Lawry, J., & Lee, C. (2020). Probability Pooling for Dependent Agents in Collective Learning. *Artificial Intelligence*, 288, [103371].
<https://doi.org/10.1016/j.artint.2020.103371>

Peer reviewed version

License (if available):
CC BY-NC-ND

Link to published version (if available):
[10.1016/j.artint.2020.103371](https://doi.org/10.1016/j.artint.2020.103371)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Elsevier at https://www.sciencedirect.com/science/article/pii/S0004370220301211?dgcid=rss_sd_all . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Probability Pooling for Dependent Agents in Collective Learning

Jonathan Lawry*, Chanelle Lee

Department of Engineering Mathematics, University of Bristol, Bristol BS8 1UB, UK

Abstract

The use of copulas is proposed as a way of modelling dependencies between different agents' probability judgements when carrying out probability pooling. This is combined with an established Bayesian model in which pooling is viewed as a form of updating on the basis of probability values provided by different individuals. Adopting the Frank family of copulas we investigate the effect of different assumed levels of comonotonic dependence between individuals, in the context of a collective learning problem in which a population of agents must reach consensus on which of two mutually exclusive and exhaustive hypotheses is true. In this scenario agents receive evidence from two sources; directly from the environment and also from other agents in the form of probability judgements. They then apply Bayesian updating to the former and probability pooling to the latter. We carry out multi-agent simulation experiments and show that optimal population level performance is obtained under the assumption of some degree of comonotonicity between agents, and consequently show that the standard assumption of agent independence is suboptimal. This is found to be particularly true of scenarios where there is a large amount of noise and very low amounts of direct evidence. Finally, we investigate dynamic environments in which the true state of the world changes and show that identifying the optimal level of agent dependency has an even greater effect on performance than for static environments in which the true state remains constant.

Keywords: Probability pooling, copulas, collective learning, dependent agents

1. Introduction

The assumption of independent sources of evidence or of independent experts is very common in information fusion and opinion pooling. For example, in Dempster-Shafer theory most combination operators model different sources of evidence as independent random sets [5], and the well-known logarithmic probability pooling

*Corresponding author.

Email address: j.lawry@bristol.ac.uk (Jonathan Lawry)

operator makes the inherent assumption that the experts' probability judgements are independent random variables [19]. Such assumptions bring clear practical benefits since they only require that we capture the beliefs and competence of individuals without needing to explicitly characterise what may be complex inter-relationships between them. It can also be argued that we should in fact seek out groups of independent experts, since dependency between experts results in redundancy when pooling opinions, meaning that a pool of dependent experts tends to provide less aggregated information than an equal-sized pool of independent experts [3]. In this paper we will present a Bayesian approach to probability pooling using copulas to capture dependence between agents [24]. Using multi-agent simulation experiments we will then show that in a particular distributed learning scenario, modelling positive dependence between individuals results in better performance at the macro or population level than can be achieved when agents are assumed to be independent.

In collective learning a population of agents repeatedly receives evidence about the state-of-the-world from two distinct sources; by interacting directly with the environment they obtain sensory data and then through interactions with other agents they learn about their peers' current beliefs and opinions. Douven and Kelp [6] consider the combination of these two processes from the perspective of social epistemology. For example, they argue that modern scientific investigation proceeds not only on the basis of experimentation but also through extensive dialogue and collaboration between scientists. They propose to model these two processes in terms of belief updating and belief aggregation respectively, and they then use agent simulations to investigate a scenario in which a population of individuals learn the true value of a noisy real-valued parameter by applying linear updating and the Heggelmann–Krause aggregation operator to combine estimates. Their results show that for this simple scenario, the population reaches a more accurate consensus about the true parameter value more quickly when combining belief updating and aggregation, than when only belief updating is used. An overview of similar studies using Heggelmann–Krause aggregation is given in [7].

There are a number of related studies in the artificial intelligence and swarm robotics literature using various different updating and pooling models. These mostly focus on a class of distributed learning problems called the best-of- n [31] in which a swarm of robots must identify the best out of n options on the basis of both sensor data and information exchanged during local interactions between individuals. For example, in [20] a robot swarm applied probabilistic pooling and updating based on negative evidence eliminating certain states as not being the best. Negative updating is also used in [18] but where beliefs are represented as sets of possible options and pairwise fusion corresponds to taking the intersection of two sets when they overlap and the union otherwise. In both cases we see that the combination of updating and pooling is more effective and accurate than if only updating alone is used. In general, collective learning has considerable potential across a range of application domains in robotics and autonomous systems. For example, in decentralised search and rescue a robot swarm could be used to identify

the location of casualties within a search area [26]. Another application of this kind is pollution treatment swarms which could be deployed after an oil spill and which would need to identify the region where there is highest concentration of pollutants [16].

Instead of the best-of- n problem we focus here on a simple binary learning problem, and apply a probabilistic approach to both updating and pooling along similar lines to that described in [19]. More specifically, we consider two mutually exclusive and exhaustive hypotheses, denoted $\mathcal{H}_1$ and $\mathcal{H}_2$, so that $\mathcal{H}_2$ is equivalent to $\sim \mathcal{H}_1$ and hence the belief of an agent A_i can be characterised by a real number $x_i \in [0, 1]$ indicating that $P_{A_i}(\mathcal{H}_1) = x_i$ and $P_{A_i}(\mathcal{H}_2) = 1 - x_i$. Given a pool of k agents $A_1, \dots, A_k$ then in its most generic form a probability pooling operator takes the k probability values as inputs and returns a single aggregate probability value.

Definition 1. *Pooling Operator*

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

A number of well-known probability pooling operators are given in the following example.

Example 2.

- *The Linear Operator:* $\pi(x_1, \dots, x_k) = \frac{\sum_{i=1}^k w_i x_i}{\sum_{i=1}^k w_i}$ for $w_i \geq 0$.
- *The Log-Linear Operator:*

$$\pi(x_1, \dots, x_k) = \frac{\prod_{i=1}^k x_i^{w_i}}{\prod_{i=1}^k x_i^{w_i} + \prod_{i=1}^k (1-x_i)^{w_i}}$$
 for $w_i \geq 0$.
- *The Product Operator:* This is a special case of the log-linear operator in which $w_i = \dots = w_k = 1$ so that $\pi(x_1, \dots, x_k) = \frac{\prod_{i=1}^k x_i}{\prod_{i=1}^k x_i + \prod_{i=1}^k (1-x_i)}$
- *The Scaled Product Operator:* This is a special case of the log-linear operator in which $w_i = \dots = w_k = w > 0$ so that $\pi(x_1, \dots, x_k) = \frac{(\prod_{i=1}^k x_i)^w}{(\prod_{i=1}^k x_i)^w + (\prod_{i=1}^k (1-x_i))^w}$ for $w_i \geq 0$.

