Search Task Performance Using Subtle Gaze Direction with the presence of Distractions

Ann McNamara[1]
Texas A & M University
and
Reynold Bailey[2]
Rochester Institute of Technology
and
Cindy Grimm[3]
Washington University in Saint Louis

A new experiment is presented which demonstrates the usefulness of an image space modulation technique called Subtle Gaze Direction (SGD) for guiding the user in a simple searching task. SGD uses image space modulations in the luminance channel to guide a viewer's gaze about a scene without interrupting their visual experience. The goal of SGD is to direct a viewer's gaze to certain regions of a scene without introducing noticeable changes in the image. Using a simple searching task we compared performance using no modulation, using subtle modulation and using obvious modulation. Results from the experiments show improved performance when using subtle gaze direction, without affecting the user's perception of the image. We then extend the experiment to evaluate performance with the presence of distractors. The distractors took the form of extra modulations which do not correspond to a target in the image. Experimentation shows, that, even in the presence of distractors, more accurate results are returned on a simple search task using SGD, as compared to results returned when no modulation at all is used. Results establish the potential of the method for a wide range of applications including gaming, perceptually based rendering, navigation in virtual environments and medical search tasks.

Categories and Subject Descriptors: I.3.3 [**Computer Graphics**]: Display algorithms

Additional Key Words and Phrases: eye-tracking, psychophysics, image manipulation, gaze direction, luminance

1. INTRODUCTION

Humans routinely perform visual search tasks such as searching for a familiar face in a crowd or scanning a document for some important information. Although such searches are a natural part of our visual processing, there are situations in which the task becomes

ann@viz.tamu.edu
rjb@cs.rit.edu
cmg@wustl.edu

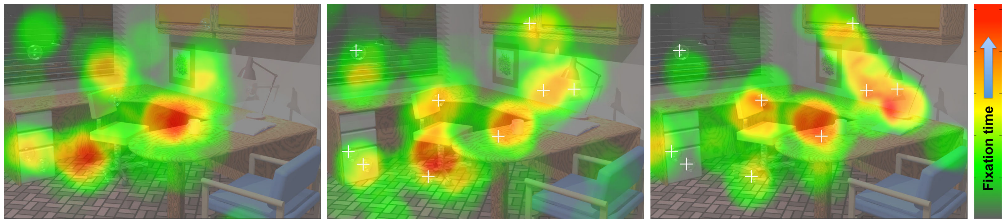

Fig. 1: *From left to right: distribution of fixation time under normal viewing conditions, using Subtle Gaze Direction modulations, and obvious modulations*

quite complex and demanding. There are numerous factors which impact the difficulty of a visual search. For example, size of the scene, size of the target, subtlety of the target, contrast, number of objects in the scene, etc. Some types of visual searches may even require specialized training and significant experience in order for the viewer to become proficient. In the medical profession, for example, deciphering x-rays while searching for abnormalities is a demanding search task [Schwaninger et al. 2007] [Schwaninger et al. 2004].

Fig. 2: *Example of an image used in this study. The search targets are the transparent spheres "bubbles" in the image.*

One way to improve performance in such tasks is to develop a technique to guide the viewer's gaze toward the regions of a scene that are important for successful completion of the task. To date several researchers have focused on *following* the viewer's gaze pattern to gain efficiencies in rendering and presentation [Luebke et al. 2002] [Duchowski 2002] [O'Sullivan et al. 2003]. However, research is beginning to emerge which looks at *directing* a viewer's gaze about a scene [Kim and Varshney 2006] [Mitchell 2004]

[Kosara et al. 2001] [DeCarlo and Santella 2002]. This paper focuses on the simple task of counting targets in an image (see Figure 2). Accurately counting targets efficiently is a necessary task for many applications. For example, military personnel searching for camouflaged targets in a scene [Neider and Zelinsky 2006], [Boot et al. 2009]. The gaze directing technique used in this paper, which we call Subtle Gaze Direction [Bailey et al. 2009], combines real-time analysis of eye movement data with subtle image space modulation to *direct* the viewer's gaze towards selected targets of known location.

One might argue that if the information regarding important regions is available why not simply present that to the user. We agree that in many cases this would be the correct solution, however, cases exist where an algorithm can be beneficial in "suggesting" places to look without disturbing the visual experience of the viewer.

Subtle Gaze Direction [Bailey et al. 2009] depends on the well established fact that the peripheral vision processes stimuli more quickly than the foveal vision [Ogden and Miller 1966]. When viewing a scene for the first time, the low acuity peripheral vision of the Human Visual System (HVS) locates areas of interest. The slower, high-acuity foveal vision is then involuntarily directed to fixate on these regions. By modulating regions of the scene that appear only to the peripheral vision we can force the peripheral vision to locate areas of interest, which are subsequently focused on. This causes the eyes to move involuntarily (saccade) to focus the foveal vision on the modulated region in an attempt to identify the stimuli detected. Luminance modulations were chosen because the HVS is very sensitive to luminance changes [Spillmann 1990].

The modulations used in the Subtle Gaze Direction technique are made by alternately blending the pixels in a small area with some amount of black, then some amount of white. Bailey et. al. [Bailey et al. 2009] describe several measures taken to ensure that the modulations are subtle. For example the intensity of the modulations are adjusted so that they are just intense enough to be detected by the peripheral vision. Additionally, a Gaussian fall off function is used to ensure that the edges of the modulated regions are not prominent. The size of the modulated region is kept very small, approximately 1 cm diameter circular region on the screen in the initial set up. The modulation blends from black to white at a rate of 10 Hz.

Finally, the viewer's foveal vision is never allowed to fixate on the modulated region. This is achieved by monitoring the direction component of the saccade velocity vector, to determine if the foveal vision is about to enter the modulated region. If this is the case, the modulation is immediately terminated.

Experiments have demonstrated that Subtle Gaze Direction is an effective approach to gaze direction. As shown in Figure 3 gaze patterns can be altered in the presence of Subtle Gaze Direction. This technique has several important benefits [Bailey et al. 2009]. The impact on the viewing experience is minimized as images are not permanently altered. The technique is fast, viewers typically attend to the modulated regions within 0.5 seconds of modulation onset. And, despite the measures taken to ensure subtlety, the technique is still quite accurate, typically resulting in fixations within a single perceptual span of the modulated region.

Figures 1 and 3 illustrate the results of applying Subtle Gaze Direction to an image, and

