Mobility Data: Modeling, Management and Understanding

Edited by
Chiara Renso, Stefano Spaccapietra, Esteban Zimányi

Editors

Contents

Illustrations

Tables

Contributors

Natalia Andrienko *Fraunhofer Institute IAIS, Schloss Birlinghoven, Sankt-Augustin, D-53754 Germany, natalia.andrienko@iais.fraunhofer.de*
Gennady Andrienko *Fraunhofer Institute IAIS, Schloss Birlinghoven, Sankt-Augustin, D-53754 Germany, gennady.andrienko@iais.fraunhofer.de*

PART ONE

MOBILITY DATA MODELING AND REPRESENTATION

PART TWO

MOBILITY DATA UNDERSTANDING

9

Visual analytics of movement: a rich palette of techniques to enable understanding

Natalia Andrienko and Gennady Andrienko

9.1 Introduction

Visual analytics develops knowledge, methods, and technologies that exploit and combine the strengths of human and electronic data processing (Keim et al., 2008). Technically, visual analytics combines interactive visual techniques with algorithms for computational data analysis. The key role of the visual techniques is to enable and promote *human understanding* of the data and *human reasoning* about the data, which are necessary, in particular, for choosing appropriate computational methods and steering their work. Visual analytics approaches are applied to data and problems for which there are (yet) no purely automatic methods. By enabling human understanding, reasoning, and use of prior knowledge and experiences, visual analytics can help the analyst to find suitable ways for data analysis and problem solving, which, possibly, can later be fully or partly automated. In this way, visual analytics can drive the development and adaptation of computational analysis and learning algorithms.

Visualization is particularly essential for analyzing phenomena and processes unfolding in geographical space. Since the heterogeneity of the space and the variety of properties and relationships occurring in it cannot be adequately represented for fully automatic processing, exploration and analysis of geospatial data and the derivation of knowledge from it needs to rely upon the human analyst's sense of the space and place, tacit knowledge of their inherent properties and relationships, and space/place-related experiences. This applies, among others, to movement data.

To support understanding and analysis of movement, visual analytics researchers leverage the legacy of cartography , with its established techniques for representing movements of tribes, armies, explorers, hurricanes, etc.; time geography (a branch of human geography), with its revolutionary idea of considering space and time as dimensions of a unified continuum (space-time cube) and representation of behaviours of individuals as paths in this continuum; information visualization, with its techniques for user-display interaction supporting exploratory data analysis; and geovisualization, with its interactive maps and associated methods enabling exploration of spatial information.

This chapter gives a glimpse of the variety of the existing visual analytics methods for analyzing movement data. We group the methods into four categories according to the analysis focus:

1. Looking at trajectories : The focus is on trajectories of moving objects considered as wholes. The methods support exploration of the spatial and temporal properties of individual trajectories and comparison of several or multiple trajectories.
2. Looking inside trajectories: The focus is on variation of movement characteristics along trajectories. Trajectories are considered at the level of segments and points. The methods support detecting and locating segments with particular movement characteristics and sequences of segments representing particular local patterns of individual movement.
3. Bird's-eye view on movement : The focus is on the distribution of multiple movements in space and time. Individual movements are not of interest; generalization and aggregation are used to uncover overall spatio-temporal patterns.
4. Investigating movement in context : The focus is on relations and interactions between moving objects and the environment (context) in which they move, including various kinds of spatial, temporal, and spatio-temporal objects and phenomena. Movement data are analyzed together with other data describing the context. Computational techniques are used to detect occurrences of specific kinds of relations or interactions and visual methods support overall and detailed exploration of these occurrences.

We demonstrate the capabilities of the visual analytics by examples using a dataset consisting of GPS tracks of 17,241 cars collected during one week in Milan (Italy). The data were provided by Comune di Milano (Municipality of Milan).

9.2 Looking at trajectories

In this section, we consider, first, the techniques for visual representation of trajectories and interaction with the representations; second, the use of clustering methods for comparative studies of multiple trajectories; and, third, the time transformations supporting exploration of temporal properties of trajectories and comparison of dynamic properties of multiple trajectories.

9.2.1 Visualizing trajectories

The most common types of display for the visualization of movements of discrete entities are static and animated maps and interactive space-time cube (STC), STC is a unified representation of space and time as a 3-dimensional cube in which two dimensions represent space and one dimension represents time. Spatio-temporal positions can be represented as points in an STC and trajectories as three-dimensional lines When multiple trajectories are shown, the displays may suffer from visual clutter and occlusions. The drawback of STC, besides occlusion, is distortion of both space and time due to projection. It is also quite limited with respect to the length of the time interval that can be effectively explored. To compensate for these limitations, map and STC displays are often complemented with other types of graphs and diagrams.

Common interaction techniques facilitating visual exploration of trajectories and related data include manipulations of the view (zooming, shifting, rotation, changing the visibility and rendering order of different information layers, changing opacity levels, etc.), manipulations of the data representation (selection of attributes to represent and visual encoding of their values, e.g. by colouring or line thickness), manipulations of the content (selection or filtering of the objects that will be shown), and interactions with display elements (e.g. access to detailed information by mouse pointing, highlighting, selection of objects to explore in other views, etc.). Multiple co-existing displays are visually linked by using consistent visual encodings (e.g. same colours) and exhibit coordinated behaviours by simultaneous consistent reaction to various user interactions.

Figure 9.1 gives examples of map and STC displays and demonstrates some basic interaction techniques. The map in Figure 9.1A shows a subset of the Milan dataset consisting of 8206 trajectories that began on Wednesday, 4th of April 2007. To make the map legible, the trajectory lines are drawn with only 5% opacity. A temporal filter, as in Figure 9.1C, can be used to limit the map view to showing only the positions and movements within a selected time interval. Thus, the display state in Figure 9.1B corresponds to the 30-minutes time interval from 06:30 till 07:00. The time filter can also be used for map animation: the limiting time interval is moved (automatically or interactively) forward or backward in time making the map and other displays dynamically update their content according to the current start and end of the interval.

Figure 9.1B also demonstrates the access to various attributes associated

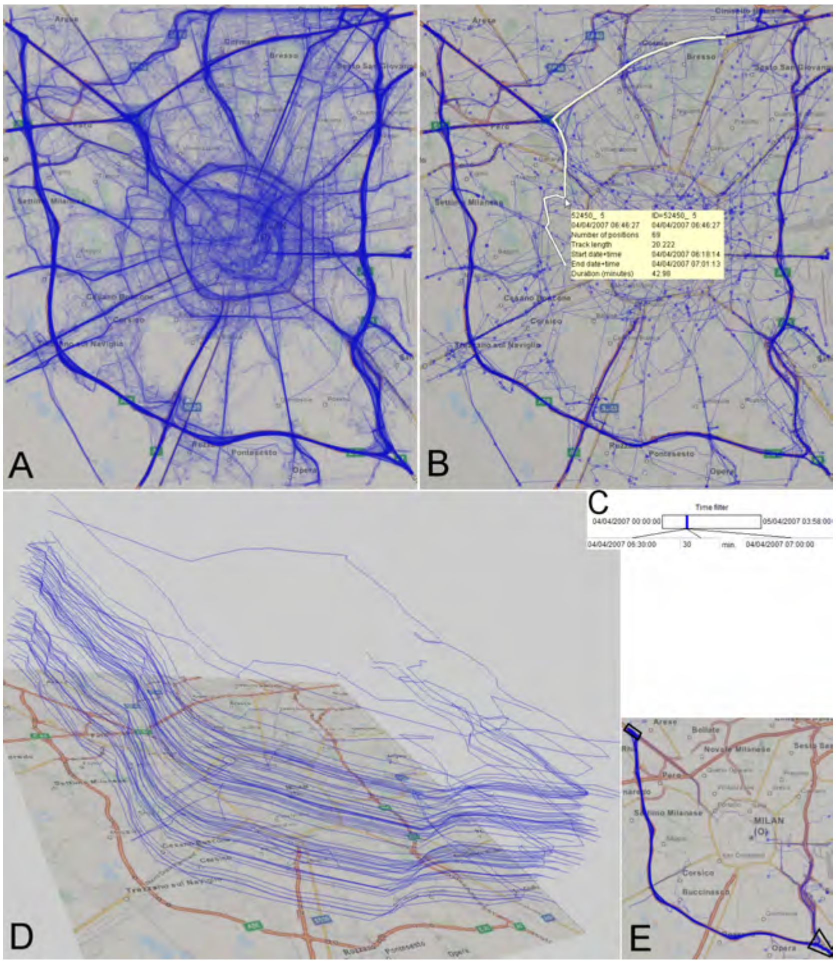

Figure 9.1 A: 8206 trajectories of cars are shown on a map as lines drawn with 5% opacity. B: The map shows only positions and movements from a 30-minutes time interval selected by means of a temporal filter (C). D: A space-time cube (STC) shows a subset of trajectories selected by means of a spatial filter (E).

