Joint EUROGRAPHICS - IEEE TCVG Symposium on Visualization (2002)
D. Ebert, P. Brunet, I. Navazo (Editors)

Useful Properties of Semantic Depth of Field for Better F+C Visualization

Robert Kosara Silvia Miksch
Vienna University of Technology, Vienna, Austria
http://www.asgaard.tuwien.ac.at/
{rkosara,silvia}@asgaard.tuwien.ac.at

Helwig Hauser
VRVis Research Center, Vienna, Austria
http://www.VRVis.at/vis/
Hauser@VRVis.at

Johann Schrammel Verena Giller Manfred Tscheligi
Center for Usability Research and Engineering (CURE), Vienna, Austria
http://www.cure.at/
{schrammel, giller, tscheligi}@cure.at

Abstract

This paper presents the results of a thorough user study that was performed to assess some features and the general usefulness of Semantic Depth of Field (SDOF). Based on these results, concrete hints are given on how SDOF can be used for visualization. SDOF was found to be a very effective means for guiding the viewer's attention and for giving him or her a quick overview of a data set. It can also very quickly be perceived, and therefore provides an efficient visual channel.

Semantic Depth of Field is a focus+context (F+C) technique that uses blur to point the user to the most relevant objects. It was inspired by the depth of field (DOF) effect in photography, which serves a very similar purpose.

Categories and Subject Descriptors (according to ACM CCS): I.3.36 [Computer Graphics]: Methodology and Techniques; H.5.2 [Information Interfaces and Presentation]: User Interfaces

1. Introduction

Like few other areas in computer science, visualization involves the user as the most important part. No matter how good a visualization technique is in terms of its computational cost, its clever design, or the pretty pictures it produces – if it does not convey information to the user efficiently and effortlessly, it is useless. Visualization therefore lacks elegant, formal proofs for its methods, and instead requires the "dirty work" of user studies and psychological tests to asses which methods and techniques are useful, and which are not. Such studies have been neglected in the past, but the awareness of the need of proper evaluation of methods is slowly growing[3, 5, 10].

This paper reports the results of a study that was performed to evaluate a new method called Semantic Depth of Field. We also present conclusions we drew about how and where this method can be used.

1.1. Semantic Depth of Field (SDOF)

Semantic Depth of Field (SDOF)[6, 7] 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 (e.g., certain chess figures in Figure 1). It is based on the depth of field (DOF) effect known from photography and cinematography[8], which depicts objects sharply or blurred depending on their distance from the lens. This is used to guide the viewer's attention, and is quite effective and intuitive. SDOF extends this effect to decide for every object whether to display it sharply or blurred, not based on geometry, but on the object's current relevance.

We measure blur 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.

Because blur is generally known to be slow in computer graphics, we developed a fast implementation that uses texture mapping on commodity graphics hardware[6] to make interactive applications possible on state-of-the-art PCs.

1.2. Study Goals

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 in two ways: a) by testing the ability to preattentively (i.e., very quickly and without serial search, see Section 3) detect and locate objects, as well as estimate the percentage of sharp objects; and b) by comparing search times for different cues, i.e., sharpness versus color and orientation (Section 4), and also checking for the interplay between SDOF and those other cues. We also tested the thresholds necessary to tell different blur levels apart, and also the blur levels necessary for an object to appear sharp or blurred (Section 5). Applications [7] were also tested, but are not presented here because of space constraints.

The following section presents some high-level results we obtained from the study; the sections after it go into the details – they first present the hypothesis to be tested, then the test method, and finally the results. Technical details of the study (sample, etc.) are given in the appendix.

2. Results – How to Use SDOF

The following points are the key findings of our study:

- SDOF can be used to quickly and effectively guide the user's attention.
- SDOF makes it possible to discriminate between a small number (about two to four) of object groups.
- Interaction is very important, because people do not like looking directly at blurred objects (if they do so, the application is badly designed).
- SDOF enables the user to get a quick overview of data by letting him or her ask questions quickly and efficiently.
- Blur levels have to be chosen carefully. For normal viewing conditions, we found a blur of 7 pixels too small, and a value of 11 pixels sufficient.
- Things that don't need to be blurred shouldn't be.

3. Preattentivity

Preattentive processes take place within about 200 ms after a stimulus is presented[1, 4, 9], and are performed in parallel, without the need for serial search. Such processes involve a limited set of features (e.g., orientation, closure, color, proximity, etc.) for which certain tasks (e.g., detection, location, count estimation, recognition of groups, etc.) can be performed with ease. Using preattentive features for visualization makes the information easier to see in order to get an overview. Especially methods for pointing out information have to make the relevant objects immediately stand out.

