

Real-time and archival data visualisation techniques in city dashboards

Abstract

City dashboards have become a common smart city technology, emerging as a key means of sharing and visualising urban data for the benefit of the public and city administrations. Operating as the front-end of many cities' data stores, dashboards display and benchmark indicators relating to city operations, characteristics, and trends, displayed through interactive visual representations of spatial and temporal patterns. Many dashboards collect, archive, and present data collected in real-time, as well as more traditional time-sliced administrative data. In this paper, we evaluate the techniques that dashboards employ to present real-time data to dashboard users. Our analysis identifies two factors that shape and differentiate real-time visual analytic tools: the dynamic nature of the data, how they are refreshed, and how the realtimeness of the data is communicated to the user; and how the tool enables archival comparison. We assess dashboard design according to the strategies used to address specific challenges associated each factor, specifically change blindness and temporal pattern detection. We conclude by proposing effective techniques for city dashboard design.

Key words: city dashboards, real-time data, visualisation, time series, real-time, urban analytics

1. Introduction

City dashboards are public-facing collections of front-end visualisations for presenting spatio-temporal data about the characteristics, administration, and operations of a city (Kitchin, et al., 2015; Mattern, 2015). Typically, dashboards employ a suite of dynamic and/or interactive graphics (e.g., gauges, traffic lights, meters, arrows, bar charts, graphs) and maps to display information about the performance, structure, pattern and trends of cities. The graphics are dynamic in the sense that they seek to update as new data are released, and

are interactive through operations such as selecting, filtering and querying data, zooming in/out, panning, and overlaying. The utility of the dashboard format is that it enables a user to gain a ‘span of control’ over a large amount of varied and quickly transitioning data (Brath & Peters, 2004). In particular, dashboards allow a user to track and compare over time and space, and in the case of real-time data, the here-and-now, of different phenomena. In some cases, key data are ‘consolidated and arranged on a single screen so the information can be monitored at a glance’ (Few, 2006, p. 36). Here, a city dashboard operates like a car dashboard or plane cockpit display providing critical information in a single view (Dubriwny & Rivards, 2004; Gray, et al., 2013). Recently, dashboards are being made accessible to the general public, partly as a means of enabling transparency and accountability as a component of an open government agenda, and partly for place promotion (Behn, 2014; Kitchin, et al., 2015). Dashboards are frequently used toward the same goals in the context of smart cities for both management of city resources and public outreach (Kitchin, et al., 2015).

A crucial element of a city dashboard is that it seeks to document change over time with respect to a defined geography (e.g., the city as a whole, neighbourhoods). That is, the underlying datasets are explicitly longitudinal in nature. To have utility, the desired temporal resolution of data is usually annual or sub-annual. In many instances, phenomena are tracked with respect to defined targets and milestones in order to be able to measure performance (Behn, 2014). Increasingly, city dashboards are incorporating high-velocity streaming data that are produced in real-time by sensors, actuators, meters, transponders, cameras, and computational devices, but also through crowdsourcing and locative and social media (Kitchin, et al., 2015). These real-time data enable a highly granular temporal view of the city, which is often also spatially granular, being associated with particular sites (e.g., bike stations, car parks, weather stations, pollution sensor stations). Being able to monitor and act on real-time information in the form of sensor readings and the locations of moving objects

such as public transit vehicles, has become critical for those charged with managing city systems and infrastructures (Coletta & Kitchin, 2017), but also for the public in planning their spatial behaviour on-the-fly in relation to present conditions (Kitchin, 2017; de Lange, 2018).

Heim (1994, p. 49) defines real-time as ‘simultaneity in the occurrence and registering of an event.’ In reality, not only is there a small latency between data generation and display, but the data are often temporally sampled, albeit with a relatively fine temporal scale (e.g., the number of cars in a car park is measured every 30 seconds). The temporal rate of data measurement and sharing is in part chosen and in part imposed: how a system is configured involves making decisions about balancing data resolution and noise (data quality) with respect to the task requirements against system configuration and performance (e.g., life of batteries, costs of data transmission/storage) (Kitchin, 2017). The system components and architecture also affect temporality due to network capacity, memory buffering, CPU scheduling, and process interrupts (Mackenzie, 2007, de Lange, 2018).

Weltevrede, Helmond, & Gerlitz (2014, p. 127) thus note that there are varying forms of ‘realtimeness’, which produce ‘real-time cultures’ within platforms and systems, including city dashboards (Kitchin, 2017). Real-time data and their usefulness in a city dashboard context are then highly mediated by the technologies which collect and digitise them, expose them to analysis, and display them graphically. Realtimeness is a property of a system which appears to provide feedback concurrently with the collection of new information, while recognising the mediation which reduces the concurrency of real-world observations.

Techniques and issues with regards to representing both spatial and temporal variation are well-documented (Fabrikant, et al., 2008; Fish, et al., 2011). However, to date, there has been no systematic evaluation of the specific mechanisms by which spatio-realtime data are visualized in dashboards of city data, or of their attendant issues. Thus, there are no standards or guides for visualisation design or user ability considerations to help implement usable

dashboards, despite their widespread use in cities around the world. In this paper, we present a novel analysis of city dashboards and use geovisual analytics to assess the creation of design principles for communicating real-time data, including archived time-series real-time data, drawing on an analysis of bespoke city dashboards. We examine the practices of city dashboard development with respect to real-time data, which include two primary themes: dynamic data, refresh, and reatimeness; and archival real-time data.

2. Real-time visualisation challenges

Technical issues around implementation of temporal visualisations is a significant concern to those developing city dashboards, but except in specific circumstances concerning data storage, technical challenges have been little discussed to date. In this section, we identify and discuss two common challenges in the visualisation of spatio-temporal data that are crucial to address to create effective city dashboards: perceptual issues of change blindness and communication of spatio-temporal variability.

2.1 Change blindness

Small, brief, or dispersed changes to a visual display introduce opportunities for missed perception of the change and potentially introduce the phenomena of change blindness; that is, a diminishing ability of viewers to identify changes to a visualisation when there is a transition between states (Goldsberry & Battersby, 2009; Rensink, 2002). Change blindness has been measured in animated cartographic maps, where several properties of dynamic displays impact the perceptual ability to recognise and enumerate changes (Fish, et al., 2011). Specifically, dynamic displays introduce the potential for three cognitive issues of display interpretation: inability to recognise that changes have occurred (Simons & Chabris, 1999),

inability to locate the elements which have changed, and inability to recount the amount of change which occurred.

Change blindness is more pronounced when the transition is not a gradual change but a rapid one, and when more of the visualisation changes simultaneously (Fish, et al., 2011). Brief periods of inconsistency or disappearance in the display is enough to disrupt the concentration needed to comprehend a pattern in the values as they shift between consecutive states of representation (Goldsberry & Battersby, 2009). Although display disruptions shift attention away from particular dynamic parts of the display, the absence of disruptions have also proven to increase change blindness in parts of the display where attention is not focussed (Simons, et al., 2000). This is a particularly important issue in web-based streaming dashboards, where browser pages refresh, creating disruption in the entire display, and parts of the display may change at different rates while others remain unchanged.

