Visual Exploration of the Spatial Distribution of Temporal Behaviors

Gennady Andrienko and Natalia Andrienko
Fraunhofer Institute for Autonomous Intelligent Systems
Schloss Birlinghoven, D-53754 Sankt Augustin, Germany
http://www.ais.fraunhofer.de/and
{gennady,natalia}.andrienko@ais.fraunhofer.de

Abstract

The paper elaborates on the previous research on the analysis of temporal and spatio-temporal data done in statistical graphics and geo-visualization. We focus on the exploration of spatially distributed time-series data, i.e. values of numeric attributes referring to different moments in time and locations in space. After considering appropriate interactive visualization techniques, we propose several methods of exploration based on user-controlled data transformation and aggregation, easy-to-understand calculations, and dynamic linking of data displays. The proposed methods are potentially scalable, i.e. they can be applied to large data sets without overcrowding the displays and loosing interactivity.

Keywords: **visual data mining, spatial time-series analysis, data visualization, dynamic and coordinated data displays**

1. Introduction

Analysis of time series data is one of the most important topics in several research areas. Thus, a number of computational methods have been recently developed in data mining [1]. Researchers in statistical graphics, information visualization, and human-computer interaction propose various interactive visualization techniques for time-series. The papers [2,3] introduce basic graphical and interaction facilities for enhancing analytical capabilities of time series plots, or time graphs:

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

Most of these tools are applicable to a relatively small number of time series but unsuitable for studying large collections of time series data. In most cases, such tools are implemented as stand-alone prototypes with a limited number of available complementary displays. The authors do not propose any approaches to dealing with the spatial aspect of the data except for linking time graphs with simple maps, which show only the locations of the time series and provide quite restricted visualization and interaction opportunities.

Cartographic and geo-visualization research communities pay much attention to visualization and analysis of spatial time series data [4-11]. Since [12], different types of spatio-temporal data are distinguished according to the type of changes that occur over time:

1. Existential changes: appearing, disappearing, reviving of objects or/and relationships;
2. Changes of spatial properties of objects (location, size, shape);
3. Changes of thematic properties, i.e. values of attributes.

Sometimes only one type of change takes place or is of interest for an analyst, but in many cases one needs to deal with several types simultaneously.

In our previous publications, we suggested several new methods supporting the exploration of various types of spatio-temporal data and various aspects of them. Thus, we designed and implemented techniques for the exploration of trajectories of moving objects [13] and spatio-temporal distribution of point events [14], detection of changes and analysis of variance in spatial time series data [15,16], characterization and comparison of spatial development scenarios [17]. We have recently undertaken a survey [18] of existing methods and software tools for the exploratory spatio-temporal visualization in order to build an inventory of approaches and to identify the necessary directions for further work. In paper [19] we have studied the impact of data and task characteristics on design of spatio-temporal data visualization tools.

In this paper, we focus on the exploration of changes of thematic properties of spatial objects. Our objective is to find suitable approaches and suggest tools for analyzing multiple spatially distributed time series.

Proceedings of the Ninth International Conference on Information Visualisation (IV'05)
1550-6037/05 $20.00 © 2005 **IEEE**

2. Local behaviors and their visualization

We shall use the term behavior, or local behavior, to denote the temporal variation of attribute values in a particular place. One of the possible views of a dataset consisting of attribute values referring to a set of spatial locations and a sequence of time moments is to view it as a collection of local behaviors, with each behavior referring to one of the locations. Other views are also possible, for example, treating the data as a sequence of spatial distributions of the attribute values, where each distribution refers to one of the time moments. In this paper, we take the former view. We limit our focus to behaviors formed by changeable values of a single numeric attribute.

In the course of the exploration of a collection of spatially distributed behaviors, 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 behaviors in this area and how do they compare to the behaviors on the remaining territory?
- Find locations with the behaviors having specific features and check whether these locations are neighbors, or, in other words, form a spatial cluster. Here are some examples of the features that may be looked for:
  - persistently low (or high) values;
  - high value fluctuations;
  - continuous increase (or decrease) of values during a given time period.
- Identify spatial clusters of similar behaviors.

Let us look what tools and techniques may be helpful to the data analyst in finding the answers. For our investigation, we shall use the dataset concerning the statistics of crimes in the USA published by the U.S. Department of Justice. The data are 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 behavior, the visualization of the data on a time graph is typically used. The behavior 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 behaviors 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 behaviors 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.

