

# On the Use of Social Trajectory-Based Clustering Methods for Public Transport Optimization

Jordi Nin<sup>1</sup>(✉), David Carrera<sup>1</sup>, and Daniel Villatoro<sup>2</sup>

<sup>1</sup> Barcelona Supercomputing Center (BSC),  
Universitat Politècnica de Catalunya (BarcelonaTech), Barcelona, Catalonia, Spain  
nin@ac.upc.edu, dcarrera@ac.upc.edu

<sup>2</sup> Barcelona Digital Technology Centre, Barcelona, Catalonia, Spain  
dvillatoro@bdigital.org

**Abstract.** Public transport optimisation is becoming everyday a more difficult and challenging task, because of the increasing number of transportation options as well as the increase of users. Many research contributions about this issue have been recently published under the umbrella of the smart cities research. In this work, we sketch a possible framework to optimize the tourist bus in the city of Barcelona. Our framework will extract information from Twitter and other web services, such as Foursquare to infer not only the most visited places in Barcelona, but also the trajectories and routes that tourist follow. After that, instead of using complex geospatial or trajectory clustering methods, we propose to use simpler clustering techniques as  $k$ -means or DBScan but using a real sequence of symbols as a distance measure to incorporate in the clustering process the trajectory information.

**Keywords:** Smart cities · Geospatial clustering · Metric spaces · OSA distance · Cloud computing · High performance computing

## 1 Introduction

Trajectory clustering algorithms [24] group similar trajectories into groups (clusters), thus discovering common trajectories. These methods are different from geospatial clustering [38]. This latter clustering methods group similar objects (points in an Euclidian or geodesic space) based on their distance, connectivity, or relative density in the space.

Since geospatial clustering methods are in general easier than trajectory clustering algorithms, they has been employed in the field of spatial analysis for years. Spatial clustering is the process of grouping similar objects based on their distance, connectivity, or relative density in space. Thanks to that, such methods prevail over trajectory clustering algorithms in pattern recognition smart cities applications, such as public transport optimization. The main problem behind

the use of geospatial clustering methods for this purpose is that patterns (most of times trajectories) must be built over geospatial clustering results. To do that, expert knowledge is usually required making this approach less useful for big cities and highly dynamic city environments.

Since trajectory clustering algorithms include by definition all the trajectory information, it is possible to adjust these algorithms to reduce the expert's time to take the decision about which global trajectories are the most interesting ones. In addition, trajectory clustering algorithms can help the data practitioners to discover common sub-trajectories inside a cluster. This information can be very valuable in many applications, especially if we have regions of special interest for analysis inside a big city, such as, the city centre or certain regions with a high number of touristic attractions or city services.

All trajectory clustering algorithms are based on computing a distance assuming that trajectory elements are represented by means of space coordinates. This lack of element semantics makes that clusters construction only considers space similarities. In this paper, we would like to study how to convert trajectory coordinates to symbols, in such a way, we can include semantics inside the trajectory elements. Therefore, it will be possible to group similar trajectories considering the nature of their elements, for instance if they are touristic attractions, city services, restaurants, etc. The main drawback of this approach is that it is required to define an appropriate distance for this type of sequences.

In this paper we propose to use the Optimal Symbol Alignment (OSA) distance [18] for *semantic-aware* trajectory clustering. To do that, we would like to integrate it into the clustering framework ELKI [1] to perform some experiments using trajectories about tourists visiting Barcelona.

## 1.1 Paper Organization

The rest of this paper is organized as follows. Firstly, in Sect. 2 we introduce some basic concepts about sequences of symbols distances. We also provide a complete definition of the Edit and OSA distances. Then, in Sect. 3 we provide a taxonomy for the currently used clustering algorithms, as well as, some implementation decisions we have already taken. In Sect. 4 we describe how we have obtained the data required for our clustering analysis using several social network services. Later, in Sect. 5, we mention several computational issues we must consider due to the high amount of data available. Finally, Sect. 6 depicts the following steps in our proposal.

