

The Product Explorer: Decision Making with Ease

Patrick Riehmann
patrick.riehm@uni-weimar.de

Jens Opolka
jens.opolka@uni-weimar.de

Bernd Froehlich
bernd.froehlich@uni-weimar.de

Virtual Reality Systems Group
Bauhaus-Universität Weimar
Bauhausstrasse 11, 99423 Weimar
Germany

ABSTRACT

We present the Product Explorer, an interactive parallel coordinates display for facilitating the selection process of typical products offered in online shops. Users can quickly narrow down the product search to a small subset or even a single product by using our visual query interface. Our study confirms that our interactive Product Explorer is a fast and easy-to-use tool for the product selection process of casual users.

Categories and Subject Descriptors

H.5.2 [User Interfaces]: Graphical user interfaces

General Terms

Parallel coordinates, product visualization, categorical data.

1. INTRODUCTION

The major domains of parallel coordinates are still limited to the academic world. We found that occasionally companies utilize them for a specific purpose, but for the public this representation is still largely unknown. Nevertheless, the search and exploration of multi-dimensional data sets is the basis of most decisions that have to be made nowadays, such as when buying a new car or washing machine or looking for a specific building material. For these tasks, however, text-based Web interfaces for selecting products in shops are most common despite several issues: Users are not able to view all products at a glance nor estimate the overall number of products which they can choose from. In most stores just a few products are preselected. The search masks consist mainly of combo boxes or drop down lists. The search does not instantly begin while typing the search items. The resulting product list is not sortable by arbitrary attributes. Empty result sets often occur after searching a certain configuration, but most interfaces lack a possibility to inform you about the attributes that caused that situation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AVI '12, May 21-25, 2012, Capri Island, Italy
Copyright 2012 ACM 978-1-4503-1287-5/12/05 ...\$10.00.




Figure 1: The Product Explorer showing a notebook data set. The user has constrained various attributes. The blue lines represent the remaining notebooks fulfilling the user's requirements. On the left a list of all notebooks is displayed showing the selected ones on top. Each list entry is linked to the parallel coordinates display. On the bottom left the decision bar contains an item, which represents an exclusive decision with respect to the CPU vendor for narrowing down the search. On the right there is an attribute repository for adding and removing axes from the display. Highlighting (in red) allows users to explore the attributes of one or more products in detail.

These limitations motivated the development of the Product Explorer, a tool that uses an interactive parallel coordinates display to keep users well informed throughout the entire product selection process. All products and important attributes are visible at a glance and all the time. All interactions can be performed with immediate feedback and users do not have to wait until a query response appears. Our prototype is particularly suited for coping with product data by providing effective drawing of polylines and axes for easier product recognition in the parallel coordinates display. We provide a simple and intuitive visual interface and an appropriate logic for specifying attribute values and ranges for quickly narrowing down the product search to a small subset or even a single product. Product Explorer also provides effective means to cope with multiple dimensions by supporting an attribute repository and a decision bar for storing axis with already finalized configuration decisions.

The main contributions of this work are: the visual query specification along with the exclusive decisions technique; the axes repository for making parallel coordinates an effective tool for the product selection process; a guideline for selecting the appropriate drawing styles and enhancements. Our user study confirms that the Product Explorer is indeed an excellent tool for its intended purpose for casual users.

2. RELATED WORK

Inselberg [8] invented parallel coordinates many years ago and the earliest ideas of it dating back to the 1960. A two-dimensional mapping of the multi-dimensional space is introduced by drawing each dimension as a vertical axis on a regular two-dimensional canvas. Each point in the multi-dimensional space is depicted as a sequence of line segments connecting the different vertical axes. In recent years, many important improvements with respect to interaction, drawing and data organization of parallel coordinates have been made. Siirtola [13] proposed a variety of ideas for directly manipulating parallel coordinates, for example lines can be selected and grouped with the help of logical operations. The lines within a group are visually represented by a single line, which is defined by the mean of all included lines.

Graham et al. [5] proposed using a combination of quadratic and cubic curves to represent a data line. Higher order parallel coordinates [14] place invisible axes between two adjacent axes to control the curvature of the line segment. Yuan et al. [18] presented another way of controlling the curvature of a line by combining parallel coordinates, scatter plots and multidimensional scaling. Here, the points of a scatter plot are printed between two adjacent axes and used as control points for the curve. Illustrative parallel coordinates [10] go even further: McDonnell et al. focus on artistic drawing techniques to convey as much information as possible. They use various cluster techniques and adapt the edge bundles presented by Holten et al. [7]. Fanea et al. [3] combine parallel coordinates with star glyphs, by creating a three-dimensional rendering.

