Machine Vision and Applications (2000) 11: 313–320

Machine Vision and Applications

Automatic detection of vehicle occupants: the imaging problem and its solution

I. Pavlidis[1], P.Symosek[1], B.Fritz[1], M.Bazakos[1], N. Papanikolopoulos[2]

[1] Honeywell Technology Center, 3660 Technology Dr., MN65-2500, Minneapolis, MN 55418, USA;
e-mail: {pavlidis, symosek, fritz, bazakos}@htc.honeywell.com
[2] University of Minnesota, Computer Science Department, 200 Union St. S.E., Minneapolis, MN 55455, USA;
e-mail:npapas@cs.umn.edu

Abstract. The automatic detection and counting of vehicle occupants is a challenging research problem that was given little attention until recently. An automated vehicle-occupant-counting system would greatly facilitate the operation of freeway lanes reserved for car pools (high occupancy vehicle lanes or HOV lanes). There are three major aspects of this problem: (a) the imaging aspect (sensor phenomenology), (b) the pattern recognition aspect, and (c) the system architecture aspect. In this paper, we present a solution to the imaging aspect of the problem. We propose a novel system based on fusion of near-infrared imaging signals and we demonstrate its adequacy with theoretical and experimental arguments. We also compare our solution to other possible solutions across the electromagnetic spectrum, particularly in the thermal infrared and visible regions.

Key words: Near-infrared – Sensor fusion – Sensor phenomenology – Detection of vehicle

1 Introduction

There are compelling reasons for the existence of an automatic vehicle-occupant-counting system in the high-occupancy vehicle (HOV) lane. In particular, such a system will be useful in the following respects.

1. It will facilitate the gathering of statistical data for road construction planning. The gathering of usage statistics in the HOV lane is mandated by the U.S. Federal Highway Administration. Currently, the gathering of data is performed manually. This is obviously laborious, inefficient, and prone to error.
2. It will facilitate law enforcement in the HOV lane. Currently, HOV lane enforcement requires substantial commitments of State Highway Patrol personnel and equipment. HOV lane enforcement has other costs as well. These include the risks of high-speed pursuit in lanes adjacent to stop-and-go traffic and the deterioration of traffic flow when tickets are issued during peak commuting periods.
3. It will enable departments of transportation to offer the option to single drivers to use some HOV lanes for a nominal monthly fee.

A complete HOV monitoring system suitable for the above applications will consist of a vehicle occupant detector and a license plate reader. Although, substantial work has been reported in the technical literature regarding license plate readers, work on automated vehicle occupant detectors is still in its infancy. There are three major technical challenges in the development of an automatic vehicle occupant detector.

1. The imaging signal should provide a clear picture of the interior of the vehicle. The contrast between the human silhouettes and the background should be sufficient to provide for reliable image processing.
2. The pattern recognition algorithm that performs the vehicle occupant detection should exhibit high recognition rates and robust behavior. Of course, its performance depends to a significant degree on the quality of the imaging signal. Even the best pattern recognition algorithm cannot perform reliably when the imaging signal is corrupted with noise.
3. The system architecture should be designed in such a way that will ensure real-time operation, vehicle occupant detection in both the front and back seats, and protection from the weather elements.

In this paper, we address the first of the three technical challenges. We describe a novel near-infrared fusion system that provides a high-quality imaging signal during daytime and nighttime operation and in adverse weather conditions. In particular, in Sect. 2 we give an overview and justification of our approach and compare its features to those of other approaches. In Sect. 3, we describe in detail the theoretical computations that support our assertions. In Sect. 4, we present the experimental validation of our hypotheses. Finally, in Sect. 5, we conclude the paper and briefly mention our ongoing and future work.

Correspondence to: I. Pavlidis

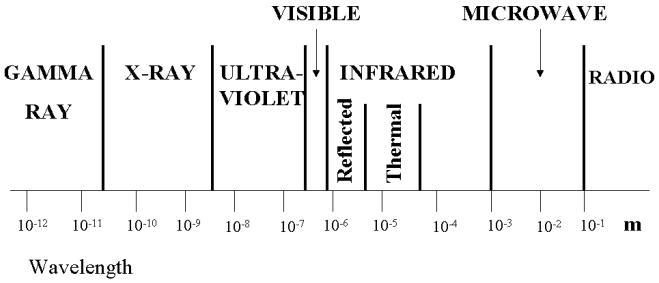

Fig. 1. Electromagnetic (EM) spectrum

2 Overview

We embarked on research for a solution to the imaging aspect of an HOV system (sensor phenomenology) to resolve the following questions.

1. Is there a band in the electromagnetic(EM) spectrum that can penetrate through the vehicle's window glass, during day and night and in adverse weather conditions? Do the objects of interest (vehicle occupants) have a consistent appearance in this EM band, irrespective of their physical characteristics?
2. If there is more than one band, can we fuse the multiple bands in a meaningful way to increase the detecting power and reliability of the system?
3. Are there appropriate cameras for these bands that have the necessary resolution and speed?

