

Tracing Networks of Knowledge in the Digital Age

Mirco Musolesi
University College London and The Alan Turing Institute

Abstract

The emergence of new digital technologies has allowed the study of human behaviour at a scale and at level of granularity that were unthinkable just a decade ago. In particular, by analysing the digital traces left by people interacting in the online and offline worlds, we are able to trace the spreading of knowledge and ideas at both local and global scales.

In this article we will discuss how these digital traces can be used to map knowledge across the world, outlining both the limitations and the challenges in performing this type of analysis. We will focus on data collected from social media platforms, large-scale digital repositories and mobile data. Finally, we will provide an overview of the tools that are available to scholars and practitioners for understanding these processes using these emerging forms of data.

Keywords: knowledge networks, information dissemination, social media, spatial networks, network science.

1 Overview

Thanks to the existing and emerging digital technologies, nowadays it is possible to trace networks of knowledge at a scale and at a level of granularity that were unimaginable just a few years ago. In this article we will discuss a series of studies concerning the analysis and mapping of networks of knowledge at different scales through a variety of sources of digital data. We will focus in particular on data from social media, large-scale digital repositories and mobile data.

For instance, networks of knowledge can be extracted by analysing the interactions happening in social network platforms. In particular, models of information diffusion can be defined to represent and understand the spreading of knowledge in groups of users in time and space (Nekovee et al. 2007, Liben-Nowell & Kleinberg 2008, Lazer et al. 2009, Lerman & Ghosh 2010, Kitsak et al. 2010, Miritello et al. 2011, Dodds et al. 2011, Karsai et al. 2011, Pastor-Satorras et al. 2015). A typical way is to describe the process of dissemination of the information using epidemic models, where the process of diffusion does not involve a disease (Anderson & May 1991), but a piece of information. In this way, it is possible to represent knowledge dissemination over networks as epidemic processes and formalise them by means of mathematical models (Keeling & Eames 2005). One of the most fascinating aspects related to tracing networks of knowledge in the digital world is that it is possible to extract data in a digital format and use them to develop and test models of dissemination at a scale. Indeed, this is one of the most compelling examples of applications of big data outside the commercial world (Musolesi 2014).

Diffusion models can be devised to understand how knowledge ‘spreads’ through a network of people in a community, in a society, in a city, in a country or even at planetary scale. Knowledge is usually defined as information acquired through experience or education. Therefore, tracing the diffusion of knowledge corresponds in a sense to tracing the diffusion of information (Hidalgo 2015). We have witnessed the development of a variety of models for describing the spreading and the adoption of knowledge, such as, for example, the generation and acceptance of innovative ideas, starting from the marketing community (Goldenberg et al. 2001). One of the most famous models is indeed the Bass model (Bass 1969), which describes the adoption of innovation in a community by means of a very parsimonious mathematical model.

An interesting class of emerging models is that describing temporal (Tang, Scellato, Musolesi, Mascolo & Latora 2010, Holme & Saramäki 2012) and spatio-temporal networks (Williams & Musolesi 2016). These models can be used for example to identify influencers and mediators in time-varying social networks (Tang, Musolesi, Mascolo, Latora & Nicosia 2010). Spatio-temporal network models can also be used to represent the spreading of information in time and space through networks of people (or places) and study its dynamics.

2 Social media

Social media can be used to track the spreading of information as it happens in real-time at a very large scale. They are extremely popular, involving thousands of millions of users. The first example of information diffusion in the digital world indeed happens through Web blogs, which predate the advent of social media outlets such as Facebook (Ellison et al. 2007) and Twitter (Kwak et al. 2010, Cha et al. 2010). One of the first study of information diffusion in these media can be found in (Gruhl et al. 2004, Adar & Adamic 2005). The geography of Twitter has been analysed in several works, such as (Leetaru et al. 2013). Social media can contribute substantially to opinion formation, as studied for example in (Watts & Dodds 2007).