The linear operator can be understood as being based on the veridical assumption, meaning that each agent A_i is viewed as a distinct random process generating the truth values of the two hypotheses according to the probability x_i [30, 13]. However, from the decision maker's perspective it is uncertain which process is being applied and she quantifies the probability that truth values are being generated according to A_i 's probability distribution, as being proportional to w_i . The linear pool then gives the expected truth value of $\mathcal{H}_1$ under these assumptions. The log-linear operator has been widely studied and a number of possible justifications for its use have been proposed; see [10] for an early review. From an information theoretic

perspective, it is known to be the operator which minimizes the Kullback-Leibler divergence between the aggregated probability distribution and the distributions in the pool [1]. It preserves Markov independence as is often required in probabilistic graphical models [27], and is consistent when renormalising incoherent probability distributions in that the same result is obtained if renormalisation is performed before or after pooling [28]. The log-linear operator belongs to a class of pooling operators referred to as supra-Bayesian [19]. These assume that pooling corresponds to a type of Bayesian updating in which evidence takes the form of probability values provided by the agents in the pool [32, 21, 22]. In this context, it makes the explicit assumption of agent independence. We will explore this general approach to pooling in more detail in section 2 and consider ways in which dependence between agents can be incorporated into the model. The product and scaled product operators are simplified special cases of the log-linear operator and as such make the same independence assumption. The product operator in particular has been widely discussed with applications in machine learning [11] and in management sciences [2]. It is also known to have desirable evidence preservation properties in which the result of updating given a single piece of evidence is the same if the updating takes place before or after pooling [4].

In the remainder of this section we will introduce a particular method for modelling dependence between random variables based on the notion of a copula. In section 2, we will then use this method to model the dependence between agents when pooling probabilities. A copula is a function which relates a joint probability distribution to its associated marginal distributions [24], and which satisfies the following properties.

Definition 3. *Copula*

An k -dimensional copula is a function $C : [0, 1]^k \rightarrow [0, 1]$ satisfying the following: $\forall \vec{y} \in [0, 1]^k$

- If $\exists i \in \{1, \dots, k\}$, $y_i = 0$ then $C(y_1, \dots, y_k) = 0$.
- If $\exists i \in \{1, \dots, k\}$, $\forall j \neq i$, $y_j = 1$ then $C(y_1, \dots, y_k) = y_i$.
- For any k -box $[a_1, b_1] \times \dots \times [a_k, b_k]$ where $[a_i, b_i] \subseteq [0, 1]$ for $i = 1, \dots, k$

$$\Delta_{a_k}^{b_k} \Delta_{a_{k-1}}^{b_{k-1}} \dots \Delta_{a_1}^{b_1} C(y_1, \dots, y_k) \geq 0$$

where for any function $h : [0, 1]^k \rightarrow [0, 1]$,

$$\Delta_{a_i}^{b_i} h(y_1, \dots, y_k) = h(y_1, \dots, y_{i-1}, b_i, y_{i+1}, \dots, y_k) - h(y_1, \dots, y_{i-1}, a_i, y_{i+1}, \dots, y_k)$$

By a well known result due to Sklar [29], it holds that any joint distribution is related to its marginal distributions by means of some copula, where the latter then represents the dependencies between the random variables. For instance, there is an established link between copulas and statistical measures of association such as Spearman's rho and also his dependence index Φ^2 (see [24] for an overview).

Theorem 4. *Sklar's Theorem [29]*

Let $z_1, \dots, z_k$ be random variables with joint cumulative distribution H and marginal cumulative distributions $F_1, \dots, F_k$. Then there exists a copula $C : [0, 1]^k \rightarrow [0, 1]$ such that;

$$H(z_1, \dots, z_k) = C(F_1(z_1), \dots, F_k(z_k))$$

From theorem 4 it follows that if H has density function h and F_i has density function f_i for $i = 1, \dots, k$ these densities are related according to;

$$h(z_1, \dots, z_k) = c(F_1(z_1), \dots, F_k(z_k)) \prod_{i=1}^k f_i(z_i)$$

where c is the copula density for C as given by;

$$c(y_1, \dots, y_k) = \frac{\partial^k C(y_1, \dots, y_k)}{\partial y_1 \dots \partial y_k}$$

The following result due to Fréchet and Hoeffding establishes bounds for copulas.

Theorem 5. *Fréchet-Hoeffding Inequality [24]*

Let $C : [0, 1]^k \rightarrow [0, 1]$ be a k -dimensional copula then $\forall \vec{z} \in [0, 1]^k$,

$$\max\left(\sum_{i=1}^k y_i - k + 1, 0\right) \leq C(y_1, \dots, y_k) \leq \min(y_1, \dots, y_k)$$

In general, for $k > 2$, only the upper bound is realisable since the lower bound is not a copula [24]. Notice however that the independent copula $C(y_1, \dots, y_k) = \prod_{i=1}^k y_i$ is a valid copula in this range for all values of k . The min copula characterises a particularly strong form of positive dependence between the random variables, referred to as comonotonicity. Intuitively, comonotonicity means that the random variables vary according to a monotonic relationship in which they all either increase or decrease together.

Definition 6. *The Frank Family of Copula*

$$C(y_1, \dots, y_k) = -\frac{1}{s} \ln \left(1 + \frac{\prod_{i=1}^k (e^{-sy_i} - 1)}{(e^{-s} - 1)^{k-1}} \right)$$

where $s > 0$. The Frank copula tends to the independent copula as s tends to 0 and to the minimum copula as s tends to infinity.

The Frank copula parameter s represents different degrees of comonotonicity between the variables, ranging from complete independence to full comonotonicity. In the following section we will exploit this to provide a basic one parameter model for agent dependence in probability pooling.

Figure 1: Two normal marginal distributions and the associated joint distribution generated using the Frank copula with $s = 8$.

Example 7. Let z_1 and z_2 be normally distributed random variables with standard deviation 1 and means 5 and 7 respectively. Assuming a Frank copula C we obtain the following copula density:

$$c(x, y) = \frac{s(1 - e^{-s})e^{-sx}e^{-sy}}{(e^{-s} - 1 + (1 - e^{-sx})(1 - e^{-sy}))^2}$$

The joint density of z_1 and z_2 is then given by:

$$h(z_1, z_2) = c(F_1(z_1), F_2(z_2))f_1(z_1)f_2(z_2)$$

Figure 1 shows a contour map of the resulting joint distribution when $s = 8$. Notice that in this case there is clear comonotonicity between the two random variables.

