

Visualising official statistics

Olav ten Bosch and Edwin de Jonge

Statistics Netherlands, Henri Faasdreef 312, 2492 JP The Hague, The Netherlands

Tel.: +31 70 3375025; Fax: +31 70 3877429; E-mail: o.tenbosch@cbs.nl/e.dejonge@cbs.nl

Abstract. Data are everywhere: on the internet, in newspapers, in computer games, at school, at work and on TV. Data are essential to our daily information needs and people are increasingly becoming used to reading and interpreting data. Companies use advanced technologies to present their data in fancy and visually attractive ways. Animated and interactive graphs, mapping tools, charting components are now common techniques, and many examples of attractive and well-designed data visualisation are to be found on the internet in particular. Many of these concepts can be used to visualise official statistics. Indeed, we think people expect official statistics to follow these trends so that they can use the same tools and concepts to consult official statistics. This paper presents an overview of visualisation trends on the internet. It highlights visualisation initiatives that can currently be found on the web and are likely to be useful for official statistics. It also describes the status of visualisation activities at Statistics Netherlands, in the field of regional statistics as well as other more specific data visualisations, and touches on some of the more experimental features currently being developed.

Keywords: Statistics, data visualisation, dissemination, internet trends, animation, charts, maps

1. Introduction

The speed at which statistics are flooding our life seems to increase every year, as if our hunger for information cannot be satisfied. These numbers are increasingly presented in a visual way. Take sports statistics for example, fifteen years ago a television presenter would just talk to viewers, giving the most interesting background information verbally. Today viewers can read the latest match statistics and details about earlier achievements of players at the bottom of their screen. During a soccer match, the numbers of passes, shots at goal, yellow and red cards and substitutions are displayed on screen. And many more live statistics about other matches can be found on the accompanying website. A similar information explosion has affected our telephone bills. Twenty years ago telephone bills used to be a sheet of paper specifying the total amount of money to be paid for phone calls in the previous two months. Nowadays, people have access to their personalised billing page via the internet, where each call is specified to the millisecond, colour-coded for the phone rate applicable at that moment. Going even further, phone companies and some other independent websites present visual advice on op-

timum spending patterns. Airline company websites help people choose the cheapest available flight using an easy to grasp colouring scheme based on date, time, fare and availability. People are becoming increasingly used to visual representations of information. This is especially true for young people: children play computer games intuitively, and are not even aware of the visual feedback they get.

One important enabler for the large increase in graphical presentations of everyday information is the rise of graphical technologies on the internet in combination with the smooth integration of websites with their information systems. It is easy for companies to present information from their backend systems in real-time to the public. The resulting information needs to be understandable for users. Websites have evolved from mainly static communication channels into more mature media-based output media with all the possibilities for visual interactivity. Therefore websites are increasingly becoming the primary dissemination channel for organisations. Over the years most websites have evolved from simple pages with text to interactive visualisations presenting information attractively.

Official statistics have a much longer history of public dissemination compared with the existence of the

internet. Official statistics have been around for more than one hundred years, and for most of that time have been published on paper. The most common publication form was a statistical yearbook, full of tables with reliable, undisputed and readable material. Although organisations publishing official statistics have made the transition to the web, their tradition of table-based dissemination is still very visible. In many cases statistical publications look like a more or less direct translation from paper to web. For the last ten years, national statistical institutes (NSIs) have been exploring modern internet techniques to publish statistics on the web [9]. We think this is a very good direction and believe that NSIs should continue to explore the visualisation possibilities to improve the presentation of their output to the public. In the past few years Statistics Netherlands has started activities in this direction, and our experiences are described in this paper. It presents an overview of the state of visualisation and trends on the internet. It further highlights some excellent visualisation initiatives to be found on the web these days that may well be applicable in the field of official statistics. It also describes some of the activities at Statistics Netherlands and touches on some of the more experimental features currently being developed.

2. Visualisation of statistical data

Publishing statistics is not just about supplying numbers; how the numbers are presented is important. Tables are common and often used to present data, but tables have weaknesses alongside their strengths. They are very useful for looking up individual values of data, and finding out precise values. They can contain a huge amount of information and can be interpreted easily. However, trends, value distributions and correlations are not easy to detect in tabular data. Visualisations like charts are often more helpful for this purpose. A good visualisation utilises the human visual system to highlight patterns in the data. These patterns can be outliers, trends or the distribution of the data.

Traditionally, visualisations have been used in statistics for two purposes: *analysis* and *communication*. In the case of analysis, visualisation is used to assist the statistical analyst to detect relevant patterns in his data. Visualisation techniques for analysis typically give the analyst many options to explore and visualise the data. The analyst adjusts the options in the visualisation until he finds the data characteristic he is looking for and

completes his analysis. Detailed descriptions and use of these techniques can be found in [5,6].

In the case of communication, visualisation is used to tell the story or show the pattern that has already been discovered in the data. This type of visualisation tries to highlight the relevant parts of a dataset to users in such a way that they immediately get “the message”. Visualisation techniques for communication therefore focus on design: it must be an honest, unambiguous and effective presentation of (a pattern in) the data. Detailed design principles for communication (charts and tables) can be found in [7,10].

Official statistics use visualisations for both purposes. Analysis of statistical data is part of the core business of statistics offices and visual analysis is one way to do this. In the dissemination of official statistics visualisation is mostly used for communication: press releases and web articles often are accompanied by a chart or map to tell the story more effectively. However, as we shall see later in this paper, visualisation for analysis can increasingly be found on official statistics websites.

One very important type of visualisation is thematic cartography. In fact, historically, cartography was the mother of most data visualisations as we know them [7]: e.g. the concept of a graph with two axes is derived from the concept of a map. Since many important official statistics are regional statistics or have a regional breakdown, thematic cartography is an important means to visualise these kinds of data.

3. Trends in dissemination of statistics

In the following sections we describe some trends in visualisation we have found on the web; it is by no means intended to be an exhaustive list of current trends.

3.1. Dynamic charts

When internet started, most statistical charts found on the web were bitmaps. These resembled charts in paper publications, in that users could not interact with them. Over years websites made their charts more “dynamic”, by adding mouseover effects. Current technology makes it possible to let users interact with each graphical object in a chart. A dynamic chart embraces two different concepts: *interactivity* and *animation*. Interaction means that users have some control over the visualisation: for example a user can click on a data

point to get extra information about that point. The chart comes to life more because it reacts to users' actions. Animation means that data changes are shown as an animated picture. In many dynamic charts interaction results in animation. Animation is a powerful visualisation technique that should be used with caution. Examples of good animations are those of time series that show the evolution of data in time. A good accompanying user control is the play and rewind button, by which users can control the animation. However, many examples of animation can be found that do not add any value to the visualisation. They distract the attention of the users. In our opinion, animation should be used only when it is meaningful and/or functional.

