

Article

Not peer-reviewed version

Quantifying the Cross-Sectoral Intersecting Discrepancies Within Multiple Groups Using Latent Class Analysis Towards Fairness

[Yingfang Yuan](#) , Kefan Chen , Mehdi Rizvi , Lynne Baillie , [Wei Pang](#) *

Posted Date: 12 June 2024

doi: 10.20944/preprints202406.0700.v2

Keywords: Machine Learning; AI for Social Science; AI bias



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Quantifying the Cross-sectoral Intersecting Discrepancies within Multiple Groups Using Latent Class Analysis Towards Fairness

Yingfang Yuan [†], Kefan Chen [†], Mehdi Rizvi, Lynne Baillie and Wei Pang ^{*}

School of Mathematical and Computer Sciences Heriot-Watt University

^{*} Correspondence: w.pang@hw.ac.uk

[†] These authors contributed equally to this work.

Abstract: The growing interest in fair AI development is evident. The “Leave No One Behind” initiative urges us to address multiple and intersecting forms of inequality in accessing services, resources, and opportunities, emphasising the significance of fairness in AI. This is particularly relevant as an increasing number of AI tools are applied to decision-making processes, such as resource allocation and service scheme development, across various sectors such as health, energy, and housing. Therefore, exploring joint inequalities in these sectors is significant and valuable for thoroughly understanding overall inequality and unfairness. This research introduces an innovative approach to quantify cross-sectoral intersecting discrepancies among user-defined groups using latent class analysis. These discrepancies can be used to approximate inequality and provide valuable insights to fairness issues. We validate our approach using both proprietary and public datasets, including EVENS and Census 2021 (England & Wales) datasets, to examine cross-sectoral intersecting discrepancies among different ethnic groups. We also verify the reliability of the quantified discrepancy by conducting a correlation analysis with a government public metric. Our findings reveal significant discrepancies between minority ethnic groups, highlighting the need for targeted interventions in real-world AI applications. Additionally, we demonstrate how the proposed approach can be used to provide insights into the fairness of machine learning.

Keywords: machine learning; AI for social science; AI bias

1. Introduction

As AI systems become increasingly prevalent, ensuring fairness in their design and implementation is crucial [1]. AI fairness research spans various sectors, including healthcare [2–4], finance [5], and education [6]. However, studies on cross-sectoral intersecting fairness in AI are limited. Stiglitz [7] highlights that economic inequalities extend beyond income and wealth, affecting sectors such as education and health, implying correlated inequalities. At the same time, the “Leave No One Behind” (LNOB) principle urges addressing multiple, intersecting inequalities that harm individuals’ rights [8]. Therefore, we believe that research exploring joint inequalities across sectors is necessary to comprehensively understand overall inequality and gain insights into unfairness.

Bias, Inequality, Fairness Many obstacles to accessing services, resources, and opportunities result from discriminatory laws, policies, and societal norms that marginalise specific groups [8]. Recent studies highlight concerns that AI-supported decision-making systems may also be influenced by biases [9], which can unfairly impact vulnerable groups, such as ethnic minorities, emphasising the need for research on AI system fairness [10]. A primary obstacle in advancing and implementing fair AI systems is the presence of bias [11]. In AI, bias can originate from various sources, including data collection, algorithmic design, and user interaction, as illustrated in Figure 1.




Figure 1. The loop of bias placed in the data, algorithm, and user interaction feedback [1].




Figure 2. The correlations between fairness, bias, inequality, and discrepancy in the context of AI [11,12].

Most AI systems heavily rely on data for their operations. This close connection means that any biases in the training data can be embedded into the algorithms, leading to biased predictions. Even if the data itself is not inherently biased, algorithms can still exhibit biased behaviour due to inappropriate design choices. These biased outcomes can influence AI systems in real-world applications, creating a feedback loop where biased data from user interactions further trains and reinforces biased algorithms, resulting in a vicious cycle [1].

From our perspective, bias stemming from data plays a crucial role in achieving fairness, as it can trigger a cascade of other biases, exacerbating fairness issues. Additionally, data bias may reflect various forms of social inequality in the real world. Inequality refers to disparities in opportunities and rewards based on different social positions or statuses within a group or society [13]. Inequality is considered unfair due to its causes, such as discrimination or failures in providing equal opportunities, and its consequences, such as resulting in objectionable disparities in status or power [13]. As previously mentioned, cross-sectoral intersecting inequality is crucial. Therefore, quantifying these disparities within the data is essential for understanding inequality and bias, ultimately leading to the development of fairer AI systems.

Discrepancy According to the LNOB principle of equal opportunities, everyone should **ideally** have **equal** access to public services and resources without discrepancies. We prefer the term “discrepancy” over “disparity” or “difference” because it suggests an unexpected difference. Discrepancies in data may indicate unfairness, biases, and inequalities, given that individuals should **ideally** be treated **equally**. It is worth noting that bias can arise from various inequalities, though not all inequalities are biases [13]. However, in the context of LNOB, for this study, “discrepancy” is more appropriate when discussing **equal opportunity**, as it can be seen as a specific type of inequality causing unfairness.

