

Interpretable Predictions of Tree-based Ensembles via Actionable Feature Tweaking

Gabriele Tolomei

Yahoo Research
London, UK
gtolomei@yahoo-inc.com

Andrew Haines

Yahoo Research
London, UK
haines@yahoo-inc.com

Fabrizio Silvestri*

Facebook
London, UK
fsilvestr@fb.com

Mounia Lalmas

Yahoo Research
London, UK
mounia@acm.org

ABSTRACT

Machine-learned models are often described as “black boxes”. In many real-world applications however, models may have to sacrifice predictive power in favour of human-interpretability. When this is the case, feature engineering becomes a crucial task, which requires significant and time-consuming human effort. Whilst some features are inherently static, representing properties that cannot be influenced (e.g., the age of an individual), others capture characteristics that could be *adjusted* (e.g., the daily amount of carbohydrates taken). Nonetheless, once a model is learned from the data, each prediction it makes on new instances is irreversible - assuming every instance to be a static point located in the chosen feature space. There are many circumstances however where it is important to understand (i) why a model outputs a certain prediction on a given instance, (ii) which adjustable features of that instance should be modified, and finally (iii) how to alter such a prediction when the mutated instance is input back to the model.

In this paper, we present a technique that exploits the internals of a tree-based ensemble classifier to offer *recommendations* for transforming true negative instances into positively predicted ones. We demonstrate the validity of our approach using an online advertising application. First, we design a Random Forest classifier that effectively separates between two types of ads: *low* (negative) and *high* (positive) quality ads (instances). Then, we introduce an algorithm that provides recommendations that aim to transform a low quality ad (negative instance) into a high quality one (positive instance). Finally, we evaluate our approach on a subset of the active inventory of a large ad network, *Yahoo Gemini*.

*The author has contributed to this work while he was employed at Yahoo Research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '17, August 13-17, 2017, Halifax, NS, Canada

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-4887-4/17/08...\$15.00

<https://doi.org/10.1145/3097983.3098039>

CCS CONCEPTS

• **Theory of computation** → **Machine learning theory**; • **Computing methodologies** → **Ensemble methods**; • **Information systems** → **Online advertising**;

KEYWORDS

Model interpretability; Actionable feature tweaking; Recommending feature changes; Altering model predictions; Random forest

ACM Reference format:

Gabriele Tolomei, Fabrizio Silvestri, Andrew Haines, and Mounia Lalmas. 2017. Interpretable Predictions of Tree-based Ensembles via Actionable Feature Tweaking. In *Proceedings of KDD '17, Halifax, NS, Canada, August 13-17, 2017*, 10 pages.
<https://doi.org/10.1145/3097983.3098039>

1 INTRODUCTION

An increasing number of organisations and governments rely on *Machine Learning* (ML) techniques to extract knowledge from the large volumes of data they collect every day to optimise their operational effectiveness.

ML solutions are usually considered as “black boxes”; they take some inputs and produce desired outputs, especially when their ultimate goal is *prediction* rather than *inference*. As long as ML models work properly, “everybody” is happy and little attention is devoted to understand why such surprisingly good results are obtained. Still, it is beneficial to have available techniques supporting humans in interpreting and “debugging” these models, particularly when they fail [23] or lead to some oddities.¹

Excluding recent trends in ML such as Deep Learning [13], typically the initial effort when designing an ML solution consists in modelling the objects of a given domain of interest, *i.e.*, *feature engineering*. This step aims to describe each object in the domain using an appropriate set of properties (*features*), which define a so-called *feature space*. For a given dataset, each object can be considered as a static point located in the feature space since each feature value is deemed to be fixed; once a model is learned from the data, each prediction it makes on new objects is irreversible.

Let us assume that we disagree with a prediction that the model returns for a given object or that we would like to enforce switching

¹ <http://www.telegraph.co.uk/technology/2016/03/24/microsofts-teen-girl-ai-turns-into-a-hitler-loving-sex-robot-wit/>

such a prediction. The research question we ask in this work is *how can we understand what can be changed in the feature vector in order to modify the prediction accordingly?*

To better understand this challenge with an example, consider an ML application in the healthcare domain, where patients (objects) are mapped to a vector of clinical indicators (features), such as age, blood pressure, daily carbohydrates taken, *etc.* Assume next that an ML model has been designed to accurately predict from these features whether a patient is at risk of a heart attack or not. If for a given patient our model predicts that there is a high risk of a heart attack it would be of great advantage for medical physicians to also have a tool that suggests the most appropriate clinical treatment by offering targeted adjustments to specific indicators (*e.g.*, reducing the daily amount of carbohydrates). In other words, to recommend the clinical treatment to switch a patient from being of high risk (negative instance) to low risk (positive instance).

In this work, we propose an algorithm for *tweaking* input features so as to change the output predicted by an existing machine-learned model. Our method is designed to operate on top of any tree-based ensemble binary classifier, although it can be extended to multi-class classification. Our proposed algorithm exploits the internals of the model to generate recommendations for transforming true negative instances into positively predicted ones (or vice versa).

We describe the theoretical framework along with experiments designed to validate the proposed algorithm. Our approach is then evaluated in the commercial and more implementable setting of online advertisement recommendations to illustrate the generic nature of our framework and the many and varied domains it can be applied to.

After presenting an effective Random Forest classifier that is able to separate between *low* and *high* quality advertisements [12], we show how our algorithm can be used to automatically generate “interpretable” and “actionable” suggestions on how to convert a low quality ad (negative instance) into a high quality one (positive instance). Such insights can be provided to advertisers who may turn them into actual changes to their ad campaigns with the aim of improving their return on investment. Finally, we assess the quality of recommendations that our algorithm generates out of a dataset of advertisements served by the *Yahoo Gemini* ad network.

2 PROBLEM STATEMENT

We start by considering the typical binary classification problem and focus on an ensemble of tree-based classifiers as an effective solution to the above problem. Additionally, we define how the internals of an existing ensemble of trees can be used to derive a feedback loop for recommending how true negative instances can be turned into positively predicted ones (or vice versa).

The approach we propose can be easily extended to the more general multi-class classification problem. We plan to present this result in future extended work.

2.1 Notation

Let $\mathcal{X} \subseteq \mathbb{R}^n$ be an n -dimensional vector space of real-valued features. Any $\mathbf{x} \in \mathcal{X}$ is an n -dimensional feature vector, *i.e.*, $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$, representing an object in the vector space \mathcal{X} . Suppose that each \mathbf{x} is associated with a binary *class label* – either

neg (*negative*) or pos (*positive*) – and let $\mathcal{Y} = \{-1, +1\}$ be the set encoding all such possible class labels.

