Combining Opinion Pooling and Evidential Updating for Multi-Agent Consensus

Chanelle Lee[1,2,3], **Jonathan Lawry**[2], **Alan Winfield**[1,3]
[1] Bristol Robotics Laboratory
[2] University of Bristol
[3] University of the West of England, Bristol
{c.l.lee, Alan.Winfield}@brl.ac.uk, J.Lawry@bristol.ac.uk

Abstract

The evidence available to a multi-agent system can take at least two distinct forms. There can be direct evidence from the environment resulting, for example, from sensor measurements or from running tests or experiments. In addition, agents also gain evidence from other individuals in the population with whom they are interacting. We, therefore, envisage an agent's beliefs as a probability distribution over a set of hypotheses of interest, which are updated either on the basis of direct evidence using Bayesian updating, or by taking account of the probabilities of other agents using opinion pooling. This paper investigates the relationship between these two processes in a multi-agent setting. We consider a possible Bayesian interpretation of probability pooling and then explore properties for pooling operators governing the extent to which direct evidence is diluted, preserved or amplified by the pooling process. We then use simulation experiments to show that pooling operators can provide a mechanism by which a limited amount of direct evidence can be efficiently propagated through a population of agents so that an appropriate consensus is reached. In particular, we explore the convergence properties of a parameterised family of operators with a range of evidence propagation strengths.

1 Introduction and Background

In agent-based systems individuals receive information from at least two distinct sources; either as direct evidence from their environment, e.g. directly from different sensor modalities or by running tests or experiments, or from other individuals in the population. More specifically, we envisage agents as holding beliefs about a set of hypotheses of interest and then adapting those beliefs, either on the basis of direct evidence, or by taking account of the beliefs of the other agents that they are interacting with. Assuming that the agents' beliefs take the form of probability distributions over the hypotheses, we can model evidential updating in a Bayesian manner and then allow a form of probability pooling so that agents can combine their distributions with those of others.

This paper investigates the relationship between the two processes, updating based on direct evidence and probability pooling, in a dynamic setting in which agents adapt their beliefs over time. The motivating assumption is that pooling operators should provide a mechanism by which a limited amount of direct evidence can be efficiently propagated through the agent population so that an appropriate consensus is reached. This effect has already been studied in the context of social epistemology where the focus is on the overall knowledge of the group or population [Douven *et al.*, 2017]. For example, Douven and Kelp [Douven and Kelp, 2011] argue that dialogue between scientists is an important aspect of scientific research which complements experimental work. They support their case with simulation studies in which a population of agents receive occasional evidence, in this case the true value of a real-valued parameter, as well as pooling opinions by taking weighted linear combinations of values from other agents. The results suggest that the population converges much faster to the true parameter value when both pooling and evidential updating are employed, than it does when there is only evidential updating.

In the sequel we investigate probability pooling combined with updating based on direct evidence. We also adopt a Bayesian approach to pooling in which it is viewed as a form of second order updating on the basis of evidence consisting of the probability values of agents in the pool. From this perspective we provide a clear interpretation of a particular pooling operator and show the assumptions underlying it. We then introduce the properties of evidence dilution, preservation and amplification for pooling operators as a way of gauging the extent to which they enable direct evidence to be propagated across the agent population. Simulation experiments will be carried out to investigate a single parameter family of pooling operators which can be evidence diluting, preserving or amplifying depending on the parameter value. We consider the effectiveness of these operators at evidence propagation for different parameter values, different pool sizes and where the evidence is received at different rates. We focus on two mutually exclusive and exhaustive hypotheses denoted $\mathcal{H}_1$ and $\mathcal{H}_2$ so that each agent's beliefs can be characterised by a real number $x \in [0, 1]$ indicating that $P(\mathcal{H}_1) = x$ and $P(\mathcal{H}_2) = 1 - x$.

Probability pooling operators have been proposed as a means for combining a group of expert opinions to obtain

a single probability distribution, with early work dating, at least, back to DeGroot [DeGroot, 1974] and Stone [Stone and others, 1961]. As in both of these papers a common approach is simply to take the pooled distribution as a weighted linear combination of all the distributions of the agents in the cohort. For the two hypotheses case, if there are k agents with probabilities $x_1, \ldots, x_k$, then the pooled probability of $\mathcal{H}_1$ is $\sum_{i=1}^{k} w_i x_i$ where $w_i > 0 : i = 1, \ldots, k$ and $\sum_{i=1}^{k} w_i = 1$. The weights w_i then quantify a relative level of confidence in the different agents. A general definition of a pooling operator is simply as follows:

Definition 1.1 (Pooling Operator)**:** A pooling operator for k agents is a function $c : [0,1]^k \rightarrow [0,1]$, so that for agents $A_1, \ldots, A_k$ with probabilities $P_{A_i}(\mathcal{H}_1) = x_i$ for $i = 1, \ldots, k$ then $c(x_1, \ldots, x_k)$ is the pooled probability of $\mathcal{H}_1$.

In this paper we focus on three related geometric pooling operators which have clear Bayesian justifications and which it turns out have interesting evidence propagation properties. A summary of these is given in definition 1.2.

