SpiderVision: Extending the Human Field of View for Augmented Awareness

Kevin Fan[1], Jochen Huber[2,3], Suranga Nanayakkara[2], Masahiko Inami[1]

[1] Keio Media Design
4-1-1 Hiyoshi, Kohoku,
Yokohama, 223-8526 Japan

[2]Singapore University of
Technology and Design, 20
Dover Drive, 138682 Singapore

[3] MIT Media Lab
75 Amherst Street
Cambridge, MA 02139 USA

{kevinfan | inami}@kmd.keio.ac.jp, jhuber@mit.edu, suranga@sutd.edu.sg

ABSTRACT

We present SpiderVision, a wearable device that extends the human field of view to augment a user's awareness of things happening behind one's back. SpiderVision leverages a front and back camera to enable users to focus on the front view while employing intelligent interface techniques to cue the user about activity in the back view. The extended back view is only blended in when the scene captured by the back camera is analyzed to be dynamically changing, e.g. due to object movement. We explore factors that affect the blended extension, such as view abstraction and blending area. We contribute results of a user study that explore 1) whether users can perceive the extended field of view effectively, and 2) whether the extended field of view is considered a distraction. Quantitative analysis of the users' performance and qualitative observations of how users perceive the visual augmentation are described.

Author Keywords

Visual Augmentation; Field of View Extension; Augmented Awareness; Visual Blending; Human Behavior.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

General Terms

Design; Human Factors;

INTRODUCTION

We live in a meaningful world. There are constantly interesting things happening around us which may only last an instant. However, we are ever so often focused on a certain task at hand and therefore unaware of things happening beyond our field of view. In particular, the human visual sense, being the most dominant one [14], is limited on the horizontal axis at 200 degrees [4, 16]. According to a US study [13], the number of pedestrians treated in the emergency room due to focusing on their mobile phone while walking has increased in the last years. We thus have to rely on a fusion of other senses in order to detect surrounding activities. However, this can also be problematic: (i) Senses can be temporarily occupied, e.g. when listening to music through headphones, and (ii) a non-visual sensory cue provides only a hint as to what is happening around us and lacks precise localization.

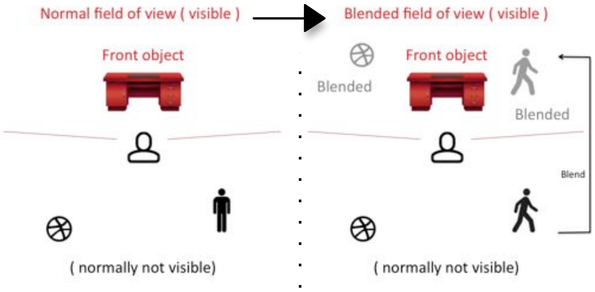

Figure 1. Extending the field of view for augmented awareness by blending.

Science fiction literature discusses these limitations by equipping protagonists with additional "super powers" that can overcome these limitations. For instance Spider-Man has a special sensory skill called "spider sense" that makes him aware of things happening beyond his visual field of view. In this paper, we draw on these ideas and contribute "SpiderVision", a wearable device that extends the human field of view to augment a user's awareness of things beyond the visual field of view.

SpiderVision has two main design goals. First, users should not be distracted while interacting within their ordinary field of view. Second, users should be able to detect activities outside the field of view without extra effort. SpiderVision leverages a video-see-through system to provide the user with an ordinary front view, while the field of view is extended through blending both back view and front view only when an activity is detected (Figure 1). Through extending the field of view only when necessary, users have the advantage of being able to interact normally with their ordinary front view and still be aware of areas beyond their field of view.

We conducted a user study to investigate two major research questions: 1) can users perceive the blended view from the back as extended field of view and react accordingly, and 2) is the blended view distracting to users focusing on a primary task in the front view.

Our contribution is four-fold:

- Technical implementation of a system that extends the field of view without inducing extra visual load.
- A method that augments humans' awareness visually.
- A user study to understand how people perceive "having eyes at the back of their heads" and examine how different blending proportions of both front and back affect the perception.
- Qualitative observations of users that effectively adapt to using the back view are discussed.

