

©Copyright 2023

Connie Ebling Okasaki

Mechanistic Statistical Models of the Environment

Connie Ebling Okasaki

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

Doctor of Philosophy

University of Washington

2023

Reading Committee:
Andrew Berdahl, Chair
Mevin Hooten
Sándor Tóth

Program Authorized to Offer Degree:
Quantitative Ecology and Resource Management

University of Washington

Abstract

Mechanistic Statistical Models of the Environment

Connie Ebling Okasaki

Chair of the Supervisory Committee:
Assistant Professor Andrew Berdahl
School of Aquatic and Fisheries Sciences

Statistical models are often abstract in nature. However, in environmental contexts, data are often limited and important insight can be gained by applying knowledge of real-world mechanisms. In this dissertation, I present three mechanistic statistical models, applied to the environment.

In my second chapter, I model the effect of sociality on the migration of adult Pacific salmon past large-scale dams in the Columbia River Basin, USA. I explicitly break down and model the process by which a salmon passes a dam. In all three steps of this process, I evaluate the effect of the density of conspecifics, to determine whether sociality plays a role in dam passage.

In my third chapter, I present a method for inferring the source of a signal which has been deformed by well-understood linear dynamics. I use as an example the case of a pollutant which, upon entering the environment, is subject to advective-diffusive transport. I show how to incorporate a mechanistic linear partial differential equation (PDE) model into the classic stochastic PDE (SPDE) method from spatial statistics, and how to invert the transport dynamics within a statistical model.

In my fourth chapter, I present a mixed integer linear program (MILP) model for constructing optimal sampling design under complex logistical or budgetary constraints. I use as an example the case of the US Forest Service (USFS) Forest Inventory and Analysis

(FIA) program in Tanana, Alaska. I compare solutions of this model with three randomized, design-based benchmarks based on MSE and feasibility.

TABLE OF CONTENTS

	Page
List of Figures	iii
List of Tables	vi
Chapter 1: Introduction	1
Chapter 2: Dam Passage Analysis for Pacific Salmon	7
2.1 Abstract	8
2.2 Introduction	9
2.3 System & Methods	11
2.4 Statistical Analysis	16
2.5 Results	17
2.6 Discussion	22
Chapter 3: Source Reconstruction for Spatio-Temporal Physical Statistical Models	33
3.1 Abstract	34
3.2 Introduction	34
3.3 Background	37
3.4 Methods	40
3.5 Results	50
3.6 Conclusion	58
Chapter 4: Optimal Sampling Design Under Logistical Constraints with Mixed Integer Programming	67
4.1 Abstract	68
4.2 Introduction	68
4.3 Background	71

4.4	Methods	76
4.5	Results	86
4.6	Conclusions	92
Chapter 5:	Discussion	99
Appendix A:	Appendices to Ch. 2	104
A.1	Statistical Analysis	105
A.2	Model Summaries	111
Appendix B:	Appendices to Ch. 3	127
B.1	General Operator	128
B.2	Dirichlet Boundary Conditions	128
Appendix C:	Appendices to Ch. 4	130
C.1	Optimal Sampling Design with Linear Regression	131
C.2	Discrete Uncertainty in the Precision Matrix	132
C.3	Using the Covariance Matrix	132
C.4	Design-Based Knapsack Sampling	133

LIST OF FIGURES

Figure Number	Page
2.1 Schematic of John Day Dam. Layout of The Dalles Dam is qualitatively similar. The John Day facility is 2,327-m long and 56-m high and has two adult fishways: one on each shoreline. The tailrace antennas were 1.8 km downstream from the dam (3.2 km at The Dalles Dam). Multiple underwater antennas were used to monitor fish passage into and through the fishways. Components of this schematic are not to scale. We modeled three processes: the “finding” time from when a fish enters a tailrace (blue) to when it first enters a fishway (red); the “fishway” time from when a fish last enters a fishway to when it exits into the upstream reservoir; and the “commit” probability that a fish passes all the way through a fishway on its first attempt.	12
2.2 The results of our null distribution simulations. Histograms represent fitted density coefficients from over 2000 simulations of the best fitting null model (with no density effect). Red lines represent the density coefficient fitted to our actual data. Chin refers to models of Chinook salmon; Sock refers to models of sockeye salmon. Find , Fishway , and Commit refer to our three different process models (see Figure 2.1). TD and JD refer to The Dalles and John Day dams. Summer and Spring refer to models in which summer-run and spring-run Chinook were separated. East refers to the east fishway at The Dalles Dam. Separation of models beyond species and process were the result of model diagnostic procedures.	20
2.3 Predictions from our two Chinook “finding” models under various fish density scenarios. Densities were chosen to be a factor (see legend) multiplied against measured counts, to simulate a realistic scenario of higher or lower densities. Factors ranged from 0, to simulate near-extirpation up to twice current levels. Curves at each density factor were generated by simulating from the fitted model 192 times and calculating the median probability of passage at each time point.	21

3.1	The results of a numerical simulation of the proposed spatio-temporal Matérn distribution. Note that this function is both the solution to an SPDE in one context ($u(s, t)$ in Equation 3.15) but that we propose to use it as a source function in another context ($f(s, t)$ in Equation 3.10). Parameters used were: $\tau = 2, \alpha = 4, \kappa = 1$, corresponding to $\rho \approx 7$. A spatial buffer zone of length 15 on each side was excluded from this figure. No temporal buffer zone was removed. Space-time was discretized using 751 spatial nodes, and 10^4 time steps.	46
3.2	The results of a numerical experiment for conducting source reconstruction on the advection-diffusion-reaction equation, with a linear regression on spatial factors. Panel (a) shows the spatially-variable flow rate (always to the right). Panel (b) shows inference on \mathbf{u} . Panel (c) shows inference on $\mathbf{f} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$. Note the discontinuities in \mathbf{f} generated by discontinuities in \mathbf{X} . Dots show positions of measurements and their observed values.	51
3.3	Empirical vs. approximate plots showing the convergence of error for our method when estimating both the source and solution of a PDE from solution-only measurements in Case Study 1. Note that our analytical approximations underestimate the error initially, but become accurate at a sample size of $N < 10$. In the first simulation (a), $D = 0.75, r = 0.2$, and $\sigma_{\epsilon}^2 = 5$, while in the second (b) $D = 0.075, r = 0.02$, and $\sigma_{\epsilon}^2 = 0.5$	52
3.4	The results of Case Study 1's MCMC analysis. Prior PDF had been scaled up by the number of samples to be comparable to the histogram.	53
3.5	The results of Case Study 1's MCMC analysis. Black dotted lines represent the true simulated source and solution functions. Red dotted lines represent the pointwise posterior mean. Background lines represent the kriging estimate from each MCMC sample. Note that the solution (or concentration) is estimated precisely but the source (or emissions) is more uncertain. Both functions lie largely within the range of variation among samples.	54
3.6	The generated distributions for the source, solution, and regression coefficients given the posterior samples of the parameters. The vertical line represents the true value of the coefficient. Note that all three coefficients are underestimated on average, perhaps due to oversmoothing in the kriging step.	54
3.7	The results of Case Study 2. The source process was spatio-temporal Matérn as defined in Section 3.4.4 with parameters: $\tau = 2, \alpha = 4, \kappa = 1$. Flow rate was constant in time, and varied in space as shown in figure 3.2. Red points denote sampling locations. A spatial buffer zone of length 5 on each side is not included in the plot. No temporal buffer zone was removed.	57

3.8	The main results of Case Study 3, parameterized with the NOMADs climate model (Rutledge et al., 2006), and using EPA PM 2.5 measurements (US Environmental Protection Agency, 2021). Advection field (a; upper left); advection speed (b; upper right); PM 2.5 map (c; middle left); PM 2.5 map posterior standard deviation (d; middle right); PM 2.5 emissions map (e; lower left); PM 2.5 source map posterior standard deviation (f; lower right). The buffer region (approximately 4 degrees wide) is set to zero to avoid showing boundary anomalies. Dots represent locations of data points.	59
4.1	Example sample for the Knapsack scenario. Sample locations shown in circles (blue). Pilot samples shown in diamonds (magenta). Domain and strata outlined in dashed line (black).	87
4.2	Example sample for the Helipad scenario. Sample locations shown in circles (blue). Pilot samples shown in diamonds (magenta). Helipad locations shown in Xs (red). Helipad ranges shown in large transparent circles (green). Domain and strata outlined in dashed line (black). Choice of helipads for (a)-(e) determined by cost-minimization MILP.	88
4.3	Example sample for the Tanana scenario. Sample locations marked with 0 or 1 depending on which sampling season the sample is to be collected. Base ranges shown in large transparent circles (green with hatching toward upper-right for season 0; blue with hatching toward upper-left for season 1). Base used to reach each sample shown in radiating lines (blue). Domain (the Tanana region of Alaska) outlined in dashed line (black), and also outlined (red) in a map of Alaska in panel (e). Axes in panels (a)-(d) and (f) are in NAD83 / Alaska Albers coordinates; panel (e) is in latitude and longitude.	89
4.4	Results of our benchmarking simulations. Left-side plots contain the same lines as the right-side, but without the addition of bootstrapped error bars which obscure the underlying lines. Red numbers display the sample size for each point, and the characters “B&B” and “G” stand for the branch and bound/greedy methods respectively.	91
A.1	A diagram of the models we fit. Originally we planned to fit the 6 models shown in blue on the left. However, due to model diagnostic violations (specifically violations detected by <code>cox.zph</code>) we were forced to split further to allow baseline hazard functions to differ between factors. These more specific final models are shown in red. Our final set of models which passed all model diagnostics are shown in green.	109

LIST OF TABLES

Table Number	Page	
2.1	Results for the 13 models considered after model splitting. Final model was not fit due to small sample size. All p values calculated using parametric bootstraps from the AICc-selected null model. All 95% CIs calculated using parametric bootstrap from the full model (selected null model plus a density effect). Sample size reflects the number of radio-tagged fish. Note that the p values were not Bonferroni corrected; see ESM for more details.	19
3.1	The true parameters used for simulation in Case Studies 1 and 2.	49
3.2	Posterior mean and standard deviations for each parameter of the advection-diffusion-reaction source model of U.S. air pollution. Also the priors, and prior quantiles for those parameters that had them. The bottom three parameters are generated quantities.	58
4.1	Logistic factors modeled in the Tanana scenario.	82
4.2	The parameter values used for the three different scenarios of our benchmarking study. Greek letters refer to parameters used to simulate the data-generating process. We translated between κ and ρ using the empirical equation from (Lindgren et al., 2011). Time Limit refers to the time limit at which the Gurobi solver terminates the B&B procedure and returns the best known solution.	83
4.3	The number of benchmarking simulations run to estimate MSE and probability of feasibility for each method, by scenario.	86
4.4	Proportional improvements in MSE demonstrated by our B&B solution as a percentage of our three benchmarks' MSEs. For SBRS in the Knapsack scenario, shows the proportional increase in MSE. StratRS was not used in the Tanana scenario.	90

ACKNOWLEDGMENTS

I'd like to express my heartfelt appreciation to my family and friends: particularly Leanna, James, Dylan, and Skyler for seeing me through these last few challenging years; my parents for support, encouragement, and reminiscent commiseration; and my labmates, particularly Anne and Maria for sharing the struggles of grad school. Thanks also to my advisor Andrew Berdahl, and the QERM department as a whole. Thanks to my committee members Sándor Tóth and Mevin Hooten for their support and encouragement, and to my coauthors Matthew Keefer and Peter Westley.

DEDICATION

to Dylan

Chapter 1
INTRODUCTION

Model selection is a core problem for ecological data analysis. In other fields, manipulative experiments are commonly used to rigorously infer cause and effect, but such experiments are often infeasible on the spatio-temporal scales relevant in ecology (Tredennick et al., 2021). Even when they are feasible, such as in whole-lake manipulations, thorny ethical questions can be raised by such direct manipulation of the natural world. As a result, scientists often must rely on statistical models to draw inference about the processes at play in the environment. Nor can ecologists rely solely on recent advances in machine learning (Baker et al., 2018), since frequently the goal of such analysis is understanding of a particular phenomenon rather than exploration or prediction. Rather, models that explicitly account for mechanism can both leverage scientific knowledge (e.g., Wikle and Hooten 2010) to improve inferences and can also shed light on the mechanisms themselves.

There are a wide range of applications in which mechanistic models offer notable advantages. In this dissertation, I develop three such applications, along the whole cycle of statistical modeling from sampling to model building to data analysis.

My first chapter deals with animal movement models, which are frequently at least pseudo-mechanistic. These models incorporate many different behaviors, from home ranges associated with foraging and territoriality (Börger et al., 2008) to task switching behavior (e.g. Hooten et al. 2019) to flocking (e.g. Ballerini et al. 2008). Mechanisms can be modeled in a variety of way, from hierarchical statistical models with discrete variables determining behavior; to bias terms in stochastic differential equations; to individual based models of emergent group behaviors. When models such as these are fit to data, they frequently rely on having individual-level position data over time. However, when data are only available on the aggregate level, or is only at a very coarse grain on the individual level, these models can become more difficult to apply. In my first chapter, I explore such a case, where coarse-grain data are available for a small fraction of radio-tagged fish and only for a handful of time points when these fish pass fixed antennas (Keefer et al., 2004; Caudill et al., 2007). Even coarser data was also available in the form of aggregate counts (CBR, 2015). In my first chapter I construct a generalized linear models and two proportional hazards models

to model these data, and use a parametric bootstrap method to account for the feedback loop wherein high movement rates reduce the density of fish, which itself is hypothesized to influence movement rate.

My second chapter deals with a source reconstruction problem, which is more common in the inverse problems literature than in statistics *per se*. Inverse problems deal in loose terms with determining the causes of a set of measured effects, as opposed to forward problems which predict the effects given a set of causes, although less emphasis is generally given to formal inference and uncertainty quantification, since inverse problems are often formulated in the case of ideal, unlimited, noise-free data (Evans and Stark, 2002). Nevertheless, inverse problems are quite amenable to statistical approaches in the case of finite, noisy data, and indeed many regularization techniques used in inverse problems methods admit statistical interpretations, often as maximum *a posteriori* estimates. Close ties with mechanistic forward problems mean that inverse problems tend to use mechanistic information heavily. In my second chapter I explore a statistical approach to the inverse problem of source reconstruction. This problem involves measuring a signal subject to some dynamics, often a partial differential equation (PDE), and attempting to determine the input that would generate such a signal. However rather than seeking a single estimate, I approach this inverse problem from a statistical perspective by modifying the stochastic partial differential equation (SPDE) method of Lindgren et al. (2011) to allow spatial statistical inference of the source term using Gaussian processes.

My third chapter deals with optimal sampling design. This field is less commonly mechanistic in nature, dealing more frequently with methods of quantifying the quality of a design, and of finding an approximately optimal solution to such a problem. However, in this case, mechanism may be introduced from the perspective of logistics and budget. Environmental surveys, particularly large-scale spatial surveys, are frequently subject to logistical and budgetary constraints such as transportation, sample processing costs, and accessibility (Hughes and Peck, 2008). Traditional methods of sampling design do not account for these logistics, which can make it difficult to feasibly collect the prescribed sample. On the flip side, many

ecologists collect convenience samples without any formal design process at all (Anderson, 2001). My methods provides a way of collecting a rigorous model-based design constrained by a mechanistic model of the logistics involved with collecting such a design. By using a mixed integer linear program (MILP) to express the design problem, the rich literature of MILP problems and industrial solvers can be brought to bear on the optimal sampling problem.

BIBLIOGRAPHY

- Anderson, D. R. (2001). The need to get the basics right in wildlife field studies. *Wildlife Society Bulletin*, pages 1294–1297.
- Baker, R. E., Pena, J.-M., Jayamohan, J., and Jérusalem, A. (2018). Mechanistic models versus machine learning, a fight worth fighting for the biological community? *Biology Letters*, 14(5):20170660.
- Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, E., Giardina, I., Lecomte, V., Orlandi, A., Parisi, G., Procaccini, A., et al. (2008). Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences*, 105(4):1232–1237.
- Börger, L., Dalziel, B. D., and Fryxell, J. M. (2008). Are there general mechanisms of animal home range behaviour? A review and prospects for future research. *Ecology Letters*, 11(6):637–650.
- Caudill, C. C., Daigle, W. R., Keefer, M. L., Boggs, C. T., Jepson, M. A., Burke, B. J., Zabel, R. W., Bjornn, T. C., and Peery, C. A. (2007). Slow dam passage in adult Columbia River salmonids associated with unsuccessful migration: delayed negative effects of passage obstacles or condition-dependent mortality? *Canadian Journal of Fisheries and Aquatic Sciences*, 64(7):979–995.
- CBR (2015). Columbia River DART (data access in real time).
- Evans, S. N. and Stark, P. B. (2002). Inverse problems as statistics. *Inverse Problems*, 18(4):R55.

- Hooten, M. B., Scharf, H. R., and Morales, J. M. (2019). Running on empty: Recharge dynamics from animal movement data. *Ecology Letters*, 22(2):377–389.
- Hughes, R. M. and Peck, D. V. (2008). Acquiring data for large aquatic resource surveys: the art of compromise among science, logistics, and reality. *Journal of the North American Benthological Society*, 27(4):837–859.
- Keefer, M. L., Peery, C. A., Bjornn, T. C., Jepson, M. A., and Stuehrenberg, L. C. (2004). Hydrosystem, dam, and reservoir passage rates of adult Chinook salmon and steelhead in the Columbia and Snake rivers. *Transactions of the American Fisheries Society*, 133(6):1413–1439.
- Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498.
- Tredennick, A. T., Hooker, G., Ellner, S. P., and Adler, P. B. (2021). A practical guide to selecting models for exploration, inference, and prediction in ecology. *Ecology*, 102(6):e03336.
- Wikle, C. K. and Hooten, M. B. (2010). A general science-based framework for dynamical spatio-temporal models. *Test*, 19(3):417–451.

Chapter 2

DAM PASSAGE ANALYSIS FOR PACIFIC SALMON

Publication history: A version of this chapter is published as: Okasaki, Connie, et al. "Collective navigation can facilitate passage through human-made barriers by homeward migrating Pacific salmon." Proceedings of the Royal Society B 287.1937 (2020): 20202137 (<https://doi.org/10.1098/rspb.2020.2137>), where it is subject to a Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>). The following is modified to fit the formatting and language of a dissertation chapter. This version is subject to a CC-BY 4.0 license

2.1 Abstract

The mass migration of animals is one of the great wonders of the natural world. Although there are multiple benefits for individuals migrating in groups, an increasingly recognized benefit is collective navigation, whereby social interactions improve animals' ability to find their way. Despite substantial evidence from theory and lab-based experiments, empirical evidence of collective navigation in nature remains sparse. Here we used a unique large-scale radiotelemetry dataset to analyze the movements of adult Pacific salmon (*Oncorhynchus sp.*) in the Columbia River Basin, USA. These salmon face substantial migratory challenges approaching, entering, and transiting fishways at multiple large-scale hydroelectric mainstem dams. We assess the potential role of collective navigation in overcoming these challenges and show that Chinook salmon (*O. tshawytscha*), but not sockeye salmon (*O. nerka*) locate fishways faster and pass in fewer attempts at higher densities, consistent with collective navigation. The magnitude of the density effects were comparable with major established drivers such as water temperature, and model simulations predicted that major fluctuations in population density can have substantial impacts on key quantities including mean passage time and fraction of fish with very long passage times. The magnitude of these effects indicate the importance of incorporating conspecific density and social dynamics into models of the migration process. Density effects on both ability to locate fishways and number of passage attempts have the potential to enrich our understanding of migratory energetics and success of migrating anadromous salmonids. More broadly, our work reveals a potential role of

collective navigation, in at least one species, to mitigate the effects of anthropogenic barriers to animals on the move.

2.2 Introduction

Long distance migration is an iconic and threatened behavior (Wilcove and Wikelski, 2008). Migratory species navigate with incredible precision to and from highly spatially restricted locations (Bingman and Cheng, 2005). To solve these challenging navigational problems, species use mechanisms ranging from an innate sun compass (Reppert et al., 2010) and magnetic maps (Putman et al., 2014) to learned olfactory cues (Dittman and Quinn, 1996; Keefer and Caudill, 2014). However, long-distance migrations are threatened by human influence (Berger, 2004). Fences, highways, and other developments block terrestrial migration pathways (Seidler et al., 2015), light pollution interferes with aerial migrations (Gauthreaux Jr et al., 2006) and dams and de-watering impede passage both upstream and downstream (e.g. Norrgård et al. (2013)). In light of these impacts, it is essential to understand the mechanisms of navigation in order to predict and mitigate human impacts on migratory populations.

Many species migrate in groups, and this is thought to aid navigation (Berdahl et al., 2018). Theory suggests that such *collective navigation* may be the result of a number of mechanisms. For example, groups can increase accuracy by averaging over error-prone individual directional estimates — known as the “many wrongs” principle (Codling et al., 2007). Even without individual directional estimates, accurate directional responses can emerge through social interactions — known as “emergent sensing” (Berdahl et al., 2013). We direct interested readers to Box 1 of Berdahl et al. (2018) for an overview of these and other mechanisms including leadership (Couzin et al., 2005), and social and collective learning (Fagan et al., 2012; Kao et al., 2014). A growing body of empirical literature lends compelling support to these hypothesized mechanisms, but evidence from wild populations is rare (Berdahl et al., 2018).

One iconic example of a migratory species is salmon, which home in large numbers back

to their natal spawning grounds. Salmon navigation is not fully understood but is known to include an inherited magnetic map (Putman et al., 2014) and olfactory recognition of natal water (Dittman and Quinn, 1996). Berdahl et al. (2016b) hypothesized that salmon use collective navigation, based on seven independent studies reporting positive associations between homing accuracy and run size. However, like many migratory species, salmon face an anthropogenic barrier: dams. For example, the impassable Swan Falls Dam on the Snake River in Idaho rendered approximately 25% of mainstem riverine habitat inaccessible to Snake River Chinook salmon (*Oncorhynchus tshawytscha*). A study of fall-run Chinook spawning habitat on the Columbia River found that between impassable dams and altered flow regimes, less than 20% of historical spawning habitat for fall-run Chinook remained available (Dauble et al., 2003). Many dams have fishways, which allow salmon to pass by these barriers. However, locating a fishway entrance is non-trivial, since salmon use rheotaxis to move upstream, and the main source of flow at most dams is the spillway or turbines. Adult salmon migrations are often delayed in dam tailraces, and this delay may bear important costs, including increased exposure to predation (e.g. Keefer et al. (2012)) and increased energy output which can lead to greater mortality (e.g. Burnett et al. (2014)). Thus, if collective navigation eases the dam passage process, it may have an important effect on salmon survival and reproductive success, and therefore important implications for population conservation.

Here, we use radiotelemetry data on adult Chinook and sockeye salmon (*Oncorhynchus nerka*) navigating upstream past dams on the Columbia River, combined with daily fish counts at the dams, to evaluate the hypothesis that collective navigation helps salmon overcome the navigational challenges posed by fishways. We find strong evidence that Chinook salmon find and commit to fishways more rapidly on higher-density days. Evidence for density effects in sockeye salmon, and for Chinook salmon navigating within fishways, was weak. Although some prior evidence exists for negative density effects in similar contexts (e.g. Gorig and Castro-Santos (2017)), we did not find any prior evidence of negative density effects in either sockeye or Chinook salmon.

2.3 System & Methods

2.3.1 Study System

The Columbia River drains $>600,000$ km² of seven western U.S. states and two Canadian provinces and historically supported some of the most abundant Pacific salmon and steelhead runs in the world (Chapman, 1986). The basin has been transformed by hydroelectric development, with 14 large dams on the main stem Columbia River and 20 dams on the main stem Snake River, the Columbia's largest tributary by area. The dams, along with overharvest, habitat loss, and artificial propagation, contributed to steep declines in Columbia River salmon populations (Lichatowich, 2001) and subsequent threatened or endangered status under the U.S. Endangered Species Act (NFSC, 2015).

Upstream-migrating adult salmonids can pass many of the Columbia basin dams via pool-and-weir fishways (Katopodis and Williams, 2012) that rise ≈ 17 – 56 m per dam. To navigate past the dams, adults first pass through turbulent, high-velocity tailraces that are several kilometers long and >1 km wide. Fish must then locate low-volume fishway openings sited near powerhouses or adjacent to spillways, move through a series of collection channels and junction pools, ascend a fish ladder, and then exit into the upstream reservoir (Figure 2.1). The spatial scale and hydraulic complexity of dam tailraces and fishways present several navigational and physiological challenges. The combination of searching for passage routes and fishway exit and re-entry behaviors, for example, is energetically demanding, particularly when fish make multiple passage attempts (Brown et al., 2006; Crozier et al., 2017). Typical upstream migration rates for Chinook and sockeye salmon in undammed sections of the Columbia basin range from ≈ 18.5 – 52.7 km/day (Fryer et al., 2012). In contrast, adult salmonids take ≈ 1 – 3 days to pass each main stem dam along their Columbia River migration route (Keefer et al., 2004). These tailrace and dam reaches range in length from 0.5–3.2 km, such that, on average, the fish are travelling 0.17–3.2 km/day during dam passage (a ≈ 10 – 100 -fold reduction in up-stream passage speed), potentially delaying timely arrival at spawning sites.




Figure 2.1: Schematic of John Day Dam. Layout of The Dalles Dam is qualitatively similar. The John Day facility is 2,327-m long and 56-m high and has two adult fishways: one on each shoreline. The tailrace antennas were 1.8 km downstream from the dam (3.2 km at The Dalles Dam). Multiple underwater antennas were used to monitor fish passage into and through the fishways. Components of this schematic are not to scale. We modeled three processes: the “finding” time from when a fish enters a tailrace (blue) to when it first enters a fishway (red); the “fishway” time from when a fish last enters a fishway to when it exits into the upstream reservoir; and the “commit” probability that a fish passes all the way through a fishway on its first attempt.

2.3.2 Data Collection

Data used in this study were from salmon collected and radio-tagged at Bonneville Dam in 2013 and 2014 using previously-described methods (Keefer et al., 2004; Caudill et al., 2007). Bonneville Dam is at Columbia River kilometer (rkm) 235.1 and is the first dam returning adult salmon encounter during their upstream migration. Telemetered fish were monitored at multiple dams and in tributaries, but analyses focused on data collected as fish entered tailraces and passed through fishways at The Dalles Dam (TD; rkm 308.1) and John Day Dam (JD; rkm 346.9), the second and third dams from the Pacific. Monitoring arrays at the two projects included aerial 9-element Yagi antennas sited on tailrace shorelines 1.8–3.2 rkm downstream from the dams and underwater coaxial cable antennas at fishway openings, inside fishway collection channels and junction pools, and in fish ladders (Keefer et al., 2004, 2013b). Detection ranges for the aerial antennas ranged from 100s of m to >1 km, depending on fish depth (Melnychuk, 2012). Ranges for the underwater antennas were 5–15 m. The raw telemetry data from all antennas were assembled, filtered, and coded using established methods (Keefer et al., 2004; Caudill et al., 2007); files of coded records were then used to identify when salmon entered and exited tailraces and fishways.

We obtained daily counts of salmon from the Columbia Basin Research Data Access in Real Time (DART) database (CBR, 2015) and environmental data including water temperature and spillway discharge from the US Geological Survey (USGS) National Water Inventory System (NWIS) database (<https://waterdata.usgs.gov/nwis>). This synthesis of several different sources of data was crucial for our analysis. Daily counts are collected by human observers and are able to capture a large fraction of passing fish. These data are therefore suited to measuring density, but not to modeling individual passage rates. On the other hand, telemetry arrays are more expensive and technical, capturing only a small fraction of fish, but allow much more detailed observations of those fish which are tagged. These data allowed us to model passage rates for individual fish. Although the synthesis of these datasets was of great benefit, it also introduced several challenges, discussed in the

following section. Most notably, counts were *not* a direct measurement of fish density.

2.3.3 Modeling Approach

We split the process of passing a dam into a sequence of three distinct stages: i) the “finding” process which starts at a fish’s first detection in a tailrace and continues until its first detection at a fishway opening; ii) the “committing” process which models the probability that a fish actually passes all the way through a fishway to the upstream exit on its first attempt, rather than exiting back into the tailrace; and iii) the “fishway” process which starts when a fish enters a fishway and commits (i.e. the last fishway entrance) and continues until the fish exits from the top of the ladder.

We modeled the “finding” and “fishway” processes using a time-to-event analysis (also known as “survival” analysis in the medical literature) framework (for a general introduction, see Kleinbaum and Klein 2010). We modeled the “committing” process using a logistic regression. A detailed technical description of our models and diagnostics can be found in the electronic supplementary materials (ESM).

A time-to-event dataset consists of a series of observations through time of an individual along with a set of covariates for that individual. We were interested in two events — that of entering a fishway, and that of exiting a fishway into the upstream reservoir — and we wished to understand how long it took for these events to occur. For each event we must also define a starting time: respectively, the entrance to the tailrace and the last entrance to the fishway. In other words, our first model answers the question: “how quickly does a fish find and enter a fishway after entering a tailrace?” Our second model answers the question: “how quickly does a fish navigate and exit a fishway after entering it and committing to passage?” More complex “multi-state” models tie events together into a sequence (Andersen et al., 2002), but to simplify our analysis we used more tractable single-process models.

We used a common time-to-event model known as the proportional hazards (PH) model. A hazards model assumes that the time-to-event process is essentially an exponential decay, with an event rate $\lambda(t|X)$ (the “hazard function”) which changes depending on both how

long it has been since the fish entered the system, t , and with a set of covariates X . The PH model assumes that this rate depends on covariates according to the equation

$$\lambda(t|X) = \lambda_0(t) \exp(X\beta).$$

The function $\lambda_0(t)$ is known as the baseline hazard function. In our case, covariates such as temperature, spillway discharge, and fish density were time-varying and the rate at which fish passed through the dam was assumed to also vary in time proportionately.

Our focus was on the effect of density, which we quantified using daily counts of individuals obtained from the DART database (CBR, 2015). This use of counts introduced two major problems. First, our count covariate was a product of density and hazard: when fish were passing through the fishway more quickly, a higher proportion of them would be counted over the course of a day. This introduced reverse causality: a high rate (the response) caused our counts (a predictor) to be higher. Moreover, since hazard varied substantially over the course of a day, the relationship between counts and density was inconsistent. Second, the fish observation windows are placed near the middle of fish ladders. Thus there was an unknown amount of time between when a fish triggered a telemetry reading and when it was counted by the observers. A third concern, not related to our use of counts, was that density is related to hazard through time: when fish are passing through a fishway more quickly, fewer of them remain behind to contribute to density. In other words, high passage rates deplete density over time. This introduced another source of reverse causality.

