23 Multimodal Analytical Visualisation of Spatio-Temporal Data

Gennady Andrienko and Natalia Andrienko

Fraunhofer Institute for Autonomous Intelligent Systems, Sankt Augustin, Germany

23.1 Introduction

The concept of multimedia cartography has appeared as a natural response to the new opportunities provided by computer technologies as compared to the paper-based technologies of the past. It is generally believed that multimedia can convey the multifaceted and dynamic character of the spatial environment much more adequately than paper (Peterson 1999). The scope of multimedia cartography includes not only combination of maps with images, sound, movies, etc. but also dynamic (animated) maps, interactive maps, linked maps (e.g. in electronic atlases), three-dimensional images and virtual reality. In principle, combination of maps with statistical graphics can also be viewed as a kind of multimedia cartography; however, this does not seem to be a quite traditional view.

By looking at the current topics and trends in the research on multimedia cartography, one may observe a shift towards one of the major roles played by maps, namely, communication of spatial information. The other important role of maps, to model the reality for the purposes of exploration and analysis, seems to get less attention.

An opposite shift can be seen in another part of the cartographic research community, which is lead by the ICA Commission on Visualisation (see http://www.cartoweb.nl/icavis/) with the research agenda stressing the role of highly interactive maps in hypothesis generation, data analysis, and decision support (MacEachren and Kraak 1997). Combination of interactive maps with other instruments for data analysis, including statistical graphics and other visualisation techniques as well as data manipulation and computational tools, is among the mainstream topics in geovisualisation research. It is generally recognised that exploration of non-trivial datasets requires the data to be viewed from diverse perspectives, which is supported by multiple dynamically coordinated displays of various types.

Currently, it is not usual to apply the term “multimedia” to such display combinations (which indispensably include maps when spatial data are dealt with). Therefore, we have used the term “multimodal visualisation” in the title of our paper, which tells about using diverse types of interactive displays for exploration of spatio-temporal data, i.e. data having both spatial and temporal components. Although the research topic represented in the paper cannot be viewed as mainstream in multimedia cartography, we hope to provide a useful complement to the traditional themes. At least, we do not see any particular reason for the interests of multimedia cartography to be limited to presentation and communication of known information and not include the use of multiple media to support exploration of spatial data and discovery of previously unknown patterns, trends, and relationships. So, we consider ourselves as representatives of an ‘exploratory trend’ in multimedia cartography, and our paper is about the use of interactive and dynamic maps in combination with other types of data displays for exploration of spatio-temporal data.

Cartographic representation of time attracted the attention of cartographers and map designers for a long time (Bertin 1983, Vasiliev 1997). Starting from the pioneering work of Tobler (1970), computers are used for creation and distribution of dynamic maps, in which display time is used to represent real time. However, it is now generally recognised that dynamic maps alone are insufficient for a comprehensive exploration of spatio-temporal data. Hence, there is a general research problem of devising combinations of tools that could adequately support data exploration. The choice of tools depends on the structure and characteristics of data to be explored.

There are three basic types of spatio-temporal data according to the type of changes that occur in time (Blok 2000):

Existential changes: appearing, disappearing, reviving of objects or/and relationships;

Changes of spatial properties of objects (location, size, shape); and

Changes of thematic properties, i.e. values of attributes.

In this chapter, we focus on the exploration of changes of thematic properties of spatial objects. More precisely, we consider data consisting of multiple spatially referenced time series of numeric values. Visualisation of such data in a dynamic map supports perceiving the pattern of spatial distribution of the values at any moment in time and the evolution of the pattern over time. However, this technique is inadequate for such exploratory tasks as analysing the temporal behaviours of the values in various places, understanding the spatial distribution of these behaviours and detecting behaviours with particular characteristics. The objective of this paper is to suggest visualisation techniques and tools that can support the exploration of the temporal behaviours. In combination with dynamic maps, these techniques and tools enable comprehensive exploration of the various aspects of spatially referenced time-series data.

In the next section, we briefly review the methods for the visualisation and exploration of time series developed in cartography and other disciplines. In the sections following this, we define the exploratory tasks relevant to the exploration of temporal behaviours and propose ways to support them.

23.2 Visualisation of Spatial Time-Series in Computer Cartography and Statistical Graphics

Visualisation of spatio-temporal data has been in the focus of many researchers in computer cartography. Thus, Peuquet (1994) proposed a conceptual framework for the representation of spatio-temporal dynamics. Koussoulaku and Kraak (1992) refined the theory of Bertin concerning the types of questions that can be asked about spatio-temporal data. Harrower et al (1999) considered human-computer interaction issues and developed a powerful and convenient user interface for temporal focusing and temporal brushing that accounts for linear and cyclical nature of time. Some researchers study the opportunities provided by emerging technologies such as Tcl/Tk, Java, and SVG (see, for example, Dykes et al 2005). Current approaches to visualisation and exploration of spatio-temporal data are surveyed in (Andrienko et al 2003). The most commonly used techniques include

1. Representation of several time moments/intervals on multiple maps;
2. Dynamic maps, with various possibilities for user control;
3. Change maps, which represent differences between situations at two time moments (Slocum *et al* 2000); and
4. Space-time cube, a 3D visualisation where two dimensions represent geographical space and the third dimension represents time (Hedley et al 1999, Kraak 2003).

