Swarm Intelligence
https://doi.org/10.1007/s11721-024-00238-7

Imprecise evidence in social learning

Zixuan Liu[1] · Michael Crosscombe[2] · Jonathan Lawry[1]

Received: 23 February 2023 / Accepted: 27 March 2024

Abstract
Social learning is a collective approach to decentralised decision-making and is comprised of two processes; evidence updating and belief fusion. In this paper we propose a social learning model in which agents' beliefs are represented by a set of possible states, and where the evidence collected can vary in its level of imprecision. We investigate this model using multi-agent and multi-robot simulations and demonstrate that it is robust to imprecise evidence. Our results also show that certain kinds of imprecise evidence can enhance the efficacy of the learning process in the presence of sensor errors.

Keywords Social learning · Distributed decision-making · Multi-agent systems · Imprecise evidence

1 Introduction

In distributed autonomous systems, individuals learn about the state of the world by both individually gathering evidence directly from the environment, and by collectively fusing their own beliefs with those of their peers; the latter is usually referred as social learning and can be observed in many social animals (Heyes, 1994). Social learning is inherently decentralised, involving local interactions between agents to share information and to reach consensus. For example, in one simple form of social learning, a population of agents attempt to learn the true state of the world as represented by the truth values of several propositions. A proposition in this context could assert whether or not a certain property holds. For example, in a decentralised wildfire detection mission it

Michael Crosscombe and Jonathan Lawry have contributed equally to this work.

✉ Zixuan Liu
z.liu@bristol.ac.uk

Michael Crosscombe
cross@sacral.c.u-tokyo.ac.jp

Jonathan Lawry
j.lawry@bristol.ac.uk

1 School of Engineering Mathematics and Technology, University of Bristol, Queens Road, Bristol BS8 1QU, UK

2 Graduate School of Arts and Sciences, The University of Tokyo, Komaba, Tokyo 153-8902, Japan

Published online: 16 April 2024

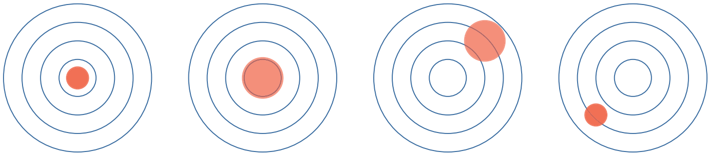

Fig. 1 Accuracy and precision. From left to right: accurate and precise; accurate and imprecise; inaccurate and imprecise; inaccurate and precise

may be necessary to determine whether or not fires are present in each of the different grid locations overlaying a map of the search area.

In the social learning literature, noise is usually modelled as Gaussian variation from a true value (Lee et al., 2021; Talamali et al., 2019) or as a probability of receiving false evidence (Lawry & Lee, 2020). In this paper we differentiate noise of this form, which we hereafter refer to as inaccuracy, from imprecision and we show that the latter also has an important and potentially useful effect on social learning. Specifically, noise as outlined above, relates to the accuracy of the evidence obtained with reference to the true state of the world. In other words, evidence is accurate if it is consistent with the true state of the world and inaccurate otherwise. In contrast, evidence is imprecise if it fails to identify a single state of the world. From this perspective it is possible for evidence to be both accurate and imprecise since this simply means that the evidence is consistent with the true state of the world but does not uniquely identify it. In other words, the evidence is consistent with a number of possible states including the true state. In this case the degree of precision of the evidence is dependent on the number of states that are consistent with it.

From this perspective we consider noise to consist of both inaccuracy and imprecision, where the former describes the difference between the evidence gathered and the ground truth, and the latter is where evidence fails to identify a single state of the world. Imprecise evidence is inherently less informative than precise evidence, but this means that in the presence of errors it is more likely to be consistent with the true state of the world. Figure 1 illustrates the distinction between accuracy and precision for noisy evidence in relation to the ground truth. The centre of the rings represents the true state of the world and an orange circle depicts the evidence gathered by an agent. The accuracy of the evidence then depends on the distance of the centre of the circle to the true state, while the imprecision of the evidence depends on the area of the circle. Hence, imprecision and inaccuracy are independent features of the evidence. Increasing/decreasing the radius of the circle results in an increase/decrease in the precision of the evidence, respectively, but the accuracy of the evidence remains unchanged. On the other hand, the distance from the centre of the orange circle to the inner-most ring is inversely proportional to the accuracy of the evidence (i.e., the greater the distance, the lower the accuracy). In this paper, *imprecision* is the size of the evidence set or the number of possible states consistent with the evidence, and can be a parameter of our model. Accuracy, on the other hand, is simply a measure of the distance from the evidence set to the true state of the world and is assumed to be a feature of the interaction between the agents' sensors and the environment.

It has been suggested in the literature that some social insects may be able to adapt and even take advantage of imprecision to make more robust collective decisions (De Marco et al., 2008; Meyer et al., 2017; Dussutour et al., 2009). In this context it is therefore interesting to investigate if artificial multi-agent systems can adapt to imprecise evidence and whether some forms of imprecise evidence may even enhance learning in the presence of sensor error or other types of environmental noise.

The contributions of this paper are twofold; 1) we show that social learning can be robust to imprecise evidence that arises naturally when evidence gathering provides only partial information about the state of the world; 2) we show that in the presence of inaccuracy, updating based on a certain type of imprecise evidence can actually improve social learning performance. This type of imprecise evidence takes the form of a neighbourhood in state space centred around an agent's estimated state of the world as obtained from the environment. More specifically, we focus on the state-of-the-world problem in which agents attempt to learn the truth values of a set of propositions based on direct but inaccurate and imprecise evidence and on local interactions with other agents with whom they fuse their beliefs. In particular, using agent-based and robot simulation experiments, we show that the proposed set-based social learning model is robust to different degrees of evidential imprecision where the evidence received by the agents is inherently imprecise. We then propose a novel form of evidential imprecision, accompanied by an evidence-updating strategy. This strategy transforms the precise evidence obtained by agents into imprecise evidence, allowing us to regulate the desired degree of imprecision. In this case simulation experiments suggest that systems using more imprecise representations of evidence can be more robust and more accurate than those using precise representations for certain evidence and fusion rates in some circumstances.

An outline of the remainder of the paper is as follows: In Sect. 3 we describe a set-based model for social learning, including approaches to belief fusion and evidential updating. In Sect. 4, we describe a series of simulations experiments that investigate the robustness of social learning to imprecise evidence occurring as the result of partial information. As well as agent-based simulations this also includes multi-robot simulation experiments of a location classification task. Then in Sect. 5 we introduce imprecise evidence in the form of a Hamming distance neighbourhood of an estimated state of the world, where the latter is obtained as direct evidence from the environment. In Sect. 6 we explicitly consider robustness of performance to varying rates at which evidence is obtained and at which fusion is carried out, and show that imprecise evidence neighbourhoods can lead to more robust performance. Section 7 applies imprecise evidence neighbourhoods to the location classification problem introduced in Sect. 4. Finally, we give some conclusions and outline possible future directions in Sect. 8.

2 Background

Social learning is common in social animals, like ants and bees, where individuals learn collectively by both observation and imitation of others (Heyes, 1994). Decentralised social learning and decision-making are also well studied in swarm robotics (Brambilla et al., 2013). One common family of consensus formation problems are best-of-n problems, in which a swarm of robots collectively identifies the best option out of n possible options on the basis of local interactions and limited feedback (Parker & Zhang, 2009). An extensive overview of these problems can be found in (Valentini et al., 2017). Furthermore, social learning is not just limited to consensus formation; it can also shape individual behaviours. For instance, social feedback from peers can influence the altruistic attitude of individuals in a resource sharing scenario (Rausch et al., 2020a). This paper will instead focus on a different social learning task in which the system must learn a complete description of the true state of its environment.

We consider social learning in terms of two distinct processes; evidential updating and belief fusion (Crosscombe & Lawry, 2021). The former is the the process by which individuals learn directly from the environment, by updating their current beliefs based on evidence received from the environment. In robotic applications evidence might take the form of signals received by various sensor modalities, such as cameras, microphones, and ultra-sound sensors. 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 based on the well known bounded confidence model (Hegselmann et al., 2006), which suggests that the whole society can end up reaching a consensus on the truth assuming an appropriate confidence level and evidence rate. In addition, in the context of social epistemology, Douven and Kelp (2011) argue that the communication between agents can significantly enhance the performance in truth approximation due to the ability to correct errors while propagating information across the population. For problems of this type it has been argued that approaches combining individual evidence collection and local fusion of beliefs between individuals are more robust to noise and more efficient than those that rely on evidence collection only (Douven & Kelp, 2011; Crosscombe et al., 2017; Lee et al., 2021). In this paper we investigate the impact of evidential imprecision on social learning by proposing a model in which the level of imprecision of the evidence obtained by the agents varies. We then implement this model using both agent-based and robotic simulations.

The idea that imprecise evidence can sometimes be beneficial in social learning is also found in the study of social insects. The honeybee's waggle dance to communicate information about potential food sources is imprecise in its direction indication, and this results in variation in the angles indicated when the dance is repeated. This imprecision can be observed in either a series of waggle runs by a single dancing bee or between individuals with there typically being a variation of between 10° and 15°(De Marco et al., 2008). The "tuned error" hypothesis (Towne & Gould, 1988) states that this imprecision is used for spreading recruits over a certain spatial configuration and is selected for by natural selection. Weidenmüller and Seeley (1999) reported experiments that supported this hypothesis by finding smaller divergence angles in dances indicating potential home sites than in dances indicating food sources. The former are always point locations and the latter are often patches. However, other studies shed doubt on this hypothesis by suggesting that the imprecision of honey bee dances is the result of physical constraints and the honeybees' limited capacity to perform what is a complex sensory task (Couvillon et al., 2012; Preece & Beekman, 2014; Tanner & Visscher, 2006). Despite the debate concerning this hypothesis, simulation studies have suggested that an imprecision level of 10° is beneficial while a higher degree of imprecision is only beneficial when the food source is scarce (Okada et al., 2014).

Ants have also been reported to take advantage of noise in environmental information in order to adapt to dynamic environments. Dussutour et al. (2009) suggested that the imprecision caused by variation in ants' trail following behaviour between individuals, as well as by individuals' behaviour over time, plays an important role in enabling the effective tracking of changes in the environment. Meyer et al. (2017) suggests that the crucial role of imprecision is not tied to a particular organism or species but can be relevant across a wide range of systems.

