Paper:
Wilkie, RM, Wann, JP and Allison, RS (2011) *Modelling Locomotor Control: the advantages of mobile gaze.* ACM Transactions on Applied Perception, 8 (2). 1 - 18 (18).

http://dx.doi.org/10.1145/1870076.1870077

Modelling Locomotor Control: the advantages of mobile gaze

RICHARD M. WILKIE
Institute of Psychological Sciences, University of Leeds, UK
JOHN P. WANN
Department of Psychology, Royal Holloway, University of London, UK
and
ROBERT S. ALLISON
Centre for Vision Research, University of York, Toronto, Canada

In 1958 JJ Gibson put forward proposals on the visual control of locomotion. Research in the last 50 years has served to clarify the sources of visual and non-visual information which contribute to successful steering, but has yet to determine how this information is optimally combined under conditions of uncertainty. Here we test the conditions under which a locomotor robot with a mobile camera can steer effectively using simple visual and extra-retinal parameters to examine how such models cope with the noisy real-world visual and motor estimates that are available to humans. This applied modelling gives us an insight into both the advantages and limitations of using active gaze to sample information when steering.

Categories and Subject Descriptors: ... [**...**]: ...

General Terms: Vision, Perception, Action, Cognition

Additional Key Words and Phrases: Robot, Locomotion, Steering, Gaze, Eye Movements, Active Vision

1. INTRODUCTION

For most animals successful locomotor control is fundamental for survival. For many predators it is an integral part of catching prey, and likewise, for prey species, it can be essential for avoiding capture. While the underlying anatomical structures and perceptual information systems vary dramatically between animal species the requirements for steering control remain similar: it is generally true that mobile

ACKNOWLEDGMENTS

Thanks to the Centre for Vision Research, University of York, Toronto for their support during the collaborative research visit that contributed to the work presented here. Thanks also to Georgios Kountouritos for his technical help with the manuscript. Research supported by the UK EPSRC GR/S86358 and EP/D055342/1.

animals change their physical orientation to be aligned with the steering goal. This reorientation can sometimes be performed by rotating on the spot before locomotion begins (e.g. jumping spiders, Land [1971]) but in some situations this has to occur gradually during locomotion (e.g. bats in flight, Ghose and Moss [2006]). Humans use both these strategies: when stationary they tend to orient in the direction of the target before walking a straight-line path, whereas when travelling at higher speeds (e.g. running or cycling) the direction of motion is adjusted gradually to align with the target. Humans will initially direct their gaze (eyes and head) in the direction they wish to steer and then gradually rotate their body so that head and eyes return to centre [Land 1992; 2004; Wilkie and Wann 2005]. A flying bat is similar in that there is a tight relationship between the direction of acoustic gaze (sonar beam direction) and the direction of flight as it searches for and intercepts insect prey [Ghose and Moss 2006]. The natural locomotor speed of humans would not usually reach that of a flying bat, however humans have developed vehicles that significantly increase speed of locomotion (e.g. bicycles and cars.) At higher speeds gaze direction seems to be particularly important for ensuring successful and effortless control of steering [Land and Lee 1994; Wilkie and Wann 2005]. The robust and effective control of high-speed locomotion seems to be supported by the ability of the visual-motor system to integrate multiple sources of information even under conditions of uncertainty. The sources of information that humans use to control active locomotor behaviours have been investigated in tasks such as walking [Rushton et al. 1998; Warren et al. 2001], driving [Land and Tatler 2001; Salvucci and Gray 2004; Wilkie and Wann 2002]; and cycling [Wilkie et al. 2008] and it has now been shown that there are a number of visual and non-visual sources that inform steering [Rushton et al. 1998; Warren et al. 2001; Wilkie and Wann 2003a; 2005]. The availability of multiple information sources has even been used to explain the ability of humans to control steering when travelling high speeds when driving in difficult visual conditions, such as at night [Wilkie and Wann 2002].

