16 A Framework for Using Coordinated Displays for the Analysis of Multidimensional, Spatial, and Temporal Data

Natalia Andrienko, and Gennady Andrienko

Abstract. In geographic hypermedia, maps and other information displays are jointly used for the purpose of communicating information. We consider another role of maps: maps as instruments for data exploration and analysis. In this function, maps are also used in combination with other information displays. In order to establish links between multiple displays, various mechanisms have been developed. Some of these mechanisms might be useful in geographic hypermedia in addition to classical hyperlinks.

We propose a taxonomy of generic mechanisms for linking complementary information displays and, in broader terms, complementary tools for data exploration and analysis involving maps and graphics as well as querying, data transformation, and computation-based analysis techniques. We give an example of exploration of geographically referenced data in which different mechanisms are used in cooperation.

16.1 Introduction

Geographic hypermedia is conventionally defined as integration of maps with other media (text, graphics, animation, sound, images, and video) through hyperlinks. Geographic hypermedia is used primarily as a way to organise complex information for the purposes of communication. The role of any component of a hypermedia presentation is to carry a certain portion of information or a specific message the designer of the presentation wants to convey to the audience. This refers, in particular, to maps, the most appropriate medium for communicating geographical information.

However, the functions of maps are not restricted to communication only. Maps are also used for exploration and analysis of geographical information (MacEachren and Kraak 1997). In these functions, maps are also used together with other information displays, most often with graphics such as scatterplots, parallel coordinate plots, time graphs, histograms, etc. A general requirement is that multiple displays must be somehow interlinked in order to help the user to reconstruct the whole picture from the partial views. Not only hyperlinks but also other ways of display linking are used in such systems, which are conventionally called "geographic visualisation systems " rather than "geographic hypermedia".

This chapter is about various ways of linking displays in geographic visualisation systems and, more generally, about linking tools that are jointly used for visual exploration of geographic as well as non-geographic data. While it is not strictly related to the main topic of the book, it can still be interesting to the readers by broadening their eyesight to other uses of maps and related displays and other ways of display linking than they are accustomed to. It is quite possible that some of the techniques applied in geographic visualisation can be also effectively used in geographic hypermedia for information communication.

The ideas concerning the use of multiple linked displays for data communication and exploration come from the pre-computer era. Thus, the well-known representation of Napoleon's Russian campaign of 1812 created by Minard as early as in 1861 combines a map and a time graph[1]. In order to refer locations on the map to the marks on the graph showing the temperatures at the time moments when these locations were visited, Minard connected them with lines.

It is needless to say that computer technologies provide much broader opportunities for display linking than what was possible to do on paper. This encourages the designers and developers of visualization tools to build various display combinations and look for novel ways of linking them. Display coordination (i.e. linking between interactive displays simultaneously present on a computer screen) is currently a hot topic in information visualisation, geographic visualisation, and statistical graphics. Dedicated international conferences are convened (CMV 2003-2005) and special journal issues published such as (InfoVis 2003).

The most popular method of display coordination is known as "brushing" (this term seems to originate from the paper (Newton 1978)). The basic idea is that the user selects some items in one of the displays by direct manipulation through the mouse, and this results in the corresponding items being highlighted in other displays present on the screen. Since this idea was introduced, various more sophisticated forms of brushing have been suggested, for instance, multi-colour brushing: the explorer can use distinct colours for marking different selections.

[1] This visualisation is frequently cited; see, for example, (Tufte 1983). It can also be found in the Web, e.g. at http://www.edwardtufte.com/tufte/posters. The site http://www.math.yorku.ca/SCS/Gallery/re-minard.html contains many references to various citations of Minard's graphics as well as suggested revisions of it with the use of modern technologies.

Besides numerous variants of brushing, there are many other ways of display linking. This prompts some researchers to systemise the existing approaches and develop general frameworks of data visualisation through multiple views. Thus, Buja et al. (1996) classify visualization techniques into three categories: focusing, linking, and arranging views. Focusing includes the selection of subsets and variables (projections) for viewing and various manipulations of the layout of information on the screen: choosing an aspect ratio, zooming and panning, 3-D rotations etc. Focusing results in conveying only partial information and, therefore, must be compensated by showing different aspects of data in multiple views. These multiple views need to be linked so that the information contained in individual views can be integrated into a coherent image of the data as a whole. The method of linking depends on whether the views are displayed in sequence over time or in parallel. In the first case, linking is provided by smooth animation; in the second case, brushing may be used. The purpose of arranging multiple views is to facilitate comparisons. A possible approach is to display each view in a separate window and allow the user to arbitrarily arrange the windows.

