

Multivariate Visualization of Socio-Economic Indicators Using Geographic Maps

Mike Patterson and Suresh Lodha
Department of Computer Science
University of California, Santa Cruz
{mike,lodha}@soe.ucsc.edu
Tech Report Number UCSC-SOE-08-24

Abstract—National data on several socio-economic indicators are being collected and disseminated by various agencies including World Bank and United Nations. These indicators are typically visualized using well-known bivariate techniques including geographic maps, line graphs, bar graphs, and scatter plots, with some notable exceptions such as the visualization of four variables on a scatter plot by adding glyph sizes and colors. Judicious use of multivariate visualization of several global indicators may allow useful insights. In this work, we introduce a novel visualization technique that allows pseudo-coloring of the boundary of a country based on an indicator. Boundaries of neighboring countries are separated using a displacement method along normals. This visualization coupled with three standard techniques of pseudo-coloring a country, glyph size, and glyph color, and judicious use of opacity allows us to display four variables simultaneously on a geographic map effectively. Working with sociologists, we applied our techniques to investigate global socio-economic indicators available in reliable databases. We have built a flexible, interactive, real-time visualization system that allows extensive investigation of these indicators by sociologists and computer scientists working together. Flexibility of our visualization platform allows us to map a fifth variable using cartograms. Our system also supports real-time animation of indicators over time. We have demonstrated that the multivariate visualization of global indicators provides valuable insights in many cases not easily available otherwise.

Index Terms—Geographic visualization, multivariate, global inequality, socio-economic indicators.

1 INTRODUCTION

Large amounts of national data on socio-economic indicators are being collected and disseminated by a number of international agencies including World Bank [1], United Nations [22], UNESCO [5], and OECD [14]. Several websites [20, 6, 13, 21, 3] are allowing general users to visualize these data sets using classical visualization techniques including line graphs, bar graphs, scatter plots, and pseudo-coloring of world maps.

The global socio-economic data set is a three-dimensional matrix, with the three axes being countries, indicators, and time. Most of the existing techniques allow visualization of a few rows or columns of this three-dimensional data set. For example, a time series graph can display data for a few countries for one indicator over time or a bar graph can display a few indicators for a few countries for a fixed time period. Most popular is perhaps mapping of an indicator to pseudo-colors on a world map which is then animated over time. More recently, Gapminder [6] introduced a 4-variate display using scatter plots where x-axis and y-axis are mapped to two different indicators and two additional indicators are displayed using glyph size and colors for a fixed time period. These scatter plots can be animated over time and allows zooming features to explore further relationship between indicators and countries. Judicious use of these visualizations have been very helpful in investigating underlying associations in an intuitive manner.

In this work, we present a multivariate visualization technique for displaying geographic data that supports display of five variables simultaneously. In addition, the sixth variable, time is displayed using animation. Although several novel techniques have been introduced to display multivariate and spatio-temporal geographic data that we discuss later in Section 2, we have leveraged and focused primarily on techniques that are simple to understand, an important characteristic for visualizations to become adopted and attractive for experimental investigation. The simplicity and familiarity of visualizations is also important while working collaboratively with social scientists to explore socio-economic data.

To keep the visualization simple and effective at the same time, we introduce a simple novel visualization technique where an indicator is mapped to the pseudo-color of the boundary of a country.

Boundaries of neighboring countries are separated using a displacement method along normals. In addition, we use three well-known visualization methods, namely pseudo-coloring of the country, glyph size, and glyph color. Furthermore, we have employed cartograms to display a fifth variable to establish the versatility and flexibility of our system to incorporate additional visualization techniques. This static visualization can be animated in real-time to display the evolution of these indicators.

We believe that judicious use of this multivariate display allows users to explore underlying associations between variables easily. The reverse is true as well, that is, mapping four arbitrary variables to this visualization may produce some plethora of colors, that may be hard to understand. As we will later demonstrate, opacity also plays an important role in the proposed multivariate display. We have built a flexible, interactive, real-time visualization system that allows extensive investigation of the socio-economic indicators from a variety of databases.

To investigate the usefulness of this multivariate display, we applied our visualization system to examine the relationship between health indicators and variables that impact them. We chose to investigate the relationship between these variables to gain better understanding of the causality model recently proposed by Cornia *et al.* [2] to explain the evolution of health indicators. We describe the results of our investigation in Section 5.

2 BACKGROUND AND RELATED WORK

We now describe some geographic visualization techniques for various combinations of multivariate, spatial, and temporal data. We then also discuss related work on geographic visualizations supported by web sites and contrast these with our work.

2.1 Geographic Visualization Techniques

A common technique for geographic visualization is to encode a spatially-mapped variable as the polygon fill color. Although this is traditionally univariate, many researchers have explored clever color combinations that compress two or more variables in the same area. Guo *et al.* [9] demonstrate the use of systematically designed color

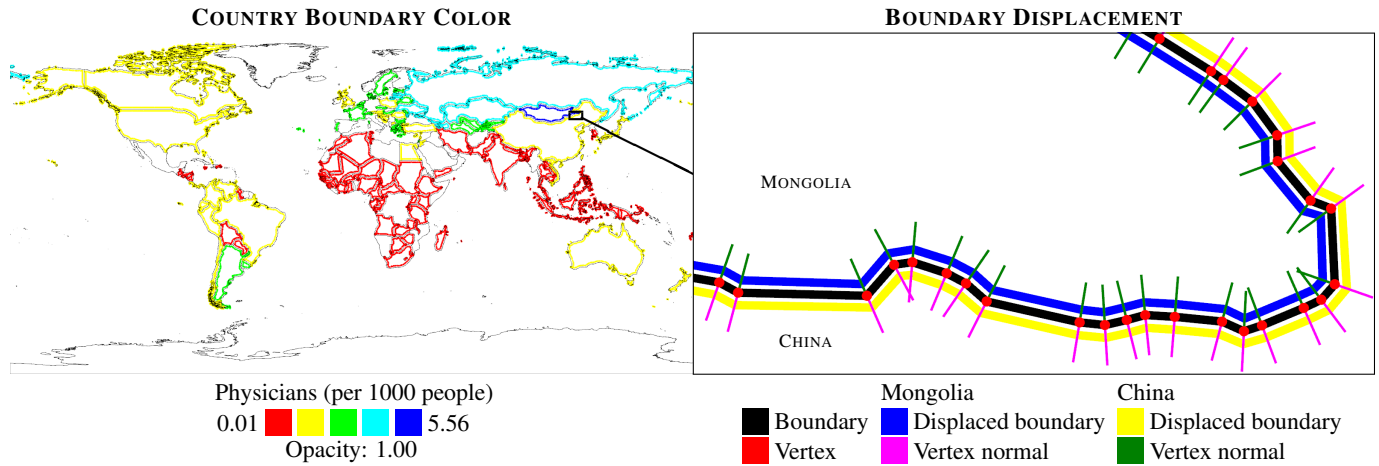


Fig. 1. The political boundaries are used to visualize GHND data from 1986 and are offset from one another to avoid ambiguities between adjacent countries. The enlarged Mongolia/China border on the right illustrates the technique used to achieve this effect: Mongolia's vertex normals (magenta) are used to displace its border (blue) and similarly China's vertex normals (green) are used to displace its border (yellow).

