

FlowSampler: Visual Analysis of Urban Flows in Geolocated Social Media Data

Alvin Chua^{1(✉)}, Ernesto Marcheggiani^{2,3}, Loris Servillo¹,
and Andrew Vande Moere¹

¹ Department of Architecture, KU Leuven,
Kasteelpark Arenberg, 1 – Bus 2431, 3001 Heverlee, Belgium
{alvin.chua, loris.servillo, andrew.vandemoere}@asro.kuleuven.be

² Department of Earth and Environmental Sciences, KU Leuven,
Celestijnenlaan 200E – Bus 2411, 3001 Heverlee, Belgium

³ Department of D3A, University of Marche,
Via Breccie Bianche 10, 60131 Ancona, Italy
ernesto.marcheggiani@ees.kuleuven.be

Abstract. Analysis of flows such as human movement can help spatial planners better understand territorial patterns in urban environments. In this paper, we describe FlowSampler, an interactive visual interface designed for spatial planners to gather, extract and analyse human flows in geolocated social media data. Our system adopts a graph-based approach to infer movement pathways from spatial point type data and expresses the resulting information through multiple linked multiple visualisations to support data exploration. We describe two use cases to demonstrate the functionality of our system and characterise how spatial planners utilise it to address analytical task.

Keywords: Social media analytics · Geovisualisation · Spatio-temporal analysis · Data mining · Flow maps

1 Introduction

Urban (or inter-urban) flow analysis is a particularly important subject in spatial planning that identifies territorial patterns in human movement to inform policymaking. Although many techniques have been devised to carryout such analysis [1-3], the growing volume of geolocated social media data presents spatial planners with new opportunities to formulate evidence based policies that could lead to improvements in the urban environment. A key component in analysing human movement is the notion of trajectory. A trajectory provides information about the position of a person through space and time. By analysing patterns in aggregated trajectories, spatial planners aim to identify pathways where important movement or flows occur. The insights that they gain from analysis are used to conceptualise territorial structures, such as functional urban areas [4], that ultimately determine where and how policies are enacted.

Geolocated social media data is a source of publicly accessible data that contains information, which may be extracted to study urban flows. Since such data typically

contain a timestamp and can be referenced to specific user identifiers, it is reasonable to construct a social media user's trajectory based on a chronologically ordered set of geolocated data records. While this task may appear to be outwardly trivial, it can be rather challenging for spatial planners to accomplish with generic GIS software, as the data tend to be large and ill structured.

We present FlowSampler, a visual analytics system designed for spatial planners to gather and analyse urban flows in geolocated social media data. This work was motivated by the need for an interactive visual interface that would extract trajectories out of geolocated social media data and summarise them in a flow map [5]. The strength of this system is that it enables spatial planners to formulate and subsequently verify research questions by reconfiguring the interface for analysis at various spatial and temporal granularities. The interactions are carried out through a series of integrated control widgets that allows the spatial planners to directly manipulate the visualisation. We make three contributions in this paper: First, we propose a graph-based approach to construct a flow map from the trajectories extracted from a geolocated social media dataset. Then, we describe a visual analytics procedure to identify pathways with significant movement between sets of locations. Finally, we demonstrate the functionality and scalability of our system with two use cases that characterise the task this system addresses at different spatial and temporal granularities.

2 Data

Recent growth in smart phones usage [6] and emergence of location aware services has enabled large-scale data collection [7] through participatory sensor networks [8]. A key feature that makes such systems particularly relevant for urban informatics [9] is the ubiquity of the sensors, and the existence of infrastructures that enable sensing. Twitter is an example of a participatory sensor network. It is a microblogging service that allows people to share events and news or have conversations in real time [10]. Empirical studies have shown that people generally use Twitter to describe what they are doing or express how they are feeling [11]. Apart from text content, each tweet is accompanied by a range of meta data such as timestamp and geographic location. We refer to geographically referenced tweets as geolocated tweets. Geographic referencing is not exclusive to Twitter but has been a popular concept, implemented in many other social media services. Depending on individual preferences, Twitter users may decide to publically share their activities on other social media. When they do so, the information posted on those services are also publicised on Twitter. Foursquare, an online service for users to share their whereabouts is an example of such a network. Because Twitter offers a relatively simple protocol to access such information, other studies in literature [12, 13] have also collected geolocated data from other social media through Twitter. For these reasons, we developed our system based on geolocated twitter data. Yet, the concepts we describe can be generalizable to a wider class of geolocated social media data with similar characteristics.

