Built Environment. Special issue on 'Big Data'. London, United Kingdom: Alexandrine Press, Volume 42, Number 3, Autumn 2016, pp. 474-497.

Wise Cities: 'Old' big data and 'slow' real time

FABIO CARRERA
Worcester Polytechnic Institute, Interdisciplinary and Global Studies Division

'Big Data' typically refers to large datasets, mined in bulk from modern electronic devices, that can be analyzed to extract patterns of behavior at both the macro and micro level. More often than not, big data is derived from major systems, such as social networks or telephone providers, or from other sensor networks that compose the big data firehose. Humans have also been recruited to act as sentient sensors that can contribute richer and finer-grained crowdsourced data to this big pile. Nevertheless, big data is not only found in the form of (near) real-time information extracted through reality mining. While the activities that take place in a city are fast and mobile, the immobile structures that make up the physical city evolve at a much slower pace, creating a sedimentary layer of 'old big data' that is possibly more useful -- certainly as useful -- for city planning as the information derived from real-time sensing.Using our efforts in Venice (Italy) to record the historical evolution and current state of the city through many kinds of data, we show how smart-city approaches can be used to capture the backlog of data that predates the current age of the internet, and also to intercept gradual structural changes to the built environment that happen in 'slow real time'. Such techniques are applicable to all cities, and especially to world heritage sites like Venice where we can accumulate longitudinal big data that can gradually transform a smart city into a wise city.

Since its appearance, in the beginning of the twenty-first century, the 'big data' moniker has evolved into a veritable meme (Lohr, 2012). Despite losing some of its luster of late (Bean, 2014), it has been making steady inroads into the mainstream of today's technology *Zeitgeist*. The primary connotation of Big data refers to the sheer size of the datasets, and in the early days the phrase was used to describe the data-mining techniques invented by the Human Genome Projects to sort through the myriad of patterns in DNA (Howe *et al.*, 2008). More recently, the term has acquired a strong association with *real-time* data collected automatically from the web (Eagle and Pentland, 2006): from networks of cameras, sensors, turnstiles and card readers; from online websites and social networks; or in general from smartphone users connected to internet or cellular providers (Ratti *et al.*, 2006). In common parlance, the term big data is used broadly to refer to vast troves of bits, gathered quickly and automatically, from which we can extract useful information about patterns of behavior. Nevertheless, as urban scholars, we should remind ourselves that some data is big but not new, and that old data can become very big, if we collect it and organize it using modern data mining and crowdsourcing techniques.

Smart cities are beginning to use big data to some advantage, by analyzing real-time streams that measure such activities as daily commutes, financial transactions, social relationships and other human events (Batty *et al.*, 2012). In comparison, much less attention is being paid to urban change that happens in 'slow real time', namely the gradual creation and modification of a city's physical structures, mostly governed by municipalities through building permits, inspections, approvals, licenses and the like. These slow but steady changes have created the current built environment where today's real-time big data is produced and collected, yet, ironically, these sedimentary layers of old big data are almost never organized and analyzed using traditional big data techniques.

This paper illustrates some of the innovative approaches that were employed to collect, organize and re-circulate 'old big data' about the city of Venice collected by cohorts of engineering students from Worcester Polytechnic Institute (WPI) over many years. We hope this will inspire other cities and world heritage sites to adopt similar techniques to mine their own deep-rooted repositories of past information and to intercept slow urban change as it happens in piecemeal fashion day after day.

Structures and Activities: Physical Context and Digital Content

One useful way of thinking about our cities -- smart or not -- is to divide them neatly into two halves as did Italo Calvino in his *Sophronia* (Calvino, 1972): the ephemeral half-city of movement and activities and the solid half-city of the physical structures. This dichotomy represents 'the content and the context' of city life (Carrera, 2004), which Kevin Lynch (1984) called 'the spatiotemporal distribution of human actions and the physical things which are the context of those actions'. Even though today's big data technologies are almost exclusively applied to activities, we argue here that they could (and should) also be applied to physical structures.

The endless streams of digital content mined from the web or from other networks are frequently mapped against their physical contexts, which are deemed immutable at the time-scale of today's instant data mining, but are in actuality changing as well, albeit much more slowly and almost imperceptibly. The material makeup of the city is taken somewhat for granted, but it is itself a consequence of past activities which are less about the short-term daily use of the city and more about the long-term creation of the city's physical infrastructures through structural city-making activities such as land development, construction projects, and conversions of extant buildings. When a citizen goes to city hall to apply for a permit to build a house in a vacant lot, he or she is creating a single transactional data point somewhere in some municipal database. This single event may seem unremarkable vis a vis the myriad of activities and transactions that happen in the vicinity of that vacant lot on a daily basis: from phone calls to internet logins, from Facebook posts to Instagram uploads, from vehicular traffic to transit commutes, from car accidents to crime, from theatre shows to music concerts, and much more. However, in short order, this one boring permit event will engender a new building, which in turn will change the physical configuration of this block for a long, long time, and affect all activities in the area as well. This single transaction will create a new context for future activities and the cumulative effect of many such apparently insignificant transactions is what created the here in which we all live now.

Mapping the present-day layout of cities has been the focus of the past two decades of Geographical Information System (GIS) efforts at the municipal, regional and national levels. Despite these ongoing endeavors, detailed GIS layers of each city are not always accessible to outsiders. OpenStreetMap, which is crowdsourcing the creation of a single, GIS-grade, world map thanks to the efforts of tens of thousands contributors, is the major example of big data about the physical infrastructure that comes to mind. However, it only portrays a limited subset of the layers of urban elements that make up the city. Furthermore, it does not capture urban change as it happens by intercepting government processes -- such as a new construction permit -- but relies instead on *ex-post* updates after the building has been constructed. Moreover, online mapping services like Google Maps or OpenStreetMap capture mainly the geometry of urban change, but not the attributes associated with such change, which is where the real big data lies.

One of the arguments of this paper is that a smart city can become wise by coupling the here-and-now of G.I.S. with fast real-time data and also with the there-and-then of past records about modifications of the built environment over time. Big data approaches to mine activities as they occur should be applied not only to everyday patterns of human behavior, but also adopted to capture the more mundane trickle of events and transactions that signal the slow but steady change in the fabric of our cities.

When we collect dynamic data about human activities, we should always be cognizant of the exact configuration of the physical context at the time of the data collection. In Venice, we manually gathered boat traffic data between 1998 and 2003, and the methods we developed were later adopted by the municipal government who continued to conduct regular boat counts until 2009 as we show in (Figure 1).

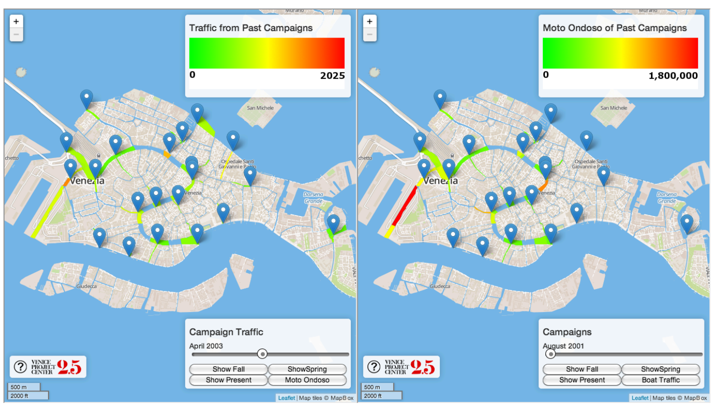

While hundreds of trained data collectors were conducting dozens of day-long field campaigns, some of the city's canals were undergoing regular maintenance and therefore were closed to traffic. We knew about the closures at the time of the boat counts, yet we did not save a precise map of the network situation on the exact days of the campaigns, foregoing a great opportunity to record patterns of behavior around closed canals, and in fact missing the chance to measure driver adaptation in the real world. Our data was skewed by these closures. We may have encountered a higher-than-usual number of boats on a specific canal segment simply because an adjoining segment was closed for maintenance. This mistake of ours is probably very common, since it is unlikely that government agencies or traffic consultants, as part of their standard data collection procedure, routinely keep a record of the road network condition at the time of their traffic counts. Yet this approach would be obviously useful, if not necessary, in the long run.

This is just one of the many examples that we have worked with in Venice, where a snapshot of activity in time should have been associated with the corresponding snapshot of space at that same time. We have grown so accustomed to separating our research on physical structures from our research on activities that it does not even occur to us that recording one without the other would negatively affect the quality of the data collected. We make the mistake of thinking of structures as fixed and immutable backgrounds just because they change so slowly and gradually. Yet maintenance and construction activities that are constantly taking place around physical structures always affect other human activities in that area, temporarily as well as permanently. We should accept that there is no such thing as a 'steady state' in the built environment of our cities, and collect and organize our data accordingly.

Traffic is a great example of an activity which is totally affected by structures, since structures (e.g. roads) physically restrict and guide driving behavior in major ways, but also because structures (e.g. buildings) are occupied and used by people who in turn generate traffic. The construction of new homes not only adds new children to the school system, increases water supply requirements, requires larger water treatment capacity, but also generates extra vehicular and pedestrian traffic, by bringing in more patrons to local restaurants and to public events. New construction always affects and tends to increase, human interactions in one way or another. Structures and activities are inextricably intertwined and we should teach ourselves to record the relevant structural context every time we study the activities at a particular location.

