

Domain Knowledge Alleviates Adversarial Attacks in Multi-Label Classifiers

This is the peer reviewed version of the following article:

Original:

Melacci, S., Ciravegna, G., Sotgiu, A., Demontis, A., Biggio, B., Gori, M., et al. (2022). Domain Knowledge Alleviates Adversarial Attacks in Multi-Label Classifiers. IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, 44(12), 9944-9959 [10.1109/TPAMI.2021.3137564].

Availability:

This version is available <http://hdl.handle.net/11365/1206734> since 2023-04-27T09:53:52Z

Published:

DOI:10.1109/TPAMI.2021.3137564

Terms of use:

Open Access

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. Works made available under a Creative Commons license can be used according to the terms and conditions of said license.

For all terms of use and more information see the publisher's website.

(Article begins on next page)

Domain Knowledge Alleviates Adversarial Attacks in Multi-Label Classifiers

Stefano Melacci, *Member, IEEE*, Gabriele Ciravegna, Angelo Sotgiu, Ambra Demontis, *Member, IEEE*, Battista Biggio, *Member, IEEE*, Marco Gori, *Fellow, IEEE*, and Fabio Roli, *Fellow, IEEE*

Abstract—Adversarial attacks on machine learning-based classifiers, along with defense mechanisms, have been widely studied in the context of single-label classification problems. In this paper, we shift the attention to multi-label classification, where the availability of domain knowledge on the relationships among the considered classes may offer a natural way to spot incoherent predictions, i.e., predictions associated to adversarial examples lying outside of the training data distribution. We explore this intuition in a framework in which first-order logic knowledge is converted into constraints and injected into a semi-supervised learning problem. Within this setting, the constrained classifier learns to fulfill the domain knowledge over the marginal distribution, and can naturally reject samples with incoherent predictions. Even though our method does not exploit any knowledge of attacks during training, our experimental analysis surprisingly unveils that domain-knowledge constraints can help detect adversarial examples effectively, especially if such constraints are not known to the attacker. We show how to implement an adaptive attack exploiting knowledge of the constraints and, in a specifically-designed setting, we provide experimental comparisons with popular state-of-the-art attacks. We believe that our approach may provide a significant step towards designing more robust multi-label classifiers.

Index Terms—Learning from constraints, domain knowledge, adversarial machine learning, multi-label classification.

1 INTRODUCTION

Despite the impressive results reported in many different application domains, machine-learning algorithms have been shown to be easily misled by adversarial examples, i.e., input samples carefully perturbed to cause misclassifications at test time [1], [2]. In the last few years, a growing number of studies on the properties of adversarial attacks¹ and of the corresponding defenses have been produced by the scientific community [3], [4], [5], [6], [7] (see [8], [9] for an in-depth review on this topic). Most of the existing approaches work in the fully-supervised learning setting, either proposing methods for improving classifier robustness by modifying the learning algorithm to explicitly account for the presence of adversarial data perturbations [10], [11], [12], or developing specific detection mechanisms for adversarial examples [8], [13], [14], [15], [16], [17]. Those approaches are developed against a specific

class of attacks and usually are not robust against adversarial examples generated with different techniques [18]. Only a few approaches leverage also unlabeled data to improve adversarial robustness [19], [20], [21], [22], [23], [24], [25], [26], although the semi-supervised learning setting provides a natural scenario for real-world applications in which labeling data is costly while unlabeled samples are readily available. More importantly, the case of multi-label classification, in which each sample can belong to more classes, is only preliminarily discussed in the context of adversarial learning in [27], while using adversarial examples to improve the accuracy on legitimate (non-adversarial) samples of some multi-label classifiers is studied in [28], [29].

In this paper, we focus on multi-label classification and, in particular, in the case in which domain knowledge on the relationships among the considered classes is available. Such knowledge can be naturally expressed by First-Order Logic (FOL) clauses, and, following the learning framework of [30], [31], it can be used to improve the classifier by enforcing FOL-based constraints on the unsupervised or partially labeled portions of the training set. A well-known intuition in adversarial machine learning suggests that a reliable model of the distribution of the data could be used to spot adversarial examples, being them not sampled from such distribution, but it is not a straightforward procedure [32]. We borrow such intuition and we intersect it with the idea that semi-supervised examples can help learn decision boundaries that better follow the marginal data distribution, coherently with the available knowledge [31], [33], and we investigate the role of such knowledge in the context of data generated in an adversarial manner. While the generic idea of considering domain information in adversarial attacks has been recently followed by other authors to different extents [34], [35], [36], to the best of

- Stefano Melacci and Marco Gori are with the Department of Information Engineering and Mathematics, University of Siena, Siena, Italy. Marco Gori is also with MAASAI, Université Côte d’Azur, Nice, France. Gabriele Ciravegna is with the Department of Information Engineering, University of Florence, Florence, Italy. Angelo Sotgiu, Ambra Demontis, Battista Biggio, and Fabio Roli are with the Department of Information and Electric Engineering, University of Cagliari, Cagliari, Italy. E-mail: mela@diism.unisi.it, gabriele.ciravegna@unifi.it, {angelo.sotgiu, ambra.demontis, battista.biggio}@unica.it, marco.gori@unisi.it, roli@unica.it

Accepted for publication in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, DOI: 10.1109/TPAMI.2021.3137564. © 20XX IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

1. Even though every attack is adversarial by definition, we use the term *adversarial attacks* here to refer to the attack algorithms used to generate adversarial examples against machine-learning models.

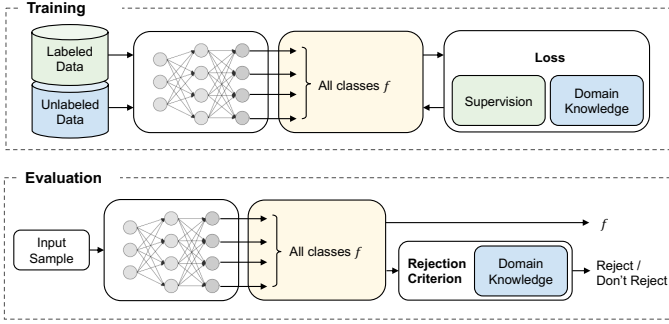


Fig. 1: Leveraging domain knowledge to improve robustness of multi-label classifiers. At training time, domain knowledge is used to enforce constraints on the learning process using unlabeled or partially-labeled data. At evaluation time, domain-knowledge constraints are used to detect and reject samples outside of the training data distribution.

our knowledge we are the first to use domain knowledge expressed in FOL and converted into polynomial constraints to improve adversarial robustness of multi-label classifiers.

In detail, this paper contributes in showing that domain knowledge is a powerful feature (i) to improve robustness of multi-label classifiers, and (ii) to help detect adversarial examples. The underlying idea of our approach is conceptually represented in Fig. 1. At training time, domain-knowledge constraints are enforced on the unlabeled (or partially-labeled) data to learn decision boundaries which better align with the marginal distributions. At test time, the same constraints can be efficiently evaluated on the test samples to identify and reject incoherent predictions, ideally outside of the training data distribution, potentially including adversarial examples. Our approach can be also used in single-classification tasks where domain knowledge and auxiliary classes are present, and can be exploited internally by the classifier to implement the rejection mechanism based on domain-knowledge constraints. We will show some concrete examples of this latter setting in our experiments, reporting comparisons with state-of-the-art adversarial attacks and concurrent defenses developed for single-classification tasks. To properly evaluate the robustness of our approach, which remains one of the most challenging problems in adversarial machine learning [9], [37], [38], we propose a novel multi-label attack that can implement both black-box and white-box adaptive attacks, being driven by the domain knowledge in the latter case. While we show that an adaptive attack having access to the domain knowledge exploited by our classifier can bypass it, even though at the cost of an increased perturbation size, it remains an open issue to understand how hard for an attacker would be to infer such knowledge in practical cases. For this reason, we believe that our work can provide a significant contribution towards both evaluating and designing robust multi-label classifiers.

This paper is organized as follows. Section 2 introduces the notion of learning with domain knowledge, emphasizing its effects in the input space. Section 3 shows how domain knowledge can be used to defend against adversarial attacks, together with a knowledge-aware attack procedure. A detailed experimental analysis is reported in

Section 4, evaluating the quality of our defense mechanisms, also considering state-of-the-art attacks and existing defense schemes. Related work is discussed in Section 5 and, finally, conclusions are drawn in Section 6.

2 LEARNING WITH DOMAIN KNOWLEDGE

In the context of this paper we focus on multi-label classification problems with c classes, in which each input $x \in \mathcal{X}$ is associated to one or more classes. We consider the case in which additional *domain knowledge* is available for problem at hand, represented by a set of relationships that are known to exist among (a subset of) the c classes. Exploiting such knowledge when training the classifier is the main subject of this section, and it has been shown to improve the generalization skills of the model [30], [31], [39]. The introduction of domain knowledge in the learning process provides precious information only when the training data are not fully labeled, as in the classic semi-supervised framework. Some examples might be partially labeled (i.e., for each data points a subset of the c classes participates to the ground truth) or a portion of the training set might be unsupervised. Of course, if the data are fully labeled then all the class relationships are already encoded in the supervision signal. However, in this paper we also consider domain knowledge as a mean to define a criterion that can spot potentially adversarial examples at test time, as we will discuss in Section 3, and that is also feasible in fully-supervised learning problems.

Notation. Formally, we consider a vector function $f : \mathcal{X} \rightarrow \mathbb{R}^c$, where $f = [f_1, \dots, f_c]$ and $\mathcal{X} \subseteq \mathbb{R}^d$. Each function f_i is responsible for implementing a specific task on the input domain \mathcal{X} .² In a classification problem, function f_i predicts the membership degree of x to the i -th class. Moreover, when we restrict the output of f_i to $[0, 1]$, we can think of f_i as the fuzzy logic predicate that models the truth degree of belonging to class i . In order to simplify the notation, we will frequently make no explicit distinctions between function names, predicate names, class names or between input samples and predicate variables. Whenever we focus on the predicate-oriented interpretation of each f_i , First-Order Logic (FOL) becomes the natural way of describing relationships among the classes, i.e., the most effective type of domain knowledge that could be eventually available in a multi-label problem; e.g., $\forall x \in \mathcal{X}, f_v(x) \wedge f_z(x) \Rightarrow f_u(x)$, for some v, z, u , meaning that the intersection between the v -th class and the z -th class is always included in the u -one. Table 1 reports a summary of the notation described so far and of the main symbols that will be introduced in what follows.

Learning from Constraints. The framework of Learning from Constraints [30], [31], [39] follows the idea of converting domain knowledge into constraints on the learning problem and it studies, amongst a variety of other knowledge-oriented constraints (see, e.g., Table 2 in [30]), the process of handling FOL formulas so that they can be both injected into the learning problem or used as a knowledge verification measure [31], [39]. Such knowledge

2. This notion can be trivially extended to the case in which the task functions operate in different domains.

TABLE 1: List of the main symbols and notations.

Notation	Description
$x \in \mathcal{X} \subseteq \mathbb{R}^d$	Input sample.
c	Number of classes.
$f_i(x)$	Membership score to the i -th class/fuzzy logic predicate.
$\text{NAME}(x)$	Same as f_i , using the name of the class instead of f_i .
$f = [f_1 \dots f_c]$	Vector function collecting all the f_i 's.
$f(\cdot, W)$	Network that implements f , with weights collected in W .
\mathcal{K}	Domain knowledge (FOL formulas).
ℓ	Number of formulas in \mathcal{K} .
$\wedge, \vee, \neg, \Rightarrow$	Logical connectives.
$\phi(f(x)) = 1$	T-Norm constraint from a generic FOL formula (in x).
$\hat{\phi}(f(x))$	Penalty (≥ 0) associated to constraint $\phi(f(x)) = 1$.
$\phi_h(f(x))$	Penalty (≥ 0) associated to the T-Norm-based constraint from the h -th formula in \mathcal{K} .
$\varphi(f, \mathcal{Z}, \mathcal{K}, \mu)$	Average penalty associated to the violations of the formulas in \mathcal{K} weighed by the scalars in μ , and evaluated on data \mathcal{Z} .
$\mathcal{L}, \mathcal{V}, \mathcal{T}$	Training, validation, and test data (respectively).
\mathcal{S}	Supervision attached to (some) of the data in \mathcal{L} .
suploss	Loss function, supervised data only.
\bar{x}	Test sample (unused in training).
$\mu^{\mathcal{L}}, \mu^{\mathcal{T}}$	Weights of the FOL formulas when φ is evaluated on training or test data, respectively.
λ	Scalar (≥ 0) that weighs φ in the learning criterion.
Ω	Rejection criterion.
τ	Rejection threshold (> 0).
ϵ	Upper bound on the norm of the difference between a clean example and its perturbed instance.
$l_p(x)$	Logits ($> -\kappa$) of the positive class with the tiniest output.
$l_n(x)$	Logits ($< \kappa$) of the negative class with the largest output.
κ	Threshold (≥ 0) used to bound the logits.
α	Scalar (≥ 0) that weighs φ in our white-box attack.

is enforced on those training examples for which either no information or only partial/incomplete labeling is available, thus casting the learning problem in the semi-supervised setting. As a result, the multi-label classifier can improve its performance and make predictions on out-of-sample data that are more coherent with the domain knowledge (see, e.g., Table 4 in [30]). In particular, FOL formulas that represent the domain knowledge of the considered problem are converted into numerical constraints using Triangular Norms (T-Norms, [40]), binary functions that generalize the conjunction operator \wedge . Following the previous example, $f_v(x) \wedge f_z(x) \Rightarrow f_u(x)$ is converted into a bilateral constraint $\phi(f(x)) = 1$ that, in the case of the product T-Norm, is $1 - f_v(x)f_z(x)(1 - f_u(x)) = 1$. The 1 on the right-hand side of the constraint is due to the fact that the numerical formula must hold true (i.e., 1), while the left-hand side is in $[0, 1]$. We indicate with $\hat{\phi}(f(x))$ the loss function associated to $\phi(f(x))$. In the simplest case (the one followed in this paper) such loss is $\hat{\phi}(f(x)) = 1 - \phi(f(x))$, where the minimum value of $\hat{\phi}(f(x))$ is zero. The quantifier $\forall x \in \mathcal{X}$ is translated by enforcing the constraints on a discrete data sample $\mathcal{Z} \subset \mathcal{X}$. The loss function $\varphi(f, \mathcal{Z})$ associated to all the available FOL formulas \mathcal{K} is obtained by aggregating the losses of all the corresponding constraints and averaging over the data in \mathcal{Z} . Since we usually have $\ell > 1$ formulas whose relative importance could be uneven, we get

$$\varphi(f, \mathcal{Z}, \mathcal{K}, \mu) = \frac{1}{|\mathcal{Z}|} \sum_{j=1}^{|\mathcal{Z}|} \left(\sum_{h=1}^{\ell} \mu_h \hat{\phi}_h(f(x_j)) \right) \in [0, \gamma] \quad (1)$$

where μ is the vector that collects the scalar weights $\mu_h > 0$ of the FOL formulas, and $\gamma = \sum_{h=1}^{\ell} \mu_h$.

In this paper, f is implemented with a neural architecture with c output units and weights collected in W . We distinguish between the use of Eq. (1) as a loss function in the training stage and its use as a measure to evaluate the constraint fulfillment on out-of-sample data. In detail, the classifier is trained on the training set \mathcal{L} by minimizing

$$\min_W \left[\text{suploss}(f(\cdot, W), \mathcal{L}, \mathcal{S}) + \lambda \cdot \varphi(f(\cdot, W), \mathcal{K}, \mu^{\mathcal{L}}) \right] \quad (2)$$

where $\mu^{\mathcal{L}}$ is the importance of the FOL formulas at training time, and $\lambda > 0$ modulates the weight of the constraint loss with respect to the supervision loss suploss , being \mathcal{S} the supervision information attached to some of the data in \mathcal{L} . The optimal λ is chosen by cross-validation, maximizing the classifier performance. When the classifier is evaluated on a test sample \bar{x} , the measure

$$\varphi(f(\cdot, W), \{\bar{x}\}, \mathcal{K}, \mu^{\mathcal{T}}) \in [0, \gamma^{\mathcal{T}}], \quad (3)$$

with weights $\mu^{\mathcal{T}}$ and $\gamma^{\mathcal{T}} = \sum_{h=1}^{\ell} \mu_h^{\mathcal{T}}$, returns a score that indicates the fulfillment of the domain knowledge on \bar{x} (the lower the better). Note that $\mu^{\mathcal{L}}$ and $\mu^{\mathcal{T}}$ might not necessarily be equivalent, even if certainly related. In particular, one may differently weigh the importance of some formulas during training to better accommodate the gradient-descent procedure and avoid bad local minima.

It is important to notice that Eq. (2) enforces domain knowledge only on the training data \mathcal{L} . There are no guarantees that such knowledge will be fulfilled in the whole input space \mathcal{X} . This suggests that optimizing Eq. (2) yields a stronger fulfillment of knowledge \mathcal{K} over the space regions where the training points are distributed (low values of φ), while φ could return larger values when departing from the distribution of the training data. The constraint enforcement is soft, so that the second term in Eq. (2) is not necessarily zero at the end of the optimization.

3 EXPLOITING DOMAIN KNOWLEDGE AGAINST ADVERSARIAL ATTACKS

The basic idea behind this paper is that the constraint loss of Eq. (1) is not only useful to enforce domain knowledge into the learning problem, but also (i) to gain some robustness with respect to adversarial attacks and (ii) as a tool to detect adversarial examples at no additional training cost, that are the main directions of this paper, as anticipated in Section 1.