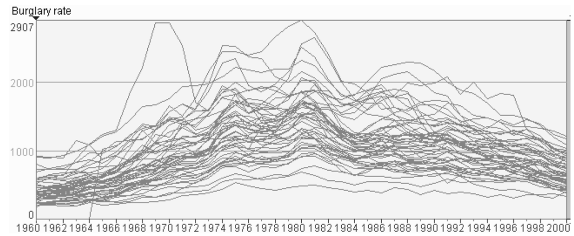

Figure 1. Behaviors of the burglary rates in the states of the USA are represented on a time graph display as multiple overlaid lines.

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 behaviors in this or that area. However, since the user can only see the behaviors corresponding to selected locations, it is very difficult to understand how the character and distinctive features of the local behaviors vary over the whole territory.

Figure 2. A cartographic representation of the spatial distribution of the behaviors of burglary rates over the USA

A suitable graphical representation for the latter task is shown in Figure 2: the local behaviors 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.

With such a map, we can see how different behaviors are distributed over the territory of USA. Thus, we can observe that the states in the north-central part of the country had smaller 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 behaviors of the burglary rate: (1) west-southwest-south, (2) middle north, (3) the area around the Great Lakes, (4) center and southeast, except for Florida. The clusters are roughly outlined in Figure 3.

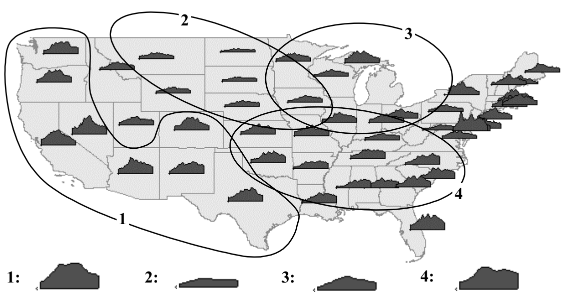

Figure 3. Spatial clusters of states with similar temporal behaviors of the burglary rate. Below the map, the typical behavioral patterns for each cluster are schematically shown.

By visual inspection of the map, an explorer can find locations with the corresponding behaviors having specific features and, naturally, immediately see whether these locations form a spatial cluster. It is also possible to detect the essential common features of behaviors in a certain area.

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, quite short time series, and a small number of spatial locations. Therefore, we have been able to review and compare all the local behaviors 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 out of the scope of this paper. 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 stick to the same example data, the methods are also applicable to substantially larger data volumes.

3. Combining tools for behavior 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, there are several problems with this solution. 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.

3.1 Getting the general picture of the behavior on the entire territory

In Figure 1, we have demonstrated a time graph display with multiple overlaid lines. Already with fifty local behaviors, 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 the time graph display representing multiple behaviors.

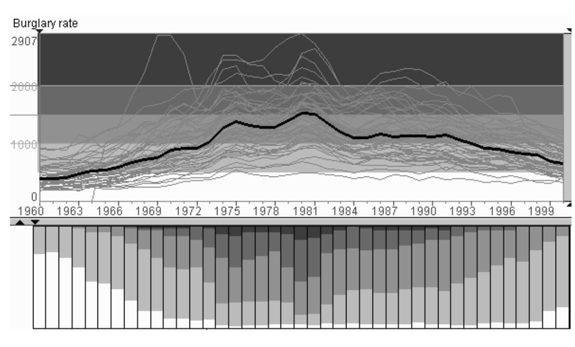

Figure 4. The value range of the attribute has been divided into intervals. The segmented bars show for each year the proportions of values over the whole country fitting in these intervals.

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

Proceedings of the Ninth International Conference on Information Visualisation (IV'05)
1550-6037/05 $20.00 © 2005 **IEEE**
IEEE COMPUTER SOCIETY

2000. For each year, the display contains a segmented bar with the differently colored 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 gray – to values between 500 and 1000, and so on (the higher the values, the darker the color). The background coloring 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 behavior of the burglary rates on the whole territory over the time. Thus, we can see that in early sixties 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 mid-sixties, 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 slightly improved and remained more or less stable until the beginning of nineties, when it started to gradually improve and preserved this trend 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.

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 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 represents the transformed values instead of the original ones.

Like in the previous case, we have applied the division of the whole range of the transformed values into intervals. This time, chose 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 for finding years and periods of significant countrywide increase or decrease of values. The light gray color 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.

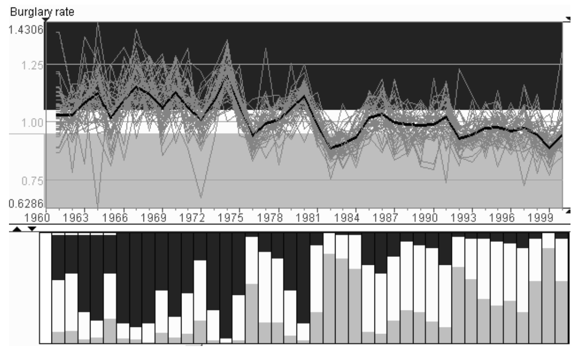

Figure 5. The time graph has been transformed to show the annual changes in the individual states. The lower image aggregates the changes for each year.

