

Interest Networks (iNETs) for Cities: Cross-Platform Insights and Urban Behavior Explanations

Gustavo H. Santos

Universidade Tecnológica Federal do Paraná
Curitiba, Brazil
gustavohenriquesantos@alunos.utfpr.edu.br

Myriam Delgado

Universidade Tecnológica Federal do Paraná
Curitiba, Brazil
myriamdelg@utfpr.edu.br

Thiago H. Silva

Universidade Tecnológica Federal do Paraná
Curitiba, Brazil
thiagoh@utfpr.edu.br

Abstract

Location-Based Social Networks (LBSNs) provide a rich foundation for modeling urban behavior through iNETs (Interest Networks), which capture how user interests are distributed throughout urban spaces. This study compares iNETs across platforms (Google Places and Foursquare) and spatial granularities, showing that coarser levels reveal more consistent cross-platform patterns, while finer granularities expose subtle, platform-specific behaviors. Our analysis finds that, in general, user interest is primarily shaped by geographic proximity and venue similarity, while socioeconomic and political contexts play a lesser role. Building on these insights, we develop a multi-level, explainable recommendation system that predicts high-interest urban regions for different user types. The model adapts to behavior profiles—such as *explorers*, who are driven by proximity, and *returners*, who prefer familiar venues—and provides natural-language explanations using explainable AI (XAI) techniques. To support our approach, we introduce *h3-cities*, a tool for multi-scale spatial analysis, and release a public demo for interactively exploring personalized urban recommendations. Our findings contribute to urban mobility research by providing scalable, context-aware, and interpretable recommendation systems.

CCS Concepts

• **Information systems** → **Information systems applications; Information retrieval**; • **Computing methodologies** → **Machine learning**; • **Applied computing** → *Law, social and behavioral sciences*.

Keywords

Location-Based Social Networks, Google Places, Foursquare, User Interest, Urban Areas, Recommendation

ACM Reference Format:

Gustavo H. Santos, Myriam Delgado, and Thiago H. Silva. 2025. Interest Networks (iNETs) for Cities: Cross-Platform Insights and Urban Behavior Explanations. In *Proceedings of 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '25)*. ACM, New York, NY, USA, 8 pages. <https://doi.org/XXXXXXX.XXXXXXX>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '25, Toronto, Canada

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-XXXX-X/2018/06
<https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Informative urban data obtained by Location-Based Social Networks (LBSNs) can reflect user preferences and mobility behaviors [2, 16, 30, 32, 33, 35, 38]. By capturing geolocated activities such as check-ins and reviews, these platforms enable the construction of Interest Networks (iNETs)—undirected graphs where nodes represent urban areas (e.g cities, neighborhoods, ZIP codes, etc...) and edges link areas co-visited by the same users. Our proposed iNETs allow exploring urban dynamics and can be used in applications such as recommendation systems, mobility prediction, and policy design. Our research comprises three phases.

Formation of iNETs: We construct Interest Networks (iNETs) from user activity data drawn from LBSNs. These iNETs are built by connecting spatial regions whenever the same users interact with places in multiple areas. To support more comprehensive and interpretable analyses, we enrich each region with a set of contextual attributes, including socioeconomic indicators, cultural profiles, political leanings, and geographic features. This enriched representation allows for deeper investigation into the factors shaping user mobility, regional interest patterns, and cross-platform consistency.

Validation across granularity levels and platforms: We systematically evaluate how iNETs perform across varying spatial granularities and urban configurations. Our results reveal that coarser levels of granularity lead to greater alignment between iNETs across platforms, while finer levels highlight subtle behavioral differences. To reconcile these discrepancies, we propose clustering high user interest - such clusters (namely, Urban Preference Zones - UPZones) appear more consistently across different platforms and spatial scales. The clusters offer a stable basis for comparison and act as interpretable reference points for further modeling.

Understanding user interests: We leverage the validated iNETs to model user interests at both aggregated and individual levels. At the aggregated level, we analyze iNET structures to assess general mobility patterns. Our findings indicate that geographic proximity and venue similarity modestly influence user movement between areas, while socioeconomic and political factors show limited impact. At the individual level, we distinguish between high- and low-interest regions based on user behavior. In addition to general mobility traits, we examine returner vs. explorer behavioral patterns [23], and introduce a personalized multi-level recommendation model grounded in users' previously visited high-interest areas. Explainability is a key feature of this model: we apply XAI techniques to produce interpretable, natural-language explanations that connect individual preferences to spatial and behavioral characteristics. A demo showcasing this capability is available online¹.

¹<https://cityhood.vercel.app/>

2 Related Works

Prior studies have examined the complementarity of LBSN data, comparing platforms like Instagram and Foursquare [31], proposing frameworks for urban dynamics [20], and exploring their role in neighborhood and tourism studies [21]. While some validate LBSN data against official sources (e.g., World Values Survey²) [29], our work focuses on how different platforms independently model urban behavior and where discrepancies emerge. Building on multi-scale approaches such as micro–meso–macro frameworks [26] and wavelet transformations [36], we examine how granularity influences urban preference modeling. Additionally, while prior work has used clustering techniques on LBSN and POI data to define urban zones, such as Livehoods [5] and functional topic models [8], our grid-based method aggregates spatial cells based on user behavior to ensure both structural and behavioral consistency.

