotam_gl_0111	--	012_023788_IW2	221, 657	121_258661_IW1	233, 693
kodud_gl_0239	East Hongu 2	012_023789_IW1	222, 660	121_258662_IW2	233, 690
gakal_gl_0022	Tallo Kekyap	158_337393_IW3	221, 657	092_196363_IW1	222, 660
gakal_gl_0023	Kekyap	158_337393_IW3	221, 657	092_196363_IW1	222, 660
kodud_gl_0049	Dig	012_023790_IW1	220, 654	121_258662_IW2	233, 690
kodud_gl_0193	Sabai	012_023789_IW1	222, 660	121_258662_IW2	233, 690
kodud_gl_0205	--	012_023789_IW1	222, 660	121_258662_IW2	233, 690
kodud_gl_0229	Hongu 1	012_023789_IW1	222, 660	121_258662_IW2	233, 690

Table C.2: High coherence and low-coherence periods for each moraine dam.

ICIMOD ID	common name	Impacted by snow?	High-coherence period (DOY)	Low-coherence period (DOY)	Snowy period (DOY)
kotak_gl_0009	Rolpa	yes	100-160	220-280	340-100
koaru_gl_0009	Barun	yes	120-180	220-280	340-120
kodud_gl_0184	Imja	no	0-100	220-280	--
kodud_gl_0036	Lumding	yes	160-200	230-290	320-160
kodud_gl_0242	Chamlang South	no	0-100	230-290	--
gamar_gl_0018	Thugali	yes	120-180	220-280	340-120
kotam_gl_0135	Nagma	yes	130-200	240-300	340-130
kodud_gl_0241	Chamlang North	no	0-100	230-290	--
gakal_gl_0004	Muli Tal	yes	120-200	220-280	260-120
kodud_gl_0012	Barun Pokhari	yes	120-180	230-290	340-120
kodud_gl_0238	East Hongu 1	no	0-100	230-290	--
gabud_gl_0009	Birendra	yes	140-200	230-290	330-140
kodud_gl_0220	Mera	yes	120-180	230-290	340-120
koaru_gl_0016	Upper Barun	yes	120-180	230-290	300-100
gakal_gl_0008	--	yes	120-180	230-290	300-100
kotam_gl_0111	--	yes	150-210	230-290	330-140
kodud_gl_0239	East Hongu 2	no	0-100	230-290	--
gakal_gl_0022	Tallo Kekyap	no	0-180	230-290	--
gakal_gl_0023	Kekyap	no	0-180	230-290	--
kodud_gl_0049	Dig	yes	100-180	230-290	330-120
kodud_gl_0193	Sabai	yes	120-180	230-290	--
kodud_gl_0205	--	yes	130-190	230-290	330-130
kodud_gl_0229	Hongu 1	yes	100-180	230-290	330-100

Table C.3: Characteristics of moraine dam moving areas.

ICIMOD ID	Common name	Moraine dam area (km^2)	Moving area (km^2)	Percentage moving (%)
kotak_gl_0009	Rolpa	0.17	0.01	6
koaru_gl_0009	Barun	0.22	0.08	38
kodud_gl_0184	Imja	0.50	0.29	58
kodud_gl_0036	Lumding	0.06	0.02	33
kodud_gl_0242	Chamlang South	0.56	0.24	42
gamar_gl_0018	Thugali	0.32	0.13	41
kotam_gl_0135	Nagma	0.28	0.07	25
kodud_gl_0241	Chamlang North	0.25	0.14	53
gakal_gl_0004	Muli Tal	0.02	0.01	84
kodud_gl_0012	Barun Pokhari	0.11	0.03	28
kodud_gl_0238	East Hongu 1	0.43	0.29	67
gabud_gl_0009	Birendra	0.10	--	0
kodud_gl_0220	Mera	0.11	0.04	33
koaru_gl_0016	Upper Barun	0.05	0.02	33
gakal_gl_0008	--	0.11	0.09	78
kotam_gl_0111	--	0.10	--	0
kodud_gl_0239	East Hongu 2	0.08	0.05	68
gakal_gl_0022	Tallo Kekyap	0.07	0.06	88
gakal_gl_0023	Kekyap	0.03	0.01	33
kodud_gl_0049	Dig	0.23	0.01	3
kodud_gl_0193	Sabai	0.22	--	0
kodud_gl_0205	--	0.23	0.03	12
kodud_gl_0229	Hongu 1	0.05	--	0

Table C.4: Moraine dam topographic form and buried ice area.

Common name	ICIMOD ID	Topographic form	Buried ice area (km^2)	Buried ice area (% of moraine dam area)	Rounce et al. (2017) ice core
Rolpa	kotak_gl_0009	narrow, sharp-crested	0.0148	8.8	Yes
Barun	koaru_gl_0009	wide, rounded	0.0836	37.2	Yes
Imja	kodud_gl_0184	wide, rounded	0.2976	59.3	Yes
Lumding	kodud_gl_0036	wide, rounded	0.0324	52.0	No
Chamlang South	kodud_gl_0242	narrow, sharp-crested	0.2692	48.5	Yes
Thugali	gamar_gl_0018	wide, rounded	0.1000	31.5	Yes
Nagma	kotam_gl_0135	narrow, sharp-crested	0.0548	19.9	No
Chamlang North	kodud_gl_0241	wide, rounded	0.0836	32.9	No
Muli Tal	gakal_gl_0004	wide, rounded	0.0096	56.8	No
Barun Pokhari	kodud_gl_0012	narrow, sharp-crested	0.0076	7.0	Yes
East Hongu 1	kodud_gl_0238	wide, rounded	0.1344	31.6	Yes
Birendra	gabud_gl_0009	narrow, sharp-crested	0.0004	0.4	No
Mera	kodud_gl_0220	narrow, sharp-crested	0.0052	4.8	Yes
Upper Barun	koaru_gl_0016	wide, rounded	0.0072	14.2	No
--	gakal_gl_0008	wide, rounded	0.0700	61.8	No
--	kotam_gl_0111	narrow, sharp-crested	0.0028	2.8	No
East Hongu 2	kodud_gl_0239	wide, rounded	0.0156	19.6	No
Tallo Kekyap	gakal_gl_0022	wide, rounded	0.0000	0.0	No
Kekyap	gakal_gl_0023	wide, rounded	0.0000	0.0	No
Dig	kodud_gl_0049	narrow, sharp-crested	0.0000	0.0	No
Sabai	kodud_gl_0193	narrow, sharp-crested	0.0312	14.4	No
--	kodud_gl_0205	narrow, sharp-crested	0.0080	3.5	Yes
Hongu 1	kodud_gl_0229	wide, rounded	0.0004	0.8	Yes

Table C.5: Summary statistics for vertical and East/West velocities in moraine dam moving areas and corresponding stable areas over the 7.37-year period from January 1, 2017 to May 15, 2024. Directional coherence is the proportion of pixels sharing the majority velocity sign. Primary movement direction is assigned based on this majority sign; when directional coherence is below 0.6, movement is classified as "Mixed." Stable-area mean and standard deviation provide estimates of residual bias and measurement dispersion.

