Finding REMO - Detecting Relative Motion Patterns in Geospatial Lifelines

Patrick Laube[1], Marc van Kreveld[2], Stephan Imfeld[1]

Department of Geography, University of Zurich, Winterthurerstrasse 190, CH–8057 Zurich, Switzerland, {plaube,imfeld}@geo.unizh.ch[1]
Department of Computer Science, Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, The Netherlands, marc@cs.uu.nl[2]

Abstract

Technological advances in position aware devices increase the availability of tracking data of everyday objects such as animals, vehicles, people or football players. We propose a geographic data mining approach to detect generic aggregation patterns such as flocking behaviour and convergence in geospatial lifeline data. Our approach considers the object's motion properties in an analytical space as well as spatial constraints of the object's lifelines in geographic space. We discuss the geometric properties of the formalised patterns with respect to their efficient computation.
Keywords: Convergence, cluster detection, motion, moving point objects, pattern matching, proximity

1 Introduction

Moving Point Objects (MPOs) are a frequent representation for a wide and diverse range of phenomena: for example animals in habitat and migration studies (e.g. Ganskopp 2001; Sibbald et al. 2001), vehicles in fleet management (e.g. Miller and Wu 2000), agents simulating people for modelling crowd behaviour (e.g. Batty et al. 2003) and even tracked soccer players on a football pitch (e.g. Iwase and Saito 2002). All those MPOs share motions that can be represented as *geospatial lifelines*: a series of observations consisting of a triple of *id*, *location* and *time* (Hornsby and Egenhofer 2002).

Gathering tracking data of individuals became much easier because of substantial technological advances in position aware devices such as GPS receivers, navigation systems and mobile phones. The increasing number of such devices will lead to a wealth of data on space-time trajectories documenting the space-time behaviour of animals, vehicles and people for off-line analysis. These collections of geospatial lifelines present a rich environment to analyse individual behaviour. (Geographic) data mining may detect patterns and rules to gather basic knowledge of dynamic processes or to design location based services (LBS) to simplify individual mobility (Mountain and Raper 2001; Smyth 2001; Miller 2003).

Knowledge discovery in databases (KDD) and data mining are responses to the huge data volumes in operational and scientific databases. Where traditional analytical and query techniques fail, data mining attempts to distill data into information and KDD turns information into knowledge about the monitored world. The central belief in KDD is that information is hidden in very large databases in the form of interesting patterns (Miller and Han 2001). This statement is true for the spatio-temporal analysis of geospatial lifelines and thus is a key motivator for the presented research. Motion patterns help to answer the following type of questions.

- Can we identify an alpha animal in the tracking data of GPS-collared wolves?
- How can we quantify evidence of 'swarm intelligence' in gigabytes of log-files from agent-based models?
- How can we identify which football team played the more catching lines of defense in the lifelines of 22 players sampled at seconds?

The long tradition of data mining in the spatio-temporal domain is well documented (for an overview see Roddick et al. (2001)). The Geographic Information Science (GISc) community has recognized the potential of Geographic Information Systems (GIS) to 'capture, represent, analyse and explore spatio-temporal data, potentially leading to unexpected new knowledge about interactions between people, technologies and urban infrastructures (Miller 2003). Unfortunately, most commercial GIS are based on a static place-based perspective and are still notoriously weak in providing tools for handling the temporal dimensions of geographic information (Mark 2003). Miller postulates expanding GIS from the place-based perspective to encompass a people-based perspective. He identifies the development of a formal representation theory for dynamic spatial objects and of new spatio-temporal data mining and exploratory visualization techniques as key research issues for GISc.

In this paper work is presented which extends a concept developed to analyse relative motion patterns for groups of MPOs (Laube and Imfeld 2002) to also analyse the object's absolute locations. The work allows the identification of generic formalised motion patterns in tracking data and the extraction of instances of these formalised patterns. The significance of these patterns is discussed.

