Photogrammetric Week '15 Dieter Fritsch (Ed.) Wichmann/VDE Verlag, Belin & Offenbach, 2015

Interpretation of Moving Point Trajectories

Monika Sester, Udo Feuerhake, Colin Kuntsch, Stefania Zourlidou, Hannover

ABSTRACT

Moving point data are a new type of data sources which is growingly available with the advent of new mass GPS sensors and other positioning devices. The trajectories capture the movement and thus reflect the behaviour of the tracked object, however they also reflect the underlying movement constraints. Revealing this information requires automatic analysis procedures. In the paper a short overview of the characteristics of trajectory data are given, followed by examples for sophisticated data interpretation applications, which demonstrate the potential of this new data source.

1. INTRODUCTION

Moving point data sets are a new type of spatial data source. The amounts of these are rapidly growing with the increasing availability of (cheap) sensors with positioning capabilities. Such trajectory data are collected by mobile phones, registering the radio cell, or smartphones, cars, quadcopters equipped with GPS sensors, just to name a few. A trajectory in its pure form just consists of a sequence of points, possibly with time stamps (*x,y,t*). Points may be enriched by additional information, e.g. velocity, temperature, or even a camera picture acquired from the quadcopter or a car sensor. Trajectories are mostly analyzed with respect to the underlying movement or movement behavior. Thus they are often applied in human or animal behavior studies. Another important application is the extraction of underlying movement constraints – the most important application being the extraction of a road network.

The fundamental challenges lie in all aspects of geospatial data processing, namely the storage and efficient access of the data, interpretation and visualization. In the paper, there will be a concentration on interpretation in order to extract relevant, possibly implicit information. Relevant applications will be presented, ranging from the extraction of traffic network via movement behavior of traffic participants to the detection of group movement patterns.

The paper is structured as follows: after a description of the characteristics of moving point data, a brief review of the state of the art is given. Then examples for different trajectory analysis approaches are described.

2. PROPERTIES OF MOVING POINT DATA

A trajectory consists of a (temporal) sequence of points (*x,y,t,a*), and can possibly contain additional attributes *a* measured at each point. Trajectories can be generated by moving objects (animals, cars, football players) but also by moving phenomena, e.g. measurement points on a hill slide. The points can be captured at regular intervals or irregularly; they can be acquired by the moving object itself (if it is equipped with a positioning device) or by an external observing sensor (e.g. a camera at a public space, at checkpoints using RFID). The quality of the positioning can also vary considerably, depending on the positioning method and the environment.

The trajectories reflect the movement behavior of the tracked object or phenomenon. This behavior can be constrained by the underlying environment, e.g. a road network, the infrastructure in a public building (e.g. booths and banks in a railway station), or by given rules (e.g. restricted access, restricted movements), see Figure 1a),c),d) for examples.

The trajectories are also “constrained” by the behavior of the object creating it. This aspect is especially important in the context of animal ecology, where such observations are often the only means to get new insight into the animals’ behavior (see Figure 1b), which shows trajectories of sea gulls.

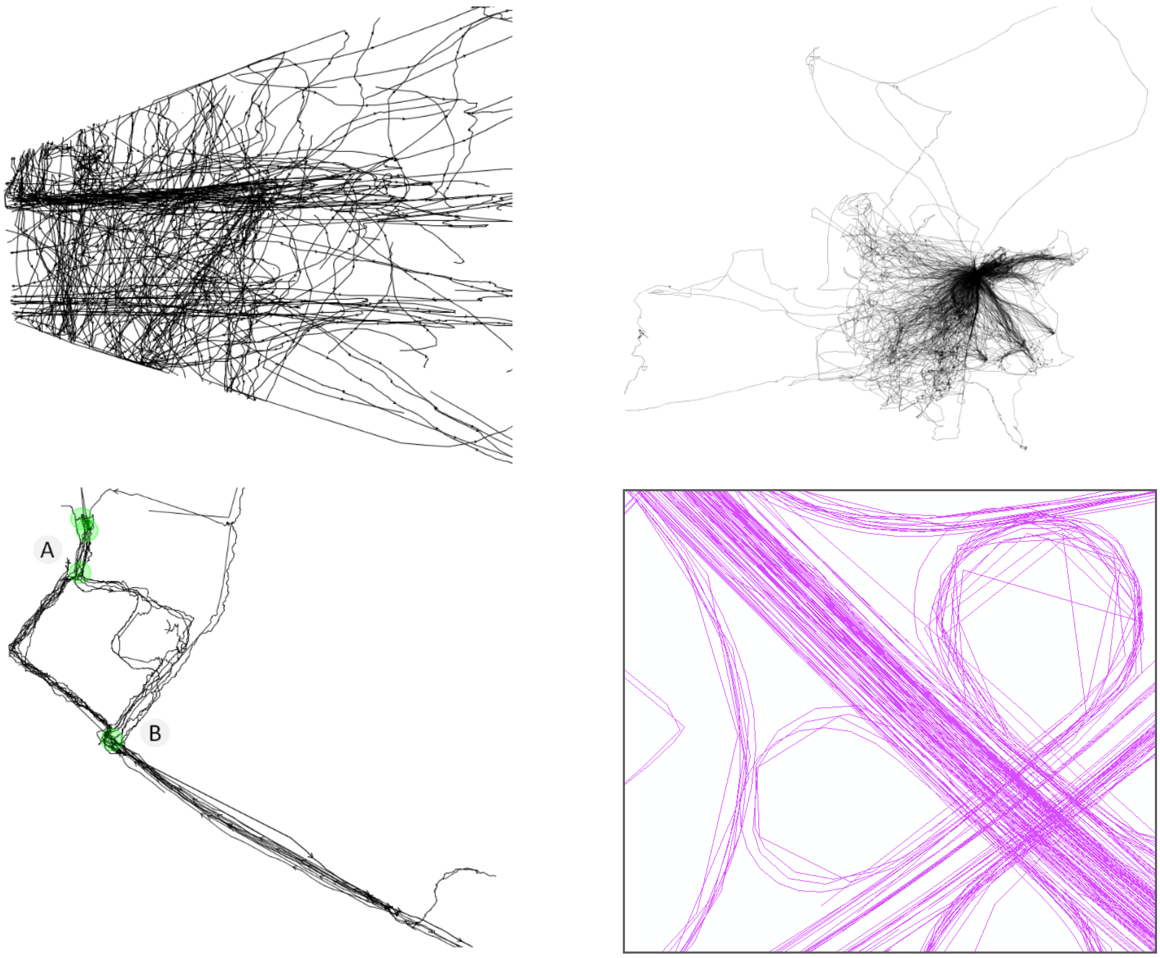

Figure 1: Examples for trajectories: a) (top left): indoor: people moving in an open space with some obstacles; b) (top right) birds movement; c) (bottom left) pedestrians; d) (bottom right): cars on a motorway.

