

FairTest: Discovering Unwarranted Associations in Data-Driven Applications*

Florian Tramèr¹, Vaggelis Atlidakis², Roxana Geambasu², Daniel Hsu²,
Jean-Pierre Hubaux³, Mathias Humbert⁴, Ari Juels⁵, Huang Lin³

¹Stanford, ²Columbia University, ³EPFL, ⁴Saarland University, ⁵Cornell Tech, Jacobs Institute

Abstract—In a world where traditional notions of privacy are increasingly challenged by the myriad companies that collect and analyze our data, it is important that decision-making entities are held accountable for unfair treatments arising from irresponsible data usage. Unfortunately, a lack of appropriate methodologies and tools means that even identifying unfair or discriminatory effects can be a challenge in practice.

We introduce the *unwarranted associations (UA) framework*, a principled methodology for the discovery of unfair, discriminatory, or offensive user treatment in data-driven applications. The UA framework unifies and rationalizes a number of prior attempts at formalizing algorithmic fairness. It uniquely combines multiple investigative primitives and fairness metrics with broad applicability, granular exploration of unfair treatment in user subgroups, and incorporation of natural notions of utility that may account for observed disparities.

We instantiate the UA framework in *FairTest*, the first comprehensive tool that helps developers check data-driven applications for unfair user treatment. It enables scalable and statistically rigorous investigation of associations between application outcomes (such as prices or premiums) and sensitive user attributes (such as race or gender). Furthermore, *FairTest* provides *debugging capabilities* that let programmers rule out potential confounders for observed unfair effects.

We report on use of *FairTest* to investigate and in some cases address disparate impact, offensive labeling, and uneven rates of algorithmic error in four data-driven applications. As examples, our results reveal subtle biases against older populations in the distribution of error in a predictive health application and offensive racial labeling in an image tagger.

1. Introduction

Today’s applications collect and mine vast quantities of personal information. Such data can boost applications’ utility by personalizing content and recommendations, increase business revenue via targeted product placement, and improve a wide range of socially beneficial services, such as healthcare, disaster response, and crime prevention.

The collection and use of such data raise two important challenges. First, massive data collection is perceived by many as a major threat to traditional notions of individual privacy. Second, the use of personal data for algorithmic

decision-making can have unintended and harmful consequences, such as unfair or discriminatory treatment of users.

In this paper, we deal with the latter challenge. Despite the personal and societal benefits of today’s data-driven world, we argue that companies that collect and use our data have a responsibility to ensure equitable user treatment. Indeed, European and U.S. regulators, as well as various policy and legal scholars, have recently called for increased *algorithmic accountability*, and in particular for decision-making tools to be audited and “tested for fairness” [1], [2].

There have been many recent reports of unfair or discriminatory effects in data-driven applications, mostly qualified as unintended consequences of data heuristics or overlooked bugs. For example, Google’s image tagger was found to associate racially offensive labels with images of black people [3]; the developers called the situation a bug and promised to remedy it as soon as possible. In another case [4], *Wall Street Journal* investigators showed that Staples’ online pricing algorithm discriminated against lower-income people. They referred to the situation as an “unintended consequence” of Staples’s seemingly rational decision to adjust online prices based on user proximity to competitors’ stores. This led to higher prices for low-income customers, who generally live farther from these stores.

Staples’ intentions aside, it is evidently difficult for programmers to foresee all the subtle implications and risks of data-driven heuristics. Moreover, these risks will only increase as data is passed through increasingly complex machine learning (ML) algorithms whose associations and inferences may be impossible to anticipate.

We argue that such algorithmic biases are new kinds of *bugs*, specific to modern, data-driven applications, that programmers should proactively check for, debug, and fix with the same rigor as they apply to other security and privacy bugs. Such bugs can offend and even harm users, and cause programmers and businesses embarrassment, mistrust, and potentially loss of revenue. They may also be symptoms of a malfunction of a data-driven algorithm, such as a ML algorithm exhibiting poor accuracy for minority groups that are underrepresented in its training set [5].

We refer to such bugs generically as *unwarranted associations*. Understanding and identifying unwarranted associations is an important step towards holding automated decision-making entities *accountable* for unfair practices, thus also providing incentive for the adoption of corrective measures [1], [2], [6], [7].

The Unwarranted Associations Framework. In order to

*Work done while the first author was at EPFL.

reason about this new class of bugs, and propose preliminary means to fixing them, we first need an appropriate definitional framework in which to express them.

Substantial prior work exists on algorithmic fairness [8]–[20]. The general approach is to characterize as “fairness bugs” any sufficiently strong statistical dependency between an algorithm’s outputs (*e.g.*, prices or labels) and protected user groups (*e.g.*, those defined by race or gender). While the overall approach is sound, and partially underlies our own definitional framework, we find previously proposed definitions and metrics for fairness to be systematically lacking in at least one dimension of interest: (1) Applicability to a broad range of settings, applications or data types, with rigorous statistical assessments; (2) Efficient and scalable assessment of disparate effects on a granular level, such as to discover biases affecting specific user subpopulations; and (3) Inclusion of natural explanatory factors (*e.g.*, application or user “utility”) that account for apparent unfair effects.

To unify and rationalize the substantial list of prior work, and extend it to broader settings, we introduce the *unwarranted associations (UA) framework*. Informally, we define an *unwarranted association*, also called an *association bug*, as any strong association between the outputs of an algorithm and features defining a protected user group, where the association arises in a meaningful subset of users and has no explanatory factors. (We give more details below.)

The UA framework proposes a principled methodology for testing for—and possibly debugging—unwarranted associations in a series of five steps: (1) Data collection, and identification of user features and algorithmic outputs of interest; (2) Integration of *explanatory factors*, typically special user or application requirements that justify ostensible unwarranted associations; (3) Selection of appropriate *statistical metrics* for assessing the strength of unfair associations; (4) Testing for association bugs over semantically meaningful *user subpopulations*, while ensuring *statistical validity* and *interpretability* of results; and finally (5) All or part of steps 1-4 may be repeated for *debugging purposes*, *e.g.*, through the addition of explanatory factors.

The UA framework specifically addresses three pressing needs noted by Kroll et al. in their recent survey article on accountable algorithms [2]: (1) “Expanding the repertoire of definitions of group fairness that can be usefully applied in practice”; (2) “providing better exploratory and explanatory tools for comparing different notions of fairness” and (3) enabling “system designers to have a set of rules, standards, or best practices that explain what notions of fairness are best used in specific real-world applications.”

The FairTest Testing Toolkit. We instantiate our UA framework in *FairTest*, the first system that helps developers test for and to some extent debug unwarranted associations in data-driven applications. Designed for ease of use, FairTest aggregates significant unwarranted associations in an easily interpretable *bug report*, filtered and ranked by association strength, while accounting for known explanatory factors.

FairTest integrates a set of carefully chosen metrics, each suitable for testing for associations in particular applications and data types. FairTest further offers *debugging primitives*

that let a developer rule out potential confounders for observed unfair effects. Our experience with four data-driven applications suggests that FairTest is effective at revealing and diagnosing subtle association bugs, such as skewed error against older patients in a predictive healthcare system and offensive racial labeling in an image tagger.

To efficiently identify semantically meaningful subpopulations affected by association bugs (step 4 in the UA framework), FairTest uses a novel heuristic we call *association-guided tree construction*. Inspired by decision-tree classifiers, our algorithm recursively splits the user space into smaller subsets so as to maximize some metric of association between algorithm outputs and protected user attributes. Each step yields subpopulations of decreasing size and increasingly strong disparate effects. We experimentally show that our tree-based algorithm finds highly affected subpopulations by searching through up to 8 times fewer candidates than previous “unguided” approaches [9]–[11].

Contributions. Our work offers the following contributions:

- (1) *The UA framework*, and its principled methodology for discovering and analyzing association bugs in data-driven applications. Our framework uniquely combines: Multiple primitives and metrics with broad applicability, explanatory factors, and fine-grained testing over user subpopulations, with rigorous statistical assessments, even for multiple adaptively chosen experiments.
- (2) *The association-guided tree construction algorithm*, for efficiently finding meaningful and interpretable user subsets strongly affected by algorithmic bias.
- (3) *The design, implementation, and evaluation of FairTest*, a system that instantiates the UA framework, and the first testing and debugging tool for unwarranted associations. The source code and datasets used in our experiments are available at <https://github.com/columbia/fairtest>.
- (4) *Thorough experimentation with FairTest* across four data-driven applications, revealing the widespread occurrence of association bugs, as well as FairTest’s effectiveness in detecting them.

2. Preliminaries

Developers routinely test and debug their programs for functionality, security, and privacy bugs. Many supporting tools exist for these purposes. Unfortunately, unwarranted associations constitute a newly emerging class of bugs in modern, data-driven applications, and no good testing tools for them currently exist. Our aim is to demonstrate the importance of proactively testing for unwarranted associations in data-driven applications, and to provide appropriate definitional foundations and tools for investigating them.

2.1. Motivating Examples

We present three examples that underscore the diverse contexts in which unwarranted associations can arise, and illustrate the testing capabilities needed to detect them:

- *Google Photos*. Google’s recently released Photos application includes an ML-based image tagging sys-

tem. Users found that Photos produced offensive labels, tagging black people in photos as “gorillas”. Google promptly apologized for the bug, saying that “This is 100% not OK,” and promised to fix it [3].

- *Staples’ differential pricing scheme.* The office retailer Staples implemented a seemingly rational differential pricing scheme for online purchases: users located within about 20 miles of a rival store (e.g., Office Depot) were often offered discounted prices. The *Wall Street Journal* (WSJ) found that the pricing scheme had a negative disparate impact on low-income customers, with results varying widely between states. WSJ called the situation an “unintended side-effect” [4].
- *Healthcare prediction.* Machine learning algorithms may tend to favor (i.e., provide higher prediction accuracy to) statistically dominant groups, at the expense of various minority groups. Indeed, a classifier may exhibit high overall accuracy, while performing poorly for certain demographics. Disparity in classification accuracy among groups has a high potential for unfair treatment in automated decision-making pipelines [5].

To illustrate this issue, we used real-world data and a winning approach from the Heritage Health Competition [21] to build a model using past healthcare claims to predict a user’s number of hospital visits in a year. Using FairTest, we found that while the model has high accuracy overall, its errors are unevenly concentrated on elderly users, especially in subpopulations with certain pre-existing conditions. An insurance company that uses this algorithm to tune insurance premiums might involuntarily discriminate against these elderly patients.

These examples show the wide variety of unwarranted associations, and the importance of proactively testing for and remedying them before they can harm users. These settings also illustrate specific, significant advances needed to address such concerns. Google’s Photos case shows that unwarranted associations can be difficult to anticipate, for algorithms with large output spaces, and therefore that we need sound methodologies and tools for *Discovery* of such bugs. Staples’ mishap illustrates the need, given a suspected discriminatory effect, for methods that enable rigorous *Testing* of its presence, extent, and impact, across different meaningful user subpopulations. Finally, the healthcare prediction example shows that rather than yielding outcomes that are discriminatory in their content, an application can provide disparate levels of *utility* to certain subpopulations, such as uneven error rates [5]. This scenario highlights the necessity to account for differences in the utility perceived by users, when assessing unwarranted associations in algorithmic decision making. In particular, the healthcare example calls for principled *Error Profiling* of ML algorithms.

