Chapter 2 - Basic Concepts of Movement Data

Natalia Andrienko[1], Gennady Andrienko[1], Nikos Pelekis[2], and Stefano Spaccapietra[3]

[1] Fraunhofer Institute, Germany {gennady.andrienko, natalia.andrienko}@iais.fraunhofer.de
[2] Computer Technology Institute (CTI) and Department of Informatics, University of Piraeus, Greece npelekis@unipi.gr
[3] Database laboratory, École Polytechnique Fédérale de Lausanne, Switzerland stefano.spaccapietra@epfl.ch

1 Introduction

Perhaps, ever since people exist, they have observed various moving entities, from insects and fishes to planets and stars, and investigated their movement behaviours. While methods that were used in earlier times for observation, measurement, recording, and analysis of movements are very different from modern technologies, there is still much to learn from past studies. First, this is the thorough attention paid to the multiple aspects of movement. These include not only the trajectory (path) in space, characteristics of motion itself such as speed and direction, and their dynamics over time but also characteristics and activities of the entities that move. Second, this is the striving to relate movements to properties of their surroundings and to various phenomena and events.

As an illustration, let us take the famous depiction of Napoleon's march on Moscow published by Charles Joseph Minard in 1861 (this representation is reproduced in Fig. 1; a detailed description can be found in Tufte [15]). The author engages the readers in the exploration of the fate of Napoleon's army in the Russian campaign of 1812-1813. Beginning at the Polish-Russian border, the thick band shows the size of the army at each position. The path of Napoleon's retreat from Moscow in the cold winter is depicted by the dark lower band, which is tied to temperature and time scales. Tufte [15] identified six separate variables that were shown within Minards drawing. First, the line width continuously marked the size of the army. Second and third, the line itself showed the latitude and longitude of the army as it moved. Fourth, the lines themselves showed the direction that the army was travelling, both in advance and retreat. Fifth, the location of the army with respect to certain dates was marked. Finally, the temperature along the path of retreat was displayed. It can also be noted that, despite of the schematic character of the drawing with its rudimentary cartography, Minard depicted some features of the underlying territory (specifically, rivers and towns) he deemed essential for the understanding of the story.

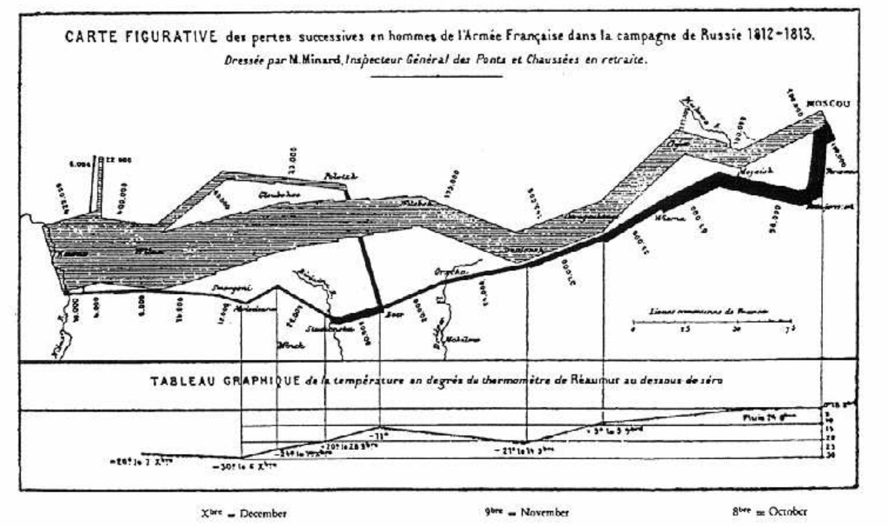

Fig. 1. The representation of Napoleon's Russian campaign of 1812 produced by Charles Joseph Minard in 1861.

Since the environment in which movements take place and the characteristics of the moving entities may have significant influence on the movements, they need to be considered when the movements are studied. Moreover, movements themselves are not always the main focus of a study. One may analyse movements with the aim to gain knowledge about the entities that move or about the environment of the movements. Thus, in the research area known as time geography, the observation of everyday movements of human individuals was primarily the means of studying activities of different categories of people. On an aggregate level, time geography looks for trends in society.

The ideas of time geography originate from Torsten Hagerstrand [6]. A prominent feature of time geography is the view of space and time as inseparable. Hagerstrands basic idea was to consider space-time paths in a 3D space where horizontal axes represent geographic space and the vertical axis represents time. This representation is known as space-time cube. The idea is illustrated in Fig. 2 left. The line represents the movements of some entity, for example, a working person who initially was at home, then travelled to his workplace and stayed there for a while, then moved to a supermarket for shopping and, having spent some time there, returned back home. Vertical lines stand for stays at a certain location (home, workplace, or supermarket). The workplace is an example of a station, that is a place where people meet for a certain activity. The sloped line segments indicate movements. The slower the movement, the steeper the line will be. The straightness of the lines in our drawing assumes that the person travels with constant speed, which is usually just an approximation of the real behaviour. The space-time path can be projected on a map, resulting in the path's footprint.

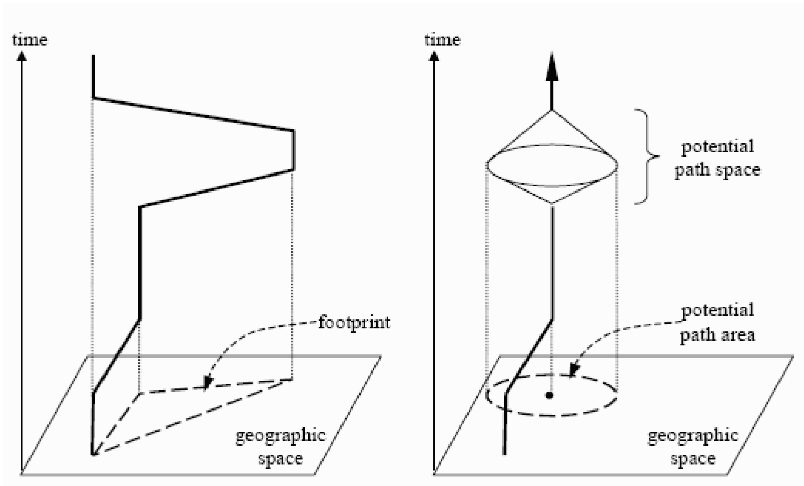

Fig. 2. An illustration of the notions of space-time path and space-time prism

Another important concept of time geography is the notion of space-time prism, which is schematically illustrated in Fig. 2 right. In the 3D representation, this is the volume in space and time a person can reach in a particular time interval starting and returning to the same location (for instance, where a person can get from his workplace during lunch break). The widest extent is called the potential path space and its footprint is called potential path area. In Figure 2 right it is represented by a circle assuming it to be possible to reach every location within the circle. In reality, the physical environment will not always allow this. In general, the space-time paths of individuals are influenced by constraints. One can distinguish between capability constraints (for instance, mode of transport and need for sleep), coupling constraints (for instance, being at work or at the sports club), and authority constraints (for instance, accessibility of buildings or parks in space and time).

