

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

Johann Schrammel Verena Giller Manfred Tscheligi

Cure - Center for Usability Research and Engineering
Vienna, Austria

Robert Kosara Helwig Hauser

VRVis Research Center
Vienna, Austria

Silvia Miksch

Vienna University of Technology
Vienna, Austria

Abstract: This paper presents the results of a thorough user study that was performed to assess special features and the general usefulness of Semantic Depth of Field (SDOF). Semantic Depth of Field is a focus+context (F+C) technique that uses blur to point the user to the most relevant objects. SDOF was found to be an effective means for guiding the viewer's attention and for giving him or her a quick overview of a data set.

Keywords: Information Visualization, Focus+Context, Empirical Evaluation

1 Introduction

SDOF (Kosara et al, 2001, 2002) is a focus+context (F+C) technique that uses selective blur to make less important objects less prominent, and thus point out the more relevant parts of the display to the user. It is based on the depth of field (DOF) effect known from photography and cinematography (Lee, 1990), which depicts objects sharply or blurred depending on their distance from the lens. SDOF extends this effect and displays an object sharply or blurred based on the object's current relevance.

Blur is measured as the diameter of a circle over which the information from one pixel is spread when it is blurred. Thus, a blur diameter of 1 means a perfectly sharp image, with larger values creating more and more blurred depictions.

The overall goal of the study was to find out if SDOF is an effective means of guiding the user's attention, and if it supports the user in applications.

Effectiveness was assessed by testing the ability to preattentively perceive sharp objects and by comparing search times for different cues.

Application tests were done with two different applications: a text viewer and a map viewer.

The sample size was 16 individuals, which we recruited from different universities in Vienna. Each test session took about two hours.

We used chi-square tests and ANOVA for significance testing, and Scheffé tests for post-hoc analyses where needed. All results that are described as significant in this paper were tested for with a probability for error of $p < 0.05$.

2 Preattentivity

Preattentive processes take place within about 200 ms after a stimulus is presented (Bartram, 1997, Healey et al, 1999, Treisman, 1985) and are performed in parallel, without the need for serial search. Such processes involve a limited set of features (e.g., orientation, color, etc.) for which certain tasks (e.g., detection, location, etc.) can be performed without effort.

We studied two preattentive abilities: being able to detect and locate a sharp object, and being able to estimate the percentage of targets among distractors.




Figure 1: Sample image for target detection and location with 32 distractors of the highest blur level, and a target.

The images for target detection and location showed ellipses which were scattered over the image (Figure 1). The reason for choosing ellipses was that we needed objects that would not change their shape significantly when blurred to rule out shape perception effects. Ellipses don't change their shape (in contrast to e.g. rectangles, which look more elliptical the more they are blurred), and they can also be rotated (which was needed in the interplay trial). Participants were shown images with 3, 32, or 63 distractors, with or without a target (50% with, 50% without a target) and one of the seven combinations of three different blur levels (7, 11, and 15 pixels) resulting in 42 different combinations. For each combination, each participant was shown five images (i.e. 210 total).

The test procedure consisted of four steps: First, an empty screen was shown for 300 ms, followed by the image, which was shown for 200 ms. After that, an answer screen was presented, which gave the participant the choice between clicking on one of four quadrants or buttons for “no target” and “target not locatable”.

For percentage estimation, the sequence was identical, except that the answer screen contained only three buttons for the estimated number of targets: “few” (up to 19 targets), “intermediate” (20 to 45) and “many” (more than 45 targets). The images shown in this trial only used one blur level per image, and always contained 64 objects, with 5% to 95% of objects being targets (in steps of 10%).

Finding sharp targets among blurred distractors is performed preattentively. Figure 2 shows the accuracies for correct location of targets, which were very high (> 90%) or high (> 60%) depending on the blur level. When the lowest blur level (7 pixels)

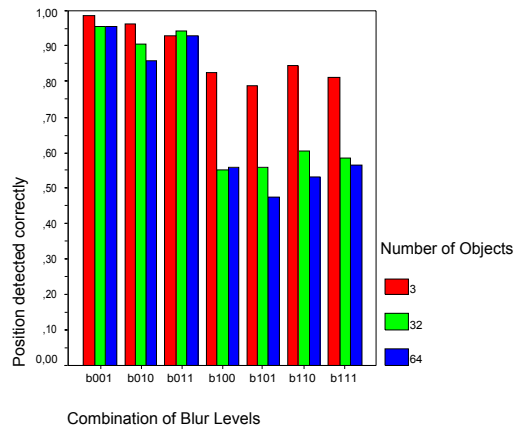


Figure 2: Correct answers by blur level and number of objects. The labels under the chart encode which blur levels were present (increasing left to right).