## 2 Distances for Sequences of Symbols

Comparison functions for sequences (of symbols) are important components of many applications, for example clustering, data cleansing and integration.

For this reason, there is a lot of work in computing similarities among sequences of symbols [8, 15, 29]. However, the similarity measures presented in most of those works either do not fulfil the mandatory conditions to be a real

distance (most of the times, because the triangular inequality does not hold) or do not contain a proof for that. Formally, a distance function  $d$  must satisfy the following properties:

1. Symmetry:  $d(A, B) = d(B, A)$  for all sequences  $A, B$
2. Positivity:  $d(A, B) \geq 0$  for all sequences  $A, B$
3. Reflexivity:  $d(A, A) = 0$  for all sequence  $A$
4. Triangular Inequality:  $d(A, B) \leq d(A, C) + d(B, C)$  for all sequences  $A, B, C$

To the best of our knowledge, there are only three sequence measures that fulfil these conditions: the Hamming distance [17], the Levenshtein (Edit) distance [25] and the OSA Distance [18]. The remaining measures are similarity functions instead of real distances because they do not comply with the triangular inequality (or this is not proved). For this reason the application of such measures to the scenarios where having a metric space is a must, such as metric spaces [5], clustering [19] or k-nearest neighbors algorithms [6], becomes unfeasible from a theoretical point of view.

The Hamming and Edit distances present also some problems. For instance, the Hamming distance can only be applied to sequences of the same length, while the Edit distance has a large, both practical and theoretical, complexity ( $O(n^2)$ ). For these reasons many similarity measures have been developed, albeit sacrificing some of the mandatory properties of a distance. For example, the Jaro-Winkler distance [21] is very efficient in terms of practical computational cost when the compared strings are not too large. Therefore, it saves execution time (when compared to Edit distance) in applications where there are many comparisons to be done.

## 2.1 Edit Distance

The Edit distance [25, 32] measures the difference between two sequences, given by the minimum number of edit operations needed to transform one sequence into the other. An edit operation can be either an insertion, deletion or substitution of a single symbol, although many variations exist in which the set of allowed operations is larger or more restricted. In some way, Edit distance assumes that the differences between two sequences are due to typos or spelling errors.

The Edit distance has found a large variety of applications in many scenarios and has achieved very good results [31]. However, the Edit distance has a large complexity: its computation using classical algorithms [36] based on dynamic programming has a complexity equal to  $O(n^2)$ , where  $n$  is the size of the shortest string. Other algorithms to compute the Edit distance exist [4, 34], having a lower complexity of  $O(dn)$ , for example, where  $d$  is the real Edit distance. Note that, when two completely different strings have to be compared, the complexity of these variants is the same as that of the classical algorithm.

## 2.2 Optimal Symbol Alignment (OSA) Distance

The intuition behind the Optimal Symbol Alignment (OSA) distance [18] is that strings are close if they have many common symbols, and in addition their common symbols are placed in similar positions, in the strings being compared.

Given a finite alphabet of symbols  $X$ , let  $A = (a_1, \dots, a_{n_A})$  and  $B = (b_1, \dots, b_{n_B})$  be two sequences of symbols, where  $a_i, b_j \in X$ , for  $i = 1, \dots, n_A$ ,  $j = 1, \dots, n_B$ . For any sequence of symbols  $A$ , we define as  $X_A \subseteq X$  the subset of symbols that appear in  $A$ ; that is,  $X_A = \{x \in X \text{ s.t. } \exists i \in \{1, 2, \dots, n_A\} \text{ with } a_i = x\}$ . For a symbol  $x \in X_A$ , we also define the subset of positions  $A_x = \{i \in \{1, \dots, n_A\} \text{ s.t. } a_i = x\}$ .