We assume there exists an *unknown target* function $f : \mathcal{X} \mapsto \mathcal{Y}$ that maps any feature vector to its corresponding class label. In addition, we let $\hat{f} \approx f$ which is learned from a labelled dataset of m instances $\mathcal{D} = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_m, y_m)\}$. More specifically, \hat{f} is the estimate that best approximates f on \mathcal{D} , according to a specific *loss function* ℓ . Such a function measures the “cost” of prediction errors we would make if we replaced the true target f with the estimate \hat{f} .

The flexibility vs. interpretability of \hat{f} depends on the *hypothesis space* which \hat{f} has been picked from by the learning algorithm. In this work, we focus on \hat{f} represented as an *ensemble* of K tree-based classifiers, $\hat{f} = \phi(\hat{h}_1, \dots, \hat{h}_K)$. Each $\hat{h}_k : \mathcal{X} \mapsto \mathcal{Y}$ is a base estimate, and ϕ is the function responsible for combining the output of all the individual base classifiers into a single prediction.

A possible implementation of ϕ could use a *majority voting* strategy. In this setting, a given instance \mathbf{x} would obtain a predicted class label $\hat{f}(\mathbf{x})$ based on the result of the majority of the base classifiers; this is the *mode* of the base predictions. Although other strategies may be used, this does not impact our proposed approach.

2.2 Enforcing Positive Prediction

In any ensemble of tree-based classifiers, each base estimate \hat{h}_k is encoded by a decision tree T_k , and the ensemble is represented as a forest $\mathcal{T} = \{T_1, \dots, T_K\}$.

Our aim is to identify how to transform a true negative instance into a positively predicted one. Let $\mathbf{x} \in \mathcal{X}$ be a true negative instance such that $f(\mathbf{x}) = \hat{f}(\mathbf{x}) = -1$. The task can now be defined as transforming the original input feature vector \mathbf{x} into a new feature vector $\mathbf{x}' (\mathbf{x} \rightsquigarrow \mathbf{x}')$ such that $\hat{f}(\mathbf{x}') = +1$. Moreover, we accomplish an optimised form of the problem by choosing \mathbf{x}' as the best transformation among all the possible transformations \mathbf{x}^* , according to a *cost function* $\delta : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$. This is defined as follows:

$$\mathbf{x}' = \arg \min_{\mathbf{x}^*} \left\{ \delta(\mathbf{x}, \mathbf{x}^*) \mid \hat{f}(\mathbf{x}) = -1 \wedge \hat{f}(\mathbf{x}^*) = +1 \right\}$$

The cost function measures the “effort” of transforming \mathbf{x} into \mathbf{x}' . A possible choice of such a function is the number of features affected by the transformation or the Euclidean distance between the original and the transformed vector.

2.3 Positive and Negative Paths

Any root-to-leaf path of a single decision tree can be interpreted as a cascade of *if-then-else* statements, where every internal (non-leaf) node is a boolean test on a specific feature value against a threshold. We restrict the tree decisions to be binary representations as any multiway decision can be represented in a binary form and there is little performance benefit in n -ary splits. An instance’s feature value is then evaluated at each node to determine which branch to traverse. This is repeated until the leaves are reached whereby the pos/neg classification labels are defined and assigned.

Given a forest of K decision trees $\mathcal{T} = \{T_1, \dots, T_K\}$, we denote by $p_{k,j}$ the j -th path of the k -th tree T_k . We refer to $p_{k,j}^+$ (or $p_{k,j}^-$) as the j -th path of T_k that leads to a leaf node labelled as pos (or neg) – a *positive* (or *negative*) path. For simplicity, we assume that each

path of a decision tree contains at most n non-leaf nodes, which correspond to n boolean conditions, one for each distinct feature.² We thus represent a root-to-leaf path as follows:

$$p_{k,j} = \{(x_1 \lesseqgtr \theta_1), (x_2 \lesseqgtr \theta_2), \dots, (x_n \lesseqgtr \theta_n)\} \quad (1)$$

Let $P_k^+ = \bigcup_{j \in T_k} p_{k,j}^+$ describe the set of all positive paths, and $P_k^- = \bigcup_{j \in T_k} p_{k,j}^-$ the set of all the negative paths in T_k . Also, let $P_k = P_k^+ \cup P_k^-$ be the set of *all* the paths in T_k .

We thus enumerate the possible paths in a single decision tree. Even under the assumption that each $p_{k,j} \in P_k$ is at most a length- n path then T_k is a depth- n binary tree, whose number of leaves is therefore bounded to 2^n . As the total number of leaves coincides with the total number of possible paths, we obtain $|P_k| \leq 2^n$.

In general, we cannot ensure a bound on each $|P_k|$, as there might exist some paths $p_{k,j} \in P_k$ whose length is greater than n . In practice though, we can specify the maximum number of paths at training time by bounding the depth of the generated trees to the number n of features. Even with such a relaxed condition, the total number of possible paths encoded by the forest \mathcal{T} is equal to $\sum_{k=1}^K |P_k| \leq K2^n$, therefore still exponential in n . We see later how this does not disrupt computational efficacy in practice, as our algorithm operates on a subset of those paths.

2.4 Tweaking Input Features

Given our input feature vector \mathbf{x} , we know from our hypothesis that $f(\mathbf{x}) = \hat{f}(\mathbf{x}) = -1$. If the overall prediction is obtained using a majority voting strategy, it follows:

$$\hat{f}(\mathbf{x}) = -1 \iff \left(\sum_{k=1}^K \hat{h}_k(\mathbf{x}) \right) \leq 0$$

Furthermore, there must be at least $\lceil \frac{K}{2} \rceil$ decision trees (base classifiers) of the forest \mathcal{T} whose output is -1 . That is, there exists $K^- \subseteq \{1, \dots, K\}$ with $|K^-| \geq \lceil \frac{K}{2} \rceil$, such that:

$$\hat{h}_k(\mathbf{x}) = -1, \forall k^- \in K^-$$

As we are operating in a binary classification setting, there must also exist $K^+ = \{1, \dots, K\} \setminus K^-$, which denote the set of classifier indices that output a positive label when input with \mathbf{x} , i.e., $\hat{h}_k(\mathbf{x}) = +1, \forall k^+ \in K^+$.

Our goal is to tweak the original input feature vector \mathbf{x} so as to adjust the prediction made by the ensemble from negative (-1) to positive ($+1$). We can skip all the trees indexed by K^+ , as these are already encoding the (positive) prediction we ultimately want. We therefore focus on each tree T_k where $k \in K^-$, and consider the set P_k^+ of all its positive paths. With each $p_{k,j}^+ \in P_k^+$ we associate an instance $\mathbf{x}_j^+ \in \mathcal{X}$ that *satisfies* that path – i.e., an instance whose adjusted feature values meet the boolean conditions encoded in $p_{k,j}^+$ to finish on a pos-labelled leaf, and therefore $\hat{h}_k(\mathbf{x}_j^+) = +1$.

