

Chapter 1

Introduction

Abstract This chapter provides an informal introduction of the main concepts related to analysis of movement. The concepts are introduced by illustrated examples, which also demonstrate some techniques that may be used for visual exploration and analysis of movement data. The examples show how the capabilities of the computer and human can be combined to extract knowledge from movement data. This sets the stage for introducing the concept of visual analytics. The chapter also explains the objectives and the structure of the book.

Let us begin by considering a simple example of movement data got from a person who installed a GPS device in her car to record the geographical positions of the car as it moves. Figure 1.1 demonstrates what the *position records* made by the device look like. The main components of the records are the geographical coordinates (X denotes the longitudes and Y the latitudes) and the times when the positions were measured; we shall call them “timestamps”. This is the most typical structure of position records. In a general case, the coordinates are not necessarily geographical and the timestamps do not necessarily consist of Gregorian calendar dates and times of the day. Generally, movement data include positions of some moving objects in a certain space: the geographical space, the entire space of the universe, the internal space of a building, the space of a football field, etc. The positions may be represented by coordinates in a suitable spatial reference system. The timestamps of the position records may be expressed in any system of temporal measurement. These may be, in particular, relative times, such as counts of seconds that passed from the beginning of the observation. The precision of timestamps may also vary: nanoseconds, hours, days, years, centuries, etc. The sequence of position records representing the movement of one object is called the *movement track*.

Movement data represent paths of moving objects through space over time. These paths are usually continuous, that is, a moving object occupies a certain spatial position at any time moment. However, for technical reasons, movement

Fig. 1.1 An example of movement data: GPS-measured positions of a car

pidx	X	Y	time
1	7.207383	50.776790	24-04-07 09:47:55
2	7.207997	50.776703	24-04-07 09:47:58
3	7.208042	50.776783	24-04-07 09:47:59
4	7.208030	50.776830	24-04-07 09:48:00
5	7.208205	50.776817	24-04-07 09:48:01
6	7.208652	50.776836	24-04-07 09:48:04
7	7.208812	50.776910	24-04-07 09:48:10
8	7.208987	50.776863	24-04-07 09:48:52
9	7.209077	50.776670	24-04-07 09:48:56
10	7.209118	50.776394	24-04-07 09:48:59
11	7.209210	50.776093	24-04-07 09:49:02
12	7.209407	50.775810	24-04-07 09:49:05
13	7.209600	50.775566	24-04-07 09:49:08
14	7.209807	50.775303	24-04-07 09:49:11
15	7.209980	50.775043	24-04-07 09:49:14
16	7.210142	50.774796	24-04-07 09:49:17
17	7.210277	50.774560	24-04-07 09:49:20
18	7.210373	50.774323	24-04-07 09:49:23
...

data are discrete. They reflect the spatial positions only at some time moments. There is an inevitable uncertainty concerning the positions of the moving objects in the times between the timestamps. The longer the intervals between the timestamps are, the higher the uncertainty.

The path of a moving object made during the whole time of its existence or movement observation is usually divided into meaningful parts, called *trajectories*. For example, a trajectory may represent a single trip of a person or person's movement during one day. A trajectory of a migratory animal may represent its movement during one migration season. These examples demonstrate that movement of an object may be divided into trajectories in many different ways. The choice of a suitable division depends on the nature of the movement and goals of the analysis.

1.1 A Single Trajectory

Let us return to the GPS track of the personal car. The car owner used the positioning device for almost one year, although not every day, and collected 112,890 position records. We shall take one of the days of the observation period and consider the records made during this day. Thus, on 24 April 2007, the device recorded 458 positions; a few of them are shown in Fig. 1.1. Can we extract meaningful information from these records?

Of course, not much can be gained by just looking at the recorded numbers. However, since these are geographical positions, they can be represented on



Fig. 1.2 The trajectory of the car represented by a *line* on a map. Here and throughout the book, we use a cartographical background from OpenStreetMap (<http://www.openstreetmap.org/>)

a map, which will provide us with the spatial context and aid in interpreting the data. Figure 1.2 demonstrates a map in which the trajectory of the car on 24 April 2007 is represented by a line. The first position of the trajectory is marked by a small hollow square and the last position by a filled square; both squares are on the north close to each other. Note that the line is a continuous representation of the discrete data. It has been constructed by connecting consecutive positions by straight line segments. Since the time intervals between the position records are quite short, this does not introduce much error. We can see that the reconstructed car trajectory quite nicely fits in streets represented in the map.

The map shows us that the trajectory corresponds to a round trip. The start and end positions, marked by the hollow and filled squares, respectively, are approximately in the same place on the north of the area. We can identify the geographical location of the trip and the streets that were used. However, we cannot learn much more about this trip using only the map. The main problem is that the map does not show us the temporal component of the data. As a result, we cannot determine

whether the car moved clockwise or counter-clockwise and whether it stopped on the way. Hence, we need a display that represents not only the space but also the time, such as the space–time cube in Fig. 1.3. This is a perspective view of a three-dimensional representation where the two horizontal dimensions represent the space and the vertical dimension, the time. The temporal axis is oriented from the bottom to the top; the represented time interval is from 09:47:55 till 19:42:29 on 24 April 2007, as can be seen at the bottom of the cube.

The space–time cube gives us additional information. We see that the car moved clockwise and that there was a long stop, which is indicated by a long vertical line segment. Generally, a vertical line segment in a space–time cube means that the spatial position remained the same during a time interval.

To explore parts of the trajectory in more detail, we can apply temporal focusing, which limits the time interval of the data that are visible in the displays. Thus, Fig. 1.4 contains two screenshots of the space–time cube display (top) and two screenshots of the map display (bottom). The screenshots on the left show the part of the trajectory made during the first 12 min of the trip. The screenshots on the right represent the last 50 min of the trip. The cube has been rotated so that it is viewed from the east, that is, the south is on the left and the north on the right. On the left image of the space–time cube, we can detect short stops (vertical line segments), which, probably, occurred at street crossings. On the right image, we see a quite long stop.

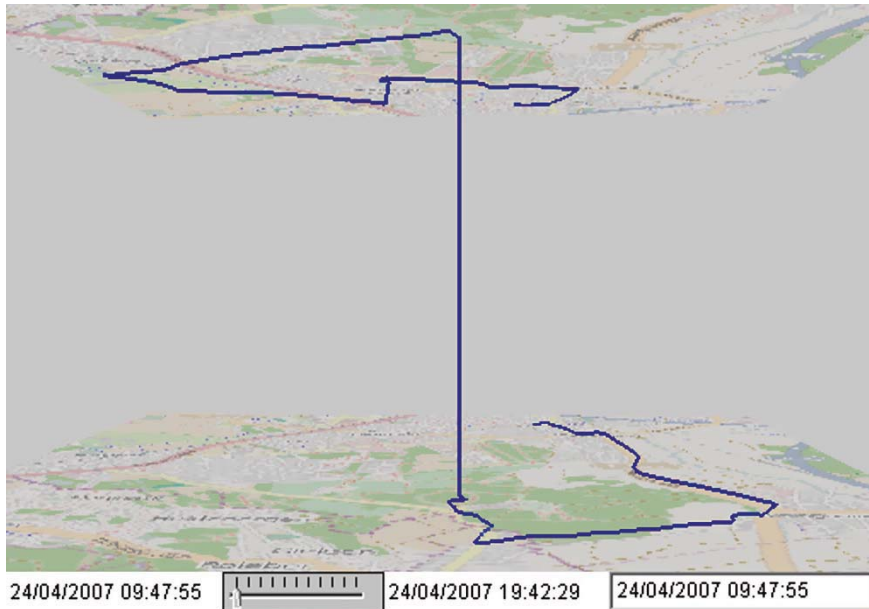


Fig. 1.3 The trajectory of the car is represented by a *line* in a space–time cube

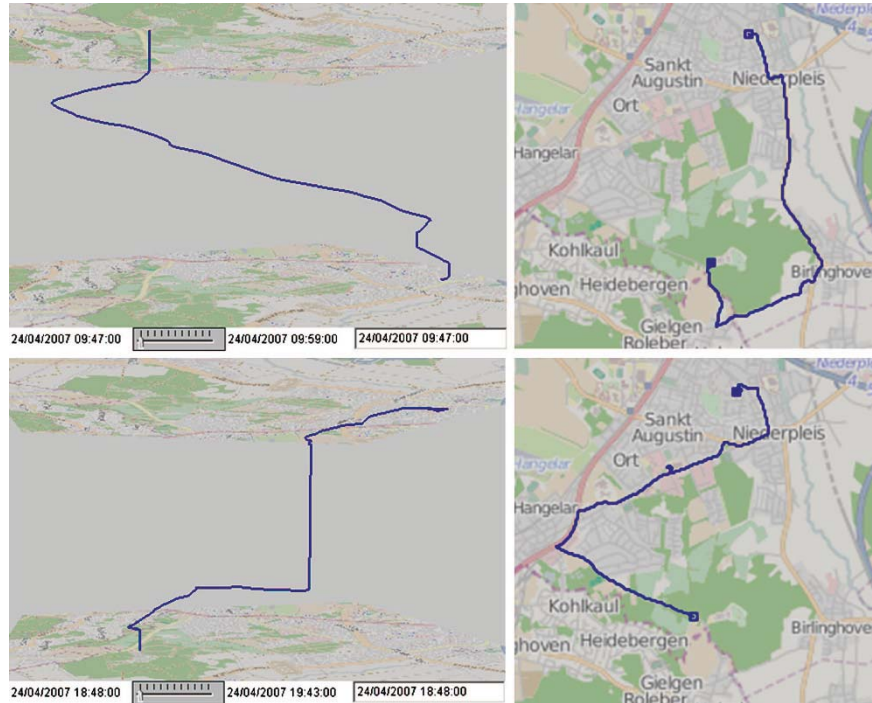


Fig. 1.4 Temporal focusing in the space-time cube display (left) and map display (right). Upper part the first 12 min of the trip. Lower part the last 50 min of the trip

Although the perspective view of the space-time cube can be interactively zoomed, rotated, and moved, it is not very convenient for locating trajectory segments in space and in time. Thus, it is not easy to see at what times and in what places the stops occurred and how long they lasted. There is no ideal display showing space and time together. Therefore, it may be reasonable to use two displays: the map, which is good at conveying the spatial information, and some other display that would be good at conveying the temporal and time-dependent information, such as the time graph in Fig. 1.5. The horizontal axis of the display represents the temporal range of the data. The vertical axis can represent the values of any time-dependent numeric attribute. This may be an attribute computed from the trajectory, such as movement speed, direction, travelled distance, distance to a particular location. The label of the vertical axis (i.e. the attribute name) is shown in the top left corner of the display.

