# IF-City: Intelligible Fair City Planning to Measure, Explain and Mitigate Inequality

Yan Lyu, Hangxin Lu, Min Kyung Lee, Gerhard Schmitt, and Brian Y. Lim

**Abstract**—With the increasing pervasiveness of Artificial Intelligence (AI), many visual analytics tools have been proposed to examine fairness, but they mostly focus on data scientist users. Instead, tackling fairness must be inclusive and involve domain experts with specialized tools and workflows. Thus, domain-specific visualizations are needed for algorithmic fairness. Furthermore, while much work on AI fairness has focused on predictive decisions, less has been done for fair allocation and planning, which require human expertise and iterative design to integrate myriad constraints. We propose the Intelligible Fair Allocation (IF-Alloc) Framework that leverages explanations of causal attribution (Why), contrastive (Why Not) and counterfactual reasoning (What If, How To) to aid domain experts to assess and alleviate unfairness in allocation problems. We apply the framework to fair urban planning for designing cities that provide equal access to amenities and benefits for diverse resident types. Specifically, we propose an interactive visual tool, Intelligible Fair City Planner (IF-City), to help urban planners to perceive inequality across groups, identify and attribute sources of inequality, and mitigate inequality with automatic allocation simulations and constraint-satisfying recommendations (IF-Plan). We demonstrate and evaluate the usage and usefulness of IF-City on a real neighborhood in New York City, US, with practicing urban planners from multiple countries, and discuss generalizing our findings, application, and framework to other use cases and applications of fair allocation.

Index Terms—Fairness, Intelligibility, Explainable Artificial Intelligence, Resource Allocation, Urban Planning

## **1** INTRODUCTION

The increasing pervasiveness of Artificial Intelligence (AI) has raised concerns of algorithmic bias, unfairness and discrimination, since AI benefits some segments of society while disadvantaging others [1]. Current methods to assess fairness in model outcomes and decisions focus on visualizing tabular data with bar charts and scatter plots [2], [3], which are well-suited for data scientists and machine learning engineers. However, assessing fairness requires domain experts and stakeholders to consider relevant factors and balance multiple criteria [4]. For example, assessing fairness in urban planning needs to satisfy different residents with diverse needs, and requires domain-specific visualizations, such as maps with 3D buildings [5], [6]. Hence, instead of having just data scientists assess model fairness, fairness visualizations should be integrated into domain-specific tools to involve other stakeholders to participate collaboratively [7], [8].

Such participatory design will require users to understand how the intelligent system determines fairness in order for them to mitigate unfairness. Fortuitously, explanations are an effective way for users to understand and trust system outcomes [9], [10], [11]. Recently, many explainable AI (XAI) techniques have been developed, such as attribution [12], [13], contrastive [14] and counterfactual [15], [16] explanations. This motivates us to apply XAI techniques to help planners to understand how fairness is computed and how to improve fairness in their planned solutions.

Much recent progress on fair machine learning focuses on algorithmic bias in inference-based decisions [17], [18], yet there is also a strong need for fairness in resource allocation for domains such as urban planning, organizational management, public services, and healthcare [7]. Automatic algorithms can find optimal allocations, but they may overlook nuances and normative decision points that are part of complex real-world systems [8]; instead, the participatory involvement of stakeholders and domain experts is necessary in fair allocation decisions [7].

Hence, we propose the Intelligible Fair Allocation (IF-Alloc) Framework that leverages explainable AI to help users assess fairness, identify sources of inequality, and mitigate inequality. Specifically, we adopt Lim and Dey's user-centered intelligibility framework of question types [19] to satisfy user reasoning needs for causal attribution (Why), contrastive (Why Not) and counterfactual reasoning (What If, How To) [11], [14]. We use the Generalized Entropy Index as the fairness metric and leverage its additive decomposability to attribute inequality to constituent components.

We applied the framework to develop the Intelligible Fair City Planner (IF-City) for fair urban planning to balance the benefits from building locations across residents with different preferences. For example, having large parks in the city is valuable to nature lovers, but may cause schools or offices to be placed farther away and disadvantage students and office workers, respectively. IF-City is an urban planning visualization and 3D model design tool that helps urban planners to design fairer cities by placing buildings on a map, simulating the allocation of the resident population

<sup>•</sup> Yan Lyu is with the School of Computer Science and Engineering, Southeast University, Nanjing, China. E-mail:lvyanly@gmail.com.

<sup>•</sup> Hangxin Lu and Gerhard Schmitt are with Future Cities Laboratory, Singapore-ETH Centre, Singapore, and Department of Architecture, ETH Zurich, Switzerland. E-mail: luhangxin@gmail.com, schmitt@arch.ethz.ch.

Min Kyung Lee is with the School of Information, The University of Texas at Austin, US. E-mail: minkyung.lee@austin.utexas.edu

Brian Y. Lim (Corresponding) is with the School of Computing, National University of Singapore, Singapore. E-mail: brianlim@comp.nus.edu.sg.

IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS



Fig. 1: Dashboard user interface of IF-City for fair urban planning with main 3D map view, side bar with several indicators, and pop-up controls. Urban planners can add or edit buildings of different function types, and evaluate the impact on resident benefits and inequality. Users can also query for recommendations to improve fairness and save design iterations.

with iterative proportional fitting (IPF) [20], and computing the fairness indicator for the urban plan based on accessibility to and preferences for different amenities. IF-City helps users to understand the fairness calculation by visualizing where benefits were unfairly allocated to different resident types across locations (Why). Urban planners can update the urban design to explore how inequalities can be reduced (What If), and compare between design iterations (Why Not). IF-City also provides recommendations for editing the quantity and distribution of amenities to optimize fairness (How To), and predicts their partial contributions, as calculated by Shapley values, towards improving overall fairness (Why of How To). In summary, our contributions are:

- Intelligible Fair Allocation framework (IF-Alloc) to perceive fairness, explain causes and mitigate inequality for fair urban planning.
- Intelligible Fair Planning recommendation (IF-Plan) technique to generate recommendations for design edits to conveniently improve fairness in urban planning.
- Intelligible Fair City Planner (IF-City), an interactive visualization tool for urban planners to design fair neighborhoods by iteratively assessing and improving fairness. It simulates the allocation of a resident population, computes an inequality indicator, explains possible causes visually based on resident type and location.

We demonstrated IF-City with a use case for designing a neighborhood in New York City to improve fairness, and evaluated it in a formative study with practicing urban planners and urban designers to investigate how the explanation features are used to understand and improve fairness in urban designs. We conclude with a discussion to generalize IF-City and the IF-Alloc framework.

#### **RELATED WORK** 2

We introduce notions of fairness for predictive and allocation systems, and relate them to urban planning. We then describe visualization tools to assess and improve fairness, explain data and models, and analyze and plan cities.

#### Fairness Definitions and Metrics 2.1

## 2.1.1 Metrics of Fair Outcomes

Recently, research on bias and fairness in machine learning has been very active [18]. Such works frame fairness by individual or group fairness. Individual fairness requires that similar individuals should have similar outcomes (e.g., people with similar credit scores should have similar loan approval chances). However, defining what makes individuals similar may be difficult. Instead, individuals can be grouped based on having similar attributes or demographics. In socially-sensitive applications, groups are defined by protected attributes like age and gender. Group fairness requires that different groups have similar outcomes (e.g., younger and older workers with similar other attributes should have similar hiring chances). Many fairness metrics (e.g., statistical parity [21], equal opportunity [22], conditional equality [23], intersectional bias [3]) have been proposed to evaluate the balance in decision outcomes. In this work, we focus on supporting fairness across groups, but focus on resource allocation instead of predictive decisions.

## 2.1.2 Metrics of Fair Allocation

Unlike algorithmic fairness that considers how fair a prediction result is, fair resource allocation seeks to allocate limited resources fairly to individuals or organizations with

varying demands or needs [7], [24], [25], [26]. Metrics to measure allocation fairness include Jain's index [27], Max-min/min-max [28], proportional fairness [29], and entropy [30]. Other popular metrics measure inequality as opposed to fairness, especially to quantify income inequality. For example, the Gini coefficient [31] and Hoover index [32] determine inequality by measuring the deviation of the income distribution of the Lorenz curve from the diagonal line of perfect equality [31]. These metrics provide a single measure to indicate the inequality in a population, but have limited interpretability to discern whether the inequality comes from high- or low-income groups. The Atkinson index [33] includes an inequality aversion parameter  $\varepsilon$  that can be used to increase sensitivity towards changes in low incomes. However, for populations with diverse subgroups, this metric does not explain which groups are most or least disadvantaged. This requires a metric to be *additively decomposable* to attribute inequality to specific groups. One such metric is the Generalized Entropy (GE) index [34], [35], derived from information theory to measure redundancy. It is often used to measure the diversity of incomes. It also has a sensitivity parameter  $\alpha$  to adjust sensitivity towards low or high incomes. Special cases of the GE index are the mean log deviation ( $\alpha = 0$ ), Theil index ( $\alpha = 1$ ), and half of the squared coefficient of variation ( $\alpha = 2$ ). Although other inequality measures such as Variance of Logarithms and Coefficient of Variation are also decomposable [36], GE index has been applied to evaluating the fair allocation of urban resources, including healthcare [37], transportation [38], and ecosystem [39] services. In this work, we leverage the additive decomposability of the GE index to satisfy a common axiom in explainability [40] to attribute inequality to specific groups with between-group and within-group inequalities, on their accessibility to different building types. Hence, from these decomposed inequalities, the distribution of benefits can be compared across different groups.

## 2.1.3 Fair Allocation in Urban Planning

The fair access to scarce public resources, such as parks, hospitals and schools, is a paramount goal in urban planning. Urban and transportation planners typically examine the accessibility to these amenities by the distance from housing to the amenity. Accessibility differs by transport modes, including walking [41] and public transit [42], and can be calculated using various metrics [43]. While visualizing accessibility can indicate if benefits are unequally distributed, it is difficult to see which regions have more inequality. Hence, several methods have proposed measuring various inequality metrics, such as the Gini coefficient [44], and local indicators of spatial autocorrelation (LISA) [45]. However, these inequality metrics are typically reported for each city globally, rather than by sub-regions, though some recent works visualize them geospatially (e.g., [44]). In this work, we visualize housing and amenity locations, and the accessibility and local inequality of each housing location.

## 2.2 Visualization and Analytical Tools

## 2.2.1 Visualization for Algorithmic Fairness

Many visualization tools have been developed to examine fairness or bias in machine learning. Commercially available

ones include IBM AI Fairness 360 [46], Google Tensorflow Fairness Indicators [47], and Microsoft Fairlearn [48]; these provide basic bar charts and scatter plots to show predictive parity in datasets and model predictions. More sophisticated visualization interfaces help to examine other forms of unfairness, such as intersectional bias in FairVis [3], and causal fairness in Silva [49]. However, these only provide a high-level, global view of the model fairness, and do not identify which subgroups are particularly unfairly treated. DiscriLens [4] can identify itemsets that are discriminatory using rule mining, though these data-driven sets may not align with domain-specific groups and be less interpretable. FairSight [2] can examine individual and group fairness, and explain the influence of each feature based on feature perturbation. This explainability method is common for black boxes, but suffers from approximation errors. In this work, we visualize inequality across subgroups in a domainspecific visualization for urban planning, and explain the sources of inequality based on the additive decomposability of the Generalized Entropy Index, which is a white box method that is deterministically calculated. Though, we compute feature attributions for our planning recommendations using Shapley values, which is similar to perturbation.

#### 2.2.2 Visualization for Fair Resource Allocation

Unlike research on visualizing unfairness in datasets and predictive models, visualizations of fair resource allocation are sparser. AlgoCrowd [50] visualized the AI-driven matching between workers and tasks in a dashboard, showing the distributions of worker reputation and productivity, the Jain fairness index over rounds of task allocation, and an argumentation-based explanation of why each worker was allocated. VisMatchmaker [51] supports the comparison of trade-offs between two solutions for a job allocation task using novel but unfamiliar visualizations (number lines and glyphs). To promote adoption, it is preferable to use visualizations that are familiar to each domain. Talen visualized the distribution of accessibility using geographical maps that are familiar to urban planners [52]. Similar to [51], FairVizARD [53] enables comparing the outcomes between two matching algorithms for ride-sharing, visualizing with a map view of taxi and request locations, and a graph view showing the time series variation of several indicators. In IF-City, we leverage well-known geographical map, heatmap, and bar chart visualizations to convey information about inequality, benefits, and accessibility. Furthermore, while the aforementioned visualizations support fairness assessments, they do not directly support mitigating unfairness through redesigns, which IF-City does.