Among all the possibly infinite instances satisfying $p_{k,j}^+$, we restrict to $\mathbf{x}_{j(\epsilon)}^+$ to be feature value changes with a “tolerance” of at most ϵ . We call it the ϵ -satisfactory instance of $p_{k,j}^+$. We consider

$p_{k,j}^+$ containing at most n boolean conditions, as specified by Equation 1. Therefore, for any (small) fixed $\epsilon > 0$, we build a feature vector $\mathbf{x}_{j(\epsilon)}^+$ as follows:

$$\mathbf{x}_{j(\epsilon)}^+[i] = \begin{cases} \theta_i - \epsilon & \text{if the } i\text{-th condition is } (x_i \leq \theta_i) \\ \theta_i + \epsilon & \text{if the } i\text{-th condition is } (x_i > \theta_i) \end{cases} \quad (2)$$

Having a single *global* tolerance ϵ for all the features works as long as we standardise every feature (e.g., using z-score or min-max). If we use standard z-score for each feature, the actual magnitude of the change cannot just depend on ϵ but instead needs to be considered a multiple of a unit of standard deviation from the feature mean. Let $\theta_i = \frac{t_i - \mu_i}{\sigma_i}$ be the z-score of the threshold on the i -th feature value, where t_i is the non-standardised value, μ_i and σ_i are the mean and standard deviation of the i -th feature respectively.³ Now, suppose that $x_i = \theta_i \pm \epsilon$ is the tweaked value of the i -th feature according to the ongoing transformation of the input vector \mathbf{x} . Therefore, $x_i = \frac{t_i - \mu_i}{\sigma_i} \pm \epsilon$. Returning to the original (i.e., non-standardised) feature scale, we obtain $x_i = t_i - \mu_i \pm \epsilon \sigma_i$. Depending on the sign, the tweaked feature is either moving closer to or farther away from the original feature mean μ_i , pivoting around t_i .

For each $p_{k,j}^+ \in P_k^+$ we transform our input feature vector \mathbf{x} into the ϵ -satisfactory instance $\mathbf{x}_{j(\epsilon)}^+$ that validates $p_{k,j}^+$. This leads us to a set of transformations $\Gamma_k = \bigcup_{j \in P_k^+} \mathbf{x}_{j(\epsilon)}^+$, associated with the k -th tree T_k .

Each resulting transformation in Γ_k may have an impact on other trees of the forest. There might exist $l \in K^+$ whose corresponding tree already provides the correct prediction when this is input with \mathbf{x} , $\hat{h}_l(\mathbf{x}) = +1$. It may also happen that by changing \mathbf{x} into \mathbf{x}' the prediction of the l -th tree is incorrect and now $\hat{h}_l(\mathbf{x}') = -1$. In other words, by changing \mathbf{x} into another instance $\mathbf{x}' \in \Gamma_k$ we are only guaranteed that the prediction of the k -th base classifier is correctly fixed, i.e., from $\hat{h}_k(\mathbf{x}) = -1$ to $\hat{h}_k(\mathbf{x}') = +1$. The overall prediction for \mathbf{x}' may or may not be fixed, where $\hat{f}(\mathbf{x}')$ may still output -1 , exactly as $\hat{f}(\mathbf{x})$ did.

If the change from \mathbf{x} to \mathbf{x}' also leads to $\hat{f}(\mathbf{x}') = +1$, then \mathbf{x}' will be a *candidate* transformation for \mathbf{x} . More formally, let $\Gamma = \bigcup_{k=1}^K \Gamma_k$ be the set of *all* the ϵ -satisfactory transformations of the original \mathbf{x} from the positive paths of all the trees in the forest. Our feature tweaking problem can then be generally defined as follows:

$$\mathbf{x}' = \arg \min_{\mathbf{x}_{j(\epsilon)}^+ \in \Gamma \mid \hat{f}(\mathbf{x}_{j(\epsilon)}^+) = +1} \left\{ \delta(\mathbf{x}, \mathbf{x}_{j(\epsilon)}^+) \right\}$$

In [6], it has already been proven that a problem similar to the one we define above is NP-hard as it reduces to DNF-MAXSAT. Our version, in fact, introduces an additional constraint (ϵ) on the possible way features can be tweaked and thus it is itself NP-hard. This problem definition is still valid for the base case when $K = 1$. There, the additional condition requiring $\hat{f}(\mathbf{x}_{j(\epsilon)}^+) = +1$ is not necessary because it is implicitly true by definition. In that scenario, the ensemble is composed of a single base classifier – i.e., the forest contains a single decision tree and tweaking its prediction also results in changing the overall prediction. Note that when there is only one decision tree, our problem can be solved *optimally*:

²In general, there can be multiple boolean conditions associated with a single feature.

³In practice, μ_i and σ_i are often unknown, so the sample mean and the sample standard deviation are used instead.

Input :

- ▷ An estimate function \hat{f} resulting from an ensemble of decision trees $\mathcal{T} = \{T_1, \dots, T_K\}$, each one associated with a base estimate \hat{h}_k , $k = 1, \dots, K$
- ▷ A feature vector \mathbf{x} representing a *true negative* instance, such that $f(\mathbf{x}) = \hat{f}(\mathbf{x}) = -1$
- ▷ A cost function δ
- ▷ A (small) threshold $\epsilon > 0$

Output:

- ▷ The optimal transformation \mathbf{x}' with respect to δ , such that $\hat{f}(\mathbf{x}') = +1$

```

begin
   $\mathbf{x}' \leftarrow \mathbf{x}$ ;
   $\delta_{min} \leftarrow +\infty$ ;
  for  $k = 1, \dots, K$  do
    if  $\hat{f}(\mathbf{x}) == \hat{h}_k(\mathbf{x})$  and  $\hat{h}_k(\mathbf{x}) == -1$  then
      /* retrieve the set of positive paths of the  $k$ -th decision tree */
       $P_k^+ \leftarrow \text{getPositivePaths}(T_k)$ ;
      foreach  $p_{k,j}^+ \in P_k^+$  do
        /* generate the  $\epsilon$ -satisfactory instance associated with
           the  $j$ -th positive path of the  $k$ -th decision tree */
         $\mathbf{x}_{j(\epsilon)}^+ \leftarrow \text{buildPositiveInst}(\mathbf{x}, p_{k,j}^+, \epsilon)$ ;
        if  $\hat{f}(\mathbf{x}_{j(\epsilon)}^+) == +1$  then
          if  $\delta(\mathbf{x}, \mathbf{x}_{j(\epsilon)}^+) < \delta_{min}$  then
             $\mathbf{x}' \leftarrow \mathbf{x}_{j(\epsilon)}^+$ ;
             $\delta_{min} \leftarrow \delta(\mathbf{x}, \mathbf{x}_{j(\epsilon)}^+)$ ;
          end
        end
      end
    end
  end
  return  $\mathbf{x}'$ ;
end

```

Algorithm 1: The FEATURE TWEAKING ALGORITHM

We can enumerate all the positive paths, choose the one with the minimum cost, and check if the threshold of tolerance ϵ is satisfied on each feature. Because base trees are interconnected through the features they share, simply enumerating positive paths does not work for an ensemble of trees since the output of a base tree may affect outputs of its sibling trees. The presence of the condition $\hat{f}(\mathbf{x}_{j(\epsilon)}^+) = +1$ circumvents this issue by querying the model itself on the correctness of the ϵ -transformation of an instance.

