BARD: A structured method and application for collaboratively building Bayesian networks to support human reasoning and decisions

Erik P. Nyberg[1] **Ann E. Nicholson**[1] **Kevin B. Korb**[1] **Michael Wybrow**[1] **Ingrid Zukerman**[1]
Steven Mascaro[1] **Ross Pearson**[1]

[1]Department of Data Science & AI, Monash University, Melbourne, Australia

In many complex real-world situations, problem solving and decision making require effective reasoning about causation and uncertainty. However, human reasoning in these cases is notoriously prone to confusion and error [Kahneman et al., 1982]. One way to support better reasoning is to employ Bayesian networks (BNs) [Pearl, 1988] to model and represent uncertain situations clearly for the user, and make complex calculations quickly and accurately on demand. BNs have been deployed for this purpose in diverse domains such as medicine, education, engineering, surveillance, the law, weather forecasting, and the environment.

We present our new methodology and application, Bayesian ARgumentation via Delphi (BARD), which combines BNs with a Delphi social process: a systematic method for combining multiple perspectives in a democratic, reasoned, iterative manner [Linstone and Turoff, 1975]. The initial motivation for BARD was to extend the use of BNs to a new domain: intelligence analysis. Development began as part of the Crowdsourcing Evidence, Argumentation, Thinking and Evaluation (CREATE) program funded by the Intelligence Advanced Research Projects Activity (IARPA). CREATE sought to develop, and experimentally test, systems that use crowdsourcing and structured analytic techniques to improve analytic reasoning, and help users to communicate it, without requiring prior expertise or expert assistance.

For domain experts to construct their own BNs, current software has major deficiencies: (A) it usually requires substantial upfront training, not included in the software, (B) it does not provide much guidance on the model building process, or (C) on using the resulting model for reasoning and reporting, and (D) it does not provide any support for collaboratively building BNs. BARD addresses each of these deficiencies by integrating the following key novel features:

(A) Short, high-quality e-courses, tips and help on demand. BARD provides compressed, high-quality, training to allow novices to start using the system as soon as possible, and then receive further help as needed. All key elements of the BARD approach are condensed into 4 hr of short, interactive, modular e-courses. These are augmented by embedded help components that include training problems with ideal solutions, optional product tours, context-specific tips, and guidance on what to do next.

(B) A stepwise, iterative and incremental BN construction process. BARD breaks down a given task into six steps that are performed by the analysts: (1) pre-modeling exploration of the problem to be solved, (2–4) building the components of the BN, (5) exploring the BN's reasoning on specific scenarios (Figure 1), and (6) report writing with BARD's support. However, progress needn't be linear: BARD encourages analysts to incrementally and iteratively build their individual BNs, and seek regular feedback through communication with other group members and the facilitator.

(C) Analytical report templates and an automated explanation tool. BARD guides verbal reporting with an analytical template, designed to elicit relevant points in a logical and thorough way that is consistent with general good reasoning guidelines [e.g., Clapper, 2015]. BARD also auto-generates from the BN model many key points, in English, organized according to the same template—such as the diagnosticity of evidence and critical uncertainties—which users can easily incorporate into their reports.

(D) A multi-user web-based software platform and Delphi-style social processes. Analysts in small groups, optionally assisted by a facilitator, are guided through a structured Delphi-style elicitation protocol to consider and represent their problem-relevant knowledge in a causal BN augmented by descriptive annotations. BARD provides tools to assist with the elicitation of BN structure and parameters, with review and consensus building in the group, and with evaluation of the results. Delphi protocols are designed to avoid common pitfalls of group deliberation, e.g., overweighting the first opinion expressed. Analysts are required to first develop an answer on their own, in a private phase, then BARD shows other group members' contributions after

Accepted for the 15th Bayesian Modeling Applications Workshop (BMAW-2021).

