# Imprecise Fusion Operators for Collective Learning

Zixuan Liu[1], Michael Crosscombe[1] and Jonathan Lawry[1]

Department of Engineering Mathematics, University of Bristol,
Bristol BS8 1UB, United Kingdom
{z.liu, m.crosscombe, j.lawry}@bristol.ac.uk

## Abstract

A decentralised collective learning problem is investigated in which a population of agents attempts to learn the true state of the world based on direct evidence from the environment and belief fusion carried out during local interactions between agents. A parameterised fusion operator is introduced that returns beliefs of varying levels of imprecision. This is used to explore the effect of fusion imprecision on learning performance in a series of agent-based simulations. In general, the results suggest that imprecise fusion operators are optimal when the frequency of fusion is high relative to the frequency with which evidence is obtained from the environment.

## Introduction

Decentralised collective learning is a key task in multi-agent systems and swarm robotics, in which agents learn from each other as well as from evidence directly obtained from the environment. These two processes can naturally be modelled using belief fusion operators for the former and evidential updating operators for the latter. In this context we investigate the effect of imprecision in the fusion process on collective learning in scenarios where direct evidence can be both sparse and noisy. More specifically, by proposing a pairwise fusion operator with a single parameter controlling the level of precision of fused beliefs, we will use agent-based simulations to show that more imprecise fusion is optimal for scenarios in which, from an agent perspective, fusion is more frequent than belief updating from evidence.

We will focus on a propositional model in which agents attempt to collectively identify the true state of the world in the form of the correct allocation of truth values to a set of predefined propositions. In particular, we assume that agents' beliefs take the form of sets of possible states of the world where each state is a complete allocation of truth values to the propositions under consideration. For example, consider a search and rescue task in which five possible locations for survivors have been identified and the presence or absence of survivors at each site is asserted by a binary proposition so that there are $2^5 = 32$ states. Each state then corresponds to an allocation of truth values to the $5$ propositions, only one of which is the ground truth.

Representing beliefs by sets of states provides a natural way of modelling imprecision in terms of cardinality. The higher the cardinality of a belief, i.e., the more states it deems to be possible, the more imprecise it is. Hence, in this context a fusion operator is imprecise if it tends to result in beliefs of high cardinality. The importance of imprecision has long been considered in philosophy, with many arguing that imprecision in the form of vague expressions has great utility in natural language (Parikh, 1994; van Deemter, 2009). Despite autonomous agents being – by their very nature – inherently precise in how they store and retrieve information, for decades researchers in artificial intelligence have sought to provide agents with belief representations that capture the notions of uncertainty or imprecision (Zadeh, 1988; Shafer, 1976).

The propositional model of collective learning can be applied to a broad class of applications. For example, we could model the COVID-19 risk for areas around the world with propositions of the form "this area is high/low risk". The true state of the world would then correspond to the actual risk profile across all areas. Agents might then be different health authorities, which would share information between them and obtain direct evidence based on COVID testing in areas for which they are responsible.

An outline of the rest of the paper is as follows: In the next section, we give an overview of some relevant existing literature on collective learning and decision making. We then describe a propositional model and introduce an imprecise operator for belief fusion and evidential updating, before detailing simulation experiments in which this model is applied to collective learning. The results from these simulation experiments are then considered in detail with a particular focus on how they pertain to the use of imprecision in belief fusion for collective learning. Finally, we give some conclusions and outline possible future directions.

## Related Work

Decentralised collective learning and decision-making are well studied tasks in swarm robotic systems (Brambilla et al., 2013). One common family of problems considered are best-of-$n$ problems, in which a population of robots must collectively identify the best out of $n$ distinct options on the basis of local interactions and limited feedback (Parker and Zhang, 2009). See (Valentini et al., 2017) for an extensive overview of this class of problems. In this paper we will instead focus on a more challenging collective learning problem in which a population of agents must learn a complete description of the true state of the world rather than simply identifying the best out of a set of $n$ alternatives. We consider collective learning in terms of two distinct processes; evidential updating and belief fusion. Evidential updating is the process by which an agent updates its current beliefs based on evidence received from the environment. In robotics applications evidence might take the form of signals received by various sensor modalities. In a more general multi-agent context evidence could take the form of data received directly relating to a particular instance or set of instances. The interaction between evidential updating and belief fusion has been studied in the context of social epistemology by Douven and Kelp (2011). The authors argue that agent-to-agent communication would significantly enhance the performance in truth approximation due to the ability to correct errors and spread information across the population. Recent work along these lines has shown the efficacy of combining both processes on the best-of-$n$ problem (Crosscombe et al., 2017; Lawry et al., 2019; Lee et al., 2018).

Information fusion has been widely studied in the areas of image fusion (Bloch, 2015), robotic sensor fusion (Murphy, 1998) and opinion diffusion logics (Konieczny and Pérez, 2002; Cholvy, 2018). In general, the fusion process provides a means of resolving inconsistencies between different sources and hence achieving consensus. The fusion of agents' beliefs can be modelled by pairwise operators. Several such fusion operators have been proposed and Dubois et al. (2016) introduces a number of desirable properties that any information fusion process should satisfy.