Common name	ICIMOD ID	Vertical velocity (cm/yr)								East/West velocity (cm/yr)							
		Moving area						Stable area		Moving area						Stable area	
		Primary direction	Directional coherence	Mean	Std dev	5th percentile	95th percentile	Mean	Std dev	Primary direction	Directional coherence	Mean	Std dev	5th percentile	95th percentile	Mean	Std dev
Rolpa	kotak_gl_0009	Subsidence	1.00	-1.5	1.1	-3.0	-0.5	0.1	0.5	West	1.00	-3.7	1.5	-5.0	-1.7	0.7	0.5
Barun	koaru_gl_0009	Subsidence	1.00	-13.9	8.0	-29.6	-4.1	-2.1	0.4	West	0.67	-1.3	7.6	-11.5	13.9	-2.3	0.6
Imja	kodud_gl_0184	Subsidence	0.96	-13.2	9.4	-28.7	-0.3	1.2	0.3	Mixed	0.58	5.3	11.0	-5.5	30.5	0.4	0.4
Lumding	kodud_gl_0036	Subsidence	1.00	-7.2	4.6	-14.7	-1.0	1.7	0.6	East	0.91	3.3	2.7	-0.4	8.1	1.5	0.8
Chamlang South	kodud_gl_0242	Subsidence	0.94	-4.4	3.4	-10.9	0.1	-1.2	0.4	West	0.80	-1.9	2.6	-7.2	1.5	-1.3	0.7
Thugali	gamar_gl_0018	Subsidence	0.91	-11.3	10.8	-28.9	6.1	0.1	0.5	West	0.85	-16.2	17.2	-51.8	6.5	-1.7	0.6
Nagma	kotam_gl_0135	Subsidence	0.89	-4.3	5.3	-16.7	1.9	0.4	0.2	West	0.87	-5.2	8.1	-27.8	1.2	0.2	0.3
Chamlang North	kodud_gl_0241	Subsidence	0.92	-3.3	3.3	-9.2	0.4	-1.1	0.4	East	0.87	2.4	3.4	-3.3	7.1	1.4	0.6
Muli Tal	gakal_gl_0004	Subsidence	0.85	-5.3	4.8	-11.1	4.3	-0.2	0.8	Mixed	0.52	-0.2	4.8	-7.7	6.7	2.2	0.5
Barun Pokhari	kodud_gl_0012	Subsidence	1.00	-4.2	2.5	-8.8	-0.8	-0.4	0.6	Mixed	0.60	0.0	3.0	-3.9	5.3	0.0	1.0
East Hongu 1	kodud_gl_0238	Subsidence	1.00	-6.2	3.1	-11.2	-1.7	-1.1	0.4	East	0.94	5.7	5.0	-0.3	16.0	1.4	0.6
Mera	kodud_gl_0220	Subsidence	1.00	-1.6	0.8	-3.3	-0.6	-1.2	0.6	West	0.96	-1.5	0.9	-3.0	-0.3	-0.5	0.7
Upper Barun	koaru_gl_0016	Subsidence	1.00	-4.1	1.1	-6.0	-2.4	-2.0	0.6	East	0.69	1.3	2.0	-1.4	4.6	0.8	0.8
--	gakal_gl_0008	Subsidence	0.82	-2.1	2.5	-6.6	1.8	1.0	0.3	Mixed	0.54	1.1	5.5	-8.0	10.8	0.5	0.3
East Hongu 2	kodud_gl_0239	Subsidence	1.00	-3.2	0.8	-4.3	-1.9	-2.3	0.5	East	0.75	0.8	1.3	-0.8	3.0	-1.3	0.8
Tallo Kekyap	gakal_gl_0022	Subsidence	0.97	-1.2	0.7	-2.5	-0.2	0.2	0.3	West	0.95	-1.5	1.0	-3.0	0.0	0.1	0.3
Kekyap	gakal_gl_0023	Subsidence	0.97	-0.7	0.4	-1.3	-0.1	-0.1	0.5	West	1.00	-1.5	1.1	-4.0	-0.5	-1.0	0.6
Dig	kodud_gl_0049	Subsidence	1.00	-2.8	0.8	-4.2	-2.0	-0.4	0.2	East	0.89	1.0	0.7	-0.3	1.8	-0.4	0.4
--	kodud_gl_0205	Subsidence	1.00	-5.7	3.1	-10.3	-1.8	-2.3	0.5	West	0.94	-2.4	1.9	-6.0	0.0	-0.1	0.7
Minimum		--	0.82	-13.9	0.4	-29.6	-4.1	-2.3	0.2	--	0.52	-16.2	0.7	-51.8	-1.7	-2.3	0.3
Mean		--	0.96	-5.1	3.5	-11.1	-0.1	-0.5	0.4	--	0.79	-1.5	4.1	-10.1	5.7	-0.1	0.6
Maximum		--	1.00	-0.7	10.8	-1.3	6.1	1.7	0.8	--	1.00	5.7	17.2	-0.3	30.5	2.2	1.0

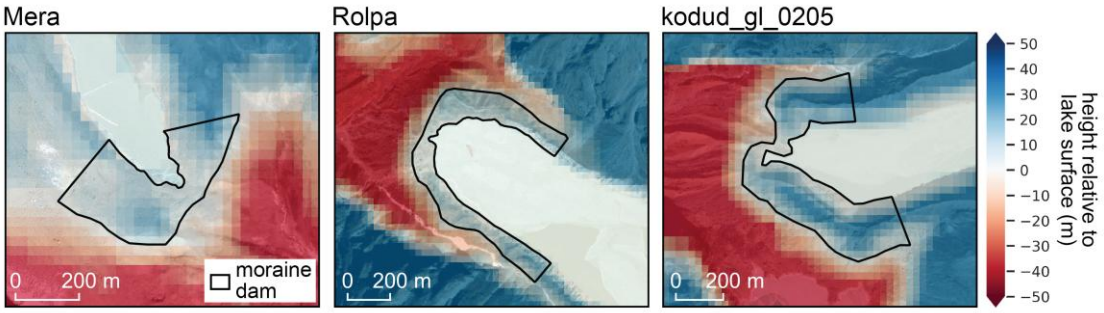

Figure C.1: Example maps of height relative to lake elevation from the GLO-30 DEM (European Space Agency, 2021). These maps, alongside satellite basemaps (Microsoft Bing; Google Satellite; ESRI World Imagery; all accessed May-June 2024) were used to identify the non-bedrock material separating each lake from the downstream area.

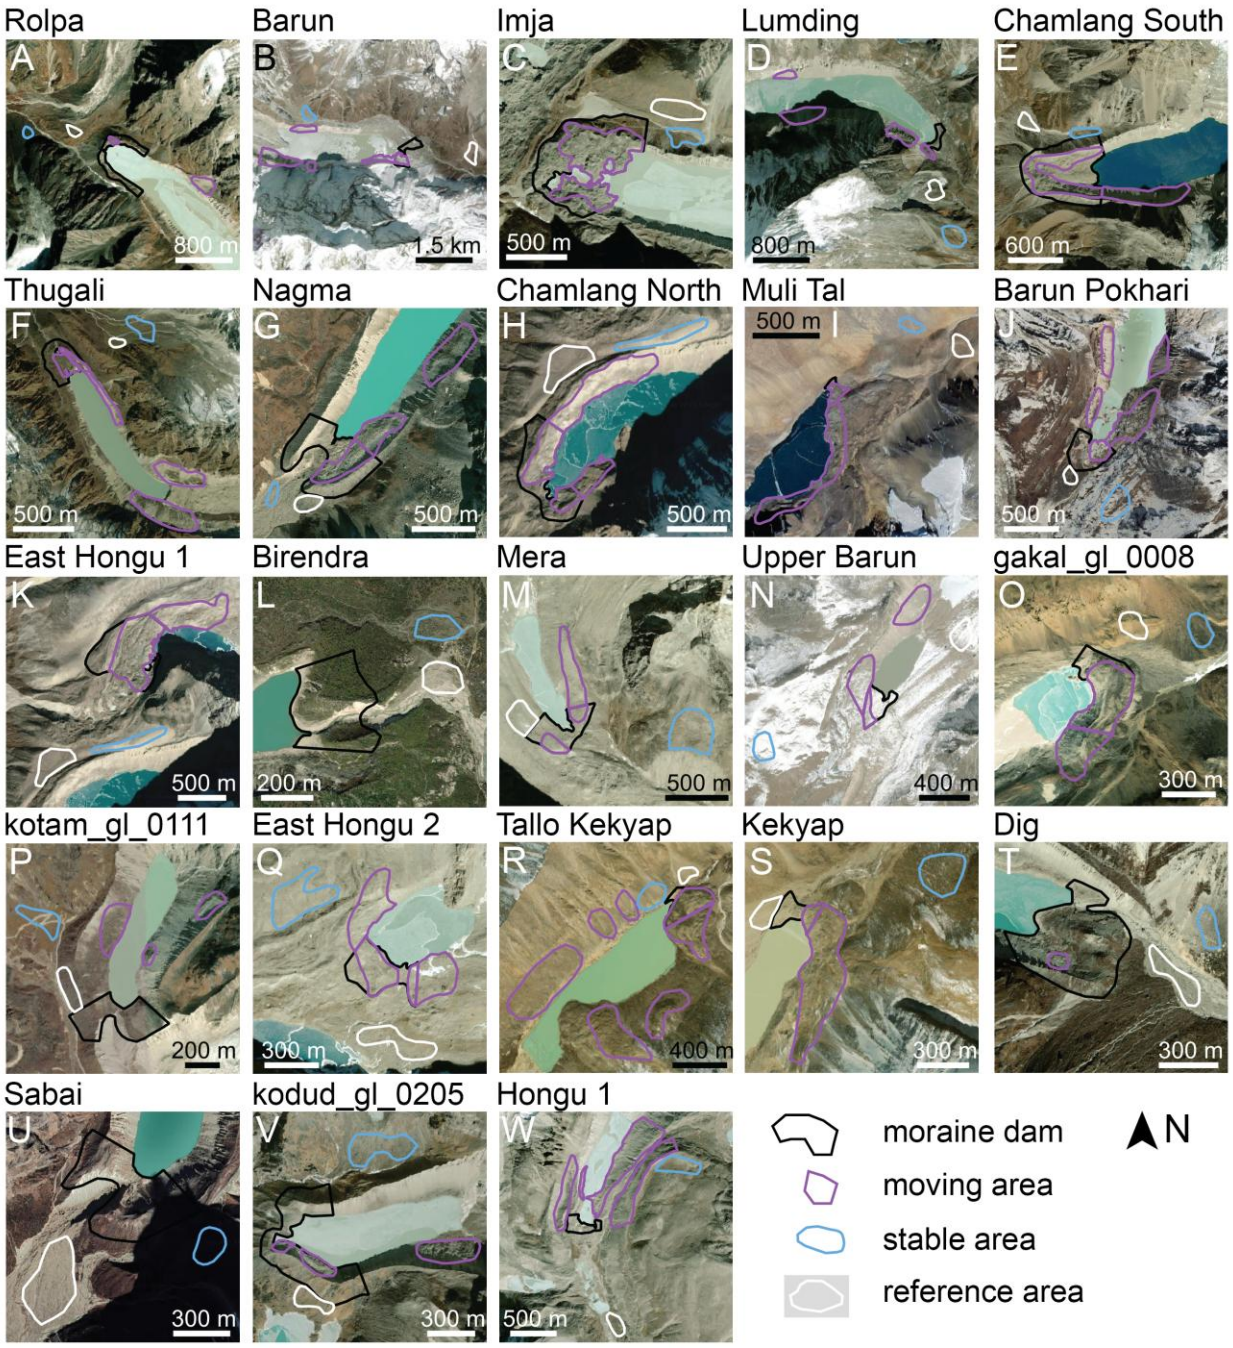

Figure C.2: Moraine dam, moving, stable, and reference areas for each moraine dam. Moving areas were defined where coherent LOS surface velocity was greater than or equal to 1 cm/yr. Here, we include moving areas along the lake bank. Reference and stable areas were selected over relatively flat surfaces with high year-round coherence and near-zero LOS velocity, as close as possible to the moraine dam.

