Computational Movement Analysis

Joachim Gudmundsson, NICTA, Australia
Patrick Laube, University of Zurich, Switzerland
Thomas Wolle, NICTA, Australia

Recent advances in tracking technologies result in geographic information representing the movement of individuals at previously unseen spatial and temporal granularities. This new, inherently spatiotemporal, kind of geographic information offers new insights into dynamic geographic processes but also challenges the traditionally rather static spatial analysis toolbox. This chapter presents an introductory overview to movement data in general, the theory for modeling and analyzing movement, as well as a set of key application fields of movement analysis. Finally, the chapter addresses privacy concerns relevant to the analysis of human movement.

1 Introduction

Mobility is key in a globalized world, where people, goods, data and even ideas move in increasing volumes at increasing speeds over increasing distances. Moving objects are the building blocks of dynamic geographic processes in natural and built environments. Recent advances in localization (GPS, RFID), wireless communication, mobile computing, and environmental sensing technologies allow for a near ubiquitous tracking of moving individuals in spacetime, resulting in large information volumes capturing dynamic processes of high socio-economic relevance.

Be it mobile and wireless communication [42], transportation [55], video surveillance for security and sports applications [4], or even fundamental behavioral ecology research [49], movement data is accumulated in previously unseen volumes and granularities. However, whereas the technologies for capturing movement are intensely researched and hence reach considerable sophistication, the analytical toolbox to aid enriching the thereby captured data to higher-level exploitable information and process knowledge is much less developed.

This chapter first portrays movement traces as a new and intrinsically spatiotemporal kind of geographic information. Then it discusses the scientific fundamentals required for analyzing such information, introducing techniques from a wide range of contributing science fields. The following part is structured as a grand tour of application fields dealing with the analysis of spatiotemporal movement. After a short introduction of privacy issues inseparable from movement analysis, the chapter concludes with final remarks and an outlook.

2 Movement Traces – A New Kind of Geographic Information

The computational analysis of movement data is for two main reasons a relatively young and little developed research field. First, emerging from static cartography, geographical information systems and its underlying GIScience struggled for a long time with the admittedly substantial challenges of handling dynamics. For many years, occasional changes in a cadastral map were challenging enough, not to mention the constant change of location as is needed for modelling movement. Second, for many years, the tracking of movement entities has been a very cumbersome and costly undertaking. Hence, movement analysis could only address single individuals or very small groups.

Figure 1: Grazing sheep near lake Taupo, NZ. Photo: Patrick Laube.

Only in recent years has the technological advancement in tracking technology reached a level that allowed the fine-grained seamless tracking of sufficient numbers of individuals needed for studying the dynamism of individual moving objects and collectives on a larger scale [39]. With GPS, RFID, and various other outdoor and indoor tracking technologies movement pattern research entered a new era, moving beyond 'thread trailing' and 'mark and recapture' apporaches [24]. Current tracking technology allows for low cost, almost continuous capture of individual trajectories with possibly sub-second sampling rates [91]. Within a few years the situation completely reversed from a notorious data deficit to data overload, with a lack of suited analytical concepts coping with the sudden surge of movement data [71].

Movement traces, capturing the spatiotemporal footprint of the movement of individual moving entities, emerged as a new kind of dynamic geographic information along the established rather static points, lines, and polygons. So far, most applications and related analysis techniques choose to model moving entities as point objects, and their movement as a sequence of time-stamped observation points. However, also more complex spatial entities such as linear or areal objects can move, producing much more complex movement data. For instance, the path and constantly changing shape of a tropical cyclone can be derived from image analysis of multi-temporal remote sensing data. It cannot be stressed enough that movement data is inherently spatiotemporal and hence cannot be represented just by its geometric footprint of a set of static points and a connecting line.

The potential and challenge of providing the new tools and techniques required for analyzing movement traces as a new kind of geographic information has recently attracted the interest of many research fields, both in theory and applications, as is outlined in the next two sections.

3 Scientific Fundamentals of Computational Movement Analysis

Assume that the entities in Figure 1 are sheep on a pasture and that they are observed by a database expert, a computational geometer, and a geographer. Even though all three experts see the very same sheep, they may all perceive totally different things. For the database expert each sheep may represent a leaf in a dynamic tree optimised for fast queries. In contrast, the geographer might want to segment the trajectories into semantically meaningful segments or interpolate a sheep density surface of the pasture. Finally, the computational geometer might build the convex hull of the sheep in order to data mine a flocking movement pattern.

Even though the sheep will not care, their grazing challenges various research fields interested in movement analysis. The following overview bundles the scientific fundamentals underpinning the analysis of movement data, drawn from as diverse research fields as database research, geography and geographic information science, computer science and computational geometry, as well as data mining and knowledge discovery.

3.1 Modeling Movement and Movement Spaces

Modeling movement means modeling the moving entities, but equally modeling the space they move in [61]. The chosen conceptual data model embedding the movement – for instance a 2D or 3D Euclidean space, a network space or a space partition – rules how the entities can move, and consequently impacts on what analysis tools are required to understand that movement. Figure 2 illustrates basic movement spaces and how movement can be modeled therein.

In the most simple case, space is modeled as a 2D Euclidean space (Fig.2(a)). Entities are free to move around and are only limited by potential obstacles, such as, for instance, buildings b_1 and b_2. The space can be modeled using an entity or field conceptual data model. The movement of an entity is then modeled as a set of locations (x, y) that the entity has visited over time (in 3D a location is a triple: (x, y, z)). Even though the actual movement may be a smooth curve, for simplicity reasons, the movement in between known *fixes* is typically modeled as straight line connectors (see e_1). Such a

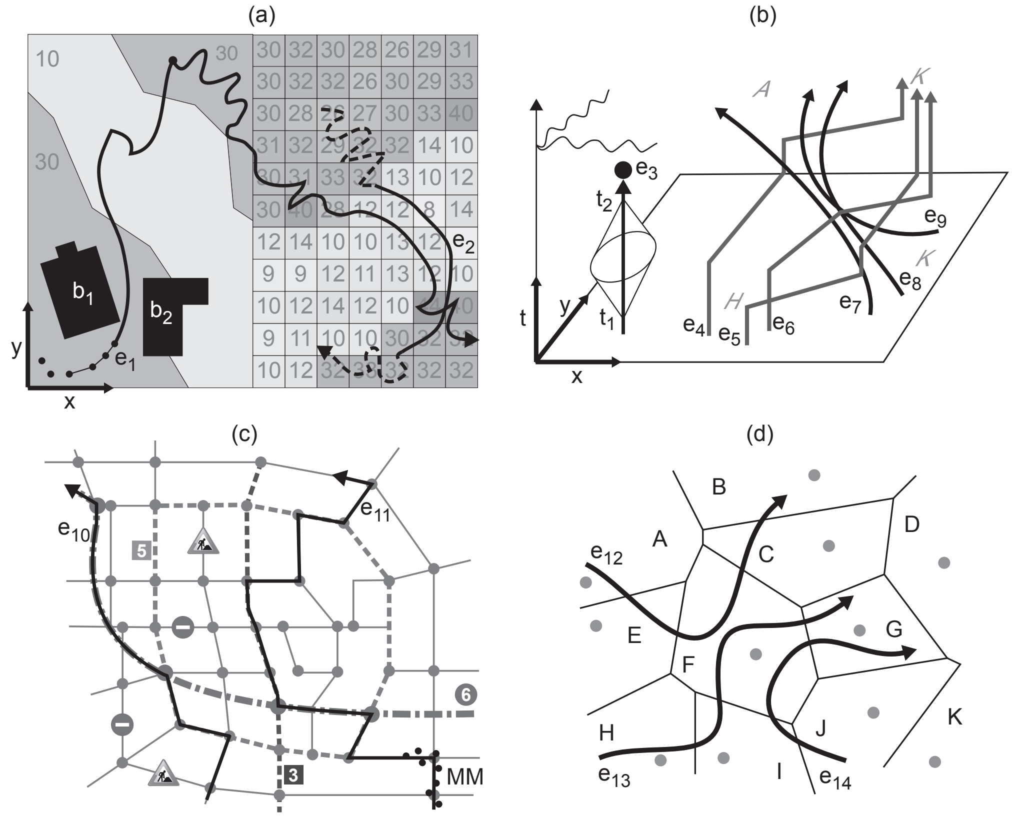

Figure 2: Four basic movement spaces. (a) 2D Euclidean space, (b) 3D Space-time cube, (c) network space, (d) irregular tesselation, e.g. phone tower cells.

movement trace is called a trajectory. Formally, a *trajectory* is defined as a sequence of time-stamped locations $(x, y)_{T_1}, \ldots, (x, y)_{T_t}$, where $T_1, \ldots, T_t$ are t consecutive time steps. Connecting the locations of a trajectory in temporal order yields a polygonal line that can self-intersect [45].