Heinrich and Weiskopf [6] developed a method for mapping continuous scatterplots into a density model for parallel coordinates. When using large amounts of data, it becomes necessary to structure the data. Binning can be used to gather single lines into groups. Outlier detection can further enhance the quality of the rendering [11]. Fua et al. implemented a hierarchical structure for parallel coordinates. The data is organized in a hierarchy of clusters. Advantages of structured data are an adjustable level of detail and improved performance. Dimensional reordering techniques can help to reduce over-plotting in cluttered parallel coordinate plots. Based on Peng et al. [12] dimension reorder has a significant effect on the visual expressiveness of a visualization.

Yang et al. [17] designed an interactive method for dimensional reordering, spacing and filtering based on dimension hierarchies. Reordering and spacing puts similar dimensions next to each other and reveals the structures within the data. Wong et al. developed an edge lens that allows the user to bend nearby lines in the parallel coordinate plot away from a point of interest, to reveal underlying structures in dense areas [16]. Another way of clutter reduction is random sampling of the data. Ellis et al. propose an automatic sampling lens to subsample dense areas [1]. This way trends in the data can be discovered with reduced over-plotting.

Product data attributes are partially categorical. An early

approach for visualizing categorical data was introduced by Hartigan and Kleiner [4]. Mosaic plots map categories into square tiles, where the size of a tile depends on the frequency in the corresponding class. This technique was extended to three-way mosaic plots and combinations of mosaic plots and scatter plots [15]. Finally, Bendix and Kosara introduced Parallel Sets, a new technique which focuses on categorical data [9]. Relationships between different attributes are depicted by ribbons passing from one category in one axis to another category on the next axis. Categories can be arranged on each axis by the user and can be combined into new meta-categories. Elmqvist [2] offers a new presentation of scatterplot matrices by mapping the individual scatterplots on an n-dimensional cube. The user is able to navigate (like turning a dice) around this cube and visually sculpt a suitable subset of a product data base. Our approach also focuses on product data bases, but it always provides an overview of the data in one view.

3. VISUALIZING PRODUCT DATA

Product data bases are rarely perfect and often quite complex. In particular, if the data base is constructed by crawling the web, various attributes of individual items may be missing. A simple solution is to remove all dimensions that are not complete, but some of the removed attributes could be crucial for choosing a product. Thus, the visualization needs to be able to cope with incomplete data sets.

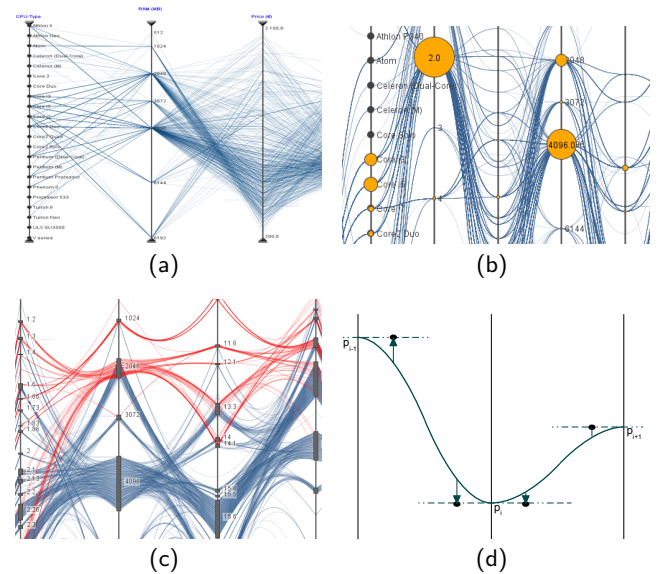


Figure 2: (a) Different kinds of axes in a classical manner of drawing: (From left to right) Categorical data, pseudo-continuous (ordinal) and continuous data axes. (b) Histograms summarize the paths sharing the same value. (c) Drawing paths as bundles. (d) Construction of a cubic curve.

The attributes of a product may be continuous, categorical or ordinal. Even continuous attributes are often closer to an ordinal data type since only certain values are assumed, e.g. for the display size of a notebook. We refer to these data types as being pseudo-continuous. Binary data shows whether or not a feature exists. Generally, binary data can

be treated as a special kind of categorical data (see Figure 2(a)). In general, parallel coordinates treat each dimension as essential, but, according to product data, the attributes a product possesses are not equally important. Moreover, various attributes may be more or less important to each user for making a decision. Therefore, a global importance of an axis has to be taken into account, which is suitable for most users as well as an individual assessment needs to be enabled.

The terms clutter and over-plotting are often not very clearly distinguished, so we define clutter as a large number of axes and lines drawn together in the same system, such that it is usually overwhelming for users to recognize features and structures in the data or to interact with the display. With plain parallel coordinate drawings this issue cannot be avoided. As discussed in the related work, several techniques have been developed that try to overcome this problem.