In collective intelligence, noise is usually studied in terms of the speed-accuracy trade-off of collective decision-making tasks (Valentini et al., 2017; Douven, 2019). Collective decision-making tasks usually require a system-level balance between speed and accuracy. Methods for making reliable decisions are usually slower than those for which a greater degree of inaccuracy is tolerated. For example, using a larger neighbourhood when applying majority rule on social networks can speed up the process of reaching consensus, however this can also reduce accuracy in collective decision-making (Valentini et al., 2016). Khaluf (2022) proposed a weighted fusion model where the weights are determined by the robots' self-classification, which is based on the divergence of their beliefs compared to their neighbours. This method can enhance consensus accuracy, especially in the presence of sensor noise, albeit at the expense of additional time. In this paper we investigate how explicitly incorporating a certain kind of imprecise evidence into the evidential updating process can lead to improved performance in social learning when there is sensory error. The idea is that agents obtain an estimate of the true state of the world as direct evidence from interacting with the environment. However, they then update so as to also allow states that are sufficiently similar to the estimated state to be taken into account. Sufficient similarity between states in this context will then be modelled as a threshold on distance between different states, and form a parameter of the proposed social learning method.

Belief sets are a direct way to represent beliefs and evidence with varying levels of imprecision, in which an agent's belief is represented by the set of states or worlds that it regards as being possible (Dubois et al., 2016). Therefore, the cardinality of the set is a straightforward measure of the imprecision of an agent's belief or evidence. As such, we adopt a set-based learning model in which beliefs about the world and evidence collected by agents are represented by sets of states. Belief sets are one of the simplest formalisms for representing uncertainty and imprecision in AI. The use of sets of states to represent beliefs, sometimes referred to as 'epistemic sets', dates back to Hintikka's possible worlds semantics (Hintikka, 1962), with some early applications in AI and computer science found in Vardi (1989) and Ruspini (1987). More recently, epistemic sets have also been applied to a simple social learning problem using an abstract agent-based simulation (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). Opinion diffusion logic is relevant in this context as it sometimes employs a semantic model of belief equivalent to epistemic sets (Schwind et al., 2015). In particular, an agent's belief or evidence can be represented by a logical formula F which has a model theoretic representation equivalent to the epistemic set consisting of those states (or interpretations) in which F is true (Cholvy, 2018). The imprecision of an epistemic set then depends on the generality of the associated formula. For example, let l_1 and l_2 denote propositions asserting that there is a wildfire in locations 1 and 2, respectively. In this context a state of the world is a pair of truth values, each of which is either 0 or 1, indicating whether l_1 and l_2 are false or true. In this case evidence $l_1 \wedge l_2$ asserts that there are wildfires in both locations, while evidence $l_1 \vee l_2$ asserts that there is a wildfire in at least one of the two locations. Clearly the second formula is more general than the first and this is reflected in the relative imprecision of the associated epistemic sets as given by $\{\langle 1, 1\rangle\}$ and $\{\langle 1, 1\rangle, \langle 1, 0\rangle, \langle 0, 1\rangle\}$, respectively.

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 Lawry et al. (2019) suggested the fusion process can help to both propagate correct information while facilitating error correction provided that belief imprecision is linked to evidence collection. 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, an overview of which is given by Osswald and Martin (2006). The robustness to noise of several of these operators applied to the best-of-n problem has been compared by 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 & Prade, 1988). Bartashevich and Mostaghim (2021) expanded their exploration to include a broader range of operators from the existing literature, further substantiating the potential of fusion rules in evidence theory for robust collective intelligence. In general, the fusion process provides a means of resolving inconsistencies between different sources and hence achieving consensus. Several pairwise operators for fusing beliefs have been proposed and Dubois et al. (2016) suggest a number of desirable properties that any information fusion process should satisfy.

3 Agent-based social learning model

Consider a population of agents cooperating in an attempt to learn the state of their environment. We assume that the state of their environment can be described by a finite set of propositions $\mathcal{P} = \{p_1, \dots, p_n\}$. From this perspective a state s is the allocation of Boolean truth values to each of the propositions. Let $\mathbb{S}$ denote the set of all states so that $|\mathbb{S}| = 2^n$. An agent's belief $B \subseteq \mathbb{S}$ is the set of states which they believe can possibly be the true state of the world. Imprecise beliefs are then the subsets of $\mathbb{S}$ with cardinality $|B| > 1$, while a singleton belief $B = \{s\}$ indicates 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. Therefore, agents' beliefs are constrained 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 otherwise *uncertain*. Hence, the more imprecise an agent's belief, the more propositions the agent believes to be uncertain. This indicates a natural relationship between the epistemic model of beliefs and three-valued approaches (Crosscombe et al., 2017). For example, consider a search and rescue scenario with 8 locations where p_i denotes the proposition 'casualties are in location i' for $i = 1, \dots, 8$. Now consider the belief B given by:

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

In this case 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, $p_3, \dots, p_8$ are classified as being certainly false, and p_1 and p_2 are uncertain.

We now introduce the pairwise fusion operator for combining agents' belief sets proposed by Dubois and Prade (1988, eq. 56):

$$B_1 \odot B_2 = \begin{cases} B_1 \cap B_2 & : \; B_1 \cap B_2 \neq \emptyset \\ B_1 \cup B_2 & : \; B_1 \cap B_2 = \emptyset \end{cases}. \tag{1}$$

In addition to fusing with other agents, agents receive direct information from the environment as evidence. For evidential updating, we assume that the evidence takes the form of a set of assertions about the true state of the world. Suppose, for example, agents receive a set of states that can possibly be the true state. Given the form of E we then propose that an agent updates its belief B to $B|E$ by applying the operator as follows:

$$B|E = \begin{cases} B \cap E & : B \cap E \neq \emptyset \\ B & : \text{otherwise} \end{cases}. \tag{2}$$

This method of evidence updating in which certain states are ruled out as part of the learning process has already been applied effectively in social learning (Lawry et al., 2019). Since evidence corresponds to a set of states, then we can naturally measure its imprecision by its cardinality, i.e. the greater $|E|$ the more imprecise the evidence. For instance, consider a search and rescue scenario involving 8 distinct locations. If we have $E_1 = \{\langle 1,0,0,0,0,0,0,0\rangle, \langle 0,1,0,0,0,0,0,0\rangle\}$ and $E_2 = \{\langle 0,1,0,0,0,0,0,0\rangle\}$, then $|E_1| = 2 > |E_2| = 1$. In other words, E_1 is evidently more imprecise than E_2. In this context:

- E_1 suggests that there are casualties either in location 1 or location 2 (but not in both), and no other locations have casualties.
- E_2, on the other hand, is a more precise evidence set and indicates that there are casualties only in location 1, with no casualties reported in any other location.

In order to investigate the impact of imprecision on social learning we will focus on the state-of-the-world problem in which agents attempt to learn the truth values of a set of propositions. This will be based both on direct but potentially inaccurate and imprecise evidence, in conjunction with local interactions between the agents whereby they fuse their beliefs. In Sect. 4 we now investigate how tolerant this form of social learning is to the presence of imprecise evidence which arises naturally as the result of only partial information being obtained during evidence gathering.

4 Social learning with imprecise evidence

In this section, we introduce two simulation scenarios where imprecise evidence is naturally encountered by the agents/robots. In Sect. 4.1, the agents receive error-free but imprecise evidence. In other words, the cardinality of the evidence set will vary while the true state of the world will always be included. We then introduce a robotic classification task where the imprecision of evidence is dependent on the number of locations visited by robots during a single evidence collection episode.

4.1 Agent-based simulation with random imprecision

This section describes simulations of a multi-agent system with $k = 100$ agents attempting to socially learn the truth values of $n = 8$ propositions. The agents are initialised without any prior knowledge, i.e., they deem every state as being equally possible so that $B = \mathbb{S}$. In order to investigate the model's robustness to imprecise evidence, the degree of evidence imprecision received, as measured by $|E|$, is varied. For example, we might think of this as occurring when agents gradually learn facts about the world from different sources and expressed as logical formulas with different levels of generality. In this section, we only consider an error-free scenario such that the evidence obtained by the agents will be guaranteed to include the true state of the world s^*, while each of the other states has equal probability of being sampled without replacement. The number of states that are not s^* included in the evidence set will depend on a model parameter controlling the level of imprecision.

In practice, both evidence and agent interactions may be sparse or limited. We model this probabilistically as follows: Each agent conducts evidential updating and belief fusion iteratively. During each iteration, every agent has probability ρ (the *evidence rate*) of successfully obtaining evidence from the environment, and the agent will stop collecting evidence if it is certain about every proposition, i.e. it has a singleton belief. Agents also learn from the evidence being gathered by their peers using belief fusion. Every agent in the population has probability σ (the *fusion rate*) of being placed in a pool to fuse their belief with the belief of another agent. Each agent within the pool will be randomly paired with another agent and then each pair will combine their beliefs using the fusion operator in Equation (1). For every pair, both agents will adopt the result of this fusion as their new belief. If the number of agents in the pool is odd, then one agent will not take part in fusion.

The evidence received by the agents will be modelled as follows. The degree of the imprecision of evidence is pre-defined and varied to investigate the robustness of the set-based social learning model to different degrees of evidential imprecision. The evidential imprecision is constant across all iterations and across the population. In this section, we present error-free simulation results of which the true state of the world is guaranteed to be included in the evidence sets and the rest of the states will be randomly selected until the pre-defined evidential imprecision $|E|$ is satisfied.

We evaluate performance on this learning task using 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:

$$H(s_1, s_2) = \sum_{i=1}^{n} |s_1(p_i) - s_2(p_i)|. \tag{3}$$

.

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

$$H(B, s^*) = \frac{1}{|B|}\frac{1}{n}\sum_{s \in B} H(s, s^*). \tag{4}$$

Furthermore, we evaluate the performance at the population level, α, as the average Hamming distance between the population of agents $\mathcal{A}$ of size k, and s^*, re-scaled as an accuracy measure such that:

$$\alpha(\mathcal{A}, s^*) = 1 - \frac{1}{k}\sum_{B \in \mathcal{A}} H(B, s^*). \tag{5}$$

This accuracy measure α is ranging from 0 to 1. $\alpha = 0$ represents the belief is completely incorrect and $\alpha = 1$ indicates the belief is completely correct about every proposition. The population's initial belief is set to be completely uncertain and therefore the system starts from an accuracy of $\alpha = 0.5$.

