# 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 advent of massively interconnected technologies over cities raises equally big challenges regarding interfaces to enable citizens to make sense of urban data for improving their daily life and for fostering online participation. The existing dashboards include only pre-defined and limited use cases which can only address the most common needs of citizens, but do not allow for personalization. As consequence, the great effort of cities to make data widely available has still scarce capacity to get an impact for the public good. In this work, we propose an open source dashboard with a set of tools and services to enable citizens to easily explore city-related data and create interactive visualizations. Moreover, users can personalize the dashboard based on their individual, local, task-specific goals and interests, and share the resulting dashboard to promote co-creation. In this way, citizens can build up a data-driven public awareness, supporting an open, transparent, and collaborative city where they are actively involved in local activities and campaigns.

## 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 *Proceedings of The Web Conference (TheWebConf 2018).* ACM, New York, NY, USA, 4 pages. https://doi.org/XX.XXX/XX_XX

## 1 INTRODUCTION

The proliferation of massively interconnected objects, devices and sensors around citizens raises equally big challenges regarding the end-user interfaces that enable cities to manage the complexity of the data they produce and exploit the opportunities that such data open up to support citizens in meaningful decisions and actions [1]. Many architectures are emerging to unlock this hidden potential and make sense of this big data [2–5], while urban dashboards from all over the world are enabling citizens to visualize data to better understand the operating dynamics of cities, including noise pollution [6], mobility patterns [7, 8], and electricity usage [9].

However, the existing dashboards include pre-defined and limited use cases which can only address the most typical needs of citizens, but do not allow for personalization. Consequently, the impact of open data on the public good is still limited [10]. On the other hand, most people are familiar with the basic interfaces, but customizing or modifying applications that can effectively support their goals still requires considerable expertise, that cannot be expected of everyone. As envisioned and confirmed in [11], developing environments that enable users with little or no experience to interactively arrange and personalize their applications has shown promising results to reduce this gap and leveraging it into the urban domain can greatly empower citizens to make sense of city data.

In this paper, we propose an open source dashboard which enables citizens to create and explore interactive visualizations of city-related data and arrange them into dashboard pages tailored to their goals. The environment is composed of a back-end aggregator built upon state-of-the-art Big Data technologies which provide efficient normalized access to heterogeneous data sources, and a front-end interface through which users can leverage the available data sources to create their compositions of interactive visualizations and, eventually, share them with other people, building up a data-driven public awareness. The contribution is threefold:

- Distribute an easy-to-deploy framework capable of effectively collecting, processing, and visualizing spatio-temporal data on different types of visualizations and dashboard pages. The code is released under open source licence[1].
- Develop an editor which enables citizens to arrange, explore, and share visualizations about facts that interest them, and cross and compare information from different sources.
- Implement a prototype dashboard operating on real-life data from open and public repositories and systems in Barcelona. The demo is accessible online[2].

Section 2 highlights how our proposal differs from the existing urban dashboards. Then, Section 3 describes the proposed dashboard and Section 4 summarizes its capabilities under real-world settings. Finally, Section 5 depicts conclusions and future research.

*TheWebConf 2018, 23-27 April 2018, Lyon, France*
2018. ACM ISBN XXX-XXXX-XX-XXX/XX/XX…$15.00
https://doi.org/XX.XXX/XX_XX

[1]Available at the address: https://github.com/DECODEproject/bcnnow
[2]Accessible at the address: http://bcnnow.decodeproject.eu

## 2 RELATED WORK

In the light of the huge quantities of city data, both academia and companies have invested in ways to analyze and operationalize such data on urban dashboards. Here, we discuss the most prominent alternatives and highlight how they differ from our proposal. See [12] for a more detailed survey on the existing dashboards.

