

Mining Group Movement Patterns

Udo Feuerhake and Monika Sester
Institute for Cartography and Geoinformatics
Appelstrasse 9a
30167 Hannover
+49 511 762 19369

{Feuerhake, monika.sester}@ikg.uni-hannover.de

ABSTRACT

In this paper we aim to recognize a priori unknown group movement patterns. We propose a constellation-based approach to extract repetitive relative movements of a constant group, which are allowed to be rotated, translated or differently scaled. To this end, we record a sequence of constellations, which are used for describing the movements relatively. We deal with uncertainties, and similarities of constellations respectively, by clustering the constellations. Further, we have developed a sequence mining algorithm, which uses the clustering results and tree-like data structures to extract the requested patterns from the sequence. Finally, this approach is applied to different datasets containing real trajectory data provided by different tracking devices. By this way, we want to show its portability to different use cases.

Categories and Subject Descriptors

I.5.3 [Pattern Recognition]: *Clustering, Algorithms, Similarity measures*

General Terms

Algorithms.

Keywords

Pattern Mining, Spatio-Temporal Analysis, Movement Patterns, Clustering, Constellation

1. INTRODUCTION

Analyzing the behavior of moving objects provides important knowledge which is needed for tasks from a large variety of application areas. For instance, with the help of this knowledge urban planners can provide better designs for highly frequented buildings (shopping malls, stations, etc.), surveillance systems can be adapted and become more intelligent, or (social) behavior studies of human beings or animals can be supported. Often the research focuses on movement patterns, especially on group patterns. Based on such patterns movement models can be developed which represent the actual movement behavior and enable subsequent more advanced analysis tasks like movement prediction, outlier detection, etc.

In our work we aim to recognize group movement patterns. In contrast to other research, which deals with the detection or identification of pre-defined generic or behavioral movement patterns, such as the encounter, flock or leadership patterns, we aim to extract a priori unknown patterns. We consider a pattern to be 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 observation of soccer players during a match. The players represent a constant group of

objects, which moves in space. We can find typical movement 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. Since those patterns may occur at different places in the field and, furthermore, may be rotated or in a different scale, but still remain the same pattern, we have to consider these transformation invariance while designing our mining method.

The approach described in the paper consists of two parts: in a first step, typical constellations of the group members are identified using clustering. In this way, static patterns at temporal snapshot are found. In the second step, typical sequences of the clusters are searched in the data, leading to typical movement patterns.

The paper is organized as follows: after a discussion of related work, the approach is described in detail, using an illustrative example from soccer. Then results of experiments with several trajectory data sets are presented. The dataset are generated by different kinds of tracking devices; thus they show its portability to different use cases. The paper ends with a conclusion and an outlook on future work.

2. RELATED WORK

There are many approaches that deal with the identification of movement patterns. Most of these approaches are looking for given movement patterns, as e.g. classified by [7], such as the encounter [10], flock [4],[10],[15] or leadership [3],[10] pattern. In contrast to these approaches, the work in this paper aims at detecting typical behavior of groups of a given number of members. The work of [14] is similar to ours. The authors also detect relative movement patterns, which, however, are a priori known. For this, they calculate a matrix of motion attributes based on the objects' trajectories. With the help of this matrix they create a so-called analysis matrix, in which they look for patterns. To this end, they match generic motion pattern with the patterns generated by the matrix coefficients. In contrast to the REMO-approach we do not know what kind of patterns we are look for. So, an approach which matches patterns is not useful for us.

In the research field of sequence mining many approaches have been proposed. There is a group of approaches dealing with mining for patterns in itemsets. Han et al. [11] give a comprehensive overview of current and future approaches. The problem solved by the latter is not exactly matching our problem, because we search for patterns in one sequence instead of dealing with several itemsets.

Another group presents approaches to mine for patterns in strings. There is also an overview of methods given by Abouelhoda and Ghanem in [1]. In the section dealing with repeat-related problems they have created a taxonomy of the major repeat-related problems. Considering their taxonomy our problem belongs to the domain of dispersed repeats with a variable length. According to this classification they propose the algorithm of Gusfield which is based on a suffix tree. Cao et al. [5] present two variants of a pattern mining algorithm, STPMine1 and STPMine2ki. In [12] Han et al. describe their algorithm to mine partial periodicity with a max-subpattern hit set method.

