

Urban Data for Everybody: Visual Comparison & Multi-Attribute Ranking for Casual Users

Joshua Pola; TU Wien and Halo GmbH; Vienna, Austria
Krešimir Matković; VRVis Research Center, Vienna, Austria

Abstract

Urban data is increasingly available to the public, but making it accessible and useful for non-expert users remains challenging. Interactive visualization offers a powerful means to explore such data, particularly for tasks like comparing districts within one or between two cities, a common scenario for students relocating, tourists planning a trip, or citizens considering moving. Due to the varied interests and priorities of individuals, developing a universally applicable solution is demanding.

In this work, we propose an approach that enables visual, weighted comparison of city attributes. We first abstract analysis tasks and derive corresponding design requirements, and then present a coordinated multiple views system that allows users to compare different weighted attributes across districts of two cities. The system integrates linked views that support selection, weighting, ranking, and detailed exploration of urban data.

We demonstrate the usefulness of our approach through a comparison of Vienna, Austria, and Berlin, Germany, and validate it with a pilot study, which received very positive feedback. Our results indicate that the proposed approach supports casual users in exploring urban data effectively while allowing flexible, personalized comparisons.

Introduction

Today's cities generate data at an unprecedented pace, reflecting the complex dynamics of urban life. This wealth of information is valuable not only for domain experts—such as planners, policymakers, and service providers—but increasingly also for laypersons, journalists, and citizen scientists. Visualization has long been established as a powerful tool for data exploration, especially in contexts where problems are not clearly defined and insights must emerge through interactive engagement with heterogeneous datasets. The diversity of urban data, coupled with the differing needs and motivations of its users, makes cities a prime case for visualization-based approaches.

Recent advancements in urban visual analytics have made significant progress in supporting expert users across different domain problems, as highlighted by Deng et al. [1]. These solutions provide city planners, policymakers, businesses, and public service providers with tailored tools for analyzing urban data more efficiently and effectively. However, they often remain domain-specific, highly specialized, and predominantly designed for experts. Everyday citizens, who are often non-experts and not necessarily tech-savvy, are largely excluded from these innovations. Their access to urban information is typically limited to static resources such as online city plans (e.g., Vienna [2], Berlin [3], London [4]), or statistical platforms such as Statistik Berlin-Brandenburg [5] and Statistik Austria [6]. While such sources

allow users to look up data, they provide little support for deeper exploration, discovery, and analysis.

The visualization research community increasingly acknowledges the importance of empowering the general public to engage with complex datasets [7]. Yet, unlike experts whose problems and goals are often well-defined, non-expert users approach urban data with diverse and sometimes loosely defined motivation. A one-size-fits-all solution is therefore unlikely to succeed: aspects of the data that are critical for some citizens may be irrelevant to others. Despite these challenges, enabling non-experts to analyze and interact with urban data holds significant value, not only for fostering data literacy and civic participation, but also for addressing everyday concerns that lie outside the scope of domain-specific expert tools [8].

To address the gap in supporting casual, non-expert users, we present a system designed specifically for casual exploration of urban data and for comparing districts across two cities. The system enables users to specify their own priorities—or explore different preferences interactively—when comparing a familiar city with an unfamiliar one, or contrasting two cities regardless of prior knowledge. By supporting such flexible and user-driven comparisons, the system empowers citizens to gain meaningful insights into urban environments in a way that is both accessible and engaging.

Based on prior research on casual, non-expert user visualization [9, 10, 11], we formulate five overarching goals that guided the design of our system:

- **G1: Provide conceptual orientation.** Offer overview and orientation.
- **G2: Support exploratory search.** Since casual users often do not know what they are looking for, enabling independent exploration is crucial.
- **G3: Increase user experience.** Create an inviting and engaging experience for users.
- **G4: Keep cognitive load low.** Provide a general overview while allowing users to adjust detail granularity as needed.
- **G5: Enable intuitive modes of interaction.** Interaction techniques should feel natural and remain simple.

To achieve this, we first identified analysis tasks and requirements specific to casual users, with a focus on individual user preference and data comparison on the district level. Based on these requirements, we designed a visual exploration and comparison dashboard that allows non-expert users to interactively explore urban data. Finally, we conducted a pilot study that provided very positive feedback, indicating that our approach supports casual users in making sense of urban environments and comparing

districts within and among cities effectively. Therefore, the contributions of this paper can be summarized as:

1. a task analysis and requirements identification for casual user urban data visualization,
2. the design of a visual exploration and comparison dashboard including a multi-attribute ranking approach suitable for casual users, and
3. an initial pilot evaluation demonstrating the feasibility and value of the approach.

