

Knowledge-Driven Active Learning

Original

Knowledge-Driven Active Learning / Ciravegna, G., Precioso, F., Betti, A., Mottin, K., Gori, M.. - 14169:(2023), pp. 38-54.
(European Conference, ECML PKDD 2023 Turin (ITA) September 18–22, 2023) [10.1007/978-3-031-43412-9_3].

Availability:

This version is available at: 11583/2988536 since: 2024-05-12T18:35:36Z

Publisher:

Springer

Published

DOI:10.1007/978-3-031-43412-9_3

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

Springer postprint/Author's Accepted Manuscript

This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use, but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: http://dx.doi.org/10.1007/978-3-031-43412-9_3

(Article begins on next page)

Knowledge-driven Active Learning

Gabriele Ciravegna^{1,2}[0000-0002-6799-1043]✉,
 Frédéric Precioso²[0000-0001-8712-1443], Alessandrio Betti²[0000-0002-9052-8743],
 Kevin Mottin²[0009-0003-9277-0749], and Marco Gori^{2,3}[0000-0002-9052-8743]

¹ Politecnico di Torino, DAUIN, Torino, Italy
 gabriele.ciravegna@polito.it

² Université Côte d’Azur, Inria, CNRS, I3S, Maasai, Nice, France

³ Università di Siena, DIISM, Siena, Italy

Abstract. The deployment of Deep Learning (DL) models is still precluded in those contexts where the amount of supervised data is limited. To answer this issue, active learning strategies aim at minimizing the amount of labelled data required to train a DL model. Most active strategies are based on uncertain sample selection, and even often restricted to samples lying close to the decision boundary. These techniques are theoretically sound, but an understanding of the selected samples based on their content is not straightforward, further driving non-experts to consider DL as a black-box. For the first time, here we propose to take into consideration common domain-knowledge and enable non-expert users to train a model with fewer samples. In our Knowledge-driven Active Learning (KAL) framework, rule-based knowledge is converted into logic constraints and their violation is checked as a natural guide for sample selection. We show that even simple relationships among data and output classes offer a way to spot predictions for which the model need supervision. We empirically show that KAL (i) outperforms many active learning strategies, particularly in those contexts where domain knowledge is rich, (ii) it discovers data distribution lying far from the initial training data, (iii) it ensures domain experts that the provided knowledge is acquired by the model, (iv) it is suitable for regression and object recognition tasks unlike uncertainty-based strategies, and (v) its computational demand is low.

Keywords: Active Learning · Knowledge-aided Learning · Neurosymbolic Learning.

1 Introduction

Deep Learning (DL) methods have achieved impressive results over the past decade in fields ranging from computer vision to text generation [31]. However, most of these contributions relied on overly data-intensive models (e.g. Transformers), trained on huge amounts of data [34]. With the advent of Big Data, sample collection does not represent an issue any more, but, nonetheless, in some contexts the number of *supervised* data is limited, and manual labelling can be

expensive [58]. Therefore, a common situation is the unlabelled pool scenario [36], where many data are available, but only some are annotated. Historically, two strategies have been devised to tackle this situation: semi-supervised learning which exploits the unlabelled data to enrich feature representations [62], and active learning which selects the smallest set of data to annotate to improve the most model performances [49].

The main assumption behind active learning strategies is that there exists a subset of samples that allows to train a model with a similar accuracy as when fed with all training data. Iteratively, the strategy indicates the optimal samples to be annotated from the unlabelled pool. This is generally done by ranking the unlabelled samples w.r.t. a given measure, usually on the model predictions [49,38,54], or on the input data distribution [61,44] and by selecting the samples associated to the highest rankings [42,59]. While being theoretically sound, an understanding of the selected samples based on their content is not straightforward, in particular to non-ML experts. This issue becomes particularly relevant when considering that Deep Neural Networks are already seen as black box models [19,11]. On the contrary, we believe that neural models must be linked to Commonsense knowledge related to a given learning problem. Therefore, in this paper, we propose for the first time to exploit this symbolic knowledge in the selection process of an active learning strategy. This not only lower the amount of supervised data, but it also enables domain experts to train a model leveraging their knowledge. More precisely, we propose to compare the predictions over the unsupervised data with the available knowledge and to exploit the inconsistencies as a criterion for selecting the data to be annotated. Domain knowledge, indeed, can be expressed as First-Order Logic (FOL) clauses and translated into real-valued logic constraints (among other choices) by means of T-Norms [28] to assess its satisfaction [21,14,37].

In the experiments, we show that the proposed Knowledge-driven Active Learning (KAL) strategy (i) performs better (on average) than several standard active learning methods, particularly in those contexts where domain-knowledge is rich. We empirically demonstrate (ii) that this is mainly due to the fact that the proposed strategy allows discovering data distributions lying far from the initial training data, unlike uncertainty-based approaches. Furthermore, we show that (iii) the KAL strategy can be easily employed also in regression and object-detection contexts, where standard uncertainty-based strategies are not-straightforward to apply [25], (iv) the provided knowledge is acquired by the trained model, (iv) KAL can also work on domains where no knowledge is available if combined with a XAI technique, and, finally, (vi) KAL is not computationally expensive unlike many recent methods.

The paper is organized as follows: in Section 2 the proposed method is explained in details, with first an example on inferring the XOR operation and then contextualized in more realistic active learning domains; the aforementioned experimental results on different datasets are reported in Section 3, comparing the proposed technique with several active learning strategies; in Section 4 the related work about active learning and about integrating reasoning with ma-

chine learning is briefly resumed; finally, in Section 5 we conclude the paper by considering possible future work.

2 Knowledge-driven Active Learning

In this paper, we focus on a variety of learning problems, ranging from classification to regression and also object-detection. Therefore, we consider the problem $f: X \rightarrow Y$, where $X \subseteq \mathbb{R}^d$ represents the feature space which may also comprehend non-structured data (e.g., images) and d represents the input dimensionality and Y the output space. More precisely, in classification problems we consider a vector function $f = [f_1, \dots, f_c]$, where each function f_i predicts the probability that x belongs to the i -th class. When considering an object-detection problem, instead, for a given class i and a given image $x \in X$, we consider as class membership probability f_i the maximum score value among all predicted bounding boxes around the objects belonging to the i -th class. Formally, $f_i(x) = \max_{s \in \mathcal{S}^i(x)} s(x)$ where $\mathcal{S}^i(x_j)$ is the set of the confidence scores of the bounding boxes predicting the i -th class for sample x . Finally, in regression problems the learning function f_i represents the predicted value for the i -th class and takes values outside the unit interval, i.e. $f_i(x) \in \mathbb{R}$.

In the Active Learning context, we also define $X_s \subset X$ as the portion of input data already associated to an annotation $y_i \in Y_s \subset Y$ and n the dimensionality of the starting set of labelled data. At each iteration, a set of p samples $X_p \subset X_u \subset X$ is selected by the active learning strategy to be annotated from X_u , the unlabelled data pool, and be added to X_s . This process is repeated for q iterations, after which the training terminates. The maximum budget of annotations b therefore amounts to $b = n + q \cdot p$.

Let us also consider the case in which additional *domain knowledge* is available for the problem at hand, involving relationships between data and classes. By considering the logic predicate \mathbf{f} associated to each function f , First-Order Logic (FOL) becomes the natural way of describing these relationships. For example, $\forall x \in X, \mathbf{x}_1(x) \wedge \mathbf{x}_2(x) \Rightarrow \mathbf{f}(x)$, meaning that when both predicates are true also the output function $f(x)$ needs to be true and where $\mathbf{x}_1(x), \mathbf{x}_2(x)$ respectively represent the logic predicates associated to the first and the second input features. Also, we can consider relations among classes, such as $\forall x \in X, \mathbf{f}_v(\mathbf{x}) \wedge \mathbf{f}_z(\mathbf{x}) \Rightarrow \mathbf{f}_u(\mathbf{x})$, meaning that the intersection between the v -th class and the z -th class is always included in the u -th one. Finally, we can consider predicates defined over open $\mathbf{f}(x) > k$ or closed intervals, $k_1 < \mathbf{f}(x) < k_2$

2.1 Converting Domain-Knowledge into loss functions

Among different approaches that allow to inject domain knowledge into a learning problem (see [20] for a complete review of approaches), in this work we employ the Learning from Constraints framework [21,14] which converts domain knowledge into numerical constraints. Among a variety of other type of constraints (see, e.g., Table 2 in [21]), it studies the process of handling FOL




Fig. 1. A visual example of KAL working principles on the *XOR-like* problem. We depict network predictions with different colour degrees. Also, we depict in orange the samples selected by the active strategy in the current iteration and in blue those selected in previous iterations (or initially randomly annotated). Notice how the proposed method immediately discovers the data distribution not covered by the initial random sampling (right-bottom quadrant).

formulas so that they can be either injected into the learning problem (in semi-supervised learning [35]) or used as a knowledge verification measure (as in [37] and in the proposed method). Going into more details, the FOL formulas representing the domain knowledge are converted into numerical constraints using the Triangular Norms (T-Norms, [28]). These binary functions generalize the conjunction operator \wedge and offer a way to mathematically compute the satisfaction level of a given rule.

Following the previous example, $\mathbf{x}_1(x) \wedge \mathbf{x}_2(x) \Rightarrow \mathbf{f}(x)^4$ is converted into a bilateral constraint $\phi(f(x)) = 1$. By first rewriting the rule as a conjunction of terms $\neg((\mathbf{x}_1 \wedge \mathbf{x}_2) \wedge \neg \mathbf{f})^5$ and by employing the product T-Norm which replaces the \wedge with the product operators and $\neg \mathbf{x}$ with $1 - \mathbf{x}$, the bilateral constraint becomes $1 - (\mathbf{x}_1 \mathbf{x}_2 (1 - \mathbf{f})) = 1$. With $\varphi(f(x)) = 1 - \phi(f(x))$ we indicate the loss function associated to the bilateral constraints, which measures the level of satisfaction of the given constraints and has its minimum value in zero. Again, recalling the previous example, the associated loss function would be $\varphi(f(x)) = \mathbf{x}_1 \mathbf{x}_2 (1 - \mathbf{f})$, which indeed is satisfied when either \mathbf{x}_1 or \mathbf{x}_2 is zero or \mathbf{f} is approximately one. For further detail on how to convert FOL formulas into numerical constraints see Appendix A.1 and [35] which also proposed an automatic computation of the loss function φ associated to a rule.

Based on this assumption, we can detect whether the predictions made by the model on unlabelled data are coherent with the domain knowledge, and we select the data associated to the highest violations as those to be annotated. More precisely, considering the set \mathcal{K} of all available FOL formulas for the given

⁴ Practically, the predicate $\mathbf{x}_i(x)$ is obtained applying a steep logistic function over the i -th input feature: $\mathbf{x}_i = \sigma(x_i) = 1/(1 + e^{-\tau(x_i - h)})$, where τ is a temperature parameter and h represents the midpoint of the logistic function ($h = 0.5$). For predicates expressing inequalities, e.g., $\mathbf{f}(x) > k$ we simply need to set $h = k$.

⁵ For the sake of simplicity, we drop the argument (x) of the logic predicates.

problem, we select the points x^* which violate the most the constraints as follows:

$$KAL : \quad x^* = \arg \max_{x \in X_u} \sum_{\varphi \in \mathcal{K}} \varphi(f(x)) \quad (1)$$

At each iteration, the KAL strategy selects p samples x^* to annotate from the unlabelled pool X_u .