RELATED WORK

Our research builds on the following research areas: 1) augmenting the field of view through extending or substituting the visual sense, and 2) augmented modalities in awareness perception.

Augmented Field of View

The emergence of head mounted displays (HMD) provides intriguing ways for visual immersion. Combined with cameras, video-see-through HMDs are widely used in augmenting the user's visual perception of the surrounding reality. While most video-see-through HMDs are used in Augmented Reality systems to superimpose information on the display, systems such as FlyVIZ suggest combining an omnidirectional camera and an HMD to extend a user's field of view [1]. In FlyVIZ, Ardouin et al. mount the omnidirectional camera on the top of the user's head and map the captured image to the HMD so that the user can observe a horizontal view of 360 degrees without turning. Mizuno et al. developed the Virtual Chameleon, which allows the users to operate two cameras separately as free-moving eyes to observe the environment with a different field of view [12]. In the art project, although not using a HMD to augment visual sense in real time, Bilal equipped a camera to the back of the head to record images that are usually not seen with normal eyes [2].

These systems require cognitive adaption effort and a certain training period in order to be used effectively. This is mostly due to being based on aspects not akin to our usual way of perceiving the world, i.e. they require users to process omnidirectional or visually separated images, as studied in neuroscience researches [7]. In our system, we allow users to see through the HMD just as though they are seeing with their normal eyes, and only blends in the extended view, in a one-to-one mapping of front and back, when needed. In this way, there will be less cognitive effort to adapt to the constantly altered view.

Augmented Awareness

Given the limited senses, humans cannot be fully aware of our surroundings if we cannot see, hear, or feel the change. This is even more of an obstacle for people who are impaired in one of the senses. There are many projects that aim at solving the awareness issue by augmenting one of the modalities. Heun et al. developed the Perifoveal Display, which augments the user's awareness in the peripheral vision [8]. Projected peripheral displays are also studied [3]. To augment the hearing impaired, a system that visualizes sound from the user's back is developed [9]. In the haptic sensation, Mateevitsi et al. equipped a user's whole body with tactile transducers to notify the user as objects come near [11].

The goal of our research is similar in concept. We aim to help users to be more aware by augmenting one of the senses, namely the most dominant one, the visual. While previous works focused mainly on the peripheral vision, our work goes beyond that and extends a user's awareness beyond the normal field of view.

DESIGN

An extended field of view is beneficial in that we can see more of our surroundings. However, this could be problematic as well, as an extended field of view also means an increased amount of visual information that a person has to process. Naturally, humans are not accustomed to seeing more than what they can with the normal field of view. Also, it would require more cognitive effort to adapt to an extended field of view if presented as-is. Moreover, we often just want to focus on working with our foveal vision. In that case, a constantly extended field of view would actually be a distraction.

Our approach therefore is to extend the field of view if, and only if, activities that are relevant to the user happen outside the normal field of view. This way, users are not disturbed by a constant additional information feed. We define the activities as e.g. any object movement because the human vision is particularly sensible to objects in the peripheral view, even when they are focusing on a primary task at hand. The system automatically handles both detection and blending and the user can leverage the benefit of an extended field of view.

SPIDERVISION

We developed a system that extends a user's field of view horizontally beyond that of the normal front view. We do so by blending in information that is contained in the back view, i.e. the portion that we cannot see with our normal eyes. The additional information is overlaid over the front view that we are accustomed to see with our eyes. The combined information is displayed to the user. As a result, a user can retain the normal vision, and has the benefit of an extended field of view (Figure 2).

