ELSEVIER

Applied Ergonomics 31 (2000) 523–530

APPLIED ERGONOMICS

www.elsevier.com/locate/apergo

Visibility of road hazards in thermal, visible, and sensor-fused night-time imagery

Jason S. McCarley*, William K. Krebs

Operations Research Department, Naval Postgraduate School, Monterey, CA, USA

Received 3 September 1999; accepted 22 January 2000

Abstract

Sensor fusion combines the output of multiple imaging sensors within a single composite display. Ideally, a fused image will retain important spatial information provided by individual input images, and will convey useful spatial or chromatic emergent information derived from the contrast between input images. The present experiment assessed the potential benefits of sensor fusion as a method of enhancing drivers' night-time detection of road hazards. Observers were asked to detect a pedestrian within thermal and visible images of a night-time scene, and within chromatic and achromatic renderings created by sensor fusion of grayscale thermal and visible images. Results indicated that fusion can both improve spatial image content, and can effectively embellish spatial content with emergent chromatic information. The benefits of both sensor fusion and of color rendering, however, were inconsistent, varying substantially with quality of input images submitted for fusion.

Keywords: Night vision; Sensor fusion

Though traffic accidents can occur under any conditions, evidence indicates that drivers' risks increase substantially under conditions of low light. Under night-time illumination, drivers suffer decreased acuity and contrast sensitivity (Sturgis and Osgood, 1982; Sturr et al., 1990; Andre, 1996), and may have these difficulties exacerbated by glare such as that produced by headlights of oncoming vehicles (Sturgis and Osgood, 1982; Andre, 1996). Further, the illumination provided by a driver's own headlights is often insufficient to compensate for this visual loss (Owens and Sivak, 1983). Accordingly, Owens and Sivak (1996) found that traffic fatalities are disproportionately likely to occur under low illumination, and that this tendency is exaggerated among collisions involving inconspicuous obstacles such as pedestrians or pedalcyclists rather than larger, more visible obstacles such as motor vehicles.

Efforts to increase the visibility of night-time road hazards have led to development of long-wave infrared (IR) imaging systems for use onboard vehicles (Piccione and Ferrett, 1998). Because long-wave IR sensors respond to emitted thermal energy, rather than reflected light, their performance is unhindered by darkness or glare. Obstacles that are undetectable to unaided vision can therefore be readily visible in a thermal image. Nonetheless, thermal imagery may not be wholly optimal as an aid to night driving. Since the luminance contrast of an object in an emitted-IR image is determined by the difference in temperature between the distal object and its background, obstacles of low thermal contrast may be indiscernible in an IR image. Furthermore, since thermal imaging systems respond to energy within a single waveband, they provide only monochromatic imagery.

An alternative form of imagery potentially useful as an aid to night-time driving might come from sensor fusion of multiple images. Through sensor fusion, images collected with sensors of differing spectral sensitivities are combined and presented within a unitary display. Spatially registered thermal and visible renderings of scene, for example, might be merged within a single composite image. Such processing could improve image quality at least two ways. First, sensor fusion could allow observers to easily and simultaneously view information provided by more than one sensor. Thus, an obstacle visible in any single component image might be visible in the fused image, even if it is invisible in alternative component

* Corresponding author. Current address: Beckman Institute, University of Illinois at Urbana-Champnign, 405 North Mathews Avenue, Urbana, IL 61801, USA.

E-mail address: mccarley@uiuc.edu (J.S. McCarley).

0003-6870/00/$ - see front matter
PII: S0003-6870(00)00010-7

images. More intriguingly, fusion algorithms could exploit differences between component images to derive emergent information not available within any component image singly. Contrast between component images, for example, could be used to sharpen spatial detail within a fused image (Waxman et al., 1997), or could provide the basis for chromatic rendering of a fused image much as differences in the output of retinal receptors allow for biological color vision (Waxman et al., 1997; Scribner et al., 1996).

There is little guarantee, however, that sensor fusion will enhance image quality. The emergent information derived through fusion could be of little salience to a human observer, or of little relevance to an observer's visual task. Worse, fusion might degrade task-relevant information provided by individual input sensors (e.g., by reducing contrast or resolution) (Toet and Walraven, 1996), and thus could actually impair visual performance. These concerns mandate that extensive psychophysical testing precede deployment of any sensor fusion system as an aid to human perception. Unfortunately, past research has at best produced equivocal evidence for benefits of sensor fusion. Although a number of studies have assessed the value of sensor fusion using various measures of image quality (e.g., Ryan and Tinkler, 1995; Steele and Perconti, 1997; Toet et al., 1997; Essock et al., 1997; Essock et al., 1999), their results have been largely inconsistent – sometimes favoring sensor-fusion, other times not – and differences in their methods and stimuli make them difficult to compare. Further, a number of studies have failed to distinguish effects of chromatic information within color-fused images from the effects of spatial information, or have used psychophysical methods that were not criterion-free. Even the results of studies which have found benefits of sensor fusion may be therefore difficult to interpret.