2 Aggregation in Space and Time

Following Waldo Tobler's first law of geography, near things are more related than distant things (Tobler 1970). Tolber's law is often referred to as being the core of spatial autocorrelation (Miller 2004). Nearness as a concept can be extended to include both space and time. Thus analysing geospatial lifelines we are interested in objects near in space-time. Objects that are near at certain times might be related. Although correlation is not causality, it provides evidence of causality that can (and should) be assessed in the light of theory and/or other evidence. Since this paper focuses on the formal and geometrical definition and the algorithmic detection of motion patterns we use geometric *proximity* in euclidian space to avoid the vague term nearness.

To analyse geospatial lifelines this could mean that MPOs moving within a certain range influence each other. E.g. an alpha wolf leads its pack by being seen or heard, thus all wolves have to be located within the range of vision or earshot respectively. Analysing geospatial lifelines we are interested in first identifying motion patterns of individuals moving in proximity. Second we want to know how, when and where sets of MPOs aggregate, converge and build clusters respectively.

Investigating aggregation of point data in space and time is not new. Most approaches focus on detecting localized clustering at certain time slices (e.g. Openshaw 1994; Openshaw et al. 1999). This concept of spatial clusters is static, rooted in the time sliced static map representation of the world. With a true spatio-temporal view of the world aggregation must be perceived as the momentary process convergence and the final static cluster as its possible result. The opposite of convergence, divergence, is equally interesting. Its possible result, some form of dispersal, is much less obvious and thus much harder to perceive and describe.

A cluster is not the compulsory outcome of a convergence process and vice versa. A set of MPOs can very well be converging for a long time without building a cluster. The 22 players of a football match may converge during an attack without ever forming a detectable cluster on the pitch. In reverse, MPOs moving around in a circle may build a wonderful cluster but never be converging. In addition the process of convergence and the final cluster are in many cases sequential. Consider the lifelines of a swarm of bees. At sunset the bees move back to the hive from the surrounding meadows, showing a strong convergence pattern without building a spatial cluster. In the hive the bees wiggle around in a very dense cluster, but do not converge anymore. In short, even though convergence and clustering are often spatially and/or temporally tied up, there need not be a detectable relation in an individual data frame under investigation.

3 The Basic REMO–Analysis Concept

The basic idea of the analysis concept is to compare the motion attributes of point objects over space and time, and thus to *relate* one object's motion to the motion of all others (Laube and Imfeld 2002). Suitable geospatial lifeline data consist of a set of MPOs each featuring a list of fixes. The REMO concept (RElative MOtion) is based on two key features: First, a transformation of the lifeline data to a REMO matrix featuring motion attributes (i.e. speed, change of speed or motion azimuth). Second, matching of formalized patterns on the matrix (Fig. 1).

Two simple examples illustrate the above definitions: Let the geospatial lifelines in Fig. 1a be the tracks of four GPS-collared deer. The deer O_1 moving with a constant motion azimuth of 45° during an interval t_2 to t_5, i.e. four discrete time steps of length ∂t, is showing *constance*. In contrast, four deer performing a motion azimuth of 45° at the same time t_4 show *concurrence*.

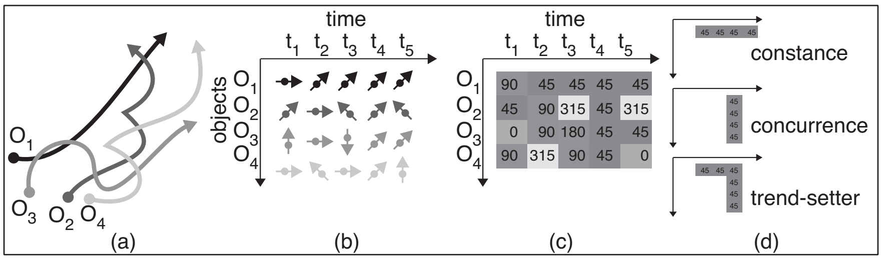

Fig. 1. The geospatial lifelines of four MPOs (a) are used to derive in regular intervals the motion azimuth (b). In the REMO analysis matrix (c) generic motion patterns are matched (d).

