# Culture Fingerprint: Identification of Culturally Similar Urban Areas Using Google Places Data

Fernanda R. Gubert[1], Gustavo H. Santos[1], Myriam Delgado[1], Daniel Silver[2], and Thiago H Silva[1]

[1] Universidade Tecnológica Federal do Paraná, Curitiba, Brazil
[2] University of Toronto, Toronto, Canada
fernandagubert,gustavohenriquesantos}@alunos.utfpr.edu.br,
{myriamdelg,thiagoh}@utfpr.edu.br, dan.silver@utoronto.ca

**Abstract.** This study investigates methods using a global data source, Google Places, to identify culturally similar urban areas without relying on difficult-to-access data like user preferences shown through check-ins. We propose and assess a simple method requiring only information about place types and their frequency in the studied areas, and a more advanced method that enhances venue categories using Scenes Theory - it helps us understand the cultural significance of everyday urban life. We tested our methods in 14 cities worldwide and all US states. The results suggest that a straightforward approach based on category frequencies can highlight major cultural differences. However, the Scenes Theory-based method provides a better understanding of cultural nuances, as the ones supported by survey data.

**Keywords:** Cultural signature · large scale assessment · Google Places

## 1 Introduction

Traditional methods like surveys and interviews are important data sources for studying culture in its complexity. However, these methods have drawbacks (e.g. high costs and time-consuming), which limit their scalability. To remedy this situation, some works evaluate alternative geolocalized data sources from the web to study culture. These sources exist on a global scale and are faster to obtain. Studies have shown the usefulness of these data sources in several domains [6, 16, 19, 21], including the cultural ones [2, 3, 8, 15, 17, 18].

Bancilhon et al. [2] explore an approach to quantifying a society's culture through city street names, revealing that these names reflect cultural values. Using Foursquare data, Senefonte et al. [15] examine how regional and cultural characteristics affect the mobility patterns of both tourists and residents. The results indicate that the tourist's origin significantly influences their behavior, especially in large cultural differences between the origin and destination. Silva and Silver [18] introduce a graph neural network method for predicting local culture. They evaluate their approach on Yelp data, showing that it could help predict local culture even when traditional local information is unavailable.

When aiming to provide methods based on geolocalized web data to describe local culture, some research indicates that eating and drinking habits can be a valuable option [3, 8, 17]. These studies illustrate promising approaches to identifying cultural boundaries and similarities between different societies at different scales. However, they rely on user preferences, typically manifested through check-in data, which is challenging to obtain in practice for many users or with global coverage. Another perspective follows the argument presented in [11], which suggests that the availability of resources and services that meet the population's needs contributes to forming a local identity. What is notable about this approach is the opportunity to consider multiple aspects of culture, as the resources of a region can be associated with various categories like religion, cuisine, and arts, providing a format that is still little explored. Our approach aligns with this direction by exploring Scenes Theory [20], which captures local public cultural dimensions embodied in venues such as cafes, churches, restaurants, and nightclubs. This enables the creation of a cultural description of local areas, allowing comparison with other areas—a step we perform in this study to identify cultural similarities. This differs from previous studies [1, 4, 9, 14], which tend to disregard the cultural component in their analyses.

Extending previous studies, the approaches proposed here to describing local culture rely on simple data from the Google Places API. One can provide an expressive cultural abstraction of any covered urban area, thanks to the mapping to Scenes Theory – see Section 2. Unlike studies that explored the cultural characteristics of regions using eating habits and user mobility, this study aims to derive such characteristics from the categories of venues present in a city. This allows us to evaluate whether our proposed approaches can adequately express key cultural aspects without relying on user actions, such as check-ins and evaluations.

We evaluate the approaches using data from 14 cities on different continents and all states of the United States. The results indicate that a simple approach, *Frequency*, can capture significant cultural differences satisfactorily. However, a more sophisticated approach, *Scenes*, can add extra semantic expressiveness in capturing cultural characteristics. This added expressiveness is evident in our outcomes and survey data comparison, indicating that *Scenes* better captures cultural nuances.

