CHAPTER

1

Setting the Stage for Causal AI

The ability to understand information in the context of solving complex problems is not new. From the earliest days of artificial intelligence, scientists and mathematicians have tried to find new ways to understand the world through models and data. The promise of artificial intelligence (AI) is to reach the point where machines could think and provide answers to some of the most challenging problems of our world. There are a huge number of sophisticated analytics tools that provide significant help in understanding what has occurred in the past and predict a possible future from that data. However, one element that has been missing from the analyses is understanding the cause and effect of the observed and unobserved interactions. The dynamic of understanding why events happen and what can be done to change the outcomes is the power and opportunity of causal AI.

This chapter will put causal AI in perspective and set the stage for our exploration of the evolution of the field of AI.

Why Causality Is a Game Changer

Why is there a sudden explosion in interest in causal AI? The answer is both complex and simple. Causal AI enables us to move beyond the predictive modeling capabilities of traditional AI to understand and predict causal relationships between variables in a system. Here are some of the most salient topics that outline the value of causal AI:

- **Understanding causality:** Traditional AI models can make predictions based on observed correlations between variables but cannot tell us why a particular outcome occurred. Causal AI, on the other hand, can identify the causal relationships between variables and help us understand why a particular outcome occurred. Causality and understanding the dynamics of causality can be particularly important in fields such as healthcare, where understanding the causal relationships between risk factors and health outcomes can help identify new interventions and treatments.
- **Identifying interventions:** Causal AI can help us identify interventions that can change outcomes. For example, causal AI provides a graphical technique to pinpoint the most relevant variables needed to understand specific objective or estimate the consequences of a given intervention. The goal of causal AI is to help an organization assess the possible cause and effects of various policy actions. Causal AI has the potential to enable a team to have a common understanding of a problem so that they can work together to determine why a situation has occurred and establish a plan to arrive at the best next actions.

- **Predicting counterfactuals:** Causal AI can predict the effect of a particular variable on an outcome of interest in an alternative scenario. This is especially useful when the variable of interest is not directly observable or measurable, as it allows the estimation of the causal effect on the outcome. For example, it can help predict what would have happened if a particular intervention or policy had not been implemented.
- **Avoiding bias:** Traditional AI models can be biased if they are trained on biased data or if they do not account for all relevant variables. Causal AI, on the other hand, can help avoid bias by identifying and accounting for all the relevant variables in a system. This can help ensure that the predictions and decisions made using AI are fair and unbiased.
- **Improving decision-making:** Causal AI can help make better decisions by providing a deeper understanding of the causal relationships between variables. This can be particularly useful in fields like business, where understanding the causal relationships between different variables can help businesses make more informed decisions and achieve better outcomes. Causal AI provides us with a deeper understanding of the causal relationships between variables in a system and can help us identify interventions, predict alternative choices and actions, avoid bias, and make better decisions.

The next generation of artificial intelligence can benefit from a deeper level of collaboration between data experts, business leaders, and subject-matter experts. While AI has long been used to solve complex problems, in this new era of expanded AI, hybrid teams of business and analytics professionals can include an examination of why problems happen and what alternate approaches can help a business move forward to gain a sustainable and measurable advantage when faced with increasingly sophisticated and aggressive competition. Therefore, we are in an interesting and complicated transition in the evolution of artificial intelligence. While the focus of many traditional AI approaches focuses on data and feature engineering, there is a revolutionary trend emerging. Causal AI uses causal inference as the underlying math of cause and effect. The focus of causal AI is on business outcomes. Causal inference is the science of why, as explained so well by Judea Pearl in *The Book of Why*: "The ideal technology causal inference strives to emulate resides within our own minds. Some tens of thousands of years ago, humans began to realize that certain things cause other things and that tinkering with the former can change the latter."[1]

While it is time-consuming and challenging to examine all the possible answers, we can easily ask why events happen and what are the primary the cause-and-effect factors of a problem we are trying to solve.

We have many years of experience working with business and technology leaders who are grappling with some of the most complex problems that our current and traditional technologies are designed to solve. We have seen hundreds of emerging technologies come and go that promise to turn human knowledge into packaged solutions that are easy to understand and implement. The history of technology has proven repeatedly that there are no simple solutions to complex problems. However, each technology takes us a step further to addressing issues. For the past 10 years the focus of AI and advanced analytics has been on analyzing massive amounts of data to understand the answers to difficult problems. Big data was the silver bullet that offered some success but did not go far enough.