We define the OSA distance  $d(A, B)$  between the sequences  $A$  and  $B$  as

$$d(A, B) = \sum_{x \in X_A \cup X_B} d(x, A, B),$$

where the value  $d(x, A, B)$  is defined as

$$d(x, A, B) = \begin{cases} |A_x| & \text{if } x \in X_A - X_B \\ |B_x| & \text{if } x \in X_B - X_A \\ f(x, A, B) & \text{if } x \in X_A \cap X_B \end{cases}$$

Finally, we have to define the value of  $f(x, A, B)$ , which is the contribution of the symbol  $x$  to the distance  $d(A, B)$ , when this symbol  $x$  is included in both sequences  $A$  and  $B$ . Let us assume without loss of generality that  $|A_x| \leq |B_x|$ . The idea is to select the subset of  $|A_x|$  positions  $j$ , from the set  $B_x$ , which are globally closest to the set of  $|A_x|$  positions in  $A_x$ . Namely, if  $i_1 < i_2 < \dots < i_{|A_x|}$  are the positions in  $A_x$ , then we select  $|A_x|$  positions  $j_1 < j_2 < \dots < j_{|A_x|}$  in  $B_x$  minimizing the global distance  $|i_1 - j_1| + \dots + |i_{|A_x|} - j_{|A_x|}|$ . We use notation  $j_h = \mathbf{pj}(i_h, A, B)$ , for  $h = 1, \dots, |A_x|$ , to denote the position in  $B_x$  that optimally matches position  $i_h \in A_x$ . We say that  $j_h$  is the *projection* of position  $i_h$  from sequence  $A$  to sequence  $B$ . For completeness, we also use the symmetric notation  $i_h = \mathbf{pj}(j_h, B, A)$ .

Each of these common symbols  $a_{i_h} = b_{j_h} = x$ , for  $h = 1, \dots, |A_x|$ , will contribute with  $\frac{|i_h - j_h|}{n_{AB}}$  to the value  $f(x, A, B)$ , where  $n_{AB} = \max\{n_A, n_B\}$ . In this way, we ensure that these contributions are bounded by 1. The remaining  $|B_x| - |A_x|$  symbols will be considered as non-common symbols, so each of them will contribute with a 1 to the global distance  $d(A, B)$ .

Taking all these facts into account, we finally have

$$f(x, A, B) = (|B_x| - |A_x|) + \frac{1}{n_{AB}} \sum_{i_h \in A_x} |i_h - \mathbf{pj}(i_h, A, B)|.$$

Depending on the differences between the two sequences to be compared (more or less repeated symbols, more or less transpositions, etc.) the OSA distance  $d_{\text{OSA}}(A, B)$  will be more or less similar to the Edit distance  $d_{\text{Edit}}(A, B)$ . But in any case, they will not be very far, because it is easy to prove that  $\frac{d_{\text{Edit}}(A, B)}{2} \leq d_{\text{OSA}}(A, B) \leq 2 \cdot d_{\text{Edit}}(A, B)$ , for any two sequences  $A, B$ .


### 3 Clustering

Clustering is the task of relating similar elements in a dataset to build groups and create general representations (centroids). It is widely used in many fields as data mining, machine learning, image analysis, etc. Since clustering is a very general concept, there are a great variety of algorithms that apply clustering, focusing, for example, on possible centroids, density, data structures, etc.


Clustering algorithms are divided into different categories: space partitioning, also called top-down methods, hierarchical methods, known as bottom-up methods as well or Density-Based Clustering Methods. Now, we review three well-known clustering algorithms which illustrate these categories:  $k$ -means (a space partitioning method), the agglomerative hierarchical clustering and DBScan

*$k$ -means Algorithm* [26]. It is one of the most commonly used clustering techniques. It is an algorithm to cluster  $n$  objects into  $k$  partitions ( $k < n$ ).  $K$ -means starts by partitioning randomly the input objects into  $k$  initial sets. Then, it calculates the centroid of each set. Following, it constructs a new partition by associating each object with the closest centroid. Finally, the centroids are recalculated for the new clusters. This algorithm is repeated until it converges, i.e. there is no changes in its centroids. Figure 1 shows an example of  $k$ -means over a dataset of 100 gaussian random pairs.