In (De Domenico et al. 2013) the authors examine the spreading of rumours related to the discovery of the Higgs boson (The ATLAS Collaboration 2012, The CMS Collaboration 2012) in real-time from Twitter data. Twitter is a social network where a user can write microblog posts composed of 140 characters. A post can then be *retweeted* (i.e., reposted) by other users that are following a person or an organisation (*followers*). The availability of data from the platform also allows for the extraction of mathematical models of diffusion in the social network of followers across the globe. Furthermore, the study of the topics discussed in Twitter has attracted the attention of a vast number of researchers in the past years. For example, in (Romero et al. 2011) the authors analyse the adoption of hashtags and they show that for example those that are politically controversial are particularly persistent in the network over time.

Another interesting example of tracing knowledge using digital data is indeed the analysis of software developers' activities in GitHub, an online software development collaborative tool. In (Lima et al. 2014) the authors examine the network of collaborations on different software projects at a global scale, looking at how collaborations unfold, including their geographic characteristics. In GitHub users can create code repositories; every repository has a list of collaborators, who can make changes to the content of the repository. A user can submit changes to the codebase or if they do not want to be a collaborator, there is the possibility of 'forking' a project. The action of fork creates a duplicate of the repository for independent work. By analysing the interactions on GitHub, it is possible to observe patterns of collaborations over time and over space. Indeed, software production is a creative process: collaboration on software projects can be seen as a new way of sharing knowledge and participating to the co-creation of intellectual artefacts.

One of the limitations of this body of work is the fact that often external factors cannot be captured, such as for example the influence of other social media or press outlets, e.g., television and newspapers and other online and offline media. An interesting study about evaluating the external influence of other sources on information spreading can be found in (Myers et al. 2012).

The diffusion of social media has also led scholars to ask questions related to the importance of distance in a world connected by global online networks (Mok et al. 2010). In (Scellato et al. 2010), by analysing four social networks (Brightkite, Foursquare, LiveJournal and Twitter), the authors show that distance still matters in the establishment of social links: we cannot talk about *death of distance* (Cairncross 2001), at least yet, in the case of social media. The resulting geo-social networks show heavy-tailed degree distributions, a tendency to exhibit users with high node locality and social triangles on a local geographic scale. Hristova et al. showed that geo-social networks in different platforms are also highly correlated (Hristova, Noulas, Brown, Musolesi & Mascolo 2016) and this fact can be used for effective link prediction.

Other interesting problems include the identification of influential spreaders of information (Kitsak et al. 2010), the maximisation of the spreading process itself (Kempe et al. 2003), and the impact of homophily in the diffusion of information (Aral et al. 2009). Thanks to the availability of geo-social networks, it is also possible to identify the key spreaders in space (Lima & Musolesi 2012): this allows us to understand how specific people in a spatial network contribute to the dissemination of information and ideas in certain geographic areas, such as neighbourhood, cities or states. It is worth noting that the study of information diffusion in social media has not been limited to text but also focussed on other multimedia content such as photos in Flickr (Cha et al. 2009).

Information can also be political in nature such as in the case of protests (González-Bailón et al. 2011, 2013, Margetts et al. 2015). For example, by retrieving information through social media, it is possible to map the evolution of protests, including the exchange of information, of groups operating at local and a global level. This offers an unprecedented opportunity for studying and understanding political phenomena in real-time. Another very interesting aspect is that, for the first time, researchers are able to follow the spreading of information involving single actors and their thoughts at a very fine-grained granularity both in space and time. Moreover, it is worth noting that this information is not mediated through recollection, but it is instantaneous in nature, i.e., it is generated (and can be collected) as events unfold.

It is worth noting that social media data have been used for a very large number of studies, for example for understanding not only the spreading of information itself, but also the inherent contextual factors influence it, such as demographics and culture. For example, one of the topics of interest of social media studies has been the analysis of surnames, which can be used to understand the geo-demographics of different areas of large cities, such as London (Longley & Adnan 2016). In this way, it is possible to trace back global networks of people and migrations reflected in the ethnic composition of large-scale cities. Geo-social networks can also

be used effectively to trace and study diversity and gentrification phenomena (Hristova, Williams, Musolesi, Panzarasa & Mascolo 2016). Finally, in (Silva et al. 2014) the authors discuss how data from the Foursquare platform can be used to understand cultural characteristics of geographic areas by looking at the presence of specific food and drinks venues.

