Analyzing 'Near Me' Services: Potential for Exposure Bias in Location-based Retrieval

Ashmi Banerjee
Technical University of Munich

Gourab K Patro
IIT Kharagpur

Linus W. Dietz
Technical University of Munich

Abhijnan Chakraborty
MPI for Software Systems

Abstract—The proliferation of smartphones has led to the increased popularity of location-based search and recommendation systems. Online platforms like Google and Yelp allow location-based search in the form of a *nearby* feature to query for hotels or restaurants in the vicinity. Moreover, hotel booking platforms like Booking[dot]com, Expedia, or Trivago allow travelers searching for accommodations using either their desired location as a search query or near a particular landmark. Since the popularity of different locations in a city varies, certain locations may get more queries than other locations. Thus, the exposure received by different establishments at these locations may be very different from their intrinsic quality as captured in their ratings.

Today, many small businesses (shops, hotels, or restaurants) rely on such online platforms for attracting customers. Thus, receiving less exposure than what is expected can be unfavorable for businesses. It could have a negative impact on their revenue and potentially lead to economic starvation or even shutdown. By gathering and analyzing data from three popular platforms, we observe that many top-rated hotels and restaurants get less exposure vis-a-vis their quality, which could be detrimental for them. Following a meritocratic notion, we define and quantify such exposure disparity due to location-based searches on these platforms. We attribute this exposure disparity mainly to two kinds of biases — *Popularity Bias* and *Position Bias*. Our experimental evaluation on multiple datasets reveals that although the platforms are doing well in delivering distance-based results, exposure disparity exists for individual businesses and needs to be reduced for business sustainability.

Index Terms—Location-based Search, Local Recommendation, Near Me, Meritocratic Fairness, Popularity Bias, Position Bias

I. INTRODUCTION

Location-based search and recommendation have become very popular nowadays, where a user can find businesses (e.g., hotels, restaurants, bars, cafes, or shops) based on their physical location. The query is often made using the user's current location or some other location explicitly entered by the user. The proliferation of smartphones with GPS availability has made these services very convenient and fast. Consequently, 30% of all mobile searches today are location-based, and every month people visit around 1.5 billion destinations using location-based results[1]. Since 2011, *"near me"* or location-based searches have increased by an astounding $3,400\%$[1]. In the case of local recommendations, the search query is implicit and the recommendations are based on the user's inferred location.

Multiple online platforms and mobile applications like Google, Yelp, etc. allow location-based searches and recommendations in the form of *nearby* feature to query restaurants, shops, hotels, etc. nearby. Moreover, hotel booking platforms like Booking[dot]com or Expedia allow travelers looking for hotels to explicitly specify a particular area of a city, and as a search result, hotels near that particular location are shown. Our analysis over a publicly available dataset by a hotel booking platform Trivago[2] reveals that 26% of users use a particular point of interest in a city to search for hotels at their destination, while the remaining 74% of users use generic city names as queries. Many other online platforms like TripAdvisor, Zomato, Airbnb, etc.[3] also feature location-based searches and recommendations.

With many people relying on such online platforms to choose places for their offline experiences[1] [1], these places, in order to be successful, need to get proper exposure for getting the attention of the users. Specifically, they need to appear in the top section of search results to gain attention, as users are susceptible to *position bias*, which makes them pay the most attention to the top results only [2], [3]. Appearing in lower ranks of search results usually translates to lower exposure, thereby causing loss of business opportunity for these places. Moreover, lack of proximity to popular locations, which is termed as *Popularity Bias* [4], [5], can also hamper the exposure different establishments get. This is especially undesirable if the received exposure does not corroborate with the intrinsic quality of these places.

In this work, we attempt to investigate whether the above concern actually gets manifested in real-world online platforms. By gathering and analyzing data from three popular platforms — Google, Yelp, and Booking[dot]com, we observe that many of the top-rated establishments indeed receive lower exposure compared to their lower-rated counterparts, due to their (lack of) proximity to popular locations. In the offline scenario, it is well-known that location impacts the success of different businesses, and based on footfalls estimated by market research agencies, businesses often try to offer superior services and other incentives to overcome the location factor [6]. Moreover, different users have different tolerances to distance — some users may prefer to go to a farther place to enjoy superior quality, whereas, for some other users, the

---

[1]https://www.thinkwithgoogle.com/consumer-insights/mobile-search-trends-consumers-to-stores

[2]https://recsys.trivago.cloud/challenge/dataset
[3]www.tripadvisor.com, www.zomato.com, www.airbnb.com

distance may be a strict constraint. Present location-based search and recommendations do not consider these nuances, and ranking different places solely in the ascending order of their distances can create the unintended exposure bias.

Our work doesn't question the idea of providing nearby establishments as results for location-based searches, rather it questions the ranking of those nearby establishments in the sole order of distance. This means that, even though the location is an important factor to determine the prosperity of an establishment, it should not be the only factor affecting its economic success. Establishments should ideally have the opportunity to compensate for a bad location with better quality, and the platforms should factor this into the ranking, thereby providing the opportunity to achieve more exposure. In future work, we plan to devise algorithms to possibly reduce the exposure bias by re-ranking search and recommendation output to effectively balance the trade-off between fairness and user satisfaction.

II. RELATED WORK

In recent times, virtually all types of online markets have observed a colossal increase in the number of users [7]. While there are concerns about these platforms behaving explicitly or implicitly biased manner, they undoubtedly exert a great deal of control over how their users behave and which decisions they make. With the possibility of the existence of bias, discrimination, and fairness concerns on these platforms, researchers have started investigating such online markets [8]–[12]. We study the literature along three dimensions: first, we talk about the existing biases on the online platforms; followed by potential mechanisms of fairness and generating the top-K best recommendations following a notion of fairness.