*Agglomerative Hierarchical Clustering* [20]. This method builds a hierarchy tree, called dendrogram, from the individual elements by progressively merging clusters. Note that, at the beginning each element is considered as an independent cluster. The algorithm starts computing a distance matrix among all the elements to be clustered, where the distance in the  $(i, j)$  position corresponds to



**Fig. 1.** Example of clustering using  $k$ -means algorithm with  $k = 5$



**Fig. 2.** Example of dendrogram after applying hierarchical clustering


the distance between the  $i$ th and  $j$ th elements. Then, when clustering progresses, the corresponding rows and columns have to be also merged. This algorithm does not explicitly builds a number of clusters, instead, we must decide the number of clusters and where we split them within the dendrogram. An example, over the same dataset as Fig. 1, is illustrated by Fig. 2, which suggests to cut at height 3 or 4.

*DBScan* [9]. It starts with an arbitrary starting point  $r$  that has not been visited.  $r$ 's close neighbours are retrieved, and if this set contains sufficiently many points, a cluster is started. Otherwise,  $r$  is labeled as noise. Note that this point might later be found in a sufficiently sized environment of a different point and hence be made part of a cluster. If a point is found to be a dense part of a cluster, its neighbours are also part of that cluster. Hence, all points that are found within the neighborhood limit are added to the dense cluster. This process continues until the density-connected cluster is completely found. Then, a new unvisited point is retrieved and processed, leading to the discovery of a further cluster or noise. Figure 3 shows the output of DBScan within a two dense clusters dataset.

### 3.1 Geospatial Clustering

Geospatial clustering is a special kind of clustering, its main goal is to group similar objects based on their distance, connectivity, or relative density in space. General clustering methods can be used for geospatial clustering, however due to its inherent spatial nature, specific clustering algorithms categories have been defined in the last years:

*Grid-Based Clustering Methods.* Such methods divide the clustering space into a finite number of cells and then perform all the clustering operations on the grid. Cells containing more than a certain number of points are considered to be dense. Contiguous dense cells are connected to form clusters. One examples of grid-based clustering methods is CLIQUE [2].



**Fig. 3.** Example of DBScan with 2 clusters

*Constraint-Based Clustering Methods.* These clustering methods add spatial constraints, such as obstacles, to obtain more desirable clusters for real geospatial information systems. Depending on the nature of the constraints and applications, constraint-based clustering methods includes four different types of constraints: constraints on individual objects, obstacle objects as constraints, clustering parameters as constraints, and constraints imposed on each individual cluster [33]. DBRS+ [37] is one example of these clustering methods.

### 3.2 Trajectory Clustering

A number of trajectory clustering methods have been proposed. A common characteristic of first trajectory clustering methods is that they use the shapes of whole trajectories to do the clusters, using for instance hidden Markov models (HMM). Later, more complex clustering algorithms [23] were proposed. In these papers authors divide the space into a grid to allow to compute sub-trajectories and consider more complex patterns than using only the whole trajectory.

Our approach is far from these ideas. Here we would like to adapt classical clustering methods, which has been largely studied to cluster trajectories using well-defined distances for sequences of symbols.

### 3.3 Clustering Tools

Nowadays, it is possible to find a large variety of clustering tools, such as WEKA [16]. However, most of these tools are not able to work with big data or cannot be parameterized using an external distance.

To solve the first problem, MongoDB [28] offers special data structures and indices called Geohash [27]. When you create a geospatial index on legacy coordinate pairs, MongoDB computes geohash values for the coordinate pairs within the specified location range and then indexes the geohash values. Geohash values, recursively divide a two-dimensional map into quadrants until a desired level of specificity is reached. We have used MongoDB and geohash data structures to create our trajectory dataset.

In order to easily execute several clustering algorithms coming from the different families described in this section, we will use the ELKI framework [1]. ELKI separates data mining algorithms and data management tasks to allow for including different components inside the library. Additionally, ELKI allows for the use of external distances as the OSA distance described in Sect. 2.2. For these reasons we have selected ELKi for our experiments.