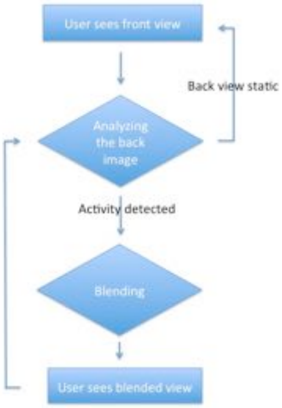

Figure 2. Blending the back view only when necessary

Overview

The system (Figure 3) consists of a head-mounted display (Oculus Rift Development Kit), two cameras (Logitech HD Pro C920), and two wide-angle lenses. The HMD has a field of view of 100 degrees. The two cameras have a diagonal field of view of 78 degrees, and have an aspect ratio of 16:9. The formula for calculating the horizontal field of view from a diagonal one is:

$$W = \frac{rD}{\sqrt{1+r^2}}$$

where W is the horizontal field of view, r is the aspect ratio, and D is the diagonal field of view. The two cameras are calculated to be around 68 degrees in the horizontal field of view. This is not close to the human's field of view; therefore, we modified the lenses of the cameras with commercial wide-angle lenses, which increase a camera's field of view by a scale of 1.5, to allow the cameras to capture around 102 degrees horizontally (Figure 4). Although this is still smaller than human's field of view, we do not increase the field of view further with fisheye lenses because studies have revealed that when there is a mismatch between the field of view of the HMD and the field of view of the image presented, users have a greater tendency to have sickness while wearing the HMD [5]. The two cameras, with modified lenses, are each installed at the front and the back of the HMD. The front camera constantly streams image to the HMD to act as the eye in a video see-through HMD. The back camera acts as the "eye" to the back of the head and constantly analyzes the captured image for detection of scene change, for example, a moving object. We opted for monocular vision and refrained from implementing stereoscopic vision because we focus on the movement of the scene in an extended field of view rather than the depth information. Movement can be detected with just a monocular vision, and we can have a simple sense of depth information through cues such as occlusion. The HMD either displays the front image only or a blended image of the front and the back to the user.

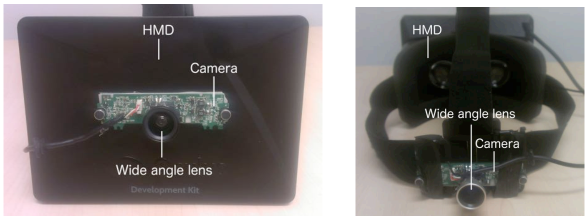

Figure 3. The SpiderVision System: front view (left), rear view (right).

Figure 4. Modifying the camera with wide-angle lens for wider field of view

Analyzing the Back Image

Back activities are only blended in when there is arbitrary object movement. Therefore, we need to constantly analyze the image captured by the back camera.

To detect an object movement, we first need to know the locations of the objects in the scene, after which we track to see if the locations of the objects move, or if new objects enter the scene captured by the back camera.

We implemented the Shi-Tomasi corner detection algorithm to extract the features of the objects presented [15]. The features extracted are from foreground objects that could possibly move, such as a standing person, and from background objects such as a painting on the wall. When new objects enter the scene, their features are extracted and tracked.

To actually track the extracted features, we leverage the optical flow of the scene, which is the apparent motion of objects induced by the relative motion between the camera and the objects. We implemented the Lucas-Kanade method, which combines information from nearby pixels of an extracted feature, and compares the current frame and the previous frame to calculate the optical flow of the feature to see if the feature has moved [10]. In a static scene when the camera is not moving, a moved feature would indicate that the object this feature is associated with is moving, which we could then detect and blend the information to the front image (Figure 5).

In a dynamic scene, when the back camera is moving, such as when the user wearing the device is walking, every extracted feature has an optical flow as the scene captured is constantly changing. In this case, we need to differentiate

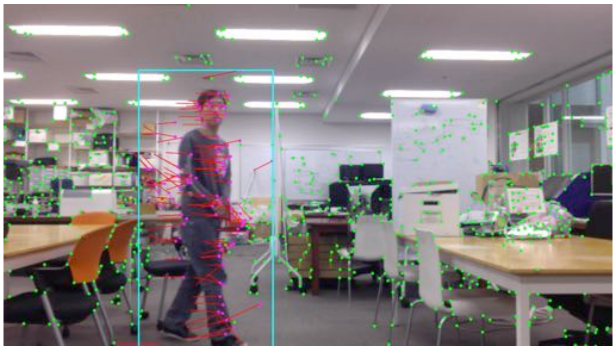

Figure 5. Using optical flow to detect and track arbitrary moving objects.