These issues posed a substantial challenge. The first and last were particularly worrisome, since reverse causality has the potential to bias model fits and produce spurious results. Even if there was no causal effect of density, we might observe statistically significant density effects in our models. For example, suppose that due to other unmeasured covariates or random chance, some days have a higher than average rate of fish passage into a fishway. On those days we would observe a higher count, and empirically the rate of passage would positively correlate with count, despite having no causal effect of density on passage.

We took two major steps to account for these modeling challenges. First, although hourly

count data were available, we used daily count data to estimate density. Since a consistently high proportion of radio-tagged fish in our dataset passed within 24 hours ($>97\%$), and relatively few fish passed overnight ($<10\%$ entering the fishway and $<5\%$ exiting to the reservoir), the daily count provided a roughly accurate assessment of daily density while smoothing out any within-day variation in passage rate. On the other hand, hourly counts are far more dependent on random chance and are systematically biased due to daily variations in passage rate. Moreover, hourly counts increase the risk of other biases due to the time lag between a tagged fish triggering a telemetry event and being counted by an observer.

Second, we used a parametric bootstrap to protect against reverse causality. Our bootstrap analysis was justified using the following chain of logic. First, our system violated modeling assumptions and our models may therefore produce biased inference. Second, if density truly has no effect then counts have no effect. In this case we need not include counts as a covariate, and our models should provide valid unbiased fits. Therefore, simulations from models with no count covariates produce valid null distributions. By fitting models *with* count covariates to these null model simulations, we obtained a null distribution for the count coefficient. Even if the above problems induced reverse causality in our dataset, our null distribution remains valid, since the null model has no reverse causality. Due to reverse causality, the power of our tests remains unknown, but we may nevertheless correctly calculate p -values. For a more detailed introduction to bootstraps in general see Efron and Tibshirani (1994). Generally we use *null model* to mean a model without a count covariate.

2.4 Statistical Analysis

To satisfy model assumptions, we split our initial 4 PH models (2 species \times 2 processes) into 11 models. We did not split our 2 logistic regression models. By splitting a model, we mean we split the dataset for that model into disjoint subsets and fit separate but structurally identical models to each set of subsets. We split all PH (finding; fishway) models by dam (TD; JD). We further split our Chinook fishway models by run (spring; summer) and our

Sockeye fishway TD model by specific fishway (East; North). These submodels are referenced in Table 2.1 and are explained in further detail in the ESM. One PH model (Sockeye fishway TD North) was discarded due to low sample size, leaving us with 12 final models. In each of those 12 models we conducted one set of null model simulations (bootstraps) to test for the presence of a density effect as well as a set of secondary bootstraps to calculate confidence intervals for the count coefficients.

After conducting model diagnostics, we added a threshold density effect ($\text{count} \leq 150$) to our Chinook finding model. This successfully accounted for a series of large positive residuals, possibly modeling a strong crowding effect (i.e. negative density effect). However, this effect was highly confounded with fishway identity, with 88 out of 89 fish under the threshold passing at TD's east fishway, and because of this no strong conclusions can be made about this effect. We used this effect as part of the null model, prior to testing for a linear density effect.

To ensure that our fitted density effects were biologically important as well as statistically significant, we produced simulations for a variety of density scenarios ranging from zero density (near-extirpation) to twice current density, which can be found in Figure 2.3. Each density scenario is obtained by multiplying observed counts by a fixed factor, and then producing model simulations. We did this for both the PH models and the logistic regression models.

For additional technical details regarding our statistical analysis, please refer to the ESM.

2.5 Results

Of our twelve tests, three detected significant ($p < 0.001$) density effects. Two were Chinook PH models for finding the fishway, and the third was a Chinook logistic regression model for committing to a fishway. Two other effects were significant at the $p < 0.05$ or 0.1 levels (Summer-run Chinook navigating the TD fishway; and Sockeye navigating the TD east fishway). However, since these two effects were statistically weak and did not display any consistent pattern with other models of the same species and process, we concluded that

these associations were spurious. All significant density effects were positive, meaning that an increase in density facilitated faster predicted completion of the modeled process, and therefore faster overall passage.

Our simulated null distributions for the count coefficients are shown in Figure 2.2, compared to the actual count coefficient fitted to our true (non-simulated) dataset. These results along with our confidence intervals, are shown in Table 2.1. Coefficients are reported after standardizing all covariates so that magnitudes can be compared. Full model summaries are available in the ESM.

These results provide strong evidence for the existence of positive density effects, particularly in Chinook salmon. No evidence was found for negative density effects. Evidence for sockeye salmon did not support a density effect.

Our density scenario simulations for the “finding” models can be found in Figure 2.3. These simulations predict a 20-30% increase in finding times for Chinook under the extirpation scenario compared to present conditions, and 50-115% more fish with high finding times of >24 hours. Conversely, if salmon densities were to double compared to 2013/2014 levels we predicted a 15-19% decrease in finding times and 31-46% fewer fish with high (>24 hr) finding times. Our density scenario simulations for the “committing” model predicted a linear increase in commit rates with densities, with doubling densities bringing an increase of almost 5% commit rate from the current level of 43% to 47.7%. However, lower density predictions from this model were unreliable, due to the threshold density effect in our null model. Because this effect was only evident in one model, and was confounded with fishway identity we were not able to fully investigate this possible density effect, which has a profound effect on our predictions at low densities.

Consistent with previous studies in the same system (Keefer et al., 2004, 2013a), we also observed strong temperature ($0.22 < |\beta| < 1.29$; negative in 2/4 Chinook fishway models, otherwise positive) and very strong diel ($1.22 < \beta < 2.96$; higher passage during the day) effects in our PH models. We observed a strong temperature effect ($\beta = -1.23$) in our Chinook committing model. Our significant density coefficients were of roughly the

Species	Process	Dam	Subset	p -value	95% CI	Sample Size	
Chinook	Find	JD	.	$< 5e-04$	(0.36 , 0.75)	804	
		TD	.	$< 5e-04$	(0.16 , 0.49)	751	
	Commmmit	.	.	$< 5e-04$	(0.29 , 1.2)	930	
	Fishway	JD	Spring Run	.	> 0.1	(-0.088 , 0.49)	439
			Summer Run	.	> 0.1	(-0.59 , 1.1)	345
		TD	Spring Run	.	> 0.1	(-0.17 , 0.2)	415
			Summer Run	.	0.061	(-0.054 , 1.2)	332
Sockeye	Find	JD	.	> 0.1	(-0.032 , 0.04)	609	
		TD	.	> 0.1	(-0.023 , 0.042)	616	
	Commmmit	.	.	> 0.1	(-0.053 , 0.075)	678	
	Fishway	JD	.	> 0.1	(-0.047 , 0.028)	605	
		TD	East Fishway	.	0.037	(0.00067 , 0.039)	553
			North Fishway	.	.	.	59

Table 2.1: Results for the 13 models considered after model splitting. Final model was not fit due to small sample size. All p values calculated using parametric bootstraps from the AICc-selected null model. All 95% CIs calculated using parametric bootstrap from the full model (selected null model plus a density effect). Sample size reflects the number of radio-tagged fish. Note that the p values were not Bonferroni corrected; see ESM for more details.

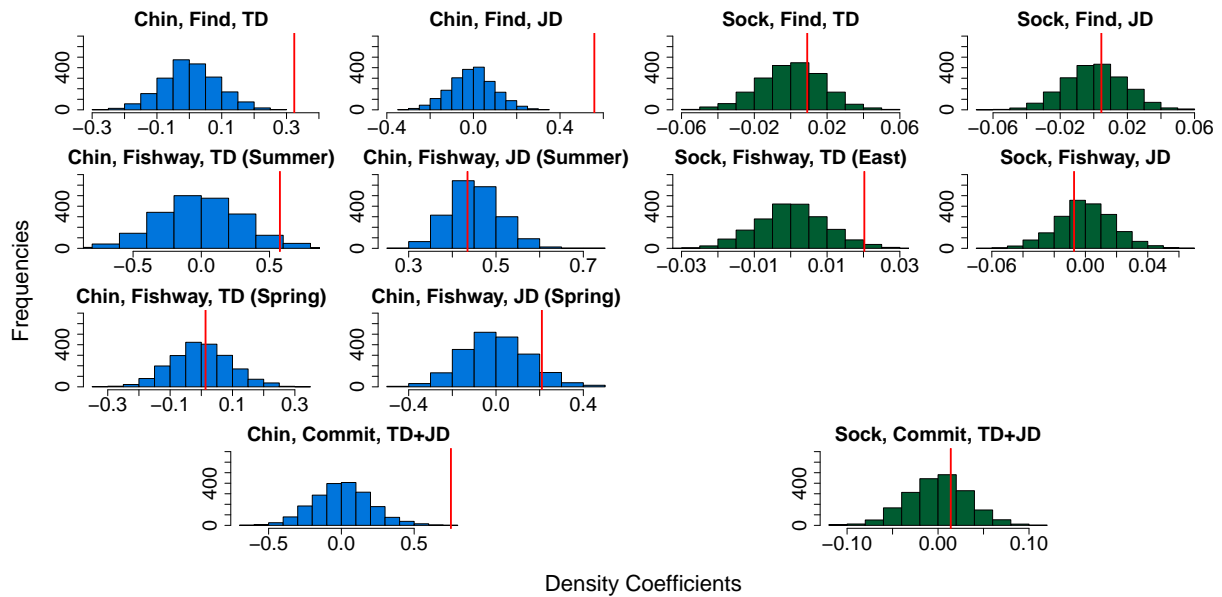


Figure 2.2: The results of our null distribution simulations. Histograms represent fitted density coefficients from over 2000 simulations of the best fitting null model (with no density effect). Red lines represent the density coefficient fitted to our actual data. **Chin** refers to models of Chinook salmon; **Sock** refers to models of sockeye salmon. **Find**, **Fishway**, and **Commit** refer to our three different process models (see Figure 2.1). **TD** and **JD** refer to The Dalles and John Day dams. **Summer** and **Spring** refer to models in which summer-run and spring-run Chinook were separated. **East** refers to the east fishway at The Dalles Dam. Separation of models beyond species and process were the result of model diagnostic procedures.

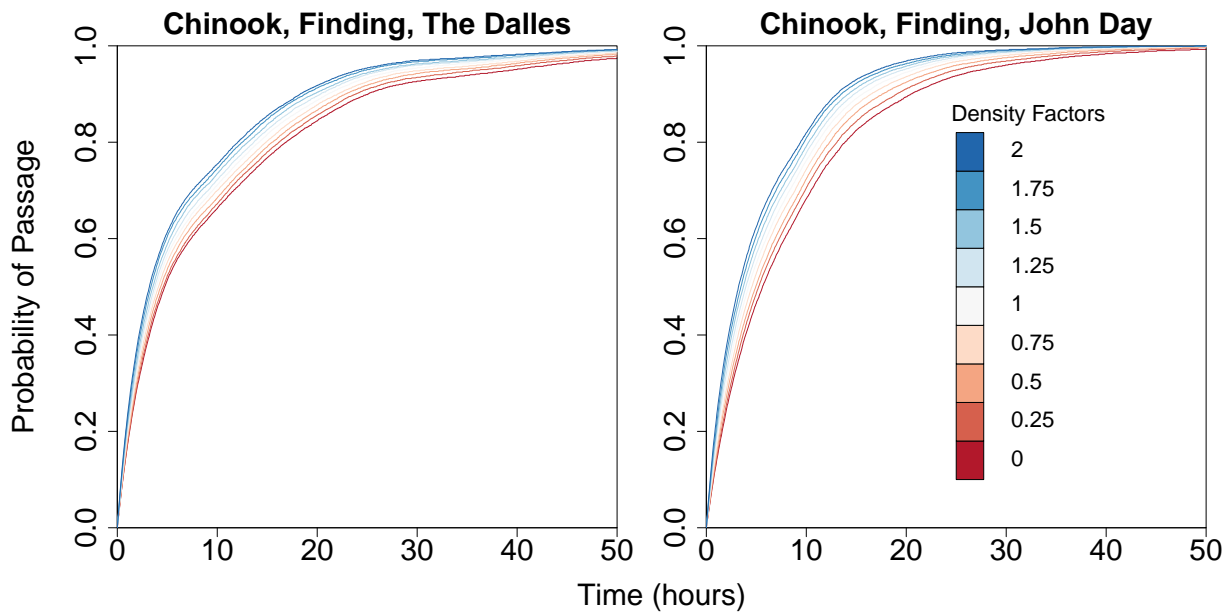


Figure 2.3: Predictions from our two Chinook “finding” models under various fish density scenarios. Densities were chosen to be a factor (see legend) multiplied against measured counts, to simulate a realistic scenario of higher or lower densities. Factors ranged from 0, to simulate near-extirpation up to twice current levels. Curves at each density factor were generated by simulating from the fitted model 192 times and calculating the median probability of passage at each time point.

same order of magnitude as our temperature coefficients, indicating that density may have a substantial impact, comparable to other established drivers. Since these other variables were not our focus, we did not conduct any further analysis of these environmental effects. Model summaries are available in ESM.

2.6 Discussion

Due to human modifications of the global landscape, animals face increased migratory challenges (Hardesty-Moore et al., 2018). Salmon returning to spawning grounds in the Columbia River basin are confronted with multiple main stem dams. While many dams have fishways, passage remains a daunting challenge as fish need to locate a fishway entrance and then ascend a ladder to continue their migration. Here we reveal evidence that Chinook, but not sockeye, salmon appear to benefit from social interactions during this challenge. We demonstrated, using a bootstrap method to account for several sources of reverse causality, a positive effect of density on two key quantities for Chinook salmon: rates of locating fishway entrances, and probability of committing to passage of a fishway. In contrast, we found no density effect for passage rates through a fishway, or for any aspect of dam passage for sockeye salmon.

Our results can alternatively be explained by an effect of migratory motivation or run-timing, as other studies have noted (Caudill et al., 2007). Fish on peak density days may be more motivated, or may have timed their migration better with respect to unmeasured environmental conditions than earlier- or later-migrating fish. Controlling for these effects without directly manipulating density is difficult. Given the consistency of variation between species, we feel that the most plausible explanation is the presence of a true density effect in Chinook salmon, rather than an effect of motivation or timing, which we might expect to see more strongly in both species. Generally, we see our lack of statistically “significant” results for sockeye salmon as adding support to our results for Chinook salmon. We do not expect density effects to exist in all processes or for all species, but do expect these effects to apply consistently.

It is nevertheless counter-intuitive that the species with a stronger tendency to school might benefit less from collective navigation. Here we provide two plausible, yet speculative, explanations. First, having a strong tendency to migrate in groups (Berdahl et al., 2017), sockeye salmon might form sizeable groups regardless of population density — i.e. their local density may not be highly dependent on global density. In contrast, the less-social Chinook salmon’s local density might be more governed by global density. Thus for Chinook the number of social interactions may scale with daily fish density whereas for sockeye social interactions may be relatively constant. Second, casual observations by the authors suggest another potential mechanism. In tailraces, Chinook salmon tend to hold near fishway entrances before entering. This holding behaviour seems to result in loose shoal-like aggregations near fishway entrances, potentially leaving a sort of social signpost that may draw subsequent Chinook towards an entrance. On the other hand, sockeye salmon’s increased tendency to school might inhibit their ability to take advantage of such a social effect (as in Lemasson et al. (2014)). By travelling in large, cohesive, polarized schools, sockeye salmon may have trouble transitioning from a large tailrace to a relatively small fishway. The social “momentum” of the school may prevent individuals who do spot a fishway entrance from stopping to explore it further. This possible mechanism also explains why we observed no strong effects within the fishways, although it does not explain the effect we observed for commit probabilities.

The effect of density on fishway passage rates was first studied over 60 years ago, but has largely focused on the effects of overcrowding (i.e., deleterious effects of density). The possibility of a positive effect of density on fishway passage rates was discussed by Lander (1959) in the context of a model for overcrowding, and saw experimental work in a study of alewives (*Alosa pseudoharengus*), also in the context of overcrowding (Dominy, 1973). Unfortunately, although Dominy (1973) reported a positive density effect, they analyzed their data in a way that was inconsistent with our analysis¹. More recent work in the

¹Dominy (1973) used a numerical response (number passing per unit time) rather than the rate response (proportion passing per unit time) used here. Although they reported a sigmoidal (positive at low densities,

context of culvert passage has found negative (Goerig and Castro-Santos, 2017), and no (Johnson et al., 2012) density effects for brook trout (*Salvelinus fontinalis*) and Coho salmon (*O. kisutch*) respectively. These study systems were most similar to our “fishway” system, involving little difficulty in finding the entrance to the culvert (e.g. in both studies fish were confined at the downstream extremity, as opposed to our open system). Thus our null results for “fishway” are consistent with Johnson et al. (2012), while our positive result for “finding” involves a fundamentally different, relatively unstudied system. A study by Caudill et al. (2007) on Chinook salmon and steelhead (sea-run *O. mykiss*) reported several positive associations with density, but density effects were not the focus of their analysis, and their model encompassed the entire dam passage process from tailrace entry to fishway exit. A study of juvenile palmetto bass (*Morone saxatilis x chrysops*), Lemasson et al. (2014) found that fish in schools took much longer than lone individuals to pass an artificial barrier when moving downstream and thus showed a negative effect of density on passage rate.

Given the diversity of findings in this literature, and in our own results, there is scope for additional work in this area. Our methods can be applied to other systems and species, where “finding” and “committing” behaviors have rarely been isolated for analysis. For example, controlled navigation experiments investigating these behaviors, where density can be manipulated systematically, are a promising avenue. Emerging technologies including sonic tags, acoustic cameras, and computer vision will make more detailed analysis of individual and collective movement around dams possible and also yield local (i.e. actual group size, rather than estimated fish density in tailraces) measures of conspecific density (Hughey et al., 2018). Such studies could also shed light on the specific mechanisms driving any collective navigation. Revealing individual and collective search algorithms may contribute general principles to the fields of animal movement and bio-inspired engineering, but may also contribute insights for dam management of fish passage. For example, mechanisms such as the

then negative at high densities) density effect in alewives, a visual analysis of their data appears more consistent (in our framework) with either no density effect or a negative effect, although there is a large group of outliers that could indicate a positive effect.

“social signpost” hypothesis for Chinook salmon might motivate management interventions designed to promote more efficient salmon migration past fishways. One simple hypothetical intervention might simulate a holding pattern of fish about fishway entrances using model fish. These decoys could provide a social signpost, attempting to activate our hypothesized social behaviors even at low densities.

When conducting our analysis we added an additional threshold density effect to our Chinook “commit” model, possibly modeling a strong crowding effect. However, we did not conduct an in-depth analysis of this variable, since it was confounded with use of the TD east fishway. One possible explanation for this effect is as a density effect associated with overcrowding in the fishway — however it is unclear why overcrowding would be represented by a threshold effect rather than a more gradual decline in commit probabilities. Furthermore, even if this density effect were associated with overcrowding in one particular fishway, we would expect such an effect to vary among fishways, making it difficult to generalize this effect to make predictions elsewhere. We recommend this as an avenue for further research.

Our analysis faced three substantial challenges: counts were confounded with hazard, counts were a time-delayed estimate of density, and density was confounded with hazard across time. We addressed these challenges using our bootstrap approach to circumvent concerns about reverse causality, and by using smoother daily counts rather than hourly counts which are more subject to concerns about time-delay and inconsistent relationships with density. These were the more unique roadblocks we encountered, but like any statistical analysis, there are some other caveats to consider. For example, other covariates such as flow velocity may have played an important role but were not included, and some variables such as migratory motivation are all but impossible to measure quantitatively in any case. Like any model, ours was an imperfect representation of reality. We feel it was close enough to reality to provide useful insight. Another caveat of our analysis is that we dropped fish with missing measurements from our dataset, rather than conducting a so-called “censoring analysis” of missing data. Since we used only “entry” and “exit” data points, a proper censoring analysis was impossible. Furthermore, since data points were likely missing independently at random

due to missed radiotelemetry signals, we incurred little bias in this manner.

To our knowledge, our results provide the first demonstration of a positive density effect for fish passing riverine obstacles. This is in contrast to the prevailing wisdom that overcrowding is the dominant effect. More broadly, it is one of only a handful of examples of collective navigation in freely-migrating populations. Given the ubiquity of social movement during migration, we expect many more examples of collective navigation to be uncovered, especially as new technologies improve our ability to quantify collective movement *in situ* (Hughey et al., 2018). Such studies underscore the need for further investigation of density effects, since these processes could have important ecological implications. For example, if populations decline, densities at dams will decline: will this alleviate overcrowding or reduce collective benefits? In the first case, population decline is buffered, while in the other it is magnified, potentially generating critical transitions that may lead to sudden population collapses (Fagan et al., 2012; Berdahl et al., 2016a). As anthropogenic disturbances simultaneously increase navigational difficulty in a variety of contexts and decrease population densities, understanding the role of density-dependent processes, such as collective navigation, may therefore yield critical insights for sound management and conservation (Westley et al., 2018).

Acknowledgments

Telemetry data were collected with funding from the US Army Corps of Engineers, Portland District and with assistance from the NW Fisheries Science Center (NOAA-Fisheries) and the University of Idaho Fish Ecology Research Lab. This material is based upon work supported by the NSF-GRFP under Grant No. DGE-1762114. We thank David Smith and Christa Woodley for assistance obtaining data and to Christy Contreras for preliminary analysis as part of an NSF REU project at the Santa Fe Institute. Finally, we thank Chris Caudill for invaluable support.

BIBLIOGRAPHY

- Andersen, P. K., Abildstrom, S. Z., and Rosthøj, S. (2002). Competing risks as a multi-state model. *Statistical Methods in Medical Research*, 11(2):203–215.
- Berdahl, A., Torney, C. J., Ioannou, C. C., Faria, J. J., and Couzin, I. D. (2013). Emergent sensing of complex environments by mobile animal groups. *Science*, 339(6119):574–576.
- Berdahl, A., van Leeuwen, A., Levin, S. A., and Torney, C. J. (2016a). Collective behavior as a driver of critical transitions in migratory populations. *Movement Ecology*, 4(1):18.
- Berdahl, A., Westley, P. A., Levin, S. A., Couzin, I. D., and Quinn, T. P. (2016b). A collective navigation hypothesis for homeward migration in anadromous salmonids. *Fish and Fisheries*, 17(2):525–542.
- Berdahl, A., Westley, P. A., and Quinn, T. P. (2017). Social interactions shape the timing of spawning migrations in an anadromous fish. *Animal Behaviour*, 126:221–229.
- Berdahl, A. M., Kao, A. B., Flack, A., Westley, P. A., Codling, E. A., Couzin, I. D., Dell, A. I., and Biro, D. (2018). Collective animal navigation and migratory culture: from theoretical models to empirical evidence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1746):20170009.
- Berger, J. (2004). The last mile: how to sustain long-distance migration in mammals. *Conservation Biology*, 18(2):320–331.
- Bingman, V. P. and Cheng, K. (2005). Mechanisms of animal global navigation: comparative perspectives and enduring challenges. *Ethology Ecology & Evolution*, 17(4):295–318.
- Broström, G. and Broström, M. G. (2019). Package ‘eha’.

- Brown, R., Geist, D., and Mesa, M. (2006). Use of electromyogram telemetry to assess swimming activity of adult spring Chinook salmon migrating past a Columbia River dam. *Transactions of the American Fisheries Society*, 135(2):281–287.
- Burnett, N. J., Hinch, S. G., Braun, D. C., Casselman, M. T., Middleton, C. T., Wilson, S. M., and Cooke, S. J. (2014). Burst swimming in areas of high flow: delayed consequences of anaerobiosis in wild adult sockeye salmon. *Physiological and Biochemical Zoology*, 87(5):587–598.
- Burr, D. (1994). A comparison of certain bootstrap confidence intervals in the Cox model. *Journal of the American Statistical Association*, 89(428):1290–1302.
- Caudill, C. C., Daigle, W. R., Keefer, M. L., Boggs, C. T., Jepson, M. A., Burke, B. J., Zabel, R. W., Bjornn, T. C., and Peery, C. A. (2007). Slow dam passage in adult Columbia River salmonids associated with unsuccessful migration: delayed negative effects of passage obstacles or condition-dependent mortality? *Canadian Journal of Fisheries and Aquatic Sciences*, 64(7):979–995.
- Caudill, C. C., Keefer, M. L., Clabough, T. S., Naughton, G. P., Burke, B. J., and Peery, C. A. (2013). Indirect effects of impoundment on migrating fish: temperature gradients in fish ladders slow dam passage by adult Chinook salmon and steelhead. *PLoS One*, 8(12):e85586.
- CBR (2015). Columbia River DART (data access in real time).
- Chapman, D. (1986). Salmon and steelhead abundance in the Columbia River in the nineteenth century. *Transactions of the American Fisheries Society*, 115(5):662–670.
- Codling, E., Pitchford, J., and Simpson, S. (2007). Group navigation and the “many-wrongs principle” in models of animal movement. *Ecology*, 88(7):1864–1870.
- Couzin, I. D., Krause, J., Franks, N. R., and Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513–516.

- Crozier, L. G., Bowerman, T., Burke, B. J., Keefer, M. L., and Caudill, C. C. (2017). High stakes steeplechase: a behavior-based model to predict individual travel time through diverse migration segments. *Ecosphere*, 8(10):e01965.
- Dauble, D. D., Hanrahan, T. P., Geist, D. R., and Parsley, M. J. (2003). Impacts of the Columbia River hydroelectric system on main-stem habitats of fall Chinook salmon. *North American Journal of Fisheries Management*, 23(3):641–659.
- Dittman, A. and Quinn, T. (1996). Homing in Pacific salmon: mechanisms and ecological basis. *Journal of Experimental Biology*, 199(1):83–91.
- Dominy, C. (1973). Effect of entrance-pool weir elevation and fish density on passage of alewives (*Alosa pseudoharengus*) in a pool and weir fishway. *Transactions of the American Fisheries Society*, 102(2):398–404.
- Efron, B. and Tibshirani, R. J. (1994). *An Introduction to the Bootstrap*. CRC press.
- Fagan, W. F., Cantrell, R. S., Cosner, C., Mueller, T., and Noble, A. E. (2012). Leadership, social learning, and the maintenance (or collapse) of migratory populations. *Theoretical Ecology*, 5(2):253–264.
- Fryer, J. K., Whiteaker, J., and Kelsey, D. (2012). Upstream migration timing of Columbia Basin Chinook and sockeye salmon and steelhead in 2010. *Columbia River Inter-Tribal Fish Commission Technical Report*, pages 12–02.
- Gauthreaux Jr, S. A., Belser, C. G., Rich, C., and Longcore, T. (2006). Effects of artificial night lighting on migrating birds. *Ecological consequences of artificial night lighting*, pages 67–93.
- Goerig, E. and Castro-Santos, T. (2017). Is motivation important to brook trout passage through culverts? *Canadian Journal of Fisheries and Aquatic Sciences*, 74(6):885–893.

- Hardesty-Moore, M., Deinet, S., Freeman, R., Titcomb, G. C., Dillon, E. M., Stears, K., Klope, M., Bui, A., Orr, D., Young, H. S., et al. (2018). Migration in the Anthropocene: how collective navigation, environmental system and taxonomy shape the vulnerability of migratory species. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1746):20170017.
- Hughey, L. F., Hein, A. M., Strandburg-Peshkin, A., and Jensen, F. H. (2018). Challenges and solutions for studying collective animal behaviour in the wild. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1746):20170005.
- Johnson, G. E., Pearson, W. H., Southard, S. L., and Mueller, R. P. (2012). Upstream movement of juvenile coho salmon in relation to environmental conditions in a culvert test bed. *Transactions of the American Fisheries Society*, 141(6):1520–1531.
- Kao, A. B., Miller, N., Torney, C., Hartnett, A., and Couzin, I. D. (2014). Collective learning and optimal consensus decisions in social animal groups. *PLoS Computational Biology*, 10(8).
- Katopodis, C. and Williams, J. (2012). The development of fish passage research in a historical context. *Ecological Engineering*, 48(November):8–18.
- Keefer, M. L. and Caudill, C. C. (2014). Homing and straying by anadromous salmonids: a review of mechanisms and rates. *Reviews in Fish Biology and Fisheries*, 24(1):333–368.
- Keefer, M. L., Caudill, C. C., Peery, C. A., and Moser, M. L. (2013a). Context-dependent diel behavior of upstream-migrating anadromous fishes. *Environmental Biology of Fishes*, 96(6):691–700.
- Keefer, M. L., Jepson, M. A., Clabough, T. C., and Caudill, C. C. (2013b). Fishway passage bottleneck identification and prioritization: a case study of Pacific lamprey at Bonneville dam. *Canadian Journal of Fisheries and Aquatic Sciences*, 63(10):1752–1762.

- Keefer, M. L., Peery, C. A., Bjornn, T. C., Jepson, M. A., and Stuehrenberg, L. C. (2004). Hydrosystem, dam, and reservoir passage rates of adult Chinook salmon and steelhead in the Columbia and Snake rivers. *Transactions of the American Fisheries Society*, 133(6):1413–1439.
- Keefer, M. L., Stansell, R. J., Tackley, S. C., Nagy, W. T., Gibbons, K. M., Peery, C. A., and Caudill, C. C. (2012). Use of radiotelemetry and direct observations to evaluate sea lion predation on adult Pacific salmonids at Bonneville Dam. *Transactions of the American Fisheries Society*, 141(5):1236–1251.
- Kleinbaum, D. G. and Klein, M. (2010). *Survival Analysis*, volume 3. Springer.
- Lander, R. H. (1959). *The problem of fishway capacity*, volume 301. US Department of Interior, Fish and Wildlife Service.
- L’ecuyer, P., Simard, R., Chen, E. J., and Kelton, W. D. (2002). An object-oriented random-number package with many long streams and substreams. *Operations Research*, 50(6):1073–1075.
- Lemasson, B. H., Haefner, J. W., and Bowen, M. D. (2014). Schooling increases risk exposure for fish navigating past artificial barriers. *PLoS One*, 9(9).
- Lichatowich, J. (2001). *Salmon without rivers: a history of the Pacific salmon crisis*. Island Press, Washington DC.
- Melnichuk, M. (2012). Detection efficiency in telemetry evaluations: definitions and evaluation methods. In NS Adams, J. B. and Eiler, J., editors, *Telemetry Techniques: A User Guide for Fisheries Research*, pages 339–357. American Fisheries Society, Bethesda, MD.
- NFSC (2015). *Status review update for Pacific salmon and steelhead listed under the Endangered Species Act: Pacific Northwest*. Northwest Fisheries Science Center, NOAA-Fisheries.

