



Nakao, Y., Strappelli, L., Stumpf, S. , Naseer, A., Regoli, D. and Del Gamba, G. (2022) Towards responsible AI: a design space exploration of human-centered artificial intelligence user interfaces to investigate fairness. *International Journal of Human–Computer Interaction*, 39(9), pp. 1762-1788. (doi: [10.1080/10447318.2022.2067936](https://doi.org/10.1080/10447318.2022.2067936))

The material cannot be used for any other purpose without further permission of the publisher and is for private use only.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<https://eprints.gla.ac.uk/269193/>

Deposited on 14 April 2022

Enlighten – Research publications by members of the University of
Glasgow
<http://eprints.gla.ac.uk>

Towards Responsible AI: A Design Space Exploration of Human-Centered Artificial Intelligence User Interfaces to Investigate Fairness

Yuri Nakao^a, Lorenzo Strappelli^b, Simone Stumpf^b, Aisha Naseer^c, Daniele Regoli^d and Giulia Del Gamba^d

^aArtificial Intelligence Laboratory, Fujitsu Laboratories Ltd., Kawasaki-city, Japan; ^bCentre for HCI Design, City, University of London, London, UK; ^cFujitsu Laboratories of Europe, Hayes, UK; ^dIntesa Sanpaolo S.p.A., Turin, Italy

ARTICLE HISTORY

Compiled April 14, 2022

ABSTRACT

With Artificial intelligence (AI) to aid or automate decision-making advancing rapidly, a particular concern is its fairness. In order to create reliable, safe and trustworthy systems through human-centred artificial intelligence (HCAI) design, recent efforts have produced user interfaces (UIs) for AI experts to investigate the fairness of AI models. In this work, we provide a design space exploration that supports not only data scientists but also domain experts to investigate AI fairness. Using loan applications as an example, we held a series of workshops with loan officers and data scientists to elicit their requirements. We instantiated these requirements into FairHIL, a UI to support human-in-the-loop fairness investigations, and describe how this UI could be generalized to other use cases. We evaluated FairHIL through a think-aloud user study. Our work contributes better designs to investigate an AI model's fairness—and move closer towards responsible AI.

KEYWORDS

Human-centred AI; fairness; user interfaces; visualizations; human-in-the-loop; data scientists; domain experts

1. Introduction

Artificial intelligence (AI) to aid or automate decision-making is advancing rapidly, and there are now many systems in use in jurisprudence, medicine and finance (Barocas, Hardt, & Narayanan, 2019; Barocas & Selbst, 2016; Chouldechova, 2017; Corbett-Davies, Pierson, Feller, Goel, & Huq, 2017). However, there have been increasing calls for a human-centred artificial intelligence (HCAI) design approach that is a part of a responsible AI development process to address concerns that automated decision-making is not very reliable, safe or trustworthy because they can ignore human values or contexts that change depending on the circumstances (Shneiderman, 2020, 2021). HCAI proposes that considerations are made for how to integrate user control into tasks, how to involve stakeholders in responsible AI development, and how to design user interfaces (UIs) that allow users to interact with AI systems.

In recent years, a number of UIs and technologies have been proposed that aim to involve stakeholders in investigating fairness of machine learning (ML) models by making them more transparent, such as AI Fairness 360 (Bellamy et al., 2018), the What-If Tool (Wexler et al., 2020), FairSight (Ahn & Lin, 2020), FairVis (Cabrera et al., 2019), SILVA (Yan, Gu, Lin, & Rzeszotarski, 2020), and Fairlearn (Bird et al., 2020). However, this strand of HCAI research has mostly focused on enabling data scientists or ML experts to inspect and assess their models, rather than involving other stakeholders such as domain experts or end-users. This is especially important for AI fairness, as stakeholder groups may vary in human values underlying fairness, information needs, practices, and technical abilities (Saxena et al., 2020; Veale, Van Kleek, & Binns, 2018), and these differences need to be reflected in the design of UIs that allow users to inspect fairness of AI (Lee, Kim, & Lizarondo, 2017; Veale et al., 2018).

In this work, we provide a design space exploration of UI components that allow both domain experts and data scientists to investigate fairness. We investigated this by: 1) conducting a series of workshops to better understand domain experts' and data scientists' practices and information needs for investigating fairness and produce a set of requirements; 2) designing FairHIL which instantiates a suite of 7 UI components based on requirements that can support these stakeholders in investigating fairness through a human-in-the-loop approach; and 3) evaluating FairHIL through a user study. We situated our work in the loan application domain, working directly with loan officers and data scientists employed by Intesa Sanpaolo, a bank who partnered with us in these studies.

Our aims were as follows:

- to understand how these different groups of users assess AI fairness in loan applications in terms of process, criteria, informational needs, and transparency, and develop requirements for UI design;
- to design a suite of UI components to support data scientists as well as loan officers in assessing fairness and finding potential fairness issues in datasets that underlie ML models and also ML models developed from these datasets. We show how these UI components could be generalized for other domains;
- to evaluate the usefulness and usability of these UIs through a user study involving data scientists and loan officers. We reflect on what these results mean for implementing these UIs into an AI development process.

We contribute to an evolving space of HCAI design with a focus on developing responsible AI. Our work concentrates specifically on fairness which is of great concern in increasingly automated decision-making. We advance the current state-of-the-art by:

- identifying the requirements of domain experts and data scientists for investigating fairness;
- providing a set of UI components which support the processes and practices of these stakeholder groups;
- clarifying the space of design options for developing UIs that allow fairness to be investigated in other domains;
- evaluating an instantiation of these UI components in the loan application domain.

Our paper is structured as follows: we first review the related work in ML fairness, paying particular attention to UIs that have been developed to interactively explore fairness in ML models. We then describe our first study with data scientists and loan

officers and show their UI requirements for investigating fairness. We then present in detail UI components that we developed in order to help data scientists and domain experts to investigate fairness and find potential fairness issues. We provide the findings of a second user study that evaluates these UI components with data scientists and loan officers. We conclude with a discussion of the implications and limitations of our work and future research that is warranted.

2. Related Work

2.1. Assessing AI Fairness

Fairness is often equated with justice which comprises two main principles, liberty and equality, in which all people should be entitled to basic rights and freedoms, and should have equal opportunities (Rawls, 1958). Scholars have attempted to differentiate and taxonomize notions of justice (Binns et al., 2018; Cook & Hegtvedt, 1983; Eckhoff, 1974; Lee et al., 2017), including equity (i.e. equal distribution), procedural (i.e. relating to a fair decision-making process), interactional (i.e. relating to people’s treatment), and informational (i.e. that a decision is well-justified) fairness.

A common way of perceiving fairness is at the group and individual level (Hegtvedt, 2005). For example, in US law, group fairness is embodied through the concepts of disparate treatment and disparate impact (Barocas & Selbst, 2016). The former is an intentional discriminatory treatment of groups or members of groups, and the latter, the process that causes a dissimilar impact on groups. These groups are often defined using ‘protected attributes’ which are enshrined in law. For example, the United Kingdom Equality Act 2010 states that it is illegal to discriminate against someone because of their age, disability, gender reassignment, marriage and civil partnership, pregnancy and maternity, race, religion or belief, sex and sexual orientation (Equality & Commission, 2020). Other countries have also moved to protect certain groups from unfair treatment. However, discrimination against a group can be less obvious. Even if protected attributes are not explicitly used in decision-making, discrimination might be indirect through the link of a protected attribute (e.g. race) with non-protected attributes (e.g. zip code). A different way of judging fairness is through individual fairness, which relates to the similar treatment or outcomes of two similar individuals (Dwork, Hardt, Pitassi, Reingold, & Zemel, 2012). Often, one difficulty that arises in individual fairness is defining how two people are similar. Additionally, the trade-off between individual and group fairness widely used in legal situations also has been considered as a difficult issue to solve (Dwork et al., 2012).

Unfairness can manifest in ML systems through biases. The ML development process presents numerous opportunities for bias to infiltrate, and biases can eventually play a role in making a decision (Barocas & Selbst, 2016; Caliskan, Bryson, & Narayanan, 2017; Dodge, Liao, Zhang, Bellamy, & Dugan, 2019; Friedman & Nissenbaum, 1996; Hajian, Bonchi, & Castillo, 2016; Olteanu, Castillo, Diaz, & Kiciman, 2019), for example, through the choice and characteristics of the dataset, feature engineering and selection, or choice of learning model. A well-known study underscores this point: COMPAS (Correctional Offender Management Profiling for Alternative Sanctions), a system used in the US justice system, was found to predict Black defendants as a much higher risk of violent recidivism than their White counterparts (Jeff Larson & Angwin, 2016). Because a lot of these ML systems are not accountable or transparent, there is a strong need to investigate and then potentially mitigate fairness issues in ML.

To measure fairness in AI systems, a number of fairness measures have been proposed. In the last few years over twenty different types have been identified (Narayanan, 2018; Verma & Rubin, 2018), for example:

- Individual fairness: similar outcomes for similar individuals (Dwork et al., 2012).
- Group fairness (statistical/demographic parity): people in both protected and unprotected groups have equal probability of having the positive outcome (Kamiran & Calders, 2009; Verma & Rubin, 2018).
- Subgroup fairness (intersectional bias): a combination of multiple sensitive features lead to an unfair result, which would have otherwise been considered fair when looking at them individually (Kearns, Neel, Roth, & Wu, 2018; Yang, Cisse, & Koyejo, 2020).
- Counterfactual fairness: a comparison between predictions concerning an individual and its “counterfactual self” with different values for protected attributes (Kusner, Loftus, Russell, & Silva, 2017).
- Equalized Odds: same error rates (false positive rate and false negative rate) in groups with different values of protected attributes (Hardt, Price, & Srebro, 2016).
- Fairness Through Unawareness: sensitive features are not explicitly employed to make decisions (Kusner et al., 2017).
- Predictive Parity: same precision for groups with different values of protected attributes (Verma & Rubin, 2018).

These fairness measures can then be used to develop ways to mitigate or remove bias or unfairness. There are a number of toolkits available to investigate and mitigate fairness issues, for example, IBM’s AI Fairness 360 (Bellamy et al., 2018) and Fairlearn (Bird et al., 2020). There are various approaches for bias mitigation, for example, manipulating or cleaning observed data, reweighting observations, suppressing features correlated with protected attributes, or trying to remove from the data all the information of protected attributes while trying to keep as much information as possible from other variables (Calmon, Wei, Vinzamuri, Ramamurthy, & Varshney, 2017; Kamiran & Calders, 2009; Zemel, Wu, Swersky, Pitassi, & Dwork, 2013). Bias mitigation can also include adopting training approaches such as adversarial training (Zhang, Lemoine, & Mitchell, 2018), regularization (Kamishima, Akaho, & Sakuma, 2011), or reductions approach (Agarwal, Beygelzimer, Dudík, Langford, & Wallach, 2018). Finally, it is also possible to tweak the final outcomes in order to satisfy the desired fairness metric (Barocas et al., 2019; Hardt et al., 2016).

Despite the exploration of fairness-aware ML, there has been scant agreement which measures to use (Verma & Rubin, 2018) because of trade-offs between them (Barocas et al., 2019; Barocas & Selbst, 2016; Chouldechova, 2017; Verma & Rubin, 2018). It has also been argued that fairness measures are poor proxies for identifying discrimination (Corbett-Davies & Goel, 2018; Narayanan, 2018; Veale et al., 2018), and that therefore statistical bias, based on fairness measures, and societal bias, based on human values, need to be differentiated (Mitchell, Potash, Barocas, D’Amour, & Lum, 2018; Narayanan, 2018).

Whether a decision or a system is judged as fair depends very much on human values embedded in individuals and society (Saxena et al., 2020), and, increasingly, there have been research efforts to investigate human fairness values empirically. For example, it has been found that while fairness outputs were perceived by participants as statistically fair, some participants were not comfortable with an AI system making high-stakes decisions without a human in the loop (Binns et al., 2018). In addition,

stakeholders can vary in what they perceive as fair (Narayanan, 2018). Many variables have been identified that influence fairness judgements in individuals. These include education, computer literacy (Wang, Harper, & Zhu, 2020), gender (Mallari et al., 2020; Wang et al., 2020), culture and region (Berman, Murphy-Berman, & Singh, 1985; Blake et al., 2015; Bolton, Keh, & Alba, 2010; Hofstede, 2011; Kim & Leung, 2007; Mattila & Choi, 2006), and environment and stakeholder role (Lee et al., 2017; Veale et al., 2018). Hence, it is very important to study how stakeholders investigate fairness in different domains, how stakeholder groups differ within a domain, and also how to better involve them through human-in-the-loop UIs, as we do in this work.

2.2. User Interfaces to Investigate Fairness

There has been a recent influx of fairness UIs that introduce a human-in-the-loop approach in the development of ML models. Many of these efforts are based on work in Explainable AI (XAI) to communicate models and decisions to users, for example, through model-agnostic tools such as LIME (Ribeiro, Singh, & Guestrin, 2016) and SHAP (Lundberg & Lee, 2017). LIME looks to explain model predictions on a local level, identifying the features that are driving the prediction (Ribeiro et al., 2016). SHAP can explain the contribution of features to the model (Lundberg & Lee, 2017). Coupled with data visualisation, XAI approaches are often used to identify potential bias, for example, Dodge et al. (Dodge et al., 2019) investigated global and local explanation styles in fairness assessments.

There are also now a number of toolkits being developed to support interactive exploration of fairness in datasets and ML models. We review five of the most well-known and recent ones in detail.

2.2.1. What-If Tool

The What-If Tool (Wexler et al., 2020) has been developed for data scientists to visualise fairness in ML (Figure 1 a). The UI is comprised of three screens: datapoint editor, performance and fairness, and features. The datapoint editor facilitates on-the-fly visualizations of the dataset or model. Users are easily able to bin, position, color and label data points. Features can be sliced to investigate group fairness and combined for subgroup fairness. Data points can be directly investigated and compared through counterfactual reasoning; by allowing users to select a data point, visualise its counterfactual and adjust its values, they are able to see ‘what could have happened’ and subsequently infer a reason for an outcome. Through color encoding, users are also able to see decision boundaries between the data points. The performance and fairness screen provides the user with a selection of six performance and fairness metrics. The features screen provides users with more descriptive statistics of the features in the dataset, such as their count, mean, median, and standard deviation.

