A General Framework for Using Aggregation in Visual Exploration of Movement Data

Gennady Andrienko, Natalia Andrienko

Fraunhofer Institute IAIS (Intelligent Analysis and Information Systems), Sankt Augustin, Germany

http://geoanalytics.net/and

Abstract

To be able to explore visually large amounts of movement data, it is necessary to apply methods for aggregation and summarization of the data. The goal of our research has been to systemize the possible approaches to aggregation of movement data into a framework clearly defining what kinds of exploratory tasks each approach is suitable for. On the basis of a formal model of movement of multiple entities, we consider two possible views of movement data, situation-oriented and trajectory-oriented. For each view, we discuss the appropriate methods of data aggregation and the visualization techniques representing the results of aggregation and supporting data exploration. A special attention is given to dynamic aggregation working in combination with interactive filtering and classification of movement data.

CR Categories and Subject Descriptors: H.1.2 [User/Machine Systems]: Human information processing – Visual Analytics; I.6.9 [Visualization]: information visualization.

Additional Keywords: Movement data, spatio-temporal data, aggregation, scalable visualization, geovisualization.

1 INTRODUCTION

Visual representations of data are essential for enabling a human analyst to understand the data, extract relevant information, and derive knowledge. It is generally recognized that visual displays facilitate effective perception and cognition (McCormick et al. 1987), promote ideation (Dykes et al. 2005) and support analytical thinking (Thomas & Cook 2005).

One of the strengths of visual representations lies in aiding abstraction and generalization (Thomas & Cook 2005). Thus, appropriate positioning and/or appearance of visual elements representing data items can stimulate holistic perception of multiple data items as a unit. This mechanism fails, however, when a large size and/or complex structure of a dataset does not permit representing all data items in a sensible way, such that none of the visual elements is fully or partly covered by others. A common approach that helps in such situations is data aggregation.

Essentially, aggregation is combining several data items in a single unit. This, on the one hand, reduces the size of the data (which makes it easier to visualize and to perceive), on the other hand, promotes abstraction and generalization. Certainly, aggregation involves information loss. A positive side of this is an opportunity to distill general features out of fine-detail particulars. However, when particulars are (also) relevant, it is necessary to use (additionally) approaches not involving aggregation.

Data resulting from tracking of various kinds of moving agents and objects (animals, vehicles, vessels, aircrafts, pedestrians, visitors of public buildings, players and balls in sport games, etc.) are currently collected in growing amounts. Such data, further referred to as *movement data*, may be useful for many purposes. However, in order to understand how to use movement data and before doing any sort of computational analysis and/or deriving predictive models, an analyst needs first to look at the data and explore them. In other words, an analyst needs an appropriate visualization of movement data.

The visualization of large amounts of movement data is a hard problem. Traditional approaches such as delineating the traces of moving objects on a map or in a space-time cube (Kraak 2003) do not work because of tremendous overlapping and cluttering. Hence, aggregation of movement data is necessary for supporting visual exploration.

2 OBJECTIVES OF THIS STUDY

The topic of this paper is the use of aggregation for visual exploration of movement data. More precisely, we deal with data about multiple discrete entities changing their spatial positions over time while preserving their integrity and identity (i.e. the entities do not split or merge). Movement data consist of items called *position records*. A position record specifies the spatial position of some entity at some time moment. It may also specify the values of other attributes characterizing the movement (e.g. speed, course, transportation mode, etc.) and/or the state of the moving entity (e.g. heartbeat and blood pressure). It should be borne in mind that, for obvious reasons, position records can only exist for sampled time moments rather than all time moments. If the sampling is sufficiently fine, the intermediate positions and attribute values can be approximated by means of interpolation.

In our study we sought answers to two major questions:

(1) How movement data can be aggregated? This means: What approaches to aggregating movement data exist or may exist? What are the principal differences? When each of the approaches is appropriate?

(2) How the results of aggregating movement data can be visualized, depending on the way of aggregation?

By answering these questions we aimed at creating a systematic and comprehensive framework for the use of aggregation for visual analysis of movement data. The results of our study are presented in this paper, which extends an earlier conference paper (Andrienko & Andrienko 2008). We make a reservation that our framework is not meant to cover everything related to analysis of movement. It includes only aggregation methods for position data and visualization techniques suitable for aggregated position data. It does not include methods for statistical analysis and modeling. Such methods, in which visualization does not play a significant role, have been out of the scope of our research.

The body of the paper is structured as follows. In the next section we briefly overview the relevant literature concerning aggregation of movement data. After that we describe the formal model that helped us to organize the variety of approaches into a system. Then we introduce two example datasets, which are used in the remainder of the paper to facilitate the presentation and explanation of our framework.

3 RELATED WORK

One of the observations that can be made by studying the literature is that movement data are often aggregated using the same approaches as Fredrikson et al. (1999) suggested for another type of spatio-temporal data, namely, discrete events such as traffic incidents. Fredrikson et al. introduced three basic types of aggregation, spatial (S), temporal (T), and categorical, or attributive (A). For the spatial aggregation, the space is divided into suitable compartments. The events that occurred in the same compartment are united in an aggregate. For the temporal aggregation, the time is divided into suitable intervals. The events that occurred during the same interval are put together. The attributive aggregation unites events characterized by the same or close values of analysis-relevant attributes. For numeric attributes, the closeness of values is defined by dividing the value ranges into intervals so that all values within an interval are considered to be close. These three basic types of aggregation can be used in various combinations.

A number of research papers describe the application of these types of aggregation to the position records in movement data. In (Dykes & Mountain 2003, Mountain 2005), T-aggregation appears in the form of temporal histogram where the bars correspond to time intervals and their heights are proportional e.g. to the number of locations visited or the distance traveled. For S-aggregation, the territory is divided into compartments by means of a regular grid. The results of aggregation, such as density counts, are represented by coloring or shading of the grid cells on a map display. S×T-aggregation is done by the grid cells and consecutive time intervals. The results are shown on an animated map display. Analogously to densities, other aggregated characteristics can be computed and visualized. Thus, Forer and Huisman (2000) compute the total number of person/minutes spent in each cell. A sophisticated S×T×A-aggregation is suggested in (Wood et al. 2008). First, position records are grouped spatially by cells of a regular grid. Then, temporal (e.g. by days of the week) and attributive (e.g. by vehicle types) aggregation is applied to each group. The results are represented by treemaps (Shneiderman 1992) placed inside each cell.

In all these aggregations, the movement is, in fact, viewed as a set of independent discrete events, i.e. each position record is treated as representing an event of presence of some entity in some position at some time. This view does not fully capture the essence of movement as continuous change of spatial position; hence, it cannot be sufficient for a comprehensive analysis of movement data. Still, it may be helpful for certain types of analysis (sub)tasks, which need to be explicitly defined.

Another way of aggregating movement data described in the literature is based on considering the data as a set of *moves* between predefined places (typically, the places are spatial compartments). Each move is treated as a vector characterized by its origin and destination (i.e. the places where it starts and ends), by the start and end times, and, possibly, by additional attributes such as duration and length (traveled distance). Moves with coinciding origins and destinations are united into *aggregate moves*, which are characterized by the number of the original moves and by other derived attributes such as minimal, maximal, and average duration and length. This kind of aggregation may be represented by the formula S×S, where the two symbols 'S' stand for the place of origin and place of destination. The results may be visualized as a transition matrix where the rows and columns correspond to the places and symbols in the cells or cell coloring or shading encode the derived attribute values (Guo et al. 2006, Guo 2007). A disadvantage of such visualization is the lack of spatial context. Aggregate moves can also be visualized on a map by bands or arrows connecting pairs of locations (Tobler 1987, 2005). The widths of the bands or arrows are proportional to the volumes moved between these locations. Unfortunately, such a map may be illegible because of intersecting and overlapping symbols. Therefore, Tobler suggests a specific method for spatial smoothing of aggregate moves and generation of continuous flow maps.

In the S×S-aggregation, the moves are grouped only according to their origin and destination places irrespective of the time when they occurred. To take into account also the time, the whole time span is divided into intervals. Aggregates are built from moves having common origin, common destination, and common time interval when they occurred (which means that the start and end time of each move lie within this interval). This aggregation may be denoted by the formula S×S×T×T as it is done according to the place of origin, place of destination, time of start, and time of end of each move. The results of the S×S×T×T-aggregation can be represented, in principle, by a sequence of transition matrices or flow maps, one matrix or map per time interval. The matrices or maps can be put on the screen side by side, which facilitates comparisons but requires much screen space. Another possibility is to present them one by one in the animation mode, which may be less effective (Tversky et al. 2002).

A different approach to representing aggregate moves can be seen in the visualization of the movement of tourists in New Zealand (Drecki & Forer 2000) (discussed in Andrienko & Andrienko 2007). The moves are shown in a perspective view where the vertical dimension represents time divided into daily intervals. For each break between two successive time intervals, there is a horizontal plane where a map of New Zealand is drawn. The movements of the tourists are represented as lines connecting the locations of the major tourist destinations on successive planes. The brightness of a line corresponds to the number of people that moved from the origin location (on the upper plane) to the destination location (on the lower plane) during the respective time interval. To make the view clearer, the authors omitted minor flows.

In all aggregations discussed so far the results are numeric values such as counts, sums, statistical means, etc. Buliung & Kanaroglou (2004) derive a kind of geometric summary of several trajectories. The authors use functions of ArcGIS to build a convex hull containing the trajectories, compute the central tendency and dispersion of the paths, and represent the results on a map as the averaged path. Such geometric summarization can work well only when the trajectories are similar in shape and close in space. It can be applied, for example, to groups of similar trajectories resulting from clustering. Grouping of trajectories by similarity and/or closeness of the routes followed by geometric and/or numeric summarization may be called R-aggregation (i.e. route-based). The paper (Andrienko et al. 2007) contains examples of combining route-based grouping of trajectories by means of clustering with S×S- and S×S×T×T-aggregation. It should be noted that route-based grouping does not guarantee that each trajectory is put in some group as there may be trajectories whose routes significantly differ from all others. Hence, a result of R-aggregation may consist of aggregates and solitary trajectories. For the sake of uniformity, the latter can be represented as aggregates of magnitude one.

Hence, there is a variety of approaches to aggregating movement data:

- S-, T-, and A-aggregation and their combinations, in particular, S×T and S×T×A, which are applied to the position records treated as independent events;
- S×S- and S×S×T×T-aggregation, where the data are treated as straight moves between predefined places while the actual paths are ignored;
- R-aggregation, which is applied to trajectories with close and similar routes.

The formal model presented in the next section is meant to help in defining the applicability conditions of each approach and understanding what kind of analysis each approach permits.

4 THEORETICAL BACKGROUND

In our earlier papers (Andrienko & Andrienko 2007, Andrienko et al. 2008) we introduced a formal model of collective movement of multiple entities as a function μ: $E \times T \rightarrow S$ where E is the set of moving entities, T (time) is the continuous set of time moments and S (space) is the set of all possible positions. Another representation is $\mu(e,t)=s$, $e \in E$, $t \in T$, $s \in S$, which shows that e and t are independent variables and s is a dependent variable. The function μ can be extended to include also various movement attributes (speed, direction, etc.) and/or movement-related attributes of the entities such as heartbeat: μ: $E \times T \rightarrow S \times A_1 \times A_2 \times \ldots \times A_N$ or $\mu(e,t)=(s, a_1, a_2, \ldots, a_N)$. These additional attributes are dependent variables, analogously to the space. For the sake of simplicity, we shall use the basic form of the function μ in the following discussion; however, everything what is stated below applies also to the extended form.

As it is argued in (Andrienko & Andrienko 2006), in order to analyze data having two or more independent components, one may need to decompose the original function of several variables into multiple single-variable functions (of course, this should not be understood literally but as a metaphor for what is done in practice). For a function of two independent variables, there are always two possible decompositions.

Hence, as a function of two independent variables, μ can be decomposed in two complementary ways:

- $\{\mu_e: T \rightarrow S \mid e \in E\}$, where each function μ_e: $T \rightarrow S$ describes the movement of a single entity. We shall call the function μ_e the *trajectory* of the entity e. The decomposition of μ into a set of μ_e may thus be called *trajectory-oriented view*;

– $\{\mu_t: E \rightarrow S \mid t \in T\}$, where each function $\mu_t: E \rightarrow S$ describes the *situation* at a time moment t, consisting of the spatial positions (and, possibly, additional attributes) of all entities. The decomposition of μ into a set of μ_t may be called *situation-oriented view*.