between the optical flow of the static background objects, and the possible optical flow if dynamic foreground objects are present. The optical flow of the static objects are only induced by the moving camera, while the optical flow of the dynamic objects are induced by both the moving camera and the dynamic objects. Therefore, we can extract and detect a moving object in a dynamic scene by comparing the difference in the velocity of the optical flow to see if an object's optical flow is different than that of the background. We calculate the optical flow into clusters based on difference in velocity. If the optical flow of all features are constant and belong in the same cluster, we can regard that as caused solely by camera movement, which should not trigger the blending. However, if a few flow are different, which would be caused by the camera movement and the self-motion, we can regard that as object movement to trigger the blend.

Blending

We exemplarily illustrate two methods of blending the information from the back to the front image: (i) visualization of the actual activity video feed as an overlay on the camera image and (ii) visualization as an abstract wedge [6].

Camera image

We can directly blend the back image with the front image by adjusting the alpha values of each image so that the images are translucent. In this method, the users could observe both the front image and the back image simultaneously, just as though they have eyes on the back of their heads. The factors of blending the images we can adjust include the area of the blended image, such as blending only the 25 degrees to the left and right while keeping the center 50 degrees to be front image only so we can still focus, and the visual dominance of the images, such as increasing the clearness of the back image and decreasing that of the front image so we can perceive the back image more easily when attention is needed.

The key point of blending the back image is that the back and the front should be of one-to-one mapping, so that a 100-degree back image is mapped to a 100-degree field of view of the normally observed front image. If we choose to only blend a portion of the back image, for example, 20 degrees, then only the same 20 degrees of the front image will be blended. This is crucial in keeping the field of view of both front and back constant, so that the users can perceive the "back eye" as the same as the front eye and can spend less cognitive adaption effort.

Abstract wedge

Extended field of view not only can be represented by blending the camera image directly, but can also be represented by an abstract notice displayed to the user's front view. The abstract notice does not display directly the back image so the users cannot actually see the back without turning their heads, but has the same benefit of an extended field of view by allowing the users to be aware of changes that they cannot observe with their normal eyes.

For the abstract representation, Gustafson et al. suggested an effective wedge method to visualize both the direction and distance of an off-screen object [6]. We display the wedge in our front view, as we humans are quite capable of noticing movements in that area. We use the wedge angle to suggest to the users which way they should turn their heads, and wedge size to indicate by how much they should turn in order to see the change as analyzed by the back camera.

USER STUDY

We conducted a user study to examine the effect of different factors of the extended field of view on augmented awareness. Our main goal was to study 1) whether users perceive the blended view from the back and react accordingly, and 2) whether the blended view is distracting to users focusing on a primary task in the front view. We opted for a setup with a primary and secondary task to better investigate cognitive demand. We observed how participants performed in a front view primary task while simultaneously being aware of the secondary task in the back.

Participants

We recruited 8 male participants. They are aged between 22 to 38 years old (avg: 28, SD:6.2). All participants had no previous knowledge of using the system. 2 participants had no previous experience with an HMD.

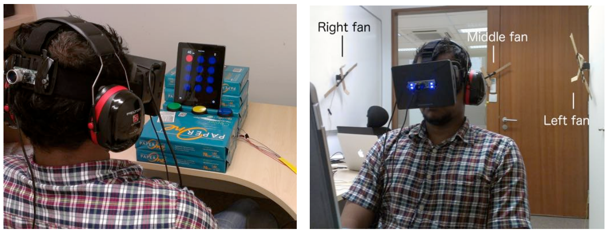

Figure 6. User study: primary task (left), secondary task (right).

Environmental Setup

The study was conducted in a small office as to provide a closed environment. The participants were seated in front of a desk. An iPad and 3 push buttons were placed on the desk. To the back of the participants are 3 fans. One was placed at 2 meters behind to the left, one at 2.2 meters behind to the right, and one at 2.5 meters directly behind.

Experimental Design and Method

We designed the user study as a 2x2x2 within subjects factorial design. The factors were blended view abstraction level (camera image (C), abstract wedge (W)), the blending level (foveal (F), peripheral (P)), and the visual noise level of the environment (quiet (Q), noisy (N)). The participants wore a HMD and a noise-cancelling headphone while completing two tasks.