2.2. Candidate Approaches and Related Work

A number of intuitive approaches may appear to allow discovery or prevention of association bugs. We review potential solutions and highlight some common limitations.

Information Removal and Tracking. It is well known that explicit omission of protected user attributes (e.g., gender, race) from an application’s inputs is insufficient to avoid discrimination along these axes. Indeed, subtle associations between protected attributes and other program inputs (e.g., location) may result in indirect biases. For example, Staples’ algorithm did not explicitly ingest information about its customers’ incomes; yet location surfaced as an unanticipated *proxy* for this sensitive information. For the same reason, traditional information-flow tracking [22] is also insufficient to detect proxy-based discriminatory effects.

Algorithmic Fairness. The literature on *algorithmic fairness* has proposed valuable techniques for both *preventing* and *detecting* unfairness in data-driven applications. However, our review of them reveals a fragmented set of definitions and concepts lacking unifying definitional foundations, and little work on development and evaluation of usable tools for data-driven programmers. We have carefully studied 13 representative works on algorithmic fairness: [8]–[20]. We observe four deficiencies in the field’s current state:

- *No generic, broadly-applicable criteria or metrics:*

The range of fairness criteria proposed in the literature is large and fragmented. Most works [9]–[17] use ratio and difference metrics, introduced by Pedreschi et al. [8]. Other criteria include a-protection [9], [10], [16], e-fairness [17], statistical parity [19], [20] and mutual information [18]. All of these criteria measure some form of association of *program outputs* (e.g., prices) on *protected user attributes* (e.g., race, gender). Our review suggests that each of these metrics is relevant to specific situations and data types, yet none is universally applicable. Proposed mechanisms typically apply only to specific chosen metrics, and thus cover a limited range of cases. For instance, most prior works only consider binary protected features and outputs [8]–[17].

Our UA framework unifies this fragmented landscape, by proposing a principled mapping from scenarios and data types of interest to appropriate statistical metrics.

- *Limited consideration of fine-grained discrimination contexts:* Most prior works [12]–[15], [17], [18] consider algorithmic fairness solely at full user population level. Yet, effects in a population may differ from, and even contradict, those exhibited in smaller subsets, an effect known as Simpson’s paradox [23]. The 1973 Berkeley admissions [24] are a famous example: admission rates seemed to disfavor women, yet individual departments showed either no bias or a reverse bias. Some works [9]–[11] recognized the need to proactively search for biases in user subsets, but do so through exhaustive enumeration of meaningful population subsets, leading to a number of contexts that is exponential in the feature space [9], [10] or linear in the user space [11].

Fine-grained bias discovery is a core component of the UA framework, and we propose a novel algorithm, called association-guided tree construction, to efficiently search for user subpopulations exhibiting strong unfair effects.

- *Disregard for utility and plausible explanations for observed biases:* Biases observed in a population may disappear when accounting for specific differences between

groups of individuals. In many situations, apparent discrimination can be explained by inherent differences in *utility*, such as naturally opposing inclinations between protected groups, or fundamental application requirements that favor certain user demographics. Some prior works do account for genuine application requirements [9]–[11], [15], [19]. Yet all without exception explicitly define fairness as a difference in *outcome proportions* between protected groups, thus rejecting the possibility that users may receive varying levels of utility from the same algorithmic outcomes.

For instance, when profiling the error of an ML algorithm (one of our motivating examples), it is implicitly understood that an algorithm can be unfair, even if it provides all of its users with the same outcome, if this outcome happens to be the “correct” one for only a fraction of them.

- *Limited experimentation*: Prior works describe limited experimental results, usually over two or fewer datasets [9], [10], [12]–[16], [18], [19], and omit system design issues such as scalability and usability. We believe that extensive experience with many applications and datasets is crucial for developing a robust and flexible system, such as FairTest, that can address real-world algorithmic fairness issues.

Web Transparency. Our work also relates to the field of *web transparency* [7], [25]–[32]. While some works touch on discrimination and fairness (e.g., [7], [26], [27], [30]), their setting is different: These works rely on *controlled, randomized experiments* that *probe* a service with different inputs (generally fake user profiles) and observe the effects on outputs, so as to quantify Web services’ use of personal data to target, personalize, and tune prices. Detection of unwarranted associations, as in FairTest, requires making inferences from application behavior on *real user profiles*, which may contain attribute correlations that do not manifest in fabricated user profiles. For instance, a controlled experiment of Staples’ algorithm (*i.e.*, with fabricated profiles of low-income and high-income users) would not uncover associations between prices and location, unless correlations between location and income are explicitly accounted for.

Debugging of Machine Learning Applications. A number of methods exist in ML to understand data features and modeling errors [33]–[36]. These mechanisms are insufficient to identify and investigate unwarranted associations: It remains a challenge to systematically inspect the large number of combinations of features that may define meaningful subpopulations. Existing tools for deriving subpopulations (e.g., clustering algorithms [35]) are not guided by measures that define discriminatory effects, and may hence miss important effects.

2.3. Goals and Assumptions

Our Goals. In defining unwarranted associations and the UA framework (see §3), our goal is not only to unify and rationalize the various definitions and metrics introduced in prior works, but also to extend algorithmic accountability and fairness to a broader range of applications. For instance, by including the Discovery and Error Profiling primitives, consideration of utility, and support for adaptive debug-

ging, we aim to cover scenarios that have received little attention in the literature. In designing FairTest, our goal is to instantiate our proposed methodology in an efficient, broadly applicable, and easy-to-use system. FairTest should help developers test for—and to some extent debug—the types of bugs or unintended effects motivated in §2.1.

Assumptions and Limitations. FairTest is intended for use by honest developers, willing to find unwarranted associations in their applications, e.g., for accountability purposes. We do not consider detection of unfair effects in applications that intentionally seek to conceal them.

When debugging uncovered association bugs, we restrict ourselves to the identification of *confounding* factors that might explain (and maybe justify) a bug’s presence (e.g., in the Berkeley admissions [24], the department to which a person applied to was a confounder). While this is valuable, confounders are not the only possible “explanation” for association bugs. Other potential causes, such as insufficient training data for an ML algorithm, are out of scope. While our framework and FairTest offer methodologies and tools for association bug *detection* and to some extent *debugging, remediation* is not explicitly covered. In practice, with FairTest—as with other testing tools—fixes often become apparent once a developer understands a bug. In §5.3 we show cases where debugging yields clear remediation paths.

A key assumption in our framework is that a developer has at her disposal a set of protected attributes (e.g., gender or race) for her users. In some cases, this information may be available from user profiles. Alternatively, public datasets (e.g., census data) can be leveraged to test for unwarranted associations on attributes that a developer lacks (see §4.1). Finally, one could imagine a deployment where a trusted auditor (e.g., the EFF) collects protected attributes from a user population and runs FairTest on developers’ behalf.

3. The UA Framework

We now give a more detailed definition of unwarranted associations, inspired by legal and statistical practices. We then define a generic methodology for the discovery of such bugs in data-driven applications, and show how the investigative primitives motivated in §2.1 fit into this framework.

3.1. Unwarranted Associations

Unwarranted associations are, as noted above, a new type of bug that encompasses a broad class of unfair effects of data-driven applications. To be able to test for such a bug, we need to define it. In doing so, our goal is *not* to give a new mathematical specification of algorithmic discrimination; many characterizations are possible and useful, as we show. Instead, we propose a definition flexible enough to encompass and extend a number of existing definitions and metrics, yet underpin a principled framework for identifying and ultimately resolving unwarranted associations.

We define an *unwarranted association* as any **statistically significant association**, in a semantically meaningful **user subpopulation**, between a **protected attribute** and an

algorithmic output, where the association has no accompanying **explanatory factor**.

This definition is broadly applicable, and rooted in legal discrimination practice. For example, in U.S. law, a ratio of 4/5 for hiring rates of two groups is generally considered discriminatory, but lower effects may qualify if statistically significant (and higher effects may be ignored if statistically insignificant) [37]. Furthermore, the notion of explanatory factors reflects, e.g., the legal notion of business necessity, which stipulates that differential treatment of two groups may be acceptable if that difference can be shown to arise as a consequence of fundamental business needs. Whether factors that account for observed effects are deemed as acceptable or not remains foremost a matter of context and policy, which we do not attempt to objectify here.

Statistical association is a broad notion: it encompasses any relationship between two measured quantities that renders them statistically dependent [38]. This captures direct (causal) dependencies, indirect dependencies via proxies, and relationships between any types of variables. For each variable type, there are specific metrics known in the statistics literature that can be applied to measure the degree of dependency. For each metric, there are also known statistical tests to quantify the statistical significance of an association.

In the following, we propose a principled methodology for the discovery and analysis of association bugs, which formally specifies the various components that make up the above definition (*i.e.*, protected features and outputs, statistical associations, user subsets, explanatory factors).

3.2. The Methodology

Our methodology consists of five generic steps for discovering and investigating unwarranted associations. We detail these steps below, and highlight design considerations for the FairTest system that implements this methodology.

1) Data Collection and Pre-Processing. We consider applications that take inputs on users, such as location or age, and return user-specific outputs. To check an application for unwarranted associations, one begins by collecting a dataset containing attributes from application users, along with the outputs (or some derived quantity) for those users.

The output to be tested for associations is denoted O . The choice of O may reflect notions of *user utility*, a point we expand on below. To ensure broad applicability, we require that our methodology and FairTest remain agnostic to the *semantics* of the quantity O to be investigated. Following [8]–[10], we consider three types of user attributes:

- (1) *Protected attributes*, S , are discrimination-sensitive features (*e.g.*, race or gender) on which to look for associations with O . What constitutes a protected attribute is foremost a matter of context, policy and law [37].
- (2) *Contextual attributes*, X , are dimensions along which the user population can be split into semantically meaningful subpopulations (*e.g.*, location, profession, etc.), in which to search for associations O and S . These include user attributes that are knowingly used in an application and may also include protected attributes.

Metric	Description	When to Use
Binary Ratio & Difference [8]–[17]	compare probability of an output for groups	binary S, O ; often for <i>Testing</i>
Mutual Information (MI) [18]	dependence measure for discrete variables	categorical S, O ; often for <i>Testing</i>
Pearson Correlation (CORR)	linear dependence measure for scalar variables	scalar S, O ; often for <i>Error Profiling</i>
Regression	for labeled outputs, measure each label’s association	high dimensional O ; always for <i>Discovery</i>

TABLE 1. Association Metrics for the UA Framework.

- (3) *Explanatory attributes*, E , are user properties on which it is deemed acceptable to differentiate, even if that leads to apparent discrimination on protected attributes. Explanatory attributes play a critical role in expressing fundamental application requirements and in the debugging of previously found association bugs.