Definition 1.2 (LogOp, SProdOp and ProdOP Pooling Operators)**:** The following are the pooling operators that we consider in this paper;

- The Log-Linear Operator (LogOp):
$$c(x_1, \ldots, x_k) = \frac{\prod_{i=1}^{k} x_i^{w_i}}{\prod_{i=1}^{k} x_i^{w_i} + \prod_{i=1}^{k} (1 - x_i)^{w_i}},$$
for $w_i \geq 0$.
- The Scaled Product Operator (SProdOp): This is a special case of LogOp in which $w_i = \ldots = w_k = w > 0$ so that,
$$c(x_1, \ldots, x_k) = \frac{(\prod_{i=1}^{k} x_i)^w}{(\prod_{i=1}^{k} x_i)^w + (\prod_{i=1}^{k} (1 - x_i))^w},$$
for $w_i \geq 0$.
- The Product Operator (ProdOp): This is a special case of SProdOp in which $w = 1$ so that,
$$c(x_1, \ldots, x_k) = \frac{\prod_{i=1}^{k} x_i}{\prod_{i=1}^{k} x_i + \prod_{i=1}^{k} (1 - x_i)}.$$

LogOp has been widely discussed in the literature (see Genest and Zidek [Genest and Zidek, 1986] for an early overview), with a variety of possible justifications for the operator being presented. For instance, it is known to be the operator which minimizes the weighted Kullback-Leibler divergence between the pooled probability distribution and the agents' distributions [Abbas, 2009]. Pennock [Pennock and Wellman, 1999] shows that LogOp preserves Markov independence and can therefore be useful in graphical probability models. Most recently Pettigrew [Pettigrew, 2017] has considered LogOp as a way of pooling incoherent probabilities. More specifically, agents are deemed to be incoherent if their probabilities do not sum to one. Such distributions can be 'fixed' by renormalisation and [Pettigrew, 2017] suggests that fixing before pooling and fixing after pooling should yield the same result. Indeed if standard renormalisation is employed (i.e. dividing by the sum of the non-normalised probabilities) then LogOp is shown to satisfy this property.

ProdOp, a special case of LogOp, has been proposed as an effective tool for classifier combination [Hinton, 1999] and also for expert pooling in the management sciences [Bordley, 1982]. Other applications include natural language processing [Osborne and Baldridge, 2004].

An outline of the remainder of the paper is as follows. Section 2 will consider a Bayesian interpretation for LogOp under certain assumptions including agent independence. In section 3 we consider the combination of pooling and direct evidence and in particular, we introduce the properties of evidence preservation, amplification and dilution as means of classifying the extent to which the pooling process propagates evidence. In section 4 we present simulation experiments into the convergence properties of SProdOp for different values of w and k, and at different evidence rates. Finally, in section 5 we give some discussions and conclusions.

2 A Bayesian Interpretation of Opinion Pooling

In this section we consider pooling as being a type of conditioning based on second order evidence in the form of agents' beliefs. For this we adopt a Bayesian approach and show that it provides clear justifications for LogOp. In this context we introduce the notion of an oracle as an arbitrator who performs Bayesian updating given the probabilities from all k agents in the pool.

Given a pool of k agents suppose that the aggregated probability corresponds to the conditional probability of $\mathcal{H}_1$ of an 'oracle' O, given the evidence provided by the probabilities of the agents in the pool. Here the oracle is an abstract entity which we might choose to interpret in a number of different ways. For instance, O might be a kind of independent arbitrator tasked with identifying a single shared probability which takes account of the beliefs of the other agents. Alternatively, we could think of O as an aggregate representation of the whole pool. The idea of opinion pooling as based on the judgement of an oracle is well-known, with [Hogarth, 1975] and [Keeney and Raiffa, 1976] referring to O as a 'synthetic personality' and a 'supra Bayesian' respectively. Furthermore, early work by [Winkler, 1968] and [Morris, 1974] shows that from a Bayesian perspective, the pooling operator c can then be understood to be O's posterior distribution determined along the following lines. Suppose that O has a prior probability of $\mathcal{H}_1$, denoted $P_O(\mathcal{H}_1)$, then this is conditioned on the evidence $\mathcal{B} = \bigwedge_{i=1}^{k} (P_{A_i}(\mathcal{H}_1) = x_i)$ representing the beliefs of the k agents, according to Bayes' theorem as follows;

$$\begin{aligned} c(x_1, \ldots, x_k) &= P_O(\mathcal{H}_1 | \mathcal{B}) \\ &= \frac{P_O(\mathcal{B}|\mathcal{H}_1) P_O(\mathcal{H}_1)}{P_O(\mathcal{B}|\mathcal{H}_1) P_O(\mathcal{H}_1) + P_O(\mathcal{B}|\mathcal{H}_2) P_O(\mathcal{H}_2)}. \end{aligned}$$

We now consider two assumptions that O might make concerning the distribution of $P_{A_i}(\mathcal{H}_1)$ for $i = 1, \ldots, k$.