Note that the set of μ_t is ordered according to the time. Theoretically speaking, it is continuous: a certain situation exists, in principle, for any element of the continuous set T. In practice, however, only a finite set of different situations can be retrieved from the data and explored.

Figure 1 gives a graphical illustration of the two possible views of the movement of multiple entities.

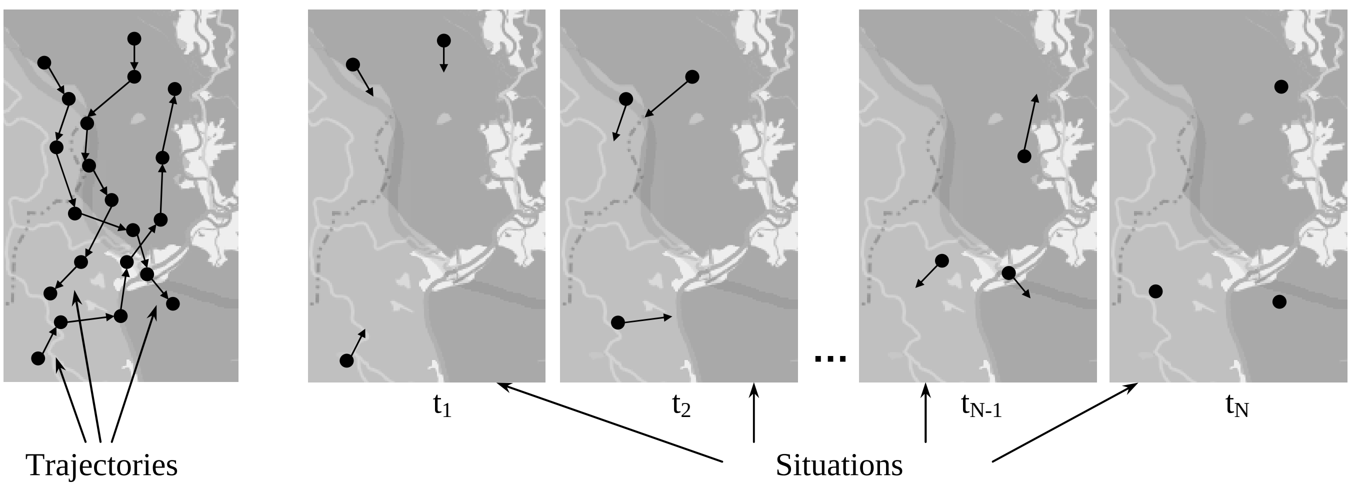

Figure 1. A graphical illustration of the two possible views of the movement of multiple entities.

Each of the views permits a different kind of analysis. Let us consider first the situation-oriented view. Each situation is a set of discrete entities (possibly, with some attached characteristics) distributed over space. In analyzing this spatial distribution and its change over time, one can be interested in two distinct things:

– the absolute positions of the entities with respect to the space, i.e. the presence and density of entities in different parts of the space, and, possibly, general characteristics of the movement in different parts;

- the relative spatial positions of the entities with respect to others, i.e. whether the entities form spatial clusters, alignments, or other arrangements and what entities are isolated.

These two foci of interest correspond, respectively, to the absolute and relative views of space (Peuquet 1994, 2002). According to the *absolute view*, space is an independently existing container where the entities are placed. According to the *relative view*, space is a positional attribute attached to the entities.

The absolute view focuses on space as the subject matter. Movement that occurs in space is thus considered as a property of the space. Respective analysis tasks are, for example, studies of the use of space, its accessibility and permeability (the ease or difficulty of moving in different directions), in particular, in urban planning. We shall use the term *space-centered* to denote the class of analysis tasks supposing the absolute view of space. Note that the identities of the moving entities are irrelevant for space-centered analyses and may be ignored. This means, in particular, that the data may be aggregated in such a way that multiple entities are handled together as a unit. Another implication is that the positions of the entities at different time moments (together with other attached characteristics) may be treated as *independent discrete events*. Hence, the aggregation methods suitable for independent discrete events (i.e. the S-, T-, and A-aggregation and their combinations) are applicable in this case.

The relative view, in contrast, focuses on the entities as the subject matter. Movement is considered as a property of the entities while space is treated as a collection of relationships between the entities. Analysis tasks requiring such a view will be called *entity-centered*. An example is the investigation of movement patterns in a population of animals: whether the animals tend to move in large or small groups, in pairs, or separately from others, whether the groups have leaders, whether they are arranged in a particular way (such as the V-shape of a flock of flying gees), etc. The identities of the moving entities are important in this case since it is necessary to trace the groupings and arrangements over time.

Hence, any kind of aggregation where multiple entities are merged into a single unit would not be suitable.

Let us now switch to the trajectory-oriented view of movement. The distinction between the absolute and relative views of space and, respectively, between space-centered and entity-centered analysis tasks also applies here. In the space-centered tasks, the analyst investigates the connectivity of the space (the ease or difficulty of getting from one place to another), the major flows, typical routes, the use of existing pathways, and so on. Examples of such tasks can be found in urban and transport planning and in ecological studies (landscape connectivity). As in the respective class of tasks within the situation-oriented view, the identities of the moving entities are irrelevant and, hence, it is suitable to apply aggregation techniques where multiple entities are merged. Thus, S×S- and S×S×T×T-aggregation can support the analysis of the space connectivity and its variation over time. R-aggregation can be instrumental in studies of the major flows, typical routes, and the use of pathways.

In entity-centered tasks, the analyst is interested in similarities and differences among the movement behaviors of the entities and in various kinds of *relative movements*: convergence, divergence, following, joint movement, parallel movements, opposite movements, and so on. Such tasks arise, for example, in studies of animal behaviors, in particular, interactions between different animal species. As with the situation-oriented view, aggregations where multiple entities are considered jointly are not suitable for entity-centered analysis tasks. However, as a trajectory of an entity is a complex spatio-temporal construct, it may be appropriate to aggregate in some way the trajectory of each entity and compare the aggregated trajectories of different entities, for example, traveled distances by time intervals (T-aggregation) or amounts of time spent in different places (S-aggregation). With the help of such aggregates, it is possible to investigate certain aspects of similarity or difference of the movement behaviors, but, apparently, analysis of relative movements requires other methods.

Table 1 summarizes our argumentation concerning the possible views of movement, respective analysis tasks, and suitable aggregation techniques.

Table 1. Classes of analysis tasks and applicability of aggregation techniques.

Class of tasks **View of movement**	**Space-centered** (movement as a property of space)	**Entity-centered** (movement as a property of entities)
Situation-oriented	Tasks: exploration of space use, accessibility, permeability Aggregation: S, T, A, S×T, S×A, T×A, S×T×A; applied to the position records treated as independent discrete events	Tasks: exploration of collective movement patterns (grouping, separation, spatial arrangement, etc.) Aggregation: current methods are not applicable
Trajectory-oriented	Tasks: exploration of space connectivity, major flows, routes, use of pathways Aggregation: S×S, R, R×S×S; applied to the entire trajectories or to fragments connecting predefined places - S×S×T×T, R×S×S×T×T; applied to the fragments made during the chosen time intervals	Tasks: comparison of the movement behaviors Aggregation: S, T, A, S×T, S×A, T×A, S×T×A; applied separately to the position records of each entity Tasks: exploration of relative movements and interactions Aggregation: current methods are not applicable

In the next section, we shall introduce two example datasets that will be used for the illustration of the aggregation methods listed in Table 1.

5 EXAMPLE DATASETS

The first example dataset results from GPS-tracking of 17,241 cars in Milan (Italy) during one week from Sunday to Saturday (the data have been kindly provided by Comune di Milano (Municipality of Milan) for the use within the project GeoPKDD). The dataset consists of more than 2 million records, which is too much not only for effective visualization and interactive exploration without the use of aggregation but even just for loading in the computer main memory. Each record includes car identifier, time stamp (date and time of the day), geographical coordinates, and speed. The time intervals between the records of the same car are irregular, mostly ranging from 30 to 45 seconds while there are also larger intervals ranging from several minutes to several days. We used the large intervals as dividers of the position sequence of each car into subsequences, which are treated as trajectories (more precisely, as discrete representations of trajectories). Generally, there is no unique way to build trajectories from a set of position records. The paper (Andrienko et al. 2007) discusses several possible methods; the division by temporal gaps is one of them. A temporal gap is a time interval between consecutive position records of the same entity with the length exceeding a specified threshold. For the Milan dataset, we chose the threshold 30 minutes, which produced about 176,000 trajectories.

The second example dataset is not so challenging in terms of the size but is suitable for illustrative purposes. The dataset consists of 1886 position records describing the seasonal migration of 19 white storks during the period from August 1998 till May 2006. The data have been collected in Vogelwarte Radolfzell (http://www.orn.mpg.de/vwrado_html/), a department of the Max Planck Institute for Ornithology, Germany. The time intervals between the records range from 1 day to 220 days, while the median time intervals are 1 or 2 days (the time spacing differs from animal to animal). To build trajectories, we divided the record sequences according to the migration cycle, so that a trajectory describes the path of a stork to the south in late summer and autumn and return back in spring next year, plus major movements in the southern regions during the winter period. As a result, we obtained 34 trajectories. The maximum number of trajectories per animal is 5 (2 of the 19 animals have been tracked during 5 consecutive years).

In the following sections we demonstrate the use of different aggregation methods and the visualization techniques suitable for viewing and exploring the outcomes of the aggregation.

6 SUPPORTING THE SITUATION-ORIENTED VIEW

We use the term “situation” to denote the spatial positions of all moving entities and the values of the movement-related attributes including speed, direction, acceleration (change of speed) and turn (change of direction) at some time moment. In the situation-oriented view, an analyst looks at situations at different time moments and considers the evolution of a situation over time. For practical reasons, the analyst cannot explore the situation of each second. On the one hand, this would require too much time and effort; on the other hand, the available data may not allow this because of larger time intervals between the measurements. A reasonable approach is to aggregate the data by time intervals of appropriate lengths.

S×T-aggregation can adequately support the consideration of aggregate situations on time intervals. Besides dividing the time into intervals, the space (i.e. the territory where the entities move) is divided into appropriate compartments. Then, various aggregates are computed for each pair of space compartment and time interval from the position records fitting in this compartment and this interval: number of different entities, number of visits, total time spent, statistics of the movement-related attributes (minimum, maximum, average, median, etc.). The aggregation can be done in the database and only the results are loaded in the main memory. In our prototype implementation, spatial compartments are defined by building a regular rectangular grid of a desired resolution. This is sufficient for our experimental and illustrative purposes, and the extension to arbitrarily shaped areas, such as traffic analysis zones (used in transportation planning) or road segments, is quite straightforward.

The results of the S×T-aggregation are time series of attribute values related to the spatial compartments. They can be visualized using various techniques suitable for time series, for example, animated map or time graph. For example, Figure 2 and Figure 3 demonstrate different visualizations of the Milan data aggregated by grid cells and hourly time intervals.

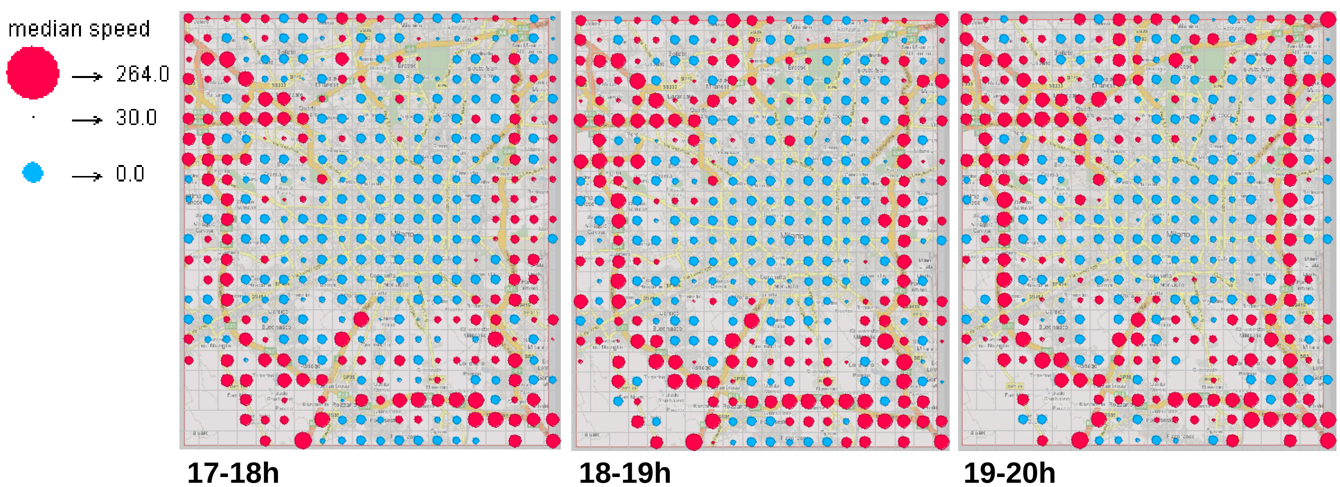

Figure 2. Three screenshots of an animated map representing results of S×T-aggregation.

