

Received July 20, 2020, accepted August 26, 2020, date of publication August 31, 2020, date of current version September 14, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.3020638

Copying Machine Learning Classifiers

IRENE UNCETA^{1,2}, JORDI NIN³, AND ORIOL PUJOL²

¹BBVA Data & Analytics, 08018 Barcelona, Spain

²Department of Mathematics and Computer Science, Universitat de Barcelona, 08007 Barcelona, Spain

³Department of Operations, Innovation and Data Sciences, Universitat Ramon Llull, ESADE, 08172 Sant Cugat del Vallès, Spain

Corresponding author: Irene Unceta (irene.unceta@bbvadata.com)

This work was supported in part by the Spanish Project (MICINN) under Grant PID2019-105093GB-I00, and in part by the AGAUR of the Generalitat de Catalunya through the Industrial Ph.D., under Grant 2017-DI-25.

ABSTRACT We study copying of machine learning classifiers, an agnostic technique to replicate the decision behavior of any classifier. We develop the theory behind the problem of copying, highlighting its properties, and propose a framework to copy the decision behavior of any classifier using no prior knowledge of its parameters or training data distribution. We validate this framework through extensive experiments using data from a series of well-known problems. To further validate this concept, we use three different use cases where desiderata such as interpretability, fairness or productivization constrains need to be addressed. Results show that copies can be exploited to enhance existing solutions and improve them adding new features and characteristics.

INDEX TERMS Applied machine learning, classification, copying, differential replication, fidelity.

I. INTRODUCTION

In many every-day examples, performance of state-of-the-art machine learning is held back by operational constraints that appear along a system's life-cycle. Either the data or the models themselves are subject to privacy restrictions [1]–[3] or new specific regulations apply that require models to be self-explanatory [4]–[6] or fair with respect to sensitive data attributes [7]–[9]. Other issues include time or space limitations for deployment, and production bottlenecks in delivering certain models to the market [10]. To the best of our knowledge, these issues have been traditionally addressed by means of re-training tailored solutions. As a result, off-the-shelf machine learning techniques often yield only sub-optimal results or can only be exploited during a limited period of time.

Under such circumstances, training a new model may seem straightforward. However, a re-training is not always possible, nor advisable. This may be, for example, because production protocols require the maintenance of predictive performance over time, because the specifics of the model are unknown or even because the training data are no longer available. What is more, re-training is timely and often costly too, as it may require a non-negligible amount of human and material resources. Whatever the cause, the impossibility of re-training calls for new ways to address this situation.

The associate editor coordinating the review of this manuscript and approving it for publication was Yongming Li¹.

In this article we study copying, the problem of building a new model that replicates the decision behavior of another. The idea of approximating a model's decision boundary can be found in the literature under different topics, including model extraction [11], [12], knowledge distillation [13], [14], or adversarial learning [15], [16]. All these notions refer to scenarios where the knowledge acquired by one model is used to build another. Specifically, we here envisage the most agnostic scenario, where we make the minimum number of required assumptions about the amount of information available during the process. We assume access to the model is limited to a membership query interface. In addition and unlike previous articles, where the training data distribution is directly [13] or indirectly [14] known, we also assume the training data to be unknown or, simply, lost. Finally, we also assume the query interface to produce only hard predictions, as opposed to scenarios where rich information outputs can be used as soft targets for the new model [17], [18]. This scenario can be understood as a form of zero-knowledge distillation, where the decision behavior of a larger model is transferred to a simpler one in circumstances where no knowledge is assumed about the training data or the model internals. Effectively, this corresponds to an scenario where the larger model is a black-box and distillation is conducted in a data-free way.

In this context, we propose copying as a methodology to project the decision function learned by a model onto a new hypothesis space that enables the same decision behaviour,

while incorporating new features and properties. This process is one of *differential replication* [19]. Copies not only retain the original accuracy, but can also be used to endow classifiers with new characteristics, including interpretability, online learning or equity features. Hence, copying can be exploited to overcome the aforementioned limitations by building next generation models that are fitter under the new conditions.

We summarize the main contributions of this paper as:

- We formalize the problem of building a copy that replicates the decision behavior of a machine learning model in the most agnostic setting.
- We explore the theoretical framework and implications of this problem, and study how these can be exploited when building a copy.
- We put this theory into practice to highlight the specific characteristics of copying and validate this proposal on a series of well known problems.
- We further illustrate the value of copying for differential replication in three real use cases. First, we address the issues of non-decomposability and delayed time-to-market delivery in non-client mortgage risk scoring. Second, we build an online copy that recovers a critical operating point in a loan default prediction problem. Finally, we use copies to ensure a fair classification of superhero alignment.

The rest of this article is organized as follows. Sec. II presents a literature survey of related work. The theoretical basis for copying is introduced in Sec. III, while Sec. IV extracts meaningful insights for a practical implementation. In Sec. V we validate copies on various UCI problems. In Sec. VI we consider the advantages and limitations of this methodology and present three real applications. The paper concludes with a brief summary of our findings and an outline of future research.

II. RELATED WORK

The idea of leveraging the knowledge of one classifier to train another has been explored under different forms and scenarios in the literature. We find this notion in early works on concept extraction, where trained artificial neural networks are compiled into a set of representative rules [20]–[23]. Of particular interest to this article is TREPAN [11], a query algorithm that extracts tree-structured representations of trained neural networks. Following these early ideas, in 2006 Bucilua *et al.* proposed a method to compress the knowledge acquired by large classifier ensembles into more compact models that were better fit to meet the requirements of production deployment at test time [14]. Ever since, this notion has been popularized under the name of knowledge distillation to study how the knowledge acquired by a complex model, the teacher, can be exploited to guide training of a simpler model, the student [13]. Papers in this field have explored different forms of supervision from the teacher [24], training the same network in generations [25] or inducing teacher signals with a softened label distribution to convey useful task-dependent information to students [18].

An important degree of freedom in distillation is the transfer set used to train the simpler model. Traditionally, knowledge transfer has been treated as a standard learning process, where the training data are relabelled and extended to learn an alternative model [26]. Most papers use the same set to train teacher and student, either in its raw form [13], [26], [27] or enriched with additional synthetic data [11], [28], [17]. Besides, researchers have also studied cases where teachers and students faced with the same task have different access to the training data [29]. Or even situations where the training data are not accessible and distillation is conducted under the use of unlabelled data [14], [30]. Generating such unlabelled data, however, is generally expensive as well as complex. Over the years, different disciplines have evolved in relation to this issue. See, for example, works on machine teaching, where a human teacher hand-picks as small a training set as possible to train a machine learning system [31]. Or, alternatively, the numerous contributions to the field of active learning, where a desired hypothesis is learned by reducing the number of queries to a human oracle [32]. In this context, different query optimization techniques have been proposed to obtain a reduced set of highly informative samples [33]–[35], [36].

Distillation has been found to work well across a wide range of applications, including mutual learning [37], distributed learning [38], learning from noisy labels [39] or training stabilization [40]. In a few cases, it has also been extended to other tasks, such as data augmentation [41] or data privacy [42], [43]. In particular, distillation has been exploited for *transferability*-based adversarial learning [15], [44], [45]–[47], where a malicious adversary exploits samples crafted from a local substitute of a model to compromise it.

Despite this success, however, there is still a very limited understanding of the theoretical and empirical foundations behind knowledge distillation. Lopez-Paz *et al.* [48] related distillation to a form of learning using privileged information, while Phuong and Lambert [49] proposed described a series of factors that determine the success of distillation. To the best of our knowledge, these are the two only contributions aimed at revealing the mechanisms underlying knowledge distillation.

III. COPYING

Copying refers to the process of building a functional model which is equivalent in its decision behaviour to another. During this process, the knowledge acquired by the first model is transferred to a copy, in circumstances where both the internals and the training data of the former are unknown, and access to its knowledge is only possible through a membership query interface.

Let us take a classifier $f_{\mathcal{O}} : \mathcal{X} \rightarrow \mathcal{T}$, where \mathcal{X} and \mathcal{T} correspond to the input and label spaces, respectively. We define the set $\mathcal{D} = \{(x_i, t_i)\}_{i=1}^M$ as the training data, for M the total number of instances, and restrict to the case of classification, where $\mathcal{T} \in \mathbb{Z}_k$ for k the number of classes.

Copying is defined as the problem of finding a model $f_C(\theta) \in \mathcal{H}_C$, parameterized by θ , such that given a new sample \mathbf{x}^* it predicts the output $y^* = f_C(\mathbf{x}^*)$. Our objective is therefore to obtain a new model, the copy, whose decision function mimics that of f_O all over the space.

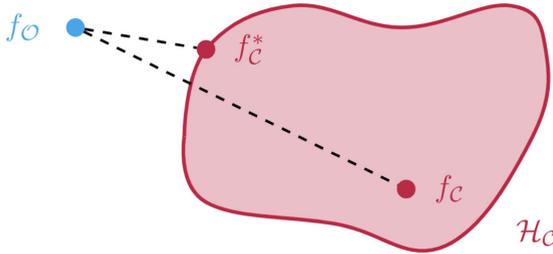


FIGURE 1. Copying as a projection of a decision function onto a new hypothesis space \mathcal{H}_C . This space need not coincide with that of the classifier, i.e. f_O and f_C need not belong to the same family of models, and they most usually don't. The optimal copy f_C^* is the closest to f_O .

The process of copying can be interpreted as projecting the decision function f_O onto the new hypothesis space \mathcal{H}_C the copy belongs to. A graphical illustration of this is shown in Fig. 1. As we will later explain in more detail, this new hypothesis space need not coincide with that of f_O . On the contrary, we can exploit to our advantage the fact that both spaces are different to endow the model with new features, not present in the original hypothesis space. This *differential replication* process is the crucial characteristic of copying.

The problem of copying is characterized by the predictive distribution $P(y^*|f_O, \mathbf{x}^*)$. Marginalizing with respect to the copy parameters θ

$$P(y^*|f_O, \mathbf{x}^*) = \int_{\theta \in \mathcal{H}_C} P(y^*|\theta, f_O, \mathbf{x}^*)P(\theta|f_O, \mathbf{x}^*)d\theta,$$

for \mathcal{H}_C the complete parameter space for the copy. We simplify this expression by making two basic assumptions.

First, when building the copy, knowledge about the unseen data point \mathbf{x}^* is not available, so that $P(\theta|f_O, \mathbf{x}^*) = P(\theta|f_O)$. Second, once having built the copy, i.e. fixed the value of θ , interaction with the classifier f_O is no longer required, so that $P(y^*|\theta, f_O, \mathbf{x}^*) = P(y^*|\theta, \mathbf{x}^*)$. On this basis, we rewrite the expression above as

$$P(y^*|f_O, \mathbf{x}^*) = \int_{\theta \in \mathcal{H}_C} P(y^*|\theta, \mathbf{x}^*)P(\theta|f_O)d\theta.$$

We take a *winner takes it all* approach and force the posterior to have the form of a point mass density, $P(\theta|f_O) = \delta(\theta - \theta^*)$, for $\delta(\cdot)$ the Dirac delta function and θ^* the optimal parameter set. All the probability mass is then placed onto θ^* , so that

$$P(y^*|f_O, \mathbf{x}^*) = P(y^*|\theta^*, \mathbf{x}^*).$$

Hence, the problem of copying can be understood as that of finding the optimal parameter values θ^* to maximize the posterior probability

$$\theta^* = \arg \max_{\theta} P(\theta|f_O). \tag{1}$$

A. THE NEED FOR UNLABELLED DATA

We study the most general scenario, where the training data \mathcal{D} is assumed to be lost. Solving (1) therefore requires that we generate new data in order to gain information about the form of f_O throughout the input space \mathcal{X} . We introduce unlabelled data points $\mathbf{z} \in \mathcal{X}$ and rewrite (1) as

$$\theta^* = \arg \max_{\theta} \int_{\mathbf{z} \sim P_Z} P(\theta|f_O(\mathbf{z}))dP_Z, \tag{2}$$

for an arbitrary generating probability distribution P_Z from which the new samples are independently drawn. This distribution defines the spatial support for the copy, i.e. its plausible operational space. In the existing literature, the training data distribution, P , is directly [13] or indirectly [14] accessible. Here we completely lack this information, so that we cannot match P_Z to our estimate of P . Nonetheless, note that despite P_Z could be related to the training distribution, this is not mandatory for our purposes.

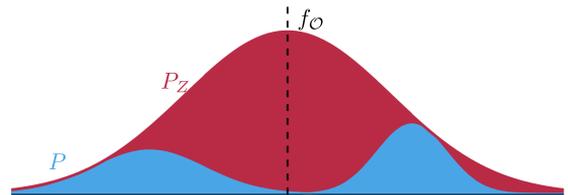


FIGURE 2. Gaussian training data distribution P and learned decision boundary f_O . Alternative gaussian distribution for P_Z .

Take for example the completely separable binary problem in Fig. 2, where each class comes from a Gaussian distribution and the decision boundary lies in a low density area of the space. Further assume that we are in a production setting, so that we have full knowledge of the system. In principle, in this scenario it would be possible, and even desirable, to match P_Z with P . Indeed, by forcing $P_Z = P$ we ensure that the copy replicates the learned decision behaviour in those areas where the training data lie. However, the copy may display a completely different behaviour around the boundary, where these data are scarce. An interesting modelling question in this scenario would be: *what should the copy do in corner cases?* Another extreme case is that of counterfactuals, which include operation regimes even in front of impossible events and data values.