2.2.2. FairVis

FairVis (Cabrera et al., 2019) (Figure 1 b) is a UI developed for data scientists to audit their models before deployment. In addition to comparing fairness between two or more groups, it also allows users to explore intersectional bias, i.e., identifying subgroup discrimination. The UI is made up of four interconnected panels. The first panel enables users to see all features (and their distributions), select them, or drill down further to select specific target classes. The second panel provides users with the

ability to apply a number of predefined metrics such as accuracy, precision and recall from a drop-down list. Strip plots for each metric are displayed below. By clicking or hovering on subgroups, the user is able to see, through color encoding, the subgroup’s position across the selected metrics and the four panels, therefore comparing and contrasting the subgroup across the dataset. To explore further, the user is able to click and pin their subgroups of interest and compare metric outputs in the third panel. The fourth and last panel facilitates further exploration through a carousel of suggested subgroups, which showcases features and their classes (sorted by chosen metric) as percentile bar charts.

2.2.3. *FairSight*

FairSight (Ahn & Lin, 2020) provides users familiar with ML with a UI that embeds a workflow to ensure fair decision-making. The FairSight UI (Figure 1 c) is made up of six panels; the first, ‘generator’, supports the user in setting up a model by defining sensitive features, features (including showing their distribution and correlation with the target), and target, and selecting a prediction algorithm from a pre-defined list. After running the model, group and instance-based performance and fairness metrics are shown in the Ranking View, Rankings, and Global Inspector panels. The Local Inspector and Feature Inspector panels facilitate the investigation of a specific data point and its nearest neighbor, and the features that might contribute to bias.

2.2.4. *Silva*

Silva (Yan et al., 2020) is a UI to support the AI development workflow. Crucially it is claimed that this UI is useful for both data scientists and ‘novices’ unfamiliar with fairness analysis. Its main feature is a causal graph generated by structural learning algorithms (Figure 1 d). The graph enables exploration of the features (nodes) and the relationships (edges) between them to identify potential bias in a classifier. The interface is made up of four panels; the first enables users to select and group features of interest and toggle their sensitiveness. The second panel shows the causal graph. Groups and their fairness metrics are shown in the third panel. The user is also able to access the dataset. The final panel displays the fairness metrics scores of each group for each of the machine learning algorithms.

2.2.5. *Fairness Elicitation Interface*

A UI, coupled with an interview protocol, has been proposed to elicit fairness notions from stakeholders who have no technical background in ML (Cheng et al., 2021) (Figure 1 e). The investigations for this UI were focused on predicting child maltreatment cases, and involved social workers and parents. The UI consists of a group view which shows differences between variable groups along fairness metrics and allows division into subgroup by a further variable; a case-by-case view which shows a pair of cases and their prediction and AI model features; and a similarity comparison view which shows a scatterplot that compares a selected case with all other cases, calculated through feature similarity using weighted Euclidean distance. While this UI is encouraging as it targets stakeholders other than data scientists, its aim is only to function as a supportive tool to elicit fairness notions instead of providing comprehensive support to investigate fairness of an AI model.

[Figure 1 near here]

2.2.6. Overview of Emerging Fairness UI Design Space

While every tool works in different ways to support a user to explore the fairness of an AI model, some common UI functionality is emerging designed to cover various functional and informational needs. For example, all tools allow the data to be grouped in some way. Many tools either display performance or fairness metrics, or both. Table 1 shows a mapping of functionality to the tools we have described in 2.2.

[Table 1 near here]

Our work explores a design space of UI components and associated functionality to investigate fairness. We aim to support a variety of stakeholder groups that differ in their technical ability with respect to AI, and in their experience with formalized notions of fairness. In order to ground our exploration, we started with an empirical investigation of stakeholder functional and informational needs to investigate fairness in the loan applications domain. This then led to a design of a tool based on these requirements. The coverage of the design space and the differences of our tool compared to the ones reviewed in section 2.2 are discussed in detail in section 6.2.

3. Study 1: Investigating stakeholder needs to assess fairness

To explore user needs, we worked with two stakeholder groups, loan officers and data scientists, employed in a partner bank, Intesa Sanpaolo based in Italy. Loan officers act as intermediaries between the bank and the customer, and are domain experts in this area but typically have little formal knowledge of AI or fairness metrics. Data scientists have experience in developing ML models for banking systems and processes and will typically have come across some fairness metrics.

We conducted two workshops with each stakeholder group to inform the exploration of this design space and uncover functional and informational needs. Due to Covid-19 restrictions, all research was conducted remotely using online meeting software.

Workshop 1 focused on how stakeholders investigate fairness, and what information that they consider important and need to assess fairness. It also exposed them to some initial potential UI components, based on our literature review, as design provocations. Workshop 1 was analyzed, and the insights informed the development of an initial prototype. This prototype was then discussed in Workshop 2, in order to obtain early feedback on more detailed designs while also being open to extending functionality that we had missed previously or needs that we had misunderstood in Workshop 1. Based on the requirements and feedback from the workshops, we produced a final set of UI components (section 4) for the evaluation study (section 5).

3.1. Participants

We recruited 6 loan officers (5 men, 1 woman, mean age = 36.5, stdev = 3.39) and 6 data scientists (3 men, 3 women, mean age = 29.7, stdev = 4.13). For further background details, see Table 2. All participants were over 18 years old and fluent in English. We excluded participants with any severe visual or hearing impairments, as it would have been difficult to make accommodations for these impairments in the study setup. Five of the 6 loan officers had previous experience with granting retail loans in bank branches across Italy, one of which had served as branch manager. The remaining participant had previously worked in Zagreb, Croatia, in another role but moved to Italy where he worked in retail and wealth management for a year. Four of

the 6 data scientists were very involved in considering AI fairness as part of their role. [Table 2 near here]

No incentives were offered for participation. This study was approved by the City, University of London Research Ethics Committee.

3.2. Setup

Separate workshops were held for data scientists and loan officers. All 6 data scientists were in one online workshop, and the 6 loan officers were in another. The sessions were run as semi-structured discussions, with the flexibility to explore relevant areas in more detail. This meant that data scientists and loan officers worked together in their respective groups and discussed their approaches and views for investigating fairness.

3.2.1. Workshop 1

The aim of workshop 1 was to conduct user research into how fairness was investigated by these data scientists and loan officers, their functional and informational requirements, and to expose them to initial conceptual designs. Workshop 1 took place in September 2020 for each group and lasted 2 hours. These 2 hours were structured into 3 parts: 1) a general discussion about fairness in AI, 2) a detailed investigation of the fairness of decisions in a given dataset, and 3) initial reactions to causal reasoning graphs.

In part 1 (approx. 30 minutes), participants had a general discussion about fairness, for example, what makes decisions in loan applications fair or unfair, and the role of ML in loan decision-making and fairness. This was to familiarize participants to the topic at hand, and to discover any general fairness values that they paid attention to.

In part 2 (approx. 1 hour), participants were asked to investigate fairness for a dataset we provided to them using their own preferred way. The dataset was developed by Intesa Sanpaolo, a partner bank in this project. The dataset contained 1000 anonymized but real loan applications and their decisions, with 26 features. These features include demographic information of the applicant (age, gender, nationality, etc), financial information (household income, insurance, etc), loan information (amount of loan requested, purpose of loan, loan duration, monthly payments, etc), as well as some information of their financial and banking history (years of service with the bank, etc). There were also some features that related to internal bank procedures, such as a money laundering check and a credit score developed by the bank. The dataset was sent ahead of the workshop so that participants had time to familiarize themselves with it (they could but did not have to have a look at the dataset before the workshop). We created a data visualization tool, providing the ability to slice the features and present them using various chart types such as histograms, scatter plots, bar graphs and a strip plot, in order to answer questions that participants had about the data on-the-fly. We screenshared this dataset and the visualisation tool so that any questions they had could be answered on the fly, in order to further investigate the fairness of the dataset. They were asked to discuss their individual approaches of how they would examine this dataset. As such, there was no set flow; it was all directed by the participants.

In part 3 (approx. 30 minutes), we aimed to understand how these users would react to some initial design provocations. Based on our literature review, casual graphs, as used in (Yan et al., 2020), seemed an intuitive and well-received yet much less common

way to explore fairness that could be useful for both data scientists and loan officers. We wanted to see the reactions of our participants to this. Hence, we introduced the concept of a causal graph to support fairness investigations.

Design provocations are often employed to explore the use and tensions of designs with users in order to speculate about future technologies. Consequently, the choice of these provocations can influence the future design space for the UI and visualizations. We only provided a simplified example of a small aspect of the UI (Figure 2) to gauge their initial reactions, so as not to bias stakeholders towards a particular design too early.

[Figure 2 near here]

After the workshop, all data was transcribed and de-identified. Thematic coding (Clarke, Braun, & Hayfield, 2015) identified themes around:

- Fairness in AI: Opinions on the limitations and benefits of AI and human involvement in loan application decision-making and fairness.
- Fairness Criteria: The criterion that participants use to judge and decide fairness within the dataset.
- Fairness Assessments: How participants make fairness assessments within the dataset. This includes what they look at, how they manipulate data to explore and troubleshoot, and the decision points they encounter along their journey.
- Causal Graph: Comments on the limitations, benefits and improvements of the causal graph presented during the workshop.

The insights from workshop 1 guided the design of an initial prototype which was discussed in Workshop 2.

3.2.2. Workshop 2

The second workshop aimed to gain early feedback on an initial prototype that was developed based on insights from workshop 1, to deepen our understanding of the informational and functional requirements needed to support stakeholders to explore fairness, and to extend and change our designs based on requirements that we had missed in workshop 1 or needs that we had misunderstood. These workshops took place with the same groups and participants in November 2020, two months after Workshop 1, lasting 2 hours each.

The initial prototype contained several different UI components (Figure 3): a system overview (A), a causal graph showing the relationships between five features (B), a dataset view (C), comparing nearest application with a different outcome (D). There was also an opportunity to see information on subgroups. Components within the UI worked in conjunction with one another, for example, selecting an application in component C highlighted this application in component D.

[Figure 3 near here]

We took participants through the UI to gather feedback. The prototype was screen-shared; participants did not have access to it directly, instead the researcher clicked through it and described what would happen in a fully developed prototype.

As in workshop 1, data was collected through screen and audio recordings, transcribed and analysed through thematic coding (Clarke et al., 2015). Themes relating to participants' reactions were developed around the screens and interface areas, and the clarity and usefulness of information and functionality offered in the prototype UI.

3.3. Insights into Stakeholder Requirements

3.3.1. Human-in-the-loop Fairness Process

There were some differences between data scientists and loan officers in perceptions about the role that ML should play and how human-in-the-loop ML fairness assessments need to be situated in the decision-making process. Both stakeholder groups stated that ML systems do not currently provide the ability to explain their decisions. Participants were aware of the impact biased ML systems could have on individuals and that human intervention is needed to ensure that outputs were fair. However, they also stated that it is difficult for users to control the ML system’s decision-making.

Loan officers saw ML systems as additional support for human decision-making. Rather than replacing human expertise, they thought that a ML system does not see the whole picture of the decision-making. Loan officer thought that the information that the ML system is based on might sometimes be subjective, or that important information is not available to the system at all. For example:

“I spent two years in a very small branch in very little town, the knowledge of the customer is something that is very valuable and maybe you as an employee could provide something extra - soft information - that would help the system to get the right decision in the end.” - WL01

In their view, this inability of the ML system to work with tacit information and take into account ‘grey areas’ may then lead to unfairness. This reflects earlier findings (Lee et al., 2017) that pointed out that some tasks are seen to require human skill.

Data scientists on the other hand leaned towards a process where the ML model was checked by humans before adoption.

What this means for user interfaces: The context in which human-in-the-loop fairness tools are used will matter for their design. Because data scientists would use these tools for model fairness assurance, we expect that they focus more on fairness measures of the model and relationships across attributes. Loan officers, on the other hand, want support for fair decision-making where humans have high control (Shneiderman, 2020), and hence they would focus more on information and functions that apply to individual applications. We do not suggest that these specific foci are exclusive, i.e. information about model fairness might also be useful to loan officers. What we take away from this though is that the UI should be adaptable to the needs of these user groups, or be equally usable by both.

3.3.2. Fairness Criteria and Metrics

Our investigation found that loan officers and data scientists use different fairness criteria.

Loan officers agreed that group fairness was important:

“I think there is a problem with adverse selection, often related to age, for example, younger people are disadvantaged, or nationality - foreign people are also penalised and another issue that I saw in the past is the absence of a credit history.” - WL04

However, they focused mainly on individual fairness, often wrestling with unintended bias creeping in. For example:

“I’d like to have the same treatment if I asked for money for a car and for healthcare. Why should I receive a different result if I can afford something?” - WL02

Perhaps because of the danger of bias based on the ‘worthiness’ of loan purpose, loan officers focused most frequently on *affordability*, by trying to understand how income and out-goings contributed to the application decisions. For example:

“We would have to first define what we intend for fairness, so fairness from an economic point of view from the sustainability of the numbers, or fairness from the point of view that this person is good enough to have the loan, maybe despite some short fall in the economic data, so I think this is important to set what we are looking at.” – WL01

In contrast, for data scientists group fairness played a significant role, while they did not mention individual fairness frequently. In fact, it seemed that they shied away from using individual fairness, possibly because it is more subjective, for example:

“How do we assess that two people are similar?” - WD06

Data scientists preferred to use metrics that captured group fairness objectively. There was also an expectation that a range of different metrics should be available in the UI, to cover different ways of assessing fairness.

“Okay, yeah, one thing is, if this should be a fairness explorer, there must be a lot of other metrics of fairness.” – WD02

What this means for user interfaces: Designs should support a range of popular fairness metrics, and support adding new ones if they are not already included. As previously discussed, there are quite a number of ways to measure fairness and it is sometimes unclear what measures to apply, therefore, it might also be necessary to describe these metrics and how they are derived in more detail.

While there are readily available metrics for capturing fairness that can be applied in AI systems and which would be useful for data scientists, there is currently a lack of ways to support the fairness assessments of loan officers. This could be overcome by allowing loan officers to specify ‘custom’ attributes, for example, to capture the notion of affordability. These attributes can then be used in ensuring fairer decision-making.