3 Large-scale data repositories

Wikipedia, an online encyclopaedia that is completely edited by volunteers, is probably the most successful example of crowdsourced work. Through the analysis of Wikipedia usage in different languages and the application of data mining techniques to the content of the various articles, it is possible to reconstruct how different topics are of interest and relevant in different part of the globe (Brandes et al. 2009, Yasseri et al. 2012). Moreover, the articles themselves can be geo-located, i.e., mapped to different areas of the planet. In other words, it is truly possible to create maps of knowledge that connect people, objects, facts and places around the world. In order to extract information from the text usually natural language techniques (NLP) techniques (Manning & Schütze 1999) are applied. With respect to the extraction of geographic information, geo-localisation techniques are used to map names to geographic named entities. The mapping is performed by trying to map names to geographic locations that are available in geographic gazetteers such as Geonames (Ahlers 2013).

As more general trend, we observe that citizens are getting more and more involved in the creation of knowledge, both contributing through text and multimedia content (images, videos, etc.) but also by acting as sensors and collecting information about the environments where they live (Goodchild 2007). One of the most interesting examples is OpenStreetMap (*OpenStreetMap* 2016). Through the collaboration of a large number of individuals, knowledge is created, shared and disseminated. Several researchers have been investigated how this collaboration takes place, characterising his dynamics (Hristova et al. 2013).

Other repositories that allow for the analysis of networks of knowledge include online literary repositories, such as for example Gutenberg (*The Gutenberg project* 2016). The analysis of the networks emerging in the literature have been a theme of the work of Moretti among the others (Moretti 2005). Indeed, one possibility is to try to extract names of geographic locations as discussed above for Wikipedia from text in order to reconstruct how literary works are mapped in the geographic space, in order to understand how the fictional worlds are related to the real ones.

Another interesting example is the tracing of influence through citation networks (Bilke & Peterson 2001). These can be traced both over time and space and they provide a way for mapping out collaborations among individuals in different countries in a very detailed way considering both exchange information across disciplines (Porter & Rafols 2009) and in specific communities (Lehmann et al. 2003).

4 Mobility data

Another interesting source of data is related to the increasing availability of mobility traces datasets, from flight (Guimera et al. 2005) to migration (Davis et al. 2013) data, from social media (Noulas et al. 2012, Hawelka et al. 2014) to mobile data collected directly by means of users' smartphones (Campbell et al. 2008) or by the cellular phone operators through the network infrastructure (Laurila et al. 2012).

In this case, the network of knowledge emerging from the data is in a sense indirect, but through this information we can map how people move at different geographic scale and spread ideas. Indeed, as discussed above, some people have talked about the death of distance (Cairncross 2001) as a consequence of the emergence of digital technologies, including for example the use of smartphones and global network connectivity. However, this claim is disputed by other researchers that are stressing the fact that distance is still a major factor in human relationships and, therefore, in spreading of ideas (Scellato et al. 2010).

In general, one of the most difficult aspects of using mobility data for understanding spreading processes is the fact that data provide only a partial view of the interactions between people. Indeed, it is extremely difficult to collect ground-truth information for reconstructing interactions from these datasets. For these reasons, it is always very important to be extremely cautious in deriving universal models from these datasets in terms of human behaviour and interactions.

5 Tools

In this section, we will briefly review the tools that are used for tracing networks from digital data. We can broadly divide them into 3 main categories:

Data Collection Data collection is a fundamental phase in the process of tracing networks of knowledge. Data is usually collected from digital sources such as the Web through scraping tools (Mitchell 2015) or from digital

platforms by means of Application Programming Interfaces (APIs) such as those provided by Twitter (*Twitter REST API* 2016). A limited amount of data is usually available for free, such as in the case of Twitter.

Data Analysis Data analysis is performed with a variety of tools from stand-alone applications such as Gephi (*Gephi - The Open Graph Viz Platform* 2016) to purpose-built software developed using existing libraries such as NetworkX (for Python) (*Networkx Library* 2016) or igraph (available for a variety of programming languages) (*igraph* 2016).