In Fig. 2 we show the average accuracy α of the agents' beliefs against time t for different values of evidential imprecision $|E|$ and different combinations of evidence rate, ρ, and fusion rate, σ. For each combination of parameters we ran the experiments 50 times and

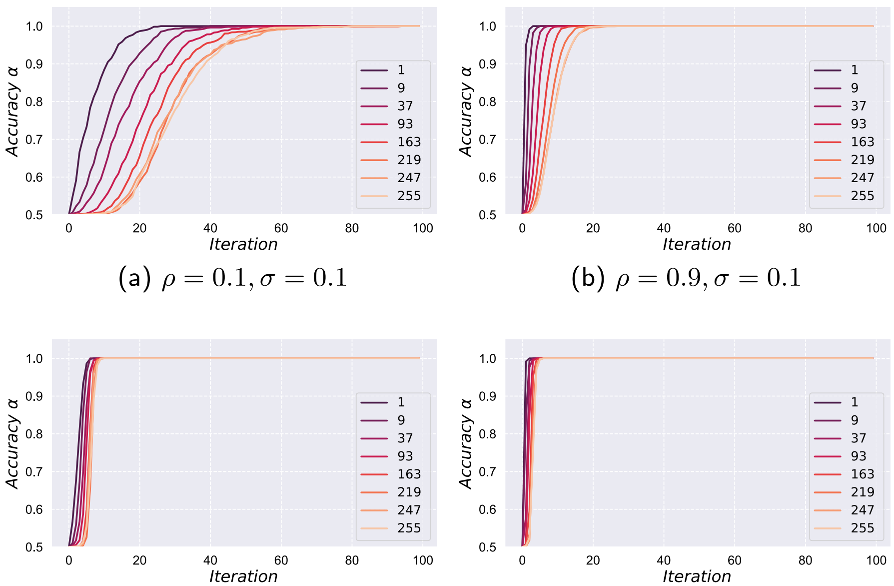

(a) $\rho = 0.1, \sigma = 0.1$

(b) $\rho = 0.9, \sigma = 0.1$

(c) $\rho = 0.1, \sigma = 0.9$.

(d) $\rho = 0.9, \sigma = 0.9$.

Fig. 2 Average accuracy at steady state against time t for evidence rate $\rho \in \{0.1, 0.9\}$, fusion rate $\sigma \in \{0.1, 0.9\}$, and evidential imprecision $|E| = 1, \ldots, 255$

the results are averaged over those 50 runs. We see that the population converges to $\alpha = 1$ at steady state for every combination of ρ, σ, and $|E|$. Since the only belief for which $\alpha = 1$ is $B = \{s^*\}$ it follows that the population has reached consensus by correctly and precisely identifying the true state of the world. Furthermore, this suggests that this social learning model is robust to various levels of evidential imprecision, provided that $s^* \in E$.

In Fig. 2a we see that the population converges more slowly as evidence imprecision increases, such as when evidence is sparse ($\rho = 0.1$) and fusion is infrequent ($\sigma = 0.1$). For $|E| = 1$, the population converges quickly to consensus by around $t = 25$. Convergence time increases to around 80 time steps when $|E| = 255$, this corresponding to the most imprecise non-vacuous evidence for $n = 8$ propositions since only one state of the world is emitted. In a scenario where evidential updating is frequent while fusion is infrequent, as shown in Fig. 2b, the population converges faster for all levels of evidence imprecision than when evidential updating is infrequent; within 10 time steps for $|E| = 1$ and within 25 iterations for $|E| = 255$. When evidence is sparse and fusion is more frequent, Fig. 2c shows that the population also reaches consensus faster than in Fig. 2a whereas all levels of imprecision converge after around the same number of time steps. In Fig. 2d, both evidential updating and fusion are frequent, resulting in fast convergence to the true belief, with little difference between the different levels of evidential imprecision. The effect of evidence imprecision on the convergence time decreases as we increase ρ or σ, as seen in Fig. 2b–d. Futhermore, from Fig. 2b, c we see that both more frequent evidence and more frequent fusion can reduce the differences in the learning speed of the population. Furthermore, the frequency of belief fusion has the greater effect. Comparing Fig. 2b, c, we see that the convergence speeds are less affected by the level of evidential imprecision when agents learns more socially ($\rho = 0.1, \sigma = 0.9$) than when they learn more independently ($\rho = 0.9, \sigma = 0.1$).

So far, the evidence received by agents in this section has been error-free. However, in real-world situations it is highly likely that the evidence received could be inaccurate; often due to noisy sensing equipment or even environmental noise. In the following section we will investigate social learning in a multi-robot scenario in which both evidential inaccuracy and imprecision are present.

4.2 Location classification task by multi-robot system

In practice, a potential source of imprecision is when agents receive only partial information about the state of the world during the evidence collection process. In this section we consider a particular form of imprecise evidence arising from partial knowledge in a multi-robot system in which robots are investigating properties of a number of different locations. During an evidence collection episode individual robots are only able to visit a limited number of locations and, hence, obtain only partial information about the full state of the world. In this context our results show that for a simple robot arena environment, social learning is robust to a range of different evidence collection bounds while being more effective than an approach based on individual, evidence-based learning alone.

We conduct simulation experiments for a multi-robot system needing to make a collective decision about the true state of the world by attempting to classify n locations as either being red or green. Specifically, the proposition p_i asserts that location *i is red* and $\neg p_i$ asserts that location *i is green*. In this case a state s_i is the Boolean allocation of the two colours to each location and beliefs are the set of the allocations deemed possible.

We now describe a series of experiments in which $k = 5$ or 10 e-puck robots investigate $n = 4$, 8 or 12 locations arranged in a circle, each coloured either red or green. This is a similar configuration to that used by Lee et al. (2021). Evidence is collected using the e-puck's in-built camera and with an additional error imposed on the classification process. To model this we introduce an error rate ϵ corresponding to the probability of receiving evidence that is inconsistent with the true state of the world. More specifically, let the true state be denoted by s^*, then in the case that an agent receives information about location $s(p_i)$ they will receive $s(p_i) = s^*(p_i)$ with probability $1 - \epsilon$, and $s(p_i) = 1 - s^*(p_i)$ with probability ϵ. Between evidence collection episodes individuals move to the centre of the arena for fusion. The goal is for the whole system to reach consensus by identifying the true state of the world, i.e., the correct colour of each location from the $2^n = 16$, 256, or 4096 possibilities for $n = 4, 8$ or 12, respectively.

We simulate e-puck robots (Mondada et al., 2009) which are well-suited to a classification task of this kind since they are equipped with a range of sensors. Experiments were performed in the CoppeliaSim[1] simulation environment which models the physical characteristics of the e-pucks, including motion, communication and sensing. Figure 3 shows the experimental arena which has $n = 8$ locations equally distributed around a 1.12 m circle with the $k = 5$ e-pucks' fusing positions spaced evenly around a 0.3 m disc at the centre. Each e-puck returns to their fusing positions after an evidence collection episode to fuse their beliefs with another robot also at their positions at that time. The robots will only be able to communicate when they are at the fusing positions, i.e., around the central area of the arena, this restriction providing a basic model of a scenario in which communication is limited.

[1] https://www.coppeliarobotics.com/

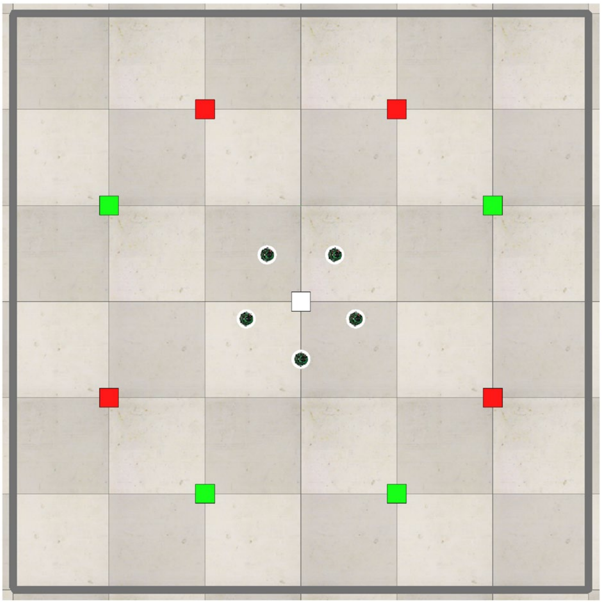

Fig. 3 Top-down view of the experimental setup for $n = 8$ sites. The red/green squares indicate the location of the sites, and the white circles show the fusion positions. An e-puck resides at each fusion position

A group of e-pucks are initialised without any prior knowledge and hence with initial beliefs corresponding to the set of all states, representing total ignorance. They are given the coordinates of all locations and relevant fusing positions and apply simple path planning to travel between sites and the fusing positions. The e-pucks featured have 8 IR proximity sensors and therefore we apply Braitenberg-based collision avoidance algorithms (Braitenberg, 1986). The system learns the environment iteratively with individuals alternating between episodes of evidence collection and fusion. During an evidence collection episode a robot visits some of the sites about which they are currently uncertain, where the number of locations visited is less than a pre-specified upper bound N_u.

This parameter N_u is defined as the maximum number of locations that a robot is able to visit during each evidence collection episode, this modelling plausible real-life constraints imposed by both robot hardware, and the scale and complexity of the environment, on distance that can be travel during evidence collection. A natural consequence of this bound is that evidence collected during each episode will typically only provide partial information in the form of noisy data concerning the class of only some of the locations. The representation of this partial information will take the form of imprecise evidence. We can then determine the impact that imprecise evidence has on the population by varying N_u. Robots will obtain precise evidence if they visited a relatively high numbers of locations since the evidence E they received would then identify a small number of possible states, i.e., have low cardinality. More specifically, the relationship between the number of sites visited v and the cardinality of the evidence set is given by $|E| = 2^{n-v}$. Each robot selects up to N_u locations from the set of locations for which either class is still possible according to their current belief B. At each location an e-puck uses its camera to capture a colour value indicating the class to which that particular site belongs. After visiting all the locations that they have selected, they update their belief according to Equation (2) and return to their pre-specified central location. Each robot then broadcasts its belief along with a flag message identifying the broadcaster as being ready to fuse. In our model we simulate the

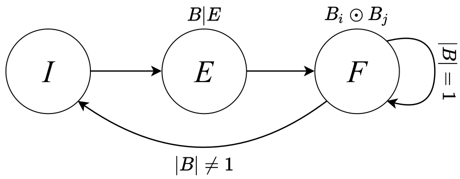

Fig. 4 The robot's state transition model. (*I*) When the robot visits all of the locations it has planned to visit. Evidential updating (*E*), and Belief fusion (*F*)

communication between e-pucks using the built-in CoppeliaSim functions while physical e-pucks can achieve the same communication via on-board Wi-Fi or Bluetooth functionality in practical applications. They also listen for any other robots currently broadcasting and fuse their belief with the transmitted belief in the first such message they receive.

Figure 4 shows a state transition diagram for the above process. Robots have 3 internal states: Investigation (*I*); Evidential updating (*E*); and Belief fusion (*F*). Robots are initialised in state *I* before they select locations to investigate and visit them in order, collecting class information as they go. Once all relevant locations have been visited the robot transitions to state *E* and updates their belief based on the evidence collected. They then transition to the state *F*, return to their fusing positions and perform fusion. If their beliefs now identify a single complete set of classifications for all the locations they remain in their fusing positions in state *F* and carry out another fusion, otherwise they transition to state *I* for another evidence collection episode.