The REMO concept allows construction of a wide variety of motion patterns. See the following three basic examples:

- *Constance:* Sequence of equal motion attributes for *r* consecutive time steps (e.g. deer O_1 with motion azimuth 45° from t_2 to t_5).
- *Concurrence:* Incident of *n* MPOs showing the same motion attributes value at time *t* (e.g. deer O_1, O_2, O_3, and O_4 with motion azimuth 45° at t_4)
- *Trend-setter:* One trend-setting MPO anticipates the motion of *n* others. Thus, a trend-setter consists of a *constance* linked to a *concurrence* (e.g. deer O_1 anticipates at t_2 the motion azimuth 45° that is reproduced by all other MPOs at time t_4)

For simplicity we focus in the remainder of this paper on the motion attribute azimuth, even though most facets of the REMO concept are equally valid for speed or change of speed.

4 Spatially Constrained REMO patterns

The construction of the REMO-matrix is an essential reduction of the information space. However it should be noted that this step factors out the absolute locations of the fixes of the MPOs. The following two examples illustrate generic motion patterns where the absolute locations must be considered.

- Three geese all heading north-west at the same time – one over London, one over Cardiff and one over Glasgow are unlikely to be influenced by each other. In contrast, three geese moving north-west in the same gaggle are probably influenced. Thus, for flocking behaviours the spatial proximity of the MPOs has to be considered.
- Three geese all heading for Leicester at the same time – one starting over London, one over Cardiff and one over Glasgow show three different motion azimuths, not building any pattern in the REMO matrix. Thus, convergence can only be detected considering the absolute locations of the MPOs.

The basic REMO concept must be extended to detect such *spatially constrained* REMO patterns. In Section 4.1 spatial proximity is integrated and in Section 4.2 an approach is presented to detect convergence in relative motion. Section 4.3 evaluates algorithmic issues of the proposed approaches.

4.1 Relative Motion with Spatial Proximity

Many sheep moving in a similar way is not enough to define a flocking pattern. We expect additionally that all the sheep of a flock graze on the same hillside. Formalised as a generic motion pattern we expect for a flocking the MPOs to be in spatial proximity. To test the proximity of *m* MPOs building a pattern at a certain time we can compute the spatial proximity of the *m* MPO's fixes in that time frame. Following Tobler's first law, proximity among MPOs can be considered as impact ranges, or the other way around: a spatio-temporally clustered set of MPOs is evidence to suggest an interrelation among the involved MPOs.

The meaning of spatial constraint in a motion pattern is different if we consider the geospatial lifeline of a single MPO. The consecutive observations (fixes) of a single sheep building a lifeline can be tested for proximity. Thus the proximity measure constrains the spatial extent of single object's motion pattern. A constance for a GPS-collared sheep may only be meaningful if it spans a certain distance, excluding pseudo-motion caused by inaccurate fix measurements.

Different geometrical and topological measures could be used to constrain motion patterns spatially. The REMO analysis concept focuses on the following open list of geometric proximity measures.

- A first geometric constraint is the maximal length of the cumulated distances to the mean or median center (length of star plot).
- Another approach to indicate the spatial proximity of points uses the Delaunay diagram, applied for cluster detection in 2-D point sets (e.g. Estivill-Castro and Lee 2002) or for the visualisation of habitat-use intensity of animals (e.g. Casaer et al. 1999). According to the cluster detection approach two points belong to the same cluster, if they are connected by a small enough Delaunay edge. Thus, adapted to the REMO concept a second distance proximity measure is to limit the average length of the Delauney edges of a point group forming a REMO pattern.
- Proximity measures can have the form of bounding boxes, circles or ellipses (Fig. 2). The simplest way of indicating an impact range would be to specify a maximal bounding box that enclosed all fixes relevant to the pattern. Circular criteria can require enclosing all relevant fixes within radius *r* or include the constraint to be spanned around the mean or median center of the fixes. Ellipses are used to rule the directional elongation of the point cluster (major axis *a*, minor axis *b*).
- Another areal proximity measure for a set of fixes is the indication of a maximal border length of the convex hull.

Using these spatial constraints the list of basic motion patterns introduced in Section 3 can be amended with the spatially constrained REMO patterns (Fig. 2).

