Capturing Urban Dynamics with Scarce Check-In Data

An analysis of three months' worth of Foursquare check-in data for a city in Greece shows that although the data generated by the citizens is scarce, it's sufficient to build a good model of the city's dynamics.

Andreas Komninos
Glasgow Caledonian University

Vasilios Stefanis, Athanasios Plessas, and Jeries Besharat
University of Patras

Location-based services are an important part of our daily interactions with mobile and desktop services. Services such as Foursquare and Facebook let users geo-annotate information about a venue, using spatial and temporal context dimensions. These interactions result in large exploitable datasets describing patterns of human interactions with their environment. Researchers have been conducting studies into the use of such datasets and recent work in this area suggests that check-ins convey not just that a user has visited a place but that he or she finds it to be a place of interest and worth mentioning (for more information, see the sidebar).[1,2]

Given the potential of such datasets to convey areas of interest, an application that mines check-in data, automatically extracting "local knowledge" in the form of urban rhythms and locations of true interest, might prove invaluable to tourists visiting a city. Here, we describe our analysis of Foursquare check-in data for a city in Greece, showing that although the data generated by the citizens is scarce, we can use it to build a good model of the city's dynamics. We further discuss how this information can be used to guide visitors in the city or provide innovative services for city inhabitants using the cloud.

The Importance of Local Knowledge

In the tourism domain, local knowledge has long been branded the ultimate source of information for visitors, and its sharing has been the subject of discussion for years. Visitors and tourists rely heavily on third-party information relating to sights and areas of interest. Traditional information sources include guidebooks, where the subjective opinion of an expert is used to describe a location of interest. Malin Zillinger, however, has argued that guidebooks restrict tourists, artificially popularizing locations and discouraging visits to unmentioned sites.[3] Furthermore, Benjamin Lucca Iaquinto found that this effect is worsened by editorial interventions that influence the guidebooks' final content.[4]

Online user-generated content websites, such as Wikitravel and Trip Advisor, also reflect the subjective opinions of participants that might well have been influenced by guidebooks, given that 60.6 percent of online users have been found to use guidebooks during a visit.[5] Even though online media is regarded as collaborative and thus likely to present a more accurate picture, Ulrike Gretzel and Kyung Hyan Yoo found that reviews therein play a much less significant role in terms of determining what to do (32.5 percent), where to go (27.7 percent), and when to go (26 percent).[5] Their results show that the most popular use of online reviews was

 Published by the IEEE CS ■ 1536-1268/13/$31.00

Related Research in Location Sharing and Checking In

Anastasios Noulas and his colleagues conducted a large-scale study involving 700,000 users, collecting data for 100 days that represented approximately 12 million Foursquare check-ins.[1] Their data is indirectly observed, because it's accessed through Twitter messages generated from the Foursquare application and thus only reflects the behavior of Foursquare users who have connected the application to a Twitter account. Their work focuses on the diurnal breakdown of check-in times and venue categories, as well as temporal dynamics and spatial variances between check-in actions. The authors demonstrate distinct potential in inferring information such as transitions between venues.

Zhiyuan Chen and his colleagues performed a similar study, analyzing location-tagged check-ins from Twitter.[2] They found that users tend to exhibit periodic behaviors, while the social linkage, geographic, and economic constraints seem to affect user mobility patterns.

Justin Cranshaw and his colleagues attempted to cluster such data into areas of particular social activity in cities.[3] The data collected in their analysis was again indirectly observed through Twitter. Through interviews with 22 location-based-services users, they found that check-in data could be used to represent known divisions in communities and reveal subtle changes in the local social patterns.

Jessica Benner and Cristina Robles used Foursquare data to explore the trending behavior of places in three cities in the US.[4] They uncovered distinct patterns of trending in each city and argue that such analysis can help urban analysts understand what's happening in different cities, thus helping business owners better manage their businesses.

Jonathan Chang and Eric Sun used Facebook Places check-ins to build a predictive model based on previous check-ins, friends' check-ins and day, demographics, and time, with a target of predicting the user's likely next check-in location.[5]

The semantics of check-ins have only recently been investigated in academic literature. Janne Lindqvist and his colleagues studied the context under which people check in and discovered that participants don't check in to places they consider embarrassing (such as fast-food restaurants) or uninteresting or that they visit frequently.[6] This shows that Foursquare check-ins indicate positive disposition toward a location and signify its importance (it's a place to be seen or is an interesting place). Sameer Patil and his colleagues found location sharing was driven strongly by projections of personal taste and image, such as a desire to indicate that a user likes a place or wants to appear cool and interesting.[7] Other reasons, such as financial incentives or promoting events, were much less frequently stated as reasons for sharing. Similar findings were also discovered by Henriette Cramer, Mattias Rost, and Lars Erik Holmquist, who, in a qualitative study of in-depth interviews and surveys, report the emergence of "social and identity-driven uses, such as sharing lifestyle, events, and information that's interesting and enhances self-presentation."[8]