schemes to encode two variables in a two-dimensional color table. Taking this one step further, Hagh-Shenas *et al.* [10] use noise images and color weaving to encode up to six variables in the same area. However, as noted by Shanbhag *et al.* [19], both of these techniques push human visual perception to the limit and heavily rely on multiple, unique color schemes which can inhibit understanding.

Another common visualization technique places glyphs on the map and exploits properties such as size and fill color to represent variables. Glyphs typically consist of simple shapes, such as circles, and are placed at centroid of the associated region. Shanbhag *et al.* [19] use specialized glyphs (wedges, rings, and slices) to encode temporal data in a single view, and thus avoid the inattention blindness and change blindness associated with animation-based solutions. Spatio-temporal visualization is also supported using mashups by Wood and Dykes [23]. Dykes and Brundson [4] have also proposed geographically weighted visualizations for scale exploratory analysis.

Cartograms, *i.e.*, maps in which the sizes of geographic regions such as countries or provinces appear in proportion to their population or some other analogous property, have been popularized in recent years by Gastner and Newman [8]. This technique produces attractive global comparisons that are easily understood, yet are limited to visualizing only one or two variables. Combination of cartograms with pixelmaps have been used by Panse *et al.* [15].

Additional efforts include use of textures in visualization [12, 11] and two-tone pseudo coloring [18] which maps one-dimensional data into compact rectangles suitable for use in tables and as glyphs.

Although outlining regions as a means to distinguish one from another is commonplace in spatial visualizations, it is problematic in cartography because ambiguities arise on the boundaries between adjacent regions. We propose a new boundary displacement technique which eliminates the adjacency ambiguities and works well in conjunction with preexisting visualization techniques.

2.2 Geographic Visualization Applications

We now discuss some applications that support and promote geographic visualizations. One approach, used by the Worldmapper Project [3], places the burden of creating effective geographic visualizations in the hands of software developers with some cartographic experience. These maps are further transformed into publication quality posters with the addition of supporting text and figures. Though they provide excellent global comparisons, the typical consumer is limited to using preexisting static maps which illustrate only one variable. Our application supports the development of publication quality maps without sacrificing interactivity and multivariate visualization.

Web applications have gained momentum in recent years and use either Flash, Java applets, or JavaScript to provide interactivity. Gap-

minder [6] developed a powerful Flash application that is capable of displaying up to four variables over time on a two-dimensional scatterplot and up to two variables over time on a geographic map. Their scatterplot's major limitation is that it requires significant user interaction to differentiate each of the countries and their map is restricted to displaying only two variables. Our application merges the effectiveness of multivariate visualization with the familiarity of geographic maps.

We have developed Tart as a native application and describe our reasons for this choice. The tradeoffs associated with using web applications often overshadow its benefits of accessibility and ubiquity. Input data is generally limited to that which is hard-coded into the application back-end and output options primarily consist of screen captures and textual data dumps, neither of which are suitable for publication quality maps. Developing publication quality geographic visualizations also requires intense graphical computation, which can benefit from code optimizations designed for the CPU and GPU; this luxury is typically unavailable to web applications. Professional visualization products, such as *Tableau* and *Autodesk Map 3D*, are almost exclusively developed as native applications and it is for these reasons that we chose to implement our visualization system as a native application.

3 TART VISUALIZATION

We refer to our visualization system as *Tart*. Tart provides the ability to display up to five variables simultaneously using a geographic map. In addition, this visualization can be animated over time using a slider. The system includes many options to customize the resultant visualization so that it is both functional and insightful. Each visualization method – country boundary color, country fill color, glyph fill color, glyph size, and cartogram – can be individually bound to any socio-economic variable chosen from a database. The variables are typically divided amongst uniformly-sized bins, represented by either a color or a size. Currently, the three color-based methods can each be mapped to one of eight preset color legends and have an opacity slider to permit appropriate blending with the background. The preset color legends are derived from commonly used linear RGB gradients, such as rainbow coloring. An image of Earth can be displayed at various opacities in the background. The following subsections describe the individual components in further detail.

3.1 Country Boundary Color

The first visualization method relies on colored political boundaries that are displaced inwards to disambiguate adjacent countries. We believe that this technique is novel and is being used for the first time in

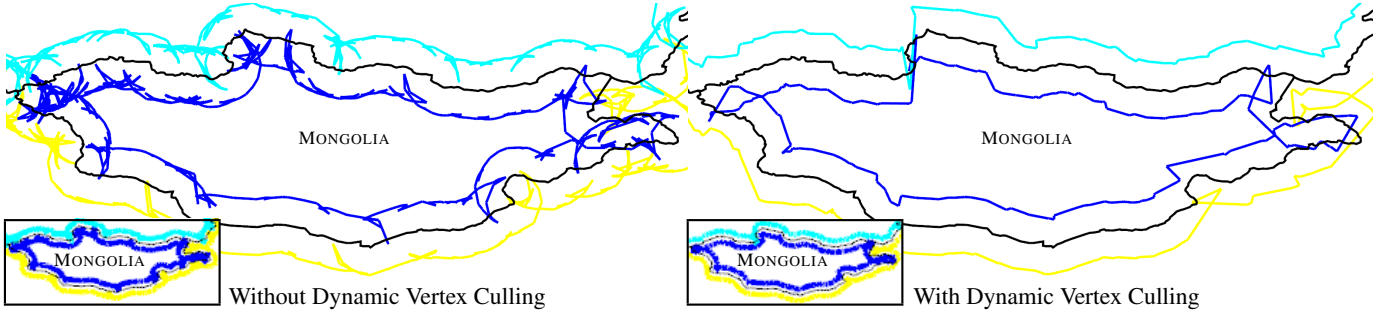


Fig. 2. These images represent significantly enlarged views of the displaced boundaries that are rendered at the lowest zoom level, showing the advantage of using dynamic vertex culling. On the left, all displaced vertices are drawn. This results in egregious self-intersections. The right image illustrates the significant visual improvement gained by dropping vertices that are too close in pixel space to the previously drawn vertex. In this case, Mongolia’s displaced boundary is simplified to four self-intersections.