2.2 An intuitive example: the XOR-like problem

A well-known problem in machine learning is the inference of the eXclusive OR (XOR) operation. To show the working principles of the proposed approach, we propose a variant of this experiment, in which a neural network learns the XOR-like operation from a distribution of non-boolean samples. Specifically, we sampled 10^5 points $x \in [0, 1]^2$, and we assigned a label $y(x)$ as following: $y(x) = 1$ if $(x_1 > 0.5 \wedge x_2 \leq 0.5) \vee (x_1 \leq 0.5 \wedge x_2 > 0.5)$ else $y(x) = 0$. Also, we express the XOR operation through a FOL formula $(\mathbf{x}_1 \wedge \neg \mathbf{x}_2) \vee (\neg \mathbf{x}_1 \wedge \mathbf{x}_2) \Leftrightarrow \mathbf{f}$. As seen before, through the T-Norm operation we can convert the logic rule into a numerical constraint, compute its violation as:

$$\begin{aligned} \varphi_{\mathbf{x}_1 \oplus \mathbf{x}_2 \rightarrow \mathbf{f}} &= (\mathbf{x}_1 + \mathbf{x}_2 - 2\mathbf{x}_1\mathbf{x}_2)(1 - \mathbf{f}), \\ \varphi_{\mathbf{f} \rightarrow \mathbf{x}_1 \oplus \mathbf{x}_2} &= \mathbf{f}(1 - (\mathbf{x}_1 + \mathbf{x}_2 - 2\mathbf{x}_1\mathbf{x}_2)) \end{aligned} \quad (2)$$

In Fig. 1, we reported an example of the proposed strategy starting from $n = 10$ randomly selected labelled data and by selecting $p = 5$ samples at each iteration violating the most Eq. 2, and for $q = 100$ iterations. We can appreciate how, as is often the case, the initial random sampling (blue points—figure on the left) does not well represent the whole data distribution: no samples drawn from the bottom-right quadrant. Nonetheless, the proposed method immediately discovers the data distribution not represented by the initial sampling (orange points—figure on the left), by selecting the samples violating $\mathbf{x}_1 \oplus \mathbf{x}_2 \rightarrow \mathbf{f}$. After 5 iterations (figure at the centre) the network has mostly learnt the correct data distribution. Later, the proposed strategy refines network predictions by sampling along the decision boundaries (blue points—figure on the right), allowing the network to almost already solve the learning problem (accuracy $\sim 100\%$) in just 10 iterations. As it will be seen in the next section, standard random selection (but also uncertainty-based ones) will require many more iterations.

2.3 Real-life scenario: partial knowledge and different type of rules

It is clear that, in the case of the XOR-like problem, the knowledge is complete: if we compute the predictions directly through the rule, we already solve the learning problem. However, the purpose of this simple experiment is to show the potentiality of the proposed approach in integrating the available symbolic knowledge into a learning problem. In real-life scenarios, such a situation is unrealistic, but still we might have access to some partial knowledge that may allow

solving more quickly a given learning problem. Also, it may facilitate domain experts to accept and understand the active learning labelling process, since here the samples to label are the ones violating the knowledge they provided.

More precisely, when we consider structured data (e.g., tabular data), a domain expert may know some simple relations taking into consideration few features and the output classes. This knowledge may not be sufficient to solve the learning problem, but a KAL strategy can still exploit it to drive the network to a fast convergence, as we will see in Section 3. On the opposite, when we consider unstructured data (e.g., images or audio signals) the employed knowledge cannot directly rely on the input features. Nonetheless, in multi-label learning problems, a user may know in advance some relations between the output classes. Let us consider, as an example, a Dog-vs-Person classification: we might know that a main a dog is composed of several parts (e.g., a muzzle, a body, a tail). A straightforward translation of this compositional property into a FOL rule is $\mathbf{Dog} \Rightarrow \mathbf{Muzzle} \vee \mathbf{Body} \vee \mathbf{Tail}$. Formulating the composition in the opposite way is correct as well i.e., $\mathbf{Muzzle} \Rightarrow \mathbf{Dog}$. Also, in all classification problems, at least one of the main classes needs to be predicted, i.e., $\mathbf{Dog} \vee \mathbf{Person}$, with main classes being mutually exclusive in standard multi-class problems, i.e., $\mathbf{Dog} \oplus \mathbf{Man}$. Finally, we can always incorporate an uncertainty-like rule requiring each predicate to be either true or false, i.e., $\mathbf{Dog} \oplus \neg\mathbf{Dog}$.

3 Experiments

In this work, we considered six different learning scenarios, comparing the proposed technique with several standard active strategies. We evaluated the proposed method on two standard classification problems [4], the inference of the *XOR-like* problem (already introduced in Section 2.2), and the classification of *IRIS* plants given their characteristics. To assess the validity of the proposed method on regression tasks, we experimented on the Insurance dataset⁶, which requires to model insurance charges based on insured persons features. We also considered two standard image-classification tasks: the *ANIMALS* dataset, representing 7 classes of animals extracted from ImageNet [13], and the Caltech-UCSD Birds-200-2011 dataset (*CUB200*, [53]), a fine-grained classification dataset representing 200 bird species. At last, as a proof of concept, we analysed the performances of the KAL in the simple *DOGvsPERSON* object recognition task, a novel publicly available dataset that we extracted from *PASCAL-Part* [7]. For more details regarding the latter, please refer to Appendix A.2. For each dataset, n, p, q , as well as the number of training epochs and the network structure are arbitrarily fixed in advance according to the number of classes, the dataset size and the task complexity. Reported average results are computed on the test sets of a k -fold Cross Validation (with $k = 10$ in the first three tasks and $k = 5$ in the computer vision ones). More details regarding each experimental problem, as well as the tables reporting all the rules employed, are available in Appendix A.3.

⁶ Available from Kaggle <https://www.kaggle.com/datasets/teertha/ushealthinsurancedataset>

The code to run all the experiments is published on a public GitHub repository⁷. A simple code example is also reported in Appendix B showing how to solve the *XOR-like* problem with the KAL strategy. All experiments were run on an Intel i7-9750H CPU machine with an NVIDIA 2080 RTX GPU and 64 GB of RAM.

Compared methods We compared KAL with 12 active learning strategies commonly considered in literature [42,60]. As representatives of uncertainty-based strategies, we considered **Entropy** [49] selecting samples associated to predictions having maximum entropy, **Margin** [38] predictions with minimum margin between the top-two classes, and **LeastConf** [54] predictions with the lowest confidences, together with their Monte Carlo Dropout versions [3] (respectively **Entropy_D**, **Margin_D**, **LeastConf_D**), which, by applying dropout a test time, compare the predictions of Monte Carlo sampled networks to better assess uncertain predictions. As more recent uncertainty-based strategies, we compared with Bayesian Active Learning by Disagreements **BALD** [17], with two strategies computing the margin by means of adversarial attacks **ADV_{DEEPPFOOL}** [16], **ADV_{BIM}** [60] and with **SupLoss** a simplified upper bound of the method proposed in [57] employing the actual labels (available only on benchmarks). As Diversity-based methods, we selected **KMeans** [61] and **KCenter** a greedy version of the CoreSet method [47]. More details regarding are reported in Appendix A.4, together with a table resuming the associated losses.

3.1 KAL provides better performance than many active strategies

For a quantitative comparison of the different methods, we evaluated the network accuracy when equipped with the different active learning strategies. In Figure 2 we reported the average F1 scores (R score for regression) budget curves when increasing the number of selected labelled data. In Table 1 we also report the Area Under the Budget Curves (AUBC), as defined in [59].

XOR-like, IRIS In both standard machine learning problems, we can observe how KAL and KAL_D (the corresponding Monte-Carlo Dropout version) reach the highest performance with a 4-10 % higher AUC over standard uncertainty-based strategies in both cases. The only competitive methods in both cases are the CoreSet-based approach KCenter and the SupLoss method. This behaviour will be better analysed in Section 3.3. Interestingly, when analysing the corresponding plots in Fig. 2 we can appreciate how the proposed methods not only allows to reach a higher overall accuracy, but it also enables the network to learn more quickly the given tasks w.r.t. the other ones. While this was an expected behaviour on the *XOR-like* task since the provided rules completely explain the learning problem, on the *IRIS* classification task it is surprising since only 3 simple rules are given, considering a maximum of 2 features each (e.g., $\neg\text{Long_Petal} \Rightarrow \text{Setosa}$).

⁷ KAL repository: github.com/gabrieleciravegna/Knowledge-driven-Active-Learning




Fig. 2. Average test performance growth when increasing the number of labelled samples in terms of F1 or R score (%) in the regression task. Confidence intervals are not reported for better readability. Method variants (e.g., the Monte-Carlo Dropout versions) are displayed with the same colour.

Insurance (R) Also in the regression scenario, KAL results to be the most effective active learning strategy, with only the CoreSet-based approach KCenter reaching similar performance (top-right plot in Fig. 2). Other methods, instead, report average performance at least 10 % lower than KAL. Furthermore, also in this case, KAL employs simple relations like $\neg\text{Smoker} \wedge \text{Age} < 40 \Leftrightarrow \text{Charge} < 7500$. Uncertainty-based strategies are not reported in this case, as they cannot be applied in regression problems (unless using auxiliary models to estimate confidence over open intervals as in [10]).

ANIMAL, CUB200 A slightly-different situation can be observed in the image classification tasks (bottom plots in Fig. 2). Here we notice the importance of employing well-structured knowledge. In the *ANIMALS* task, indeed, only 17 rules are provided relating animal species and their characteristics (e.g., **Fly** \Rightarrow \neg **Penguin**). In this case, the results with KAL are only on average w.r.t. uncertainty-based approaches (better than LeastConf, BALD and ADV but worse than Margin and Entropy). On the contrary, KAL performs much better than KCENTER and KMEANS which are unable to correctly represents data distributions in complex scenarios even though being applied in the network latent space. In the *CUB200* task instead, where 311 rules are employed in the KAL strategy considering bird species and their attributes (e.g.,

Table 1. Comparison of the methods in terms of the mean F1 score (R score on the regression dataset) AUBC and standard deviation when increasing the number of labelled points. On top, starting and ending labelling budget for each dataset. The two best results are reported in bold. Uncertainty strategies are reported with – in the regression task, since they cannot be applied in this context.

Dataset	XOR	IRIS	Insurance (R)	Animals	CUB200
Strategy Budget	10-100	10-50	10-300	100-2500	2000-7000
KAL	93.44 ± 3.39	92.05 ± 4.19	67.52 ± 7.28	55.57 ± 1.55	51.98 ± 0.29
KAL _D	93.38 ± 3.31	91.47 ± 4.51	66.54 ± 9.04	55.52 ± 1.68	52.10 ± 0.24
ADV _{BIM}	82.77 ± 11.06	90.60 ± 3.92	–	53.93 ± 0.42	50.77 ± 0.50
ADV _{DEEFOOL}	83.79 ± 9.49	90.45 ± 4.59	–	54.35 ± 1.02	50.41 ± 0.29
BALD	78.13 ± 9.75	75.21 ± 10.18	–	53.87 ± 1.36	51.17 ± 0.62
KCENTER	91.84 ± 1.73	90.55 ± 4.76	66.04 ± 6.04	43.37 ± 3.29	48.90 ± 0.39
KMEANS	83.53 ± 3.89	86.49 ± 9.80	53.63 ± 12.99	52.87 ± 1.18	49.90 ± 0.33
Entropy	81.98 ± 10.77	75.04 ± 10.05	–	56.85 ± 0.92	51.91 ± 0.12
Entropy _D	83.21 ± 10.98	75.07 ± 9.88	–	57.68 ± 0.95	51.92 ± 0.39
LeastConf	83.12 ± 11.31	80.21 ± 15.74	–	53.95 ± 1.99	50.17 ± 0.28
LeastConf _D	84.77 ± 10.95	80.50 ± 15.91	–	54.31 ± 2.17	50.07 ± 0.45
Margin	83.12 ± 11.31	81.85 ± 16.90	–	57.59 ± 1.24	51.64 ± 0.30
Margin _D	84.77 ± 10.95	80.88 ± 16.24	–	56.88 ± 1.33	51.54 ± 0.48
Random	88.96 ± 2.90	88.08 ± 6.39	56.66 ± 12.87	54.22 ± 1.11	50.63 ± 0.25
SupLoss	90.81 ± 2.13	90.24 ± 3.60	42.44 ± 5.65	54.14 ± 1.48	49.42 ± 0.23

WhitePelican \Rightarrow **BlackEye** \vee **SolidBellyPattern** \vee **SolidWingPattern**), the proposed approaches are once again the best two strategies. SupLoss, instead, provide low performances in the computer vision problems. We believe that selecting samples with high supervision loss is not an optimal active strategy in this scenario, as it might mostly select outliers.