North and Shneiderman (1997) suggest a "taxonomy of multiple window coordinations", which organises the strategies for display coordination along two dimensions. First, the user can perform two types of actions in one of the displays, selecting items or navigating the view. These actions can trigger automatic selection or navigation operations in another display, with three different combinations being possible: selecting items ↔ selecting items, navigating views ↔ navigating views, and selecting items ↔ navigating views. Second, two or more displays can represent either the same collection of information items or different collections. This gives a 3×2 matrix of possible variants of display coordination.

Recently, Roberts has published a comprehensive survey of the state-of-the-art in display coordination (Roberts 2005). The author touches upon various aspects of coordination:

- what tools for interaction are available to the user, e.g. sliders and buttons or direct manipulation facilities;
- how a new view resulting from user's actions is positioned with respect to the previous view, with three possible strategies: replacement, overlay, and replication;
- what conceptual models and architectures for coordination exist[2];

[2] Roberts himself advocates a layered approach based on the dataflow model. Visualisation is considered as the flow data → selected feature set → abstract visualisation object (data features mapped onto visual features) → rendered im-

- how multiple views can be managed in order to avoid overwhelming and disorienting the user;
- what role multiple linked views play in the exploration process.

It should be noted that Roberts, like many other researchers dealing with display coordination, presents software developer's perspective on the topic and concentrates primarily on technical issues. In our book (Andrienko and Andrienko 2006), we have considered the topic in a more user-oriented way. Our goal has been to instruct the potential users on how they can consciously combine various tools for the purposes of data exploration. We have extended the scope from "pure" coordination of data displays to the combined use of various exploratory tools including, in addition to visualisation, data transformation, querying, and computation-based analysis techniques.

In the present paper, we briefly introduce our taxonomy of tool linking modes and demonstrate an example of exploration of geographically referenced data in which different modes and mechanisms of tool linking are used complementarily.

16.2 Taxonomy of Tool Linking Modes

There is no tool for data exploration and analysis capable of everything. Each tool has different capabilities and, hence, a variety of tools have to be applied in the process of exploring a dataset. On the other hand, datasets that need to be analysed are often very large and consist of many components. It may be necessary to process data piecewise and link the fragmental information thus obtained into a coherent (mental) model of the data and the underlying phenomenon. Therefore, an explorer does not only need to apply different tools but also to apply one and the same tool several times, sequentially or concurrently, to different portions of data.

There are two basic modes of linking tools or multiple "instances" of the same tool in data analysis:

1. Sequential mode: a tool is applied to outcomes of another tool. The second tool starts its operation and produces results only after receiving the output of the first tool.
2. Concurrent mode: two or more tools or tool instances run simultaneously and independently from each other; the analyst needs to compare

age → window to manage the image. Coordination may occur at any level of the visualisation flow.

and/or relate their results. The tools or tool instances may be applied to the same data portion or to (partially) different data portions.

These two basic modes can be combined in various “hybrid” constructions; a detailed example is given in the next section. From the two modes, the sequential mode seems less relevant to geographic hypermedia except for cases when one or more of sequentially linked tools is *dynamic*, i.e. can change its results, for example, in response to user’s actions or updates of the data. Such a dynamic tool is used in the example presented in the next section: the output of the tool is dynamically modified as the user interactively changes the tool parameters.

In general, if “Tool 1” is dynamic (i.e. may change its outputs) and “Tool 2” is applied to its results, it is necessary that any changes of the results of “Tool 1” are properly accounted for in “Tool 2” and further along the chain. This means that “Tool 2” must be re-applied to the modified results of “Tool 1”, the tool following “Tool 2” in the chain of tools must be re-applied to the new results of “Tool 2”, and so on. Modern software packages often do such re-application automatically. When some tool modifies its results, it notifies all other tools using these results about the change occurred. In response, these other tools automatically update their own results and, in turn, notify the following tools in the sequence, and so on. In particular, result displays are also updated.

Automatic tool re-application is not always a benefit. It may be easier to compare results of several tool runs with different input settings when the results of the previous runs remain unchanged (until the user explicitly performs certain actions for changing them) than when all the tools are very reactive, so that all the results along a chain immediately change after even a slight interaction of the user with the first tool.