with a trajectory, such as start and end time, number of positions, length, duration, etc. When the mouse cursor points on a trajectory line, the attributes of this trajectory are shown in a pop-up window as well as the time when the car was in the position at the cursor.

Figure 9.1D demonstrates the space-time cube (STC) display where two dimensions represent the space and the third dimension time. The time axis is oriented from the bottom of the cube, where the base map is shown, to the top. When all trajectories are included in the STC, the view is illegible due to overplotting. In our example, the STC shows 63 trajectories selected by means of a spatial filter (Figure 9.1E). For the filter, we have outlined on the map two areas on the northwest and southeast of the city and set the filter so that only the trajectories that visited both areas in the given order are visible. There are also many other interactive techniques for data querying and filtering, e.g., the ones suggested by Bouvier and Oates (2008) and Guo et al. (2011).

9.2.2 Clustering of trajectories

Clustering is a popular technique used in visual analytics for handling large amounts of data. Clustering should not be considered as a standalone analysis method whose outcomes can be immediately used for whatever purposes. An essential part of the analysis is interpretation of the clusters by a human analyst; only in this way they acquire meaning and value. To enable the interpretation, the results of clustering need to be appropriately presented to the analyst. Visual and interactive techniques play here a key role. Visual analytics usually does not invent new clustering methods but wraps existing ones in interactive visual interfaces supporting not only inspection and interpretation but often also interactive refinement of clustering results.

Trajectories of moving objects are quite complex spatio-temporal constructs. Their potentially relevant characteristics include the geometric shape of the path, its position in space, the life span, and the dynamics, i.e. the way in which the spatial location, speed, direction and other point-related attributes of the movement change over time. Clustering of trajectories requires appropriate distance (dissimilarity) functions which can properly deal with these non-trivial properties. However, creating a single function accounting for all properties would not be reasonable. On the one hand, not all characteristics of trajectories may be simultaneously relevant in practical analysis tasks. On the other hand, clusters produced by means of such a universal function would be very difficult to interpret.

A more reasonable approach is to give the analyst a set of relatively simple distance functions dealing with different properties of trajectories and provide the possibility to combine them in the process of analysis. The simplest and most intuitive way is to do the analysis in a sequence of steps. In each step, clustering with a single distance function is applied either to the whole set of trajectories or to one or more of the clusters obtained in the preceding steps. If the purpose and work principle of each distance function is clear to the analyst, the clusters obtained in each step are easy to interpret by tracking the history of their derivation. Step by step, the analyst progressively refines his/her understanding of the data. New analytical questions arise as an outcome of the previous analysis and determine the further steps. The whole process is called ”progressive clustering” (Rinzivillo et al., 2008) .

There is an implementation of the density-based clustering algorithm OPTICS in which the process of building clusters is separated from measuring the distances between the objects. This allows clustering with the use of diverse distance functions. Hence, the procedure of progressive clustering is done as follows: The user chooses a suitable distance function and applies the clustering tool first to the whole set of trajectories. Then the user interactively selects one or more clusters and applies the clustering algorithm to this subset using a different distance function or different parameter settings. The latter step is iterated. In this way, the user may (a) refine clustering results, (b) combine several distance functions differing in semantics, and (c) gradually build comprehensive understanding of different aspects of the trajectories.

The procedure of progressive clustering is illustrated in Figure 9.2. The first image (A) shows the result of clustering of the same subset of the car trajectories as in Figure 1 using the distance function ”common destinations”, which compares the spatial positions of the ends of trajectories. From the 8206 trajectories, 4385 have been grouped into 80 density-based clusters and 3821 treated as noise. Figure 9.2B shows the clusters without the noise. We have selected the biggest cluster consisting of 590 trajectories that end on the northwest (Figure 9.2C) and applied clustering with the distance function ”route similarity” to it. This distance function compares the routes followed by the moving objects. Figure 9.2D presents the 18 clusters we have obtained; the noise consisting of 171 trajectories is hidden. The largest cluster (in red) consists of 116 trajectories going from the city centre and the next largest cluster (in orange) consists of 104 trajectories going from the northeast along the northern motorway. The orange cluster and the yellow cluster (68 trajectories) going from the southeast along the motorway on the south and west are, evidently, trajectories of transit cars. The clusters by route similarity are also shown in the STC in Figure 9.2E . This display involves time transformation, which is discussed in the next subsection.

9.2.3 Transforming times in trajectories

Comparison of dynamic properties of trajectories using STC , time graph , or other temporal displays is difficult when the trajectories are distant in time, because their representations are located far from each other in a display. This

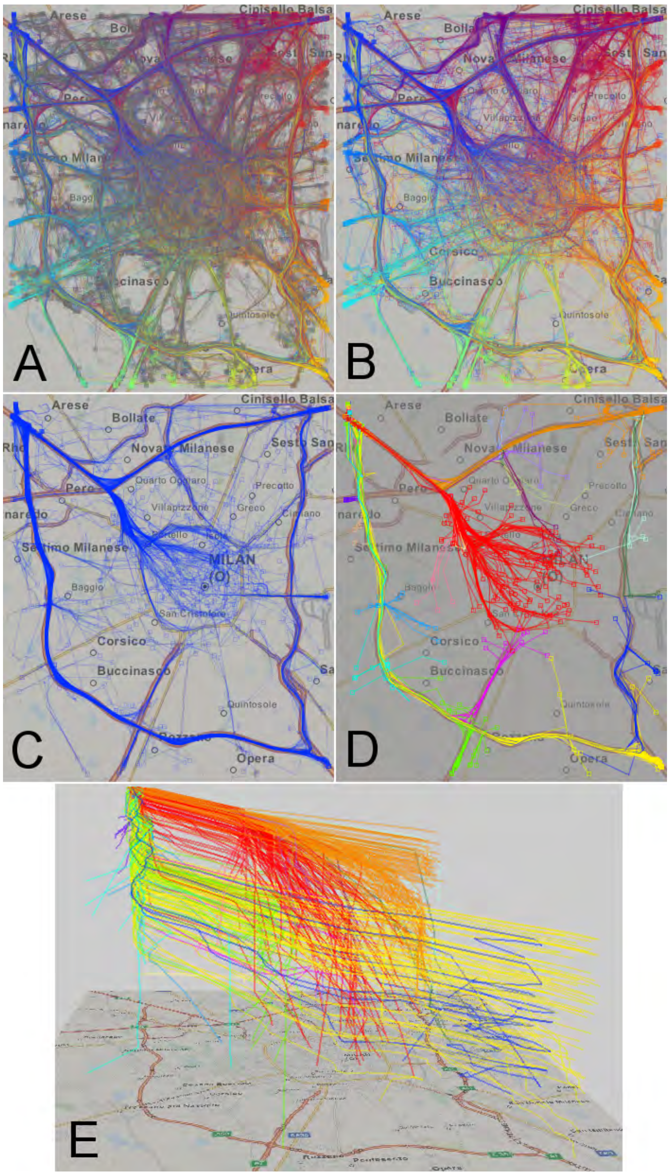

Figure 9.2 Interactive progressive clustering of trajectories. A: The car trajectories have been clustered according to the destinations. B: The noise is hidden. C: One of the clusters is selected. C: Clustering by route similarity has been applied to the selected cluster; the noise is hidden. D: The clusters by route similarity are shown in an STC.