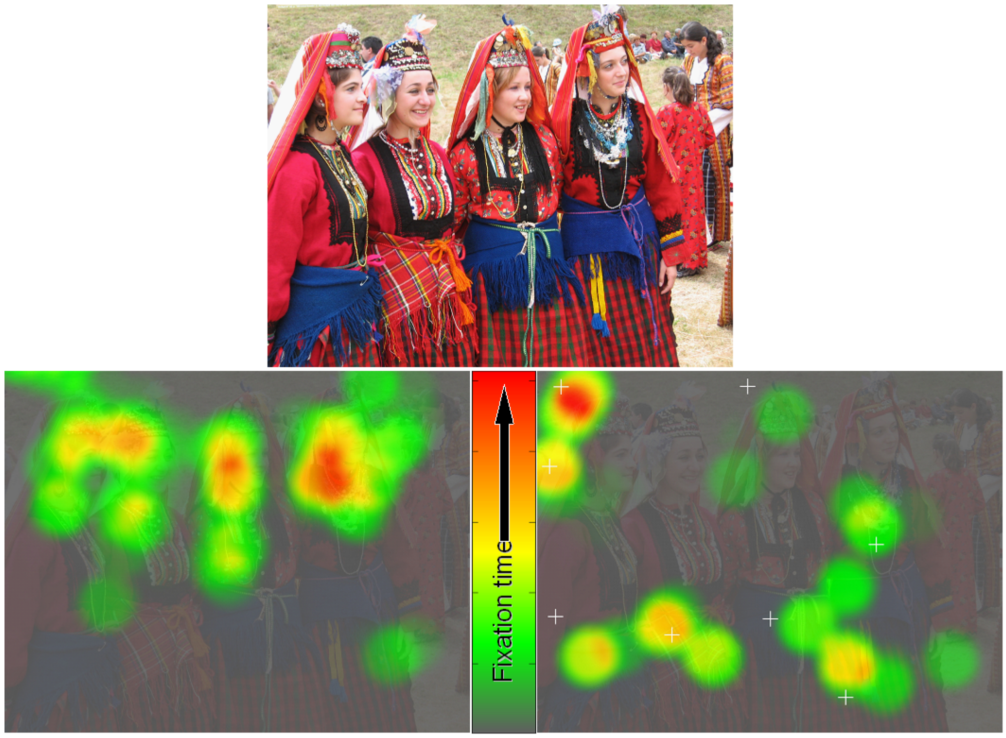

Fig. 3: *Gaze distributions for an image under static and modulated conditions. Input image (top). Gaze distribution for static image (bottom left). Gaze distribution for modulated image (bottom right). White crosses indicate locations preselected by researchers for modulation.*

the affect on the viewer's gaze pattern. The Figures shows a heat map of the average scan patterns over the image for 6 observers. In Figure 1 the image on the left shows the scan pattern resulting from normal viewing. The middle and right images were altered by adding modulations at the target regions indicated by white crosses. In the middle image we applied the subtle modulations, in the right image larger modulations were used. In Figure 3 the top image is for reference, the left image shows the normal gaze pattern over the image while the image on the right shows the gaze pattern with Subtle Gaze Direction applied to the image. Again the white crosses indicate where the modulations were invoked. As can be seen from both these images gaze can be influenced in a prescribed manner using Subtle Gaze Direction. Further details can be found in [Bailey et al. 2009].

This paper presents a psychophysical experiment that explores the impact of Subtle Gaze Direction on performance during a visual search task. The results show that this method works well without introducing visible disruption into the image.

2. INITIAL EXPERIMENT

Twenty-four images served as stimuli for the experiment (see Figure 4). Six environments were chosen and populated with four different target counts, ranging from 4 targets to 12

Fig. 4: *Scenes used in the experiments. All images were 693 by 1024 pixels except the bathroom scene and the interior scene which were 691 x 1024 and 797 x 1024 respectively.*

targets for a total of 24 images. The targets were small transparent spheres roughly uniformly distributed within the scene. Some spheres were deliberately placed so as to be difficult to resolve. The reason for this was to allow us to investigate if those hard to see targets were more easily resolved using modulation. All of the models used to create the images were taken from the RADIANCE web site [Larson and Shakespeare 2004]. Presentation order was randomized to eliminate any learning effects. Images were presented for 14 seconds. A black screen with a white cross at the center was presented between each image to allow the participant to refocus on the center of the screen. Participants were explicitly instructed to focus on the center of the cross. This means the initial viewing position for each image is the same i.e. viewing begins in the center of the images. An example of an image used in this study is shown in Figure 2. Image sizes varied as shown in Figure 4. In cases where image size was smaller than the viewing screen the image was centered on a black background.

Participants were seated in front of the computer screen in a well lit room with their chin comfortably resting on a chin-rest to reduce head movement. Using an infrared camera-based eye-tracking system[4], data pertaining to fixation position and saccades were recorded for the dominant eye of each participant. A fixation is defined as any pause in gaze $\geq 150ms$. Participants were instructed to remain as still as possible while the eye-tracker was calibrated and the experiment was conducted. The chin-rest was positioned $75cm$ from the screen. At this distance, the actual perceptual span (area of high acuity) of the observer occupies a circular region of diameter $5cm$ on the screen [Rayner 1975]. The subtle modulations were presented in a smaller ($1cm$ diameter) circular region.

Eye-tracking was employed to record the viewer's fixation and saccades while counting targets in the various images. Eye-tracking information also served as input to trigger the modulations on targets that were not attended to, in an effort to highlight them so the user could identify them and include them in their count. Image complexity varied as did the number of targets. The behavior of the targets was also varied as follows:

—**GROUP 1, NO MODULATION:** no behavioral actions applied to the targets, so images were viewed normally with no modulations.

—**GROUP 2, SUBTLE MODULATION:** subtle image modulations were used to highlight the target regions in an effort to aid in counting. Modulations happen on target regions, selected at random that are not being attended to by the viewer. Modulation was never applied to targets while they were being directly viewed. Any modulation was applied in the periphery only, and modulations were terminated as the user moved their gaze toward the modulated region. Thus the viewer was never allowed to directly view the modulated region. A modulation radius of 0.04 degrees of visual angle was defined to ensure subtlety.

—**GROUP 3, OBVIOUS MODULATION:** subtle behavior was exaggerated so that the modulations were clearly visible by increasing the size of the modulation. Modulation was similar to the modulation applied in Group 2, however, in this condition the modulations were deliberately set to be more obvious. A modulation radius of 0.125 degrees of visual angle was defined to ensure visibility.

The targets subtended visual angles ranging from 0.05 to 0.08 degrees, depending on their location in the scene. Therefore subtle modulations subtended ≤ 0.5 the size of the targets, while the obvious modulations subtended a visual angle of between 1.5 to 2 times the size of the target.

Eighteen participants were assigned randomly to one of the three groups. Participants volunteered from a group of undergraduates. All had normal or correct-to-normal vision and were naive to the purpose of the experiment. Participants viewed the images on a 22" LCD Screen at a resolution of 1200 X 1600 from a distance of 75cm. Head position was held constant using a chin rest for support, as shown in Figure 5. The eye movements of each participant were recorded along with a count of the targets found and the time to

[4] ViewPoint EyeTracker® by Arrington Research, Inc.

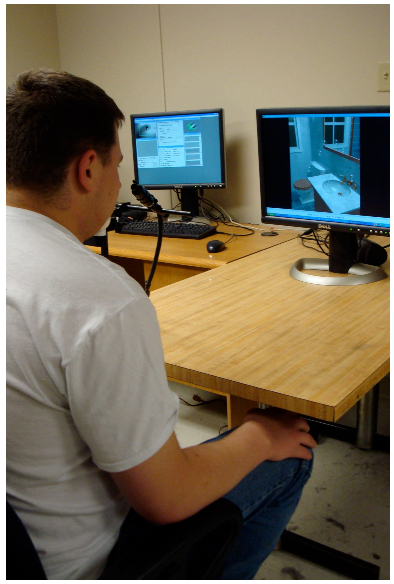

Fig. 5: *Experimental Set Up*

respond for each image. An informal exit interview questioned the participants about the quality of the images to determine if the modulations in conditions 2 and 3 were disturbing to the viewer. Participants in condition 2 reported nothing unusual, whereas in condition 3 participants reported seeing the modulations.

The task involved viewing each scene and counting the number of targets present. Participants verbally reported the number of targets counted on completing the task.