In the case of concurrent tool linking, an analyst can compare and relate results of two or more tools if they are appropriately visualised and, moreover, the visual displays can be seen simultaneously. This possibility to view several displays in parallel provides an elementary level of support to the analyst’s work on comparing the results and linking the information portions provided by the distinct tools into a coherent mental image.

Various mechanisms and tools can raise this level of support. We suggest the following taxonomy of the most common mechanisms for concurrent tool linking:

- Display coordination on the basis of a data *subset selection*: Several displays show information related to a selected subset of data records in such a way that the user can readily discern it from the rest. Different methods may be used to achieve this:

- Highlighting, i.e. special marking of display items corresponding to the selected records to make them easily discernible from the remaining items, e.g. by changing their colour or increasing the size.
- Focusing or zooming: a display is adjusted so that the information relevant to the selected subset is shown with the maximum possible expressiveness and legibility at the cost of the rest of the information being omitted or reduced in its conspicuousness.
- Filtering, i.e. removing the display items that do not correspond to the selected records or “muting” the visual appearance of these items and restricting the user interaction with them.

For example, classical brushing involves a query tool allowing the user to perform various selections through direct manipulation with display items. Highlighting is applied to the user-selected items in this display and to the items in the other displays corresponding to the same data records. Another example is an animation tool, which selects a particular time moment and, consequently, all data records involving this time moment. Thus, if the data are spatial time series, the animation tool selects all such pairs <*location, time*> in which *time* equals the currently chosen time moment.

- Display coordination on the basis of a data *set division*: the data set is divided into two or more non-overlapping subsets (for example, using a classification tool), and, in response, several displays show the information relevant to each subset so that it is easily recognisable and distinguishable from the information related to the other subsets. This can be achieved in following ways:
  - Multi-colour marking: each subset receives a certain unique earmark (typically a colour), which is used for marking the display elements corresponding to this subset in all coordinated views.
  - Display multiplication: a display is replaced or supplemented by several displays of the same type so that each display represents information related to one of the subsets.
  - Re-arrangement of display items: display items are positioned within the display space in such a way that the items corresponding to the same subset are situated close to each other.

- Linking on the basis of a *common visual encoding* of data in several displays such as:
  - Common scales along the display dimensions;
  - Common meanings of colours, sizes, symbols, etc. throughout the displays.

  When the encoding is changed by means of display manipulation, the changes must affect all the linked displays.

For example, data about earthquake occurrences can be represented on a map and in a space-time cube (Gatalsky *et al*. 2004) by identical circles with the sizes being proportional to the earthquake magnitudes. When the user switches from the linear to logarithmic or exponential scale of encoding of the magnitudes by circle sizes, the new scale is applied both in the map and in the space-time cube.

- Linking on the basis of a *common data transformation*: one and the same transformation technique is applied to data represented on several displays. When the user changes the transformation, this affects all the linked displays.
  For example, several maps may represent values of various attributes referring to a particular time moment. The user may apply a tool transforming the original values to differences with respect to the values for the previous time moment. The transformation takes place simultaneously on all the maps. When the user switches the tool to computing the ratios to the previous moment instead of the differences, each map is updated to reflect the change.

Besides these generic methods of tool linking, there are also methods specific for particular display types. For example, some visualisation techniques such as a map with bar charts may use colours to distinguish between attributes. If there are several displays using colours for attributes, a special coordination mechanism can maintain the consistency of the assignment of the colours to the attributes throughout these displays.

Not only sequential and concurrent modes of tool linking are often used together but also different mechanisms of sequential and concurrent linking can work simultaneously, for example, filtering together with multi-colour marking and common transformation of attribute values.

16.3 Cooperation between Sequential and Concurrent Linking: An Example

In this example, an analyst explores data concerning the health care in different counties of the state of Idaho (USA)[3] in order to determine which counties are the most in need of support for improving the availability and accessibility of health care facilities for the population. The explorer needs to evaluate the counties on the basis of multiple attributes, specifically,

[3] The example dataset was provided by Prof. Piotr Jankowski from the University of Idaho, USA. The data are described in more detail in (Jankowski et al. 2001).

- *N of estimated unmet visits*: the estimated number of unmet doctor visits when people coming to see a doctor cannot be attended due to doctor's overload;
- *Low-weight birth rate*: the percent of infants born with insufficient body weight averaged from a multi-year interval;
- *Burden on on-call providers*: the number of hours on call for each provider;
- *Population in >35 miles from hospital*: the number of individuals residing outside the influence zone (i.e. the 35-mile radius, according to the national standard for rural areas) of the nearest hospital.