The primary task of the participants was a visually demanding task where they had to click on as many red dots as possible on the iPad in a time frame of 1 minute. A red dot does not disappear unless clicked. Once clicked, a new red dot would appear at a random position selected from 12 blue dots. We measured the total number of red dots clicked as the performance of the primary task.

While focusing on the primary task, the participants also had a concurrent secondary task. In the task, the participants were instructed to monitor 3 fans, which were initially at rest. The participants had to press the corresponding left, middle, right buttons placed before them if they notice a fan is activated, which would deactivate the fan. The fan would deactivate automatically after 5 seconds if not deactivated by the participant. We measured the rate of correct detection and also the response time.

The combination of the factor abstraction level and blending level give a total of 4 kinds of representation in the HMD (Figure 7). When the abstraction level is real camera image, the back image presented is either blended fully as a translucent image that would obstruct the front foveal view, or blended to only replace the peripheral areas of the front view. When the abstraction level is abstract wedge, the wedge is either blended in the middle, or in the corresponding peripheral area of the direction. The wedges are colored red or blue. We intentionally choose the same colors as the primary task where red indicates an activating fan, and blue indicates noise. The difference in the wedges is explained to the participants beforehand. The visual noise factor is that no movements except the fans would be present in a quiet trial, and the experimenter would walk behind the participants to one end of the room and back, which takes 5 seconds, in a noisy trial.

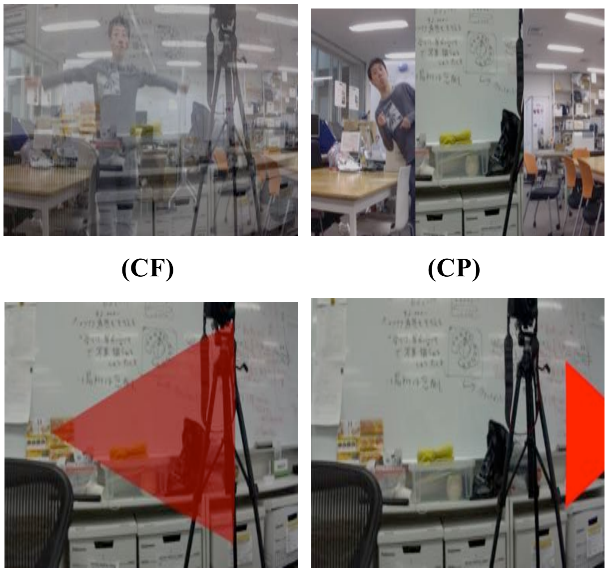

(CF) (CP)

(WF) (WP)

Figure 7. Blending methods: foveal camera (CF), peripheral camera (CP), foveal wedge (WF), peripheral wedge (WP).

The factorial design gives a total of 8 experimental conditions. The order at which the conditions are experimented on the participants is balanced according to a balanced Latin Square design. In each condition there are 4 trials, corresponding to each of the fans activating plus one where no fan is activated. The order of the 4 trials in a condition, as well as the activation time of the only activated fan, are randomized. Each trial is one minute, and the participants can rest after each trial if needed.

The participants began with a base trial where each of the fans was randomly activated for 5 seconds but the back camera did not blend the back view. The base trial was to measure the participants' primary task performance with no distraction and also to confirm that they could not hear the activation of the fans with the noise-cancelling headphone, since our focus is on augmented awareness through visual. The participants then conducted 4 trials in each 8 conditions. After all trials were conducted, a trial with no secondary task was conducted to measure the primary task performance again to account for the learning effect. The participants were asked to fill a questionnaire after the experiment.

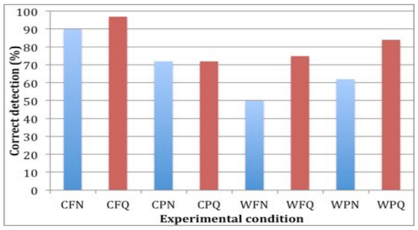

Figure 8. Correct detection of the activating fans.