The other major issue, particularly for product database visualization, is the handling of categorical and ordinal data. An ever-increasing amount of lines is plotted one above the other which is caused by sharing the same values in succeeding axes. Sometimes lines continue on the same path even over many axes. Thus, users are not able to recognize how many lines (i.e. items) share the same value(s) in one or more dimensions (see Figure 2(a)).

3.1 Drawing Axes and Extended Areas

A simple possibility to help users to clearly understand the number of products with an identical attribute is the use of round histograms, which summarize the paths sharing the same value within an axis. This technique provides a visual representation of the quantitative distribution and relations within a dimension as well as over the entire plot. Additional information can be easily displayed within, if there is enough space. However, all overlapping path segments still pass through the center point of the histogram (Figure 2(b)).

Thus, we prefer another approach to visualize the number of lines going through a point. The aim is to spread a particular point on an axis to a vertical area. The vertical height of a point is determined relative to the number of lines that pass through the point. This configuration also shows how many products have a certain value and reduces the over-plotting situation at crowded points by distributing the start and end coordinate of line segments between the axes to a vertical range (see Figure 1). The fraction of an axis that can be used for extended areas is globally adjustable.

However, we believe that there is an optimal range for this fraction. Either the areas are too small, so users are still not able to recognize the amount of lines going through or the areas are too large such that the users lose the mental association of an area as a point. Graham [5] proposed a similar possibility to vertically extend points. Contrary to our aim, which is the visualization of frequencies along an axis, his solution is motivated by aesthetic reasons to improve the curvature of a line across the axes. Also the interaction is used to adjust the curves by spreading the already extended points further. Instead, we use the extended areas for an explicit selection of an attribute value.

3.2 Visualizing Missing Data

As previously mentioned, product data bases are rarely complete with respect to the set of potential attributes. Therefore, we have to visually represent gaps within the data

set. Instead of introducing a pseudo-category *undefined* on each axis and avoid giving the false impression that there exists a value on an axis, we do not draw a line or a curve across the axes which do not contain a value for the current line.

We propose two simple possibilities to circumvent this problem (Figure 3). The first one just indicates a line by drawing a stub, which does not cross the axis where the value is missing. In combination with extended areas, the occurrence and the amount of missing values can be easily estimated, but which stubs correspond with which counterpart on the next axis can only be assumed (Figure 3 left).

Thus, we propose a second possibility in which a translucent line with decreasing opacity links the two stubs (Figure 3 right). The barely-visible line spans between the inconsecutive axes and connects at the appropriate values, but it disappears behind the axis that does not contain a value for the given line.




Figure 3: Two possibilities of visualizing gaps in the dataset. The stubs on the left indicate that the path continues, but there is no indication where it reconnects. On the right the path disappears behind the axis that does not contain a value.

3.3 Drawing in-between Axes

Several improvements have been made during recent years for drawing lines between axes. We experimented with several kinds and combinations of bundling techniques as well as with specific curves for supporting better recognition and improving aesthetics. Utilizing bundles of curves (Figure 2(c)) instead of straight lines visually smoothes the parallels coordinates and enables easier pattern recognition by tending a bundle lines to its center. Therefore bundles can be recognized and can be distinguished between each other much more easily than in common drawings.

As mentioned in the related work section, using curves instead of lines facilitates the tracking capabilities of a single data item across all axes, thus making it easier to distinguish between different paths. Usually, cubic and quadratic splines are combined to avoid oscillation above and below the vertical endings of the axes and to guarantee an aesthetic curvature.

We implemented polylines, a combination of quadratic and cubic splines and bundling for drawing lines, and showed them to users. However, they remarked that these approaches do not work well in combination with the extended areas, since they increase occlusion in the direct proximity of an extended area. Simple cubic functions offer an alternative strategy for drawing lines. They are fully defined by two points and two derivatives on two successive axes. We found them particularly effective in the combination with extended areas if the curves between two axes intersect both axes in

a right angle since this aids the overall legibility by creating a steadier path (Figure 2(d) and 1).

3.4 List

All the products are displayed in a list which is shown on the left side (Figure 1). To reinforce the visual metaphor of one path being one product, every path originates from the corresponding product entry within the list panel. The list can be understood as a special type of axis. However, some users also appreciate a familiar interface as a helpful start for dealing with parallel coordinates.

4. VISUAL QUERY GENERATION

How do people usually search for products? Often they start with some requirements they are sure about, some others that are not very specific and finally a few that are less or not at all important. They start searching for products that match these fuzzy requirements and put them into an imaginary subset. If no item appropriately matches their requirements, users then have to modify their search. On the other hand, if there are too many matches, they must narrow down the features that should be covered.