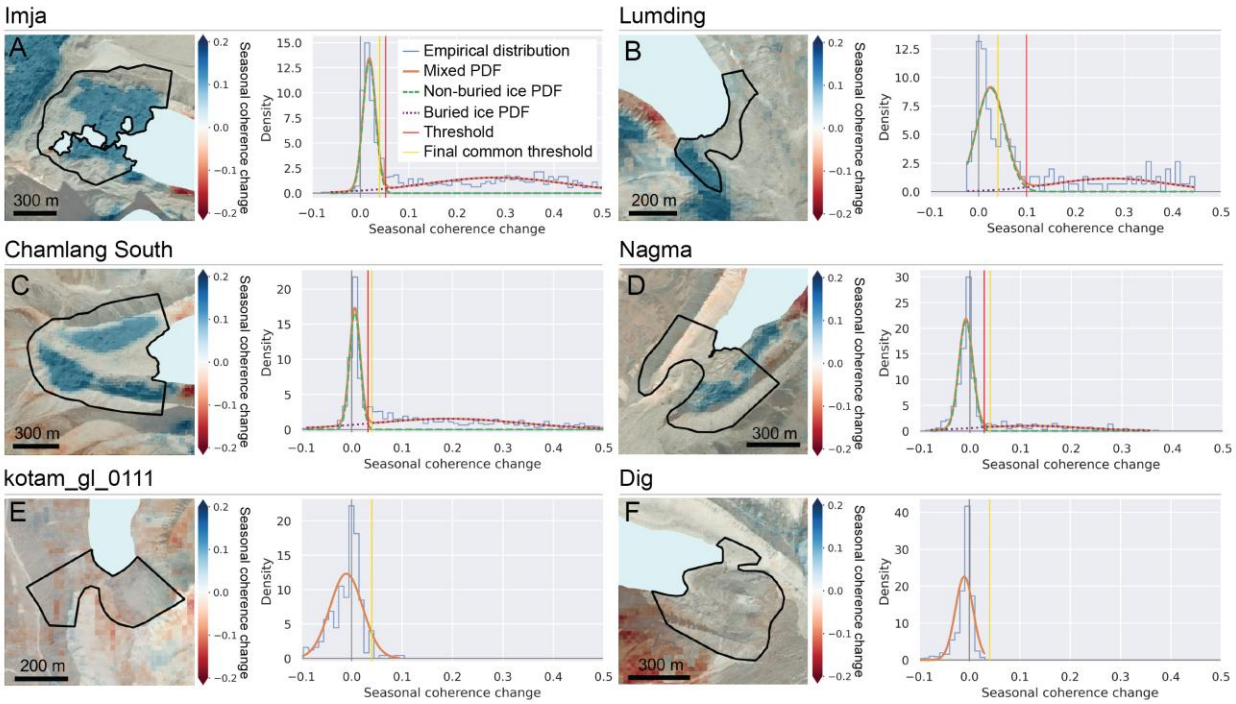

Figure C.3: Example of how seasonal coherence-change thresholds were identified to delineate buried-ice extent using Gaussian mixture modeling (GMM). The maps show the change in median coherence between the cold and warm seasons (cold - warm). Buried-ice pixels are expected to have positive values corresponding to a decrease in warm-season coherence, while non-buried-ice pixels follow an approximately Gaussian distribution centered near zero. Histograms show the empirical distribution of seasonal coherence change (blue) together with the Gaussian components representing the non-buried-ice and buried-ice subpopulations (green and orange) and their summed mixture (red). For moraine dams where a buried-ice component was appropriate, the seasonal coherence threshold (red vertical line) was defined as the smallest value where the probability of belonging to the buried-ice component exceeded that of the non-buried ice component. The final common threshold (yellow line) is the median of the moraine-specific thresholds and was used to map buried-ice pixels across all sites.

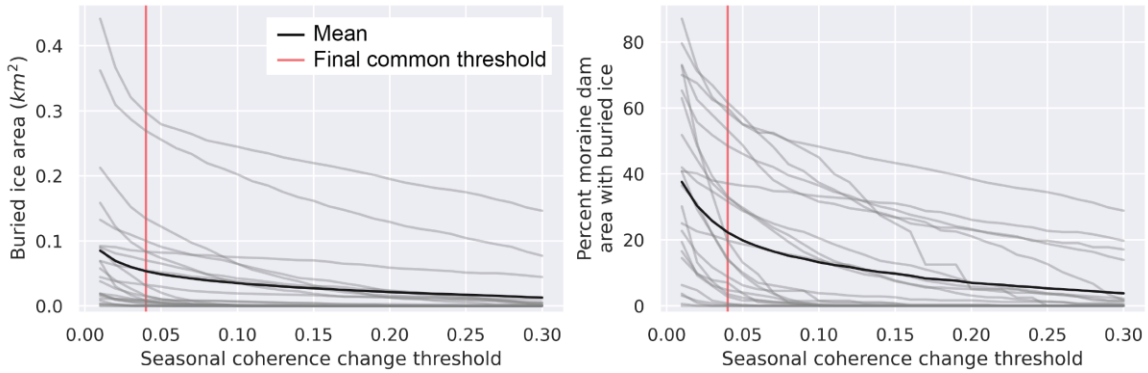

Figure C.4: Sensitivity analysis for seasonal coherence change threshold for identifying buried ice. Buried ice area and percent of moraine dam with buried ice are sensitive to seasonal coherence threshold choice. They are more sensitive to decreases in the threshold than increases in the threshold.

