Applied Causal Inference Powered by ML and AI

Victor Chernozhukov* Christian Hansen† Nathan Kallus‡
Martin Spindler§ Vasilis Syrgkanis¶

February 5, 2026

Publisher: Online
Version 0.1.2

* MIT
† Chicago Booth
‡ Cornell University
§ Hamburg University
¶ Stanford University

Preface

This book aims to provide a working introduction to the emerging fusion of modern statistical inference – aka machine learning (ML) or artificial intelligence (AI) – and causal inference methods. The book is aimed at upper level undergraduates and master's-level students as well as doctoral students focusing on applied empirical research. A sufficient background for the core material is one semester of introductory econometrics and one semester of machine learning. We hope the book is also useful to empirical researchers looking to apply modern methods in their work.

The book provides an overview of key ideas in both predictive inference and causal inference and shows how predictive tools are key ingredients to answering many causal questions. We use the term predictive inference to refer to settings where prediction or description is the main goal such that models and estimates do not need a causal interpretation. ML/AI tools are largely designed to answer predictive inference questions, and we provide a high-level overview of popular ML/AI methods (such as Lasso, random forests, and deep neural networks, among others) to provide background for readers less familiar with these methods.

On the causal inference side, we introduce foundational ideas that provide the underpinning to attaching causal interpretations to statistical estimates. We discuss these ideas using the language of potential outcomes, directed acyclical graphs (DAGs), and structural causal models (SCMs). We view the language of potential outcomes, DAGs, and SCMs as complementary. We recognize that readers coming from different backgrounds may be more familiar or disposed to one of potential outcomes, DAGs, or SCMs, but we strongly believe that individuals interested in causal inference should be familiar with each of these frameworks. We find that they all offer useful insights and being able to communicate using each framework allows one to communicate with audiences interested in understanding causality coming from many different backgrounds.

The book has two main sections: Core Material and Advanced Topics. The Core Material provides the main content of the book. After concluding the Core Material, a reader should have an idea of the key ideas underlying both predictive and causal inference and how to wed these ideas to learn canonical objects in causal inference settings. The Core Material is made up of chapters that move between predictive inference and causal inference, typically by first introducing tools developed for predictive inference and then showing how these tools can be used as inputs to answering causal inference questions. The Advanced Topics then provide extensions of the Core Material to settings with more complicated causal structures, such as instrumental variables models, to settings where understanding heterogeneity in causal effects is the goal, and to specific popular settings in empirical work such as Difference-in-Differences.

Within sections, blocks marked with ⋆ require more substantial preparation in mathematical statistics. We recommend that the reader looking to apply machine learning methods in their work skim or pass them on their first reading and return to them at their leisure.

Short lists of references and study problems are included after each chapter to offer the reader opportunities to investigate further and consolidate their knowledge.

We would like to also acknowledge the tremendous and exceptional help and expertise provided by Philipp Bach, Wenxuan Guo, Andy Haupt, Shunzhuang Huang, David Hughes, Jannis Kück, Malte Kurz, Sven Klassen, Oliver Schacht, Sophie Sun, Vira Semenova, Gulin Tuzcuoglu, Suhas Vijaykumar, John Walker, Thomas Wiemann, Justin Young, and Dake Zhang with both writing and developing supporting Notebooks in R and Python. We are also grateful to Alexander Quispe and Anzony Quispe for developing a Bookdown version of the notebooks and providing other complementary topics and great examples.

Chernozhukov, Hansen, Kallus, Spindler & Syrgkanis