3.3.3. *Exploring Features and Relationships*

To make fairness assessments, both groups looked in detail at what they considered sensitive features. Gender and citizenship were discussed the most in both groups and both considered them unfair to be used in AI decision-making. However, age as a sensitive feature proved to be more contentious. Loan Officers in particular pointed out that young people often could not provide a long credit history, so suffered from a ‘cold start’ problem in loan applications. Similarly, older people might be deemed too much of a risk if their repayment terms were long.

Both groups requested information on the acceptance rates for loan applications related to feature values, particularly for the ones they considered sensitive. In general, investigating relationships between attributes was very important to them, through exploring the distribution of attribute values to the target values. This allowed them to compare whether there was a bias against certain applicant groups, and to determine if there was a cause for concern in terms of group fairness:

“Perfect, so it’s not so small the difference between 70/30 and 40/60, about 10% higher the rate of rejected proposal of non-national customers.” – WL05

Data scientists also wanted to explore subgroup fairness, in addition to simple group fairness. This meant combining a number of feature values and looking at the acceptance rates.

As already highlighted, participants identified that outcomes might be related to features and feature combinations, however the stakeholders were also interested in investigating relationships between any features to see if they were linked to each other, to see how bias might have crept in unintentionally:

“The first thing that I did was to check the correlations, mainly between the target variable, which is the result... the most correlated variable is of course, the risk levels. But for example, I didn’t find any correlation among the gender or the age or something like that... [it would be] probably nice to deeply analyse [the dataset] better and check also the correlations between the other variables to understand if there is something biased as we said at the beginning of the discussion... I also did some plots to check for example, the percentage of accepted or rejected for each variable.” – WD01

Loan officers often leveraged their domain knowledge to explore fairness. For example, they were already aware how credit risk could impact the outcome, and brought in the citizenship feature in their exploration:

“Yes, this is pretty similar, yes, it’s more approval in the national but it’s maybe usual because national people usually earn more money than people from abroad. Nothing unusual.” - WL03

In our workshops, we provided Sankey diagrams (Figure 4) as design provocations for communicating relationships between features values. While both stakeholder groups stated that visualizing relationships between features is useful, it became clear that this choice of visualization was too complex for participants to grasp, especially for continuous, binned feature values. Instead, they requested simplified visualizations or numerical information.

[Figure 4 near here]

We also showed causal relationships underlying the data, showing these in a simple node-edge graph (Figure 2), to allow an alternative and easy way for stakeholder to investigate relationships between all features. Both stakeholder groups saw the value in having causal relationships shown:

“I think the crucial point is at the beginning to decide what is fair and what is not in that specific domain. And in this respect, maybe the graph that you’re showing could help. Let’s take for example, gender. This graph is telling us that gender impacts the decision, at least through the credit risk level, then you have to ask about this impact; so, the impact that goes from gender through credit risk level to result is something that I’m willing to accept as a correlation from gender to result or not? So, you have to ask yourself, how is credit risk levels calculated?” – WD04

They also appreciated how these graphs could be integrated in the process of decision-making in loan applications and mitigating fairness issues:

“We cannot modify the graph? So for example, if we don’t want gender, we can remove it?” - WL06

What this means for user interfaces: Although protected features by law will need to be taken into account, there is considerable variability between the laws of countries in terms of which features count as sensitive. Further, other features might also be considered sensitive even though they are not protected by law. Hence, these features need to be selectable by the user, and then highlighted as especially important in the UI.

Outcomes related to features values and their ratios and distributions formed a core part of the assessments of the stakeholders. This information will need to be made avail-

able in an easily understandable format, perhaps through appropriate visualizations that can support comparisons between applicant groups, or ratio percentages.

Exploring relationships between features should also be extended to features that are not considered sensitive and features that do not directly influence the outcome. In addition, information about combination of features needs to be provided to investigate issues of subgroup fairness. This also allows intersectional issues of fairness to be explored, and would allow stakeholders to find biases that are perhaps less obvious.

Exposing causal relationships can add information that would otherwise not be available when looking at a dataset. For loan applications this is important, as the features used in calculations are often interrelated.

3.3.4. Exploring Individual Cases and Their Similarities

While a lot of the exploration of the data focused on investigating group fairness, we showed earlier (3.3.2) that loan officers also considered individual fairness. In the workshops we also investigated how to support this.

Feedback indicated that both stakeholder groups liked the ability to see and compare individual applications:

“For me it’s cool. Maybe the coolest part” – WD05.

However, the visualizations we presented proved to be problematic as participants did not understand how similarity was calculated and how the data was mapped to the axes. What they wanted to compare was the similarity of applications, especially if they led to different outcomes, and what attributes caused this.

What this means for user interfaces: Comparing similar applications was appreciated and should be supported in fairness exploration, in order to investigate individual fairness. These should provide information about similarity of cases, including their attributes, across decision boundaries.

3.3.5. Transparency and Explanations

Both groups wanted more transparency about how metrics or other information was derived or calculated. For example, they wanted information on how the causal graph was created, and how the ‘similarity’ score was calculated.

In particular, data scientists wanted a clear distinction between what belongs to the ML model and what belongs to the dataset from which the ML model is derived:

“The other thing that I noticed is I see a system overview accepted and rejected. I don’t know if it’s the model or the original, maybe I guess their model, but I’m not sure.” – WD06

It was obvious that loan officers struggled with ML concepts that were very familiar to data scientists. For example, they had difficulty with concepts such as target, precision, etc. and also how ML algorithms worked.

What this means for user interfaces: Much effort has been directed at providing explanations of ML systems and making them more transparent (Kulesza, Burnett, Wong, & Stumpf, 2015; Lim & Dey, 2009). Our findings echo this, showing that users require detailed descriptions so they can understand how information and visualizations have been generated, in order to interpret them effectively. Contextual help could aid in this. In addition, explanations need to be carefully targeted at different stakeholders to extend users’ understanding and mental models (Kulesza, Stumpf, Burnett, & Kwan, 2012).

3.3.6. Summary

Table 3 shows a list of requirements that we gathered through the workshops. Participants in our study supplied many informational and functional requirements that need to be supported to explore fairness. These requirements provide a much richer design space than the functionality provided by existing tools described in section 2.2.6. We found that data scientists and loan officers have slightly different priorities but overall their needs are surprisingly similar. Based on these requirements, we started to develop a prototype UI that would support data scientists and loan officers needs. We describe this UI next.

[Table 3 near here]

4. FairHIL, a UI to Support Investigating Fairness

Our design space exploration that started with a literature review and design workshops continued with the implementation of a UI of 7 components. We now present a detailed description of these UI components we developed; for each component we reference the corresponding requirements we captured from the design workshops (see Table 3). Please note that we focused on informational and functional requirements that allow stakeholders to investigate fairness. Requirements that focus on mitigating fairness by adjusting the model were considered out of scope for the purposes of our research. For each of these UI components we also show how it could be generalized for *any* dataset and model.

The UI is divided into a setup process UI and the main UI screen to explore the input dataset and the model, respectively. The setup comprises five steps for data scientists and four for loan officers, inspired by a wizard approach which is frequently implemented in UI setup processes (e.g. Fairlearn (Bird et al., 2020)) and involves loading the dataset, setting the target feature, selecting the prediction model, marking up sensitive features, and choosing fairness metrics, including the ability to define a 'custom metric' which leads to a function builder. (In order to simplify the interface, we did not include a choice of prediction model or metrics for loan officers.) We will focus on the *metrics choice* UI component in section 4.1.

[Figure 5 near here]

The main UI screen (Figure 5) allows users to rapidly switch between information about the dataset and the model through tabbing, and thus it supports them to easily compare the fairness of underlying loan decisions as well as the resulting predictions. The UIs to explore the dataset and the model are very similar, each with six main components that relate to each other. The *system overview* (Figure 5 A) provides overall information about the data and target, described in section 4.2. The main component supporting users to explore fairness is a *causal graph* (Figure 5 B) which shows features and causal relationships between them and the target. We discuss this UI component in detail in section 4.3. Detailed *information about features and relationships*, such as fairness metrics, distributions, and acceptance ratios in the causal graph, is displayed in a component underneath the graph (Figure 5 C). This UI component is described in section 4.4. Users are also able to see how intersections of feature values affected the acceptance ratios by exploring *feature combinations* (Figure 5 D), as we will describe in section 4.5. We also provide access to the underlying *dataset* (Figure 5 E) which allows filtering and more detailed explorations, see section 4.6. Finally, users are able to select an application in the dataset component and compare it with other

applications in the *compare applications* component (Figure 5 F), which we will detail in section 4.7.

Throughout we focus on providing transparency and explanations in UI components, to satisfy user requirements uncovered in section 3.3.5. We add detailed explanations, particularly targeting loan officers who struggled in the workshops with ML terminology and concepts. We concentrate on explaining the target and the algorithm, sensitive features, and the fairness metrics, and their calculation. Furthermore, we explain how the causal graph was developed and the information it showed in detail. We also provide contextual information on individual UI elements, for example when fairness metrics are shown in the UI components.

4.1. Metrics Choice

A main way of exploring fairness is through the use of fairness metrics (REQ1.3, REQ2.8). Providing users with a range of popular metrics is important as fairness criteria are both subjective and contextual to the problem at hand. To support this, data scientists are able to choose one or more of five metrics: statistical parity difference (SPD), equality of opportunity difference, average odds difference, disparate impact, and Theil index (see Figure 6). These metrics are also provided in IBM’s AI Fairness 360 (Bellamy et al., 2018), and we followed their metric definitions in our instantiation. For example, in our work, SPD was calculated as the difference between acceptance rate of the (grouped) values of an attribute. In cases where an attribute had continuous values, e.g. age, we split these values into ‘bins’ to represent groups, (e.g. 10 years, 20 years, etc.), calculated the acceptance rate for each bin, and then provided the difference between the highest and the lowest acceptance rate for the bins. For loan officers, we provide only one metric, SPD, because it is relatively simple and most widely used in different contexts. To extend and generalize this UI component, further metrics could be added to the selection, e.g. drawn from (Mitchell et al., 2018) or (Verma & Rubin, 2018).

[Figure 6 near here]

We also provide the ability to specify a custom attribute which could be used as an informal ‘metric’ (REQ2.7), in order to support alternative user-generated or domain specific notions of fairness, as seen in section 3.3.2. For example, during the workshop, stakeholders repeatedly raised notions of affordability which could be added in addition to the predefined fairness metrics. The custom metric builder (Figure 7) allows users to give the metric a name, to select from existing features along with their value distributions against the target, and a work area where they are able to construct a metric using features and mathematical operators, inspired by other formula builders. Further extensions could obviously be made to the formula builder, for instance to show example output, or to incorporate more complex functions than the basic mathematical operators.

[Figure 7 near here]

4.2. System Overview

This component provides important contextual information around fairness of the data or the model to the user by showing the number of instances in the dataset (REQ1.1), an overall target ratio, shown as a percentage and a visual pie slice, which in our use case was the acceptance rate (REQ1.4). For the model, we show the same information

calculated on the test data.

The system overview could of course be extended with other global information. For example, for the model, we could add a positive prediction rate, alongside other accuracy metrics. Other forms of representing and visualizing this data could also be explored.

4.3. Causal Graph

The main way for a user to explore features and relationships between features, a key need that was highlighted in the workshops, is through a causal graph (Figure 8) (REQ2, REQ3). The causal graph is made up of nodes representing features and edges representing causal relationships between features, presented in a circular layout. Causal graphs can be inferred by applying causal discovery algorithms (Zheng, Aragam, Ravikumar, & Xing, 2018) and then followed up and validated by domain experts. Sensitive features marked up during the setup process are highlighted in gold (REQ1.5); relationship strengths are indicated by line thickness (REQ3.2); out-degree is represented by node size. The target is shown as a circle node to differentiate it from the other features.

[Figure 8 near here]

In the dataset view, each feature node contains a bar graph which shows the target distribution, i.e. acceptance rate, for each feature value (REQ2.2). In the case of categorical data this is fairly straightforward, while binning is used for continuous features. We use a simple and commonly used square root approach for the binning process but other ways of determining bin size, such as Sturges approach, could be explored. Above this information, we show a small horizontal bar: its length indicates the relative value of the statistical parity difference metric for this feature, mapped to a 0 to 1 bar chart. This allows users to quickly identify features for further exploration, either through an imbalance in target distributions for feature values, or a high fairness metric value, as discussed in section 3.3.3. We choose percentage bar graphs as they normalize the data and allow easy comparison to enable the identification of trends or possible relationships.

Using a contextual menu, users are able to add a feature to a combination (see section 4.5) (REQ2.9), toggle the feature’s status as sensitive (REQ1.5), or mark the feature as potentially ‘unfair’ for further discussion (REQ1.7, REQ2.16), which set the feature color to red.

To ease visual processing load, clicking or hovering over a feature makes all connected relationships more prominent. Outgoing edges are colored blue to encode the direction of causality (REQ3.2). In addition, if a relationship is clicked or hovered over, the relationship’s strength is shown in numerical form, generated from the causal discovery algorithm.

Initially, all features are displayed in the causal graph to get an overall sense of the data, based on feedback from the workshops (REQ1.1, REQ1.3). Users are also allowed to temporarily filter features and their relationships that interest them most for further exploration, reducing the visual ‘clutter’ of the graph. To do so, users are able to enter a drill-down mode and select the features of interest (REQ2.1, REQ2.3). Only the selected features and their relationships are then shown in this UI component. The user is able to easily exit this drill-down view and return to the full causal graph.

When investigating the model, feature importance is encoded as the feature’s color saturation (Figure 9) (REQ3.1). Feature importance in the model represents the ab-

solute value of the feature weight and therefore can be mapped to a saturation scale between 0 to 1. The less saturated a node, the smaller of an impact it has on the model, giving users an indication of how much features contribute to the model’s predictions.

[**Figure 9 near here**]

The causal graph offers opportunities for extension based on domain and information requirements. Obviously, different choices could have been made for encoding information in the causal graph, by choosing different mappings of information to visual attributes such as colors and size, by selecting other graph types, or by highlighting other important information such as different feature types. However, we encourage designers to favor simplicity and ease of understanding in order to communicate this information to data scientists but particularly to domain experts. We also implemented a common ‘overview and drill-down’ data visualisation strategy which was suited to our stakeholders.

