Urban Perception Extraction from Texts Shared on Social Media: Framework and Applications

Frances Albert Santos[1], **Thiago Henrique Silva**[2], **Leandro Aparecido Villas**[1]

[1]Institute of Computing – University of Campinas (Unicamp)
Campinas – SP – Brazil

[2]Department of Informatics – Federal University of Technology of Paraná (UTFPR)
Curitiba – PR – Brazil

frances.santos.27@gmail.com, lvillas@unicamp.br, thiagoh@utfpr.edu.br

Abstract. *This thesis presents an automatic, generic framework for extracting urban perceptions from Location-Based Social Network (LBSN) data. The framework is organized into five key layers: Data Collection, Preprocessing and Embeddings, Model Training, Knowledge Extraction, and Applications. By leveraging deep learning techniques, including advanced sentence embedding methods, the framework captures both lexical and semantic nuances in textual data, thereby efficiently extracting user perceptions of urban environments. This approach eliminates the need for labor-intensive field surveys and manual data extraction, allowing scalable real-time analysis. We validated the framework by applying it to selected urban areas in Chicago, New York City, and London, demonstrating its effectiveness in uncovering valuable insights about urban perceptions. Furthermore, a comparative evaluation using a public dataset derived from volunteers' perceptions in a controlled experiment revealed a high level of agreement between the two sets of results. As a proof-of-concept, we introduce Real-Estate Urban Perceptions (REAL-UP), an innovative tool designed to enhance the real estate marketplace. REAL-UP provides interactive 2D maps that integrate traditional real-estate data (e.g., rent prices and property types) with enriched information on neighborhood emotions, sentiments, and brief narrative reviews generated by a Large Language Model (LLM) based on LBSN messages.*

Resumo. *Esta tese apresenta um framework automático e genérico para a extração de percepções urbanas a partir de dados de Redes Sociais Baseadas em Localização (LBSN). O framework é estruturado em cinco camadas principais: Coleta de Dados, Pré-processamento e Embeddings, Treinamento de Modelos, Extração de Conhecimento e Aplicações. Utilizando técnicas de deep learning, incluindo métodos avançados de incorporação de sentenças, o framework captura tanto nuances lexicais quanto semânticas nos dados textuais, permitindo a extração eficiente das percepções dos usuários sobre o ambiente urbano. Essa abordagem elimina a necessidade de pesquisas de campo intensivas e extração manual de dados, viabilizando análises escaláveis em tempo real. Validamos o framework aplicando-o a áreas urbanas selecionadas em Chicago, Nova York e Londres, demonstrando sua eficácia na revelação de insights valiosos sobre percepções urbanas. Além disso, uma avaliação comparativa utilizando um conjunto de dados público, derivado das percepções de voluntários* *em um experimento controlado, revelou um alto nível de concordância entre os dois conjuntos de resultados. Como prova de conceito, introduzimos o Real-Estate Urban Perceptions (REAL-UP), uma ferramenta inovadora projetada para aprimorar o mercado imobiliário. O REAL-UP oferece mapas interativos em 2D que integram dados imobiliários tradicionais (por exemplo, preços de aluguel e tipos de propriedade) com informações enriquecidas sobre emoções e sentimentos dos bairros, além de breves resenhas narrativas geradas por um Modelo de Linguagem de Grande Escala (LLM) com base em mensagens de LBSN.*

1. Introduction

Cities are not merely collections of buildings, streets, and residents – they are dynamic spaces where diverse experiences converge. The aesthetic quality of urban areas, crime rates, and ambient noise all shape how people perceive their surroundings. For example, locals may avoid tourist spots that attract visitors during their daily routines. Human perception is multifaceted, influenced by factors such as culture, age, and the way we interpret sensory impressions. In this context, the term "urban perception" refers to how individuals experience and interact with outdoor spaces, such as parks, streets, and plazas, areas that can influence behaviors related to health and safety.

Traditional methods for capturing urban perception, such as field surveys and sensory walks, are time-consuming and labor-intensive. Although crowdsourcing offers an alternative for gathering data, maintaining volunteer participation can be challenging. In contrast, Location-Based Social Networks (LBSNs) provide a rich source of urban perception data through their extensive user base, offering insights into aspects such as safety, traffic, and aesthetics. However, extracting meaningful information from the vast and varied content on LBSNs remains a complex task. Most studies focus on a single data format, typically text, using techniques ranging from simple keyword extraction, which can lead to false positives, to more sophisticated methods like topic modeling, both of which struggle with the brevity and density of social media texts.