Thus trajectories can be interpreted with different focus, e.g.

- determination of underlying infrastructure (e.g. road, path)
- determination of rules constraining the movement (e.g. traffic rules)
- determination of movement behavior of individuals or groups
- determination of unusual behavior of individuals or groups
- identification of underlying phenomenon (e.g. part of hill which is sliding)

On the algorithmic side, the challenge of processing large quantities of data can be tackled in a centralized or decentralized way: decentralized algorithms are favorable, if the individual datum is not of interest and can be aggregated, if there are limitations in communication capabilities of the sensors, or if privacy is an issue (Duckham, 2012).

3. RELATED WORK

There is a large amount of work of researchers dealing with trajectory interpretation. Dodge et al. (2008) provide a taxonomy of movement patterns, distinguishing on the highest level generic and behavioral patterns. Movement ecology is mainly interested in processes that cause and influence movement (e.g. Soleymani et al., 2015). A review on interdisciplinary approaches on movement analysis is given in Demšar et al. (2015).

The increasing ubiquity of low-cost localization sensors, mainly GPS, has stimulated research towards extracting map information. The approaches have to deal with problems introduced by heterogeneous coverage of the real road network by trajectory data as well as with outliers and systematic GPS errors. With the general goal of generating an underlying map from raw trajectory data, a number of approaches have been proposed which can be categorized in terms of their algorithmic foundation: one type of approaches tries to reconstruct road geometries from the raw trajectory samples by spatial aggregation, e.g. clustering of GPS samples, sometimes using additional information such as derived heading (Edelkamp and Schrödl 2003). Other researchers exploit the clustering properties of kernel density estimation (KDE) in order to reconstruct the full road network (Davies et al., 2006, Biagioni and Eriksson 2012). The trajectories are rasterized in 2D space and algorithms from image processing are applied in order to extract 2D shapes from the KDE histogram, from which centre lines for roads representing the road network geometry are calculated. Another group of clustering methods detects intersections first, leading to an implicit clustering of incident sub-trajectories between junction locations (Fathi and Krumm 2010, Karagiorgou and Pfoser 2012).

A different algorithmic approach is to integrate several tracks using map matching, i.e. finding common segments among multiple trips, mainly based on spatial proximity and similarity of additional features such as heading. Tracks are sequentially integrated into an initially empty map, leading to modifications to the existing network whenever a new track is inserted (Ahmed and Wenk 2012). Cao and Krumm (2009) propose a model for integrating input trajectories by applying a physical attraction model between spatially close segments before sequentially integrating them into a common map.

Under a broader view, the trajectories not only reflect the movement restrictions that the road network imposes, but also the driving behavior limitations that road regulations introduce. Thus, not only the geometrical and topological features of the road network but also the local maneuver restrictions that regulations initiate (e.g. do not turn left, one-way road, compulsory stop, etc.) are implicitly contained in the trajectories. Thus, in addition to geometry and topology, researchers have also tried to extract additional spatial features of road networks such as the width and number of lanes for each road, turning restrictions and other traffic related features (Cao and Krumm 2009, Chen and Krumm 2010, Zhang et al. 2010).

Traffic navigation systems do contain this kind of information; however the problem is to keep it up-to-date. Consequently, there is a need to discover and update modifications in the underlying network. Furthermore, crowdsourcing this kind of information allows to also reveal typical behavior of travelers, e.g. slowing down in a curve to a moderated velocity. This could be used as a recommendation in navigation systems, leading to more adequate and situation adapted regulations. But also unexpected behavior due to the distinct environment can be detected (e.g. a tree preventing a good overview of the next section of a road).

The problem of recognition, prediction and modeling of driver's behavior or intent has been addressed with many different techniques. Torkkola et al. (2004) use Hidden Markov Models (HMMs) for modeling sensor sequence for maneuver classification, Oliver & Pentland (2000) use Coupled HMMs (CHMMs), an extension of HMMs, to create models of seven different driver maneuvers. Aoude et al. (2012) classify behavior at intersections using Support Vector Machines.

The determination of group movement patterns has been dealt with in several ways: there are approaches, which aim at the detection of pre-defined generic or behavioral movement patterns, such as the encounter, flock or leadership (Benkert et al., 2008, Laube et al., 2008, Andersson et al., 2008). On the other hand, the goal can be to determine a priori unknown patterns. Examples are proposed by Laube et al., (2004) and Tsai et al., (2011), where the authors detect relative movement patterns.

4. EXTRACTION OF ROAD STRUCTURE AND ROAD REGULATIONS

The underlying idea of trajectory analysis is that they record a spatio-temporal phenomenon. Whereas a single trajectory alone may be inaccurate due to noise in the measurement instrument or even malicious handling, the information is getting more reliable, the more measurements vote or confirm it. Confirmation can be both spatially (i.e. overlap of trajectories indicating a travel path) or contextually (e.g. similar sequence of activities, indicating an underlying regulation). Examples for both aspects are shown in the following.

4.1. Three-step process for traffic network reconstruction

Trajectories are an interesting data source to capture and update the road traffic network. There are, however, some considerations to take: roads obviously can only be detected, when they are used; also, the more trajectories are created on a road, the better is the reliability of the detection. The properties of road trajectories can further be characterized as follows: the inaccuracy of the GPS-coordinates and the often irregular spacing of the track points lead to wide, blurred, bands of trajectory collections, which occur especially at junctions, whereas the roads in between are relatively well delineated. This can be seen in Figure 2.