Figure 2 illustrates the correlations between fairness, bias, inequality, and discrepancy. Essentially, fairness is indirectly linked with discrepancy, and discrepancy can contribute to unfairness. The difference between fairness and bias is that the former can be viewed as a technical issue, while the latter can be viewed as a social and ethical issue [11]. Furthermore, bias is a problem caused by historical and current social inequality [12], and inequality can manifest as discrepancies. Figure 2 starts with discrepancy and moves through inequality and bias to fairness. Therefore, *our research focuses on quantifying cross-sectoral intersecting discrepancies between different groups, uncovering insights or patterns related to inequality, bias, and AI fairness.*

Background and Motivation Currently, there is limited research focusing on quantifying discrepancies. Most recent research on quantifying bias and/or inequality primarily revolves around resource allocation strategies and generally relies on objective data (e.g., [14]). However, these approaches have limitations and face challenges in effectively assessing and measuring biases in datasets unrelated to resource allocation. For example, in social science, much data is collected through questionnaires, which often include binary, categorical, or ordinal data types related to subjective responses and user experiences. These questionnaires may cover various aspects, resulting in intersecting, cross-sectoral, and high-dimensional data. Analysing data based on no more than two dimensions or sectors may overlook important information or patterns. Therefore, we believe that quantifying cross-sectoral intersecting discrepancies is valuable, as it can provide comprehensive insights and joint information.

The *Anonymous project*¹ aims to establish safer online environments for minority ethnic (ME) populations in the UK. Its survey questionnaire covers five key aspects: demography, energy, housing,

¹ Project name removed to maintain anonymity

health, and online services. Notably, the data collected in this context does not directly pertain to resource allocation, making it challenging to explicitly define and detect bias within the data using current methods, despite the presence of discrepancies. These discrepancies may arise from various factors, including culture, user experience, and discrimination, potentially contributing to bias or unfairness. This research is mainly motivated by the *Anonymous* project, so the datasets we used primarily cover the health, energy, and housing sectors, with our research targeting ME groups.

In our preliminary research (see Appendix A) in England, based on a *Anonymous* project survey question regarding health and digital services², we observed a notable discrepancy in the Chinese group: 30.16% lacked English proficiency and 26.98% struggled to use the online system, while most other ethnic groups reported no concerns. These discrepancies, stemming from cultural differences, user experiences, or discrimination, contribute to inequality and may affect AI fairness. For instance, an AI system might wrongly assume Chinese individuals need more English support, leaving other ME groups with less assistance. Ideally, the percentages should be similar across ethnic groups in the LNOB context, indicating equal treatment or experience. However, investigating multiple and cross-sectoral questions simultaneously is challenging. For instance, the *Anonymous* project's data contains similar questions for the energy and housing sectors, and current methods struggle to analyse these sectors jointly. Therefore, we propose an approach to quantify intersecting and cross-sectoral discrepancies for multiple ethnic groups by leveraging latent class analysis (LCA) [15].

LCA is a popular method in social science [16] because it identifies latent groups within a population based on observed characteristics or behaviours. LCA offers a flexible framework for exploring social phenomena and integrating with other analytical techniques. In this research, we use LCA to cluster intersecting and cross-sectoral data, encompassing questions across the health, energy, and housing sectors. This method enables us to derive latent classes and outcomes, with each class describing a distinct cross-sectoral user profile. This approach moves beyond defining user classes solely based on individual questions. More details of our proposed approach are presented in Section 2, with experiments and results reported in Section 4.

The main contributions of this research are as follows: (1) we propose a novel and generic approach to quantify intersecting and cross-sectoral discrepancies between user-defined groups; (2) our findings reveal that ME groups cannot be treated as a homogeneous group, as varying discrepancies exist among them; and (3) we demonstrate how the proposed approach can be used to provide insights to AI fairness.

2. Quantifying the Cross-Sectoral Intersecting Discrepancies

The overall workflow of the proposed approach is shown in Figure 3, using a binary-encoded survey data as an example. It is noted that our approach is not limited to this specific format and can be applied to a wide range of similar problems. We will present more experiments with other datasets in Section 4 to further validate and showcase the features of our approach.

In Figure 3, Stage (1) shows the binary-encoded data D , where $\{Q_1, Q_2, \dots\}$ represents the selected questions from the survey data, covering user experiences across sectors. Meanwhile, $\{u_1, u_2, u_3, \dots\}$ represents the collection of survey respondents. Here, q represents the response options, for example, $q_{2,1}$ denotes the first option for Q_2 . A '1' indicates a selected option, while '0' indicates it was not selected. For other datasets, the encoding method should be chosen based on the data format and type. The blue arrow in Figure 3 illustrates the LCA process³, which includes hyperparameter selection and model fitting.

² Question 21: Which of the following concerns do you have about communicating with your general practice (GP) through apps, websites, or other online services?

³ StepMix (<https://stepmix.readthedocs.io/en/latest/index.html>) Python repository is used to implement LCA in this research.

The advantage of using LCA is straightforward: it considers the joint probability distribution of all variables. This means potential inequalities or discrepancies can be analysed jointly. Once we obtain the distributions of latent classes $\{c_1, c_2, c_3, \dots\}$ over user-defined groups $\{e_1, e_2, e_3, \dots\}$ (as shown in Figure 3 Stage (2)), we can calculate the discrepancy Δ .




Figure 3. The process of the proposed approach for quantifying cross-sectoral discrepancies within different groups.

Algorithm 1 Quantifying the Intersecting Discrepancies within Multiple Groups

```

1: Input:  $D$  and  $G$   $\triangleright G$  denotes a set of user-defined groups
2:
3: Initialise  $M$   $\triangleright$  Create LCA model
4:
5: Estimate  $M$  based on  $D$ 
6:
7: for  $e$  in  $G$  do
8:   for  $c$  in  $C$  do
9:      $r_c^e = N_c^e / N^e$ 
10:   end for
11: end for
12:
13: for  $e$  in  $G$  do
14:   for  $e'$  in  $G$  do
15:      $\Delta_{ee'} = 1 - \frac{\mathbf{R}_e \cdot \mathbf{R}_{e'}}{\|\mathbf{R}_e\| \|\mathbf{R}_{e'}\|}$   $\triangleright$  Pair-wise Calculation
16:   end for
17: end for
18: Output: Discrepancy matrix  $S$  with size  $|G| \times |G|$ 

```

Quantification of Discrepancy Let us denote the size of the dataset as N , where $i \in 1, 2, \dots, N$ represents an individual sample. Concurrently, let $c \in C$ denote a latent class, with the total number of classes being $|C|$, and let N_c represent the count of samples classified into latent class c . To quantify the discrepancies, it is necessary to establish a grouping variable G , which can be defined based on factors such as ethnicity, age, or income level. Here, $|G|$ denotes the total number of user-defined groups, $e \in G$ represents one specific group within this set, N^e denotes the number of individuals from group e , and N_c^e denotes the number of individuals from group e assigned to the class c .