To address these challenges, this study introduces a novel framework that leverages deep learning algorithms to extract urban perceptions from LBSN data. By employing sentence embeddings for robust text representations, our framework effectively captures a range of perceptions, including sentiments and emotions, without the need for extensive field surveys. As a practical application, we present Real-Estate Urban Perceptions (REAL-UP), a pioneering tool that enhances the real-estate marketplace by integrating urban perception insights. To the best of our knowledge, REAL-UP is the first tool to utilize users' urban perceptions to improve real estate market analysis.

1.1. Adherence to SBRC

This thesis focused on exploring Natural Language Processing (NLP) in texts from social media to extract urban perceptions. Thus, this thesis helps to leverage Urban Computing by providing a semantic understanding of areas that could enable new services, as well as help in understanding user behavior in urban spaces, which could be useful, for example, to network capacity planning and vehicular routing, among others. Topics addressed in this thesis are very related to many topics of interest in SBRC, such as: (i) Urban

Computing; (ii) Smart Cities; (iii) Social Networks; (iv) Data Mining and Analysis; and (v) Social Computing. Evidence of the strong adherence to SBRC, as described in the by-products of the work, a part of this thesis was presented as a short course in 2019 and as a full paper in 2018.

2. Methodology

2.1. Objective

The primary objective of this thesis is to extract and analyze urban perceptions from public social media content, thereby deepening our understanding of urban environments and enabling the development of innovative services and applications.

2.2. Framework

Our framework is composed of five major layers and leverages multiple online data sources to uncover latent urban perceptions, providing valuable insights into the intrinsic characteristics of urban areas that can enhance smart services and applications, as shown in Figure 1.

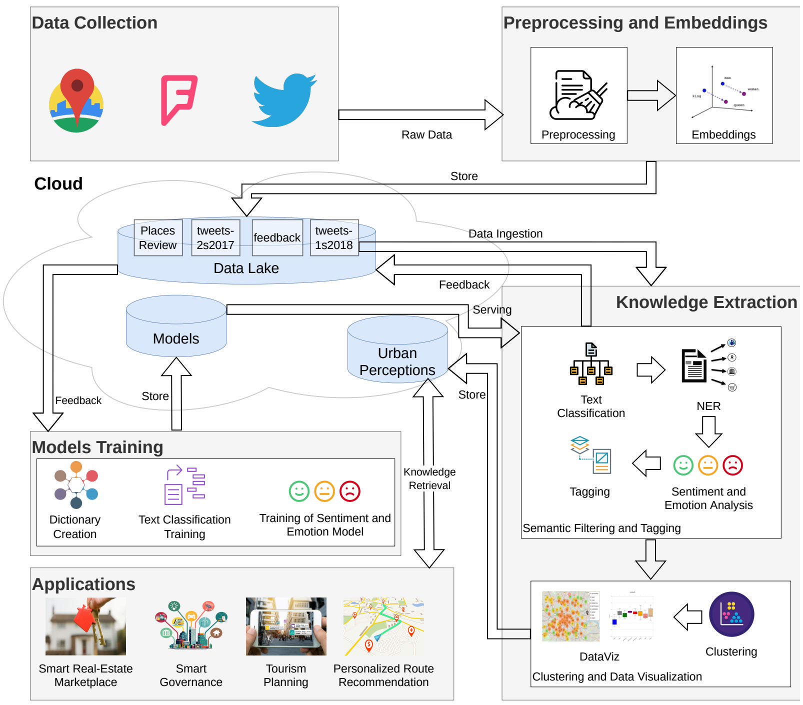

Figura 1. Urban Perception framework.

A key component of our framework is the creation of the UOP-dictionary, a lexicon that organizes the primary descriptive adjectives used by individuals to characterize their experiences in urban outdoor spaces. Unlike previous approaches that require an exact match between text and dictionary entries, our method assigns tags based on the similarity of text embeddings. This enables our system to capture relevant perceptions even when users employ vocabulary not explicitly listed in the UOP-dictionary. Another crucial element is the Semantic Filtering and Tagging module, which performs subject classification and assigns semantic tags to each tweet. These tags convey sentiment, emotion, recognized entities, and the corresponding UOP-dictionary category, thereby enriching our understanding of urban environments.