## 2 Cultural Signatures obtained from Google Places (GP)

### 2.1 Data From GP

GP is a location-based social network that allows users to discover and share information about local venues, geographic locations or points of interest, such as universities, cafes, bus stations, and parks. No type of location was disregarded. GP API provides geolocated venue data, resulting in one of the world's most accurate, up-to-date, and comprehensive venue models. In addition to latitude and longitude coordinates, venues are associated with at least one category designed to describe the venue type. In this study, we consider two datasets from GP, States and Cities, as described next.

The *Dataset States*, presented in [10, 22], includes business metadata (geographic info, category information, etc.) from GP up to September 2021 for all U.S. states. The dataset is composed of 4,963,111 unique venues and has 4,501 unique categories. The District of Columbia has the lowest number of distinct venues, totaling 11,003, while California has the highest count at 513,134 unique venues. We explore this dataset to study states focusing on geographic and category information.

For *Dataset Cities* we have collected data from a set of cities. GP API provides, by default, 141 unique categories. However, these categories do not provide the level of specificity necessary in this study. For example, the API assigns the category "restaurant" to all venues of this type, but it does not offer more specific categories related to cuisine, such as Italian or Japanese, which is necessary for this work. The optional "keyword" parameter is used in requests to the GP API aiming to overcome the limitation. The API documentation[3] guarantees valid results when inputs to this parameter are categories of venues, making it a convenient option for the desired purposes. The categories chosen to use in this parameter are those from the Yelp database due to the higher specificity, e.g., Yelp offers specific types of restaurants, such as Italian Restaurants. Yelp categories have a four-level hierarchical structure, making it suitable for our work to adopt only those at the last level. Some of them were excluded because they were not relevant to the purpose of the study, such as Provencal and Northeastern Brazilian, resulting in a total of 888 categories.

Using the proposed strategy, we have collected data from 14 cities, namely: Curitiba and Rio de Janeiro in Brazil; Toronto and Vancouver in Canada; Chicago and Los Angeles in the USA; Berlin and Frankfurt in Germany; Paris and Lyon in France; Seoul and Busan in South Korea; and Nairobi and Mombasa in Kenya. These cities are important in their respective countries and cover regions with different cultural characteristics. A publicly available tool[4] details the data acquisition process and clarifies the need for a balance between costs and data volume, which leads us to have a summarized set of venues. This tool aids in reproducing our study [5].

### 2.2 Urban Areas' Cultural Dimensions

Following research on local "scenescapes," we measure local scenes for the urban areas by aggregating the set of available venue categories in terms of qualitative meanings they express. To translate these concepts into measurements, for each venue category (e.g., restaurant, university, or bar), a team of trained coders has assigned a score of 1-5 on a set of 15 cultural dimensions $s_i \in S = \{s_1, s_2, ..., s_{15}\}$, such as transgression, tradition, local, authenticity, or glamour. Each area then receives a score for each of the 15 dimensions, calculated as a weighted average.

[3] https://developers.google.com/maps/documentation/places/web-service/overview.
[4] https://github.com/FerGubert/google_places_enricher.

Detailed descriptions of the theoretical meaning of each dimension can be found in [20].

### 2.3 Transfer Knowledge Procedure

The categories retrieved from GP need to be mapped to the appropriate set of 15 dimensions scores of the Scenes Theory. Without trained coders for our particular areas (States and Cities) we examine the Scenes' dimension scores of the Yelp categories presented in [19]. This knowledge is then adapted for use with GP/ categories, a transferring knowledge outlined in Figure 1. It illustrates an example for two different venues, each provided by a different dataset, venue A from *Dataset States* and venue B from *Dataset Cities* .

