

Towards a Hazard Perception Assistance System using Visual Motion

Chris McCarthy^{1,2}, Nick Barnes^{1,2}, Kaarin Anstey³, and Mark Horswill⁴

¹ Embedded Systems Theme, NICTA, Canberra, ACT, Australia

² Dept. Information Engineering, Australian National University, Canberra, ACT, Australia

³ Centre for Mental Health Research, Australian National University, Canberra, ACT, Australia

⁴ School of Psychology, University of Queensland, St. Lucia, Queensland, Australia

Abstract. We report on preliminary work in the application of low-level visual motion cues to identify potential hazards during on-road driving. In conjunction with a clinical study of hazard perception in older age drivers, we consider the detection of a range of hazardous scenarios identified as particularly challenging for older drivers in video sequences used in the clinical study. Central to our approach is the use of visual motion as a means of estimating self motion, from which we identify optical flow due to other motions in the scene. We report results obtained using the same hazard perception test used in clinical trials.

1 Introduction

There is growing evidence that a driver's ability to perceive hazards declines with age. The likely cause of this is age related decreases in cognitive and visual functions [1]. Population and case-control studies have found that reaction time, speed of processing, visual selective attention, executive function, eye disease and poor contrast sensitivity are associated with increased crash risk and poorer on-road driving performance [2]. Increased response time for hazard perception in older drivers has been most strongly linked to a loss of contrast sensitivity, and a reduced *useful field of view* [1]. Such visual and cognitive deficits can force older adults to cease driving, despite being otherwise capable. Forced cessation of driving can be especially difficult where public transport is not readily available (particularly rural and outer-suburban communities) and has also been linked to depression in older adults [3].

A possible alternative is to develop intervening hazard detection technologies that may compensate for the specific visual deficits causing decreased hazard perception ability. This in turn may allow otherwise capable drivers to keep driving safely, longer. To this end, the work presented in this paper forms part of a larger collaborative project investigating the effects of cognitive and visual ageing on hazard perception in older drivers. Through clinical studies, this project aims to investigate age-related differences in hazard perception, and provide insight into what factors dictate performance in hazard perception. An additional aim of the

project is to pilot possible automated hazard perception interventions that may alert a driver to specific classes of hazards in the scene. Such interventions may then be tested and validated via clinical trials.

The detection of independently moving objects in a scene has been an active field of research in computer vision for some time. Given two or more time-separated views of the scene, the problem becomes one of motion segmentation, whereby regions of homogeneous motion are grouped together. Geometric methods such as [4, 5] attempt to optimally recover the motion of all objects in the scene from matched points in two or more views. While often highly accurate, these techniques are currently not feasible for real-time use. Other approaches examine the apparent motion of objects in the scene via explicit tracking of feature points [6, 7], or via the optical flow field generated by the motion of objects [8, 9]. A particularly challenging problem is the segmentation of optical flow due to independently moving objects from flow due to self-motion. Despite much attention, the problem remains an active area of research.

In the context of road-based hazard detection, a common approach is to apply models of the expected optical flow due to self-motion, thereby identifying regions of the flow field that violate this model. This is often achieved via motion models of the road plane [10, 11]. Braillon *et al.* [10] derive a motion model for the ground plane from odometric information obtained from the vehicle. Using this, they extract the ground plane via a correlation-based generative method. Suzuki and Kanade [11] apply a parametric estimation model to obtain a camera-mounted vehicle’s ego-motion parameters. Song and Chen [12] propose a system for detecting moving vehicles entering regions on either side of a car. They detect objects that lie on the road plane via feature-based motion estimation and segmentation of flow on the road plane. An issue with road-based motion models is that local intensity variation is often too small to reliably compute optical flow, or extract feature points. In addition, obstacle detection based on the violation of motion models on the road plane ignores the possibility of objects entering from the side, that may not occupy the road plane.