Some strategies have been proposed to counteract the multiple types of change blindness in digital displays. Some initial strategies propose changes to the representation of dynamic data within displays, such as time-series glyphs (Thakur & Hanson, 2010). Time series symbology can make streaming data prominent, while including previous measurements, making change detectable from a static representation. However, time series visualisation introduces additional visual clutter, especially to a spatial map. Cybulski and Medyńska-Gulij (2018) tested the effect of using redundant cartographic principles to represent dynamic data, showing that multiple visual indicators of weight can improve change perception. Many have suggested that interactivity in the form of control over the temporality of data allows a user to explore change at a pace which removes imperceptiveness to visual changes and enables multiple attempts to perceive specific changes (Fabrikant, et al., 2008; Harrower, 2007). DuBois and Battersby (2012) suggest a raster-based approach has advantages over vector representations for the perceptiveness of

change. The latter methods require a conceptual redesign of what the dashboard and its real-time components do, so may not be effective solutions within that environment.

Change blindness has been measured in animated displays, which is a property of real-time dashboards. However, real-time visual change requires new ways of reducing change blindness given the ways that real-time dashboards are used. The replacement of old values with new, current ones means that a review of previous measurements is not possible. The animations used to test change blindness assume rapid transition between temporal states of the display, which may be true of some dashboards, but many operate with long latency between new real-time measurements on the order of minutes. The challenge of change blindness cannot be completely overcome, but strategies common to some dashboard designs, as we will explore, reduce its negative effect on the ability to observe important temporal patterns in city data.

2.2 Variability in time and space

As historical data is integrated into spatial displays, representation becomes a greater challenge. Historical time series data, or longitudinal compilations of measurements collected from a single source at regular intervals between measurements, is important for representing urban patterns (Monmonier, 1990; Silva & Catarci, 2000) measured and regulated constantly and over the long term by Internet of Things (IoT) sensors and their accompanying governance structures (Batty, et al., 2012; Coletta & Kitchin, 2017). Representing time series information on spatial maps is an ongoing challenge, with solutions including small multiples of time sliced data (Fabrikant, et al., 2008), animations (Harrower, 2007; Robertson, et al., 2008), flow maps (Andrienko, et al., 2009), radial/polar plots (Draper, et al., 2009), and the 3D space-time cube (Bach, et al., 2014). Many types of patterns are relevant for urban and smart city analysis, and so visual analytics is frequently used to facilitate new pattern discovery through interactivity and question-based exploration of big temporal data (Dykes &

Mountain, 2003; Guo, et al, 2006). Still, visualisation may present only part of a comprehensive effort to understand temporal dynamics, as Peuquet et al (2015) state by integrating computational pattern discovery with interactive visual analysis.

In city dashboard applications, time series data serves primarily to provide a historical record of evolution and change with regards to a specific phenomenon. Temporal data is maintained and compared against present data in order to track variation and progress toward future goals (Kitchin, et al., 2015). Benchmarking is a common city strategy for comparing against other cities while tracking statistics within a single city, partly for city transparency, but primarily for competitive reasons and for monitoring policy performance (Huggins, 2010).

Geovisual analytics methodologies have rarely been used to assess the design function of interfaces to streaming smart city data. Social media has been one target of such research (Pezanowski, et al., 2018). Highly interactive spatio-temporal interfaces must balance the complexity of the data being explored and the needs and abilities of users (Wisniewski, et al., 2009). The use of high-velocity data in many formats including text, imagery, time series, and others means that data complexity is already a barrier to novice use of dashboards. Such properties of big data require new forms of visual exploration (MacEachren, 2017), but complexity remains a barrier to novice users. The methods explored in the following section add minimal visual complexity to the interface while reducing barriers to the usability of real-time data.

3. Real-time practices in city dashboards

These two challenges to temporal visualisation play a role in city dashboard design, but no standardised strategies have been established to counteract the negative effects in temporal visualisation design. To examine the strategies used to date to resolve these challenges, we

conducted a survey of techniques used in public, real-time dashboards and similar real-time tools. We used targeted web searches, browsing of known research groups, and institutional knowledge to survey dashboard tools and their techniques of visual presentation. Our survey was not designed to be an exhaustive evaluation of every city dashboard created to date, nor does it intend to assess individual dashboards for any particular purpose besides documenting their various visualisation techniques used for displaying real-time data. Our analysis consisted of systematically examining and deciphering real-time tools with regards to their objectives, underlying data, mode of presentation, and specific techniques for minimising the challenges identified in Section 2. In every example from which the strategies in table 1 are derived, the data used are real-time and the visualisation is dynamic and interactive; that is, the visualisation updates as new data are parsed in and users can click on or hover over the graphic to gain more information, zoom in/out, pan, etc. From our analysis, we identified several strategies which respond to real-time visual challenges, and have provided an example of a dashboard which implements the strategy (see Table 1). The examples given are illustrative of the strategy, but other dashboards may exist which do the same, and many of the examples here use more than one technique.

<u>Challenge</u>	<u>Method to address</u>	<u>Example</u>	<u>URL/reference</u>
	Manual refresh	Best et al	Best, D. M., et. al. (2012) Web-based visual analytics for social media
	Automatic refresh	MarineTraffic	https://www.marinetraffic.com/en/ais/home/centerx:-123.090/centery:49.304/zoom:13
Change	Indicator of refresh type	Human/Best et al	http://cities.human.co/
	Time since last update	Toronto	http://totransit.ca/#
blindness	Time until next update	London	http://citydashboard.org/london/
	Indicator of refresh rate	CEI Moncloa	https://ceiboard.dit.upm.es/dashboard/
	Updates in clock time	Boston	https://boston.opendatasoft.com/map/+45a1c7b1ec524cee/edit/
	Updates in local time	Berlin	http://fahrinfo.vbb.de/bin/help.exe/dn?L=vs_mobilitymap&tpl=fullmap&tabApp=show
Temporal	Archival display	Columbus	http://curio.osu.edu/
archive	Control over archival timing	Plume Labs	https://air.plumelabs.com/en/
	Temporal benchmarking	Oberlin	https://buildings.com/s/oberlincity/storyboard314/?chapterId=1395
	Compare sites in archive	Dublin	http://dubindashboard.ie/EnvironmentTransport/stats

Table 1. Major visualisation challenges and the strategies which dashboards have used to try and resolve them

3.1 Refresh and indicators of realtimeness

Real-time data visualisation methods necessitate dynamic displays in dashboards. A dynamic display reflects the evolving nature of urban data and social structure through continuous updating, something a static display does not do or only does periodically. An overly dynamic display might reduce the capacity to observe meaningful change if new data replaces old data at undesirable times and change over time is not displayed. Most importantly, dynamic displays leave users vulnerable to change blindness. Here, we are concerned with strategies for reducing the effect of change blindness in dynamic displays (see Table 2), as well as the positives and potential negative implications of each strategy.

<u>Method</u>	<u>Pros</u>	<u>Cons</u>
Manual refresh	New data on demand	Lack of realtimeness
Automatic refresh	No user input necessary	Change blindness
Indicator of refresh type	No unnecessary interaction	Display clutter
Time since last update	Data currency	No indication of next data
Time until next update	Preparedness for data replacement	Depends on server structure
Indicator of refresh rate	Data frequency	Display clutter
Updates in clock time	No data translation necessary	No indication of ‘now’
Updates in local time	Coincides with local function	Sourced external to browser

Table 2. Visual strategies of display refresh and indicators of refresh.