The distinction between protected, contextual and explanatory attributes ultimately depends on the social norms, policies, or legal guidelines governing the application at hand. For instance, a genuine explanatory attribute in one setting may not be acceptable in another. Therefore, the UA framework makes no attempts at objectifying how this distinction should be made. Rather, our aim is to propose principled techniques for analyzing associations, given an established categorization of available user attributes.

2) Integration of Explanatory Factors. Associations between protected features and application outputs may not always be indications of a bug, but rather natural consequences of utility requirements. As noted above, the outputs of interest O can (and should) be chosen so as to reflect the utility perceived by application users. On the one hand, all users may perceive similar utility from receiving a discount price or a positive hiring decision. On the other, a health prediction may only be useful to a user if it is accurate.

Some associations may also be deemed acceptable or necessary if they are induced by fundamental application requirements, *e.g.*, a hiring decision based on a person’s qualifications, which could be an inadvertent proxy for gender or age. Such explanatory factors can be expressed through explanatory attributes E on which the statistical association between S and O can be conditioned.

3) Selection of Appropriate Metrics. Given a chosen quantity O to be tested for associations against a protected attribute S , the next step is to select a statistical metric that is appropriate for measuring the strength or extent of the association between O and S in a given context.

As noted in §2.2, prior work has identified a number of metrics for measuring associations between certain types of quantities O and S . Table 1 shows a set of four representative metrics that we selected, two of which are new to the algorithmic fairness space. This small set of metrics appears sufficient to cover a vast variety of applications and data types, as we illustrate in §5. Organized based on S and

O types, the metrics can be classified in three categories:

- *Frequency Distribution Metrics*: Association between nominal S and O (with few possible values) can be represented as a *contingency table* that displays the variables’ frequency distribution. Prior work used such tables to define *ratio and difference metrics* for binary variables [8]–[17]. For $O = \{o_1, o_2\}$ and $S = \{s_1, s_2\}$, the *ratio metric* is $\Pr(o_1|s_1)/\Pr(o_1|s_2) - 1$, and the *difference metric* is $\Pr(o_1|s_1) - \Pr(o_1|s_2)$. These are difficult to extend to non-binary protected attributes [19], such as race.

In general cases, we measure association with *mutual information* (MI), given by $\sum_{o,s} \Pr(o,s) \ln \left(\frac{\Pr(o,s)}{\Pr(o)\Pr(s)} \right)$. MI is well-defined for arbitrary output and protected feature spaces, and thus more widely applicable than binary metrics. A natural *normalization* of MI (NMI) divides the measure by the minimum of the Shannon entropies of S and O . MI can be defined for continuous O and S but is more expensive to estimate, especially in high dimensions [39], [40]. In these cases, we instead opt for a different metric: correlation.

- *Correlation*: Pearson’s correlation measures the strength of a linear relationship between continuous O and S , which may exist even if O and S are non-linearly related; such linear relationships are typically robust and broadly interpretable [41]. Unlike MI (and some non-linear dependence measures [42], [43]), a finding of zero correlation does *not* imply that O and S are independent. However, because our aim is not to verify independence, and because we prefer interpretable findings of dependence, Pearson’s correlation is a natural fit in many applications. We can use correlation, for instance, to test for associations between a patient’s age and the error incurred by a healthcare predictor.

- *Regression*: High-dimensional output spaces occur in many use-cases, such as for applications that assign tags or labels to users, where it is not known *a priori* which specific tag/label to test for associations (see the Discovery investigation examples in §2.1). For these, we introduce a metric based on *regression*. We model the relationship between the protected attribute S and a large number of dependent output labels O with a regression model (logistic or linear). This yields a *regression coefficient* for each label, with which we can estimate that label’s association with S .

We detail our approach for binary S and applications that output t labels, each in $\mathcal{L} = \{l_1, \dots, l_d\}$ (O takes values in \mathcal{L}^t). Let b_1, \dots, b_d be indicator variables for each label ($b_i = 1 \Leftrightarrow l_i \in O$). We model the distribution of S given O by $\Pr[S = 1 | b_1, \dots, b_d] = \text{logistic}(\beta_0 + \sum_{i=1}^d \beta_i b_i)$, where the β_i are regression coefficients that measure the association between the labels and the protected feature [44].

If an *explanatory attribute* E is specified, the metrics are adapted to measure the *conditional association* of S and O , given E . Conditional metrics quantify the association between S and O that remains after controlling for E . Given some association metric $\mathcal{M}(S; O)$, the conditional association between S and O , given E , is $\mathbb{E}_E(\mathcal{M}(S; O) | E)$.

4) Granular and Statistically Rigorous Testing across User Subpopulations. As argued in §2.2, testing for associations solely in the full user population is insufficient,

as strong unfair effects may manifest only in specific user subsets. The core testing step in our framework consists in finding strong associations between protected features and outputs, in specific subpopulations of application users.

How should a “subpopulation” be defined? If we consider arbitrary subsets of users, it is trivial to group users so as to obtain maximal associations between a protected feature and output. To ensure that subpopulations are semantically meaningful and interpretable, we restrict ourselves to subsets obtained by partitioning users based on the value of a few contextual features X (e.g., the subset of users with ‘Job = Researcher’ and ‘Age ≤ 50 ’). Such “association contexts” enable an analysis of algorithmic fairness among users that share similar characteristics.

Efficiently searching for subpopulations with high associations is non-trivial: Any user subset may exhibit very different effects, a consequence of Simpson’s paradox. Where prior work has considered unguided, exhaustive search of potential contexts [9]–[11], FairTest uses a novel heuristic algorithm, called *association-guided tree construction*, to successively split users into smaller subsets with increasingly strong associations (§4.2.2). We experimentally compare our algorithm to prior work in §5.2.

When assessing associations in multiple subpopulations, care has to be taken to ensure the results remain statistically valid. To this end, FairTest applies well-known techniques, such as hypothesis testing, cross-validation and rigorously correcting for multiple statistical inferences (§4.2.3).

5) Adaptive Debugging. Following the previous steps, a developer may find evidence of association bugs affecting certain application users. The question remains on how such bugs may be *debugged* or *fixed*. Maybe the presence of the bug itself is sufficient to alert a developer to a specific issue in her application, allowing her to correct the flaw. Alternatively, a discovered bug could hint to an overlooked explanatory feature that accounts for its presence. In any case, a new investigation would have to be run (following the above methodology) to verify that the bug was effectively removed. As these debugging investigations are inherently informed by results from previous investigations, ensuring statistical validity of discoveries requires extra care. We discuss the system design implications for FairTest, as well as our preliminary solution to this problem in §4.2.4.

3.3. Three Core Investigation Primitives

We show how the core investigative primitives identified in §2.1 map onto the UA framework methodology:

Testing. This is the simplest investigative primitive in the UA framework. It tests for suspected associations (under a suitable metric such as MI or correlation) between application outputs O and protected user features S , possibly conditioned on an appropriate explanatory feature E .

Discovery. This primitive applies to scenarios in which a developer may not know the specific association bugs to test for. An example are applications that tag users with labels from a large space. None of the metrics proposed in prior work are appropriate for extracting biases between

protected features and a large set of candidate outputs. We thus propose a new regression-based metric that efficiently estimates the strength of the association between protected attributes and each output label and allows the identification of the most highly associated labels. In this sense, *Discovery* requires little *a priori* knowledge of what could constitute an association bug or what subpopulations it might affect.

Error Profiling. The relationship between Error Profiling and fairness is best seen under the lens of *user utility*: For many applications where fairness is commonly a concern (e.g., hiring), the outcomes O (the hiring decision) provide similar utility to all users (the candidates). Testing for unfairness can thus be reduced to testing for associations between O and S . In contrast, an output from a predictive classifier is often only as useful to a user as it is reasonably correct. Thus, a classifier may unfairly impact a given subpopulation by consistently providing users with *wrong* predictions.

Error Profiling can be seen as an instance of a more general form of Testing, where the quantity tested for associations is not the algorithm’s output, but each user’s utility. This primitive illustrates how the UA framework can be applied to a varied range of fairness issues, simply by adapting the semantics of the values tested for associations.

4. The FairTest Design

We describe the design of FairTest, a system that realizes the UA framework methodology. Fig.1 shows the FairTest architecture. To test her application for unwarranted associations, a developer supplies FairTest with a dataset consisting of attributes from application users, along with the outputs (or derived quantities) for those users. This corresponds to the first step of the UA methodology. FairTest analyzes this data and returns an *association report* that lists strong associations that FairTest has found between protected attributes S and outputs O . The developer inspects the report and determines which reported associations are real bugs that require fixing and which are admissible effects in her context. After giving a concrete example of association report, we detail FairTest’s architecture and illustrate how it realizes the remaining steps of the UA methodology.

4.1. Association Report Example

Suppose that Staples’ developers wished to inspect their pricing scheme’s impact on users before deploying it (e.g., to avoid bad publicity or for accountability purposes). To do so, they could use U.S. census statistics [45] to emulate users with realistic demographics visiting their website from various locations. They would run their location-based pricing scheme for those users and use FairTest’s *Testing investigation* to test for disparate impact on race or income.

We ran an investigation on a simulated pricing scheme akin to Staples’, which gives discounts to users located within 20 miles of competing OfficeDepot stores. Fig.2 shows part of FairTest’s bug report, generated by testing for suspected differential pricing based on income. The report lists statistically significant associations discovered between

Report of associations of $O=Price$ on $S_i=Income$: Assoc. metric: norm. mutual information (NMI).			
Global Population of size 494,436			
p-value=3.34e-10 ; NMI=[0.0001, 0.0005]			
Price	Income <\$50K	Income >=\$50K	Total
High	15301 (6%)	13867 (6%)	29168 (6%)
Low	234167 (94%)	231101 (94%)	465268 (94%)
Total	249468 (50%)	244968 (50%)	494436 (100%)
1. Subpopulation of size 23,532			
Context={State: CA, Race: White}			
p-value=2.31e-24 ; NMI=[0.0051, 0.0203]			
Price	Income <\$50K	Income >=\$50K	Total
High	606 (8%)	691 (4%)	1297 (6%)
Low	7116 (92%)	15119 (96%)	22235 (94%)
Total	7722 (33%)	15810 (67%)	23532 (100%)
2. Subpopulation of size 2,198			
Context={State: NY, Race: Black, Gender: Male}			
p-value=7.72e-05 ; NMI=[0.0040, 0.0975]			
Price	Income <\$50K	Income >=\$50K	Total
High	52 (4%)	8 (1%)	60 (3%)
Low	1201 (96%)	937 (99%)	2138 (97%)
Total	1253 (57%)	945 (43%)	2198 (100%)
...more entries (sorted by decreasing NMI)...			

Figure 2. **Sample Association Bug Report.** Testing investigation of disparate impact in Staples pricing simulation.

protected attribute ‘income’ and output ‘price’ in various subpopulations. Association strength is measured with *normalized mutual information* (NMI), one of the canonical measures of statistical dependence implemented in FairTest.