Prior to the availability of geolocated social media data, large-scale studies of mobility were mainly based on cellular activity logs [14-17] that track the spatial position

of people at different moments in time. To analyse movement in cellular datasets, analysts rely on techniques that partition a given territory into subspaces based on the locations of cellular base stations. The position of a cell phone is then approximated to the location of the base station responsible for routing its signal. An estimated trajectory can then be constructed by chronologically ordering the locations of the base stations that served the cell phone. While this approach has revealed valuable insights about human mobility [14, 15], the spatial resolution in which studies can be conducted depends on the physical geometry of the infrastructure. In comparison to geolocated social media datasets that offer spatial information of up to street level precision, the space partitioning technique implies that studies conducted in territories with sparsely distributed base stations will be limited to relatively low resolution spatial analysis. Moreover, cellular datasets are proprietary in nature. In most cases, obtaining such data tends to require a long time to accomplish due to complicated procedures and long discussions with stakeholders.

3 Related Work

There are many examples that take advantage of the fine spatial granularity offered by geolocated social media data to study cities in greater detail. In popular culture, such datasets have been used to create casual visualisations [18] to engage the lay audience [19]. Several prominent examples include maps that show key paths in transport infrastructure [20], track the use of different languages in cities [21, 22] and reveal the distribution of urban wealth [23]. Previous work in literature has also made use of geolocated datasets for a multitude of purposes such as studying or developing technologies to support land use analysis, crisis management and mobility.

Land Use Analysis. Applications that use geolocated social media data for land use analysis are generally concerned with identifying the type of activities that are most common in specific urban areas. Frias-Martinez *et al* described a straightforward procedure that combines a space partitioning technique with human deduction to identify changes in land use over time [24]. Livehoods, a project by Cranshaw *et al* [12] addresses the same issue but adopts an automatic technique to draw alternative neighbourhood boundaries by clustering nearby locations with similar social activities. Their approach illustrates how the fine spatial resolution offered by geolocated tweets can be used to reveal social-spatial divisions in cities. Kling and Pozdnoukhov [13] developed a more sophisticated system that addresses the same issue. However, their work differs from the former in that they extract a chronologically ordered set of keywords to provide analysts with time stamped contextual information of activity on the ground.

Crisis Management Systems. Apart from land use analysis, geolocated social media data also serve as a source of information in crisis management systems. The task addressed by analysts in this domain involves extracting information to monitor situations and explain how they evolve. Studies such as De Longueville *et al*'s analysis of a forest fire near Marseille [25], and Prasetyo *et al*'s investigation of how a severe haze affected the residents of Singapore [26], act as some instances to characterise how such data can be used as a quantifiable source of information in times of crisis.

Mobility Analysis. While there exists a diverse range of work that made use of social media data for land use analysis and crisis monitoring, relatively little has been done to tap its potential for understanding human mobility. Traffic and navigation is one application area where such data have been exploited for mobility analysis. Wei *et al* describes an approach for constructing routes that navigate popular landmarks in cities [27]. Likewise, Pan *et al* address the problem of detecting and describing traffic anomalies by monitoring changes in mobility behaviour [28]. Pan's approach however, does not infer routing information from geolocated social media data but analyses it for contextual information that may be useful to describe events occurring on the ground. Character profiling is another application area where geolocated social media data has been applied. Fuchs *et al* presents an analytical approach to extract knowledge about personal behaviour from geolocated social media data by classifying profiles based on movement trajectories [29]. Andrienko *et al* addresses a similar challenge but classifies profiles based on venues instead [30]. A similarity between the existing works that make use of geolocated social media data to derive mobility information is that they focus on very precise patterns. Our work differs from existing applications in that we are more concerned with identifying general flow pathways between locations in a territory rather than the actual transit route or specific points of interest. In this respect, recent work by Gabrielli et al [31] addresses a similar topic as us yet their intent was to identify semantic rather than spatial patterns.