The present research was conducted to address these concerns, and more specifically to assess the prospective utility of sensor-fused thermal/visible imagery as an aid to detection of night-time road hazards. Observers were presented night-time images of a road viewed from a driver's perspective, and were asked to detect a pedestrian within the depicted scene. Performance was measured with thermal and visible images, and with composite images derived from fusion of thermal and visible images. To allow the effects of spatial and chromatic information within fused images to be disentangled, fused images were presented in both color and grayscale formats. Finally, to begin delineating variables which might mediate quality of sensor-fused images, performance was tested with images characterized by varying levels of illumination (low, moderate, and excessive). Two methods of fusion were tested, one of which (principal components fusion) has been assessed both favorably and unfavorably in past research (Steele and Perconti, 1997; Essock et al., 1999).

1. Method

1.1. Observers

Nine male observers, age 30–38 yr, were recruited from among the students and faculty of the Operations Research Department at the Naval Postgraduate School. All observers had normal or corrected-to-normal visual acuity, and had normal color vision as tested with pseudoisochromatic plates. All observers granted informed consent prior to participation. All observers but one were naïve to the experimental hypothesis.

1.2. Apparatus

Stimuli were displayed by a VisionWorks computer graphics system (Vision Research Graphics, Inc., Durham New Hampshire; Swift et al., 1997) on a Nanao Flexscan F2.21 monitor. The monitor had a resolution of 800×600 pixels, a frame rate of 98.9 Hz, and a maximum luminance of 100 cd/m^2. Luminance was linearized by means of a look-up table. Observers viewed the screen from a distance of approximately 1.5 m.

1.3. Stimuli

Stimuli were visible, thermal, and sensor fused visible/thermal images of an outdoor night-time scene arranged to simulate the forward view of a road seen through the windshield of a vehicle. Imagery was collected with a visible (Nikon Action8 — VN750 Hi8 video recorder with a $10\times$ zoom lens) and a thermal camera (LMIRS LTC-500 uncooled forward long-wave infrared, with a 327×245 pixel element and a spectral sensitivity between 8 and 12 μm). Sensors were mounted vertically with matched fields-of-view of 33.7° horizontal and 26° vertical. Imagery was originally recorded to Hi-8 video. Single frames were later digitized for use as experimental stimuli.

Stimuli were collected on the campus of the Naval Postgraduate School, Monterey, CA, between the hours of 9:00 and 11:00 PM, April 12, 1999. As reported by the National Climatic Data Center, temperature at time of data collection was 10.6° C and skies were overcast. Images depicted a segment of road roughly parallel to the sensors' line of sight, with an opposing vehicle stationed at a distance of approximately 68.5 m and facing directly into the sensors' field of view. Throughout stimulus collection, the imaged scene was illuminated by the low-beam headlights of a sport utility vehicle located directly behind the sensors. The intensity of the opposing vehicle's headlights was varied to allow stimulus collection at three levels of illumination: low, moderate, and excessive (i.e., glare). Imagery of low illumination was collected with the opposing vehicle's headlights off, and only its parking lights illuminated. Imagery of moderate and excessive illumination, respectively, was collected with the opposing vehicle's headlights on low-beam and high-beam. At each level of illumination, twenty-three pairs of synchronized visible and thermal frames were digitized for use as experimental stimuli. Of the images chosen at each level of illumination, ten depicted a pedestrian at a distance of 22.9 m perpendicular from the sensors, ten depicted the pedestrian at a distance of 38.1 m perpendicular from the sensors, and three depicted no pedestrian. The pedestrian, when visible, could be located at any of several lateral positions across the width of the road. Visible images were converted to grayscale during digitization. After digitization, any spatial misregistration that had obtained between pixels in raw thermal and visible images despite the sensors' matched fields-of-view was corrected, and images were contrast-enhanced.

