causaLens

CLX-TR-004

Comparative Analysis of Bayesian Networks and Structural Causal Models

The power of science is its discovery of causal law.

Bertrand Russell [Rus48].

This paper is part of our tech reports series. For more papers in the series see causalens.com/research. In our tech reports we share some of the most significant concepts in the causality literature and we point our readers to the most important publicly accessible academic references. The topics are chosen based on our frequent discussions with data science practitioners, and the reports are written primarily for a technical audience. While much of our original research forms part of our proprietary technology which we do not publicly disclose, these reports are part of our commitment to contribute to the wider research community and share the benefits of Causal AI.

Introduction

Bayesian networks and Structural Causal Models are often conflated, however the two modeling approaches are very different. Bayesian networks are purely statistical models that summarize probabilistic relationships between variables. Structural Causal Models represent the underlying causal processes that are responsible for these probabilistic relationships. This gives Causal Models superior functionality, that includes the capability to design optimal interventions and to reason about counterfactual scenarios. This report unpacks the ways in which Structural Causal Models outrun standard Bayesian networks.

Bayesian Reasoning

Bayesian reasoning supports quantification of belief hence expressing an epistemic state of an agent [Pea88]. Let S and T be two events, with S preceding T. Bayes' rule postulates that

$$P(S|T)P(T) = P(T|S)P(S), \tag{1}$$

where $P(T|S)$ is a probability of T happening given S, and conversely, $P(S|T)$ is the likelihood that S happened if we know that T happened. $P(T)$ ($P(S)$) is an a-priori probability, or likelihood, of T (S) happening.

Bayes' rule is used to compute *forward probabilities*, that is, to predict the probability of T occurring given that we observed S, as well as *backward likelihoods*, that is, when observing T, to estimate how likely it is that S had happened.

Bayesian inference is a method of statistical inference in which Bayes' theorem is used to update the probability for a hypothesis as more evidence or information becomes available. It is an important technique in statistics, and especially in mathematical statistics. Bayesian "backwards updating" is particularly important in the dynamic analysis of a sequence of data. Bayesian inference has found application in a wide range of activities, including science, engineering, philosophy, medicine, sport, and law.

In the philosophy of decision theory, Bayesian inference is closely related to subjective probability, often called "Bayesian probability". In this setting, the probability is interpreted as our belief or knowledge about the events, and hence forward and backward probabilities represent the estimated likelihood of events based on the epistemic state of an agent.

Bayesian networks (BN) are *directed acyclic graphs* (DAGs) representing our belief about the dependencies between the variables. The arrows are meant to represent dependencies based on the statistical correlations between the variables. The direction of the dependency is defined based on the epistemic state (knowledge or belief) of an agent constructing the network; interpreted as such, these connections are in particular not necessarily transitive, that is, if $A \to B$, and $B \to C$, it does not imply that $A \to C$.

Interpretation and Applications

Bayesian networks are useful for structuring and concisely representing information, and are often used in the legal system to support inquest and inquiry. The key difference between Bayesian and standard statistical approach is that in Bayesian reasoning, probabilities are interpreted as *degrees of belief*, expressing an epistemic state of an agent, rather than probabilities in the traditional sense. In particular, a variable whose value is unknown is represented as a random variable in a Bayesian network.

This approach, of interpreting probabilities as a knowledge or belief, allows to combine the values of variables as weighted probabilities, hence paving the way to multi-layered Bayesian networks. At the same time, the connections in a Bayesian network reflect the belief or knowledge of an agent (i.e. an expert constructing the network), hence there is a significant amount of manual effort in constructing the network.

In practical applications, Bayesian networks are typically used to present unstructured and imprecise data and a human expert's knowledge and opinion graphically. Bayesian networks do not consider hidden confounders, and the probability distributions are typically considered to be identical for all involved variables, since they represent the expert's confidence, rather than true probabilities.