A. Bias in the Web

The sheer abundance of data have lead to the compromise of its quality while extracting it from its source. A severe quality issue is data bias, which has been studied by several researchers. The existence of biases affects machine learning algorithms designed to enhance the user experience, eventually exacerbating the outcome of these algorithms, particularly in the context of recommendations and personalization systems [13]. There is extensive prior work that has explored various forms of biases in social media [14]–[16], polarization in the society [17], and in linguistics [18]. Evidence that Google's search algorithm along with its auto-complete function systematically advocates information which is either false or inclined to an extreme right wing bias, on a diverse range of subjects, has also been reported [19]. Furthermore, Nunez' article [20] indicates instances of Facebook suppressing new stories of interest to conservative readers from the social network's influential "trending", eventually giving rise to biases. However, most literature focuses on the existence of biases in social media platforms. Our contribution lies in identifying the exposure bias existing on the sharing economy platforms. A high correlation between online exposure of the results generated on these sharing economy platforms and

their offline popularity reinforces the fact that it is important to diagnose and mitigate such biases in exposure to ensure conformity between their online received exposure and their actual quality.

B. Fairness

Data bias is inadvertently introduced into data-driven decision-making systems, as these systems operate by learning from historical decisions, often taken by humans [21], [22]. The goal of such systems should be to achieve maximum utility by minimizing the errors or misclassifications over given historical data. However, in such data, bias is already manifested and it could be that an optimally trained classifier does not make an impartial decision for all social groups and thereby introducing a notion of unfairness in decision making. Zafar et al. [22] address this unfairness in biased data in terms of misclassification rates, as disparate mistreatment and introduces intuitive measures of this disparate mistreatment for decision boundary-based classifiers with a slight compromise on accuracy. Another possible approach for the aforementioned problem could be by quantifying the measures of algorithmic unfairness [21].

Yao and Huang illustrate a notion of fairness in the context of collaborative recommender systems [23]. Experiments using their metrics, addressing different forms of unfairness, on various real data further establishes their claim of their method performing better than the baseline approaches. On the other hand, recent works [10]–[12], [24]–[26] bring to light the fact that fairness may be a multi-sided concept, in the context of recommendations, where the outcome should be fair for multiple stakeholders. Depending on these considerations, several fairness-aware recommender systems as well as suggestions for possible fairness-aware recommendation architectures have been proposed [10]–[12], [24], [26].

The study by Dwork et al. introduces the notions of individual and group fairness [27]. This is analogous to our approach, where we claim that items with similar quality metrics should receive similar attention on the sharing economy platform. Even though Chakraborty et al. [28] also emphasize on a novel framework regarding fairness in the matching mechanisms of online sharing economy platforms from the providers' perspective, the novelty of our work lies in the fact that we analyze the results on real-world platforms taking into consideration the inherent position bias of results and the popularity of the location query.

C. Top-K Fair Recommendations

The prolific need for personalized recommendations have led to the emergence of recommender systems where a filtering strategy is used to rank items that might be of interest to the user [29]–[31]. Thus, the result of the recommendation is not a single item, but a ranked list of length K, where the user still has the opportunity to choose the final item to be consumed. In multi-stakeholder platforms, where there are several parties involved, several notions of fairness need to be addressed by a recommendation algorithm [24], [26]. To achieve fairness,

Chakraborty et al. [30] envisioned top-K non-personalized crowd-sourced recommendations, such as the outcome of a multi-winner election which is periodically repeated, while selecting items at a particular time from popular websites, like Twitter, Yelp, TripAdvisor, or NYTimes [30]. The Fair Top-K Ranking problem of selecting a subset of K items from a large pool of N candidate items by maximizing the utility (i.e., select the best" candidates) can also be performed by subjecting to group fairness criteria such that it is capable of mitigating biases in the representation of an under-represented group along a ranked list [31].

III. DATASETS

The experimental evaluation has been carried out on three datasets obtained by querying for businesses — hotels and restaurants — from popular locations on Booking[dot]com, Yelp, and Google Places. This situation is analogous to generating a simulation of searching for the nearby establishments from the current location of a particular user. To demonstrate our analyses, we chose New York City (NYC) and San Francisco (SF), as these cities are both very large, and sufficiently different in terms of their neighborhood structure. We look into three platforms — Yelp and Google Places for the restaurant searches, and Booking[dot]com for accommodation searches. The rankings were constructed using the order in which each result appeared in the respective searches and we use the mean ratings as a measure of the quality of each establishment.

A. Location Queries

We used publicly available Foursquare check-in information of respective 37,470 (SF) and 572,338 (NYC) total check-ins at different venues, comprising of the geographic coordinates and timestamps [32], [33]. We used these check-ins as proxies for user interest locations, and use the check-in counts as an approximate measure of the popularity of places in a city where most location-based searches or nearby searches are made.

We clustered the Foursquare check-ins using K-Means clustering [34] and with a value of K=1000, thus, having the mean cluster radii for NYC to be 499 meters and 199 meters for SF. The cluster centroids obtained for both cities and a list of popular landmarks were used as location queries for generating results on the said platforms. We used 561 landmarks in New York[4], and 272 landmarks in San Francisco [5]. The location queries dataset for both cities contained the location of each query, followed by the name of the place and its address as on Google Maps. Here, the term location query and search query is used interchangeably within the purview of the work.

We assigned the check-ins to their nearest location query to calculate the popularity of each location query. Haversine distance [35] was used to compute the distance between the check-ins and their nearest location query. The popularity of a location query was calculated as the fraction of check-ins

mapped to the location query to the total number of unique check-ins in the city. We use this popularity score later as the weighing measure in our scoring model to calculate the exposure scores assigned to each establishment. Figure 1 shows the overall distribution of the establishments with respect to the check-ins and location queries on the different platforms for NYC and SF. It is evident from Figure 1 that our data is distributed all over the two cities.

B. Yelp

Yelp is a platform powered by a crowd-sourced review forum which lists local places. We used the official Yelp Fusion API[6] to query the different locations and collect the ranked results for restaurants. A total of 5,821 restaurants in NYC and 9,929 restaurants in SF were obtained in this way.

C. Google Places

Using the aforementioned location queries, we also gathered the nearby search results on the Google Places API. The first 20 results spanning over a radius of 8km were considered for every location query. A total of 31,220 restaurants in NYC and 56,541 restaurants in SF could be extracted in this way. This covers 1,761 and 4,252 unique restaurants in NYC and SF, respectively.