geographic maps. As shown in Figure 1, the boundary is inset by moving each vertex in the polygon in the opposite direction of the normal. The amount of displacement is calculated based on the zoom level. However, this algorithm alone is unsuitable for concave polygons and polygons with low vertex counts, which together include most of the world’s polygons. At low zoom levels, these polygons contain sections where the distance in pixel space between edges with nearly-opposite normals is less than two times the width of the rendered line and the displacement, and thus exhibit the self-intersection artifacts shown in the left diagram of Figure 2.

To address this, the algorithm is modified to drop vertices that are too close in pixel space to the previously drawn vertex. The results of applying this algorithm to the boundary of Mongolia are shown in the right diagram of Figure 2. This dynamic vertex culling algorithm is still susceptible to self-intersections by non-neighboring edges, but in practice, the remaining artifacts are negligible at low zoom levels and disappear as the user zooms in. A robust solution would need to consider a larger subset of, if not all, vertices per vertex to determine each of their displacement ranges as a function of the zoom level. Such an algorithm would approach quadratic complexity and require additional data per vertex. The vertex culling algorithm does not require additional precomputed data and runs in real-time with linear complexity.

3.2 Country Fill Color

As part of Tart’s initialization, each polygon is tessellated using a standard algorithm [16]. The country fill color method colors the interior of each polygon using OpenGL’s triangle rasterizer.

3.3 Glyph Fill Color

As part of Tart’s initialization, the glyph geometry is tessellated using a standard algorithm [16]. The glyph fill color method colors the interior of each glyph using OpenGL’s triangle rasterizer.

3.4 Glyph Placement and Size

The fourth method places a circle (i.e., glyph) with size proportional to a socio-economic indicator. Although drawing the glyph is trivial, finding the ideal location to draw it is more challenging. It is difficult to ensure that users will correctly identify a glyph with its corresponding country and, depending on the current zoom level, some polygons may be obscured by their own glyph. This problem is further complicated by the fact that some countries, such as the United States and Indonesia, are represented by multiple polygons.

The glyph placement algorithm is composed of two steps, as shown in Figure 3. First, the axis-aligned bounding box (AABB) and centroid for each polygon is calculated. Then for each country, a subset of polygons is selected using the following criterion: a selected polygon has the largest AABB from the set of polygons that are either within a centroidal radius of 15 decimal degrees (an aesthetically-determined distance) or have intersecting AABB. Glyphs are drawn at the centroid of each polygon in the aforementioned subset.

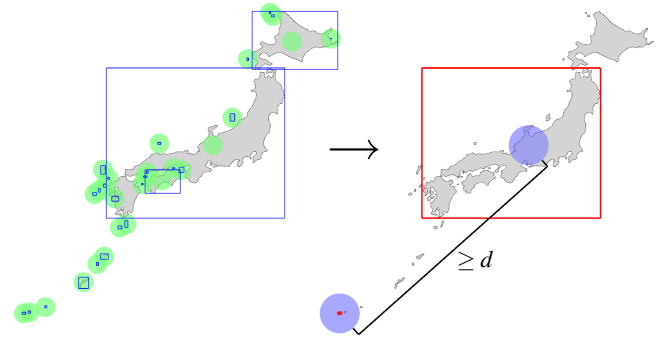


Fig. 3. This pair of Japan images illustrates the glyph placement algorithm. Let S (\square) be the set of all polygons in a country. Let the aesthetically-determined distance between two polygon centroids $d = 15$ decimal degrees. Let P (\square) be the set described by the relation $P \subset S \mid \forall q, r \in P : q \neq r$ the following holds: q_{AABB} does not intersect r_{AABB} and the distance between $q_{centroid}$ and $r_{centroid} \geq d$. Glyphs (\bullet) are drawn at the centroid (\bullet) of each polygon in P .

3.5 Time

A time slider at the bottom allows one to see how these variables change over time. Years with missing data are filled by linearly sampling between the nearest-known data points on either end. When missing data cannot be sampled, the associated visualization method is disabled for that year. Linear interpolation between bins for color and size can be collectively toggled on or off. This is most useful when creating smooth time-series animations with the time slider.

3.6 Cartograms

Cartogram support was added by integrating Michael Gastner’s C implementation [7] of the Gastner and Newman diffusion method [8] into Tart’s processing pipeline. Tart sends input geometry and indicator data to the cartogram code via in-memory data streams. The computed displacement field is then used to translate all of the world’s vertices. Finally, the characteristics of the distorted polygons – centroids, axis-aligned bounding boxes, vertex normals, and glyph locations – are recomputed.

On a modern CPU, computing a cartogram with Gastner’s implementation generally takes at least thirty seconds. Because of this limitation, cartograms are not automatically synchronized with the time slider and socio-economic indicator; the user must manually instruct Tart to recompute a cartogram for the currently selected time and indicator. Preliminary attempts to reduce the computation time by porting the algorithm to CUDA have shown promise and it may be possible to develop a comparable implementation that runs in real-time on a modern GPU.

4 REAL-TIME FLEXIBLE VISUALIZATION SYSTEM

4.1 Real-Time Image Generation

To achieve WYSIWYG results at interactive frame rates, Tart strikes a balance between one-time, static initialization and dynamic rendering. On startup, the Earth image is loaded into texture memory and the associated triangle mesh is generated. This is followed by generation of the glyph geometry. Then, the country geometry, which consists of vertices grouped by polygon and country, is loaded. Several important values are calculated during this time, including axis-aligned bounding boxes, vertex normals, centroids, and triangle meshes. Finally, the country indicators are loaded, during which min/max values are calculated and missing data is filled in by linearly sampling the neighboring known values. The remainder of the image generation follows a dynamic rendering pipeline, which is controlled by various user options, and is executed whenever the window is invalidated. This dynamic pipeline allows functionality such as zooming, full screen anti-aliasing, and the boundary visualization to provide vector-quality output.