Madrid Dashboard[3] aims at experimenting with smart city services and fine-tuning them before deploying them at large scale. The platform currently offers two services, namely people flow monitoring based on wireless tracking and environmental monitoring. For instance, the people flow monitoring information and the environmental information can be considered to smartly control the university heating, ventilation, and air conditioning systems. The dashboard cannot be personalized and the limited use cases reduce the number of potentially supported citizens. UK City Dashboard[4] 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. Although 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. Moreover, it does not give them context with data over recent hours, days, and months. Edmonton Dashboard[5] enables citizens to visualize pre-defined snapshots of city data integrated from various official sources together with simplified and descriptive indicators under consistent color-coding, iconography and fonts. Furthermore, the dashboard enables users to do a deeper historical analysis with interactive tools to filter data. However, no real-time operational data is available and it consequently misses a lot of potential use cases. Amsterdam Dashboard[6] is able to show data on a map view, displaying points representing discrete information , and a partition view, where each partition displays a certain category on which city elements are projected. The statistics are presented on blocks of 24 hours. However, it does not handle massive amounts of streaming data and dashboard personalization per user. Dublin Dashboard[7] pulls together data from the city council, the government departments and several existing smart city and social applications. The dashboard contains hundreds of data representations grouped in different modules, including overall statistics from the city and information from key points in the city. However, it does not enable users to modify the default visualizations. Moreover, it is not an open-source initiative and no way to create key performance indicators is presented. Boston Dashboard[8] helps residents to track the city progress with a set of baseline metrics and plans on several areas, including education, job creation, infrastructure, housing and commuting. The platform does not have a real user dashboard, but provides a smart tool to enable users to develop their own visualizations. However, the tool is not intended for users with no or little technical skills, such as citizens, since it embraces advanced operations like manual raw data querying.

[3]http://ceiboard.dit.upm.es/dashboard
[4]http://citydashboard.org/choose.php
[5]https://dashboard.edmonton.ca/
[6]http://citydashboard.waag.org
[7]http://www.dublindashboard.ie
[8]https://boston.opendatasoft.com

Comparing to these solutions, the proposed dashboard has advantages in terms of flexibility, usability and personalization, so it can be employed for wider scenarios and use cases.

## 3 THE BARCELONANOW DASHBOARD

The dashboard is presented in the form of a web application composed of a back-end subsystem acting as data aggregator accompanied with a front-end subsystem which enables citizens to build interactive visualizations and arrange them on personalized dashboard pages. Figure 1 depicts the reference architecture, whose main components are described below in more technical detail.

**Back-end Subsystem.** On the back-end side, the dashboard is written in Python and runs on AWS[9] (Amazon Web Services) cloud computing platforms. The subsystem collects streaming data coming from heterogeneous data sources, translates them to a common data format based on international standards, stores such data into a database and provides simple access to them through APIs. Each data source is collected and integrated by its own customized data collector using either periodic processes or push notifications. First, an instance of an Apache Kafka[10] producer is used to efficiently handle data feeds. Then, the data are translated from the original schema to a common schema used to easily access and filter data, based on standard temporal and geographic formats. It is worth noticing that the subsystem does not require pure real-time processing, because most of the operations are done in small batches. Subsequently, the data need to be stored in order to enable citizens to have access to historical data. Because the environment receives high amounts of data frequently queried with geo-temporal tags, MongoDB[11] database is exploited since it supports geo-operations. Database indexes are created on geographic and temporal attributes, to efficiently search for data within certain coordinates or time intervals. The stored data sources are made externally available in JSON format via REST APIs developed in Flask[12].

**Front-end Subsystem.** On the front-end side, the dashboard consists of a web interface implemented via HTML and JavaScript

[9]https://aws.amazon.com/it
[10]https://kafka.apache.org
[11]https://www.mongodb.org/
[12]http://flask.pocoo.org/