4.4. Features and Relationships Information

The causal graph works in conjunction with a supporting information component. This displays additional contextual information triggered when selecting features and relationships (REQ1.1, REQ 1.3, REQ2.2), satisfying the main requirements outlined in section 3.3.3.

In the dataset view when clicking on a feature (Figure 10), the component shows in-degree and out-degree values derived from the causal graph, a bar graph of fairness metric scores, and a table which shows, for each categorical or binned feature value, the number of instances in each target category – here ‘accepted’ and ‘rejected’ – and the acceptance rate in the same way as in the system overview. In the model view, we add prediction confidence, i.e. the certainty for making this prediction represented by its closeness to the decision boundary (Kulesza et al., 2015), as an additional column to this table, and display this as pie slices. This table provides the same information as the feature bar graph in the causal graph and also allows easy, visual comparison between feature values and against the overall target ratio.

[**Figure 10 near here**]

When clicking on a relationship in the causal graph, the component presents a stacked bar graph (Figure 11), showing the percentage of the intersection of the two related feature values (REQ2.1). The outgoing feature’s value (the cause) is mapped to the x-axis of the graph, while the other feature’s values are shown on the y-axis. This simplified the Sankey diagrams originally shown in the workshops (see section 3.3.3) while supporting easier comparisons.

[**Figure 11 near here**]

This UI component could be extended to suit other use cases. The bar graph of fairness metrics is extensible to show more metrics, if desired. The table of detailed feature information could also include other information that is of importance to the user group. As always, information could be encoded differently. In particular, while we chose to show relationships between two features in a very simplistic way, in preference to Sankey diagrams, there might be other way of better visualizing relationships for specific user group.

The bar graphs within this component are also connected to the dataset component (see section 4.6); clicking on a particular value or bin filters applications shown in the dataset component. This allows users to further explore applications that might be problematic (REQ4).

4.5. Feature Combinations

To support the investigation of intersectional bias and subgroup fairness, as evidenced in the workshops (see section 3.3.3), users are able to add features to a combination (REQ2.9). Subgroups (Figure 12) are presented as ‘cards’ which communicate the subgroup’s acceptance rate, number of application instances in this subgroup and the combination’s feature values. To focus attention on possible issues of fairness, cards were ordered by low to high acceptance ratio. Users were able to highlight combinations as ‘unfair’ for further exploration and discussion.

[Figure 12 near here]

In the implemented prototype, users have to select subgroups of interest. In a real working system, this could be extended through automatic calculation of all possible subgroup combinations. This might be computationally expensive, however. On the other hand, more support for focused user-driven exploration could also be provided, such as filtering against a minimum number of instances that makes up a subgroup.

4.6. Dataset

This component enables users to see the underlying raw data and investigate individual applications (REQ4). The dataset component (Figure 13) presents all features in a spreadsheet-like table, familiar to many users (REQ4.1). It allows users to filter and sort by feature and value, and is also connected to the feature information component described in section 4.4. The target feature (here, the result column) and its values are always visible, with additional color-coding to help with rapid visual processing.

[Figure 13 near here]

The dataset table in the model view (Figure 14) also includes a column for prediction confidence. i.e. the application’s certainty or closeness to the decision boundary, which is shown as a percentage and a pie slice visualization. In addition, if a user selects an application we show the corresponding feature values and their relative feature weights. The weights are shown as bi-directional bars, either negative (red) or positive (blue). In addition, the criticality of the feature value, i.e. the feature value multiplied by the weight of the feature, is encoded by color depth; the deeper the color, the more critical the feature value is in determining the target value. This enables users to understand how features contribute to the model.

[Figure 14 near here]

Additionally, this component is also connected to the compare applications component (section 4.7). By selecting an individual application, a similarity comparison graph is generated.

4.7. Compare Applications

This component is designed to give users the ability to explore similar applications that might have very different outcomes, hence uncovering potential biases or unfairness at the individual level (see section 3.3.3) (REQ4.3, REQ4.4). This is initiated by selecting an application in the dataset component to show a similarity comparison graph.

In this graph in the dataset view (Figure 15), we plot all applications as data points. The target categories are encoded with color (accepted: blue, rejected: red) to be consistent with the dataset table’s encoding of the target. On the y-axis, the plot shows similarity of all applications to the currently selected application, shown in a square outline (REQ4.3). Different similarity metrics could be applied, we chose the Pearson

correlation coefficient. When investigating similarity in the dataset, the target categories are presented on the x-axis. In the model view, the x-axis represents prediction confidence (or certainty) of the target outcome for each application; either to be accepted or rejected. This allows users to understand how close a selected application is to a decision boundary, and also its relationship to all other applications (REQ4.5).

[Figure 15 near here]

The user can select another application on the plot, which shows the compared features side by side, along with the feature similarity score represented as a bar chart (REQ4.4). This information is ordered from least similar values to most similar.

Constructing the plot relies to a great extent on the binary target which aids mapping this information on a 2D visualisation. However, if the target had been multi-categorical or continuous, different visualization choices are needed while still encoding similarity and prediction confidence as important information.

5. Study 2: FairHIL evaluation with stakeholders

The aim of this study was to evaluate FairHIL in order to understand if the UI components we instantiated were useful and usable by loan officers and data scientists, our two stakeholder groups, to investigate fairness. As fairness is such a subjective notion, we did not measure any 'task success', instead we paid attention to qualitative feedback which point the way to improvements of the UI or further requirements.

5.1. Participants

We conducted one-to-one 1-hour sessions with 17 participants, 8 loan officers (mean age=38, stdev=4.54) and 9 data scientists (mean age=31.8, stdev=3.56). Table 4 shows the background details of the participants in more detail. Most data scientists were aware of AI fairness while most loan officers were not. We did not capture participants' familiarity with similar tools in our background questions, however, currently there are no such tools in systematic use at the partner bank. Due to COVID-19 restrictions all sessions were held online through remote meeting software. Recruitment was facilitated by Intesa Sanpaolo, a partner bank involved in the project. This study was approved by City, University of London Ethics Committee. No incentives were offered.

[Table 4 near here]

5.2. Setup

We focused in detail on the dataset UI components, as they were very similar to the AI model UI components. We probed participants' reactions to each UI component as they used the prototype: first we gathered opinions about the setup process of loading the dataset and initialising the model, focusing in detail on selection of sensitive features and metrics. We then explored how they would use the prototype to investigate fairness of a dataset, which consisted of using the causal graph, the feature and relationship information, feature combinations, and investigating the dataset and comparing applications. We repeated this for the model exploration where we focused also on the use of prediction confidence and feature weights.

As they interacted with the prototype, we encouraged the participants to 'think aloud', and verbalize if they would use the UI component, and if so, how. The study

concluded with two post-session questionnaires to evaluate users' experience in investigating fairness using the UI. First, we asked participants to answer the following questions around the *usefulness* of the prototype, using a 7-point Likert scale:

- The prototype supports my ability to assess fairness effectively
- The interface provides the information I need to assess fairness effectively
- The interface provides a sufficient amount of detail needed to assess fairness effectively
- The interface provides the functionality I need to assess fairness effectively
- The interface supports the way I reason when making a decision

We also allowed participants to add free-text comments on their experience of the prototype. Second, we employed the simplified NASA-TLX questionnaire (Hart & Staveland, 1988), which is widely used in assessing UIs and XAI to understand workload.

Sessions were screen and audio recorded and analyzed qualitatively through thematic analysis (Clarke et al., 2015) focusing on themes of understanding, use and usefulness for each UI component; in addition, we also analyzed the usefulness questionnaire's free comments qualitatively. We report the quantitative results of the questionnaires through simple descriptive statistics for each stakeholder group.

5.3. Findings

Overall, the prototype was received equally well by both sets of stakeholders, as evidenced in the usefulness questionnaire responses (Figure 16). All responses were above 3.5 on a 7-point scale, with responses mostly at the upper end of the rating scale. This indicates that both stakeholder sets were happy that the tool supported them in assessing fairness by providing necessary functionality, that it provided information necessary and in sufficient detail to make assessments, and that it supported their reasoning when making decisions. Loan officers viewed the prototype slightly more positively on the whole, but we only found significant differences between data scientists and loan officers in two of the responses, for 'This interface provides a sufficient amount of detail needed to assess fairness effectively' (Mann-Whitney U, $U=58.00$, $p=0.036$) and 'The interface supports the way I reason when making a decision' (Mann-Whitney U, $U=57.00$, $p=0.046$).

[Figure 16 near here]

Regarding workload as measured through the NASA-TLX questionnaire, there were no significant differences between data scientists and loan officers when investigated through Mann-Whitney U tests (Figure 17). Note that ratings for both stakeholder sets on mental demand and effort were relatively high. This is likely due to the complexity of the UI and the high amount of information that needs to be processed by the user. However, this is balanced by the perceived performance which indicates that users also felt that the fairness exploration through the prototype was paying off and led to success.

These results show that the UI components were perceived as useful for performing their tasks by both sets of stakeholders, and that felt that they were usable.

[Figure 17 near here]

Our efforts of explaining the UI concepts in AI fairness were mainly successful, with very few areas that caused confusion. The ML and fairness concepts that were particularly difficult for loan officers to grasp were targets (5 out of 8 loan officers) and sensitive features (4 out of 8 loan officers). Frequently, targets and sensitive features

were seen as *any* features that mattered in decision-making, e.g. that were strongly related to the target or to sensitive features. Possibly, this would require further clarifications about input and output in ML predictive models, and what features are protected by law.

We now turn to investigating responses to individual UI design components.

5.3.1. Metrics Choice

An area that requires additional transparency to increase understanding is fairness metrics. Both sets of stakeholders, but in particular data scientists, requested more information about how they were calculated, which ones should be selected, and how they should be applied. This reflects the fact that in general they were not entirely trusted as a method for assessing fairness, for example:

“These five metrics would have trouble letting me understand, won’t allow me to understand well if I am being discriminatory because I don’t see in these descriptions the way they are used.” - ED03

The custom metric builder in the UI was very popular: 7 data scientists and 4 loan officers stated that they would use it to create metrics. While data scientists stated that they would use this to create metrics that were not included in the UI by default, loan officers wanted to create a new *feature* to measure affordability for *each* application; it seems that they were trying to define a similarity metric in the sense of (Dwork et al., 2012). For this, they wanted to define and compute affordability by selecting other features, and then, if the applicant was able to afford the loan, use this as an acceptance criterion, independent of other (sensitive) features. Of course, a challenge with this approach is to define affordability in a fair way.

5.3.2. Causal Graph

Responses indicated that there were usually no problems in understanding the causal graph; all users considered it very intuitive, for example:

”There are also indirect effects, indirect ways that age can affect the result. I think that it is a really interesting picture” - ED03

Encoding important relationships as thicker lines helped users to focus on the problematic relationships, such as ones connected to sensitive features, which possibly cause fairness issues, or to the target, which directly affected the final loan decisions. Highlighting sensitive features was mentioned positively by all data scientists and 6 loan officers, while highlighting important relationships were mentioned by 8 data scientists and 7 loan officers. Users also said that the causal graph encouraged exploration of relationships between features that were not directly affecting the target, or that it encouraged exploring features that were not sensitive.

Using the casual graph, loan officers were able to confirm or reject hypotheses they had about the data. For example:

“It’s clearer for me that age is an important feature more than gender or citizenship, that’s an idea I had before because of my experience, and I told you before that age is more important but now, I can realise it even looking at the chart, it’s like a confirmation.” – EL03

This can be useful in investigating decision-making biases, and possible discriminatory practices, based on domain expertise.

Users were able to simplify the causal graph to drill down into exploring features of interest in more detail, and all participants found this useful. Participants either concentrated on sensitive features and their relationships to the target, or to hide features that were considered irrelevant or unimportant, usually because they did not have strong relationships with other features.

5.3.3. Features and Relationships Information

Both sets of stakeholders paid attention to the fairness metrics in the information component or metric indicator in the causal graph. However, possibly because it was not clear what these fairness metrics captured or how they could be used, they were much more interested in exploring the relationships. This was useful to explore possible biases and discrimination in decision-making. For example:

“ah that’s interesting to see the acceptance rating on the age, so we can appreciate that it’s growing, there’s just one that is decreasing here. Also, we can see the number of people for each bucket, this is useful. Yeah, the majority are in their 40s and 50s.” – ED09

When exploring the model, both sets of stakeholders found the information provided useful for exploring possible issues in the decision-making process and identify fairness issues. Nine data scientists and 6 loan officers stated that feature importance was useful; 4 data scientists and 3 loan officers commented on the usefulness of the prediction confidence indicators; 7 data scientists and 5 loan officers mentioned the usefulness of seeing feature weightings for individual applications.

5.3.4. Feature Combinations

Users were able to combine features and explore their relationship to the target through the Group Combination UI component. Nine data scientists and 7 loan officers stated that they would use the combination feature to help them investigate bias. For example:

“We can understand the relation between age, citizenship, net monthly income in order to obtain a loan. And as I expected the rate is growing when net monthly income is growing, citizenship national and the age is around 40. I expected this result, this acceptance rate. I believed this acceptance rate could have been higher than 54.2%... I think this tool is fantastic because you can understand where, in what kind of area we have to concentrate in order to avoid a lot of rejected, to understand why there are a lot of rejected, and to help avoid this result.” – EL08

However, there were some concerns voiced by participants around scalability of this functionality to include more feature intersections (we limited this to 3) or how to facilitate easy comparisons between subgroups (we showed a percentage).

5.3.5. Dataset

Having access to the raw data was appreciated by all participants because it facilitated the investigation of individual fairness. This was especially surprising since the workshops reported in section 3.3.2 seemed to show that data scientists were not as interested in exploring the dataset from this perspective.

Linking the dataset component to the feature and relationship information component was especially liked:

“[It would be nice to] click on this grey area and click to see how many male or how many females are in this grey area [the credit risk level group in the bar]. [...] Ahh ok the dataset updates, cool, cool, this is very useful. This is useful because here I can see the records and understand why, and if there is any error or something like that... to understand better the data.” - ED06

Participants stated that they liked how the dataset information was interconnected with other UI components which allowed them to ‘trace’ the decision-making through to the raw data.