- *Track:* Consists of the REMO pattern constance and the attachment of spatial constraint. Definition: *constance* + spatial constraint *S*.
- *Flock:* Consists of the REMO pattern concurrence and the attachment of a spatial constraint. Definition: *concurrence* + spatial constraint *S*.
- *Leadership:* Consists of the REMO pattern trend-setter and the attachment of a spatial constraint. For example the followers must lie within the range *(∂x, ∂y)* when they join the motion of the trend-setter. Definition: *trend-setter* + spatial constraint *S*.

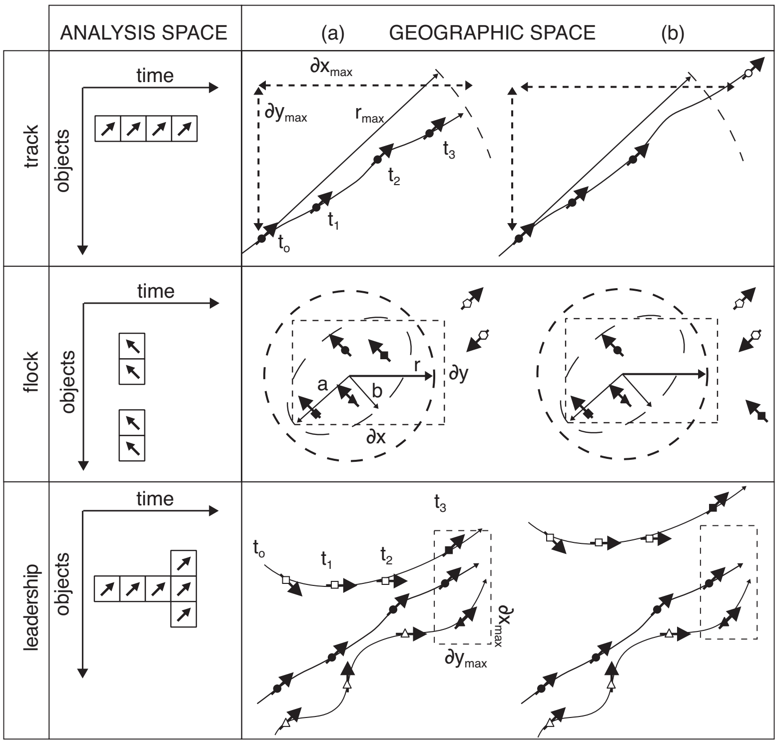

Fig. 2. The figure illustrates the constraints of the patterns track, flock and leadership in the analysis space (the REMO matrix) and in the geographic space. Fixes matched in the analysis space are represented as solid forms, fixes not matched as empty forms. Some possible spatial constraints are represented as ranges with dashed lines. Whereas in the situations (a) the spatial constraints for the absolute positions of the fixes are fulfilled they are not in the situations (b): For track the last fix lies beyond the range, for flock and leadership the quadratic object lies outside the range.

4.2 Convergence

At the same time self-evident and fascinating are groups of MPOs aggregating and disaggregating in space and time. An example is wild animals suddenly heading in a synchronised fashion for a mating place. Wildlife biologists could be interested in the *who*, *when* and *where* of this motion pattern. Who is joining this spatio-temporal trend? Who is not? When does the process start, when does it end? Where lies the mating place, what spatial extent or form does it have? A second example comes from the analysis of crowd behaviour. Can we identify points of interest attracting people only at certain times, events of interest rather than points of interest, losing their attractiveness after a while? To answer such questions we propose the spatial REMO pattern *convergence*.

The phenomenon aggregation has a spatial and a spatio-temporal form. An example may help to illustrate the difference. Let A be a set of n antelopes. A wildlife biologist may be interested in identifying sets of antelopes heading for some location at certain time. The time would indicate the beginning of the mating season, the selected set of m MPOs the ready-to-mate individuals, and the spot might be the mating area. This is a *convergence* pattern. It is primarily spatial, that means the MPOs head for an area but may reach it at different times. On the other hand the wildlife biologist and the antelopes may share the vital interest to identify MPOs that head for some location and actually meet there at some time extrapolating their current motion. Thus, the pattern *encounter* includes considerations about speed, excluding MPOs heading for a meeting range but not arriving there at a particular time with the others.