Those algorithms cannot be applied to our problem directly, because we, instead of searching the sequence of constellations, look for patterns in a resulting sequence of clusters. After this we have to remap the cluster sequence patterns to the original sequence of constellations to extract the constellation sequence pattern, which we are looking for. The methods described above do not provide a feature which allows such a mapping between two sequences.

3. APPROACH

The approach consists of two steps: first, typical group patterns are identified using clustering. A prerequisite is that the patterns are invariant to certain transformations, namely translation, rotation and scale. Thus, the parameters used to describe potential clusters have to take these invariants into account. Each of the clusters can be seen as a spatial sub-pattern. Clustering is applied to temporal snapshots of the data set. It leads to an assignment of cluster-identifications to each of these snapshots, e.g. cluster A is followed by clusters B, C, A, etc., yielding a sequence of sub-patterns.

In the second step, typical sequences of sub-patterns have to be identified in the whole sequence, leading to typical movement patterns. To this end, sequence analysis is applied.

3.1 Identification of typical (static) group patterns using clustering

In order to find typical group patterns, a method which is based on object constellations is designed. A constellation, known from the astronomy, describes the positions of objects relative to each other by position relations. There are several ways to describe the relative position relationships between a set of points: e.g. using coordinates, distances, angles, etc. Depending on this selection, there are $\binom{n}{2}$ position relations stored in a constellation; where n is the number of observed objects.

In the case of the soccer example, a constellation represents a formation of the players at a certain point in time. Due to the transformation invariance requirement given by the application scenario we have to choose, which position relations describe a constellation. 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. If we replace the distances by the relationships of the distances $b_i = \frac{a_i}{\sum_j a_j}$, constellations will be also scale-invariant (see Figure 1).




Figure 1: An example for a constellation. On the right the position relations are given in terms of distances and relative distances.

Figure 2 visualizes the requirements concerning transformation invariance. Depending on the application, the three scenarios are considered as equal. Thus, the parameters describing the constellations have to be invariant.




Figure 2: Equality of constellations. Depending on the desired invariance regarding translation, rotation (b) and scale (c) the shown constellations will be treated to be equal to (a).

In order to take small variations in the relative positions of the objects into account, the approach has to deal with uncertainties. This is on the one hand due to the inherent inaccuracies of the measurement devices (e.g. GPS or UWB-Tracking).

Further, we do not expect to find exactly matching subsequences of constellations in a limited observation time span: because the more objects are observed and the more precise their positions can be determined the lower the probability is equal constellations can be found. Because of that reasons we search for similar constellations instead of requiring them to be equal.

As a measure of similarity, we use the distance in the space, which is set up by the position relation vectors of the constellations. Similar constellations are clustered. To this end, we are able to use either a density-based clustering, like DBSCAN [8], or a centroid-based clustering method, like k-means [16]. In both cases reasonable parameters have to be found to prevent an over- or an under-classification.

The DBSCAN-algorithm requires a range parameter, which describes the affiliation of a point to a cluster and has to be given also in the $\binom{n}{2}$ -dimensional space of the constellations. The k-means-algorithm needs the number of resulting clusters as input, which is not known a priori. In the current approach, these values have been set using empirical investigations with the data we carried out the experiments. An algorithmic approach is planned in our future work.

A typical result of the clustering leads to similar constellations as shown in Figure 3. The similarity is expressed by the colors.




Figure 3: Clustering results. Constellations shown in the same color are part of the same cluster and are considered to be similar.

Clustering is applied to the data at a given sampling frequency. This can be the frequency, in which the data was recorded, or also depend on the application scenario and the expected object velocities. The lower the latter are, the less significant the objects' movements are and the lower the frequency can be.

The result of the first step is a sequence of constellations with corresponding cluster-names. It contains consecutive constellations recorded in 19 time steps. The clusters are classified with capital letters A-E (see Figure 6).

