Editorial to the Special Issue on Evidence-based Reasoning and Applications, Computing in Science and Engineering, Vol. 20, Issue 6, pp. 1-2, November/December 2018

Evidence-based Reasoning and Applications

Gheorghe Tecuci
Computer Science Department and Learning Agents Center, George Mason University

This special CiSE issue features recent work on the theory and applications of evidence-based reasoning. Different approaches to this type of reasoning, including interactive and automatic ones, are discussed and illustrated in intelligence analysis, medicine, and cybersecurity.

Evidence is any observable sign, datum, or item of information that is relevant in deciding whether a statement or hypothesis we are considering (e.g., a scientific or medical claim) is true or false. It is important to realize that almost any item of information we are dealing with is not a fact, but evidence about that fact. Indeed, measurement instruments are imperfect and may not even be well-calibrated, our senses as well are imperfect and our observations may not entirely reflect the reality, statements from people may be biases or sometimes even deliberately deceptive, and so on.[1]

Evidence-based reasoning is the type of reasoning that explicitly treats information as evidence, when assessing the truthfulness of hypotheses. In particular, it considers evidence as always being incomplete and with various degrees of credibility, and that it may also be ambiguous (not clear what it says), inconclusive (consistent with more than one hypothesis), and contradictory (some favoring one hypothesis while other evidence favoring other hypotheses).[2]

Evidence-based reasoning is at the core of many problem solving and decision making tasks in a wide variety of domains, including law, intelligence analysis, forensics, cybersecurity, medicine, physics, chemistry, history, archaeology, and many others.[3] This is not surprising because, as Jeremy Bentham stated over two centuries ago, “The field of evidence is no other than the field of knowledge.”[4]

This special issue features recent work on the theory and applications of evidence-based reasoning. Three of the papers are by performers in the Crowdsourcing Evidence, Argumentation, Thinking and Evaluation (CREATE) program of the Intelligence Advanced Research Projects Activity (IARPA). Launched in January 2017, CREATE is a 4.5-year effort to develop and experimentally test systems that use crowdsourcing and structured analytic techniques to improve

analytic reasoning. These papers present applications to intelligence analysis. Another paper presents a general cognitive view on reasoning with evidence, and focuses on its application to medicine. While all these approaches view evidence-based reasoning as a collaborative human-machine effort, the fifth paper in this special issue views evidence-based reasoning as an automatic process performed by intelligent agents, in the area of cybersecurity.

In "Evidence-based Reasoning in Intelligence Analysis: Structured Methodology and System," the authors introduce a systematic methodology for "connecting the dots" in intelligence analysis, based on the scientific method, that encompasses the processes of hypotheses generation, evidence collection, and hypotheses analysis. This approach is implemented in the Cogent cognitive assistant that facilitates a synergistic integration of analyst imaginative reasoning and expertise with agent knowledge and critical reasoning, to develop Wigmorean argumentations[2] for answering intelligence questions. The analysis methodology and the capabilities of Cogent are illustrated with a detailed example of answering the question "Which surface-to-air missile system is Manada selling Sindia?" based on imperfect information.

In "SWARM: Cultivating Evidence-Based Reasoning," the authors describe an online collaboration platform supporting evidence-based reasoning, focusing on its three main design principles: cultivating user engagement, exploiting natural expertise, and supporting rich collaboration. SWARM attempts to cultivate user engagement by providing opportunities for users to satisfy three deep needs: autonomy, competence, and relatedness. Exploiting natural expertise means setting up the conditions for users to freely exercise their reasoning abilities in ways they feel are most productive to them. Finally, the system is designed to provide rich mechanisms of collaboration in developing reports for answering intelligence questions. Although the application focus of SWARM is intelligence analysis, the platform is generic enough to be used in other domains as well.

In "User-Centered Design and Experimentation to Develop Effective Software for Evidence-Based Reasoning in the Intelligence Community: The TRACE Project," the authors describe their work on experimentally evaluating existing structured analytic techniques in order to determine the most effective ones. The best methods are integrated into the TRACE system together with additional techniques, including nudging, checklists, structured debate, and pros and cons. To help reasoners process and understand information from multiple sources, TRACE includes tools to enable highlighting, commenting, and tagging. A final checklist of attributes of a completed analytic report helps improve its overall quality of reasoning.

In "Evidence and Argument: A Cognitive View," the author proposes a unified perspective on evidence and evidence-based reasoning, as a generalization of the evidentiary practice in law, science, and medicine. Based on the observation that the degree to which evidence supports hypotheses in law or in science is not usually quantified, he presents an alternative to the evidentiary reasoning approaches that are based exclusively on probability concepts. His approach is based on logical theories of argumentation, relying on qualitative and logical representations of knowledge when uncertainty cannot be quantified. This offers a natural, qualitative way of reasoning with evidence, the practicality of which was demonstrated in a wide range of deployed medical applications.

Finally, in "Evidence-based Detection of Advanced Persistent Threats," the authors present an innovative approach to automatically detecting sophisticated cyberattacks that currently can only be detected manually by cyber analysts. First, a cybersecurity expert teaches a learning agent how to detect such attacks, through explained examples of analysis that follow the systematic approach presented in the first paper of this special issue. The trained learning agent is then customized into a team of specialized autonomous collaborative agents that are integrated into a cybersecurity operations center to defend against cyberattacks by applying the learned expertise. As opposed to the approaches presented in the other papers of this issue, this paper shows how one can automate the entire process of evidence-based reasoning, including the most challenging one of hypotheses generation that involves abductive (imaginative) reasoning. The paper also discusses the broader applicability of this approach to other areas, such as intelligence, surveillance and reconnaissance, or automatic monitoring of industrial installation or patients.

Papers like those from this special issue are steps in the development of the emerging Science of Evidence,[5] a new discipline of a the highest importance given the pervasiveness and importance of evidence and evidence-based reasoning, not only in so many scientific or engineering disciplines, but also in our daily lives. Indeed, we would be well-advised to apply evidence-based reasoning in all our understanding of the information we encounter, whether from the Internet, the media, or our direct communication.

ACKNOWLEDGEMENTS

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA) under contract number 2017-16112300009, by the Air Force Research Laboratory (AFRL) under contract number FA8750-17-C-0002, and by George Mason University. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, AFRL, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

REFERENCES

1. Schum, D. A., *The Evidential Foundations of Probabilistic Reasoning*. Northwestern University Press, 2001.
2. Tecuci, G., Schum, D. A., Marcu, D., and Boicu, M., *Intelligence Analysis as Discovery of Evidence, Hypotheses, and Arguments: Connecting the Dots*, Cambridge University Press, 2016.
3. Tecuci G., Marcu D., Boicu M., Schum D.A., *Knowledge Engineering: Building Cognitive Assistants for Evidence-based Reasoning*, Cambridge University Press, 2016.
4. Bentham, J. (1810). *An Introductory View of the Rationale of the Law of Evidence for Use by Non-lawyers as well as Lawyers* (vi works 1-187 (Bowring edition, 1837-43) originally edited by James Mill circa 1810).
5. Schum, D. A., (2009). A science of evidence: contributions from law and probability, *Law, Probability and Risk*, 8, 197−231, doi:10.1093/lpr/mgp002

ABOUT THE AUTHOR

Gheorghe Tecuci is Professor of Computer Science and Director of the Learning Agents Center at George Mason University, member of the Romanian Academy, and former Chair of Artificial Intelligence at the U.S. Army War College. Email: tecuci@gmu.edu