## 2.2.3 Visualizations for Model Explainability

Other than being fair, intelligent systems need to be understandable, thus there has been significant recent research on explainable AI (XAI) [9], [10]. Drawing from psychology and philosophy, Miller [14] argued that explanations should be causally attributional, contrastive, and counterfactual. These correspond to explanations of feature attributes [12], [13], contrastive explanations [54], [55], and counterfactual explanations [15], [16], [56], [57]. Other explanations leverage sophisticated visualizations, such as partial dependence

IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS

plots [58], [59], network activations [60], [61], network summaries [62], saliency maps [63], feature visualizations [64]. The large variety of XAI techniques makes it hard for developers to prototype with them. ExplAIner [65] proposed a framework to unify how XAI techniques are used, and defined a pipeline to understand, diagnose and refine models.

Some research has focused on defining workflows and developing rich dashboards combining multiple visualizations for explanations. Krause et al. [66] visualized the model performance on groups of instances with similar feature importance. To make sense of the high number of features or patterns in models, matrix layouts have been used to visualize extracted rules with RuleMatrix [67] and important features in random forests with ExMatrix [68]. Supporting interactivity by investigating counterfactual cases has been particularly popular. Prospector [59] visualizes with partial dependence plots and sliders, allowing users to investigate how predictions change with different feature values. Gamut [69] provides similar capabilities for generalized additive models (GAM). DECE [56] provides an interactive visualization to refine counterfactual queries for subgroups of instances. The What-If Tool [16] provides visualizations of confusion matrices, scatter plots, histograms for iterative counterfactual explorations. These tools were designed primarily for data scientists and engineers familiar with statistical concepts and graphs. For domain experts or lay users, simpler interactions will be more accessible. To support the user-centered design of explanations, several frameworks have been proposed based on the human reasoning process [11], and patterns of user in inquiry [19], [70], [71]. In this work, we implemented Lim and Dey's intelligibility taxonomy of question types [19] to understand and improve fairness. While some XAI tools can be used to investigate AI fairness (e.g., with the What-If Tool [16]), they only support viewing predictive parity or equality [72], and not fair allocation, which is our focus in this work.

## 2.2.4 Visualization for Urban Analysis and Planning

To understand complex urban environments, many visualizations have been developed to view spatio-temporal movement patterns [73], [74], [75], [76], [77], [78], transit networks [79], air pollution [80], activities [81] and events [82]. But visualizing inequality in cities has been rudimentary with locations and simple metrics on maps [44], [52].

Urban planning requires more sophisticated interactive visualizations to design the urban layout, simulate stakeholder and resident behaviors in the urban environment, and visualize various urban indicators (metrics) [83]. Popular tools used by professional urban planners, such as ArcGIS [84] and UrbanSim [83], use 2D and 3D maps as the main view. Recently, tools have been designed to inform and elicit feedback from collaborators and stakeholders [85], [86]. These tools calculate urban indicators to visualize geospatially, and can provide AI suggestions to optimize planning. However, they currently do not allow planners to assess, analyze, and mitigate inequality. In this work, we extend the visual paradigm of urban planning tools to support these new capabilities.

# **3** INTELLIGIBLE FAIR ALLOCATION FRAMEWORK

We focus on designing an urban plan to be more fair to diverse groups of residents. This can be treated as a fair allocation problem to allocate building resources such that different residents equally benefit despite differing preferences. We consulted with 5 urban planning experts in an urban planning research lab (two are our coauthors) to propose an iterative 3-step framework for fair urban planning: i) perceive the fairness of the design, ii) identify causes for the unfairness, and iii) mitigate the inequality by identifying opportunities for change. Repeat as necessary. This differs from prior fairness visualization approaches for data scientists (e.g., [2]) by focusing on fair resource allocation and targeting for use by urban planners. Our application towards urban planning presents two further challenges that our approach addresses: 1) Planning is an iterative process where urban planners trial each design, evaluate the outcomes with a simulation, and redesign newer plans. 2) Users are domain experts (urban planners), not technical experts (data scientists), so the system predictions need to be explainable and intuitive. This clarifies the need for *intelligible fairness* in urban planning.

## 3.1 Design Requirements

From discussions with our domain experts, we identified design requirements for base urban planning and for intelligible fair design. To support urban design, our visualization should support navigating on a map, viewing building heights and functions, and adding, editing or deleting buildings. Thus, basic requirements are to provide a 3D map of buildings, colors for different building functions, drawing facilities, and popups to change building parameters.

To display and explain fairness indicators in urban design, we drew inspiration from the Intelligibility Taxonomy of Lim and Dey [19]. It defines explanations for intelligent systems for question types (e.g., Why, Why Not, What If, How To), and has been effective in providing useful explanations for non-technical users (i.e., not data scientists). We identified requirements for Intelligible Fair (IF) allocation:

- **IF.1 Inequality Indicator**. Provide a quantitative indicator of *What* inequality score was calculated to represent the state of the urban design.
- **IF.2 Inequality Attribution by Resident Type**. Explain *Why* the inequality was high/low by showing which resident types were advantaged or disadvantaged. This helps users know who to satisfy more/less.
- **IF.3 Inequality Attribution by Location**. Explain *Why* inequality was high/low by showing which locations had higher or lower benefit than average. This helps users know where to change buildings to improve access to amenities.
- **IF.4 Inequality Trace**. Explain *Why* each location had high (or low) benefits by tracing the calculation from resident preferences of and accessibility to each amenity.
- **IF.5 Population Simulation to Allocate Residents.** Simulate the outcome (*What If*) of the inequality with iterative design changes. For each urban plan redesign, this requires calculating i) the allocation of residents based on diverse preferences, and ii) the consequent benefits and reporting the resulting inequalities.

#### IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS



Fig. 2: IF-Alloc Framework to perceive fairness, identify causes, and mitigate inequality for fair urban planning. On perceiving fairness by viewing the Inequality score (Step 1), the user flow to identify the causes for inequality starts with noting uneven benefits across resident types (2), seeing how differently benefited one group is (2a), and noting its large inequality (2b). Next, identify the location of inequality (3) by either filtering for buildings with higher occupancy and deviating benefits (3a), or locating dark blocks and buildings in the relative benefits heatmap (3b). Then trace the cause for low or high benefits due to accessibility to various amenities and resident preferences (4). The user flow to mitigate inequality iteratively has two approaches — through simulation trials by adding or editing buildings (5), or using the recommendation system (6) and editing buildings to satisfy the recommendations. The user can track improvements in inequality through the design iteration timeline (7), and iteratively improve the design (back to Step 1).

- IF.6 Intelligible Fair Planning Recommendation. Recommend How To change the urban plan to improve fairness, informing how much improvement to expect (prediction of What If) and Why each change is beneficial *(attribution towards changes in inequality). This helps* the user know where to make changes, which building functions to change, by how much floor area. Instead of manually and tediously trying designs, this accelerates designs for the user to try. To support design priorities, recommendations should be constrained by construction budgets, height limits, To retain prior benefits, this should also limit whether benefits to specific resident types should only go up, go down, or stay the same.
- IF.7 Design Iteration Timeline. Provide a timeline to compare between design iterations, showing changes in various indicators (e.g., Inequality, Population, Planning Area). Users should also be able to load previous urban plans for comparison (Why Not).

These requirements support the goals for intelligible fairness in urban planning to perceive fairness (IF.1, IF.7), identify causes (IF.2, IF.3, IF.4), and mitigate inequality (IF.5, IF.6).

## 3.2 IF-Alloc Framework

Based on the design requirements, we introduce the IF-Alloc framework (Fig. 2) for Intelligible Fair Allocation in urban planning to: I) perceive unfairness with the Inequality Indicator; II) *identify causes* of inequalities by attributing towards resident types (2) based on uneven benefits (2a, 2b), or location (3) by filtering problematic areas (3a) or viewing

heatmaps (3b), and tracing to root causes in accessibility being too high or poor (4); III) mitigate inequality by manually editing buildings (5) or following automatic fair planning recommendations (6), and reviewing changes in inequality iteratively in a timeline (7). The user can iteratively repeat these steps towards a desired fairness level.

#### 4 **IF-CITY: TECHNICAL APPROACH**

We describe our technical approach (Fig. 3) to implement IF-Alloc for fair urban planning in the Intelligible Fair City Planner (IF-City). Specifically, we calculate: 1) the accessibility to buildings of various amenities, 2) the benefit of each resident based on their preferences for different amenities and their accessibility, and 3) the inequality based on inputting the benefits into the Generalized Entropy Index. The inequality is additively decomposable by groups of residents and location. Computing these indicators first depends on a predetermined building layout, then the subsequent allocation of residents, which we describe in Subsection 4.4. Finally, we describe our recommendation system for fairer edits to accelerate the design iteration.

## 4.1 Accessibility to Building Type (IF.4)

In urban planning, buildings are defined by their functions or amenities. Residential buildings provide homes for people; commercial buildings support commercial activities like retail and restaurants; educational buildings include schools and libraries; cultural buildings promote culture in museums and theaters; and parks support outdoor recreation. These are defined as building **function types** *f*.

IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS



Fig. 3: System overview of data preparation and calculations. Map and building data (1a) are used to extract building function types (1b) and their existing locations. User trajectories (2a) are extracted from location-based social media data and clustered (2b) to identify resident types (2c) and their average preferences for different building functions (2d). The resident allocation (3) simulates the migration of residents into the designed city. The benefits for the population are calculated by resident type (4a) and location (4b), which are used to calculate the group inequality for each resident type (4c). To support understanding the fairness indicator (4c), the user can view which resident types had unfair benefits (4a) or where benefits were unfairly located (4b), redesign the urban plan to change benefits and reduce inequalities through direct manipulation (5a) or run the recommendation to receive fairer design proposals (5b).

Accessibility is the ease of reaching spatially dispersed resources with a certain function. It depends on the distance to and the amount of each resource. The farther a building is from the resident, the lower the accessibility. The larger the floor area of the function in the building, the higher the accessibility (e.g., adding offices will raise the likelihood of having a job nearby). We compute accessibility with the gravity model [87] which treats resources farther away as having inversely smaller effects. For a specific function type f, the accessibility at a residential building *location* l is:

$$a_{l,f} = \sum_{\ell \in L} \frac{v_{\ell,f}}{\rho_f} e^{-\kappa_f d_{l,\ell}},\tag{1}$$

where  $\ell \in L$  is a nearby building,  $v_{\ell,f}$  is its floor area for function type f,  $d_{l,\ell}$  is the distance of  $\ell$  to residential building l,  $\kappa_f$  is the impedance to travel for the function f[88], and  $\rho_f$  is the planning priority weight for function type f to prioritize for equity. For example, residents typically need more space for a park than for retail for similar utilities, so the priority weight of parks should be higher.

#### 4.2 Resident Preference, Utility, and Benefit (IF.4)

Cities attract diverse residents who have various interests or **preferences**. Even in the same residential building, different residents will value their location differently, i.e., they will have varying **utilities** at their location due to differing preferences to nearby amenities. For example, placing a large park near a residential building will be more appreciated by residents who love outdoor activities than those who prefer cultural facilities. We calculate the utility of individual resident *i* living at residential building location *l* as the preference-weighted sum of accessibility, i.e.,

$$u_{i,l} = \sum_{f \in F} \pi_{i,f} a_{l,f},\tag{2}$$



Fig. 4: Group benefit calculation for one example resident type. For a residential building (1a), within a distance threshold radius (1b), measure the floor area of each building type (1c), and multiply by decaying distance to calculate the accessibility (1d). For a resident type (2a), multiply their preferences (2b) with accessibility to get the weighted accessibility (2c). For residents of the type in the residential building (3a), for each person (3b), calculate the mean weighted accessibility to get the total utility (3c). The difference from the original utility is the benefit (3d).

where *F* is the set of non-residential building function types, and  $\pi_{i,f}$  is the resident's preference for each function type *f*. The preferences can be acquired subjectively using preference elicitation surveys [89], or implicitly with objective data-driven measures from their visit trajectories in location-based social media [90]. In IF-City, we obtained preferences with the latter approach (details in Appendix C).