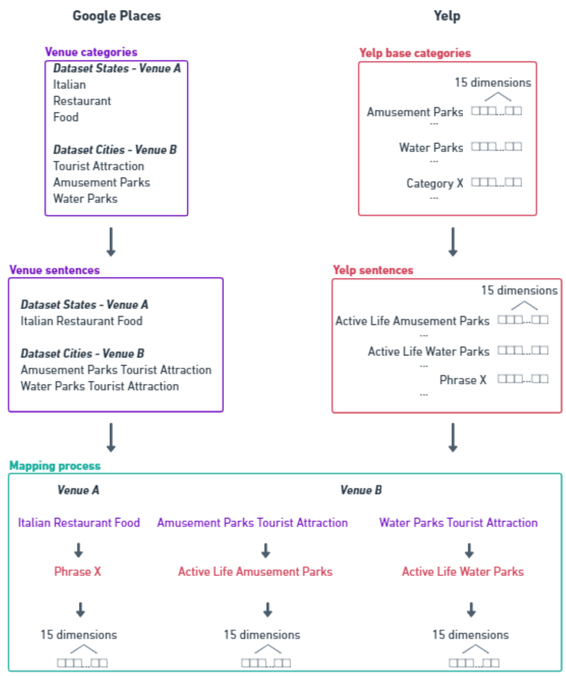


**Fig. 1.** Overview of mapping GP categories to the local cultural dimensions.

As depicted in Figure 1, for a better description of the venues, both the selected Yelp categories used in the requests and the broader categories made available by GP are used. To increase semantic capacity and mapping accuracy, sentences are created for each venue, following the procedures for each dataset.

In *Dataset States*, one sentence is created per venue, combining all associated categories. For example, if the venue has the categories "Italian", "Restaurant" and "Food", the sentence is: "Italian Restaurant Food". *Dataset Cities* on the other hand, lacks specific categories by default. Therefore, sentences include a requested Yelp category and all GP categories associated with that venue. For example, if a venue has “Amusement Parks” and “Water Parks” due to Yelp requests and “Tourist Attraction” as a default GP category, the sentences are: “Amusement Parks Tourist Attraction” and “Water Parks Tourist Attraction”.

Yelp categories are organized in a 4-level hierarchical structure. To expand semantic capacity, Yelp sentences are created using all hierarchical levels. In other words, for each category at the last level, the associated sentence returns to the first level. This is why “Active Life” was added to the Yelp sentences in Figure 1; these Yelp categories are immediately below its root category.

In possession of sentences describing the venue, the mapping process is carried out with SBERT, using the Sentence Transformers framework, in which several pre-trained models with a large and diverse dataset of more than 1 billion training pairs are made available and can be used to calculate embeddings from sentences and texts to more than 100 languages [12]. The cosine similarity compares the generated embeddings, and for each sentence related to the venues, the Yelp sentence with the highest score is retrieved. With this mapping, each venue is associated with one or more vectors (depending on the number of related sentences) containing the 15 dimensions of the Scenes Theory.

### 2.4 Cultural Signatures

We propose two approaches for creating cultural signatures, *Scenes-based approach*and *Frequency-based approach*.

For a particular urban area, the *Scenes-based approach*considers a vector $S_{area} = \{s_1^{area}, s_2^{area}, ..., s_{15}^{area}\}$, where $s_i^{area} = \frac{1}{\omega}\sum_{v=1}^{\omega}\left(\frac{1}{m}\sum_{\phi=1}^{m} S_i^{v,\phi}\right)$, with $\omega$ representing the number of unique venues in an urban area, $m$ is the number of categories a venue has, and $S_i^{v,\phi}$ is the $i$-th element of the vector of cultural dimensions for a certain venue $v$ and one of its category $\phi$; thus, $s_i^{area}$ represents the average score of all venues in the urban area for a specific cultural dimension, considering the average scores of all categories for each venue.

We also present an alternative approach, *Frequency*, aimed at creating cultural signatures that disregard Scenes information, using only location categories. This approach considers the frequency of the category in the area, i.e., for a particular urban area, a vector describes it by all the unique categories found in that area. For example, an area could be described by the categories [University, Restaurant, Coffee Shop, American Restaurant] and another by [Italian Restaurant, Wine Shop]. The frequency values are normalized per category.

*Frequency* helps answer the question: Are the existence and the number of certain types of venues in two different urban areas enough to explain their cultural differences?

## 3 Cultural Signatures Identify Culturally Similar Areas

### 3.1 Cities Worldwide

***Scenes* for *Dataset Cities*** First, we evaluate the results of the cultural signatures generated by the *Scenes-based approach*. We perform hierarchical clustering using Ward's linkage method and Euclidean distance, with the 15 dimensions of Scenes Theory as features. The results are represented in the dendrogram depicted at the top of Figure 2, where a division into six clusters is identified.

