



Evidential uncertainty sampling strategies for active learning

Arthur Hoarau¹ · Vincent Lemaire² · Yolande Le Gall¹ · Jean-Christophe Dubois¹ · Arnaud Martin¹

Received: 22 November 2023 / Revised: 9 April 2024 / Accepted: 4 May 2024 /

Published online: 27 June 2024

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2024

Abstract

Recent studies in active learning, particularly in uncertainty sampling, have focused on the decomposition of model uncertainty into reducible and irreducible uncertainties. In this paper, the aim is to simplify the computational process while eliminating the dependence on observations. Crucially, the inherent uncertainty in the labels is considered, i.e. the uncertainty of the oracles. Two strategies are proposed, sampling by Klir uncertainty, which tackles the exploration–exploitation dilemma, and sampling by evidential epistemic uncertainty, which extends the concept of reducible uncertainty within the evidential framework, both using the theory of belief functions. Experimental results in active learning demonstrate that our proposed method can outperform uncertainty sampling.

Keywords Active Learning · Uncertainty sampling · Belief Functions

1 Introduction

Active Learning (AL) For reasons of efficiency, cost or energy reduction in machine learning or deep learning, one of the important issues is related to the amount of data and in some cases, to the amount of labeled data. Active learning (Settles, 2009) is a part of

Editor: Myra Spiliopoulou.

✉ Arthur Hoarau
arthur.hoarau@univ-rennes.fr

Vincent Lemaire
vincent.lemaire@orange.com

Yolande Le Gall
yolande.le-gall@univ-rennes.fr

Jean-Christophe Dubois
jean-christophe.dubois@univ-rennes.fr

Arnaud Martin
arnaud.martin@univ-rennes.fr

¹ CNRS, IRISA, DRUID, University of Rennes, 35000 Rennes, France

² Orange Innovation, 22300 Lannion, France

machine learning in which the learner can choose which observation to label in order to work with only a fraction of the labeled dataset to reduce the labeling cost. While primarily used for cost reduction (Hacohen et al., 2022), active learning finds application in various domains like anomaly detection, as seen in (Abe et al., 2006) and (Martens et al., 2023). The essence of active learning lies in empowering the learner to strategically label selected observations. Among all the proposed strategies in the literature (Settles, 2009; Aggarwal et al., 2014), one of the most recognized is uncertainty sampling (Lewis and Gale, 1994; Nguyen et al., 2022).

Uncertainty Quantification (UQ) finds applications across various fields, including medical image analysis, as discussed in a review by (Huang et al., 2023), and in-depth exploration of deep learning applications and techniques (Abdar et al., 2021), such as recent evidential deep learning (Sensoy et al., 2018). Concerning frameworks for uncertainty, numerous methods exist for quantifying uncertainty, with many applied to credal sets as reviewed by (Hüllermeier et al., 2022), or evidential entropies (Deng, 2020). Despite being described several years ago (Hora, 1996), recent literature (Hüllermeier and Waegeman, 2021; Kendall and Gal, 2017; Senge et al., 2014; Charpentier et al., 2020) distinguishes two main types of uncertainty: epistemic and aleatoric. Aleatoric uncertainty arises from the stochastic property of the event and is therefore not reducible, whereas epistemic uncertainty is related to a lack of knowledge and can be reduced. Most proposed calculations hinge not only on the model predictions but also on parameter estimations derived directly from the observations themselves.

(AL) \cup (UQ)—What is the issue? In uncertainty sampling, the learner selects the instances for which it is most uncertain. Until recently, the literature has mostly proposed measures to quantify this uncertainty, such as entropy, in a probabilistic form. But this kind of uncertainty cannot exploit and capture the difference between a label given by someone who has hesitated for a long time and a label given by someone who has no doubt, and therefore uncertainty that may already exist in the labels. In this paper, we propose to use evidential reasoning within the context of active learning. Moreover we propose eliminating direct dependence on the observations and advocating for solely utilizing the model output to achieve a comparable decomposition (*i.e.* epistemic-aleatoric) of uncertainty. This also tackles the exploration–exploitation issue in active learning, with the possibility of choosing one or the other, or even a compromise as suggested by (Bondu et al., 2010). This paper will use the theory of belief functions that generalises probabilities and will use it in the context of active learning.

Note on exploration–exploitation dilemma In Fig. 1, we present a visual depiction of the exploration–exploitation dilemma. In a 2D classification task, the left panel portrays a scenario where all observations are labeled, while the right panel reflects the outcome after iterative labeling rounds, resulting in sparse observations. Should the sampling strategy heavily favor exploitation, as denoted by the blue line representing the classifier, it is evident that the classifier will neglect further investigation into the top left corner. Consequently, upon encountering the red examples in this unexplored territory later in the process, the classifier’s performance undergoes a significant decline, necessitating extensive parameter adjustments. Conversely, a purely exploratory strategy prolongs the duration required to unveil critical patterns, including the “two red patterns” mentioned. Thus, an optimal strategy must delicately balance exploration and exploitation to navigate this trade-off effectively.

Contributions of this paper The primary goal of this paper is to consider the uncertainty inherent in the labels (introduced by the entities labeling observations, referred to as “Oracles” in active learning), and to address the exploration–exploitation dilemma, during

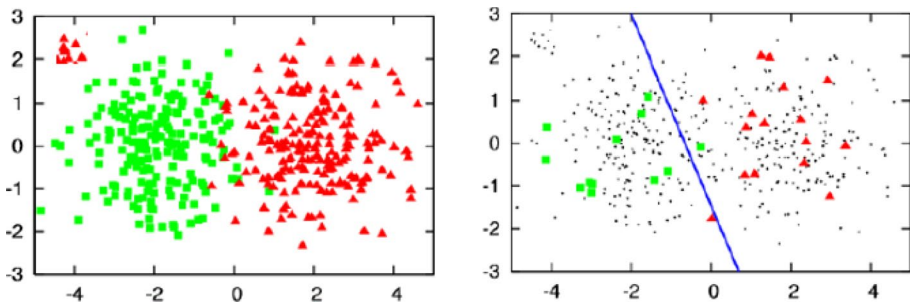


Fig. 1 Illustrating the exploration–exploitation dilemma in active learning: complete dataset versus active learning iterations

sampling. To this end we propose two uncertainty sampling strategies capable of representing a decomposition of the model uncertainties with regard to the uncertainty already present in the labels: (i) a first strategy which is based upon two different uncertainties, discord—how self-conflicting the information is—and non-specificity—how imprecise the information is—in the model output; and (ii) a second strategy which extends the epistemic uncertainty to the evidential framework and to several classes, thus simplifying the computation. To succeed in this challenge we use evidential models able to handle such uncertain labels, such as (Deceux, 1995; Elouedi et al., 2001; Denoeux and Bjanger, 2000; Yuan et al., 2020). By doing this, one can effectively distinguish and account for the difference between a label provided by someone who has hesitated extensively and a label given by someone who has no doubts. As a result, we can identify and quantify the uncertainty inherent in the labels themselves.

Organization The paper is organized as follows: Sect. 2 introduces some important notions of imperfect labeling and the modeling of these richer labels using the theory of belief functions. The conventional uncertainty sampling approach is also recalled and Sect. 3 describes the separation between aleatoric and epistemic uncertainties. Section 4 introduces the two new proposed strategies and Sect. 5 presents the experiments,¹ first on a real world dataset with rich labels and then in active learning to highlight the relevance and efficacy of the proposed method. Section 6 discusses the encountered limits and Sect. 7 concludes the article.

2 Preliminaries

In this section, we provide foundational knowledge essential for understanding the rest of the paper, beginning with rich labels, which are characterized by the theory of belief functions, and concluding with the classical method of uncertainty sampling.

¹ For details on experiments conducted in theoretical sections, visit: <https://anonymous.4open.science/r/evidential-uncertainty-sampling-D453>.

2.1 Imperfect labeling

Most of the datasets used for classification consider hard labels, with a binary membership where the observation is either a member of the class or not. In this paper, we refer as rich labels the elements of response provided by a source that may include several degrees of imprecision (i.e. “*This might be a cat*”, “*I don’t know*” or “*I am hesitating between dog and cat, with a slight preference for cat*”). Such datasets, offering uncertainty already present in the labels, exist (Thierry et al., 2022) but are not numerous. These labels are called rich in this paper since they provide more information than hard labels and can be modeled using the theory of belief functions.