More generally, defining P_Z to resemble the form of P might help in ensuring that the copy generalizes well in the training domain. However, this can also be achieved by other methods, such as updating the form of P_Z as we gain more information about f_O , or choosing a P_Z that adapts to the form of the copy hypothesis space. Indeed, choosing P_Z adequately can be difficult, given that we have no intuition about where the training data are located or which specific regions the copy should focus on. In Sec. IV we study this problem in more depth.

B. INTRODUCING THE DUAL OPTIMIZATION

Let us then assume an arbitrary form for the probability distribution P_Z . Because maximizing the posterior is equal

to maximizing the log-posterior, we rewrite (2) as

$$\begin{aligned} \theta^* &= \arg \max_{\theta} \left[\log \left(\int_{\mathbf{z} \sim P_Z} P(\theta | f_{\mathcal{O}}(\mathbf{z})) dP_Z \right) \right] \\ &= \arg \max_{\theta} \left[\log \left(\int_{\mathbf{z} \sim P_Z} \frac{P(f_{\mathcal{O}}(\mathbf{z}) | \theta) P(\theta)}{P(f_{\mathcal{O}}(\mathbf{z}))} dP_Z \right) \right] \end{aligned}$$

where we apply *Bayes' rule* to the terms inside the integral. Using *Jensen's inequality*¹ we can then provide a lower bound for θ^* of the form²

$$\begin{aligned} \theta^* &= \arg \max_{\theta} \int_{\mathbf{z} \sim P_Z} \log \left(\frac{P(f_{\mathcal{O}}(\mathbf{z}) | \theta) P(\theta)}{P(f_{\mathcal{O}}(\mathbf{z}))} \right) dP_Z \\ &= \arg \max_{\theta} \left[\int_{\mathbf{z} \sim P_Z} \log P(f_{\mathcal{O}}(\mathbf{z}) | \theta) dP_Z \right. \\ &\quad \left. - \int_{\mathbf{z} \sim P_Z} \log P(f_{\mathcal{O}}(\mathbf{z})) dP_Z + \log P(\theta) \right] \\ &= \arg \max_{\theta} \left[\int_{\mathbf{z} \sim P_Z} \log P(f_{\mathcal{O}}(\mathbf{z}) | \theta) dP_Z + \log P(\theta) \right] \quad (3) \end{aligned}$$

where we drop the term $\int_{\mathbf{z} \sim P_Z} \log P(f_{\mathcal{O}}(\mathbf{z})) dP_Z$, which has no dependence on θ .

The solution to (3) depends on the form of the considered models. In this seminal article we study hard decision copies. Under this framework, we can recover regularized empirical risk minimization models [50] if we approximate the distributions above with an exponential family

$$P(f_{\mathcal{O}}(\mathbf{z}) | \theta) \propto e^{-\gamma_1 \ell_1(f_{\mathcal{C}}(\mathbf{z}, \theta), f_{\mathcal{O}}(\mathbf{z}))}; P(\theta) \propto e^{-\gamma_2 \ell_2(\theta, \theta^+)}$$

for $\ell_i(a, b)$ a measure of disagreement between a and b , and θ^+ our prior about θ . Using this approximation we can rewrite (3) as

$$\theta^* = \arg \min_{\theta} \left[\int_{\mathbf{z} \sim P_Z} \gamma_1 \ell_1(f_{\mathcal{C}}(\mathbf{z}, \theta), f_{\mathcal{O}}(\mathbf{z})) dP_Z + \gamma_2 \ell_2(\theta, \theta^+) \right] \quad (4)$$

The first term in this expression is the expected value of the disagreement between model and copy, which has the form of empirical risk minimization. The expected loss particularized to our copying problem can be defined as

$$R^{\mathcal{F}}(f_{\mathcal{C}}(\mathbf{z}, \theta), f_{\mathcal{O}}(\mathbf{z})) = \mathbb{E}_{\mathbf{z} \sim P_Z} [\ell_1(f_{\mathcal{C}}(\mathbf{z}, \theta), f_{\mathcal{O}}(\mathbf{z}))] \quad (5)$$

over the probability distribution P_Z . We refer to this value as the *fidelity error*. This error captures all the loss of copying. In the general form, it corresponds to the integral $\int_{\mathbf{z} \sim P_Z} \log P(f_{\mathcal{O}}(\mathbf{z}) | \theta) dP_Z$ in (3), *i.e.* the probability that the copy resembles the model.

The second term in (4) refers to the fit of the parameters to the prior and can be identified as the regularization term

$$\Omega(\theta) = \ell_2(\theta, \theta^+).$$

¹Jensen's inequality states that for any concave function f it holds that $E[f(X)] \leq f(E[X])$. In particular, for the $\log(x)$ function.

²Maximization of the lower bound also maximizes the original function. However, the optimal value of the lower bound may differ from that of the original objective function.

Under the empirical risk minimization framework we approximate the expected loss by the empirical risk. The particularization of the empirical risk to the copying setting corresponds to the *empirical fidelity error*, $R_{emp}^{\mathcal{F}}$. We define this value as the empirical version of the fidelity error

$$R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}(\mathbf{z}, \theta), f_{\mathcal{O}}(\mathbf{z})) = \frac{1}{N} \sum_{j=1}^N \ell_1(f_{\mathcal{C}}(\mathbf{z}_j, \theta), f_{\mathcal{O}}(\mathbf{z}_j)) \quad (6)$$

and rewrite (4) for the discrete case as follows

$$\begin{aligned} (\theta^*, \mathbf{Z}^*) &= \arg \min_{\theta, \mathbf{z} \sim P_Z} \left[R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}(\mathbf{z}, \theta), f_{\mathcal{O}}(\mathbf{z})) + \Omega(\theta) \right] \\ &= \arg \min_{\theta, \mathbf{z} \sim P_Z} \left[\frac{1}{N} \sum_{j=1}^N \gamma_1 \ell_1(f_{\mathcal{C}}(\mathbf{z}_j, \theta), f_{\mathcal{O}}(\mathbf{z}_j)) \right. \\ &\quad \left. + \gamma_2 \ell_2(\theta, \theta^+) \right], \quad (7) \end{aligned}$$

where \mathbf{Z} corresponds to the set of synthetic samples $\mathbf{z} \sim P_Z$. We refer to the set of labelled synthetic pairs $\mathcal{Z} = \{(\mathbf{z}_j, f_{\mathcal{O}}(\mathbf{z}_j))\}_{j=1}^N$ as the *synthetic dataset*. The expression above is a dual optimization, where we simultaneously optimize the copy parameters θ and the synthetic set \mathcal{Z} . This duality results from referring to the decision function $f_{\mathcal{O}}$ instead of exploiting the training data \mathcal{D} , and it fundamentally shapes how copying works.

C. SOLVING THE COPYING PROBLEM

The class membership predictions of $f_{\mathcal{O}}$ define a hard classification boundary. The resulting problem has two important characteristics: (i) the synthetic dataset is always separable and (ii) a potentially infinite stream of synthetic data is accessible. These features change the basic assumptions of traditional machine learning and can be exploited for solving the optimization problem in (7).

Because the synthetic set is separable, if we assume a copy with enough capacity, it is always possible to achieve zero empirical error, $R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}(\mathbf{z}, \theta), f_{\mathcal{O}}(\mathbf{z})) = 0$. The error then only depends on the generalization gap for the synthetic dataset. And since we can generate infinite synthetic data, this value can be asymptotically reduced to zero. Hence, in theory, copying can be performed without loss and redefined as the unconstrained optimization problem

$$\text{minimize}_{\theta, \mathbf{Z}} R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}(\mathbf{z}, \theta), f_{\mathcal{O}}(\mathbf{z})). \quad (8)$$

Yet, in practice, the synthetic set is finite. It therefore stands to reason to impose that the copy have small capacity, $\Omega(\theta)$, and rewrite the copying problem as

$$\begin{aligned} &\text{minimize}_{\theta, \mathbf{Z}} \Omega(\theta) \\ &\text{subject to } \|R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}, f_{\mathcal{O}}) - R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}^{\dagger}, f_{\mathcal{O}})\| < \epsilon, \quad (9) \end{aligned}$$

for $f_{\mathcal{C}}^{\dagger}$ the solution to (8) and ϵ a defined tolerance.³ The term $\|R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}, f_{\mathcal{O}}) - R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}^{\dagger}, f_{\mathcal{O}})\| < \epsilon$ defines a feasible

³In what follows, we favour a more concise notation and drop the explicit dependence on the synthetic data \mathbf{z} and copy parameters θ .

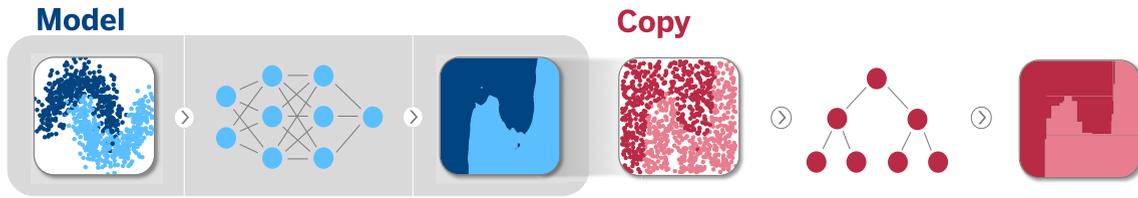


FIGURE 3. Example of the *single-pass copy*.

set of parameters. The solution to (9) achieves the smallest capacity while keeping $R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}, f_{\mathcal{O}})$ within a tolerance of the unconstrained optimal value of the empirical fidelity error, $R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}^{\dagger}, f_{\mathcal{O}})$. We argue that there exists a set of parameters θ that fulfill this constraint.

In some cases the optimal loss value is known in advance. Consider, for example, the hinge-loss in SVMs, where $R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}^{\dagger}, f_{\mathcal{O}}) = 0$. However, this is not always the case, e.g. least-square errors in classification.⁴ Copying is different from the standard multi-objective optimization in a pure learning setting, where the optimal values of both the loss and the regularization term are unknown. Instead of having a *Pareto's surface* of plausible optimal solutions, as long as $\Omega(\theta)$ is convex, the solution to (9) is unique.

This optimization can be straightforwardly solved in cases where the capacity is directly modelled, such as those of SVMs and neural networks, using a regularization function, or Bayesian models, selecting the priors. For other models, such as trees, the complexity control must be done by either early stopping or by an external process, such as post- or pre-pruning. Finally, techniques such as boosting or deep learning may exhibit a delayed overfitting effect [51]–[53]. A property that can be exploited to our advantage to directly solve (8) instead of (9).

D. THE SINGLE-PASS COPY

Conducting a simultaneous optimization of the synthetic data and the copy parameters requires the copy hypothesis space to have certain properties, such as online updating. This challenging issue is out of the scope of this paper and requires further research. Hence, for the sake of simplicity, in the rest of this article we consider the simplest approach to solving the dual copying problem: the *single-pass copy*. We cast the simultaneous optimization problem into one where only a single iteration of an alternating projection optimization scheme is used. This effectively splits the problem in two independent sub-problems:

Step 1: Synthetic sample generation. The first step is to find the optimal set \mathcal{Z}^* . This set is that for which the empirical fidelity error, $R_{emp}^{\mathcal{F}}$, is minimal

$$\mathcal{Z}^* = \arg \min_{\mathcal{Z}} R_{emp}^{\mathcal{F}}$$

As a result, we obtain the optimal synthetic dataset \mathcal{Z}^* .

⁴Instead of tracking the empirical risk we can track the empirical error, which can be set to zero due to the separability property.

Step 2: Building the copy. Once having generated and labelled the set \mathcal{Z}^* , the next step is to find θ^* such that

$$\begin{aligned} & \text{minimize } \Omega(\theta) \\ & \text{subject to } \|R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}, f_{\mathcal{O}}) - R_{emp}^{\mathcal{F}}(f_{\mathcal{C}}^{\dagger}, f_{\mathcal{O}})\| < \epsilon, \end{aligned}$$

or its simplified version (8), provided that the adequate conditions hold.

An example of the single-pass copy is shown in Fig. 3, where the binary decision function learned by a fully-connected neural network is copied with a decision tree classifier. The tree-based copy is built using a set of synthetic samples drawn from a uniform distribution and labelled according to the hard predictions output by the neural net.

IV. MEANINGFUL INSIGHTS

In what follows, we bridge the gap between theory and practice by using toy problems to draw relevant conclusions from the derivation above. We focus on the two steps of the single-pass copy: we begin by studying the synthetic sample generation process and then show how the particularities of copying problem can be exploited during copy building.