Data Visualisation Various tools can be used for visualising the networks of knowledge extracted from the data at local, national and global scales. An emerging class of tools is represented by open source tools for visualisation. Examples include again Gephi (*Gephi - The Open Graph Viz Platform* 2016) or Cytoscape (*Cytoscape: An Open Source Platform for Complex Network Analysis* 2016). Network data visualisation is also performed through tools developed by means of 'general-purpose' programming languages and platforms for data manipulation and visualisation such as Processing (Reas & Fry 2014). Other data visualisation tools include classic Geographic Information Systems (Longley et al. 2015), which provide functionalities for mapping networks in the geographic space and more advanced digital cartographic visualisations (Cheshire & Uberti 2014).

6 Outlook

In this paper, we have examined new emerging and fascinating ways of tracing networks of knowledge in real-time using the digital traces that we leave in our everyday lives. We have explored the new possibilities offered by the analysis of a variety of data sources, including social media, large-scale data repositories and mobile data.

The availability of these new sources of data allows for new types of investigation that were not possible just a few years ago. At the same time, it is important to consider the limitations that are inherent in using these new forms of data for mapping social phenomena. One of the key issues is indeed associated to the intrinsic bias that is associated to these data sources. For example, the demographics of social media users might not be representative of the general population (Mislove et al. 2011). However, it is worth noting that these sources are still valuable for understanding the behaviour of particular subsets of the population. In general, it is still possible to derive population-level observations taking into consideration at the same time the limitations of the data sources used for the analysis. Other important considerations are related to the potential privacy issues associated with these new emerging sources of data (Musolesi 2014). It is worth noting that most of the data cited in this article are the results of aggregation, i.e., they preserve the identity of the single users. Most of the existing work has been based on extracting, mapping and visualising networks of knowledge. We believe that modelling and predicting the dissemination of knowledge represent a very interesting and still open areas for researchers, also considering the recent results in both network science (Boccaletti et al. 2006, Barrat et al. 2008, Newman 2010, Pastor-Satorras et al. 2015) and machine learning (Murphy 2012). This is indeed an area where truly interdisciplinary work is essential given the fact that the research questions in this field are really at the interface of the humanities, social sciences, computational and mathematical sciences.

Biographical Note

Mirco Musolesi is a Reader in Data Science at the Department of Geography at University College London and a Faculty Fellow at the Alan Turing Institute. At UCL he leads the Intelligent Social Systems Lab. He received a PhD in Computer Science from University College London and a Master in Electronic Engineering from the University of Bologna. After postdoctoral work at Dartmouth College and Cambridge, he held academic posts at St Andrews and Birmingham. The research focus of his lab is on sensing, modelling, understanding and predicting human behaviour in space and time, at different scales, using the 'digital traces' we generate daily in our online and offline lives. He is interested in developing mathematical and computational models as well as implementing real-world systems based on them. This work has applications in a variety of domains, such as ubiquitous systems design, healthcare and security&privacy.

References

- Adar, E. & Adamic, L. A. (2005), Tracking information epidemics in blogspace, in 'Proceedings of the 2005 IEEE/WIC/ACM International Conference on Web Intelligence', IEEE Computer Society, pp. 207–214.
- Ahlers, D. (2013), Assessment of the accuracy of geonames gazetteer data, in 'Proceedings of the 7th Workshop on Geographic Information Retrieval', ACM, pp. 74–81.
- Anderson, R. M. & May, R. M. (1991), *Infectious diseases of humans: dynamics and control*, Vol. 28, Oxford University Press, Oxford and New York.