An outline of the remainder of the paper is as follows. Section 2 introduces the oracle or supra-Bayesian approach to probability pooling and describes how copulas can be used in this model to represent dependence between agents in the pool. In section 3 we apply supra-Bayesian pooling in conjunction with the Frank family of copulas to collective learning. The results of multi-agent simulation experiments are presented which suggest that assuming some comonotonic dependence between agents when pooling can significantly improve macro-level learning in scenarios where there is significant noise and very little evidence. Section 4 extends this study to dynamic environments whereby the true hypothesis changes at some time point during the simulation. This increases the difficulty of the collective learning problem since the population of agents must adapt to a new state-of-the-world after potentially already having reached consensus. Finally, section 5 gives some conclusions and discussion of possible future work.

2. Bayesian Probability Pooling

In this section we describe the so-called supra-Bayesian approach to probability pooling [32, 23, 21, 22]. This adopts a Bayesian perspective and views pooling as a type of conditional updating on the basis of evidence in the form of probability values provided by the agents in the pool. More specifically, 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 that $P_{A_i}(\mathcal{H}_1) = x_i$ for the agents $A_1, \dots, A_k$. Here the oracle is an abstract entity which we might choose to interpret in a number of different ways. For instance, O might be an independent arbitrator or decision maker 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 [12] and [17] referring to O as a ‘synthetic personality’ and a ‘supra Bayesian’ respectively. Furthermore, early work by [32] and [23] shows that from a Bayesian perspective, the pooling operator π can then be understood as O ’s posterior distribution determined along the following lines. Let the random variable X_i denote $P_{A_i}(\mathcal{H}_1)$ for $i = 1, \dots, k$ and suppose that O has a prior probability of $\mathcal{H}_1$, denoted $P_O(\mathcal{H}_1)$. This is then conditioned on the evidence $\bigwedge_{i=1}^k (X_i = x_i)$, representing the beliefs of the k agents, according to Bayes theorem as follows;

$$\begin{aligned} \pi(x_1, \dots, x_k) &= P_O(\mathcal{H}_1 | X_1 = x_1, \dots, X_k = x_k) \\ &= \frac{P_O(X_1 = x_1, \dots, X_k = x_k | \mathcal{H}_1) P_O(\mathcal{H}_1)}{P_O(X_1 = x_1, \dots, X_k = x_k | \mathcal{H}_1) P_O(\mathcal{H}_1) + P_O(X_1 = x_1, \dots, X_k = x_k | \mathcal{H}_2) P_O(\mathcal{H}_2)} \end{aligned}$$

If we further assume that O is a priori unbiased so that $P_O(\mathcal{H}_1) = \frac{1}{2}$ then the above can be simplified to:

$$\begin{aligned} \pi(x_1, \dots, x_k) &= \\ &= \frac{P_O(X_1 = x_1, \dots, X_k = x_k | \mathcal{H}_1)}{P_O(X_1 = x_1, \dots, X_k = x_k | \mathcal{H}_1) + P_O(X_1 = x_1, \dots, X_k = x_k | \mathcal{H}_2)} \end{aligned}$$

This expression can be further simplified if we also assume a form of symmetry between the two likelihood functions given $\mathcal{H}_1$ and given $\mathcal{H}_2$, as follows:

Definition 8. *Likelihood Symmetry [19]*

$$P_O(X_1 = x_1, \dots, X_k = x_k | \mathcal{H}_2) = P_O(X_1 = 1 - x_1, \dots, X_k = 1 - x_k | \mathcal{H}_1)$$

Hence, according to definition 8 learning that $\mathcal{H}_2$ holds provides the oracle O with the same information about the agents’ probabilities of $\mathcal{H}_1$, as learning $\mathcal{H}_1$ provides them about the agents’ probabilities of $\mathcal{H}_2$. Making this assumption we obtain:

$$\begin{aligned} \pi(x_1, \dots, x_k) &= \\ &= \frac{P_O(X_1 = x_1, \dots, X_k = x_k | \mathcal{H}_1)}{P_O(X_1 = x_1, \dots, X_k = x_k | \mathcal{H}_1) + P_O(X_1 = 1 - x_1, \dots, X_k = 1 - x_k | \mathcal{H}_1)} \end{aligned}$$

Based on this model we now propose to use a copula to represent the dependence between the agents' probabilities, i.e. between the random variables $X_1, \dots, X_k$, given $\mathcal{H}_1$. This is a variant of the approach proposed in [14] in which agents opinions were estimates of a numerical quantity, adapted to our current context in which the oracle is trying to determine the truth value of a hypothesis given a set of probability estimates for that hypothesis. More formally, suppose that given $\mathcal{H}_1$, X_i has distribution function F_i and density f_i for $i = 1, \dots, k$, then by Sklar's theorem there is a copula C with associated copula density c such that:

$$P_O(X_1 = x_1, \dots, X_k = x_k | \mathcal{H}_1) = c(F_1(x_1), \dots, F_k(x_k)) \prod_{i=1}^k f_i(x_i)$$

Hence, we obtain the following expression for the pooling function π :

$$\pi(x_1, \dots, x_k) = \frac{c(F_1(x_1), \dots, F_k(x_k)) \prod_{i=1}^k f_i(x_i)}{c(F_1(x_1), \dots, F_k(x_k)) \prod_{i=1}^k f_i(x_i) + c(F_1(1-x_1), \dots, F_k(1-x_k)) \prod_{i=1}^k f_i(1-x_i)}$$

Note that if we take C to be the independent copula then we are making the standard assumption of agent independence formalised as in the following definition.

Definition 9. *Independent Agents*

The agents $A_1, \dots, A_k$ are independent if,

$$P_O(X_1 = x_1, \dots, X_k = x_k | \mathcal{H}_j) = \prod_{i=1}^k P_O(X_i = x_i | \mathcal{H}_j)$$

Furthermore, if we assume independent agents as in definition 9 then since the copula density for the independent copula is a constant, the pooling function becomes:

$$\pi(x_1, \dots, x_k) = \frac{\prod_{i=1}^k f_i(x_i)}{\prod_{i=1}^k f_i(x_i) + \prod_{i=1}^k f_i(1-x_i)}$$

Following [19] we now assume that each random variable X_i is distributed according to a beta distribution with parameters a_i and b_i so that $f_i(x_i) \propto x_i^{a_i-1}(1-x_i)^{b_i-1}$. For a beta distribution with parameters $a_i > b_i$ 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. Negative skewness in this context, is arguably an indicator of the competence or reliability 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 competent agent would tend to be close to 1. For example, a competent agent may have received a significant amount of correct evidence and hence as a consequence of Bayesian updating their probability value for the true hypothesis will be close to 1. In view of this argument