A. STEP 1: SYNTHETIC SAMPLE GENERATION

For the sake of this discussion, let us consider a binary classification problem and let $f_{\mathcal{O}}(\mathbf{z}) \in \{-1, +1\}$ and $f_{\mathcal{C}}(\mathbf{z}, \theta) \in \{-1, +1\}$, for any $\mathbf{z} \in \mathcal{X}$. Let us also consider the case where ℓ_1 corresponds to the 0/1 loss. For this case, the empirical fidelity error in (6) can be rewritten as

$$\begin{aligned} R_{emp}^{\mathcal{F}} &= \frac{1}{2N} \sum_{j=1}^N |f_{\mathcal{O}}(\mathbf{z}_j) - f_{\mathcal{C}}(\mathbf{z}_j, \theta)| \\ &= \frac{1}{2N} \sum_{j=1}^N |f_{\mathcal{O}}(\mathbf{z}_j)| \left| 1 - \frac{f_{\mathcal{C}}(\mathbf{z}_j, \theta)}{f_{\mathcal{O}}(\mathbf{z}_j)} \right| \\ &= \frac{1}{2N} \sum_{j=1}^N \left(1 - \frac{f_{\mathcal{C}}(\mathbf{z}_j, \theta)}{f_{\mathcal{O}}(\mathbf{z}_j)} \right) \\ &= \frac{1}{2N} \sum_{j=1}^N 1 - \frac{1}{2N} \sum_{j=1}^N \frac{f_{\mathcal{C}}(\mathbf{z}_j, \theta)}{f_{\mathcal{O}}(\mathbf{z}_j)} \\ &= \frac{1}{2} - \frac{1}{2N} \sum_{j=1}^N f_{\mathcal{C}}(\mathbf{z}_j, \theta) f_{\mathcal{O}}(\mathbf{z}_j), \quad \mathbf{z}^{(N)} \sim P_{\mathcal{Z}} \end{aligned}$$

Let us now define a partition of the space such that $\mathcal{X} = \mathcal{X}_+ \cup \mathcal{X}_-$ and $\mathcal{X}_+ \cap \mathcal{X}_- = \emptyset$, where $\mathcal{X}_+ = \{\mathbf{z} | \mathbf{z} \in \mathcal{X},$

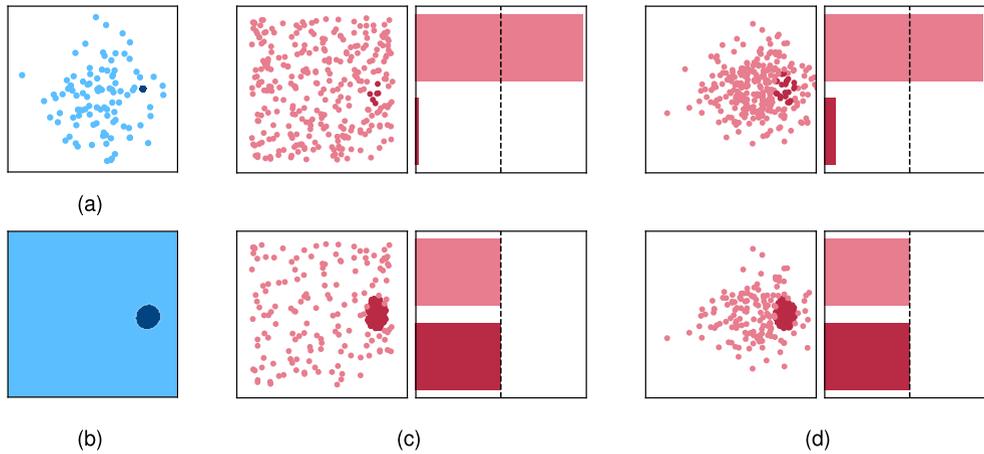


FIGURE 4. (a) Training dataset. (b) Decision boundary learned by a Gaussian Process classifier. From top to bottom raw and balanced synthetic datasets generated from (c) a uniform distribution and (d) a standard normal distribution.

$f_{\mathcal{O}}(z) = 1\}$ and $\mathcal{X}_- = \{z|z \in \mathcal{X}, f_{\mathcal{O}}(z) = -1\}$ are the two sub-spaces defined by the model. We rewrite the equation above in terms of this partition as

$$R_{emp}^{\mathcal{F}} = \frac{1}{2} - \frac{1}{2N_+} \sum_{j=1}^{N_+} f_{\mathcal{C}}(z_j, \theta) + \frac{1}{2N_-} \sum_{j=1}^{N_-} f_{\mathcal{C}}(z_j, \theta)$$

for N_+ and N_- the number of samples lying in \mathcal{X}_+ and \mathcal{X}_- , respectively.

We define the probability of a sample lying in \mathcal{X}_+ as $p_+ = \mathbb{P}(z \in \mathcal{X}_+)$ and the probability of a sample lying in \mathcal{X}_- as $p_- = \mathbb{P}(z \in \mathcal{X}_-)$. These two probabilities depend on the *size* of the positive and negative domains. In particular,

$$p_+ = \int_{z \in \mathcal{X}_+} P_Z(z) dz, \quad p_- = \int_{z \in \mathcal{X}_-} P_Z(z) dz.$$

With these quantities, we can see that $N_+ = Np_+$ and $N_- = Np_-$. Thus,

$$R_{emp}^{\mathcal{F}} = \frac{1}{2} - \frac{1}{2Np_+} \sum_{j=1}^{Np_+} f_{\mathcal{C}}(z_j, \theta) + \frac{1}{2Np_-} \sum_{j=1}^{Np_-} f_{\mathcal{C}}(z_j, \theta).$$

Minimization of this expression explicitly depends on the form of P_Z . In the simplest case, we can assume this distribution to be flat on the domain \mathcal{X} , so that $z \sim \mathcal{U}(\mathcal{X})$. Under this assumption, p_+ and p_- correspond to the fraction of volume for each of the classes. Recalling the form of the error for the Monte Carlo estimator under this distribution, we can express the standard error for $R_{emp}^{\mathcal{F}}$ as

$$\sigma(R_{emp}^{\mathcal{F}}) \propto \mathcal{O}\left(\frac{1}{\sqrt{Np_+}} + \frac{1}{\sqrt{Np_-}}\right).$$

We exploit this expression to extract relevant insights for the synthetic sample generation process. First, we confirm the need to define an attribute representation \mathcal{X} . This is a reasonable assumption, since we need to have an approximate

idea of the dynamic range of all variables in order to build meaningful queries.

Second, we note that in some situations there might be a mismatch between the decision boundary achievable by the copy and $f_{\mathcal{O}}$. As a consequence, a given synthetic dataset may not perform equally for different copy hypotheses. Consider a non-linear decision function and a linear copy model. Exploring the twists of the decision boundary during the synthetic sample generation process may not be relevant in this situation. Thus, we should consider the properties and assumptions of the copy hypothesis space to effectively exploit each generated sample.

Another important issue is that of volume imbalance, which arises when one or more of the classes occupy a region of the space much smaller than the rest.

1) THE ISSUE OF VOLUME IMBALANCE

The empirical fidelity error depends on the fraction of volume occupied by each decision region. If the spatial support of one class is small with respect to the total volume, it may be difficult to have a meaningful number of samples on that region, resulting in large approximation errors.

In Fig. 4(a), we show a binary dataset with a balanced label distribution. Points belonging to the class depicted in lighter blue are spread out throughout the plot, while those corresponding to the darker class are packed more densely in a smaller area. Despite the number of instances per class being equal, note that there are notable differences in the volume of each of the classes. The resulting decision function is displayed in Fig. 4(b).

To copy this model, we assay two different forms for P_Z . In a preliminary approach, we generate samples at random until we reach a desired number of points. In Fig. 4(c) and Fig. 4(d) we plot the sets that result for a uniform distribution and for a standard normal distribution, respectively. The resulting data, shown together with their corresponding

label distribution, are notably imbalanced: there is one class for which we only recover a few number points. This result is unrelated to class distribution.

Fortunately, the volume imbalance effect can be alleviated either by a good choice of P_Z or by imposing that the resulting set be balanced. For example, we can try to infer a sampling distribution that allocates a large amount of the probability mass around the unknown decision boundary. Due to its complexity, we believe the problem of finding an optimal P_Z to be out of the scope of this work. This issue will be subject to further analysis in future contributions. Indeed, in a recent paper [54] we have studied different sampling algorithms for the copying setting, including a technique that focuses on boundary exploration, a Bayesian-based optimizer, a modified version of the Jacobian approach proposed by [15] and raw random sampling.

Alternatively, we can overcome the issue of volume imbalance using heuristics that balance a general exploration of the space with exploitation around the areas of interest. Hence, we impose that the resulting set be balanced with respect to the class labels. We force the data generator to focus on those areas where the misrepresented class is located, to ensure that all labels are well represented in the resulting set, as shown in Fig. 4.

B. STEP 2: BUILDING THE COPY

The second part of the alternating projection scheme corresponds to finding the optimal parameters for the copy. For illustration purposes, consider a radial basis function kernel SVM. This model is defined by a kernel function of the form $\mathcal{K}(\mathbf{x}, \mathbf{x}') = e^{-\gamma \|\mathbf{x} - \mathbf{x}'\|^2}$, where $\|\mathbf{x} - \mathbf{x}'\|^2$ corresponds to the squared Euclidean distance, and γ is the inverse of the radius of influence of the support vectors, i.e. the width of the kernel. This means, in essence, that γ controls the capacity: the larger its value, the higher the complexity. In other words, minimizing the model capacity in (9) amounts to minimizing γ . In Fig. 5 we show how this can be exploited in practice to copy the neural net in Fig. 3 using synthetic samples drawn at random from a uniform distribution.

In particular, Fig. 5(a) shows the copy decision function for a maximal value of γ , such that the second term in (9) is satisfied and the empirical error is zero. Fig. 5(b) shows the decision boundary for a copy with optimal capacity γ , computed for a tolerance $\epsilon = 1e - 4$. This solution results from sequentially reducing the value of γ and monitoring the change in accuracy until the error deviation is greater than ϵ . When comparing both plots we observe the improvement in generalization performance. This improvement is also seen in Fig. 5(c), where train and generalization errors of the copy are shown for decreasing values of γ . For a bounded value of the empirical error, the generalization error is reduced as we decrease the capacity of the copy.

Unlike the classical machine learning, where capacity is optimized during the validation step, this result shows that it is possible to optimize the capacity of a copy during training.

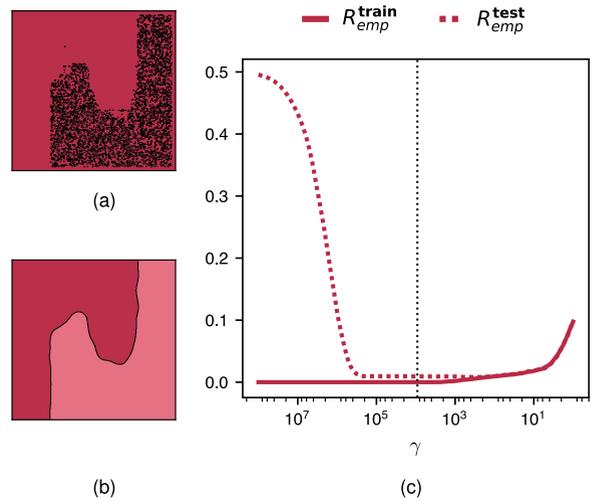


FIGURE 5. Decision boundaries learned by copies with (a) a maximal and (b) an optimal γ . (c) Empirical risk and generalization error for decreasing values of γ . The dotted black line refers to the optimal γ .

This has a profound impact on how copying is performed and shows that copying can further exploit the particularities of this problem allowing for a more effective and efficient building of algorithms to handle the copying process than using standard machine learning pipelines and assumptions.

1) CAPACITY ERROR

Lastly, note that the specific choice of copy hypothesis has a significant impact on performance. Different capacity copies may behave very differently when confronted with the same set of synthetic data points.

We refer to the capacity of a classifier as a measure of its complexity. A mismatch of capacity between model and copy can lead to poor performance results, even in cases where the synthetic dataset properly covers the input space. Take for example the case of a linear logistic regression and a support vector machine. The decision functions resulting from building copies based on these two architectures are notably different. Given the same set of synthetic points, the logistic model may not be able to fully recover the form of the considered decision boundary if this is non-linear. This is because the original classifier, is not contained in the new hypothesis space. In the case of the SVM, the mismatch in capacity is presumably not so pronounced and therefore the copy decision boundary may be much more precise.

V. EMPIRICAL VALIDATION

In this section we present our experiments to empirically validate copies in a variety of well-known problems that include a diverse selection of UCI datasets with different number of classes and dimensions. We begin by proposing a set of performance metrics.

A. PERFORMANCE METRICS

When evaluating copies, we may ask questions of the form: “what does the performance on a synthetic validation set tell

us about the generalization of the copy?”, “does the copy have enough capacity to replicate the decision function?” or, more generally, “what metrics should we use to evaluate copies in terms of the available information?”. In what follows we introduce a set of definitions aimed at answering these questions.

1) EMPIRICAL FIDELITY ERROR

We particularize the empirical fidelity error in (6) to the 0/1 loss and measure it over the synthetic set \mathcal{Z} as

$$R_{emp}^{\mathcal{F}, \mathcal{Z}} = \frac{1}{N} \sum_{j=1}^N \mathbb{I}[f_{\mathcal{O}}(z_j) \neq f_{\mathcal{C}}(z_j)] \quad (10)$$

for \mathbb{I} the indicator function. In resorting to Monte Carlo integration we here necessarily incur in an approximation error that depends, among other things, on the quality of the set \mathcal{Z} . As a result, a low $R_{emp}^{\mathcal{F}, \mathcal{Z}}$ is no absolute guarantee of a good copy. For this value to be a valid assessment of the total error, the synthetic dataset must be large enough to ensure coverage of the input space and the volume imbalance effect needs to be controlled for.

In cases where the constraints of the copying scenario are relaxed and the training data \mathcal{D} is accessible, we could also evaluate the empirical fidelity error over this set as

$$R_{emp}^{\mathcal{F}, \mathcal{D}} = \frac{1}{M} \sum_{i=1}^M \mathbb{I}[f_{\mathcal{O}}(x_i) \neq f_{\mathcal{C}}(x_i)] \quad (11)$$

For validation purposes, in the following we assume these data to be known. In general, $R_{emp}^{\mathcal{F}, \mathcal{D}}$ and $R_{emp}^{\mathcal{F}, \mathcal{Z}}$ yield very different values. This difference arises from the mismatch between the probability density functions P and P_Z .