Such methods have been applied, in particular, in electronic atlases (Cartwright *et al* 1999) providing access to historical data to a wide community of users and stimulating users to explore such data.

The cartographic community is not the only group of scientists interested in methods to represent and explore time series data. Thus, numerous computational methods have been developed in data mining (Keogh and Kasetty 2003). Researchers in statistical graphics, information visualisation, and human-computer interaction proposed some interactive visualisation techniques for time series. Papers by Unwin and Wills (1988) and Hochheiser and Shneiderman (2004) introduced basic graphical and interaction facilities for enhancing analytical capabilities of time series plots, or time graphs:

1. Interactive access to values through graphics by pointing on a line;
2. Tools for overlaying lines to enable their comparison, which allow the user to align, smooth, stretch, and shrink the lines for a better correspondence;
3. Possibilities to select lines having particular characteristics, such as specific values at a given time moment or interval, specific profiles, etc. ; and
4. Dynamic linking between a time graph and other information displays (scatter-plots, histograms, maps etc.) by identical marking of visual items corresponding to user-selected objects.

The objective of this chapter is to find such a combination of techniques that could adequately support the exploration of a spatially referenced collection of time series (temporal behaviours). To do this, we apply a task-analytical approach, i.e. start with identifying the potential exploratory tasks pertinent to such data.

23.3 Visualisation of Local Behaviours

We use the term 'local behaviour', or simply 'behaviour', to denote the temporal variation of attribute values in a particular place. A dataset consisting of attribute values referring to a set of spatial locations and a sequence of time moments can be viewed as a collection of local behaviours, where each behaviour refers to one of the locations. This is not the only possible view. The data can also be treated as a sequence of spatial distributions of the attribute values, where each distribution refers to one of the time moments. While proper consideration of both aspects is required for a comprehensive analysis of the data (Andrienko and Andrienko 2005), this paper focuses on the former view. Moreover, we further limit our focus to behaviours formed by values of a single numeric attribute.

In the course of the exploration of a collection of spatially distributed behaviours, an analyst may be interested to find answers to various questions:

- What is the general dynamics of values over the entire territory?
- What are the general features of the local behaviours in a given area and how do they compare to the behaviours on the remaining territory?
- What spatial clusters have similar behaviours.
- What locations with the behaviours having specific features? Are these locations neighbours, or, in other words, do they form a spatial cluster? (Here some examples of the features that may be looked for are: persistently low (or high) values; high value fluctuations; continuous increase (or decrease) of values during a given time period).

Let us look what tools and techniques may be helpful to a data analyst in finding the answers. For our investigation, we shall use the dataset containing the statistics of crimes in the USA published by the U.S. Department of Justice. The data is available at the URL http://bjsdata.ojp.usdoj.gov/dataonline/. This dataset contains values of 21 numeric attributes referring to 51 states of the USA and to 41 time moments, specifically, the years from 1960 to 2000. The attributes include the population number and the absolute numbers and rates of the crimes of different types, for example, burglaries, motor vehicle thefts, murders, etc.

For studying a single temporal behaviour, the visualization of the data on a time graph is typically used. The behaviour of a numeric attribute often appears on such a graph as a line, which results from connecting the positions corresponding to the values at consecutive time moments. When multiple behaviours are explored irrespectively of the geographic space, they can be represented in a single display as lines drawn in a common coordinate framework. Figure 1 shows a time graph with 51 lines corresponding to the behaviours of the attribute ‘Burglary rate’ in all the states. This representation can be used for getting the first rough idea about the dynamics of the burglary rates over the entire country: overall increase until 1980 followed by a period of gradual decrease. However, the cluttering and overlapping of the lines make the display hardly legible.

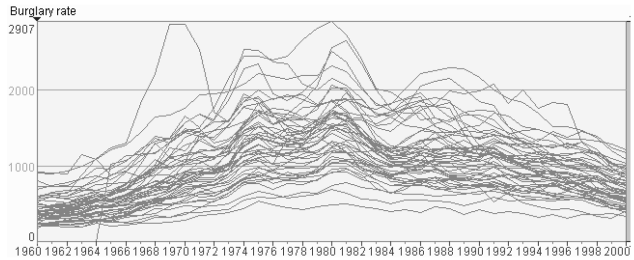

Fig. 1. Behaviours of the burglary rates in the states of the USA are represented on a time graph display as multiple overlaid lines.