3.1.1 Manual vs automatic refresh

The dynamic nature of the data, and particularly its presentation, can be controlled through refresh, and its realtimeness communicated to users. In rare cases, real-time data integration

is implemented such that real-time data is compiled external to the display and only integrated at the command of a manual refresh action taken by the user (See Figure 1). Here, data is only updated at the instigation of the user, enabling them to control the replacement of data on the display and to concentrate their perception to detect changes. However, preliminary conversations with users of several dashboard designs indicates that users prefer updating to occur passively and automatically, despite the heightened risk of change blindness inherent in automatically refreshing displays. The effect of change blindness cannot be removed from a dynamic display, but strategies can be used to increase the comprehension of changing visual presentation. Manual incorporation of new real-time updates is not necessarily a solution to change blindness, since display effects persist regardless of how new data is integrated as a result, for example, of a temporarily blank display and the overwriting of outdated data. Manual refresh does, however, allow for a degree of preparation prior to an interruption to the viewing process and provides a user-defined amount of time by which to observe the current state of the display. Additional control over the temporal aspect of the display would provide greater comprehension of change through non-linear navigation of the temporality of the data (Harrower & Fabrikant, 2008), although importantly, the sense of real-time presence may be lost.

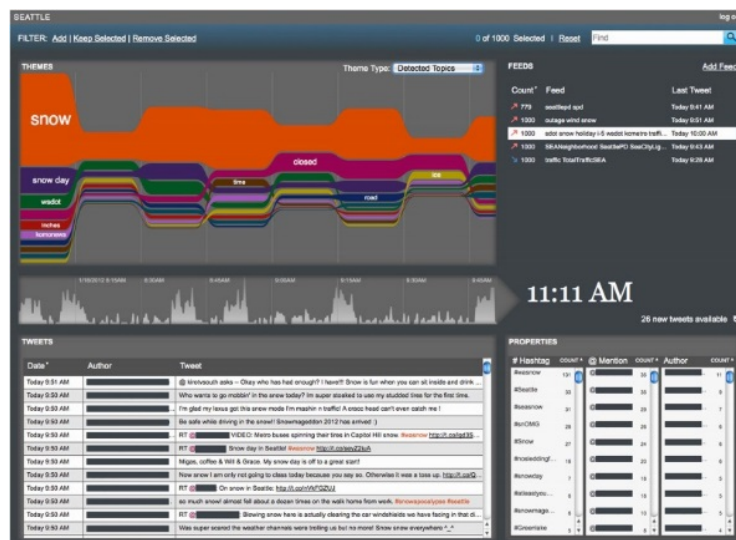


Figure 1: Screen capture from Best et al's dashboard built for the city of Seattle. Manual update is implemented via clickable text next to the current time display. Original image from (Best, et al., 2012).

Automatically-refreshing displays, on the other hand, are less predictable, not containing an interactive component to initialise a refresh. Automatic updates are a function of timed information flows, sensor calibration, and time required for visual comprehension of the display, which restrict a fully on-demand data integration strategy. The city dashboard has the primary function of showing the status of the city it represents via a set of information which is constantly changing and may become out of date if a user is required to anticipate new information and change the display manually. In order to maintain a real-time display, visualisations must automatically update when new information becomes available. This is especially true of IoT sensor networks which operate by reporting new updated measurements to a server only when updates are available (Hunkeler, et al., 2008). Automatic updates ensure that a dashboard operates as closely to real-time as possible given the time from sensor measurement to server compilation, and rendering of the visualisation. The risk of change blindness is enhanced by this rapidly-adapting display, but other visual methods can be implemented to reduce that risk.

Regardless of the method used to incorporate new data into the display, the state of the visualisation at any given point in time will appear equivalent, therefore it is important that the display indicate its data refresh style. Aside from any obvious changes in the display when incorporating new real-time data, the display should indicate the refresh type – whether it is automatic or manual. Above all else, this ensures that a viewer is aware of their role in changing the content of the page. A manually updating display without the indication to perform the necessary action results in a static, non-real-time display, while a manual refresh action on an automatically updating page results in frustration at the lack of change when

action is taken. Necessarily, a manually refreshing visualisation must include a clickable object which initiates the display refresh. Automatic indicators may be textual, explaining the automatic nature of the update or the expected occurrences of new updates. It may also appear as a visual suggestion, such as representative or flashing icons, as the Human fitness tracker's live activity graphs use (for example, see Human's activity in Tokyo, <http://cities.human.co/details/Japan/Tokyo>).

3.1.2 Indicators of refresh timing

Interaction with the visualisation can be performed without the aid of visual indicators but the techniques discussed here illustrate that many displays are enhanced by providing additional information about the realtimeness of the data. Indicators of the refresh timing, where it is out of the control of the user, as in automatically updating dashboards, serves the purpose of indicating the realtimeness of data while also potentially reducing change blindness by preparing users for impending updates. Although change blindness was still measured in studies of users who knew that a change was to occur (Fish, et al., 2011), efforts to recognise patterns are impossible when updates cannot be anticipated.

A time since last update visual indicator above all else provides a notification of the age or extent to which the data being viewed is out of date. In a real-time display, this is critical, as data quickly becomes out of date, as in the case of real-time public transit tracking. Time since last update is not sufficient in itself to identify a manual or automatic refresh style, since a user would need to observe the display for a length of time to determine if the visualisation was going to transition. For example, see Toronto's live map of bus locations with updates given in a seconds since the last reported location (Figure 2).



Figure 2: Screen capture of the Toronto transit dashboard and MarineTraffic dashboard. A popup for each bus and ship displays the time since last update in accumulated time.

A time until next update method is much less frequently utilised in real-time dashboards and cannot be used with manual updating since update frequencies are user-defined in manual refresh dashboards. This strategy was only implemented in one illustrative dashboard. The London Dashboard uses a countdown method with second precision to alert viewers of the next display refresh for each separate module on the dashboard (see Figure 3). The consistent countdown makes for a very dynamic and potentially distracting display, but reduces the interaction necessary to understand the displayed data. Interaction is, in fact, discouraged in this view except on specific items regarding transit suspensions and headlines of news articles. The next update method also gives a very strong indication of the realtimeness of the data and the sampling used to generate real-time urban data.

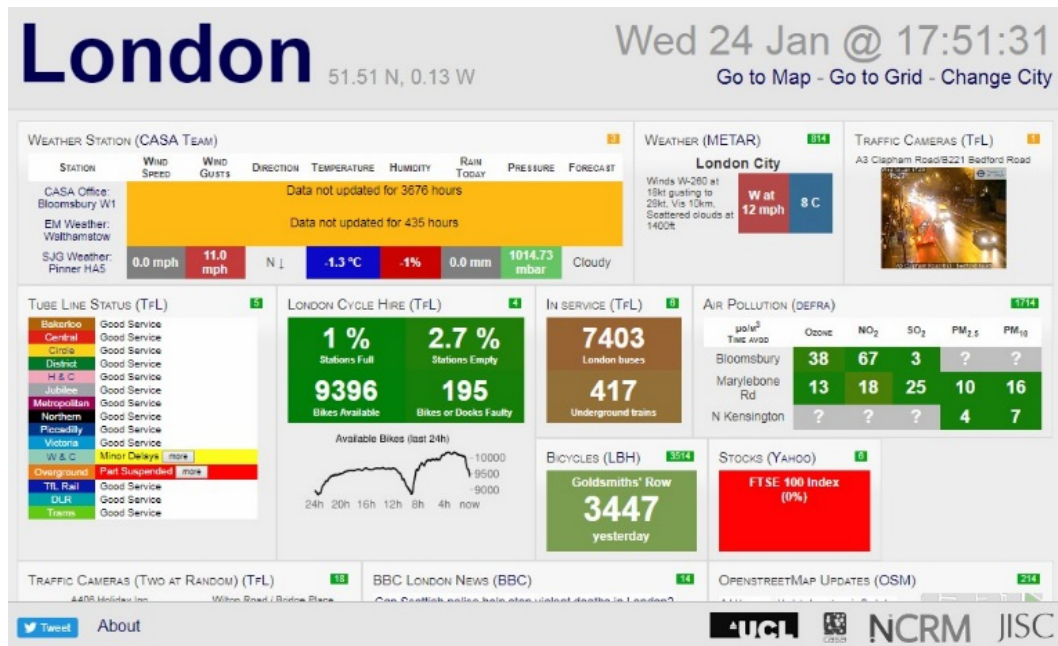


Figure 3: Screen capture of the London dashboard. The counter in the top-right corner of each individual module is dynamic, counting down each second until the data in the module refreshes.