## 4 Social Sensing Trajectory Data Extraction

In order to capture citizens' mobility traces of citizens, different data sources have been used in the literature such as GPS traces [39], call detail records (CDR) from mobile phones [7], and even geopositioned social media [30].

In this work we focus on geopositioned tweets. Twitter allows developers to obtain all the geopositioned tweets within a certain bounding box that covers a certain city. We have used the data-acquisition architecture depicted in [35]. The dataset generated contains the geopositioned tweets generated in Barcelona during six months from July 2012 to the end of December 2012, resulting in approximately 1.160.000 geopositioned data traces.

To transport the geopositioned tweets into actual trajectories, we borrow a methodology presented elsewhere [12]: a trajectory can be understood as the set described by at least two consecutive tweets with a minimum distance of 100 m amongst them and published in less than 75 min of difference.

Moreover, we also profit from a semantically-enhanced module that allows to detect the origin country of each user leaving a digital mobility trace. Each tweet contains the user-specified origin location, allowing users to specify free text values. Other than handling fake values, such as "From heaven", user-specified locations can be referred at different levels of granularity (GPS coordinate, neighborhood, city, region, or country level) and in different languages (e.g. *London*, *Londra*, *Londres*, etc...). In order to deal with this ambiguity, we use the GeoNames service [13], which provides a unique country identifier per any input, being also successful dealing with fake values.<sup>1</sup>

Therefore, our technological infrastructure allows us to capture user trajectories and classify them per user origin country.

By applying clustering algorithms on the digital traces of our 6 months dataset, we obtain clusters of activity whose cluster identifier provides us with a unique tag to identify users origins and destinations.

<sup>1</sup> A human-supervised test on 100 random user-specified locations obtained from our dataset obtained an 72% accuracy rate.



## 5 High Performance Computing Challenges in Geospatial Clustering

To avoid the problems associated to persisting data to disks, persistent memories are leveraged (Flash memory at the current time, any kind of Storage Class Memories in the near future).

To increase the address space made available to a particular workload, the memory of all nodes involved in the execution of such workload are given access to the memory of all other compute nodes through software APIs.

The major problem from the infrastructure point of view for Geospatial Clustering is dealing with large amounts of data. It is a common case that the performance of clustering algorithms is I/O bound because they are data-intensive and process very large data sets. Therefore, novel ways of managing data are required to deliver scalability and performance. The common approach nowadays is to distribute data across many compute nodes and provide applications with a flat and shared name space to access data, independently if it is stored locally or remotely.

The single namespace can be provided either by a distributed file system (the case of MapReduce Distributed File System [3,14]) or using key/value pairs (the case for most NoSQL databases [22,28] and key/value stores [11]).

Over the last years an intersection of different technologies is gaining momentum performance-wise: RDMA-enabled network technologies and persistent memories. The goal of combining these technologies is to provide low latency and high bandwidth all-to-all in a network topology, and therefore fast access to all data in a distributed dataset.

High speed networks require lightweight protocol stacks and CPU offloading to move data between nodes at high speeds (e.g. using Infiniband verbs instead of a heavyweight TCP/IP stack to achieve 56 Gbps bandwidth in Infiniband FDR networks), what is achieved using RDMA-enabled networks such as Infiniband or iWARP.

At the same time, to achieve high bandwidth to store data, fast memories are used instead of slow rotational disks. Such memory can be accessed in different ways: through conventional disk interfaces, what is the case of Solid State Disks (SSD); through conventional PCIe buses, what is the case of PCIe Flash boards; or directly through the memory buses in the processors, what is the case of Phase Change Memories (PCM) or any generic Storage Class Memory (SCM) technology.

The outcome of this technology trend is middlewares that allow for transparent access to remote or local data at the same speeds and with the same latencies, unifying the memory space of all nodes involved in the execution of a workload, and simplifying the programmability of the applications by providing simple user APIs, such as the Blue Gene Active Storage [10] (BGAS) platform does.