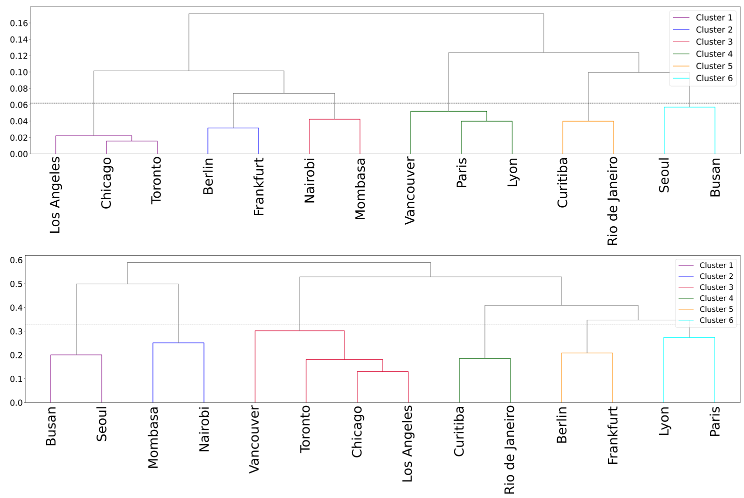


**Fig. 2.** Hierarchical clustering dendrogram of cities resulting from *Scenes* (top) and *Frequency* (bottom).

The result aligns with what is expected concerning the cultural characteristics of the areas studied. Most of the clusters coherently grouped cities from the same country - in general, countries have distinct cultural characteristics; the exceptions in this sense are clusters 1 and 4. In cluster 1, Toronto was grouped with Chicago and Los Angeles; note also that Los Angeles is the most dissimilar city in the grouping. The result of Chicago and Toronto being together and more similar makes sense, in that they are often considered to be culturally similar to one another, even compared to Los Angeles. Regarding cluster 4, Vancouver was grouped with Paris and Lyon. We found significant similarities between the most recurrent categories of French cities and Vancouver, such as "Art galleries," which could help explain this result. Although German cities (Berlin and Frankfurt) and French cities (Paris and Lyon) are on the same continent, they are quite distinct culturally, and so their location in separate clusters seems reasonable.

To facilitate a comparative analysis by contrasting the values of each cluster dimension with its corresponding overall average, we calculate the Z-Score, as shown in Figure 3. The Z-Score is the number of standard deviations concerning the average of what is being observed. This facilitates comparing clusters by extracting the characteristics that stand out in each, compared with a general overview, i.e., the centroid of clusters' centroids. For example, cluster 3, representing Kenya, has one of the lowest values for Tradition. In contrast, for cluster 4 with the cities Vancouver, Paris, and Lyon, this dimension represents one of the most important characteristics. Looking at cluster 1, composed of Chicago, Los Angeles and Toronto, we see that Tradition is not as predominant as in cluster 4. This highlights the potential to identify cultural signatures and provide an overview of geographic areas by extracting their key dimensions.

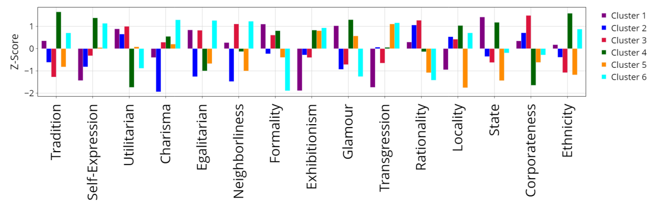


**Fig. 3.** Z-Score values of Scenes dimensions per cluster.

***Frequency* for *Dataset Cities*** For *Frequency*, we perform hierarchical clustering using the Complete linkage criteria and Cosine distance – the best combination tested. As depicted at the bottom of Figure 2, the results for *Frequency*, as with *Scenes*, align with what is expected when grouping cities of the same country. However, using *Frequency* differently, Chicago is more similar to Los Angeles, and Vancouver is more related to Toronto than to the French cities.

The results obtained demand reflection because although Toronto and Vancouver are in the same country, they are not necessarily similar in terms of immigration patterns, governance, geography, ecology, and cultural style. Toronto and Chicago, on the other hand, have much in common: they are both Great Lakes cities, with strong industrial heritages and are now in the midst of a post-industrial transformation. Hence, they are often compared as similar cases [7,13].