2) COPY ACCURACY

To evaluate the copy generalization performance over \mathcal{D} we introduce the *copy accuracy*, $\mathcal{A}_{\mathcal{C}}$, as follows

$$\mathcal{A}_{\mathcal{C}} = \frac{1}{M} \sum_{i=1}^M \mathbb{I}[t_i = f_{\mathcal{C}}(\mathbf{x}_i)], \quad (12)$$

for $t \in \mathcal{T}$ the true labels. The performance of the copy on \mathcal{D} is bounded by $\mathcal{A}_{\mathcal{O}}$, the accuracy of $f_{\mathcal{O}}$ on these data. In the ideal case the fidelity error is zero, so that $\mathcal{A}_{\mathcal{C}} = \mathcal{A}_{\mathcal{O}}$. In general, we can use the empirical fidelity error over the synthetic set to approximate $\mathcal{A}_{\mathcal{C}}$ by means of the *estimated copy accuracy*, $\widehat{\mathcal{A}}_{\mathcal{C}}$, as follows

$$\widehat{\mathcal{A}}_{\mathcal{C}} = \mathcal{A}_{\mathcal{O}}(1 - R_{emp}^{\mathcal{F}, \mathcal{Z}}) \quad (13)$$

B. EXPERIMENTS

We use 60 datasets from the UCI Machine Learning Repository database [55]. We refer the reader to [56] for a specific description of initial data selection and preprocessing. We select those datasets with more than 100 samples and a frequency above 10% for all class labels. We also require the number of inputs to be greater than double the number

of attributes. Among the selected datasets 42 correspond to binary classification problems and 18 are multiclass.

1) EXPERIMENTAL SET UP

We convert nominal attributes to numerical and re-scale variables to zero mean and unit variance. We split data into stratified 80/20 training and test sets. We use 6 state-of-the-art classification algorithms, including adaboost (*adaboost*), an artificial neural network (*ann*), a random forest (*random_forest*), a linear SVM (*linear_svm*), a SVM with a radial basis function kernel (*rbf_svm*) and a gradient-boosted tree (*xgboost*). To avoid bias regarding the algorithm choice, we sort datasets in alphabetical order, group them in sets of 10 and randomly assign a classifier to each group.

We build a generic pipeline and train all models using a cross-validated grid-search over a fixed parameter grid. Three classifiers learn decision functions that exclude at least one of the class labels. This occurs for *pittsburg-bridges-REL-L*, for which only two of the three classes are learned, and *planning* and *statlog-australian-credit*, for which a single class label is assigned to all data points. Besides, because we use a fixed pipeline, not all models yield an optimal performance. See, for example, the case of *echocardiogram*, where accuracy is equal to 0.3.

We keep this result for two reasons. First, we want the experimental setup to be as agnostic as possible and hence the random pairing of models and datasets. Second, it reinforces an important idea: a copy can only be as good as the model it aims to replicate. Or in the other words, the baseline for the copy performance is the original model performance. Non-optimal models lead to poorly performing copies. We stress, nonetheless, that in a real setting one would be interested in copying only those models that perform reasonably well.

We draw $1e6$ random samples from a uniform distribution to generate balanced synthetic sets. We identify three cases of volume imbalance: *congressional-voting*, *ilpd-indian-liver* and *statlog-image*. Despite the training data being balanced with respect to class distribution, we only recover a small fraction of samples for one or more of the labels. As previously mentioned, this could lead to sub-optimal results, given that the copy tends to wrongly classify points that belong to the subsampled classes. Imposing that the synthetic dataset be balanced mitigates this issue to a great extent and ensures that the copy treats all labels equally.

To evaluate the impact of heuristics, we assay different copy model hypotheses. We use decision trees because they are easily interpretable, logistic regression because it is a linear model and random forest as an example of a bagging method. We copy using no cross-validation or hyperparameter tuning: trees are grown until each leaf contains a single sample and neural networks and boosting methods are trained with no regard for generalization. For validation purposes, we run each experiment 100 times and report averages over all repetitions for the true and the estimated

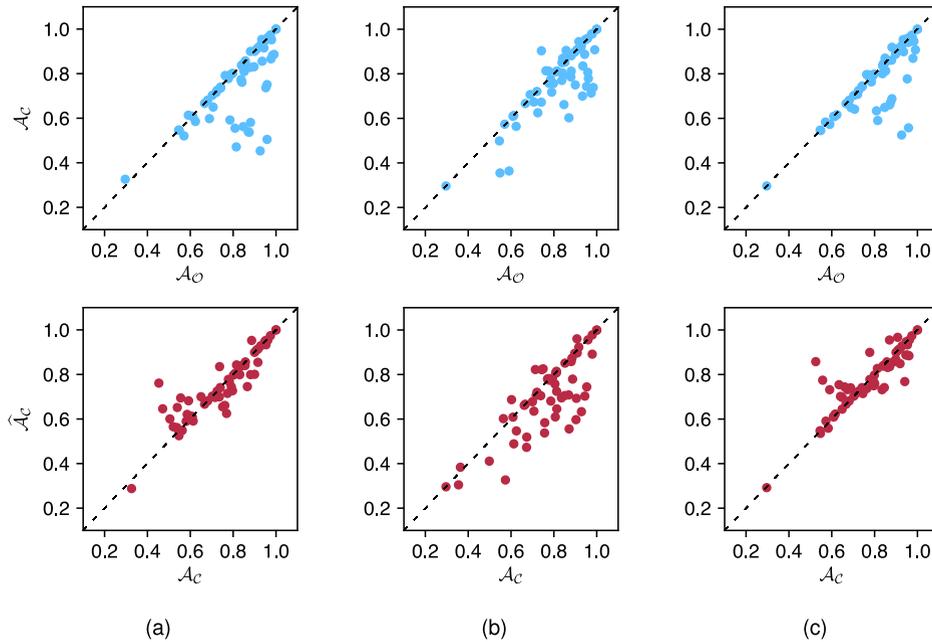


FIGURE 6. From top to bottom, distribution of average copy accuracy against original accuracy and distribution of average estimated copy accuracy against average true copy accuracy for all datasets and for copies based on (a) decision trees, (b) logistic regression and (c) random forest.

copy accuracy. We also report the mean empirical fidelity error measured over both training and synthetic data.

2) RESULTS

The measured performance metrics are shown in Fig. 6. In particular, Fig. 6(a), Fig. 6(b) and Fig. 6(c) show the distribution of the mean copy accuracy \mathcal{A}_C against the original accuracy \mathcal{A}_O and the estimated copy accuracy $\widehat{\mathcal{A}}_C$ for all datasets and copies based on decision trees (*decision_tree*), logistic regression (*logistic_regression*) and random forest (*random_forest*) classifiers, respectively.

Results for both *decision_tree* and *random_forest* are scattered around the main diagonal, whereas copies based on *logistic_regression* show a greater dispersion; especially when comparing \mathcal{A}_C to $\widehat{\mathcal{A}}_C$. In general, the value of $\widehat{\mathcal{A}}_C$ is smaller than \mathcal{A}_C , which means that the empirical fidelity error over the synthetic data overestimates the real error. This is in part due to the difference in the distributions P and P_Z . When evaluating $R_{\mathcal{F}}^{\mathcal{F}}$, we measure the performance of the copy in the space defined by P_Z , so that we may penalize the copy for errors in regions where there are no actual training data.

The complete summary of results for all problems and copy algorithms is shown in Table 3 in the Appendix. In most problems, results show the ability of copies to replicate the target decision behaviour. Overall, copy accuracy is competitive for the proposed synthetic dataset size and the estimated copy accuracy provides a reliable approximation to the accuracy of the copy in real data. The empirical fidelity error generally yields values close to 0, which indicates that copies are correctly built.

Table 1 shows a selected set of results. There are several datasets where there is no degradation when using a *logistic_regression* to copy higher capacity models such as *ann* or *xgboost*. This is the case, for example, with *breast-cancer-wisc* and *wine*, where \mathcal{A}_C is reasonably close to \mathcal{A}_O , even while the logistic model can only learn linear relations among attributes. We take this as an indication that the initial classifiers were too complex for the relatively simple problems. Copying here allows us to move to a more suitable solution, with less parameters and training requirements.

On the other hand, we identify a number of cases where copies based on *decision_tree* and *random_forest* clearly outperform *logistic_regression*. See, for example, *energy-y1* and *iris*. This is because when the decision function is not linear⁵, non-linear copies are needed. Here, the error due to a mismatch of capacity dominates, because the copy hypothesis space, the logistic family, does not contain f_O .

Finally, in some instances the copy hypothesis space is well chosen and yet the empirical fidelity error is high. See for example *musk_1* and *musk_2*, which are both high dimensional problems where a *linear_svm* is copied using a *random_forest*. In both cases, \mathcal{A}_C is notably lower than \mathcal{A}_O . This happens in complex datasets, where $1e6$ synthetic data points are probably not enough to ensure a small $R_{emp}^{\mathcal{F}}$.

C. DISCUSSION

The different error contributions are collectively defined by the fidelity error and approximated through the empirical

⁵Despite the training data being linearly separable, the learned decision boundary may be non-linear.

TABLE 1. Subset of Relevant Results for the UCI Experiment. (*) Multiclass. (†) More than 10 dimensions. (§) Standard deviation above 0.05.

Dataset	\mathcal{H}_O	\mathcal{A}_O	decision_tree			logistic_regression			random_forest		
			\mathcal{A}_C	$\hat{\mathcal{A}}_C$	$R_{emp}^{\mathcal{F},\mathcal{D}}$	\mathcal{A}_C	$\hat{\mathcal{A}}_C$	$R_{emp}^{\mathcal{F},\mathcal{D}}$	\mathcal{A}_C	$\hat{\mathcal{A}}_C$	$R_{emp}^{\mathcal{F},\mathcal{D}}$
breast-cancer-wisc	<i>adaboost</i>	0.93	0.93	0.9286	0.00	0.93	0.6333	0.06	0.93	0.9286	0.00
chess-krvkp [†]	<i>ann</i>	0.99	0.89	0.9527	0.11	0.91	0.9603	0.10	0.91	0.9670	0.09
echocardiogram	<i>ann</i>	0.3	0.33	0.2879	0.05	0.30	0.2960	0.00	0.30	0.2922	0.00
energy-y1*	<i>ann</i>	0.96	0.96	0.9537	0.00	0.78	0.7744	0.23	0.96	0.9551	0.00
iris*	<i>random_forest</i>	0.93	0.95	0.9332	0.02	0.70	0.6778	0.30	0.95	0.9333	0.02
musk-1 [†]	<i>linear_svm</i>	0.88	0.54	0.5620	0.46 [§]	0.88	0.8732	0.01	0.67 [§]	0.7323 [§]	0.32
musk-2 [†]	<i>linear_svm</i>	0.96	0.50 [§]	0.6005	0.50 [§]	0.96	0.9556	0.00	0.56	0.7745	0.44
oocytes_merluccius_nucleus_4d [†]	<i>linear_svm</i>	0.82	0.47 [§]	0.6460	0.52 [§]	0.81	0.8144	0.00	0.59	0.7317	0.38
wine* [†]	<i>xgboost</i>	0.92	0.92	0.9147	0.00	0.94	0.7031	0.08	0.92	0.9147	0.00

fidelity error. However, the condition that empirical fidelity error be small is necessary, but not sufficient. Having significant errors in certain regions and none in others may lead to a low error, while altogether not ensuring a good generalization performance. The opposite is also true: a large empirical fidelity error may not lead to a low copy accuracy. Take, for example, errors distributed around the boundary. This may happen when trying to copy a smooth function using linear decision cuts. If errors are very substantial, this may be seen as a problem. However, if the training data are distributed far away from the boundary, errors in this region would have no real impact. No effective error would therefore be measured when substituting the model with the copy.

To a large extent, copy evaluation depends on the available information. The more information we have, the more reliable our estimates will be. If the training data were accessible, we could obtain a direct estimate of the copy generalization performance. Furthermore, we could choose P_Z to be as close to P as possible, *i.e.* redefine the copy operation space to match P . If the form of the model was also known, we could refine the choice of copy hypothesis. In those cases where model and copy have similar decision boundary shapes, copying is conducted with greater ease. That is, when the decision function is formed of cuts perpendicular to the axes, *i.e.* it is a random forest, it is easier to copy with a decision tree than it is with a radial basis kernel SVM. Conversely, those models with smooth decision functions are better copied using classifiers other than trees.

At this stage, we may ask ourselves the question: *if the training data are available why copy instead of learning a new classifier?* There exist scenarios where a new training may not be advisable. A new model may display very different behaviour and decision properties. This is unacceptable in production environments where performance has to be preserved and controlled. Moreover, training a new classifier with the training data involves having to take care of the overfitting effect. As shown in Sec. IV, when copying we can avoid the hyper-parameter optimization step.

Another reason to use copies is that when training a new model, we might not be able to recover the same operation point as before. In contrast, as explained in Sec. VI, a copy

can help bias the parameter optimization process towards a desired solution.

In general, copies can be understood as a tool to bridge the gap between accuracy and any other desired property. Copying helps in breaking the trade-offs we face in training high-performance models when characteristics such as interpretability, simplicity or compliance are required.