In the era of pre-computer graphics, it was time-consuming and expensive to produce space-time cube visualisations to support the exploration of movement behaviours. However, with the rise of new visualisation technology and interactivity, researchers revisited this concept ([8], [13]). Moreover, modern time geography is not entirely based on visual representations and qualitative descriptions. Thus, Miller [10] suggests a measurement theory for its basic entities and relationships, which includes formal definitions of the space-time path, space-time prism, space-time stations as well as fundamental relationships between space-time paths and prisms. This provides foundations for building computational tools for time geographic querying and analysis.

Whatever tools and technologies have been used for the collection, representation, exploration, and analysis of movement data, the underlying basic concepts related to the very nature of movement in (geographical) space remain stable and the characteristics of movement examined in past studies do not lose their relevance. In Section 2, we present a synthesis from existing literature concerning the basic concepts and characteristics of movement. Movement occurs in space and in time; so we discuss the possible ways of spatial and temporal referencing and relevant properties of space and time. We also briefly mention other matters that may have an impact on movement and therefore need attention in analysis. These include properties and activities of moving entities and various space- and/or time-related phenomena and events.

Data analysis is seeking answers to various questions about data. In Section 3, we define the types of questions that can arise in analysis of movement data. For the question types to be independent of any analysis methods and tools, we define them on the basis of an abstract model of movement data, which involves three fundamental components: population of entities, time, and space. We distinguish between elementary questions, which refer to individual data items, and synoptic questions, which refer to the data as a whole or to data subsets considered in their entirety. Synoptic questions play the primary role in data analysis. At the end, we relate the tool-independent taxonomy of analytical questions to the established typology of data mining tasks.

2 Movement data and their characteristics

This section presents a synthesis from the current literature talking about movement and movement data: what is movement? How can movement be reflected in data? How can movement be characterised? What does it depend on?

2.1 Trajectories

A strict definition of *movement* relates this notion to change in the physical position of an entity with respect to some reference system within which one can assess positions. Most frequently, the reference system is geographical space.

A *trajectory* is the path made by the moving entity through the space where it moves. The path is never made instantly but requires a certain amount of time. Therefore, time is an inseparable aspect of a trajectory'. This is emphasised in the term 'space-time path' [6][10][11], one of the synonyms for ‘trajectory’. Another well-known term, ‘geospatial lifeline’ introduced by Hornsby and Egenhofer [7], also refers to time although less explicitly (through the notion of ‘life’).

If t_0 is the time moment when the path started and t_{end} is the moment when it ended, for any moment t_i between t_0 and t_{end} there is a position in space that was occupied by the entity at this moment (although in practice this position is not always known). Hence, a trajectory can be viewed as a function that matches time moments with positions in space. It can also be seen as consisting of pairs (time, location). Since time is continuous, there are an infinite number of such pairs in a trajectory. For practical reasons, however, trajectories have to be represented by finite sequences of time-referenced locations. Such sequences may result from various ways used to observe movements and collect movement data:

- time-based recording: positions of entities are recorded at regularly spaced time moments, e.g. every 5 minutes;
- change-based recording: a record is made when the position of an entity differs from the previous one;
- location-based recording: records are made when an entity comes close to specific locations, e.g. where sensors are installed;
- event-based recording: positions and times are recorded when certain events occur, in particular, activities performed by the moving entity (e.g. calling by a mobile phone);
- various combinations of these basic approaches.

Typically, positions are measured with uncertainty. Sometimes it is possible to refine the positions taking into account physical constraints, e.g. the street network.

In studying movements, an analyst attends to a number of characteristics, which can be grouped depending on whether they refer to states at individual moments or to movements over time intervals. *Moment-related characteristics* include:

- time, i.e. position of this moment on the time scale;
- position of the entity in space;
- direction of the entity’s movement;
- speed of the movement (which is zero when the entity stays in the same place);
- change of the direction (turn);
- change of the speed (acceleration);
- accumulated travel time and distance.

Overall characteristics of a trajectory as a whole or a trajectory fragment made during a subinterval $[t_1, t_2]$ of the entire time span $[t_0, t_{end}]$ include:

- geometric shape of the trajectory (fragment) in the space;
- travelled distance, i.e. the length of the trajectory (fragment) in space;

- duration of the trajectory (fragment) in time;
- movement vector (i.e. from the initial to the final position), or major direction;
- mean, median, and maximal speed;
- dynamics (behaviour) of the speed:
  - periods of constant speed, acceleration, deceleration, and stillness;
  - characteristics of these periods: start and end times, duration, initial and final positions, initial and final speeds, etc.;
  - arrangement (order) of these periods in time;
- dynamics (behaviour) of the directions:
  - periods of straight, curvilinear, circular movement;
  - characteristics of these periods: start and end times, initial and final positions and directions, major direction, angles and radii of the curves, etc.;
  - major turns ('turning points') with their characteristics: time, position, angle, initial and final directions, and speed of the movement in the moment of the turn;
  - arrangement (order) of the periods and turning points in time.

Besides examining a single trajectory, an analyst is typically interested in *comparison* of two or more trajectories. These may be trajectories of different entities (e.g. different persons), trajectories of the same entity made at different times (e.g. trajectories of a person on different days), or different fragments of the same trajectory (e.g. trajectories of a person on the way from home to the workplace and on the way back). Generally, the goal of comparison is to establish *relations* between the objects that are compared. Here are some examples of possible relations:

- equality or inequality;
- order (less or greater, earlier or later, etc.);
- distance (in space, in time, or on any numeric scale);
- topological relations (inclusion, overlapping, crossing, touching, etc.).

Many other types of relations may be of interest, depending on the nature of the things being compared. In comparing trajectories, analysts are most often interested in establishing the following types of relations:

- similarity or difference of the overall characteristics of the trajectories, which have been listed above (i.e. shapes, travelled distances, durations, dynamics of speed and directions, and so on);
- spatial and temporal relations:
  - co-location in space, full or partial (i.e. the trajectories consist of the same positions or have some positions in common):
    - ordered co-location: the common positions are attained in the same order;
    - unordered co-location: the common positions are attained in different orders;

- co-existence in time, full or partial (i.e. the trajectories are made during the same time period or the periods overlap);
- co-incidence in space and time, full or partial (i.e. same positions are attained at the same time);
- lagged co-incidence, i.e. entity e_1 attains the same positions as entity e_0 but after a time delay Δt;
- distances in space and in time.