2.2 Theory of belief functions

The theory of belief functions introduced by (Dempster, 1967) and (Shafer, 1976), is used in this study to model uncertainty and imprecision for labeling and prediction. Let $\Omega = \{\omega_1, \dots, \omega_M\}$ be the frame of discernment for M exclusive and exhaustive hypotheses. In supervised learning, this refers to the labels (i.e., classes), or the output space. It is assumed that only one element of Ω is true (closed-world assumption (Smets and Kennes, 1994)). The power set 2^Ω is the set of all subsets of Ω . A mass function assigns the belief that a source may have about the elements of the power set of Ω , such that the sum of all masses is equal to 1.

$$m : 2^\Omega \rightarrow [0, 1], \sum_{A \in 2^\Omega} m(A) = 1. \quad (1)$$

Each subset $A \in 2^\Omega$ such as $m(A) > 0$ is called a *focal element* of m . The uncertainty is therefore represented by a mass $m(A) < 1$ on a focal element A and the imprecision is represented by a non-null mass $m(A) > 0$ on a focal element A such that $|A| > 1$.

A mass function m is called *categorical mass function* when it has only one focal element such that $m(A) = 1$. In the case where A is a set of several elements, the knowledge is certain but imprecise. For $|A| = 1$, the knowledge is certain and precise.

On decision level, the pignistic probability *BetP* of (Smets and Kennes, 1994) helps decision making on singletons:

$$BetP(\omega) = \sum_{A \in 2^\Omega, \omega \in A} \frac{m(A)}{|A|}. \quad (2)$$

It is also possible to combine several mass functions (beliefs from different sources) into a single body of evidence. If the labels and therefore the masses are not independent, a simple average of the mass functions m_j derived from N sources can be defined as follows:

$$m(A) = \frac{1}{N} \sum_{j=1}^N m_j(A), \quad A \in 2^\Omega. \quad (3)$$

There are other possible combinations that are more common than the mean, many of which are listed by (Martin, 2019).

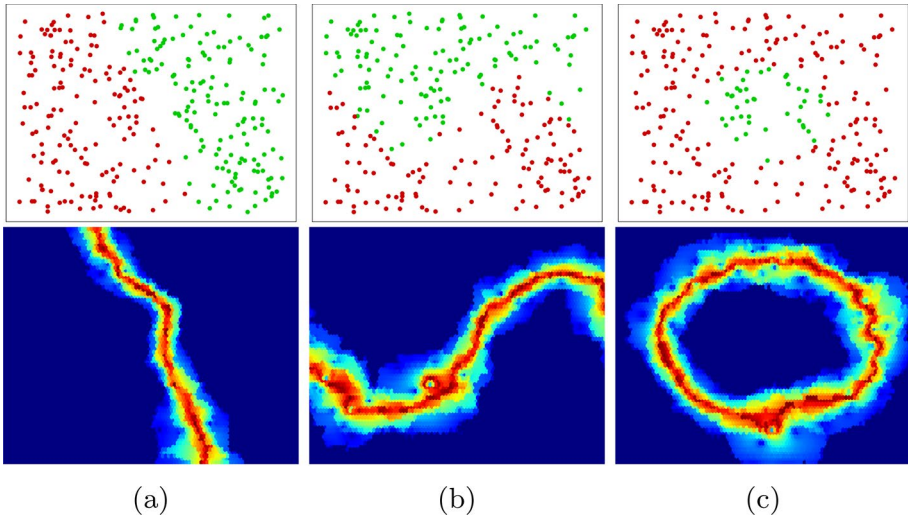


Fig. 2 Visualization of uncertainty areas in two-dimensional datasets

Example 1 Let $\Omega = \{Cat, Dog\}$ be a frame of discernment. An observation labeled “Cat” by a source can be modeled in the framework of belief functions by the mass function m_1 such as: $m_1(\{Cat\}) = 1$ and $m_1(A) = 0, \forall A \in 2^\Omega \setminus \{Cat\}$.

Example 2 An observation labeled “Cat or Dog” by a source can be modeled by the mass function m_2 such as: $m_2(\{Cat, Dog\}) = 1$ and $m_2(A) = 0, \forall A \in 2^\Omega \setminus \{Cat, Dog\}$.

Example 3 The average mass function \bar{m} of m_1 and m_2 is: $\bar{m}(\{Cat\}) = 0.5, \bar{m}(\{Cat, Dog\}) = 0.5$ and $\bar{m}(A) = 0$ for all other subsets A in 2^Ω . Its pignistic probability $BetP$, used for decision making, is given as follows: $BetP(\{Cat\}) = 0.75$ and $BetP(\{Dog\}) = 0.25$.

2.3 Uncertainty sampling

Active learning iteratively builds a training set by selecting the best instances to label. The principle is to label as few observations as possible for a given performance or to achieve the best possible performance within a given budget. Among all the strategies proposed in the literature (Settles, 2009) one of the best known methods is uncertainty sampling (Lewis and Gale, 1994), where the function that defines the instances to be labeled maximizes the uncertainty related to the model prediction as described below.

Let \mathcal{U} be the uncertainty to label a new observation x for a given model and $\Omega = \{\omega_1, \dots, \omega_M\}$ the set of the M possible classes. The uncertainty \mathcal{U} can be calculated in several ways, a classical approach is to use Shannon’s entropy:

$$\mathcal{U}(x) = - \sum_{\omega \in \Omega} p(\omega|x) \log[p(\omega|x)], \tag{4}$$

Fig. 3 Illustration of reducible and irreducible uncertainties in a coin toss experiment (and a Finnish word representation)



Klaava

(a) Aleatoric uncertainty (b) Epistemic uncertainty

with $p(\omega|x)$ the probability for x to belong to the class ω , given by the model. Other common uncertainty criteria include the least confidence measure:

$$\mathcal{U}(x) = 1 - \max_{\omega \in \Omega} [p(\omega|x)]. \quad (5)$$

Measuring the uncertainty of a model to predict the class of some observations can be useful to identify the areas of uncertainty in a space.

Figure 2 represents three two-dimensional datasets, the classes are perfectly separated.

Given the model and one of the uncertainty criteria,² we can compute the uncertainty of any point in space. For each dataset, the areas of uncertainty of the model are represented, with more red for more uncertainty. It is remarkable that these uncertainty areas can be compared to the decision boundaries of the model. Often, the closer the observation is to the decision boundary, the less confident the model is about its prediction.

Uncertainty sampling consists of choosing the observation for which the model is the least certain of its prediction. This is one of the basis of active learning, however, other methods allow to extract more information about this uncertainty which leads to the decomposition into epistemic and aleatoric uncertainties.

3 On the interest and limits of epistemic and aleatoric uncertainties for active learning

In this section, we introduce additional elements to decompose the uncertainty of the model so it can focus, in active learning, on the observations that will make it rapidly gain in performance.

The uncertainty $\mathcal{U}(x)$ can be divided into two types, as outlined by (Hora, 1996): one is reducible, and the other is irreducible. The example provided in Fig. 3 illustrates these distinctions. In Fig. 3a, the outcome of a coin toss is uncertain, and it is impossible to gain further knowledge to predict whether the coin will land heads or tails. This lack of knowledge is referred to as aleatoric uncertainty. On the other hand, in Fig. 3b, a word in Finnish,³ representing either heads or tails is shown. This uncertainty can be resolved by learning the language, making it epistemic uncertainty.

Being able to model these two uncertainties can help to delimit where it is more interesting to provide knowledge and where it is useless. The total uncertainty $\mathcal{U}(x)$ is often

² From now on, the model used is K -NN (K -Nearest Neighbors), with a probabilistic output and on the distance-weighted version available with scikit-learn (Pedregosa et al., 2011), every other parameters are scikit-learn default parameters. The uncertainty used is the least confidence measure given in Eq. (5).

³ In the example, the word “tails” is written in Finnish, the word “heads” is called *Kruuna*.