VI. APPLICATIONS AND LIMITATIONS

Having demonstrated the feasibility of copying and discussed its main characteristics, in this section we elaborate on its utility in a wide variety of scenarios. We present three use cases with real-life applications of copying. Further, we analyse shortcomings and discuss different approaches to overcoming the identified barriers.

A. APPLICATIONS

One of the main benefits of copying is that it enables differential replication of models. This means that copies can be used to enhance existing solutions. They can, for example, be used to evolve from batch to online learning schemes [57]. This extends a model's lifespan as it enables adaptation to data drifts or performance deviations. Equivalently, when new class labels appear during a model's deployment in the wild, copies can account for the new data points and evolve from binary to multiclass classification settings [58]. More generally, there are numerous examples where differential replication can be applied to solve specific problems. In the following lines, we describe some of them and discuss how copies could be useful in addressing these issues.

1) INTERPRETABILITY

Recent advances in the field of machine learning have led to increasingly sophisticated models, capable of learning ever more complex problems to a high degree of accuracy. This comes at the cost of simplicity [59], [60], a situation that stands in contrast to the growing demand for transparency in automated processing [4]–[6]. Recent papers have shown that the knowledge acquired by black-box solutions can be transferred to interpretable models such as trees [27], [28], [61], rules [62] and decision sets [63]. In the copying scenario

models of any arbitrary type can be substituted by copies specifically designed to be globally self-explanatory.

2) PRODUCTION

Model deployment is often costly in company environments [10], [64], [65], [66]. Common issues include the inability to maintain the technological infrastructure up-to-date with latest software releases, conflicting versions or incompatible research and deployment environments. Consider the case of neural network library Tensorflow. Despite the library itself provides detailed instructions on how to serve models in production [67], this typically requires several third-party components for docker orchestration, such as Kubernetes or Elastic Container Service [68], which are seldom compatible with on-premise software infrastructure. Moving to a copy in a less demanding environment helps bridge the gap between the data science and engineering departments.

3) FAIRNESS AND AUDITING

Machine learning models can reproduce existing patterns of discrimination [7], [9]. Some algorithms have been reported to be biased against people with protected characteristics like race [69]–[72], gender [73], [74] or sexual orientation [75]. Under these circumstances distillation has been shown to be useful for model auditing [76] and so have copies. Upon them, *desiderata* such as equity of learning can be directly imposed to, for example, reduce the biased of trained classifiers.

B. USE CASES

In what follows we demonstrate some of these non-trivial applications in real-life scenarios. First, we derive regulatory-compliant high-performing copies for non-client mortgage loan default prediction in a private dataset from BBVA. Second, we use copies to recover the operation point of a model trained on borrower information from the Lending Club website [77]. Lastly, we study how copies can be applied to obtain a fair classification of alignment in the superheroes dataset [78].

1) RISK SCORING FOR NON-CLIENT MORTGAGE LOANS

Logistic regression is a widely established technique for credit risk scoring. Mainly because it performs relatively well on credit prediction settings. But also because it offers the additional advantage of a relative ease of interpretation to comply with regulatory requirements. Even so, models based on logistic regression fail to account for non-linearities in the data, which are usually modelled during an increasingly complex preprocessing step.

During this step, which is critical to maximize business objectives, domain knowledge is exploited to artificially generate a set of highly predictive attributes. Here, a qualified risk analyst is required to conduct a tedious process of trial and error to find an optimal set of variables. This incurs in a large economical cost and a delayed time-to-market delivery. Even worse, preprocessing largely reduces interpretability:

new variables often reflect complex relations among attributes and therefore remain non-decomposable [60] as far as the regulators are concerned.

In what follows, we tackle these issues in two different scenarios. In the first, we use a set of hand-crafted attributes to predict credit default using a logistic regression. We then build a copy that remains interpretable while retaining predictive performance. In the second, we decrease time-to-market delivery by training a high capacity model that avoids the preprocessing step. We copy this model with a simpler architecture that is nonetheless compliant with production and regulatory requirements.

In both cases, we use a private dataset of non-client⁶ mortgage loan applications recorded during 2015 all over Mexico [79]. This dataset consists of 19 attributes for 1.328 loan applicants, among which only 77% paid it off.

a: DEOBFUSCATED RISK SCORING MODELS

We emulate a standard production pipeline and preprocess the data to obtain 6 carefully crafted variables. We then train a logistic regression that achieves an accuracy of 0.77. We copy this whole predictive system, composed of both the preprocessing module and the logistic model, using a decision tree classifier. Fig. 7 shows the distribution of scores for this experiment. We obtain an averaged copy accuracy of 0.71 ± 0.04 and an estimated copy accuracy of 0.74314 ± 0.00018 . The mean empirical fidelity errors over \mathcal{L} and \mathcal{D} are 0.03488 ± 0.00018 and 0.15 ± 0.05 , respectively.

The empirical fidelity error over the synthetic data is small. However, when computed over the original test set this error grows. We argue that if we were to increase the number of synthetic samples, and better explore the boundaries, the approximation error would converge to a more reliable value and the overall error would be reduced.

In this example, the copy uses the deobfuscated 19 variables. Thus, the problem of non-decomposability is effectively solved. For validation purposes, in Fig. 7(a) we show the accuracy of a decision tree classifier trained directly on the training data. Note that it is smaller than that of our copy. This shows an additional advantage of copying: it can be used to guide a certain model to a more optimal solution in its parameter space.

b: HIGH-PERFORMANCE REGULATORY COMPLIANT COPIES

In this scenario, we use a high capacity model without any preprocessing. We train a gradient-boosted tree with all the 19 attributes in the training dataset. This model achieves an original accuracy of 0.79. We copy it using a decision tree classifier and report the results in Fig. 8. The mean copy accuracy averaged over all runs is 0.74 ± 0.02 and the accuracy estimated using (13) is equal to 0.7194 ± 0.0003 . Thus, the average empirical fidelity error is 0.09 ± 0.0003 and the average empirical fidelity error over \mathcal{D} is 0.09 ± 0.02 .

⁶The term non-client refers to those individuals who had no previous contractual relation with the bank at the time of loan application.

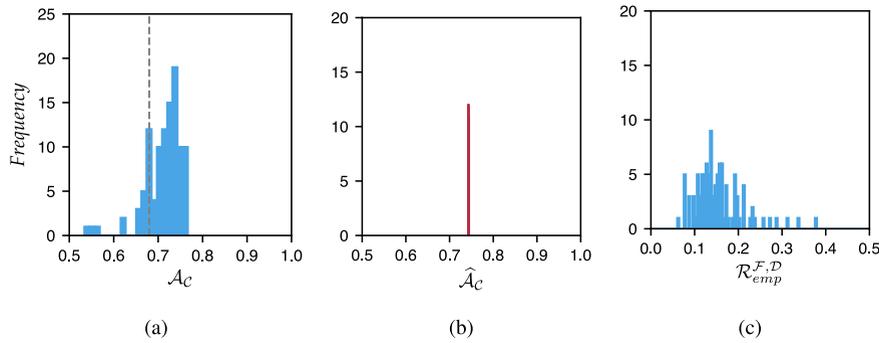


FIGURE 7. Distribution of values computed for the *scenario 1* for (a) the true copy accuracy, (b) the estimated copy accuracy and (c) the empirical fidelity error over the training data. For comparison purposes, the accuracy of a decision tree trained on original data is shown in black in (a).

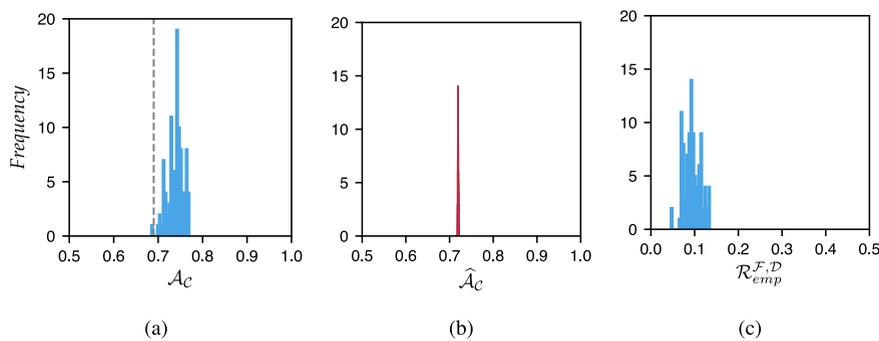


FIGURE 8. Distribution of values computed for the *scenario 2* for (a) the true copy accuracy (b) the estimated copy accuracy and (c) the empirical fidelity error over the training data. For comparison purposes, the accuracy of a decision tree trained on original data is shown in black in (a).

Note that while final model attributes differ from this application to that of *scenario_1*, the same samples are shared in both cases, so as to minimize any bias regarding the specific choice of data.

The difference in performance between the preprocessed logistic model in *scenario_1* and the copy decision trees in *scenario_2* is minor when tested against the test data. In Fig. 8(a) we display the accuracy achieved by a decision tree trained directly on the training data. This value is equal to 0.69 ± 0.01 . Comparison between this result and the mean true copy accuracy for this problem provides further evidence for the benefits of using copies in this context.

2) RESTORING FULL OPERATIONAL POTENTIAL IN ONLINE LOAN DEFAULT PREDICTION

For predicting whether a potential borrower will repay a loan, the Lending Club website publishes statistics about individual loan applicants [77]. We use these data to show how copies can be used to move a trained classifier to an online setting and recover the original operation point.

The complete dataset contains a comprehensive list of attributes for all loans issued through the 2007-2015 period, including loan status, latest payment information, number of finance inquires, borrower's annual income or zip code, among others. We remove null and missing values and drop

all fields which provide no useful information for inference. We also identify and drop all variables that cause data leakage as those that are typically not available at the time of prediction [80]. Finally, we label instances by classifying all loans identified as defaulted, charged off or late as *bad*. The resulting database consists of 50 attributes for 887,379 loans, divided into two classes.

We train a denseNet neural network [81] consisting of 5 hidden layers with 256, 128, 64, 32 and 16 neurons. We use self-normalizing units [82] to avoid internal covariate shift, a dropout rate of 10% and a least squared loss optimized using Adam. Because training data are highly imbalanced, with *bad* loans accounting only for 8% of the data, we use balanced batches. We choose our operation point to be that for which the recall values for both classes are closer to each other. Accuracy is equal to 0.63 and recall is 0.59 and 0.63 for the *bad* and *good* classes, respectively.

We copy this model using a neural net with a much simpler architecture, consisting of five fully connected layers with 256, 128, 64, 32 and 16 selu neurons, no dropout and a least-square loss with a default parameter Adam optimizer. We obtain a mean copy accuracy of 0.63 ± 0.07 . The estimated copy accuracy is 0.603 ± 0.009 , the empirical fidelity error is 0.042 ± 0.009 and the empirical fidelity error over the training data is 0.45 ± 0.07 . The copy recall distribution over these data

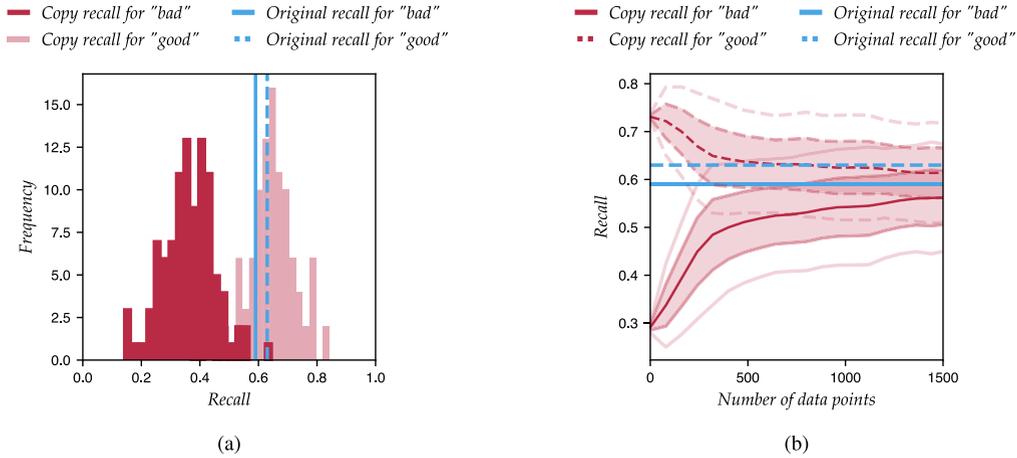


FIGURE 9. Distribution of (a) recall scores for model and copy and (b) recall scores over the number of new training data points for both classes.

is shown in Fig. 9(a), for both classes. We correctly recover the recall operating point for one of the classes, but suffer a loss of around 20% for the other.

We conclude that we can build copies with online capabilities, while retaining most of the accuracy and reaching a reasonably close operating point. Moreover, in the presence of new data points, copies can be fine tuned to achieve a new desirable operating point, as shown in Fig. 9(b). Here, we recover an equal rate of 59% after visiting a few hundred examples of the training data. It is worth noting that this example also shows that copies can serve as analysis tools for other models. In particular, we observe that the denseNet and the fully connected architectures both have very similar operation points.

3) A FAIR CLASSIFICATION OF SUPERHERO ALIGNMENT

In this use case we exploit a fictitious example that nonetheless represents a use case common to many real scenarios. We assume a model has been trained using protected data attributes and that it cannot be modified to correct for any bias. Instead, we build a copy that reproduces the learned decision function, while excluding these attributes.