Previous work in road-based hazard detection typically reports performance results using specific hand-picked scenarios. While advances have been made in the detection of moving obstacles, there has not been any significant study of how such subsystems may actually address the specific needs of a hazard perception assistance systems. Results reported typically do not consider performance over large video sequence sets depicting real, unscripted hazardous scenarios.

In this paper, we present preliminary work in the development of a system for detecting potential hazards. As a first step, we focus exclusively on the detection of one class of potential traffic conflict involving side-moving objects which enter the field of view in the periphery. We refer to these as *side-entering hazards*. To facilitate a more general frame-work for hazard detection, we do not incorporate contextual information such as the road plane, or other environmental assumptions. Instead, we focus specifically on the use of early vision cues such as optical flow. Through the estimation and subtraction of optical flow due to self-motion, we identify side-entering hazards with respect to the current direction of motion,

from the residual motion due to independently moving objects. Through this, we seek to identify regions of heightened crash risk in the periphery of the image. We examine the performance of the hazard detector over video sequences used in the clinical study of older drivers. These sequences contain real, unscripted footage of hazardous scenarios selected in consultation with road accident experts. These scenes form part of a hazard perception test used in the clinical study of hazard perception in older age drivers currently being undertaken. This allows for a thorough assessment of performance, and provides relevant feedback informing the inclusion of higher-level processing to improve performance.

The paper is organised as follows. Section 2 provides a brief background and description of the clinical study of hazard perception in older drivers, and outlines the hazard perception test used in clinical trials, and in results presented in this paper. Section 3 outlines our approach to the detection of side-entering hazards from optical flow. Section 4 provides results obtained using the system over video sequences used in the hazard perception study, and discussion of these results. Section 5 presents conclusions and future work.

2 Hazard perception and cognitive ageing in older drivers: investigation and development

The work presented in this paper forms a part of a larger clinical study, aiming to characterise the effects of cognitive ageing on hazard perception in older drivers. In so doing, this study aims to inform the development of interventions to improve hazard perception and road safety in older drivers.

To this end, the study focuses on identifying and quantifying the relationships between types of visual cognitive abilities and types of hazards, thus forming a typology that can be used in the development of computer vision technology to improve hazard perception in older drivers.

The investigation is split into three studies:

- *Study 1* is a large scale investigation of the cognitive abilities underlying hazard perception, as well as potential mediating effects of visual function and state anxiety. A sample of 300 community dwelling older drivers have been selected to sit the hazard perception test (discussed below).
- *Study 2* validates the hazard perception test as a measure of driving skill in older drivers with an on-road driving test.
- *Study 3* is a pilot study aimed at trialling possible interventions, including the use of computer vision technology to compensate for the deficits in cognitive and visual ability that affect hazard perception performance.

2.1 The hazard perception test

The hazard perception test used in the present study is a video test incorporating local (Canberra, Australia) road hazards. The test is an adaptation of a previously used technique developed by Horswill *et al.* [13]. Participants view

video footage of a driver’s eye view of various genuine traffic hazards. They are instructed to respond by pressing a touch-screen when they detect a potential traffic conflict (where the camera car might have to brake or take evasive action to avoid a collision). Reaction times to selected incidents on the video are measured and averaged to give an overall hazard perception reaction score. The test has been shown to differentiate between novice and experienced drivers. Previous versions of this test have correlated with accident involvement and driving instructor ratings. For the present study, the test has been modified specifically for older drivers. Consultation with focus groups and road safety experts has been conducted and used to identify scenes to be included in the video footage.

The hazard perception scenes used in the present study were filmed using a video camera attached to the inside of the windscreen, giving a driver’s eye view of the traffic ahead. All footage was filmed around the local area of the study in normal traffic. Whenever a potential traffic conflict was encountered, this was indexed on the tape. The clips were then rated for suitability for inclusion in a local hazard perception test (for example, the predictability of conflicts was rated). To validate the clips, their ability to distinguish between younger novice drivers (a high crash risk group) and older experienced drivers (a lower crash risk group) was examined (with the prediction that the lower risk experienced drivers should be faster at detecting the hazards). It was found that older experienced drivers responded significantly faster than younger novices. It is important to emphasize that all the clips depicted genuine, unstaged hazardous events, unlike the staged events used in other hazard perception tests.