We can reveal specific characteristics of each cluster by extracting the five most distinct categories for each of them – we do that by calculating the distance of the category from its cluster centroid. After that, we calculate the Z-Score for the selected categories against the overall average. The result of this process is illustrated in Figure 4. Certain categories in some clusters stand out so notably that they not only significantly deviate from their overall average, but also emerge as the sole positive value compared to others. For example, in French cities, "municipality", in Brazilian cities, "hang gliding", and in Korean cities, "face painting" exhibits this distinct characteristic. Making a comparison with the Z-Score values illustrated in Figure 3, we can relate these specific findings depicted in Figure 4 to the aspects highlighted in Tradition for cluster 4 (predominantly French), Transgression for cluster 5 (Brazil) and Self-Expression and Charisma for cluster 6 (South Korea).

To analyze the clusters that differ between the *Scenes* and *Frequency* approaches, we examine the most evident characteristics in each. For *Scenes*, we focus on clusters 1 and 4, selecting the three most prominent dimensions in each and retrieving the most important sentences for those dimensions. For *Frequency*, we look at cluster 3 and identify the 50 most frequent categories. For example, Los Angeles, Chicago, and Toronto have "Business Consulting", "Libraries" and "Gastropubs" in common, whereas Vancouver, Paris, and Lyon are

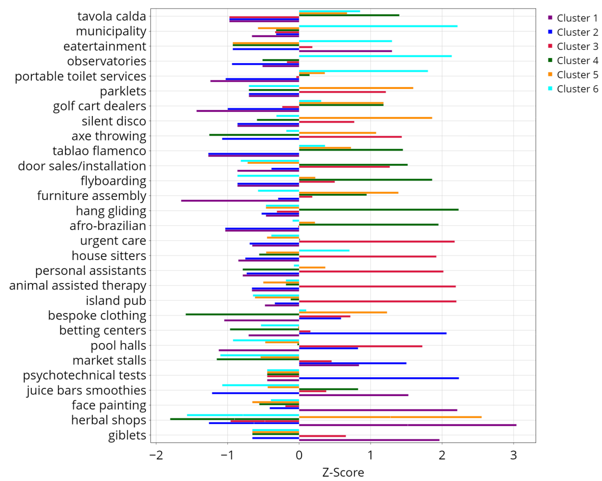


**Fig. 4.** Z-Score values for the most distinct categories per cluster (*Frequency*).

marked by "Antiques Book Store", "Art Gallery", "Comedy and Night Club" and gastronomic diversity, such as "Portuguese Bakery", "Spanish Meal Delivery", "Sushi Bars" and "Tapas Bars". In *Frequency*, many categories can be found that summarize these characteristics, such as "Gastropubs", "Art Installation", "Imported Food", "Meal Takeaway" and "Souvenir Shops". The result indicates that, unlike *Frequency*, through human knowledge in its dimensions, *Scenes* can detect subtle differences among categories with similar meanings.

### 3.2 All States in the USA

Using *Dataset States*, we apply the transfer knowledge methodology (Section 2.3) and create cultural signatures for all states in the country.

**Evaluating *Scenes-based approach* for *Dataset States*** To analyze cultural signatures in this dataset using *Scenes*, we also perform hierarchical clustering with 15 dimensions of the Scenes Theory as features, Ward linkage criteria, and Euclidean distance. By inspecting the dendrogram, we observe a tendency to group regions by geographic proximity. By mapping one of the clearest cuts in the dendrogram, we obtain Figure 5 (right). It shows that culturally similar regions, such as the US South, are grouped. These results reinforce the effectiveness of the proposed method in identifying culturally similar regions.

**Evaluating *Frequency-based approach* for *Dataset States*** For this case, we perform hierarchical clustering using the Ward linkage criterion and Euclidean distance. Other combinations were experimented with, but none proved superior. We observe difference between this approach and the results obtained with *Scenes*. Figure 5 (left) illustrates the mapped clusters provided by *Frequency-based approach*.

