

Towards Logical Specification of Statistical Machine Learning [★]

Yusuke Kawamoto¹[0000–0002–2151–9560]

AIST, Tsukuba, Japan

Abstract. We introduce a logical approach to formalizing statistical properties of machine learning. Specifically, we propose a formal model for statistical classification based on a Kripke model, and formalize various notions of classification performance, robustness, and fairness of classifiers by using epistemic logic. Then we show some relationships among properties of classifiers and those between classification performance and robustness, which suggests robustness-related properties that have not been formalized in the literature as far as we know. To formalize fairness properties, we define a notion of counterfactual knowledge and show techniques to formalize conditional indistinguishability by using counterfactual epistemic operators. As far as we know, this is the first work that uses logical formulas to express statistical properties of machine learning, and that provides epistemic (resp. counterfactually epistemic) views on robustness (resp. fairness) of classifiers.

Keywords: Epistemic logic · Possible world semantics · Divergence · Machine learning · Statistical classification · Robustness · Fairness

1 Introduction

With the increasing use of machine learning in real-life applications, the safety and security of learning-based systems have been of great interest. In particular, many recent studies [36, 8] have found vulnerabilities on the robustness of deep neural networks (DNNs) to malicious inputs, which can lead to disasters in security critical systems, such as self-driving cars. To find out these vulnerabilities in advance, there have been researches on the formal verification and testing methods for the robustness of DNNs in recent years [22, 25, 33, 37]. However, relatively little attention has been paid to the formal specification of machine learning [34].

To describe the formal specification of security properties, logical approaches have been shown useful to classify desired properties and to develop theories to compare those properties. For example, security policies in temporal systems

[★] This work was supported by JSPS KAKENHI Grant Number JP17K12667, by the New Energy and Industrial Technology Development Organization (NEDO), and by Inria under the project LOGIS.

have been formalized as trace properties [1] or hyperproperties [9], which characterize the relationships among various security policies. For another example, epistemic logic [39] has been widely used as formal policy languages (e.g., for the authentication [5] and the anonymity [35, 20] of security protocols, and for the privacy of social network [32]). As far as we know, however, no prior work has employed logical formulas to rigorously describe various statistical properties of machine learning, although there are some papers that (often informally) list various desirable properties of machine learning [34].

In this paper, we present a first logical formalization of statistical properties of machine learning. To describe the statistical properties in a simple and abstract way, we employ *statistical epistemic logic* (StatEL) [26], which is recently proposed to describe statistical knowledge and is applied to formalize statistical hypothesis testing and statistical privacy of databases.

A key idea in our modeling of statistical machine learning is that we formalize logical properties in the syntax level by using logical formulas, and statistical distances in the semantics level by using accessibility relations of a Kripke model [28]. In this model, we formalize statistical classifiers and some of their desirable properties: classification performance, robustness, and fairness. More specifically, classification performance and robustness are described as the differences between the classifier’s recognition and the correct label (e.g., given by the human), whereas fairness is formalized as the conditional indistinguishability between two groups or individuals by using a notion of counterfactual knowledge.

Our contributions. The main contributions of this work are as follows:

- We show a logical approach to formalizing statistical properties of machine learning in a simple and abstract way. In particular, we model logical properties in the syntax level, and statistical distances in the semantics level.
- We introduce a formal model for statistical classification. More specifically, we show how probabilistic behaviors of classifiers and non-deterministic adversarial inputs are formalized in a distributional Kripke model [26].
- We formalize the classification performance, robustness, and fairness of classifiers by using statistical epistemic logic (StatEL). As far as we know, this is the first work that uses logical formulas to formalize various statistical properties of machine learning, and that provides epistemic (resp. counterfactually epistemic) views on robustness (resp. fairness) of classifiers.
- We show some relationships among properties of classifiers, e.g., different strengths of robustness. We also present some relationships between classification performance and robustness, which suggest robustness-related properties that have not been formalized in the literature as far as we know.
- To formalize fairness properties, we define a notion of certain counterfactual knowledge and show techniques to formalize conditional indistinguishability by using counterfactual epistemic operators in StatEL. This enables us to express various fairness properties in a similar style of logical formulas.

Cautions and limitations. In this paper, we focus on formalizing properties of classification problems and do not deal with the properties of learning algorithms

(e.g., fairness through unawareness of sensitive attributes in data preparation), quality of training data (e.g., sample bias), quality of testing (e.g., coverage criteria), explainability, temporal properties, system level specification, or process agility in system development. It should be noted that most of the properties formalized in this paper have been known in literatures on machine learning, and the novelty of this work lies in the logical formulation of those statistical properties.

We also remark that this work does not provide methods for checking, guaranteeing, or improving the performance/robustness/fairness of machine learning. As for the satisfiability of logical formulas, we leave the development of testing and (statistical) model checking algorithms as future work, since the research area on the testing and formal/statistical verification of machine learning is relatively new and needs further techniques to improve the scalability. Moreover, in some applications such as image recognition, some formulas (e.g., representing whether an input image is panda or not) cannot be implemented mathematically, and require additional techniques based on experiments. Nevertheless, we demonstrate that describing various properties using logical formulas is useful to explore desirable properties and to discuss their relationships in a framework.

Finally, we emphasize that our work is the first attempt to use logical formulas to express statistical properties of machine learning, and would be a starting point to develop theories of specification of machine learning in future research.

Paper organization. The rest of this paper is organized as follows. Section 2 presents background on statistical epistemic logic (StatEL) and notations used in this paper. Section 3 defines counterfactual epistemic operators and shows techniques to model conditional indistinguishability using StatEL. Section 4 introduces a formal model for describing the behaviours of statistical classifiers and non-deterministic adversarial inputs. Sections 5, 6, and 7 respectively formalize the classification performance, robustness, and fairness of classifiers by using StatEL. Section 8 presents related work and Section 9 concludes.

2 Preliminaries

In this section we introduce some notations and recall the syntax and semantics of the *statistical epistemic logic* (StatEL) introduced in [26].

2.1 Notations

Let $\mathbb{R}^{\geq 0}$ be the set of non-negative real numbers, and $[0, 1]$ be the set of non-negative real numbers not greater than 1. We denote by \mathbb{DO} the set of all probability distributions over a set \mathcal{O} . Given a finite set \mathcal{O} and a probability distribution $\mu \in \mathbb{DO}$, the probability of sampling a value y from μ is denoted by $\mu[y]$. For a subset $R \subseteq \mathcal{O}$ we define $\mu[R]$ by: $\mu[R] = \sum_{y \in R} \mu[y]$. For a distribution μ over a finite set \mathcal{O} , its *support* is defined by $\text{supp}(\mu) = \{v \in \mathcal{O} : \mu[v] > 0\}$.

The *total variation distance* of two distributions $\mu, \mu' \in \mathbb{DO}$ is defined by: $D_{\text{tv}}(\mu \parallel \mu') \stackrel{\text{def}}{=} \sup_{R \subseteq \mathcal{O}} |\mu(R) - \mu'(R)|$.

2.2 Syntax of StatEL

We recall the syntax of the statistical epistemic logic (StatEL) [26], which has two levels of formulas: *static* and *epistemic formulas*. Intuitively, a static formula describes a proposition satisfied at a deterministic state, while an epistemic formula describes a proposition satisfied at a probability distribution of states. In this paper, the former is used only to define the latter.

Formally, let \mathbf{Mes} be a set of symbols called *measurement variables*, and Γ be a set of atomic formulas of the form $\gamma(x_1, x_2, \dots, x_n)$ for a predicate symbol γ , $n \geq 0$, and $x_1, x_2, \dots, x_n \in \mathbf{Mes}$. Let $I \subseteq [0, 1]$ be a finite union of disjoint intervals, and \mathcal{A} be a finite set of indices (e.g., associated with statistical divergences). Then the formulas are defined by:

$$\begin{aligned} \text{Static formulas: } \psi &::= \gamma(x_1, x_2, \dots, x_n) \mid \neg\psi \mid \psi \wedge \psi \\ \text{Epistemic formulas: } \varphi &::= \mathbb{P}_I \psi \mid \neg\varphi \mid \varphi \wedge \varphi \mid \psi \supset \varphi \mid \mathbf{K}_a \varphi \end{aligned}$$

where $a \in \mathcal{A}$. We denote by \mathcal{F} the set of all epistemic formulas. Note that we have no quantifiers over measurement variables. (See Section 2.4 for more details.)

The *probability quantification* $\mathbb{P}_I \psi$ represents that a static formula ψ is satisfied with a probability belonging to a set I . For instance, $\mathbb{P}_{(0.95, 1]} \psi$ represents that ψ holds with a probability greater than 0.95. By $\psi \supset \mathbb{P}_I \psi'$ we represent that the conditional probability of ψ' given ψ is included in a set I . The *epistemic knowledge* $\mathbf{K}_a \varphi$ expresses that we knows φ with a confidence specified by a .