Definition 2.1 (Likelihood Symmetry)**:** $\forall x_i$,
$P_O(\bigwedge_{i=1}^{k}(P_{A_i}(\mathcal{H}_1) = x_i)|\mathcal{H}_2) =$

$$P_O(\bigwedge_{i=1}^{k}(P_{A_i}(\mathcal{H}_1) = 1 - x_i)|\mathcal{H}_1).$$

This assumes that learning $\mathcal{H}_1$ holds provides O with the same information about the agents' probabilities of $\mathcal{H}_2$ as learning $\mathcal{H}_2$ holds provides about their probabilities of $\mathcal{H}_1$.

Definition 2.2 (Independent Agents)**:** The agents $A_1, \ldots, A_k$ are independent if $\forall x_i$,
$P_O(\bigwedge_{i=1}^{k}(P_{A_i}(\mathcal{H}_1) = x_i)|\mathcal{H}_j) =$

$$\prod_{i=1}^{k} P_O(P_{A_i}(\mathcal{H}_1) = x_i|\mathcal{H}_j).$$

The agents are deemed to be independent if the random variables $P_{A_1}(\mathcal{H}_1), \ldots, P_{A_k}(\mathcal{H}_1)$ are conditionally independent of each other given either $\mathcal{H}_1$ or $\mathcal{H}_2$.

In the following theorem we consider the situation in which given $\mathcal{H}_1$, O considers each agents' probability to be distributed according to a Beta distribution with parameter values a_i and b_i for $i = 1, \ldots, k$.

Theorem 2.1. *If given $\mathcal{H}_1$, $\forall i$, $P_{A_i}(\mathcal{H}_1)$ has a Beta distribution with parameters a_i and b_i where $a_i \geq b_i$ so that $P_O(P_{A_i}(\mathcal{H}_1) = x_i|\mathcal{H}_1) \propto x_i^{a_i-1}(1-x_i)^{b_i-1}$, and if we assume likelihood symmetry, independent agents and that $P_O(\mathcal{H}_1) = 0.5$ then c is LogOp with $w_i = a_i - b_i$ for $i = 1, \ldots, n$.*

Proof. By independence and assuming

$$P_O(P_{A_i}(\mathcal{H}_1) = x_i|\mathcal{H}_1) \propto x_i^{a_i-1}(1-x_i)^{b_i-1},$$

it follows that;

$$P_O(\mathcal{B}|\mathcal{H}_1) \propto \prod_{i=1}^{k} x_i^{a_i-1}(1-x_i)^{b_i-1}.$$

Also, by independence and likelihood symmetry it holds that;

$$P_O(\mathcal{B}|\mathcal{H}_2) \propto \prod_{i=1}^{k}(1-x_i)^{a_i-1}x_i^{b_i-1}.$$

Hence, since $P_O(\mathcal{H}_1) = \frac{1}{2}$ then substituting into Bayes' theorem gives,
$P_O(\mathcal{H}_1|\mathcal{B}) =$

$$\frac{\prod_{i=1}^{k} x_i^{a_i-1}(1-x_i)^{b_i-1}}{\prod_{i=1}^{k} x_i^{a_i-1}(1-x_i)^{b_i-1} + \prod_{i=1}^{k}(1-x_i)^{a_i-1}x_i^{b_i-1}}.$$

Dividing top and bottom by $\prod_{i=1}^{k} x_i^{b_i-1}(1-x_i)^{b_i-1}$ gives
$P_O(\mathcal{H}_1|\mathcal{B}) =$

$$\frac{\prod_{i=1}^{k} x_i^{a_i-b_i}}{\prod_{i=1}^{k} x_i^{a_i-b_i} + \prod_{i=1}^{k}(1-x_i)^{a_i-b_i}},$$

as required. □

For a Beta distribution with parameters $a_i \geq b_i$ the skewness value is negative. This means that the conditional probability density for $P_{A_i}(\mathcal{H}_1)$ given $\mathcal{H}_1$ is skewed towards 1.

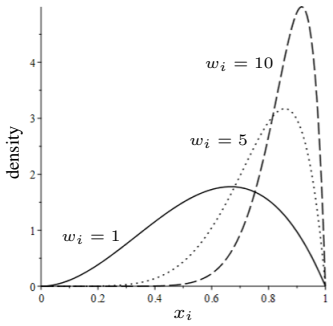

Figure 1: Beta distributions for $P_O(P_{A_i}(\mathcal{H}_1) = x_i|\mathcal{H}_1)$ for a fixed b_i and increasing w_i exhibiting increasing skewness towards 1. The density functions shown are for $b_i = 2$ and $w_i = 1, 5$ and 10.

Furthermore, the skewness can be expressed as a function of $w_i = a_i - b_i$ and b_i, which for any fixed b_i is a strictly decreasing function of w_i. Hence, in this case as w_i increases the distribution of $P_{A_i}(\mathcal{H}_1)$ becomes increasingly skewed towards 1. Negative skewness and hence w_i, in this context, is arguably an indicator of the capability of A_i when predicting which hypothesis holds, since if $\mathcal{H}_1$ is true then we would expect that the probability $P_{A_i}(\mathcal{H}_1)$ of a capable agent would tend to be close to 1 (see figure 1).

