Preface

THIS BOOK, AS well as my broader interest in ecological forecasting, is the result of a bit of academic serendipity. In the summer of 2001, less than a year into grad school and at that perfect stage for latching onto an idea, Clark et al. (2001) published "Ecological Forecasts: An Emerging Imperative," arguing that to make ecology more relevant to society, and the problems we claim to care about, requires that we make our discipline more predictive. The young, nature-loving, fledgling ecologist in me and the quantitatively focused former engineering student in me both embraced this call to action. By a bit of luck the lead author happened to be my graduate advisor, and the rest, so they say, is history.

Or is it? In the intervening years since Clark et al. (2001), to what degree has ecological forecasting emerged as a distinct discipline unto itself? Is the supply of ecological forecasts keeping up with the demand? Is ecology becoming a more predictive science? I believe the time is ripe to assess where we have come as a field. What have been the successes of ecological forecasting, and what is holding us back? Having spent the last 15 years working and teaching in this area, I am convinced one of the things holding us back is a shortage of training in the set of statistical, modeling, and informatics tools necessary to tackle the challenges before us in an era of big data and even bigger environmental problems. This book aims to help fill that gap.

In many ways this book describes an emerging new way of doing science in ecology that involves a much closer connection between models and data. As such the skills and techniques in an ecological forecaster's toolbox include, but go beyond, those of a traditional modeler's. Forecasting a real place and time is a much more data-intensive process than theoretical modeling, and draws heavily upon skills from statistics and informatics. At the same time it relies on more sophisticated modeling than is often employed for pure data analysis. Indeed, most of this book is focused on the statistics and informatics of model-data fusion and data assimilation. The literature on these topics has largely been scattered across other disciplines, and adoption by ecologists has been encumbered by discipline-specific jargon and mathematical notation that is difficult to penetrate. This has been unfortunate, because while many of these techniques require a change in how one perceives a problem, they are not unreasonably complicated or difficult to implement.

As noted earlier, the backbone of this book covers topics related to the statistics and informatics of model-data fusion: data management, workflows, Bayesian statistics, uncertainty analysis, fusing multiple data sources, assessing model performance, and a suite of data assimilation techniques. Throughout these chapters the aim is to provide enough technical detail to enable the reader to understand how the methods work and to be able to implement them. Subtleties will be highlighted when they will help readers avoid common pitfalls or misinterpretations, but otherwise the goal is

for the text to remain accessible and not get bogged down in esoteric details. The focus is very much on application not theory. That said, *this is definitely a book about concepts, not a training manual.*

Interspersed among these concept chapters are "case study" chapters on specific ecological subdisciplines that aim to highlight the successes and failures of ecological forecasting in each: natural resources, endangered species, epidemiology, and the carbon cycle. These chapters also aim to make it clear that forecasting is useful across broad swaths of ecology, and not limited to specific subdisciplines. Indeed, to date expertise in many of the topics outlined here has been distributed across subdisciplines, with some fields further along in adopting one set of approaches, while another field might excel in some other area. The book also begins and ends with a discussion of decision support, as improving decision making is often the primary driver of many ecological forecasts.

Finally, as with all technical topics, we inevitably need to "get our hands dirty" to really understand what's going on. To facilitate this hands-on learning, while also aiming to keep material up to date, I have made code and tutorials available online as supplemental material. Rather than providing static code within the book itself, the intention is to keep this material up to date and publicly available as the field evolves, something that would not be possible if this book was tied to specific code. This material draws heavily from my graduate course on this same topic and my experience teaching this material at a variety of workshops and summer courses. All code and tutorials are currently hosted at https://github.com/EcoForecast. In addition, this site includes a second repository dedicated to the code used to generate the figures, tables, and analyses described in this book. More generally, the site ecoforecast .org provides information not only about this book and its tutorials but also other literature, workshops, and new developments in the field.