This strategy can be compared to an iterative process of sending queries to a database. Queries have to be relaxed if no results have been received, or new clauses are added to the query, or the existing ones are varied until the user is satisfied with the received subset.

The aim is to visually generate and refine a query, which specifies the user's requirements for the various attributes of a product. In the context of parallel coordinates this has to be expressed by selecting paths, for which several approaches can be combined: (1) The user can define a range on the axis with the help of two sliders. (2) Another possibility is to group items by brushing over certain paths. (3) We suggest to utilize the extended areas, which also work for histograms, for directly selecting certain attribute values. This mechanism is intended for axes of categorical and ordinal data. The sliders work equally well in all the mentioned kind of axes, thus it is possible to mix attribute range selection and direct attribute selection in a single query.

Early user feedback on these basic approaches eventually resulted in a consistent mechanism for choosing products rapidly and comfortably. Our prototype combines a slider-based interaction along with the possibility of selecting extended areas. Within an axis all selected attribute values and ranges are aggregated like an *OR* operation regardless of whether they have been selected by extended areas or by the sliders. The subset of selected products on each axis are intersected like an *AND* operation with all the other axes' selections. Products that fulfill the specified query are drawn in a salient color (Figure 5).

As mentioned before, all lines in the parallel coordinates display originate from the product list on the left. All products that match the entire query are at the top of the list, followed by partially satisfying products ordered by their rank and finally by the products that do not match up at all. All list items are drawn with the same color as the related paths. The list is reorganized instantly with respect to the product ranking depending on chosen ranges and values. The explicit marking of products is also possible (Figure 1).

In addition to the introduced selection mechanisms, our system provides a set of common parallel coordinates techniques that help users to organize the display. The axis

direction can be switched from top down to bottom up. For comparing each axis with others, all axes can be horizontally dragged and dropped in arbitrary positions. The lines will instantly be redrawn during movement. Another capability allows users to zoom into the space between two axes with the mouse scroll wheel, making it possible to get a better overview of the path segments and their connection between two axes. As a result the other axes are temporarily squeezed together.

5. EXCLUSIVE DECISIONS

The ranges and sizes of displays people have nowadays are vast, including high resolution monitors, projection systems and smart phone displays. For making our Product Explorer available on these various platforms for different types of products, we have to particularly consider the number of axes that can be drawn at once since a useful drawing must provide appropriate space between axes so that the paths can be easily followed. As mentioned in the section about data issues, the axes' order should reflect the importance of the individual axis. Therefore, the most important axes should be preferably drawn first as long as sufficient gaps between axes can be provided.

For the remaining attributes, we propose an attribute repository, which helps the user to organize the axes that do not fit in the drawing area. With the repository, the user can decide to add an attribute to the drawing area or exchange axes between the repository and the drawing area. As depicted in Figure 7, the attributes can be easily chosen (via drag-and-drop) from the repository which contains only the hidden axes. The repository is implemented as a drawer and is only opened if necessary. The attributes themselves are depicted as miniature axes to maintain a similar aesthetics and to aid the user when dealing with attributes that are almost as important as the previously selected ones.

Moreover, we propose an exclusive decisions technique to reduce the number of axes in the drawing area. Additionally, this technique supports the process of fast decision-making that our prototype is intended for. We assume that according to the individual importance of each axis, the major attributes will be chosen first and the user is very certain about these decisions. Therefore, there is no need to waste space by continuing to draw such axes. After selecting an extended area or adjusting a range on an axis, the axis will disappear and thus all lines that do not belong to the current selection will be hidden. However, sometimes it may be necessary to step back and to change a decision already made.

Therefore, a history below the parallel coordinates display shows the decisions (the chosen axes **and** its chosen values) the user has had already made and they are then able to reselect an axis that needs to be taken back into consideration. All these techniques are facilitated by appropriate animated transitions, which help the users in recognizing the appearance and disappearance of axes and paths.

Unfortunately, because of the general orientation of the parallel coordinates, scalability issues with the vertical space cannot be solved with the same methods for the horizontal direction. Our aim is to depict all products at once, so pre-computing techniques like binning or clustering are not suitable for this purpose. Instead, we expect that a Focus+Context technique along the axes might be useful.

6. USER FEEDBACK

Our application prototype is implemented in Java and we employed Java2D for drawing operations. Java2D provides convenient functions for drawing curves and the graphics update rates are still smooth on high resolution displays.

We performed an initial user study to assess the basic usability of our implementation with a fixed set of axes and without the axis repository on the bottom of the display against a typical Web interface consisting of entry fields, drop boxes and option buttons and containing a submit button, which invokes the search by requesting a result page.