A visually elegant way of incorporating time in the depiction of 2D-movement is the use of a 3D space. Here, the two spatial dimensions x and y are combined with an orthogonal temporal axis t (Fig. 2(b)). Such a space-time cube (sometimes called a space-time aquarium) underpins Hägerstrand's time geography [48]. This concept furthermore allows the modeling of potential mobility, that is where an entity could have been between two known fixes. The space-time prism illustrated for e_3 indicates as a volume where e_3 could have been between t_1 and t_2, given the constraint of a maximal speed [72].

Most human movement is restricted to some form of network. For instance, when tracking the movement of a commuter in an urban area, the movement is most adequately modeled as the progression of the visited edges and nodes of the transit network [83]. Figure 2(c) depicts the movement of entities e_{10} and e_{11}, using a transit network featuring streets, and train lines. Even when the actual movement is captured with a GPS tracking device and deviates from the network, it is semantically given that the movement is bound to the network. Map matching is the preprocessing step matching the inaccurate fixes to the respective edges and nodes [14].

Finally, when tracking a cell phone, the movement trace is only captured in the sequence of identifiers of the cell phone towers the phone was connected to [30]. For example, in Fig. 2(d) entity e_{12} has been connected to the cell phone towers A, E, F, C, and B. Even though this form of movement information does not offer precise point locations, in many application contexts such vague traces are sufficient for the task at hand (e.g. a proximity query in a location-based service).

All four movement spaces can accommodate the movement of individuals, but they are so different that they need different movement analysis tools and techniques. The following subsections present various analysis techniques for a variety of movement forms.

3.2 Indexing Trajectories

In the database community, considerable research has been focussing on spatial and temporal databases. Research in the spatiotemporal area in many ways started with the dissertations by Lorentzos [65] in 1988 and Langran [59] in 1989. Not surprisingly research has mainly focussed on indexing databases so that basic queries concerning the data can be answered efficiently. The most common queries considered in the literature are variants of nearest neighbour queries and range searching queries, such as:

- Spatiotemporal range query, e.g. 'Report all entities that visited region S during the time interval $[t_1, t_2]$.'
- Spatial nearest neighbours given a time interval, e.g. 'Report the entity closest to point p at time t.'
- Temporal nearest neighbours given a spatial region, e.g. 'Report the first entity visiting region S.'

In general one can classify indexing methods used for spatiotemporal data into Parametric Space Indexing methods (PSI) and Native Space Indexing methods (NSI). The PSI method uses the parametric space defined by the movement parameters, and is an efficient approach especially for predictive queries. A typical approach, described by Šaltenis et al. [87] is to represent movement defined by its velocity and projected location along each spatial dimension at a global time reference. The parametric space is then indexed by a new index structure referred to as the TPR-tree (Time Parameterised R-tree). The TPR-tree is a balanced, multi-way tree with the structure of an R-tree. Entries in leaf nodes are pairs of the position of a moving point and a pointer to the moving point, and entries in internal nodes are pairs of a pointer to a subtree and a rectangle that bounds the positions of all moving points or other bounding rectangles in that subtree. The position of a moving point is represented by a reference position and a corresponding velocity vector. To bound a group of d-dimensional moving points, d-dimensional bounding rectangles are used that are also time parameterised, i.e. their coordinates are functions of time. A time-parameterised bounding rectangle bounds all enclosed points or rectangles at all times not earlier than the current time. The search algorithm for a range query also performs computation on the native space by checking the overlap between the range of the query and the trapezoid representation of the node.

The NSI methods represent movement in d dimensions as a sequence of line segments in $d+1$ dimensions, using time as an additional dimension, see for example the work by Hadjieleftheriou et al. [47]. A common approach is to use a multi-dimensional spatial access method like the R-tree. An R-tree would approximate the whole spatiotemporal evolution of an entity with one Minimum Bounding Region (MBR) that tightly encloses all the locations occupied by the entity during its lifetime. An improvement for indexing movement trajectories is to use a multi version index, like the Multi Version R-tree (MVR-tree), also known as a persistent R-tree. This index stores all the past states of the data evolution and allows updates to the most recent state. The MVR-tree divides long-lived entities into smaller intervals by introducing a number of entity copies. A query is directed to the exact state acquired by the structure at the time that the query refers to; hence, the cost of answering the query is proportional to the number of entities that the structure contained at that time.

3.3 Segmenting Trajectories and Trajectory Simplification

An important piece of information in many application fields is the activity the object is performing during the time it is tracked. Can, for example, in a transportation context the mode of transport be inferred from the trajectory? And in behavioral ecology, is it possible to extract an animal's activity by only studying its trajectory? As an example consider seagulls. They can perform several different activities when they are flying. This could involve playing with other seagulls (their flight is then highly sinous with highly variable speed and moving in a small spatial area), flying towards their nest (straight flight with an almost fixed energy-conserving speed) or following a fishing vessel (piecewise straight flights over water with occasional changes in altitude and a steady speed over the time interval of about 7m/s). This example indicates that some information about a seagull's activity can be extracted automatically from the trajectory.

Segmenting a trajectory into relevant segments representing specific activities is clearly highly dependent on the context. It would not be possible to perform the above segmentation without expert domain knowledge and context information. In this section we will only discuss the general version of the problem.

The segmentation problem for a trajectory is to partition it into a number of pieces, which are called segments (subtrajectories). The movement characteristics inside each segment is uniform in some sense. Movement characteristics are for example, speed, heading, altitude or sinousity. The first$^{PL:}$ [Are you sure that this is really the first time, in such an absolute form? What about Dykes and Mountain 2003?] approach to specifically focus on segmenting the trajectory into segments of similar movement characteristics was by Buchin et al. [20]. Their approach defines an attribute function for each movement characteristic that specifies a value at each point on the trajectory (assuming the characteristic attribute can be calculated locally). They then define criteria that specify that within a segment the attribute values at any location within a segment are sufficiently similar. This guarantees a similarity of each incorporated attribute within each segment. They further show that if the characteristic attributes are *monotone* then a simple approach will generate an optimal segmentation in linear time.

A very different trajectory segmentation approach but also with a very different objective was considered by Anagnostopoulos et al. [5], Rasetic et al. [80] and, Yoon and Shahabi [102]. They consider segmenting the trajectory using minimum bounding boxes which leads to simplification of the original objects into smaller, less complex primitives that are better suited for storage and retrieval purposes. These segmentations are specifically effective for index structures that use minimum bounding boxes to store the trajectories. Usually the objective is to find segmentation for all the trajectories such that the cost is minimized, that is, the total number of segments is bounded by some number k. Many different cost functions have been suggested ranging from simply minimizing the volume of the bounding boxes to highly complex combinations of many different attributes, see [5, 80, 102] for more details.

Note that many trajectory simplification algorithms produce a segmentation of the trajectory, usually using a simple distance measure to "shortcut" the trajectory. A common approach is to use an algorithm based on the Douglas-Peucker algorithm [29], see for example the papers by Cao et al. [21], Gudmundsson et al. [44] and, Meratnia and de By [70]. For a real number $\varepsilon > 0$, the polygonal path $P = \langle p_1, \ldots, p_n \rangle$ is simplified by the Douglas-Peucker algorithm as follows. Start with the segment (p_1, p_n). If the distance between the segment (p_1, p_n) and P is at most ε, accept the line segment (p_1, p_n) as an approximation for the whole path. Otherwise, split the path at a point and recursively approximate the two pieces. A crucial aspect of path-simplification algorithms is the tolerance criterion used: When is a line segment accepted as an appropriate approximation for a subpath? Originally in the Douglas-Peucker algorithm, the Euclidean distance between the points on the given subpath and line through the endpoints of the subpath. However, in principle any tolerance criterion could be used.

3.4 Trajectory Similarity

Many fundamental problems in trajectory analysis have one issue in common, namely calculating the similarity between two trajectories. How do we measure the similarity between two curves, or trajectories?[PL:] [Could we here in the intro work out clearly the crucial difference between a curve (not time), and a trajectory (with time). I'd rather have not the two mentioned just like this in a sentence, as if they were the same.]

