Experimental evaluation of semantic depth of field, a preattentive method for focus+ context visualization

Manfred Tscheligi
University of Salzburg

Johann Schrammel
AIT Austrian Institute of Technology

Silvia Miksch
TU Wien

Experimental Evaluation of Semantic Depth of Field, a Preattentive Method for Focus+Context Visualization

Verena Giller[1], Manfred Tscheligi[1], Johann Schrammel[1], Peter Fröhlich[1], Birgit Rabl[1], Robert Kosara[2], Silvia Miksch[2], and Helwig Hauser[3]

[1] CURE: Center for Usability Engineering, Austria, http://www.cure.at/

[2] Vienna University of Technology, Austria, http://www.asgaard.tuwien.ac.at/, { rkosara, silvia } @ asgaard.tuwien.ac.at

[3] VRVis Research Center, Austria, http://www.VRVis.at/vis/, mailto:Hauser@VRVis.at

Abstract

We introduce the Semantic Depth of Field (SDOF) technique, which is an alternative focus+context method for information visualisation. SDOF blurs objects which are currently out of focus, i.e., not interesting for the user. In the experimental study described in this paper, we found that SDOF supports the preattentive perception of sharp targets and the numerical estimation of the amount of data being in focus. We investigated the influence of distracting encodings (color and orientation), as well as threshold values for human perception. We also tested applications of SDOF (a textviewer and a scatterplot) with regard to their effectiveness in guiding the user's attention and with highlighting of data items.

Keywords: visualization, information visualization, focus+context visualization; blur, depth of field; preattentivity; experimental user study

1 Introduction

Information visualization (InfoViz) is the use of computer-supported, interactive, and visual representations of abstract data to facilitate cognition. The goal of InfoViz is to ease understanding, to promote a deeper comprehension of the data under investigation, and to foster new insights into the underlying processes. Because most data lack an inherent structure that could be directly understood as spatial, it is important to find a mapping of data dimensions to space, color, etc., so that the user can understand the data. It is also important not to change this mapping too often, because that requires the user to learn a new mapping (because his or her mental map [9] is destroyed).

When a lot of data is shown, it is necessary to be able to zoom in on certain parts to get more detail. At the same time however, the user must be able to retain an idea of where in the data the zoomed data is, so as to understand it in relation to the rest of the data. This is called a *focus+context* (F+C) approach, and has been the subject of a considerable amount of work in visualization (see the section 2 for an overview).

Semantic Depth of Field (SDOF) is an F+C technique that is based on an effect known from photography and cinematography, called *depth of field* (DOF). A lens or lens system does not depict all objects equally sharp, but only those points which are in a plane parallel to the film plane at a certain distance from the lens – all other objects are blurred [7].

We use the fact that the human eye is used to ignoring blurred objects (because the eye also has a limited depth of field) by blurring currently irrelevant objects. This way, the user's attention is guided to the most relevant objects without losing the (blurred and thus less prominent) context.

In this paper, we present the results of a user study that examined the following hypotheses: exploring if SDOF supported preattentive perception of sharp target items, the numerical estimation of presented data (values, amounts, etc.), work in combination with another coding as an orthogonal (additional) dimension, the guidance of the user's attention and the highlighting of data in a concrete application context. The study consisted of two parts in which the test participants were confronted with visual non-meaningful images (experimental study) and with meaningful application context tasks (evaluation of applications).

In the following section, we discuss some existing F+C methods and uses of blur in visualization. In the subsequent section, we present the idea of SDOF, describe the design of our study, discuss its results, and finally provide some conclusions and future plans.

2 Related Work

Most focus+context methods are *distortion-oriented* [3] (or, as we call them, *spatial*) ones, e.g. fisheye views or transformations into hyperbolic space [3]. These methods magnify important areas, while at the same time shrinking less relevant ones. This way, it is possible to put more objects into the same amount of screen space without losing the ability to recognize details for the more important ones.

The methods we call *dimensional* methods do not change the layout or space allocation of a visualization, but show different data where needed. The user moves a focus window over the visualization, inside which different or additional information is displayed for the same objects. It is thus possible to use fewer visual cues and display more information without cluttering the display. Examples for these methods are Magic Lenses and Toolglasses [2].

What we call *cue* methods, finally, are used to point the user to a part of the displayed information that is currently most relevant. An example of such a method is a GIS (graphical information system) visualization that allows the user to find hospitals, streets, etc. [8]. The result is pointed out by drawing it with a higher color saturation and brightness than the surrounding objects. With respect to this classification, SDOF is a cue method.

