Exploring Changes in Census Time Series with Interactive Dynamic Maps and Graphics

Natalia Andrienko, Gennady Andrienko, and Peter Gatalsky

GMD – German National Research Center for Information Technology, Institute for Autonomous Intelligent Systems, Knowledge Discovery Research Group, Schloss Birlinghoven, Sankt Augustin, D-53754, Germany

Summary

Interactive graphical tools are available in a number of statistical packages. However, none of them provide interactive tools for visual exploration of time series data that have spatial reference (spatio-temporal data), such as census data. Analysis of spatio-temporal data requires simultaneous representation of their spatial, temporal, and thematic aspects. We propose a system that combines interactive thematic maps, dynamic statistical graphics, and advanced controls for manipulating them. These components as well as interaction between them are specifically designed in order to support visual investigation of changes in data, the ultimate goal being to help users to reveal spatio-temporal trends. The suggested tools facilitate also detecting errors and anomalies in data.

Keywords: Spatio-temporal data, Interactive dynamic graphics, Linked displays, Exploratory data analysis

1 Motivation

Unwin and Wills (1998) demonstrated that interactive graphics could greatly facilitate analysis of time-series data. They listed several analytical operations supported by software for time-series analysis called Diamond Fast:

- scaling of graphics in attribute and time dimensions (vertical and horizontal components of a time plot);
- overlaying time series;
- simultaneous consideration of multiple time series with common scaling and dynamic linking.

Census data (aggregated) provide an example of spatially referenced time series data, or spatio-temporal data. They usually consist of values of multiple attributes referring both to units of territory division and to time moments. Exploratory analysis of spatio-temporal data requires tools for visualisation and manipulation of all 3 aspects of the data: thematic (values of attributes), temporal and spatial. For dealing with the spatial aspect involvement of maps is essential.

Traditional cartography has been always concerned about possible methods of representation of spatio-temporal data on maps. However, this proved to be a difficult task, and only a few methods were invented (Vasiliev 1997). One of them is map iteration, i.e. juxtaposition of several maps where each map represents data referring to a different time moment. Another method is representation of changes, i.e. calculated absolute or relative differences between data at two time moments, rather than the source data. Both these methods have evident restrictions regarding the number of time moments that can be represented and analysed. Besides these generic techniques, there are also methods designed for specific types of data. For example, arrows can be applied to represent movement of entities in space.

Recently cartographers have paid much attention to the opportunities offered by modern computers. These opportunities are:

- interaction with the data display for getting answers to various queries;
- dynamic linking, or “brushing”, between parallel views representing different aspects or projections of the same data set (Buja et al. 1991);
- fast transformations of data and graphics;
- animation when a display dynamically changes to reflect changes of data.

Researchers in cartographic visualisation strive to combine these new facilities with the established cartographic techniques. A number of techniques and tools for visualisation of spatio-temporal data have been suggested recently: temporal focusing and temporal “brushing” (Monmonier 1990), interaction with various aggregations of data (Frederikson, Plaisant and Shneiderman 1999), map animation (Tobler 1970; Slocum et al. in press), “active legends” for controlling animation (Kraak, Edsall and MacEachren 1997; Edsall and Peuquet 1997, Buziek 1999), synchronisation of several animations (Blok et al. 1999), animated display of differences between successive moments (Oberholzer and Hurni, 2000).

One can observe that the map animation technique has got particular attention of researchers and tool developers. It is true that this technique is visually very appealing. However, in most cases the animation alone cannot be a sufficient tool for data analysis, even when the user is provided a high level of control over its parameters. As we found in our research, it may work well in investigation of phenomena that allow representation of states at multiple successive moments in a single map, for example, spatial movement of objects (Andrienko, Andrienko and Gatalsky 2000).

In other cases the animation does not provide to the analyst an opportunity to compare directly states of the phenomenon under study at different time moments. In order to detect and evaluate changes, she/he has to compare the state viewed at the current moment with the mental image of the previous state. Investigation of temporal trends would require memorising of a large number of consecutive states. This is practically unworkable unless the trends are as apparent as urban growth (see, for example, Tobler 1970 or Slocum et al. in press).

Hence, to design appropriate tools for exploratory analysis of spatio-temporal data, one needs to employ the other facilities enabled by computers: multiple linked views, display transformations, and interactivity. All this can, of course, be combined with map animation.

The goal of this paper is to demonstrate how a transformable time-series plot dynamically linked to a map can help in discovering interesting, significant features in data. This is done on an example of time-referred aggregated census data. The data set we investigated contains annual records about population of 45 municipalities of the Overijssel region in the Netherlands from 1988 to 1998. It was kindly provided by Prof. M.-J.Kraak from ITC, Enschede. We should make a reservation that we did not pursue the goal of comprehensive analysis of population changes in Overijssel but rather tested the tools we had developed on data that were previously unknown to us.