3 Detecting peripheral side-moving hazards

We employ the use of visual motion to highlight regions of the image representing a heightened risk of a side-entering hazard occurring. To distinguish optical flow due to self-motion from flow due to independently moving objects, we apply a model of visual motion due to self-motion over the whole scene. We do not base segmentation on models of visual motion on the ground plane. Through the estimation and subtraction of optical flow due to self-motion, we identify side-entering hazards with respect to the observer’s current direction of motion using the residual motion after self-motion subtraction. We assume a forward facing camera undergoing predominantly translational motion. Note that while we do not apply de-rotation to the flow field, de-rotation algorithms such as [14] can be applied to eliminate rotational flow.

We divide the task into two stages. Firstly, the removal of self-motion from the peripheral image motion field. Secondly, the segmentation of regions of residual flow (after self-motion removal) associated with moving objects deemed to be side-entering hazards. Each is discussed below.

3.1 Stage 1: Removing self-motion

We seek to remove the components of flow due to the motion of the camera. To achieve this, we consider only the direction of flow vectors in the image.

By generating an expected pattern of self-motion induced flow directions across the image, we may then subtract this model from unit vectors of the estimated optical flow field. Residual flow indicates the existence of independent motion and thus can be examined further in the second stage.

Let $\dot{\mathbf{t}}_c = [t_x \ t_y \ t_z]$ be the velocity of a camera-mounted vehicle. Assuming translational motion only, the optical flow due to camera motion is given by:

$$u_c = \frac{t_z}{z} \left(\frac{f_x t_x}{t_z} - x \right), \quad v_c = \frac{t_z}{z} \left(\frac{f_y t_y}{t_z} - y \right), \quad (1)$$

where f_x and f_y are focal lengths (in pixels), and (x, y) is the projected image location of a point $P = [x \ y \ z]$ in a camera-centred coordinate system. Assuming motion is predominantly on the ground plane, and is approximately along the optical axis, we set t_x and $t_y = 0$. Note that while we assume motion is, on average, in the direction of the optical axis, we account for small shifts of the FOE inevitably introduced under real-world conditions. This is primarily due to small rotational effects during ego-motion, causing the direction of motion to move with respect to the optical axis of the camera. We therefore incorporate the location of the FOE, (x', y') into (1). Considering only the direction, we then obtain:

$$u_c = -\frac{(x-x')}{\|\mathbf{u}\|}, \quad v_c = -\frac{(y-y')}{\|\mathbf{u}\|}, \quad (2)$$

where $\|\mathbf{u}\|$ is the norm of the flow vector. The above equations can then be used to generate a motion template for self motion. Notably, the above template does not account entirely for the perspective effects introduced by shifts of the translational axis with respect to the optical axis. FOE tracking should, however, provide a sufficiently accurate template for small rotational effects.

It should be noted that this technique does not differentiate flow due to self-motion from all possible independently moving objects in the scene. Specifically, objects moving in the exact opposite direction to the camera will generate the same directional flow pattern as that generated from self-motion. This problem, however, does not arise for side-entering hazards, where the optical flow generated by the object's motion will always be significantly different to that induced by self-motion.

3.2 Stage 2: Identifying side-entering hazards

The subtraction of the self-motion template from the unit vectors of the estimated flow field yields residual motion directions due to independently moving objects. We therefore seek to identify those residual flow regions corresponding to *side-entering* hazards.