In particular, we were interested in the following aspects: (1) How does the representation of a product by a single path across all the axes work for the participants? (2) Is an extended area still recognized as a single point or attribute? (3) Which possibilities for drawing axes and lines are generally preferred? (4) Does the combination of different kinds of selection techniques work well? (5) Is the Product Explorer appreciated by the participants or do they prefer a classic web interface.

None of the eight participants were familiar with our work. The study comprised six test product assignments that were selected for variety and with two different difficulties (three each). For example:

- **Easy** Look for a notebook under 1000 Euros, which is as light as possible and has an AMD processor.
- **Difficult** You want a notebook with an Intel CPU and an Intel graphics chip. The CPU should be a Core i3 or i5. Find the fastest notebook with the lowest price.

After a brief introduction to the parallel coordinates interface, the participants were asked to search for and to select the most suitable product(s) for the given assignments. To minimize learning influences and other effects, the participants were divided into two subgroups: one subgroup began the trial with the Product Explorer and the other group started with the Web interface. Subsequently, the participants were asked to answer a questionnaire. It consisted of two parts. Part one asked about the interaction in general (e.g. how natural it was to use, preference for an interface etc.). The second one was designed to gather feedback about the aesthetics decisions that have been made during development. Rating questions used a Likert scale from one to six.

A *t*-test was conducted to compare the results of both interfaces. We used an alpha level of .01 for all statistical tests. As expected, the Product Explorer was significantly faster for the difficult tasks ($M = 39.12$, $SE = 9.14$) vs. ($M = 92.00$, $SE = 9.14$), $t(7) = 6.161$, $p < .001$ as well as for the easy tasks ($M = 32.20$, $SE = 4.19$) vs. ($M = 72.00$, $SE = 3.83$), $t(7) = 8.67$, $p < .001$ (Figure 4). The subjective ratings confirm these results. Overall, the participants prefer the Product Explorer over the Web interface; ($M = 5.50$, $SE = .19$) vs. ($M = 3.25$, $SE = .16$), $t(7) = 13.74$, $p < .001$, which reveals a high acceptance for our interface that uses novel interaction techniques and visual metaphors that none of the participants had ever seen before.

The “one path is one product” dualism seems to be very comprehensible for the users with a rating of 5.8 as well as the “one extended area represents one value” metaphor with a rating of 6. The range slider was judged as helpful by nearly half of the participants, but the area selection




Figure 4: Mean task completion times over all participants for both levels of difficulty along with their standard errors.

technique was considered much more effective. The tutor’s observations during the assignments confirmed the predominant usage of the area selection.

It might be that Grahams method [5] provides an aesthetic curvature for a single path. However, for the visual impression of the entire plot, the users favor the simple cubic curve approach (7 of 8 participants) which they found to be more steady and legible. Furthermore, the participants preferred the extended area instead of histograms (6 against 2) if they were forced to choose between the two. Moreover, those who voted for the histogram would actually like to see a combination of both.

The participants of our study suggested several features to improve the prototype. The most frequently mentioned ones were: (1) to reconsider the ordering of the values along the axis (most wanted by frequency) and (2) to start with only a few axes and to slide in additional axes on demand. In light of this feedback, the latter feature has already been implemented and discussed in this paper.

Not surprisingly, all of the participants of our study would like to use our Product Explorer for choosing products on the web. This was also the general opinion of people who used our system during public demonstrations. The most interesting finding during these events was that other researchers, such as engineers, architects and archaeologists, are generally not familiar with parallel coordinates. Nevertheless, they are faced with similar problems (e.g. finding the best-suited material for a bridge, searching for a glass with certain requirements or tagging and organizing thousands of small archaeological pieces, which have numerous attributes). On the whole, after a brief introduction, they are excited about how helpful such a tool can be for their own field of research. It can help them to organize, present (especially with capabilities to mix categorical and continuous data) and find items by visually generating a query.

Our user study is not particularly fair regarding the most advanced web shop search interfaces, which already utilize all new capabilities that are given by the so called “WEB 2.0”. We believe that a web form made of sliders, ranges and also checkboxes which realizes an instant approach to display the result list by any change of the user’s requirements and without reloading the entire website (for example using the XML-HTTP request) can be nearly as fast as the Product Explorer for certain queries.

But even with these improvements a list-based results view cannot show all items at once, it cannot display whether there are other possibilities in the vicinity of a user's requirements that do not exactly match. On the contrary, with the Product Explorer, the user can see what configurations are offered by the market, which influences further expectations. Moreover, if the users' wishes cannot be fulfilled, they will be able to see why and what alternative options are available. Additionally, before users make their next decision, they can anticipate what options will be available.