Blur has been little used in visualization, which seems to be mostly due to its computational resources with existing methods. One example is a multi-layer geographical visualization [4] that allows the user to select which layers to view by changing their blur level and transparency.

When appropriately used, graphical features such as shape, size, color, and position have proved to be effective in information visualization because they are mentally economical, and rapidly and efficiently processed by the *preattentive visual system* rather than with cognitive effort [1].

Preattentive processing has been surveyed in a limited scope regarding visualization techniques. Behavioral experiments [9] indicate that accurate target and boundary detection is possible with both *hue and curvature*.

Furthemore, preattentive and interpretative perceptual properties have been tested on *motion (animation)* for complex information visualization [l].

3 Semantic Depth of Field (SDOF)

The basic idea behind SDOF is to point out relevant objects in a visualization by using blur. This is useful whenever the user needs to select between groups of objects, find objects that meet certain criteria, or get an overview of a data set by querying it in different ways.

Relevance

SDOF is based on the notion of relevance, which is a measure for the current importance of an object (relevance is also used in other F+C methods, but not explicitly). The relevance of objects is calculated by the application for every object, and the user has to be able to change the used relevance function at any time. The relevance r of an object is a value between 0 and 1, where 1 denotes an object of maximum relevance, and 0 an object that is currently completely irrelevant.

Blurring

The relevance r is translated to a blur factor b, which describes the amount of blur in terms of a factor. A b-value of 1 produces a perfectly sharp image, while any larger value means that every point in an object is enlarged by this factor (i.e, is spread over a so-called *circle of confusion* [7]). On a computer display, this means that the image is drawn and then blurred over b pixels in each direction.

The mapping of r- to b-values is done by means of the blur function. Any function can be used here, but we found a linear function to be sufficient that allowed for a threshold to be set up to which no blurring takes place at all, and that then steps up to a value where blurring can be easily perceived.

In a practical application, the blur level is not set as a number, but by showing the user sharp and blurred objects from the application, and letting her or him choose from those. This also makes it possible to adapt to different uses that require different settings – like screens, projection displays, etc. (see the section on Output Sensitivity for details).

Implementation

Blurring is an operation that is inherently slow on current computer architectures. In order to make SDOF usable in real applications, we developed a method that makes use of texture mapping hardware found on modern consumer 3D graphics cards to provide fast blurring of any object [5]. The method uses texture mapping to draw several slightly displaced copies of the image over each other. This way, it is possible to achieve high frame rates of up to 80 frames per second (and always at least close to 20 fps), even when a large part of the display is blurred [6].

Properties

SDOF does not change the layout of a visualization – like distortion-oriented methods do. It also differs from other methods in that it changes the appearance of irrelevant objects, rather than of relevant ones (which are magnified, for example). It is therefore useful whenever objects contain a lot of information that would be harder to read if the relevant objects were changed.

SDOF is also independent of color. It does not change the color or saturation of objects, and can thus be used as an additional cue to color, and also be used by color-blind users.

SDOF is intuitive. Because users know depth of field from photographs and movies, and because the human perceptual system is used to masking out blurred parts of the visual field, users immediately understand which objects are the relevant ones. As will be shown below, this distinction can be made preattentively and therefore provides a very efficient means of conveying information.

Figure 1: A screenshot of lesSDOF

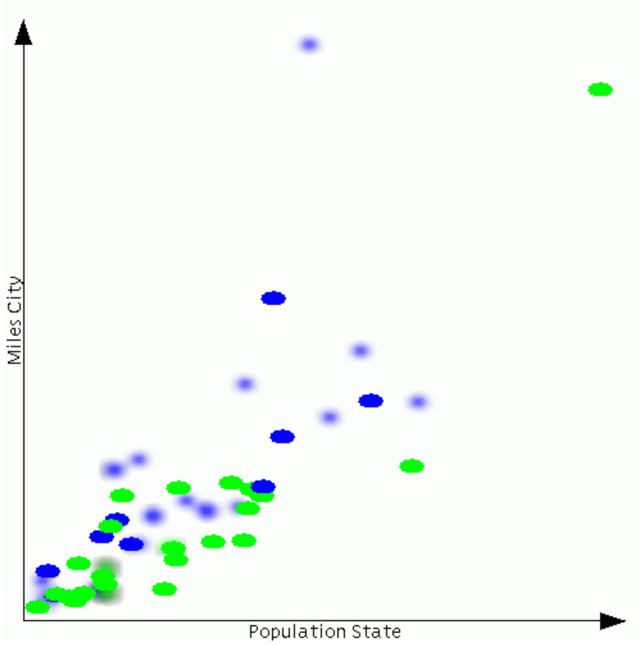

Figure 2: An image of sScatter

Distortion-oriented methods create side-effects where objects that are less relevant are magnified because they are close to a relevant object, or because they are in between areas of relevance. SDOF allows direct control of every single object.