The time graph in Fig. 1.5 represents the cumulative path length (travelled distance) from the starting point of the trajectory. The stops appear in this display as horizontal lines. The labels along the time axis allow approximate temporal positioning. Thus, for the longest stop, we can see that the car stopped at about 10 in the morning and resumed the movement at about 19 o'clock. When we move the mouse cursor over the graph area, the time corresponding to the mouse position is shown above the graph. Hence, by mouse-pointing on the left and right ends of the horizontal line, we can

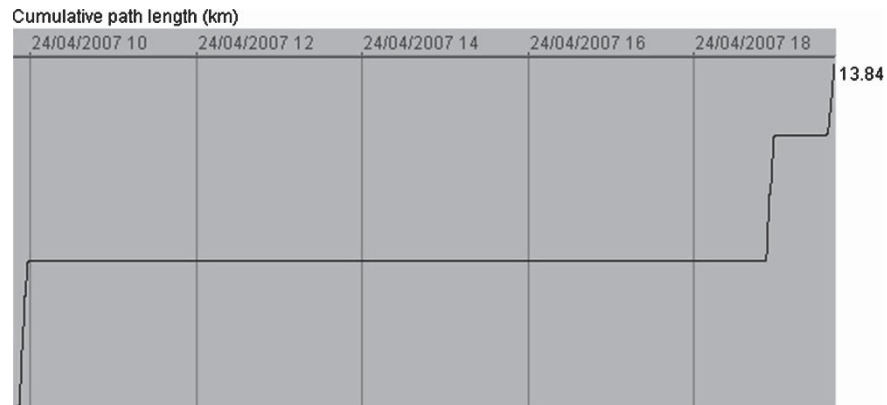


Fig. 1.5 A time graph representing the cumulative path travelled by the car

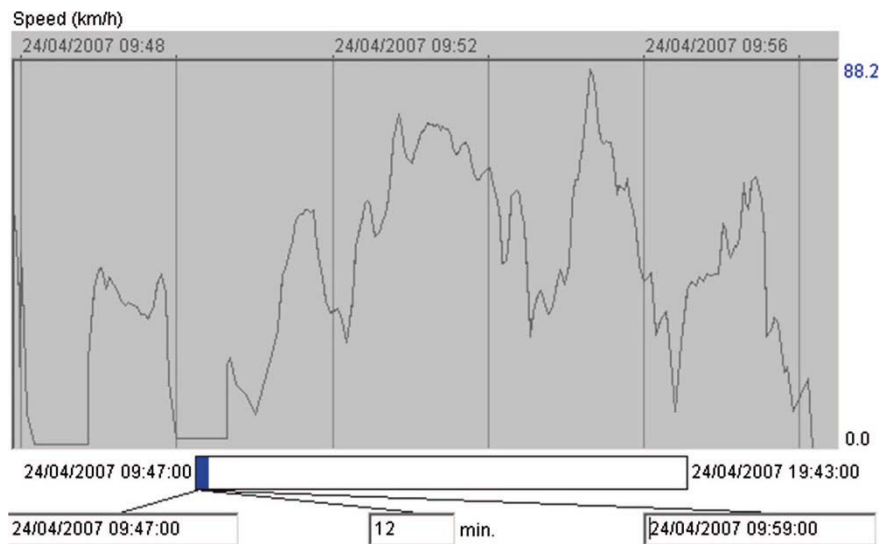


Fig. 1.6 A time graph representing the variation in the instant speed of the car during the first 12 min of the trip

ascertain the start and end times of the stop more precisely: 09:58 and 18:52. The stop visible in the time-focused space-time cube in Fig. 1.4 (top right) appears as a horizontal line segment at the right end of the graph. We can find out that the stop occurred from 18:59 to 19:36. Besides, we can learn that the total length of the trajectory is 13.84 km and the longest stop occurred at 6.44 km from the beginning of the trip.

Figure 1.6 contains the time graph representing the variation in the instant speed of the car over time. We can apply temporal focusing to explore parts of the trajectory in more detail. In Fig. 1.6, we have focused on the first 12 min of the trip. We see that there were two short time intervals of constant very low speed manifested

by horizontal line segments at the bottom of the graph. They correspond to the two vertical line segments in the time-focused space-time cube in Fig. 1.4 (top left), which mean that the car was not really moving during these time intervals. One could expect that the corresponding speed is zero; however, the graph shows positive (although very low) speed values. The reason is that position measurements are never absolutely accurate, and neither is the instant speed computed from them. Hence, in analysing movement data, one should not expect that stops will be always manifested by zero speed values but rather should take a reasonable threshold, that is, a minimum value such that all speed values below it are considered as absence of movement.

The time graph alone is insufficient for exploring spatio-temporal data since it does not represent the spatial aspect. Thus, we can easily find out when the stops occurred in time, but we do not know where they occurred in space. We need a link between the time graph and the map, like the one demonstrated in Fig. 1.7.

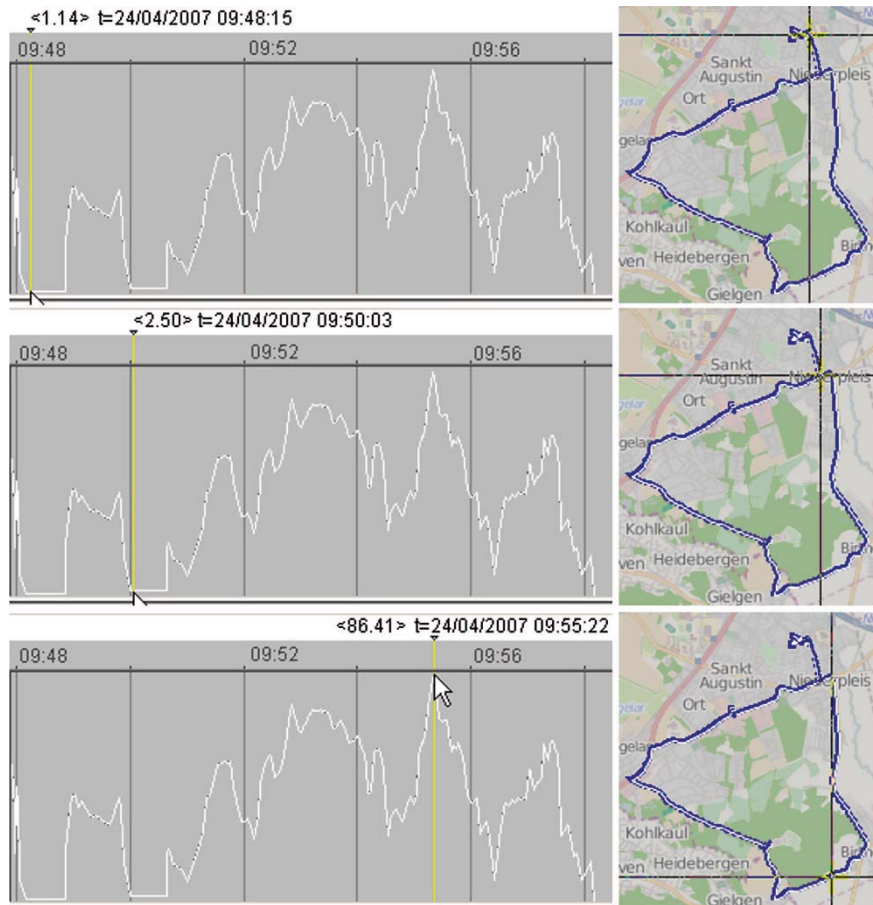


Fig. 1.7 Investigating the speed variation in time and space using dynamically linked map and time graph

When we put the mouse cursor on some point of the line in the time graph, the corresponding spatial position is marked on the map by a cross formed by the intersection of two lines, horizontal and vertical. The screenshots on the top and in the middle of Fig. 1.7 show how we determine the spatial positions of the two short stops at the beginning of the trip (only the tip of the mouse cursor is visible at the bottom of each image). On the bottom left, we have put the mouse cursor on the point in the time graph corresponding to the highest speed value attained during the first 12 min of the trip (the text above the time graph says that the value was 86.41 km/h and that it occurred at 09:55:22). On the bottom right, the cross cursor shows us the respective spatial position. The link between the two displays works also in the opposite direction: when we put the mouse cursor on some point of the trajectory represented on the map, the corresponding temporal position is shown in the time graph by a yellow vertical line.

Unfortunately, a single trajectory does not give us much information. We can guess that the person drove from home to work in the morning, stayed in the working place until evening, and then drove home by another route in order to visit some place on the way, probably, a shop. However, the single trajectory does not give us enough data to check our guesses.

1.2 Multiple Trajectories of a Single Object

Let us now consider the whole dataset that we received from the car owner. This is a single very long sequence of position records. It is not feasible to explore the data in fine detail as we did with the one-day trajectory. If we try to represent all the data in visual displays in the same way as we did for the small subset before, we discover that the displays are not very useful.

For example, Fig. 1.8 presents a fragment of the map display (top) and the space–time cube (bottom). The overlapping lines and the visual clutter do not allow any useful findings. The time graph looks even worse: the whole time span of the data is 27,021,154 s (i.e. 450,354 min or 7,506 h), and there are not enough pixels on the screen to represent the temporal variation in movement attributes with a reasonable resolution. The movement during a whole day has to be squeezed into just two or three pixels along the horizontal dimension representing time. Of course, temporal focusing, as in Fig. 1.6, is applicable; however, we would have to consider, for example, 7,506 hourly intervals, which would be very time-consuming and tiresome. This shows the limitations of purely visual and interactive techniques. Also, note that this dataset is not really big. Much larger amounts of movement data are more usual. Very often the datasets are so large that they cannot even be fully loaded in the computer’s main memory.

To explore large datasets, we need to involve to a greater extent the power of computers. Interactive visualization needs to be combined with computational processing and/or database operations. One of the common approaches for dealing with large datasets is computational aggregation.