From this point onward, we will address non-expert users with no in-depth domain knowledge as casual users.

Related Work

Our approach touches on several complementary areas of research, including visualization for casual users, urban data dashboards, multi-attribute ranking, comparative visualization, and coordinated multiple views (CMV). In the following, we briefly review prior research in each of these areas to situate our work within the existing literature.

Tajja and Martin [8, 12] analyze **user requirements for urban data platforms** targeting non-technical citizens (casual users) and present the implementation and evaluation of a prototype to validate their approach. Their work focuses on design patterns and a general unified approach for urban data platforms for the general public. Our approach can be understood as a sub-area of their research, as we focus specifically on visualization, interaction, and attribute ranking techniques that can also be integrated into urban data platforms.

BarcelonaNow [13] is an interactive dashboard for exploring city-related data. Marras et al. provide not only the dashboard itself but also a broader framework for collecting, processing, exposing, and visualizing heterogeneous data, including both streaming and non-streaming sources. The dashboard successfully enabled “citizens with little or no expertise to create their custom visualizations by leveraging the provided functionalities.” Our approach differs from BarcelonaNow in that we omit point of interest (POI) data and instead focus on visualizing data attributes per district. This abstraction allows for a simpler overview on a city map, eliminating the need for zooming or filtering to access district-level insights.

Fundamental to our approach for multi-attribute ranking and comparison is **LineUp** [14]. In their paper, *LineUp: Visual Analysis of Multi-Attribute Rankings*, Gratzl et al. [14] introduce a scalable visualization technique that allows users to explore data by combining attributes and refining parameters—essentially letting users select which attributes are relevant and how they are weighted in the ranking calculation. Building on this approach, our system allows users to choose attributes, assign custom weights to each, and see the ranking results in real time. To maintain a simple and accessible interface suitable for casual users, we omit most of LineUp’s more complex elements or provide simplified alternatives, ensuring the dashboard meets the requirements of non-expert user visualization.

For the active research area of **comparative visualization**, Gleicher et al. [15] introduce a taxonomy for visual comparison with three categories: juxtaposition (where objects are placed separately), superposition (where objects are placed overlaying within the same space), and explicit encoding (in the context of

the relationship between objects). Placing objects in Juxtaposition, for example, is also applied by Munzner et al. [16] when comparing large phylogenetic trees and Tufte [17] refers to it as small multiples. Taking these approaches into account, our system uses side-by-side, district-level visualizations combined with interactive multi-attribute ranking to support users in efficiently comparing multiple cities, enabling a clear and intuitive overview even for casual users.

Coordinated multiple views and dashboards are a common approach to support complex exploration tasks. Systems like *Improvise* [18] and *VisBricks* [19] allow users to explore multiple linked views simultaneously, making it easier to compare data across dimensions. Such coordinated dashboards are particularly relevant in urban data exploration, as they allow citizens to interactively combine and prioritize attributes across spatial and non-spatial visualizations.

Our approach is inspired by BarcelonaNow’s accessible dashboard framework and LineUp’s multi-attribute ranking capabilities. We use the coordinated multiple views principles to provide a user-friendly system for casual urban data exploration and comparison based on individual user preferences.

Tasks and Design Requirements

In this section, we discuss the identified tasks and the corresponding design requirements for the effective exploration and comparison of urban data by casual users. The following tasks summarize the main goals casual users may pursue when exploring and comparing urban data:

Tasks

- **T1: City overview.** Gain an overview of the districts of a city to learn about its structure.
- **T2: Comparative city overview.** Gain an overview of districts of two cities to learn about their respective structures.
- **T3: Identify interesting districts.** Find districts of potential interest based on user-defined preferences to guide further exploration.
- **T4: District insight.** Explore data of a single district for purposes such as: a citizen scientist learning more about urban characteristics, a tourist preparing for a visit, a journalist gathering context for a story.
- **T5: Comparative district insight.** Explore and compare data of two districts to: learn about districts in relation to each other (citizen scientist), choose between districts to visit or stay in (tourist), gather additional context (journalist), find a similar district to move to (citizen), compare cost of living between current and future residence (student), identify a district similar to a district in another city (tourist).
- **T6: Drill-down analysis.** Analyze underlying data of a district to gain a deeper understanding at a fine-grained level.
- **T7: Comparative drill-down analysis.** Analyze detailed underlying data of two districts to gain insights into similarities and differences at a fine-grained level.