We use superheroes dataset [78], which describes characteristics such as powers and physical attributes of 660 superheroes in SuperHeroDb [83]. We choose alignment as the target attribute to label all superheroes as either *good* or *bad*. We use these data to train a fully-connected artificial neural network with 4 hidden layers, each consisting of 128, 64, 32 and 16 neurons with *SeLu* activation, a softmax cross entropy loss optimized using *Adam* optimizer and a drop-out equal to 0.6. This model yields an accuracy of 0.65

Among the 177 input attributes, *gender* and *race* may be deemed sensitive. The differences in accuracy by the *gender* and *race* groups are shown in Table 2. In both cases, the resulting decision boundary leads to biased predictions. To overcome this issue, we propose to build a copy that does not include this information.

TABLE 2. Accuracy by *gender* and *race* groups for model and copy.

Attribute	Value	Model	Copy
<i>gender</i>	<i>female</i>	0.73	0.69
	<i>male</i>	0.64	0.66
<i>race</i>	<i>human</i>	0.78	0.76
	<i>mutant</i>	0.75	0.75
	<i>robot</i>	0.67	0.5
	<i>extraterrestrial</i>	0.25	0.5
	<i>other</i>	0.59	0.64

As a first step, we check that no other variable is correlated with *gender* and *race* and can leak this information into the copy. We train different models to predict *gender* or *race* using the rest of the variables. We average over 100 runs and obtain a mean balanced accuracy over classes of 0.42 ± 0.08 when predicting *gender* and of 0.28 ± 0.03 when predicting *race*. We also compute the one-to-one correlation for all attributes. At most, this correlation is equal to 0.18 in the case of *gender* and to 0.35 in the case of *race*. We conclude that the remaining attributes are very weakly correlated with these two, so that we can safely remove them without incurring in any leakage of information.

Hence, we extract these two attributes from the synthetic set and build a copy based on the existing network architecture. The mean copy accuracy is 0.66 ± 0.01 , the estimated copy accuracy is 0.61 ± 0.02 , and the empirical fidelity error is 0.059 ± 0.003 . The mean empirical fidelity error over the test data is 0.22 ± 0.01 . While this value may seem high, we stress that the removal of two variables results in a certain shift of the decision function. As shown in Table 2, this shift accommodates those instances that are unfairly classified by the model and reduces the overall bias in the copy.

C. LIMITATIONS

Despite its flexibility and large range of applications, copying has several limitations, for example, when it comes to dealing with high-dimensional data, or with certain problem environments. We highlight some of them.

TABLE 3. Experimental Results for the 60 Datasets in the UCI Experiment. Blank spaces correspond to cases where models learn a single class label.

Dataset	Classes	Samples	Features	Original	decision_tree			logistic_regression			random_forest		
					A_C	\widehat{A}_C	$R_{\mathcal{F}}$	A_C	\widehat{A}_C	$R_{\mathcal{F}}$	A_C	\widehat{A}_C	$R_{\mathcal{F}}$
abalone	3	3341	8	adaboost	0.57	0.5653 ± 0.002	0.27 ± 0.02	0.57 ± 0.00	0.3066 ± 0.0001	0.41 ± 0.00	0.58 ± 0.01	0.5590 ± 0.0109	0.27 ± 0.01
acute-inflammation	2	96	6	adaboost	1	1.0000 ± 0.0000	0.00 ± 0.00	1.00 ± 0.00	1.0000 ± 0.0000	0.00 ± 0.00	1.00 ± 0.00	1.0000 ± 0.0000	0.00 ± 0.00
acute-nephritis	2	96	6	adaboost	1	1.0000 ± 0.0000	0.00 ± 0.00	1.00 ± 0.00	0.9987 ± 0.0001	0.00 ± 0.00	1.00 ± 0.00	1.0000 ± 0.0000	0.00 ± 0.00
bank	2	3616	16	adaboost	0.85	0.82 ± 0.03	0.10 ± 0.03	0.87 ± 0.00	0.5558 ± 0.0001	0.12 ± 0.00	0.87 ± 0.00	0.8353 ± 0.0026	0.07 ± 0.00
blood	2	598	4	adaboost	0.71	0.65 ± 0.04	0.14 ± 0.05	0.67 ± 0.00	0.1590 ± 0.0001	0.38 ± 0.00	0.74 ± 0.00	0.7018 ± 0.0226	0.16 ± 0.04
breast-cancer	2	228	9	adaboost	0.74	0.74 ± 0.01	0.00 ± 0.00	0.74 ± 0.00	0.4731 ± 0.0001	0.31 ± 0.00	0.74 ± 0.00	0.7412 ± 0.0000	0.00 ± 0.00
breast-cancer-wisc	2	559	9	adaboost	0.93	0.93 ± 0.00	0.00 ± 0.00	0.93 ± 0.00	0.6333 ± 0.0001	0.06 ± 0.00	0.93 ± 0.00	0.9286 ± 0.0000	0.00 ± 0.00
breast-cancer-wisc-diag	2	455	30	adaboost	0.95	0.95 ± 0.00	0.00 ± 0.00	0.95 ± 0.00	0.4444 ± 0.0001	0.04 ± 0.00	0.95 ± 0.00	0.9473 ± 0.0000	0.00 ± 0.00
breast-cancer-wisc-prog	2	158	33	adaboost	0.73	0.72 ± 0.00	0.00 ± 0.00	0.72 ± 0.00	0.5468 ± 0.0001	0.35 ± 0.00	0.72 ± 0.00	0.7249 ± 0.0000	0.00 ± 0.00
breast-tissue	6	84	9	adaboost	0.59	0.61 ± 0.02	0.16 ± 0.04	0.36 ± 0.00	0.3838 ± 0.0001	0.59 ± 0.00	0.57 ± 0.02	0.5909 ± 0.0223	0.18 ± 0.01
chess-krvkp	2	2556	36	am	0.61	0.61 ± 0.00	0.11 ± 0.02	0.91 ± 0.00	0.9603 ± 0.0000	0.10 ± 0.00	0.91 ± 0.01	0.9670 ± 0.0091	0.09 ± 0.01
congressional-voting	2	348	16	am	0.61	0.61 ± 0.00	0.00 ± 0.00	0.61 ± 0.00	0.6087 ± 0.0000	0.00 ± 0.00	0.61 ± 0.00	0.6092 ± 0.0000	0.00 ± 0.00
conn-bench-sonar-mines-rocks	2	166	60	am	0.88	0.58 ± 0.08	0.5897 ± 0.0004	0.81 ± 0.01	0.8079 ± 0.0000	0.12 ± 0.01	0.69 ± 0.03	0.7405 ± 0.0327	0.21 ± 0.04
conn4	2	54045	42	am	0.87	0.54 ± 0.04	0.6518 ± 0.0003	0.46 ± 0.04	0.871 ± 0.0000	0.39 ± 0.00	0.66 ± 0.00	0.7442 ± 0.0031	0.32 ± 0.00
connect	3	1178	9	am	0.55	0.55 ± 0.01	0.07 ± 0.02	0.50 ± 0.00	0.4111 ± 0.0000	0.40 ± 0.00	0.55 ± 0.00	0.5357 ± 0.0043	0.03 ± 0.01
credit-approval	2	552	15	am	0.79	0.80 ± 0.02	0.7268 ± 0.0003	0.10 ± 0.02	0.7066 ± 0.0000	0.19 ± 0.00	0.79 ± 0.01	0.7527 ± 0.0071	0.03 ± 0.01
cylinder-bands	2	409	35	am	0.69	0.60 ± 0.05	0.6083 ± 0.0002	0.37 ± 0.05	0.6354 ± 0.0000	0.25 ± 0.00	0.65 ± 0.01	0.6445 ± 0.0121	0.35 ± 0.01
echocardiogram	2	104	10	am	0.3	0.33 ± 0.04	0.2879 ± 0.0001	0.05 ± 0.03	0.2960 ± 0.0000	0.00 ± 0.00	0.30 ± 0.00	0.2922 ± 0.0000	0.00 ± 0.00
energy-y1	3	614	8	am	0.96	0.96 ± 0.01	0.9537 ± 0.0001	0.00 ± 0.01	0.7744 ± 0.0001	0.23 ± 0.00	0.96 ± 0.00	0.9511 ± 0.0039	0.00 ± 0.00
energy-y2	3	614	8	am	0.84	0.84 ± 0.00	0.8403 ± 0.0001	0.00 ± 0.00	0.7835 ± 0.0001	0.09 ± 0.00	0.84 ± 0.01	0.8410 ± 0.0065	0.01 ± 0.00
fertility	2	80	9	am	0.9	0.90 ± 0.00	0.8993 ± 0.0000	0.00 ± 0.00	0.8255 ± 0.0001	0.15 ± 0.00	0.90 ± 0.00	0.8993 ± 0.0000	0.00 ± 0.00
haberman-survival	2	244	3	random_forest	0.61	0.61 ± 0.01	0.6125 ± 0.0000	0.01 ± 0.01	0.4888 ± 0.0001	0.26 ± 0.00	0.61 ± 0.01	0.6123 ± 0.0074	0.00 ± 0.01
heart-hungarian	2	235	12	random_forest	0.76	0.79 ± 0.04	0.7445 ± 0.0002	0.07 ± 0.03	0.6457 ± 0.0001	0.05 ± 0.00	0.80 ± 0.00	0.7498 ± 0.0000	0.03 ± 0.00
hepatitis	2	124	19	random_forest	0.74	0.74 ± 0.05	0.6991 ± 0.0002	0.05 ± 0.04	0.5973 ± 0.0001	0.16 ± 0.00	0.75 ± 0.01	0.7145 ± 0.0097	0.00 ± 0.01
ilpd-indian-liver	2	466	9	random_forest	0.62	0.59 ± 0.02	0.6207 ± 0.0001	0.36 ± 0.02	0.6020 ± 0.0000	0.47 ± 0.00	0.62 ± 0.01	0.6224 ± 0.0089	0.37 ± 0.01
ionosphere	2	280	33	random_forest	0.94	0.92 ± 0.02	0.8543 ± 0.0004	0.07 ± 0.03	0.7920 ± 0.0000	0.11 ± 0.00	0.95 ± 0.01	0.8867 ± 0.0138	0.05 ± 0.01
iris	3	120	4	random_forest	0.88	0.83 ± 0.01	0.9332 ± 0.0003	0.02 ± 0.02	0.6809 ± 0.0001	0.30 ± 0.00	0.95 ± 0.02	0.9333 ± 0.0163	0.02 ± 0.02
magic	2	15216	5	random_forest	0.8	0.80 ± 0.00	0.8024 ± 0.0003	0.10 ± 0.01	0.6809 ± 0.0001	0.17 ± 0.00	0.86 ± 0.00	0.8316 ± 0.0015	0.06 ± 0.00
mammographic	2	768	5	random_forest	0.88	0.88 ± 0.01	0.7974 ± 0.0001	0.01 ± 0.01	0.5377 ± 0.0000	0.19 ± 0.00	0.80 ± 0.00	0.7974 ± 0.0039	0.01 ± 0.00
mole-biol-splice	2	6499	50	random_forest	0.84	0.86 ± 0.01	0.8402 ± 0.0004	0.12 ± 0.01	0.6951 ± 0.0004	0.14 ± 0.00	0.80 ± 0.00	0.8679 ± 0.0040	0.15 ± 0.00
muskmushroom	3	2552	60	linear_svm	0.98	0.77 ± 0.01	0.7151 ± 0.0002	0.17 ± 0.01	0.7801 ± 0.0000	0.00 ± 0.00	0.87 ± 0.05	0.7590 ± 0.0040	0.11 ± 0.06
muskmushroom	3	6499	21	linear_svm	0.98	0.95 ± 0.02	0.9526 ± 0.0002	0.03 ± 0.02	0.8732 ± 0.0000	0.00 ± 0.00	0.95 ± 0.02	0.9549 ± 0.0544	0.32 ± 0.04
muskmushroom	3	380	166	linear_svm	0.88	0.54 ± 0.04	0.5620 ± 0.0003	0.46 ± 0.05	0.9556 ± 0.0000	0.01 ± 0.00	0.67 ± 0.06	0.7323 ± 0.0375	0.44 ± 0.04
muskmushroom	3	817	41	linear_svm	0.96	0.50 ± 0.05	0.6005 ± 0.0004	0.50 ± 0.06	0.8144 ± 0.0000	0.00 ± 0.00	0.59 ± 0.03	0.7543 ± 0.0296	0.32 ± 0.03
oocytes_merluccius_nucleus_4d	2	729	25	linear_svm	0.81	0.47 ± 0.06	0.6946 ± 0.0002	0.43 ± 0.07	0.8184 ± 0.0000	0.00 ± 0.00	0.63 ± 0.03	0.7097 ± 0.0050	0.02 ± 0.02
oocytes_trisopterus_nucleus_2f	2	156	22	linear_svm	0.9	0.56 ± 0.05	0.6946 ± 0.0002	0.11 ± 0.06	0.8950 ± 0.0000	0.00 ± 0.00	0.89 ± 0.01	0.8562 ± 0.0126	0.02 ± 0.01
pima	2	614	8	linear_svm	0.72	0.72 ± 0.01	0.6967 ± 0.0001	0.04 ± 0.01	0.7203 ± 0.0000	0.00 ± 0.00	0.71 ± 0.01	0.9091 ± 0.0050	0.02 ± 0.01
pittsburg-bridges-MATERIAL	3	84	7	linear_svm	0.91	0.91 ± 0.00	0.9090 ± 0.0000	0.00 ± 0.00	0.8968 ± 0.0000	0.00 ± 0.00	0.91 ± 0.00	0.9091 ± 0.0050	0.00 ± 0.00
pittsburg-bridges-REI-L	3	82	7	linear_svm	0.67	0.67 ± 0.00	0.6667 ± 0.0000	0.00 ± 0.00	0.6667 ± 0.0000	0.00 ± 0.00	0.67 ± 0.00	0.6667 ± 0.0000	0.00 ± 0.00
pittsburg-bridges-TOR-D	3	81	7	linear_svm	0.86	0.86 ± 0.00	0.8562 ± 0.0000	0.00 ± 0.00	0.6932 ± 0.0001	0.24 ± 0.00	0.86 ± 0.00	0.8568 ± 0.0000	0.00 ± 0.00
planning	2	145	12	rlf_svm	0.7	0.70 ± 0.00	0.7027 ± 0.0000	0.00 ± 0.00	0.7049 ± 0.0000	0.26 ± 0.00	0.70 ± 0.00	0.7027 ± 0.0000	0.00 ± 0.00
ringnorm	2	5920	18	rlf_svm	0.98	0.88 ± 0.01	0.7992 ± 0.0005	0.11 ± 0.01	0.74 ± 0.00	0.00 ± 0.00	0.95 ± 0.00	0.8877 ± 0.0022	0.05 ± 0.00
seeds	3	168	7	rlf_svm	0.88	0.80 ± 0.02	0.8006 ± 0.0003	0.06 ± 0.03	0.88 ± 0.00	0.00 ± 0.00	0.92 ± 0.02	0.8487 ± 0.0158	0.04 ± 0.02
seeds	3	168	7	rlf_svm	0.93	0.45 ± 0.10	0.7610 ± 0.0002	0.55 ± 0.11	0.92 ± 0.00	0.02 ± 0.00	0.53 ± 0.08	0.8577 ± 0.0824	0.48 ± 0.09
shuttle-australian-credit	2	3680	57	rlf_svm	0.68	0.68 ± 0.00	0.6812 ± 0.0000	0.00 ± 0.00	0.9227 ± 0.0000	0.00 ± 0.00	0.68 ± 0.00	0.6812 ± 0.0000	0.00 ± 0.00
shuttle-german-credit	2	552	14	rlf_svm	0.79	0.59 ± 0.03	0.6820 ± 0.0003	0.36 ± 0.04	0.7817 ± 0.0000	0.02 ± 0.00	0.81 ± 0.01	0.7369 ± 0.0121	0.07 ± 0.01
shuttle-german-credit	2	800	24	rlf_svm	0.85	0.81 ± 0.02	0.8051 ± 0.0002	0.05 ± 0.01	0.8510 ± 0.0000	0.00 ± 0.00	0.86 ± 0.01	0.8271 ± 0.0056	0.04 ± 0.01
shuttle-image	7	216	13	rlf_svm	0.95	0.74 ± 0.01	0.8351 ± 0.0003	0.25 ± 0.02	0.8218 ± 0.0001	0.23 ± 0.00	0.78 ± 0.01	0.8993 ± 0.0076	0.21 ± 0.01
shuttle-vehicle	4	1848	18	rlf_svm	0.85	0.56 ± 0.05	0.5489 ± 0.0002	0.40 ± 0.04	0.66 ± 0.00	0.66 ± 0.00	0.65 ± 0.02	0.6971 ± 0.0186	0.29 ± 0.02
shuttle-vehicle	4	676	18	rlf_svm	0.85	0.56 ± 0.05	0.5489 ± 0.0002	0.23 ± 0.03	0.81 ± 0.00	0.3047 ± 0.0001	0.94 ± 0.01	0.7683 ± 0.0099	0.02 ± 0.01
synthetic-control	6	480	60	sgboost	0.96	0.55 ± 0.02	0.6547 ± 0.0007	0.01 ± 0.01	0.35 ± 0.00	0.3047 ± 0.0001	0.55 ± 0.01	0.5476 ± 0.0099	0.02 ± 0.01
teaching	2	120	5	sgboost	0.55	0.55 ± 0.02	0.5476 ± 0.0000	0.00 ± 0.00	0.8224 ± 0.0002	0.16 ± 0.00	0.97 ± 0.00	0.9740 ± 0.0000	0.00 ± 0.00
tic-tac-toe	2	766	9	sgboost	0.97	0.78 ± 0.00	0.9740 ± 0.0000	0.00 ± 0.00	0.5837 ± 0.0000	0.20 ± 0.00	0.97 ± 0.00	0.7778 ± 0.0000	0.00 ± 0.00
tic-tac-toe	2	1760	3	sgboost	0.78	0.78 ± 0.00	0.7778 ± 0.0000	0.00 ± 0.00	0.8913 ± 0.0000	0.14 ± 0.01	0.88 ± 0.00	0.8842 ± 0.0021	0.02 ± 0.01
twonorm	2	5920	20	sgboost	0.98	0.87 ± 0.01	0.7454 ± 0.0003	0.13 ± 0.01	0.98 ± 0.00	0.00 ± 0.00	0.96 ± 0.00	0.9647 ± 0.0000	0.00 ± 0.00
vertebral-column-2classes	2	248	6	sgboost	0.77	0.78 ± 0.02	0.7669 ± 0.0001	0.05 ± 0.02	0.81 ± 0.01	0.7244 ± 0.0001	0.84 ± 0.01	0.7687 ± 0.0000	0.02 ± 0.01
vertebral-column-3classes	3	248	6	sgboost	0.84	0.84 ± 0.02	0.8379 ± 0.0000	0.02 ± 0.01	0.80 ± 0.00	0.7053 ± 0.0000	0.84 ± 0.00	0.8378 ± 0.0000	0.00 ± 0.00
waveform	2	4000	21	sgboost	0.84	0.77 ± 0.01	0.6250 ± 0.0005	0.18 ± 0.01	0.85 ± 0.00	0.7575 ± 0.0000	0.83 ± 0.00	0.7318 ± 0.0044	0.08 ± 0.00
waveform-noise	3	4000	40	sgboost	0.84	0.76 ± 0.01	0.6610 ± 0.0004	0.19 ± 0.01	0.87 ± 0.00	0.7083 ± 0.0000	0.85 ± 0.00	0.7415 ± 0.0043	0.07 ± 0.00
wine	3	142	11	sgboost	0.92	0.92 ± 0.00	0.9147 ± 0.0000	0.00 ± 0.00	0.7031 ± 0.0001	0.08 ± 0.00	0.92 ± 0.00	0.9147 ± 0.0000	0.00 ± 0.00