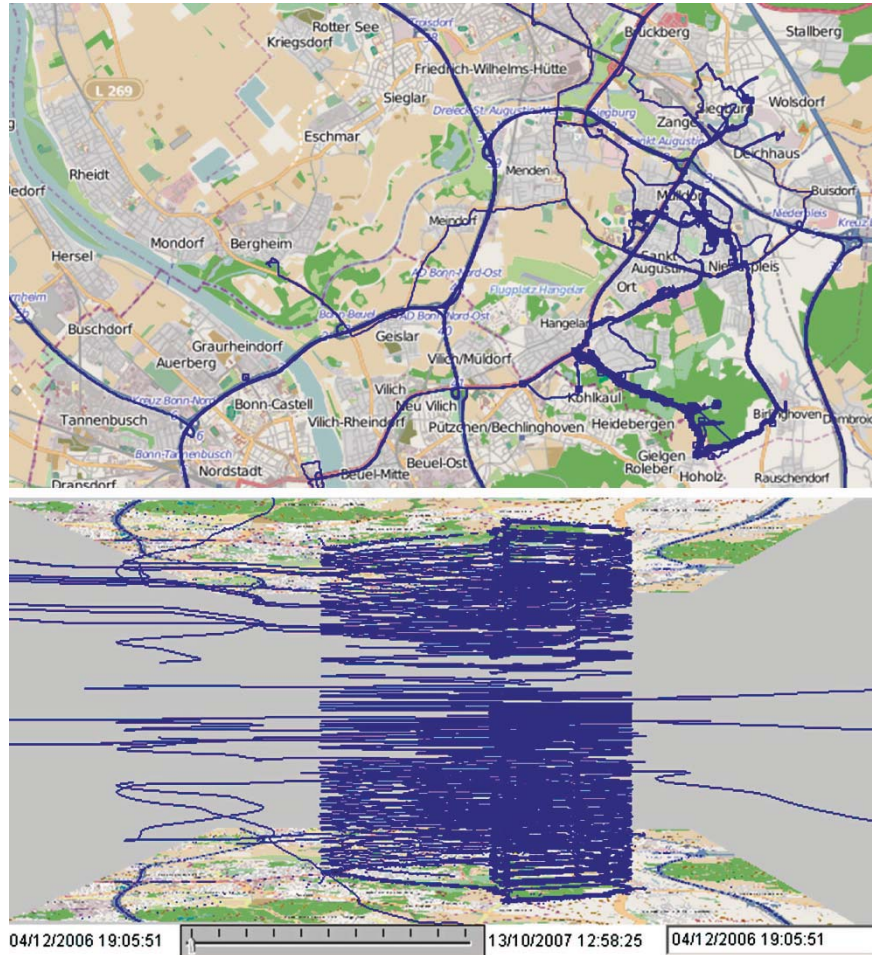


Fig. 1.8 A fragment of the map display (*top*) and the space-time cube display (*bottom*) representing the whole dataset with the positions of the car

For example, in Fig. 1.9, the car movement data have been spatially aggregated into *flows* representing the intensity of the movement. For this purpose, the territory has been divided into compartments. For each pair of compartments, the number of times the car moved from the first to the second compartment has been counted. The results of the aggregation are represented in the map by half-arrow symbols with the widths proportional to the counts. The symbols point in the direction of the movement. This is done in a generalized way, that is, the orientation of a symbol does not necessarily coincide with the actual heading of the car in the respective place but is an aggregate of multiple headings of the car and represents the major movement direction. The half-arrow rather than full-arrow symbols are used in order to show flows in two opposite directions.



Fig. 1.9 A flow map representing the movements of the car in an aggregated form

As compared to the map in Fig. 1.8, the flow map in Fig. 1.9 gives us more information. We see where the car owner moved frequently and where the car owner moved occasionally. The part of the territory where the car owner moved most frequently is shown in more detail in Fig. 1.10. The thickest symbol represents 172 moves (i.e. times when the car owner drove through the respective place). The corresponding symbol oriented in the opposite direction represents 149 moves. In some places, we see that the car owner moved significantly more often in one direction than in the other.

When we explored a one-day trajectory of the car owner in Sect. 1.1, we guessed about the purposes of the trip and the meanings of the places of the start/end and the stops, but we were not much confident of our guesses. Analysing the movement over a longer time period may allow us to come to more definite conclusions. First of all, we would like to find out the significant places of the car owner, that is, the places of her home, work, regularly visited shops, and, possibly, places of other frequent activities. A place significant for a person can be recognized from the number and duration of stops. Home and work places are places where a person usually stops often and stays for quite a long time. Hence, to find these places, we need to retrieve the positions of long stops, say, 3 h or longer. In our data, a stop is manifested by a temporal gap between consecutive position records since the recording was only done when the car actually moved. This may be different in other data: stops may appear as sequences of records with very low speed values, or consecutive spatial positions may form a dense spatial cluster (remember that position measurements are never perfectly accurate; several measurements taken in the same point in space typically do not coincide).



Fig. 1.10 A fragment of the map display showing in more detail the part of the territory where the car owner moved most frequently

Using a database operation, we retrieve the stops for 3 h or more from the car movement data. Each stop has a certain position in space and a certain position in time. We use the term *spatial event* to refer to any discrete physical or abstract object that has a certain position in space and time. Stops are just one example of spatial events that can be extracted from movement data. It is possible to extract many other kinds of spatial events, such as high-speed events, acceleration events, significant turn events, events of passing a street crossing, and so on.

The long stop events we have retrieved from the car movement data are shown as dots on a map in Fig. 1.11. When several stops occurred at the same place, the dots on the map overlap. It is hard to see how many dots are in a place. To distinguish the places of frequent stops from those of occasional stops, we apply a clustering tool, which groups the stops according to the spatial distances between them. It detects two spatially dense clusters of stop points. The corresponding dots are shown in red and blue. The stops that do not belong to the clusters (i.e. the car only occasionally stopped in those places) are shown in dark grey. On the right of Fig. 1.11, a space–time cube shows the spatial and temporal positions of the stops. The cube is slightly rotated, so that we are viewing it from the southeast. The red and blue clusters appear as columns formed by many dot symbols. Precisely, the red cluster contains 220 stops, and the blue cluster contains 135 stops; 10 stops are out of the clusters.

We can be more or less confident that the places of the clusters are the home and work places of the person. To find out which of them is home and which is work, we should look at the times when the person stopped there. Figure 1.12 contains two two-dimensional frequency histograms showing the temporal distribution of the stops in the red and blue clusters by the days of the week and hours of the day. The columns of the histograms correspond to the hours of the day, from 0 to 23, and the rows correspond to the days of the week, from 1 (Monday) to 7

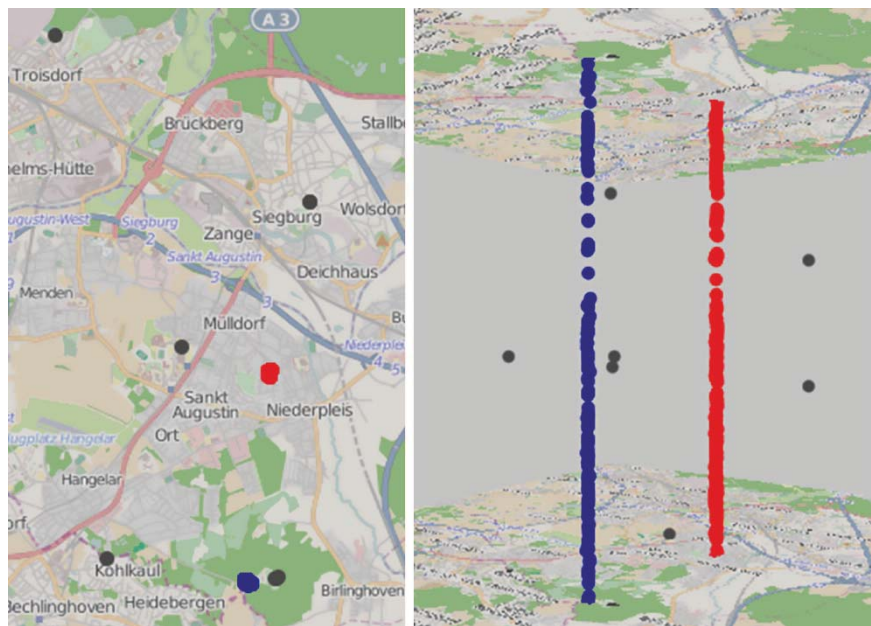


Fig. 1.11 *Left* a fragment of the map display with the spatial positions of the stops for 3 or more hours. *Right* a fragment of the space–time cube display with the spatio-temporal positions of the stops

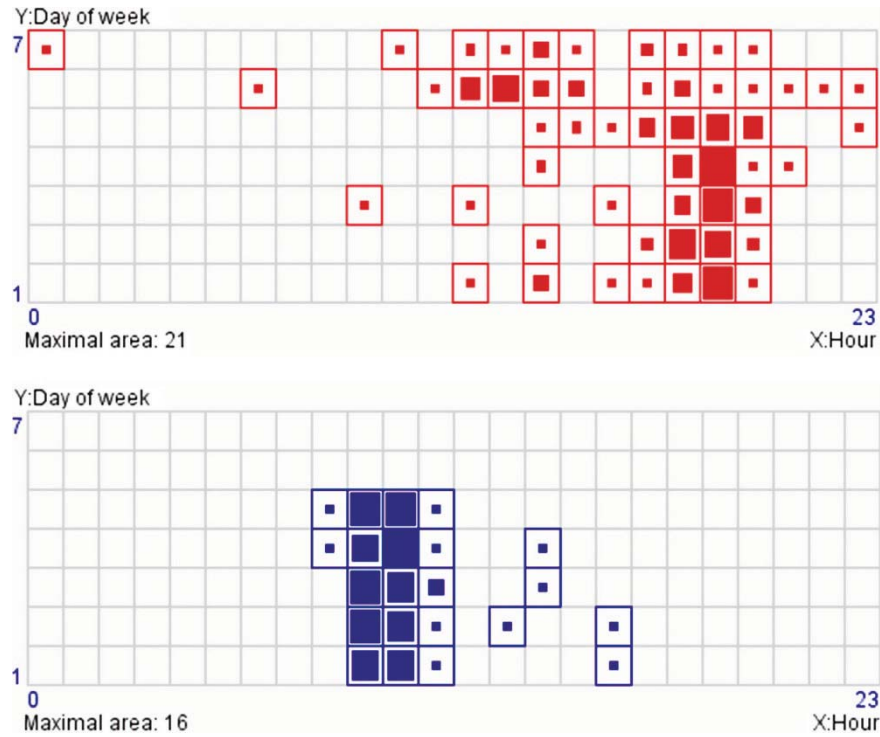


Fig. 1.12 The temporal distribution of the stops from the *red* (top) and *blue* (bottom) clusters by the days of the week and hours of the day

(Sunday). Note that the vertical axis, corresponding to the days of the week, is oriented upwards. The square symbols in the cells represent the frequencies of the respective combinations of day and hour, so that the filled areas inside the squares are proportional to the frequencies. The text below the histogram says what frequency value corresponds to the maximal filled area, that is, the full area of each square.

The upper histogram, which represents the red cluster, tells us that the stops occurred on all days of the week. On the working days (from 1 to 5), the stops mostly occurred in the evening; the maximal frequencies are at about 19 o'clock. On the weekend (days 6 and 7), the stops are more spread over a day; the highest frequencies are on Saturday at 12 and 13 o'clock. The times of the stops in the blue cluster, represented in the lower histogram, are quite different: the stops occurred only on the working days and mostly in the morning; the maximal frequencies are attained at about 9 and 10 o'clock. From these statistics, we can quite confidently conclude that the red cluster represents the home place of the person and the blue cluster, the work place.