The use of sets of states to represent beliefs, referred to as epistemic sets, dates back to Hintikka's possible worlds semantics (Hintikka, 1962) and some early applications in AI and computer science can be found in (Vardi, 1989) and (Ruspini, 1987). More recently, epistemic sets have been applied to the best-of-$n$ collective learning problem (Lawry et al., 2019). Using epistemic sets allows for agents to hold beliefs of varying levels of precision which has the potential to improve system-level robustness to noise compared with other simpler models, e.g., the weighted voter model (Valentini et al., 2014; Crosscombe et al., 2017).

For epistemic sets we can define fusion operators that work on the principle that disagreement or inconsistency between agents results in more imprecise beliefs, while agreement between agents increases precision. In this way the fusion process can help to both propagate correct information while also correcting errors (Lawry et al., 2019). There is a strong relationship between epistemic sets and Dempster-Shafer theory since in the latter beliefs can be thought of as being characterised by mass functions defined over epistemic sets. In Dempster-Shafer theory a large number of fusion operators have been proposed in the literature as overviewed by Osswald and Martin (2006). The robustness to noise of several of these operators applied to the best-of-$n$ problem has been compared in (Crosscombe et al., 2019). In particular, best performance in noisy conditions was achieved by Yager's operator (Yager, 1992) and by Dubois & Prade's operator (Dubois and Prade, 1988).

## Belief Representation and Fusion Operators

Consider a population of agents attempting to collectively learn the state of their environment which we assume can be described by a finite set of propositions $\mathcal{P} = \{p_1, ..., p_n\}$. From this perspective a state $s$ is the allocation of Boolean truth values to each of the propositions. In other words, a state is a function $s : \mathcal{P} \rightarrow \{0, 1\}^n$. For notational convenience we represent a state $s$ by the $n$-tuple $\langle s(p_1), \ldots, s(p_n)\rangle$. Let $\mathbb{S}$ denote the set of all states so that $|\mathbb{S}| = 2^n$. An agent's belief $B \subseteq \mathbb{S}$ is then the set of states which the agent believes can possibly be the true state. We therefore represent uncertain beliefs as being subsets of $\mathbb{S}$ with cardinality $|B| > 1$ while a singleton belief $B = \{s\}$ means that the agent is certain that $s$ is the true state. We assume that agents adopt a closed-world assumption which in this context means assuming that $\mathbb{S}$ covers all possible states of the world. Consequently agents beliefs' are such that $B \neq \emptyset$ since it cannot be the case that all states in $\mathbb{S}$ are impossible. Note that a given belief $\emptyset \neq B \subseteq \mathbb{S}$ classifies each proposition $p_i$ as being either true, if $s(p_i) = 1$ for all $s \in B$, false, if $s(p_i) = 0$ for all $s$ in $B$, or uncertain otherwise. Hence, the more imprecise an agent's belief the more propositions they will tend to be uncertain about. This indicates a natural relationship between the epistemic model of beliefs and three-valued approaches (Crosscombe et al., 2017). For example, consider the search and rescue scenario with 5 location outlined above where $p_i$ denotes the proposition 'casualties are in location $i$ for $i = 1, \ldots, 5$. Now consider the belief $B$ given by,

$$B = \{\langle 1, 0, 0, 0, 0\rangle, \langle 0, 1, 0, 0, 0\rangle\}$$

In case this $B$ corresponds to the belief that there are casualties either in location 1 or location 2 but not both, and there are no casualties in any other location. Therefore, according to $B$ no propositions are classified as being certainly true, $p_3$, $p_4$ and $p_5$ are classified as being certainly false, and $p_1$ and $p_2$ are uncertain.

We now introduce a parameterized fusion operator for combining epistemic sets which returns beliefs of varying levels of imprecision. This requires a measure of the similarity between epistemic sets for which we use the well-known

Jaccard similarity (Jaccard, 1912) defined as follows: For $B_1, B_2 \subseteq \mathbb{S}$,

$$J(B_1, B_2) = \frac{|B_1 \cap B_2|}{|B_1 \cup B_2|} \quad (1)$$

We now define the similarity threshold operator as follows: for $\gamma \in [0, 1]$,

$$B_1 \odot_\gamma B_2 = \begin{cases} B_1 \cap B_2 & : \ J(B_1, B_2) > \gamma \\ B_1 \cup B_2 & : \ J(B_1, B_2) \leq \gamma \end{cases} \quad (2)$$

For example, let $B_1 = \{s_1, s_2, s_3\}$ and $B_2 = \{s_2, s_3, s_4, s_5\}$ then,

$$J(B_1, B_2) = \frac{|\{s_2, s_3\}|}{|\{s_1, s_2, s_3, s_4, s_5\}|} = \frac{2}{5}$$

and hence $B_1 \odot_\gamma B_2 = \{s_2, s_3\}$ if $\gamma < \frac{2}{5}$ and $B_1 \odot_\gamma B_2 = \{s_1, s_2, s_3, s_4, s_5\}$ for $\gamma \geq \frac{2}{5}$.