7. CONCLUSIONS AND FUTURE WORK

Our work on the Product Explorer and in particular our pilot study clearly confirm that parallel coordinates with a number of crucial improvements and enhancements can be successfully used to implement product search interfaces on the Web that are intuitive, fast and provide a user with much more information than the commonly used text-based interfaces. The following layout decisions and enhancements contributed significantly to the usability of our tool:

- Extended areas instead of histograms or even points facilitate the tracing of product attributes.
- Cubic curves in combination with extended areas lead to a tidier display of a large number of products by avoiding occlusions near an axis.
- An intuitive visual interface based on a set of simple rules is the key for quickly narrowing down the product search.
- The attribute repository is important for the scalability to various display sizes. Also, this way one may start with the most important axes.
- The decision bar reduces the complexity of the display by reducing the number of considered attributes and products, which have already been considered.


There are many further possibilities to refine and optimize our tool. Missing attribute data is currently only treated in the visualization, but it also needs to be handled during the visual query specification by adding, for example, a selectable pseudo-value or pseudo-category on each axis. Support for multiple range specifications on a single axis might be useful when working with data sets containing mainly continuous axes. We also see some potential for reducing the overall number of axes by merging related axes, for example, more specific categories could be generated by combining CPU vendor and CPU type or hard disk size and hard disk velocity.

People working in various companies asked for a version of our Product Explorer prototype to present and search their own product catalog. Additionally, and perhaps of greater interest, they would like to visualize their respective markets with a comparison of all products from the competing companies. Thus, our tool needs to support different product types and their relationships and dependencies.

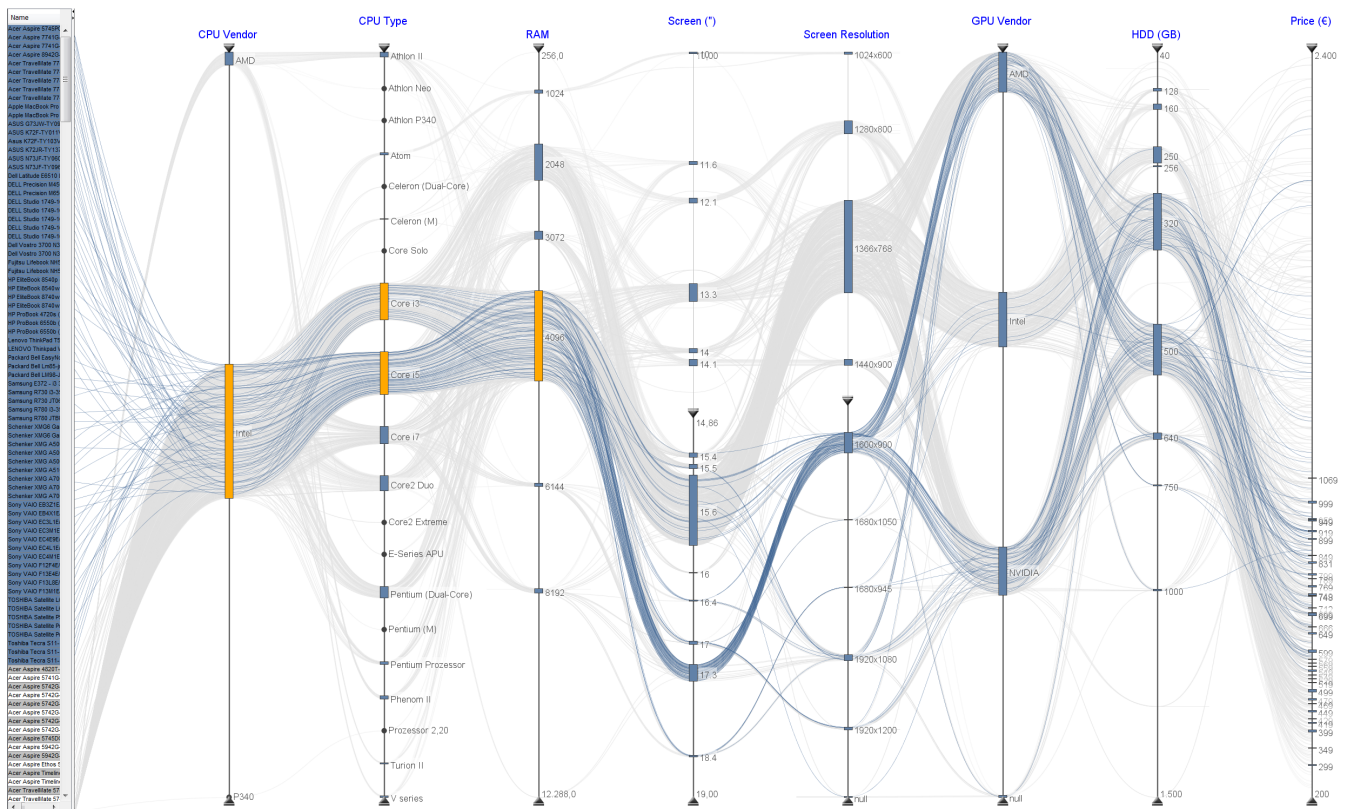
Our most important next step is the development of an HTML-only implementation without the need for plugins. We plan to integrate this new version in a web shop interface or a product search site to gain more experience with a larger audience and further advance the visual Product Explorer interface.