One of the biggest stumbling blocks to making AI and advanced analytics solutions work effectively is the complexity of the underlying technologies. Typically, business leaders want to be able to visualize the outcomes from the data buried inside applications and from both internal and external data sources. Business managers and leaders want to not only understand what the data tells them about their current situations but what actions they can take to protect and advance their future goals and objectives. Business leaders look to data scientists who employ statistical and computational techniques to determine insights from big data. Many data scientists use correlation and machine learning techniques to identify patterns and anomalies to predict outcomes. Increasingly, business leaders are beginning to understand that there is tremendous potential to leverage AI to solve complex business problems. The greatest potential for AI is to create a way to abstract the complexity from the underlying technology so that data scientists, subject-matter experts, data specialists, and business leaders can collaborate to solve business problems. Therefore, one of our goals with this book is to bridge the gap between the data scientist and the business leader so that it opens the door to use the power of causal AI and traditional AI to create a competitive advantage.

However, there are no silver bullets or simple answer. Many business and technology leaders are either adopting or evaluating AI-based solutions to automate processes and determine how to improve their businesses. The promise of AI is tantalizing—organizations can use algorithms to analyze their data in context to anticipate changes in customer requirements and prepare for the future.

In competitive markets, it is imperative to be able to understand what and why situations are happening within the business. How can leadership within a business ensure that revenue can grow? When looking into the future, organizations need to be able to understand the impact of the decisions they make. What happens if a product price is reduced by 10 percent? Will a lower price cause more customers to buy? Will the price increase entice more new customers to buy? If revenue suddenly decreases, does management understand why this has happened and what can be done to change the current course of business? Are customers leaving because of a quality issue triggered because the business began using a new supplier or because of an emerging competitor? Understanding the cause and effect from processes and data is one of the primary reasons why causal AI is emerging as such a critically importatnt approach across the fields of academia and business.

The most common techniques that have been used by data scientists are correlation-based techniques that are common in the field of traditonal AI. While correlation and causality are related approaches, they are not the same. In brief, correlation is a technique for establishing the statistical relationship between variables. In contrast, causality refers to how one variable has an impact on other variables. In the case of causality, one variable might have a direct impact on a second variable, there could be an indirect effect, or there could be a confounding effect. Causal AI is the art and science of understanding the myriad of relationships between variables that drive relevant causes and effects in a system that we are seeking to understand and manage.

Understanding the difference between correlation-based statistical analysis and causality-based analytics is key to being able to employ and deploy the power of causal AI. Therefore, in the next section we will explore the broad area of analytics. Applying a causal AI approach to analytics will help guide organizations to focus on the assumptions and knowledge that we have about how the world works. If we can answer the complex questions about why an issue occurs, we can adapt our approach to solve problems. The goal of causal AI is to be able to understand the story of the data.

Causal AI in Perspective with Analytics

Analytics is one of the most widely used, and often misused, terms today. The discussion of analytics is widespread. The term *analytics* often refers to dashboards and historical reports. In addition, analytics includes collections of data, information, applications, and analytical models related to work with descriptive statistics. Analytics encompasses work products resulting from predictive, prescriptive, simulation, and optimization projects and programs.

There are many different perceptions and assumptions about what it means to conduct an analytics project. One group may assume a focus on a historical dashboard, while others creating a simulation model might work in the operational area of the business. To make matters worse, there is likely to be a different vision for how to approach analytics. Typically, organizations and departments will be trying to solve very different problems depending both on the problems they need to address and on the stage of their projects. Approaching analytics in the context of causal AI requires a common understanding of the types and approaches to analytics.

At a conceptual level, correlation-based AI and causal-based AI approaches have the same roots; they come from the same family/branch/category of advanced analytics. So, before we move forward with our discussion and description of causal AI, let's define the broader term *analytics* to ensure that we have a common understanding as we move forward in our dialogue.

Analytics is one of the most widely used, and often misused, terms today. The discussion of analytics is widespread. Routinely, we talk about analytics with academics, researchers, government officials, university administrators, scientists, business executives, subject-matter experts, data scientists, and more.

The term *analytics* is employed when referring to dashboards and historical reports. Also, analytics is used when referring to collections of data, information, applications, and analytical models related to work in and with descriptive statistics. And the term is deployed when referring to work related to and the effort/products resulting from predictive, prescriptive, simulation, and optimization projects and programs.