Note that for $\gamma = 0$ this operator corresponds to the intersection-union operator (Dubois et al., 2016, eq. 23) as given by:

$$B_1 \odot_0 B_2 = \begin{cases} B_1 \cap B_2 & : \ B_1 \cap B_2 \neq \emptyset \\ B_1 \cup B_2 & : \ B_1 \cap B_2 = \emptyset \end{cases} \quad (3)$$

On the other hand, for $\gamma = 1$ we have that $B_1 \odot_1 B_2 = B_1 \cup B_2$. In general, $\gamma$ controls the level of generality or precision of the operator such that for $\gamma \leq \gamma'$, $B_1 \odot_\gamma B_2 \subseteq B_1 \odot_{\gamma'} B_2$ for all sets $B_1, B_2 \subseteq \mathbb{S}$. The use of Jaccard similarity to deal with inconsistency has also been proposed by (Schockaert and Prade, 2010) who applies similarity-based enlargement of the sets of interpretations to resolve inconsistencies in fusion problems.

In addition to combining their beliefs with others, agents also gather evidence from the environment. Here we assume that this evidence $E$ takes the form of an assertion of the truth value of one proposition $p_1, \ldots, p_n$, i.e., $E = \{s \in \mathbb{S} : s(p_i) = v\}$ for some $p_i$ and where $v \in \{0, 1\}$. Given $E$ we then propose that an agent updates their belief $B$ to $B|E$ such that:

$$B|E = \{s|E : s \in B\}. \quad (4)$$

and where,

$$s|E(p_j) = \begin{cases} v : j = i \\ s(p_j) : \text{otherwise} \end{cases} \quad (5)$$

In contrast to more established belief updating methods such as where

$$B|E = \begin{cases} B \cap E : B \cap E \neq \emptyset \\ B : \text{otherwise} \end{cases} \quad (6)$$

as applied in (Lawry et al., 2019), the above approach has the advantage that it preserves beliefs about the propositions which are consistent with $E$. To see this consider the case where we have two propositional variables and where $B = \{\langle 1, 0\rangle, \langle 0, 1\rangle\}$, i.e., in this case both $p_1$ and $p_2$ are uncertain since for both there are states in $B$ where they are true and also states where they are false. Suppose evidence $E$ is "$p_1$ is true", then applying intersection-based updating results in $B|E = \{\langle 1, 0\rangle\}$ and therefore removes the uncertainty about $p_2$. More specifically, according to $B$, $p_2$ is now certainly false. However, since $E$ makes no reference to $p_2$ this seems counter-intuitive. In contrast applying our proposed updating method results in $B|E = \{\langle 1, 0\rangle, \langle 1, 1\rangle\}$, hence preserving the uncertainty about $p_2$. In general, provided that the proposition $p_i$ to which the evidence pertains is classified uncertain by belief $B$ then updating on the basis of equation 4 will result in more imprecise beliefs than updating on the basis of equation 6. Less restrictive updating of this kind can be advantageous in collective learning scenarios in which evidence is noisy and where agents focus on investigating propositions about which they are currently uncertain.

## Simulation Experiments

In this section, we present the results for a series of agent-based simulation experiments in which a population of agents attempts to reach a consensus about a propositional state description of the world. At each time step, each agent randomly chooses a proposition about which it is uncertain to investigate. During each iteration, every agent has probability $\rho$ of successfully obtaining evidence from the environment. We refer to $\rho$ as the evidence rate. In addition, there is some probability $\epsilon$ that the evidence received could be incorrect. More specifically, let the true state be denoted by $s^*$ then if an agent receives evidence about proposition $p_i$, they will receive $E = \{s : s(p_i) = s^*(p_i)\}$ with probability $1 - \epsilon$ and $E = \{s : s(p_i) = 1 - s^*(p_i)\}$ with probability $\epsilon$. Hence, $\epsilon$ provides a simple model of environmental or sensor noise. Should an agent become certain in its belief about the truth values of each proposition $p_1, \ldots, p_n$, then the agent stops seeking to obtain additional evidence directly from its environment. Agents also exploit the evidence being gathered by the other agents by means of belief fusion. At every time step a pool of agents is generated at random and agents in the pool fuse their beliefs with others. Each agent in the population has probability $\sigma$ (the fusion rate) of entering the pool in any given time step, and once in the pool, agents are randomly paired to combine their beliefs using the fusion operator in Equation (2). Both agents then adopted the result of this fusion as their new belief. If the number of agents in pool is odd, there will be one agent left without combining its belief.