Figure 2: $F_i(x_i)$ plotted against x_i for $w_i = 0.1$, $w_i = 1$, $w_i = 10$ and $w_i = 100$.

we propose to consider beta distributions for which $a_i = w_i$ and $b_i = 0$ where $w_i > 0$ for $i = 1, \dots, k$. In this case $F_i(x_i) = x_i^{w_i+1}$ and $f_i(x_i) = (w_i + 1)x_i^{w_i}$ and hence we obtain pooling functions of the following form.

$$\pi(x_1, \dots, x_k) = \frac{c(x_1^{w_1+1}, \dots, x_k^{w_k+1}) \prod_{i=1}^k x_i^{w_i}}{c(x_1^{w_1+1}, \dots, x_k^{w_k+1}) \prod_{i=1}^k x_i^{w_i} + c((1-x_1)^{w_1+1}, \dots, (1-x_k)^{w_k+1}) \prod_{i=1}^k (1-x_i)^{w_i}}$$

Notice, that for values of $x_i \in (0, 1)$, $F_i(x_i)$ is a strictly decreasing function of w_i , and since $F_i(0) = 0$ and $F_i(1) = 1$ for all $w_i > 0$, it follows that for $x_i < 0$, $\lim_{w_i \rightarrow \infty} F_i(x_i) = 0$ and $\lim_{w_i \rightarrow \infty} F_i(x_i) = 1$. In other words, in the case that $\mathcal{H}_1$ is true the assumption is that as $w_i \rightarrow \infty$ the agent A_i is so competent that they will always attribute probability 1 to $\mathcal{H}_1$. For example, figure 2 shows F_i for increasing values of w_i . Notice that making this assumption about the distributions of X_i in combination with agent independence leads us to recover the log-linear pooling operator as introduced in definition 1.

The next example illustrates the use of the Frank family of copulas to model comonotonic dependence between agents for a simple two agent pool. In particular, it compares the likelihood function given $\mathcal{H}_1$ for independent agents with the likelihood when there is comonotonicity between the agents' probability values as modelled by a Frank copula with $s = 5$.

Example 10. Consider the case in which we have a pool of 2 agents such that $P_{A_1}(\mathcal{H}_1) = x_1$ and $P_{A_2}(\mathcal{H}_1) = x_2$. Adopting the above Bayesian model we have that the oracle's likelihood of agents A_1 and A_2 holding these beliefs given $\mathcal{H}_1$ is:

$$P_O(X_1 = x_1, X_2 = x_2 | \mathcal{H}_1) = c(x_1^{w_1+1}, x_2^{w_2+1})(w_1 + 1)x_1^{w_1}(w_2 + 1)x_2^{w_2}$$

and hence:

$$\pi(x_1, x_2) = \frac{c(x_1^{w_1+1}, x_2^{w_2+1})x_1^{w_1}x_2^{w_2}}{c(x_1^{w_1+1}, x_2^{w_2+1})x_1^{w_1}x_2^{w_2} + c((1-x_1)^{w_1+1}, (1-x_2)^{w_2+1})(1-x_1)^{w_1}(1-x_2)^{w_2}}$$

Then using a Frank copula, figure 3 shows contour plots of the above likelihood function for different weight values w_1 and w_2 and also for different values of the Frank parameter s . In these plots we can see the combined effects of the assumption of comonotonic dependence between agents for positive values of s and the beta distribution model of agent competence described above. For example, given uniform competence weights $w_1 = w_2 = 1$, figure 3a shows the contour plot for the likelihood assuming independent agents, i.e. $s = 0$, while figure 3a shows the likelihood when $s = 5$. In both cases the function is skewed towards $(1, 1)$ as is consistent with the linear competency model where $f_1(x) = f_2(x) \propto x$. In figure 3a we also see the effect of comonotonicity since the likelihood function is concentrated close to the line $x_1 = x_2$. Figures 3c and 3d show contour plots of the likelihood for independent and dependent agents respectively, when $w_1 = 0.5$ and $w_2 = 3$ i.e. when A_2 is considerably more competent than A_1 . In both cases we see that the likelihood is now skewed towards higher values of x_2 , this then being combined with dependence between X_1 and X_2 in figure 3d.

Copulas are of course not the only way of modelling agent dependence in this context. Work by French [9], Lindley [21, 22] and Clemen [3] reformulated supra-Bayesian pooling in terms of log-odds for the different agents, and then used joint normal distributions to capture the dependencies between them. More specifically, for agent A_i we define:

$$Q_i = \ln \left(\frac{X_i}{1 - X_i} \right)$$

Pooling a specific set of log-odds $q_i = \ln \left(\frac{x_i}{1-x_i} \right)$ for $i = 1, \dots, k$, then involves the oracle O evaluating the following log likelihood ratio:

$$\ln \left(\frac{P_O(Q_1 = q_1, \dots, Q_k = q_k | \mathcal{H}_1)}{P_O(Q_1 = q_1, \dots, Q_k = q_k | \mathcal{H}_2)} \right)$$

It is then assumed that $P_O(Q_1, \dots, Q_k | \mathcal{H}_1)$ and $P_O(Q_1, \dots, Q_k | \mathcal{H}_2)$ are both multivariate normally distributed with different means but the same covariance matrix Σ , where the latter encodes any dependencies between the agents $A_1, \dots, A_k$. This approach is natural and interesting but the model does require that we have sufficient knowledge of agent dependencies so as to completely specify Σ . Instead, in this paper we focus on the copula approach as a way of modelling agent dependence using a small number of parameters. In particular, in the next section we will use simulation experiments to investigate the effect on collective learning of assuming different levels of agent comonotonicity as modelled by the single Frank parameter s .

(a) Likelihood function when $s = 0$ and $w_1 = w_2 = 1$.

(b) Likelihood function when $s = 5$ and $w_1 = w_2 = 1$.

(c) Likelihood function when $s = 0$, $w_1 = 0.5$ and $w_2 = 3$.

(d) Likelihood function when $s = 5$, $w_1 = 0.5$ and $w_2 = 3$.

Figure 3: Contour plots of the likelihood functions $P_O(X_1 = x_1, X_2 = x_2 | \mathcal{H}_1)$ for the Frank copula with different s values and agent weights w_1 and w_2 .