The long tail of old data

Intuitively, it makes logical sense that there is more big data about the past than about the present, just as there are cumulatively more people who have died before us than there are living people on earth today. The past is therefore where the 'big' big data really lies. A lot of old big data is sitting in some archive waiting to be unearthed and digitized, and some may be literally buried underground in the form of archeological artifacts. But a lot of old big data is already in electronic form, stored on various hard drives in local government offices. Lately, in the wake of the burgeoning open data movement, more and more government records about the recent past are now being posted on the web. Additionally, all of the data that is produced by automatic devices ages quickly and therefore continuously enriches the old data pile, day after day.

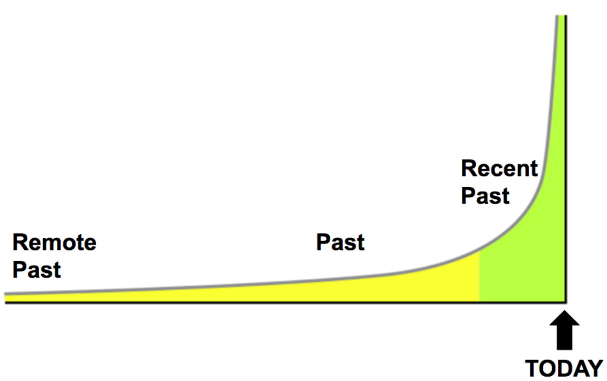

Old data can be schematically arrayed in a long tail that stretches backwards from today and thins out as it recedes to the remote past as shown schematically in Figure 2.
The exponential curve resembles the one that describes network effects in complex adaptive systems and it is the mirror image (if we rotated it along the Y axis) of the typical long-tailed depictions of rank-frequency distributions.

The 'old data' power curve starts in the remote past and ends with a high spike at today. It precipitously declines like a hockey-stick as it moves backwards in time. Like all power laws, we can hypothesize that it is Pareto-distributed according to the typical 80-20 rule, which means that the terabytes of data collected in the 16 years since the start of the new millennium are probably matched by the total number of records available from the last century, yet the sum of all data points from the distant past is larger than all the data points collected recently, no matter what time span we use to delimit recent and past events. Organizing data about the pre-internet era is perhaps harder than it is to intercept streams of real-time data today, and it gets progressively harder to access and organize data the farther we dig into the past. Yet we need to address old data systematically if we want to be wise users of city knowledge.

In short, old data is big data and we should begin to treat it as such. It is big in terms of the number of records but perhaps even more so when it comes to its attributes. Old data can surprisingly be deeper and richer than new real-time data because each old record enshrines lots of attributes accrued over time and from different sources, whereas new real-time big data is often shallower, perhaps consisting of more records, but with potentially fewer attributes which are typically derived from a single data source.

The Long Tail of Old Structures

Much time is spent by researchers and professionals to gather rich information about specific characteristics of the built environment in an area of interest, but unfortunately very little of this hard-earned data is re-usable by the next researcher studying the same area. Old big data about structures can be quite rich and it is often redundantly collected -- in *ad-hoc* fashion -- for very impellent needs, often without a strategic systematic approach. Some 'old big data' are found on government open-data web sites, yet, despite a recent focus on open data, these government datasets are still few and far between. So, old data is still being re-exhumed periodically by scholars and professionals, only to go back into obscurity again shortly thereafter.

In historical cities like Venice, information about the existing infrastructure can come from many sources: from published books, studies or reports; from government records; from medieval manuscripts; and even from archeological reports about the ancestral layers that underlie the current city. Since many ancient buildings are still part of today's built environment, the need for detailed information about our heritage goes beyond the mere collection of historical or architectural trivia, but requires a skillful technical approach in order to adequately support present-day maintenance, management and planning activities concerned with the urban infrastructure of historic places.

Since 1988, teams of WPI students at the Venice Project Center (VPC) have conducted extensive inventories of all the major urban infrastructure in Venice, including: canals and canal segments (Chen *et al.*, 2011), bridges (Bossalini *et al.*, 2013), docks (Bennett *et al.*, 2013), streets (Flaxington *et al.*, 2015), wellheads (Blackwell *et al.*, 2000), fountains (Kelley, *et al.*, 2004), clocks (Chakuroff *et al.*, 2014), bell towers (Baruffi *et al.*, 2012), bells (Bove *et al.*, 2015), churches (Dechaine *et al.*, 2012), palaces (Fletcher *et al.*, 2002), convents (Heinricker *et al.*, 2013), public art (Elbag *et al.*, 2003) stores (Bruso *et al.*, 2012), hotels (Connor *et al.*, 2015), and many others (see veniceprojectcenter.org). Every dataset about structures includes attribute fields that reflect the 'permanent' features of each urban element, such as its date of construction, or its address, or its material, its dimensions, and others. But a 'wise' dataset also should keep track of information about more 'dynamic' features, such as the object's state of conservation, which changes continuously and can only be captured with time-stamped snapshots.

An example of an application that uses modern big data techniques to capture information about urban elements from past centuries is PreserVenice (http://preservenice.org) which employs crowdsourcing and crowdfunding techniques to catalog and preserve the over 7,500 public artifacts that decorate the streets and canals of Venice (Ascare *et al.*, 2010).

The long old tail of these decorative urban elements stretches back over one thousand years, and it took over a decade for our project teams to map out the whole collection, which includes: coats of arms, statues, wellheads, fountains, street altars, flagpoles, inscriptions, reliefs and numerous other types of public art of various shapes and sizes. It was a prolonged and time-consuming process, but it is now completed. We have mapped the long tail of public art in Venice, and there will be very few modifications to be made in the future. The big data has been collected and now all we have to worry about is the current condition of each artifact to preserve the most 'endangered pieces' before they deteriorate irreparably (Carrera, 1997).

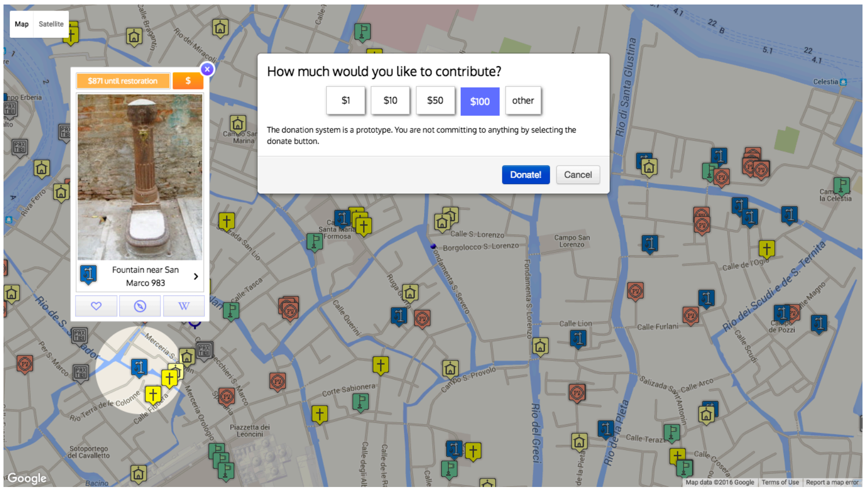

The PreserVenice app, funded by UNESCO, employs advanced data management techniques, borrowed from the world of complexity science, incorporated into our City Knowledge (CK) platform, a proprietary technology that assigns software agents to each element of the built environment. These software agents act as 'guardian angels' that oversee and mediate the interactions between each public art object and the mobile app (Carrera, 2012), which in turn collects micro-donations to preserve the pieces, by soliciting contributions while the donor is face-to-face with the artifact as shown in Figure 3.

PreserVenice is an example of an app that deals with old data yet uses modern big data techniques to crowdsource the collection of updated condition information and to crowdfund the preservation efforts. There are lots of opportunities to apply similar approaches to the big old data backpile in every city, and there are more big old data in heritage sites than anywhere else.

The Long Tail of Old Activities

Activities that happened in the past, when collected at all, were probably only used for a brief period of time and were then probably forgotten after the issue-at-hand has been dealt with. Because of their ephemeral nature, past events are typically harder to mine and organize for future reuse, making it hard to conduct longitudinal comparisons of activities over time. Past activities that involve human behavior can be divided into two big categories:

1. Anonymous events, such as vehicular traffic or transit commutes;
2. Individual transactions generated by citizens, government or businesses.

Some metrics of mass behavior -- such as passages at toll booths or turnstiles -- actually consist of frequent and numerous transactions, but they are typically analyzed anonymously and in the aggregate, thus they belong more to category 1 than 2. The rest of big behavioral data in the first category is collected automatically by sensors and counters, which capture events rather than transactions, as in the case of traffic volumes, which are routinely collected using pneumatic tubes or induction loops unbeknownst to the individual drivers.

When it comes to past human activities, it is typically easier to backtrack through logs and records of individual (category 2) transactions, for licenses, permits, deeds and the like, than it is to reconstruct and reuse massive event-based datasets from the past, which are rarely publicly available and require an understanding of the methods used to collect the data for apple-to-apple comparisons which alone can allow to accurately track the evolution of a phenomenon over time. An example of a big-data approach to old transaction records is the application called uScript (Cafferelli *et al.*, 2007), an online tool for the cumulative digitization of historical documents which is shown in figure 4. uScript crowdsources the efforts of paleographers, who manually transcribe ancient manuscripts as part of their research (see Figure 4 and Cafferelli *et al.*, 2007), and then uses these manual transcriptions to learn how to interpret automatically vast troves of historical records of government transactions from the Middle Ages and before (Carrera, 2005). uScript's ability to tag people's names, dates, currencies, place names, professions and the like, makes the entirety of the archive eminently searchable in many dimensions. uScript is a concrete example of modern techniques applied to old data. The widespread utilization of a tool like uScript would guarantee that all of the history of the world would be digitized after some time, making it the historical equivalent of the biological genome projects of the 90's. uScript is a big data application for historical knowledge, just like the human genome project was a big data application for bio-genetic knowledge.