2 Census data: possible analysis tasks

In analysis of census data one considers values of different attributes referring to territory units and variation of these values over time. The questions an analyst may seek to answer (with regard to a single attribute) may be summarised in the following table:

Space → **Time ↓**	**Elementary level** (individual locations)	**Regional level** (the whole territory under study and its parts)
Single moment *t*	What value had the attribute *a* at *t* at a given location?	How were the values of the attribute distributed at *t* over the studied region?
Two moments *t1* and *t2* (change)	What is the difference between values of the attribute at *t1* and *t2* at a given location?	**a)** What is the difference between the patterns of spatial distribution of values at *t1* and *t2*? **b)** How are the changes between *t1* and *t2* distributed over the region? E. g. where did the maximum changes occur?
Interval [*t1*,*t2*] (trend)	What happened during [*t1*,*t2*] at a given location? For example, did the values monotonously grow or decrease? Are the changes periodic? When did the maximum change occur?	**a)** How did the spatial pattern of value distribution evolve over time? For example, did areas of concentration of high (low) values keep in their place or move with the time over the territory? Did the patterns repeat periodically? When did the maximum changes of distribution pattern occur? **b)** Is there a common temporal trend for the whole region or for parts of the region, e.g. monotonous increase or decrease? Did changes in adjacent areas occur synchronously or with a time lag? What locations had an outstanding behaviour?

This table can be extended to the case of simultaneous consideration of several attributes. However, the tools we developed so far support investigation of a single numeric attribute at a time

3 Instruments for exploration of changes

As an instrument for exploratory analysis of time-series data referring to units of territory division we suggest a map dynamically linked to a time-series plot. In such a plot the horizontal axis (X) corresponds to the temporal dimension of data while the vertical axis (Y) represents the range of values of the numeric attribute under investigation. The plot can show evolution of attribute values for one or several objects. Each data item is mapped onto a point on the plot with the Y-position corresponding to the value of the attribute and the X-position corresponding to the time moment this value refers to. Pairs of points referring to the same object and successive time moments are connected with lines. As a result, for each object there is a corresponding trajectory through the graph (see Figure 1a).

Linking between the time-series plot and the map works in the following way. Pointing with the mouse on any object of the map highlights the corresponding trajectory in the plot, and putting the mouse over a line of the plot highlights the corresponding object in the map. Another kind of interaction between the time plot and the map occurs when the user clicks on some position on the time axis of the plot. This leads to selection of the corresponding time moment for current viewing, and the map is immediately updated to display data referring to this moment. The system enables user's access to exact attribute values both through the map and the plot.

For displaying attribute values (as well as various derived figures) on the map we use the cartographic representation method known as "choropleth map": territory segments are painted with varying degrees of darkness depending on the values of the attributes. Typically higher values are shown by darker shades. Often a double-ended, or diverging colour scale is applied in choropleth maps (Brewer 1994). Such a scale effectively represents deviations from a certain reference value: shades of one colour are used for values above the reference value (positive deviations) and shades of another colour for negative deviations. The magnitudes of the deviations are shown by degrees of darkness.

In our selection of the "choropleth map" representation method we were aware of its disadvantages. In particular, sizes of territory segments usually differ, sometimes very significantly, and larger segments tend to dominate visually smaller ones, especially when being painted in darker shades. The handbooks on cartography warn against the use of choropleth maps for representation of absolute numbers that can be related to sizes of areas, such as population number. Bertin (1967/1983) even considers this method inappropriate for numeric data at all and recommends instead to use symbols varying in size.

On the other hand, the method has an undeniable advantage of supporting perception of the whole map as a single image. This is especially beneficial during the animation: instead of trying to trace changes of numerous symbols scattered over the map the analyst can observe changes of the image as a whole. Therefore we think that in combination with other analysis instruments choropleth maps can be used. However, all findings originating from observation of such a map should be carefully checked, and no conclusions should be made on the basis of the map alone.