To initiate the quantification process, the proportions r of samples from each user-defined group within each latent class need to be calculated, as detailed in Stage (3) in Figure 3. This calculation can be performed using $r_c^e = N_c^e / N^e, \forall e \in G, c \in C$. The reason for calculating r is that user-defined groups may have different numbers of samples; therefore, using percentages for subsequent analyses ensures fairness and consistency.

Subsequently, we can derive a matrix of results characterised by dimensions $|G| \times |C|$, as shown in Figure 3 (3). Within this matrix, each row corresponds to the proportions of samples from a specific group within each latent class. It is important to note that in this context, each latent class effectively represents an individual user profile and can be viewed as a distinctive feature. Consequently, each row within the matrix may be employed as a feature vector denoted as R_e , serving as a representation of a specific group within the feature space.

In the assessment of discrepancy between two vectors, various methods may be employed, including the Euclidean distance, Kullback-Leibler Divergence, Earth Mover’s Distance, and Manhattan Distance, among others. In our approach, we propose the utilisation of Cosine Similarity to calculate

the discrepancy, which is defined as $\Delta = 1 - \cos(\theta) = 1 - \frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\| \|\mathbf{B}\|}$. Finally, we can iteratively calculate Δ between any pairs of vectors R and obtain the discrepancy matrix S (as shown in Stage (4) of Figure 3). The AVG column in Stage (4) contains the mean discrepancy values for e , which can be viewed as an approximation for how each e is different from others.

We suggest the use of Cosine Similarity due to its inherent characteristics, including a natural value range spanning from 0 to 1. Importantly, it does not necessitate additional normalisation procedures. This metric, possessing with a fixed value range, enhances comparability and offers support to subsequent AI fairness research. The proposed approach is summarised in Algorithm 1.

3. Related Work

Quantifying and improving AI fairness As AI technologies are used more and more frequently in real life, people's concerns about the ethics and fairness of AI have always existed, especially when AI is increasingly used in problems with sensitive data [17]. Morley et al. [10] and Garattini et al. [18] noticed that an algorithm "learns" to prioritise patients it predicts to have better outcomes for a particular disease. And they also noticed that AI models have discriminatory potential when facing ME groups on health. Therefore, people are paying more and more attention on the impact and mitigating methods of AI bias.

Wu et al. [14] proposes the allocation-deterioration framework for detecting and quantifying health inequalities induced by AI models. This framework quantifies inequalities as the area between two allocation-deterioration curves. They conducted experiments on synthetic datasets and real-world ICU datasets to assess the framework's performance and applied the framework to the ICU dataset and quantified the unfairness of AI algorithms between White and Non-White patients. So et al. [19] explores the limitations of fairness in machine learning and proposes a reparative approach to address historical housing discrimination in the US. In that work, they used contemporary mortgage data and historical census data to conduct case studies to demonstrate the impact of historical discrimination on wealth accumulation and estimate housing compensation costs. They then proposed a remediation framework that includes analysing historical biases, intervening in algorithmic systems, and developing machine learning processes that reduce correct historical harms.

Latent Class Analysis (LCA) is a statistical method based on mixture models and often used to detect potential or unobserved heterogeneity in samples [20]. By analysing response patterns of observed variables, LCA can identify potential subgroups within a sample set [21]. The basic idea of LCA is that some parameters of a postulated statistical model differ across unobserved subgroups, forming the categories of a categorical latent variable [22]. In 1950, Lazarsfeld [23] introduced LCA as a means of constructing typologies or clusters using dichotomous observed variables. Over two decades later, Goodman [24] enhanced the model's practical applicability by devising an algorithm for obtaining maximum likelihood estimates of its parameters. Since then, many new frameworks have been proposed, including models with continuous covariates, local dependencies, ordinal variables, multiple latent variables, and repeated measures [22].

Because LCA is a person-centered mixture model, it is widely used in sociology and statistics to interpret and identify different subgroups in a population that often share certain external characteristics from data [25]. However, in social sciences, LCA is used in cross-sectional and longitudinal studies. For example, in relevant studies in psychology [26], social sciences [27], and epidemiology [28], mixed models and LCA can be used to establish probabilistic diagnoses when no suitable gold standard is available [15].

In [26], the relationship between cyberbullying and social anxiety among Hispanic adolescents was explored. The sample consisted of 1,412 Spanish secondary school students aged 12 to 18 years. There were significant differences in cyberbullying patterns across all social anxiety subscales after applying LCA. Compared with other profiles, students with higher cyberbullying traits scored higher on social avoidance and distress in social situations, as well as lower levels of fear of negative evaluation and distress in new situations. Researchers in [27] developed a tool, using LCA, to characterise

energy poverty without the need to arbitrarily define binary cutoffs. The authors highlight the need for a multidimensional approach to measuring energy poverty and discuss the challenges of identifying vulnerable consumers. The research in [28] aimed to identify subgroups in COVID-19-related acute respiratory distress syndrome (ARDS) and compare them with previously described ARDS subphenotypes by using LCA. The study found that there were two COVID-19-related ARDS subgroups with differential outcomes, similar to previously described ARDS subphenotypes.