REFERENCES

1. A. Noulas et al., "An Empirical Study of Geographic User Activity Patterns in Foursquare," *Proc. 5th Int'l AAAI Conf. Weblogs and Social Media,* AAAI press, 2011, pp. 150–573.
2. Z. Cheng et al., "Exploring Millions of Footprints in Location Sharing Services," *Proc. 5th Int'l AAAI Conf. Weblogs and Social Media,* AAAI Press, 2011, pp. 81–88.
3. J. Cranshaw et al., "The Livehoods Project: Utilizing Social Media to Understand the Dynamics of a City," *Proc. 6th Int'l AAAI Conf. Weblogs and Social Media,* AAAI Press, 2012, pp. 81–88.
4. J. Benner and C. Robles, "Trending on Foursquare: Examining the Location and Categories of Venues that Trend in Three Cities," *Proc. Workshop GIScience in the Big Data Age—7th Int'l Conf. Geographic Information Science,* 2012, pp. 27–34; http://stko.geog.ucsb.edu/gibda2012/gibda2012_submission_3.pdf.
5. J. Chang and E. Sun, "Location3: How Users Share and Respond to Location-Based Data on Social Networking Sites," *Proc. 5th Int'l AAAI Conf. Weblogs and Social Media,* AAAI Press, 2011, pp. 74–80.
6. J. Lindqvist et al., "I'm the Mayor of My House: Examining Why People Use Foursquare," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 11), ACM, 2011, pp. 2409–2418; doi: 10.1145/1978942.1979295.
7. S. Patil et al., "'Check Out Where I Am!': Location-Sharing Motivations, Preferences, and Practices," *Proc. CHI '12 Extended Abstracts on Human Factors in Computing Systems* (CHI EA 12), ACM, 2012, pp. 1997–2002.
8. H. Cramer, M. Rost, and L.E. Holmquist, "Performing a Check-in: Emerging Practices, Norms and 'Conflicts' in Location-Sharing Using Foursquare," *Proc. 13th Int'l Conf. Human Computer Interaction with Mobile Devices and Services* (Mobile HCI 11), ACM, 2011, pp. 57–66.

for deciding where to stay (77.9 percent).[5] Just under half of online site users look for information from local destination (44.6 percent) and state tourism websites (29.7 percent), although these might more accurately reflect a local's knowledge. Ryen White and Georg Buscher showed that relying on information from other tourists isn't optimal, because nonlocals tend to select venues that lead to lower-quality experiences.[6]

The reviews likely to be shared on tourism websites do little to help tourists uncover hidden gems or learn about what's truly worth visiting. The sharing of local knowledge is thus still elusive, despite advances

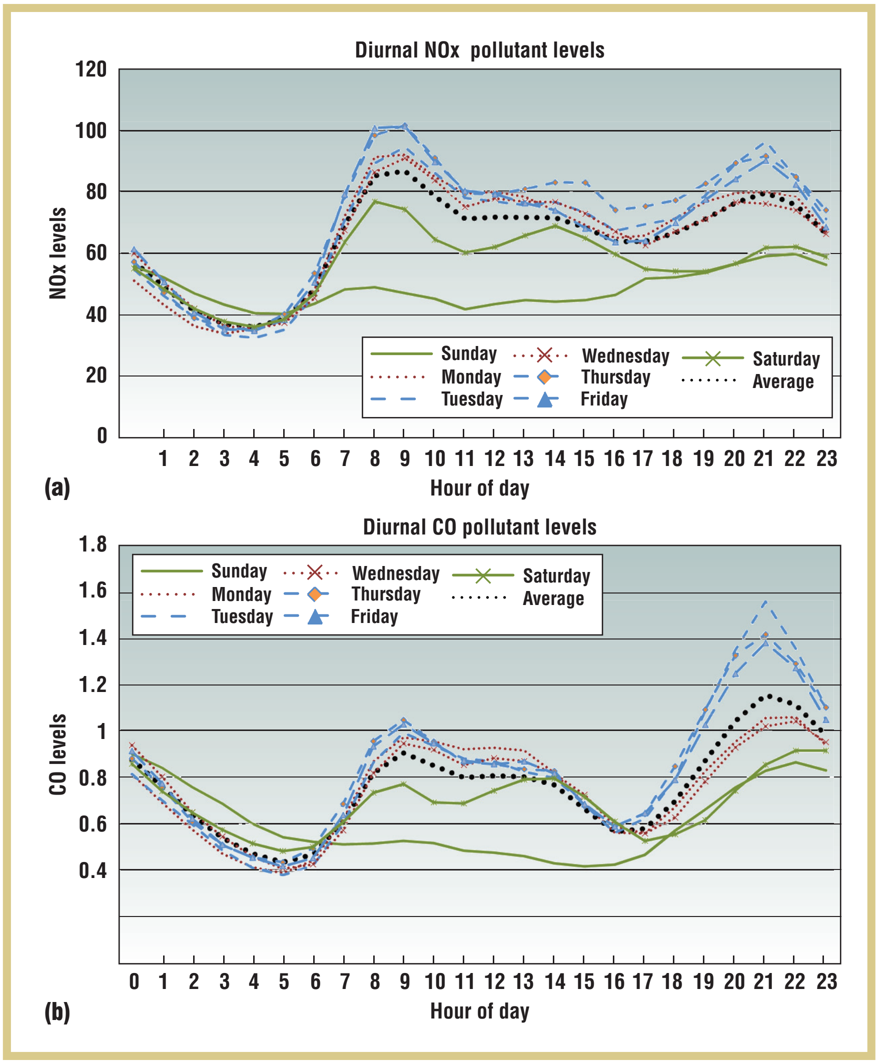

Figure 1. Diurnal cumulative (a) nitrogen (NO_2 and NO) pollutant and (b) carbon monoxide (CO) level averages.

