Evidence Propagation and Consensus Formation in Noisy Environments

Extended Abstract

Michael Crosscombe
University of Bristol
United Kingdom
m.crosscombe@bristol.ac.uk

Jonathan Lawry
University of Bristol
United Kingdom
j.lawry@bristol.ac.uk

ABSTRACT

We study the effectiveness of consensus formation in multi-agent systems where belief updating is an iterative two-part process, consisting of both belief updating based on direct evidence and also belief combination between agents, within the context of a best-of-n problem. Agents' beliefs are represented within Dempster-Shafer theory by mass functions and we investigate the macro-level properties of four well-known belief combination operators: Dempster's rule, Yager's rule, Dubois & Prade's operator and the averaging operator. Simulation experiments are conducted for different evidence rates and noise levels. Broadly, Dubois & Prade's operator results in better convergence to the best state, and is more robust to noisy evidence.

KEYWORDS

Consensus formation; evidence propagation; noisy decision-making; emergent behaviour; distributed problem solving

ACM Reference Format:
Michael Crosscombe and Jonathan Lawry. 2019. Evidence Propagation and Consensus Formation in Noisy Environments. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION AND BACKGROUND

Agents operating in noisy and complex environments will receive evidence from a variety of different sources, many of which will be at least partially inconsistent. We therefore investigate the interaction between two broad categories of evidence; direct evidence from the environment and evidence received from other agents with whom an agent is interacting. For example, robots engaged in a search and rescue mission will receive data directly from sensors as well as information from other robots in the team.

The efficacy of combining these two types of evidence in multi-agent systems has been studied from a number of different perspectives. In social epistemology [5] has argued that agent-to-agent communications has an important role to play in propagating locally held information widely across a population. Simulation results are then presented which show that a combination of direct evidence and agent interaction, within the Hegselmann-Krause opinion dynamics model [8], results in faster convergence to the true state than updating based solely on direct evidence. A probabilistic agent-based model combining Bayesian updating and probability pooling of beliefs has been proposed in [12]. An alternative methodology exploits three-valued logic to combine both types of evidence [2] and has been effectively applied to distributed decision-making in swarm robotics [3].

In this current study we exploit the capacity of Dempster-Shafer theory (DST) to fuse conflicting evidence, in order to investigate how direct evidence can be combined with a process of iterative belief aggregation in the context of the best-of-n problem [14, 18]/ It is not our intention to study the axiomatic properties of particular operators at the local level (see [6] for an overview of such properties). Instead, our main contribution is a study of the macro-level convergence properties of several established operators when applied iteratively by a population of agents over long timescales and in conjunction with a process of evidential updating.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

2 AN OVERVIEW OF DEMPSTER-SHAFER THEORY

Given a set of states or frame of discernment $\mathbb{S} = \{s_1, ..., s_n\}$, let $2^{\mathbb{S}}$ denote the power set of $\mathbb{S}$. An agent's belief is then defined by a basic probability assignment, or *mass function* $m : 2^{\mathbb{S}} \to [0, 1]$, where $m(\emptyset) = 0$ and $\sum_{A \subseteq \mathbb{S}} m(A) = 1$. The mass function then characterises a belief and a plausibility measure defined on $2^{\mathbb{S}}$ such that for $A \subseteq \mathbb{S}$;

$$Bel(A) = \sum_{B \subseteq A} m(B) \text{ and } Pl(A) = \sum_{B : B \cap A \neq \emptyset} m(B)$$

and hence where $Pl(A) = 1 - Bel(A^c)$.