D. Booking[dot]com

Booking[dot]com[7] is a meta-search engine for making lodging reservations all across the world. We used the results displayed on the first 10 pages of Booking[dot]com for the location queries. Each page displayed a minimum of 15 search results. This way, we obtained 1,547 unique lodging establishments for NYC and 4,334 for SF between 25th December 2018 and 31st December 2018.

E. Dataset Statistics

In this section, we illustrate the relevant statistics of the datasets that were gathered. Table I summarises the information gathered in the previous sections for a better overview of the datasets.

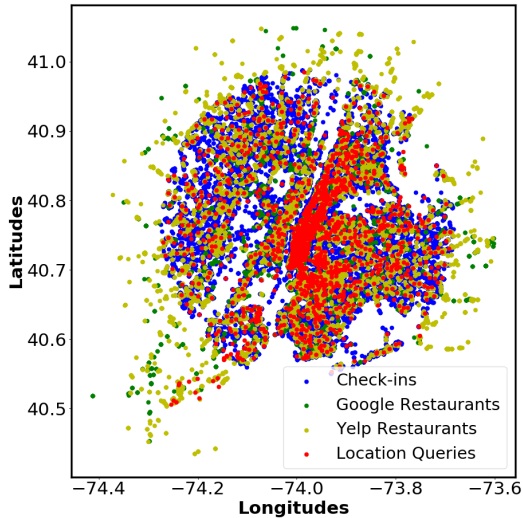|	Yelp		Google		Booking[dot]com	
	NYC	SF	NYC	SF	NYC	SF
# Search Queries	1560	899	1560	899	1560	899
#Unique Establishments	4554	3587	827	5702	1547	4334
Minimum Rating	1.50	1.00	1.00	2.30	4.10	4.30
Mean Rating	4.20	4.19	3.91	3.87	8.35	7.76
Maximum Rating	5.00	5.00	5.00	5.00	10.00	10.00

TABLE I: **Basic statistics of the gathered datasets for the three different platforms for each of the two cities — NYC and SF.**
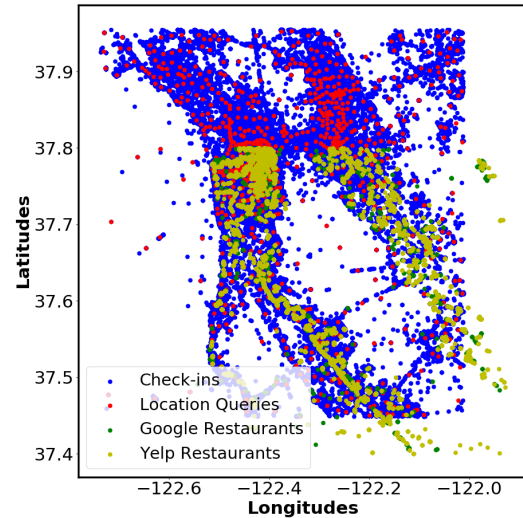
---

[4]NYC Open data: https://opendata.cityofnewyork.us
[5]Wikipedia page: https://en.wikipedia.org/wiki/List_of_San_Francisco_Designated_Landmarks

[6]https://www.yelp.com/developers/documentation/v3
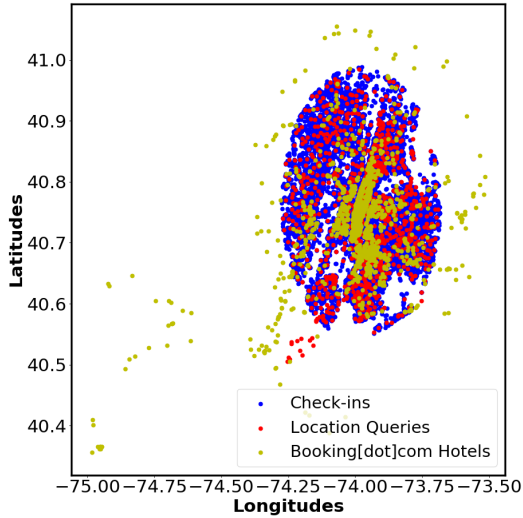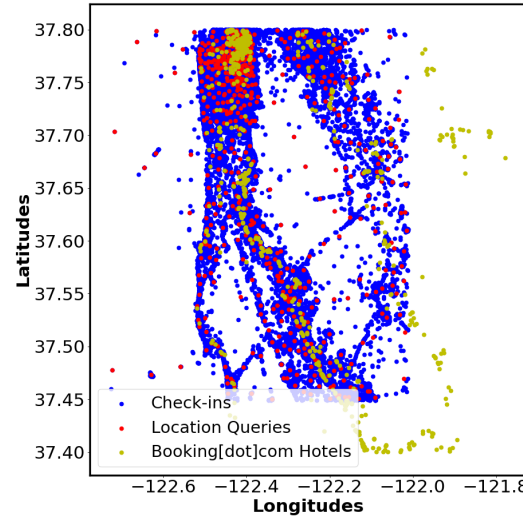[7]https://www.booking.com

(a) NYC: Google & Yelp

(b) SF: Google & Yelp

(c) NYC: Booking[dot]com

(d) SF: Booking[dot]com

Fig. 1: **Figure (a) and (b) show the distribution of check-ins, location queries, and establishments for NYC and SF respectively on Google and Yelp. Figure (c) and (d) show a similar distribution for Booking[dot]com. Each point is a coordinate represented by longitude and latitude showing the locations of check-ins, location queries, and establishments on the platforms. It can be seen that our location queries, which represent the cluster centroids, are spread across the two cities. We obtained average cluster radii of 499 meters for NYC and 199 meters for SF.**

F. Exploratory Analyses

Several exploratory analyses have been performed on the three aforementioned datasets to observe the underlying trends in the data. Figure 2 shows the mean distance of each establishment from the query location versus their ranks when displayed in the location-based search for three different platforms — Yelp, Google, and Booking[dot]com. Here, rank denotes the position at which the establishment appeared in the search result. The mean of the various positions at which a particular establishment appeared during all the searches has been denoted as the rank, for each establishment.