3.2. Mapping/geotagging tools

Today many people use map-based applications for common tasks, such as checking for traffic jams or local weather forecasts. Even if they want to buy a used car, a house or book a meal in the nearest good restaurant, people in the Netherlands can choose between a number of popular map-driven websites. In general, then, people are familiar with mapping engines used by popular Dutch websites for everyday tasks. Usually, these sites use technologies such as Google maps [13] or the more advanced Google Earth [12] to display their information. The maps presented are created from satellite photos or aerial photos that display the area in more or less detail. Because of the huge number of photos needed to display the almost infinite number of possible maps, these services use streaming technologies, where only the appropriate photos are sent to the user and put together on the client's computer. We call this type of map generation *streaming photo-based mapping* technology. Users can also add their own layers of information; they can annotate or "tag" geographic objects or coordinates with extra information and photos. Examples of the use of Google Earth for statistical information are the *Google Earth layer* and *Statistics Netherlands in your Neighbourhood* which will be presented later on in this paper.

In addition to these easy to use *streaming photo-based mapping* tools, statistical organisations also use more professional Geographical Information Systems (GIS). These systems were already being used in statistics offices, as huge amounts of statistical information were usually available at a detailed regional level and the offices needed powerful geographical systems to analyse them. Traditionally, these regional statistics were presented in the form of tables and maps. This

works reasonably well for people who know their way around these kinds of mapping modules: they typically have the skills and patience to browse the datasets and generate the maps of their choice. However, the average internet user required an easier and more agile user interface. And that is what he got: the user interfaces of the GIS tools are becoming simpler to use and are starting to use Google Earth defaults. It should be noted that the design of the web interfaces on GIS tools is influenced by Google tools. ESRI [21] and GeoClip [19,22] are examples of other GIS systems that have a simplified web based user interface.

3.3. Bubble charts

Recently there has been a lot of interest in the so-called bubble charts, popularised by Hans Rosling of Gapminder [18]. Bubble charts are scatter plots where the area of the dot is used to plot a third variable. Bubble charts are used to present multivariate data with three or more variables. Bubble charts are aesthetically pleasing. When combined with animation, in particular, bubble charts are an excellent tool for communicating statistics. Animated bubble charts, also known as motion charts, can be used to annotate a statistical trend as the moving and changing bubble sizes "tell a story". However, they also have a disadvantage. They may suffer from occlusion: for datasets with approximate data points, bubbles overlap or hide other bubbles, thereby making the visualisation less reliable and less clear. This is almost always the case for large datasets. Attractive examples of bubble charts can be found at Gapminder [18], the OECD Explorer [17] and the IMF data mapper [16].

3.4. "Analysis to the people"

As stated above, most visualisations developed today are dynamic. Adding interaction to a visualisation may result in applications that are more analytical by their nature than communicative. Instead of just "reading a story" through the visualisation, users are invited to detect patterns themselves by interacting with the visualisation. An analytic visualisation typically contains a large number of user controls and works well for statistically literate users. An example of this approach is the OECD Explorer [17]. Using this well-designed tool, people can browse through a dataset of regional statistics via a map and a bubble chart. It also contains a parallel coordinate chart. This type of chart is very useful for analysing multivariate data, but the

average user is not familiar with it and needs to learn its concepts before he can use it effectively [8].

Applications that are more analytical than communicative may be very powerful for advanced and skilled users. They provide them with the means to analyse official statistics, thus democratising the analysis of statistical output. However, a couple of remarks should be made in this respect. Publishing data using an analysis tool does not relieve the NSI from its duty to tell an important story with the data: just as newspapers filter and highlight the news, so NSIs should still highlight important statistical patterns themselves. Furthermore, even with an excellent visual analysis tool a user might create a misleading visualisation and make a wrong analysis. Of course, that is the responsibility of the user himself. So although additional usability testing on the actual use of these applications would be useful, we think that applications that give analytic power to the users are a very promising direction.

3.5. Generic versus tailor-made visualisations

In general we can distinguish two types of data visualisation in statistics: generally applicable methods and tailor-made approaches. In the first case the visualisation algorithm is generally applicable to statistical data, whatever the topic: for example, the motion charts in Google charts and in the data mapper of IMF. They can be used for any type of multidimensional data and the user has to specify which data are presented on which visualisation axis (x , y , bubble size, colour). They are optimised for time series, using a play facility. A tailor-made visualisation approach has a different starting point. Based on a specific dataset, or the metadata of a specific dataset, a corresponding visualisation is designed. A well-known example of this approach is a population pyramid, presenting the characteristics of a population (number of people by age by sex). Since the gender axis has only two categories, the data can be presented in a pyramid. But this presentation method is suited for this kind of data only. Data that have more categories in one of their dimensions could not be presented in such a way. For example, to visualise the number of people by age and region we would need another type of visual representation. Nevertheless the population pyramid is very understandable for this specific dataset, which stresses the power of a tailor-made visualisation design.

3.6. Zoomable User Interfaces (ZUI)

Another promising and interesting type of visualisation is the *Zoomable User Interface (ZUI)* [1,2]. The ZUI shows data at a certain level of detail and users can zoom into part of the chart to show these data in greater detail. ZUIs are useful for data with different levels of detail: they allow users to magnify parts of the data or, vice versa, hide details of the data. An example of a ZUI is Google Finance [14], where users can choose the level of detail in time of stock exchange rates. ZUIs may also implement *semantic zooming*: different representations of the data depending on the level of detail. Examples of ZUIs with semantic zooming are Google Earth and Google Maps: depending on the level of detail, roads and highways are hidden or shown. Mapping tools are a specific form of ZUI where the zoomable dimension is geography. Many official statistics use hierarchical classifications and have data at different levels of detail. ZUIs are very interesting for these types of statistics. There are not many examples of zoomable visualisations yet, but we think they are an interesting direction for visualisation and could be used more in the field of official statistics.

3.7. Tree-maps

Another example for the visualisation of hierarchical data are tree-maps. A tree-map is a method for displaying tree-structured data using nested rectangles. The area of each rectangle expresses the variable to be visualised. Colour schemes may be applied to display additional data characteristics. Tree-maps are increasingly to be found on the web. For example, tree-maps are one of the types of charts that can be created via the excellent visualisation site of IBM ManyEyes [15].

As most of the statistical classifications are hierarchical by nature, this kind of visualisation could be used for disseminating or analysing official statistics, as mentioned in [3]. However, up to now there have not been very many examples of the use of tree-maps for statistical data. This may be because they are relatively unknown, but it may also be because they are not intuitive and clear enough for inexperienced users. There seems to be some evidence for the latter. However, we still think that tree-maps might be useful for visualising certain hierarchically organised statistics, and hope that they will be tested further.