in technology. Bearing in mind the difficulty in capturing and sharing knowledge from locals, we explore a system for automatically inferring such knowledge by mining check-in data.

Capturing the Rhythm of a City

Our work focuses on Patras, Greece—a medium-sized Mediterranean city with approximately 200,000 inhabitants. Here, we provide some information about the local context and known urban dynamics, which vary distinctly from the typical 9-to-5 work week. This background information is helpful in interpreting the results that emerged from our dataset.

Local Context

Most employees in the public and private sector work Monday through Friday. Working hours are 7 a.m. to 3 p.m. for the public sector. Banks are open for business between 9 a.m. and 3 p.m. (although employees stay later). Most private-sector office workers work from 9 a.m. to 5 p.m. Shop opening hours are regulated by law and are Monday through Friday, 9 a.m. to 9 p.m., with a break between 2 and 5:30 p.m. Shops are open Saturdays (9 a.m. to 2:30 p.m.) but are closed on Monday and Wednesday evenings. Several self-employed professionals (lawyers, civil engineers, and so on) also follow the shop opening hours, although they regularly work on Monday and Wednesday evenings and not on Saturdays. Finally, most shops and businesses are closed on Sundays.

Although the population's work schedule is widely fragmented, busy days in the city are typically Tuesday, Thursday, and Friday, when shops are open in the morning and evening. We refer to these as "full days," and we refer to Monday and Wednesday as "half days" and Saturday and Friday as the "weekend."

Objective Data Verification

Although local expert knowledge is a good starting point, we wanted to ensure that we could validate our analysis of data captured through Foursquare against an objective baseline. We sought other datasets indicative of human activity in an urban environment and considered data on vehicular traffic and air pollution.

Air pollution measurements have been found to coincide well with known city patterns.[7] Our data was obtained from the public repository of the Hellenic Ministry of Environment, Energy and Climate Change (atmospheric pollution data—www.ypeka.gr/Default.aspx?tabid=492&language=el-GR [in Greek]). Although this data includes diurnal hourly measurements of several pollutants, we considered only carbon monoxide (CO) and nitrogen oxides (NOx), because they're the two pollutants most closely related to traffic.[8] We analyzed data from 2009, obtained from the city center air quality monitoring station, so the data (Figure 1) would match the available traffic volume data, which we gathered from a 2009 feasibility study for adding a tram system to the public structure infrastructure in the city of Patras.[9]

Figure 2a shows the locations of measurements taken during the tram

Figure 2. Objective data for verification and Foursquare check-in data in Patras: (a) traffic volume measurement stations and the air quality monitoring station in the area of coverage of (b) Foursquare data-collection listening posts.

feasibility study and the location of the air quality monitoring station. The traffic data covers a much smaller period than the atmospheric data but is still useful. The blue markers indicate locations where measurements were taken during single 24-hour periods, while pink markers indicate locations where measurements were taken during an entire week (stations 140 and 141). The air quality monitoring station is marked with an "X". The average diurnal volume measurements are shown in Figure 3, with the volumes for stations 140 and 141 depicted separately.

As is evident from the analysis of this data, the city exhibits a measurable rhythm, which coincides with our local knowledge as described in the previous section. Distinct activity peaks are noted in the morning and afternoons, coinciding with shop opening hours. Half-days display less pollution in the evenings compared to full days, when the shops are open. On weekends, Saturday is busy in the morning but less so than other full or half-days. Sunday seems quiet throughout, showing some increased activity in the evening.

The atmospheric pollution data is partially backed up by the traffic volume data, although the traffic dataset is much less extensive, because the coverage period is just 24 hours for all measurement stations (Thursdays and Fridays during the months of March and April 2009) except two, which were measured over a week. Still, the twin peaks expected from the analysis of atmospheric data are also present in the traffic volume data.

Foursquare Data Analysis

In contrast with other studies,[10,11] we aimed to focus on a particular location, explore the quality of data that could be obtained, and consider how this data could be used to describe and share users' interactions with their environment and with others.

