The Book of Why, by Judea Pearl & Dana Mackenzie. A book about how we ask questions and how to ask (and solve) better questions.

Adrian Soto-Mota

Review of: *The Book of Why* (Basic Books: New York, 2018)
Written by *Judea Pearl and Dana Mackenzie*

"If our conception of causal effects had anything to do with randomised experiments, the latter would have been invented 500 years before Fisher." (Pearl 2016).

Engaging in the cause-effect analysis is something we cannot escape doing. Regardless of what you study, regardless of where you live, regardless of when you grew up, we all make causal inferences every day of our lives. From the African plains to Twitter feeds, almost seamlessly, the question is asked: *why* do we perform thought (or actual) experiments before we act accordingly, both intentionally and unconsciously?

Everyone has a personal stance in debates about politics, the environment or the consequences of vaccinating children. Yet despite having different opinions, our perspectives are connected by an underlying belief in causal relationships; a concept examined by *The Book of Why* which speaks to all audiences. Reading Pearl's text has made me question and reconsider my own arguments from a new perspective. I invite you to test the logic of your own arguments by doing the same.

Judea Pearl is a distinguished computer scientist and philosopher who was among the first to mathematise causal modelling. During the last decade, he has written many "technical" books about causal inference and, in *The Book of Why* explores the very intuitive but slightly daunting concepts of causal modelling and its algebra. While doing it, he also tells the story of their development and the implications of their current and future implementation in empirical sciences and daily life.

"Correlation is not causation" but… What is causation?

Anyone who has taken a Statistics course knows and lives by this mantra. No one can dispute its veracity or relevance, and it is a guiding light for data analysts. However, these words obscure the meaning of causation in our approach to modern scientific investiagtion.

What is causation? Have you ever asked yourself that after hearing or repeating "correlation is not causation"? Honestly, I had not.

Definitions by negation (cold as the absence of heat, darkness as the absence of light) are usually considered imperfect because they do not actually explain what the

subject means or is (Macagno & Walton 2013). It is likely that your life currently revolves around elucidating one or more causal relationships through this process of negation. But, if correlation is not causation, how do you know when you have done your job? How do you know when you already found a causal relationship? Indulge me by engaging in a thought experiment:

Imagine you stopped for a takeaway chicken wrap on your way home. You are about to get to your place when a cosmic accident throws you to a parallel reality.

You arrive at a primitive Earth populated by farmers. They have developed statistical knowledge because of its usefulness in crop and cattle management but have no clue about Physics, Biology or Geography. Of course, this civilisation is highly dependent on agriculture and therefore, the Sun.

These people are convinced that rooster's crow makes the Sun rise and therefore, worship them as deities. Remember that chicken wrap in your bag? Guess what…? They found it! Logically, you are facing death accused of the ultimate form of blasphemy in this world.

Your first reaction is, of course, telling them that they got it mixed-up, however, they quickly refute your claims by showing you their "Big-data". For the last two millennia, they have carefully collected detailed datasets, and they proudly show their beautiful plots showing almost perfect correlations between the rooster's crow and the Sun rising.

To save your life, you would correctly discard their 0.99 correlated data directly on the grounds of being ludicrous and would immediately propose experiments blinding or muting the roosters to prove the Sun would rise anyway.

My example posits how we are inclined to agree that a carefully designed and simple experiment can unveil cause-effect relationships more effectively than large and immaculate datasets. However, many of our methods, and even further, many of our questions focus on correlations and elude the true notion of causality. How did this happen? As Pearl asks, *why* did this happen?

Asking 'why' is a much more complex process than simply identifying input-response mechanisms and attempting to manipulate them (there are a lot of beings capable of similar feats). When we ask why, we engage in the abstract exercise of explaining exact causes and their direct effects.

Despite being quite complex, asking why comes naturally to us. Perhaps you haven't payed much attention to how you (or anyone) assess cause-effect relationships? This human trait changed the world because it allowed us to outperform all other species in manipulating the environment.

Our brains excel at finding patterns or explanations. This explanation thirst is so ingrained in our nature that we feel anxious around uncertainty and can be susceptible to accept almost any explanation rather than none. As a result, we are ironically prone to make wrong causal-effect conclusions (particularly in complex scenarios).

Pearl's investigation reveals that it was not until we started attempting to teach machines and artificial systems how to learn and how to produce cause-effect

conclusions, that we were forced into thinking about how we think, and into formalising the language and steps involved in causal inference.

About the Book

The book has ten chapters and three main objectives:
1. Explaining the core concepts of causal modelling in non-mathmatical language.
2. Reviewing the history of how empirical science went astray from causal inference and the recent efforts made towards mathematising causality.
3. Analysing the consequences of implementing causality algorithms in machine learning as we stand in the dawn of the Big-data/machine-learning/AI era.

A Glossary of Key Terms:
Causation: a relationship that connects one element (the cause) with another element (the effect), where the first is at least partly responsible for the second, and the second is at least partly dependent on the first.
Correlation: The statistical association between two variables.
Counterfactuals: a conditional containing an if-clause which is contrary to fact.
Spontaneous generation: an obsolete body of thought on the ordinary formation of living organisms without descent from similar organisms.
Randomised controlled trial: A research model where participants are randomised to receive the treatment that is being studied or a placebo. Differences between groups are assessed statistically at specific times.
Bradford-Hill's Criteria: A group of nine criteria a proposed "cause" should fulfil in order to be accepted as such.