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: 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.

3.2 Finding spatial patterns of similar local behaviors

Another way of aggregating the same data is demonstrated in Figure 6.

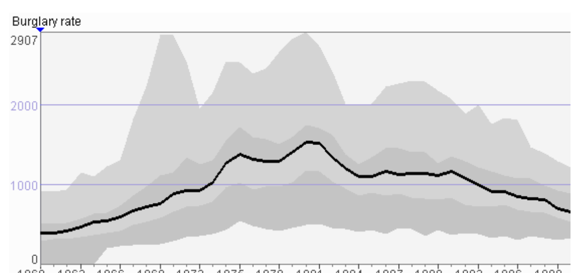

Figure 6. The aggregate time graph represents the changes of the value range and the quartiles over time.

The display contains a so-called "envelope", that is, a polygon enclosing all the lines representing the local behaviors. Actually, the lines themselves are not needed for building the envelope but only the minimum and maximum values over the country in each year. The envelope thus shows how the range of the 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 painted using alternating light and dark shades of gray, 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 from 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.

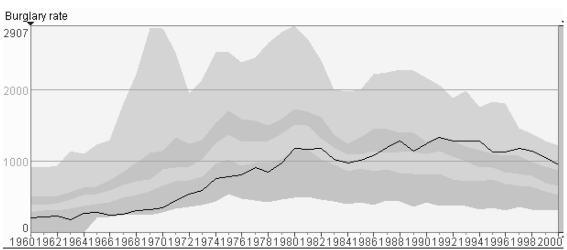

Figure 7. The line representing the local behavior in Mississippi is superimposed on the aggregate time graph.

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 sub-intervals. 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 particularize this observation by giving some numeric estimations: 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.

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 being proportional to the respective counts. Additionally, the coloring 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.

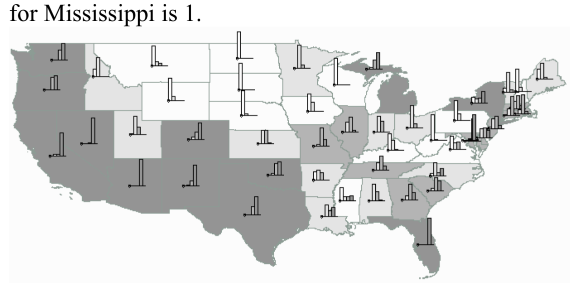

Figure 8. Computationally supported detection of clusters of similar local behaviors

It may be noted that, in the result of this automated procedure, we have got spatial clusters (groups of neighboring identically colored states) very similar to the clusters in Figure 3 revealed "manually".

The transformed data can be also analyzed in another way. For each local behavior, 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 behaviors with particular characteristics in terms of the average and variance. Thus, in Figure 9, the scatterplot has been used to find the behaviors 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.

IEEE COMPUTER SOCIETY

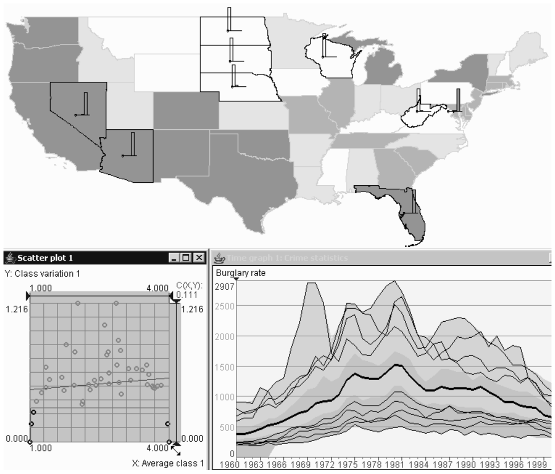

Figure 9. Selection of the states with consistently low (painted in light color) and consistently high (shown in dark) burglary rates

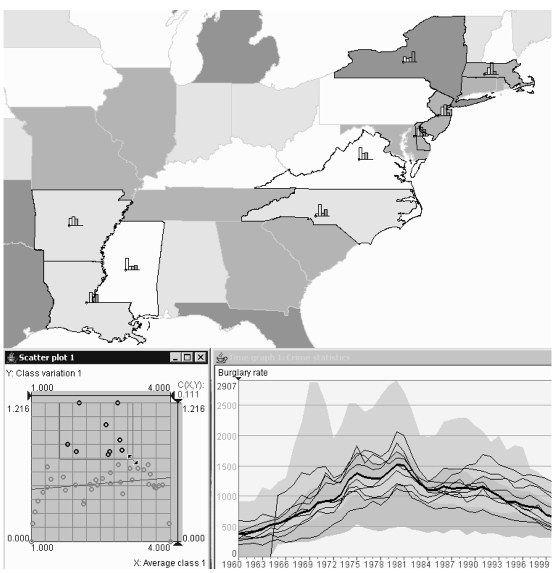

Figure 10. Selection of the states with high variability of values

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 behaviors and a detailed representation of selected behaviors. This technique compensates for the deficiencies of both representations being used separately. Despite of 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.