Data Capture

To ensure we didn't miss check-ins that weren't tweeted (see the sidebar for information about Twitter check-ins), we devised a way to collect data directly from Foursquare's API. For 100 days between July and September 2012, we set up "listening posts"—fixed locations in the city, covering the commercial and social areas of interest, based on our local knowledge (Figure 2b). For each listening post, we queried the Foursquare API every 30 minutes to retrieve the names of nearby businesses and their check-in data. The information for each venue (time of query, current check-ins, and total check-ins) was saved in a database. By querying every 30 minutes, we were able to approximate check-in time since the API didn't provide this information. This temporal resolution seemed adequate, considering Foursquare's check-in timeout policy, which keeps a user checked into a place for a maximum of three hours or until he or

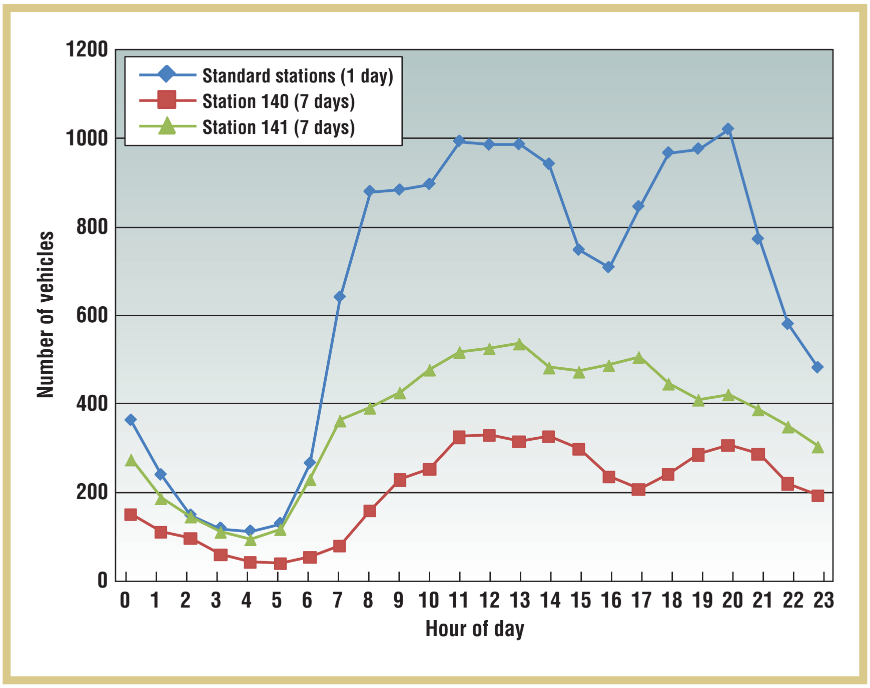

Figure 3. Traffic volume in 2009 for the city of Patras.

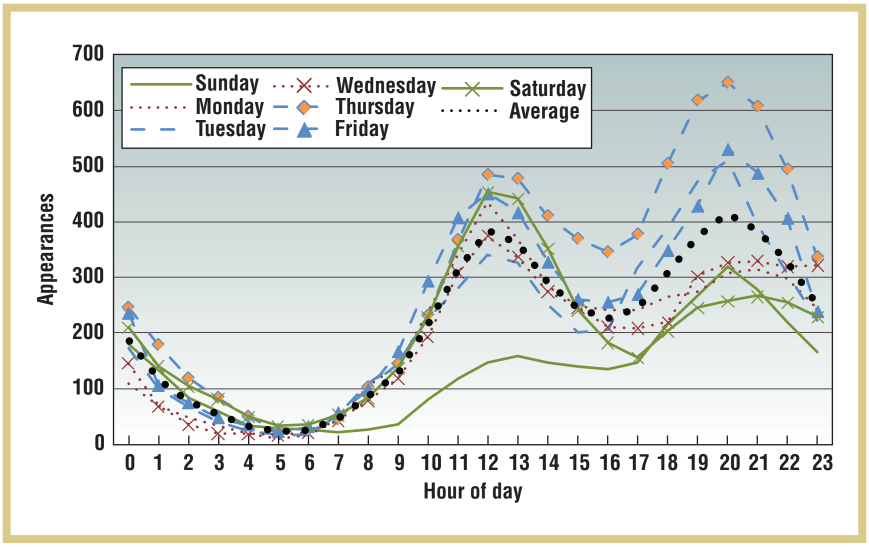

Figure 4. Daily diurnal breakdown of total appearances. The emergent pattern seems realistic—that is, activity is reduced in the early hours of the day, and peaks appear around mid-day and mid-evening.

she checks in to another venue. As such, our data doesn't show distinct check-ins but rather how many people appear to be checked into a venue at any point in time.

Using a discovery process (and not a static places list) meant that if a new place was added in the period of study, it would also start to be included in the results. In total, we included 282 venues, of which 249 actually showed check-ins during our data-collection period. The remaining 33 venues were present in Foursquare; however, no check-ins had occurred until the end of the data-collection period of our study.

In total, we collected 889,043 "appearance" reports for the 249 venues. We calculated an estimate of average check-ins per day by dividing the sum of differences in the total number of check-ins for each venue at the start and end of the period by the number of days (m = 145.82 check-ins per day). This estimate shows that use of Foursquare isn't widespread in the city and that the check-ins can be considered scarce.