Output Sensitivity

Blur (and other techniques in computer graphics, like photo-realistic rendering) is inherently dependent on the viewing circumstances. A blur disc is perceived as a sharp point if it is too small for the eye of the viewer to be seen as more than a point. Therefore, the magnification of an image and also the distance of the viewer play a crucial role for SDOF (this problem is also known in photography). This makes SDOF hard to use when people are viewing an image from many different distances (e.g., in a large audience). For this reason, an SDOF application must enable the user to adjust the parameters in the blur function for every session, so that these differences can be accounted for.

Applications

A number of application prototypes have already been implemented. Here, we only describe the two that were used in the user study.

lesSDOF: text display and keyword search – many applications provide functions to search text for keywords. Examples for this are word processors, web browsers, and simple text viewing programs like the UNIX program *less*. These programs usually show the identified word by inverting it (i.e., dark background, light text). But the user does not get any help in finding the relevant context for the word, and has to search for the borders of the containing sentence.

lesSDOF not only searches for the keyword, but also determines the surrounding sentence by looking for punctuation marks. It can show the context it has found this way using three different methods. The keyword is always inverted, and in the simplest mode, the rest of the page is simply displayed as usual. In *gray* mode, the sentence containing the keyword is displayed with a light gray background, and is thus quite easy to see. And in *blur* mode (Fig. 1), the containing sentence is displayed sharply, while the rest of the page is blurred.

The user interacts with the program by entering a keyword and then using the cursor up and down keys to navigate between hits.

sScatter: scatter-plots – in scatter-plots, it is often difficult to display more than two variables. When adding color, orientation, or other cues, it is possible to extend this number. Blur is probably a valuable addition to this as an additional and effective visual cue.

sScatter is a program that draws scatter plots from arbitrary data sets. The user can select which data dimensions to map to which visual attributes – x-position, y-position, color, orientation, size, shape, transparency, and blur. It is thus easily possible to find certain data points which have a certain set of features (Fig. 2).

4 User Study

An in-depth study was carried out in order to verify (or falsify) both the principal theory which SDOF is built upon, and specific questions in an application context. It provides interesting results and allows insights for the use of SDOF as a visualization technique, which can be transferred for other domains, applications, and techniques.

Research Questions

The study was targeted at a defined set of hypotheses concerning the SDOF visualization method, which assess the two issues just mentioned above. The following questions summarize the hypotheses:

Does SDOF (1) support preattentive perception of sharp target items, (2) support the numerical estimation of presented data (values, amounts, etc.), (3) work in combination with other visual codings as an orthogonal (additional) dimension, (4) support the guidance of the user's attention in the context of a concrete application, and (5) support the highlighting of data in the context of a concrete application?

Participants

We chose to test a very narrow defined user group, so we would minimize any (inter- or intra-personal) bias and gain most valid results. The user group therefore consisted of 18 male students (including 2 pre-tests), all with extensive Internet and computer experience, aged between 18 and 25 years and with very good eye sight. The latter was defined as a requirement for two reasons: Firstly, because the study focuses on issues of visual perception and bad eye sight could cause a critical bias. And secondly, because our eye tracking system works best for users without lenses or glasses and we included eye-tracking observations.

Test Plan of User Study

The underlying concept for this user study was to confront the test participants with numerous images – one at a time and for most tasks only for a period of less than 200 ms – and then make him answer a question. The 200 ms were used because this is the time span in which it is widely agreed that preattentive perception takes place [10]. The analysis and interpretation of the samples' answers was then based on the correctness of the answer as well as on the response time.

The overall test plan was composed into two phases in which the test participants were confronted with (1) visual non-meaningful images (experimental study) as well as with (2) meaningful tasks in an application context (evaluation of applications). Each of both phases was split into multiple series which are listed below:

Experimental study – series 1: target detection;
series 2: count estimation; series 3: interplay;
series 4: perceptual thresholds.