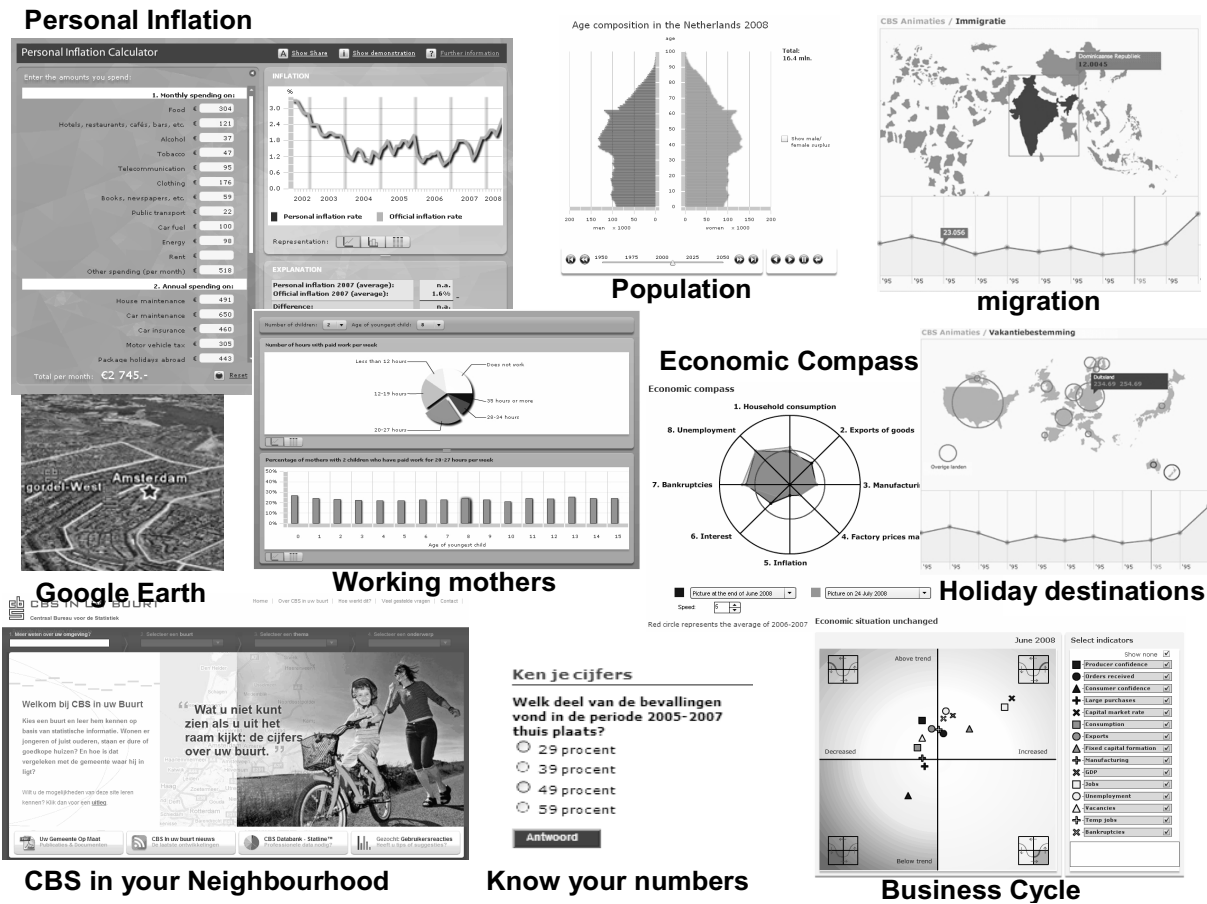


Fig. 1. Examples of alternative presentation methods of statistics.

3.8. Sparklines

Sparklines are data-intense, design-simple, word-sized graphics, invented by Edward Tufte [11]. These small charts or other graphical symbols can be a concise representation of some statistical trend, to be used within the body of a table or within blocks of text. Like tree-maps, the concept has not been used very much in the field of official statistics, but in principle sparklines could be the perfect way to present the main results of some important statistics very concisely, possibly integrated within a text or table.

4. The statistics Netherlands approach

For many years now, Statistics Netherlands has been disseminating statistical information via its main website www.cbs.nl. Press releases, short statistical articles in its “web magazine”, paper publications, data quizzes

(“know your numbers”) and scientific reports can all be found there. People looking for detailed information can browse through the huge statistics database StatLine. This has worked reasonably well for the more internet-aware and professional visitors. However, other user groups have not always been using the available information effectively. Therefore Statistics Netherlands has started to focus more on visualising statistics in more attractive and interactive ways than statistical tables. Figure 1 shows a number of examples of these alternative ways of presenting statistics. We shall now illustrate the trends described in the previous section with visualisations created by Statistics Netherlands.

4.1. Regional statistics

Statistics Netherlands holds large datasets at detailed regional levels. Users can create maps from these datasets from the statistics database of the Netherlands, StatLine. Figure 2 shows two examples of such maps.

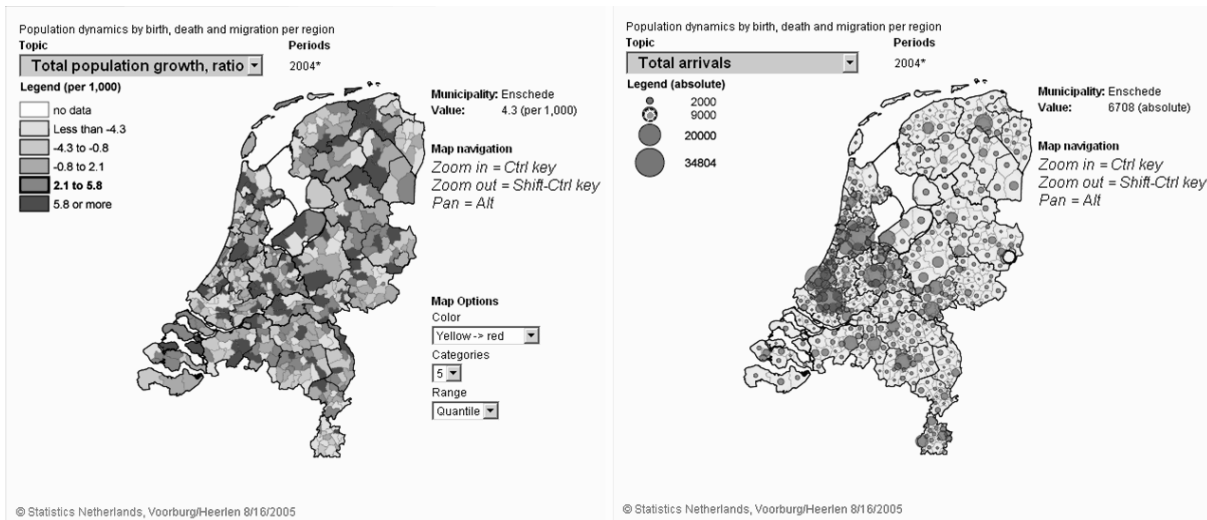


Fig. 2. Municipality maps generated by StatLine.