A Paradigmatic Example. The example in Fig. 2 illustrates the main principles followed in this work, in a multi-label classification problem with 4 classes (cat, animal, motorbike, vehicle) for which the following domain knowledge is available, together with labeled and unlabeled training data:

$$\forall x, \text{CAT}(x) \Rightarrow \text{ANIMAL}(x), \quad (4)$$

$$\forall x, \text{MOTORBIKE}(x) \Rightarrow \text{VEHICLE}(x), \quad (5)$$

$$\forall x, \text{VEHICLE}(x) \Rightarrow \neg \text{ANIMAL}(x), \quad (6)$$

$$\forall x, \text{CAT}(x) \vee \text{ANIMAL}(x) \vee \text{MOTORBIKE}(x) \vee \text{VEHICLE}(x). \quad (7)$$

Such knowledge is converted into numerical constraints, as described in Section 2, while the loss function φ is enforced on the training data predictions during classifier training (Eq. 2). Fig. 2 shows two examples of the learned classifier.

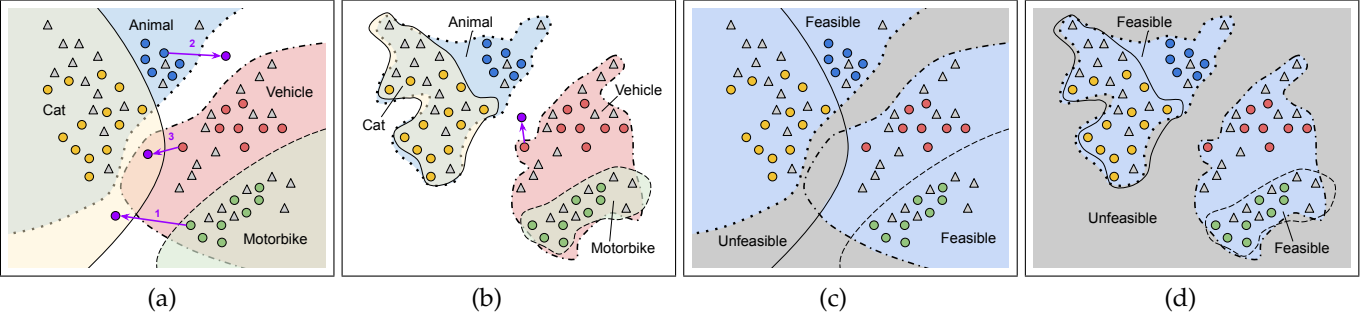


Fig. 2: Toy example using the domain knowledge of Eqs. (4-7) on 4 classes: cat (yellow), animal (blue), motorbike (green), vehicle (red). Labeled/unlabeled training data are depicted with rounded dots/gray triangles. (a,b) The decision regions for each class are shown in two sample outcomes of the training procedure: (a) open/loose decision boundaries; (b) tight/closed decision boundaries. The white area is associated with no predictions. Some adversarial examples (purple arrows/dots) are detected as they end up in regions that violate the constraints. Moreover, in (c,d) The feasible/unfeasible regions (blue/gray) that fulfill/violate the constraints for (a,b) are shown. Decision boundaries of the classes in (a,b) are also depicted in (c,d).

Considering point (i), in both cases, the decision boundaries are altered on the unlabeled data, enforcing the classifier to take a knowledge-coherent decision over the unlabeled training points and to better cover the marginal distribution of the data. This knowledge-driven regularity improves classifier robustness to adversarial attacks, as we will discuss in Section 4. Going into further details to illustrate claim (ii), in (a) we have the most likely case, in which decision boundaries are not always perfectly tight to the data distribution, and they might be not closed (ReLU networks typically return high-confidence predictions far from the training data [41]). Three different attacks are shown (purple). In attack 1, an example of motorbike is perturbed to become an element of the cat class, but Eq. (4) is not fulfilled anymore. In attack 2, an example of animal is attacked to avoid being predicted as animal. However, it falls in a region where no predictions are yielded, violating Eq. (7). Attack number 3 consists of an adversarial attack to create a fake cat that, however, is also predicted as vehicle, thus violating Eq. (4) and Eq. (6). In (b) we have an ideal and extreme case, with very tight and closed decision boundaries. Some classes are well separated, it is harder to generate adversarial examples by slightly perturbing the available data, while it is easy to fall in regions for which Eq. (7) is not fulfilled. The pictures in (c-d) show the unfeasible regions in which the constraint loss φ is significantly larger, thus offering a natural criterion to spot adversarial examples that fall outside of the training data distribution.

Domain Knowledge-based Rejection. Following these intuitions, and motivated by the approach of [42], [43], we define a rejection criterion Ω as the Boolean expression

$$\Omega(\bar{x}, \tau | f(\cdot, W), \mathcal{K}, \mu^T) = \varphi(f(\cdot, W), \{\bar{x}\}, \mathcal{K}, \mu^T) > \tau \quad (8)$$

where $\tau > 0$ is estimated by cross-validation in order to avoid rejecting (or rejecting a small number of³) the examples in the validation set \mathcal{V} . Eq. (8) evaluates the constraint loss on the validation data \mathcal{V} , using the importance weights μ^T (that we will discuss in what follows), as in Eq. (3). The rationale behind this idea is that those samples for which the

constraint loss is larger than what it is on the distribution of the data that are available when training/tuning the classifier, should be rejected. The training samples are the ones over which domain knowledge was enforced during the training stage, while the validation set represents data on which knowledge was not enforced, but that are sampled from the same distribution from which the training set is sampled, making them good candidates for estimating τ . Notice that Ω is measured at test time on an already trained classifier, and it can be used independently on the nature of the training data (fully or partially/semi-supervised). Differently from ad-hoc detectors, that usually require to train generative models, this rejection procedure comes at no additional training cost.⁴

Pairing Effect. The procedure is effective whenever the functions in f are not too strongly paired with respect to \mathcal{K} , and we formalize the notion of “pairing” as follows.

Definition 3.1. Pairing. We consider a classification problem whose training data are distributed accordingly to the probability density $p(x)$. Given \mathcal{K} and μ^T , the functions in f are strongly paired whenever $\zeta(\mathcal{H}, \mathcal{L}) = \|\varphi(f(\cdot, W), \mathcal{H}, \mathcal{K}, \mu^T) - \varphi(f(\cdot, W), \mathcal{L}, \mathcal{K}, \mu^T)\| \approx 0$, being \mathcal{H} a discrete set of samples uniformly distributed around the support of $p(x)$.

This notion indicates that if the constraint loss is fulfilled in similar ways over the training data distribution and space areas close to it, then there is no room for detecting those examples that should be rejected. While it is not straightforward to evaluate pairing before training the classifier, the soft constraining scheme of Eq. (2) allows the classification functions to be paired in a less strong manner that what they would be when using hard constraints.⁵ Note that a multi-label system is usually equipped with activation functions that do not structurally enforce any dependencies among the classes (e.g., differently from what happens with softmax), so it is naturally able to respond without assigning the

4. Generative models on the fulfillment of the single constraints could be considered too.

5. See [44] for a discussion on hard constraints and graphical models in an adversarial context.

3. 10% in our experiments.

input to any class (white areas in Fig. 2). This property has been recently discussed as a mean for gaining robustness to adversarial examples [6], [45]. The formula in Eq. (7) is what allows our model to spot examples that might fall in this “I don’t know” area. Dependencies among classes are only introduced by the constraint loss φ in Eq. (2) on the training data.

The choice of μ^T is crucial in the definition of the reject function Ω . On the one hand, in some problems we might have access to the certainty degree of each FOL formula, that could be used to set μ^T , otherwise it seems natural to select an unbiased set of weights μ^T , $\mu_h = 1, \forall h$. On the other hand, several FOL formulas involve the implication operator \Rightarrow , that naturally implements if-then rules (if class v then class z) or, equivalently, rules that are about hierarchies, since \Rightarrow models an inclusion (class v included in class z). However, whenever the premises are false, the whole formula holds true. It might be easy to trivially fulfill the associated constraints by zeroing all the predicates in the premises, eventually avoiding rejection. As rule of thumb, it is better to select μ_h ’s that are larger for those constraints that favor the activation of the involved predicates.

Single-label Classifiers. The type of domain knowledge described so far usually involves logic formulas that encode relationships among multiple classes, thus it is naturally associated with multi-label problems. Let us focus our attention on multi-label scenarios in which there exists a subset of categories that are known to be mutually exclusive, that we will refer to as *main classes*, while the remaining categories will be referred to as *auxiliary classes*. If we restrict the original classification problem to the main classes only, we basically end-up in a single-label scenario. Let us assume that the available logic formulas introduce relationships between (some of) the main classes and (some of) the auxiliary ones. As a result, in order to setup our defense mechanism (Eq. 8) or to learn with domain knowledge (Eq. 2), predictions on both the main and auxiliary classes must be available, so that the truth degree of the logic formulas can be evaluated. This consideration can be exploited to design classifiers that expose single label predictions on the main classes, thus acting as single-label classifiers, and include predictions on the auxiliary classes that are not exposed to the user at all, but that are internally used to setup our defense mechanism or to improve the quality of whole classifier when learning in a semi-supervised context. Formally, let us assume that the components $\{f_i, i = 1, \dots, c\}$ of the vector function f are partitioned into two disjoint subsets, where the first one considers the components about the mutually-exclusive main classes and the second subset is about the auxiliary classes. We define with f^v the vector function with the elements in the first subset, while f^h is the vector function based on the elements of the second one, as shown in Fig. 3. The system only exposes to the user predictions computed by means of f^v , while the computations of f^h are hidden. Overall, the system can still exploit domain knowledge that consists in relationships between the classes associated to f^v (main classes) and the ones associated to f^h (auxiliary classes), or among the ones in f^h only, thus leveraging the learning principles that were described in Section 2. Moreover, the system can exploit the hidden predictions




Fig. 3: Single-label classifier on a set of mutually exclusive classes (*main classes*), computing the class activations by f^v and exposing them to the user (red path). It internally computes by f^h additional predictions over *auxiliary classes* that are involved in the domain knowledge (together with the main classes). Training considers *all* the classes, Fig. 1.

and the available knowledge to implement the knowledge-based rejection mechanism that we proposed in this section, as sketched in Fig. 3.

Due to the single-label nature of the visible portion of the classifier, existing state-of-the-art attacks, specifically designed for single-label models, can be used to fool the classifier in a black-box scenario. In Section 4, when the consider data are compatible with this special setting, we will exploit recent attack procedures to generate adversarial examples and evaluate the proposed knowledge-based rejection mechanism. Of course, differently from what we previously stated about the real multi-label setting, we cannot consider the cost of the rejection mechanism negligible in this case, since the system must learn the functions in f^h in order to be able to compute the rejection criterion.

3.1 Attacking Multi-label Classifiers

Robustness against adversarial examples is typically evaluated against *black-box* and *white-box* attacks [8], [9]. In the black-box setting, the attacker is assumed to have only black-box query access to the target model, ignoring the presence of any defense mechanisms and without having access to any additional domain knowledge and related constraints. However, a surrogate model can be trained on data ideally sampled from the same distribution of that used to train the target model. Within these assumptions, gradient-based attacks can be optimized against the surrogate model, and then transferred/evaluated against the target one [3], [46]. In the white-box setting, instead, the attacker is assumed to know everything about the target model, including the defense mechanism. White-box attacks are thus expected to also exploit the available domain knowledge to try to bypass the knowledge-based defense.

The existing literature on the generation of adversarial examples is strongly focused on single-label classification problems (see [8] and references therein). In such context, the classifier is expected to take a decision that is only about one of the c classes, and, in a nutshell, attacking the classifier boils down to perturb the input in order to make the classifier predict a wrong class. The whole procedure is subject to constraints on the amount of perturbation that the system is allowed to apply. Formally, given $x \in \mathcal{T}$, being \mathcal{T} the test set, the attack generation procedures in single-label classification commonly solve the following problem,

$$\begin{aligned} x^* &= \arg \min_{x'} [-\text{suploss}(f(x', W), \mathcal{L}, \mathcal{S})], \\ \text{s.t. } \|x - x'\| &< \epsilon, \end{aligned} \quad (9)$$

being $\|\cdot\|$ an L_p -norm and $\epsilon > 0$. Each x has a unique class label/index attached to it and stored in \mathcal{S} , and suploss is usually the cross-entropy loss. Different attacks and optimization techniques for solving the problem of Eq. (9) have been proposed [47]. While there are no ambiguities on the class on which we want the classifier to reduce its confidence, i.e., the ground-truth (positive) class of the given input x , the class that the classifier will predict in input x^* might be given or not, thus each of the remaining $c - 1$ classes could be a valid option. When moving to the multi-label setting, each $x \in \mathcal{T}$ is associated to multiple ground-truth positive classes, collected in set P_x , and we indicate with N_x the set of ground-truth negative classes of x . Differently to the previous case, due to the lack of mutual-exclusivity of the predictions, creating an adversarial example out of x is more arbitrary. For example, the optimization procedure could focus on making the classifier not able to predict any of the classes in P_x , or a subset of them. Similarly, the optimization could focus on making the classifier positively predict one or more classes of N_x .

Departing from the overwhelming majority of existing attacks for single-label classifiers, we propose a multi-label attack that focuses on the classes on which the classifier is less confident (thus easier to attack), that are selected and re-defined during the optimization procedure in function of the way the predictions of the classifier progressively change. Of course, in the *black-box* case, this attack is not considering that classes are related, and it is not taking care that, perhaps, changing the prediction on a certain class should also trigger a coherent change in other related classes. Differently, in the *white-box* setting, the previously introduced domain knowledge and, in particular, the corresponding loss of Eq. (1) is what encodes such relationships in a differentiable way, so that we can easily exploit it when crafting attacks. We first introduce the proposed multi-label attack in a *black-box* setting, in which domain knowledge is not available. To make gradient computation numerically more robust, as in [48], we consider the activations (logits) of the last layer of f to compute the objective function, instead of using the cross-entropy loss. Let us define $p = \arg \min_i [f_i(x), i \in P_x]$, and $n = \arg \max_i [f_i(x), i \in N_x]$, i.e., p (n) is the index of the positive (negative) class with the smallest (largest) output score. These are essentially the indices of the classes for which x is closer to the decision boundaries. Our attack optimizes the following objective,

$$\begin{aligned} x^* &= \arg \min_{x'} [\max(l_p(x'), -\kappa) - \min(l_n(x'), \kappa)] \\ \text{s.t. } &\|x - x'\| < \epsilon, \end{aligned} \quad (10)$$

where l_j is the value of the logit of f_j , $\|\cdot\|$ is an L_p -norm (L_2 in our experiments), and in the case of image data with pixel intensities in $[0, 1]$ we also have $x' \in [0, 1]$. The scalar $\kappa \geq 0$ is used to threshold the values of the logits, to avoid increasing/decreasing them in an unbounded way (in our experiments, we set $\kappa = 2$). Optimizing the logit values is preferable to avoid sigmoid saturation. While the definition of Eq. (10) is limited to a pair of classes, we *dynamically* update p and n whenever logit l_p (l_n) goes beyond (above) the threshold $-\kappa$ (κ), thus multiple classes are considered by the attack, compatibly with the maximum number of iterations of the optimizer. This strategy resulted to be more

effective than jointly optimizing all the classes in P_x and N_x . Moreover, the classes involved in the attack can be a subset of the whole set, as in [27]. In a *white-box* scenario, when the attacker has the use of the domain knowledge, the information in \mathcal{K} provides a comprehensive description on how the predictions of the classifier should be altered over several classes in order to be coherent with the knowledge. In such a scenario, we enhance Eq. (10) to implement what we refer to as multi-label knowledge-driven adversarial attack (MKA), including the differentiable knowledge-driven loss φ in the objective function,

$$\begin{aligned} x^* &= \arg \min_{x'} [\max(l_p(x'), -\kappa) - \min(l_n(x'), \kappa) + \\ &\alpha \cdot \varphi(f, \{x'\}, \mathcal{K}, \mu^T)], \quad \text{s.t. } \|x - x'\| < \epsilon \end{aligned} \quad (11)$$

in which we set $\alpha > 0$ to enforce domain knowledge and avoid rejection. When crafting adversarial examples, MKA softly enforces the fulfillment of domain knowledge by means of the loss function φ . For *black-box* attacks, instead, we set $\alpha = 0$ to recover Eq. (10). MKA naturally extends the formulation of single-label attacks (when P_x is composed of a single class) and it allows staging both *black-box* and *white-box* (adaptive) attacks against our approach. Eq. (11) is minimized via projected gradient descent (1000 samples and 50 iterations in our experiments).

3.2 Impact of Domain Knowledge and Main Issues

Our approach is built around the idea of exploiting the available domain knowledge \mathcal{K} on the target classification problem, both in the cases of rejection and multi-label attack. Several existing works use additional knowledge on the learning problem with different goals, being it represented by logic [30], [31], [39], [49], inherited by knowledge graphs or other external resources [50], [51], and encoded in multiple ways to face specific tasks [50], [52], [53]. For instance, Semantic-based Regularization [31] and the theory formalized in [30] focus on the same approach we use here to convert generic FOL knowledge. On one hand, \mathcal{K} might not always be available, thus limiting the applicability of what we propose and of the other aforementioned approaches. On the other hand, \mathcal{K} is about relationships among classes that, in the case of the universal quantifier, hold $\forall x$. As a result, such knowledge is more generic than specific example-level supervisions. Human experts can produce FOL rules to a lesser effort than what is needed to manually label large batches of examples, since \mathcal{K} naturally represents the type of high-level knowledge on the target domain that a human would develop during a concrete experience on the considered task (e.g., *if A happens, then also B or C are triggered, but not D*). Moreover, we are currently working on methods to extract the type of knowledge that we consider in this paper by means of special neural architectures, with clear connections to Explainable AI [54], [55].