2.5 The Feature Tweaking Algorithm

Our approach takes as input 4 key components: (i) The trained ensemble model \hat{f} ; (ii) A feature vector \mathbf{x} that represents a true negative instance; (iii) A cost function δ measuring the “effort” required to transform the true negative instance into a positive one; and (iv) A positive threshold ϵ that bounds the tweaking of each single feature to pass every boolean test on a positive path of each tree. The result being the transformation \mathbf{x}' of the original \mathbf{x} that exhibits the minimum cost according to δ . The detailed description is presented in Algorithm 1.

In the worst case, our algorithm examines all K trees of the forest, although it investigates only those trees whose base predictions are negatives (neg-labelled). Then, all the positive paths of each tree are considered, and for each of those a potential candidate ϵ -transformation is built according to the scheme proposed in Equation 2. As such, the number of steps depends on the number of positive paths on each tree, which in turn is related to the number of leaves.

We stated in Section 2.3 that we cannot provide *apriori* any limit to the depth of a decision tree, and therefore to its number of leaves. However, in practice these can be bounded whilst training the model. We therefore set the maximum length of each root-to-leaf path, *i.e.*, the depth of each tree, to be at most equal to the number of input features n . Considering our goal, this is not a limitation as each transformation should at most affect each feature exactly once. The total number of positive paths to be examined will then be limited by $K2^n$, and so the worst case complexity is $O(2^n)$. It follows that our method might be unsuitable for dealing with high-dimensional datasets, such as text, images, or videos. In reality such a setting would not really make sense when transforming input instances in the original feature space (*e.g.*, change a few words on a document or a few pixels of an image).

Although the search space can be exponential in the number of features, in practice the number of positive paths of each tree is significantly smaller. This makes our method feasible on average input sizes (*i.e.*, below 100 features). For example, in our experiments we found that the maximum depth of each tree leading to the best classification results is significantly smaller than the total number $n = 45$ of features (see Table 2). In addition, many positive paths may share several boolean conditions, especially when extracted from the same decision tree. This allows us to avoid tweaking the same input feature multiple times according to the same condition by using some caching mechanism.

Finally, our algorithm can be easily parallelised since each tree can be explored independently from the others for any given instance – *i.e.*, we can adjust the k -th tree while keeping the remaining $K \setminus k$ trees simultaneously fixed for all $k \in \{1, \dots, K\}$.

3 USE CASE: IMPROVING AD QUALITY

We demonstrate the utility of our method when applied to a real-world use case in online advertising. We investigate how our algorithm can be used to improve the *quality* of advertisements served by the *Yahoo Gemini* ad network.

3.1 Why Ad Quality?

A main source of monetisation for online web services comes in the form of advertisements (*ads* for short) impressed in dedicated real estate units of rendered web pages. Online publishers operating these web services typically reserve predefined slots within their streams, utilising a third party ad network to deliver ad inventory to impress within them. Ad networks free publishers from running their own ad servers as they decide for them which ads should be placed at which slots, when and to whom. Advertisers rely on ad networks to optimise their *return on investment* – for example through targeting the right audience according to the advertiser budget and marketing strategy.

A trustworthy relationship between advertisers, publishers and ad networks is instrumental to the success of online advertising. On the advertiser side, ad networks provide advertisers with tools for monitoring key performance indicators (e.g., number of ad *impressions*, *click-through rate* or CTR, *bounce rate*) as well as effective mechanisms to overcome fraudulent activities such as *click spam* [7, 22]. On the publisher side, ad networks provide mechanisms that allow users to hide ads they dislike and indicate the reasons for doing so. This information can be used by ad networks to ensure that the ads they serve do not negatively affect user engagement on a publisher website, as well as reporting to advertisers on the *quality* of their ads.

As with any self served content delivery platform, an ad network has a varying distribution of quality with many ads being of low quality. Not serving them may not be an option when supply is exhausted. Therefore, another approach to positively shift the inventory quality distribution is to leverage the interpretability of the internal machinery of existing ad quality prediction models – i.e., binary classifiers – so as to offer actionable *recommendations*. The intention of these programmatically-computed recommendations is to provide advertisers with guidance on how they can improve their ad quality at scale. Such a system yields value for all beneficiaries in this advertising ecosystem ultimately culminating with a better user experience.

Returning to our work, by applying our feature tweaking algorithm introduced in Section 2.5, we show in the rest of this paper how we can transform a low quality ad into a set of new “proposed” high quality ads by shifting their position in an ad quality feature space. The algorithm employs the internals of a learned binary classifier to *tweak* the feature-based representation of a low quality ad so that the new “proposed” ads are promoted to high quality ones when re-input to the classifier. Each transformation is associated with a *cost*, allowing us to generate actionable suggestions from the “proposed” instances with the least cost, so as to improve low quality ads. We validate our approach on a dataset of mobile native ads served by the *Yahoo Gemini*⁴ ad network.

3.2 A Definition of Ad Quality

Many factors can affect the quality of an ad: its *relevance*, i.e., whether the ad matches the user interest [19]; the *pre-click* experience, i.e., whether the ad annoys a user [27]; and finally the *post-click* experience, i.e., whether the *ad landing page*⁵ meets the user click intent that brought them to the landing page [12]. We focus on the latter, the post-click experience, following from [1, 12].

Inspired by these studies, we define and measure the quality of ads using the time spent on their landing pages as a proxy, referred to as *dwell time*. We know from [12] that ad landing pages exhibiting long dwell times promote a positive long-term post-click experience. Based on this definition of ad quality, we design a binary classifier that effectively separates between *low* and *high* quality ads, i.e., ad landing pages whose dwell time is below or above a threshold τ , respectively. We compute τ as the median of all the sample means of dwell times observed for a large set of ad landing pages. Intuitively,

an ad landing page is of high quality if its average dwell time is greater than τ – i.e., if the average time users spent on the page is greater than the average time users spent on at least 50% of any other ad landing page. Although more sophisticated approaches can be designed [1], this is not the main goal of this research, and we leave it for future work.

3.3 Predicting Ad Quality

To apply our algorithm to our use case, we first need to learn a binary classifier that predicts whether an ad is of high quality or not, given a feature-based representation of the ad creative and the landing page.

3.3.1 Ad Feature Engineering. A sample of the set of features used in this work is listed in Table 1. This is based on the same set used in [1], with the additional “Language” category (marked with “†” in the Table). Due to space constraints we do not show all the features, and invite the reader to refer to [1] for a complete description of them.