As the Bayesian networks approach represents belief and reasoning about belief, they are inherently subjective. In other words, different experts can construct different Bayesian networks given the same events. While this can be seen as a disadvantage, as the results are not guaranteed to be objectively correct, Bayesian networks offer a way to incorporate domain knowledge and opinions into the reasoning, structuring the reasoning that is otherwise represented as a text. Bayesian networks are used in legal applications, most notably in the work of A.P. Dawid [DMV07], who is a leading proponent of Bayesian reasoning and Bayesian statistics. While imprecise, Bayesian networks offer a way to organize and present evidence in an easy-to-understand graphical form.

Another prominent area of application of Bayesian networks is biology and neuroscience [BL14]. As biological data is inherently complex and hard to classify, the flexibility of Bayesian networks offers a route to presenting complex and not completely understood dependencies between variables in a semi-formal graphical way. We note that inference done with Bayesian networks does not result in computation of probabilities per se, but rather composite beliefs, and hence cannot be verified or disproved by the evidence. Finally, a relatively new area of application for Bayesian networks is expert systems, where there is a need to represent expert opinion in presence of a large number of variables, often precluding exact statistical reasoning [WKB10].

Structural Causal Models

Structural causal models (SCMs) are mathematical constructs that allow us to reason about true causality. An SCM is defined as a set of variables with a domain of each variable and functional equations describing the dependencies between variables. Some salient features of SCMs can be represented as DAGs, where an arrow from A to B indicates a causal link, which is supported by the functional equation for B. The functional equations involved in the causal reasoning are typically omitted from the graph.

do-calculus is an instrument to reason about the outcomes of interventions. The *do*-operator sets the value of a variable X to x, regardless of its value as calculated by the functional equations of the SCM. We can then reason about the values of other variables under this intervention. Essentially, the *do*-operator is a detector of correlation as opposed to true causality, as $P(y|x)$ is different from $P(y|do(x))$ if y causally depends on x and is the same if x and y are merely correlated. We note that since the *do*-operator represents an intervention, $do(x)$ would, in particular, cut the ties between x and its hidden causes, hence eliminating the correlation between y and x if both causally depend on a common hidden confounder. Furthermore, the functional equations allow us to estimate the *causal effects*, that is, the degree of influence of x on the value of y.

Judea Pearl describes the process that led him to propose SCMs as the correct mathematical framework to support reasoning about causality as follows: *"We played around with the possibility of replacing the parents-child relationship $P(X_i|P_i)$ with its functional counterpart $X_i = f_i(P_i, U_i)$ and, suddenly, everything began to fall into place: we finally had a mathematical object to which we could attribute familiar properties of physical mechanisms instead of those slippery epistemic probabilities $P(X_i|P_i)$ with which we had been working so long in the study of Bayesian networks."* [Pea09].

Comparison between Bayesian Networks and Structural Causal Models

The concept of belief propagation and the direction of causality in Bayesian networks are important concepts and are carried over to the domain of structural causal models. However, at their core, Bayesian networks are graphical representations of probability tables (and usually subjective ones) representing correlations between variables and do not offer any deeper understanding or ability to reason about causal dependence. An arrow from A to B means that the epistemic probability of B (which is an agent's confidence in the truth of B) depends on the epistemic probability of A. Conversely, there can be an arrow from B to A as well, as there is no way to determine the direction of causation from correlation. In particular, whereas Bayesian networks superficially represent DAGs, they have no capability to compute the important measures which are needed for decision making.

In contrast, structural causal models (SCMs) can answer *interventional* and *counterfactual* questions. Namely, what happens if we intervene on the structure of the network by changing the value of one of its internal nodes, disregarding the equations, and what would have happened if one of the leaves of the network would have had another value. As a side remark, note that these questions are different: interventions include cutting causal dependencies and replacing calculated values with values that are set artificially; counterfactual questions ask what would be the values of nodes in the network if some of the leaves had different values, without changing the equations. Bayesian networks lack a mechanism for introducing and estimating the effect of interventions and counterfactuals, hence making them unsuitable for precise causal reasoning.

The following simple example demonstrates some of the differences between BNs and SCMs.