In the next section we consider the combination of probability pooling with evidential updating on the basis of direct evidence e.g. from sensors or other sources, either obtained passively or through exploration and experimentation.

3 Combining Pooling Operators with Direct Evidence

Direct evidence is assumed to correspond to an assertion that one of the hypotheses holds i.e. evidence E is a variable taking value $\mathcal{H}_1$ or $\mathcal{H}_2$. An agent then updates their probability of $\mathcal{H}_1$ from x to $x|E$ using Bayes' theorem as follows:

Definition 3.1 (Evidential Updating)**:** For $x \in [0, 1]$, $E \in \{\mathcal{H}_1, \mathcal{H}_2\}$ and $\alpha \in [0, \frac{1}{2}]$, then $x|E \in [0, 1]$ such that;

$$x|E = \frac{\delta_E x}{\delta_E x + (1-\delta_E)(1-x)},$$

where $\delta_{\mathcal{H}_1} = 1 - \alpha$ and $\delta_{\mathcal{H}_2} = \alpha$.

The rationale behind definition 3.1 is that, according to Bayes' theorem, an agent with prior probability x should update their beliefs given evidence E such that;

$$\begin{aligned} x|E &= P(\mathcal{H}_1|E) \\ &= \frac{P(E|\mathcal{H}_1)P(\mathcal{H}_1)}{P(E|\mathcal{H}_1)P(\mathcal{H}_1) + P(E|\mathcal{H}_2)P(\mathcal{H}_2)} \\ &= \frac{P(E|\mathcal{H}_1)x}{P(E|\mathcal{H}_1)x + P(E|\mathcal{H}_2)(1-x)}. \end{aligned}$$

We then make the further assumption that the likelihood of E is given by:

$$P(E|\mathcal{H}_i) = \begin{cases} 1-\alpha : E = \mathcal{H}_i, \\ \alpha : E \neq \mathcal{H}_i. \end{cases}$$

Here α quantifies the (agent's belief about the) reliability of the source of the evidence. For $\alpha = 0$ the source is absolutely reliable and will only provide E if and only if E is true. For $\alpha = 0.5$ the source is totally unreliable and is just as likely to provide E if it is false as if it is true. Notice that when $\alpha = 0.5$, $x|E = x$. This seems to be intuitive since if an agent has absolutely no faith in a source of evidence, they should not change their beliefs on the basis of evidence that it provides. Notice in the case that $\alpha = 0$, then $1|\mathcal{H}_2$ and $0|\mathcal{H}_1$ are undefined.

We now introduce three possible properties concerning the extent to which existing direct evidence is propagated through the pooling process. In each case only probabilities in the open interval $(0, 1)$ are considered so as to avoid any cases in which conditional probabilities (definition 3.1) are undefined.

Definition 3.2 (Evidence Preservation)**:** [Dietrich *et al.*, 2016] For $E \in \{\mathcal{H}_1, \mathcal{H}_2\}$, $\forall x_i \in (0, 1) : i = 1, \dots, k$, $\forall i \in \{1, \dots, k\}$, $\forall \alpha \in [0, \frac{1}{2}]$;
$c(x_1, \dots, x_{i-1}, x_i|E, x_{i+1}, \dots, x_k) =$
$c(x_1, \dots, x_i, \dots, x_k)|E$.

Definition 3.3 (Evidence Dilution)**:** $\forall x_i \in (0, 1) : i = 1, \dots, k$, $\forall i \in \{1, \dots, k\}$, $\forall \alpha \in [0, \frac{1}{2})$;
$c(x_1, \dots, x_{i-1}, x_i|\mathcal{H}_1, x_{i+1}, \dots, x_k) \leq$
$c(x_1, \dots, x_i, \dots, x_k)|\mathcal{H}_1$ and
$c(x_1, \dots, x_{i-1}, x_i|\mathcal{H}_2, x_{i+1}, \dots, x_k) \geq$
$c(x_1, \dots, x_i, \dots, x_k)|\mathcal{H}_2$.
Furthermore, there exists $(x_1, \dots, x_k) \in (0, 1)^k$ such that the above inequalities are strict.

Definition 3.4 (Evidence Amplification)**:** $\forall x_i \in (0, 1) : i = 1, \dots, k$, $\forall i \in \{1, \dots, k\}$, $\forall \alpha \in [0, \frac{1}{2})$;
$c(x_1, \dots, x_{i-1}, x_i|\mathcal{H}_1, x_{i+1}, \dots, x_k) \geq$
$c(x_1, \dots, x_i, \dots, x_k)|\mathcal{H}_1$ and
$c(x_1, \dots, x_{i-1}, x_i|\mathcal{H}_2, x_{i+1}, \dots, x_k) \leq$
$c(x_1, \dots, x_i, \dots, x_k)|\mathcal{H}_2$.
Furthermore, there exists $(x_1, \dots, x_k) \in (0, 1)^k$ such that the above inequalities are strict.