Design Requirements

Based on the identified tasks, we define design requirements for casual users. Each requirement is motivated by the corresponding tasks and supports one or more overarching goals.

- **R1: Show City Structure.** Display spatial features of a single city to provide an overview and support orientation. (Task: T1; Goal: G1 Orientation)
- **R2: Compare Cities.** Show spatial features of two cities side by side to support comparative exploration. (Task: T2; Goals: G1 Orientation, G2 Exploratory search)
- **R3: District Information.** Provide accessible data for a single district, allowing casual users to explore characteristics in context. (Task: T4; Goals: G1 Orientation, G2 Exploratory search)
- **R4: Detailed Drill-down.** Allow users to drill down into additional detailed district-level data while not overwhelming them. (Task: T6; Goals: G2 Exploratory search, G4 Low cognitive load)
- **R5: District Comparison.** Present data of multiple districts on the same view for direct comparison. (Task: T5; Goals: G2 Exploratory search, G5 Intuitive interaction)
- **R6: Comparative Drill-down.** Enable detailed comparison of two districts at a fine-grained level. (Task: T7; Goals: G2 Exploratory search, G4 Low cognitive load)
- **R7: Attribute Filtering.** Allow users to filter attributes according to their preferences to identify districts which are interesting to them. (Task: T3; Goals: G2 Exploratory search, G5 Intuitive interaction)
- **R8: Attribute Weighting and Ranking.** Allow users to assign custom weights to attributes in addition to filtering, compute resulting rankings, and display them in real time. (Task: T3; Goals: G2 Exploratory search, G4 Low cognitive load, G5 Intuitive interaction)
- **R9: Reproducibility.** Enable users to save the current state of the dashboard and reload it later, ensuring continuity of exploration. (Tasks: T3, T5; Goal: G3 User experience)

Data Collection and Pre-processing

Collecting data for urban data comparison is challenging because cities provide different types of data, especially at the district and more detailed level. Data might be unavailable on the district level, sub-categories definition (of, e.g., different types of green area) and data classification rules (what exactly is counted as a car within one district) might be different as well. In addition, districts can differ in size and population within and across cities, making direct comparison difficult. To enable meaningful comparison, it is necessary to unify the data in a structured way.

In order to guide data collection towards a more diverse and heterogeneous data set, we define five data categories in correspondence with what may be of interest to our users: Demographics, Environment, Livability, Tourism, and Traffic. This categorization is not intended as a universal scheme, but to exemplify our approach given the high amount of available data for the cities we consider. For the current study and evaluation, we use data from Vienna (Austria) and Berlin (Germany). Although we examined several other cities, differences in the available data (especially on the detailed level) were often too large to allow meaningful comparison. To tackle modern challenges regarding urban areas we argue it is necessary that data becomes more unified, available and is reported in more detail in the future.

Vienna and Berlin provide sufficiently similar data, resulting in a total of 13 attributes across the five main categories for our initial evaluation. The number of districts differs between the

cities (Vienna: 23; Berlin: 12). The attributes used in this study are as follows:

- **Demographics**
 - Population Density [people per km²]
 - Share of residents without national citizenship [% of district population]
- **Environment**
 - Green Area [% of district area]
 - Trees in Streets [trees per ha]
 - Dogs [per 1,000 district capita]
- **Livability**
 - Average Living Space [m² per apartment]
 - Cost of Buying Apartments [€ / m²]
 - Average Rent [€ / m²]
- **Tourism**
 - Overnight Stays [# of days per 1,000 district capita]
 - Average Length of Stay [# of days]
- **Traffic**
 - Accident Victims [per 100,000 district capita]
 - Street Space [% of district area]
 - Cars [per 1,000 district capita]

In correspondence with our requirements R4 and R6 we also collect detailed level data (e.g., age distribution regarding the population density, types of tree planted by the city among the street trees, average living space per resident in addition to the average apartment size) wherever possible. Keep in mind that differences within the provided detail level data result in further pre-processing steps where we have to unify deviating categories.