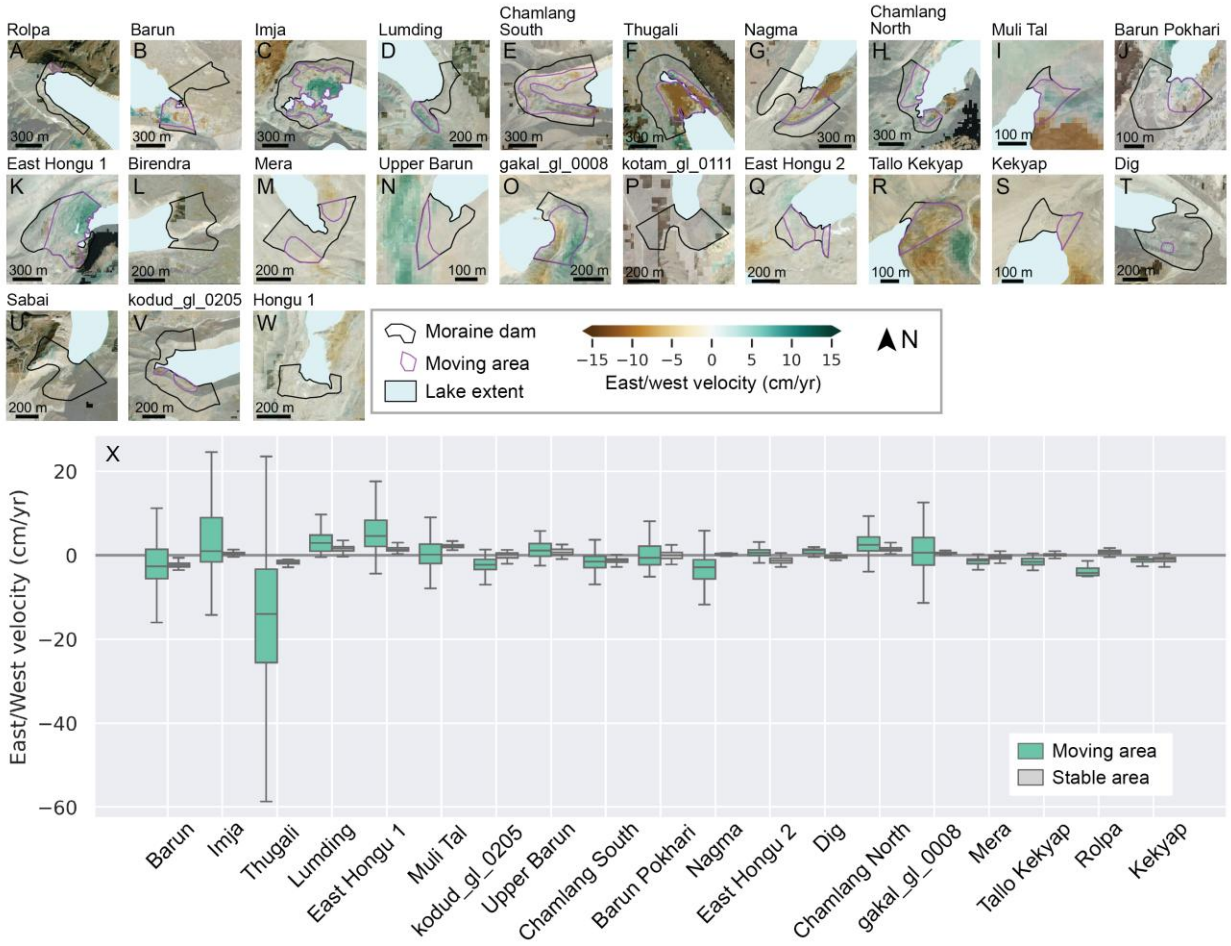

Figure C.5: Median horizontal velocity of moraine dams over the 7.37-year period from January 1, 2017 to May 15, 2024. The magnitude of horizontal velocity is often similar to the magnitude of subsidence (Figure C.5).

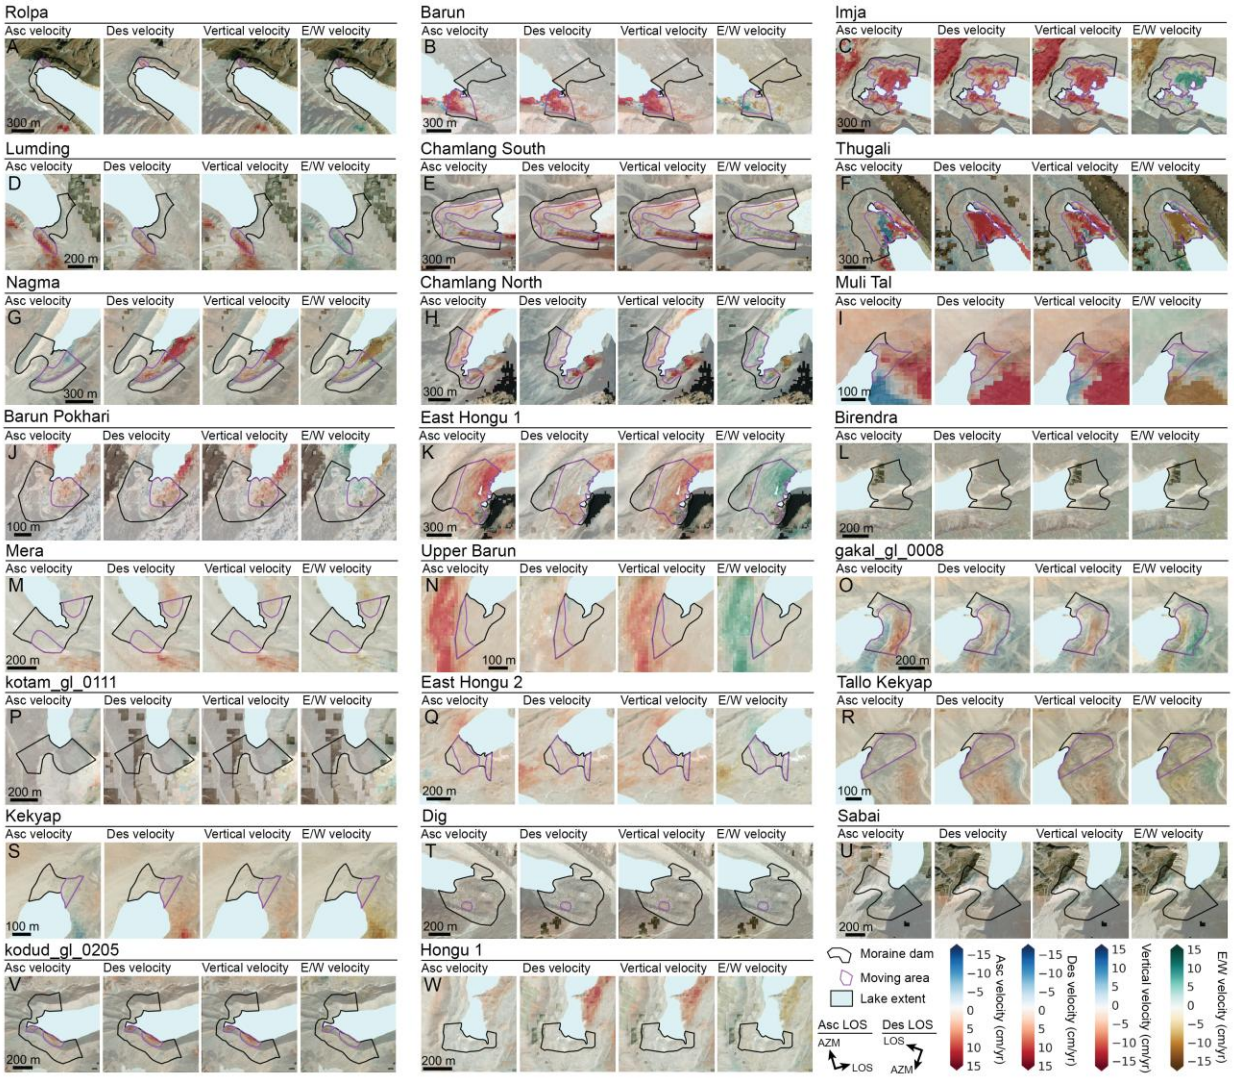

Figure C.6: Median ascending LOS, descending LOS, vertical, and East/West velocity of high-priority moraine dams in Nepal over the 7.37-year period from January 1, 2017 to May 15, 2024. Median vertical and East/West velocity are calculated from the median LOS velocity.

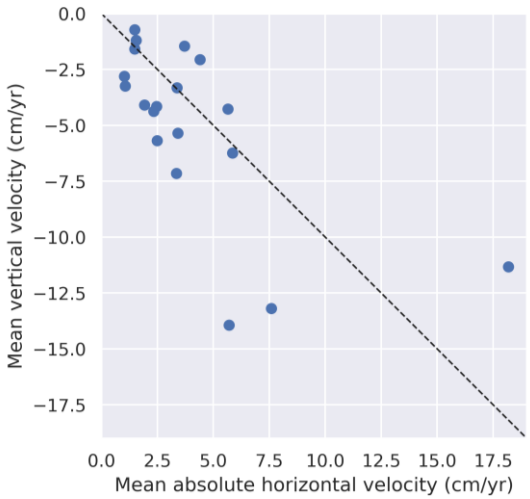

Figure C.7: The magnitude of velocity over the moraine dam moving areas is similar for the vertical and East/West components during the 7.37-year study period from January 1, 2017 to May 15, 2024.