From Figure 2, it is evident that the farther away a particular establishment is from the location query the lower it is ranked on average. Although, it is not certain whether a platform ranks its results purely based on distance, the strong correlation suggests that the distance from the location query could have been taken as a very important feature for the ranking.

Table II brings out the ranks versus average distance (in kilometers) for the top three ranked establishments for the three platforms. The possible reason behind having higher distances in SF than NYC could be potentially attributed to the higher population density of NYC [36] than SF.
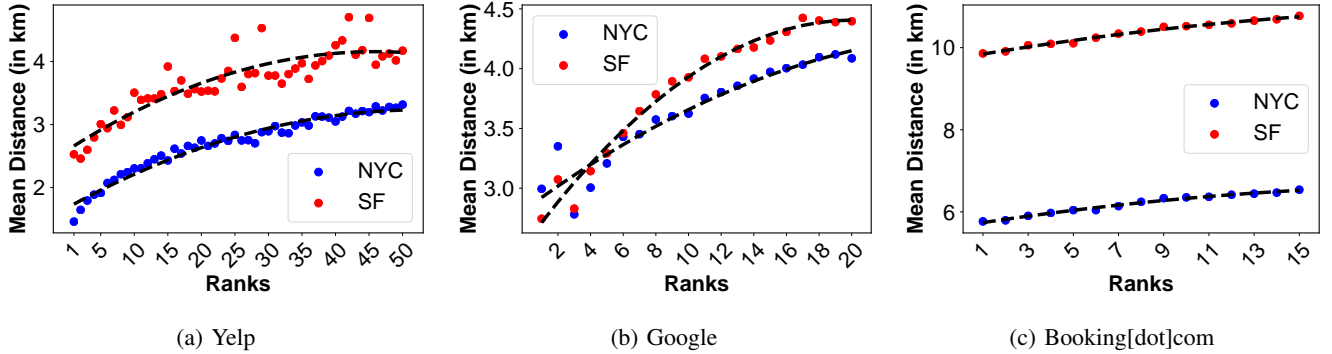
(a) Yelp   (b) Google   (c) Booking[dot]com

Fig. 2: **Ranks vs. average distance from the location queries for all three platforms. Rank here denotes the mean of the various positions at which a particular establishment appeared during all the searches. An increase in the mean distance from the location query has been observed with the increase in the ranks of the establishments. This strengthens the fact that the distance from the location query plays an important role in generating the results on these platforms.**

| Ranks | NYC | | | SF | | |
	Yelp	Google	Booking	Yelp	Google	Booking
1	1.45	2.99	5.67	2.52	2.74	9.85
2	1.64	3.35	5.79	2.45	3.07	9.90
3	1.79	2.78	5.90	2.59	2.82	10.05

TABLE II: **Comparison of ranks vs. mean distance (in km) from location queries on different platforms. The table shows that establishments with lower ranks tend to be farther away from the location queries.**

IV. EXPOSURE BIAS IN SEARCH AND RECOMMENDATION

Subsequently, we introduce the concepts and methods we propose to capture and analyze biases in these platforms. Even though, in this study, we focus on the local businesses like hotels and restaurants, our methods can be employed in any such sharing-economy platform that presents its recommendations in the form of ranked lists. To establish a generic terminology, we use the term *"businesses"* for establishments. Otherwise, the terms can be used interchangeably.

A location-based search query necessarily comprises of the location it is made from, but it can have various other parameters specifying e.g., types of places such as restaurants, bars, monuments, or specific types of restaurants pertaining to specific cuisine such as Italian or Asian restaurants as filtering criteria. These filtering criteria are usually binary, thus, independent of the ranking mechanism of the platform. Hence, we ignore the filtering criteria in our analysis and consider only the base query locations as the search queries.

In the following, we define relevance and exposure in location-based searches. We assume a city to have a limited number of landmarks based on which users perform nearby searches and that the landmarks have different popularities among the users.

A. Relevance of Businesses

The relevance of a business represents how *"worthy"* it is to be included as result to a user query. We can formalize relevance as a function from the set of locations L, and set of businesses S to the set of non-negative real numbers.

$$V : (L, S) \to \mathbb{R}_{\geq 0} \qquad (1)$$

$V(l_i, s_j)$ is the normalized relevance of business s_j when location l_i is searched.

The output of the relevance function can be considered as an indicator for the *desired exposure* of a business in a specific instance of location-based retrieval. Typically platforms, however, devise their own function for relevance according to their goals and applications. For our study, we use the mean rating of a business as the relevance function. However, more complex functions can also be used to capture more complex requirements.

B. Exposure from Search and Recommendation

The exposure of a business is related to the amount of attention it receives from the users, which depends explicitly on the position it appears in search results, for any ranked retrieval system. Let $\omega(t)$ be the attention a business gets when it appears in rank t, then the total exposure gained by a business s_j over k searches can be written as below.

$$E_j = \sum_{i=1}^{k} \omega(\rho_i(s_j)) \qquad (2)$$

Here, $\rho_i(s_j)$ denotes the rank of business s_j in result (permutation) ρ_i. Its value can be any integer from 1 to n. The computation of $\omega(t)$ has been elaborated in subsection V-B.

C. Fair Exposure for Businesses

Our notion of fair exposure in this work is analogous to the ideas of *individual fairness* and *distributive fairness*. We follow *individual fairness* proposed by Dwork et al. [27] stating that *"similar individuals should be treated similarly."* The similarity between individuals needs to be established based on similar values of a suitable metric specific to the task at hand. Here, we consider businesses having similar relevance scores as similar businesses and we proclaim that

two businesses are treated similarly if they receive similar exposure scores. We consider *exposure given by an online platform in total* as a readily available, yet limited resource. How this exposure is distributed is up to the discretion of the platform's recommendation algorithm. Potential distributive norms could be equity, equality, power, or need.

Following works on meritocracy [37], in this work, we use *merit* as the distributive norm which requires the distribution of resources in proportion to the relevance of the businesses. This idea is also inspired by the concept of *"distributive fairness"* proposed by Rawls [38], [39]. In this regard, we use the term *"meritocratic fairness"* for businesses to keep their exposures in proportion to the corresponding values of the relevance function.