For data collection, stimulus images were presented to observers in visible and black-hot thermal formats, in chromatic composite formats derived from sensor fusion of visible and thermal images, and in achromatic composite formats derived from grayscale rendering of chromatic fused images. All images had dimensions of 560×432 pixels, and subtended a visual angle of $11.2^\circ \times 8.64^\circ$ from a viewing distance of 150 cm. Fig. 1 presents a pair of corresponding thermal and visible stimulus images at the highest level of illumination. Fusion of paired visible and thermal images was performed through the simple and principal components fusion algorithms of Scribner and colleagues. Algorithms are described briefly here and explained in detail elsewhere (Scribner et al., 1993, 1996). Both algorithms map pairs of grayscale visible and thermal images onto a two-dimensional color space with principal axes corresponding to the complementary colors red and cyan. Simple fusion does this by merely assigning pixel values from a thermal component image to the red phosphor of a fused image, and assigning pixel values from a visible component image to the green and blue phosphors of a fused image. The result is a composite image wherein colors range from saturated red (created by illumination of only the red phosphor) through gray to saturated cyan (created by illumination of only the green and blue phosphors), and orthogonally, range from dim to bright; pixels that are bright only in the thermal input appear red in the fused image, pixels that are bright only in the visible input appear cyan, and pixels whose values are approximately the same in both input images appear achromatic, while pixels whose mean thermal/visible input values are low appear dim and those whose mean thermal/visible input values are large appear bright. The brightness of a fused pixel, that is, is determined by the weighted mean value of corresponding visible and thermal pixels, and hue is determined by the difference in value between corresponding visible and thermal pixels.

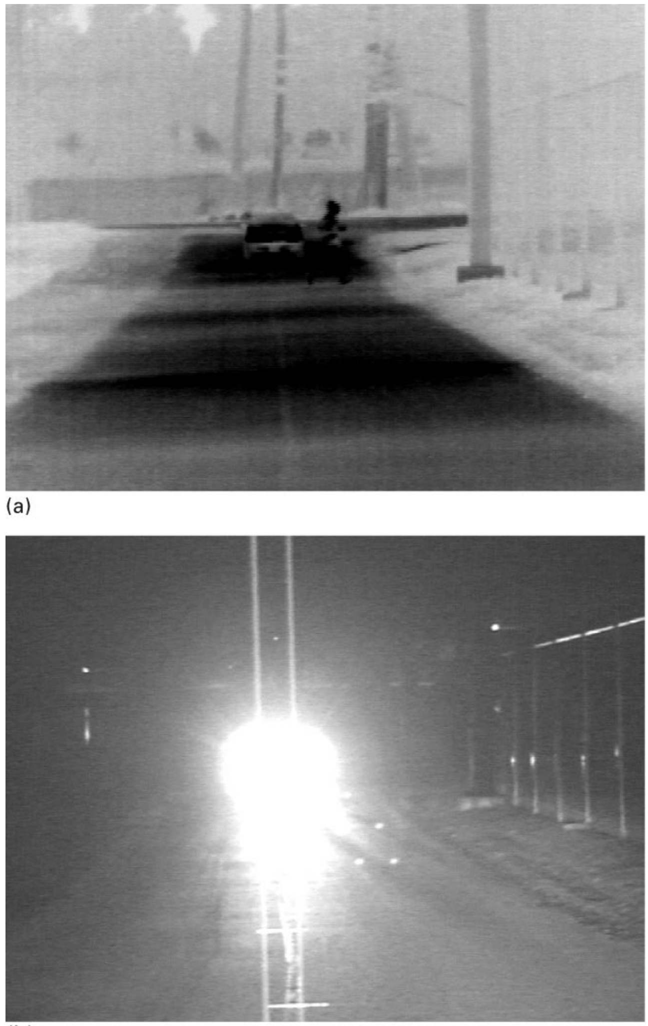

(a) (b)

Fig. 1. A pair of corresponding thermal (top) and visible (bottom) component images of excessive illumination.

Principal components fusion differs from simple fusion in attempting to improve image quality by manipulating the relationship between red and cyan pixel values. More specifically, the principal components fusion first maps input pixel values onto a red/cyan color space, taking thermal pixel values as distances in the red direction (red phosphor values) and visible pixel values as distances in the cyan direction (green and blue phosphor values). The algorithm then normalizes the distribution of pixel values in the direction orthogonal to brightness direction of the red/cyan space, compressing or expanding the distribution to either increase or decrease the correlation between red and cyan values at each pixel. The amount and direction by which the distribution of pixel values is manipulated is determined manually for a given set of images, and varies with the waveband of input images. Within the set of images employed here, principal components fusion generally decreased the strength of correlation between paired input pixels, but

Table 1
Mean correlations between red and cyan pixel values in fused images of varying illumination[a]

Illumination	Simple fusion	Principle components fusion
Low	− 0.50	0.17
Moderate	− 0.65	0.20
Excessive	− 0.69	0.24

[a]Note. $n = 23$ for each cell. SD = 0.01 for each cell.

reversed the direction of that correlation from negative to positive. Table 1 presents mean values of the correlation between red and cyan pixel values for simple fused and principal components fused images of varying illumination.

Grayscale versions of color-fused images were created with commercial image processing software (Adobe Photoshop 5.0). Chromatic and achromatic versions of each image were matched in pixel-by-pixel luminance, and thus contained identical spatial information. Fig. 2 presents grayscale simple fused and principal components fused images derived from the thermal and visible images presented in Fig. 1.