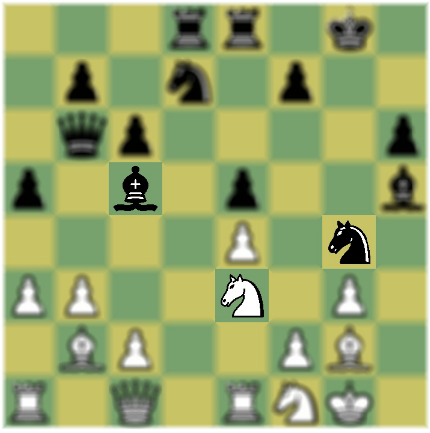

Figure 1: *Chess board application, with the chessmen threatening the knight on e3 in focus (from Kosara et al.[6])*

Experience suggests that sharp objects can be preattentively recognized among blurred ones: depth of field is a very effective means in photography and also cinematography, where the eye can be guided from one object to the other with focus changes. And blur is also present in the human eye, which also only has a limited depth of field (like a camera lens), but we hardly notice that – we simply ignore blurred areas.

3.1. Test Procedure

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

The images for target detection and location showed ellipses whose main axes were horizontal, and which were scattered over the image (Figure 2b). The reason for choosing ellipses was that we needed objects that would not change their shape drastically when blurred to rule out shape perception effects. Ellipses seemed perfect for this, because they don't change, and they can also be rotated (which was needed in the interplay trial, Section 4). 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, participants were shown five images (randomly picked from 30 generated ones), resulting in 210 images per participant.

The test procedure consisted of four steps (Figure 2a): First, an empty screen was shown for 300 ms, followed by the image, which was shown for 200 ms. After that, the answer screen was presented, which gave the participant the choice between clicking on one of the four quadrants or buttons for “no target” and “target not locatable”. After the answer was given, a screen with the German word for “Next” was shown, which required a key-press to continue with the next iteration. This was done to provide the participants with a means to control the speed of the test. Additionally, a screen encouraging the subject to take a short break was shown after every 30 images. 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 targets (in steps of 10%), and the rest distractors (Figure 2c).

3.2. Results

Finding sharp targets among blurred distractors is indeed performed preattentively. Figure 3a 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) 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 (Figure 3b), 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. This is most likely due to the participants not being able to distinguish between the slightly blurred distractors and the sharp target, and thus not finding it – so stronger blur than 7 pixels is needed (this is also quite apparent from Figure 3c).

Estimation of the percentage of sharp objects can also be done preattentively. The accuracy for all blur levels is significantly better than chance. When analyzed by number of targets, the accuracy is lowest close to the borders of the intervals (“few”, “many”, “intermediate”), and slightly higher on the low and high end than in the middle – which is not surprising, because for these numbers, the participants can make the decision more easily. The dependence on blur levels is weaker than for target location, and does not differ significantly between the lowest one and the stronger two. For the smallest blur level, more objects were perceived as sharp (Figure 3d), which led to more errors.

The results clearly show that SDOF is an effective method that can draw the user’s attention to objects quickly. Getting a first idea of data (e.g., in a scatter plot) seems also possible. The smallest blur level (7 pixels) clearly was too small for these viewing conditions, because it seriously impeded the subjects’ performance. Proper parameterization of the method for the user’s viewing conditions is therefore necessary.