3. RESULTS AND DISCUSSION

IMAGE	NONE	SUBTLE	OBVIOUS
Image A: Soda Hall	25.0%	54.2%	66.0%
Image B: Conference	79.2%	79.2%	66.7%
Image C: Interior	12.5%	29.2%	58.3%
Image D: Office	20.8%	62.5%	54.2%
Image E: Bathroom	37.5%	50.0%	29.2%
Image F: Counter	70.8%	62.5%	66.7%
AVERAGE	40.97%	56.25%	56.94%

Table I: *This table shows the percentage of accurate detection of all the targets in an image. 100% means all of the targets were found. Each column is the average over 4 cases. Standard Deviations were of the order of 2%.*

The number of targets reported and the time to respond was recorded for each image. The average response times were consistent across all three conditions, 6.272, 6.495 and 6.570 seconds (with standard deviations of 0.55, 0.94 and 0.96) for the no modulation, subtle modulation and obvious modulation conditions respectively.

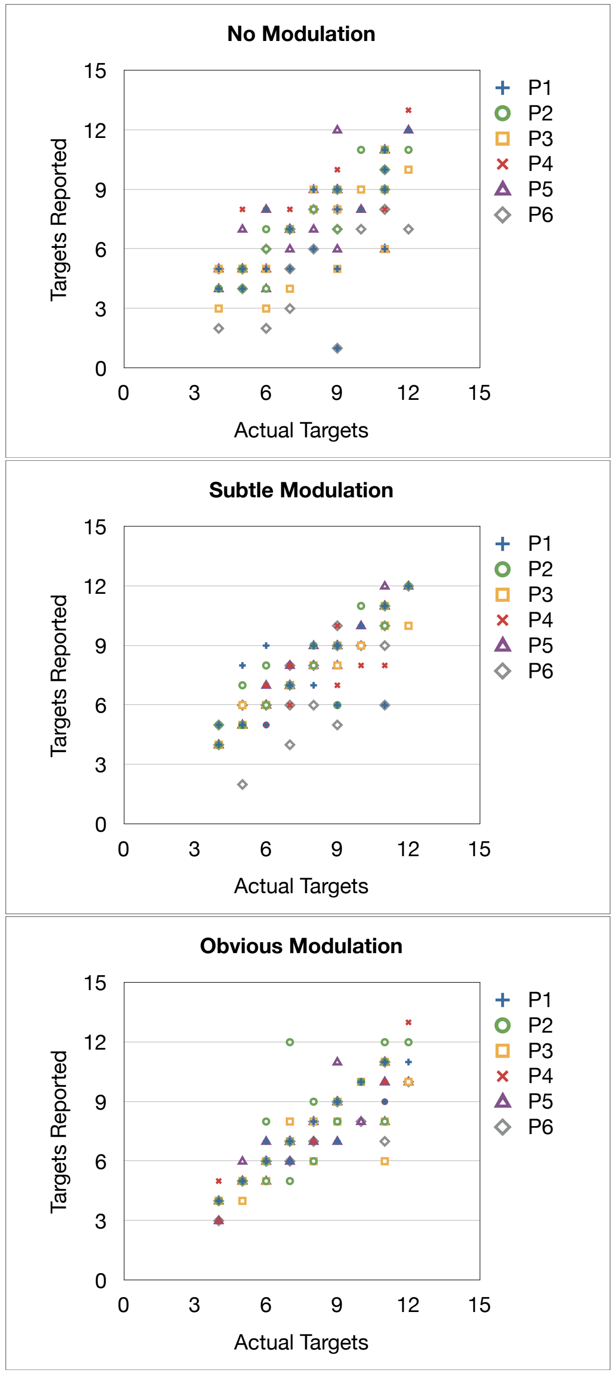

Fig. 6: *The graphs show the correlations between the actual number of targets present in each image and the number of targets reported by the observers for groups 1, 2 and 3 respectively.*

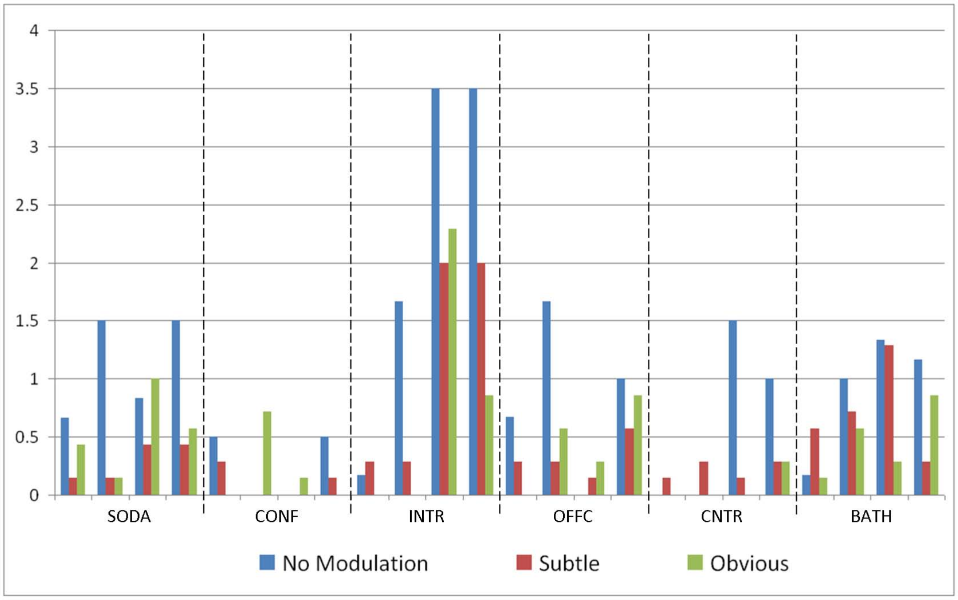

Fig. 7: *Experimental Results: This chart shows the sum of differences between the actual number of targets and the number of targets reported for each condition.*

We compared the reported number of targets to the actual number of targets and used this to define a correlation. These correlations are graphed in Figure 6. The correlation values represent how close the reported number of targets were to the actual number of targets. Higher correlation corresponds to more accurate task performance. Correlation is higher for the modulated conditions than in the static image with values of 0.80, 0.89 and 0.90 for groups 1, 2 and 3 respectively. This indicates that participants did slightly better on task with the aid of modulation as opposed to normal viewing.

Another interpretation of the results is to simply compare the number of targets missed during counting in each case. This data is shown in Figure 7. Each bar represents the absolute difference between the actual number of targets and the number reported. Smaller bars indicate more accurate counts, no bar indicates 100% precision. The blue bars show the differences in the normal no modulation viewing condition, while the orange and green bars show the modulated images, subtle and obvious respectively. The four bars represent the number of targets for each image (the number of targets ranged from four to 12). Each cluster is one image, as labeled on the x-axis. As the data shows, in most images, the modulation aids in the accuracy of the results. In some cases the number of targets reported does not increase as a function of the number of targets, one reason for this may be the participant's failure to see the targets, and subsequent failure to include them in their count.

The percentage of correct counts reported also reveals that a higher percentage of counts returned in the modulated imagery were accurate compared to the imagery with no modulation. What is interesting in this analysis is that the percentage correct in both modulated cases is higher than in the static imagery. However, it is important to note that while no obvious distractions were noted by participants viewing the subtle modulation case, all participants in the obvious modulation condition reported seeing the modulations. They noted that while the modulations did distract their gaze, it also helped them to identify targets they may not have otherwise counted. This indicates that simple region highlighting, even if noticeable, can contribute to improving task accuracy. The data suggests that subtle gaze direction, where the highlighting is sufficiently faint so as to go unnoticed, successfully guides the viewer's gaze to the target regions, thereby improving task performance. Data is tabulated in Table I.