We use the accuracy metric defined in Equation (5) to measure the performance of the whole population of robots. Notice that for evidence-only learning, that is when the belief fusion process is disabled and robots learn independently based *only* on the evidence they receive, the simulations lead to Bernoulli experiments with the success probability $p = 1 - \epsilon$. The expected value of the Hamming distance between a robot's belief and the true state is then as follows:

$$\mathbb{E}(H(B, s^*)) = \frac{1}{n} \sum_{i=0}^{n} \binom{n}{i} i(1-\epsilon)^{(n-i)} \epsilon^i = \epsilon. \tag{6}$$

We therefore employ $1 - \epsilon$ as a benchmark to illustrate the improvement in performance which results from robots interacting with one another and fusing their beliefs to achieve social learning.

We now present results from simulation experiments conducted with $k = 5$ or 10 e-pucks, $n = 4$, 8, or 12 sites and different upper bounds N_u on the number of sites visited during an evidence collection episode. For each combination of parameters we ran the experiments 20 times. The results presented herein are then averaged across runs with error bands corresponding to the 5th and 95th percentiles.

Figure 5 shows the log average cardinality $log_2|B|$ (blue line) and the average accuracy α (purple line) of the robots' beliefs against time t for $n \in \{8, 12\}$ locations with an upper bound $N_u = 3$. For all plots in Fig. 5 the dotted orange lines indicate the value of $1 - \epsilon$; i.e. the expected accuracy when agents learn individually from direct evidence alone. For instance, Fig. 5a shows results from an evidence-only learning scenario where we see that, for $n = 12$ locations, $k = 10$ e-pucks and error rate $\epsilon = 0.2$, the average cardinality of the population's beliefs decreases to 0 over time while the average accuracy of the system converges to the expected accuracy of $1 - \epsilon$. This is because, without belief fusion, agents are dependent on the evidence that they receive directly, without any communication with other agents. This leads to agents adopting erroneous beliefs about the state of the world.

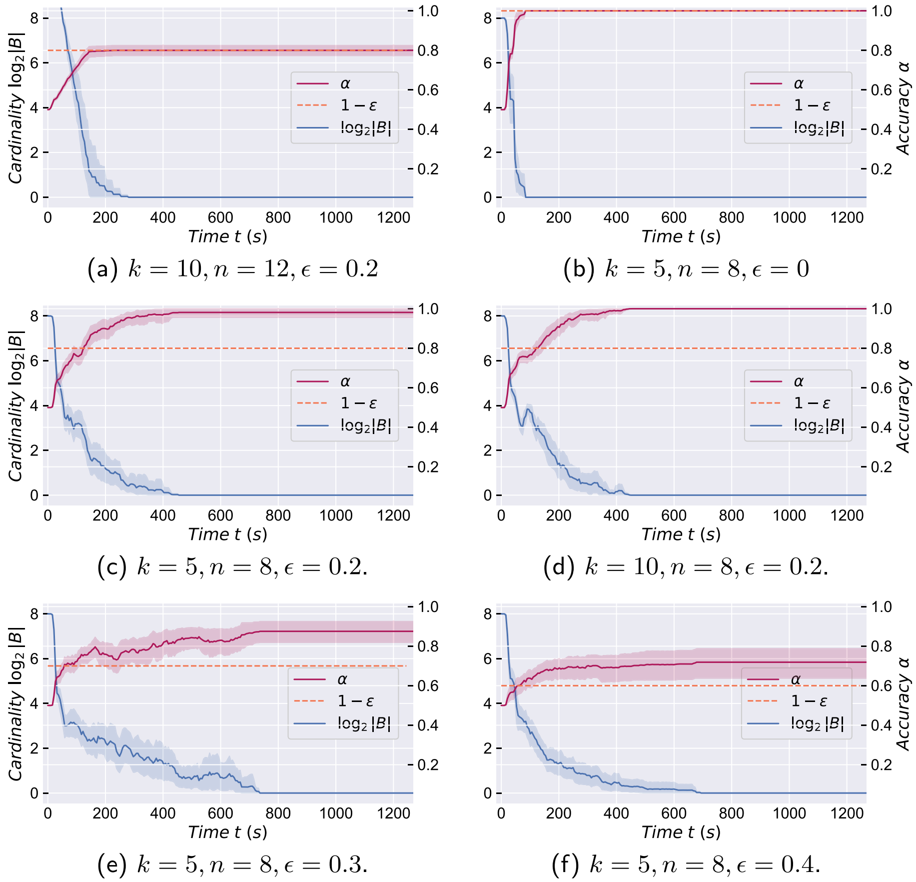

(a) $k = 10, n = 12, \epsilon = 0.2$

(b) $k = 5, n = 8, \epsilon = 0$

(c) $k = 5, n = 8, \epsilon = 0.2.$

(d) $k = 10, n = 8, \epsilon = 0.2.$

(e) $k = 5, n = 8, \epsilon = 0.3.$

(f) $k = 5, n = 8, \epsilon = 0.4.$

Fig. 5 Average log cardinality $log_2|B|$ and average accuracy α plotted against iteration for $N_u = 3$ with different numbers of agents k, locations n and error rate ϵ. **a** Evidence-only learning (without fusion). **b–f** Social learning (with fusion)

An accuracy value above the dotted line is therefore indicative of the system's ability to correct for errors as a direct result of social learning.

From Fig. 5b–f we see that, across all runs and for all parameter combinations, the robots converge to a belief of cardinality $|B| = 1$, i.e., a singleton belief, as shown by the blue lines decreasing to $log_2|B| = 0$. This means that, under our model, all of the e-pucks eventually remain stationary at their central fusing positions having reached a consensus about the state of the world that they believe to be true. Of course, the primary purpose of location classification is for the system to *accurately* identify the true state of the world. The purple lines in Fig. 5b–f show the average accuracy of the population plotted against time. Broadly, we see that the system is able to correctly classify locations with high accuracy. Starting with an error-free scenario in Fig. 5b with 8 locations and an upper bound of $N_u = 3$ as well as $k = 5$ e-pucks, the robots always identify the true state of the world s^* with accuracy $\alpha = 1$ in under 100 seconds across all 20 experiments. This is perhaps unsurprising when robots obtain perfect information during evidence collection. In Fig. 5d we show a moderately noisy scenario with $\epsilon = 0.2$ and an upper bound $N_u = 3$. With noisy

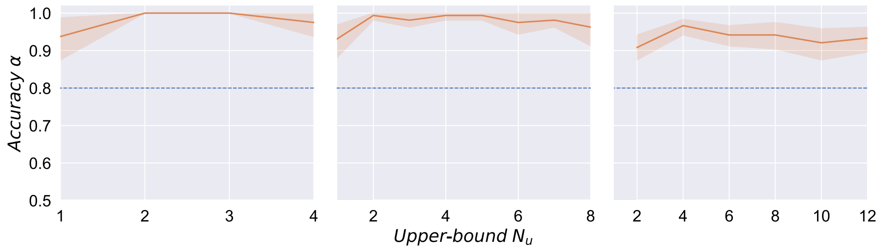

(a) $k = 5$.

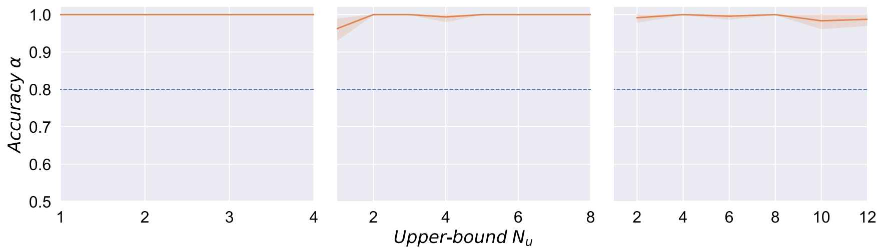

(b) $k = 10$.

Fig. 6 Average accuracy at steady state for various upper bounds N_u and error rate $\epsilon = 0.2$. **a** $n \in \{4, 8, 12\}$ from left to right with $k = 5$. **b** $n \in \{4, 8, 12\}$ from left to right with $k = 10$

evidence the system is slower to reach a consensus but still manages to converge on a belief with an avearage accuracy of 0.98 after roughly 400 seconds. With higher noise levels $\epsilon = 0.3$ and $\epsilon = 0.4$, the robots converges with a lower accuracy $\alpha = 0.85$ and $\alpha = 0.75$ and more variation across multiple runs, as shown in Fig. 5e, f. However, the systems still manage to reach consensus at a singleton belief in these higher noise level scenarios, with higher time costs, around 700 seconds.

The upper bound on the number of locations N_u also has an impact on accuracy when the population size k is small, with higher accuracy being achieved for intermediate values of N_u. For example, Fig. 6a shows the average accuracy against the upper bound N_u taking values from $[1, n]$, for $k = 5$ and $n \in \{4, 8, 12\}$, with moderate error rate $\epsilon = 0.2$. In Fig. 6a the best performance is obtained when the visit bound $N_u = 2, 3$ where, at steady state, the system achieves an accuracy $\alpha = 1$ on average. That means the system reaches consensus on s^* across all the 20 experiments. For $N_u = 1$ and $N_u = 4$, however, the accuracy falls to around 0.87 and 0.95, respectively. More generally, good performance across all N_u values shows that this form of social learning is still accurate under imprecise evidence. Although, there is some reduction in performance for the most imprecise case when a maximum of only 1 location is visited in each evidence collection episode.

Increasing the number of robots k can improve the overall performance of social learning and decrease the variance in accuracy, as shown in Fig. 6a, b. The system achieves greater accuracy for all upper bounds with $k = 10$ robots than $k = 5$ robots. The performance is also less variant across simulations and across different N_u.

Overall, the results of the robot simulation experiments show that this form of multi-robot social learning achieves high accuracy across all different imprecision levels. The approach therefore has the potential to be effectively applied to location classification tasks

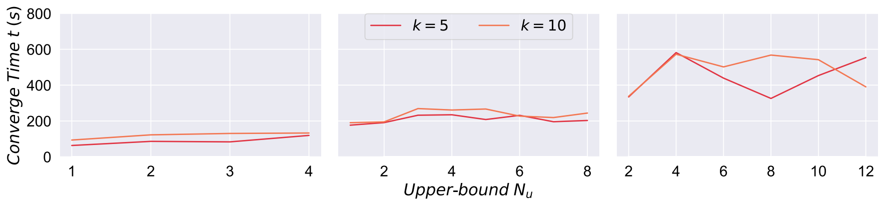

Fig. 7 Average time to convergence against upper bound N_u for $\epsilon = 0.2$, $k \in \{5, 10\}$ with $n \in \{4, 8, 12\}$ from left to right

conducted by multi-robot systems. In this approach, each robot does not have to investigate every location for the system to reach consensus. There can be good performance in scenarios in which access to some locations is restricted, either by the range of the robots (e.g., due to power constraints) or a heterogeneous system possessing different levels of access or capabilities. Furthermore, the approach scales well to scenarios in which the number of locations is greater than the number of robots.