Since residents would have had a prior residence before being placed in the new urban design, they would have a prior utility  $u_i^{(0)}$  that is non-zero. This can be estimated from location-based activity data. Hence, we seek to balance the change in utility, which we denote as **benefit**, i.e.,

$$b_{i,l} = u_{i,l} - u_i^{(0)}.$$
(3)

We assume that residents would not relocate to the new urban design if their benefits were negative. Hence, allocations will not be made for such cases, and we can balance the benefits instead new utility scores.

Balancing resident benefits improves equality, but some groups, such as the elderly, children, and low-income workers, may still be disadvantaged and need further help. This would improve fairness through *equity*. To support this, we define a *population priority weight*  $\rho_g$  for resident type g of resident i to prioritize or penalize his or her benefit, i.e.,

$$b_{i,l} \leftarrow \rho_g b_{i,l},\tag{4}$$

where  $\rho_g \ge 0$ . Setting  $\rho_g < 1$  for a disadvantaged resident type will make the effective benefit lower. The weighted benefits will then be used to the calculations for inequality.

Fig. 4 summarizes the steps to calculate resident benefit from accessibility and preferences. These satisfy requirement IF.4 (Inequality Trace). Next, we calculate a fairness score for all residents from their benefits.

© 2023 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS

## 4.3 Generalized Entropy for Fairness Attribution (IF.2)

We employ the Generalized Entropy (GE) index [35] to calculate the extent that the residents' benefits are unevenly distributed. It is popularly used to measure income inequality and is the generalization of information entropy [34], [35]. The GE index of the set of benefits *B* of all residents is

$$\varepsilon^{\alpha}(B) = \frac{1}{n\alpha(\alpha - 1)} \sum_{i=1}^{n} \left[ \left( \frac{b_i}{\overline{b}} \right)^{\alpha} - 1 \right], \tag{5}$$

where  $b_i$  is the benefit of individual resident i, b is the global mean benefit, n is the number of residents (population size), and  $\alpha$  is a sensitivity parameter,  $\alpha \in \mathbb{R}$ . We choose  $\alpha = 2$ , for  $\varepsilon^{\alpha}$  to be more sensitive towards individuals with smaller benefits (i.e., disadvantaged groups). The GE index has nice explainable properties of being additively decomposable.

Additive. The GE index is a simple sum of terms for each individual's benefit  $b_i$ , so decomposing it into all individuals can identify which individuals have above or below average benefit, and hence affect inequality.

**Subgroup decomposability**. By partially aggregating GE index inequalities of several individuals, we can calculate the subgroup inequality of non-overlapping groups g within the population. To aid interpretability, these terms can be divided into between-group inequality  $\varepsilon_b^{\alpha}(B)$  and within-group inequality  $\varepsilon_w^{\alpha}(B)$ . The total inequality is also conveniently just the sum of these components, i.e.,

$$\varepsilon^{\alpha}(B) = \varepsilon^{\alpha}_{b}(B) + \varepsilon^{\alpha}_{w}(B). \tag{6}$$

The between-group inequality measures how different the mean benefits of different groups are, i.e.,

$$\varepsilon_b^{\alpha}(B) = \sum_{g \in G} \frac{n_g}{n\alpha(\alpha - 1)} \left[ \left( \frac{\bar{b}_g}{\bar{b}} \right)^{\alpha} - 1 \right], \tag{7}$$

where  $\bar{b}_g$  is the mean group benefit,  $n_g$  is the number of individuals in group g, and G is the set of all groups.

The within-group inequality measures how different the individual benefits within each group are, i.e.,

$$\varepsilon_w^{\alpha}(B) = \sum_{g \in G} \frac{n_g}{n} \left(\frac{\bar{b}_g}{\bar{b}}\right)^{\alpha} \varepsilon^{\alpha}(B_g), \tag{8}$$

where  $B_g$  is the set of individual benefits of residents in group g. We demonstrate how total inequality, between- and within-group inequalities can vary with different benefit distributions and group divisions in Appendix B. In IF-City, we grouped residents with similar preferences using k-means clustering on their frequency of visiting various building types (details in Appendix C). These calculations identify which resident types raised or lowered inequality, and satisfy IF.2 (Inequality Attribution by Resident Type).

## 4.4 Population Simulation to Allocate Residents (IF.5)

Changes in an urban plan affects the location l where each resident i will live, thereby affecting the benefit and inequality. We determine these locations by solving a population allocation problem based on the user's urban plan. This satisfies IF.5 (Population Simulation to Allocate Residents).

We first estimate the marginal probability  $p_i$  of resident *i* moving to the city.  $p_i$  should be monotonic to the average

benefit  $\bar{b}_i = \sum_{l \in L_{\text{Res}}} b_{i,l} / |L_{\text{Res}}|$  of the resident living in all residential building locations  $L_{\text{Res}}$  in the city, where  $b_{i,l}$  is the resident's benefit of living at building *l*. We normalize this as a probability with a tanh transform function, i.e.,

$$p_i = \max(\tanh(\gamma b_i), 0), \tag{9}$$

7

where  $\gamma > 0$  is a normalization parameter that is determined by the constraint that the total probability of all residents should be equal to the sum of occupancy of all residential buildings in the city, i.e.,  $\sum_i p_i = \sum_l o_l$ , where  $o_l$ denotes the occupancy of residential building l. Note that  $\bar{b}_i$ could be negative because the resident may not benefit from moving to the city. In this case, the probability  $p_i$  is 0.

Given the marginal distribution  $p_i$  for the resident *i* for the whole city, we seek to estimate probability  $p_{i,l}$  of the resident moving into each specific building *l*. We perform iterative proportional fitting (IPF) [20], [91] to estimate  $p_{i,l}$ by treating  $\{p_{i,l}\}_{I,L}$  as a matrix of rows *i* and columns *l*. At each iteration step  $\eta$ , for each residential building location *l* with vacancy, IPF estimates  $p_{i,l}^{\eta}$  with the following two equations for the odd and even steps, respectively:

$$\hat{p}_{i,l}^{(2\eta-1)} = \frac{\hat{p}_{i,l}^{(2\eta-2)} p_i^{(2\eta-2)}}{\sum_l \hat{p}_{i,l}^{(2\eta-2)}}, \quad \hat{p}_{i,l}^{(2\eta)} = \frac{\hat{p}_{i,l}^{(2\eta-1)} o_l^{(2\eta-1)}}{\sum_i \hat{p}_{i,l}^{(2\eta-1)}}, \quad (10)$$

where  $p_i^{(2\eta-2)} = \sum_l p_{i,l}^{(2\eta-2)}$  and  $o_l^{(2\eta-1)} = \sum_i p_{i,l}^{(2\eta-1)}$  are the marginal distributions across individuals and locations, respectively. The iterations terminate when  $p_{i,l}$  converges to a small threshold. The resident *i* is thus allocated to building *l* with the probability  $p_{i,l}$ . Residents are allocated in a random order until all buildings are fully occupied.

IPF can be generalized to consider other real-world factors, such as diverging preferences among residents at the household level [92], and resident income and property value [92], [93].

#### 4.5 IF-Plan: Fair Planning Recommendation (IF.6)

While population allocation simulation helps to provide feedback on the user's design, it can be tedious for users to iterate designs on their own, leading to less productive trialand-error. To speed up the design iterations, we provide a Fair Planning Recommendation system to satisfy IF.6.

Our approach to generate the counterfactual explanations is distinct from current approaches in machine learning by: 1) recommending a numeric outcome and optimizing its value, rather than an alternative target label for categorical classification [15]; 2) supporting userconstrained recommendations by solving a constraint optimization problem, rather than using a slow brute-force search method [57]; 3) recommending in terms of coarser user-manageable features, rather than raw fine-grained features; and 4) indicating the importance of each change by attributing their contribution towards the improved outcome using Shapley value calculations (inspired by SHAP [12]). This is the first work to formulate *fair* urban planning as an optimization problem. We propose a heuristic algorithm to approximate the optimal fair urban design by adapting the Frank-Wolfe conditional gradient method [94].

## 4.5.1 Problem Formulation

An urban design is the set of floor areas for all buildings of different function types in the city. Given the current design, the user's goal is to propose a design change that decreases the inequality index. Instead of recommending changes to specific buildings, we recommend changes to a coarser unit of geography — the census block  $k \in K$ . This reduces the search complexity for an optimal solution, aligns with planning practice to consider planning at coarser granularity, and allows for more freedom in designing specific buildings.

Let  $\delta_{k,f} = \Delta v_{k,f}$  denote the floor area change of census block k and building function type f,  $v_{k,f}$  denotes the floor area of existing block k and type f. The new design  $v_{k,f} + \delta_{k,f}$  determines the allocation of residents (i.e.,  $p_{i,l}$ ) and thus a distribution of benefits B for all the allocated residents. We seek the optimal  $\delta_{k,f}$ , such that the inequality index  $\varepsilon^{\alpha}(B)$ defined by Eq. 5 is minimized, i.e.,

minimize 
$$\varepsilon^{\alpha}(B)$$
 (11)

subject to 
$$\delta_{k,f} + v_{k,f} \ge 0,$$
 (12)

$$\sum_{f} \left( \delta_{k,f} + v_{k,f} \right) / s_k - \bar{h}_k \le h_{\Delta}^{(\max)}, \qquad (13)$$

$$\sum_{k} \sum_{f} |\delta_{k,f}| \le \delta^{(\max)}, \tag{14}$$

$$\left|\sum_{k} \delta_{k, f=\text{Res}}\right| \le \delta_{f=\text{Res}}^{(\text{max})}.$$
(15)

Eq. 12 constrains that the final floor area cannot be negative, since that is physically impossible. Eq. 13 constrains the increase in average height to less than a small threshold  $h_{\Delta}^{(\max)}$ , where  $s_k$  is the total building footprint (ground floor area) in the block k and  $\bar{h}_k$  is the current average height in the block; this preserves the skyline of blocks. Eq. 14 constrains how much the floor area in the whole city can change with a construction budget  $\delta^{(\max)}$ . Eq. 15 constrains the total change of residential floor area to be smaller than a threshold  $\delta_{f=\text{Res}}^{(\max)}$  to stabilize population changes.

Furthermore, as urban planners may want to control the increase or decrease of group mean benefits for some resident types, we add an objective to constrain the generated design recommendations. Given any arbitrary new design, let  $\Delta \bar{b}_{g^+}$ ,  $\Delta \bar{b}_{g^-}$  and  $\Delta \bar{b}_{g^0}$  denote the group average benefit difference between the new design and the current design for the groups whose mean benefits are to be increased  $(g^+ \in G^+)$ , decreased  $(G^-)$ , and kept unchanged  $(G^0)$ , respectively. The objective is to limit  $\Delta \bar{b}_{g^+} > 0$ ,  $\Delta \bar{b}_{g^-} < 0$ , and  $|\Delta \bar{b}_{g^0}| \approx 0$ . We define a solution penalty as

$$\phi(G^+, G^-, G^0) = \sum_{g^+ \in G^+} -\min(\Delta \bar{b}_{g^+} + \tau, 0) + \sum_{g^- \in G^-} \max(\Delta \bar{b}_{g^-} - \tau, 0) + \sum_{g^0 \in G^0} \max(|\Delta \bar{b}_{g^0}| - \tau, 0), \quad (16)$$

where  $\tau$  is a small threshold with  $\tau \ge 0$ . Adding this penalty function to Eq. 11 gives the new objective function

minimize 
$$\varepsilon^{\alpha}(B) + \lambda \phi(G^+, G^-, G^0),$$
 (17)

which searches for  $\delta_{k,f}$  that minimizes inequality while satisfying the design constraints.  $\lambda$  is a penalty hyperparameter, which we calibrated to be large, i.e.,  $\lambda \gg 1$ .