As syntax sugar, we use *disjunction* \vee , *classical implication* \rightarrow , and *epistemic possibility* \mathbf{P}_a , defined as usual by: $\varphi_0 \vee \varphi_1 ::= \neg(\neg\varphi_0 \wedge \neg\varphi_1)$, $\varphi_0 \rightarrow \varphi_1 ::= \neg\varphi_0 \vee \varphi_1$, and $\mathbf{P}_a \varphi ::= \neg \mathbf{K}_a \neg\varphi$. When I is a singleton $\{i\}$, we abbreviate \mathbb{P}_I as \mathbb{P}_i .

2.3 Distributional Kripke Model

Next we recall the notion of a distributional Kripke model [26], where each possible world is a probability distribution over a set \mathcal{S} of states and each world w is associated with a stochastic assignment σ_w to measurement variables.

Definition 1 (Distributional Kripke model). Let \mathcal{A} be a finite set of indices (typically associated with statistical tests and their thresholds), \mathcal{S} be a finite set of states, and \mathcal{O} be a finite set of data. A *distributional Kripke model* is a tuple $\mathfrak{M} = (\mathcal{W}, (\mathcal{R}_a)_{a \in \mathcal{A}}, (V_s)_{s \in \mathcal{S}})$ consisting of:

- a non-empty set \mathcal{W} of probability distributions over a finite set \mathcal{S} of states;
- for each $a \in \mathcal{A}$, an accessibility relation $\mathcal{R}_a \subseteq \mathcal{W} \times \mathcal{W}$;
- for each $s \in \mathcal{S}$, a valuation V_s that maps each k -ary predicate γ to a set $V_s(\gamma) \subseteq \mathcal{O}^k$.

The set \mathcal{W} is called a *universe*, and its elements are called *possible worlds*. All measurement variables range over the same set \mathcal{O} in every world.

We assume that each $w \in \mathcal{W}$ is associated with a function $\rho_w : \mathbf{Mes} \times \mathcal{S} \rightarrow \mathcal{O}$ that maps each measurement variable x to its value $\rho_w(x, s)$ observed at a state s . We also assume that each state s in a world w is associated with the assignment $\sigma_s : \mathbf{Mes} \rightarrow \mathcal{O}$ defined by $\sigma_s(x) = \rho_w(x, s)$.

Since each world w is a distribution of states, we denote by $w[s]$ the probability that a state s is sampled from w . Then the probability that a measurement variable x has a value v is given by $\sigma_w(x)[v] = \sum_{s \in \text{supp}(w), \sigma_s(x)=v} w[s]$. This implies that, when a state s is drawn from w , an input $\sigma_s(x)$ is sampled from the distribution $\sigma_w(x)$.

2.4 Stochastic Semantics of StatEL

Now we recall the *stochastic semantics* [26] for the StatEL formulas over a distributional Kripke model $\mathfrak{M} = (\mathcal{W}, (\mathcal{R}_a)_{a \in \mathcal{A}}, (V_s)_{s \in \mathcal{S}})$ with $\mathcal{W} = \mathbb{DS}$.

The interpretation of static formulas ψ at a state s is given by:

$$\begin{aligned} s \models \gamma(x_1, x_2, \dots, x_k) & \text{ iff } (\sigma_s(x_1), \sigma_s(x_2), \dots, \sigma_s(x_k)) \in V_s(\gamma) \\ s \models \neg\psi & \text{ iff } s \not\models \psi \\ s \models \psi \wedge \psi' & \text{ iff } s \models \psi \text{ and } s \models \psi'. \end{aligned}$$

The *restriction* $w|_\psi$ of a world w to a static formula ψ is defined by $w|_\psi[s] = \frac{w[s]}{\sum_{s': s' \models \psi} w[s']}$ if $s \models \psi$, and $w|_\psi[s] = 0$ otherwise. Note that $w|_\psi$ is undefined if there is no state s that satisfies ψ and has a non-zero probability in w .

Then the interpretation of epistemic formulas in a world w is defined by:

$$\begin{aligned} \mathfrak{M}, w \models \mathbb{P}_I \psi & \text{ iff } \Pr \left[s \stackrel{\mathbb{S}}{\leftarrow} w : s \models \psi \right] \in I \\ \mathfrak{M}, w \models \neg\varphi & \text{ iff } \mathfrak{M}, w \not\models \varphi \\ \mathfrak{M}, w \models \varphi \wedge \varphi' & \text{ iff } \mathfrak{M}, w \models \varphi \text{ and } \mathfrak{M}, w \models \varphi' \\ \mathfrak{M}, w \models \psi \supset \varphi & \text{ iff } w|_\psi \text{ is defined and } \mathfrak{M}, w|_\psi \models \varphi \\ \mathfrak{M}, w \models \mathbb{K}_a \varphi & \text{ iff for every } w' \text{ s.t. } (w, w') \in \mathcal{R}_a, \mathfrak{M}, w' \models \varphi, \end{aligned}$$

where $s \stackrel{\mathbb{S}}{\leftarrow} w$ represents that a state s is sampled from the distribution w .

Then $\mathfrak{M}, w \models \psi_0 \supset \mathbb{P}_I \psi_1$ represents that the conditional probability of satisfying a static formula ψ_1 given another ψ_0 is included in a set I at a world w .

In each world w , measurement variables can be interpreted using σ_w . This allows us to assign different values to different occurrences of a variable in a formula; E.g., in $\varphi(x) \rightarrow \mathbb{K}_a \varphi'(x)$, x occurring in $\varphi(x)$ is interpreted by σ_w in a world w , while x in $\varphi'(x)$ is interpreted by $\sigma_{w'}$ in another w' s.t. $(w, w') \in \mathcal{R}_a$.

Finally, the interpretation of an epistemic formula φ in \mathfrak{M} is given by:

$$\mathfrak{M} \models \varphi \text{ iff for every world } w \text{ in } \mathfrak{M}, \mathfrak{M}, w \models \varphi.$$

3 Techniques for Conditional Indistinguishability

In this section we introduce some modal operators to define a notion of “counterfactual knowledge” using StatEL, and show how to employ them to formalize conditional indistinguishability properties. The techniques presented here are used to formalize some fairness properties of machine learning in Section 7.

3.1 Counterfactual Epistemic Operators

Let us consider an accessibility relation \mathcal{R}_ε based on a statistical divergence $D(\cdot \parallel \cdot) : \mathbb{D}\mathcal{O} \times \mathbb{D}\mathcal{O} \rightarrow \mathbb{R}^{\geq 0}$ and a threshold $\varepsilon \in \mathbb{R}^{\geq 0}$ defined by:

$$\mathcal{R}_\varepsilon \stackrel{\text{def}}{=} \{(w, w') \in \mathcal{W} \times \mathcal{W} \mid D(\sigma_w(y) \parallel \sigma_{w'}(y)) \leq \varepsilon\},$$

where y is the measurement variable observable in each world in \mathcal{W} . Intuitively, $(w, w') \in \mathcal{R}_\varepsilon$ represents that the probability distribution $\sigma_w(y)$ of the data y observed in a world w is *indistinguishable* from that in another world w' in terms of D .

Now we define the complement relation of \mathcal{R}_ε by $\overline{\mathcal{R}_\varepsilon} \stackrel{\text{def}}{=} (\mathcal{W} \times \mathcal{W}) \setminus \mathcal{R}_\varepsilon$, namely,

$$\overline{\mathcal{R}_\varepsilon} = \{(w, w') \in \mathcal{W} \times \mathcal{W} \mid D(\sigma_w(y) \parallel \sigma_{w'}(y)) > \varepsilon\}.$$

Then $(w, w') \in \overline{\mathcal{R}_\varepsilon}$ represents that the distribution $\sigma_w(y)$ observed in w *can be distinguished* from that in w' . Then the corresponding epistemic operator \mathbb{K}_ε , which we call a *counterfactual epistemic operator*, is interpreted as:

$$\mathfrak{M}, w \models \mathbb{K}_\varepsilon \varphi \text{ iff for every } w' \text{ s.t. } (w, w') \in \overline{\mathcal{R}_\varepsilon}, \text{ we have } \mathfrak{M}, w' \models \varphi \quad (1)$$

$$\text{iff for every } w' \text{ s.t. } \mathfrak{M}, w' \models \neg \varphi, \text{ we have } (w, w') \in \mathcal{R}_\varepsilon. \quad (2)$$

Intuitively, (1) represents that if we were located in a possible world w' that looked distinguished from the real world w , then φ would always hold. This means a *counterfactual knowledge*¹ in the sense that, if we had an observation different from the real world, then we would know φ . This is logically equivalent to (2), representing that all possible worlds w' that do not satisfy φ look indistinguishable from the real world w in terms of D .