1.4. Procedure

The experimental task asked observers to search stimulus images for a pedestrian. Each trial began with an alerting tone and presentation of a fixation cross. After 500 ms, the fixation cross was replaced by a stimulus image which remained visible until the observer's response. The observer's task was to provide a manual response indicating whether or not a target, the pedestrian, was present in the image. Observers were asked to press '1' on the numeric keypad of a standard PC keyboard if a target was present, and to press '2' if a target was absent. Error rates and reaction times (RTs) were recorded. Observers were asked to provide responses as quickly as possible while maintaining a high level of accuracy. At the beginning of an experimental session, the observer performed a block of twenty practice trials with images drawn at random from the pool of stimuli. Thereafter, the observer performed two blocks of 180 trials each, with each block comprising 120 target-present and 60 target-absent trials. Within a block, each of twenty target-present images appeared once in every stimulus format, and each of three target-absent images appeared multiple times. To prevent observers from learning to make responses based on idiosyncratic image features, no feedback was provided following any response. Individual trials were separated by intervals of approximately 1000 ms. Observers were allowed periodic rest throughout the experimental session.

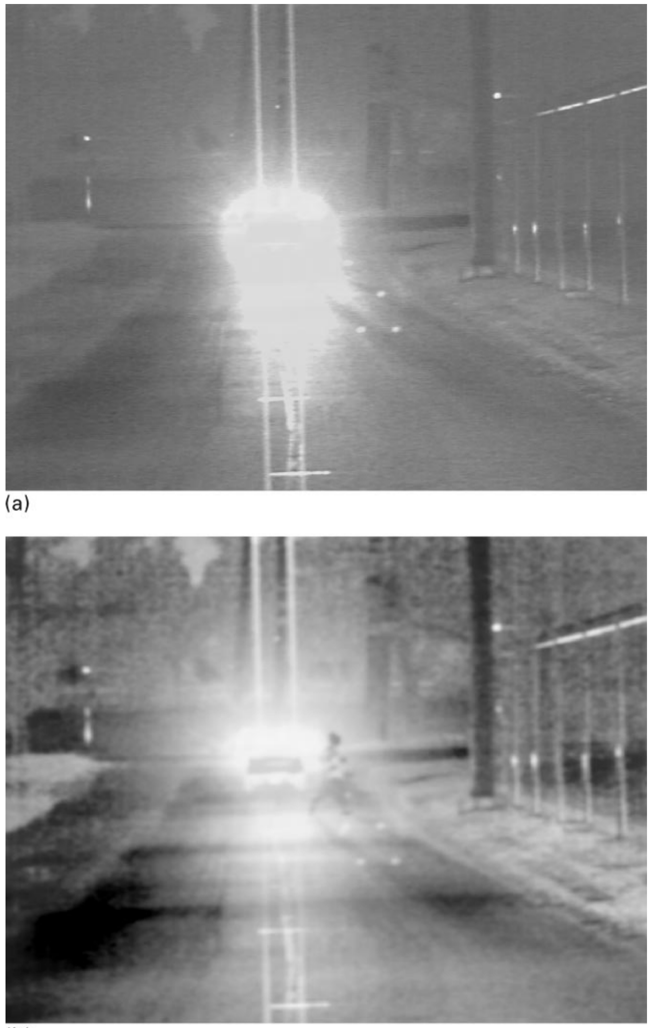

(a)

(b)

Fig. 2. Composite images derived through simple fusion (top) and principal components fusion (bottom) of the thermal and visible images presented in Fig. 1.

1.5. Results

Mean RTs and error rates were calculated for all combinations of image format and level of illumination. Overall accuracy for many conditions, however, was too low to allow for meaningful interpretation of RTs. Accuracy data were therefore recast as hit rates and false alarm rates, and subjected to signal detection analysis[1] (Green and Swets, 1966).

Because receiver operating characteristic curves were not available to allow a check of the statistical assumptions necessary for calculating d' (Gescheider, 1997), the traditional measure of sensitivity, sensitivity was instead quantified with the measure A' (Pollack and Norman, 1964). Mean A' scores are presented in Fig. 3 as a function of image format for all three levels of illumination.

[1] We thank Dr. Colin Drury and an anonymous reviewer for recommending this analysis.