## Algorithm 1 Inequality Mitigation Algorithm

**Input:** Initial design  $v = \{v_{k,f}\}_{K,F}$ , design constraints P**Output:**  $\delta_c = \{\Delta v_{k,f}\}_{K,F}$  8

1: function INEQUALITY MITIGATION

2: Randomly initialize  $\delta_c \in P$ , c = 0

3: while  $\Delta m > \epsilon \operatorname{do}$ 

4:  $\boldsymbol{\delta} = \operatorname{argmin}_{\boldsymbol{\delta} \in \boldsymbol{P}} \boldsymbol{\nabla}_{\boldsymbol{\delta}} m(\boldsymbol{\delta}_c + \boldsymbol{v})^\top (\boldsymbol{\delta} - \boldsymbol{\delta}_c)$ 5:  $\boldsymbol{\delta}_{c+1} = \boldsymbol{\delta}_c + \zeta(\boldsymbol{\delta} - \boldsymbol{\delta}_c), \ \zeta = \frac{2}{2+c}$ 

 $egin{aligned} oldsymbol{v}_{c+1} &= oldsymbol{b}_c + \zeta (oldsymbol{b} - oldsymbol{b}_c), \, \zeta = rac{1}{2+c} \ oldsymbol{v} \leftarrow oldsymbol{\delta}_{c+1} + oldsymbol{v} \end{aligned}$ 

6:  $\boldsymbol{v} \leftarrow \boldsymbol{\delta}_{c+1} + \boldsymbol{v}$ 7:  $m_{c+1} = m(\boldsymbol{v})$ 

8:  $\Delta m = m_{c+1} - m_c$ 

9: c = c + 1

10: end while

11: return  $\delta_c$ 

12: end function

```
13: function m(v)
```

```
14: Run IPF to allocate residents to urban design v
```

15: Get benefit *B* of allocated residents by Eqs. 3 and 4

16: Get inequality 
$$\varepsilon^{\alpha}(B)$$
 and penalty  $\phi(G^+, G^-, G^0)$  by

Eqs. 5 and 16

17: **return**  $\varepsilon^{\alpha}(B) + \lambda \phi(G^+, G^-, G^0)$ 

18: end function

## 4.5.2 Heuristic Solution

We denote the objective function in Eq. 17 as  $m(\delta + v)$ , where matrix  $v = \{v_{k,f}\}_{k \in K, f \in F}$  denotes the current floor areas of each function type f in each census block k, and matrix  $\delta = \{\delta_{k,f}\}_{k \in K, f \in F}$  denotes the changes in floor areas to recommend. Note that for each potential  $\delta$ value, the total resident benefits B needs to be computed using IPF (described in Section 4.4), which is stochastic. This makes optimizing m inefficient. To efficiently find the optimal  $\delta$ , we propose an inequality mitigation algorithm (Algorithm 1) by employing the Frank-Wolfe conditional gradient method [94] to iteratively estimate  $\delta$  by linearly approximating the objective function m.

Algorithm 1 starts with the initial design v with a random design change  $\delta$  that is within the search space P defined by the constraints Eqs. 12 to 15. In each iteration c, it approximates the objective function  $m(\delta + v)$  with the first-order Taylor series expansion at the point  $\delta_c + v$ , i.e.,  $m(\delta_c + v) + \nabla_{\delta} m(\delta_c + v)^{\top} (\delta - \delta_c)$ . Minimizing this linear equation finds  $\delta$  that points to the direction towards the optimum (Line 4). We then update  $\delta_{c+1}$  with step size  $\zeta = \frac{2}{c+2}$  (Line 5). The objective function m then updates by allocating residents to the updated urban design  $v + \delta_{c+1}$  and calculating benefits, inequality and the object function value (Lines 13-18). The algorithm terminates when the difference between consecutive objective values is smaller than a threshold, i.e.,  $m_{c+1} - m_c \leq \epsilon$ .

## 4.5.3 Fairness Attribution to Edited Blocks

The recommendation will propose edits of multiple function types for several census blocks. It may recommend editing many blocks, which is tedious for the user. To help users prioritize which blocks to edit, we calculate another attribution explanation to indicate the partial contributions of each block towards improving the overall fairness. Inspired by SHAP [12] for explaining machine learning classifiers, we calculate these attributions as Shapley values [95], which fairly measure the contributions independent of calculation



Fig. 5: Basic features of IF-City for urban design. a) Charts showing Planning Indicators for floor areas for each function type, and Population Indicators for number of residents of each type with their respective preferences for function types. b) 3D map showing buildings that are color-coded by building type as indicated in the Planning Indicators. Users can navigate around the city by zooming, panning, and rotating the 3D visualization. c) Users can view details by selecting blocks or buildings and open a pop-up dialog. This shows floor area, height, building type; for residential buildings, this further shows population and benefit indicators, and accessibility to other building types. Users can edit the urban design by adding, editing or deleting buildings.

order. For a block k among the set of recommended blocks to edit R,  $k \in R$ , we calculate its attribution towards decreasing inequality as

$$\frac{1}{|R|} \sum_{S \subseteq R \setminus \{k\}} \underbrace{(\operatorname{contribution towards fairness}}_{-(\varepsilon^{\alpha}(B_R) - \varepsilon^{\alpha}(B_S))} / \underbrace{(|R| - 1, |S|)}_{C(|R| - 1, |S|)}$$
(18)

where *S* is a subset of *R* without block *k*; *B<sub>S</sub>* is the set of benefits after editing blocks *S* as recommended,  $\varepsilon^{\alpha}(B_S)$  is the inequality due to benefits *B<sub>S</sub>*; *B<sub>R</sub>* and  $\varepsilon^{\alpha}(B_R)$  are the benefits and inequality due to all recommended block edits. C(|R| - 1, |S|) denotes the number of combinations *S* that exclude block *k* chosen from set *R*. Since permuting all combinations is time-consuming, we adopt a sampling technique [96] that uses the average attribution of a few permutation samples to approximate the Shapley value to speed up calculations. Note that the attributions are only representative if all the recommended changes are executed.

## 5 IF-CITY: VISUAL DESIGN

Having derived the indicators and explanations to interpret fairness in urban planning, we next describe how we convey them through visual components of IF-City. Fig. 1 shows the whole application dashboard. Fig. 5 shows the basic urban design and Figs 6, 7 and 8 show intelligible fairness features. Inspired by Qua-Kit [85], we developed IF-City as a web app with 3D map and colored blocks for buildings.

## 5.1 Basic Features for Urban Design

**Urban (Planning and Population) indicators.** Urban planners track indicators to understand various characteristics of an urban design. We present these as numbers (e.g., total population, site area) and with charts to indicate subgroup information. Planning and population indicators (Fig. 5a) are presented on the right sidebar of the dashboard (Fig. 1), and indicate the total floor area of each building function type and population of each resident type, respectively. The different preferences of each resident type are also presented as vertical Preference Charts with colors corresponding to each function type. Planning and population priority weights (for equity) can be set with sliders (Appendix A, Fig. A1). Tooltip label hints are also provided.

9

**3D** map view and navigation. Users can examine the urban design by panning, zooming, and rotating a 3D map (see Fig. 5b). We use 3D instead of 2D so that users can view top-down to understand the location context, perceive building shapes and heights from various angles, and view at street-level for immersion. We also exploited 3D to visualize building occupancy in the Benefit Heatmap with cylinder heights. A satellite view (top of Fig. 1) is also provided to show the context of the neighborhood. The color of each building corresponds to its function type.

Editing buildings. Users can edit the design by adding, changing, or deleting buildings. IF-City supports planning at the granularity of specifying building footprints, heights, and function types. This is suitable for neighborhood planning, is more fine-grained than zoning, and more coarse-grained than specifying building tenants (e.g., retail or restaurants can occupy commercial spaces). Clicking a building shows a dialog (Fig. 5c) with floor area, number of floors, function type, and residential population and accessibility. Mixed building types include Residential, Office and Commercial functions, and users can set their ratios. Users can edit the details and press the "Calculate Benefit" button to rerun the resident allocation simulation and recalculate the benefits and inequality indicators; this will change the population distributions in the building and the wider city. Users can also draw new buildings or delete existing ones.

## 5.2 Intelligible Features for Fair Urban Design

IF-City has specialized features to indicate and explain fairness in urban design. Though central to the contributions of the paper, to ensure usability, we designed these features to blend in sensibly with the primary task of the dashboard, i.e., viewing and designing the city. Therefore, the benefit and inequality features supplement the basic urban planning features, rather than taking central focus in the UI design. We highlight these intelligibility features that occur across various parts of the dashboard (Figs. 6-8).

**Inequality indicators (IF.1).** Users can perceive fairness by the Total Inequality indicator (side panel in Fig. 1).

**Inequality attributions (IF.2, IF.3).** Users can identify sources of inequality by *Resident Type* or *Location*. Fairness across resident types can be checked by perceiving whether levels are flat in the Benefits Chart with small error bars and whether between- and within-group indices are small in the Inequality Chart (Fig. 6a). Fairness across locations can be examined in two ways: 1) by using Highlight Buildings to

© 2023 IEEE, Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.



Fig. 6: Intelligible fairness features of IF-City to identify causes of inequality. a) Inequality Attribution by Resident Type: View benefits of each resident type for the whole city, per block and per building. Benefits by resident type showing group benefit mean (bars) and standard deviation (error bar) and global mean (horizontal line). Inequality by resident type (red diamond) calculated as Generalized Entropy Index decomposable into between- (dark red bars) and within-group (pink bars) inequalities per group. b) Inequality Attribution by Location: Highlight (cyan outline) buildings filtered by occupancy and average benefit, or view relative benefit of each block and building as a Heatmap (green, red, white for above-, below-, average, respectively; darker colors for farther from average; colored ground areas for block-level benefits, 3D cylinders for building-level benefits with height for number of occupants). c) Inequality Trace: trace accessibility for different function types at each residential building for each resident type.

filter buildings with specific occupancy and average benefit ranges (Fig. 6b, Left), or 2) viewing the Benefits Heatmap to see above-average (green) or below-average (red) blocks and buildings (3D cylinders). (Fig. 6b, Right).

Inequality traces (IF.4). Focusing on a building, users can trace its source of inequality by clicking it to examine its Accessibility Circle and its dialog popup (Fig. 6c). Clicking on a block will also show similar information as for buildings. The blue Accessibility Circle shows which nearby buildings were included to calculate accessibility and utility, so users may want to edit them to improve inequality. The building-specific Benefit Chart shows benefit distribution for occupants within the building. The Accessibility Chart shows the accessibility for each function type, and has the option to show preference-weighted accessibility (utility) for each resident type. For example, in Fig. 6c, Outdoor Recreationalist residents at the selected building have very high utility due to the nearby park in the northwest, leading to above-average benefits compared to other resident types and contributing to between-group inequality.

**Population simulation to allocate residents (IF.5).** As an extension of the editing capabilities of the baseline interface, recalculating and visualizing the benefits and inequality indicators support users to understand how design changes can impact the fairness outcome (Fig. 7a). After editing buildings and pressing "Calculate Benefit", the user will see updates to population, benefit and inequality indicators per building and for the whole city.

Intelligible fair planning recommendation (IF.6). To let urban planners know what amenities would be most needed in which block and to accelerate the design iterations, users can query for recommendations on which census blocks to edit, and how much floor area of each building function type to change (Fig. 7b). The user can then choose how they want to edit buildings in recommended census blocks. Users can constrain the recommended changes that can be proposed, such as locking the benefits of resident types to remain unchanged, restricting benefits of specific resident types to only increase or decrease, and limiting the percentage change in total floor area. The recommended blocks are highlighted with thick colored outlines on the map with corresponding bars in an attribution chart with square icons of the same outline color. The green attribution bar chart shows the importance of each block towards reducing inequality (Eq. 18) and ranks them by decreasing attribution. On expanding each attribution bar, users can view a table of recommended floor area change for each function type for that block. Below the attribution chart, users can see the expected benefits and inequalities if the recommendation is followed. After accepting a recommendation, the user needs to edit the city design and recalculate (population simulation) to assess the design change.

10

**Design iteration timeline (IF.7).** For complex iterative design tasks, it is important to provide feedback across iterations for users to assess what is improving, and what may be traded-off. Fig. 8 shows a Timeline View of saved designs for users to track and compare planning indicators, including inequality. Clicking on a specific time point will load that design in a new browser window, showing the full dashboard with a map and charts. Using multiple screens, users can compare between urban designs for the overall city, per building, location, or per resident type.

## 6 SYSTEM IMPLEMENTATION

IF-City was developed as an interactive web app with HTML5 and JavaScript front-end, and Python back-end with a Flask web server and custom code for the application logic. The 3D map and navigation interface was built using the ArcGIS JavaScript API and the charts were implemented using HighCharts. For the fair planning recommendation engine, we utilized PyTorch to speed up gradient calculations and Gurobi as the optimization solver. Resident type clustering was performed using scikit-learn.

#### IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS



Fig. 7: Intelligibility fairness features of IF-City to mitigate inequality. a) Population Simulation to Allocate Residents: Change urban design by drawing new buildings, deleting existing ones, or editing them. Building heights and types can be edited. Residents are reallocated to the new design, benefits and inequality are recalculated after changes. b) Planning Recommendation to Mitigate Inequality: Request recommendation for which blocks to edit and how much floor area to change for each function type to reduce inequality. Estimated benefits and inequalities along with their attribution towards fairness improvement are shown. Users can constrain recommendations to specify whether benefits should increase, decrease, or be fixed for each resident type, and the percentage of floor areas allowed to be changed.



Fig. 8: Intelligibility fairness feature of IF-City Timeline View to show changes in indicators across designs. Clicking a time point opens the saved design for detailed comparison.

# 7 CASE STUDY: NEIGHBORHOOD RE-PLANNING

## 7.1 Background and Design Brief

We present a *synthetic* case study of the neighborhood of Southern Boulevard (Fig. 9a) in Bronx, New York City, to redesign it for fairer benefits across diverse resident types. Southern Boulevard is home to almost 60,000 residents. Urban planners have identified a number of gaps and opportunities to improve this neighborhood<sup>1</sup>. To examine how re-designing this neighborhood affects the re-allocation of residents, we simulated a *synthetic* population based on the check-in data of 4247 users from Foursquare collected

1. https://www1.nyc.gov/site/planning/plans/southernblvd/southern-blvd-updates.page



11

Fig. 9: Case study information. We used building and planning data regarding the South Boulevard neighborhood in the Bronx of New York City (NYC), United States. Building specifications follow the NYC Department of City Planning Documentation, and building data is from the PLUTO dataset [97]. Six resident types with distinct preferences for different building types were derived from a cluster analysis of check-in trajectory data of 4247 NYC Foursquare users.

in NYC (technical details in Appendix C). We identified six resident types, based on their preference for different building function types: Outdoor Recreationalists, General Consumers, Culture Consumers, Commercial Consumers, Office Workers, and Educators & Students (Fig. 9b). We pose a design brief with a planning goal to decrease the inequality indicator from 95 to  $\leq$ 60 and increase the average benefit from 190 to  $\geq$ 220, while maintaining the population to within 10% of the original indicator. This planning goal was also used later in the user study.

## 7.2 Walkthrough

We demonstrate IF-City with a user flow (Fig. 2) to design the neighborhood to be fairer. We articulate steps to perceive inequality (Steps 1, 7), identify their causes (2-4), and mitigate inequality (5-6). Actions are of a hypothetical user<sup>2</sup>.

**Perceive Inequality**. The user starts by looking at the inequality indicator and realizes the current indicator (94.98) is far from the planning goal ( $\leq 60$ ) (Step 1). This indicator later helps the user to compare between design iterations in the Timeline View (Step 7).

**Identify Causes of Inequality**. The user then wants to know who are advantaged and disadvantaged (*Step 2: Attribution by Resident Type*). She starts by looking at the Inequality Chart and finds that Office Workers, General Consumers and Culture Consumers have negative betweengroup inequality (dark red), while Outdoor Recreationists

<sup>2.</sup> Derived from our user study with real urban planning experts described in Section 8

12

have much higher positive between-group inequality. The within-group inequalities are relatively small for all resident types. She then examines the Benefit Chart and realizes that benefits are not evenly distributed among resident types. This explains the big variance in between-group inequality. She notices that Office Workers have negative betweengroup inequality (Step 2a) as they have below-average benefits (Step 2b). Next, she looks for the locations of inequality.

Step 3: Attribution by Location. There are two methods to identify the locations of greatest inequality — highlighting filtered buildings or viewing the benefits heatmap. Using the Highlight Building feature (Step 3a), she selects settings to highlight buildings which have high occupancy ( $\geq$ 80) of Office Workers with low average benefit ( $\leq$ 160 points). These buildings are highlighted in cyan on the 3D map. The user can also review her selections with the Benefits Heatmap and see that these buildings have tall green cylinders. Alternative to filtering, using the Benefits Heatmap (Step 3b), the user can perceive the uneven distribution of benefits across locations, and focus on the blocks with the lowest benefit (darkest red) and high occupancy to look for opportunities to add or remove Office spaces.

Step 4: Inequality Trace. For locations with below-average benefits for Office Workers, she may view their accessibility circle, change buildings to be used for Office functions, or add a new one. Conversely, for locations with aboveaverage benefits, locate Office buildings and decide whether to delete them or edit their function.

Mitigate Inequality. There are two approaches to mitigate inequality — trial-and-error with population simulation (Step 4) or following planning recommendations (Step 5). We describe a use case of editing different building function types to improve the average benefit from 196.5 and decrease total inequality from 94.98. Having found an empty land plot in the south, the user draws a new Mixed building (Step 5) with Office and Commercial floor areas to help disadvantaged Office Workers and General Consumers living nearby. After recalculating benefits, the user can perceive changes to the benefits and inequality in the respective charts (Step 2) and heatmap (Step 3b). In the Timeline View (Step 7), the user can track a decrease in the Inequality Indicator score (iteration 1 to 2). She continues to make other edits (Step 5). However, after a few iterations, the decrease plateaus (iteration 4). Next, the user edits Cultural buildings to help Cultural Consumers, for example, increasing the height of one to 15 floors. This decreases Inequality somewhat, and is repeatable for two more iterations until plateauing again (iteration 5 to 7).

At this stage, the user changes the strategy to use planning recommendations (Step 6). She requests a recommendation with constraints to not increase benefits to Outdoor Recreationalists any higher, not decrease benefits to General Consumers, Cultural Consumers, and Office Workers, and limit floor area changes to 30%. She is recommended to edit 9 blocks (4 shown in Fig. 2(6)), ranked in decreasing attribution to improve fairness. For Block 014, if she adds 23,533m<sup>2</sup> and 11,233m<sup>2</sup> of Office and Commercial floor areas, respectively, she could decrease Inequality by 18. Following this recommendation, she draws a Mixed building at an empty space and achieved a significant decrease in Inequality (Fig. 2(7), iteration 8). She runs the recommendation one more time and finishes. Ultimately, she raised the average benefit to 280.7 and reduce total inequality to 33.13.

**Findings:** The analysis recommends that to decrease inequality, the neighborhood needs more (i) office floor area in the south of the neighborhood; (ii) cultural area in the mid-west; and (iii) cultural and commercial areas in the west. We present analysis details in Appendix D.

## 8 EVALUATING WITH DOMAIN EXPERTS

We conducted a qualitative study with six domain experts separately to 1) confirm the scope and priority of fairness in urban planning objectives, and 2) evaluate the usefulness of IF-City in helping experts to understand the causes of inequalities and mitigate them. Three experts are urban planners (E1 and E2 have experience in China, E3 has experience in Europe) and three are urban designers (E4, E5 and E6 have experience in the United States). All experts had 3-10 years of experience.

Lasting two hours, the study procedure included: i) a short interview where the expert was asked about his own understanding of and experience with fairness in urban planning, ii) an introduction to the target neighborhood Southern Boulevard its planning goals, followed by iii) a description of our fairness definition, iv) a briefing on two tasks to identify the causes of inequalities and mitigate them, and v) a tutorial on using IF-City. With consent, we recorded audio and screen captured interactions with IF-City. We analyzed the recordings in terms of using intelligibility within the three stages of the IF-Alloc framework. We present our findings on the experts' understanding on fairness in urban planning, their strategies to complete the tasks, and the reported usefulness and usability of IF-City.

#### 8.1 Priority and Scope of Fairness in Urban Planning

All experts believed fairness is important. Some (E1, E2, E3 and E6) have previously integrated fairness in their designs qualitatively, while others (E4 and E5) have quantified fairness in terms of accessibility and walkability to evaluate existing urban designs. They all agreed that measuring and visualizing fairness in IF-City made it "easy to compare the fairness of planning outcomes" [E4], and it was "nice to see the intermediate effect of fairness after every single change of the design" [E1]. However, they also argued that fairness is not the first planning priority; there were "many other important criteria, such as environmental effect and economic effect" [E1].

Other than equality, the experts also considered equity as an important aspect of fairness. For example, E6 mentioned that "designs should guarantee the benefit of low-income population, and also the elderly and disabled people". Regarding grouping residents into types, all experts would typically segment the population by demographics rather than activity preference, but E3 and E5 felt that segmenting by the latter was useful to inform "what are their [population group] needs" [E3] and to examine the needs of non-traditional groups such as outdoor recreationalists and educators.

## 8.2 Perceiving and Identifying the Inequality Sources

Overall, experts employed common steps to find the sources of inequality (numbered as in Fig. 2): 1) perceive the inequality indicator, 2) identify one disadvantaged resident type

#### IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS

at a time from the Benefit and Inequality Charts, 3) locate where they live by filtering for buildings with such residents (3a) and identifying buildings with high occupancy (i.e., tall cylinders in the benefit heatmaps; 3b), 4) check accessibility by selecting the building or block and examining the function type floor area distribution of nearby buildings, 5) edit building properties (e.g., height, function type) to explore whether the changes could decrease inequality indicator. Next, we report how the visual components in IF-City helped the experts to perform these steps.

Interpreting benefit and inequality charts (IF.2). Most experts preferred to use the Benefit Chart to identify which resident type is disadvantaged, as "*it is simple and very easy to understand*" [E1]. However, due to positivity bias, they tended to overlook that advantaged resident types contributed to inequality too. Though somewhat less intuitive than the Benefit Chart, experts could learn more insights from the Inequality Chart. E5 found that, in the Benefit Chart, both Outdoor Recreationalists and Offices Workers "deviate from the average with the same amount" above and below mean, respectively; but in the Inequality Chart the more populous "Recreationalists have the largest between-group inequality score, so they contribute most to the total inequality". Moreover, while the experts could perceive within-group inequality, they did not investigate their cause further.

Locating disadvantaged resident types (IF.3). Most experts preferred to use the Benefit Heatmap to locate disadvantaged resident types on the map, as they could "quickly locate the blocks in red color and with high cylinders" [E1]. In contrast, Building Filtering "[took] more time" [E2] and required more interactions to tune settings for a suitable selection threshold. Nevertheless, it provided detailed distributions of benefit and occupancy, allowing experts to select the buildings based on perceived quantiles.

Identifying causes of inequality due to accessibility (IF.4). On finding where disadvantaged residents were on the map, experts further examined the accessibility at the affected buildings. Most experts only quickly viewed the blue Accessibility Circle to identify what function types it encloses; only a few experts studied the Accessibility Chart to examine the floor area distribution of function types. This helped them to clearly see how much specific function types could be added or removed, e.g., E1 pointed out "the bar chart helps me to compare floor area of different function types and find out which function type should be added".

#### 8.3 Mitigating and Verifying Inequality

Here, we describe how experts used two approaches to mitigate inequality, and verified improvements while iterating.

Trial and error strategy (IF.5). All experts first chose to manually improve fairness by editing buildings or blocks identified as disadvantaged, and running the population simulation to recalculate the benefit and inequality indicators. Most experts added new Office, Culture and Mixed buildings to improve the benefit for Office Workers, Culture Consumers and General Consumers. They tended to add new buildings than edit existing ones to avoid disrupting existing activities. The edits were interspersed with studying the Benefit and Inequality Charts to identify which resident types had inequality, their Preference Charts of function types, and the Benefit Heatmap to see where they lived. This helped E2 and E4 to retarget their edits. With their independent effort, four experts (E1, E4, E5, E6) achieved the planning goal of sufficiently low inequality and high total benefit, but they took a long time with many edit iterations to do so. They performed 4 to 14 iterations (M = 7) lasting 30-90 minutes. Experts E2 and E3 gave up editing further after only 2 to 3 edits, and only decreased inequality by less than 15 points. We report detailed editing actions and planning outcomes in Appendix D. To further improve fairness, the experts subsequently used the planning recommendations.

Finally, enunciating the wicked nature of urban planning, the experts reported that mitigating inequality was complex and challenging because improving benefits for one resident type may hurt others. It would be "tedious to check which building I can change so that no one's benefit gets hurt" [E2]. E6 found that the benefit of General Consumers "is hard to improve because their preference is multi-fold", to include Office, Commercial and Cultural buildings, but changing them will also affect the benefits to Office Workers, Commercial Consumers and Culture Consumers.