In a similar way, we retrieve the stops for at least 30 min. Naturally, they include the stops for 3 h or longer, which we have considered before. The map

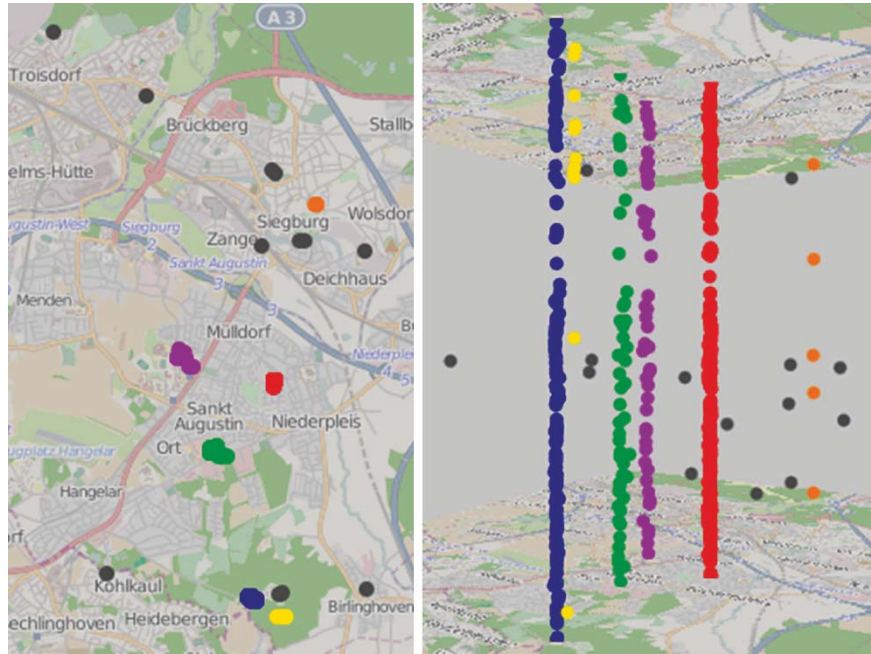


Fig. 1.13 Positions of stops for 30 min or longer in a map (*left*) and space–time cube (*right*)

and space–time cube in Fig. 1.13 represent the stops clustered spatially by means of the same method. Besides the home and work clusters, which are shown in red and blue, there are two other large clusters, green with 51 stops and purple with 46 stops. The two-dimensional histograms in Fig. 1.14 show the distribution of the stops in these clusters by the days of the week and hours of the day. We see that the stops in the green cluster occurred most often in the middle of the day on Saturday and in the evenings of the working days. The stops in the purple cluster occurred most often in the middle of the day on Saturday. For other days and times, the stops were occasional: the filled areas represent one or two stops, except for the square at 18 o'clock on Thursday (day 4), which represents four stops.

The times of the stops in the green and purple clusters suggest that these may be the places of the person's shopping. To check this hypothesis, we zoom in on the places of these clusters in the map (Fig. 1.15) and find out that, indeed, the clusters are located in shopping areas.

For the yellow cluster, consisting of 11 stops, which are quite irregular in time, the map in Fig. 1.13 indicates that it is located in a forest. This may mean that the stops are related to recreational activities of the car owner. All but one stop occurred in the months from May to September in the morning hours (from 8 to 11 o'clock). The remaining stop occurred in December at noon time. On a satellite image from Google Maps, we recognize a tennis ground near the location of the cluster. Perhaps, the person sometimes plays tennis (in warm months of the year) or goes for a walk in the forest.

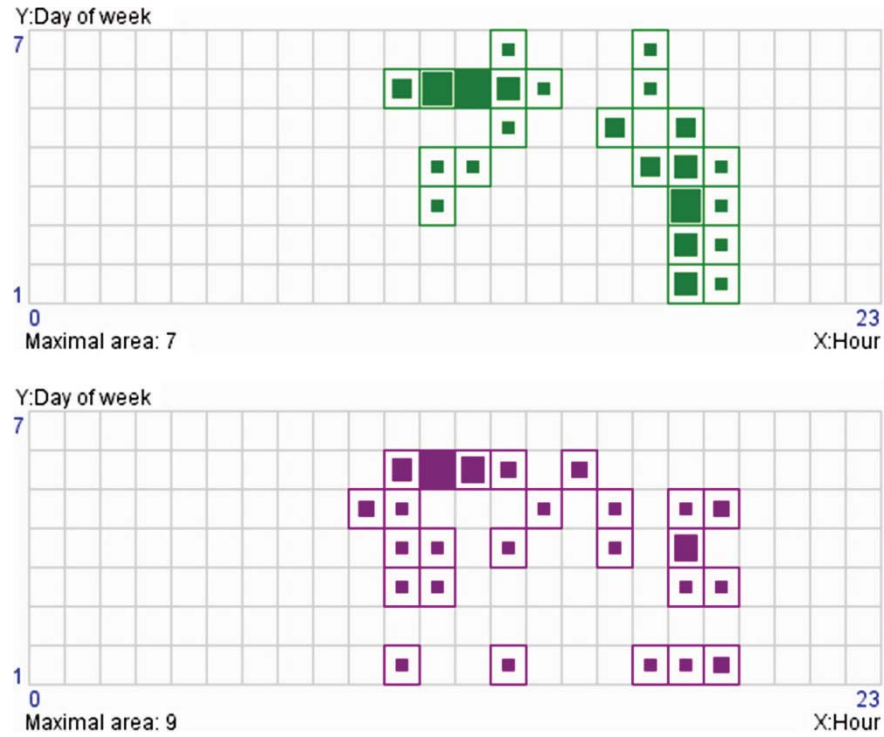


Fig. 1.14 The temporal distribution of the stops from the *green* (top) and *purple* (bottom) clusters by the days of the week and hours of the day

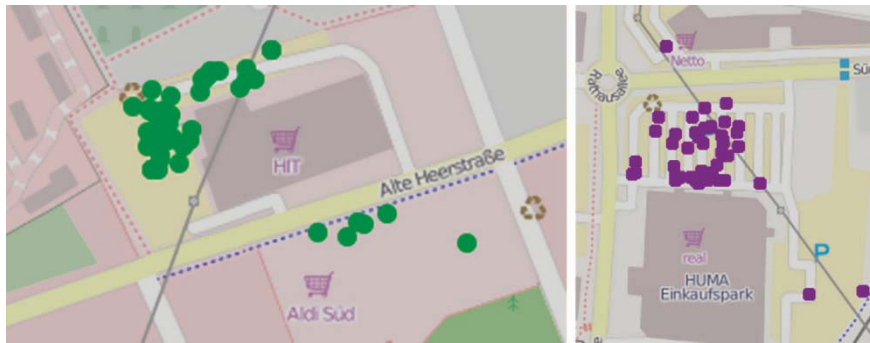


Fig. 1.15 The *green* and *purple* clusters of stops are located in shopping areas

Hence, by analysing the car movement data, we have discovered and interpreted the significant places of the car owner. Now, we are interested in the routes of the movement. However, there is a problem: the dataset that we have is a single movement track. For our analysis, we need it to be divided into trajectories representing different trips. One possible solution to this problem is to divide the track

by stops: an occurrence of a long stop is treated as the end of the previous trip, and the resumption of the movement after the stop is treated as the beginning of the next trip. We need to select a suitable minimum duration of a stop. Selecting different values will divide the track differently. Thus, the single trajectory we have considered in Sect. 1.1 would be divided into two pieces if we choose the minimum stop duration of 3 h and into three pieces if we choose 30 min.

In this example, we choose the minimum stop duration of 3 h. Hence, if the car owner made a stop for shopping on the way from work to home, we consider this as one trajectory rather than as two trajectories. Using this approach, we obtain 365 trajectories. To find repeatedly used routes, we use the same approach as we did for finding the places of frequent stops: we apply clustering. However, this time we apply clustering to the trajectories rather than to stop events. The clustering groups together trajectories following similar routes. We obtain nine groups of similar trajectories varying in size from 4 to 105 and a set of 121 trajectories that do not belong to clusters (this means that their routes are not similar enough to the routes of other trajectories). In terms of clustering, objects that are not assigned to any cluster are called “noise”.

In Fig. 1.16, the clusters of trajectories are shown on a map; the “noise” is hidden by unselecting the respective checkbox in the legend on the right of the map. The clusters are represented by different colours of the trajectory lines. Since overlapping of the lines makes the clusters hard to distinguish, we have to look at each cluster separately.

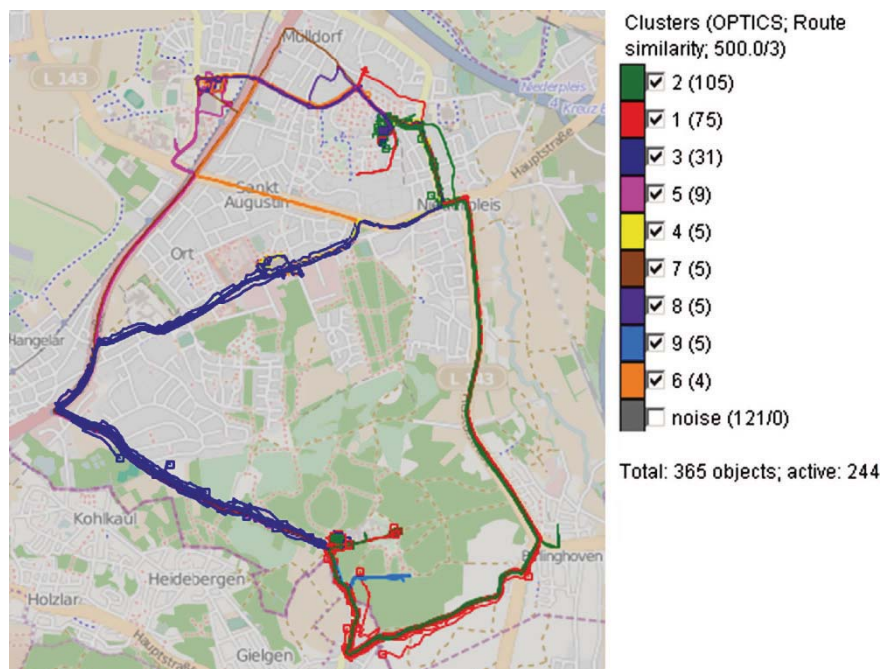


Fig. 1.16 Clusters of car trajectories by route similarity are represented on a map by differently coloured lines

In Fig. 1.17, each cluster is shown separately in a summarized form of flow map, similar to Figs. 1.9 and 1.10. Knowing the person's significant places, we can easily interpret the routes. Cluster 2 (green) consists of trips from home to work. The route represented by this cluster was followed 105 times. Clusters 1 (red), 3 (blue), and 5 (purple) are trips from work to home following three different routes.