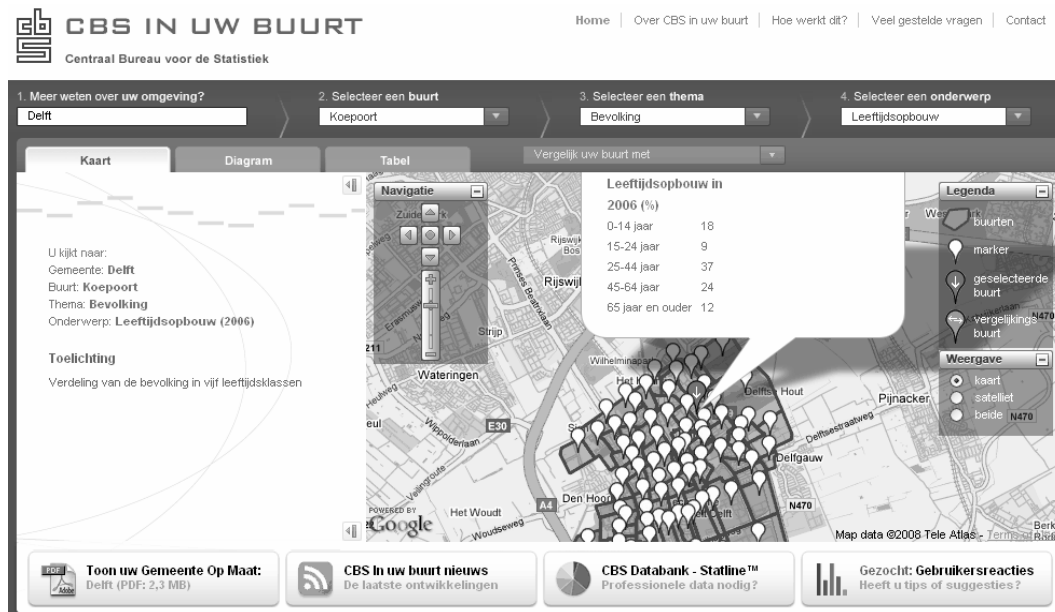


Fig. 3. Website CBS in your neighbourhood.

These maps show total population and total arrivals at municipality level. Although this type of mapping engine is useful for frequent visitors, generating such a map can be difficult for other users. First of all, not all data available are suitable for presentation in a map. Only data that use one of the standardised geographic dimension formats can be visualised in a map. Secondly, the mapping engine generates maps in scalable vector graphics (SVG). This technique is technically advanced, but most people do not have the appropri-

ate plug-in installed on their computer, which creates another threshold. A third disadvantage of these mapping engines is that users must tune the maps generated via a number of parameters. These parameters may complicate the mechanism for ordinary users. So, in practice there are a number of problems with this type of traditional mapping engines.

Since 2000 there has been a tremendous increase in the use of Google Earth and Google Maps mapping technology in the Netherlands. Combining the fact that

Statistics Netherlands has very many detailed data to neighbourhood level with the familiarity of the public with streaming photo-based mapping engines, the use of these technologies for statistical dissemination is a next obvious step.

In 2007, a subset of some 20 variables for about 11,000 neighbourhoods was selected from the statistics database StatLine. This resulted in a dataset ranging from population characteristics such as age, household composition, etc., to income and employment. These data were presented using streaming photo-based mapping technology. The resulting service is a combination of a website, called “CBS in your neighbourhood” (see <http://www.cbsinuwbuurt.nl> (in Dutch only)) and a statistical layer on Google Earth. The website went online on 12 February 2008, and attracted 20,000 visitors on its first day. A screen dump of the website is shown in Fig. 3. The Google Earth layer for this application was launched at the same time. This dual approach was necessary to combine ease of use for the public with powerful facilities such as combining statistical information with other layers. The website “CBS in your neighbourhood” serves the first use; the Google Earth layer facilitates more advanced use.

The statistics layer on Google Earth is offered to the public from the main website of Statistics Netherlands. Users must install Google Earth on their computer to use this layer. Google Earth can be downloaded from Google without charge. Once it is installed, selecting the link on the webpage of Statistics Netherlands will do the job. Google Earth will start up with detailed borders of all 11,000 neighbourhoods in the Netherlands, names of the neighbourhoods and a Statistics Netherlands (CBS) icon in each neighbourhood. This information becomes visible if the user zooms in to a specific region. An example is displayed in Fig. 4 on the left-hand side. The Statistics Netherlands icons are clickable, and when clicked will open to show all variables an info-window. The right-hand side of Fig. 4 gives an example for the neighbourhood “Zierikzee binnen de vesten” of municipality “Zierikzee”.

The layer Statistics Netherlands provided on Google Earth uses the underlying streaming technology of Google Earth. The data and borders to be displayed in the map are described in KML (keyhole markup language) [20]. The streaming mechanism is parameterised so that neighbourhood borders and statistical values are transferred “just in time”. Data and border coordinates are transferred only if the user zooms in on this area. For a more detailed explanation of the technology used we refer to [4]. The visualisation of

statistics on Google Earth is an example of the use of a zoomable user interface for official statistics.

The Google Earth layer currently published on the website of Statistics Netherlands is only a first step on the road of new ways of disseminating statistics. For example, we have noticed that the current layer makes it really easy to get a quick view of the characteristics of a specific neighbourhood, but it does not help users much to get a whole view of the status of a certain variable for a municipality, district or the Netherlands as a whole. This would require a more advanced mechanism possibly using heat maps to display the status of the Netherlands in one view. Early experiments with this facility are promising. Figure 5 gives an example of a heat map displaying the population density of the Netherlands on Google Earth.

4.2. Tailor-made visualisations

At Statistics Netherlands, certain subjects have recently been selected as candidates for visualisation. First of all, some well known examples from other statistical offices, like the “population pyramid” and various “personal inflation calculators” were studied and implemented accordingly (population pyramid) or converted to a Dutch equivalent (personal inflation calculator). Moreover, in dialogue with in-house statisticians certain datasets were selected from the statistics database and ideas were formulated for a visual presentation of these datasets. This resulted in a list of about 25 ideas for specific visualisations. Statistics Netherlands is currently implementing the first ideas. The common principle behind all these initiatives is to take a dataset from our overall statistics database StatLine as a starting point. Based on the dataset, a corresponding visual representation is developed in close cooperation with the statisticians involved.