## 6 Conclusions

In this abstract we have described all the components of a possible framework for public transport optimisation in big cities, in our case Barcelona. We have

described several distances and clustering algorithms to illustrate that combining a distance for sequences of symbols with classical clustering algorithms is possible to create a clustering algorithm for this goal. We have also introduced some performance problems and how we can overcome them. Finally, we have explained how to collect the required data to convert this framework in a real tool during the next year.

**Acknowledgments.** This work is partially supported by the Ministry of Science and Technology of Spain under contract TIN2012-34557 and by the BSC-CNS Severo Ochoa program (SEV-2011-00067) and with the support of ACCIÓ, the Catalan Agency to promote applied research and innovation; and by the Spanish Centre for Development of Industrial Technology under the INNPRONTA program, project IPT-20111006, “CIUDAD2020”.

## References

1. Achtert, E., Kriegel, H.-P., Zimek, A.: ELKI: a software system for evaluation of subspace clustering algorithms. In: Ludäscher, B., Mamoulis, N. (eds.) *SSDBM 2008*. LNCS, vol. 5069, pp. 580–585. Springer, Heidelberg (2008)
2. Agrawal, R., Gehrke, J., Gunopulos, D., Raghavan, P.: Automatic subspace clustering of high dimensional data for data mining applications. *SIGMOD Rec.* **27**(2), 94–105 (1998)
3. Apache Software Foundation. Hadoop Distributed File System (HDFS) Architecture.
4. Berghel, H., Roach, D.: An extension of Ukkonen’s enhanced dynamic programming asm algorithm. *ACM Trans. Inf. Syst.* **14**(1), 94–106 (1996)
5. Chávez, E., Navarro, G., Baeza-yates, R., Marroquín, J.L.: Searching in metric spaces. *ACM Comput. Surv.* **33**, 273–321 (1999)
6. Cover, T., Hart, P.: Nearest neighbor pattern classification. *IEEE Trans. Inf. Theor.* **13**(1), 21–27 (1967)
7. de Montjoye, Y.A., Hidalgo, C.A., Verleysen, M., Blondel, V.D.: Unique in the crowd: the privacy bounds of human mobility. *Sci. R.* **3**, 1376 (2013)
8. Dong, G., Pei, J.: *Sequence Data Mining*. Springer, Heidelberg (2007)
9. Ester, M., Kriegel, H., Sander, J., Xu, X.: A density-based algorithm for discovering clusters in large spatial databases with noise. In: *2nd International Conference on Knowledge Discovery and Data Mining*, pp. 226–231 (1996)
10. Fitch, B.G., Rayshubskiy, A., Pitman, M.C., Christopher Ward, T.J., Germain, R.S.: Using the active storage fabrics model to address petascale storage challenges. In: *Proceedings of the 4th Annual Workshop on Petascale Data Storage, PDSW ’09*, pp. 47–54. ACM, New York (2009)
11. Fitzpatrick, B.: Distributed caching with memcached. *Linux J.* **124**, 5 (2004)
12. Gabrielli, L., Rinzivillo, S., Ronzano, F., Villatoro, D.: From tweets to semantic trajectories: mining anomalous urban mobility patterns. In: *Proceedings European Conference on Complex Systems ECCS: Barcelona*. Springer, Spain (2013)
13. GeoNames. GeoNames geographical database. <http://www.geonames.org/> (2010). Accessed July 2013
14. Ghemawat, S., Gobioff, H., Leung, S.-T.: The google file system. *SIGOPS Oper. Syst. Rev.* **37**(5), 29–43 (2003)