Also, to allow meaningful comparison between districts of different sizes and populations, we normalized the data according to the units of each attribute. For example, population density is expressed in people per km², the number of dogs is scaled per 1,000 district inhabitants, trees are expressed per hectare, and traffic accidents are scaled per 100,000 inhabitants. Financial and living space data (rent, purchase cost, average apartment size) are reported per m² to ensure comparability. These normalization steps make the data directly comparable across districts and cities despite differences in population and area.

Multi-Attribute Ranking Implementation

As our ranking approach is heavily influenced by the ranking implementation of **LineUp** by Gratzl et al. [14] we use a similar formal notation to make the comparison of our approaches easier.

Given are a set of items $x_i \in X = \{x_1, \dots, x_m\}$, which in our case corresponds to the districts of *both* cities, and a heterogeneous set of attribute vectors $\mathbf{a}_j \in A = \{\mathbf{a}_1, \dots, \mathbf{a}_n\}$, where each vector contains the values of an attribute for all districts.

Because of their heterogeneity the attributes need to be mapped to a common scale. The mapping function itself is defined by $m_j(\mathbf{a}_j) = \mathbf{a}' \in A'$ with $m_j : a_{j_i} \rightarrow [m_{j_{\min}}, m_{j_{\max}}]$ for all attributes with $0 \leq m_{j_{\min}} \leq m_{j_{\max}} \leq 1$. Like Gratzl et al. [14], we refer to the mapped attribute values a'_{j_i} as *attribute scores* and to A' as the *mapped attributes*. Now we simplify this approach by




Figure 1. A screenshot of the implemented dashboard. The data selection view (A) is used to select attributes of interest, set weights, and keep track of findings, the city map views (B) show each of the cities and the districts. Colors encode single-attribute values while three gold, silver, and bronze icons indicate the top ranked districts. The multi-attribute ranking comparison view (C) displays the result of the rank computation for each district based on the selected weights and mapping function within the data selection view. The data comparison view (D) shows all selected attributes including some sub-categories (if available) of two districts for further comparison. These two districts can be selected in the city map views.

not allowing our users to define values for m_{jmin} and m_{jmax} . By assigning $m_{jmin} = 0$ and $m_{jmax} = 1$ we end up with all attribute scores having values between 0 and 1 based on the lowest and highest district value. This leads to our new definition for the mapping function with $m_j : a_{ji} \rightarrow [0, 1]$.

To rank our districts (or items) we need to assign a score $s_i \in S \subset \mathbb{R}_0^+$ to each district x_i . Given our prior definition for attribute scores a'_{ji} , a simple function for the item score could look like

$$s_i = s(x_i) = \sum_{a'_{ji}} a'_{ji}$$

basically the score for a district x_i is the sum of all attributes scores a'_{ji} of the district. Our users have two possibilities to influence how attribute scores are considered during ranking. Now to allow our users to influence how much each mapped attribute in A' is considered during ranking we introduce a weight $w_{a'_{ji}}$. Also, we multiply the weighted attribute scores by 100 for readability in case we display the score to our users. This leads to our new scoring function

$$s_i = s(x_i) = \sum_{a'_{ji}} w_{a'_{ji}} a'_{ji} * 100 \mid 0 \leq w_{a'_{ji}} \leq 1$$

Compared to LineUp, the sum of all weights $\sum w_{a'_{ji}}$ does not have to be equal to 1. Weights are not normalized because we want our users to encode preference in a way that is directly tied to how the

scores are calculated. Users can encode as many attributes as very important ($w_{a'_{ji}} = 1$) without having to worry about the underlying proportions between individual weights.

The rank $r_i \in \mathbb{N}^+$ of a district x_i is the result of the item's score s_i . Ranks are calculated for each city individually. Since we allow the sum of all weights $\sum w_{a'_{ji}}$ to be greater than one and multiply each attribute score by 100 the $max(S) = 100 * n$.