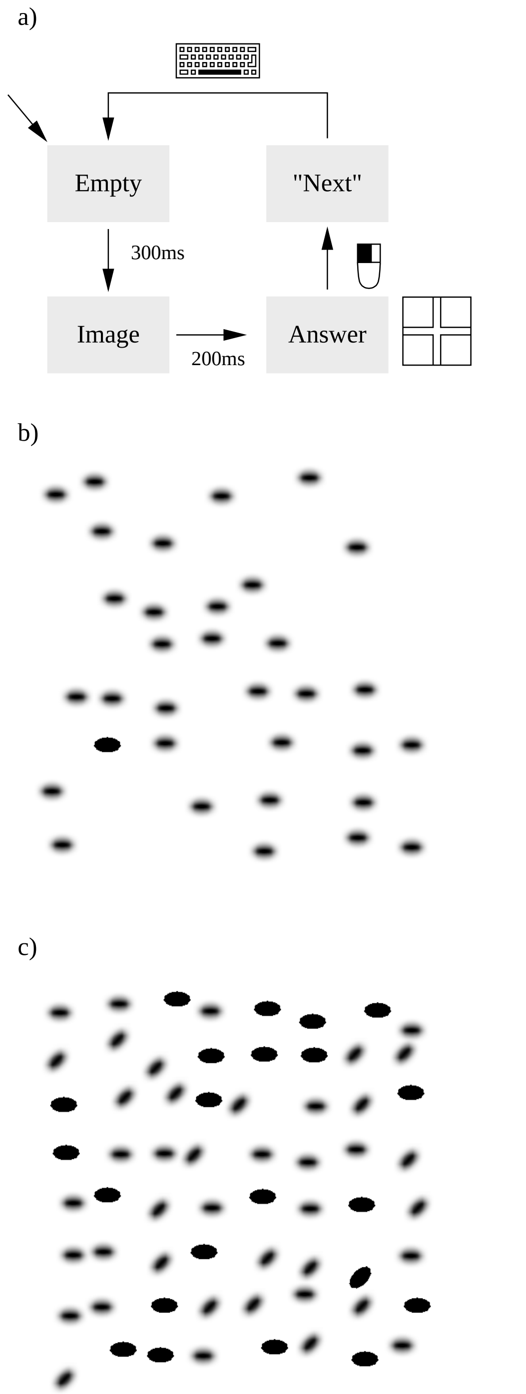

Figure 2: *Test sequence and sample images.* ***a)*** *The sequence of screens for testing preattentive location of objects (Section 3) and interplay (Section 4);* ***b)*** *Sample image for target detection and location with 32 distractors of the highest blur level, and a target;* ***c)*** *Example image for interplay: Find the rotated, sharp object.*

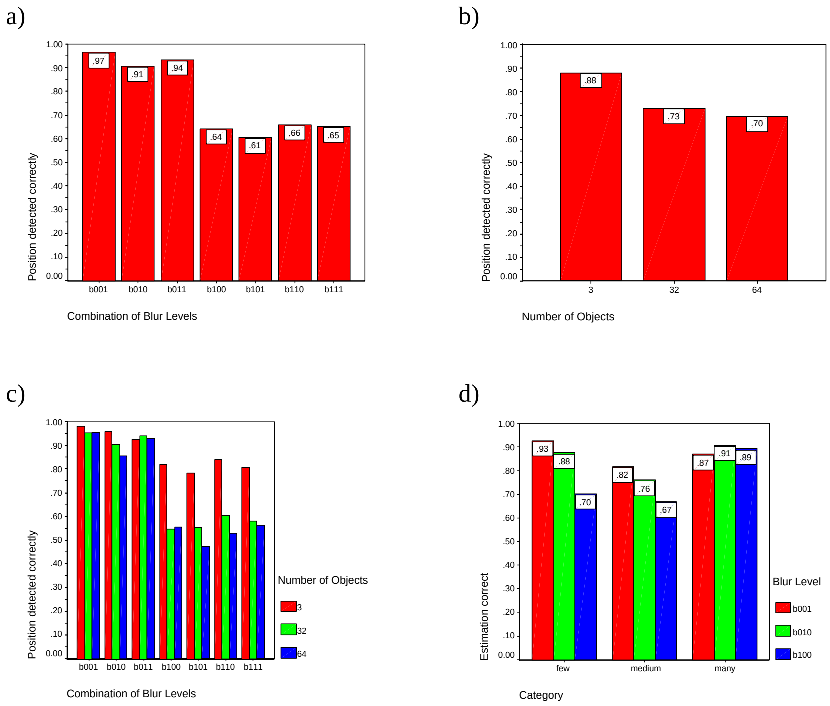

Figure 3: *Results for preattentiveness:* ***a)*** *Ratio of correctly located targets depending on blur levels used (encoding of blur levels see below);* ***b)*** *Ratio of correct answers by number of objects;* ***c)*** *Correct answers by blur level and number of objects;* ***d)*** *Ratio of correct estimations by blur level and number of targets.* ***Encoding of blur levels:*** *for each of the three blur levels, a 1 indicates that it is present, and a 0 that it is not. So for "b011", the lowest blur level was not present, the higher ones were.*

4. Interplay

SDOF will very likely not be used without any other visual cues, which is why we were interested in its interaction with other features; for this test we selected color and orientation.

4.1. Test Procedure

For this part, images similar to the ones used for the preattentiveness test were used, with the additional properties color (red, black) and orientation (main axis horizontal or at 45°).

The user interaction was similar to the first blocks, only 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. 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 sharp object), while the distractors could have any other combination of the two.