15. Gómez-Alonso, C., Valls, A.: A similarity measure for sequences of categorical data based on the ordering of common elements. In: Torra, V., Narukawa, Y. (eds.) MDAI 2008. LNCS (LNAI), vol. 5285, pp. 134–145. Springer, Heidelberg (2008)
16. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The weka data mining software: an update. *ACM SIGKDD Explor. Newslett.* **11**(1), 10–18 (2009)
17. Hamming, R.W.: Error detecting and error correcting codes. *Bell Syst. Tech. J.* **26**(2), 147–160 (1950)
18. Herranz, J., Nin, J., Solé, M.: Optimal symbol alignment distance: a new distance for sequences of symbols. *IEEE Trans. Knowl. Data Eng. (TKDE)* **23**(10), 1541–1554 (2011)
19. Jain, A.K., Murty, M.N., Flynn, P.J.: Data clustering: a review. *ACM Comput. Surv.* **31**(3), 264–323 (1999)
20. Jardine, N., Sibson, R.: The construction of hierarchic and non-hierarchic classifications. *Comput. J.* **11**(2), 177–184 (1968)
21. Jaro, M.A.: Advances in record-linkage methodology as applied to matching the 1985 census of Tampa, Florida. *J. Am. Stat. Assoc.* **84**, 414–420 (1989)
22. Lakshman, A., Malik, P.: Cassandra: a structured storage system on a p2p network. In: Proceedings of the Twenty-First Annual Symposium on Parallelism in Algorithms and Architectures, SPAA '09, p. 47. ACM, New York (2009)
23. Lee, J.G., Han, J., Li, X., Gonzalez, H.: Traclass: trajectory classification using hierarchical regionbased and trajectorybased clustering. In: *ACM Very Large Data Base (VLDB)* (2008)
24. Lee, J.G., Han, J., Whang, K.Y.: Trajectory clustering algorithms. In: International Conference on Management of Data (SIGMOD), pp. 593–604 (2007)
25. Levenshtein, V.I.: Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Phys. Doklady* **10**, 707–710 (1966)
26. Lloyd, S.: Least squares quantization in PCM. *IEEE Trans. Inf. Theor.* **28**, 129–137 (1982)
27. Geohash, <http://docs.mongodb.org/manual/core/geospatial-indexes>
28. MongoDB, <http://www.mongodb.org>
29. Navarro, G.: A guided tour to approximate string matching. *ACM Comput. Surv.* **33**(1), 31–88 (2001)
30. Noulas, A., Scellato, S., Lambiotte, R., Pontil, M., Mascolo, C.: A tale of many cities: universal patterns in human urban mobility. *PloS One* **7**(5), e37027 (2012)
31. Ristad, E., Yianilos, P.: Learning string edit distance. *IEEE Trans. Pattern Recogn. Mach. Intell.* **20**(5), 522–532 (1998)
32. Sellers, P.: The theory and computation of evolutionary distances: pattern recognition. *J. Algorithms* **1**(4), 359–373 (1980)
33. Tung, A.K.H., Han, J., Lakshmanan, L.V.S., Ng, R.T.: Constraint-based clustering in large databases. In: Van den Bussche, J., Vianu, V. (eds.) *ICDT 2001*. LNCS, vol. 1973, pp. 405–419. Springer, Heidelberg (2000)
34. Ukkonen, E.: On approximate string matching. In: Karpinski, M. (ed.) *FCT 1983*. LNCS, vol. 158, pp. 487–495. Springer, Heidelberg (1983)
35. Villatoro, D., Serna, J., Rodríguez, V., Torrent-Moreno, M.: The TweetBeat of the city: microblogging used for discovering behavioural patterns during the MWC2012. In: Nin, J., Villatoro, D. (eds.) *CitiSens 2012*. LNCS, vol. 7685, pp. 43–56. Springer, Heidelberg (2013)
36. Wagner, R.A., Fischer, M.J.: The string-to-string correction problem. *J. ACM* **21**(1), 168–173 (1974)

37. Wang, X., Hamilton, H.J.: Dbrs: a density- based spatial clustering method with random sampling. In: 7th Pacific-Asia Conference on Knowledge Discovery and Data Mining, pp. 563–575 (2003)
38. Wang, X., Wang, J.: Using clustering methods in geospatial information systems. *Geomatica* **64**(3), 347–361 (2010)
39. Zheng, Y., Zhang, L., Xie, X., Ma, W.-Y.: Mining interesting locations and travel sequences from gps trajectories. In: Proceedings of the 18th International Conference on World Wide Web, pp. 791–800. ACM, New York (2009)