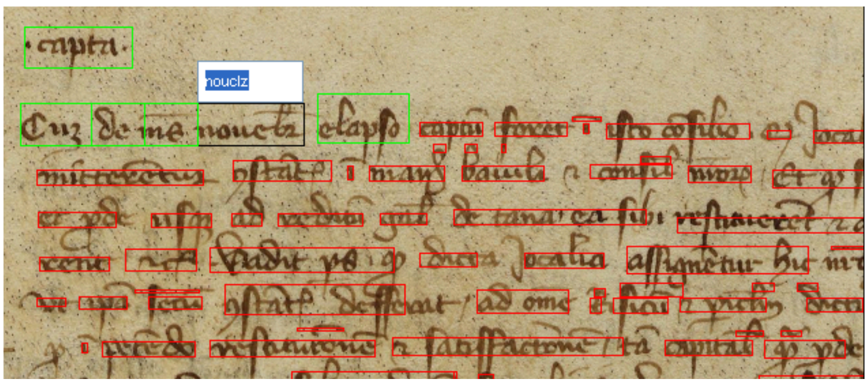

The rest of this paper explores how both structures and activities are composed of old and new data, which can be collected and consumed at rates that are fast or slow, depending on the situation. In the sections that follow, the term 'new' refers to data produced today and more generally includes all data about the present, whereas 'old data' refers to data records about past events which may be still relevant to today's situation, but may also provide a comparative record of change over time for both structures and activities.

Slow old data: catching up with the backlog of history

Given the preponderance of old data vs. new, one would expect that the backlog of past information would have been farmed and reused with the same systematic mindset that has been applied to the harvesting of big data about real-time urban activity metrics. But unfortunately that is not the case.

The cities we live in are the result of years of piecemeal changes by countless individuals and institutions, both private and public. For practical present-day reasons, we need to know a lot about the physical characteristics of historical structures that survive today, and often we would like to know about their past histories as well. Longitudinal data about structures can be gleaned, even in the absence of documentary records, by decoding the historical information physically embedded into the construction materials and techniques that are layered into heritage sites.

Physical structures change slowly and gradual modifications over long time periods leave visible traces that can help us date the various alterations since their original inception.
Some of these changes were also captured in maps and paintings, which depict buildings and monuments through the ages. In 2013, a Venice Project Center team developed a web platform for the annotation of ancient maps (Chines *et al.*, 2013), which allows the users to highlight structures on the maps, name them and link them to representations of the same structure in other maps in Figure 5, thus creating an interactive portfolio of historical images of the built environment over time (cartography.veniceprojectcenter.org).

Each structure that is identified in this cartography system contains a link to a Wiki page that gives detailed information about that structure. In fact, we have automated the creation of individual Wiki pages for each urban object, which are populated on-the-fly with data from our underlying platform. All of these individual pages about each object of the built environment are found in Venipedia Figure 6 (*Bobel et al.*, 2012), a specialized Wiki which showcases all of the slow old data about structures that we have collected in Venice since 1988 (Finelli *et al.*, 2010).

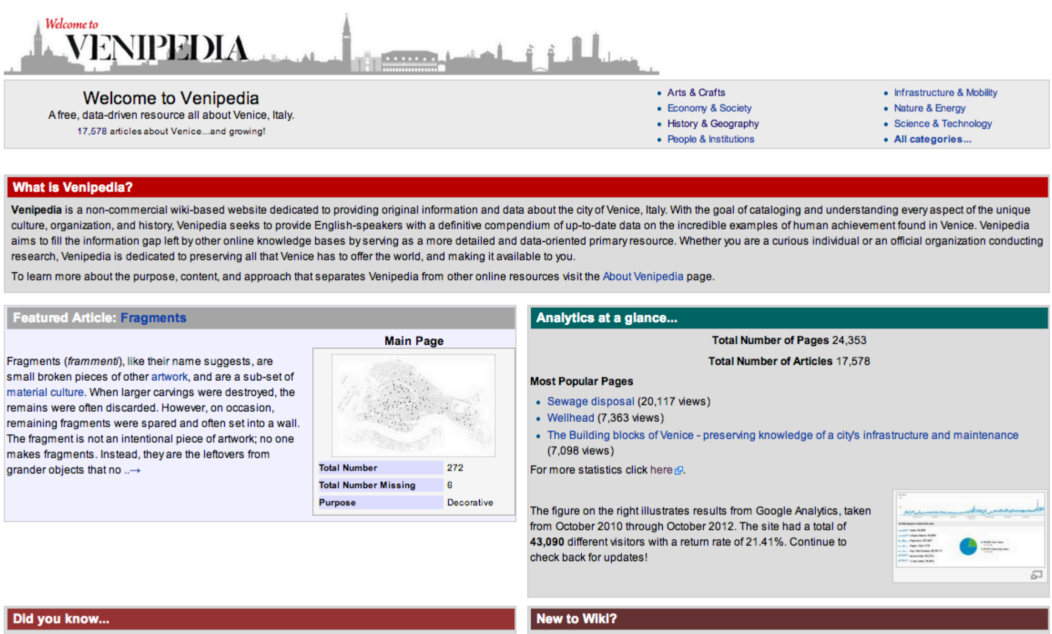

As of the summer of 2016, Venipedia contains over 24,000 pages, most of which are automatically generated from our platform (venipedia.org).

On the wake of these academic projects, we are now developing a scalable commercial product, the City Knowledge (CK) platform, which links all available attributes -- including both current as well as historical information -- to each urban element, using space as the index in figure 7 shows. Akin to what uScript does for ancient manuscripts, the CK platform makes it possible to gradually accrue richer and richer datasets about the built environment. The CK platform is a flexible framework for the organization and cumulative accrual of rich attribute data about the physical elements that make up the city and it is being developed as part of an all-Venetian startup incubator aimed at repopulating the city by providing good-paying jobs for young Venetians who are not interested in pursuing a career in tourism (Carrera, 2016a).
All of the visualizations here utilize the underlying CK platform, which demonstrates the reusability of the system in a variety of contexts as shown in figure 7.

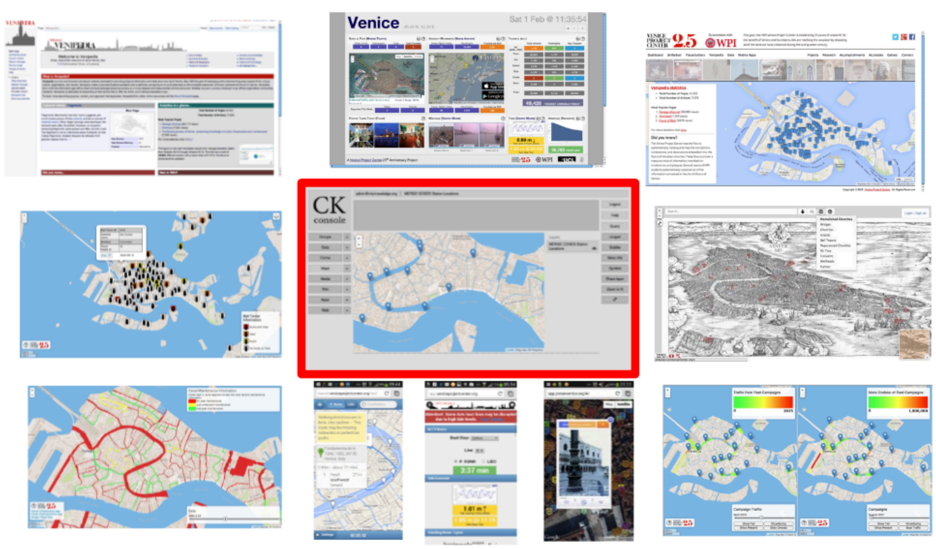

All these projects prove that it is possible to gradually fill-in the long tail of old structures and to preserve and reutilize the resulting big old data for a variety of useful modern-day purposes. What is a lot harder to do is to backtrack and follow the trail of activities that took place in and around those structures over their lifespan. By and large, past records of human activities are not as useful *ex-post* except for comparative purposes or to reconstruct a chain of events that led to today's situation.

The digitalization of old data is never truly over, yet it asymptotically approaches zero as time passes and we stretch backwards into the past. The next section explains what we should be doing to prevent the hunting-and-gathering of past data going forward, by intercepting the old data produced by future change when it happens.

Slow New Data: reality mining in slow real time

Governments are mostly concerned about the present-day *status quo* within their jurisdiction. They keep fairly good records of past permits, deeds, plans, and other documents, but these are not as easily accessible as data that is the daily focus of municipal maintenance, management and planning. While they perfunctorily archive past documentation of the change of 'place over time', cities are more likely to have up-to-date digital records about the 'permanent present' for their day-to-day duties and services, and yesterday's situation is typically replaced with today's changes in a never-ending cycle of constant renewal, which always keeps the present in the foreground. Although GIS and computerized databases are fairly widespread in local governments, the priority is still to cater to today's needs, therefore digital and accessible longitudinal datasets are few and far between.

As explained earlier, the daunting task of attaching rich descriptive attributes to the mapped geometries of physical structures is guaranteed to be attainable over time, in a cumulative manner, since cities are fortunately finite both in time and in space. As old as Venice is, its historical records only go back so far into the past and even the spatial extent of Venice is geographically circumscribed, so the effort of catching up with the backlog of past data is definitely finite in scope, thus achievable in large part.