- *Convergence:* Heading for R. Set of m MPOs at interval i with motion azimuth vectors intersecting within a range R of radius r.
- *Encounter:* Extrapolated meeting within R. Set of m MPOs at interval i with motion azimuth vectors intersecting within a range R of radius r and actually meeting within R extrapolating the current motion.

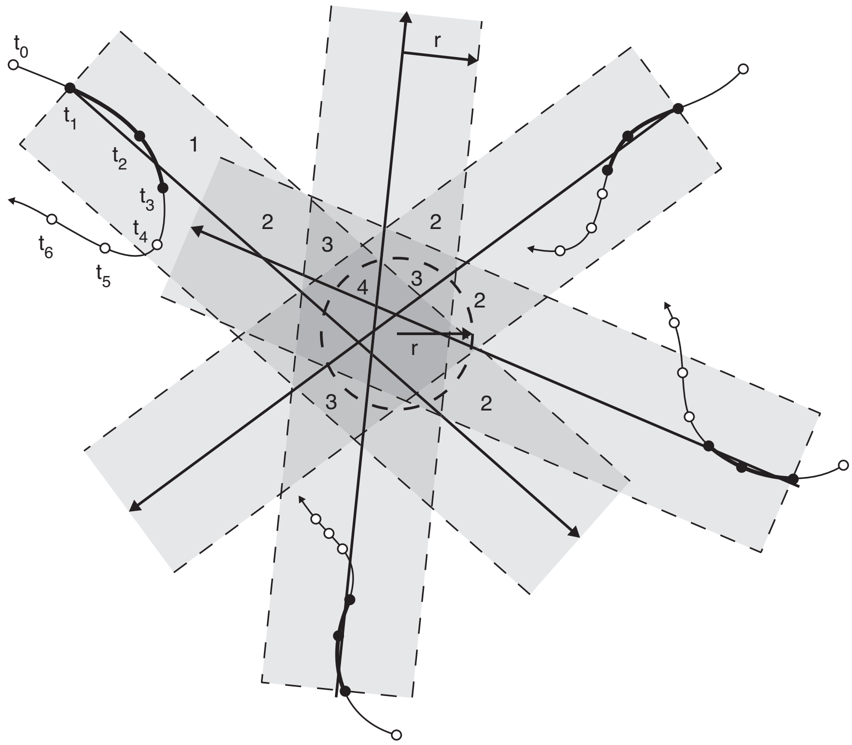

Fig. 3. Geometric detection of convergence. Let *S* be a set of 4 MPOs with 7 fixes from t_0 to t_6. The illustration shows a convergence pattern found with the parameters *4 MPOs* at the temporal interval t_1 *to* t_3. The darkest polygon denotes an area where all 4 direction vectors are passing at a distance closer than *r*. The pattern convergence is found if such a polygon exists. Please note that the MPOs do not build a cluster but nevertheless show a convergence pattern.

The convergence pattern is illustrated in Figure 3. Let *S* be a set of MPOs with *n* fixes from t_0 to t_{n-1}. For every MPO and for every interval of length *i* an azimuth vector fitting in its fixes within *i* represents the current motion. The azimuth vector can be seen as a half-line projected in the direction of motion. The convergence is matched if there is at any time a circle of radius *r* that intersects *n* directed half-lines fitted for each MPO in the fixes within *i*. For the encounter pattern whether the objects actually meet in future must additionally be tested.

The opposites of the above described patterns are termed *divergence* and *breakup*. The latter term integrates a spatial divergence pattern with the temporal constraint of a precedent meeting in a range *R*. The graphical representation of the divergence pattern is highly similar to Fig. 3. The only difference lies in the construction of the strips, heading backwards instead of forwards, relative to the direction of motion.

4.3 Algorithms and Implementation Issues

In this section we develop algorithms to detect the above introduced motion patterns and analyse their efficiency.

The basic motion patterns in the REMO concept are relatively easy to determine in linear time. The addition of positions requires more complex techniques to obtain efficient algorithms. We analyse the efficiency of pattern discovery for track, flock, leadership, convergence, and encounter in this section. We let the range be a circle of given radius R. Let n denote the number of MPOs in the data set, and t the number of time steps. The total input size is proportional to nt, so a linear time algorithm requires $O(nt)$ time. We let m denote the number of MPOs that must be involved in a pattern to make it interesting. Finally, we assume that the length of a time interval is fixed and given.