Gibson [1958] first put forward a conceptual framework for the control of locomotion based on the use of optic flow (the pattern of relative motion between an observer and surface textures in the world). Gibson suggested that a simple but effective control strategy would be to ensure that the focus of expansion (FoE) of the optic flow field emanates from the chosen direction of travel (i.e. aligning 'heading' with the goal). Eye movements introduce additional rotation components to optic flow at the retina [Royden et al. 1992] which results in a 'retinal flow' pattern where the FoE is no longer a simple feature available within the flow field. It is sometimes assumed that these rotation components need to be removed in order that the FoE can be recovered, however, we believe there are at least two problems with this view of locomotor control [Wilkie and Wann 2006]: 1) it presumes that simply having access to current heading is sufficient for controlling locomotion when this does not necessarily specify the steering required to reach the goal, and 2) it treats eye movements as a complication without which steering would be simpler. We believe that active gaze behaviours are crucial for carrying out effective highly skilled real-world actions (for a variety of examples see Land and Furneaux [1997]) and we have shown that mobile gaze results in fewer steering errors than when gaze is fixed [Wilkie and Wann 2003c]. It has even been suggested that mobile gaze directed towards the steering target performs an essential transformation to the retinal flow field that specifies whether current steering behaviour is sufficient to produce a course that will pass through the fixation point [Kim and Turvey 1999; Wann and Swapp 2000]. Because orienting gaze towards the desired steering goal supplies non-visual direction information about that target (from the motor commands to and proprioception from muscles controlling eye and head) it is possible that a flexible weighted combination of visual and non-visual sources act to underpin successful steering [Wilkie and Wann 2005; Wilkie et al. 2008]. The examination of eye movements when steering can, therefore, expand our understanding of what information is used and when [Land and Hayhoe 2001; Land and Lee 1994; Robertshaw and Wilkie 2008; Wilkie and Wann 2003c].

1.1 Why build a Locomotor Robot?

Bicho and Schoner [1997] and Murray et al. [1997] have demonstrated the power of using robotic platforms to examine the problem of locomotor control, and these studies have helped to shape our understanding of human steering (e.g. Fajen and Warren [2003]). A considerable body of research accumulated over the past forty years has centred around how optic flow can be used by humans during locomotion, with less attention being paid to sources of direction information (with some notable exceptions, e.g. Llewewllyn [1971]). While recent studies have demonstrated that both visual direction and extra-retinal direction can influence steering behaviour (e.g. Wilkie and Wann [2002; 2005]), it can be difficult to directly test the use of specific informational sources in humans. Here we will be primarily considering whether steering strategies that have been put forward for human locomotion [Wilkie and Wann 2002; Wilkie et al. 2008] are able to steer a robot in real-world settings. There are clear differences in the motor output of a robot in comparison to a running human, but human steering control readily adapts to a range of contexts that involve different output devices with varying vehicle dynamics (e.g. on foot, on a bike, or in a wheelchair.) We feel, therefore, that the differences in the motor action of a human and a wheeled robot are not critical. For our purposes we only need a robot with a simple mobile camera to act as an input device, to determine the requirements for successfully directing motor output. We propose that a simple robotic platform can be an important test of control models that are suggested to generalize across a wide range of human locomotor capabilities (in line with Schoner et al. [1995]). We can take what we believe to be critical parameters for the control of visually guided action and study manipulations of the camera image in an unforgiving real-world environment. This introduces real noise, latencies, complex interactions, imprecision and problems that are sometimes unanticipated. By comparison, pure simulations and ideal-observer analyses are totally determined by the input parameters. In this respect designing a robot that will move autonomously through the world focuses the issue of what informational variables are critical to the visual control of locomotion and serves to outline what is necessary and what is sufficient [Rushton et al. 2002].

1.2 Building a Locomotor Robot

The natural starting point when building a locomotor robot is to mount a single static camera on a mobile base. The camera provides rich visual information about the illuminated scene, though not without limitations stemming from a restricted field of view and finite spatial resolution of the sensor. Having the camera fixed on the locomotor platform in a known position has a number of advantages: all transformations to the camera image can be assumed to be due to rotations or translations of the robot (ignoring for the sake of simplicity other mobile objects) and therefore the flow of elements on the camera image indicates the locomotor activity of the robot.