By adopting a 'reality mining' mindset for structural changes to the built environment, we can avoid creating more backlog for future researchers to sort through, using is a big data approach for the sustainable and perpetual updating of the knowledge base, as new change happens from here on down (Carrera, 2004). Ever since the term 'reality mining' was coined (Eagle and Pentland, 2006), it has been associated with the real-time logging of human activities over some digital network, the moment they occur. This paper simply proposes to extend this concept to structures, instead of relegating it to just activities and behaviors.

Reality mining, when applied to structures, implies a systematic approach above and beyond the gradual and total accumulation of knowledge about the physical elements of the built environment already in existence (slow old data), by simultaneously capturing sporadic structural transactions as they happen in 'slow real time', in order to maintain an accurate and updated record of the physical makeup of the city as it changes -- ever so slowly -- over the years.

The structures that comprise today's built environment have reached their current status after a series of changes punctuated by government permits and other recorded events. Many of these transactions -- which, over time, change the face of the cities in which we live -- are part of the public record and are therefore available for consultation, albeit still mostly in paper form.

More and more, cities are publishing logs of past transactions, like permits and licenses, as part of the government open data movement. For the most part, though, the data is published 'posthumously' -- months, if not years, after the transactions have occurred -- and often as aggregate information, which is of more limited use than would be the 'atomic' record of each transaction. Government open data is updated infrequently and typically represents 'delayed real time', yet it is still useful for tracking urban change over time albeit with some delay.

In Venice, we recently acquired public records,, dating back as far as the early 1900's, from the Chamber of Commerce (LaRovere *et al.*, 2015). The dataset contains records of all economic activities including all of the retail shops in the city. This sort of 'delayed real time' data is extremely useful to show spatial patterns and temporal trends of all Venice shops over time (Figure 8). If these data were made public on a daily basis, the information would be constantly up-to-date and could assist government officials and decision-makers in the formulation of plans to manage, for instance, the spread of bars and tourist accommodations, which subtract retail space and residences from local citizens.

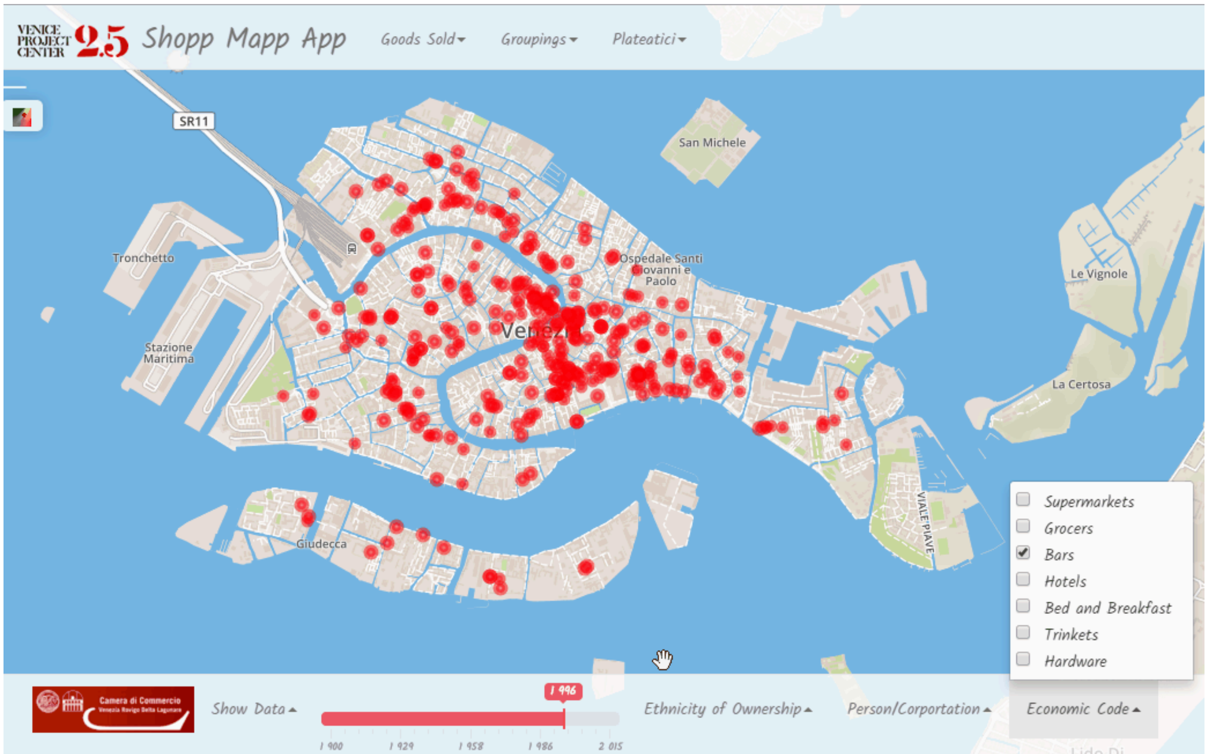

Given the current state of technology, mining the infrequent, yet relentless, trickles of activity that impact the built environment on a daily basis would be entirely feasible by employing some of the same big data techniques used for fast real-time data mining. Using a big-data mindset to intercept structural change at the source --. when the paperwork that precedes construction is filed and approved -- would have significant benefits for the present as well as for the immediate future (Carrera, 2005).

In fact, any predictive model of future patterns of behavior would greatly benefit from incorporating known changes that are bound to manifest themselves in the not-too-distant future, even though at the present they may just exist on paper as approved construction permits or the like. Since these approved projects are going to change the material fabric of the city, any future model should include them as base data about a future that is almost guaranteed to happen, and which will, in turn, affect the rest of future trends all around them. Alas, planning agencies do not routinely include approved urban change in their forecasts, even though the data to do so exists somewhere in a server, under the same roof, in the same municipal government.

The closest our teams have come to incorporating plausible future changes into a planning exercise was an analysis of the maximum build-out that current zoning regulations would allow in Worcester, Massachusetts (Carrera, 2004). As sobering as it is to see one's town spreading out and filling-in with more and more construction, these analyses portray purely hypothetical situations, whereas actually capturing data from development plans (Finzel *et al.*, 2007), environmental impact assessments (Gomes *et al.*, 2005) and actual construction permits or zoning waivers would give us a concrete basis for predicting the probable future.

When it comes to accessing slow new data about structures in (slow) real time, our WPI teams have explored what could be possible if such data were indeed available, but without ever really gaining direct access to it. It is easier for municipal governments to periodically publish these data in batches as static files, than to create machine-accessible endpoints to provide real-time access to their legacy systems, which are typically old enough to make such efforts quite an uphill battle. It is comparatively easier for progressive governments to make available streams of data which employ modern, internet-enabled technologies than to try and jerry-rig an old system to do the same for permit data, but it is a fair bet to expect that things will change in the foreseeable future, as older technologies are replaced with cloud-based systems which in turn will make it possible to harvest urban change in 'slow real time' as well.

Fast Old Data: the impermanence of ephemeral activities

Real time data about activities gets old fast. Yet, even old data about activities can be used in real-time to enable fast choices in the present day. As we become more cognizant of the fact that the real-time data of today quickly becomes the old data of tomorrow, we are confronted with hard choices about what data to keep and what to throw away (cf. Carrera, 1997).

Because data about activities are ephemeral by nature, we have been treating them as essentially disposable. Dynamic, real-time data are typically produced for immediate consumption or for later comparisons as an aggregate statistic but, typically, little thought is given to potential reuse of the same exact data at a later date. Raw data is hardly ever retained or published, and fine-grained information is often lost in the publication of overall aggregate statistics about the phenomenon.

A big data approach to fast-aging, real-time data was employed by a team of WPI students who studied noise at the Venice Project Center (Calamari *et al.*, 2011). The system consists of a mobile app that allows citizens to report noise complaints using a smartphone and an associated website which displays all of the noise complaints as they occur. The app records an audio sample of the noise, calculates the dB level of the recording, and allows the user to take a picture of the alleged source of the noise. The submission of the noise report instantly creates a new datapoint on a companion web site (http://www.venicenoise.org) in real-time. On the website, the data points are immediately processed and a heat map is produced automatically, as are timelapses that can show patterns of noise at each hour of the day, each day of the week and each month of the year.

The ephemeral nature of a noise event is captured in-the-moment, but logs of all noise reports are also maintained and can be searched and visualized over time as shown in Figure 9. Each point is matched to Venice's 'noise zoning maps', and samples that exceed the permitted dB level are highlighted in red on the online the map. The red dots can be clicked to see dB level of the noise complaint and to listen to its recording.