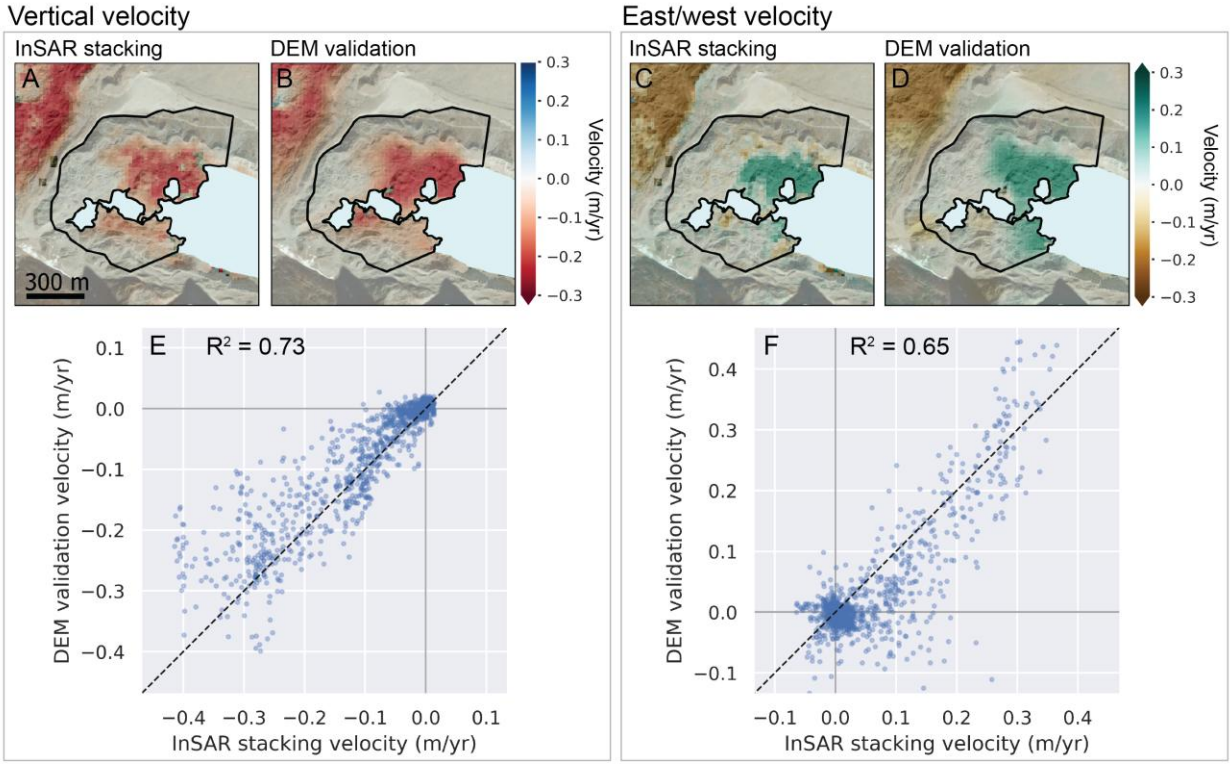

Figure C.8: Validation of the InSAR stacking approach employed to calculate median moraine dam velocity over the study period. Maps show the median velocity derived from InSAR stacking and digital elevation model (DEM) velocity products for the corresponding period from Brencher et al. (2025) at the Imja Lake moraine dam. The bottom panels show the agreement between pixels in the InSAR and DEM products.

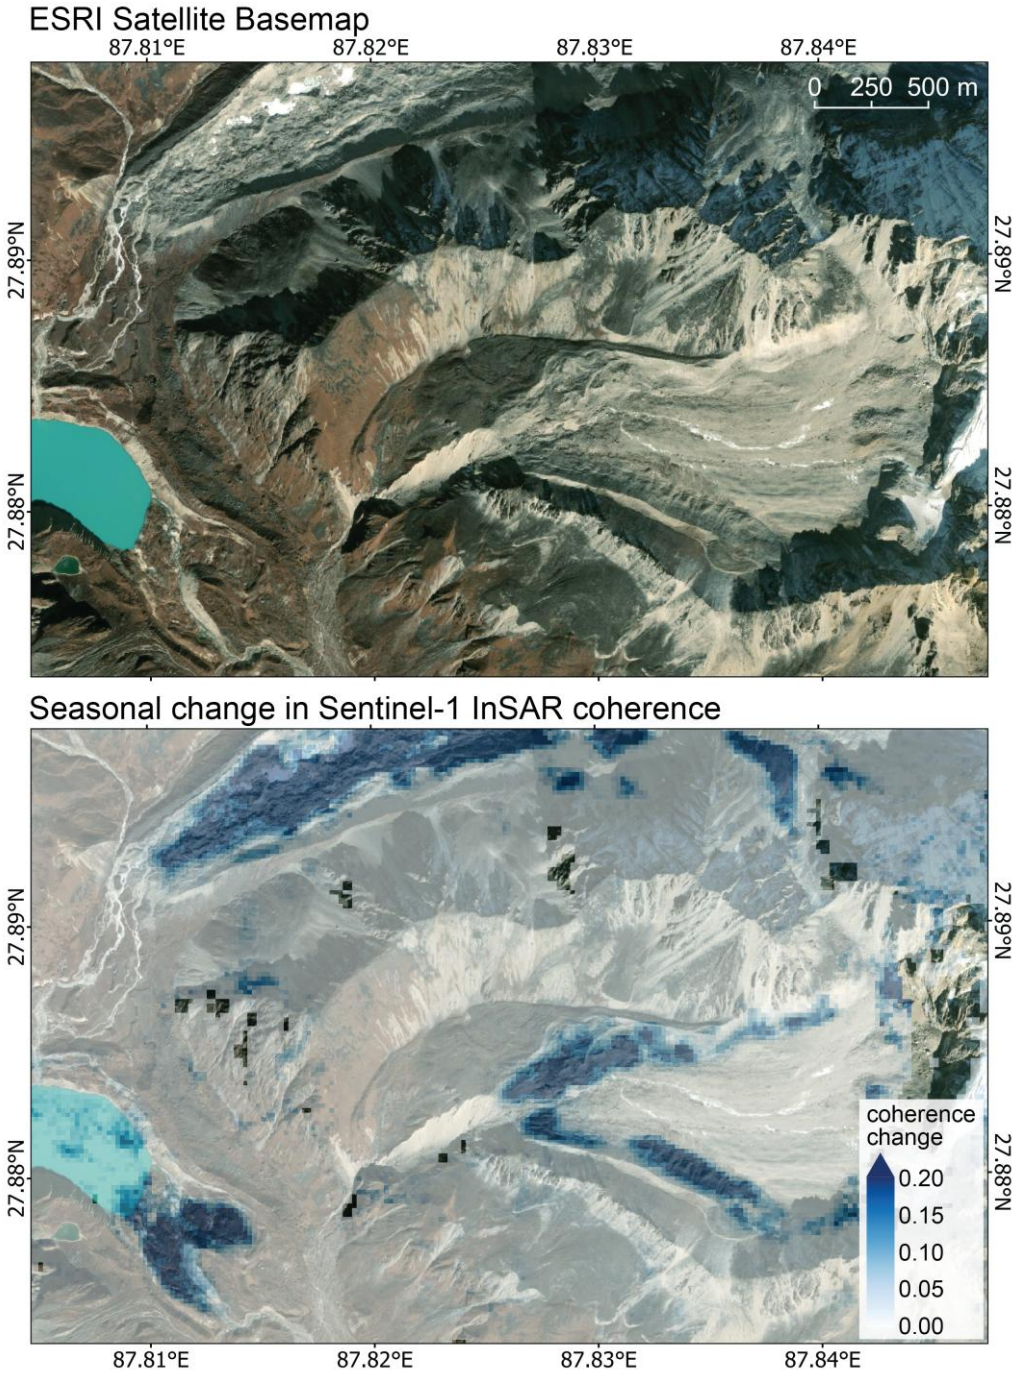

Figure C.9: Seasonal InSAR coherence change reveals the extent of buried ice in moraines, moraine dams, and debris-covered glaciers. The extent of buried ice is not obvious based on topography and surface morphology alone. Areas that do not contain buried ice do not typically experience a strong decrease in InSAR coherence during the warmer months.

Appendix D

SUPPLEMENTARY MATERIAL FOR CHAPTER 5: ACCURATE SNOW DEPTH PREDICTIONS ACROSS THE WESTERN U.S. USING A DEEP LEARNING MODEL TRAINED ON 7 YEARS OF AIRBORNE LIDAR SNOW DEPTH MEASUREMENTS

Table D.1: Hyperparameters for five models included in the CNN ensemble.

Learning rate	Weight decay	Epochs	Minimum test loss
0.00046	0.00011	62	0.00091
0.00025	0.00020	63	0.00091
0.00016	0.00013	98	0.00090
0.00012	0.00036	63	0.00092
0.00027	0.00009	85	0.00090

Table D.2: Test performance of individual CNNs trained without features from a single data source.

Data source removed	R^2	MSE	MAE	Mean of residuals (m)	STD of residuals (m)	Median relative error (%)	NMAD relative error (%)
None	0.87	0.31	0.36	0.01	0.55	20	21
COP30	0.83	0.39	0.41	0.03	0.63	23	24
Sentinel-2	0.83	0.40	0.42	0.06	0.63	26	28
Sentinel-1	0.86	0.32	0.36	0.10	0.55	21	22
SNODAS	0.79	0.48	0.46	0.25	0.65	27	24
FCF	0.86	0.31	0.36	0.08	0.55	21	22
DOWY	0.85	0.35	0.38	0.11	0.58	22	23