4. Experiments

4.1. The Anonymous project

Within the *Anonymous* project⁴, a systematic online survey was conducted to explore the experiences of individuals from ME groups in the UK, focusing on digitalised health, energy, and housing aspects. To examine cross-sectoral intersecting discrepancies among the seven ethnic groups (as shown in Table 1), we selected three questions from the *Anonymous* survey data related to these sectors. The answers from ME participants formed the dataset for this experiment. For England and Scotland, we have 594 and 284 samples, respectively. Due to varying sample sizes across ethnicities, we calculated discrepancies separately for each region.

Selecting the number of latent classes in LCA requires presetting. To address this, we conducted hyperparameter optimisation for each experiment to find the elbow point, applying this optimisation to all subsequent experiments.

The discrepancy results for England are presented in the left part of Table 1. The table shows that the Chinese group has the largest average (AVG) discrepancy value compared to other groups. Meanwhile, the Indian group exhibits the smallest discrepancy with the Bangladeshi group and also shows similarity to the Pakistani group. This is likely due to their close geographical locations and similar cultural backgrounds and lifestyles.

The right part of Table 1 presents the results for Scotland, showing similar outcomes: the Chinese group is distinct from others, while the Bangladeshi group is similar to the Pakistani group. The differences between England and Scotland may be attributed to their different policies and circumstances. We hypothesise that the primary reason for the Chinese group standing out is a lack of English proficiency. This is supported by our preliminary research mentioned in Section 1, which shows a significant number of Chinese participants expressing this concern. Additionally, BBC News [29] reports that the Chinese community experiences some of the highest rates of racism among all ethnic groups in the UK. Our discrepancy values may help explain this, as the Chinese group shows different experiences in digitalised online services.

Table 1. The matrix of discrepancies between 7 ethnic groups for *Anonymous* project’s England and Scotland data. The AVG denotes the average discrepancy value for one group.

	England								Scotland							
	African	Bangladeshi	Caribbean	Chinese	Indian	Mixed Group	Pakistani	AVG	African	Bangladeshi	Caribbean	Chinese	Indian	Mixed Group	Pakistani	AVG
African	0.0000	0.0899	0.0383	0.1810	0.0547	0.0590	0.0517	0.0678	0.0000	0.0666	0.0296	0.1956	0.0342	0.0118	0.0227	0.0515
Bangladeshi	0.0899	0.0000	0.1359	0.3734	0.0200	0.2738	0.0308	0.1320	0.0666	0.0000	0.0324	0.3043	0.0989	0.0563	0.0118	0.0815
Caribbean	0.0383	0.1359	0.0000	0.2456	0.0764	0.1131	0.0951	0.1006	0.0296	0.0324	0.0000	0.3430	0.0191	0.0546	0.0159	0.0706
Chinese	0.1810	0.3734	0.2456	0.0000	0.3700	0.1201	0.2459	0.2194	0.1956	0.3043	0.3430	0.0000	0.3717	0.1334	0.2438	0.2274
Indian	0.0547	0.0200	0.0764	0.3700	0.0000	0.2139	0.0311	0.1094	0.0342	0.0989	0.0191	0.3717	0.0000	0.0821	0.0575	0.0948
Mixed Group	0.0590	0.2738	0.1131	0.1201	0.2139	0.0000	0.1987	0.1398	0.0118	0.0563	0.0546	0.1334	0.0821	0.0000	0.0209	0.0513
Pakistani	0.0517	0.0308	0.0951	0.2459	0.0311	0.1987	0.0000	0.0933	0.0227	0.0118	0.0159	0.2438	0.0575	0.0209	0.0000	0.0532

4.2. EVENS

The Centre on the Dynamics of Ethnicity (CoDE), funded by the Economic and Social Research Council (ESRC), conducted “The COVID Race Inequalities Programme”. As part of this project, CoDE

⁴ More dataset details can be found in the Appendix A.

carried out the Evidence for Equality National Survey (EVENS)⁵, which documents the lives of ethnic and religious minorities in Britain during the coronavirus pandemic. The EVENS dataset comprises 14,215 data points and categorizes participants into 18 different ethnic minority groups. To facilitate a comparative analysis with the *Anonymous* project’s experiment, we focused on the same seven ethnic groups from *Anonymous* project, resulting in a filtered dataset of 4,348 participants from England and 253 participants from Scotland. We excluded entries with missing values to ensure the robustness of our analysis.

It is important to note that the EVENS dataset uses a different definition for mixed or multiple ethnic groups compared to *Anonymous* project. To avoid inconsistencies that could affect the final analysis, we excluded mixed or multiple ethnic groups from the EVENS calculations. Due to the different questions in the EVENS and *Anonymous* project surveys, we selected three types of cross-sectoral questions from EVENS for analysis: housing, experiences of harassment, and financial situation. These questions were chosen to provide a broad view of participants’ living conditions, social experiences, and economic status during the pandemic.

We applied LCA separately to the data from England and Scotland to uncover patterns and discrepancies within and between these regions. This approach allows us to identify distinct subgroups within the ethnic communities based on their responses to the selected questions, providing deeper insights into the intersectional and cross-sectoral experiences of these groups.

Seeing the discrepancies in the EVENS England data, shown in the left part of Table 2, it is clear that all discrepancy values are relatively small compared to the *Anonymous* project results. Notably, there are large discrepancies between the Indian and Bangladeshi groups, and between the Indian and Chinese groups. This differs slightly from the conclusions of the *Anonymous* project’s experiment, likely due to the different types of cross-sectoral questions selected. In the AVG column, the Bangladeshi group has the highest values, indicating they experienced COVID-19 differently. The Census 2021 reported that COVID-19 mortality rates were highest for the Bangladeshi group, for both males and females [30], supporting our findings. The right part of Table 2 describes the discrepancies for the EVENS Scotland data. There are significant discrepancies between the Caribbean and other ethnic groups, likely related to the unique background of the Caribbean group.

