Bayesian Networks vs. Conditional Trees for Creating Questions for Forecasting Tournaments

Ross Gruetzemacher[1]

[1]W. Frank Barton School of Business, Wichita State University, Wichita, Kansas, USA

Abstract

We evaluate the strengths and weaknesses of Bayesian networks and a specific subset of Bayesian networks—conditional trees—for creating questions for forecasting tournaments, particularly in the context of high uncertainty forecasts such as for technological forecasting or for the forecasting of existential risk. A conditional tree takes the probability of a final outcome and divides it among a sequence of conditioning nodes that would have the maximal impact on the final outcome. This framework can be used to identify the maximally impactful questions to ask in forecasting tournaments through the evaluation of conditional trees using an evidence ratio. However, conditional trees are not without limitations. In this study, we describe applications in which Bayesian networks resolve some of the limitations of conditional trees, but often at additional costs. We further discuss when it is best to use each of the methods, and we make some suggestions for the particular contexts of generating forecasting questions for technological forecasting and the forecasting of existential risk.

1 INTRODUCTION

Second generation forecasting tournaments seek to address challenges that first generation tournaments, like the ACE-style tournaments described by Tetlock and Gardner [2016], did not address. Examples include question generation or expert recruitment [Karger et al., 2022]. Question generation is of particular importance, especially when forecasts concern longer-range topics with inherently high degrees of uncertainty like technology forecasting or existential risk[1].

Karger et al. [2022] have recently proposed conditional trees as an effective means for breaking down complex questions involving extreme outcomes and high uncertainty that are beyond the 12-18 month horizon for which forecasting tournaments have been demonstrated to be maximally effective [Tetlock and Gardner, 2016], e.g., questions regarding existential risk or technological progress. Conditional trees take the probability of a final outcome and divide it among a sequence of conditioning nodes that would have the maximal impact on the final outcome. Conditional trees are a special case of Bayesian networks (BNs) that are equivalent to event trees, and which have applications in risk analysis [Marsh and Bearfield, 2008], physics, policy analysis and biological regulation [Smith and Anderson, 2008]. In such applications, event trees are often elicited instead of BNs due to simplicity [Bearfield and Marsh, 2005].

BNs can also be used forecasting applications [Abramson et al., 1996]. Tetlock has previously proposed full-inference-cycle tournaments as one form of second generation forecasting tournament. These include four phases for scenario generation, question generation, a first generation forecasting tournament and a postmortem. Tetlock has proposed generating question clusters in the second phase, but we believe that this is better-suited for BNs.

Oftentimes BN elicitation is not as straightforward as conditional tree elicitation [Bearfield and Marsh, 2005], and collaborative elicitation is even more entailed. Recent work has described a framework for Bayesian Argumentation via Delphi (BARD) that is suitable for collaborative elicitation of BNs [Nyberg et al., 2021]. While BARD has been demonstrated to be very effective at helping teams improve BN generation over individuals for intelligence analysis tasks [Korb et al., 2020], collaborative elicitation of BNs does not appear to have been used for forecasting applications.

[1]Technological forecasting attempts to forecast future technologies and their impacts to inform decision makers and policy makers. Existential risk forecasting attempts to identify and forecast potential global catastrophic or existential risks for informing policy makers. This could include forecasting artificial intelligence (AI) progress, its impacts and its risk of global catastrophe or human extinction.

Accepted for the 16th Bayesian Modelling Applications Workshop (@UAI2022) (BMAW 2022).

In the context of elicitation for forecasting question generation, little previous work exists—in general and specifically concerning the use of BNs or conditional trees. However, for longer-range forecasting questions, it is more tractable to forecast early warning indicators that have causal relationships with longer-range outcomes. Relationships of this nature are easily represented through BNs, or, through conditional trees.

2 DISCUSSION

Here, we examine the advantages and disadvantages of both BNs and conditional trees in the context of creating forecasting questions for technological forecasting or forecasting existential risk.

One of the foremost advantages of conditional trees is that they are more easily evaluated, using an evidence ratio as a metric [Karger et al., 2022]. In competitive creation processes, this helps to identify which trees are superior. Further, conditional trees are more easily merged with other conditional trees due to structural similarities.