Considering the penetration of smartphones in Greece (approximately 25 percent),[12] the local cost of 3G connections (made significantly less affordable because of the economic crisis and cuts), the lack of widespread adoption of free Wi-Fi, and the general low adoption of Foursquare usage worldwide (31 percent of mobile users active on social networks are on Foursquare),[13] the low number of daily check-ins isn't surprising. Nevertheless, we wanted to explore whether urban dynamics could be uncovered from such scarce data.

Diurnal and Daily Analysis

Figure 4 shows the total appearances, broken down by hour of day and by weekday in the city. The emergent pattern seems realistic—that is, activity is much reduced in the early hours of the day, while peaks around mid-day and mid-evening are in line with the well-known busy periods when the shops are open and the city is buzzing.

In terms of daily breakdowns, we can see that the data follows the atmospheric pollution patterns quite closely. Correlations are found between diurnal Foursquare check-ins and traffic volume ($R_{(24)} = 0.757$, $p < 0.001$), NOx ($R_{(24)} = 0.621$, $p < 0.001$), and CO ($R_{(24)} = 0.688$, $p < 0.001$) pollutants. Full days exhibit greater-than-average activity, while half days are on par with the average in the mornings and are below average in the evening. Saturdays show increased activity in the morning, and Sundays are quiet, apart from evenings. Peaks occur slightly later than pollution and traffic data in the morning and earlier in the evening, which is expected,

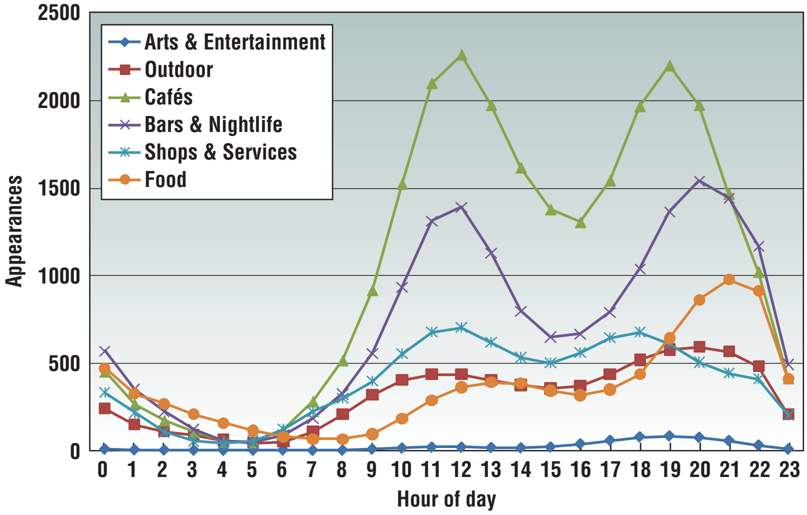

Category Legend
Cafés *Café, coffee place, Internet café*
Bars & Nightlife *Bar, wine bar, nightclub*
Food *Fast food, ice cream shop, Greek restaurant*
Outdoors *Plaza, marina, park*
Shops & Services *Clothing store, electronics store, bookstore*
Arts and Entertainment *Movie theater, theater, music venue*

Figure 5. Venue categories and total appearances by day of the week.

because people converge in the city first, then start "checking in" to places; in the evening, check-in activity gives way to transport activity.

The three-hour check-in timeout might delay the decline of check-in volumes by some amount. However, because the API doesn't provide the precise time that each check-in was cancelled, we can't definitely tell what effect this has on our data. Still, a resolution of three hours should be appropriate for describing urban dynamics in larger daily segments, instead of hourly granularity, particularly for tourism-related applications (discussed later).

Based on the pollution data, we expected full days to display similar levels of check-in activity. However, we can clearly see that Thursday shows distinctly increased activity in the evenings compared to other full days. A possible explanation is that as the week draws closer to the weekend, people are more inclined to go out, so they check-in more often. Indeed, a careful look into our data shows that Thursdays exhibit greater activity throughout the day, while Friday only shows increased activity in the morning compared to Tuesday.

Weekend behavior varies distinctly from full and half days. Saturday and Sunday seem less busy than expected in the evenings. This is in line with our atmospheric pollution findings. It can also be attributed to the summer weather when our study took place, because people are known to gather outside the city center for nights out (the city has two popular seaside summer "evening" suburbs, Rio and Vrachneika, which are approximately 10 km east and west of the city center).

Venues and Business Types

Foursquare uses a multilevel category system to let users characterize venues. Our data clustered around five default top-level categories (Food, Nightlife, Outdoors, Shops & Services, and Arts & Entertainment). When analyzing the data per category, we found that the top-level categories were too generic for meaningful clustering (for example, Food contains both restaurants and coffee places, but activity in these venue types is very different). As such, we decided to use our own top-level categories to group the low-level categories for each venue into clusters.