Figure 2 presents three screenshots from an animated map. Each state of the map corresponds to one hourly interval and portrays the median speeds in the cells during that interval. We use the representation by graduated symbols in the mode of “visual comparison” (Andrienko and Andrienko 2006), in which the symbol size is proportional to the difference between the represented value and a chosen reference value (we have chosen 30 km per hour) and the color denotes the sign of the difference (in our example, red means positive and blue means negative). The animated map can be used for studying how the median speeds in different parts of the territory vary over time. Thus, it can be noted that there is a system of belt roads around the city where the median speeds are generally much higher than in the other parts of the city. However, in the interval 17-18 hours the speeds on the northern and eastern belt roads are quite low, which may indicate traffic congestions. On the east, the situation improves in the interval 18-19h. In the interval 19-20h, the median speeds increase on all belt roads and adjacent streets, but still there are places on the northeast where the speeds remain low.

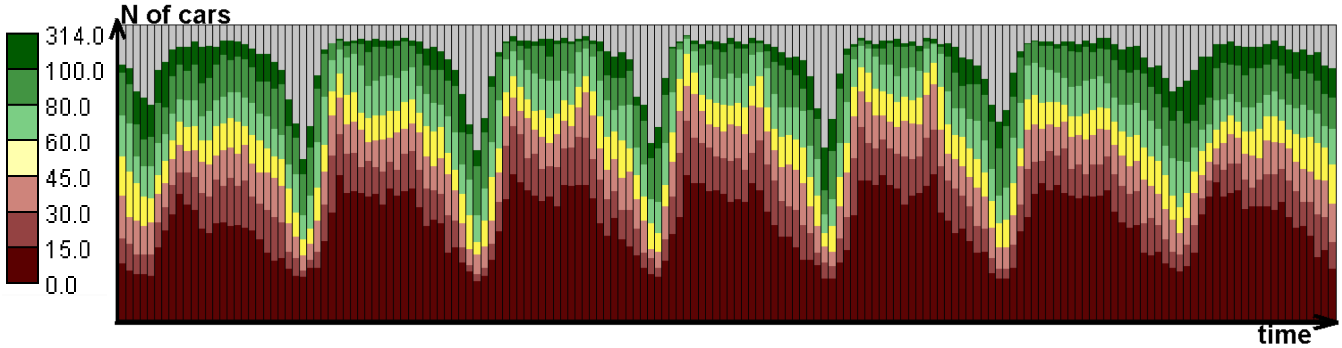

Figure 3. A temporal histogram showing the temporal variation of the frequency distribution of the values of median speeds by hourly time intervals.

In Figure 3, the aggregated data from all cells are put together in a temporal histogram, where each bar shows the frequency distribution of the values of median speed for one hourly time interval. The whole diagram depicts the variation of the frequency distributions over the week, from Sunday to Saturday. The colors of the bar segments correspond to intervals of the values of the summary attribute "median speed". The chosen breaks are 15, 30, 45, 60, 80, and 100 km/h. Yellow is assigned to the interval from 45 to 60, the shades of red represent median speeds below 45, and the shades of green are used for median speeds over 60 km/h (the color legend can be seen on the left of the histogram). The heights of the bar segments are proportional to the numbers of the compartments where the median speeds fitted in the respective intervals. Gray segments show the numbers of the compartments with no occurrences of tracked cars during the corresponding time intervals. The temporal histogram demonstrates a clear periodic pattern. The number of compartments with low speeds (red segments) is especially high in the morning hours, then slightly decreases towards the midday, but increases again in the afternoon, and then rapidly decreases. It may be noted that the speed variation patterns of the days from Monday to Thursday (days 2-5) are similar to each other and different from the patterns of Sunday, Friday, and Saturday (days 1, 6, and 7). The Friday pattern is more similar to the weekend pattern than to the workday pattern. After detecting this unexpected peculiarity, we found out that that particular Friday was a Friday before Easter, which may explain its similarity to a weekend.

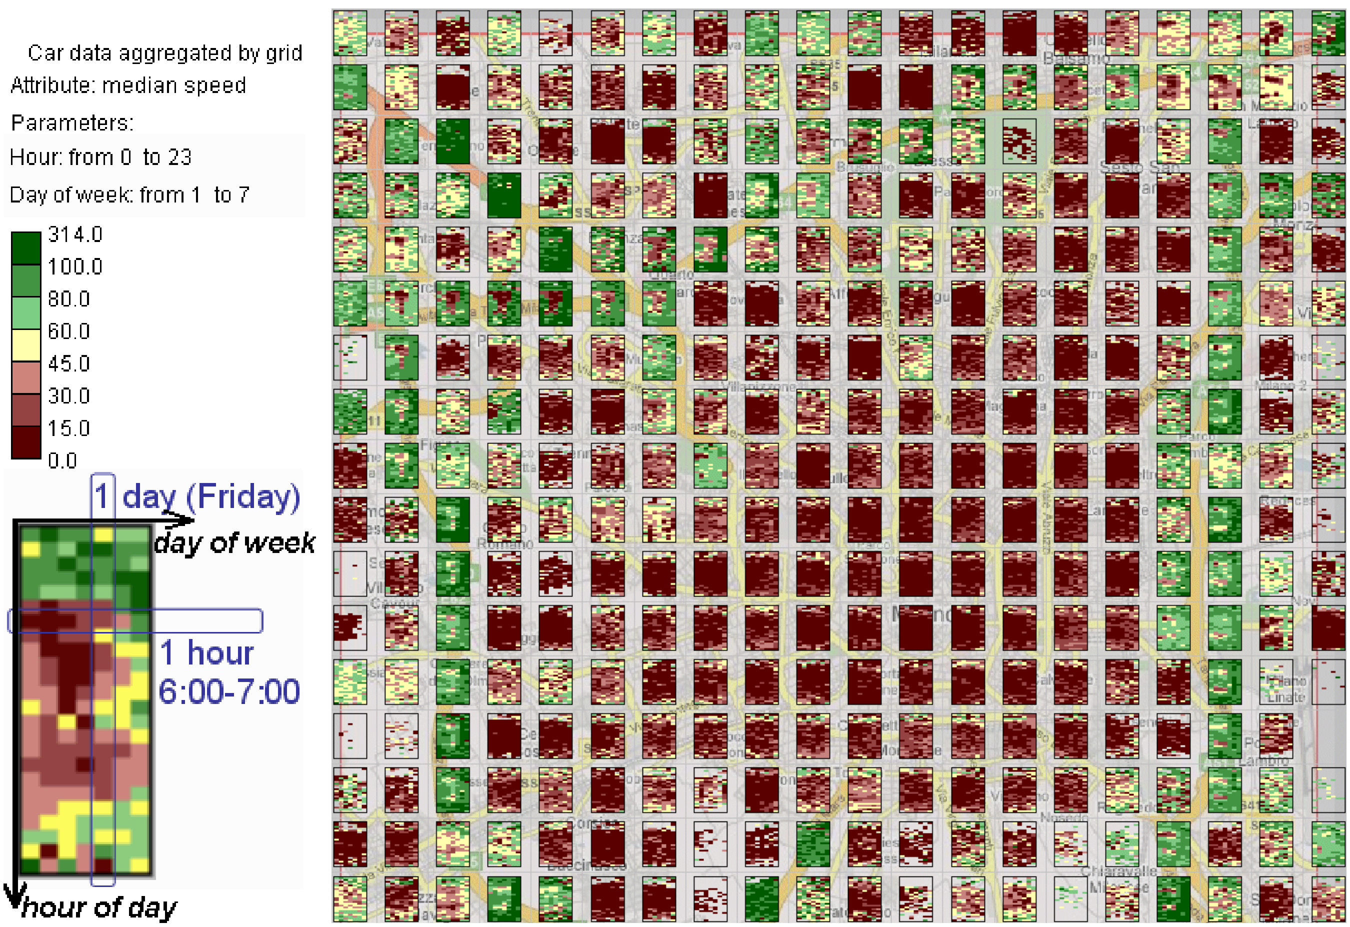

Figure 4.   The mosaic diagrams show the variation of the median speeds in spatial compartments by days of the week (columns of the diagrams) and hours of the day (rows of the diagrams). The cells

are colored according to the speeds. The breaks and colors for the speed intervals are the same as in Figure 3. Slow speeds are shown in shades of red and fast speeds in shades of green.

Since the temporal variation of the traffic characteristics clearly depends on the daily and weekly temporal cycles, it is reasonable to explicitly take this into account in the aggregation. For such a kind of aggregation, the time span of the data is divided into intervals according to one or more temporal cycles: times of the day, days of the week, and/or months of the year, whatever is appropriate. Thus, Figure 4 represents aggregates obtained with the use of two temporal divisions: according to the days of the week and according to the hours of the day. The first division groups together data referring to the same day of the week irrespective of the date. The second division groups together data from different days referring to the same hour of the day. As a result, aggregated values have been computed for each combination of space compartment, day of the week, and hour of the day. The visualization by "mosaic" diagrams shown in Figure 4 is especially suitable for data aggregated according to two temporal cycles. Each diagram summarizes the daily and weekly patterns of the traffic in the corresponding compartment. The columns of the diagrams correspond to the days of the week and the rows to the hourly intervals of the day. The colors of the "tiles" of the "mosaics" encode the values of the median speed in the respective days of the week and hours of the day. The color encoding is the same as in Figure 3: green corresponds to fast speeds and red to slow speeds (in particular, dark red represents speeds below 15km/h). It may be seen that the speeds are always very low in the inner city except for the northwestern part. The diagrams in the compartments on the western and eastern belt roads are mostly green but contain red or yellow spots indicating decreased speeds. Thus, in the northern part of the eastern belt road the red spots in the diagram indicate very low speeds in the morning hours of the work days. In some compartments on the northwest the speeds are also low in the mornings of the work days but also in the midday and afternoon on Wednesday and Thursday. On the northern belt road, especially the eastern part of it, the situation appears quite bad. In some compartments, the speeds are very low during the whole day and increase only in the night.

Figure 5 demonstrates the application of the S×T-aggregation to the data about the seasonal migration of the white storks. The data have been aggregated by months irrespective of the years, i.e. the records for the same month from different years have been put together. In the bar diagrams, each bar corresponds to one month, and the height is proportional to the number of birds that visited the respective compartment in that month. The bars of different months have different colors. In particular, the shades of orange and red correspond to months from August to October (beginning of the migration season), the shades of blue and violet correspond to the winter months, and the shades of cyan and green to the spring months.