Research on urban mobility has leveraged large-scale data to uncover patterns such as Lévy flight dynamics from Twitter traces [3], migration flows through temporal graphs [12], class-based transit differences using smart card data [27], and formulate models to explain urban change [33]. Moving beyond trajectory modeling, our work maps user interests across cities and neighborhoods. We extend the returner/explorer framework [23] by integrating contextual, geographic, and cultural factors into an explainable, multi-scale recommendation system that predicts urban interest and justifies suggestions via user behavior and regional signals—advancing transparency beyond POI-level systems [4, 6, 7].

3 Data

3.1 LBSN Datasets

Foursquare (FS): We use public check-in data shared on Twitter, covering different cities in different continents. Each dataset entry, curated by [29], includes check-in time, venue name/category, and user ID, for a total of 7,537,057 check-ins made by 1,403,808 users on 2,739,548 venues.

Google Places (GP): This dataset, provided by [11, 24] for academic use, contains 11,453,845 reviews from 5,054,567 users of Google Plus for 3,114,353 venues on Google Maps around the world. It includes user details (name, education, employment), review text (multilingual), ratings, timestamps, and unique user IDs, along with venue information (name, category, hours, contact, address, coordinates).

Extended Google Places (GP+): We also consider a bigger dataset from Google Places, for user-level interest analysis, as it provides a larger volume of reviews per user, essential to our high vs low interest analysis. Originally introduced in [15, 37], the dataset contains 666,324,103 geo-localized reviews of 4,963,111 venues, written by 113,643,107 users across the United States up to September 2021. Notably, 98.9% of the data covers the period from 2015 to 2021.

3.2 Cultural, Socioeconomic, Electoral Datasets

To enrich the experiments, we collected diverse contextual attributes across cities in different countries. In Brazil, we used data from the 2010 Brazilian Demographic Census (IBGE)³, including

average monthly income, literacy rate, population size, and racial composition (White, Black, Brown, Yellow, Indigenous) for each city neighborhood. Electoral data for the 2014 Brazilian presidential election (second round) are obtained from the Regional Electoral Court (TRE)⁴, following the approach of [17] and [12], which uses voting discrepancies to assess political polarization. The TRE dataset provides polling place-level results across different electoral zones. Since direct neighborhood mappings are unavailable, we match voting location addresses to neighborhoods using Google Maps API, ensuring alignment with the city’s limits. This method connects most polling places to neighborhoods, though some neighborhoods can lack data, mainly due to shared or misassigned polling locations, particularly in peripheral or commercial areas.

For U.S. regions, socio-demographic indicators—racial composition (White, Black, Asian, Hispanic), median household income, population size, higher education attainment, and employment rate—are obtained from the National Historical Geographic Information System (NHGIS)[19], based on the 2017–2021 American Community Survey. Political orientation is estimated using precinct-level data from the 2020 U.S. presidential election⁵, aggregated at the ZIP code and county levels. Finally, we assess cultural characteristics using Scenes Theory [10], deriving 15-dimensional scene vectors based on venue types and an alternative proxy based on venue category frequency distributions.

3.3 Urban Areas

For each addressed city, we use specific delimitations like census tracts and ZIP codes, neighborhoods, or boroughs. For more coverage of different and consistent area sizes across countries, we developed h3-cities⁶, a tool that integrates OpenStreetMap with Uber’s Hexagonal Hierarchical Geospatial Indexing System to subdivide cities into hexagons of multiple resolutions. This enables consistent, multi-scalar urban analysis, as emphasized by [26] and [22]. We tested four different hexagonal grid resolutions (*hr*): *h6* (average area of 36.12km²), *h7* (5.16km²), *h8* (0.74km²), and *h9* (0.11km²). Figure 1 shows a size and shape comparison of the different urban areas in New York.

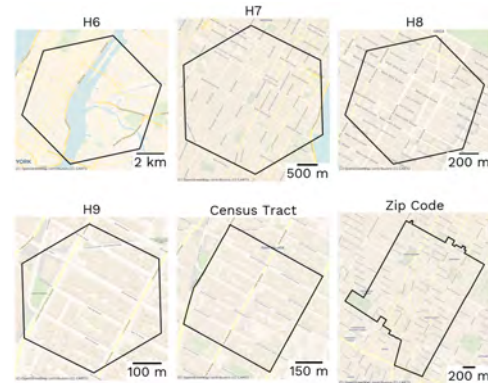


Figure 1: Different granularity levels in New York

²<https://www.worldvaluessurvey.org>

³<https://sidra.ibge.gov.br>

⁴<https://www.tre-pr.jus.br/eleicoes/eleicoes-antiores/eleicoes-2014>

⁵<https://github.com/TheUpshot/presidential-precinct-map-2020>

⁶<https://h3-cities.vercel.app>

4 Methodology

4.1 LBSN Datasets

In the experiments, we addressed LBSN data from two datasets (Google Places and Foursquare) on three different areas: Curitiba, Brazil, London, UK, and several cities and counties in the USA. Tables 1 and 2 summarize the datasets from Google Places and Foursquare, respectively. Overall, the Google Places dataset offers a larger data volume in most regions, except for Curitiba, where Foursquare has more check-ins. Foursquare uses broader categories (e.g., “Food”), whereas Google Places provides more specific labels (e.g., “Italian Restaurant”). Additionally, Google Places data span a longer period, but sparse pre-2010 data have been excluded to match Foursquare’s coverage (2010–2014).

