Introduction

Motivation

Richard Feynman describes one of the "most dramatic moments in science" as when two great fields suddenly come together and are unified [23]. I cannot think of a single new scientific method more revolutionary over the past 30 years than Causal Inference and Reinforcement Learning have been: Both have sparked widespread conversations and inventions that have completely overhauled approaches to problems such as personalized medicine, landing such a massive impact that they garnered widespread media attention. The Reinforcement Learning revolution, in particular, has overhauled many disparate fields, giving rise to the invention of technology such as autonomous robots and self-driving cars[1].

Certainly one of the most formative moments in my own education was when I realized that in the application I was working on - development of automated personalized treatment plans in immuno-radiobiology - the theories of Causal Inference and Reinforcement Learning were complementary, rather than contradictory[10]. My Ph.D program was fairly unique in that I was not only entrenched in the classical statistics taught in a traditional academic setting at Southern Methodist University, but also worked at the Medical AI and Automation lab at UT Southwestern Medical Center's department of Medical Physics. My footing in these two programs was relatively equal as I had two advisors: A physicist from UTSW, from whom I learned about Reinforcement Learning, and a biostatistician from SMU, from whom I learned about Causal Inference. While their general approaches to my dissertation research were very similar, the language surrounding each approach was very different, which introduced substantial barriers to communication. Moreover, the fields of Causal Inference and Reinforcement Learning have very different standards and conventions when it comes to validation and publication of results.

In light of the language mismatch, it is wholly unsurprising to me that the intimate relationship between Causal Inference and Reinforcement Learning is not widely known: In 2023 on PubMed, for example, 3939 articles mentioned Causal Inference but NOT Reinforcement Learning, 3700 articles mentioned Reinforcement Learning but NOT Causal Inference, while only 194 articles mentioned "Causal Reinforcement Learning", or both "Causal Inference" as well as "Reinforcement Learning". This unfortunate dichotomy drives a wedge into the middle of the densely connected web of research that gave rise to both fields, unnecessarily partitioning the knowledge graph into two smaller - and thus less powerful - components (see Figure 1-2).

The ability to see how a specific problem sits within the larger scientific context is imperative for impactful breakthroughs and discovery. The same is true of the computational and mathematical tools used to solve scientific problems, especially for those of us - statisticians and computer scientists - whose job function is to improve those tools: Awareness of

[1] There have also been many books on both topics. Perhaps the most famous is Judea Pearl's nontechnical guide to Causal Inference titled *The Book of Why*[55]. *What If..?* is the standard technical reference book for Causal Inference[28]. The standard text on Reinforcement Learning is *An Introduction to Reinforcement Learning*[74].

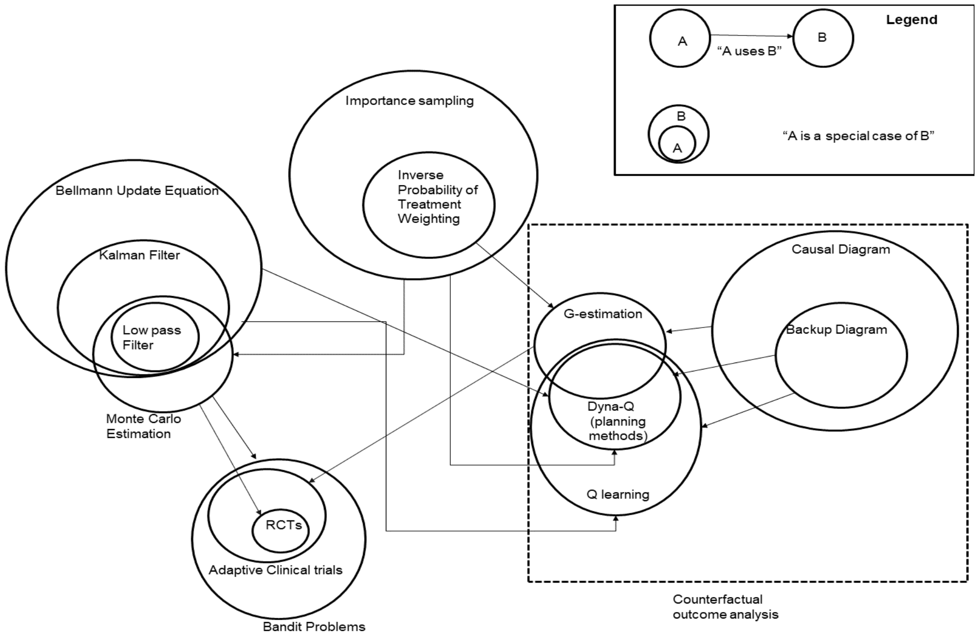

FIGURE 1: Illustration of connections in the Causal Reinforcement Learning methodological landscape.