When the number ℓ of FOL formulas in \mathcal{K} is large, a larger number of penalty terms $\hat{\phi}_h$ will be considered in φ of Eq. (1). Of course, every approach that exploits additional knowledge usually incurs in increased complexity when the knowledge base is large [30], [31], [39], [49]. In our case, the T-Norm-based conversion does not represent an issue, since it is computed only once in a pre-processing stage, and, similarly, the output of the network $f(x, W)$ is

computed only once in order to evaluate φ for a certain sample x and for given weights W , independently on the size of \mathcal{K} . However, the computation of φ must be repeated at each iteration of the optimization of Eq. (2) or Eq. (11), and when evaluating whether an input should be rejected or not, Eq. (8). From the practical point of view, the computational complexity scales almost linearly with ℓ , but each $\hat{\phi}_h$ has a different structure depending on the FOL formula from which it was generated—roughly speaking, formulas involving more predicates usually yield more complex T-Norm-based polynomials. Several heuristic solutions are indeed possible to overcome these issues. For example, the knowledge base could be sub-selected in order to bound the number of rules in which each class is involved, or a stochastic optimization could be devised to sample the rules included in φ at each iteration of the optimization process. However, we remark that, in the experimental activities of this paper, none of the mentioned issues arose.

The way we convert FOL rules into polynomial constraints, described in Section 2, inherits the flexibility of logic in terms of knowledge representation capabilities. Of course, the concrete impact of \mathcal{K} in the rejection mechanisms or in MKA depends on the specific information that is encoded by the FOL rules. For instance, suppose that $f_i(x) = 1$ for a certain x . The formula $f_i(x) \Rightarrow f_v(x) \vee f_z(x) \vee \dots \vee f_u(x)$ is “more likely” to be fulfilled than the formula with an analogous structure in which \vee ’s are replaced by \wedge ’s. In the former, it is enough for a predicate in the conclusions to be 1, while in the latter, all the predicates of the conclusions must be jointly true. The rejection criterion or MKA are likely to be more effective in the latter case, but it cannot be strongly stated in advance, since it depends on the concrete way in which $f(\cdot, W)$ is developed by the learning procedure, as discussed in Section 3, and, in the case of MKA, on the difficulty in optimizing Eq. (11).

When restricting our attention to the rejection function of Eq. (8), a key element to the success of the proposed criterion is the choice of τ . In Section 3 we suggested using data in \mathcal{V} to tune τ , that is a valuable solution, but, of course, it strongly depends on the quality of \mathcal{V} , similarly to what happens when tuning other hyper-parameters. More generally, a too small τ will result in a reject-prone system that does not reject only those inputs that are strongly coherent with the domain knowledge. A too large τ would end up in not rejecting inputs, being them coherent with \mathcal{K} or not. If further information on the formulas in \mathcal{K} is available, such as their expected importance with respect to the considered task, one could avoid computing an averaged measure as φ , and evaluate the penalty term $\hat{\phi}_h$ of each single formula against its own reject threshold (i.e., multiple τ ’s), that might be selected accordingly to the importance of the formula itself (i.e., smaller τ ’s in more important formulas).

4 EXPERIMENTS

In this section, we report our experimental analysis, discussing the experimental setup in Section 4.1, and the results of standard and adversarial evaluations for multi-label classifiers in Section 4.2. We then show in Section 4.3 how our multi-label classifiers can also be adopted to mitigate the

TABLE 2: Datasets and details on the experimental setting. “Classes” reports the total number of categories, specifying the number of main classes in parentheses. The fraction of labeled ($\%L$) samples, the level of partial labeling ($\%P$), along with the number of training ($|\mathcal{L}|$), validation ($|\mathcal{V}|$), and test ($|\mathcal{T}|$) examples are also reported.

Dataset	Classes	%L	%P	$ \mathcal{L} $	$ \mathcal{V} $	$ \mathcal{T} $
ANIMALS	33 (7)	30%	90%	5808	1244	1243
CIFAR-100	120 (100)	30%	0%	40000	10000	10000
PASCAL-Part	64 (20)	30%	70%	7072	1515	1515

TABLE 3: Values of the hyperparameter λ selected via cross-validation in our experiments. Note that baseline models TL and FT do not exploit domain knowledge ($\lambda = 0$).

Model	ANIMALS	CIFAR-100	PASCAL-Part
TL+C	10^{-2}	3	10^{-1}
TL+CC	1	10	1
FT+C	10^{-2}	3	10^{-1}

impact of adversarial examples in single-label classification tasks, when auxiliary classes are exploited. This allows us to highlight that our approach exhibits competitive performances with respect to other baseline defense methods designed under the same assumptions (i.e., without assuming any specific knowledge of the attacks) and against state-of-the-art attacks that are developed for single-label classification tasks.

4.1 Experimental Settings

Datasets. We considered three image classification datasets, referred to as ANIMALS, CIFAR-100 and PASCAL-Part respectively. The first one is a collection of images of animals, taken from the ImageNet database,⁶ the second one is a popular benchmark composed of RGB images (32×32) belonging to different types of classes (vehicles, flowers, people, etc.),⁵ while the last dataset is composed of images in which both objects (Man, Dog, Car, Train, etc.) and object-parts (Head, Paw, Beak, etc.) are labeled.⁷ All datasets are used in a multi-label classification setting, so that the ground truth of each example is composed by a set of binary class labels. In the case of ANIMALS there are 33 categories, where the first 7 ones, also referred to as “main” classes, are about specific categories of animals (albatross, cheetah, tiger, giraffe, zebra, ostrich, penguin) while the other 25 classes are about more generic features (mammal, bird, carnivore, fly, etc.). The CIFAR-100 dataset is composed of 120 classes, out of which 100 are fine-grained (“main” classes) and 20 are superclasses. In the PASCAL-Part dataset, after having processed data as in [56], we are left with 64 categories, out of which 20 are objects (“main” classes) and the remaining 44 are object-parts. We have the use of domain knowledge that holds for all the available examples. In the case of ANIMALS, it is a collection of FOL formulas

6. ANIMALS <http://www.image-net.org/>, CIFAR-100 <https://www.cs.toronto.edu/~kriz/cifar.html>

7. PASCAL-Part: <https://www.cs.stanford.edu/~roozbeh/pascal-parts/pascal-parts.html>

TABLE 4: Values of the constraint loss φ on the test data \mathcal{T} .

Model	ANIMALS	CIFAR-100	PASCAL-Part
TL	0.5833 \pm 0.0316	1.4440 \pm 0.0087	2.7286 \pm 0.0853
TL+C	0.2134 \pm 0.0160	1.0406 \pm 0.0020	1.7422 \pm 0.0516
TL+CC	0.2004 \pm 0.0097	0.7267 \pm 0.0020	0.7387 \pm 0.0151
FT	0.3751 \pm 0.0169	0.9603 \pm 0.0041	2.4478 \pm 0.0723
FT+C	0.0897 \pm 0.0113	0.4449 \pm 0.0068	0.8434 \pm 0.0471

TABLE 5: Multi-label classification results in \mathcal{T} , for different models, averaged across different repetitions (standard deviations are $< 1\%$). The second row-block is restricted to the main classes (Accuracy or F1). See the main text for details.

Metric	Dataset	TL	TL+C	TL+CC	FT	FT+C
F1 (%)	ANIMALS	98.3	98.6	98.1	98.6	99.2
	CIFAR-100	52.0	55.1	53.1	59.3	64.0
	PASCAL-Part	69.5	70.0	69.4	69.1	71.0
AccMain (%) ¹	ANIMALS ¹	98.8	99.2	99.2	98.5	99.1
F1Main (%) ²	CIFAR-100 ¹	53.3	55.6	52.8	60.5	61.6
	PASCAL-Part ²	73.8	75.9	69.5	70.4	75.0

that were defined in the benchmark of P.H. Winston [57], and they involve relationships between animal classes and animal properties, such as $\forall x \text{ FLY}(x) \wedge \text{LAYEGGS}(x) \Rightarrow \text{BIRD}(x)$. In CIFAR-100, FOL formulas are about the father-son relationships between classes, while in PASCAL-Part they either list all the parts belonging to a certain object, i.e., $\text{MOTORBIKE}(x) \Rightarrow \text{WHEEL}(x) \vee \text{HEADLIGHT}(x) \vee \text{HANDLEBAR}(x) \vee \text{SADDLE}(x)$, or they list all the objects in which a part can be found, i.e., $\text{HANDLEBAR}(x) \Rightarrow \text{BICYCLE}(x) \vee \text{MOTORBIKE}(x)$. we also introduced a disjunction or a mutual-exclusivity constraint among the main classes, and another disjunction among the other classes. See Table 2 and the supplementary material for more details. Each dataset was divided into training and test sets (the latter indicated with \mathcal{T}). The training set was divided into a learning set (\mathcal{L}), used to train the classifiers, and a validation set (\mathcal{V}), used to tune the model parameters. We defined a semi-supervised learning scenario in which only a portion of the training set is labeled, sometimes partially (i.e., only a fraction of the binary labels of an example is known), as detailed in Table 2. We indicated with $\%L$ the percentage of labeled training data, and with $\%P$ the percentage of binary class labels that are unknown for each labeled example.⁸

Classifiers. We compared two neural architectures, based on the popular backbone ResNet50, trained using ImageNet data. In the first network, referred to as TL, we transferred the ResNet50 model and trained the last layer from scratch in order to predict the dataset-specific multiple classes (sigmoid activation). The second network, indicated with FT, has the same structure of TL, but we also fine-tuned the last convolutional layer. Each model is based on the product T-Norm, and it was trained for a number of epochs e that

we selected as follows: 1000 epochs in ANIMALS, 300 (TL) or 100 (FT) epochs in CIFAR-100, and 500 (TL) or 250 (FT) in PASCAL-Part, using minibatches of size 64. We used the Adam optimizer, with an initial step size of 10^{-5} , except for FT in CIFAR-100, for which we used 10^{-4} to speedup convergence. We selected the model at the epoch that led to the largest F1 in \mathcal{V} . We considered unconstrained ($\lambda = 0$) and knowledge-constrained ($\lambda > 0$) models. The latter are indicated with the +C (and +CC) suffix.

Evaluation Metrics. To evaluate performance, we considered the (macro) F1 score and a metric restricted to the main classes.⁹ For ANIMALS and CIFAR-100, the main classes are mutually exclusive, so we measured the accuracy in predicting the winning main class (*AccMain*), while in PASCAL-Part we kept the F1 score (*F1Main*) as multiple main classes can be predicted on the same input.

Hyperparameter Tuning. In Table 3 we report the optimal value of $\lambda \in \{10^{-2}, 10^{-1}, 1, 3, 5, 8, 10, 10^2\}$ for the TL+C and FT+C models used in our experiments, selected via a 3-fold cross-validation procedure. In the case of TL, we also considered a strongly-constrained (+CC) model with inferior performance but higher coherence (greater λ) among the predicted categories (that might lead to a worse fitting of the supervisions).¹⁰ Table 4 reports the value of the constraint loss φ measured on the test set \mathcal{T} . We used $\mu^{\mathcal{L}} = \mu^{\mathcal{T}}$, setting each component μ_h to 1, with the exception of the weight of the mutual exclusivity constraint or the disjunction of the main classes, which was set to 10 to enforce the classifier to take decisions.

4.2 Experimental Results on Multi-label Classifiers

We discuss here the main experiments related to the evaluation of the considered multi-label classifiers.

Standard Evaluation. In order to assess the behaviors of the classifiers in the considered datasets and the available domain knowledge, we compared classifiers that exploit domain knowledge with the ones that do not exploit it. The results of our evaluation are reported in Table 5, averaged over the 3 training-test splits. For each of them, 3 runs were considered, using different initialization of the weights. The introduction of domain knowledge allows the constrained classifiers to slightly outperform the unconstrained ones.

Adversarial Evaluation. To evaluate adversarial robustness, we used the MKA attack procedure described in Section 3. and we restricted the attack to work on the already introduced main classes, being them associated to the most important categories of each problem. In ANIMALS and CIFAR-100 we assumed the attacker to have access to the information on the mutual exclusivity of the main classes, so that p in Eq. (11) is not required to change during the attack optimization. We also set $\kappa = \infty$ to maximize confidence of misclassifications at each given perturbation bound ϵ . All the following results are averaged after having attacked twice the model obtained after each of the 3 training runs.

In the *black-box* setting, we assumed the attacker to be also aware of the network architecture of the target classifier, and attacks were generated from a surrogate model trained

8. When splitting the training data into \mathcal{L} and \mathcal{V} , we kept the same percentages of unknown binary class labels per example ($\%P$) in both the splits. Of course, in \mathcal{V} there are no fully-unlabeled examples ($\%L$ is 100). Moreover, when generating partial labels, we ensured that the percentages of discarded positive (i.e., 1) and negative (i.e., 0) class labels were the same.

9. We compared the outputs against 0.5 to obtain binary labels.

10. FT+C has more learnable weights: constraint loss is already small.




Fig. 4: Black-box attacks. Classification quality of vanilla and knowledge-constrained models in function of ϵ . Dotted plots include rejection (Rej) of inputs that are detected to be adversarial.




Fig. 5: White-box attacks in the case of the FT classifiers. Classification quality of vanilla and knowledge-constrained models in function of ϵ . Dotted plots include rejection (Rej) of inputs that are detected to be adversarial.

on a different realization of the training set. Fig. 4 shows the classification quality as a function of the data perturbation bound ϵ , comparing models trained with and without constraints against those implementing the detection/rejection mechanism described in Eq. (3). When using such mechanism, the rejected examples are marked as correctly classified if they are adversarial ($\epsilon > 0$), otherwise ($\epsilon = 0$) they are marked as points belonging to an unknown class, slightly worsening the performance. The +C/+CC models show larger accuracy/F1 than the unconstrained ones. Despite the lower results at $\epsilon = 0$, models that are more strongly constrained (+CC) resulted to be harder to attack for increasing values of ϵ . When the knowledge-based detector is activated, the improvements with respect to models without rejection are significantly evident. No model is specifically designed to face adversarial attacks and, of course, there are no attempts to reach state-of-the-art results.¹¹ However,

11. Recall that our rejection mechanism is completely agnostic to the attack; it neither assumes any knowledge of the attack algorithm nor is retrained on adversarial examples. Nevertheless, it can be used as a complementary defense mechanism.

the positive impact of exploiting domain knowledge can be observed in all the considered models and datasets, and for almost all the values of ϵ , confirming that such knowledge is not only useful to improve classifier robustness, but also as a mean to detect adversarial examples at no additional training cost. In general, FT models yield better results, due to the larger number of optimized parameters. In ANIMALS the rejection dynamics are providing large improvements in both TL and FT, while the impact of domain knowledge is mostly evident on the robustness of FT. In CIFAR-100, domain knowledge only consists of basic hierarchical relations, with no intersections among child classes or among father classes. By inspecting the classifier, we found that it is pretty frequent for the fooling examples to be predicted with a strongly-activated father class and a (coherent) child class, i.e., we have strongly-paired classes, accordingly to Def. 3.1. Differently, the domain knowledge in the other datasets is more structured, yielding better detection quality on average, remarking the importance of the level of detail of such knowledge to counter adversarial examples. In the case of PASCAL-Part, the detection mechanism turned out

to better behave in unconstrained classifiers, even if it has a positive impact also on the constrained ones. This is due to the intrinsic difficulty of making predictions on this dataset, especially when considering small object-parts. The false positives have a negative effect in the training stage of the knowledge-constrained classifiers.

To provide a comprehensive, worst-case evaluation of the adversarial robustness of our approach, we also considered a *white-box* adaptive attacker that knows everything about the target model and exploits knowledge of the defense mechanism to bypass it. Of course, this attack always evades detection if the perturbation size ϵ is sufficiently large. We evaluated multiple values of α of Eq. (11), selecting the one that yielded the lowest values of such objective function. In Fig. 5 we report the outcome of this analysis for FT models, showing that, even if the accuracy drop is obviously evident for all datasets, in ANIMALS the constrained classifiers require larger perturbations than the unconstrained ones to reduce the performance of the same quantity. A similar behavior is shown in CIFAR-100, even though only at small ϵ values. Accordingly, fooling the proposed detection mechanism is not always as trivial as one might expect, even in this worst-case setting. The impact of the rejection mechanisms is significantly reduced, as expected, but still having a positive impact. We refer the reader to the supplementary material for more details about these attacks and their optimization. Finally, let us point out that the performance drop caused by the white-box attack is much larger than that observed in the black-box case. However, since domain knowledge is not likely to be available to the attacker in many practical settings, it remains an open challenge to develop stronger, practical black-box attacks that are able to infer and exploit such knowledge to bypass our defense mechanism.

4.2.1 In-depth Analysis

Incomplete Domain Knowledge. We investigated in more detail the relative impact of the domain information on a target problem, simulating the availability of differently sized knowledge bases, $\mathcal{K}_1, \mathcal{K}_2, \mathcal{K}_3, \mathcal{K}_4$, where each $\mathcal{K}_j \subseteq \mathcal{K}$. In particular, we considered the ANIMALS dataset, and we generated $\mathcal{K}_1, \mathcal{K}_2, \mathcal{K}_3$ by removing some of the FOL formulas of the original \mathcal{K} that was used in the previous experiments (i.e., the one of Table 8, supplementary material), while $\mathcal{K}_4 = \mathcal{K}$. This means that some information that belongs to \mathcal{K}_4 is actually missing in the other knowledge sets. In detail, we created \mathcal{K}_1 by removing the rules that either include the ‘mammal’ or the ‘bird’ categories, while \mathcal{K}_2 is the outcome of discarding from \mathcal{K} the rules including the ‘mammal’ category. Similarly, \mathcal{K}_3 is obtained removing the rules of the ‘bird’ category. We executed a batch of independent experiments, each of them using only one of the generated knowledge bases, and focusing on the same models of Fig. 4 (bottom) and Fig. 5, that were retrained from scratch. Fig. 6 (a,c) shows the classification quality we obtained in the black (a) and white box (c) settings, using the MKA attack and $\epsilon = 0.5$ (almost the mid of the plots in Figs. 4-5). In the black-box case, focusing on the models that include rejection (Rej), it is evident how larger knowledge bases yield better results. Interestingly, comparing the

outcome of such models with the ones without rejection, we can see that our defense makes the classifier more robust to attacks even when using the smallest amount of knowledge (\mathcal{K}_1), confirming the versatility of what we propose. In the white-box setting there are still changes in the accuracy when varying \mathcal{K}_j , but they are not so evident and they lack a clear trend. This was expected, since, in this case, the attack procedure is aware of the domain knowledge. However, this result confirms the capability of MKA to craft adversarial examples that lead to knowledge-coherent predictions (to a certain extent) even when varying the level of detail of the knowledge sets.