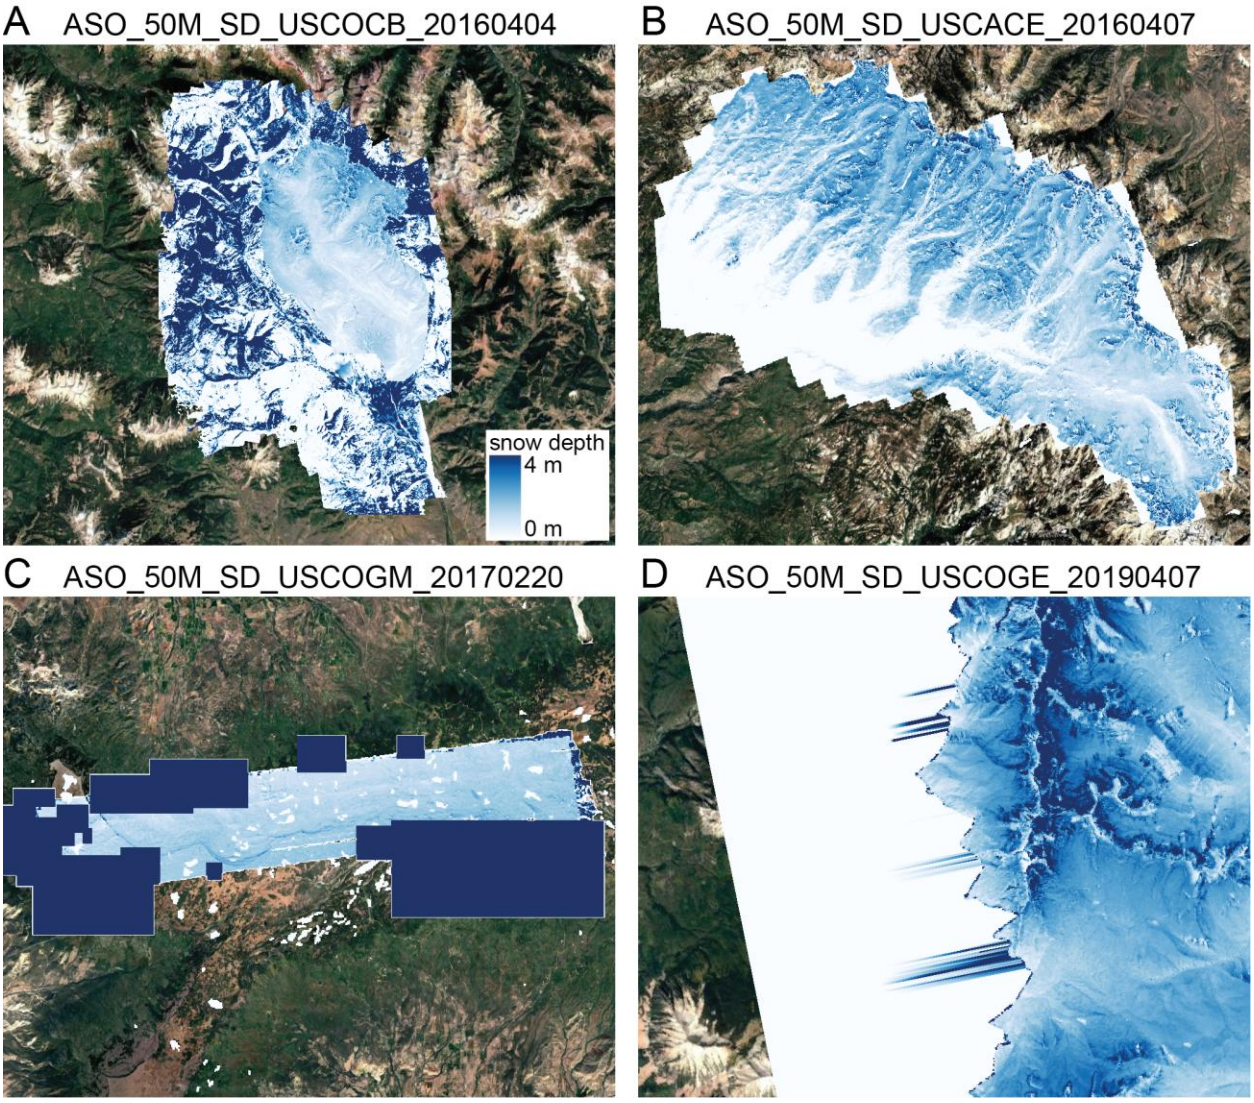

Figure D.1: Examples of artifacts observed impacting a subset of airborne lidar snow depth data sets. Snow depth products with clear, extensive artifacts from early ASO flights were removed from the training dataset (e.g. A, C). We manually defined clipping masks for other products with limited edge artifacts to preserve valid data regions (e.g. B, D).

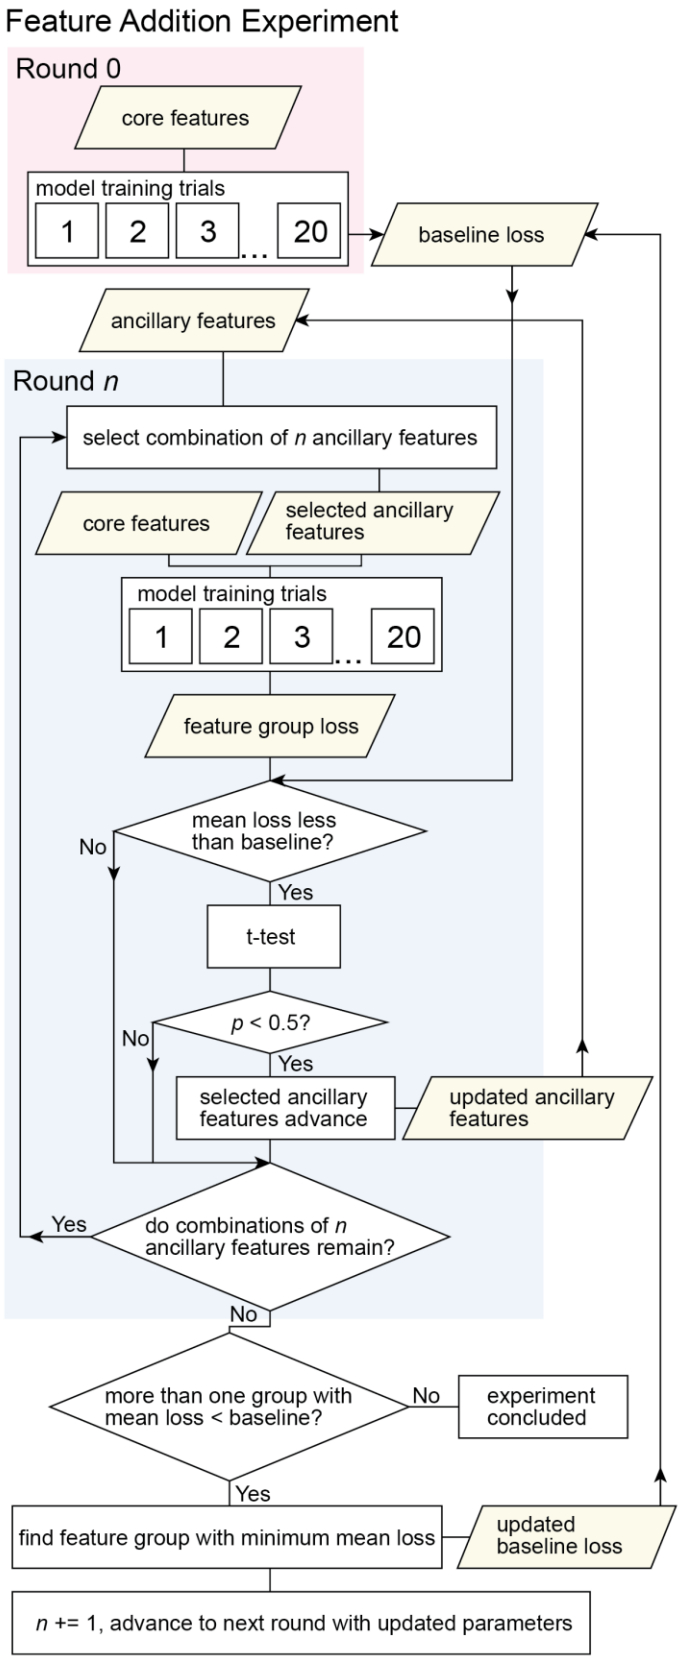

Figure D.2: Schematic showing feature addition experiment approach.

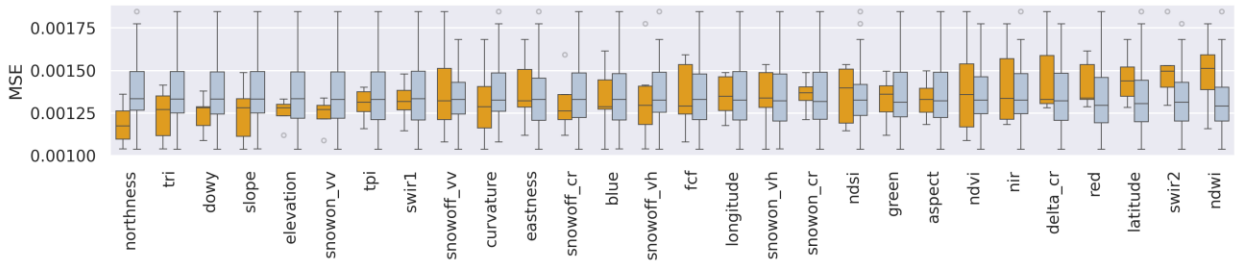

Figure D.3: Results of feature selection Monte Carlo experiment, where only groups already including SNODAS snow depth are included. When SNODAS snow depth is included, a different group of features appear to have high importance compared to when SNODAS snow depth may not be included (Figure 5.2).