That is, given two (polygonal) curves P and Q in $\mathcal{R}^d$, one would like to establish a mapping between the two curves that is associated with an appropriate metric. The mapping should match up the two curves in an optimal way (i.e., find a continuous mapping from the points of P to the points of Q, so that the mapping maps the endpoints of one curve to the endpoints of the other curve). However, the choice of distance measure is problematic and crucially dependant on the context. Here we will only consider the most general attempts to define similarity measures.

The perhaps simplest approach to define similarity between two trajectories is to map each sequence into a vector and then use an L_p-distance metric to define the similarity measure. Various early work extended and generalised this distance metric, see for example [1, 23, 79, 100]. The main drawback using an L_p-metric as a distance measure is that it is not robust against noisy data or outliers.

In an attempt to develop more robust measures researchers, considered edit distances, such as Dynamic Time Warping (DTW), which was first used to match signals in speech recognition. Berndt and Clifford [13] were the first to apply the technique to measure the similarity in time-series data. The technique was later refined in several follow-up papers [56, 57, 82, 101].

A similar measure, and also an edit distance, is to find the Longest Common Subsequences (LCSS) between two sequences and then define the distance using the length of the subsequence. The LCSS shows how well two sequences can match one another if one is allowed to stretch them without rearranging (very similar to DTW). Typically one would allow approximate matchings between the values. Several variants have been proposed, see Agrawal et al. [2], Das et al [25] and Vlachos et al. [97]. One drawback with the edit distance measures is that they are not metric, which usually makes them less attractive to use for certain problems, such as clustering.

Buchin et al. [19] claimed that these and many other measures for trajectory similarity are too shape-dependent and do not consider temporal aspects such as speed enough. Sinha and Mark [85] consider the trajectory of a geospatial lifeline which is the set of discrete space-time observations of an individuals residence history. As distance measure they use the average distance between residences weighted by the length of residence. Trajcevski et al. [92] also propose a distance measure for trajectories that takes the temporal aspect into account. They give algorithms for optimal matching under rotations and translations. They consider the maximum distance of geographic locations at equal points in time. These approaches provide a natural time-dependent trajectory similarity measure: the average or maximum distance at corresponding times. Nanni and Pedreschi [73] handled this problem by using the average Euclidean distance between two trajectories as the distance measure. Note that this requires a temporal domain common to the two objects. Van Kreveld and Luo [95] generalised their definition allowing for time shifts and subtrajectory similarity by parameterizing the trajectories over time. This model was further studied in [19] where it was shown how the distance measure can be effectively approximated.

The most fundamental and successful distance measure to this date is probably the Fréchet metric, which is one of the most natural measures of this type. The Fréchet metric between two curves P and Q was first defined by Fréchet [37] in 1906. It requires finding a continuous mapping f between the curves P and Q so that $W(f) = \max_{x \in P} |x, f(x)|$ is minimized. Note that the first location of P has to be mapped to the first location of Q, and the last location of P has to be mapped to the last location of Q. This is also known as the person-dog metric: imagine a person walking on P and a dog walking on Q. The Fréchet distance between P and Q is the length of a shortest leash that enables both the person and the dog to travel along P and Q, respectively, with a leash connecting them (person or dog are not allowed to backtrack; a location on P and its mapped location on Q correspond to the locations of the person and the dog connected by a leash). Alt and Godau [3] introduced the free space diagram of two curves as a data structure to compute the Fréchet distance between them. This structure was later used by Buchin et al. [17] (see also [67]) to extend the concept of the Fréchet distance to not only find a mapping that minimizes the Euclidean distance but in principle to any distance function. In their paper they state several examples, for example when the distance depends on the visibility between the two moving objects or the relative speed between the objects.

Closely related to similarity is clustering, which heavily depends on the distance measure used. In general clustering can be performed more efficiently if the distance measure is a metric, for example, clustering using the Fréchet distance [18] or the average Euclidean distance [73]. While using non-metric measures such as DTW or LCSS usually complicates the clustering step [103].

3.5 Mining for Movement Patterns

This section will focus on mining trajectories for movement patterns, which has mainly been done using algorithmic or data mining approaches.

Some of the most interesting spatiotemporal patterns are periodic patterns. For example, consider a trajectory obtained by tracking an elephant; it is easy to detect which areas are important for the elephant, i.e. where it spends a certain amount of its time. However, ideally one would like to be able to detect if this area is visited with some regularity which might indicate that it could be used as a grazing or mating area during certain times of year. Note however, that the visits may be regular even though the region is not visited every night, every week or every year. It might be a regular event even though it only takes place 50% of the times. Djordjevic et al. [26] gave a generic tool that given a sequence of events with time stamps detects if there is a periodic pattern of the events. They consider several settings, depending on the available information.

Mamoulis et al. [68] considered the special case when the period is given in advance. They partition space into a set of regions which allows them to define a pattern P as a τ-length sequence of the form $r_0, r_1, \ldots, r_{\tau-1}$, where r_i is a spatial region or the special character *, indicating the whole spatial universe. If the entity follows the pattern enough times, the pattern is said to be *frequent*. However, this definition imposes no control over the density of the regions, i.e. if the regions are too large then the pattern may always be frequent. Therefore an additional constraint is added, namely that the points of each subtrajectory should form a cluster inside the spatial region.

Chawla and Verhein [96] defined spatiotemporal association rules (STARs) that describe how entities move between regions over time. They assume that space is partitioned into regions, which may be of any size and shape. The aim is to find interesting regions and rules that predict how entities will move through the regions. A region is interesting when a large number of entities leaves (sink), a large number of entities enters (source) or a large number of entities enters and leaves (thoroughfare).

In 2004 Laube et al. [64] defined a collection of spatiotemporal patterns based on direction of movement and location, e.g. flock, leadership, convergence and encounter, and they gave algorithms to compute them efficiently. As a result there were several subsequent articles studying the discovery of these patterns. Benkert et al. [12] modified the original definition of a flock to be a set of entities moving close together during a time interval, see Fig. 3. Note that in this definition the entities involved in the flock must be the same during the whole time interval. Benkert et al. [12] observed that a flock of m entities moving together during k time steps corresponds to a cluster of size m in $2k$ dimensional space. Thus the problem can be restated as clustering in high dimensional space. To handle high dimensional space one can use well-known dimensionality reduction techniques. There are several decision versions of the problem that have been shown to be NP-hard, for example deciding if there exists a flock of a certain size, or of a certain duration. The special case when the flock is stationary is often called a *meeting* pattern.[PL:] [Do we need a reference here?]

Kalnis et al. [53] define and compute moving clusters where entities might leave and join during the existence of a moving cluster. For each fixed discrete time-step t_i they use standard clustering algorithms to find clusters with a minimum number of entities and a minimum density. Then they compare any cluster c found for t_i with any (moving) cluster c' found for time-step t_{i-1}. If c and c' have enough entities in common, which is formally specified by a threshold value, then c' can be extended by c, which results in a *moving cluster*.[PL:] [Do we want to add how this relates to the last paragraph on clustering in Sec. 3.4?] They propose several ideas to increase the speed of their method, e.g. by avoiding redundant cluster comparisons, or approximating moving clusters instead of giving exact solutions, and they experimentally analyse their performance.

While not discussed in this chapter, additional movement patterns have been studied such as leadership patterns [6], popular places [11], single file [17] and commuting patterns [18].

For the detection of a flock the chosen disk size has a substantial effect on the results of the discovery process. And the selection of a proper disc size can turn out to be difficult as situations can occur where objects that intuitively belong together or do not belong together are not quite within any disk of the given size or are within such a disk. To get around this problem Jeung et al. [50, 51] employ the notion of density connection, which enables the formulation of arbitrary shapes of groups. They define a *convoy* as a group of objects each of which has at least m objects and they are so-called densityconnected. Intuitively, two objects in a group are densityconnected if a sequence of objects exists that connects the two objects and the distance between consecutive objects is small.

Be it exploratory analysis approaches, indexing techniques or data mining algorithms, all effort put in theory ultimately leads to more advanced ways of inferring high level process knowledge from low level tracking data. The following section will illustrate a wide range of fields where such fundamentals underlie various powerful applications.