- Aral, S., Muchnik, L. & Sundararajan, A. (2009), 'Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks', *Proceedings of the National Academy of Sciences* **106**(51), 21544–21549.
- Barrat, A., Barthelemy, M. & Vespignani, A. (2008), *Dynamical Processes on Complex Networks*, Cambridge University Press.
- Bass, F. M. (1969), 'A new product growth for model consumer durables', *Management Science* **15**(5), 215–227.
- Bilke, S. & Peterson, C. (2001), 'Topological properties of citation and metabolic networks', *Physical Review E* **64**(3), 036106.
- Boccaletti, S., Latora, V., Moreno, Y., Chavez, M. & Hwang, D.-U. (2006), 'Complex networks: Structure and dynamics', *Physics reports* **424**(4), 175–308.
- Brandes, U., Kenis, P., Lerner, J. & van Raaij, D. (2009), Network analysis of collaboration structure in Wikipedia, in 'Proceedings of WWW'09', ACM, pp. 731–740.
- Cairncross, F. (2001), *The death of distance: How the communications revolution is changing our lives*, Harvard Business Press.
- Campbell, A. T., Eisenman, S. B., Lane, N. D., Miluzzo, E., Peterson, R., Lu, H., Zheng, X., Musolesi, M., Fodor, K. & Ahn, G.-S. (2008), 'The rise of people-centric sensing', *IEEE Internet Computing Special Issue on Mesh Networks*.
- Cha, M., Haddadi, H., Benevenuto, F. & Gummadi, P. K. (2010), 'Measuring user influence in twitter: The million follower fallacy.', *Proceedings of ICWSM'10* **10**(10-17), 30.
- Cha, M., Mislove, A. & Gummadi, K. P. (2009), A measurement-driven analysis of information propagation in the Flickr social network, in 'Proceedings of WWW'09', ACM, pp. 721–730.
- Cheshire, J. & Uberti, O. (2014), *London The Information Capital*, Particular Books.
- Cytoscape: An Open Source Platform for Complex Network Analysis* (2016). <http://www.cytoscape.org> [Accessed on 1 October 2016].
- Davis, K. F., D'Odorico, P., Laio, F. & Ridolfi, L. (2013), 'Global spatio-temporal patterns in human migration: a complex network perspective', *PLoS ONE* **8**(1), e53723.
- De Domenico, M., Lima, A., Mougél, P. & Musolesi, M. (2013), 'The anatomy of a scientific rumor', *Scientific Reports* **3**(02980).
- Dodds, P. S., Harris, K. D., Kloumann, I. M., Bliss, C. A. & Danforth, C. M. (2011), 'Temporal patterns of happiness and information in a global social network: Hedonometrics and Twitter', *PLoS ONE* **6**(12), e26752.
- Ellison, N. B., Steinfield, C. & Lampe, C. (2007), '“the benefits of facebook “friends”: Social capital and college students? use of online social network sites”', *Journal of Computer-Mediated Communication* **12**(4), 1143–1168.
- Gephi - The Open Graph Viz Platform* (2016). <https://gephi.org> [Accessed on 1 October 2016].
- Goldenberg, J., Libai, B. & Muller, E. (2001), 'Talk of the network: A complex systems look at the underlying process of word-of-mouth', *Marketing letters* **12**(3), 211–223.
- González-Bailón, S., Borge-Holthoefer, J. & Moreno, Y. (2013), 'Broadcasters and hidden influentials in online protest diffusion', *American Behavioral Scientist* p. 0002764213479371.
- González-Bailón, S., Borge-Holthoefer, J., Rivero, A. & Moreno, Y. (2011), 'The dynamics of protest recruitment through an online network', *Scientific Reports* **1**(197).
- Goodchild, M. F. (2007), 'Citizens as sensors: the world of volunteered geography', *GeoJournal* **69**(4), 211–221.
- Gruhl, D., Guha, R., Liben-Nowell, D. & Tomkins, A. (2004), Information diffusion through blogspace, in 'Proceedings of WWW'04', ACM, pp. 491–501.
- Guimera, R., Mossa, S., Turtschi, A. & Amaral, L. N. (2005), 'The worldwide air transportation network: Anomalous centrality, community structure, and cities' global roles', *Proceedings of the National Academy of Sciences* **102**(22), 7794–7799.
- Hawelka, B., Sitko, I., Beinatz, E., Sobolevsky, S., Kazakopoulos, P. & Ratti, C. (2014), 'Geo-located Twitter as proxy for global mobility patterns', *Cartography and Geographic Information Science* **41**(3), 260–271.