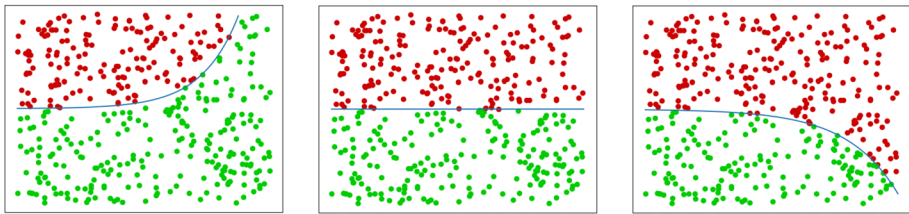
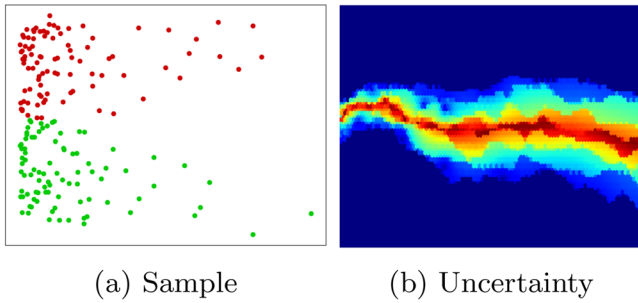


Fig. 4 Visualization of model uncertainty and sample evolution in two-class datasets

represented as the sum of the epistemic uncertainty $\mathcal{U}_e(x)$ and the aleatoric uncertainty $\mathcal{U}_a(x)$: $\mathcal{U}(x) = \mathcal{U}_e(x) + \mathcal{U}_a(x)$.

In a two-class problem where $\Omega = \{0, 1\}$, it is suggested by (Senge et al., 2014) to model this uncertainty, here under the formalism of (Nguyen et al., 2022), by computing the plausibility π of belonging to each of the two classes with the following formula, based on a probabilistic model θ :

$$\begin{aligned} \pi(1|x) &= \sup_{\theta \in \Theta} \min[\pi_{\Theta}(\theta), p_{\theta}(1|x) - p_{\theta}(0|x)], \\ \pi(0|x) &= \sup_{\theta \in \Theta} \min[\pi_{\Theta}(\theta), p_{\theta}(0|x) - p_{\theta}(1|x)], \end{aligned} \tag{6}$$

with $\pi_{\Theta}(\theta)$ depending on the likelihood $L(\theta)$ and the maximum likelihood $L(\hat{\theta})$:

$$\pi_{\Theta}(\theta) = \frac{L(\theta)}{L(\hat{\theta})}. \tag{7}$$

The epistemic uncertainty is then high when the two classes are very plausible⁴ while the aleatoric uncertainty is high when the two classes are implausible:

$$\begin{aligned} \mathcal{U}_e(x) &= \min[\pi(1|x), \pi(0|x)], \\ \mathcal{U}_a(x) &= 1 - \max[\pi(1|x), \pi(0|x)]. \end{aligned} \tag{8}$$

This calculation depends not only on the prediction of the model but also on the observations. To summarize, the fewer observations there are in a region, or the fewer decision

⁴ The notion of plausibility within the theory of belief functions used in the proposed methods differs from the one presented here and will be discussed in greater detail in Sect. 4.

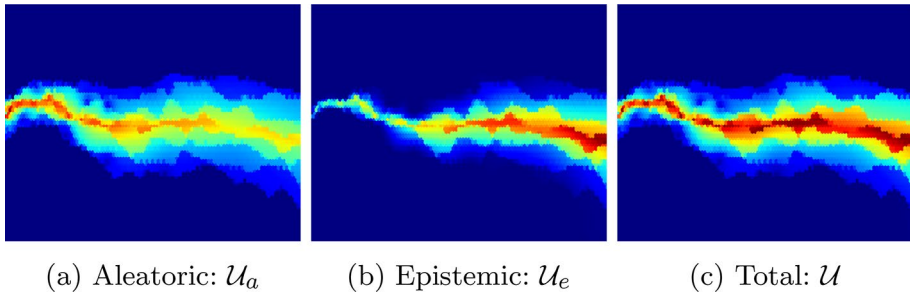


Fig. 5 Representation of aleatoric and epistemic uncertainties in model predictions according to Fig. 4a

elements there are to strongly predict a class, the higher the plausibility of the two classes, and the more reducible (and thus epistemic) the uncertainty is by adding knowledge.

An example is shown in Fig. 4, a two-class dataset is shown in Fig. 4a and the areas of model uncertainty are shown in Fig. 4b according to the uncertainty sampling presented in the previous section.

An horizontal line can be distinguished where the model uncertainty is the highest. However, the sample represented in Fig. 4a, shows that part of the uncertainty can be removed more easily by adding observations. In the same figure, three different datasets show how the sample can evolve by adding observations. Whatever the final distribution is, the uncertainty on the left is not very reducible, while the uncertainty on the right can be modified by adding knowledge.

These two uncertainties can be calculated using Eq. (8), and are shown in Fig. 5. The aleatoric uncertainty, and therefore irreducible, is represented in Fig. 5a and the epistemic uncertainty, reducible, is represented in Fig. 5b. The total uncertainty is then the sum of the two, Fig. 5c. Here the goal is to only use the epistemic uncertainty, to know the areas where the model can learn new knowledge and where it will have more impact.

Using epistemic uncertainty as a sampling strategy is not reductive since it theoretically provides similar areas of uncertainty to those used previously when epistemic and aleatoric uncertainties are indistinguishable. But this statement is based on the sum decomposition of total uncertainty, which has recently been questioned.

This information is valuable for identifying areas of reducible uncertainty. However, it is not compatible with richer labels containing uncertainty. Computing this epistemic uncertainty also relies on observations in addition to the model. Essentially, the model defines its zones of uncertainty and seeks locations with the fewest observations to define the reducible uncertainty. Moreover, the exploration–exploitation problem is not fully addressed. This leads to the next section where two uncertainty sampling strategies for rich labels are proposed, extending to multiple classes.

Fig. 6 Rich label representation: observations on two dimensions with varying ignorance



4 Richer labels and multiple classes

In this section, we propose two uncertainty sampling strategies with a simplified calculation phase, capable of handling richer labels. These strategies are no longer directly dependent on observations but only on the model prediction.⁵ We also propose a natural extension for a number of classes higher than two. The first method uses discord and non-specificity to map uncertainty in order to address the exploration–exploitation problem. The second method extends the epistemic and aleatoric uncertainties to rich labels, also simplifying the computation phase.

From there, a label can be uncertain and imprecise, which means that additional information on ignorance is represented. Figure 6 illustrates how labels are represented in this document: the darker the dot, the less ignorance the label contains (e.g., *I'm sure this is a dog*); the lighter the dot, the more ignorance it contains (e.g., *I have no idea between dog and cat*). It is important to note that labels are no longer “hard”, but modeled by a belief function, which allows such a representation.

4.1 Discord and non-specificity: Klir uncertainty

In the framework of belief functions, discord and non-specificity are tools that allow to model uncertainty, we propose to use (Klir and Wierman, 1998)’s representation for uncertainty sampling, with potential connections to epistemic and aleatoric uncertainty.

4.1.1 Discord

It is here applied to the output of a model capable of making an uncertain and imprecise prediction.⁶ Discord represents the amount of conflicting information in the model’s prediction. It is, for most models, at its maximum closest to the decision boundary and is calculated using the following formula:

$$D(m) = - \sum_{A \subseteq \Omega} m(A) \log_2(\text{Bet}P(A)), \quad (9)$$

with m a mass function, or the output of the model (see Sect. 2.2).

Figure 7 illustrates three different cases where discord varies: from high discordance, where labels around the central point (the observation to label) highly disagree in Fig. 7a, to low discordance, where each label is in agreement in Fig. 7c.

⁵ The uncertainty no longer depends on observations, but the model does.

⁶ From now, the Evidential K -nearest Neighbors model of (Deœux, 1995) is considered.