Evidence preservation was recently proposed by Dietrich and List [Dietrich *et al.*, 2016] who refer to it as 'Individual Bayesianism'. It requires that if a single agent receives evidence then that evidence is preserved by the pooling process. More specifically, it should make no difference whether the evidence is presented to any one of the agents prior to pooling or if it is presented after pooling. This property could be appropriate if different agents receive evidence from independent sources. To see this notice that for an evidence preserving operator;

$$c(x_1|E, \dots, x_k|E) = c(x_1, \dots, x_k)|E \dots |E.$$

In other words, if each agent receives the same evidence this has a strong reinforcement effect on the pooled probabilities. The evidence dilution and amplification properties then claim that pooling should respectively dilute or amplify the effect of evidence presented to any one of the agents. We now consider evidence dilution, preservation and amplification in detail so as to obtain an insight into the evidence propagation properties of the pooling operators introduced in section 1. Indeed, the following theorem due to [Dietrich *et al.*, 2016] shows that evidence preservation is a sufficiently strong property to characterise ProdOp.

Theorem 3.1. *[Dietrich* et al.*, 2016] A pooling operator c satisfies evidence preservation and $c(0.5, \dots, 0.5) = 0.5$ if and only if $\forall x_i \in (0, 1) : i = 1, \dots, k$, c is ProdOp.*

We now extend this result to show that SProdOp is evidence diluting, preserving or amplifying depending only on the value of the parameter w. In section 4 we will use this property of SProdOp to investigate optimal levels of evidence propagation for varying evidence rates.

Theorem 3.2. *SProdOp is evidence diluting if $w < 1$, evidence preserving if $w = 1$ and evidence amplifying if $w > 1$.*

Proof. W.l.o.g let $E = \mathcal{H}_1$. Then for the scaled product operator we have that;
$c(x_1, \dots, x_i|E, \dots, x_k) =$

$$\frac{(x_i|E)^w \left(\prod_{j\neq i} x_j\right)^w}{(x_i|E)^w \left(\prod_{j\neq i} x_j\right)^w + (1 - x_i|E)^w \left(\prod_{j\neq i}(1 - x_j)\right)^w}$$
$$= \frac{(1-\alpha)^w \left(\prod_{j=1}^k x_j\right)^w}{(1-\alpha)^w \left(\prod_{j=1}^k x_j\right)^w + \alpha^w \left(\prod_{j=1}^k (1 - x_j)\right)^w}.$$

Also,
$c(x_1, \dots, x_k)|E =$

$$\frac{(1-\alpha)c(x_1, \dots, x_k)}{(1-\alpha)c(x_1, \dots, x_k) + \alpha\,(1 - c(x_1, \dots, x_k))}$$
$$= \frac{(1-\alpha)\left(\prod_{j=1}^k x_j\right)^w}{(1-\alpha)\left(\prod_{j=1}^k x_j\right)^w + \alpha\left(\prod_{j=1}^k (1 - x_j)\right)^w}.$$

For evidence dilution we require that for some $(x_1, \dots, x_k) \in (0, 1)^k$, $c(x_1, \dots, x_i|E, \dots, x_k) < c(x_1, \dots, x_k)|E$. From the above this implies that;

$$\frac{(1-\alpha)^{w-1}}{(1-\alpha)^w \left(\prod_{j=1}^k x_j\right)^w + \alpha^w \left(\prod_{j=1}^k (1 - x_j)\right)^w}$$
$$< \frac{1}{(1-\alpha)\left(\prod_{j=1}^k x_j\right)^w + \alpha\left(\prod_{j=1}^k (1 - x_j)\right)^w} \Rightarrow$$
$$(1-\alpha)^{w-1}\left[(1-\alpha)\left(\prod_{j=1}^k x_j\right)^w + \alpha\left(\prod_{j=1}^k (1 - x_j)\right)^w\right]$$
$$< (1-\alpha)^w \left(\prod_{j=1}^k x_j\right)^w + \alpha^w \left(\prod_{j=1}^k (1 - x_j)\right)^w \Rightarrow$$
$$(1-\alpha)^{w-1}\alpha \left(\prod_{j=1}^k (1 - x_j)\right)^w < \alpha^w \left(\prod_{j=1}^k (1 - x_j)\right)^w$$
$$\Rightarrow (1-\alpha)^{w-1}\alpha < \alpha^w \Rightarrow (1-\alpha)^{w-1} < \alpha^{w-1}.$$

Now $\alpha < \frac{1}{2} \Rightarrow 1-\alpha > \alpha \Rightarrow (1-\alpha)^{w-1} < \alpha^{w-1}$ if and only if $w - 1 < 0$ if and only if $w < 1$. Similarly, for evidence

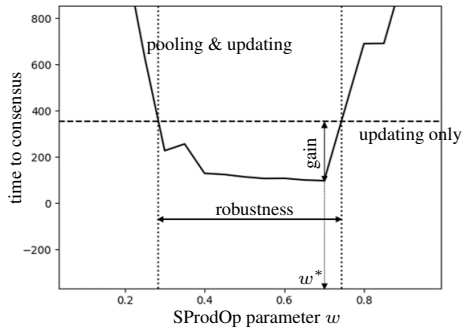

Figure 2: Plot of average time to consensus against w for $k = 3$ and $\epsilon = 2\%$.