In addition to the weight for each attribute the users can influence how an attribute score a'_{ji} is considered during ranking is by inverting the attribute score by choosing a different mapping function. In consequence, a lower attribute value results in a higher score. We end up with two mapping functions, one where the highest value results in the highest score

$$a'_{ji} = m_j(a_{ji}) = \begin{cases} \frac{a_{jmax} - a_{jmin}}{a_{jmax} - a_{jmin}} & \mid a_{jmax} > a_{jmin} \\ 1 & \mid a_{jmax} = a_{jmin} \end{cases}$$

and one where the lowest value results in the highest score

$$a'_{ji} = m_j(a_{ji}) = \begin{cases} 1 - \frac{a_{jmax} - a_{jmin}}{a_{jmax} - a_{jmin}} & \mid a_{jmax} > a_{jmin} \\ 0 & \mid a_{jmax} = a_{jmin} \end{cases}$$

As whether more or less of any given attribute is better is very subjective, this mapping function allows the users to decide this individually for each attribute score computation.




Figure 2. City map view for Vienna, with a different attribute selected than in Figure 1 for color coding, to facilitate comparison of districts based on different attributes.

Visualization and Interaction Design

In this section, we present the visual mappings and interaction design developed to address the requirements outlined in the Tasks and Design Requirements section. Our system follows a coordinated multiple views (CMV) approach (see Figure 1), where several linked views work together to support exploration and comparison. In the following subsections, we describe each view in detail and highlight its role in addressing the identified requirements. Together, these coordinated views provide an integrated environment that addresses the identified tasks and design requirements, enabling casual users to efficiently explore and compare urban data.

Data Selection View

The data selection view (marked with A in Figure 1) provides the interface for choosing, weighting, and organizing attributes of interest (Requirements R7 and R8). Since the number of available attributes typically does not exceed a few dozen, we provide a simple list interface that supports straightforward selection. For each selected attribute only, the view offers sliders for adjusting weights, a option to switch between the two available mapping functions (by clicking on the "more/less is preferred" label), a distinct color (from the Tableau 20 color palette) that is consistently applied across all views and can be quickly verified within the data selection view, and arrows to modify the attribute order. This can be seen in more detail in Figure 4.

To address reproducibility (R9), the data selection view further allows users to save and reload interesting exploration states. At any point, a user may store the current configuration under a chosen name. A list of saved sessions is provided, and a single click is sufficient to restore a previous state, enabling users to continue their exploration seamlessly without losing prior progress.




Figure 3. Multi-attribute ranking comparison view showing the same data as in Figure 1, but with a different attribute order to facilitate easier comparison of the attribute at the base of the bars.

City Map View


The city map view (marked with B in Figure 1) serves multiple purposes. First, it provides an overview of the districts and spatial structure of a city (R1). By including two city map views within the same dashboard, users can compare the spatial features of two cities (R2). This follows the juxtaposition strategy for comparison, where entities are placed in parallel visual contexts to facilitate direct inspection. The well-known map metaphor ensures that the visualization is immediately understandable.

Users can actively select an attribute in the data selection view, which is then used for color coding the map (R3). A linear color gradient from white to the attribute color encodes the values of this single attribute, scaled from the lowest to the highest district values across both cities. This also enables district comparison (R5) within the city map view as seen in more detail in Figure 2. In contrast, to highlight overall results across all selected attributes, the top three districts in each city based on the multi-attribute ranking (R8) are marked with gold, silver, and bronze icons. This distinction between single-attribute visualization and multi-attribute ranking allows users to explore both focused and aggregated perspectives simultaneously. Finally, the city map views allow users to select two districts (from the same or from different cities) for more detailed comparison within the data comparison view.


Multi-Attribute Ranking Comparison View

The multi-attribute ranking comparison view (marked with C in Figure 1) is the central visualization component for comparison and exploration, driven by user-defined weights and mapping functions for the selected attributes (R8). Positioned between the two city map views, it simultaneously shows the weighted attribute scores for both cities. The computation of the ranking scores is discussed in detail in the previous section.

The visualization employs stacked bar charts, with one bar representing each district. Districts of the city displayed on the left map are shown to the left of the vertical axis, while districts



(a)



(b)

Figure 4. (a): Data selection view showing the initial data selection and preference settings of the use case. The green area attribute is selected and therefor also shown on the map views as seen in Figure 2 and Figure 5. (b): Data selection view showing the adjusted data selection and preference settings of the use case. Different weights and preference mappings are applied among the attributes leading to the dashboard state depicted in Figure 1.

of the right city are displayed on the right. Within each city, the highest-ranked district is placed closest to the axis, followed by lower-ranked districts further outward. As a result, the view is not

symmetrical when the two cities differ in the number of districts. This view ultimately also allows for district comparison (R5) in the context of ranking scores.

