

SPATIO-TEMPORAL CLUSTERING OF MOVEMENT DATA: AN APPLICATION TO TRAJECTORIES GENERATED BY HUMAN-COMPUTER INTERACTION

G. McArdle^{a,*}, A. Tahir^b, M. Bertolotto^b

^aNational Centre for Geocomputation, National University of Ireland Maynooth, Maynooth, Co. Kildare, Ireland
gavin.mcardle@nuim.ie

^bSchool of Computer Science and Informatics, University College Dublin, Dublin 4, Ireland
(ali.tahir, michela.bertolotto)@ucd.ie

Commission II/6

KEY WORDS: data mining, web based, visualisation, spatial, tasks

ABSTRACT:

Advances in ubiquitous positioning technologies and their increasing availability in mobile devices has generated large volumes of movement data. Analysing these datasets is challenging. While data mining techniques can be applied to this data, knowledge of the underlying spatial region can assist interpreting the data. We have developed a geovisual analysis tool for studying movement data. In addition to interactive visualisations, the tool has features for analysing movement trajectories, in terms of their spatial and temporal similarity. The focus in this paper is on mouse trajectories of users interacting with web maps. The results obtained from a user trial can be used as a starting point to determine which parts of a mouse trajectory can assist personalisation of spatial web maps.

1 INTRODUCTION

One of the main processes in our physical environment is movement. The increasing mobility of individuals coupled with advances in ubiquitous positioning technologies has created a growing volume of data describing movement in many domains. The predominant challenge concerns the quantity of data contained in massive movement datasets. The research presented in this paper is concerned with the analysis of movement data generated through Human-Computer Interaction (HCI), in order to identify and interpret patterns, which can be used to predict future behaviour. Traditional data mining techniques can be applied to movement data, however as the data corresponds to processes taking place in physical environment, the analysts' knowledge regarding the underlying spatial region can greatly assist in interpreting the data. Geovisual analysis is a powerful technique which facilitates this. Stemming from visual analytics, it aims to support the human reasoning and cognitive process through interactive visual interfaces (Thomas and Cook, 2005). In many cases, such geovisual analysis of movement data, can detect outliers or unusual behaviour which data mining approaches miss. For example, trajectory pattern mining is a specialised emerging field within data mining which utilises complex spatio-temporal databases (Jeung et al., 2011), to discover unusual patterns and outliers (such as unusual mouse movements).

Movement data is often described by trajectories which represent the locations of objects over a period of time. Such trajectories are not limited to physical movement; in scenarios such as HCI, both mouse and eye movements on a computer screen generate trajectories. Here, we are interested in the special form of movement which is generated through the mouse interactions of human-users with web maps. The goal is to identify patterns which can be used for personalising such web maps (MacAoidh et al., 2008). A mouse trajectory on a web map describes the mouse position on the screen after it has been translated to valid geographical coordinates. Additional information such as timestamps, speed, acceleration and map scales are also recorded. Mouse movements on web maps differ from movement in the physical world which makes analysis more complex. For example, mouse trajectories are a true form of free movement (Jensen et al., 2009)

and are not constrained by a road network or certain rules of a physical environment. Unlike navigation in the real world, web map users can alter the scale at which they are interacting with spatial content which is a significant indicator of interest and important for the personalisation process. Geographical Positioning Systems (GPS) are generally used to track the physical position of objects in Location based Social Networks (LBSN). The nature of this technology makes it prone to erroneous and missing data. In contrast, every mouse movement is recorded precisely by the computer, but like movement in a physical space, the challenge is to determine which movements are important. We have developed a geovisual analysis tool for studying movement data. The tool incorporates visualisation techniques for understanding user behaviour with special emphasis on techniques for interpreting mouse trajectories. Spatial and behavioural clustering are used to determine similarities between trajectories. Ordering Points To Identify the Clustering Structure (OPTICS), a density-based clustering algorithm is applied to measure spatial similarity while the velocity and acceleration of mouse movements are considered as a technique for determining similar behaviour (Ankerst et al., 1999). Progressive clustering (Rinzivillo et al., 2008) is also employed to identify mouse trajectories which occur in the same spatial region and behave similarly. We have used this tool to analyse the mouse trajectories captured during a user trial. The data is used as a starting point to determine what important elements of mouse interactions should be considered for personalisation of spatial data, web maps and associated tools. The remainder of this paper is organised as follows: Section 2 discusses the relevant literature. A description of our geovisual analysis tool is provided in Section 3. Section 4 presents details on a user trial and the results obtained. Finally, a discussion and future directions are presented in Section 5.