A number of operators have been proposed in DST for combining or fusing mass functions [16]. In this paper we will compare four operators:

Definition 2.1. Combination Operators
Let m_1 and m_2 be mass functions on $2^{\mathbb{S}}$. Then the combined mass function $m_1 \odot m_2$ is a function $m_1 \odot m_2 : 2^{\mathbb{S}} \to [0, 1]$ such that for $\emptyset \neq A, B, C \subseteq \mathbb{S}$;

- *Dempster's Rule [15]*:
  $m_1 \odot m_2(C) = \frac{1}{1-K} \sum_{A \cap B = C \neq \emptyset} m_1(A) \cdot m_2(B)$, where $K = \sum_{A \cap B = \emptyset} m_1(A) \cdot m_2(B)$.
- *Dubois & Prade's Operator [7]*
  $m_1 \odot m_2(C) = \sum_{A \cap B = C \neq \emptyset} m_1(A) \cdot m_2(B)$
  $+ \sum_{A \cap B = \emptyset, A \cup B = C} m_1(A) \cdot m_2(B)$.
- *Yager's Operator [20]*
  $m_1 \odot m_2(C) = \sum_{A \cap B = C \neq \emptyset} m_1(A) \cdot m_2(B)$
  if $C \neq \mathbb{S}$ and $m_1 \odot m_2(\mathbb{S}) = 1 - \sum_{C \neq \mathbb{S}} m_1 \odot m_2(C)$

- *Averaging Operator*
  $m_1 \odot m_2(C) = \frac{1}{2}(m_1(C) + m_2(C))$

The first three operators all make the assumption of independence between the sources of the evidence to be combined but then employ different techniques for dealing with the resulting inconsistency.

In the agent-based model of the best-of-n problem proposed in the following section, agents are required to make a choice as to which state they should investigate at any particular time. To this end we utilise the notion of *pignistic distribution* proposed by Smets and Kennes [17]:

Definition 2.2. Pignistic Distribution
Given a mass function m, the corresponding pignistic distribution on $\mathbb{S}$ is given by;

$$P(s_i|m) = \sum_{A:s_i \in A} \frac{m(A)}{|A|}.$$

Relevant applications of DST (in its various forms) or the operators to dynamic multi-agent belief revision include [1, 2, 4, 9–11, 13, 19].

3 THE BEST-OF-n PROBLEM

We take the n choices to be the states $\mathbb{S}$. Each state $s_i \in \mathbb{S}$ is assumed to have an associated quality value q_i which we take to be in the interval $[0, 1]$ with 0 and 1 corresponding to minimal and maximal quality, respectively.

In the best-of-n problem agents explore their environment and interact with each other with the aim of identifying which is the highest quality (or true) state. Agents sample states and receive evidence in the form of the quality, so that in the current context evidence E_i regarding state s_i takes the form of the mass function $m_{E_i} = \{s_i\} : q_i, \mathbb{S} : 1 - q_i$. Hence, q_i is taken as quantifying both the evidence directly in favour of s_i provided by E_i, and also the evidence directly against any other state s_j for $j \neq i$. Given evidence E_i an agent updates her belief by combining her current mass function m with m_{E_i} using a combination operator so as to obtain the new mass function given by $m \odot m_{E_i}$.

A summary of the process by which an agent might obtain direct evidence in this model is then as follows: Based on her current mass function m, an agent stochastically selects a state $s_i \in \mathbb{S}$ to investigate, according to the pignistic probability distribution for m as given in Definition 2.2. She will update m to $m \odot m_{E_i}$ with probability $P(s_i|m) \times r$ for $i = 1, \ldots, n$ and leave her belief unchanged with probability $(1 - r)$, where $r \in [0, 1]$ is a fixed evidence rate. We also allow for the possibility of noise in the evidential updating process. This is modelled by a random variable ϵ associated with each quality value such that the evidence E_i received by an agent has the form $m_{E_i} = \{s_i\} : q_i + \epsilon, \mathbb{S} : 1 - q_i - \epsilon$, where ϵ is a normally distributed random variable with mean 0 and standard deviation σ[1]. Overall, the process of updating from direct evidence is governed by the two parameters, r and σ, quantifying the availability of evidence and the level of associated noise, respectively.

[1]We normalise so that if $q_i + \epsilon < 0$ then it is set to 0, and if $q_i + \epsilon > 1$ then it is set to 1.