problem can be solved or alleviated by transforming times in trajectories. Two classes of time transformations are possible:

1. Transformations based on temporal cycles. Depending on the data and application, trajectories can be projected in time onto a single year / season / month / week / day etc. This allows the user to uncover and study movement patterns related to temporal cycles, e.g., find typical routes taken in the morning and see their differences from the routes taken in the evening.
2. Transformations with respect to the individual lifelines of trajectories. Trajectories can be shifted in time to a common start time or a common end time. This facilitates the comparison of dynamic properties of the trajectories (particularly, spatially similar trajectories), for example, the dynamics of the speed. Aligning both the start and end times supports comparison of internal dynamics in trajectories irrespective of the average movement speed.

An example of time-transformed trajectories is shown in Figure 9.2E. The STC shows the route-based clusters of car trajectories ending on the northwest. The times in the trajectories have been transformed so that all trajectories have a common end time. This allows us to see that, although the routes within each cluster are similar, the dynamics of the movement may differ greatly. The speeds can be judged from the slopes of the lines. Fast movement is manifested by slightly inclined lines (which means more distance travelled in less time) while steep lines signify slow movement. Vertical line segments mean staying in the same place. In the STC in Figure 9.2 we can very clearly observe the movement dynamics in the red cluster: the cars moved slowly while being in the city centre but could move fast after reaching the diagonal motorway. The orange cluster is divided in two parts. One part consists of nearly straight slightly tilted lines indicating uniformly high speed along the whole route. The other part consists of trajectories with steep segments at the beginning. This means that there were times when the movement in the eastern part of the northern motorway was obstructed and the cars could not reach high speed. We can interactively select the trajectories with the steep segments and find out the times of the obstructed traffic: from about 06 till 13 o'clock; the most difficult situation was after 10:30. Making such observations could be hardly possible with the trajectories positioned in the STC according to their original times.

9.3 Looking inside trajectories: attributes, events and patterns

The methods described in the previous section deal with trajectories as wholes, i.e., treat them as atomic objects. Here we consider methods operating on the level of points and segments of trajectories. They visualize and analyze the variation of movement characteristics (speed, direction, etc.) and other dynamic attributes associated with trajectory positions or segments. The most obvious way to visualize position-related attributes is by dividing the lines or bands representing trajectories on a map or in a 3D display into segments and varying the appearance of these segments. Attribute values are usually represented by colouring or shading of the segments.

Position-related dynamic attributes can also be visualized in separate temporal displays such as a time graph or a time bars display. An example of a time bars display is given in Figure 9.3A. The horizontal axis represents time. Each trajectory is represented by a horizontal bar such that its horizontal position and length correspond to the start time and duration of the trajectory. Note that temporal zooming has been applied: a selected interval from 06:30 till 12:00 is stretched to the full available width. The vertical dimension is used to arrange the bars, which can be sorted based on one or more attributes of the trajectories (start time in our example). Colouring of bar segments encodes values of some user-selected dynamic attribute associated with the positions in the trajectories. This may be an existing (measured) attribute or an attribute derived from the position records, i.e., coordinates and times. Examples of such derivable attributes are speed, acceleration, direction, etc. To represent attribute values by colours, the value range is divided into intervals and each interval is assigned a distinct colour or shade. In Figure 9.3A, shades of red and green represent speed values; red is used for low speeds and green for high. The legend on the left explains the colour coding. Interactive linking between displays allows the user to relate attribute values to the spatial context: when the mouse cursor points on some element within the time bars display, the corresponding spatial position is marked in the map by crossing horizontal and vertical lines and the trajectory containing it is highlighted (Figure 9.3B). In this example we see that the car whose trajectory is highlighted moved at 06:54 on the northeast with the speed 1.2km/h.

The use of this kind of dynamic link is limited to exploration of one or a few particular trajectories. To investigate position-related dynamic attributes in a large number of trajectories, the analyst can apply filtering of trajectory segments according to attribute values. Figure 9.3(C-D) illustrates how such filtering can be done in a highly interactive way. The colour legend on the left

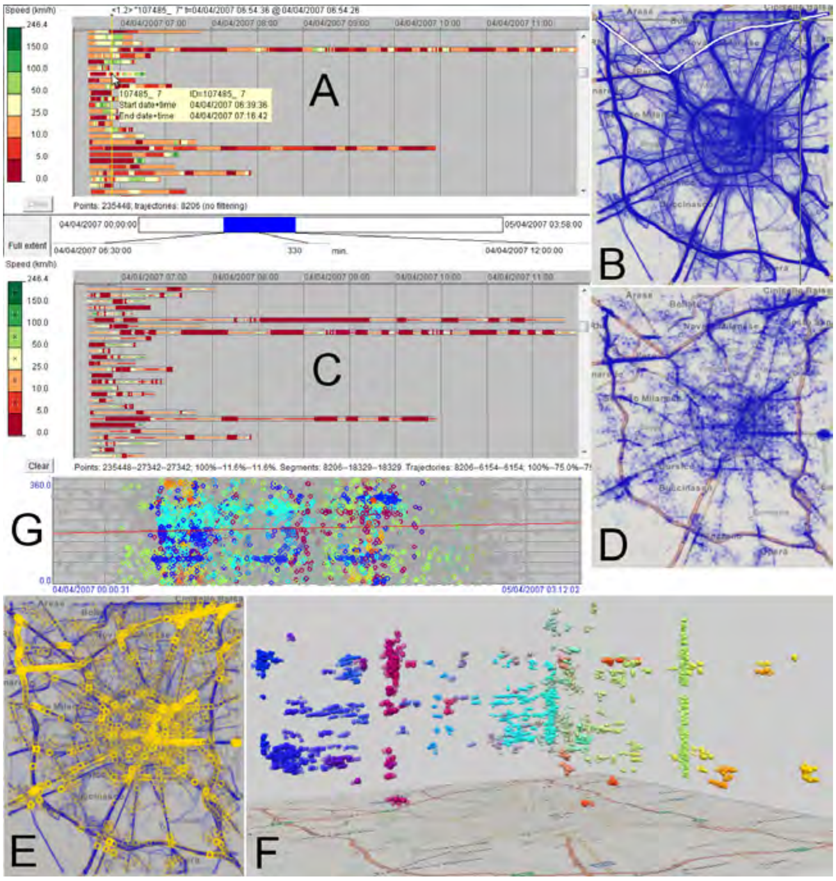

Figure 9.3 A: A time bars display shows the speeds by colour-coding. Mouse-pointing highlights the trajectory and marks the pointed position in a map (B). C: Trajectory segments are filtered according to the speed values. D: Only the segments satisfying the filter are visible on the map. E: Low speed events have been extracted from the trajectories according to the segment filter. F: Density-based spatio-temporal clusters of the low speed events are shown in a space-time cube. G: A scatterplot shows the times (horizontal dimension) and movement directions (vertical dimension) of the low speed events.

of the time bars display is simultaneously a filtering device: the user can switch off and on the visibility of any value interval by clicking on the corresponding coloured rectangle in the legend. In Figure 9.3C, the user has switched off all intervals except for that with the speeds from 0 to 5km/h. As a result, the trajectory segments with the speed values higher than 5km/h have been hidden.