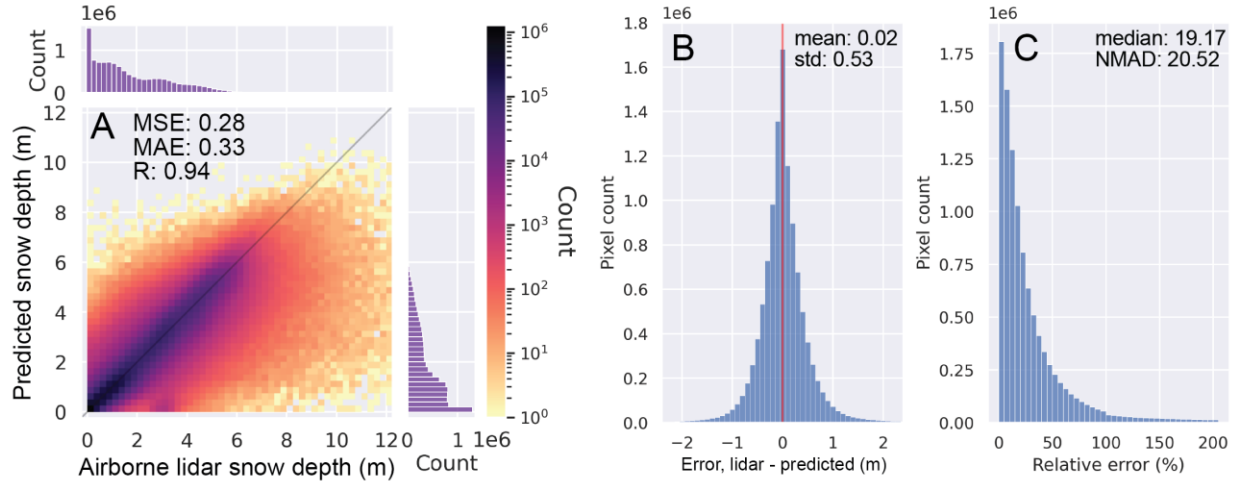

Figure D.4: CNN ensemble (Section 5.4.6) performance for snow-covered grid cells in the test dataset (n≈11.2M grid cells covering ~28,000 km^2). See Fig 5. Caption for details.

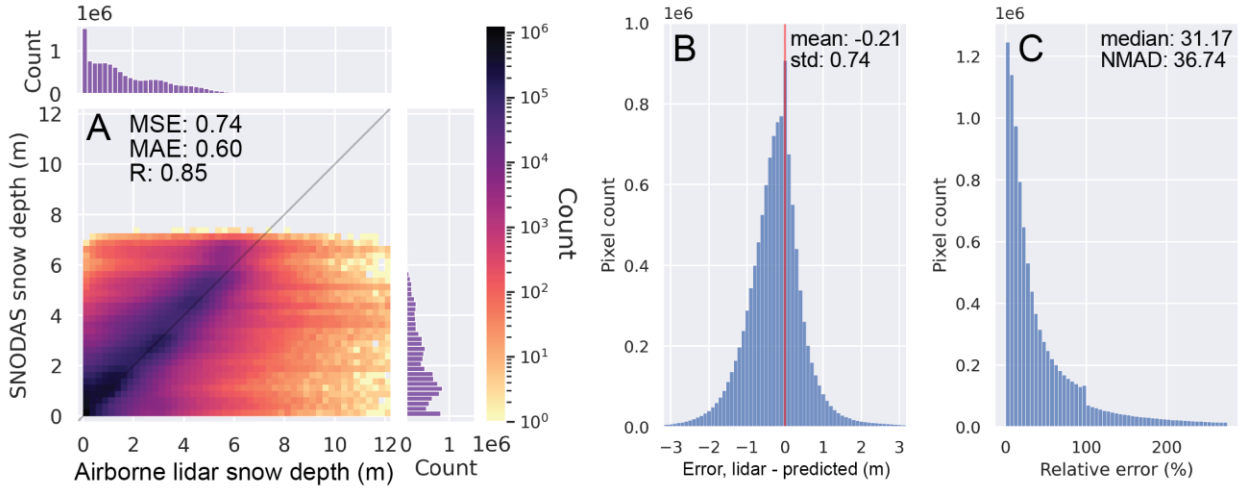

Figure D.5: Upsampled SNODAS snow depth performance on snow-covered pixels in the test dataset (n≈11.2M grid cells covering ~28,000 km^2). See Fig 5. Caption for details.

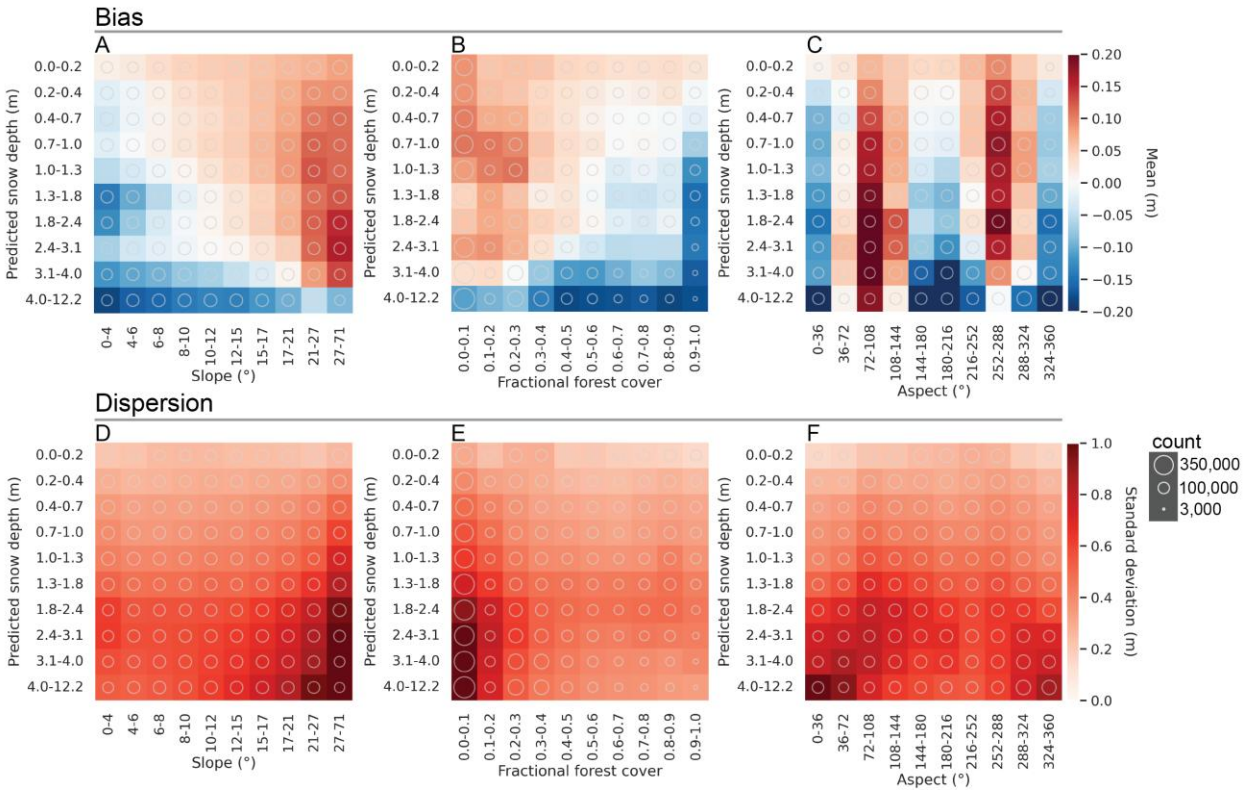

Figure D.6: Two-dimensional histograms showing bias (top) and dispersion (bottom) of prediction error for groups of A, D) predicted snow depth vs. terrain slope, B, E) predicted snow depth vs. fractional forest cover, and C, F) predicted snow depth vs. terrain aspect. Circle size corresponds to pixel count in each bin.