So, it is not surprising to see confused expressions on people's faces when someone talks about analytics. One person is talking about a historical dashboard, while others are trying to convey how a simulation model might work in the same operational area of the business. It can be frustrating for people not to be able to connect on a topic in which they are all deeply interested.

We have found that beginning a discussion by defining the analytical maturity model or framework to be used in the ongoing dialogue helps in reducing confusion and ensures that everyone in the conversation has a clearer understanding about what is being said or what others are trying to convey given the wide range of definitions individuals hold in their minds about the term and areas of analytics.

Analytical Sophistication Model

Various analytical maturity models have been developed, including one by the staff at Gartner. Using an approach similar to Gartner's, we have created an extended sophistication model to clarify and further define *analytics enablers*, *analytics*, and *advanced analytics* (see Figure 1.1).

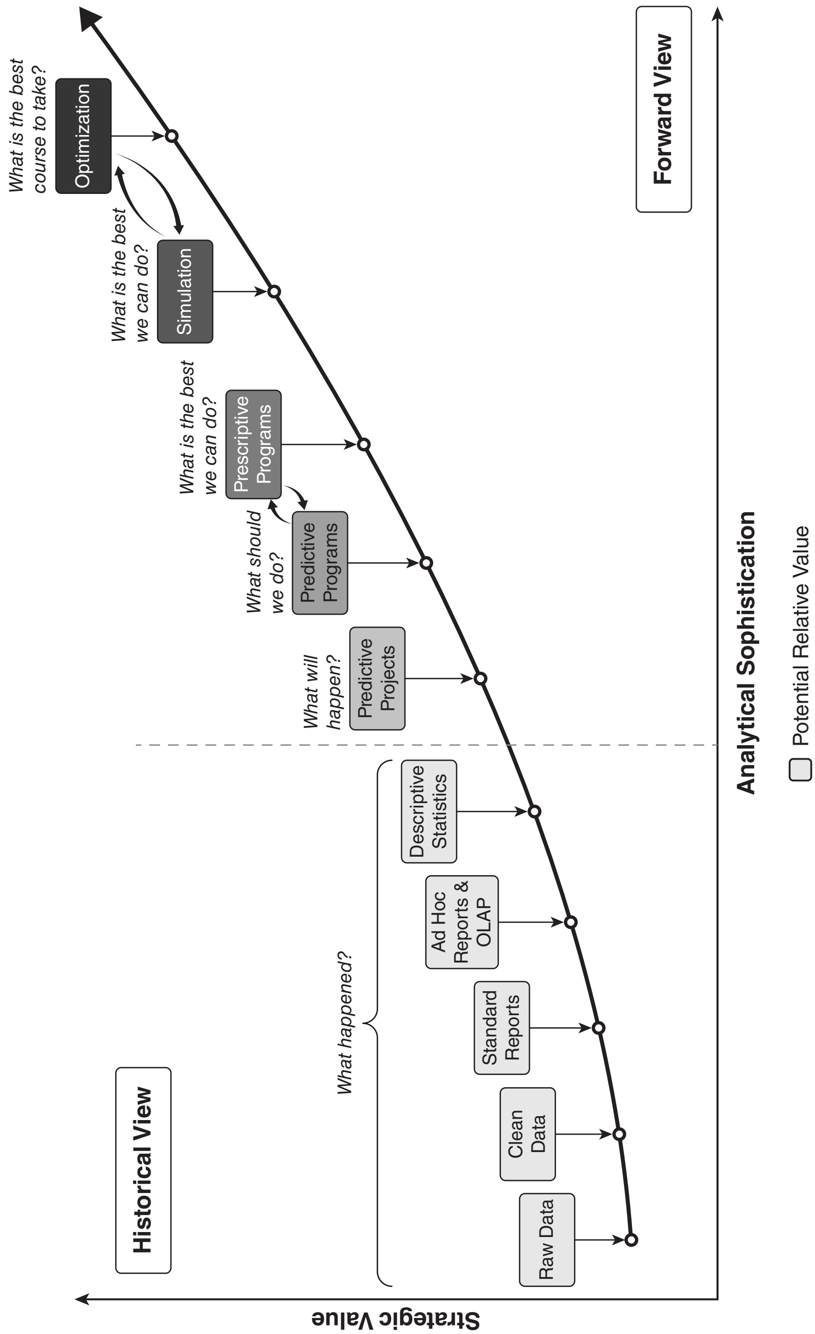

FIGURE 1.1 Analytical sophistication model

Analytics Enablers

In this model, the initial categories to the left of Descriptive Statistics (e.g., Raw Data, Clean Data, Standard Reports, and Ad Hoc Reports & OLAP) are related to data, dashboards, and historical reporting. These categories are not included in our definition of analytics. These categories are enablers of analytics but not analytics in and of themselves.