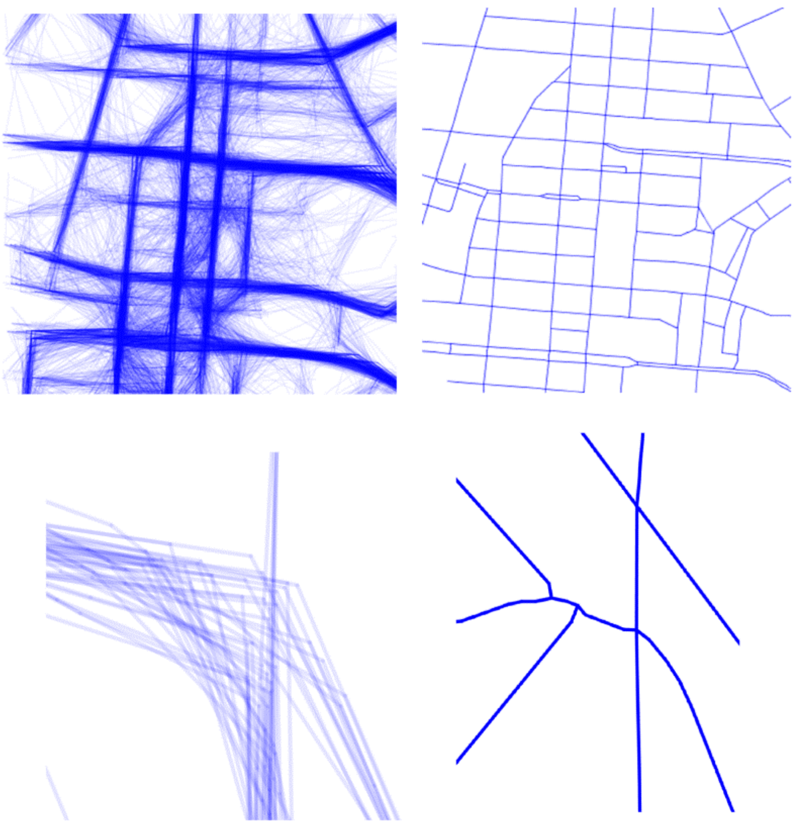

Figure 2: Examples of trajectories (left) and ground truth (right).

Considering these characteristics lead to the design of a three-step approach, which combines a raster based method to estimate the road segments, with a vector based optimization which estimates type and location of the junctions linking the road segments.

In a first step, the existence and reliability of roads is evaluated using a raster based aggregation with KDE (Kernel Density Estimation), followed by a thinning. KDE produces good road segments and is able to reduce or ignore GPS errors; however, it loses the topology and has problems in reconstructing junctions. Thus, KDE leads to a general delineation of the road network, resulting in a set of unconnected road segments in raster representation (see Figure 3b). In step 2 the original trajectories are matched to the partial road network obtained in step 1. For each KDE road segment the set of corresponding samples is obtained, allowing a segmentation of the original trajectories into a disjoint sequence of sub-trajectories, each corresponding to either one of the previously identified road segments or a transition between two adjacent road segments, indicating missing parts within the reconstructed road network (see Figure 3c). The transitional sub-trajectories are aggregated into sets of transitions corresponding to the same junction between a set of road segments. This leads to both coarse locations, as well as a full topological description of missing junctions.
In the final step, the junctions are reconstructed based on this information and additional prior information about typical junction types. This is achieved using a generative modeling approach (similar to Chai et al., 2013). It samples junction instances created from pre-defined parameterized junction models, optimizing the reconstruction towards more desirable configurations (details of the approach are described in Kuntzsch et al., 2015).

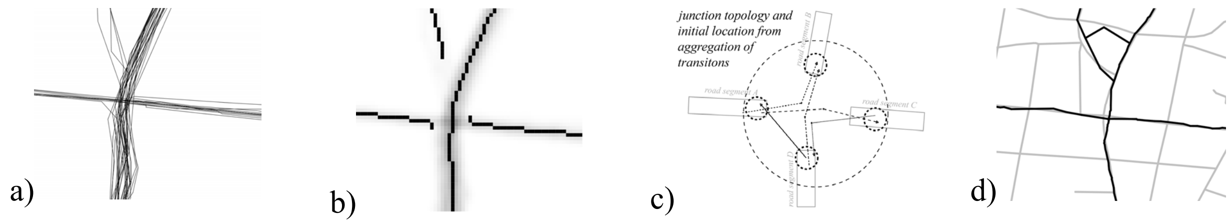

Figure 3: processing steps of the approach: a) raw trajectories, b) KDE histogram and road segment blobs, c) trajectory linking at junctions, d) final reconstruction result on top of ground truth map).

The approach was applied to a benchmark data set described in Ahmed et al. (2014), and compared to other methods described in the literature.

The results in Figure 4 show an overlay of raw trajectories and reconstruction results on OSM used as ground truth. Row (a) shows that the approach is able to extract and identify the correct road, although the trajectories are very noisy. In (b) a low-error scene is shown, which is also nicely recovered. The examples show that the approach is tolerant to a high level of noise, yielding a good representation of the road network. Row (c) shows effects due to not filtering out low representation road segments: a single trajectory sample leads to an instance of a hypothetical road. Such isolated occurrences (outliers) can be filtered out during step 2, based on the number of trajectories associated with road segments.