Rejection Threshold. In the same experimental setting we also explored the sensitivity of the system to the rejection threshold τ , using the whole knowledge set \mathcal{K} . We compared different τ 's that are smaller or greater than the one we selected using validation data (Section 3), indicated here with τ^* . In particular, we evaluated $\tau \in \{10^k \tau^*, \kappa \in [-5, 5], \text{integer}\}$, and we measured both the classification quality on the perturbed data, as we did so far, and the rejection rate on the clean test data, where no samples should be rejected. Fig. 6 (b,d) reports these two measures on the y-axis and x-axis, respectively, and each marker is about a specific τ , considering black (b) and white (d) box settings. The value of τ^* has been highlighted with a red circle, and only the significant portions of the plots are shown. Too small thresholds (rightmost areas of each plot) lead to systems that frequently reject also clean examples, while too large τ 's (leftmost areas) do not improve the quality of the classifiers, that are fooled by some of the data generated in an adversarial manner within the ϵ -bound. Overall, the τ^* we selected represents a pretty appropriate trade-off between the two measures.

Noisy Domain Knowledge. We further extended our analysis by considering the case in which the available domain knowledge is noisy, thus including a small percentage of information that is incorrect in the ANIMALS domain. We simulated three scenarios by altering the original knowledge base \mathcal{K} using three different criteria, yielding the noisy knowledge bases $\tilde{\mathcal{K}}_a, \tilde{\mathcal{K}}_b, \tilde{\mathcal{K}}_c$, respectively. The chosen criteria either modify four of the existing rules, making them not coherent with the clean knowledge, or they add four new rules that are not correct in the considered domain, as shown in detail in Table 12, Table 13, and Table 14 of the supplementary material, and for which we provide a brief description in the following. In the case of $\tilde{\mathcal{K}}_a$, we selected four existing implications of \mathcal{K} , and we altered the premises of two of them and the conclusions of the other two ones. We ensured to inject noise in the main and auxiliary classes in a balanced manner. The same balancing is also kept when generating $\tilde{\mathcal{K}}_b$, that, however, was obtained by adding four new rules to \mathcal{K} . Finally, $\tilde{\mathcal{K}}_c$ is the result of augmenting each main-class-oriented conclusion of four existing implications with a disjunction involving a different, randomly selected, main class. For example, the clean rule $\text{BLACKSTRIPES}(x) \wedge \text{UNGULATE}(x) \wedge \dots \Rightarrow \text{ZEBRA}(x)$ is altered by replacing the conclusion $\text{ZEBRA}(x)$ with $\text{ZEBRA}(x) \vee \text{TIGER}(x)$. These rules are fulfilled both for configurations that make true their original/clean counterparts, and for other configurations that are actually wrong in

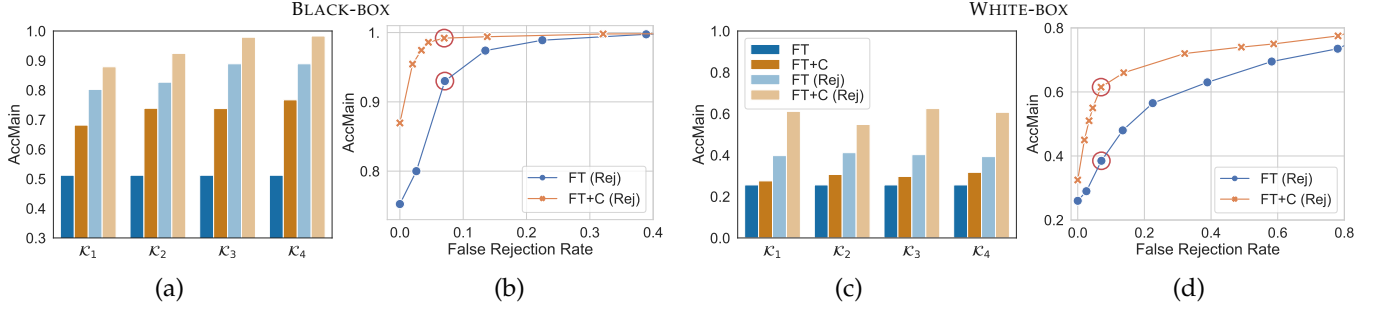


Fig. 6: Further analysis of the proposed approach in the ANIMALS dataset ($\epsilon = 0.5$). The first two plots are about the black-box setting, and the last two ones are about the white-box case; (a,c)—legend reported only on the latter, for better readability; increasing amounts of domain knowledge $\mathcal{K}_1, \dots, \mathcal{K}_4$; (b,d): different values of the rejection threshold τ (from large to smaller values, left-to-right).

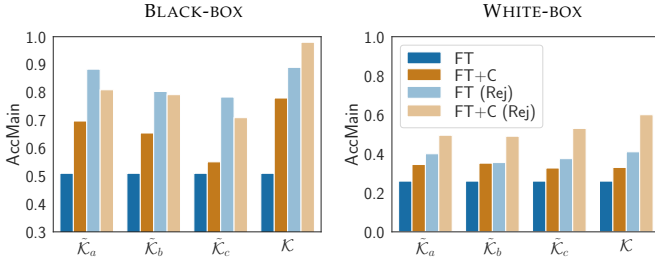


Fig. 7: Noisy domain knowledge. Analysis of the proposed approach in the same setup of Fig. 6, when exploiting different noisy knowledge bases $\tilde{\mathcal{K}}_a, \tilde{\mathcal{K}}_b, \tilde{\mathcal{K}}_c$ (see the paper text for details). \mathcal{K} is the original noise-free knowledge.

the ANIMALS domain. Focusing on the same experimental setup we defined when testing differently sized knowledge bases, we investigated the effects of using each noisy knowledge base both to learn the classifier and/or as a rejection criterion. Fig. 7 shows the results of our experience. As expected, learning with a noisy knowledge base reduces the accuracy of the classifier, since the network learns from FOL rules that are partially in contrast with some of the available labeled examples. It is the case of FT+C trained on any noisy $\tilde{\mathcal{K}}$, when compared with the case of the clean \mathcal{K} (rightmost set of bars). However, when adding the rejection criterion, we still observe a significant improvement in the performance, even if slightly smaller than in the case of \mathcal{K} . Rejection is the outcome of evaluating the average violations of all the available rules, so that the effect of the noisy portion of the knowledge base is partially compensated by the other non-noisy rules. Of course, the final outcome depends on the type of noise we injected into the knowledge base. In the considered experience, adding wrong rules ($\tilde{\mathcal{K}}_b$) led to lower accuracies than when perturbing some of the existing rules ($\tilde{\mathcal{K}}_a$). The case of $\tilde{\mathcal{K}}_c$ is the one that most evidently impacted the classifier performance, with or without rejection. On one hand, we simply introduced more error-tolerant conclusions; on the other hand, we did it by altering FOL rules whose conclusions are all about main classes, significantly compromising the way they are related. Interestingly, for all the noisy knowledge bases, adding the rejection module (Rej) to FT turned out to be better than adding it to FT+C, suggesting that a rejection

criterion based on noisy knowledge could be more effective in classifiers that have not been already exposed to such noisy information during the learning stage. As expected, results collected in the white-box case show that whenever the attacker has full access to the knowledge, being it noisy or not, he can craft MKA attacks with similar outcomes in terms of performance drops. Overall, this experience confirms that what we propose can indeed be applied also in case of partially noisy domain knowledge, still increasing the robustness of the classifier, even if to a smaller extent than in case of clean knowledge.

4.3 Experimental Results on Single-label Classifiers

The focus of this paper is on multi-class classification paired with domain knowledge. However, as anticipated in Section 3 and qualitatively shown in Fig. 3, we can consider a special setting in which a single-label classifier internally includes predictors over auxiliary classes that are involved in the knowledge constraints. We experimentally evaluate this setting in the context of the ANIMALS and CIFAR-100 datasets, where the respective main classes (described in Section 4.1) are mutually exclusive (which is not the case of PASCAL-Part), thus well suited to simulate the setting of Fig. 3. We compared the proposed rejection mechanisms to a concurrent defense mechanism developed under the same assumptions (i.e., without assuming any knowledge of the attack algorithm), known as Neural Rejection (NR) [58], [59], and against the state-of-the-art attacks included in the AutoAttack framework [47], developed for single-label classification tasks.

Compared Defense and Attack Strategies. The NR defense mechanism, proposed in [58], [59], aims to reject inputs that are far from the training data in a given representation space. The rationale is that points with low support from the training set cannot be reliably classified, and should be thus rejected. To this end, the output layer of the deep network is replaced with a Support Vector Machine trained using the RBF kernel (SVM-RBF), which enforces the prediction scores to be proportional to the distance between the input sample and the reference prototypes (i.e., the support vectors) in the representation space. Samples are rejected if the prediction scores do not exceed the rejection threshold. Similarly to our approach, this defense mechanism does not make any assumptions on the attack to be detected, other

than assuming an anomalous behavior with respect to the observed training data.

To compare our defense with NR, we have considered four different state-of-the-art evasion attacks: APGD-CE, APGD-T, FAB-T, and Square, implemented within the framework of AutoAttack [47]. APGD-CE (APGD-T) is an indiscriminate (targeted) step-free variant of the famous attack called PGD [60]. Unlike PGD, the step size reduction is not scheduled a priori but instead governed by the optimization function trend. Moreover, both APGDs attack use momentum. FAB-T is the targeted version of an attack called Fast Adaptive Boundary Attack (FAB) [61], which tries to find the minimum distance sample beyond the boundary of the desired class. The Square attack [62], differently from the previously mentioned ones, is a black-box attack; namely, it can query the classifier obtaining the predicted scores without exploiting any knowledge of the model architecture. By default, APGD-CE makes five random restarts, whereas the targeted versions of the attacks, i.e., APGD-T and FAB-T, run the attack nine times, each setting the target class as one of the nine top classes except the true class.

Adversarial Evaluation. In our experiments, we fixed the maximum allowed perturbation ϵ to 0.5 and 0.03 on the ANIMALS and CIFAR-100 datasets, respectively – the values in the middle of x-axis of Figs. 4-5 – and we used the default value for all the other attacks’ parameters. Table 6 and Table 7 report the results of this analysis, showing the classification quality (same measure of Figs. 4-5) on the clean (unmodified) test set \mathcal{T} ($\epsilon = 0$) and on the attacked instances of \mathcal{T} generated in the same white-box and black-box scenarios described in Section 4.2. As expected, white-box attacks are more effective than the black-box ones, reducing the model accuracy in a more evident manner. Results confirm that both the domain knowledge introduced at training time (+C and +CC) or exploited to implement the proposed rejection mechanism (Rej) improve the model robustness against all the considered attacks, and jointly using these strategies further improves it. On average, the performances of the unconstrained classifiers (TL and FT) paired with the proposed knowledge-based rejection are comparable with the ones paired with NR, even though they clearly behave in different manners across the datasets/attacks. Differently, when also considering constrained models (+C and +CC), in most of the cases we can find a classifier with rejection (Rej) that outperforms the unconstrained classifiers equipped with NR. On the clean samples ($\epsilon = 0$), the knowledge-based rejection criterion resulted more aggressive than NR.

MKA and APGD-CE/T are more effective than the other attacks. On average, their performances are comparable, and they depend on both the considered model and the dataset. In CIFAR-100, MKA outperforms APGD-CE/T on the black-box transfer scenario and against the model equipped with the proposed rejection mechanism (Rej), whereas on the ANIMALS dataset, APGD-CE usually obtained better results. APGD-CE/T leverages an optimization strategy that is more advanced than the one of MKA that, differently from APGD-CE/T, is designed to be used in multi-label problems too. For example, APGD-CE/T makes several attack restarts and uses a special type of adaptive step size. In the white-box setting, attacks yield a larger reduction of the perfor-

mances. However, in the case of ANIMALS, the proposed rejection mechanism is still robust to all the attacks, with the exceptions of MKA, that is knowledge aware, and of APGD-CE. The fine-tuned optimization procedure in APGD-CE allows the attack to create samples that are confidently misclassified, and they end-up in belonging to space regions in which the classification functions are paired (Def. 3.1). In CIFAR-100, the rejection mechanism still has a positive impact, even if it is less significant than in ANIMALS.

4.3.1 In-depth Analysis

We further analyzed our results, visualizing the behavior of all the compared attacks in terms of value of the constraint loss of Eq. (3) and of the supervision loss – first term of Eq. (2). Figs. 8-9 show each generated adversarial example, highlighting them with different markers/colors in function of the corresponding attack procedure, on the ANIMALS and CIFAR-100 datasets, respectively, black-box (i.e., the constraint loss is measured for the purpose of determining whether to reject or not an example). Samples that are rejected are indicated with crosses, while circles represent the non-rejected ones. The dotted line is about the rejection threshold τ from Eq. (8).




Fig. 8: Adversarial data generated ($\epsilon = 0.5$) by different attacks – ANIMALS, TL+C(Rej), black-box. Examples that are rejected/not-rejected by the proposed knowledge-based criterion are depicted with crosses/circles (“Clean” indicates unaltered examples from the test set; the vertical line is the reject threshold).

In line with what we observed in the numerical results, in the case of ANIMALS, Fig. 8, it is evident how APGD-CE is actually able to craft attacks that strongly increase the supervision loss, still fulfilling the constraints (top-left area). Differently, the other attacks are not able to reach such result, so that their data is localized in high-constraint loss regions, easily rejected by the proposed technique, especially FAB-T, while Square actually fails in generating evident attacks. It is interesting to notice the D-shaped white region over the origin. It is an area in which constraints are almost fulfilled and the loss function can reach significantly non-null values, but no attacks fall there. This suggests that it is not straightforward to increase the supervision loss without violating the constraints. However, there are more extreme

TABLE 6: ANIMALS dataset. Vulnerability analysis of the classifiers against MKA and state-of-art attacks—classification quality is reported, the same of Figs. 4-5 (first column). For each type of classifier (TL,FT), rows are organized into three groups, that are: models without rejection, with rejection (Rej), classifier equipped with Neural Rejection (NR). For each attack (columns—see [47] for a description of the compared attacks), the result of the most robust classifier in the group is highlighted in bold. Models exploiting the proposed rejection (Rej) that overcome NR are marked with *, and vice-versa.

White-box attacks ($\epsilon = 0.5$)							Black-box transfer attacks ($\epsilon = 0.5$)				
Model	$\epsilon = 0$	MKA	APGD-CE	APGD-T	FAB-T	Square	MKA	APGD-CE	APGD-T	FAB-T	Square
TL	99.0	25.0	17.5	14.9	20.0	96.6	45.3	29.4	29.1	83.1	98.4
TL+C	99.3	25.0	19.0	15.4	21.5	98.0	47.7	29.8	30.6	87.7	98.9
TL+CC	99.3	24.5	18.1	15.5	22.7	98.2	48.0	30.2	32.0	89.5	99.0
TL (Rej)	91.8	49.8	43.3	97.0*	100.0*	100.0*	85.6	53.7	98.4	99.9	99.9
TL+C (Rej)	92.3	56.8*	47.8	97.7*	100.0*	100.0*	91.2	56.0	98.6	99.9	99.9
TL+CC (Rej)	92.7	57.5*	45.8	98.1*	100.0*	100.0*	93.4	55.4	98.8	100.0*	100.0*
TL (NR)	99.2	55.5	58.5*	96.8	99.6	99.8	98.0*	71.7*	99.3*	100.0*	100.0*
FT	98.6	25.6	21.9	12.9	20.0	96.3	51.2	47.2	75.6	95.7	98.2
FT+C	99.1	31.7	51.1	18.0	29.5	97.7	76.7	57.5	88.8	98.3	98.9
FT (Rej)	92.7	39.3*	36.8	90.0*	99.7*	99.7*	88.9	66.9	99.6*	99.7*	99.8*
FT+C (Rej)	93.2	60.7*	66.6*	97.3*	99.8*	99.9*	98.3*	82.2*	99.9*	99.9*	99.9*
FT (NR)	98.6	37.3	38.3	87.3	97.0	99.6	91.0	79.2	98.7	99.2	99.5

TABLE 7: CIFAR-100 dataset. Vulnerability analysis of the classifiers against MKA and state-of-art attacks—classification quality is reported, the same of Figs. 4-5 (second column). Refer to the caption of Table 6 for more details (see [47] for a description of the compared attacks).