Evaluation of applications – series 5: lesSDOF;
series 6: sScatter.

For each test participant, the test started with a general briefing and the setup of the eye tracking system. Then all series were carried out sequentially (but in different order), and finally the test was finished with qualitative questions. It is important to mention, that the sequence of the images within a series was randomly generated in order to prevent any sequence-based bias. The test series were changing systematically for every participant. Each test took approximately two hours – including approximately 15 minutes of regeneration time for the test participant, which was consumed in one to two minute slots (mostly after a set of images/tasks was finished). The study was carried out in the Usability Engineering Labs in Vienna (http://www.cure.at/) and took place within a time span of 14 days in August 2001.

Test Materials

In order to have images which fully allow valid interpretations, comparative analysis, and images which avoided all thinkable bias, we generated approximately 2000 images which were checked by hand for overlaps and targets too close to the quadrant borders (see below) before being included in the test. The images mainly consisted of ellipses in various numbers, colors, blur levels, and orientations as well as combinations of those. We chose ellipses as items for two reasons: we needed an item which (1) could be rotated, changing its visual appearance accordingly (e.g., a circle remains identical) and (2) could be blurred without resulting in any unplanned visual misconceptions (e.g., due to changes in shape). Images are described in closer detail in each series. In addition, two applications were developed (lesSDOF and sScatter) and respective data sets were located and texts were produced.

Technical Environment

In order to provide an image presentation of less than 200 ms, we had to use an LCD screen (15" TFT, 1024x768, 16.8 million colors, 75 Hz), because any other screen type would have caused a serious bias due to its line-by-line refresh. The computer was a Dual-Celeron 433 MHz with 128 MB RAM; the input (user feedback) was given with a mouse and the space-bar of a keyboard. To support the analysis of the gathered data we automatically logged time on task, reaction times, location of mouse clicks and any other system or user behavior which we required.

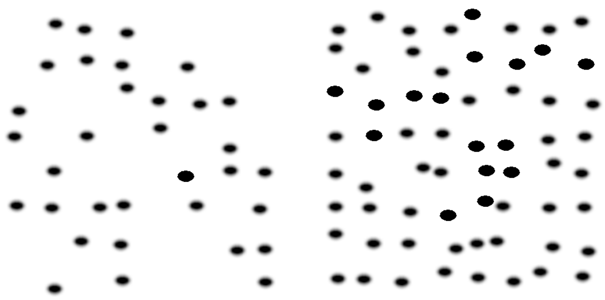

(a) (b)

Figure 3: Images from series 1 & 2: (a) locate the sharp object; (b) estimate the amount of sharp objects

5 Experimental Evaluation

During the first phase of this user study the basic theory which SDOF is based upon was tested. In total, four series were tested with each of the test participants. In the following details about this phase of the used study are presented.

5.1 Series 1 & 2: Target Detection & Count Estimation

The first two series were designed to test the preattentiveness for the two tasks target location (series 1) and count estimation (series 2).

Tasks

For target location, the test participants had to identify the position of a sharp target within an array of blurred objects (see Fig. 3(a) for an example). To test the ability of count estimation, the test participants had to estimate the number of sharp objects in three categories (few, medium, many; see Fig. 3(b) for an example). For both tasks, we used the same interaction procedure and screen design, only the respond screen was slightly modified to suit the different needs.

Screen Design and Procedure

The test images were presented on a centered and 500x500 pixels wide array (white background) within a gray frame, the transition between array and frame was smoothened to ensure that no interfering perception effects were produced by sharp edges.

The test screen appeared for 200 ms, after which the test participants had to provide their response by mouse click on an respond screen.

After answering a question, the test participants had to press the space bar to get the next image. After a delay of 300 ms, the next image was shown for 200 ms within the test area and was then automatically replaced by the respond screen. Entering the answer caused a blank gray screen to appear and the whole procedure was started anew by pressing the space bar.

Materials

In general, the test screens consisted of black ellipses arranged in an 8x8 array where every item was placed randomly in direction and distance, but restricted to positions that would not produce overlaps.

For target detection, the screens showed 3, 32, or 64 items placed randomly in the array. The used blur levels for the non-target objects consisted of 3 different grades (7, 11, and 15) and their combinations (resulting in 7 classes).

For each combination of the steps of the two factors, 60 images were created, 30 including a sharp target object, 30 without. From these 30 images, 5 were chosen randomly for each test participant and the total of 210 images was presented to the participant in random order.