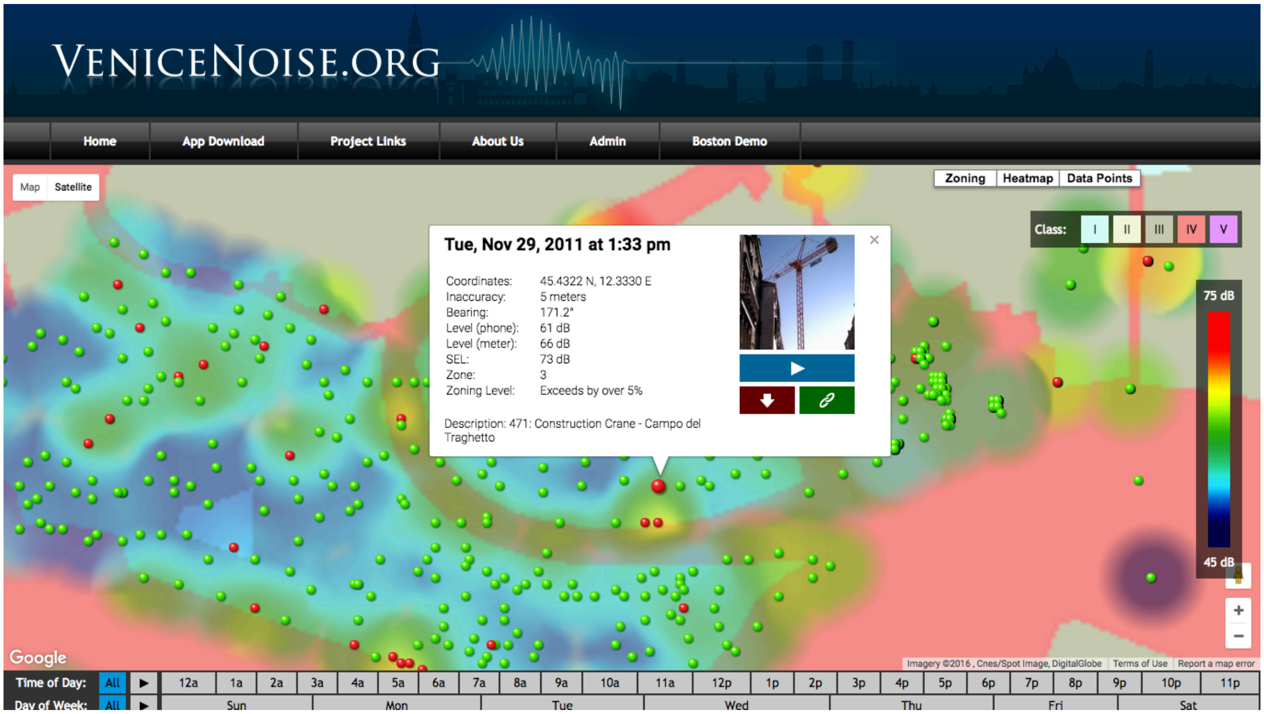

VeniceNoise is an example of a modern app that uses big data crowdsourcing techniques to capture events as they happen, yet it is mindful of the historical importance of these samples over time, so it provides tools to archive, search and visualize noise patterns over time. The broad lesson for all cities -- old and new -- is that, even though real-time data can be useful in the moment, it also gets old almost instantly, so it is important to log it and retain it for future use as VeniceNoise does. The patterns that emerge over time are typically more useful than an instant snapshot, so it pays to be mindful about the long-lasting aspects of any fast data collected in real-time.

Old data often needs to be accessed quickly and repeatedly to support real-time visualizations for a variety of useful purposes. Many VPC applications use old data for fast processes, such as autonomous agent models, or real time smartphone applications. For example, a major dynamic phenomenon that needs to be taken into account in a variety of present-day contexts in Venice is the ebb and flow of tides through the lagoon and the city's canals. The hydrodynamic behavior of currents is important, for example, to determine the 'flushing capacity' of a canal, given that a good percentage of Venice's homes still flush sewage into the city's waterways (Felices *et al.*, 1997).

In the 1990's many teams of WPI students conducted extensive and repeated campaigns to manually the currents in the inner canals of Venice (DeMaio *et al.*, 2012), which allowed us to define the 'typical hydrodynamic behavior' of each canal segment in various tidal and lunar phases (Carrera, 1999). Our data, in turn, was used to design and calibrate a numerical model of the inner canal flows (Umgiesser in Caniato *et al.*, 1999). More recently, a WPI team (Angelo *et al.*, 2014) applied big data principles to our past datasets and models, and produced an interactive visualization which connects the hydrodynamic model to real-time data about current tide levels (from our Dashboard) to produce a tide latency model shown in Figure 10, which allows us to predict with good approximation the tide levels in every canal of the city hours ahead of time.

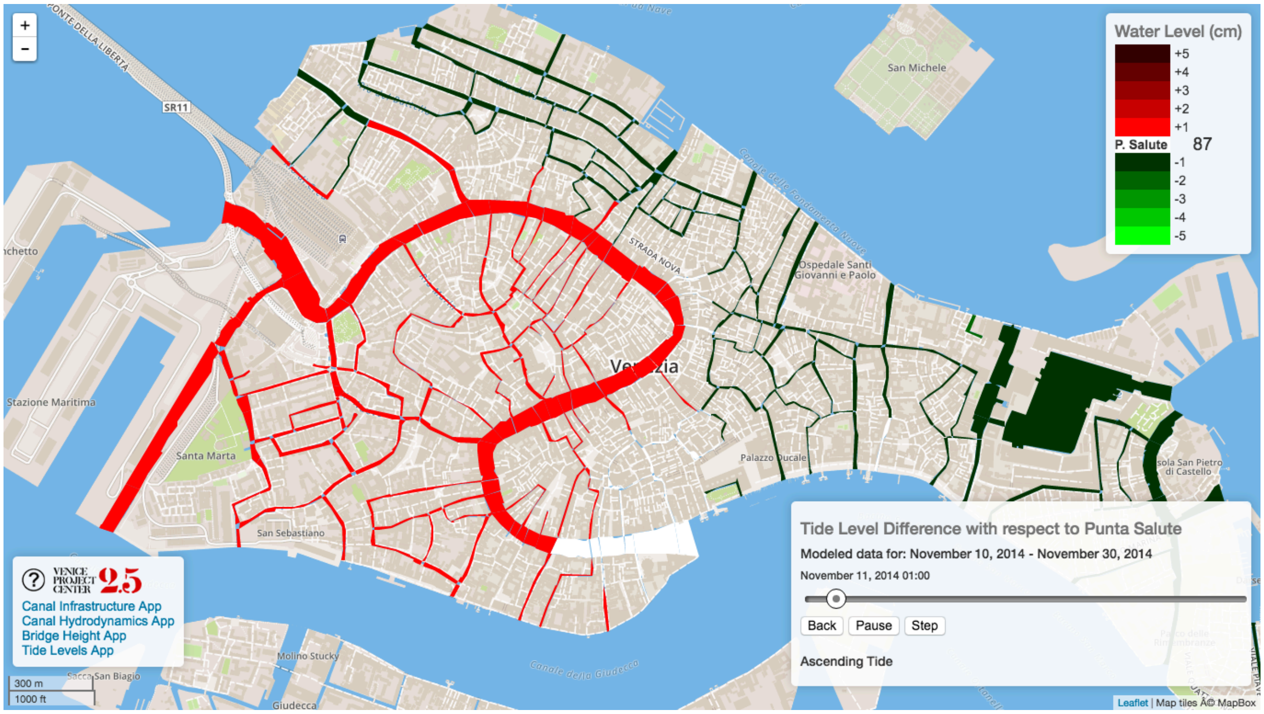

In Venice, one can already download smartphone apps that show the current and forecasted tides in the city, as does the Tides widget in our Venice Dashboard. The cleverest of these existing apps (W.V.F. by ArMa Informatica) basically shows coloured maps of pavement elevation contours, zoomed to your location, based on the current tide at the main tide gauge in Venice (*punta Salute*). While not useless, this 'smart' app is good for the here and now, but not that useful for planning ahead for a trip during a high-tide event.

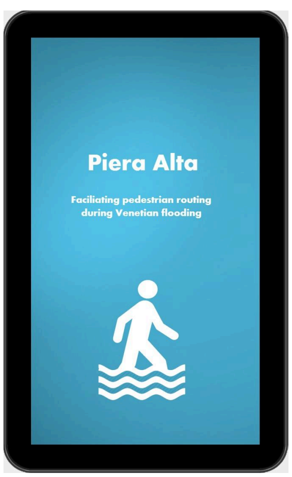

Merging the past with the present and future, we connected the past data to the present tide to the future model, and developed an app prototype called *Piera Alta* (shown in figure 11) which shows how a city can be can go beyond 'smart' and actually become 'wise'.

This smartphone app blends the fast and old to produce a useful predictive tool for the navigation of pedestrian paths in Venice, during high tides (Angelo *et al.*, 2014). It employs sophisticated hydrodynamic models, together with real-time water levels to produce location-specific forecasts of inundation levels and assists in trip planning while the city is being flooded, all well ahead of time. Although it is still in development, *Piera Alta* is perhaps the best example of an app which combines old data together with real time data, and links them both to a future-looking model to go beyond the seemingly smart, and actually provide wise advice that travelers can use to plan their trips during *acqua alta* in Venice.

Applications like *Piera Alta* show what's possible when we apply big data approaches to extract fast information from old data, and we hope that these examples can inspire other cities to evolve from smart to wise with similar big data approaches blending old data about the built environment with real-time events and predictive models.

Fast New Data: Real-time for old cities

Big Data has acquired a broad appeal in the urban technology circles in part because it often focuses on what is fast and new, like commuting patterns in London recorded by the Oyster card system (Batty *et al.*, 2013), or visualizations of phone calls made and received over the cellular network (Ratti *et al*, 2006). While collected in real time by public or private companies, (mostly for accounting purposes), these types of big data are not usually displayed in real time, but frequently represent the post-processing of a day, week, month or year of transaction data, typically obtained from some major network operator, such as a public transit agency or a telecommunications company. The patterns that emerge from these *ex-post* visualizations of big data can be useful to analyze both large and small scale processes, which in turn can lead to better planning or more efficient management of our cities.

Conversely, 'dashboard systems', really harvest and visualize live data streams and other internet data in true real time. While the data instantly published by these dashboards are arguably less 'big' than a year's worth of commute data for a city like London, they can, if logged over time, add up to really big data that can reveal interesting aggregate patterns of behavior.