Let H be an independently moving object with velocity $\dot{\mathbf{h}} = [h_x \ h_y \ h_z]$. Let θ_h be the direction of motion of H on the ground plane, with respect to the Z axis, such that:

$$\theta_h = \arctan\left(\frac{h_x}{h_z}\right). \quad (3)$$

We consider H to be a *side-entering* hazard if $h_x < 0$ and $0 \leq \theta_h \leq \frac{\pi}{2}$, or $h_x > 0$ and $-\frac{\pi}{2} \leq \theta_h \leq 0$. That is, H is a side-entering hazard if there exists a

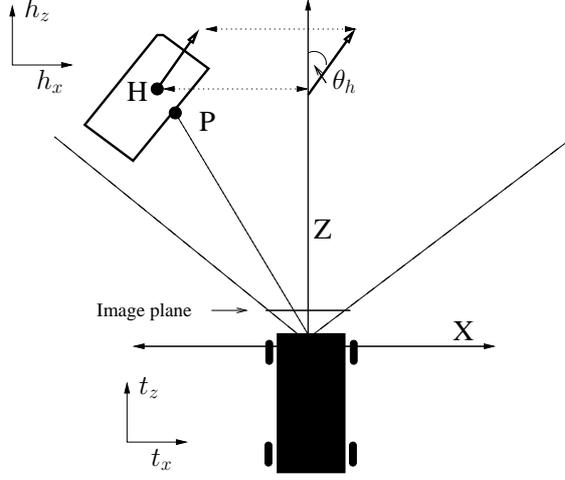


Fig. 1. Geometric framework for side-entering hazard detection.

component of horizontal motion towards the Z axis. Figure 1 shows the geometric framework used.

We seek to infer the direction of motion of H in the $X - Z$ plane from the projection of its apparent motion in the image plane. Let $P = [p_x \ p_y \ p_z]$ be a point on H . Assuming a pinhole camera model, we project P into the image plane such that:

$$p_x = f_x \frac{p_x}{p_z}, \quad p_y = f_y \frac{p_y}{p_z}, \quad (4)$$

where (p_x, p_y) is projected location of P in the image plane (with origin at image centre), and (f_x, f_y) are focal lengths in pixels. Setting both focal lengths to 1, and considering only translational motion in the $X - Z$ plane, we obtain the following equations for the image velocity of P :

$$u_p = \frac{h_z}{p_z} \left(\frac{h_x}{h_z} - p_x \right), \quad (5)$$

$$v_p = -\frac{p_y h_z}{p_z}, \quad (6)$$

where (u_p, v_p) are the horizontal and vertical components of the image velocity of P .

Notably, u_p and v_p provide a linear system of equations relating the unknown object velocity direction components: h_x and h_z . While obtaining h_x and h_z directly from these equations is not possible, the ratio of these components can be obtained. Re-arranging (6) such that:

$$\frac{h_z}{p_z} = -\frac{v_p}{p_y}, \quad (7)$$

and substituting into (5), after simple re-arrangement we obtain:

$$\frac{h_x}{h_z} = -\frac{p_y u_p + p_x v_p}{v_p}. \quad (8)$$

Substituting back into (3), we obtain an equation for the direction of motion of H :

$$\theta_h = \arctan\left(-\frac{p_y u_p + p_x v_p}{v_p}\right), \quad (9)$$

in terms of known and measurable visual quantities. From this we can identify regions of motion corresponding to side-entering hazards, as defined earlier.

4 Implementation and results

We implement the two stages using the estimated direction of optical flow vectors computed in both peripheral regions of the image. These *side-hazard regions* are placed on either side of the estimated location of the FOE, at a preset horizontal distance, d , from the FOE. For the trials presented here, $d = 0.1 \times \text{image_width}$.

To estimate the location of the FOE, we employ a Hough-based voting approach to find the intersection point of computed optical flow vectors. The estimated FOE location is then used to generate the self-motion template as defined in (2). This template is then subtracted from the estimated unit vector flow field, thus leaving only flow due to the motion of independently moving objects in the scene. This residual motion is then passed to the second stage.

We implement Stage 2 by convolving a 5×5 weighted window over u and v separately to obtain the relative support of visual motion in both directions. This is then used to compute θ_h as defined in (9). By considering the computed θ_h at each image location, side-entering hazard regions are constructed via a simple region-growing technique, whereby neighbouring pixels also classified as side-entering hazards are grouped together. From this, a bounding box is computed. Figure 2 gives a sample frame showing the estimated FOE, the peripheral regions used for detection, and a region of the image identified as a potential hazard (with optical flow).