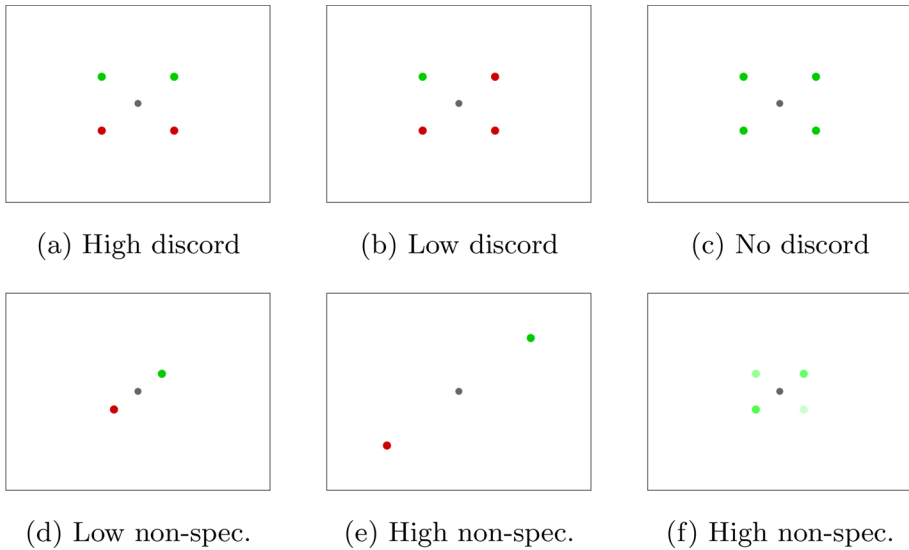


Fig. 7 Quantifying discord and non-specificity in model uncertainty at the central point

4.1.2 Non-specificity

Non-specificity quantifies the degree of imprecision of the model (Dubois and Prade, 1987). This information may be inferred because the model lacks data or because the oracle labeling the instances is itself ignorant. The higher it is, the more imprecise the model's response, it is calculated with:

$$N(m) = \sum_{A \subseteq \Omega} m(A) \log_2(|A|). \quad (10)$$

The same Fig. 7 also represents three different cases of non-specificity, in Fig. 7d the non-specificity is low as there are relevant sources of information next to the observation to be labeled, in Fig. 7e the non-specificity increases the further away the elements are from the observation and in Fig. 7f the non-specificity is also high because the nearby sources of information are themselves ignorant.

4.1.3 Klir uncertainty

This uncertainty is derived from discord and non-specificity, it is used here for uncertainty sampling by adding the two previous formulas:

$$\mathcal{U}_m(x) = N(x) + D(x), \quad (11)$$

with $N(x)$ and $D(x)$ respectively the non-specificity and discord of the model in x . (Klir and Wierman, 1998) propose to use the same weight for discord and non-specificity, but (Denoeux and Bjanger, 2000) introduce a parameter $\lambda \in [0, 1]$ that allows to bring more weight to non-specificity (we propose to use it for more exploration) or to discord (for more exploitation):

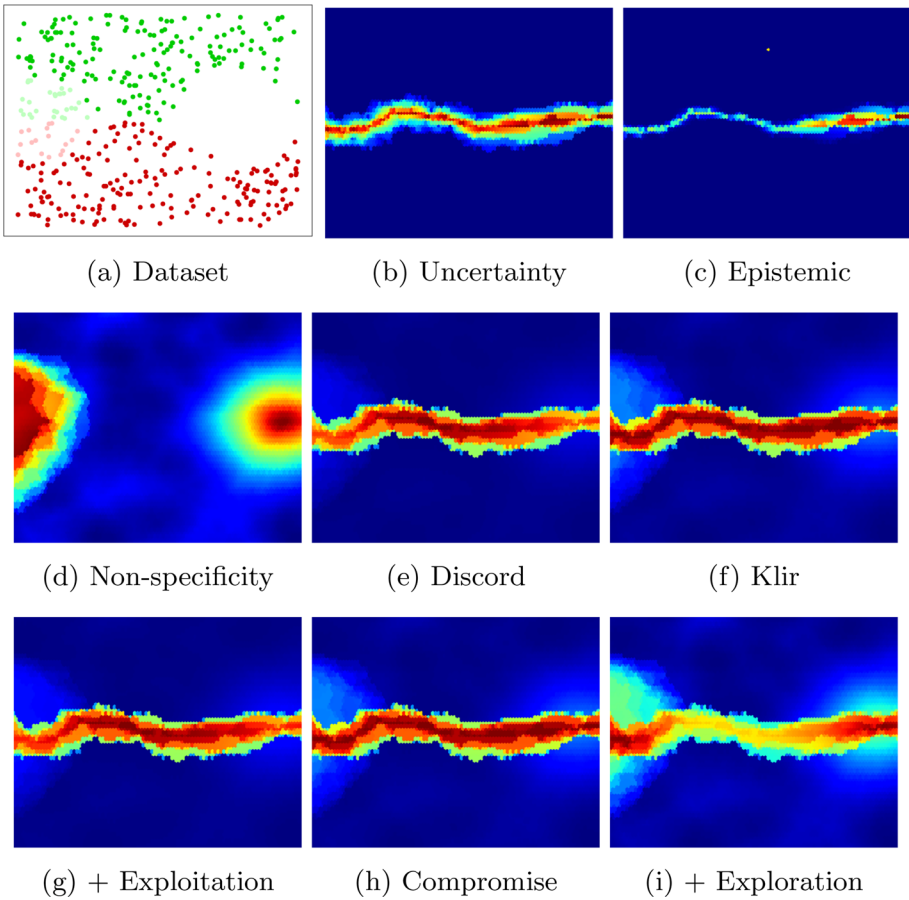


Fig. 8 An imperfectly labeled dataset: exploring uncertainty areas through sampling strategies, epistemic uncertainty, and proposed non-specificity, discord, and total Klir uncertainty, alongside the potential for exploration–exploitation compromise

$$U_m(x) = \lambda N(x) + (1 - \lambda)D(x). \tag{12}$$

Note that this uncertainty is naturally extended to $|\Omega| \geq 2$ classes.

This formula has the advantage of identifying the total uncertainty as well as the reducible one, but also of taking into account the uncertainty already present in the labels and of being adjustable for more exploration or exploitation. Figure 8 depicts a dataset with two areas of uncertainty: on the right, an area with a lack of data, and on the left, an area where labels are more ignorant. The uncertainty sampling, using Shannon’s entropy (4) or the least confidence measure (5) is not able to see either of these two areas, Fig. 8b. The epistemic uncertainty (8) is able to distinguish the uncertainty related to the arrangement of the observations in space (i.e. the uncertainty on the right) but not the uncertainty related to the ignorance of the sources, Fig. 8c.

The proposal of using Klir uncertainty for sampling (discord and non-specificity) allows to represent each of these uncertainties. The second line shows the areas of non-specificity Fig. 8d, of discord Fig. 8e and Klir uncertainty Fig. 8f.

Klir uncertainty can then be used for uncertainty sampling in active learning, it is also possible to vary the result for more exploration or more exploitation by modifying λ . The last line shows the areas of uncertainty for different values of λ , more discord on the left Fig. 8g to more non-specificity on the right Fig. 8i.

We have proposed here to use Klir’s uncertainty in sampling, which allows to represent some uncertainties areas in active learning related to rich labels. The method is no longer dependent on the observations, but only on the prediction of the model and the exploration–exploitation problem is addressed thanks to the λ parameter. Even though discord may recall aleatoric uncertainty (non-reducible) and non-specificity may recall epistemic uncertainty (reducible), these notions are not equivalent. Therefore, in the following section we also propose an extension of epistemic (and aleatoric) uncertainty for rich labels and for several classes.

4.2 Evidential epistemic uncertainty

We propose here to extend the notion of epistemic uncertainty (*c.f.* Sect. 3) to rich labels (*c.f.* Sect. 2.1), by removing the dependence on observations, simplifying the computational phase, and allowing the model to detect new areas of uncertainty.