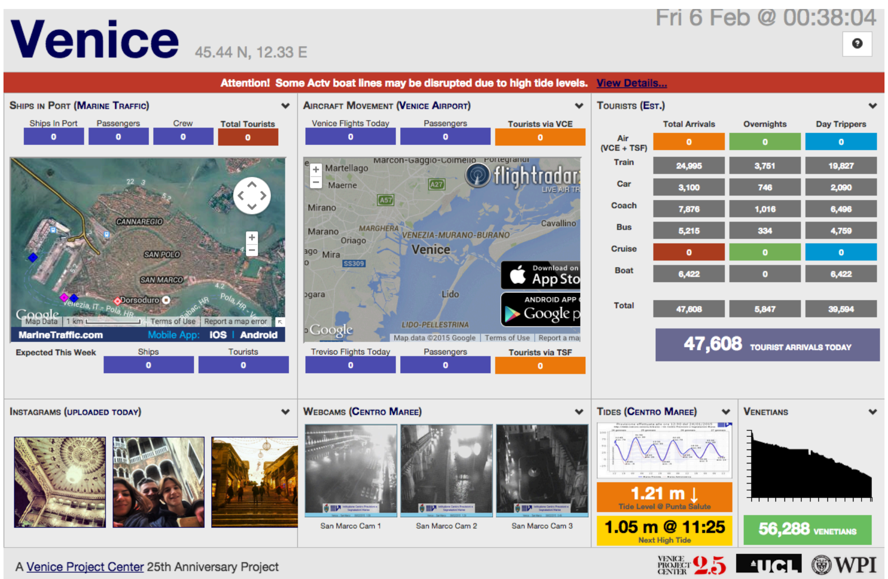

In Venice, we have implemented a real-time dashboard based on the initial efforts conducted at the Centre for Advanced Spatial Analysis (CASA) at University College London (UCL), adapting the type of data collected and visualized by the 'widgets' that compose the dashboard to the peculiar needs of an old heritage city which we show in figure 12 (Brann *et al.*, 2013).

While Venice has many of the same needs and characteristics of all modern cities, such as daily commuter traffic, air pollution and the like, we chose to focus our Venice Dashboard (http://dashboard.veniceprojectcenter.org) on specific phenomena that affect Venice in very unique ways, namely the 'dual tides' of water and people which flood the city and negatively affect the normal functioning of the city.

For the water tides, our dashboard taps into the city's tidal forecasting website to provide real-time updates every 5 minutes of the current tide level, using color codes to signify moderate to severe high-tides. The Venice Dashboard even shows a red banner at the top of the page to warn of possible disruptions to the waterbus service, when the levels reach certain critical thresholds that make some of the bridges too low for waterbuses to pass under, as indicated in figure 12 . Another series of widgets collect data about the human tide of tourist arrivals by airplane, cruise ships and trains, using custom-designed widgets that constantly harvest internet data, such as FlightStats, Maritime Traffic and Railroad reservation systems.

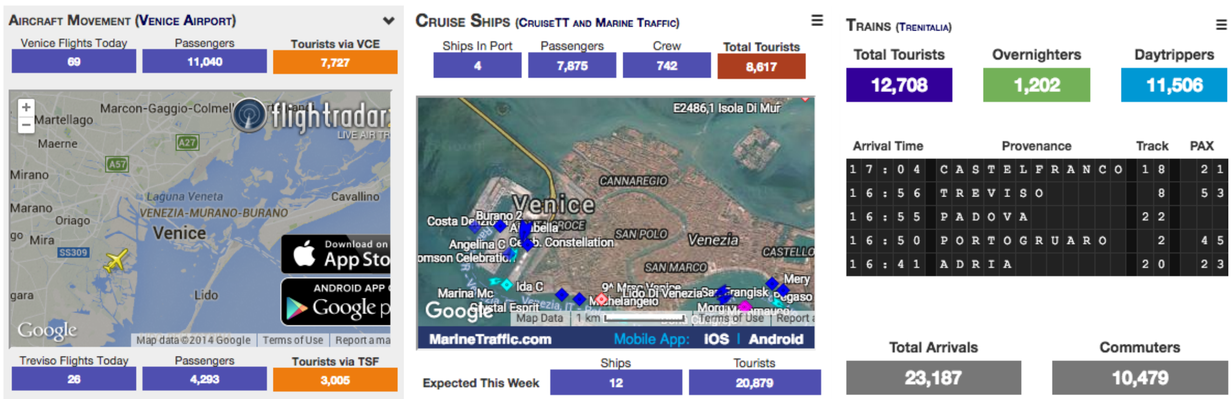

All three widgets in Figure 13 extract tourist arrival estimates by continuously mining big data from the web in real time.

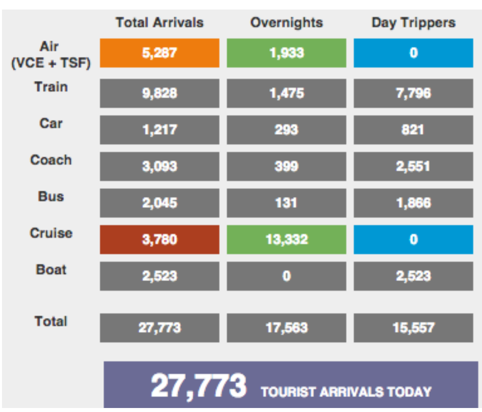

Together, these numbers add up to a plausible estimate of the number of tourists arriving in Venice every single day, as displayed in the Tourist Arrivals widget in figure 14 which summarizes the arrivals and splits them statistically into Overnight visitors and Day Trippers.

Similarly, we also harvest hotel reservations (www.Booking.com), hostels (www.Hostels.com) and Airbnb (using its APIs) to determine the number of beds that are occupied by tourists every night as shown in figure 15, giving us the total presence of tourists in the city at any moment.

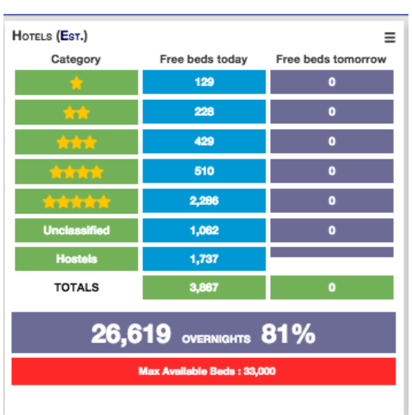

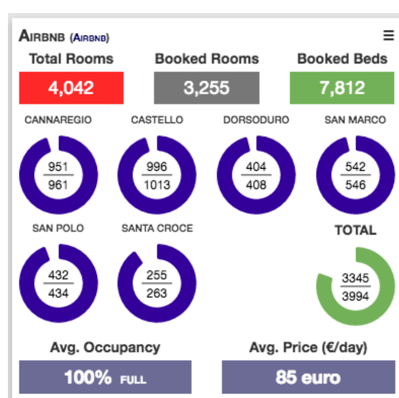

The number of tourists thus obtained can then be juxtaposed with the declining number of Venetians in figure 16 published at midnight every day on the City's website (Comune, 2016).

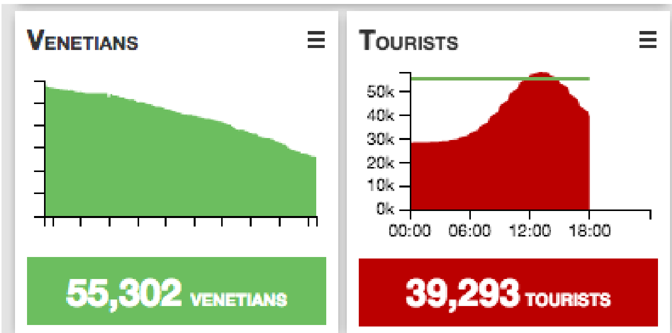

In the foreseeable future, it is likely that the City of Venice will need to restrict access to tourists in order to maintain a sustainable balance between its inhabitants and its visitors and our system is already accurate enough to support such tourist management measures (Carrera, 2016*b*). The Venice Dashboard could be made even more precise if the city's transportation agencies released transit ticketing and parking data, which currently lies hidden in their database systems. The availability of such data would allow us to track these phenomena even more precisely, which in turn would enable a more balanced management of the visitors, who outnumber the citizens on any average day of the year.

Big data, captured in real time, is useful in all cities, old and new. Older heritage cities are no different than newer ones in their need for fast new data to manage activities on a daily basis. If anything, heritage cities have more big data needs that newer ones, due to the simple fact that they to have to service and manage large numbers of visitors in addition to their own citizens.