To make the values of all attributes comparable, the analyst has transformed them into z-scores, which express the relative deviations of the values from the mean values of the respective attributes. Positive z-scores signify the original values being worse than the average values for the state of Idaho.

Now, the transformed values of the multiple attributes need to be somehow combined into integral scores characterising the situation in each county in general. For this purpose, the analyst decides to apply a special evaluation support tool capable of attribute integration. We shall not describe the specific algorithm of value combination applied in the tool since this is not relevant to our discussion. It is only important that, first, the tool produces new attributes and, second, the tool is dynamic, i.e. changes the values of these new attributes when the user modifies tool settings.

One of the new attributes produced by the tool is the integrated scores of all counties expressed as real numbers from 0 to 1; the higher the score, the more problematic is the situation in the county. Additionally to the scores, the tool produces an attribute reflecting the ranking of the counties from the most problematic (i.e. with the highest score) to the least problematic (i.e. with the lowest score). The ranks are specified as integer numbers from 1 to 44, which is the number of the counties in the state of Idaho.

The settings influencing the tool outputs are relative weights assigned to the attributes being integrated. The weights are specified as real numbers between 0 and 1. The sum of the weights of all attributes participating in the computation must equal 1. By default, all attributes receive equal weights. In our example, there are four attributes; accordingly, each of them receives the weight 0.25. The evaluation tool provides a direct manipulation interface for changing the weights. When the user modifies the weights, the tool dynamically re-computes the values of the derived attributes.

The table display in Fig. 16.1 shows the evaluation scores and ranks of the Idaho counties obtained with the attribute weights 0.3, 0.3, 0.2, and

	N of estimated unmet visits (z)	Low-weight birth rate (z)	Burden on on-call providers (z)	Population in >35 miles from hospital (z)	Evaluation score	Ranking
Washington	0.556	1.997	2.071	0.556	0.7354	1
Payette	0.650	1.099	2.071	0.570	0.6847	2
Jerome	0.678	0.873	2.071	-0.182	0.6436	3
Latah	1.172	1.617	-0.293	-0.182	0.5941	4
Madison	0.912	-0.138	-0.341	3.246	0.5926	5
Clearwater	0.540	-1.934	1.008	4.926	0.5913	6
Gooding	-1.398	2.774	1.157	-0.182	0.5849	7
Twin_Falls	1.198	1.010	-0.065	-0.182	0.5694	8
Clark	0.632	-0.356	2.071	-0.411	0.5538	9
Gem	-1.079	1.124	2.071	-0.182	0.5493	10
Power	0.358	1.997	-0.766	-0.182	0.5416	11
Blaine	0.629	1.431	-0.554	-0.411	0.5253	12
…	…	…	…	…	…	…
Caribou	-0.313	-0.567	-0.894	-0.182	0.3291	38
Bear_Lake	-0.136	-1.101	-0.453	-0.411	0.3214	39
Boise	-0.375	-0.914	-0.686	-0.182	0.3144	40
Nez_Perce	-1.340	0.185	-0.899	-0.411	0.3037	41
Adams	-2.338	0.088	-0.261	-0.411	0.2694	42
Elmore	-3.380	0.040	-0.777	-0.411	0.1729	43
Custer	-2.903	-1.214	-0.479	-0.411	0.1391	44

Fig. 16.1. The table display represents the results of evaluating the counties of Idaho on the basis of four attributes with the weights 0.3, 0.3, 0.2, and 0.2. The table columns show the values of the source four attributes (previously transformed into z-scores), the evaluation scores derived, and the ranking of the counties according to the scores. The table rows are arranged in the order of decreasing evaluation scores and, consequently, increasing ranks.

0.2, respectively. The table rows are arranged in the order of decreasing evaluation scores and, consequently, increasing ranks. Additionally to showing the attribute values as numbers, the same values are also represented visually by dark grey shading within the cells. The left edge of the shaded area indicates the relative position of the value contained in the cell between the minimum and the maximum values of the respective attribute: the higher the value, the larger the shaded area. To save the space, we did not include all the 44 rows of the table in Fig. 16.1 but only the top and bottom parts of the table. However, it may be clearly seen how the ranking is related to the evaluation scores.