Analytics

In Figure 1.1, analytics begins with Descriptive Statistics. Analytics, as we are defining the term in this model, extends to, and ends at, the vertical dotted line encompassing Predictive Projects. Descriptive statistics provide insights into historical and current data and are valuable and useful tools in analytics projects and work. When an analytics team creates an exploratory data analysis (EDA) to investigate and gain a deeper understanding of a problem space, descriptive statistics play an integral role in examining and outlining how a business scenario operates in the real world and how that operation is described through data. An EDA and the descriptive statistics used provide a view into the factors that are at work and the relationship of those factors.

Analytics includes prescriptive projects as well. Organizations begin to experiment with their journey into more sophisticated analytics via prescriptive projects. These projects may illustrate that an organization has an appetite for advanced analytics, and it may show that this type of work is too costly, complicated, and difficult for the organization to embark upon as a sustained activity.

Advanced Analytics

Beginning at the vertical dotted line in the diagram and continuing to the right, all of the subsequent areas (i.e., predictive programs, prescriptive programs, simulation, and optimization) are encompassed in the category of advanced analytics. For the purposes of our discussion, all of AI, including both families of correlation and causal AI, exists in the portion of the analytical maturity model that is to the right of the vertical dotted line. In our discussion, advanced analytics and AI are synonymous.

Scope of Services to Support Causal AI

Causal AI must be viewed as part of the overall computing infrastructure for businesses. Therefore, it needs to be a consideration in the approaches to managing data and cloud services.

The hybrid cloud is a critical asset when moving to causal AI. One can argue that analytics can be managed locally. However, in the complex world of advanced analytics and more specifically causal AI, you have to assume that relevant data will come from a variety of sources—some will be third-party data sources managed in a public or private cloud. Other data may be generated by the Internet of Things (IoT) devices at the edge of the network. Creating a federated data approach that takes into account the fact that there are myriad data sources needed is critical to the success of advanced analytics in the context of causal AI.

In any discussion of advanced analytics, the scope of the data is critical to successfully approaching causal AI. For the past decade, data scientists have assumed that if you can collect enough data to create a model, you can make well-informed decisions.

However, one of the distinctions between a data-first approach and a model-first approach is the ability to anticipate the type and scope of the data needed to understand relationships between variables and the data that supports establishing causal relationships. Being successful with causal AI requires the ability to discover the relevant data whether it is internal or external. Some of these data sources may be extremely large, and it would be impractical to move that much data to an internal data center.

Some of the early experiments with causal AI have failed because the business simply did not collect enough of the right data to answer why a problem occurred and what approaches will help to understand appropriate next actions. The data that you will need to create a causal AI solution will be varied. For example, there is considerable data that may be stored in data lakes. There will be data from systems of record as well as third-party data sources that are relevant to market trends. In some situations, your organization may have correct information, but there is simply not a large enough corpus of data to make relevant decisions. If an analytics team using a small data set draws conclusions from that data, the answers may be incorrect.

Keep in mind that building a causal AI model and solution is an iterative process requiring incrementally adding data that corresponds to variables and relationships modeled. For example, what data might you discover that is relevant to understanding a complex problem? If you were going to try to determine (in an agricultural setting) the optimal time to plant crops, you would need a variety of data ranging from the chemical construct of the soil, the changing weather patterns, and the past planting history. You might also need to understand what crops are selling better in a changing economic environment. If one of these factors is ignored, the resulting analysis will be useless. Data discovery and incrementally adding newly discovered data is critical to building effective causal AI solutions.

The hope has been that if we are able to analyze this data, we will be able to understand our world and understand how to transform our businesses. However, harnessing and making sense of data is not simple. As with any emerging technology innovation, moving from the idea of managing data to the reality of discovering solutions to complex problems is harder than anyone could have imagined.

There are many practical techniques for analyzing massive amounts of data that has helped organizations in many ways. Tools are available that help bring together many different types of structured and unstructured data to better understand the meaning of information, which has been a tremendous help to businesses. Data warehouses and data marts, for example, have enabled organizations to analyze business data to help make decisions, such as tracking transactions and managing operational data.