The examination of the diurnal breakdown of appearances by category depicts a realistic pattern of results (see Figure 5). The coffee and nightlife venues follow a similar two-peak pattern with nightlife spots peaking later in the evening. Their morning popularity can be explained as most bars and cafés in the city operate as all-day bar-café venues. The appearances in shops peak and fade as might be expected based on the shops' opening hours and the restaurant peaks in the evening (Greeks are known for eating later at night). Outdoor locations seem more popular in the evenings, which is natural in the summer. The infamous Mediterranean "coffee culture" seems to be well captured in our data, with Cafés topping the chart (25,663 appearances), followed by Bars & Nightlife (17,921).

We were also surprised to find that in the Food category (11,758 appearances), fast food outlets were by far the most popular (3,174) subcategory. This could be indicative of the low average age of location-based-services users or even a byproduct of the current economic crisis the country faces. Outdoor locations (7,602), Shop & Service (6,673), and Arts & Entertainment venues (662) were the least popular categories for checking in.

The distribution of appearances for all venues in our study exhibits

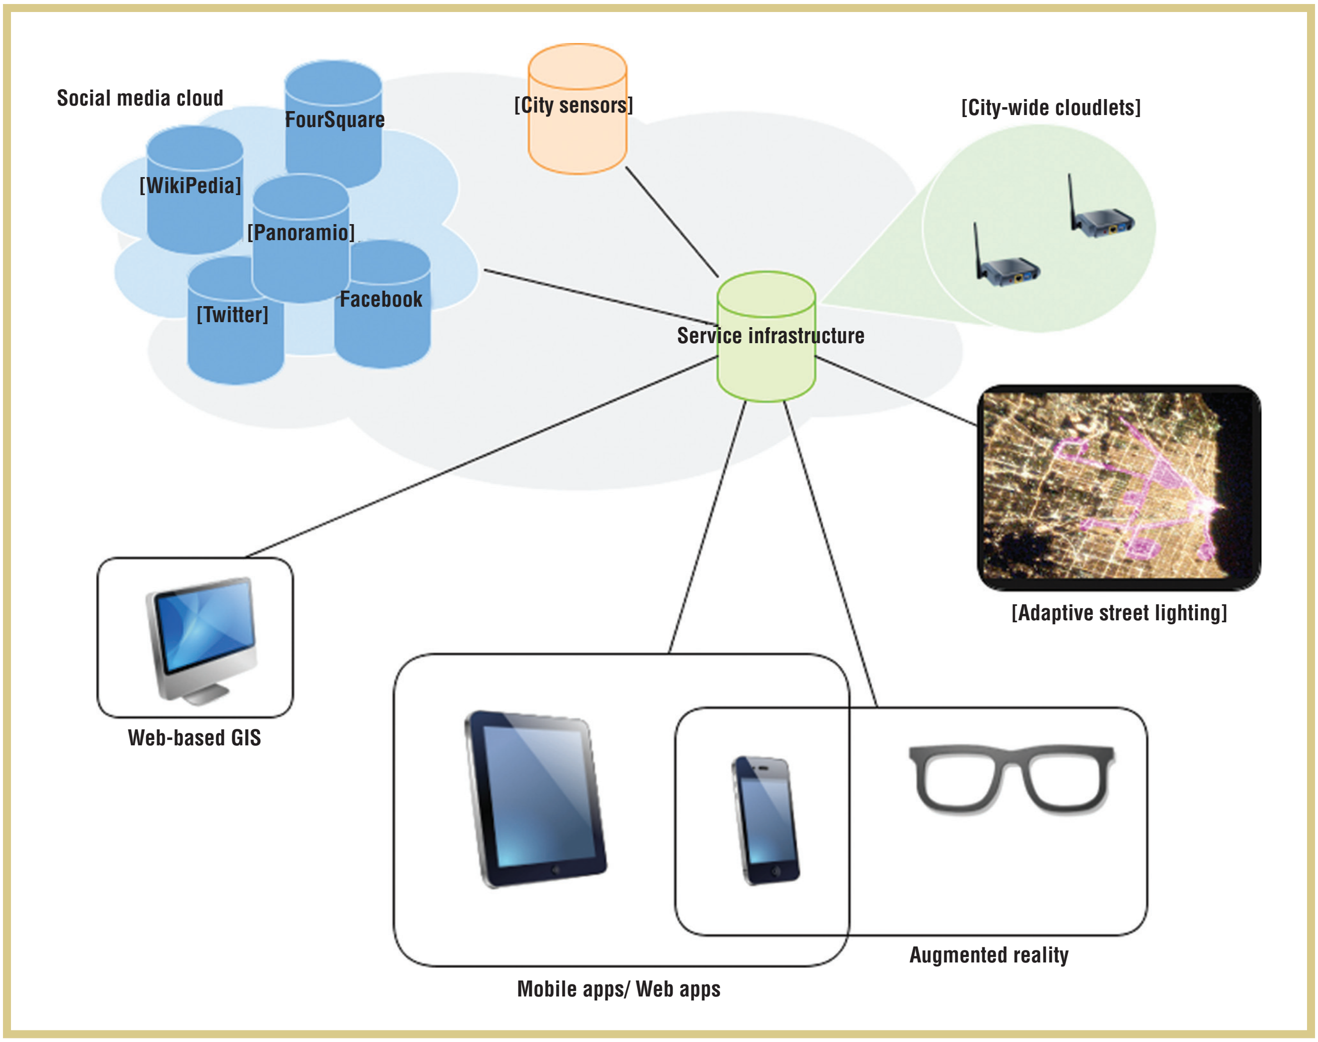

Figure 6. A cloud-oriented system approach for sharing captured local context. Elements yet to be implemented appear in brackets.