Visual Analytics. There are two broad approaches to conduct data analysis. Automatic algorithms can be used to address well-defined task with a known set of steps [32] while visual analysis is often required to support explorative task that require human deduction and reasoning. Visual analysis is not new to spatial planning as the discipline has a tradition of using maps for thinking and reasoning [33, 34]. There are several visualisation techniques that are relevant for urban flow analysis. Minard's map of Napoleon's Russian campaign [35] is one of the earliest attempts at visualising flows. The map depicts the size of the French army by the width of a band on the map, and depicts the change in its numbers in relation to air temperature throughout the duration of the campaign. Tobler [36] provides some early examples of computer generated flow maps. Flow maps are maps that show the movement of objects from one location to another [37]. The objects that are represented vary by theme. Flow maps rely on a node-link type representation where lines of different widths are used to represent the direction and quantity of objects being moved. An alternative to the node link representation is an origin destination map [38]. Origin destination maps comprises of a set of origin destination matrices arranged in geographic order. The map is interpreted by tracing a point of origin to a corresponding destination in one of the other matrices. While benchmarks [39] have shown that the matrix representation outperforms the node-link representation in task such as search and quantity estimation, node-link representations are reported to be more effective at path finding, an important task in interpreting the direction and sequence of flows.

4 Design

FlowSampler is a visual analytics system that comprises of four visualisations components each highlighting a different attribute in the data. We have chosen a visual analytics approach, as the task we address is exploratory and requires human judgement for pattern analysis and evaluation. By providing a visual interface, spatial planners can establish and fine tune their analytical procedure, identify uncertainty and biases in the data, and communicate their findings in an interactive visual environment. Our system should allow spatial planners to identify significant flow pathways that connect various locations in a given territory. Specifically, we are interested in (1) flow patterns that exhibit characteristics of routine behaviour and (2) sequence of movements that can be used to characterise spontaneous, unexpected, behaviour. Our model considers a flow path to be routine if it contains a large number of trajectories that are made by few people. Conversely, a flow pathway is considered to exhibit spontaneous characteristics when it is infrequently traversed by a large number of people (fig. 1).

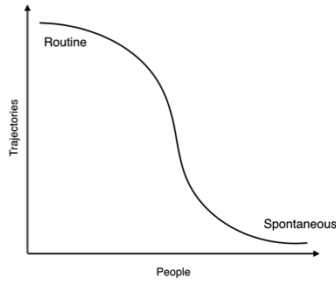


Fig. 1. Model to distinguish between routine and spontaneous characteristics in flow pathways

4.1 Data Transformation

To model urban flows (fig. 1), we propose a data transformation procedure for which we can transform spatial point type data into a graph structure suitable for expressing flow pathways [40]. The procedure consists of two steps. We begin by discretising the territory with an n^2 grid where n^n possible trajectories may occur. Next, we propagate a directed graph $G(V, E)$, where nodes $v_i \in V$ are cells in the grid generated in step 1. A directed edge $E(i, j)$ represents a movement trajectory from node v_i to node v_j if a tweet has been made by a user in cell v_i and cell v_j in chronological order. We adopt two attributes as edge weights: The number of aggregated trajectories $T(i, j)$ and the number of unique people $P(i, j)$ that move between cell v_i and cell v_j . Looping or self directed edges, for instance $E(i, i)$, are also accounted for in the same manner.

4.2 Interface Components

Our interface comprises of four components in a linked, integrated view (fig. 2). It contains a flow map, a time selection widget, a trajectory selection widget and a headcount selection weight. We implemented dynamic zooming in the flow map. Both weight and time selection widgets have dynamic filter ranges for the user to define filter boundaries. Inline with the visual information seeking mantra [41], all filter ranges are set to span the entire distribution while the map is set to the furthest zoom level on initialisation. This provides spatial planners with an overview of the data before further analytical task are carried out. The zoom level of the flow map can be modified with the mouse wheel and the filter range by manipulated by moving the interactive range sliders or selecting individual bars.