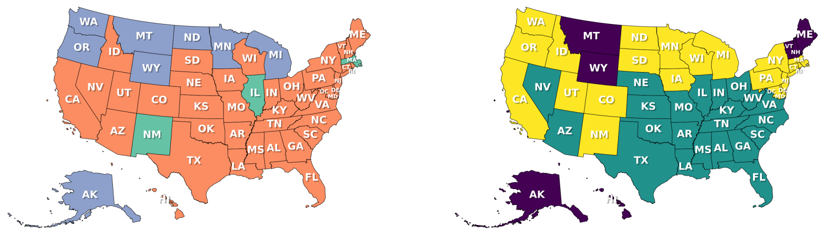


**Fig. 5.** Results of hierarchical clustering considering all states in the USA represented by *Frequency* (left) and *Scenes* (right).

It is not possible to detect clear patterns in the *Frequency* results, at least as clear as identified by *Scenes*, regardless of the number of clusters adopted. Surprisingly, Alaska and Maine are positioned within clusters larger than with *Scenes*. Alaska is situated among states such as Washington, Oregon, North Dakota, Minnesota, and Michigan. Maine is part of the largest cluster, which includes most of the remaining states. Thus, *Scenes* provides extra semantic expressiveness in smaller dimensions.

## 4 Comparing with Survey Data

There is no clear way to access the ground truth of our results. However, we explore in this work a source where we expect some correlation: the American Value Survey (AVS, access https://www.prri.org). The survey was conducted among a representative sample of 5,031 adults (age 18 and up) living in all 50 states in the United States, having a statistically valid representation of the USA population, including many minorities or hard-to-reach populations. Interviews were conducted online between September 16-29, 2021 and September 1-11, 2022. Additional details about the methodology can be found on the Ipsos website[5]. The survey questions include political aspects and basic beliefs. We represent these questions as features to describe states, where the values are the mean answers of all participants for each state. We exclude political questions and focus solely on basic beliefs[6].

To assess the relationship between the results of the AVS and our proposals (*Scenes* and *Frequency*), we use the Pearson correlation for the Euclidean distance between all pairs of states when describing them by AVS and our approaches. By doing that, we got a moderate correlation of 0.51 ($p < 10^{-4}$) for

[5] https://www.ipsos.com/en-us/solutions/public-affairs/knowledgepanel.

[6] The complete list of questions used can be found at: https://sites.google.com/view/neighbourhood-change.

*Scenes*. Using *Frequency*, on the other hand, resulted in a Pearson correlation of $-0.06$ ($p < 10^{-1}$) for the Euclidean distance between all pairs of states.

To better understand the correlation results individually we calculated the Euclidean distance of each state in comparison to all others, considering its descriptions using AVS and each of our proposals. Then, we calculate the Pearson correlation ($\rho$) of these values. For *Scenes*, $\rho \in [-0.221, 0.709]$ and approximately 75% of all states exhibit either a moderate or high correlation. Alaska is the only state with a negative correlation. By looking at the results for *Frequency*, with $\rho \in [-0.257, 0.149]$, it is clear that it shows a worse association with another source (AVS) regarding cultural beliefs.

## 5 Conclusion

In the present work, we examined data from Google Places (GP) and developed two methods to establish cultural signatures of urban areas. The proposals (*Frequency* and *Scenes*) were then assessed for their effectiveness in cities worldwide and all states in the United States. We obtained evidence that the proposed approaches, even a simple one based on frequency, could capture the cultural character of geographic areas. We gathered evidence based on a comparison with survey data that one of the approaches, based on the Scenes Theory, could capture better cultural nuances. Unlike other approaches that demand proxy data for users' preferences, e.g., user check-ins, our approach only demands simple data, i.e., categories of venues, which are easily obtainable in GP for almost any urban area. Hence, there is significant potential to utilize the proposed methodology for identifying cultural similarities between different locations. This could facilitate the development of numerous new services and applications, such as innovative location recommendation systems based on cultural criteria.

There are several ways to expand this work, such as expanding the dissimilarity analysis to both approaches, *Frequency* and *Scenes*, or testing the proposed methodology with other data sources. Since GP data is not free and acquiring a considerable amount can be costly, this could also allow for expanding the set of venues. Another possibility is to evaluate different levels of granularity, such as neighborhoods and countries.

## Acknowledgment

SocialNet project (process 2023/00148-0 of FAPESP) and CNPq (processes 313122/2023-7, 314603/2023-9 and 441444/2023-7).