Let us suppose that some limited funding for health improvement is available and can be divided between at most five counties in need. Some extra funding is expected in near future, which will allow the state administration to support additional five counties. Accordingly, the task of the analyst is to divide the 44 counties of Idaho into three classes:

1. The most needy counties, which will receive an immediate financial support.
2. The counties that will be supported later, when the additional funding comes.

3. The counties where the state of the health care is satisfactory so that they will not be funded.

A convenient way to do this division is to apply a tool for classification on the basis of a numeric attribute to the results of ranking of the counties. The counties with the ranks from 1 to 5 will be the candidates for immediate funding, the counties with the ranks from 6 to 10 will be included in the waiting list, and the remaining counties will not be funded.

This classification is shown on the map display in Fig. 16.2. The class breaks specified by the analyst are 5.01 and 10.01. The analyst uses these values rather than 5 and 10 in order to ensure that the county with the rank 5 is included in the first class and the county with the rank 10 in the second class. The counties on the map are coloured according to the classes they belong to. The black colour corresponds to the first class, the dark grey colour – to the second class, and the light grey – to the third class.

Fig. 16.3 demonstrates the effect of the classes being propagated to the table display. The display responds to the classification by appropriate colouring of its rows. Additionally, the rows of counties belonging to the same class are put together. At the top of the table, we can see 5 black rows, which correspond to the counties included in the first class. These

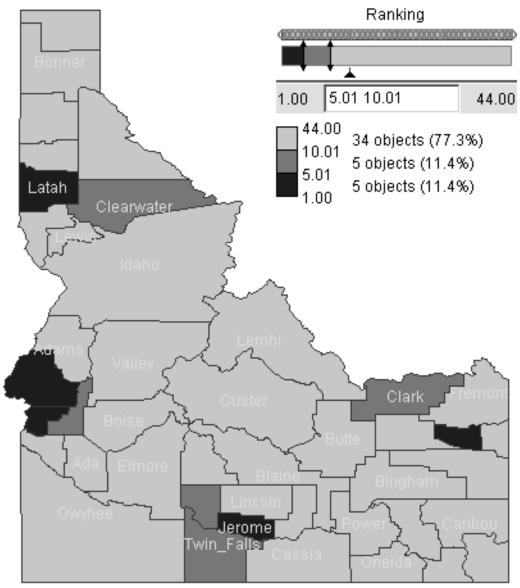

Fig. 16.2. Counties of Idaho are classified according to their ranking into the counties with the ranks from 1 to 5 (the candidates for receiving an immediate support), those with the ranks from 6 to 10 (to receive a support later), and the rest, which will not be supported.