The highest discrepancies occurred in the image of the interior scene (Image C). Here the composition of the scene may have influenced the visibility of the bubbles, with only one person getting 100% accuracy. The placement of the targets was also made deliberately difficult. One reason for the poor results in the bathroom scene may be due to the fact that participants reported being confused regarding the inclusion or exclusion of targets reflected in the mirrors, despite explicitly being told not to count reflections.

Further statistical analysis of the data was conducted. The percentage of targets reported, that is, the number of targets reported as a percentage of the actual targets present, is taken as the dependent variable. Independent variables are the modulation (none, subtle, obvious) and the image (6 scenes). To show a significant effect of condition is present between the 3 groups a repeated measures ANOVA was performed with 3 levels of modulation factor resulting in $F(2, 15) = 2.004; p \leq 0.001$.

This gives evidence that a significant effect of condition is present between the tasks i.e. performance differed in each group. Follow up comparisons in the form of a Bonferroni pairwise comparison shows that the differences between groups are significant. The full results are given in Table II. The results indicated that group 1 (no modulation) differs from group 2 (subtle) and group 3 (obvious) with $p \leq 0.001$ and $p \leq 0.05$ respectively. However the results showed no significant difference between performance in groups 2 and 3 (subtle and obvious). This supports our hypothesis that people are more accurate at counting targets with the aid of modulation, even if that modulation is obvious.

Group (I)	Group (J)	Mean Difference (I-J)	Std. Error	Sig.
1	2	-.069*	.016	.001
	3	-.051*	.015	.008
2	1	.069*	.016	.001
	3	0.18	.015	.717
3	1	.051*	.015	.008
	2	-.018	.015	.717

Table II: *Bonferroni pairwise comparison for each groups, showing significant differences between group 1 and group 2/group 3, but no significant differences between group 2 and group 3.*

Including the scene as an independent variable gives a 3 (modulation level) X 6 (scene) ANOVA with modulation as a between-subjects variable and scene as within-subjects variable. The result from the ANOVA shows no effect of scene $F(5, 75) = 1.738; p \geq 0.05$, suggesting that scene did not have an overall effect on performance in this case.

This means that scene did not have an effect on performance in this case.

In summary the results show slight improvement in task performance when modulation is employed to direct gaze to target regions. This seems to hold true whether or not the viewer notices the modulations. Some applications may elect to include obvious modulations whereas in other applications subtlety may play a key role.

4. EXTENDING SIMPLE SEARCH TO INCLUDE DISTRACTORS

4.1 Motivation

Many vision algorithms are known to return false positives. In medical applications, pattern classifiers and image processing techniques are routinely used to identify information anomalies, for example in the diagnosis of certain cancers [Tourassi et al. 2003] [Zhuowen et al. 2006]. Data from images used to diagnose such diseases can be hard to read and interpret by the human eye, and so vision techniques are employed to *guide* diagnosis, [El-Baz et al. 2006] [Mancas et al. 2004]. Generally these techniques come with the caveat that they will performs automatic suspect localization, feature extraction, and diagnosis of a particular pattern-class. But the final diagnosis must be ascertained by a trained professional. The overall aim of such algorithms is a higher rate of of true-positive detection and low false-positive detection. This means that some false-positive information may still be present in the results. However, the information returned by these algorithms still maintains a high level of usefulness for certain applications, especially when post processed by a human observer.

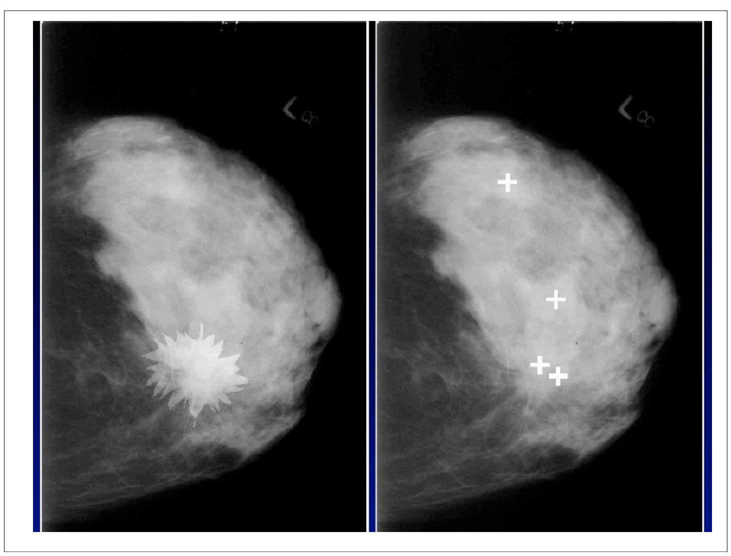

Fig. 8: *False Postives: This image shows the results of a mammogram. There is an irregular mass on the breast which has been outlined by a surgeon (left) while a computer vision algorithm correctly predicted the irregular mass it also generated 2 other false positive marks as indicated by the white crosses (right).*

Figure 8 shows two mammogram films. On the left image an irregular mass on the breast has been outlined by experienced personnel simply drawing on film. The image on the right shows the results generated using a computer vision technique. It correctly identifies the breast tumor, but also returns two false positives [Tourassi et al. 2003] [Bilska-Wolak et al. 2003]. The algorithm correctly predicted the irregularity, but also highlighted 2 potential irregularities. Hence such a technique is useful as a first iteration of the search task.

We were interested to see how Subtle Gaze Direction would perform if the modulations were driven by the results of such an algorithm. To test this hypothesis we re-ran experiment 2 (*the subtle modulation group*) from the initial set of experiments, with the addition of extra modulations which did not pertain to target location i.e. distractors. In this experiment participants were exposed to modulations not only on the targets but also in random locations away from targets i.e. the false positives.

For this follow on experiment the set up was essentially identical to the initial experiment, with some alterations to allow for inclusion of the distractor modulations. The initial 24 images were used as stimuli. This time modulations appeared not only on the targets but also in random locations away from the targets. 3 additional *distractor* modulations were added to each image. 6 new participants were included in the study. All had normal or corrected-to-normal vision. Each image was viewed for 15 seconds. Participants were asked to report the number of targets in the scene. Scenes were presented in a random order, as were the modulations within each image. Fixations were recorded while the counting task was performed.

4.2 Experimental Results

Taking a correlation between the actual number of targets and the number of targets reported gives an initial metric of search performance in the images with additional distractors (false positives). The average correlation over all images was 0.93 (and ranged over participant from 0.89 to 0.96). These correlations are notably higher than those reported in the initial experiment. Figure 9 shows the correlations for each participant, with the results from the initial experiment shown on the right.

As before we report the percentage of times that participants were accurate on the task, i.e. the number of times that the count returned was equal to the actual number of targets present, Table III. As can be seen overall performance was enhanced in the distractor condition.