In this sequence of clusters we search for repeating subsequences. This is achieved by a sequence mining approach.

3.2 Identification of movement patterns in cluster sequences

In our case a pattern is a repetitive subsequence in the total sequence. 'Repetitive' means that a subsequence occurs at least twice. Further, we demand the subsequences to contain more than one element ($\text{length} \geq 1$). In order to find patterns with certain significance, which means a certain frequency and/or length, we parameterize the algorithm as follows: we introduce the parameters n , which describes the minimal frequency (min. frequency condition) a subsequence has to occur, and l , which describes the minimal length of subsequence. We use this as a basis for our algorithm, which builds up tree-like data structures containing the patterns, in several steps. In our case, the tree contains the clusters as nodes. In addition, the nodes also carry sets of references to the original constellations. Each tree contains

Figure 6: The example sequence has been clustered. The affiliation to the clusters is symbolized by the letters in the additional row.

T	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Id	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Cluster	A	B	C	B	E	A	D	E	A	B	D	C	A	B	C	D	B	C	D	A

sequences of potential movement patterns starting with the same root (constellations). Those data structures own the typical tree-structure, where the clusters are arranged at the leaf nodes, but are extended by the possibility to carry sets of objects (in our case the original constellations) at the nodes (see Figure 4).




Figure 4: A scheme of the tree-like data structure, which is extended by object sets at the nodes.

In addition, 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 we determine the roots of the trees. 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 . In our example, we set $n \geq 2$ and $l = 2$, which leads to the result shown in Figure 5. In the given sequence there are 4 rows containing more than $n=2$ occurrences, thus four trees will be created with their roots AB, BC, EA and CD. The remaining subsequences are discarded, because they cannot be part of a pattern due to their count, which is lower than n .

Figure 5: Occurrences of subsequences with the length 2 in the original sequence.

Subsequence	Indices(t)	n
AB	1,9,13	3
BC	2,14,17	3
CB	3	1
BE	4	1
EA	5,8	2
AD	6	1
DE	7	1
BD	10	1
DC	11	1
CA	12	1
CD	15,18	2
DB	16	1
DA	19	1

In the next step we build up the trees with the help of sequence parts, which result by segmenting the original sequence at the corresponding occurrence indices. The example sequence is cut into the parts presented in Figure 7.

Figure 7: Cutting sequence into parts at given occurrence indices: resulting sequence parts.

Subsequence	Sequence parts
AB	ABCBEADE ABDC ABCDBCDA
BC	BCBEADEABDCA BCD BCDA
EA	EAD EABDCABDCBCDA
CD	CDB CDA

Each sequence part is then attached to the tree started at the root element wise. Our example is illustrated in Figure 8. At the beginning our tree structure only consists of the sequence ABCB (first row). If we add the next sequence part ABCD (second row), we will start at the root element (AB) and add the corresponding indices (13, 14) to the index sets ({1}, {2}). After that we distinguish two cases: if the current node (B) already has an equal child node (C), we will add the corresponding sequence index (15) to the existing set of indices ({3}), otherwise, if there is no equal child node (like at node C), we will start a new branch (D) taking the sequence index as the first entry in the index set. In Figure 9(top) the tree structure is shown, which uses the sequence AB as a root.




Figure 8: Adding the sequence parts to the tree structure.

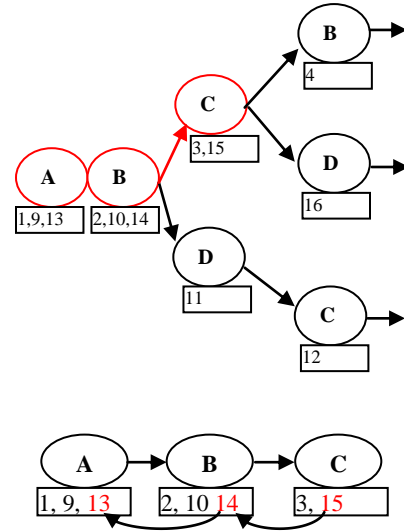


Figure 9: (top) An example for a resulting tree using the sequence AB as a root. The boxes at the nodes contain the stored constellation indices. The red colored path fulfills the min. frequency condition. (bottom) The path is run through to extract the constellation sequences (in reverse order).