The advantages and disadvantages of indicators of the last and the next update are complementary. Time since last update indicates the age of real-time data and the extent to which it is out of date – its reatimeness. However, it cannot provide indication of impending updates or the frequency of updates without prolonged observation of the display to observe precisely when the indicator changes. Conversely, a time until the next update indicator provides information about the next planned update to the display, but contains no information about how long the existing data has been displayed. In indicating the time of the next update, a display can more prepare a user for expected changes, thus reducing, but not resolving, the issue of change blindness. As a front-end to data servers, which are often independent of the dashboard itself, the visual components of dashboards have no control over when updates are communicated between sensors and servers and from servers to their public APIs. Thus, time until next update may not be possible in many dashboards.

The singular methods of update indicators alone provide no information about the refresh rate of the dashboard visualisation. Without an indication of the refresh rate, it might be assumed that the various modules on the dashboard possess the same realtimeness, whereas different types of real-time data vary significantly in their temporality, for example the less frequently updated weather information on the London dashboard. The timing of the data visible is a critical part of the metadata necessary for inferences from observed patterns. However, refresh rate can be gleaned from the combination of last update and next update indicators. And as a standalone textual indicator, refresh rate may introduce unnecessary clutter to the display.

3.1.3 Timing indication with clocks

Although a fairly innocuous difference, dashboards are divided on the question of reporting update indicators with clocks or accumulated time. Accumulated time is shown in Figure 2 of the Toronto transit dashboard. Previous updates are displayed as seconds since the previously known bus location, rather than given by the time of day. Clock time is simpler to create, as IoT sensors report precise times along with new measurements. Additionally, the display will not require frequent animation to change an indicator displaying accumulated time, which changes as time passes. However, unless the viewer has a keen sense of time, clock time updates require them to compare their display to an external clock to ground the observation against the actual time. For this reason, most dashboards include a clock on their display to indicate current time. Most users can interpret an accumulated time indicator as a function of their own embodiment of time, and it provides a better sense of realtimeness as it changes in reference to the present moment.

An inconsistency exists among dashboard designs with regards to displaying clock time in local time or the user's own time zone. For example, in Figure 4, where the Boston and Berlin transit dashboards were captured at 12:17 GMT, local to the authors and indicated

by the clock widget, the clock time of the last update should indicate the local time of 07:17 EST (GMT -5) for Boston and 13:17 CET (GMT +1) for Berlin. While Berlin displays the correct time, the Boston display indicates an impossible time which has not yet occurred because it is based on the user's time zone. There should be no assumption that dashboard users only engage with dashboards within the time zone of the data which is represented. Information about transportation systems and schedules make little sense outside of the local time of the city, but a user-centric design would cater the display to the reference point of the viewer. A simple implementation would extract a timestamp from the browser client, local to the dashboard user, but would require the user to then translate to the local time zone of the dashboard.

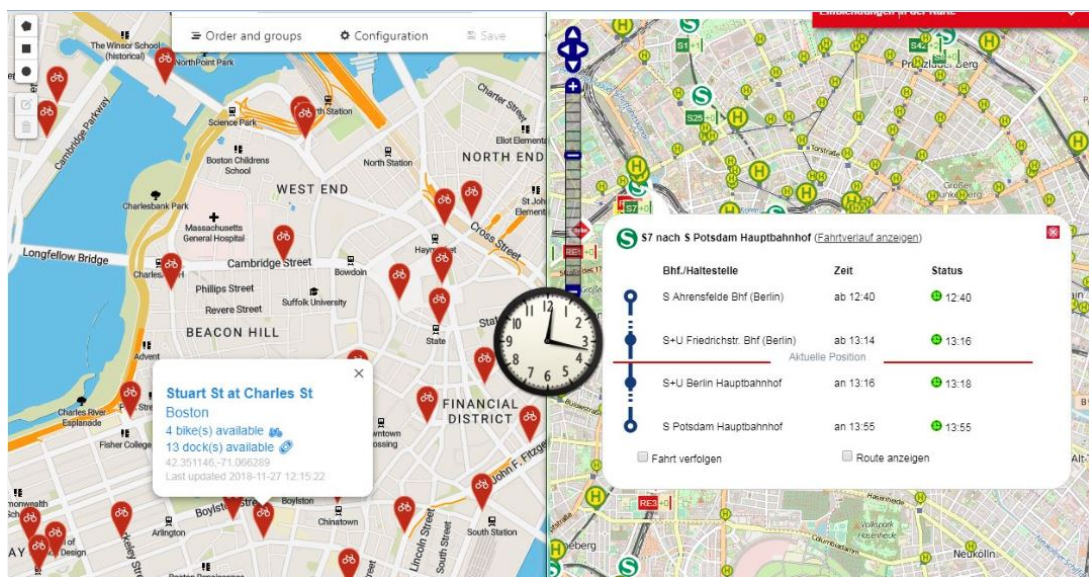


Figure 4. Screen capture of the Boston (left) and Berlin (right) transit dashboards. Local time indicators are consistent with user time of 12:17, despite different time zones.

3.2 Visual analytics of archived real-time data

Archival real-time displays capture real-time information at regular intervals over time and maintain previous measurements for retrieval and visualisation. Although previously

observed real-time data becomes out of date in rapidly-changing displays quickly and storing longitudinal information is not a trivial task, the added context of previous measurements is useful for analyses of changing cities and for contextually understanding current observations. Table 3 highlights a few strategies used by dashboards and the advantages and disadvantages that they provide with respect to archival real-time data, each of which is then explained in greater detail.

<u>Method</u>	<u>Pros</u>	<u>Cons</u>
Archival display	Long-term patterns	Data archive
Control over archival timing	Trend and seasonality visibility	False patterns
Temporal benchmarking	Intuitive comparison	Interpretability
Compare sites in archive	No interactivity necessary	Display clutter

Table 3. List of visualisation methods which enable archival real-time data exploration, with primary pros and cons.

3.2.1 Archival display

While spatial displays are ideal for showing differences across geography at a particular time, they are often ill-suited for showing time-series data as each refresh erases the previously viewable pattern. As such, many dashboards also employ temporal techniques to show trends in measurements over time. Because time-series visual methods do not animate through time, this method displays changes while minimising the effects of change blindness. In some cases, these time-series are relatively short in time frame, perhaps over the past hour or day. In other cases, the time-series form a longitudinal archive of data that can be examined over different timeframes, such as week, month, and year. Many dashboards consider archival data

important in establishing trends, discovering rhythms of regular variation in the city, and providing contextual information for interpreting present real-time data. However, significant technical implementation is necessary to build an archive to store and recall the time-series generated by prolonged data capture. Additionally, the display of longitudinal data requires new visual designs which may be incompatible with existing interfaces and require some specialised knowledge to interpret.