Figure 7 shows the average convergence time for the system against N_u. This is the time it takes for the robots to reach a consensus. For $\epsilon = 0.2$ we see a relatively consistent convergence time for different N_u, which demonstrates the robustness of the social learning model to noisy evidence and insensitivity to the level of evidential precision. In general, the time cost of the classification task is insensitive to both the population size and different constraints on the number of locations that can be visited in an evidence collecting episode. In other words, the convergence time is insensitive to the different levels of evidential imprecision.

In Sect. 4 we have shown that social learning can perform well when evidence is imprecise. This is the case even when the level of imprecision is very high or if imprecision is combined with noise. In other words, social learning can be robust to evidence that is both imprecise and noisy. In the following sections we show that social learning performance can actually be improved by incorporating imprecise evidence of a particular type and level into the belief updating process.

5 The benefits of imprecise evidence in social learning

In this section, we introduce a method to use evidence imprecision as a design parameter for social learning and conduct agent-based simulations to investigate the resulting performance in social learning. Here we assume the information received by the agent is precise, identifying a single state s_e which may or may not deviate from the true state of the world, depending on whether or not there is environmental noise i.e. $\epsilon > 0$. Let $H_e = H(s^*, s_e)$ denote the Hamming distance between the true state of the world and the state identified during evidence collection. In the case that the noise $\epsilon > 0$ then H_e is a random variable with the following probability distribution:

$$P(H_e = i|\epsilon) = \binom{n}{i} \epsilon^i (1 - \epsilon)^{n-i}. \tag{7}$$

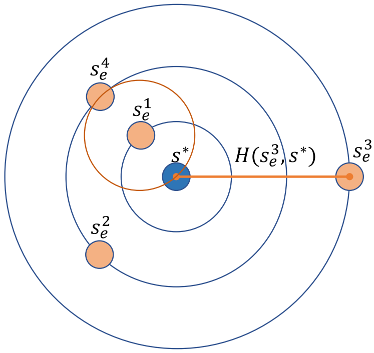

Fig. 8 Diagram showing the relationship between the observed state s_e and the true state s^*

Based on Equation (7), the maximum likelihood estimation of H_e is $\hat{H}_e = 0$, $\hat{H}_e = 1$, and $\hat{H}_e = 2$ for $\epsilon = 0.1, 0.2,$ and 0.3 respectively if $n = 8$. In other words, the probability that $s^* = s_e$, i.e. $H_e = 0$, can be relatively small compared to $s^* \neq s_e$ ($H_e > 0$) if the values of n and ϵ are high. In Fig. 8, s_e^1, s_e^2, and s_e^3 are three states independently collected as evidence assuming $\epsilon > 0$. All three states are different from the true state s^* with the distance shown as concentric circles centred on s^*. However, if we consider imprecise evidence in the form of a neighbourhood of the evidence states, then as the radius of that neighbourhood increases so will the probability that s^* is contained in the evidence set (see the brown circle around s_e^1). We hypothesize that this increase in probability of the evidence being consistent with the true state of the world can be potentially beneficial in social learning.

In general, given an estimated state of the world s_e obtained from evidence then an associated imprecise evidence set $E(s_e, \tilde{H})$ can be defined as a Hamming distance neighbourhood of s_e based on distance threshold $\tilde{H}$ in the following way:

$$E(s_e, \tilde{H}) = \{s : H(s, s^e) \leqslant \tilde{H}\}, \text{where } \tilde{H} \in \{0, \dots, n-1\}. \tag{8}$$

We use $\tilde{H}$ as a representation parameter to control the imprecision of the evidence received by the agents. An agent receives precise evidence s_e and uses it as an estimate of the state of the world. This is generated by independently sampling the truth value of each proposition and then recording the value 1 with probability $1 - \epsilon$ and 0 with probability ϵ, based on the assumption that the actual truth value is 1 for each proposition. The cardinality of the evidence increases with $\tilde{H}$, for example, for n propositions, the cardinality of the evidence is $|E(s_e, \tilde{H})| = \sum_{i=0}^{\tilde{H}} \binom{n}{i}$.

We now describe a number of agent-based simulation experiments to investigate the effect of varying the threshold $\tilde{H}$ on social learning. Here we will assume a population of $k = 100$ agents investigating the truth-values of $n = 8$ propositions. The evidential updating and belief fusion method as defined in Equation (1) and Equation (2) are also applied in this simulation with the evidence rate $\rho \in [0.02, 1)$ and the fusion rate $\sigma \in [0.02, 1)$. The population of agents are initialised as having no prior knowledge about the world and hence hold completely ignorant beliefs, i.e., $B = \mathbb{S}$, at time $t = 0$. Experiments are run 50 times to account for variation in performance.

Figure 9 shows heat maps of average accuracy for varying fusion and evidence rates, at different levels of evidential imprecision as parameterised by different Hamming distance thresholds and different error rates. For each single heat map, with a step

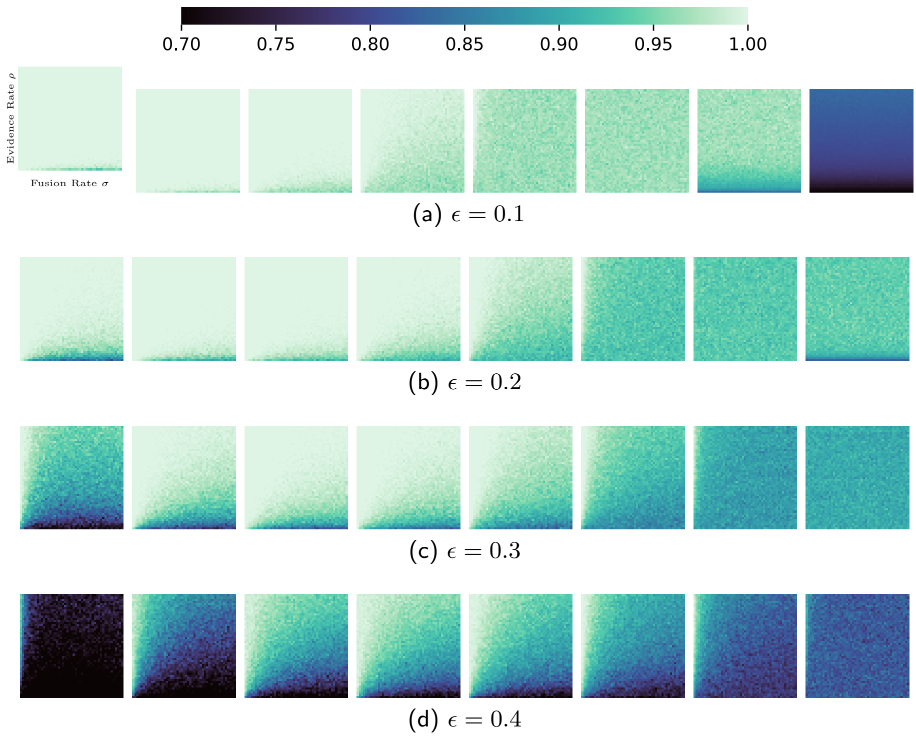

Fig. 9 Average accuracy α at steady state for different evidence imprecision for different error rates $\epsilon \in \{0.1, 0.2, 0.3, 0.4\}$, fusion rates $\sigma \in [0.02, 1)$, and evidence rates $\rho \in [0.02, 1)$. From left to right: $\tilde{H} \in \{0, \ldots, 7\}$

size of 0.02 we have evidence rates $\rho \in [0.02, 1)$ for the vertical axis and fusion rates $\sigma \in [0.02, 1)$ for the horizontal axis. The top-left plot in Fig. 9 includes axis labels for clarity. The Hamming distance threshold increases from left to right for Fig. 9a–d. Here lighter and darker colours indicate higher and lower accuracy, respectively. Across the heat maps, when evidence is of low to intermediate imprecision, we see that the system is more accurate when the evidence rate is relatively high in relation to the fusion rate. In other words, lower fusion rates increase the system's robustness across various evidence rates, error rates, and levels of evidential imprecision. On the other hand, in areas where the evidence rate is relatively low compared to the fusion rates, the system requires significantly low fusion rates to reach higher accuracy. In regions of the parameter space in which the fusion rate is higher than the evidence rate, the learning accuracy is improved by increasing the Hamming threshold $\tilde{H}$ to different levels, depending on different error rates. Beyond the certain points, e.g. $\tilde{H} = 3$ for $\epsilon = 0.3$, additional increases in the Hamming threshold tend to reduce overall accuracy while making it more consistent across various regions of the parameter space. As the level of evidential imprecision increases, the performance across the parameter space becomes more uniform, displaying a moderate level of accuracy. There is a unique case in Fig. 9a when $\tilde{H} = 7$, where the accuracy is uniformly low. This particular scenario will be further discussed and analysed in conjunction with the subsequent figures.

From Fig. 9 we see that the optimal Hamming thresholds for achieving the best performance vary depending on the error rates. In particular, when $\epsilon = 0.3$, from Fig. 9c we see that accuracy is highest across the whole fusion and evidence rate parameter space when $\tilde{H} = 2$ or $\tilde{H} = 3$. For higher levels of imprecision there is then a relatively uniform decrease in performance across the parameter space; see particularly the first column from the right when $\tilde{H} = 7$. In general, these results indicate that highest overall accuracy is obtained for moderate levels of imprecision when using the proposed neighbourhood approach. For lower error rates, Fig. 9a, b show that optimal performance is achieved with relatively low Hamming thresholds of $\tilde{H} = 1$ and $\tilde{H} = 2$ for $\epsilon = 0.2$, and $\tilde{H} = 0$ and $\tilde{H} = 1$ for $\epsilon = 0.1$. This suggests that at lower error rates, more precise evidence can result in high overall accuracy. In contrast, as shown in Fig. 9d, at a higher error rate of $\epsilon = 0.4$, the model suggests that best performance is achieved with considerably higher Hamming thresholds of $\tilde{H} = 3$ and $\tilde{H} = 4$, i.e. in scenarios where the noise level or error rate is higher, more imprecise evidence can be advantageous for improving the overall accuracy. Hence, in general these results suggest a pattern of performance in which the higher the noise, the higher the level of imprecision at which the best accuracy is obtained. In summary, as observed from Fig. 9, the highest overall accuracy across different evidence and fusion rates, and in the presence of significant inaccuracy, is achieved at intermediate levels of evidential imprecision, for which we see more lighter colours. Furthermore, the optimal level of imprecision increases with higher error rates.