Table 2. The matrix of discrepancies for EVENS England and Scotland data.

	England							Scotland						
	African	Bangladeshi	Caribbean	Chinese	Indian	Pakistani	AVG	African	Bangladeshi	Caribbean	Chinese	Indian	Pakistani	AVG
African	0.0000	0.0112	0.0014	0.0038	0.0062	0.0018	0.0040	0.0000	0.0246	0.5408	0.0300	0.0637	0.1800	0.1398
Bangladeshi	0.0112	0.0000	0.0102	0.0031	0.0227	0.0090	0.0094	0.0246	0.0000	0.5283	0.0522	0.1539	0.2697	0.1714
Caribbean	0.0014	0.0102	0.0000	0.0047	0.0030	0.0002	0.0032	0.5408	0.5283	0.0000	0.3946	0.4505	0.2506	0.3608
Chinese	0.0038	0.0031	0.0047	0.0000	0.0138	0.0048	0.0050	0.0300	0.0522	0.3946	0.0000	0.0662	0.1036	0.1078
Indian	0.0062	0.0227	0.0030	0.0138	0.0000	0.0040	0.0083	0.0637	0.1539	0.4505	0.0662	0.0000	0.0589	0.1322
Pakistani	0.0018	0.0090	0.0002	0.0048	0.0040	0.0000	0.0033	0.1800	0.2697	0.2506	0.1036	0.0589	0.0000	0.1438

4.3. Census 2021 (England and Wales)

To further test our approach, we applied it to the Census 2021 dataset [31], which gathers information on individuals and households in England and Wales every decade. These data help plan and finance essential local services. We compared our results with the UK deprivation indices data from 2019 [32], which classify relative deprivation in small areas. We hypothesised that discrepancy values should correlate with the deprivation indices, reflecting discrepancies across energy, health, housing, and socioeconomic sectors. The key difference is that our discrepancy values are data-driven, while deprivation indices are based on human-centred assessments, suggesting that our approach is complementary.

⁵ <https://www.evensusurvey.co.uk/>

For our experiments, we selected four cross-sectoral questions from the census related to energy (type of central heating), health (general health), housing (occupancy rating for bedrooms), and socioeconomic status (household deprivation). Note that the socioeconomic data in Census 2021 differ from the 2019 deprivation indices due to different definitions and coverage [31,32].

Additionally, for Census 2021, we selected Lower Layer Super Output Areas (LSOAs) [31] as samples instead of individuals, as individual data were not accessible. After cleaning the data and removing unmatched LSOAs, we had 31,810 LSOAs in total. Unmatched LSOAs, which appear only in either Census 2021 or Deprivation 2019, were removed. In the Deprivation 2019 dataset, each LSOA is labelled with a deprivation level from 1 to 10 (1 being the most deprived). As our samples are LSOAs, we quantified discrepancies between different LSOAs. Since the raw data does not include group attributes, we classified LSOAs into five groups based on the percentage of the population from ME groups: [0%, 20%), [20%, 40%), [40%, 60%), [60%, 80%), and [80%, 100%].

The proposed approach quantifies the discrepancies between the defined ME population-related groups based on the selected Census 2021 data. The results are shown in the left part of Table 3. It is evident that the discrepancies between LSOAs increase as differences in ME population percentages increase, indicating significant disparities in living conditions for ME individuals across different LSOAs, particularly in terms of energy, housing, and health aspects. Notably, the 0%-20% group shows the largest AVG discrepancy compared to other groups, suggesting that White individuals in those LSOAs experience significantly different living conditions. Since the 0%-20% group constitutes a large portion of the UK (see Appendix B), this finding suggests potential unequal treatment and possible neglect of other LSOAs. Additionally, the 40%-60% group has the smallest AVG discrepancy value, likely due to its intermediate position among the ME groups, sharing characteristics with the 0%-20% and 20%-40% groups, as well as 60%-80% and 80%-100% groups.

Table 3. The matrices for the discrepancies of Census 2021 and the deprivation discrepancies between LSOAs across user-defined ME population percentage groups based on Deprivation 2019.

	Census						Deprivation					
	0-20%	20-40%	40-60%	60-80%	80-100%	AVG	0-20%	20-40%	40-60%	60-80%	80-100%	AVG
0-20%	0.0000	0.4865	0.6783	0.8603	0.9347	0.5920	0.0000	0.2001	0.2896	0.3877	0.5064	0.2768
20-40%	0.4865	0.0000	0.1371	0.3934	0.5565	0.3147	0.2001	0.0000	0.0313	0.1123	0.2203	0.1128
40-60%	0.6783	0.1371	0.0000	0.1173	0.2744	0.2414	0.2896	0.0313	0.0000	0.0314	0.0963	0.0897
60-80%	0.8603	0.3934	0.1173	0.0000	0.0445	0.2831	0.3877	0.1123	0.0314	0.0000	0.0283	0.1119
80-100%	0.9347	0.5565	0.2744	0.0445	0.0000	0.3620	0.5064	0.2203	0.0963	0.0283	0.0000	0.1703

Correlation Analysis Furthermore, based on the deprivation indices from Deprivation 2019 [32] and the defined groups, we calculated the percentages of LSOAs in each deprivation-labelled group across various ME population groups. The results are shown in Appendix B, and we treated each row as a feature vector representing one group of LSOAs. We then iteratively calculated the deprivation discrepancies for each pair of rows, with the results presented in the right part of Table 3. We observed similar patterns (color change) to those in the left part of Table 3, which can verify the reliability of our proposed approach.