The report shows three populations: the global population is first, followed by the two subpopulations exhibiting the strongest disparities (highest NMI). Subpopulations are defined by user attributes: white people in California (first subpopulation) and black men in New York (second subpopulation). For each (sub)population, FairTest reports varied statistical information: a p-value (a measure of statistical significance, with a value below 5% generally considered statistically significant), a confidence interval for the NMI metric, and a *contingency table* that summarizes the frequency distribution of the outputs over the (sub)population.

Results can be read from these tables as follows. For the first subpopulation: among California’s white population (23,532 people in our test set), 7,722 (or 33%) have an income below \$50K. Out of these 7,222 users, 606 (or 8%) got the high price and the rest (92%) got a discount.

The report can be interpreted as follows: “At global U.S. population level, the disparate impact of the pricing scheme against lower-income users is nearly zero (NMI is close to zero for the global population; low-income and high-income users receive high prices in similar proportion, 6%). Yet, disparities are much stronger among white people in California (first subpopulation), where 8% of lower-income people get higher prices vs. only 4% of higher-income people. Strong disparities also exist for black men in New York (second subpopulation): 4% of lower-income black men get higher prices vs. 1% for higher-income black men.” As remediation, a programmer may decide to alter her pricing scheme, e.g., to disable price tuning in affected

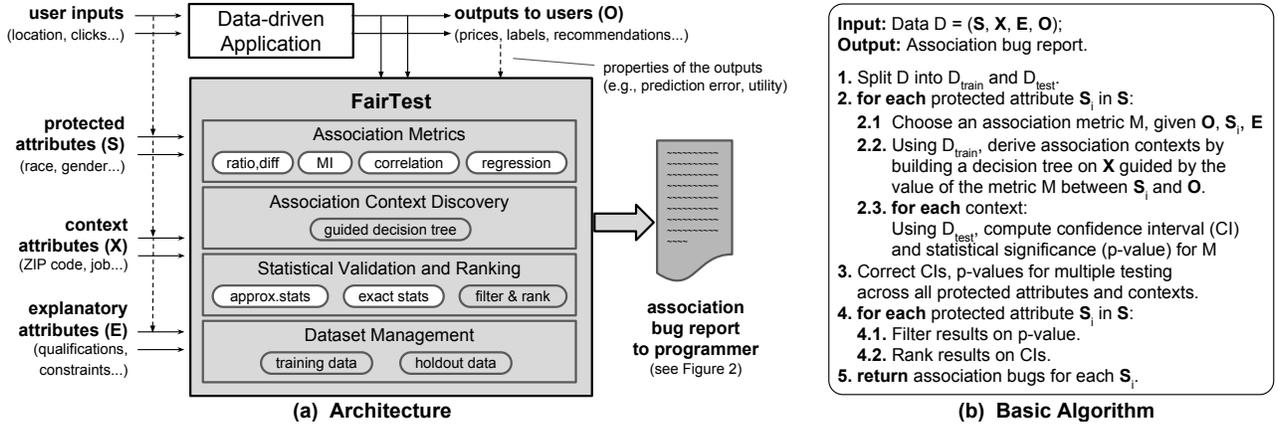


Figure 1. **FairTest Architecture.** (a) Grey boxes denote FairTest components. Rounded boxes denote specific mechanisms. White rounded boxes denote extensibility points; transparent rounded boxes denote core, generic mechanisms. (b) FairTest’s basic algorithm, realizing the UA framework methodology. S, X, E denote protected, context, and explanatory attributes, respectively; O denotes the output quantity of interest.

regions, or to take into account public statistics for various areas when deciding prices for people in those areas.

4.2. Architectural Components

Fig.1(a) shows FairTest’s inputs and architectural components, which map onto the UA framework methodology introduced previously. FairTest’s core consists of four components (*Association Metrics*, *Association Context Discovery*, *Statistical Validation and Ranking*, and *Dataset Management*), detailed in forthcoming sections. At a high level, FairTest generates reports using the algorithm in Fig.1(b), which makes use of these four components as follows.

- The dataset $D = \{(S, X, E, O)\}$ is first split into a *training set*, D_{train} and a *testing set*, D_{test} (Step 1).
- The *Association Metric* module, which incorporates the association metrics from Table 1, chooses a metric that will be used to quantify the dependency of the output on a protected attribute (Step 2.1).
- For each protected attribute S_i in S , the *Association Context Discovery* module uses D_{train} to split the user population into meaningful subpopulations based on contextual attributes, X (Step 2.2). Its goal is to maximize the association between S_i and O in the discovered subpopulations. For that, it uses our new association-guided decision tree construction heuristic (§4.2.2).
- For each discovered subpopulation, the *Statistical Validation and Ranking* module assesses the bug’s validity on D_{test} by applying an appropriate test statistic (Steps 2.3 and 3). To prioritize developers’ efforts, the module also filters and ranks bugs according to the strength of the measured associations (Steps 4-5).
- The *Dataset Management* module manages the testing sets to guarantee statistical validity of FairTest’s results across multiple investigations that follow Steps 1-5.

4.2.1. Association Metrics

This module incorporates a canonical (yet extensible) set of statistical association metrics (Table 1) and selects the appropriate metric for each investigation, as prescribed by the third step of the UA methodology. Unless specified by

Algorithm 1: Association-Guided Decision-Tree Construction We build increasingly specific contexts (subpopulations) with increasingly stronger associations, by recursively splitting the data D upon the user attribute X_i that maximizes the average association over derived contexts \mathbb{D} . Contexts are defined by predicates \mathcal{P} over attributes X (tree paths).

Params : `MIN_SIZE` // Minimum size of a context
`MAX_DEPTH` // Maximum tree depth
`Metric` // Association metric

```

Function findContexts( $D = \{S, X, E, O\}, \mathcal{P} = \emptyset$ )
  Create a subpopulation defined by predicates  $\mathcal{P}$ 
  if  $|D| < MIN\_SIZE$  or  $|\mathcal{P}| \geq MAX\_DEPTH$  then return
  for  $X_i \in X$  do
     $\mathbb{D} = \{D_1, D_2, \dots\} \leftarrow$  partition of  $D$  based on  $X_i$ 
    if  $\exists D_k \in \mathbb{D} : Metric(D_k) > Metric(D)$  then
      // Avg. association for this split
       $Score_i \leftarrow \sum_{D_k \in \mathbb{D}} Metric(D_k) / |\mathbb{D}|$ 
    else  $Score_i \leftarrow 0$ 
  if  $\forall i : Score_i = 0$  then
    return // No split yields higher assoc.
   $X_{\text{best}}, \mathbb{D}_{\text{best}} \leftarrow$  partition with highest score
  for  $D_k \in \mathbb{D}_{\text{best}}$  do
     $V \leftarrow$  values taken by  $X_{\text{best}}$  in  $D_k$ 
    findContexts( $D_k, \mathcal{P} \cup \{X_{\text{best}} \in V\}$ )

```

the user, a metric is chosen automatically based on the data types of protected attributes S and outputs O , the presence of explanatory attributes E , and the investigation type.

4.2.2. Association Context Discovery

This module efficiently “zooms into” a user population, to discover subpopulations (a.k.a. contexts) strongly affected by association bugs. In prior work, finding such bugs has required exhaustive enumeration of user subsets, leading to a number of contexts either exponential in the feature space [9], [10] or linear in the user space [11]. These methods raise two concerns: (1) They require making a large number of statistical inferences, thus providing only weak guarantees on the false discovery rate. (2) They sacrifice the ability to discover small subpopulations (e.g., a few hundred users), even if these exhibit the highest associations.

To efficiently identify strong hidden associations (the fourth step in the UA methodology), we develop a novel heuristic partitioning scheme, called *guided decision-tree construction*. In contrast to prior work, our method generates only a *constant* number of contexts, while actively searching for smaller, highly affected populations. Alg.1 shows our algorithm. Similarly to how decision-tree learning greedily optimizes some measure of target *homogeneity* (e.g., Gini impurity) [46], our algorithm actively maximizes some *association metric* between protected attributes and outputs.

The algorithm works by selecting a splitting rule, based on an attribute $X_i \in X$, so as to split the dataset D into subsets $\mathbb{D} = \{D_1, D_2, \dots\}$ with highest average association between S and O . The types of splits we consider are inspired by the C4.5 algorithm [46]: If X_i is categorical, the data is split into one subset per value of X_i . For continuous X_i , we consider binary splits of D based on some threshold t , i.e., $\mathbb{D} = \{D_1, D_2\}$ where D_1 contains all elements in D for which $X_i \leq t$. To choose the threshold t , we sort D according to the value of X_i and test each unique value taken by X_i in D as the potential threshold. We retain the value that maximizes the average association in D_1 and D_2 .

We say a split is *valid* if at least one of the resulting sub-contexts D_i exhibits a higher association than the one measured over the current population D . We ultimately choose the *valid* split (if one exists) for which the average association between S and O is maximal. We then recursively apply this process on each subset derived from the highest scoring split. The advantages of this approach are that it: (1) permits use of *any* association metric; (2) produces simply-defined and interpretable subpopulations; and (3) aggressively searches for subpopulations with strong associations using just scalable/distributable computations [47].

We additionally employ well-known techniques for preventing this tree construction from overfitting the training data [46], such as bounding the tree’s depth and pruning very small subpopulations (< 100 members).

4.2.3. Statistical Validation and Ranking

After discovering contexts that exhibit potential association bugs, we validate and rank them before reporting them to developers. Validation is needed because the contexts were explicitly built so as to maximize associations over a finite user sample (D_{train}). We validate bugs on an independent test sample (D_{test}). We suggest setting $|D_{\text{train}}| = |D_{\text{test}}|$, to maximize the likelihood that effects discovered during training have similar magnitude over the testing set.

We use distinct notions of significance for bug validation and ranking. For validation, we use statistical significance based on *hypothesis testing*: a bug is significant if its manifestation in the test set is unlikely under the “null hypothesis” (i.e., the association between S and O is null). This is quantified by the *p-value* for a test. For ranking, we use *effect size*, i.e., the actual value of the association metric, estimated by means of a *confidence interval* (CI). FairTest incorporates statistical methods for computing p-values and CIs for all metrics in §4.2.1; for small samples, we use generic *permutation tests* [48] and bootstraps [49] instead of

```

class DataSource(D, budget) # hold out 1 test set per budget unit

class Investigation(DataSource,S,X,E,O,M={}) # base investigation class
# M stands for association metrics
class Testing(DS,S,X,E,O,M) # Investigation subclass for testing
class Discovery(DS,S,X,E,O,M,top_k) # Subclass for discovery
# Takes in number of outputs to consider in each context.
class ErrorProfiling(DS,S,X,E,O,M,truth) # Subclass for error profile
# Takes ground truth for a predictive output

# Derive putative contexts for one/more investigations
train(Investigations,maxDepth=5,minLeafSize=100)
test(Investigations,conf=0.95) # Test and correct associations
report(Investigations,conf=0.95,outDir) # Filter, rank, save reports

class Metric # Abstract class for association metrics
Metric.computeStats(data, conf) # Calculate p-value and CI

```

Figure 3. FairTest API. Data holdout (1st), investigation types (2nd), methods to run investigations (3rd), API to implement for metrics (4th).

approximations. We apply Holm-Bonferroni corrections [50] to the p-values and CIs to ensure their simultaneous validity.