3. Collective Learning with Dependent Agents

In this section we use finite time agent-based simulation to investigate collective learning where agents are assumed to have different degrees of dependence as modelled using Frank copulas with different parameter values. Agents are assumed to receive information of two distinct types; there is direct evidence obtained through direct interaction with the environment. For example, this may be as the result of conducting an experiment or simply from sensor data. Here we assume that evidence takes the form of an assertion that either $\mathcal{H}_1$ or $\mathcal{H}_2$ is true. Evidence is also assumed to be rare and this is modelled by means of an evidence rate $\rho \in [0, 1]$, corresponding to the probability that an individual agent will receive evidence in any given time step. Furthermore, we assume that the process of obtaining evidence is inherently noisy with a fixed probability $\epsilon \in [0, 0.5]$ of receiving false evidence. As well as direct evidence agents also learn from each other by applying the probability pooling operator described in section 2. At each time step a small number of agents

are selected at random, modelling random interactions in a well-mixed population [25], and their beliefs are aggregated with each agent in the pool then adopting the consensus belief. For the experiments described below we assume a population of 100 agents and that at each time step a random 5% of the population pool their beliefs.

Pooling is conducted by applying a Bayesian pooling operator with a Frank copula and with a small weighting toward ignorance as represented by a uniform distribution over the two hypotheses. More specifically, each agent in the pool updates their probability of $\mathcal{H}_1$ to:

$$\lambda \frac{1}{2} + (1 - \lambda) \pi(x_1, \dots, x_k)$$

where λ is a dampening term typically taking a low value in the range $[0, 0.5]$. The inclusion of a small weighting for ignorance allows for more robust convergence by preventing agents reaching absolute certainty, i.e. probability values of either 0 or 1, and then being unable to update when they receive new evidence which conflicts with their current beliefs. This is particularly important in the context of a dynamic environment in which the truth value of the various hypotheses may suddenly change. Such environments will be investigated in section 4.

Evidential updating is performed in the standard Bayesian manner as outlined in definition 11 below.

Definition 11. *Bayesian Evidential Updating*

For $E \in \{\mathcal{H}_1, \mathcal{H}_2\}$ we have that:

$$P_{A_i}(\mathcal{H}_j|E) = \frac{P_{A_i}(E|\mathcal{H}_j)P_{A_i}(\mathcal{H}_j)}{P_{A_i}(E|\mathcal{H}_1)P_{A_i}(\mathcal{H}_1) + P_{A_i}(E|\mathcal{H}_2)P_{A_i}(\mathcal{H}_2)}$$

In this case the likelihood $P_{A_i}(E|\mathcal{H}_j)$ captures the agent's judgement about the reliability of the evidence E .

A simplified version of definition 11 is where the likelihood function has the following form:

$$P_{A_i}(E|\mathcal{H}_j) = \begin{cases} 1 - \alpha : E = \mathcal{H}_j \\ \alpha : \text{otherwise} \end{cases}$$

Here $\alpha \in [0, 0.5]$ is a parameter quantifying the agent's general level of doubt in evidence. Hence, for $P_{A_i}(\mathcal{H}_1) = x_i$ we have that:

$$P_{A_i}(\mathcal{H}_1|E = \mathcal{H}_1) = \frac{(1 - \alpha)x_i}{(1 - \alpha)x_i + \alpha(1 - x_i)}$$

and

$$P_{A_i}(\mathcal{H}_1|E = \mathcal{H}_2) = \frac{\alpha x_i}{\alpha x_i + (1 - \alpha)(1 - x_i)}$$

In collective learning, the effect of α tends to be to slow convergence by reducing the change in probability values resulting from any given update. In this respect the overall effect of increasing α is similar to that of lowering the evidence rate ρ . In the sequel we will assume that agents do not doubt the evidence they receive so that $\alpha = 0$, and we will focus instead on understanding the effect of varying ρ .

While in this paper we adopt the standard Bayesian updating model given in definition 11 there are alternatives. For example, [8] considers several probability updating methods some of which emphasise the explanatory power of different hypotheses. Furthermore, [8] also introduces an evolutionary approach to identify the optimal evidential updating method in a particular collective learning setting.

We now present simulation results supporting the broad claim that in situations where evidence is very sparse, i.e. very low ρ , and where there is significant noise present, i.e. high ϵ , then it is optimal to assume a level of comonotonic dependence between agents when pooling, i.e. Frank parameter $s > 0$. Furthermore, system level performance assuming such dependence can be significantly better than if independent agents are assumed.

Throughout this section we assume without loss of generality that $\mathcal{H}_1$ is the true hypothesis, and hence the average value of $P_{A_i}(\mathcal{H}_1)$ across the population provides an appropriate measure of system level performance. Initially we will assume that all agents A_i are equally reliable as quantified by weight $w_i = 1$. Each experiment, as characterised by parameter values s , ρ , ϵ , and λ , is run 100 times and results are then averaged across the different runs with error bars showing 90% percentiles. In all cases we initialize by allocating the agents' probability values at random.

Now consider figure 4 where $\epsilon = 0.3$, $\rho = 0.001$ and $\lambda = 0.01$. Figure 4a shows the average values of $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations plotted against the Frank parameter s . We see that performance is optimal for Frank's parameter values around $s = 4$ (average $P_{A_i}(\mathcal{H}_1) > 0.9$), and if independent agents are assumed then performance is worse (average $P_{A_i}(\mathcal{H}_1) < 0.6$). Figure 4b then shows time series trajectories for the $s = 0$ and $s = 4$ agent dependency models. Here we see that the assumption of agent independence means that true evidence fails to propagate across the population in this sparse evidence, high noise scenario. We can unpick these results further by considering histograms of average $P_{A_i}(\mathcal{H}_1)$ values at 1500 iterations, across the 100 simulation runs of each experiment, as shown in figure 5. For the independent model, figure 5a, there is an almost even distribution across the runs of experiments in which agents reach the consensus that $P_{A_i}(\mathcal{H}_1) \approx 0$ and experiments where the consensus is $P_{A_i}(\mathcal{H}_1) \approx 1$. This is close to the distribution of runs obtained when there is no evidence i.e. when $\rho = 0$. In contrast, if $s = 4$ as in figure 5b then the distribution of runs is highly skewed towards those in which the consensus is $P_{A_i}(\mathcal{H}_1) \approx 1$ i.e. in the vast majority of runs the population correctly identifies $\mathcal{H}_1$ as the true hypothesis.

An extensive study of the parameter space was conducted by varying values of s , ρ , ϵ and λ , and the results are shown as heat maps in figures 6 and 7. In

(a) Average value of $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations plotted against the Frank parameter s .

(b) Trajectories showing average value of $P_{A_i}(\mathcal{H}_1)$ against time for $s = 0$ (blue line) and $s = 4$ (black line).