- Hidalgo, C. (2015), *Why information grows: The evolution of order, from atoms to economies*, Basic Books.
- Holme, P. & Saramäki, J. (2012), 'Temporal networks', *Physics Reports* **519**(3), 97–125.
- Hristova, D., Noulas, A., Brown, C., Musolesi, M. & Mascolo, C. (2016), 'A multilayer approach to multiplexity and link prediction in online geo-social networks', *EPJ Data Science* **5**(24).
- Hristova, D., Quattrone, G., Mashhadi, A. J. & Capra, L. (2013), The life of the party: Impact of social mapping in openstreetmap., in 'Proceedings of ICWSM'13'.
- Hristova, D., Williams, M. J., Musolesi, M., Panzarasa, P. & Mascolo, C. (2016), Measuring urban social diversity using interconnected geo-social networks, in 'Proceedings of WWW'16'.
- igraph (2016). <http://igraph.org/> [Accessed on 1 October 2016].
- Karsai, M., Kivela, M., Pan, R. K., Kaski, K., Kertész, J., Barabási, A.-L. & Saramäki, J. (2011), 'Small but slow world: How network topology and burstiness slow down spreading', *Physical Review E* **83**(2), 025102.
- Keeling, M. & Eames, K. (2005), 'Networks and epidemic models', *Journal of the Royal Society Interface* **2**(4), 295–307.
- Kempe, D., Kleinberg, J. & Tardos, É. (2003), Maximizing the spread of influence through a social network, in 'Proceedings of KDD'03', ACM, pp. 137–146.
- Kitsak, M., Gallos, L., Havlin, S., Liljeros, F., Muchnik, L., Stanley, H. & Makse, H. (2010), 'Identification of influential spreaders in complex networks', *Nature Physics* **6**(11), 888–893.
- Kwak, H., Lee, C., Park, H. & Moon, S. (2010), What is Twitter, a social network or a news media?, in 'Proceedings of WWW'10', pp. 591–600.
- Laurila, J. K., Gatica-Perez, D., Aad, I., Bornet, O., Do, T.-M.-T., Dousse, O., Eberle, J., Miettinen, M. et al. (2012), The mobile data challenge: Big data for mobile computing research, in 'Pervasive Computing'.
- Lazer, D., Pentland, A. S., Adamic, L., Aral, S., Barabasi, A. L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M. et al. (2009), 'Life in the network: the coming age of computational social science', *Science* **323**(5915), 721.
- Leetaru, K., Wang, S., Cao, G., Padmanabhan, A. & Shook, E. (2013), 'Mapping the global Twitter heartbeat: The geography of Twitter', *First Monday* **18**(5).
- Lehmann, S., Lautrup, B. & Jackson, A. (2003), 'Citation networks in high energy physics', *Physical Review E* **68**(2), 026113.
- Lerman, K. & Ghosh, R. (2010), 'Information contagion: An empirical study of the spread of news on Digg and Twitter social networks', *ICWSM* **10**, 90–97.
- Liben-Nowell, D. & Kleinberg, J. (2008), 'Tracing information flow on a global scale using internet chain-letter data', *Proceedings of the National Academy of Sciences* **105**(12), 4633–4638.
- Lima, A. & Musolesi, M. (2012), Spatial dissemination metrics for location-based social networks, in 'Proceedings of UbiComp'12', ACM, pp. 972–979.
- Lima, A., Rossi, L. & Musolesi, M. (2014), Coding Together at Scale: GitHub as a Collaborative Social Network, in 'Proceedings of ICWSM'14', Ann Arbor, Michigan, USA.
- Longley, P. A. & Adnan, M. (2016), 'Geo-temporal Twitter demographics', *International Journal of Geographical Information Science* **30**(2), 369–389.
- Longley, P. A., Goodchild, M. F., Maguire, D. J., Rhind, D. W. et al. (2015), *Geographic Information Systems and Science.*, John Wiley & Sons Ltd.
- Manning, C. D. & Schütze, H. (1999), *Foundations of Statistical Natural Language Processing*, MIT Press.
- Margetts, H., John, P., Hale, S. & Yasseri, T. (2015), *Political Turbulence: How Social Media Shape Collective Action*, Princeton University Press.
- Miritello, G., Moro, E. & Lara, R. (2011), 'Dynamical strength of social ties in information spreading', *Physical Review E* **83**(4), 045102.