Figure 1 shows the EM spectrum. We have limited our sensor phenomenology investigation to the infrared and visible spectrum wavebands. Nature constrains our choices from the regions below the visible spectrum, since, gamma rays, X-rays and, ultraviolet radiation are harmful to the human body. Therefore, the typical active systems in these ranges cannot be employed in the HOV lane. Technology constrains our choices beyond the infrared region, since millimeter wave and radio wave imaging sensors are very expensive, bulky, and achieve insufficient resolution [7]. Still, the visible plus the infrared range is a huge area of the EM spectrum and we had to identify narrow bands within this area that are appropriate for the task.

2.1 The visible and thermal infrared spectrum solutions

We know from experience as humans that the visible spectrum has certain disadvantages for the purpose of the particular application. A visible-range sensor (like the human eye) cannot easily see at night unless it is aided by an artificial illumination source. Employing a visible-range flashlight to illuminate the passing vehicles is definitely not an option, since it will distract the drivers with probably fatal results. Figure 2 shows three images of vehicles taken at night with three different imaging sensors. The first image was taken with a passive thermal-infrared camera system. The second image was taken with an active near-infrared camera system. The quality in the first two images is high and one can readily detect the human occupants in the vehicles. The third image was taken with a passive visible-spectrum camera system. Although no occupant appears to be in sight,

a

b

c

Fig. 2a–c. a Thermal image. **b** Near-infrared image. **c** Visible spectrum image. All the images were captured during nighttime

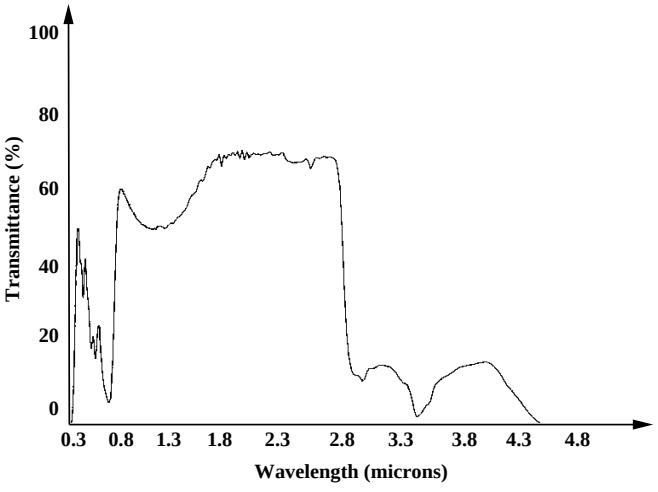

Fig. 3. Transmittance of a typical tinted side window

there were actually two occupants seated in the interior of the vehicle. Due to darkness, the occupants cannot be seen in the image. Tinted window glass (now common in certain vehicle types) also prohibits a clear view of the vehicle's interior to visible-range sensors (see Fig. 3). Also, visible-range sensors are incapacitated during foul weather conditions. Finally, vehicle occupants produce variable patterns in the visible range, depending on their physical characteristics, time of day, and the illumination conditions. This variability makes the machine vision task much more difficult.

From the above discussion, it is apparent that only the infrared range holds promise for a solution to the problem.