These results prove that KAL is a very effective active learning strategy when the provided knowledge sufficiently represents the given task, both in standard and in computer vision problems. In the *ANIMALS* task, instead, where the provided knowledge is scarce, KAL performance are only on average w.r.t. uncertainty strategies.

3.2 Ablation Studies

Amount of knowledge directly proportional to performance improvement To further show the importance of having a diverse and rich set of rules as introduced in Section 3.1, we performed here an ablation study. Table 2 reports the performance of the network when equipped with a KAL strategy considering only 0%, 25%, 50%, 75% or 100% of the available knowledge. The results show evidently that the amount of knowledge is directly proportional to the performance improvement, up to +1.8%. In the 0% scenario, the only rule employed is the uncertainty-like rule, which was always retained. Notice how 50.13 is similar to

Table 2. Ablation study on the quantity of knowledge employed to support the KAL strategy on the *CUB200* dataset. The amount of knowledge is directly proportional to the increase of performance.

KAL _{0%}	KAL _{25%}	KAL _{50%}	KAL _{75%}	KAL _{100%}
50.13	50.19	50.22	51.28	51.98

the LeastConf result (50.20), suggesting that KAL without any further knowledge results in an uncertainty-based strategy. In Appendix A.5, we report the complete table showing that this result is valid for all experimented scenarios.

Selecting diverse constraint violations and employing uncertainty-like rule improves the performance Given a set of rules \mathcal{K} , the proposed method might in theory select p samples all violating the same rule $\phi_k(f(x))$. To avoid this issue, we select a maximum number r of samples violating a certain rule k , similarly to [5] introducing diversity in margin-based approaches. Specifically, we group samples $x \in X_u$ according to the rule they violate the most, and we allow a maximum number of $p/2$ samples from each group (still following the ranking given by Eq. 1). In Appendix A.6, we report a table showing how requiring samples violating diverse constraints improves the overall quality of the KAL selection process. Also, we show the importance of adding the uncertainty-like rule $\bigwedge_i \mathbf{f}_i \oplus \neg \mathbf{f}_i$ introduced at the end of Section 2.3. Together, these two features allow improving the average performances of the network up to 2 %.

3.3 KAL discovers novel data distributions, unlike uncertainty strategies

To further analyse the results obtained, in Figure 3 we report the samples selected by some compared strategies at the last iteration on the *XOR-like* task (100 labelled data), starting from the same randomly selected samples of Figure 1. As introduced in Section 2.2, the KAL strategy enables to discover novel




Fig. 3. Comparison of the sample selection process on the *XOR-like* task after 100 labelled samples (starting from the same points as in Figure 1). Notice how uncertainty-based strategies (BALD, Margin) have not discovered the novel data distribution (right-bottom quadrant).

Table 3. Violation of the \mathcal{K}_{CUB-S} knowledge computed as the increased percentage over the violation of a model actively trained to respect this knowledge (KAL_{small}). The lower, the better. The proposed method ensures domain experts that their knowledge is acquired by the model.

KAL_{small}	Random	BALD	Entropy	LeastConf
+0.00%	+483.10%	+720.25	+861.74%	+1334.50%

data distribution (leftmost figure) even when they are not represented by the initial random sampling. On the contrary, uncertainty-based strategies (like Margin but also BALD, central figures) are unable to discover new data distributions. Indeed, all the data required to label is selected along the decision boundaries of already known distributions. For this reason, they provide mediocre results on average on the *XOR-like* and *IRIS* tasks and very high variance ($> 10 - 15\%$ on *IRIS*). The CoreSet representative strategy, instead, has covered the four quadrants. However, by only working on input features statistics and without notion on the predictions, this strategy does not choose points along the decision boundaries, preventing the network from reaching high accuracy performances. More figures are reported in Appendix A.7.

3.4 KAL ensures domain experts that their knowledge is acquired

It may be the case that domain experts are provided with a small corpus of rules which is crucial to be respected by the trained model, e.g., because it has to be deployed in a sensitive context. By always selecting the data that violate this corpus of rules, KAL ensures them that their knowledge is acquired by the model. To simulate this scenario, we computed the argument of Eq. 1 over a small part of the CUB knowledge \mathcal{K}_{CUB-S} (where $-S$ stands for small) on the test data X_T for the f^b model trained with all the budget: $\varphi(\mathcal{K}_{CUB-S}, f^b, X_T) = \sum_{x \in X_T} \sum_{k \in \mathcal{K}_{CUB-S}} \varphi_k(f^b(x))$. In Table 3 we report the increased percentage of the violation by models trained following a few compared methods w.r.t. the violation of a model trained following the KAL strategy and equipped with the small corpus of rules (KAL_{small}). The complete table together with more experimental detail is reported in appendix A.8. For the sake of completeness, this model reaches a lower test F1 AUBC (49.04). Nonetheless, it ensures domain experts that the provided knowledge is respected significantly more than using Random selection, or, worse, standard active learning strategies.

3.5 KAL can be used even without domain-knowledge

It might be argued that the proposed strategy can be employed only when a domain knowledge is available. However, recent works in the eXplainable AI (XAI) field [23,43,2,9] have shown that we can extract the same knowledge from a trained model. In general, they achieve this by training a white-box model (e.g. a decision tree) to globally explain the behaviour of a neural network. Here, we

Table 4. Accuracy of the KAL strategy coupled with a XAI method, extracting the knowledge from the same network. Notice how the AUCB reduction of performance is always smaller than 1 – 2 %.

	XOR	IRIS	Animals	CUB
KAL	93.44 ± 3.39	92.05 ± 4.19	55.57 ± 1.55	51.98 ± 0.29
KAL _{XAI}	92.18 ± 2.64	90.00 ± 5.99	54.07 ± 2.21	50.33 ± 0.47

propose to employ these FOL-based explanations (\mathcal{K}_{XAI}) as the base knowledge of the proposed strategy when no other knowledge is available (KAL_{XAI}). More precisely, after each iteration, we employ a simple decision tree as proposed in [23] to extract the knowledge. More details on how we trained the XAI method are reported in Appendix A.9. However, the knowledge may be partial, particularly during the first iterations, since it is extracted on the training distribution only. Therefore, we use Eq. 1 to select only 60% of the samples, with the remaining randomly selected. This allows to eventually recover the complete knowledge. In Table 4, we report the performance of the network when equipped with this strategy (KAL_{XAI}), together with the performance of the standard strategy. Notice how the reduction of performance is less than 1-2%, confirming the validity of the proposed approach even in this scenario. The amount of randomly chosen samples has not been cross-validated, therefore we expect to get even higher results by fine-tuning this parameter.

3.6 KAL can be employed in object recognition tasks

To test the proposed method in an object recognition context, as a proof of concept, we experimented on the simple *DOGvsPERSON* dataset. On this task, we compute the AUCB of the mean Average Precision curves. Also in this case, the network increases more its performances when equipped with the KAL strategy (55.90 ± 0.39) with respect to standard random sampling (51.41 ± 1.25) but also compared to the SupLoss method (55.30 ± 0.54), proving the efficacy of the KAL strategy also in this context. A figure showing the three budget curves is reported in appendix A.10. Reported results are averaged over 3 initialization seeds. We only compared with Random selection and the simplified version of [57], since uncertainty-based strategies are not straightforward to apply in this context [25]. Finally, we want to highlight that the SupLoss performance reported here is an upper bound of the performance of the method proposed in [57]. Particularly in this context, we believe that the object recognition loss might not be easily learnt by an external model, thus reducing the performance of the SupLoss method.

3.7 KAL is not computationally expensive

When devising novel active learning techniques, of crucial importance is also the computational effort. Indeed, since re-training a deep neural network already requires a substantial amount of resources, the associated active strategy should

Table 5. Computational demand of some of the compared methods computed as the proportional increase over the time required for random sampling as defined in [60]. The lower, the better. Notice how the proposed method is less computationally expensive than many recent methods. Standard deviation is not reported for better readability, but it is reported in Appendix A.11 with the average time for all methods.

Strategy	XOR	Iris	Insurance (R)	Animals	CUB200
KAL	5.22	16.92	20.52	41.34	180.22
KAL _D	15.79	23.80	31.43	53.41	197.05
ADV _{BIM}	36.93	657.37	–	5600.65	7440.67
ADV _{DEEFOOL}	401.73	6451.91	–	57950.46	188435.28
BALD	20.70	21.50	–	157.47	613.48
KCENTER	31.82	42.89	158.28	2379.57	8713.37
KMEANS	7.90	142.36	28.75	718.70	4724.60
Entropy	4.00	13.07	–	14.58	39.02
Entropy _D	14.51	18.79	–	22.83	52.77

be as light as possible. In [60], authors used as a term of comparison the average time needed to randomly sample a novel batch of data. In Table 5, we report the proportional increased computational time w.r.t. random sampling. KAL strategies are not computationally expensive (5-180 times slower than random sampling). On the contrary, BALD (20-613) and, more importantly, KMEANS (8-4724), KCENTER (30-8713) and ADV-based (37-188435) strategies demand considerable amounts of computational resources, strongly reducing the usability of the same methods. Standard uncertainty-based techniques like Entropy, instead, are not computationally demanding, with only the Dropout versions increasing 14-52 times the computational demand of random sampling (similarly to KAL_D). The complete table with all methods is reported in Appendix A.11.

4 Related work

Active Learning In the literature, two main approaches have been followed: uncertainty sampling which selects the data on which the model is the least confident; curriculum learning which focuses first on easy samples and then extends the training set to incorporate more difficult ones while also targeting more diversity. Standard uncertainty-based strategies choose samples associated to maximal prediction entropy [26] or at minimum distance from the hyperplane in SVM [45] or with the highest variation ratio in Query-by-committee with ensemble methods [15,3]. Establishing prediction uncertainty is more difficult with DL models. Indeed, they tend to be over-confident, particularly when employing softmax activation functions [52]. Furthermore, as there is no easy access to the distance to the decision boundary, it needs to be computed. This problem has been tackled by devising different uncertain strategies, such as employing Bayesian Neural Network with Monte Carlo Dropout [17], by calculating the minimum distance required to create an adversarial example [16], or even predicting

the loss associated to unlabelled sample [57]. As pointed out by [40], however, uncertain strategy may choose the same categories many times and create unbalanced datasets. To solve this, uncertain sample selection can be coupled with diversity sampling strategies. Diversity can be obtained by preferring batches of data maximizing the mutual information between model parameters and predictions [27], or selecting core-set points [46], samples nearest to k-means cluster centroids [61], or even by learning sample dissimilarities in the latent space of a VAE with an adversarial strategy [50] or by means of a GCN [6].

Hybrid Models It has been pondered that human cognition mainly consists in two different tasks: perceiving the world and reasoning over it [51]. While these two tasks in humans take place at the same times, in artificial intelligence they are separately conducted by machine learning and logic programming. It has been argued that joining these tasks (to create a so-called hybrid model) may overcome some of the most important limits of deep learning, among which the “*data hungry*” issue [34]. In the literature, there exists a variety of proposals aiming to create hybrid models, ranging from Statistical Relational Learning (SRL) [30] and Probabilistic Logic Programming [12] which focuses on integrating learning with logic reasoning, to enhanced networks focusing on relations or with external memories [44,22]. Recently, several approaches have been devised to compute and enforce the satisfaction of a given domain knowledge within DL models [33,1,56], (see survey [20] for a complete list of works in this domain). Among these options, in this work we chose to employ the learning from constraints framework [21,14] since it provides the great logical expressivity (both universal and existential quantifier) and a straightforward implementation.