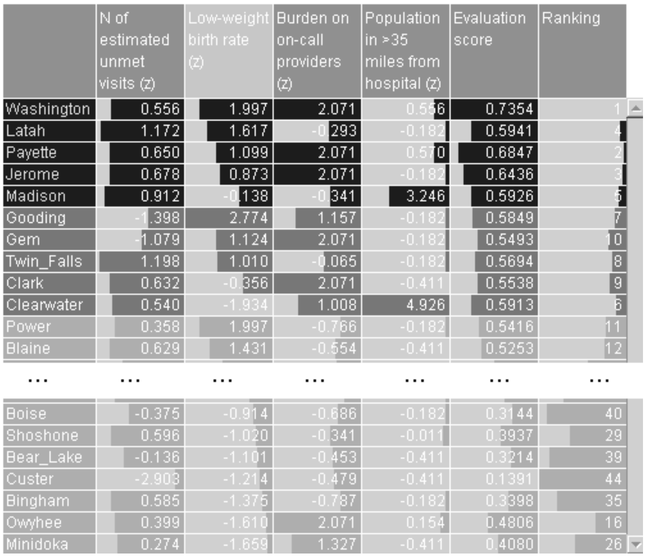

	N of estimated unmet visits (z)	Low-weight birth rate (z)	Burden on on-call providers (z)	Population in >35 miles from hospital (z)	Evaluation score	Ranking
Washington	0.556	1.997	2.071	0.556	0.7354	1
Latah	1.172	1.617	-0.293	-0.182	0.5941	4
Payette	0.650	1.099	2.071	0.570	0.6847	2
Jerome	0.678	0.873	2.071	-0.182	0.6436	3
Madison	0.912	-0.138	-0.341	3.246	0.5926	5
Gooding	-1.398	2.774	1.157	-0.182	0.5849	7
Gem	-1.079	1.124	2.071	-0.182	0.5493	10
Twin_Falls	1.198	1.010	-0.065	-0.182	0.5694	8
Clark	0.632	-0.356	2.071	-0.411	0.5538	9
Clearwater	0.540	-1.934	1.008	4.926	0.5913	6
Power	0.358	1.997	-0.766	-0.182	0.5416	11
Blaine	0.629	1.431	-0.554	-0.411	0.5253	12
…	…	…	…	…	…	…
Boise	-0.375	-0.914	-0.686	-0.182	0.3144	40
Shoshone	0.596	-1.020	-0.341	-0.011	0.3937	29
Bear_Lake	-0.136	-1.101	-0.453	-0.411	0.3214	39
Custer	-2.903	-1.214	-0.479	-0.411	0.1391	44
Bingham	0.585	-1.375	-0.787	-0.182	0.3398	35
Owyhee	0.399	-1.610	2.071	0.154	0.4806	16
Minidoka	0.274	-1.659	1.327	-0.411	0.4080	26

Fig. 16.3. The classes of counties from Fig. 16.2 have been propagated to the table display. The rows of the tables are grouped according to the classes the corresponding counties belong to. The grouping is accompanied by colouring of the rows.

are followed by five dark grey rows corresponding to the second class of counties. The remaining rows of the table are coloured in light grey.

Since the effect of grouping is hardly visible when the rows are ordered according to decreasing evaluation scores or increasing ranks, another attribute, “Low-weight birth rate”, has been used in Fig. 16.3 for sorting the table rows. This resulted in a different order of the rows as compared to Fig. 16.1. Grouping of table rows has a priority over ordering: the rows are first grouped according to the current division of the set of counties, and then the specified method of ordering is applied to each group individually.

In our case, the rows in each section of the table (i.e. black, dark grey, and light grey) are arranged in the order of decreasing low-weight birth rates. It can be seen that the row with the highest value of this attribute, 2.774, is put on the sixth position in the table, after a row with a much lower value, -0.138. This can be explained by the priority of grouping over ordering: the row with the value -0.138 corresponds to the county Madison, which belongs to the first class, and is put together with the rows of the other counties from the same class. The row with the value 2.774 corresponds to the county Gooding, which belongs to the second class. Accordingly, this row is put at the top of the dark grey section of the table.

Let us now suppose that, after some deliberation, the analyst decides to change the weights of the attributes so as to increase the influence of the attribute “Population in >35 miles from hospital”. She increases the weight of this attribute to 0.25, and the evaluation tool decreases automatically the weights of the other attributes proportionally to the values they had before the operation. The resulting weights are 0.28, 0.28, 0.19, and 0.25.

In response, the evaluation tool immediately re-computes the integrated scores and re-ranks the counties. In the result, the values of the derived attribute “Ranking” change, and the classification tools needs to re-classify the counties. The results are shown in Fig. 16.4. It may be seen that the class breaks have not changed as compared to Fig. 16.2 but the content of the classes is slightly different: the neighbouring counties Latah and Clearwater on the north have exchanged their classes.

In Fig. 16.5, we can see the effect of propagating the new division of the set of counties to the table display. In the result, the rows of the counties Latah and Clearwater have changed their colours and have been moved to other sections.

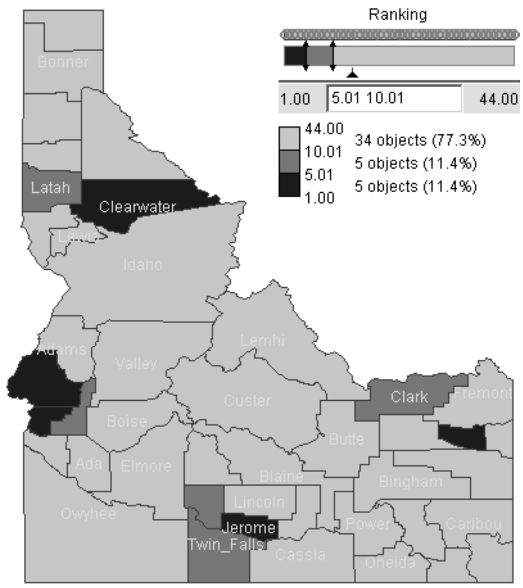

Fig. 16.4. After changing the attribute weights to 0.28, 0.28, 0.19, and 0.25, the counties have been re-evaluated and re-ranked, and the classification re-applied to the new ranks. In the result, the content of the classes has slightly changed.