5.3.6. Compare Applications

The compare applications component supported the exploration of individual fairness further. All data scientists and 6 loan officers stated that they would use this component. Participants contrasted individual applications in an attempt to explain alternative outcomes and to identify the features and their values that could have impacted the decisions, for example:

“I can compare the selected one that has been rejected with this one that has been approved, in order to see if there is some difference that has affected the application. This is important because for example the first thing that I see is that the loan amount is very high in the one that has been rejected, while this one is not so high, probably one of the most important reasons that the application has been rejected for the loan amount.” – EL06

6. Discussion

Our work has investigated the design space to support domain experts and data scientists to investigate fairness, and how a UI could be implemented for the loan applications domain and more generally to fulfill their requirements. We will now discuss, first, the limitations of our study, second, how current tools and our prototype cover this design space, and third, implications for further extensions to FairHIL. We conclude by discussing three areas of supporting human-in-the-loop fairness that warrant further attention.

6.1. Limitations

The work presented here has some limitations in its scope. First, our workshops and evaluation drew on a relatively small sample of participants, and thus cannot generate any robust quantitative findings. However, the value of a qualitative approach is to provide rich insights into users’ needs that is able to generate novel UI designs. A related issue is that we have not implemented a fully working prototype that could be evaluated in a larger field trial, embedded in the practices of domain experts and data scientists. Both of these concerns would need to be addressed in further work that goes beyond the scope of this paper.

Second, our investigations were only situated in the loan applications domain. We have already highlighted that this might affect the datasets used, the ML models and hence also the UI design and visualizations. Furthermore, we involved loan officers and data scientists from only one organization. While we do not anticipate that practices differ significantly between banking institutions in terms of loan applications, we do not

know how fairness investigations extend to other stakeholders in loan applications, e.g. loan applicants or regulatory bodies. Similarly, domain experts and data scientists in other domains might have different criteria and ways to assess fairness which warrants further investigation.

6.2. Coverage of design space by existing tools

Table 5 shows how requirements that we uncovered during our workshops—which constitute the design space for tools to support investigating fairness—map to functionality by existing tools, including FairHIL. Currently, there are no tools that fulfill *all* of the requirements that we uncovered during our work. However, FairHIL is so far the first to integrate most of these requirements to support both stakeholder groups to investigate fairness. Still, particular gaps exist in providing advanced data manipulation functionality, comparing individual cases, and mitigating fairness by adjusting feature weights or ‘deleting’ features. These are areas that future design endeavours could target.

[Table 5 near here]

6.3. Design Implications for Human-in-the-loop Fairness Tools

Our findings have concrete ramifications for the design of human-in-the-loop fairness UIs. We here summarize these insights.

While metrics are useful for assessing fairness, data scientists sometimes had difficulties in understanding fairness metrics, as found in section 5.3.1. Hence, in addition to describing the metrics on a high level, as we did in FairHIL, the details of how they are calculated should be given. Furthermore, users will need to be given guidance as to the expected ranges of these metrics, or when a metric indicates that there are discriminatory outcomes. This information could be integrated into visualizations. Loan officers would also benefit from this improvement, since they have less background knowledge about fairness and ML concepts.

From 5.3.3, we found that allowing users to see the features’ relationship to the target and to each other was pivotal. We suggest that this could be used to identify potential biases in decision-making, and that this could lead to further discussions within the organization for changing policy or applying fairness mitigation.

As we have seen in 5.3.4, exploring subgroup fairness was viewed positively by both stakeholder groups and could be a powerful way to explore intersectional bias (Kearns et al., 2018). However, further work needs to consider how to extend combinations of features, possibly automatically as discussed in section 4.5, while easing comparisons between subgroups through visual means. To make the assessment of feature combination more effective, it is necessary to extend the number of features taken into account when comparing the combinations, and to show the number of applications included in a specific subgroup.

Users want a way to explore the raw data and its relationship to the target, as we found in 5.3.5. This allows users to trace decisions back to individual data instances, and gain a deeper understanding of the data as it relates to fairness. Current interfaces rarely support this (Ahn & Lin, 2020) yet its inclusion, especially to drive further detailed comparisons with other applications is warranted. Our work has shown that comparing applications was intuitive and useful for assessing fairness in loan applications, as we found in 5.3.6. Yet there are few human-in-the-loop fairness UIs, FairSight

(Ahn & Lin, 2020) being a notable exception, that currently support this approach. This component also clearly supports the exploration of individual fairness in loan decisions, helpful for both stakeholder groups. Yet, the Compare Applications component could be further extended. While only data scientists asked to compare clusters' i.e. more than one application, this ability may be useful for both stakeholder groups. Identifying these subgroups could also be automatically supported through, for example, k-nearest neighbor approaches. Comparisons could be extended even further, by letting users choose similarity metrics, or by letting the user build a custom similarity metric by re-weighting features.

6.4. Supporting the Human-in-the-loop Fairness Process

Our work also raises three wider questions.

6.4.1. How to Explain ML and Fairness Concepts?

Our findings confirm that transparency and explanations of ML and fairness concepts are necessary (Hohman, Head, Caruana, DeLine, & Drucker, 2019; Lim & Dey, 2009). Adopting common XAI approaches (Kulesza et al., 2015; Lundberg & Lee, 2017; Ribeiro et al., 2016) to explain models in a way that is understood by non-technical domain experts and technically-savvy data scientists alike might be one way forward. To ensure understandability, information needs to be carefully encoded to support intuitive exploration and easy visual processing. While general principles for explaining ML concepts have been suggested e.g. 'be iterative, sound and complete but do not overwhelm' (Kulesza et al., 2015), there is still much work to be done to translate these principles into concrete design guidelines or design patterns.

A lack of understanding around fairness measures often meant that they are not transparent as to what they capture, how they are calculated and how they should be used. For example, in our studies, both domain experts and data scientists wanted more information and explanations of when measures indicated unfairness or discrimination. There also seemed confusion when decision-making bias tipped over into unfairness, for example, older applicants had an easier time getting a loan because their credit rating was higher due to a longer credit history yet does this mean it is unfair? Fairness measures could be made more transparent through a UI but solving these issues might also involve education of stakeholders around ML concepts and fairness principles.

6.4.2. What is fair?

It is becoming clear from our as well as others' work (Kasinidou, Kleanthous, Barlas, & Otterbacher, 2021; Lee, Jain, Cha, Ojha, & Kusbit, 2019; Saxena et al., 2020; Woodruff, Fox, Rousso-Schindler, & Warshaw, 2018) that different stakeholders might have different conceptualizations of what fairness is, how it should be measured, and how fairness considerations should be integrated into the wider decision-making process when using ML. We found that in our domain, loan officers preferred to pay attention to aspects related to individual fairness while data scientists tend to employ notions of group fairness. It could be argued that existing fairness measures capture both of these perspectives, and thus could be integrated in a UI to support both domain experts and data scientists. However, as previously covered in section 2.1, often these measures are not compatible or there needs to be a trade-off between fairness measures and accuracy. So who gets to decide which fairness measures to apply? We

suggest that this is a question that urgently needs attention within an organizational and also a socio-technical context.

6.4.3. *How and when to investigate fairness?*

Our study suggests that requirements to support investigating fairness are very rich and varied. While we found that domain experts and data scientists might go about investigating fairness slightly differently, it was surprising to us that in the main these two stakeholder groups agreed on the information and functionality required for them to ably investigate fairness. This begs the question whether different toolkits are indeed necessary for these user groups, or whether we can find a way to communicate effectively across different user types. Current perspectives in XAI argue that different UIs might be necessary for different stakeholder groups (Gunning et al., 2019) but our work suggests that it might be possible to craft UIs that suit a variety of users.

An interesting yet open question is how fairness investigations should be integrated into the AI model development. Most existing tools, such as What-if Tool and AI Fairness 360, adopt a mitigation strategy after AI models have been developed. While tools such as ours and Silva can be used on datasets as well as AI models, urgent discussion is needed on how to integrate fairness in *all* steps in the AI development process, from data collection to after-deployment.

There are several opportunities to actively involve humans in the design process of an AI-enabled system. A tool such as ours could be used once data has been collected to explore bias and potential unfairness in a dataset, even before a model is developed. This could involve business stakeholders but also people affected by the decisions of such an AI system, to better understand whether the collected data is acceptable. Furthermore, these stakeholders could again be involved once models have been developed, before the deployment. Their input could contribute to the choice between models for deployment, or to assess whether a model is indeed ready to be deployed. We also see the potential for use of our tool *after* an AI system has been deployed, to keep checks on the continuing fairness of the AI system.

7. Conclusion

In this paper, we presented our findings from a series of investigations to understand the design space to support investigating AI fairness for domain experts and data scientists. We ran design workshops to gather requirements for human-in-the-loop fairness UIs. We then presented 7 UI components that support data scientists as well as loan officers in assessing fairness and finding potential fairness issues, instantiated in FairHIL, a prototype UI. We addressed how these UI components could be adapted and extended for other domains. Finally, we evaluated FairHIL through a user study.

Our findings show that:

- Requirements underlying the design space are rich and varied. Surprisingly, domain experts and data scientists share many requirements to investigate fairness.
- AI and fairness concepts need to be carefully explained, especially how they should be applied and what might indicate fairness issues.
- Causal graphs, alongside relevant and easy-to-process information about features and relationships, were an intuitive way to support fairness investigations by both stakeholder groups.
- Both stakeholder groups appreciated access to the underlying data and the abil-

ity to explore and compare individual applications.

- Comparing subgroups is important to users.

This work will help researchers, designers and developers to better understand how domain experts and data scientists investigate fairness, and to build better human-centred AI tools to support them. We hope that this will lead to AI systems that are fairer and more transparent for everyone.

Acknowledgement

Work by City, University of London was supported by Fujitsu Limited under a research contract.

Disclosure statement

We don't have a conflict of interest.

References

- Agarwal, A., Beygelzimer, A., Dudík, M., Langford, J., & Wallach, H. (2018). A reductions approach to fair classification. *arXiv preprint arXiv:1803.02453*.
- Ahn, Y., & Lin, Y. R. (2020). Fairsight: Visual analytics for fairness in decision making. *IEEE Transactions on Visualization and Computer Graphics*, 26(1), 1086-1095.
- Barocas, S., Hardt, M., & Narayanan, A. (2019). *Fairness and machine learning*. fairml-book.org. (<http://www.fairmlbook.org>)
- Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671-732. Retrieved from <http://www.jstor.org/stable/24758720>
- Bellamy, R. K. E., Dey, K., Hind, M., Hoffman, S. C., Houde, S., Kannan, K., ... Zhang, Y. (2018, October). *AI Fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias*. Retrieved from <https://arxiv.org/abs/1810.01943>
- Berman, J. J., Murphy-Berman, V., & Singh, P. (1985). Cross-cultural similarities and differences in perceptions of fairness. *Journal of Cross-Cultural Psychology*, 16(1), 55-67. Retrieved from <https://doi.org/10.1177/0022002185016001005>
- Binns, R., Van Kleek, M., Veale, M., Lyngs, U., Zhao, J., & Shadbolt, N. (2018). 'it's reducing a human being to a percentage': Perceptions of justice in algorithmic decisions. In *Proceedings of the 2018 chi conference on human factors in computing systems* (p. 1-14). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3173574.3173951>
- Bird, S., Dudík, M., Edgar, R., Horn, B., Lutz, R., Milan, V., ... Walker, K. (2020, May). *Fairlearn: A toolkit for assessing and improving fairness in AI* (Tech. Rep. No. MSR-TR-2020-32). Microsoft. Retrieved from <https://www.microsoft.com/en-us/research/publication/fairlearn-a-toolkit-for-assessing-and-improving-fairness-in-ai/>
- Blake, P. R., McAuliffe, K., Corbit, J., Callaghan, T. C., Barry, O., Bowie, A., ... Warneken, F. (2015, November). The ontogeny of fairness in seven societies. *Nature*, 528, 258-261.
- Bolton, L. E., Keh, H. T., & Alba, J. W. (2010). How do price fairness perceptions differ across culture? *Journal of Marketing Research*, 47(3), 564-576. Retrieved from <https://doi.org/10.1509/jmkr.47.3.564>

- Cabrera, A. A., Epperson, W., Hohman, F., Kahng, M., Morgenstern, J., & Chau, D. H. (2019). Fairvis: Visual analytics for discovering intersectional bias in machine learning. In *2019 IEEE Conference on Visual Analytics Science and Technology (VAST)* (p. 46-56).
- Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, *356*(6334), 183–186. Retrieved from <https://science.sciencemag.org/content/356/6334/183>
- Calmon, F., Wei, D., Vinzamuri, B., Ramamurthy, K. N., & Varshney, K. R. (2017). Optimized pre-processing for discrimination prevention. In *Advances in neural information processing systems* (pp. 3992–4001).
- Cheng, H.-F., Stapleton, L., Wang, R., Bullock, P., Chouldechova, A., Wu, Z. S. S., & Zhu, H. (2021, May). Soliciting Stakeholders’ Fairness Notions in Child Maltreatment Predictive Systems. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (pp. 1–17). New York, NY, USA: Association for Computing Machinery. Retrieved 2021-08-18, from <https://doi.org/10.1145/3411764.3445308>
- Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, *5*(2), 153–163.
- Clarke, V., Braun, V., & Hayfield, N. (2015). Thematic analysis. *Qualitative psychology: A practical guide to research methods*, 222–248.
- Cook, K. S., & Hegtvedt, K. A. (1983). Distributive justice, equity, and equality. *Annual Review of Sociology*, *9*(1), 217-241. Retrieved from <https://doi.org/10.1146/annurev.so.09.080183.001245>
- Corbett-Davies, S., & Goel, S. (2018). *The measure and mismeasure of fairness: A critical review of fair machine learning*.
- Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., & Huq, A. (2017). Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (p. 797–806). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3097983.3098095>
- Dodge, J., Liao, Q. V., Zhang, Y., Bellamy, R. K. E., & Dugan, C. (2019). Explaining models: An empirical study of how explanations impact fairness judgment. In *Proceedings of the 24th International Conference on Intelligent User Interfaces* (p. 275–285). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3301275.3302310>
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (p. 214–226). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/2090236.2090255>
- Eckhoff, T. (1974). *Justice: Its determinants in social interaction*. Rotterdam University Press.
- Equality, & Commission, H. R. (2020). *Protected characteristics*. Retrieved from <https://www.equalityhumanrights.com/en/equality-act/protected-characteristics>
- Friedman, B., & Nissenbaum, H. (1996, July). Bias in computer systems. *ACM Trans. Inf. Syst.*, *14*(3), 330–347. Retrieved from <https://doi.org/10.1145/230538.230561>
- Gunning, D., Stefik, M., Choi, J., Miller, T., Stumpf, S., & Yang, G.-Z. (2019, December). XAI—Explainable artificial intelligence. *Science Robotics*, *4*(37). Retrieved 2020-02-05, from <https://robotics.sciencemag.org/content/4/37/eaay7120>
- Hajian, S., Bonchi, F., & Castillo, C. (2016). Algorithmic bias: From discrimination discovery to fairness-aware data mining. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (p. 2125–2126). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/2939672.2945386>
- Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. In *Advances in neural information processing systems* (pp. 3315–3323).
- Hart, S. G., & Staveland, L. E. (1988). Development of NASA-TLX (Task Load Index): Results