IMAGE	NONE	SUBTLE	OBVIOUS	DISTRACTOR
Image A: Soda Hall	25.0%	54.2%	66.0%	54.1%
Image B: Conference	79.2%	79.2%	66.7%	62.5%
Image C: Interior	12.5%	29.2%	58.3%	75%
Image D: Office	20.8%	62.5%	54.2%	37.5%
Image E: Bathroom	37.5%	50.0%	29.2%	100%
Image F: Counter	70.8%	62.5%	66.7%	54%
AVERAGE	40.97%	56.25%	56.94%	65.8%

Table III: *This table shows the percentage of accurate detection of all the targets in an image along side additional data for the distractor image. 100% means all of the targets were found. Each column is the average over 4 cases. Again Standard Deviations were of the order of 2%.*

Following the same analysis as outlined in section 3 for the initial experiment, we take a one-way ANOVA over modulations to test our hypothesis that there is a difference between groups. This time there is another condition for modulation, making 4 levels (none, subtle, obvious, distractor) giving a one way ANOVA with 4 levels of modulation $F(3, 20) = 6.337; p \leq 0.001$.

Again a post-hoc Bonferroni pairwise analysis shows significant differences lie between the no modulation and modulated conditions, that is there is a difference between group 1 (no modulation) when compared against groups 2, 3 and 4 (subtle, obvious and distractors) but no difference among groups 2, 3 and 4. Again this supports our hypothesis that there is difference in performance when modulations are present. Full Bonferoni results are given in Table IV.

Group(I)	Group(J)	Mean Difference (I-J)	Std. Error	Sig.
1	2	-.069*	.016	.002
	3	-.051*	.015	.017
	4	-.057*	.018	.028
2	1	.069*	.016	.002
	3	0.018	.015	1.000
	4	.012	.010	1.000
3	1	.051*	.015	.017
	2	-.018	.015	1.000
	4	-.005	.015	1.000
4	1	.057*	.018	.028
	2	-.012	.010	1.000
	3	-.005	.015	1.000

Table IV: *Bonferroni pairwise comparison for each pair of groups, showing significant differences between group 1 and group 2/group 3/group 4, but no significant differences between group 2/group 3/group 4.*

Including the scene as an independent variable gives a 4 (modulation level) X 6 (scene) ANOVA with modulation as a between-subjects variable and scene as within-subjects variable. The result from the ANOVA gives $F(5; 100) = 5.861; p \leq 0.001$.

Table III indicates that distractors may actually increase performance when compared to subtle or obvious modulation, especially in the interior and bathroom scenes. Significance tests, again in the form of Bonferroni post-hoc pairwise comparisons on the interaction between modulation and scene indicate that the highest number of significant differences occur in the distractor condition. The image which differs most from all others across all four conditions is the interior scene (Image C). This is especially evident in Group 1 (no modulation) and Group 2 (subtle). However in Group 4, the image which differs most from the other images is the office scene (Image D). Significance for Image C and D are summarized in table V. This interaction between modulation and scene may lead to an interesting follow on experiment in which scene context is examined, for example there may be certain scene characteristics which favor the presence of distractors in modulation for certain tasks.

This can also be seen in Figure 10, which displays the average distribution of fixation time across the image. The white crosses indicate the pixel areas which were modulated on target, while the white crosses bounded by a rectangle show the distractor modulations.

Image	Image	Group 1 None	Group 2 Subtle	Group 3 Obvious	Group 4 Distractor
C	A	1.000	0.050*	1.000	1.000
	B	0.003*	0.002*	0.031*	0.067
	D	0.008*	0.001*	1.000	0.000*
	E	0.019*	0.013*	0.405	1.000
	F	0.000*	000*	0.024*	0.289
D	A	0.569	1.000	1.000	0.000*
	B	1.000	1.000	1.000	1.000
	C	0.008*	0.001*	1.000	0.000*
	E	1.000	1.000	1.000	0.002*
	F	0.234	1.000	1.000	453

Table V: *Significance values for Images C (interior) and D (office) across modulation. A significance value $\leq$ 0.05 means that significant differences.*

Notice how the gaze is drawn toward the distractor modulations, indicating that they successfully drew the viewers gaze away from their original focus. Representative images for the other stimuli are shown in Figures 11, 12 and 13.

One possible reason for this increase in accuracy is that the distractors attracted gaze and forced viewers to distribute their gaze more completely over the image. It is well established that people can miss even very salient features in an image if their attention is focussed elsewhere (Inattentional blindness [Mack and Rock 1998] [Cater et al. 2002] [Simons and Chabris 1995]). By re-focussing attention, through Subtle Gaze Direction, the viewer is prevented from focussing on a single region for long periods of time and avoid some inattentional blindness by moving their gaze over larger portions of the image.

5. CONCLUSIONS AND FUTURE WORK

We presented an experiment to compare task performance in digital images across three sets of stimuli. The results indicate that using either subtle or obvious image modulations on the target regions improves the precision of a simple counting task. The difference between using subtle and obvious modulations is the level of disruption to image viewing. With subtle image modulation none of the participants reported noticing the modulations, whereas with the obvious image modulation all participants reported seeing the modulations.

The success of employing Subtle Gaze Direction on targets to guide viewers in a simple search task led us to investigate the extent to which this success could carry over into images where modulations were presented not only on targets, but also on non-target regions. The reason for investigating the degree to which the technique could improve efficiency in a search task in even with the presence of distractors is important. As mentioned earlier many image processing and vision algorithms are known to return false positives. The presence of false positives does not render the algorithm useless. Rather, false positives can be used as an initial guide for the search task. Subtle Gaze direction may also find application in interactive environments such as games and animation. Inherently in dynamic scenes such as animation, or navigation through virtual worlds several objects may be moving. To successfully incorporate Subtle Gaze Direction into such environments we

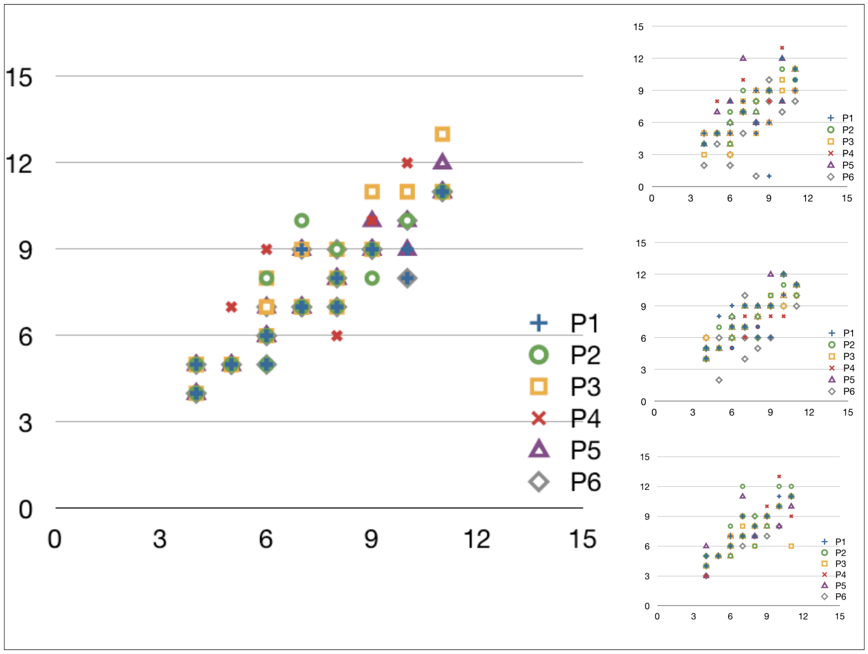

Fig. 9: *Experimental Results: This graph charts the correlation between the actual number of targets and the number of targets reported for each image in the distractor condition. Notice there is higher correlation than in the earlier experiment (as shown on the right; from the top, no modulation, subtle modulation and obvious modulation.*

need further studies which gauge the performance of the technique in the presence of such distractors in the presence of motion.