We remark that the dual operator $\overline{\mathbb{P}}_\varepsilon$ is interpreted as:

$$\mathfrak{M}, w \models \overline{\mathbb{P}}_\varepsilon \varphi \text{ iff there exists a } w' \text{ s.t. } (w, w') \notin \mathcal{R}_\varepsilon \text{ and } \mathfrak{M}, w' \models \varphi. \quad (3)$$

This means a counterfactual possibility in the sense that it might be the case where we had an observation different from the real world and thought φ possible.

3.2 Conditional Indistinguishability via Counterfactual Knowledge

As shown in Section 7, some fairness notions in machine learning are based on conditional indistinguishability of the form (2), hence can be expressed using counterfactual epistemic operators.

Specifically, we use the following proposition, stating that given that two static formulas ψ and ψ' are respectively satisfied in worlds w and w' with probability 1, then the indistinguishability between w and w' can be expressed as $w \models \psi \supset \overline{\mathbb{P}}_a \mathbb{P}_1 \psi'$. Note that this formula means that there is no possible world where we have an observation different from the real world w (satisfying ψ) but we think ψ' possible; i.e., the formula means that if ψ' is satisfied then we have an observation indistinguishable from that in the real world w .

¹ Our definition of counterfactual knowledge is limited to the condition of having an observation different from the actual one. More general notions of counterfactual knowledge can be found in previous work (e.g., [38]).

Proposition 1 (Conditional indistinguishability) *Let $\mathfrak{M} = (\mathcal{W}, (\mathcal{R}_a)_{a \in \mathcal{A}}, (V_s)_{s \in \mathcal{S}})$ be a distributional Kripke model with the universe $\mathcal{W} = \mathbb{D}\mathcal{S}$. Let ψ and ψ' be static formulas, and $a \in \mathcal{A}$.*

- (i) $\mathfrak{M} \models \psi \supset \neg \overline{\mathbb{P}}_a \mathbb{P}_1 \psi'$ iff for any $w, w' \in \mathcal{W}$, $\mathfrak{M}, w \models \mathbb{P}_1 \psi$ and $\mathfrak{M}, w' \models \mathbb{P}_1 \psi'$ imply $(w, w') \in \mathcal{R}_a$.
- (ii) If \mathcal{R}_a is symmetric, then $\mathfrak{M} \models \psi \supset \neg \overline{\mathbb{P}}_a \mathbb{P}_1 \psi'$ iff $\mathfrak{M} \models \psi' \supset \neg \overline{\mathbb{P}}_a \mathbb{P}_1 \psi$.

See Appendix A for the proof.

4 Formal Model for Statistical Classification

In this section we introduce a formal model for statistical classification by using distributional Kripke models (Definition 1). In particular, we formalize a probabilistic behaviour of a classifier C and a non-deterministic input x from an adversary in a distributional Kripke model.

4.1 Statistical Classification Problems

Multiclass classification is the problem of classifying a given input into one of multiple classes. Let \mathbf{L} be a finite set of *class labels*, and \mathcal{D} be a finite set of *input data* (called *feature vectors*) that we want to classify. Then a *classifier* is a function $C : \mathcal{D} \rightarrow \mathbf{L}$ that receives an input datum and predicts which class (among \mathbf{L}) the input belongs to. Here we do *not* model how classifiers are constructed from a set of training data, but deal with a situation where some classifier C has already been obtained and its properties should be evaluated.

Let $f : \mathcal{D} \times \mathbf{L} \rightarrow \mathbb{R}$ be a *scoring function* that gives a score $f(v, \ell)$ of predicting the class of an input datum (feature vector) v as a label ℓ . Then for each input $v \in \mathcal{D}$, we denote by $H(v) = \ell$ to represent that a label ℓ maximizes $f(v, \ell)$. For example, when the input v is an image of an animal and ℓ is the animal's name, $H(v) = \ell$ may represent that an oracle (or “human”) classifies the image v as ℓ .

4.2 Modeling the Behaviours of Classifiers

Classifiers are formalized on a distributional Kripke model $\mathfrak{M} = (\mathcal{W}, (\mathcal{R}_a)_{a \in \mathcal{A}}, (V_s)_{s \in \mathcal{S}})$ with $\mathcal{W} = \mathbb{D}\mathcal{S}$ and a real world $w_{\text{real}} \in \mathcal{W}$. Recall that each world $w \in \mathcal{W}$ is a probability distribution over the set \mathcal{S} of states and has a stochastic assignment $\sigma_w : \mathbf{Mes} \rightarrow \mathbb{D}\mathcal{O}$ that is consistent with the deterministic assignments σ_s for all $s \in \mathcal{S}$ (as explained in Section 2.3).

We present an overview of our formalization in Fig. 1. We denote by $x \in \mathbf{Mes}$ an input datum given to the classifier C (and to the oracle H), by $y \in \mathbf{Mes}$ a correct label given by the oracle H , and by $\hat{y} \in \mathbf{Mes}$ a label predicted by C . We assume that the input variable x (resp. the output variables y, \hat{y}) ranges over the set \mathcal{D} of input data (resp. the set \mathbf{L} of labels); i.e., the deterministic assignment σ_s at each state $s \in \mathcal{S}$ has the range $\mathcal{O} = \mathcal{D} \cup \mathbf{L}$ and satisfies $\sigma_s(x) \in \mathcal{D}$ and $\sigma_s(y), \sigma_s(\hat{y}) \in \mathbf{L}$.




Fig. 1: A world w is chosen non-deterministically. With probability $w[s_i]$, the world w is in a deterministic state s_i where the classifier C receives the input value $\sigma_{s_i}(x)$ and returns the output value $\sigma_{s_i}(\hat{y})$.

A key idea in our modeling is that we formalize logical properties in the syntax level by using logical formulas, and statistical distances in the semantics level by using accessibility relations \mathcal{R}_a . In this way, we can formalize various statistical properties of classifiers in a simple and abstract way.

To formalize a classifier C , we introduce a static formula $\psi(x, \hat{y})$ to represent that C classifies a given input x as a class \hat{y} . We also introduce a static formula $h(x, y)$ to represent that y is the actual class of an input x . As an abbreviation, we write $\psi_\ell(x)$ (resp. $h_\ell(x)$) to denote $\psi(x, \ell)$ (resp. $h(x, \ell)$). Formally, these static formulas are interpreted at each state $s \in \mathcal{S}$ as follows:

$$\begin{aligned} s \models \psi(x, \hat{y}) &\text{ iff } C(\sigma_s(x)) = \sigma_s(\hat{y}). \\ s \models h(x, y) &\text{ iff } H(\sigma_s(x)) = \sigma_s(y). \end{aligned}$$

4.3 Modeling the Non-deterministic Inputs from Adversaries

As explained in Section 2.3, when a state s is drawn from a distribution $w \in \mathcal{W}$, an input value $\sigma_s(x)$ is sampled from the distribution $\sigma_w(x)$, and assigned to the measurement variable x . Since x denotes the input to the classifier C , the input distribution $\sigma_w(x)$ over \mathcal{D} can be regarded as the *test dataset*. This means that each world w corresponds to a test dataset $\sigma_w(x)$. For instance, $\sigma_{w_{\text{real}}}(x)$ in the real world w_{real} represents the actual test dataset. The set of all possible test datasets (i.e., possible distributions of inputs to C) is represented by $\Lambda \stackrel{\text{def}}{=} \{\sigma_w(x) \mid w \in \mathcal{W}\}$. Note that Λ can be an infinite set.

For example, let us consider testing the classifier C with the actual test dataset $\sigma_{w_{\text{real}}}(x)$. When C assigns a label ℓ to an input x with probability 0.2, i.e., $\Pr\left[v \stackrel{\$}{\leftarrow} \sigma_{w_{\text{real}}}(x) : C(v) = \ell\right] = 0.2$, then this can be expressed by:

$$\mathfrak{M}, w_{\text{real}} \models \mathbb{P}_{0.2} \psi_\ell(x).$$

We can also formalize a non-deterministic input x from an adversary in this model as follows. Although each state s in a possible world w is assigned the probability $w[s]$, each possible world w itself is not assigned a probability. Thus, each input distribution $\sigma_w(x) \in \Lambda$ itself is also not assigned a probability, hence our model assumes no probability distribution over Λ . In other words, we assume that a world w and thus an adversary’s input distribution $\sigma_w(x)$ are non-deterministically chosen. This is useful to model an adversary’s malicious inputs in the definitions of security properties, because we usually do not have a prior knowledge of the distribution of malicious inputs from adversaries, and need to reason about the worst cases caused by the attack. In Section 6, this formalization of non-deterministic inputs is used to express the robustness of classifiers.