Table 1: Description of Google Places dataset

	Curitiba	London	USA	
			cities	counties
Reviews	8,372	178,231	1,191,934	1,632,165
Users	4,909	75,897	394,588	486,393
Venues	2,213	31,075	186,639	286,075
Categories	685	1,81	3,439	3,793
Period		from 2010 to 2014		

Table 2: Description of Foursquare dataset

	Curitiba	London	USA	
			cities	counties
Reviews	53,253	27,088	398,805	495,698
Users	5,116	9,128	83,414	85,946
Venues	8,523	11,104	122,808	164,655
Categories	368	427	562	567
Period	2014*	from 2012 to 2014		

* from April to June

A comparison of top categories in Curitiba (Figure 2) highlights key distinctions—e.g., Foursquare includes venues like homes and workplaces, which Google Places omits.




Figure 2: 50 most popular categories in Curitiba [28]

Since our models focus on areas frequented by people, we also analyze user activity frequency and intervals between interactions. In Google Places, 25.7% of users posted at least 2 reviews, 6.7% posted 5, and 2.5% posted 10 or more. Foursquare shows higher engagement: 61.3% performed 2 check-ins, 25.3% performed 5, and

9.6% performed 10 or more. Regarding the time between interactions, roughly half are under 6 hours on both platforms. Longer gaps differ: one-third of Google Places users wait over a week between posts, while Foursquare users often check in every 6–24 hours (one-fourth of cases) or every 1–7 days (another fourth). In Curitiba, Foursquare users rarely go a week without checking in. These patterns reveal that Google Places users engage less regularly, often in short bursts or after long gaps, whereas Foursquare users are more consistent. This difference and category specificity shape how each platform models urban activity.

This exploratory analysis raises a central question in our study: Can models derived from each platform yield comparable insights despite the differences? If so, may models reveal generalizable urban behavior patterns across platforms with distinct user dynamics? In the next sections, we describe how we derive the interest Networks (iNETs), the basis of our models, from LBSN data and how we explore them to provide insights into mobility patterns and user interests in complex urban environments.

4.2 Formation of iNETs

iNETs model user interests in urban areas as weighted graphs. Starting from an individual user iNET (Fig.3a), each user is linked to the regions they reviewed, weighted by the number of reviews. Regions—whether cities, counties, or neighborhoods—are of the same class. For every two regions we compute the inter-regional interest subgraph (Fig. 3b) as the number of users who reviewed both regions. This number becomes the edge weight of the aggregated iNET (Fig. 3c). We also define self-loops in a region as the number of users who have made at least two reviews in the same region.

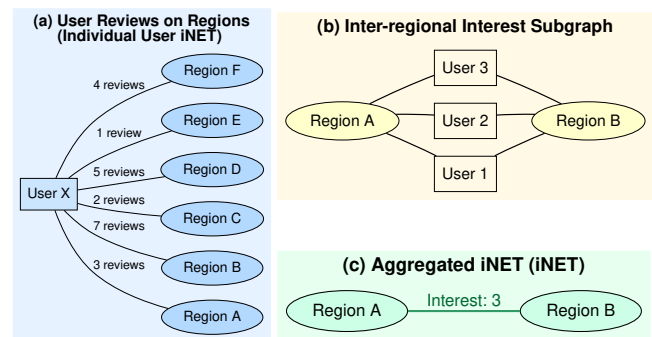


Figure 3: iNETs modelling

4.3 Validating iNETs across granularity levels

In the second phase of our research, we assess the similarity between aggregated iNETs modeled by different LBSNs across varying granularity levels. To this end, we use the FS and GP datasets (Section 3.1), comparable in size and period, but with several user behavior differences as shown below, covering Curitiba (Brazil), London (UK), and the top 20 most popular U.S. cities and counties.

We apply specific metrics across multiple spatial granularities to assess the similarity between the iNETs modeled by different LBSNs. Then we analyze how area size and shape affect modeling results and their consistency. The similarity metrics used are:

- **Edge weight** similarity: We calculate Pearson and Spearman correlations between corresponding edge weights in the Google Places and Foursquare networks.
- **Node importance** similarity: We evaluate eigenvector centrality rankings using Kendall's Tau to measure agreement in node importance.

These metrics help assess whether different LBSNs capture similar spatial interest patterns (edges) and consistently identify key urban areas (nodes).

To address the differences observed between iNETs, we propose a new approach that defines Urban Preference Zones (UPZones) based on iNET data. By analyzing edge weights, we find that the strongest user interests typically occur between neighboring areas. This indicates that zones should be shaped by user behavior and spatial proximity rather than arbitrary boundaries.

Our approach identifies densely connected, geographically adjacent zones using the $h9$ granularity level (see an example in Figure 1) as the regions in the iNET. We detect tightly-connected communities by applying the Leiden algorithm [34] ($\gamma = 1$, run until convergence). Finally, we employ TF-IDF to generate word clouds highlighting distinctive POI categories in each urban preference zones. To compare urban interest zones across LBSNs, we align the Leiden-derived clusters using the Hungarian algorithm [14], and evaluate their similarity through Normalized Mutual Information (NMI) for shared cluster information and Rand Index [13] for pairwise sample agreement.