Copying is highly dependent on the synthetic data generation process. The complexity of this process grows with increasing dimensionality. Hence, while the copying methodology itself remains valid in this context, its performance may be affected. Mostly because sampling an unknown decision function is hard. More so, because we have no information about the training data distribution and lack any insight on how the different classes may be distributed throughout the space. In theory, we could overcome this problem by generating infinite query points. Yet, this is not tractable in practice, since we are limited by our computational resources.

In our experience, when considering large dimensionality data it is worth replacing uniform sampling distributions with normal distributions. The first conduct an arbitrary exploration of the space, whereas the second better characterize the typicality⁷ of a standardized dataset. This is because, as the number of dimensions increases, so do the regions of the space where there are no data present. By using a normal distribution to guide sampling we focus only on those areas that could potentially contain data.

Not only the amount of data but also their structure can be problematic. In structured environments, such as those of images or text, data tend to lie on top of a variety. Finding the optimal synthetic dataset therefore requires sampling the appropriate manifold. While this may be doable, it is not straightforward. In general, copying in such domains would require access to the training data to generate synthetic data with a suitable representation. This could be done, for example, using an autoencoder that ensures image invariance.

An additional limitation is choosing P_Z . As shown above, blindly exploring the input space works well for simple cases. As the complexity of the problem grows, however, so does the intricacy of the decision function and more *ad hoc* techniques are needed to appropriately sample the input space. See for example [54], where we assay uncertainty based methods to guide sampling,

Lastly, many local minima exist. This is because an infinite number of different synthetic sets can be used to replicate a given decision boundary. In theory, the empirical error is known and equal to zero, so that all sets should converge to the same result. Due to training variability, however, this is not always the case.

VII. CONCLUSION AND FUTURE WORK

In this paper we propose and validate a model-agnostic framework to copy machine learning classifiers. Copying refers to the process of creating an exact replica of a classifier's decision boundary (or the most similar one if this can not be achieved). As such, this process can be understood as a projection operator of a decision function onto a target model space. The resulting copy optimizes the fidelity measure to preserve the original predictive performance.

⁷The concept of typicality refers to properties holding for the vast majority of cases [84].

We derive the theory for copying and highlight its differences with learning, as traditionally understood by the machine learning community. The process of building a copy does not require access to training data. Moreover, we consider the most general case, where the original model is treated as a black-box whose internals remain unknown.

We exploit the concept of differential replication as the property of endowing copies with new features by adequately selecting the target projection space. This enables copies to provide reliable solutions to many open issues in machine learning. We also discuss the implications of building copies in practice and introduce a set of performance metrics assuming access to different levels of information. Our experiments demonstrate that our approach is feasible. Moreover, the case studies presented show the potential of copies to ensure interpretability, fairness or productivization of machine learning models.

The problem of representing the decision behaviour of a machine learning model using a finite number of samples is far from being solved. Notably, an in-depth study should be conducted to evaluate methods to sample closed domains where class distribution is governed by an unknown decision function. Much research also remains to be done on how to solve the dual optimization problem. While the single pass-copy provides a reasonable approximation, more general approaches should be studied.

In this article we restrict ourselves to exploring the application of copies to specific areas such as interpretability, fairness and general enhancement. Nonetheless, there exist other fields where copies are potentially useful. Particularly that of privacy, where copies could be specifically built to be privacy-preserving with respect to the training data. This wide range of applications is ensured by the differential replication property of copies, which enables adaptation to new needs and requirements. This characteristic should be the subject of further research.

APPENDIX RESULTS FOR UCI CLASSIFICATION

See Table 3.

REFERENCES

- [1] M. Fredrikson, S. Jha, and T. Ristenpart, "Model inversion attacks that exploit confidence information and basic countermeasures," in *Proc. Conf. Comput. Commun. Secur.*, 2015, pp. 1322–1333.
- [2] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2017, pp. 3–18.
- [3] C. Song, T. Ristenpart, and V. Shmatikov, "Machine learning models that remember too much," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, Oct. 2017, pp. 587–601.
- [4] *Legislation*, OJ, EU Commission, Brussels, Belgium, 2016.
- [5] B. Goodman and S. Flaxman, "European union regulations on algorithmic decision-making and a right to explanation," *AI Mag.*, vol. 38, no. 3, pp. 50–57, Oct. 2017.
- [6] A. D. Selbst and J. Powles, "Meaningful information and the right to explanation," *Int. Data Privacy Law*, vol. 7, no. 4, pp. 233–242, Nov. 2017.
- [7] S. Barocas and A. D. Selbst, "Big data's disparate impact," *California Law Rev.*, vol. 104, no. 671, pp. 671–732, 2016.