3.6 Visual and Exploratory Movement Analysis

Given Geography's legacy in cartography, it is not surprising that movement analysis is often addressed by a combination of geovisualisation and exploration. Visual and exploratory movement analysis approaches combine the speed and patience of computers with the excellent capability of humans to detect the expected and discover the

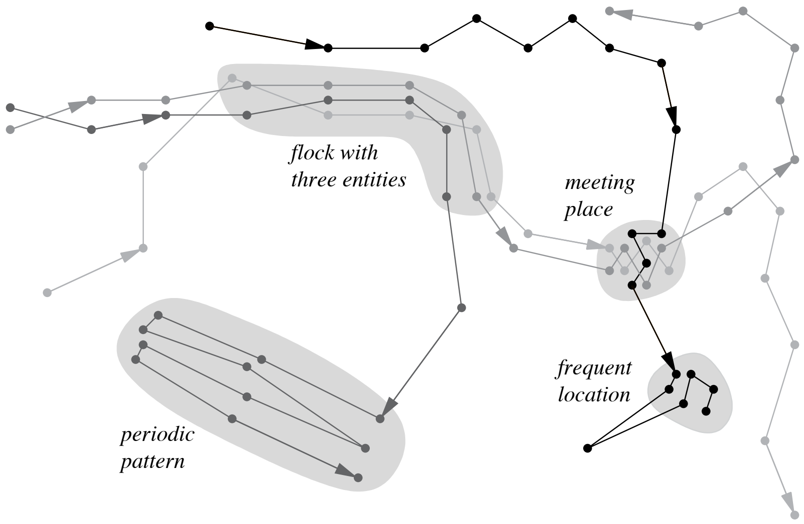

Figure 3: Illustrating the trajectories of four entities moving over 20 time steps. The following patterns are highlighted: a *flock* of three entities over five time-steps, a *periodic pattern* where an entity shows the same spatiotemporal pattern with some periodicity, a *meeting place* where three entities meet for four time steps, and finally, a *frequently visited location* which is a region where a single entity spends a lot of time.

unexpected given an appropriate graphical representation.

Some basic movement patterns become obvious from simple *mapping* of movement trajectories on a two-dimensional map. Trajectories bundled in narrow, directed bottlenecks represent often used corridors. Less focussed trajectory footprints represent more arbitrary movement, such as in grazing animals or visitors at a sports event strolling around a stadium. The application of conventional GIS analysis tools on points and lines representing moving entities has proven to be a very effective approach. For example, GIS tools for generalisation, interpolation and surface generation may be applied to support the analysis of movement data. Brillinger et al. [16] use a regularly sampled vector field to illustrate the overall picture of animals moving in their habitat, with each vector coding in orientation and size for mean azimuth and mean speed at that very location. Dykes and Mountain [36] use a continuous density surface and a 'spotlight' metaphor for the detection of activity patterns. Again, common GIS tools such as algorithms initially designed for the analysis of digital terrain models can easily be adopted, for instance to identify 'peaks' of frequent visitation and 'ridges' of busy corridors [36].

Animation is suited to uncover specific movement behaviours of individuals and groups. Animating moving entities with a constant moving time window in the so-called dynamic view uncovers speed patterns of individuals [9, 36]. Flocking or converging are more complex patterns of coordination in groups. Such group patterns are very striking when animating even large numbers or individuals in a movie-like animation.

The 3D spacetime cube is an often used representation for the visual exploration of movement information. In the specific geometry in such a *three-dimensional space-time aquarium* episodes of immobility and certain speed behaviours produce distinctive patterns of vertical and inclined time lines, respectively. It has to be noted, however, that the concept of lifelines in the spacetime cube are most efficient with small numbers of moving entities. The larger the number of entities populating a spacetime aquarium, the more trajectories overlap and occlude each other, and hence the less clarifying the 3D representation becomes.

Exploration and analysis of typically very large movement data sets requires often more than just visual approaches. Most recently, a growing number of authors suggest the use of visual analytics for the analysis of massive collections of movement data [7]. *Visual analytics* combines database technology, computational data processing and analysis techniques with visualization, facilitating synergetic collaboration between computers and human explorers whereby the strengths of each 'partner' can be utilized [8].

As a further development of existing exploratory data analysis approaches, visual analytics especially involves methods suited for data of multiple types and multiple sources, and even data that is conflicting and incomplete [90]. With respect to movement, visual analytics tools include methods for data aggregation, grouping/dividing movement data, transformations of space and time, trajectory clustering [8]. Such methods are typically linked in an interactive setting, where the results of one step can build the input of a subsequent step [81].

4 Application Fields of Movement Analysis

4.1 Behavioral Ecology

The observation of behavioural patterns is crucial to animal behaviour science. So far, individual and group patterns are rather directly observed than derived from tracking data. However, there are more and more projects that collect animal movement information by equipping them with GPS-GSM collars [91]. In a comprehensive study in ecology, Holyoak *et al.* found nearly 26000 published articles in the last decade referring to the movement of organisms [49]. It is even possible to track the positions of insects, e.g. butterflies or bees. However, most of the times non-GPS based technologies are used that allow for very small and light sensors or transponders.

Analysing movement patterns of animals can help to understand their behaviour in many different aspects. In behavioral ecology, scientist especially try to understand from movement traces how internal and external factors rule movement behavior [74]. Biologists can learn about places that are popular for individual animals, or 'hot spots' that are frequented by many animals. It is possible to investigate social interactions, ultimately revealing the social structure within a group of animals. A major focus lies on the investigation of leading and following behaviour in socially interacting animals, such as in a flock of sheep or a pack of wolves [35]. On a larger scale, animal movement data reflects very well the seasonal or permanent migration behaviour. In the animation industry, software agents implement movement patterns in order to realistically mimic the behaviour of animal groups. Most prominent is the flocking model implemented in NetLogo which mimics the flocking of birds [98].

4.2 Mobility and Transportation

Human mobility in urban spaces can be collected and exploited in several ways. For instance, using mobile phones that communicate with a base station is one way to gather data about the approximate locations of people [42]. A profound understanding about how humans move about in an urban context is key to urban planning, e.g. to plan where to build new roads or where to extend public transport.

The analysis of human urban mobility can furthermore be used to optimise the design of location-based-services (LBS). The services offered to a moving user could not only be dependent on the actual position, but also on the estimated current activity, which may be derived from a detected movement pattern [36]. Once identified as either 'pedestrian', 'cyclist' or 'motorist', a user could be offered tailored services.

Automatic analysis of the flows in an urban movement network are used for traffic management in order to detect undesirable or even dangerous constellations of moving entities, such as traffic jams or aeroplane course conflicts. Traffic management applications may require basic Moving Object Database queries, but also more sophisticated movement patterns involving not just location but also speed, movement direction and other activity parameters [86]. Computational movement analysis if furthermore an important corner stone method underpinning the emerging field of *computational transportation science* (CTS), the computer science foundation of intelligent transportation systems [40]. CTS aims at improving the safety, mobility, and sustainability of the transportation system by taking advantage of information technologies and ubiquitous computing.

4.3 Surveillance and Security

Surveillance and intelligence services might have access to more detailed data sets capturing the movement of people, e.g. coordinates from mobile phones or credit card usage, video surveillance camera footage or maybe even GPS data. Apart from analysing the movement data of a suspect to help prevent further crime, it is an important task to analyse the entire data set to identify suspicious behaviour in the first place. This leads to define 'normal behaviour' and then search the data for any outliers, i.e. entities that do not show normal behavior, which often appears to be challenging [75]. Some specific activities and the corresponding movement patterns of the involved moving entities express predefined signatures that can be automatically detected in spatiotemporal or footage data. One example is that fishing boats in the sea around Australia have to report their location in fixed intervals. This is important for the coast guards in case of an emergency, but the data can also be used to identify illegal fishing in certain areas. Another example is that a car thief is expected to move in a very characteristic and hence detectable way across a surveilled car park. Movement patterns have furthermore attracted huge interests in the field of spatial intelligence and disaster management. Batty et al. [10] investigated local pedestrian movement in the context of disaster evacuation where movement patterns such as congestion or crowding are key safety issues.

4.4 Marketing

In recent development, supermarkets have been equipped with tracking technology. Shopping trolleys and baskets are tracked using Ultra Wide Band technology with the aim to analyse the visibility of product categories, movement of different shopper categories and establish the connection between product penetration (how many customers pass a product), dwell time (amount of time spent in front of the product) and conversion (how many customers buy the product). There are many more questions such as how the layout affects the movement patterns and in the end how this affects what the customers buys.

Until now very little research has been done in this area. Larson et al. [60] analysed a large set of shoppers paths to understand actual travel patterns within a supermarket. The analysis was performed using multivariate clustering algorithms. They observed that the time spent in the store plays an important role, leading to different cluster configurations for short, medium, and long trips. The resulting three sets of clusters identify a total of 14 canonical path types that are typical of grocery store travel. These results dispel a few common myths about shopper travel behavior that common intuition perpetuates, including behavior related to aisles and end-cap displays.