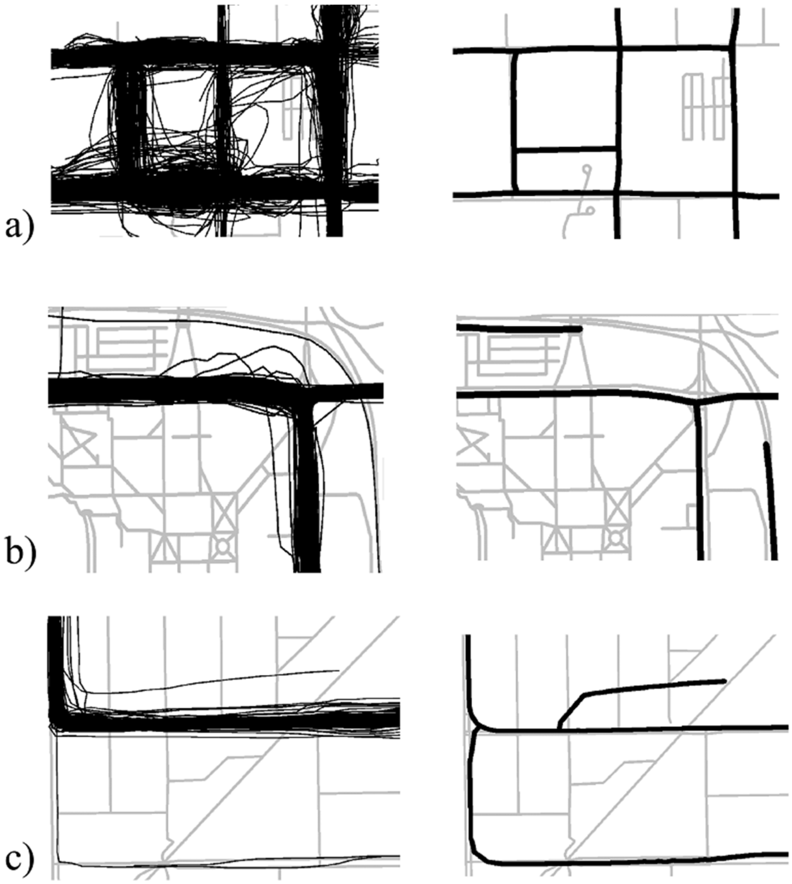

Figure 4: detailed screenshots of results: left: original trajectories, right: result of road extraction; both overlaid on OSM ground truth.

4.2. Extracting traffic regulations from trajectories

Cars are equipped with a lot of sensors. The data can be stored on the CAN-bus and exploited for further analysis, especially in the context of behavior recognition. The idea for using blinker and brake sensors is quite intuitive: driving maneuvers are composed of a sequence of actions which shows great similarity among drivers, with small variance in regard to the context (similar sequence of actions for the same maneuver at different locations). For example, when someone wants to turn, he applies the brakes for reducing vehicle's speed, indicates the intended direction (blinkers) and then turns. Figure 5 shows four instances of driving maneuvers (turn left, turn right, compulsory stop), where the spatiotemporal features of the blinkers and brake are presented. Consecutive samples from an activated sensor (red for blinkers and yellow for brake) form a trajectory which overlays with vehicle's trajectory (blue). The order of these two actions can be reversed, which means that depending on the surrounding traffic, a driver might first brake, e.g. if other vehicles are stopped in a close distance from the intersection spot, or first indicate his intention to turn before he starts decreasing the speed, e.g. if there is no vehicle in front to enforce him to brake earlier from the intersection. It can also be the case, a short time interval to mediate after braking until direction signaling be activated (Figure 5, bottom left) because the speed has been already regulated for the intended maneuver and not further braking is needed. These observations motivate the usage of data from blinker and brake sensors, as their activation "constitutionally" accompanies such driving maneuvers.

A two-step method is proposed for identifying potential valid road rules:

i. *Clustering of the trajectories:* we find clusters of similar trajectories, that is, trajectories that show the same maneuver pattern.
ii. *Analysis of the extracted clusters:* given a standard 4-way intersection of two roads, the clusters extracted from the previous step can be examined as follows. Entering the intersection from an incoming road:
   1. If only a right turn maneuver pattern has been identified (cluster of right turn), a compulsory *right turn* rule is valid.
   2. If only a left turn maneuver pattern has been identified, a compulsory *left turn* rule is valid.
   3. If neither right *or* left turn pattern is observed, a compulsory *drive through the intersection* rule is valid.
   4. If 1, 2 and 3 are not valid, and no turn-right cluster exists, then a *do not turn right* rule is valid.
   5. If 1, 2 and 3 are not valid, and no turn-left cluster exists, then a *do not turn left* rule is valid.
   6. If 1, 2 and 3 are not valid, and no go-through- the-intersection cluster exists, then a *do not drive intersection* rule is valid.
   7. If a stop maneuver pattern is detected, a compulsory stop rule is valid.

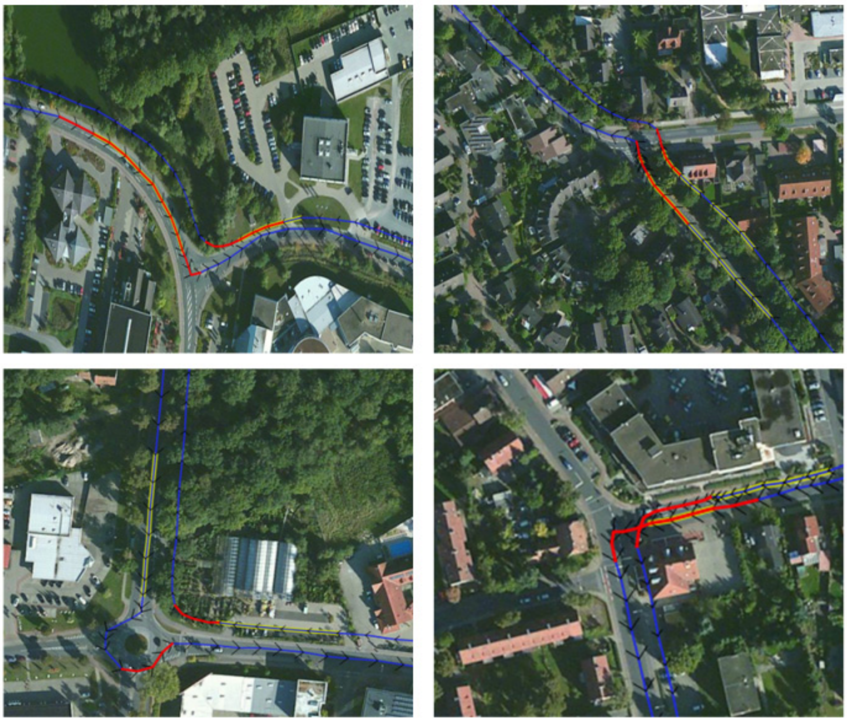

Figure 5: Examples for driving maneuvers visible in blinker (red) and brake (yellow) sensors.