Finally, it should be noted that we cannot enumerate all possible adversarial inputs, hence cannot construct \mathcal{W} by collecting their corresponding worlds. Since \mathcal{W} can be an infinite set and is unspecified, we do not aim at checking whether or not a formula is satisfied in all possible worlds of \mathcal{W} . Nevertheless, as shown in later sections, describing various properties using StatEL is useful to explore desirable properties and to discuss relationships among them.

5 Formalizing the Classification Performance

In this section we show a formalization of classification performance using StatEL (See Fig. 2 for basic ideas). In classification problems, the terms *positive/negative* represent the result of the classifier’s prediction, and the terms *true/false* represent whether the classifier predicts correctly or not. Then the following terminologies are commonly used:

- (*tp*) *true positive* means both the prediction and actual class are positive;
- (*tn*) *true negative* means both the prediction and actual class are negative;
- (*fp*) *false positive* means the prediction is positive but the actual class is negative;
- (*fn*) *false negative* means the prediction is negative but the actual class is positive.

These terminologies can be formalized using StatEL as shown in Table 1. For example, when an input x shows true positive at a state s , this can be expressed as $s \models \psi_\ell(x) \wedge h_\ell(x)$. True negative, false positive (Type I error), and false negative (Type II error) are respectively expressed as $s \models \neg\psi_\ell(x) \wedge \neg h_\ell(x)$, $s \models \psi_\ell(x) \wedge \neg h_\ell(x)$, and $s \models \neg\psi_\ell(x) \wedge h_\ell(x)$.

Then *precision (positive predictive value)* is defined as the conditional probability that the prediction is correct given that the prediction is positive; i.e., $\text{precision} = \frac{tp}{tp+fp}$. Since the test dataset distribution in the real world w_{real} is expressed as $\sigma_{w_{\text{real}}}(x)$ (as explained in Section 4.3), the precision being within an interval I is given by:

$$\Pr \left[v \stackrel{\$}{\leftarrow} \sigma_{w_{\text{real}}}(x) : H(v) = \ell \mid C(v) = \ell \right] \in I,$$

which can be written as:

$$\Pr \left[s \stackrel{\$}{\leftarrow} w_{\text{real}} : s \models h_\ell(x) \mid s \models \psi_\ell(x) \right] \in I.$$

Table 1: Logical description of the table of confusion

	Actual class		Prevalence $_{\ell,I}(x) \stackrel{\text{def}}{=} \mathbb{P}_I(h_\ell(x))$	Accuracy $_{\ell,I}(x) \stackrel{\text{def}}{=} \mathbb{P}_I(\psi_\ell(x) \leftrightarrow h_\ell(x))$
	positive $h_\ell(x)$	negative $\neg h_\ell(x)$		
Positive prediction $\psi_\ell(x)$	$tp(x) \stackrel{\text{def}}{=} \psi_\ell(x) \wedge h_\ell(x)$	$fp(x) \stackrel{\text{def}}{=} \psi_\ell(x) \wedge \neg h_\ell(x)$	Precision $_{\ell,I}(x) \stackrel{\text{def}}{=} \psi_\ell(x) \supset \mathbb{P}_I h_\ell(x)$	FDR $_{\ell,I}(x) \stackrel{\text{def}}{=} \psi_\ell(x) \supset \mathbb{P}_I \neg h_\ell(x)$
Negative prediction $\neg \psi_\ell(x)$	$fn(x) \stackrel{\text{def}}{=} \neg \psi_\ell(x) \wedge h_\ell(x)$	$tn(x) \stackrel{\text{def}}{=} \neg \psi_\ell(x) \wedge \neg h_\ell(x)$	FOR $_{\ell,I}(x) \stackrel{\text{def}}{=} \neg \psi_\ell(x) \supset \mathbb{P}_I h_\ell(x)$	NPV $_{\ell,I}(x) \stackrel{\text{def}}{=} \neg \psi_\ell(x) \supset \mathbb{P}_I \neg h_\ell(x)$
	Recall $_{\ell,I}(x) \stackrel{\text{def}}{=} h_\ell(x) \supset \mathbb{P}_I \psi_\ell(x)$	FallOut $_{\ell,I}(x) \stackrel{\text{def}}{=} \neg h_\ell(x) \supset \mathbb{P}_I \psi_\ell(x)$		
	MissRate $_{\ell,I}(x) \stackrel{\text{def}}{=} h_\ell(x) \supset \mathbb{P}_I \neg \psi_\ell(x)$	Specificity $_{\ell,I}(x) \stackrel{\text{def}}{=} \neg h_\ell(x) \supset \mathbb{P}_I \neg \psi_\ell(x)$		

By using StatEL, this can be formalized as:

$$\mathfrak{M}, w_{\text{real}} \models \text{Precision}_{\ell,I}(x) \text{ where } \text{Precision}_{\ell,I}(x) \stackrel{\text{def}}{=} \psi_\ell(x) \supset \mathbb{P}_I h_\ell(x). \quad (4)$$

Note that the precision depends on the test data sampled from the distribution $\sigma_{w_{\text{real}}}(x)$, hence on the real world w_{real} in which we are located. Hence the measurement variable x in $\text{Precision}_{\ell,I}(x)$ is interpreted using the stochastic assignment $\sigma_{w_{\text{real}}}$ in the world w_{real} .

Symmetrically, *recall (true positive rate)* is defined as the conditional probability that the prediction is correct given that the actual class is positive; i.e., $\text{recall} = \frac{tp}{tp+fn}$. Then the recall being within I is formalized as:

$$\text{Recall}_{\ell,I}(x) \stackrel{\text{def}}{=} h_\ell(x) \supset \mathbb{P}_I \psi_\ell(x). \quad (5)$$

In Table 1 we show the formalization of other notions of classification performance using StatEL.

6 Formalizing the Robustness of Classifiers

Many studies have found attacks on the robustness of statistical machine learning [8]. An input data that violates the robustness of classifiers is called an *adversarial example* [36]. It is designed to make a classifier fail to predict the actual class ℓ , but is recognized to belong to ℓ from human eyes. For example, in computer vision, Goodfellow et al. [18] create an image by adding undetectable noise to a panda's photo so that humans can still recognize the perturbed image as a panda, but a classifier misclassifies it as a gibbon.

In this section we formalize robustness notions for classifiers by using epistemic operators in StatEL (See Fig. 2 for an overview of the formalization). In addition, we present some relationships between classification performance and robustness, which suggest robustness-related properties that have not been formalized in the literature as far as we know.

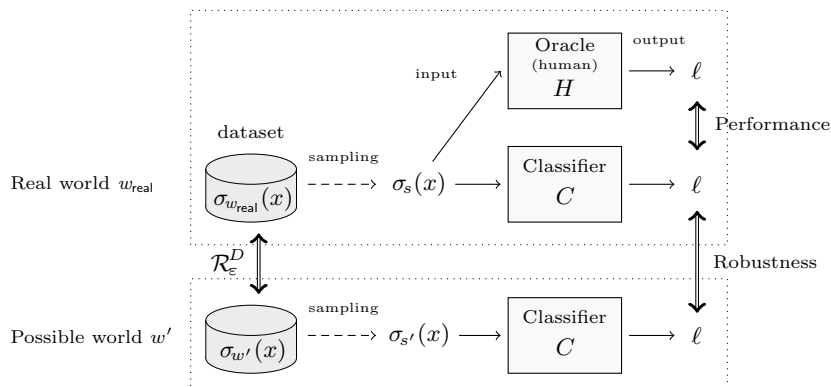


Fig. 2: The classification performance compares the conditional probability of the human H 's output with that by the classifier C 's. On the other hand, the robustness compares the conditional probability in the real world w_{real} with that in a possible world w' that is close to w_{real} in terms of \mathcal{R}_ϵ^D . Note that an adversary's choice of the test dataset $\sigma_{w'}(x)$ is formalized by the non-deterministic choice of the possible world w' .

6.1 Total Correctness of Classifiers

We first note that the *total correctness* of classifiers could be formalized as a classification performance (e.g., precision, recall, or accuracy) in the presence of all possible inputs from adversaries. For example, the total correctness could be formalized as $\mathfrak{M} \models \text{Recall}_{\ell, I}(x)$, which represents that $\text{Recall}_{\ell, I}(x)$ satisfies in all possible worlds of \mathfrak{M} .