To improve robustness to false positives, temporal support is also included. A hazard detection alert is not issued unless the region associated with the possible hazard has been identified as a hazard in the last two updates. If no additional support is received after three frames, the hazard region is considered invalid, and thrown away.

Optical flow is computed using Lucas and Kanade's [15] gradient-based method with eigenvalue thresholding to discard flow in regions of low intensity variation. Flow vectors were computed for every 8th pixel, over images of resolution 360×288 pixels.

4.1 Results

To assess the accuracy and robustness of the *side-entering* hazard detector, we employ the same video stimuli used in the study of hazard perception abilities

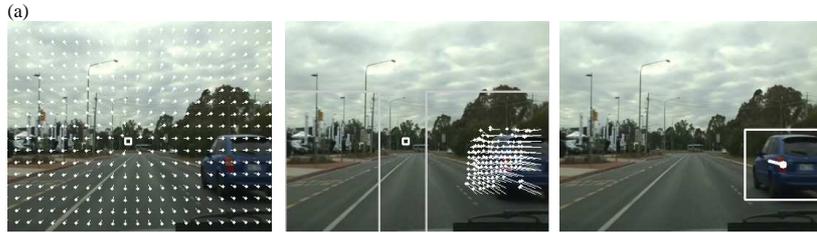


Fig. 2. A sample side-entering hazard scenario showing from left to right: 1. the computed self-motion template and estimated location of the FOF, 2. the estimated optical flow for the segmented region, and the peripheral regions of interest, and 3. the marked hazard, and estimated direction of motion in the image plane.

in older drivers. Six video segments containing a number of indexed, unstaged driving hazards were used. From the set of all indexed hazards across the video segments, those fitting the description of side-entering were marked as hazards to detect. Start and end time-codes listed for each of the indexed side-entering hazards were used to define the duration of time in which the detector must locate the hazard. Detected hazards were indicated by a bounding box drawn around the image region associated with the potential threat.

Table 1 provides a full list of all indexed side-entering hazards. A brief description of the side-entering scenario is given, along with the time interval defined for the hazard. Note that the time interval is the same as that used in human trials with the same footage. The right two columns provide results obtained from the application of the hazard detector over the video segments. Where the detector was successful, the time-code of the initial detection, and the time difference of this detection with respect to the indexed time-code are given. Figure 3 shows sample hazard detections recorded over the video segments. Each corresponds to a successfully detected indexed hazard in Table 1.

Across the video segments, the detector identified 24 of the 30 (80%) indexed side-entering hazards. Of the total number of hazard alerts issued, 41% were observed to be false positives. A false positive was deemed to be any hazard alert not involving a moving vehicle or person.

4.2 Discussion

Table 2 shows a breakdown of performance statistics into the major classes of side-entering hazards. The strongest results achieved for speed of detection involve situations where a vehicle enters the field of view moving. This is in contrast to the worst performing scenario for detection time involving vehicle's pulling out from an initially stationary position. It should be noted that the indexed start time of these hazards is significantly earlier than when the vehicle starts to move. Arguably, these hazards constitute a *stopped vehicle in lane*