Visually, archival time-series plots are easily interpretable as horizontally-oriented linear displays. For example, the smart wifi dashboard of the Moncloa Campus of International Excellence in Madrid (Alvarez-Campano et. al 2017) uses an archival time series view to compare the number of current real-time wifi users to those of the last seven days (see Figure 5). Patterns are detectable in this visual representation from daily events (the weekend is obvious with fewer people using wifi on campus), sub-daily patterns appearing which correspond to the regular work day, and smaller regular variations within that.

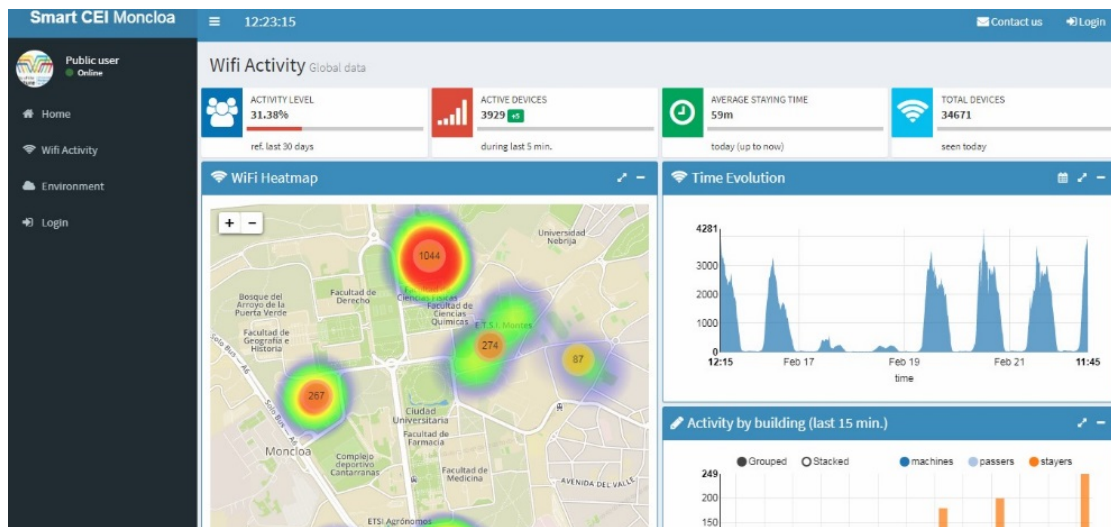


Figure 5. Wifi dashboard for the Moncloa Campus of International Excellence.

3.2.2 Control of archive timing

Those dashboards that incorporate archival comparison have a widely variable range of time scales on which the archival data is displayed. The CEI Moncloa wifi dashboard uses a weekly archival scale and uses data with fifteen-minute precision, providing no means to view patterns which occur outside of these parameters. In other cases, dashboards provide methods to control the timing of the archive to which the real-time information is compared. Visual analytics research has considered interactive ways to adjust a temporal display to account for temporal patterns in multiple scales. By controlling the intervals of daily historic hotel visits, Weaver et al (2007) illuminated patterns of traveling salesperson behaviour that were not visible in the typical seven-day weekly-repeating display by making the temporal display interactive. Similar methods in city dashboards facilitate exploration of temporal patterns at various and irregular intervals.

In city dashboards, the ability to control the timing of archives is important because of the prevalence of non-cyclical patterns in social activity and infrastructure. Most social activity on a city scale is somewhat consistent – daily patterns in traffic flows, for example, show typical rush hour patterns when large volumes of vehicles move in similar directions. But temporal scales of multiple days create inconsistencies which are visible in simple cases involving weekends and holidays. Archival comparison that does not consider these expected temporal variations are insufficient analytical tools. Although the ability to change the visible extent of data can help users discover new periodic patterns, it also provides more control to find spurious patterns. A relevant periodicity for one piece of archival data may indicate only randomness in another, so pre-set controls may be misleading at the same time that they facilitate new exploration.

Dashboards which provide the capability to control temporal framing provide users with greater control over the comparisons that they can make between real-time and archival

data. The Plume Labs world air map displays pollution levels in major cities by using both real-time and archival records over the course of hours, the most recent week, most recent month, and most recent year. The example of Istanbul is shown in Figure 6. The interactive display allows the viewer to trace back along the time series, changing the numerical display to reflect the selected time.



Figure 6. Screen capture of Plume Labs' city air quality dashboard with Istanbul selected. Tabs at the top allow for changing between current, weekly, monthly, and yearly scales of archival comparison to real-time measurements of air quality.

3.2.3 Temporal benchmarking

The city dashboard for Oberlin, Ohio utilises a similar anthropomorphic animation to help represent water and electricity usage throughout the city. Its more innovative data visualisation strategy is the additional feature of time-series benchmarking comparison, displaying the most recent full temporal cycle (i.e., the previous day's data) (see Figure 7), as well as a "Typical use" series based on longitudinal data collection. Benchmarking is an important task in city management for tracking progress toward goals and is not impossible with a simple, well-labelled time series plot. For example a weekly time series plot allows for comparison between real-time data and the data from the previous day at the same time of

day. But by superimposing the last full day's cycle alongside today's, the Oberlin example also allows for trend comparison in a way that requires no special ability to decipher a time series diagram.

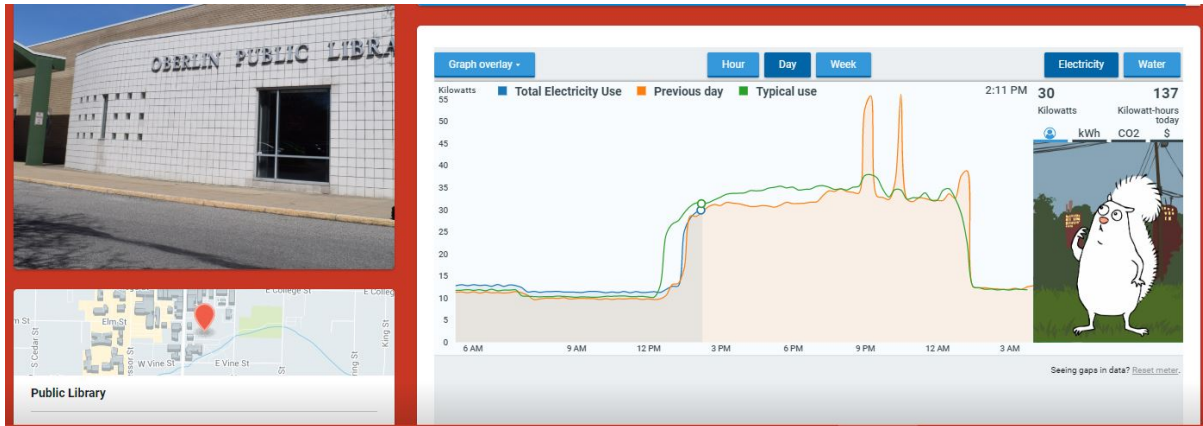


Figure 7. Screen capture from the Oberlin city dashboard's City Water Flows module, with daily real-time data selected. Real-time data is displayed, along with archived data from the pre day and the previous day's time series.