Most researchers dealing with movement data agree in recognising the necessity to consider not only trajectories with their spatial and temporal characteristics but also the structure and properties of the space and time where the movement takes place as having a great impact upon the movement behaviour. The concepts and characteristics related to space and time are briefly discussed below.

2.2 Space

Space can be seen as a set consisting of *locations*, or *places*. An important property of space is the existence of distances between its elements. At the same time, space has no natural origin and no natural ordering between the elements. Therefore, in order to distinguish positions in space, one needs to introduce in it some reference system, for example, a system of coordinates. While this may be done, in principle, quite arbitrarily, there are some established reference systems such as geographical coordinates.

Depending on the practical needs, one can treat space as two-dimensional (i.e. each position is defined by a pair of coordinates) or as three-dimensional (each position is defined by a triple of coordinates). In specific cases, space can be viewed as one-dimensional. For example, when movement along a standard route is analysed, one can define positions through the distances from the beginning of the route, i.e. a single coordinate is sufficient.

Theoretically, one can also deal with spaces having more than three dimensions. Such spaces are abstract rather than physical; however, movements of entities in abstract spaces may also be subject to analysis. Thus, Laube et al. [9] explore the movement (evolution) of the districts of Switzerland in the abstract space of politics and ideology involving three dimensions: left vs. right, liberal vs. conservative, and ecological vs. technocratic.

The physical space is continuous, which means that it consists of an infinite number of locations and, moreover, for any two different locations there are locations "in between", i.e. at smaller distances to each of the two locations than the distance between the two locations. However, it may also be useful to treat space as a discrete or even finite set of locations. For example, in studying the movement of tourists over a country or a city, one can "reduce" space to the set of points of interest visited by the tourists. Space discretisation may be even indispensable, in particular, when positions of entities cannot be measured precisely and are specified in terms of areas such as cells of a mobile phone network, city districts, or countries.

The above-cited examples show that space may be *structured*, in particular, divided into areas. The division may be hierarchical; for instance, a country is divided into provinces, the provinces into municipalities, and the municipalities into districts. Areas can also be derived from a geometric decomposition (e.g. 1sq.km. cells), with no semantics associated to the decomposition. A street (road) network is another common way of structuring physical space.

Like coordinate systems, space structuring also provides a reference system, which may be used for distinguishing positions, for instance, by referring to streets or road fragments and relative positions on them (house numbers or distances from the ends). The possible ways of specifying positions in space can be summarised as follows:

- coordinate-based referencing: positions are specified as tuples of numbers representing linear or angular distances to certain chosen axes or angles;
- division-based referencing: referring to compartments of an accepted geometric or semantic-based division of the space, possibly, hierarchical;
- linear referencing: referring to relative positions along linear objects such as streets, roads, rivers, pipelines, etc.; for example, street names plus house numbers or road codes plus distances from one of the ends.

Since it is often the case that positions of entities cannot be determined accurately, they may be represented in data with uncertainty, for example, as areas instead of points.

Sometimes, an analyst is not so much interested in absolute positions in space as in relative positions with regard to a certain place. For example, the analyst may study where a person travels with regard to his/her home or movements of spectators to and from a cinema or a stadium. In such cases, it is convenient to define positions in terms of distances and directions from the reference place (or, in other words, by means of polar coordinates). The directions can be defined as angles from some base direction or geographically: north, northwest, and so on.

Comprehensive analysis may require consideration of the same data within different systems of spatial referencing and, hence, transformation of one reference system to another: geographical coordinates to polar (with various origins), coordinate-based referencing to division-based or network-based, etc.

It may also be useful to disregard the spatial positions of locations and consider them from the perspective of their domain-specific semantics, e.g. home, workplace, shopping place, etc.

It should be noted that space (in particular, physical space) is not uniform but heterogeneous, and its properties vary from place to place. These properties may have a great impact on movement behaviours and, hence, should be taken into account in analysis. The relevant characteristics of individual locations include:

- altitude, slope, aspect and other characteristics of the terrain;
- accessibility with regard to various constraints (obstacles, availability of roads, etc.);
- character and properties of the surface: land or water, concrete or soil, forest or field, etc.;
- objects present in a location: buildings, trees, monuments, etc.;
- function or way of use, e.g. housing, shopping, industry, agriculture, or transportation;
- activity-based semantics, e.g. home, work, shopping, leisure, and so on.

When locations are defined as space compartments (i.e. areas in two-dimensional space or volumes in three-dimensional space) or network elements rather than points, the relevant characteristics also include:

- spatial extent and shape;
- capacity, i.e. the number of entities the location can simultaneously contain;
- homogeneity or heterogeneity of properties (listed above) over the compartment.

It should be noted that properties of locations may change over time. For example, a location may be accessible on weekdays and inaccessible on weekends; a town square may be used as a marketplace in the morning hours; a road segment may be blocked or its capacity reduced because of an accident or reparation works.

Similarly to space, there are different ways of defining positions in time, and time may also be heterogeneous in terms of properties of time moments and intervals.

2.3 Time

Mathematically, time is a continuous set with a linear ordering and distances between the elements, where the elements are moments, or positions in time. Analogously to positions in space, some reference system is needed for the specification of moments in data. In most cases, temporal referencing is done on the basis of the standard Gregorian calendar and the standard division of a day into hours, hours into minutes, and so on. The time of the day may be specified according to the time zone of the place where the data are collected or as Greenwich Mean Time (GMT). There are cases, however, when data refer to relative time moments, e.g. the time elapsed from the beginning of a process or observation, or abstract time stamps specified as numbers 1, 2, and so on. Unlike the physical time, abstract times are not necessarily continuous.

Like positions in space, moments may be specified imprecisely, i.e. as intervals rather than points in time. But even when data refer to points, they are indispensably imprecise: since time is continuous, the data cannot refer to every possible point. For any two successive moments t_1 and t_2 referred to in the data there are moments in between for which there are no data. Therefore, one cannot know definitely what happened between t_1 and t_2 but can only estimate this by means of interpolation.