The Value of the Hybrid Team

While taking advantage of a team of professionals when successfully managing a data-focused project is critical to any project, there are requirements when approaching causal AI. We will cover the process of collaboration in a causal AI project in Chapter 5. Many early AI projects have failed when data scientists work in isolation from representatives from different areas of the business. It is particularly important in causal AI to involve subject-matter experts very early in the design process. These professionals understand the content and details of the operation that will become instrumental in creating a model that matches the problem as well as having hands-on experience with the most important data sources. Business strategists will help focus the team on the most important business problems that the company needs to address.

The Promise of AI

While AI was conceived of more than 60 years ago, it has been slow to realize true benefits for industry and business. The promise of artificial intelligence for decades has been the ability of machines to mimic the capacity of the human brain to understand the context of data and make quick decisions even when all the data you need is not available. One aspect of cognition that sets humans apart from machines is that humans typically maintain a mental model of our environment. Creating mental and computing models is one of the key processes to transforming data into a meaningful understanding of phenomena that we seek to explain or understand. A goal of AI is to be able to predict the behavior of individuals. It is no wonder that the best minds in artificial intelligence haven't been able to build reliable and accurate models of similar complexity to those that are inherent in the human brain.

It is now possible to develop sophisticated models; deploy those models in a range of operational environments; and treat, integrate, and ingest numerous types of data to produce reliable, scalable, and accurate results. However, none of these approaches go far enough to understand why events happen and how predictions are made. The value of causal AI is that it provides a technique for helping organizations to leverage the power of AI and advanced analytics so that organizations understand why events happen and how to come up with better solutions.

Traditionally, AI has been thought of as one large category. However, the reality is that to be successful, there needs to be a more nuanced and multifaceted approach to AI. Commercial firms, regulators, and academic organizations—to name a few—are interested in and eager to work with a new class of AI. This new category of AI needs to provide for transparency in data transformation, model development, and, most importantly, the interpretation of the results produced by AI systems. This is one of the primary reasons why we see that causal AI and causal inference are moving out of research labs and the halls of academia into the commercial market.

Almost since the beginning of time, humans have been searching for the reasons why things happen. Once statistics became associated with correlation, there was an effort to disregard the power of causality. However, with the requirement to understand why an event has happened, what it means, and how to change outcomes—causality has emerged as a transformative trend.

As we develop a broader and deeper view of causality throughout the book, we will provide definitions and employ a business-oriented use case so that data scientists, subject-matter experts, data analysts, managers, and business leaders can have a common vocabulary to understand causal AI and how it is important to implement business solutions that will help the organization move forward.

Understanding the Core Concepts of Causal AI

As we move through the chapters in this book, we will get into more and more detail about how causal AI works and the elements that are important to understand. In this next section, we will provide an overview of the key concepts that are required to understand causal AI.

Explainability and Bias Detection

As AI has evolved over the last several decades, models have become more sophisticated, and, in some cases, multiple models are implemented in an analytics pipeline where the models interact and influence the results of the overall pipeline and the individual models. Monitoring and understanding multiple models is challenging and can be difficult to discern the influence of any one model on the overall results or outcomes.

In addition to ensembles of models running simultaneously and, in some cases, in parallel, certain analytical techniques in traditional AI rely on algorithms that are very difficult to examine, discern, and understand why certain results or outcomes are produced. Algorithms such as neural networks belong to a class of analytical technical that are referred to as *black boxes*. These class models have been given this label because the designs of the resulting models and processes running within the models are so complicated that it has been nearly impossible for humans to understand how the input variables are combined to reach the ultimate conclusions.

Academics, researchers, and commercial technology companies have been working to develop solutions to provide explainability solutions for neural networks. Over the past 5 years, we have seen solid progress made on providing clear explanations of how neural networks process data and why those models generate the outcomes and results that are seen. It will be a few years before we see and can use explainability technology that will be acceptable to government regulators and governance professionals.

A black-box model executes its analytical process based on how it interprets the data. The black-box model will not be able to provide business managers with an explanation about how the model makes decisions as to why a campaign did not work. As a result, data business managers cannot explain to shareholders or government regulators why certain business decisions were made. This traditional approach to AI lacks two important factors that are imperative for business success: explainability and fairness or identification of bias.

Explainability It seems logical, and it is mandated by national laws and regulations around the world, that if you are using machine learning techniques to help your organization make decisions based on analyzing data, you should be able to defend the results of your analysis. One of the problems with many approaches to deep learning is that they are built with complex neural networks that have hidden logic that humans find difficult to understand. This becomes a problem when a model is being used to make decisions that impact the fairness of actions taken by an organization such as decisions about who to hire or when a policy discriminates against a specific group of people. To address this issue, explainable Artificial Intelligence (XAI) is a set of processes and methods that allows human users to comprehend and trust the results and output created by machine learning algorithms. XAI is used to describe an AI model, its expected impact, and its potential biases.