Example. *Consider an expert who constructs a Bayesian network of the connection between weather, arson, and forest fires, where weather is measured in degrees, arson has several levels of severity, and the fire variable has several values corresponding to the size of the fire. Forest fires correlate with weather and with arson, and the expert reasonably believes or knows, based on her knowledge of how the world works, that hot weather and arson cause forest fires, and not the other way around. The Bayesian network in the figure on the left captures the knowledge or belief of the expert. We note that, while it captures the dependency of forest fires on hot weather and arson, it does not provide any means to quantify this dependency or to capture the exact formula that ties weather, arson, and forest fires.*

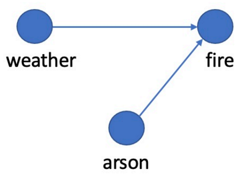

A structural causal model, on the other hand, can provide the means for quantifying the dependency and confirming its causal nature by analysing the existing data and introducing interventions to estimate the causal effect of changes in different variables on the result. Let us assume that the functional equation for the forest fire is

$$\textit{fire} = 1 \textit{ if weather } > 30°C \textit{ and arson } \geq 2; 0 \textit{ otherwise.}$$

Given the typical probability distribution of seasonal temperatures and the frequency of arson, we can calculate the probability distribution for the forest fires.

Moreover, we can ask intervention-related questions; for example, "what if the temperature in this region will rise 5°C degrees on average due to global warming", or "what is the probability distribution of forest fires in this model for regions where the temperature is consistently 10% lower than in the current region (see the figure on the left).

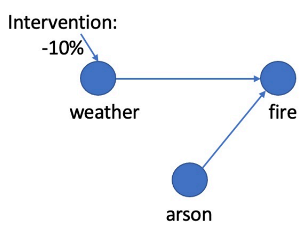

Finally, in the absence of expert knowledge, the Bayesian network can attribute causal dependency to events that have no dependency at all, such as, for example, the frequency *of shark attacks and the consumption of ice-cream. These events have a very strong correlation, with the causal explanation being that both rise during the summer. However, if the expert has no knowledge of the hidden confounder (summer), they may decide that there is an arrow between sharks and ice-cream in the diagram, pointing in either direc-*

tion.

We summarize the differences between Bayesian networks and Structural Causal Models in the following table. While Bayesian networks remain a useful tool for graphically representing experts' knowledge and opinions and are undoubtedly better than writing these opinions in plain unstructured text, it is clear that Bayesian networks are not formal or expressive enough to capture true causality. Prediction of possible outcomes in the current model and under interventions and supporting exact answers to "what if" questions are central to the required functionality of causal reasoning frameworks. Both are supported naturally by Structural Causal Models, making them vastly superior to the Bayesian networks for practical applications. On the philosophical side, Bayesian networks merely express the expert's epistemic state (knowledge and belief), whereas Structural Causal Models capture the true causal connections.

Functionality	**Bayesian networks**	**Structural Causal Models**
Graphical	✓	✓
Expert knowledge	✓	✓
Arrow=causation	✗	✓
Counterfactual reasoning	✗	✓
Interventions	✗	✓
Confounders	✗	✓
Causal effects	✗	✓
Quantitative predictions	✗	✓

About causaLens

causaLens are the pioneers of Causal AI — a giant leap in machine intelligence. We build Causal AI powered products that empower humans to make superior decisions. We are creating a world in which humans can trust machines with the greatest challenges in the economy, society, and healthcare.

causaLens

Bibliography

[Rus48] Bertrand Russell. *Human Knowledge: Its Scope and Limits*. Simon and Schuster, 1948.

[Pea88] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Series in Representation and Reasoning. Morgan Kaufmann, 1988.

[DMV07] A. P. Dawid, J. Mortera, and P. Vicard. "Object-oriented Bayesian networks for complex forensic DNA profiling problems." In: *Forensic Science International* (2007), pp. 195–205.

[Pea09] Judea Pearl. *Causality*. 2nd ed. Cambridge University Press, 2009.

[WKB10] Wim Wiegerinck, Bert Kappen, and Willem Burgers. "Bayesian Networks for Expert Systems: Theory and Practical Applications." In: *Interactive Collaborative Information Systems. Studies in Computational Intelligence* 281 (2010).

[BL14] Concha Bielza and Pedro Larrañaga. "Bayesian networks in neuroscience: a survey." In: *Front. Comput. Neurosci.* (2014).