The filter affects not only the time bars display but also the map (Figure 9.3D). It is possible to combine several segment filters based on values of different attributes.

The points satisfying filter conditions can be extracted from the trajectories into a separate dataset (information layer) consisting of spatial events, i.e., objects located in space and time. This dataset can be visualized and analyzed independently from the original trajectories or in combination with them. In Figure 9.3E, the yellow circles represent 19339 spatial events constructed from the points of the car trajectories where the speeds did not exceed 5km/h. The filtering of the trajectory segments has been cancelled so that the whole trajectory lines are again visible. As could be expected, there are many low speed events in the centre of the city. However, there are also visible concentrations of such events in many places on the motorways and their entrances/exits. These events are very probable to have occurred due to traffic congestions.

To investigate when and where traffic congestions occurred, we apply density-based clustering to the set of extracted events in order to find spatio-temporal clusters of low speed events. We look for dense spatio-temporal clusters because standalone low speed events may be unrelated to traffic jams. The distance function we use is spatio-temporal distance between events. The STC in Figure 9.3F displays the clusters we have obtained; the noise (15554 events) is hidden. The clusters are coloured according to the geographical positions. We see a vertically extended cluster in light green on the east of the city. More precisely, it is located at the Linate airport. Most probably, the reason for these low speed events is not traffic congestions but car parking or disembarking/embarking of passengers. The clusters in the other locations are more probable to be related to traffic jams. Some clusters on the northwest (blue) and northeast (cyan) are quite extended spatially, which means that the traffic was obstructed on long parts of the roads. The existence times of the clusters can be more conveniently seen in a two-dimensional display like the scatterplot in Figure 9.3G, where the times of the events (horizontal axis) are plotted against the movement directions. It is possible to select the clusters one by one and see when they occurred and in which direction the cars were moving. For instance, two large clusters of slow movement westwards occurred on the far northeast in the time intervals 05:38-06:50 and 10:20-12:44.

Generally, there are many possible ways how events extracted from trajectories can be further analyzed and used. Interested readers are referred to papers (Andrienko et al., 2011b) and (Andrienko et al., 2011c).

9.4 Bird's-eye on movement: generalization and aggregation

Generalization and aggregation enables an overall view of the spatial and temporal distribution of multiple movements, which is hard to gain from displays showing individual trajectories. Besides, aggregation is helpful in dealing with large amounts of data. There are two major groups of analysis tasks supported by aggregation:

- Investigation of the presence of moving objects in different locations in space and the temporal variation of the presence.
- Investigation of the flows (aggregate movements) of moving objects between spatial locations and the temporal variation of the flows.

9.4.1 Analyzing presence and density

Presence of moving objects in a location during some time interval can be characterized in terms of the count of different objects that visited the location, the count of the visits (some objects might visit the location more than once), and the total time spent in the location. Besides, statistics of various attributes describing the objects, their movements, or their activities in the location may be of interest. To obtain these measures, movement data are aggregated spatially into continuous density surfaces or discrete grids. Density fields are visualized on a map using colour coding and/or shading by means of an illumination model (Figure 9.4A). Density fields can be built using kernels with different radii and combined in one map to expose simultaneously large-scale patterns and fine features, as demonstrated in Figure 9.4A.

An example of spatial aggregation using a discrete grid is given in Figure 9.4B. The irregular grid has been built according to the spatial distribution of points from the car trajectories. The darkness of the shading of the grid cells is proportional to the total number of visits. Additionally, each cell contains a circle with the area proportional to the median duration of a visit. It can be observed that the median duration of staying in the cells with dense traffic (dark shading) is mostly low. Longer times are spent in the cells in the city centre and especially at the Linate airport on the east. There are also places around the city where the traffic intensity is low while the visit durations are high.

To investigate the temporal variation of object presence and related attributes across the space, spatial aggregation is combined with temporal aggregation, which can also be continuous or discrete. The idea of spatial density can be extended to spatio-temporal density: movement data can be aggregated into density volumes in three-dimensional space-time continuum, which can be represented in an STC.

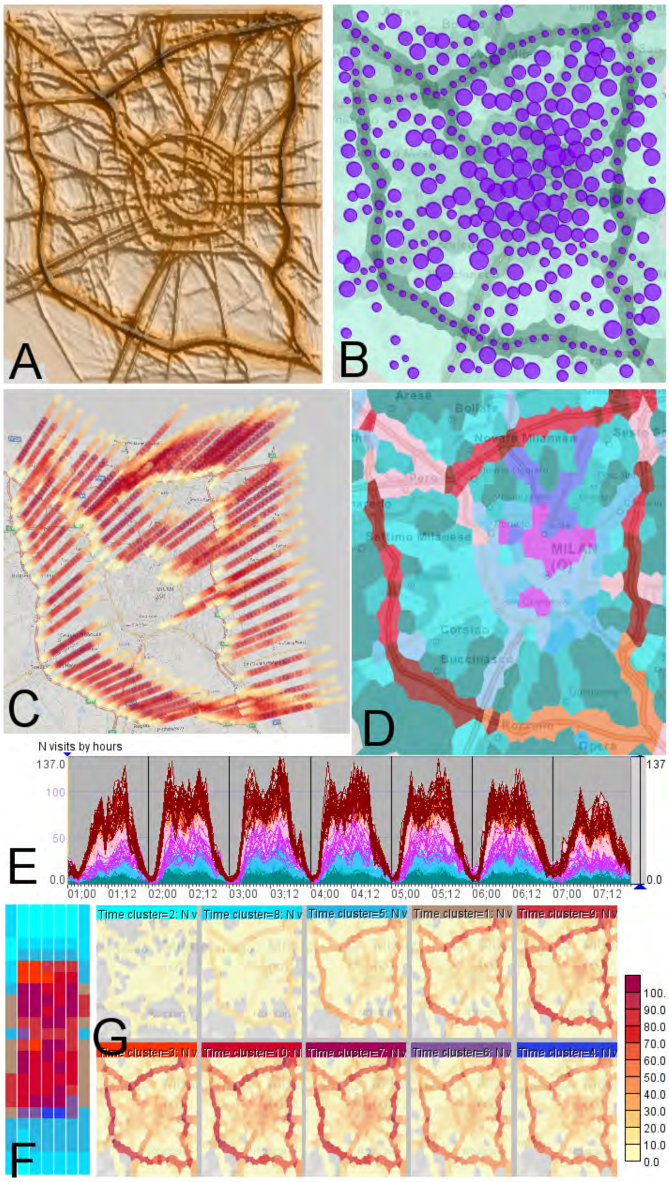

Figure 9.4 A,B: Car tracks aggregated in a continuous density surface (A) and by discrete grid cells (B). C: STC shows the variation of car presence over a day in the most visited cells. D: The cells clustered by similarity of the presence time series shown on a time graph in (E). F: Hourly time intervals clustered by similarity of the spatial distributions of car presence, which are summarized in (G).

For discrete temporal aggregation, time is divided into intervals. Depending on the application and analysis goals, the analyst may consider time as a line (i.e. linearly ordered set of moments) or as a cycle, e.g., daily, weekly, or yearly. Accordingly, the time intervals for the aggregation are defined on the line or within the chosen cycle. The combination of discrete temporal aggregation with continuous spatial aggregation gives a sequence of density surfaces, one per time interval, which can be visualized by animated density maps. It is also possible to compute differences between two surfaces and visualize them on a map, to see changes occurring over time (this technique is known as change map).

The combination of discrete temporal aggregation with discrete spatial aggregation produces one or more aggregate attribute values for each combination of space compartment (e.g., grid cell) and time interval. In other words, each space compartment receives one or more time series of aggregate attribute values. Visualization by animated density/presence maps and change maps is possible as in the case of continuous surfaces. There are also other possibilities. The time series may be shown in an STC by proportionally sized or shaded or coloured symbols, which are vertically aligned above the locations; Figure 9.4C gives an example; the colour legend is given in the lower right corner of Figure 9.4. Occlusion of symbols is often a serious problem in such a display; therefore, we have applied interactive filtering so that only the data for the most intensively visited cells (1000 or more visits per day) are visible.