Fig. 2. FlowSampler interface components. (A) Flow map with flow pathways represented as arcs. Red arcs represent incoming flows while blue arcs represent outgoing flows. (B) Time selection widget. (C) Trajectory selection widget (D) Headcount selection widget.

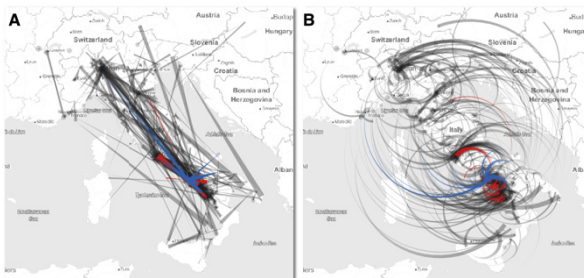


Fig. 3. Comparison between two flow representations. (A) Depending on the physical geometry of the territory and flow patterns in the data, polyline representations may create distracting crossings that are confusing to interpret. (B) Arc representations address this problem by separating the short flow from the long flows.

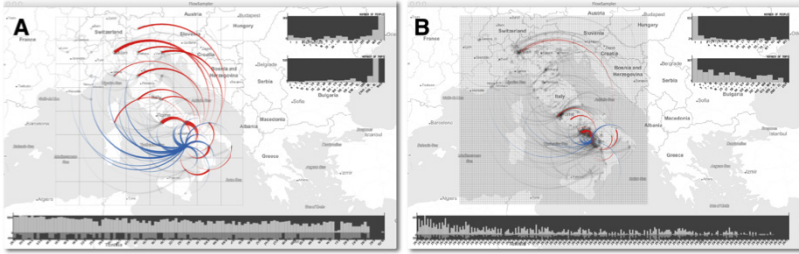


Fig. 4. FlowSampler configured to show flow pathways at two levels of spatial-temporal resolution. (A) Coarse granularity reveals general trends in the data. (B) Fine granularity allows spatial planners to identify outliers and nuances such as movements between locations that only occur at specific hours of the day.

Flow Map. The flow map provides a spatial view of the data. The principle behind flow map is based on a node-link type representation where trajectories are plotted as lines that link origin to destination. While this approach produces maps that are familiar to many people, it does not scale well to large numbers of trajectories. To reduce visual clutter, flow maps merge trajectories that share similar origins and destinations. A line of varying thickness is then used to express the number of trajectories that have been aggregated. Similarly, ellipses of varying diameter are used to represent self-directed flows. We adopt a node-link representation that can be super imposed onto a variety of base maps depending on various planning needs. To reduce ambiguity during interpretation, we represent flows with tapered polylines as recommended in literature [42]. While polylines are effective when flow distances are relatively short, they can become problematic when connections between two distant cells create long diagonal lines that may cause overlap or distracting crossings. To address the problem, we curve the polyline to form an arc. Arcs were chosen for the alternative representation, as they are a computationally cheap solution to avoid distracting crossings by separating the short flows from long flows. A disadvantage of the arc representation however is the added visual complexity it introduces to the display. There are two interactive features that support data exploration. Spatial planners may filter the flow to focus on a specific range of flow pathways in the data or select a cell to highlight the incoming and outgoing flows related to it. We tint selected flows with divergent colours to emphasize directionality.

Time Selection Widget. The time selection widget is an interactive split bar chart where each bar represents a time unit (i.e. week/day/hour) predefined by the spatial planner. Bars are arranged in a chronological order. The height of each bar in the upper half of the bar chart is used to encode the number of trips that occurred during that period while the lower half of the bar chart is used to indicate the total number of flows related to a selected cell. The lower half of the bar chart will be empty if no cells are selected.

Trajectory and Headcount Selection Widget. The trajectory selection widget and the headcount selection widget are histograms that visualise the distribution of the aggregated trajectories and unique people who travelled along a certain flow pathway respectively. The intervals of the histogram are determined by a linear interpolation

by default but spatial planners may dynamically switch the display to a logarithmic interpolation in the event of a highly skewed distribution.