3.2.4 Comparison of multiple sites' archives

In addition, some dashboards facilitate the exploration of multiple sensors' archival records simultaneously. In the methods observed so far, individual time series have been compared against real-time data for the purpose of understanding patterns in a single sensor or spatial scale's measurements. In this task, not only are the patterns in individual time series important to visualise, but displays must also facilitate the ability to compare those patterns in the data collected from other sensors. Comparison is made optimal by the simultaneous viewing of archives (Andrienko, et al., 2009; Steiger, et al., 2014).

Display space is a limiting factor for comparing multiple archived time series. Archival display frequently uses linear representations of time, which demand greater display space to produce an interpretable visual presentation. Thus, to compare archival information from multiple sensors, spatial organisation is typically removed in dashboards in favour of

multiple stacked or overlaid time series charts. So visual comparison of these sensors' time series can optimise display space and comparability by plotting them on the same chart, as in the 'Tweets %', 'Air Quality Index', 'Operating Busses', and 'Bus Delays' modules of the Curio dashboard of Columbus, Ohio, created by the Center for Urban and Regional Analysis (CURA) at Ohio State University (Xiao, 2017) in Figure 8a. This strategy may be unusable when the time series have different ranges, necessitating a separation of the display into individual plots, as in the Dublin dashboard (McArdle and Kitchin, 2016) in Figure 8b. Even though archival analysis is the primary objective of these comparisons, real-time information is still typically a component of them. In the linear displays of the Curio dashboard, the current real-time values are situated at the end of the horizontal lines beginning near the plot's origin. The Columbus dashboard provides limited interactivity to compare the plotted values at selected points in time along the plot, including the most recent, real-time figure, which affords greater analytical potential than simply observing general trends.



(a)



(b)

Figure 8 (a). Screen capture of the Curio dashboard. The ‘Air Quality’, ‘Tweets (%)’, ‘Operating Busses’, and ‘Bus Delays’ modules show separate sensors on the same Cartesian plot. (b) Screen capture of the Dublin dashboard. Time series are split into separate displays.

4. Principles of dashboard design

The observations we have made about displaying spatio-realtime data and facilitating real-time and archival temporal analysis reveals a lack of standardised principles for dashboard design, though despite their non-uniformity, visualisations are largely based on established fundamentals of temporal visualisation. We identified two challenges that shaped and differentiated spatio-realtime dashboard tools: the dynamic nature of the data, how they are refreshed, and how the realtimeness of the data is communicated to the user; and how the tool enables archival analysis. The variety of spatio-temporal methods employed is indicative of the sources and forms of data generated by smart cities. However, the methods we have

documented have different strengths and weaknesses that seem to have been little evaluated in the design and building of some dashboards.

Dashboards, as the front end to a variety of both official and other data which can be informative for day-to-day urban tasks such as commuting, are important tools for communicating city data to both decision makers and casual data explorers. Thus, these design principles facilitate interpretability as well as advanced pattern analysis, and provide a novel effort to establish much-needed standards of real-time dashboard design.

4.1 Refresh and realtimeness

Among existing city dashboards, the preference is for automatically refreshing visual tools. Automatic refresh dashboards better reflect the always-on nature of smart city technologies than the requirement to manually update data. Automatic refresh is able to adapt to data availability in such a way as to require minimal interaction to obtain a real-time perspective of city function. The disadvantage of an automatic refresh rate is a greater likelihood of change blindness being introduced. Consequently, there must also be a method put in place to communicate the key properties of refresh rate, last update, and next update.

Although Fish et al. (2011) still measured change blindness in subjects who were aware that a change was occurring in a spatial display, we maintain that communication of the display's refresh and the realtimeness of the data currently on display enables users to anticipate and focus on data transition. A dynamic display which has a multitude of individual dynamic elements (such as the London dashboard, Figure 3), as opposed to a single updating feature, is vulnerable to greater degrees of change blindness due to the difficulty of monitoring multiple elements (O'Regan, et al., 1999). Thus, a single indicator of planned updates to the display is necessary for user preparation for that update.

Although we observed no dashboards with indicators of both last and next update time, both are necessary to begin to counteract the effects of change blindness. With both

indicators, a viewer can prepare for an impending update as well as know the data refresh rate through their combination. Thus, no explicit indicator of refresh rate is necessary. Using both indicators also allows for displaying data with different temporal granularities, as in the case of sensors with non-uniform activation rates. Each sensor must thus have its own indicator of previous update contained within a popup, which reduces display clutter by keeping the accompanying text behind interactive components.

Finally, with both refresh indicators in use, a combination of clock time and self-referential accumulated time make for an ideal display. Time until next update indicators should involve a countdown in regularly-decreasing time to emphasise the currency of the data and the display's reactivity. The animated nature of a consistently-updating indicator of the next update ensures that users will not perform unnecessary manual interaction with the assumption that their actions only will create a change. A time of last update indicator provides no advantage by being dynamic, as with an accumulating time indicator, thus the time of last update should be recorded to the display as the local time at the moment when the data refreshed. Clock time indicators and digital clock readouts are critical to real-time dashboards, but only if the clock appears in local time and not in the time zone of the viewer. Sensor times, scheduled public transport, and metadata should all be displayed in the local time of the dashboard, so their display does not appear temporally independent of associated activities.

Human perception reveals that change blindness is an issue in dynamic displays, and we acknowledge that further testing is necessary to measure the impact that it has on real-time interpretation. We also have shown that example methods in existing dashboard environments may address some of the issues associated with change blindness and its associated loss of pattern comprehension. Our categorisation of visual indicators of realtimeness contributes to a body of literature testing strategies for interactive data

exploration among varying user types (Roth, et al., 2015), but which has not yet been applied to the open city dashboard. Since public engagement and interpretation are primary goals of interfaces to the data generated by smart cities (Degbelo, et al., 2016), indication of the realtimeness of that data is critical to the openness and usability of dashboards.

4.2 Archival time series

Although not all dashboards which utilise real-time data also make use of archived records of those data, all that prevents a long term comparison between real-time and archival information is data storage. The advantages of archiving real-time data are vast for the purpose of trend analysis, benchmarking current conditions against previous ones, establishing regular patterns of city activity, and being able to create predictive models. Archival methods are of little public utility where the long term pattern is designed and remains fairly rigid, such as in GPS traces of public transportation routes, but long term measurement data is invaluable for understanding urban patterns. With fairly simple additions to temporal visualisation strategies, archival displays facilitate a range of useful tasks, so we suggest that archival visualisation methods be used whenever the data exists to populate them.

Even simple archival time series with real-time data at one end of a trend line (e.g., the Columbus dashboard in Figure 8) overcomes the issue of change blindness with respect to losing sight of previous data. The addition of temporal benchmarking against which to compare records with other times (e.g., yesterday, week, month, quarter, year) provides a means to evaluate the progress of city initiatives and track change at meaningful intervals. Because a previous data cycle will not largely vary in scope from real-time measurements, adding a second series to a linear time series plot, as the example from Oberlin City in Figure 7 does, is a straightforward process and one that does not introduce significant complexity to the interpretation of the visualisation. Simplified sparklines can provide an effective means to

allow comparison between real-time and recorded data, but there must be an option to show temporal granularity to enable contextual interpretation.

Ideally, users should be allowed to customise the temporal scale of any benchmark comparison. Again, as Plume Labs and Oberlin City demonstrate, changing the duration of the time series on display no further complicates the visualisation while allowing new patterns to emerge from the different time scales being compared. This selectivity is particularly important in urban temporal pattern analysis, where some functions are better compared at weekly or seasonal intervals than daily ones, or against specific policy-related target dates. Short term changes and trends can guide decision making at the city scale, while long term patterns are useful for comprehensive analysis of sustainability and growth.