2 RELATED WORK

Moving objects are generally described in terms of a trajectory, i.e. a sequence of positions in a two-dimensional geographic environment with associated time stamps (Laube et al., 2005). Often there are also associated attribute data, which can be either static (the attribute has the same value for the object regardless

of its position, e.g. object type) or dynamic (the attribute changes over time, e.g. attributes which describe the physical properties of movement, such as velocity and acceleration). Various geovisual tools have been developed to display and interpret such trajectories and associated data. For example, CommonGIS (Andrienko et al., 2003), GIIViZ (MacAoidh et al., 2008), The Animal Ecology Explorer (Spretke et al., 2011), TrajVis (McArdle and Demšar, 2011) and COMPASS (Doyle et al., 2010). Below, we describe the Space Time Cube (STC), a classic visualisation technique for movement data and present techniques for determining the spatial and temporal similarity of trajectories.

While 2D map-based geographic representations are appropriate for analysing the spatial aspects of movement, without animation or the use of individual visualisations for each time period, the temporal component of movement is lost. The STC is a geovisual analysis approach in which the temporal and spatial components of a trajectory are visualised simultaneously (Hägerstrand, 1970). In a STC, the x and y planes represent the spatial context while the z plane represents the temporal component. Trajectories are represented by a 3D polyline in a space-time cube. Using this approach, it is clear to see the routes and relative speeds of objects. Sophisticated computer technology has advanced the STC. For example, Eccles et al. (2007) have developed a highly interactive STC for displaying human movement data, the visualisation is augmented by showing connections, such as phone calls, between those being tracked. Vrotsou et al. (2010) compared the STC approach with a traditional 2D approach. The results indicate that 3D views improved task performance in some situations; however for complex data the 2D visualisation produced more accurate results. As with most visual analysis tools, issues related to complexity and visualising large volumes of data in a STC are well documented (Andrienko et al., 2007). In such cases additional analysis techniques such as clustering and aggregation are essential. Clustering data involves identifying similarities between data points and using this as a basis to group them (Jain et al., 1999; Xu et al., 2005). The approach can be used to summarise patterns in the data. In order for clustering of movement data to be effective, it requires an appropriate similarity measure to compare trajectories. Spatial and geometric similarity, temporal similarity and attribute similarity are the principle techniques and below some common approaches are discussed.

Methods for identifying groups of similar trajectories are commonly based on their geometric similarity in two-dimensional geographic space. Several trajectory distance measures have been classified into global and local distance measures (Zheng and Zhou, 2011). A global measure computes the distances between two trajectories with respect to all points in a trajectory while a local measure calculates the similarity between sub-trajectories. Global measure distances include euclidean distance, alignment-based distance, Dynamic Time Warping (DTW), Edit Distance on Real Sequence (EDR), Longest Common SubSequence Measure (LCSS) and Edit Distance with Real Penalty (ERP). Local measures include Minimum Bounding Rectangles (MBR)-based distance, trajectory Hausdorff and trajectory segment. Morris and Trivedi (2009) evaluated widely used distance measures based on fixed length measures (Hu Euclidean and Principal Component Analysis) as well as time-normalised measures (modified Hausdorff, Piciarelli and Foresti, LCSS and DTW). Common destination and route similarity (Rinzivillo et al., 2008) are other approaches for calculating trajectory similarity which we have adopted to compute the distance between mouse trajectories while carrying out spatial tasks. An appropriate distance or similarity measure is required in all clustering techniques. Han (2005) classified clustering techniques into partitioning, hierarchical, density-based, grid-based, model-based, constraint-based and clustering