4.2. Results

In terms of search time, SDOF is not significantly worse than color – this is perhaps the most interesting and surprising finding of this study. There is no significant difference between a simple search for colored or for sharp objects (Figure 4a). The conjunctive searches for color and blur, orientation and blur, and color and orientation differ significantly from each other, with color and orientation being the slowest – each of these two features combined with SDOF is faster. Also, the conjunctive search for color and blur is not significantly slower than the simple and disjunctive searches, which is quite contrary to what we expected, because conjunctive searches usually are slower[9].

Search times were longer when no target was present (Figure 4b), which is not surprising, because it takes longer for subjects to make sure they have not overlooked a target[2]. The total number of errors in this block was only 10 (i.e., less than 0.7%) for the whole test (90 images per participant, 1440 in total). This shows that subjects took the tests seriously, and did not sacrifice accuracy for speed.

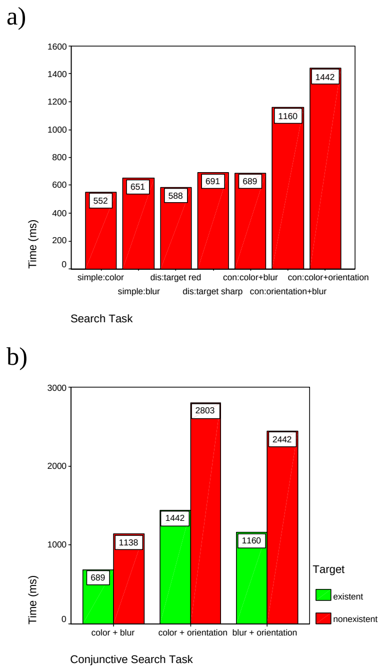

Figure 4: *Results for interplay of features.* ***a)*** *time needed for search when target present ("simple": only look for one feature, with no other feature present; "dis": disjunctive search for one feature with two distractor features; "con": conjunctive search for combination of features;* ***b)*** *search times for conjunctive search by search task and existence of target.*

5. Blur Perception

One aspect of SDOF we were planning was to use it as a fully-fledged, separate visualization dimension that could be used in addition to the existing such as space, color, etc. In order to do this, we needed to assess the minimal difference in blur that can be perceived, and the rate at which "steps" in blur are perceived. Our original hypothesis was that there would be an exponential relationship between the blur level and the perceived blur (similar to the way luminance is perceived, for example).

5.1. Test Procedure

This test consisted of several parts. In the 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 decided it was equal, the blur of one of the objects was increased (starting with sharp objects), and the objects were shown again. Another part started with strong blur and decreased the blur level. In a third part, participants had to decide, which of the two was sharper, or if they were equal.

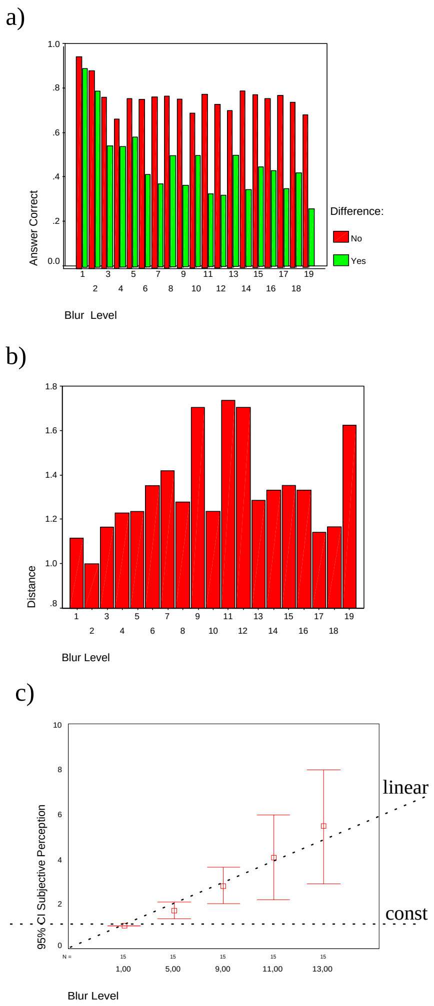

Figure 5: *Results for blur perception.* ***a)*** *Correct answers for identical ("no") and different ("yes") objects, by blur level;* ***b)*** *Distance needed to detect difference, by blur level;* ***c)*** *Numerical answer to perception of absolute blur value, by displayed blur value (error bars for 95% of values).*

We also 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 (until it was perceived as sharp).

In the final part, participants had to tell the perceived relation in blur in terms of a ratio of two numbers. They were free to use any numbers they wanted (i.e., not restricted to "1:x"), which were later normalized. These numbers were given orally, and recorded by the test supervisor.