Each feature in the table is associated with a *category* and a *source*. The former indicates the type of features whilst the latter specifies whether the feature is computed from the ad *landing page* (LP), the ad *creative* (CR), or a combination of the two (CR-LP). Although our focus is on the post-click experience (i.e., ad landing page), we aim to obtain the best performing model for predicting the quality of the ads, and hence include both pre-click (i.e., ad creative) and historical features; the latter not participating in the tweaking process as – by definition – they cannot be altered.

3.3.2 Learning Binary Classifiers. As our feature tweaking algorithm is designed to work on tree-based ensemble classifiers, we train the following learning models to find our best estimate \hat{f} : Decision Trees (DT) [18], which can be thought of as a special case of an ensemble with a single tree; Gradient Boosted Decision Trees (GBDT) [9], and Random Forests (RF) [2].

The original dataset $\mathcal{D} = \{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_m, y_m)\}$ is split into two datasets $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{test}}$ using stratified random sampling. $\mathcal{D}_{\text{train}}$ is used for training the models and accounts for 80% of the total number of instances in \mathcal{D} , whilst $\mathcal{D}_{\text{test}}$ contains the remaining held-out portion used for evaluating the models. $\mathcal{D}_{\text{train}}$ is also used to perform model selection, which is achieved by tuning the hyperparameters specific for each.

With every combination of model and corresponding hyperparameters, we run a 10-fold cross validation to find the best settings for each model – i.e., the one with the best cross validation performance. We measure this performance using the *Area Under the Curve of the Receiver Operating Characteristic* (ROC AUC). Each model is in turn re-trained on the whole $\mathcal{D}_{\text{train}}$ using the best hyperparameter setting. Finally, the overall best model is deemed to be the one that best performs on the test set $\mathcal{D}_{\text{test}}$.

3.3.3 Labelled Dataset of Ads. We collect a random sample of 1,500 ads served by the *Yahoo Gemini* ad network on a mobile app during one month. To ensure reliable estimates of dwell time, we only consider ads clicked at least 500 times. We saw that the distribution was skewed with around 80% of the instances having an average dwell time within approximately 100 seconds, whereas the remaining 20% sat in the long tail of very long dwell times. The

⁴<https://gemini.yahoo.com>

⁵We refer to *ad landing page* as the web page of the advertiser that a user is redirected to after clicking on an *ad*.

Table 1: Set of features used to characterise ad *creative* (CR), ad *landing page* (LP), or both (CR-LP).

Category	Source	Description
Language [†]	CR	This set of features capture the extent to which the text of the ad creative may include adult, violent, or spam content (e.g., ADULT_SCORE, HATE_SCORE, and SPAM_SCORE)
DOM	LP	This set of features are derived from the elements extracted from the HTML DOM of the ad landing page, such as the main textual content (LANDING_MAIN_TEXT_LENGTH), the total number of internal and external hyperlinks (LINKS_TOTAL_COUNT), the ratio of main text length to the total number of hyperlinks on the page (LINKS_MAIN_LENGTH_TOTAL_RATIO), etc.
Readability	CR-LP	These features range from a simple count of tokens (words) in the text of the ad creative and landing page to well-known scores for measuring the summarisability/readability of a text (e.g., READABILITY_SUMMARY_SCORE), etc.
Mobile Optimising	LP	This set of features describe the degree of mobile optimisation of the ad landing page by measuring the ability of it to be tuned to different screen sizes (VIEW_PORT), testing for the presence of a click-to-call button (CLICK_TO_CALL), etc.
Media	LP	These features refer to any media content displayed within the ad landing page, such as the number of images (NUM_IMAGES), etc.
Input	LP	This set of features represent all the possible input types available on the ad landing page, such as the number of checkboxes, drop-down menus, and radio buttons (NUM_INPUT_CHECKBOX, NUM_INPUT_DROPDOWN, NUM_INPUT_RADIO), etc.
Content & Similarity	CR-LP	These features extract the set of Wikipedia <i>entities</i> from the ad creative and landing page (NUM_CONCEPT_ANNOTATION), and measure the <i>Jaccard</i> similarity between those two sets (SIMILARITY_WIKI_IDS), etc.
History	LP	These features measure historical indicators, such as the median <i>dwelt time</i> as computed from the last 28 days of observed ad clicks (HISTORICAL_DWELLTIME), and the <i>bounce rate</i> – i.e., the proportion of ad clicks whose dwell time is below 5 seconds (HISTORICAL_BOUNCE_RATE), etc.

median τ of those averages, calculated as described in Section 3.2, is equal to ≈ 62.5 seconds. We therefore reach an evenly balanced ground truth, where 50% of the instances have an average dwell time at most equal to τ and the remaining 50% above τ .

To build our labelled dataset \mathcal{D} , for each ad we extract the features listed in Section 3.3.1. As these are a mix of categorical (i.e., discrete) and continuous features, we apply *one-hot encoding* to transform each k -valued categorical feature into a k -dimensional binary vector. The i -th component of such a vector evaluates to 1 if and only if the value of the original feature is i , and 0 otherwise. We also standardise continuous features by transforming their original values into their corresponding z -scores [11]. Finally, we obtain a set of 45 features.

3.3.4 Offline Evaluation. From the balanced labelled dataset above we derive two random partitions $\mathcal{D}_{\text{train}}$ and $\mathcal{D}_{\text{test}}$, which contain 80% and 20% of the total samples respectively. We again validate our performance by running a 10-fold cross validation on $\mathcal{D}_{\text{train}}$ selecting between different models and their varying hyperparameter settings.

For DT, we test two different node-splitting criteria s : *Gini index* and *entropy*. We also set the maximum depth of the tree d to the total number of features. For GBDT, we use four values of the number K of base trees in the ensemble, with $K = \{10, 100, 500, 1,000\}$ in combination with the learning rates α , 0.001, 0.01, 0.05, 0.1, 1. Finally, for RF we test the same number of base trees as for GBDT whilst again bounding the maximum depth of each base tree to the total number of features. For each learning model we report

Model	s	α	d	K	ROC AUC
RF	–	–	16	1,000	0.93
GBDT	–	0.1	–	100	0.92
DT	<i>entropy</i>	–	3	–	0.84

Table 2: Best cross validation hyperparameter settings for each learning model (“–” if the hyperparameter is not considered).

	RF	GBDT	DT
F_1	0.84	0.81	0.75
MCC	0.66	0.63	0.49

Table 3: Evaluation of best performing models on the test set $\mathcal{D}_{\text{test}}$.

in Table 2 the hyperparameter settings leading to the best cross validation ROC AUC. The overall best performing model is RF with an ensemble of 1,000 base trees and maximum depth 16.