An early example of this approach is the Dutch business cycle tracer [23]. The first edition was launched in 2005. It is a tool to assist in the analysis of the state and the course of the Dutch economy by visualising the cyclical nature of economic developments. The target audience is therefore mainly people professionally interested in economic developments, but the visualisation has also been used in education as part of Dutch economics exams. It is updated every month and annotated with an analysis of the state of the Dutch economy. Figure 6 shows the business cycle tracer reflecting the state of economy in October 2008. The business cycle tracer is clearly a tailor-made visualisation as it uses the inherent characteristic of the indicators (above

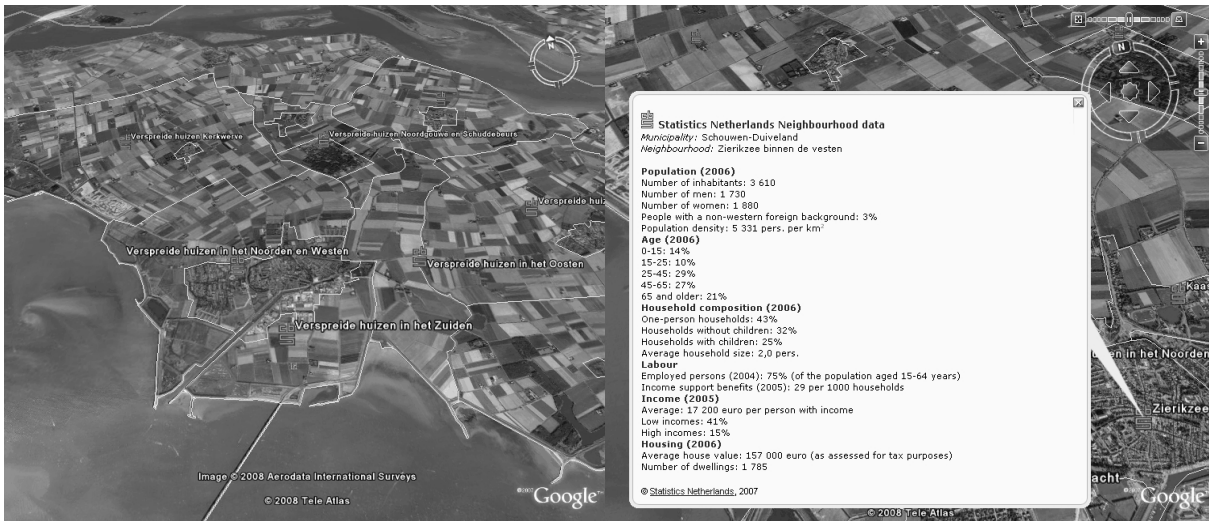


Fig. 4. Neighbourhood borders and names on Google Earth.

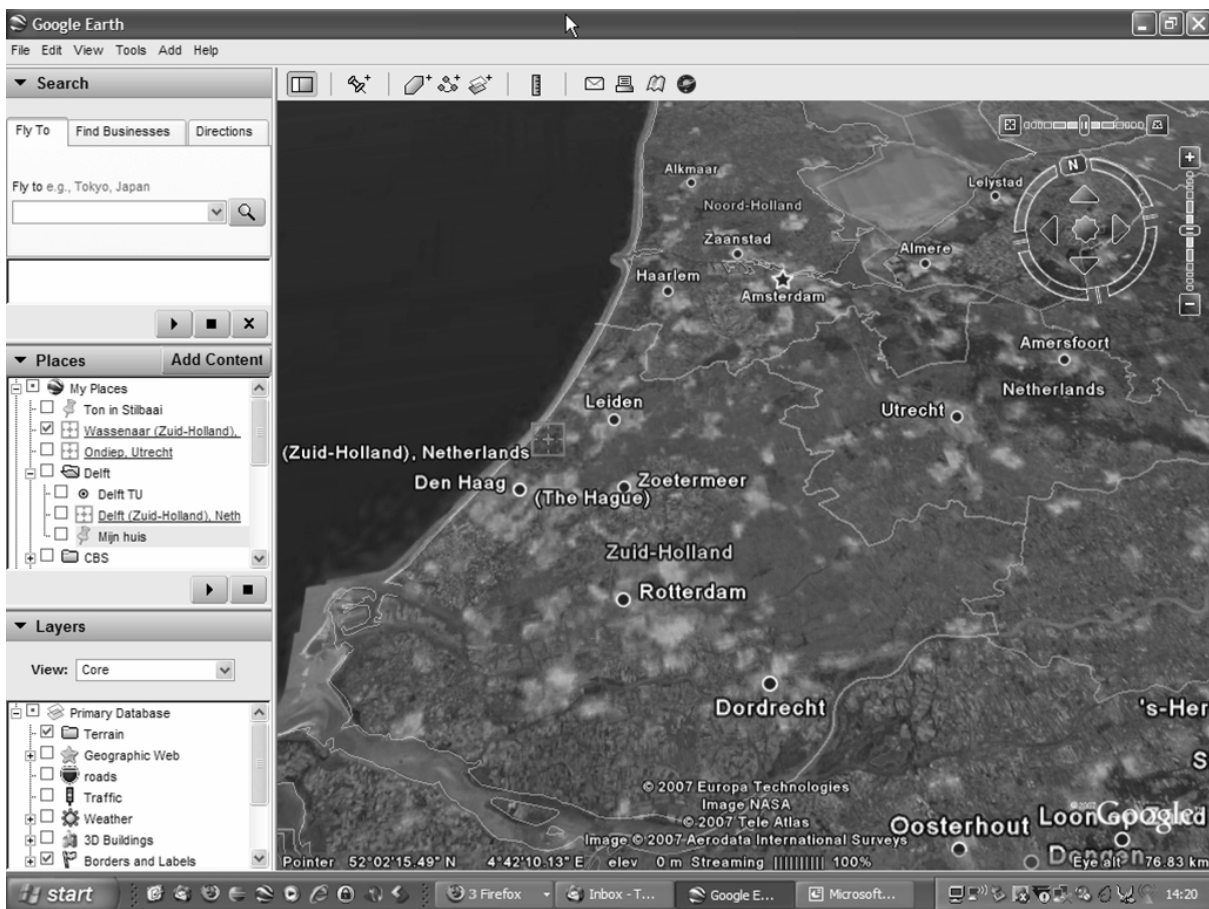


Fig. 5. Population density on Google Earth.