5 Conclusions

In this paper, we proposed an active learning strategy leveraging available domain knowledge to select the data to label. The performance of a model equipped with such a strategy outperforms standard uncertainty-based approaches in context where the domain knowledge is sufficiently rich, without being computationally demanding. Furthermore, we think that KAL could induce more trust in DL, since it enables non-expert users to train models leveraging their domain knowledge and ensuring them that it will be acquired by the model. A main limitation of the proposed approach is in computer vision contexts, if no attributes about main classes are known. A possible solution could be to automatically extract such concepts from the latent space of the network, as proposed in [18,8]. Also, if the domain knowledge is highly complex, FOL may not be able to fully express it and higher-order logic may be required.

Acknowledgments This work was supported by the EU Horizon 2020 project AI4Media, under contract no. 951911 and by the French government, through the 3IA Côte d’Azur, Investment in the Future, project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002.

References

1. Badreddine, S., Garcez, A.d., Serafini, L., Spranger, M.: Logic tensor networks. *Artificial Intelligence* **303**, 103649 (2022)
2. Barbiero, P., Ciravegna, G., Giannini, F., Lió, P., Gori, M., Melacci, S.: Entropy-based logic explanations of neural networks. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 36, pp. 6046–6054 (2022)
3. Beluch, W.H., Genewein, T., Nürnberger, A., Köhler, J.M.: The power of ensembles for active learning in image classification. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 9368–9377 (2018)
4. Bishop, C.M.: *Pattern Recognition and Machine Learning*, vol. 4. Springer (2006)
5. Brinker, K.: Incorporating diversity in active learning with support vector machines. In: *ICML*. pp. 59–66 (2003)
6. Caramalau, R., Bhattarai, B., Kim, T.K.: Sequential graph convolutional network for active learning. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9583–9592 (2021)
7. Chen, X., Mottaghi, R., Liu, X., Fidler, S., Urtasun, R., Yuille, A.: Detect what you can: Detecting and representing objects using holistic models and body parts. In: *CVPR*. pp. 1971–1978 (2014)
8. Chen, Z., Bei, Y., Rudin, C.: Concept whitening for interpretable image recognition. *Nature Machine Intelligence* **2**(12), 772–782 (dec 2020). <https://doi.org/10.1038/s42256-020-00265-z>, <https://doi.org/10.1038/2Fs42256-020-00265-z>
9. Ciravegna, G., Barbiero, P., Giannini, F., Gori, M., Liò, P., Maggini, M., Melacci, S.: Logic explained networks. *Artificial Intelligence* **314** (2023). <https://doi.org/10.1016/j.artint.2022.103822>, <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85141927580&doi=10.1016%2fj.artint.2022.103822&partnerID=40&md5=e7b7bbb63e2e654e00e7195706cad1c6>, cited by: 3; All Open Access, Green Open Access
10. Corbiere, C., Thome, N., Saporta, A., Vu, T.H., Cord, M., Perez, P.: Confidence estimation via auxiliary models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021)
11. Das, A., Rad, P.: Opportunities and challenges in explainable artificial intelligence (xai): A survey. *arXiv preprint arXiv:2006.11371* (2020)
12. De Raedt, L., Kimmig, A.: Probabilistic (logic) programming concepts. *Machine Learning* **100**(1), 5–47 (2015)
13. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 248–255. Ieee (2009)
14. Diligenti, M., Gori, M., Sacca, C.: Semantic-based regularization for learning and inference. *Artificial Intelligence* **244**, 143–165 (2017)
15. Ducoffe, M., Precioso, F.: Active learning strategy for cnn combining batchwise dropout and query-by-committee. In: *ESANN* (2017)
16. Ducoffe, M., Precioso, F.: Adversarial active learning for deep networks: a margin based approach. *arXiv:1802.09841* (2018)
17. Gal, Y., Islam, R., Ghahramani, Z.: Deep bayesian active learning with image data. In: *International Conference on Machine Learning*. pp. 1183–1192. PMLR (2017)
18. Ghorbani, A., Wexler, J., Zou, J.Y., Kim, B.: Towards automatic concept-based explanations. *Advances in Neural Information Processing Systems* **32** (2019)

19. Gilpin, L.H., Bau, D., Yuan, B.Z., Bajwa, A., Specter, M., Kagal, L.: Explaining explanations: An overview of interpretability of machine learning. In: 2018 IEEE 5th International Conference on data science and advanced analytics (DSAA). pp. 80–89. IEEE (2018)
20. Giunchiglia, E., Stoian, M.C., Lukasiewicz, T.: Deep learning with logical constraints. arXiv preprint arXiv:2205.00523 (2022)
21. Gnecco, G., Gori, M., Melacci, S., Sanguineti, M.: Foundations of support constraint machines. *Neural computation* **27**(2), 388–480 (2015)
22. Graves, A., Wayne, G., Reynolds, M., Harley, T., Danihelka, I., Grabska-Barwinska, A., Colmenarejo, S.G., Grefenstette, E., Ramalho, T., Agapiou, J.P., Badia, A.P., Hermann, K.M., Zwols, Y., Ostrovski, G., Cain, A., King, H., Summerfield, C., Blunsom, P., Kavukcuoglu, K., Hassabis, D.: Hybrid computing using a neural network with dynamic external memory. *Nature* **538**, 471–476 (2016)
23. Guidotti, R., Monreale, A., Ruggieri, S., Pedreschi, D., Turini, F., Giannotti, F.: Local rule-based explanations of black box decision systems. arXiv preprint arXiv:1805.10820 (2018)
24. Hájek, P.: *Metamathematics of fuzzy logic* (1998)
25. Haussmann, E., Fenzi, M., Chitta, K., Ivanecky, J., Xu, H., Roy, D., Mittel, A., Koumchatzky, N., Farabet, C., Alvarez, J.: Scalable active learning for object detection. In: IEEE IV Symposium. pp. 1430–1435. IEEE (2020)
26. Houthby, N., Huszár, F., Ghahramani, Z., Lengyel, M.: Bayesian active learning for classification and preference learning (2011)
27. Kirsch, A., Van Amersfoort, J., Gal, Y.: Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. *NeurIPS* **32**, 7026–7037 (2019)
28. Klement, E., Mesiar, R., Pap, E.: *Triangular norms*, vol. 8. Springer Science & Bus. Media (2013)
29. Koh, P.W., Nguyen, T., Tang, Y.S., Mussmann, S., Pierson, E., Kim, B., Liang, P.: Concept bottleneck models. In: ICML. pp. 5338–5348 (2020)
30. Koller, D., Friedman, N., Džeroski, S., Sutton, C., McCallum, A., Pfeffer, A., Abbeel, P., Wong, M.F., Meek, C., Neville, J., et al.: *Introduction to statistical relational learning*. MIT press (2007)
31. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *nature* **521**(7553), 436–444 (2015)
32. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv:1711.05101 (2017)
33. Manhaeve, R., Dumancic, S., Kimmig, A., Demeester, T., De Raedt, L.: Deep-problog: Neural probabilistic logic programming. *advances in neural information processing systems* **31** (2018)
34. Marcus, G.: Deep learning: A critical appraisal. arXiv:1801.00631 (2018)
35. Marra, G., Giannini, F., Diligenti, M., Gori, M.: Lyrics: A general interface layer to integrate logic inference and deep learning. In: ECML/PKDD (2019)
36. McCallumzy, A.K., Nigamy, K.: Employing em and pool-based active learning for text classification. In: ICML. pp. 359–367. Citeseer (1998)
37. Melacci, S., Ciravegna, G., Sotgiu, A., Demontis, A., Biggio, B., Gori, M., Roli, F.: Domain knowledge alleviates adversarial attacks in multi-label classifiers. *IEEE PAMI* pp. 1–1 (2021). <https://doi.org/10.1109/TPAMI.2021.3137564>
38. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning. In: NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011 (2011), http://ufldl.stanford.edu/housenumbers/nips2011_housenumbers.pdf

39. Novák, V., Perfilieva, I., Mockor, J.: *Mathematical Principles of Fuzzy Logic*, vol. 517. Springer Science & Business Media (1999)
40. Pop, R., Fulop, P.: Deep ensemble bayesian active learning: Addressing the mode collapse issue in monte carlo dropout via ensembles. *arXiv:1811.03897* (2018)
41. Redmon, J., Farhadi, A.: Yolo v3: An incremental improvement. *CoRR abs/1804.02767* (2018), <http://arxiv.org/abs/1804.02767>
42. Ren, P., Xiao, Y., Chang, X., Huang, P.Y., Li, Z., Gupta, B.B., Chen, X., Wang, X.: A survey of deep active learning. *ACM Computing Surveys (CSUR)* **54**(9), 1–40 (2021)
43. Ribeiro, M.T., Singh, S., Guestrin, C.: Anchors: High-precision model-agnostic explanations. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 32 (2018)
44. Santoro, A., Raposo, D., Barrett, D.G., Malinowski, M., Pascanu, R., Battaglia, P., Lillicrap, T.: A simple neural network module for relational reasoning. *NeurIPS* **30** (2017)
45. Schohn, G., Cohn, D.: Less is more: Active learning with support vector machines. *ICML (10 2000)*
46. Sener, O., Savarese, S.: Active learning for convolutional neural networks: A core-set approach. In: *ICLR* (2018)
47. Sener, O., Savarese, S.: Active learning for convolutional neural networks: A core-set approach. *arXiv preprint arXiv:1708.00489* (2017)
48. Serafini, L., d’Avila Garcez, A.: Logic tensor networks: Deep learning and logical reasoning from data and knowledge. *arXiv:1606.04422* (2016)
49. Settles, B.: *Active learning literature survey* (2009)
50. Sinha, S., Ebrahimi, S., Darrell, T.: Variational adversarial active learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5972–5981 (2019)
51. Solso, R., MacLin, M., MacLin, O.: *Cognitive psychology*. Pearson Education New Zealand (2005)
52. Thulasidasan, S., Chennupati, G., Bilmes, J., Bhattacharya, T., Michalak, S.: On mixup training: Improved calibration and predictive uncertainty for deep neural networks. *NeurIPS* **32** (2019)
53. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: *The Caltech-UCSD Birds-200-2011 Dataset*. Tech. Rep. CNS-TR-2011-001, CalTech (2011)
54. Wang, D., Shang, Y.: A new active labeling method for deep learning. In: *2014 International joint conference on neural networks (IJCNN)*. pp. 112–119. IEEE (2014)
55. Winston, P., Horn, B.: *Lisp* (1986)
56. Xu, J., Zhang, Z., Friedman, T., Liang, Y., Broeck, G.: A semantic loss function for deep learning with symbolic knowledge. In: *International conference on machine learning*. pp. 5502–5511. PMLR (2018)
57. Yoo, D., Kweon, I.: Learning loss for active learning. In: *IEEE CVPR*. pp. 93–102 (2019)
58. Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., Xiao, J.: Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv:1506.03365* (2015)
59. Zhan, X., Liu, H., Li, Q., Chan, A.B.: A comparative survey: Benchmarking for pool-based active learning. In: Zhou, Z.H. (ed.) *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*. pp.

- 4679–4686. International Joint Conferences on Artificial Intelligence Organization (8 2021). <https://doi.org/10.24963/ijcai.2021/634>, <https://doi.org/10.24963/ijcai.2021/634>, survey Track
60. Zhan, X., Wang, Q., Huang, K.h., Xiong, H., Dou, D., Chan, A.B.: A comparative survey of deep active learning (2022). <https://doi.org/10.48550/ARXIV.2203.13450>, <https://arxiv.org/abs/2203.13450>
61. Zhdanov, F.: Diverse mini-batch active learning. arXiv:1901.05954 (2019)
62. Zhu, X., Goldberg, A.: Introduction to semi-supervised learning. Synthesis lectures on artificial intelligence and machine learning **3**(1), 1–130 (2009)

A Experimental details and further results

A.1 Converting Logic Formulas into Numerical Constraints

To better understand how the KAL framework works, in this section we will briefly go through some basic principles on FOL and later we will focus on how to convert rules into numerical constraints by means of different T-Norms.