5.2. Results

SDOF cannot be used as a full visualization dimension. Participants were able to tell the difference between objects of different blur levels (Figure 5a) with a good accuracy (which even stayed quite constant even for strong blur). But they were not able to correctly identify objects of the same blurriness, and did not better than chance for blur over 7 pixels.

The differences in blur needed to tell the blur levels apart (Figure 5b) do not show a clear trend. The distances are quite small (less than 1.8), and overall appear quite constant, which suggests a good differentiation between blur levels – in accordance with the above results. 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 already judged as blurred when starting out with a sharp object.

When judging the ratio of blur of two objects, subjects reported very small numbers compared to the real ratio (Figure 5c). Their answers also differed very much, so that no clear trend could be made out. This is quite contrary to the above results about being able to differentiate between blur levels. So while subjects were able to see a difference, they were not able to quantify it – another peculiar similarity to color perception.

The quantitative results of this part of the study form a consistent image with the participants' comments, that they disliked having to look at blurred objects and to compare them. It therefore appears to be necessary to make sure that no important parts of the display are blurred, and that the user can switch to a different view, or back to a completely sharp image at any time.

6. Conclusions and Future Plans

This study has shown that SDOF is, indeed, an effective and efficient method for guiding the user's attention. We were surprised to find the similarities with color (even though they were not significant), which we had not expected. We now also have some data for parameterization of SDOF for the use in standard desktop environments, and can perhaps extend this to other viewing conditions as well.

This study has revealed a lot of interesting information about SDOF, but it was only a first step. We want to continue investigating SDOF properties and parameters in new studies which we want to design based on this one.

Acknowledgments

We would like to thank Peter Fröhlich and Birgit Rabl for their work to make this study possible.

This work is part of the Asgaard Project, which is supported by *Fonds zur Förderung der wissenschaftlichen Forschung* (Austrian Science Fund), grant P12797-INF. Parts of this work have been carried out in the scope of the basic research on visualization at the VRVis Research Center (`http://www.VRVis.at/vis/`) in Vienna, Austria, which is funded by an Austrian governmental research program called Kplus. Furthermore, this work is partly supported by the City of Vienna.

Appendix: The Gory Details

The test setup was designed in workshop sessions between the computer scientists at the University of Technology/VRVis and the usability experts at the Center for Usability Research and Engineering (CURE); the test software was developed by Robert Kosara. All tests were performed in August 2001 in the CURE usability lab in Vienna, Austria.

For significance testing, we used chi-square tests and ANOVAs with Scheffé tests for post-hoc analyzes. All results that are described as significant in this paper were tested for with a probability for error of $p < 0.001$. The base level for the whole study was $p < 0.05$.

To rule out large differences in perception between test participants, and to allow for a rather small sample size due to financial and time constraints, we selected a rather narrow group of participants who all fulfilled the following requirements: male, aged 18–25, very good vision (no contact lenses or glasses), student at university, basic computer knowledge.

The sample size was 16 individuals, which we recruited from different universities in Vienna. Each participant was paid a small amount of money for taking part. Each test session took about two hours.

References

1. Lyn Bartram. Perceptual and interpretative properties of motion for information visualization. In *ACM Workshop on New Paradigms in Information Visulization and Manipulation*, 1997.
2. E. Bruce Goldstein. *Sensation and Perception*. Brooks/Cole Publishing Company, 5th edition, 1998.
3. Christopher G. Healey, Kellogg S. Booth, and James T. Enns. Visualizing real-time multivariate data using preattentive processing. *ACM Transactions on Modeling and Computer Simulation*, 5(3):190–221, July 1995.
4. Christopher G. Healey and James T. Enns. Large datasets at a glance: Combining textures and colors in scientific visualization. *IEEE Transactions on Visualization and Computer Graphics*, 5(2):145–167, April 1999.
5. Ivan Herman, Guy Melançon, and M. Scott Marshall. Graph visualization and navigation in information visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 6(1):24–43, January-March 2000.
6. Robert Kosara, Silvia Miksch, and Helwig Hauser. Semantic depth of field. In *IEEE Symposium on Information Visualization 2001 (InfoVis 2001)*, pages 97–104. IEEE Computer Society Press, 2001.
7. Robert Kosara, Silvia Miksch, and Helwig Hauser. Focus and context taken literally. *IEEE Computer Graphics & Applications, Special Issue on Information Visualization*, 22(1):22–29, January/February 2002.
8. Hsien-Che Lee. Review of image-blur models in a photographic system using principles of optics. *Optical Engineering*, 29(5):405–421, May 1990.
9. Anne Treisman. Preattentive processing in vision. *Computer Vision, Graphics, and Image Processing*, 31:156–177, 1985.
10. Colin Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers, 2000.