Physical time is not only a linear sequence of moments but includes inherent cycles resulting from the earths daily rotation and annual revolution. These natural *cycles* are reflected in the standard method of time referencing: the dates are repeated in each year and the times in each day. Besides these natural cycles, there are also cycles related to people's activities, for example, the weekly cycle. Various domain- and problem-specific cycles exist as well, for example, the revolution periods of the planets in astronomy or the cycles of the movement of buses or local trains on standard routes.

Temporal cycles may be nested; in particular, the daily cycle is nested within the annual cycle. Hence, time can be viewed as a hierarchy of nested cycles. Several alternative hierarchies may exist, for example, year/month/day-in-month and year/week-in-year/day-in-week.

It is very important to know which temporal cycles are relevant to the movements under study and to take these cycles properly into account in the analysis. For this purpose, it is necessary that the cycles were reflected in temporal references of the data items. Typically, this is done through specifying the cycle number and the position from the beginning of the cycle. In fact, the standard references to dates and times of the day are built according to this principle. However, besides the standard references to the yearly and daily cycles, references to other (potentially) relevant cycles, e.g. the weekly cycle of people's activities or the cycles of the movement of satellites, may be necessary or useful. Hence, an analyst may need to transform the standard references into references in terms of alternative time hierarchies.

Temporal cycles may have variable periods. For example, the cycle of El Nino and La Nina climatic events, which influences the movement of air and water masses in the Pacific Ocean, has an average return period of four and a half years but can recur as little as two or as much as ten years apart. In order to make data related to different cycles comparable, one needs to "standardise" somehow the time references, for example, divide the absolute time counts from the beginning of a cycle by the length of this cycle.

Transformation of absolute time references to relative is also useful when it is needed to compare movements that start at different times and/or proceed with different speeds. The relative time references would in this case be the time counts from the beginning of each movement, possibly, standardised in the way of dividing them by the duration of the movement.

As we have noted, the properties of time moments and intervals may vary, and this variation may have significant influence on movements. For example, the movements of people on weekdays notably differ from the movements on weekends; moreover, the movements on Fridays differ from those on Mondays and the movements on Saturdays from those on Sundays. In this example, we have a case of a regular difference between positions within a cycle. Another example of the same kind is the difference between times of a day: morning, midday, evening, and night. However, the regularity in the variation of properties of time moments may be disrupted, for example, by an intrusion of public holidays. Not only the intrusions themselves but also the preceding and/or following times may be very different from the "normal" time; think, for example, of the days before and after Christmas. Such irregular changes should also be taken into account in the analysis of time-dependent phenomena, in particular, movements.

The regularity of changes may itself vary, in particular, owing to interactions between larger and smaller temporal cycles. Thus, the yearly variation of the duration of daylight has an impact on the properties of times of a day, which, in turn, influence movements of people and animals. In the results, movements at the same time of the day in summer and in winter may substantially differ.

Typically, the heterogeneity of properties of time is not explicitly reflected in data and, hence, cannot be automatically taken into account in data analysis. Much depends on the analyst's ability to involve his/her background knowledge. Hence, the methods and tools used for the analysis must allow the analyst to do this.

2.4 Moving entities and their activities

Like locations in space and moments in time, the entities that move have their own characteristics, which may influence the movement and, hence, need to be taken into account in the analysis. Thus, the movements of people may greatly depend on their occupation, age, health condition, marital status, and other properties. It is also relevant whether an entity moves by itself or by means of some vehicle. The way and means of the movement pose their constraints on the possible routes and other characteristics of the movement.

People are an example of entities that typically move purposely. The purposes determine the routes and may also influence the other characteristics, in particular, the speed. For other types of entities, for example, tornadoes or elementary particles, one needs to attend to the causes of the movement rather than the purposes.

Movement characteristics may also depend on the activities performed by the entities during their movement. For example, the movement of a person in a shop differs from the movement on a street or in a park. The characteristics of the movement may change when the person starts speaking by a mobile phone.

2.5 Related phenomena and events

Any movement occurs in some environment and is subject to the influences from various events and phenomena taking place in this environment. Thus, Minard included a graph of winter temperatures in his depiction of Napoleon's Russian campaign since he was sure that the temperatures produced a great influence on the movement and fate of the army. Movements of people are influenced by the climate and current weather, by sport and cultural events, by legal regulations and established customs, by road tolls and oil prices, by shopping actions and traffic accidents, and so on. In order to detect such influences or to take them into account in movement data analysis, the analyst needs to involve additional data and/or background knowledge.

We have reviewed thus far what characteristics and aspects of movement are considered in the analysis of movement data and what other types information are relevant. However, we did not define what it means, "to analyse movement data", and for what purposes such an analysis is done. Let us now try to do this.

3 Analytical questions

One can hardly find a strict definition of the term "data analysis" in handbooks or research literature. However, most of the writers agree in the view of data analysis as an iterative process consisting of the following activities:

- formulate questions;
- choose analysis methods;
- prepare the data for application of the methods;
- apply the methods to the data;
- interpret and evaluate the results obtained.

In short, data analysis is formulating questions and seeking answers. In this section, we try to define the types of questions that can arise in analysis of movement data. Examples of various questions concerning moving entities can be easily found in literature, for instance, in (Guting and Schneider [5]).

- How often do animals stop?
- Which routes are regularly used by trucks?
- Did the trucks with dangerous goods come close to a high-risk facility?
- Were any two planes close to a collision?
- Find "strange" movements of ships, indicating illegal dumping of waste.

However, we did not find a systematic taxonomy of the *types of questions* relevant to the analysis of movement data. Therefore, we try to build such a taxonomy by applying and adapting the general framework suggested by Jacques Bertin [3] and extended by Andrienko and Andrienko [2].

Jacques Bertin is a French cartographer and geographer, who was the first in articulating a coherent and reasoned theory for what is now called Information Visualisation. Bertin has developed a comprehensive framework for the design of maps and graphics intended for data analysis, where the function of a graphic is answering questions. Logically, a part of Bertin's theory deals with the types of questions that may need to be answered. The question types, as Bertin defines them, have no specific "graphical flavour" and no influence of any other method for data representation or analysis. Questions are formulated purely in the "language" of data, and hence have general relevance. Therefore, we can use Bertin's framework to define the types of questions that arise in analysis of movement data irrespectively of what analysis methods are chosen.

To achieve this independence, we define the question types on the basis of an abstract view of the structure of movement data, which is presented next. In our typology, we distinguish between elementary questions, which refer to individual data items, and synoptic questions, which refer to the data as a whole or to data subsets considered in their entirety. Synoptic questions play the primary role in data analysis. We consider various types of elementary and synoptic questions. At the end, we relate the tool-independent taxonomy of analytical questions to the established typology of data mining tasks.