- of Empirical and Theoretical Research. In Peter A. Hancock and Najmedin Meshkati (Ed.), *Advances in Psychology* (Vol. Volume 52, pp. 139–183). North-Holland. Retrieved 2013-11-06, from <http://www.sciencedirect.com/science/article/pii/S0166411508623869>
- Hegtvedt, K. A. (2005). Doing justice to the group: Examining the roles of the group in justice research. *Annual Review of Sociology*, *31*, 25–45. Retrieved from <http://www.jstor.org/stable/29737710>
- Hofstede, G. (2011). Dimensionalizing cultures: The Hofstede model in context. *Online readings in psychology and culture*, *2*(1), 2307–0919.
- Hohman, F., Head, A., Caruana, R., DeLine, R., & Drucker, S. M. (2019, May). Gamut: A Design Probe to Understand How Data Scientists Understand Machine Learning Models. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). New York, NY, USA: Association for Computing Machinery. Retrieved 2021-09-19, from <https://doi.org/10.1145/3290605.3300809>
- Jeff Larson, L. K., Surya Mattu, & Angwin, J. (2016, May). *How we analyzed the compas recidivism algorithm*. Retrieved from <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>
- Kamiran, F., & Calders, T. (2009). Classifying without discriminating. In *2009 2nd international conference on computer, control and communication* (pp. 1–6).
- Kamishima, T., Akaho, S., & Sakuma, J. (2011). Fairness-aware learning through regularization approach. In *2011 IEEE 11th International Conference on Data Mining Workshops* (pp. 643–650).
- Kasinidou, M., Kleanthous, S., Barlas, P., & Otterbacher, J. (2021). I agree with the decision, but they didn’t deserve this: Future developers’ perception of fairness in algorithmic decisions. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency* (pp. 690–700).
- Kearns, M., Neel, S., Roth, A., & Wu, Z. S. (2018). Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *International conference on machine learning* (pp. 2564–2572).
- Kim, T.-Y., & Leung, K. (2007). Forming and reacting to overall fairness: A cross-cultural comparison. *Organizational Behavior and Human Decision Processes*, *104*(1), 83–95. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0749597807000076>
- Kulesza, T., Burnett, M., Wong, W.-K., & Stumpf, S. (2015). Principles of Explanatory Debugging to Personalize Interactive Machine Learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces* (pp. 126–137). New York, NY, USA: ACM. Retrieved 2015-12-14, from <http://doi.acm.org/10.1145/2678025.2701399>
- Kulesza, T., Stumpf, S., Burnett, M., & Kwan, I. (2012). Tell me more? the effects of mental model soundness on personalizing an intelligent agent. In *Proceedings of the sigchi conference on human factors in computing systems* (p. 1–10). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/2207676.2207678>
- Kusner, M. J., Loftus, J., Russell, C., & Silva, R. (2017). Counterfactual fairness. In *Advances in neural information processing systems* (pp. 4066–4076).
- Lee, M. K., Jain, A., Cha, H. J., Ojha, S., & Kusbit, D. (2019). Procedural justice in algorithmic fairness: Leveraging transparency and outcome control for fair algorithmic mediation. *Proceedings of the ACM on Human-Computer Interaction*, *3*(CSCW), 1–26.
- Lee, M. K., Kim, J. T., & Lizarondo, L. (2017). A human-centered approach to algorithmic services: Considerations for fair and motivating smart community service management that allocates donations to non-profit organizations. In *Proceedings of the 2017 CHI conference on human factors in computing systems* (p. 3365–3376). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3025453.3025884>
- Lim, B. Y., & Dey, A. K. (2009). Assessing Demand for Intelligibility in Context-aware Applications. In *Proceedings of the 11th International Conference on Ubiquitous Computing* (pp. 195–204). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/1620545.1620576>
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions.

- In I. Guyon et al. (Eds.), *Advances in neural information processing systems* (Vol. 30). Curran Associates, Inc. Retrieved from <https://proceedings.neurips.cc/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf>
- Mallari, K., Inkpen, K., Johns, P., Tan, S., Ramesh, D., & Kamar, E. (2020). Do i look like a criminal? examining how race presentation impacts human judgement of recidivism. In *Proceedings of the 2020 chi conference on human factors in computing systems* (p. 1–13). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3313831.3376257>
- Mattila, A. S., & Choi, S. (2006). A cross-cultural comparison of perceived fairness and satisfaction in the context of hotel room pricing. *International Journal of Hospitality Management*, 25(1), 146 - 153. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0278431904001240>
- Mitchell, S., Potash, E., Barocas, S., D’Amour, A., & Lum, K. (2018). Prediction-based decisions and fairness: A catalogue of choices, assumptions, and definitions. *arXiv preprint arXiv:1811.07867*.
- Narayanan, A. (2018). Translation tutorial: 21 fairness definitions and their politics. In *Proc. conf. fairness accountability transp., new york, usa* (Vol. 2, pp. 6–2).
- Olteanu, A., Castillo, C., Diaz, F., & Kicman, E. (2019). Social data: Biases, methodological pitfalls, and ethical boundaries. *Frontiers in Big Data*, 2, 13. Retrieved from <https://www.frontiersin.org/article/10.3389/fdata.2019.00013>
- Rawls, J. (1958). Justice as fairness. *The philosophical review*, 67(2), 164–194.
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). ”why should i trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining* (p. 1135–1144). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/2939672.2939778>
- Saxena, N. A., Huang, K., DeFilippis, E., Radanovic, G., Parkes, D. C., & Liu, Y. (2020). How do fairness definitions fare? testing public attitudes towards three algorithmic definitions of fairness in loan allocations. *Artificial Intelligence*, 283, 103238.
- Shneiderman, B. (2020, April). Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy. *International Journal of Human-Computer Interaction*, 36(6), 495–504. Retrieved 2021-08-18, from <https://doi.org/10.1080/10447318.2020.1741118> (Publisher: Taylor & Francis eprint: <https://doi.org/10.1080/10447318.2020.1741118>)
- Shneiderman, B. (2021, August). Responsible AI: Bridging From Ethics to Practice. *Communications of the ACM*, 64(8), 32–35. Retrieved 2021-08-18, from <https://cacm.acm.org/magazines/2021/8/254306-responsible-ai/fulltext?mobile=false>
- Veale, M., Van Kleek, M., & Binns, R. (2018). Fairness and accountability design needs for algorithmic support in high-stakes public sector decision-making. In *Proceedings of the 2018 chi conference on human factors in computing systems* (p. 1–14). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3173574.3174014>
- Verma, S., & Rubin, J. (2018). Fairness definitions explained. In *2018 ieee/acm international workshop on software fairness (fairware)* (pp. 1–7).
- Wang, R., Harper, F. M., & Zhu, H. (2020). Factors influencing perceived fairness in algorithmic decision-making: Algorithm outcomes, development procedures, and individual differences. In *Proceedings of the 2020 chi conference on human factors in computing systems* (p. 1–14). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3313831.3376813>
- Wexler, J., Pushkarna, M., Bolukbasi, T., Wattenberg, M., Viégas, F., & Wilson, J. (2020). The what-if tool: Interactive probing of machine learning models. *IEEE Transactions on Visualization and Computer Graphics*, 26(1), 56–65.
- Woodruff, A., Fox, S. E., Rousso-Schindler, S., & Warshaw, J. (2018). A qualitative exploration of perceptions of algorithmic fairness. In *Proceedings of the 2018 chi conference on human factors in computing systems* (pp. 1–14).
- Yan, J. N., Gu, Z., Lin, H., & Rzeszotarski, J. M. (2020). Silva: Interactively assessing ma-

- chine learning fairness using causality. In *Proceedings of the 2020 chi conference on human factors in computing systems* (p. 1–13). New York, NY, USA: Association for Computing Machinery. Retrieved from <https://doi.org/10.1145/3313831.3376447>
- Yang, F., Cisse, M., & Koyejo, O. O. (2020). Fairness with overlapping groups; a probabilistic perspective. *Advances in Neural Information Processing Systems*, 33.
- Zemel, R., Wu, Y., Swersky, K., Pitassi, T., & Dwork, C. (2013). Learning fair representations. In *International conference on machine learning* (pp. 325–333).
- Zhang, B. H., Lemoine, B., & Mitchell, M. (2018). Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 aai/acm conference on ai, ethics, and society* (pp. 335–340).
- Zheng, X., Aragam, B., Ravikumar, P. K., & Xing, E. P. (2018). Dags with no tears: Continuous optimization for structure learning. In *Advances in neural information processing systems* (pp. 9472–9483).

Table 1. A comparison of UI functionality among existing tools described in section 2.2.

Functionality	What-if Tool (Wexler et al., 2020)	FairVis (Cabrera et al., 2019)	FairSight (Ahn and Lin, 2020)	Silva (Yan et al., 2020)	Fairness Elicitation Tool (Cheng et al., 2021)
Show performance metrics	✓	✓	✓		
Show statistical fairness metrics	✓		✓	✓	✓
Select sensitive/protected features			✓	✓	
Show data in groups	✓	✓	✓	✓	✓
Show sub-groups	✓				✓
Show counterfactuals to data points	✓	✓			
Show decision boundaries	✓				
Provide descriptive stats for features	✓		✓		
Show feature contributions		✓	✓		✓
Show causal graph				✓	
Show data similarity			✓		✓

Table 2. Overview of workshop participants' details.

ID	Gender	Age	Education	Role
WL01	Male	36	Bachelor	Loan Officer
WL02	Male	33	Master	Loan Officer
WL03	Male	39	Master	Loan Officer
WL04	Male	42	Master	Loan Officer
WL05	Male	34	Master	Loan Officer
WL06	Female	35	Master	Loan Officer
WD01	Female	28	Master	Data Scientist
WD02	Male	38	Doctorate	Data Scientist
WD03	Female	27	Master	Data Scientist
WD04	Male	28	Master	Data Scientist
WD05	Male	28	Master	Data Scientist
WD06	Female	29	Master	Data Scientist

Table 3. Requirements arising from design workshops. In the Stakeholder row, L stands for the requirements from loan officers, D from data scientists, and B from both stakeholders.

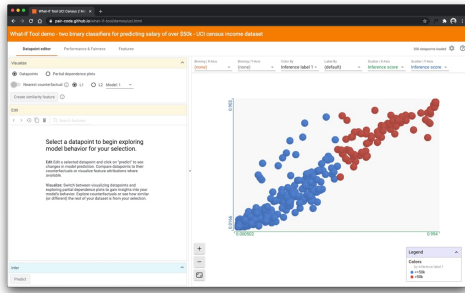
Area	Use	Requirement	Stakeholder
1. Attribute overviews	Informational	1.1 attributes, number of records and attribute value distributions	B
		1.2 amount of missing data	B
		1.3 fairness metrics for model and individual protected attributes	B
		1.4 target distribution	B
		1.5 protected attributes	D
		1.6 explanation how attribute values are calculated or derived	D
		1.7 Identify if the data is subjective/ objective	D
		1.8 Identify where the data has come from (applicant, bank, third party) / attribute provenance	D
2. Investigate Relationships between attributes	Informational	2.1 distribution of protected attributes with other attributes	B
		2.2 distribution of user-selected attribute values (e.g. credit risk ratings) and target values	B
		2.3 distribution of two user-selected attributes' values	B
		2.4 Credit risk rating traffic light system	D
	Functional	2.5 Support data transformations (e.g. categorical into numerical, binning)	D
		2.6 support filling in missing values	D
		2.7 Support creation of new attributes (i.e. calculated from other attributes e.g. affordability)	B
		2.8 Ability to create/include own fairness metric (if not already in system)	B
		2.9 Allow creation of subgroups based on a combination of attributes and see their distribution on target	B
		2.10 Input custom thresholds to affect AI model	B
	Adjust model	2.11 Change weights on attributes to adjust AI model	B
		2.12 Remove attributes	B
		2.13 Change weightings of attribute variables (the variables that make up the attribute) on attributes	D
		2.14 Identify similar attributes which do not contain protected attributes and substitute attribute from these choices	D
		2.15 Optimise model against fairness metrics and accuracy automatically	D
		2.16 Feedback to data scientists on 'questionable' attributes that should not be used for decision-making	L
3. Causal Graph	Informational	3.1 Node weight and impact on target	B
		3.2 Relationships between nodes, their 'strength' and direction	B
		3.3 Explanation of how the graph was derived	B
	Adjust model	3.4 The ability to remove nodes and relationships to adjust AI model	B
4. Individual cases	Informational	4.1 specific application and attribute values	B
		4.2 fairness metric for individual case	B
		4.3 Level of similarity between cases	B
		4.4 Select specific cases to compare and show which attributes are similar	B
		4.5 Show decision boundaries	B
	Functional	4.6 See 'What If' results on target based on changes to attribute values	L
5. Model	Informational	5.1 how model works	B
		5.2 how it was created, rationale for decisions in modelling	B
		5.3 who created it	B

Table 4. Overview of evaluation participants' details.