The time graph can be dynamically linked to a map display: positioning the mouse cursor on a line or clicking on it can result in this line being specially marked (highlighted) as well as the corresponding state on the map. The link can also work in the opposite direction, i.e. selection of a state on the map highlights the corresponding line on the graph. By selecting several states, one can find answers to questions concerning the general features of the local behaviours in this or that area. However, since the user can only see which behaviours correspond to the selected locations but not where the remaining behaviours are situated in space, it is very difficult to understand how the character and distinctive features of the temporal behaviours vary over the whole territory.

A suitable graphical representation for the latter task is shown in Figure 2: the local behaviours are represented by symbols superimposed on a map according to their spatial references, i.e. at the locations of the respective states. The form of the symbols is a modification of the time graph technique: the coordinate frame is omitted and the lines are complemented to closed shapes with internal filling for a better visibility against the cartographical background.

In such a map, we can see how different behaviours are distributed over the territory of USA. Thus, we can observe that the states in the north-central part of the country had lower burglary rates than in the other states during the whole time period from 1960 to 2000. Another observation is that the states on the west and southwest have higher peaks in the middle of the time interval than the states on the east (with a few exceptions). It is possible to see some spatial clusters of states with similar temporal behaviours of the burglary rate: (1) west-southwest-south, (2) middle north, (3) the area around the Great Lakes, (4) centre and southeast, except for Florida. The clusters are roughly outlined in Figure 3. By visual inspection of the map, an explorer can also find locations with the corresponding behaviours having specific features and, naturally, immediately see whether these locations form a spatial cluster. It is also possible to detect essential common features of behaviours in a certain area.

Fig. 2. A cartographic representation of the spatial distribution of the behaviours of burglary rates over the USA

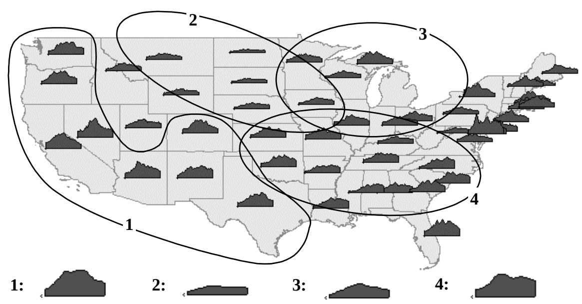

Fig. 3. Spatial clusters of states with similar temporal behaviours of the burglary rate. Below the map, the typical behavioural patterns for each cluster are schematically shown.

In the examples we have considered in this section, a purely visual exploration has been quite effective due to the small dimensionality of the data in the spatial, temporal, and thematic aspects. Thus, we have dealt with a single attribute, short time series, and a small number of spatial locations. Therefore, we have been able to review and compare all the local behaviours represented on the map. However, for the exploration of more complex data, it is necessary to use more sophisticated methods or combinations of techniques. Consideration of multiple attributes is not within the scope of this chapter. For the remaining complexities, i.e. the length of the time series and the number of different locations, we propose some new methods that shall be discussed in the next section. Although we use the same example data, the methods are also applicable to substantially larger data volumes.

23.4 Combining Tools for Behaviour Exploration

One possibility to overcome the dimensionality problem is to use methods of computational statistics and data mining. For example, an explorer can use cluster analysis for replacing numerous spatial locations by groups made of locations with close characteristics. Spatially aware clustering methods can create geographically continuous aggregates. However, this solution entails several problems. First, such methods usually require significant efforts for data preparation. Second, the outcomes are very sensitive to method parameters. An explorer needs to do multiple trials with assigning different values to the parameters and then choose the most appropriate result. Third, execution time may be rather long. Finally, the results are often difficult to interpret.

The approaches we are going to suggest are based on simple, easily understandable calculations enhanced by interactive dynamically linked data displays. We shall demonstrate that, despite of the simplicity, these techniques can be quite effectively applied to large datasets.

23.4.1 Getting the General Picture of the Behaviour on the Entire Territory

In Figure 1, we have demonstrated a time graph display with multiple overlaid lines. Already with fifty local behaviours, this display suffers from cluttering and overlapping. It becomes completely unusable when this number increases to hundreds or thousands. In Figure 4, bottom, we propose an aggregation-based alternative to a time graph display representing multiple behaviours.

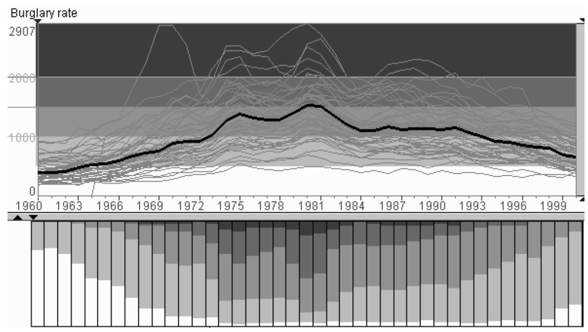

Fig. 4. The value range of the attribute has been divided into intervals. The segmented bars show for each year the proportions of values over the whole country fitting in these intervals.