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 n 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 (see Figure 9 (bottom)). Those constellation sequences, which are longer than 1, are stored. All constellation sequences that resulted by running through the same path belong to the same pattern: in the small example, two patterns of length 3 have been found, each has been observed twice in the data set (see Figure 10).

Figure 10: Results of the pattern mining algorithm applied to our example.

Pattern tree root	Pattern	Constellation sequences
AB	ABC	1,9,13 13,14,15
BC	BCD	14,15,16 17,18,19

3.3 Performance of the Algorithm

The performance analysis of our algorithm can be divided into the clustering step and the pattern mining step. Ester et al. [8] have determined a $O(N^2)$ complexity for DBSCAN. Inaba et al. [13] have shown that the k-means clustering, if k and d (the dimension) are fixed, has a complexity of $O(ndk + \log n)$, where n is the number of entities to be clustered. The performance (runtime) of the proposed sequence mining algorithm (the second part of our approach) is shown in the figure below. The values have been

calculated in a benchmark test by using randomly generated sequences with an increasing length. For the moment this approach has moderate runtimes even when it is applied to big datasets, but it does not have the capability to work at runtime. Therefore, only a post-processing analysis is possible.




Figure 11: Runtime [ms] of the mining algorithm related to the sequence length (number of elements).

4. Experiments and Discussion

We apply our mining algorithm to datasets recorded by different tracking devices and have different scales. We selected two datasets from the sports domain, soccer and handball. Another contains movement data from animal behavior studies, namely sea gulls movement recorded near Texel (Netherlands).

4.1 Experiment with soccer data set

The soccer dataset [17] has been recorded at the FraunhoferIIS with a Real-Time Locating System, which has been deployed on a football field of the Nuremberg Stadium in Germany. The dataset originally contains data of 16 soccer players (two teams with 8 players), one referee and four balls during a match of about 60 minutes. The positions of the objects have been recorded at a sampling rate of about 200Hz. Based on this dataset we have made different experiments. In each of it we have built a different observation group. In the first we have observed all the players but created a sequence of constellations for each team separately, instead of creating one common sequence. In the second case we have selected only a subgroup of objects to extract patterns exclusively originated by them: as the goal keeper is not really involved in the certain patterns, he was removed from the team groups, thus resulting in groups of 7 objects each.

For each of those experiments we have used a sampling rate of 1Hz to record our constellation sequence. This leads to 3560 constellations for the whole sequence of the roundabout 60 minute game.

In order to determine the parameters for the clustering, we conducted several tests. Figure 12 presents a parameter study of the corresponding clustering parameters for the two clustering methods used. It can be observed that a parameter setting, which leads to a higher number of clusters, also leads to lower max. and avg. lengths and frequencies. However, the sequences of each pattern become more similar to each other.

Figure 12: The clustering parameter study for the k-means and the DBSCAN algorithm using one sequence with 1765 constellations.

k-means			
# clusters	# patterns	Max./avg. length	Max./avg. frequency
5	566	69/11	34/15
10	653	35/6	32/16
50	532	21/5	12/7
100	331	21/5	7/5
1000	14	5/4	2/2
DBSCAN			
Range/ clusters	# patterns	Max./avg. length	Max./avg. frequency
1/4	0	-/-	-/-
3/28	12	5/4	3/2
10/105	272	16/9	6/5
12/44	315	17/9	7/5
15/4	358	19/10	8/6

For the first scenario (all 8 players, two teams), in which we used a k-means clustering (100 clusters), the algorithm determined over 300 movement patterns. We exemplarily show some of them in Figure 13.

The first two rows show a movement pattern, consisting of a sequence of 3 group patterns (green, dark blue, light blue): it shows, how five players in the group move top down. It occurs twice during the observation time. Of course, there are longer and more frequent patterns found, but this one's feature is easy to interpret. We have picked the second pattern in rows 3 and 4 (consisting of a sequence of 8 constellations in 3 groups): this example clearly shows the invariance regarding rotation: both sequences are part of the same pattern and show that both teams conduct similar movements patterns.