To statistically verify our proposed approach, we ran Pearson and Spearman row-wise correlation analyses for Census 2021 discrepancies (the left part of Table 3) and Deprivation 2019 discrepancies (the right part of Table 3). The detailed results are shown in Appendix B. All rows exhibit very strong correlations, implying that our approach can draw conclusions very similar to those of experts. Furthermore, we also flattened both matrices to run a one-time correlation analysis. The Pearson correlation coefficient is 0.9797 with a *p*-value of 1.4437e-17, and the Spearman correlation coefficient is 0.9872 with a *p*-value of 7.436e-20. Both *p*-values are far less than 0.001, indicating a strong correlation.

Discrepancy for AI fairness Now, we will show how discrepancies relate to potential AI bias. We selected logistic regression to classify deprivation indices for LSOAs using the Census 2021 data

from previous experiments. To simplify the classification task, we redefined the deprivation indices as deprived (indices 1-5, labelled as 0) and not deprived (indices 6-10, labelled as 1). We split the dataset into training and validation sets in an 8:2 ratio. The prediction accuracy on the Census 2021 dataset, with an overall accuracy of 90.35% and a standard deviation (STD) of 4.20 across five group results, is shown in the Census column of Table 4. Meanwhile, the accuracy for each group is displayed on the left in Figure 4. We noticed that the accuracy varies across different groups, with the 0-20% group showing 100% accuracy. In the context of LNOB and AI fairness, we consider this biased and problematic. Additionally, in this research, the STD is considered an important indicator of fairness; smaller STD values imply that the model treats each group more equally. In the following paragraph, we will discuss how the discrepancy relates to AI fairness based on two undersampled datasets.

For the Census 2021 data, we noticed two data imbalance issues likely affecting accuracy. Firstly, the 0-20% group constitutes 82.59% of the LSOA samples and has the largest AVG discrepancy and relatively low accuracy, potentially indicating distinct features compared to other groups. Generally, in machine learning, imbalanced data can negatively impact the majority, leading to issues like overfitting. To address this, we used random undersampling [33], resampling all classes except the minority 80-100% group, which had 125 samples. After undersampling, all groups had an equal number of samples, and we split them into training and test sets with a ratio of 8:2. This method aimed to observe changes in discrepancies, accuracy, and STD to gain insights into bias. The results are shown in the middle of Figure 4 and in the Undersampled Census (ME) column of Table 4.

We found that all AVG discrepancies increased along with the STD (from 4.2 to 4.83), and the prediction accuracy decreased for three groups. This may indicate that our approach efficiently detects data discrepancies that can exacerbate bias issues. Notably, the 40-60% group showed the most significant increase in discrepancy, accompanied by a dramatic decrease in accuracy. We believe that this phenomenon is due not only to the increase in discrepancies but also to the reduced data size, which may not provide enough data for effective model training. When samples from different groups are far apart, the model may require more data to train effectively; otherwise, it may focus on some subareas of the feature space. It is worth noting that while the overall accuracy increases, the accuracy of most groups decreases due to a statistical artifact related to changes in test data size. In this research, our primary focus is on the STD and discrepancies.

Furthermore, another data imbalance issue may be related to deprivation labels. Specifically, in the original Census 2021 data, all groups except the 0-20% group have more data labelled as 0 than 1. Additionally, the ratio of data labelled as 0 to 1 increases with the percentage of the ME population. For example, in the 40-60% group, the number of samples for 0 and 1 labels are 1,014 and 215, respectively, and 537 and 54 for the 60-80% group. We believe this may contribute to bias issues. Therefore, we conducted undersampling on the original Census 2021 data, targeting the labels. After undersampling, the dataset had an equal number of samples for 0 and 1 labels. Then, we split the dataset into training and test sets following the same rule. The results are shown on the right of Figure 4 and in the Undersampled Census (deprivation) column of Table 4.

Overall, we found that all AVG discrepancies decreased, indicating that samples from different groups became closer to each other. This suggests that the model is likely to treat them more similarly in the feature space. On the right of Figure 4, we observed that accuracy across different groups became more similar compared to the other two figures (left and middle), supported by the STD dramatically decreasing to 1.82. In the context of LNOB, the model is now less biased. Meanwhile, the overall accuracy is better than the original one. In general, this set of experiments demonstrated that data discrepancies obtained from our proposed approach can potentially indicate the degree to which the model treats each group equally.




Figure 4. Prediction accuracy for classification tasks on the Census 2021 data (left), undersampled data targeting ME groups (middle), and undersampled data targeting deprivation labels (right).

Table 4. The results of discrepancies and model performance on the original Census dataset and two undersampled Census datasets.

Model		Census	Undersampled Census (ME)	Undersampled Census (deprivation)
	Accuracy ↑	90.35	91.20	90.75
Discrepancy AVG	STD ↓	4.20	4.83	1.82
	0-20%	0.5920	0.6398	0.4966
	20-40%	0.3147	0.3242	0.1428
	40-60%	0.2414	0.3340	0.1394
	60-80%	0.2831	0.3278	0.1584
	80-100%	0.3620	0.4340	0.1787

5. Conclusions and Limitations

In conclusion, the issue of AI fairness is of paramount importance and warrants attention from all stakeholders. In our research, we addressed this challenge by focusing on quantifying the discrepancies present in data, recognising that AI models heavily rely on data for their performance. Our proposed data-driven approach is aligned with the LNOB initiative, as it aids in discovering and addressing discrepancies between user-defined groups, thus contributing to efforts to mitigate inequality. Moreover, we believe that our proposed approach holds promise for applications across a broad spectrum of tasks, offering insights to develop fair AI models. Through testing on three datasets, we have demonstrated the efficacy and informativeness of our approach, yielding satisfactory results. Our proposed approach can be considered as an approximation of bias, as selecting different parameters for LCA may yield slightly varying results, to address this we have done hyperparameter optimisation.