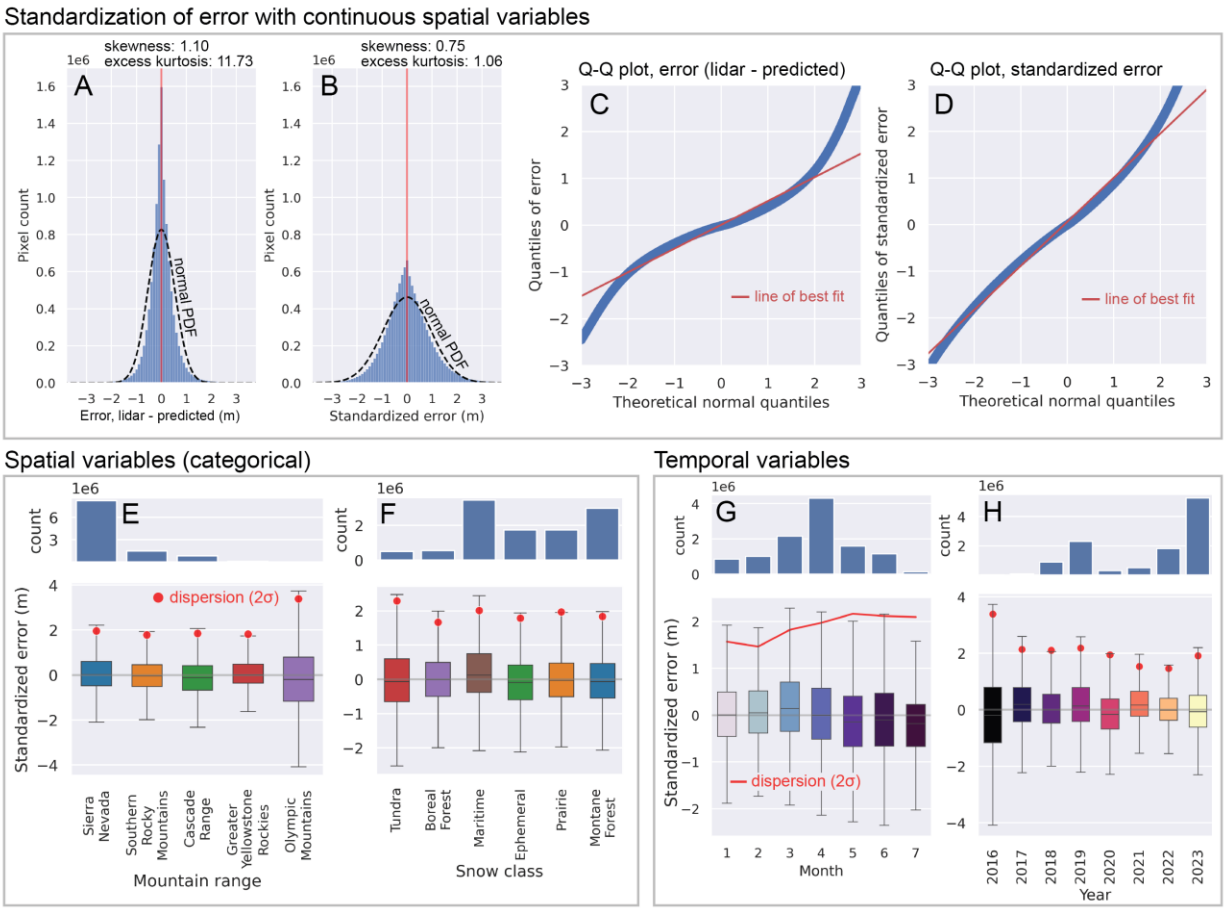

Figure D.7: Heteroscedasticity largely explains the departure of snow depth prediction error (airborne lidar minus predicted) from normality. After prediction error is standardized with slope, aspect, forest cover fraction, and predicted snow depth magnitude, the error distribution is close to normal. A, B) Distribution of prediction error before (A) and after (B) standardization. Note that the magnitude of skewness and excess kurtosis decrease following standardization. C, D) Quantile-Quantile (Q-Q) plot of prediction error distribution before (C) and after (D) standardization. The red lines show best fit. Following standardization, the best fit line is much closer to a 1:1 line. E, F) Standardized error for different mountain ranges (E) and snow classes (F), with accompanying histograms showing the number of pixels in each category. G, H) Standardized error for different months of the year (G) and for different years (H), with accompanying histograms showing the number of pixels in each time period. Note that standardizing the error eliminates many apparent differences in dispersion between groups, indicating that these differences can be explained by differences in slope, aspect, forest cover, and snow depth (Figure 5.6).

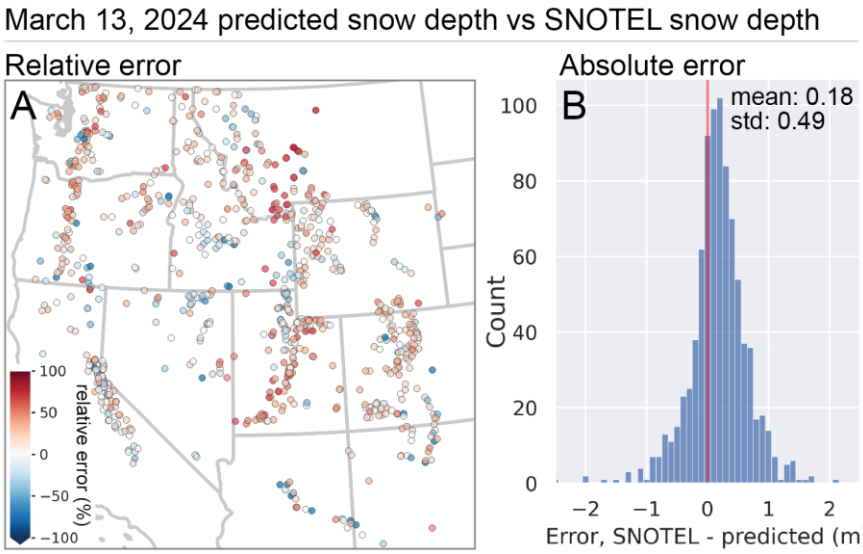

Figure D.8: A) Relative error for the March 13, 2024 predicted snow depth (Figure 5.7) vs. in-situ snow depth validation data from all available SNOTEL and CCSS stations. B) Distribution of prediction error (in situ minus predicted) for all snow monitoring stations.

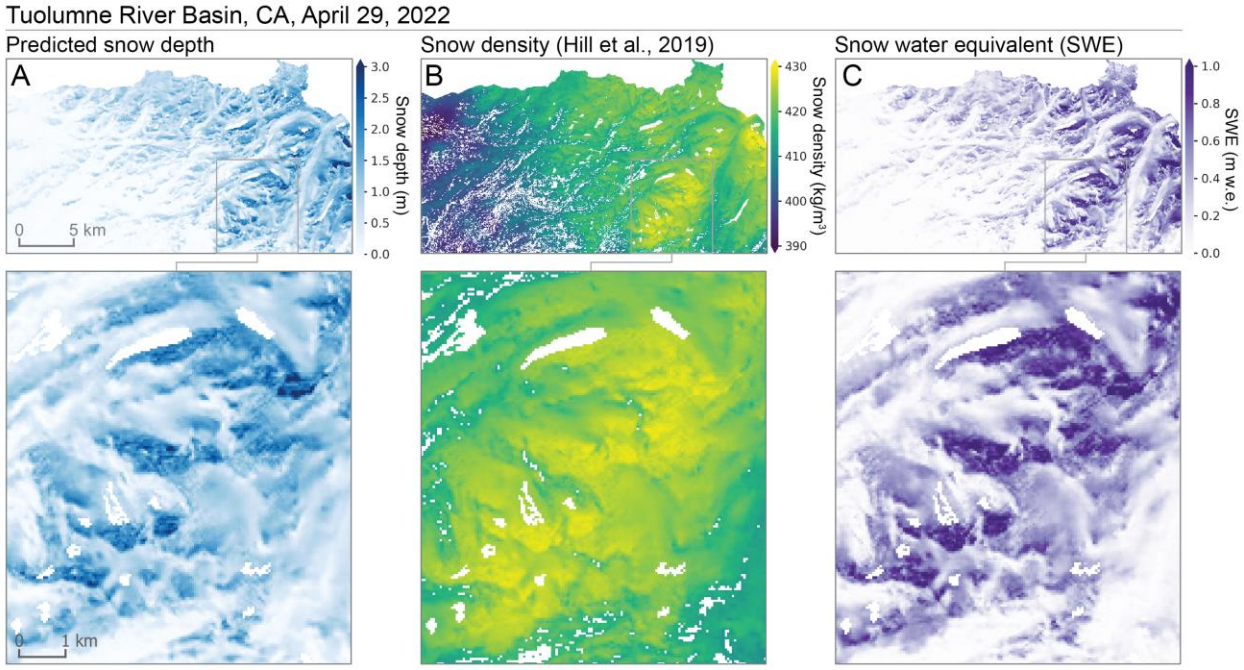

Figure D.9: Example snow water equivalent (SWE) calculation based on CNN snow depth prediction for an unseen test tile in the upper Tuolumne River Basin, CA on April 29, 2022 (See Figure 5.1 for context). A) Predicted CNN snow depth, B) Distributed snow density from the Hill et al. (2019) empirical density model and C) corresponding snow water equivalent estimate.

VITA

George (Quinn) Brencher is from Killingworth, CT, USA. His parents are both lawyers. Killingworth is a rural town boasting glaciated granite outcrops, an extensive network of streams, ponds, and wetlands, and an abundance of public land. Quinn developed an affinity for nature at an early age. He went to the public middle school, then to the Choate Rosemary Hall, a fancy boarding school. After high school, he and some friends went on a road trip to the West. The vast and rugged landscapes he saw there inspired him to think about the forces that shape the earth.

Quinn went to Middlebury College, which is a liberal arts college in Vermont. He studied biology and geology. He also really loved cartography. In his junior year, he met his now fiancé Micaela. After college, they moved to Bend, OR, then to Golden, CO. Micaela started medical school in Colorado, and Quinn moved to Seattle to work on a PhD at the University of Washington, marking the beginning of about four years of long-distance relationship. Now that the PhD is finished, they're reunited, getting married, and getting ready to move to wherever it is Micaela will be matched to for her residency. Quinn is excited to keep studying mountain environments for as long as they let him.