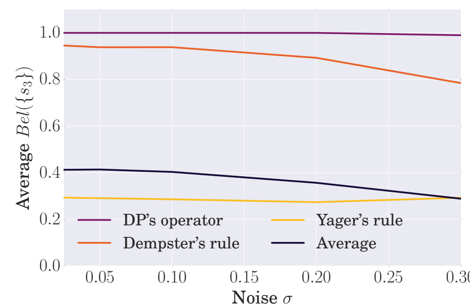

Figure 1: Average $Bel(\{s_3\})$ for all four operators plotted against $\sigma \in [0, 0.3]$ and evidence rate $r = 0.05$.

In addition to receiving direct evidence we also include belief combination between agents in this model. This is conducted in a pairwise symmetric manner in which two agents are selected at random to combine their beliefs, such that if the two agents have beliefs m_1 and m_2, respectively, then they both replace these with $m_1 \odot m_2$.

4 SIMULATION EXPERIMENTS

We consider a population $\mathcal{A}$ of $k = 100$ agents with beliefs initialised so that $m_i^0 = \mathbb{S} : 1$ for $i = 1, \ldots, 100$. In other words, at the beginning of each simulation every agent is in a state of complete ignorance. Each experiment is run for a maximum of 10 000 iterations, or until the population converges.

For a given set of parameter values the simulation is run 100 times and results are then averaged across these runs. Quality values are defined so that $q_i = \frac{i}{n+1}$ for $i = 1, \ldots, n$ and consequently s_n is the best state. Hence, in the following, $Bel(\{s_n\})$ provides a measure of convergence performance for the two operators.

Initially we consider the best-of-n problem where $n = 3$ with quality values $q_1 = 0.25$, $q_2 = 0.5$ and $q_3 = 0.75$. For an evidence rate of $r = 0.05$ Figure 1 shows the average value of $Bel(\{s_3\})$ at steady state plotted against $\sigma \in [0, 0.3]$. We see that Dubois & Prade operator is the most robust operator to increased noise. Specifically, for $\sigma = 0$, Dubois & Prade's operator converges to an average value of $Bel(\{s_3\}) = 1$ and for $\sigma = 0.3$ this only decreases to 0.98. On the other hand, the presence of noise at this evidence rate has a much higher impact on the performance of Dempster's rule and the averaging operator. Yager's rule is the exception in this context since for $r = 0.05$ the average value of $Bel(\{s_3\})$ remains constant at approximately 0.3.

ACKNOWLEDGMENTS

This work was funded and delivered in partnership between the Thales Group and the University of Bristol and with the support of the UK Engineering and Physical Sciences Research Council Research Grant Award Reference EP/R004757/1 entitled "Thales-Bristol Partnership in Hybrid Autonomous Systems Engineering (T-B PHASE)."

REFERENCES

[1] Jin-Hee Cho and Ananthram Swami. 2014. Dynamics of Uncertain Opinions in Social Networks. *2014 IEEE Military Communications Conference* (2014), 1627–1632. https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6956988

[2] Michael Crosscombe and Jonathan Lawry. 2016. A Model of Multi-Agent Consensus for Vague and Uncertain Beliefs. *Adaptive Behavior* 24, 4 (2016), 249–260. https://journals.sagepub.com/doi/abs/10.1177/1059712316661395

[3] Michael Crosscombe, Jonathan Lawry, Sabine Hauert, and Martin Homer. 2017. Robust distributed decision-making in robot swarms: Exploiting a third truth state. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 4326–4332. https://doi.org/10.1109/IROS.2017.8206297

[4] Ranga Dabarera, Rafael Nú, ñez, Kamal Premaratne, and Manohar N. Murthi. 2014. Dynamics of Belief Theoretic Agent Opinions Under Bounded Confidence. *17'th International Conference on Information Fusion (FUSION)* (2014), 1–8. https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6916216