An online platform's search or recommendation results $\rho_1, \rho_2, \ldots, \rho_k$ for respective base locations $l'_1, l'_2, \ldots, l'_k$ will be said to follow the notion of *meritocratic fairness* if the following condition is satisfied.

$$\frac{E_i}{V_i} = \frac{E_j}{V_j}, \forall s_i, s_j \in S \tag{3}$$

where $V_j = \sum_{t=1}^{k} V(l'_t, s_j)$ is the sum of all the normalized relevance scores of business s_j over k customer instances and $\rho_1, \rho_2, ..., \rho_k$ denotes the results of location-based searches on locations $l'_1, l'_2, ..., l'_k$ correspondingly as provided by the platform. The equality constraint defined above is a hard constraint, but we can also relax it by replacing it with soft constraints using any kind of tolerance formulation.

D. Exposure Bias on Real-World Platforms

In location-based searches, establishments located in the proximity of the location query are usually observed to be ranked higher and establishments located farther are often ranked lower irrespective of their intrinsic quality. Here, we consider the ratings of the establishments as the measure of their intrinsic quality. As different parts of a city have very different popularity among customers, establishments located closer to the popular parts of the city get to be in the top results for a number of searches. It has been often observed that even poor quality establishments located in proximity to popular parts of the city get much more exposure than many superior quality ones located slightly farther. Such discrepancy between the *actual exposure* and the *deserved exposure* can be interpreted as *Exposure Bias*. We define deserved exposure as the intrinsic quality of the establishment which might be specific to different platforms or different applications, and might also depend on various factors including the ratings of place, but also the distance from the base location. The popularity of different parts of the city also varies over time, thereby adding another level of complexity in conceptualizing the *exposure bias over time*. However, currently we do not consider the temporal dimensions for this work. Moreover, the usual practice in the case of location-based searches is to sort the establishments, which are in the same vicinity, based on the distance and the difference in the popularity of different location queries. This can cause high levels of exposure bias

for some establishments and can potentially result in economic starvation or shutdown of the establishments which is something a recommendation platform should actively counteract. Unfortunately, little research has been done on quantifying this kind of exposure bias, despite the clear need to carefully analyze and mitigate it for the benefit of all stakeholders: users, businesses, and platforms.

V. CHARACTERIZING EXPOSURE BIAS

Essentially, the exposure bias for any business can be described as the difference between its relevance to the query and the user attention it receives. Hence, the exposure bias B_j for business s_j when query with base location l'_i gives ranking ρ_i, can be written as below.

$$B_j(\rho_i, l'_i) = V(l'_i, s_j) - \omega(\rho_i(s_j)) \tag{4}$$

The overall exposure bias can be computed in the following ways:

$$\text{Utilitarian Measure} \quad {}^1B = \sum_{j=1}^{n} |B_j|$$

$$\text{Egalitarian Measure} \quad {}^2B = \max_{j \epsilon [1,n]} |B_j|$$

In this paper, we focus on the utilitarian measure. As it is difficult to achieve zero exposure bias in each search ranking, it is more useful to quantify the bias in cumulative form over a number of searches. We utilize the idea of *amortized fairness* as proposed by Biega et al. [40] to define cumulative exposure bias here. The exposure bias introduced by an online platform's k search results $\rho_1, \rho_2, ..., \rho_k$ for respective user locations $l'_1, l'_2, ..., l'_k$, can be seen as below.

$$B = \sum_{j=1}^{n} |V_j - E_j| = \sum_{j=1}^{n} \left| \sum_{i=1}^{k} \{V(l'_i, s_j) - \omega(\rho_i(s_j))\} \right| \tag{5}$$

Results of location-based search on a platform are said to follow *meritocratic fairness* if exposure bias approaches zero.

We consider position bias and popularity bias to be responsible for above exposure bias. We define them in the following sections, as two components of this exposure bias. Even though position bias has been studied earlier [27], [40], [41], fewer focus has been given to popularity bias [42], [43].

A. Popularity Bias

By gathering and analyzing data from three popular platforms, namely Google, Yelp, and Booking[dot]com, we observe that many of the top-rated places indeed receive lower exposure compared to their lower-rated counterparts due to their lack of proximity to popular locations. Such popularity bias has been previously studied under the context of recommender systems and social media [5], [43]. In our analyses, we seek to provide evidence for the existence of unintended exposure bias due to the popularity of location queries.

The frequencies of location queries are different in different parts of the city, with more popular and known districts and landmarks being queried more often (Figure 3a). Thus, establishments situated near the popular locations of the city, tend to be in the top results more often. It has been observed many
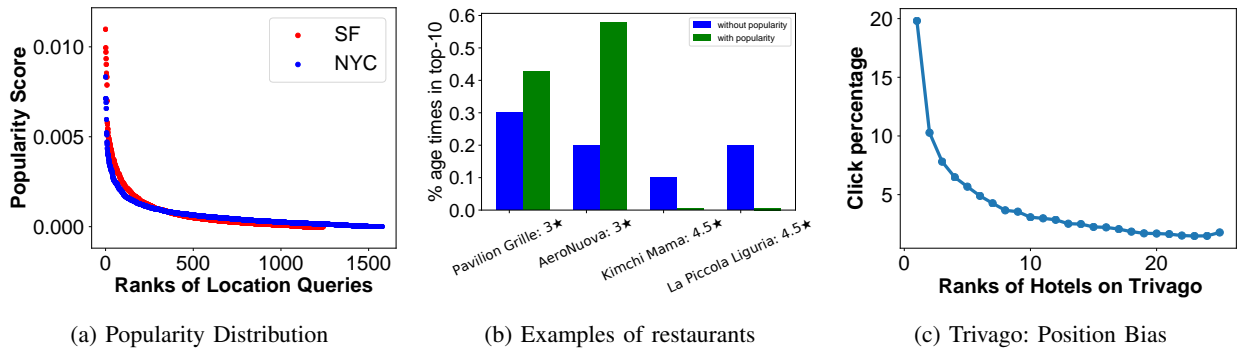
(a) Popularity Distribution   (b) Examples of restaurants   (c) Trivago: Position Bias