The addition of geographic position to the REMO framework requires the addition of geographic tests or the use of geometric algorithms. The *track* pattern can simply be tested by checking each basic constance pattern found for each MPO. If the constance pattern also satisfies the range condition, a track pattern is found. The test takes constant additional time per pattern, and hence the detection of track patterns takes $O(nt)$ time overall.

Efficient detection of the *flock* pattern is more challenging. We first separate the input data by equal time and equal motion direction, so that we get a set of $n' \leq n$ points with the same direction and at the same time. The $8t$ point sets in which patterns are sought have total size $O(nt)$. To discover whether a subset of size at least m of the n points lie close together, within a circle of radius R, we use higher-order Voronoi diagrams. The m-th order Voronoi diagram is the subdivision of the plane into cells, such that for any point inside a cell, some subset of m points are the closest among all the points. The number of cells is $O(m(n'-m))$ (Aurenhammer 1991), and the smallest enclosing circle of each subset of m points can be determined in $O(m)$ time (de Berg et al. 2000, Sect. 4.7). If the smallest enclosing circle has radius at most R, we have discovered a pattern. The sum of the n' values over all $8t$ point sets is $O(nt)$, so the total time needed to detect these patterns is $O(ntm^2 + nt \log n)$. This includes the time to compute the m-th order Voronoi diagram (Ramos 1999).

Leadership pattern detection can be seen as an extension of flock pattern detection. The additional condition is that one of the MPOs shows constance over the previous time steps. Leadership detection also requires $O(ntm^2+nt \log n)$ time.

For the *convergence* pattern, consider a particular time interval. The n MPOs give rise to n azimuth vectors, which we can see as directed half-lines. To test whether at least m MPOs out of n converge, we compute the arrangement formed by the thickened half-lines, which are half-strips of width $2r$. For every cell in the arrangement we determine how many thickened half-lines contribute, which can be done by traversing the arrangement once and maintaining a counter that shows in how many half-strips the current cell is. If a cell is contained in at least m half-strips, it constitutes a pattern. Computing the arrangement of n half-strips and setting the counters can be done in $O(n^2)$ time in total; the algorithm is very similar to computing levels in arrangements (de Berg et al. 2000, Chap. 8). Since we consider t different time intervals, the total running time becomes $O(n^2t)$.

The *encounter* pattern is the most complex one to compute. The reason is that extrapolated meeting times must also match, which adds a dimension to the space in which geometric algorithms are needed. We lift the problem into 3-D space, where the third dimension is time. The MPOs become half-lines that go upward from a common horizontal plane representing the beginning of the time interval; the slope of the half-lines will now be the speed. The geometric problem to be solved is finding horizontal circles of radius R that are crossed by at least m half-lines, which can be solved in $O(n^4)$ time with a simple algorithm. For all time intervals of a given length, the algorithm needs $O(n^4t)$ time.

5 Discussion

The REMO approach has been designed to analyse motion basing on geospatial lifelines. Since motion is expressed by a change in location over time the REMO patterns intrinsically span over space and time. Our approach thus overcomes the limitation of only either detecting spatial clusters on snapshots or highlighting temporal trends in attributes of spatial units. It allows pattern detection in space-time.

REMO patterns rely solely on point observations and are thus expressible for any objects that can be represented as points and leave a track in a euclidean space. Having translated the expected behaviours into REMO patterns, the detection process runs unsupervised, listing every pattern occurrence. The introduced patterns can be detected within reasonable time. Many simple patterns can be detected in close to linear time if the size of the subset m that constitutes a pattern is a constant, which is natural in many situations. The encounter pattern is quite expensive to compute, but since we focus on off-line analysis, we can still deal with data sets consisting of several hundreds of MPOs. Note that the dependency on the number of time steps is always linear for fixed length time intervals. The most promising way to obtain more efficient algorithms is by using *approximation algorithms*, which can save orders of magnitude by stating the computational problem slightly less firm (Bern and Eppstein 1997). In short, the REMO concept can cope with the emerging data volumes of tracking data.