power-law behavior, with only a few venues taking up the majority of check-ins. In fact, the top 20 percent of all venues in terms of appearances (50 venues) take up 69.2 percent of the appearance distribution. A correlation exists between the sum of appearances in a place and the precise total number of its check-ins as reported by Foursquare at the end of our study ($R_{(249)} = 0.78$, $p < 0.01$). This shows that a popular place will generally continue to be popular, although we'll need to re-examine this behavior once we've collected a year's worth of data, which will show behavior change over time, because it's well known in the city that the popular areas for hanging out change between winter and summer.

Practical Implications

Encouraged by the fact that even such scarce data can indeed provide an accurate depiction of the city's rhythms, we started considering the practical implications in making this knowledge ubiquitously available. Our aim was to make this information practically digestible by the average person in the form of a useful service that would support interactions with their surrounding urban environment, regardless of the available computing equipment (or lack of) in the context of a user. As such, we designed a cloud-based service approach (Figure 6), which removes the computation, connectivity, and data management load from ubiquitous clients such as a mobile device.

Our architecture is inspired by the application of cloud computing and cloudlets in ubicomp scenarios, as envisioned by Mahadev Satyanarayanan.[14] In this architecture, our cloud-based service aggregates information from other cloud sources, such

(a) (b) (c)

Figure 7. Cloud-driven services: (a) a map-based temporally contextual visualization of city data on a regular website (venue suggestions are on the right); (b) a mobile Web app showing recommendations for venues in the current spatial and temporal context; and (c) a prototype dynamic augmented reality interface to the city data built on an Android device (with the Unity + Vuforia software development kit).

as social media and city-wide sensors, and produces knowledge based on it. The aggregated knowledge can be passed on as a variety of services—GIS-based websites, context-aware mobile apps and mobile Web apps, augmented-reality apps, and even adaptive infrastructure (such as street lighting that can act as low-fidelity displays). Locally placed cloudlets, distributed throughout the city as part of the public service infrastructure, can act as proxies to minimize latency and communication issues with the aggregation service.

We've already started implementing this architecture. The first step after collecting our data and implementing a knowledge-generation service was to create a webpage displaying all known venues in the city and graphically depicting their "hotness." We're currently just using Foursquare data, but our code also supports the gathering of other information, such as Facebook likes and tags. Thus far, we've refrained from using other data, because we're still investigating how it relates to Foursquare data and the general perception of venue "importance."

Using a heatmap overlaid on the map, the "hot" areas in the city can be made dynamically visible, and this service can show current information (where people are currently checked in) and historically derived information at the same time (which areas are typically hot on a Tuesday afternoon). Furthermore, we can identify and present users' trending venues (that is, those with the greatest rate of increase in total check-ins) or the most popular places, either globally or for the current spatial and temporal context (see Figure 7a). We've also developed a mobile Web app, affording a list-based instead of map-based approach to disseminating information, which is more suitable for small screen devices (Figure 7b).

More interestingly, our current work focuses on a technique to visualize this information in augmented reality (see Figure 7c), with the goal of investigating how the combination of paper maps and information on local social context (real-time check-ins and historical data visualization) can help city visitors uncover and explore stimulating venues. In addition, we're expanding our work to create automated and adaptable guided tours, based on such collected data, which will take tourists through the most interesting parts of a city, given the current day and time. We also plan to provide navigation guidance using the concept of the "most interesting" route instead of the current "shortest path."

Finally, we're interested in applying this information to an adaptive infrastructure (such as city lighting), which can guide people toward the "hot" areas of a city at night (or, conversely, help them avoid it) by colorfully illuminating various streets and areas or modulating light output. Performing the heavy computation and data aggregation in the cloud lets us focus on designing interaction and information retrieval modalities for the various platforms, without worrying about resources and performance issues.

Our research is still underway, but clearly, cloud-sourced, scarce human-physical environment interaction data—such as Foursquare check-in information—can be effectively used to represent the social "buzz" of an urban environment. Scarce check-in data is adequate for building an accurate picture of urban dynamics over time. To our knowledge, we're the first researchers to examine check-in data mined directly from Foursquare instead of indirectly through other social media, such as Twitter. We're also the first to objectively validate the correlation of check-in data against other datasets, instead of relying on subjective user feedback.

A cloud-oriented infrastructure for disseminating this knowledge can drive a variety of services from the same data, in both high- and low-resource devices and device ecologies. We now aim to explore how the inferred knowledge can affect citizens' interactions with the city, when presented to the users in a ubiquitously available manner, such as on desktop and mobile applications, through augmented reality, or through large projections in the urban environment.