To avoid mixing model selection with model evaluation, we re-train each model on $\mathcal{D}_{\text{train}}$ using its best hyperparameter setting, and assess its validity on the held-out and unseen test set $\mathcal{D}_{\text{test}}$. We measure two standard quality metrics, F_1 and Matthews Correlation Coefficient (MCC) [17]. Table 3 shows the results. RF is the best performing model also with respect to the ability of generalising its predictive power to previously unseen examples.

Compared with the results reported in [12], we notice a remarkable increase of ROC AUC (+10.7%) and a small improvement of F_1 (+1.2%). In their work, Logistic Regression (LogReg) [26] was the




Figure 1: Top-20 most important features of our RF model.

best model (ROC AUC = 0.84 and $F_1 = 0.83$). Our increased performance comes from a combination of a more rigorous procedure for determining the threshold τ ,⁶ a more effective learning model, RF, and a larger set of features.

We also calculate the “importance” of each feature from the learned RF model. Figure 1 lists the top-20 most important ones. As in [1, 12], historical features have significant predictive power. We keep historical features because we want the learned model for which we run our feature tweaking algorithm to be the most effective possible. However, our feature tweaking algorithm will ignore historical features when generating recommendations, as it only considers *adjustable* ad features, *i.e.*, features that advertisers can actively alter to improve the quality of their ads.

From now on, we use RF as our learned model to generate actionable suggestions on which features to tweak to turn a low quality ad into a high quality one.

4 EXPERIMENTS: AD FEATURE RECOMMENDATIONS

We validate the recommendations generated with our approach, by applying our feature tweaking algorithm to our learned RF model. Any \mathbf{x}' that results from a valid (*i.e.*, positive) ϵ -transformation of the original negative instance \mathbf{x} encapsulates a set of directives on how to positively change the ad features. We compute the vector \mathbf{r} resulting from the component-wise difference between \mathbf{x}' and \mathbf{x} , which is $\mathbf{r}[i] = \mathbf{x}'[i] - \mathbf{x}[i]$. Then for each feature i , such that $\mathbf{r}[i] \neq 0$ (*i.e.*, $\mathbf{x}'[i] \neq \mathbf{x}[i]$), this vector provides the *magnitude* and the *direction* of the changes that should be made on feature i . The magnitude denotes the absolute value of the change (*i.e.*, $|\mathbf{x}'[i] - \mathbf{x}[i]|$), whilst the direction indicates whether this is an *increase* or a *decrease* of the original value of feature i (*i.e.*, $\text{sgn}(\mathbf{x}'[i] - \mathbf{x}[i])$). Finally, to derive the final list of recommendations, we sort \mathbf{r} according to the feature ranking, as shown in Figure 1.

4.1 The impact of hyperparameters δ and ϵ

Our approach depends on a *tweaking cost* (δ) associated with transforming a negative instance (low quality ad) into a positive instance (high quality ad), and a *tweaking tolerance* (ϵ) used to change each individual ad feature. We first explore how ϵ impacts on the ad

⁶In their work, the threshold was set as the *flat* median of the observed dwell times of all ads.

ϵ	0.01	0.05	0.10	0.50	1.00
ad coverage (%)	58.5	64.2	72.3	77.4	63.2

Table 4: The impact of tolerance threshold ϵ on ad coverage.

coverage, which is the percentage of ads for which our approach is able to provide recommendations. We experiment with five values of ϵ : 0.01, 0.05, 0.1, 0.5, and 1. These values can be thought of as multiples of a unit of standard deviation from each individual feature mean, as discussed in Section 2.4. Table 4 shows the highest coverage is when $\epsilon = 0.5$.

Although some low quality ads cannot be transformed, those that can are often associated with multiple transformations. Figure 2 shows the distribution of ϵ -transformations across the set of ads, generated using different values of ϵ (except $\epsilon = 1$ which is similar to $\epsilon = 0.05$). All the distributions are skewed offering a high number of transformations proposed for few ads. Interestingly, the number of transformations is more evenly distributed across the ads when ϵ increases. This is in agreement with the finding above, where larger values of ϵ result in a higher coverage before decreasing again between $\epsilon = 0.5$ and 1.

To choose the most appropriate transformation for an ad, we experiment with several tweaking cost functions δ , each taking as input the original (\mathbf{x}) and the transformed (\mathbf{x}') feature vectors:

- *tweaked_feature_rate*: proportion of features affected by the transformation of \mathbf{x} into \mathbf{x}' (range = $[0, 1]$);
- *euclidean_distance*: Euclidean distance between \mathbf{x} and \mathbf{x}' (range = \mathbb{R});
- *cosine_distance*: 1 minus the cosine of the angle between \mathbf{x} and \mathbf{x}' (range = $[0, 2]$);
- *jaccard_distance*: one’s complement of the Jaccard similarity between \mathbf{x} and \mathbf{x}' (range = $[0, 1]$);
- *pearson_correlation_distance*: 1 minus the Pearson’s correlation coefficient between \mathbf{x} and \mathbf{x}' (range = $[0, 2]$).

Up to a certain value, the tolerance ϵ is positively correlated with the ad coverage. We explore how it impacts the five tweaking cost functions. Figure 3(a) plots the micro-average costs and Figure 3(b) shows the median of all the individual per-ad average costs. In general, the greater the tolerance the higher the cost (except for *tweaked_feature_rate* and *jaccard_distance* when $\epsilon = 1$); thus a trade-off between ϵ (*i.e.*, ad coverage) and the cost of ad transformations δ is desirable.

4.2 Evaluating Recommendations

We first present descriptive statistics on the recommendations obtained with our approach on a set of 100 low quality ad landing pages, the true negative instances in $\mathcal{D}_{\text{test}}$. Each recommendation either suggests to *increase* or *decrease* the value of a given feature. Overall, the recommendations are almost evenly distributed over the two cases above.

In Figure 4, we list the top-5 most frequent features recommended to be tweaked according to the top-1, top-2 and top-3 proposed ϵ -transformations by measuring the relative frequency of each feature appearing among each of the ϵ -transformations. The most frequent feature is LINKS_TEXT_LENGTH_TOTAL_RATIO in all settings, which measures the ratio of text length to the total number of hyperlinks




Figure 2: Distribution of per-ad ϵ -transformations.




Figure 3: The impact of tolerance threshold ϵ on costs δ .




Figure 4: Top-5 most frequent features appearing in the top-1, top-2 and top-3 ϵ -transformations ($\epsilon = 0.05$).

in the ad landing page. Interestingly, *all* the recommendations concerning this feature suggest to *decrease* its value. This indicates that low quality ad landing pages generally exhibit an unbalanced ratio of text to hyperlinks suggesting that saturating a page in links rather than content has negative effects on dwell time.