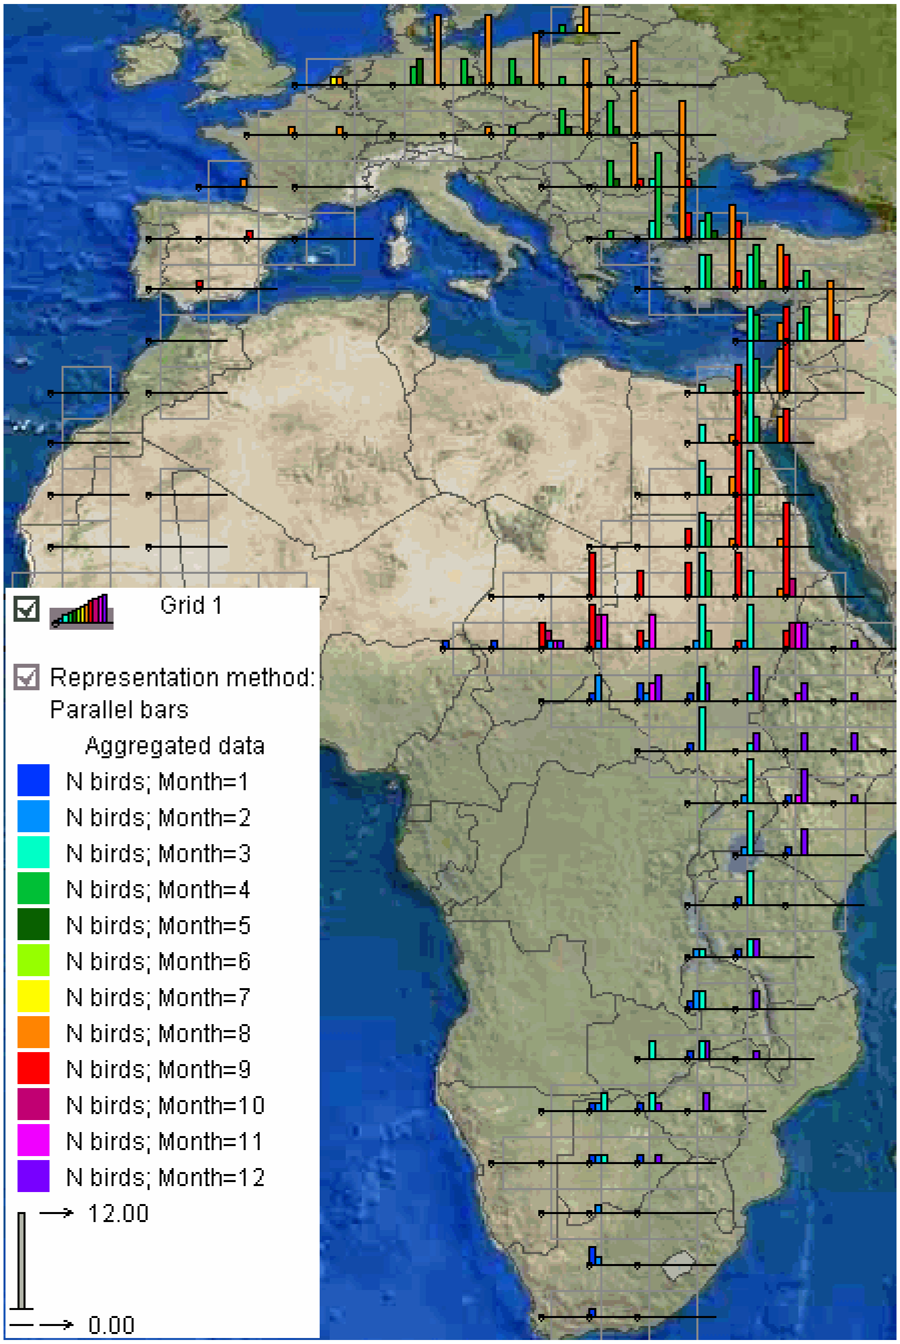

Figure 5. The bar diagrams show the presence of white storks in spatial compartments (grid cells) by months irrespective of the year.

The aggregations we have demonstrated so far do not capture such an important characteristic of movement as the direction of the movement. We suggest a method where the data are aggregated not only by space and time but also by the direction (course). This kind of aggregation can be denoted by the formula S×T×D, where D stands for “direction”. Movement directions are often indicated in the original track records. If this is not the case, they can be computed from pairs of consecutive positions of the same entity.

The directions are specified in movement data as numeric values typically representing angular degrees from 0 to 359. For the S×T×D-aggregation, we suggest to divide this range into intervals corresponding either to four main compass directions (north, east, south, and west) or to four main and four intermediate directions. Position records fitting in the same spatial compartment and temporal partition are additionally grouped by the movement directions. A separate group is made from records where the speed is below a chosen threshold. This is treated as the absence of movement (in real data consecutive position measurements are never exactly the same even then the object does not move, which means speed values of non-moving objects may differ from zero). Then, various counts and statistics of attribute values are computed for the groups.

To visualize the resulting aggregate data, we suggest a special technique in which the data are represented on a map by directional bar diagrams. Analogously to the *wind rose* used in meteorology, the bars are oriented in four or eight compass directions and their lengths are proportional to the values of the currently selected summary attribute corresponding to the respective directions. We demonstrate this technique by the example of the stork data (Figure 5). As in the previous example, the data have been aggregated by grid cells and by months, and additionally by eight compass directions. The speeds below 1 km/h have been treated as absence of movement. Figure 5 contains screenshots from an animated map corresponding to four months: August, September, March, and April (these were the months of the most intensive movement). The represented summary attribute is the number of distinct birds. The bars in the directional bar diagrams are colored depending on their orientation; a particular color is assigned to each direction. This helps in gaining an overall view of the prevailing movement directions in different parts of the territory. Besides the directional bars, some diagrams include gray circles representing the groups of records with the speeds below the chosen threshold. The radii of the circles are proportional to the values of the currently selected summary attribute computed for these groups of records. The radii can be easily compared with the lengths of the bars.

From the screenshots in Figure 5 it can be easily seen that in August most movements occurred over the territory of Europe and Middle East towards southeast and south. In September most movements were over Middle East and northeastern Africa towards south and southwest. In the areas closer to the equator there were more movements towards the west while some birds stayed without major movements. In March we can see massive movement over Africa towards the north; there are also northeastern and northwestern movements. In April we see movements towards the north in northeastern Africa and movements towards the west and northwest in Turkey and Eastern Europe.

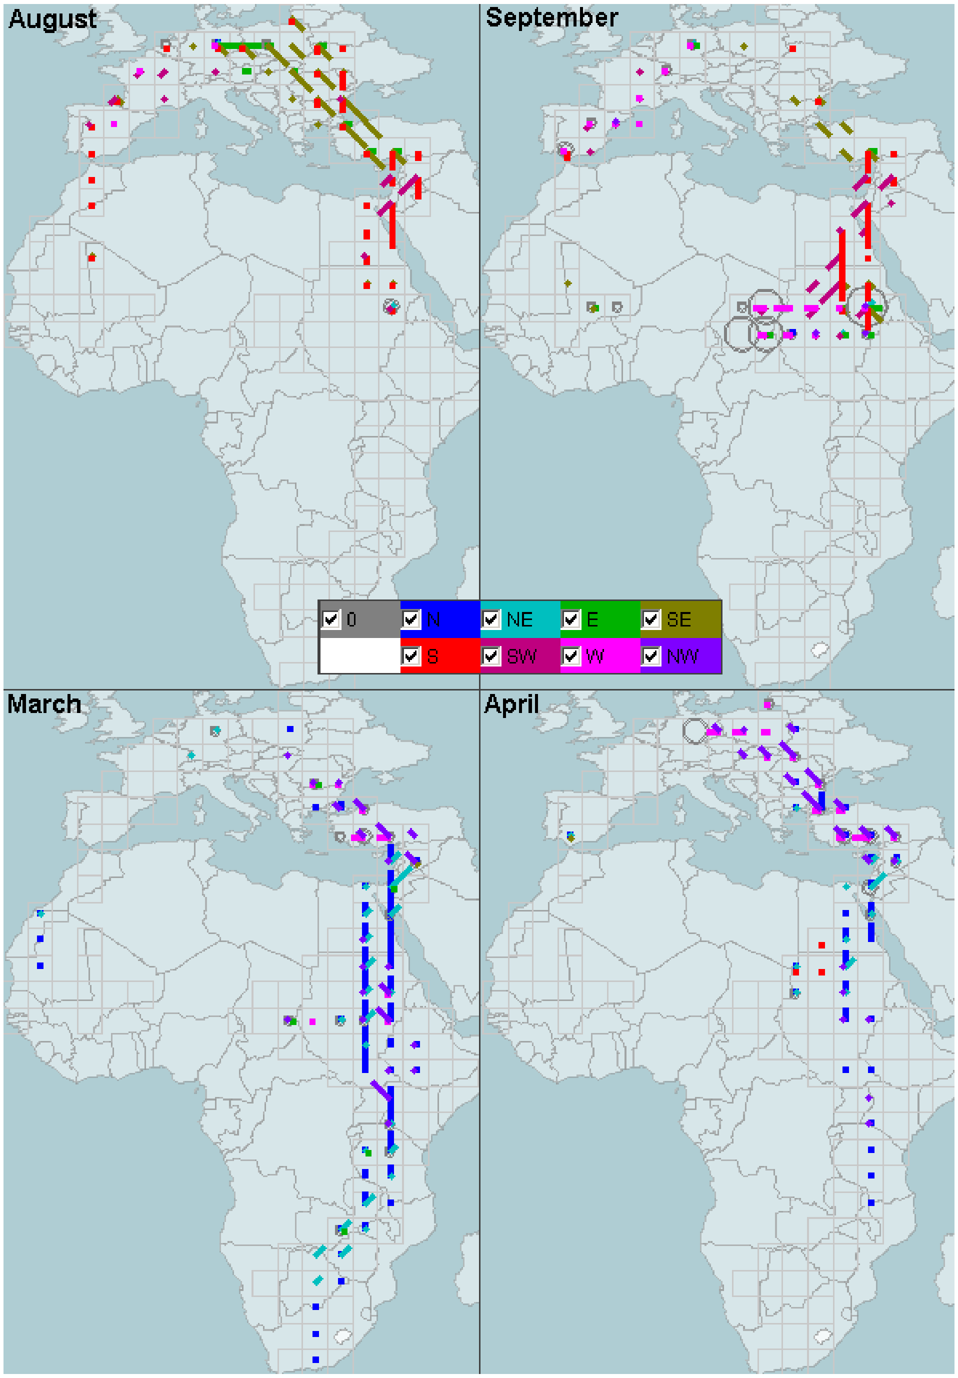

Figure 6. The directional bar diagrams show the stork movement data aggregated by grid cells (S), months (T), and eight compass directions (D). The lengths of the bars are proportional to the numbers of the birds that moved in the respective directions. The radii of the circles are proportional to the numbers of the birds with the speeds below a selected threshold (here 1km/h).

Visual exploration of movement with the use of this kind of display technique can be supported by a number of interactive facilities:

– switch from one summary attribute to another, e.g. from the number of the moving entities to the average or median speed;

– select another temporal partition, i.e. another interval, month, day of the week, time of the day, etc., depending on how the data have been aggregated;

– hide some directions in order to focus on the remaining direction(s), e.g. to see where northward movement occurs;

– choose the mode of presenting only the dominant direction(s) in each spatial compartment. A direction is treated as dominant when the corresponding value of the current summary attribute exceeds the highest value among the remaining directions by a chosen threshold, which may be either absolute (i.e. minimum difference between the values) or relative (i.e. minimum ratio).

In the case of city traffic data, the S×T×D-aggregation together with the visualization can support a detailed exploration of the traffic flow along a selected street. An example is shown in Figure 7. To look at the traffic on the northern belt road of Milan, we have selected only the space compartments (grid cells) covering this road. The data have been aggregated according to the four main compass directions. The bar diagrams represent the median speeds in the eastern (green) and western (purple) directions. The diagrams are substantially asymmetric, which means different speeds of the movement in the eastern and in the western directions. Lower speeds, in turn, signify higher obstruction to the movement.

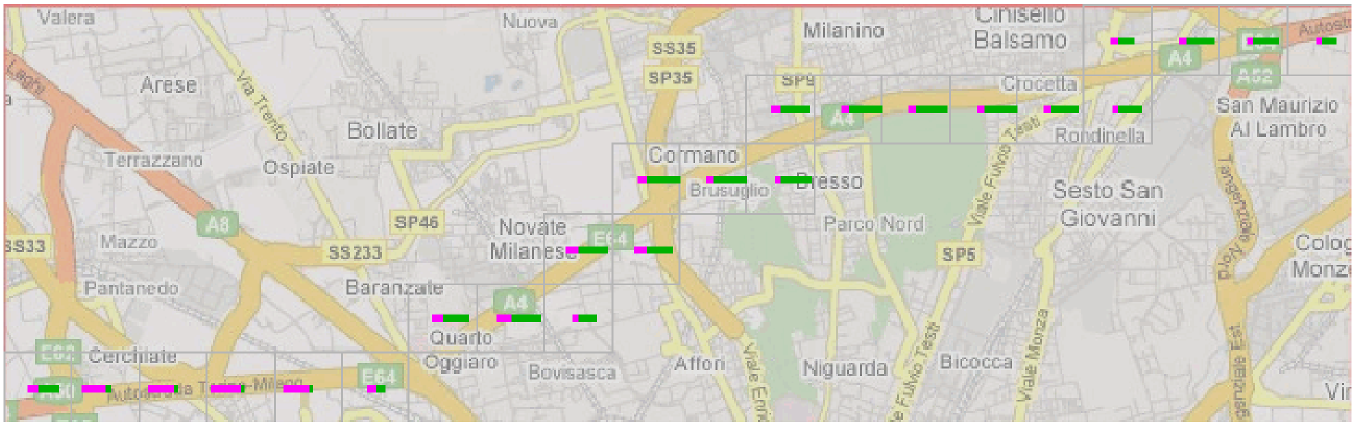

Figure 7. The bars represent the median speeds of the movement toward the east (green) and west (purple) between 11 and 12 AM on Wednesday along a motorway on the north of Milan.