3.1 Data structure

According to the general framework, the types of questions are defined on the basis of the structure of the data under analysis, i.e. what components the data consist of and how they are related. On an abstract level, movement data can be viewed as consisting of the three principal components:

- time: a set of moments;
- population (this term is used in statistical rather than demographic sense): a set of entities that move;
- space: a set of locations that can be occupied by the entities.

As noted above, a trajectory may be viewed as a function mapping time moments onto positions in space. Analogously, movement of multiple entities may be seen as a function mapping pairs <time moment, entity> onto positions. This is a very abstract data model, which is independent of any representation formalism (of course, there may be other models; for example, a database-oriented view would consider the same data as a table of tuples with at least three attributes Entity, Time, and Space). The time and population of entities play the role of "independent variables", or *referential components*, according to the terminology suggested by Andrienko and Andrienko [2] and the space plays the role of "dependent variable", or *characteristic component.*

A combination of values of the referential components is called a *reference.* In our case, a reference is a pair consisting of a time moment and an entity. The set of all possible references is called the *reference set.* Values of the characteristic components corresponding to the references are called *characteristics* of these references.

As it was mentioned in the previous section, the state of a moving entity at a selected time moment can be characterised not only by its position in space but also by additional characteristics such as speed, direction, acceleration, etc. These characteristics can be viewed as secondary since they can be derived from the values of the principal components. Nevertheless, we can extend our concept of movement data and see it as a function mapping references <time moment, entity> onto combinations of characteristics (position, speed, direction, ...).

We have also mentioned in the previous section that locations, time moments, and entities may have their own characteristics. For example, locations may be characterised by altitude, slope, character of the surface, etc.; entities may be characterised by their kind (people, vehicles, animals, ...), age, gender, activity, and so on. Such characteristics are independent of the movement, that is, do not refer to pairs <time moment, entity> but to individual values of the three principal components, time, population, and space. Note that the space plays the role of a referential component for altitude, slope, and so on. The characteristics of time moments, entities, and locations will be further called *supplementary characteristics*. The characteristics of the pairs <time moment, entity> (including the secondary ones) will be called *characteristics of movement*.

Analytical questions arising in the analysis of movement data address first of all the references (i.e. times and entities) and the characteristics of movement. However, they may also involve supplementary characteristics.

3.2 Elementary and synoptic questions

The types of questions are differentiated first of all according to their *level*: whether they address individual references or sets of references. Questions addressing individual references are called *elementary*. The term ‘elementary’ means that the questions address *elements* of the reference set. Questions addressing sets of references (either the whole reference set or its subsets) are called *synoptic*. The word ‘synoptic’ is defined in a dictionary (Merriam-Webster [1], p.1197) as follows:

1. affording a general view of a whole
2. manifesting or characterized by comprehensiveness or breadth of view
3. presenting or taking the same or common view; specifically often capitalized : of or relating to the first three Gospels of the New Testament
4. relating to or displaying conditions (as of the atmosphere or weather) as they exist simultaneously over a broad area

The first interpretation is the closest to what we mean by synoptic questions, which assume a general view of a reference (sub)set as a whole, as will be clear from examples below. Interpretations 2 and 4 are also quite consistent with our usage of the term.

When there are two referential components, like in movement data, a question may be elementary with respect to one of them and synoptic with respect to the other. Examples are given in Table 1. Note that these examples are templates rather than specific questions since they contain slots, or variables.

		Population	
		Elementary	Synoptic
Time	Elementary	Where was entity e at time moment t?	What was the spatial distribution of all entities at time moment t?
	Synoptic	How did entity e move during the time period from t_1 to t_2?	How did all entities move during the time period from t_1 to t_2?

Table 1. Different levels of questions about movement data.

The difference between elementary and synoptic questions is not merely the number of elements involved. It is more fundamental: a synoptic question requires one to deal with a set as a whole, in contrast to elementary questions addressing individual elements. Although an elementary question may address two or more elements, it does not require these elements to be considered all together as a unit. Compare, for instance, the questions:

- What were the positions of entities e_1, e_2, ..., e_n at time moment t?
- What was the spatial distribution of the set of entities e_1, e_2, ..., e_n at time moment t?

The first question is elementary with respect to the population although it addresses multiple entities. However, each entity is addressed individually, and the question about n entities is therefore equivalent to n questions asking about each of the entities separately (i.e. the same answer can be given in both cases: entity e_1 was in place p_1, e_2 was in p_2, and so on). The second question does not ask about the individual positions of all entities but about the spatial distribution of the set of entities as a whole. The possible answers could be "the entities are distributed evenly" (or randomly, or concentrated in some part of the territory, or aligned, etc.).

In our examples, the elementary questions ask about locations of entities at time moments. They may also ask about the secondary characteristics of movement corresponding to references *t*ime moment, entity>, e.g. 'What was the speed of entity e at moment t?' Supplementary characteristics may also be involved, as in the question 'Describe the location where entity e was at moment t'. To answer this question, one needs, first, to determine the spatial position of entity e at moment t, second, to ascertain the supplementary characteristics of the location thus found.

What do synoptic questions ask about? What is in common between "how did the entity (entities) move?" and "what was the spatial distribution of the entities" (see Table 1)?

3.3 Behaviour and pattern

We introduce the notion of *behaviour*: this is the configuration of characteristics corresponding to a given reference (sub)set. The notion of behaviour is a generalisation of such notions as distribution, variation, trend, dynamics, trajectory, etc. In particular, a trajectory of a single entity is a configuration of locations (possibly, in combination with the secondary characteristics of movement) corresponding to a time interval. We say "configuration" rather than "set" meaning that the characteristics are arranged in accordance with the structure and properties of the reference (sub)set and the relations between its elements. Thus, since a time interval is a continuous linearly ordered set, a trajectory is a continuous sequence of locations ordered according to the times they were visited.

The term "behaviour" is used here in a quite general sense and does not necessarily mean a process going on in time. Thus, the spatial distribution of a set of entities at some time moment is also a kind of behaviour although it does not involve any temporal variation.

Since a population of entities is a discrete set without natural ordering and distances between the elements, it does not impose any specific arrangement of the corresponding characteristics. Still, the corresponding behaviour is not just a set of characteristics. Thus, one and the same characteristic or combination of characteristics can occur several times, and these occurrences are treated as different while in a set each element may occur only once. A behaviour over a set of entities may be hence conceptualised as the frequency distribution of the characteristic values over this set of entities.