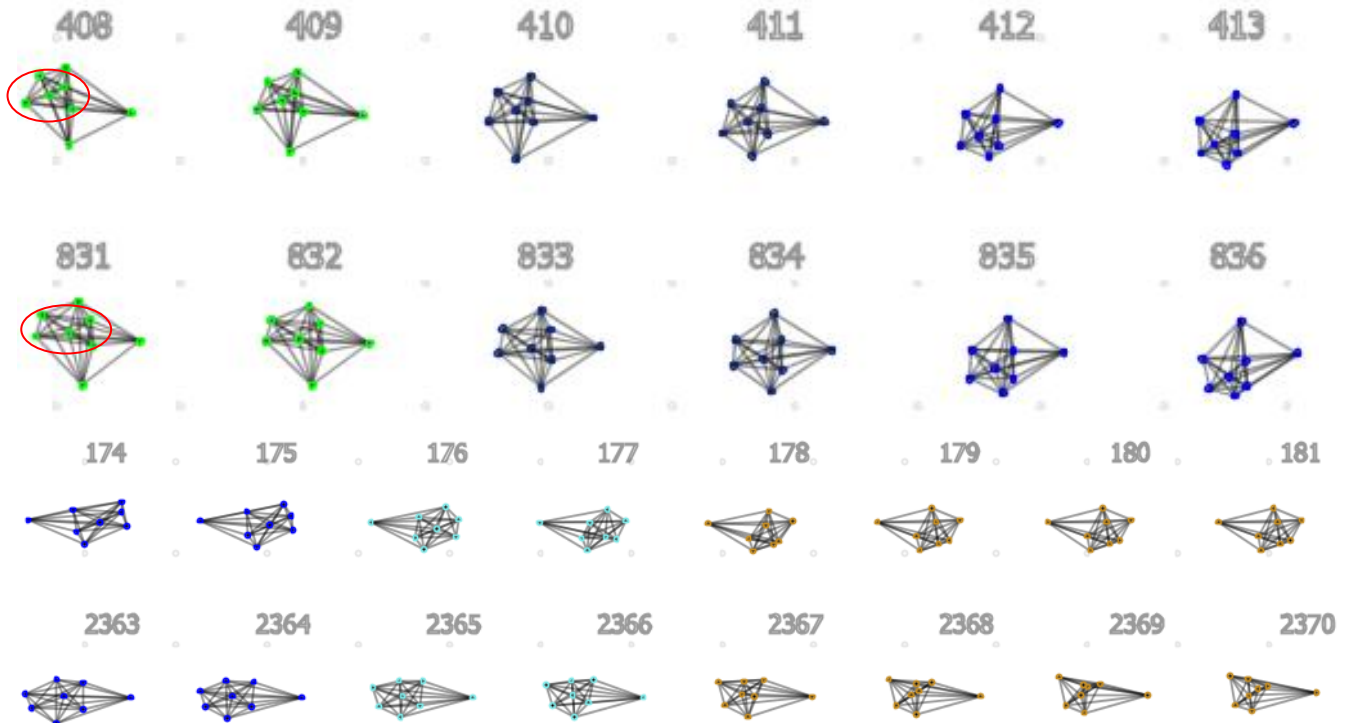


Figure 13: Resulting pattern for the scenario of observing a whole team. In this example the rotation invariance is shown.



Figure 14: Movement pattern found in subgroup (i.e. only 7 field players): this pattern has been found twice in the data set and describes the uniform movement of the majority of the group from top to down.

A result of the experiments, where only the field players were analyzed, is shown in Figure 14. The pattern consists of two sequences, which last for 11 time steps (constellations), assigned to 4 clusters (light blue, purple, red, green). In both sequences the encircled players uniformly move top down.

4.2 Experiment with handball data set

Our second experiment we have based on another trajectory dataset [6] of the sports domain. This time we have observed a handball team, which consists of seven players and is video-tracked for ten minutes.