In these experiments we focussed only on modulating the luminance channel. It may be that the modulations would be more successful in certain conditions if we used other channels. We have experimented with the warm-cool channel, but found luminance to be slightly more efficient [Bailey et al. 2009]. It may be that the most effective modulation should be a function of the image itself, and may be dynamic depending on the behavior in the scene.

The results from this initial study are promising and several follow up experiments are imminent. There are several avenues open for future investigation. In the follow up experiment we showed that the presence of distractors does not hinder task performance, and in fact could be an accurate visual aid since it encourages people to look over the entire image. It is well known that we as humans suffer from *inattentional blindness* a phenomenon which essentially renders us blind to objects in our visual field which would otherwise be obvious due to the fact our attention is focussed elsewhere. We could use Subtle Gaze Direction to force attention away from the object currently focussed on in an effort to distribute gaze over an entire image and hence provide a more robust gaze

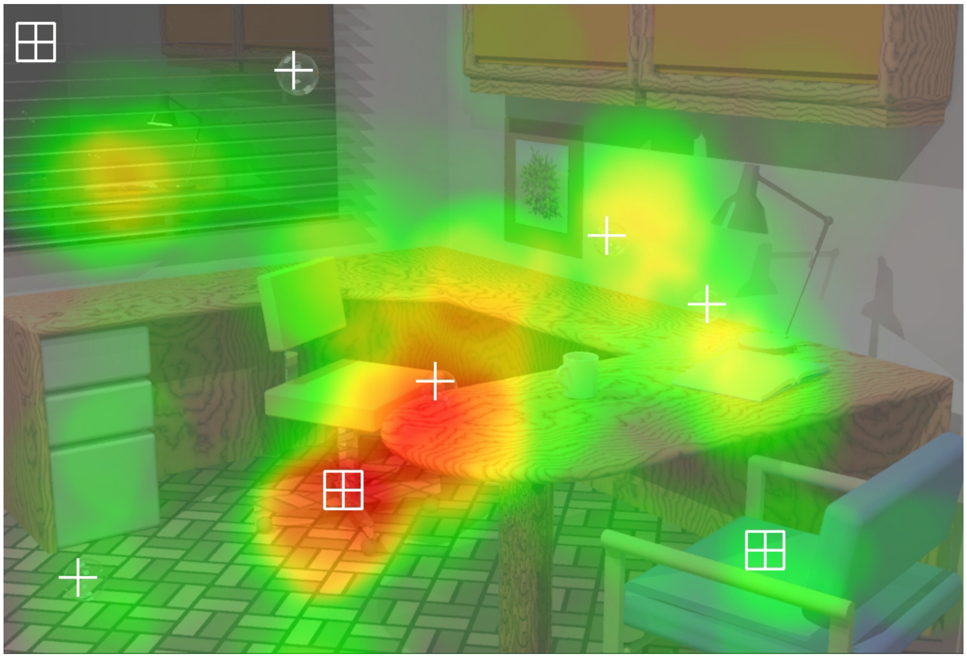

Fig. 10: *Experimental Results: This chart shows fixation distribution for one image in the distractor condition (the office scene). Fixation data is averaged over all particiants When compared to earlier results it can be seen that image coverage is greater, and this may contribute to greater accuracy on task. White crosses indicate the modulations on target, while the bounded white crosses indicate a* distractor *modulation.*

guidance technique.

We have only just begun to explore the potential use of Subtle Gaze Direction in the medical arena. In future experiments we would like to include data from medical applications and focus on the detection of breast cancer, tumors, or other anomalies that may be hard to detect in medical imaging. Coupling Subtle Gaze Direction with robust vision algorithms, that occasionally return false positives, could prove an invaluable tool for training for medical personnel. This will be particularly useful for complex or noisy images where features are difficult to decipher. The results from vision algorithms could be used to select the modulated regions in the image, and even with false positives (or distractors) present, performance on search tasks could potentially be enhanced.

One further possible line of inquiry would be to examine the usefulness of Subtle Gaze Direction in imagery where the task involves identifying or separating targets from non-targets. By applying Subtle Gaze Direction in target regions, gaze could be directed only to the targets, making them more distinguishable from non-target regions. Another interesting problem is that of moving targets, or moving imagery in general. Future experiments will focus on the performance of Subtle Gaze Direction in dynamic environments, such as animations and interactive environments. Subtle Gaze Direction could be used to help guide a user's navigation, or to highlight those parts of an animation that are relevant to the application. We would expect that stronger modulations would be necessary in order for Subtle Gaze Direction to be effective in dynamic scenes.

We have shown that Subtle Gaze Direction can improve people's performance on a counting task without noticeably changing the image. Possible uses of this technique might be to guide gaze in more complex visual search tasks where targets are numerous or difficult to identify.

REFERENCES

BAILEY, R., MCNAMARA, A., SUDARSANAM, N., AND GRIMM, C. 2009. Subtle gaze direction. To appear.

BILSKA-WOLAK, A., FLOYD, C., NOLTE, L., AND LO, J. 2003. Application of likelihood ratio to classification of mammographic masses; performance comparison to case-based reasoning. *Medical Physics 30*, 949–958.

BOOT, W. R., NEIDER, M. B., AND KRAMER, A. F. 2009. Training and transfer of training in the search for camouflaged targets. *Attention, Perception and Psychophysics 71,* 4, 950–963.

CATER, K., CHALMERS, A., AND LEDDA, P. 2002. Selective quality rendering by exploiting human inattentional blindness: looking but not seeing. In *VRST '02: Proceedings of the ACM symposium on Virtual reality software and technology*. ACM, New York, NY, USA, 17–24.

DECARLO, D. AND SANTELLA, A. 2002. Stylization and abstraction of photographs. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*. ACM Press, New York, NY, USA, 769–776.

DUCHOWSKI, A. T. 2002. A breadth-first survey of eye-tracking applications. *Behavioral Research Methods 34,* 4 (November), 455–470.

EL-BAZ, A., FARAG, A., GIMELFARB, G., FALK, R., AND ELDIASTY, T. 2006. A framework for automatic segmentation of lung nodules from low dose chest ct scans. *Pattern Recognition, International Conference on 3*, 611–614.

KIM, Y. AND VARSHNEY, A. 2006. Saliency-guided enhancement for volume visualization. *IEEE Trans. on Visualization and Computer Graphics (IEEE Visualization 2006) 12, No. 5*, 925–932.

KOSARA, R., MIKSCH, S., AND HAUSER, H. 2001. Semantic depth of field. In *IEEE Symposium on Information Visualization 2001 (InfoVis 2001)*. IEEE, San Diego, CA, USA.

LARSON, G. W. AND SHAKESPEARE, R. 2004. *Rendering With Radiance: The Art And Science Of Lighting Visualization*. Booksurge Llc, New York.

LUEBKE, D., WATSON, B., COHEN, J. D., REDDY, M., AND VARSHNEY, A. 2002. *Level of Detail for 3D Graphics*. Elsevier Science Inc., New York, NY, USA.

MACK, A. AND ROCK, I. 1998. *Inattentional Blindness*. MIT Press, MIT.

MANCAS, M., GOSSELIN, B., AND MACQ, B. 2004. Automatic fast detection of tumor suspect areas on ct scan. *Visualization Conference, IEEE 0*, 33p.