We measure the Pearson’s correlation coefficient (ρ) between feature rankings appearing in the top-1, top-2 and top-3 ϵ -transformations. All three rankings are strongly related with each other, with top-1 reaching $\rho = 0.93$ and 0.81 when compared to top-2 and top-3, respectively. Similarly, top-2 is highly correlated to top-3 ($\rho = 0.79$).⁷ We also compute the correlation coefficient between

⁷All values are statistically significant at $\alpha = 0.01$.

top-1 ϵ -transformations for all values of ϵ , and $\delta = \text{cosine_distance}$. The top-1 rankings derived from $\epsilon = 0.05$ and 0.1 are the highest correlated ($\rho = 0.92$). However, there is no statistical significant correlation between top-1 rankings when $\epsilon = 0.1$ and 0.5 , indicating that higher values of tolerance may impact more on the features requiring change.

We now perform a *qualitative* and *quantitative* assessment of such recommendations. For each low quality landing page, we focus on its top- k ϵ -transformations, *i.e.*, the k less costly transformations according to the cost function δ . In turn, each transformation contains a list of recommendations, sorted by the feature rank they refer to. We set the hyperparameters to $\epsilon = 0.05$ and $\delta = \text{cosine_distance}$,

as this combination provides the best trade-off between ad coverage and average cost. We consider the top-3 ϵ -transformations suggested for each ad landing page. Around 91.0% of landing pages can be associated with all three ϵ -transformations, whereas our algorithm provides the remaining 7.5% and 1.5% with two and one ϵ -transformation, respectively.

We also asked an internal team of creative strategists (CS)⁸ to validate the recommendations generated by our approach. Each CS was assigned a set of ad landing pages with the corresponding ϵ -transformations, and additional metadata useful for assessing the recommendations within each transformation. The same set of ad landing pages – and therefore the same list of recommendations – was assessed by two CSs, who were asked to rate each recommendation as *helpful*, *non-helpful*, or *non-actionable*. A recommendation is deemed *helpful* when it is likely to help the advertiser to improve the user experience of the ad, and *non-helpful* otherwise. A *non-actionable* recommendation is one that cannot be practically implemented. Whenever a disagreement occurred, a third CS was called to resolve the conflict.

Overall, 57.3% of all the generated recommendations are rated helpful with an inter-agreement rate of 60.4% and only 0.4% result in a non-actionable suggestion. We also look at the 42.3% non-helpful recommendations, and saw that about 25% can be considered “neutral”; that is, they would not hurt the user experience if discarded as well as not adding any positive value if implemented.

Non-helpful tweaks might occur due to two reasons. First, the learned model we leverage for generating feature recommendations – no matter how accurate it is – is not perfect; therefore, a true negative instance that is transformed into a positive prediction does not necessarily mean it is *actually* positive. Second, tuning the hyperparameters (δ and ϵ) of our algorithm affects the set of candidate transformations. As such, limiting non-helpful tweaks can be achieved by improving the accuracy of the learned model and choosing values of the hyperparameters to minimize errors.

Furthermore, when we further look into the non-actionable recommendations we see that these are related to the features ADULT_SCORE and NUM_INPUT_DROPDOWN. Our algorithm suggests to decrease the value of those features; however, the ad landing pages do not contain adult words nor drop-downs. Most likely, the ad copies and landing pages used to generate recommendations have changed before the CSs performed their assessment.

Finally, we measure the “helpfulness” of each feature recommendation as follows:

$$\text{helpfulness}(i) = \frac{|\text{helpful}(i)|}{|\text{helpful}(i)| + |\neg\text{helpful}(i)|}$$

This computes the relative frequency of recommendations for feature i as being described as helpful by the CS team. In Figure 5, we report the ranked list of features involved in the top-10 most helpfulness recommendations. A similar ranking is obtained if we weight the helpfulness score on the basis of the *overall* relative recommendation frequency. The majority of the most helpful recommendations were features extracted from the DOM structure and content of the ad landing page, indicating that high quality landing pages should exhibit a good balance between textual content and

⁸Creative strategists work with advertisers’ web masters on strategic choices to help them developing effective advertising messages.




Figure 5: Top-10 most helpful feature recommendations according to the helpfulness score.

hyperlinks. Those features were the most predictive in our RF ad quality model (Figure 1).

5 RELATED WORK

The research challenge addressed in this work is largely unexplored. Although machine learning has received a lot of attention in recent years, the focus has been mainly on the accuracy, efficiency, scalability, and robustness of the proposed various techniques. Works on extracting actionable knowledge from machine-learned models have been mostly conducted within the business and marketing domains. Early works have focused on the development of interestingness metrics as proxy measures of knowledge actionability [3, 10].

Another line of research on actionable knowledge discovery concerns post-processing techniques. Liu *et al.* propose methods for pruning and summarizing learned rules, as well as matching rules by similarity [14, 15]. Cao *et al.* present domain-driven data mining; a paradigm shift from a research-centered discipline to a practical tool for actionable knowledge [4, 5]. The authors discuss several frameworks for handling different problems and applications.

Many works discuss post-processing techniques specifically tailored to decision trees [8, 16, 24, 25]. Yang *et al.* study the problem of proposing actions to maximise the expected profit for a group of input instances based on a single decision tree, and introduce a greedy algorithm to approximately solve such a problem [25]. This is significantly different from our work; in fact, our work is more related to the one presented by Cui *et al.* [6]. Here, the authors propose a method to support actionability for additive tree models (ATMs), which is to find the set of actions that can change the prediction of an input instance to a desired status with the minimum cost. The authors formulate the problem as an instance of integer linear programming (ILP) and solve it using existing techniques.

Similarly to Cui *et al.*, we also consider transforming the prediction for a given instance output by an ensemble of trees, and we introduce an algorithm that finds the *exact* solution to the problem. Our work differs from theirs in several aspects: (i) We tackle the theoretical intractability (NP-hardness) of the problem by designing an algorithm that creates a feedback loop with the original model to build a set of candidate transformations without the need, in practice, to explore the entire exponential search space; (ii) We introduce another hyperparameter (ϵ) to govern the amount of

change that each feature can sustain; (iii) We experiment with five concrete functions describing the cost of each transformation (δ); (iv) We leverage on the importance of each feature derived from the model to rank the final list of recommendations; (v) We focus on the actual recommendations generated, and how they impact in practice on a real use case, if properly implemented.

More recent work on related topics are those of Ribeiro *et al.* [20, 21]. In particular, [21] presents LIME, a method that aims to explain the predictions of any classifier by learning an interpretable model that is specifically built around the predictions of interest. They frame this task as a submodular optimization problem, which the authors solved using a well-known greedy algorithm achieving performance guarantees. They test their algorithm on different models for text (e.g., random forests) and image classification (e.g., neural networks), and validate the utility of generated explanations both via simulated and human-assessed experiments.