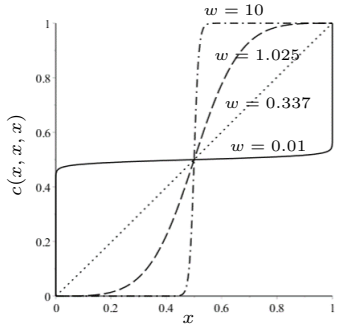

Figure 3: $c(x, x, x)$ as function of x for SProdOp with $w = 0.01$, $w = 0.337$, $w = 1.025$ and $w = 10$.

amplification we require that $c(x_1, \ldots, x_i|E, \ldots, x_k) > c(x_1, \ldots, x_k)|E$ which then holds if and only if $w > 1$. Finally, taking $w = 1$ corresponds to the product operator which is evidence preserving by theorem 3.1. □

4 Experimental Results

In this section we describe simulation experiments exploring the consensus attainment properties of SProdOp when implemented in a multi-agent system. The aim of these experiments is to investigate SProdOp using different values of w in order to better understand the system-level behaviour of evidence diluting, preserving and amplifying operators.

A population of 100 agents is initialised with probabilities of $\mathcal{H}_1$ randomly chosen from $[0, 1]$. At each iteration, the population is permuted to emulate movement, this approach being consistent with the well-stirred assumption as described in [Parker and Zhang, 2009]. A pool of k agents are then chosen at random from the population. Their probabilities are combined using SProdOp with each agent then adopting the resulting pooled probability. In addition to pooling, every agent also has a probability ϵ of directly receiving the evidence $E = \mathcal{H}_1$, in which case they use Bayesian updating to obtain a posterior probability as given in definition 3.1. We take $\alpha = 0.1$ indicating that agents have a high-level of trust in the sources of direct evidence, but which falls short of absolute confidence. In multi-agent systems or robot swarms we might envisage this set-up as modelling a scenario in which agents explore their environment and receive direct sensory evidence from time-to-time. Subsets of the population would then also regularly come together to pool their beliefs. This is a common formulation for the best-of-n problem in swarm robotics, in which the population must identify the true hypothesis in a distributed manner [Valentini *et al.*, 2017].

We judge that consensus has been reached in a simulation experiment if 90% of the agents in the population have a probability for $\mathcal{H}_1$ greater than 0.9, in which case the simulation is terminated and the time to consensus recorded. In the case that consensus is not reached then a simulation is automatically terminated after $10,000$ iterations, which is for practical purposes then recorded as the consensus time. In this context, failure to reach consensus can occur for a number of reasons. This includes polarisation in which a subset of the population converges to a probability close to 1, while the remaining agents converge to a probability close to 0. However, it also includes the possibility that all agents converge on a probability close to 0.5. Although, there is agreement about the probability of $\mathcal{H}_1$ in this case, there is clearly a failure to effectively propagate the direct evidence that $\mathcal{H}_1$ holds provided to individual agents. In the following, for each set of parameter values, 100 independent runs of the simulation were carried out. Results are averaged over those 100 runs.

We compare the time to consensus for a combination of pooling and updating with that for evidential updating alone, across a range of values for ϵ, k and w. For all values of k and ϵ we find that there is an interval of values of w in which the combination of pooling & updating outperforms updating alone. Figure 2 illustrates this, showing the average time to consensus plotted against w for $k = 3$ and $\epsilon = 2\%$. The solid line is the consensus time for pooling & updating, while the dotted horizontal line is the consensus time for updating alone. The interval in which the solid line lies below the dotted line, therefore identifies the range of w values for which pooling provides some additional advantage when used in addition to updating from direct evidence. The characteristic shape of these curves suggests a number of metrics which may provide insight into the evidence propagation properties of SProdOp. We take *gain* to be the maximum gain in convergence time resulting from using pooling in addition to updating. *Robustness* corresponds to the length of the interval of w values for which pooling provides some benefit. This gives us an indication of the sensitivity of the choice of w for a particular combination of k and ϵ values. We also identify the optimal value of w, denoted w^*, resulting in the shortest time to consensus. Taken in conjunction with theorem 3.2 this quantifies the extent to which the optimal form of SProdOp should be evidence diluting, preserving or amplifying.

Figure 4 shows *gain*, w^* and *robustness* plotted against k and ϵ. Figure 4a suggests that there is an optimal value of k for any given evidence rate ϵ, and that for larger k values performance then declines. As might be expected *gain* decreases as the evidence rate increases for all k values, since for higher evidence rates updating from direct evidence dominates (see figure 4d). It should also be noted that for higher values of k there is increased sensitivity concerning the choice of w, with a narrower interval of values in which pooling & updating outperforms updating only. This can be seen in figure 4c in which *robustness* decreases with k. *Robustness* tends to in-