For producing this visualization, the value range of the attribute "Burglary rate" has been divided into 5 intervals by introducing breaks 500, 1000, 1500, and 2000. For each year, the display contains a segmented bar with the differently shaded segments showing how many values in this year over the whole country fit in the corresponding intervals. The white segments correspond to the values below 500, the light grey – to values between 500 and 1000, and so on (the higher the values, the darker the colour). The background shading of the time graph display above the segmented bar display visually illustrates the principle of the aggregation and the current division of the attribute value range.

The aggregated display at the bottom allows us to do important observations concerning the general behaviour of the burglary rates on the whole territory over the time. Thus, we can see that in the early 1960s the burglary rates in most states were below 500 and in a part of states (from one-fourth to one-third) between 500 and 1000. Starting from the mid-1960s, the burglary rates over the whole country increased and reached their extremes in 1980 and 1981, when the proportions of the states with the rates below 500 and between 500 and 1000 were the smallest and the proportions of the states with the values between 1500 and 2000 and more than 2000 were the largest. After 1981, the situation improved slightly and remained more or less stable until the beginning of 1990s, when it started to gradually improve. This trend was preserved until the end of the time period under study. However, the low criminality level of early sixties was not reached again.

The thick black line on the time graph above the aggregated display connects the yearly countrywide median values and thereby provides an additional generalized view of the dynamics of the burglary rates over the country.

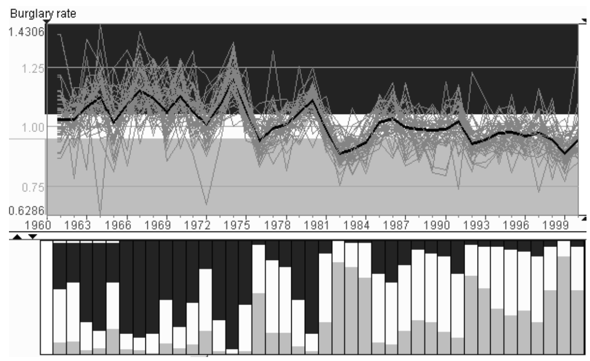

Fig. 5. The time graph has been transformed to show the annual changes in the individual states. The lower image aggregates the changes for each year.

The aggregation technique can be combined with various transformations of the original attribute values, such as computing the changes from year to year, as is shown in Figure 5. Here, the burglary rates in each year have been transformed into their ratios with respect to the values in the same states in the previous year. The resulting values higher than 1.0 correspond to annual increase and values below 1.0 – to decrease. The time graph at the top of Figure 5 represents the transformed values instead of the original ones.

As in the previous case, we have applied the division of the whole range of the transformed values into intervals. This time, choosing two breaks, 0.95 corresponding to 5% decrease and 1.05 corresponding to 5% increase, and thereby obtained 3 intervals. The resulting aggregated display is instrumental in finding years and periods of significant countrywide increase or decrease of values. The light grey colour corresponds to the decrease by more than 5%, white – to the changes between 5% decrease and 5% increase, and dark – to the increase of the values by more than 5%. It may be seen that 1974 and 1980 were the years of almost overall rapid increase and 1999 was the year of general rapid decrease. Starting from 1992, the crime rates mostly decreased or changed by less than 5% in comparison to the previous year.

Hence, this simple aggregation technique proves quite useful in application to both the original data and the transformed ones. The technique is scalable in respect to the number of objects and to the length of time series and the computational complexity is linear with regard to both dimensions. The limitation for the visualization is the screen size and resolution. The time graphs accompanying the aggregated displays in Figures 4 and 5 are not necessary for the sort of exploratory tasks considered and may be omitted.

23.4.2 Finding Spatial Patterns of Similar Local Behaviours

Another way of aggregating the same data is illustrated in Figure 6. The display contains a so-called ‘envelope’, that is, a polygon enclosing all the lines representing the local behaviours. Actually, the lines themselves are not needed for building the envelope but only the minimum and maximum values over the country in each year are needed. The envelope therefore shows how the range of burglary rates changes over time. In addition to the value range in each year, the median and quartiles are shown. The positions of the corresponding positional measures in consecutive years are connected so that the original envelope is divided into four polygons. The polygons are shaded using alternating light and dark grey, which makes them clearly visible and distinguishable.

In principle, there is a danger that a viewer may consider these polygons as containers of certain subsets of lines whereas they are just indicators of the positions of each year’s median and quartiles, and individual lines may cross the boundaries of the polygons, as may be seen in Figure 7. Perhaps, it would be less misleading to mark somehow these positions without connecting them, but such a display would be much more complex to perceive.