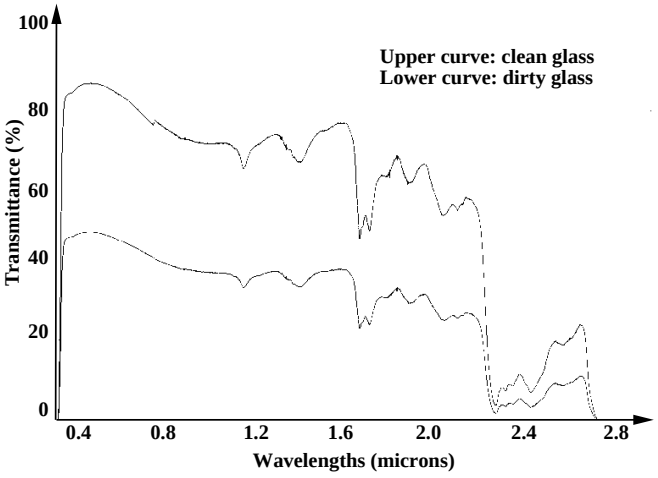

Fig. 4. Transmittance of a typical windshield

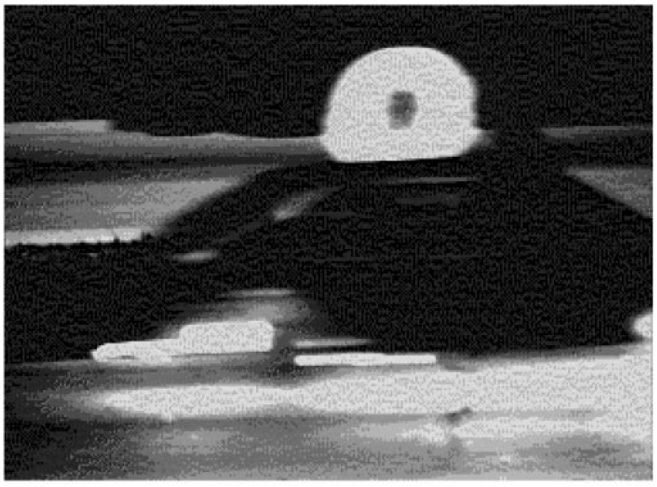

Fig. 5. Side snapshot of a fast-moving vehicle (65 mph) with a mid-infrared camera. The image appears heavily blurred, since the speed of the camera cannot keep up with the speed of the target

Within the infrared range the wavebands of particular interest are the reflected infrared (0.7–3.0 μm) and the thermal infrared (3.0–5.0 μm, 8.0–14.0 μm) wavebands. The reflected infrared band on one hand is associated with reflected solar radiation that contains no information about the thermal properties of materials. This radiation is for the most part invisible to the human eye. The thermal infrared band on the other hand is associated with the thermal properties of materials. We soon found that the thermal infrared band was difficult to exploit for HOV purposes because vehicle glass severely attenuates EM radiation beyond 2.4 μm (see Figs. 3 and 4). For all practical purposes, using an appropriate mid-infrared camera like the Mitsubishi $M700$, one can only "see through" the side window glass. This only exacerbates the problem caused by the low speed specification of typical mid-infrared cameras. For an imaging configuration where the imaging system views the side of the vehicle, the vehicle remains in the field of view of the camera for less time: for vehicles moving at freeway speeds (about 65 mph), the blurring effect on the image is overwhelming (see Fig. 5). An additional annoyance with the use of mid-infrared arises when the vehicle's defogger has been on for more than half an hour. Then, the thermal signature of the interior of the vehicle is completely dominated by the defogger heat (see Fig. 6a–c).

Contrary to the difficulties we encountered in the visible and mid-infrared spectra, a major portion of the reflected-infrared spectrum, the so-called near-infrared spectrum (0.7–2.4 μm), appeared very suitable for the application at hand. In particular, we found that

1. A camera in this range can safely operate in the HOV lane both day and night. During nighttime, we would need a matching near-infrared illumination source to enhance the scene. Provided that the spectral signature of the illumination source is deep into the near-infrared range, the light will be invisible to the human eye. Therefore, no danger of driver distraction exists.
2. A camera in this range can "see through" both the vehicle's windshield and its side windows. The transmittance of typical vehicle windows in the near-infrared spectrum is at least 40% (see Figs. 3 and 4).

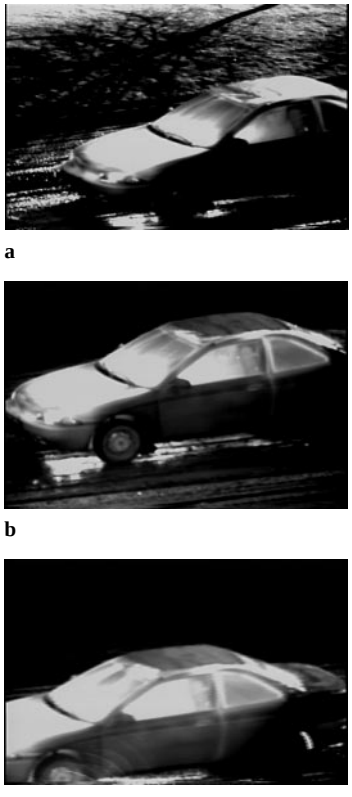

a

b

c

Fig. 6a–c. Mid-infrared snapshots of the interior of a vehicle when the defogger is on for various periods of time. **a** The defogger is on for 10 min. **b** The defogger is on for 20 min. **c** The defogger is on for 30 min

3. A camera in this range can operate in adverse weather conditions. For example, it has been established that near-infrared spectrum imaging systems are particularly good in penetrating haze.
4. If the near-infrared range is split into two bands around the threshold point of 1.4 μm, the *lower-band* (0.7–1.4 μm) and the *upper-band* (1.4–2.4 μm), then vehicle occupants will produce consistent signatures in the re-