the origin and context of any technical method is imperative to fully understanding how to improve it. By ignoring the connection to Reinforcement Learning when discussing Causal Inference, statisticians are effectively amputating a branch of the scientific literature which might contain valuable and relevant results and techniques. The converse is also true: By ignoring Causal Inference research, AI practitioners who use Reinforcement Learning miss out on powerful statistical frameworks and results. The result is an egregious inefficiency in the development of both fields.

There is a human tendency in research to cling to known tools, regardless of whether those tools are useful or correct for the problem at hand: The old adage when your only tool is a hammer, everything looks like a nail rings empirically true. For example, several studies have shown that the safest time to have a heart attack is actually during one of the major cardiology conferences, when all of the cardiologists are away[36, 35, 59]. The reason for this is because cardiologists have a preferred tool: heart stents, which they tend to apply lavishly, even if the patient's situation does not match the context in which heart stents are useful[70, 46]. In a more literal definition of the word "tool", many firefighters die from forest fires because they cannot drop their tools, subconsciously choosing to die beside their saw rather than to leave their heavy equipment behind and run[83, 63, 82].

It is difficult, if not impossible, to remove this cognitive bias. The problem can be mitigated, however, by expanding one's familiar toolset: The more - and the more diverse - tools in one's toolbox, the wider the breadth of problems which are solvable using those tools. One goal of this book is to put the tools of Causal Inference into the same toolbox as those of reinforcement learning: This unification of theory will allow the practitioner of Reinforcement Learning to view Causal Inference as a familiar tool (and vice versa) without feeling the need to learn the foundations of Computer Science or Statistics from the ground up to get there. The first four chapters of this book are oriented towards the first goal,

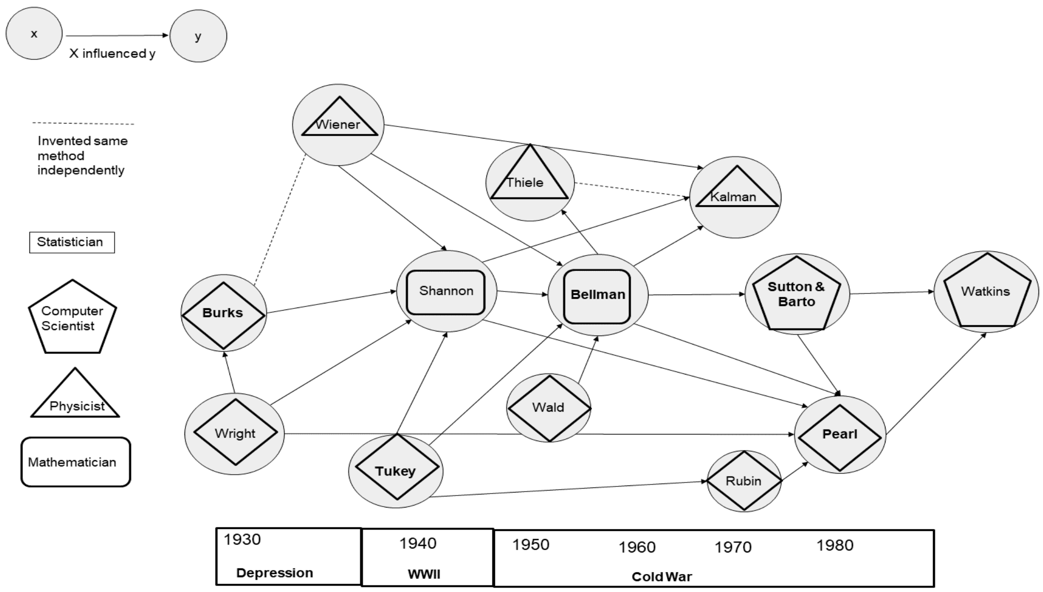

FIGURE 2: Causal diagram consisting of interactions between prominent figures in the development of Causal Inference and Reinforcement Learning.

providing an applications-oriented translation manual which focuses on the cross-talk and connections between the two fields.

