BarcelonaNow: Empowering Citizens with Interactive Dashboards for Urban Data Exploration

Mirko Marras
University of Cagliari
Cagliari, Italy
mirko.marras@unica.it

Matteo Manca
EURECAT
Barcelona, Spain
matteo.manca@eurecat.org

Ludovico Boratto
EURECAT
Barcelona, Spain
ludovico.boratto@acm.org

Gianni Fenu
University of Cagliari
Cagliari, Italy
fenu@unica.it

David Laniado
EURECAT
Barcelona, Spain
david.laniado@eurecat.org

ABSTRACT

The massive amount and variety of city-related data raise equally big challenges to enable citizens to make sense of such data for improving their daily life and fostering collective decision making. The existing dashboards include limited pre-defined use cases which can only address the most common needs of citizens, but do not allow for personalization. In this work, we propose an open source environment that enables citizens to easily explore city-related data. A backend continuously collects heterogeneous data, turns them into a unified format and exposes them through an API; a front-end allows users to create interactive visualizations combining different data sources and visual models, and to share them to engage others. In this way, citizens can build up a data-driven public awareness supporting an open, transparent, and collaborative city.

KEYWORDS

Data Exploration; Data Visualization; Urban Dashboard

ACM Reference Format:

Mirko Marras, Matteo Manca, Ludovico Boratto, Gianni Fenu, and David Laniado. 2018. BarcelonaNow: Empowering Citizens with Interactive Dashboards for Urban Data Exploration. In *WWW '18 Companion: The 2018 Web Conference Companion, April 23-27, 2018, Lyon, France*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3184558.3186983>

1 INTRODUCTION

The proliferation of technologies interconnecting objects and people raises key challenges regarding the infrastructures to manage the complexity of the generated data, and the visual interfaces to leverage such data to support users in decisions and actions [1]. Over the years, the increasing number of open datasets available in raw and heterogeneous formats (e.g., API, CSV, JSON, RDF, or plain text) has mainly fostered analysis by data scientists. Tailored architectures and urban dashboards [2-5] are just recently emerging to unlock this hidden potential to the general public and may help

to better understand city dynamics, including noise pollution [6], mobility patterns [7, 8], and electricity usage [9].

The existing dashboards include limited pre-defined use cases which can only address the most typical needs without personalization capabilities. This is one of the reasons the impact of open data on public good is still limited [10]. On the other hand, customizing applications that can support end-users' specific goals requires expertise not expected of most of them. Offering ad-hoc interfaces where non-expert users can interactively personalize applications has shown promising results to reduce the gap [11]. By following this idea in the urban domain within the DECODE project¹, we aim to empower curious citizens with little or no technical expertise to get insightful knowledge from urban data.

In this paper, we propose an environment that enables citizens to create and explore interactive visualizations of city-related data and arrange them into personalized dashboards. Our solution offers customization functionalities to different exploration scenarios according to users' specific goals. A back-end aggregator built upon state-of-the-art technologies provides unified access to heterogeneous data. On top of it, a front-end interface allows users to create, combine, share and reuse visualizations, fostering data-driven public awareness.

The contribution is threefold:

- Build an easy-to-deploy framework capable of effectively collecting, processing, exposing and visualizing data. The code is released under open source license².
- Develop a front-end interface which enables citizens to explore data, arrange and share visualizations about city issues, and compare information coming from different sources.
- Implement a prototype version of the dashboard operating on real-life data coming from open and public repositories and systems in Barcelona. The live demo is accessible online³.

The paper is organized as follows. Section 2 shows how existing urban dashboards relate to our proposal. Then, Section 3 describes BarcelonaNow and Section 4 showcases its capabilities. Section 5 depicts some conclusions and future directions.

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '18 Companion, April 23-27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5640-4/18/04.

<https://doi.org/10.1145/3184558.3186983>

¹The project website is available at <https://www.decodeproject.eu/>

²The code is available at <https://github.com/DECODEproject/bcnnow>

³The live demo is available at <http://bcnnow.decodeproject.eu>

Figure 1: The reference architecture.

Figure 2: Sample dashboard overview.

2 EXISTING URBAN DASHBOARDS

In the light of the huge quantities of city data, both academia and companies have investigated ways to analyze and operationalize such data on urban dashboards [12]. Here, we briefly discuss existing implementations and how they relate to our proposal.