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 apply another data simplification procedure that counts 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.

3.3 Detecting spatio-temporal patterns of similar changes

The geometrical representation of temporal behaviors 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.

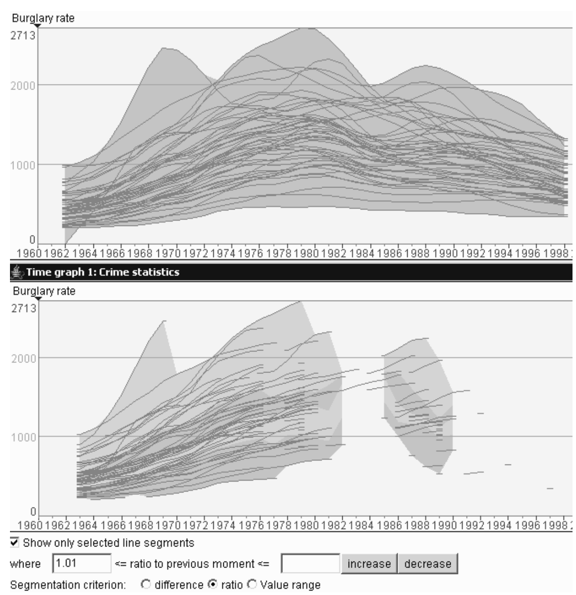

Figure 11. The time graph shows only the line fragments corresponding to value increase by 1% or more.

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. 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 centered moving averages and appeared as is shown in the upper image in Figure 11.

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 having the specified inclination in all these years.

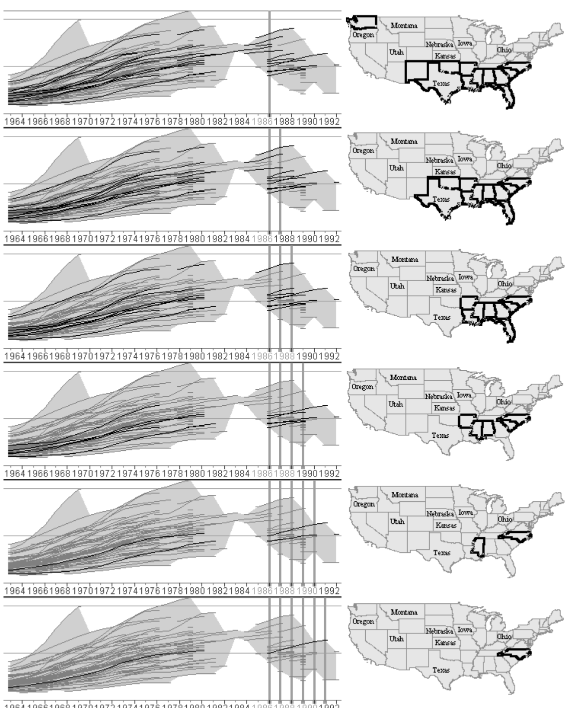

Figure 12. Looking for states with 1% or more increase of the burglary rates in consecutive years starting from 1986.

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 colors 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.

4. Discussion and conclusions

In this paper, we proposed several interactive methods for visual exploration of spatially distributed time-series data. The first method described in § 3.1 is based on data aggregation and supports the overall view of the value dynamics on the entire territory. The second method (§ 3.2) supports finding spatial patterns of similar temporal behaviors. This method is based on data simplification: the values within intervals determined by user-selected positional statistical measures are treated as equivalent. Easily interpretable 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 visualization. The third method (§ 3.3) 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 representing 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 [20]). We have an experience of applying these methods to datasets of different sizes. In all cases, the results were useful and stimulating.

There are several directions for further research. New interactive methods are needed for analyzing multiple attributes simultaneously, for dealing with very long time series, and for detecting structures in temporal behaviors (e.g. cyclical phenomena etc.) Besides developing new methods, there is a need in helping users to select proper tools and effectively operate them for solving specific problems and making informed well-grounded decisions. For this purpose, we are going to develop a knowledge-based system for assisting users in data analysis that will incorporate the generic knowledge on data analysis, visual communication, and decision making.