Clustering is particularly interesting in uncovering collective knowledge about some aspects. The reason for this is that individual opinions (or perceptions) might not reflect the real characteristics and/or situation of the place. Moreover, a person can share his/her perception regarding a place different from his/her actual position. Both cases can be considered noises, and using the right strategies, we can identify and remove the outliers. Finally, with the help of interactive maps built using a Python visualization library called Folium[1], urban perceptions are mapped into the cities according to the number of observed collective perceptions. This tool eases the process of studying the perceptions found.

Therefore, potential applications could emerge by leveraging the perceptual knowledge extracted through the proposed framework, enabling citizens and tourists to gain a deeper understanding of the hidden characteristics of urban spaces.

2.3. Experiments and Evaluation

A set of experiments to evaluate the extracted perceptions of different urban outdoor areas. We have used a Twitter dataset to demonstrate the potential of our framework to uncover the urban perceptions of outdoor spaces that emerge from LBSNs. Studying some urban areas from Chicago, New York City, and London, we demonstrate the framework's effectiveness in extracting valuable insights related to urban perceptions from LBSN data. We contrasted our results with Place Pulse 2.0, a controlled experiment expressing volunteers' perceptions in different urban outdoor areas [Dubey et al. 2016] (experiment authorized, using Amazon Mechanical Turk). Place Pulse 2.0 considers visual patches of images to identify the place's urban perception instead of natural language texts shared by individuals while visiting the area, which can potentially differ the people's perception since non-visual features can influence their perception. However, it is one of the few studies that mapped urban perception on a large scale, considering a significant number of categories and providing their data, giving us a piece of general knowledge about urban perception of outdoor areas. We observe that our approach yields results very similar to those shown by Place Pulse 2.0.

For example, considering the Chicago Downtown area, as we can see in Figure 2, it has a considerable concentration of distinct perceptions extracted by our framework, suggesting that most of them coexist in the neighborhood. Such a finding is not surprising because Downtown areas tend to present a wide variety of sounds, visual elements, and odors, among other characteristics, which can potentially cause distinct perceptions to people. Using the z-score, we conducted a temporal analysis to measure the persistence of perceptions identified over time. Such analysis is useful for ranking the perceptions in the neighborhoods, as well as to find if perceptions still are unchanged independent of the period. Thus, we can capture, for every neighborhood, the perceptions that stand out, even when the number of samples of perception categories is very unbalanced, which is our case. For instance, the categories *Neutral* and *Positive* are by far the most numerous, and the reason for that could be diverse, e.g., people tend to share more the "good" moments

[1]Folium: https://python-visualization.github.io/folium/. Last accessed January 3rd, 2023.

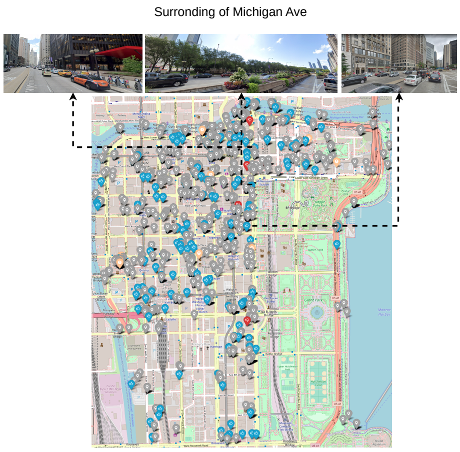

Figura 2. Perceptions map in Loop (Downtown), Chicago, IL. [Best in color]

in social media, or they only want to register this memory without sharing how it makes themselves feeling, in terms of sentiment.