ACKNOWLEDGMENTS

We thank N. Milionis for kindly providing the vehicle traffic data from his 2009 study.

REFERENCES

1. J. Lindqvist et al., "I'm the Mayor of My House: Examining Why People Use Foursquare," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 11), ACM, 2011, pp. 2409–2418; doi: 10.1145/1978942.1979295.
2. S. Patil et al., "'Check Out Where I Am!': Location-Sharing Motivations, Preferences, and Practices," *Proc. CHI '12 Extended Abstracts on Human Factors in Computing Systems* (CHI EA 12), ACM, 2012, pp. 1997–2002.
3. M. Zillinger, "The Importance of Guidebooks for the Choice of Tourist Sites: A Study of German Tourists in Sweden," *Scandinavian J. Hospitality and Tourism*, vol. 6, no. 3, 2011, pp. 229–247.
4. B.L. Iaquinto, "Fear of a Lonely Planet: Author Anxieties and the Mainstreaming of a Guidebook," *Current Issues in Tourism*, vol. 14, no. 8, 2011, pp. 705–723.
5. U. Gretzel and K.H. Yoo, "Use and Impact of Online Travel Reviews," *Information and Communication Technologies in Tourism*, Springer, 2008, pp. 35–46.
6. R. White and G. Buscher, "Characterizing Local Interests and Local Knowledge," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 12), ACM, 2012, pp. 1607–1610.
7. H. Chen, A. Namdeo, and M. Bell, "Classification of Road Traffic and Roadside Pollution Concentrations for Assessment of Personal Exposure," *Environmental Modelling & Software*, vol. 23, no. 3, 2008, pp. 282–287.
8. "Specific Air Pollutant Emissions (TERM 028)," European Environmental Agency Indicators, Jan. 2011; www.eea.europa.eu/data-and-maps/indicators/specific-air-pollutant-emissions/specific-air-pollutant-emissions-assessment-3.
9. N. Milionis and I. Aggelidis, *Feasibility Study for a Tram Service in the City of Patras*, tech. report, Development and Engineering Consultants, 2009.
10. A. Noulas et al., "An Empirical Study of Geographic User Activity Patterns in Foursquare," *Proc. 5th Int'l AAAI Conf. Weblogs and Social Media*, AAAI Press, 2011, pp. 150–573.
11. Z. Cheng et al., "Exploring Millions of Footprints in Location Sharing Services," *Proc. 5th Int'l AAAI Conf. Weblogs and Social Media*, AAAI Press, 2011, pp. 81–88.
12. B. Stlering, "42 Major Countries Ranked by Smartphone Penetration Rates," *Wired*, 16 Dec 2011; www.wired.com/beyond_the_beyond/2011/12/42-major-countries-ranked-by-smartphone-penetration-rates.
13. L. Fisher, "Social Media by the Numbers: The Latest Site Figures (May 2012)," blog, 19 May 2012; www.simplyzesty.com/Blog/Article/May-2012/Social-media-by-the-numbers-the-latest-site-figures.
14. M. Satyanarayanan et al., "The Case for VM-Based Cloudlets in Mobile Computing," *IEEE Pervasive Computing*, vol. 8, no. 4, 2009, pp. 14–23.

the AUTHORS

Andreas Komninos is a lecturer of mobile and ubiquitous computing in the School of Engineering & Computing and co-director of the Mobile & Ubiquitous Computing Group (www.mucom.mobi) at Glasgow Caledonian University. He also serves as an associate professor of pervasive computing systems at the Hellenic Open University. His research interests are mobile human computer interaction and ubiquitous context-aware systems. Komninos has a PhD in mobile computing from the University of Strathclyde, Scotland, and is a Chartered Engineer and Chartered IT Professional of the British Computer Society. Contact him at andreas.komninos@gmail.com.

Vasilios Stefanis is a PhD student in the Department of Computer Engineering and Informatics at the University of Patras, Greece. He's also a researcher at the Computer Technology Institute and Press "Diophantus" (CTI). His research interests include mobile and pervasive computing, personal information management, mobile HCI, and Internet technologies. Stefanis has an MSc in computer science and engineering from the University of Patras. Contact him at stefanis@ceid.upatras.gr.

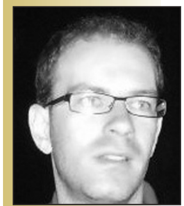

Athanasios Plessas is a PhD student in the Department of Computer Engineering and Informatics at the University of Patras, Greece. He also works as an external researcher at the Computer Technology Institute and Press "Diophantus" (CTI). His research interests include pervasive and mobile computing, Web technologies, and social networks. Plessas has an MSc in computer science and technology from the University of Patras. Contact him at plessas@ceid.upatras.gr.

Jeries Besharat is a PhD student in the Department of Computer Engineering and Informatics at the University of Patras, Greece. He's also a researcher at the Computer Technology Institute and Press "Diophantus" (CTI). His research interests include pervasive and mobile computing, augmented reality, and social networks. Besharat has an MSc in computer science and engineering from the University of Patras. Contact him at besarat@ceid.upatras.gr.