Wise Cities: the next stage in the evolution of 'smart cities'

Just like the companion term 'Big Data', the 'Smart City' moniker has been popularized throughout the world for the past several years as the new frontier for municipal governments. Some are questioning its value and are proposing a piecemeal approach they call the 'clever city' which aims at getting truly useful applications out there for the immediate benefit of citizens, while we figure out what 'smart city' really means (Atkin, 2015). While it does mean different things for different municipalities, the overall gist of the 'smart city' appellation conjures up images of a digital cosmopolis governed through intelligent systems with access to a trove of real time data. Big data is what feeds the smart city in the public's imagination and we have seen many examples of clever big data visualizations in the media as well as in academic publications. Wise cities are not only 'clever' (Atkin, 2015) and they also go beyond the glitz and wiz-bangery of smart city applications and actually embed these technologies into their day-to-day operations, as the office of New Urban Mechanics has been doing for the City of Boston for more than a decade (http://newurbanmechanics.org/). An often-cited example of their practical approach to bringing innovation to the public sector is the application called *StreetBump*, which pinpoints the location of potholes in Boston using the accelerometers inside smartphones to create a crowdsourced map of road conditions every day (Carrera *et al.*, 2013). We developed this pioneering application as the logical consequence of prior experiments we had conducted at the Venice Project Center -- to map boat wakes (Carino and Marshall, 2006) -- and in Boston, where we tested a hardware pothole-mapping system in (DeMarco and Stedman, 2007). Each of these prior research projects could be labelled as a 'smart city' application, even though they predate the widespread diffusion of the 'smart city' tagline. Learning from our past experiences, we produced a 'wise city' application with StreetBump, which represents the natural evolution of the initial smart ideas, and has been successfully embedded into city operations to prioritize road maintenance in Boston. StreetBump is a 'wise city' application because it is based on past experience, it is clever, cheaper, faster and open to inputs from citizens as well as city workers, and it actually improves city services on a daily basis (Carrera *et al.*, 2013).

Wise cities, like Boston, publish their open data frequently and reliably (http://data.cityofboston.gov); they push out real-time data when possible, as was done by Boston's MBTA (Barry and Card, 2014); they are cognizant of their past and are able to extrapolate past information into the future, while also being knowledgeable about the present. Old cities like Venice and Boston cannot afford to be just 'clever' or 'smart' and ignore their past, so they need to go beyond clever/smart city applications and leverage the same technologies to capture the backlog of old slow data from the past, and to intercept slow real time changes as they happen, while still extracting useful patterns from fast new data and logging these away for future use as fast old data.

Wise cities blend past and present data not only to manage the day-to-day, but also to foresee realistically what may happen in the future. Wise city applications provide a complete, fungible record of what led to today's situation so that the future can be more predictable. They adopt state-of-the-art technologies like autonomous agent models, social collaborative filtering, crowdsourcing, micro-funding and other modern tools and adapt them to the big old data of the past as well as to the live streams of real-time big data of the present.

Wise city applications are not rigid and fixed, but rather open and constantly improving in an agile fashion. They do not become obsolete, but keep up with changing technology, smoothly and routinely. They get better with age. They do not create walled-gardens behind proprietary fences, nor do they lock data into silos. They make all data available to other applications and monetize only the transactions that are generated by the 'agents of change', by applying to information the same pay-as-you-use standards that are applied to other municipal infrastructure such as water and sewage (Carrera, 2004).

Wise cities treat all documentation as information and cumulatively collect all plan-demanded data generated by ongoing projects, developments and contracts to arrive at a state of plan-readiness whereby all of the existing information about present and past is made easily accessible to citizens, planners, professionals and decision-makers (Carrera and Hoyt, 2006).

A wise urban app does not just collect citizen complaints (Commonwealth Connect, 2016), but actually provides accurate measurements that are instantly mapped against allowable rules, like VeniceNoise does *vis* a *vis* the official noise zoning regulations (Calamari *et al.*, 2011). Clever apps allow citizens to contribute information with their phone by actively recording a situation (Schloss *et al.*, 2011), but wise ones achieve the crowdsourcing of practical information without any conscious citizen intervention, by simply being turned on and running in the background as we go about our daily tasks, as StreetBump does (Carrera *et al.*, 2013). Wise city applications don't just assist visitors in exploring the heritage of a city with clever location-aware apps and virtual reality overlays, but go beyond the show-and-tell and involve users in the preservation of monuments and public art, as PreserVenice does with its micro-donation approach to crowdfunding the restoration of historical artifacts (Carrera, 2012).

Wise city apps combine knowledge about the structures of the past, together with real-time data from the here and now, to facilitate activities in the future, as does our bridge clearance app (bridge.veniceprojectcenter.org) which assists Venice's emergency boats during high tides by instantly identifying low-clearance bridges that could be too low to pass under at high tide, depending on the height of the boat's cabin (Figure 17).

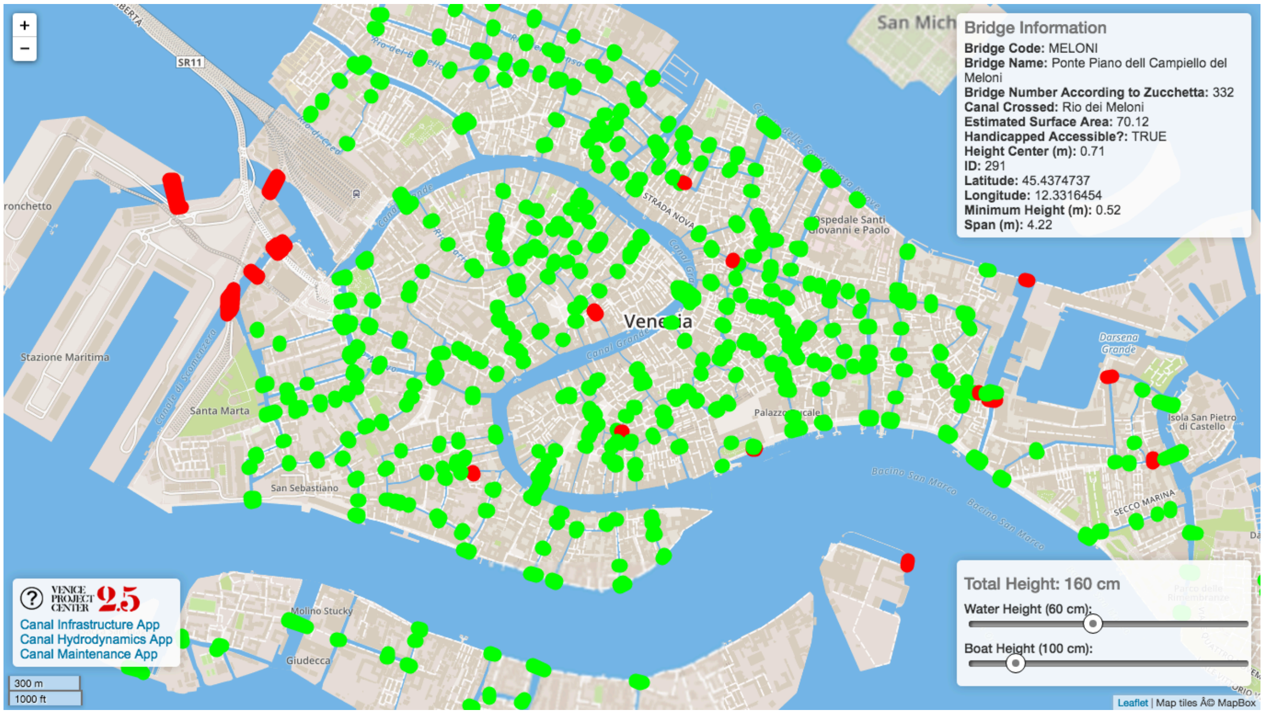

This wise app relies on old big data about bridge clearances and taps into a government tide-gauge in real time to show the current tide. Users can manipulate the slider at the bottom of the screen to set the height of their boat and instantly the map will show in red all of the bridges that are impassable by that boat, allowing ambulance, police or fire boats to re-route around them on the way to an emergency.

We hope that the examples from our work in Venice contained in this paper can point the way towards a more balanced approach to big data that focuses not just on present-day human activities and behaviors but also on the long tail of old data about the built environment that changes in slow real time, since real city knowledge is about both the old and the new, and incorporates activities as well as structures, which all together can better inform our decisions, and turn clever or smart cities into 'wise cities' that are truly effective in the maintenance, management and planning of our built environment.

NOTES

1. The examples in this paper are based on research developed and supervised by the author and conducted in Venice since 1988 by over 750 engineering and science students from the Worcester Polytechnic Institute (WPI) during their two-month research trips to the Venice Project Center

(VPC), augmented by hundreds of volunteers and funded through academic grants and professional contracts through the consulting firm Forma Urbis s.r.l. and powered by the technology spin-off called City Knowledge LLC. Some examples also come from the author's experience as director of WPI's Boston Project Center between 1997 and 2007.

REFERENCES

Angelo, R., Chan, A., Hill, D., Vasconcelos, J.M. (2014) Piera Alta: Facilitating pedestrian mobility in Venice during floods. Worcester Polytechnic Institute IQP.

Ascare, A., Fletcher, H., Mazzucotelli, A., Pierson, D. (2010). PRESERVENICE: Preserving Venetian Material Culture. Worcester Polytechnic Institute IQP.

Atkin, R. (2015) *Manifesto for a clever city*. Available at: http://theclevercity.net. Accessed on 27 July, 2016.

Baruffi, F., Boucher, J., Coryea, M., Spector, D., (2012) Venice bells and bell towers A Striking Source of Knowledge. Worcester Polytechnic Institute IQP.

Barry, M. and Card, B. (2014) *Visualizing MBTA Data*. Available at: http://mbtaviz.github.io. Accessed 27 July, 2016.

Batty, M., Axhausen, K.W., Giannotti, F., Pozdnoukhov, A., Bazzani, A., Machowicz, M., Ouzounis, G., Portugali, Y. (2012) Smart cities of the future. *European Physical Journal Special Topic*, **214,** pp. 481-518.

Bean, Randy (June 23, 2014) *Big Data Fatigue?*. MIT Sloan Management Review. Available at http://sloanreview.mit.edu/article/big-data-fatigue. Accessed 23 July, 2016.

Bennett, A., Hastings, S., Petry, S., Solomon, A. (2013) Boats are Waking Me Crazy: An Analysis of Boat Traffic and Moto Ondoso in Venice. Worcester Polytechnic Institute IQP.