- Norrgård, J. R., Greenberg, L. A., Piccolo, J. J., Schmitz, M., and Bergman, E. (2013). Multiplicative loss of landlocked Atlantic salmon *Salmo salar* L. smolts during downstream migration through multiple dams. *River Research and Applications*, 29(10):1306–1317.
- Putman, N. F., Scanlan, M. M., Billman, E. J., O’Neil, J. P., Couture, R. B., Quinn, T. P., Lohmann, K. J., and Noakes, D. L. (2014). An inherited magnetic map guides ocean navigation in juvenile Pacific salmon. *Current Biology*, 24(4):446–450.
- R Core Team (2019). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Reppert, S. M., Gegear, R. J., and Merlin, C. (2010). Navigational mechanisms of migrating monarch butterflies. *Trends in Neurosciences*, 33(9):399–406.
- Seidler, R. G., Long, R. A., Berger, J., Bergen, S., and Beckmann, J. P. (2015). Identifying impediments to long-distance mammal migrations. *Conservation Biology*, 29(1):99–109.
- Therneau, T. M. and Lumley, T. (2014). Package ‘survival’.
- Westley, P. A., Berdahl, A. M., Torney, C. J., and Biro, D. (2018). Collective movement in ecology: from emerging technologies to conservation and management. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1746):20170004.
- Wilcove, D. S. and Wikelski, M. (2008). Going, going, gone: is animal migration disappearing. *PLoS biology*, 6(7).

Chapter 3

**SOURCE RECONSTRUCTION FOR SPATIO-TEMPORAL
PHYSICAL STATISTICAL MODELS**

Publication history: A version of this chapter is published as: Okasaki, Connie, Mevin B. Hooten, and Andrew M. Berdahl. "Source reconstruction for spatio-temporal physical statistical models." Spatial Statistics 52 (2022): 100707.

3.1 Abstract

In many applications, a signal is deformed by well-understood dynamics before it can be measured. For example, when a pollutant enters a river, it immediately begins dispersing, flowing, settling, and reacting. If the pollutant enters at a single point, its concentration can be measured before it enters the complex dynamics of the river system. However, in the case of a non-point source pollutant, it is not clear how to efficiently measure its source. One possibility is to record concentration measurements in the river, but this signal is masked by the fluid dynamics of the river. Specifically, concentration is governed by the advection-diffusion-reaction PDE, with an unknown source term. We propose a method to statistically reconstruct a source term from these PDE-deformed measurements. Our method is general and applies to any linear PDE. This method has important applications in the study of environmental DNA and non-point source pollution.

3.2 Introduction

In many applications, a substance is emitted into a dynamic environment such as a stream or the atmosphere, where it is then passively transported. Measuring concentrations of such a substance is a natural monitoring tool, but inferring emissions from concentrations requires unraveling the dynamics of the transport medium. Often, concentrations can be modeled as the solution to a partial differential equation (PDE), in which case the emissions can be modeled as a source term.

We present a method using stochastic PDEs (SPDEs) to calculate the conditional distribution of the source term of a linear PDE, given measurements of the PDE solution. These calculations are efficient when the source can be represented as a Gaussian Markov Random Field (GMRF). Our technique builds on the seminal SPDE work of Lindgren et al. (2011),

which suggested advection-diffusion modeling as a potential application of their approach. Few studies have extended this application within the SPDE context. The closest in spirit are Sigrist et al. (2012), which used an equivalent integro-difference equation; and Sigrist et al. (2015), which used Fourier methods to build a spatio-temporal advection-diffusion model. However, these studies largely demonstrate the case of a single drift vector rather than an entire advection field. Outside the SPDE context, spatial and spatio-temporal statistical models of concentration within dynamic mediums are more common, but often either ignore transport and emissions effects (e.g. Moraga et al. 2017), or incorporate these effects in a non-mechanistic (e.g. Cameletti et al. 2019) or pseudo-mechanistic manner (e.g. Ver Hoef et al. 2006). Here we use explicit advection-diffusion mechanisms, within the SPDE context in order to obtain sparse precision matrices, and can incorporate a spatio-temporally varying advection field.

Our method provides a mechanistic kernel for modeling concentrations, as well as a way to infer emissions from those concentrations. The latter is novel in the field of spatial statistics, and statistically more difficult, because transport distorts the original emissions signal. We therefore focus on the latter application. In Section 3.3, we provide background information: a review of the SPDE approach for generating the Matérn covariance, a review of the finite element method as it is used in this approach, and an overview of the advection-diffusion-reaction equation. In Section 3.4, we outline the method we used for source reconstruction, as well as the method we used for quantifying the accuracy of our source reconstructions, a method for constructing spatio-temporal versions of our model, the MCMC method we used for sampling, and finally any specifics regarding our three case studies. In Section 3.5, we describe the results of these three case studies.

3.2.1 Motivation

Our focus on inferring emissions is motivated by several continuing and emerging challenges in environmental science. The first and most direct application is the problem of non-point source air and water pollution. There is a long history of investigation into the inverse

source problem in these domains (e.g. Panofsky 1969), but study has been largely focused on the more mathematically tractable point-source case (e.g., Stockie 2011). Nevertheless, non-point sources are a major concern for regulators, in part because the traditional policies used to regulate pollution are difficult to implement for these sources (Xepapadeas, 2011).

Another major motivation of this work is environmental DNA¹ (eDNA). The eDNA literature has made great strides in the last decade and techniques such as DNA metabarcoding (Deiner et al., 2017) have seen wide acceptance. More recently, scientists have investigated the use of quantitative polymerase chain reaction analysis (qPCR²) to estimate the abundance of organisms (e.g., Doi et al. 2015, Doi et al. 2017). Once the dynamic properties of eDNA (shedding rate, decay rate, etc) can be quantified in field settings, eDNA offers a promising tool for non-invasive abundance estimation (Harrison et al., 2019). Our method offers a mathematical framework in which many different dynamic properties can be incorporated to infer spatial and spatio-temporal abundance.

The inverse source problem is broad, and includes many problems beyond these two motivating applications. Other literature in this field includes inferring sources of heat (El Badia and Ha-Duong, 2002); acoustic waves (Alves et al., 2009); and electromagnetic signals, particularly EEG data (Grech et al., 2008). We do not consider these applications in this paper, but the adaptation of our method to these domains is a promising area of future study.

¹Environmental DNA is particles of DNA shed from an organism into the environment. These particles can be used to passively detect the presence of particular species (Rees et al., 2014), or in certain circumstances can be used to estimate biomass (Rourke et al., 2022). In theory, provided accurate parameters describing the physical-chemical dynamics of environmental DNA, one could use it to derive a spatio-temporal map of biomass using only non-invasive sampling.

²The polymerase chain reaction is a biological reaction used to amplify a particular DNA sequence if it is present in a sample. It is traditionally used as a detector of presence/absence. However, in recent years qPCR has been developed as a technique to estimate not just presence/absence but the *amount* of a particular DNA sequence present in the original sample.

3.3 Background

3.3.1 Statistical Background

We assume that the source, or emissions, $f(\mathbf{s}) \sim GP(\mu(\mathbf{s}), k(\mathbf{s}, \mathbf{s}'))$ is a Gaussian process with mean function μ and covariance function k . We assume that $f(\mathbf{s})$ cannot be measured directly, and is the primary target of inference. Furthermore, we assume that f can be represented using a Gaussian Markov Random Field (GMRF; for general reference see Rue and Held 2005) so that for any particular finite set of points \mathbf{s} , $\mathbf{f}(\mathbf{s}) \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{s}), \mathbf{Q}(\mathbf{s}, \mathbf{s})^{-1})$ with a sparse precision matrix \mathbf{Q} . In the special case that k is the Matérn covariance function,

$$k(\mathbf{s}, \mathbf{s}') = \frac{\sigma^2}{2^{\nu-1}\Gamma(\nu)} (\kappa \|\mathbf{s}' - \mathbf{s}\|)^\nu K_\nu(\kappa \|\mathbf{s}' - \mathbf{s}\|), \quad (3.1)$$

Lindgren et al. (2011) provided a method for directly generating a sparse \mathbf{Q} for certain integer and half-integer values of the smoothness parameter ν . Their approach for doing so was based on applying the finite element method (FEM) to a particular set of nested diffusion equations whose solutions are governed by the Matérn covariance function. We use the same method to translate from emissions to concentration, using a different PDE (see Equation 3.2).

3.3.2 The Advection-Diffusion-Reaction Equation

We now consider the advection-diffusion-reaction equation in more detail. Suppose that we seek to infer the source of a chemical, but only have measurements of chemical concentration in a fluid medium. The chemical flows with, and disperses through, that medium, and possibly decays linearly over time. The chemical concentration is then described by the equation:

$$\frac{\partial u(\mathbf{s}, t)}{\partial t} + \nabla \cdot (\mathbf{v}(\mathbf{s}, t)u(\mathbf{s}, t)) - \nabla \cdot (D\nabla u(\mathbf{s}, t)) + ru(\mathbf{s}, t) = f(\mathbf{s}, t), \quad (3.2)$$

where ∇ is the del operator for the spatial dimensions, D is the diffusion coefficient, r is the decay rate, $\mathbf{v}(\mathbf{s}, t)$ is the advection field, assumed here to be known, $f(\mathbf{s}, t)$ is the

source function, a general Gaussian process in space and time, and $u(\mathbf{s}, t)$ is the solution, or concentration function. If we assume that $f(\mathbf{s}, t) = f(\mathbf{s})$ is homogeneous in time and that $u(\mathbf{s}, t) = u(\mathbf{s})$ has reached steady state, the PDE simplifies to

$$\nabla \cdot (\mathbf{v}(\mathbf{s})u(\mathbf{s})) - \nabla \cdot (D\nabla u(\mathbf{s})) + ru(\mathbf{s}) = f(\mathbf{s}). \quad (3.3)$$

For example, this model may be appropriate when trying to estimate the origin of non-point source pollution (e.g., runoff pollution); the source of isotopes found in water samples³; or the biomass of fish from eDNA measurements. This inverse problem has been considered in applied math studies, and is usually solved either in the case where the chemical is assumed to originate with unknown strength from a small finite number of unknown point locations (El Badia et al., 2005; El Badia and Hamdi, 2007; Lushi and Stockie, 2010), or using regularization techniques from the inverse problems literature that do not admit a probabilistic interpretation and do not include an estimate of uncertainty (Porter and Devaney, 1982; Engl et al., 1996; Yan et al., 2008).

Previous statistical work has examined the advection-diffusion equation in similar contexts using integro-difference equations (Sigrist et al., 2012) and Fourier methods (Sigrist et al., 2015). Stroud et al. (2010) used the finite difference method to account for these practicalities but used an ensemble Kalman filter method to conduct inference, rather than approaching the problem from an SPDE perspective.

3.3.3 The Finite Element Method

In general, discretization methods translate partial differential equations into systems of linear equations

$$\mathbf{K}\mathbf{u} = \mathbf{L}\mathbf{f}. \quad (3.4)$$

³Many elements found in nature have multiple naturally occurring and scientifically useful isotopes, such as tritium (³H), radiocarbon (¹⁴C), and oxygen-18 (¹⁸O). Such isotopes can be naturally enriched by environmental processes such as evaporation and can therefore be used to identify the source of water flow (groundwater vs. snow pack vs. rainfall; e.g. Hao et al. 2019) or match a sample with a particular location (e.g. Rachel et al. 2008).

In the finite element method, \mathbf{K} is known as the stiffness matrix and \mathbf{L} is known as the mass matrix. In the finite difference method, often $\mathbf{L} = \mathbf{I}$. Essential boundary conditions (i.e., those that are not automatically satisfied; usually Dirichlet) modify this basic equation. We demonstrate this modification in the Appendix. We refer to \mathbf{u} as the solution vector and \mathbf{f} as the source vector, with precision matrices \mathbf{Q}_u and \mathbf{Q}_f .

Our methodology for generating a sparse precision matrix \mathbf{Q}_u given \mathbf{Q}_f is essentially the same as in Lindgren et al. (2011). We generalize their notation and define \mathcal{L} to be any linear partial differential operator and \mathcal{B} to be a general Dirichlet, Neumann, or Robin boundary condition, to produce a general formulation of a PDE

$$\begin{aligned}\mathcal{L}u(\mathbf{s}) &= f(\mathbf{s}) & \text{for } \mathbf{s} \in \Omega \\ \mathcal{B}u(\mathbf{s}) &= c(\mathbf{s}) & \text{for } \mathbf{s} \in \partial\Omega.\end{aligned}\tag{3.5}$$

Neumann boundary conditions are the natural boundary conditions in our analyses and we use them throughout, but Dirichlet boundary conditions could also be used, as could non-zero Neumann boundary conditions and Robin boundary conditions. Boundary conditions may also be specified in terms of lower-dimensional random fields (see Appendix for details).

The finite element method discretizes Equation 3.5 by finding a weak solution using a specified set of test and trial functions. In Lindgren et al. (2011), and in this study, both classes are spanned by a series of piecewise linear “bump” functions defined to be 0 over most of the domain, and 1 at one specific node chosen from a set of finite points in space. Other classes may be advantageous for particular SPDEs, including the advection-diffusion equation (which in future work may benefit from the Streamline-Upwind Petrov Galerkin bases; Brooks and Hughes 1982) and Maxwell’s equations (which benefit from the use of specialized bases such as the Nédélec elements; Monk et al. 2003). For a more comprehensive review of the FEM in general see Brenner et al. (2008), or for an accessible review of the FEM applied to the Matérn equation see Bakka (2018).

Conceptually, the finite element method serves to translate the PDE in Equation 3.5 into a matrix equation such as Equation 3.4. Classically, one then solves for $\mathbf{u} = \mathbf{K}^{-1}\mathbf{L}\mathbf{f}$.

Lindgren et al. (2011) observed that in doing so, we imply that if \mathbf{f} is multivariate normal (MVN) with mean $\boldsymbol{\mu}_f$ and precision \mathbf{Q}_f then \mathbf{u} is MVN with mean $\mathbf{K}^{-1}\mathbf{L}\boldsymbol{\mu}_f$ and precision $\mathbf{K}^T\mathbf{L}^{-T}\mathbf{Q}_f\mathbf{L}^{-1}\mathbf{K}$. To ensure sparsity of the resulting matrix, we replace \mathbf{L} with a diagonal matrix $\tilde{\mathbf{L}}$, formed by summing the rows of \mathbf{L} and placing the resulting sums along the diagonal.

3.4 Methods

3.4.1 Finite Element Calculations

Although the finite difference method would produce conceptually identical results, we use the finite element method to compute discretized spatial derivatives. Because the stiffness and mass matrices are standard components of finite element analysis in engineering, they can be calculated by leveraging existing FEM software such as the Finite Element Computational Software (FEniCS; Alnæs et al. 2015). We rely on FEM software to implement our stiffness and mass matrices, to improve reproducibility and facilitate generalizability. We assume that the FEM approach is appropriate for any SPDE in which the stochastic source term is well-behaved and in which the FEM equation converges to the true solution in the deterministic case.

3.4.2 Source Reconstruction

Our basic observation is that physical meaning can sometimes be ascribed to $f(\mathbf{s})$ and that the finite element representation \mathbf{f} can be calculated by inverting Equation 3.4 in the reverse direction $\mathbf{f} = \mathbf{L}^{-1}\mathbf{K}\mathbf{u}$. Our method for source reconstruction given a set of parameters (which generally must be estimated) therefore follows essentially the same logic as classical kriging. We assume that the solution \mathbf{u} is observed through some set of linear functionals which can be discretized to a matrix \mathbf{A} . We therefore observe $\mathbf{y} = \mathbf{A}\mathbf{u} + \boldsymbol{\epsilon}$ with i.i.d. Gaussian noise.

The conditional field $\mathbf{u}|\mathbf{y}$ is therefore

$$\mathbf{u}|\mathbf{y} \sim N \left(\boldsymbol{\mu}_u + \frac{1}{\sigma^2} (\mathbf{Q}_u + \frac{1}{\sigma_\epsilon^2} \mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T (\mathbf{y} - \mathbf{A} \boldsymbol{\mu}_u), (\mathbf{Q}_u + \frac{1}{\sigma_\epsilon^2} \mathbf{A}^T \mathbf{A})^{-1} \right). \quad (3.6)$$

This is a standard kriging computation for a GMRF (Eq. 2.15 of Rue and Held 2005). However, we propose to take the additional step of kriging $\mathbf{f}|\mathbf{y}$ by observing that $\mathbf{f}|\mathbf{y} = \mathbf{L}^{-1} \mathbf{K} \mathbf{u}|\mathbf{y}$ and therefore

$$\mathbf{f}|\mathbf{y} \sim N \left(\mathbf{L}^{-1} \mathbf{K} \boldsymbol{\mu}_{u|\mathbf{y}}, \mathbf{L}^{-1} \mathbf{K} (\mathbf{Q}_u + \frac{1}{\sigma_\epsilon^2} \mathbf{A}^T \mathbf{A})^{-1} \mathbf{K}^T \mathbf{L}^{-1} \right). \quad (3.7)$$

To facilitate numerical calculation we use the Cholesky factors of \mathbf{Q}_u and $\mathbf{Q}_u + \frac{1}{\sigma_\epsilon^2} \mathbf{A}^T \mathbf{A}$. We calculate these Cholesky factors using CHOLMOD functions (Chen et al., 2008) provided by the `sksparse` library in Python. Specifically, we calculate the Cholesky factor for $\mathbf{R}^T \mathbf{R} = \mathbf{Q}_u$, and use the sparse Cholesky update methods provided by CHOLMOD to update that factor to account for the addition of a low-rank matrix $\frac{1}{\sigma_\epsilon^2} \mathbf{A}^T \mathbf{A}$. Note that this update step is computationally difficult when implementing our model in space-time.

Linear regression on \mathbf{f} can be accommodated within this kriging framework; this calculation is shown in the Appendix.

3.4.3 Accuracy Quantification

To determine the accuracy of our method for mechanistic kriging, we quantified error using the L^2 norm over the domain (excluding any buffer regions) relative to the kriging estimator:

$$L^2(x|y) = \sqrt{\int_{\Omega} (x|y - \mu_{x|y})^2 d\Omega}. \quad (3.8)$$

In simulations, we calculated the integral over the finite element basis functions using FEniCS native integration functions. However, assuming the mesh has cells of approximately equal extent, and assuming that all hyperparameters are known perfectly so that all error is produced by Gaussian process variability and observation error, we may approximate the error as

$$L^2(x|y) \approx \sqrt{\frac{V}{M} \text{tr} (\mathbf{I}_{\text{int}} \boldsymbol{\Sigma}_{x|\mathbf{y}} \mathbf{I}_{\text{int}})}, \quad (3.9)$$

where \mathbf{x} represents either \mathbf{f} or \mathbf{u} , V is the volume of the interior region Ω , $\Sigma_{\mathbf{x}|\mathbf{y}}$ is the conditional covariance of vector \mathbf{x} given the vector \mathbf{y} of observations of \mathbf{u} , \mathbf{I}_{int} is a diagonal matrix whose entries are the fractions of each basis function contained within Ω , and M is $\text{tr}(\mathbf{I}_{\text{int}})$.

Additionally, to quantify the rate of convergence of error, we consider the possibility of observing the same set of locations $\mathbf{A}^T \mathbf{A}$ a total of n times, resulting in a precision matrix

$$\mathbf{Q}_{\mathbf{u}} + \frac{n}{\sigma_{\epsilon}^2} \mathbf{A}^T \mathbf{A}.$$

We relax the interpretation that we actually observe locations an integral number of times so as to allow a derivative with respect to n . We reparameterize (with m the number of rows in \mathbf{A}), and set $\zeta = \log(mn)$; we also reparameterize to obtain the log of the error. Differentiating the log-error with respect to ζ , we can obtain the local polynomial convergence rate. Further, if we do not use a buffer zone, thereby allowing us to observe all nodes, we can also show that as $n \rightarrow \infty$, the polynomial convergence rate approaches $-1/2$ for both the source and solution functions. This corresponds to a $1/\sqrt{N}$ type asymptotic convergence for the L^2 error. We show that this asymptotic behavior is not observed empirically, either due to the inclusion of a buffer region, or due to slow convergence to the asymptote. Instead, the local polynomial convergence rate calculation is more useful in practice. Full calculations are shown in the Appendix.

3.4.4 Spatio-Temporal Modeling

Our implementation of time-dependence relies on a finite element method in space and a finite difference method in time. This results in a sparse vector auto-regressive (VAR) model for the spatial variables (Wikle and Hooten, 2010). We discretize the time interval $[t_0, t_1]$ into small units of $\Delta t = (t_1 - t_0)/N$. We can then derive a matrix equation from any arbitrary SPDE

$$\frac{\partial u(\mathbf{s}, t)}{\partial t} + \mathcal{L}(t)u(\mathbf{s}, t) = f(\mathbf{s}, t) \quad (3.10)$$

(where $\mathcal{L}(t)$ is an arbitrary partial differential operator; for example, in Equation 3.2, we have $\mathcal{L}(t)u = \nabla \cdot (\mathbf{v}u) - \nabla \cdot (D\nabla u) + ru$) by first applying a backward Euler discretization in time

$$\begin{aligned} \frac{u(\mathbf{s}, t + \Delta t) - u(\mathbf{s}, t)}{\Delta t} + \mathcal{L}u(\mathbf{s}, t + \Delta t) &\approx f(\mathbf{s}, t + \Delta t), \\ u(\mathbf{s}, t + \Delta t) + \Delta t \mathcal{L}u(\mathbf{s}, t + \Delta t) &\approx u(\mathbf{s}, t) + \Delta t f(\mathbf{s}, t + \Delta t), \end{aligned} \quad (3.11)$$

and then applying an FEM discretization in space

$$\begin{aligned} \mathbf{L}\mathbf{u}_{t+\Delta t} + \Delta t \mathbf{K}_{t+\Delta t} \mathbf{u}_{t+\Delta t} &\approx \mathbf{L}\mathbf{u}_t + \Delta t \mathbf{L}\mathbf{f}_{t+\Delta t}, \\ (\mathbf{I} + \Delta t \mathbf{L}^{-1} \mathbf{K}_{t+\Delta t}) \mathbf{u}_{t+\Delta t} &\approx \mathbf{u}_t + \Delta t \mathbf{f}_{t+\Delta t}, \\ u(\mathbf{s}, t) &\approx (\mathbf{I} + \Delta t \mathbf{L}^{-1} \mathbf{K}_{t+\Delta t})^{-1} (\mathbf{u}_t + \Delta t \mathbf{f}_{t+\Delta t}), \end{aligned} \quad (3.12)$$

where all functions of space and time are replaced by vectors in space, subscripted by time, and the operator $\mathcal{L}(t+\Delta t)$ is discretized into $\mathbf{K}_{t+\Delta t}$. This equation can be solved inductively by

$$\begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_N \end{bmatrix} = \begin{bmatrix} \mathbf{M}_1 \\ \mathbf{M}_2 \mathbf{M}_1 \\ \vdots \\ \mathbf{M}_N \dots \mathbf{M}_1 \end{bmatrix} \mathbf{u}_0 + \Delta t \begin{bmatrix} \mathbf{M}_1 & & & \\ \mathbf{M}_2 \mathbf{M}_1 & \mathbf{M}_2 & & \\ \vdots & \vdots & \ddots & \\ \mathbf{M}_N \dots \mathbf{M}_1 & \mathbf{M}_N \dots \mathbf{M}_2 & \cdots & \mathbf{M}_N \end{bmatrix} \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \\ \vdots \\ \mathbf{f}_N \end{bmatrix}, \quad (3.13)$$

where $\mathbf{M}_i = (\mathbf{I} + \Delta t \mathbf{L}^{-1} \mathbf{K}_{t_0+i\Delta t})^{-1}$. The above Toeplitz matrix has sparse inverse

$$\mathbf{R} = \frac{1}{\Delta t} \begin{bmatrix} \mathbf{M}_1 & & & \\ -\mathbf{I} & \mathbf{M}_2 & & \\ & \ddots & \ddots & \\ & & -\mathbf{I} & \mathbf{M}_N \end{bmatrix}. \quad (3.14)$$

A forward Euler discretization yields a very similar matrix, but requires that the time step Δt be small enough to ensure numerical stability.

Suppose now that the full, discretized, spatio-temporal process \mathbf{f} has sparse precision matrix \mathbf{Q}_f and that we have a deterministic initial condition $\mathbf{u}_0 = \mathbf{0}$ (accounted for by

including a temporal buffer region before the region of interest). Then the precision matrix for \mathbf{u} is $\mathbf{Q}_u = \mathbf{R}^T \mathbf{Q}_f \mathbf{R}$. Our initial condition requires us to extend the temporal discretization prior to the first time-of-interest, to avoid boundary effects. Alternatively, the steady-state equation could be used to derive an approximate spatial prior for the initial condition, although this should still be coupled with a (perhaps smaller) buffer region.

This model uses sparse precision matrices for efficient computation. The block bi-diagonal Cholesky factors allow rapid computation in a high dimensional space, in part because the blocks are themselves sparse matrices. This makes computing likelihoods efficient, scaling linearly in the number of time-steps modeled. The computational difficulty of this model arises from the Cholesky update step (i.e., the factorization of $\mathbf{Q}_u + \frac{1}{\sigma_\epsilon^2} \mathbf{A}^T \mathbf{A}$), necessary both for inferring the distribution for the emissions or concentration maps, conditional on the observations (the kriging step) and for calculating the marginal likelihood function (the marginalization step). These conditional updates break the structure of our Cholesky factors, resulting in much greater computational burden. In essence, the problem is that information from each observation is transported over time, resulting in long-range dependencies. In densely sampled systems, this problem might be solved by calculating the temporal range of the covariance function, and truncating dependencies after a fixed number of time-steps, thus limiting the bandwidth of the resulting updated Cholesky factor.

Nested Diffusion Definition of Spatio-temporal Matérn Source

The need for a non-mechanistic spatio-temporal source distribution compatible with our method motivates the following definition of a three-parameter spatio-temporal Matérn field

$$\left(\tau \frac{\partial}{\partial t} + \kappa^2 - \Delta \right)^{\alpha/2} u(\mathbf{s}, t) = \mathcal{W}(\mathbf{s}, t). \quad (3.15)$$

This is a nested spatio-temporal diffusion equation, where the more general source term $f(\mathbf{s}, t)$ is specified to be constant-in-time spatial white noise (denoted $\mathcal{W}(x)$ for consistency with Lindgren et al. 2011) then the spatial Matérn SPDE is the steady state solution. Although the Matérn model does not necessarily require interpretation, the interpretation of

this nested diffusion model is more straightforward than an alternative model for which the spatial Matérn SPDE is also the steady state solution

$$\tau \frac{\partial u(\mathbf{s}, t)}{\partial t} + (\kappa^2 - \Delta)^{\alpha/2} u(\mathbf{s}, t) = \mathcal{W}(\mathbf{s}, t). \quad (3.16)$$

The two equations are equivalent when $\alpha = 2$; both can be seen as a diffusion-decay equation with natural time scale τ . In this circumstance, “randomness” in the form of white noise is spread out across space by means of a familiar mechanism. At higher values of α the equations and their interpretations diverge. Equation 3.15 can be seen to be a nested form of the diffusion-decay equation; converting from an Eulerian to a Lagrangian perspective, one can imagine random particles entering a system according to $\mathcal{W}(\mathbf{s}, t)$ and diffusing, only to continue emitting secondary particles that themselves diffuse, and so on. In contrast Equation 3.16, even for integer values of $\alpha/2$, involves at minimum fourth-order-in-space operators coupled to a first-order-in-time derivative. Fourth-order equations, while not at all uncommon in specialized mathematical fields, are relatively exotic. This definition aligns with that proposed in (Bakka et al., 2020), with $\alpha_s = 2$ and \mathcal{E}_Q white in both space and time.

If we wish to avoid fractional diffusion we must assume $\alpha = 2n$ for $n \in \mathbb{Z}_+$, for which we obtain

$$\mathbf{Q}_u^{(2n)} = \frac{\mathbf{R}^{Tn}}{\Delta t} (\mathbf{L} \otimes \mathbf{I}) \frac{\mathbf{R}}{\Delta t}. \quad (3.17)$$

This results in $\mathbf{Q}_f = \mathbf{L} \otimes \mathbf{I}$, which is straightforward to calculate. For example, Figure 3.1 shows the results of a numerical simulation of this distribution. Although a spatial buffer-zone was included in these calculations and removed from the plot, the temporal buffer-zone was not removed, to demonstrate the transient behavior imposed by our deterministic initial condition. A small zone with length on the order τ displays lower spatio-temporal variance than the remainder of the plot, but no differences appear to persist beyond that. For this set of parameters, we used $\Delta t = 0.05$ and 2000 time steps. With 751 spatial nodes, simulation required ≈ 21 seconds on a desktop computer with a 3.2 GHz 6-core processor and 32 GB of RAM.




Figure 3.1: The results of a numerical simulation of the proposed spatio-temporal Matérn distribution. Note that this function is both the solution to an SPDE in one context ($u(s, t)$ in Equation 3.15) but that we propose to use it as a source function in another context ($f(s, t)$ in Equation 3.10). Parameters used were: $\tau = 2, \alpha = 4, \kappa = 1$, corresponding to $\rho \approx 7$. A spatial buffer zone of length 15 on each side was excluded from this figure. No temporal buffer zone was removed. Space-time was discretized using 751 spatial nodes, and 10^4 time steps.

3.4.5 Parameter Estimation

To estimate the parameters of our air pollution model (see Section 3) we used MCMC. To facilitate a Gibbs update for the variance, we specified inverse-gamma priors for σ_f^2 the variance of the emissions function, and we specified gamma priors for ρ, v, D , and r (respectively, the range; ratio of observation variance σ_ϵ^2 to emissions variance σ_f^2 ; diffusion coefficient; and decay coefficient). All other parameters are updated in Metropolis-Hastings steps.

We marginalized over the concentration and emission vectors \mathbf{u} and \mathbf{f} , but used a second-stage sampler to draw from the full-conditional distributions of \mathbf{u} and \mathbf{f} to compute their posterior means and pointwise standard deviations.

To treat the range as a parameter we used the empirical equation $\rho = \sqrt{(8\nu)}/\kappa$ which relates range to the Matérn parameters (Lindgren et al., 2011) and transformed the log-likelihoods accordingly. To use Gibbs sampling for σ_f^2 we reparameterized to the variance ratio: $v = \frac{\sigma_\epsilon^2}{\sigma_f^2}$. We transformed the log-likelihoods accordingly. Due to the identifiability problems with the Matérn α parameter (Lindgren et al., 2011), we assumed *a priori* that $\alpha = 2$ and therefore $\nu = 1$, and that the prior mean of \mathbf{f} was everywhere zero.

To verify the performance of this MCMC algorithm, we also conducted MCMC in our 1-D non-temporal model (see Section 1). In this model we also reparameterized the hyperparameter $\sigma_\beta^2 = \sigma_f^2 v_\beta$ to allow a Gibbs update for the emissions variance. Our air pollution model did not include covariates and therefore did not require this parameter. We also used different numbers of iterations in the two examples. Otherwise the algorithms were identical. Note that we did not conduct MCMC in our 1-D spatio-temporal model (see Section 3.5.2) due to computational constraints. In this case study, we only demonstrate the mechanics of the kriging step.