The population of agents are initialised to hold completely ignorant beliefs, i.e., $B = \mathbb{S}$ at time $t = 0$. In other words, the agents are initialised as having no prior knowledge about the world. Experiments are then run multiple times to reduce random noise. Population-level performance in this collective learning task is measured by the average Hamming distance $H$ from the agents' belief values to the true state $s^*$. Furthermore, without loss of generality we assume that $s^*$ is such that $s^*(p_i) = 1$ for $i = 1, \dots, n$. In this context, the Hamming distance between states is defined as follows: Let $s_1 = \langle s_1(p_1), \dots, s_1(p_n) \rangle$ and $s_2 = \langle s_2(p_1), \dots, s_2(p_n) \rangle$ be two states, then the Hamming distance between them is given by:

$$H(s_1, s_2) = \sum_{i=1}^{n} |s_1(p_i) - s_2(p_i)| \tag{7}$$

We then extend this to give the normalised Hamming distance between an epistemic set $B \subseteq \mathbb{S}$ and the true state of the world $s^*$ as follows:

$$H(B, s^*) = \frac{1}{|B|} \frac{1}{n} \sum_{s \in B} H(s, s^*) \tag{8}$$

Furthermore, we evaluate the performance at the population level, the average Hamming distance between the population of agents $\mathcal{A}$ of size $k$ and $s^*$ is given as follows:

$$H(\mathcal{A}, s^*) = \frac{1}{k} \sum_{B \in \mathcal{A}} H(B, s^*) \tag{9}$$

In addition to Hamming distance we also evaluate performance by measuring the proportion of correct propositions in each belief and by using an F-measure. More specifically, for an epistemic set $B$ we define the following:

$$C(B) = \{p_i : s(p_i) = s^*(p_i) \text{ for all } s \in B\} \tag{10}$$
$$I(B) = \{p_i : s(p_i) = 1 - s^*(p_i) \text{ for all } s \in B\} \tag{11}$$

If agent $a_i \in \mathcal{A}$ has belief $B_i$ then let,

$$\alpha = \frac{1}{n} \frac{1}{k} \sum_{a_i \in \mathcal{A}} |C(B)| \text{ and } \theta = 1 - \frac{1}{n} \frac{1}{k} \sum_{a_i \in \mathcal{A}} |I(B)| \tag{12}$$

$C(B)$ is the number of propositions about which belief $B$ is both certain and correct, while $I(B)$ is the number of propositions about which $B$ is both certain and incorrect. Hence, $\alpha$ is the average proportion of propositions about which agents are both certain and correct, while $\theta$ is the average proportion of propositions about which agents are either correct or uncertain, i.e., not incorrect. From this we can then define the $F_\beta$ score according to:

$$F_\beta = (1 + \beta^2) \frac{\alpha \cdot \theta}{\beta^2 \cdot \alpha + \theta}. \tag{13}$$

Here $\beta$ is a parameter that allows us to give different degrees of importance to $\alpha$ and $\theta$ such that $F_\beta$ attributes $\beta$ times as much importance to not being incorrect as to being correct.

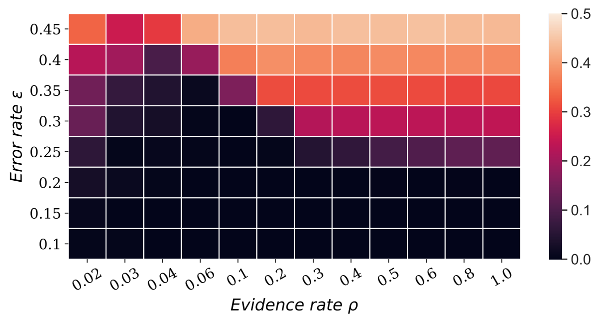


Figure 1: The average distance $H$ to the true state for evidence rate $\rho \in [0.02, 1]$ and error rate $\epsilon \in [0.1, 0.45]$ where the fusion rate $\sigma = 0.04$.

## Results

### The Intersection-Union Operator

In this section we present the results of our model for collective learning using the intersection-union operator in Equation (3) to which we shall later compare the proposed similarity threshold operator in Equation (2). The population size is $k = 50$ for all the experiments unless specified otherwise. Figure 1 shows a heat map of the average Hamming distance $H$ of the population's beliefs (from Equation (9)) to the true state $s^*$ after 3000 time steps, for varying $\rho$ and $\epsilon$. Here, a darker colour indicates that the average belief of the population is at a lower distance from the true state and therefore the population is performing better under these conditions, while a lighter colour indicates a greater distance to the true state and poorer performance. For an evidence rate $0.02 < \rho < 0.04$ we see that increasing the evidence rate also leads to an increase in robustness to noise in that good performance in maintained for higher values of $\epsilon$. However, when $\rho > 0.04$, the robustness to noise surprisingly decreases as the evidence rate increases further. In these experiments, the fusion rate $\sigma = 0.04$ is fixed and relatively low corresponding on average to only one fusion operation per time step. This suggests that in noisy scenarios the relative frequency of fusion and evidential updating has an impact on overall performance. We hypothesise that fusion plays an error-correcting role when $\epsilon$ is high (Lawry et al., 2019; Crosscombe et al., 2017), but if evidence is obtained at too high a rate then fusion is not sufficiently frequent to correct for the noise. On the other hand, if the evidence rate is too low then fusion drives the system to consensus too quickly before there has been time for agents to receive a sufficient amount of evidence to identify the true state. To investigate this directly, we now vary the fusion rate $\sigma$ and present the results in Figure 2.