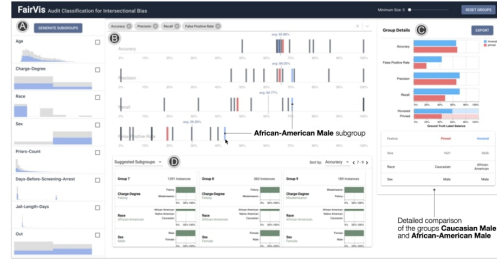
ID	Gender	Age	Education	AI Fairness Knowledge	Role
EL01	Female	37	Doctorate	I am aware of fair machine learning	Loan Officer
EL02	Female	40	Master	None	Loan Officer
EL03	Male	42	Master	None	Loan Officer
EL04	Male	43	Master	None	Loan Officer
EL05	Male	41	Master	I am aware of fair machine learning	Loan Officer
EL06	Female	30	Master	None	Loan Officer
EL07	Male	38	Master	None	Loan Officer
EL08	Male	33	Master	None	Loan Officer
ED01	Female	32	Master	I am aware of fair machine learning	Data Scientist
ED02	Female	29	Master	I actively work in fair machine learning	Data Scientist
ED03	Male	34	Doctorate	I am aware of fair machine learning	Data Scientist
ED04	Female	31	Master	None	Data Scientist
ED05	Male	34	Doctorate	I am aware of fair machine learning	Data Scientist
ED06	Male	35	Master	None	Data Scientist
ED07	Male	27	Master	I am aware of fair machine learning	Data Scientist
ED08	Female	37	Doctorate	I consider myself up to date with what's in the fair machine learning space	Data Scientist
ED09	Male	27	Master	I am aware of fair machine learning	Data Scientist

Table 5. Requirements arising from design workshops covered by conventional tools and ours. In the Stakeholder row, L stands for the requirements from loan officers, D from data scientists, and B from both stakeholders. WiT, FV, FS, SV, FET, and HFIL stand for What-if Tool (Wexler et al., 2020), FairVis Cabrera et al. (2019), FairSight (Ahn & Lin, 2020), Silva (Yan et al., 2020), Fairness Elicitation Tool (Cheng et al., 2021), and FairHIL (our tool), respectively.

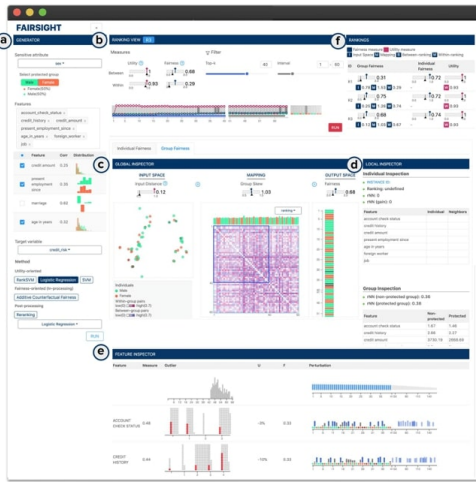
Area	Use	Requirement	Stakeholder	WiT	FV	FS	SV	FET	HFIL (ours)
1. Attribute overviews	Informational	1.1 attributes, number of records and attribute value distributions	B	✓	✓	✓		✓	✓
		1.2 amount of missing data	B						
		1.3 fairness metrics for model and individual protected attributes	B	✓	✓	✓	✓	✓	✓
		1.4 target distribution	B	✓	✓				✓
		1.5 protected attributes	D			✓	✓		✓
		1.6 explanation how attribute values are calculated or derived	D						
		1.7 Identify if the data is subjective/ objective	D						
		1.8 Identify where the data has come from (applicant, bank, third party) / attribute provenance	D						
2. Investigate Relationships between attributes	Informational	2.1 distribution of protected attributes with other attributes	B	✓		✓			✓
		2.2 distribution of user-selected attribute values (e.g. credit risk ratings) and target values	B	✓		✓			✓
		2.3 distribution of two user-selected attributes' values	B	✓					✓
		2.4 Credit risk rating traffic light system	D						
	Functional	2.5 Support data transformations (e.g. categorical into numerical, binning)	D	✓					
		2.6 support filling in missing values	D	✓					
		2.7 Support creation of new attributes (i.e. calculated from other attributes e.g. affordability)	B	✓					✓
		2.8 Ability to create/include own fairness metric (if not already in system)	B						✓
2. Investigate Relationships between attributes	Functional	2.9 Allow creation of subgroups based on a combination of attributes and see their distribution on target	B	✓	✓			✓	✓
		2.10 Input custom thresholds to affect AI model	B						
	Adjust model	2.11 Change weights on attributes to adjust AI model	B	✓					
		2.12 Remove attributes	B	✓					
		2.13 Change weightings of attribute variables (the variables that make up the attribute) on attributes	D	✓					
		2.14 Identify similar attributes which do not contain protected attributes and substitute attribute from these choices	D						
		2.15 Optimise model against fairness metrics and accuracy automatically	D						
2.16 Feedback to data scientists on 'questionable' attributes that should not be used for decision-making	L								
3. Causal Graph	Informational	3.1 Node weight and impact on target	B						✓
		3.2 Relationships between nodes, their 'strength' and direction	B				✓		✓
		3.3 Explanation of how the graph was derived	B						
	Adjust model	3.4 The ability to remove nodes and relationships to adjust AI model	B						
4. individual cases	Informational	4.1 specific application and attribute values	B	✓		✓	✓	✓	✓
		4.2 fairness metric for individual case	B						
		4.3 Level of similarity between cases	B					✓	✓
		4.4 Select specific cases to compare and show which attributes are similar	B			✓		✓	✓
		4.5 Show decision boundaries	B	✓					✓
	Functional	4.6 See 'What If' results on target based on changes to attribute values	L	✓					
5. Model	Informational	5.1 how model works	B						✓
		5.2 how it was created, rationale for decisions in modelling	B						
		5.3 who created it	B						



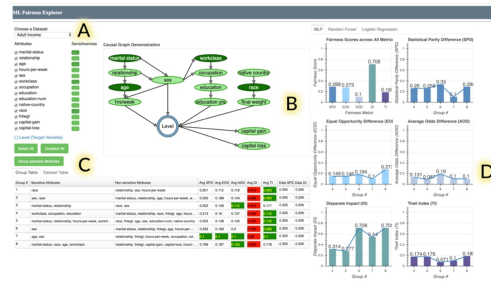
(a) The What-if Tool UI (Wexler et al., 2020)



(b) FairVis UI (Cabrera et al., 2019)



(c) FairSight UI (Ahn & Lin, 2020)



(d) Silva UI (Yan et al., 2020)



(e) Fairness Elicitation Interface (Cheng et al., 2021)

Figure 1. Human-in-the-loop Fairness Toolkits




Figure 2. Workshop causal graph example.

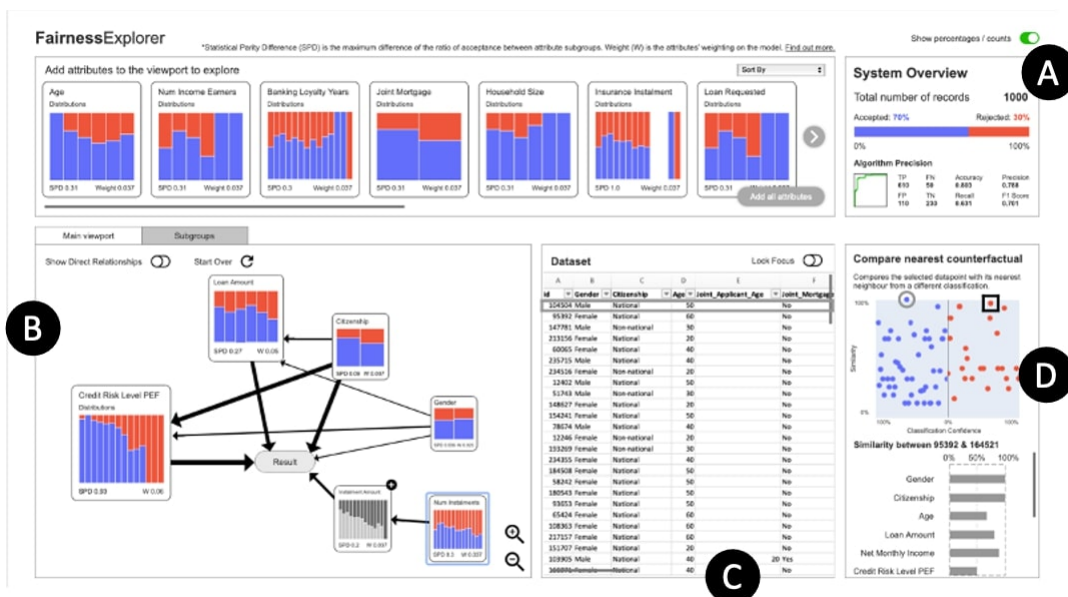


Figure 3. Initial prototype.




Figure 4. Sankey graph showing relationships between nationality and credit level.

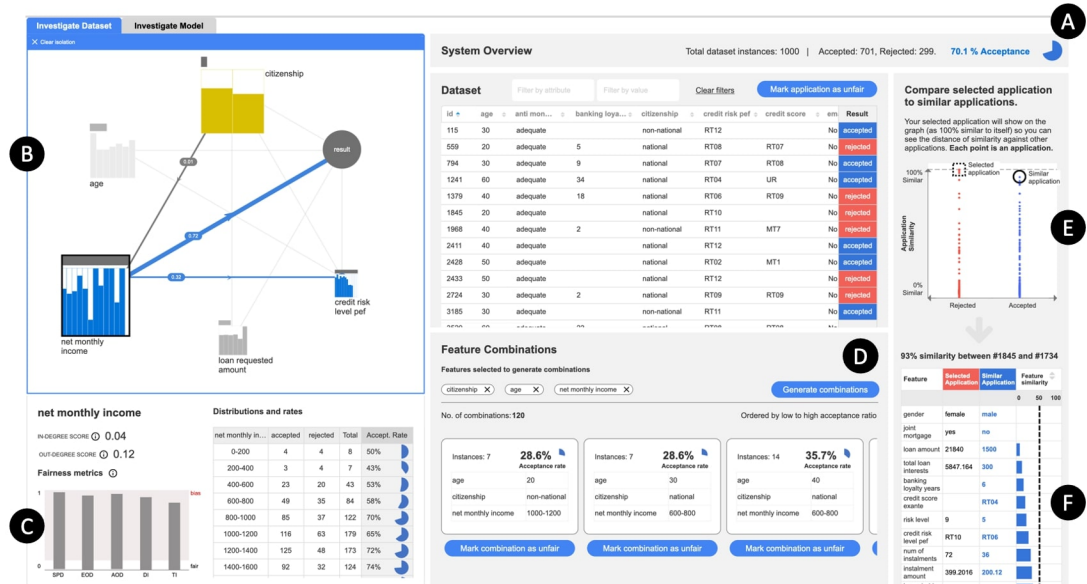


Figure 5. Main UI Components.

Data information

ISP Dataset 2020
1000 instances
26 features

How do you want to measure fairness?

Fairness metrics are statistical methods to measure fairness within datasets and models. Which fairness metric is most appropriate depends on the specifics of your context. Please select the metrics you would like to see in the interface.

Statistical Parity Difference <input type="checkbox"/> Computed as the difference of the rate of favorable outcomes received.	Equality of Opportunity Difference <input type="checkbox"/> Computed as the difference of true positive rates. The true positive rate is the ratio of true positives to the total number of actual positives for a given group.	Average Odds Difference <input type="checkbox"/> Computed as average difference of false positive rate (false positives / negatives) and true positive rate (true positives / positives).
Disparate Impact <input type="checkbox"/> Computed as the ratio of rate of favorable outcome.	Theil Index <input type="checkbox"/> Computed as the generalised measure of benefit for all individuals in the dataset. It measures the inequality in benefit allocation for individuals.	Custom <input type="checkbox"/> Use the metric builder to create your own fairness measure.

Which metric would you like as your main metric for the interface?

Please select metric

Start

Figure 6. Setup: metrics.

[Go back to setup](#)

FairnessExplorer

Metric Builder

Welcome to the metric builder. Here you can create your own metrics which you can then use to investigate fairness within datasets and AI models.

Your selected metrics

- Statistical Parity Difference
- Equality of Opportunity Difference
- Average Odds Difference
- Disparate Impact
- Theil Index

Features

Please give your metric a name

Please give your metric a description

$+$ $-$ \times \div \pm \sim $!$ ∞ $=$ $\frac{x}{y}$ e^x $\sqrt{\quad}$ $\{\}$
 $<$ $<<$ \leq \geq $>>$ $>$ α β γ fraction subscript root brackets

Clear

Create

Figure 7. Metric builder.