In practice, however, it is not possible or tractable to check whether the classification performance is achieved for all possible dataset and for all possible inputs, e.g., when \mathcal{W} is an infinite set. Hence we need a weaker form of correctness notions, which may be tested in a certain way. In the following sections, we deal with robustness notions that are weaker than total correctness.

6.2 Probabilistic Robustness against Targeted Attacks

When a robustness attack aims at misclassifying an input as a specific target label, then it is called a *targeted attack*. For instance, in the above-mentioned attack by [18], a gibbon is the target into which a panda's photo is misclassified.

To formalize the robustness, let $\mathcal{R}_\epsilon^D \subseteq \mathcal{W} \times \mathcal{W}$ be an accessibility relation that relates two worlds having closer inputs, i.e.,

$$\mathcal{R}_\epsilon^D \stackrel{\text{def}}{=} \{(w, w') \in \mathcal{W} \times \mathcal{W} \mid D(\sigma_w(x) \parallel \sigma_{w'}(x)) \leq \epsilon\},$$

where D is some divergence or distance. Intuitively, $(w, w') \in \mathcal{R}_\epsilon^D$ implies that the two distributions $\sigma_w(x)$ and $\sigma_{w'}(x)$ of inputs to the classifier C represent close datasets in terms of D (e.g., two test datasets consisting of slightly different

images that look pandas from the human’s eyes). Then an epistemic formula $K_\varepsilon^D \varphi$ represents that we are confident that φ is true as far as the classifier C classifies the test data that are perturbed by noise of a level ε or smaller².

Now we discuss how we formalize robustness using the epistemic operator K_ε^D as follows. A first definition of robustness against targeted attacks might be:

$$\mathfrak{M}, w_{\text{real}} \models h_{\text{panda}}(x) \supset K_\varepsilon^D \mathbb{P}_0 \psi_{\text{gibbon}}(x),$$

which represents that a panda’s photo x will not be recognized as a gibbon at all after the photo is perturbed by noise. However, this does not express probability or cover the case where the human cannot recognize the perturbed image as a panda, for example, when the image is perturbed by a transformation such as linear displacement, rescaling and rotation [2]. Instead, for some $\delta \in [0, 1]$, we formalize a notion of *probabilistic robustness against targeted attacks* by:

$$\text{TargetRobust}_{\text{panda},\delta}(x, \text{gibbon}) \stackrel{\text{def}}{=} K_\varepsilon^D (h_{\text{panda}}(x) \supset \mathbb{P}_{[0,\delta]} \psi_{\text{gibbon}}(x)).$$

Since L^p -norms are often regarded as reasonable approximations of human perceptual distances [6], they are used as distance constraints on the perturbation in many researches on targeted attacks (e.g. [36, 18, 6]). To represent the robustness against these attacks in our model, we should take the metric D as the ∞ -Wasserstein distance W_d (in terms of the L^p metric d) between the two distributions $\sigma_w(x)$ and $\sigma_{w'}(x)$ ³.

6.3 Probabilistic Robustness against Non-Targeted Attacks

Next we formalize *non-targeted attacks* [31, 30] in which adversaries try to misclassify inputs as some arbitrary incorrect labels (i.e., not as a specific label like a gibbon). Compared to targeted attacks, this kind of attacks are easier to mount, but harder to defend.

A notion of *probabilistic robustness against non-targeted attacks* can be formalized for some $I = [1 - \delta, 1]$ by:

$$\text{TotalRobust}_{\ell,I}(x) \stackrel{\text{def}}{=} K_\varepsilon^D (h_\ell(x) \supset \mathbb{P}_I \psi_\ell(x)) = K_\varepsilon^D \text{Recall}_{\ell,I}(x). \quad (6)$$

Then we derive that $\text{TotalRobust}_{\text{panda},I}(x)$ implies $\text{TargetRobust}_{\text{panda},\delta}(x, \text{gibbon})$, namely, robustness against non-targeted attacks is not weaker than robustness against targeted attacks.

Next we note that by (6), robustness can be regarded as recall in the presence of perturbed noise. This implies that for each property φ in Table 1, we could consider $K_\varepsilon^D \varphi$ as a property related to robustness although these have not been formalized in the literature of robustness of machine learning as far as

² This usage of modality relies on the fact that the value of the measurement variable x can be different in different possible worlds.

³ A coupling that achieves $W_d(\sigma_w(x), \sigma_{w'}(x)) \leq \varepsilon$ provides a transformation of an image in $\text{supp}(\sigma_w(x))$ to another in $\text{supp}(\sigma_{w'}(x))$ perturbed by a level ε of noise.

we recognize. For example, $\mathsf{K}_\varepsilon^D \text{Precision}_{\ell,i}(x)$ represents that in the presence of perturbed noise, the prediction is correct with a probability i given that it is positive. For another example, $\mathsf{K}_\varepsilon^D \text{Accuracy}_{\ell,i}(x)$ represents that in the presence of perturbed noise, the prediction is correct (whether it is positive or negative) with a probability i .

Finally, note that by the reflexivity of $\mathcal{R}_\varepsilon^D$, $\mathfrak{M}, w_{\text{real}} \models \mathsf{K}_\varepsilon^D \text{Recall}_{\ell,I}(x)$ implies $\mathfrak{M}, w_{\text{real}} \models \text{Recall}_{\ell,I}(x)$, i.e., robustness implies recall without perturbation noise.

7 Formalizing the Fairness of Classifiers

There have been researches on various notions of fairness in machine learning. In this section, we formalize a few notions of fairness of classifiers by using StatEL. Here we focus on the fairness that should be maintained in the *impact*, i.e., the results of classification, rather than the *treatment*⁴.

To formalize fairness notions, we use a distributional Kripke model $\mathfrak{M} = (\mathcal{W}, (\mathcal{R}_a)_{a \in \mathcal{A}}, (V_s)_{s \in \mathcal{S}})$ where \mathcal{W} includes a possible world w_d having a dataset d from which an input to the classifier C is drawn. Recall that x , y , and \hat{y} are measurement variables denoting the input to the classifier C , the actual class label, and the predicted label by C , respectively. In each world w , $\sigma_w(x)$ is the distribution of C 's input over \mathcal{D} , (i.e., the test data distribution), $\sigma_w(y)$ is the distribution of the actual label over \mathbf{L} , and $\sigma_w(\hat{y})$ is the distribution of C 's output over \mathbf{L} . For each group $G \subseteq \mathcal{D}$ of inputs, we introduce a static formula $\eta_G(x)$ representing that an input x belongs to G . We also introduce a formula ξ_d representing that all data are drawn from some subset of the dataset d . Formally, these are interpreted by:

- For each state $s \in \mathcal{S}$, $s \models \eta_G(x)$ iff $\sigma_s(x) \in G$;
- For each world $w \in \mathcal{W}$, $w \models \xi_d$ iff there exists a $\mathcal{S}' \subseteq \mathcal{S}$ s.t. $w[s] = \frac{w_d[s]}{\sum_{s' \in \mathcal{S}'} w_d[s']}$ if $s \in \mathcal{S}'$, and $w[s] = 0$ otherwise.

For two worlds w and w' , we write $w \models \mathbb{Q}_{w'} \psi$ to denote that $w \models \mathbb{P}_1 \psi$ and $s \not\models \psi$ for all $s \in \text{supp}(w') \setminus \text{supp}(w)$.

Then we obtain the following proposition on conditional indistinguishability.

Proposition 2 (Conditional indistinguishability in a world w_d) *Let $\mathfrak{M} = (\mathcal{W}, (\mathcal{R}_a)_{a \in \mathcal{A}}, (V_s)_{s \in \mathcal{S}})$ be a distributional Kripke model with the universe $\mathcal{W} = \mathbb{D}\mathcal{S}$. Let w_d be a world with a dataset d , ψ and ψ' be static formulas, and $a \in \mathcal{A}$.*

- (i) $\mathfrak{M}, w_d \models \psi \supset \neg \overline{\mathbb{P}}_a(\xi_d \wedge \mathbb{Q}_{w_d} \psi')$ iff for any $w, w' \in \mathcal{W}$, $\mathfrak{M}, w \models \xi_d \wedge \mathbb{Q}_{w_d} \psi$ and $\mathfrak{M}, w' \models \xi_d \wedge \mathbb{Q}_{w_d} \psi'$ imply $(w, w') \in \mathcal{R}_a$.
- (ii) If \mathcal{R}_a is symmetric, then $\mathfrak{M}, w_d \models \psi \supset \neg \overline{\mathbb{P}}_a \mathbb{Q}_{w_d} \psi'$ iff $\mathfrak{M}, w_d \models \psi' \supset \neg \overline{\mathbb{P}}_a \mathbb{Q}_{w_d} \psi$.