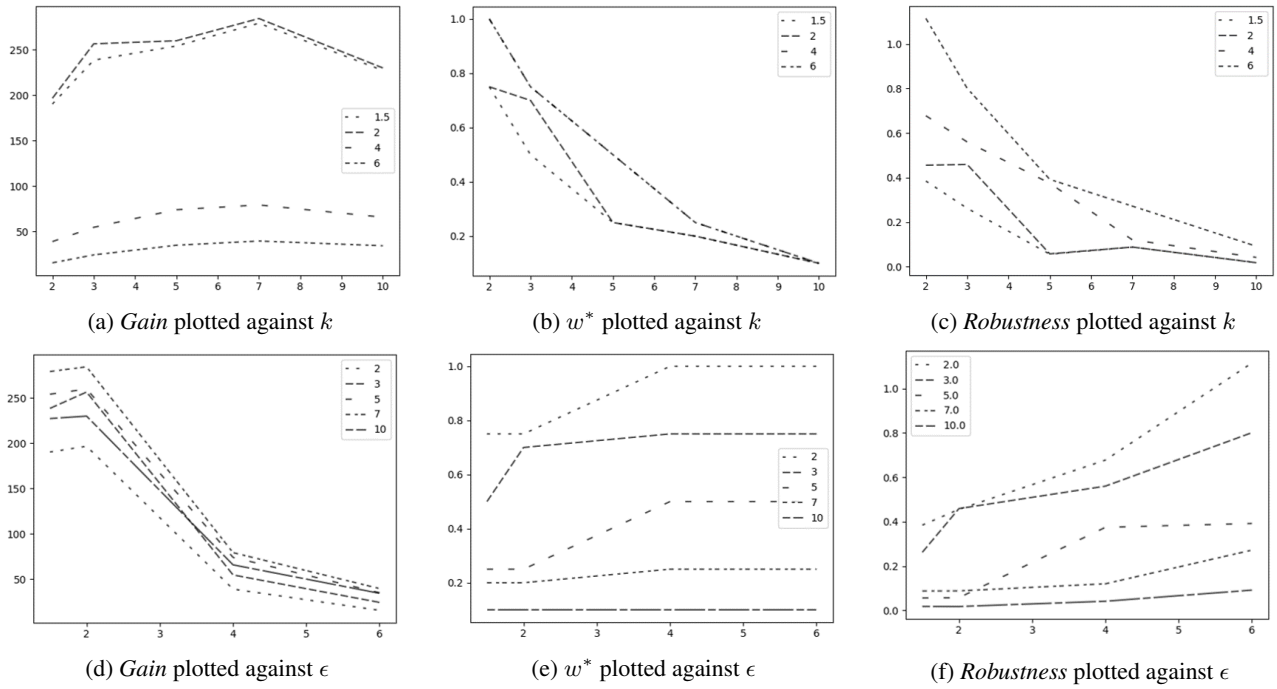

(a) *Gain* plotted against k (b) w^* plotted against k (c) *Robustness* plotted against k

(d) *Gain* plotted against ϵ (e) w^* plotted against ϵ (f) *Robustness* plotted against ϵ

Figure 4: *Gain*, *robustness* and w^* plotted for a number of different values of k and ϵ.

crease with evidence rate (figure 4f), although this is against a backdrop of decreasing gain. Balancing between *gain* and *robustness*, figures 4a and 4c would seem to provide a case for smaller pool sizes.

From figure 4b we can see that w^* decreases with k, so that the larger the size of the pool the more diluting the operator should be. In contrast w^* tends to increase with ϵ, although this increase levels off for higher evidence rates (see figure 4e). In general, all w^* values identified lie in the range $(0, 1]$ with $w^* = 1$ only for $k = 2$ and evidence rates greater than or equal to 4%. Hence, the results suggest that evidence diluting operators may be optimal for evidence propagation particularly when evidence rates are relatively low. Furthermore, for higher values of ϵ the *gain* from pooling is lower (figure 4d).

The variation of consensus time with w, as illustrated by figure 2, suggests that consensus is not achieved either for very low (close to 0) or for very high parameter values. For low evidence rates the system dynamics are significantly influenced by the properties of the pooling operator.We can gain insight into this by considering the fixed points of SProdOp i.e. those values of x for which $c(x, \ldots, x) = x$. For SProdOp with $w > 0$, the only fixed points are at $x = 0$, $x = 0.5$ and $x = 1$. For $w \geq \frac{1}{k}$ only 0 and 1 are stable fixed points and 0.5 is unstable. For $w < \frac{1}{k}$ all three fixed points are stable. Furthermore, as w tends to 0, SProdOp tends to a uniform value of 0.5, while as w tends to infinity SProdOp tends to a step function. This is illustrated in figure 3, which shows $c(x, x, x)$ for several values of w. From this we might hypothesise that for very high values of w, where SProdOp is extremely evidence amplifying, the population will quickly polarise with some holding probability 0 and others probability 1. On the other hand, for very low values of w, where SProdOp is extremely evidence diluting, we would expect the population to quickly converge on a shared probability of 0.5.

5 Conclusions

In this paper we have investigated the relationship between direct evidential updating and probability pooling in multi-agent systems. We have presented a Bayesian interpretation of LogOp which considers pooling to be a form of second order updating where the weight w_i quantifies the competency of the agent A_i by influencing the skew of an oracles' distribution on $P_{A_i}(\mathcal{H}_1)$ (theorem 2.1). We have then introduced the properties of evidence dilution, preservation and amplification as a way of gauging the evidence propagation properties of pooling operators. Furthermore, in this context we have shown that SProdOP is either diluting, preserving or amplifying depending only on the weight parameter w (theorem 3.2). This then provides a single parameter for varying the degree of evidence propagation of an operator, which we have then studied in detailed simulation experiments.