The epistemic uncertainty can be extended to rich labels by using the notion of plausibility within the framework of belief functions (which differs here from the one presented in Sect. 3). It represents the total evidence that does not support the complementary event for a class ω or more generally for an element $A \in 2^\Omega$. The plausibility Pl defines the belief that could be allocated to A :

$$Pl(A) = \sum_{A \cap B \neq \emptyset} m(B). \tag{13}$$

The plausibility being the consistent evidence, the belief function Bel defines the total evidence directly supporting A :

$$Bel(A) = \sum_{B \subseteq A, B \neq \emptyset} m(B). \tag{14}$$

We have $Pl(A) = 1 - Bel(\bar{A})$. For example, suppose the only reliable evidence is that a picture depicts an animal, and a Dog is an animal. In this scenario, it is entirely plausible that the picture is a Dog (plausibility is 1), yet there is no direct belief that the picture is a Dog (belief is 0). Analogous to Eq. (8) and for two classes $\Omega = \{0, 1\}$ the epistemic uncertainty is maximal when both classes are highly plausible. The proposed evidential epistemic and aleatoric uncertainties are defined as follows:

$$\begin{aligned} \mathcal{U}_e(x) &= \min[Pl(1|x), Pl(0|x)], \\ \mathcal{U}_a(x) &= 1 - \max[Pl(1|x), Pl(0|x)]. \end{aligned} \tag{15}$$

The equation for the aleatoric uncertainty can be rewritten depending on the belief Bel :

$$\mathcal{U}_a(x) = \min[Bel(1|x), Bel(0|x)]. \tag{16}$$

The sum of the epistemic and aleatoric uncertainties is then the total evidential uncertainty: $\mathcal{U}(x) = \mathcal{U}_e(x) + \mathcal{U}_a(x)$. However, when the number of classes exceeds 2 the equation of the epistemic uncertainty cannot be simplified by the minimum plausibility:

$$\begin{aligned}\mathcal{U}_e(x) &\neq \min([Pl(\omega|x)|\omega \in \Omega]), \\ \mathcal{U}_a(x) &\neq 1 - \max([Pl(\omega|x)|\omega \in \Omega]).\end{aligned}\quad (17)$$

It is preferable to first define the uncertainty related to one of the classes ω , rewritten with the belief Bel to avoid having to manipulate $\bar{\omega}$:

$$\begin{aligned}\mathcal{U}_e(\omega|x) &= \min[Pl(\omega|x), Pl(\bar{\omega}|x)] \\ &= \min[Pl(\omega|x), 1 - Bel(\omega|x)].\end{aligned}\quad (18)$$

The evidential extension of the epistemic and aleatoric uncertainties for $|\Omega| \geq 2$ classes is then:

$$\begin{aligned}\mathcal{U}_e(x) &= \sum_{\omega \in \Omega} \min[Pl(\omega|x), 1 - Bel(\omega|x)], \\ \mathcal{U}_a(x) &= \sum_{\omega \in \Omega} \min[Bel(\omega|x), 1 - Pl(\omega|x)].\end{aligned}\quad (19)$$

The example in Fig. 9 shows a dataset of three classes with a zone of imprecision for some labels (between the green and red classes). Probabilistic (4)–(5) and epistemic (8) uncertainties cannot model the imprecision present in the labels, this less complete uncertainty zone is represented in Fig. 9b.

The previous uncertainty resulting from the sum of the discord and the non-specificity is also presented. It manages both exploration, Fig. 9c, and exploitation, Fig. 9d, to give a better representation of the uncertainty, Fig. 9e.

The extension of the epistemic uncertainty, also introduced in this paper, is presented in the same figure. First, the evidential epistemic areas of uncertainties for each of the three classes are presented in Fig. 9f, g, h. Then, the resulting evidential epistemic uncertainty of the model is deducted from Eq. (19) in Fig. 9j along with the evidential aleatoric and total uncertainties.

5 Experiments

In this section, we conduct two types of experiments. The first is more theoretical, applying the two proposed methods to a dataset with real rich labels. The second is more traditional in active learning, comparing one of the methods with uncertainty sampling on several real-world datasets. The exploration–exploitation dilemma is addressed in this second part.

This aim of the first exploratory and non-comparative experiment is to demonstrate how information is mapped by the model. The second experiment offers a more traditional metric-based approach in active learning, allowing for a tangible comparison of methods using classical metrics used in this domain (Kottke et al., 2017)

5.1 Sampling on real world dataset

In this section we use datasets for which we have access to truly imperfectly labeled data with rich labels. This part is exploratory in nature and does not endorse the superiority of any method. Moreover, conventional methods for computing model uncertainty do not take into account the degrees of imprecision of these rich labels and only have access to

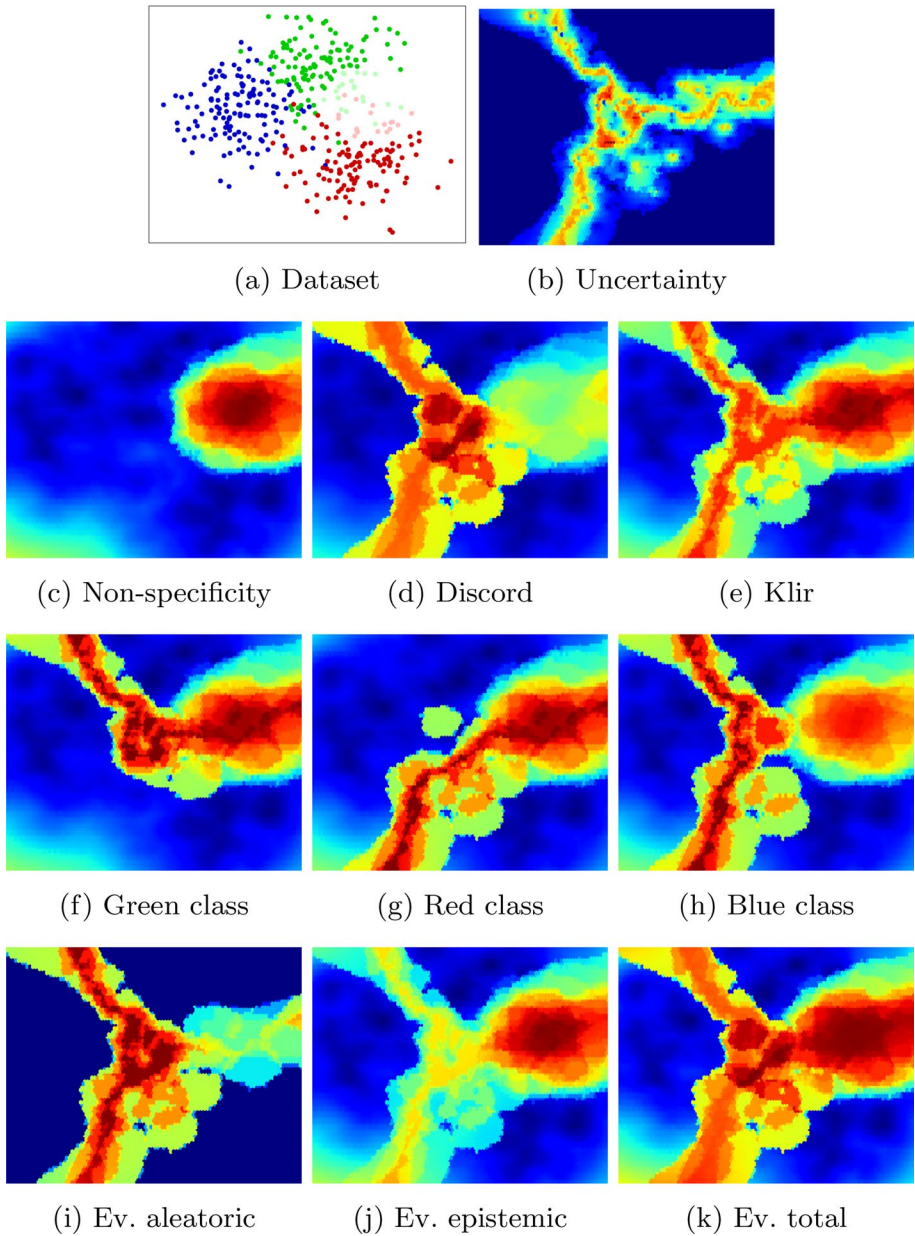


Fig. 9 A three-class dataset: representing label imprecision and uncertainty zones

hard labels. This paper proposes two methods capable of addressing this gap. They are illustrated on Credal Dog-2, a dataset labeled uncertainly and imprecisely by users during crowdsourcing campaigns (Hoarau et al., 2023b). Figure 10 shows the dataset, on the two first components of a Principal Component Analysis to represent this 42-variable dataset in 2D. This is a two-class dataset represented in Fig. 10a with true classes and in Fig. 10b