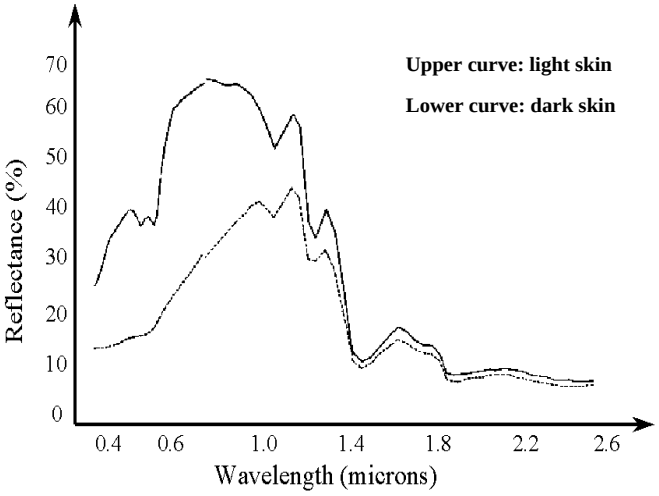

Fig. 7. Reflectance of dark skin versus light skin

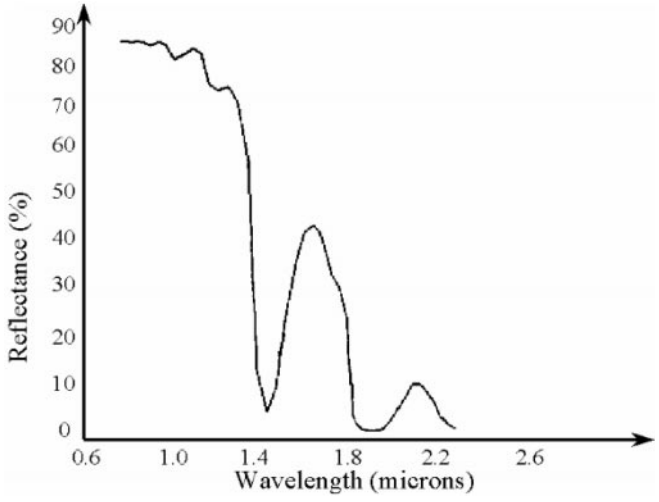

Fig. 8. Reflectance of distilled water

spective imagery. In the upper band imagery, humans will appear consistently dark, irrespective of their physical characteristics and the illumination conditions. In the lower band imagery, humans will appear comparatively lighter. This is because human skin appears to have very high reflectance just before 1.4 μm wavelength, but very low reflectance just after 1.4 μm (see Fig. 7) [1, 5].

We found that the intriguing phenomenon of the abrupt change in the reflectance of human skin around 1.4 μm is due to the water content of the human body. Water absorbs near-infrared radiation heavily above 1.4 μm, and thus has low reflectance in this region. (see Fig. 8). Humans consist of 70% water, and therefore they exhibit spectral behavior very similar to water. Interestingly, other inanimate objects in the vehicle scene maintain nearly the same reflectance levels, below and above the threshold point 1.4 μm [2, 4, 6]. For example, see Fig. 9 for the reflectance diagrams of some fabric materials commonly found in the interior of vehicles. This observation provoked the following line of thought. Ideally, everything but the human skin signature should appear proportionally the same in the HOV imagery from the two bands. Therefore, by subtracting an image from

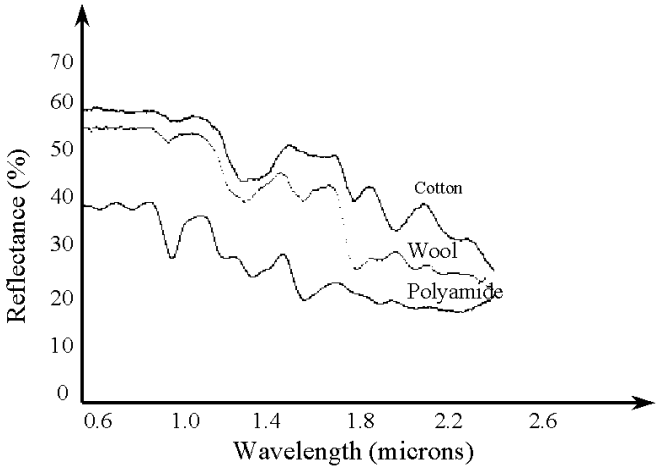

Fig. 9. Reflectance of different fabric materials

the lower band from its matching, co-registered image in the upper band, we can produce a fused image, where

1. the silhouettes of the vehicle occupants' faces will be reinforced (big difference) and more clearly stand out.
2. The background will dim away (small difference).

This increased contrast will facilitate a clean-cut thresholding of the fused image. The thresholded result will be an image where only the face blobs of the vehicle occupants remain and everything else is eliminated. A good classifier will always classify such a simple binary pattern accurately, ensuring the reliable real-time operation of the HOV system.

The near-infrared camera we found most appropriate to use for testing our ideas was the Sensors Unlimited $SU-320$ Indium Gallium Arsenide camera. In terms of spectral response, it was less than perfect because it did not cover the entire usable spectral range (0.7–2.4 μm). Instead, it covered the subrange 1.1–1.4 μm for the lower band and 1.4–1.7 μm for the upper band. The question we had to address was, given our hypotheses and the particular camera model available, could the signal-to-noise (S/N) ratio and the speed of the camera ensure good quality images for vehicle occupant detection. The complete set of calculations regarding this issue and the interpretation of their results are described in the next section.