For count estimation, we always used a fully occupied array, i.e., 64 items, and only used one blur level per image. We produced test images for every blur level and varied the number of sharp items from 5% to 95% in 5% increments.

Results

The results of this series significantly prove that the degree of blur has an effect on the visual perception of target items. In both series 1 and 2, the test participants had (statistically proven) significantly higher correct answering rates when the blurred visual context of the target item exceeded the lowest, least blurred visual presentation. This lowest blur level showed to be insufficiently visually differentiated from the sharp target item (see also series 4 on thresholds). The correct answers in series 1 (position/quadrant) where 94% for a strongly blurred context versus 64% for the insufficiently blurred context. For existence of targets (without location), the correct answers were similar (95% versus 73%).

The numbers of items on the screen had a significant effect on the correct answering behavior. Concerning position detection, the number of items and the percentage of correct answers was 88% for 3 objects, 73% for 32, and 70% for 64 objects. For existence, the respective results were 89% for 3 objects, 80% for 32, and 77% for 64 objects.

A variance analysis for both effects (blur level, number of items; dependent variable: position) as well as their interaction effects supported these results (blur level: $F(6,90) = 83.1$ at $p < 0.001$; number of items: $F(2,30) = 31.7$ at $p < 0.001$; interaction $F(12,180) = 6.0$ at $p < 0.001$). A Scheffé test concerning the blur level proved that the effects were significant (critical difference: 9%) between the resulting groups (context with least, i.e., non-differentiable blur level yes/no) but not within these groups. An analogue

Scheffé test for the number of existing items showed a significant difference between 3 versus 32 or 64 items. There was no significant difference between the 32 and 64 items.

For series 2 on count estimation the results prove that it is possible to accurately and quickly estimate the share of sharp objects among blurred ones. Taking estimation errors independent from coding into account, the values are quite high: 88% of the estimations were correct when using the highest blur level, 85% (mid), and 74% for the lowest one. A variance analysis showed significant differences for the lowest blur level compared to both, mid and high blur: $F(2, 24) = 16.2$ at $p < 0.001$.

5.2 Series 3: Interplay

In this series we were interested in the interplay effects of combining different coding methods such as SDOF, color and orientation and their impact on the speed of perception.

Tasks

The main task for all types of search was to locate the target as fast and as correctly as possible. We decided to test three types of search tasks: *simple*, *disjunctive*, and *conjunctive*, where

- *simple* means that the target is defined by one coding dimension and no other dimensions are present – e.g., find the red item within the black ones,
- *disjunctive* means that the target is defined by one coding dimension but a second, ”meaningless” though potentially distracting dimension is present – e.g., find the red item within the black ones where both the target and the distractors could be sharp or blurred, and
- *conjunctive* means that the target is defined by the combination of two coding dimensions (Fig. 4) – e.g. find the red and sharp item within items that could be red-blurred, black-sharp, black-blurred.

Screen Design and Procedure

In terms of laying out contents on the screen, series 3 was identical to series 1 and 2. Due to the different problem tested, we modified the procedure known from series one insofar as the program-controlled disappearing of the test pictures after 200 ms was removed and the test person had to press the space bar to continue to the respond screen. The subjects were requested to do so as soon as they were able to give the right answer.

The reaction time was measured from the appearing of the image until the space bar was pressed and used as the main dependent variable. Also the respond screen was modified by removing the button ”not able to locate”.

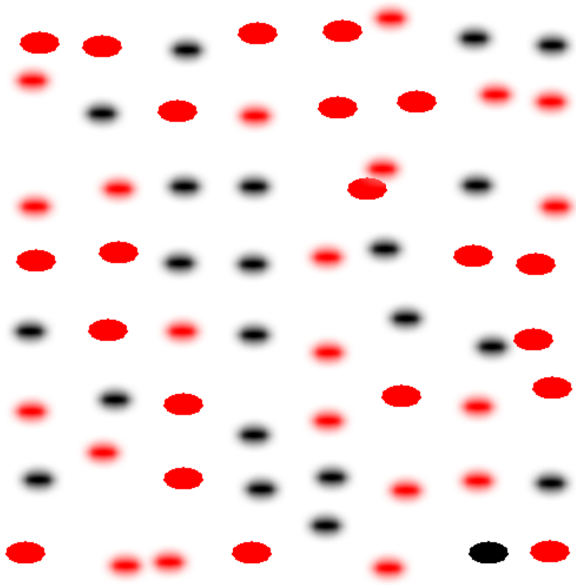

Figure 4: An image from series 3, conjunctive search: Find the black sharp object.

