

Foreword

“I’d like to use machine learning, but I can’t invest much time.” That is something you hear all too often in industry and from researchers in other disciplines. The resulting demand for hands-free solutions to machine learning has recently given rise to the field of automated machine learning (AutoML), and I’m delighted that with this book, there is now the first comprehensive guide to this field.

I have been very passionate about automating machine learning myself ever since our Automatic Statistician project started back in 2014. I want us to be really ambitious in this endeavor; we should try to automate all aspects of the entire machine learning and data analysis pipeline. This includes automating data collection and experiment design; automating data cleanup and missing data imputation; automating feature selection and transformation; automating model discovery, criticism, and explanation; automating the allocation of computational resources; automating hyperparameter optimization; automating inference; and automating model monitoring and anomaly detection. This is a huge list of things, and we’d optimally like to automate all of it.

There is a caveat of course. While *full* automation can motivate scientific research and provide a long-term engineering goal, in practice, we probably want to *semiautomate* most of these and gradually remove the human in the loop as needed. Along the way, what is going to happen if we try to do all this automation is that we are likely to develop powerful tools that will help make the practice of machine learning, first of all, *more systematic* (since it’s very ad hoc these days) and also *more efficient*.

These are worthy goals even if we did not succeed in the final goal of automation, but as this book demonstrates, current AutoML methods can already surpass human machine learning experts in several tasks. This trend is likely only going to intensify as we’re making progress and as computation becomes ever cheaper, and AutoML is therefore clearly one of the topics that is here to stay. It is a great time to get involved in AutoML, and this book is an excellent starting point.

This book includes very up-to-date overviews of the bread-and-butter techniques we need in AutoML (hyperparameter optimization, meta-learning, and neural architecture search), provides in-depth discussions of existing AutoML systems, and

thoroughly evaluates the state of the art in AutoML in a series of competitions that ran since 2015. As such, I highly recommend this book to any machine learning researcher wanting to get started in the field and to any practitioner looking to understand the methods behind all the AutoML tools out there.

San Francisco, USA
Professor, University of Cambridge and
Chief Scientist, Uber
October 2018

Zoubin Ghahramani