high-dimensional data. Commonly used clustering approaches are discussed in several studies (Rinzivillo et al., 2008; Panagiotakis et al., 2011; Lee et al., 2007, 2008). For spatial clustering of mouse trajectories, we found density-based clustering, in particular the OPTICS (Ankerst et al., 1999) algorithm from the DBSCAN family, to be the most suitable. This is due to the fact that density-based methods are efficient for finding noise and detecting outliers. Furthermore they are capable of detecting clusters of an arbitrary shape which is a desirable property when analysing mouse trajectories.

Certain techniques discussed above incorporate the temporal aspects of trajectories. For example, DTW (Sakurai et al., 2005) stretches the time axis in order to identify similarities in trajectory shape. This allows a comparison of trajectories which span different time frames. Spatial transformations can also be applied to realign trajectories for better comparison. The LCSS approach does not consider the entire trajectory as a whole but finds similarities between substrings (Vlachos et al., 2002). EDR, which measures the number of operations (insert, delete or replace) required to transform one trajectory to another, extends this approach by assigning penalties to the gaps between two matched sub-trajectories according to the lengths of gaps (Chen et al., 2005). Another study bases its similarity measures on additional attributes of movement, such as speed, acceleration, duration and direction (Dodge et al., 2009). These techniques are generally components of geovisual analysis tools. They can be used as a form of data reduction to support cognitive processes by reducing the number of trajectories displayed simultaneously. Once a suitable, similarity metric has been determined, the trajectories can be compared and ultimately clustered. This allows similar trajectories to be grouped together. Outliers and salient trends can be identified and visually analysed through aggregation.

In this paper we build on these geovisual analysis tools and techniques to analyse mouse trajectories. The goal is to use such approaches to identify usage patterns of an interactive web-based map. The new geovisual analysis tool which we have developed combines 2D map overlays, statistical analysis and STC visualisations to assist analysts with interpreting the mouse movements of users and to identify behaviour and intentions.

3 SYSTEM DESCRIPTION

Our geovisual analysis tool provides in-depth analysis of mouse movements when studying user behaviour and mining specific usage patterns. These patterns reveal user intentions and interest which offer an insight into the requirements for map personalisation. The tool enables analysts to visualise mouse movements, hesitations, clicks and trajectories. All these features reflect trends, usage patterns and behaviour and are important indicators in map personalisation. In addition to these features several visual analysis tasks can be performed by the tool. A spatial heat map can be generated based on user actions on a map that shows the regions of user interests using color intensity. Similarly, mouse speed can be visualised in the form of a trajectory which can highlight user activity in a particular portion of a map. This information can be used to classify users (for example, slow, moderate, fast, novice, experienced). The map scale is visualised in the form of a bounding box. Since a user performs multiple map operations (zooming and panning), it is vital to visualise an individual scale where an activity is taking place. Moreover, the map scale becomes significant towards the completion of a spatial task. We use a term 'prime view', which is a map scale view recorded when a user accomplishes a spatial task. The prime view can also be visualised in our tool and is often the most important

part of a particular task as it reveals user intentions. The STC is another visualisation technique that has been incorporated in our Web-based tool. As described above, the technique enables analysts to visualise the temporal ordering and sequence of mouse movements. In our tool, trajectories are draped over a virtual globe which provides access to the underlying spatial data via the 3D representation of the earth. While our system framework is presented in (Tahir et al., 2012), Figure 1 highlights the principle visual analysis functionality of the tool. The tool has been developed using open source technologies and a client-server Web architecture (Tahir et al., 2011).