A common trade-off in social learning is between speed of learning and accuracy of learning; known as the speed vs. accuracy trade-off. This has been studied extensively across the collective intelligence literature from insect swarms to swarm robotics (Valentini et al., 2016). In the context of evidence neighbourhoods, we can consider the impact of different levels of evidential imprecision on time to convergence. We define convergence as all agents reaching a consensus on a singleton belief, i.e. all agents agree that the singleton belief represents the true state of the world s^*. Figure 10 shows the time steps to consensus for different Hamming threshold values and error rates, with lighter colours indicating less convergence time steps. A maximum limit of 20, 000 time steps is set for each simulation, i.e. this will be used as the convergence time for any simulation run that fails to converge within 20, 000 time steps to get the average value. It is significant to note that the system exhibits increased robustness to varying evidence rates and error rates when the fusion rate is low, as demonstrated in Fig. 9. However, a trade-off exists, as lower fusion rates requires a greater number of time steps for agents to achieve consensus, as shown in Fig. 10. For example, for $\tilde{H} = 2$, comparing Figs. 9c to 10c, for low fusion rates $\sigma \in (0, 0.1)$, agents reach consensus after 100 to 300 time steps with an accuracy close to 1, even for very low evidence rates. Whereas for low evidence rates and the highest fusion rates $\sigma \in (0.9, 1)$, the average accuracy is much lower (around 0.8) but a consensus is reached within just 20 time steps.

From Fig. 10 we also see that the system exhibits slower convergence with increasing error rates when $\tilde{H} \leq 3$, whereas convergence accelerates with higher error rates when $\tilde{H} \geq 6$. For $\tilde{H} \in \{4, 5\}$, the most rapid convergence is observed at medium error rates, specifically when $\epsilon = 0.2$ for $\tilde{H} = 4$ and $\epsilon = 0.3$ for $\tilde{H} = 5$. One possible explanation for this phenomenon is that, under conditions of highly imprecise evidence and low error rates, agents are more likely to receive similar imprecise evidence at different time steps. Consequently, their beliefs remain imprecise following the fusion process. As a result, a greater number of time steps are required to accumulate sufficient variation through errors, which in turn drives the cardinality down and facilitates convergence. This insight highlights the

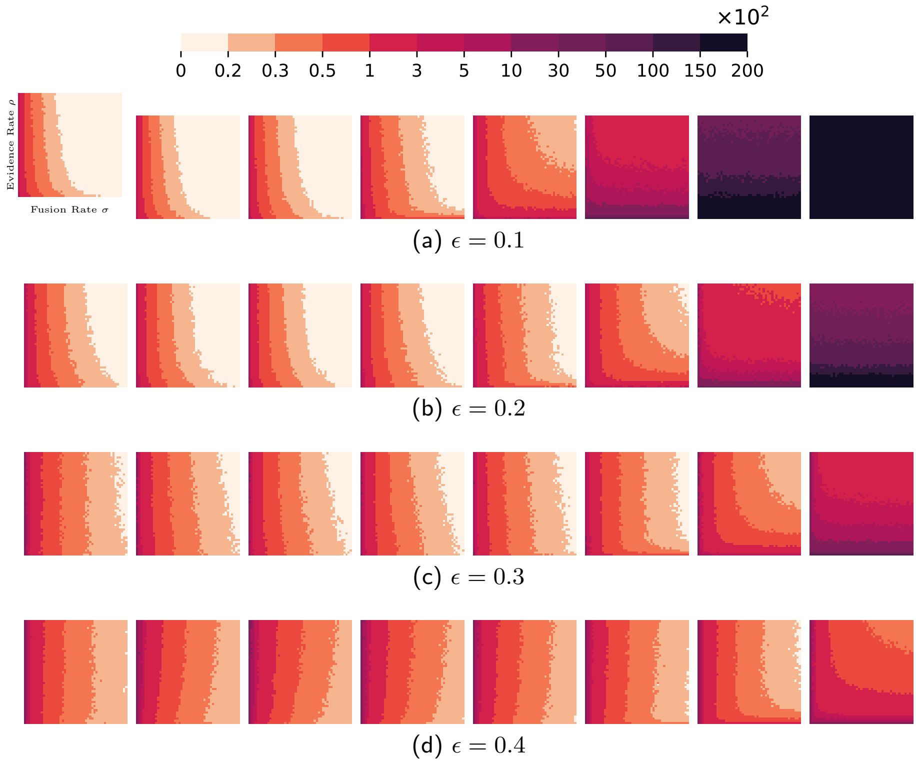

Fig. 10 Average number of time steps until convergence τ at steady state for different evidence imprecision defined by various Hamming thresholds $\tilde{H}$ and error rate $\epsilon \in \{0.1, 0.2, 0.3, 0.4\}$, fusion rates $\sigma \in [0.02, 1)$, and evidence rates $\rho \in [0.02, 1)$. From left to right: $\tilde{H} \in \{0, \ldots, 7\}$

intricate relationship between error rates, evidence imprecision, and system dynamics in the context of social learning models.

6 Imprecision and robustness in social learning

Robustness to variation in underlying conditions is important in this social learning context since environments are often dynamic and our knowledge of them is usually limited. For example, evidence and fusion rates may be varying and difficult to predict in advance since different factors may influence agents' capacity to collect evidence or interact with each other during the learning process. In order to evaluate the influence of different levels of evidential imprecision on the robustness to different evidence and fusion rates, we can use aspects of info-gap theory proposed by Ben-Haim (2006). Info-gap theory provides theoretical tools to aid decision-making under severe uncertainty, by analysing robustness to variation around a set of estimated parameter values representing the best available knowledge of the underlying conditions of the system. We apply the info-gap theory to evaluate variation in the fusion and evidence rate, σ and ρ.

Suppose we have estimates of the evidence and fusion rates for a given social learning problem denoted by $\hat{\rho}$ and $\hat{\sigma}$ respectively. Let $U(h)$ denote a neighbour of $(\hat{\sigma}, \hat{\rho})$

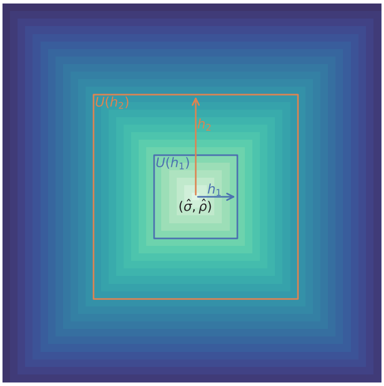

Fig. 11 Diagram showing the horizon of uncertainty in Info-Gap Theory as a neighbourhood of the estimated fusion and evidence rates $(\hat{\sigma}, \hat{\rho})$

in the parameter space of size h. This is referred to in info-gap theory as an horizon of uncertainty:

$$U(h) = \{(\sigma, \rho) \in (0,1)^2 \, : \, |\sigma - \hat{\sigma}| \leqslant h, |\rho - \hat{\rho}| \leqslant h\}. \tag{9}$$

The robustness at $(\hat{\sigma}, \hat{\rho})$ is then defined as the size of the largest horizon of uncertainty for which the average learning error $1 - \alpha$ is guaranteed to not exceed a critical maximum value δ. For different values of δ we then have the following robustness function:

$$\hat{h}(\delta) = \max\{h \, : \, m(h) \leqslant \delta\}, \tag{10}$$

where $m(h) = \max\{|1 - \alpha(\sigma, \rho)| \, : \, (\sigma, \rho) \in U(h)\}$ is the maximum error across all parameter values in the horizon of uncertainty of size h.

Figure 11 illustrates the application of info-gap theory to the current context. Suppose that $\hat{\sigma}$ and $\hat{\rho}$ are at the centre of the parameter space i.e. $\hat{\sigma} = \hat{\rho} = 0.5$. Let $\hat{h}_1$ and $\hat{h}_2$ be robustness functions for two different algorithms. The if $\hat{h}_1(\delta) \geq \hat{h}_2(\delta)$ for all δ then we say that algorithm 1 robust dominates algorithm 2 at parameter estimates $\hat{\sigma}$ and $\hat{\rho}$. In other words, for every tolerance level δ there is larger neighbourhood of $(\hat{\sigma}, \hat{\rho})$ for which the error tolerance constraint is guaranteed to be met for algorithm 1 than there is for algorithm 2. This is a clear indication that under these conditions the performance of algorithm 1 is more robust to variation in fusion and evidence rates than algorithm 2 for all tolerance levels. On the other hand, if the robustness curves $\hat{h}_1$ and $\hat{h}_2$ cross then this suggests that there are some levels of tolerance at which algorithm 1 is the most robust and some at which algorithm 2 is.

Figure 12 shows the robustness curves $\hat{h}(\delta)$ for different Hamming thresholds $\tilde{H}$, and assuming $(\hat{\sigma}, \hat{\rho}) = (0.5, 0.5)$. In both Fig. 12a, b the robustness curves for $\tilde{H} = 1$ and $\tilde{H} = 2$ (pink and purple lines) are everywhere greater than the curve for $\tilde{H} = 0$. This indicates that this moderate level of imprecise evidence is more robust for all tolerance levels δ than precise evidence. In Fig. 12b where the noise is $\epsilon = 0.2$ the robustness curve for the high imprecision level of $\tilde{H} = 7$ is dominated by the robustness curves for $\tilde{H} = 1$ and $\tilde{H} = 2$ showing that it is the least robust model of these imprecise evidence models. On the other hand, it crosses the curve for $\tilde{H} = 0$ at $\delta = 0.1$ showing that it is slightly more robust than the most precise evidence model if the tolerance to error is higher than 0.1. However, in Fig. 12a showing the higher noise level $\epsilon = 0.3$, the robustness curve for $\tilde{H} = 7$ crosses all other robustness curves so that there is a small range of higher tolerance values for which it is the most robust imprecise evidence model. In other words, in this case the highly

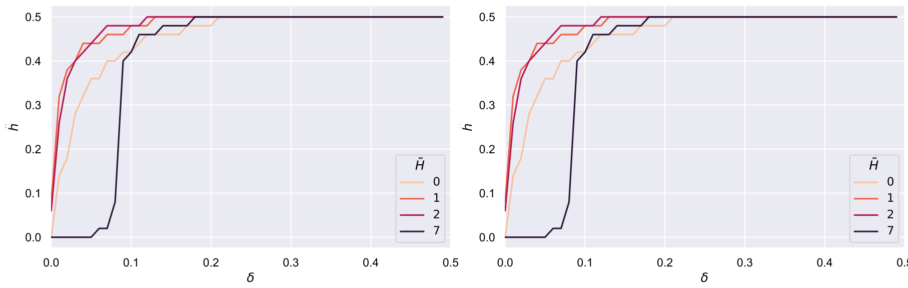

(a) $\epsilon = 0.3$ (b) $\epsilon = 0.2$

Fig. 12 Robustness curves for various evidence imprecision levels and $\hat{\rho} = 0.5, \hat{\sigma} = 0.5$, **a** $\epsilon = 0.3$; **b** $\epsilon = 0.2$

imprecise model is more robust than other more precise evidential models if the tolerance to error is relatively high. Overall, this suggests that in noisy environments the level of error that is acceptable will play a role in deciding which level of evidential imprecision is most robust.