Syntactic pattern recognition adopts a hierarchical perspective where complex patterns are viewed as being composed of simple primitives and grammatical rules (Jain et al. 2000). Sections 3 and 4 introduced a subset of possible pattern primitives of the REMO analysis concept. Using the primitives and a pattern description formalism almost arbitrary motion patterns can be described and detected. Due to this hierarchical design the concept easily adapts to the special requirements of various application fields. Thus, the approach is flexible and universal, suited for various lifelines such as of animals, vehicles, people, agents or even soccer players. The detection of patterns of higher complexity requires more sophisticated and flexible pattern matching algorithms than the ones available today.

The potential users of the REMO method know the phenomenon they investigate and the data describing it. Hence, in contrast to traditional data mining assuming no prior knowledge, the users come up with expectations about conceivable motion patterns and are able to assess the results of the pattern matching process. Therein lies a downside of the REMO pattern detection approach: It requires relatively sophisticated knowledge about the patterns to be searched for. For instance, the setting of an appropriate impact range for a flock pattern is highly dependent on the investigated process and thus dependent on the user. In general the parametrisation of the spatial constraints influences the number of patterns detected. Further research is needed to see whether autocalibration of pattern detection will be possible within the REMO concept.

Even though the REMO analysis concept assumes users specifying patterns they are interested in, the pattern extent can also be viewed as an analysis parameter of the data mining approach. One reason to do so is to detect scale effects lurking in different granularities of geospatial lifeline data. The number of matched patterns may be highly dependent on the spatial, temporal and attributal granularity of the pattern matching process. For example the classification of motion azimuth in only the two classes *east* and *west* reveals a lot of presumably meaningless constance patterns. In contrast, the probability of finding constance patterns with 360 azimuth classes is much smaller, or take the selection of the impact range *r* for the flock pattern in sheep as another example. By testing the length of the impact range *r* against the amount of matched patterns one could search for a critical maximal impact range within a flock of sheep. Future research will address numerical experiments with various data to investigate such relations.

A critical issue in detecting convergence is fitting the direction vector in a set of fixes. Only slight changes in its azimuth may have huge effects on the overlapping regions. A straightforward solution approach to this problem is to smooth the lifelines and then fitting the azimuth vector to a segment of the smoothed lifeline.

The paper illustrates the REMO concept referring to ideal geospatial lifeline data. In reality lifeline data are often imprecise and uncertain. Sudden gaps in lifelines, irregular fixing intervals or positional uncertainty of fixes requires sophisticated interpolation and uncertainty considerations on the implementation side (e.g. Pfoser and Jensen 1999).

6 Conclusions

With the technology driven shift from the static map view of the world to a dynamic process in GIScience, cluster detection on snapshots is insufficient. What we need are new methods that can detect convergence processes as well as static clusters, especially if these two aspects of space-time aggregation are separated. We propose a generic, understandable and extendable approach for data mining in geospatial lifelines. Our approach integrates individuals as well as groups of MPOs. It also integrates parameters describing the motion as well as the footprints of the MPOs in space-time.

Acknowledgements

The ideas to this work prospered in the creative ambience of Dagstuhl Seminar No. 03401 on 'Computational Cartography and Spatial Modelling'. The authors would like to acknowledge invaluable input from Ross Purves, University of Zurich.

References

Aurenhammer F (1991) Voronoi diagrams: A survey of a fundamental geometric data structure. ACM Comput. Surv. 23(3):345-405

Batty M, Desyllas J, Duxbury E (2003) The discrete dynamics of small-scale spatial events: agent-based models of mobility in carnivals and street parades. Int. J. Geographical Information Systems 17(7):673-697

Bern M, Eppstein D (1997) Approximation algorithms for geometric problems. In Hochbaum DS (ed) Approximation Algorithms for NP-Hard Problems, PWS Publishing Company, Boston, MA, pp 296-345