3 Theoretical computations

Our hypotheses, as described in the previous section, held great promise. Nevertheless, before we could proceed with any experimentation, we had to determine if, given the particular $SU-320$ camera specifications,

1. we would have an imaging signal with sufficient S/N ratio.
2. The speed of the camera would be sufficient to capture the vehicle occupants moving at an average speed of 65 mph (freeway speed).

As we stated earlier, we consider two spectral bands, one above the 1.4-μm threshold point and one below it. We assume that two time-synchronized $SU-320$ cameras

would acquire images of the same scene. One camera should be equipped with an upper band filter and one with a lower band filter. Both cameras should be equipped with a polarizer during daytime to reduce solar glare. They should also be equipped with a tele-photo lens. Because of the $SU-320$ camera's technical characteristics, we limit the upper-band to the range 1.4–1.7 μm and the lower band to the range 1.1–1.4 μm. We will demonstrate our S/N computation for the lower band only, since a similar argument applies to the upper band. The computation of the camera speed will take place as part of the S/N computation.

The first step in a S/N radiometric computation is to determine the amount of irradiation that falls upon the objects of interest [3] – in our case the vehicle occupants. The spectral irradiance of the sun (our illumination source) on a clear day at sea level is approximately I_{sunny} = 0.008 Watts/cm^2 in the 1.1–1.4-μm waveband. In our computation, however, we consider the worst case scenario of an overcast day. For an overcast day, the irradiance value is reduced by 10^{-3}, giving an irradiance at the vehicle of approximately

$$\begin{aligned} I_{overcast} &= 10^{-3} \times I_{sunny} \\ &= 10^{-3} \times 0.008 \\ &= 8\,\mu\text{Watts/cm}^2. \end{aligned} \tag{1}$$

We assume that the camera is pointed at the vehicle's windshield, not at a side window. The transmittance of the windshield of a common vehicle in the spectral band of interest is approximately 0.4. We assume the worst case scenario of a dirty window (see Fig. 4). This results in an irradiance on the vehicle occupants of

$$\begin{aligned} I_{occupant} &= 0.4 \times I_{overcast} \\ &= 0.4 \times 8 \\ &= 3.2\,\mu\text{Watts/cm}^2. \end{aligned} \tag{2}$$

The second step in a radiometric computation is to determine how much of the incident irradiation on the objects of interest is reflected back to the sensor (the $SU-320$ near-infrared camera in our case). The radiance into a hemisphere, assuming the worst case skin reradiate rate of 0.4 (see Fig. 7), would be

$$\begin{aligned} R_{occupant} &= 0.4 \times I_{occupant}/\pi \\ &= 0.4 \times 3.2/\pi \\ &= 0.4\,\mu\text{Watts/cm}^2 - ster. \end{aligned} \tag{3}$$

This represents the reflected portion of the vehicle occupant irradiation. The rest is absorbed by the occupant's body. The reflected radiation has to pass through the windshield, the camera lens, the bandpass filter, and the polarizer to reach the near-infrared sensor array. As we did earlier, we assume a 0.4 windshield transmittance in the spectral band of interest. We also assume a $f/2$ camera lens (14.32^o cone angle) with 0.8 transmittance, a polarizer with 0.4 transmittance, and a bandpass filter with 0.6 transmittance. Then, the irradiance at the sensor array of the $SU-320$ camera will be

$$\begin{aligned} I_{camera} &= 0.4 \times 0.8 \times 0.4 \times 0.6 \times \pi \times \\ &\quad R_{occupant} \times \sin^2(14.32^o) \\ &= 0.4 \times 0.8 \times 0.4 \times 0.6 \times \pi \times \\ &\quad 0.4 \times \sin^2(14.32^o) \end{aligned}$$

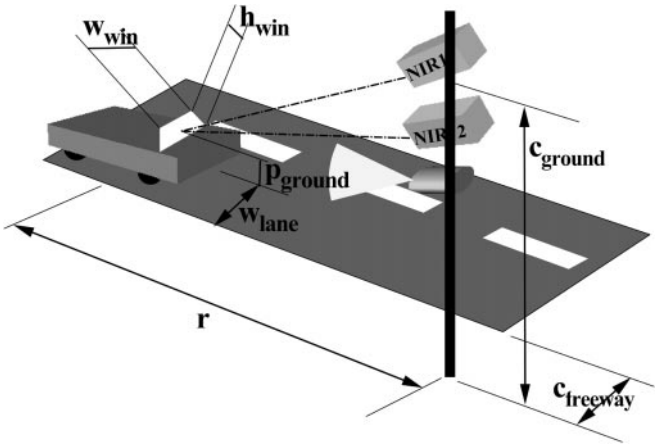

Fig. 10. Configuration of the camera set

$$= 0.006\,\mu\text{Watts/cm}^2. \tag{4}$$

The SU-320 camera has square pixels with a side dimension of 37.5×10^{-4} cm or an area

$$\begin{aligned} A &= 37.5 \times 10^{-4} \times 37.5 \times 10^{-4} \\ &= 1.40 \times 10^{-5}\,\text{cm}^2. \end{aligned} \tag{5}$$