Materials

Due to time constraints, we only tested the highest blur level and used a fully occupied array similar to the other series. For the different search tasks, we produced the following test pictures:

Simple Search: • all distractors were black & blurred, the target black & sharp; • all distractors were red & blurred, the target red & sharp; • all distractors are black & sharp, the target red & sharp; • all distractors were black & blurred, the target red & blurred;

Disjunctive Search: • the sharp object had to be located: blurred distractors were selected randomly so that 50% were red and 50% black, • the red object had to be located: half of the distractors were blurred, the other half sharp;

Conjunctive Search: distractors were arranged along 3 further combinations of • color and sharpness (for targets red & blurred, red & sharp), • color and orientation (for targets black & rotated, black & normal), • blurring and orientation (for targets normal & sharp), and orientation and sharpness (for targets rotated & sharp);

Results

This series showed statistically significant effects for the conjunctive search tasks, which each proved to be dif-

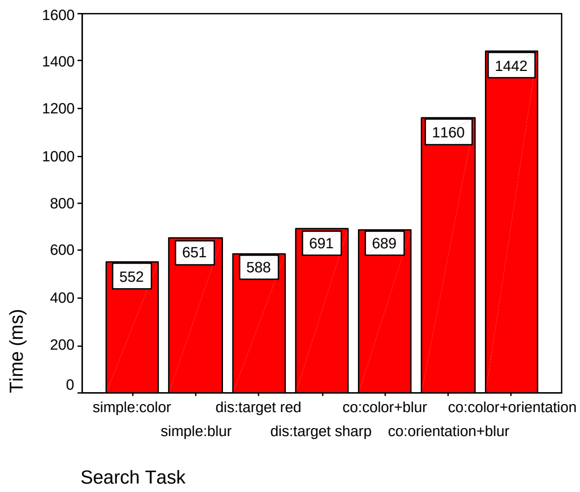

Figure 5: The results of series 3, by type of task.

ferent from all others (variance analysis $F(6, 72) = 48.3$ at $p < 0.001$; Scheffé test: critical difference 250.8).

An interesting result was that conjunctive search for color and blur cannot be statistically differentiated from any of the disjunctive or simple search tasks. It seems also notable that there is no significant difference between color and blur coding, although blur shows a higher average perception/answering time. See also Fig. 5 for a comparison of the results of series 3.

5.3 Series 4: Measurement of Threshold Value

Considering the use of SDOF in practical applications, we were interested in identifying perceptual thresholds as well as the correlation between technical blur value and perceived blur to help with deciding which blur levels to use.

Tasks

The participants were asked to state

a) if two items presented at the same time were perceived different or equal concerning their blur level,

b) if one object presented on its own was perceived as sharp or blurred, or

c) a numerical ratio of two objects shown side by side on the screen which best fits their perception of the blur levels.

Screen Design and Procedure

In this series, we used a screen design similar to the one already known. The screen showed one (centered) or two ellipses (distance: 80 pixels) within the test array.

For a), we presented the items in three varying time series sorted from highly blurred to sharp, the other way round, and in random order. The blur level of the second object was increased/decreased until a difference was subjectively perceived. To avoid adulterations, this in-/decreasing took place after one to five randomly chosen loops, i.e., the test persons wouldn't be able to guess the correct answer.

For b), a single object was presented in the center of the test area and the blur level was in-/decreased continuously until the right answer was given. Similar to a), there were random repetitions.

In c), every possible combination of the chosen blur levels (see below) were presented in random order and the participants had to state the perceived ratio verbally.

Materials

We used the already well-known ellipses blurred in a range of 1 to 19. For c), we restricted the comparison to 5 blur levels: 1, 5, 9, 11, and 15.

Results

The assessment of the test participants' subjective estimations and perception of thresholds and differences for varying blur levels provided valuable results. When test participants had to tell if a presented item was regarded as focused (sharp), the test participants showed different average thresholds depending on whether the presented item series was sorted from highly blurred to focused (average blur level 3.2) or the other way around (average blur level 1.6). This result is consistent with our expectation, considering that test participants mentally compared the respective predecessor item with the currently presented one. Another very interesting result is the observation of test participants' perception of visual differences of blur levels and the estimation of their ratio. These estimations were highly correlated to the actual mathematical ratio on which the blur levels were determined[1].