Figure 2(a) shows population performance at $t = 3000$ for varying values of $\sigma$ and $\rho$ with the fixed $\epsilon = 0.3$ and $k = 50$. The figure suggests that lower evidence rates re-

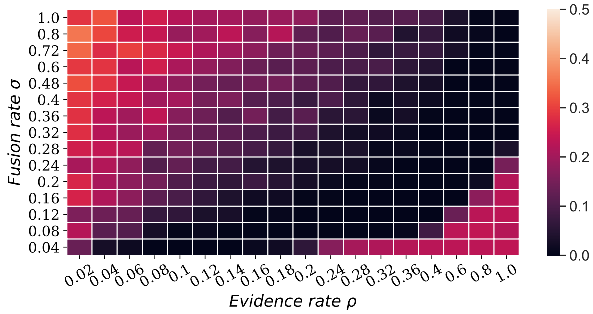


(a) Evidence rate $\rho \in [0.02, 1]$ where $\epsilon = 0.3$.

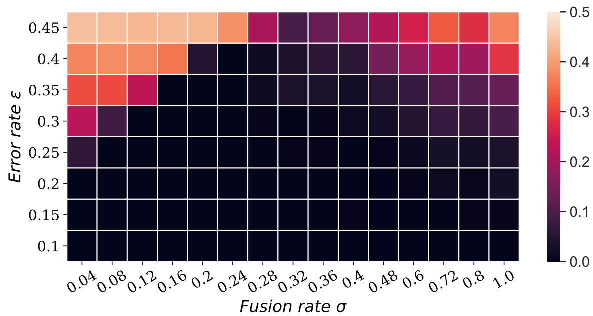


(b) Error rate $\epsilon \in [0.1, 0.45]$ where $\rho = 0.4$.

Figure 2: The average distance $H$ to the true state for fusion rate $\rho \in [0.04, 1]$.

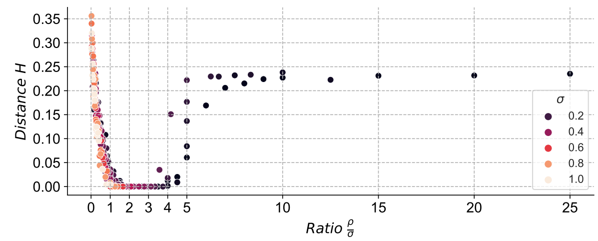


Figure 3: The average distance $H$ to the true state against the ratio of evidential updating to belief fusion $\frac{\rho}{\sigma}$, for a fixed error rate $\epsilon = 0.3$.

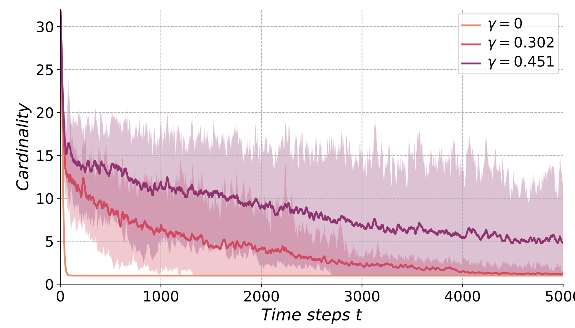


Figure 4: Average belief cardinality of the population against time steps for $\rho = 0.01, \sigma = 0.2, \epsilon = 0.3$, and $n = 5$

quire a lower fusion rate in order to achieve good performance and vice versa. We suggest that again this effect may be due to the relative frequency of updating and fusion from the perspective of an individual agent. More specifically, Figure 2(a) suggests that for a given evidence rate $\rho$ there is an interval of fusion rates $\sigma$ for which optimal performance can be obtained. If the fusion rate is too high relative to the evidence rate then the population converges too quickly before there is time to collect sufficient evidence. On the other hand, if the fusion rate is too low relative to the evidence rate then there is insufficient fusion to correct for the errors introduced by evidential noise.

Figure 2(b) shows that higher error rates lead to smaller optimal intervals for $\sigma$. The performance is poorer on the top left and right regions of the figure where the fusion rate is either too low or too high respectively. Figure 3 shows $H$ plotted against $\frac{\rho}{\sigma}$ when $\epsilon = 0.3$. This suggests that performance is best when $\frac{\rho}{\sigma} \in [2, 3]$, i.e., when the evidence rate is between 2 and 3 times the fusion rate. The average distance to the true state increases significantly more slowly when $\frac{\rho}{\sigma} > 10$ and $H(\mathcal{A}, s^*)$ tends to 0.25. We suggest that evidence updating dominates in such cases so that fusion has little impact on the population and the distance to the true state converges to $\epsilon = 0.3$ as would be expected if learning is based only on evidential updating without the error correction provided by fusion. Therefore, the performance of collective learning in this case can potentially be optimised by controlling the relative frequency at which evidential updating and fusion take place. Performance is relatively poor as $\frac{\rho}{\sigma}$ tends to 0. In such cases it is likely that beliefs across the population are converging before sufficient information becomes available to the population in the form of direct evidence. The discontinuity of the operator can also be problematic in such situations since it can lead to loss of information. For example, $B_1 \odot B_2 = \{s_i\}$ if $B_1 = \{s_1, ..., s_i\}$ and $B_2 = \{s_i, ..., s_n\}$, while if $B_1 = \{s_1, ..., s_i\}$ and $B_2 = \{s_{i+1}, ..., s_n\}$ then $B_1 \odot B_2 = \mathbb{S}$. Therefore, we hypothesise that an imprecise operator will improve the performance and we propose using the threshold operator as given in Equation (2) to vary the imprecision and to reduce the discontinuity effect of the operator.