## References

1. Arribas-Bel, D., Fleischmann, M.: Spatial signatures-understanding (urban) spaces through form and function. Habitat Int **128**, 102641 (2022)
2. Bancilhon, M., Constantinides, M., Bogucka, E.P., Aiello, L.M., Quercia, D.: Streetonomics: Quantifying culture using street names. Plos one **16**(6), e0252869 (2021)

3. de Brito, S.A., Baldykowski, A.L., Miczevski, S.A., Silva, T.H.: Cheers to untappd! preferences for beer reflect cultural differences around the world. In: Proc. of AMCIS'18. New Orleans, USA (2018)
4. Çelikten, E., Le Falher, G., Mathioudakis, M.: Modeling urban behavior by mining geotagged social data. IEEE Trans on Big Data **3**(2), 220–233 (2016)
5. Gubert, F., Silva, T.: Google places enricher: A tool that makes it easy to get and enrich google places api data. In: Proc. of WebMedia'22, Extended Proceedings. pp. 91–94. SBC, Curitiba, PR, Brasil (2022)
6. Hu, L., Li, Z., Ye, X.: Delineating and modeling activity space using geotagged social media data. Cartogr Geogr Inf Sci **47**(3), 277–288 (2020)
7. Kolpak, P., Wang, L.: Exploring the social and neighbourhood predictors of diabetes: a comparison between toronto and chicago. Prim. health care resear. & devel. **18**(3), 291–299 (2017)
8. Laufer, P., Wagner, C., Flöck, F., Strohmaier, M.: Mining cross-cultural relations from wikipedia: a study of 31 european food cultures. In: Proc. of the ACM WebSci'15. pp. 1–10. Oxford, UK (2015)
9. Le Falher, G., Gionis, A., Mathioudakis, M.: Where is the soho of rome? measures and algorithms for finding similar neighborhoods in cities. In: Proc. of the ICWSM'15. Oxford, UK (2015)
10. Li, J., Shang, J., McAuley, J.: UCTopic: Unsupervised Contrastive Learning for Phrase Representations and Topic Mining. In: Proc. of the ACL'22. pp. 6159–6169. ACL, Dublin, Ireland (2022)
11. Mehta, V., Mahato, B.: Measuring the robustness of neighbourhood business districts. J. of Urban Design **24**(1), 99–118 (2019)
12. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. In: Proc. of the EMNLP'19. ACL, Hong Kong, China (11 2019)
13. Robson, K., Anisef, P., Brown, R.S., Nagaoka, J.: A comparison of factors determining the transition to postsecondary education in toronto and chicago. Res. in Comp. Inter. Educ. **14**, 338–356 (2019)
14. Sen, R., Quercia, D.: World wide spatial capital. PloS one **13**(2), e0190346 (2018)
15. Senefonte, H., Frizzo, G., Delgado, M., Luders, R., Silver, D., Silva, T.: Regional Influences on Tourists Mobility Through the Lens of Social Sensing. In: Proc. of SocInfo'20. Pisa, Italy (2020)
16. Senefonte, H.C.M., Delgado, M.R., Lüders, R., Silva, T.H.: Predictour: Predicting mobility patterns of tourists based on social media user's profiles. IEEE Access **10**, 9257–9270 (2022)
17. Silva, T.H., de Melo, P.O.V., Almeida, J.M., Musolesi, M., Loureiro, A.A.: A large-scale study of cultural differences using urban data about eating and drinking preferences. Information Systems **72**, 95–116 (2017)
18. Silva, T.H., Silver, D.: Using graph neural networks to predict local culture. Environment and Planning B: Urban Analytics and City Science **0**(0), 12 (0)
19. Silver, D., Silva, T.H.: Complex causal structures of neighbourhood change: Evidence from a functionalist model and yelp data. Cities **133**, 104130 (2023)
20. Silver, D.A., Clark, T.N.: Scenescapes: How qualities of place shape social life. The University of Chicago (2016)
21. Skora, L.E., Senefonte, H.C., Delgado, M.R., Lüders, R., Silva, T.H.: Comparing global tourism flows measured by official census and social sensing. Online Soc Netw Media **29**, 100204 (2022)
22. Yan, A., He, Z., Li, J., Zhang, T., McAuley, J.: Personalized showcases: Generating multi-modal explanations for recommendations. In: Proc. of the SIGIR'23. p. 2251–2255. ACM, Taipei, Taiwan (2023)