When the number of the space compartments is big and time series are long, it may be difficult to explore the spatio-temporal distribution of object presence using only visual and interactive techniques. It is reasonable to cluster the compartments by similarity of the respective time series and analyze the temporal variation cluster-wise, i.e., investigate the attribute dynamics within the clusters and do comparisons between clusters. Figure 9.4D demonstrates the outcome of k-means clustering of grid cells according to the time series of car presence obtained by aggregating the car movement data from the whole time period of one week by hourly intervals (hence, the time series consist of 168 time steps). Distinct colours have been assigned to the clusters and used for painting the cells on the map. The same colours are used for drawing the time series lines on the time graph in Figure 9.4E. The colours are chosen by projecting the cluster centroids onto a two-dimensional continuous colour map; hence, clusters with close centroids receive similar colours and, vice versa, high difference in colours signifies much dissimilarity between the clusters. Figure 9.4E shows a prominent periodic variation of car presence in the grid cells over the week. Interactive tools allow us to select the clusters one by one or pairs of clusters for comparison and see only these clusters on the displays.

We find out that the clusters differ mainly in the value magnitudes and not in the temporal patterns of value variation, with the exception of the bright red and orange clusters. The value ranges in these clusters are very close. The main difference is that the red cluster has higher values in the afternoons of Sunday and Saturday. This may have something to do with people spending their leisure time near lakes, which are located to the north from the city.

Spatially referenced time series is one of two possible views on a result of discrete spatio-temporal aggregation. The other possibility is to consider the aggregates as a temporal sequence of *spatial situations*. The term 'spatial situation' denotes spatial distribution of aggregate values of one or more attributes in one time interval. Thus, in our example, there are 168 spatial situations each corresponding to one of the hourly intervals within the week. Temporal variation of spatial situations can also be investigated by means of clustering . In this case, the spatial situations are considered as feature vectors characterizing different time intervals. Clustering groups the time intervals by similarity of these feature vectors.

In Figure 9.4F, we have applied k-means clustering to the 168 spatial situations in terms of car presence and built a time mosaic display where each hourly interval is represented by a square. As in the previous case, different colours have been assigned to the clusters. The squares in the time mosaic are painted in these colours. The squares are arranged so that the columns, from left to right, correspond to the days, from Sunday (the first day in our dataset) to Saturday, and the rows correspond to the hours of the day, from 0 on the top to 23 in the bottom. We see that the working days (columns 2-6) have quite similar patterns of colouring, which means similarity of the daily variations of the situations. The patterns on Sunday (column 1) and Saturday (column 7) are different. The multi-map display in Figure 9.4G shows summarized spatial situations: each small map represents the mean presence values in the respective time cluster (the colour coding is the same as in the STC in Figure 9.4C; see the legend in the lower right corner). It is seen that the shades of cyan, which occur in the night hours, correspond to very low car presence over the city and the shades of red, which occur in the working days from 5 till 17 o'clock, to high presence, especially on the belt roads around the city. Red also occurs in the afternoon of Sunday (from 15 till 17) and in the morning of Saturday (from 8 till 9).

To deal with very large amounts of movement data, possibly, not fitting in RAM, discrete spatio-temporal aggregation can be done within a database or data warehouse. The aggregates can then be loaded in RAM for visualization and interactive analysis.

9.4.2 Tracing flows

In the previous section, we have considered spatial aggregation of movement data by locations (space compartments). Another way of spatial aggregation is by pairs of locations: for two locations A and B, the moves (transitions) from A to B are summarized. This can result in such aggregate attributes as number of transitions, number of different objects that moved from A to B, statistics of the speed, transition duration, etc. The term "flow" is often used to refer to aggregated movements between locations. The respective amount of movement, that is, count of moving objects or count of transitions, may be called "flow magnitude".

There are two possible ways to aggregate trajectories into flows. Assuming that each trajectory represents a full trip of a moving object from some origin to some destination, the trajectories can be aggregated by origin-destination pairs, ignoring the intermediate locations. A well-known representation of the resulting aggregates is the origin-destination matrix (OD matrix) where the rows and columns correspond to the locations and the cells contain aggregate values. OD matrices are often represented graphically as matrices with shaded or coloured cells. The rows and columns can be automatically or interactively reordered for uncovering connectivity patterns such as clusters of strongly connected locations and "hubs", i.e., locations strongly connected to many others. A disadvantage of the matrix display is the lack of spatial context.

Another way to visualize flows is the flow map where flows are represented by straight or curved lines or arrows connecting locations; the flow magnitudes are represented by proportional widths and/or colouring or shading of the symbols. Since lines or arrows may connect not only neighbouring locations but any two locations at any distance, massive intersections and occlusions of the symbols may occur, which makes the map illegible. Several approaches that have been suggested for reducing the display clutter either involve high information loss (e.g. due to filtering or low opacity of lesser flows) or work well only for special cases (e.g., for showing flows from one or two locations).

The other possible way of transforming trajectories to flows is to represent each trajectory as a sequence of transitions between all visited locations along the path and aggregate the transitions from all trajectories. Movement data having sufficiently fine temporal granularity or allowing interpolation between known positions may be aggregated so that only neighbouring locations (adjacent spatial compartments) are linked by flows. Such flows can be represented on a flow map without intersections and occlusions of the flow symbols. To summarize movement data in this way, the space can be tessellated into larger or smaller compartments, e.g., using the method suggested in (Andrienko and

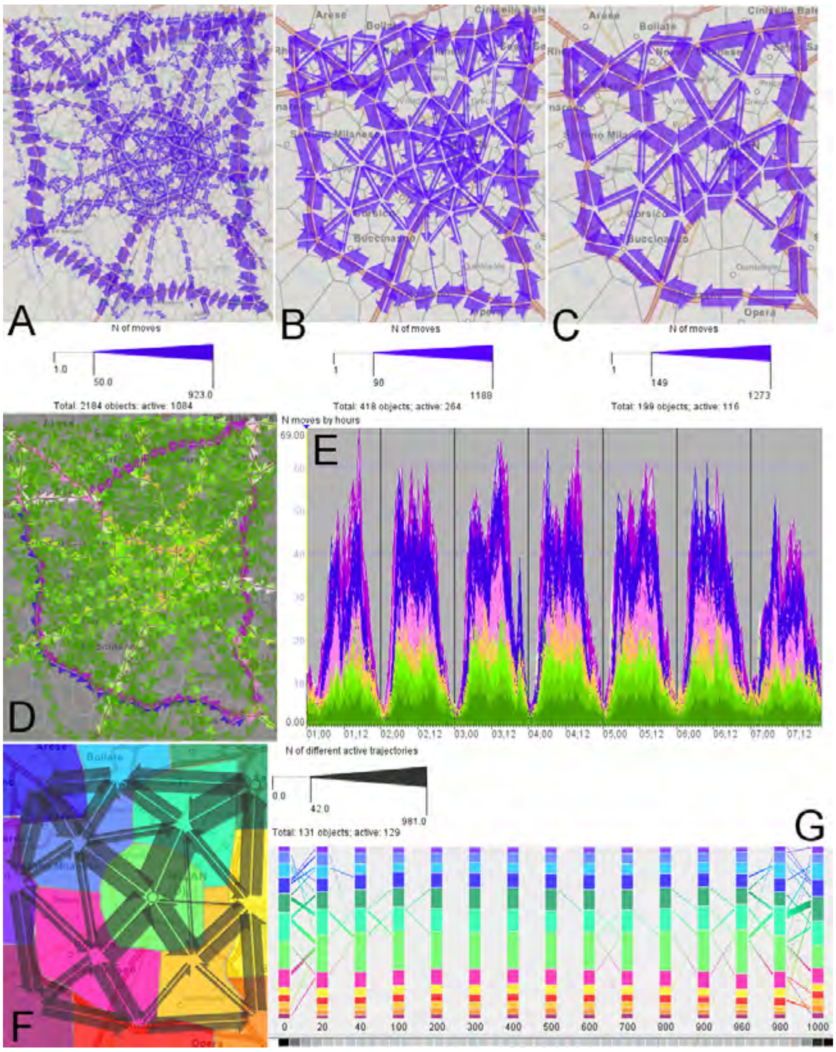

Figure 9.5 A,B,C: Flow maps based on fine, medium, and coarse territory divisions obtained automatically. D,E: Clustering of flows based on the time series of flow magnitudes. F: Flows between predefined regions. G: Investigation of movements between the regions over time adjusted to individual lifetimes of the trajectories.