The goal of the analysis was to extract attack and retreat patterns of the team. Also here, we investigated different player subgroups. Again we have removed the goal keeper from the observation in one case. For this experiment we keep the same constellation sampling rate, because we expect similar player velocities like in

the soccer example before. The clustering has been done by the k-means algorithm (20 clusters). This parameter setting leads to the following exemplary results (see Figure 15).

The first pattern (rows 1-3), which occurs three times, represents a transition between defending and attacking behavior. It shows that the player (red circle), which is in front of the defending formation, always crosses the center line first and moves to the outer left. The other players also seem to have given routes.

In the second pattern (rows 4 and 5) a special attack sequence is presented. There, starting from their standard attack formation, the center players (red circle) first move to the left and after that meet near the goal area.

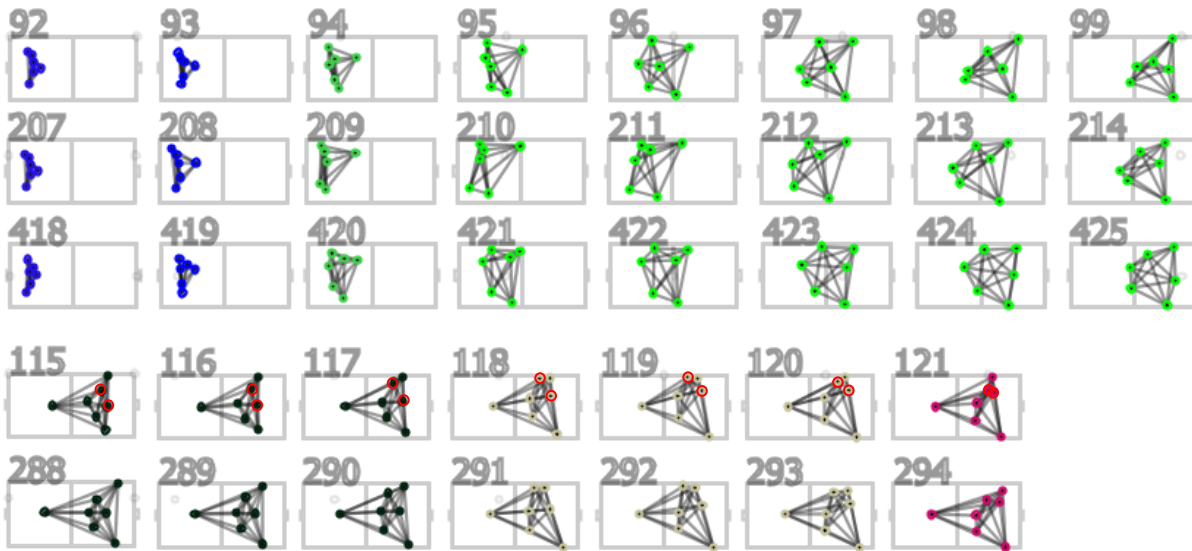


Figure 15: Resulting attack patterns extracted from the handball dataset. The first patterns (occurrence: three times) shows the transition between defending and attacking (neglecting the goal keeper), the second (occurrence twice) a special attack pattern of the whole team.

Figure 16 shows a retreat pattern performed by the whole team. This pattern occurs eight times during the ten minutes of observation. Each sequence (one per row) contains eight constellations, clustered into two groups (green, red). The green constellations represent quite extended formations (so the distances between the players are quite long), while the succeeding red constellations represent narrow formations. This is a typical defense behavior of a handball team, which is initiated after the ball possession is lost and makes it harder to score for the other team.




Figure 16: Retreat pattern of the whole team.

4.3 Experiment with gull data set

For the last experiment we have chosen a GPS-dataset. By this, we want to show that our approach also works on data, which has a significantly lower accuracy and sampling rate. There, trajectories of eight gulls are recorded with a mean gap of 10 minutes between two tracking points. The observation time of about two month leads to 4500 tracking points per gull. The accuracy is determined by standard GPS-devices, which have been attached to the animals. The gulls start each of their trips from their nests, which are situated very close to each other (red circle). These trips cover an area of about 100km x 100km, which is much larger than those in the examples before (Figure 17). The goal of the analysis is to investigate, if the gulls perform similar movements more often.