## Threshold Operator Results

In this section we investigate the effect on performance of systematically varying the imprecision of the fusion operator under different learning scenarios. More specifically, we consider the threshold fusion operator in Equation (2) for a range of different values of $\gamma$ between 0 and 1, and for different values of $\epsilon$ and $\gamma$. We show that in high noise and low evidence scenarios performance is optimal for an intermediate level of imprecision, especially when performance is measured in terms of an $F_\beta$ score which gives significantly higher importance to no being incorrect than to being correct. For this we will use the $F_3$ score. In the following results we will consider the fusion operators $\odot_\gamma$ for $\gamma \in \{\frac{j}{i} | i \in [1, 2^n], j \in [0, i], i, j \in \mathbb{Z}\}$.

Figure 4 shows the average cardinality of beliefs in the population decreasing over time for different thresholds. Recall that for epistemic sets, i.e. belief represented by sets of states, cardinality provides a measure of imprecision. Hence, we see a general trend by which the precision of the agents' beliefs increases during collective learning. The red line ($\gamma = 0.302$) declined slower than the orange line ($\gamma = 0$) which is equivalent to the intersection-union operator and they both reach a cardinality of 1 after 5000 iterations. However, the purple line ($\gamma = 0.451$) decreases even more slowly to an average cardinality of around 5 after 5000 iterations. Hence, decreasing the precision of the fusion operator has the effect of slowing convergence and also tends to result in agents holding more imprecise beliefs over time.

Let $\gamma_{F_3}$ denote the value of $\gamma$ at which $F_3$ is maximal at $t = 5000$, and let $\gamma_\alpha$ denote the value of $\gamma$ at which $\alpha$ is maximal at $t = 5000$. Figure 5 shows two heat maps of performance, i.e. values of $\alpha$ and $F_3$ respectively, with $\epsilon = 0.3$ plotted against fusion rate $\sigma$ and evidence rate $\rho$. For low evidence rate and relatively high fusion rate we have that $\gamma_\alpha > 0$ and $\gamma_{F_3} > 0$. This is consistent with the hypothesis that more imprecise operators are optimal when $\rho$ is significantly less than $\sigma$, i.e. when $\frac{\rho}{\sigma}$ is low.

Figure 6 shows $\alpha$ and $F_3$ score at $t = 5000$ plotted against different values of $\gamma$ for evidence rate $\rho = 0.01$ and fusion rate $\sigma = 0.2$. Both measures are maximal for an intermediate value of $\gamma$, and this is particularly pronounced for the $F_3$ score. Figure 6 also shows the average proportion of determined propositions as corresponding to those propositions about which an agent is certain (orange line). This is given by $\alpha + 1 - \theta$. Notice that this decreases as $\gamma$ increases thus resulting in more imprecise fusion and $\alpha = F_3$ when agents are certain on all the propositions, i.e., $\alpha = \theta$. However, at the values of $\gamma$ for which $\alpha$ or $F_3$ are maximal the number of determined propositions remains high, suggesting that the population is still learning the true state with a high degree of precision.

Figure 7 shows $\gamma_{F_3}$ (purple lines), and $\gamma_\alpha$ (red lines) for varying $\rho$, $\epsilon$ and $k$. Figure 7(a) plots $\gamma_{F_3}$ and $\gamma_\alpha$ against fusion rate $\sigma$ for $\rho = 0.01$ and $\epsilon = 0.3$, and shows that in this case both increase with $\sigma$ for $\sigma < 0.4$. $\gamma_{F_3}$ stops changing and $\gamma_\alpha$ decreases slightly with $\sigma$ for $\sigma > 0.4$.

Figure 7(a) suggests that as the fusion rate increases relative to the evidence rate then optimal performance requires an increasingly imprecise fusion operator. In other words, if fusion is frequent relative to evidence acquisition then it is better if the result of the fusion is imprecise, whereas if fusion is rare relative to evidence acquisition then it is better if the result of the fusion is precise. Figure 7(b) illustrates this relationship between $\sigma$ and $\rho$ from a different perspective by plotting $\gamma_{F_3}$ and $\gamma_\alpha$ against evidence rate $\rho$. Again we see that as the evidence rate increases relative to the fusion rate, optional performance is obtained using increasingly precise fusion operators.

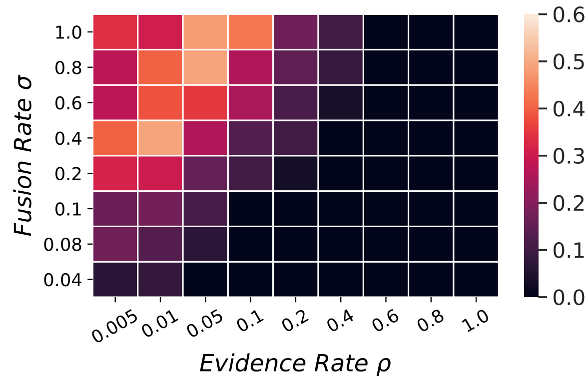


(a) Optimal thresholds $\gamma_\alpha$ for $\alpha$.