Madrid Dashboard⁴ aims to experiment with smart city services and fine-tuning them before deployment at large scale. The platform currently offers two services, namely people flow monitoring based on wireless tracking and environmental monitoring. For instance, such services can be applied to smartly control the university heating, ventilation, and air conditioning systems. UK City Dashboard⁵ summarizes quantitative data about the major United Kingdom cities on a single screen. It primarily shows weather, environmental, transportation, and energy demand with color-coded numerical values. Even though it provides a bird-eye city view, the overall selection and representation of real-time data is not translated from raw numerical values to a format that might be easier to be digested by non-technical users. Edmonton Dashboard⁶ enables citizens to visualize pre-defined snapshots of city data together with simplified and descriptive indicators under consistent color-coding, iconography and fonts. Furthermore, the dashboard allows a deeper historical analysis with interactive tools to filter data. Amsterdam Dashboard⁷ enables professionals to visualize data on a map view, displaying points representing discrete information, and geographical partition views, where each partition displays a certain category on which city elements are projected. Dublin Dashboard⁸ pulls together hundreds of data representations grouped in different modules, including overall statistics from the city and information from key points. Boston Dashboard⁹ helps users to monitor the city with a set of baseline metrics and plans on several areas, including education, infrastructure, and housing. The platform does not have a real user dashboard, but provides a smart tool to enable professionals to develop their own visualizations from a limited set of types, requiring technical operations.

⁴<http://ceiboard.dit.upm.es/dashboard>

⁵<http://citydashboard.org/choose.php>

⁶<https://dashboard.edmonton.ca/>

⁷<https://data.amsterdam.nl>

⁸<http://www.dublindashboard.ie>

⁹<https://boston.opendatasoft.com>

CityPulse¹⁰ is an open source framework which provides access to historical data that the user can visualize across multiple contexts.

With respect to these solutions, our proposal has advantages in terms of openness, flexibility, usability and personalization, so it can be employed for a wider range of scenarios and use cases.

3 THE BARCELONANOW DASHBOARD

The dashboard is composed of a back-end subsystem acting as data aggregator and manipulator, accompanied by a web-based front-end subsystem. The front-end enables citizens to leverage the data provided by the back-end to create interactive visualizations and personalized dashboards. Figure 1 depicts the architecture.

Back-end. The data aggregator and manipulator prototype is written in Python and runs on the Cloud. It collects streaming and non-streaming data coming from heterogeneous data sources and formats (e.g., JSON, CSV, RDF, API), translates such data to an internal data format based on international standards, stores them into a database and provides access to them through API. Each data source can be integrated into the back-end subsystem with the creation of its own data collector using either pull/get technology (i.e., collectors periodically request data to data sources) or push technology (i.e., data sources send data to collectors whenever available). Leveraging different pre-defined templates of collectors, developers can easily integrate a new data source by selecting the appropriate template, modifying it to set the mapping between the original data structure and the internal common data structure and adding it to the collector list. Once the collector accesses the original data records, they are translated from the original schema to a common schema based on standard time and geographic formats used to perform internal data manipulation¹¹. Subsequently, because the back-end subsystem aims to provide historical access to data and is expected to receive geo-temporal queries, each translated record is stored into an instance of MongoDB, which natively supports massive datasets and geo-operations on them. To allow efficient search within geographic coordinates or time intervals, proper database indexes are created on the corresponding data fields. Finally, the

¹⁰<https://github.com/CityPulse/CityPulse-City-Dashboard>

¹¹The documentation released together with the code repository provides more details.

data are made externally available via API using Flask. The flexibility and control given by Flask makes it easier to create a unique point for providing access to both stored data and data which require to be accessed on-the-fly (e.g., personal data). The back-end components and processes enable the simple integration of new data producers and data consumers (e.g., web/mobile applications).

Front-end. The front-end dashboard prototype currently consists of a web application implemented via HTML and JavaScript with the support of jQuery and jQueryUI libraries. The Bootstrap library is used to build the responsive layout and arrange widgets in grids. The interactive visualization of data in charts and maps is handled by D3.js, Plotly, and Leaflet libraries. The communication between the front-end and back-end subsystems is built on AJAX and uses JSON for transmitting information in both directions. The technologies we used are standard and compatible with existing browsers, ensuring platform-independence.

The primary components provided by the subsystem are a visual interface to create and explore interactive visualizations and a set of functionalities to enable citizens to organize such visualizations into different dashboards in accordance with needs and goals. Figure 2 showcases a dashboard. On the left sidebar, the interface gives access to the creation tool and the dashboards previously created.