Another advantage of conditional trees identified in previous literature was that their elicitation was straightforward [Bearfield and Marsh, 2005], and, relative to BN elicitation, simple—little to no training is required for collaborative elicitation. BN elicitation, and particularly collaborative elicitation of BNs, does require significant training for non experts [Nyberg et al., 2021]. However, for forecasting applications related to existential risk or tehcnological forecasting—topics that involve significant uncertainty—elicitation requires participants to select the most critical cruxes, a task that can be challenging. Thus, the ease of elicitation that was described for conditional trees in previous work may not hold; because BNs capture more variables with complex interactions they may be better suited for the applications of concern here.

Advantages of BNs include the ability to incorporate more complex relationships between variables, such as variable interactions, that enable the creation of more complex structures for representing the topic of interest. In the context of forecasting, due to the large degrees of inherent uncertainty, more complex structures may be appropriate.

BNs are also better suited for incorporating variables with a large number of states. Conditional trees must represent each state as a branch in the tree structure, thus requiring balancing the number of states a variable can represent with the simplicity of the structure. Including a large number of states is optimal for forecasting applications when questions attempt to elicit a year rather than whether or not an outcome will occur by a given year. However, a large number of states requires the elicitation of very large conditional probability trees. These trees can very quickly grow too large and complex for simple elicitation of probabilities, and this is a challenge that must be addressed if BNs are to be used for such applications.

Conditional trees and BNs each likely have appropriate applications in the context of forecasting. For example, conditional trees may be better suited for forecasting targets with a clear, long-term outcome involving several critical junctures, like existential risk scenarios. Alternatively, BNs may be better suited for forecasting topics like technological progress, where there are numerous more specific forecasting targets over a range of horizons, each of interest with respect to the potential transformative impact of the technology.

Experiments could be conducted to further explore the intuitions described in this discussion. Topics of interest include:

- Do BNs or conditional trees yield better forecasting questions?
- Do BNs or conditional trees yield better forecasts?
- On topics of technological forecasting or existential risk, is elicitation easier with conditional trees or BNs?
- Can BNs and conditional trees be used together in the forecasting question generation process? (For example, see Figure 1.)
- Does the use of BNs or conditional trees for question generation lead subject matter experts to deepen their explanatory models?

Exploring these questions would require running forecasting tournaments incorporating BNs or conditional trees, similar to those discussed in the introduction. To demonstrate the viability of the concept, initial efforts could simply evaluate whether very simple BNs or conditional trees improved forecasts and their rationales. Later, work could build on this to explore the questions above, and to explore the use of group elicitation processes for BNs and conditional trees.

3 CONCLUSION

This extended abstract is simply intended to raise the issue of using BNs, or simplified BNs like conditional trees, for the purpose of forecasting by incorporating them into forecasting tournaments. We feel that these techniques are particularly well-suited for applications where there is extreme uncertainty such as for mid- to long-range forecasts. The examples we discuss in this abstract are that of technological forecasting and existential risk (and specifically, AI). The analysis in this document is preliminary, but we hope that it is able to yield constructive conversations on the application of BNs to this critical area of ongoing research.

Author Contributions

Ross Gruetzemacher is responsible for this work.

Acknowledgements

We would like to thank Philip Tetlock, Josh Rosenberg, Tegan McCaslin and Benjamin Tereick for their comments.

References

Bruce Abramson, John Brown, Ward Edwards, Allan Murphy, and Robert L Winkler. Hailfinder: A bayesian system for forecasting severe weather. *International Journal of Forecasting*, 12(1):57–71, 1996.

George Bearfield and William Marsh. Generalising event trees using bayesian networks with a case study of train derailment. In *International Conference on Computer Safety, Reliability, and Security*, pages 52–66. Springer, 2005.

Ezra Karger, Pavel D Atanasov, and Philip Tetlock. Improving judgments of existential risk: Better forecasts, questions, explanations, policies. *Questions, Explanations, Policies (January 5, 2022)*, 2022.

Kevin B Korb, Erik P Nyberg, Abraham Oshni Alvandi, Shreshth Thakur, Mehmet Ozmen, Yang Li, Ross Pearson, and Ann E Nicholson. Individuals vs. bard: Experimental evaluation of an online system for structured, collaborative bayesian reasoning. *Frontiers in Psychology*, 11: 1054, 2020.