5 Use Cases

5.1 Investigating Daily Routine in Trip Making Behaviour

One purpose of FlowSampler is to experiment with alternative approaches to identify the centres of sub regions in territories based on urban flows. The significance of analysing geolocated social media data is that the information extracted may give an alternative image of how functional urban areas are shaped in comparison to existing techniques that mainly analyse home to work commuting information obtained through census data [4].

Using 734,494 geolocated tweets collected from 2,786 twitter users in Belgium over the duration of a year, we generate a flow map consisting of trajectories belonging to 2,194 users. We omit 592 users because of insufficient tweeting activity. Figure 5a. provides a visual summary of the trips that have occurred over the year. From the map, we identify four distinct clusters that reveal a polycentric distribution of movements. To obtain a map that illustrates routine trip making, we remove the flow pathways that exhibit spontaneous characteristics by filtering flows that fall into the lower percentile of the trajectory selection widget and flows that fall into the upper percentile of the headcount selection widget (fig. 5b). The resulting map characterises the routine trip making behaviour of Twitter users in Belgium. To verify the regularity of the remaining flow pathways, we inspect the frequency of these trips with the time selection widget. Through this process, we discovered that majority of the routine movements take place around local communities typically in towns and villages. Yet, we also observe routine intercity travel across contiguous urban areas between three major Belgian cities (fig. 5c).

5.2 Investigating Exceptional Trip Making Behaviour

To demonstrate the scalability of FlowSampler, we describe an orthogonal use case investigating exceptional, short-term transit behaviour that took place over the touristic season in Italy. The dataset we analyse consist of 13,953,814 tweets generated by 344,660 twitter users over a period of three months. For this study, we are specifically interested in identifying movements that converge on, and take place within, Cilento, a national park in southern Italy. From this data, we construct trajectories belonging to 78,477 twitter users. We omit 266,183 users due to insufficient tweeting activity. The map in figure 6a. illustrates that majority of movement towards Cilento originate from three major Italian cities. We refine the spatial granularity of the map to obtain a more precise boundary over the park and exclude twitter users who were in fact traveling to nearby cities. This reduces the analysis to 1,214 trajectories that transit Cilento. Figure 6b. presents a micro view of the park showing a concentration of twitter users along the coastline. This reveals the extent by which the coastal regions are perceived as privileged destinations in contrast to the inland regions. Visually inspecting the




Fig. 5. (A) Flow map aggregated trajectories that have occurred over the course of a year in Belgium. We identify four distinct clusters that show a polycentric distribution of movement. (B) Flow pathways that exhibit spontaneous characteristics are removed from the analysis. (C) Flow pathways that represent the routine trip making behaviour of Twitter users in Belgium. A particularly striking feature in this map is the connection between three major Belgian cities.

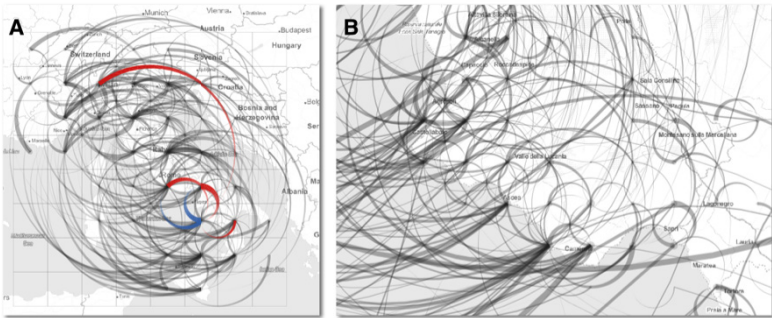


Fig. 6. (A) The arcs highlighted in red indicate that majority of the flow pathways heading towards Cilento Park originate from Milan, Rome and Naples. (B) Micro view of Cilento Park showing a concentration of activity near the coastland.