For humans most forms of locomotion align the direction of motion with the midline of the body, so that when you walk, run, cycle or drive forwards your shoulders are perpendicular to the direction of travel [1]. This is also the case for our robot since the camera is positioned in the centre of the chassis, so the direction of travel is aligned with the midline of the camera and the centre of the resulting image. This means that a target to the right half of the camera image is located in space in front and to the right of the robot and would require a clockwise turn to orient towards this target (see Figure 1). The simplest method of reaching the target would be to simply null the yaw angle between the centre of the image and the target before inititing locomotion (α nulled; Figure 2, Left Panel), and this strategy is often used by humans when walking to a single target from a standing start. When travelling at higher speeds, however, it becomes important to close down α at a controlled rate ($\dot{\alpha}$) rather than instantaneously since momentum can lead to the vehicle overturning.

1.3 Introducing a mobile eye

For the human visual system if a target is in the extreme visual periphery then we would expect the estimate of α to be more erroneous than when it is close to the fovea[2]. One function of mobile gaze in humans is to ensure that the important visual information lies upon the highest resolution part of the retina. Optical distortion in the periphery does occur with camera systems, for example when a fish-eye lens or panoramic mirror is used to increase the field of view, so both biological and machine vision systems can benefit from mobile gaze (remarkably even the tiny jumping spider benefits from a pair of mobile principal eyes [Land 1969]). Even when a camera has an equally distributed resolution across the whole image there is an advantage for mobile gaze. A sensor that can be rotated in both yaw and pitch has a larger field of regard because a greater region of the scene can be quickly sampled without resorting to whole body movements. This is an important advantage of mobile gaze in biological systems. The disadvantage for locomotion, however, is that it could complicate the extraction of visual information from the camera/eye image. Not only would it add rotation components to the retinal flow (as discussed previously) but any estimate of α would need to be computed from a combined measurement based on the direction of the target in the camera coordinate frame of reference, and the relationship of the camera coordinate frame to the robot body. Thus, uncertainty in the controlled parameter (body-centric α) would depend on

[1] There are interesting exceptions, for example when snowboarding or skateboarding, that may require alternative mappings between body orientation and locomotor direction to be learnt.
[2] When a target is visible in the higher resolution areas of the retina, eye-movements towards it are more accurate [Gnadt et al. 1991; White et al. 1994]

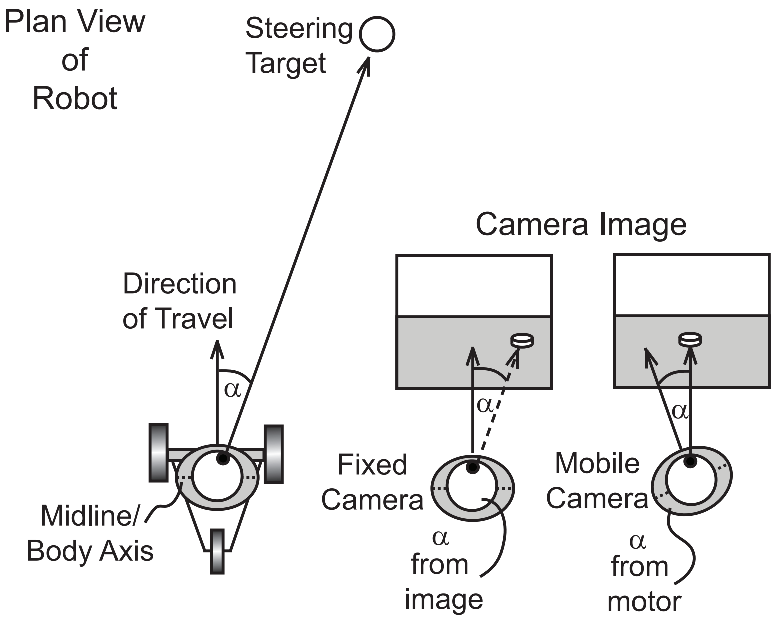

Fig. 1. A locomotor robot with either a fixed or mobile camera. The direction of travel is always fixed relative to the midline of the robot, as is usually the case in humans.