The creation tool allows users to interactively compose and manage widgets with visualizations fed by a subset of the data provided by the back-end subsystem. Each widget indicates a user-defined title, the author's name, and the last modification time. Furthermore, for each included data source, the widget lists the name together with the set of user-selected parameters. Interacting with the icons on the widget header, the user can insert one or more data sources to be visualized inside it, and specify additional parameters, such as the time range, the time granularity (e.g., daily, weekly, monthly, yearly), the visualization type (e.g., markers, heatmap), the geographical aggregation (e.g., individual, census section, neighborhood, district), and filtering keywords (Figure 3). On the center, the preview widget panel depicts how the visualization appears with the current settings. The legend shows the meaning of different colors. In addition to this, users can inspect and personalize the detailed views regarding the individual elements depicted on the visualization. For instance, markers-based visualizations provide functionalities to select markers and interact with a panel providing further exploration capabilities.

Figure 3: The interface to insert a new data source.

The environment enables users to define one or more dashboards where they can group the widgets to be monitored on a single screen based on different thematics, goals, and needs. Each widget can be directly moved inside the same dashboard or between different dashboards by drag and drop and shared via short custom links to allow other users to view and reuse them. The whole dashboard can be shared as well. The sharing methodology includes a set of actions and parameters to reproduce the visualizations without sharing the data itself. This could also support more advanced scenarios where the environment is deployed in a setting involving individual users with different access permission to one or more subsets of the same data source. By clicking on the link, the individual widget or the whole dashboard is reproduced, so that the current user can explore and modify it. We envision each user to have a personal repository where their widgets are saved to be accessed in future sessions.

4 DEMONSTRATION

The demonstration is planned as a short journey leading the audience along the core functionalities of BarcelonaNow through real-life scenarios. The data sources integrated into the prototype come from the Barcelona City Council and third-party providers: ODI¹² for public-body generated data, ASIA¹³ for events and buildings, IRIS¹⁴ for citizen complaints, incidents, and suggestions, Sentilo¹⁵ and Smart Citizen¹⁶ for sensors data, and Inside Airbnb¹⁷ for housing data. To grasp the contribution of our work, the attendees can actively interact during the process and explore the interface on laptops and mobile devices. The demonstration runs as follows.

Single-Level Visualization. First, we focus on a motivating scenario: living in a large city like Barcelona means becoming intimately acquainted with noise and noise is not equally distributed across neighborhoods, so we show an interface to understand temporal noise patterns in different points of the city. Beginning with a single data source (50 noise sensors from Smart Citizen, data updated every minute, historical data of 12 months, about 20 millions of observations in total), we demonstrate the dashboard capability of easily creating a new data visualization showing a time-evolving map of acoustic levels (Figure 4 - Left). Playing with the controls, it is possible to recognize daily and weekly patterns in different neighborhoods, and identify the areas more affected by noise pollution during nights or weekends. Other analogous visualizations could be created with any kind of data coming with a location and a timestamp, such as events in the city from ASIA, or the density of bikes into the bike-sharing stations from ODI.

Multiple-Level Visualization. In the second part, once an interesting data source combination is identified by the audience, such as the noise levels already dumped from Smart Citizen and the housing information from Inside Airbnb, we demonstrate the multiple data source combination capability that is inherent to the visualizations made available by the environment. To this end, we use the dashboard to create a widget including two data sources displayed into stacked layers with two different types of visualization

¹²<http://opendata-ajuntament.barcelona.cat/en/open-data-bcn>

¹³<http://opendata-ajuntament.barcelona.cat/data/en/dataset/agenda-mensual>

¹⁴<http://opendata-ajuntament.barcelona.cat/data/en/dataset/iris>

¹⁵<http://www.sentilo.io/wordpress/>

¹⁶<https://smartsitizen.me/>

¹⁷<http://insideairbnb.com/>

Figure 4: Noise measurements along time with temporal pattern for two sensors (left) and heat map of the Airbnb listings distribution stacked with noise sensor measurements (right). Hotter colors represent higher values of noise/listing density.

(Figure 4 - Right). In this way, the attendees can put in contrast the information to get insights and correlations, and to answer questions such as: do areas with high density of short-term rent listings have higher levels of noise during night? Furthermore, other data sources can be inserted into the debate to enrich the information available to better understand the city dynamics. One example may be showing how citizens complaints about noise from IRIS are distributed across neighborhoods in comparison with the density of short-term rent listings or the location of the noise sensors to detect possible issues, such as whether noise sensors monitor areas affected by a large number of noise complaints. Other visualization types can be leveraged in a similar way.