Economy past its peak

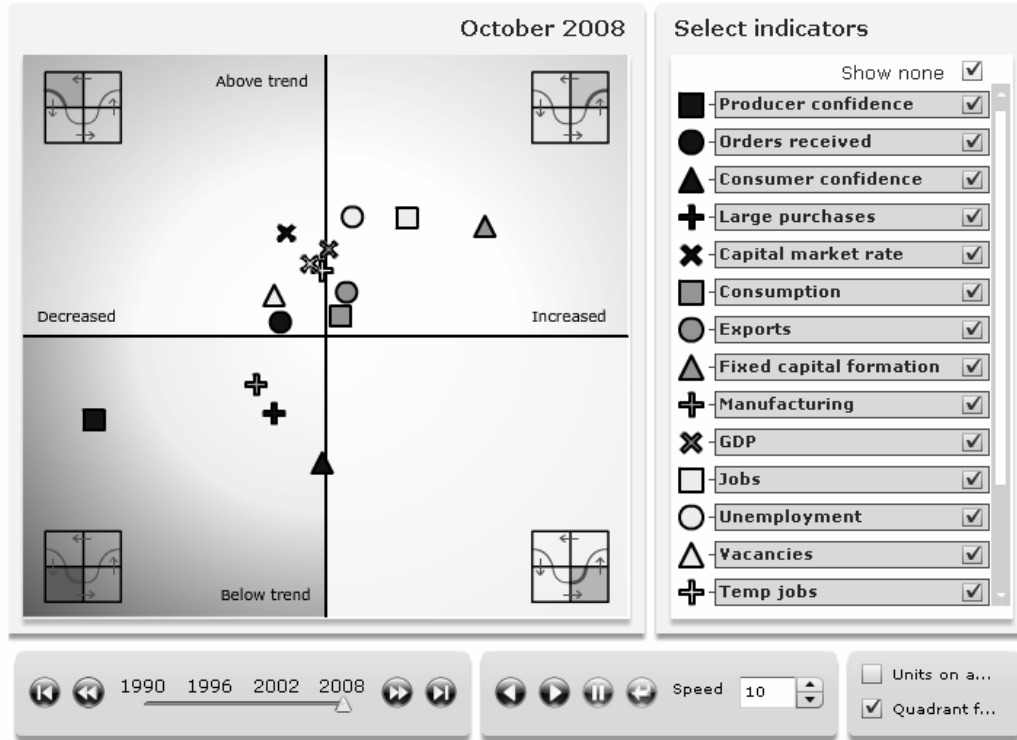


Fig. 6. Business cycle tracer.

or below trend and increased or decreased) to project the data in a four-quadrant space. The business cycle tracer visualisation has been adopted by other statistical institutes as well (Germany, Eurostat).

A completely different example of visualising statistics is the visualisation of adopted children in the Netherlands. Based on a dataset ranging back to 1995, which contains the countries of origin of children adopted in the Netherlands, a map-based representation was chosen. The map highlights the countries in this dataset and annotates them with a square representation of the number of children adopted from that country in that year. This visual is animated over time (Fig. 7). This visualisation is more generic in the sense that the dot representation could also be used to display other country-based statistical variables. However, the colouring of blue and red for boys and girls is, of course, very specific for this dataset.

As mentioned above, the idea of a personal inflation calculator has been around for quite some time in statistical circles. It is an excellent way to personalise statistics as people tend to be very interested in the positive and negative effects of inflation on their

own budget. The German office introduced a personal inflation calculator in which expenditure on specific goods and services can be adjusted using sliders. The Office for National Statistics in the UK created a personal inflation calculator where users can enter their own budget pattern. The Netherlands decided to implement its own personal inflation calculator, which is shown in Fig. 8. Based on a personal spending pattern, personal inflation is shown in relation to official inflation. It is clear that personal inflation calculators exceed the above-mentioned visualisations in complexity, as complex calculations have to be made for each user input. The challenge was to hide the complex mechanism from the user and just sum up the message in a simple chart. The personal inflation calculator is an example of an application that gives the user analysis power. Users can explore the effects of different spending patterns on their inflation and thus analyse inflation for self-defined user groups.

Another example is the visualisation of international trade (both imports and exports) of the Netherlands. Figure 9 shows the visualisation, which uses the fact that the main trade partners of the Netherlands are only

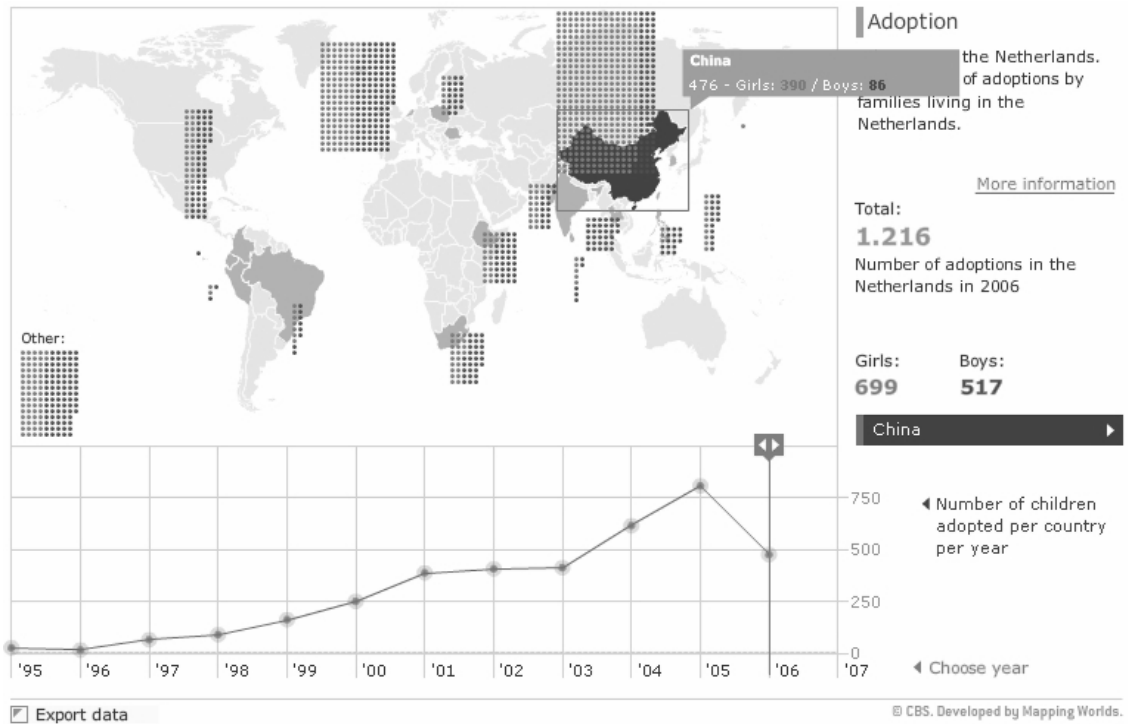


Fig. 7. Visualisation of numbers of adopted boys and girls.



Fig. 8. Personal inflation calculator.