The stacked bars enable comparison of overall ranking scores as well as of individual attribute contributions. In particular, the lowest portion of each bar is aligned across districts, which simplifies comparison for the attribute currently positioned at the base. To facilitate comparisons of different attributes, the order of stacked attributes follows the arrangement in the data selection view. Any reordering performed there is immediately reflected in the multi-attribute ranking view. Figure 3 shows an example of the same view as Figure 1 but with a different order of attributes.

Data Comparison View

To enable users to drill down and compare data between two districts in greater detail (R5), we introduce the data comparison view (marked with D in Figure 1).

For each selected attribute, the view displays a pair of bars representing the values for the two chosen districts. To distinguish the selected districts while maintaining consistency with the attribute color scheme, two shades of the attribute's color are used.

When more fine-grained data is available for an attribute (e.g., tree types for the attribute “trees per area”), additional pairs of bars are shown, one pair per subcategory. This allows users to discover valuable details. For instance, as illustrated in Figure 1, the total number of trees is higher in Berlin's sixth district (BER-6) than in Vienna's thirteenth district (VIE-13), yet the latter contains more maple (Ahorn) trees.

Displaying such levels of detail in other views would quickly become overly complex and reduce scalability. Presenting them on demand within the data comparison view, however, keeps the overall system efficient and manageable while still supporting detailed exploration when needed.

This view complements the city map and multi-attribute ranking views by focusing on detailed, local comparisons, allowing users to inspect specific district-level data that is aggregated or summarized in the other views.

Implementation

The dashboard is implemented as a webpage using React and the Ant-Design UI library. All charts and maps are created using D3 [20]. Given that our solution is very light-weight regarding computation and rendering we expect no performance issues on most hardware configurations.

Use Case

To highlight the capabilities of our approach, we provide an example use case. Afterward, we will also discuss the results of the initial evaluation of the dashboard with a test user.

The dashboard in its initial state offers many different ways in which users can begin their data exploration, the only necessary step is the selection of at least one data attribute within the data selection view. While the ranking implementation also allows us to perform simple queries (e.g. find the district with the cheapest rent) we choose multiple attributes for a more detailed comparison. One such example use case where we want to compare many attributes across districts is looking for a new district to move to. In this case we are interested in more green area, little tourism, so low number of overnight stays and higher average apartment size.




Figure 5. State of the dashboard views after three attributes are selected and a weight and preference mapping is applied in the context of the use case. The applied data selection can be seen in Figure 4a.

According to our priorities we assign preferences to each selected attribute as seen in Figure 4a. Based on the rank computation we now select the top ranking district for Vienna (VIE-13) and Berlin (BER-9). This data and preference selection leads to the state of the views depicted in Figure 5 with higher rank scores among the city's outskirts districts. Because the green area attribute is actively selected within the included data attributes, the green area values of the districts are shown on the maps as well which is shown for Vienna in Figure 2.

At this point, there is new information that we can extract from the current state of the dashboard. First, as mentioned, the outer districts tend to have much more green area and while both cities have very little green area in the central districts, the outer district of Vienna have more green area compared to Berlin. Also, the first district in Vienna is an outlier regarding average apartment size. And, to finish our first assessment, we see that Vienna's top scoring district reach a higher score compared to Berlin.

Now, based on the current state we can explore the data further and pivot to a more detailed comparison. We add the amount of street trees to additionally encode our interest in districts with more green area, but only with a ranking preference weight of 50%. Further, as we would love to find the best apartment in the city (according to our preference), we want to find something affordable. So, we include the price of rent in our selection and set a somewhat medium preference for cheaper rents. Last, we figure, that while the amount of tourism is a factor, what we really want to focus on is finding districts which are not highly populated. This leads to adjusting the ranking weight of little tourism back to a lower value but adding population density with a high factor and preference for less population density in our ranking. The adjusted data selection can be seen in Figure 4b and the resulting dashboard state is shown in Figure 1. At this point, it is also possible to reorder the selected attributes using the up and down buttons within the data selection view leading to the updated bar




Figure 6. Traffic accident victim data displayed on the Vienna city map view per district, showing the first district being an outlier and the ranking markers showing the districts with the least amount of traffic accident victims.

order within the stacked bar chart in Figure 3.