Consequently, the radiant power on the camera pixel will be

$$\begin{aligned} P_{pixel} &= A \times I_{camera} \\ &= 1.4 \times 10^{-5} \times 0.006 \\ &= 0.084 \times 10^{-12}\,\text{Watts}. \end{aligned} \tag{6}$$

The camera's detectivity, D^*, is $D^* = 10^{12}\,\text{cm}-\sqrt{Hz}/\text{Watts}$. The noise-equivalent power NEP is related to detectivity D^*, pixel area A, and electronic bandwidth Δf by the following equation:

$$NEP = (A \times \Delta f)^{1/2}/D^*. \tag{7}$$

We already know the value of D^* and A. In order to compute the NEP, we need to also know the value of Δf. The bandwidth Δf is determined by the exposure time (speed) of the camera. In turn, the exposure time depends on the vehicle speed (v), the camera's instantaneous field of view ($IFOV$), the range (r), and the footprint of the horizontal translation (f_{ht}). Based on the parameters v, $IFOV$, r, and f_{ht}, we compute the required exposure time (speed) of the camera such that the image smear is less than 1 pixel. Then, we check if the exposure time value falls within the operational range of the $SU-320$ camera. If it does, the $SU-320$ camera is adequate for the HOV task in terms of speed. We can substitute the corresponding value for the bandwidth Δf in Eq. 7 and continue the process of computing the S/N ratio.

Figure 10 shows the configuration of the $SU-320$ camera relative to the oncoming traffic. The camera is located c_{ground} = 3.6 m above the ground, $c_{freeway}$ = 7.5 m off the edge of the freeway, and at a distance of r = 40 m from the oncoming traffic. This arrangement ensures that the camera is located in a safe place and has the appropriate field of view. We assume that the camera focuses at the centerline of the incoming vehicle, at the level of the occupants' faces (p_{ground} = 1.2 m). The half-width of a standard freeway lane

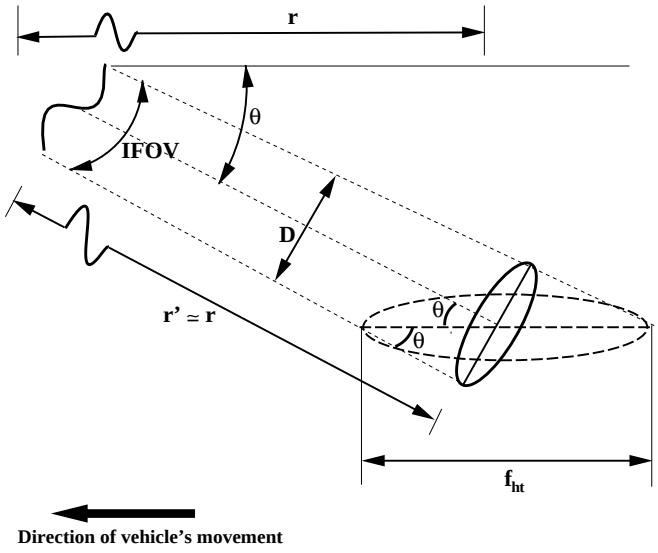

Fig. 11. Geometry for the computation of the footprint of a single pixel in a horizontal plane

is w_{lane} = 1.8 m. We assume that the vehicle travels in the middle of the freeway lane. Therefore, the lateral distance of the vehicle's centerline from the camera is

$$c_{vehicle} = c_{freeway} + w_{lane} = 7.5 + 1.8 = 8.8\,m. \quad (8)$$

Finally, we assume that, for a typical vehicle's windshield, the average width and height are w_{win} = 1.5 m and h_{win} = 0.9 m, respectively.

The IFOV is the camera's field of view with respect to a single pixel (see Fig. 11). We assume that the distance $r^{'}$ is approximately equal to the camera's range r ($r^{'} \approx r$ = 40 m). Then, the $IFOV$ can be computed from the following equation:

$$\begin{aligned} IFOV &= \frac{\arctan\left[(h_{win}/2)/r^{'}\right]}{(h_{FPA}/2)} \\ &= \frac{\arctan\left[(0.9/2)/40\right]}{(240/2)} \\ &= 0.0001\,rad, \end{aligned} \quad (9)$$

where $h_{FPA} = 240\,pixels$ is the vertical dimension of the $SU-320$ focal plane array (FPA).