In summary, our research represents a significant step towards promoting fairness in AI and offers an innovative avenue for social science research. By highlighting data-driven approaches and their alignment with broader societal initiatives, we aim to foster a more equitable and inclusive landscape for AI development and deployment.

Appendix A The Anonymous project

As part of the *Anonymous* project, a multilingual online survey, available in 10 languages, was conducted to investigate the experiences of individuals from minority ethnic groups with digitalised housing, health, and energy services. The survey contains a total of 32 questions. The survey data includes 594 responses from England and 284 responses from Scotland. In terms of respondent selection, researcher carefully determined the required number of participants from each ethnic group in England using a proportional allocation method based on their respective population percentages from the 2021 Census (England & Wales). However, due to the unavailability of Scotland’s 2021 census results during the planning phase, the project aimed to limit the number of respondents from each ME group in Scotland to a maximum of 40. The total number of survey respondents was 878, with 594 participants from England and 284 from Scotland. A detailed breakdown of the respondents’ demographic information, including their ethnicities, is provided in Table A1.

Table A1. The number of respondents from each ME group in both England and Scotland

Anonymous project’s Target Ethnic Group	England	Scotland	Total
African	176	37	213
Bangladeshi	97	41	138
Indian	93	40	133
Chinese	63	39	102
Pakistani	62	40	102
Caribbean	47	32	79
Mixed or Multiple ethnic groups	56	55	111
Total	594	284	878

We show the distribution of responses from seven ethnic groups in England regarding health and digital services based on *Anonymous* project’s data. We observed a notable discrepancy in the Chinese group, with 30.16% lacking English proficiency and 26.98% struggling to use the online system, while most other ethnic groups reported no concerns.




Figure A1. This figure gives an example of the exploratory data analysis (EDA) for the *Anonymous* project’s England data focuses on Question 21: “Which of the following concerns do you have about communicating with your general practice (GP) through apps, websites, or other online services?” The bars represent the portion of respondents from each ethnic group selecting each option.

It should be noted that the definitions of ethnicities differ between the Scottish and English censuses. For instance, while individuals in the English census may choose “Pakistani” as a sub-category under the broader category “Asian or Asian British”, in Scotland, this category translates to ‘Pakistani, Pakistani Scottish, or Pakistani British’. For consistency, this paper adopts the ethnicity naming conventions used in the English census.

The questions we selected for quantifying discrepancies include: Question 20: Which of the following concerns do you have about communicating with your GP through apps, websites or other online services? Question 27: Do you have any concerns about using an app, website or digital system to carry out energy-related activities? Question 23: Do you have any concerns about using an app, website or other digital service for these housing-related activities?

Regarding the selection of the number of latent classes, we used grid search for hyperparameter optimisation with 10-fold cross-validation. For all experiments, the search space was adjusted based on our empirical observations. In the *Anonymous* project experiments, we set the range from 2 to 10. For EVENs and Census experiments, we set the range from 2 to 30.




Figure A2. The example of hyperparameter optimisation process to seek the elbow point.

Appendix B Census 2021 (England and Wales)

Table A2. The correlation analysis between Deprivation 2019 and the results obtained from our proposed based on Census 2021.

	Pearson	Spearman
0-20%	0.9802	1
20-40%	0.9769	1
40-60%	0.9949	0.9
60-80%	0.9829	1
80-100%	0.9830	1

Table A3. The percentages of LSOAs in different ME population groups cross 10 deprivation levels based on the Deprivation 2019 dataset and the total number LSOAs in each group.

	1	2	3	4	5	6	7	8	9	10	Total LSOAs
0-20%	8.62%	7.83%	8.28%	9.14%	10.04%	10.49%	11.03%	11.32%	11.52%	11.75%	26273
20-40%	12.22%	18.93%	18.12%	14.81%	10.77%	8.05%	5.96%	4.04%	3.79%	3.31%	3592
40-60%	20.10%	23.19%	17.49%	13.02%	8.71%	6.18%	4.48%	3.58%	2.12%	1.14%	1229
60-80%	31.13%	22.67%	16.58%	12.35%	8.12%	4.91%	2.71%	1.18%	0.34%	0.00%	591
80-100%	41.60%	26.40%	20.00%	7.20%	1.60%	2.40%	0.80%	0.00%	0.00%	0.00%	125

Appendix C Experiment Details

It is worth noting that our approach is not time-consuming like deep learning models. The time required, based on the hardware shown in Table A4, ranges from 5 seconds to a maximum of 5 minutes.

Table A4. The hardware and software details of experiments.

	Hardware
CPU	12th Gen Intel(R) Core(TM) i9-12950HX 2.30 GHz
GPU	NVIDIA GeForce RTX 3080 Ti Laptop GPU
Memory	1TB
RAM	64.0 GB
OS	Windows 11 Pro

References

1. Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; Galstyan, A. A survey on bias and fairness in machine learning. *ACM computing surveys (CSUR)* **2021**, *54*, 1–35.