The absence of natural ordering and distances on a population of entities does not mean that ordering and distances between entities cannot exist at all. Thus, a set of participants of a military parade is spatially ordered and has distances between the elements. However, the ordering and distances are defined in this case on the basis of certain characteristics of the entities, specifically, their spatial positions. The characteristics that define ordering and/or distances between entities can be chosen, in principle, quite arbitrarily. Thus, participants of a parade can also be ordered according to their heights, or weights, or ages. In data analysis, it may be useful to consider different orderings of the entities and the corresponding arrangements of characteristics. In such cases, the behaviours are not just frequency distributions but more complex constructs where characteristic values are positioned according to the ordering and/or distances between the entities they are associated with.

The collective movement behaviour of a population of entities over a time period is a complex configuration built from movement characteristics of all entities at all time moments, which has no arrangement with respect to the population of entities and has a continuous linear arrangement with respect to the time.

Hence, synoptic questions address reference (sub)sets and corresponding behaviours while elementary questions address individual references and corresponding characteristics. An answer to an elementary question is (are) the value(s) of the characteristic component(s) it is asking about. An answer to a synoptic question is a description of the behaviour or, more generally, a representation of this behaviour in some language, e.g. natural, mathematical, graphical, etc. Such a representation will be called *pattern*. This agrees with the definition of a pattern in the data mining literature: "a pattern is an expression E in some language L describing facts in a subset F_E of a set of facts F so that E is simpler than the enumeration of all facts in F_E" ([4]). Note that the latter definition emphasises the synoptic nature of a pattern: a pattern does not simply enumerate some facts but describes them all together as a whole.

As should be clear from the definition, different patterns (e.g. focusing on different aspects) may represent one and the same behaviour. A pattern may be compound, i.e. composed of other patterns. For example, the description "most of the people tend to move towards the city centre in the morning and outwards in the evening" is a compound pattern including two simpler patterns, inward and outward movement. Patterns representing movement behaviours of individual entities (i.e. trajectories) and collective movement behaviours of sets of entities base first of all on the characteristics of movement but may also involve supplementary characteristics. Thus, our example pattern concerning the movement of people describes first of all the direction of the movement but also mentions such supplementary characteristics as the character of the moving entities (people), the character of a location (city centre), and the character of the times (morning or evening).

In a pattern describing the movement behaviour on a set of references, one may include various summary values derived from the individual characteristics of the references, for instance, the average speed, prevailing direction, or frequency of turns.

3.4 Structure of a question

Any question contains some information that is known to the person who asks the question and aims at gaining some new information, which must be somehow related to the known information. The expected new information will be called the *target* of the question while the known information will be called the *constraint* (since it sets certain requirements to the content of the new information being sought). Thus, in a question asking about the characteristic corresponding to a given reference, the characteristic is the target while the reference is the constraint. For example, in the question "Where was the entity e at time moment t?" the reference, i.e. the pair (e, t), is the constraint and the target is the location corresponding to this pair. There are also inverse questions, which ask about references corresponding to given characteristics, for example, "What entities visited place p and when?" In this question, the target is the unknown pair consisting of an entity and a time moment that corresponds to the given place p, which is the constraint of this question.

When references consist of two components, as time and entity in the case of movement data, one of the components may be included in the question constraint while the other being the target:

- What entities were present in place p at time t?
- At what moments (if any) did entity e visit place p?

These are examples of elementary questions. Synoptic questions, which deal with reference sets and behaviours, have the same structure, i.e. include targets and constraints. In the examples of synoptic questions in Table 1, the constraints are (sub)sets of references and the targets are the behaviours corresponding to these (sub)sets. There are also questions where behaviours (described by means of appropriate patterns) are the constraints and reference (sub)sets are the targets, for example, "What group(s) of entities and in what time period(s) moved as specified by pattern P?"

Like in elementary questions, one of the components defining the references (i.e. set of entities or time interval) may appear in the constraint of a sentence while the other being the target:

- What entities moved as specified by pattern P during the time interval from t_1 to t_2?
- In what time period(s) did the group of entities $e_1, e_2, , e_n$ move as specified by pattern P?

Synoptic questions requiring the search for occurrences of specified patterns, as in the above-presented examples, may be called *pattern search tasks*. We highlight this question type and give it a special name since it plays a prominent role in visual data exploration, which is generally viewed as being based on pattern recognition.

3.5 Comparison questions

In the examples considered so far, the questions were targeted at (i.e. asking about) characteristics, or behaviours, or references, or reference sets. Let us give a few examples of a different kind:

1. What were the relative positions of entities e_1 and e_2 at time t?
2. How did the location of entity e change from time t_1 to time t_2?
3. What is the difference in the times when entity e visited places p_1 and p_2?
4. What are the commonalities and differences between the movement behaviours of entities e_1 and e_2 (or groups of entities E_1 and E_2) on the time interval from t_1 to t_2?
5. How does the movement behaviour of entity e (or group of entities E) in time interval from t_1 to t_2 differ from the behaviour in interval from t_3 to t_4? What is in common?
6. Compare the time intervals when entity e (or group of entities E) moved according to pattern P_1 and according to pattern P_2.

These questions are targeted at *relations* between characteristics (questions 1 and 2), between behaviours (questions 4 and 5), between references (question 3), or between reference sets (question 6). Such questions are called comparison questions. Questions 1 to 3 are *elementary* comparison questions while questions 4 to 6 are synoptic comparison questions. The term 'comparison' is used in a quite broad sense as establishing relations between things. The nature of the things determines what relations are possible. Let us list the relations relevant to movement data.

1. Relations between characteristics (including both characteristics of movement and auxiliary characteristics)
   a) Positions: spatial relations including distance, direction, and topological relations such as touch, inside, overlap, etc.
   b) Numeric characteristics, e.g. speed, acceleration, angle of turn: equality (equal or not equal), order (greater than or less than), and distance (difference).
   c) Qualitative characteristics, e.g. direction of movement or character of a location: equality.
2. Relations between references
   a) Time moments: equality, order, distance (amount of time between two moments).
      i) Additionally, relations between the auxiliary characteristics of time moments (numeric or qualitative): see 1b and 1c.
   b) Entities: equality.
      i) Additionally, relations between the auxiliary characteristics of entities (numeric or qualitative): see 1b and 1c.
3. Relations between behaviours: equality (equal or not equal); similarity (similar or dissimilar); conformity (conformal or opposite).
   a) Additionally, relations between the summary characteristics of the behaviours such as the average speed (numeric) or prevailing direction (qualitative): see 1b and 1c.
4. Relations between sets of references
   a) Time intervals: temporal order, distance, topological relations such as touch, inside, overlap, etc.
      i) Additionally, relations between the auxiliary characteristics of the intervals such as length (numeric) or character of the times (qualitative): see 1b and 1c.
   b) Groups (subsets) of entities: equality, inclusion, overlap or absence of overlap.
      i) Additionally, relations between the auxiliary characteristics of the groups such as size (numeric) or character of the entities (qualitative): see 1b and 1c.