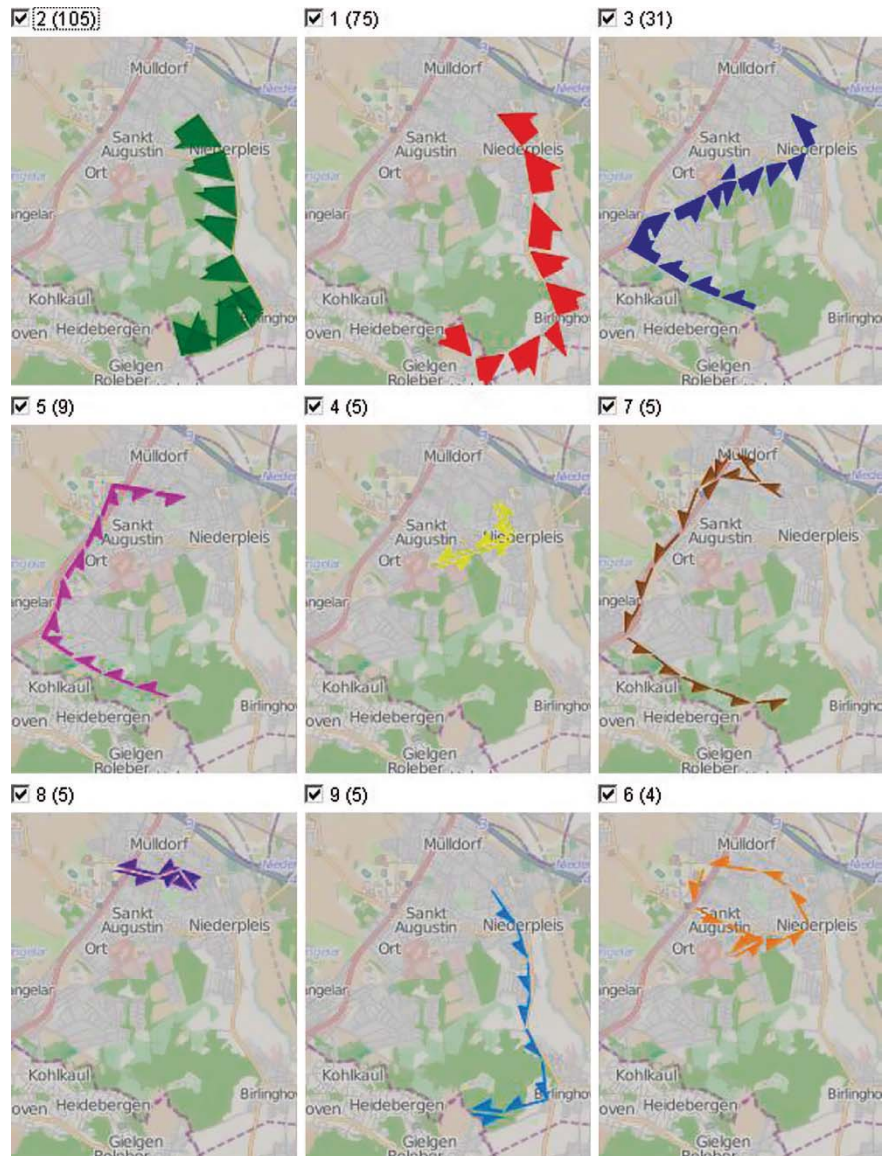


Fig. 1.17 Clusters of car trajectories by route similarity are represented separately in a summarized form (as flow maps)

The first route is opposite to that of cluster 2, and the latter two routes pass the two shopping areas we have discovered before. The first route was followed much more often than the routes through the shops. Cluster 7 (brown) includes five trips from home to work through one of the shopping areas. Clusters 4 (yellow) and 8 (violet) consist of trips from home to these shopping areas and back, and cluster 6 (orange) consists of round trips passing both shopping areas. The trajectories of cluster 9 are similar to those of cluster 2. The difference is that they visit the place in the forest near the work where the tennis ground is located.

Now, we would like to examine and compare the temporal characteristics of the clusters of trajectories using a space–time cube. However, Fig. 1.8 demonstrates us that a space–time cube representing a long time interval may be not very effective. To improve the view and at the same time gain additional information about

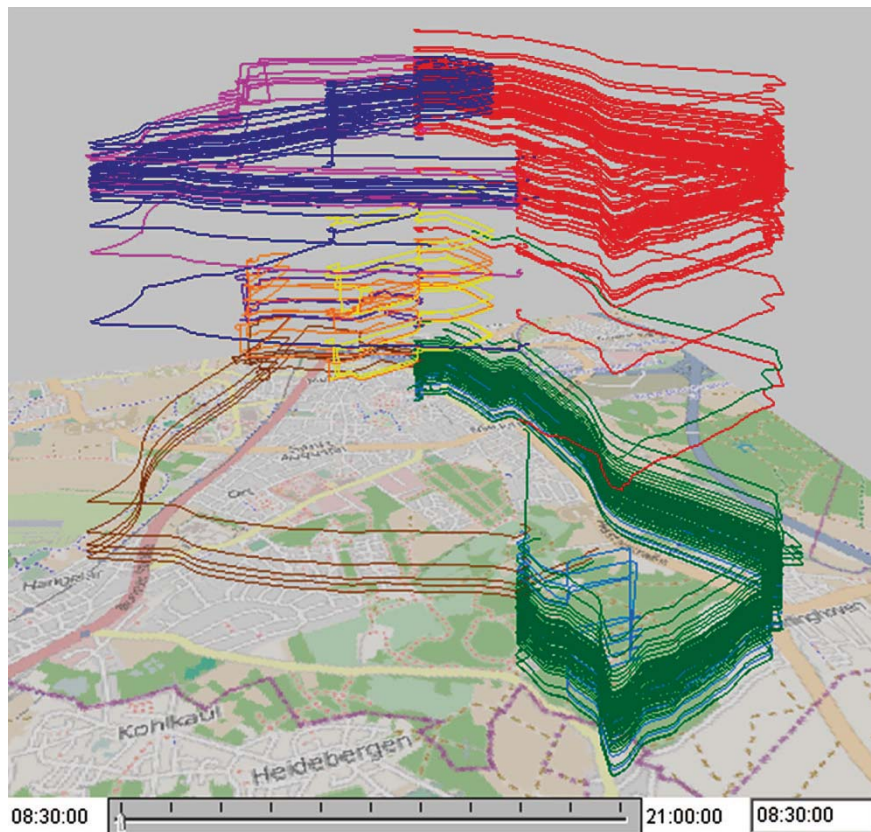


Fig. 1.18 Clusters of car trajectories by route similarity are represented in a space–time cube. The time references in the trajectories have been transformed to times of the same day. Hence, the trajectories are vertically positioned in the cube according to the times of the day when they occurred

temporal characteristics of the trajectories, we can transform the temporal references in the trajectories. One possibility is to shift the trajectories in time to a single day. This means that the dates in the temporal references are replaced by one and the same date, while the times of the day are preserved. The result can be seen in Fig. 1.18. The trajectories are positioned in the cube according to the times of the day in which they took place. We remind the reader that the temporal axis of the display is oriented upwards. As one could expect, the trips from the home to the work (green, brown, and light-blue clusters) occurred mostly in the morning and the trips from the work to the home (red, blue, and purple clusters), mostly in the evening. The trips from the home to the shopping areas (yellow, violet, and orange) occurred mostly in the middle of the day.

Another useful transformation of the time references is demonstrated in Fig. 1.19. The original dates have been transformed to relative position in a weekly cycle starting from Monday and ending with Sunday. Hence, the trajectories are

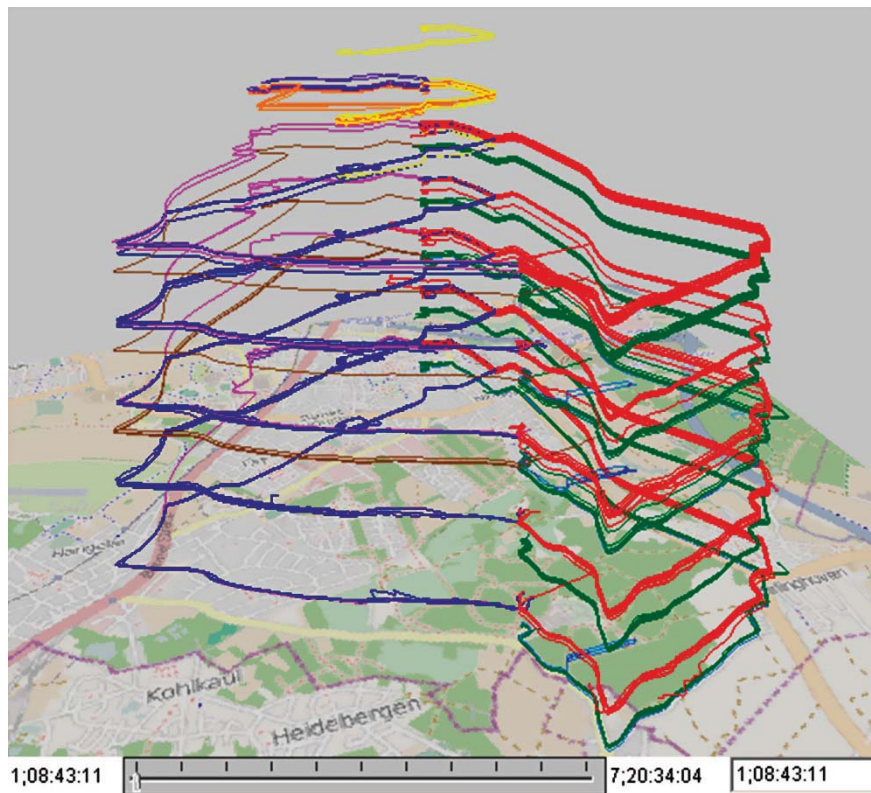


Fig. 1.19 Clusters of car trajectories by route similarity are represented in a space–time cube. The time references in the trajectories have been transformed to relative positions in a weekly cycle from Monday to Sunday. Hence, the trajectories are vertically positioned according to the days of the week when they took place

vertically positioned in the cube according to the days of the week, with Monday at the bottom and Sunday at the top. It can be seen that the trajectories linking home with work occurred on the working days from Monday to Friday and the trajectories from home to shopping areas occurred on the weekend.