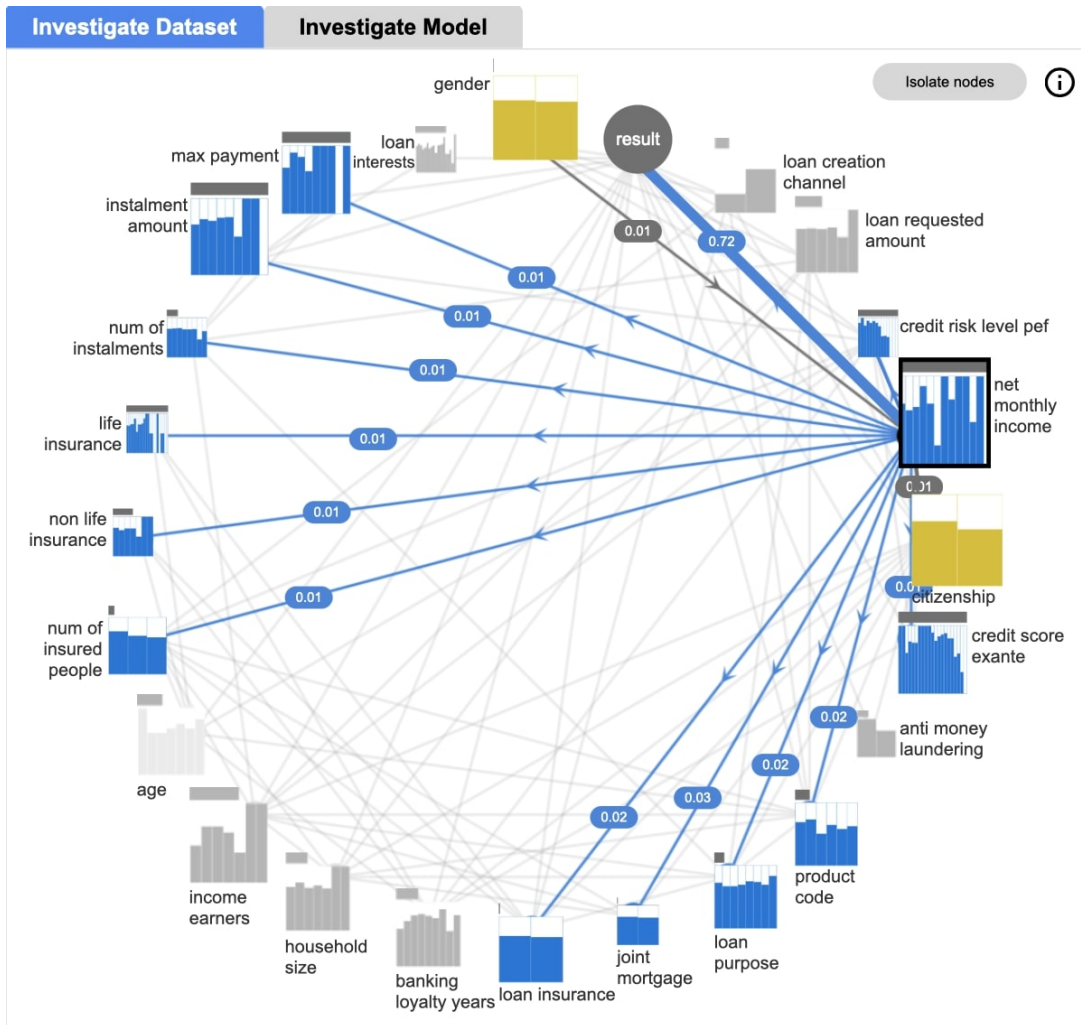


Figure 8. Causal graph in the dataset view showing a selected node: "net monthly income"

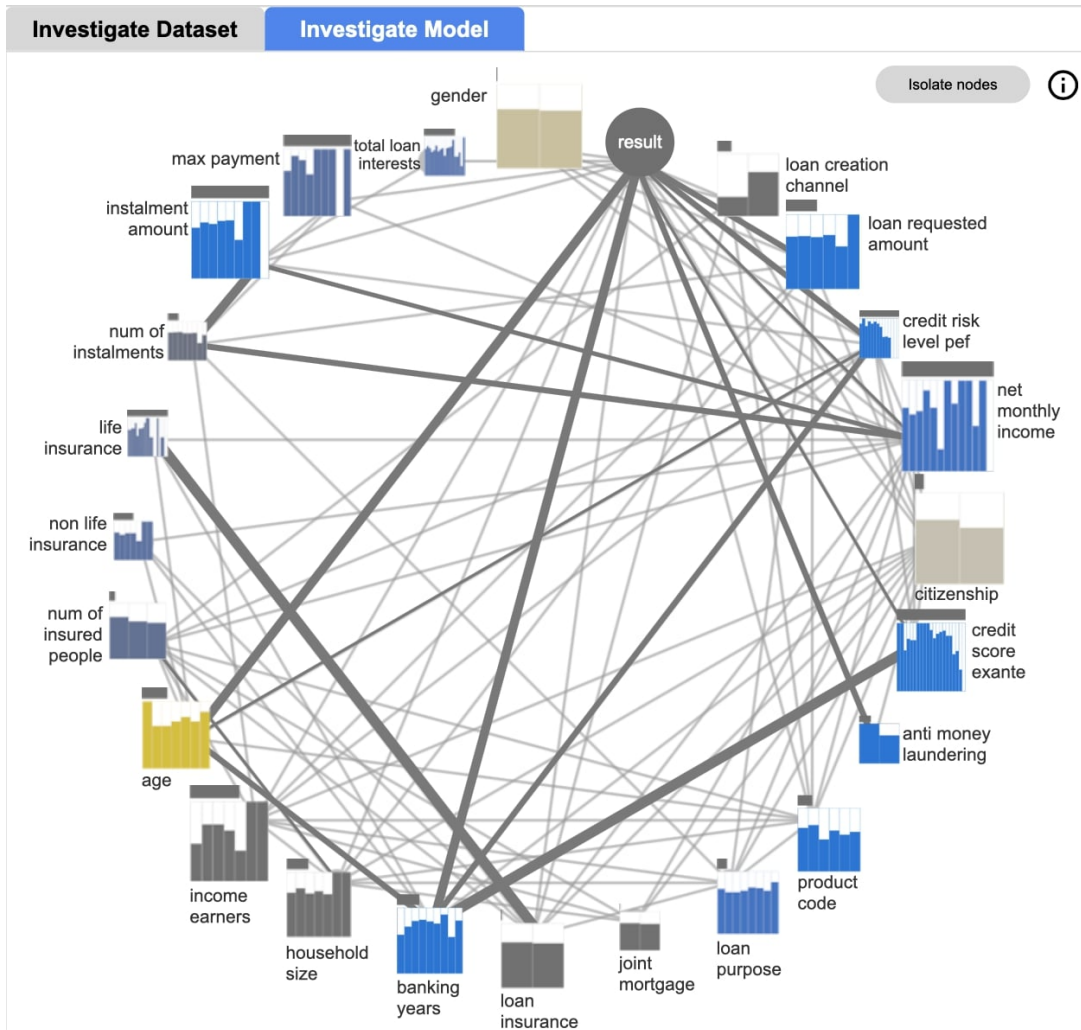


Figure 9. Causal graph in the model view showing feature importance (saturation).



Figure 10. Detailed Information on Features.




Figure 11. Detailed Information on Relationships.

Feature Combinations

Features selected to generate combinations

Generate combinations

No. of combinations: **120** Ordered by low to high acceptance ratio

<p>Instances: 7 28.6%</p> <p style="text-align: right; font-weight: bold;">Acceptance rate</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td>age</td><td>20</td></tr> <tr><td>citizenship</td><td>non-national</td></tr> <tr><td>net monthly income</td><td>1000-1200</td></tr> </table> <p style="text-align: center; border: 1px solid #007bff; border-radius: 15px; padding: 5px; color: white;">Mark combination as unfair</p>	age	20	citizenship	non-national	net monthly income	1000-1200	<p>Instances: 7 28.6%</p> <p style="text-align: right; font-weight: bold;">Acceptance rate</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td>age</td><td>30</td></tr> <tr><td>citizenship</td><td>national</td></tr> <tr><td>net monthly income</td><td>600-800</td></tr> </table> <p style="text-align: center; border: 1px solid #007bff; border-radius: 15px; padding: 5px; color: white;">Mark combination as unfair</p>	age	30	citizenship	national	net monthly income	600-800	<p>Instances: 14 35.7%</p> <p style="text-align: right; font-weight: bold;">Acceptance rate</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr><td>age</td><td>40</td></tr> <tr><td>citizenship</td><td>national</td></tr> <tr><td>net monthly income</td><td>600-800</td></tr> </table> <p style="text-align: center; border: 1px solid #007bff; border-radius: 15px; padding: 5px; color: white;">Mark combination as unfair</p>	age	40	citizenship	national	net monthly income	600-800
age	20																			
citizenship	non-national																			
net monthly income	1000-1200																			
age	30																			
citizenship	national																			
net monthly income	600-800																			
age	40																			
citizenship	national																			
net monthly income	600-800																			

Figure 12. Feature Combinations Component.

Dataset								
Filter by attribute		Filter by value		Clear filters		Mark application as unfair		
id	age	anti mon...	banking loya...	citizenship	credit risk pef	credit score	em	Result
115	30	adequate		non-national	RT12		No	accepted
559	20	adequate	5	national	RT08	RT07	No	rejected
794	30	adequate	9	national	RT07	RT08	No	accepted
1241	60	adequate	34	national	RT04	UR	No	accepted
1379	40	adequate	18	national	RT06	RT09	No	rejected
1845	20	adequate		national	RT10		No	rejected
1968	40	adequate	2	non-national	RT11	MT7	No	rejected
2411	40	adequate		national	RT12		No	accepted
2428	50	adequate		national	RT02	MT1	No	accepted
2433	50	adequate		national	RT12		No	rejected
2724	30	adequate	2	national	RT09	RT09	No	rejected
3185	30	adequate		non-national	RT11		No	accepted

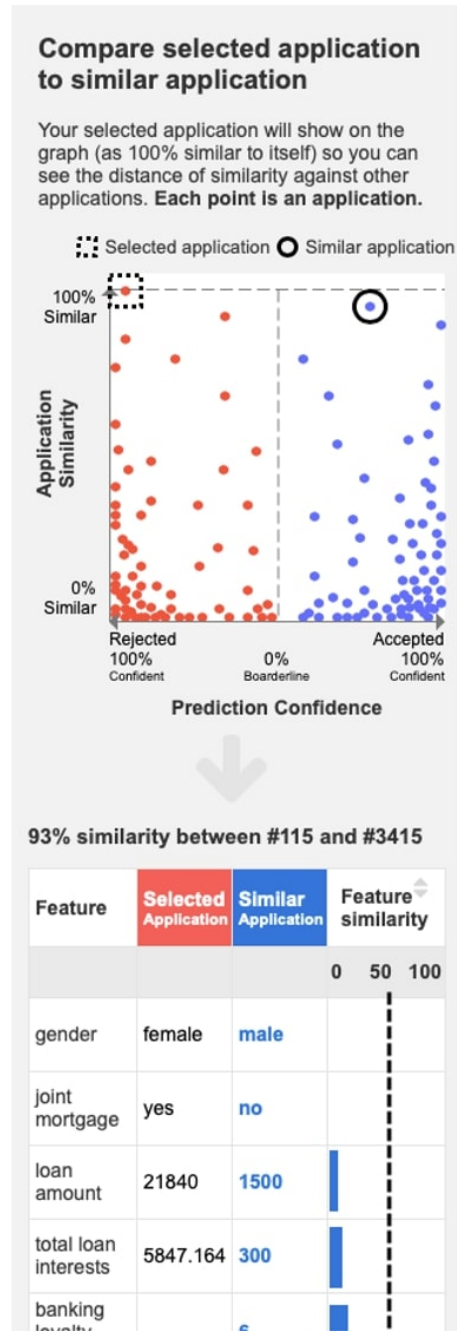
Figure 13. Dataset Component in Dataset View.

Dataset				Clear filters		Mark application as unfair	
Filter by attribute		Filter by value					
id	prediction confidence	prediction result		Feature	Value	Weight	
115	73%	rejected				-1	0
559	61%	rejected					1
794	89%	accepted		age	30		
1241	36%	rejected		anti money laundering	adequate		
1379	52%	rejected		banking loyalty years			
1845	87%	accepted		citizenship	non-national		
1968	72%	rejected		credit risk level pef	RT12		
2411	67%	rejected		credit score exante			
2428	98%	rejected		email contact	No		
2433	91%	rejected		gender	female		
2724	82%	rejected		household size	4		
3185	54%	accepted					

Figure 14. Dataset Component in Model View.



(a) Dataset View.



(b) Model View.

Figure 15. Compare Application component.




Figure 16. Mean Post-session Usefulness Questionnaire Responses. Error bars show standard deviation.




Figure 17. Mean Post-session NASA-TLX Responses. Error bars show standard deviation.

- Figure 1.** Human-in-the-loop Fairness Toolkits
- Figure 2.** Workshop causal graph example.
- Figure 3.** Initial prototype.
- Figure 4.** Sankey graph showing relationships between nationality and credit level.
- Figure 5.** Main UI Components.
- Figure 6.** Setup: metrics.
- Figure 7.** Metric builder.
- Figure 8.** Causal graph in the dataset view showing a selected node: "net monthly income"
- Figure 9.** Causal graph in the model view showing feature importance (saturation).
- Figure 10.** Detailed Information on Features.
- Figure 11.** Detailed Information on Relationships.
- Figure 12.** Feature Combinations Component.
- Figure 13.** Dataset Component in Dataset View.
- Figure 14.** Dataset Component in Model View.
- Figure 15.** Compare Application component.
- Figure 16.** Mean Post-session Usefulness Questionnaire Responses. Error bars show standard deviation.
- Figure 17.** Mean Post-session NASA-TLX Responses. Error bars show standard deviation.

8. Bios

Mr. Yuri Nakao is a Researcher at the Research Center for AI Ethics, Fujitsu Limited, and a Ph.D. student at the Graduate School of Arts and Sciences, University of Tokyo. His research focus is on human autonomy in the development of responsible technology, combining perspectives of Human-Computer Interaction, and Science, Technology and Society (STS).

Mr. Lorenzo Strappelli is a User Experience Designer with a Master of Science in Human-Computer Interaction Design, which he completed at City, University of London. He currently works at the BBC on projects exploring ethical design and has an interest in AI fairness and the challenges of emerging technology.

Dr. Simone Stumpf is a Reader in Responsible and Interactive AI at University of Glasgow, UK, with a research focus on user interactions with AI systems. Her work in Explanatory Debugging for interactive machine learning has provided design principles for enabling better Human-Computer Interaction and AI transparency.

Dr. Aisha Naseer is a Research Manager at Fujitsu Research of Europe Limited UK, leading AI Ethics activities. Dr. Naseer is a founding Editorial Board member of the Springer Journal of AI and Ethics. Recently, she received a global recognition of being nominated and selected as the top "100 Brilliant Women in AI Ethics List 2022".

Dr. Daniele Regoli is a senior Data Scientist at the Big Data Lab of Intesa Sanpaolo. He has a background in theoretical physics, with a Ph.D. in cosmology from the University of Bologna and a Masters in mathematical finance. He worked for several years as an academic researcher, mainly on mathematical models applied to finance and economics.

Ms. Giulia Del Gamba, a law graduate, works at Intesa Sanpaolo as an AI ethicist and legal advisor. In 2017, she worked at the European Data Protection Supervisor. Currently she sits on the Founding Editorial Board for Springer's AI & Ethics Journal and she deals with European Digital Policies.