A bug report is generated as follows: We filter out contexts with corrected p-values > 0.05 (the choice of significance level α is a matter of policy. The default level of 5% is configurable). We rank the remaining contexts by the lower bounds of their corrected effect-size CIs. We only include a context (e.g., white males in NY) if it exhibits a stronger effect than the larger populations (e.g., males in NY) that contain it. The report prioritizes debugging efforts by (1) listing only statistically significant associations and (2) first displaying the most strongly affected subpopulations.

4.2.4. Dataset Management

To debug an association bug, developers may run multiple investigations, each informed by previous ones: A bug is detected in a first investigation, after which a series of other analyses (e.g., with explanatory attributes) are run to narrow down its source. Yet, while the test set D_{test} is independent from hypotheses formed in the first investigation, it is *not* independent from hypotheses formed in subsequent adaptively-chosen ones. Such adaptive debugging is the challenge raised by the final step of the UA methodology.

The system-design implications are significant: FairTest cannot be a stateless library; it must manage the dataset it is given across multiple investigations. We achieve this by having developers specify a budget B (number of adaptive investigations they plan to run) upfront, when they supply a dataset; FairTest then splits the dataset into B testing sets, each used for a single investigation. Only B investigations are allowed on the same dataset. The *DataSource* abstraction in Fig.3 implements this functionality. More efficient approaches have been proposed [51], but they all lead to similar systems implications, restrictions, and interfaces.

4.3. Testing and Debugging Primitives

To test and debug unwarranted associations, FairTest provides the three core investigative primitives defined in §3.3. §5.3 shows how these primitives can be combined together to support quite powerful analyses of unwarranted associations, from discovery to diagnosis and remediation.

4.3.1. Investigation API

Fig. 3 shows FairTest’s API. A core *Investigation* class is subclassed by three specific investigation types. To run these, a developer first collects user attributes and application outputs. This data is enclosed in a *DataSource*, which holds out test sets for successive experiments (see §4.2.4).

Testing. This investigation type was used in §4.1 to test for disparate impact in Staples’ pricing scheme. A developer provides a dataset $D = \{(S, X, E, O)\}$, and FairTest first finds contexts with potential associations between S and O (possibly conditioned on E), using the association-guided tree construction. The developer calls the *train* method to initiate context discovery, optionally tuning the size and complexity of the resulting contexts. She calls *test* and *report* to validate discovered bugs and produce reports.

Discovery. This lets developers search for associations over a large number of output labels, using the regression metric from §3.2 to efficiently estimate the strength of the association between protected attributes and each output label. We then select the *top_k* labels that exhibit the strongest associations, and test each label individually using an appropriately chosen metric. These regressions are executed at every step of the guided decision-tree recursion. In the end, *Discovery* finds subpopulations in which a particular label is disproportionately associated with a protected group.

Error Profiling. This searches for populations strongly affected by algorithmic mistakes. This may help improve both the algorithm’s accuracy and fairness, as we show for our healthcare predictor (§5.3.1). *Error Profiling* takes as inputs algorithm predictions and ground truth, and computes a suitable measure of prediction error, to be tested for associations. Note that in our implementation, Error Profiling computes a new outputs O' representing the prediction error and then simply calls the Testing procedure over O' .

4.3.2. Explanatory Attributes

Following the use-cases detailed in §3.2, FairTest lets developers use explanatory attributes in two ways: (1) After an initial investigation that reveals apparent unfair effects, she may *debug* these associations by specifying explanatory attributes that she believes are responsible (*i.e.*, confounders) for the observed behavior. FairTest then recomputes conditional association metrics over the *same* contexts discovered in the first investigation. (2) The developer can define user properties E that are necessary for an application. FairTest then explicitly avoids deriving associations that are accounted for by these attributes, by measuring the dependence of protected attributes S and outputs O conditioned on E .

To illustrate the first use-case of explanatory attributes, we examine the Berkeley graduate admissions dataset, which contains admission decisions and gender for 4,425 applicants [24]. As mentioned in §2.2, this data exhibits a paradoxical effect: at full university level, admissions appear to disfavor women, yet this bias is not reflected in any department. We show how an analyst could use FairTest to measure gender-disparities in admission rates,

Report of assoc. of O=Admitted on S _i =Gender, conditioned on attribute E=Department:			
Global Population of size 2,213			
p-value=7.98e-01 ; COND-DIFF=[-0.0382, 0.1055]			
Admitted	Female	Male	Total
No	615 (68%)	680 (52%)	1295 (59%)
Yes	295 (32%)	623 (48%)	918 (41%)
Total	910 (41%)	1303 (59%)	2213 (100%)
* Department A: Population of size 490:			
p-value=4.34e-03 ; DIFF=[0.0649, 0.3464]			
Admitted	Female	Male	Total
No	9 (15%)	161 (37%)	170 (35%)
Yes	51 (85%)	269 (63%)	320 (65%)
Total	60 (12%)	430 (88%)	490 (100%)
* Department B: Population of size 279:			
p-value=1.00e+00 ; DIFF=[-0.4172, 0.3704]			
Admitted	Female	Male	Total
No	3 (30%)	93 (35%)	96 (34%)
Yes	7 (70%)	176 (65%)	183 (66%)
Total	10 (4%)	269 (96%)	279 (100%)
* ... Departments C-F, with high p-values ...			

Figure 4. Disparate Admission Rates in the Berkeley Dataset. Shows a Testing investigation with explanatory attribute $E = \text{Department}$. COND-DIFF is the binary difference metric (DIFF), conditioned on E .

while allowing that each department may have different gender demographics and admission rates.

The analyst defines ‘department’ as an *explanatory attribute*, to instruct FairTest to look for associations only among applicants of a same department. She then runs a *Testing* investigation. The report (Fig.4) clearly highlights the paradox: Over the full population, only 32% of female applicants are admitted versus 48% for males. Yet, the only department with a significant disparity in admission rates (department ‘A’) actually favors women! Incidentally, the difference in admission rates *conditioned* on an applicant’s department is not statistically significant (p-value of 0.798).

4.4. Summary

With this design, FairTest realizes the UA framework methodology as follows: **1)** By keeping our design aware of the different types of user features (S, X, E), yet agnostic to the semantics of these features and of the algorithm’s outcomes, we ensure FairTest’s compatibility with the data-collection stage proscribed by the UA framework. **2)** By providing debugging primitives based on the use of explanatory features as confounders, FairTest instantiates the second step of the UA methodology. **3)** By making our design extensible with multiple metrics, we ensure FairTest’s applicability to a broad range of applications and scenarios. **4)** FairTest’s novel association-guided tree construction algorithm enables efficient, scalable, and statistically rigorous detection of strongly affected subpopulations. **5)** By keeping state of the data used in successive debugging investigations, FairTest provides basic support for adaptive data analysis.

5. Evaluation

We implemented FairTest in Python, to be used either as a standalone library or as an online service. As a library,

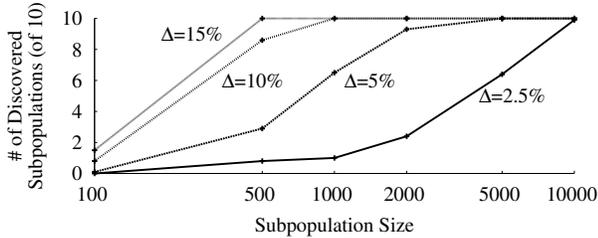


Figure 5. **FairTest Effectiveness with Affected Subpopulation Size and Effect Strength (Δ)**. Number of contexts discovered out of the ten we artificially inserted in 1M-user population. Average over 10 trials. Note that the x -axis is log-scale.

FairTest is designed to integrate with Pandas, SciPy’s data analysis library. Our online service prototype enables multiple users to asynchronously conduct Fairtest investigations. Users post Fairtest investigations through a web interface and access the respective bug reports once the experiments are completed. Our implementation is based on python job queues and scales beyond a single machine. Each Fairtest investigation is abstracted into a job and is dispatched for asynchronous execution into a poll of workers. Once a job is complete, the respective bug report is placed into a shared directory and is available online. We focus our evaluation on the standalone library, the more mature of our two prototypes. We aim to address three questions: **(Q1)** Is FairTest effective at detecting association bugs? **(Q2)** Is it fast enough to be practical? and **(Q3)** Is it useful to identify and to some extent debug association bugs in a variety of applications? We use seven workloads:

- *One tightly controlled microbenchmark*, which we use to evaluate FairTest’s bug detection abilities with a priori known ground truth for the associations.
- *Four data-driven applications* fed by public datasets: (1) a simulator of Staples’ pricing scheme (described in [4]) fed by U.S. census data; (2) a predictive healthcare app, based on a winning method and data from the Heritage Health Competition [21]; and (3) an image tagger based on Caffe [52], fed by ImageNet [53]. and (4) a movie recommender trained on MovieLens data [54].
- *Two social datasets* – the Adult Census dataset [55] and the 1973 Berkeley Admissions dataset [24].

Table 2 shows workload information: number of users/attributes, investigations we ran, and metrics we used.

5.1. Detection Effectiveness (Q1)

Microbenchmark. Inspired by the Staples case, we create a microbenchmark that lets us control the strength and span of association bugs. We use U.S. Census [45] data for gender, race, and income to generate $\approx 1M$ synthetic users. We start from a “fair” algorithm that randomly provides users with $\{0, 1\}$ -output, independent of income. We then plant disparities in certain subpopulations (determined by location and race), so that income level (high or low) implies a difference in output proportions of size 2Δ . E.g., for $\Delta = 10\%$, we might give output “1” to 60% of high-income users and 40% of low-income users, in the subpopulation of white users in California. For various subpopulation sizes and effect sizes,

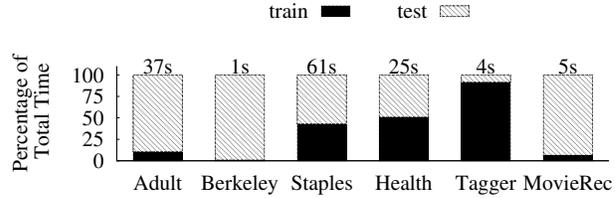


Figure 6. **FairTest Performance**. Total analysis time (labels above bars) broken down into training and testing times (bars).

we inject 10 randomly chosen discrimination contexts into our data and measure how many are discovered by FairTest.

Fig.5 shows FairTest’s discovery rate as we increase population size and Δ . FairTest reliably detects strong disparities that affect a few hundred users, as well as effects as low as 2.5% in larger contexts. Low effects in small contexts often go undetected due to the limited statistical power. This issue is not unique to FairTest, but rather an inherent limitation of any statistical inference procedure in the absence of sufficient data. In all cases, FairTest made zero false discoveries (finding a disparity that we did not introduce). Statistical testing lets us tightly control the false discovery rate: at a confidence level of 95%, we expect at most 5% false discoveries.