- [8] B. Friedman and H. Nissenbaum, "Bias in computer systems," *ACM Trans. Inf. Syst.*, vol. 14, no. 3, pp. 330–347, Jul. 1996.
- [9] M. Hardt. (2014). *How Big Data is Unfair*. Medium. [Online]. Available: <https://medium.com/mrtz/how-big-data-is-unfair-9aa544d739de>
- [10] D. Sculley, G. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, M. Young, J.-F. Crespo, and D. Nison, "Hidden technical debt in machine systems," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2015, pp. 2503–2511.
- [11] M. W. Craven and J. W. Shavlik, "Extracting tree-structured representations of trained networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 1995, pp. 24–30.
- [12] F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, "Stealing machine learning models via prediction APIs," in *Proc. USENIX Secur. Symp.*, 2016, pp. 601–618.
- [13] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," in *Proc. Workshop Deep Learn. Represent. Learn.*, 2015, pp. 1–9.
- [14] C. Bucila, R. Caruana, and A. Niculescu-Mizil, "Model compression," in *Proc. ACM Int. Conf. Knowl. Discovery Data Mining*, 2006, pp. 535–541.
- [15] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, "Practical black-box attacks against machine learning," in *Proc. ACM Asia Conf. Comput. Commun. Secur.*, Apr. 2017, pp. 506–519.
- [16] D. Lowd and C. Meek, "Adversarial learning," in *Proc. ACM Int. Conf. Knowl. Discovery Data Mining*, 2005, pp. 641–647.
- [17] R. Liu, N. Fusi, and L. Mackey, "Teacher-student compression with generative adversarial networks," 2018, *arXiv:1812.02271*. [Online]. Available: <http://arxiv.org/abs/1812.02271>
- [18] N. Frosst and G. Hinton, "Distilling a neural network into a soft decision tree," 2017, *arXiv:1711.09784*. [Online]. Available: <http://arxiv.org/abs/1711.09784>
- [19] I. Unceta, J. Nin, and O. Pujol, "Differential replication in machine learning," 2020, *arXiv:2007.07981*. [Online]. Available: <http://arxiv.org/abs/2007.07981>
- [20] R. Andrews, J. Diederich, and A. B. Tickle, "Survey and critique of techniques for extracting rules from trained artificial neural networks," *Knowl.-Based Syst.*, vol. 8, no. 6, pp. 373–389, Dec. 1995.
- [21] M. W. Craven and J. W. Shavlik, "Learning symbolic rules using artificial neural networks," in *Proc. Int. Conf. Mach. Learn.*, 1993, pp. 73–80.
- [22] L. Fu, "Rule learning by searching on adapted nets," in *Proc. Nat. Conf. Artif. Intell.*, 1991, pp. 590–595.
- [23] S. Thrun, "Extracting rules from artificial neural networks with distributed representations," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, vol. 7, 1995, pp. 505–512.
- [24] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2818–2826.
- [25] T. Furlanello, Z. Lipton, M. Tschannen, and A. Anandkumar, "Born again neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2018, pp. 1607–1616.
- [26] G. Bologna and Y. Hayashi, "A comparison study on rule extraction from neural network ensembles, boosted shallow trees, and SVMs," *Appl. Comput. Intell. Soft Comput.*, vol. 2018, pp. 1–20, Jan. 2018.
- [27] Z. Che, S. Purushotham, R. Khemani, and Y. Liu, "Interpretable deep models for ICU outcome prediction," in *Proc. AMIA Annu. Symp.*, 2016, pp. 371–380.
- [28] O. Bastani, C. Kim, and H. Bastani, "Interpreting blackbox models via model extraction," 2018, *arXiv:1705.08504*. [Online]. Available: <https://arxiv.org/abs/1705.08504>
- [29] S. Hong, C. Xiao, T. N. Hoang, T. Ma, H. Li, and J. Sun, "RDPD: Rich data helps poor data via imitation," in *Proc. 28th Int. Joint Conf. Artif. Intell.*, Aug. 2019, pp. 5895–5901.
- [30] X. Zeng and T. R. Martinez, "Using a neural network to approximate an ensemble of classifiers," *Neural Process. Lett.*, vol. 12, pp. 225–237, Dec. 2000.
- [31] X. Zhu, "Machine teaching: An inverse problem to machine learning and an approach toward optimal education," in *Proc. 29th AAAI Conf. Artif. Intell.*, Palo Alto, CA, USA, 2015, pp. 4083–4087.
- [32] B. Settles, "Active learning literature survey," Dept. Comput. Sci., Univ. Wisconsin-Madison, Madison, WI, USA, Tech. Rep. 1648, 2009.
- [33] A. Fujii, T. Tokunaga, K. Inui, and H. Tanaka, "Selective sampling for example-based word sense disambiguation," *Comput. Linguist.*, vol. 24, no. 4, pp. 573–597, 1998.
- [34] D. Lewis and W. Gale, "A sequential algorithm for training text classifiers," in *Proc. ACM Conf. Res. Dev. Inf. Retr.*, 1994, pp. 3–12.
- [35] M. Lindenbaum, S. Markovitch, and D. Rusakov, "Selective sampling for nearest neighbor classifiers," *Mach. Learn.*, vol. 54, no. 2, pp. 125–152, Feb. 2004.
- [36] T. Scheffer, C. Decomain, and S. Wrobel, "Active hidden Markov models for information extraction," in *Proc. Int. Conf. Adv. Intell. Data Anal.*, 2001, pp. 309–318.
- [37] Y. Zhang, T. Xiang, T. M. Hospedales, and H. Lu, "Deep mutual learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 4320–4328.
- [38] A. Polino, R. Pascanu, and D. Alistarh, "Model compression via distillation and quantization," in *Proc. 6th Int. Conf. Learn. Represent. (ICLR)*, 2018, pp. 1–21.
- [39] Y. Li, J. Yang, Y. Song, L. Cao, J. Luo, and L.-J. Li, "Learning from noisy labels with distillation," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 1928–1936.
- [40] A. Romero, N. Ballas, S. Kahou, A. Chassang, C. Gatta, and Y. Bengio, "FitNets: Hints for thin deep nets," in *Proc. 3th Int. Conf. Learn. Represent. (ICLR)*, 2015, pp. 1–13.
- [41] H. Lee, S. J. Hwang, and J. Shin, "Rethinking data augmentation: Self-supervised label augmentation via input transformations," 2019, *arXiv:1910.05872*. [Online]. Available: <http://arxiv.org/abs/1910.05872>
- [42] Z. B. Celik, D. Lopez-Paz, and P. McDaniel, "Patient-driven privacy control through generalized distillation," in *Proc. IEEE Symp. Privacy-Aware Comput. (PAC)*, Aug. 2017, pp. 1–12.
- [43] J. Wang, W. Bao, L. Sun, X. Zhu, B. Cao, and P. S. Yu, "Private model compression via knowledge distillation," in *Proc. 33rd AAAI Conf. Artif. Intell.*, 2019, pp. 1190–1198.
- [44] N. Papernot, P. McDaniel, and I. Goodfellow, "Transferability in machine learning: From phenomena to black-box attacks using adversarial samples," 2016, *arXiv:1605.07277*. [Online]. Available: <http://arxiv.org/abs/1605.07277>
- [45] Y. Liu, X. Chen, C. Liu, and D. Song, "Delving into transferable adversarial examples and black-box attacks," in *Proc. Int. Conf. Represent.*, 2017, pp. 1–24.
- [46] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," in *Proc. Int. Conf. Learn. Represent.*, 2014, pp. 1–10.
- [47] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrnđić, P. Laskov, G. Giacinto, and F. Roli, "Evasion attacks against machine learning at test time," in *Proc. Eur. Conf. Mach. Learn. Knowl. Discovery Databases*, 2013, pp. 387–402.
- [48] D. Lopez-Paz, L. Bottou, B. Scholkopf, and V. Vapnik, "Unifying distillation and privileged information," in *Proc. Int. Conf. Learn. Represent.*, 2016, pp. 1–10.
- [49] M. Phuong and C. Lampert, "Towards understanding knowledge distillation," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 5142–5151.
- [50] V. N. Vapnik, *The Nature of Statistical Learning Theory*. Berlin, Germany: Springer, 2000.
- [51] R. E. Schapire, Y. Freund, P. Bartlett, and W. S. Lee, "Boosting the margin: A new explanation for the effectiveness of voting methods," *Ann. Statist.*, vol. 26, no. 5, pp. 1651–1686, Oct. 1998.
- [52] B. Neyshabur, R. Tomioka, R. Salakhutdinov, and N. Srebro, "Geometry of optimization and implicit regularization in deep learning," 2017, *arXiv:1705.03071*. [Online]. Available: <http://arxiv.org/abs/1705.03071>
- [53] A. Brutzkus, A. Globerson, E. Malach, and S. Shalev-Shwartz, "SGD learns over-parameterized networks that provably generalize on linearly separable data," in *Proc. Int. Conf. Learn. Represent.*, 2017, pp. 1–17.
- [54] I. Unceta, D. Palacios, J. Nin, and O. Pujol, "Sampling unknown decision functions to build classifier copies," 2019, *arXiv:1910.00237*. [Online]. Available: <http://arxiv.org/abs/1910.00237>
- [55] D. Dheeru and E. K. Taniskidou, "UCI machine learning repository," School Inf. Comput. Sci., Univ. California Irvine, Irvine, CA, USA, Tech. Rep., 2017.
- [56] M. Fernández-Delgado, E. Cernadas, S. Barro, D. Amorim, and D. A. Fernández-Delgado, "Do we need hundreds of classifiers to solve real world classification problems?" *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 3133–3181, 2014.
- [57] L. Bottou and Y. Le Cun, "Large scale online learning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2004, pp. 217–234.
- [58] S. Escalera, D. Masip, E. Puertas, P. Radeva, and O. Pujol, "Online error correcting output codes," *Pattern Recognit. Lett.*, vol. 32, no. 3, pp. 458–467, 2009.

- [59] F. Doshi-Velez and B. Kim, "Towards a rigorous science of interpretable machine learning," 2017, *arXiv:1702.08608*. [Online]. Available: <http://arxiv.org/abs/1702.08608>
- [60] Z. C. Lipton, "The myths of model interpretability," in *Proc. Workshop Hum. Interpret. Mach. Learn.*, 2016, pp. 96–100.
- [61] N. Probst and G. Hinton, "Distilling a neural network into a soft decision tree," 2017, *arXiv:1711.09784*. [Online]. Available: <http://arxiv.org/abs/1711.09784>
- [62] M. T. Ribeiro, S. Singh, and C. Guestrin, "Anchors: High-precision model-agnostic explanations," in *Proc. Assoc. Adv. Artif. Intell.*, 2018, pp. 1527–1535.
- [63] H. Lakkaraju, S. H. Bach, and J. Leskovec, "Interpretable decision sets," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2016, pp. 1675–1684.
- [64] I. Flaounas, "Beyond the technical challenges for deploying machine learning solutions in a software company," in *Proc. Workshop Hum. Loop Mach. Learn.*, 2017, pp. 1–4.
- [65] A. Spector, P. Norvig, and S. Petrov, "Google's hybrid approach to research," *Commun. ACM*, vol. 55, no. 7, pp. 34–37, 2012.
- [66] A. Zheng and S. Raman, "The challenges of bringing machine learning to the masses," in *Proc. Workshop Softw. Eng. Mach. Learn.*, 2014.
- [67] *Deploy TensorFlow*. Accessed: Jan. 1, 2019. [Online]. Available: <https://www.tensorflow.org/deploy/>
- [68] W. C. Yau, *How Zendesk Serves TensorFlow Models in Production*. Medium, 2017.
- [69] J. Angwin, J. Larson, S. Mattu, and L. Kirchner, "Machine bias: There's software used across the country to predict future criminals and it's biased against blacks," *ProPublica*, 2016.
- [70] J. Buolamwini and T. Gebru, "Gender shades: Intersectional accuracy disparities in commercial gender classification," in *Proc. Mach. Learn. Res.*, vol. 81, 2018, pp. 1–15.
- [71] B. F. Klare, M. J. Burge, J. C. Klontz, R. W. V. Bruegge, and A. K. Jain, "Face recognition performance: Role of demographic information," *IEEE Trans. Inf. Forensics Security*, vol. 7, no. 6, pp. 1789–1801, Dec. 2012.
- [72] A. B. Popejoy and S. M. Fullerton, "Genomics is failing on diversity," *Nature*, vol. 538, no. 7624, pp. 161–164, Oct. 2016.
- [73] T. Bolukbasi, K. W. Chang, J. Zou, V. Saligrama, and A. Kalai, "Man is to computer programmer as woman is to homemaker? Debiasing word embeddings," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2016, pp. 4356–4364.
- [74] A. Caliskan, J. J. Bryson, and A. Narayanan, "Semantics derived automatically from language corpora contain human-like biases," *Science*, vol. 356, no. 6334, pp. 183–186, 2017.
- [75] S. Guha, B. Cheng, and P. Francis, "Challenges in measuring online advertising systems," in *Proc. 10th Annu. Conf. Internet Meas. (IMC)*, 2010, pp. 81–87.
- [76] S. Tan, R. Caruana, G. Hooker, and Y. Lou, "Distill-and-compare: Auditing black-box models using transparent model distillation," 2017, *arXiv:1710.06169*. [Online]. Available: <http://arxiv.org/abs/1710.06169>
- [77] *Lending Club Loan Data*. Accessed: Sep. 15, 2018. [Online]. Available: <https://www.kaggle.com/wendykan/lending-club-loan-data>
- [78] *Super Heroes Dataset*. Accessed: Jun. 15, 2018. [Online]. Available: <https://www.kaggle.com/clauidodavi/superhero-set>
- [79] I. Unceta, J. Nin, and O. Pujol, "Towards global explanations for credit risk scoring," 2018, *arXiv:1811.07698*. [Online]. Available: <http://arxiv.org/abs/1811.07698>
- [80] A. Namvar, M. Siami, F. Rabhi, and M. Naderpour, "Credit risk prediction in an imbalanced social lending environment," *Int. J. Comput. Intell. Syst.*, vol. 11, no. 1, p. 925, 2018.
- [81] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4700–4708.
- [82] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, "Self-normalizing neural networks," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 971–980.
- [83] *Superhero Database*. Accessed: Jun. 15, 2018. [Online]. Available: <https://www.superherodb.com/>
- [84] S. B. Volchan, "Probability as typicality," 2007, *arXiv:physics/0611172*. [Online]. Available: <https://arxiv.org/abs/physics/0611172>



IRENE UNCETA received the degree in physics from the University of Barcelona and the M.Sc. degree in computational science from the University of Amsterdam. She is currently pursuing the Ph.D. degree in industrial with the University of Barcelona and BBVA Data & Analytics. Equal parts a Social Scientist and a Data Scientist, her professional career has been a journey through the intersection of both worlds, with a focus mainly on innovation and transparency. She also works on the interpretability of machine learning models applied to financial risk with the University of Barcelona and BBVA Data & Analytics.



JORDI NIN received the degree in computer science from the Universitat Autònoma de Barcelona (UAB) and the Ph.D. degree (Hons.) in computer science from the Artificial Intelligence Research Institute, Spanish National Research Council (IIIA-CSIC), in 2008, on a work on how to apply machine learning models to improve several statistical disclosure control methods. He is currently an Assistant Professor with ESADE, Universitat Ramon Llull. In recent years, his research has focused on the application of machine learning models to financial risk evaluation from different perspectives.



ORIO PUIJOL received the Ph.D. degree from the Universitat Autònoma de Barcelona, in 2004, on a work on deformable models applied to medical imaging and fusion of supervised and unsupervised learning. He is currently an Associate Professor with the Department of Mathematics and Computer Science, Universitat de Barcelona. In recent years, his research has focused on ensemble learning methods, kernel machines and online methods, and applications to problems in finance, computer vision, signal processing, and natural language processing.

• • •