4.4 Understanding User Interests

In the third phase, we focus on predicting if a region is highly interesting to a user based on other visited regions. Therefore, we use the GP+ dataset (see Section 3.1) for the larger review volume per user. We analyze the users who made reviews on at least six distinct places. We use the number of reviews to rank the places visited into low- and high-interest regions, but the ranks can blur for some users. To address the blurring of rank differences as users visit more places, particularly at the county level, the analysis focuses on a subset of users whose k -th most reviewed place has a higher review count than the $(k+1)$ -th. For the analysis where $k = 3$, this selection criterion retains 3.07M users and their 210.7M reviews, with 65.9% of users posting all their reviews within four years and averaging 68 (median: 49) reviews each. To enhance this third phase analysis, we incorporate cultural, socioeconomic, and political contextual attributes (see Section 3.2).

Several factors are evaluated for the third research phase to understand aggregated and individual user interests. We consider geographic distance to identify mobility patterns, such as a tendency toward proximate destinations, and population size to capture preferences for residential or business-dense areas. Income and employment are used to assess economic alignment. Educational attainment, reflected in the proportion of residents holding bachelor's degrees or literacy (in Curitiba), proxies for intellectual engagement. Racial composition reveals demographic affinity. Political orientation provides insight into the role of ideological alignment. Venue categories evaluate the similarity of the experiences that the places offer. Scenes help us understand the importance of cultural aspects in urban areas.

We used several distance and similarity metrics to better assess how regions resemble each other based on each characteristic.

On the aggregated level of iNETs, we correlate the edge weights with the distance and similarity metrics for each feature analyzed. For the individual user level of iNETs, we propose a method for personalized user interest modeling and recommendation. We first differentiate a user's visited regions (e.g., cities, neighborhoods) into high-interest (top) and low-interest (bottom) sets based on review counts. Specifically, we rank regions by review frequency and select the top k as the high-interest set (e.g., for six visited cities and $k = 3$, the top 3 by reviews are high-interest, the rest are low-interest). This can be further applied at multiple scales. For instance, using $m = 2$, we could identify the top 2 most-reviewed ZIP codes within each top k city to form a finer-grained high-interest set.

We then train a LightGBM classifier to predict whether a region belongs to the user's high-interest set, based on the mean distances/similarities to the other top and bottom regions. To understand the factors driving high interest, we analyze feature importance using permutation importance and SHAP values [18] on the test set (derived from an 80/20 train-test split).

The model assigns greater weight to high-interest regions, based on the ratio of low- to high-interest samples in the training data. Hyperparameters are tuned using a randomized grid search.

Finally, we show the power of this modeling by analyzing the 'returners' vs. 'explorers' dichotomy [23]. Returners are identified as users where the radius of gyration of their top k locations exceeds half the radius of gyration of all their visited locations, demonstrating the method's ability to capture distinct behavioral patterns.

5 Results

5.1 The impact of different Granularity Levels and Platforms

In the U.S., we modeled 240 iNETs per LBSN (GP or FS), i.e., 20 cities (or counties), each evaluated using six granularities. The smallest is New York County (GP, $h6$): 7 nodes, 28 edges. The largest is Los Angeles County (GP, $h9$): 10,363 nodes, 455,976 edges. On all these iNETs, we perform similarity analyses between LBSNs (see Section 4.3). The results are presented in Figure 4.

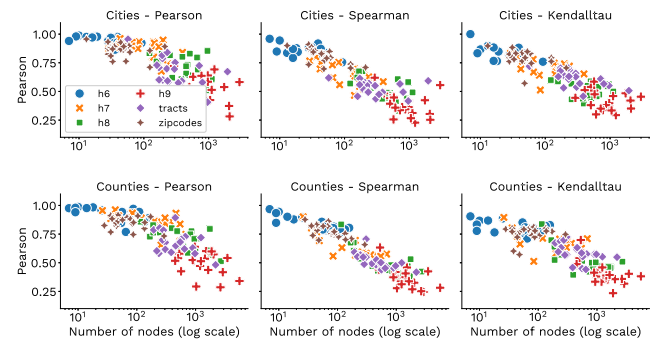


Figure 4: Results of correlation between both iNETs in USA

In this figure, the y-axis indicates the correlation value analyzed for each granularity level, while the x-axis presents, on a logarithmic

scale, the number of nodes in both iNETs (GP and FS). Observing Figure 4, there is an evident pattern that the more nodes an iNET has, the greater the differences between the iNETs modeled by the studied LBSNs. It is also important to highlight that counties vary in area from 87 km² to 23,895 km² while cities vary from 177 km² to 1,740 km². As a result, some regions have more nodes at the *h7* granularity level than other regions at the *h8* level, even if the *h8* nodes are smaller than the *h7* nodes.

Figure 5 presents similarity correlations between iNETs from Google Places and Foursquare in Curitiba, segmented by neighborhoods, London segmented by boroughs, and both using *h6* to *h9* granularities. Notably, neighborhood sizes in Curitiba typically fall between *h7* and *h8*, and London boroughs fall between *h6* and *h7*. The findings suggest that finer spatial granularity reveals greater differences between iNETs.




Figure 5: Results of correlation between both iNETs

The granularity level *h6* (36.12 km²) provides consistent insights across LBSNs when examining central urban areas and their broader relationships. Finer granularities (*h7*, *h8*) yield more detailed analyses but may introduce diverse iNETs, capturing LBSN-specific behaviors rather than general trends. The *h9* level, while enabling sharper differentiation between LBSNs, suffers from increased noise due to sparse data and higher node counts.

Government-defined divisions (e.g., census tracts, ZIP codes, neighborhoods) align with certain granularities: census tracts and neighborhoods resemble *h7/h8*, while ZIP codes and boroughs parallel *h6* in homogeneity. Socioeconomic data linked to administrative divisions aid in investigating interest factors, but hierarchical granularities offer flexibility, from broad patterns (*h6*) to micro-scale interactions (*h9*). These findings align with [36], confirming that granularity choice shapes complementary urban insights.