- Mislove, A., Lehmann, S., Ahn, Y.-Y., Onnela, J.-P. & Rosenquist, J. N. (2011), 'Understanding the Demographics of Twitter Users', *ICWSM* **11**, 5th.
- Mitchell, R. (2015), *Web Scraping with Python: Collecting Data from the Modern Web*, O'Reilly Media, Inc.
- Mok, D., Wellman, B. & Carrasco, J. (2010), 'Does distance matter in the age of the internet?', *Urban Studies* **47**(13), 2747–2783.
- Moretti, F. (2005), *Graphs, Maps, Trees: Abstract Models for a Literary History*, Verso.
- Murphy, K. P. (2012), *Machine learning: a probabilistic perspective*, MIT press.
- Musolesi, M. (2014), 'Big mobile data mining: Good or evil?', *IEEE Internet Computing* **18**(1), 78–81.
- Myers, S., Zhu, C. & Leskovec, J. (2012), Information diffusion and external influence in networks, in 'Proceedings of the 18th International Conference on Knowledge Discovery and Data Mining', ACM, pp. 33–41.
- Nekovee, M., Moreno, Y., Bianconi, G. & Marsili, M. (2007), 'Theory of rumour spreading in complex social networks', *Physica A: Statistical Mechanics and its Applications* **374**(1), 457–470.
- Networkx Library* (2016). <https://networkx.github.io> [Accessed on 1 October 2016].
- Newman, M. (2010), *Networks: an Introduction*, Oxford University Press.
- Noulas, A., Scellato, S., Lambiotte, R., Pontil, M. & Mascolo, C. (2012), 'A tale of many cities: universal patterns in human urban mobility', *PloS ONE* **7**(5), e37027.
- OpenStreetMap* (2016). <http://www.openstreetmap.org> [Accessed on 1 October 2016].
- Pastor-Satorras, R., Castellano, C., Van Mieghem, P. & Vespignani, A. (2015), 'Epidemic processes in complex networks', *Reviews of modern physics* **87**(3), 925.
- Porter, A. & Rafols, I. (2009), 'Is science becoming more interdisciplinary? measuring and mapping six research fields over time', *Scientometrics* **81**(3), 719–745.
- Reas, C. & Fry, B. (2014), *Processing: a Programming Handbook for Visual Designers and Artists*, MIT Press.
- Romero, D., Meeder, B. & Kleinberg, J. (2011), Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on Twitter, in 'Proceedings of WWW'11', ACM, pp. 695–704.
- Scellato, S., Mascolo, C., Musolesi, M. & Latora, V. (2010), Distance matters: Geo-social metrics for online social networks., in 'Proceedings of WOSN'10'.
- Silva, T., Vaz De Melo, P., Almeida, J., Musolesi, M. & Loureiro, A. (2014), You are What you Eat (and Drink): Identifying Cultural Boundaries by Analyzing Food & Drink Habits in Foursquare, in 'Proceedings of ICWSM'14', Ann Arbor, Michigan, USA.
- Tang, J., Musolesi, M., Mascolo, C., Latora, V. & Nicosia, V. (2010), Analysing information flows and key mediators through temporal centrality metrics, in 'Proceedings of WOSN'10', ACM, p. 3.
- Tang, J., Scellato, S., Musolesi, M., Mascolo, C. & Latora, V. (2010), 'Small-world behavior in time-varying graphs', *Physical Review E* **81**(5), 055101.
- The ATLAS Collaboration (2012), 'Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC', *Physics Letters B* **716**(1), 1 – 29.
- The CMS Collaboration (2012), 'Observation of a new boson at a mass of 125 gev with the cms experiment at the lhc', *Physics Letters B* **716**(1), 30 – 61.
- The Gutenberg project* (2016). <https://www.gutenberg.org> [Accessed on 1 October 2016].
- Twitter REST API* (2016). <https://dev.twitter.com/rest/public> [Accessed on 1 October 2016].
- Watts, D. J. & Dodds, P. S. (2007), 'Influentials, networks, and public opinion formation', *Journal of consumer research* **34**, 441–458.
- Williams, M. J. & Musolesi, M. (2016), 'Spatio-temporal networks: reachability, centrality and robustness', *Royal Society Open Science* **3**(6).
- Yasserli, T., Sumi, R., Rung, A., Kornai, A. & Kertész, J. (2012), 'Dynamics of conflicts in Wikipedia', *PloS ONE* **7**(6), e38869.