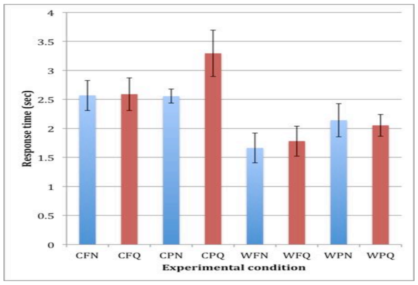

Figure 9. Response time of the correct detections (95% confidence interval).

RESULTS

Correct Detection

To visualize how humans can benefit from augmented awareness through an extension of the field of view, we measured and compared the correct detection of the activating fans that were at the back and out of the normal field of view. We observed that the participants were most successful in detection when presented with a full translucent image of the back camera that also blended the foveal area in a visually quiet environment. The condition with a wedge in the foveal area was the least successful, with only half correct detection. ANOVA analysis revealed that there is a main effect for abstraction level ($F(1, 248) = 8.29$, $p < 0.05$) and for visual noise level ($F(1,248) = 6.65$, $p < 0.05$). We noted that camera view generally yields a better detection than wedges as the participants could actually see and perceive the image captured by the camera. During the interview after the experiment, 2 participants noted that when the wedges were displayed in the foveal area, the participants might have occasionally pressed wrong buttons because the wedges were equilateral triangles and were hard to perceive the direction. There is also an interaction between abstraction level and blending level ($F(1, 248) = 10.12$, $p < 0.05$), and between abstraction level and visual noise level ($F(1, 248) = 3.88$, $p < 0.05$). When the abstraction level is camera view and blending level is peripheral, the participants could not see their direct back without turning their heads and utilize the back camera. Most of the participants continued to focus on the primary task and ignored the blended view even though the direct back fan was activated. A significant difference in detection was in the conditions of wedges in visually noisy trials. We learned that participants were confused when there was visual noise and were too many constantly changing wedges to perceive.

Response Time

We measured the response time of the correct detections from when a fan is activated to the time of the button press to analyze the time frame that a participant needs to perceive the extended view. ANOVA on correct detected trials revealed a significant main effect for abstraction level ($F(1, 135) = 33.72$, $p < 0.001$). We observed that the camera images have an average slower response time. This could be accounted to the fact that the participants had to perceive both the front and back images. While in wedges, the participants could treat the wedges as augmented instructions and react intuitively to the wedge directions.

Primary Task Performance

To measure the distraction induced by the extended field of view, we compared the number of red dots clicked in conditions camera view (mean: 43.22, SD: 8.70), abstract wedge (mean: 44.72, SD: 8.56), and an absence of the secondary task (mean: 45.90, SD: 8.35). ANOVA analysis did not reveal a significant main effect of the conditions. Comparison between visually noisy trials and quiet trials did not reveal a significant main effect. This suggests that the participants could handle an additional low-occurrence secondary task to the primary task with reasonable focus when awareness is augmented. The slightly higher mean performance in an absence of secondary task could be inferred, as the participants did not have to react by pressing the buttons in this condition.

Task Load

We asked the participants to complete a NASA-TLX questionnaire after the experiments. The scale was from 0 (least load) to 10 (highest load). We compared the load between the four representation conditions: foveal camera (CF), peripheral camera (CP), foveal wedge (WF), and peripheral wedge (WP). The foveal camera representation required a decent amount of mental load, as the participants had to perceive two fully blended translucent front and back images. The participants were also most confident in their performance in this case because they noted that they could

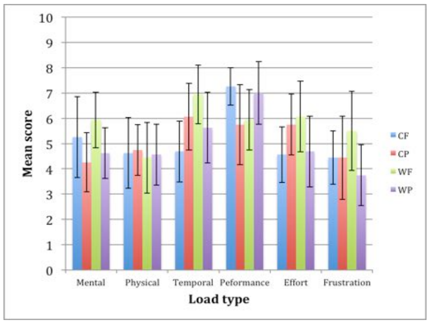

Figure 10. Subjective task load questionnaire (95% confidence interval).