Counting on crowdsourced data for rule extraction relies on the assumption that drivers respect the rules, so the samples are expected to be in agreement with each local regulation context. Nevertheless, the dataset might still contain “violating” samples; either due to low attention or deliberate action (e.g. TomTom reports that a one way road is only accepted as such, when 97% of all trajectories go in one direction). These samples need to be identified as such and their behavioural pattern not to be taken into account. Recognising anomalous behaviour patterns such as detour (see next paragraph and Huang et al., 2014) due to loss of way or low attention level and reckless driving (aggressive

acceleration, braking, left-right and u-turns) could explain individual samples or clusters composed by few members and consequently they could be excluded from the rule mining process.
An initial evaluation of the approach shows promising results. In Figure 6 a scene with two trajectories is shown, where the spatial correlation of similar driving actions can be observed. Different drivers apply alike sequences of actions in the same spatial context and by clustering the multi-dimensional samples, the repetitive location-based behavior as posed by the valid regulation set can be revealed. Solid circles indicate a coincidence of the same behavior, dashed circles show an event which is observed only in one trajectory. Green circles show a turn action whereas violet circles indicate a braking in the middle of a road segment. Such a behavior might be very interesting for the analysis of accident prone sections of a road and might give hints to possible improvement of the roads and its environment.

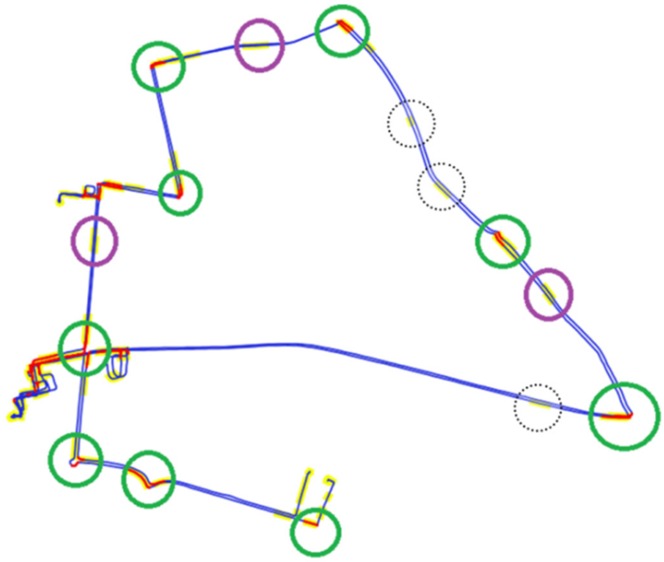

Figure 6: Example for behavior patterns: green circles: turn-situation (braking and blinking); violet circles: braking situation; dashed circles: only one trajectory shows braking behavior.

The next step is to test the proposed methods with larger amounts of data. Also, instead of defining rules beforehand, it is also possible to learn the rules from the data. This can be done both based on spatial clustering to detect rules, which are valid at a certain location, but also based on clustering of the attributes alone, in order to detect general rules. The future aims at a general framework of regulation-aware navigation. Distinctive challenges are the extraction of the local regulatory context, proposing a 2-step clustering-based method and the real-time prediction of the intended driving maneuvers using Hidden Markov Models (Zourlidou and Sester, 2015).

5. EXTRACTION OF ANOMALOUS BEHAVIOUR

Anomalous pattern detection is of great interest for navigation, driver assistance systems, surveillance as well as crisis management. An approach using a Bayesian filter has been proposed by (Huang et al., 2014, Huang, 2015). The idea is to determine where the behavior of a driver shows anomalies. Behavior is considered anomalous, when the driver encounters problems, such as taking a wrong turn, performing a detour or repeating a part of a road. In the approach, these anomalies are described in terms of three high-level features, i.e., turns and their density, detour factor and route repetition. These features are extracted from the given trajectory geometry, for which a long-term perspective is required to observe data sequences of a significant length instead of individual time stamps. High-

order Markov chains with a “dynamic memory" are employed to model the trajectory integrating these long-term features. The Markov model is processed by a recursive Bayesian filter to infer an optimal probability distribution of the potential anomalous driving behaviors dynamically over time. The filter performs an unsupervised detection in single trajectories based on local features only. As the behavior probabilities of the local features are given, no training process is required to characterize the anomalous behavior. By analyzing the detection results of individual trajectories collective behaviors can be derived indicating traffic issues such as congestions and turn restrictions.
The following figure shows the result of the analysis of a trajectory (Figure 7). It visualizes the belief in an anomalous behavior of a trajectory going from bottom to top. The belief in anomaly increases as the driver makes a U-turn; after returning to the standard route, the belief in anomalous behavior decreases again.

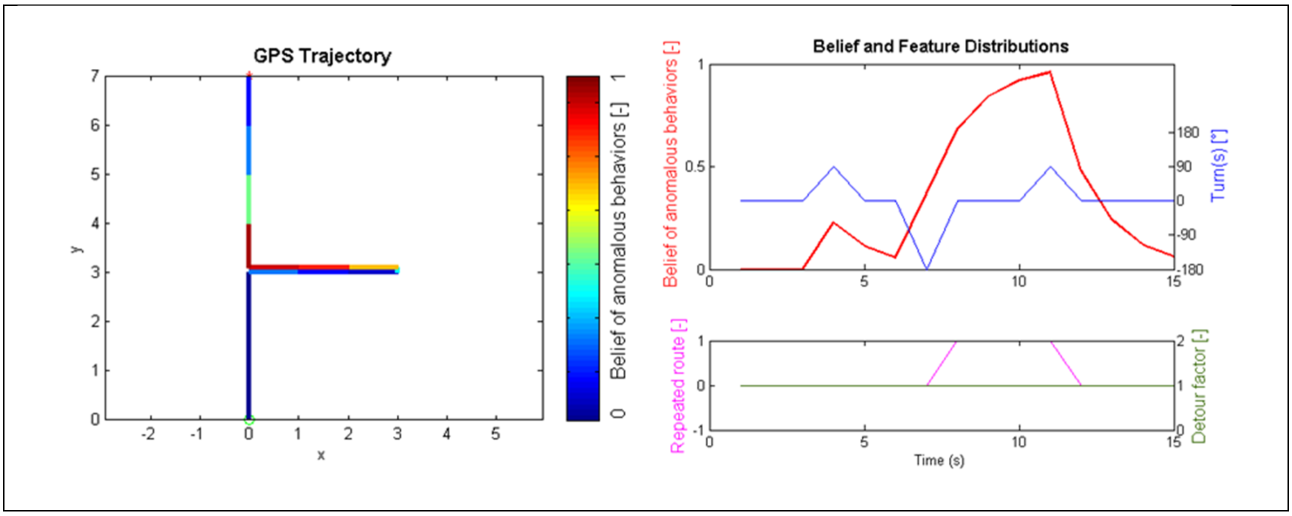

Figure 7: Belief in anomalous behavior of route from bottom to top.