Fig. 3: **Demonstration of Exposure Bias: (a) Popularity Distribution of different location queries across the city. It can be observed that certain areas in the city have much higher popularity than its counterparts. Therefore establishments situated closer to these popular locations tend to appear more frequently in the top results even if they do not have very high ratings. This leads to Popularity Bias where establishments located farther away from popular locations of the city are deprived of their deserved exposure. (b) Examples of some top-rated restaurants from Yelp with good exposure without popularity (i.e., considering the same popularity for all the location queries) versus the same when the popularity of location queries have been considered. It is evident that top-rated restaurants, farther from the popular locations end up receiving lesser exposure compared to their lower-rated counterparts. (c) Percentage of clicks on the ranked lists of results following a geometric distribution for Trivago data, bolstering the existence of Position Bias on these platforms.**

times that even lower-rated establishments or establishments of poor quality, when located in the proximity of popular regions, receive much more exposure than their higher-rated counterparts, which are located slightly farther. Figure 3b shows examples of how the popularity of location queries has caused top-rated restaurants—farther from the popular locations—to end up receiving lesser exposure compared to their lower-rated equivalents.

B. Position Bias

Position Bias or the bias induced due to the ranking of the businesses have been studied in earlier works [40], [41]. Owing to position bias, disproportionately less attention is being paid to businesses that are ranked lower in the search results obtained in a ranked retrieval system. Our work is inspired by the work by Biega et al. [40], which focuses on the relationship between relevance and attention. Since relevance can be imagined as a proxy for worthiness in the context of a given search task, the attention received by a business from searchers should ideally be proportional to their relevance. Studies have shown that establishments need to appear in the top results of a search query in order to gain attention, as users are susceptible to position bias, which makes them pay the most attention to the top results [**?**], [2]. Appearing in lower ranks of search results usually translates to lower exposure, thereby causing loss of business opportunity for these places. It is especially undesirable if the received exposure does not corroborate with the intrinsic quality of these places, which we approximate using their user ratings.

Our proposed models require a weight assigned to each position of the ranked list, denoting the fraction of the total attention the position gets. The higher the position of the establishment is on the list, the higher weight will be assigned. To

calculate exposure scores, we require the fraction of attention each position in the search result gets, to be the normalized amount of users' attention to the position. These values can be calculated from the previous user behavior.

In this regard, we used a publicly accessible dataset provided by Trivago (made available during RecSys 2019 Challenge[8]) to study the user behavior on the Trivago hotel search platform. It was observed from Figure 3c that the percentage of clicks on the ranked lists of results geometrically decreased with the increase in the position at which they appeared as a result of the search. Therefore, for our analysis, we have considered the following Geometric-κ attention distribution for assigning the exposure scores to the establishments.

Geometric-κ Attention: The weights are distributed geometrically with parameter p starting from position 1 to κ. It remains 0 for positions beyond κ. The distribution of attention is given by:

$$\omega_t = \begin{cases} \frac{p(1-p)^{t-1}}{\sum_{j=1}^{\kappa} p(1-p)^{j-1}} & 1 \leq t \leq \kappa \\ 0 & t > \kappa \end{cases}$$

The value of κ has been considered to be the number of results returned on the first page of the search query on each platform. Therefore, we have considered a value of κ equals 50, 20 and 15 for Yelp, Google, and Booking[dot]com respectively. The value p equals 0.144 has been computed from Figure 3c as the best fitting value for the curve. By adding up the individual attention scores received over different user instances, we compute the total exposure for every business on a platform (as given in Equation 2), and then, the total

---

[8]https://recsys.trivago.cloud

exposure bias for each business and for the particular platform (using Equation 4 and Equation 5).

C. Exposure Scores vs. Ratings

Each of the establishments obtained from the various platforms — Yelp, Google, and Booking[dot]com were scored against the Geometric-κ attention mechanism scheme elaborated in section V-B, based on their appearance in the search result. The mean exposure score for every establishment was then computed under this scheme.

To assess the impact of the popularity of the search query on the exposure scores, first, we assigned the scores to the establishments depending on their appearance in the search result. Then they were weighed against their popularity score. The mean of these weighted exposure scores was calculated for every establishment. Figure 4 shows the variation of exposure scores with respect to the establishment ratings for the Geometric Attention distribution scheme for the different platforms. It is evident from Figure 4 that certain establishments despite having a high rating (4.0 or above for Yelp and Google and greater than 6.0 for Booking[dot]com) receive lower exposure than their lower-rated counterparts and vice-versa, thereby proving the existence of exposure bias on these platforms. However, the range of exposure values across different ratings are similar indicating that it doesn't follow a meritocratic notion.

D. Observed Exposure Bias

Exposure bias seems to exist when low rated establishments have higher exposure than their higher-rated counterparts. In other words, when the percentage of establishments with high ratings receiving low exposure exceeds the ones with high ratings and high exposure and vice versa. To calculate the existing exposure disparity with respect to the ratings we define the following terms —

- **High Exposure (HE)**: Exposure scores belonging to the top 25 percentile.
- **Low Exposure (LE)**: Exposure scores belonging to the bottom 25 percentile.
- **High Rating**: For Google and Yelp, since the rating distribution is discrete, any restaurant having a mean rating of 4.0 or higher has been considered to fall in the range of highly-rated restaurants. However, for Booking[dot]com, since the rating distribution was continuous, therefore the establishments with ratings in the top 25 percentile have been considered to be highly-rated.
- **Low Rating**: Similarly, for Google and Yelp, any restaurant having a mean rating lower than 4.0 has been considered to fall in the range of low rated restaurants. For Booking[dot]com, the establishments with ratings in the bottom 25 percentile have been considered to be low rated.