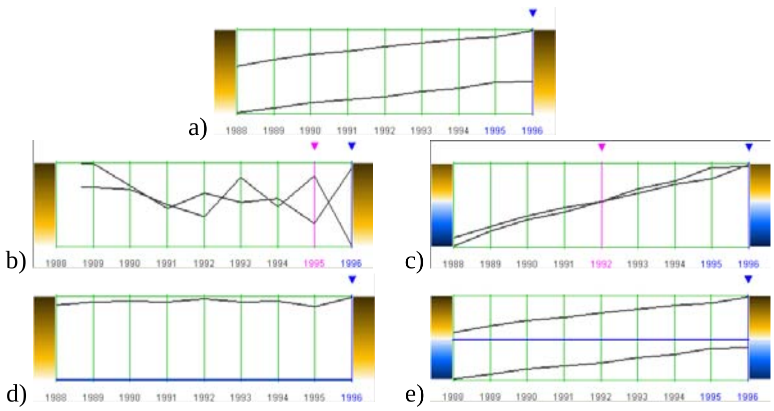

	1988	1989	1990	1991	1992	1993	1994	1995	1996
a) Source data									
A	4696	4794	4889	4961	5014	5127	5195	5311	5320
B	5639	5773	5874	5940	6029	6104	6185	6228	6355
b) Comparison to the previous year									
A		+98	+95	+72	+53	+113	+68	+116	+9
B		+134	+101	+66	+89	+75	+81	+43	+127
c) Comparison to the year 1992									
A	-318	-220	-125	-53	0	+113	+181	+297	+308
B	-390	-256	-155	-89	0	+75	+156	+199	+326
d) Comparison to a reference object (A)									
A	0	0	0	0	0	0	0	0	0
B	+943	+979	+985	+979	+1015	+977	+990	+917	+1035
e) Comparison to a reference value (5500)									
A	-804	-706	-611	-539	-486	-373	-305	-189	-180
B	+139	+273	+374	+440	+529	+604	+685	+728	+855

Figure 1. Time plot representing data for 2 objects (a) and its transformations enabling comparison with previous year (b), with year 1992 (c), with one of the objects (d), and with a reference value (e). Absolute changes (i.e. differences) are shown.

In our system we combine the choropleth map with the time-series plot. Both displays can show not only the original attribute values but also various derived figures representing magnitudes of change. The user may choose to study absolute or relative changes between two time moments. Accordingly, the system calculates differences or ratios of values at these moments.

The user may also choose to what the data at each moment are compared:

- to data at the previous moment of time (i.e. values in 1989 to those in 1988, 1990 to 1989, and so on);

- to data at a selected moment of time (e.g. data in each year are compared to values in 1990);
- to values associated with a selected spatial object (i.e. in each year values for all objects are compared to the value for the selected object in the same year);
- to a specified constant number (the same for each year).

In order to represent calculated values characterising changes, we apply in the map a diverging colour scale. In this scale white colour depicts absence of changes (i.e. when the difference is equal to 0 or the ratio is equal to 1). Shades of brown are used to represent increase and shades of blue for decrease. The degree of darkness represents the magnitude of change. The advantage of such a representation is that neighbouring districts painted in the same colour tend, despite differences in darkness, to associate visually. The observer can easily find areas of coherent increase or decrease of values (if they exist).

The time-series plot is also transformed depending on the currently selected variant of comparison. Figure 1 (b-e) illustrates the four variants of comparison on a simplified example with only two objects. It may be seen what figures are represented on the map and on the plot in each case, and how the plot is transformed in relation to the plot of original values.

The user can interactively change selection of the current moment (i.e. the moment that is compared). In response to this the differences or ratios are immediately recomputed, and the displays change accordingly. It is also possible to animate the display of changes. During the animation the time plot indicates the currently represented time moment and the reference moment for comparison (when it differs from the current moment).

One can find similarity between appearance and behaviour of the interactive time-serious plot (in particular, highlighting of lines when the mouse points on them) and the interactive plot of parallel co-ordinates (Inselberg 1985). In fact, some people were mislead by this similarity and took the former for the latter. Therefore we find it necessary to explain the difference between these two types of graphs. The parallel co-ordinates plot represents values of different attributes, usually with no reference to time. It has several parallel axes corresponding to the attributes. The time plot shows values of the same attribute at different moments of time. The vertical lines present in our time-series plots are not axes but a grid simplifying finding positions corresponding to specific time moments.

However, some of the interactive techniques proposed for manipulation of the parallel co-ordinates plot (e.g. Avidan and Avidan 1999) are applicable also to the time-series plot and may be useful in analysis of spatio-temporal data:

- querying data by imposing constraints on the value range of some attribute(s) or on the angle of lines between a pair of attributes;

- removing selected lines from the plot;
- scaling the plot horizontally and vertically (in our case – in the temporal and attribute spaces);
- classifying the objects by breaking the value range of some attribute into subintervals, then assigning a distinct colour to each subinterval, and brushing the plot (and other linked displays) using these colours.