Indexed Hazards			Detector Results	
Hazard description	Time duration		Time detected (sec)	Response time (sec)
	start	end (sec)		
Video segment 1				
merge in front from right	2.96	11.75	2.80	-0.16
merge in front from right	21.60	32.00	24.39	+2.79
car turns out from left	49.72	55.96	51.91	+2.19
pedestrians crossing road	140.84	157.36	153.85	+13.01
bus pulls out	161.839	178.05	178.32	+16.92
truck merges from left	251.04	283.60	271.45	+20.41
Video segment 2				
pedestrian crossing from right	8.95	20.20	15.72	+6.77
van swerves right from left	56.24	64.94	57.47	+1.23
pedestrian crossing from right	145.83	156.76	-	-
car turns out from left	202.80	209.16	205.08	+2.28
bus turns out from left	253.16	264.14	255.21	+2.05
Video segment 3				
car merges right	22.91	39.65	16.67	-6.24
pedestrians crossing from left	130.82	141.88	-	-
pedestrian crossing from right	240.86	245.52	243.59	+2.73
Video segment 4				
truck pulling out from right	0	13.08	9.25	+9.25
car turns out from left	75.64	83.24	80.17	+4.53
car merges from left	153.793	169.24	152.92	-0.87
car on round-about from right	170.28	179.36	172.18	+1.9
pedestrian moving from left	261.73	274.04	-	-
Video segment 5				
car merges from left	23.79	39.08	26.64	+2.85
pedestrian crossing from right	54.88	63.61	-	-
pedestrians crossing from left	98.69	105.93	101.78	+3.09
car crosses road from right	107.32	114.60	109.84	+2.52
pedestrian crossing from right	157.32	165.13	162.20	+4.88
pedestrian crossing from right	204.20	215.6	-	-
Video segment 6				
car turns out from right	11.32	19.18	13.40	+2.08
bus starts pulling out from left	65.49	77.20	75.21	+9.72
pedestrians crossing from right	84.05	96.30	91.36	+7.31
van pulls out from left	126.16	139.68	138.46	+12.3
bus pulls out from left	152.84	174.5	166.93	+14.09
truck enters round-a-bout from left	186.72	226.39	-	-

Table 1. Results obtained using the side-entering hazard detector across all video segments used in hazard perception testing of older drivers.



Fig. 3. Sample hazard detections from video segment (vs) testing: (a) merge from right (vs 1, 29.36 sec), (b) car turns out from left (vs 1, 52.79 sec), (c) truck merges from left (vs 1, 271.45 sec), (d) pedestrians crossing (vs 5 101.78 sec), (e) bus pulls out from left (vs 6 166.93 sec), (f) car crosses road from right (vs 5, 109.84 sec)

Hazard class	Detections	Avg response time (secs)
Vehicle side road entry	7/8	+2.27
Vehicle merge (or swerve) to front	6/6	+3.4
Pedestrian(s) crossing	6/11	+6.3
Vehicle pull out from curb	5/5	+12.47
Total	24/30	+5.65

Table 2. Hazard detection results broken down to major side-entering hazard classes.

hazard rather than side-entering when they initially enter view. All such hazards were still detected within the allotted time.

Pedestrian-related hazards posed the greatest challenge (6 from 11 were identified). This is in contrast to the very strong results achieved for side-entering vehicles, a result also reflected in average response times. Successful pedestrian-related detections took, on average, twice as long as side-entering vehicle detections. The likely cause of this discrepancy is the relatively slow apparent motion, and small size, of pedestrians as compared with vehicles. In addition, pedestrians often entered the field of view as stationary objects waiting to cross the road, making their detection difficult.

While the number of recorded false positives was significant, their occurrence was predictable, and limited to specific environmental scenarios. Of the total number of false positives recorded, 83% were found to be the result of lines, shadows, and other features on the surface of the road. In many cases, these

features would remain for a significant time, thus causing repeated detections. Given an extraction of the road-plane, such false detections should easily be filtered out, thus reducing false-positives to 7%.

4.3 Future work

The hazard detection system presented represents preliminary work in the development of a full hazard perception assistance system. Future work will consider other classes of hazards, and the use of other visual information in which to detect hazards. Cues such as flow field magnitude and divergence (or looming) provide a direct gauge of the relative proximity of objects in the scene. Such cues may also provide a means of gauging the level of threat posed by environmental conditions in general. For example, increasing flow magnitude in the periphery would suggest conditions are narrowing, thus increasing the risk of pedestrians or other objects entering from the side. To facilitate more pre-emptive hazard perception, the inclusion of subsystems to detect more contextual cues such as road signs, flashing indicators and stop lights will also be considered.