FOL Domains, Individuals, Functions, Predicates and Constraints In First Order Logic, a *Domain* D is a data space representing *Individuals* x which share the same representation space X . As an example, a domain can be composed of images representing birds, as in the CUB200 dataset. Each bird is an *Individual*, which is represented in a certain Domain by its features (e.g., in this case, by its image). A *Function* is a mapping of individuals between an input and an output domain. In this paper, we only focus on unary-function, i.e., functions that take only one individual in input and transform it into an individual of an output domain. An example of a function is $Age(x)$, which returns the age of a bird given its image representation. N-ary functions, taking more than one individual in input (e.g., $Relationship(x, x)$, returning the kind of relationship given two bird images) are also supported by the framework but are not used in this case. The output domain can be the same or a different domain w.r.t the input one. A *Predicate* is a special type of function returns as output a truth value 0, 1 if we consider boolean predicate, or in $[0, 1]$ if also consider fuzzy values as in this work. An example of predicate is $Hummingbird(x)$, which tells you whether the considered bird belongs to the Hummingbird species. Both predicates and functions can be parametrized and learnt. In our case, all considered predicates were modelled by means of a neural network. At last, we can provide our knowledge about a certain domain by means of a set of *Constraints*. A constraint is a FOL rule defined on functions and predicates (which are the atoms of the rule). An example of a constraint could be $\forall x, Hummingbird(x) \rightarrow Bird(x)$. Finally, existential quantifier $\exists x \in X$ are also supported by the framework but we have not used them in this work.

Converting FOL rules into numerical constraints To convert a FOL formula into a numerical constraints we need a way to convert connectives and quantifiers into numerical operators. To do so, we employ the fuzzy generalization of FOL that was first proposed by [39]. More precisely, T-norm fuzzy logic [24] generalize Boolean Logic to continuous values in $[0, 1]$. T-norm fuzzy logics are defined by the operator modelling the AND logic operator. All the other operators are generally derived from it. In Table 6 (reported from [35]), some possible implementations of common connectives when using the Product (the one used in this paper), the Lukasiewicz and the Gödel Logics.

Finally, to evaluate the violation of each numerical constraints, a loss function has to be chosen (also called *generator*). These functions need to be strictly decreasing $g : [0, 1] \rightarrow [0, +\infty]$ and such that $g(1) = 0$. Possible choices are $g(x) = 1 - x$ (as used in this paper), or $g(x) = -\log(x)$.

Table 6. Some of the most used example of T-Norm with their translation of some logic operators. Table extracted from [35]

T-norm Op.	Product	Lukasiewicz	Gödel
$x \wedge y$	$x \cdot y$	$\max(0, x + y - 1)$	$\min(x, y)$
$x \vee y$	$x + y - x \cdot y$	$\min(1, x + y)$	$\max(x, y)$
$\neg x$	$1 - x$	$1 - x$	$1 - x$
$x \rightarrow y$	$1 - (x \cdot (1 - y))$	$\min(1, 1 - x + y)$	$x \leq y ? 1 : y$

A.2 The *DOGvsPERSON* dataset

The *DOGvsPERSON* dataset is a publicly available dataset that we have created for showing the potentiality of the proposed method on a simple object-recognition problem, where well-defined relations are present among the classes, and we can employ common-knowledge rule to easily relate them. It is extracted from the *PASCAL-Part* dataset by considering only the Dog and Person main classes and their corresponding parts. In the original *PASCAL-Part* dataset, labels are given in the form of segmentation masks. We extracted a bounding box from each mask by considering the leftmost and highest pixel as the first coordinate and the rightmost and lowest pixel as the second one. Very specific parts are merged into a single class, following the approach of [48] (e.g., **LeftLowerArm**, **LeftUpperArm**, **RightLowerArm**, **RightUpperArm** becomes **Arm**). Differently from the standard *PASCAL-Part* dataset, however, the parts in common to different objects are considered as different classes (**head** becomes **DogHead**, **PersonHead**). Furthermore, we only consider masks having areas $\geq 1\%$ of the whole image areas as valid label. At last, only classes appearing at least 100 times are retained. This lead to a total of 20 classes with 2 main classes (**Dog** and **Person**) and 18 parts (**DogEar**, **DogHead**, **DogLeg**, **DogMuzzle**, **DogNeck**, **DogNose**, **DogPaw**, **DogTail**, **DogTorso**, **PersonArm**, **PersonFoot**, **PersonHair**, **PersonHand**, **PersonHead**, **PersonLeg**, **PersonNeck**, **PersonNose**, **PersonTorso**) displayed in a total of 4304 samples. Final classes are distributed in the samples as shown in Figure 4.

A.3 Network architectures, hyperparameters, and domain knowledges

XOR-like The problem of inferring the *XOR-like* operation has been already introduced in Section 2.2: it is an artificial dataset consisting of 100000 samples $x \in X \subset \mathbb{R}^2$, mapped to the corresponding label $y \in Y \subset [0, 1]$ as defined in Section 2.2. A Multi-Layer Perceptron (MLP) $f: X \rightarrow Y$ is used to solve the task⁸. It is equipped with a single hidden layer of 100 neurons and Rectified Linear Unit (ReLU) activation, and a single output neuron with sigmoid activation. It

⁸ For all strategies requiring a multi-class output, we considered the network output as $\hat{f} = [f, 1 - f]$.




Fig. 4. Class distribution in the *DOGvsPERSON* dataset.

has been trained with an AdamW optimizer [32] for 250 epochs at each iteration, with a learning rate $\eta = 10^{-3}$. Standard cross-entropy loss has been used to enforce f to learn the available supervisions. By starting from $n = 10$ samples, we added $p = 5$ labelled samples at each iteration for a total of $q = 78$ iterations, resulting in a total final budget of $b = 400$ labelled samples. As anticipated, in the *XOR* problem, the rule employed for the KAL strategy is $\forall x \mathbf{x}_1 \oplus \mathbf{x}_2 \Leftrightarrow \mathbf{f}$, as also reported in Table 7.

Table 7. Domain knowledge on the *XOR-like* dataset.

$$\forall x \mathbf{x}_1 \oplus \mathbf{x}_2 \Leftrightarrow \mathbf{f}$$

For more real-life style problems, we have considered four more datasets where the domain-knowledge is partial or only related to the class functions.

IRIS The *IRIS*⁹ dataset is the standard iris-species classification problem. More precisely, the task consists in classifying $c = 3$ Iris species (Iris Setosa, Iris Versicolour, Iris Virginica) starting from $d = 4$ features (sepal length, sepal width,

⁹ *Iris*: <https://archive.ics.uci.edu/ml/datasets/iris>

petal length, petal width). To solve the learning problem, an MLP $f: X^d \rightarrow Y^c$ is employed, with one hidden layer composed of 100 neurons equipped with ReLu activation functions. It has been trained again with AdamW optimizer for 200 epochs at each iteration and learning rate $\eta = 3 \cdot 10^{-3}$. A cross-entropy loss is employed to enforce the supervisions. By starting from $n = 5$ points and by adding $p = 5$ labelled samples at each iteration for $q = 14$ iterations, leading to a total budget of $b = 75$ labelled samples. The knowledge employed in this case consists of 3 very simple rules (one per class) based on the two predicates **LongPetal** and **WidePetal** (built on the 3rd and 4th features of the dataset, respectively, as explained in Section 2.2). In addition, in this case, a mutually exclusive rule on the classes is also considered $\forall x, \mathbf{Setosa} \oplus \mathbf{Versicolour} \oplus \mathbf{Virginica}$, as reported in Table 8.

Table 8. Domain knowledge on the *IRIS* dataset.

$\forall x \neg \text{Long_Petal} \Leftrightarrow \text{Setosa}$ $\forall x \text{Long_Petal} \wedge \neg \text{Wide_Petal} \Leftrightarrow \text{Versicolour}$ $\forall x \text{Long_Petal} \wedge \text{Wide_Petal} \Leftrightarrow \text{Virginica}$ $\forall x \text{Setosa} \oplus \text{Versicolour} \oplus \text{Virginica}$

Insurance (R) In the Insurance dataset available from Kaggle¹⁰, the proposed task is to model personal insurance charges based on 6 features regarding the insured persons (Age, Sex, BMI, Number of Children, Smoker and Region).

The knowledge employed in this case consists of 4 rules working on 3 features and defining 4 intervals over the output space, as reported in Table 9.

ANIMALS The *Animals*' dataset is a collection of 8287 images of animals, taken from the ImageNet database¹¹. The task consists in the classification of 7 main classes (**Albatross, Giraffe, Cheetah, Ostrich, Penguin, Tiger, Zebra**) and 26 animal attributes (e.g., **Mammal, Fly** or **LayEggs**), for a total of $c = 33$ classes. In this case, a Resnet50 CNN has been employed to solve the task $f: X^d \rightarrow Y^c$. Going into more details, a transfer learning strategy has been employed: the network f has been pretrained on the ImageNet dataset [13], and two fully connected layers (the first one equipped with 100 neurons) have been trained (from scratch) on the *ANIMALS* dataset. Again, an AdamW optimizer is considered with a learning rate $\eta = 10^{-3}$ employed for 250 epochs of training at each iteration, with binary cross-entropy loss since we deal with a multi-label problem. We started with $n = 100$ labelled samples, and we added $p = 50$ samples each time for $q = 48$ iterations, for a final budget of $b = 2500$ labelled

¹⁰ <https://www.kaggle.com/datasets/teertha/ushealthinsurancedataset>

¹¹ *Animals (Imagenet)*: <http://www.image-net.org/>, released with BSD 3-Clause "New" or "Revised" License

Table 9. Domain knowledge on the Insurance dataset.

$\forall x \neg \text{Smoker} \wedge \text{Age} < 40 \Leftrightarrow \text{Charge} < 7500$ $\forall x \neg \text{Smoker} \wedge \text{Age} > 40 \Leftrightarrow \text{Charge} > 7500 \wedge \text{Charge} < 15000$ $\forall x \text{Smoker} \wedge \text{BMI} < 30 \Leftrightarrow \text{Charge} > 15000 \wedge \text{Charge} < 30000$ $\forall x \text{Smoker} \wedge \text{BMI} > 30 \Leftrightarrow \text{Charge} > 30000$
--

samples. In the case of *Animals*, the employed knowledge is a simple collection of 16 FOL formulas, defined by [55] as a benchmark. They involve relationships between animals and their attributes, such as $\forall x \text{Fly} \wedge \text{LayEggs} \Rightarrow \text{Bird}$. To this collection of rules, we have also added a mutual exclusive disjunction among the animal classes (only one animal is present in each image) and a standard disjunction over the animal attributes (each animal may be associated to many attributes). The complete list of rules employed is reported in Table 10.

Table 10. Domain knowledge on the *Animals* dataset.