7 SUPPORTING THE TRAJECTORY-ORIENTED VIEW

In the trajectory-oriented view, collective movement of multiple entities is considered as a set of trajectories of the entities. In practical tasks, the entire trajectory of each entity made during the whole period of the observation is usually divided into parts representing different trips of this entity; the term "trajectory" is also applied to such a part.

In analyzing trajectories, one may be interested in the origins and destinations of the trips, routes, start and end times, durations, distances, variation of the speeds along the routes, intermediate stops, etc. When trajectories are numerous, it is impracticable to examine each of them in detail. They need to be aggregated in such a way that the distribution of the relevant properties over the set of trajectories could be seen. For certain properties, the aggregation may be quite traditional. Thus, a frequency histogram can appropriately represent the distribution of the trip durations or distances. More specific aggregation and visualization techniques are required for the spatial properties (origins, destinations, and routes) and for the spatio-temporal properties (speed variation and intermediate stops).

The general approach is to group the trajectories by similarity in terms of the properties relevant to the current focus of the exploration. Then, the groups need to be represented in a summarized way, which appropriately conveys the relevant properties. The easiest case is when the analyst focuses only on the origins and destinations of the trips, for example, for exploring space connectivity. In such a case, the trajectories may be grouped by their origins and destinations.

7.1 Aggregation by origins and destinations

In this method, which may be called S×S-aggregation, two approaches are possible. One is to refer the starts and ends of the trajectories to predefined areas of interest, for example, city districts. Then, for each pair of areas, the trajectories starting in the first area and ending in the second area are grouped together. This applies also to the pairs where the first element coincides with the second one. The other approach is to define areas on the basis of spatial clustering of the start and end points of the trajectories. It is reasonable to assign meaningful names to the resulting clusters so that they could also be used as the names of the origins and destinations of the trips.

For each group of trajectories with a common origin and a common destination, the group size and the statistics (minimum, maximum, mean, median, etc.) of the numeric properties of the trajectories such as trip durations and distances are computed. The results may be displayed in the form of origin-destination matrix where the rows correspond to the origins, columns to the destinations, and the cells contain the values of the computed aggregates. The values in the cells may be represented visually by graduated symbols or diagrams. In our experimental software, the matrix display is linked to a map: clicking on a row, column, or cell highlights the corresponding areas on the map.

Another possibility for the aggregation is to account not only for the areas where a trajectory starts and ends but also for all intermediate areas visited by the trajectory. This means that each trajectory is generalized into a sequence of *moves* between areas. A *move* is defined as a tuple (p_1, t_1, p_2, t_2), $t_1 < t_2$, where p_1 and t_1 are the place and time of the start and p_2 and t_2 are the place and time of the end. An *aggregate move* combines moves with the same place of the start and the same place of the end. It is characterized by the number of the elementary moves it combines and various statistics of the duration, distance, speed, time, etc. computed from the respective trajectory fragments. These characteristics can be visualized in a matrix display like in the case of complete trajectories. The aggregate moves can also be shown on a map display as directed lines (vectors) between areas. The widths of the lines may represent the values of a selected summary attribute.

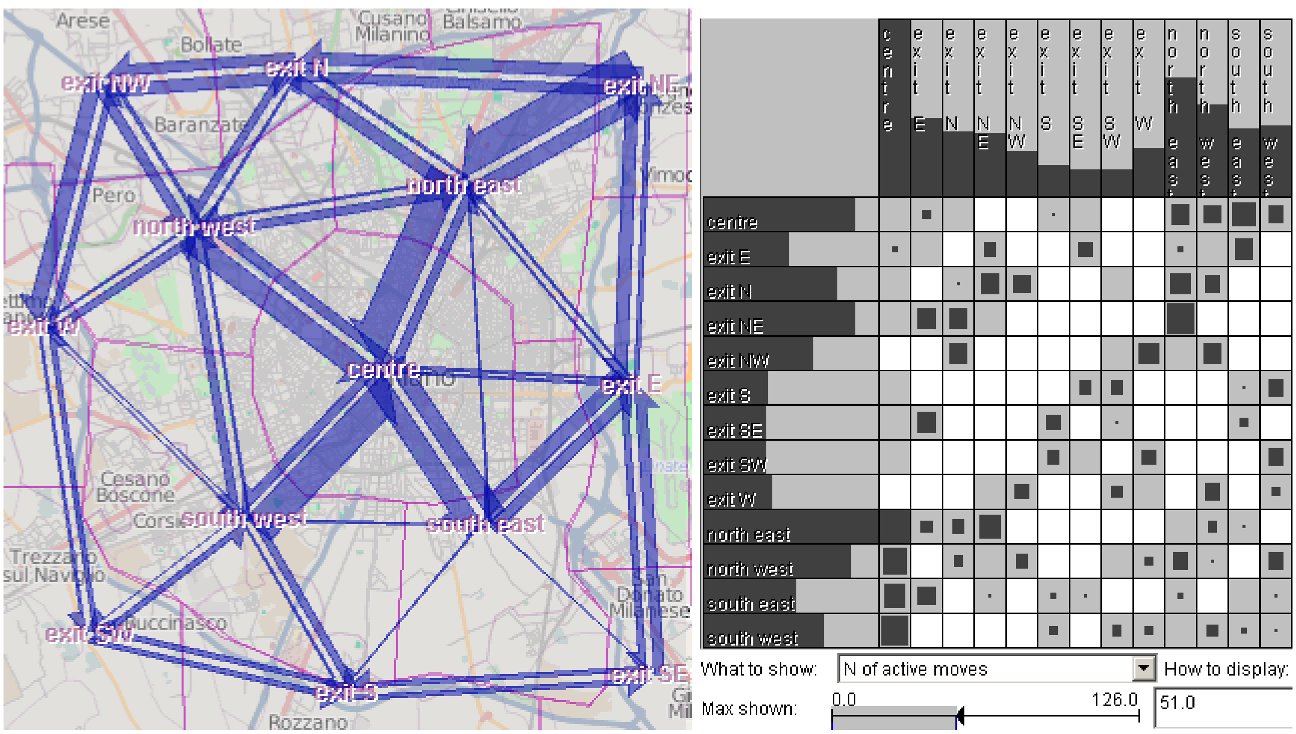

Figure 8. Summarization of trajectories into aggregate moves.

To investigate and compare the movement between the places in different time periods, the S×S×T×T-aggregation is used. The time is partitioned into linearly ordered intervals or according to temporal cycles. The S×S-aggregation, as described above, is then applied separately to the trajectories or fragments of trajectories made during each of the temporal partitions. For example, Figure 8 presents the screenshots of a map and an origin-destination matrix representing the movement between regions in Milan in the time from 5 to 6 AM on Wednesday. The regions “centre”, “north west”, “exit NW”, “exit N” and so on have been defined interactively by drawing their outlines on the map display. On the map, the widths of the arrow symbols represent the numbers of the moves between the regions connected by the arrows (the particular shape of the arrow symbols is used for a better differentiation of opposite moves). The same information is shown in the cells of the origin-destination matrix by sizes of the squares. It can be seen that the highest number of moves in the selected time period is from “north east” to “centre”; close to it are the numbers of moves from “south west” to “centre” and “exit NE” to “north east”. The dark grey shading in the first column of the matrix display, which contains the names of the regions, represents the total number of the moves originating in each region. The highest number of moves originates from “north east”. Analogously, the shading in the caption of the matrix display represents the total numbers of the moves ending in each of the regions. The most frequent destination of the moves in the selected time period is “centre”.

The S×S×T×T-aggregation described in this subsection supports the analysis tasks where the routes used for getting from place to place are irrelevant, such as analyses of space connectivity and flows among places. For analysis tasks where routes are relevant, it is necessary to have methods for grouping trajectories according to the routes and for presenting the routes in a summarized way.

7.2 Aggregation by routes

In all aggregations discussed so far it is possible to specify in advance the groups to be produced in terms of the properties of their members. Thus, in S×T-aggregation of track records, the groups are defined in terms of the spatial positions (which must fit in predefined space compartments) and time references (which must fit in predefined temporal partitions). In S×T×D-aggregation, the intervals for the values of movement direction are additionally specified. In S×S-aggregation of trajectories, the groups are defined through pairs of areas in which the origins and destinations of the trajectories or their fragments must fit. In S×S×T×T-aggregation, predefined temporal partitions for the start and end times of the trajectories or fragments are added.

In grouping by routes, it may not be possible to pre-specify a finite number of prototype routes for putting trajectories into groups based on their similarity to this or that prototype route (unless the use of existing pathways is investigated). When the possible routes are not known in advance, the trajectories may be grouped by means of clustering. In (Andrienko et al. 2007) a clustering tool is described that is capable of using different measures of similarity between trajectories, also called distance functions. One of the distance functions described in the paper computes the average distance between corresponding points of two trajectories. It can be used for clustering of trajectories by similarity of their routes. It should be noted that a clustering algorithm does not necessarily put each trajectory into some group. When a trajectory is not similar enough (in terms of a given distance threshold) to a certain minimum number of other trajectories, it may be treated as "noise", i.e. not included in any cluster, which seems quite natural. Hence, clustering is not a tool for aggregating all trajectories but a tool for finding groups of similar trajectories.

A detailed discussion of the use of clustering in analyzing trajectories can be found in (Rinzivillo et al. 2008). Here we shall focus on the topic of representing a group of similar trajectories in a summarized way for enabling interpretation by an analyst. A possible approach, which has been mentioned in the literature review, is to build an envelope around a group and, additionally, show the central tendency, i.e. a kind of "average trajectory" (Buliung & Kanaroglou 2004). A disadvantage of this approach is that it does not provide information about the intra-group variance.

Our idea is to represent groups of trajectories by aggregate moves between small areas. This is similar to what is described in the previous subsection except that the areas are not pre-specified but defined automatically using characteristic points of the trajectories, i.e. starts, ends, turns, and stops. The areas are built as circles around clusters of characteristic points from multiple trajectories and around isolated points. The radii vary within a user-specified range. Note that the areas so produced play an auxiliary role and do not need to be visualized (usually they are numerous and clutter the display). Let us demonstrate by an example how this approach can allow the analyst to see both the commonality and the variance among trajectories of a group.

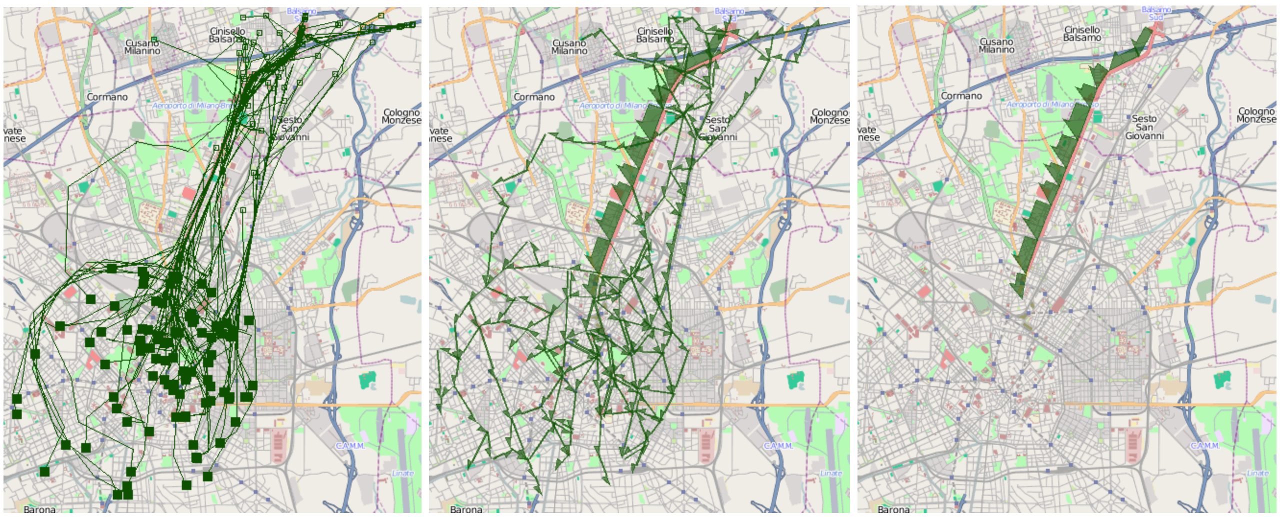

Figure 9. Left: an example of a cluster of trajectories grouped by the routes. Middle: Aggregate moves summarizing the cluster. Right: Only the major flows (specifically, the moves occurring in at least 30 trajectories) are visible.