[5] Igor Douven and Christoph Kelp. 2011. Truth Approximation, Social Epistemology, and Opinion Dynamics. *Erkenntnis* 75 (2011), 271–283. https://link.springer.com/article/10.1007/s10670-011-9295-x

[6] Didier Dubois, Weiru Liu, Jianbing Ma, and Henri Prade. 2016. The basic principles of uncertain information fusion. An organised review of merging rules in different representation frameworks. *Information Fusion* 32 (2016), 12 – 39. https://doi.org/10.1016/j.inffus.2016.02.006

[7] Didier Dubois and Henri Prade. 2008. A set-theoretic view of belief functions. *Classic Works of the Dempster-Shafer Theory of Belief Functions. Studies in Fuzziness and Soft Computing* 219 (2008), 375–410. https://doi.org/10.1007/978-3-540-44792-4_14

[8] Rainer Hegselmann and Ulrich Krause. 2002. Opinion dynamics and bounded confidence: models, analysis and simulation. *J. Artificial Societies and Social Simulation* 5 (2002).

[9] Audun Jøsang. 2002. The consensus operator for combining beliefs. *Artificial Intelligence* 141, 1-2 (oct 2002), 157–170. https://doi.org/10.1016/S0004-3702(02)00259-X

[10] Orakanya Kanjanatarakul and Thierry Denoux. 2017. Distributed data fusion in the dempster-shafer framework. In *2017 12th System of Systems Engineering Conference (SoSE)*. IEEE, 1–6. https://doi.org/10.1109/SYSOSE.2017.7994954

[11] S. Kornienko, O. Kornienko, C. Constantinescu, M. Pradier, and P. Levi. 2005. Cognitive Micro-Agents: Individual and Collective Perception in Microrobotic Swarm. In *IJCAI-05 Workshop on Agents in Real-Time and Dynamic Environment*. 33–42. http://cybertronica.co/sites/default/files/publications/IJCAI-workshop-ARTDE.pdf

[12] Chanelle Lee, Jonathan Lawry, and Alan Winfield. 2018. Combining Opinion Pooling and Evidential Updating for Multi-Agent Consensus. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18) Combining*. 347–353.

[13] Xi Lu, Hongming Mo, and Yong Deng. 2015. An evidential opinion dynamics model based on heterogeneous social influential power. *Chaos, Solitons and Fractals* 73 (2015), 98–107. https://doi.org/10.1016/j.chaos.2015.01.007

[14] C. A. C. Parker and H. Zhang. 2009. Cooperative Decision-Making in Decentralized Multiple-Robot Systems: The Best-of-N Problem. *IEEE/ASME Transactions on Mechatronics* 14, 2 (April 2009), 240–251. https://doi.org/10.1109/TMECH.2009.2014370

[15] Glenn Shafer. 1976. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton. 314 pages.

[16] Philippe Smets. 2007. Analyzing the combination of conflicting belief functions. *Information Fusion* 8, 4 (2007), 387–412. https://doi.org/10.1016/j.inffus.2006.04.003

[17] Philippe Smets and Robert Kennes. 1994. The Transferable Belief Model. *Artificial Intelligence* 66 (1994), 387–412.

[18] Gabriele Valentini, Eliseo Ferrante, and Marco Dorigo. 2017. The Best-of-n Problem in Robot Swarms: Formalization, State of the Art, and Novel Perspectives. *Frontiers in Robotics and AI* 4, March (2017). https://doi.org/10.3389/frobt.2017.00009

[19] Thanuka L Wickramarathne, Kamal Premaratine, Manohar N Murthi, and Nitesh V. Chawla. 2014. Convergence Analysis of Iterated Belief Revision in Complex Fusion Environments. *IEEE Journal of Selected Topics in Signal Processing* 8, 4 (2014), 598–612. https://doi.org/10.1109/JSTSP.2014.2314854

[20] Ronald R. Yager. 1992. On the specificity of a possibility distribution. *Fuzzy Sets and Systems* 50, 3 (1992), 279 – 292. https://doi.org/10.1016/0165-0114(92)90226-T