Classification in general is a very useful instrument of analysis of spatially referenced data (Yamahira, Kasahara and Tsurutani 1985). The idea is that the analyst uses some method to divide the whole set of geographical objects into classes. Then a distinct colour is assigned to each class, the objects on the map are painted in the corresponding colours, and the analyst observes the spatial distribution of the classes. Our system enables multiple methods of classification. Here are just a few examples:

- small, average, and big values of the attribute in a specific year;
- decrease or increase of values between two selected years;
- faster or slower growth or decline than in a selected reference district.

A link between our system and the system for data mining Kepler (Wrobel et al., 1996) allows the analyst to use the classes thus generated as an input for some data mining algorithms, for example, the C4.5 classification learning method (Quinlan, 1993). The data mining methods may help to reveal interesting relationships between attributes. The link between the systems is described in (Andrienko and Andrienko 1999a). Andrienko and Andrienko (1999b) describe various techniques for interactive classification and application of the C4.5 method to the classed produced.

The interactive graphical tools for spatio-temporal data analysis are implemented in Java and are to be included in the Descartes system. An interested reader can run them in the Internet at http://borneo.gmd.de/and/time/itc.html.

The rest of the paper demonstrates the use of our instruments for visual exploration of spatio-temporal data by an example data analysis session.

4 A scenario of data analysis

In order to test our analytical instruments, we undertook an exploration of a previously unknown data set with population statistics in a certain region. The data set contained values of several attributes for 11 years referring to different municipalities. To represent the data on a map, we also got geometric data specifying the actual boundaries of these municipalities. We didn’t get data about the boundaries in the previous years and had to assume that they were stable. As we demonstrate later, this assumption proved to be wrong. However, it is

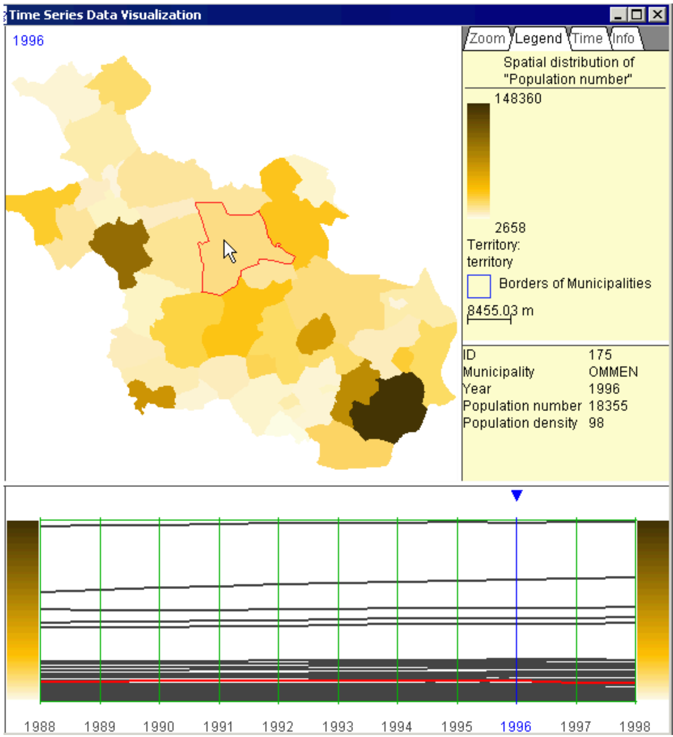

Figure 2. Initial view on data: population number in 1996 is shown on the map.

interesting that we were able to detect a change of a boundary only on the basis of analysis of the attribute data.

We started with visualisation of the absolute population numbers and tried to view their development using an animated map presentation (a momentary screenshot is shown in Figure 2). The animation did not demonstrate any noticeable changes as well as the time-series plot of the original data that contained only parallel lines. We switched to the comparison mode and looked at presentation of absolute changes in comparison to the previous time moment (Figure 3).

The time plot immediately attracted our attention by demonstrating a very strange fluctuation of population number in some municipality during 3 successive years: a sharp decrease of population in 1995 followed by an even more drastic increase in 1996. The fluctuation was even more vividly seen in the mode of comparison with a fixed year, for example, 1994 (Figure 3a).