The functionality discussed above represents the visualisation of a single user session. However, for recommendation and personalisation purposes, the history of multiple users needs to be considered. Multiple user sessions produce a large quantity of trajectories. This results in cluttering and occlusion of the visualisation components which can be resolved by applying appropriate spatio-temporal clustering of trajectories. For spatial clustering, our geovisual analysis tool supports OPTICS (Ankerst et al., 1999), a density-based clustering algorithm to find clusters of arbitrary shaped trajectories. The OPTICS algorithm can be described in three steps which search for a core and a reachability distance. First of all, a random object x is chosen from the full dataset. At the next level, at each iteration i , the next object y is selected from the dataset with the smallest reachability distance with respect to the already visited core objects. Finally, the process is repeated until all objects in the dataset have been considered. The output of the OPTICS algorithm is a 2-dimensional plot that shows the number of trajectories on the x-axis while the y-axis plots a suitable reachability distance. From this plot, a clustering structure can be obtained by choosing an appropriate threshold value of reachability distance. The valleys which appear on such a graph signify the gaps between clusters, see Figure 4 for an example.

In order to support behavioural analysis based on the clustering of temporal components of a trajectory, we have developed an algorithm, which considers the speed and acceleration at each location in a trajectory to describe behaviour. Assuming that two trajectories have the same number of points (This can be achieved via sub sampling of the shorter trajectory and interpolation of missing locations), the trajectories have a similar shape if the mathematical slope of both functions is similar at all locations and the rate of change of mathematical slope of both functions is similar at all locations. These values correspond to the speed and acceleration of the trajectory. Trajectories can therefore be grouped by performing clustering on the dataset of numerically calculated first and second derivatives of each trajectory. Such slope based similarity computation is a well known approach for clustering of time series (Altıparmak et al., 2006), however, only the first derivative of the time series is considered, while we add the second derivative (the rate of change of mathematical slope), for a more detailed approach and a more complex description of the movement being studied.

Once each trajectory is described in terms of its behaviour, clustering techniques can be applied. We opted to use Spectral Clustering (Song et al., 2008), specifically the approach developed by Chen et al. (2011), as it is faster than other methods such as simple k-means. Furthermore, it can detect clusters that k-means would not recognise such as non-convex clusters. Spectral Clustering has been effective for trajectory analysis (Atev et al., 2010) however, the similarity measure was based on trajectory location, rather than behaviour. In our case, it has proved successful at detecting temporal patterns and grouping trajectories of a similar duration with a similar number of stops.

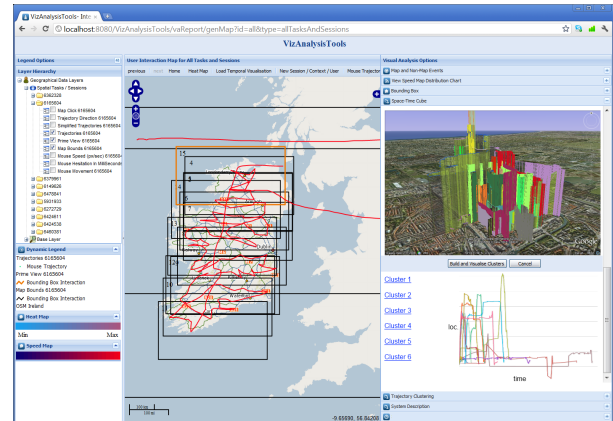


Figure 1: Geovis tool showing different functionality

4 EXPERIMENT AND RESULTS

In order to demonstrate the power of our geovisual analysis tool and gather some useful insight into user behaviour, a user trial was conducted. The trial involved participants interacting with a web map using a mouse. 12 participants (10 males and 2 females) volunteered and took part in these trials which took place in an unsupervised environment.




Figure 2: Trajectories in 2D




Figure 3: Trajectories in 3D

The majority of the participants (with the exception of 2) had previous experience with interactive maps. Due to the authors familiarity with Ireland, this was selected as the study area as it facilitated the design of 10 meaningful tasks. Each spatial task corresponds to one user session. A web interface was designed with a mapping component and tasks were clearly presented at the top of the web page with the map below. Users were required to complete each task by answering the question at the top of the page before they could proceed to the next task. The mapping component consisted of basic map operations (zooming and panning) however no search facility was provided as most tasks were based on scanning operations. The 10 spatial tasks are listed below.