Figure 8 shows a larger scene, where collective behavior is observed. In the scene, a temporary road block leads to the fact that cars cannot pass the road from north to south, but had to go back (2) and take a detour (1). Similarly, also the access from the south (3) is blocked and leads the driver to take another route to his/her destination.

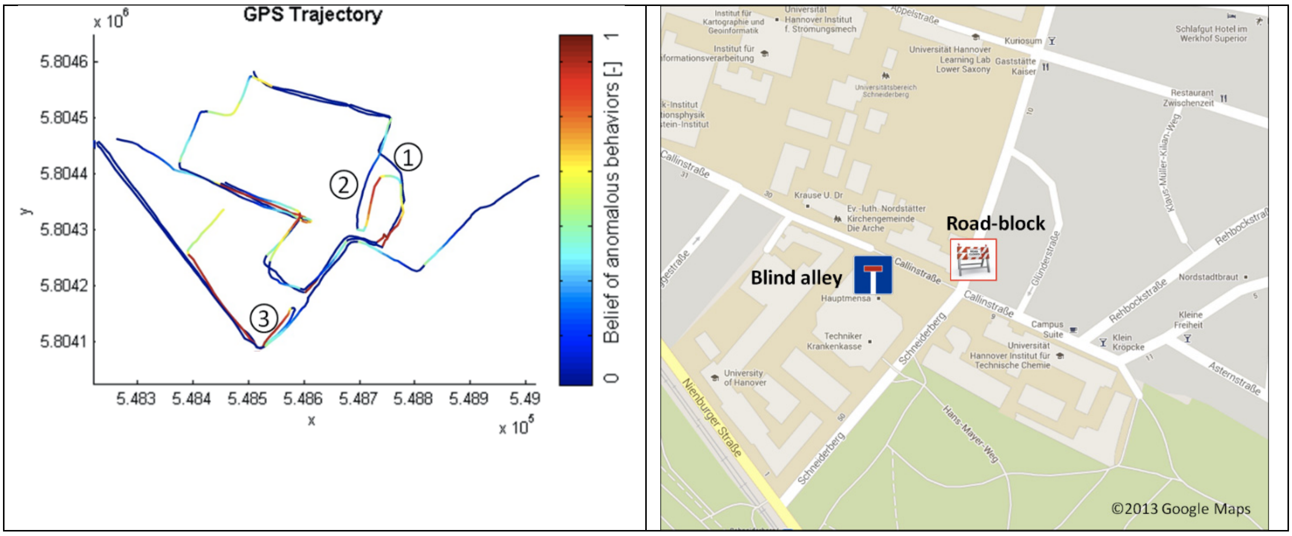

Figure 8: Example for collective behavior detected at road block.

6. DETECTION OF GROUP PATTERNS – IN SOCCER

A general pattern can be defined as a series of repetitive constellations of a given group of objects. Thus, any pattern can be detected, as long as it occurs more often in the data set. Consider as an example the behavior of soccer players during a match. The players represent a constant group of objects, which moves in space. We can observe typical movements patterns, for example, repetitive attack or defense patterns or cycles of offense and defense formations. During these situations, the players form spatial constellations (e.g. distribution in own field vs. around the adversary's goal); sequences of constellations are characteristic for a certain tactic maneuver. As such patterns may occur at different locations in the field and may be rotated or in a different scale, transformation invariance required while designing an interpretation method. The method presented in the following is based on a cluster analysis followed by a sequence analysis method and is described in detail in Feuerhake and Sester, 2013). It is a generic method to analyze groups of a constant size, which can be adapted to the special case of analysis of soccer players.

The algorithm is based on so-called object constellations, which describe the positions of objects relative to each other by position relations, e.g. distances, angles, etc. This leads to $\binom{n}{2}$ position relations stored in a constellation; where n is the number of observed objects. In our case a constellation represents a formation of a team at a certain point in time. In order to be able to also detect transformed (translated, rotated or scaled) patterns, suitable invariant relations to be stored in a constellation have to be chosen. The constellations are described by a vector of these relation values. Using the distances between the positions makes a constellation invariant regarding rotation and translation. Replacing the distances by the relationships of the distances to the sum of all distances $d_i = \frac{a_i}{\sum_j a_j}$, turns the constellations also scale-invariant. Figure 9 shows three constellations, where b) is a rotated and c) a scaled version of the original a).

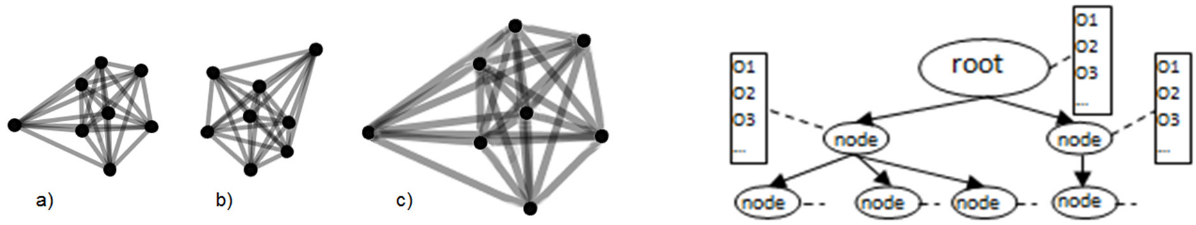

Figure 9: Left: Depending on the desired invariance regarding translation, rotation (b) and scale (c) the shown constellations will be treated to be equal to (a). Right: A scheme of the tree-like data structure, which is extended by object sets at the nodes.

Due to the inherent inaccuracies of the measurement devices (e.g. GPS), it is not expected to find exactly matching sequences of constellations in a limited observation timespan. Therefore, the goal is to find similar constellations – instead of equal ones. As a similarity measure, the distance of the position relation vectors of the constellations is used.