We divided the set of establishments on the three platforms into two groups — establishments with *low ratings* and *high ratings*. For each of these groups, we calculated the percentage of establishments receiving *high exposure* (**HE)** and *low*

		NYC		SF	
	Rating	LE	HE	LE	HE
Yelp	Low	29.05	70.94	35.42	64.57
	High	10.65	89.34	11.81	88.18
Google	Low	60.38	39.61	43.28	56.71
	High	51.69	48.30	39.03	60.96
Booking[dot]com*	Low	24.54	25.58	34.63	18.64
	High	16.00	32.00	25.00	28.12

TABLE III: **Percentage of establishments showing exposure disparity owing to Position Bias. HE denotes high exposure (top 25 percentile), LE denotes low exposure (bottom 25 percentile). Low and High denote low and high ratings respectively. The values in red indicate the observed exposure disparity. The exposure disparity values indicate that exposure bias does not follow a meritocratic notion.**

exposure (**LE**). The whole experiment was done under two scenarios —

- when all location queries were considered to have equal popularities (Position Bias)
- when the popularity of the location queries was considered (Popularity Bias)

The results of the fraction of establishments exhibiting exposure bias on the different platforms were summarized in Table III. The range of exposure values across different ratings — both high and low are similar, indicating that it does not follow a meritocratic notion.

VI. CONCLUSION

We theoretically formulate the idea of exposure bias due to location-based retrieval. In our experimental study using real-world data, we show that substantial discrepancy exists between the actual exposure received by the establishments and their deserved exposure, despite several establishments having equal relevance scores. We attribute this exposure disparity to mainly two kinds of biases — *Popularity Bias* and *Position Bias*. While Position Bias has been previously addressed by several studies, investigating Popularity Bias in the domain of multi-sided market platforms remains a relatively less explored problem. Our work makes the novel contribution of identifying the existence of unintended exposure bias due to *Popularity Bias* and *Position Bias* in location-based retrievals. To the best of our knowledge, we are the first to quantify such bias in the case of location-based retrieval.

Exposure preservation is necessary, so that establishments are not deprived from their deserved exposure both online and offline. There is evidence that online exposure can lead to offline economic opportunity for the establishments as many people rely on online platforms for their offline experiences [1]. Hence, this bolsters our notion of fairness, according to which all the establishments should receive online exposure in proportion to their quality or merit.

Our work is focused on the real-world data gathered from Yelp, Google and Booking[dot]com for two cities, namely
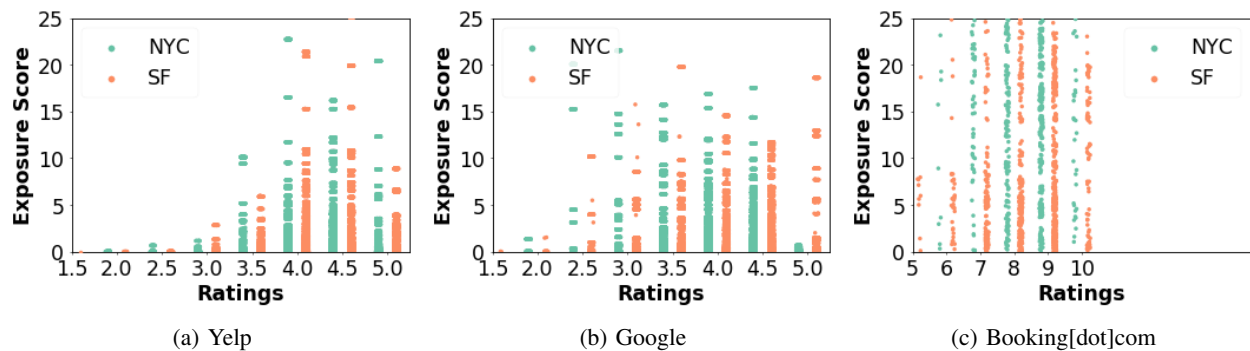
| (a) Yelp | (b) Google | (c) Booking[dot]com |

Fig. 4: **Ratings vs. exposure scores for all three platforms — Yelp, Google, and Booking[dot]com. It can be seen that there are several establishments on these platforms which have higher ratings but receive lower exposure and vice-versa. This further reinforces the existence of exposure bias on these platforms.**

New York City and San Francisco. However, the analysis can be extended to other cities in the world and on other platforms such as Airbnb or Trivago as well. Various demographic and economic influences of the neighborhoods in which these establishments are located and their impact on their exposure can also be explored in future work.

Finally, while we have discussed the potential causes of exposure bias and established the existence of considerable discrepancy in exposure among various establishments, methods to mitigate this still remain unexplored. In future we plan to design an online optimization framework to minimize exposure bias over time and mitigate the effects of popularity bias and position bias.

ACKNOWLEDGEMENT

REFERENCES

[1] H. Huang, G. Gartner, J. M. Krisp, M. Raubal, and N. V. de Weghe, "Location based services: ongoing evolution and research agenda," *Journal of Location Based Services*, vol. 12, no. 2, pp. 63–93, 2018.

[2] N. Craswell, O. Zoeter, M. Taylor, and B. Ramsey, "An experimental comparison of click position-bias models," in *ACM WSDM*, 2008, pp. 87–94.

[3] R. M. Ursu, "The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions," *Marketing Science*, vol. 37, no. 4, pp. 530–552, 2018.

[4] G. L. Ciampaglia, A. Nematzadeh, F. Menczer, and A. Flammini, "How algorithmic popularity bias hinders or promotes quality," *Scientific reports*, vol. 8, no. 1, p. 15951, 2018.

[5] H. Steck, "Item popularity and recommendation accuracy," in *ACM RecSys*, 2011, pp. 125–132.

[6] H. Erkuş-Öztürk and P. Terhorst, "Innovative restaurants in a mass-tourism city: Evidence from antalya," *Tourism management*, vol. 54, pp. 477–489, 2016.

[7] P. N. Danziger, "6 global consumer trends for 2019, and the brands that are out in front of them," Online, Jan. 2019, https://www.forbes.com/sites/pamdanziger/2019/01/13/6-global-consumer-trends-and-brands-that-are-out-in-front-of-them-in-2019/.

[8] M. Lin and S. Viswanathan, "Home bias in online investments: An empirical study of an online crowdfunding market," *Management Science*, vol. 62, no. 5, pp. 1393–1414, 2015.