Providing that the meaning of the polygons is understood correctly, one can effectively use them for data analysis. The display gives us a summarized picture of the countrywide situation with the burglary rates in each particular year and allows us to compare the situations in different years. We can also get an idea of the overall trend of the burglary rates over the country during the whole period from 1960 to 2000 or any of its subintervals. Thus, an increase of median and quartile values indicates the overall increasing trend of the burglary rates, and the same for decrease. For example, a clear decreasing trend can be observed on the interval from 1991 to 2000. Moreover, using the properties of the positional measures, we can isolate this observation by giving some numeric estimation. In 1991, more than a half of the states had burglary rates over 1000, whereas in 2000, the burglary rates in more than 75% of the states were below 1000. We can easily see the period of the highest burglary rates from 1977 to 1982, when the rates in at least 75% of states were over 1000. The synchronous peaks of the values of the median and the quartiles in 1975 and 1980-1981 may be also worth attention as well as the rather steep decrease from 1981 to 1984.

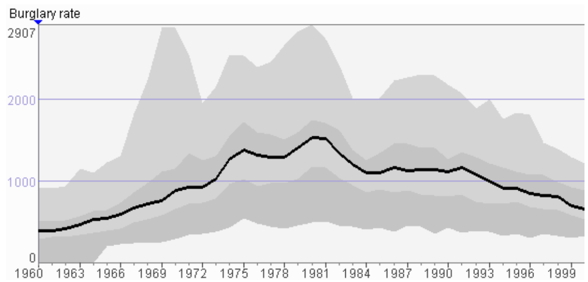

Fig. 6. The aggregate time graph represents the changes of the value range and the quartiles over time.

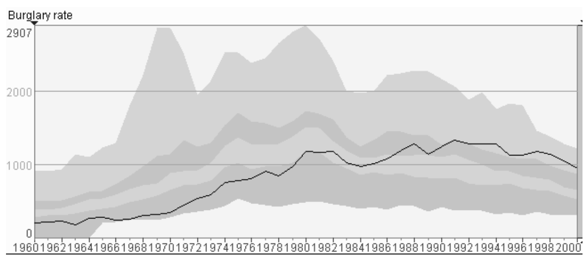

Fig. 7. The line representing the local behaviour in Mississippi is superimposed on the aggregate time graph.

The computed yearly medians and quartiles can be further used for data simplification. The original numeric values can be replaced by the numbers of the respective quarters. Thus, the line of Mississippi in Figure 7 was in the lower quarter of the value distribution from the beginning of the time period until 1980, i.e. for 21 years. Accordingly, the original values of the burglary rates in this state for the years from 1960 until 1980 are encoded into the value 1. Then, from 1981 up to 1986, the burglary rate values were in the second quarter and, hence, are encoded into 2. Analogously, the values for the years 1987-1991 are encoded into 3, and starting from 1992 – into 4. In the same way, the values for all other states are transformed.

For each so transformed time series, the occurrences of each quarter number are then counted. For example, the counts for Mississippi are 21, 6, 5, and 9, respectively. The counts are presented in Figure 8 by bar charts positioned on a map. Each symbol consists of four bars corresponding to the quarters with the sizes of the bars being proportional to the respective counts. Additionally, the shading of each state corresponds to the dominating quarter, i.e. having the maximum number of occurrences among the four quarters. Thus, the dominating quarter for Mississippi is 1.

Fig. 8. Computationally supported detection of clusters of similar local behaviours

It may be noted that, in the result of this automated procedure, we have received spatial clusters (groups of neighbouring identically shaded states) very similar to the clusters in Figure 3 revealed 'manually'.

The transformed data can also be analysed in another way. For each local behaviour, the average (mean) quarter and the variance can be computed. In Figure 9, bottom left, these values are represented on a scatterplot. Using the dynamic link between the scatterplot and the time graph, one may find and explore behaviours with particular characteristics in terms of the average and variance. Thus, in Figure 9, the scatterplot has been used to find the behaviours with low variance and the values being mostly either in the lowest or in the highest quarter. The corresponding lines are highlighted on the time graph. Similarly, in Figure 10, the states with high variability of values have been selected through the scatterplot. The map fragment shows only the eastern part of the country where all the selected states are located.

It may be noted that the time graph displays in Figures 7, 9, and 10 combine an aggregated representation of the entire set of local behaviours and a detailed representation of selected behaviours. This technique compensates for the deficiencies of both representations being used separately. Despite the aggregation, individual data can be viewed in sufficient detail and compared with the general characteristics of the entire dataset. Since the detailed representation is only applied to selected items, cluttering and overlapping of display elements is considerably reduced.