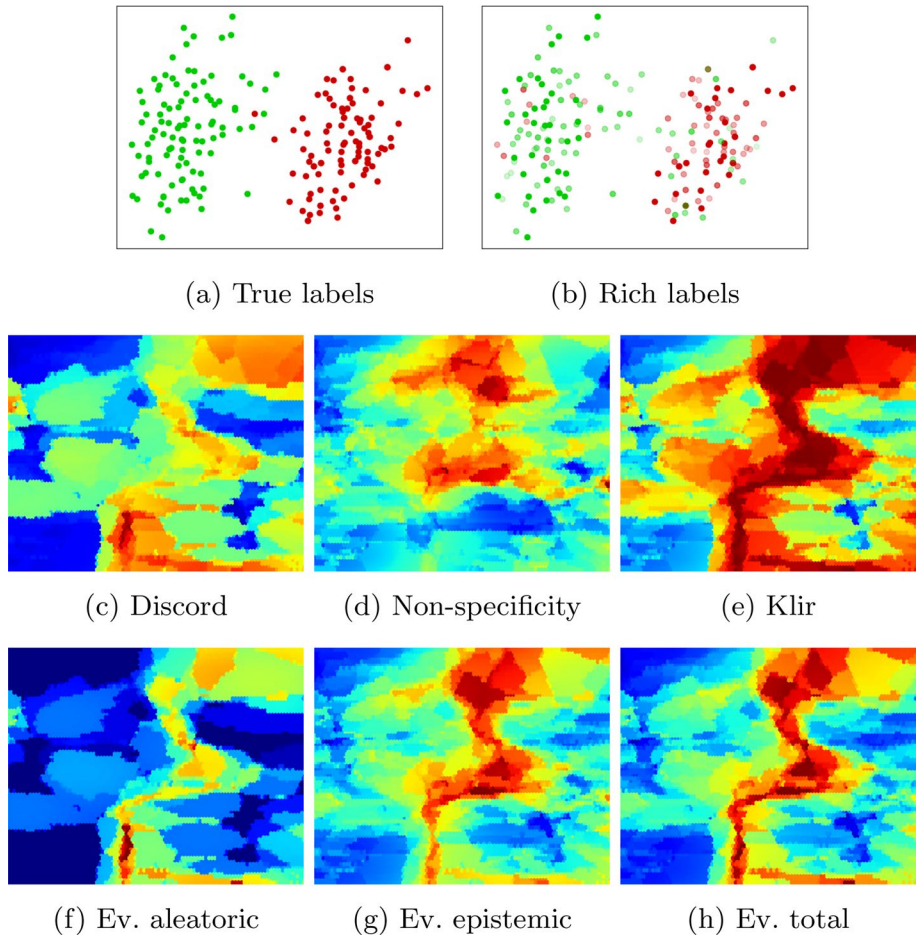


Fig. 10 Ignorance mapping in the Credal Dog-2 dataset: Brittany breed (green) and Beagle breed (red) (Color figure online)

with uncertain and imprecise rich labels given by contributors. Darker dots indicate higher certainty, and *vice versa*.

The first proposed method, sampling by Klir uncertainty, is demonstrated on the dataset with rich labels, Fig. 10b. The non-specificity is presented in Fig. 10d, and can be interpreted as the zones of imprecision of the model, either because it has not had enough access to information (lack of data) or because the users who labeled these instances are themselves ignorant. Discord is also represented in Fig. 10c, these are the areas where the model's prediction is conflicting, meaning it is closest to its decision boundary. The total uncertainty in Fig. 10e is the sum of the two, it is this latter information that is used to sample on the model uncertainty.

The second proposed method extends epistemic uncertainty, which is a reducible uncertainty applied to evidential reasoning. The irreducible aleatoric evidential uncertainty in Fig. 10f is presented along with the reducible epistemic evidential uncertainty in Fig. 10g. The total uncertainty in Fig. 10h is the sum of the reducible and irreducible uncertainties.

Table 1 Datasets description, with number of observations, classes, features and class distribution entropy

Dataset	Observations	Classes	Features	Entropy
Bank	1372	2	4	0.99
Qsar	1055	2	41	0.92
Blod	748	2	4	0.79
Breast Cancer	569	2	30	0.95
Ionosphere	351	2	34	0.94
Heart	303	2	7	1.00
Liver	345	2	6	0.98
Sonar	208	2	60	1.00
Parkinson	195	2	22	0.81
Dog-2	200	2	42	1.00
Seeds	210	3	7	1.00
Iris	150	3	4	1.00
Wine	178	3	13	0.99
Glass	214	6	9	0.83
Ecoli	336	8	7	0.73

For active learning, it is not the total uncertainty, but the epistemic reducible uncertainty that is used. Similarities can be noted between discord and aleatoric uncertainty and between non-specificity and epistemic uncertainty. Additionally, the areas of total uncertainty are also similar.

One advantage of the methods proposed in this paper is their ability to account for the uncertainty already present in the labels (i.e. the uncertainty of the oracles). During labeling by human⁷ (as in the Cedal Dog-2 dataset), a user may hesitate between class 1 and 2. In such cases, it is preferable to model their uncertainty rather than forcing them to provide a wrong label for one of the classes, which would introduce noise into the dataset. The advantage of belief functions is that for multiple classes, the user can respond with various degrees of ignorance. For instance, they may indicate that they are uncertain about the true class but are relatively confident that it is not class 4 or 2. Therefore, being able to represent the imperfections of the oracle can also lead to improved results in machine learning. Unfortunately, extremely expressive visual results are hard to obtain (and interpret) since these datasets are rare, difficult to collect, and very noisy. However, the following experiment, in Sect. 5.2, demonstrates that the proposed methods perform very significantly on this dataset.

5.2 Application to active learning

Sampling by Klir uncertainty was chosen for this series of experiments, and the only parameter λ in the proposed method is at the heart of this study. A λ that tends towards 0 implies more exploitation, whereas a λ that tends towards 1 implies more exploration. For these experiments, we set λ to 0.2, which means that the model will opt for more exploitation

⁷ This representation also applies to labeling performed by a machine.

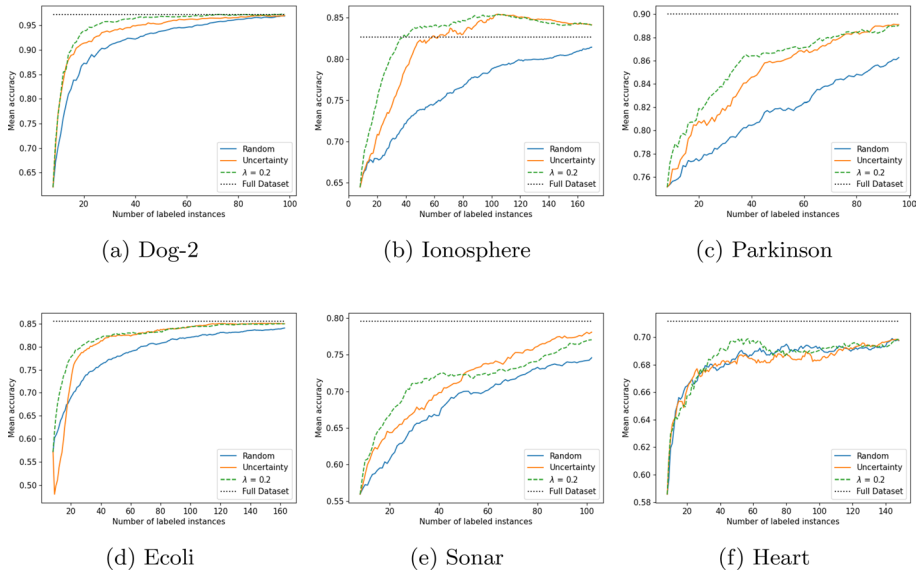


Fig. 11 Mean accuracy versus number of labeled instances for Random Sampling, Uncertainty Sampling, and the proposed method with $\lambda = 0.2$ for 6 datasets

than exploration (this choice is motivated in Sect. 6). Below, the results are compared with random sampling (the baseline) and the very popular uncertainty sampling (5).