3.4.6 Case Studies

Case Study 1: Inferring a 1-Dimensional Spatial Source

We conducted an analysis in one-dimension of space, for a system at equilibrium, with a source governed by a linear model.

Suppose we make measurements of runoff pollution $u(s)$ at various points in a 1-D stream. The land alongside the stream may be used for several different purposes; for example, it may be agricultural, residential, or industrial. The primary source of variation in pollutant source $f(s)$ may be due to variations in land use (represented by a factor variable), so that $f(s) = \mathbf{x}(s)^T \boldsymbol{\beta} + \epsilon(s)$ where ϵ can be represented as a mean-zero GMRF and $\boldsymbol{\beta}$ has a normal prior. In this case study, to demonstrate the kriging step, we assumed all dynamic parameters were known and just calculated the kriging step. We applied the finite element method to Equation 3.3 to obtain \mathbf{K} and \mathbf{L} . Then, we calculated the joint conditional distribution for \mathbf{f} and $\boldsymbol{\beta}$ given measurements of $\mathbf{y} = \mathbf{A}\mathbf{u} + \boldsymbol{\epsilon}$ as described in Section 3.4.2 by first calculating the joint conditional distribution of \mathbf{u} and $\boldsymbol{\beta}$ then inverting Equation 3.4 to obtain the conditional distribution for \mathbf{f} .

To demonstrate the “kriging step” we first made predictions assuming the dynamics were known. Then, to demonstrate the accuracy of this kriging calculation, we calculated the average L^2 error across 30 simulations for each of 30 sampling densities ranging from a single data point up to 200,000 data points. For each sampling density, samples were evenly distributed across the interior of the fixed domain, but not into the buffer region. For the purposes of approximating the error we assumed that we observed nodes uniformly $\mathbf{A}^T \mathbf{A} = \mathbf{I}_{\text{int}}$ a fractional number of times equal to $n = N/m$, the true sample size divided by the number of rows of \mathbf{A} . Note that the relative error that we estimate here is purely interpolation error – all parameters are assumed known.

Finally, we conducted an MCMC analysis (described in Section 3.4.5) to demonstrate the efficacy of our inference in the air pollution case study. In this case study, we ran 4 chains for 5,000 iterations each, with a thinning rate of 5 for a total of 4,000 samples. We also

Parameter	α	ρ	D	r	σ_f^2	σ_ϵ^2	σ_β^2
Case Study 1	2	2	0.75	0.2	10	10	25
Case Study 2	4	5.3	0.25	0.05	10	10	NA

Table 3.1: The true parameters used for simulation in Case Studies 1 and 2.

used 5,000 iterations of burn-in. To match our air pollution case study, we specified weakly informative priors. Since we marginalized out the source, solution, and regression coefficients, we calculated these posterior distributions as generated quantities after completing the main MCMC sampling.

The true parameter values used for simulation in this case study can be found in Table 3.1.

Case Study 2: Inferring a 1-Dimensional Spatio-Temporal Source

In the previous section, we assumed a PDE in steady-state. Steady-state PDEs assume that dynamics and emissions are constant, or change on a time scale much longer than the time scale over which the steady-state solution is reached. However, in many applications these assumptions do not hold. In these cases, a full spatio-temporal model may be necessary. We provide a demonstration of this model, but we note that spatio-temporal models are computationally much more challenging, and we defer a more intensive spatio-temporal application for future work. In particular we demonstrate only the kriging step that is the foundation of these calculations, and not a more intensive MCMC analysis.

In this case study, we simulate a 1-D domain, with small enough dimensionality that computation remains feasible (101 spatial nodes over 2000 time steps, with simulation taking ≈ 1 second on a desktop computer with a 3.4 GHz 6-core processor and 32 GB of RAM. Full spatio-temporal emissions estimation in higher-dimensional spaces is a promising area of future study.

The true parameter values used for simulation in this case study can be found in Table 3.1.

Case Study 3: Inferring Sources of Air Pollution in the U.S.

Finally, we implemented the steady-state advection-diffusion-reaction equation to conduct source inference on $\text{PM}_{2.5}$ over the United States for January 1st 2019. We obtained $\text{PM}_{2.5}$ data from the EPA (US Environmental Protection Agency, 2021) and parameterized our wind-speed field using NOMADS (Rutledge et al., 2006). We assumed a homogeneous prior mean for \mathbf{f} , although population density, automobile density, land use, and many other covariates could also be included. We included a buffer region of approximately 4 degrees on all sides, and projected locations from degrees latitude/longitude to kilometers, preserving the distance between longitudes at all latitudes. We also projected advection velocities to ensure that the travel time between two points was preserved, although distances between two points and advection speeds were in general not preserved.

We assumed that the emissions (source function; $f(\mathbf{s})$) were governed by a Matérn Gaussian process and that the variance and range of this process were unknown, along with the diffusion and decay coefficients. We fit this model using an MCMC algorithm described in Section 3.4.5. In this case study we ran 4 chains for 25,000 iterations each, with a thinning rate of 5 for a total of 20,000 samples. We also used 25,000 iterations of burn-in.

We specified weakly informative priors based on existing literature (ρ (Song et al., 2018); D (Byun and Schere, 2006); r (Watson et al., 2000; Cao et al., 2013); σ_f^2 (Tucker, 2000); v (Peters et al., 2001)) and understanding of the mechanisms involved. Prior parameters and quantiles, and posterior statistics can be found in Table 3.2.

3.5 Results

3.5.1 1-Dimensional Spatial Source Results

Figure 3.2 shows the results of our kriging analysis, with three land-use zones applying differing constant contributions to the source. This demonstrates the ability to incorporate a linear regression model into our method and to obtain analytical conditional distributions when dynamics are perfectly known.

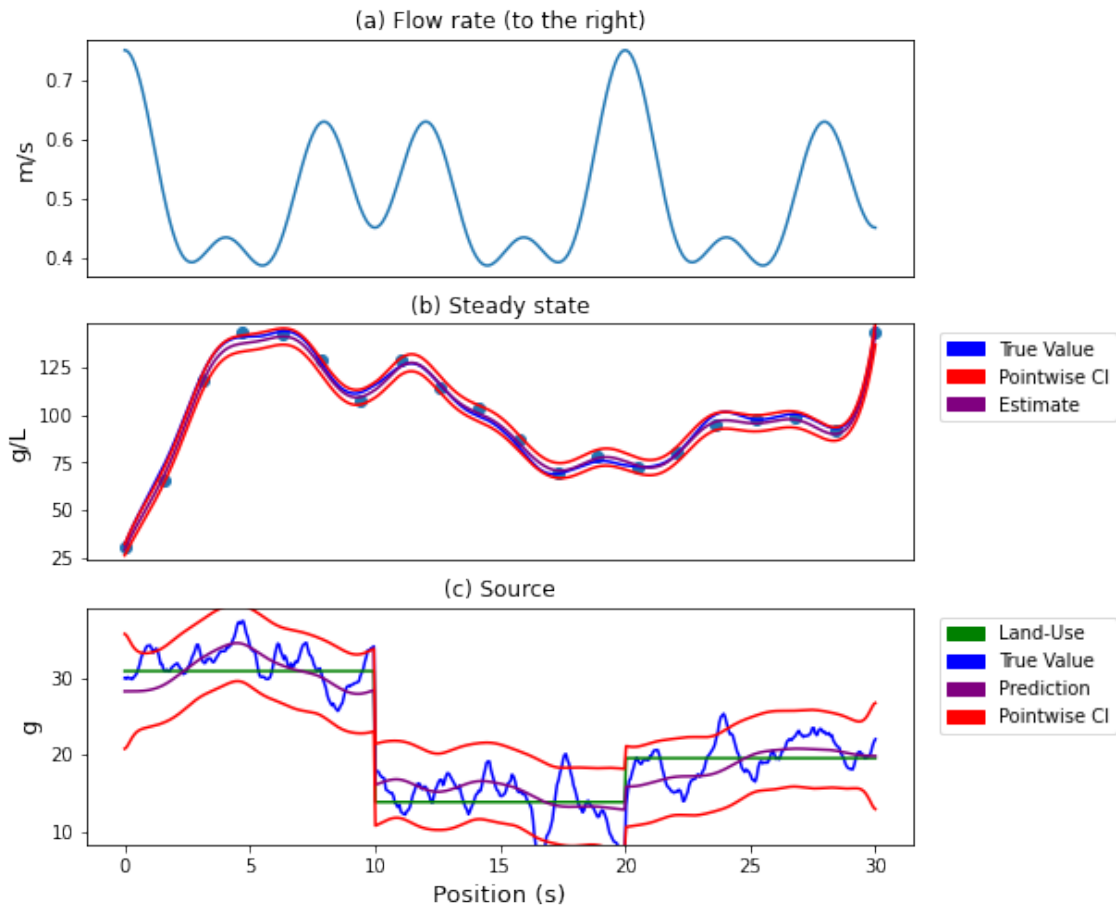


Figure 3.2: The results of a numerical experiment for conducting source reconstruction on the advection-diffusion-reaction equation, with a linear regression on spatial factors. Panel (a) shows the spatially-variable flow rate (always to the right). Panel (b) shows inference on \mathbf{u} . Panel (c) shows inference on $\mathbf{f} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$. Note the discontinuities in \mathbf{f} generated by discontinuities in \mathbf{X} . Dots show positions of measurements and their observed values.

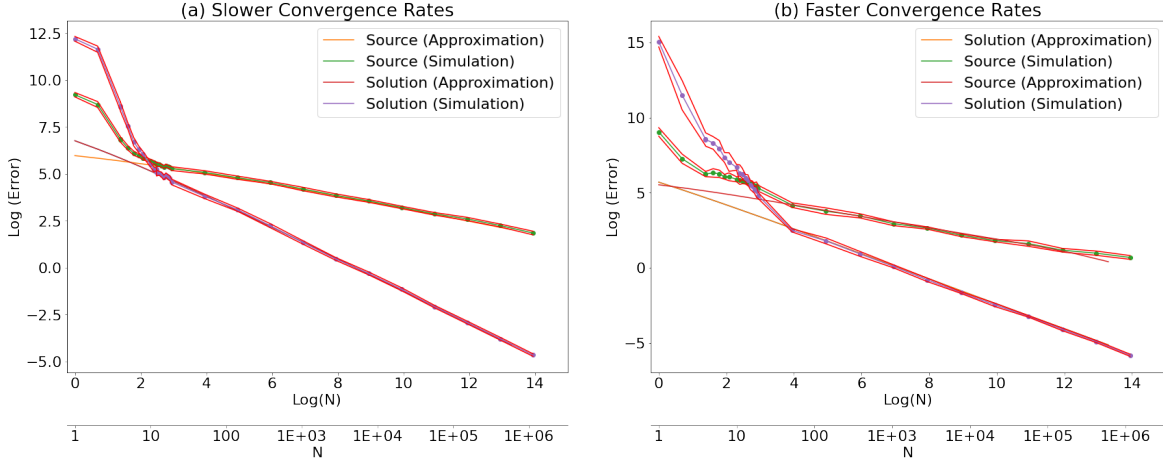


Figure 3.3: Empirical vs. approximate plots showing the convergence of error for our method when estimating both the source and solution of a PDE from solution-only measurements in Case Study 1. Note that our analytical approximations underestimate the error initially, but become accurate at a sample size of $N < 10$. In the first simulation (a), $D = 0.75, r = 0.2$, and $\sigma_\epsilon^2 = 5$, while in the second (b) $D = 0.075, r = 0.02$, and $\sigma_\epsilon^2 = 0.5$.

When quantifying the error incurred by kriging in this case study, we found that the local polynomial convergence rate approximations for the L^2 error agreed with our empirical calculations at all but the lowest sample sizes ($\lesssim 10$). We also found that across a wide range of moderate sample sizes, the error converges at a rate slower than $1/\sqrt{N}$. Specifically, we observed a convergence rate of approximately $N^{-0.4}$ when estimating concentrations (i.e., PDE solutions), and an error convergence rate of approximately $N^{-0.1}$ when estimating emissions (i.e., PDE sources). Hypothesizing that lower rates of convergence are the result of information being obscured by the dynamics, particular by diffusion and decay, and of observation error being too large to effectively “see past” these obfuscations, we performed the same set of simulations with the diffusion rate (D), decay rate (λ), and observation variance (σ_ϵ^2) all reduced by a factor of 10. In this case, we observed faster convergence rates of $N^{-0.41}$ and $N^{-0.19}$ respectively.

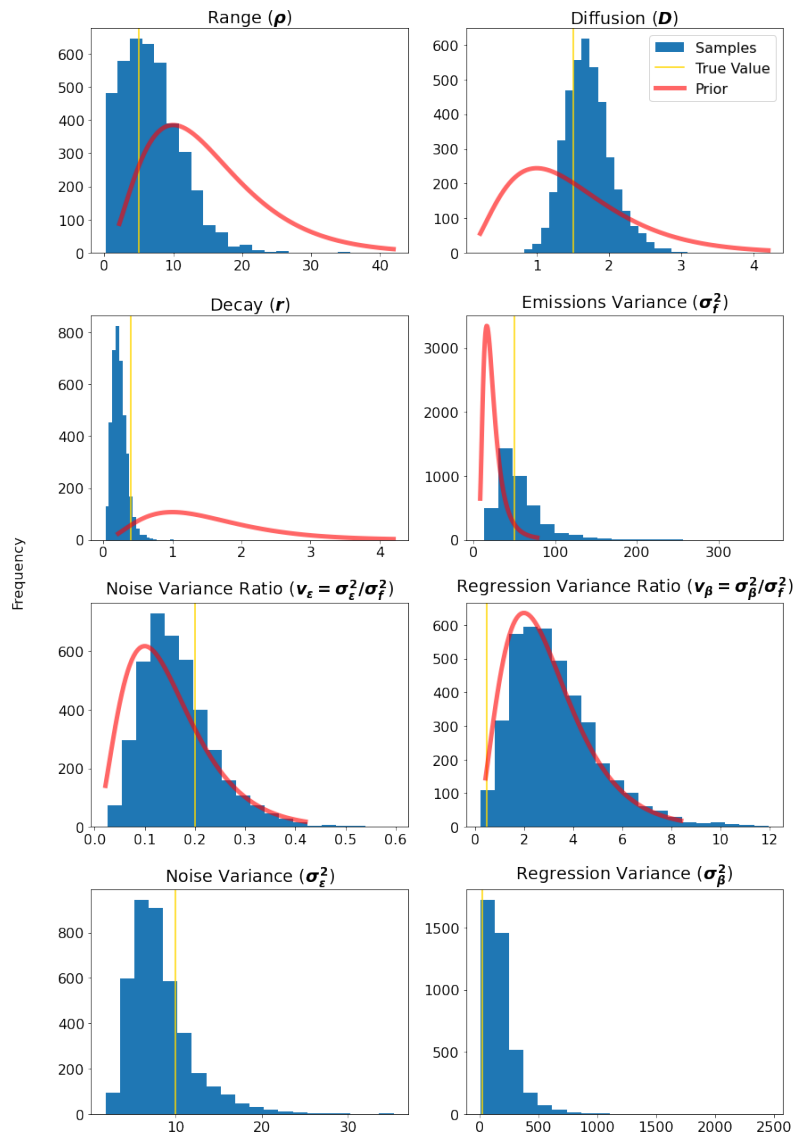


Figure 3.4: The results of Case Study 1's MCMC analysis. Prior PDF had been scaled up by the number of samples to be comparable to the histogram.




Figure 3.5: The results of Case Study 1's MCMC analysis. Black dotted lines represent the true simulated source and solution functions. Red dotted lines represent the pointwise posterior mean. Background lines represent the kriging estimate from each MCMC sample. Note that the solution (or concentration) is estimated precisely but the source (or emissions) is more uncertain. Both functions lie largely within the range of variation among samples.




Figure 3.6: The generated distributions for the source, solution, and regression coefficients given the posterior samples of the parameters. The vertical line represents the true value of the coefficient. Note that all three coefficients are underestimated on average, perhaps due to oversmoothing in the kriging step.

Upon conducting MCMC we found that some parameters could be inferred well while others suffered from issues of identifiability. The posteriors for the diffusion, decay, range, and emissions variance (shown in Figure 3.4) improved noticeably over the priors, although the decay showed a slight tendency toward underestimation. Range, diffusion, decay, and emissions variance were fairly well identified, while the two variance ratios were poorly identified. The unidentifiability of the regression variance ratio had little effect on the regression coefficients however. These three regression coefficients (shown in Figure 3.6) were somewhat underestimated, likely due to oversmoothing in the kriging estimates.

Due to the underestimation issues in the regression coefficients, we recommend interpreting model output largely in terms of overall trends and predictions of the source/solution functions. For example, in Figure 3.5, the model successfully identifies that the left-most zone has generally higher emissions than two zones on the right. In cases where accurate estimation of regression coefficients is a high priority, we recommend specifying non-Gaussian priors such as the Laplace or Horseshoe priors, and updating the coefficients in the MCMC step. Ultimately, the model allows us to infer the solution with high accuracy, and although the source is more uncertain, the inference remains accurate, with the true source falling well within the range of variation among the samples.

3.5.2 1-Dimensional Spatio-Temporal Source Results

Figure 3.7 shows the results of our analysis when we add a temporal dimension to the 1-D system analyzed in Numerical Case Study 1. We simulated the advection-diffusion equation with a spatio-temporal Matérn source as defined in equation (3.15). We simulated 10 sensors equidistant along the s -axis, and simulated 20 measurements each, equidistant in time.

The solution (concentration) estimate is more accurate than the source (emissions) estimate, because this is the scale at which we are able to take direct measurements. The source estimate displays a greater degree of smoothing due to the obscuring effects of diffusion and advection. Because the material detected in a concentration measurement could have originated from any point nearby or upstream, a larger number of measurements are required to

provide equivalent resolution on the source function. As more measurements are recorded, the reconstruction converges to the simulated source. Our method can be thought of as a filter that treats concentration measurements as an integral over previous, mostly upstream sources. That is, a measurement of concentration at time t only provides information about the source process up to time t , and provides much stronger information from upstream locations (due primarily to advection) than from downstream locations (due to diffusion). Information about future or downstream source values is primarily obtained from the correlation structure of the source process. Figure 3.7 panel (f) demonstrates this behavior on its margins, where downstream and future uncertainties are higher than upstream and past uncertainties.

3.5.3 Air Pollution Results

The results of this analysis can be visualized in Figure 3.8 and Table 3.2. Uncertainty was quantified using the square root of the diagonal elements of the posterior covariance matrix (i.e., the pointwise posterior standard deviations). Posterior uncertainty was lower around the main landmass of the continental US, because this is where the data were collected, and regions of high estimated emissions and concentration were clustered around some regions of high population including Texas, California, Florida, and the Pacific Northwest. Some anomalies are present in the results, notably some rare negative predictions for PM 2.5 concentrations. These are non-physical, but occur almost entirely outside the sampling region, likely due to the artificial oscillatory behavior induced by the FEM discretization of advection. These could possibly be reduced or removed by using Petrov-Galerkin elements or a full spatio-temporal model, but we leave these possibilities for future work.

The units of emissions in Case Study 3 are $\mu\text{g}/\text{m}^3/\text{hour}$, because concentration is measured as a density, in units of $\mu\text{g}/\text{m}^3$. Emissions measured in $\mu\text{g}/\text{hour}$ would be more intuitive, but would require making a homogeneity assumption to account for the vertical distribution of $\text{PM}_{2.5}$ (we use this approach to approximate a prior distribution for σ_f^2), or directly modeling vertical distribution with a fully 3-D model. Such a model is not difficult

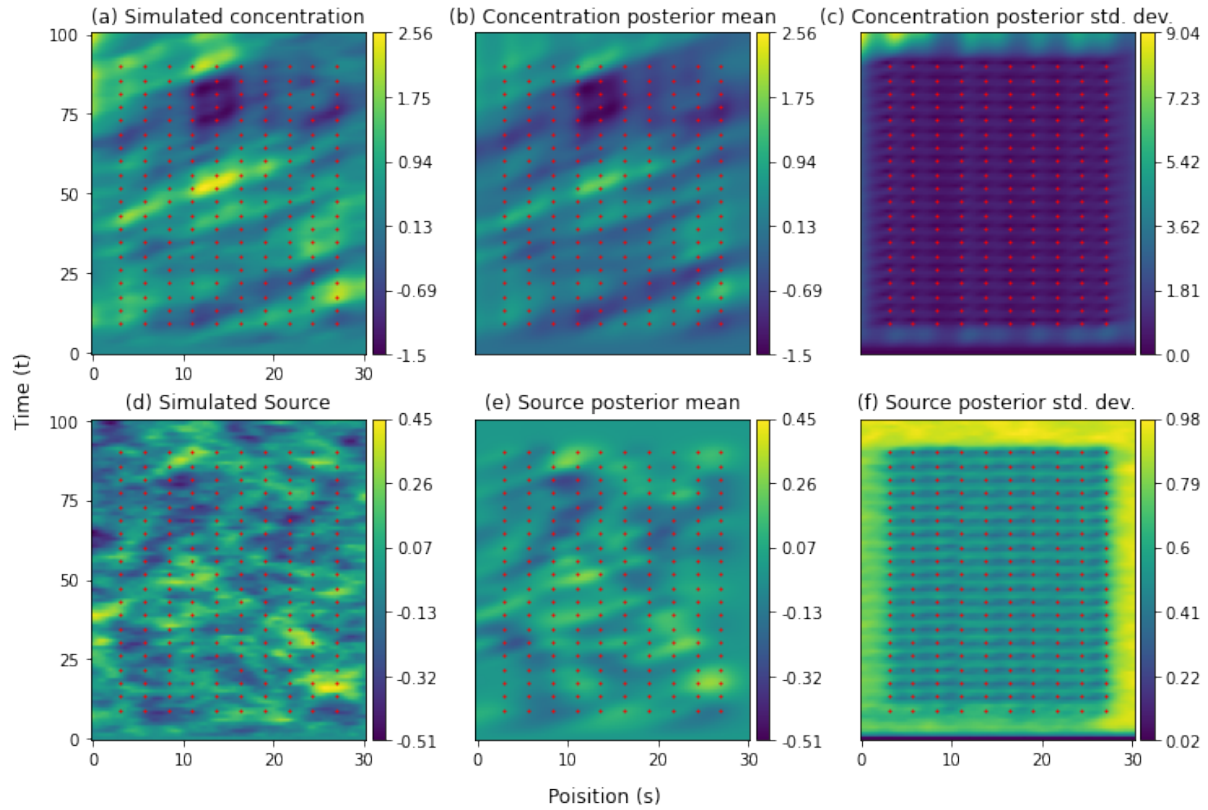


Figure 3.7: The results of Case Study 2. The source process was spatio-temporal Matérn as defined in Section 3.4.4 with parameters: $\tau = 2, \alpha = 4, \kappa = 1$. Flow rate was constant in time, and varied in space as shown in figure 3.2. Red points denote sampling locations. A spatial buffer zone of length 5 on each side is not included in the plot. No temporal buffer zone was removed.

Parameter	Post. Mean	Post. Std. Dev.	Prior	Prior Quantiles (.025,.975)
ρ	456 km	38 km	$\alpha = 4, \beta = 0.1$	(10,80)
D	8.2km ² /hr	2.8km ² /hr	$\alpha = 8.5, \beta = 1$	(3.6,14.4)
r	0.77/hr	0.42/hr	$\alpha = 1.36, \beta = 2.94$	(0.027,1.5)
σ_f^2	32 ($\mu\text{g}/\text{m}^3$) ²	44 ($\mu\text{g}/\text{m}^3$) ²	$\alpha = 1.1, \beta = 3.9$	(1,100)
v	2.1	1.7	$\alpha = 1.1, \beta = 0.13$	(0.33,30)
σ_ϵ^2	29.6 ($\mu\text{g}/\text{m}^3$) ²	1.7 ($\mu\text{g}/\text{m}^3$) ²	NA	NA
σ_f	5.0 $\mu\text{g}/\text{m}^3$	2.7 $\mu\text{g}/\text{m}^3$	NA	NA
σ_ϵ	5.4 $\mu\text{g}/\text{m}^3$	0.16 $\mu\text{g}/\text{m}^3$	NA	NA

Table 3.2: Posterior mean and standard deviations for each parameter of the advection-diffusion-reaction source model of U.S. air pollution. Also the priors, and prior quantiles for those parameters that had them. The bottom three parameters are generated quantities.

to mesh using FEniCS but is much larger and more computationally intensive, and we leave this extension for future work.

3.6 Conclusion

We presented a method for SPDE-based inference on the source of a linear PDE. We demonstrated the use of this method for reconstructing the source of an advection-diffusion-reaction process in a 1-D non-temporal setting, a 1-D spatio-temporal setting, and in a 2-D non-temporal case study of air pollution in the United States. We found that the method worked well to estimate concentrations, exhibiting close to a standard $1/\sqrt{N}$ convergence rate. We also found that the method worked well to estimate emissions, but with a reduced error rate that we hypothesized to be due to the destruction of information by diffusion processes. When applied in a case study of realistic complexity, we found that reasonable predictions were made, but we occasionally observed anomalies such as negative concentration. Although these were limited in spatial extent and had a negligible impact on the overall

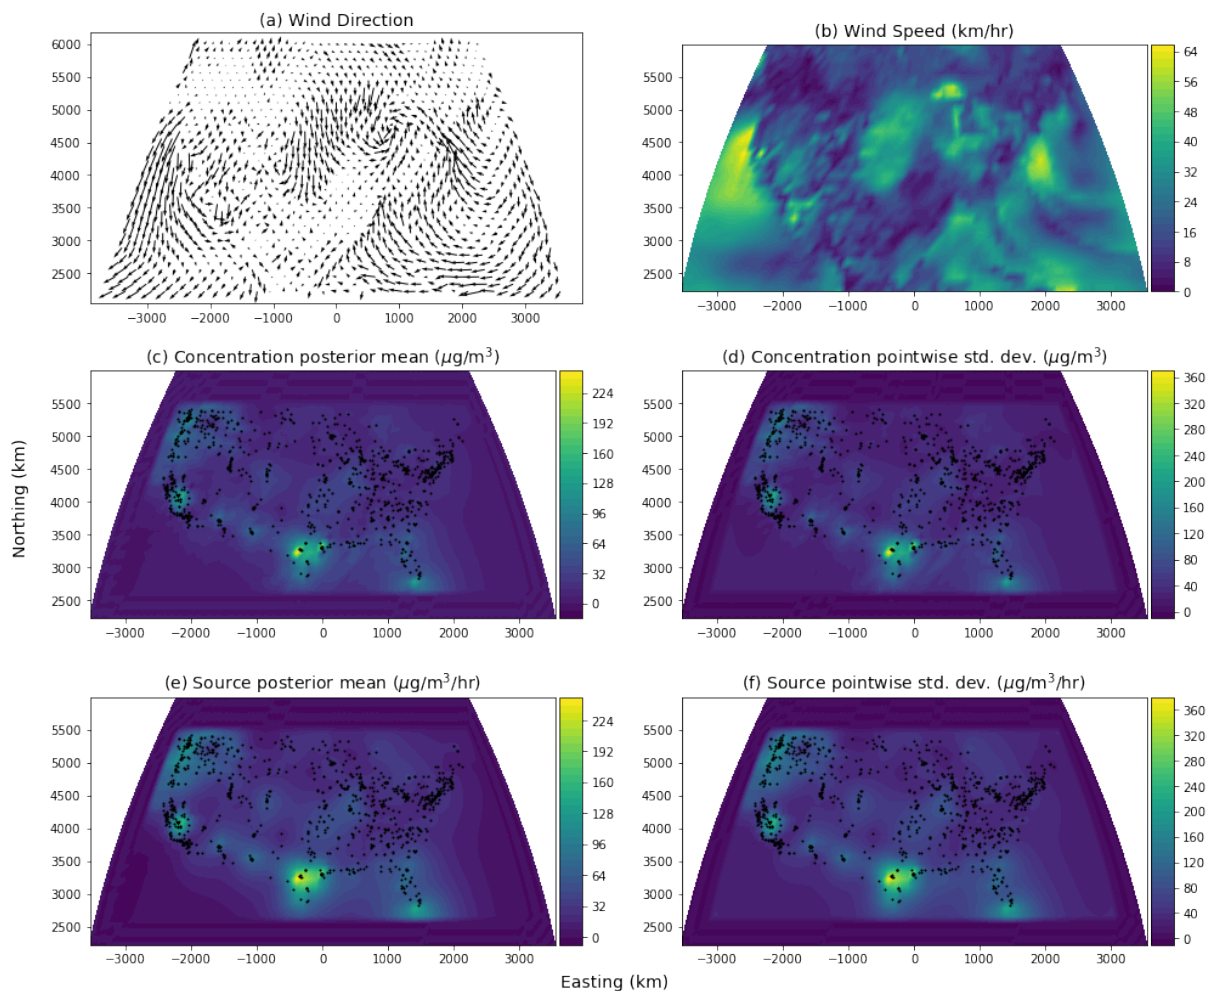


Figure 3.8: The main results of Case Study 3, parameterized with the NOMADS climate model (Rutledge et al., 2006), and using EPA PM 2.5 measurements (US Environmental Protection Agency, 2021). Advection field (a; upper left); advection speed (b; upper right); PM 2.5 map (c; middle left); PM 2.5 map posterior standard deviation (d; middle right); PM 2.5 emissions map (e; lower left); PM 2.5 source map posterior standard deviation (f; lower right). The buffer region (approximately 4 degrees wide) is set to zero to avoid showing boundary anomalies. Dots represent locations of data points.

analysis, we hypothesize that these artifacts could be reduced or eliminated in future work by using Petrov-Galerkin elements, or a full spatio-temporal model. We provide a framework through which our model can be extended to space-time, and through which efficient likelihood calculations can be conducted in space-time, but due to the challenges of conducting the marginalization step in space-time (specifically, maintaining sparsity when updating the Cholesky factor of \mathbf{Q}_u to add in $\frac{n}{\sigma_\epsilon^2} \mathbf{A}^T \mathbf{A}$), leave further spatio-temporal modeling to future work.