Task1: How many motor ways are there in Ireland? 2. Find the total number of exits on M50 motorway in Dublin. 3. Name the southern cities of Ireland starting with letter 'W'. 4. Write the names of areas close to the starting and ending of M50 motorway in Dublin. 5. Which is the nearest hospital to Heuston train station in Dublin? 6. Name the possible tourist attraction closest to Saint Stephen's Green Park in Dublin city center. 7. What is the name of the largest park in Dublin? 8. Write the total number of bridges on Liffey River in Dublin starting from Grattan bridge towards the eastern coast. 9. How many vehicle entrances are there on the University College Dublin (UCD) campus? 10. Locate Connolly train station in Dublin. What is the name of the nearest canal to this station.




Figure 4: An OPTICS plot showing clustering structure



Figure 5: The trajectories representing task 9 shown in the Space Time Cube. The trajectories all converge on one location which represents the focus of this particular task and therefore the prime view

Based on these 10 spatial tasks and 12 users, a total of 120 trajectories were collected. However, only 117 trajectories were used for analysis purposes as some participants answered the questions without performing spatial analysis. As an initial analysis, the 117 trajectories were visualised with our geovisual analysis tool, in 2D and in the STC (Figures 2 and 3 respectively). Trajectories related to a specific task can be analysed separately. For example, task 9 is visualised in the STC in Figure 5 which shows the trajectories of individual users performing the same task. Each trajectory has a different colour to aid identification. By examining the height of the trajectories, it is notable that the highest points, which correspond to the location where the last interaction took place, are located in the same geographic region (the prime view). In this figure it is evident that all trajectories for this task converge on the same location, despite some initial seemingly random movement.

In terms of interaction, each trajectory can be queried by clicking on it. This provides additional information such as the task and user which the trajectory corresponds to, similarly, trajectories can be removed and added depending on the analyst's focus. While individual trajectories can be removed and added, cluttering and occlusion is a problem. Therefore, the clustering techniques mentioned in Section 3 were applied in order to extract usage patterns from the large set of trajectories.

The OPTICS algorithm was used to find spatial similarity between mouse trajectories. For spatial clustering only the end points of trajectories were considered within trajectories in a given prime view. All those trajectories whose destinations were within a specified distance threshold were grouped together to form a cluster. There are two inputs which are required by the OPTICS algorithm: a distance threshold which was chosen as 100 kilometres and minimum number of neighbours as 5. The algorithm was run several times in order to obtain the correct combination of in-




Figure 6: Spatial task validation relative to cluster cardinalities

put values. Based on the above parameters, an OPTICS plot was obtained as shown in Figure 4. This graph indicates 8 clusters including a noise cluster (cluster 8 in Figure 6). The noise appeared as the spatial tasks were spread across a large geographical area. This classification successfully grouped trajectories into the correct task cluster for nine of the ten spatial tasks as shown in Figure 6, however some overlaps were observed.

In order to determine the similarity among the behaviour of users over multiple tasks and to identify groups of users whose mouse behaviour was similar, the temporal clustering described above was applied to 115 trajectories (2 were removed as visual analysis showed them to be of a very short duration). The speed and acceleration of the mouse at each point on each trajectory was calculated. The results were then clustered using spectral clustering. As we were interested in users, 12 (the number of users) clusters were generated. The results are presented in Figure 7. They show that while some user trajectories (user6 and user10) were predominantly placed in a single cluster and signify a consistent behaviour over various tasks, for the majority this is not the case. This can be attributed to the fact that the tasks were very varied and called for different types of behaviour and different forms of interaction with the web map.