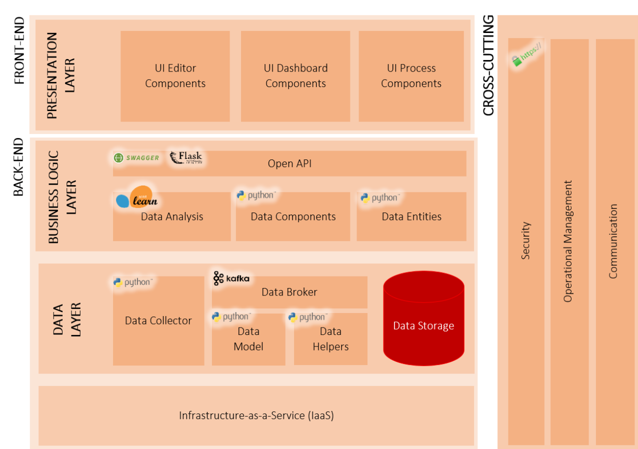


**Figure 1: The architecture of the proposed dashboard.**

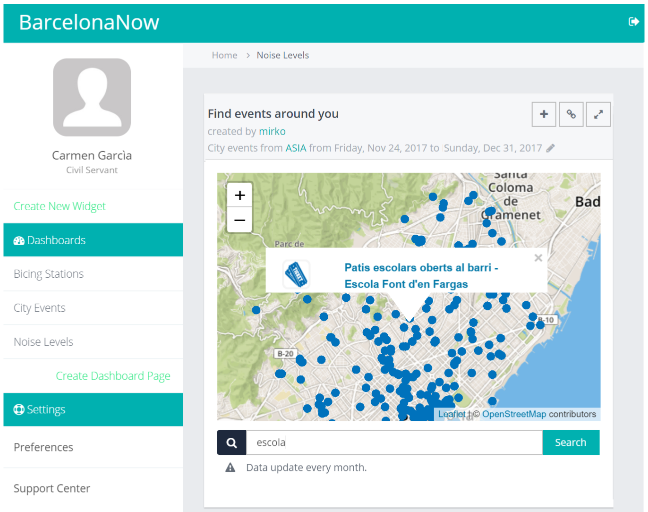


**Figure 2: Sample overview of the dashboard. On the center, a widget showing events distribution, filtered by a keyword.**

with the support of jQuery[13] and jQueryUI[14] libraries. The Bootstrap[15] library is used for the responsive layout and for displaying widgets in grids. The two primary purposes that are served are providing an editor to create and explore interactive visualizations and defining a set of functionalities to arrange different dashboard pages where citizens can organize the visualizations previously created in accordance with their goals. On the left sidebar, the editor and the list of the current dashboard pages are accessible through different sub-menus. The technologies we used are standard and compatible with existing browsers, ensuring platform-independence.

The editor allows users to interactively create and manage widgets with visualizations fed by a subset of the available data. Each widget is defined by (i) a title set at creation time, (ii) an author identified by the username of the creator, (iii) the creation and last modification time, (iv) the original data source names, and (v) the set of application user-selected parameters. On the center, the preview panel depicts how the visualization appears with the current settings. Interacting with the icons on the widget header, the user can specify one or more data sources to be visualized inside it by selecting the name, the type of visualization (e.g. map points, heat-map), the geographical aggregation (e.g. per point, per neighborhood, per district) and the time interval (e.g daily, weekly, monthly, yearly, custom). The visualization preview changes in real time. On the right side, whenever a map point is selected, a panel provides information for further exploration. The users can inspect and personalize the numerical statistical views regarding a specific point. The interactive visualization of data in charts and maps is handled by the D3.js[16] and Leaflet[17] JavaScript libraries which use a combination of HTML, CSS, and SVG to display data. The communication between the front-end and back-end subsystems is built on AJAX and uses JSON for transmitting information in both directions (i.e, input query on data sources and output data sources). Examples of the editing commands and a sample map visualization are shown in Figure 2.

Keeping in mind that the environment can be used by several users, each subscriber has a personal repository in order to save their resulting widgets. Each citizen can define one or more dashboard pages where they can arrange related widgets to be monitored on a single screen. To this end, dashboard pages support direct manipulation of interactive widgets inside the same page or between different ones. An additional functionality enables citizens to share visualization through social networks via custom links in order to allow other users to benefit from them. Users can also clone and modify an existing widget created by another user.