Planning recommendations (IF.6). After manual attempts to improve fairness, the experts investigated further improvements with automatic recommendations. Most experts were conservative and wanted to see the recommendation details to examine whether they were realistic and feasible. E2, E3, E5 and E6 were interested in how the recommended edits were automatically calculated. Conversely, E1 and E4 were initially skeptical and assumed that the calculations were too simple; they believed that planning is complicated and many factors should be considered such as street views and height control. After following and executing the top few recommendations, all experts found the recommended edits were feasible and led to large decreases in the inequality indicator. They "trust[ed] the recommendations" [E5] and found that the recommendation table was "very transparent" [E3] as they learned "which block needs what [to edit]" [E2]. The recommendations helped them to "narrow down the inequality problems to certain blocks" [E4]. They also appreciated the coarse, block-level recommendations as they "[had] the freedom to make detailed design within a block" [E6]. E3 was an exception who wanted the recommendation "automated into buildings" so that he did not need to make any edit by himself. Regarding constraining the recommendations, the experts appreciated the control of prioritizing resident types, as "sometimes urban planners need to make sure the benefit of certain resident types should be *improved by some policy*" [E3]. They also liked the setting to threshold the floor area change because they could easily limit a planning budget. Overall, the experts thought the planning recommendation feature was "a good guide" [E3] and wanted to use it rather than solving the inequality problems by themselves manually.

**Verifying and inspecting improvements (IF.7).** Most experts used the Timeline View to check whether the inequality indicator improved with each design iteration. E1 liked that he could perceive "the intermediate effect of fairness after every single change of the design". E2 appreciated perceiving the "magnitude and scale" of his edits on the inequality index. This indicates the importance of the comparative use of What, to contextualize the meaning of the inequality index that would otherwise be too abstract. Furthermore,

IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS

E1 clicked on her past saved designs and compared details with the current version. This helped her to understand that her adding of Office buildings in the south and Commercial buildings in the middle of the neighborhood had significantly improved fairness.

## 8.4 Feedback on Usability and Usefulness

**Usability**. All the experts found the Benefit and Inequality charts, the Heatmap, and the Cylinders easy to understand: *"These features are very helpful to see the inequality situation. I like this visualization and I think it is very presentable to clients and stakeholders"* [E4]. E6 liked the design of IF-City and felt it easy to get familiar with it as *"this tool is quite intuitive"*. E5 commented that the *"information is well organized, and the concept of fairness is clearly conveyed"*, and believed planners could learn it in a short time. However, she would prefer less scrolling to see all the charts and tables.

**Usefulness**. The experts reported IF-City would be useful for "overall high-level land use planning" [E4], "before the block subdivision process" [E6] and for "[re-evaluating] land use" [E4]. E4 also commented that having the quantitative tool is helpful to explicitly articulate fairness as "most of the time the fairness was discussed in conversation with photos". E5 thought IF-City could be integrated with the existing professional design tools such as Autodesk and Rhino.

## **9** EVALUATION OF RECOMMENDED MITIGATIONS

We conducted a quantitative study to evaluate how much our inequality mitigation recommendation Algorithm 1 could improve fairness 1) with different constraints on how much floor area to change and 2) in comparison to the performance of the six experts with the trial and error strategy. We ran the algorithm on the neighborhood (Fig. 9) with an initial inequality indicator of 94.98. In each run with a certain floor area change limit, the algorithm returns a list of census blocks with floor area of each building function type to change (Fig. 7b). The blocks are sorted by decreasing Shapley attribution. We manually edited the floor areas (e.g., adding office space) to follow recommendations for each block in the sorted order, and measured the inequality indicators. We also measured the manual edits and resulting inequality scores for our domain experts from user study.

Fig. 10 shows the inequality indicators resulting from manual edits by domain experts and automated recommendation by IF-Plan. IF-Plan edits generally resulted in at least the same or lower inequality scores than human experts. Our domain experts stopped early with relatively high inequality scores, while the recommendation system could iterate more to achieve even lower inequality.

## 10 DISCUSSION

We have demonstrated the usability and usefulness of intelligibility for fair urban design. Here, we discuss 1) the need for domain-specific fairness visualizations, 2) the usefulness of intelligibility in fairness understanding, 3) using IF-City to analyze fairness for different neighborhoods or cities, 4) generalizing IF-Alloc framework beyond urban planning, and 5) limitations and future work.



Fig. 10: Inequality mitigation performances of the 6 human experts versus Intelligible Fair Planning Recommendation (IF-Plan). Each black marker with a different shape represents the lowest inequality achieved by a different human expert. Colored lines represent iterations by IF-Plan for different floor area change limit (10%, 30%, and 50%).

Domain-specific fairness visualizations. Unlike data scientists whose role is to understand and model data, expert users have domain-specific roles and tasks with specialized tools and workflows. Regarding domain experts, we have identified fair allocation as a concern for urban planners, defined and implemented inequality indicators and explanations for resident types and locations, designed visualizations of group fairness, and integrated fairness visualizations seamlessly into a workflow for urban design. In our evaluation, users were comfortable using IF-City, appreciated the need for fairness, and were effective to reduce inequality. Though important, fairness is not the primary goal in many applications, so its visualization needs to be carefully designed and integrated. We drew from existing urban planning tools to leverage geospatial visualizations, 3D models, and accessibility calculations to develop a domain-specific intelligible fairness tool. Other applications that can benefit from further study include social network analysis [98] or computer networking [99] by visualizing graph networks, and fair scheduling [100] by visualizing timetables.

Intelligibility for fair allocation. We have developed various intelligibility features to support three stages to iteratively perceive, understand, and mitigate inequality in resource allocation. Attribution by Resident Type (IF.2) and Location (IF.3) to identify the cause of inequality were useful, especially with heatmaps to locate inequality, perhaps due to the familiarity with maps in urban planning. Providing both population simulation (IF.5) and planning recommendation (IF.6) to mitigate inequality helps to empower expert designers to make fine-grained decisions, while supporting efficient suggestions on-demand. This grants control to domain experts to integrate other concerns or constraints in their solutions. In our study, experts first explored via trial-and-error (IF.5), before exploring automatic recommendations. However, for cases where users are not expert planners (e.g., [7]), it may be better to prioritize recommendations (IF.6). Finally, although contrastive explanations are key for explainability goals [14], we found limited use of the design iteration timeline (IF.7). Our experts did use the timeline to track their progress, but rarely compared details between past and current designs. Perhaps, due to

#### IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS

the easily remembering their earlier designs, the tediousness of examining many buildings between urban plans, and the brief duration of the experiment session. We expect the timeline to be more useful in usages spanning hours or days.

Generalizing to analyze fairness for other neighborhoods or cities. The IF-Alloc framework (Fig. 2) can be directly applied to other cities. Using IF-City, urban planners can upload the designs of other cities (as Geo-JSON) and their resident demographic or preference data (JSON) to assess the fairness of those cities. Moreover, by loading the dashboard in multiple web browser windows, users can examine detailed reasons why one city may be fairer than another. We have evaluated IF-City on a small neighborhood and note the need to scale the calculations and memory requirements for much larger cities. Furthermore, while we focused on fairness, IF-City can be extended to other objectives, such as green space and economic sustainability.

Generalizing to other fairness applications beyond urban planning. IF-Alloc can be applied to other fair allocation applications, such as job allocation for the gig economy [101], product placement in retail or online stores [102], worker shift allocation (e.g., in restaurants, hospitals) [103], and donation division [7]. Following the steps in Fig. 3, we describe an example application for fair driver job allocation in ride-sharing applications [104], [105]. 1) Extract map and road network data to determine routes and distances. 2) Extract driver background and behavioral information of drivers from worker surveys, app usage, and driving trajectories (2a) to cluster (2b) them into driver types (2c) based on their demographics, preferred driving times and locations, etc. (2d). 3) Simulate a job request scenario (multiple riders' requests, given multiple drivers) and automatically allocate jobs [106]. 4) Calculate the benefits (earnings) for each driver given their job allocation, aggregate benefits across driver types (4a) and current locations (4b), and calculate group inequalities using Generalized Entropy (4c). Steps 4a and 4b support explaining Why about the fairness. 5) Enable job planners to adjust planning parameters (e.g., commission fee [107], surge pricing rate [108], incentive threshold [109], working speed [110]) and run the allocation engine to simulate What If the parameters were different (5a), or request recommendations of fairer settings (5b). In general, IF-Alloc can be applied to applications involving resource or task allocation, planning parameters, and diverse stakeholders.

**Limitation and future work.** Since urban planning is a complex problem, in this work, we simplified the planning process which may result in some limitations to be addressed in the future work.

- a) More inclusive demographic data. We mined resident preferences from social media which might be biased, as it excludes residents who are non-users. Our preference clustering was also based on location visits, but excluded less tangible activities. Nevertheless, our activity-based preferences can be complemented with contemporary approaches that collect demographicbased preferences from surveys of urban residents.
- b) More practical population simulation. We allocated residents individually in the population simulation, and neglected modeling families or households that move together. Future work can perform allocation at the household level to balance the diverse preferences of

different family members. Also, the residents are allocated by probabilities that are proportional to benefits, neglecting the dynamic market prices and competition for space. Future work will integrate more practical factors, such as resident income and house property values, ages and types [92], [93].

- c) *More transport modes for accessibility modeling.* We calculated accessibility based on straight-line distances and walkability, which is a common approach. Extending this work to include transportation planning can consider transportation networks (roads, pathways, tracks) and transportation modes (e.g., cars, bicycles, trains).
- d) *More local urban planning knowledge.* We evaluated with domain experts who were not US-based planners, focusing on evaluating the method of using IF-City. For actual urban planning, local planners should be engaged.

Note that the Inequality Score depends on planning and population priority weights,  $\rho_f$  and  $\rho_g$ . Urban planners will need to determine these for each city and population based on their planning goals. Thus, users of IF-City should not naively use the Inequality Score as a generic benchmark across cities, but consider assumptions encoded with the priority weights.

Our formulation of resident types assumes that each resident belongs to only one type. However, in reality, the types may overlap, e.g., residents belonging to both disabled and elderly populations. Future work should treat resident types as attribute-based rather than categorically. Population priority weights would have to be modeled to accommodate this too, e.g., by averaging across the types that a resident belongs to.

## 11 CONCLUSION

We have proposed Intelligible Fair City Planner (IF-City) an interactive visualization tool to support the design of fair urban plans. In consultation with urban planners, we identified requirements to perceive fairness, identify causes, and mitigate inequality. We formalized this the generalizable IF-Alloc framework for intelligible fair allocation. We further proposed an intelligible fair planning recommendation (IF-Plan) method to automatically recommend fairer urban plans to accelerate design iterations. We demonstrated and evaluated the usage and usefulness of IF-City in a real neighborhood in New York City, US, with practicing urban planners and urban designers. Using various intelligibility features, urban planners could perceive and identify causes of inequality, and mitigate inequality. This works sheds light on how to carefully design detailed user interactions for fair design in collaborative human-AI planning.

## ACKNOWLEDGMENTS

This work was supported in part by the National Key Research and Development Program of China under grant No. 2019YFB2102200; in part by the Singapore Ministry of Education (MOE) Academic Research Fund Tier 2 T2EP20121-0040; in part by the Natural Science Foundation of China under Grant 62102082; in part by the Jiangsu Natural Science Foundation of China under Grant BK20210203.