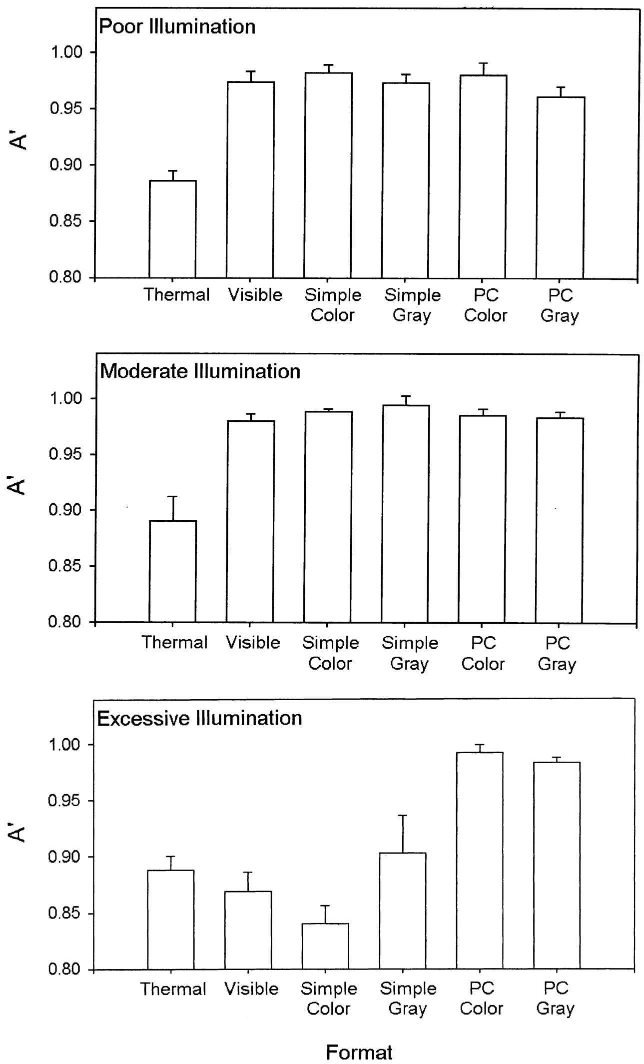

Fig. 3. Sensitivity for target detection as a function of image format at three levels of illumination.

False alarm rates were generally low, ranging between 0.00 and 0.08. In contrast, hit rates varied between 0.55 and 0.99. Thus, changes in sensitivity across various image formats and levels of illumination were manifest largely as increases in the frequency of missed targets, rather than the frequency of false alarms. In general, sensitivity with thermal imagery was relatively low across all levels of illumination.[2] Sensitivity with visible and simple fused imagery, conversely, was relatively high under low and moderate illumination, but was degraded under the highest level of illumination. Sensitivity with principal components fused imagery was relatively high across all levels of illumination.

For statistical analysis, sensitivity scores were first submitted to an omnibus 6×3 repeated measures ANOVA with image format (thermal, visible, achromatic simple fused, chromatic simple fused, achromatic principal components fused, chromatic principal components fused) and level of illumination (low, moderate, excessive) as factors. Results of this analysis, after Geisser–Greenhouse (1958) correction, were a reliable main effect of format, $F(1, 8) = 45.86, p < 0.01$, a reliable main effect of illumination, $F(1, 8) = 37.78, p < 0.01$, and a reliable interaction of format by glare, $F(1, 8) = 13.21, p = 0.01$. Subsequent simple effects tests of image format were reliable at all three levels of illumination, all p's < 0.01. Conversely, simple effects tests of illumination were significant for visible and simple fused image formats, p's < 0.01, but not for thermal imagery or for either form of principal components fused imagery, p's > 0.125. Variations in illumination, that is, appeared to reliably affect performance with visible and simple fused images, but not with thermal or principal components fused images.

Further analysis was conducted to examine the effects of image format within each level of illumination, and in particular to address two questions. First, did sensor fusion either degrade or improve image quality? Second, was color rendering of sensor-fused imagery either more or less useful than grayscale rendering? Analysis entailed a set of five orthogonal planned comparisons conducted at each level of glare. Individual planned comparisons were designed to (1) compare performance with component thermal imagery to performance with component visible imagery, (2) compare performance with the better form of component imagery (thermal or visible) to performance with achromatic simple fused imagery, (3) compare performance with the better form of component imagery to performance with achromatic principal components fused imagery, (4) compare performance with chromatic simple fused imagery to performance with achromatic simple fused imagery, and (5) compare performance with chromatic principal components fused imagery to performance with achromatic principal components fused imagery. If sensor fusion improved spatial image quality relative to that of component images, these comparisons should have revealed higher sensitivity with grayscale fused images than with the better form of component imagery. If chromatic rendering of fused imagery improved image quality, these comparisons should have revealed higher sensitivity with chromatic fused images than with grayscale images created through the same fusion algorithm.