The industry is highly likely to benefit from further research since studies [84] have documented the high incidence of consumers' in-store purchase decision making and the corresponding importance of *point of purchase* (PoP) materials in these purchase decisions. It is estimated that in-store decisions represent as much as 70% of supermarket purchase decisions and 74% of the decisions in mass merchandiser stores.

4.5 Sports Scene Analysis

Advancements in many different areas of technology are also influencing professional sports. For example, some of the major tennis tournaments provide three-dimensional reconstructions of every single point played, tracking the players and the balls. It is furthermore known that, e.g. football coaches routinely analyse match video archives to learn about an opponents behaviours and strategies. Software is already available that automatically provides basic statistical information about the match and the performance of the players. Typical information provided is distance covered by players, top/average speed of players, number of passes performed by players, number of shots on goal, free kick ball speed, heat maps of players, player average position, offside margin and so on.

Attempts have been made to develop more sophisticated analytical tools. Taki and Hasegawa [89, 88] defined the dominant regions where a player has priority over others, and analyzed the areas of distribution of dominant regions. At each instant of a game, each player[TW:]'s dominant region is defined as the region the player can access before any other players. Of course, this requires a knowledge about a player[TW:]~~s~~'s movement capabilities in order to compare their minimum time to access points on the field. Taki and Hasegawa used a very simple model of movement which assumed that a player[TW:]~~s~~'s acceleration is constant. That was later extended by Fujimura and Sugihara [38] and by Kang et al. [54] to include an advanced modelling of human movement.

Very recent development include Memmert and Perl [69] and Grunz et al. [43] proposing to use neural networks to analyze complex game scenarios and interactions, such as group and team tactics. Gudmundsson and Wolle [46] applied trajectory clustering techniques developed by Buchin et al. [18] to the football application to study frequent movements of an individual player and groups of players. This is still an area very much in its infancy and much research is needed to make useful tools widely available.

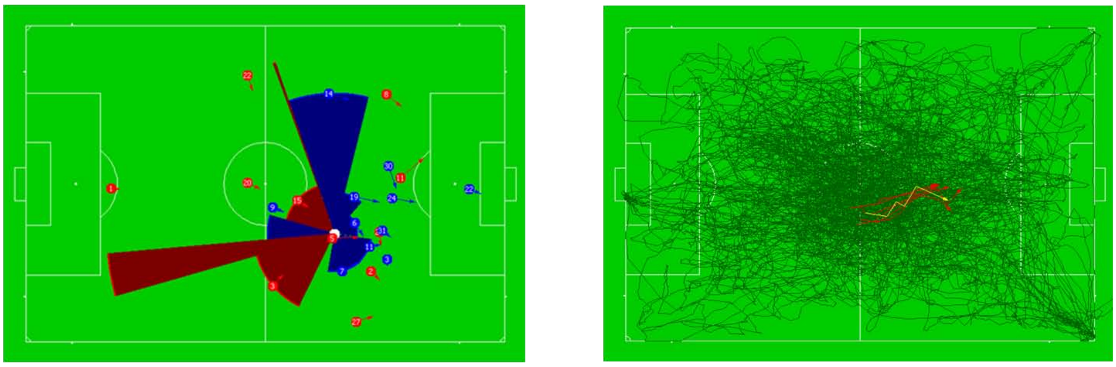

Figure 4: The figure shows screenshots from [46]. On the left the passable regions are visualized with red (to players in the same team as the ball holder) and blue (to players in the opposite team). The picture on the right illustrates the movement of a player during a football game, and a cluster of frequent movements is highlighted in red.

4.6 Movement in Abstract Spaces

In contrast to tracking and analysing the movement of animals and people on the surface of the earth, it is also possible to obtain and analyse spatiotemporal data in abstract spaces also in higher dimensions. Every scatter plot that constantly updates the changes in the x and y values, produces individual trajectories open for movement analysis. Two stock exchange series plotted against each other could build such a dynamic scatter-plot.

For example, basic ideological conflicts can be used to construct abstract ideological spaces. Performing factor analysis on referendum data, researchers hypothesised a structure of mentality consisting of dimensions such as 'political left vs. political right' or 'liberal vs. conservative'. Whole districts or even individuals such as members of parliament could now be localised and re-localised in such ideological space depending on their voting behaviour and its change over time, respectively. The analysis of the districts moving in this abstract movement space can reveal interesting patterns of political change [63].

4.7 Ubiquitous computing and assisted living

Recent technological advancements in mobile computing and wireless communication are resulting in a new generation of integrated spatial systems, where miniaturized computing devices are embedded in everyday objects and environments. Such embedded systems provide data capture and processing services in a wide range of applications, from traffic management to policing and emergency response, from assisted living to health care, from environmental monitoring to smart farming.

For geographic applications, *geosensor networks* (wireless networks of miniaturized computing platforms sensing geographic variables [76]) are a key technology for implanting spatial computing in geographical environments. Not only is the geographic information emerging such embedded systems very dense and complex, but also the distinction between data capture and data processing is increasingly blurred. Instead of processing spatial information in powerful omniscient databases or GISystems, new decentralized information processing approaches are required, where individual computing units cooperatively address tasks without a single node having access to the entire system state [66]. In such systems the spatial computing can happen *somewhere*, and spatial intelligence could be said to become ambient (*Ambient Spatial Intelligence*, AmSI, [31]).

Mobility is a key aspect of many embedded spatial systems. Imagine an assisted living system in health care, where mobile sensors implanted in the shoes of an elder person track the gait pattern and alert a control center in case of a tumble. Or in transportation management, imagine an intelligent transportation system (ITS), where vehicles track the locations of their immediate neighbors and autonomously adjust their speeds in order to avoid congestion. In this wider context, Laube et al. [62] presented an approach for the detection of the movement pattern *flock*. They present algorithms where mobile objects, each of which only possessing partial knowledge of the global system state, collaborate locally and detect the movement pattern in a decentralized way.

5 Privacy

From the spatiotemporal footprint of a person's movement behavior one can potentially infer where that person lives, where the person works, places that person regularly visits, when that person visits those places, and where that person meets other people. Since spatial information systems have the capacity for rapid integration of spatiotemporal information and personal information, privacy issues are becoming an important aspect of geographic information science in general [27], and has outmost relevance in movement analysis.

Fine grained individual movement data is especially privacy sensitive information. Even though as a user of a smart mobile ICT device, I may want to benefit from spatially smart, location-aware services or applications somehow making use of my locational information, I still may want to safeguard my personal privacy. The crux with movement data is, that conventional anonymization does not work. Detailed trajectories typically contain personal points of interest (home, work) and personal movement patterns (daily commuting patterns) that can act as *quasi identifier* and allow for re-identification of anonymized data [15].

There are several strategies for safeguarding location privacy. The different strategies for protecting a mobile individual's location privacy can be categorized into regulatory, policy, anonymity and obfuscation strategies. The following list summarizes an extensive overview of privacy protection strategies to be found in Duckham and Kulik, 2005 [34].

- *Regulatory* approaches to privacy develop rules to govern fair use of personal information. Such regulation will ensure transparency (declaration of who is collecting information), consent and use limitation (individuals consent to information collection for a given limited use), access (to stored information), integrity and security (that collectors must ensure), and accountability (collectors are accountable for misuse) [93]. However strict regulation may be, it does not prevent invasion of privacy [34].

- *Privacy policies* are trust-based mechanisms prohibiting misuse of location information [52]. Whereas regulatory approaches aim at global guarantees, privacy policies shall be flexible to be adapted to individual users or even individual transactions. A privacy rule could, for instance, describe, after what time information about a transaction must be discarded. Just as regulatory approaches, also privacy policies cannot enforce privacy.

- *Anonymity* approaches aim at separating the identity of an individual from its location information. For instance, k-anonymity hides each sensitive release in at least k equally matching individuals [15]. Such additional, similar but fake trajectories are typically randomly generated [58]. Anonymity approaches rely on a trusted anonymity broker and don't work for applications requiring authentication.

- In the context of this chapter, *obfuscation* is probably the most promising approach for safeguarding location privacy. Obfuscation is the process of deliberately degrading the quality of personal location information[TW:] [33]. For many location-based services it will, for example, be sufficient to reveal only the mobile phone cell ID. The actual whereabouts of the user can be hidden. Obfuscation can be based on any combination of inaccuracy, imprecision, and vagueness – these are all concepts spatial information science is very familiar with. For example, for a proximity query ('where is the closest starbucks...'), an individual can obfuscate its location by sending the service provider a list of locations, only one of which is the individual's location [32].