Real-World Apps and Datasets. Table 2 reports the number of association contexts found by FairTest in each application. We show the number of potential bugs found by the guided decision-tree mechanism, the number of associations that are statistically significant after correcting for multiple testing, and the number of bugs reported to the developer (recall that we only report a context if it exhibits a higher unfair effect than the larger subpopulations that contain it). The size of the smallest reported context is also shown (the largest reported context is the full test set). We do not have ground truth for these workloads, but our experience inspecting the reports (detailed in §5.3) suggests that FairTest detects discrimination contexts of a variety of sizes, which appear accurate and revelatory for an investigator.

Results for the predictive healthcare application are for an experiment with a follow-up debugging investigation (see §5.3.1). FairTest thus splits the dataset in three: a train set and two test sets. We found that FairTest would have reported the same bugs, had we used all the data for a single investigation (*i.e.*, with no debugging). We further analyzed the effect of the debugging budget B on the number of discovered bugs for the Staples application. For budgets B of 2 and 3, we discover 168 and 125 contexts, respectively; of these, we report 15 and 13 contexts, respectively. In both cases, the most affected subpopulation is the same as the one found for a budget $B = 1$. Thus, for this application, FairTest can allow at least one or two follow-up analyses, while preserving the main results reported to developers.

5.2. Performance (Q2)

Timing. Although its building blocks (decision trees, statistical tests) admit efficient and scalable implementations, our prototype does not incorporate all available optimizations. Still, FairTest appears fast enough for practical use. Fig.6

shows the analysis time for each application (top numbers), broken down into: (1) the time spent on training to form association hypotheses, and (2) the time spent on testing and correcting these hypotheses. On a commodity laptop (4-core Intel CPU @1.7GHz, 8GB RAM), the total execution time ranges from 1-5 seconds for the smallest datasets to 60 seconds for the largest (Staples, with 1M users). For small datasets (Adult, Berkeley, Movies) we often use bootstraps and permutation tests to compute CIs and p-values in small contexts (≤ 1000 users); these are expensive and subsume the training cost. For datasets that yield larger contexts, we use faster, approximate methods, making the testing phase fast and the training phase proportionally more expensive.

Subpopulation Discovery. To discover subpopulations with high associations, Ruggieri et al. [9], [10] use a simple data mining approach: They list all *frequent itemsets* over contextual features, and consider each itemset as a potential discrimination context. As this enumeration of contexts is not informed by any association metric, it can be viewed as a brute-force analog to the guided tree mechanism used in FairTest. Their approach suffers from two main caveats. First, it requires that all features be discretized. Second, it leads to an exponential increase in the number of contexts to consider, unless subpopulations of small size are neglected.

For a fair comparison between our guided-tree mechanism and the data mining approach, we analyze their performance when processing the *same* number of contexts. We proceed as follows: We run FairTest, and record the number N_{train} of contexts examined in the training phase (*i.e.*, the number of metric computations in Alg.1) and the number N_{test} retained for testing. FairTest considers contexts with at least `MIN_SIZE` users (*e.g.*, 100 or 500), and defined by at most `MAX_DEPTH` features (*e.g.*, 5). For the data-mining approach, we list all itemsets with support at least `MIN_SIZE` and at most `MAX_DEPTH` items. We choose N_{train} such itemsets *at random*, compute the association metric in those contexts over the training set, and retain the N_{test} contexts with highest effect for testing. We compare effect-sizes of the associations found by FairTest, to the ones found with the “brute-force” approach, over the testing set. Limiting the unguided approach to only N_{train} random contexts may appear arbitrary. We thus further compare the two methods for larger numbers of considered contexts, by simply repeating the above process with fresh random choices. Both FairTest and the unguided approach then consider the same total number $\gg N_{\text{train}}$ of potential contexts.

We use the Adult Census data, with protected feature “Gender”, target “Income”, and the binary difference metric. We add random permutations of the remaining features to obtain a dataset with either 15 or 40 contextual features (to see how the two methods scale with the number of features). The number of itemsets with at least 100 users, and defined by at most 5 features, grows from 84,000 (for 15 features) to over 200 million for 40 features. For efficiency reasons, we thus limit our search to subpopulations of at least 500 users, of which there are roughly 3 million (for 40 features).

In Fig.7, we plot the average association (as per the

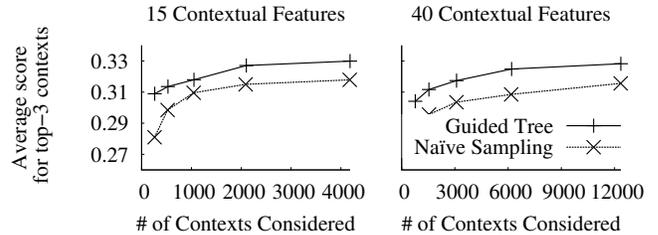


Figure 7. **Comparison between FairTest’s guided tree mechanism and naïve subpopulation sampling.** Shows the average of the 3 highest associations scores uncovered in the testing phase, for an increasing number of candidates considered in the training phase.

difference metric) for the 3 most affected subpopulations discovered by FairTest, compared to those found by the unguided approach. The process is repeated 2, 4, 8 and 16 times, resulting in both methods observing more potential contexts (x-axis) and discovering increasingly strong associations (y-axis). In summary, FairTest finds similarly high associations as with the itemset mining approach, by considering $4\times$ to $8\times$ fewer potential contexts. Furthermore, it is always more optimal to simply re-run FairTest than it is to search for more potential contexts in an unguided fashion.

5.3. Investigation Experience (Q3)

To assess FairTest’s usefulness, we searched for unwarranted associations in all real-world applications in Table 2. The Staples scheme was described in §4.1. The others are detailed below. Our experience shows that FairTest: (1) discovers meaningful bugs; (2) it lets developers include acceptable explanations for associations; and (3) it assists in debugging and even remedying discovered bugs.

More specifically, we find insightful association bugs in three of the four applications and showcase how to debug and remedy them in the case of the predictive healthcare application. For the movie recommender application, we show how apparent bugs discovered in a preliminary investigation can be explained by natural notions of user utility.

5.3.1. Predictive Healthcare

Our predictive health application uses methods and data from winners of the Heritage Health Prize Competition [21], [56]. The random-forest based algorithm uses past healthcare claims to predict a user’s number of hospital visits in the next year (predictions are for $\log(1 + \text{number of visits})$). The algorithm has low error overall (the average difference between the true and predicted number of visits is 0.42), but we want to study the error’s distribution among users.

Our study gives an end-to-end view of how FairTest can be used to detect and debug unwarranted associations, but also obtain hints for potential fixes. (1) We first discover an association bug: the application has much higher error rates for older than for younger users. (2) We investigate the bug’s source: the bias can be explained by lower prediction confidence for older people. (3) From there, we suggest potential fixes, such as only using high-confidence predictions.

Detection. We first use FairTest’s *Error Profiling* to examine associations between the algorithm’s prediction error and a

Application	Invest.	Users	Attr.	Metric(s)	Association Contexts			
					Discovered	Validated	Reported	Smallest Reported
Microbenchmark	T	988871	4	NMI	n/a	n/a	n/a	n/a
Staples Pricing	T	988871	4	NMI	224	100	21	211
Predictive Healthcare	EP	86359	128	CORR	33	33	2	91
Image Tagger	D,T	2648	1	REG,DIFF	1	1	1	1324
Movie Recommender	T	6040	3	CORR	15	10	7	511
Adult Census	T	48842	13	NMI	108	57	10	104
Berkeley Admission	T	4425	2	DIFF	1	0	1	2213

TABLE 2. **Workloads.** Investigations: *Discovery* (D), *Testing* (T), *ErrorProfiling* (EP). Metrics: normalized mutual information (NMI), correlation (CORR), binary difference (DIFF), regression (REG). For each application, we report the number of potential association contexts found by FairTest’s guided-tree construction, the number that were found to be statistically significant (p-value < 5%), and the number of reported bugs.

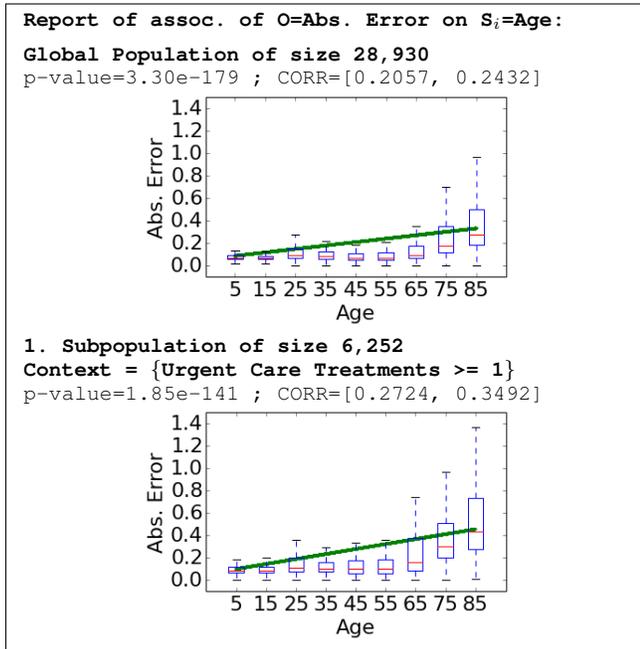


Figure 8. **Error Profile for Health Predictions.** Shows the global population and the subpopulation with highest effect (correlation). Plots display correlation between age and prediction error, for predictions of $\log(1 + \text{number of visits})$. For each age-decade, we display standard box plots (box from the 1st to 3rd quantile, line at median, whiskers at 1.5 IQRs). The straight green line depicts the best linear fit over the data.

user’s age (scalar quantities, hence we use correlation). The report (Fig.8) shows the error/age correlations for the full user population and one subpopulation with higher effect. We visualize correlation with plots instead of contingency tables. Globally, prediction error grows with age (correlation is positive and the data shows a clear positive linear trend). This effect is strongest for patients with prior urgent-care treatments. In that context, the average error for patients of age 61-99 is 1.07, compared to 0.33 for younger patients.

This finding is alarming, as such disparities could cause quantifiable harms if, e.g., the algorithm is used to adjust insurance premiums (one of the competition’s motivations [21]). Hence we wish to further investigate the *causes* of this accuracy loss for older patients, and get insights into how to fix this fairness (and accuracy) bug.