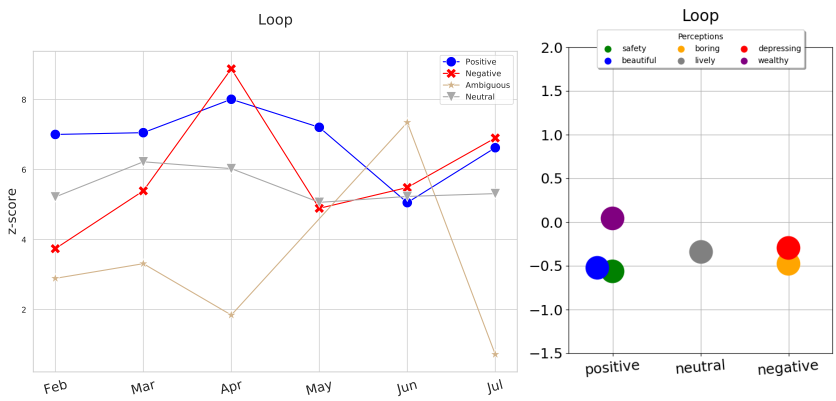

(a) Our approach **(b)** Place Pulse 2.0 [Dubey et al. 2016]

Figura 3. The perception strength of downtown Chicago. [Best in color]

Figure 3(a) shows the results for downtown Chicago. As we can see, several perceptions occurred in the Downtown area with high intensity (z-score around 4 or more) regardless of the month, where *Positive* and *Neutral* had low variation over time (less than one standard deviation in most cases). On the other hand, *Negative* and *Ambiguous* had a higher variation. In the first case, we can see two peaks, one in April and another in July, indicating that visitors might have some negative experiences when visiting the neighborhood, generating this class of perception. However, the most significant perceptions that tend to occur with high strength in this area are *Positive* and *Neutral*. Besides sentiments, we could identify the top three UOP-tags in Loop (also in terms of z-score), which are **unforgettable&moving** (z-score $= 8.152$), **crowded&public** (z-score $= 7.958$), and **multiple&single** (z-score $= 7.955$), respectively. According to Place Pulse data, Loop is mainly *wealthy*, but is also *lively* and *depressing*, with less intensity as shown in Figure 3(b). The similarity between the perception identified by our approach and this result is striking, where the most intense perception is also positive, followed by neutral and negative perceptions. For more results, refer to Chapter 5 of the thesis.

2.4. Application

The real estate marketplace has been providing websites and apps for people to consult detailed property information, such as price, parking spaces, description of rooms, photos, etc. However, this market still does not explore other relevant information about the property's surrounding area. To address this gap, we propose REAL-UP, an interactive tool designed to enrich real-estate marketplaces. In addition to the commonly provided data—e.g., rent price—REAL-UP integrates insights from urban computing to offer a deeper understanding of neighborhoods. By analyzing messages from Location-Based Social Networks (LBSNs), our tool extracts subjective perceptions of urban areas, such as safety, vibrancy, and overall sentiment, which could significantly enhance the property search process.

REAL-UP leverages urban perception models extracted using our framework to present subjective neighborhood information through interactive 2D maps, specifically highlighting perceived emotions and sentiments. Furthermore, by integrating a Large Language Model (LLM), REAL-UP generates meaningful neighborhood reviews based on crowdsourced knowledge from LBSN users. REAL-UP also uses data collected from the initiative Inside Airbnb[2] to obtain accurate information on properties. Inside Airbnb provides data and advocacy about Airbnb's impact on residential communities in several cities around the world, including those evaluated in this work: Chicago, IL; New York City, NY; and Greater London, London. Among the data available on this website, we used the file called "listings.csv," which has summary information (quarterly data for the last 12 months) and metrics for properties of each city, such as latitude, longitude, price, and room type (entire home/apt, private room, and shared room). We collected this file for three cities evaluated on January 25th, 2022.

This approach not only improves the user experience in real-estate searches but also aligns with research topics in Urban Computing and Smart Cities, which have been actively discussed at the SBRC [Krohling et al. 2024, Almeida et al. 2022]. Given the increasing relevance of human-centric urban analytics, REAL-UP contributes to emerging studies that explore how computational methods can support urban planning, mobility analysis, and smart decision-making.