To conclude this short discussion about privacy issues in movement information: Protecting privacy when dealing with movement information is a non-trivial problem [41]. Spatially informed professionals dealing with movement information must be aware of the thin line between opportunities and risks [94]. Movement analysis is not only a technical or algorithmic problem, but also an ethical challenge [77].

6 Conclusion and Outlook

In the mid-1990s, after spatial and geographic information systems had been successfully established as tools handling spatial information, critique arose about the static nature of such systems and important voices claimed that "it's about time" [78] and challenged the GIScience community to "move beyond the snapshot" [22]. Now, one and a half decades later, it can be noted that indeed the static legacy of cartography has been overcome and geographic information systems and respective analysis concepts have in many respects integrated dynamic aspects. The most progress can be found with respect to the analysis of movement information, capturing with the movement of point objects the arguably simplest form of dynamism.

In the last decade, the analysis of movement traces has built up ample momentum as a subfield of geographic information science with high socio-economic relevance. Various science fields contributed to a constantly growing collection of basic methods for handling and analyzing movement data, including inputs from geography, computational geometry, information visualization, database research, as well as data mining and knowledge discovery. The movement analysis toolbox includes concepts for modeling, indexing, and segmenting trajectories, as well as basic analysis tools addressing movement similarity and a range of basic movement patterns. Also in the last decade, movement analysis experienced its own quantitative revolution. Just as other science fields, also geography has moved from a data-poor and computation-poor to a data-rich and computation-rich environment [71]. However, movement analysis was from the beginning to be a data-intensive task and the capacity of the movement analysis toolbox grew with its data sources.

A large portion of the existing analytical tools and techniques developed so far are rather technology-driven, deterministic, and often require some threshold or strict parameter, resulting in a set of discrete and crisp patterns. With initial analytical tools at hand, shedding light on dynamic geographic phenomena in various application fields, we start to apprehend their complexity, and experience the limits of the analytical toolkit developed so far. We have also learned that the simplicity of movement data – just points in a row – bears the danger of a simplicity fallacy. With more and more studies investigating movement, we realize that the closer we look, the less we can really say with certainty.

Technological advances promote the tracking of individuals in many application fields, including ecology, transportation, security, marketing and even sports analysis. However, most application studies focus on pragmatic approaches and produce their own solutions to their very specific problems. Hence, it is difficult to find a common strategy in the community that would help in sharing results, exchanging methods, as well as heading towards what would be an established theory on movement analysis. For example, despite several attempts establishing a categorization of movement patterns [99, 28, 8], an ontology or even just a generally agreed on collection of basic movement patterns is still missing.

A further challenge lies in relating movement to the underlying geography, in order to understand where, when and ultimately why the objects move the way they do. Grazing sheep, for example, may perform a certain movement pattern only when they are on a certain vegetation type. Sea gulls may show certain flight patterns only when close to a salient landscape feature such as a river or a highway. And, the movement patterns expressed by a tracked vehicle will obviously be very dependant on the environment the vehicle is moving in, be it in a car park, in a suburb or on a highway. Whereas most current movement analysis techniques focus on characteristics of and interactions between moving entities (*second order effects*), a true understanding of movement information must involve the variability of the local context embedding the movement (*first order effects*).

All in all, the massive volumes of movement data currently emerging our spatially-enabled built and natural environments, offer fascinating new insights in many dynamic processes of high socio-economic relevance. Geography and computer science have successfully embarked on providing the toolbox required for exploiting movement information. The now required refinement of methods can be achieved through traditional strengths of the contributing disciplines: Expertise in data integration, uncertainty, and scale-effects from Geography, and the handling of large data volumes and efficient algorithm design from computer science.

References

[1] R. Agrawal, C. Faloutsos, and A. N. Swami. Efficient similarity search in sequence databases. In *Proceedings of the 4th International Conference on Foundations of Data Organization and Algorithms*, pages 69–84, 1993.

[2] R. Agrawal, K.-I. Lin, H. S. Sawhney, and K. Shim. Fast similarity search in the presence of noise, scaling, and translation in time-series databases. In *Proceedings of the 21th International Conference on Very Large Data Bases*, pages 490–501, 1995.

[3] H. Alt and M. Godau. Computing the fréchet distance between two polygonal curves. *International Journal on Computational Geometry and Applications*, 5:75–91, 1995.

[4] Amisco. Web resource - verified 10-01, 2011. http://www.sport-universal.com.

[5] A. Anagnostopoulos, M. Vlachos, M. Hadjieleftheriou, E. Keogh, and P. S. Yu. Global distance-based segmentation of trajectories. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 34–43, 2006.

[6] M. Andersson, J. Gudmundsson, P. Laube, and T. Wolle. Reporting leadership patterns among trajectories. In *Proceedings of the 22nd ACM Symposium on Applied Computing.* ACM, 2007.

[7] G. Andrienko, N. Andrienko, U. Demsar, D. Dransch, J. Dykes, S. I. Fabrikant, M. Jern, M. J. Kraak, H. Schumann, and C. Tominski. Space, time and visual analytics. *International Journal of Geographical Information Science*, 24(10):1577–1600, 2010.

[8] N. Andrienko and G. Andrienko. Designing visual analytics methods for massive collections of movement data. *Cartographica*, 42(2):117–138, 2007.

[9] N. V. Andrienko and G. L. Andrienko. Interactive maps for visual data exploration. *International Journal of Geographical Information Science*, 13(4):355–374, 2003.

[10] M. Batty, J. Desyllas, and E. Duxbury. The discrete dynamics of small-scale spatial events: agent-based models of mobility in carnivals and street parades. *International Journal of Geographical Information Science*, 17(7):673–697, 2003.

[11] M. Benkert, B. Djordjevic, J. Gudmundsson, and T. Wolle. Finding popular places. *International Journal on Computational Geometry and Applications*, 20(1):19–42, 2010.

[12] M. Benkert, J.Gudmundsson, F. Hübner, and T. Wolle. Reporting flock patterns. *Computational Geometry - Theory and Applications*, 41(3):111–125, 2008.

[13] D. J. Berndt and J. Clifford. Using dynamic time warping to find patterns in time series. In *Proceedings of the Knowledge Discovery in Database Workshop*, pages 359–370, 1994.

[14] D. Bernstein and A. Kornhauser. An introduction to map matching for personal navigation assistants. Technical report, New Jersey TIDE Center, 1996.

[15] C. Bettini, X. Wang, and S. Jajodia. Protecting privacy against location-based personal identification. In W. Jonker and M. Petkovic, editors, *Secure Data Management*, volume 3674 of *Lecture Notes in Computer Science*, pages 185–199. Springer, Heidelberg, 2005.

[16] D. R. Brillinger, H. K. Preisler, A. A. Ager, and J. G. Kie. An exploratory data analysis (EDA) of the paths of moving animals. *Journal of statistical planning and inference*, 122(2):43–63, 2004.

[17] K. Buchin, M. Buchin, and J. Gudmundsson. Constrained free space diagrams: a tool for trajectory analysis. *International Journal of Geographical Information Science*, 24(7):1101–1125, 2010.

[18] K. Buchin, M. Buchin, J. Gudmundsson, M. Löffler, and J. Luo. Detecting commuting patterns by clustering subtrajectories. *International Journal on Computational Geometry and Applications*, To appear.

[19] K. Buchin, M. Buchin, M. van Kreveld, and J. Luo. Finding long and similar parts of trajectories. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 296–305, 2009.

[20] M. Buchin, A. Driemel, M. J. van Kreveld, and V. Sacristan. An algorithmic framework for segmenting trajectories based on spatio-temporal criteria. In *Proceedings of the 18th ACM SIGSPATIAL International Symposium on Advances in Geographic Information Systems*, pages 202–211, 2010.

[21] H. Cao, O. Wolfson, and G. Trajcevski. Spatio-temporal data reduction with deterministic error bounds. *The VLDB Journal*, 15(3):211–228, 2006.

[22] N. R. Chrisman. Beyond the snapshot: Changing the approach to change, error, and process. In M. J. Egenhofer and R. G. Golledge, editors, *Spatial and Temporal Reasoning in Geographic Information Systems*, pages 85–93. Oxford University Press, Oxford, UK, 1998.

[23] K. Chu and M. Wong. Fast time-series searching with scaling and shifting. In *Proceedings of the 18th ACM Symposium on Principles of Database Systems*, pages 237–248, 1999.