Possible applications of our method include inferring the source of environmental DNA measurements, non-point source pollution, or stable isotopes. Other important linear PDEs include Maxwell’s equations and the wave equation. Using these equations, our method could be applied to MRI measurements such as in Bao et al. (2002) or earthquake reconstruction such as in Das and Suhadolc (1996). Our approach uses an SPDE method similar to Lindgren et al. (2011) to solve the non-point source problem in a Bayesian framework, providing the opportunity to explicitly incorporate linear regression models on the source function. This framework also allows us to incorporate covariance functions for the source, and to calculate uncertainties, in contrast to many existing inverse problems approaches, which often operate under either a prediction-only framework (e.g., Ayvaz 2010), or under more complex frameworks such as the ensemble Kalman filter (e.g., Stroud et al. 2010). By incorporating our methodology and source-inference perspective into existing hierarchical frameworks (Wikle and Hooten, 2010), process uncertainty and non-Gaussian observations could also be accommodated, although the use of a link-function would warp the mechanistic interpretations we seek to preserve in this study. Our method has a variety of possible applications, and we have shown that it is both accurate, and implementable using reproducible general-purpose FEM software (Alnæs et al., 2015) that can easily accommodate arbitrary changes to the underlying PDE structure.

BIBLIOGRAPHY

- Alnæs, M., Blechta, J., Hake, J., Johansson, A., Kehlet, B., Logg, A., Richardson, C., Ring, J., Rognes, M. E., and Wells, G. N. (2015). The FEniCS project version 1.5. *Archive of Numerical Software*, 3(100).
- Alves, C., Kress, R., and Serranho, P. (2009). Iterative and range test methods for an inverse source problem for acoustic waves. *Inverse Problems*, 25(5):055005.
- Ayvaz, M. T. (2010). A linked simulation–optimization model for solving the unknown groundwater pollution source identification problems. *Journal of Contaminant Hydrology*, 117(1-4):46–59.
- Bakka, H. (2018). How to solve the stochastic partial differential equation that gives a Matérn random field using the finite element method. *arXiv preprint arXiv:1803.03765*.
- Bakka, H., Krainski, E., Bolin, D., Rue, H., and Lindgren, F. (2020). The diffusion-based extension of the Matérn field to space-time. *arXiv preprint arXiv:2006.04917*.
- Bao, G., Ammari, H., and Fleming, J. L. (2002). An inverse source problem for Maxwell’s equations in magnetoencephalography. *SIAM Journal on Applied Mathematics*, 62(4):1369–1382.
- Brenner, S. C., Scott, L. R., and Scott, L. R. (2008). *The mathematical theory of finite element methods*, volume 3. Springer.
- Brooks, A. N. and Hughes, T. J. (1982). Streamline upwind/Petrov-Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier-Stokes equations. *Computer Methods in Applied Mechanics and Engineering*, 32(1-3):199–259.

- Byun, D. and Schere, K. L. (2006). Review of the governing equations, computational algorithms, and other components of the Models-3 Community Multiscale Air Quality (CMAQ) modeling system. *Applied Mechanics Reviews*, 59(2):51–77.
- Cameletti, M., Gómez-Rubio, V., and Blangiardo, M. (2019). Bayesian modelling for spatially misaligned health and air pollution data through the INLA-SPDE approach. *Spatial Statistics*, 31:100353.
- Cao, J., Chow, J. C., Lee, F. S., Watson, J. G., et al. (2013). Evolution of PM_{2.5} measurements and standards in the US and future perspectives for China. *Aerosol and Air Quality Research*, 13(4):1197–1211.
- Chen, Y., Davis, T. A., Hager, W. W., and Rajamanickam, S. (2008). Algorithm 887: Cholmod, supernodal sparse cholesky factorization and update/downdate. *ACM Transactions on Mathematical Software (TOMS)*, 35(3):1–14.
- Das, S. and Suhadolc, P. (1996). On the inverse problem for earthquake rupture: The Haskell-type source model. *Journal of Geophysical Research: Solid Earth*, 101(B3):5725–5738.
- Deiner, K., Bik, H. M., Mächler, E., Seymour, M., Lacoursière-Roussel, A., Altermatt, F., Creer, S., Bista, I., Lodge, D. M., de Vere, N., et al. (2017). Environmental DNA metabarcoding: Transforming how we survey animal and plant communities. *Molecular Ecology*, 26(21):5872–5895.
- Doi, H., Inui, R., Akamatsu, Y., Kanno, K., Yamanaka, H., Takahara, T., and Minamoto, T. (2017). Environmental DNA analysis for estimating the abundance and biomass of stream fish. *Freshwater Biology*, 62(1):30–39.
- Doi, H., Uchii, K., Takahara, T., Matsushashi, S., Yamanaka, H., and Minamoto, T. (2015). Use of droplet digital pcr for estimation of fish abundance and biomass in environmental DNA surveys. *PloS one*, 10(3):e0122763.

- El Badia, A. and Ha-Duong, T. (2002). On an inverse source problem for the heat equation. application to a pollution detection problem. *Journal of inverse and ill-posed problems*, 10(6):585–599.
- El Badia, A., Ha-Duong, T., and Hamdi, A. (2005). Identification of a point source in a linear advection–dispersion–reaction equation: Application to a pollution source problem. *Inverse Problems*, 21(3):1121.
- El Badia, A. and Hamdi, A. (2007). Inverse source problem in an advection-dispersion-reaction system: Application to water pollution. *Inverse Problems*, 23(5):2103.
- Engl, H. W., Hanke, M., and Neubauer, A. (1996). *Regularization of Inverse Problems*, volume 375. Springer Science & Business Media.
- Grech, R., Cassar, T., Muscat, J., Camilleri, K. P., Fabri, S. G., Zervakis, M., Xanthopoulos, P., Sakkalis, V., and Vanrumste, B. (2008). Review on solving the inverse problem in EEG source analysis. *Journal of Neuroengineering and Rehabilitation*, 5(1):25.
- Hao, S., Li, F., Li, Y., Gu, C., Zhang, Q., Qiao, Y., Jiao, L., and Zhu, N. (2019). Stable isotope evidence for identifying the recharge mechanisms of precipitation, surface water, and groundwater in the ebinur lake basin. *Science of the Total Environment*, 657:1041–1050.
- Harrison, J. B., Sunday, J. M., and Rogers, S. M. (2019). Predicting the fate of edna in the environment and implications for studying biodiversity. *Proceedings of the Royal Society B*, 286(1915):20191409.
- Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498.

- Lushi, E. and Stockie, J. M. (2010). An inverse Gaussian plume approach for estimating atmospheric pollutant emissions from multiple point sources. *Atmospheric Environment*, 44(8):1097–1107.
- Monk, P. et al. (2003). *Finite element methods for Maxwell's equations*. Oxford University Press.
- Moraga, P., Cramb, S. M., Mengersen, K. L., and Pagano, M. (2017). A geostatistical model for combined analysis of point-level and area-level data using INLA and SPDE. *Spatial Statistics*, 21:27–41.
- Panofsky, H. A. (1969). Air pollution meteorology. *American Scientist*, 57(2):269–285.
- Peters, T. M., Norris, G. A., Vanderpool, R. W., Gemmill, D. B., Wiener, R. W., Murdoch, R. W., Mcelroy, F. F., and Pitchford, M. (2001). Field performance of PM_{2.5} federal reference method samplers. *Aerosol Science & Technology*, 34(5):433–443.
- Porter, R. P. and Devaney, A. J. (1982). Holography and the inverse source problem. *Journal of the Optical Society of America*, 72(3):327–330.
- Rachel, B.-J., Pearson, T. E., Ramos, F. C., Grimes, C. B., and Bruce MacFarlane, R. (2008). Tracking natal origins of salmon using isotopes, otoliths, and landscape geology. *Limnology and Oceanography*, 53(4):1633–1642.
- Rees, H. C., Maddison, B. C., Middleditch, D. J., Patmore, J. R., and Gough, K. C. (2014). The detection of aquatic animal species using environmental DNA—a review of eDNA as a survey tool in ecology. *Journal of applied ecology*, 51(5):1450–1459.
- Rourke, M. L., Fowler, A. M., Hughes, J. M., Broadhurst, M. K., DiBattista, J. D., Fielder, S., Wilkes Walburn, J., and Furlan, E. M. (2022). Environmental DNA (eDNA) as a tool for assessing fish biomass: A review of approaches and future considerations for resource surveys. *Environmental DNA*, 4(1):9–33.

- Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*. CRC Press.
- Rutledge, G. K., Alpert, J., and Ebisuzaki, W. (2006). NOMADS: A climate and weather model archive at the National Oceanic and Atmospheric Administration. *Bulletin of the American Meteorological Society*, 87(3):327–342.
- Sigrist, F., Künsch, H. R., and Stahel, W. A. (2012). A dynamic nonstationary spatio-temporal model for short term prediction of precipitation. *The Annals of Applied Statistics*, pages 1452–1477.
- Sigrist, F., Künsch, H. R., and Stahel, W. A. (2015). Stochastic partial differential equation based modelling of large space–time data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77(1):3–33.
- Song, W., Jia, H., Li, Z., and Tang, D. (2018). Using geographical semi-variogram method to quantify the difference between NO_2 and $\text{PM}_{2.5}$ spatial distribution characteristics in urban areas. *Science of the Total Environment*, 631:688–694.
- Stockie, J. M. (2011). The mathematics of atmospheric dispersion modeling. *SIAM Review*, 53(2):349–372.
- Stroud, J. R., Stein, M. L., Lesht, B. M., Schwab, D. J., and Beletsky, D. (2010). An ensemble Kalman filter and smoother for satellite data assimilation. *Journal of the American Statistical Association*, 105(491):978–990.
- Tucker, W. G. (2000). An overview of $\text{PM}_{2.5}$ sources and control strategies. *Fuel Processing Technology*, 65:379–392.
- US Environmental Protection Agency (2021). Air Quality System Data Mart. Internet Database available at <http://www.epa.gov/ttn/airs/aqsdatamart>. Accessed: 2021-02-24.

- Ver Hoef, J. M., Peterson, E., and Theobald, D. (2006). Spatial statistical models that use flow and stream distance. *Environmental and Ecological Statistics*, 13(4):449–464.
- Watson, J. G., Chow, J. C., and Pace, T. G. (2000). Fugitive dust emissions. *Crops*, 3(14):7.
- Wikle, C. K. and Hooten, M. B. (2010). A general science-based framework for dynamical spatio-temporal models. *Test*, 19(3):417–451.
- Xepapadeas, A. (2011). The economics of non-point-source pollution. *Annual Review of Resource Economics*, 3(1):355–373.
- Yan, L., Fu, C.-L., and Yang, F.-L. (2008). The method of fundamental solutions for the inverse heat source problem. *Engineering Analysis with Boundary Elements*, 32(3):216–222.

Chapter 4

OPTIMAL SAMPLING DESIGN UNDER LOGISTICAL CONSTRAINTS WITH MIXED INTEGER PROGRAMMING

Publication history: This chapter is to be submitted for review at a peer-reviewed journal.

4.1 Abstract

The goal of survey design is often to minimize the errors associated with inference: the total of bias and variance. Random surveys are common because they allow the use of theoretically unbiased estimators. In practice however, such design-based approaches are often unable to account for logistical or budgetary constraints. Thus, they may result in samples that are logistically inefficient (i.e., more or more informative samples could be collected at no extra cost), or infeasible to implement. Various balancing and optimal sampling techniques have been proposed to improve the statistical efficiency of such designs, but few studies have investigated improving the logistical efficiency of such designs. We introduce a flexible mixed integer linear program (MILP) for optimal sampling design, capable of incorporating a variety of logistical constraints and a wide class of Bayesian regression models. We demonstrate the use of our model on three spatial sampling problems of increasing complexity, culminating in a scenario developed to model real logistics in conjunction with the US Forest Service Forest Inventory and Analysis survey of Tanana, Alaska. Our methodological contribution to survey design is significant because the proposed modeling framework makes it possible to generate high-quality sampling designs and inferences while guaranteeing the satisfaction of practical constraints that can be flexibly defined by the user. The technical novelty of the method is the integration of explicit statistical modeling into combinatorial optimization. This integration might allow a paradigm shift toward high-quality survey designs that explicitly incorporate practical constraints.

4.2 Introduction

Since the early 1900s, random sampling methods have been preferred for their unbiased estimators, derived from the controlled introduction of randomness through the design process (Rao and Bellhouse, 1990). However, these unbiased estimators rely on the assumption that any random sample generated by the design process would in fact be possible to collect.

Under complex logistical and budgetary constraints, this assumption may not always hold. Some methods, such as cluster sampling, have arisen to handle particular types of logistical constraints (Lohr, 2019), but we are not aware of a randomized method capable of flexibly guaranteeing satisfaction of a wide range of constraints. Practical, consistent, and rigorous implementation of complex surveys subject to budgetary constraints can therefore require a great deal of effort, and sometimes require making compromises between these factors (Hughes and Peck, 2008). In large-scale spatial surveys in particular, heterogenous costs, access to sampling sites, transportation of crews and samples, and other factors can complicate the implementation of a design. In this chapter, we use spatial sampling as an example to demonstrate and quantify the performance of our novel method for conducting a model-based optimal sampling design, in which logistics are explicitly modeled and feasibility of the resulting design is guaranteed.

Apart from complex logistical and budgetary constraints, spatial sampling poses a second major challenge, spatial autocorrelation, which can be summarized by the proverbial Tobler’s Law that “everything is related to everything else, but near things are more related than distant things” (Tobler, 1970). In the absence of independent samples, the quality of inference is therefore dependent upon the entire set of sampled locations, rather than a summary statistic such as the sample size. This adds an additional layer of complexity. The premise of our model is that combinatorial optimization can serve as a mathematical frame within which statistical models as well as the practical constraints of sampling can be captured simultaneously and explicitly, thereby providing high-quality designs that are guaranteed to be feasible on the ground.

In particular, we propose to use a C -criterion as the objective function of a discrete mathematical program. This criterion can be expressed as $\mathbf{v}^T \mathbf{M}^{-1} \mathbf{v}$ for a fixed coefficient vector \mathbf{v} and a matrix \mathbf{M} which is optimized by the model. In the proposed application of spatial sampling, the vector \mathbf{v} can be chosen in such a way that this criterion measures uncertainty for an estimate of an areal mean, which is a common target of inference for randomized designs (Wang et al., 2012). Furthermore, because we express our method as a

MILP, logistics can be modeled with a great deal of complexity and sophistication. MILP models are a well-developed branch of operations research, and have for decades been used to model a wide range of industrial applications (Dantzig, 1957).

4.2.1 Motivation

Our method is motivated by the following case study. The United States Forest Service (USFS) conducts a regular Forest Inventory and Analysis (FIA) program across the United States, in which the report forest area, tree growth and mortality, and land ownership, among many other metrics. Sampling forests in remote regions such as interior Alaska, however, poses a logistical challenge. The USFS has developed a complex two-phase design and post-stratified estimation procedure for these regions (Chapter 4, Bechtold and Patterson 2005). Although the resulting data are flexible and may be used for a wide variety of purposes, the design has become increasingly difficult to implement as more remote regions need to be sampled. This challenge has motivated our development of an alternative, model-based method of design that can guarantee feasibility.

Our model can also be applied in a broad range of other scenarios. MILPs can be used to generate detailed optimal routes for vehicles such as remote sensing aircraft (e.g. Avellar et al. 2015), or fleets of autonomous aquatic robots (e.g. Yilmaz et al. 2008), or to plan the placement of stationary sensors (e.g. Kwak et al. 2022) or watercraft inspection sites (e.g. Fischer et al. 2021). Under these and other sufficiently complex logistics, randomized designs may be difficult to practically collect (e.g. Hughes and Peck 2008), and even if the design is feasible to collect, it may ultimately give less information than an optimized design.

4.2.2 Summary

The paper is organized as follows. Section 4.3 is an overview of the relevant background material in sampling design, spatial statistics, and optimization. We review other optimal sampling methods; the spatial statistical methods that we rely on in our simulations; and optimization methods that can be used to solve our model. In Section 4.4 we construct

our model and outline how it differs from the existing literature. We show how our design criterion can be calculated as a MILP objective function, and discuss the design-based benchmarks that we use to test the efficacy of our model. We also discuss the details of the three increasingly complex logistics scenarios we use in the benchmarking study. In Section 4.5, we present the results of our benchmarking study, and show example designs for each of the three scenarios, under each of the simulated design methods. Finally in Section 4.6 we draw conclusions about the performance of our method and outline next steps.

4.3 Background

Spatial sampling design is a well-studied problem with a variety of existing methods. To place our method in context, we discuss: (1) the set of design-based sampling alternatives; (2) the set of model-based optimal sampling alternatives; and (3) MILP models, their capabilities, and common methods of solving them. We then discuss in Section 4.4 our specific model choices.

4.3.1 Random Sampling Methods

Simple random sampling (SRS) is the foundation of most random sampling methods. It consists of selecting samples independently at random from the set of all possible samples with equal probability of inclusion although unequal probabilities of inclusion is a natural extension (Lohr, 2019). However, one of the core principles of spatial statistics is the presence of autocorrelation (Tobler, 1970). Spatial surveys are therefore often designed to account for this autocorrelation by ensuring that points are spatially balanced. Design-based inferences from SRS are still valid when locations are autocorrelated, but are generally less precise than if autocorrelation was accounted for.

An alternative is stratified random sampling (StratRS). In this method, the space is divided into a set of disjoint subspaces (called strata), and each of these is sampled separately. This forces spatial balance on the large scale by ensuring points are evenly distributed between regions, but does not enforce spatial balance on the small scale, because balance

within strata is not necessarily controlled. A third method, systematic sampling, similarly enforces large scale spatial balance, by sampling on a regular grid, often square, hexagonal, or triangular (White et al., 1992) and often with a randomized starting point (Bellhouse, 2005).

A more complex method of achieving spatial balance in a randomized method is through any of a number of so-called spatially balanced random sampling (SBRS) methods. SBRS methods include generalized random tessellation stratified (GRTS) sampling (Stevens Jr and Olsen, 2004) and balanced acceptance sampling (BAS; Robertson et al. 2013). GRTS sampling operates by mapping two-dimensional space onto one-dimensional space, selecting a systematic sample of 1-D space with a random starting point, then mapping the sample back to two-dimensional space, while BAS leverages the Halton sequence to balance location in any number of dimensions. Both GRTS and BAS generally account for spatial autocorrelation equally well (Kermorvant et al., 2019).

None of these methods take account of logistical or budgetary constraints, although all except systematic sampling are able to accommodate unequal inclusion probabilities, which can be used to adjust for (but not guarantee satisfaction of) a knapsack constraint (a single constraint assigning a cost to each sample and an overall budget). In the presence of more complex constraints feasibility cannot necessarily be guaranteed, and even if it can, the main way feasibility can be attained is by decreasing sample size and therefore quality of inference.

4.3.2 Optimal Sampling Methods

An optimization problem in general consists of two components: an objective function that one seeks to either minimize or maximize; and a set of constraints that must be met by the optimum. Although a problem can be defined with only these two components, a third component is also essential in practice: namely the choice of method by which the optimal solution will be found or approximated.

Because optimization problems frequently include a set of non-trivial constraints, it is natural to consider using an optimization method when feasibility is a primary concern.

However, in the case of optimal sampling design, most existing methods treat constraints as only a secondary consideration.

Previous studies have focused on the choice of objective function (frequently termed “design criterion” in this context). Some common criteria are the A -, D -, E -, and T -criteria (Ch. 6, Pukelsheim 2006). These objective functions can be interpreted as minimizing, for example, the average variance of a set of model parameters (the A -criterion), or the determinant of a dispersion matrix (the D -criterion). They can also be interpreted more abstractly as ways of mapping matrices to “sizes.” The D -criterion can for example be defined in terms of an $n \times n$ matrix \mathbf{M} as $\det(\mathbf{M})^{1/n}$, while the A -criterion can be defined as $(\frac{1}{n}\text{trace } \mathbf{M}^{-1})^{-1}$. Each of these criteria is distinct and will in general imply a different optimal sample.

Our method relies on a variant of the C -criterion, which can be interpreted as minimizing the conditional variance of an average of model parameters, and can be expressed as an inner product $\mathbf{v}^T \mathbf{M}^{-1} \mathbf{v}$ for a constant vector \mathbf{v} . One challenge when applying the C -criterion is how to choose the vector \mathbf{v} . We propose that when a spatial function is represented by a vector of model parameters, such as in the SPDE approach, \mathbf{v} can be chosen naturally to minimize the variance of an estimate of an areal mean.

The set of constraints that can be applied depends upon the method being used to find or approximate a solution. For example, the popular Broyden-Fletcher-Goldfarb-Shanno algorithm requires modification to be able to incorporate even simple box-type constraints (Byrd et al., 1995). In general, integer constraints are particularly challenging to apply, because the domain becomes discontinuous. In the context of optimal sampling, a variety of methods have been applied. For example, in a small design space (such as when only a single location is to be sampled) the optimal sample may be found by enumeration (e.g. Leach et al. 2022). Such a method can be extended to produce a greedy algorithm which sequentially selects the best single location to add to a sample (e.g. Wynn 1970; Diaz et al. 2018). In certain cases, the greedy algorithm may even have provably near-optimal performance, although in these cases the set of modelable constraints is often limited, for

example to knapsack and matroid-type constraints (Krause et al., 2008; Lee et al., 2009). Other models also admit knapsack-type constraints, such as the Elfving method (Elfving, 1952; Stoica and Babu, 2010). For complex models, methods such as genetic algorithms (e.g. Dupont et al. 2021) or simulated annealing (e.g. Van Groenigen et al. 1999) can be used to produce high-quality designs, under logistical constraints, although these generally provide no guarantee as to the quality of the design relative to the optimum.

4.3.3 *Mixed Integer Linear Programming*

Problem Formulation

Linear programming encompasses the methods used to solve optimization problems of the following form, termed linear programs (LPs):

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \mathbf{c}^T \mathbf{x}, \\ & \text{subject to} && \mathbf{A}^T \mathbf{x} \leq \mathbf{b}. \end{aligned} \tag{4.1}$$

Many common optimization problems can be cast as LPs, including the knapsack problem, the diet problem, the multicommodity flow problem, and problems involving manufacturing of multiple types of goods (Chapter 1; Bertsimas and Tsitsiklis 1997). Linear programs are therefore used frequently in industrial applications, where they are solved using industrial solvers such as CPLEX (CPLEX, 2020) and Gurobi (Gurobi Optimization, LLC, 2023). In practice, they are often solved by the simplex method.

Logistics involving integer or binary constraints in addition to linear constraints cannot be modeled using linear programs. These problems may instead be modeled using mixed integer linear programs (MILPs). The addition of integer and binary constraints allows a wider range of logistics to be modeled including among many others the traveling salesman problem, and a variety of more complex manufacturing, scheduling, and transportation problems. Our model is a MILP and can therefore accommodate these and many other logistical constraints.

Branch and Bound

MILPs are difficult to solve computationally (being in general NP-hard), but due to their importance in industry, techniques have been developed to find high-quality solutions nevertheless. By default industrial solvers such as CPLEX and Gurobi use the branch and bound algorithm (Land and Doig, 2010), which allow these solvers to both find high-quality solutions, and calculate bounds on how much better the true optimal solution could be.

Branch and bound (B&B) is the main algorithm used in the solution of NP-hard combinatorial optimization problems (Clausen, 1999). In general, B&B can be seen as more of a paradigm than a specific algorithm, however in the case of MILPs, the industrial solvers have developed tunable B&B algorithms that work well on a wide variety of problems. The B&B paradigm generally works by dynamically generating a tree whose root is the set of all possible solutions and whose branches are subsets thereof. Each time a branch is generated (often by fixing the value of a particular discrete variable), a function is evaluated to bound the objective function values that may be obtained within that branch. In the case of MILPs, the bounding function is usually a relaxation of the integrality constraints to produce a more easily solved LP, the solution of which provides a bound on the MILP solution. Periodically, heuristic methods are used to generate solutions from within the various branches as they are explored, and the best solution is maintained at any given time. Often, a so-called “warmstart” solution which is expected to be of reasonably high-quality is provided at the start of the algorithm. Solutions are compared to bounds, and when a solution is found of higher quality than a branch’s bound, the branch is pruned from the tree. Over time it is often possible to prune large portions of the tree, and to dynamically prove a global bound on the optimal solution (Land and Doig, 2010). Comparing this global bound with the current candidate solution provides a dynamically updated “optimality gap” often expressed in both absolute and relative terms. This ability to improve the optimality gap over time, until either a certain gap is achieved or until a time limit is reached, is an appealing property of the B&B paradigm.

Linearizing a Binary-Continuous Product

Our formulation involves calculating the product z_i of a binary decision variable x_i with a linear combination of continuous variables $\mathbf{a}_i^T \mathbf{y}$. This operation, although non-linear, can be carried out within a MILP through the use of additional constraints (Chapter 9; Williams 2013), so long as the range of $\mathbf{a}_i^T \mathbf{y}$ is bounded both above and below. This is accomplished by prescribing the following set of standard constraints:

$$\begin{aligned} z_i &\leq \mathbf{a}_i^T \mathbf{y} - (1 - x_i)L, \\ z_i &\geq \mathbf{a}_i^T \mathbf{y} - (1 - x_i)U, \\ z_i &\leq x_i U, \\ z_i &\geq x_i L. \end{aligned}$$

When x_i is equal to 1 the latter two constraints simply enforce the *a priori* upper and lower bounds $L \leq \mathbf{a}_i^T \mathbf{y} \leq U$, while the former two enforce the desired equality $z_i = \mathbf{a}_i^T \mathbf{y}$. Meanwhile, when x_i is equal to 0, the latter two enforce the desired equality $z_i = 0$ while the former two bounds are loose because again, by the *a priori* bounds $0 \leq \mathbf{a}_i^T \mathbf{y} - L$ and $0 \geq \mathbf{a}_i^T \mathbf{y} - U$. This technique further expands the set of possible problems which may be modeled using MILPs to include those with bounded binary-linear products.

4.4 Methods

4.4.1 Design Criterion

Consider a set of parameters \mathbf{u} , jointly modeled as multivariate normal with precision matrix \mathbf{Q} . We assume without loss of generality that the prior mean is 0. Suppose that a set of linear combinations $\mathbf{a}_i^T \mathbf{u}$ are candidates to be sampled, and that the set of observations $d_i = \mathbf{a}_i^T \mathbf{u} + \epsilon_i$ are obscured by a vector of i.i.d Gaussian noise ϵ , with known variances σ_i^2 . Finally, we assume that the object of inference is a known linear combination $\mathbf{v}^T \mathbf{u}$. Our design criterion is the conditional variance C^2 of $\mathbf{v}^T \mathbf{u}$ given a the chosen set of observations \mathbf{d} .

Under this formulation, our model can be written as follows, where x_i is a set of indicator variables denoting whether the observation d_i is sampled:

$$\mathbf{d} = \mathbf{A}^T \mathbf{u} + \boldsymbol{\epsilon} \quad (4.2)$$

$$\mathbf{u} | \mathbf{d} = N \left(0, \mathbf{Q} + \sum_{i=1}^n \frac{x_i}{\sigma_i^2} \mathbf{a}_i \mathbf{a}_i^T \right) \quad (4.3)$$

$$C^2 = \mathbf{v}^T \left(\mathbf{Q} + \sum_{i=1}^n \frac{x_i}{\sigma_i^2} \mathbf{a}_i \mathbf{a}_i^T \right)^{-1} \mathbf{v}. \quad (4.4)$$

In the case of spatial sampling, the vector \mathbf{u} may represent a Gaussian random field $u(\mathbf{s})$ over a spatial coordinate \mathbf{s} , discretized according to the SPDE method (Lindgren et al., 2011), in which case the vector \mathbf{v} may be chosen to approximate a spatial integral over $u(\mathbf{s})$, and the precision matrix \mathbf{Q} is sparse. Modifications to our method to make use of covariance matrices and to allow uncertainty in the covariance function are discussed in Appendix 3.

Our objective function is a Bayesian formulation of the C -criterion, which would more typically use the Moore-Penrose pseudoinverse of an information matrix (e.g. Sagnol et al. 2010) rather than a precision matrix. Using our formulation, we may linearize our objective function by defining the following set of intermediate variables:

$$\mathbf{y} = \left(\mathbf{Q} + \sum_{i=1}^n \frac{x_i}{\sigma_i^2} \mathbf{a}_i \mathbf{a}_i^T \right)^{-1} \mathbf{v} \quad (4.5)$$

$$z_i = x_i (\mathbf{a}_i^T \mathbf{y}), \quad (4.6)$$

so that our objective function may be written as $\mathbf{v}^T \mathbf{y}$, under the constraint

$$\mathbf{Q} \mathbf{y} + \frac{1}{\sigma^2} \mathbf{A}^T \mathbf{z} = \mathbf{v}. \quad (4.7)$$

This allows us to optimize the decision variables x_i using a MILP.

4.4.2 Bounds

As outlined in Section 4.3, to calculate the product variables \mathbf{z} it is necessary for all elements of $\mathbf{A}^T \mathbf{y}$ to be bounded above and below. Such bounds may be derived using the Cauchy-Schwartz inequality, because the precision matrix $\mathbf{Q} + \sum \frac{x_i}{\sigma_i^2} \mathbf{a}_i \mathbf{a}_i^T$ defines an inner product

$\langle \mathbf{u}, \mathbf{v} \rangle_X = \mathbf{u}(\mathbf{Q} + \sum \frac{x_i}{\sigma_i^2} \mathbf{a}_i \mathbf{a}_i^T)^{-1} \mathbf{v}$ such that $\langle \mathbf{u}, \mathbf{u} \rangle_{\mathbf{x}+\boldsymbol{\zeta}} \leq \langle \mathbf{u}, \mathbf{u} \rangle_{\mathbf{x}}$, where $\boldsymbol{\zeta}$ and $\mathbf{x} + \boldsymbol{\zeta}$ are both 0/1-vectors. Thus,

$$\mathbf{z}_i^2 = \langle \mathbf{a}_i, \mathbf{v} \rangle_{\mathbf{x}}^2 \leq \langle \mathbf{a}_i, \mathbf{a}_i \rangle_{\mathbf{x}} \langle \mathbf{v}, \mathbf{v} \rangle_{\mathbf{x}} \leq \langle \mathbf{a}_i, \mathbf{a}_i \rangle_0 \langle \mathbf{v}, \mathbf{v} \rangle_0.$$

In fact we may strengthen these bounds by noting that $\langle \mathbf{v}, \mathbf{v} \rangle$ is our objective function, so that at optimality

$$\mathbf{z}_i^2 \leq \langle \mathbf{a}_i, \mathbf{a}_i \rangle_{\mathbf{x}_{\text{opt}}} \langle \mathbf{v}, \mathbf{v} \rangle_{\mathbf{x}_{\text{opt}}} \leq \langle \mathbf{a}_i, \mathbf{a}_i \rangle_0 \langle \mathbf{v}, \mathbf{v} \rangle_{\mathbf{x}_{\text{feas}}}.$$

Thus the bounds on \mathbf{z}_i can be strengthened whenever an improved feasible solution is found. To allow a single model to be run, we use our greedy algorithm solution for \mathbf{x}_{feas} and do not update this further.