5. Acknowledgements

This work was partly supported in projects SPIN!, GIMMI, EuroFigures, EFIS, NEFIS, SILVICS. We are grateful to our partners and colleagues for friendly support and fruitful discussions. We would especially like to thank Peter Gatalsky for many years of cooperation in spatio-temporal research.

References

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

[2] Unwin, A. R., and Wills, G., Eyeballing Time Series. Proceedings of the 1988 ASA Statistical Computing Section, 1988, pp.263-268

[3] H.Hochheiser and B.Shneiderman, Dynamic query tools for time series data sets: Timebox widgets for interactive exploration, Information Visualization, 2004, 3 (1), pp.1-18.

[4] J. Bertin, "Semiology of Graphics Diagrams Networks Maps", The Univ. of Wisconsin Press, Madison, 1983.

[5] A. Koussoulakou, M.J. Kraak, "Spatio-temporal maps and cartographic communication", The Cartographic Journal, 29 (1992) 101–108.

[6] I.R. Vasiliev, Mapping time, Cartographica 34 (2) (1997) 1–51.

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

[8] Tobler, W., (1970) "A computer movie simulating urban growth in the Detroit region", Economic Geography, 46(2), 234-240.

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

[10] M. Harrower, A.L. Griffin, A.M. MacEachren, Temporal focusing and temporal brushing: assessing their impact in geographic visualization, in: Proceedings of 19th International Cartographic Conference, Ottawa, Canada, Vol. 1, 1999, pp. 729–738.

[11] A. Fredrikson, C. North, C. Plaisant, B. Shneiderman, Temporal, geographic and categorical aggregations viewed through coordinated displays: a case study with highway incident data, in: Proceedings of Workshop on New Paradigms in Information Visualization and Manipulation (NPIVM '99), ACM Press, New York, 1999, pp. 26–34.

[12] C. Blok, "Monitoring Change: Characteristics of Dynamic Geo-spatial Phenomena for Visual Exploration", in Ch.Freksa et al. (eds) Spatial Cognition II, LNAI 1849, Berlin, Springer-Verlag pp.16-30

[13] Andrienko, N., Andrienko, G., and Gatalsky, P. Supporting Visual Exploration of Object Movement. In V. Di Gesù, S. Levialdi, L. Tarantino (eds.) Proceedings of the Working Conference on Advanced Visual Interfaces AVI 2000, Palermo, Italy, May 23-26, 2000, ACM Press, 2000, pp.217-220, 315

[14] P. Gatalsky, N. Andrienko, and G. Andrienko, Interactive Analysis of Event Data using Space-Time Cube, In Banissi, E. et al (Eds.) IV 2004. 8th International Conference on Information Visualization, Proceedings, 14-16 July, 2004, London, UK, IEEE Computer Society, Los Alamitos, California, 2004, pp. 145-152

[15] Andrienko, N., Andrienko, G., and Gatalsky, P., Exploring Changes in Census Time Series with Interactive Dynamic Maps and Graphics, Computational Statistics, 2001, v.16 (3), pp.417-433

[16] Andrienko, N., and Andrienko, G., Interactive visual tools to explore spatio-temporal variation, In M.F.Coastabile (Ed.) Proceedings of the Working Conference on Advanced Visual Interfaces AVI 2004, Gallipoli, Italy, May 25-28, 2004, ACM Press, 2004, pp.417-420

[17] N. Andrienko, G. Andrienko, and P. Gatalsky, Tools for Visual Comparison of Spatial Development Scenarios, In Banissi, E. et al (Eds.) IV 2003. Seventh International Conference on Information Visualization, Proceedings, 16-18 July, 2003, London, UK, IEEE Computer Society, Los Alamitos, California, 2003, pp. 237-244

[18] Andrienko, N., Andrienko, G., and Gatalsky, P., "Exploratory Spatio-Temporal Visualization: an Analytical Review", Journal of Visual Languages and Computing, 2003, v.14 (6), pp. 503-541

[19] Andrienko, N., Andrienko, G., and Gatalsky, P., Impact of data and task characteristics on design of spatio-temporal data visualization tools, In J. Dykes, A. MacEachren and M.-J. Kraak (Eds.) Exploring Geovisualization, Oxford: Elsevier, 2004 (in press)

[20] G. Andrienko and N. Andrienko, "Blending Aggregation and Selection: Adapting Parallel Coordinates for the Visualisation of Large Datasets", The Cartographic Journal, 2005 (accepted)

IEEE COMPUTER SOCIETY