Fig. 7. Comparing the connectivity of coastal regions to inland regions. We observe that flows occurring in the inland regions are limited to adjacent localities while regions along the coastland are better connected. We tint unrelated flow pathways in a lighter shade to improve the legibility of the image.

incoming and outgoing edges of each node reveals an asymmetry in trip making behaviour. Whereas coastal regions appear well connected to other locations, trip making within the inland regions of Cilento tend to be limited to adjacent localities (fig. 7). This finding corresponds to the availability of transport infrastructure as well as to how the coastline is marketed as a key touristic attraction. Filtering the time selection widget further reveals an “inland, coastal, inland” travel pattern between two disjoint inland regions showing that the settlements along the coastline serve as important hubs for transit between locations.

6 Limitations, Uncertainty and Bias

We have described FlowSampler, a visual analytics system that supports the extraction and exploration of urban flows in geolocated social media data. A key advantage of this system is that it enables planners to interactively reconfigure the interface to explore and detect patterns in the data at various spatial and temporal granularities. Furthermore, our system allows spatial planners to include external geographic information in form of base maps to evaluate the significance of patterns that they have identified. To show the functionality and scalability of our system, we presented two use cases that investigate urban flows at different spatial and temporal scales. We identified pathways of routine movement that occur within Flanders, a region in Belgium, over the duration of a year, and traced exceptional transit activity converging on and subsequently occur within Cilento, a national park in Italy over the touristic season spanning three months. While initial deployment of FlowSampler with the spatial planners in our department has resulted in positive feedback, several discussion topics have been raised.

Skewed Demographic. As existing studies indicate that majority of the online social media users are young adults [43], there is concern that the flow patterns we detect only represent a partial slice of the actual population on the ground. While we acknowledge this limitation, we would like to point out that our approach provides equally valuable and alternative insights that are complimentary to the results derived from other urban flow analysis techniques.

Sporadic Activity. We observe a non-linear distribution of tweeting activity in the form of a long tail where a handful of highly active users are trailed by a substantially larger number of people who tweet sporadically. Because highly active users have trajectories that comprise of many more trips than sporadic user, the uneven distribution implies that certain movement pathways will be over emphasized thus skewing the overall representation. We address this challenge by allowing spatial planners to interactively modify the flow map to determine which attribute edge thickness encodes (i.e. the number of trips or the number of people). This facilitates visual comparison between both attributes in order to identify bias in the representation. Another feasible solution is to pre-filter overly active users to remove the bias entirely from the analysis however this narrows the slice of the population being studied.

Privacy. Our system is designed to present information about aggregated movement behavior yet we provide functionalities for the information to be disaggregated. While we acknowledge that it may be difficult to prevent the recovery of personal information under such circumstances, imposing control measures to displace or distort the data may be counter productive for the spatial planners.

7 Future Work

There are several avenues for future work. We plan to conduct a comparative study with existing urban flow analysis techniques in order to evaluate and better understand the added value and potential pitfalls that may occur when using geolocated social media data to inform spatial planning. To optimise our system, we will experiment with visualisation techniques such as interactive clustering [44-46] to address challenges with visual clutter. Finally, feedback from spatial planners suggests that contextual data such keywords could be useful for characterising flow patterns. The occurrence of special events such as festivals or strikes can be better understood by combining what people say with what they do.