We strengthen the bounds further by using two different sets of bounds in the linearization equations

$$\begin{aligned} L_i^{(0)} \mathbf{x}_i &\leq \mathbf{z}_i \leq U_i^{(0)} \mathbf{x}_i, \\ \mathbf{a}_i^T \mathbf{y} - U_i^{(1)} &\leq \mathbf{z}_i \leq \mathbf{a}_i^T \mathbf{y} - L_i^{(1)} \mathbf{x}_i, \end{aligned}$$

namely, an upper and lower bound when $\mathbf{x}_i = 0$ and an upper and lower bound when $\mathbf{x}_i = 1$. Specifically, when $x_i = 1$ we may strengthen our bounds using the Sherman-Morrison formula

$$\langle \mathbf{a}_i, \mathbf{a}_i \rangle_{\mathbf{e}_i} = \langle \mathbf{a}_i, \mathbf{a}_i \rangle_0 - \frac{\langle \mathbf{a}_i, \mathbf{a}_i \rangle_0^2}{1 + \langle \mathbf{a}_i, \mathbf{a}_i \rangle_0} = \frac{\langle \mathbf{a}_i, \mathbf{a}_i \rangle_0}{1 + \langle \mathbf{a}_i, \mathbf{a}_i \rangle_0}.$$

where \mathbf{e}_i is an index vector. Therefore our bounds $L_i^{(1)}$ and $B_i^{(1)}$ can be strengthened by a factor of $1 + \langle \mathbf{a}_i, \mathbf{a}_i \rangle_0$ to

$$\mathbf{z}_i^2 \leq \langle \mathbf{a}_i, \mathbf{a}_i \rangle_{\mathbf{e}_i} \langle \mathbf{g}, \mathbf{g} \rangle_{\text{feas}} = \frac{\langle \mathbf{a}_i, \mathbf{a}_i \rangle_0 \langle \mathbf{g}, \mathbf{g} \rangle_{\text{feas}}}{1 + \langle \mathbf{a}_i, \mathbf{a}_i \rangle_0}.$$

Empirically these bounds appear to be relatively weak, and strengthening them further is a promising area of future study.

4.4.3 Benchmarking

Logistics Scenarios

For the purposes of benchmarking, we formulated a series of three increasingly complex logistics models, which we call Knapsack, Helipad, and Tanana. Our Knapsack model assigns to each possible observation a fixed cost c_i , and imposes the constraint $\mathbf{c}^T \mathbf{x} \leq B$, for budget B . In our demo, these costs are generated based on two-dimensional location (s, r) from the linear function $c(r, s) = r + s + 1$. The budget is set to be 100. Both the Knapsack and Helipad scenarios were simulated on the unit square.

Our Helipad model maintains the knapsack constraint, but additionally assumes that observations must be accessed from a series of 9 equally spaced helipads, each with a limited range. Each location is therefore only accessible from a small number of nearby helipads, and to access a point from a given helipad, the helipad must be maintained for the season, incurring a fixed cost. It is therefore possible to save costs by using fewer helipads overall, thereby forgoing the locations that are accessible from those locations, but allowing more locations to be sampled from the helipads which are maintained.

The helipad constraint can be modeled by an adjacency matrix \mathbf{D} describing whether point i is reachable from helipad j . Then the following set of constraints can be used, using the vector \mathbf{h} for the helipad binary decision variables, and the vector \mathbf{f} for the fixed costs of maintenance:

$$\begin{aligned} \mathbf{c}^T \mathbf{x} + \mathbf{f}^T \mathbf{h} &\leq B \\ \mathbf{x}_i &\leq \sum D_{ij} \mathbf{h}_j. \end{aligned}$$

In our demonstration, the cost of each helipad is set to be a constant $f = 10$ for all helipads. The cost of sampling was again set to be $c(r, s) = r + s + 1$. The range of each helipad was set to be $\frac{1}{3\sqrt{2}}$ to allow only a small amount of overlap between ranges on the unit square.

Finally, our Tanana model was developed in collaboration with the logistics team for the US Forest Service's (USFS) Forest Inventory Analysis (FIA) project in Alaska; in which a major goal is the estimation of total carbon sequestration. This model matched actual on-the-ground logistics as closely as possible, using the actual geometry of the region. Our Tanana logistics model included: the actual set of available helicopter bases, fuel costs, costs of food and lodging, variable efficiency in number of plots that can be sampled per day from each base, a fixed maximum number of days in the sampling season of each of multiple years, and a fixed cost of lost sampling days incurred by traveling to a new base mid-season. The complete set of logistics modeled are described in Table 4.1.

The Tanana logistics motivate the following constraints:

$O_i + F_i \leq B_i$	operating cost + flight cost \leq budget
$H_i + S_i \leq D_i$	hub relocation + sampling \leq season length
$O_i = \sum_j c_j s_{i,j}$	operating cost = \sum cost of site \times days at site
$F_i = \sum_j f_j d_{i,j}$	flight costs = \sum fuel cost at site \times distance traveled
$H_i = 2 * \left(\sum_j h_{i,j} - 1 \right)$	hub relocation days = 2 per hub, except the first
$S_i = \sum_j s_{i,j}$	sampling days = \sum sampling days per site
$p_{i,j} = \sum_k x_{i,j,k}$	plots sampled at site = \sum plot indicator variables
$s_{i,j} = p_{i,j} / e_j$	sampling days = plots sampled / plots per day
$d_{i,j} = \sum_k 4d'_{j,k} x_{i,j,k}$	distance traveled = \sum 4 \times distance to sites sampled
$h_{i,j} \leq M_{k,j} x_{i,j,k} \forall i, j$	a hub must be used if a site is sampled from it
$x_{i,j,k} \leq M_{k,j} h_{i,j} \forall i, j, k$	if a site is sampled, it must be from some hub

where B_i and D_i describe the budget and days in season i , which are generally constant but are reduced by the time and money spent conducting the pilot study in the first year; O_i

represents the operating cost in year i ; F_i represents the annual flight costs; H_i represents the number of days spent transferring hubs in year i ; S_i describes the number of days spent sampling in year i ; c_j describes the daily cost of operations at hub j ; $s_{i,j}$ denotes the number of days spent sampling at hub j in year i ; f_j denotes the fuel costs at hub j ; $d_{i,j}$ denotes the distance flown from hub j in year i ; $h_{i,j}$ denotes whether hub j was used in year i ; $x_{i,j,k}$ denotes whether plot k was sampled from hub j in year i ; $p_{i,j}$ denotes the number of plots sampled from hub j in year i ; e_j denotes the efficiency of plot j measured in plots/day; and $d'_{j,k}$ denotes the distance from hub j to plot k . To facilitate subsampling of the actual set of actual sampled locations, we used only two sampling seasons, each of reduced length and budget.

Quantifying Efficacy

To estimate the efficacy of our method, we simulated the full process by which such a method would be implemented in practice. First we generated the underlying spatial field \mathbf{u} as a Matérn Gaussian process with parameters as outlined in Table 4.2. We chose the α parameter to ensure our simulations reflected the reality that α is frequently unidentifiable (i.e. because our models assumed $\alpha = 2$ instead of 4; Lindgren et al. 2011). We set σ_ϵ^2 to one for all scenarios. We converted between κ and ρ using the empirical relationship $\kappa = \sqrt{8\nu}/\rho$ (Lindgren et al., 2011) and chose both so that the range was smaller than but on the same order of magnitude as the spatial extent of the domain. In this regime, spatially balancing can be expected to perform reasonably well, allowing us to treat our results as a sort of baseline level of performance for our method relative to SBRS. If range was either very large or very small, it would be advantageous to cluster points either in the cheapest locations, or nearby to the edges, which would tend to widen the gap between our method and the benchmarks. We chose σ^2 to reflect both a high-uncertainty regime for the Knapsack and Helipad scenarios and a low-uncertainty regime for the Tanana scenario, in which great expense is taken to reach a location and therefore high quality data are taken at each location.

To model the requisite pilot study to estimate the covariance function hyperparameters,

Annual budget	Each season must be individually within budget. We used 1/5th the true budget to accommodate sub-sampling the actual locations.
Length of sampling season	Each season must be completed in a fixed number of days. We used 1/5th the true season to accommodate sub-sampling the actual locations.
Helicopter hubs	Teams of samplers lodge near hubs, and are flown by helicopter to nearby sites within range.
Multi-team routing	Multiple teams share a single helicopter, which must drop off a single team before returning to transport the next. The helicopter then must pick up each team at the end of the day.
Variable efficiencies	Due to a wide variety of factors, hubs have a variable average number of sites visited per day
Variable cost of operation	Due to a wide variety of factors, hubs have a variable daily cost of operation
Cost of fuel	Each mile traveled incurs a 4× fuel cost to drop off the team, return, pick them up, and return.
Hub switching	Every time the teams move to a new hub they lose 2 days of productivity

Table 4.1: Logistic factors modeled in the Tanana scenario.

Model	α	σ^2	κ	ρ	σ_ϵ^2	Time Limit
Knapsack	4	20	2	≈ 0.24	1	5min
Helipad	4	20	2	≈ 0.24	1	20min
Tanana	4	1000	$\approx 1.63 \times 10^{-5}$	300 km	1	4hr

Table 4.2: The parameter values used for the three different scenarios of our benchmarking study. Greek letters refer to parameters used to simulate the data-generating process. We translated between κ and ρ using the empirical equation from (Lindgren et al., 2011). Time Limit refers to the time limit at which the Gurobi solver terminates the B&B procedure and returns the best known solution.

we used sampling routines designed to balance estimation of these hyperparameters with budget and feasibility considerations for the pilot study, and with optimality of the eventual sample for estimating the areal mean. For the Helipad and Tanana scenarios, we chose two bases with which to collect the pilot study. In the Tanana scenario, due to wide variation in the area covered by each base, these bases were fixed to be the two with the largest covered area. In the Knapsack and Helipad scenarios, pilot samples were then chosen according to an incremental routine designed to produce some clusters and some distant points, by weighting points according to their distance from existing points, then sampling randomly. In the Tanana scenario, due to the complexity of the logistics, the need to overestimate the range was more pressing, and we therefore conducted a SBRS procedure to collect the pilot study, because precise estimation of the hyperparameters was less of a priority.

In all cases, the selected pilot sample was then used to produce simulated data, by adding i.i.d Gaussian noise to the simulated spatial field. We used these data to fit the hyperparameters using maximum *a posteriori* estimation, assuming for the sake of testing robustness an *incorrect* value for the hyperparameter α , which is difficult to identify in practice, and is frequently fixed *a priori* (Lindgren et al., 2011). It was necessary to obtain a point estimate in order to have a single matrix \mathbf{Q} to use in the final optimization model.

An important hyperparameter for our model is the range (determined by κ in the Matérn model), because if the range is small the model attempts to save costs by moving observations closer together, which it can do while sacrificing only very little information, because when the range is small, observations are approximately uncorrelated even at low distances. When the range is underestimated, however, this can be detrimental to the resulting inference, thus to produce a conservative design we imposed informative priors overestimating the range.

After simulating a pilot study, we generate a warmstart sample using a greedy algorithm. A warmstart is an initial guess at a near-optimal solution, often one which is known to be feasible, and provides the B&B algorithm a heuristic from which to begin solving. In our case we use a greedy algorithm, in which we evaluate the decrease in our objective function for each candidate location then sort the candidates by the size of this decrease, and evaluate feasibility until we find the most informative feasible location. Feasibility is determined using an auxiliary MILP model, generated by removing the objective function and the variables \mathbf{y} and \mathbf{z} used to calculate it. When no locations are feasible the algorithm terminates. This warmstart was provided to Gurobi as a heuristic solution, then Gurobi was allowed to run for a fixed period of time, which was longer for more complex scenarios. At the end of this period of time, the algorithm outputs the best sample found. Due to the size of our simulation, the B&B algorithm was necessarily constrained by a relatively short time limit. If a more optimal solution was able to be found after a longer period of time, we expect the gap between our method and the benchmark comparisons would improve, as would the reported optimality gap.

For comparison, we also conducted simple random sampling, stratified random sampling (except for Tanana, which featured a less regular domain), and spatially-balanced random sampling (SBRS). In the Knapsack and Helipad models, these methods were allowed to choose points continuously from the domain, and SBRS was accomplished using balanced acceptance sampling (BAS; Robertson et al. 2013). In these scenarios, candidate locations for optimization were generated as a dense square grid. In the Tanana scenario, the set of possible observations for the design-based methods were chosen from the actual, hexagonally

systematic sample used by the US Forest Service (USFS) to sample the region, and we used Generalized Random Tessellation Stratified (GRTS) sampling for our SBRS method to accommodate the discrete set of candidate points (Stevens Jr and Olsen, 2004). For each design-based method, we used a series of sample sizes which ranged from near-certain infeasibility to near-certain feasibility. We also evaluated the performance of our greedy algorithm warmstart samples, as well as our B&B samples, optimized using the Gurobi solver (Gurobi Optimization, LLC, 2023).

For each method we then conducted the appropriate respective calculations to infer the areal mean of the simulated function. In the case of the greedy and optimized designs we used kriging to estimate the areal mean; we did not update the fitted autocorrelation model, but when not analyzing large numbers of simulations, a less programmatic analysis would be appropriate. We did not model covariates in any of our simulations to have a cleaner comparison to the base versions of our design based benchmarks, but show how one could do so in the Appendix 1. In the case of the design based methods we used the mean of the measurements, because we used equal weights across the whole spatial domain, or in the case of stratified random sampling the mean of the strata estimates (which were themselves of equal area).

Finally, we calculated the squared error for each simulation, as well as the feasibility. We bootstrap resampled these two benchmarks 1000 times to produce confidence intervals of mean squared error (MSE) and probability of feasibility. Our total number of simulations for each benchmark and scenario can be found in Table 4.3. For purposes of computation, we used fewer simulations for the more complex logistics models, because these took substantially longer to run. We allowed the Knapsack optimization routine to run for 5 minutes per simulation; the Helipad routine to run for 20 minutes per simulation; and the Tanana routine to run for 4 hours per simulation.

	Knapsack	Helipad	Tanana
Design	3000	1500	1000
Optimized	300	100	50

Table 4.3: The number of benchmarking simulations run to estimate MSE and probability of feasibility for each method, by scenario.

4.5 Results

Example designs produced by each of our benchmark methods, our pilot methods, our greedy warmstart, and our B&B solution can be found in Figures 4.1-4.3. Visually, our methods appear to space points out comparably to stratified sampling or SBRS. However, by accounting for logistics, we hypothesized that our method would be able to select a higher sample size of points while maintaining feasibility.

We found that, for the Knapsack scenario, our model performed comparably to stratified random sampling, but underperformed relative to SBRS, which we hypothesize to be due to the non-optimality of the pilot study. Knapsack logistics are well-studied and generally simpler logistics than our method is intended for; other optimal design methods such as submodular optimization may be more appropriate in these scenarios because the simpler structure allows for better optimization properties, such as optimality approximation guarantees (Lee et al., 2009) on greedy samples. Indeed for knapsack logistics, it is even possible to produce a set of design-based optimal weights, provided that spatial covariance is not modeled (see Appendix 4).

For more complex logistics, including both the Helipad and Tanana scenarios, both our greedy and optimized designs outperformed the design-based benchmarks. For the Tanana scenario, the logistics were sufficiently complex that even at small sample sizes it was difficult for the design-based methods to achieve near-certain feasibility. In contrast, our model produced a high-quality MSE comparable to relatively large sample-size SBRS, while guar-

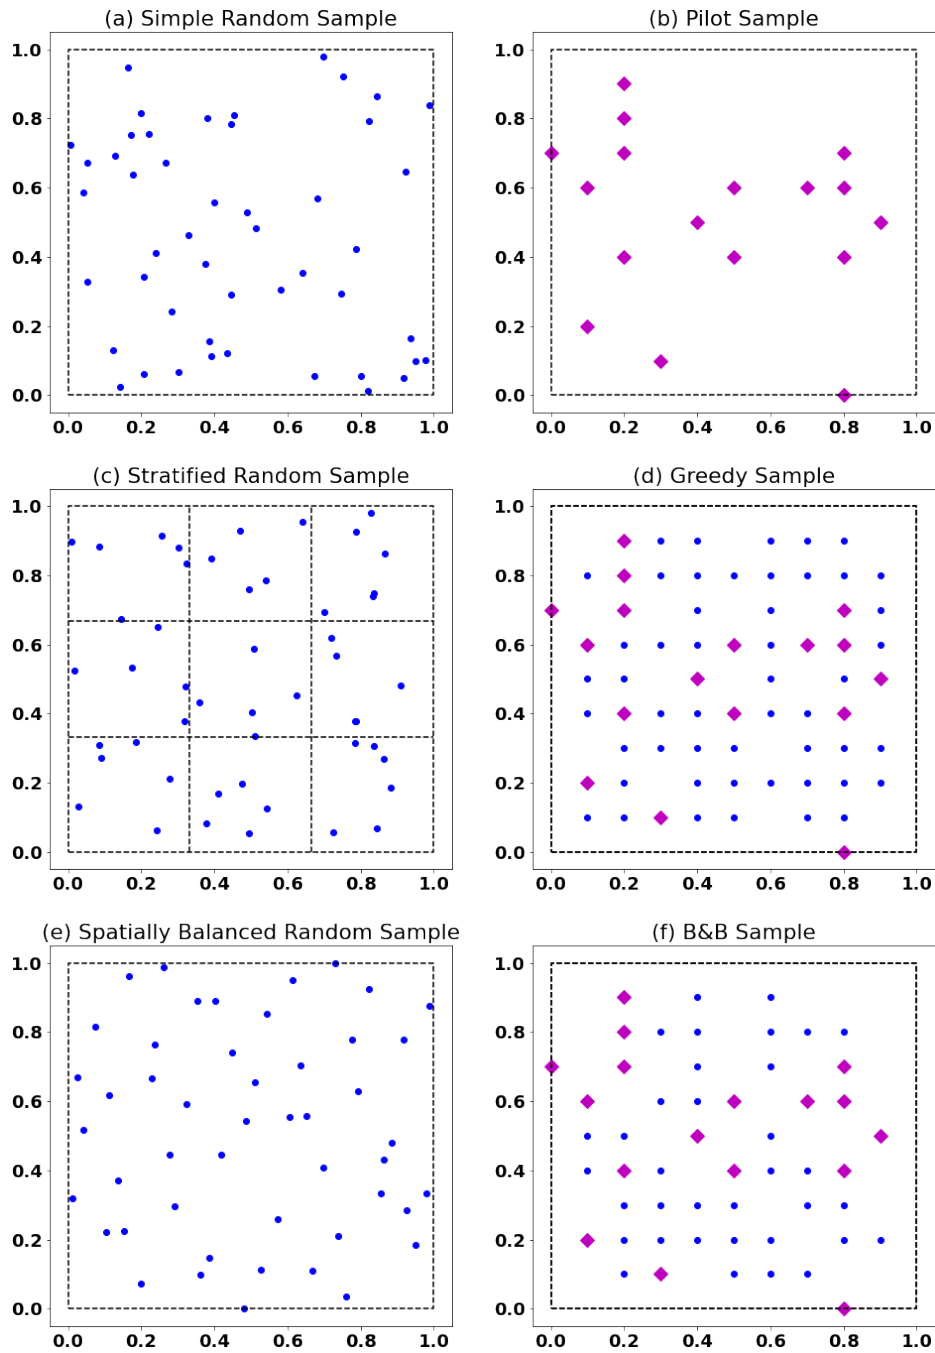


Figure 4.1: Example sample for the Knapsack scenario. Sample locations shown in circles (blue). Pilot samples shown in diamonds (magenta). Domain and strata outlined in dashed line (black).

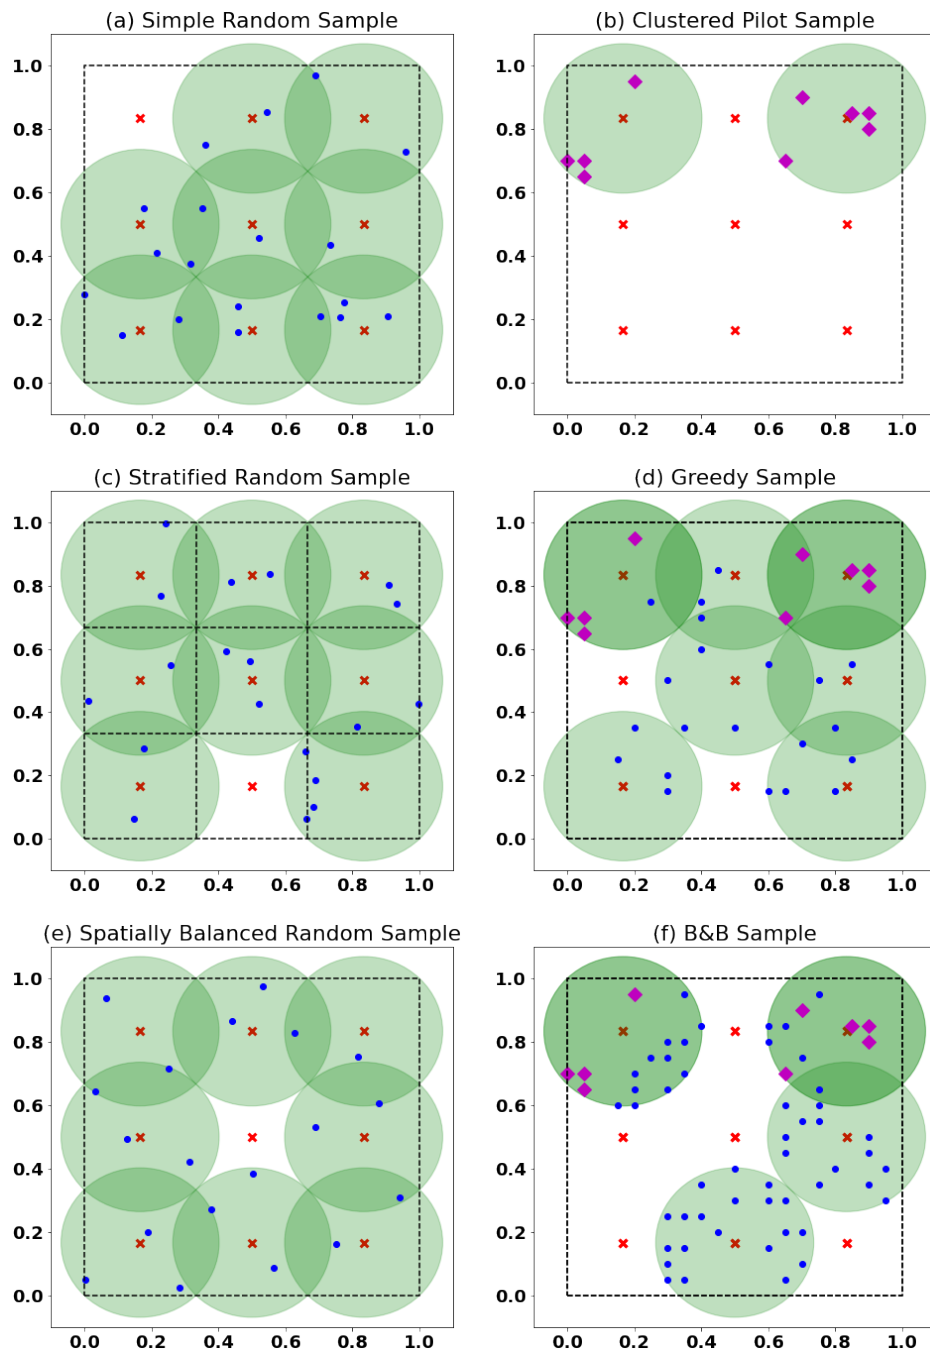


Figure 4.2: Example sample for the Helipad scenario. Sample locations shown in circles (blue). Pilot samples shown in diamonds (magenta). Helipad locations shown in Xs (red). Helipad ranges shown in large transparent circles (green). Domain and strata outlined in dashed line (black). Choice of helipads for (a)-(e) determined by cost-minimization MILP.

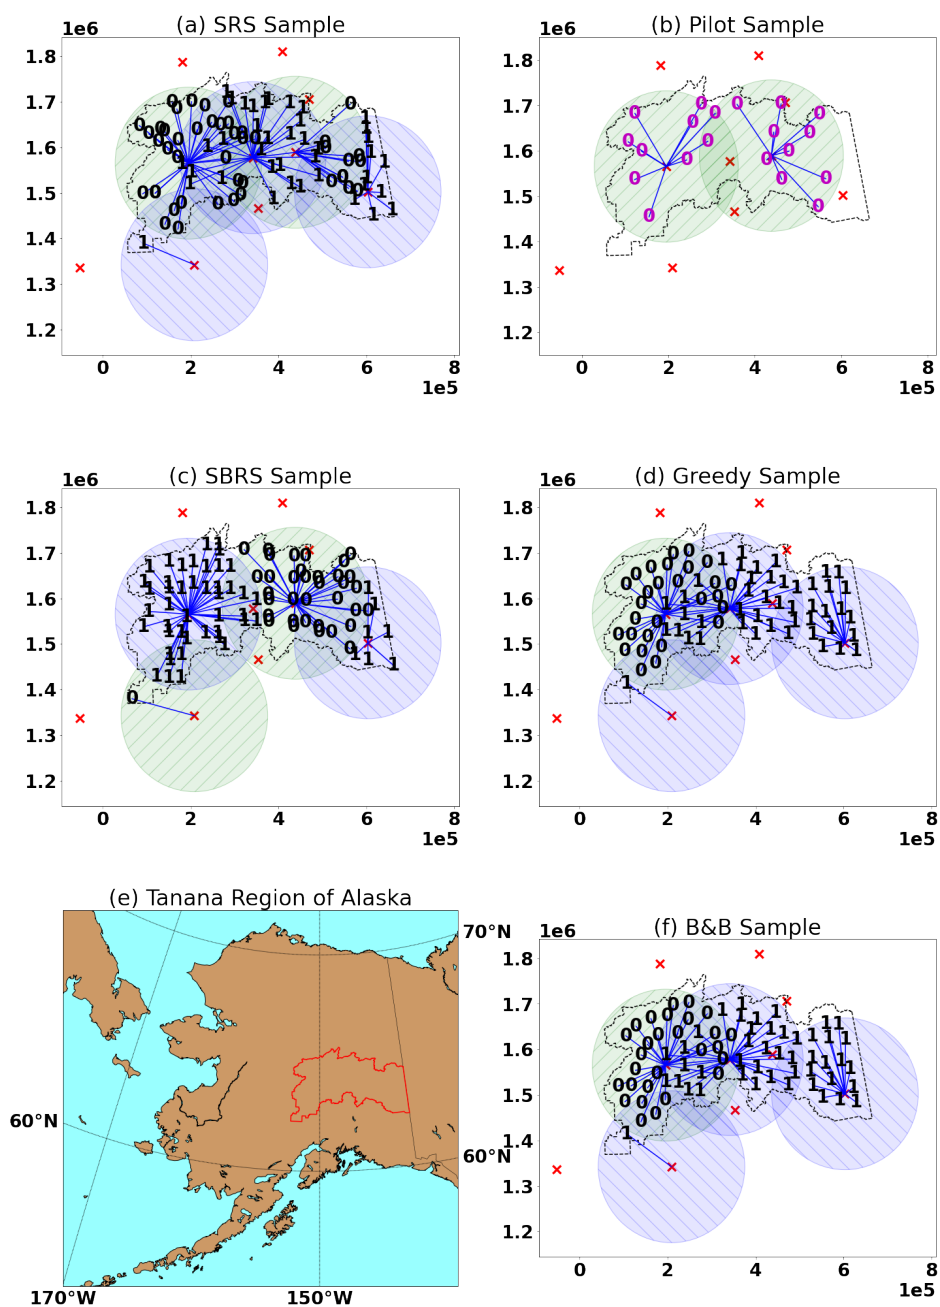


Figure 4.3: Example sample for the Tanana scenario. Sample locations marked with 0 or 1 depending on which sampling season the sample is to be collected. Base ranges shown in large transparent circles (green with hatching toward upper-right for season 0; blue with hatching toward upper-left for season 1). Base used to reach each sample shown in radiating lines (blue). Domain (the Tanana region of Alaska) outlined in dashed line (black), and also outlined (red) in a map of Alaska in panel (e). Axes in panels (a)-(d) and (f) are in NAD83 / Alaska Albers coordinates; panel (e) is in latitude and longitude.

	Knapsack	Helipad	Tanana
SRS	44-54%	56-67%	68-79%
StratRS	19-32%	56-67%	NA
SBRS	-4-25%	45-59%	45-64%

Table 4.4: Proportional improvements in MSE demonstrated by our B&B solution as a percentage of our three benchmarks' MSEs. For SBRS in the Knapsack scenario, shows the proportional increase in MSE. StratRS was not used in the Tanana scenario.

anteeing feasibility. Total simulation results can be found in Figure 4.4, where we show how the MSE and probability of feasibility change with sample size for each of our competing methods, as well as bootstrapped confidence intervals showing the uncertainty in both metrics. The greedy and branch and bound solution metrics are displayed as points, because we cannot vary the sample size.

Proportional improvements for the branch and bound solution relative to benchmarks are shown in Table 4.4 as percentages of the benchmark MSE, calculated from bootstrapped confidence intervals. Note additionally that only in the Knapsack scenario did all benchmarks achieve 100% simulated feasibility. Although it is more difficult to calculate precise estimates, it is illustrative also to consider the loss of feasibility that would be required to match the MSE of our model solutions. Comparing horizontally the extremes of our bootstrap intervals, for the Helipad scenario, requisite losses appear to range from $\approx 30-95\%$ and for the Tanana scenario $\approx 5-100\%$ probability of feasibility would need to be sacrificed.

We also observed relatively little difference between our B&B solution and our greedy warmstart solution. The MILP optimization routine outperformed the greedy heuristic routine on average, but by a relatively small margin. Meanwhile, the B&B routine required a substantially larger computational burden.