Debugging. We use FairTest’s debugging abilities (explanatory attributes) to verify a plausible cause for the observed bias: Higher errors for older users could be due to a higher

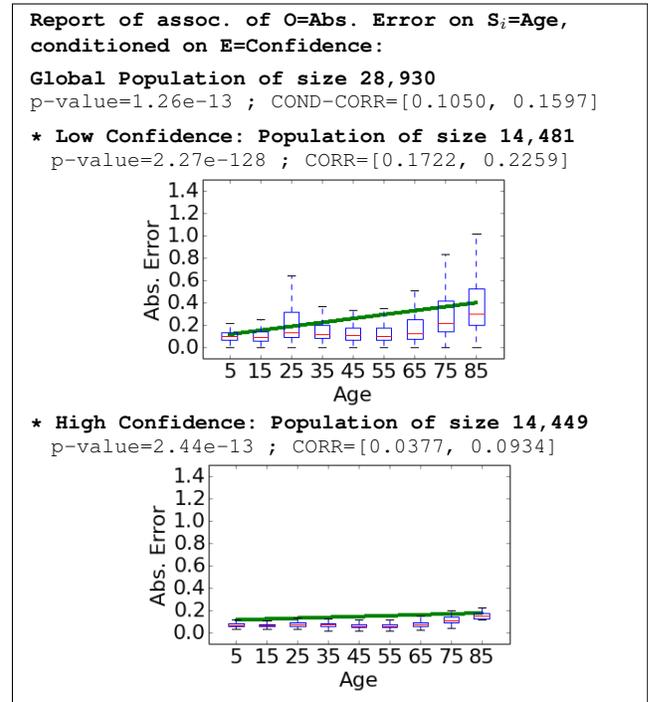


Figure 9. **Error Profile for Health Predictions using prediction confidence as an explanatory attribute.** Shows correlations between prediction error and user age, broken down by prediction confidence.

variance of the prediction target (the number of hospital visits). To estimate the variance in a patient’s target value, we train multiple predictors over random data subsets, and infer prediction intervals for our algorithm’s outputs [57]. The width of this interval is our estimate of the target’s variance. Low variance means high prediction confidence.

We run a new *Error Profiling*, with prediction confidence as an explanatory attribute (Fig.9). Conditioning on prediction confidence reduces the correlation in the full population. For low confidence predictions, the correlation of error on age remains positive and significant, but for high confidence users, the effect is close to zero. Similar results (omitted here) are obtained for users with an urgent-care history.

Remediation. These results imply an immediate remediation strategy: when using this algorithm to, say, tune insurance premiums, one should consider the predictions’ confidence. For example, one might decide to automatically tune premiums only for high-confidence predictions. This would

result in about half of the users in our dataset receiving customized premiums. One could also develop a scheme that weighs any price increase by prediction confidence. FairTest can then be used to test either of these approaches for disparate impact on the population.

5.3.2. Movie Recommender

This application highlights another interesting aspect about FairTest: not all associations it discovers are true bugs that need remediation. Some may simply be due to variations in *utility* perceived by different users for some outcomes.

We train a movie recommender using the ALS algorithm [58] and MovieLens data [54] (1M ratings provided by 6,040 users on a total of 3,900 movies). The ratings take values in [1, 5], and each user has rated at least 20 movies. The dataset includes user demographics (*e.g.*, age, gender) and movie metadata (*e.g.*, release date, genre). The system is trained to model the kinds of movies users generally like, and it is configured to recommend 50 new movies that a user has not yet seen, but is likely to rate highly.

Detection. We test for disparities in the “quality” of movies recommended to users. Let a movie’s *score* be the average rating given to that movie by all users. We view a movie’s score as an indicator of movie quality. We then use a *Testing* investigation to find associations between the average score of a user’s 50 new recommended movies, and that user’s age or gender. FairTest’s report (omitted for lack of space) shows that men and older users receive recommendations for “better” movies than women and younger users respectively (correlation of 0.02 for the gender disparity and of 0.10 for the age disparity). This suggests that the recommendation system might provide higher utility to certain user demographics, by recommending movies with higher scores.

Debugging. Although a low score indicates that a movie is generally disliked, it may be that certain users actually enjoy movies regarded as “bad” by others. Indeed, certain movie genres (*e.g.*, horror movies or comedies) are consistently rated lower than others (*e.g.*, documentaries or dramas). As we do not know the ratings that a user would give to the 50 movies recommended by our system, we cannot directly assess the “utility” provided to each user. However, we can test whether users that get recommendations for “bad” (low-scoring) movies are the ones that also watched “bad” movies in the past. In other words, we want to use FairTest’s debugging capabilities to investigate whether the discovered associations can be explained by some users’ natural inclinations towards low-rated movies.

For each user, we compute the average score of all movies rated by that user. We then add a new attribute that splits users into two groups: those that watched movies with below (respectively above) average scores in the past. When added to E , *i.e.*, treated as an *explanatory attribute*, we find that a user’s taste for low- or high-scoring movies accounts for the disparities found previously: After controlling for the scores of previously-rated movies, the difference in quality of recommended movies across gender is no longer significant. Slight age-disparities remain, but the disparate effect is much lower than before (correlation of 0.04).

Report of associations of O=Labels on S _i =Race:				
Global Population of size 1,324				
* Labels associated with Race=Black:				
Label	Black	White	DIFF	p-value
Cart	4%	0%	[0.014, 0.065]	3.31e-05
Drum	4%	0%	[0.010, 0.060]	3.83e-04
Helmet	8%	3%	[0.010, 0.089]	2.34e-03
Cattle	2%	0%	[0.0037, 0.0432]	4.73e-03
* Labels associated with Race=White:				
Label	Black	White	DIFF	p-value
Face Powder	1%	10%	[-0.134, -0.053]	5.60e-12
Maillot	4%	15%	[-0.159, -0.058]	3.46e-10
Person	96%	99%	[-0.056, -0.004]	6.06e-03
Lipstick	1%	4%	[-0.062, -0.003]	1.03e-02

Figure 10. **Racial Label Associations in the Image Tagger.** Shows partial report of a *Discovery* (top_k=35); the four most strongly associated labels (for the binary difference metric DIFF) are shown for each race.

Remediation. To conclude, we highlighted a situation in which apparent association bugs can be explained by simple and natural differences in user inclinations. We showed a successful use of FairTest’s debugging features, in dismissing associations found by a previous *Testing* investigation.

5.3.3. Image Tagger

We next show FairTest’s *Discovery* capability from the perspective of the developer of an image tagging system, who is willing to search for offensive labeling in racial groups. To illustrate, we inspect the labels produced by Caffe’s [52] implementation of R-CNN [59], a ready-to-use image tagger, when applied to photos of people from ImageNet [53]. The tagger was trained on images from ImageNet with 200 tags, including images of people. We tag 1,405 images of black people and 1,243 images of white people with 5 labels each, and run a *Discovery* to find the 35 (top_k) labels most strongly associated with each race.

Fig.10 shows FairTest’s report. It lists the labels most disparately applied to images of black people (first table) and white people (second table); we show only 4 (of 35) labels per race. A developer could inspect all top_k labels and judge which ones deserve further scrutiny. In Fig.10, the ‘cattle’ label might draw attention due to its potentially negative connotation; upon inspection, we find that none of the tagged images depict farm animals. Moreover, black people receive the ‘person’ tag less often, implying that the model is less accurate at detecting them. Further work is needed to understand these errors. While such analyses currently fall outside FairTest’s scope, this example shows that FairTest is effective at providing “leads” for investigations. It will also help test the effectiveness of a remediation.

5.3.4. Adult Income Census Dataset

A second use case for FairTest is studying discrimination in social datasets: We use the Adult dataset [55], which reports census data and income (under or over \$50K) for 48,842 U.S. citizens. Some disparate effects were found in prior algorithmic fairness works [11], [15]–[18], [20].

Fig.11 shows partial bug reports for *Testing* for income biases on race (top) and gender (bottom). We make three observations. First, FairTest confirms previously known bi-

Report of associations of O=Income on S _i =Race:				
Global Population of size 24,421				
p-value=1.39e-53 ; NMI=[0.0063, 0.0139]				
Income	Asian	Black	White	Total
<=50K	556 (73%)	2061 (88%)	15647 (75%)	18640 (76%)
>50K	206 (27%)	287 (12%)	5238 (25%)	5781 (24%)
Total	762 (3%)	2348 (10%)	20885 (86%)	24421 (100%)
1. Subpopulation of size 341				
Context={Age <= 42, Hours <= 55, Job: Fed-gov}				
p-value=3.24e-03 ; NMI=[0.0085, 0.1310]				
Income	Asian	Black	White	Total
<=50K	10 (71%)	62 (91%)	153 (63%)	239 (70%)
>50K	4 (29%)	6 (9%)	91 (37%)	102 (30%)
Total	14 (4%)	68 (20%)	244 (72%)	341 (100%)
2. Subpopulation of size 14,477				
Context={Age <= 42, Hours <= 55}				
p-value=7.50e-31 ; NMI=[0.0070, 0.0187]				
Income	Asian	Black	White	Total
<=50K	362 (79%)	1408 (93%)	10113 (83%)	12157 (84%)
>50K	97 (21%)	101 (7%)	2098 (17%)	2320 (16%)
Total	459 (3%)	1509 (10%)	12211 (84%)	14477 (100%)
Report of associations of O=Income on S _i =Gender:				
Global Population of size 24,421				
p-value=1.44e-178 ; NMI=[0.0381, 0.0540]				
Income	Female	Male	Total	
<=50K	7218 (89%)	11422 (70%)	18640 (76%)	
>50K	876 (11%)	4905 (30%)	5781 (24%)	
Total	8094 (33%)	16327 (67%)	24421 (100%)	
1. Subpopulation of size 1,371				
Context={9 <= Education <= 11, Age >= 47}				
p-value = 2.23e-35 ; NMI = [0.0529, 0.1442]				
Income	Female	Male	Total	
<=50K	423 (88%)	497 (56%)	920 (67%)	
>50K	57 (12%)	394 (44%)	451 (33%)	
Total	480 (35%)	891 (65%)	1371 (100%)	
2. Subpopulation of size 6,791				
Context={Education >= 12}				
p-value=3.71e-124 ; NMI=[0.0517, 0.0883]				
Income	Female	Male	Total	
<=50K	1594 (76%)	2156 (46%)	3750 (55%)	
>50K	492 (24%)	2549 (54%)	3041 (45%)	
Total	2086 (31%)	4705 (69%)	6791 (100%)	

Figure 11. Disparate Impact Reports on Race (top) and Gender (bottom) in the Adult Income Dataset. Shows the full population and two subpopulations with higher disparate effects. Additional races were omitted for clarity.

ases in the full dataset: 88% of blacks have <\$50K-income compared to 75% of whites and 73% of Asians. Similarly, 89% of women have low income compared to 70% of men.

Second, FairTest reveals new insights into these biases. For race (top), black people are strongly disfavored among people younger than 42 working fewer than 55 hours a week – especially for federal government employees. For gender (bottom), the groups where women are most disadvantaged are: (1) older people with 9-11 years of education and (2) (perhaps surprisingly) people with a higher education (≥ 12 years of education). We are unaware of any prior works in the algorithmic fairness area that have reported these particularly strong biases upon inspecting this dataset.

Third, as shown by the first context in Fig.11, FairTest

is capable of revealing even small contexts that show particularly strong disparate effects.