#### REFERENCES

- J. Buolamwini and T. Gebru, "Gender shades: Intersectional [1] accuracy disparities in commercial gender classification," in Conference on fairness, accountability and transparency. PMLR, 2018, pp. 77–91.
- [2] Y. Ahn and Y.-R. Lin, "Fairsight: Visual analytics for fairness in decision making," IEEE transactions on visualization and computer graphics, vol. 26, no. 1, pp. 1086-1095, 2019.
- [3] A. A. Cabrera, W. Epperson, F. Hohman, M. Kahng, J. Morgenstern, and D. H. Chau, "Fairvis: Visual analytics for discovering intersectional bias in machine learning," in IEEE Conference on Visual Analytics Science and Technology, 2019, pp. 46-56.
- Q. Wang, Z. Xu, Z. Chen, Y. Wang, S. Liu, and H. Qu, "Visual [4] analysis of discrimination in machine learning," IEEE Transactions on Visualization and Computer Graphics, 2020.
- W. Zeng and Y. Ye, "Vitalvizor: A visual analytics system for studying urban vitality," *IEEE Computer Graphics and Applications*, vol. 38, no. 5, pp. 38–53, 2018. [5]
- D. Weng, H. Zhu, J. Bao, Y. Zheng, and Y. Wu, "Homefinder [6] revisited: Finding ideal homes with reachability-centric multicriteria decision making," in Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, 2018, pp. 1-12.
- M. K. Lee, D. Kusbit, A. Kahng, J. T. Kim, X. Yuan, A. Chan, D. See, R. Noothigattu, S. Lee, A. Psomas *et al.*, "WeBuildAI: Par-[7] ticipatory framework for algorithmic governance," Proceedings of the ACM on Human-Computer Interaction, vol. 3, pp. 1-35, 2019.
- S. Robertson, T. Nguyen, and N. Salehi, "Modeling assumptions [8] clash with the real world: Transparency, equity, and community challenges for student assignment algorithms," in Proceedings of the CHI Conference on Human Factors in Computing Systems, 2021.
- [9] A. Abdul, J. Vermeulen, D. Wang, B. Y. Lim, and M. Kankanhalli, "Trends and trajectories for explainable, accountable and intelligible systems: An hci research agenda," in Proceedings of the CHI conference on human factors in computing systems, 2018, pp. 1–18. R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and
- [10] D. Pedreschi, "A survey of methods for explaining black box models," ACM computing surveys, vol. 51, no. 5, pp. 1–42, 2018. D. Wang, Q. Yang, A. Abdul, and B. Y. Lim, "Designing theory-
- [11] driven user-centric explainable ai," in Proceedings of the 2019 CHI conference on human factors in computing systems, 2019, pp. 1-15.
- S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting [12] model predictions," in Proceedings of the 31st International Conference on Neural Information Processing Systems, 2017, pp. 4768-4777.
- M. T. Ribeiro, S. Singh, and C. Guestrin, ""why should i trust you?" Explaining the predictions of any classifier," in *Proceedings* [13] of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, 2016, pp. 1135-1144.
- T. Miller, "Explanation in artificial intelligence: Insights from the [14] social sciences," Artificial Intelligence, vol. 267, pp. 1–38, 2019.
- S. Wachter, B. Mittelstadt, and C. Russell, "Counterfactual expla-[15] nations without opening the black box: Automated decisions and the gdpr," Harv. JL & Tech., vol. 31, p. 841, 2017.
- J. Wexler, M. Pushkarna, T. Bolukbasi, M. Wattenberg, F. Viégas, [16] and J. Wilson, "The what-if tool: Interactive probing of machine learning models," IEEE transactions on visualization and computer graphics, vol. 26, no. 1, pp. 56-65, 2019.
- [17] S. Corbett-Davies and S. Goel, "The measure and mismeasure of fairness: A critical review of fair machine learning," arXiv preprint arXiv:1808.00023, 2018.
- N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Gal-[18] styan, "A survey on bias and fairness in machine learning," ACM Computing Surveys (CSUR), vol. 54, no. 6, pp. 1-35, 2021.
- [19] B. Y. Lim and A. K. Dey, "Assessing demand for intelligibility in context-aware applications," in Proceedings of the 11th international conference on Ubiquitous computing, 2009, pp. 195–204. K. Müller and K. W. Axhausen, "Population synthesis for mi-
- [20] crosimulation: State of the art," Arbeitsberichte Verkehrs-und Raumplanung, vol. 638, 2010.
- C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, [21] "Fairness through awareness," in Proceedings of the 3rd innovations in theoretical computer science conference, 2012, pp. 214-226.
- [22] M. Hardt, E. Price, and N. Srebro, "Equality of opportunity in supervised learning," Advances in neural information processing systems, vol. 29, pp. 3315–3323, 2016.
- A. Beutel, J. Chen, T. Doshi, H. Qian, A. Woodruff, C. Luu, P. Kre-[23] itmann, J. Bischof, and E. H. Chi, "Putting fairness principles into practice: Challenges, metrics, and improvements," in Proceedings

of the AAAI/ACM Conference on AI, Ethics, and Society, 2019, pp. 453-459.

- [24] S. K. Baruah, N. K. Cohen, C. G. Plaxton, and D. A. Varvel, "Proportionate progress: A notion of fairness in resource allocation," Algorithmica, vol. 15, no. 6, pp. 600-625, 1996.
- [25] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica, "Dominant resource fairness: Fair allocation of multiple resource types." in Nsdi, vol. 11, no. 2011, 2011, pp. 24-24.
- [26] C. Joe-Wong, S. Sen, T. Lan, and M. Chiang, "Multiresource allocation: Fairness-efficiency tradeoffs in a unifying framework," IEEE/ACM Transactions on Networking, vol. 21, no. 6, pp. 1785-1798, 2013.
- R. K. Jain, D.-M. W. Chiu, W. R. Hawe et al., "A quantitative mea-[27] sure of fairness and discrimination," Eastern Research Laboratory, Digital Equipment Corporation, Hudson, MA, 1984.
- B. Radunovic and J.-Y. Le Boudec, "A unified framework for [28] max-min and min-max fairness with applications," IEEE/ACM Transactions on networking, vol. 15, no. 5, pp. 1073-1083, 2007.
- [29] F. Kelly, "Charging and rate control for elastic traffic," European transactions on Telecommunications, vol. 8, no. 1, pp. 33-37, 1997.
- A. Rényi et al., "On measures of entropy and information," [30] in Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics. The Regents of the University of California, 1961.
- M. O. Lorenz, "Methods of measuring the concentration of [31] wealth," Publications of the American statistical association, vol. 9, no. 70, pp. 209-219, 1905.
- E. M. Hoover, "The measurement of industrial localization," The [32] Review of Economic Statistics, pp. 162-171, 1936.
- A. B. Atkinson et al., "On the measurement of inequality," Journal [33] of economic theory, vol. 2, no. 3, pp. 244-263, 1970.
- A. F. Shorrocks, "The class of additively decomposable inequality [34] measures," Econometrica: Journal of the Econometric Society, pp. 613-625, 1980.
- T. Speicher, H. Heidari, N. Grgic-Hlaca, K. P. Gummadi, [35] A. Singla, A. Weller, and M. B. Zafar, "A unified approach to quantifying algorithmic unfairness: Measuring individual & group unfairness via inequality indices," in Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2018, pp. 2239-2248.
- [36] T. Das and A. Parikh, "Decomposition of inequality measures and a comparative analysis," Empirical Economics, vol. 7, no. 1, pp. 23-48, 1982.
- M. Liu, S. Luo, and X. Du, "Exploring equity in healthcare [37] services: Spatial accessibility changes during subway expansion, ISPRS International Journal of Geo-Information, vol. 10, no. 7, p. 439, 2021.
- B. Santos, A. Antunes, and E. J. Miller, "Integrating equity objectives in a road network design model," *Transportation Research* [38] *Record*, vol. 2089, no. 1, pp. 35–42, 2008. C. Nyelele and C. N. Kroll, "The equity of urban forest ecosystem
- [39] services and benefits in the bronx, ny," Urban Forestry & Urban Greening, vol. 53, p. 126723, 2020.
- [40] A. Datta, S. Sen, and Y. Zick, "Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems," in 2016 IEEE symposium on security and privacy (SP). IEEE, 2016, pp. 598-617.
- M. Reyes, A. Páez, and C. Morency, "Walking accessibility to [41] urban parks by children: A case study of montreal," Landscape and Urban Planning, vol. 125, pp. 38-47, 2014.
- N. Foth, K. Manaugh, and A. M. El-Geneidy, "Towards equitable [42] transit: examining transit accessibility and social need in toronto, canada, 1996-2006," Journal of transport geography, vol. 29, pp. 1-10, 2013.
- [43] T. Neutens, T. Schwanen, F. Witlox, and P. De Maeyer, "Equity of urban service delivery: a comparison of different accessibility measures," Environment and Planning A, vol. 42, no. 7, pp. 1613-1635, 2010.
- P. Cheng, M. Min, W. Hu, and A. Zhang, "A framework for [44] fairness evaluation and improvement of urban green space: A case of wuhan metropolitan area in china," Forests, vol. 12, no. 7, p. 890, 2021.
- L. Anselin, "Local indicators of spatial association-lisa," Geo-[45] graphical analysis, vol. 27, no. 2, pp. 93–115, 1995. R. K. Bellamy, K. Dey, M. Hind, S. C. Hoffman, S. Houde,
- [46] K. Kannan, P. Lohia, J. Martino, S. Mehta, A. Mojsilovic et al.,

"AI Fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias," *arXiv preprint arXiv:1810.01943*, 2018.

- [47] C. Xu, C. Greer, M. N. Joshi, and T. Doshi, "Fairness indicators demo: Scalable infrastructure for fair ml systems," 2020.
- [48] S. Bird, M. Dudík, R. Edgar, B. Horn, R. Lutz, V. Milan, M. Sameki, H. Wallach, and K. Walker, "Fairlearn: A toolkit for assessing and improving fairness in ai," *Microsoft, Tech. Rep. MSR-TR*-2020-32, 2020.
- [49] J. N. Yan, Z. Gu, H. Lin, and J. M. Rzeszotarski, "Silva: Interactively assessing machine learning fairness using causality," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–13.
- [50] H. Yu, Y. Liu, X. Wei, C. Zheng, T. Chen, Q. Yang, and X. Peng, "Fair and explainable dynamic engagement of crowd workers." in *IJCAI*, 2019, pp. 6575–6577.
- [51] P.-M. Law, W. Wu, Y. Zheng, and H. Qu, "Vismatchmaker: Cooperation of the user and the computer in centralized matching adjustment," *IEEE transactions on visualization and computer graphics*, vol. 23, no. 1, pp. 231–240, 2016.
- [52] E. Talen, "Visualizing fairness: Equity maps for planners," *Journal of the American Planning Association*, vol. 64, no. 1, pp. 22–38, 1998.
- [53] A. Kumar, S. Shah, M. Lowalekar, P. Varakantham, A. Ottley, and W. Yeoh, "FairVizARD: A visualization system for assessing fairness of ride-sharing matching algorithms," in *Proceedings of the International Conference on Automated Planning and Scheduling* (ICAPS), 2021.
- [54] A. Dhurandhar, P.-Y. Chen, R. Luss, C.-C. Tu, P. Ting, K. Shanmugam, and P. Das, "Explanations based on the missing: Towards contrastive explanations with pertinent negatives," in *Advances in Neural Information Processing Systems*, 2018.
- [55] T. Miller, "Contrastive explanation: A structural-model approach," The Knowledge Engineering Review, vol. 36, 2021.
- [56] F. Cheng, Y. Ming, and H. Qu, "DECE: Decision explorer with counterfactual explanations for machine learning models," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 1438–1447, 2020.
- [57] O. Gomez, S. Holter, J. Yuan, and E. Bertini, "Vice: visual counterfactual explanations for machine learning models," in *Proceedings* of the 25th International Conference on Intelligent User Interfaces, 2020, pp. 531–535.
- [58] A. Abdul, C. von der Weth, M. Kankanhalli, and B. Y. Lim, "Cogam: Measuring and moderating cognitive load in machine learning model explanations," in *Proceedings of the CHI Conference* on Human Factors in Computing Systems, 2020, pp. 1–14.
- [59] J. Krause, A. Perer, and K. Ng, "Interacting with predictions: Visual inspection of black-box machine learning models," in *Proceedings of the 2016 CHI conference on human factors in computing* systems, 2016, pp. 5686–5697.
- [60] M. Kahng, P. Y. Andrews, A. Kalro, and D. H. P. Chau, "Activis: Visual exploration of industry-scale deep neural network models," *IEEE transactions on visualization and computer graphics*, vol. 24, no. 1, pp. 88–97, 2017.
- [61] Z. J. Wang, R. Turko, O. Shaikh, H. Park, N. Das, F. Hohman, M. Kahng, and D. H. Chau, "CNN explainer: Learning convolutional neural networks with interactive visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 1396–1406, 2020.
- [62] F. Hohman, H. Park, C. Robinson, and D. H. P. Chau, "Summit: Scaling deep learning interpretability by visualizing activation and attribution summarizations," *IEEE transactions on visualization and computer graphics*, vol. 26, no. 1, pp. 1096–1106, 2019.
- [63] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," in *In Workshop at International Conference on Learning Representations*, 2014.
- [64] C. Olah, A. Mordvintsev, and L. Schubert, "Feature visualization," Distill, vol. 2, no. 11, p. e7, 2017.
- [65] T. Spinner, U. Schlegel, H. Schäfer, and M. El-Assady, "explAIner: A visual analytics framework for interactive and explainable machine learning," *IEEE transactions on visualization and computer* graphics, vol. 26, no. 1, pp. 1064–1074, 2019.
- [66] J. Krause, A. Dasgupta, J. Swartz, Y. Aphinyanaphongs, and E. Bertini, "A workflow for visual diagnostics of binary classifiers using instance-level explanations," in 2017 IEEE Conference on Visual Analytics Science and Technology (VAST), 2017, pp. 162–172.