5.2 Urban Preference Zones (UPZones)

This section uses London as a case study for the urban preference zones, although similar results are found in Curitiba. In London (using *h9* cells), GP iNET includes 123,292 reviews from 67,276 users, forming 396,876 edges (1,533 self-loops) across 5,654 of 6,600 cells. FS iNET has 22,738 check-ins from 8,178 users, yielding 34,075 edges (1,209 self-loops) across 2,899 cells.

Using the method from Section 4.3, we identify 1,760 UPZones from Google Places and 1,023 from Foursquare. Comparing these zones to the city's boroughs [28] shows that urban preference zones offer finer-grained insights into user interests than traditional administrative divisions.

Figure 6 shows how UPZones constructed from the iNET at *h9* granularity capture semantically meaningful areas in London. Each colored and numbered region represents a UPZone derived from Google Places, with numbers indicating the top 6 zones by activity. Word clouds illustrate the most prominent categories in each zone (see Section 4.3), emphasizing the dominant place types in these areas.



Figure 6: Characterization of the main UPZones found

The figure shows that UPZone 1 includes Soho, a vibrant area known for nightlife and dining; UPZone 2 captures an artistic hub hosting performances and creative activity; and UPZone 3 corresponds to Mayfair, noted for high-end retail. UPZone 6 covers the South Bank, home to cinemas, galleries, and landmarks like the London Eye. Other UPZones similarly align with popular areas frequented by users with shared demographic interests. While Foursquare data also identifies central zones, its smaller dataset may limit comparability. Overall, the proposed method effectively detects urban preference zones, revealing areas that transcend administrative boundaries and offer deeper insight into urban user behavior.

Using the metrics from Section 4.3, we compare UPZones derived from Google Places and Foursquare iNETs, obtaining an NMI of 0.6364 and a Rand Index of 0.7378. These results suggest that, despite differences in the underlying iNETs, both LBSNs yield moderately similar UPZones, supporting the robustness of our modeling approach.

5.3 Influence of Cultural, Socioeconomic, Political, and Geographic Factors

For the aggregate level iNET for Google Places users in Curitiba, we analyze 1,127 users (with reviews ≥ 2) and 4,590 reviews, while Foursquare's iNET encompasses 3,933 users and 52,033 check-ins. The resulting Google Places graph covers 69 of 75 neighborhoods, with 1,287 edges (53 self-loops), while Foursquare's has 75 nodes and 1,618 edges (58 self-loops). Both networks show concentrated activity in central commercial and leisure districts, with Google Places extending further into peripheral areas. We use the neighborhoods due to their association with socioeconomic data.

As per Section 4.3, we correlate the edge weights of the modeled networks with the pairwise node distances across key characteristics for filtered iNETs. Using Spearman correlation, the results for Google Places and Foursquare networks are shown in Table 3. We kept only the top 75% of edges by weight, emphasizing stronger interest connections.


The analysis shows that geographic distance has the most strongly negative correlation (Table 3), reflecting known mobility constraints where long-distance visits are less common [1, 3, 9, 25]. Then, education levels and racial composition showed a lower but interesting

Table 3: Spearman Correlations with the Analyzed Factors in Curitiba

	Google Places	Foursquare
Population	-0.14	-0.03
Income Differences	-0.17	-0.06
Cultural Affinity	-0.28	0.00
Political Polarization	-0.26	-0.14
Venues Similarity	0.33	0.12
Racial Composition	-0.26	-0.14
Education Levels	-0.33	-0.23
Geographic Distance	-0.4	-0.42

correlation in Curitiba. After that, venue similarity emerges as the second strongest correlated in Google Places, highlighting the relevance of shared venue experience traits in driving urban interest, but lower in Foursquare, likely due to having categories like the user’s own home. Finally, political polarization, cultural affinity (Except by GP, where a more venue-centered data better fits with Scenes Theory), income differences, and population size exhibited weak correlations with people’s interest patterns, suggesting users do not strongly prefer areas with similar or dissimilar aspects based on these traits.

For the user-level analysis on U.S. counties (see Section 4.4), we perform a preliminary test using $k = 3$, which yields a recall of 0.68 and an F1-score of 0.55 on the test set. When applied to the corresponding ZIP codes using $m = 2$, the model achieves a recall and F1-score of 0.71. To identify the most influential features for distinguishing high-interest areas, we use permutation feature importance and mean absolute SHAP values. After applying min-max normalization, the top 10 features—ranked by their mean normalized importance across scenarios—are presented in Figure 7.

**Figure 7: Most influential features for high-interest areas**

Geographic distance to top regions is the strongest predictor, indicating that users favor venues near those they’ve previously visited more. Similarity in venue categories to top regions helps identify high-interest areas, while dissimilarity suggests low-interest zones. Population size primarily influences preferences at the county level, reflecting urban vs. rural appeal. Lastly, race and political alignment show moderate importance at both the county and ZIP code levels.

Analyzing the dichotomy between the returners and explorers reveals a notable pattern. When predicting regions of interest separated for these classes of users, the model considering explorers outperforms the one focusing on returners, achieving a recall of 0.80 and an F1-score of 0.63, compared to returners’ 0.71 recall and 0.54 F1-score. For $k = 3$, 47% of users are classified as explorers. The key difference between groups lies in feature importance. For explorers, proximity to other high-interest cities is the most influential factor—nearly twice as important as the second-ranked feature, venue similarity. In contrast, returners are more influenced by venue category similarity, with geographic distance playing a minor role. This distinction appears only at the city level; at the neighborhood level, both groups show similar behavior in selecting high-interest areas.