When the trajectories which represent consistent users are visually analysed, it was seen that their behaviour was somewhat similar, with a consistent speed over different parts of the task. Faster speeds were seen at the start of a task while the speed of the mouse movements were slower near the end of the task. Similarly, the duration of the tasks was comparable. These findings are in contrast to the bulk of the users whose speed and duration differed greatly depending on the task they were completing. In order to identify similar behaviour among different users when completing the same task, progressive clustering was applied. Progressive clustering involves applying clustering to a complete dataset and then applying further clustering to each of the resulting clusters. In our case, spatial clustering was initially applied. This essentially returns the clusters corresponding to the original tasks as evident from Figure 6. The trajectories corresponding to the predominant task of each cluster were then extracted. Behavioural clustering, based on the speed and acceleration of the

	1	2	3	4	5	6	7	8	9	10	11	12
user1	2					2		2		1		3
user2	1					1		3		2		3
user3	3		1		2	2			1		1	
user4	2	1		1	1	1			4			
user5	3				1	6						
user6	1						9					
user7	1				2	2		2		2		2
user8	1			4	2						3	
user9				3	4						2	
user10	1											
user11	3				3				3	1		
user12	1				1			1	1	2		3

Figure 7: Table showing users frequency based on spatial tasks appearing in clusters. The red color shows maximum occurrences while green color shows minimum occurrences

	1	2	3	4	5	6	7	9	10
user 1		2	1	1	1	1	1	1	1
user2	1	2	2	2	2	2	2	2	2
user3	1		1	1	1	1	1	1	1
user4			2			2	2	2	2
user5				1				1	1
user6		1	1	1	1	1	1	1	1
user7	2							2	1
user8	1	1		1		1		1	1
user9	1		1						
user10	1							1	1
user11	2					1	1	2	1
user12			1	2	2	1	1	1	1

Figure 8: Table showing the cluster assignment of the user trajectories corresponding to tasks: 1,2,3,4,5,6,7,9 and 10. Task 8 has been removed as the quantity of data was insufficient. Outliers are highlighted in red

mouse movements, was then applied to each of these clusters to determine similarity within the tasks.

The rule of thumb, $k \approx \sqrt{\frac{n}{2}}$, which determines an appropriate but approximate value for the number of clusters present in a dataset of size n , was used. In this case, given that the clusters contained at most 10 trajectories, 2 clusters were requested. The results presented in Figure 8 show that the majority of users behaved similarly to each other in each task (placed in same cluster). Conversely, the same users consistently appeared as outliers for each task and are highlighted in red in the figure. For example, user2 and user4 are consistently outliers as they appear in a cluster alone. When visually analysed using the STC component, it was seen that these users tend to move the mouse slower and for a shorter duration than the other users for each task. These were novice users in terms of web map experience.

5 DISCUSSION AND FUTURE WORK

This paper has described a new visualisation environment for analysing movement data. The tool includes interactive visualisation and analysis tools. While the environment is suitable for any movement data, we concentrate on analysing mouse movements from users interacting with a web-based map. To study this in more detail, a series of web-based map tasks were devised and carried out by 12 users. All interaction data was recorded and used to produce trajectories which were visualised and analysed in the environment. The environment enables the complete data to be visualised simultaneously in a flat 2D display or in a STC to analyse the temporal aspects. Similarly, a subset of the trajectories can be visualised in this way. The geographic extent and the temporal complexity and the data can be extracted from such visualisations.

In order to provide additional analysis of the trajectories, clustering techniques also form part of the tool. Spatial clustering identifies trajectories which have a similar geographic pattern. This proved beneficial for identifying the tasks which formed part of the user study. Clustering based on the temporal aspects of the trajectories revealed that velocity and acceleration vary from task

to task and as a result, an individual's behaviour is dependent on the type of task. Both clustering techniques were combined to determine if users consistently behave the same in different tasks. Firstly, spatial clustering extracted spatially similar trajectories into clusters. Temporal or behavioural clustering was then applied to each of these clusters. The results revealed that certain users are consistently outliers and perform the task using a different type of mouse behaviour (in terms of speed, acceleration and duration) than other users.

The results highlight the dangers of stereotyping users and reinforce the need to provide personalisation, not just at the user level but also at the task level. The experiment was carried out with a relatively low number of trajectories (120) and this must be considered when examining the results of the clustering. A larger user base can provide more sophisticated analysis in terms of creating similarity and clustering metrics and would give greater weight and importance to the results obtained. At present, the tool provides no methods for cleaning the data and this obviously affects the results of the experiment. For example, the visual analysis which was carried out clearly showed that some mouse movements were associated with accessing the map tools and not necessarily relevant to the underlying spatial content. The challenge is to detect such moves and remove them from the similarity analysis where tool usage is not of concern.