References

1. Asakura, Y., Hato, E.: Tracking survey for individual travel behaviour using mobile communication instruments. *Transportation Research Part C: Emerging Technologies* **12**, 273–291 (2004)
2. Witlox, F.: Evaluating the reliability of reported distance data in urban travel behaviour analysis. *J. Transp. Geogr.* **15**, 172–183 (2007)
3. O’Neill, E., Kostakos, V., Kindberg, T., Schiek, A., Penn, A., Fraser, D.S., Jones, T.: Instrumenting the city: Developing methods for observing and understanding the digital cityscape. In: Dourish, P., Friday, A. (eds.) *UbiComp 2006*. LNCS, vol. 4206, pp. 315–332. Springer, Heidelberg (2006)
4. Servillo, L., Atkinson, R., Smith, I., Russo, A., Sýkora, L.k., Demazière, C., Hamdouche, A.-I.: TOWN, small and medium sized towns in their functional territorial context, Final Report, Espon (2014)
5. Slocum, T.A.: Flow mapping. In: *Thematic cartography and visualization*, pp. 360–370. Prentice hall, Upper Saddle River (1999)
6. Demirbas, M., Bayir, M.A., Akcora, C.G., Yilmaz, Y.S., Ferhatosmanoglu, H.: Crowdsourced sensing and collaboration using twitter. In: *IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM) 2010*, pp. 1–9. IEEE (2010)
7. Silva, T.H., Vaz de Melo, P.O.S., Almeida, J.M.d., Loureiro, A.A.F.: Uncovering properties in participatory sensor networks. In: *Proceedings of the 4th ACM International Workshop on Hot Topics in Planet-Scale Measurement*, pp. 33–38. ACM (2012)
8. Burke, J.A., Estrin, D., Hansen, M., Parker, A., Ramanathan, N., Reddy, S., Srivastava, M.B.: *Participatory sensing*. Center for Embedded Network Sensing (2006)
9. Foth, M., Choi, J.H.-j., Satchell, C.: Urban informatics. In: *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work*, pp. 1–8. ACM (2011)
10. Twitter. <https://business.twitter.com/twitter-101>

11. Java, A., Song, X., Finin, T., Tseng, B.: Why we twitter: understanding microblogging usage and communities. In: Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 Workshop on Web Mining and Social Network Analysis, pp. 56–65. ACM (2007)
12. Cranshaw, J., Schwartz, R., Hong, J.I., Sadeh, N.M.: The livehoods project: utilizing social media to understand the dynamics of a city. In: ICWSM (2012)
13. Kling, F., Pozdnoukhov, A.: When a city tells a story: urban topic analysis. In: Proceedings of the 20th International Conference on Advances in Geographic Information Systems, pp. 482–485. ACM (2012)
14. Gonzalez, M.C., Hidalgo, C.A., Barabasi, A.-L.: Understanding individual human mobility patterns. *Nature* **453**, 779–782 (2008)
15. Song, C., Qu, Z., Blumm, N., Barabási, A.-L.: Limits of predictability in human mobility. *Science* **327**, 1018–1021 (2010)
16. Sevtsuk, A., Ratti, C.: Does urban mobility have a daily routine? Explorations using aggregate mobile network data. In: Proceedings of the 11th International Conference on Computers in Urban Planning and Urban Management (2009)
17. Doyle, J., Hung, P., Farrell, R., McLoone, S.: Population Mobility Dynamics Estimated from Mobile Telephony Data. *Journal of Urban Technology* **21**, 109–132 (2014)
18. Pousman, Z., Stasko, J.T., Mateas, M.: Casual information visualization: Depictions of data in everyday life. *IEEE Trans. Visual Comput. Graphics* **13**, 1145–1152 (2007)
19. Grammel, L., Tory, M., Storey, M.: How information visualization novices construct visualizations. *IEEE Trans. Visual Comput. Graphics* **16**, 943–952 (2010)
20. Citylab. <http://www.citylab.com/commute/2012/02/map-day-how-people-travel-around-city/1131/>
21. Twitter Tongues. <http://twitter.mappinglondon.co.uk>
22. Twitter NYC. <http://ny.spatial.ly>
23. Business Insider. <http://www.businessinsider.com/android-is-for-poor-people-maps-2014-4>
24. Frias-Martinez, V., Soto, V., Hohwald, H., Frias-Martinez, E.: Characterizing urban landscapes using geolocated tweets. In: Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on Social Computing (SocialCom), pp. 239–248. IEEE (2012)
25. De Longueville, B., Smith, R.S., Luraschi, G.: OMG, from here, I can see the flames!: a use case of mining location based social networks to acquire spatio-temporal data on forest fires. In: Proceedings of the 2009 International Workshop on Location Based Social Networks, pp. 73–80. ACM (2009)
26. Prasetyo, P.K., Gao, M., Lim, E.-P., Scollon, C.N.: Social sensing for urban crisis management: The case of singapore haze. In: Jatowt, A., Lim, E.-P., Ding, Y., Miura, A., Tezuka, T., Dias, G., Tanaka, K., Flanagan, A., Dai, B.T. (eds.) SocInfo 2013. LNCS, vol. 8238, pp. 478–491. Springer, Heidelberg (2013)
27. Wei, L.Y., Zheng, Y., Peng, W.C.: Constructing popular routes from uncertain trajectories. In: Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 195–203. ACM (2012)
28. Pan, B., Zheng, Y., Wilkie, D., Shahabi, C.: Crowd sensing of traffic anomalies based on human mobility and social media. In: Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, pp. 334–343. ACM (2013)
29. Fuchs, G., Andrienko, G., Andrienko, N., Jankowski, P.: Extracting personal behavioral patterns from geo-referenced tweets. AGILE (2013)