To better demonstrate the strength of our user-level modeling, we propose the CityHood demo⁷, which allows users to explore personalized city and neighborhood recommendations based on their past preferences, enhanced with AI-powered explainability.

6 Conclusion

This work introduces *Interest Networks (iNETs)* to model urban preferences using Location-Based Social Network (LBSN) data from Google Places and Foursquare. Our analysis reveals that coarser spatial granularities (e.g., city or borough levels) yield consistent cross-platform insights, while finer resolutions (e.g., neighborhoods or hexagons) expose platform-specific behaviors. Key findings highlight that, in general, geographic proximity and venue similarity dominate user interest patterns, while socioeconomic and political factors have a comparatively smaller influence.

To bridge Google Places and Foursquare platform discrepancies, we propose *Urban Preference Zones (UPZones)*—behaviorally coherent clusters that transcend administrative boundaries. We further develop a scalable, explainable recommendation system that adapts to user profiles (e.g., *explorers* vs. *returners*) and provides interpretable justifications via XAI. Although evaluating such a system is challenging—e.g., varying k and m or testing its predictions on various subsets and permutations of user history—we plan to explore this further in future work. Our public demo demonstrates practical applications for urban planning and mobility research.

While the study uses partially dated datasets, the methodology remains applicable to modern platforms and contexts. In future work, we aim to better investigate how the sampling biases in the LBSNs studied impact the results, integrate textual reviews through natural language processing, incorporate time-aware mobility patterns, and validate our findings through local expert feedback. Our tools, datasets, and public demo provide a foundation for further research and applications in urban computing, mobility modeling, and recommender systems.

Acknowledgments

This research is partially supported by the SocialNet project (grant 2023/00148-0 of FAPESP), by CNPq (grants 314603/2023-9, 441444/2023-7, 409669/2024-5, and 444724/2024-9) and the INCTs ICoNiOT (CNPq grant 405940/2022-0 and CAPES FinanceCode 88887.954253/2024-00) and TILD-IAR (CNPq grant 408490/2024-1).