Andrienko, 2011), to achieve higher or lower degree of generalization and abstraction. This is illustrated in Figure 9.5A-C. The same trajectories of cars (a one-day subset from Wednesday) have been aggregated into flows using fine, medium, and coarse territory tessellations. The flows are represented by ”half-arrow” symbols, to distinguish movements between the same locations in the opposite directions. Minor flows have been hidden to improve the display legibility; see the legends below the maps. The exact values of the flow magnitudes and other flow-related attributes can be accessed through mouse-pointing on the flow symbols. Flow maps can also be built using predefined locations or space partitioning, as demonstrated in Figure 9.5F, where the flow map is built based on a division of the territory of Milan into 13 geographic regions.

Flow maps can serve as expressive visual summaries of clusters of similar trajectories. To obtain such summaries, aggregation is applied separately to each cluster.

When movement data are aggregated into flows by time intervals, the result is time series of flow magnitudes. These can be visualized by animated flow maps or by combining flow maps with temporal displays such as a time graph. Flows may be clustered by similarity of the respective time series (Figure 9.5D,E) and the temporal variation analyzed cluster-wise, as was suggested for time series of presence indicators in the previous section. Note that the spatial patterns visible on the map and the periodic patterns of flow variation visible on the time graph are similar to those that we observed for the presence (Figure 9.4D,E). However, we see that symmetric flows (i.e., flows between the same locations in opposite directions) may have different patterns of temporal variation. Thus, on the east and south of the city symmetric flow symbols are coloured in blue and in magenta, i.e., the respective time series belong to different clusters. The flows in the magenta cluster achieve higher magnitudes in the afternoons of all days, except Friday (day 6).

Aggregation of movement data into transitions between locations does not allow investigation of paths and movement behaviors where more than two locations are visited. The visualization technique demonstrated in Figure 9.5G aggregates trajectories in such a way that movement behaviors can be traced (Bremm et al. (2011)). This is an abstract display where the horizontal axis represents time and colours represent different locations. The map in Figure 9.5F shows the geographic regions of Milan filled in different colours. The same colours are used in Figure 9.5G.

In this example, we investigate the movements of 4634 cars who spent at least 6 hours on the territory under study on Wednesday (i.e., we have selected the trajectories with the duration of at least 6 hours); the flow map in Figure 9.5F summarizes the movements of these cars. The trajectories have been aligned in time to common start and end times, as mentioned in section 9.2.3. The resulting time units are thousandths (also called ’per mill’) of the total trajectory duration. Then the transformed time has been divided into 50 intervals of the length 20 per mills, or 2 percents. The temporal display in Figure 9.5G represents time intervals by vertical bars divided into coloured segments proportionally to the number of cars that visited the regions in these intervals. Aggregated transitions between the regions are represented by bands drawn between the bars. The widths of the bands are proportional to the counts of the objects that moved. Gradient colouring is applied to the bands so that the left end is painted in the colour of the origin location and the right end in the colour of the destination location.

The coloured bars are shown not for all time intervals but for a subset of intervals selected interactively or automatically. In our example, we have selected the first 3 intervals, the last 3 intervals, and each 10th interval (i.e., 100 per mills, 200 per mills, and so on). The small rectangles at the bottom of the display represent all time intervals. The greyscale shading encodes the amount of change in each interval with respect to the previous interval, i.e., how many objects moved to different locations. We can observe that the most intensive movements of the selected cars occurred in the first 2% and in the last 2% of the total trajectory lifetime. Between the time intervals 100 and 900 the cars mostly stayed in the same regions. The most visited region was centre. There were higher presence and more movements in the northern part of the city than in the southern part. The most intensive flows at the beginning of the trips were to the centre and inner northeast and at the end to the outer northeast.

By interacting with the display, it is possible to explore not only direct transitions between locations but also longer sequences of visited locations. When the user clicks on a bar segment, the movements of the corresponding subset of objects are highlighted in the display (i.e., shown by brighter colours). It is possible to see which locations were visited and when. Thus, we can learn that from the 994 cars that were in the centre in the interval 500 (i.e., in the middle of the trip time) 489 cars were in this region during the whole time and the remaining cars came to the centre mainly from the northeast (133), southwest (132), northwest (74) and southeast (62) in the first 2% of the time. At the end, these cars moved back. Analogously, the user can click on bands connecting segments to select the objects participating in the respective transitions and trace their movements.

9.5 Investigation of movement in context

The spatio-temporal context of the movement includes the properties of different locations (e.g., land cover or road type) and different times (e.g. day or night, working day or weekend) and various spatial, temporal, and spatio-temporal objects affecting and/or being affected by the movement. The methods discussed so far seem to deal with movement data alone and not address the context of the movement, at least in an explicit way. However, the context is always involved in the process of interpreting what is seen on visual displays. Thus, the analyst always tries to relate visible spatial patterns to the spatial context (e.g., the highest car traffic density is on motorways) and visible temporal patterns to the temporal context (e.g., the traffic decreases on weekends).

The cartographic map is a very important provider of information about spatial context; therefore, maps are essential in analyzing movement data. It is not very usual although possible to include information about temporal context in temporal displays such as a time graph. A space-time cube may show spatio-temporal context, but occlusions and projection effects often complicate the analysis. Besides the context items that are explicitly represented on visual displays, the analyst also takes relevant context information from his/her background knowledge. Visual displays, especially maps, help the analyst in doing this since things that are shown can facilitate recall of related things from analyst's mind. After noticing a probable relationship between an observed pattern and some context item, or group of items, or type of items, the analyst may wish to check it, which can be supported by interactive visual tools.

The analyst may not only attend to the movement context for interpreting results of previously done analysis. It may also be a primary goal of analysis to detect and investigate particular relationships between the movement and a certain specific context item or group of items. For example, the goal may be to investigate how cars move on motorways or in traffic congestions. To do the analysis, one may need special techniques that support focussing on the context items and relationships of interest.

Position records in movement data may include some context information, but this is rarely the case. In any case, movement data cannot include all possible context information. Typically, the source of relevant context information is one or more additional datasets describing some aspect(s) of the movement context. We shall shortly call such data "context data". Context data may result from previous analyses of movement data. In our previous examples we have demonstrated derivation of spatial events, event clusters, as well as classes (clusters) of locations and of time moments. Such derived data can be considered as context data and used in further analysis of movement data.

The general approach is to derive contextual attributes for trajectory positions by joint processing of movement data and context data and then visualize the attributes to observe patterns and determine relationships. The derived attributes may characterize the environment (such as weather conditions) at the positions of the moving objects or relations (such as spatial distance) between the positions and context items in focus. Values of these attributes are defined, as a rule, for all trajectory positions. The analyst looks for correlations, dependencies, or, more generally, stable or frequent correspondences between the contextual attributes and movement attributes.