8. REFERENCES

- [1] G. Ellis and A. Dix. Enabling automatic clutter reduction in parallel coordinate plots. *IEEE TVCG*, 12:717–724, September 2006.
- [2] N. Elmquist, P. Dragicevic, and J.-D. Fekete. Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. *IEEE TVCG*, (14):1141–1148, 2008.
- [3] E. Fanea, S. Carpendale, and T. Isenberg. An interactive 3d integration of parallel coordinates and star glyphs. In *Proc.*, INFOVIS'05, pages 20–28, 2005.
- [4] M. Friendly. A brief history of the mosaic display. *Journal of Computational and Graphical Statistics*, 11:89–107, 2001.
- [5] M. Graham and J. Kennedy. Using curves to enhance parallel coordinate visualisations. In *Proc. of the Seventh International Conference on Information Visualization*, pages 10–16, 2003.
- [6] J. Heinrich and D. Weiskopf. Continuous parallel coordinates. *IEEE TVCG*, 15:1531–1538, November 2009.
- [7] D. Holten. Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *IEEE TVCG*, 12:741–748, September 2006.
- [8] A. Inselberg. *Parallel Coordinates, Visual Multidimensional Geometry and Its Applications*. Springer, 2009.
- [9] R. Kosara, F. Bendix, and H. Hauser. Parallel sets: Interactive exploration and visual analysis of categorical data. *IEEE TVCG*, (12):558–568, 2006.
- [10] K. T. McDonnell and K. Mueller. Illustrative parallel coordinates. *Computer Graphics Forum*, (27):1031–1038, 2008.
- [11] M. Novotny and H. Hauser. Outlier preserving focus+context visualization in parallel coordinates. *IEEE TVCG*, (12):893–900, 2006.
- [12] W. Peng, M. O. Ward, and E. A. Rundensteiner. Clutter reduction in multi-dimensional data visualization using dimension reordering. In *Proc.*, INFOVIS'04, pages 89–96, 2004.
- [13] H. Siirtola. Direct manipulation of parallel coordinates. In *CHI '00 extended abstracts on human factors in computing systems*, CHI EA '00, pages 119–120, 2000.
- [14] H. Theisel. Higher order parallel coordinates. *Proc. of VMV '00*, pages 415–420, 2000.
- [15] M. Theus. Interactive data visualization using mondrian. *Journal of Statistical Software*, 7(11):1–9, 11 2002.
- [16] N. Wong, S. Carpendale, and S. Greenberg. Edgelens: an interactive method for managing edge congestion in graphs. In *Proc.*, INFOVIS'03, pages 51–58, 2003.
- [17] J. Yang, W. Peng, M. O. Ward, and E. A. Rundensteiner. Interactive hierarchical dimension ordering, spacing and filtering for exploration of high dimensional datasets. In *Proc.*, INFOVIS'03, pages 105–112, 2003.
- [18] X. Yuan, P. Guo, H. Xiao, H. Zhou, and H. Qu. Scattering points in parallel coordinates. *IEEE TVCG*, 15:1001–1008, November 2009.




(a)