[24] D. L. Claussen, M. S. Finkler, and M. M. Smith. Thread trailing of turtles: methods for evaluating spatial movements and pathway structure. *Canadian Journal of Zoology-Revue Canadienne De Zoologie*, 75(12):2120–2128, 1997.

[25] G. Das, D. Gunopulos, and H. Mannila. Finding similar time series. In *Proceedings of the 1st European Symposium on Principles of Data Mining and Knowledge Discovery*, pages 88–100, 1997.

[26] B. Djordjevic, J. Gudmundsson, A. Pham, and T. Wolle. Detecting regular visit patterns. In *Proceedings of the 16th annual European Symposium on Algorithms*, pages 344–355, 2008.

[27] J. E. Dobson and P. F. Fisher. Geoslavery. *IEEE Technology and Society Magazine*, 22(1):47–52, 2003.

[28] Somayeh Dodge, Robert Weibel, and Anna-Katharina Lautenschutz. Towards a taxonomy of movement patterns. *Inf Visualization*, 7(3-4):240–252, 2008.

[29] D. H. Douglas and T. K. Peucker. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *Canadian Cartographer*, 10(2):112–122, 1973.

[30] C. Du Mouza and P. Rigaux. Mobility patterns. *Geoinformatica*, 9(4):297–319, 2005.

[31] M. Duckham and R. Bennett. Ambient spatial intelligence. In B. Gottfried and H Aghajan, editors, *Behaviour Monitoring and Interpretation - BMI - Smart Environments*, volume 3 of *Ambient Intelligence and Smart Environments*, pages 319–335. IOS Press, Amsterdam, NL, 2009.

[32] M. Duckham and L. Kulik. A formal model of obfuscation and negotiation for location privacy. In H. W. Gellersen, R. Want, and A. Schmidt, editors, *Pervasive Computing, Proceedings*, volume 3468 of *Lecture Notes in Computer Science*, pages 152–170. Springer-Verlag Berlin, Berlin, 2005.

[33] M. Duckham and L. Kulik. Simulation of obfuscation and negotiation for location privacy. In *Spatial Information Theory (COSIT 2005)*, volume 3693 of *Lecture Notes in Computer Science*, pages 31–48. Springer, Heidelberg, 2005.

[34] M. Duckham and L. Kulik. Location privacy and location-aware computing. In J. Drummond, R. Billen, D. Forrest, and E. Joao, editors, *Dynamic and Mobile GIS*. CRC Press, Boca Rator, FL, 2006.

[35] B. Dumont, A. Boissy, C. Achard, A. M. Sibbald, and H. W. Erhard. Consistency of animal order in spontaneous group movements allows the measurement of leadership in a group of grazing heifers. *Applied Animal Behaviour Science*, 95(1–2):55–66, 2005.

[36] J. A. Dykes and D. M. Mountain. Seeking structure in records of spatio-temporal behaviour: visualization issues, efforts and application. *Computational Statistics and Data Analysis*, 43(4):581–603, 2003.

[37] M. Fréchet. Sur quelques points du calcul fonctionnel. *Rendiconti del Circolo Mathematics di Palermo*, 22:1–74, 1906.

[38] A. Fujimura and K. Sugihara. Geometric analysis and quantitative evaluation of sport teamwork. *Systems and Computers in Japan*, 36:49–58, 2005.

[39] A. Galton. Dynamic collectives and their collective dynamics. In A.G. Cohn and D. M. Mark, editors, *Spatial Information Theory, Proceedings*, volume 3693 of *Lecture Notes in Computer Science*, pages 300–315. Springer, Heidelberg, 2005.

[40] G. Geers, M. Sester, S. Winter, and O. E. Wolfson. 10121 report – towards a computational transportation science. In G. Geers, M. Sester, S. Winter, and O. E. Wolfson, editors, *Computational Transportation Science*, number 10121 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany, 2010. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany.

[41] F. Giannotti and D. Pedreschi. Mobility, data mining and privacy: A vision of convergence. In F. Giannotti and D. Pedreschi, editors, *Mobility, Data Mining and Privacy*, pages 1–11. Springer, Berlin Heidelberg, 2008.

[42] M. C. Gonzalez, C. A. Hidalgo, and A. L. Barabasi. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 2008.

[43] A. Grunz, D. Memmert, and J. Perl. Analysis and simulation of actions in games by means of special self-organizing maps. *International Journal of Computer Science in Sport*, 8:22–36, 2009.

[44] J. Gudmundsson, J. Katajainen, D. Merrick, C. Ong, and T. Wolle. Compressing spatio-temporal trajectories. *Computational Geometry – Theory & Applications*, 42:825–841, 2009.

[45] J. Gudmundsson, M. J. van Kreveld, and B. Speckmann. Efficient detection of patterns in 2d trajectories of moving points. *GeoInformatica*, 11(2):195–215, 2007.

[46] J. Gudmundsson and T. Wolle. Towards automated football analysis: Algorithms and data structures. In *Proceedings of the 10th Australasian Conference on Mathematics and Computers in Sport*, 2010.

[47] M. Hadjieleftheriou, G. Kollios, V. J. Tsotras, and D. Gunopulos. Indexing spatio-temporal archives. *The VLDB Journal*, 15(2):143–164, 2006.

[48] T. Hägerstrand. What about people in regional science. *Papers of the Regional Science Association*, 24:7–21, 1970.

[49] M. Holyoak, R. Casagrandi, R. Nathan, E. Revilla, and O. Spiegel. Trends and missing parts in the study of movement ecology. *Proceedings of the National Academy of Sciences*, 105(49):19060–19065, 2008.

[50] H. Jeung, H. T. Shen, and X. Zhou. Convoy queries in spatio-temporal databases. In *Proceedings of the 24th International Conference on Data Engineering*, pages 1457–1459, 2008.

[51] H. Jeung, M. L. Yiu, X. Zhou, C. S. Jensen, and H. T. Shen. Discovery of convoys in trajectory databases. *Proceedings of the VLDB Endowment*, 1:1068–1080, 2008.

[52] Eija Kaasinen. User needs for location-aware mobile services. *Personal and Ubiquitous Computing*, 7(1):70–79, 2003.

[53] P. Kalnis, N. Mamoulis, and S. Bakiras. On discovering moving clusters in spatio-temporal data. In C. B. Medeiros, M. J. Egenhofer, and E. Bertino, editors, *Proceedings of the 9th International Symposium on Advances in Spatial and Temporal Databases*, volume 3633 of *Lecture Notes in Computer Science*, pages 364–381. Springer, 2005.

[54] C.-H. Kang, J.-R. Hwang, and K.-J. Li. Trajectory analysis for soccer players. In *Proceedings of the 6th IEEE International Conference on Data Mining Workshop*, pages 377–381, 2006.

[55] W. Kellerer, C. Bettstetter, C. Schwingenschlogl, P. Sties, and K. E. Steinberg. (auto) mobile communication in a heterogeneous and converged world. *IEEE Personal Communications*, 8(6):41–47, 2001.

[56] E. Keogh and C. A. Ratanamahatana. Exact indexing of dynamic time warping. *Knowledge and Information Systems*, 7:358–386, 2005.

[57] E. J. Keogh and M. J. Pazzani. Scaling up dynamic time warping for datamining applications. In *Proceedings of the 6th ACM International Conference on Knowledge Discovery and Data Mining*, pages 285–289, 2000.

[58] H. Kido, Y. Yanagisawa, and T. Satoh. An anonymous communication technique using dummies for location-based services. In *International Conference on Pervasive Services (ICPS '05)*, pages 88–97, 2005.

[59] G. Langran. *Time in Geographic Information Systems.* PhD thesis, University of Washington, 1999.

[60] J. S. Larson, E. T. Bradlow, and P. S. Fader. An exploratory look at supermarket shopping paths. *International Journal of Research in Marketing*, 22:395–414, 2005.

[61] P. Laube. Progress in movement pattern analysis. In B. Gottfried and H Aghajan, editors, *Behaviour Monitoring and Interpretation - BMI - Smart Environments*, volume 3 of *Ambient Intelligence and Smart Environments*, pages 43–71. IOS Press, 2009.

[62] P. Laube, M. Duckham, and T. Wolle. Decentralized movement pattern detection amongst mobile geosensor nodes. In T. J. Cova, K. Beard, M. F. Goodchild, and A. U. Frank, editors, *GIScience 2008*, volume 5266 of *LNCS*, pages 199–216. Springer, 2008.

[63] P. Laube, S. Imfeld, and R. Weibel. Discovering relative motion patterns in groups of moving point objects. *International Journal of Geographical Information Science*, 19(6):639–668, 2005.