4.2 I/O Control

Flexible I/O control is one advantage that Tart has over popular web-based solutions such as Gapminder World 2006 [6]. In Tart, an arbitrary geometry file and a cross-referenced database file that have been properly formatted are read in from the local filesystem and within seconds the data is ready to be analyzed. No web-based system that we know of allows the user specify the geometry on which the data is visualized. Furthermore, the set of indicators that can be visualized is typically also fixed. Tart also allows the user to capture high resolution PNGs of the current visualization that are devoid of unnecessary UI controls and needn't be restricted to the user's screen resolution. Web-based solutions generally do not provide any output options beyond links or exported data tables. Clever users can resize the browser, capture the screen, and finally crop/edit the image to remove the UI, but this is clearly more effort and the final image is of limited use in published works.

4.3 User Interface

Tart provides basic GIS functionality such as panning and zooming via the mouse. Users map socio-economic indicators to visualization methods by cycling through them with specific keys on the keyboard. Figure 4 shows an example of mapping youth female illiteracy rate using country fill color (in this case, with the use of { and }). Each visualization method has several relevant options to customize the final image. The four legend-based methods – country fill color, country boundary color, glyph fill color, and glyph size – each have their values separated into one of five uniformly sized bins. For each color-based visualization, users can select one of eight preset color legends. They can also adjust the color opacity, which is particularly useful in multivariate visualization for blending colors one on top of other as shown later in Figure 5. The glyph size legend can be inverted to place emphasis on either small or large glyphs.




Fig. 4. The visualization control allows one to bind a variable to a visualization method, select an appropriate coloring scheme for the range bins, and adjust the color opacity.

4.4 Underlying Technology

Tart is written in C++ and relies on several statically linked third-party libraries, in addition to an internally developed application framework. Boost libraries are used extensively to support such issues as exception safety, dynamic memory management, and type-safe callbacks.

FreeImage is used to load the Blue Marble Next Generation image and to save output images. The polygon tessellation is handled by a small class library available from FlipCode [16]. Cartograms are generated by a modified implementation of Gastner's C code [7].

5 APPLICATIONS

We present two examples of application of Tart.

Example 1: Life Expectancy, Birth Rate, Contraceptive Prevalence, and Female Illiteracy

We have applied our visualization system Tart to investigate relationship between female illiteracy rate, birth rate, life expectancy at birth, and contraceptive prevalence. The data is drawn from World Health Organization (WHO). Figure 5 displays the mapping of these four variables onto country fill color, glyph fill color, glyph size, and country boundary color respectively for the year 1995.

Please note that although these variables are mapped to rainbow coloring, the mapping is not necessarily one where high values are shown in red. Rather, value mappings are chosen so that the expected relationships between different variables coincide with the relationships between the colors. For example, in Figure 5, *high* illiteracy rate and *high* birth rates are shown in red while *low* contraceptive prevalence is shown red. This makes it easier to explore the expected relationship between these variables.

Furthermore, opacity is used at different levels to discern the three pseudo-color mappings. The country fill color is set to low opacity so that glyphs show up clearly. It also helps to avoid excessive contrast with the predominantly white background. The glyph color is set at medium opacity to differentiate it from the country fill color without completely occluding the information underneath. This is particularly useful for visually preserving black political boundaries in areas with small polygons. The country boundary color is drawn fully opaque so that boundaries may be easily identified.

Figure 5 highlights the observations on Niger, and the pink-red hue confirms strong relationship between these variables for Niger in 1995. Interestingly, many countries display these color similarities. India has the green hue for the three variables and mid-level glyph size again confirming the relationship between all these four variables. China has the blue hue indicating low illiteracy rate, high contraceptive prevalence, but slightly higher birth rate. There are several countries with yellow hue for all the variables confirming this relationship further.

This display also brings out those countries where this relationship does not exist or is not strong. An example is the West African country, Mali, where the red boundary and the red glyph displays low contraceptive prevalence and high birth rate, and the small glyph size indicates low life expectancy but the green country fill color shows that the illiteracy rate is medium in comparison to other countries. This observation leads one to investigate further the national policies of the country on accessibility and education on contraceptive prevalence and health policies.

Example 2: Life Expectancy, GDP, DPT Immunization and Physicians

In our second example, we have drawn the data from the Global Health Nexus Database (GHND) [17]. This is a freely available, highly reliable database that has recently been used to study impact of several independent variables on global health indicators [2]. The study examines five different pathways for impacting health and uses a fixed effect model to estimate their impact on health indicators for all countries. Health indicators included in this study are life expectancy at birth (LEB), under 5 mortality rate (u5MR), and infant mortality (IMR) over a period of more than 40 years from 1960-2005. Several independent variables are considered including DPT immunization rate and Log (Physicians/1000 people)/ Log (GDP per capita), the two that we consider in our example. The reason for considering the modified latter variable rather than simple number of physicians is to obtain an index of availability of health personnel relative to the GDP/c norm. The study also estimated the impact of the independent variables on the health indicators for four groupings of countries – high income, middle income, low income, and transitional countries – based on GDP. Dependence of LEB and u5MR on DPT immunization