In Fig. 1.20, we have transformed the time references to relative times with respect to the trajectories, that is, the starting times of the trajectories have been set to one and the same time moment, and the remaining time references have been adjusted so that the lengths of the time intervals between them are preserved. The trajectories appear in the space–time cube as if they start simultaneously. This transformation allows us to compare the durations of the trajectories, including the stops, and the durations of the stops. Now, we can clearly see that the red and green clusters consist mainly of fast direct trips from home to work and vice versa without intermediate stops (there are only a couple of trajectories in the green cluster that slightly deviate from the main route and have stops). The blue, purple, and brown routes between home and work were usually used for visiting shops

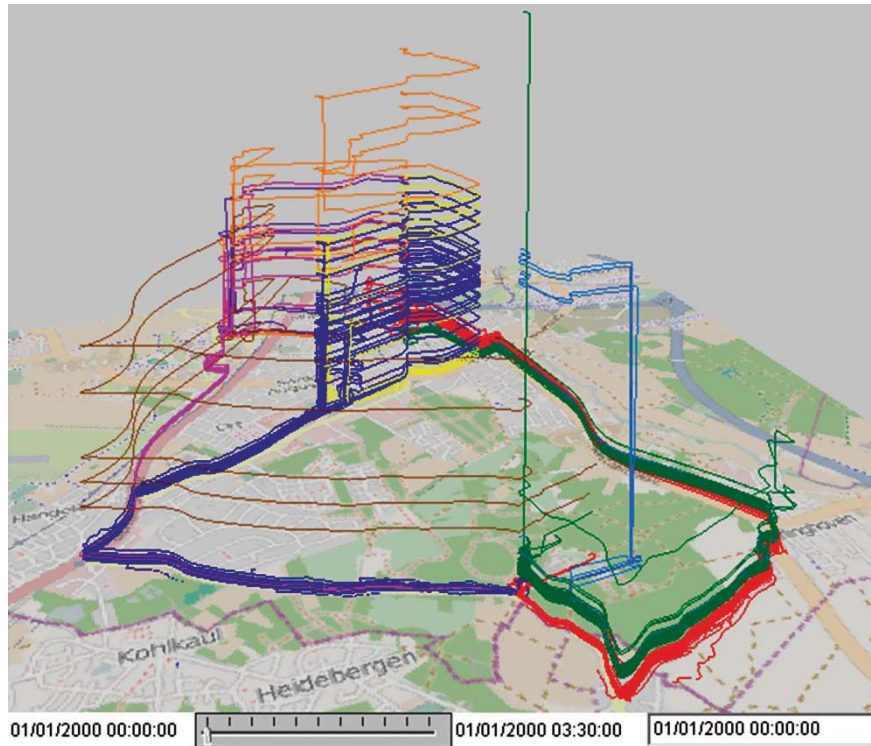


Fig. 1.20 Clusters of car trajectories by route similarity are represented in a space–time cube. The time references in the trajectories have been transformed to relative positions with respect to the starting moments of the trajectories. Hence, the trajectories appear as if they started simultaneously

since the trajectories have stops in the shopping areas. Quite naturally, the round trips from home to shopping areas (yellow, violet, and orange clusters) also have stops in the shopping areas. The trajectories of the orange cluster, which visit both shopping areas, have longer durations than the trajectories that visit only one of the areas. The cube also clearly shows that the trajectories of the light-blue cluster had quite long stops in the forest. This supports our hypothesis that the person might do sports there.

The example dataset we have analysed is just a sequence of time-referenced positions of a car. However, we have managed to learn a lot about the person who drove the car. We now know her home and work place, the places where she usually shops, the times when she does this, and how long it takes. We know the usual routes of the person and the reasons for choosing among them. We know the usual times of driving to work and back home. This knowledge has been obtained by combining computational processing of the data with interactive visual interfaces, which allows us to relate the data to the spatial and temporal contexts and involve our previous knowledge and common-sense reasoning.

1.3 Simultaneous Movements of Many Objects

So far, we have considered movements of a single object (car). Let us take another example dataset with positions of many cars. The data consist of GPS tracks of 17,241 cars in Milan (Italy) collected during one week from Sunday to Saturday. The dataset consists of more than 2 million records each including a car identifier, timestamp (date and time of the day), geographical coordinates, and movement speed. Dividing the movement tracks of the cars by the minimum stop duration of 30 min, as described in Sect. 1.2, produces about 176,000 trajectories. This dataset is much bigger than the one we considered previously. The whole dataset is too large for the kind of analysis we did before. We cannot consider all car trajectories individually. The visual displays turn out to be ineffective even for small subsets of the trajectories. Thus, Fig. 1.21 shows less than 10 % of the car trajectories. The tools for display interaction, for example, zooming on a map or manipulation of the view in a space–time cube, work with significant delays impeding the analysis. The available clustering tools cannot be straightforwardly applied to this amount of data because clustering works in main memory of the computer, whereas the data do not fit there. Therefore, we cannot group all car trajectories by similarity in order to consider and compare the groups as we did in the previous example.

To analyse large datasets, it is necessary to use special analysis techniques relying on database processing. As we have already mentioned, one possible approach to dealing with large amounts of data is aggregation. We have applied spatial aggregation to the movements of the single personal car (see Figs. 1.9 and 1.10). In this case, there are multiple cars that moved simultaneously. One can expect, however, that their collective movements could be different in different times. In order to investigate the differences, we apply spatio-temporal aggregation.

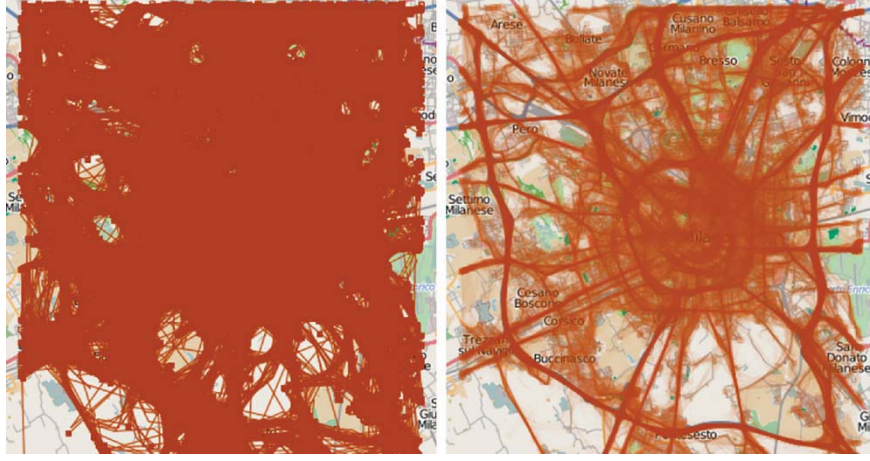


Fig. 1.21 About 14,100 trajectories of cars in Milan made on Monday, 2 April 2007. *Left* the trajectories are drawn in a fully opaque mode. *Right* the trajectories are drawn with 5 % opacity

We divide the space (i.e. the territory of Milan) into compartments and the time span of the data into intervals. For this example, we have chosen hourly intervals; hence, the whole time span of one week has been divided into 168 hourly intervals. Then, we use database operations to compute statistics for the compartments and intervals. Thus, we can ask:

- how many different cars visited each compartment in each interval;
- how many moves (transitions) occurred between two neighbouring compartments in each direction in each interval.

It is also possible to compute other aggregate statistics such as the average (or the minimum, maximum, median, etc.) speed or average time spent in each compartment. Aggregates computed in the database can be loaded in main memory and visualized. In particular, one can use animated maps: one step of the animation corresponds to one time interval in the aggregated data. Thus, Fig. 1.22 presents four screenshots of an animated map showing the counts of different cars that visited the spatial compartments in different time intervals. The screenshots correspond to the intervals 03–04, 04–05, 05–06, and 22–23 h on Monday.

The counts are represented by circles with proportional areas. We can observe how the presence of cars, which reflects the intensity of the city traffic, increases in the morning hours from 03–04 to 05–06 h. In the evening hours, the intensity of the traffic decreases. We have included in Fig. 1.22 only one screenshot from the evening hours. We have selected the interval 22–23 h, in which the counts of the visits were close to those in the early morning interval 03–04 h. However, the overall spatial distributions of the cars are different. In the evening, there were notably more cars in the centre of the city than in the early morning.

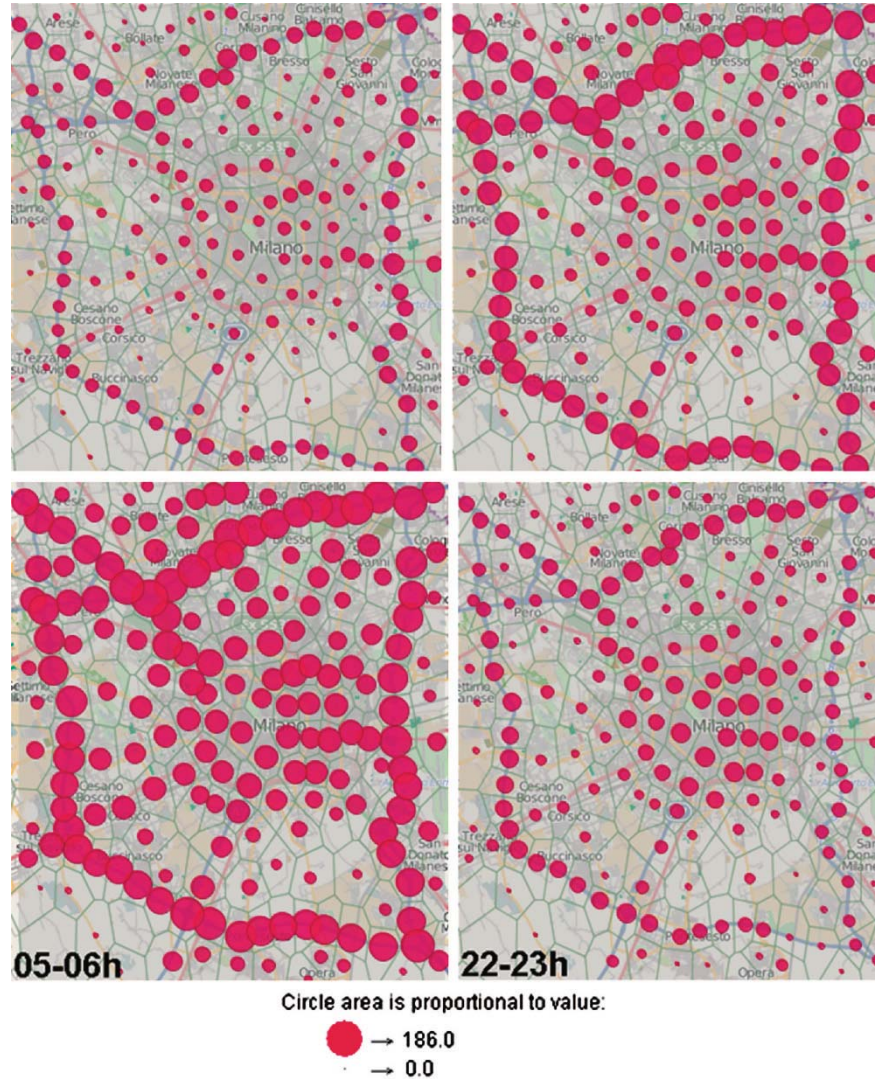


Fig. 1.22 Movement of multiple cars in Milan: counts of the presence of cars by spatial compartments and hourly intervals