⁷<https://cityhood.vercel.app/>

References

- [1] D. Brockmann, L. Hufnagel, and T. Geisel. 2006. The scaling laws of human travel. *Nature* 439, 7075 (Jan. 2006), 462–465. doi:10.1038/nature04292
- [2] Yong Chen, Xiqun Michael Chen, and Ziyu Gao. 2024. Toward equitable, transparent, and collaborative human mobility computing for smart cities. *The Innovation* 5, 5 (2024).
- [3] Zhiyuan Cheng, James Caverlee, Kyumin Lee, and Daniel Sui. 2011. Exploring Millions of Footprints in Location Sharing Services. *Proceedings of the International AAAI Conference on Web and Social Media* 5, 1 (Aug. 2011), 81–88. doi:10.1609/icwsm.v5i1.14109
- [4] Zhiyuan Cheng, James Caverlee, Kyumin Lee, and Daniel Z Sui. 2012. Fused matrix factorization with geographical and social influence in location-based social networks. *Proc. of the AAAI* (2012).
- [5] Justin Cranshaw, Raz Schwartz, Jason Hong, and Norman Sadeh. 2012. The Livehoods Project: Utilizing Social Media to Understand the Dynamics of a City. *Proceedings of the International AAAI Conference on Web and Social Media* 6, 1 (2012), 58–65. doi:10.1609/icwsm.v6i1.14278 Number: 1.
- [6] Luis M De Campos, Juan M Fernández-Luna, and Juan F Huete. 2024. An explainable content-based approach for recommender systems: a case study in journal recommendation for paper submission. *User Modeling and User-Adapted Interaction* 34, 4 (2024), 1431–1465.
- [7] Keyan Gao, Xu Yang, Chenxia Wu, Tingting Qiao, Xiaoya Chen, Min Yang, and Ling Chen. 2020. Exploiting Location-Based Context for POI Recommendation When Traveling to a New Region. *IEEE Access* 8 (2020), 52404–52412.
- [8] Song Gao, Krzysztof Janowicz, and Helen Couclelis. 2017. Extracting urban functional regions from points of interest and human activities on location-based social networks. *Transactions in GIS* 21, 3 (2017), 446–467. doi:10.1111/tgis.12289 eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/tgis.12289>
- [9] Marta C. González, César A. Hidalgo, and Albert-László Barabási. 2008. Understanding individual human mobility patterns. *Nature* 453, 7196 (June 2008), 779–782. doi:10.1038/nature06958
- [10] Fernanda R. Gubert, Gustavo H. Santos, Myriam Delgado, Daniel Silver, and Thiago H. Silva. 2025. Culture Fingerprint: Identification of Culturally Similar Urban Areas Using Google Places Data. In *Proc. of ASONAM*. Rende, Italy, 286–297.
- [11] Ruining He, Wang-Cheng Kang, and Julian McAuley. 2017. Translation-based Recommendation. In *Proceedings of the Eleventh ACM Conference on Recommender Systems (RecSys '17)*. ACM. doi:10.1145/3109859.3109882
- [12] Peng Huang and Carter T. Butts. 2023. Rooted America: Immobility and Segregation of the Intercounty Migration Network. *American Sociological Review* 88, 6 (Nov. 2023), 1031–1065. doi:10.1177/00031224231212679
- [13] Lawrence Hubert and Phipps Arabie. 1985. Comparing partitions. *Journal of Classification* 2, 1 (Dec. 1985), 193–218. doi:10.1007/BF01908075
- [14] H. W. Kuhn. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly* 2, 1-2 (1955), 83–97. doi:10.1002/nav.3800020109 eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/nav.3800020109>
- [15] Jiacheng Li, Jingbo Shang, and Julian McAuley. 2022. UCTopic: Unsupervised Contrastive Learning for Phrase Representations and Topic Mining. In *Proc. of the ACL '22*. ACL, Dublin, Ireland.
- [16] Qingsong Liu, Zheyang Wang, and Xinyue Ye. 2018. Comparing mobility patterns between residents and visitors using geo-tagged social media data. *Transactions in GIS* 22, 6 (2018), 1372–1389.
- [17] Xi Liu, Clio Andris, and Bruce A. Desmarais. 2019. Migration and political polarization in the U.S.: An analysis of the county-level migration network. *PLOS ONE* 14, 11 (Nov. 2019), e0225405. doi:10.1371/journal.pone.0225405
- [18] Scott M. Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Proc. of NIPS*. Red Hook, NY, USA, 4768–4777.
- [19] Steven Manson, Jonathan Schroeder, David Van Riper, Katherine Knowles, Tracy Kugler, Finn Roberts, and Steven Ruggles. 2024. National Historical Geographic Information System: Version 19.0.
- [20] Pablo Martí, Leticia Serrano-Estrada, and Almudena Nolasco-Cirugeda. 2019. Social media data: Challenges, opportunities and limitations in urban studies. *Computers, Environment and Urban Systems* 74 (2019), 161–174.
- [21] Almudena Nolasco-Cirugeda and Clara García-Mayor. 2022. Social dynamics in cities: Analysis through LBSN data. *Procedia Computer Science* 207 (2022), 877–886.
- [22] Elek Pafka. 2022. Multi-scalar urban densities: from the metropolitan to the street level. *URBAN DESIGN International* 27, 1 (March 2022), 53–63. doi:10.1057/s41289-020-00112-y
- [23] Luca Pappalardo, Filippo Simini, Salvatore Rinzivillo, Dino Pedreschi, Fosca Giannotti, and Albert-László Barabási. 2015. Returners and explorers dichotomy in human mobility. *Nature Communications* 6, 1 (Sept. 2015).
- [24] Rajiv Pasricha and Julian McAuley. 2018. Translation-based factorization machines for sequential recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys '18)*. ACM. doi:10.1145/3240323.3240356
- [25] I. Rhee, M. Shin, S. Hong, K. Lee, and S. Chong. 2008. On the Levy-Walk Nature of Human Mobility. In *IEEE INFOCOM 2008 - The 27th Conference on Computer Communications*. IEEE. doi:10.1109/infocom.2008.145
- [26] Mikhail Rogov and Céline Rozenblat. 2018. Urban Resilience Discourse Analysis: Towards a Multi-Level Approach to Cities. *Sustainability* 10, 12 (Dec. 2018), 4431. doi:10.3390/su10124431 Number: 12 Publisher: Multidisciplinary Digital Publishing Institute.
- [27] Priscila Santin, Fernanda R. Gubert, Mauro Fonseca, Anelise Munaretto, and Thiago Henrique Silva. 2020. Characterization of Public Transit Mobility Patterns of Different Economic Classes. *Sustainability* 12, 22 (2020). doi:10.3390/su12229603
- [28] Gustavo H. Santos, Fernanda R. Gubert, Myriam Delgado, and Thiago H. Silva. 2025. Modeling Interest Networks in Urban Areas: A Comparative Study of Google Places and Foursquare Across Countries. *Journal of Internet Services and Applications* 16, 1 (Mar. 2025), 25–42. doi:10.5753/jisa.2025.5152
- [29] Thiago H. Silva, Pedro O.S. Vaz de Melo, Jussara M. Almeida, Mirco Musolesi, and Antonio A.F. Loureiro. 2017. A large-scale study of cultural differences using urban data about eating and drinking preferences. *Information Systems* 72, Supplement C (2017), 95 – 116. doi:10.1016/j.is.2017.10.002
- [30] Thiago H Silva and Daniel Silver. 2024. Using Graph Neural Networks to Predict Local Culture. *arXiv* (2024). arXiv:2402.17905 [cs.LG]
- [31] Thiago H. Silva, Pedro O. S. Vaz de Melo, Jussara M. Almeida, Juliana Salles, and Antonio A. F. Loureiro. 2013. A comparison of Foursquare and Instagram to the study of city dynamics and urban social behavior. In *Proc. ACM SIGKDD Int. Workshop on Urban Computing (UrbComp '13)*. Chicago, USA.
- [32] Thiago H. Silva, Aline Carneiro Viana, Fabrício Benevenuto, Leandro Villas, Juliana Salles, Antonio Loureiro, and Daniele Quercia. 2019. Urban Computing Leveraging Location-Based Social Network Data: A Survey. *Comput. Surveys* 52, 1 (Feb. 2019), 1–39. doi:10.1145/3301284
- [33] Daniel Silver and Thiago H. Silva. 2023. Complex causal structures of neighbourhood change: Evidence from a functionalist model and yelp data. *Cities* 133 (2023), 104130. doi:10.1016/j.cities.2022.104130
- [34] V. A. Traag, L. Waltman, and N. J. van Eck. 2019. From Louvain to Leiden: guaranteeing well-connected communities. *Scientific Reports* 9 (March 2019), 5233. doi:10.1038/s41598-019-41695-z
- [35] David A. M. Veiga, Gabriel B. Frizzo, and Thiago H. Silva. 2019. Cross-cultural study of tourists mobility using social media. In *Proceedings of the 25th Brazilian Symposium on Multimedia and the Web (Rio de Janeiro, Brazil) (WebMedia '19)*. Association for Computing Machinery, New York, NY, USA, 313–316. doi:10.1145/3323503.3360620
- [36] Dr. Qiong Wu, Jinxiang Tan, Fengxiang Guo, Hongqing Li, Shengbo Chen, and Sheng Jiang. 2020. Multi-Scale Identification of Urban Landscape Structure Based on Two-Dimensional Wavelet Analysis: The Case of Metropolitan Beijing, China. *Ecological Complexity* 43 (Aug. 2020), 100832. doi:10.1016/j.ecocom.2020.100832
- [37] An Yan, Zhankui He, Jiacheng Li, Tianyang Zhang, and Julian McAuley. 2023. Personalized Showcases: Generating Multi-Modal Explanations for Recommendations. In *Proc. of the SIGIR*. Taipei, Taiwan, 5 pages.
- [38] Chao Yang, Meng Xiao, Xuan Ding, Wenwen Tian, Yong Zhai, Jie Chen, Lei Liu, and Xinyue Ye. 2019. Exploring human mobility patterns using geo-tagged social media data at the group level. *Journal of Spatial Science* 64, 2 (2019), 221–238.