2. Cirillo, D.; Catuara-Solarz, S.; Morey, C.; Guney, E.; Subirats, L.; Mellino, S.; Gigante, A.; Valencia, A.; Rementeria, M.J.; Chadha, A.S.; others. Sex and gender differences and biases in artificial intelligence for biomedicine and healthcare. *NPJ digital medicine* **2020**, *3*, 81.
3. Celi, L.A.; Cellini, J.; Charpignon, M.L.; Dee, E.C.; Dernoncourt, F.; Eber, R.; Mitchell, W.G.; Moukheiber, L.; Schirmer, J.; Situ, J.; others. Sources of bias in artificial intelligence that perpetuate healthcare disparities—A global review. *PLOS Digital Health* **2022**, *1*, e0000022.
4. Byrne, M.D. Reducing bias in healthcare artificial intelligence. *Journal of PeriAnesthesia Nursing* **2021**, *36*, 313–316.
5. Zhang, Y.; Zhou, L. Fairness assessment for artificial intelligence in financial industry. *arXiv preprint arXiv:1912.07211* **2019**.
6. Fenu, G.; Galici, R.; Marras, M. Experts' view on challenges and needs for fairness in artificial intelligence for education. International Conference on Artificial Intelligence in Education. Springer, 2022, pp. 243–255.
7. Stiglitz, J.E. *The price of inequality: How today's divided society endangers our future*; WW Norton & Company, 2012.
8. UNSDG. Universal values principle two: leave no one behind, 2022.
9. Leslie, D.; Mazumder, A.; Peppin, A.; Wolters, M.K.; Hagerty, A. Does “AI” stand for augmenting inequality in the era of covid-19 healthcare? *bmj* **2021**, 372.
10. Morley, J.; Machado, C.C.; Burr, C.; Cows, J.; Joshi, I.; Taddeo, M.; Floridi, L. The ethics of AI in health care: a mapping review. *Social Science & Medicine* **2020**, *260*, 113172.
11. Ferrara, E. Fairness and bias in artificial intelligence: A brief survey of sources, impacts, and mitigation strategies. *Sci* **2023**, *6*, 3.
12. Data, I.; Team, A. Shedding light on AI bias with real world examples. *Security Intelligence* **2023**.
13. Satz, D.; White, S. What is wrong with inequality. *Inequality: The IFS Deaton Review. The IFS* **2021**.
14. Wu, H.; Wang, M.; Sylolypavan, A.; Wild, S. Quantifying health inequalities induced by data and AI models. *arXiv preprint arXiv:2205.01066* **2022**.
15. Morin, S.; Legault, R.; Bakk, Z.; Giguère, C.É.; de la Sablonnière, R.; Lacourse, É. StepMix: A Python Package for Pseudo-Likelihood Estimation of Generalized Mixture Models with External Variables. *arXiv preprint arXiv:2304.03853* **2023**.
16. Collins, L.M.; Lanza, S.T. *Latent class and latent transition analysis: With applications in the social, behavioral, and health sciences*; Vol. 718, John Wiley & Sons, 2009.
17. Trocin, C.; Mikalef, P.; Papamitsiou, Z.; Conboy, K. Responsible AI for digital health: a synthesis and a research agenda. *Information Systems Frontiers* **2023**, *25*, 2139–2157.
18. Garattini, C.; Raffle, J.; Aisyah, D.N.; Sartain, F.; Kozlakidis, Z. Big data analytics, infectious diseases and associated ethical impacts. *Philosophy & technology* **2019**, *32*, 69–85.
19. So, W.; Lohia, P.; Pimplikar, R.; Hosoi, A.; D'Ignazio, C. Beyond Fairness: Reparative Algorithms to Address Historical Injustices of Housing Discrimination in the US. Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, 2022, pp. 988–1004.
20. Hagenaars, J.A.; McCutcheon, A.L. *Applied latent class analysis*; Cambridge University Press, 2002.
21. Muthén, B.; Muthén, L.K. Integrating person-centered and variable-centered analyses: Growth mixture modeling with latent trajectory classes. *Alcoholism: Clinical and experimental research* **2000**, *24*, 882–891.
22. Vermunt, J.K.; Magidson, J. Latent class analysis. *The sage encyclopedia of social sciences research methods* **2004**, *2*, 549–553.
23. Lazarsfeld, P.F. The logical and mathematical foundation of latent structure analysis. *Studies in social psychology in world war II Vol. IV: Measurement and prediction* **1950**, pp. 362–412.
24. Goodman, L.A. The analysis of systems of qualitative variables when some of the variables are unobservable. Part IA modified latent structure approach. *American Journal of Sociology* **1974**, *79*, 1179–1259.
25. Weller, B.E.; Bowen, N.K.; Faubert, S.J. Latent class analysis: a guide to best practice. *Journal of Black Psychology* **2020**, *46*, 287–311.
26. Martínez-Monteagudo, M.C.; Delgado, B.; Inglés, C.J.; Escortell, R. Cyberbullying and social anxiety: a latent class analysis among Spanish adolescents. *International journal of environmental research and public health* **2020**, *17*, 406.
27. Bardazzi, R.; Charlier, D.; Legendre, B.; Pazienza, M.G. Energy vulnerability in Mediterranean countries: A latent class analysis approach. *Energy Economics* **2023**, *126*, 106883.

28. Sinha, P.; Furfaro, D.; Cummings, M.J.; Abrams, D.; Delucchi, K.; Maddali, M.V.; He, J.; Thompson, A.; Murn, M.; Fountain, J.; others. Latent class analysis reveals COVID-19-related acute respiratory distress syndrome subgroups with differential responses to corticosteroids. *American journal of respiratory and critical care medicine* **2021**, *204*, 1274–1285.
29. J. Sarpong. Bame we're not the same: Chinese, 2024. URL <https://www.bbc.com/creativitydiversity/nuance-in-bame/chinese>.
30. Drummond, R.; Pratt, M. Updating ethnic and religious contrasts in deaths involving the coronavirus (covid-19), England: 24 January 2020 to 23 November 2022, 2023.
31. ONS. Census 2021 data and analysis from census 2021, 2022. URL <https://www.ons.gov.uk/>.
32. GOV.UK. National statistics english indices of deprivation 2019, 2019. URL <https://www.gov.uk/government/statistics/english-indices-of-deprivation-2019>.
33. Lemaître, G.; Nogueira, F.; Aridas, C.K. Imbalanced-learn: A Python Toolbox to Tackle the Curse of Imbalanced Datasets in Machine Learning. *Journal of Machine Learning Research* **2017**, *18*, 1–5.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.