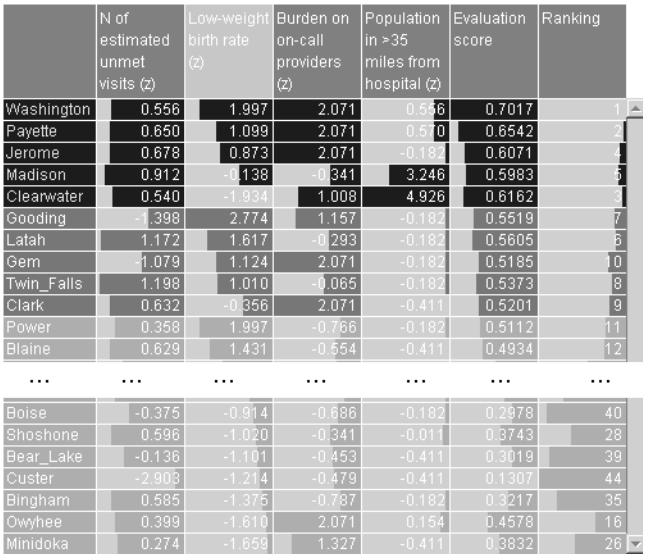

	N of estimated unmet visits (z)	Low-weight birth rate (z)	Burden on on-call providers (z)	Population in >35 miles from hospital (z)	Evaluation score	Ranking
Washington	0.556	1.997	2.071	0.556	0.7017	1
Payette	0.650	1.099	2.071	0.570	0.6542	2
Jerome	0.678	0.873	2.071	-0.182	0.6071	4
Madison	0.912	-0.138	-0.341	3.246	0.5983	5
Clearwater	0.540	-1.934	1.008	4.926	0.6162	3
Gooding	-1.398	2.774	1.157	-0.182	0.5519	7
Latah	1.172	1.617	-0.293	-0.182	0.5605	6
Gem	-1.079	1.124	2.071	-0.182	0.5185	10
Twin_Falls	1.198	1.010	-0.065	-0.182	0.5373	8
Clark	0.632	-0.356	2.071	-0.411	0.5201	9
Power	0.358	1.997	-0.766	-0.182	0.5112	11
Blaine	0.629	1.431	-0.554	-0.411	0.4934	12
…	…	…	…	…	…	…
Boise	-0.375	-0.914	-0.686	-0.182	0.2978	40
Shoshone	0.596	-1.020	-0.341	-0.011	0.3743	28
Bear_Lake	-0.136	-1.101	-0.453	-0.411	0.3019	39
Custer	-2.903	-1.214	-0.479	-0.411	0.1307	44
Bingham	0.585	-1.375	-0.787	-0.182	0.3217	35
Owyhee	0.399	-1.610	2.071	0.154	0.4578	16
Minidoka	0.274	-1.659	1.327	-0.411	0.3832	26

Fig. 16.5. The altered classes of counties have been propagated to the table display, which resulted in re-colouring and re-grouping of the rows.

This example demonstrates a joint use of four different tools:

- the evaluation tool;
- the classification tool, which is applied to the output of the evaluation tool;
- the cartographic visualisation tool, which shows the output of the classification tool on a map display;
- the visualisation tool producing the table display, which shows, along with other attributes, the results of the evaluation tool. The table visualisation tool is linked to the classification tool by means of the class propagation mechanism.

The links between the tools are schematically represented in Fig. 16.6. The solid arrows represent the sequential tool linking, when results of one tool are used as an input to another tool. In this case, the evaluation tool produces new attributes, which are supplied to the table visualisation tool and to the classification tool. The latter two tools are used concurrently and coordinated by means of propagation of the division of the set of data records into subsets, or classes (in this case, there are 44 data records corresponding to the 44 counties of Idaho). This link, which is signified by the dotted arrow, allows the analyst to compare and relate the information provided in the table display (i.e. the characteristics of the counties) to the