6 CONCLUSIONS

Machine-learned models are often designed to favour accuracy of prediction at the expense of human-interpretability. However, in some circumstances it becomes important to understand why the model returns a certain prediction on a given instance and how such an instance could be transformed in such a way that the model changes its original prediction. We investigate this problem within the context of general ensembles of tree-based classifiers, which has been proven to be NP-hard. We then introduce an algorithm that is able to transform a true negative instance into a set of new “proposed” positive instances by shifting their position in the feature space. The algorithm leverages the internals of the learned ensemble to *tweak* the feature-based representation of a true negative instance so that the new “proposed” ones are promoted to a positive classification when re-input to the classifier.

Despite computationally intractable in the worst case, we demonstrate the applicability of our approach on a real-world use case in online advertising. The feasibility of our approach has been achieved by (i) setting an upper bound to the maximum number of changes affecting each instance (i.e., at most equivalent to the number of features), which can be controlled at training time, and by (ii) creating a feedback loop with the original model to build a set of candidate transformations without the need, in practice, to explore the entire exponential search space.

After designing an effective Random Forest classifier able to separate between *low* and *high* quality ads – our application scenario – we automatically provide “actionable” suggestions on how to optimally convert a low quality ad (negative instance) into a high quality one (positive instance) using our approach. To illustrate the outcomes of our algorithm, we assess the quality of the recommendations that our method generates from a dataset of ads served by a large ad network, *Yahoo Gemini*. An evaluation conducted by an internal team of creative strategists shows that 57.3% of the provided recommendations are indeed helpful, and likely to improve the ad quality, if implemented.

In future work, we plan to extend the approach presented in this work to multi-class setting as well as to other learning models, and to encapsulate it into a reinforcement learning framework.

ACKNOWLEDGMENTS

The authors would like to thank Huw Evans, Mahlon Chute, and all the Yahoo’s internal team of creative strategists for their invaluable contributions in evaluating the quality of the method presented in this work.

REFERENCES

- [1] Nicola Barbieri, Fabrizio Silvestri, and Mounia Lalmas. 2016. Improving Post-Click User Engagement on Native Ads via Survival Analysis. In *WWW ’16. International World Wide Web Conferences Steering Committee*, 761–770.
- [2] Leo Breiman. 2001. Random Forests. *Machine Learning* 45, 1 (Oct. 2001), 5–32.
- [3] Longbing Cao, Dan Luo, and Chengqi Zhang. 2007. Knowledge Actionability: Satisfying Technical and Business Interestingness. *International Journal of Business Intelligence and Data Mining* 2, 4 (Dec. 2007), 496–514.
- [4] Longbing Cao and Chengqi Zhang. 2006. Domain-Driven Actionable Knowledge Discovery in the Real World. In *PAKDD ’06. Springer-Verlag*, 821–830.
- [5] Longbing Cao, Huaifeng Zhang, E.K. Park, Dan Luo, Yanchang Zhao, and Chengqi Zhang. 2013. Flexible Frameworks for Actionable Knowledge Discovery. *IEEE TKDE* 22, 9 (Sep. 2013), 1299–1312.
- [6] Zhicheng Cui, Wenlin Chen, Yujie He, and Yixin Chen. 2015. Optimal Action Extraction for Random Forests and Boosted Trees. In *KDD ’15. ACM*, 179–188.
- [7] Neil Daswani and Michael Stoppelman. 2007. The Anatomy of Clickbot.A. In *HotBots ’07. USENIX Association*, 11–11.
- [8] Jianfeng Du, Yong Hu, Charles X. Ling, Ming Fan, and Mei Liu. 2011. Efficient Action Extraction with Many-to-Many Relationship between Actions and Features. In *LORI ’11. Springer Berlin Heidelberg*, 384–385.
- [9] Jerome H. Friedman. 2002. Stochastic Gradient Boosting. *Computational Statistics and Data Analysis* (2002).
- [10] Robert J. Hilderman and Howard J. Hamilton. 2000. Applying Objective Interestingness Measures in Data Mining Systems. In *PKDD ’00. Springer Berlin Heidelberg*, 432–439.
- [11] E. Kreyzig. 1979. *Introductory Mathematical Statistics: Principles and Methods*. Wiley.
- [12] Mounia Lalmas, Janette Lehmann, Guy Shaked, Fabrizio Silvestri, and Gabriele Tolomei. 2015. Promoting Positive Post-Click Experience for In-Stream Yahoo Gemini Users. In *KDD ’15. ACM*, 1929–1938.
- [13] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep Learning. *Nature* 521 (May 2015), 436–444. Issue 7553.
- [14] Bing Liu and Wynne Hsu. 1996. Post-analysis of Learned Rules. In *AAAI ’96. AAAI Press*, 828–834.
- [15] Bing Liu, Wynne Hsu, and Yiming Ma. 1999. Pruning and Summarizing the Discovered Associations. In *KDD ’99. ACM*, 125–134.
- [16] Karim Masud and M. Rahman Rashedur. 2013. Decision Tree and Naïve Bayes Algorithm for Classification and Generation of Actionable Knowledge for Direct Marketing. *Journal of Software Engineering and Applications* 6, 4 (2013), 196–206.
- [17] B. W. Matthews. 1975. Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochimica et Biophysica Acta* 405 (Oct. 1975), 442–451.
- [18] J. R. Quinlan. 1986. Induction of Decision Trees. *Machine Learning* 1, 1 (March 1986), 81–106.
- [19] Hema Raghavan and Dustin Hillard. 2009. A Relevance Model Based Filter for Improving Ad Quality. In *SIGIR ’09. ACM*, 762–763.
- [20] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. Model-agnostic interpretability of machine learning. In *2016 ICML Workshop on Human Interpretability in Machine Learning (WHI 2016)*, 91–95.
- [21] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. In *KDD ’16. ACM*, 1135–1144.
- [22] Brett Stone-Gross, Ryan Stevens, Apostolis Zarras, Richard Kemmerer, Chris Kruegel, and Giovanni Vigna. 2011. Understanding Fraudulent Activities in Online Ad Exchanges. In *IMC ’11. ACM*, 279–294.
- [23] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2013. Intriguing properties of neural networks. *CoRR* (2013).
- [24] Qiang Yang, Jie Yin, Charles Ling, and Rong Pan. 2007. Extracting Actionable Knowledge from Decision Trees. *IEEE TKDE* 19, 1 (Jan. 2007), 43–56.
- [25] Qiang Yang, Jie Yin, Charles X. Ling, and Tielin Chen. 2003. Postprocessing Decision Trees to Extract Actionable Knowledge. In *ICDM ’03. IEEE Computer Society*, 685–688.
- [26] Hsiang-Fu Yu, Fang-Lan Huang, and Chih-Jen Lin. 2011. Dual Coordinate Descent Methods for Logistic Regression and Maximum Entropy Models. *Machine Learning* 85, 1–2 (Oct. 2011), 41–75.
- [27] Ke Zhou, Miriam Redi, Andrew Haines, and Mounia Lalmas. 2016. Predicting Pre-click Quality for Native Advertisements. In *WWW ’16. International World Wide Web Conferences Steering Committee*, 299–310.