see all of their back. Their confidence was reflected in the high detection of the activating fans. The peripheral camera representation had the least mental load as most participants said that they could intuitively see the peripheral back images clearly when compared to the translucent blending. However, the participants felt they had to spend effort to see their direct back and therefore felt least confident in their performance. The foveal wedge representation was the most difficult in terms of highest mental load, highest effort, highest frustration, and low confidence in performance. Most participants noted that the wedges occluded their view of the iPad and the colors of the wedges, being almost the same as the dots in the primary task, seemed to be difficult to distinguish. This was reflected in the low detection of the activating fans. The peripheral wedge representation did not have the same defect as the wedges were placed in the peripheral areas of the view. Many participants also noted that the placement of the wedges (left, top, right) was easy to perceive and differentiate.

DISCUSSION

Feedback

From the feedbacks of the 8 users in the user study, 5 prefer the foveal camera representation, 2 prefer the peripheral camera representation, and 1 prefers the peripheral wedge representation. People who prefer the foveal camera representation state that they feel as though they could see all around them and have a better advantage in awareness. The 2 who prefer the peripheral camera representation are adept in utilizing the back camera movements and are not affected by the occluded direct back, and they prefer seeing in complete opacity. The user who prefers the peripheral wedge representation perceives the wedges as augmented reality.

Eyes on the Back

The back camera in SpiderVision acts as an “eye” on the back of the head. A question of high interest, therefore, is to see whether the users can gain the most benefit through utilizing this “eye” the same as the normal eyes. During our user study, we observed an interesting phenomenon that, although the users were not taught that they could turn their heads to see the originally not visible middle fan in the peripheral camera representation, 4 users were adapted to utilizing the back camera in such a technique after only a few trials; among them, 2 were adept in detecting when the middle fan is activated (Figure 11). This observation suggests that humans could adapt to having and utilizing eyes on the back of our heads, and could get accustomed to using the back eyes to observe the back instead of turning 180 degrees to use the front eyes.

Augmented Awareness

As each of the human’s senses is limited, humans usually use a combination of multiple sensory feedbacks to gain an awareness of the surrounding. However, in the case where humans are focused in their visual sense, or when senses such as auditory are blocked, humans are not capable of being aware of the activities to their back. SpiderVision takes the approach of augmenting the visual sense for augmented awareness only when necessary. This approach is suggested to be effective from the user study in that the users performed similarly on the primary task while at the same time were able to detect and perceive the correct activities more than half of the time, which would not be possible had their visual senses not been augmented. For example, imagine the users having to constantly take their eyes off the primary task to turn and check their back and then turn to focus on the primary task again, which would result in either lower performance or missed activities in the back, or both.

Figure 11. Turning the head to use the back camera to see the originally occluded middle fan on the door.

Limitation and Future Work

Our system uses the cameras extensively. One limitation is that just as humans have eyesight that limits how far and clear they can see, cameras also have characteristics such as resolution that limits the quality of the image to be analyzed, which in turn affects the correct detection of activities by the system. An object, with the size of 5 cm in diameter, moving at 5 meters away was not detected by the system. Also, activities that do not involve changes in visual, such as only changes in sound, are not detectable. We are looking to add more detection methods, such as detecting audio change, as a trigger to the visual blend. Another limitation is that since the back camera is fixed at the back of the head, the extended field of view is invert from head motion. For example, to see more of the left using the back view, the user would have to turn right. However, as we observed in the user study, humans could adapt to this technique and utilize the back camera just as an eye on the back.

One interesting idea could be to combine the benefits of both the camera image and abstract wedges. Blending camera images provide a detailed view but occlude the front view more. Wedges are efficient in representing awareness out of the field of view with direction but the users cannot grasp exactly what the activities are. We could implement eye tracking so that the system could switch between representations of camera image and wedge as users gaze at the extended field of view. Also, the visualization of the abstract wedges can be improved to be silhouette as not to obstruct the normal view, and can be even more effective if the silhouettes are the shape of the actual objects in interest.