Finally, after further exploration, we can assess the result again. Although the thirteenth district of Vienna remains the highest scoring district, the scores of other districts have changed more significantly. The ranks of the other top districts of both cities have changed and in addition, the top Berlin districts score much higher compared to the top Vienna districts. To point out the repeating outlier again, we can also see that the first district of the Vienna scored no points because of the expensive rent but scoring the maximum amount of points regarding the apartment size.

Ultimately, while we could, again, adjust our selection and preferences and explore further, we have answered our starting question while gaining insight in many different aspects and details regarding the selected attributes across both cities.

Initial Evaluation

For a first evaluation, we conducted a usability test with one casual user. The user was presented with eleven different tasks: five basic usability tasks, five more complex exploration tasks, and one open exploration task. The user was asked to think aloud during the tasks, allowing us to draw conclusions not only about usability but also about general enjoyment and potential knowledge generation while interacting with the dashboard.

The user successfully completed all tasks. The first five tasks focused on basic interactions with the dashboard, such as locating a district within the city map view, saving and loading the current dashboard state, identifying a district based on specific data criteria, and comparing two districts directly.

The next five tasks emphasized exploration, allowing the user to investigate one or more data attributes across multiple districts independently. The final task allowed the user to explore the dashboard freely for five minutes, with the option to end the session at any time.




Figure 7. Resulting multi-attribute ranking comparison view after data of green area and average cost of real estate is included with a preference for higher values for both attributes. The view is showing that while some greener districts are more expensive regarding real estate cost, there is no clear relationship between these two attributes.

During two tasks, the user was provided with help: in one case, guidance was given regarding the unit of a selected data attribute, and in another, the user was reminded of the possibility to view district data in more detail. As a result, we improved the contrast of the unit text in the data selection view and made the optional interaction with the detail-level district data more visible within both the city map and data comparison views.

While performing the tasks, the user made numerous remarks about their findings during data exploration. For example, they pointed out clear outliers (such as the high number of traffic accidents in Vienna's first district as seen in Figure 6) and drew conclusions about the relationship between two data attributes (such as the amount of green area compared to housing costs as seen in Figure 7). Furthermore, during the open exploration task, the user chose to continue interacting with the dashboard beyond the suggested time limit, indicating engagement and interest.

To further evaluate our work, a more extensive usability study including both casual and expert users is planned.

Conclusion

In this paper, we address the challenging problem of urban data exploration for a broad audience, recognizing that users have diverse interests and priorities. By enabling user-specified weighting of attributes, our system supports customized, personalized comparisons of city districts across multiple criteria.

Given the size and complexity of urban data, we provide a design for a multi-level solution that integrates coordinated views, providing both overview and detailed exploration. Special attention was given to an accessible multi-attribute ranking approach, the arrangement of views and the use of consistent color scales to ensure intuitive interpretation and effective comparison.

The first pilot study yielded very positive feedback, confirming the usability and usefulness of the system and motivating a larger, more comprehensive evaluation in the future. These results

demonstrate the potential of our approach to make urban data exploration accessible and meaningful for a wide range of users.

Acknowledgments

The VRVis GmbH is funded by BMIMI, BMWET, Tyrol, Vorarlberg and Vienna Business Agency in the scope of COMET – Competence Centers for Excellent Technologies (904918, 911654) which is managed by FFG.