6. Conclusion

In an effort to rationalize and unify prior definitional foundations for algorithmic fairness, we have introduced *unwarranted associations*, a generic type of bug that characterizes disparate user treatment in data-driven applications. To reason about these bugs, we have proposed the *UA framework and methodology* which uniquely combine multiple primitives and metrics with broad applicability, fine-grained exploration of biases in user subsets, incorporation of natural utility notions, and rigorous statistical assessments.

As an instantiation of our UA framework, we have designed and presented *FairTest*, a tool that helps responsible developers to thoroughly check data-driven applications for unfair, discriminatory, or offensive user treatment. Designed for ease-of-use, FairTest enables scalable, statistically rigorous investigation of *unwarranted associations* between application outcomes and protected user groups. Our study of four applications shows the broad utility of FairTest’s three investigation types: *Discovery* of association bugs, *Testing* of suspected bugs, and *Error Profiling*.

We hope that our proposed tools will help developers answer recent calls from regulatory bodies, asking for data-driven algorithms to be tested and audited for fairness, in an effort to promote algorithmic accountability [1].

References

- [1] J. Angwin, “Make algorithms accountable,” <http://www.nytimes.com/2016/08/01/opinion/make-algorithms-accountable.html>, Aug 2016.
- [2] J. A. Kroll, J. Huey, S. Baroas, E. W. Felten, J. R. Reidenberg, D. G. Robinson, and H. Yu, “Accountable algorithms,” *University of Pennsylvania Law Review*, vol. 165, 2017, forthcoming.
- [3] J. Guynn, “Google photos labeled black people ‘gorillas’,” <http://www.usatoday.com/story/tech/2015/07/01/google-apologizes-after-photos-identify-black-people-as-gorillas/29567465/>, July 2015.
- [4] J. Valentino-Devries, J. Singer-Vine, and A. Soltani, “Websites vary prices, deals based on users’ information,” <http://www.wsj.com/articles/SB10001424127887323777204578189391813881534>, Dec 2012.
- [5] M. Hardt, “How big data is unfair. Understanding sources of unfairness in data driven decision making,” <https://medium.com/@mrtz/how-big-data-is-unfair-9aa544d739de>, Sep 2014.
- [6] F. Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information*. Harvard University Press, 2015.
- [7] A. Datta, S. Sen, and Y. Zick, “Algorithmic transparency via quantitative input influence,” in *S&P*. IEEE, 2016.
- [8] D. Pedreschi, S. Ruggieri, and F. Turini, “Discrimination-aware data mining,” in *KDD*. ACM, 2008, pp. 560–568.
- [9] S. Ruggieri, D. Pedreschi, and F. Turini, “Integrating induction and deduction for finding evidence of discrimination,” *Artificial Intelligence and Law*, vol. 18, no. 1, pp. 1–43, 2010.
- [10] —, “Data mining for discrimination discovery,” *TKDD*, vol. 4, no. 2, 2010.
- [11] B. T. Luong, S. Ruggieri, and F. Turini, “k-NN as an implementation of situation testing for discrimination discovery and prevention,” in *KDD*. ACM, 2011, pp. 502–510.

- [12] F. Kamiran and T. Calders, "Classifying without discriminating," in *IC4*. IEEE, 2009, pp. 1–6.
- [13] F. Kamiran, T. Calders, and M. Pechenizkiy, "Discrimination aware decision tree learning," in *ICDM*. IEEE, 2010, pp. 869–874.
- [14] T. Calders and S. Verwer, "Three naive Bayes approaches for discrimination-free classification," *Data Min. Knowl. Discov.*, vol. 21, no. 2, pp. 277–292, 2010.
- [15] I. Zliobaite, F. Kamiran, and T. Calders, "Handling conditional discrimination," in *ICDM*. IEEE, 2011, pp. 992–1001.
- [16] S. Hajian and J. Domingo-Ferrer, "A methodology for direct and indirect discrimination prevention in data mining," *TKDE*, vol. 25, no. 7, pp. 1445–1459, 2013.
- [17] M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian, "Certifying and removing disparate impact," in *KDD*. ACM, 2015, pp. 259–268.
- [18] T. Kamishima, S. Akaho, H. Asoh, and J. Sakuma, "Fairness-aware classifier with prejudice remover regularizer," in *ECML PKDD*. Springer, 2012, pp. 35–50.
- [19] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, "Fairness through awareness," in *ITCS*. ACM, 2012, pp. 214–226.
- [20] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork, "Learning fair representations," in *ICML*, 2013, pp. 325–333.
- [21] Heritage Provider Network, "Heritage Health Prize Competition," <http://www.heritagehealthprize.com/c/hhp>, 2012.
- [22] S. Sen, S. Guha, A. Datta, S. Rajamani, J. Tsai, and J. Wing, "Bootstrapping privacy compliance in big data systems," in *S&P*. IEEE, 2014, pp. 327–342.
- [23] E. H. Simpson, "The interpretation of interaction in contingency tables," *J. R. Stat. Soc. Series B*, pp. 238–241, 1951.
- [24] P. Bickel, E. Hammel, and J. O'Connell, "Sex bias in graduate admissions: Data from Berkeley," *Science*, vol. 187, no. 4175, pp. 398–404, 1975.
- [25] A. Hannak, P. Sapiezynski, A. M. Kakhki, B. Krishnamurthy, D. Lazer, A. Mislove, and C. Wilson, "Measuring personalization of web search," in *WWW*. IW3C2, May 2013.
- [26] A. Hannak, G. Soeller, D. Lazer, A. Mislove, and C. Wilson, "Measuring Price Discrimination and Steering on E-commerce Web Sites," in *IMC*. ACM, 2014.
- [27] T. Vissers, N. Nikiforakis, N. Bielova, and W. Joosen, "Crying Wolf? On the Price Discrimination of Online Airline Tickets," *HotPETS*, pp. 1–12, Jun. 2014.
- [28] P. Barford, I. Canadi, D. Krushevskaja, Q. Ma, and S. Muthukrishnan, "Adscape: Harvesting and Analyzing Online Display Ads," in *WWW*. IW3C2, Apr. 2014.
- [29] B. Liu, A. Sheth, U. Weinsberg, J. Chandrashekar, and R. Govindan, "AdReveal: improving transparency into online targeted advertising," in *HotNets-XII*. ACM, Nov. 2013.
- [30] A. Datta, M. C. Tschantz, and A. Datta, "Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination," in *PETS*, 2015.
- [31] M. Lecuyer, G. Ducoffe, F. Lan, A. Papancea, T. Petsios, R. Spahn, A. Chaintreau, and R. Geambasu, "XRay: Enhancing the Web's Transparency with Differential Correlation," in *USENIX Sec*, 2014.
- [32] M. Lecuyer, R. Spahn, Y. Spiliopoulos, A. Chaintreau, R. Geambasu, and D. Hsu, "Sunlight: fine-grained targeting detection at scale with statistical confidence," in *CCS*. ACM, 2015.
- [33] A. Ng, "Advice for applying machine learning," <http://cs229.stanford.edu/materials/ML-advice.pdf>, 2011.
- [34] —, "Machine learning system design," <https://d396qusza40orc.cloudfront.net/ml/docs/slides/Lecture11.pdf>, 2013.
- [35] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, 2nd ed. Springer, 2009.
- [36] M. Kabra, A. Robie, and K. Branson, "Understanding classifier errors by examining influential neighbors," in *CVPR*, 2015.
- [37] EEOC, "Information on impact (§ 1607.4), Uniform Guidelines on Employee Selection Procedure," 1978.
- [38] G. Upton and I. Cook, "A dictionary of statistics," 2008.
- [39] L. Paninski, "Estimation of entropy and mutual information," *Neural computation*, vol. 15, no. 6, pp. 1191–1253, 2003.
- [40] B. Poczos, L. Xiong, and J. Schneider, "Nonparametric divergence estimation with applications to machine learning on distributions," in *Uncertainty in Artificial Intelligence*, 2011.
- [41] J. L. Rodgers and W. A. Nicewander, "Thirteen ways to look at the correlation coefficient," *Am. Stat.*, vol. 42, no. 1, pp. 59–66, 1988.
- [42] G. J. Székely, M. L. Rizzo, and N. K. Bakirov, "Measuring and testing dependence by correlation of distances," *The Annals of Statistics*, vol. 35, no. 6, pp. 2769–2794, 2007. [Online]. Available: <http://www.jstor.org/stable/25464608>
- [43] A. Gretton, O. Bousquet, A. Smola, and B. Schölkopf, "Measuring statistical dependence with hilbert-schmidt norms," in *Algorithmic Learning Theory*, ser. Lecture Notes in Computer Science, S. Jain, H. Simon, and E. Tomita, Eds. Springer Berlin Heidelberg, 2005, vol. 3734, pp. 63–77. [Online]. Available: http://dx.doi.org/10.1007/11564089_7
- [44] D. W. Hosmer Jr and S. Lemeshow, *Applied logistic regression*. John Wiley & Sons, 2004.
- [45] U.S. Census Bureau, <http://www.census.gov/easystats/>, Sep 2015.
- [46] R. Quinlan, *C4.5: programs for machine learning*. Elsevier, 2014.
- [47] X. Meng, J. Bradley, B. Yavuz, E. Sparks, S. Venkataraman, D. Liu, J. Freeman, D. Tsai, M. Amde, S. Owen *et al.*, "Mllib: Machine learning in apache spark," *arXiv preprint arXiv:1505.06807*, 2015.
- [48] M. D. Ernst, "Permutation methods: A basis for exact inference," *Statistical Science*, vol. 19, no. 4, pp. 676–685, 2004.
- [49] B. Efron, "Bootstrap methods: Another look at the jackknife," *The Annals of Statistics*, vol. 7, no. 1, pp. 1–26, 1979.
- [50] S. Holm, "A simple sequentially rejective multiple test procedure," *Scandinavian Journal of Statistics*, vol. 6, no. 2, pp. 65–70, 1979.
- [51] C. Dwork, V. Feldman, M. Hardt, T. Pitassi, O. Reingold, and A. Roth, "The reusable holdout: Preserving validity in adaptive data analysis," *Science*, vol. 349, no. 6248, pp. 636–638, 2015.
- [52] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," *arXiv preprint arXiv:1408.5093*, 2014.
- [53] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *CVPR*. IEEE, 2009, pp. 248–255.
- [54] F. M. Harper and J. A. Konstan, "The MovieLens datasets: History and context," *TiiS*, vol. 5, no. 4, p. 19, 2016.
- [55] M. Lichman, "UCI machine learning repository," <http://archive.ics.uci.edu/ml>, 2013.
- [56] P. Brierley, D. Vogel, and R. Axelrod, "Round 1 milestone prize," <https://www.kaggle.com/wiki/HeritageMilestonePapers/file/Market%20Makers-Milestone1DescriptionV21.pdf>.
- [57] S. Wager, T. Hastie, and B. Efron, "Confidence intervals for random forests: The jackknife and the infinitesimal jackknife," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1625–1651, 2014.
- [58] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, vol. 42, no. 8, pp. 30–37, 2009.
- [59] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *CVPR*. IEEE, 2014.