Figure 17: The trips of the eight gulls cover an area of about 100km x 100km. The red circle points at their nests.

This time we have to adjust the parameters to the low sampling rate and accuracy. We have chosen a frequency of six measurements per hour for recording and have used a k-means clustering with 500 clusters to calculate the similarities.




Figure 18: Two movement patterns of the eight gulls, which occur twice in the data set.

Figure 18 shows two extracted movement patterns, each of which has been observed twice. In some constellations it looks as if less than eight gulls have been observed, however, this is due to the fact, that some members of the group are very close to each other. Both patterns contain sequences, which last one hour and directly follow each other. So, it is quite impressive, how similar the movements are. As we have not included any context information and we are no domain experts concerning gull behavior, we are not able to interpret these movements.

5. Conclusion and Outlook

In this paper, a novel approach has been presented that allows an extraction of group movement patterns from trajectory data. The proposed approach is based on constellations (which describe group movements relatively) and their similarities (which handle uncertainties and inaccuracies). Typical constellations occurring more often in sequences are considered as movement pattern. The algorithms that are used to determine the similarities and to mine for patterns in the constellations sequences are applied to datasets containing the movements of human beings and animals acquired by different kinds of tracking devices. We have showed the results of three experiments, which stem from different applications areas and have different analysis goals. This shows that the approach can easily be transferred to different application areas.

Most of the results contain interesting patterns, which are easy to understand. But sometimes the feature of the pattern may be hard to recognize at the first glance, because it is produced by small repetitive movements of single objects, which, nevertheless, may be not less important. In other cases, like in our last experiment, the pattern may be quite obvious, but it is hard to interpret without any further information or being a domain expert. However, if the patterns are enriched by context information or the knowledge of a domain expert, they can provide useful knowledge about the objects' movement behavior: for instance, a football/handball coach might be able to extract the usual opponent's strategies.

Nevertheless, some problems are not solved in this approach. The most obvious one, which has already been mentioned, is the determination of the clustering parameters. At the moment we use the experience we collected during the experiments. A good solution would be an algorithm, possibly based on histograms of distances between the constellations, to determine a priori reasonable clustering parameters or an approach similar to [2]. Of course, in some cases, depending on the scenario, they have to be adjusted manually, to enable different analysis tasks like we have shown in our soccer and handball experiment. Although we have just removed the goal keeper there, it is possible to create a behavior analysis of certain team parts (such as forwards' or defenders' behavior).

Besides that, a future challenge is to avoid the restriction to analyze only given groups of observed objects. It is already possible to manually select subgroups, as we have done in the experiments, but this approach does not analyze possible subgroups automatically. This would be a useful feature when analyzing larger groups of objects without having any knowledge about certain semantically meaningful group divisions. A possibility would be to use existing movement analysis methods to get a coarse interpretation of the scene in terms of object behaving in a similar way; based on this coarse grouping, our method could be applied to analyze the movement behavior within the group. E.g. in the soccer-examples, possibly a flock-approach with conservative parameter setting would be able to distinguish field players and goal keeper.

Furthermore, we plan to make our approach work incrementally that it can be used at runtime. Up to now we are just able to analyze the movements in a post processing step. Doing it in real-time enables further analysis tasks like movement prediction, which surely are possible as post processing step but are not that useful. An incrementally working approach requires incremental algorithms. There already are incremental versions of the DBSCAN [8] or the k-means [18] clustering. However, our sequence mining method has to be adapted to it.