The application of our results for supporting map personalisation is the next step on our research agenda. Classifying users according to their intentions and behaviours can help generate more usable maps that simplify user tasks. Furthermore, identifying user preferences allows recommendation techniques to be developed. For example, these can be used to recommend points of interest or accommodation venues to tourists, specific commercial outlets to shoppers and restaurants to people living in or visiting an area.

Finally, in future trials, an individual will be asked to perform the same task type repeatedly in order to see specific patterns for a single user. Such patterns can help in recommending and personalising map content for individuals but can also be expanded and used with collaborative filtering to recommend content to groups of users. One major challenge is to incorporate map scale with the analysis. For example, a map scale could be omitted or included in a user session to facilitate task completion based on the fast or slow user respectively. The scale at which mouse events occur is an extremely important indicator of user interest in a spatial region. Presently, map scale forms part of the visual analysis. Techniques for incorporating it with the similarity analysis are currently being developed so that map scale can be utilised as part of a larger study in the future.

ACKNOWLEDGEMENTS

Research presented in this paper was funded by a Strategic Research Cluster grant (07/SRC/I1168) by Science Foundation Ireland under the National Development Plan. The authors gratefully acknowledge this support.

References

- Altıparmak, F., Ferhatosmanoglu, H., Erdal, S. and Trost, D., 2006. Information mining over heterogeneous and high-dimensional time-series data in clinical trials databases. *IEEE Transactions on Information Technology in Biomedicine* 10(2), pp. 254–263.
- Andrienko, G., Andrienko, N. and Voss, H., 2003. *Gis for everyone: the commongis project and beyond*. Maps and the Internet. Elsevier Science pp. 131–146.