⁴ For instance, *fairness through unawareness* requires that protected attributes (e.g., race, religion, or gender) are not explicitly used in the prediction process. However, StatEL may not be suited to formalizing such a property in treatment.

See Appendix A for the proof.

Now we formalize three popular notions of fairness of classifiers by using counterfactual epistemic operators (introduced in Section 3) as follows.

7.1 Group Fairness (Statistical Parity)

The *group fairness* formulated as *statistical parity* [13] is the property that the output distributions of the classifier are identical for different groups. Formally, for each $b = 0, 1$ and a group $G_b \subseteq \mathcal{D}$, let μ_{G_b} be the distribution of the output (over \mathbb{L}) of the classifier C when the input is sampled from a dataset d and belongs to G_b . Then the statistical parity up to bias ε is formalized using the total variation D_{tv} by $D_{\text{tv}}(\mu_{G_0} \parallel \mu_{G_1}) \leq \varepsilon$.

To express this using StatEL, we define an accessibility relation $\mathcal{R}_\varepsilon^{\text{tv}}$ in \mathfrak{M} by:

$$\mathcal{R}_\varepsilon^{\text{tv}} \stackrel{\text{def}}{=} \{(w, w') \in \mathcal{W} \times \mathcal{W} \mid D_{\text{tv}}(\sigma_w(\hat{y}) \parallel \sigma_{w'}(\hat{y})) \leq \varepsilon\}. \quad (7)$$

Intuitively, $(w, w') \in \mathcal{R}_\varepsilon^{\text{tv}}$ represents that the two probability distributions $\sigma_w(\hat{y})$ and $\sigma_{w'}(\hat{y})$ of the outputs by the classifier C respectively in w and in w' are close in terms of D_{tv} . Note that $\sigma_w(\hat{y})$ and $\sigma_{w'}(\hat{y})$ respectively represent μ_{G_0} and μ_{G_1} .

Then the statistical parity w.r.t. groups G_0, G_1 means that in terms of $\mathcal{R}_\varepsilon^{\text{tv}}$, we cannot distinguish a world having a dataset d and satisfying $\eta_{G_0}(x) \wedge \psi(x, \hat{y})$ from another satisfying $\eta_{G_1}(x) \wedge \psi(x, \hat{y})$. By Proposition 2, this is expressed as:

$$\mathfrak{M}, w_d \models \text{GrpFair}(x, \hat{y})$$

where $\text{GrpFair}(x, \hat{y}) \stackrel{\text{def}}{=} (\eta_{G_0}(x) \wedge \psi(x, \hat{y})) \supset \neg \overline{\text{P}}_\varepsilon^{\text{tv}}(\xi_d \wedge \mathbb{Q}_{w_d}(\eta_{G_1}(x) \wedge \psi(x, \hat{y})))$.

7.2 Individual Fairness (as Lipschitz Property)

The *individual fairness* formulated as a Lipschitz property [13] is the property that the classifier outputs similar labels given similar inputs. Formally, for $v, v' \in \mathcal{D}$, let μ_v and $\mu_{v'}$ be the distributions of the outputs (over \mathbb{L}) of the classifier C when the inputs are v and v' , respectively. Then the individual fairness is formalized using a divergence $D : \mathbb{DL} \times \mathbb{DL} \rightarrow \mathbb{R}^{\geq 0}$, a metric $r : \mathcal{D} \times \mathcal{D} \rightarrow \mathbb{R}^{\geq 0}$, and a threshold $\varepsilon \in \mathbb{R}^{\geq 0}$ by $D(\mu_v \parallel \mu_{v'}) \leq \varepsilon \cdot r(v, v')$.

To express this using StatEL, we define an accessibility relation $\mathcal{R}_\varepsilon^{r,D}$ in \mathfrak{M} for the metric r and the divergence D as follows:

$$\mathcal{R}_\varepsilon^{r,D} \stackrel{\text{def}}{=} \left\{ (w, w') \in \mathcal{W} \times \mathcal{W} \mid \begin{array}{l} v \in \text{supp}(\sigma_w(x)), v' \in \text{supp}(\sigma_{w'}(x)), \\ D(\sigma_w(\hat{y}) \parallel \sigma_{w'}(\hat{y})) \leq \varepsilon \cdot r(v, v') \end{array} \right\}. \quad (8)$$

Intuitively, $(w, w') \in \mathcal{R}_\varepsilon^{r,D}$ represents that, when inputs are closer in terms of the metric r , the classifier C outputs closer labels in terms of the divergence D .

Then the individual fairness w.r.t. r and D means that in terms of $\mathcal{R}_\varepsilon^{r,D}$, we cannot distinguish between the two worlds w and w' where $\psi(x, \hat{y})$ is satisfied (i.e., C outputs \hat{y} given an input x). By Proposition 2, this is expressed as:

$$\mathfrak{M}, w_d \models \text{IndFair}(x, \hat{y})$$

where $\text{IndFair}(x, \hat{y}) \stackrel{\text{def}}{=} \psi(x, \hat{y}) \supset \overline{\mathbf{P}_\varepsilon^{r,D}}(\xi_d \wedge \mathbb{Q}_{w_d} \psi(x, \hat{y}))$.

This represents that by observing the classifier's output \hat{y} , we can less distinguish two worlds w and w' when their inputs $\sigma_w(x)$ and $\sigma_{w'}(x)$ are closer.

7.3 Equal Opportunity

Equal opportunity [21, 40] is the property that the recall (true positive rate) is the same for all the groups. Formally, given an advantage class $\ell \in \mathbf{L}$ (e.g., not defaulting on a loan) and a group $G \subseteq \mathcal{D}$ of inputs with a protected attribute (e.g., race), a classifier C is said to satisfy equal opportunity of ℓ w.r.t. G if it holds for each $\hat{\ell} \in \mathbf{L}$ that:

$$\Pr[C(x) = \hat{\ell} \mid x \in G, H(x) = \ell] = \Pr[C(x) = \hat{\ell} \mid x \in \mathcal{D} \setminus G, H(x) = \ell]. \quad (9)$$

If we allow the logic to use the universal quantification over the probability value i , then the case of $\hat{\ell} = \ell$ in (9) could be expressed as:

$$\forall i \in [0, 1]. (\xi_d \wedge \eta_G(x) \supset \text{Recall}_{\ell,i}(x)) \leftrightarrow (\xi_d \wedge \neg\eta_G(x) \supset \text{Recall}_{\ell,i}(x)).$$

However, instead of allowing for this universal quantification, we can use the modal operators $\overline{\mathbf{P}_\varepsilon^{\text{tv}}}$ (defined by (7)) with $\varepsilon = 0$, and represent equal opportunity as the fact that we cannot distinguish a world having a dataset d and satisfying $\eta_G(x) \wedge \psi(x, \hat{y}) \wedge h_\ell(x)$ from another satisfying $\neg\eta_G(x) \wedge \psi(x, \hat{y}) \wedge h_\ell(x)$ as follows:

$$\text{EqOpp}(x, \hat{y}) \stackrel{\text{def}}{=} (\eta_G(x) \wedge \psi(x, \hat{y}) \wedge h_\ell(x)) \supset \overline{\mathbf{P}_0^{\text{tv}}}(\xi_d \wedge \mathbb{Q}_{w_d}(\neg\eta_G(x) \wedge \psi(x, \hat{y}) \wedge h_\ell(x))).$$

8 Related Work

In this section, we provide a brief overview of related work on the specification of statistical machine learning and on epistemic logic for describing specification.

Desirable properties of statistical machine learning. There have been a large number of papers on attacks and defences for deep neural networks [36, 8]. Compared to them, however, not much work has been done to explore the formal specification of various properties of machine learning. Seshia et al. [34] present a list of desirable properties of DNNs (deep neural networks) although most of the properties are presented informally without mathematical formulas. As for robustness, Dreossi et al. [11] propose a unifying formalization of adversarial input generation in a rigorous and organized manner, although they formalize and classify attacks (as optimization problems) rather than define the robustness notions themselves. Concerning the fairness notions, Gajane [16] surveys the formalization of fairness notions for machine learning and present some justification based on social science literature.

Epistemic logic for describing specification. Epistemic logic [39] has been studied to represent and reason about knowledge [14, 19, 20], and has been applied to describe various properties of systems.

The *BAN logic* [5], proposed by Burrows, Abadi and Needham, is a notable example of epistemic logic used to model and verify the authentication in cryptographic protocols. To improve the formalization of protocols’ behaviours, some epistemic approaches integrate process calculi [23, 10, 7].

Epistemic logic has also been used to formalize and reason about privacy properties, including anonymity [35, 20, 17, 27], receipt-freeness of electronic voting protocols [24], and privacy policy for social network services [32]. Temporal epistemic logic is used to express information flow security policies [3].