Figure 1.23 presents screenshots from an animated flow map showing aggregated moves between neighbouring compartments. The screenshots have been taken for the same time intervals as those in Fig. 1.22. As with the previous map, we can observe a substantial increase in the movement intensity from the interval 03–04 to 05–06 h. The intensity increases first on the belt roads surrounding the city and later in the centre. In the evening, the intensity of the movements in the

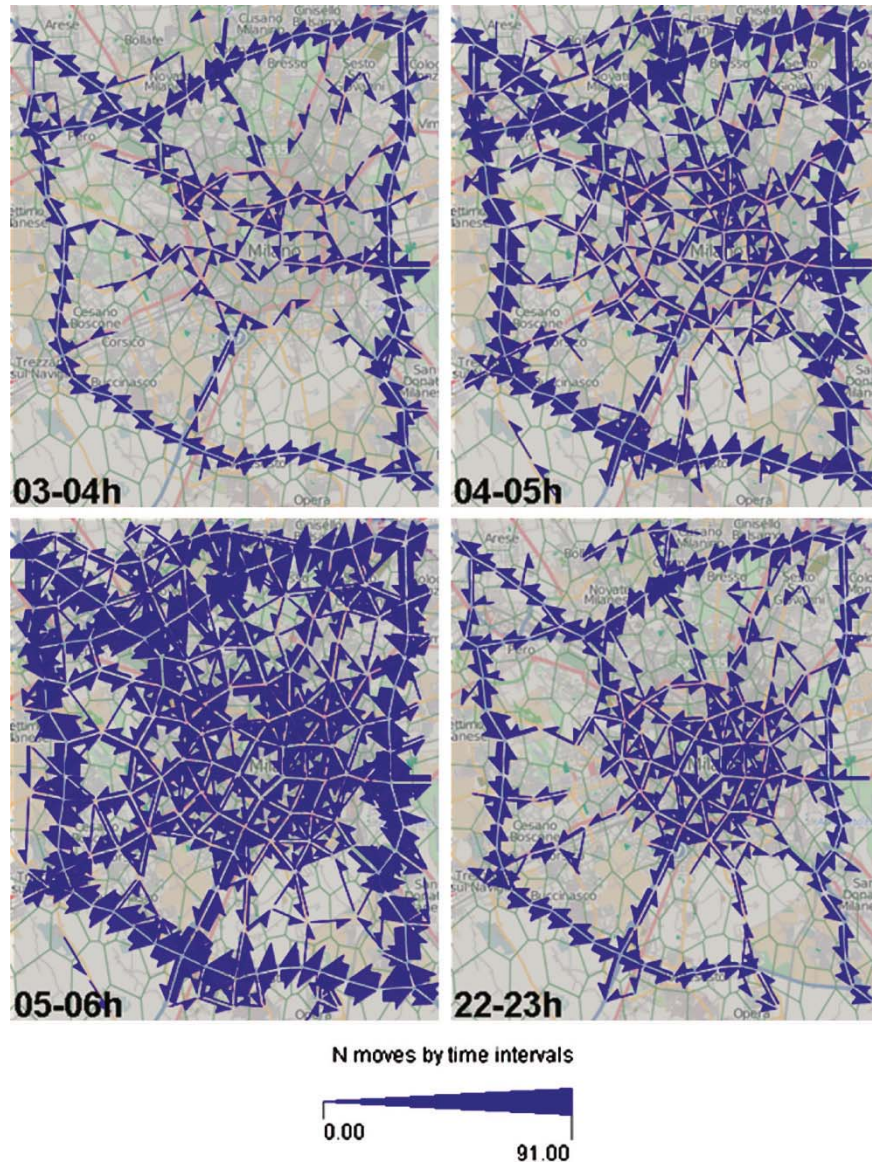


Fig. 1.23 Movement of multiple cars in Milan: counts of the moves of cars between spatial compartments by hourly intervals

centre is higher than in the early morning. Similar observations can be made for other working days.

The dynamics of the movement on the weekend is different. We shall not include more screenshots of animated maps showing the movements on Sunday

and Saturday since they take considerable page space. There is another method to visualize spatio-temporal aggregates: to draw in each spatial compartment a diagram representing the temporal variation in the aggregate values in this compartment. In our case, we have 168 hourly intervals; hence, each diagram should represent 168 different values. We use diagrams in which the values are represented by colouring of small rectangles (pixels); see Fig. 1.24. We use a diverging colour scale blue–yellow–red, where shades of blue are used for low values and shades of red for high values. The pixels are arranged in 24 columns corresponding to the 24 hourly intervals of a day and seven rows corresponding to 7 days from Sunday to Saturday. The row for Sunday is on the top and the row for Saturday on the bottom of the diagrams. The columns are arranged from left to right; the leftmost column represents the interval 00–01 h and the rightmost column, the interval 23–24 h. Hence, each diagram tells us how the car presence varied in the respective place over days of the week and times of the day.

Figure 1.24 includes two map fragments. The upper fragment represents the eastern part of the northern belt road (A4). The lower fragment is taken from the city centre. It can be noticed that all diagrams have blue colours at the left and right edges, which reflects low traffic intensity in the nights. It is also noticeable that the top and bottom rows of pixels, which correspond to Sunday and Saturday, differ from the remaining five rows corresponding to the working days. On the belt road, the morning period of low traffic intensity is longer on the weekend than on the working days. However, the intensities in the afternoons and evenings of Sunday and Saturday are close to those on the working days. The dynamics in the



Fig. 1.24 Movement of multiple cars in Milan: two map fragments with diagrams showing the variation in the presence of cars in the spatial compartments by hourly intervals. The columns of the diagrams correspond to 24 h of a day, and the rows correspond to 7 days from Sunday to Saturday

centre differs from those on the belt road. The presence of cars remains quite low during the entire day on Sunday and Saturday, and it is notably lower than on the working days. In addition, in the mornings of the working days, the presence of cars starts to increase later than on the belt road.

After we have acquired an overall picture of car traffic in Milan, we would like to learn how certain places in the city are connected. In particular, we are interested in how people get from the suburbs to the city centre. We outline the city centre and the major crossings on the belt roads as shown in Fig. 1.25 and again use database operations to compute the total numbers of moves among the areas of interest we have defined. We also compute the numbers of moves by hourly intervals. Figure 1.25 presents the total counts of moves. We see that there were many more cars that moved on the belt roads without going to the city centre than cars that moved to and from the centre to the belt roads. More specifically, the highest number of moves between two crossings on a belt road is 3,245 (from crossing N to crossing NW3), while the highest number of moves between the centre and one of the crossings is 1,794 (from crossing NW3 to the centre). The flows between the centre and crossings NW3 and E are more intensive than between the centre and the other crossings.

Figure 1.26 presents three selected hourly intervals to provide an idea of the temporal variation in the aggregated movements among the areas of interest.

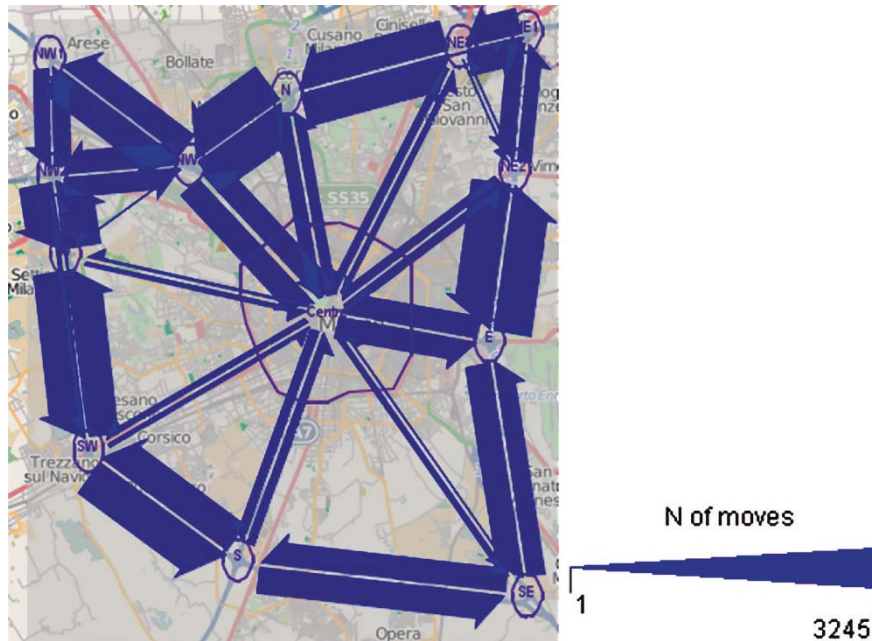


Fig. 1.25 Flows among selected areas of interest in Milan, including the city centre and major crossings on the belt roads around the city

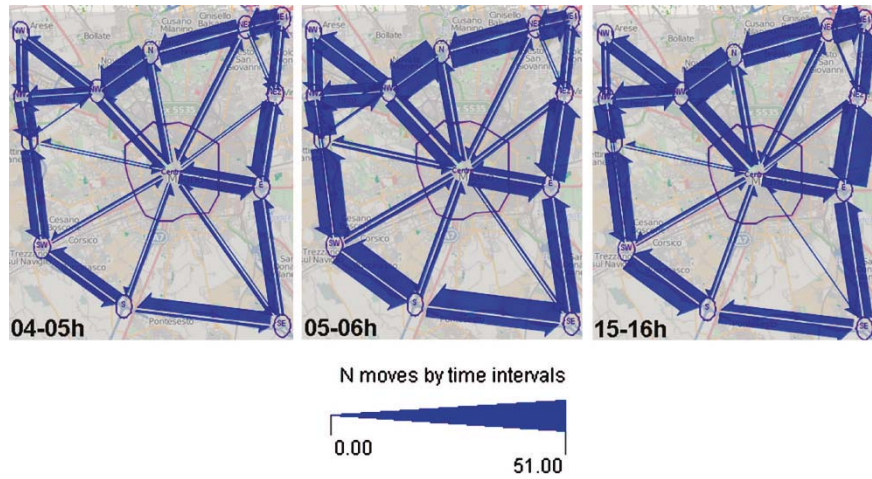


Fig. 1.26 Flows among selected areas of interest in Milan by hourly intervals

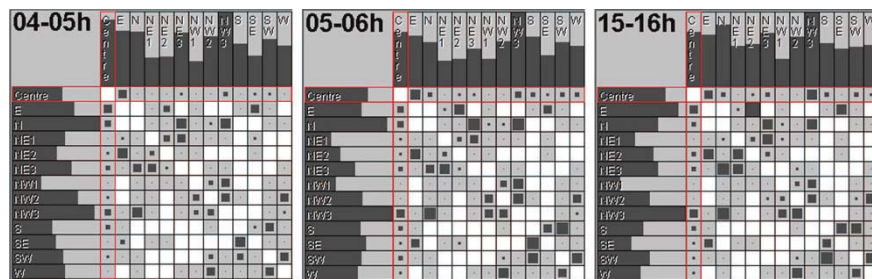


Fig. 1.27 Moves among selected areas of interest in Milan by hourly intervals represented in the form of origin–destination matrix

Another way to represent information about movements among places is to utilize an origin–destination matrix, as shown in Fig. 1.27. The rows and columns of the matrix correspond to the places of interest, in our case, the city centre and major crossings on the belt roads. A cell shows the amount of movement from the place corresponding to the row to the cell corresponding to the column. The numbers can be visually encoded, in our case, by filled squares with areas proportional to the values. The dark-grey bars in the leftmost column (containing the place labels) represent the total amounts of movement from the respective places. The dark-grey bars in the column headers represent the total amounts of movement to the respective places. The three screenshots of the matrix display correspond to the same time intervals as represented by the maps in Fig. 1.26.