Figure 4: Average values of $P_{A_i}(\mathcal{H}_1)$ for different values of the Frank parameter s . The error probability is $\epsilon = 0.3$, the evidence rate is $\rho = 0.001$ and $\lambda = 0.01$.

each heat map two parameters are varied while the others are fixed. Each cell then shows the average value of $P_{A_i}(\mathcal{H}_1)$ across agents after 1500 iterations and then averaged across 100 independent runs with the same parameter settings. Figure 6 focuses on the effect of evidence rate in a setting where there is a high amount of noise i.e. $\epsilon = 0.3$. In figure 6b we consider the case where the evidence rate is very low, with ρ ranging between 0.0005 and 0.005, and for different degrees of comonotonic agent dependence as characterised by the Frank parameter s . Good performance is obtained for a range of positive dependence values centred around $s = 4$ or $s = 5$, and as the evidence rate increases the robustness to the choice of s value also increases. For example, when $\rho = 0.001$ good performance is only achieved for $s = 4$ or $s = 5$, while when $\rho = 0.005$ good performance is achieved for a range of s values between 2 and 5. This increasing trend in robustness to the choice of agent dependence degree continues for higher evidence rates as shown in figure 6a. For instance, at $\rho = 0.05$ good performance is achieved for all $s \leq 6$. In summary, optimal performance is always achieved for $s = 4$ or $s = 5$ but as the evidence rate increases this level of performance can be achieved for a much broader range of values of the Frank parameter. For very low evidence rates, i.e. $\rho < 0.005$, the assumption of independent agents is clearly sub-optimal, whereas for higher rates, i.e. $\rho > 0.015$, optimal performance can be obtained for a range of dependency models, including for independent agents.

Figure 7 shows the effects of varying the Frank parameter s , in conjunction with either the error probability ϵ or the dampening parameter λ . In both cases we assume an evidence rate of $\rho = 0.001$. In figure 7a we see that the best performance for all error probability values is achieved for $s = 4$ or $s = 5$ although for low error the optimal performance can be obtained for a broader range of s values. Furthermore, good performance with values of average $P_{A_i}(\mathcal{H}_1)$ close to 1 can be achieved for

(a) Histogram for $s = 0$

(b) Histogram for $s = 4$

Figure 5: Histograms showing the distribution of $P_{A_i}(\mathcal{H}_1)$ at 1500 iterations across 100 independent runs. Here $\epsilon = 0.3$, $\rho = 0.001$ and $\lambda = 0.01$.

error probabilities up to and including 0.3, when $s = 4$ or $s = 5$. Again we see that pooling under the assumption of independent agents, i.e $s = 0$, does not perform well in these high error, low evidence settings.

Figure 7b shows the case when $\epsilon = 0.3$ and $\rho = 0.001$ under different values of s and the dampening parameter λ . Broadly speaking performance is best for low values of λ , i.e. $\lambda < 0.05$. In addition, as λ increases the optimal value of s slightly decreases, although in all cases considered the independence assumption is sub-optimal.

We now consider the effect of changing agent reliability on the optimal agent dependency value. To do so we continue to assume that all agents are equally reliable but investigate different levels of reliability. In other words, we take $w_i = w > 0$ for all agents A_i and then run simulation experiments for different values of w . Figure 8 shows average values of $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations plotted against s for 5 different reliability weight values i.e $w = 0.5, 0.75, 1, 1.25$ and 1.5 . Once again we focus on the low evidence, high error scenario where $\rho = 0.001$ and $\epsilon = 0.3$. Here there is a clear trend that as w increases so does the optimal value of s . In other words, the more reliable that the agents are assumed to be, then the more comonotonicity we should assume there to be between them when pooling probability values. Indeed this result is perhaps intuitive given the interpretation of reliability weights proposed in section 2. More specifically, here we are assuming that given $\mathcal{H}_1$ is true, the probability distributions of $P_{A_i}(\mathcal{H}_1)$ for all agents A_i are identical, corresponding to $F_i(x_i) = x_i^{w+1}$. As discussed in section 2 and illustrated in figure 2 this means that as w increases then the closer we are assuming the values of $P_{A_i}(\mathcal{H}_1)$ to be clustered to 1, across the population of agents. Now if this were to be the case then we would also expect for these values to be increasingly dependent as w increases. Consequently, an assumption of high reliability for agents should naturally be combined with an assumption of strong dependence between their probability values.

(a) Average $P_{A_i}(\mathcal{H}_1)$ for varying s and ρ between 0.005 and 0.05.

(b) Average $P_{A_i}(\mathcal{H}_1)$ for varying s and ρ between 0.0005 and 0.005.

Figure 6: Heat maps showing average values of $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations for different values of the Frank parameter s and for different evidence rates ρ . In this case $\epsilon = 0.3$ and $\lambda = 0.01$.

4. Learning in a Changing Environment

In the simulation experiments described in section 3 it was assumed that the state-of-the-world remains constant. More specifically, we assumed that $\mathcal{H}_1$ was the true hypothesis throughout, and that this was reflected in the evidence received by the agents at any time, all be it with some probability of error. However, in general collective learning will take place in dynamic environments in which the truth values of hypotheses change over time. This poses new challenges in which the requirement to reach a definitive consensus must be balanced against the need for flexibility and a degree of open-mindedness. In this section we investigate the effectiveness of the proposed Bayesian pooling model in dynamic environments and consider the role of dependency assumptions in this context. More specifically, we will employ the same simulation set-up as described in section 3 but where the truth-values of the two hypotheses switch suddenly at a single time point; that is $\mathcal{H}_2$ is the true hypothesis upto time $t = 700$ and $\mathcal{H}_1$ is true after that time. We once again focus on low evidence and high noise scenarios and show that an assumption of moderate comonotonic dependence between agents is once again optimal. Furthermore, we will highlight the importance of the dampening term λ as a means of ensuring that agents preserve a small element of doubt in the beliefs which then allows them to adapt to new evidence which is inconsistent with what they currently believe to be the most probable hypothesis.

Figure 9 shows heat maps of average values of $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations for varying values of s and λ . Since the switch in state is from $\mathcal{H}_1$ being false to true at $t = 700$, then the average $P_{A_i}(\mathcal{H}_1)$ at $t = 1500$ remains a suitable performance metric, with a value of 1 being optimal. In both heat maps we assume uniform competence across agents with $w_i = 1$ for all A_i , and the evidence rate is $\rho = 0.001$. Figures 9a and 9b show the results for $\epsilon = 0.3$ and $\epsilon = 0.2$ respectively. In both

(a) Average $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations for varying s and ϵ . In this case $\lambda = 0.01$.