both noise in the 'retinal' α estimate derived from the camera image (before the camera moved) and the 'extra-retinal'[3] sensor reading of the position of the camera. Noise may therefore propagate through the system, with errors in retinal estimates of α being compounded by additional noise in the motor output that orients the camera to fixate the target, as well as further potential inaccuracy in sensing the final position of the camera. The extra-retinal angle must be estimated relative to a known midline and because proprioceptive signals are susceptible to bias or drift and they may require calibration (e.g. Wann and Ibrahim [1992]). If the accuracy of the retinal signal varies with eccentricity (as in the human visual system as described above) then reorienting gaze may result in an improved composite signal, provided the extra-retinal signal is not poorer than peripheral retinal signals. In all cases of mobile gaze, however, the precision of the controlled variable depends on both the retinal and extra-retinal signals. It could be argued, therefore, that a mobile gaze fosters greater opportunities for introducing noise into estimates of target direction, when compared with fixed gaze. Here we build a robot with a mobile camera to investigate how well active gaze models of human locomotion (c.f. Wilkie et al. [2008]) perform under noisy real-world conditions.

[3]Extra-retinal is used here to refer to non-visual sensing of the camera pose to stress the analogy with biological vision where the qualifier 'extra-retinal' has been used to refer to non-visual information about the position or orientation of the eye such as efference copy or sensory afference from eye-muscle proprioceptors. In the human system an extra-retinal signal would usually be needed to keep gaze fixed (to prevent drift), whereas in a robot gaze could usually be fixed without the requirement for an extra-retinal signal.