3.6 Relation search

In opposite to comparison questions, in which relations are unknown and need to be ascertained, there are questions requiring the search for occurrences of specified relations. In such questions, it is typically necessary to determine and describe the characteristics or behaviours linked by the specified relations and the corresponding references or reference subsets, i.e. where these relations occur. For example:

- Find all cases when two or more entities met in the same location (In what locations? What entities did meet? At what time moments?)
- Find all cases when two or more entities moved together, i.e. simultaneously passed the same locations (What sequences of locations, i.e. paths in space? What entities? On what time intervals?)
- Find all cases when an entity repeatedly made the same path in space (What path in space? What entity? On what time intervals?)
- Finds groups of entities that had similar movement behaviours (What is the common pattern for these behaviours? What entities? On which time intervals?)

From the examples given above, the first is an elementary question since it addresses individual characteristics (locations) and references (entities and time moments). The remaining examples are synoptic since they involve behaviours (in particular, paths in space) and reference subsets (in particular, time intervals). It may be noted, however, that both elementary and synoptic questions are not atomic but involve several tasks:

1. Detect an occurrence of the specified relation, i.e. at least two characteristics or at least two behaviours related in this way.
2. Find out what references or reference subsets correspond to the characteristics or behaviours thus detected.
3. For the answer to be complete, the characteristics or behaviours should also be described, in particular, the behaviours represented by suitable patterns.

3.7 Building an overall pattern

One of the major goals of the analysis of movement data is to characterise the overall movement behaviour of the whole set of entities over the entire time period the data refer to, or, in other words, to build an appropriate pattern representing this overall behaviour (in data mining and statistics, a pattern describing the entire set of facts has the special name *model*). "Appropriate" means adequate to the further goals, which may be, for example, prediction of the future behaviour or optimisation of the road network. The overall pattern (model) needs to be sufficiently comprehensive and precise. Typically, the required precision cannot be achieved in a simple ("atomic") pattern but rather in a compound pattern built from sub-patterns, which, in turn, may also be compound.

Compound patterns result from decomposing the overall behaviour into parts, representing these partial behaviours by sub-patterns, and, finally, bringing the sub-patterns together into an overall pattern. The decomposition is required because the movement behaviour is not uniform throughout the reference set. The decomposition is based on detecting similarities and differences, i.e. involves relation search and comparison questions. The following synthesis of the compound overall pattern involves pattern search and comparison questions (Where else does this sub-pattern occur? What is the relative position of these sub-patterns in time and in space? etc.). Since the data analysis aims first of all at building patterns and models, elementary questions play a marginal role in it, as compared with synoptic questions.

There are several approaches to the decomposition of the overall movement behaviour of a population of entities E over a time period $[t_0, t_{end}]$:

- Divide E into subsets of entities with similar behaviours; build a pattern for each subset; describe the subset each pattern is valid for.
- Divide the period $[t_0, t_{end}]$ into intervals where the behaviour can be regarded as homogeneous; build a pattern for each interval; describe the intervals and relations between them; describe the temporal arrangement of the patterns.
- Factorise the time into its component parts, i.e. the linear component and one or more cycles (yearly, weekly, daily, or other, domain-specific cycles); build a pattern for the behaviour with respect to each component.

In practice, these approaches are usually combined for the resulting model to be more precise. However, a full precision is hardly reachable. First, any pattern is a result of abstraction and simplification; the real data it represents usually slightly deviate from it. Second, extraordinary values and unusual value combinations may occur in a dataset or particular entities may behave in an uncommon way. Such outliers usually need to be analysed and described separately.

The division of the set of entities and/or of the time period may be done either on the basis of observed (or somehow else detected, e.g. computed) differences between the respective behaviours or according to expected differences, where the expectations come from the background knowledge. For example, one can expect that children behave differently from adults and from elderly people and that movements in the morning differ from those in midday and evening. Divisions according to expected differences often base on supplementary characteristics of entities and time moments.

3.8 Connection discovery

When studying a phenomenon, an analyst is interested not only in describing or summarising its behaviour but also in explaining it. The analyst wishes to find out the reasons or driving forces that make the phenomenon behave in the way observed. These forces may be internal or external. Internal forces originate from the inherent structure of the phenomenon and interactions between its structural components. External forces originate from interactions between the phenomenon and other phenomena. Hence, the goal is to determine what components and/or phenomena interact and how they interact. Thus, concerning the movement of entities, an analyst may be interested to know whether and how the movement is related to various spatial, temporal, and spatio-temporal phenomena such as weather, events (e.g. traffic jams or accidents), opening hours of shops, activities of people, etc. The analyst may also wish to detect interactions between parts of the overall movement behaviour, e.g. between the behaviours of traffic and of pedestrians, or between properties of movement, e.g. direction and speed.

We use the term *connection discovery tasks* to denote seeking for indications of possible interactions between phenomena or between different aspects of the same phenomenon. A result of such a task (or, in other words, an answer to a question about interactions) is a description in some language of the connection that has been discovered. We call such a description a *connection pattern* while a connection, or interaction, may be viewed as a 'mutual behaviour' of two or more phenomena or parts of the same phenomenon.

In data analysis, the following types of connections are typically looked for:

- *Correlation*: an undirected, or symmetrical, connection. This includes not only the statistical correlation between two numeric variables but also all cases of regular co-occurrence of characteristics or behaviours, possibly, with a temporal and/or spatial lag. For example, working in the centre of a city may correlate with using the public transport or a bike for getting to the workplace.
- *Dependency, or influence*: a directed connection; for example, the use of a car or a bike for getting to the workplace depends on the weather (or, in other words, the weather influences whether a car or a bike is used).
- *Structural connection*: an observed movement behaviour results from a composition of two or more different movements performed simultaneously, like the observed movement of the planets is the result of a combination of their own movement and the movement of the Earth.

Connection discovery tasks are synoptic since they require dealing with sets and behaviours rather than with elements and individual characteristics.

3.9 Taxonomy as a whole

Table 2 summarises the taxonomy of the analysis questions concerning movement data. The connection discovery tasks are listed separately below the table.