As we can see in Figure 4(a), REAL-UP's homepage displays a short description and requires users to input two information: select one of the cities available, by default Chicago is selected; the period granularity, where "Last month" (default value) consider data from July 2018 to August 2018, "Last 3 months" consider data from May 2018 to August 2018, and "Whole period" consider data from January 2018 to August 2018. After providing this information and clicking on the button "Go," users are redirected to a new web page, where an interactive 2D map, built using the Folium library[3], is displayed. Also, a message informing the number of properties available for rent is displayed on top of the page, as shown in Figure 4(b) for Chicago, IL. As we can see, the neighborhoods' boundaries are highlighted with strong black lines. Small colorful circles represent the properties, where each color refers to the room type: red for shared rooms, yellow for

[2]http://insideairbnb.com/get-the-data.html.

[3]https://python-visualization.github.io/folium/latest/.

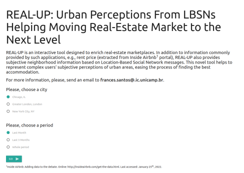

(a) Home-page.

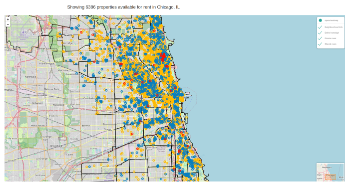

(b) Interactive 2D map.

(c) Property information.

(d) Neighborhood information.

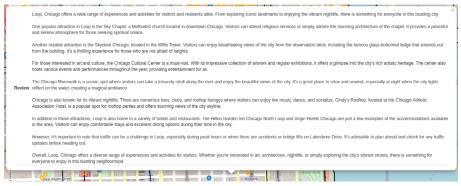

(e) Full review.

Figura 4. System interfaces. [Best in color]

the entire home/apartment, and blue for private rooms. Moreover, we can use the small map in the right-bottom corner to facilitate our localization in the whole city, which can be helpful when we increase the zoom. Also, we can apply filters to remove or add information on the map at any time by disabling/enabling them using the top-right menu.

To know more information about any property, we just need to point the mouse at the circle; then a pop-up window will open with the following information: *Location id*, a unique identifier of property in Inside Airbnb data; *Room type*; and, *Price*, as shown in Figures 4(c). Similarly, as Figures 4(d) shows, we can point the mouse at any neighborhood to obtain the following information regarding it: *Neighborhood*; *Emotions*; *Sentiment*; and, *Review*. As we can see, all sentiments and emotions occur in Loop, Chicago, and are displayed in an orderly manner, from the most relevant to the least relevant. Then, as shown in Figure 4(d), we can click on the left mouse button to access the full review of the neighborhood, as shown in Figure 4(e).

REAL-UP's approach extends beyond real estate applications. By providing real-time urban sentiment analysis, our tool can support municipal decision-makers, urban planners, and researchers interested in understanding social dynamics in cities. For instance, city officials could leverage the tool to identify areas where residents feel unsafe, detect trends in urban vitality, or assess the impact of policy interventions. Additionally, integrating REAL-UP with smart city initiatives could enhance transportation planning, tourism recommendations, and even crisis management. Future developments include expanding the tool's dataset to incorporate multi-modal urban data, such as traffic patterns, environmental conditions, and infrastructure quality, further strengthening its potential as a decision-support system for urban intelligence.

3. Main Results

In this thesis, we presented a novel framework to support the learning and mapping of the perception of urban outdoor areas from an extensive collection of noisy data expressing users' opinions in LBSNs. Our results suggest that it is possible to identify perceptions reflected in urban areas in a scalable way. This is useful for supporting mechanisms to help people better understand the semantics in different city regions. To better discuss our results, we answer the research questions that have guided the development of this thesis:

- **Research Question 1:** *How can we extract urban perceptions from natural language texts shared by people on social media?*

Initially, we explored a vast amount of social media data to discover which subjects we could find on it. Based on the well-known unsupervised clustering algorithm K-means, we could find 17 main subjects existing in the Twitter data, among them the "Urban Perception" subject. Then, we proposed a framework composed of 5 major layers (Data Collection, Preprocessing and Embeddings, Models Training, Knowledge Extraction, and Applications) to support the urban perception extraction from natural language texts. Our framework has three major advantages:

- **Scalability:** our framework eliminates the need for time-consuming traditional methods like sensory walks and does not rely on specific social media features (e.g., hashtags, emoticons). It can integrate various data sources, such as IoT devices and government data, as long as they provide natural language text and spatiotemporal information.
- **Flexibility:** our framework is adaptable to any English-speaking city using only public LBSN content. It demonstrates generalization by training models with Chicago data while successfully applying them to New York City and London, showing its ability to analyze unseen urban areas without modification.
- **Dynamism:** since language and urban perceptions evolve due to global events (e.g., COVID-19), the framework includes a feedback mechanism to identify misclassified text messages. This ensures future model updates can adapt by discovering new topics or refining existing classifications.