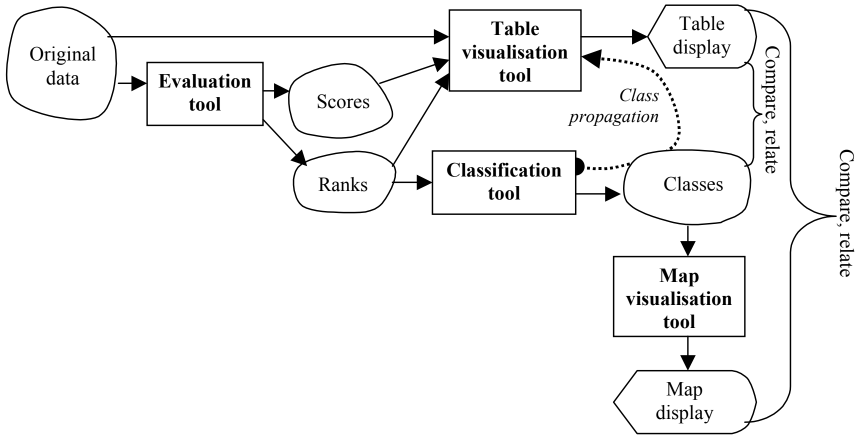

Fig. 16.6. The links between the tools in the example with evaluating the counties of Idaho. The solid arrows represent the sequential mode of tool linking, and the dotted arrow indicates the coordination of the concurrently used tools.

results of the classification. Besides, this link also allows the analyst to compare and relate the information provided by the table to the information contained in the map display, which shows the geographical positions of the counties belonging to the different classes.

16.4 Conclusion

Comprehensive data exploration and analysis essentially requires multiple tools to be used in combination. This includes diverse visual displays of various aspects of data as well as tools for data transformation, querying, computational analysis, etc. In this paper, we have tried to define the generic mechanisms of how different tools can be used together. We expect this taxonomy to be helpful not only for users of data analysis tools and developers of such tools but also for designers of geographic hypermedia.

In the suggested general framework of tool linking, we pay a particular attention to dynamic tools, which may change their results in response to user's actions or other events. If the results are used or reflected in other tools, it is necessary that the changes were properly accounted for. Tool or hypermedia designers should care about this by devising appropriate mechanisms of change propagation.

The suggested framework does not only address spatial and temporal data but is more general and applicable to other types of data as well.

While spatial and temporal data may require specific analysis tools (e.g. maps, animated displays, time graphs, space-time cubes, etc.), the modes and mechanisms of tool linking and coordination are basically the same as for tools oriented to other data types. Moreover, the generic tool linking mechanisms allow specific tools for analysis of spatial and/or temporal data to be used together with tools that do not assume the spatial or temporal nature of the data. For example, maps or time series graphs can be used in combination with scatterplots, histograms, table displays, and so on. This provides complementary views of various data aspects and features and increases the comprehensiveness of analysis or information presentation.

References

Andrienko N, Andrienko G (2006) Exploratory analysis of spatial and temporal data, Springer, Berlin

Buja A, Cook D, Swayne DF (1996) Interactive high-dimensional data visualization. Journal of Computational and Graphical Statistics, vol 5, pp 78-99

CMV (2003) Proceedings of the 1st International Conference on Coordinated and Multiple Views In Exploratory Visualization. IEEE Computer Society, Los Alamitos, London, England

CMV (2004) Proceedings of the 2nd International Conference on Coordinated and Multiple Views In Exploratory Visualization. IEEE Computer Society, Los Alamitos, London, England

CMV (2005) Proceedings of the 3rd International Conference on Coordinated and Multiple Views In Exploratory Visualization, IEEE Computer Society, Los Alamitos, London, England

Gatalsky P, Andrienko N, Andrienko G (2004) Interactive analysis of event data using space-time cube. In: Banissi E et al. (eds) Proceedings of IV 2004 – 8th International Conference on Information Visualization, IEEE Computer Society, Los Alamitos, London, UK, pp 145-152

InfoVis (2003) Information Visualization. Special Issue on Coordinated and Multiple Views in Exploratory Visualization, Roberts JC (ed), vol 2(4)

Jankowski P, Andrienko N, Andrienko G (2001) Map-centered exploratory approach to multiple criteria spatial decision making. International Journal Geographical Information Science, vol 15(2), pp 101-127

MacEachren AM, Kraak MJ (1997) Exploratory cartographic visualization: advancing the agenda. Computers and Geosciences, vol 23(4), pp 335-344

Newton CM (1978) Graphics: from alpha to omega in data analysis. In: Wang PCC (ed) Graphical representation of multivariate data, Academic Press, New York, pp 59-92

North C, Shneiderman B (1997) A taxonomy of multiple-window coordinations. Technical Report CS-TR-3854, University of Maryland Computer Science Department

Roberts JC (2005) Exploratory visualization with multiple linked views. In: Dykes J, MacEachren AM, Kraak M-J (eds) Exploring geovisualization. Elsevier, Oxford, UK, pp 159-180