MITCHELL, G. E. 2004. *Taking Control Over Depth of Field: Using the Lens Blur Filter in Adobe Photoshop CS*. The Light's Right Studio.

NEIDER, M. B. AND ZELINSKY, G. J. 2006. Searching for camouflaged targets: effects of target-background similarity on visual search. *Vision Res 46,* 14 (July), 2217–2235.

OGDEN, T. E. AND MILLER, R. F. 1966. Studies of the optic nerve of the rhesus monkey: Nerve fiber spectrum and physiological properties. *Vision Research 6*, 485–506.

O'SULLIVAN, C., DINGLIANA, J., AND HOWLETT, S. 2003. *Eye-movements and Interactive Graphics*. Elsevier, Netherlands. J. Hyona, R. Radach, and H. Deubel (Eds.).

RAYNER, K. 1975. The perceptual span and peripheral cues in reading. *Cognitive Psychology 7*, 65–81.

SCHWANINGER, A., HARDMEIER, D., AND HOFER, F. 2004. Measuring visual abilities and visual knowledge of aviation security screeners. *Proceedings IEEE ICCST 38*, 258–264.

SCHWANINGER, A., MICHEL, S., AND BOLFING, A. 2007. A statistical approach for image difficulty estimation in x-ray screening using image measurements. In *APGV '07: Proceedings of the 4th symposium on Applied perception in graphics and visualization*. ACM, New York, NY, USA, 123–130.

SIMONS, D. AND CHABRIS, C. 1995. Gorillas in our midst: Sustained inattentional blindness for dynamic events. *British Journal of Developmental Psychology 13*, 113–142.

SPILLMANN, L. 1990. *Visual Perception: The Neurophysiological Foundations*. Academic Press, San Diego.

TOURASSI, G., VARGAS-VORACEK, R., CATARIOUS, D., AND FLOYD, C. 2003. Computer-assisted detection of mammographic masses: a template matching scheme based on mutual information. *Medical Physics 30*, 2123–2130.

ZHUOWEN, T., ZHOU, X., BOGONI, L., BARBU, A., AND COMANICIU, D. 2006. Probabilistic 3d polyp detection in ct images: The role of sample alignment. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on 2*, 1544–1551.

A. APPENDIX

The following instructions were given to participants in the study before they began the trial.

—You are going to see a series of images. These images have target regions in them that I would like you to count.

—The targets are transparent spheres, or bubbles, in the image. (note I showed them an example image with targets)

—After each scene you will see a black screen with a white cross in the middle. Take a moment to focus on the white cross.

—Do not count targets in reflections.

—If you feel tired or need a break at any time let me know.

—Do you have any questions?

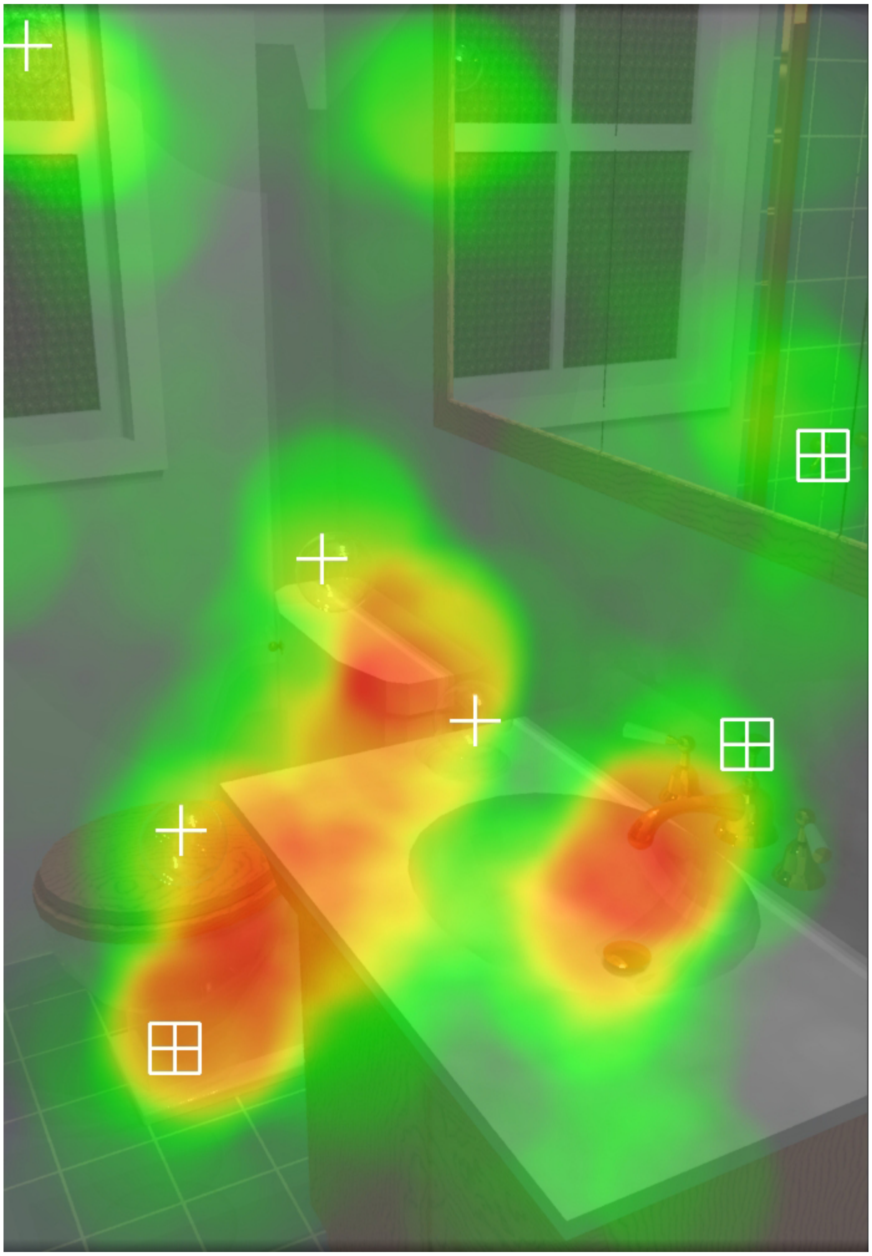

Fig. 11: *Representative Results [Bath]: These images represent average fixation distribution for each image in the distractor condition. White crosses indicate the modulations on target, while the bounded white crosses indicate a distractor modulation.*