- **Research Question 2:** *What is the level of agreement among extracted perceptions with respect to a "ground truth"?*

It is very hard, if not impossible, to uncover the urban perceptions that correspond accurately to the ground truth. For this reason, our work, aligned with some others, focuses on finding general urban perception, which can better define the latent aspects of urban areas. To validate our results, we conducted a comparative analysis based on Place Pulse 2.0 data, which has volunteers' opinions about urban areas expressed in a controlled experiment. We observe that both results yield a very similar level of agreement.

- **Research Question 3:** *How can we combine multiple layers of perception to obtain unified knowledge regarding urban areas?*

Toward this goal, our framework combines sentiment and emotion analyses with the UOP-dictionary, enabling us to extract multiple aspects of perceptions of urban areas. For instance, in our results, we could find that downtown areas favor a mix of sentiments to occur, besides the UOP-tag **crowded&public** being very common in such places. Despite this finding being straightforward, our results also revealed less obvious insights, as, for example, citizens and visitors in the Streeterville neighborhood enjoy maritime activities, once the UOP-tags **maritime&titanic** and **great&good** are outstanding in this area. The major advantage of our approach to give a tag from UOP-dictionary to text messages (in this case, tweets), compared to others, is that such messages don't need to have the same adjectives that comprise the UOP-dictionary, which is a clear disadvantage of previous works using dictionaries to extract and/or label text messages (e.g., [Aiello et al. 2016, Quercia et al. 2015]). Instead of that, we compare the similarity between embeddings to define the proper tag. In this way, our approach is more robust, since people might choose words that do not belong to UOP-dictionary vocabulary, but still, our approach will return relevant results. With this, we presented evidence in Chapter 5 of the thesis that our framework can extract diverse knowledge about urban areas derived from social media data.

- **Research Question 4:** *How can urban perception be exploited to leverage new services and applications?*

To show the benefits of a better understanding of urban perceptions, we developed a demo that aims to enhance the real-estate marketplace as a proof-of-concept for our work. Our demo provides rich knowledge about urban areas in the form of interactive 2D maps, more specifically, the emotion and sentiment perceived and UOP tag from UOP-dictionary, for every city's neighborhood, in addition to information commonly provided by such applications, as rent price, property type, and so on. Thus, we can illustrate the potential of our framework to help people in the decision-making process in real-world applications very present in the lives of millions of people worldwide, such as Airbnb, Booking.com, Trivago, and Yelp, to name a few.

4. Main contributions

In total, **nine** peer-reviewed publications were produced to share the results obtained from this work over the doctorate program in relevant conferences and a prestigious journal in the Computer Science field, in particular, on tracks such as social network analysis and mining. In specific, one journal article [Santos et al. 2020] (SNAM - A2), six conference papers [Santos et al. 2024, Santos et al. 2018a, Santos et al. 2018b, Santos et al. 2018c, Santos et al. 2017a, Santos et al. 2017b] (WWW - A1; two works in ICC - A1; WI-IAT - A3; SBRC - A4; WGRS-SBRC – Best Paper Award – B4), and two book chapters [Santos et al. 2022, Rodrigues et al. 2019] (WebMedia - A4; SBRC - A4). Other **two** works were developed in collaboration with our research group [Ladeira et al. 2022, Santos et al. 2016] (VTC - A1; ICSC - A3). Besides the publications, parts of this thesis were conducted in collaboration with Prof. Dr. Azzedine Boukerche at SITE (School of Information Technology and Engineering), University of Ottawa, Canada, as part of the Sandwich Doctorate Program (PDSE/CAPES), and with Prof. Dr. Richard Pazzi at Ontario Tech University, Canada, as visiting research financial supported by Global Affairs Canada, via Emerging Leader in the Americas Program (ELAP).