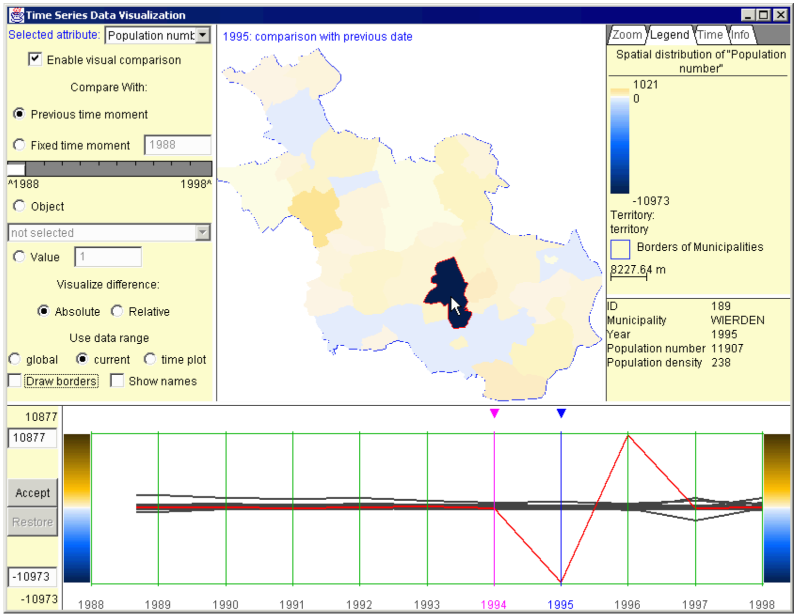

Figure 3. Presentation of changes in comparison to the previous year exposes an error in data.

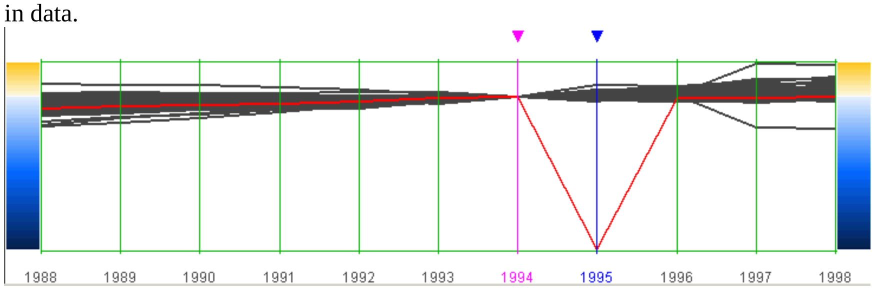

Figure 3a. The same fluctuation in data represented in the mode of comparison with a fixed year.

We looked at the source data and noticed that the population of this municipality was around 22,000 in all years except for 1995 when it happened to be 11,000. This was apparently an error in the data[2]. Since we did not know the real number,

[2] We assure that the error reported here was really present in the data set and not artificially introduced. Data sets very often contain errors, even after careful verification. An important feature of the system is that it facilitates detecting errors in data by providing special displays and interactive tools.

we replaced the wrong value by the average value between the neighbouring years. After we had reloaded the system the appearance of the time plot changed (see Figure 4).

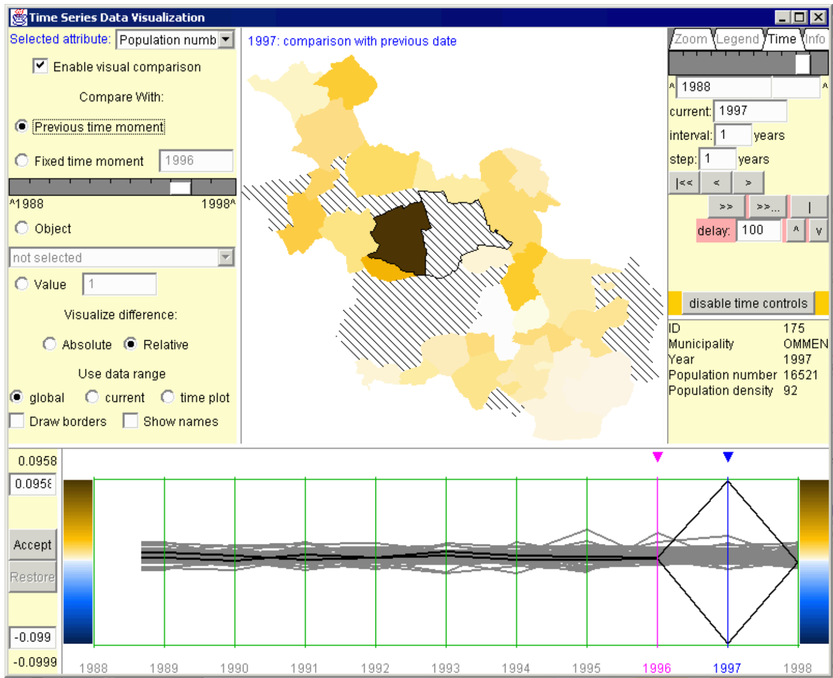

Figure 4. The display of changes after correction of the data shows opposite changes in 1997 in two neighbouring districts.

