In *Working Notes of the AAAI 1994 Fall Symposium Series: Relevance,* page 100, New Orleans, Luisiana, November 4-6, 1994

Moisés Goldszmidt's "Belief-Based Irrelevance and Networks: Toward Faster Algorithms for Prediction"

Discussion by Marek J. Druzdzel[1]

I am pleased with the invitation from Russ Greiner and Devika Subramanian to be the discussant of Moisés Goldszmidt's excellent paper. This short note will, as they suggested, place the paper in a broader context, familiarizing the audience with related work and open problems. Given the size restriction, I will limit my discussion to four topics that I believe to touch both the foundations of the paper and the main theme of the symposium: qualitative reasoning, relevance, conditioning, and causality.

Qualitative Reasoning One of the points that come across well in the paper is that we can perform certain tasks efficiently by sacrificing some precision. Belief-based relevance, based on the order of magnitude foundations of kappa calculus, identifies effectively those interactions in a model that make a large difference while ignoring variables with weak impact. Effectively, the scheme proposed in the paper is of polynomial complexity, contrasting with the NP-hardness of general inference in belief networks. One application of this work, interesting for researchers working in the area of non-monotonic reasoning, is modeling belief revision. Another application, that I find particularly promising, is an approximate algorithm for belief updating, mentioned in the paper. As I have demonstrated recently,[2] the author's implicit assumption that a small number of states of a model will account for most of the total probability mass is justified both theoretically and empirically.

Relevance There is a considerable agreement between the author's view of relevance and what Jaap Suermondt and I present in another paper in this symposium.[3] We both reason about relevant parts of an existing model. One might argue that this is a restrictive view of relevance, asking "Where does the model come from?" and "How do you determine what is relevant to be included in the model in the first place?" I believe that these questions cut both ways. In the process of constructing a model, the reasoner had some set of variables to choose from — why would we not call the body of knowledge that was available to the reasoner a model as well? We agree with Goldszmidt about the importance of the focus of reasoning and conditioning in the definition of relevance. While he is concerned with computation, our work includes subjects as hard to formalize as explanation of reasoning. Consequently, while he can afford a crisp definition, we prefer to leave it slightly imprecise to accommodate various purposes of reasoning about relevance.

Conditioning Dependency networks, on which the author's work is based, are a member of the family of directed probabilistic graphs (with another notable member being Bayesian belief networks[4]). Directed probabilistic graphs satisfy *Markov condition*, also called "screening off," which is informally speaking independence of a node from its non-successors conditioned on its parents. It is useful to realize that Markov condition is quite prevalent in our thinking about the world. We can derive it from a simple assumption that all interactions in a system that we model are deterministic. (As shown elsewhere,[5] this is the case for directed probabilistic graphs, which can be represented by systems of simultaneous deterministic equations.) Markov condition is believed to hold in causal interactions.[4] [6] Because of its "screening off" effect, I believe that conditioning is of vital importance for the concept of relevance. Any definition of relevance is bound to include conditioning, as things may become relevant or cease to be relevant given new observations. I believe that I am here in accord with the author.

Causality There seems to be little doubt about the usefulness of the concept of causality in artificial intelligence: we need causal information in knowledge representation for autonomous agents who reason about the effects of their actions. The task of prediction, for example, addressed by Goldszmidt's paper, requires that the model contain causal information. Causality is also useful in interacting with humans, as it plays an important part in our reasoning. I believe that causality lies at the foundation of the notion of relevance. The concept of a causal mechanism,[7] that has a sound representation in probabilistic graphs,[5] is closely related to relevance. Mechanisms determine how variables in a model will be connected and how they will interact. This provides a basis for constructing models and for reasoning about relevance. Constructing a model of a real world system amounts to capturing formally the mechanisms acting in that system. The most likely place to find variables that are relevant for a given variable are the causal mechanisms in which this variable participates.

---

[1] University of Pittsburgh, Department of Information Science, Pittsburgh, PA 15260, *marek@lis.pitt.edu*

[2] M.J. Druzdzel. Some properties of joint probability distributions. In *Proceedings of the Tenth Annual Conference on Uncertainty in Artificial Intelligence (UAI–94)*, pp 187–194, Seattle, WA, July 1994.

[3] M.J. Druzdzel and H.J. Suermondt. Relevance in probabilistic models: "backyards" in a "small world".

[4] J.Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann Publishers, Inc., San Mateo, CA, 1988.

[5] M.J. Druzdzel and H.A. Simon. Causality in Bayesian belief networks. In *Proceedings of the Ninth Annual Conference on Uncertainty in Artificial Intelligence (UAI–93)*, pp 3–11, Washington, D.C., 1993.

[6] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search.* Springer Verlag, 1993.

[7] H.A. Simon. Causal ordering and identifiability. In W.C. Hood & T.C. Koopmans (Eds.) *Studies in econometric method.* New York: John Wiley & Sons, 1953.