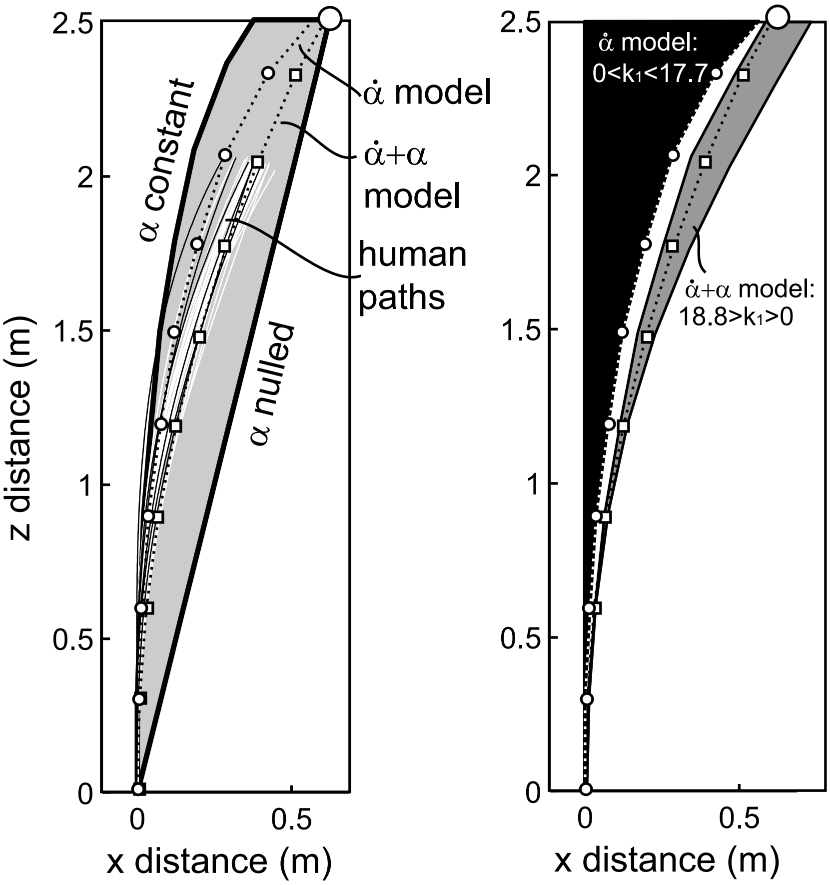

Fig. 2. Steering paths taken when steering to a target (the large open circle) offset by 14°. Left Panel: the thick solid black lines indicate the extreme strategies of keeping α constant or nulling α. The dotted lines represent paths taken by the steering models (open circle symbols = $\dot{\alpha}$ model; open square symbols = $\dot{\alpha} + \alpha$ model). Both models had the damping parameter $b = 1$; $\dot{\alpha}$ was scaled by $k_1 = 4.7$ and α was scaled with $k_2 = 1$ in the $\dot{\alpha} + \alpha$ model; $k_1 = 17.7$ in the α model. The thin curved solid paths that terminate at 2m are a series of human trajectories. Two sets of data are shown: dark lines are trajectories from Wilkie and Wann [2003a] where retinal flow and extra-retinal direction information was available; white lines are trajectories from six participants during the night-time condition of Wilkie and Wann [2002] where only extra-retinal direction information was available. The human paths have been scaled to fit the dimensions of the robotic environment as detailed in the text. Right Panel: Zones defined when using a variety of values for k_1. The black zone shows the possible stable trajectories of the α model with $b = 1$. The grey zone shows the possible stable trajectories of the $\dot{\alpha}+\alpha$ model with $b = 1$ and $k_2 = 1$. Other stable trajectories are possible with the $\dot{\alpha}+\alpha$ model by reducing k_2 but these trajectories naturally start to overlap the black zone as α contributes less.

1.4 Active gaze models of steering

To implement an autonomous steering robot requires clearly identified perceptual inputs that can be converted into a steering response. The Wilkie and Wann [2002] model combined information from a number of perceptual variables to provide a rotation input ($\dot{\alpha}$) for the steering system. In this model each input provides a weighted estimate that is correlated with the rate of change of the visual angle of the fixation target (Equation 1, referred to henceforth as the $\dot{\alpha}$ model).

$$\ddot{\theta} = k_1\dot{\alpha} - b\dot{\theta} \tag{1}$$

where $\dot{\theta}$ is the rate of current steering response damped by the parameter b, and k_1 scales the combined perceptual inputs ($\dot{\alpha}$); these terms are used to generate an acceleration in the steering response ($\ddot{\theta}$). Equation 1 is also equivalent to the steering system proposed in Fajen and Warren [2004] which nulled the rate of rotation between the goal and instantaneous heading. Wilkie and Wann [2002] proposed that the perceptual inputs could be made available from the rotation of the retinal flow field, or a retinal estimate of changing target direction. For the purposes of our robot, however, this will be supplied by an extra-retinal estimate of the rate of change of target direction (equivalent to gaze rotation for a fixated target) which has been shown to provide a powerful source of information when steering [Wilkie and Wann 2005].

When walking, with only a small amount of forward momentum, it is possible to turn on the spot. This means that the target angle (α) can be instantaneously nulled, and a direct path to target can be taken (Figure 2, α nulled). When travelling at speed it is not safe to instantaneously null α, so the rate at which α is reduced ($\dot{\alpha}$) must be controlled. One strategy would be to null $\dot{\alpha}$, thereby keeping α constant as shown in Figure 2. This will result in a spiral path that turns towards and passes close to the target. It would, in principle, never completely intersect the target, so the final part of the trajectory would need to be taken using a simple direct path [Lee 1998]. Maintaining a constant target angle is a strategy used by raptors when approaching prey [Tucker 2000] who break out of the spiral near the end of the flight to take a rapid straight path to their prey. A constant α model has also been demonstrated to be effective for controlling a locomotor robot [Rushton et al. 2002]. In humans, when steering to a target, the rate of steering depends upon the quality of information that is made available (as discussed later), however humans tend to follow a course that turns more quickly than merely keeping α constant or Equation 1 (Figure 2, Left Panel). Steering control effected using Equation 1 relies upon the $\dot{\alpha}$ signal which remains small when the steering target is distant, so most of the steering occurs towards the second half of the trajectory (when $\dot{\alpha}$ has increased). Simply increasing the response stiffness (k_1) does not solve this issue, since it makes the model prone to extreme oscillations (Figure 2, Right Panel). To capture the tighter trajectories taken by humans Wilkie et al. [2008] proposed an extension to Equation 1, whereby a reduction of the gross angular offset of the target (α) also contributes to the steering output (Equation 2, referred to henceforth as the $\dot{\alpha} + \alpha$ model):