Critique and Conclusion

I confess that, at the beginning of the book, I felt it was a bit over-enthusiastic with some of the ideas and that adjectives such as "revolution", "paradigm-shift" or "ground-breaking" were used lightly. However, after a few chapters and particularly after the historical review concerning how statistics regarded (or should I say disregarded) causality in chapter 2, I could not agree more with the use of these words.

The Book of Why has changed my view on the true relevance and role that clinical trials such as double-blind, placebo controlled, and randomised studies have in the advancement of medical knowledge. Pearl develops the ideas of Fisher, Pearson and Bradford-Hill who view causality not as a direct object of study, but as a logic automatically implied when we find strong correlations that correspond with our theoretical background.

After reading the first two chapters, the idea of using counterfactuals (asking "what if?") as a better tool for approaching causality seemed so obvious and intuitive that made me feel slightly embarrassed of having accepted and defended Bradford-

Hill's 'Nine Criteria for Causality' (Bradford-Hill 1965) as the gold standard for causal relationships in my field.

As Pearl and Mackenzie identify in chapter five while discussing the causal relationship between smoking and cancer, Bradford-Hill's criteria are still useful as a description of how a discipline comes to accept a causal hypothesis in the light of new evidence. However, they are limited by the lack of methodology for their implementation or quantification. In other words, because the rely on the plausibility or coherence of an idea, they are still too subjective to be scientifically useful (or coded into a machine).

What has been the biggest debunked myth in Biology? I propose *spontaneous generation* (living creatures arising from non-living matter, as in maggots from rotten meat). There was a time when challenging this idea could bring one professional, social and even legal problems. Today, even the most recalcitrant creationists would agree that "maggots come from flies, not from rotten meat".
How did this happen? A very simple counterfactual is enough (What if I isolate the rotten meat?). Perhaps you remember Francesco Redi's experiment from high school, in which he placed a piece of rotten meat in a closed jar and a piece of rotten meat in an open jar to test spontaneous generation. Redi shocked his contemporaries' central belief about life by finding that maggots only appeared in the meat in the open jar.

Again, imagine that you are asked to prove the same idea "maggots come from flies, not from rotten meat" without using counterfactuals, but just the statistical methods you typically read in scientific journal of choice. How would you do it?

I knew about Redi's experiment before reading this book, I knew correlation is not causation, I knew I am supposed to find causal relationships in my work and somehow, I totally ignored counterfactuals in my methods or while I read scientific papers. Of course, there are limitations (practical, ethical and legal) to asking, "what if?" particularly in Biomedical Science. However, acknowledging that there is an "ideal" way of proving a certain idea can improve our methods even if we can't actually test it.

To conclude, I genuinely think that amidst the era of supercomputing and Big Data, the mathematisation of causality and the concepts contained in the 'the ladder of causation' (Pearl & Mackenzie 2018, Figure 1) are paradigm shifting and entail a real revolution for Science.
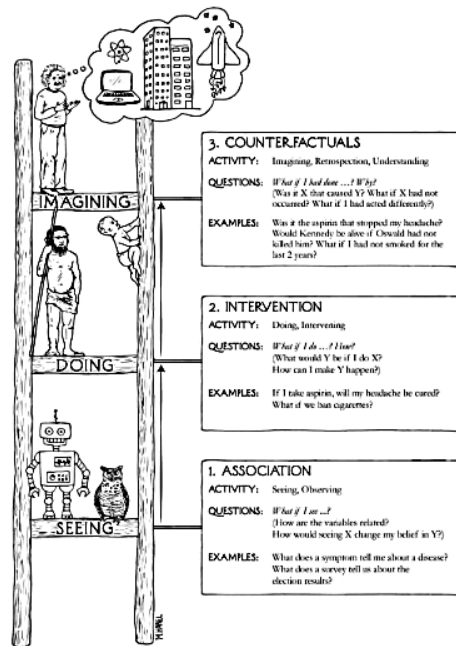
Figure 1: The Ladder of Causation (from Pearl and Mackenzie, 2018, p. 465).

Written in an accessible style, *The Book of Why* is worth reading regardless of your main academic focus or level of expertise in computer science, statistics or philosophy. From Mount Intervention to Mediation, Pearl succeeds chapter after chapter in giving the reader a new topic to consider. I would not be surprised if *The Book of Why* becomes a regular reading in many Science and Philosophy syllabi. The experience of reading and conceptualising the book's hypotheses is enriched by Pearl's narrative talent; historical events unfold before the reader's eyes, glimpsed through the authors' vivid descriptions as if they had seen them happen in the flesh. In 2019, we can manipulate our environment as was never possible in the past. Reflecting upon causal inferences is therefore more important today than ever before.

Department of Physiology, Anatomy and Genetics

Bibliography

Pearl, J. and Mackenzie, D. 2018. *The Book of Why*. London: Penguin.

Macagno, F. & and Walton, W. 2013. *Emotive Language in Argumentation*. Cambridge: Cambridge University Press.

Bradford-Hill, A. 19645. *The Environment and Disease: Association or Causation*. Proc R Soc Med. https://www.edwardtufte.com/tufte/hill

Francesco Redi, F. 1668. *Esperienze Intorno alla Generazione degli Insetti*. Public Domain.