The map and the matrices tell us that in the morning, there is more movement to the centre than from the centre, except for the link centre—E (east). In the interval 05–06 h,

there are more movements from the centre to the east than in the opposite direction. Perhaps, many cars go to the airport Linate, which is located on the east. In the afternoon, the flows from the centre increase, especially the flow to the crossing NW3.

Hence, by aggregating the data and exploring the aggregates with the help of interactive visual displays, we could learn a lot about the car traffic in Milan. We have learned how the spatial distribution of the cars and the intensity of movements vary over time. We have investigated the variations in the presence of car in different places by hours of the day and days of the week and discovered differences between the centre and the belt roads. We have studied connections and flows between selected areas of interest. Although we do not know the territory of Milan, maps have provided us with the spatial context and allowed us to use our general knowledge of geographical space, which includes such concepts as city centre, belt roads, and crossings. We have also used our general knowledge of time, in particular temporal cycles (daily and weekly), and differences between day and night, working days and weekends, and so on.

1.4 What Should Have Been Achieved by These Examples

The examples allowed us to introduce informally the major concepts we shall be dealing with throughout the book:

- position records and movement tracks;
- trajectories;
- dynamic (time-dependent) attributes of movement, such as speed;
- properties of trajectories: start and end positions in time and space, route, stops on the way, and speed variation;
- spatial events, such as stops;
- flows (summarized movements) between places;
- spatial situations: spatial distribution of multiple moving objects at different times and aggregate characteristics of their movement, such as intensity of flows among places;
- local dynamics (temporal variations) of presence and movements in places;
- spatial context of the movement, which was conveyed by the maps;
- temporal context of the movement, in particular, daily and weekly cycles.

We have also demonstrated a number of transformations of movement data:

- division of movement tracks into trajectories representing different trips;
- extraction of events, such as stops;
- spatial and spatio-temporal aggregation;
- transformations of time references.

We have touched upon the use of clustering in analysis of movement-related data. Clustering of events allowed us to find significant places, and clustering of trajectories uncovered habitual routes.

In our example analyses, we have used a variety of interactive visualization techniques. The most common techniques for visualizing trajectories and events are the map and space–time cube. These can be complemented by time graphs and other temporal displays, which are more effective in representing time. Diverse displays can be dynamically linked, which means that interactive operations performed by the user on one of the displays are somehow reflected in the others. For example, in Fig. 1.7, the map display marks the spatial position corresponding to the temporal position of the mouse cursor within a time graph. We have shown which techniques can be used to visualize aggregated movement data; in particular, we have introduced flow maps and origin–destination matrices showing summarized movements among places.

Besides introducing major concepts and demonstrating some of the analytical techniques used for movement data, the role of the examples was to show how the capabilities of the computer and human can be combined for extracting knowledge from data. Movement data are usually semantically poor as they basically consist of coordinates and timestamps. This was the case in our examples. However, by analysing the datasets, we have learned much about the life and habits of the car owner in the first example and about the city traffic in Milan in the second example. The computer helped us to generate data abstractions, to find similar occurrences and repeated patterns, to extract what we deemed potentially interesting, and to transform the data for considering them from multiple perspectives. The computer also did an extremely important thing: it represented the data and their derivatives on visual displays and allowed us to interact with the displays. This enabled us to use our human-specific capabilities to perceive patterns and grasp their meaning, to establish associations (link data and patterns with the context, link different perspectives to an integral mental picture, link new information to previous knowledge, etc.), to generate hypotheses, to reason, and to make conclusions.

Such human–computer analytical processes in which computers not only process data but also enable humans to involve their unique capabilities to perceive, associate, hypothesize, reason, and comprehend are a major topic of *visual analytics*.

1.5 Visual Analytics

Visual analytics is a relatively new term; it has been in use only since 2005 when the book “Illuminating the Path” was published (Thomas and Cook 2005). However, the kinds of ideas, research, and approaches that are now termed visual analytics emerged much earlier. The main idea of visual analytics is to develop knowledge, methods, technologies, and practice that exploit and combine the strengths of human and electronic data processing (Keim et al. 2008, 2010). Visualization is the means through which humans and computers cooperate using their distinct capabilities for the most effective results. This idea

has penetrated many research efforts in the areas of information visualization, GIScience, geovisualization, and data mining long before 2005 (Andrienko et al. 2010).

Since 2005, an attempt has been made to establish visual analytics as a specific scientific discipline in order to consolidate the relevant research that has been conducted within different disciplines and to stimulate its further progress. The distinctive features of visual analytics research are as follows:

- emphasis on data analysis, problem solving, and/or decision making;
- leveraging computational processing by applying automated techniques for data processing, knowledge discovery algorithms, etc.;
- active involvement of a human in the analytical process through interactive visual interfaces;
- support of the information provenance, that is, how each piece of information and knowledge has been obtained;
- support for the communication of analytical results to relevant recipients.

As a science, visual analytics develops its theoretical foundations. Since visual analytics is largely about transforming data to information and knowledge, the theoretical part of visual analytics describes the possible types of data, as well as the types of things or phenomena that can be represented by the data, and determines the types of information and knowledge that can be extracted from the data. The theory of visual analytics also grounds the possible approaches to extracting knowledge and information from the data. In these approaches, it defines the distribution of the workload between the computer and the human analyst so as to relieve the human from routine operations but utilize the human capabilities of abstractive perception and creative analytical thinking.

Space and time are considered as key topics in visual analytics research (Keim et al. 2010; Andrienko et al. 2010). Data with spatial and temporal components (including movement data) are inherently complex as a result of the complexities of space and time, in particular, their heterogeneity, the abundance and diversity of objects populating them, events and processes occurring in them, and the variety and multitude of spatial, temporal, and spatio-temporal properties and relations. Spatial and temporal data need to be analysed with a proper consideration of the spatial and temporal *context*, which includes all these complexities. It is hardly possible to formalize all aspects of the context and feed them to computers for fully automatic processing. Therefore, exploration and analysis of spatial and temporal data rely on the human analyst's tacit knowledge of space and time and space-/time-related experiences. These are incorporated in the analysis through the use of appropriate visual representations and interactive facilities.

The specifics and complexities of space and time and the directions for the visual analytics research related to space and time are considered in the dedicated chapter of the book by Keim et al. (2010) and in the paper by Andrienko et al. (2010). In our book, we shall consider the specifics and complexities of movement data and visual analytics approaches to *analysing* the data and extracting various kinds of knowledge.

1.6 Structure of The Book

[Chapter 2](#) presents the conceptual framework for the analysis of movement. It describes the types of information contained in movement data and defines the possible types of tasks in analysing movement. To enable extraction of various types of information, movement data may need to be converted to different forms. [Chapter 3](#) deals with the possible transformations, which can adapt available movement data to the analysis goals or to specific requirements of the methods that the analyst wants to apply, extract relevant parts of the data, or reduce irrelevant details.

[Chapter 4](#) describes basic visualization and interaction techniques that enable viewing and exploration of movement data and other types of spatio-temporal data and facilitate data transformations and joint analysis of different data types. These techniques provide general infrastructure for applying specific visual analytics methods and procedures and for method combination.

[Chapters 5, 6, 7, 8](#) are dedicated to the analytical methods and procedures that can be used for analysing movement data. Besides the state-of-the-art methods that have been previously published by the book authors and other researchers, there are a number of new methods that have not been published before. The methods are presented in a systematic way, being grouped according to the possible foci in movement analysis: movers ([Chap. 5](#)), spatial events ([Chap. 6](#)), places ([Chap. 7](#)), and times ([Chap. 8](#)). Most of the methods combine visual and computational techniques. The latter are typically not our original inventions but state-of-the-art techniques from statistics, data mining, and database processing. We have integrated them with interactive visual interfaces to support synergistic work of the computer and human. The work of each method is explained by richly illustrated examples, for which we have used a number of interesting and challenging datasets. The datasets are introduced in [Chap. 2](#).

We conclude in [Chap. 9](#) by showing the connections between the pieces presented in the previous chapters and presenting a general methodological framework for analysing movement behaviours in all their aspects.

References

- Andrienko, G., Andrienko, N., Demšar, U., Dransch, D., Dykes, J., Fabrikant, S., et al. (2010). Space, time, and visual analytics. *International Journal Geographical Information Science*, 24(10), 1577–1600.
- Keim, D., Andrienko, G., Fekete, J-D., Görg, C., Kohlhammer, J., & Melancon, G. (2008). Visual analytics: Definition, process, and challenges. In A. Kerren, J.T. Stasko, J-D. Fekete, C. North (Eds.), *Information visualization: Human-centered issues and perspectives* (Vol. 4950, pp. 154–175) of LNCS state-of-the-art survey, Berlin: Springer.
- Keim, D., Kohlhammer, J., Ellis, G., & Mansmann, F. (Eds.). (2010). *Mastering the information age. Solving problems with visual analytics*. Eurographics Association, Goslar, Germany. Electronic version: <http://diglib.eg.org>.
- Thomas, J. J., & Cook, K. A. (2005). *Illuminating the path: The Research and development Agenda for Visual Analytics*. New York: IEEE Computer Society Press.

Visual Analytics of Movement

Andrienko, G.; Andrienko, N.; Bak, P.; Keim, D.; Wrobel, S.

2013, XVIII, 387 p. 200 illus., 178 illus. in color.,

Hardcover

ISBN: 978-3-642-37582-8