And here are the connection discovery tasks:

Time	Population: Elementary	Population: Synoptic
Elementary	• For given references (which include entities and time moments), find the positions and other movement characteristics. • For given movement characteristics, find the corresponding references. • Compare the movement characteristics of given references. • Find occurrences of given relations between movement characteristics and determine the references they correspond to.	• Describe the spatial distribution of the set of entities and the spatial and statistical distributions of the movement characteristics at a given moment. • Find time moments when the entities and/or their movement characteristics were distributed according to a given pattern (spatial or statistical). • Compare the distributions (spatial or statistical) of the entities and/or movement characteristics at given time moments. • Find time moments with similar distributions (spatial or statistical) of the entities and/or movement characteristics.
Synoptic	• Describe the movement behaviour of a given entity. • Find entities with the movement behaviour corresponding to a given pattern. • Compare the movement behaviours of given entities. • Find occurrences of similar movement behaviours. Find entities with behaviours similar to the behaviour of a given entity.	• Describe the collective movement behaviour of a given set of entities during a given time interval. • Find the entity subsets and time periods where the collective movement behaviour corresponds to a given pattern. • Compare the collective movement behaviours a) of different groups of entities during the same time interval, b) of the same entities during different time intervals. • Divide the data into groups of entities and/or time intervals so that the behaviours are sufficiently homogeneous within the divisions and substantially differ between the divisions; find outstanding behaviours.

Table 2. Types of analytical questions about movement data.

- Detect correlations and dependencies between different characteristics of the movement.
- Detect correlations and dependencies between the movement and various supplementary characteristics of the locations, time moments, and entities and/or various external phenomena and events.
- Represent the observable movement as a composition of several interacting movements of different kinds.

We have defined these question types purely by reasoning about movement data, irrespectively of any methods of analysis. It may be interesting to see how these types are related to the established typology of the tasks of data mining.

3.10 Relation to the data mining tasks

The left column of Table 3 lists the types of data mining tasks as defined in [4][12]. On the right, the corresponding general types of analytical questions are indicated.

As may be seen, the data mining tasks correspond to synoptic questions, in particular, behaviour characterisation (i.e. representation of a behaviour by a pattern), pattern search, relation search, and connection discovery. There are no specific data mining tasks for the synoptic comparison questions. This may be an indication of the need in complementing data mining with other methods for data analysis such as visual analytics methods [14].

It is also interesting to note a clear correspondence between most of the elementary question types and database queries; see, for instance, Guting and Schneider [5]. Again, there are no specific query constructs for elementary comparison questions, and hence additional methods are needed.

4 Conclusion

In this chapter, we have discussed the structure of movement data, the nature and properties of their component parts, and the things that may have an influence on movement and hence need to be accounted for in the analysis of movement data. Based on the treatment of data analysis as seeking answers to questions about data and underlying phenomena, we have also tried to define the possible types of questions about movement of a set of entities in space. The question types have been defined purely on the basis of the structure and characteristics of movement data, irrespectively of any existing methods and tools for data analysis. The resulting taxonomy of the question types should therefore be seen as a requirement for the set of methods needed for the analysis of movement data. This means that researchers should suggest appropriate methods and tool developers implement tools that will allow analysts to find answers to these types of queries. Such methods and tools are considered in the remainder of the book.

Data mining tasks	General question types
Clustering: Determining a finite set of implicit classes that describe the data.	Divide the data into subsets of entities and/or time intervals so that the behaviours are sufficiently homogeneous within the divisions and substantially differ between the divisions.
Classification: Finding rules to assign data items to pre-existing classes.	Detect dependencies between different characteristics of the movement. Detect dependencies between the movement and various supplementary characteristics of the locations, time moments, and entities and/or various external phenomena and events. Note: the definition of the classes may be based on movement characteristics (e.g. according to the movement direction), on supplementary characteristics (e.g. according to the activities of the entities), or on the variation of the external phenomena (e.g. according to the weather).
Dependency analysis: Finding rules to predict the value of an attribute on the basis of the values of other attributes.	Detect dependencies between different characteristics of the movement. Detect dependencies between the movement and various supplementary characteristics of the locations, time moments, and entities and/or various external phenomena and events.
Deviation and outlier analysis: Searching for data items that exhibit unexpected deviations or differences from some norm.	Find entities or subsets of entities and/or time intervals with behaviours notably differing from the rest of the entities and/or the time.
Trend detection: Fitting lines and curves to data in order to summarise the database.	Find the entity subsets and time periods with the collective movement behaviours corresponding to given patterns, which are specified through algebraic formulae.
Generalisation and characterisation: Obtaining a compact description of the database, for example as a relatively small set of logical statements that condense the information in the database.	Describe the collective movement behaviour of the population of entities during the whole time period.

Table 3. Comparison between the general question types and the types of data mining tasks.

References

1. *Merriam-Websters Collegiate Dictionary, Tenth Edition.* Merriam-Webster, Incorporated, 1999.
2. N. Andrienko and G. Andrienko. *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach.* Springer, 2006.
3. J. Bertin. *Semiology of Graphics. Diagrams, Networks, Maps.* University of Wisconsin Press, 1983.
4. U. M. Fayyad, G. Piatetsky-Shapiro, and P. Smyth. From data mining to knowledge discovery: An overview. In *Advances in Knowledge Discovery and Data Mining*, pages 1–34. 1996.
5. R. H. Gueting and M. Schneider. *Moving Objects Databases.* Elsevier, 2005.
6. T. Hagerstrand. What about people in regional science? *Papers of the Regional Science Association*, 24:7–21, 1970.
7. K. Hornsby and M. J. Egenhofer. Modeling moving objects over multiple granularities. *Ann. Math. Artif. Intell.*, 36(1-2):177–194, 2002.
8. M.-J. Kraak. The space-time cube revisited from a geovisualization perspective. In *21st International Cartographic Conference, Durban, South-Africa*, pages 1988–1995, 2003.
9. 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.
10. H. Miller. A measurement theory for time geography. *Geographical Analysis*, 37:17–45, 2005.
11. H. Miller. Modeling accessibility using space-time prism concepts within geographical information systems: Fourteen years on. In *Classics of IJGIS*, pages 177–182. CRC Press, 2006.
12. H. Miller and J. Han. Geographic data mining and knowledge discovery: an overview. In *Geographic Data Mining and Knowledge Discovery*, pages 3–32. Taylor and Francis, 2001.
13. A. Moore, P. Whigwham, A. Holt, C. Alridge, and K. Hodge. A time geography approach to the visualization of sport. In *Proc. of the 7th Intern. Conference on Geocomputation*, 2003.
14. J. Thomas and K. Cook. *Illuminating the Path. The Research and development Agenda for Visual Analytics.* IEEE Computer Society, 1983.
15. E. Tufte. *Visual Display of Quantitative Information.* Graphics Press, 1983.