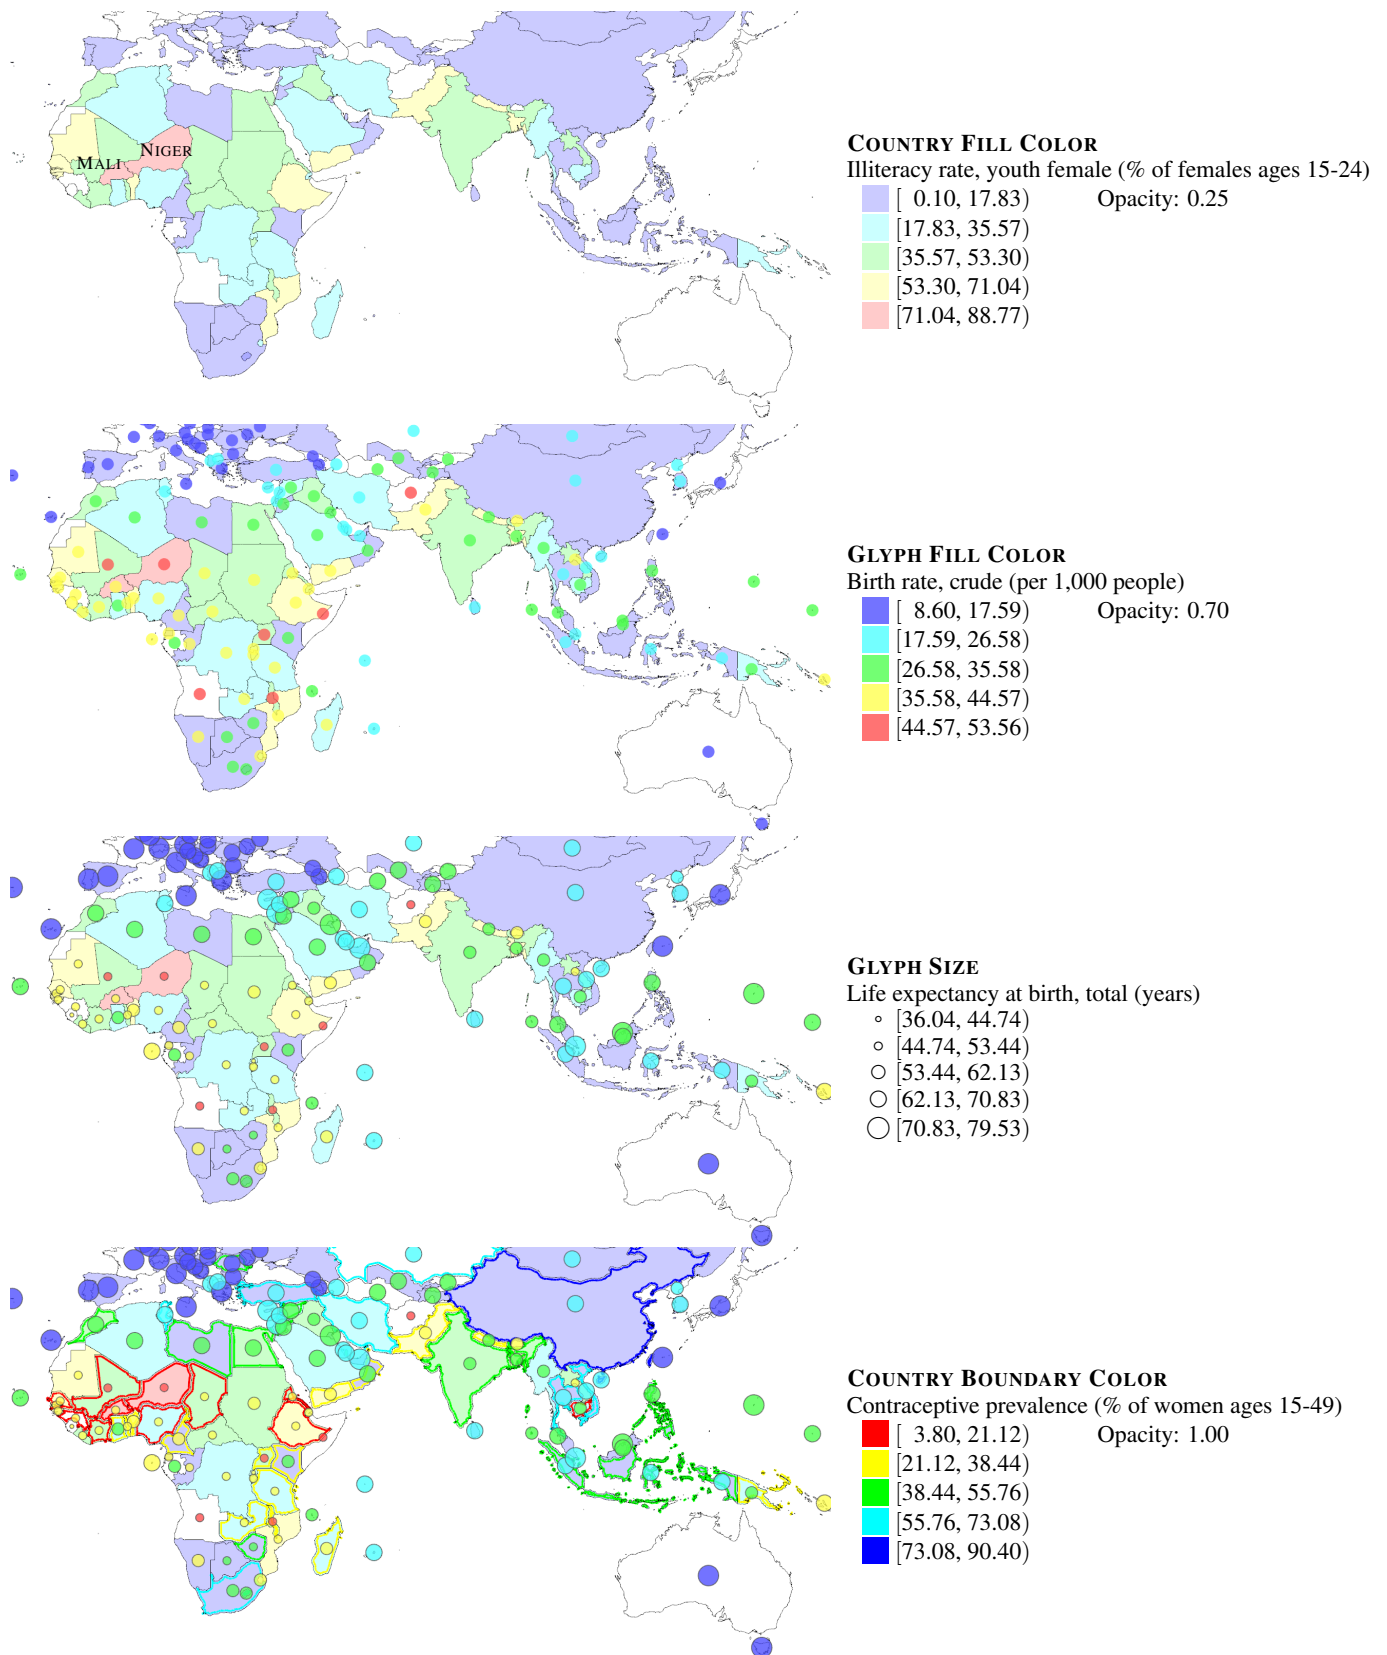


Fig. 5. This sequence of images, representing WHO data from 1995, illustrates the benefits of visualizing one to four variables simultaneously. By manipulating the legends so that red and small circles indicate problems areas, we can quickly isolate countries and identify correlations between variables. In the first image, the country fill color visualization reveals that Niger has a high youth female illiteracy rate. From the glyph fill color visualization in the second image, one can also see that Niger has a high birth rate. In the third image, the glyph size visualization further highlights Niger as having low life expectancy at birth and provides some insight regarding its high birth rate. And finally, the country boundary color visualization in the last image shows low contraceptive prevalence in Niger — additional rationalization for its high birth rate.