7 The benefits of imprecision in multi-robot systems

The agent-based simulations described in Sect. 5 have shown that evidential imprecision defined by a Hamming distance-based method can improve the overall accuracy of social learning. In this section we conduct simulation experiments for the multi-robot location classification task described in Sect. 4.2 using the imprecise evidence model defined by Equation (8) in Sect. 5. In order to obtain an estimate of the state of the world, we let the e-pucks visit all the locations during every evidence gathering episode, i.e. during an episode a robot visits all locations, including those about which they are currently certain, to collect a sensor reading. Hence, using the notation from Sect. 5, during an evidence gathering episode a robot visits each location to obtain an estimate of its colour resulting in a n-dimensional binary vector s_e corresponding to their estimate of the colours of all n locations. The robot then updates their beliefs by conditioning on imprecise evidence in the form of the Hamming neighbourhood $E(s_e, \tilde{H})$ for some threshold $\tilde{H}$. For instance, if $\tilde{H} = 1$, $E(s_e, \tilde{H})$ then includes the estimated classes s_e together with all states in which the class of exactly one of the locations is changed from that given in s_e. Then for $\tilde{H} = 2$ states with exactly two location classes switched are also included and so on. In light of the fact that Sect. 5 suggests that high levels of evidential imprecision result in much slower convergence, in the following experiments we set a maximum simulation time of 4000 seconds.

Figure 13 shows the average learning accuracy of the multi-robot system for hamming threshold $\tilde{H} \in \{0, 1, \ldots, 7\}$, and error rates $\epsilon = 0.3$ and $\epsilon = 0.4$. We run 10 independent simulations for each combination of the parameters. Error bars represent the width of the 95% confidence interval. In Fig. 13a where $\epsilon = 0.3$, peak accuracy is observed for $\tilde{H} = 2$. Compared with the most precise evidence, the optimal thresholds not only result in a higher average accuracy but also less variation in accuracy across runs. Beyond this range, from $\tilde{H} \geq 3$, the learning accuracy starts to decrease. For a higher error rate, $\epsilon = 0.4$, the highest accuracy is achieved by $\tilde{H} = 3$, as shown in Fig. 13b. For both error rates shown in

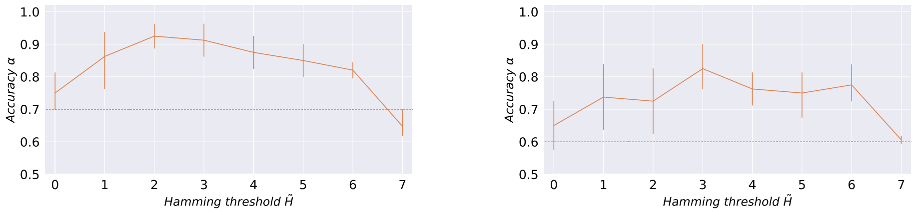

(a) $\epsilon = 0.3$ (b) $\epsilon = 0.4$

Fig. 13 Average accuracy at steady state for various Hamming thresholds $\tilde{H}$, error rate $\epsilon \in \{0.3, 0.4\}$, and $k = 5, n = 8$

Fig. 13, we see a significant improvement in the system's learning accuracy when comparing the optimal thresholds, $\tilde{H} = 2$ or 3, with the most precise evidence, $\tilde{H} = 0$. Therefore, a moderate level of imprecision introduced by the Hamming neighbourhood approach can significantly improve the system's accuracy, depending on the degree of inaccuracy of the evidence. These results are surprisingly consistent with the agent-based simulation results in Sect. 5 and particularly the heat maps in Fig. 9a, where, in the presence of error, optimal Hamming thresholds are also between $\tilde{H} = 2$ and $\tilde{H} = 3$ across varying evidence and fusion rates.

In Sect. 5, Fig. 10 suggests that for the most imprecise evidence, the speed of the social learning model may be significantly slower than models with lower levels of evidential imprecision. We therefore show Fig. 14 for the average accuracy and log cardinality against time. In Fig. 14g, h we see that for the most imprecise evidence, $\tilde{H} = 7$, the system fails to reach consensus within 4000 seconds. For $\tilde{H} = 6$, as depicted in Fig. 14e, f, all 10 runs achieve consensus within 4000 seconds at an error rate of $\epsilon = 0.4$, whereas the system does not reach consensus at the lower error rate of $\epsilon = 0.3$. In other word, in this case higher error rates result in higher speed of social learning. This is consistent with the trend shown in Fig. 10. At lower levels of evidential imprecision, the duration required for consensus is not largely affected by variations in error rates. Specifically, the system reaches consensus in 1200 seconds for $\tilde{H} = 0$ and in 800 seconds for $\tilde{H} = 3$, for both error rates investigated. Notably, the imprecise evidence model with $\tilde{H} = 3$, not only converges more rapidly than the precise evidence model with $\tilde{H} = 0$, but also attains superior learning accuracy. In other words, in scenarios with low evidential imprecision, consensus time cost remains relatively stable across error rates, and intriguingly, the moderate imprecision model with $\tilde{H} = 3$ outperforms the most precise model with $\tilde{H} = 0$ in terms of both convergence speed and learning accuracy.

8 Conclusions and future work

In this paper we have introduced a set-based model of belief for social learning. This has enabled us to clearly distinguish between imprecision and inaccuracy (error). Specifically, in this setting imprecision can be quantified by the cardinality of the evidence sets, while error refers to the difference between the evidence gathered and the true state of the world s^*, and can be modelled by a probabilistic parameter ϵ (i.e. the error rate). We have then used this model to show that social learning can be robust to imprecise evidence and furthermore that certain types of imprecision can result in improved performance.

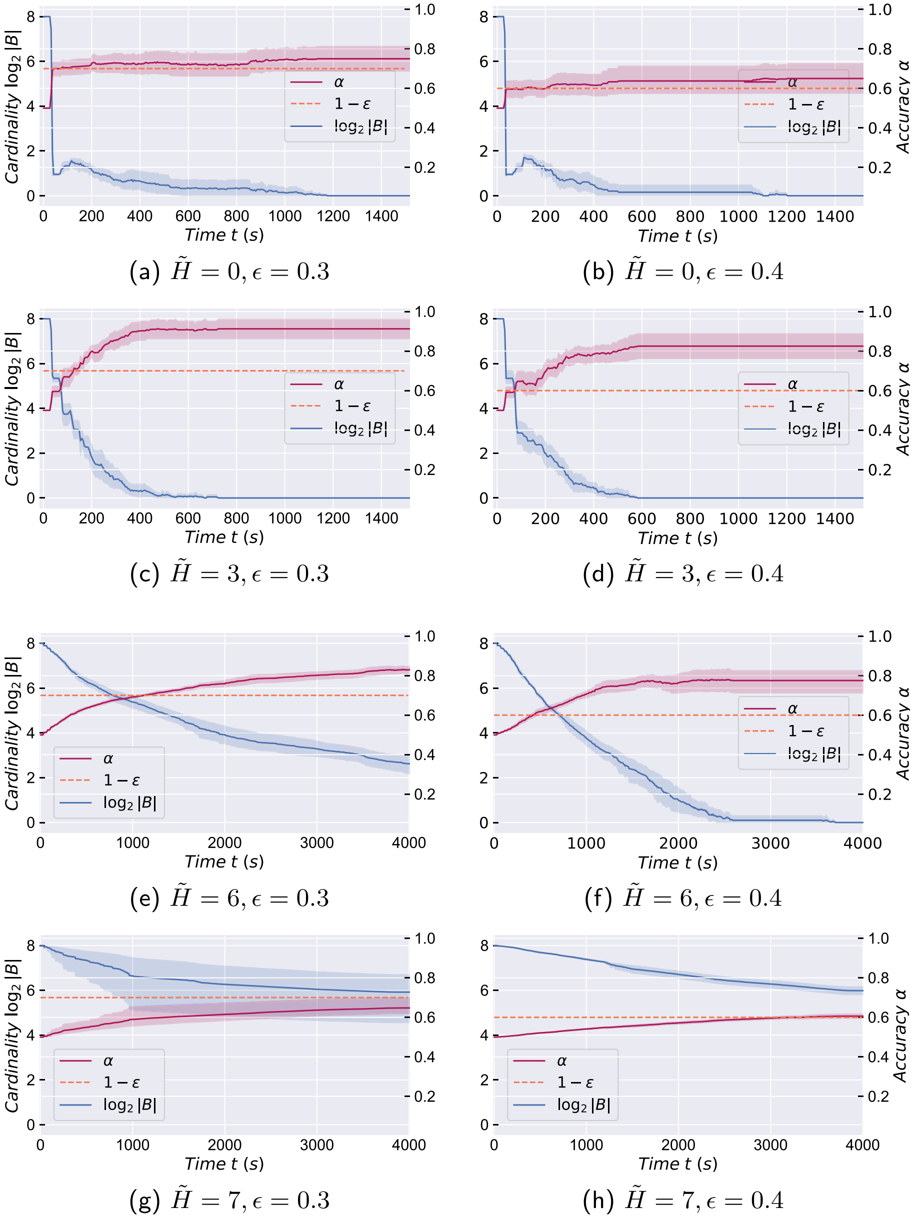

(a) $\tilde{H} = 0, \epsilon = 0.3$ (b) $\tilde{H} = 0, \epsilon = 0.4$

(c) $\tilde{H} = 3, \epsilon = 0.3$ (d) $\tilde{H} = 3, \epsilon = 0.4$

(e) $\tilde{H} = 6, \epsilon = 0.3$ (f) $\tilde{H} = 6, \epsilon = 0.4$

(g) $\tilde{H} = 7, \epsilon = 0.3$ (h) $\tilde{H} = 7, \epsilon = 0.4$

Fig. 14 Average log cardinality $log_2|B|$ and average accuracy α plotted against time for various Hamming distance thresholds with $k = 5$ e-pucks, $n = 8$ locations, and error rate $\epsilon \in \{0.3, 0.4\}$)

We have presented simulation results suggesting that social learning is robust to different levels of imprecise evidence arising naturally as a result of evidence gathering only providing partial information. In this context, we have also investigated a multi-robot location classification problem with a small population size in which the imprecision of evidence gathered by the robots may vary because of hardware or environmental constraints. For this scenario the average accuracy is robust to different levels of imprecise evidence. We have also found that increasing the number of robots can decrease the variation of system's average accuracy at different levels of imprecise evidence. The results of our robot simulation experiments show that our approach has strong potential to be applied to location classification tasks conducted by multi-robot systems.

For a certain type of imprecise evidence, where agents obtain an estimate of the true state of the world and then take a neighbourhood of that estimate, there can actually be benefits of imprecise evidence. The results showed that the overall best accuracy across various evidence and fusion rate combinations and fastest convergence speed can be obtained with an intermediate level of precision with the presence of error. Furthermore, certain levels of imprecision can enable more robustness to variations in fusion and evidence rate than can be obtained from a precise evidence model. Indeed in high error scenarios there are even some levels of tolerance at which the most imprecise evidential model is also the most robust. Evidence neighbourhoods of this kind were then applied to the multi-robot location classification problems and results once again suggest that best performance is achieved when there is a moderate level of evidential imprecision.