The simulation experiments presented in section 4 suggest that probability pooling using an appropriate operator can significantly improve the propagation of direct evidence across an agent population. As is to be expected this improvement is greatest for relatively low evidence rates. Results suggest that there is an optimal pool size k for any given evidence rate, but as k increases robustness decreases, implying that it is more difficult to identify a suitable value of w. For example, in the experiments in section 4, the optimal value of k is approximately 7 for all of the error rates investigated. However, for $k = 7$ *robustness* is relatively low i.e. between 0.1 and 0.3.

Hence, when taking account of both gain and robustness, our results suggest that lower values of k are preferable.

For all the error rates considered, the optimal value of w is between 0 and 1 and hence evidence diluting operators are more effective at evidence propagation for the type of agent-based system considered. This is partly due to the fact that SProdOp with large w tends towards a step function (figure 3), resulting in too fast convergence to either 0 or 1. In accordance with theorem 2.1, this corresponds to an assumption that agents are only moderately competent, since assuming higher levels of competence does not allow sufficient time for the population to take account of direct evidence before reaching an agreed belief.

Future work will attempt to generalise the results presented from two to multiple hypothesis i.e. to where opinions are probability distributions over $\mathcal{H}_1, \ldots, \mathcal{H}_n$ for $n \geq 2$. In this case, pooling operators are functions defined on a cartesian product of k $(n-1)$-dimensional probability simplexes. Furthermore, more detailed analytical studies should be conducted into the underlying dynamics of agent-based systems which combine both pooling and updating. This should then provide more insight into the simulation results described in section 4. We will also consider other operators such as linear pooling. Furthermore, we will consider non-Bayesian approaches such as imprecise probabilities and Dempster-Shafer theory.

Acknowledgements

This research was partially funded by an EPRSC PhD studentship as part of the Centre for Doctoral Training in Future Autonomous and Robotic Systems (grant number EP/L015293/1).

All underlying data is included in full within the paper.

References

[Abbas, 2009] A. E. Abbas. A Kullback-Leibler view of linear and log-linear pools. *Decision Analysis*, 6(1):25–37, 2009.

[Bordley, 1982] R. F. Bordley. A multiplicative formula for aggregating probability assessments. *Management science*, 28(10):1137–1148, 1982.

[DeGroot, 1974] M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

[Dietrich *et al.*, 2016] F. Dietrich, C. List, A. Hájek, and C. Hitchcock. Probabilistic opinion pooling. *Oxford Handbook of Probability and Philosophy, Oxford University Press, Oxford*, 2016.

[Douven and Kelp, 2011] I. Douven and C. Kelp. Truth approximation, social epistemology, and opinion dynamics. *Erkenntnis*, 75(2):271, 2011.

[Douven *et al.*, 2017] I. Douven, M. Ficker, P. Graham, D. Henderson, N. J. Pedersen, and J. Wyatt. Computational models of social epistemology. *Handbook of Social Epistemology*, 2017.

[Genest and Zidek, 1986] C. Genest and J. V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, pages 114–135, 1986.

[Hinton, 1999] G. E. Hinton. Products of experts. *Proceedings of the Ninth International Conference on Artificial Neural Networks (ICANN 99)*, pages 1–6, 1999.

[Hogarth, 1975] R. M. Hogarth. Cognitive processes and the assessment of subjective probability distributions. *Journal of the American statistical Association*, 70(350):271–289, 1975.

[Keeney and Raiffa, 1976] R. L. Keeney and H. Raiffa. *Decisions with multiple objectives: preferences and value trade-offs*. Harvard University, 1976.

[Morris, 1974] P. A. Morris. Decision analysis expert use. *Management Science*, 20(9):1233–1241, 1974.

[Osborne and Baldridge, 2004] M. Osborne and J. Baldridge. Ensemble-based active learning for parse selection. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, 2004.

[Parker and Zhang, 2009] C. A. C. Parker and H. Zhang. Cooperative decision-making in decentralized multiple-robot systems: The best-of-n problem. *IEEE/ASME Transactions on Mechatronics*, 14(2):240–251, 2009.

[Pennock and Wellman, 1999] D. M. Pennock and M. P. Wellman. Graphical representations of consensus belief. In *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, pages 531–540. Morgan Kaufmann Publishers Inc., 1999.

[Pettigrew, 2017] R. Pettigrew. Aggregating incoherent agents who disagree. *Synthese*, pages 1–40, 2017.

[Stone and others, 1961] M. Stone et al. The opinion pool. *The Annals of Mathematical Statistics*, 32(4):1339–1342, 1961.

[Valentini *et al.*, 2017] G. Valentini, E. Ferrante, and M. Dorigo. The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4:9, 2017.

[Winkler, 1968] R. L. Winkler. The consensus of subjective probability distributions. *Management Science*, 15(2):B–61, 1968.