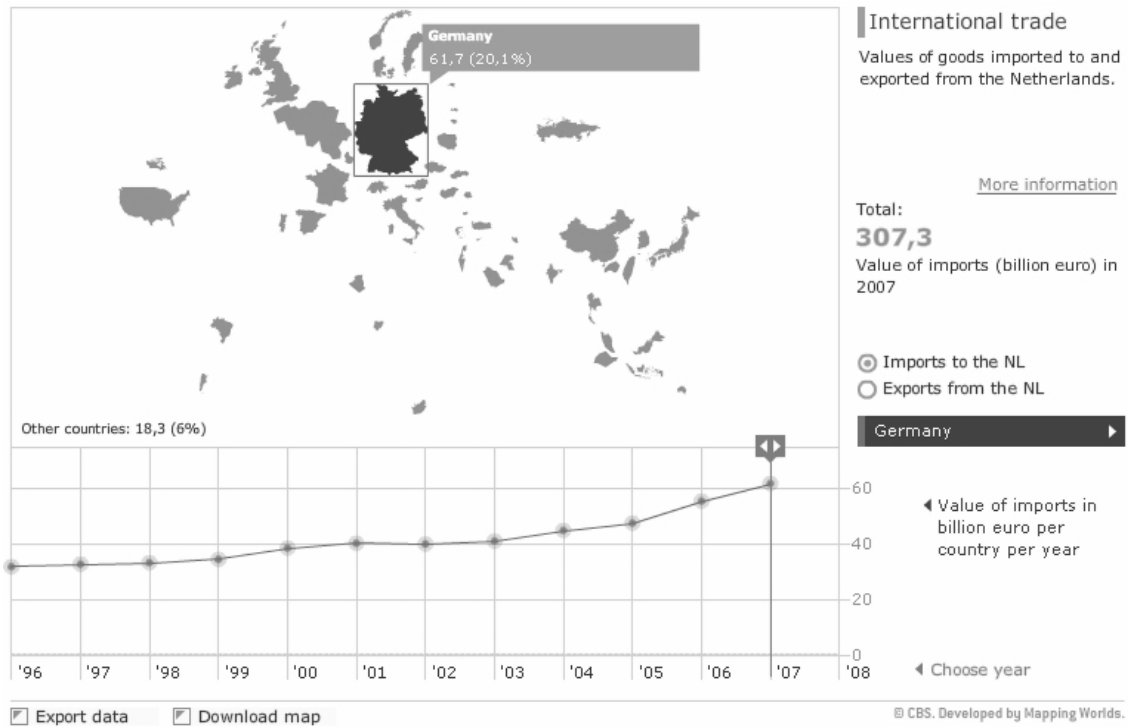


Fig. 9. Visualization of International trade.

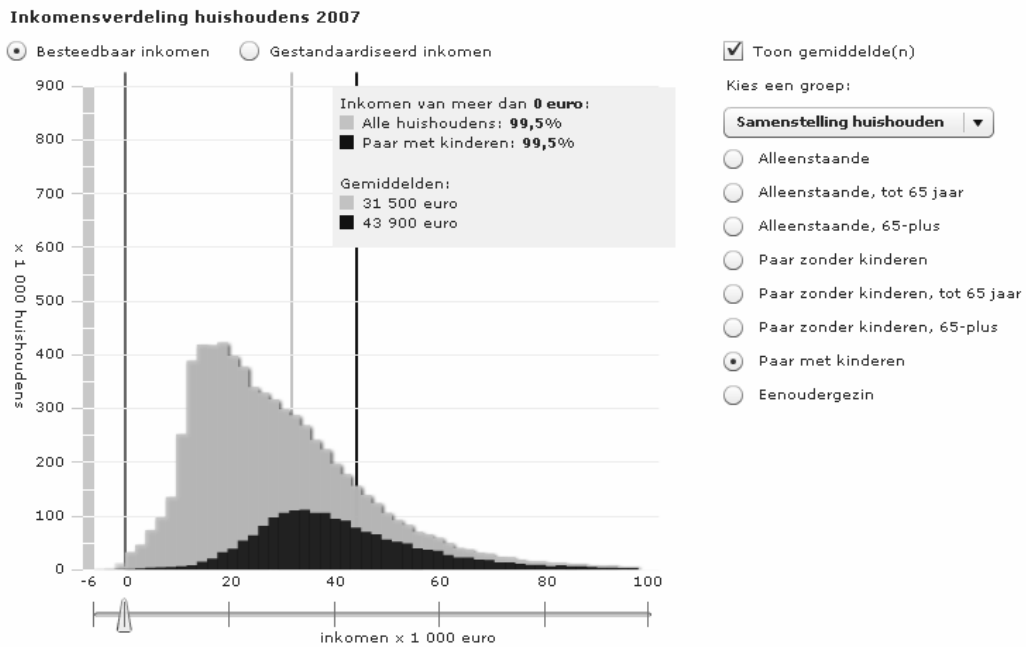


Fig. 10. Income distribution of all households (light) and households with children (dark).

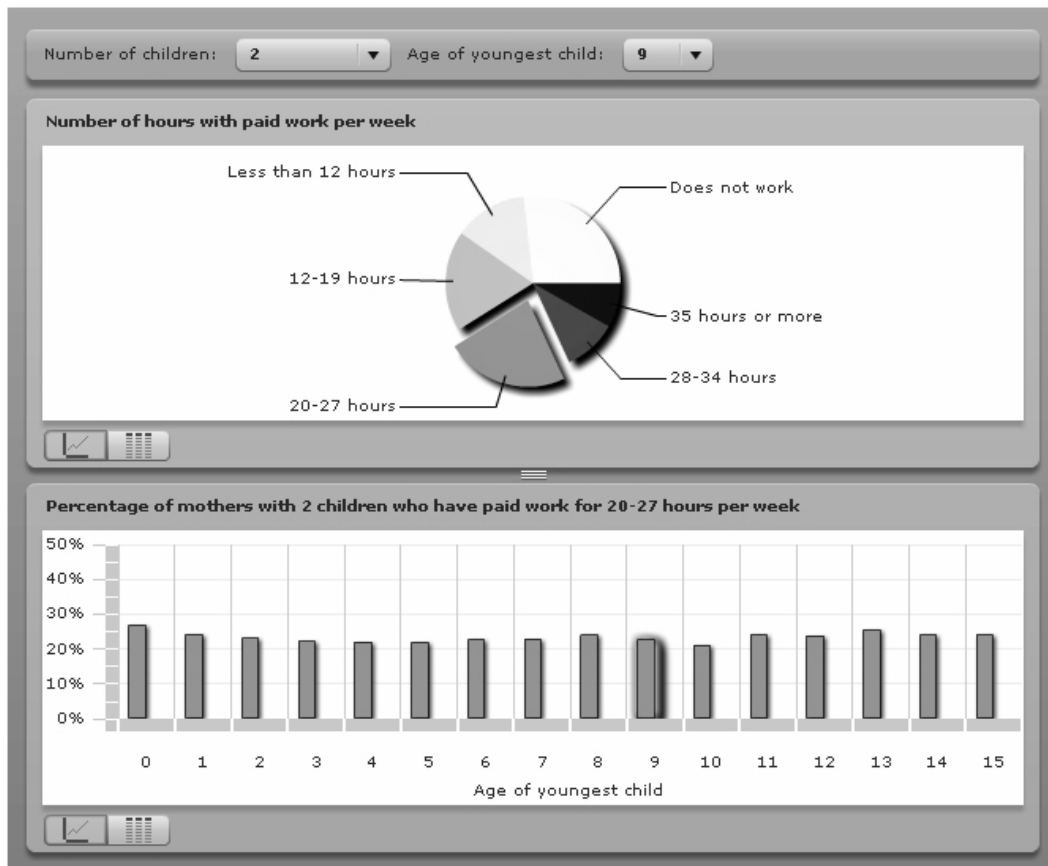


Fig. 11. Working mothers.

a few specific countries. The total values of goods imported or exported is visualised by rescaling the country shape on the world map with respect to its normal size. The example shows that Germany is an important importing trade partner of the Netherlands. The visualisation is an example of using extensive interaction capabilities of dynamic charts, as the rescaling of countries is continuously adapted when the user moves the slider at the bottom to another year.