At time t, the camera's IFOV sees a small portion of the occupant's face of diameter D. This small face area is what is imaged into a single pixel. We can determine the value of D from the following equation:

$$\begin{aligned} D &\approx IFOV \times r \\ &= 0.0001 \times 40 \\ &= 0.004. \end{aligned} \quad (10)$$

The angle θ in Fig. 11 is the angle between a horizontal plane and the optical axis of the camera. Because we have assumed that the camera is focused at the level of the occupants' faces (see the geometry in Fig. 10), the angle θ is

$$\begin{aligned} \theta &= \arctan\left[\frac{c_{ground} - p_{ground}}{r^{'}}\right] \\ &= \arctan\left[\frac{3.6 - 1.2}{40}\right] \\ &= 3.43^0 \end{aligned} \quad (11)$$

Since we know the values for D and θ, from the geometry of Fig. 11 we can compute the footprint f_{ht} of a single pixel's horizontal translation

$$\begin{aligned} f_{ht} &= D/\sin(\theta) \\ &= 0.004/\sin(3.43^0) \\ &= 0.067\,m. \end{aligned} \quad (12)$$

We assume that the vehicle occupants travel at the nominal freeway speed of v = 65 mph or v = 29.3 m/s. At this freeway speed the footprint f_{ht} is covered during time t_f

$$\begin{aligned} t_f &= f_{ht}/v \\ &= 0.067/29.3 \\ &= 2.28\,\text{ms}. \end{aligned} \quad (13)$$

Therefore, the exposure time $t_{exposure}$ of the camera should be $t_{exposure} < 2.28$ ms if we would like to have image smear of not more than 1 pixel. The operational range of the $SU-320$ camera in terms of exposure time is 127 μs–16.3 ms. Therefore, the required exposure time of $t_{exposure} < 2.28$ ms is within the camera's operational range or, in other words, the speed of the $SU-320$ is up to the HOV task. We choose to set the exposure time of the camera to 1 ms ($t_{exposure}$ = 1 msec < 2.28 ms) which corresponds to a bandwidth of Δf = 1 kHz.

Now, that we have addressed the speed question and we know the value of Δf, we can substitute the values for A, Δf, and D^* in Eq. 7 and calculate the NEP

$$NEP = 1.18 \times 10^{-13}\,\text{Watts}. \quad (14)$$

Therefore, the camera S/N ratio will be

$$S/N = P_{pixel}/NEP = 0.7. \quad (15)$$

In conclusion, assuming a worst case scenario (overcast day, dirty windshield, dark occupant skin), we determined that the $SU-320$ camera, equipped with a $f/2$ lens, a 1.1–1.4-μm filter, and a polarizer (if it is positioned at a range of r = 40 m from the incoming vehicle and at a height of c_{ground} = 3.6 m above the ground), will achieve:

1. an acceptable smear of less than one pixel, because the required exposure time of 2.28 ms is within the camera's speed capabilities.
2. A poor S/N = 0.7. To boost the S/N to a higher value during overcast days we need to employ an illumination source. This illumination source will also be useful during nighttime. If we operated in the visible spectrum, the use of an illuminator for this application would be hazardous. Fortunately, in our case, the spectral signature of the illuminator should match the range 1.1–1.7 μm. Since this range is deep into the near-infrared spectrum, there is no danger of distracting the driver and an illuminator for this range can be safely employed in the HOV lane. Contrary to the overcast sky scenario, the S/N ratio is exceptionally good in the case of a sunny day in both bands (see Table 1). Therefore, in bright sunlight conditions, the system can operate without the requirement for an artificial light source.

4 Experimental validation

The above theoretical scenario was subjected to an empirical test at one of the Minnesota Department of Transportation's traffic monitoring and research facilities. The experiment lasted for a week in late September 1998. Traffic was diverted from one area of the research facility and this area was reserved for our exclusive use. We set up the $SU-320$ camera with all its accessories above the test lane of the freeway. In the absence of a permanent installation device, we installed the sensor suite in the basket of a cherry-picker. We implemented the experiment exactly as it was specified in Sect. 3. The only deviation was that we used one $SU-320$ camera instead of two for budgetary reasons. To simulate the operation of the two-camera system, we first equipped the camera with the lower band filter. We videotaped some carefully arranged scenes. Then, we exchanged the lower band filter with the upper band filter and again videotaped the exact same scenes. Later in the lab, we co-registered the imagery from the two bands using a warping transformation. In this way, we simulated the operation of two co-registered cameras in the lower and upper band, respectively.

We used two cars for testing. The cars made successive passes through the field of view of the near-infrared camera. The passes were done at speed increments of 10 mph, ranging from 10 to 50 mph. One of the test cars was representative of the compact category (Mitsubishi Mirage) and the other of the luxury category (Oldsmobile Aurora). The experiment had both day and night sessions. During nighttime, we used an artificial near-infrared illumination source in the range 1.0–2.0 μm that covered the illumination needs of both the lower and upper band camera versions.