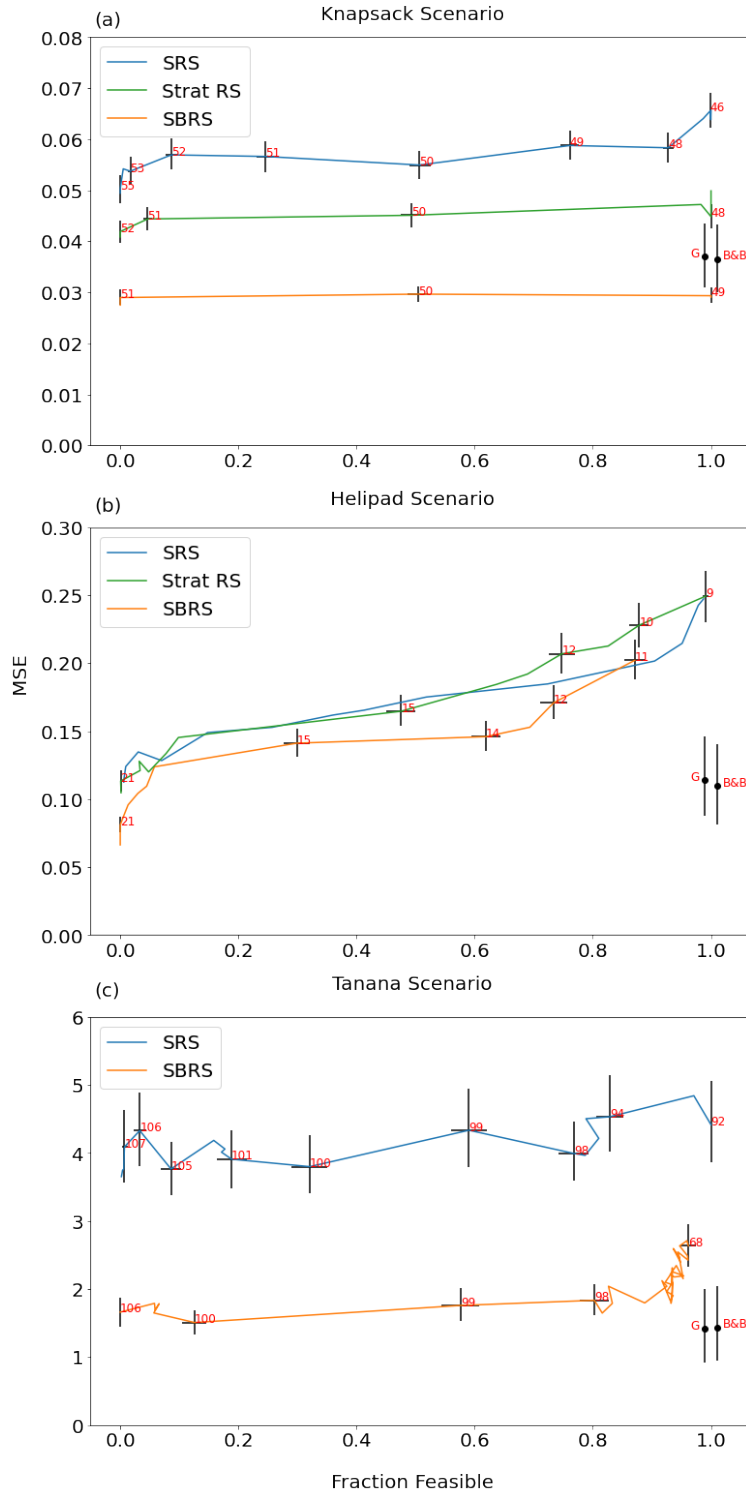


Figure 4.4: Results of our benchmarking simulations. Left-side plots contain the same lines as the right-side, but without the addition of bootstrapped error bars which obscure the underlying lines. Red numbers display the sample size for each point, and the characters “B&B” and “G” stand for the branch and bound/greedy methods respectively.

4.6 Conclusions

For even moderately complex logistics scenarios, our two optimization routines outperformed all of our design-based benchmarks, despite making model choices that would disfavour our routines, including: assuming an incorrect value for the α parameter, frequently overestimating the range, and not updating the hyperparameters after sampling. These results are in line with a recent study comparing design-based and model-based methods for spatial sampling and inference more broadly, which found that even when model assumptions are violated, model-based approaches slightly outperformed SBRS and greatly outperform SRS (Dumelle et al., 2022). The use of a logistics-aware sampling design method has the potential to ease implementation of large-scale complex designs, by shifting complexity from the often time-sensitive implementation stage to the planning stage. We expect logistics-aware sampling design to be advantageous in remote regions such as Alaska, or in surveys involving marine or aerial support, where transport and routing may account for a substantial fraction of overall cost. Using a model-based optimal design approach can ensure that limited time and money are spent efficiently to obtain the highest amount of information possible.

In the three scenarios we evaluated, the greedy heuristic performed comparably to the MILP optimization routine, while requiring substantially lower computation burden. In this particular set of scenarios therefore, the greedy heuristic is the more practical optimization routine. A number of factors in our benchmarks contributed to the high performance of the greedy heuristic. To provide a fair comparison with the randomized methods, we chose a moderate range for the generating process, and a reasonably homogeneous pricing structure for the Knapsack and Helipad scenarios, which had the effect that spacing sampling points out evenly over the domain was an approximately optimal strategy. In scenarios with more highly variable or non-linear logistics, or with ranges that are small compared to the domain, both the greedy heuristic and the various design-based methods would be expected to underperform, because they would include points without considering that doing so might exclude a large range of other important points, or that a much cheaper location might supply a

nearly equal amount of information, and would therefore space points out farther than is optimal. To provide a conservative comparison, we chose scenarios and generating processes that tended to favor the greedy heuristic. In the future, these methods can be compared under a wider variety of logistical scenarios, such as transect sampling, vehicle routing, and spatio-temporal sampling. They can also be evaluated under further scenarios, such as more severely mis-specified covariance structures. By testing a more diverse set of scenarios, it should be possible to elucidate under what conditions the computational burden of B&B optimization is beneficial, and when the greedy heuristic can be expected to suffice.

Notably, the greedy heuristic does not yet have a proven optimality gap for this problem. The B&B algorithm also did not provide useful optimality gaps, because the LP relaxations largely provided negative variances in the bounding steps. We hypothesize that this is due to the looseness of our Cauchy-Schwartz derived bounds. If tighter bounds can be proven for this problem, we expect that the B&B algorithm will be able to provide more useful optimality gaps, and will also provide improved solutions more quickly because stronger bounds would improve the efficiency with which branches may be pruned.

We also hope to tackle more complex statistical models using our method in future applications. For example, we hypothesize that non-linear objective functions such as those that would be required for generalized linear mixed models can be approximately accommodated through Taylor series when moderately informative prior information is available or assumed, and that when such prior information is lacking for a small number of variables it may be possible to nevertheless approximately optimize sampling designs for these models by using mixture priors. Such an approach may be particularly advantageous in an adaptive design scenario, when data are available from previous years (e.g. Leach et al. 2022).

Overall, our method formulates a common optimization criterion for spatial surveys – the uncertainty in the areal mean – as a MILP so that it can be optimized under complex logistical constraints. We tested this method using a variety of scenarios and found that under simple logistical constraints it was outperformed by the best design-based methods, but for even moderately complex logistical constraints it outperformed our design-based

benchmarks. At present the greedy warmstart algorithm provides comparable results more quickly, but we expect that with the development of tighter bounds, and longer run times for a single model, that the B&B algorithm will provide better results as well as a dynamically improvable optimality gap.

Acknowledgements

We gratefully acknowledge the support of Hans-Erik Andersen and Sarah Ellison for contributing their knowledge of the logistics used in the USFS FIA survey of Alaska, and providing the sampling locations used.

BIBLIOGRAPHY

- Avellar, G. S., Pereira, G. A., Pimenta, L. C., and Iscold, P. (2015). Multi-UAV routing for area coverage and remote sensing with minimum time. *Sensors*, 15(11):27783–27803.
- Bechtold, W. A. and Patterson, P. L. (2005). *The enhanced forest inventory and analysis program—national sampling design and estimation procedures*. Number 80. USDA Forest Service, Southern Research Station.
- Bellhouse, D. (2005). Systematic sampling methods. *Encyclopedia of Biostatistics*, 8.
- Bertsimas, D. and Tsitsiklis, J. N. (1997). *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA.
- Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208.
- Clausen, J. (1999). Branch and bound algorithms—principles and examples. *Department of Computer Science, University of Copenhagen*, pages 1–30.
- CPLEX (2020). IBM ILOG CPLEX Optimization Studio User’s Manual, version 12 release 6.
- Dantzig, G. B. (1957). Discrete-variable extremum problems. *Operations Research*, 5(2):266–288.
- Diaz, P., Doostan, A., and Hampton, J. (2018). Sparse polynomial chaos expansions via compressed sensing and D-optimal design. *Computer Methods in Applied Mechanics and Engineering*, 336:640–666.

- Dumelle, M., Higham, M., Ver Hoef, J. M., Olsen, A. R., and Madsen, L. (2022). A comparison of design-based and model-based approaches for finite population spatial sampling and inference. *Methods in Ecology and Evolution*, 13(9):2018–2029.
- Dupont, G., Royle, J. A., Nawaz, M. A., and Sutherland, C. (2021). Optimal sampling design for spatial capture–recapture. *Ecology*, 102(3):e03262.
- Elfving, G. (1952). Optimum allocation in linear regression theory. *The Annals of Mathematical Statistics*, pages 255–262.
- Fischer, S. M., Beck, M., Herborg, L.-M., and Lewis, M. A. (2021). Managing aquatic invasions: optimal locations and operating times for watercraft inspection stations. *Journal of Environmental Management*, 283:111923.
- Gurobi Optimization, LLC (2023). Gurobi Optimizer Reference Manual.
- Hughes, R. M. and Peck, D. V. (2008). Acquiring data for large aquatic resource surveys: the art of compromise among science, logistics, and reality. *Journal of the North American Benthological Society*, 27(4):837–859.
- Kermorvant, C., Caill-Milly, N., Bru, N., and d’Amico, F. (2019). Optimizing cost-efficiency of long term monitoring programs by using spatially balanced sampling designs: The case of manila clams in Arcachon Bay. *Ecological Informatics*, 49:32–39.
- Krause, A., Singh, A., and Guestrin, C. (2008). Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(Feb):235–284.
- Kwak, D., Jeong, J., Shin, Y., Lee, N., and Shin, D. (2022). Sensor placement optimization for fence-line monitoring of toxic gases considering spatiotemporal risk of the plant-urban interface. *Journal of the Taiwan Institute of Chemical Engineers*, 130:103858.

- Land, A. H. and Doig, A. G. (2010). An automatic method for solving discrete programming problems. In *50 Years of Integer Programming 1958-2008*, pages 105–132. Springer.
- Leach, C. B., Williams, P. J., Eisaguirre, J. M., Womble, J. N., Bower, M. R., and Hooten, M. B. (2022). Recursive Bayesian computation facilitates adaptive optimal design in ecological studies. *Ecology*, 103(2):e03573.
- Lee, J., Mirrokni, V. S., Nagarajan, V., and Sviridenko, M. (2009). Non-monotone submodular maximization under matroid and knapsack constraints. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, pages 323–332.
- Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498.
- Lohr, S. L. (2019). *Sampling: Design and Analysis: Design and Analysis*. Chapman and Hall/CRC.
- Pukelsheim, F. (2006). *Optimal Design of Experiments*. SIAM.
- Rao, J. and Bellhouse, D. (1990). History and development of the theoretical foundations of survey based estimation and analysis. *Survey Methodology*, 16(1):3–29.
- Robertson, B., Brown, J., McDonald, T., and Jaksons, P. (2013). BAS: Balanced acceptance sampling of natural resources. *Biometrics*, 69(3):776–784.
- Sagnol, G., Gaubert, S., and Bouhtou, M. (2010). Optimal monitoring in large networks by successive C-optimal designs. In *2010 22nd International Teletraffic Congress (ITC 22)*, pages 1–8. IEEE.
- Stevens Jr, D. L. and Olsen, A. R. (2004). Spatially balanced sampling of natural resources. *Journal of the American Statistical Association*, 99(465):262–278.

- Stoica, P. and Babu, P. (2010). Algebraic derivation of Elfving theorem on optimal experiment design and some connections with sparse estimation. *IEEE Signal Processing Letters*, 17(8):743–745.
- Tobler, W. R. (1970). A computer movie simulating urban growth in the Detroit region. *Economic Geography*, 46(sup1):234–240.
- Van Groenigen, J. W., Siderius, W., and Stein, A. (1999). Constrained optimisation of soil sampling for minimisation of the kriging variance. *Geoderma*, 87(3-4):239–259.
- Wang, J.-F., Stein, A., Gao, B.-B., and Ge, Y. (2012). A review of spatial sampling. *Spatial Statistics*, 2:1–14.
- White, D., Kimerling, J. A., and Overton, S. W. (1992). Cartographic and geometric components of a global sampling design for environmental monitoring. *Cartography and Geographic Information Systems*, 19(1):5–22.
- Williams, H. P. (2013). *Model Building in Mathematical Programming*. John Wiley & Sons.
- Wynn, H. P. (1970). The sequential generation of D -optimum experimental designs. *The Annals of Mathematical Statistics*, 41(5):1655 – 1664.
- Yilmaz, N. K., Evangelinos, C., Lermusiaux, P. F., and Patrikalakis, N. M. (2008). Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming. *IEEE Journal of Oceanic Engineering*, 33(4):522–537.

Chapter 5
DISCUSSION

In this dissertation, I outlined three new mechanistic models for collective movement, source reconstruction, and optimal sampling design. In each case, the results motivate further study, development, and application of these models.

In my first chapter, I built a series of regression models based upon a mechanistic representation of how a salmon navigates a fish ladder, and used a parametric bootstrap approach to test for the effect of conspecific density on passage rates, even in the presence of possible reverse causality effects. In doing so, I uncovered a positive density effect, in contrast to the prevailing wisdom that overcrowding would dominate. Furthermore this positive density effect was consistently associated with Chinook, which is thought to be the less social of the two species we tested (i.e. than sockeye).

Future study is therefore warranted. Follow-up work is needed to confirm the replicability of this result in other systems, as well as to distinguish between possible hypotheses for why Chinook may benefit from this effect while sockeye seemingly do not. This work also motivates further analyses for other species, particularly given that the expectation that the more gregarious species would benefit from density was not borne out. The methods used here, particularly the bootstrap method used to overcome reverse causality, are of potentially wider importance in the study of social movement effects, where changes in behavior induced by density have the possibility of themselves influencing density. I am aware of no other analysis of animal movement accounting for this feedback loop.

In my second chapter, I showed how to invert the FEM equation first applied in spatial analysis by Lindgren et al. (2011). This novel approach allows the source term of an SPDE to be invested with real physical meaning and interpretation, which is a previously unexplored aspect of this method.

By imbuing the source term of an SPDE with physical meaning, a variety of opportunities for application are opened. This method could be applied to the detection of non-point source pollution or the inference of biomass through non-invasive eDNA monitoring. It might also be applicable outside of ecology and the environment to the inference of earthquake source locations through use of the wave equation; of heat sources using the heat equation; or

of electromagnetic activity using Maxwell's equations. Many of these are existing fields of research into inverse problems (e.g. Lindgren et al. 2011; Ikehata 2006; El Badia and Ha-Duong 2000). However, my method provides an approach that allows the application of Gaussian process models as a regularizing term that is novel in that field. In addition to applying this model to new dynamics, there is also the possibility of further extension and development of this model to more complex hierarchical statistical models.

In my third chapter, I showed how to formulate a MILP for a Bayesian C -optimal sampling design. After testing this model in three different scenarios, I found that it performed comparably to existing stratified and spatially balanced random sampling approaches under the simplest logistics scenario (Knapsack), and outperformed all design-based benchmarks by margins of $\approx 20 - 80\%$ for more complex scenarios (Helipad and Tanana).

This analysis motivates further study of the performance of this model and possible extensions. Now that performance has been demonstrated in non-temporal case studies for kriging-type analyses, it should also be confirmed in the presence of spatio-temporal variation and covariates. Case studies can also be carried out to demonstrate performance on actual data. For example, follow-up work could extend our Tanana case study to incorporate not just the complete logistics but also the full post-stratified estimation procedure (Bechtold and Patterson, 2005) and the actual data obtained from the Tanana sampling project. Because large-scale spatial surveys may need to serve a wide variety of stakeholders and may therefore value the flexibility of a design-based approach, future efforts might also extend this model to allow a design-based approach. For example, one might consider using our model or a similar model to generate a large number of high-quality feasible samples, and then to select one of these samples at random. In such a method it would be possible to calculate first- and second-order inclusion probabilities for use in design-based analyses. Further study should also focus on extending the model to the optimization of more complex statistical models. One such possibility, motivated by the application of adaptive design in which a sample is built incrementally (Leach et al., 2022) is to consider a Taylor approximation to a design criterion which might be possible to evaluate within a MILP. In the case of adaptive

design, substantial prior information may be available with which to justify the use of a Taylor approximation, and it may thus be possible to select a sample that approximately maximizes the expected improvement in the design criterion. Finally, it is possible to combine my methods from Chapters 2 and 3, to obtain a sampling design optimized for inferring a source term. In fact, such an application would be a prime candidate for optimal sampling since locations could be expected to differ quite widely in their information content.

Each of these three models brings new possibilities to light through the use of mechanistic modeling. The technical novelty of this work is in the synthesis and creation of new methods, models, and mechanisms along all steps of the statistical life cycle from data collection to model building to actual data analysis. With the techniques in this dissertation, insights can be drawn about the sociality of Pacific salmon; inferences can be made about the sources of dynamically distorted data; and feasible designs can be created in the face of complex logistics and budgetary constraints. In these and many other cases, great advantages can be gained through the judicious incorporation of known mechanisms.

BIBLIOGRAPHY

- Bechtold, W. A. and Patterson, P. L. (2005). *The enhanced forest inventory and analysis program—national sampling design and estimation procedures*. Number 80. USDA Forest Service, Southern Research Station.
- El Badia, A. and Ha-Duong, T. (2000). An inverse source problem in potential analysis. *Inverse Problems*, 16(3):651.
- Ikehata, M. (2006). An inverse source problem for the heat equation and the enclosure method. *Inverse Problems*, 23(1):183.
- Leach, C. B., Williams, P. J., Eisaguirre, J. M., Womble, J. N., Bower, M. R., and Hooten, M. B. (2022). Recursive Bayesian computation facilitates adaptive optimal design in ecological studies. *Ecology*, 103(2):e03573.
- Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498.

Appendix A
APPENDICES TO CH. 2

A.1 Statistical Analysis

We fit two proportional hazards (PH) models and one logistic model. These models are termed “finding,” “ladder,” and “committing” respectively. Each model was fit for Chinook and sockeye separately, because we expected inter-species differences in most or all coefficients. The PH models were fit in R (R Core Team, 2019) using the `phreg` function from the `eha` package (Broström and Broström, 2019). All models included daily counts of conspecifics. Because counts are available for each ladder, we used ladder-specific counts for the “ladder” and “commit” processes; however because a fish might approach either ladder upon entering the tailrace we used dam-level counts for the “finding” process.

All models controlled for temperature (Caudill et al., 2013), dam identity, and a binary diel effect (Keefer et al., 2013). For the finding and committing processes we included a spill effect, but because fish cannot detect spill from within the fishway, no spill effect was included in the ladder models.

For all models we included a binary covariate (`middle`) indicating whether a fish’s tailrace entry was in the middle (2nd and 3rd) quartiles. This variable controlled for the effect of run timing on fish condition and passage speed, because timing is highly correlated with density (very high densities near the modal date) and might otherwise confound our analysis.

For Chinook we included a binary variable (`spring`) indicating whether a fish was labeled spring- or summer-run. For sockeye we included a three-level factor variable indicating whether a fish was sourced to the Wenatchee, the Okanogan, or to any other site. The Wenatchee and Okanogan tributaries each accounted for about a third of the sockeye in our sample.

Two of our predictors — `temperature` and `day` — had strong prior evidence from the literature (Caudill et al., 2013; Keefer et al., 2013). To reduce the number of variables being selected and the risks associated with model dredging, we included these variables *a priori*. The remainder of the predictors we selected using AICc. For consistency, however, we wished to fit structurally identical models to each set of species/process submodels. We examined

the AICc for each set of models and subjectively chose what seemed to be the best-fitting overall model. Due to the extremely long tails of the count predictor, we then considered adding either an un-transformed or a log-transformed count variable to the best-fitting null model. In most cases there was no discernible difference in AICc between the two, but in several cases the un-transformed variable was selected. For consistency, despite the long tails, we used the un-transformed version to conduct the remainder of our analysis.

The classical form of the PH model uses a partial likelihood function and does not specify the baseline hazard function, but to simulate from our bootstrapped distributions we assumed a constant baseline hazard $\lambda_0(t) = \lambda_0$. This was visually consistent with the empirical hazard functions. However, because there is a small but non-trivial minimum time to complete each task, the empirical hazard functions are zero for a short period of time. To account for this discrepancy, we truncated each fish’s timeline by a constant amount (per model) such that the minimum passage time was less than 1 minute, effectively assuming a piecewise constant baseline hazard function with the first section fixed at zero hazard. In several models there were outlier individuals detected as passing through the fishway in as little as 2 minutes. Because such passage times are implausible (fast passage times are ≈ 1 hour), we attribute this to data error and remove all records which show passage in <10 minutes. This approach conservatively includes several extraordinarily fast passage times of ≈ 30 minutes.

For each model, we dropped any fish with censored measurements. Censoring is a standard concern for time-to-event models, and occurs when an event is not observed to occur because an individual leaves observation. This could occur when a fish passed but was not detected by a telemetry antenna, or because a fish died in a tailrace. Incorrect assumptions regarding censoring are known to bias inference in PH models Kleinbaum and Klein (2010). We did not account for censoring because, for each system, we used essentially only two measurements for each fish: “entry” and “exit.” A censoring analysis would require follow-up measurements between entry and exit to verify that an individual is still being detected but has not yet exited. Although the aerial tailrace antenna misses a large number of measurements, the underwater fishway antennas miss far fewer. Only a very small por-

tion ($\approx 1\%$) of fish are observed to enter but not exit the fishway. Because a very similar portion of fish are observed to exit but not enter, we attribute this censoring almost entirely to randomly missed telemetry detections rather than mortality. Because individuals were therefore dropped from the dataset independently of hazard we lose little information and incur very little bias in this manner.

Finally, we conducted a parametric bootstrap test on the count coefficient to detect density effects, using our best-fitting null models. Because these models do not contain reverse causality, they provide valid, unbiased fits under the null hypothesis. To each simulation, we fit a model with an effect of count. Assuming our null model is correctly specified, comparing our fitted count coefficients to these null distributions provide valid p -values to test for density effects. We do not correct these p -values for multiple testing because by the nature of our split models a single false-positive result is not statistically convincing, as we note when discussing our two marginally significant results. Instead we demand consistency between matching models, and indeed we found that several models were consistently highly significant (and would remain so even with a Bonferroni correction).

To ensure that our null models were well-specified, we conducted standard model diagnostics. For our PH models we tested the proportional hazards assumption using the `cox.zph` function in the `survival` package (Therneau and Lumley, 2014), applied to the partial-likelihood model fits. Because the `cox.zph` function tests the proportional hazards assumption for each covariate as well as globally for each model we made 36 tests. Using a Bonferroni correction, this test identified strong differences (corrected $p < 10^{-5}$) in baseline hazard functions between dams, likely caused by differences in the layout of the monitoring arrays. We thus split our models by dam. Retesting these sub-models, we made an additional 59 tests. These tests caused us to split our Chinook ladder models by run ($p < 10^{-5}$) and to split our `sockeye.ladder.TD` models by site ($p < 0.005$; East (TDE) vs. North (TDN)). Because all of these were individual-level variables, we could not assign interactions with time, because this could (and did) result in predictions of effectively infinite survival time. We also added interaction effects between temperature and time ($p < 0.001$) in all of

the Chinook ladder models. Finally we made an additional 41 tests after conducting this second level of model-splitting, with no significant results. After splitting, the Chinook ladder models each contained between 330 and 440 fish. All other models contained over 500 fish, with the exception of `sockeye.ladder.TDN` which had only 59. Our bootstrap simulations frequently produced ill-conditioned fits for this model, and we subsequently removed it from our analysis. A diagram of our analysis can be found in Figure A.1. Although we hypothesized positive density effects, we conducted two-sided tests in all cases to allow for the possibility of negative effects (e.g. overcrowding in the fishway).

For the `commit` models, we fit one logistic regression model per species, without splitting by dam. We plotted binned residuals from the null model, which motivated the addition of a threshold variable ($\text{count} \leq 150$) to the Chinook model, due to substantially higher-than-expected commit rates for these low-density fish. Of these 89 individuals, 88 of them were detected at The Dalles's east ladder. Elevated commit probability was not observed at TDE at higher densities.

Because of the reverse causality in our system, the null modeling approach we took was the most reliable way we knew to detect a density effect. However, this approach provides only p -values, and not confidence intervals for the true magnitude. To quantify magnitude, we conducted another parametric bootstrap to produce 95% confidence intervals for the count coefficient. To produce these intervals we sampled from the full model (i.e. the model with a count coefficient) and fit to these simulations new count coefficients. We calculate the confidence intervals using the hybrid method, recommended in Burr (1994) for use when estimating regression coefficients in Cox models. In all cases, the bootstrapped distribution was symmetric and centered on the fitted coefficient. This secondary bootstrap also roughly reproduced the p -values of our primary bootstrap. These two metrics lend credibility to this bootstrap, despite possible biases that might enter due to reverse causality. Thus, we also report approximate 95% confidence intervals for the magnitude of the density coefficient, calculated from these bootstrapped distributions. All bootstraps included >2000 simulations.

To speed up calculation we did bootstrap simulations in parallel using 24 cores of a 48

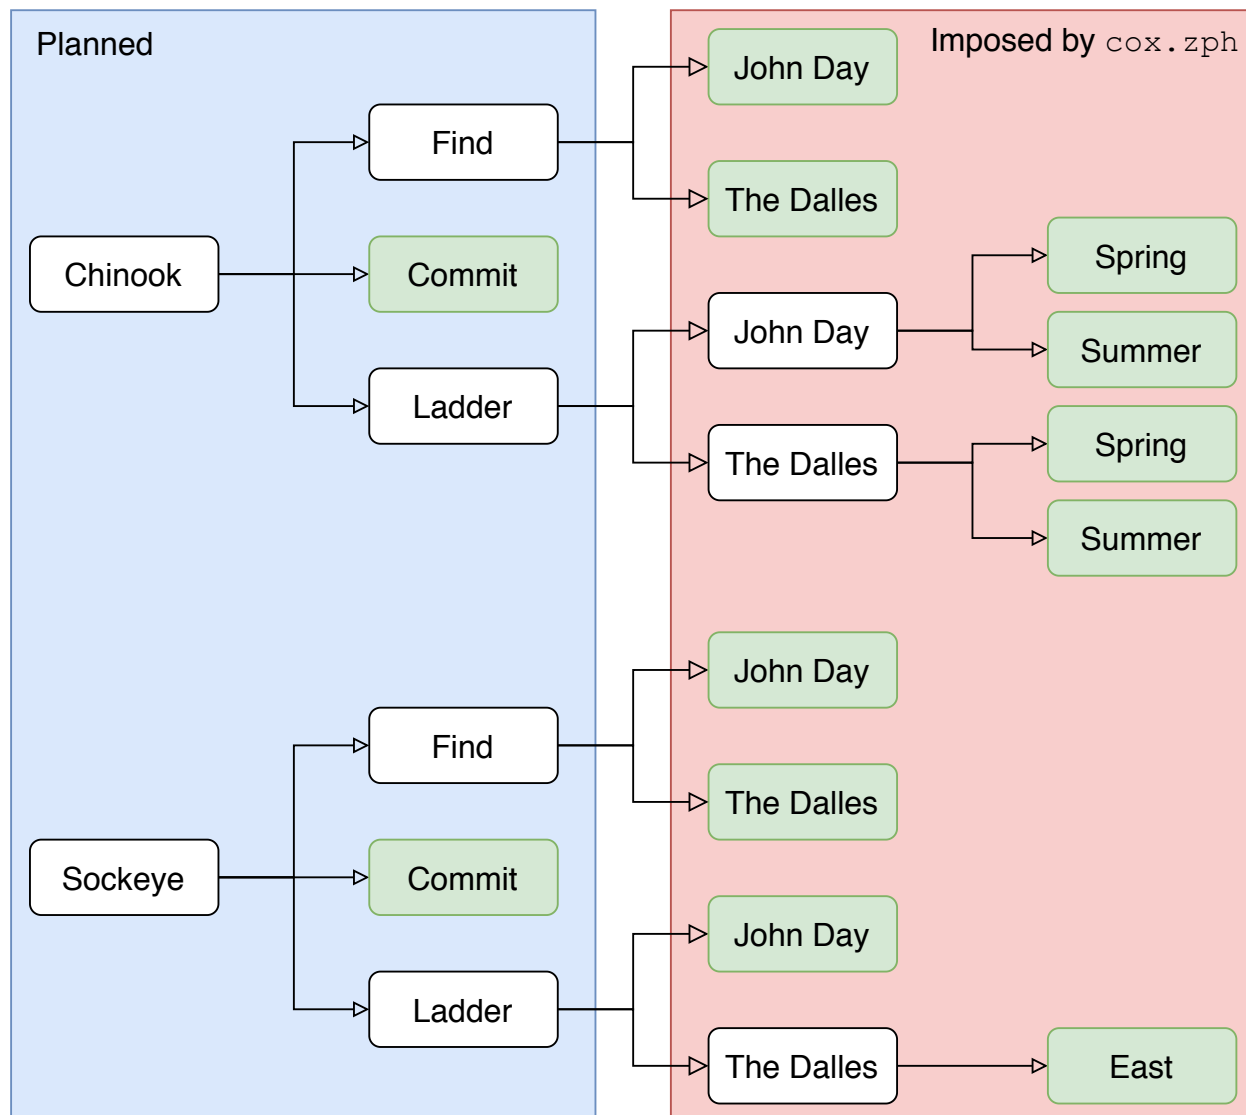


Figure A.1: A diagram of the models we fit. Originally we planned to fit the 6 models shown in blue on the left. However, due to model diagnostic violations (specifically violations detected by `cox.zph`) we were forced to split further to allow baseline hazard functions to differ between factors. These more specific final models are shown in red. Our final set of models which passed all model diagnostics are shown in green.

core Windows compute server. We parallelized using the R package `parallel` (R Core Team, 2019) and set our random seeds using L'ecuyer et al. (2002)'s method.

A.2 Model Summaries

Below, we include the model output for our final selected models. For the three models where a highly significant density effect was detected we include that density effect. Otherwise we include the best-fitting null model.