(b) Average $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations for varying s and λ . In this case $\epsilon = 0.3$.

Figure 7: Heat maps showing average values of $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations for different values of the Frank parameter s and either for different error rates ϵ or different values of λ . In this case $\rho = 0.001$.

Figure 8: Average $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations with $\rho = 0.001$ and $\epsilon = 0.3$. The lines are for weights values $w = 0.5$ (green), 0.75 (red), 1 (black), 1.25 (blue), and 1.5 (orange)

cases we see that optimal performance is when $s = 3$ or $s = 4$, and $\lambda = 0.04$ or $\lambda = 0.05$. Furthermore, performance is poor under the assumption of independent agents, i.e. for $s = 0$, and also if there is no dampening term i.e. for $\lambda = 0$.

(a) Average $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations for varying s and λ . In this case $\epsilon = 0.3$.

(b) Average $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations for varying s and λ . In this case $\epsilon = 0.2$.

Figure 9: Heat maps of average $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations for varying s and λ with a switch in the true state the world at 700 iterations. In this case $\rho = 0.001$

Figure 10 focuses on the case when $\epsilon = 0.2$, $\lambda = 0.05$ and $\rho = 0.001$. More specifically, figure 10a shows average values of $P_{A_i}(\mathcal{H}_1)$ at $t = 1500$ plotted against s . This suggests that optimal performance is at $s = 4$ but also that there is a narrower range of s values at which performance is good compared to when the environment is static; see figure 4a. Figure 10b shows the time series trajectories for when $s = 0$ and $s = 4$. For $s = 0$ there is very little propagation of evidence from the start and the adaptation to the switch in the state of the world is minimal. In contrast, for $s = 4$ we see that upto $t = 700$ the population is gradually converging on the consensus that $P_{A_i}(\mathcal{H}_1) \approx 0$, but after $t = 700$ there is adaptation to the new state of the world and the population then converges on a consensus where $P_{A_i}(\mathcal{H}_1) \approx 1$. Figure 11 shows the same results from a different perspective, as histograms of the distribution of population average values of $P_{A_i}(\mathcal{H}_1)$ at $t = 1500$ across 100 independent runs of the simulation experiments. In figure 11a we see that the assumption of independence between agents leads to a bimodal distribution where for each run there is either consensus that $P_{A_i}(\mathcal{H}_1) \approx 0$ or $P_{A_i}(\mathcal{H}_1) \approx 1$ and with some skew to the former. On the other hand, figure 11b shows that assuming dependent agents with $s = 4$ results in a distribution of probability values heavily skewed towards 1.

The importance of the dampening parameter λ in dynamic environments is illustrated by the comparison between figure 11b where $\lambda = 0.05$ and figure 12 where $\lambda = 0$. In the latter the distribution of probability values is heavily skewed towards 0 suggesting strong convergence to the initial state of the world but then little or no adaptation when the state changes. We hypothesise that this failure to adapt is a result of the population reaching a consensus that $\mathcal{H}_2$ is certainly true, i.e. with

probability 1, before $t = 700$. This means that any subsequent evidential updating given new evidence that it is in fact $\mathcal{H}_1$ that is true will not result in any change to probability values. The dampening parameter acts to prevent complete convergence of probability values to either 0 or 1 by adding a degree of uncertainty to the pooling. This then allows for continued probability updating throughout the simulation which in turn facilitates adaptation when the environment changes.

(a) Average value of $P_{A_i}(\mathcal{H}_1)$ after 1500 iterations plotted against the Frank parameter s .

(b) Trajectories showing average value of $P_{A_i}(\mathcal{H}_1)$ against time for $s = 0$ (blue line) and $s = 4$ (black line).

Figure 10: Average values of $P_{A_i}(\mathcal{H}_1)$ for different values of the Frank parameter s . The error probability is $\epsilon = 0.2$, evidence rate is $\rho = 0.001$, $\lambda = 0.05$ and the switch in the true state of the world occurs at iteration 700.

(a) Histogram for $s = 0$

(b) Histogram for $s = 4$

Figure 11: Histograms showing the distribution of $P_{A_i}(\mathcal{H}_1)$ at 1500 iterations across 100 independent runs. Here $\epsilon = 0.2$, $\rho = 0.001$, $\lambda = 0.05$ and the switch in the true state of the world occurs at iteration 700.

Figure 12: Histogram showing the distribution of $P_{A_i}(\mathcal{H}_1)$ at 1500 iterations across 100 independent runs when $s = 4$. Here $\epsilon = 0.2$, $\rho = 0.001$, $\lambda = 0$ and the switch in the true state of the world occurs at iteration 700.

5. Discussions and Conclusion

In this paper we have proposed the use of copulas in probabilistic opinion pooling as part of a general Bayesian approach. By means of agent-based simulation experiments we have then studied the application of this model to a collective learning problem in which a population of agents aim to discover which of two mutually exclusive and exhaustive hypotheses, $\mathcal{H}_1$ or $\mathcal{H}_2$, is true. Over time agents both receive evidence directly from the environment and pool their beliefs with those of other agents. In this context the results suggest that optimal performance is obtained under the assumption of some comonotonic dependence between agents when pooling, and in particular we see that the assumption of independent agents is suboptimal.

The use of families of copulas in collective learning provides a model of dependence with a low number of parameters, leaving open the possibility of optimizing the dependency between agents at the macro-level or possibly even locally in real-time. There are however some disadvantages to the approach which may limit both scalability and generalizability. For the former we should note that even for the Frank copula there is no simple general form for the copula density at any pool size. Consequently, there is no general form of the pooling operator which agents can deploy no matter how many agents are currently in communication with them. Of course, we can in principle determine an operator for any given pool size although the form of the copula density rapidly becomes complex. This problem may be less important if the optimal pooling size is relatively small as is suggested in [19]. Another limitation of parameterized families of copula is that they only tend to model a particular type of dependence, i.e. comonotonic dependence in the case of the Frank family. In more complex collective learning scenarios the dependence between individuals may well take a much richer and varied form. Of course, we know from Sklar’s theorem that we can in principle capture any dependence between ran-

dom variables using a copula, but this is not practical without using copula families. Nonetheless, the results described above suggest that the use of even simple dependency models can potentially give better performance than that obtained under the independent agents assumption.

Although we have proposed a model of probability pooling in which operators are defined with reference to a supra-Bayesian or an oracle, this is an abstraction and there is no need in practice that such an entity actually exists. For AI agents, if a pooling operator is pre-defined for all individuals in the population and if at the time of pooling agents transmit all the required information, e.g. probability values and competency information, then the learning process can take place in a decentralised manner. Nonetheless, for an entirely decentralised approach it would arguably be more natural that agents' competency weights should be calculated locally and dynamically as the knowledge and experience of different individuals evolves over time, and also that a pool of agents should be able to collectively evaluate an appropriate level of dependency between their probability judgements. This is a potentially rich area of future research as outlined below.