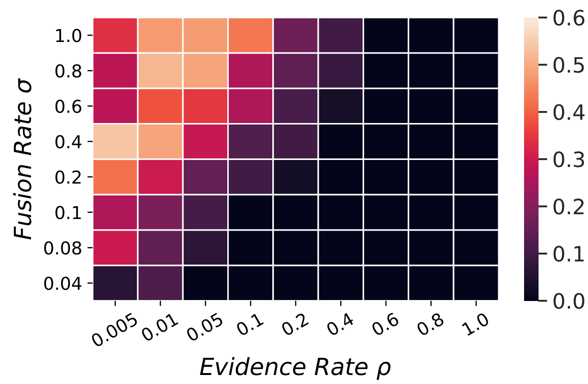


(b) Optimal thresholds $\gamma_{F_3}$ for $F_3$.

Figure 5: For $\rho \in [0.005, 1]$ and fusion rate $\sigma \in [0.04, 1.0]$ where $\epsilon = 0.3$.

Figure 7(c) shows optimal threshold values for fixed $k$ and $\rho$ plotted against the noise probability $\epsilon$. This suggests a subtle relationship between imprecision and noise. For example, as the error rate increases the optimal imprecision value $\gamma_{F_3}$ increases. In other words, as noise increases then more imprecise fusion operators are required in order to optimise the $F_3$ score. On the other hand, optimising the proportion of correctly classified propositions $\alpha$ requires more precise operators in high noise scenarios, although these are still more imprecise than the intersection-union operator. These results are likely to be due to a trade-off between error and uncertainty. For very imprecise fusion operators agents will

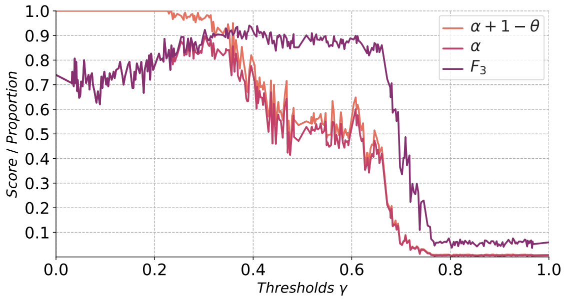


Figure 6: Average score ($F_3$) or proportion ($\alpha$, $\alpha + 1 - \theta$) for a population of $k = 50$ agents, language size $n = 5$, fusion rate $\sigma = 0.2$, evidence rate $\rho = 0.01$, and error rate $\epsilon = 0.3$.

tend to hold beliefs with higher cardinality (see Figure 4) and therefore have a higher proportion of uncertain propositions. Ideally, these uncertain propositions should be those for which a given agent has received conflicting information and would therefore have had a higher chance of incorrectly classifying them if they held more precise beliefs. The $F_3$ score takes into account this potential trade-off between accuracy and uncertainty, while on the other hand, $\alpha$ will tend to decrease as $\gamma$ increases since the proportion of correctly classified propositions is bounded by the proportion of propositions which are determined, i.e., which are classified as being true or false rather than uncertain (see Figure 6).

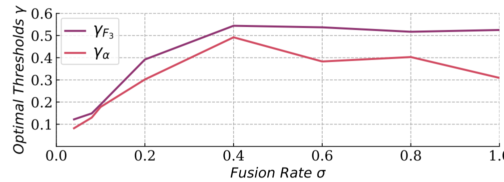


(a) Against fusion rate $\sigma$ for evidence rate $\rho = 0.01$ and error rate $\epsilon = 0.3$

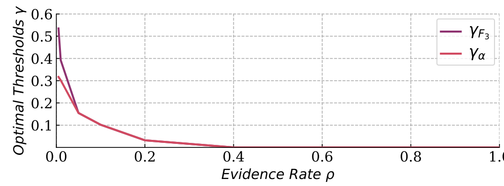


(b) Against evidence rate $\rho$ for fusion rate $\sigma = 0.2$ and error rate $\epsilon = 0.3$

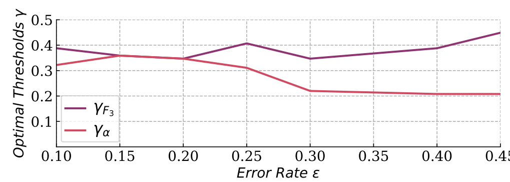


(c) Against error rate $\epsilon$ for fusion rate $\sigma = 0.2$ and evidence rate $\rho = 0.01$

Figure 7: Optimal thresholds $\gamma$ according to $F_3$ ($\gamma_{F_3}$) and $\alpha$ ($\gamma_\alpha$).

## Discussion and Conclusion

In this paper we have investigated collective learning of the state of the world in a propositional model. Agents' beliefs are represented by epistemic sets corresponding to the set of states that each agent believes are possibly the true state. Agents learn from two distinct sources: directly from the environment using belief updating, and from other agents by applying a fusion operator to combine their beliefs with those of other agents.