Concerning the formalization of fairness notions, previous work in formal methods has modeled different kinds of fairness involving timing by using temporal logic rather than epistemic logic. As far as we know, no previous work has formalized fairness notions of machine learning using counterfactual epistemic operators.

Formalization of statistical properties. In studies of philosophical logic, Lewis [29] shows the idea that when a random value has various possible probability distributions, then those distributions should be represented on distinct possible worlds. Bana [4] puts Lewis’s idea in a mathematically rigorous setting. Recently, a modal logic called statistical epistemic logic [26] is proposed and is used to formalize statistical hypothesis testing and the notion of differential privacy [12]. Independently of that work, French et al. [15] propose a probability model for a dynamic epistemic logic in which each world is associated with a subjective probability distribution over the universe, without dealing with non-deterministic inputs or statistical divergence.

9 Conclusion

We have shown a logical approach to formalizing statistical classifiers and their desirable properties in a simple and abstract way. Specifically, we have introduced a formal model for probabilistic behaviours of classifiers and non-deterministic adversarial inputs using a distributional Kripke model. Then we have formalized the classification performance, robustness, and fairness of classifiers by using StatEL. Moreover, we have also clarified some relationships among properties of classifiers, and relevance between classification performance and robustness. To formalize fairness notions, we have introduced a notion of counterfactual knowledge and shown some techniques to express conditional indistinguishability. As far as we know, this is the first work that uses logical formulas to express statistical properties of machine learning, and that provides epistemic (resp. counterfactually epistemic) views on robustness (resp. fairness) of classifiers.

In future work, we are planning to include temporal operators in the specification language and to formally reason about system-level properties of learning-based systems. We are also interested in developing a general framework for

the formal specification of machine learning associated with testing methods and possibly extended with Bayesian networks. Our future work also includes an extension of StatEL to formalize machine learning other than classification problems. Another possible direction of future work would be to clarify the relationships between our counterfactual epistemic operators and more general notions of counterfactual knowledge in previous work such as [38].

A Proofs for Propositions 1 and 2

Proposition 1 (Conditional indistinguishability) *Let $\mathfrak{M} = (\mathcal{W}, (\mathcal{R}_a)_{a \in \mathcal{A}}, (V_s)_{s \in \mathcal{S}})$ be a distributional Kripke model with the universe $\mathcal{W} = \mathbb{D}\mathcal{S}$. Let ψ and ψ' be static formulas, and $a \in \mathcal{A}$.*

- (i) $\mathfrak{M} \models \psi \supset \neg \overline{\mathcal{P}}_a \mathbb{P}_1 \psi'$ iff for any $w, w' \in \mathcal{W}$, $\mathfrak{M}, w \models \mathbb{P}_1 \psi$ and $\mathfrak{M}, w' \models \mathbb{P}_1 \psi'$ imply $(w, w') \in \mathcal{R}_a$.
- (ii) If \mathcal{R}_a is symmetric, then $\mathfrak{M} \models \psi \supset \neg \overline{\mathcal{P}}_a \mathbb{P}_1 \psi'$ iff $\mathfrak{M} \models \psi' \supset \neg \overline{\mathcal{P}}_a \mathbb{P}_1 \psi$.

Proof. We first prove the claim (i) as follows. We show the direction from left to right. Assume that $\mathfrak{M} \models \psi \supset \neg \overline{\mathcal{P}}_a \mathbb{P}_1 \psi'$. Let $w, w' \in \mathcal{W}$ satisfy $\mathfrak{M}, w \models \mathbb{P}_1 \psi$ and $\mathfrak{M}, w' \models \mathbb{P}_1 \psi'$. Then $w|_\psi = w$. By $\mathfrak{M}, w \models \psi \supset \neg \overline{\mathcal{P}}_a \mathbb{P}_1 \psi'$, we obtain $\mathfrak{M}, w|_\psi \models \neg \overline{\mathcal{P}}_a \mathbb{P}_1 \psi'$, which is logically equivalent to $\mathfrak{M}, w|_\psi \models \overline{\mathcal{K}}_a \neg \mathbb{P}_1 \psi'$. By the definition of $\overline{\mathcal{K}}_a$, for every $w'' \in \mathcal{W}$, $\mathfrak{M}, w'' \models \mathbb{P}_1 \psi'$ implies $(w|_\psi, w'') \in \mathcal{R}_a$. Then, since $w|_\psi = w$ and $\mathfrak{M}, w' \models \mathbb{P}_1 \psi'$, we obtain $(w, w') \in \mathcal{R}_a$.

Next we show the other direction as follows. Assume the right hand side. Let $w \in \mathcal{W}$ such that $\mathfrak{M}, w \models \mathbb{P}_1 \psi$. Then for every $w' \in \mathcal{W}$, $\mathfrak{M}, w' \models \mathbb{P}_1 \psi'$ implies $(w, w') \in \mathcal{R}_a$. By the definition of $\overline{\mathcal{K}}_a$, we have $\mathfrak{M}, w \models \overline{\mathcal{K}}_a \neg \mathbb{P}_1 \psi'$, which is equivalent to $\mathfrak{M}, w \models \neg \overline{\mathcal{P}}_a \mathbb{P}_1 \psi'$. By $\mathfrak{M}, w \models \mathbb{P}_1 \psi$, we have $w|_\psi = w$, hence $\mathfrak{M}, w|_\psi \models \neg \overline{\mathcal{P}}_a \mathbb{P}_1 \psi'$. Therefore $\mathfrak{M}, w \models \psi \supset \neg \overline{\mathcal{P}}_a \mathbb{P}_1 \psi'$.

Finally, the claim (ii) follows from the claim (i) immediately. \square

Proposition 2 (Conditional indistinguishability in a world w_d) *Let $\mathfrak{M} = (\mathcal{W}, (\mathcal{R}_a)_{a \in \mathcal{A}}, (V_s)_{s \in \mathcal{S}})$ be a distributional Kripke model with the universe $\mathcal{W} = \mathbb{D}\mathcal{S}$. Let w_d be a world with a dataset d , ψ and ψ' be static formulas, and $a \in \mathcal{A}$.*

- (i) $\mathfrak{M}, w_d \models \psi \supset \neg \overline{\mathcal{P}}_a (\xi_d \wedge \mathbb{Q}_{w_d} \psi')$ iff for any $w, w' \in \mathcal{W}$, $\mathfrak{M}, w \models \xi_d \wedge \mathbb{Q}_{w_d} \psi$ and $\mathfrak{M}, w' \models \xi_d \wedge \mathbb{Q}_{w_d} \psi'$ imply $(w, w') \in \mathcal{R}_a$.
- (ii) If \mathcal{R}_a is symmetric, then $\mathfrak{M}, w_d \models \psi \supset \neg \overline{\mathcal{P}}_a \mathbb{Q}_{w_d} \psi'$ iff $\mathfrak{M}, w_d \models \psi' \supset \neg \overline{\mathcal{P}}_a \mathbb{Q}_{w_d} \psi$.

Proof. We first prove the claim (i) as follows. We show the direction from left to right. Assume that $\mathfrak{M}, w_d \models \psi \supset \neg \overline{\mathcal{P}}_a (\xi_d \wedge \mathbb{Q}_{w_d} \psi')$. Let $w, w' \in \mathcal{W}$ satisfy $\mathfrak{M}, w \models \xi_d \wedge \mathbb{Q}_{w_d} \psi$ and $\mathfrak{M}, w' \models \xi_d \wedge \mathbb{Q}_{w_d} \psi'$. Then $w_d|_\psi = w$ and $w_d|_{\psi'} = w'$. By $\mathfrak{M}, w_d \models \psi \supset \neg \overline{\mathcal{P}}_a (\xi_d \wedge \mathbb{Q}_{w_d} \psi')$, we obtain $\mathfrak{M}, w_d|_\psi \models \neg \overline{\mathcal{P}}_a (\xi_d \wedge \mathbb{Q}_{w_d} \psi')$, which is logically equivalent to $\mathfrak{M}, w_d|_\psi \models \overline{\mathcal{K}}_a \neg (\xi_d \wedge \mathbb{Q}_{w_d} \psi')$. By the definition of $\overline{\mathcal{K}}_a$ and $\mathfrak{M}, w' \models \xi_d \wedge \mathbb{Q}_{w_d} \psi'$, we have $(w_d|_\psi, w') \in \mathcal{R}_a$. Therefore, by $w|_\psi = w$, we obtain $(w, w') \in \mathcal{R}_a$.