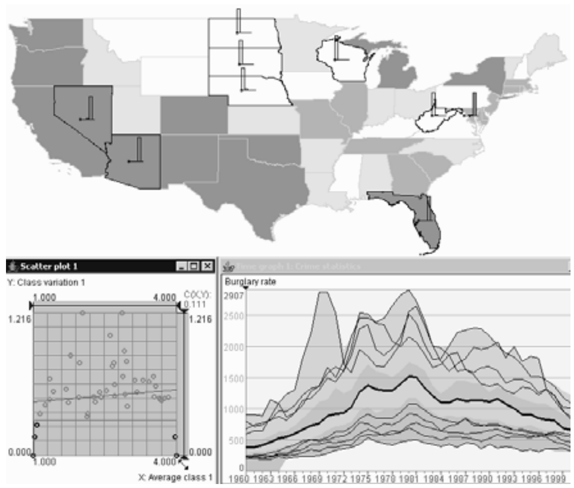

Fig. 9. Selection of the states with consistently low (shaded in light colour) and consistently high (shown in dark) burglary rates

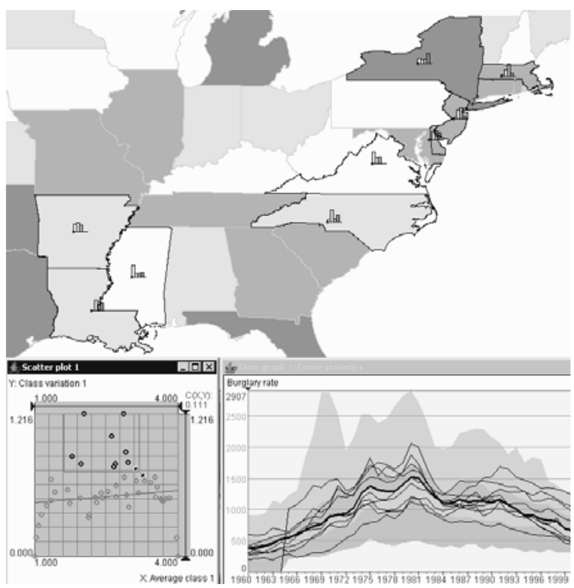

Fig. 10. Selection of the states with high variability of values.

The analytical instrument we have described allows the user to choose arbitrary percentiles for the value range division and subsequent replacement of the original values by the interval numbers. It is possible to apply the computations to the complete time period or to any sub-interval. The user can dynamically change the selected time interval. In response, the computations are automatically re-applied and the visualization of the results updated. It is also possible to use the provided counts of increases and decreases of the interval numbers inside the time series.

The proposed method can be applied to rather large data sets. It has a linear complexity in respect to time and *n*log(n)* complexity in respect to the number of locations. The visual representation is based on aggregated characteristics, and only selected individual data with specific characteristics are shown in detail.

23.4.3 Detecting Spatio-Temporal Patterns of Similar Changes

The geometrical representation of temporal behaviours provides interesting possibilities for data exploration through dynamic querying. The idea, which is demonstrated in Figure 11, is to show only line segments having a certain inclination. The inclination may be specified through setting the lower or/and upper limits for the degree of absolute or relative change (i.e. difference or ratio) in comparison to the previous moment.

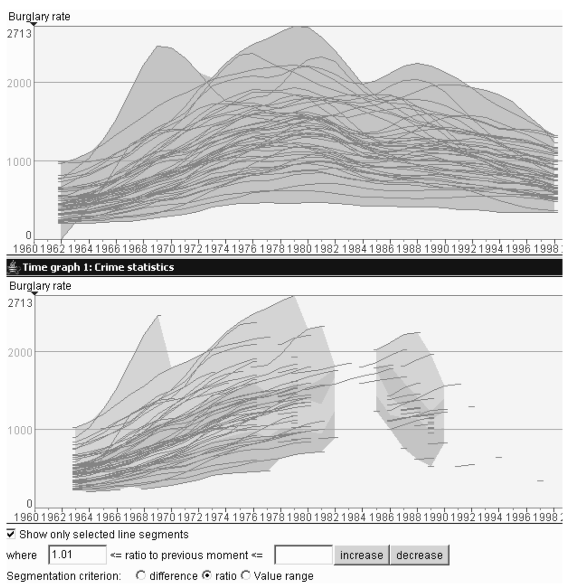

Fig. 11. The time graph shows only the line fragments corresponding to value increase by 1% or more.

Thus, in Figure 11 (lower image), the user has specified the lower limit for the relative change to be 1.01, which corresponds to 1% or more increase in a current year in comparison to the previous year. In response, the tool shows only the line fragments complying with this specification; all other line fragments have been hidden. Prior to the filtering, the lines had been smoothed using the 5-year centred moving averages and appeared as is shown in the upper image in Figure 11.