In light of previous work on imprecise fusion (Liu et al., 2021), as future work we intend to investigate whether there are advantages in combining imprecision in both the fusion and evidential updating processes. Another avenue of future research will be to add communication constraints to the model, such as network connectivity or physical distance range. There are recent studies showing that limited connectivity can improve the performance of social learning (Crosscombe & Lawry, 2021) and that constrained communication of multi-agent systems can be more robust to the environment changes (Talamali et al., 2021). Additionally, it's worth noting that different types of network connectivity may have varied performance outcomes. For instance, scale-free networks have been shown to enable faster information propagation in the nest, leading to quicker collective responses compared to proximity networks (Rausch et al., 2020b).

Author Contributions This research paper was a collaborative effort between the three authors. Zixuan Liu was responsible for carrying out the research, running simulation experiments, collecting and analysing the data, preparing all the figures, and drafting the manuscript. Jonathan Lawry and Michael Crosscombe provided critical guidance and supervision throughout the research process, offering insights and advice on study design, data analysis, and interpretation of results. Specifically, they contributed to the development of the research ideas, helped to refine the study design, provided valuable feedback on data analysis and interpretation, and assisted in revising the manuscript.

Funding Zixuan Liu is supported by a University of Bristol and China Scholarship Council joint studentship. Michael Crosscombe is a JSPS International Research Fellow supported by a Grant-in-Aid for JSPS Fellows (23KF0108).

Availability of data and materials Raw data were generated from simulation experiments as described in the paper. The data that support the findings of this study are available from the corresponding author, Zixuan Liu, upon reasonable request and can be reproduced based on the model's description provided herein.

Declarations

Conflict of interest The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Ethical approval Not applicable, no human and/ or animal studies conducted.

References

Bartashevich, P., & Mostaghim, S. (2021). Multi-featured collective perception with evidence theory: Tackling spatial correlations. *Swarm Intelligence, 15*(1–2), 83–110. https://doi.org/10.1007/s11721-021-00192-8

Ben-Haim, Y. (2006). *Info-gap decision theory: Decisions under severe uncertainty*. Elsevier.

Braitenberg, V. (1986). *Vehicles: Experiments in synthetic psychology*. MIT Press.

Brambilla, M., Ferrante, E., Birattari, M., et al. (2013). Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence, 7*, 1–41. https://doi.org/10.1007/s11721-012-0075-2

Cholvy, L. (2018). Opinion diffusion and influence: A logical approach. *International Journal of Approximate Reasoning, 93*, 24–39. https://doi.org/10.1016/j.ijar.2017.10.014

Couvillon, M. J., Phillipps, H. L., Schürch, R., et al. (2012). Working against gravity: Horizontal honeybee waggle runs have greater angular scatter than vertical waggle runs. *Biology Letters, 8*(4), 540–543. https://doi.org/10.1098/rsbl.2012.0182

Crosscombe, M., & Lawry, J. (2021). The impact of network connectivity on collective learning. In *International symposium distributed autonomous robotic systems*. Springer (pp. 82–94). https://doi.org/10.1007/978-3-030-92790-5_7

Crosscombe, M., Lawry, J., & Hauert, S., et al. (2017). Robust distributed decision-making in robot swarms: Exploiting a third truth state. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE (pp. 4326–4332). https://doi.org/10.1109/IROS.2017.8206297

Crosscombe, M., Lawry, J., & Bartashevich, P. (2019). Evidence propagation and consensus formation in noisy environments. In *International conference on scalable uncertainty management*. Springer (pp. 310–323). https://doi.org/10.1007/978-3-030-35514-2_23.

De Marco, R. J., Gurevitz, J. M., & Menzel, R. (2008). Variability in the encoding of spatial information by dancing bees. *Journal of Experimental Biology, 211*(10), 1635–1644. https://doi.org/10.1242/jeb.013425

Douven, I. (2019). Optimizing group learning: An evolutionary computing approach. *Artificial Intelligence, 275*, 235–251. https://doi.org/10.1016/j.artint.2019.06.002

Douven, I., & Kelp, C. (2011). Truth approximation, social epistemology, and opinion dynamics. *Erkenntnis, 75*(2), 271. https://doi.org/10.1007/s10670-011-9295-x

Dubois, D., & Prade, H. (1988). Representation and combination of uncertainty with belief functions and possibility measures. *Computational Intelligence, 4*(3), 244–264. https://doi.org/10.1111/j.1467-8640.1988.tb00279.x

Dubois, D., Liu, W., Ma, J., et al. (2016). The basic principles of uncertain information fusion. An organised review of merging rules in different representation frameworks. *Information Fusion, 32*, 12–39. https://doi.org/10.1016/j.inffus.2016.02.006

Dussutour, A., Beekman, M., Nicolis, S. C., et al. (2009). Noise improves collective decision-making by ants in dynamic environments. *Proceedings of the Royal Society B: Biological Sciences, 276*(1677), 4353–4361. https://doi.org/10.1098/rspb.2009.1235

Hegselmann, R., Krause, U., et al. (2006). Truth and cognitive division of labor: First steps towards a computer aided social epistemology. *Journal of Artificial Societies and Social Simulation, 9*(3), 10.

Heyes, C. M. (1994). Social learning in animals: Categories and mechanisms. *Biological Reviews, 69*(2), 207–231. https://doi.org/10.1111/j.1469-185X.1994.tb01506.x

Hintikka, J. (1962). *Knowledge and belief. An introduction to the logic of the two notions*. Ithaca: Cornell University Press.

Khaluf, Y. (2022). Robot swarms decide under perception errors in best-of-n problems. *Applied Sciences, 12*(6), 2975. https://doi.org/10.3390/app12062975

Lawry, J., & Lee, C. (2020). Probability pooling for dependent agents in collective learning. *Artificial Intelligence, 288*(103), 371. https://doi.org/10.1016/j.artint.2020.103371

Lawry, J., Crosscombe, M., & Harvey, D. (2019). Epistemic sets applied to best-of-n problems. In *European conference on symbolic and quantitative approaches with uncertainty*. Springer (pp. 301–312). https://doi.org/10.1007/978-3-030-29765-7_25.

Lee, C., Lawry, J., & Winfield, A. F. (2021). Negative updating applied to the best-of-n problem with noisy qualities. *Swarm Intelligence*. https://doi.org/10.1007/s11721-021-00188-4

Liu, Z., Crosscombe, M., & Lawry, J. (2021). Imprecise fusion operators for collective learning. In *ALIFE 2021: The 2021 conference on artificial life*. MIT Press. https://doi.org/10.1162/isal_a_00407.

Meyer, B., Ansorge, C., & Nakagaki, T. (2017). The role of noise in self-organized decision making by the true slime mold *Physarum polycephalum. PLoS ONE 12*(3):e0172,933. https://doi.org/10.1371/journal.pone.0172933.

Mondada, F., Bonani, M., & Raemy, X., et al. (2009). The e-puck, a robot designed for education in engineering. In *Proceedings of the 9th conference on autonomous robot systems and competitions, vol 1. IPCB*: Instituto Politécnico de Castelo Branco (pp. 59–65).

Okada, R., Ikeno, H., Kimura, T., et al. (2014). Error in the honeybee waggle dance improves foraging flexibility. *Scientific Reports, 4*(1), 1–9. https://doi.org/10.1038/srep04175

Osswald, C., & 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*. IEEE (pp. 1–7). https://doi.org/10.1109/ICIF.2006.301631.

Parker, C., & 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. https://doi.org/10.1109/TMECH.2009.2014370

Preece, K., & Beekman, M. (2014). Honeybee waggle dance error: Adaption or constraint? Unravelling the complex dance language of honeybees. *Animal Behaviour, 94*, 19–26. https://doi.org/10.1016/j.anbehav.2014.05.016

Rausch, I., Nauta, J., & Simoens, P., et al. (2020a). Modeling the influence of social feedback on altruism using multi-agent systems. In *The 2020 conference on artificial life*. MIT Press (pp. 727–735). https://doi.org/10.1162/isal_a_00256.

Rausch, I., Simoens, P., & Khaluf, Y. (2020). Adaptive foraging in dynamic environments using scale-free interaction networks. *Frontiers in Robotics and AI, 7*, 86. https://doi.org/10.3389/frobt.2020.00086

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* (pp .924–931). https://doi.org/10.1007/978-3-540-44792-4_17.

Schwind, N., Inoue, K., Bourgne, G., et al. (2015). Belief revision games. *Proceedings of the AAAI Conference on Artificial Intelligence*. https://doi.org/10.1609/aaai.v29i1.9415

Talamali, M. S., Marshall, J. A., & Bose, T., et al. (2019). Improving collective decision accuracy via time-varying cross-inhibition. In *2019 International conference on robotics and automation (ICRA)*. IEEE (pp. 9652–9659). https://doi.org/10.1109/ICRA.2019.8794284.

Talamali, M. S., Saha, A., Marshall, J. A. R., et al. (2021). When less is more: Robot swarms adapt better to changes with constrained communication. *Science Robotics, 6*(56), 1416. https://doi.org/10.1126/scirobotics.abf1416

Tanner, D. A., & Visscher, K. (2006). Do honey bees tune error in their dances in nectar-foraging and house-hunting? *Behavioral Ecology and Sociobiology, 59*(4), 571–576. https://doi.org/10.1007/s00265-005-0082-z

Towne, W. F., & Gould, J. L. (1988). The spatial precision of the honey bees' dance communication. *Journal of Insect Behavior, 1*(2), 129–155. https://doi.org/10.1007/BF01052234

Valentini, G., Hamann, H., & Dorigo, M., et al. (2014). Self-organized collective decision making: The weighted voter model. In *AAMAS* (pp. 45–52). https://doi.org/10.5555/2615731.2615742.

Valentini, G., Ferrante, E., Hamann, H., et al. (2016). Collective decision with 100 Kilobots: Speed versus accuracy in binary discrimination problems. *Autonomous Agents and Multi-Agent Systems, 30*(3), 553–580. https://doi.org/10.1007/s10458-015-9323-3

Valentini, G., Ferrante, E., & 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. https://doi.org/10.3389/frobt.2017.00009

Vardi, M. (1989). On the complexity of epistemic reasoning. In *Proceedings. Fourth annual symposium on logic in computer science* (pp. 243–244). https://doi.org/10.1109/LICS.1989.39179.

Weidenmüller, A., & Seeley, T. D. (1999). Imprecision in waggle dances of the honeybee (Apis mellifera) for nearby food sources: Error or adaptation? *Behavioral Ecology and Sociobiology, 46*(3), 190–199. https://doi.org/10.1007/s002650050609

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

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.