The screenshot on the left of Figure 9 shows an example of a cluster of trajectories. The trajectories are represented as lines; the hollow small rectangles mark their beginnings and the bigger filled rectangles mark the ends. A high variability of the trajectories in the cluster is noticeable whereas the main characteristic features of the cluster as a whole are difficult to grasp.

In the middle of Figure 9 a result of summarization of the cluster is shown. The trajectories have been summarized into a collection of aggregate moves between small areas, which have been automatically built around the characteristic points of the trajectories (the areas have been hidden to reduce the clutter). The aggregate moves are represented by arrows, which indicate the directions of the movement. The thickness of an arrow is proportional to the number of trajectories in which the respective move occurs. In this summarized representation we can notice that a large part of the movements occurred along one of the major streets from the northeast to the center. However, the numerous thin arrows clutter and confuse the view.

In visualization of movement data, it is common to omit minor flows for the sake of clarity (e.g. Drecki & Forer 2000, Tobler 1987). In a computer-based visualization, this can be done in an interactive way. By hiding minor aggregate moves through interactive filtering, one can see more clearly what is in common between the trajectories in the group. Thus, on the right of Figure 9 only the aggregate moves occurring in at least 30 trajectories are shown. One can clearly see that the major part of the cluster of trajectories go from the northeast of Milan towards the center along the same street.

By interactively changing the filter, the analyst may control the amount of visible detail and thereby gradually build a comprehensive understanding of the cluster. Thus, by lowering the limit for the minimum number of trajectories in which a move must occur, the analyst may detect smaller flows going parallel to the principal flow as well as branching and merging of flows (Figure 10).

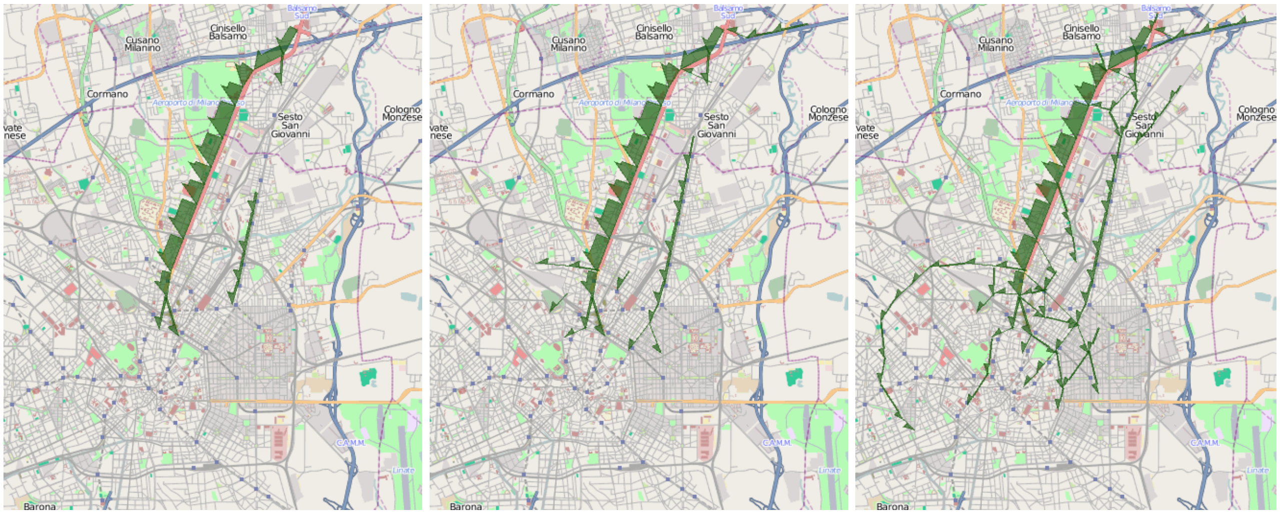

Figure 10. Interactive filtering of aggregate moves allows the analyst to control the level of detail in the summarized representation of the cluster. From left to right: the moves occurring in at least 15, 8, and 4 trajectories, respectively.

As may be noted from the screenshots in Figure 9 and Figure 10, the arrows representing the aggregate moves not always exactly fit in the streets of the city (although the fit is mostly quite good). For a better fit, it might be reasonable to take into account the background geographical information in building the areas to be used as the starts and ends of the aggregate moves. Thus, in exploring the city traffic, it is appropriate to build areas around street crossings and highway exits. The development of an automated procedure for generating such areas is out of the scope of our current research. We would only like to note that it would not be suitable to use all existing street crossings of a city, which are usually very numerous and very dense in space. An attempt to take all street crossings into account would result in a very low level of aggregation and generalization and excessive level of detail. Hence, an intelligent method for selecting appropriate street crossings would be needed.

Figure 11 demonstrates that the suggested summarization method works well also for the trajectories of the white storks. Note that this is quite a different kind of movement than in the case of the cars, which only move along the streets. Hence, an indubitable advantage of this approach is its wide range of applicability, including constrained and free movement.

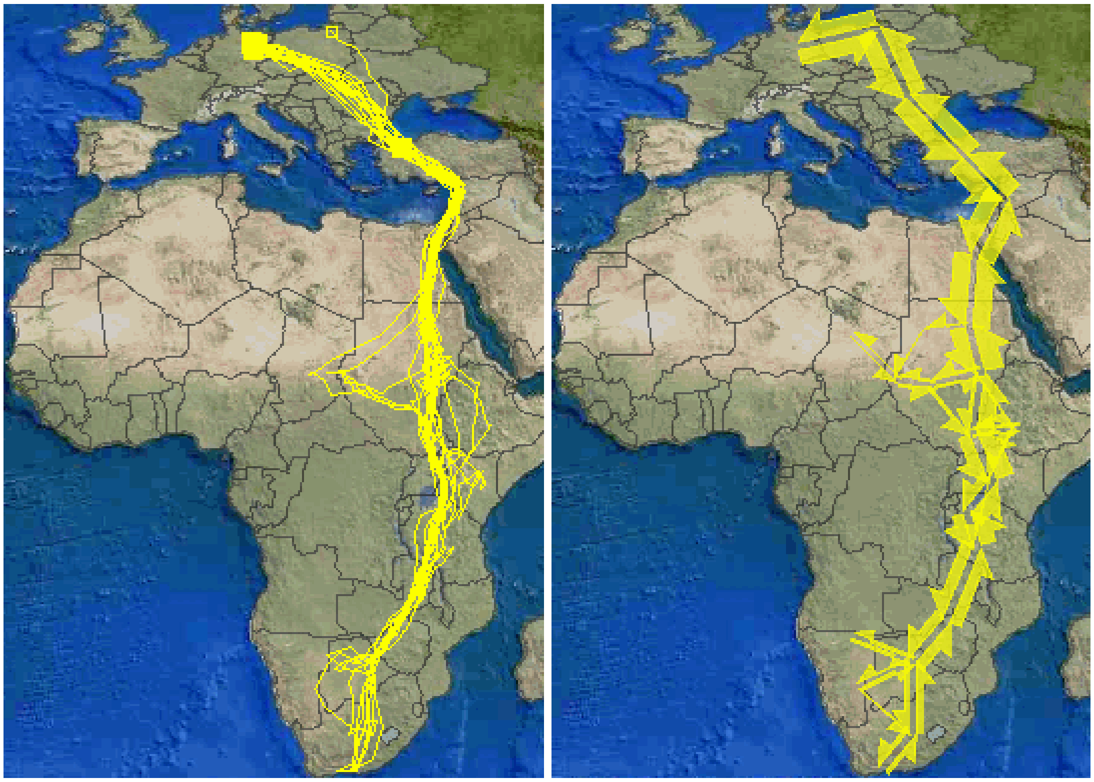

Figure 11. A cluster of similar trajectories of white storks. Left: the trajectories comprising the cluster. Right: the cluster in the summarized form.

The approach to summarizing groups of similar trajectories presented in this subsection is, essentially, S×S-aggregation. The combination of route-based grouping and the S×S-aggregation may be denoted as R×S×S-aggregation. To take into account also the temporal aspect, the approach can be extended to R×S×S×T×T-aggregation, where the time span is divided into appropriate intervals, the trajectories are divided into fragments made during these intervals, and then R×S×S-aggregation is applied to the fragments corresponding to each of the intervals.

7.3 Some open problems in aggregation by routes

The illustrations in Figure 9 and Figure 10 show how a single group of similar trajectories can be explored. It is yet an open problem how to support the exploration of multiple groups of trajectories. The source of the problem is that trajectories are not disjoint in space; they intersect and overlap. As a consequence, summarized representations of groups of trajectories also intersect and overlap when drawn on the same map display. In an arbitrary set of trajectories, like in the Milan dataset, there may be multitudes of different routes and, hence, numerous groups of similar trajectories. Putting them all together on the same map or in a space-time cube results in a completely incomprehensible picture. In fact, even a few groups of trajectories may be hard to explore together.

On the other hand, it does not seem realistic that an analyst can consider hundreds of clusters one by one. A more reasonable scenario is that the analyst has a certain focus of interest, for example, typical routes towards the centre of Milan or between two city districts. The analyst would apply clustering only to the trajectories corresponding to his/her focus and then explore only big clusters.

It would be very helpful for the analyst to have an overview display of all clusters in a "small multiples" style, where each cluster is represented on a separate map. Since each of these maps has to be quite small, the clusters need to be represented in such a way that only the principal features of each cluster are visible, but these features are very easy to grasp. The idea demonstrated in Figure 9 does not suit well to the purpose: when a map is zoomed out, the arrows representing aggregate moves become too short to be legible. A more schematic representation with much lower level of detail is needed. Automated generation of such highly schematic representations is one of the topics of our ongoing research.

Another open problem is clustering of large amounts of trajectories. All aggregations suggested for supporting the situation-oriented view can be done in a database by means of standard database operations. This approach works quite well for the whole Milan dataset. However, operations suitable for grouping and aggregation of trajectories according to the routes are not yet available in databases. Clustering of trajectories is now done in the main computer memory, in which the whole Milan dataset simply does not fit. Therefore, we can apply clustering only to subsets of trajectories. Moreover, clustering is quite a time-consuming procedure, which additionally limits the size of a subset suitable for an interactive analysis session. The limit depends on the length of trajectories in terms of the number of points. In case of the Milan dataset, about 2000 trajectories can be analyzed relatively comfortably and 6000 trajectories require certain patience. It should be noted, however, that even these numbers of trajectories are much more than can be analyzed by means of only visual and interactive techniques without clustering and aggregation.

At present, we together with our partners are developing a method for clustering large sets of trajectories. The idea of the approach is following. The analyst first selects a sample of trajectories from the database and applies the available clustering techniques to this sample. For each cluster of interest, one or several representative trajectories are chosen. Then, the remaining trajectories from the database are compared to the representative trajectories by means of the same distance function as has been used for the clustering. If the distance of a new trajectory to one of the representative trajectories is below the distance threshold (which is one of the parameters of the clustering), the new trajectory is attached to the respective cluster. This procedure works much faster than the initial clustering, in which every trajectory needs to be compared with every other trajectory. Our pilot experiments yielded promising results. However, a full implementation of the approach requires solving two sub-problems: selection of an appropriate sample and selection of appropriate cluster representatives.

7.4 Dynamic aggregation

As mentioned in the previous subsection, clustering and aggregation of trajectories is now done in the main memory, which limits the amount of data that can be analyzed. On the other hand, this offers interesting opportunities for interactive analysis. One of them is dynamic aggregation. A *dynamic aggregator* is a special kind of object which is linked to several other objects (members of an aggregate) and computes certain summary attributes such as the member count and the minimum, maximum, average, median, etc. from the attribute values of the members. The aggregator reacts to various kinds of interactive filtering applied to its members by adjusting the values of the summary attributes: only the members that pass through the current filter are taken into account in computing the values.