The experiments⁸ have been carried out on datasets containing between 2 and 8 classes with a number of observations in different ranges. We used well-known datasets available on the UCI Machine Learning Repository (Dua and Graff, 2017), and very often used in active learning, as well as *Dog-2*, the dataset presented in the previous section. Table 1 describes these datasets, with the number of observations (or instances), the number of classes, the number of features, and the entropy for the distribution of classes.⁹

Since the goal of active learning is to reduce the cost of labeling, one experiment involves evaluating the model's performance as observations are progressively labeled. Experiments are arbitrarily stopped once the dataset has been labeled at 60% (it will be clear from the graphs that there is no point in going any further). The model is the same for each method, the Evidential K -Nearest Neighbors of (Decoux, 1995), with $K = 7$ neighbours (see (Hoarau et al., 2022) for parameter selection). Each experiment is performed 100 times to obtain an estimation of the actual mean accuracy of the model for each dataset. Several criteria are used to compare the results, including accuracy, the area under the accuracy curve (AUAC) and the rank obtained for each dataset. For evaluation, several statistical tests are conducted, including Student's t -test for AUACs, Friedman's test, and the Wilcoxon-Holm method for critical difference diagrams (Demšar, 2006).

⁸ Experiments were conducted according to the following code: <https://anonymous.4open.science/r/evidential-active-learning-B266>.

⁹ An entropy of 1 means that the classes are perfectly equidistributed and an entropy of 0 would indicate the total over-representation of one of the classes.

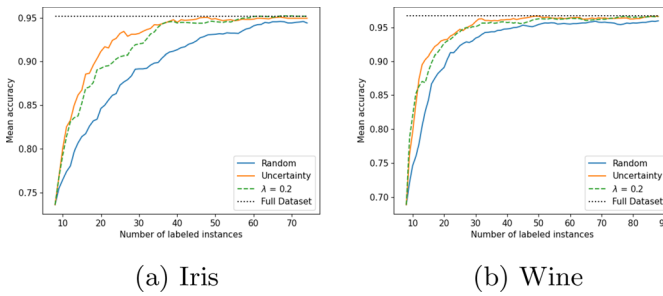


Fig. 12 Mean accuracy versus number of labeled instances for Random Sampling, Uncertainty Sampling, and the proposed method with $\lambda = 0.2$ for Iris and Wine

Figure 11 shows 6 of the 15 datasets where the proposed method offers the most significant performance, the final performance on the full labeled dataset (i.e. if there is no active learning) is represented by the dotted curve and the dashed curve represents the proposed method with $\lambda = 0.2$. In each graph, the superiority of the proposed method over uncertainty sampling is evident, particularly in terms of AUAC. Notably, for the Sonar and Heart datasets, this superiority is only temporary, observed mainly at the beginning of active learning for Sonar and in the middle for Heart. Nevertheless, it will be demonstrated that this domination is not always statistically confirmed, especially for the Parkinson dataset. Assuming an identical labelling cost for each observation, some insights can be reported below.

5.2.1 Some active learning insights on Fig. 11

- On Dog-2: When reaching 99% of the full dataset performance, uncertainty sampling manages to reduce labeling cost by 62% whereas the proposed method manages to reduce the cost by 82%.
- On Ionosphere: Using the proposed method, the labeling costs can be reduced by a factor of 9 with 0% of accuracy loss with respect to the full dataset, whereas with uncertainty sampling, to allow labeling 9 times cheaper, the model would lose 6% of accuracy (for this dataset, the reduction in labeling cost can improve the performance of the model, a phenomenon which sometimes occurs in active learning, represented by the active learning curve which exceeds the full dataset performance horizontal line).
- On Ecoli: It takes 10 steps for uncertainty sampling to reach the performance of the proposed method after only 3 steps.

Even on datasets where the performance of the proposed method is lower, the gap is not always wide. Figure 12 shows two such datasets where the method performs less well. The difference is slightly greater between the proposed method and uncertainty sampling on the Iris dataset.

Since the objective here is to minimize labeling costs,¹⁰ one can set a performance threshold and focus on the actual reduction in costs. By accepting a 2% loss in performance (threshold of 98% performance on the full dataset), conventional uncertainty sampling

¹⁰ Although it can also be to maximize performance given a cost.

Table 2 Mean AUAC for random sampling, uncertainty sampling and the proposed method with $\lambda = 0.2$ on each dataset

Dataset	Method			t-test	
	Random	Uncertainty	$\lambda = 0.2$	statistic	p value
Bank	81.17	81.49	82.23	2.74	0.0067
Qsar	97.69	99.15	99.16	0.35	0.7271
Blod	76.09	76.85	76.85	0.00	0.9965
Breast cancer	93.87	94.96	95.31	1.84	0.0669
Ionosphere	75.77	81.06	82.40	2.33	0.0210
Heart	67.89	67.72	68.08	0.29	0.7741
Liver	58.07	57.37	58.02	0.07	0.9415
Sonar	67.79	70.67	70.94	0.37	0.7089
Parkinson	80.60	83.93	84.85	1.64	0.1034
Dog-2	90.94	93.06	94.10	3.24	0.0014
Seeds	88.70	89.93	89.48	0.84	0.4010
Iris	88.22	91.23	90.60	1.19	0.2373
Wine	91.66	93.55	93.27	0.86	0.3920
Glass	57.33	58.32	59.05	0.87	0.3829
Ecoli	78.59	80.89	81.98	2.19	0.0300

Bold indicates best method

Student’s t-test is also performed to determine the significance of the best method



Fig. 13 Critical difference diagrams for different values of λ on the proposed method 13a and for random sampling, uncertainty sampling and the proposed method 13b

reduces the number of labels in the Dog-2, Ionosphere, and Heart datasets by 76%, 86%, and 43% respectively. In contrast, the proposed method reduces the number of labels in the respective datasets by 88%, 91%, and 83%. However, a performance threshold has been set, which is why the area under the curve is a good indicator, as it captures the reduction in costs for all possible thresholds. Table 2 shows the mean areas under the curve for the three methods studied and for each dataset. A statistical t-test is also performed between the first and second best methods for each value. Random sampling performs best on the Liver dataset, uncertainty sampling performs best on Seeds, Iris and Wine and the proposed method performs best on the other 10 datasets, except Blod where there is a tie.

To statistically find the best model, a critical difference diagram is drawn up. The first diagram in Fig. 13a is a comparison of the proposed method with different values of λ . On all datasets, $\lambda = 0.2$ ranks on average at position 2.13 out of 4 and $\lambda = 0.5$ (which is equivalent to as much exploration as exploitation) ranks on average at position 3.47. If a line connects two methods, this means that despite the better performance of one, the methods are not statistically differentiable. In the example, $\lambda = 0.2$, $\lambda = 0.3$ and $\lambda = 0.4$ are not statistically different. Now, Fig. 13b is obtained by comparing the proposed method with

random sampling and sampling by uncertainty. In average, the proposed method ranks 1.33 out of 3 and the significance of this result is demonstrated by the absence of a line linking the methods. It may also be interesting to note that uncertainty sampling and random sampling are not connected by a line either.

6 Discussion

Calculating epistemic uncertainty (non-evidential) is demanding and not always accessible. It depends on the observations, requiring several phases of computation, including likelihood estimation, maximum likelihood, and optimization.

The two proposed methods offer simplicity, but there is a counterpart: the model must be capable of delivering a mass function to represent uncertainty and imprecision in the output. Such models exist but are not abundant. Among them are Evidential K -Nearest Neighbors (Deceux, 1995), Evidential Decision Trees (Elouedi et al., 2001; Denoeux and Bjanger, 2000), Evidential Random Forests (Hoarau et al., 2023a), and even Evidential Neural Networks (Yuan et al., 2020). The proposed methods are compatible with probabilistic models (since a probability is a special belief function), but they may not capture the full depth of evidence modeling.

In the experiments above, λ was set at 0.2, meaning that the model prioritizes exploitation over exploration. This is a value that gives good results. Our studies to determine when one is more relevant than the other are illustrated in Fig. 13a. The results indicate that several lambda values yield fairly similar outcomes. The value of 0.2 is the one that gives the best performance in general: for the majority of datasets, it is more interesting to do more exploitation, without exceeding a certain limit, otherwise model performance will drop. For future work, it would be interesting to modify the value of λ as the labeling process progresses. This adjustment could lead to a more powerful model capable of dynamically balancing between exploration and exploitation.