[9] K. Levy and S. Barocas, "Designing against discrimination in online markets," *Berkeley Tech. LJ*, vol. 32, p. 1183, 2017.

[10] 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 *ACM KDD*, 2019, pp. 3082–3092.

[11] G. K. Patro, A. Biswas, N. Ganguly, K. P. Gummadi, and A. Chakraborty, "Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms," in *WWW*, 2020, pp. 1194–1204.

[12] G. K. Patro, A. Chakraborty, N. Ganguly, and K. Gummadi, "Incremental fairness in two-sided market platforms: On smoothly updating recommendations," in *AAAI*, vol. 34, no. 01, 2020, pp. 181–188.

[13] R. Baeza-Yates, "Data and algorithmic bias in the web," in *WebScience*, 2016.

[14] A. Chakraborty, S. Ghosh, N. Ganguly, and K. P. Gummadi, "Dissemination biases of social media channels: On the topical coverage of socially shared news," in *AAAI ICWSM*, 2016.

[15] A. Chakraborty, J. Messias, F. Benevenuto, S. Ghosh, N. Ganguly, and K. P. Gummadi, "Who makes trends? understanding demographic biases in crowdsourced recommendations," in *AAAI ICWSM*, 2017.

[16] H. Abdollahpouri, "Popularity bias in recommendation: A multi-stakeholder perspective," *arXiv preprint arXiv:2008.08551*, 2020.

[17] P. Dandekar, A. Goel, and D. T. Lee, "Biased assimilation, homophily, and the dynamics of polarization," *Proceedings of the National Academy of Sciences*, vol. 110, no. 15, pp. 5791–5796, 2013.

[18] J. Otterbacher, "Linguistic bias in collaboratively produced biographies: crowdsourcing social stereotypes?" in *AAAI ICWSM*, 2015.

[19] O. Solon and S. Levin, "How google's search algorithm spreads false information with a rightwing bias," 2016. [Online]. Available: https://bit.ly/3nAtepk

[20] M. Nunez, "Former facebook workers: We routinely suppressed conservative news," https://gizmodo.com/former-facebook-workers-we-routinely-suppressed-conser-1775461006, 2016.

[21] T. Speicher, H. Heidari, N. Grgic-Hlaca, K. P. Gummadi, A. Singla, A. Weller, and M. B. Zafar, "A unified approach to quantifying algorithmic unfairness: Measuring individual & group unfairness via inequality indices," in *ACM SIGKDD*, 2018.

[22] M. B. Zafar, I. Valera, M. Gomez Rodriguez, and K. P. Gummadi, "Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment," in *WWW*, 2017, pp. 1171–1180.

[23] S. Yao and B. Huang, "Beyond parity: Fairness objectives for collaborative filtering," in *NIPS*, 2017.

[24] R. Burke, "Multisided fairness for recommendation," *arXiv preprint arXiv:1707.00093*, 2017.

[25] G. K. Patro, A. Chakraborty, A. Banerjee, and N. Ganguly, "Towards safety and sustainability: Designing local recommendations for post-pandemic world," in *ACM RecSys*, 2020, pp. 358–367.

[26] A. S. Mondal, R. Bal, S. Sinha, and G. K. Patro, "Two-sided fairness in non-personalised recommendations," 2020.

[27] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, "Fairness through awareness," in *ACM ITCS*, 2012.

[28] A. Chakraborty, A. Hannak, A. J. Biega, and K. P. Gummadi, "Fair sharing for sharing economy platforms," 2017.

[29] M. Deshpande and G. Karypis, "Item-based top-n recommendation algorithms," *ACM TOIS*, vol. 22, no. 1, 2004.

[30] A. Chakraborty, G. K. Patro, N. Ganguly, K. P. Gummadi, and P. Loiseau, "Equality of voice: Towards fair representation in crowd-sourced top-k recommendations," 2019.

[31] M. Zehlike, F. Bonchi, C. Castillo, S. Hajian, M. Megahed, and R. A. Baeza-Yates, "Fa*ir: A fair top-k ranking algorithm," in *CIKM*, 2017.

[32] D. Yang, D. Zhang, V. W. Zheng, and Z. Yu, "Modeling user activity preference by leveraging user spatial temporal characteristics in lbsns," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 45, no. 1, pp. 129–142, 2014.

[33] D. Yang, B. Qu, J. Yang, and P. Cudre-Mauroux, "Revisiting user mobility and social relationships in lbsns: A hypergraph embedding approach," in *The World Wide Web Conference*, 2019, pp. 2147–2157.

[34] J. MacQueen *et al.*, "Some methods for classification and analysis of multivariate observations," in *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, vol. 1, no. 14, 1967, pp. 281–297.

[35] R. Sinnott, "Virtues of the haversine. sky and telescope," 1984.

[36] D. T. Cohen, G. W. Hatchard, S. G. Wilson *et al.*, *Population trends in incorporated places: 2003 to 2013*. US Department of Commerce, Economics and Statistics Administration, 2015.

[37] A. W. Cappelen, E. Ø. Sørensen, and B. Tungodden, "Responsibility for what? fairness and individual responsibility," *European Economic Review*, vol. 54, no. 3, pp. 429–441, 2010.

[38] J. Rawls, *A theory of justice*. Harvard university press, 2009.

[39] ——, *Justice as fairness: A restatement*. Harvard University Press, 2001.

[40] A. J. Biega, K. P. Gummadi, and G. Weikum, "Equity of attention: Amortizing individual fairness in rankings," in *ACM SIGIR*, 2018, pp. 405–414.

[41] A. Singh and T. Joachims, "Fairness of exposure in rankings," in *ACM KDD*, 2018, pp. 2219–2228.

[42] H. Abdollahpouri, R. Burke, and B. Mobasher, "Managing popularity bias in recommender systems with personalized re-ranking," *arXiv preprint arXiv:1901.07555*, 2019.

[43] A. Nematzadeh, G. L. Ciampaglia, F. Menczer, and A. Flammini, "How algorithmic popularity bias hinders or promotes quality," *arXiv preprint arXiv:1707.00574*, 2017.