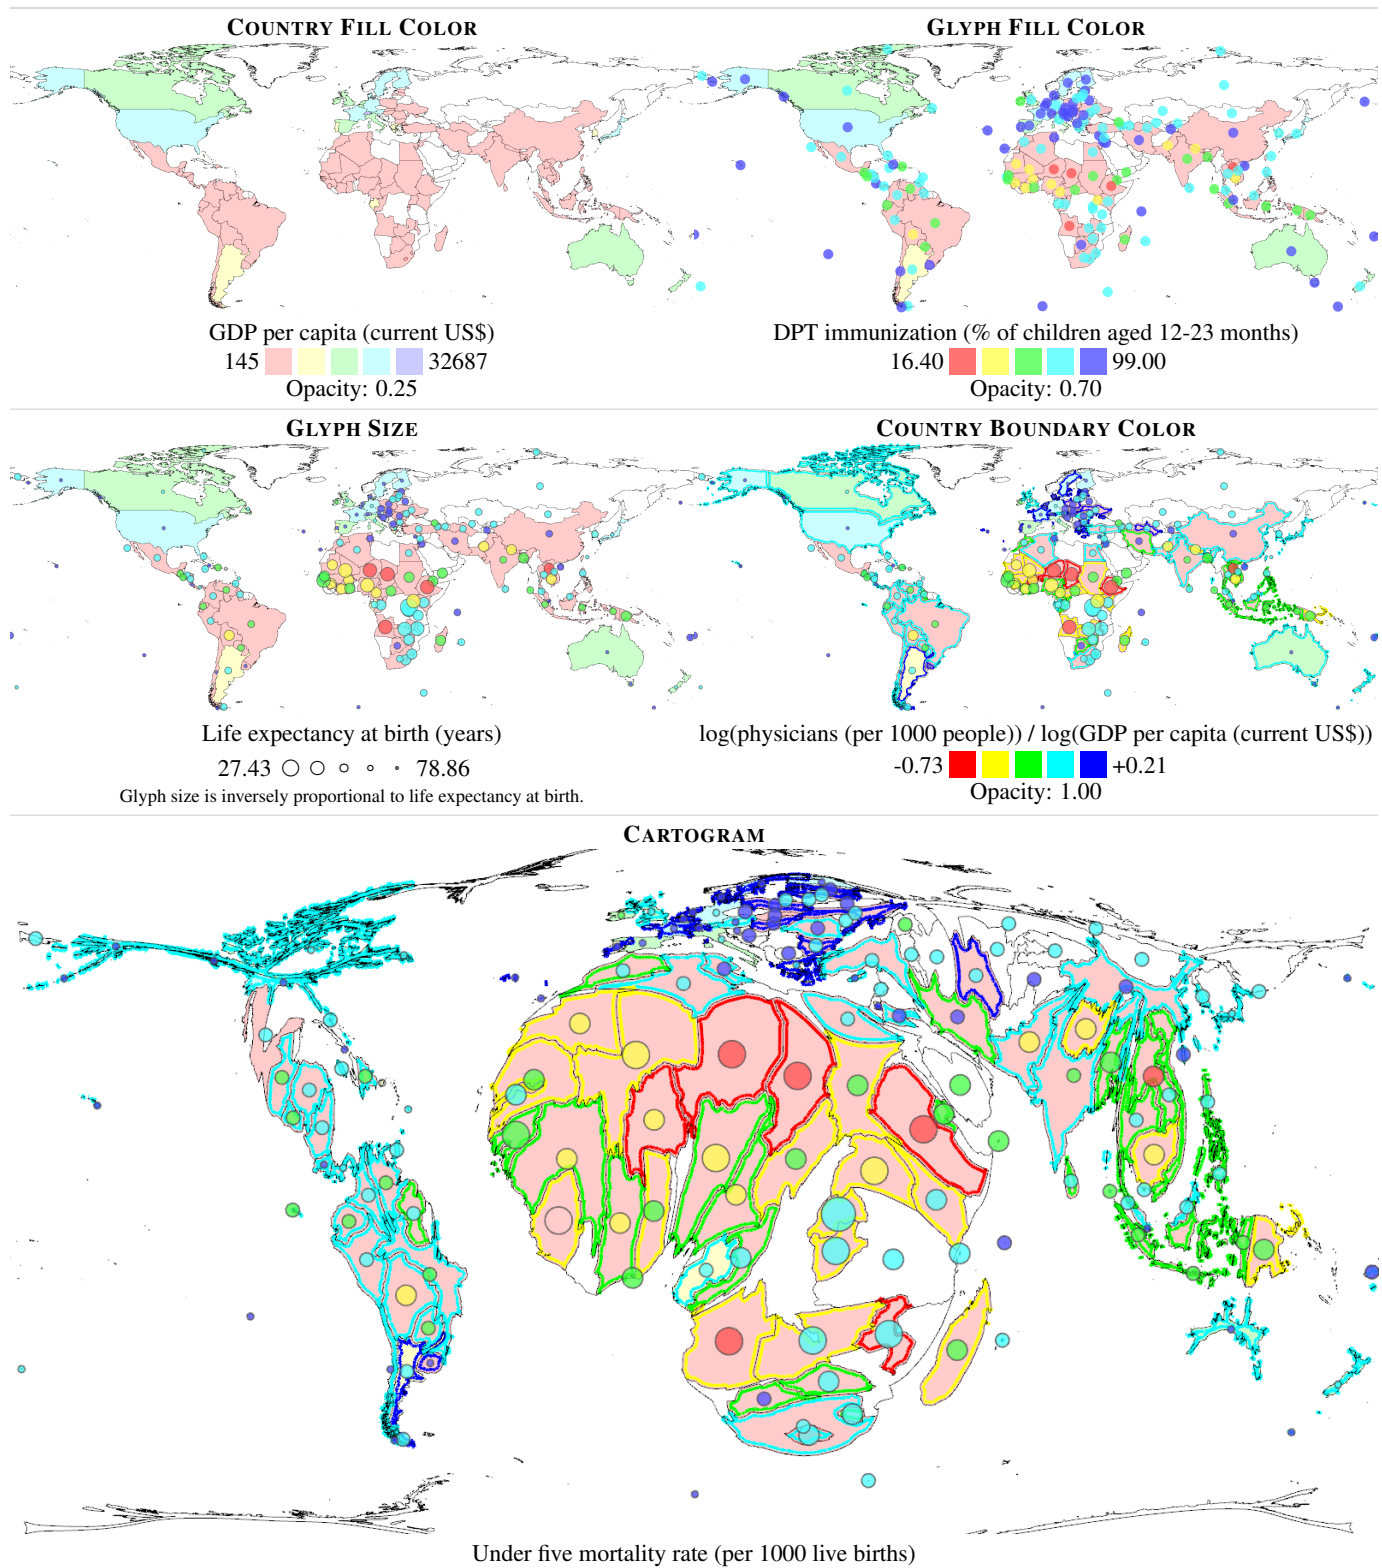


Fig. 6. This sequence of images, representing GHND data from 1990, illustrates the benefits of visualizing one to five variables simultaneously. By manipulating the legends so that red, large circles, and relatively larger regions indicate problems areas, we can quickly isolate countries and identify correlations between variables.

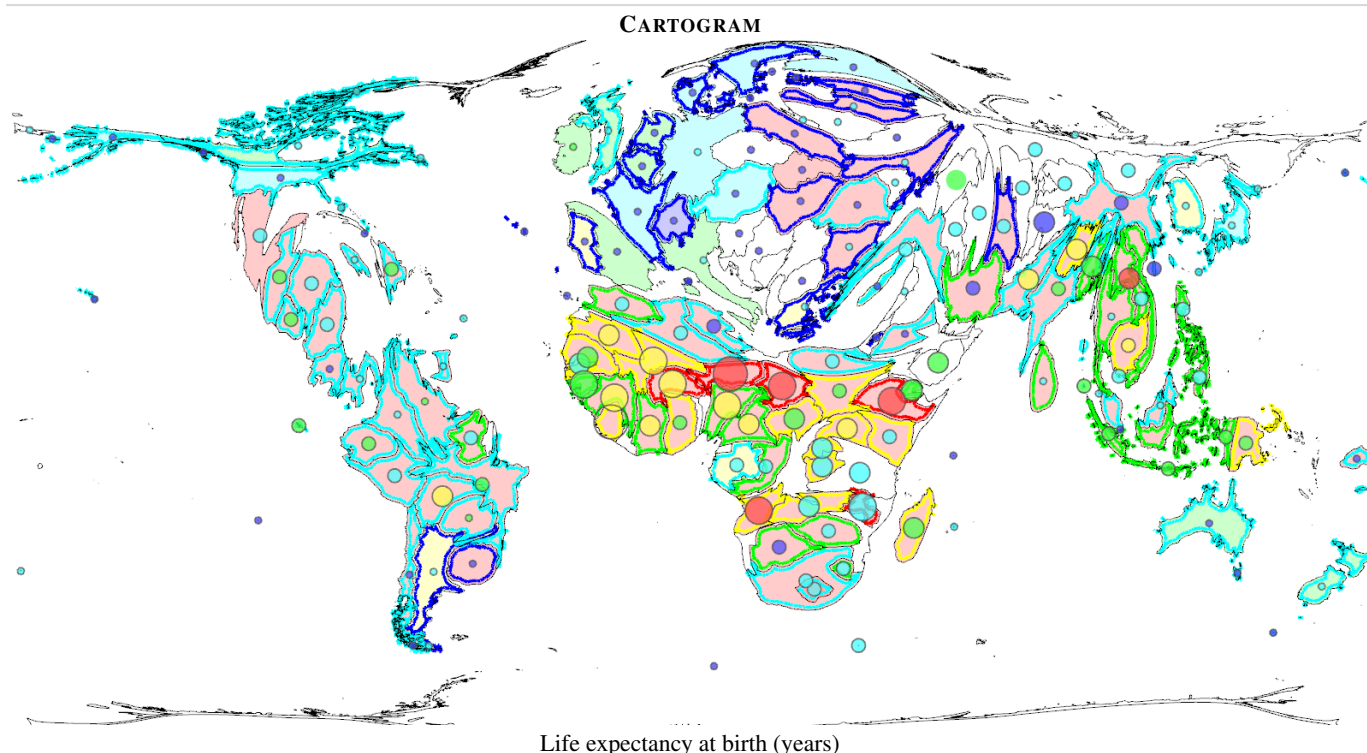


Fig. 7. This cartogram is also generated with GHND data from 1990 where the mappings of GDP, DPT immunization, and physicians variable is the same as in Figure 6. Mapping of glyph size and cartogram is switched from the previous figure to u5MR and LEB respectively. u5MR is mapped proportional to the glyph size.

and physicians variables is found to be statistically significant at either 1% or between 1 to 5%.

The study firmly establishes the impact of DPT immunization on reduction of u5MR and increase in LEB. However, the verdict on the impact of number of physicians on health indicators is somewhat mixed because although statistically significant, the impact of increasing the physicians variable on LEB is *negative* for high income countries, a surprising result. One plausible explanation for this result is that the increase in LEB accounted for by other variables was already so high that it needed to be mitigated by this variable in this group model. We chose to investigate the relationship between these variables due to these interesting and surprising results obtained in this study.

Figure 6 displays the results of mapping GDP, DPT immunization, life expectancy, physicians, and under 5 mortality variables to country fill color, glyph fill color, glyph size, country boundary color and cartogram distortion respectively. In this example, low values of GDP, DPT immunization, and physicians variables are all mapped to the red end of the color spectrum. However, glyph size is mapped inversely proportional to the life expectancy in order to keep the maps uncluttered and allowing the users to focus on countries with low LEB. Similarly, countries with high u5MR are blown up in the cartogram. As in example 1, low opacity for country fill color, intermediate opacity for the glyph color, and high opacity for the boundary color works well.