White-box attacks ($\epsilon = 0.03$)							Black-box transfer attacks ($\epsilon = 0.03$)				
Model	$\epsilon = 0$	MKA	APGD-CE	APGD-T	FAB-T	Square	MKA	APGD-CE	APGD-T	FAB-T	Square
TL	51.0	21.9	22.2	21.6	22.3	51.4	23.1	23.7	23.9	39.5	52.7
TL+C	52.9	27.4	24.6	24.2	25.1	53.3	32.3	35.5	37.9	48.4	54.5
TL+CC	50.5	27.1	25.0	24.7	25.3	49.5	35.5	38.6	40.4	46.9	51.5
TL (Rej)	48.1	26.9	33.2*	34.3*	34.4	59.2*	27.6	35.1	36.9	49.1	60.2*
TL+C (Rej)	49.4	31.8*	35.0*	35.6*	36.2	60.6*	40.7	44.8	47.0	56.0*	61.0*
TL+CC (Rej)	46.1	30.8*	34.0*	34.7*	35.4	55.7*	45.4	46.3	47.6	53.5*	57.0*
TL (NR)	49.0	30.5	30.1	24.5	39.6*	45.6	49.0*	48.3*	49.0*	51.3	53.3
FT	59.4	29.0	26.4	26.0	26.7	57.2	48.4	49.1	49.7	55.5	59.5
FT+C	60.0	31.4	29.6	28.3	30.6	60.1	51.6	52.2	52.8	57.8	61.0
FT (Rej)	57.4	31.1	37.5*	42.0*	41.1	66.1*	55.1	57.2*	58.4*	62.7*	66.2*
FT+C (Rej)	56.7	37.6*	37.8*	41.1*	44.6	67.0*	60.2*	59.5*	60.3*	64.4*	67.1*
FT (NR)	59.7	36.5	35.3	30.4	50.9*	55.1	58.0	54.2	55.7	60.0	62.7

APCG-CE configurations with the largest supervision losses that also fulfill the knowledge (Fig. 8, top-left area). Of course, this depends on several factors, such as the type of domain knowledge that is available, the way we selected to convert it into polynomial constraints, and the constraint enforcement scheme, thus opening to future improvements. Moving to the CIFAR-100 dataset, Fig. 9, we observe different patterns with respect to the case of ANIMALS. This was clearly expected, since the two datasets differ both in terms of the problem they consider, the number of classes and in terms of the known relationships among such classes, described by the dataset-specific domain knowledge and embedded into the constraint loss. However, we can still observe the D-shaped region over the origin, even if in a less significant manner. On this dataset, the rejection rates are generally lower than ANIMALS. This is mostly due to the fact the constraint loss is larger also on the unaltered data, due to the already mentioned different problem and different type of domain knowledge. As a matter of fact,

we have also a larger reject threshold τ . In this case, the behavior of the different attack strategies is more coherent, remarking previous considerations on the role of knowledge in shaping the distribution of the attacks.

5 RELATED WORK

In addition to the literature described in Section 1, we further emphasize the main differences of what we propose with respect to the most strongly related approaches.

Multi-label adversarial perturbations. Most of the work in the adversarial ML area focuses on single-label classification problems. To the best of our knowledge, the first and only study on this problem is the one in [27], in which the authors focus on targeted multi-label adversarial perturbations defining in advance the set of classes on which the attack is targeted (being them positive or negative) and also introducing another set of classes for which the attack is expected not to change the classifier predictions. The framework described in [27] is only experimented in




Fig. 9: Adversarial data generated ($\epsilon = 0.03$) by different attacks – CIFAR-100, TL+C(Rej), black-box. See Fig. 8.

a *static/targeted* context, i.e., by selecting in advance the sets mentioned above using custom criteria to simulate the attacking scenario artificially. The multi-label attack that we propose in this work is instead *dynamic/untargeted* and without the need of defining in advance what are the classes to be considered. Regarding the defenses, to our best knowledge, none of the previously proposed ones leverage multi-label classification outputs.

Semi-supervised Learning and Adversarial Training. In the context of adversarial machine learning, unlabeled data are usually employed to improve the robustness of the classifier by performing adversarial training. The rationale behind such training scheme is that if the available unlabeled samples are perturbed, then the predicted class should not change. Miyato et al. [19], [23] and Park et al. [20] exploit adversarial training (virtual adversarial training and adversarial dropout, respectively) to favor regularity around the supervised and unsupervised training data, and to improve the classifier performance. The work in [21] develops an anomaly detector using adversarial training in the semi-supervised setting. Self-supervised learning is exploited in [22], [25] to gain stronger adversarial robustness. Stability criteria are enforced on unlabeled training data in [24], whereas the work in [26] specifically focuses on an unsupervised adversarial training procedure in the context of semi-supervised classification. Our model neither exploits adversarial training nor any adversary-aware training criteria aimed at gaining intrinsic regularity. We focus on the role of domain knowledge as an indirect means to increase adversarial robustness and, afterward, to detect adversarial examples. Therefore, the proposed work is not attack-dependent, and it is faster at training time as it does not require generating adversarial examples. We believe that using unlabeled data also to simulate attacks and incorporate them into the training process may further improve robustness. All the described methods could also be applied jointly with what we propose.

Rejection-based Approaches for Adversarial Examples. A different line of defenses, complementary with adversarial training, is based on detecting and rejecting samples suffi-

ciently far from the training data in feature space. Our approach differs from other adversarial-example detectors [8], [13], [14], [15] as it has no additional training cost and negligible runtime cost. We are the first to show that domain knowledge can be used to reject adversarial examples and also to propose a detector that exploits unlabeled data.

Domain-Agnostic Methods and Semantic Attacks. Recent work in adversarial attacks considers the role of the learning domain and of additional semantic information, even if with different goals to the ones of this paper. The way the learning domain is related to the generation of attacks was recently studied in [34], that is based on the idea of developing generative adversarial perturbations that turn out to be easily transferable from the source domain (where the attack function is modeled) to another domain. Differently, we focus on knowledge that is domain specific and used both for defending and creating more informed attacks. The knowledge of a set semantic attributes is used to implement the threat model of semantic adversarial attacks in [35]. A generative network is considered, and the attack procedure focuses on altering the activation of such human-understandable attributes, that, in turn, yield visible changes in the input image (e.g., adding glasses to the input face). Differently, our work is built on an L_p -norm-bounded perturbation model that does not enforce the input image to change in a human-understandable manner. Our approach considers a more generic notion of knowledge, that includes information also on the relationships within subsets of logic predicates, and that exploits the power of FOL. Predicate activations are modeled by neural networks and not by scalar variables as for the attributes of [35].

6 CONCLUSIONS AND FUTURE WORK

In this paper we investigated the role of domain knowledge in adversarial settings. Focusing on multi-label classification, we injected knowledge expressed by First-Order Logic in the training stage of the classifier, not only with the aim of improving its quality, but also as a mean to build a detector of adversarial examples at no additional cost. We proposed a multi-label attack procedure and showed that knowledge-constrained classifiers can improve their robustness against both black-box and white-box attacks, and, using the same knowledge, they can detect adversarial attacks. We believe that these findings will open the investigation of domain knowledge as a feature to further improve the robustness of multi-label classifiers against adversarial attacks.

The proposed adversarial example rejection scheme is based on the idea of dealing with classifiers that fulfill the knowledge-related constraints over the space regions in which the non-malicious data are distributed, not guaranteeing such fulfillment in the rest of the space. While this is experimentally evaluated to be a key ingredient to profitably build a rejection strategy, we showed that advanced optimization strategies can fool the defense injecting a stronger perturbation than the one used to fool the undefended system. In future work we will consider intermixing adversarial training with knowledge constraints, to strengthen the violation of the constraints out of the distribution of the real data. We also plan to design a learnable model that decides whether to reject or not in function of the fulfillment of each

specific logic formula, going beyond a simple-but-effective threshold on the cumulative constraint loss.

ACKNOWLEDGMENT

This work was partly supported by the PRIN 2017 project RexLearn, funded by the Italian Ministry of Education, University and Research (grant no. 2017TWNMH2).

REFERENCES


- [1] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," in *International Conference on Learning Representations*, 2014. [Online]. Available: <http://arxiv.org/abs/1312.6199>
- [2] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrđić, P. Laskov, G. Giacinto, and F. Roli, "Evasion attacks against machine learning at test time," in *Machine Learning and Knowledge Discovery in Databases (ECML PKDD), Part III*, ser. LNCS, H. Blockeel, K. Kersting, S. Nijssen, and F. Železný, Eds., vol. 8190. Springer Berlin Heidelberg, 2013, pp. 387–402.
- [3] N. Papernot, P. McDaniel, and I. Goodfellow, "Transferability in machine learning: from phenomena to black-box attacks using adversarial samples," *arXiv preprint, arXiv:1605.07277*, 2016.
- [4] I. Goodfellow, P. McDaniel, and N. Papernot, "Making machine learning robust against adversarial inputs," *Communications of the ACM*, vol. 61, no. 7, pp. 56–66, 2018.
- [5] N. Carlini, A. Athalye, N. Papernot, W. Brendel, J. Rauber, D. Tsipras, I. Goodfellow, A. Madry, and A. Kurakin, "On evaluating adversarial robustness," *arXiv preprint, arXiv:1902.06705*, 2019.
- [6] A. Shafahi, W. R. Huang, C. Studer, S. Feizi, and T. Goldstein, "Are adversarial examples inevitable?" in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=r1WUoA9FQ>
- [7] A. Sotgiu, A. Demontis, M. Melis, B. Biggio, G. Fumera, X. Feng, and F. Roli, "Deep neural rejection against adversarial examples," *EURASIP Journal on Information Security*, vol. 2020, pp. 1–10, 2020.
- [8] D. J. Miller, Z. Xiang, and G. Kesidis, "Adversarial learning targeting deep neural network classification: A comprehensive review of defenses against attacks," *Proceedings of the IEEE*, vol. 108, no. 3, pp. 402–433, 2020.
- [9] B. Biggio and F. Roli, "Wild patterns: Ten years after the rise of adversarial machine learning," *Pattern Recognition*, vol. 84, pp. 317–331, 2018.
- [10] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *International Conference on Learning Representations*, 2015. [Online]. Available: <http://arxiv.org/abs/1412.6572>
- [11] N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami, "Distillation as a defense to adversarial perturbations against deep neural networks," in *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2016, pp. 582–597.
- [12] A. Sinha, H. Namkoong, and J. Duchi, "Certifying some distributional robustness with principled adversarial training," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=Hk6kPgZA->
- [13] N. Carlini and D. Wagner, "Adversarial examples are not easily detected: Bypassing ten detection methods," in *Proceedings of the ACM Workshop on Artificial Intelligence and Security*, 2017, pp. 3–14.
- [14] X. Ma, B. Li, Y. Wang, S. M. Erfani, S. Wijewickrema, G. Schoenebeck, M. E. Houle, D. Song, and J. Bailey, "Characterizing adversarial subspaces using local intrinsic dimensionality," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=B1gJlL2aW>
- [15] P. Samangouei, M. Kabkab, and R. Chellappa, "Defense-GAN: Protecting classifiers against adversarial attacks using generative models," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=Bk3ibb0->
- [16] T. Pang, C. Du, Y. Dong, and J. Zhu, "Towards robust detection of adversarial examples," in *Advances in Neural Information Processing Systems*, vol. 31. Curran Associates, Inc., 2018.
- [17] K. Lee, K. Lee, H. Lee, and J. Shin, "A simple unified framework for detecting out-of-distribution samples and adversarial attacks," in *Advances in Neural Information Processing Systems*, vol. 31. Curran Associates, Inc., 2018.
- [18] A. Araujo, L. Meunier, R. Pinot, and B. Negrevergne, "Robust Neural Networks using Randomized Adversarial Training," *arXiv preprint, arXiv:1903.10219*, 2020.
- [19] T. Miyato, S.-i. Maeda, M. Koyama, K. Nakae, and S. Ishii, "Distributional smoothing with virtual adversarial training," in *International Conference on Learning Representations*, 2016. [Online]. Available: <https://arxiv.org/pdf/1507.00677.pdf>
- [20] S. Park, J. Park, S.-J. Shin, and I.-C. Moon, "Adversarial dropout for supervised and semi-supervised learning," in *Thirty-Second AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, Apr. 2018.
- [21] S. Akcay, A. Atapour-Abarghouei, and T. P. Breckon, "Ganomaly: Semi-supervised anomaly detection via adversarial training," in *Asian Conference on Computer Vision*. Springer, 2018, pp. 622–637.
- [22] Y. Carmon, A. Raghunathan, L. Schmidt, J. C. Duchi, and P. S. Liang, "Unlabeled data improves adversarial robustness," in *Neural Information Processing Systems*, 2019, pp. 11 190–11 201.
- [23] T. Miyato, S.-i. Maeda, M. Koyama, and S. Ishii, "Virtual adversarial training: a regularization method for supervised and semi-supervised learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 8, pp. 1979–1993, 2018.
- [24] R. Zhai, T. Cai, D. He, C. Dan, K. He, J. Hopcroft, and L. Wang, "Adversarially robust generalization just requires more unlabeled data," *arXiv preprint, arXiv:1906.00555*, 2019.
- [25] A. Najafi, S.-i. Maeda, M. Koyama, and T. Miyato, "Robustness to adversarial perturbations in learning from incomplete data," in *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc., 2019.
- [26] J.-B. Alayrac, P.-S. Huang, A. Fawzi, R. Stanforth, and P. Kohli, "Are labels required for improving adversarial robustness?" in *Advances in Neural Information Processing Systems*, vol. 32. Curran Associates, Inc., 2019.
- [27] Q. Song, H. Jin, X. Huang, and X. Hu, "Multi-label adversarial perturbations," in *2018 IEEE International Conference on Data Mining (ICDM)*, 2018, pp. 1242–1247.
- [28] Y. Wu, D. Bamman, and S. Russell, "Adversarial training for relation extraction," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2017, pp. 1778–1783.
- [29] R. Babbar and B. Schölkopf, "Adversarial extreme multi-label classification," *arXiv preprint, arXiv:1803.01570*, 2018.
- [30] G. Gnecco, M. Gori, S. Melacci, and M. Sanguineti, "Foundations of support constraint machines," *Neural computation*, vol. 27, no. 2, pp. 388–480, 2015.
- [31] M. Diligenti, M. Gori, and C. Sacca, "Semantic-based regularization for learning and inference," *Artificial Intelligence*, vol. 244, pp. 143–165, 2017.
- [32] K. Grosse, P. Manoharan, N. Papernot, M. Backes, and P. McDaniel, "On the (statistical) detection of adversarial examples," *arXiv preprint, arXiv:1702.06280*, 2017.
- [33] S. Melacci and M. Belkin, "Laplacian Support Vector Machines Trained in the Primal," *Journal of Machine Learning Research*, vol. 12, pp. 1149–1184, March 2011.
- [34] M. M. Naseer, S. H. Khan, M. H. Khan, F. Shahbaz Khan, and F. Porikli, "Cross-domain transferability of adversarial perturbations," *Neural Information Processing Systems*, vol. 32, 2019.
- [35] A. Joshi, A. Mukherjee, S. Sarkar, and C. Hegde, "Semantic adversarial attacks: Parametric transformations that fool deep classifiers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 4773–4783.
- [36] R. Sheatsley, B. Hoak, E. Pauley, Y. Beugin, M. J. Weisman, and P. McDaniel, "On the robustness of domain constraints," *arXiv preprint, arXiv:2105.08619*, 2021.
- [37] N. Carlini, A. Athalye, N. Papernot, W. Brendel, J. Rauber, D. Tsipras, I. Goodfellow, A. Madry, and A. Kurakin, "On evaluating adversarial robustness," *arXiv preprint, arXiv:1902.06705*, 2019.
- [38] A. Athalye, N. Carlini, and D. Wagner, "Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples," in *Proceedings of the 35th International Conference on Machine Learning*, vol. 80. PMLR, 10–15 Jul 2018, pp. 274–283.
- [39] M. Gori and S. Melacci, "Constraint verification with kernel machines," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 24, no. 5, pp. 825–831, 2013.
- [40] E. P. Klement, R. Mesiar, and E. Pap, *Triangular norms*. Springer Science & Business Media, 2013, vol. 8.
- [41] M. Hein, M. Andriushchenko, and J. Bitterwolf, "Why relu networks yield high-confidence predictions far away from the training data and how to mitigate the problem," in *Proceedings of the*

- IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 41–50.
- [42] D. Hendrycks and K. Gimpel, “A baseline for detecting misclassified and out-of-distribution examples in neural networks,” 2017. [Online]. Available: <https://arxiv.org/pdf/1610.02136.pdf>
- [43] —, “Early methods for detecting adversarial images,” in *International Conference on Learning Representations*, 2017. [Online]. Available: <https://openreview.net/forum?id=B1dexpDug>
- [44] S. Teso, “Does symbolic knowledge prevent adversarial fooling?” *arXiv preprint, arXiv:1912.10834*, 2019.
- [45] A. Bendale and T. E. Boult, “Towards open set deep networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 1563–1572.
- [46] A. Demontis, M. Melis, B. Pintor, M. Jagielski, B. Biggio, A. Oprea, C. Nita-Rotaru, and F. Roli, “Why do adversarial attacks transfer? Explaining transferability of evasion and poisoning attacks,” in *28th USENIX Security Symposium (USENIX Security 19)*. USENIX Association, 2019.
- [47] F. Croce and M. Hein, “Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks,” in *International Conference on Machine Learning*, 2020, pp. 1–12.
- [48] N. Carlini and D. A. Wagner, “Towards evaluating the robustness of neural networks,” in *IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2017, pp. 39–57.
- [49] A. S. d’Avila Garcez, M. Gori, L. C. Lamb, L. Serafini, M. Spranger, and S. N. Tran, “Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning,” *Journal of Applied Logics - IfCoLog Journal*, vol. 6, no. 4, pp. 611–632, 2019.
- [50] S. Melacci, A. Globo, and L. Rigutini, “Enhancing modern supervised word sense disambiguation models by semantic lexical resources,” in *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, 2018.
- [51] M. Yu and M. Dredze, “Improving lexical embeddings with semantic knowledge,” in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 2014, pp. 545–550.
- [52] T. Pi, X. Li, and Z. M. Zhang, “Boosted zero-shot learning with semantic correlation regularization,” in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 2017, pp. 2599–2605. [Online]. Available: <https://doi.org/10.24963/ijcai.2017/362>
- [53] P. Morgado and N. Vasconcelos, “Semantically consistent regularization for zero-shot recognition,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6060–6069.
- [54] G. Ciravegna, F. Giannini, M. Gori, M. Maggini, and S. Melacci, “Human-driven fol explanations of deep learning,” in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, C. Bessiere, Ed. International Joint Conferences on Artificial Intelligence Organization, 7 2020, pp. 2234–2240, main track.
- [55] P. Barbiero, G. Ciravegna, F. Giannini, P. Li’o, M. Gori, and S. Melacci, “Entropy-based logic explanations of neural networks,” *arXiv preprint, arXiv:2106.06804*, 2021.
- [56] I. Donadello, L. Serafini, and A. D. Garcez, “Logic tensor networks for semantic image interpretation,” in *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, ser. IJCAI’17. AAAI Press, 2017, p. 1596–1602.
- [57] P. H. Winston and B. K. Horn, “Lisp,” 1986.
- [58] M. Melis, A. Demontis, B. Biggio, G. Brown, G. Fumera, and F. Roli, “Is deep learning safe for robot vision? Adversarial examples against the iCub humanoid,” in *ICCVW Vision in Practice on Autonomous Robots (ViPAR)*. IEEE, 2017, pp. 751–759.
- [59] A. Sotgiu, A. Demontis, M. Melis, B. Biggio, G. Fumera, X. Feng, and F. Roli, “Deep neural rejection against adversarial examples,” *EURASIP J. Information Security*, vol. 2020, no. 5, 2020.
- [60] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=rjzIBfZAb>
- [61] F. Croce and M. Hein, “Minimally distorted adversarial examples with a fast adaptive boundary attack,” in *Proceedings of the 37th International Conference on Machine Learning*, vol. 119. PMLR, 13–18 Jul 2020, pp. 2196–2205.
- [62] M. Andriushchenko, F. Croce, N. Flammarion, and M. Hein, “Square Attack: A Query-Efficient Black-Box Adversarial Attack

via Random Search,” in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, pp. 484–501.