Results confirmed that sensitivity was reliably higher with component visible imagery than with thermal

[2] A partial replication of the current experiment, employing a subset of the present stimuli but presenting them on a non-linearized monitor of higher resolution than that used here, found performance with thermal imagery consistently better than that reported here. Additional results of this follow-up study, however, were wholly consistent with the present data. Given that performance with thermal imagery was relatively consistent between-subjects and within-subjects across levels of illumination in both experiments, and that data from both studies were otherwise in full agreement, this discrepancy suggests that the quality of thermal imagery may vary more with changes in display hardware than does the quality of alternative image formats, but does not otherwise undermine conclusions drawn from the present data.

imagery under low illumination, $F(1, 8) = 87.32$, $p < 0.01$, and under moderate illumination, $F(1, 8) = 16.28$, $p < 0.01$, but that performance under excessive illumination failed to differ reliably between thermal and visible formats, $F < 1$. Subsequent tests to assess the benefits of sensor fusion therefore compared performance with fused imagery to that with component visible imagery. Under low illumination, mean performance with grayscale simple fused imagery failed to differ reliably from that with visible imagery alone, $F < 1$, suggesting that simple fusion did little to either enhance or degrade the quality of information conveyed by visible component images. Performance with grayscale principal components imagery under low illumination, however, was reliably *worse* than performance with visible imagery, $F(1, 8) = 7.20$, $p = 0.028$, indicating that principal components fusion under these conditions actually degraded the spatial content of visible input images. Thus, under poor illumination, simple fusion produced imagery of better spatial quality than did principal components fusion. Results under moderate illumination were somewhat analogous, suggesting that here simple fusion reliably improved spatial image quality while principal components fusion did not; whereas performance with achromatic simple fused images exceeded that with component visible imagery, $F(1, 8) = 7.20$, $p = 0.03$, performance with achromatic principal components images did not differ reliably from that with visible images, $F < 1$. Under excessive illumination, however, where the quality of component visible imagery was degraded by glare, results differed dramatically. Here, performance with grayscale simple fused images was again better than that visible images, but only with marginal reliability, $F(1, 8) = 4.23$, $p = 0.07$. Both of these image formats, notably, produced performance reliably worse than that which obtained under moderate illumination, $t(8) = 5.58$, $p < 0.01$ for simple fused imagery, $t(8) = 6.32$, $p < 0.01$ for visible imagery. In contrast, performance with grayscale principal components imagery under excessive illumination was not worse than that which obtained under moderate illumination, $t(8) = 0.03$, $p = 0.97$, and was reliably better than performance with component visible imagery, $F(1, 8) = 60.71$, $p < 0.01$.

Tests of the effects of color rendering indicated that the benefits of chromatic information derived through image fusion also varied with the level of illumination. Clear benefits of chromatic information were evident only under conditions of low illumination, where color rendering modestly but reliably improved performance both with simple fused imagery, $F(1, 8) = 5.27$, $p = 0.05$, and with principal components fused imagery, $F(1, 8) = 9.67$, $p = 0.01$. Under moderate illumination, color rendering had little effect on performance with either simple fused imagery, $F(1, 8) = 1.37$, $p = 0.28$, or principal components fused imagery, $F < 1$. Performance with principal components imagery under excessive illumination was likewise not reliably improved by chromatic rendering, $F(1, 8) = 3.01$, $p = 0.12$ (though a ceiling effect here might have obscured some benefits of color rendering). More remarkably, however, performance with simple fused imagery under excessive glare was degraded by color rendering, an effect which fell just short of the conventional level of statistical reliability, $F(1, 8) = 4.98$, $p = 0.06$.

Analysis of RTs for correct responses was conducted to ensure that the reliable effects of image format and illumination on sensitivity data were not the result of a speed–accuracy trade-off. Tables 2 and 3 present mean RTs and standard errors for target-present and target-absent responses, respectively. At each level of illumination, target-present and target-absent RTs were examined separately by comparisons identical to those with which sensitivities were analyzed. These analyses produced no results to contradict conclusions drawn from analysis of sensitivity. Target-present RTs were reliably faster to visible than to thermal imagery under low illumination, $F(1, 8) = 6.96$, $p = 0.03$, and under moderate illumination, $F(1, 8) = 15.39$, $p < 0.01$. Target-present RTs to grayscale principal components imagery were reliably faster than those to visible imagery under conditions of excessive illumination, $F(1, 7) = 7.90$, $p = 0.02$. Similar analysis of target-absent RTs at all three levels of illumination revealed only one reliable effect, a tendency for shorter RTs with grayscale

Table 2
Mean RTs and standard errors for target present responses

Image format	Low illumination	Moderate illumination	Excessive illumination
Thermal	653 ± 52	668 ± 50	677 ± 62
Visible	545 ± 32	530 ± 32	711 ± 96
Simple fused, grayscale	513 ± 26	520 ± 28	644 ± 50
Simple fused, color	518 ± 25	529 ± 36	659 ± 135
Principal components fused, grayscale	530 ± 25	556 ± 28	473 ± 24
Principal components fused, grayscale	525 ± 28	515 ± 29	491 ± 27