In our experimental software, we have two types of dynamic aggregators of movement data: *aggregate moves* and *summation places*. Aggregate moves have been introduced before. An aggregate move is defined by two places (areas) A and B and is "aware" of all trajectories that visit A and B in this specific order without passing any intermediate place. The members of the move are the respective trajectory fragments. An aggregate move produces the count of its members and statistical summaries of the lengths of the fragments, their durations, and speeds. A *summation place* is an area in space which is "aware" of all trajectories passing through it and stores the positions and times of entering and leaving the area. The fragments of the trajectories lying inside the area are the members of the place. A summation place counts its members and computes such summary attributes as the minimum, maximum, average, median, and total time spent in the area, statistics of the speeds, etc., as well as the number of starts (first points of trajectories) and the number of ends (last points). In particular, the places where aggregate moves originate or end are summation places. It is also possible to create other summation places, for example, in cells of a regular grid.

Aggregate moves are visually represented on a map by arrows, as shown in Figures 8-11. They can also be represented by symbols in cells of an origin-destination matrix (Figure 8). Summation places may optionally be visualized on a map by drawing the outlines of the areas and/or filling the interiors. Irrespective of whether the places themselves are shown or not, current values of selected summary attributes may be visualized on a map by graduated symbols or by diagrams. It is also possible to use non-cartographic displays such as scatterplot, parallel coordinates plot, or frequency histogram.

The interactive filters, which make dynamic aggregates re-compute the values of the summary attributes, include temporal filter (selection of a time interval), spatial filter (selection of a “window” in space), attribute filter (selection of trajectories by their attributes such as duration and length), and cluster filter (selection of clusters). In re-computing, the aggregates take into account only the active members, i.e. the members that have passed through *all* currently set filters. When an aggregate move has no active members, it does not appear on a map; otherwise, the thickness of the corresponding vector is adjusted to the current value of the represented summary attribute. The same happens to the symbols or diagrams representing the summary attributes attached to the summation places. In particular, when some cluster of trajectories is selected, the summation places and aggregate moves summarize only the trajectories belonging to this cluster.

We would like to stress the difference between the filtering of aggregate moves, which has been demonstrated in Figure 9 and Figure 10, and dynamic reaction of aggregate moves to filtering of trajectories. These are two independent mechanisms, which can be used separately or in combination. The filtering of aggregate moves can hide some of them from the view but does not change the values of their attributes. When filtering is applied to trajectories, the values of the attributes of some aggregate moves change. This may affect the visual appearance of the aggregate moves; in particular, some moves may become invisible because none of their members is active. If, additionally to this, some filtering is applied to the aggregate moves, it may happen that the new attribute values change the satisfaction of the filter conditions: a previously active move may turn into inactive and be hidden and vice versa.

Let us now present an example of using the dynamic aggregation for the exploration of the paths people take for getting from one region to another. In this example, we would like to detect the frequently used paths from the outskirts of Milan to the center. We select from the database a subset of trajectories ending in the centre and starting elsewhere (for technical reasons, which have been explained before, we load a subset of trajectories made during one day, specifically, on Wednesday). We apply the clustering tool to the trajectories, which gives us 37 clusters of trajectories and quite large proportion of “noise”, i.e. trajectories not included in any cluster. Many of the clusters are very small. We exclude them from the view and look only at the clusters containing at least 20 trajectories. The result appears on the map display as shown in Figure 12, right. To make the view clearer, we have filtered the aggregate moves: only the moves occurring in at least 5 trajectories are visible.

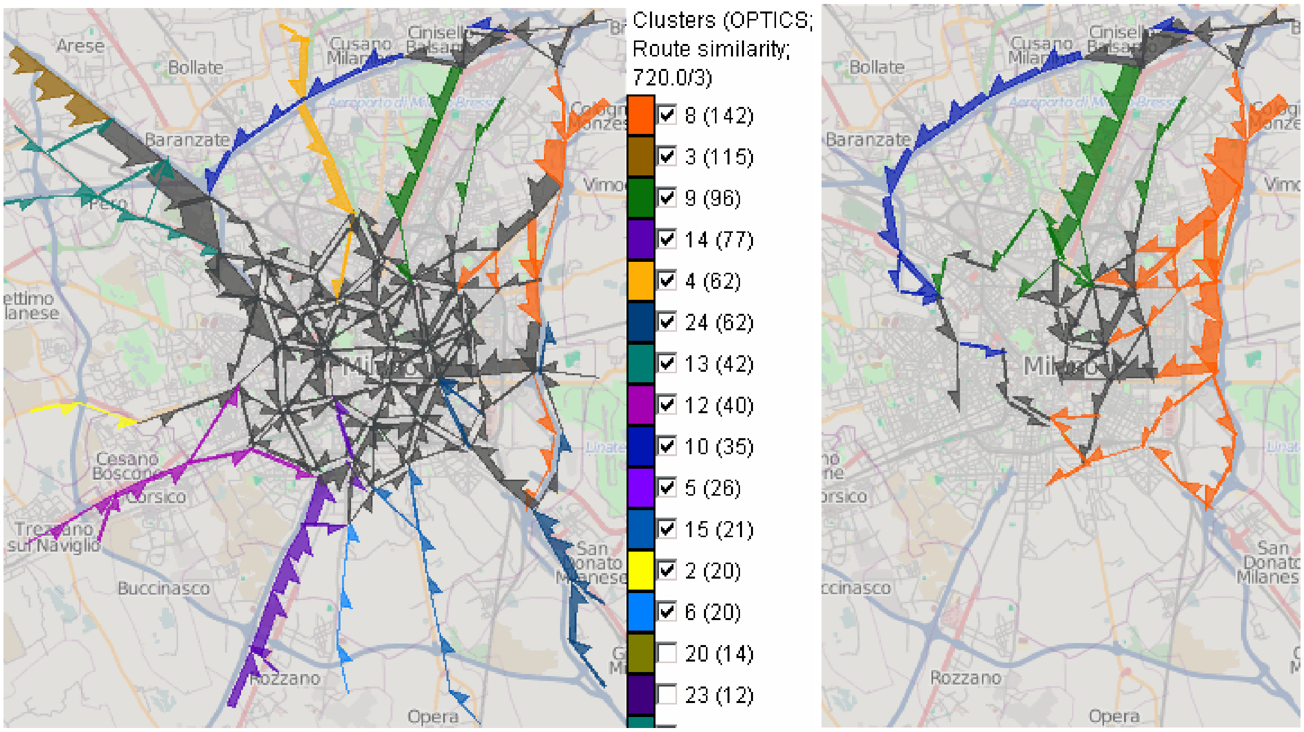

Figure 12. Left: the major routes to the center of Milan. Right: three alternative routes to the centre starting in the same area.

It may be seen that the arrows representing the aggregate moves on the map are differently colored. This is because the aggregate moves, as dynamic aggregators of fragments of trajectories, are "aware" not only about the filtering applied to the trajectories (in this case, cluster selection) but also about the classification of the trajectories and the colors assigned to the classes (in this case, the classes are the clusters). If all active members of an aggregate move belong to the same class, the corresponding arrow will have the color of this class. The dark gray arrows correspond to the moves whose active members belong to two or more classes. Thus, in Figure 12 we see that the major routes "mix" in the centre; besides, there are some overlaps of routes in other parts of the city.

By exploring the map, we detect that three different routes are used for getting from the northeast of the city to the centre. They are shown in orange (cluster 8), dark green (cluster 9), and dark blue (cluster

10). Note that this is the only case of the use of several alternative routes from the same region to the centre. To focus on these routes, we deselect other all clusters. The result is shown in Figure 12 (left).

It may be noted that, apparently, not all trajectories included in the three clusters really start in the same area. Thus, a great part of the orange cluster comes from a place situated more to the south than the origins of the majority of the trajectories. We apply additional spatial filtering to the trajectories in order to consider only the trajectories with close origins on the northeast of Milan. The aggregate moves adjust the appearance of the arrows to account for the additional filter.

After the filtering by the origin, we can see that the "green" route, which is the most direct among the three, is taken more frequently than the others. This is what could be expected based on the common sense. But does this proportion keep over the whole day? And what makes some people choose the longer routes?

We apply the temporal filter in order to investigate the use of the three routes in different time intervals of the day. Figure 13 presents the screenshots of the map made for 2-hour time intervals starting from the interval 5-7h and ending with the interval 19-21h (there was only very little movement before 5h and after 21h). For each time interval, the aggregate moves summarize only those fragments of trajectories which fit in this interval.

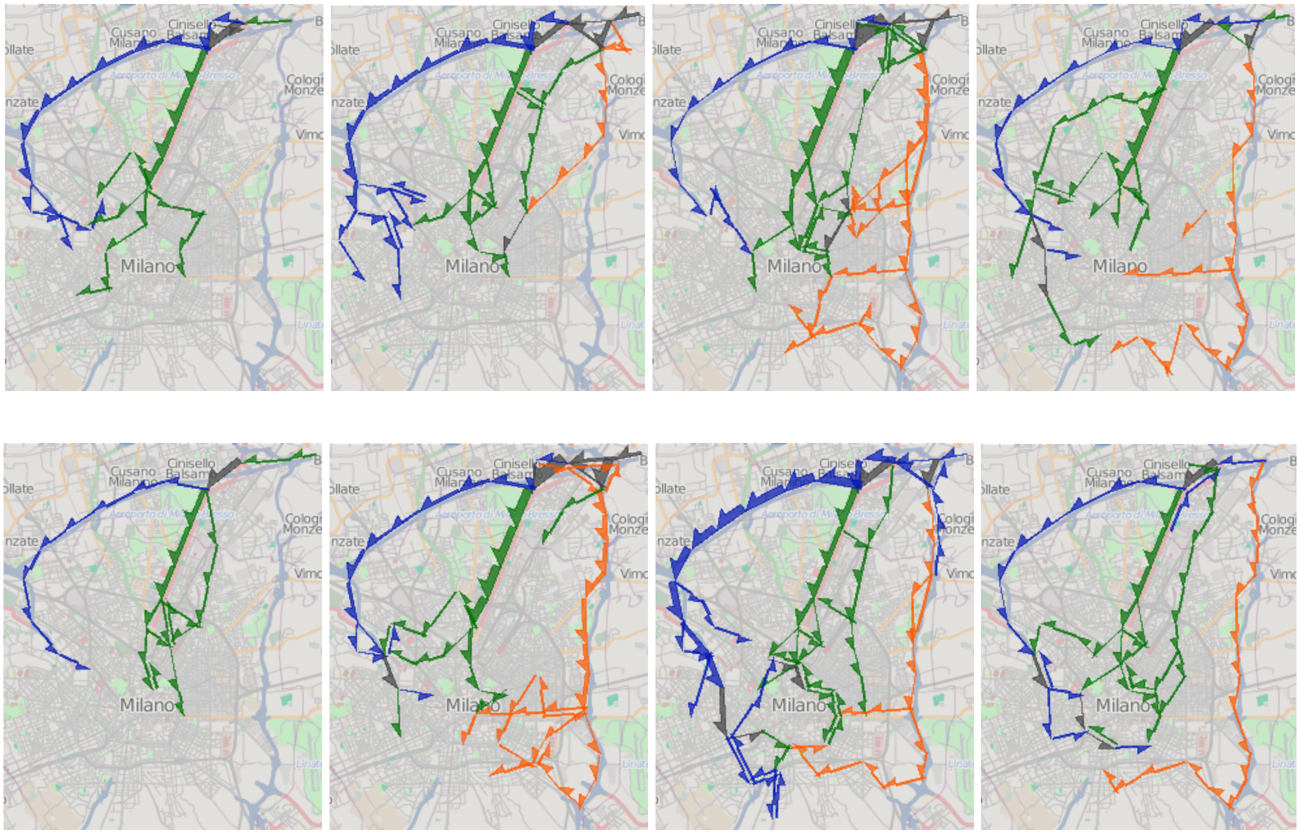

Figure 13. The use of the three alternative routes from the northeast to the center of Milan in different times of the day. Upper row, from left to right: 5-7h, 7-9h, 9-11h, and 11-13h. Lower row, from left to right: 13-15h, 15-17h, 17-19h, and 19-21h.

It can be seen that the relative frequencies of the use of the three routes vary over time. Thus, the “blue” route is as frequent as the “green” route in the interval 17-19h and only slightly less frequent in the intervals 5-7h, 7-9h, and 15-17h. In the remaining times, the “blue” route is used much less frequently than the “green” one. The frequency of the use of the “orange” route is quite low throughout the day but somewhat increases in the intervals 9-11h and 15-17h.