## 4 DEMONSTRATION

The demonstration is designed as a walk-through tour which guides the attendees along the core functionalities of the dashboard through real-life scenarios. To grasp the contribution of our work, the audience can interact during the process and explore the interface on their laptops and mobile devices. The demonstration is based on a set of data sources from the Barcelona City Council and third-party providers: ODI[18] for public-body generated data, ASIA[19] for events and buildings, IRIS[20] for citizen claims, incidents, and suggestions, Sentilo BCN[21] and Smart Citizen[22] for sensors data, and Inside Airbnb[23] for housing data. The demonstration runs as follows.

**Single-Data-Source Visualization Creation.** 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 update every minute, historical data of 12 months, about 20 millions of observations in total), we demonstrate the environment capability for easily creating a new data visualization showing a time-evolving map of acoustic levels (Figure 3 - Left). Hotter colors represent higher levels of noise, while colder colors depict more peaceful levels. Playing with the controls, it is possible to recognize both daily and weekly patterns in different neighborhoods, and to identify the areas more affected by noise pollution during nights or weekends. Other analogous visualizations could be created with any other kind of data coming with a geolocation and a timestamp, such as events in the city (from ASIA), or the density of bikesharing bikes available (from ODI).

**Multiple-Data-Source Visualization Creation.** 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 (Figure 3 - Right), 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 visualization with two data sources inserted on the same widget,

[13]https://jquery.com/
[14]https://jqueryui.com/
[15]http://getbootstrap.com/
[16]https://d3js.org/
[17]http://leafletjs.com/
[18]http://opendata-ajuntament.barcelona.cat/en/open-data-bcn
[19]http://opendata-ajuntament.barcelona.cat/data/en/dataset/agenda-mensual
[20]http://opendata-ajuntament.barcelona.cat/data/en/dataset/iris
[21]http://www.sentilo.io/wordpress/
[22]https://smartcitizen.me/
[23]http://insideairbnb.com/


**Figure 3: Sample visualizations showing noise along time with detailed temporal pattern for the selected sensor (left) and a heat map showing the distribution of Airbnb listings (right). Hotter colors represent higher values of noise/listing density.**

but in different layers. In this way, the attendees can put in contrast the information coming from the data sources to get insights and correlations, and to answer questions such as: do areas with high density of Airbnb 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 dynamics behind noise pollution and short-term house rental. One example could be a layer showing how citizen complaints about noise from IRIS are distributed across neighborhoods, in comparison with the density of Airbnb listings.

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

**Data Catalog Exploration.** Finally, we guide the attendees through the catalog of all the available data sources in order to get insights on the possibilities offered by such data. For each data source, we briefly show a short description together with the set of attributes it embraces and the update rate. The exploration can also stimulate discussions about ideas and hypotheses regarding visualizations potentially useful to citizens, or to unveil patterns and bring evidence of relevant city issues. The attendees will be invited to play with the dashboard to explore data of interest.

**Hardware Requirements.** Since the demonstration is realized via a web interface to a remote server application, a stable Internet connection and a laptop are entirely sufficient, although a larger screen can enhance the experience. Additionally, the audience is encouraged to explore the interface on their mobile devices.

## 5 CONCLUSIONS

With the proposed dashboard, we have presented a novel and comprehensive way of enabling citizens to define and tailor city-related data visualizations to their specific needs. Instead of relying on prepackaged solutions that can only address the most common goals, we enable citizens with little or no technical skills to build their custom visualizations, leveraging the functionalities of the urban dashboard pages. The framework also constitutes an ecosystem in which citizens and policy makers can share and co-create visualizations to increase public awareness on city issues, to support the development of an open, transparent and democratic city.

As next steps, we plan to extend the number of available data sources and to enable citizens to perform more advanced explorations, also with the help of automated data mining tasks running on the back-end and providing enriched data for visualization.

**Acknowledgments.** The European Union Horizon 2020 Programme under grant agreement N.732546 has partially supported this work during DECODE (DEcentralised Citizens Owned Data Ecosystem) project. 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 - Goal 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 Ernoe Kovacs. Building a big data platform for smart cities: Experience and lessons from santander. In *Big Data (BigData Congress), 2015 IEEE International Congress on*, pages 592–599. 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 *Big Data and Smart Computing (BigComp), 2017 IEEE International Conference on*, pages 13–20. 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.