Next we show the other direction as follows. Assume the right hand side. Let $w \in \mathcal{W}$ such that $\mathfrak{M}, w \models \xi_d \wedge Q_{w_d} \psi$. Then for every $w' \in \mathcal{W}$, $\mathfrak{M}, w' \models \xi_d \wedge Q_{w_d} \psi'$ implies $(w, w') \in \mathcal{R}_a$. By the definition of \overline{K}_a , we have $\mathfrak{M}, w \models \overline{K}_a \neg(\xi_d \wedge Q_{w_d} \psi')$, which is equivalent to $\mathfrak{M}, w \models \neg \overline{P}_a(\xi_d \wedge Q_{w_d} \psi')$. By $\mathfrak{M}, w \models \xi_d \wedge Q_{w_d} \psi$, we have $w_d|_\psi = w$, hence $\mathfrak{M}, w_d|_\psi \models \neg \overline{P}_a(\xi_d \wedge Q_{w_d} \psi')$. Therefore $\mathfrak{M}, w_d \models \psi \supset \neg \overline{P}_a(\xi_d \wedge Q_{w_d} \psi')$.

Finally, the claim (ii) follows from the claim (i) immediately. \square

References

1. Alpern, B., Schneider, F.B.: Defining liveness. *Inf. Process. Lett.* **21**(4), 181–185 (1985). [https://doi.org/10.1016/0020-0190\(85\)90056-0](https://doi.org/10.1016/0020-0190(85)90056-0)
2. Athalye, A., Engstrom, L., Ilyas, A., Kwok, K.: Synthesizing robust adversarial examples. In: Proc. ICML. pp. 284–293 (2018)
3. Balliu, M., Dam, M., Guernic, G.L.: Epistemic temporal logic for information flow security. In: Proc. of PLAS. p. 6 (2011). <https://doi.org/10.1145/2166956.2166962>
4. Bana, G.: Models of objective chance: An analysis through examples. In: Making it Formally Explicit. pp. 43–60. Springer International Publishing (2017). https://doi.org/10.1007/978-3-319-55486-0_3
5. Burrows, M., Abadi, M., Needham, R.M.: A logic of authentication. *ACM Trans. Comput. Syst.* **8**(1), 18–36 (1990). <https://doi.org/10.1145/77648.77649>
6. Carlini, N., Wagner, D.A.: Towards evaluating the robustness of neural networks. In: Proc. S&P. pp. 39–57 (2017). <https://doi.org/10.1109/SP.2017.49>
7. Chadha, R., Delaune, S., Kremer, S.: Epistemic logic for the applied pi calculus. In: Proc. of FMOODS/FORTE. pp. 182–197 (2009). https://doi.org/10.1007/978-3-642-02138-1_12
8. Chakraborty, A., Alam, M., Dey, V., Chattopadhyay, A., Mukhopadhyay, D.: Adversarial attacks and defences: A survey. *CoRR* **abs/1810.00069** (2018), <http://arxiv.org/abs/1810.00069>
9. Clarkson, M.R., Schneider, F.B.: Hyperproperties. In: Proc. of CSF. pp. 51–65. IEEE (2008). <https://doi.org/10.1109/CSF.2008.7>
10. Dechesne, F., Mousavi, M., Orzan, S.: Operational and epistemic approaches to protocol analysis: Bridging the gap. In: Proc. of LPAR. pp. 226–241 (2007)
11. Dreossi, T., Ghosh, S., Sangiovanni-Vincentelli, A.L., Seshia, S.A.: A formalization of robustness for deep neural networks. In: Proc. VNN (2019)
12. Dwork, C.: Differential privacy. In: Proc. of ICALP. pp. 1–12 (2006)
13. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.S.: Fairness through awareness. In: Proc. of ITCS. pp. 214–226. ACM (2012)
14. Fagin, R., Halpern, J., Moses, Y., Vardi, M.: Reasoning about Knowledge. The MIT Press (1995)
15. French, T., Gozzard, A., Reynolds, M.: Dynamic aleatoric reasoning in games of bluffing and chance. In: Proc. AAMAS. pp. 1964–1966 (2019)
16. Gajane, P.: On formalizing fairness in prediction with machine learning. *CoRR* **abs/1710.03184** (2017), <http://arxiv.org/abs/1710.03184>
17. Garcia, F.D., Hasuo, I., Pieters, W., van Rossum, P.: Provable anonymity. In: Proc. of FMSE. pp. 63–72 (2005). <https://doi.org/10.1145/1103576.1103585>
18. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. In: Proc. of ICLR (2015)

19. Halpern, J.Y.: Reasoning about uncertainty. The MIT press (2003)
20. Halpern, J.Y., O’Neill, K.R.: Anonymity and information hiding in multiagent systems. *J. of Comp. Security* **13**(3), 483–512 (2005)
21. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. In: *proc. NIPS*. pp. 3315–3323 (2016)
22. Huang, X., Kwiatkowska, M., Wang, S., Wu, M.: Safety verification of deep neural networks. In: *Proc. CAV*. pp. 3–29 (2017). https://doi.org/10.1007/978-3-319-63387-9_1
23. Hughes, D., Shmatikov, V.: Information hiding, anonymity and privacy: a modular approach. *J. of Comp. Security* **12**(1), 3–36 (2004)
24. Jonker, H.L., Pieters, W.: Receipt-freeness as a special case of anonymity in epistemic logic. In: *Proc. Workshop On Trustworthy Elections (WOTE’06)* (June 2006)
25. Katz, G., Barrett, C.W., Dill, D.L., Julian, K., Kochenderfer, M.J.: Reluplex: An efficient SMT solver for verifying deep neural networks. In: *Proc. CAV*. pp. 97–117 (2017). https://doi.org/10.1007/978-3-319-63387-9_5
26. Kawamoto, Y.: Statistical epistemic logic. *CoRR* **abs/1412.4451** (2019), <https://arxiv.org/pdf/1907.05995.pdf>
27. Kawamoto, Y., Mano, K., Sakurada, H., Hagiya, M.: Partial knowledge of functions and verification of anonymity. *Transactions of the Japan Society for Industrial and Applied Mathematics* **17**(4), 559–576 (2007). <https://doi.org/10.11540/jsiamt.17.4.559>
28. Kripke, S.A.: Semantical analysis of modal logic i normal modal propositional calculi. *Mathematical Logic Quarterly* **9**(5-6), 67–96 (1963)
29. Lewis, D.: A subjectivist’s guide to objective chance. In: *Studies in Inductive Logic and Probability, Volume II*, pp. 263–293. Berkeley: University of California Press (1980)
30. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: *Proc. ICLR* (2018)
31. Moosavi-Dezfooli, S., Fawzi, A., Frossard, P.: Deepfool: A simple and accurate method to fool deep neural networks. In: *Proc. CVPR*. pp. 2574–2582 (2016). <https://doi.org/10.1109/CVPR.2016.282>
32. Pardo, R., Schneider, G.: A formal privacy policy framework for social networks. In: *Proc. SEFM*. pp. 378–392 (2014). https://doi.org/10.1007/978-3-319-10431-7_30
33. Pei, K., Cao, Y., Yang, J., Jana, S.: Deepxplore: Automated white-box testing of deep learning systems. In: *Proc. SOSP*. pp. 1–18 (2017). <https://doi.org/10.1145/3132747.3132785>
34. Seshia, S.A., Desai, A., Dreossi, T., Fremont, D.J., Ghosh, S., Kim, E., Shivakumar, S., Vazquez-Chanlatte, M., Yue, X.: Formal specification for deep neural networks. In: *Proc. ATVA*. pp. 20–34 (2018). https://doi.org/10.1007/978-3-030-01090-4_2
35. Syverson, P.F., Stubblebine, S.G.: Group principals and the formalization of anonymity. In: *World Congress on Formal Methods (1)*. pp. 814–833 (1999). https://doi.org/10.1007/3-540-48119-2_45
36. Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I.J., Fergus, R.: Intriguing properties of neural networks. In: *Proc. ICLR* (2014)
37. Tian, Y., Pei, K., Jana, S., Ray, B.: Deeptest: automated testing of deep-neural-network-driven autonomous cars. In: *Proc. ICSE*. pp. 303–314 (2018). <https://doi.org/10.1145/3180155.3180220>
38. Williamson, T.: Philosophical knowledge and knowledge of counterfactuals. *Grazer Philosophische Studien* **74**, 89 (2007)
39. von Wright, G.H.: *An Essay in Modal Logic*. Amsterdam: North-Holland Pub. Co. (1951)

40. Zafar, M.B., Valera, I., Gomez-Rodriguez, M., Gummadi, K.P.: Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In: Proc. WWW. pp. 1171–1180 (2017). <https://doi.org/10.1145/3038912.3052660>