was present, the accuracy dropped significantly – this is most likely due to the fact that participants were not able to differentiate between sharp and slightly blurred objects. There is also a significant difference in accuracy between the cases with three distractors and those with 32 or 64, which was to be expected. Accuracies were almost identical for cases with and without targets, only for the case with only the smallest blur level present, it was much higher in the no-target case. We also presume this to be a result of subjects mistaking blurred objects for sharp ones.

Estimation of the percentage of sharp objects can also be done preattentively. The accuracy for all blur levels is significantly better than chance and does not differ significantly between the different blur levels.

4 Interplay

SDOF will very likely not be used without any other visual cues. Therefore we also studied its interaction with other features.

Images similar to the ones used for the preattentiveness test were used, with the additional features color (red or black) and orientation (main axis of ellipsis horizontal or at 45°).

The user interaction was slightly modified compared to the preattentivity test. This time subjects could look at the image as long as they wanted to find the answer – they were, of course, encouraged to answer as quickly as possible.

We tested simple, disjunctive, and conjunctive searches. Simple searches are based on the presence of one feature in the target, with the distractors not being different from one another.

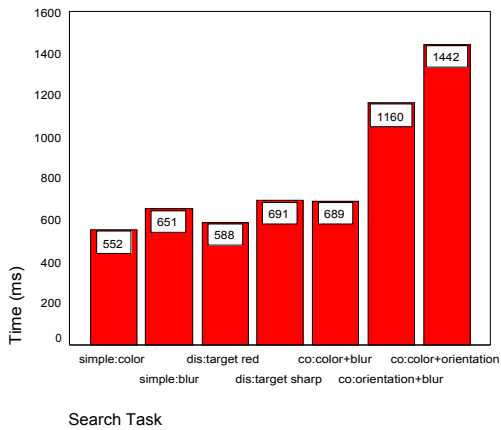


Figure 3: Time needed for search by search task (“simple”: simple search; “dis”: disjunctive search; “con”: conjunctive search).

In a disjunctive search, the subjects looked for one feature in the targets, but the distractors could also differ in another one (e.g., if the red object is the target, all distractors were black but could be sharp or blurred).

Conjunctive searches required the participant to look for a combination of two features in the targets (e.g., the red and sharp object), while the distractors could have any other combination of the two.

In terms of search time (Figure 3), SDOF is not significantly different from color - we could not verify a difference between simple searches for colored or for sharp objects.

The conjunctive searches for targets marked by color & blur, orientation & blur, and color & orientation differ significantly from each other, with color & orientation being the slowest – each of these two features combined with blur is faster.

Also, the conjunctive search for color & blur coded targets is not significantly slower than the simple and disjunctive searches, which is quite contrary to what we expected, because conjunctive searches usually are noticeably slower (Treisman, 1985).

5 Blur thresholds

To use SDOF as a separate visualization dimension we need to know threshold values such as the minimal difference in blur that can be perceived.

First, we tested the ability to tell whether or not two objects had the same blur level. For this, we showed the subjects two objects next to each other with equal or different blur. Subjects had to decide whether the blur was equal or different – if they

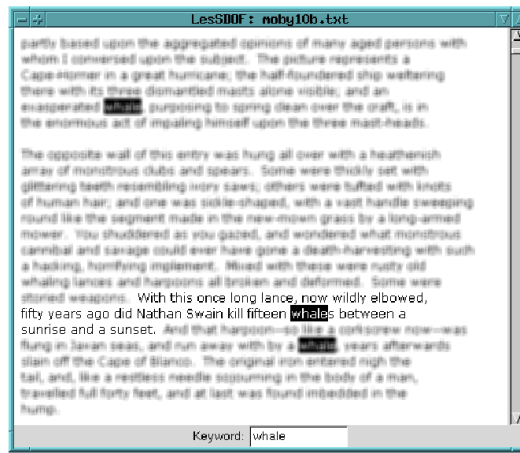


Figure 4: Screenshot of the text display application.

decided it was equal, the blur of one of the objects was increased, and the objects were shown again.