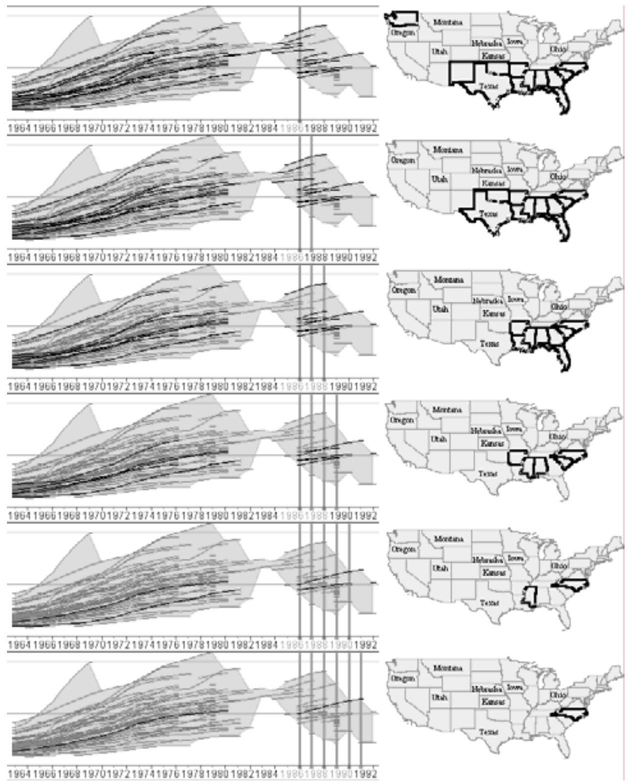

Fig. 12. Looking for states with 1% or more increase of the burglary rates in consecutive years starting from 1986.

Now, it is possible to apply direct manipulation querying in order to find out which states had no less than 1% increase of the burglary rate in particular years. This can be done by clicking on the line segments, but there is also another opportunity: the user can click on the years, that is, on the positions corresponding to different years below the horizontal axis of the graph. In the result, the lines with no less than 1% increase in these years become marked, and so do the corresponding graphical elements on other displays, in particular, the outlines of the states on a map display. Selection of two or more years marks the lines with specified inclination in all years.

Figure 12 demonstrates the effect of a series of successive selections. The vertical lines on the time graph indicate the selected years. Marked by black colour are the lines with the specified degree of increase in these years. On the right of each time graph, there is a map with the corresponding states being marked by thick black boundaries.

First, we selected the year 1986 and observed on the map what states had the specified increase of the burglary rates in this year. After that, we clicked on the next year. In the result, marking of two states disappeared, and only the states with increase in both 1986 and 1987 remained marked. Then, we clicked on the year 1988, and so on. With each subsequent selection, the number of marked states decreased. At the end, when six years from 1986 to 1991 were selected, only one state remained marked. It may be seen from the graph at the bottom that the corresponding line fragment (shown in black) ends in the year 1991, and, hence, the selection of the year 1992 would remove the last marking.

This technique is useful for answering analytical questions like: "Find spatial clusters of states with continuous increase of values during a given time period". In this method, a direct manipulation user interface supports interactive testing of the sensitivity to variation of the time period and the threshold for increase. The technique can also be applied to much larger data sets: due to the filtering, the overlapping is significantly reduced.

23.5 Discussion and Conclusions

In this paper, we proposed several interactive methods for visual exploration of spatially distributed time-series data. The first method is based on data aggregation and supports the overall view of the value dynamics on the entire territory. The second method supports finding spatial patterns of similar temporal behaviours. This method is based on data simplification: the values within intervals determined by user-selected positional statistical measures are treated as equivalent. Understandable calculations such as counting of value occurrences and finding the most frequently occurring value in a time series are then applied to the transformed data. The results are presented on dynamically linked displays that allow easy selection of objects with specific characteristics for a detailed visualisation. The third method is instrumental for detection of spatial patterns of similar temporal changes and for studying sensitivity of the procedure. This method is based on data filtering and interactive manipulation of multiple coordinated visual displays.

The essence of our approach is in combining easily understandable methods of data transformation and aggregation with interactive manipulation of linked data displays that represent the results of these methods. In the displays, the representation of aggregated characteristics of the entire dataset can be combined with viewing detailed information for selected individual instances.

An important feature of the proposed methods is their potential scalability. Although we use a small dataset in the examples presented here, almost all of the proposed methods can be applied to large data sets as well (see a related discussion in Andrienko and Andrienko 2005a). We have an experience of applying these methods to datasets of different sizes. In all cases, the results were useful and stimulating.

Besides presenting the tools, one of our objectives has been to demonstrate a task-analytical approach to choosing techniques for exploratory data analysis. There is a general problem: how to determine what combination of techniques would be necessary and sufficient for the exploration of a particular dataset or a class of datasets with similar structures and properties. Task analysis is a way to solve this problem: before starting to choose and link tools, it is necessary to identify the range of the exploratory tasks that need to be supported. The task-analytical approach is also relevant to the design of various multimedia presentations (in the more conventional sense of this term).

In the book (Andrienko and Andrienko 2005b), we present theoretical and methodological foundations for defining the tasks pertinent to the exploration of a dataset or a class of datasets, choosing appropriate tools and methods to combine the tools, and performing the exploratory analysis of the data in a systematic, comprehensive way.