Stefano Melacci received the M.S. Degree in Computer Engineering (cum laude) and the PhD degree in Computer Science (Information Engineering) from the University of Siena, Italy, in 2006 and 2010, respectively. He worked as Visiting Scientist at the Computer Science and Engineering Department of the Ohio State University, Columbus, USA, and he is currently Associate Professor of the Department of Information Engineering and Mathematics, University of Siena. His research interests include machine learning and pattern recognition, mainly focused on Neural Networks and Kernel Machines, with applications to Computer Vision and Natural Language Processing tasks. He serves as Associate Editor of the IEEE Transactions on Neural Network and Learning Systems.



Gabriele Ciravegna is a PhD student at the University of Siena since 2018 under the supervision of Professor Marco Gori. In 2018 he received the master's degree in Computer Engineering with honors at the Polytechnic University of Turin. He has always been interested in the machine learning field. Nowadays, he is focused on overcoming the intrinsic limits of machine learning and neural networks, especially in the context of Explainable AI. He presented his works in several international venues such as AAAI, IJCAI, IJCNN. He also serves as reviewer in conferences and journals that are about Neural Networks, such as IEEE Transactions on Neural Networks and Learning Systems.



Angelo Sotgiu is a Ph.D. student in Electronic and Computer Engineering at the University of Cagliari, Italy. He received his M. Sc. in Telecommunication Engineering with honors from the University of Cagliari, Italy, in 2019. His research interests include secure machine learning and computer security. He serves as a reviewer for several journals and conferences in the machine learning and computer security area.



Ambra Demontis is an Assistant Professor at the University of Cagliari, Italy. She received her M.Sc. degree (Hons.) in Computer Science and her Ph.D. degree in Electronic Engineering and Computer Science from the University of Cagliari, Italy, respectively, in 2014 and 2018. Her research interests include secure machine learning, kernel methods, and computer security. In particular, she has provided contributions to the design of secure machine learning systems in the presence of intelligent attackers and highlighted interesting trade-offs between their complexity and security. She co-organizes the AISec workshop, serves on the program committee of different conferences and workshops, such as IJCAI and DLS, and as a reviewer for several journals, such as TNNLS, TOPS Machine Learning, and Pattern Recognition. She is a Member of the IEEE and the IAPR.



Battista Biggio (MSc 2006, PhD 2010) is Assistant Professor at the University of Cagliari, Italy, and co-founder of Pluribus One (<https://www.pluribus-one.it>). His research interests include machine learning, biometrics and cybersecurity. He has provided pioneering contributions in the area of adversarial machine learning, demonstrating gradient-based evasion and poisoning attacks, and how to mitigate them, playing a leading role in the establishment and advancement of this research field. He regularly serves

as a program committee member for the most prestigious conferences and journals in the area of machine learning and computer security (ICML, NeurIPS, ICLR, IEEE SP, USENIX Sec.). He chaired the IAPR TC on Statistical Pattern Recognition Techniques, co-organized the S+SSPR, AISEC and DLS workshops, and serves as Associate Editor for Pattern Recognition, IEEE TNNLS and IEEE CIM. Dr. Biggio is a senior member of the IEEE and member of the IAPR and of the ACM.



Marco Gori received the Ph.D degree in 1990 from University of Bologna, Italy, working partly at the School of Computer Science (McGill University, Montreal). He is currently full professor at the University of Siena and he is mostly interested in machine learning with applications to pattern recognition, Web mining, and game playing. He has recently published the monograph "Machine Learning: A constraint-based approach," (MK, 560 pp., 2018), which contains a unified view of his approach to machine learning.

His pioneering role in neural networks has been emerging especially from the recent interest in Graph Neural Networks, that he contributed to introduce in the paper "Graph Neural Networks," IEEE-TNN, 2009. Professor Gori has been the chair of the Italian Chapter of the IEEE Computation Intelligence Society and the President of the Italian Association for Artificial Intelligence. He is a Fellow of IEEE, EurAI, and IAPR. He was one the first people involved in European project on Artificial Intelligence CLAIRE, and he is currently a Fellow of Machine Learning association ELLIS. He is in the scientific committee of ICAR-CNR and is the President of the Scientific Committee of FBK-ICT. Dr. Gori is currently holding an international 3IA Chair at the Université Cote d'Azur.



Fabio Roli received his Ph.D. in Electronic Engineering from the University of Genoa, Italy. He was a research group member of the University of Genoa ('88-'94), and adjunct professor at the University of Trento ('93-'94). In 1995, he joined the Department of Electrical and Electronic Engineering of the University of Cagliari, where he is now Full Professor of Computer Engineering and Director of the Pattern Recognition and Applications laboratory (<https://pralab.diee.unica.it/>). He is partner and R&D manager of the company

Pluribus One that he co-founded (<https://www.pluribus-one.it>). He has been doing research on the design of pattern recognition and machine learning systems for thirty years. He was a very active organizer of international conferences and workshops, and established the popular workshop series on multiple classifier systems. Dr. Roli is Fellow of the IEEE and of the IAPR.

Supplementary Material

We report here some additional details on the attack optimization process and on the parameter settings used in our experiments, along with the complete list of the domain-knowledge constraints available for the considered datasets.

APPENDIX A

ATTACK OPTIMIZATION

Our attack optimizes Eq. (11) via projected gradient descent. Black-box attacks are non-adaptive, and thus ignore the defense mechanism. For this reason, the constraint loss term φ in our attack is ignored by setting its multiplier $\alpha = 0$ and $\kappa = \infty$. For white-box attacks on ANIMALS and PASCAL-PART, we set $\alpha = 0.1$ and $\alpha = 1$, respectively, while setting $\kappa = 2$. These values are chosen to appropriately scale the values of the constraint loss term φ w.r.t. the logit difference (i.e., the first term in Eq. 11, lower bounded by -2κ). This is required to have the sample misclassified while also fulfilling the domain-knowledge constraints. The process is better illustrated in Figs. 10 and 11, in which we respectively report the behavior of the black-box and white-box attack optimization on a single image from the ANIMALS dataset, with $\epsilon = 1$. In particular, in each Figure we report the source image, the (magnified) adversarial perturbation, and the resulting adversarial examples, along with some plots describing the optimization process, i.e., how the attack loss of Eq. (11) is minimized across iterations, and how the softmax-scaled outputs on the main classes and the logarithm of the constraint loss φ change accordingly.

In both the black-box and white-box cases, the attack loss is progressively reduced during the iterations of the optimization procedure. While the *albatross* prediction is progressively transformed into *ostrich*, the constraint loss increases across iterations, exceeding the rejection threshold. Thus, the adversarial example is correctly detected. Similarly, the white-box attack is able to initially flip the prediction from *albatross* to *ostrich*, allowing the constraint loss to increase. However, after this initial phase, the attack correctly reduces the constraint loss after its initial bump, bringing its value below the rejection threshold. The system thus fails to detect the corresponding adversarial example. Finally, it is also worth remarking that, in both cases, the final perturbations do not substantially compromise the source image content, remaining essentially imperceptible to the human eye.

APPENDIX B

DOMAIN KNOWLEDGE

Each dataset is composed of a set of classes that, for convenience, we associate to logic predicates. Such predicates participate in First-Order Logic (FOL) formulas that model the available domain knowledge. The FOL formulas that define the domain knowledge of the ANIMALS, CIFAR-100 and PASCAL-Part data are reported in Table 8, Table 9, and Table 11, respectively, where each predicate is indicated with capital letters. In each table (bottom part) we also report those rules that are about activating at least one of the classes of each level of the hierarchy. Following the

nomenclature used in the paper, the main classes of the ANIMALS dataset are ALBATROSS, GIRAFFE, CHEETAH, OSTRICH, PENGUIN, TIGER, ZEBRA, while the other categories are MAMMAL, HAIR, MILK, FEATHERS, BIRD, FLY, LAYEGGS, MEAT, CARNIVORE, POINT-EDTEETH, CLAWS, FORWARDERS, HOOF, UNGULATE, CUD, EVENTOED, TAWNY, BLACKSTRIPES, LONGLEGS, LONGNECK, DARKSPOTS, WHITE, BLACK, SWIM, BLACKWHITE, GOODFLIER. In the case of the CIFAR-100 dataset, the main classes are the ones associated with the predicates of Table 9 that belong to the premises of the shortest FOL formulas (i.e., the formulas in the form $A(x) \Rightarrow B(x)$, where the main class is A). Formulas in PASCAL-Part are relationships between objects and object-parts. The same part can belong to multiple objects, and in each objects several parts might be visible. See Table 11 for the list of classes (main classes are in the premises of the second block of formulas).

In ANIMALS and CIFAR-100, a mutual exclusion predicate is imposed on the main classes. As a matter of fact, in these two datasets, each image is only about a single main class. The $\text{mutual_excl}(p_1, p_2, \dots, p_n)$ predicate defined below, can be devised in different ways. The first, straightforward approach consists in considering the disjunction of the true cases in the truth table of the predicate:

$$\text{mutual_excl}(p_1, p_2, \dots, p_n) = \bigvee_{i=0}^n \left(p_i(x) \wedge \bigwedge_{j=0, j \neq i}^n \neg p_j(x) \right), \quad i, j \in M, \quad (12)$$

where M is the set of the main classes, with cardinality n and $p_i(x)$ is the logic predicate corresponding to the i -th output of the network $f_i(x)$. This formulation of the mutual_excl predicate is what we used in the ANIMALS dataset. When there are several classes, as in CIFAR-100, this formulation leads to optimization issues, since it turned out to be complicated to find a good balance between the effect of this constraint and the supervision-fitting term. For this reason, the mutual exclusivity in CIFAR-100 was defined as a disjunction of the main classes followed by a set of implications that are used to implement the mutual exclusion of the predicates,

$$\text{mutual_excl}(p_1, p_2, \dots, p_n) = \begin{cases} \bigvee_{i=0}^n p_i(x), \\ p_i(x) \Rightarrow \bigwedge_{j=0, j \neq i}^n \neg p_j(x), \quad \forall i \in M \end{cases} \quad (13)$$

that resulted easier to tune, since we have multiple soft constraints that could be eventually violated to accommodate the optimization procedure.

In the case of the ANIMALS dataset, we also considered a noisy setting in which we artificially altered the FOL rules of Table 8 in order to make them not fully coherent with the (real) domain knowledge. We describe the resulting noisy knowledge bases in Table 12, Table 13, and Table 14, reporting only the changes with respect to Table 8. The knowledge base of Table 12 has been obtained by altering four of the existing rules, while knowledge of Table 13 is the outcome of *adding* four new rules. In both the cases, we considered two implications whose conclusions are about main classes and two other implications whose conclusions are about auxiliary classes. Finally, Table 14 is about a noisy

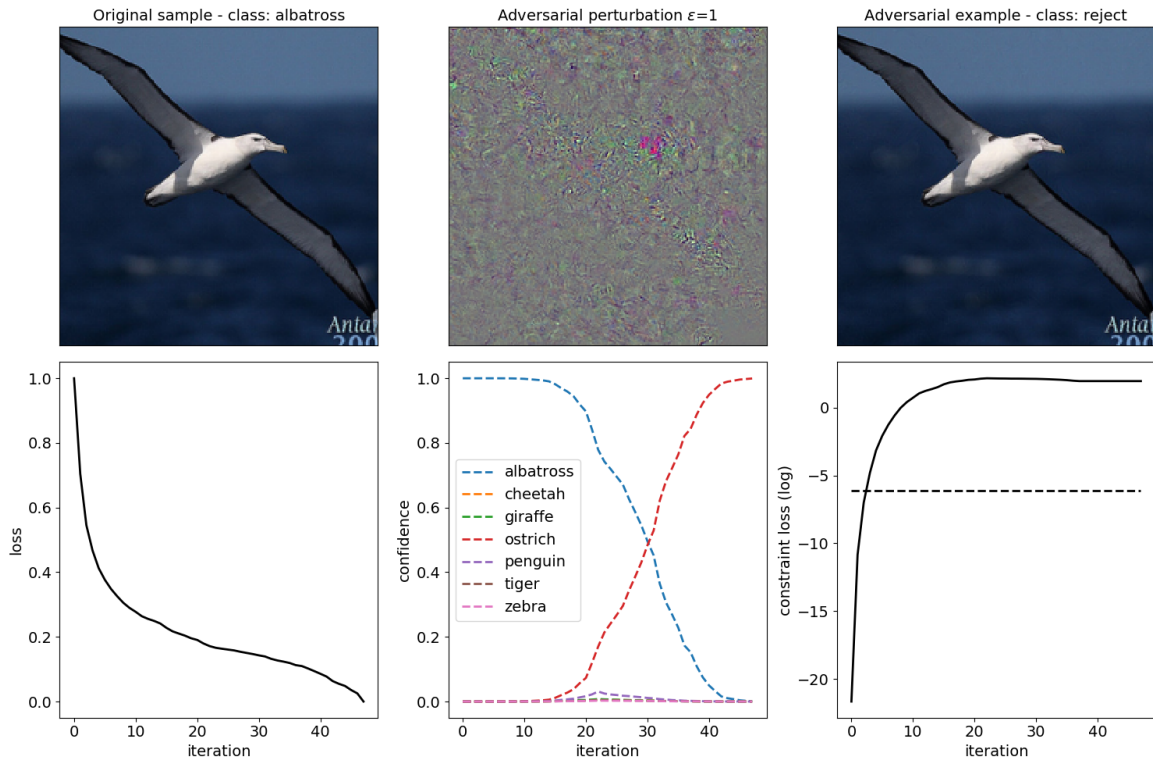


Fig. 10: Black-box attack on the ANIMALS dataset. While the attack is able to flip the initial prediction from *albatross* to *ostrich*, the attack is eventually detected as the constraint loss remains above the rejection threshold (dashed black line).




Fig. 11: White-box attack on the ANIMALS dataset. The attack is able to flip the initial prediction from *albatross* to *ostrich*, and then starts reducing the constraint loss which eventually falls below the rejection threshold (dashed black line). The attack sample remains thus undetected.

knowledge base where we relaxed the main-class-oriented conclusions of four implications. Such knowledge has been created by manually extending the conclusions using the disjunction operator, thus tolerating multiple configurations of the main classes.

APPENDIX C

RESULTS WITH SOTA ATTACK STRATEGIES

In order to better support the experimental analysis of Section 4.3, we report some examples of adversarial examples generated by the APGD-CE algorithm of the AutoAttack [47] library, ANIMALS dataset. In particular, in Fig. 12 and Fig. 13, we plot the two adversarial examples with highest supervision loss (and low constraint loss) and with highest constraint loss (and low supervision loss) (see also Fig. 8 of the main paper). No evident visual pattern is noticeable to distinguish the two cases.




Fig. 12: Adversarial examples with highest supervision loss (low constraint loss), APGD-CE attack, ANIMALS dataset.




Fig. 13: Adversarial examples with highest constraint loss (low supervision loss), APGD-CE attack, ANIMALS dataset.

TABLE 8: Domain knowledge, ANIMALS dataset.