30. Andrienko, G., Andrienko, N., Fuchs, G., Raimond, A.-M.O., Symanzik, J., Ziemlicki, C.: Extracting semantics of individual places from movement data by analyzing temporal patterns of visits. In: *Proceedings of The First ACM SIGSPATIAL International Workshop on Computational Models of Place (COMP 2013)* (2013)
31. Gabrielli, L., Rinzivillo, S., Ronzano, F., Villatoro, D.: From tweets to semantic trajectories: Mining anomalous urban mobility patterns. In: Nin, J., Villatoro, D. (eds.) *CitiSens 2013. LNCS*, vol. 8313, pp. 25–34. Springer, Heidelberg (2014)
32. Sedlmair, M., Meyer, M., Munzner, T.: Design study methodology: Reflections from the trenches and the stacks. *IEEE Trans. Visual Comput. Graphics* **18**, 2431–2440 (2012)
33. Thomas, J.J., Cook, K.A.: A visual analytics agenda. *IEEE Comput. Graphics Appl.* **26**, 10–13 (2006)
34. DiBiase, D.: Visualization in the earth sciences. *Earth and Mineral Sciences* **59**, 13–18 (1990)
35. Robinson, A.H.: *The thematic maps of Charles Joseph Minard** (1967)
36. Tobler, W.R.: Experiments in migration mapping by computer. *The American Cartographer* **14**, 155–163 (1987)
37. Phan, D., Xiao, L., Yeh, R., Hanrahan, P.: Flow map layout. In: *IEEE Symposium on Information Visualization 2005*, pp. 219–224. IEEE (2005)
38. Wood, J., Dykes, J., Slingsby, A.: Visualisation of origins, destinations and flows with OD maps. *The Cartographic Journal* **47**, 117–129 (2010)
39. Ghoniem, M., Fekete, J., Castagliola, P.: A comparison of the readability of graphs using node-link and matrix-based representations. In: *IEEE Symposium on Information Visualization*, pp. 17–24. IEEE (2004)
40. Guo, D.: Flow mapping and multivariate visualization of large spatial interaction data. *IEEE Trans. Visual Comput. Graphics* **15**, 1041–1048 (2009)
41. Shneiderman, B.: The eyes have it: A task by data type taxonomy for information visualizations. In: *Proceedings: IEEE Symposium on Visual Languages*, pp. 336–343. IEEE (1996)
42. Holten, D., van Wijk, J.J.: A user study on visualizing directed edges in graphs. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 2299–2308. ACM (2009)
43. Duggan, M., Smith, A.: *Social media update 2013*. Pew Internet and American Life Project (2013)
44. Rinzivillo, S., Pedreschi, D., Nanni, M., Giannotti, F., Andrienko, N., Andrienko, G.: Visually driven analysis of movement data by progressive clustering. *Information Visualization* **7**, 225–239 (2008)
45. Andrienko, G., Andrienko, N., Rinzivillo, S., Nanni, M., Pedreschi, D., Giannotti, F.: Interactive visual clustering of large collections of trajectories. In: *IEEE Symposium on Visual Analytics Science and Technology*, pp. 3–10. IEEE (2009)
46. Andrienko, N., Andrienko, G.: Spatial generalization and aggregation of massive movement data. *IEEE Trans. Visual Comput. Graphics* **17**, 205–219 (2011)