In a first step, similar constellations are clustered using a density-based clustering, like DBSCAN, or a centroid-based clustering method, such as k-means. The clustering is applied to the data at a given sampling frequency, which can be chosen according to the type data set. The required clustering parameter, e.g. the number of clusters in case of the k-means algorithms, can be used to control the similarity of the clustered constellations. After clustering, all the constellations are assigned to clusters with a corresponding ID. This leads to temporal sequences of constellations with corresponding cluster-IDs.

As the goal is to identify repetitive subsequences of cluster-constellations, in the second step a sequence mining approach is applied (Han et al., 2007). This approach has two input parameters: the minimal subsequence length l and the minimal count k this subsequence occurs.
The algorithm builds up tree-like data structures, containing the clusters as nodes, which carry sets of references to the original constellations. Each tree contains sequences of potential movement patterns starting with the same root (constellation). The typical tree structure is extended by the possibility to carry sets of objects (in our case the original constellations) at the nodes (Figure 9, right).
Subsequently, each constellation is linked with its predecessor. This feature is used in the last step of our algorithm, in order to trace back a found movement pattern in the real data. In the first step the roots of the trees are determined. This is achieved by iterating over the elements of the original sequence and enumerating the occurrence indices of equal subsequences with the length of l. If the minimum frequency condition is fulfilled, trees will be created with those equal subsequences as their roots. The remaining subsequences are discarded, because they cannot be part of a pattern due to their count, which is lower than k.
In the next step the trees are built up with the help of sequence parts, which result by segmenting the original sequence at the corresponding occurrence indices. The original sequence is cut into parts. Each resulting sequence part is then attached to the tree started at the root element wise. If a sequence part is added, it starts at the root element and the corresponding indices to the index sets are added. While doing that two cases are distinguished: if the current node already has an equal child node, the corresponding sequence index is added to the existing set of indices, otherwise, if there is no equal child node, a new branch is started taking the sequence index as the first entry in the index set.
The final step consists of extracting the patterns from the trees. This is done by extracting all paths (patterns) through the trees, which exclusively consist of nodes storing at least k constellations. Thus, we do not look for patterns in the cluster sequence but in the corresponding sequence of constellations, we have to remap the found cluster patterns to the original constellation subsequences to get the actual movement patterns. To this end, we run through the resulting tree paths, starting at the end, using the constellations' pointers to jump to the previous until the beginning of a path is reached. Those constellation sequences, which are longer than l, are stored. All constellation sequences that resulted by running through the same path belong to the same pattern.

6.1. Example: Soccer data set (small field with only 8 players)

In the following example, the approach is applied to a soccer data set, consisting of a small team with only 8 players. As described in Feuerhake and Sester (2013), the resulting patterns depend on the selection of the input parameters. If we are interested in a few, highly similar patterns, this can be achieved by selecting a high count of clusters.

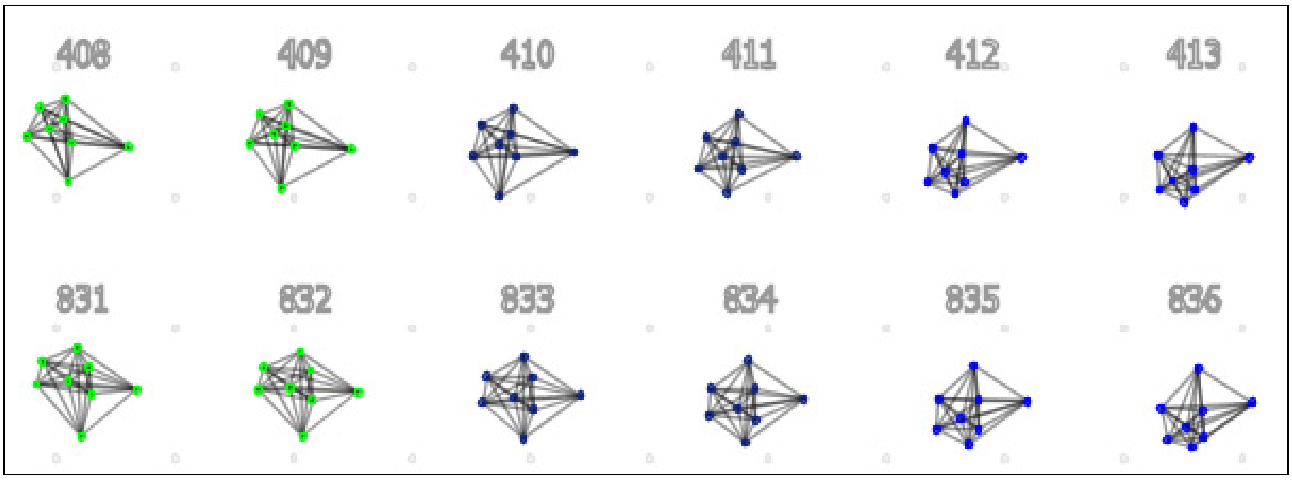

Figure 10: Movement pattern (sequence of 6 clusters GGBBLL) was observed twice in the data set: sequence 408-413, and sequence 831-836.

Figure 10 shows an example of one movement pattern (constellations 408-413), which lasts 6 seconds and consists of a sequence of 3 different group patterns (2 Green, 2 Blue, 2 Light blue clusters: GGBBLL). This GGBBLL pattern occurred twice in the whole data set, namely a second time in the cluster sequence 831-836). The interpretation of the pattern is that three out of 8 players (top, right and bottom) stay relatively constant, where as the other 5 players move from top to bottom.

7. CONCLUSIONS

In the paper trajectories as new data source have been described and possible applications have been shown. A major advantage of the data is that they are massively available and delineate movement paths and behavior in high detail. It also reflects the temporal situation and thus has the potential to detect also high frequency phenomena and patterns. A popular example is the extraction of traffic situations in Google or TomTom navigation systems.

To reveal this information requires sophisticated analysis procedures, which take the characteristics of the data into account. The paper presented three types of applications which need different algorithmic approaches. The examples showed that GPS trajectories produced by cars have the potential to delineate not only the road geometry, but also reveal behavioral patterns, which can be used to infer traffic regulations or behavior indicating when a driver needs assistance or help. The last example showed that not only individual behavior is of interest, but also group patterns can be extracted from trajectories. They can be beneficially exploited in the sports context, but also in other applications such as animal ecology.