Casaer J, Hermy M, Coppin P, Verhagen R (1999) Analysing space use patterns by thiessen polygon and and traingulated irregular network interpolation: a non-parametric method for processing telemetric animal fixes. Int. J. Geographical Information Systems 13(5):499-511

de Berg M, van Kreveld M, Overmars M, Schwarzkopf O (2000) Computational Geometry - Algorithms and Aplications. Springer, Berlin, 2nd edn

Estivill-Castro V, Lee I (2002) Multi-level clustering and its visualization for exploratory data analysis. GeoInformatica 6(2):123-152

Ganskopp, D. (2001) Manipulating cattle distribution with salt and water in large arid-land pastures: a GPS/GIS assessment. Applied Animal Behaviour Science 73(4):251-262

Hornsby K, Egenhofer M (2002) Modeling moving objects over multiple granularities. Annals of Mathematics and Artificial Intelligence 36(1-2):177-194

Iwase S, Saito H (2002) Tracking soccer player using multiple views. In IAPR Workshop on Machine Vision Applications, MVA Proceedings, pp 102-105

Jain A, Duin R, Mao J (2000) Statistical pattern recognition: A review. IEEE Transactions on Pattern Recognition and Machine Intelligence 22(1):4-37

Laube P, Imfeld S (2002) Analyzing relative motion within groups of trackable moving point objects. In: Egenhofer M, Mark D (eds), Geographic Information Science, Second International Conference, GIScience 2002, Boulder, CO, USA, September 2002, LNCS 2478, Springer, Berlin, pp 132-144

Mark, D (2003) Geographic information science: Defining the field. In: Duckham M, Goodchild M, Worboys M (eds), Foundations of Geographic Information Science, chap. 1, Taylor and Francis, London New York, pp 3-18

Miller H (2003) What about people in geographic information science? Computers, Environment and Urban Systems 27(5):447-453

Miller H (2004) Tobler's first law and spatial analysis. in preparation.

Miller H, Han J (2001) Geographic data mining and knowledge discovery: An overview. In: Miller H, Han J (eds) Geographic data mining and knowledge discovery, Taylor and Francis, London New York, pp 3-32

Miller H, Wu Y (2000) GIS software for measuring space-time accessibility in transportation planning and analysis. GeoInformatica 4(2):141-159

Mountain D, Raper J (2001) Modelling human spatio-temporal behaviour: A challenge for location-based services. Proceedings of GeoComputation, Brisbane, 6

Openshaw S (1994) Two exploratory space-time-attribute pattern analysers relevant to GIS. In: Fotheringham S, Gogerson P (eds) GIS and Spatial Analysis, chap. 5, Taylor and Francis, London New York, pp 83-104

Openshaw S, Turton I, MacGill J (1999) Using geographic analysis machine to analyze limiting long-term illness cencus data. Geographical and Environmental Modelling, 3(1):83-99

Pfoser D, Jensen C (1999) Capturing the uncertainty of moving-object representations. In: Gueting R, Papadias D, Lochowsky, F (eds) Advances in Spatial Databases, 6th International Symposium, SSD'99, Hong Kong, China, July 1999. LNCS 1651, Springer, Berlin Heidelberg, pp 111-131

Ramos E (1999) On range reporting, ray shooting and k-level construction. In: Proc. 15th Annu. ACM Symp. on Computational Geometry, pp 390-399

Roddick J, Hornsby K, Spiliopoulou M (2001) An updated bibilography of temporal, spatial, and spatio-temporal data mining research. In: Roddick J, Hornsby K (eds), Temporal, spatial and spatio-temporal data mining, TSDM 2000, LNAI 2007, Springer, Berlin Heidelberg, pp 147-163

Sibbald AM, Hooper R, Gordon IJ, Cumming S (2001) Using GPS to study the e ect of human disturbance on the behaviour of the red deer stags on a highland estate in Scotland. In: Sibbald A, and Gordon IJ (eds) Tracking Animals with GPS, Macaulay Insitute, pp 39-43

Smyth C (2001) Mining mobile trajectories. In: Miller H, Han J (eds) Geographic data mining and knowledge discovery, Taylor and Francis, London New York, pp 337-361

Tobler W (1970) A computer movie simulating urban growth in the detroit region. Economic Geography 46(2):234-240