$\forall x \text{Hair} \vee \text{Mammal}$ $\forall x \text{Milk} \Rightarrow \text{Mammal}$ $\forall x \text{Feather} \Rightarrow \text{Bird}$ $\forall x \text{Fly} \wedge \text{LayEggs} \Rightarrow \text{Bird}$ $\forall x \text{Mammal} \wedge \text{Meat} \Rightarrow \text{Carnivore}$ $\forall x \text{Mammal} \wedge \text{PointedTeeth} \wedge \text{Claws} \wedge \text{ForwardEyes} \Rightarrow \text{Carnivore}$ $\forall x \text{Mammal} \wedge \text{Hoofs} \Rightarrow \text{Ungulate}$ $\forall x \text{Mammal} \wedge \text{Cud} \Rightarrow \text{Ungulate}$ $\forall x \text{Mammal} \wedge \text{Cud} \Rightarrow \text{Eventood}$ $\forall x \text{Carnivore} \wedge \text{Tawny} \wedge \text{DarkSpots} \Rightarrow \text{Cheetah}$ $\forall x \text{Carnivore} \wedge \text{Tawny} \wedge \text{BlackStripes} \Rightarrow \text{Tiger}$ $\forall x \text{Ungulate} \wedge \text{LongLegs} \wedge \text{LongNeck} \wedge \text{Tawny} \wedge \text{DarkSpots} \Rightarrow \text{Giraffe}$ $\forall x \text{Blackstripes} \wedge \text{Ungulate} \wedge \text{White} \Rightarrow \text{Zebra}$ $\forall x \text{Bird} \wedge \neg \text{Fly} \wedge \text{LongLegs} \wedge \text{LongNeck} \wedge \text{Black} \Rightarrow \text{Ostrich}$ $\forall x \text{Bird} \wedge \neg \text{Fly} \wedge \text{Swim} \wedge \text{BlackWhite} \Rightarrow \text{Penguin}$ $\forall x \text{Bird} \wedge \text{GoodFlier} \Rightarrow \text{Albatross}$ $\forall x \text{Albatross} \oplus \text{Giraffe} \oplus \text{Cheetah} \oplus \text{Ostrich} \oplus \text{Penguin} \oplus \text{Tiger} \oplus \text{Zebra}$ $\forall x \text{Mammal} \vee \text{Hair} \vee \text{Milk} \vee \text{Feathers} \vee \text{Bird} \vee \text{Fly} \vee \text{Meat} \vee \text{Carnivore} \vee \text{PointedTeeth}$ $\vee \text{Claws} \vee \text{ForwardEyes} \vee \text{Hoofs} \vee \text{Ungulate} \vee \text{Cud} \vee \text{Eventood} \vee \text{Tawny} \vee \text{BlackStripes} \vee$ $\text{LongLegs} \vee \text{LongNeck} \vee \text{DarkSpots} \vee \text{White} \vee \text{Black} \vee \text{Swim} \vee \text{BlackWhite} \vee \text{GoodFlier}$
--

CUB200 The Caltech-UCSD Birds-200-2011¹² dataset [53] is a collection of 11,788 images of birds. The task consists in the classification of 200 birds species (e.g., **BlackfootedAlbatross**) and birds attributes (e.g., **WhiteThroat**, **Medium**

¹² *CUB200*: <http://www.vision.caltech.edu/visipedia/CUB-200-2011> released with MIT License.

Table 11. Domain knowledge on the *CUB200* dataset. Only a few rules for each type have been reported for the sake of clarity. Also, due to the length of the rules, rules with more than two terms implied have been truncated.

$\forall x$ Black_footed_Albatross \Rightarrow has_bill_shape_all-purpose \wedge has_underparts_color_yellow \wedge ... $\forall x$ Laysan_Albatross \Rightarrow has_bill_shape_hooked_seabird \wedge has_breast_pattern_solid \wedge ... $\forall x$ Sooty_Albatross \Rightarrow has_bill_shape_hooked_seabird \wedge has_wing_color_black \wedge ... $\forall x$ Groove_billed_Ani \Rightarrow has_bill_shape_hooked_seabird \wedge has_breast_pattern_solid \wedge ... $\forall x$ Crested_Auklet \Rightarrow has_wing_color_black \wedge has_upperparts_color_black \wedge $\forall x$ has_bill_shape_dagger \Rightarrow Green_Kingfisher \wedge Pied_Kingfisher \wedge ... $\forall x$ has_bill_shape_hooked_seabird \Rightarrow Laysan_Albatross \wedge Sooty_Albatross \wedge ... $\forall x$ has_bill_shape_all-purpose \Rightarrow Black_footed_Albatross \wedge Red_winged_Blackbird \wedge ... $\forall x$ has_bill_shape_cone \Rightarrow Parakeet_Auklet \wedge Indigo_Bunting \wedge ... $\forall x$ has_wing_color_brown \Rightarrow Brandt_Cormorant \wedge American_Crow \wedge $\forall x$ Black_footed_Albatross \vee Laysan_Albatross \vee Sooty_Albatross \vee ... $\forall x$ has_bill_shape_dagger \vee has_bill_shape_hooked_seabird \vee has_bill_shape_all-purpose \vee ...
--

Size). Attribute annotation, however, is quite noisy. For this reason, attributes are denoised by considering class-level annotations similarly to [29]. A certain attribute is set as present only if it is also present in at least 50 images of the same class. Furthermore, we only considered attributes present in at least 10 classes after this refinement. In the end, 108 attributes have been retained, for a total of $c = 308$ classes. Images have been resized to a dimension $d = 256 \times 256$ pixels. The same network as in the *ANIMALS* case has been employed to solve the learning problem, with two fully connected layers trained from scratch (the first one equipped with 620 neurons – twice the dimension of the following layer). Again, an AdamW optimizer is considered with a learning rate $\eta = 10^{-3}$ for 100 epochs of training. Owing to the increased difficulty of the problem, we started with $n = 2000$ labelled samples, and we added $p = 200$ samples for $q = 25$ iterations, for a final budget of $b = 7000$ labelled samples. The knowledge employed in this case consider the relation between the classes and their attributes, with logic implications both from the class to the attributes (e.g., **WhitePelican** \Rightarrow **BlackEye** \vee **SolidBellyPattern** \vee **SolidWingPattern**), and the vice-versa (e.g., **StripedBreastPattern** \Rightarrow **ParakeetAuklet** \vee **BlackthroatedSparrow** \vee ...) Furthermore, a disjunction on the main classes¹³ (**BlackFootedAlbatross** \vee **LaysanAlbatross** \vee **SootyAlbatross** \vee ...) and one on the attributes are considered (**DaggerBill** \vee **HookedBill** \vee **AllPurposeBill** \vee **ConeBill** \vee ...). A few examples of the rules employed are reported in Table 11.

DOGvsPERSON This dataset has already been introduced in Appendix A.2. Since we filtered out very small object masks in this dataset, we have been able to employ a YOLOv3 model [41] to solve the object-recognition problem. The

¹³ Due to the dimensionality of the dataset, the mutual exclusion of the main classes was computationally too expensive to compute in this case.

model has been trained for 100 epochs at each iteration with an AdamW optimizer, with a learning rate $\eta = 3 * 10^{-4}$ decreasing by 1/3 every 33 epochs. For both training and evaluation, the Input Over Union (IOU) threshold has been set to 0.5, the confidence threshold to 0.01 and the Non-Maximum Suppression (NMS) threshold to 0.5. We started training with $n = 1000$ labelled examples and by adding $p = 500$ samples for $q = 4$ iterations for a final budget of $b = 2000$ labelled examples. In Section 3.6, we reported the AUBC of the mean Average Precision (mAP) of the model averaged 10 times with Intersection over Union (IoU) ranging from 0.5 to 0.95. On *DOGvsPERSON* we considered a set of rules listing the parts belonging to the dog or the person, (e.g., **Person** \Rightarrow **PersonArm** \vee **PersonFoot** \vee **PersonHair** \vee **PersonHand** $\vee \dots$), the opposite rules implying the presence of the main object given the part (e.g., **PersonFoot** \Rightarrow **Person**). Also, we considered a disjunction of all the main classes and a disjunction of all the object-parts, for a total of 22 rules employed, as reported in Table 12.

Table 12. Domain knowledge on the *DOGvsPERSON* dataset.

$\forall x \text{ Dog_ear} \Rightarrow \text{Dog}$
$\forall x \text{ Dog_head} \Rightarrow \text{Dog}$
$\forall x \text{ Dog_leg} \Rightarrow \text{Dog}$
$\forall x \text{ Dog_muzzle} \Rightarrow \text{Dog}$
$\forall x \text{ Dog_neck} \Rightarrow \text{Dog}$
$\forall x \text{ Dog_nose} \Rightarrow \text{Dog}$
$\forall x \text{ Dog_paw} \Rightarrow \text{Dog}$
$\forall x \text{ Dog_tail} \Rightarrow \text{Dog}$
$\forall x \text{ Dog_torso} \Rightarrow \text{Dog}$
$\forall x \text{ Person_arm} \Rightarrow \text{Person}$
$\forall x \text{ Person_foot} \Rightarrow \text{Person}$
$\forall x \text{ Person_hair} \Rightarrow \text{Person}$
$\forall x \text{ Person_hand} \Rightarrow \text{Person}$
$\forall x \text{ Person_head} \Rightarrow \text{Person}$
$\forall x \text{ Person_leg} \Rightarrow \text{Person}$
$\forall x \text{ Person_neck} \Rightarrow \text{Person}$
$\forall x \text{ Person_nose} \Rightarrow \text{Person}$
$\forall x \text{ Person_torso} \Rightarrow \text{Person}$
$\forall x \text{ Dog} \Rightarrow \text{Dog_ear} \vee \text{Dog_head} \vee \text{Dog_leg} \vee \text{Dog_muzzle} \vee$ $\text{Dog_neck} \vee \text{Dog_nose} \vee \text{Dog_paw} \vee \text{Dog_tail} \vee \text{Dog_torso}$
$\forall x \text{ Person} \Rightarrow \text{Person_arm} \vee \text{Person_foot} \vee \text{Person_hair} \vee \text{Person_hand} \vee$ $\text{Person_head} \vee \text{Person_leg} \vee \text{Person_neck} \vee \text{Person_nose} \vee \text{Person_torso}$
$\forall x \text{ Dog} \vee \text{Person}$
$\forall x \text{ Dog_ear} \vee \text{Dog_head} \vee \text{Dog_leg} \vee \text{Dog_muzzle} \vee \text{Dog_neck} \vee \text{Dog_nose} \vee$ $\text{Dog_paw} \vee \text{Dog_tail} \vee \text{Dog_torso} \vee \text{Person_arm} \vee \text{Person_foot} \vee \text{Person_hair} \vee$ $\text{Person_hand} \vee \text{Person_head} \vee \text{Person_leg} \vee \text{Person_neck} \vee \text{Person_nose} \vee \text{Person_torso}$

In all experiments, we employed weight decay and low learning rate to avoid overfitting, rather than employing an early stopping strategy on a separate validation set. Indeed, we argue that it is not really realistic to rely on sufficiently

Table 13. Resume of the active losses $\mathcal{L}_a(f, x)$ maximized by each of the compared methods. Variants of the same methods (e.g., employing Monte Carlo Dropouts) have not been reported for the sake of brevity. We indicate with H the Entropy, with σ the softmax activation, with $p_f(y_i, x)$ the probability of associated to the class i , with M the set of main classes, with ϵ the adversarial perturbation, $g(\cdot)$ the model used to learn the model loss, with Z_{CENTER} the CoreSet points and with Z_{MEANS} the centroids of K-Means.

Entropy	$\mathcal{L}_a = H[f(x) x] = -(\sigma(f_M(x)) \cdot \log(\sigma(f_M(x))))$
Margin	$\mathcal{L}_a = -(p_f(\hat{y}_1 x) - p_f(\hat{y}_2 x)), \quad \hat{y}_1, \hat{y}_2 = \arg \max_{i \in M} p_f(y_i, x)$
LeastConf	$\mathcal{L}_a = 1 - p_f(\hat{y} x), \quad \hat{y} = \arg \max_{i \in M} p_f(y_i, x)$
BALD	$\mathcal{L}_a = H[f(x) x] - \mathbb{E}[H[f(x) x, \omega]]$
ADV	$\mathcal{L}_a = \ \epsilon\ , \quad \epsilon: f_M(x + \epsilon) \neq f_M(x)$
SupLoss	$\mathcal{L}_a = \mathcal{L}(f(x), y) \sim g(f(x))$
KCENTER	$\mathcal{L}_a = \ x - z_k\ , \quad z_k = \arg \min_{z \in Z_{CENTER}} \ x - z\ $
KMEANS	$\mathcal{L}_a = -\ x - z_k\ , \quad z_k = \arg \min_{z \in Z_{MEANS}} \ x - z\ $
KAL	$\mathcal{L}_a = \sum_{k \in \mathcal{K}} \varphi_k(f(x))$

large validation sets in an active learning scenario, where the amount of labels is scarce, and we try to minimize it as much as possible. Also, rather than re-training the network from scratch at each iteration, we choose to keep training it to save computational time.