Another example again is a visualisation showing income distribution in the Netherlands. This is not an animation, but an interactive graph. The income distribution of all Dutch households is shown in combination with the income distribution of the category selected. Figure 10 shows the income distribution of all households in relation to the income distribution of households with children. This visualisation is a good example of analysis skills using a very simple user interaction mechanism. The user can explore differences in income between different household types based on household characteristics, such as age, number of incomes, children, elderly people etc.

A last – very simple – visualisation example is the “working mothers” visualisation, which displays the number of hours women with children work, in relation to the age and the number of children. This is a simple animation of an interactive pie chart and bar chart as shown in Fig. 11. The visualisation is part of the so-called “emancipation” dossier on Statistics Netherlands’ website and was introduced in March 2008 together with a web magazine article) entitled “Mothers continue to work part-time when their children grow older”. It is an example of what can be done by putting together some interacting standard charting components on a dataset.

These examples have taught us that a multi-disciplinary team is essential to create good statistical visualisations. The web development team must work in close cooperation with the statisticians responsible for the data involved. The statisticians are responsible for a correct use of their data. The web development team has knowledge of possible methods of visualisation on the web. It is important to realise that

certain presentation standards that are appropriate for paper publications, may be totally inappropriate on the web. Methodological knowledge is needed to make sure that the visualisations do indeed tell the statistical story to be told, and nothing more than that. Tests were performed by people from all disciplines. And lastly, these visualisations should be launched on the occasion of new data being published, or as an annotation to a web magazine article or statistical report. After all, it is not the visualisation itself that makes the difference, the main story is statistical content, the visualisation serves to make it come to life.

5. Conclusion

Trends in statistical dissemination show an increase in the visualisation of statistical data. Although tables are still essential in the dissemination process, a visual representation of statistical data can improve the understanding of patterns present in the data. Examples are the use of dynamic charts, modern web-based mapping tools, bubble charts, applications that give analysis power to the user and zoomable user interfaces. Statistics Netherlands has taken initiatives to present its data visually. These initiatives are not intended to replace traditional dissemination of statistics, but to add a new dimension. We distinguish between two types of visualisation: tools for exploration, and tools for presentation, corresponding to the difference between generic versus tailor-made applications. Both types of visualisation approaches are useful in the dissemination of official statistics.

So-called streaming photo-based mapping techniques such as Google Maps and Google Earth are becoming more popular for the presentation of official regional statistics. Statistics Netherlands uses these techniques instead of traditional mapping techniques for the dissemination of neighbourhood statistics. It has developed a website, “CBSinyourNeighbourhood”, using Google Maps. A statistical layer on Google Earth was published via the main website of Statistics Netherlands for more advanced users. The Google Earth layer was implemented so that performance is optimal. This “just in time” loading mechanism uses the underlying streaming mechanism of Google Earth. These initiatives make it possible to present neighbourhood statistics using tools and media people use in everyday life. Future plans include the addition of new elements to the Earth layer, such as projections of one or more sta-

tistical variables on “heat maps” and the possibility to visualise statistics over time.

Statistics Netherlands has recently intensified its visualisation activities. In dialogue with in-house statisticians statistical data suited for visualisation have been selected and turned into graphical representations of the trend to be shown. Examples are the business cycle tracer, the map of countries from where children were adopted, a personal inflation calculator and a map in which countries scale to the volume of goods they import from or export to the Netherlands. Lessons learned with these visualisation initiatives include the formation of a multi-disciplinary team. Also, the visualisations should ideally be launched together with a new data release.

Acknowledgments

We would like to stress that moving ahead statistics visualisation is not the work of just a few people. To bring data to life in a responsible way, statisticians and web developers have to work closely together. Many of the examples in this paper could not have been created without them. We thank all the people at Statistics Netherlands who have developed ideas and technology to advance statistics visualisation.

References

- [1] B.B. Bederson, L. Stead and J. Holan, Pad++: *Advances in Multiscale Interfaces*, ACM SIGCHI '94, 1994.
- [2] B.B. Bederson, J. Meyer and L. Good, Jaz: *An Extensible Zoomable User Interface Graphics Toolkit in Java*, ACM UIST 2000, May 2000, pp. 171–180.
- [3] O. ten Bosch and E. de Jonge, *From statistical data supplier to statistical information service provider*, paper IMAODBC 2006, Avila.
- [4] O. ten Bosch and E. de Jonge, *Visualization of Neighbourhood Statistics using Google Earth*, invited paper, Unece Meeting on the Management of Statistical Information Systems (MSIS), Luxembourg, April 2008.
- [5] W.S. Cleveland, *Visualizing Data*, Hobart Press, 1 edition, March 1, 1993.
- [6] W.S. Cleveland, *The Elements of Graphing Data*, Hobart Press, 2 edition, October 1, 1994.
- [7] S. Few, *Show Me the Numbers: Designing Tables and Graphs to Enlighten*, Analytics Press, 2004.
- [8] Stephen Few, *Multivariate Analysis Using Parallel Coordinates*, <http://www.b-eye-netwerk.com/view/3355>, September 12, 2006.
- [9] M.D. Levi, *Statistics in Cyberspace: how has the internet changed the U.S. Bureau of Labor Statistics?* Joint ECE/Eurostat/OECD meeting on the management of statistical information systems (MSIS), Sofia June 2006.

- [10] E. Tufte, *The Visual Display of Quantitative Information*, (2nd edition), Graphics Press, May 2001.
- [11] E. Tufte, *Sparklines: Intense, Simple, Word-Sized Graphics* May, 2004.
- [12] Google Website, *Google Earth documentation*, <http://earth.google.com>.
- [13] Google Website, *Google Maps documentation*, <http://maps.google.com>.
- [14] Google Website, *Google Finance*, <http://finance.google.com>.
- [15] IBM research website, *Many Eyes*, <http://manyeyes.alphaworks.ibm.com/manyeyes/>.
- [16] IMF Website, *Datamapper*, <http://www.imf.org/external/datamapper>.
- [17] OECD Website, *OECD Explorer*, <http://stats.oecd.org/OECDregionalstatistics/indexTL2.html>.
- [18] Gapminder website, *Gapminder*, <http://graphs.gapminder.org>.
- [19] Insee, *Cartographic maps*, <http://www.statistiques-locales.insee.fr/>.
- [20] Google Website, *KML Reference*, <http://earth.google.com/kml>.
- [21] ESRI, <http://www.esri.com/>.
- [22] Geoclip, <http://www.geoclip.fr/an/>.
- [23] F. van Ruth, *The Statistics Netherlands Business Cycle Tracer; Visualising the State of the Economy*, OECD / ISTAT Seminar on Dynamic Graphics for Presenting Statistical Indicators, Rome, March 2007.