Besides stable relationships between movement and its context, the analyst may also be interested in transitory spatial, temporal, and spatio-temporal relationships occurring between moving objects and context items during the movement and lasting for limited time. This includes, in particular, relative movements of two or more moving objects such as approaching, meeting, passing, following, etc., and relative movements with respect to other kinds of spatial objects. Such occurrent relationships can be regarded as spatial events since they exist only at certain positions in space and in time.

Many types of relationships can be expressed in terms of spatial and/or temporal distances. This includes proximity between moving objects, visiting of certain locations or types of locations, being in spatio-temporal neighbourhood of a spatial event, etc. Spatial and/or temporal distances from moving objects to context items can be computed and attached to trajectory positions as new attributes, which can be visualized and/or used in further analyses. Particularly, they can be used for filtering and event extraction as described in section 9.3.

As an example of analyzing movement in context, we shall investigate how the speed of car movement on motorways is related to the distances between the cars. Hence, there are two aspects of the movement context we are interested in: type of location (specifically, motorway) and other cars (specifically, distances to them). The distances between the cars can be determined directly from the trajectory data; no additional data are needed. This can be done using a computational procedure that finds for each trajectory position the closest position in another trajectory within a given time window, e.g., of 1 minute length (from -30 to +30 seconds with respect to the time of the current position).

The location types could be taken from an additional dataset describing the streets; however, we have no such dataset for Milan. We shall demonstrate the use of previously derived data. Earlier we have made a tessellation of the territory (Figure 9.4); moreover, the clustering according to the temporal variation of the car presence (Figure 9.4D) separates quite well the cells on motorways from the other cells. We create a suitable classification of the cells, as in Figure 9.6D, by editing the clusters. Here the yellow filling corresponds to the cells on motorways. We select this class of cells and compute the distances from the trajectory positions to the selected cells; for each position the nearest cell is taken. The computed distances are attached to the position records as a new attribute, which can now be used for filtering. By filtering, we extract the

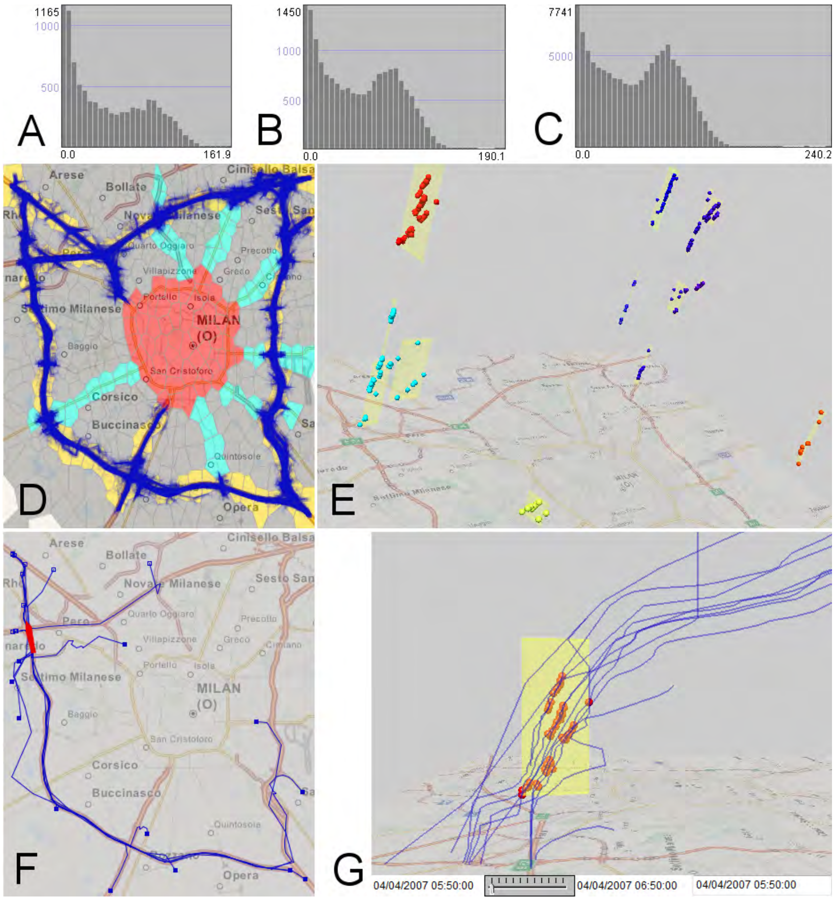

Figure 9.6 A,B,C: Frequency distributions of car speeds on motorways in different ranges of distance to the nearest neighbour car: below 20m (A), 20-50m (B), over 50m (C). D: Trajectory segments on or near motorways selected by means of segment filter. E: Spatio-temporal clusters of low speed events on motorways where the distance to the nearest neighbour is 10m or less. Yellow shapes represent spatio-temporal convex hulls of the clusters. F: Trajectories that passed through one of the convex hulls are selected by filtering. G: The selected trajectories and respective low speed events in a STC.

points and segments of the trajectories with zero distances to the selected cells (Figure 9.6D).

We compute also the distance from each position to the nearest position of another car within the 1-minute time window. This makes one more attribute attached to the position records. Then we use an additional filter according to values of this attribute to sequentially select the trajectory points with the distances to the nearest neighbour in three different ranges: below 20m, from 20 to 50m, and over 50m. For each subset of points, we produce a frequency histogram of the respective speeds. The histograms are shown in Figure 9.6A-C. They have the same height and bar width. The latter corresponds to a speed range of approximately 5km/h. Hence, despite the differing sizes of the point subsets, the shapes of the distributions can be compared. There are many points with low speeds (0-10km/h) in each subset but the relative number of such points is the highest in the first subset and the lowest in the third subset. In all subsets, there is a smaller peak of frequencies for the speeds 80-90km/h, but this peak is the lowest for the first subset and the highest for the third subset. Hence, we observe that smaller distances between cars on a motorway correspond to lower movement speeds.

To demonstrate investigation of occurrent relationships between moving objects and items of the context, we extract from the car trajectories the events where the car is on a motorway and its distance to the nearest neighbour car is at most 10m while the movement speed is not more than 10km/h. These events reflect occurrent proximity relationships of cars to motorways and other cars while the low speeds indicate that these occurrences may be related to traffic congestions. As we did in section 9.3, we find spatio-temporal clusters of these events; some of them are shown in the STC in Figure 9.6E. We build spatio-temporal convex hulls around the event clusters (the yellow shapes in Figure 9.6E). We assume that each convex hull represents a traffic jam. Hence, we have obtained an additional dataset with spatio-temporal boundaries of traffic jams on motorways. It may, in turn, be considered as context data and used in further analysis. Thus, Figure 9.6F shows selected trajectories passing through one of the traffic jams, which have been used as a filter for trajectory selection. We can closely investigate the movement of the cars affected by this traffic jam by means of a STC (Figure 9.6G).

Sections 9.2-9.4 show that movement can be analyzed at different levels: whole trajectories, elements of trajectories (points and segments), and high-level summaries (densities, flows, etc.). In principle, analyzing movement in context can also be done at these levels. A comprehensive set of visual analytics methods addressing all these levels and different types of context items does not exist yet, which necessitates further research in this direction.

9.6 Recommended reading

Keim et al. (2008) give a general definition of visual analytics and describe the scope of this research field. Andrienko et al. (2011a) suggest a conceptual framework defining the concepts of movement data, trajectories, and events, and possible relationships between moving objects, locations, and times. It shows that movement data hold valuable information not only about the moving objects but also about properties of space and time and about events and processes occurring in space and time. To uncover various types of information hidden in movement data, it is necessary to consider the data from different perspectives and to perform a variety of analytical tasks. The paper defines the possible foci and tasks in analyzing movement data. Furthermore, it defines generic classes of analytical techniques and links the types of tasks to the classes of techniques that can support fulfilling them. The techniques include visualizations, data transformations, and computational analysis methods developed in several areas: visualization and visual analytics, geographic information science, database research, and data mining.