Building on the work in this paper there are a number of clear avenues for future research. The simulation experiments could be extended to include agents with differing levels of competence as quantified by different weights w_i . These could even be dynamic so as to, for example, depend on the amount of direct evidence that an agent has received upto the current time. Another focus would be for agents to use their experience of pooling over time to determine their own approximate model of agent dependence across the population. For instance, [14] proposes methods for estimating the Frank parameter s based on data about agent probabilities over time. These could perhaps be adapted to a real-time setting in which agents gradually build up data on the probability judgements of others. Finally, another avenue for future research could be to extend the approach to situations with more than two hypotheses.

Acknowledgements

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

References

- [1] A. E. Abbas, (2009), 'A Kullback-Leibler View of Linear and Log-linear Pools', *Decision Analysis*, 6(1), 25-37.
- [2] R. F. Bordley, (1982), A Multiplicative Formula for Aggregating Probability Assessments, *Management Science*, 28(10), 1137-1148.

- [3] R. T. Clemen, R. L. Winkler, (1985), 'Limits for the Precision and Value of Information from Dependent Sources', *Operational Research*, 33, 427-442
- [4] F. Dietrich, C. List, (2016), 'Probabilistic Opinion Pooling', *Oxford Handbook of Probability and Philosophy*, Oxford University Press, (Eds. A. Hajek, C. Hitchcock), Oxford
- [5] D. Dubois, W. Liu, J. Ma, H Prade, (2016), 'The Basic Principles of Uncertain Information Fusion. An Organised Review of Merging Rules in Different Representation Frameworks', *Information Fusion*, 32, 12–39.
- [6] I. Douven, C. Kelp, (2011), Truth Approximation, Social Epistemology, and Opinion Dynamics, *Erkenntnis*, 75, 271
- [7] I. Douven, (2017), 'Computational Models of Social Epistemology', *Handbook of Social Epistemology*, (Eds M. Ficker, P. Graham, D. Henderson, N. J. Pedersen, J. Wyatt).
- [8] I. Douven, (2019), 'Optimizing Group Learning: An Evolutionary Computing Approach', *Artificial Intelligence*, 275, 235 - 251.
- [9] S. French, (1979), 'Consensus of Opinion', *European Journal of Operational Research*, 7, 332-340
- [10] C. Genest, J. V. Zidek, (1986), 'Combining Probability Distributions: A Critique and an Annotated Bibliography', *Statistical Science*, 114-135.
- [11] G. E. Hinton, (1999), 'Products of Experts', *Proceedings of the Ninth International Conference on Artificial Neural Networks (ICANN 99)*, 1-6.
- [12] R. M. Hogarth, (1975), 'Cognitive Processes and the Assessment of Subjective Probability Distributions', *Journal of the American Statistical Association*, 70(350), 271-289.
- [13] R. A. Jacobs, (1995), 'Methods For Combining Experts' Probability Assessments', *Neural Computation*, 7, 867-888.
- [14] M. N. Jouini, R. T. Clemen, (1996), 'Copula Models for Aggregating Expert Opinions', *Operations Research*, 44(4), 444-457.
- [15] J. M. Kahn, (2004), 'A Generative Bayesian Model for Aggregating Experts' Probabilities', *Proceedings of UAI*, 301-308
- [16] N. Kakalis, Y. Ventikos, (2008), 'Robotic swarm concept for efficient oil spill confrontation', *Journal of Hazardous Material*, 32, 12–39.
- [17] R. L. Keeney, H. Raiffa, (1976), *Decisions with Multiple Objectives: Preferences and Value Trade-offs*, Harvard University.

- [18] J. Lawry, M. Crosscombe, D. Harvey, (2019), ‘Epistemic Sets Applied to Best-of-n Problems’, *European Conference on Symbolic and Quantitative Approaches with Uncertainty*, Springer, 301-312
- [19] C. Lee, J. Lawry, A. Winfield, (2018), ‘Combining Opinion Pooling and Evidential Updating for Multi-agent Consensus’, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI-18)*, 347-353.
- [20] C. Lee, J. Lawry, A. Winfield, (2018), ‘Negative Updating Combined with Opinion Pooling in the Best-of-n Problem in Swarm Robotics’, *Proceedings of the International Conference on Swarm Intelligence (ANTS18)*, Springer, 97-108.
- [21] D. V. Lindley, (1985), ‘Reconciliation of Discrete Probability Distributions’, *Bayesian Statistics 2*, North-Holland, Amsterdam
- [22] D. V. Lindley, (1988), ‘The Use of Probability Statements’, In *Accelerated Life Testing*, North-Holland, Amsterdam.
- [23] P. A. Morris, (1974), ‘Decision Analysis Expert Use’, *Management Science*, 20(9), 1233-1241.
- [24] R. B. Nelson, (2005), *An Introduction to Copulas*, Springer Science and Business Media.
- [25] C. A. C. Parker, H. Zhang, (2009), ‘Cooperative decision-making in decentralized multiple-robot systems: The best-of-n problem’, *IEEE/ASME Transactions on Mechatronics*, 14(2), 240–251.
- [26] D. Peleg, (2005), ‘Distributed coordination algorithms for mobile robot swarms: new directions and challenges’, *Proceedings of IWDC 2005*, LNCS, 3741, 1–12, Springer, Heidelberg.
- [27] D. M. Pennock, M. P. Wellman, (1999), ‘Graphical representations of consensus belief’, In *Proceedings of the Fifteenth conference on Uncertainty in Artificial Intelligence*, 531–540. Morgan Kaufmann Publishers Inc.
- [28] R. Pettigrew, (2017), ‘Aggregating Incoherent Agents who Disagree’, *Synthese*, 1–40.
- [29] A. Sklar, (1973), ‘Random Variables, Joint Distribution and Copulas’, *Kybernetika*, Vol. 9, pp449-460.
- [30] M. Stone, (1961), ‘The Opinion Pool’, *The Annals of Mathematical Statistics*, 32(4), 1339-1342.
- [31] G. Valentini, E. Ferrante, M. Dorigo, (2017), ‘The Best-of-n Problem in Robot Swarms: Formalization, State of the Art, and Novel Perspectives’, *Frontiers in Robotics and AI*, 4:9.

- [32] R. L. Winkler, (1968), ‘The Consensus of Subjective Probability Distributions’,
Management Science, 15(2), B61-B75.