References

- [1] Zikun Deng, Di Weng, Shuhan Liu, Yuan Tian, Mingliang Xu, and Yingcai Wu. A survey of urban visual analytics: Advances and future directions. *Computational Visual Media*, 9(1):3–39, 2023.
- [2] Stadt Wien. Stadtplan Wien. <https://www.wien.gv.at/stadtplan/>, 2025. Accessed on 2025-03-22.
- [3] BerlinOnline GmbH. Stadtplan berlin. <https://www.berlin.de/stadtplan/>. Accessed on 2025-03-22.
- [4] Stadtplan london. <https://www.cityoflondon.gov.uk/services/planning/planning-policy/city-plan-2040>. Accessed on 2025-03-22.
- [5] Amt für Statistik Berlin-Brandenburg. Statistik berlin-brandenburg. <https://www.statistik-berlin-brandenburg.de/>, 2025. Accessed on 2025-03-22.
- [6] Bundesanstalt Statistik Österreich. Statistik Austria. <https://www.statistik.at/>, 2025. Accessed on 2025-03-22.
- [7] Regina Schuster, Kathleen Gregory, Torsten Möller, and Laura Koesten. “being simple on complex issues”—accounts on visual data communication about climate change. *IEEE Transactions on Visualization and Computer Graphics*, 30(9):6598–6611, 2024.
- [8] Yasmina Tajja and Ludger Martin. Urban data platforms as added-value systems for citizens. In *International Conference on Mobile Web and Intelligent Information Systems*, pages 175–184. Springer, 2023.
- [9] Eva Mayr, Paolo Federico, Silvia Miksch, Günther Schreder, Michael Smuc, and Florian Windhager. Visualization of cultural heritage data for casual users. In *IEEE VIS Workshop on Visualization for the Digital Humanities*, volume 1, 2016.
- [10] David Sprague and Melanie Tory. Exploring how and why people use visualizations in casual contexts: Modeling user goals and regulated motivations. *Information Visualization*, 11(2):106–123, 2012.
- [11] Zachary Pousman, John Stasko, and Michael Mateas. Casual information visualization: Depictions of data in everyday life. *IEEE transactions on visualization and computer graphics*, 13(6):1145–1152, 2007.
- [12] Yasmina Tajja and Ludger Martin. Lessons learned: A usability study of an urban data platform for citizens. In *International Conference on Mobile Web and Intelligent Information Systems*, pages 282–289. Springer, 2024.
- [13] Mirko Marras, Matteo Manca, Ludovico Boratto, Gianni Fenu, and David Laniado. Barcelonanow: Empowering citizens with interactive dashboards for urban data exploration. In *Companion Proceedings of the The Web Conference 2018*, pages 219–222, 2018.
- [14] Samuel Gratzl, Alexander Lex, Nils Gehlenborg, Hanspeter Pfister, and Marc Streit. Lineup: Visual analysis of multi-attribute rankings. *IEEE transactions on visualization and computer graphics*, 19(12):2277–2286, 2013.
- [15] Michael Gleicher, Danielle Albers, Rick Walker, Ilir Jusufi, Charles D. Hansen, and Jonathan C. Roberts. Visual comparison for information visualization. *Information Visualization*,

10(4):289–309, oct 2011.

- [16] Tamara Munzner, François Guimbretière, Serdar Tasiran, Li Zhang, and Yunhong Zhou. Treejuxtaposer: Scalable tree comparison using focus+context with guaranteed visibility. *ACM Trans. Graph.*, 22(3):453–462, jul 2003.
- [17] Edward R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 2 edition, 2001.
- [18] Jeffrey Heer, Jeffrey H. Wilkinson, and Maneesh Agrawala. Improvise: A framework for authoring interactive visualizations of structured data. In *Proceedings of the 2005 IEEE Symposium on Information Visualization (INFOVIS 2005)*, pages 39–46. IEEE Computer Society, 2005.
- [19] Wei Ren, Jeffrey Heer, and Maneesh Agrawala. Visbricks: A system for creating and manipulating Modular visualizations of tabular data. In *Proceedings of the 2011 ACM SIGCHI Conference on Human Factors in Computing Systems (CHI 2011)*, pages 1365–1374. ACM, 2011.
- [20] Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. D³ data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011.

Author Biography

Joshua Pola recieved his BSc in business informatics from the TU Wien (2019) and is currently continuing studying in this field as a masters student. Additionally, he is occupied at Halo GmbH in a software engineering and project management role (since 2020).

Krešimir Matković is a senior researcher, area coordinator of the Informed Decision Making with Machine-in-the-Loop Area, and head of the Interactive Visualization Group at VRVis Research Center in Vienna. Since 2000, he has been leading basic and applied research projects at VRVis with numerous national and international partners. He teaches at TU Wien, where he received his doctoral degree and habilitation (in 1998 and 2016) and at the University of Zagreb, where he received his graduate degree in 1994.

JOIN US AT THE NEXT EI!

electronic IMAGING

Imaging across applications . . . Where industry and academia meet!



- **SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •**
- **INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •**

www.electronicimaging.org