[64] P. Laube, M. van Kreveld, and S. Imfeld. Finding REMO - detecting relative motion patterns in geospatial lifelines. In P. F. Fisher, editor, *Developments in Spatial Data Handling*, Proceedings of the 11th International Symposium on Spatial Data Handling, pages 201–214. Springer, Berlin Heidelberg, DE, 2004.

[65] N. A. Lorentzos. *A Formal Extension of the Relational Model for the Representation and Manipulation of Generic Intervals.* PhD thesis, Birbeck College, University of London, 1988.

[66] Nancy Lynch. *Distributed Algorithms.* Morgan Kaufmann, San Mateo, CA, 1996.

[67] A. Maheshwari, J.-R. Sack, K. Shahbaz, and H. Zarrabi-Zadeh. Fréchet distance with speed limits. *Computational Geometry – Theory & Applications*, 44(2):110–120, 2011.

[68] N. Mamoulis, H. Cao, G. Kollios, M. Hadjieleftheriou, Y. Tao, and D. Cheung. Mining, indexing, and querying historical spatiotemporal data. In *Proceedings of the 10th ACM International Conference On Knowledge Discovery and Data Mining*, pages 236–245. ACM, 2004.

[69] D. Memmert and J. Perl. Game creativity analysis by means of neural networks. *Journal of Sport Science*, 27:139–149, 2009.

[70] N. Meratnia and Rolf A. de By. Spatiotemporal compression techniques for moving point objects. In *Proceedings of the 9th International Conference on Extending Database Technology*, pages 765–782, 2004.

[71] H. Miller and J. Han, editors. *Geographic Data Mining and Knowledge Discovery, Second Edition*. CRC Press, Boca Raton, FL, 2nd edition, 2009.

[72] H. J. Miller. Modelling accessibility using space-time prism concepts within geographical information systems. *International Journal of Geographical Information Systems*, 5(3):287–301, 1991.

[73] M. Nanni and D. Pedreschi. Time-focused clustering of trajectories of moving objects. *Journal of Intelligent Information Systems*, 27:267–289, 2006.

[74] R. Nathan, W. M. Getz, E. Revilla, M. Holyoak, R. Kadmon, D. Saltz, and P. E. Smouse. A movement ecology paradigm for unifying organismal movement research. *Proceedings of the National Academy of Sciences*, 105(49):19052–19059, 2008.

[75] R. T. Ng. Detecting outliers from large datasets. In H. J. Miller and J. Han, editors, *Geographic data mining and knowledge discovery*, pages 218–235. Taylor & Francis, London, UK, 2001.

[76] S. Nittel, A. Stefanidis, I. Cruz, M. J. Egenhofer, D. Goldin, A. Howard, A. Labrinidis, S. Madden, A. Voisard, and M. Worboys. Report from the first workshop on geo sensor networks. *ACM SIGMOD Record*, 33(1), 2004.

[77] S. Nouwt. Reasonable expectations of geo-privacy? *SCRIPTed*, 5(2):375–403, 2008.

[78] D. J. Peuquet. It's about time: A conceptual framework for the representation of temporal dynamics in geographic information systems. *Annals of the Association of American Geographers*, 83(3):441–461, 1994.

[79] D. Rafiei and A. O. Mendelzon. Querying time series data based on similarity. *IEEE Transactions on Knowledge and Data Engineering*, 12:675–693, 2000.

[80] S. Rasetic, J. Sander, J. Elding, and M. A. Nascimento. A trajectory splitting model for efficient spatio-temporal indexing. In *Proceedings of the 31st international conference on Very large data bases*, pages 934–945, 2005.

[81] Salvatore Rinzivillo, Dino Pedreschi, Mirco Nanni, Fosca Giannotti, Natalia Andrienko, and Gennady Andrienko. Visually driven analysis of movement data by progressive clustering. *Information Visualization*, 7(3-4):225–239, 2008.

[82] Y. Sakurai, M. Yoshikawa, and C. Faloutsos. Ftw: fast similarity search under the time warping distance. In *Proceedings of the 24th ACM Symposium on Principles of Database Systems*, pages 326–337, 2005.

[83] F. Schmid, K.-F. Richter, and P. Laube. Semantic trajectory compression, 2009.

[84] T. A. Shimp. *Advertising, Promotion, and Other Aspects of Integrated Marketing Communications*. South-Western Publisher, 2008.

[85] G. Sinha and D. M. Mark. Measuring similarity between geospatial lifelines in studies of environmental health. *Journal of Geographical Systems*, 7(1):115–136, 2005.

[86] A. P. Sistla, O. Wolfson, S. Chamberlain, and S. Dao. Modeling and querying moving objects. In *Proceedings of the 13th International Conference on Data Engineering*, pages 422–432. IEEE Computer Society, 1997.

[87] S. Săltenis, C. S. Jensen, S. T. Leutenegger, and M. A. Lopez. Indexing the positions of continuously moving objects. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 331–342, 2000.

[88] T. Taki and J.-I. Hasegawa. Visualization of dominant region in team games and its application to teamwork analysis. In *Proceedings of the International Conference on Computer Graphics*, 2000.

[89] Tsuyoshi Taki and Jun-ichi Hasegawa. A feature for group behavior analysis in team sports and its application. *The Transactions of the Institute of Electronics, Information and Communication Engineers*, J81-D-II(8):1802–1811, 1998.

[90] J. J. Thomas and K. A. Cook. A visual analytics agenda. *IEEE Computer Graphics and Applications*, 26(1):10–13, 2006.

[91] S. M. Tomkiewicz, M. R. Fuller, J. G. Kie, and K. K. Bates. Global positioning system and associated technologies in animal behaviour and ecological research. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1550):2163–2176, 2010.

[92] G. Trajcevski, H. Ding, P. Scheuermann, R. Tamassia, and D. Vaccaro. Dynamics-aware similarity of moving objects trajectories. In *Proceedings of the 15th Annual ACM International Symposium on Advances in Geographic Information Systems*, pages 11:1–11:8, 2007.

[93] U.S. Department of Justice, Office of Information and Privacy. Overview of the privacy act of 1974, May 2004 2004.

[94] A. Uteck. Ubiquitous computing and spatial privacy, anonymity, privacy and identity in a networked society. In I. Kerr, V. Steeves, and C. Lucock, editors, *Lessons from the Identity Trail*, volume Oxford, pages 83–102. Oxford University Press, 2009.

[95] M. van Kreveld and J. Luo. The definition and computation of trajectory and subtrajectory similarity. In *Proceedings of the 15th Annual ACM International Symposium on Advances in Geographic Information Systems*, pages 44:1–44:4, 2007.

[96] F. Verhein and S. Chawla. Mining spatio-temporal association rules, sources, sinks, stationary regions and thoroughfares in object mobility databases. In *Proceedings of the 11th International Conference on Database Systems for Advanced Applications*, volume 3882 of *Lecture Notes in Computer Science*, pages 187–201. Springer, 2006.

[97] M. Vlachos, D. Gunopoulos, and G. Kollios. Discovering similar multidimensional trajectories. In *Proceedings of the 18th International Conference on Data Engineering*, pages 673–682, 2002.

[98] U. Wilensky. Netlogo flocking model. Web Resource - verified 10-01, 2011. http://ccl.northwestern.edu/netlogo/models/Flocking.

[99] Z. Wood and A. Galton. Classifying collective motion. In B. Gottfried and H Aghajan, editors, *Behaviour Monitoring and Interpretation - BMI - Smart Environments*, volume 3 of *Ambient Intelligence and Smart Environments*, pages 129–155. IOS Press, Amsterdam, NL, 2009.

[100] B.-K Yi and C. Faloutsos. Fast time sequence indexing for arbitrary lp norms. In *Proceedings of the 26th International Conference on Very Large Data Bases*, pages 385–394, 2000.

[101] B.-K. Yi, H. V. Jagadish, and C. Faloutsos. Efficient retrieval of similar time sequences under time warping. In *Proceedings of the 14th International Conference on Data Engineering*, pages 201–208, 1998.

[102] H. Yoon and C. Shahabi. Robust time-referenced segmentation of moving object trajectories. In *Proceedings of the IEEE International Conference on Data Mining*, pages 1121–1126, 2008.

[103] Z. Zhang, K. Huang, and T. Tan. Comparison of similarity measures for trajectory clustering in outdoor surveillance scenes. *Proceedings of the 18th International Conference on Pattern Recognition*, pages 1135–1138, 2006.