7 Conclusion

This paper introduces two new uncertainty sampling strategies and a novel representation method for them. These two methods use Klir uncertainty and an extended evidential epistemic uncertainty. A straightforward calculation on the model output enables the extraction of uncertainties. The objective is to also take into account the uncertainty present in richer labels, which was not possible up to now. The first strategy is based on Klir's uncertainty, combining discord (how self-conflicting the information is) and non-specificity (how imprecise the information is) in the model output. The second strategy extends epistemic (reducible) uncertainty to the evidential framework and to several classes, simplifying the computational phase.

The proposed Klir uncertainty sampling is chosen for its competitiveness in active learning. Its superiority over uncertainty sampling is statistically validated across the datasets examined in the experiments. The novelty of this work lies in representing new information for uncertainty sampling, which also yields significant performance improvements in traditional active learning. The next step is to control exploration and exploitation (represented here as the λ parameter) and to determine, for each dataset, when exploration or

exploitation is more advantageous. The ability of the model to define areas of uncertainty, and to categorize these uncertainties, is then relevant information.

Author contributions Arthur Hoarau, Vincent Lemaire, Jean-Christophe Dubois, Yolande Le Gall and Arnaud Martin contributed to the manuscript equally.

Funding This work is funded by the Brittany region and the Côtes-d’Armor department. The authors also received funding from IRISA, the University of Rennes, DRUID and Orange SA.

Data availability All data used in this study are available publicly online. The datasets were extracted directly in the repositories available with the links in the following section.

Code availability The code for theoretical experiments is available at this link: <https://anonymous.4open.science/r/evidential-uncertainty-sampling-D453>. Link to the code for the experimental part in active learning: <https://anonymous.4open.science/r/evidential-active-learning-B266>.

Declarations

Conflict of interest Arthur Hoarau, Jean-Christophe Dubois, Yolande Le Gall and Arnaud Martin received research support from the University of Rennes, the IRISA laboratory and the DRUID team. Vincent Lemaire received research support from Orange SA.

Ethics approval Not applicable.

Consent to participate All authors have read and approved the final manuscript.

Consent for publication All authors approved the publication.

References

- Abdar, M., Pourpanah, F., Hussain, S., et al. (2021). A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion*, 76, 243–297.
- Abe, N., Zadrozny, B., & Langford, J. (2006). Outlier detection by active learning. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2006, 504–509.
- Aggarwal, C., Kong, X., Gu, Q., et al. (2014). *Active Learning: A Survey*, *Data Classification: Algorithms and Applications*. CRC Press.
- Bondu, A., Lemaire, V. & Boullé, M. (2010). Exploration vs. exploitation in active learning: A Bayesian approach. In *The 2010 international joint conference on neural networks (IJCNN)* (pp. 1–7).
- Charpentier, B., Zügner, D., & Günemann, S., et al. (2020). Posterior network: Uncertainty estimation without OOD samples via density-based pseudo-counts. In H. Larochelle, M. Ranzato, & R. Hadsell (Eds.), *Advances in Neural Information Processing Systems* (Vol. 33, pp. 1356–1367). Curran Associates Inc.
- Dempster, A. P. (1967). Upper and lower probabilities induced by a multivalued mapping. *The Annals of Mathematical Statistics*, 38(2), 325–339.
- Demšar, J. (2006). Statistical comparisons of classifiers over multiple data sets. *The Journal of Machine Learning Research*, 7, 1–30.
- Deng, Y. (2020). Uncertainty measure in evidence theory. *Science China Information Sciences*, 63, 210201.
- Dencœur, T. (1995). A k-nearest neighbor classification rule based on Dempster–Shafer theory. *IEEE Transactions on Systems, Man and Cybernetics*, 219.
- Denoeux, T., & Bjanger, M. (2000). Induction of decision trees from partially classified data using belief functions. *Systems, Man, and Cybernetics*, 4, 2923–2928.
- Dua, D. & Graff, C. (2017). UCI ML Repository. <https://archive.ics.uci.edu/>
- Dubois, D., & Prade, H. (1987). Properties of measures of information in evidence and possibility theories. *Fuzzy Sets and Systems*, 24(2), 161–182.
- Elouedi, Z., Mellouli, K., & Smets, P. (2001). Belief decision trees: Theoretical foundations. *International Journal of Approximate Reasoning*, 28(2), 91–124.
- Hacohen, G., Dekel, A. & Weinshall, D. (2022). Active learning on a budget: Opposite strategies suit high and low budgets. In Chaudhuri, K., Jegelka, S., Song, L., et al. (Eds.), *International conference on machine*

- learning, 2022, Baltimore, Maryland, USA, proceedings of machine learning research (vol. 162, pp. 8175–8195). PMLR.
- Hoarau, A., Martin, A., Dubois, J. C., et al. (2022). Imperfect labels with belief functions for active learning. In *Belief functions: Theory and applications*. Springer.
- Hoarau, A., Martin, A., Dubois, J. C., et al. (2023a). Evidential random forests. *Expert Systems with Applications*, 230.
- Hoarau, A., Thierry, C., Martin, A., et al. (2023b). Datasets with rich labels for machine learning. In *2023 IEEE international conference on fuzzy systems (FUZZ-IEEE)* (pp. 1–6).
- Hora, S. C. (1996). Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management. *Reliability Engineering & System Safety*, 54(2), 217–223. Treatment of Aleatory and Epistemic Uncertainty.
- Huang, L., Ruan, S. & Xing, Y., et al. (2023). A review of uncertainty quantification in medical image analysis: Probabilistic and non-probabilistic methods.
- Hüllermeier, E., & Waegeman, W. (2021). Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110, 457–506.
- Hüllermeier, E., Destercke, S. & Shaker, M.H. (2022). Quantification of credal uncertainty in machine learning: A critical analysis and empirical comparison. In Cussens, J., & Zhang, K. (Eds.), *Proceedings of the thirty-eighth conference on uncertainty in artificial intelligence, proceedings of machine learning research* (vol. 180, pp. 548–557). PMLR.
- Kendall, A. & Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? In *NIPS*.
- Klir, G. J., & Wierman, M. J. (1998). *Uncertainty-based information: Elements of generalized information theory*. Springer.
- Kottke, D., Calma, A., Huseljic, D., et al. (2017). Challenges of reliable, realistic and comparable active learning evaluation. In *Proceedings of the workshop and tutorial on interactive adaptive learning* (pp. 2–14).
- Lewis, D. D., & Gale, W. A. (1994). A sequential algorithm for training text classifiers. In *SIGIR*.
- Martens, T., Perini, L., & Davis, J. (2023). Semi-supervised learning from active noisy soft labels for anomaly detection. *Machine learning and knowledge discovery in databases: Research track: European conference, ECML PKDD 2023, Turin* (pp. 219–236). Springer-Verlag.
- Martin, A. (2019). Conflict management in information fusion with belief functions. In E. Bossé & G. L. Rogova (Eds.), *Information quality in information fusion and decision making. Information fusion and data science* (pp. 79–97). Springer.
- Nguyen, V. L., Shaker, M. H., & Hüllermeier, E. (2022). How to measure uncertainty in uncertainty sampling for active learning. *Machine Learning*, 111, 89–122.
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Senge, R., Bösner, S., Dembczynski, K., et al. (2014). Reliable classification: Learning classifiers that distinguish aleatoric and epistemic uncertainty. *Information Science*, 255, 16–29.
- Sensoy, M., Kaplan, L., & Kandemir, M., et al. (2018). Evidential deep learning to quantify classification uncertainty. In S. Bengio, H. Wallach, & H. Larochelle (Eds.), *Advances in Neural Information Processing Systems*. (Vol. 31). Curran Associates Inc.
- Settles, B. (2009). Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison.
- Shafer, G. (1976). *A mathematical theory of evidence*. Princeton University Press.
- Smets, P. & Kennes, R. (1994). The transferable belief model. *Artificial Intelligence*.
- Thierry, C., Hoarau, A., Martin, A., et al. (2022). Real bird dataset with imprecise and uncertain values. In *7th International conference on belief functions*.
- Yuan, B., Yue, X., Lv, Y., et al. (2020). Evidential deep neural networks for uncertain data classification. In: *Knowledge science, engineering and management (proceedings of KSEM 2020). Lecture notes in computer science*. Springer Verlag.

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

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.