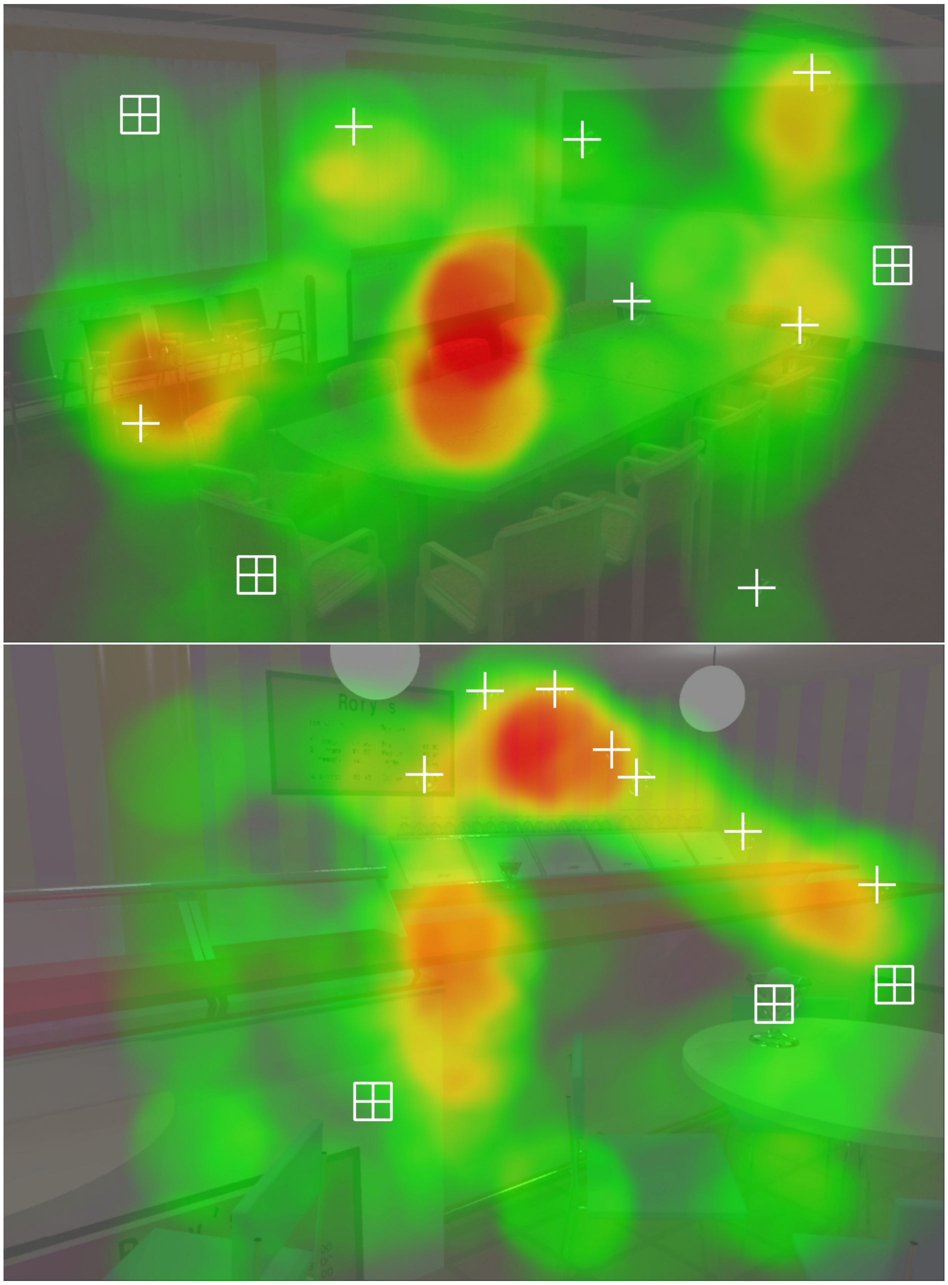

Fig. 12: *Representative Results [Conference top, Counter bottom]: These images represent average fixation distribution for each image in the distractor condition. White crosses indicate the modulations on target, while the bounded white crosses indicate a distractor modulation.*

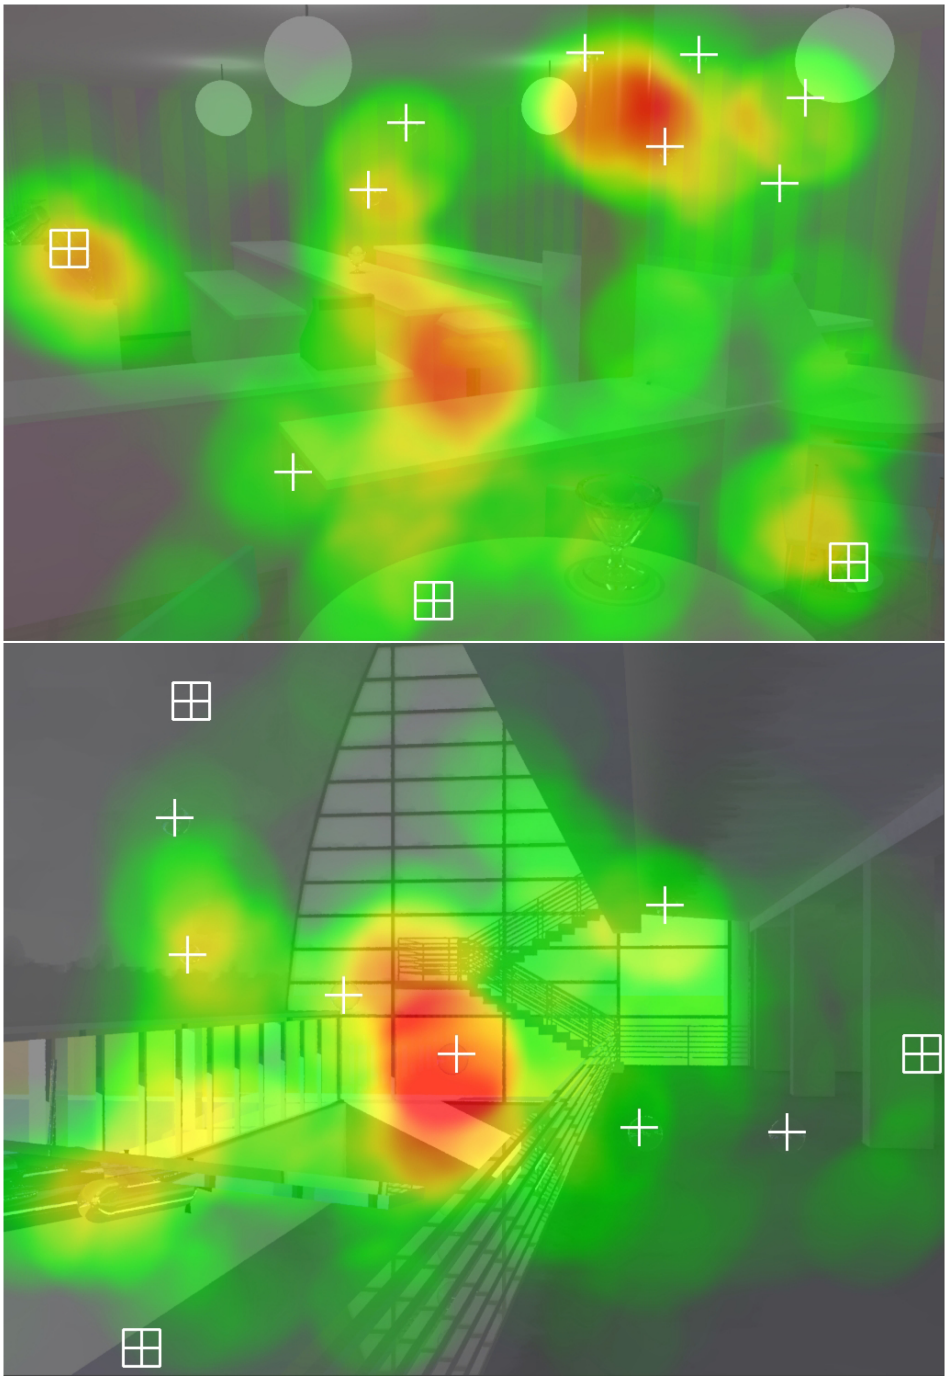

Fig. 13: *Representative Results [Soda top, Interior bottom]: These images represent average fixation distribution for each image in the distractor condition. White crosses indicate the modulations on target, while the bounded white crosses indicate a distractor modulation.*