Moving from a black-box model to an explainable model requires that the team begin by modeling to define and understand why the problem exists and how it might be addressed (see Figure 1.2). For example, why are we losing customers suddenly? What approach to marketing a product will produce the best results? What is the right dose of a medication for children under the age of 10? When an organization begins by defining the problem, it is much easier to select the variables that are needed to understand a problem before applying data. This method is designed to ensure that a model can be explained and understood.

Detecting Bias in a Model One of the risks of any technique or approach to modeling that is difficult to examine and explain is that it is difficult to detect and correct for bias. The consequences of these risks can cause an organization to be sanctioned or fined by regulatory authorities and to face a backlash from customers and activists. Bias in data can happen when an organization selects data that has not been reviewed and examined for inherent, historical patterns that may subtly but inadvertently perpetuate past inequities. Bias can creep in when data is selected that is out of step with the problem being addressed or not aligned with the current expectations of management, government regulators, and the public. For example, the data may include information about the age of individuals in the data set that can introduce bias. Therefore, in a black-box model, it will be impossible to see that the results are biased against either younger or older individuals. Imagine that a model is used to help an organization select the best candidates to hire for a job. If the model includes data about the candidate's race, age, or sex, the selection process might be inadvertently biased.

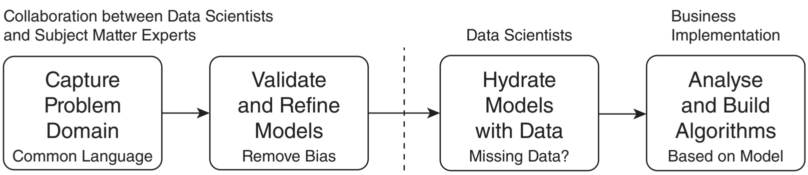

FIGURE 1.2 The collaborative process begins by articulating the problem being addressed before data is added to analyze the business problem being addressed before adding data to the model and analyzing the results.

Directed Acyclic Graphs

A directed acyclic graph (DAG) is the most common type of graphical modeling approach used to illustrate causal relationships between variables in a causal model. In a DAG, a solid forward-facing arrow indicates a subsequent variable is directly affected by the previous variable. A DAG provides process guidance as an intuitive way to represent causal relationships between variables. A DAG can represent either a direct or indirect effect, which means the sequencing of the model processing through the nodes can move only in a forward direction; there are no recurrent looping cycles. Just as in the real world, time travels only in a forward motion.

Let's examine a few examples of how variables impact and interact with each other. In this example of a causal model, one variable is warm weather; a second variable is the consumption of ice cream; and the third variable is drownings. Clearly, people eat more ice cream in warm weather, so there is a direct forward relationship between warm weather and ice cream. Likewise, there is a direct relationship between warm weather and more people drowning because people will swim in warm weather. If, in our casual models, we create a backward relationship or loop that indicates that the variable ice cream is the cause of warm weather, which is clearly wrong and misleading, and impossible, we have created a model that is invalid and is not representative of the real world and the causal relationships that naturally occur in daily life. To continue with our example, if we created a backward loop between the variable of drowning and the variable of ice cream, we would also have created an illogical conclusion. Clearly, this is a simple example that would be obvious to refute. In more complex situations, however, it will take work to understand the complex relationships and causal strengths of relationships between variables.

Structural Causal Model

A structural causal model (SCM) is a type of DAG that is designed as a mathematical framework. Once a DAG is created as a graphical model, the SCM process translates each variable or node into a mathematical equation. For example, in a complex supply chain, the SCM calculates the amount of time required to deliver a part delivered by the supplier to the customer. Therefore, moving from the graphical model of the DAG to the more precise mathematics of the SCM requires a definition of the three common types of variables. These variables include explanatory, outcome, and unobserved. In other words, there are variables that simply explain what a variable is (i.e., a metal component needed by a customer in the manufacturing process). An outcome variable is intended to evaluate the effectiveness of a process defined in the graphical DAG. The outcome variable can be calculated to determine if there is a bottleneck in a process that could impact a supply chain. As the name implies, an unobserved variable is one that is not easily identified in the DAG. For example, it may appear that everything is working well with a supply chain. However, customers are leaving the company. The unobserved variable might be lack of trust in the company or a new competitor who is doing a better job in building customer relationships. SCMs are powerful because they both indicate causal relationships between variables and can make predictions about how these relationships would change if conditions changed.