Acknowledgment. This thesis was partly supported by the project SocialNet (grant 2023/00148-0 from Sao Paulo Research Foundation - FAPESP), CNPq (process 314603/2023-9 and 441444/2023-7), and CAPES (process 88881.132016/2016-01).

Referências

Aiello, L. M., Schifanella, R., Quercia, D., and Aletta, F. (2016). Chatty maps: constructing sound maps of urban areas from social media data. *Open Science*, 3(3):150690.

Almeida, V. G., Silva, T. R., and Silva, F. A. (2022). Se for, vá na paz: Construindo rotas seguras para veículos coletivos urbanos. In *Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos (SBRC)*, pages 140–153. SBC.

Dubey, A., Naik, N., Parikh, D., Raskar, R., and Hidalgo, C. A. (2016). Deep learning the city: Quantifying urban perception at a global scale. In *Proc. of ECCV*, pages 196–212. Springer.

Krohling, B., Comarela, G., and Mota, V. F. (2024). Uma arquitetura baseada em redes neurais recorrentes para prediçao de trajetórias veiculares em ambientes urbanos. In *Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos (SBRC)*, pages 379–392. SBC.

Ladeira, L. Z., Santos, F., Cléopas, L., Buteneers, P., and Villas, L. (2022). Neo-nda: Neo natural language data augmentation. In *Proc. of ICSC*, pages 99–102. IEEE.

Quercia, D., Schifanella, R., Aiello, L. M., Kate, M., et al. (2015). Smelly maps: The digital life of urban smellscapes. In *Proc. of ICWSM*, pages 327–336, Oxford, UK. AAAI Press.

Rodrigues, D. O., Santos, F. A., Rocha Filho, G. P., Akabane, A. T., Cabral, R., Immich, R., Junior, W. L., Cunha, F. D., Guidoni, D. L., Silva, T. H., et al. (2019). Computação urbana da teoria à prática: Fundamentos, aplicações e desafios. *JAI-SBC*.

Santos, F. A., Akabane, A. T., Yokoyama, R. S., Loureiro, A. A., and Villas, L. A. (2016). A roadside unit-based localization scheme to improve positioning for vehicular networks. In *Proc of VTC-Fall*, pages 1–5. IEEE.

Santos, F. A., Kobellarz, J. K., de Souza, F. R., Villas, L. A., and Silva, T. H. (2022). Processamento de linguagem natural em textos de mídias sociais: Fundamentos, ferramentas e aplicações. *Sociedade Brasileira de Computação*.

Santos, F. A., Rodrigues, D. O., Silva, T. H., Loureiro, A. A., Pazzi, R. W., and Villas, L. A. (2018a). Context-aware vehicle route recommendation platform: Exploring open and crowdsourced data. In *Proc. of ICC*, pages 1–7. IEEE.

Santos, F. A., Rodrigues, D. O., Silva, T. H., Loureiro, A. A., and Villas, L. A. (2017a). Rotas veiculares cientes de contexto: Arcabouço e aná lise usando dados oficiais e sensoriados por usuários sobre crimes. In *Proc. of WGRS*. SBC.

Santos, F. A., Silva, T. H., Braun, T., Loureiro, A. A., and Villas, L. A. (2017b). Towards a sustainable people-centric sensing. In *Prof. of ICC*, pages 1–6. IEEE.

Santos, F. A., Silva, T. H., Loureiro, A. A., Boukerche, A., and Villas, L. A. (2018b). Identificação da reputação de áreas urbanas externas com dados de mídias sociais. In *Proc. of SBRC*, pages 810–823. SBC.

Santos, F. A., Silva, T. H., Loureiro, A. A., and Villas, L. A. (2020). Automatic extraction of urban outdoor perception from geolocated free texts. *Social Network Analysis and Mining*, 10:1–23.

Santos, F. A., Silva, T. H., Loureiro, A. A. F., and Villas, L. A. (2018c). Uncovering the perception of urban outdoor areas expressed in social media. In *2018 IEEE/WIC/ACM International Conference on Web Intelligence (WI)*, pages 120–127. IEEE.

Santos, F. A., Silva, T. H., and Villas, L. A. (2024). Real-up: Urban perceptions from lbsns helping moving real-estate market to the next level. In *Companion Proceedings of the ACM on Web Conference 2024*, pages 1071–1074.