Figure 12 shows the images from a particular scene in the upper and lower bands. The image in the upper band (Fig. 12b) looks in general darker than the image in the lower band (Fig. 12a) because the energy of the sun illumination in the upper band is less than that in the lower band. Proportionally, however, the face of the vehicle occupant looks much darker because of the marked difference in reflectance of the human skin features between the two bands. Figure 13b1 shows the result of the fusion operation among the images of Fig. 12. Figure 13b2 shows the result of the thresholding operation upon the fused image. As is apparent in the figure, only the face of the vehicle occupant has remained plus a few small noise regions. The noise would have not been there if we had two cameras that were accurately co-registered and operated simultaneously. This simple experiment demonstrates the enormous potential of data fusion for the two near-infrared bands we were using. Comparatively, Fig. 13a2 demonstrates how much more extraneous information remains after a thresholding operation for the visible range image of Fig. 13a1.

Table 1. S/N ratios for different day conditions and spectral bands

	1.1–1.4 μm	1.4–1.7 μm
Sunny day	700.0	500.0
Overcast day	0.7	0.5

a

b

Fig. 12a,b. Near-infrared day time results. **a** Image in the band 1.1–1.4 μm. **b** Image in the band 1.4–1.7 μm

a1

b1

a2

b2

Fig. 13a1–b2. Comparative results between the visible spectrum and the near-infrared fusion approach

5 Conclusions and future work

We have described a novel method for obtaining high-quality imaging signals for a system that will perform vehicle occupant detection and counting in the HOV lane. The method calls for two co-registered near-infrared cameras with spectral sensitivity above (upper band) and below (lower band) the 1.4-μm threshold point, respectively. The quality of the signal remains high even during overcast days and nighttime, because we can safely illuminate the scene with an eye-safe near-infrared illuminator [8]. The near-infrared cameras can also provide clear imaging signals even in certain foul weather situations, such as in hazy conditions.

The crown jewel of the method is the fusion of the co-registered imaging signals from the lower and upper band cameras. Because of the abrupt change in the reflectance for human skin around 1.4 μm, the fusion of the images results in the intensification of the occupant face silhouettes and the diminution of the background. This increased contrast allows for perfect segmentation that leaves only the face blobs of the vehicle occupants in the final processed

image. Evidently, such a clean-cut binary image will ensure the reliable and fast operation of the pattern classifier that will perform the vehicle occupant detection task as a post-processing calculation.

We are currently developing the vehicle occupant detection algorithm and results of its performance will be reported in another forum. We are also deriving the design for a prototype version of the HOV system that will be permanently installed in an actual freeway site. The prototype design calls for a slightly different arrangement of the camera set than the one described in this paper. Since the HOV system should count vehicle occupants in both the front and back seats, the cameras should not face only the vehicle's windshield but also part of the side windows. The theoretical and experimental results described in this paper define design restraints (to a high degree of approximation) for these arrangements as well.

Acknowledgements. We would like to extend our deep appreciation to Kevin Schwartz, the Mn/DOT HOVL program manager for his generous help and support. Many thanks to Ben Worel and Jack Herndon for accommodating our needs in the Mn/ROAD facility. We would also like to thank Joe Keller for his help during the road tests. Finally, we would like to thank Scott Nelson for a valuable discussion regarding some technical issues in this project. This work was supported by the Minnesota Department of Transportation under contract #*Q*5216211101. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the funding agency.

References

1. Anderson R, Parrish B, Parrish J (1981) The Optics of Human Skin. J Invest Dermatol 77(1): 13–19
2. Graham J, Hendra P (1995) Rapid Identification of Plastics Components Recovered from Scrap Automobiles. Plast Rubber Compos Process Appl 24(2): 55–67
3. Horn B (1986) Robot Vision. MIT Press, Cambridge, Mass., pp 202–277
4. Howell H, Davis J (1991) Qualitative Identification of Fibers Using NIR Spectroscopy. Text Chem Color 23(9): 69–73
5. Jacquez J, Huss J, McKeenan W, Dimitroff J, Kuppenheim H (1955) The spectral reflectance of human skin in the region $0.7 - 2.6\,\mu$m. Technical Report 189. Army Medical Research Laboratory, Fort Knox
6. Papini M (1997) Analysis of the Reflectance of Polymers in the Near- and Mid-Infrared Regions. J Quant Spectrosc Radiat Trans 57(2): 265–274
7. Sabins F (1997) Remote Sensing, Principles and Interpretation. W.H. Freeman, New York, N.Y.
8. Sinley D (1997) Laser and Led Eye Hazards: Safety Standards. Opt Photonics News 8(9): 32–37