A fundamental challenge in the context of trajectory analysis lies in the fact that trajectories are usually created by an individual (person). This immediately leads to privacy issues, which have to be carefully considered. To this end, several approaches have been proposed, ranging from stripping initial and final trajectory segments, to anonymizing methods such as obfuscation or cloaking (Gruteser & Grunwald, 2003; Duckham & Kulik, 2005)

8. ACKNOWLEDGEMENTS

The work is being funded by several institutions, e.g. BMBF and DFG, whose support is gratefully acknowledged.

9. REFERENCES

Ahmed, M. and Wenk, C. (2012). Constructing street networks from GPS trajectories. In Algorithms–ESA 2012 (pp. 60-71). Springer, Berlin/Heidelberg.

Ahmed, M., Karagiorgou, S., Pfoser, D. and Wenk, C. (2015). A comparison and evaluation of map construction algorithms using vehicle tracking data. GeoInformatica, 19 (3), pp. 601-632.

Andersson, M., Gudmundsson, J., Laube, P. and Wolle, T. (2008). Reporting leadership patterns among trajectories. GeoInformatica 12 (4), pp. 497-528.

Aoude, G. S., Desaraju, V. R., Stephens, L. H. and How, J. P. (2012). Driver Behavior Classification at Intersections and Validation on Large Naturalistic Data Set. Intelligent Transportation Systems IEEE Transactions. 13 (2), pp. 724-736.

Benkert, M., Gudmundsson, J., Hübner, F. and Wolle, T. (2008). Reporting flock patterns. Computational Geometry, 41 (3), pp. 111-125.

Biagioni, J. and Eriksson, J. (2012b). Map inference in the face of noise and disparity. Proc. 20th ACM SIGSPATIAL, pp. 79-88.

Cao, L. and Krumm, J. (2009). From GPS traces to a routable road map. Proc. 17th ACM SIGSPATIAL, pp. 3-12.

Chai, D., Förstner, W. and Lafarge, F. (2013). Recovering line-networks in images by junction-point processes. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Portland, US.

Chen, Y. and Krumm, J. (2010). Probabilistic Modeling of Traffic Lanes from GPS Traces. 18th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACM SIGSPATIAL GIS 2010), November 2-5, San Jose, CA.

Davies, J. J., Beresford, A. R. and Hopper, A. (2006). Scalable, distributed, real-time map generation. IEEE Pervasive Computing, 5 (4), pp. 47-54, Oct. 2006.

Demšar, U., Buchin, K., Cagnacci, F., Safi, K., Speckmann, B., Van de Weghe, N., ... and Weibel, R. (2015). Analysis and visualisation of movement: an interdisciplinary review. Movement ecology, 3 (5).

Dodge, S., Weibel, R. and Lautenschütz, A. K. (2008). Towards a taxonomy of movement patterns. Information visualization, 7 (3-4), pp. 240-252.

Duckham, M. (2012). Decentralized spatial computing: foundations of geosensor networks. Springer Science & Business Media.

Duckham, M. and Kulik, L. (2005). A formal model of obfuscation and negotiation for location privacy. In Pervasive computing (pp. 152-170). Springer Berlin Heidelberg.

Edelkamp, S. and Schrödl, S. (2003). Route planning and map inference with global positioning traces. Computer Science in Perspective. Springer, pp. 128-151.

Fathi, A. and Krumm, J. (2010). Detecting road intersections from GPS traces. Proceedings 6th international conference on geographic information science, pp. 56-69.

Feuerhake U. and Sester, M. (2013). Mining Group Movement Patterns. In Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (pp. 530-533). ACM.

Gruteser, M. and Grunwald, D. (2003). Anonymous usage of location-based services through spatial and temporal cloaking. In Proc. MobiSys '03, pp. 31-42.

Han, J., Cheng, H., Xin, D. and Yan, X. (2007). Frequent pattern mining: current status and future directions. Data Mining and Knowledge Discovery, 15 (1), pp. 55-86.

Huang, H., Zhang, L. and Sester, M. (2014). A Recursive Bayesian Filter for Anomalous Behavior Detection in Trajectory Data. In Connecting a Digital Europe Through Location and Place (pp. 91-104). Springer International Publishing.

Karagiorgou, S. and D. Pfoser. On vehicle tracking data-based road network generation. In Proceedings of the 20th International Conference on Advances in Geographic Information Systems, California, 2012.

Kuntzsch, C., Sester, M. and Brenner, C. Generative Models for Road-Network Reconstruction, Intern. Journal of Geographical Information Science, in review.

Laube P., Duckham M. and Wolle T. (2008). Decentralized Movement Pattern Detection amongst Mobile Geosensor Nodes. In 5th GIScience Conference, Park City, UT, USA, September 23-26., Springer, Heidelberg.

Laube, P., van Kreveld, M. and Imfeld, S. (2004). Finding REMO – detecting relative motion patterns in geospatial lifelines. In: 11th Int. Symp. Spatial Data Handling, pp. 201-214.

Oliver N. and Pentland, A. (2000). Graphical models for driver behavior recognition in a SmartCar. In Proceedings of the IEEE Intelligent Vehicles Symposium.

Soleymani, A., Cachat, J., Robinson, K., Dodge, S., Kalueff, A. and Weibel, R. (2015). Integrating cross-scale analysis in the spatial and temporal domains for classification of behavioral movement. Journal of Spatial Information Science, (8), pp. 1-25.

Torkkola, K., Venkatesan, S. and Huan Liu (2004). Sensor selection for maneuver classification. In Proceedings of the 7th International IEEE Conference on Intelligent Transportation Systems.

Tsai, H. P., Yang, D. N. and Chen, M. S (2011). Mining group movement patterns for tracking moving objects efficiently. IEEE Transactions on Knowledge and Data Engineering, 23 (2), pp. 266-281.

Zhang, L., Thiemann F. and Sester, M. (2010). Integration of GPS Traces with Road Map. In Proceedings of the Workshop on Computational Transportation Science in conjunction with ACM SIGSPATIAL, San Jose.

Zourlidou, S. and Sester, M. (2015). Towards regulation-aware navigation: a behavior-based mapping approach. In Proceedings of the AGILE, Lisbon, Portugal.