As to the directions for further research, it is clear that substantial progress is required in tools and methods to support exploratory analysis of spatio-temporal data, in particular, analysing multiple attributes simultaneously, dealing with very long time series, and detecting structures in temporal behaviours (e.g. cyclical phenomena). However, there is another research problem, extremely serious and urgent, which may be called 'Usability Problem' (Andrienko and Andrienko 2006). The essence of the problem is that potential users are often incapable or unwilling to use the tools and techniques offered to them.

The Usability Problem involves many aspects. One of them is the users being unfamiliar with the innovative ideas and approaches such as interactive maps or multimedia maps. Another aspect is the imperfection of the existing tools, mostly research prototypes. This refers not only to the design of the user interface but also to the functionality. On the one hand, the tools are quite complex for the user (e.g. because of involving multiple displays or novel visualization techniques); on the other hand, they are still insufficient since they do not cover the full range of potential user's tasks. Besides, the users are not given any facilities to test the hypotheses they might have generated by using the exploratory tools as well as tools to register their findings and to report the results of exploration.

Since the Usability Problem is very complex and, at the same time, very critical, coordinated efforts of researchers in different disciplines are needed to solve it. We are going to contribute to this by finding approaches to 'embed' intelligence into software tools for data exploration so that they could adapt automatically to data under analysis and user's goals and assist the user in choosing and applying appropriate techniques.

References

Andrienko N, Andrienko G, Gatalsky P (2003) Exploratory spatio-temporal visualization: an analytical review. Journal of Visual Languages and Computing 14(6):503−541

Andrienko G, Andrienko N (2005a) Blending aggregation and selection: adapting parallel coordinates for the visualisation of large datasets. The Cartographic Journal 42(1):49−60

Andrienko N, Andrienko G (2005b) Exploratory analysis of spatial and temporal data: a systematic approach. Springer, Berlin (in print)

Andrienko G, Andrienko N (2006) Visual data exploration: tools, principles, problems. In: Fisher P (ed.) Classics of IJGIS, Taylor & Francis (in print)

Bertin J (1983) Semiology of graphics: diagrams, networks, maps. The Univ. of Wisconsin Press, Madison

Blok C (2000) Monitoring change: characteristics of dynamic geo-spatial phenomena for visual exploration. In: Freksa Ch et al. (eds) Spatial Cognition II, LNAI 1849, Springer, Berlin pp.16−30

Cartwrite W, Peterson M, Gartner G (1999) Multimedia cartography. Springer, Berlin

Dykes J, MacEachren AM, Kraak MJ (2005) Exploring geovisualization, Elsevier, Oxford

Harrower M, Griffin AL, MacEachren AM (1999) Temporal focusing and temporal brushing: assessing their impact in geographic visualization. In: Proceedings of 19th International Cartographic Conference, Ottawa, Canada, Vol. 1, pp. 729–738

Hedley NR, Drew CH, Arfin EA, Lee A. (1999) Hägerstrand Revisited: Interactive Space-Time Visualizations of Complex Spatial Data. Informatica: An International Journal of Computing and Informatics, 23 (2), pp. 155-168.

Hochheiser H, Shneiderman B (2004) Dynamic query tools for time series data sets: timebox widgets for interactive exploration. Information Visualization 3(1):1−18

Keogh E, Kasetty S (2003) On the need for time series data mining benchmarks: a survey and empirical demonstration. Data Mining and Knowledge Discovery 7(4):349−371

Kraak MJ (2003) The space-time-cube revisited from a geovisualization perspective. In Proceedings of the 21st International Cartographic Conference (ICC) "Cartographic Renaissance", Durban, South Africa, August 2003, pp. 1988-1996.

Koussoulakou A, Kraak MJ (1992) Spatio-temporal maps and cartographic communication. The Cartographic Journal 29:101–108

MacEachren AM, Kraak MJ (1997) Exploratory cartographic visualization: advancing the agenda. Computers and Geosciences 23:335−344

Peterson MP (1999) Elements of Multimedia Cartography. In: Cartwright W, Peterson MP, Gartner G (eds.) Multimedia cartography. Springer, Berlin, pp. 31-40

Peuquet DJ (1994) It's about time: a conceptual framework for the representation of temporal dynamics in geographic information systems. Annals of the Association of American Geographers 84(3): 441−461

Slocum T, Yoder S, Kessler F, Sluter R (2000) MapTime: software for exploring spatiotemporal data associated with point locations. Cartographica 37(1):15–31

Tobler W (1970) A computer movie simulating urban growth in the Detroit region. Economic Geography 46(2):234−240

Unwin AR, Wills G (1988) Eyeballing time series. In: Proceedings of the 1988 ASA Statistical Computing Section, pp 263−268

Vasiliev IR (1997) Mapping time. Cartographica 34(2):1–51