Readers interested in visualization of trajectories and techniques for interaction with the displays can be referred to the papers by Kapler and Wright (2005) describing a nice implementation of the space-time cube, Bouvier and Oates (2008) suggesting original interaction techniques for marking moving objects on an animated display and tracing their movements, and Guo et al. (2011) showing the use of several coordinated displays and interactive query techniques specifically designed for trajectories, such as sketching for finding trajectories with particular shapes.

Rinzivillo et al. (2008) talk about visually supported progressive clustering of trajectories. The paper argues for the use of diverse distance functions addressing different properties of trajectories, describes several distance functions, and demonstrates the use of progressive clustering by example.

The papers Andrienko et al. (2011b) and Andrienko et al. (2011c) refer to ”looking inside trajectories” (section 9.3). The first paper describes visual displays that show temporal variation of dynamic attributes associated with trajectory positions. The second paper gives a structured list of position-related attributes that can be computationally derived from movement data alone and from a combination of movement data and context data. These attributes characterize either the movement itself or possible relationships between the moving objects and the movement context. Both papers deal with extraction of spatial events from movement data. The first paper introduces a conceptual model where movement is considered as a composition of spatial events of diverse types and extents in space and time. Spatial and temporal relations occur between movement events and elements of the spatial and temporal contexts. The model gives a ground to a generic approach based on extraction of interesting events from trajectories and treating the events as independent objects. The paper also describes interactive techniques for extracting events from trajectories. The second paper focuses more on the use of extracted events in further analysis. Thus, it considers density-based clustering of movement-related events, which accounts for their positions in space and time, movement directions, and, possibly, other attributes. The clustering allows extraction of meaningful places. The further analysis involves spatio-temporal aggregation of events or trajectories using the extracted places.

Andrienko and Andrienko (2010) give an illustrated survey of the aggregation methods used for movement data and the visualization techniques applicable to the results of the aggregation. These methods and techniques are also presented in a more formal way by Andrienko et al. (2011a). Willems et al. (2009) describe aggregation of trajectories into a continuous density surface using a specially designed kernel density estimation method, which involves interpolation between consecutive trajectory points taking into account the speed and acceleration. Density fields built using kernels with different radii can be combined into one field to expose simultaneously large-scale patterns and fine features. Andrienko and Andrienko (2011) suggest a method for the tessellation of a territory used for discrete spatial aggregation of movement data and generation of expressive visual summaries in the form of flow maps. The method divides a territory into convex polygons of desired size on the basis of the spatial distribution of characteristic points extracted from trajectories. It uses a special algorithm for spatial clustering of points that produces clusters of user-specified spatial extent (radius). Depending on the chosen radius, the data can be aggregated at different spatial scales for achieving lower or higher degree of generalization and abstraction.

An example of visualization of flows between locations in the form of an origin-destination matrix can be found in the paper by (Guo, 2007). The rows and columns can be automatically or interactively reordered for uncovering connectivity patterns such as clusters of strongly connected locations and "hubs", i.e., locations strongly connected to many others.

To deal with very large amounts of movement data, possibly, not fitting in RAM, discrete spatio-temporal aggregation can be done within a database or a data warehouse as described by (Raffaetà et al., 2011). Only aggregated data are loaded in RAM for visualization and interactive analysis. Using roll-up and drill-down operators of the warehouse, the analyst may vary the level of aggregation.

Andrienko and Andrienko (2012) give a comprehensive review and extensive bibliography of methods, tools and procedures for visual analysis of movement data.

9.7 Conclusion

Movement data link together space, time, and objects positioned in space and time. They hold valuable and multifaceted information about moving objects, properties of space and time as well as events and processes occurring in space and time. Visual analytics has developed a wide variety of methods and tools for analysis of movement data, which allow an analyst to look at the data from different perspectives and perform diverse analytical tasks. Visual displays and interactive techniques are often combined with computational processing, which, in particular, allows analysis of larger amounts of data than it would be possible with purely visual methods. Visual analytics leverages methods and tools developed in other areas related to data analytics, particularly, statistics, machine learning, and geographic information science. The main goal of visual analytics is to enable human understanding and reasoning. We have demonstrated by examples how understanding of various aspects of movement is gained by viewing visual displays and interacting with them, possibly, after appropriate data transformations and/or computational derivation of additional data.

PART THREE

MOBILITY APPLICATIONS

PART FOUR

CONCLUSION

References

Andrienko, G., and Andrienko, N. 2010. General Framework for Using Aggregation in Visual Exploration of Movement Data. *The Cartographic Journal*, **47**(1), 22–40.

Andrienko, G., Andrienko, N., Bak, P., Keim, D., Kisilevich, S., and Wrobel, S. 2011a. A conceptual framework and taxonomy of techniques for analyzing movement. *Journal of Visual Languages & Computing*, **22**(3), 213–232.

Andrienko, G., Andrienko, N., and Heurich, M. 2011b. An event-based conceptual model for context-aware movement analysis. *International Journal of Geographical Information Science*, **25**(9), 1347–1370.

Andrienko, G., Andrienko, N., Hurter, C., Rinzivillo, S., and Wrobel, S. 2011c. From Movement Tracks through Events to Places: Extracting and Characterizing Significant Places from Mobility Data. Pages 161–170 of: *IEEE Conference on Visual Analytics Science and Technology, VAST 2011*. IEEE.

Andrienko, N., and Andrienko, G. 2011. Spatial generalization and aggregation of massive movement data. *IEEE Transactions on Visualization and Computer Graphics*, **17**(2), 205–219.

Andrienko, N., and Andrienko, G. 2012. Visual analytics of movement: An overview of methods, tools and procedures. *Information Visualization*.

Bouvier, D.J., and Oates, B. 2008. Evacuation Traces Mini Challenge award: Innovative trace visualization. Staining for information discovery. Pages 219–220 of: *IEEE Symposium on Visual Analytics Science and Technology, VAST 2008*. IEEE.

Bremm, S., Andrienko, G., Andrienko, N., Schreck, T., and Landesberger, T. 2011. Interactive Analysis of Object Group Changes over Time. Pages 41–44 of: *EuroVA International Workshop on Visual Analytics*. Eurographics. Peer-reviewed short paper.

Guo, D. 2007. Visual analytics of spatial interaction patterns for pandemic decision support. *International Journal of Geographical Information Science*, **21**, 859–877.

Guo, H., Wang, Z., Yu, B., Zhao, H., and Yuan, X. 2011. TripVista: Triple Perspective Visual Trajectory Analytics and its application on microscopic traffic data at a road intersection. Pages 163 –170 of: *Pacific Visualization Symposium (PacificVis), 2011 IEEE*.

Kapler, T., and Wright, W. 2005. GeoTime information visualization. *Information Visualization*, **4**(2), 136–146.

Keim, D., Andrienko, G., Fekete, J.-D., Carsten, G., Kohlhammer, J., and Melançon, G. 2008. Visual Analytics: Definition, Process, and Challenges. Chap. 7, pages 154–175 of: Kerren, A., Stasko, J., Fekete, J.-D., and North, C. (eds), *Information Visualization*. Lecture Notes in Computer Science, vol. 4950. Berlin, Heidelberg: Springer Berlin / Heidelberg.

Raffaetà, A., Leonardi, L., Marketos, G., Andrienko, G., Andrienko, N., Frentzos, E., Giatrakos, N., Orlando, S., Pelekis, N., Roncato, A., and Silvestri, C. 2011. Visual Mobility Analysis using T-Warehouse. *International Journal of Data Warehousing & Mining*, **7**(1), 1–23.

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

Willems, N., Van De Wetering, H., and Van Wijk, J.J. 2009. Visualization of vessel movements. *Computer Graphics Forum*, **28**(3), 959–966.

Author index

Subject index