6 Evaluation of Applications

During the second phase of this user study the application of SDOF in real-world applications was tested. Since two prototype applications have been developed for this study, two series were tested with each of the test participants.

[1] At this point, an observation of test participants' subjective behavior and comments shall be mentioned: test participants found the tasks of comparing blur levels over a longer period (approx. 15 mins.) very cumbersome and annoying. These problems seem to be caused by the requirement to concentrate on the blur effect as such, in contrast to tasks where sharp objects were carrying the meaning, which obviously makes more sense.

6.1 Series 5: lesSDOF

In series 5, we tested the use of SDOF in relation to other methods in a word processor to support the users in finding relevant information.

Tasks

Main task was to answer questions, where the required information was to be found in the context, i.e., the sentence surrounding the given search item.

Screen Design and Procedure

The text display program (Fig. 1) was displayed centered on the screen; due to its smaller size a part of the desktop was visible. We used a sans-serif font (Verdana) and a font size of 11 for displaying the text.

The participant was given the task on a piece of paper, and had to enter the keyword to search for into the respective text field in the application. The screen displayed the text, and the participant used the arrow keys to navigate between hits. As soon as the answer was found, the spacebar had to be pressed. The answer was given orally. The time was recorded from the submission of the search until pressing of the space bar.

The tasks were performed with five different questions per version. We used three different versions:

a) The search item was highlighted (marked in black), the text was in conventional presentation mode.

b) The search item was highlighted (marked in black), the context is sharp, the remaining text was blurred.

c) The search item was highlighted (marked in black), the context was highlighted (marked in gray), the remaining text was presented in conventional presentation mode.

Materials

We used 15 articles downloaded from the Internet concerning subjects such as natural sciences (simple texts) or biographies. The length of the articles varied between 7 and 20 pages and there were five to seven keyword hits in each. To even out the influence of the text and searched information, we used every text for every version nested within all the test participants.

Results

The average search times for the different presentations of the target word's context were as follows: 11346 ms (context blurred), 11187 ms (context colored), and 14027 ms (without any highlighting). These values are not statistically significant (variance analysis $F(2, 30) = 1.77$ at $p = 0.19$), but the average search time without highlighting the context is higher than when the context is highlighted in either of the two ways.

These results show that in a text display program (such as lesSDOF), the blur effect works just as well as other highlighting when it comes to visually finding a searched word. But considering that blurring text the user has more effort reading it, we have to realize that SDOF does not provide (additional) user benefits.

6.2 Series 6: sScatter

The second application we tested was a scatter plot, which is able to show data in different representations (Fig. 2). We decided to compare SDOF coding with orientation due to the comparability of the character of these methods. Both seem to be suitable as an additional coding dimension but not capable of many graduations.

Tasks

According to our hypotheses from the experimental series, we think SDOF supports a better first overview and interpretation of the data and their distribution; so we selected finding the center of gravity of the target items as the task for this series.

We tested two different constellations in 2 overall tasks. One was to allocate the center of gravity for a) sharp items and b) rotated $45°$ items. The second one was to find the center of gravity of all items with a specific combination of a) sharpness and color and b) $45°$ rotation and color. Each task was repeated five times with different data distributions.

Screen Design and Procedure

As screen layout we used a white background of 500x500 pixels size. Ellipses were ploted in colors green or blue at positions according to a 2D coordinate system with its origin in the lower left corner of the display. Axes were shown together with text labels representing the data-to-axis mapping in use. See Fig. 2 for an example.

The test persons were instructed to the function and coding scheme of the scatter plot. Afterwards, the test participants had to identify the center of gravity in the experimental trials. The answer was given by clicking on the appropriate area on the screen. The series was designed nested, i.e., half of the test persons started with orientation coding and the other half with blur coding to even out learning effects.

Materials

The different test pictures were based on the same data but we did change the meaning of the displayed axes so no two pictures were similar, yet the total number of items remained unchanged. Additional to interval data there were two binary variables which were color coded and SDOF/orientation coded.

Results

The SDOF method proved to be very valuable for the perception of data and information within a scatter plot. The results show that the coding method of the target items (blur versus orientation) had a significant positive effect on the accuracy (distance to mathematically determined balance point) with which the test participants solved the tasks. It also had a positive effect on the time it took the test participants to solve the task. The accuracy differed as follows: for blur-coding of the contextual items, the average distance from the mathematical balance point was 25.5 pixel points, whereas for orientation-coded items it was 39.0 pixel points. Comparing the time it took to solve the task, blur-coding of the context was significantly better (average 4822 ms) than orientation-coding (6833 ms).