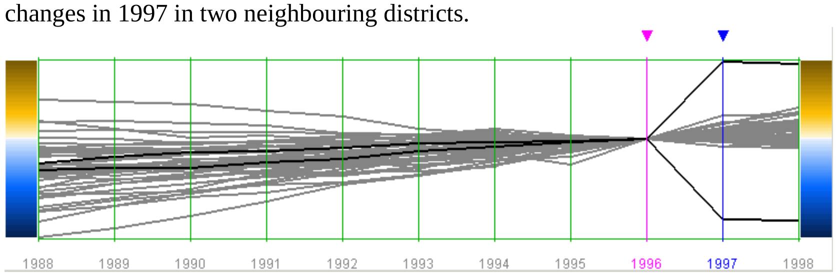

Figure 4a. The nature of the change becomes clearer in the mode of comparison with the year 1996.

Now we noticed on the plot two lines with opposite behaviours in 1997: one of them showed a sharp growth while the other an abrupt decline. In the mode of comparison with the year 1996 (Figure 4a) it may be seen that the population of these two districts was almost constant before 1997 and did not significantly change after 1997. Here are the exact values of the attribute for these two districts in the years around 1997:

	1995	**1996**	**1997**	**1998**
Ommen	18,312	18,355	16,521	16,493
Dalfsen	15,839	15,902	17,425	17,376

It looks as if a part of the population of Ommen moved in 1997 to Dalfsen. The link between the plot and the map allowed us to find easily on the map the districts corresponding to the "strange" trajectories. We found out that the districts were

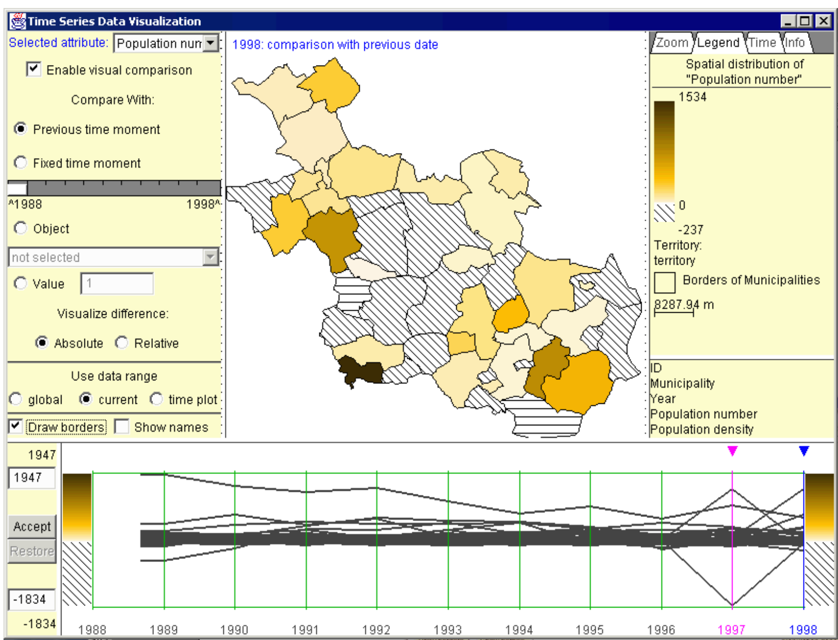

Figure 5. Classification of districts according to population change between 1997 and 1998. Diagonal hatching marks districts with decrease of population and horizontal hatching – districts without significant changes of the population number. Degrees of darkness of the remaining areas are proportional to the amounts of population increase in these areas.

neighbouring. This suggested a hypothesis that the change of population numbers actually reflects a change of boundary between the districts. Since we had only the current contours of the districts, we could not check this hypothesis by ourselves. Therefore we asked domain experts, and they confirmed our guess.

At the next stage we decided to test whether data mining techniques could help to relate changes of population number to other available characteristics of the districts. We focused on temporal trends during last years (see Figure 5). Representation of the changes on the map and the time plot demonstrated that in some municipalities the population steadily increased while in others it continuously decreased. We classified all the municipalities into 3 classes: districts with increase of population number, those with no significant changes, and districts with population number decreasing. These classes are represented on the map in Figure 5. For the black-and-white reproduction we filled the contours of the districts with population decrease using diagonal hatching. Horizontal hatching marks the districts with no changes. In the map it is possible to observe a rather prominent spatial pattern of the distribution of districts with different tendencies.