Observed and Unobserved Variables

An observed variable is a variable that can be directly measured in an experiment such as the relationship between smoking and lung cancer. If someone has lung cancer, the doctor may ask the patient if they are a smoker. If they are not a smoker, it is likely they have a form of lung cancer that isn't associated with smoking, and the treatment will be different. In contrast, an unobserved variable cannot be seen from observing a patient, for example. There are diseases such as heart disease that can be hereditary. Simply asking a patient about their lifestyle or observing their fitness level would not indicate to the doctor if the patient has underlying risks of getting lung cancer.

Counterfactuals

Merriam-Webster defines the term *counterfactual* as "contrary to fact" and "relating to or expressing what has not happened or is not the case." An example of a counterfactual conditional statement is "If kangaroos had no tails, they would topple over."

Counterfactuals are the alternate scenarios that are ingested into our causal models. Counterfactuals enable causal models to illustrate the wide range of options and alternate paths that we might want to consider. Therefore, counterfactuals are a core concept in causal AI because they help the team understand what would happen if a situation were different. Counterfactuals refer to a hypothetical scenario where the team can ask questions such as, "What would have happened if a causal relationship between two variables didn't exist?" For example, what if we had two suppliers of our products rather than one? Would we have had more sales of our new product? Would we have had better customer satisfaction if we had been able to ship products in a week rather than in three weeks? In other words, counterfactual analysis would have let us imagine a different outcome. Being able to use statistical and various machine learning algorithms to estimate causal relationships can allow the team to anticipate the impact of change in business models or actions.

Confounders

In a causal graph, a *confounding* variable is a variable that influences multiple variables and can lead to a spurious or false association.

In a causal AI system, a variable will be classified as either a cause or an effect. A cause, as the name implies, is a variable that directly influences another variable. In contrast, an effect variable is influenced by one or several variables. However, a confounder variable occurs in situations where a variable can be both a causal and an effect. Confounders are extremely important to identify in causal AI and make it difficult to identify the causal relationships between variables.

Having a confounder (also called an *observed covariate*) will likely distort the results of a model and can bias the results of the analysis. A common example would be the fact that lung cancer is caused by smoking and many low-income patients have lung cancer. Would we conclude that having less money means that a patient will have lung cancer? Of course not. However, if you look at this confounder and use it as fact, then you will likely add bias to your modeling. When creating a model, it is imperative to use statistical models such as regression analysis to remove this bias.

Colliders

The existence of a collider in a causal AI model often leads to counterintuitive and inaccurate conclusions. Colliders are also known as backdoor paths. A *collider* is a variable that has pointers to several different variables in a causal diagram or DAG. This causes confusion in understanding the actual relationships between variables. If a variable is not the cause of two conflicting variables, there is a disconnect that can lead to an erroneous conclusion and bias. Therefore, a collider is a variable in the DAG that is impacted by several different variables. These variables are not directly related to each other; therefore, these variables might lead to inaccurate conclusions. If we remember the situation where warm weather leads to an increased consumption of ice cream as well as more drowning deaths, we can't draw the conclusion that ice cream causes people to drown. This would be a collider if we tried to draw a direction relationship between the variable representing ice cream consumption and drowning.

Front-Door and Backdoor Paths

In a causal model, it is important to indicate front-door and backdoor paths to understand the cause and effect of a problem. For example, if a company wanted to determine the causal effect of customer satisfaction on sales, a typical question might be "Will increasing our customer satisfaction increase our sales?" The causal diagram would be drawn as shown in Figure 1.3.

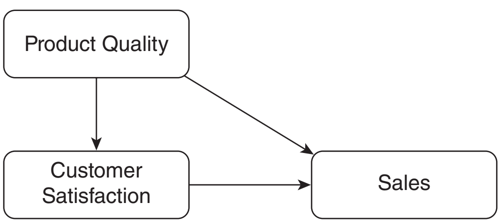

FIGURE 1.3 Causal model demonstrating the relationships between Product Quality, Customer Satisfaction, and Sales

This causal model demonstrates the different types of paths that can exist within a DAG. There is a direct causal relationship between Customer Satisfaction and Sales, which results in a front-door path between the two variables. Causal analysis would be straightforward if not for the existence of a backdoor path: Product Quality is a common cause of both Customer Satisfaction and Sales, which can be considered a confounder in this specific context. Front-door and backdoor paths are sometimes presented in a written format, as shown here:

Front-door path: Customer Satisfaction -> Sales

Backdoor path: Customer Satisfaction <- Product Quality -> Sales

The conclusion that can be drawn from our causal analysis is that to accurately determine the causal effect of customer satisfaction on sales, the company would need to control for product quality as part of their analysis. This would close the backdoor path present in the causal model and eliminate the confounding source.