A.4 Compared method details

In the experiments, several methods have been evaluated, comparing the performances both in terms of F1 score improvement and in terms of selection time. Compared techniques mostly follow two different active learning philosophies: uncertainty sample selection, which aim at estimating prediction uncertainty in deep neural networks; and diversity selection, which aims at maximally covering the input data distribution. Several standard strategies like **Entropy** [49], **Margin** [38] and **LeastConf** [54] belong to the first group, as well as some more recent methods like **BALD** [17], **ADV_{DEEPFOOL}** [16], **ADV_{BIM}** [60] and **SupLoss** [57]. On the contrary, **KMeans** [61] and **KCenter** [47], aim at reaching the highest diversity among selected samples.

In Table 13, we reported the loss $\mathcal{L}_a(f, x)$ used by each method to select active samples as following:

$$x^* = \arg \max_{x \in X_U} \mathcal{L}_a(f, x) \quad (3)$$

As the name implies, the **Entropy** method aims at measuring the uncertainty of a prediction by means of its entropy. With σ we indicate the softmax activation, that we applied on top¹⁴ of f_M to obtain a probability distribution – i.e.,

¹⁴ To avoid numerical issues, we applied it on the logits of the output of the network $\log(f(x)) - \log(1 - f(x))$.

$\sum_{i \in M} \sigma(f_i, x) = 1$. – as required to compute the Entropy. To adapt this method to the multi-label context, we restricted the computation to the M main classes (mutually exclusive) of f . Similarly, in **Margin** and **LeastConf** we compute again the softmax over the main classes f_M to obtain for each main class i the probability $p_f(y_i, x)$ of being active. In the first case, the margin between the two classes associated to the highest probability is used to estimate the uncertainty of the prediction. In the second case, the inverse of the probability associated to the most probable class is used instead. **BALD** employs as acquisition function the mutual information between the model predictions and the model parameters. The first term of the loss is again the Entropy of the model predictions; the second term of the equation, instead, represents the expected value of the entropy over the posteriori of the model parameters. Basically, the overall value is high when the predictions differ a lot, but the confidence of each prediction is rather high. Both **ADV_{DEEPFOOL}** and **ADV_{BIM}** use as a metric of uncertainty the norm of ϵ , the minimum input alteration allowing to change the prediction of the network. The difference in the two methods relies on the way ϵ is computed (i.e., on the adversarial attack employed). The stopping criterion to find ϵ in both methods consists in finding a perturbation that induce the network to predict a different class. To adapt these methods to the multi-label context, therefore, we restricted again the classes to the main ones f_M . The **SupLoss** strategy, instead, aims at approximating the supervision loss by means of a model $g(f(x))$. It receives in input the output of the f network (as well as the activation of the last hidden layers) and is trained to mimic the actual loss $\mathcal{L}(f(x), y) \sim g(f(x))$ on the supervised data. As [57] claim that $\mathcal{L}(f(x), y)$ is the actual upper bound of their method, for simplicity we employed the same \mathcal{L} to select uncertain samples. However, as we have seen in Section 3.1, even in the best case scenario this strategy does not work very well in complex problem as it mostly end up selecting outliers and making network convergence more difficult. Both **KCENTER** and **KMEANS**, instead, base their selection criteria only on distance metric on the input data distribution¹⁵. **KCENTER** aims at covering as much as possible the input data distribution, by selecting at each iteration the furthest sample to the current set of labelled samples (Z_{CENTER}). On the contrary, **KMEANS** strategy selects the closest unlabelled sample to the set of centroids (Z_{MEANS}), following a curriculum-learning strategy.

A.5 Ablation study: the amount of knowledge is proportional to performance improvement

We report here in Table 14, the complete version of Table 2 when considering all datasets. The take-home message that we reported in the main paper for the *CUB200* dataset is valid for all datasets: the amount of knowledge is proportional to the performance improvement. In the XOR dataset, since we only had the XOR rule, we only reported the performance with 0 % and 100 % knowledge.

¹⁵ In the image recognition tasks, we employed the latent distribution extracted by the convolutional features, as commonly done [42,60].

Strategy	XOR	IRIS	Animals	CUB200
KAL 00 %	93.54 \pm 0.83	81.45 \pm 18.88	53.42 \pm 1.64	50.13 \pm 0.36
KAL 25 %	–	91.72 \pm 5.03	53.21 \pm 1.39	50.22 \pm 0.40
KAL 50 %	–	93.12 \pm 4.72	54.54 \pm 1.11	50.19 \pm 0.42
KAL 75 %	–	93.84 \pm 3.94	54.61 \pm 0.47	51.28 \pm 0.29
KAL 100 %	97.73 \pm 0.83	93.60 \pm 3.87	56.15 \pm 1.06	51.98 \pm 0.31
LeastConf	89.97 \pm 10.55	84.35 \pm 13.89	53.33 \pm 2.12	50.20 \pm 0.43
Random	96.18 \pm 0.55	92.04 \pm 5.14	52.62 \pm 0.49	50.29 \pm 0.38

Table 14. Complete version of Table 2. The amount of knowledge is proportional to performance improvement. In the XOR dataset, since we only used the XOR rule, we only report the version with the rule included (KAL 100 %) and without (KAL 0 %)

We recall that in the 0 % knowledge scenario (to still evaluate our method) we always retain the uncertainty-like rule.

Strategy	XOR	Iris	Animals	CUB200
KAL	97.38 \pm 0.40	94.01 \pm 4.06	53.86 \pm 1.79	51.94 \pm 0.43
KAL Unc	97.46 \pm 0.40	93.97 \pm 3.61	54.12 \pm 0.97	51.98 \pm 0.35
KAL Div	97.62 \pm 0.79	93.44 \pm 4.50	54.56 \pm 1.09	51.97 \pm 0.39
KAL Div Unc	97.73 \pm 0.83	93.60 \pm 3.87	56.15 \pm 1.04	51.98 \pm 0.35
KAL _D	97.45 \pm 0.49	94.11 \pm 4.23	53.67 \pm 1.64	51.93 \pm 0.47
KAL _D Div	97.96 \pm 0.79	93.32 \pm 4.14	53.74 \pm 1.42	51.90 \pm 0.44
KAL _D Unc	97.57 \pm 0.51	93.51 \pm 3.80	53.92 \pm 1.28	52.02 \pm 0.33
KAL _D Div Unc	97.96 \pm 0.79	93.32 \pm 4.13	55.56 \pm 1.21	52.10 \pm 0.24

Table 15. Ablation study on diverse rule selection (**Div**) and on the uncertainty-like rule (**Unc**), both on the standard KAL and on Monte Carlo dropout version (KAL_D). In bold, the two best results. Notice how, the latter increase the AUBC by 2.3% and 1.9% respectively over KAL and KAL_D, on the Animals dataset.

A.6 Ablation study: selecting diverse constraint violations and employing an uncertainty-like rule improves the performances

In this section, we analyse the role of selecting samples violating a diverse set of rules and of employing the uncertainty-like rule introduced in Sec. 3.2. We recall that diversity sample selection is achieved by requiring a maximum of $r = p/2$ samples violating a certain rule k when selecting a new batch of p samples to be labelled. The uncertainty-like rule, instead, consists in requiring each predicate \mathbf{f}_i to be either true or false $\bigwedge_i \mathbf{f}_i \oplus \neg \mathbf{f}_i$. This way, in case many predicates for a certain sample have a not well-defined value (e.g., $f_i(x) = 0.5$), the violation of the constraint associated to this rule will be high.

We studied three different scenarios, employing the uncertainty-like rule (Unc), requiring a set of sample violating different rules (Div) or both of them (Div Unc). We compared them with the plain versions both in the case of estimating predictions with a standard classifier (KAL) and when using Monte Carlo dropout (KAL_D), leading to a total set of eight configurations. We report the results of these ablation study in Table 15. The results reported as KAL in Section 3.1 are, actually, the version with both features (KAL Div Unc). Indeed, KAL Div Unc and KAL_D Div Unc respectively increase the AUBC by 2.3% and 1.9% on the Animals dataset over KAL and KAL_D; on the *XOR-like* and on the *CUB200* dataset, the increase is as well important, although smaller than in the latter case; only in the *IRIS* case, the performances get reduced when requiring samples violating a diverse set of rules. This can be most likely explained by considering the size of the Iris dataset and the fact that we employed only 3 rules, specific for each class. In this setting, it may be more convenient for the network to gather samples related to a specific class (i.e., violating a specific rule) only. Indeed, labelling points related to a distribution already covered (i.e., where the knowledge is mostly respected) with very few examples available can slightly decrease the overall performances.

A.7 Training evolutions on the *XOR-like* problem

In Figure 5 we report further snapshots of the active selection process on the *XOR-like* dataset. They depict the model predictions similarly to Figure 3, but at different iterations and for all the compared methods. In this figure, it is even more clear that *no* uncertainty-based strategy is capable of discovering novel data distributions. More precisely, nor Entropy, nor Margin, nor LeastConf, nor BALD, nor Adv_{BIM}, nor Adv_{DEEPFOOL} can cover the data distribution in the right-bottom angle (for which no samples have been drawn during the initial random sampling) even when using all the labelling budget (20th iteration - rightmost figures). Diversity-based strategy, instead, can cover all data distribution; however, both KCENTER and even more KMEANS cannot perfectly predict samples along the decision boundaries. Finally, it is interesting to notice how the sampling selection performed by the SupLoss resemble the one made by KAL. However, even in the simplified upper case considered in this comparison, the selection process performed by SupLoss drives the network more slowly to convergences with respect to KAL. Indeed, as it can be noticed in the right-most plots, after 20 iterations the selection process in SupLoss has not covered very well yet the centre, differently from the selection process of KAL, which correctly covered this important zone.

A.8 Comparison of the knowledge violation

As introduced in Section 3.4, we tested whether the proposed method allows the model to learn to respect the knowledge provided by domain experts. In Table 16 we reported an extended version of Table 3, computing the increased percentage of violation of the \mathcal{K}_{CUB-S} knowledge by models trained following the

compared active strategies w.r.t. a model trained following the proposed strategy (KAL_{SMALL}). All the compared methods (including KALs when equipped with all the knowledge on CUB) violate the \mathcal{K}_{CUB-S} knowledge significantly more (4-13 times) than (KAL_{SMALL}).

Table 16. Violation of the \mathcal{K}_{CUB-S} knowledge computed as the increased percentage over the violation of a model trained to respect this knowledge (KAL_{small}). The proposed method ensures domain experts that their knowledge is acquired by the model on test data significantly more than using standard techniques.

Strategy	Increased Violation
KAL_{SMALL}	0.00 ± 38.09 %
ADV_{BIM}	591.03 ± 532.47 %
BALD	720.25 ± 251.86 %
Entropy	861.74 ± 519.35 %
$Entropy_D$	812.35 ± 359.23 %
KAL	863.53 ± 226.81 %
KAL_D	890.98 ± 299.79 %
KCENTER	530.13 ± 188.74 %
KMEANS	555.23 ± 327.80 %
LeastConf	1334.50 ± 697.47 %
$LeastConf_D$	934.46 ± 505.21 %
Margin	846.05 ± 341.52 %
$Margin_D$	637.26 ± 165.97 %
Random	483.10 ± 285.08 %
SupLoss	804.38 ± 76.31 %

A.9 Extracting the knowledge with an XAI method

Explainable AI techniques are more and more used in literature to mitigate the intrinsic opacity of deep neural networks. In Sec. 3.5 we employed [23] to explain the f model at each iteration, and use the explanations as the base knowledge (\mathcal{K}_{XAI}) of the proposed method (KAL_{XAI}). This is a viable solution to employ when no other knowledge is available. More precisely, we trained a decision tree h on the supervised training input data X_s to mimic the behaviour of the network f , i.e., we minimized the following loss $\mathcal{L}_{X_s}(h(x), f(x))$. From the decision tree, we extracted global explanations of each class in the form of, e.g., $\forall x \text{Setosa} \Leftrightarrow \neg \text{LongPetal}$. In computer vision tasks, however, decision trees are not suitable to be employed directly on the raw input data. For this reason, following [9], we trained the decision tree to mimic the behaviour of the model f over the main classes when receiving in inputs the attribute ones. The extracted rules, in this case, are of the type $\text{BlackfootedAlbatross} \Rightarrow \text{BillAll-purpose} \wedge \text{UnderpartsYellow} \wedge \dots$. To also explain the attribute classes, we did the reverse, i.e. we trained the decision tree to mimic the behaviour of the model

Table 17. Comparison of the methods in terms of the test mAP (%) AUBC [59] when increasing the number of labelled points on the object recognition task.