Table 3
Mean RTs and standard errors for target absent responses

Image Format	Low illumination	Moderate illumination	Excessive illumination
Thermal	1036 ± 213	958 ± 191	906 ± 149
Visible	977 ± 180	1007 ± 168	1031 ± 188
Simple fused, grayscale	953 ± 174	1007 ± 181	1035 ± 177
Simple fused, color	1019 ± 194	1064 ± 200	1022 ± 192
Principal components fused, grayscale	1003 ± 202	1056 ± 198	1034 ± 222
Principal components fused, grayscale	1043 ± 224	1020 ± 160	1150 ± 242

principal components imagery than with chromatic principal components imagery under the highest level of illumination, $F(1, 8) = 5.39, p = 0.03$. Given that no similar trend was evident in target-present RTs, however, and that sensitivities were slightly higher with chromatic than with achromatic principal components imagery under excessive illumination, this effect is unlikely to reflect true differences in quality between color and grayscale principal components imagery.

2. Discussion

Sensor fusion, as described earlier, offers to improve image quality two ways, first by combining spatial information from multiple sources within a single image, and second by deriving emergent spatial or chromatic information from contrast between input images. A sensor-fused image should therefore at minimum allow visual performance equal to that obtained with the best of its multiple component images, and will optimally allow performance superior to that obtained with any single component image. The results of the present experiment, measuring observers' ability to detect a pedestrian within component and sensor-fused composite nighttime images, demonstrate that fusion can in fact maintain, combine, and possibly enhance spatial information provided as input; performance with grayscale simple fused imagery was as good as that with component imagery under low and excessive illumination and was better under moderate illumination, while performance with grayscale and chromatic principal components fused imagery was equivalent to that with visible imagery under moderate illumination and dramatically superior under glare. Furthermore, the present results indicate that thermal/visible fusion can effectively derive chromatic information to improve image quality; chromatic rendering of fused images modestly but reliably improved performance under poor illumination. In concordance with a number of previous reports (e.g., Steele and Perconti, 1997), however, the present data also demonstrate that sensor fusion will not invariably improve or even maintain image quality. Under low illumination, performance with grayscale principal components imagery was worse than that with unfused visible images, indicating that spatial content of the visible input imagery was degraded by fusion. Unexpectedly, data also suggest that color rendering of fused imagery, even without changing the spatial content of an image, can degrade the *perceptibility* of spatial image content. Here, performance with chromatic simple fused imagery was worse under excessive glare than performance with achromatic simple fused imagery, despite the fact that achromatic and chromatic renderings conveyed identical spatial information.

In total, results demonstrate that fusion can dramatically improve image quality, but that the benefits of sensor fusion may vary with environmental conditions and/or the quality of input imagery tested, and with the form of fused imagery under consideration. This conclusion reconciles the apparently conflicting results of earlier research, and qualifies even those studies which appear to have clearly endorsed or dismissed the value of image fusion in general or any fusion algorithm in particular. More practically, this conclusion recommends sensor fusion as a prospective method of enhancing imaging systems deployed as night-vision aids, but urges caution in the design and testing of any fusion system meant to function under a broad range of environmental conditions. Considerations to be addressed will include choice of fusion algorithm, and choice of component wavebands. The present experiment employed only imagery created through the simple fusion and principal components fusion algorithms of Scribner and colleagues (Scribner et al., 1993, 1996). A number of alternative algorithms for creating chromatic and achromatic fused imagery, however, have been described by various researchers (e.g., Toet and Walraven, 1996; Therrien et al., 1997; Waxman et al., 1997). Likewise, the present work employed only composite imagery derived from fusion of thermal and visible component images, while fusion of information from any two or more wavebands is possible (e.g., Krebs et al., 1999). Extensive psychophysical research will clearly be necessary to optimize a system of sensor fusion for use as a night vision aid.

Acknowledgements

This work was sponsored by DARPA's Integrated Imaging Sensors Program, and Mr. Ray Balcerak was the Program Manager. This work was also, sponsored by Office of Naval Research contracts #N0001497WR30091 and #N0001498WR30112. Thanks to Drs. Jonathan Schuler and Dean Scribner at the Naval Research Laboratory for providing the fused images used as stimuli here, and to Dr. Colin Drury and two anonymous reviewers for helpful comments on an earlier version of this article. The views expressed in the article are those of the authors, and do not reflect the official position of the Department of the Navy, the Department of Defense, or the United States Government.

References

Andre, J.T., 1996. Visual functioning in challenging environments: effects of alcohol consumption, luminance, stimulus motion, and glare on contrast sensitivity. J. Exp. Psycho. Appl. 2, 250–269.