Blackwell, E.L.II, Fraizer Cryan, M.E., Rizzo, A.M., Vello, K.A., Wainwright, R.L. (2000) Preserving Venetian Wellheads. Worcester Polytechnic Institute IQP.

Bobell, P., McCarthy, L., Poganski, D., Tsiros, A., (2012) Venipedia, an English Language Data-Driven Wiki Dedicated to Venice. Worcester Polytechnic Institute IQP.

Bossalini, D., Gardner, B., Mathews, T., Wills, M., (2013) Canals, Bridges and Urban Maintenance. Worcester Polytechnic Institute IQP.

Bove, C., Gallagher, K., Hickey, M., Honicker, J., (2015) Preserving Venice's Bells and Their Towers. Worcester Polytechnic Institute IQP.

Brann, K.L., Giola, G.T., Kirby, D.J., Richtmyer, W.H., 2013. 25 Years of Venice Knowledge Online. Worcester Polytechnic Institute IQP.

Bruso, B., Hongling C., Olm, A., Schulman, I. (2012) The Merchants of Venice: A Look at the Changes in the Venetian Stores and Tourist Accommodations and Their Impact on the Local Population. Worcester Polytechnic Institute IQP.

Cafferelli, B., Piper, N., Sutman, E. (2007) uSCRIPT: an AJAX-based Online Manuscript Transcription Service. Worcester Polytechnic Institute MQP.

Calamari, E.S., Pomerleau, N.C., Puishys, R.E., Ripley, W. J. (2011) Noises of Venice: An Exploration of Noise in an Historic City. Worcester Polytechnic Institute IQP.

Calvino, Italo (1972) *Le Città Invisibili*. Milan: Einaudi editore.

Caniato, G., Carrera, F., Giannotti, V., Pypaert, P. (eds.) (1999) *Venezia la Città dei Rii*. Venice: Cierre.

Carino, F.J., Marshall, S.C. (2006) Mapping Turbulent Discharges in Venetian Canals. Worcester

Polytechnic Institute MQP.
Carrera, F. (1997) What Cultural Heritage do we preserve and why?, Doctoral research paper, Massachusetts Institute of Technology.
Carrera, F. (1999) Il comportamento idrodinamico dei canali interni di Venezia, in *Venezia la città dei rii*, Venice: Cierre, pp.197-207.
Carrera, F. (2004) City Knowledge: an emergent information infrastructure for sustainable urban maintenance, management and planning. PhD Dissertation, Massachusetts Institute of Technology.
Carrera, F. (2005) Making History: an Emergent System for the Systematic Accrual of Transcriptions of Historic Manuscripts, in proceedings of *IEEE's 8th International Conference on Document Analysis and Recognition*. Seoul, South Korea.
Carrera, F. and L. Hoyt. (2006). "From Plan-Demanded Data to Plan-Ready Information: A Rationale for Comprehensive Urban Knowledge Infrastructures". *Journal of Urban Technology* 13 (2) 3-23.
Carrera, F. (2012) Crowdsourcing (and Crowdfunding) the Preservation of Venice's Material Culture, *Annual Meeting of the American Institute for Conservation of Historic and Artistic Works* (AIC), Albuquerque, NM: May 8-11.
Carrera, F., Guerin, S., Thorp, J. (2013) By the People, for the People: the Crowdsourcing of "StreetBump", an Automatic Pothole Mapping App, *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* (ISPRS), **XL-4/W1**, 29th Urban Data Management Symposium, 29-31 May, 2013. London.
Carrera, F. (2016a) Future cities of the past. Presentation at *Future Cities Catapult*, London: May 23.
Carrera, F. (2016b) *Turismo sì, ma…* Book I of collection: *Promemoria per un Futuro Sindaco*. Venice: Supernova edizioni. Under revision.
Chakuroff, M.S., Powderly, K.B., Silva, D.S. (2014) Venetian Time: An Examination of Time and Space in Venice. Worcester Polytechnic Institute IQP.
Chen, J., DiNino, A., Mintz, Z., Wolf, R. (2011) An Update on the Hydrodynamics of the Venice Canals.
Chhim, S., Hoctor, T.A., Pryor, J.W., Sadowski, B.J. (2014) The Evolution of Venice Through Mapping Technology: A Venice Project Center Collaboration with the Atlante della Laguna. Worcester Polytechnic Institute IQP.
Chines, J.F.; Eoff, E., Reynolds, A.C., Weiss, A.M. (2013) Venice Through the Canals of Time: Mapping the Physical Evolution of the City. Worcester Polytechnic Institute IQP.
Commonwealth Connect. Available at: http://commonwealthconnect.io. Accessed July 25, 2016.
Comune di Venezia, Portale dei Servizi, *Mappa della popolazione residente al giorno precedente*. Available at: https://portale.comune.venezia.it/millefoglie/statistiche/home. Accessed 26 July 2016.
Connor, C.J. Jr., Hanna, T.S., Van Rensselaer, W., Wingerter, Z.R. (2015) Safe and Sustainable Tourism: Managing Venice's Millions of Visitors.Worcester Polytechnic Institute IQP.
Dechaine, D., Hennessey, M., Orszulak, J., Rullmann, K. (2012) Treasures Underfoot: Preserving Venice's Church Floor Artifacts. Worcester Polytechnic Institute IQP.
DeMaio, E., Jones, E.K., Mistretta, L-J. V, Rappoli, O.R. (2012) Streamlining Canal Hydrodynamic Measurements in Venice: Updating Measurements and Developing Plans for Extensive Data Collection. Worcester Polytechnic Institute IQP.

DeMarco, A.J. and Stedman, C.C. (2007) Automated GPS mapping of road roughness. Worcester Polytechnic Institute MQP.

Eagle, N. and Pentland, A. (2006) Reality mining: sensing complex social systems, *Pers Ubiquit Comput*, **10**, pp. 255–268. London: Springer Verlag.

Elbag, M.A.Jr., Epstein, A.M., Hannigan, J.F., Rosinha, E.M. (2003) Preservation and Restoration of Venetian Public Art: From the Completion of the Public Art Catalog to the Active Restoration Process. Worcester Polytechnic Institute IQP.

Felices, M., Goodfellow, L.M., Johnston, J.L., Maheshwary, S. (1997) Preliminary Feasibility Study of the Implementation of a HIFLO Vacuum Sewerage System within the City of Venice. Worcester Polytechnic Institute IQP.

Finelli, T., O'Brien, C., Scannell, K. (2010) Venipedia A Modern Knowledge & Data Wiki Dedicated to the City of Venice. Worcester Polytechnic Institute IQP.

Finzel, B.R., O'Toole, S.E., Perron, J.M., Russel, J.E. (2007) Noise data farming for the city of Boston. Worcester Polytechnic Institute IQP.

Flaxington, T., Mehilli, R., Pepin, A., Sturman, D. (2015) Mobility on Venetian expressways: alleviating congestion and improving navigation in Venice. Worcester Polytechnic Institute IQP.

Fletcher, A.K., Halloran, A.G., Malik, R.M., Rohleder, K.H. (2002) Integrated, Multi-Agency Approach to the Preservation of Venetian Palaces. Worcester Polytechnic Institute IQP.

Gomes, E., Griffin, J., Vaillancourt, V., Wood, M. (2005) Making environmental impact reports re-usable for the City of Boston. Worcester Polytechnic Institute IQP.

Heinricher, D., Kahn, L., Maitland, Ii., Manor, N., (2013) Ecclesiastical Architecture of Venice: Preserving Convents, Churches, Bells and Bell Towers. Worcester Polytechnic Institute IQP.

Howe, D., Costanzo, M., Fey, P., Gojobori, T., Hannick, L., Hide, W., Hill, D.P., Kania, R., Schaeffer, M. St Pierre, S., Twigger, S., White, O., Yon Rhee, S. (2008) Big data: The future of biocuration. *Nature*, **455**, pp. 47-50.

Kelley, A.E., Ratner, G.I., Schmaelzle, M.J., Thomollari, O. (2004). Public Art Preservation in Venice: Non-Public Wellheads and Fountains. Worcester Polytechnic Institute IQP.

LaRovere, F., Sawin, J., Gandomi, K., Sherifi, E. (2015) Evaluating Changes in the Venetian Retail Sector and Managing the Use of Public Space. Worcester Polytechnic Institute IQP.

Lohr, Steve (Aug 12, 2012), *How big data became so big*, New York Times: Business Day, Unboxed. Available at http://www.nytimes.com/2012/08/12/business/how-big-data-became-so-big-unboxed.html. Accessed 23 July, 2016.

Lynch, Kevin A. (1984) *Good City Form*. Cambridge: MIT Press, p. 48.

Milton, R.W., Batty, M., Manley, E., Reades, J. (2013) Imagining the Future City: London 2062. In: Bell, S. and Paskins, J., (eds.) *Imagining the Future City: London 2062*, London: Ubiquity Press, pp. 31-40.

Ratti, C., Frenchman, D., Pulselli, R.M., Williams, S.(2006) Mobile Landscapes: Using Location Data from Cell Phones for Urban Analysis, *Environ Plann B Plann Des*, **33-5**, pp. 727-748.

Schloss, A., Beaudry, J., Pickle, J., Carrera, F., Guerin, S. (2011) How mobile picture posts will expand participation in the Digital Earth Watch and Picture Post citizen environmental monitoring network. *96th ESA Annual Meeting*, Austin, TX: Aug 7-12.

SeeClickFix. Available at: https://gov.seeclickfix.com. Accessed 25 July, 2016.