A possible reason for choosing one or another route may be the traffic situations on different roads: if the movement on some road is obstructed, a driver may prefer to use another road. In order to check whether the popularity of the “green”, “blue”, and “orange” routes is related to obstructed traffic, we built a rectangular grid over the territory and generated summation places in the cells. Then we visualized the median speeds in the summation places and looked at the variation of the speeds over time. In Figure 14, the median speeds are represented by bars of proportional sizes. The screenshot fragments show the northeastern corner of the territory, since it is in this area where the decisions concerning the choice of the route are made. The time intervals are from 7-9h to 17-19h. It may be seen that in the intervals 9-11h and 15-17h (in the middle of the upper and lower rows), when the frequency of the “orange” route increases, the movement on the other two routes is very slow. This may be the reason why some drivers prefer to go by the “orange” route. In the intervals 7-9h and 17-19h, when the “blue” route gains its maximum popularity, the median speeds on this route are relatively high. It seems that the movement at the fork of the “blue” and “green” routes is quite obstructed in these time intervals. Perhaps, the highway exit, which is used in the “green” route, is very busy, and some drivers prefer to continue moving along the highway, i.e. choose the “blue” route.

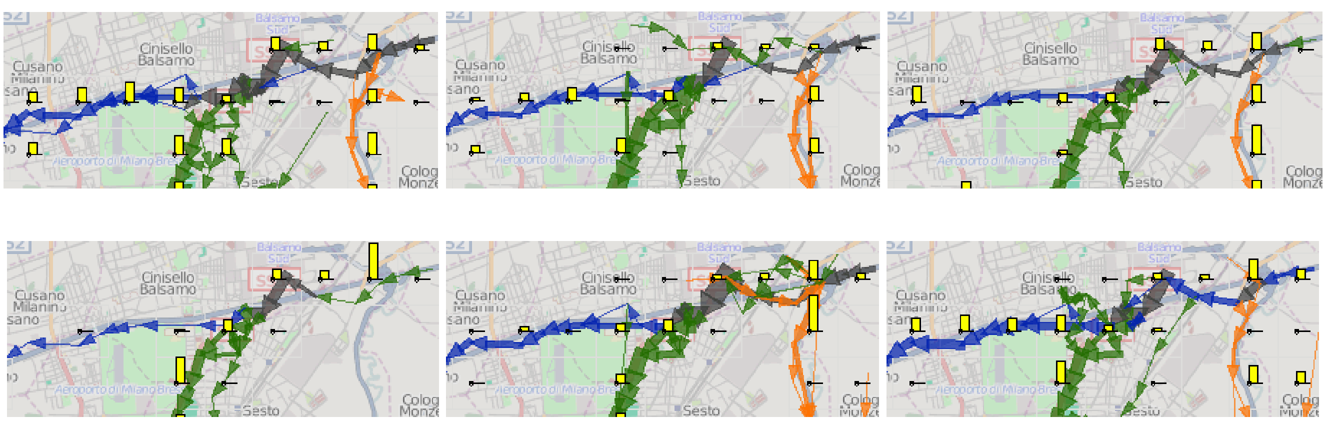

Figure 14. The median speeds of the driving on the northeast of Milan in different times of the day. Upper row, from left to right: 7-9h, 9-11h, and 11-13h. Lower row, from left to right: 13-15h, 15-17h, and 17-19h.

This example demonstrates the use of dynamic aggregation for the detailed exploration of the use of particular routes in different time intervals. Generally, the range of applicability of dynamic aggregation in terms of possible analysis tasks is rather wide. The limitation with regard to the size of data is posed by the in-memory processing; however, it may be expected that the progress of the computer hardware and database technologies will soon permit dynamic filtering and dynamic aggregation of large amounts of data without loading all data in the main memory.

8 CONCLUSION

Current positioning and tracking technologies enable collection of huge amounts of movement data. To make sense and use of such data, scalable analysis and visualization tools are very much needed. Visual exploration of massive movement data necessitates aggregation and summarization of the data.

The goal of our research has been to systemize the existing and, possibly, not yet existing approaches to aggregation of movement data into a framework clearly defining what kinds of exploratory tasks each approach is suitable for. For this purpose, we have used an abstract model of movement of multiple entities as a function of two variables. This model substantiates the possibility of considering movement data from two different perspectives, which we call situation-oriented view and trajectory-oriented view. These two views support different classes of analysis tasks. Thus, the situation-oriented view is required for such tasks as investigation of space use, accessibility, and permeability; the trajectory-oriented view is required for studies of space connectivity, movement flows, routes, etc.

Another dimension for distinguishing the possible classes of analysis tasks is according the focus of analysis: space-centered (investigation of the properties of space with respect to movement) or entity-centered (investigation of the behaviors of the movement entities and relationships between them). We have found that analysis methods based on aggregation of movement data may be inappropriate for entity-centered analysis tasks. However, it may be useful to apply aggregation to outcomes of other analysis techniques, such as automated pattern extraction, for representing these results in a summarized form.

In this paper we have considered the aggregation of original movement data, which is suitable for space-centered analysis tasks. We have investigated which aggregation methods are appropriate for the situation-oriented view and for the trajectory-oriented view of movement. The following table summarizes the aggregation methods and the possible visualization techniques applicable to the results of aggregation.

Table 2. Aggregation and visualization methods supporting two possible views of movement.

View of movement	Aggregation methods	Visualization methods
Situation-oriented: investigation of space use and accessibility	**S** (according to the presence in a space compartment), **S×T** – applied to the position records treated as independent discrete events	Animated map or map series with shading, graduated symbols, or diagrams representing the summary attributes Static map with bar diagrams or other diagrams representing local variation over time in each space compartment Temporal histogram
Situation-oriented: investigation of space permeability/ impedance	**S×T×D** (according to the movement direction) – applied to the position records treated as independent discrete events (after computing the movement directions, which requires two or more consecutive records).	Animated map or map series with directional bar diagrams
Trajectory-	**S×S** (based on the origin and destination	Flow map with arrows representing the

oriented: investigation of space connectivity and flows between places	of a trip) – applied to the entire trajectories or to fragments connecting predefined places **S×S×T×T** – applied to trajectory fragments made during the chosen time intervals	directions and amounts of movement Origin-destination matrix Animated flow map or series of flow maps Animated matrix or series of matrices
Trajectory-oriented: investigation of routes and use of pathways	**R** (route-based) – applied to the entire trajectories or to fragments connecting predefined places; **R×S×S** – summarizes results of R-aggregation **R×S×S×T×T** – applied to trajectory fragments made during the chosen time intervals	Collection of flow maps (each route on a separate map) Flow map with interactive route selection Animated flow map or series of flow maps

In presenting the aggregation and visualization methods, we have used two real example datasets describing constrained and free movements. By these two examples, we have demonstrated the generality of the principles and approaches.

ACKNOWLEDGEMENTS

The work has been done partly within the EU-funded research project GeoPKDD – Geographic Privacy-aware Knowledge Discovery and Delivery (IST-6FP-014915; http://www.geopkdd.eu) and partly within the research project ViAMoD – Visual Spatiotemporal Pattern Analysis of Movement and Event Data, which is funded by DFG – Deutsche Forschungsgemeinschaft (German Research Foundation).

References

[1] Andrienko, G., Andrienko, N., Dykes, J., Fabrikant, S.I., and Wachowicz, M.: Geovisualization of Dynamics, Movement and Change: Key Issues and Developing Approaches in Visualization Research. *Information Visualization*, 2008, v.7 (3/4), pp. 173-180

[2] Andrienko, G., & Andrienko, N.: Spatio-temporal aggregation for visual analysis of movements. In *Proceedings of IEEE Symposium on Visual Analytics Science and Technology (VAST 2008)*, IEEE Computer Society Press, 2008, pp.51-58

[3] Andrienko, G., Andrienko, N., & Wrobel, S.: Visual Analytics Tools for Analysis of Movement Data, A*CM SIGKDD Explorations*, 9(2), 2007, 38-46

[4] Andrienko, N. and Andrienko, G.: *Exploratory Analysis of Spatial and Temporal Data: A Systematic Approach*, Berlin: Springer, 2006

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

[6] Andrienko, N., Andrienko, G., Pelekis, N., & Spaccapietra, S.: Basic Concepts of Movement data. In Giannotti, F. & Pedreschi, D. (eds.): *Mobility, Data Mining and Privacy - Geographic Knowledge Discovery*, Springer, Berlin, 2008, 15-38

[7] Buliung, R.N. & Kanaroglou, P.S.: An Exploratory Data Analysis (ESDA) toolkit for the analysis of activity/travel data. In *Proceedings of ICCSA 2004*, LNCS 3044, 1016-1025

[8] Drecki, I., & Forer, P.: *Tourism in New Zealand - International Visitors on the Move* (A1 Cartographic Plate); Tourism, Recreation Research and Education Centre: Lincoln University, Lincoln, 2000

[9] Dykes, J. A., MacEachren, A. M., & Kraak, M.-J., eds. (2005): *Exploring Geovisualization*, Amsterdam: Elsevier, 710 pp.

[10] Dykes, J. A. & Mountain, D. M.: Seeking structure in records of spatio-temporal behavior: visualization issues, efforts and applications, *Computational Statistics and Data Analysis*, 43 (Data Visualization II Special Edition), 2003, 581-603.

[11] Fredrikson, A., North, C., Plaisant, C., & Shneiderman, B.: Temporal, geographical and categorical aggregations viewed through coordinated displays: a case study with highway incident data. In *Proc. Workshop on New Paradigms in information Visualization and Manipulation* (Kansas City, Nov. 1999). ACM, NY, 1999, 26-34.

[12] Forer, P., & Huisman, O.: Space, Time and Sequencing: Substitution at the Physical/Virtual Interface. In Janelle D.G. and Hodge D.C. (eds), *Information, Place and Cyberspace: Issues in Accessibility*, Springer-Verlag, Berlin, 2000, 73-90

[13] Guo, D.: Visual Analytics of Spatial Interaction Patterns for Pandemic Decision Support. *International Journal of Geographical Information Science*, 21(8), 2007, pp. 859-877

[14] Guo, D., Chen, J., MacEachren, A. M., & Liao, K.: A Visual Inquiry System for Spatio-Temporal and Multivariate Patterns (VIS-STAMP), *IEEE Transactions on Visualization and Computer Graphics*, 12(6), 2006, 1461-1474

[15] Kraak, M.-J.: *The space-time cube revisited from a geovisualization perspective*, in: Proc. of the 21st International Cartographic Conference, Durban, South-Africa, 10-16 August 2003, pp. 1988-1995

[16] McCormick, B., DeFanti, T., and Brown, M.: Definition of Visualization. *ACM SIGGRAPH Computer Graphics*, 21(6), November 1987, p.3

[17] Mountain, D.M.: Visualizing, querying and summarizing individual spatio-temporal behavior. In *Exploring Geovisualization*. (Eds, Dykes, J.A., Kraak, M.J. & MacEachren, A.M.) Elsevier, London, 2005, 181-200

[18] Ntoutsi, I., Mitsou, N., & Marketos, G.: Traffic mining in a road-network: How does the traffic flow? *Int. J. of Business Intelligence and Data Mining*, 3(1), 2008, 82-98

[19] Peuquet, D.J.: It's about time: a conceptual framework for the representation of temporal dynamics in geographic information systems. *Annals of the Association of American Geographers*, 84(3), 1994, pp.441-461

[20] Peuquet, D.J.: *Representations of Space and Time*, Guilford, New York, 2002

[21] Rinzivillo, S., Pedreschi, D., Nanni, M., Giannotti, F., Andrienko, N., & Andrienko, G.: Visually–driven analysis of movement data by progressive clustering, *Information Visualization*, 7(3/4), 2008, pp. 225-239.

[22] Shneiderman, B.: Tree visualization with treemaps: a 2-D space-filling approach. *ACM Transactions on Graphics* 11(1), 1992, 92-99

[23] Thomas, J.J., & Cook, K.A., eds.: *Illuminating the Path. The Research and development Agenda for Visual Analytics*, IEEE Computer Society, 2005

[24] Tobler, W.: Experiments in migration mapping by computer, *The American Cartographer*, 14 (2), 1987, 155-163

[25] Tobler, W.: *Display and Analysis of Migration Tables*, 2005, http://www.geog.ucsb.edu/~tobler/presentations/shows/A_Flow_talk.htm

[26] Tversky, B., Morrison, J.B., and Betrancourt, M: Animation: can it facilitate? *International Journal of Human-Computer Studies*, 57 (4), 2002, pp. 247-262

[27] Wood, J., Slingsby, A., & Dykes, J.: Using Treemaps for Variable Selection in Spatio-Temporal Visualization, *Information Visualization*, 7(3/4), 2008, pp. 210-224.