6. Acknowledgments

The seagull data set has been provided within the EU-Cost Action MOVE –projects at the occasion of a data challenge organized at the workshop on ‘*Analysis and Visualization of Moving Objects*’ which was held at the Lorentz Centre in Leiden, The Netherlands, from 27 June to 1 July 2011
<http://www.lorentzcenter.nl/lc/web/2011/453/info.php3?wsid=453&venue=Oort>

7. References

- [1] Abouelhoda, M., Ghanem, M., String Mining in Bioinformatics. In *Scientific Data Mining and Knowledge Discovery* (pp. 207-247). Springer Berlin Heidelberg, 2010.
- [2] Anders, K. H., & Sester, M., Parameter-free cluster detection in spatial databases and its application to typification. *International Archives of Photogrammetry and Remote Sensing*, 33(B4/1; PART 4), 75-83, 2000.
- [3] Andersson, M., Gudmundsson, J., Laube, P., Wolle, T.: Reporting leadership patterns among trajectories. *GeoInformatica* 12(4), 497–528, 2008.
- [4] Benkert, M., Gudmundsson, J., Hübner, F., & Wolle, T., Reporting flock patterns. *Computational Geometry*, 41(3), 111-125, 2008.
- [5] Cao, H., Mamoulis, N., Cheung, D. W., Discovery of periodic patterns in spatiotemporal sequences. *Knowledge and Data Engineering, IEEE Transactions on*, 19(4), 453-467.
- [6] CVBASE 2006 test dataset, URL <http://vision.fe.uni-lj.si/cvbase06/downloads.html>, Accessed November 12, 2008 2006.
- [7] Dodge S., Weibel R., et al.: Towards a taxonomy of movement patterns. *Information Visualization* 7, 3-4, 240–252, 2008.
- [8] Ester M., Kriegel H.-P., Sander J., Xu X.: ‘Density-Connected Sets and their Application for Trend Detection in Spatial Databases’, *Proc. 3rd Int. Conf. on Knowledge Discovery and Data Mining*, AAAI Press, 1997.
- [9] Ester, M., Kriegel, H. P., Sander, J., Wimmer, M., & Xu, X. Incremental clustering for mining in a data warehousing environment. In *VLDB* (Vol. 98, pp. 323-333), 1998.
- [10] Gudmundsson, J., Laube, P., Wolle, T.: Movement Patterns in Spatio-Temporal Data. In: *Encyclopedia of GIS*, pp. 726–732. Springer, Heidelberg, 2008.
- [11] Han, J., Cheng, H., Xin, D., & Yan, X., Frequent pattern mining: current status and future directions. *Data Mining and Knowledge Discovery*, 15(1), 55-86, 2007.
- [12] Han, J., Dong, G., & Yin, Y., Efficient mining of partial periodic patterns in time series database. In *Data Engineering*, 1999. Proceedings., 15th International Conference on (pp. 106-115). IEEE.
- [13] Inaba, M.; Katoh, N.; Imai, H. "Applications of weighted Voronoi diagrams and randomization to variance-based k-clustering". *Proceedings of 10th ACM Symposium on Computational Geometry*. pp. 332–339. doi:10.1145/177424.178042, 1994.
- [14] Laube, P., van Kreveld, M., Imfeld, S.: Finding REMO – detecting relative motion patterns in geospatial lifelines. In: *Proc. 11th Internat Symp. Spatial Data Handling*, pp. 201–214, 2004.
- [15] Laube P., Duckham M. and Wolle T., Decentralized Movement Pattern Detection amongst Mobile Geosensor Nodes // *Geographic Information Science*, 5th International Conference, GIScience 2008, Park City, UT, USA, September 23-26, Proceedings. - Heidelberg : Springer, 2008.
- [16] MacQueen, J. B., "Some Methods for classification and Analysis of Multivariate Observations". *Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability* 1. University of California Press. pp. 281–297. MR 0214227. Zbl 0214.46201. Retrieved 2009-04-07, 1967.
- [17] Mutschler, C., Fraunhofer IIS, Sensordata for ‘The ACM DEBS 2013 Grand Challenge’, <http://www.orgs.ttu.edu/debs2013/index.php?goto=cfchallengedetails>, 2013.
- [18] Pham, D. T., S. S. Dimov, and C. D. Nguyen. "An incremental K-means algorithm." *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science* 218.7 (2004): 783-795.
- [19] Shamoun-Baranes J., E.E. van Loon, R. Purves, B. Speckmann, D. Weiskopf, K. Camphuysen. 2011. Analysis and visualisation of animal movement. *Biology Letters*