Dataset	Random	SupLoss	KAL
Dog vs Person	51.27 ± 1.41	55.30 ± 0.54	55.90 ± 0.39

f over the attribute classes and receiving in input the main ones. The extracted rules in this case are of the type $\text{BillCone} \Rightarrow \text{ParakeetAuklet} \wedge \text{IndigoBunting} \wedge \dots$. Finally, we always added to this set of rule a mutual exclusion rule over the main classes and a disjunction over the attributes (where available). Indeed, this notion must be always available as it conditions the choice of the training loss (e.g. binary vs standard cross entropy).

A.10 KAL in object recognition tasks

As introduced in Section 3.6, we tested the proposed method also in an object recognition context. As a proof of concept, we experimented on the simple DOGvsPERSON dataset, described in Appendix A.2. In Table 17, we report again the AUBC of the mean Average Precision curves when increasing the budget of labelled points. The same curves are reported in Figure 7. The results in this case are averaged over three different seed initialization of the network. The network increases its more performances when equipped with the KAL strategy (55.90) with respect to standard random sampling (51.41) but also compared to the SupLoss method (55.30). Interestingly, by considering Figure 7, we can appreciate how also in this context the proposed strategy significantly improves the performance of the network already after the very first iterations. Finally, we highlight again that the SupLoss performance reported are an upper bound of the performance of the method proposed in [57]. Particularly in this context, we believe that the object recognition supervision loss might not be easily learnt by the external model g . In this proof-of-concept, we only compared with Random and our implementation of the SupLoss method, since they were the most straightforwardly adaptable method to the object recognition context. In future work, we plan to compare with the adapted uncertainty-based strategy proposed in [25].

A.11 Experimental time comparison

In Table 18, we report the complete comparison of the time required by the different strategies. More precisely, we reported the increased percentage in computational time. We included both the time required to select the data and the time required to train the network. The KAL strategy does not come at the cost of a significantly increased time (with an average +0.3-23.4 %), comparably with standard uncertainty based strategies, Entropy (+0.0-32.9%), LeastConf

Table 18. Computational time required to select the samples to annotate by all the compared methods as percentage increase w.r.t the time required for a random sampling as defined in [60]. The lower, the better. NA indicates strategies computationally too expensive to compute on certain datasets. Notice how the proposed method is not computationally expensive, contrarily many recent active learning methods proposed in the literature.

	Strategy	XOR Time	Iris Time	Insurance (R) Time	Animals Time	CUB200 Time
KAL	5.22 ± 0.97	16.92 ± 3.55	20.52 ± 1.68	41.34 ± 17.42	180.22 ± 13.81	
KAL _D	15.79 ± 1.08	23.80 ± 3.24	31.43 ± 1.95	53.41 ± 21.36	197.05 ± 14.29	
ADV _{BIM}	36.93 ± 8.61	657.37 ± 108.47	—	5600.65 ± 1544.61	7440.67 ± 1287.83	
ADV _{DEEPPFOOL}	401.73 ± 76.53	6451.91 ± 496.10	—	57950.46 ± 23069.97	188435.28 ± 25119.23	
BALD	20.70 ± 1.70	21.50 ± 4.47	—	157.47 ± 12.69	613.48 ± 33.01	
KCENTER	31.82 ± 4.47	42.89 ± 9.24	158.28 ± 18.25	2379.57 ± 212.49	8713.37 ± 981.87	
KMEANS	7.90 ± 15.68	142.36 ± 369.98	28.75 ± 52.79	718.70 ± 4753.73	4724.60 ± 22871.24	
Entropy	4.00 ± 0.06	13.07 ± 8.76	—	14.58 ± 3.63	39.02 ± 2.63	
Entropy _D	14.51 ± 0.06	18.79 ± 2.84	—	22.83 ± 3.25	52.77 ± 2.96	
LeastConf	4.38 ± 0.18	12.31 ± 2.99	—	13.87 ± 2.52	33.32 ± 1.21	
LeastConf _D	14.96 ± 0.21	18.93 ± 2.88	—	22.46 ± 3.03	47.33 ± 1.37	
Margin	4.42 ± 0.19	12.42 ± 3.11	—	14.73 ± 3.84	40.97 ± 2.69	
Margin _D	14.97 ± 0.31	19.47 ± 3.26	—	23.34 ± 4.10	55.12 ± 2.97	
Random	1.00 ± 0.37	1.00 ± 2.44	1.00 ± 1.80	1.00 ± 3.07	1.00 ± 1.77	
SupLoss	4.57 ± 0.21	12.07 ± 5.59	17.69 ± 1.30	14.53 ± 3.35	38.60 ± 2.52	

(+0.0-2.01%), and Margin (+0.6-1.8%). Only the Monte Carlo Dropout versions, require higher computational time (KAL_D +5.4-88.3%, Entropy_D +0.0-89.2%, LeastConf_D +0.3-36.6%, Margin_D +1.7-84.0%) since they repeat several times a prediction with different switched-off neurons, to better assess uncertainty. On the contrary, BALD (+44.5-191.4%) and, more importantly, KMeans (+8.3-742.5%), KCenter (+2.0-5446.3%), ADV_{BIM} (+162.4-1353.6%), and ADV_{DEEPPFOOL} (+206.8-NA) strategies demand remarkable computational resources, strongly reducing the usability of the same methods. In particular, ADV_{DEEPPFOOL} requires a huge amount of resource, since the DEEPPFOOL attack is linearly dependent in both the selected samples and the predicted classes. For this reason, it was computationally infeasible to experimented it on the Animals and CUB datasets (respectively with 7 and 200 classes). At last, we reported SupLoss with a *, since the computational time required to compute the cross-entropy loss between the labels and the prediction (as simplified in this comparison), may be significantly different from the one required to predict the loss through the model g .

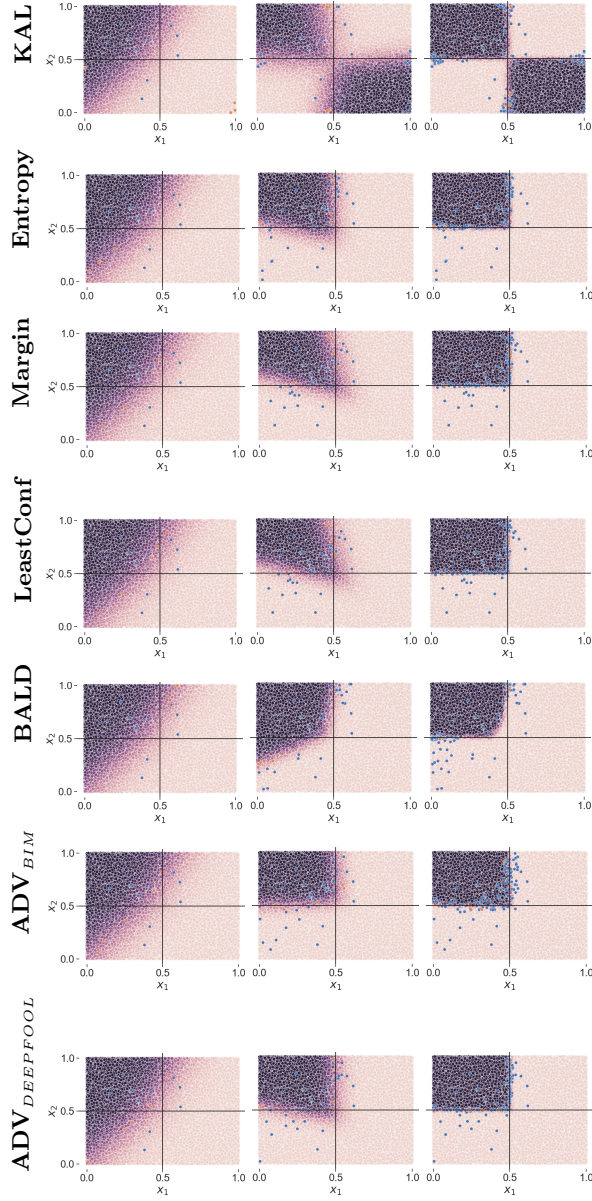


Fig. 5. An illustration of the selection process evolution for all the compared strategies on the *XOR-like* problem. We depict network predictions with different colour degrees (light colours negative predictions, dark colours positive prediction); in blue, samples selected in previous iterations, in orange those selected at the current iteration. Black lines at $x_1 = 0.5$ and $x_2 = 0.5$ are reported for visualization purposes only. From left to right, the state at the 1st, 5th, and 20th iteration.




Fig. 6. A visual example on the *XOR-like* problem, showing how the training evolves in each of the compared strategy (continued from previous page).




Fig. 7. Average test mAP performance growth on the object recognition task when increasing the number of labelled samples. Shaded areas indicate the 95 % confidence interval.

B Software

```

1 # Knowledge-drive Active Learning - Experiment on the XOR problem
2 tot_points = 10000
3 first_points = 10
4 n_points = 5
5 n_iterations = 198
6 seeds = range(5)
7 x = np.random.uniform(size=(tot_points, 2))
8 y = ((x[:, 0] > 0.5) & (x[:, 1] < 0.5)) |
9      ((x[:, 1] > 0.5) & (x[:, 0] < 0.5))
10 x_train, x_test, y_train, y_test = train_test_split(x, y)
11
12
13 # Defining constraints as product t-norm of the FOL rule expressing the XOR
14 def calculate_constraint_loss(x_continue, f):
15     # discrete_x = (x_continue > 0.5).float()
16     discrete_x = steep_sigmoid(x_continue).float()
17     x1 = discrete_x[:, 0]
18     x2 = discrete_x[:, 1]
19     c_loss1 = f * ((1 - (x1 * (1 - x2))) * (1 - (x2 * (1 - x1))))
20     c_loss2 = (1 - f) * (1 - (1 - (x1 * (1 - x2))) * (1 - (x2 * (1 - x1))))
21     return c_loss1 + c_loss2
22
23 # Constrained Active learning strategy
24 # We take the p elements that most violates the constraints and are among
25     available_idx
26 def kal_selection(labelled_idx, c_loss, n_p):
27     c_loss[torch.as_tensor(labelled_idx)] = -1
28     kal_idx = torch.argsort(c_loss, descending=True).tolist()[:n_p]
29     return kal_idx
30
31 net = MLP(2, 100)
32 accuracies = []
33 used_idx = randint(0, x_train.shape[0], first_points).tolist()
34 available_idx = [*range(tot_points)]
35 for n in range(n_iterations):
36     train_loop(net, x_train, y_train, used_idx)
37
38     with torch.no_grad():
39         preds_train = net(x_train).squeeze()
40         preds_test = net(x_test).squeeze()
41         accuracy = accuracy_score(preds_test, y_test)
42         cons_loss = calculate_constraint_loss(x_train, preds_train)
43
44     available_idx = list(set(available_idx) - set(used_idx))
45     active_idx = kal_selection(used_idx, cons_loss, n_points)
46     used_idx += active_idx

```

Listing 1.1. KAL code - Example on the XOR problem.

The Python code and the scripts used for the experiments, including full documentation, is freely available under Apache 2.0 Public Licence in a GitHub repository, and it is also provided in the supplementary material. The proposed approach only requires a few lines of code to train a model following the KAL strategy, as we sketch in the code example reported in Listing 1.1.