Dashboard Personalization. Subsequently, we show how the environment can be used to arrange the visualizations on multiple personalized dashboards. The audience can create different pages and move visualizations between them based on specific goals and needs. Furthermore, we demonstrate how the produced assets can be shared to foster co-creation or raise awareness on city issues.

Data Catalog Exploration. Finally, we guide the attendees through the rest of the catalog in order to catch interesting opportunities and get insights out of such data. For each data source, we briefly show a description together with the set of attributes it embraces. The exploration can stimulate discussions about ideas and hypotheses regarding visualizations potentially useful to citizens, or unveil patterns and bring evidence of city issues. The attendees are invited to play with the dashboard on data of interest.

Hardware Requirements. The demonstration requires a stable Internet connection and a large-screen laptop, even though an additional external wider screen can enhance the overall interaction and potentially engage more attendees together.

5 CONCLUSIONS

With the proposed dashboard, we presented a new and comprehensive way of enabling citizens to define and tailor city-related data visualizations by themselves. Instead of relying on pre-packaged solutions that can only address limited and common needs, we enable citizens with little or no expertise to create their custom visualizations by leveraging the provided functionalities. Our proposal constitutes an ecosystem in which citizens and policy makers can share and co-create visualizations to increase public awareness on city issues, supporting an open, transparent and democratic city.

In the next steps, we plan to extend the data source catalog with other open and public data and to integrate private data coming from the DECODE infrastructure in a privacy-aware manner. Moreover, we will enable citizens to perform more advanced explorations and interact with an enriched set of visualization types, also with the support of underlying automated data mining techniques.

ACKNOWLEDGMENTS

This work has been partially supported by the DECODE (DEcentralised Citizens Owned Data Ecosystem) project, funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No 732546.

Mirko Marras acknowledges Sardinia Regional Government for financial support of his PhD scholarship (P.O.R. Sardegna F.S.E., European Social Fund 2014-2020 - Axis III, TG 10, PoI 10ii, SG 10.5).

REFERENCES

- [1] Simon Elias Bibri and John Krogstie. Smart sustainable cities of the future: An extensive interdisciplinary literature review. *Sustainable Cities and Society*, 2017.
- [2] Unai Aguilera, Oscar Peña, Oscar Belmonte, and Diego López. Citizen-centric data services for smarter cities. *Future Generation Computer Systems*, 2016.
- [3] Bin Cheng, Salvatore Longo, Flavio Cirillo, Martin Bauer, and Erno Kovacs. Building a big data platform for smart cities: Experience and lessons from Santander. In *Proceedings of BigData Congress. IEEE*, 2015.
- [4] Cristian Chilipirea, Andreea-Cristina Petre, Loredana-Marsilia Groza, Ciprian Dobre, and Florin Pop. An integrated architecture for future studies in data processing for smart cities. *Microprocessors and Microsystems*, 2017.
- [5] Vladimir Zdraveski, Kostadin Mishev, Dimitar Trajanov, and Ljupco Kocarev. Iso-standardized smart city platform architecture and dashboard. *IEEE Pervasive Computing*, 16(2):35–43, 2017.
- [6] Juan M Navarro, JB Tomas-Gabarron, and Jose Escolano. A big data framework for urban noise analysis and management in smart cities. *Acta Acustica united with Acustica*, 103(4):552–560, 2017.
- [7] Andreas Kaltenbrunner, Rodrigo Meza, Jens Grivolla, Joan Codina, and Rafael Banchs. Urban cycles and mobility patterns: Exploring and predicting trends in a bicycle-based public transport system. *Pervasive and Mobile Computing*, 6(4):455–466, 2010.
- [8] Yikai Gong, Luca Morandini, and Richard O Sinnott. The design and benchmarking of a cloud-based platform for processing and visualization of traffic data. In *Proceedings of BigComp. IEEE*, 2017.
- [9] Joaquim Massana, Carles Pous, Llorenç Burgas, Joaquim Melendez, and Joan Colomer. Identifying services for short-term load forecasting using data driven models in a smart city platform. *Sustainable Cities and Society*, 28:108–117, 2017.
- [10] Ammar Gharaibeh, Mohammad A Salahuddin, Sayed Jahed Hussini, Abdallah Khreishah, Issa Khalil, Mohsen Guizani, and Ala Al-Fuqaha. Smart cities: A survey on data management, security, and enabling technologies. *IEEE Communications Surveys & Tutorials*, 19(4):2456–2501, 2017.
- [11] Fabio Paternò and Volker Wulf. New perspectives in end-user development, 2017.
- [12] Rob Kitchin. The real-time city? big data and smart urbanism. *GeoJournal*, 79(1):1–14, 2014.