The second goal of this book is to illustrate the translational landscape surrounding modern hybrid methods which leverage both causality and Reinforcement Learning. In addition to building bridges *within* Causal Reinforcement Learning, we wish to build bridges *between* Causal Reinforcement Learning and other fields, such as information theory, physics, and philosophy. These translational elements are critical to the modern decision-maker, especially to those who work in industry or academics who collaborate on large diverse teams of researchers. These external connections are most prominent in chapter 5, wherein we discuss the translational positioning of various modern hybrid methods which are known as "Causal Reinforcement Learning". However, external connections are also present in the earlier chapters, and a discussion of philosophical implications and analogies is presented throughout the work.

Structure

This book is primarily a monograph, though it contains some textbook-like elements. The first four chapters, specifically, cover a broad range of materials, some of which might be familiar to readers as undergraduate-level material. These foundations are not meant to be a comprehensive introduction; rather, they are included to facilitate discussion of the conceptual bridges between known concepts in Causal Inference and Reinforcement Learning, and to allow precise translation of notational conventions between the fields.

Additionally, the goal is for the book to be a self-contained manual for the expert in *either* Reinforcement Learning *or* Causal Inference, and the Reinforcement Learning expert may be unfamiliar with statistical fundamentals. Similarly, the professional statistician might be unfamiliar with basic Dynamic Programming concepts and notation, which form the backbone of Reinforcement Learning.

I have attempted to balance the content between Reinforcement Learning and Causal Inference. However, some asymmetry is inevitable, since there is more material in the field of Reinforcement Learning, both in terms of breadth as well as technical depth. This is apparent in the fact that the introductory textbook for Reinforcement Learning[74] is much longer than the introductory textbook for Causal Inference[28]. The discussion of Reinforcement Learning methods also typically requires more mathematical background. My approach to this issue is to keep the discussion Reinforcement Learning concepts generally broader with less technical detail. I will also keep the focus on Q-learning methods, which I believe are some of the most interconnected to Causal Inference.

A fully comprehensive academic textbook which builds the field of Causal Reinforcement Learning from the ground-up is in production by Elias Bareinboim. This book differs in several key regards. It is not a comprehensive academic textbook - indeed, there are many important methods which are not included in this book - rather it is a translational monograph directed at practitioners already familiar with Causal Inference and Reinforcement Learning. We do not attempt to provide an in-depth, comprehensive introduction to any field. Instead, we wish to highlight the position of the reader's field (Causal Inference or Reinforcement Learning) as a point in the densely connected web of science. Figure 3 highlights this idea: Bareinboim builds the mountain; this book builds the bridges.

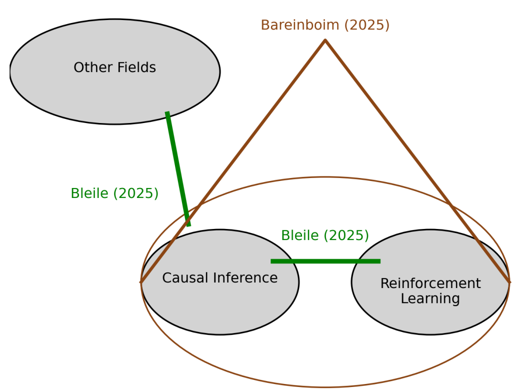

FIGURE 3: Illustration of the difference between this book and the textbook by Bareinboim (2025).

Prerequisites

Naturally, this book is written for the current practitioner of either Reinforcement Learning or Causal Inference. I will assume knowledge of basic probability (including expected values, Bayes' theorem, and conditional probabilities), calculus, and linear algebra throughout the text. Some proof sketches are included; these include some functional analysis. However, the functional analysis is kept at a conceptual level, and major theoretical concepts such as the Banach Fixed-Point Theorem are not discussed in detail.

Notation

Part of the difficulty in transferring knowledge of Reinforcement Learning to Causal Inference (and vice versa) is the differing terminology and notation between the two fields. This book introduces them both in a unified notational schema. Naturally, this unified notational schema differs from some conventional practices in each specific field.

> When I change the notational schema from convention in either field, I will indicate this by also providing the conventional notation in a shade box.

In this book, I will follow the statistical convention of denoting random variables with capital letters, e.g. A, R, X. Specific instances of these variables will always be denoted using small letters, e.g. $P(A = a) = 0.3$. I will always use j to denote the index of a Markov Decision Process. If the Markov process is episodic (such as in the case of a clinical trial with time-varying covariates), then episodes (subjects) will be indexed using i and timepoints within episodes will be indexed using t: j can then take on $i \times t$ unique values. If there is only one observation per subject, then $j = i = t$.