- Andrienko, G., Andrienko, N. and Wrobel, S., 2007. Visual analytics tools for analysis of movement data. *ACM SIGKDD Explorations Newsletter* 9(2), pp. 38–46.
- Ankerst, M., Breunig, M., Kriegel, H. and Sander, J., 1999. Optics: Ordering points to identify the clustering structure. *ACM SIGMOD Record* 28(2), pp. 49–60.
- Atev, S., Miller, G. and Papanikolopoulos, N., 2010. Clustering of vehicle trajectories. *Intelligent Transportation Systems, IEEE Transactions on* 11(3), pp. 647–657.
- Chen, L., Özsu, M. and Oria, V., 2005. Robust and fast similarity search for moving object trajectories. In: *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, ACM, pp. 491–502.
- Chen, W., Song, Y., Bai, H., Lin, C. and Chang, E., 2011. Parallel spectral clustering in distributed systems. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33(3), pp. 568–586.
- Dodge, S., Weibel, R. and Forootan, E., 2009. Revealing the physics of movement: Comparing the similarity of movement characteristics of different types of moving objects. *Computers, Environment and Urban Systems* 33(6), pp. 419–434.
- Doyle, J., Bertolotto, M. and Wilson, D., 2010. Evaluating the benefits of multimodal interface design for compass- a mobile GIS. *Geoinformatica* 14(2), pp. 135–162.
- Eccles, R., Kapler, T., Harper, R. and Wright, W., 2007. Stories in geotime. In: *Visual Analytics Science and Technology, 2007. VAST 2007. IEEE Symposium on*, IEEE, pp. 19–26.
- Hägerstrand, T., 1970. What about people in regional science? *Papers in Regional Science* 24(1), pp. 6–21.
- Han, J., 2005. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- Jain, A., Murty, M. and Flynn, P., 1999. Data clustering: a review. *ACM Computing Surveys (CSUR)* 31(3), pp. 264–323.
- Jensen, C., Lu, H. and Yang, B., 2009. Indexing the trajectories of moving objects in symbolic indoor space. *Advances in Spatial and Temporal Databases* pp. 208–227.
- Jeung, H., Yiu, M. L. and S.Jensen, C., 2011. Trajectory pattern mining. In: Y. Zheng and X. Zhou (eds), *Computing with Spatial Trajectories*, Springer-Verlag New York Inc, pp. 144–145.
- Laube, P., Imfeld, S. and Weibel, R., 2005. Discovering relative motion patterns in groups of moving point objects. *International Journal of Geographical Information Science* 19(6), pp. 639–668.
- Lee, J.-G., Han, J. and Li, X., 2008. Trajectory outlier detection: A partition-and-detect framework. In: *Proceedings of the International Conference on Data Engineering (ICDE)*, pp. 140–149.
- Lee, J., Han, J. and Whang, K., 2007. Trajectory clustering: a partition-and-group framework. In: *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data*, ACM, pp. 593–604.
- MacAoidh, E., Bertolotto, M. and Wilson, D., 2008. Understanding geospatial interests by visualizing map interaction behavior. *Information Visualization* 7(3-4), pp. 275.
- McArdle, G. and Demšar, U., 2011. Investigating similarity of trajectories through physical decomposition of movement in a space-time cube. In: *14th AGILE International Conference on Geographic Information Science, Advancing Geoinformation Science for a Changing World*.
- Morris, B. and Trivedi, M., 2009. Learning trajectory patterns by clustering: Experimental studies and comparative evaluation. In: *IEEE Conference on Computer Vision and Pattern Recognition, 2009. CVPR 2009.*, pp. 312–319.
- Panagiotakis, C., Pelekis, N., Kopanakis, I., Ramasso, E., Theodoridis, Y. et al., 2011. Segmentation and sampling of moving object trajectories based on representativeness. *IEEE Transactions on Knowledge and Data Engineering* 1, pp. 17.
- Rinzivillo, S., Pedreschi, D., Nanni, M., Giannotti, F., Andrienko, N. and Andrienko, G., 2008. Visually driven analysis of movement data by progressive clustering. *Information Visualization* 7(3), pp. 225–239.
- Sakurai, Y., Yoshikawa, M. and Faloutsos, C., 2005. Ftw: fast similarity search under the time warping distance. In: *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, ACM, pp. 326–337.
- Song, Y., Chen, W., Bai, H., Lin, C. and Chang, E., 2008. Parallel spectral clustering. In: *Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases*, Springer, pp. 374–389.
- Spretke, D., Janetzko, H. and Mansmann, F., 2011. Exploration through enrichment: A visual analytics approach for animal movement. In: *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, ACM, pp. 421–424.
- Tahir, A., McArdle, G. and Bertolotto, M., 2011. A web-based visualisation tool for analysing mouse movements to support map personalisation. In: *Proceedings of the First International Workshop on Spatial Information Modeling, Management and Mining (SIM3) (in conjunction with DASFAA2011)*, LNCS, Springer, pp. 132–143.
- Tahir, A., McArdle, G. and Bertolotto, M., 2012. Identifying specific spatial tasks through clustering and geovisual analysis. In: *Proceedings of Geoinformatics 2012, IEEE GRSS*. in press.
- Thomas, J. and Cook, K., 2005. Illuminating the path: The research and development agenda for visual analytics. Vol. 54, IEEE.
- Vlachos, M., Kollios, G. and Gunopoulos, D., 2002. Discovering similar multidimensional trajectories. In: *Proceedings of 18th International Conference on Data Engineering*, IEEE, pp. 673–684.
- Vrotsou, K., Forsell, C. and Cooper, M., 2010. 2d and 3d representations for feature recognition in time geographical diary data. *Information Visualization* 9(4), pp. 263.
- Xu, R., Wunsch, D. et al., 2005. Survey of clustering algorithms. *Neural Networks, IEEE Transactions on* 16(3), pp. 645–678.
- Zheng, Y. and Zhou, X., 2011. *Computing with Spatial Trajectories*. Springer-Verlag New York Inc.