CONCLUSION

SpiderVision is a system that extends the human field of view and thus augments a user's awareness. For this purpose, we contributed a novel head-mounted system: SpiderVision. It leverages cameras both at the front and the back. The back camera continuously analyzes the scene behind the user, and only blends in the back view as an extended field of view when necessary. The blended view is invoked by activities such as object movement. The major idea is that the field of view is only extended on demand, constituting a minimal add-on to the user's cognitive load, while the user is able to focus on the front view ordinarily. SpiderVision is different from prior work that uses omnidirectional or free moving cameras since (a) it requires less cognitive effort and (b) users do not have to constantly process images that are not akin to their world perception.

We conducted a user study that focused on the perception of the extended field of view and investigated whether users can detect activities effectively. The study also looked at whether the extended field of view hindered users from performing primary tasks. The study results show that users were able to reliably and effectively detect movements behind their back using SpiderVision. An interesting phenomenon was also observed: users were able to adapt the back camera as an "eye" on the back of their heads. One important avenue for future research is to investigate how SpiderVision can be integrated into other see-through devices such as Google Glass.

ACKNOWLEDGMENTS

This work was supported by the International Design Center of the Singapore University of Technology and Design

REFERENCES

1. Ardouin, J., Lécuyer, A., Marchal, M., Riant, C., and Marchand, E. FlyVIZ: a novel display device to provide humans with 360° vision by coupling catadioptric camera with hmd. In *Proc. VRST '12*, ACM (2012), 41-44.
2. Bilal, W. 3RDI project. http://www.3rdi.me
3. Birnholtz, J., Reynolds, L., Luxenberg, E., Gutwin, C., and Mustafa, M. Awareness beyond the desktop: exploring attention and distraction with a projected peripheral-vision display. In *Proc. GI '10*, Canadian Information Processing Society (2010), 55-62.
4. Creem-Regehr, S. H., Willemsen, P., Gooch, A. A., and Thompson, W. B. The influence of restricted viewing conditions on egocentric distance perception: Implications for real and virtual environments. *Perception*, 34(2), 2005, 191-204.
5. Draper, M. H., Viirre, E. S., Furness, T. A., and Gawron, V. J. Effects of image scale and system time delay on simulator sickness within head-coupled virtual environments. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 43(1), 2001, 129-146.
6. Gustafson, S., Baudisch, P., Gutwin, C., and Irani, P. Wedge: clutter-free visualization of off-screen locations. In *Proc CHI '08*, ACM (2008), 787-796.
7. Harris, C. S. Perceptual adaptation to inverted, reversed, and displaced vision. *Psychological review* 72, 6 (1965), 419.
8. Heun, V., Kapri, A., and Maes, P. Perifoveal display: combining foveal and peripheral vision in one visualization. In *Proc. UbiComp '12*, ACM (2012), 1150-1155.
9. Ho-Ching, F. W., Mankoff, J., and Landay, J. A. Can you see what i hear?: the design and evaluation of a peripheral sound display for the deaf. In *Proc. CHI '03*, ACM (2003), 161-168.
10. Lucas, B. D., and Kanade, T. An iterative image registration technique with an application to stereo vision. In *IJCAI*, Vol 81, 1981, 674-679.
11. Mateevitsi, V., Haggadone, B., Leigh, J., Kunzer, B., and Kenyon, R. V. Sensing the environment through SpiderSense. In *Proc. AH '13*, ACM (2013), 51-57.
12. Mizuno, F., Hayasaka, T., and Yamaguchi, T. Virtual Chameleon-A System to Provide Different Views to Both Eyes. In *World Congress on Medical Physics and Biomedical Engineering* (2009), 169-172.
13. Nasar, J. L., and Troyer, D. Pedestrian injuries due to mobile phone use in public places. *Accident Analysis & Prevention*, 2013.
14. Posner, M. I., Nissen, M. J., and Klein, R. M. Visual dominance: an information-processing account of its origins and significance. *Psychological review*, 83(2), 1976, 157.
15. Shi, J., and Tomasi, C. Good features to track. In *Proc. CVPR'94*, IEEE (1994), 593-600).
16. Wandell, B. A. Foundations of vision. Sinauer Associates, 1995.