Essock, E.A., Sinai, M.J., McCarley, J.S., Krebs, W.K., DeFord, J.K., 1999. Perceptual ability with real-world nighttime scenes: image-intensified, infrared, and fused-color imagery. Human Factors, 41, 438–452.

Geisser, S., Greenhouse, S.W., 1958. An extension of Box's results on the use of the F distribution in multivariate analysis. Ann. Math. Statist. 29, 885–891.

Geschieder, G.A., 1997. Psychophysics: The Fundamentals, 3rd edition. Erlbaum, Mahwah, New Jersey.

Green, D.M., Swets, J.A., 1966. Signal Detection Theory and Psychophysics. Wiley, New York.

Krebs, W.K., McCarley, J.S., Kozek, T., Miller, G.M., Sinai, M.J., Werblin, F.S., 1999. An evaluation of a sensor fusion system to improve drivers' nighttime detection of road hazards. Proceedings of the Human Factors Society 43rd Annual Meeting.

Owens, D.A., Sivak, M., 1983. Comparison of headlamp visibility and distance and stopping distance. Perceptual Motor Skills 57, 1177–1178.

Owens, D.A., Sivak, M., 1996. Differentiation of visibility and alcohol as contributors to twilight road fatalities. Hum. Factors 38, 680–689.

Piccione, D., Ferrett, D., 1998. Driving Miss Bradley: performance measurement to support thermal driving. In: Verly, J.G. (Ed.), Proceedings of the SPIE: Vol 3364. Enhanced and Synthetic Vision, SPIE, The International Society for Optical Engineering, Bellingham, WA, pp. 200–207.

Pollack, I., Norman, D.A., 1964. A non-parametric analysis of recognition experiments. Psychonomic Sci. 1, 125–126.

Ryan, D., Tinkler, R., 1995. Night pilotage assessment of image fusion. In: Lewandowski, R.J., Stephens, W., Haworth, L.A. (Eds.), Proceedings of the SPIE: Vol. 2465. Helmet- and Head-Mounted Displays and Symbology Design Requirements II. SPIE, The International Society for Optical Engineering, Bellingham, WA, pp. 50–67.

Scribner, D.A., Satyshur, M.P., Kruer, M.R., 1993. Composite infrared color images and related processing. Proceedings of the IRIS Specialty Group on Targets, Backgrounds, and Discrimination, January, San Antonio, TX.

Scribner, D.A., Satyshur, M.P., Schuler, J., Kruer, M.R., 1996. Infrared color vision. Proceedings of the IRIS Specialty Group on Targets, Backgrounds, and Discrimination, January, Monterey, CA.

Steele, P.M., Perconti, P., 1997. Part task investigation of multispectral image fusion using gray scale and synthetic color night vision sensor imagery for helicopter pilotage. In: Watkins, W.R., Clement, D. (Eds.), Proceedings of the SPIE: Vol. 3062. Targets and Backgrounds: Characterization and Representation III. SPIE, The International Society for Optical Engineering, Bellingham, WA, pp. 88–100.

Sturgis, S.P., Osgood, D.J., 1982. Effects of glare and background luminance on visual acuity and contrast sensitivity: implications for driver night vision testing. Hum. Factors 24, 347–360.

Sturr, J.F., Kline, G.E., Taub, H.A., 1990. Performance of younger and older drivers on a static acuity test under photopic and mesopic luminance conditions. Hum. Factors 32, 1–8.

Swift, D.J., Panish, S., Hippensteele, B., 1997. The use of Vision Works in visual psychophysics research. Spatial Vision 10, 471–477.

Therrien, C.W., Scrofani, J., Krebs, W.K., 1997. An adaptive technique for the enhanced fusion of low-light visible with uncooled thermal infrared imagery. In: Gray, R.M., Hunt, B. (Eds.), Proceedings of the IEEE International Conference on Image Processing. The Institute of Electrical and Electronics Engineers, Inc., New York, NY, pp. 405–408.

Toet, A., Ijspeert, J.K., Waxman, A.M., Aguilar, M., 1997. Fusion of visible and thermal imagery improves situational awareness. In: Verly, J.G. (Ed.), Proceedings of the SPIE: Vol. 3088. Enhanced and Synthetic Vision. SPIE, The International Society for Optical Engineering, Bellingham, WA, pp. 177–188.

Toet, A., Walraven, J., 1996. New false color mapping for image fusion. Opt. Eng. 35, 650–658.

Waxman, A.M., Gove, A.N., Fay, D.A., Racamato, J.P., Carrick, J.E., Seibert, M.C., Savoye, E.D., 1997. Color night vision: Opponent processing in the fusion of visible and IR imagery. Neural Networks 10, 1–6.