A Additional Reproducibility Details

Besides the details in the methodology section, we present some information on the implementation to aid in reproducibility tasks.

Census tract and ZIP code delimitations have been obtained from the U.S. Census Bureau’s Tiger geographic database⁸; and the neighborhood polygons from the Curitiba city hall website⁹. London’s boroughs were collected from the London Datastore¹⁰.

The h3-cities¹¹ tool makes it easy to retrieve H3 cells. Enter the name of your desired city or region, select the resolution, and the download option will appear. You’ll also see a visual preview of the area and its corresponding hexagons.

We use several distance and similarity metrics between regions, where each region belongs to the same granularity class. The metrics are as follows:

Geographic Proximity is computed as the Euclidean distance between the centroids of two regions in the ESRI:102005 (USA) and UTM22S (Curitiba) projections.

Population, income, education (bachelor or literacy rates), employment, and voting similarities were measured as the absolute difference between values in each region.

Race differences were computed as:

$$\frac{1}{2} \sum_{i=1}^n \left| \frac{P_i(r)}{P(r)} - \frac{P_i(r')}{P(r')} \right|$$

where $P(r)$ represents the population size of region r , and $P_i(r)$ is the population size of the i -th racial category in region r . This work uses the racial categories: White, Black, Asian, and Hispanic for the USA regions and White, Black, Brown, Yellow, and Indigenous for Brazil.

Cultural Similarity was calculated using the Euclidean Distance in the 15 scenes dimensions studied.

Venue similarity was computed using cosine similarity between each region’s frequency vectors of venue categories.

On the training of the LightGBM model, the following ranges of hyperparameter are tested on the random grid search: The hyperparameter were tested over the following ranges: number of trees (100–300), learning rate (0.01–0.1), maximum tree depth (3–10), number of leaves per tree (20–50), minimum child samples (10–30), and L1/L2 regularization strengths (0–0.5).

To ensure consistency in filtering the top 75% of edges by weight, we first sorted all edges in descending order of weight. For edges with weights equal to the cutoff threshold, we broke ties deterministically by sorting lexicographically using node IDs (neighborhood names), with smaller nodes first. This guarantees that the same subset of edges is selected across runs.

All analyses were run on a server with an Intel i7-10700 CPU and 128 GB RAM. After obtaining the geo-localized reviews, iNETs were built using DuckDB, which enabled efficient computation of region-to-region interest graphs through SQL-based aggregation over user co-visitation data. By joining check-in or review records on user IDs and grouping by region pairs, we calculated the number of users who interacted with both regions, forming the weighted edges of

the iNETs. This method was applied consistently across all spatial granularities ($h6$ – $h9$, census tracts, ZIP codes, neighborhoods, and boroughs) and for both LBSNs—Google Places (GP) and Foursquare (FS). Thanks to DuckDB’s optimized performance for analytical queries, the complete construction of all iNETs across all cities and granularities, for both GP and FS datasets, as well as the correlation analyses, was completed in under 5 minutes total. The construction of Urban Preference Zones (UPZones) in London using the Leiden algorithm on $h9$ granularity took about 1s, with parameters $\gamma = 1$, $n_iterations = -1$ (running until convergence), and using the iNET edge weights. LightGBM training, including SHAP and permutation importance analysis, required approximately 2 hours each for both city-level and ZIP code-level models.

Received 15 May 2025; revised 4 Jun 2025; accepted 4 Jun 2025

⁸<https://www.census.gov/cgi-bin/geo/shapefiles/index.php>

⁹<https://ippuc.org.br/geodownloads/geo.htm>

¹⁰<https://data.london.gov.uk/dataset/statistical-gis-boundary-files-london>

¹¹<https://h3-cities.vercel.app>