$\forall x$	$\text{HAIR}(x) \Rightarrow \text{MAMMAL}(x)$
$\forall x$	$\text{MILK}(x) \Rightarrow \text{MAMMAL}(x)$
$\forall x$	$\text{FEATHER}(x) \Rightarrow \text{BIRD}(x)$
$\forall x$	$\text{FLY}(x) \wedge \text{LAYEGGS}(x) \Rightarrow \text{BIRD}(x)$
$\forall x$	$\text{MAMMAL}(x) \wedge \text{MEAT}(x) \Rightarrow \text{CARNIVORE}(x)$
$\forall x$	$\text{MAMMAL}(x) \wedge \text{POINTEDTEETH}(x) \wedge \text{CLAWS}(x) \wedge \text{FORWARDEYES}(x) \Rightarrow \text{CARNIVORE}(x)$
$\forall x$	$\text{MAMMAL}(x) \wedge \text{HOOFS}(x) \Rightarrow \text{UNGULATE}(x)$
$\forall x$	$\text{MAMMAL}(x) \wedge \text{CUD}(x) \Rightarrow \text{UNGULATE}(x)$
$\forall x$	$\text{MAMMAL}(x) \wedge \text{CUD}(x) \Rightarrow \text{EVENTOED}(x)$
$\forall x$	$\text{CARNIVORE}(x) \wedge \text{TAWNY}(x) \wedge \text{DARKSPOTS}(x) \Rightarrow \text{CHEETAH}(x)$
$\forall x$	$\text{CARNIVORE}(x) \wedge \text{TAWNY}(x) \wedge \text{BLACKSTRIPES}(x) \Rightarrow \text{TIGER}(x)$
$\forall x$	$\text{UNGULATE}(x) \wedge \text{LONGLEGS}(x) \wedge \text{LONGNECK}(x) \wedge \text{TAWNY}(x) \wedge \text{DARKSPOTS}(x) \Rightarrow \text{GIRAFFE}(x)$
$\forall x$	$\text{BLACKSTRIPES}(x) \wedge \text{UNGULATE}(x) \wedge \text{WHITE}(x) \Rightarrow \text{ZEBRA}(x)$
$\forall x$	$\text{BIRD}(x) \wedge \neg \text{FLY}(x) \wedge \text{LONGLEGS}(x) \wedge \text{LONGNECK}(x) \wedge \text{BLACK}(x) \Rightarrow \text{OSTRICH}(x)$
$\forall x$	$\text{BIRD}(x) \wedge \neg \text{FLY}(x) \wedge \text{SWIM}(x) \wedge \text{BLACKWHITE}(x) \Rightarrow \text{PENGUIN}(x)$
$\forall x$	$\text{BIRD}(x) \wedge \text{GOODFLIER}(x) \Rightarrow \text{ALBATROSS}(x)$
$\forall x$	$\text{mutual_excl}(\text{ALBATROSS}(x), \text{GIRAFFE}(x), \text{CHEETAH}(x), \text{OSTRICH}(x), \text{PENGUIN}(x), \text{TIGER}(x), \text{ZEBRA}(x))$
$\forall x$	$\text{MAMMAL}(x) \vee \text{HAIR}(x) \vee \text{MILK}(x) \vee \text{FEATHERS}(x) \vee \text{BIRD}(x) \vee \text{FLY}(x) \vee \text{LAYEGGS}(x) \vee \text{MEAT}(x)$
$\forall x$	$\text{CARNIVORE}(x) \vee \text{POINTEDTEETH}(x) \vee \text{CLAWS}(x) \vee \text{FORWARDEYES}(x) \vee \text{HOOFS}(x) \vee \text{UNGULATE}(x)$
$\forall x$	$\text{CUD}(x) \vee \text{EVENTOED}(x) \vee \text{TAWNY}(x) \vee \text{BLACKSTRIPES}(x) \vee \text{LONGLEGS}(x) \vee \text{LONGNECK}(x)$
$\forall x$	$\text{DARKSPOTS}(x) \vee \text{WHITE}(x) \vee \text{BLACK}(x) \vee \text{SWIM}(x) \vee \text{BLACKWHITE}(x) \vee \text{GOODFLIER}(x)$

TABLE 9: Domain knowledge, CIFAR-100 dataset.

$\forall x$	$\text{AQUATIC_MAMMALS}(x) \Rightarrow (\text{BEAVER}(x) \vee \text{DOLPHIN}(x) \vee \text{OTTER}(x) \vee \text{SEAL}(x) \vee \text{WHALE}(x))$
$\forall x$	$\text{BEAVER}(x) \Rightarrow \text{AQUATIC_MAMMALS}(x)$
$\forall x$	$\text{DOLPHIN}(x) \Rightarrow \text{AQUATIC_MAMMALS}(x)$
$\forall x$	$\text{OTTER}(x) \Rightarrow \text{AQUATIC_MAMMALS}(x)$
$\forall x$	$\text{SEAL}(x) \Rightarrow \text{AQUATIC_MAMMALS}(x)$
$\forall x$	$\text{WHALE}(x) \Rightarrow \text{AQUATIC_MAMMALS}(x)$
$\forall x$	$\text{FISH}(x) \Rightarrow (\text{AQUARIUM_FISH}(x) \vee \text{FLATFISH}(x) \vee \text{RAY}(x) \vee \text{SHARK}(x) \vee \text{TROUT}(x))$
$\forall x$	$\text{AQUARIUM_FISH}(x) \Rightarrow \text{FISH}(x)$
$\forall x$	$\text{FLATFISH}(x) \Rightarrow \text{FISH}(x)$
$\forall x$	$\text{RAY}(x) \Rightarrow \text{FISH}(x)$
$\forall x$	$\text{SHARK}(x) \Rightarrow \text{FISH}(x)$
$\forall x$	$\text{TROUT}(x) \Rightarrow \text{FISH}(x)$
$\forall x$	$\text{FLOWERS}(x) \Rightarrow (\text{ORCHID}(x) \vee \text{POPPY}(x) \vee \text{ROSE}(x) \vee \text{SUNFLOWER}(x) \vee \text{TULIP}(x))$
$\forall x$	$\text{ORCHID}(x) \Rightarrow \text{FLOWERS}(x)$
$\forall x$	$\text{POPPY}(x) \Rightarrow \text{FLOWERS}(x)$
$\forall x$	$\text{ROSE}(x) \Rightarrow \text{FLOWERS}(x)$
$\forall x$	$\text{SUNFLOWER}(x) \Rightarrow \text{FLOWERS}(x)$
$\forall x$	$\text{TULIP}(x) \Rightarrow \text{FLOWERS}(x)$
$\forall x$	$\text{FOOD_CONTAINERS}(x) \Rightarrow (\text{BOTTLE}(x) \vee \text{BOWL}(x) \vee \text{CAN}(x) \vee \text{CUP}(x) \vee \text{PLATE}(x))$
$\forall x$	$\text{BOTTLE}(x) \Rightarrow \text{FOOD_CONTAINERS}(x)$
$\forall x$	$\text{BOWL}(x) \Rightarrow \text{FOOD_CONTAINERS}(x)$
$\forall x$	$\text{CAN}(x) \Rightarrow \text{FOOD_CONTAINERS}(x)$
$\forall x$	$\text{CUP}(x) \Rightarrow \text{FOOD_CONTAINERS}(x)$
$\forall x$	$\text{PLATE}(x) \Rightarrow \text{FOOD_CONTAINERS}(x)$
$\forall x$	$\text{FRUIT_AND_VEGETABLES}(x) \Rightarrow (\text{APPLE}(x) \vee \text{MUSHROOM}(x) \vee \text{ORANGE}(x) \vee \text{PEAR}(x) \vee \text{SWEET_PEPPER}(x))$
$\forall x$	$\text{APPLE}(x) \Rightarrow \text{FRUIT_AND_VEGETABLES}(x)$
$\forall x$	$\text{MUSHROOM}(x) \Rightarrow \text{FRUIT_AND_VEGETABLES}(x)$
$\forall x$	$\text{ORANGE}(x) \Rightarrow \text{FRUIT_AND_VEGETABLES}(x)$
$\forall x$	$\text{PEAR}(x) \Rightarrow \text{FRUIT_AND_VEGETABLES}(x)$
$\forall x$	$\text{SWEET_PEPPER}(x) \Rightarrow \text{FRUIT_AND_VEGETABLES}(x)$
$\forall x$	$\text{HOUSEHOLD_ELECTRICAL_DEVICES}(x) \Rightarrow (\text{CLOCK}(x) \vee \text{KEYBOARD}(x) \vee \text{LAMP}(x) \vee \text{TELEPHONE}(x) \vee \text{TELEVISION}(x))$
$\forall x$	$\text{CLOCK}(x) \Rightarrow \text{HOUSEHOLD_ELECTRICAL_DEVICES}(x)$
$\forall x$	$\text{KEYBOARD}(x) \Rightarrow \text{HOUSEHOLD_ELECTRICAL_DEVICES}(x)$
$\forall x$	$\text{LAMP}(x) \Rightarrow \text{HOUSEHOLD_ELECTRICAL_DEVICES}(x)$
$\forall x$	$\text{TELEPHONE}(x) \Rightarrow \text{HOUSEHOLD_ELECTRICAL_DEVICES}(x)$
$\forall x$	$\text{TELEVISION}(x) \Rightarrow \text{HOUSEHOLD_ELECTRICAL_DEVICES}(x)$
$\forall x$	$\text{HOUSEHOLD_FURNITURE}(x) \Rightarrow (\text{BED}(x) \vee \text{CHAIR}(x) \vee \text{COUCH}(x) \vee \text{TABLE}(x) \vee \text{WARDROBE}(x))$
$\forall x$	$\text{BED}(x) \Rightarrow \text{HOUSEHOLD_FURNITURE}(x)$
$\forall x$	$\text{CHAIR}(x) \Rightarrow \text{HOUSEHOLD_FURNITURE}(x)$
$\forall x$	$\text{COUCH}(x) \Rightarrow \text{HOUSEHOLD_FURNITURE}(x)$
$\forall x$	$\text{TABLE}(x) \Rightarrow \text{HOUSEHOLD_FURNITURE}(x)$
$\forall x$	$\text{WARDROBE}(x) \Rightarrow \text{HOUSEHOLD_FURNITURE}(x)$
$\forall x$	$\text{INSECTS}(x) \Rightarrow (\text{BEE}(x) \vee \text{BEETLE}(x) \vee \text{BUTTERFLY}(x) \vee \text{CATERPILLAR}(x) \vee \text{COCKROACH}(x))$
$\forall x$	$\text{BEE}(x) \Rightarrow \text{INSECTS}(x)$
$\forall x$	$\text{BEETLE}(x) \Rightarrow \text{INSECTS}(x)$
$\forall x$	$\text{BUTTERFLY}(x) \Rightarrow \text{INSECTS}(x)$
$\forall x$	$\text{CATERPILLAR}(x) \Rightarrow \text{INSECTS}(x)$
$\forall x$	$\text{COCKROACH}(x) \Rightarrow \text{INSECTS}(x)$
$\forall x$	$\text{LARGE_CARNIVORES}(x) \Rightarrow (\text{BEAR}(x) \vee \text{LEOPARD}(x) \vee \text{LION}(x) \vee \text{TIGER}(x) \vee \text{WOLF}(x))$
$\forall x$	$\text{BEAR}(x) \Rightarrow \text{LARGE_CARNIVORES}(x)$
$\forall x$	$\text{LEOPARD}(x) \Rightarrow \text{LARGE_CARNIVORES}(x)$
$\forall x$	$\text{LION}(x) \Rightarrow \text{LARGE_CARNIVORES}(x)$
$\forall x$	$\text{TIGER}(x) \Rightarrow \text{LARGE_CARNIVORES}(x)$
$\forall x$	$\text{WOLF}(x) \Rightarrow \text{LARGE_CARNIVORES}(x)$
$\forall x$	$\text{LARGE_MAN-MADE_OUTDOOR_THINGS}(x) \Rightarrow (\text{BRIDGE}(x) \vee \text{CASTLE}(x) \vee \text{HOUSE}(x) \vee \text{ROAD}(x) \vee \text{SKYSCRAPER}(x))$
$\forall x$	$\text{BRIDGE}(x) \Rightarrow \text{LARGE_MAN-MADE_OUTDOOR_THINGS}(x)$
$\forall x$	$\text{CASTLE}(x) \Rightarrow \text{LARGE_MAN-MADE_OUTDOOR_THINGS}(x)$
$\forall x$	$\text{HOUSE}(x) \Rightarrow \text{LARGE_MAN-MADE_OUTDOOR_THINGS}(x)$
$\forall x$	$\text{ROAD}(x) \Rightarrow \text{LARGE_MAN-MADE_OUTDOOR_THINGS}(x)$
$\forall x$	$\text{SKYSCRAPER}(x) \Rightarrow \text{LARGE_MAN-MADE_OUTDOOR_THINGS}(x)$
$\forall x$	$\text{LARGE_NATURAL_OUTDOOR_SCENES}(x) \Rightarrow (\text{CLOUD}(x) \vee \text{FOREST}(x) \vee \text{MOUNTAIN}(x) \vee \text{PLAIN}(x) \vee \text{SEA}(x))$
$\forall x$	$\text{CLOUD}(x) \Rightarrow \text{LARGE_NATURAL_OUTDOOR_SCENES}(x)$
$\forall x$	$\text{FOREST}(x) \Rightarrow \text{LARGE_NATURAL_OUTDOOR_SCENES}(x)$
$\forall x$	$\text{MOUNTAIN}(x) \Rightarrow \text{LARGE_NATURAL_OUTDOOR_SCENES}(x)$
$\forall x$	$\text{PLAIN}(x) \Rightarrow \text{LARGE_NATURAL_OUTDOOR_SCENES}(x)$
$\forall x$	$\text{SEA}(x) \Rightarrow \text{LARGE_NATURAL_OUTDOOR_SCENES}(x)$
$\forall x$	$\text{LARGE_OMNIVORES_AND_HERBIVORES}(x) \Rightarrow (\text{CAMEL}(x) \vee \text{CATTLE}(x) \vee \text{CHIMPANZEE}(x) \vee \text{ELEPHANT}(x) \vee \text{KANGAROO}(x))$
$\forall x$	$\text{CAMEL}(x) \Rightarrow \text{LARGE_OMNIVORES_AND_HERBIVORES}(x)$