To that point, urban data contexts may require nonstandard date range comparisons in the form of bespoke archival data query. Current monthly or quarterly durations restrict the ability to examine a date range to only those data which coincide with the dates provided. No observed dashboards have existing methods to create precise queries by date range, which would facilitate exploration of data with variable periodicities and enable specific benchmarking comparisons (Auer, et al, 2011, MacEachren, et al. 2011).

4.3. User-focussed design

Effective visualisation design often follows an iterative design principle, where feedback from users at multiple stages in the design and implementation process guides new iterations of the interface (Roth, et al., 2015). Importantly then, dashboard evaluations must consider how users with different expectations of dashboard communication and experience using visual analysis interpret and make meaning from spatio-temporal visualisations (McArdle & Kitchin, 2016). Dashboard design begins to diverge considering distinctions between user types and the anticipated users of dashboards. Internal dashboards, like those which provide interfaces to data in urban control rooms (Kitchin, et al., 2015; Mattern, 2015), are used by

what might be termed ‘power-users’ who have well developed data literacy. Public facing dashboards, however have a more diverse and less specialised user experience base who have varying levels of ability to interpret visual analytics.

User experience presents a challenge to dashboard design because overly complex visualisations will leave inexperienced and novice users unable to derive meaning in the way the display intends. Similarly, an oversimplified display risks alienating experienced users whose needs are not facilitated through basic controls. We consider the challenge of facilitating knowledge discovery through visual design for the range of users from novice to advanced to be an important factor in evaluating display effectiveness. Although visual methods have been proposed for the purpose of comparing historical time series on a spatial map, including compressed, space-saving methods (Heer, et al., 2009) and nonlinear temporal glyphs (Auer, et al., 2011, Fuchs, et al., 2013), the complex interactions and visual comprehension necessary to interpret such methods may preclude their usefulness for inexperienced users.

Many of the proposed solutions to temporal visualisation challenges are fairly complex and require some degree of advanced interpretation. As city dashboards are public-facing compilations of visual tools, advanced users are not the only anticipated users of the system. Innovative solutions to the challenges of temporal visualisation may not be viable ones if they require specialised knowledge or instruction to be used as intended. Many methods which have been proposed in visual analytics and interactive cartography literature may not satisfy the requirement of facilitating interpretation by users inexperienced with data visualisation. We believe that dashboard visualisation strategies present a unique challenge to visual analytics design by aiming to facilitate inexperienced users first. As complexity and fear of interactivity are barriers to visual analytic use by novice users (Wisniewski, et al.,

2009), visual indicators of realtimeness aim to reduce the complexity of real-time data such that the barriers to using dashboards are reduced for inexperienced users.

5. Conclusion

As dashboards become stable, public-facing gateways to technology-mediated smart cities, these principles – automatic refresh and communicating that property to viewers through indicators and making use of archival data in analytical and intuitive ways – provide, we believe, a valuable reference for creating usable data exploration tools in browsers and smartphone apps. Although modes of interactivity change how such systems are designed on different platforms (e.g., desktop, tablet, smartphone), the dashboard concept and its tools for pattern discovery and other temporal analysis are transferable.

Additional research is necessary to explore the utility of different visual methods for recognising particular urban temporal patterns. The challenges of change blindness and archival data exploration require user testing with the goal of evaluating interpretability, insights which lead to decision making, and desire to continue using the dashboard in the ways that it is intended. Critically, further research in these areas must acknowledge the multiple nature of dashboard users which includes casual users (e.g., tourists) who seek quick, current information at a glance and are less advanced at interpreting complex displays, as well as long-term decision-makers with high stakes in interpreting and obtaining comprehensive, actionable information from visualisations. Further research and testing should strike a balance between interpretability and informativeness and between familiarity and creativity to create more effective dashboards.

Dashboards remain an emerging technology, though they have quickly become vital tools for making sense of the smart city, particularly when implemented for public consumption. Dashboards are decision-making tools at both authoritative and casual levels,

and where there is a space to use them to improve public participation and data literacy in cities, design must leverage intuition and aesthetic to be usable. Improving dashboard effectiveness will ensure that more value and insight can be extracted from the spatio-temporal data of smart city technologies and increase efficiencies in the management of cities.

Acknowledgement

We gratefully acknowledge funding from Science Foundation Ireland under the Investigator's Award Program. Award number: 15/IA/3090

Works Cited

- Alvarez-Campano, M., López, G., Vázquez, E., Villagrà, V., & Berrocal, J. (2017) Smart CEI moncloa: An iot-based platform for people flow and environmental monitoring on a Smart University Campus. *Sensors*, 17(12), 2856
- Andrienko, G., Andrienko, N., Rinzivillo, S., Nanni, M., Pedreschi, D., & Giannotti, F. (2009). Interactive visual clustering of large collections of trajectories. *Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on*, (pp. 3-10).
- Auer, T., MacEachren, A. M., McCabe, C., Pezanowski, S., & Stryker, M. (2011). HerbariaViz: A web-based client–server interface for mapping and exploring flora observation data. *Ecological Informatics*, 6, 93-110.
- Bach, B., Dragicevic, P., Archambault, D., Hurter, C., & Carpendale, S. (2014). A review of temporal data visualisations based on space-time cube operations. Eurographics Conference on Visualization, June 2014, Swansea, Wales, UK
- Batty, M., Axhausen, K. W., Giannotti, F., Pozdnoukhov, A., Bazzani, A., Wachowicz, M., Bazzani, A., & Portugali, Y. (2012). Smart cities of the future. *The European Physical Journal Special Topics*, 214, 481-518.
- Behn, R. D. (2014). *The PerformanceStat potential: A leadership strategy for producing results*. Brookings Institution Press, Ash Center, New York.
- Best, D. M., Bruce, J., Dowson, S., Love, O., & McGrath, L. (2012). Web-based visual analytics for social media. *AAAI ICWSM SocMedVis: Workshop on Social Media Visualization (AAAI Technical Report WS-12-03)*, (pp. 2-5).
- Brath, R., & Peters, M. (2004). Dashboard design: Why design is important. *DM Direct*.
- Coletta, C., & Kitchin, R. (2017). Algorithmic governance: Regulating the ‘heartbeat’ of a city using the Internet of Things. *Big Data & Society*, 4.

- Cybulski, P., & Medyńska-Gulij, B. (2018). Cartographic Redundancy in Reducing Change Blindness in Detecting Extreme Values in Spatio-Temporal Maps. *ISPRS International Journal of Geo-Information*, 7, 8.
- Degbelo, A., Granell, C., Trilles, S., Bhattacharya, D., Casteleyn, S., & Kray, C. (2016). Opening up smart cities: citizen-centric challenges and opportunities from GIScience. *ISPRS International Journal of Geo-Information*, 5, 16.
- de Lange, M. (2018). From real-time city to asynchronicity: Exploring temporalities of smart city dashboards. In *Time for mapping: Cartographic temporalities*, ed. S. Lammes, C. Perkins, A. Gekker, S. Hind, C., Wilmott, and D. Evans. Manchester, UK: Manchester University Press. pp. 238-255.
- Draper, G. M., Livnat, Y., & Riesenfeld, R. F. (2009). A survey of radial methods for information visualization. *IEEE transactions on visualization and computer graphics*, 15, 759-776.
- DuBois, M. J., & Battersby, S. E. (2012). A raster-based neighborhood model for evaluating complexity in dynamic maps. *Proceedings of AutoCarto, Ohio, USA*, 1-16.
- Dubriwny, D., & Rivards, K. (2004). Are you drowning in BI reports? Using analytical dashboards to cut through the clutter. *DM Review*.
- Dykes, J. A., & Mountain, D. M. (2003). Seeking structure in records of spatio-temporal behaviour: visualization issues, efforts and applications. *Computational Statistics & Data Analysis*, 43, 581-603.
- Fabrikant, S. I., Rebich-Hespanha, S., Andrienko, N., Andrienko, G., & Montello, D. R. (2008). Novel Method to Measure Inference Affordance in Static Small-Multiple Map Displays Representing Dynamic Processes. *The Cartographic Journal*, 45, 201-215.
- Few, S. (2006). *Information dashboard design: The effective communication of data*. Sebastopol, CA: O'Reilly.