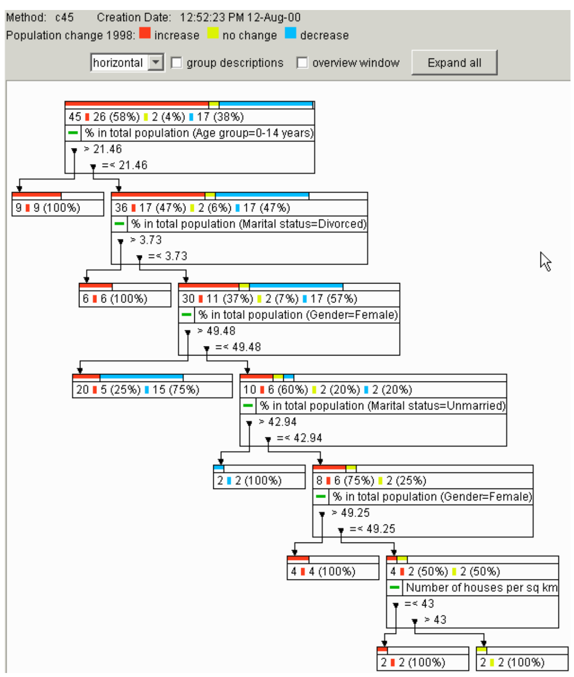

Figure 6. The decision tree produced by the C4.5 data mining algorithm relates population changes to values of other attributes.

For the classification thus made we ran the C4.5 data mining method. This algorithm tries to discriminate between the given classes on the basis of the values of available attributes and to produce a decision tree that divides the whole set of objects into groups as close as possible to the given classes. Each tree node represents a step in division based on values of one attribute. For example, the top node in the classification tree in Figure 6 divides the districts under study according to the proportion of children in population: below and above 21.46% (this value was selected by the algorithm). If the algorithm succeeds in discrimination of the classes, one can derive from the decision tree collective descriptions of the classes in terms of values of available attributes and, hence, reveal relationships between the classification and characteristics of the objects.

Thus, from the tree we got for our classification one can see that all districts with population decrease are characterised by less than 21.46% of children from 0 to 14 years in population, and most of them (15 of 17) have more than 49.48% of female. Characteristics of the districts with population increase are not so coherent, and the method divided them into smaller groups differing from each other in population structure.

5 Discussion and directions of further work

In our experiments on exploration of previously unknown data we succeeded to derive some new knowledge about the phenomenon characterised by these data, namely, population dynamics in a certain region. The instruments we developed (for a case of exploration of a single numeric attribute) proved to facilitate the following tasks from those listed in Section 2:

- all tasks involving the elementary view of data;
- view of data at the regional level at a selected time moment;
- analysis of spatial patterns of changes between two time moments;
- revealing of common temporal trends for groups of spatial objects;
- detecting of errors and anomalies in data.

Still, the instruments seem to be insufficient for analysis of change of a spatial pattern and of evolution of a spatial pattern over an interval. Therefore further research and development are required.

The important features of the system are:

- complementary role of different displays: the map represents the spatial component of data while the time plot supports study of evolution of attribute values;

- dynamic link between the displays provides easy navigation and access to data, and helps in understanding of complex spatio-temporal relationships;
- transformable displays expose changes and thus help to detect errors and anomalies in data and generate plausible hypotheses to explain findings;
- the link to data mining methods is useful for revealing relationships between changes of objects and characteristics of these objects.

The system was also tested with other census data sets on regional, national and European level (up to the NUTS-2 division with 207 different territory units). Application of the system to bigger data sets (for example, to the NUTS-3 territory division with more than 1,000 units) is restricted by current performance of Java technology. On the other hand, a human analyst can hardly discern so many trajectories represented simultaneously on a time-series plot. For such cases another solution may be suggested: the plot shows only the lines of the objects interactively selected in the map. In this way the user can investigate, for example, if adjacent districts are characterised by similar temporal trends.

In the future we shall work on enhancement of the existing system by:

- tools for interactive temporal aggregation of attribute values;
- new types of statistical graphic displays dynamically linked to maps and time plots;
- tools for multi-attribute studies;
- linking interactive analysis to other data mining methods and statistical procedures;
- combining interactive analysis of changes with map-centred exploratory spatial decision support (Jankowski, Andrienko, and Andrienko 2001).

However, we see our primary goal in supporting analysis of changes in spatial patterns and revealing of spatio-temporal trends, that is, evolution of spatial patterns over time. Currently there is no good method for this, and therefore research is needed before any tool can be created.

Acknowledgements

Our research is partly supported by the European Commission in projects EuroFigures - a digital portrait of the EU's General Statistics (Eurostat SUPCOM 97, Project n.13648, Subcontract to JRC 15089-1999-06) and SPIN ! - Spatial Mining for Data of Public Interest (IST Program, project No. IST-1999-10536). Orientation of our research interests to spatio-temporal data exploration was inspired by stimulating discussions on meetings of the ICA Visualization Commission (URL http://www.geovista.psu.edu/icavis/). We are grateful to M.-J.Kraak, C.Blok and N.Emmer (ITC) for providing data and checking our hypothesis. We thank reviewers of the paper for their useful comments.