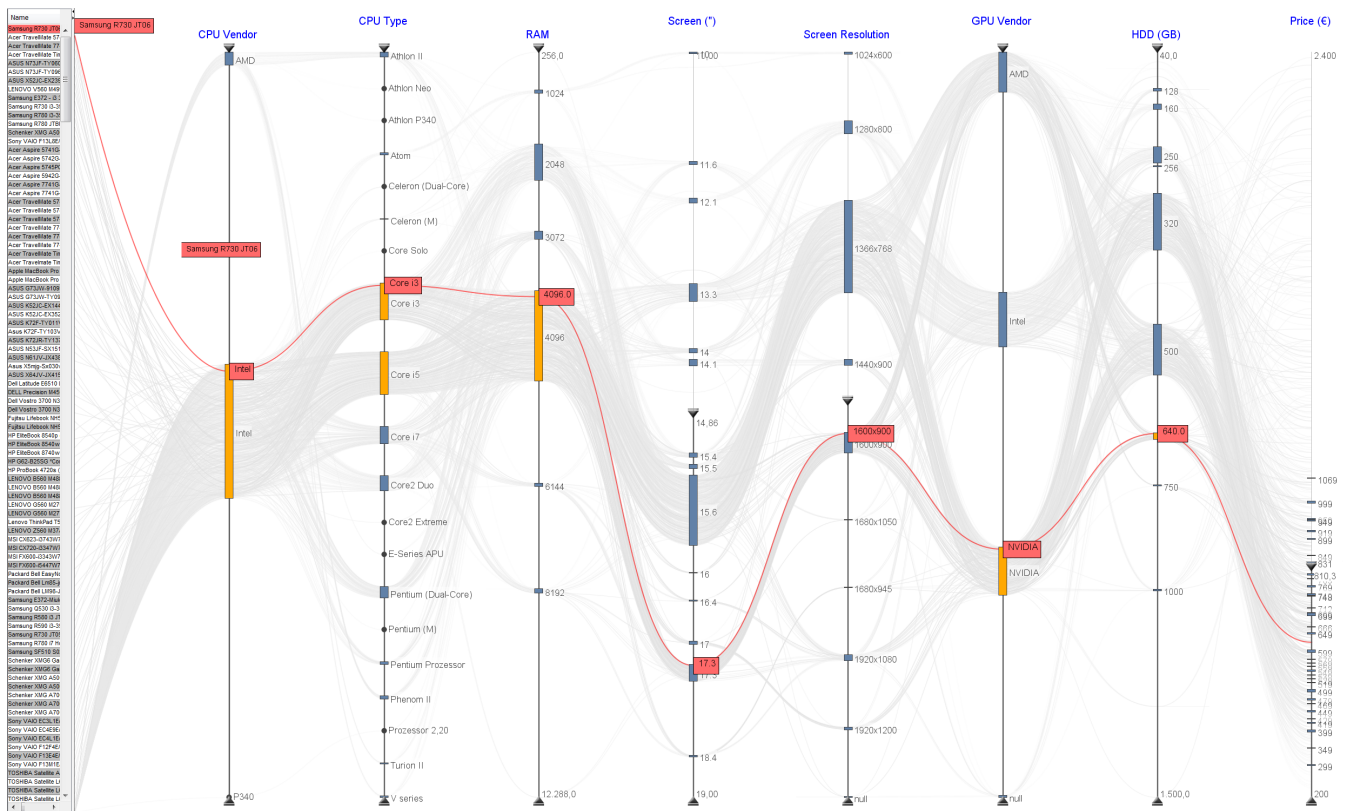


(b)

Figure 5: The visual query generation technique allows the users to express their requirements in a visual manner by selecting attribute values and by defining ranges. (a) A subset of the data is defined by selecting an Intel core i3 *OR* core i5 processor and 4 Gigabytes of RAM . All notebooks that match our constraints (and thus the query entirely) remain in blue, whilst the others are grayed out next. (b) Defining ranges for a 15 inch or larger display with a high resolution further refines the query. See next steps on Figure 6.




(c)




(d)

Figure 6: (c) To finalize the query the attributes for a dedicated graphics card, and for 640 Gigabytes harddisk size are selected. Four notebooks remain and the price may guide our final decision (c and d). If necessary, e.g. if the result set is empty or the remaining products are not desired, the user is able to relax or to adjust the criteria on each axis, either by deselecting an area, by adding another area or by moving the range sliders.




(a)




(b)

Figure 7: Exclusive decisions and the attribute repository: Only the most important axes are drawn at first. If the user is certain about a particular value for an attribute (e.g. Core i5; see (a) and (b), red rectangle), they can remove an axis by choosing an extended area of a particular axis (a, black pointer) or by selecting a range. There is no need to display this particular axis any longer and it disappears together with all paths that do not fulfill this decision. All previously chosen axes along with the selected values are depicted in the decision panel below ((a) and (b), red rectangles), so users can track the decisions they made for certain attributes. By making decisions attribute by attribute, every step releases some horizontal space which is automatically filled with newly appearing axes ((b), green pointer) that were originally of less interest but could be now important for a final decision ((b), green pointer). The axes of less importance are kept as miniatures in the attribute repository on the right. See next steps on Figure 8.



(c)



(d)

Figure 8: The miniatures can be dragged to the main drawing area and dropped at an arbitrary position where it then transforms to a regular axis which also becomes connected instantly((b),(c),(d) red pointer). The positions of all axes are readjusted according to the number of axes and the available space. If it is necessary to rethink a decision the user can drag one or more axes back into the main display area (d, purple pointer).