5 Conclusion

The ability to perceive potential hazards is crucial to safe driving. Research indicates that a driver's ability to perceive hazards declines with age. We have reported on preliminary work towards the development of a potential hazard perception intervention to assist older drivers. As part of a larger collaborative study investigating these questions, we have presented preliminary work in the development of potential interventions to improve hazard perception in older adults. A class of hazard identified as a cause of heightened crash risk are those involving side-entering objects entering the field of view in the periphery. We have proposed a simple and efficient strategy for identifying regions of the image where the likelihood of such hazard is high. Unlike previous approaches, we gauge the performance and effectiveness of the detector over six video segments also being used in concurrent clinical trials of older drivers. As part of this study, we aim to adapt what is learnt through clinical trials, and pilot an intervention to improve hazard perception in older drivers. Such a system may allow older drivers to keep driving safely, for longer.

Acknowledgements

Video sequences used in testing were filmed and captured by Chris Hatherly from the Centre for Mental Health Research, Australian National University. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program. The Hazard Perception Study was supported by the Australian Research Council and the NRMA ACT Road Safety Trust. Anstey is supported by NHMRC Fellowship # 366756

References

1. Horswill, M., Marrigton, S., McCullough, C., Wood, J., Pachana, N., McWilliam, J., Raikos, M.: The hazard perception ability of older drivers. *Journal of Gerontology: Psychological Science* (in press)
2. Anstey, K., Wood, J., Lord, S., Walker, J.: Cognitive, sensory and physical factors enabling driving safety in older adults. *Clinical Psychology Review* **25** (2005) 45–65
3. Marattoli, R., de Leon, C.M., Glass, T., Williams, C., Cooney, L., and M.E Tinetti, I.B.: Driving cessation and increased depressive symptoms: Prospective evidence from the new haven epese. established populations for epidemiologic studies of the elderly. *Journal of the American Geriatric Society* **45** (1997) 202–6
4. Li, H.: Two-view motion segmentation from linear programming relaxation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2007) 1–8
5. Vidal, R., Hartley, R.: Motion segmentation with missing data using PowerFactorization and GPCA. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (June 2004) II–310 – 316
6. Yamaguchi, K., Kato, T., Ninomiya, Y.: Moving obstacle detection using monocular vision. In: *Proceedings of the 2006 Intelligent Vehicles Symposium*. (2006) 288–93
7. Cohen, I., Medioni, G.: Detecting and tracking moving objects in video from an airborne observer. In: *Proceedings of IEEE Image Understanding Workshop*. (1998) 217–22
8. Mitiche, A., Sekkati, H.: Optical flow 3d segmentation and interpretation: A variational method with active curve evolution and level sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **28**(11) (2006)
9. Weber, J., Malik, J.: Rigid body segmentation and shape description from dense optical flow under weak perspective. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **19**(2) (1997) 139–43
10. Braillon, C., Pradalier, C., Crawley, J., Laugier, C.: Real-time moving obstacle detection using optical flow models. In: *Proceedings of the 2006 IEEE Intelligent Vehicles Symposium*. (2006) 466–71
11. Suzuki, T., Kanade, T.: Measurement of vehicle motion and orientation using optical flow. (1999) 25–30
12. Song, K.T., Chen, H.Y.: Lateral driving assistance using optical flow and scene analysis. In: *Proceedings of the 2007 IEEE Intelligent Vehicles Symposium*. (2007) 624–9
13. Horswill, M., McKenna, F.: Drivers’ hazard perception ability: Situation awareness on the road. In: *A Cognitive Approach to Situation Awareness*. Ashgate (2004) 155–75
14. Lim, J., Barnes, N.: Directions of egomotion from antipodal points. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2008) in press
15. Lucas, B., Kanade, T.: An iterative image registration technique with an application to stereo vision. In: *Proceedings of DARPA Image Understanding Workshop*. (1984) 121–30