Find Chin JD

Call:

```
phreg(formula = Surv(tstart, tstop, success) ~ day + middle +
      spring + temp.c + spill + chin, data = find.chin.JD.trunc,
      dist = "pch")
```

Covariate		W.mean	Coef	Exp(Coef)	se(Coef)	Wald p
day		0.577	2.245	9.445	0.134	0.000
middle						
	FALSE	0.375	0	1		(reference)
	TRUE	0.625	-0.523	0.593	0.075	0.000
spring						
	FALSE	0.364	0	1		(reference)
	TRUE	0.636	0.044	1.045	0.113	0.697
temp.c		-0.297	0.616	1.852	0.120	0.000
spill		0.682	0.204	1.227	0.108	0.058
chin		-0.048	0.559	1.748	0.090	0.000

Events	804
Total time at risk	5255.9
Max. log. likelihood	-1997.2

LR test statistic	632.58
Degrees of freedom	6
Overall p-value	0

Find Chin TD

Call:

```
phreg(formula = Surv(tstart, tstop, success) ~ day + middle +
      spring + temp.c + spill + chin, data = find.chin.TD.trunc,
      dist = "pch")
```

Covariate		W.mean	Coef	Exp(Coef)	se(Coef)	Wald p
day		0.575	2.889	17.978	0.187	0.000
middle						
	FALSE	0.484	0	1	(reference)	
	TRUE	0.516	0.155	1.168	0.079	0.051
spring						
	FALSE	0.345	0	1	(reference)	
	TRUE	0.655	0.654	1.923	0.164	0.000
temp.c		-0.322	1.299	3.667	0.191	0.000
spill		0.777	-0.075	0.928	0.191	0.695
chin		0.093	0.324	1.383	0.082	0.000

Events	751
Total time at risk	6445.7
Max. log. likelihood	-2030.9
LR test statistic	669.09
Degrees of freedom	6
Overall p-value	0

Fishway Chin JD Spring

Call:

```
phreg(formula = Surv(tstart, tstop, success) ~ day + site.JDN +
      temp.c + temp.ct, data = ladder.chin.JD.spring.trunc, dist = "pch")
```

Covariate		W.mean	Coef	Exp(Coef)	se(Coef)	Wald p
day		0.755	2.945	19.002	0.382	0.000
site.JDN						
	FALSE	0.738	0	1		(reference)
	TRUE	0.262	0.203	1.225	0.102	0.047
temp.c		-0.578	0.296	1.344	0.154	0.055
temp.ct		-2.409	-0.003	0.997	0.012	0.798

Events	439
Total time at risk	1707
Max. log. likelihood	-935.96
LR test statistic	198.41
Degrees of freedom	4
Overall p-value	0

Fishway Chin JD Summer

Call:

```
phreg(formula = Surv(tstart, tstop, success) ~ day + site.JDN +
      temp.c + temp.ct, data = ladder.chin.JD.summer.trunc, dist = "pch")
```

Covariate		W.mean	Coef	Exp(Coef)	se(Coef)	Wald p
day		0.797	2.360	10.594	0.282	0.000
site.JDN						
	FALSE	0.436	0	1		(reference)
	TRUE	0.564	-0.024	0.977	0.110	0.829
temp.c		0.292	-0.651	0.521	0.249	0.009
temp.ct		0.619	0.437	1.548	0.048	0.000

Events	345
Total time at risk	1101.5
Max. log. likelihood	-681.38
LR test statistic	128.25
Degrees of freedom	4
Overall p-value	0

Fishway Chin TD Spring

Call:

```
phreg(formula = Surv(tstart, tstop, success) ~ day + site.TDN +
      temp.c + temp.ct, data = ladder.chin.TD.spring.trunc, dist = "pch")
```

Covariate		W.mean	Coef	Exp(Coef)	se(Coef)	Wald p
day		0.589	2.783	16.161	0.257	0.000
site.TDN						
	FALSE	0.820	0	1		(reference)
	TRUE	0.180	0.108	1.114	0.130	0.405
temp.c		-0.520	-0.342	0.711	0.191	0.074
temp.ct		-2.105	0.063	1.065	0.021	0.003

Events	415
Total time at risk	1115.4
Max. log. likelihood	-659.5
LR test statistic	331.65
Degrees of freedom	4
Overall p-value	0

Fishway Chin TD Summer

Call:

```
phreg(formula = Surv(tstart, tstop, success) ~ day + site.TDN +
      temp.c + temp.ct, data = ladder.chin.TD.summer.trunc, dist = "pch")
```

Covariate		W.mean	Coef	Exp(Coef)	se(Coef)	Wald p
day		0.784	2.101	8.172	0.310	0.000
site.TDN						
	FALSE	0.671	0	1		(reference)
	TRUE	0.329	-0.224	0.799	0.127	0.076
temp.c		0.234	0.337	1.401	0.260	0.196
temp.ct		0.658	0.048	1.049	0.047	0.309

Events	332
Total time at risk	1017.7
Max. log. likelihood	-654
LR test statistic	99.77
Degrees of freedom	4
Overall p-value	0

Find Sock JD

Call:

```
phreg(formula = Surv(tstart, tstop, success) ~ day + middle +
       wenatchee + okanogan + temp.c + spill, data = find.sock.JD.trunc,
       dist = "pch")
```

Covariate		W.mean	Coef	Exp(Coef)	se(Coef)	Wald p
day		0.459	1.951	7.034	0.111	0.000
middle						
	FALSE	0.461	0	1		(reference)
	TRUE	0.539	-0.211	0.810	0.084	0.012
wenatchee		0.187	0.381	1.463	0.114	0.001
okanogan		0.456	0.195	1.215	0.094	0.039
temp.c		0.220	0.668	1.950	0.113	0.000
spill		0.546	-0.456	0.634	0.143	0.001

Events	609
Total time at risk	1836.4
Max. log. likelihood	-1067.7
LR test statistic	426.88
Degrees of freedom	6
Overall p-value	0

Find Sock TD

Call:

```
phreg(formula = Surv(tstart, tstop, success) ~ day + middle +
      wenatchee + okanogan + temp.c + spill, data = find.sock.TD.trunc,
      dist = "pch")
```

Covariate		W.mean	Coef	Exp(Coef)	se(Coef)	Wald p
day		0.589	2.703	14.926	0.190	0.000
middle						
	FALSE	0.525	0	1		(reference)
	TRUE	0.475	0.122	1.129	0.082	0.137
wenatchee		0.195	0.400	1.491	0.112	0.000
okanogan		0.432	0.131	1.140	0.094	0.164
temp.c		0.301	0.759	2.135	0.221	0.001
spill		0.677	-0.269	0.765	0.219	0.221

Events	616
Total time at risk	3905.7
Max. log. likelihood	-1524.1
LR test statistic	459.16
Degrees of freedom	6
Overall p-value	0

Fishway Sock JD

Call:

```
phreg(formula = Surv(tstart, tstop, success) ~ day + site.JDN +
      temp.c, data = ladder.sock.JD.trunc, dist = "pch")
```

Covariate		W.mean	Coef	Exp(Coef)	se(Coef)	Wald p
day		0.786	1.533	4.634	0.182	0.000
site.JDN						
	FALSE	0.676	0	1		(reference)
	TRUE	0.324	0.743	2.102	0.082	0.000
temp.c		0.300	0.214	1.239	0.179	0.231

Events	605
Total time at risk	2347.6
Max. log. likelihood	-1320.5
LR test statistic	209.60
Degrees of freedom	3
Overall p-value	0

Fishway Sock TD East

Call:

```
phreg(formula = Surv(tstart, tstop, success) ~ day + temp.c,
      data = ladder.sock.TDE.trunc, dist = "pch")
```

Covariate	W.mean	Coef	Exp(Coef)	se(Coef)	Wald p
day	0.784	1.199	3.315	0.156	0.000
temp.c	0.219	0.872	2.391	0.156	0.000

Events	553
Total time at risk	1871.9
Max. log. likelihood	-1169
LR test statistic	116.59
Degrees of freedom	2
Overall p-value	0

Commit Chin

Call:

```
glm(formula = success ~ day + site + spring + temp.c + thresh +
     chin.specific, family = binomial, data = prob.chin)
```

Deviance Residuals:

Min	1Q	Median	3Q	Max
-1.7907	-1.0415	-0.6673	1.1420	2.1301

Coefficients:

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	-0.03122	0.28168	-0.111	0.91176
day	-0.04247	0.23845	-0.178	0.85863
siteJDS	-1.02322	0.17067	-5.995	2.03e-09 ***
siteTDE	0.04591	0.16864	0.272	0.78544
siteTDN	0.11066	0.28167	0.393	0.69442
springTRUE	-0.31638	0.22854	-1.384	0.16625
temp.c	-1.22897	0.26553	-4.628	3.68e-06 ***
threshTRUE	1.50867	0.28215	5.347	8.94e-08 ***
chin.specific	0.75259	0.21903	3.436	0.00059 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

(Dispersion parameter for binomial family taken to be 1)

Null deviance: 2129.1 on 1554 degrees of freedom
Residual deviance: 1898.2 on 1546 degrees of freedom

AIC: 1916.2

Number of Fisher Scoring iterations: 4

Commit Sock

Call:

```
glm(formula = success ~ day + site + temp.c, family = binomial,
     data = prob.sock)
```

Deviance Residuals:

Min	1Q	Median	3Q	Max
-1.7249	-1.1053	0.7476	0.7485	1.3937

Coefficients:

	Estimate	Std. Error	z value	Pr(> z)	
(Intercept)	0.911939	0.243638	3.743	0.000182	***
day	0.319363	0.202688	1.576	0.115110	
siteJDS	-1.401290	0.183263	-7.646	2.07e-14	***
siteTDE	-0.098865	0.180317	-0.548	0.583497	
siteTDN	-0.313573	0.375813	-0.834	0.404064	
temp.c	-0.006558	0.314959	-0.021	0.983387	

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

(Dispersion parameter for binomial family taken to be 1)

Null deviance: 1564.5 on 1224 degrees of freedom
 Residual deviance: 1457.1 on 1219 degrees of freedom
 AIC: 1469.1

Number of Fisher Scoring iterations: 4

BIBLIOGRAPHY

- Broström, G. and Broström, M. G. (2019). Package ‘eha’.
- Burr, D. (1994). A comparison of certain bootstrap confidence intervals in the Cox model. *Journal of the American Statistical Association*, 89(428):1290–1302.
- Caudill, C. C., Keefer, M. L., Clabough, T. S., Naughton, G. P., Burke, B. J., and Peery, C. A. (2013). Indirect effects of impoundment on migrating fish: temperature gradients in fish ladders slow dam passage by adult Chinook salmon and steelhead. *PLoS One*, 8(12):e85586.
- Keefer, M. L., Caudill, C. C., Peery, C. A., and Moser, M. L. (2013). Context-dependent diel behavior of upstream-migrating anadromous fishes. *Environmental Biology of Fishes*, 96(6):691–700.
- Kleinbaum, D. G. and Klein, M. (2010). *Survival Analysis*, volume 3. Springer.
- L’ecuyer, P., Simard, R., Chen, E. J., and Kelton, W. D. (2002). An object-oriented random-number package with many long streams and substreams. *Operations Research*, 50(6):1073–1075.
- R Core Team (2019). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Therneau, T. M. and Lumley, T. (2014). Package ‘survival’.

Appendix B

APPENDICES TO CH. 3

B.1 General Operator

A more general formulation can be expressed as

$$\mathcal{L}u(x) = f(x), \quad \forall x \in \Omega \quad (\text{B.1})$$

$$b(u(x), \nabla u(x)) = c(x), \quad \forall x \in \partial\Omega \quad (\text{B.2})$$

where \mathcal{L} is a second-order partial differential operator and b is some function representing a boundary condition.

B.2 Dirichlet Boundary Conditions

Although Neumann boundary conditions on \mathbf{u} are most common, Dirichlet boundary conditions on \mathbf{u} can be incorporated into our approach, as can boundary conditions on \mathbf{f} . Using the basic FEM method, Neumann boundary conditions are considered “natural,” that is, the solution to $\mathbf{K}\mathbf{u} = \mathbf{L}\mathbf{f}$ will satisfy these BCs without any additional input. On the other hand Dirichlet boundary conditions are considered “essential,” meaning that we must impose these externally. Note that more sophisticated FEM methods may have different natural/essential BCs, so that these terms are not simply synonymous with Neumann/Dirichlet.

To demonstrate how to incorporate Dirichlet boundary conditions, suppose that our basic Neumann BC generates the FEM formula

$$\mathbf{K}\mathbf{u} = \mathbf{L}\mathbf{f} \quad (\text{B.3})$$

Our Dirichlet boundary condition can be expressed by splitting \mathbf{u} into component vectors $\mathbf{u}_{\partial\Omega}$ or \mathbf{u}_1 and \mathbf{u}_Ω or \mathbf{u}_2 representing the internal and external degrees of freedom (i.e., nodes/basis functions in our cases). This makes our FEM equation

$$\begin{bmatrix} \mathbf{K}_{11} & \mathbf{K}_{12} \\ \mathbf{K}_{21} & \mathbf{K}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{L}_{11} & \mathbf{L}_{12} \\ \mathbf{L}_{21} & \mathbf{L}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \end{bmatrix}. \quad (\text{B.4})$$

To impose a Dirichlet boundary condition of the form $\mathbf{B}\mathbf{u}_{\partial\Omega} = \mathbf{c}$, we replace the first block-rows of this equation

$$\begin{bmatrix} \mathbf{B} & \mathbf{0} \\ \mathbf{K}_{21} & \mathbf{K}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{L}_{21} & \mathbf{L}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{c} \\ \mathbf{f}_2 \end{bmatrix}. \quad (\text{B.5})$$

The first set of equations determines \mathbf{u}_1 , assuming that \mathbf{B} is invertible. Then we are left with

$$\mathbf{K}_{21}\mathbf{u}_1 + \mathbf{K}_{22}\mathbf{u}_2 = \mathbf{L}_{21}\mathbf{c} + \mathbf{L}_{22}\mathbf{f}_2. \quad (\text{B.6})$$

Assuming that \mathbf{c} is a constant vector, we are left with a simple offset to the mean of \mathbf{u}_2 , i.e.,

$$\mathbf{u}_2 = \boldsymbol{\mu} + \mathbf{K}_{22}^{-1}\mathbf{L}_{22}\mathbf{f}_2, \quad (\text{B.7})$$

$$\boldsymbol{\mu} = \mathbf{L}_{21}\mathbf{c} - \mathbf{K}_{21}\mathbf{u}_1 \quad (\text{B.8})$$

$$= (\mathbf{L}_{21} + \mathbf{K}_{21}\mathbf{B}^{-1})\mathbf{c}, \quad (\text{B.9})$$

In the event that $\mathbf{c} = \mathbf{0}$, this offset is also $\mathbf{0}$ and we are left with a lower-dimensional version of the same FEM equation. In the event that \mathbf{c} is itself a lower-dimensional GMRF rather than a constant, then instead of a mean offset, this results in a more complex Gaussian random field with covariance

$$(\mathbf{L}_{21} + \mathbf{K}_{21}\mathbf{B}^{-1})\boldsymbol{\Sigma}_c(\mathbf{L}_{21} + \mathbf{K}_{21}\mathbf{B}^{-1})^T + \mathbf{K}_{22}^{-1}\mathbf{L}_{22}\boldsymbol{\Sigma}_{f_2}\mathbf{L}_{22}^T\mathbf{K}_{22}^{-T}, \quad (\text{B.10})$$

This can be inverted using the Woodbury matrix formula, giving us

$$\mathbf{Q}_{f_1} = \mathbf{Q}_{f_2} - \mathbf{Q}_{f_2}(\mathbf{L}_{21} + \mathbf{K}_{21}\mathbf{B}^{-1})(\mathbf{Q}_c + (\mathbf{L}_{21} + \mathbf{K}_{21}\mathbf{B}^{-1})^T\mathbf{Q}_{f_2}(\mathbf{L}_{21} + \mathbf{K}_{21}\mathbf{B}^{-1}))^{-1}(\mathbf{L}_{21} + \mathbf{K}_{21}\mathbf{B}^{-1})^T\mathbf{Q}_{f_2}. \quad (\text{B.11})$$

With the simplifying assumption that we replace \mathbf{L} with $\tilde{\mathbf{L}}$, we arrive at $\mathbf{L}_{21} = \mathbf{0}$ so that we obtain

$$\tilde{\mathbf{Q}}_{f_2} = \mathbf{Q}_{f_2} - \mathbf{Q}_{f_2}\mathbf{K}_{21}\mathbf{B}^{-1}(\mathbf{Q}_c + \mathbf{K}_{21}^T\mathbf{B}^{-T}\mathbf{Q}_{f_2}\mathbf{B}^{-1}\mathbf{K}_{21})^{-1}\mathbf{B}^{-1}\mathbf{K}_{21}^T\mathbf{Q}_{f_2}. \quad (\text{B.12})$$

Assuming that \mathbf{c} is much lower-dimensional than \mathbf{f}_1 , this may be Cholesky decomposed by first decomposing \mathbf{Q}_{f_2} and then computing a low-rank update to this decomposition.

Appendix C

APPENDICES TO CH. 4

C.1 Optimal Sampling Design with Linear Regression

Suppose that there exist known covariates \mathbf{X} within the region of interest, and that we wish to control for the effects of these covariates on our estimates. This can be done within a hierarchical Bayesian framework without losing any of the convenient properties derived above. Specifically, by placing a normal prior (equivalent to ridge regression) upon the regression coefficients, we retain an entirely multivariate normal conditional distribution after observing the data:

$$\begin{aligned} \mathbf{d} &= \mathbf{A}\mathbf{u} + \boldsymbol{\epsilon} & \boldsymbol{\epsilon} &\sim \mathcal{N}(0, \sigma_\epsilon^2 \mathbf{I}_n) \\ \mathbf{u} &= \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\eta} & \boldsymbol{\eta} &\sim \mathcal{N}(0, \mathbf{Q}) \\ & & \boldsymbol{\beta} &\sim \mathcal{N}(0, \sigma_\beta^2 \mathbf{I}_k). \end{aligned}$$

where k is the number of covariates and n the number of data points. This gives us the following conditional distribution after observing \mathbf{d}

$$\begin{aligned} [\boldsymbol{\eta}, \boldsymbol{\beta} | \mathbf{d}] &= [\mathbf{d} | \boldsymbol{\eta}, \boldsymbol{\beta}] [\boldsymbol{\eta}] [\boldsymbol{\beta}] \\ &= \mathcal{N}(\mathbf{X}\boldsymbol{\beta} + \boldsymbol{\eta}, \sigma_\epsilon^2 \mathbf{I}_n) \mathcal{N}(0, \mathbf{Q}) \mathcal{N}(0, \sigma_\beta^2 \mathbf{I}_k). \end{aligned}$$

This distribution has a precision matrix of

$$\mathbf{Q}' = \begin{bmatrix} \mathbf{Q} + \frac{1}{\sigma_\epsilon^2} \mathbf{A}^T \mathbf{A} & \frac{1}{\sigma_\epsilon^2} \mathbf{A}^T \mathbf{A} \mathbf{X} \\ \frac{1}{\sigma_\epsilon^2} \mathbf{X}^T \mathbf{A}^T \mathbf{A} & \frac{1}{\sigma_\beta^2} \mathbf{I}_k + \frac{1}{\sigma_\epsilon^2} \mathbf{X}^T \mathbf{A}^T \mathbf{A} \mathbf{X} \end{bmatrix}.$$

This matrix can be substituted for the matrix \mathbf{Q} in the objective function and the remainder of the model building proceeds analogously. The vector \mathbf{v} will have terms for each of the covariates corresponding to the integrated value of that covariate across the spatial area. The terms involving \mathbf{X} will often be dense, but so long as the number of parameters k remains small the full matrix should remain relatively sparse, and optimization should remain computationally feasible.

C.2 Discrete Uncertainty in the Precision Matrix

There are at least two notions of uncertainty that one might apply to the precision matrix. The first is a discrete distribution for \mathbf{Q} in which several fixed options \mathbf{Q}_i are weighted with probabilities p_i . Such a situation may arise due to uncertainty in the functional form of the covariance function, or as an approximation to a continuous distribution for \mathbf{Q} . In this situation, our objective function becomes probabilistic, with mean given by:

$$\sum_i p_i \mathbf{v}^T (\mathbf{Q}_i + \sum x_i \mathbf{a}_i \mathbf{a}_i^T)^{-1} \mathbf{v}.$$

This objective function may still be used within our MILP, as follows:

$$\begin{aligned} & \underset{\mathbf{x}, \mathbf{y}_j, \mathbf{z}_j}{\text{minimize}} && \sum p_j (\mathbf{v}^T \mathbf{y}_j) \\ & \text{subject to} && \mathbf{Q}_j \mathbf{y}_j + \frac{1}{\sigma^2} \mathbf{A}^T \mathbf{z}_j = \mathbf{v}, \\ & && \mathbf{x}_i * \mathbf{A}_i \mathbf{y}_j = \mathbf{z}_j \end{aligned} \tag{C.1}$$

$$\text{subject to} \quad \mathbf{x}_i \in \{0, 1\}$$

We know of no way to represent continuous uncertainty in \mathbf{Q} will maintaining model linearity.

C.3 Using the Covariance Matrix

Some spatial statistical methods use sparse versions of the covariance matrix instead. These remain consistent with our approach. In order to utilize these methods, all one needs do is modify the constraint

$$\mathbf{Q} \mathbf{y} + \frac{1}{\sigma^2} \mathbf{A}^T \mathbf{z} = \mathbf{v}$$

to

$$\mathbf{y} + \frac{1}{\sigma^2} \mathbf{\Sigma} \mathbf{A}^T \mathbf{z} = \mathbf{\Sigma} \mathbf{v}.$$

Both the matrix $\mathbf{\Sigma} \mathbf{A}^T$ and the vector $\mathbf{\Sigma} \mathbf{v}$ may be precalculated.

C.4 Design-Based Knapsack Sampling

C.4.1 Without Spatial Covariance

Sampling under knapsack constraints can be accomplished using methods similar to those in Stoica and Babu (2010). This method may accommodate regression models, as well as spatial balancing, but may not accommodate explicit accounting for spatial covariance without additional modification. Using Stoica and Babu's notation, assume that we have a set of K regression vectors \mathbf{a}_k , each of which is sampled in proportion equal to p_k of the total number of sites (proportions are used by Stoica, but we can just as easily imagine using these "proportions" as weights). In this case the expected Fisher information matrix is given by Stoica and Babu

$$E[\mathbf{R}] = \sum_{k=1}^K p_k \mathbf{a}_k \mathbf{a}_k^T = \mathbf{A}^T \mathbf{P} \mathbf{A}$$

where

$$\mathbf{A}^T = \begin{bmatrix} \mathbf{a}_1 & \dots & \mathbf{a}_K \end{bmatrix}$$

$$\mathbf{P} = \begin{bmatrix} p_1 & 0 & \dots & 0 \\ 0 & p_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & \dots & p_K \end{bmatrix}$$

and $\sum_{k=1}^K p_k = 1$.

We wish to modify the final constraint to a more general knapsack type constraint $\sum_{k=1}^K c_k \tilde{\mathbf{P}}_k = B$, so that the expected cost is equal to a given budget B . To accomplish this, we may set $\tilde{\mathbf{P}}_k = \frac{B}{c_k} p_k$ so that

$$\sum_{k=1}^K c_k \tilde{\mathbf{P}}_k = B \sum_{k=1}^K p_k = B$$

and

$$\mathbf{R} = \sum_{k=1}^K \tilde{\mathbf{P}}_k \mathbf{a}_k \mathbf{a}_k^T = \sum_{k=1}^K p_k \left(\sqrt{\frac{B}{c_k}} \mathbf{a}_k \right) \left(\sqrt{\frac{B}{c_k}} \mathbf{a}_k \right)^T$$

Therefore whatever the minimum \mathbf{R} that can be achieved, either parameterization will be able to achieve it. One may therefore leverage the homotopy method suggested by Stoica and Babu to solve the modified problem by transforming the regression vectors to

$$\sqrt{\frac{B}{c_k}} \mathbf{a}_k$$

and then transforming the resulting p_k to obtain the knapsack constrained sampling weights. These weights may be plugged into e.g. a GRTS sampling procedure to obtain a spatially balanced, knapsack constrained sampling design.

C.4.2 With Spatial Covariance

In the event that one wishes to collect a design-based sample with respect to a particular spatial covariance function, this may be accomplished with some additional modification to the methods.

Let us suppose we wish to minimize

$$\mathbf{v}^T (\mathbf{Q} + \sum_i p_i \mathbf{a}_i \mathbf{a}_i^T)^{-1} \mathbf{v} = \mathbf{v}^T (\mathbf{Q} + \mathbf{A} \mathbf{P} \mathbf{A}^T)^{-1} \mathbf{v},$$

the conditional variance of a particular linear combination w of points (e.g. the total kriging variance in space). This equation strongly resembles that in the Elfving theorem, so we will follow an algebraic derivation to arrive at a similar result. First we will introduce the augmented problem

$$\begin{aligned} & \underset{\mathbf{p}, \mathbf{c}, \mathbf{d}}{\text{minimize}} && h = \mathbf{c}^T \mathbf{P}^+ \mathbf{c} + \mathbf{d}^T \mathbf{d} \\ & \text{subject to} && \mathbf{A}^T \mathbf{c} + \mathbf{Q}^{1/2} \mathbf{d} = \mathbf{v}, \\ & && \sum p_i = 1, \\ & && c_i = 0 \iff p_i = 0 \end{aligned}$$

We propose that, for fixed p this augmented minimization problem has solution

$$\mathbf{c}_0 = \mathbf{P}\mathbf{A}^T(\mathbf{Q} + \mathbf{A}\mathbf{P}\mathbf{A}^T)^{-1}\mathbf{v} \quad \mathbf{d}_0 = \mathbf{Q}^{1/2}(\mathbf{Q} + \mathbf{A}\mathbf{P}\mathbf{A}^T)^{-1}\mathbf{v}$$

and thus objective function $h_0 = \mathbf{v}^T(\mathbf{Q} + \mathbf{A}\mathbf{P}\mathbf{A}^T)^{-1}\mathbf{v}$ as above. If we can verify this claim then the augmented problem will have the same minimum as the simpler problem above.

We now show that this is correct:

$$\begin{aligned} h \geq h_0 &\iff \begin{bmatrix} \mathbf{c}_0^T \mathbf{P}^+ \mathbf{c}_0 + \mathbf{d}_0^T \mathbf{d}_0 & \mathbf{v}^T \\ \mathbf{v} & \mathbf{Q} + \mathbf{A}^T \mathbf{P} \mathbf{A} \end{bmatrix} \geq 0 \\ &\iff \begin{bmatrix} \tilde{\mathbf{c}}_0^T \tilde{\mathbf{P}}^+ \tilde{\mathbf{c}}_0 + \mathbf{d}_0^T \mathbf{d}_0 & \tilde{\mathbf{c}}_0^T \tilde{\mathbf{A}} + \mathbf{d}_0^T \mathbf{Q}^{1/2} \\ \tilde{\mathbf{A}}^T \tilde{\mathbf{c}}_0 + \mathbf{Q}^{1/2} \mathbf{d}_0 & \mathbf{Q} + \tilde{\mathbf{A}}^T \tilde{\mathbf{P}} \tilde{\mathbf{A}} \end{bmatrix} \geq 0 \\ &\iff \begin{bmatrix} \tilde{\mathbf{c}}_0^T & 0 \\ 0 & \tilde{\mathbf{A}}^T \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{P}}^{-1/2} \\ \tilde{\mathbf{P}}^{1/2} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{P}}^{-1/2} & \tilde{\mathbf{P}}^{1/2} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{c}}_0 & 0 \\ 0 & \tilde{\mathbf{A}} \end{bmatrix} + \begin{bmatrix} \mathbf{d}_0^T \\ \mathbf{Q}^{1/2} \end{bmatrix} \begin{bmatrix} \mathbf{d}_0 & \mathbf{Q}^{1/2} \end{bmatrix} \geq 0 \end{aligned}$$

where $\tilde{\mathbf{c}}$ and $\tilde{\mathbf{P}}$ and $\tilde{\mathbf{A}}$ all correspond to matrices with the zero elements from \mathbf{c} and \mathbf{P} removed.

This last expression is the sum of two PSD matrices, so we find that indeed, h_0 is the minimum. Now let us suppose that \mathbf{c} is fixed. Since \mathbf{Q} is nonsingular, the equality constraints gives us a solution for \mathbf{d} :

$$\mathbf{d} = \mathbf{Q}^{-1/2}(\mathbf{v} - \mathbf{A}^T \mathbf{c}).$$

Then we find that P is minimized just as if \mathbf{d} did not exist (following Stoica and Babu 2010) at

$$p_i = \frac{|c_i|}{\sum_j |c_j|}$$

This can be argued by Cauchy-Schwartz, since this choice of p_i achieves $\mathbf{c}^T \mathbf{P}^+ \mathbf{c} = (\sum_i |c_i|)^2$, but

$$\left(\sum_i |c_i| \right)^2 = \left(\sum_i \frac{|c_i|}{\sqrt{p_i}} \sqrt{p_i} \right)^2 \leq \left(\sum_i \frac{c_i^2}{p_i} \right) \left(\sum_i p_i \right) = \left(\sum_i \frac{c_i^2}{p_i} \right)$$

Now since \mathbf{c} is not fixed, we re-express \mathbf{p} and \mathbf{d} in terms of \mathbf{c} and our minimization problem becomes:

$$\underset{\mathbf{c}}{\text{minimize}} \quad \left(\sum_i |c_i| \right)^2 + (\mathbf{v} - \mathbf{A}^T \mathbf{c})^T \mathbf{Q}^{-1} (\mathbf{v} - \mathbf{A}^T \mathbf{c})$$

This is a linearly constrained convex quadratic program of the form

$$\begin{aligned} \underset{\mathbf{c}}{\text{minimize}} \quad & \mathbf{z}^T \mathbf{z} + (\mathbf{v} - \mathbf{A}^T \mathbf{c})^T \mathbf{Q}^{-1} (\mathbf{v} - \mathbf{A}^T \mathbf{c}) \\ \text{subject to} \quad & z_i \geq c_i, \\ & z_i \geq -c_i \end{aligned}$$

This problem may lose the sparsifying effect from the ℓ_1 minimization that may be used in the absence of an explicit covariance term, due to the now-quadratic objective function. Similar to the above, we may allow variable costs ($\sum q_i p_i = B$) by normalizing with respect to budget ($\sum \frac{q_i}{B} p_i = 1$) and then subsuming cost fraction into observations. That is:

$$\begin{aligned} \underset{\mathbf{p}}{\text{minimize}} \quad & \mathbf{v}^T (\mathbf{Q} + \sum p_i \mathbf{a}_i \mathbf{a}_i^T)^{-1} \mathbf{w} \\ \text{subject to} \quad & \sum q_i p_i \leq B \end{aligned}$$

is equivalent to

$$\begin{aligned} \underset{\mathbf{p}}{\text{minimize}} \quad & \mathbf{v}^T \left(\mathbf{Q} + \sum p_i \left(\sqrt{\frac{B}{q_i}} \mathbf{a}_i \right) \left(\sqrt{\frac{B}{q_i}} \mathbf{a}_i \right)^T \right)^{-1} \mathbf{v} \\ \text{subject to} \quad & \sum p_i \leq 1 \end{aligned}$$

Covariates may be included in the same way as in our main MILP model.

BIBLIOGRAPHY

Stoica, P. and Babu, P. (2010). Algebraic derivation of Elfving theorem on optimal experiment design and some connections with sparse estimation. *IEEE Signal Processing Letters*, 17(8):743–745.