Realistically, it is likely that there would be multiple other variables that would cause both customer satisfaction and sales, such as price, giving rise to more backdoor paths that would need to be closed in a similar fashion. One of the main goals of a DAG is to help identify backdoor paths that confound the relationships between variables.

Correlation

Correlation is the widely accepted statistical method for understanding the relationship between data sets. Typically, correlation assumes that you can use past outcomes to estimate future outcomes. While this may seem logical, there are some fallacies in correlation. The downside of correlation is that it only shows you results based on data but does not explain the cause of the results from that data. Typically, correlation relies on finding patterns and anomalies in data. In theory, this makes perfect sense. By applying analytics of data with machine learning techniques, correlation does a better job of using these patterns to predict outcomes. Causality, on the other hand, is designed to model relationships between a cause and an effect found in data. Unlike correlation, causal AI understands what happens and why.

Causal Libraries and Tools

Causal libraries are software tools packaged to perform key tasks in a causal AI solution. Because these methods are open-source, they provide developers with well-designed frameworks to apply methods without having to resort to proprietary approaches. These libraries abstract the complexity for leveraging causal methods such as identifying causal relationships and analyzing and visualizing data. These libraries are very important in creating DAGs. One of the most popular libraries is PyWhy, which is a Python-based library that includes an abstracted interface to a variety of existing inferences (causal methods). This library was jointly developed by Amazon and Microsoft. There are many other libraries and tools that aid developers in creating solutions. In Chapter 7 we will provide a detailed examination of the tools and libraries available to support causal AI applications.

Propensity Score

A propensity score is used to reduce bias and the effect of an observed confounding variables. The propensity score is a statistical tool that indicates differences in outcomes based on approaches to problem solving. For example, a propensity score will help understand the effect of different groups of people receiving treatment with a medication based on observed characteristics such as age or medical history.

Augmented Intelligence and Causal AI

As Judea Pearl in the introduction to *The Book of Why* stated, "Data does not understand causes and effects; humans do. . . . In the age of computers, this new understanding also brings with it the prospect of amplifying our innate abilities so that we can make better sense of data, be it big or small." The power of causal AI is not just its ability to model and analyze data but in its approach to augmenting the power of the human brain to understand the world.

Despite the speculation about the ability of machines to think, they cannot. While machines are adept at ingesting massive amounts of data, they do not have the essential causal powers of the human brain. The answer to the proposition that machines can mimic and replicate the power of human cognition is augmented intelligence. Only humans can understand the context of data to make intelligent decisions. The ability to model a business problem based on the variables that define it is a human and more importantly team process. Could a machine replace the process of modeling a problem and then understanding the counterfactuals that will change outcomes? The only way this would be possible is to have the AI environment ingest all of the data and context needed to understand and solve a problem. Such a machine would be able to anticipate changes and new data needed as circumstances change.

In reality, humans have to identify key variables and relationships based on an understanding of the issues. Data specialists must identify sources and determine if they are accurate or biased. In the context of causal AI, augmented intelligence can play an important role in helping humans identify and understand causal relationships. Humans have an ability to understand changes in the world and make decisions based on what they see and experience. It is human to ask the question, "What if things were different?"

Summary

Causal AI is the implementation of solutions based on causal inference. The power of causality comes from its ability to model a business problem based on the collaboration between team members who understand the business, the data, and the strategy. In this chapter, we provided the foundational information about what is and how it is a critical element to determining how to take actions to improve future outcomes.

In many ways, causal AI is not new. Rather, it takes advantage of the foundations of data management innovations over the past decades. It also leverages the tools and services of AI and machine learning techniques. The power of causal AI is that it goes beyond predicting outcomes and instead can identify the underlying causes of an outcome. Augmenting the power of human intuition and knowledge with the overwhelming power of cloud scalability is the benefit of causal AI.

Note

1. Judea Pearl and Dana Mackenzie. *The Book of Why*, Basic Books, 2018.