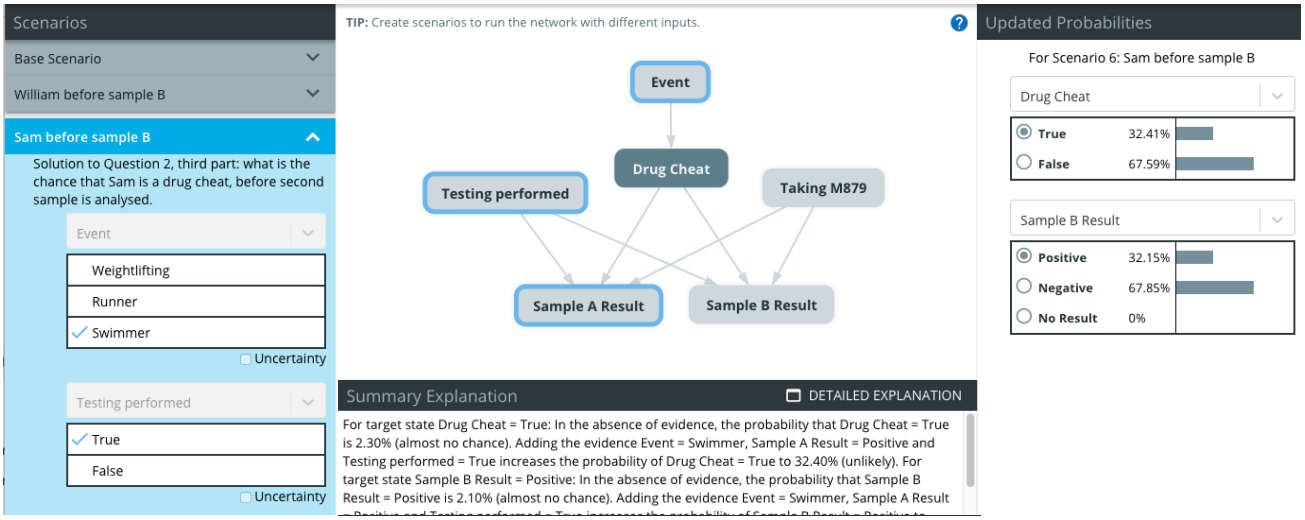

Figure 1: BARD screenshot for a training problem at Step 5: Explore Network. Evidence can be added into scenarios in the left panel, and updated probabilities for the chosen output variables are shown in the right panel. The network structure is shown in the center panel with the the evidence variables highlighted in blue, and below this is a summary verbal explanation.

the analyst has published his or her initial attempt. This approach maximizes the diversity of answers from which the group starts its work, then encourages analysts to improve their own answers, often by incorporating good features proposed by others and converging on a consensus. Other major features of Delphi utilized by BARD are anonymity and moderated discussion, both of which should help groups avoid being unduly influenced by high status or opinionated individuals rather than the knowledgeable.

Three initial experiments demonstrate that BARD aids both individuals and groups in reasoning and reporting. Our cognitive psychologists developed suitable probabilistic reasoning problems that can be solved with simple BNs but incorporate known reasoning difficulties. [Anonymized, 2020], with 59 participants, showed that BARD improves the reasoning of individuals here with a large effect size (Glass' Δ 0.8). [Anonymized, 2020], with 256 participants, showed that teams of 6–9 using BARD outperformed control individuals assisted by both generic training in good reasoning and the best available pen-and-paper tools [Gigerenzer and Hoffrage, 1995], with very large to huge effects (Glass' Δ 1.4–2.2), greatly exceededing CREATE's initial target. BARD users were also surveyed using Brooke's System Usability Scale, which showed overall positive user satisfaction. The increase in effect sizes from groups suggests BARD's BN-building and collaboration combine beneficially and cumulatively, and [Anonymized, 2020], with 75 participants, confirmed that even a minimal Delphi-style interaction improves the BN structures produced.

In future work, we plan to enhance BARD incrementally, both with industry-standard BN features and with novelties designed to support reasoning, e.g., BN "idioms" for legal arguments [Lagnado et al., 2013]. Further research comparing variations in social protocols, which can be readily configured in BARD, is needed to measure and optimize their impact. We are currently completing an experiment comparing BARD to standard BN software for teaching BN modeling to undergraduate IT students, and have commenced a spin-off XAI project to improve and test our automated explanations, by making individual elements available on demand and combining verbal with visual aids.

References

James Clapper. *Intelligence Community Directive (ICD) 203, Analytic Standards*. United States Office of the Director of National Intelligence (ODNI), 2015.

Gerd Gigerenzer and Ulrich Hoffrage. How to improve Bayesian reasoning without instruction: frequency formats. *Psychological Review*, 102(4):684, 1995.

Daniel Kahneman, Paul Slovic, and Amos Tversky, editors. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, New York, 1982.

David A Lagnado, Norman Fenton, and Martin Neil. Legal idioms: a framework for evidential reasoning. *Argument & Computation*, 4(1):46–63, 2013.

H.A. Linstone and M. Turoff. *The Delphi Method: Techniques and Applications*. Addison-Wesley, London, 1975.

Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, 1988. ISBN 0080514898.