DWR Marsh and G Bearfield. Generalizing event trees using bayesian networks. *Proceedings of the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability*, 222(2):105–114, 2008.

Erik P Nyberg, Ann E Nicholson, Kevin B Korb, Michael Wybrow, Ingrid Zukerman, Steven Mascaro, Shreshth Thakur, Abraham Oshni Alvandi, Jeff Riley, Ross Pearson, et al. Bard: A structured technique for group elicitation of bayesian networks to support analytic reasoning. *Risk Analysis*, 2021.

Jim Q Smith and Paul E Anderson. Conditional independence and chain event graphs. *Artificial Intelligence*, 172 (1):42–68, 2008.

Philip Tetlock. Full-inference-cycle tournaments (an unpublished grant proposal).

Philip E Tetlock and Dan Gardner. *Superforecasting: The art and science of prediction*. Random House, 2016.

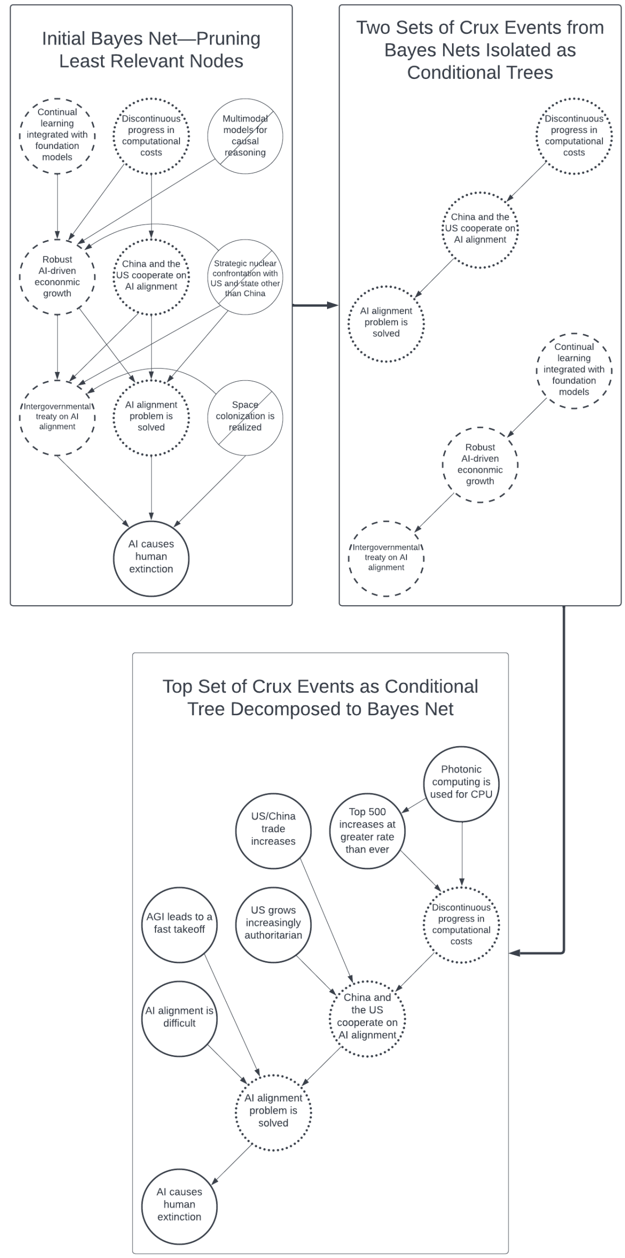

Figure 1: This depicts an example of the potential use of BNs and conditional trees together for forecasting. This depicts a proposed process that starts with a BN, the BN is then reduced to a conditional tree, then the conditional tree is decomposed back into a BN. The new BN has fewer nodes of lesser relevance and overall has a higher information value. BN reasoning through a BN tool could be used with either of the BNs. This is one possible way in which BNs and conditional trees can be used complementarily. This example deals specifically with the development of advanced AI, which is commonly thought to be capable of leading to global catastrophe or even extinction. It is thus a good real-world example of both technological forecasting and the forecasting of existential risk that is described in this document.