The quality of these results is supported by the fact that accuracy and time on tasks do not correlate with each other (Pearson, $r = 0.074$; $p = 0.632$). This shows that the factors are independent of each other, which is an important result: it would have been possible that a higher accuracy is achieved by longer time on task – but this is not the case.

These results verify the hypothesis that blurring the contextual items within a scatter plot supports the user in better perceiving the relevant (sharp) item's characteristic distribution and center of gravity.

7 Discussion of Results

With this study, we are able to prove that the SDOF concept really is preattentive, that it directly supports the perception of sharp target items when the context is blurred. SDOF can significantly support users in focusing on relevant data and guide their attention.

We also verified that the estimation of the number of sharp objects within blurred ones is possible very accurately with one glance (200 ms).

The user study also proved that blurring can be orthogonally combined with color coding, and such conjunctively coded items are still similarly fast to perceive as single coded items. This very valuable result shows that SDOF can provide additional semantic value to a visualization without increasing cognitive effort. The application context evaluations showed that – applied correctly – SDOF can support specific tasks within an application context, such as highlighting relevant data or search results. This works primarily by guiding the user's attention.

Directly regarding the blurring effect, we found that users are able to clearly distinguish 2 to 3 different blur levels, although it clearly showed that within one image, visualization, or application, only one blur level should occur simultaneously.

8 Conclusions and Future Plans

SDOF has proven to be a valuable concept worth investing more effort into to explore its potentials, application fields and extended user groups. After incorporating the results of this study into the technique, we want to carry out further studies, including broader user groups as well as extended image and application complexity.

We also want to improve the usefulness of SDOF based on the results of the user study. It has shown that SDOF is not suitable as an additional fully-fledged visualization dimension, but rather as a distinction between groups of objects. We therefore want to introduce secondary cues (similar to motion blur, for example) to improve this.

Acknowledgements

Parts of this work have been done as part of the basic research on visualization (http://www.VRVis.at/vis/) at the VRVis research center in Vienna, Austria, which is funded by an Austrian governmental research program called K plus. Parts of this work also have been funded by the Institute of Software Technology (http://www.ifs.tuwien.ac.at/) at Vienna University of Technology, Austria. This work was also supported by the Austrian Science Fund (FWF, http://www.fwf.ac.at/, grant P12797-INF). Finally, this work was also supported by the Center of Usability Engineering (CURE, http://www.cure.at/) in Vienna, Austria.

References

[1] L. Bartram. Perceptual and interpretative properties of motion for information visualization. In *Proc. of the Workshop on New Paradigms in Information Visualization and Manipulation*, 1997.

[2] E. Bier, M. Stone, K. Pier, W. Buxton, and T. DeRose. Toolglass and magic lenses: The see-through interface. *Proc. of SIGGRAPH '93*, pages 73–80, 1993.

[3] S. Card, J. Mackinlay, and B. Shneiderman, editors. *Readings in Information Visualization*. Morgan Kaufman Publishers, San Francisco, CA, 1999.

[4] G. Colby and L. Scholl. Transparency and blur as selective cues for complex visual information. *SPIE Vol. 1460, Image Handling and Reproduction Systems Integration*, pages 114–125, 1991.

[5] R. Kosara, S. Miksch, and H. Hauser. Semantic depth of field. In *Proc. of the IEEE Symp. on Information Visualization 2001 (InfoViz 2001)*, San Diego, CA, 2001.

[6] R. Kosara, S. Miksch, and H. Hauser. Focus and context taken literally. *IEEE Computer Graphics & Applications (IEEE CG&A)*, 2002.

[7] H.-C. Lee. Review of image-blur models in a photographic system using principles of optics. *Optical Engineering*, 29(5):405–421, May 1990.

[8] I. Lokuge and S. Ishizaki. Geospace: An interactive visualization system for exploring complex information spaces. In *Proc. of CHI '95*. ACM, 1995.

[9] K. Misue, P. Eades, W. Lai, and K. Sugiyama. Layout adjustment and the mental map. *Journal of Visual Languages and Computing*, 6(2):183–210, June 1995.

[10] A. Treisman. Preattentive processing in vision. *Computer Vision, Graphics, and Image Processing*, 31:156–177, 1985.