- [67] Y. Ming, H. Qu, and E. Bertini, "Rulematrix: Visualizing and understanding classifiers with rules," *IEEE transactions on visualization and computer graphics*, vol. 25, no. 1, pp. 342–352, 2018.
- [68] M. P. Neto and F. V. Paulovich, "Explainable matrix-visualization for global and local interpretability of random forest classification ensembles," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 1427–1437, 2020.
- [69] F. Hohman, A. Head, R. Caruana, R. DeLine, and S. M. Drucker, "Gamut: A design probe to understand how data scientists understand machine learning models," in *Proceedings of the CHI* conference on human factors in computing systems, 2019, pp. 1–13.
- [70] Q. V. Liao, D. Gruen, and S. Miller, "Questioning the ai: informing design practices for explainable ai user experiences," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–15.
- [71] B. Y. Lim and A. K. Dey, "Toolkit to support intelligibility in context-aware applications," in *Proceedings of the 12th ACM international conference on Ubiquitous computing*, 2010, pp. 13–22.
- [72] S. Verma and J. Rubin, "Fairness definitions explained," in 2018 IEEE/ACM International Workshop on Software Fairness (FairWare). IEEE, 2018, pp. 1–7.
- [73] N. Adrienko and G. Adrienko, "Spatial generalization and aggregation of massive movement data," *IEEE Transactions on visualization and computer graphics*, vol. 17, no. 2, pp. 205–219, 2010.
- [74] D. Holten and J. J. Van Wijk, "Force-directed edge bundling for graph visualization," in *Computer graphics forum*, vol. 28, no. 3. Wiley Online Library, 2009, pp. 983–990.
- [75] C. Hurter, O. Ersoy, and A. Telea, "Graph bundling by kernel density estimation," in *Computer Graphics Forum*, vol. 31, no. 3pt1. Wiley Online Library, 2012, pp. 865–874.
- [76] D. Liu, D. Weng, Y. Li, J. Bao, Y. Zheng, H. Qu, and Y. Wu, "Smartadp: Visual analytics of large-scale taxi trajectories for selecting billboard locations," *IEEE transactions on visualization and computer graphics*, vol. 23, no. 1, pp. 1–10, 2016.
- [77] Y. Lyu, X. Liu, H. Chen, A. Mangal, K. Liu, C. Chen, and B. Lim, "Od morphing: Balancing simplicity with faithfulness for od bundling," *IEEE transactions on visualization and computer graphics*, vol. 26, no. 1, pp. 811–821, 2019.
- [78] W. Zeng, Q. Shen, Y. Jiang, and A. Telea, "Route-aware edge bundling for visualizing origin-destination trails in urban traffic," in *Computer Graphics Forum*, vol. 38, no. 3. Wiley Online Library, 2019, pp. 581–593.
- [79] D. Weng, C. Zheng, Z. Deng, M. Ma, J. Bao, Y. Zheng, M. Xu, and Y. Wu, "Towards better bus networks: A visual analytics approach," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 817–827, 2020.
- [80] Z. Deng, D. Weng, J. Chen, R. Liu, Z. Wang, J. Bao, Y. Zheng, and Y. Wu, "Airvis: Visual analytics of air pollution propagation," *IEEE transactions on visualization and computer graphics*, vol. 26, no. 1, pp. 800–810, 2019.
- [81] F. Miranda, H. Doraiswamy, M. Lage, K. Zhao, B. Gonçalves, L. Wilson, M. Hsieh, and C. T. Silva, "Urban pulse: Capturing the rhythm of cities," *IEEE transactions on visualization and computer* graphics, vol. 23, no. 1, pp. 791–800, 2016.
- [82] H. Doraiswamy, N. Ferreira, T. Damoulas, J. Freire, and C. T. Silva, "Using topological analysis to support event-guided exploration in urban data," *IEEE transactions on visualization and computer graphics*, vol. 20, no. 12, pp. 2634–2643, 2014.
- [83] P. Waddell, "Urbansim: Modeling urban development for land use, transportation, and environmental planning," *Journal of the American planning association*, vol. 68, no. 3, pp. 297–314, 2002.
- [84] M. Law and A. Collins, Getting to know ArcGIS. ESRI press Redlands, 2015.
- [85] J. Mueller, H. Lu, A. Chirkin, B. Klein, and G. Schmitt, "Citizen design science: A strategy for crowd-creative urban design," *Cities*, vol. 72, pp. 181–188, 2018.
- *Cities*, vol. 72, pp. 181–188, 2018.
  [86] Y. Zhang *et al.*, "Citymatrix: an urban decision support system augmented by artificial intelligence," Ph.D. dissertation, Massachusetts Institute of Technology, 2017.
- [87] A. Páez, D. M. Scott, and C. Morency, "Measuring accessibility: positive and normative implementations of various accessibility indicators," *Journal of Transport Geography*, vol. 25, pp. 141–153, 2012.
- [88] M. Iacono, K. Krizek, and A. M. El-Geneidy, "Access to destinations: How close is close enough? estimating accurate distance decay functions for multiple modes and different purposes," Minnesota Department of Transportation, Tech. Rep., 2008.

#### IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS

- [89] M. Rambonilaza and J. Dachary-Bernard, "Land-use planning and public preferences: What can we learn from choice experiment method?" *Landscape and urban planning*, vol. 83, no. 4, pp. 318–326, 2007.
- [90] J. Bao, Y. Zheng, and M. F. Mokbel, "Location-based and preference-aware recommendation using sparse geo-social networking data," in ACM SIGSPATIAL, 2012.
- [91] M. Frick and K. W. Axhausen, "Generating synthetic populations using Iterative Proportional Fitting (IPF) and Monte Carlo techniques," in 3rd Swiss Transport Research Conference, Monte Verità/Ascona, 2003.
- [92] Y. Zhu and J. Ferreira Jr, "Synthetic population generation at disaggregated spatial scales for land use and transportation microsimulation," *Transportation Research Record*, vol. 2429, no. 1, pp. 168–177, 2014.
- [93] J. Liu, X. Ma, Y. Zhu, J. Li, Z. He, and S. Ye, "Generating and visualizing spatially disaggregated synthetic population using a web-based geospatial service," *Sustainability*, vol. 13, no. 3, 2021.
- [94] M. Jaggi, "Revisiting frank-wolfe: Projection-free sparse convex optimization," in *Proceedings of the 30th international conference on machine learning*, 2013, pp. 427–435.
- [95] A. E. Roth, The Shapley value: essays in honor of Lloyd S. Shapley. Cambridge University Press, 1988.
- [96] J. Castro, D. Gómez, and J. Tejada, "Polynomial calculation of the shapley value based on sampling," *Computers & Operations Research*, vol. 36, no. 5, pp. 1726–1730, 2009.
- [97] "New york city department of city planning (DCP) land use code," https://www1.nyc.gov/site/planning/data-maps/open-data/dwn-pluto-mappluto.page, accessed: 2021-03-25.
- [98] N. Henry, J.-D. Fekete, and M. J. McGuffin, "Nodetrix: a hybrid visualization of social networks," *IEEE transactions on visualization and computer graphics*, vol. 13, no. 6, pp. 1302–1309, 2007.
- [99] J. R. Goodall, A. A. Ozok, W. G. Lutters, P. Rheingans, and A. Komlodi, "A user-centered approach to visualizing network traffic for intrusion detection," in CHI'05 Extended Abstracts on Human Factors in Computing Systems, 2005, pp. 1403–1406.
- [100] G. Levin, S. Funk, C. Sadowski, I. Pye, and S. Brandt, "Dpfair: A simple model for understanding optimal multiprocessor scheduling," in 2010 22nd Euromicro Conference on Real-Time Systems. IEEE, 2010, pp. 3–13.
- [101] C. Fieseler, E. Bucher, and C. P. Hoffmann, "Unfairness by design? the perceived fairness of digital labor on crowdworking platforms," *Journal of Business Ethics*, vol. 156, no. 4, pp. 987–1005, 2019.
- [102] C. Y. Yoo, "The effects of persuasion knowledge on click-through of keyword search ads: Moderating role of search task and perceived fairness," *Journalism & Mass Communication Quarterly*, vol. 86, no. 2, pp. 401–418, 2009.
- [103] A. Uhde, N. Schlicker, D. P. Wallach, and M. Hassenzahl, "Fairness and decision-making in collaborative shift scheduling systems," in *Proceedings of the 2020 CHI Conference on Human Factors* in Computing Systems, 2020, pp. 1–13.
- [104] T. Sühr, A. J. Biega, M. Zehlike, K. P. Gummadi, and A. Chakraborty, "Two-sided fairness for repeated matchings in two-sided markets: A case study of a ride-hailing platform," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 3082–3092.
- [105] Y. Xu and P. Xu, "Trading the system efficiency for the income equality of drivers in rideshare," in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, 2020, pp. 4199–4205.
- [106] Z. Xu, Z. Li, Q. Guan, D. Zhang, Q. Li, J. Nan, C. Liu, W. Bian, and J. Ye, "Large-scale order dispatch in on-demand ride-hailing platforms: A learning and planning approach," in *Proceedings* of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2018, pp. 905–913.
- [107] N. Raman, S. Shah, and J. Dickerson, "Data-driven methods for balancing fairness and efficiency in ride-pooling," in *Proceedings* of the Thirtieth International Joint Conference on Artificial Intelligence, 2021.
- [108] A. Pandey and A. Caliskan, "Disparate impact of artificial intelligence bias in ridehailing economy's price discrimination algorithms," in *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 2021, pp. 822–833.
- [109] Y. Lu, Y. Qi, S. Qi, Y. Li, H. Song, and Y. Liu, "Say no to price discrimination: decentralized and automated incentives for price

auditing in ride-hailing services," IEEE Transactions on Mobile Computing, pp. 1–1, 2020.

[110] M. K. Lee, D. Kusbit, E. Metsky, and L. Dabbish, "Working with machines: The impact of algorithmic and data-driven management on human workers," in *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, 2015, pp. 1603– 1612.



Yan Lyu received the Ph.D. degree in computer science from City University of Hong Kong, Hong Kong, in 2016, and the M.S. degree in pattern recognition and intelligent systems from University of Science and Technology of China, China, in 2013. She was a Postdoctoral Research Fellow with Hong Kong Baptist University, Hong Kong, in 2017, and with National University of Singapore, Singapore, from 2017 to 2020. She is an Associate Professor in the School of Computer Science and Engineering at South-

east University, China. Her research interests include data analytics and visualization, spatio-temporal data mining, and smart city.



Hangxin Lu received the M.Sc. degree in Electrical Engineering from EPFL, Lausanne, Switzerland, in 2016, and the PhD degree from ETH Zurich, Switzerland. She is currently a technology consultant working on big data governance. Her research interests include interactive data visualization, participatory design, smart and responsive cities, and digital urban planning applications.



Min Kyung Lee is an assistant professor in the School of Information at the University of Texas at Austin. She received a Ph.D. in Human-Computer Interaction from Carnegie Mellon University. She has conducted some of the first studies that empirically examine the social implications of algorithms' emerging roles in management and governance in society. Her current research interests include human-centered AI with a focus on human-centered perspectives on AI fairness and equity, and participatory design

methods for community-centered AI design.



Gerhard Schmitt is Director of the Singapore-ETH Centre in Singapore, Lead PI of the ETH Future Cities Responsive Cities Scenario, Lead PI of the Cooling Singapore Project and Professor Emeritus of Information Architecture at ETH Zurich. He holds a Dipl.-Ing. and a Dr.-Ing. degree from the Technical University of Munich and a Master of Architecture from the University of California, Berkeley. His research focuses on urban climate design, big data informed urban design, urban metabolism, smart and respon-

sive cities, simulation and visualization.



Brian Y. Lim received the BS degree in engineering physics from Cornell University, Ithaca, New York, in 2006, and the PhD degree in human-computer interaction from Carnegie Mellon University, Pittsburgh, Pennsylvania, in 2012. He is currently an assistant professor with the Department of Computer Science, National University of Singapore (NUS). His research interests include ubiquitous computing, explainable artificial intelligence, interfaces, and applications for urban data analytics and smart healthcare.

18