References

Andrienko, G. and Andrienko, N. (1999a) Knowledge-Based Visualization to Support Spatial Data Mining. In Hand, D.J., Kok, J.N., and Berthold, M.R. (Eds.) *Advances in Intelligent Data Analysis*, Lecture Notes in Computer Science, vol. **1642**. Berlin: Springer-verlag, pp. 149-160

Andrienko, G. and Andrienko, N. (1999b), Data Mining with C4.5 and Cartographic Visualization. In N.W.Paton and T.Griffiths (eds*.) User Interfaces to Data Intensive Systems*, IEEE Computer Society. Los Alamitos, CA, pp. 162-165

Andrienko, N., Andrienko, G., and Gatalsky, P. (2000), Supporting Visual Exploration of Object Movement. In V. Di Gesu, 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

Avidan, T. and Avidan, S. (1999), ParallAX – a data mining tool based on parallel coordinates, *Computational Statistics*, **14** (1), 79-90.

Bertin, J. (1967/1983) *Semiology of Graphics. Diagrams, Networks, Maps* (Madison: The University of Wisconsin Press)

Blok, C., Koebben, B., Cheng, T., and Kuterema, A.A. (1999), Visualization of Relationships Between Spatial Patterns in Time by Cartographic Animation, *Cartography and Geographic Information Science*, **26** (2), pp. 139-151.

Brewer, C.A. (1994). Colour Use Guidelines for Mapping and Visualisation. In *Visualisation in Modern Cartography* (NY: Elsevier Science Inc.), pp.123-147.

Buja, A., McDonald, J.A., Michalak, J., and Stuetzle, W. (1991), Interactive data visualization using focusing and linking. In *Proceedings IEEE Visualization'91* (Washington: IEEE Computer Society Press), 156-163.

Buziek, G. (1999) Dynamic Elements of Multimedia Cartography, In *Multimedia Cartography: Cartwright,W., Peterson, M.P., Gartner, G. (Eds.)*, Springer-Verlag, Berlin Heidelberg, pp. 231-244.

Edsall, R., Peuquet, D., (1997), A graphical user interface for the integration of time into GIS. *Proceedings of the 1997 American Congress of Surveying and Mapping Annual Convention and Exhibition*, Seattle, WA, pp.182-189.

Fredrikson, A., Plaisant, C., and Shneiderman, B., (1999), 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 (NPIV'99)*, NY, ACM Press, pp.26-34

Inselberg, A. (1985). The plane with parallel coordinates. *Visual Computer*, **1** (1), pp.69-97.

Jankowski, P., Andrienko, N., and Andrienko, G. (2001), Map-Centered Exploratory Approach to Multiple Criteria Spatial Decision Making. International Journal Geographical Information Science, v.15 (2), pp.101-127

Kraak, M.J., Edsall, R., and MacEachren, A.M., 1997. Cartographic animation and legends for temporal maps: exploration and/or interaction, *Proceedings 18th International Cartographic Conference,* vol. **1**, pp. 253-261.

Monmonier, M., (1990). Strategies for the visualization of geographic time-series data. *Cartographica*, **27** (1), pp. 30-45.

Oberholzer, C., Hurni, L., (2000), Visualization of change in the Interactive Multimedia Atlas of Switzerland, *Computer & Geosciences*, **26** (1), pp. 423-435.

Quinlan, Q.R., (1993) *C4.5 Programs for Machine Learning*. Morgan Kaufmann

Slocum, T. A., S. C. Yoder, F. C. Kessler and R. S. Sluter (in press). MapTime: Software for exploring spatiotemporal data associated with point locations. Cartographica.

Tobler, W., 1970. Computer movie simulating urban growth in the Detroit region, *Economic Geography*, **46** (2), pp.234-240.

Unwin, A. and Wills, G. (1998), Exploring Time Series Graphically, *Statistical Computing and Graphics Newsletter*, (2), 13-15

Vasiliev, I.R. (1997), Mapping Time. *Cartographica*, 34 (2), pp.1-51.

Wrobel, S., Wettschereck, D., Sommer, E., and Emde, W. (1996), "Extensibility in Data Mining Systems", *Proceedings of KDD'96 2nd International Conference on Knowledge Discovery and Data Mining*, AAAI Press, pp.214-219

Yamahira, T., Kasahara, Y., and Tsurutani, T. (1985) How map designers can represent their ideas in thematic maps. *The Visual Computer*, **1**, 174-184