In this context we show that the well-known intersection-union fusion operator results in effective collective learning across a range of scenarios in which there is noise and relatively limited direct evidence. However, performance is affected by the relative relationship between the evidence rate $\rho$ and the fusion rate $\sigma$. More specifically, our results suggest that for a given value of $\rho$ there is a bounded interval of values of $\sigma$ for which performance is good.

To explore the effect of imprecision in fusion, we introduced a new parameterised fusion operator that can return beliefs of different levels of precision or imprecision. Our results then show that if $\sigma$ is high relative to $\rho$ then optimal results are obtained by a more imprecise operator. In general, the optimal level of imprecision is strongly dependent on the frequency of evidence relative to the frequency of fusion, so that the higher the frequency of fusion is compared to the frequency of evidence the more imprecise the operator should be, and much less clearly dependent of the level of evidential noise.

Following this work, we intend to investigate other imprecise fusion operators based on alternative measures of similarity or distance, e.g., a threshold operator based on the Hamming distance. Additionally, having shown that preserving imprecision during belief fusion leads to greater performance in the context of collective learning, we seek to explore whether imprecise forms of evidential updating may lead to similar improvements in the learning process. Finally, simulated robot experiments will enable us to investigate a more grounded collective learning scenario in which to better understand the implications of our model when external constraints such as communication bandwidth and noise from on-board sensors are imposed by the robot hardware.

## Acknowledgement

Zixuan Liu is supported by a scholarship granted by the China Scholarship Council (CSC, 201908150152).

## References

Bloch, I. (2015). Fuzzy sets for image processing and understanding. *Fuzzy sets and systems*, 281:280–291.

Brambilla, M., Ferrante, E., Birattari, M., and Dorigo, M. (2013). Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7:1–41.

Cholvy, L. (2018). Opinion diffusion and influence: A logical approach. *International Journal of Approximate Reasoning*, 93:24–39.

Crosscombe, M., Lawry, J., and Bartashevich, P. (2019). Evidence propagation and consensus formation in noisy environments. In *International Conference on Scalable Uncertainty Management*, pages 310–323. Springer.

Crosscombe, M., Lawry, J., Hauert, S., and Homer, M. (2017). Robust distributed decision-making in robot swarms: Exploiting a third truth state. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4326–4332. IEEE.

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

Dubois, D., Liu, W., Ma, J., and Prade, H. (2016). The basic principles of uncertain information fusion. an organised review of merging rules in different representation frameworks. *Information Fusion*, 32:12–39.

Dubois, D. and Prade, H. (1988). Representation and combination of uncertainty with belief functions and possibility measures. *Computational intelligence*, 4(3):244–264.

Hintikka, J. (1962). Knowledge and belief: an introduction to the logic of the two notions.

Jaccard, P. (1912). The distribution of the flora in the alpine zone. 1. *New phytologist*, 11(2):37–50.

Konieczny, S. and Pérez, R. P. (2002). Merging information under constraints: a logical framework. *Journal of Logic and computation*, 12(5):773–808.

Lawry, J., Crosscombe, M., and Harvey, D. (2019). Epistemic sets applied to best-of-n problems. In *European Conference on Symbolic and Quantitative Approaches with Uncertainty*, pages 301–312. Springer.

Lee, C., Lawry, J., and Winfield, A. (2018). Negative updating combined with opinion pooling in the best-of-n problem in swarm robotics. In *International Conference on Swarm Intelligence*, pages 97–108. Springer.

Murphy, R. R. (1998). Dempster-shafer theory for sensor fusion in autonomous mobile robots. *IEEE Transactions on Robotics and Automation*, 14(2):197–206.

Osswald, C. and Martin, A. (2006). Understanding the large family of dempster-shafer theory's fusion operators-a decision-based measure. In *2006 9th International Conference on Information Fusion*, pages 1–7. IEEE.

Parikh, R. (1994). Vagueness and utility: The semantics of common nouns. *Linguistics and Philosophy*, 17(6):521–535.

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

Ruspini, E. H. (1987). Epistemic logics, probability, and the calculus of evidence. In *Proceedings of the 10th international joint conference on Artificial intelligence-Volume 2*, pages 924–931.

Schockaert, S. and Prade, H. (2010). An inconsistency-tolerant approach to information merging based on proposition relaxation. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*.

Shafer, G. (1976). *A mathematical theory of evidence*, volume 42. Princeton university press.

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

Valentini, G., Hamann, H., Dorigo, M., et al. (2014). Self-organized collective decision making: the weighted voter model. In *AAMAS*, pages 45–52.

van Deemter, K. (2009). Utility and language generation: The case of vagueness. *Journal of Philosophical Logic*, 38(6):607–632.

Vardi, M. (1989). On the complexity of epistemic reasoning. In *Proceedings. Fourth Annual Symposium on Logic in Computer Science*, pages 243–244.

Yager, R. R. (1992). On the specificity of a possibility distribution. *Fuzzy Sets and Systems*, 50(3):279–292.

Zadeh, L. A. (1988). Fuzzy logic. *Computer*, 21(4):83–93.