We first observe that most high income countries including USA and several western European countries have consistently bluish hue and small circles indicating a strong relationship between high GDP, high DPT, high physicians, high LEB, and small u5MR, as is apparent in their almost disappearance from the cartogram. This observation points that in spite of the *negative* causal relationship found between physicians and LEB in the study by Cornia *et al.* [2] for high income countries, there is a strong *positive* correlation between these variables.

Strong relationship between these five variables is also clearly displayed by several African countries that appear very large in the cartogram with strong red-yellow hues indicating high u5MR, low LEB, low GDP, low DPT, and low physicians.

Figure 6 also allows quick comparisons between emerging countries. For example, India is doing better than China in terms of availability of physicians relative to GDP and LEB (smaller circle), but is doing worse than China in u5MR, DPT immunization, and GDP. This observation leads to examine the national policies of India and China where India could increase DPT immunization while China could work towards increasing accessibility to physicians to improve their health indicators.

Due to enormous distortion that can take place in a cartogram, special caution is needed to map the variables appropriately to this variable. Figure 7 shows mapping of the same five variables except that the mapping of glyph size and cartogram is swapped to u5MR and LEB respectively. This cartogram shows an interesting contrast to the cartogram in Figure 6 clearly bringing out the fact that u5MR is high in Eastern block countries in Europe, a fact that was not so obvious from the earlier figure. This observation points to a need for formulation of health policies designed to curb u5MR for these countries.

6 CONCLUSIONS AND FUTURE WORK

We have presented a multivariate visualization system for displaying five socio-economic variables simultaneously on a world map. Using examples from two different databases, we have demonstrated that the system can be used to provide valuable insights in many cases not easily available otherwise. Our visualization system supports real-time animation, provides good I/O control, and a flexible user interface to allow easy investigation of relationships between indicators and produce high quality maps. We expect to explore multivariate spatio-temporal relationships between variables using a combination of animation, principles of small multiples, and novel techniques for static spatio-temporal visualization including rings [19]. Another future direction of research is to create multivariate display of geographic data that allows visualization of raw data along with computed data such as those obtained from statistical or multi-regression causal models or projected data based on proposed national policies. We believe that these multivariate displays will allow easy comparison between indi-

cators and countries that can be a helpful aid in formulating and assessing national and global priorities.

ACKNOWLEDGEMENTS

The authors wish to thank A, B, C. This work was supported in part by a grant from XYZ.

REFERENCES

- [1] W. Bank. World development indicators. Website, 2008. <http://www.worldbank.org/data/>.
- [2] G. A. Cornia, S. Rosignoli, and L. Tiberti. Globalisation and health: impact pathways and recent evidence. In *Proceedings of Conference on Mapping Global Inequality*, 2007.
- [3] D. Dorling, A. Barford, and M. Newman. Worldmapper: The world as you've never seen it before. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):757–764, 2006.
- [4] J. Dykes and C. Brunsdon. Geographically weighted visualization: Interactive graphics for scale varying exploratory analysis. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1161–1168, 2007.
- [5] U. I. for Statistics. Global statistics. Website, 2008. <http://www.uis.unesco.org>.
- [6] Gapminder. Gapminder world 2006. Website, March 2008. <http://tools.google.com/gapminder>.
- [7] M. T. Gastner. Cartogram C code. Website, March 2005. <http://www.santafe.edu/~mgastner/>.
- [8] M. T. Gastner and M. E. J. Newman. Diffusion-based method for producing density-equalizing maps. *Proceedings of the National Academy of Sciences*, 101(20):7499–7504, 2004.
- [9] D. Guo, M. Gahegan, A. M. MacEachren, and B. Zhou. Multivariate Analysis and Geovisualization with an Integrated Geographic Knowledge Discovery Approach. *Cartography and Geographic Information Science*, 32(2):113–133, 2005.
- [10] H. Hagh-Shenas, S. Kim, V. Interrante, and C. Healey. Weaving versus blending: a quantitative assessment of the information carrying capacities of two alternative methods for conveying multivariate data with color. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1270–1277, 2007.
- [11] C. G. Healey and J. T. Enns. Large datasets at a glance: Combining textures and colors in scientific visualization. *IEEE Transactions on Visualization and Computer Graphics*, 5(2):145–167, 1999.
- [12] V. Interrante. Harnessing natural textures for multivariate visualization. *IEEE Computer Graphics and Applications*, pages 6–11, November–December 2000.
- [13] Nationmaster. Nations of the world. Website, 2008. <http://www.nationmaster.com/>.
- [14] OECD. Organisation for economic co-operation and development. Website, 2008. <http://www.oecd.org>.
- [15] M.-C. Panse, M.-M. Sips, M.-D. Keim, and S. M.-S. North. Visualization of geo-spatial point sets via global shape transformation and local pixel placement. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):749–756, 2006.
- [16] J. W. Ratcliff. Efficient polygon triangulation. Website, July 2000. <http://www.flipcode.com/archives/Efficient.Polygon.Triangulation.shtml>.
- [17] S. Rosignoli, L. Tiberti, and G. A. Cornia. The globalization-health nexus database (ghnd). Website, February 2007. <http://www.unifi.it/dpssec/sviluppo/database.html>.
- [18] T. Saito, H. N. Miyamura, M. Yamamoto, H. Saito, Y. Hoshiya, and T. Kaseda. Two-tone pseudo coloring: Compact visualization for one-dimensional data. *Proceedings of Information Visualization*, 0, 2005.
- [19] P. Shanbhag, P. Rheingans, and M. desJardins. Temporal visualization of planning polygons for efficient partitioning of geo-spatial data. *IEEE Symposium on Information Visualization*, 2005.
- [20] UCSC. UC atlas. Website, 2008. <http://ucatlas.ucsc.edu/>.
- [21] C. University and W. Bank. Global poverty mapping project. Website, 2008. <http://www.cisein.org/povmap/atlas.html>.
- [22] UNSCB. United nations common database. Website, 2008. http://unstats.un.org/unsd/cdb/cdb_help/cdb_quick_start.asp.
- [23] J. Wood, J. Dykes, A. Slingsby, and K. Clarke. Interactive visual exploration of a large spatio-temporal dataset: Reflections on a geovisualization mashup. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1176–1183, 2007.