Second, we tested for the absolute thresholds of blur perception, by showing just one object, which was sharp in the beginning and got increasingly blurred until the participant judged it as blurred. This test was also performed starting with a strongly blurred object that got increasingly sharper.

Participants were able to tell the difference between objects of different blur levels with a good accuracy. The average difference in blur needed to tell the blur levels apart is quite small (less than 1.8 pixels for all blur levels).

In terms of absolute values a blur diameter of 3.27 (on average) was already judged a sharp object, when the participant was presented a very blurred object that got sharper; but a blur level of only 1.46 was judged as blurred when starting out with a sharp object.

6 SDOF-enhanced Text Display

The first application we tested was a text viewer (Figure 4). When searching for a keyword, it not only highlights the found keyword, but can also display the immediate context (containing sentence) sharply, while the rest of the page is blurred. To be able to compare the performances, we also included a mode that only highlighted the keyword, with no additional context information; and one which highlighted the immediate context by putting it on a light gray background, while the rest of the page was displayed on white (called gray mode).

We asked participants to answer questions about selected short texts. To do this, they had to search for a specified key word, in the context of which the answer could be found.




Figure 5: Screenshot of the map viewer (without control elements).

Search times were very similar for blur and gray, and not significantly (but still slightly) better than without highlighting (approximately 11 seconds on average for blur and gray vs. approximately 14 seconds for no highlighting).

When asked to rank which method they preferred, participants ranked blur and gray equally, and no highlighting significantly worse.

The conclusion from this part of the trial is that while the method is perceived as useful by some participants, it is not measurably better. A few participants in contrary complained that they found blurred text very annoying.

7 MapViewer

The MapViewer application displays layered maps in different ways (Figure 5). One is to show the “topmost” layer sharp, and blur the others exponentially, the further “down” they are. The user can select a layer to be put on top, and thus decide which information is of more relevance. Another mode displays all layers sharply, but they can be reordered and layers can also be removed from the display. A semi-transparent display of all layers (which can be reordered) is also available.

Subjects used the program with handdrawn artificial maps with nine layers (representing rivers, roads, industrial areas, etc.), and were asked to find good spots for a factory, a technology company, and a holiday resort (the criteria were provided). The replies were noted and evaluated manually.

The response times were not significantly different; in semi-transparent mode, there were fewer clicks, because participants could see all the layers at once. This is also reflected in the number of interactions, and there is a significant correlation between the number of clicks and the response time. Observations of the tests revealed that participants needed different numbers of clicks to achieve the wanted arrangement of layers, and therefore the results are difficult to interpret.

8 Conclusions

This study has shown that SDOF is an effective and efficient method for guiding the user’s attention.

SDOF can draw the user’s attention to objects quickly. The smallest blur level was too small for these viewing conditions, it seriously impeded the subjects’ performance.

We found that SDOF can make it easier to get a first impression of data – and we also found that its use must be dosed very precisely to be useful to the user.

References

- Bartram, L. (1997), Perceptual and interpretive properties of motion in information visualization, *In: Proceedings of the Workshop on New Paradigms in Information Visualization and Manipulation, November 10-14, 1997, Las Vegas, USA*. ACM, 3-7.
- Healey, C.G. & Enns, J.T. (1999), Large datasets at a glance: Combining textures and colors in scientific visualization, *IEEE Transactions on Visualization and Computer Graphics*, 5(2): 145–167.
- Kosara, R., Miksch, S. & Hauser, H. (2001), Semantic depth of field, *In: IEEE Symposium on Information Visualization 2001 (InfoVis 2001), San Diego, CA, USA, October 22–23 2001*. IEEE Computer Society Press, 97-104.
- Kosara, R., Miksch, S. & Hauser, H. (2002), Focus and context taken literally, *IEEE Computer Graphics and its Applications, Special Issue: Information Visualization*, 22(1): 22-29.
- Lee, H.-C. (1990), Review of image-blur models in a photographic system using principles of optics, *Optical Engineering*, 29(5): 405–421.
- Treisman, A. (1985), Preattentive processing in vision, *Computer Vision, Graphics, and Image Processing*, 31: 156–177.