$\forall x$	$CATTLE(x) \Rightarrow LARGE_OMNIVORES_AND_HERBIVORES(x)$
$\forall x$	$CHIMPANZEE(x) \Rightarrow LARGE_OMNIVORES_AND_HERBIVORES(x)$
$\forall x$	$ELEPHANT(x) \Rightarrow LARGE_OMNIVORES_AND_HERBIVORES(x)$
$\forall x$	$KANGAROO(x) \Rightarrow LARGE_OMNIVORES_AND_HERBIVORES(x)$
$\forall x$	$MEDIUM_MAMMALS(x) \Rightarrow (FOX(x) \vee PORCUPINE(x) \vee POSSUM(x) \vee RACCOON(x) \vee SKUNK(x))$
$\forall x$	$FOX(x) \Rightarrow MEDIUM_MAMMALS(x)$
$\forall x$	$PORCUPINE(x) \Rightarrow MEDIUM_MAMMALS(x)$
$\forall x$	$POSSUM(x) \Rightarrow MEDIUM_MAMMALS(x)$
$\forall x$	$RACCOON(x) \Rightarrow MEDIUM_MAMMALS(x)$
$\forall x$	$SKUNK(x) \Rightarrow MEDIUM_MAMMALS(x)$
$\forall x$	$NON-INSECT_INVERTEBRATES(x) \Rightarrow (CRAB(x) \vee LOBSTER(x) \vee SNAIL(x) \vee SPIDER(x) \vee WORM(x))$
$\forall x$	$CRAB(x) \Rightarrow NON-INSECT_INVERTEBRATES(x)$
$\forall x$	$LOBSTER(x) \Rightarrow NON-INSECT_INVERTEBRATES(x)$
$\forall x$	$SNAIL(x) \Rightarrow NON-INSECT_INVERTEBRATES(x)$
$\forall x$	$SPIDER(x) \Rightarrow NON-INSECT_INVERTEBRATES(x)$
$\forall x$	$WORM(x) \Rightarrow NON-INSECT_INVERTEBRATES(x)$
$\forall x$	$PEOPLE(x) \Rightarrow (BABY(x) \vee MAN(x) \vee WOMAN(x) \vee BOY(x) \vee GIRL(x))$
$\forall x$	$BABY(x) \Rightarrow PEOPLE(x)$
$\forall x$	$BOY(x) \Rightarrow PEOPLE(x)$
$\forall x$	$GIRL(x) \Rightarrow PEOPLE(x)$
$\forall x$	$MAN(x) \Rightarrow PEOPLE(x)$
$\forall x$	$WOMAN(x) \Rightarrow PEOPLE(x)$
$\forall x$	$REPTILES(x) \Rightarrow (CROCODILE(x) \vee DINOSAUR(x) \vee LIZARD(x) \vee SNAKE(x) \vee TURTLE(x))$
$\forall x$	$CROCODILE(x) \Rightarrow REPTILES(x)$
$\forall x$	$DINOSAUR(x) \Rightarrow REPTILES(x)$
$\forall x$	$LIZARD(x) \Rightarrow REPTILES(x)$
$\forall x$	$SNAKE(x) \Rightarrow REPTILES(x)$
$\forall x$	$TURTLE(x) \Rightarrow REPTILES(x)$
$\forall x$	$SMALL_MAMMALS(x) \Rightarrow (HAMSTER(x) \vee MOUSE(x) \vee RABBIT(x) \vee SHREW(x) \vee SQUIRREL(x))$
$\forall x$	$HAMSTER(x) \Rightarrow SMALL_MAMMALS(x)$
$\forall x$	$MOUSE(x) \Rightarrow SMALL_MAMMALS(x)$
$\forall x$	$RABBIT(x) \Rightarrow SMALL_MAMMALS(x)$
$\forall x$	$SHREW(x) \Rightarrow SMALL_MAMMALS(x)$
$\forall x$	$SQUIRREL(x) \Rightarrow SMALL_MAMMALS(x)$
$\forall x$	$TREES(x) \Rightarrow (MAPLE_TREE(x) \vee OAK_TREE(x) \vee PALM_TREE(x) \vee PINE_TREE(x) \vee WILLOW_TREE(x))$
$\forall x$	$MAPLE_TREE(x) \Rightarrow TREES(x)$
$\forall x$	$OAK_TREE(x) \Rightarrow TREES(x)$
$\forall x$	$PALM_TREE(x) \Rightarrow TREES(x)$
$\forall x$	$PINE_TREE(x) \Rightarrow TREES(x)$
$\forall x$	$WILLOW_TREE(x) \Rightarrow TREE(x)$
$\forall x$	$VEHICLES1(x) \Rightarrow (BIKE(x) \vee BUS(x) \vee MOTORBIKE(x) \vee PICKUP_TRUCK(x) \vee TRAIN(x))$
$\forall x$	$BIKE(x) \Rightarrow VEHICLES1(x)$
$\forall x$	$BUS(x) \Rightarrow VEHICLES1(x)$
$\forall x$	$MOTORBIKE(x) \Rightarrow VEHICLES1(x)$
$\forall x$	$PICKUP(x) \Rightarrow VEHICLES1(x)$
$\forall x$	$TRAIN(x) \Rightarrow VEHICLES1(x)$
$\forall x$	$VEHICLES2(x) \Rightarrow (LAWN_MOWER(x) \vee ROCKET(x) \vee STREETCAR(x) \vee TANK(x) \vee TRACTOR(x))$
$\forall x$	$LAWN_MOWER(x) \Rightarrow VEHICLES2(x)$
$\forall x$	$ROCKET(x) \Rightarrow VEHICLES2(x)$
$\forall x$	$STREETCAR(x) \Rightarrow VEHICLES2(x)$
$\forall x$	$TANK(x) \Rightarrow VEHICLES2(x)$
$\forall x$	$TRACTOR(x) \Rightarrow VEHICLES2(x)$
$\forall x$	$mutual_excl(APPLE(x), AQUARIUM_FISH(x), BABY(x), BEAR(x), BEAVER(x), BED(x), BEE(x), BEETLE(x), BICYCLE(x), BOTTLE(x), BOWL(x), BOY(x), BRIDGE(x), BUS(x), BUTTERFLY(x), CAMEL(x), CAN(x), CASTLE(x), CATERPILLAR(x), CATTLE(x), CHAIR(x), CHIMPANZEE(x), CLOCK(x), CLOUD(x), COCKROACH(x), COUCH(x), CRAB(x), CROCODILE(x), CUP(x), DINOSAUR(x), DOLPHIN(x), ELEPHANT(x), FLATFISH(x), FOREST(x), FOX(x), GIRL(x), HAMSTER(x), HOUSE(x), KANGAROO(x), KEYBOARD(x), LAMP(x), LAWN_MOWER(x), LEOPARD(x), LION(x), LIZARD(x), LOBSTER(x), MAN(x), MAPLE_TREE(x), MOTORCYCLE(x), MOUNTAIN(x), MOUSE(x), MUSHROOM(x), OAK_TREE(x), ORANGE(x), ORCHID(x), OTTER(x), PALM_TREE(x), PEAR(x), PICKUP_TRUCK(x), PINE_TREE(x), PLAIN(x), PLATE(x), POPPY(x), PORCUPINE(x), POSSUM(x), RABBIT(x), RACCOON(x), RAY(x), ROAD(x), ROCKET(x), ROSE(x), SEA(x), SEAL(x), SHARK(x), SHREW(x), SKUNK(x) \vee SKYSCRAPER(x), SNAIL(x), SNAKE(x), SPIDER(x), SQUIRREL(x), STREETCAR(x), SUNFLOWER(x), SWEET_PEPPER(x), TABLE(x), TANK(x), TELEPHONE(x), TELEVISION(x), TIGER(x), TRACTOR(x), TRAIN(x), TROUT(x), TULIP(x), TURTLE(x), WARDROBE(x), WHALE(x), WILLOW_TREE(x), WOLF(x), WOMAN(x), WORM(x))$
$\forall x$	$mutual_excl(AQUATIC_MAMMALS(x), FISH(x), FLOWERS(x), FOOD_CONTAINERS(x), FRUIT_AND_VEGETABLES(x), HOUSEHOLD_ELECTRICAL(x), HOUSEHOLD_FURNITURE(x), INSECTS(x), LARGE_CARNIVORES(x), MAN-MADE_OUTDOOR(x), NATURAL_OUTDOOR_SCENES(x), OMNIVORES_AND_HERBIVORES(x), MEDIUM_MAMMALS(x),$

INVERTEBRATES(x) , PEOPLE(x) , REPTILES(x) , SMALL MAMMALS(x), TREES(x),
VEHICLES1(x), VEHICLES2(x))

TABLE 11: Domain knowledge, PASCAL-Part dataset.

$\forall x$	SCREEN(x) \Rightarrow (TVMONITOR)
$\forall x$	COACH(x) \Rightarrow (TRAIN(x))
$\forall x$	TORSO(x) \Rightarrow (PERSON(x) \vee HORSE(x) \vee COW(x) \vee DOG(x) \vee BIRD(x) \vee CAT(x) \vee SHEEP(x))
$\forall x$	LEG(x) \Rightarrow (PERSON(x) \vee HORSE(x) \vee COW(x) \vee DOG(x) \vee BIRD(x) \vee CAT(x) \vee SHEEP(x))
$\forall x$	HEAD(x) \Rightarrow (PERSON(x) \vee HORSE(x) \vee COW(x) \vee DOG(x) \vee BIRD(x) \vee CAT(x) \vee SHEEP(x))
$\forall x$	EAR(x) \Rightarrow (PERSON(x) \vee HORSE(x) \vee COW(x) \vee DOG(x) \vee CAT(x) \vee SHEEP(x))
$\forall x$	EYE(x) \Rightarrow (PERSON(x) \vee COW(x) \vee DOG(x) \vee BIRD(x) \vee CAT(x) \vee HORSE(x) \vee SHEEP(x))
$\forall x$	EBROW(x) \Rightarrow (PERSON(x))
$\forall x$	MOUTH(x) \Rightarrow (PERSON(x))
$\forall x$	HAIR(x) \Rightarrow (PERSON(x))
$\forall x$	NOSE(x) \Rightarrow (PERSON(x) \vee DOG(x) \vee CAT(x))
$\forall x$	NECK(x) \Rightarrow (PERSON(x) \vee HORSE(x) \vee COW(x) \vee DOG(x) \vee BIRD(x) \vee CAT(x) \vee SHEEP(x))
$\forall x$	ARM(x) \Rightarrow (PERSON(x))
$\forall x$	MUZZLE(x) \Rightarrow (HORSE(x) \vee COW(x) \vee DOG(x) \vee SHEEP(x))
$\forall x$	HOOF(x) \Rightarrow (HORSE(x))
$\forall x$	TAIL(x) \Rightarrow (HORSE(x) \vee COW(x) \vee DOG(x) \vee BIRD(x) \vee SHEEP(x) \vee CAT(x) \vee AEROPLANE(x))
$\forall x$	BOTTLE BODY(x) \Rightarrow (BOTTLE(x))
$\forall x$	PAW(x) \Rightarrow (DOG(x) \vee CAT(x))
$\forall x$	AEROPLANE BODY(x) \Rightarrow (AEROPLANE(x))
$\forall x$	WING(x) \Rightarrow (AEROPLANE(x) \vee BIRD(x))
$\forall x$	WHEEL(x) \Rightarrow (AEROPLANE(x) \vee CAR(x) \vee BICYCLE(x) \vee BUS(x) \vee MOTORBIKE(x))
$\forall x$	STERN(x) \Rightarrow (AEROPLANE(x))
$\forall x$	CAP(x) \Rightarrow (BOTTLE(x))
$\forall x$	HAND(x) \Rightarrow (PERSON(x))
$\forall x$	FRONTSIDE(x) \Rightarrow (CAR(x) \vee BUS(x) \vee TRAIN(x))
$\forall x$	RIGHTSIDE(x) \Rightarrow (CAR(x) \vee BUS(x) \vee TRAIN(x))
$\forall x$	ROOFSIDE(x) \Rightarrow (CAR(x) \vee BUS(x) \vee TRAIN(x))
$\forall x$	BACKSIDE(x) \Rightarrow (CAR(x) \vee BUS(x) \vee TRAIN(x))
$\forall x$	LEFTSIDE(x) \Rightarrow (CAR(x) \vee TRAIN(x) \vee BUS(x))
$\forall x$	DOOR(x) \Rightarrow (CAR(x) \vee BUS(x))
$\forall x$	MIRROR(x) \Rightarrow (CAR(x) \vee BUS(x))
$\forall x$	HEADLIGHT(x) \Rightarrow (CAR(x) \vee BUS(x) \vee TRAIN(x) \vee MOTORBIKE(x) \vee BICYCLE(x))
$\forall x$	MOTORBIKE(x) \Rightarrow (WHEEL(x) \vee HEADLIGHT(x) \vee HANDLEBAR(x) \vee SADDLE(x))
$\forall x$	WINDOW(x) \Rightarrow (CAR(x) \vee BUS(x))
$\forall x$	PLATE(x) \Rightarrow (CAR(x) \vee BUS(x))
$\forall x$	ENGINE(x) \Rightarrow (AEROPLANE(x))
$\forall x$	FOOT(x) \Rightarrow (PERSON(x) \vee BIRD(x))
$\forall x$	CHAINWHEEL(x) \Rightarrow (BICYCLE(x))
$\forall x$	SADDLE(x) \Rightarrow (BICYCLE(x) \vee MOTORBIKE(x))
$\forall x$	HANDLEBAR(x) \Rightarrow (BICYCLE(x) \vee MOTORBIKE(x))
$\forall x$	TRAIN HEAD(x) \Rightarrow (TRAIN(x))
$\forall x$	BEAK(x) \Rightarrow (BIRD(x))
$\forall x$	POT(x) \Rightarrow (POTTEDPLANT(x))
$\forall x$	PLANT(x) \Rightarrow (POTTEDPLANT(x))
$\forall x$	HORN(x) \Rightarrow (COW(x) \vee SHEEP(x))
$\forall x$	TVMONITOR(x) \Rightarrow (SCREEN(x))
$\forall x$	TRAIN(x) \Rightarrow (COACH(x) \vee LEFTSIDE(x) \vee TRAIN HEAD(x) \vee HEADLIGHT(x) \vee FRONTSIDE(x) \vee RIGHTSIDE(x) \vee BACKSIDE(x) \vee ROOFSIDE(x))
$\forall x$	PERSON(x) \Rightarrow (TORSO(x) \vee LEG(x) \vee HEAD(x) \vee EAR(x) \vee EYE(x) \vee EBROW(x) \vee MOUTH(x) \vee HAIR(x) \vee NOSE(x) \vee NECK(x) \vee ARM(x) \vee HAND(x) \vee FOOT(x))
$\forall x$	HORSE(x) \Rightarrow (HEAD(x) \vee EAR(x) \vee MUZZLE(x) \vee TORSO(x) \vee NECK(x) \vee LEG(x) \vee HOOF(x) \vee TAIL(x) \vee EYE(x))
$\forall x$	COW(x) \Rightarrow (HEAD(x) \vee EAR(x) \vee EYE(x) \vee MUZZLE(x) \vee TORSO(x) \vee NECK(x) \vee LEG(x) \vee TAIL(x) \vee HORN(x))
$\forall x$	BOTTLE(x) \Rightarrow (BOTTLE BODY(x) \vee CAP(x))
$\forall x$	DOG(x) \Rightarrow (HEAD(x) \vee EAR(x) \vee TORSO(x) \vee NECK(x) \vee LEG(x) \vee PAW(x) \vee EYE(x) \vee MUZZLE(x) \vee NOSE(x) \vee TAIL(x))
$\forall x$	AEROPLANE(x) \Rightarrow (AEROPLANE BODY(x) \vee WING(x) \vee WHEEL(x) \vee STERN(x) \vee ENGINE(x) \vee TAIL(x))
$\forall x$	CAR(x) \Rightarrow (FRONTSIDE(x) \vee RIGHTSIDE(x) \vee DOOR(x) \vee MIRROR(x) \vee HEADLIGHT(x) \vee WHEEL(x) \vee WINDOW(x) \vee PLATE(x) \vee ROOFSIDE(x) \vee BACKSIDE(x) \vee LEFTSIDE(x))
$\forall x$	BUS(x) \Rightarrow (PLATE(x) \vee FRONTSIDE(x) \vee RIGHTSIDE(x) \vee DOOR(x) \vee MIRROR(x) \vee HEADLIGHT(x) \vee WINDOW(x) \vee WHEEL(x) \vee LEFTSIDE(x) \vee BACKSIDE(x) \vee ROOFSIDE(x))
$\forall x$	BICYCLE(x) \Rightarrow (WHEEL(x) \vee CHAINWHEEL(x) \vee SADDLE(x) \vee HANDLEBAR(x) \vee HEADLIGHT(x))
$\forall x$	BIRD(x) \Rightarrow (HEAD(x) \vee EYE(x) \vee BEAK(x) \vee TORSO(x) \vee NECK(x) \vee LEG(x) \vee FOOT(x) \vee TAIL(x) \vee WING(x))
$\forall x$	CAT(x) \Rightarrow (HEAD(x) \vee EAR(x) \vee EYE(x) \vee NOSE(x) \vee TORSO(x) \vee NECK(x) \vee LEG(x) \vee PAW(x) \vee TAIL(x))
$\forall x$	MOTORBIKE(x) \Rightarrow (WHEEL(x) \vee HEADLIGHT(x) \vee HANDLEBAR(x) \vee SADDLE(x))
$\forall x$	SHEEP(x) \Rightarrow (HEAD(x) \vee EAR(x) \vee EYE(x) \vee MUZZLE(x) \vee TORSO(x) \vee NECK(x) \vee LEG(x) \vee TAIL(x) \vee HORN(x))
$\forall x$	POTTEDPLANT(x) \Rightarrow (POT(x) \vee PLANT(x))
$\forall x$	TVMONITOR(x) \vee TRAIN(x) \vee PERSON(x) \vee BOAT(x) \vee HORSE(x) \vee COW(x) \vee BOTTLE(x) \vee DOG(x) \vee AEROPLANE(x) \vee CAR(x) \vee BUS(x) \vee BICYCLE(x) \vee TABLE(x) \vee CHAIR(x) \vee BIRD(x) \vee CAT(x) \vee MOTORBIKE(x) \vee SHEEP(x) \vee SOFA(x) \vee POTTEDPLANT(x)

TABLE 12: First noisy domain knowledge ($\tilde{\mathcal{K}}_a$), ANIMALS dataset, obtained by altering the clean knowledge of Table 8. We report only the altered rules, highlighting the changes that make them not-coherent with the ANIMALS domain.

$\forall x$	$\text{FEATHER}(x) \Rightarrow \text{BIRD}(x) \text{ MAMMAL}(x)$
$\forall x$	$\text{MAMMAL}(x) \wedge \text{MEAT}(x) \text{ BIRD}(x) \Rightarrow \text{CARNIVORE}(x)$
$\forall x$	$\text{CARNIVORE}(x) \wedge \text{TAWNY}(x) \wedge \text{DARKSPOTS}(x) \Rightarrow \text{CHEETAH}(x)$
$\forall x$	$\text{BLACKSTRIPES}(x) \wedge \text{UNGULATE}(x) \wedge \text{WHITE}(x) \Rightarrow \text{ZEBRA}(x) \text{ TIGER}(x)$

TABLE 13: Second noisy domain knowledge ($\tilde{\mathcal{K}}_b$), ANIMALS dataset, obtained by adding new rules to the clean knowledge of Table 8. We report only the added rules, that were explicitly created to be not-coherent with the ANIMALS domain.

$\forall x$	$\text{FLY}(x) \Rightarrow \text{MAMMAL}(x)$
$\forall x$	$\text{MAMMAL}(x) \wedge \text{EVENTOED}(x) \Rightarrow \text{FEATHER}(x)$
$\forall x$	$\text{BLACKSTRIPES}(x) \wedge \text{WHITE}(x) \Rightarrow \text{PENGUIN}(x)$
$\forall x$	$\text{CARNIVORE}(x) \wedge \text{DARKSPOTS}(x) \Rightarrow \text{TIGER}(x)$

TABLE 14: Third noisy domain knowledge ($\tilde{\mathcal{K}}_c$), ANIMALS dataset, obtained by altering the clean knowledge of Table 8. We report only the altered rules, highlighting the changes that make them not-fully-coherent with the ANIMALS domain. They all involve main-class-oriented conclusions.

$\forall x$	$\text{CARNIVORE}(x) \wedge \text{TAWNY}(x) \wedge \text{DARKSPOTS}(x) \Rightarrow (\text{CHEETAH}(x) \vee \text{GIRAFFE}(x))$
$\forall x$	$\text{UNGULATE}(x) \wedge \text{LONGLEGS}(x) \wedge \text{LONGNECK}(x) \wedge \text{TAWNY}(x) \wedge \text{DARKSPOTS}(x) \Rightarrow (\text{GIRAFFE}(x) \vee \text{ZEBRA}(x))$
$\forall x$	$\text{BLACKSTRIPES}(x) \wedge \text{UNGULATE}(x) \wedge \text{WHITE}(x) \Rightarrow (\text{ZEBRA}(x) \vee \text{TIGER}(x))$
$\forall x$	$\text{BIRD}(x) \wedge \neg \text{FLY}(x) \wedge \text{SWIM}(x) \wedge \text{BLACKWHITE}(x) \Rightarrow (\text{PENGUIN}(x) \vee \text{OSTRICH}(x))$