- Fish, C., Goldsberry, K. P., & Battersby, S. (2011). Change Blindness in Animated Choropleth Maps: An Empirical Study. *Cartography and Geographic Information Science*, 38, 350-362.
- Fuchs, J., Fischer, F., Mansmann, F., Bertini, E., & Isenberg, P. (2013). Evaluation of alternative glyph designs for time series data in a small multiple setting. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, (pp. 3237-3246).
- Goldsberry, K., & Battersby, S. (2009). Issues of Change Detection in Animated Choropleth Maps. *Cartographica*, 44, 201-215.
- Gray, S., Milton, R., & Hudson-Smith, A. (2013). Visualising real-time data with an interactive iPad video wall. *Talisman*.
- Guo, D., Chen, J., MacEachren, A. M., & Liao, K. (2006). A visualization system for space-time and multivariate patterns (vis-stamp). *IEEE transactions on visualization and computer graphics*, 12, 1461-1474.
- Harrower, M. (2007). The cognitive limits of animated maps. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 42, 349-357.
- Harrower, M., & Fabrikant, S. (2008). The role of map animation for geographic visualization. *Geographic visualization: concepts, tools and applications*, 49-65.
- Heer, J., Kong, N., & Agrawala, M. (2009). Sizing the Horizon: The Effects of Chart Size and Layering on the Graphical Perception of Time Series Visualizations. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1303-1312). New York, NY, USA: ACM.
- Heim, M. (1994). *The metaphysics of virtual reality*. Oxford University Press on Demand.
- Huggins, R. (2010). Regional competitive intelligence: benchmarking and policy-making. *Regional Studies*, 44, 639-658.

- Hunkeler, U., Truong, H. L., & Stanford-Clark, A. (2008). MQTT-S — A publish/subscribe protocol for Wireless Sensor Networks. *2008 3rd International Conference on Communication Systems Software and Middleware and Workshops (COMSWARE '08)*, (pp. 791-798).
- Kitchin, R. (2017). The realtimeness of smart cities. *TECNOSCIENZA: Italian Journal of Science and Technology Studies*, 8, 19-42.
- Kitchin, R., Lauriault, T. P., & McArdle, G. (2015). Knowing and governing cities through urban indicators, city benchmarking and real-time dashboards. *Regional Studies, Regional Science*, 2, 6-28.
- MacEachren, A. J., Jaiswal, A., Robinson, A. C., Pezanowski S., Savalyev, A., Mitra, P., Zhang, X., & Blanford, J. (2011). SensePlace2: GeoTwitter analytics support for situational awareness. *2011 IEEE Conference on Visual Analytics Science and Technology (VAST)*, (pp. 181-190).
- MacEachren, A. M. (2017). Leveraging big (geo) data with (geo) visual analytics: Place as the next frontier. In C. Zhou, F. Su, F. Harvey, & J. Xu (Eds.), *Spatial Data Handling in Big Data Era* (pp. 139-155). Springer.
- Mackenzie, A. (2007). Protocols and the Irreducible Traces of Embodiment: The Viterbi Algorithm and the Mosaic of Machine Time. In *24/7: Time and Temporality in the Network Society*, edited by R. Hassan & RE Purser, 89--108. Stanford: Stanford University Press.
- Mattern, S. (2015). Mission control: A history of the urban dashboard. *Places Journal*.
- McArdle, G., & Kitchin, R. (2016). The Dublin Dashboard: Design and Development of a Real-Time Analytical Urban Dashboard. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*. 3. 1st International Conference on Smart Data and Smart Cities. 7-9 Sept. Split, Croatia.

- Monmonier, M. (1990). Strategies for the visualization of geographic time-series data. *Cartographica*, 27, 30-45.
- O'Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change-blindness as a result of 'mudsplashes'. *Nature*, 398, 34.
- Peuquet, D. J., Robinson, A. C., Stehle, S., Hardisty, F. A., & Luo, W. (2015). A method for discovery and analysis of temporal patterns in complex event data. *International Journal of Geographical Information Science*, 29, 1588-1611.
- Pezanowski, S., MacEachren, A. M., Savelyev, A., & Robinson, A. C. (2018). SensePlace3: a geovisual framework to analyze place--time--attribute information in social media. *Cartography and Geographic Information Science*, 45, 420-437.
- Rensink, R. A. (2002). Change detection. *Annual review of psychology*, 53, 245-277.
- Robertson, G., Fernandez, R., Fisher, D., Lee, B., & Stasko, J. (2008). Effectiveness of Animation in Trend Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14, 1325-1332.
- Roth, R. E., Ross, K. S., & MacEachren, A. M. (2015). User-Centered Design for Interactive Maps: A Case Study in Crime Analysis. *ISPRS International Journal of Geo-Information*, 4, 262-301.
- Silva, S. F., & Catarci, T. (2000). Visualization of linear time-oriented data: a survey. *Web Information Systems Engineering, 2000. Proceedings of the First International Conference on*, 1, pp. 310-319.
- Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *perception*, 28, 1059-1074.
- Simons, D. J., Franconeri, S. L., & Reimer, R. L. (2000). Change blindness in the absence of a visual disruption. *Perception*, 29, 1143-1154.

- Steiger, M., Bernard, J., Mittelstadt, S., Lucke-Tieke, H., Keim, D., May, T., & Kohlhammer, J. (2014). Visual Analysis of Time-Series Similarities for Anomaly Detection in Sensor Networks. *Computer graphics forum*, 33, pp. 401-410.
- Thakur, S., & Hanson, A. J. (2010). A 3D visualization of multiple time series on maps. *Information Visualisation (IV), 2010 14th International Conference*, (pp. 336-343).
- Weaver, C., Fyfe, D., Robinson, A., Holdsworth, D., Peuquet, D., & MacEachren, A. M. (2007). Visual Exploration and Analysis of Historic Hotel Visits. *Information Visualization*, 6, 89-103.
- Weltevrede, E., Helmond, A., & Gerlitz, C. (2014). The Politics of Real-time: A Device Perspective on Social Media Platforms and Search Engines. *Theory, Culture & Society*, 31, 125-150.
- Wisniewski, P. K., Pala, O., Lipford, H. R., & Wilson, D. C. (2009). Grounding geovisualization interface design: a study of interactive map use. *CHI'09 Extended Abstracts on Human Factors in Computing Systems*, (pp. 3757-3762).
- Xiao, N., Fontanella, S., Miller, H.J., Adair, M., & Mount, J. (2017) An Interactive Dashboard for Visualizing Big Spatio-Temporal Data in an Urban Area. *Geocomputation 2017*, Leeds, UK