$$\ddot{\theta} = k_1\dot{\alpha} + k_2\alpha - b\dot{\theta} \tag{2}$$

where k_1 and k_2 are independent scaling factors to control the rate of closure of the angular terms. A robot steering using Equation 2 will turn more quickly when the steering target has a large offset than when the target is nearly in front of it, but the rate of closure of α will still be smoothly controlled[4]. Figure 2 shows trajectories when steering to a target offset by 14°. In the human trials participants steered towards an eccentric target which was positioned 60m away (though trials were stopped at 50m to remove performance feedback). To allow comparison of this human data with the robotic paths we have scaled down the trajectories by a factor of 24 which then matches the range over which we recorded robotic steering. Human paths are shown for two visual conditions: i) the dark lines are the paths followed when both retinal flow and extra-retinal direction information was present (data from Wilkie and Wann [2003a]); ii) the white lines are paths when only extra-retinal direction information was present (data from Wilkie and Wann [2002]). We adjusted the parameters k_1, k_2 and b in Equation 1 and 2 to 'best-fit' the human paths. For the $\dot{\alpha}$ model $b = 1$ and $k_1 = 17.7$, whereas for the $\dot{\alpha} + \alpha$ model $b = 1$, $k_1 = 4.7$ and $k_2 = 1$. We retain these values throughout for both simulated and robot trajectories.

While both equations can match the shape of the human paths reasonably well (Figure 2) there is a particularly close match between the shape of the more direct human paths and the trajectory taken by the $\dot{\alpha} + \alpha$ model. Figure 2 shows that humans are able to steer smoothly curving paths when better quality information is available (dark lines) but when retinal flow information is not available the trajectories are more direct (white lines). It should be noted that while both sets of visual conditions result in similar levels of precision, the smoother paths tend to be biased towards understeer.

To further evaluate the Wilkie and Wann [2002] model (Equation 1) and compare to the Wilkie et al. [2008] model (Equation 2) we implemented steering control algorithms on a locomotor robot system that used a mobile camera to sample visual information and fixate the steering targets present in the world. We were particularly interested in comparing the use of $\dot{\alpha}$ with a combination of both $\dot{\alpha}$ and α. This evaluation would be very difficult to achieve by testing human participants since a source of $\dot{\alpha}$ would invariably supply α. In the robot we were able to retrieve both sources of information from the potentiometer controlling the mobile camera and then model steering using either Equation 1 or Equation 2. The scope of the research project did not extend to performing analysis of flow within the camera images, but the signals detected from the camera potentiometer do have direct counterparts in the retinal flow [Wann and Land 2000] and it has been shown that human steering is robust enough to cope with situations where the information from flow is attenuated [Wilkie and Wann 2002].

[4]It transpires that using both the target angle (α) and the rate of change of that angle ($\dot{\alpha}$) as control variables maps quite well to "quickened" second-order control systems where a weighted combination of position and velocity are used to anticipate and control future position [Jagacinski and Flach 2003].

2. GENERAL METHOD

We ran a series of six robotic trials, steering towards an eccentric target, with an underlying steering system based upon either Equation 1 or Equation 2. Objective motion and path of the robot were captured using an InterSense IS-900 Precision Motion Tracker (InterSense, Inc., Bedford, MA). The system is a hybrid acoustic-inertial 6 degree-of-freedom position and orientation tracking system. The room was configured with ultrasonic acoustic beacon transmitters fixed on a set of 'Sonistrips' mounted on the ceiling so that position and orientation could be tracked over the entire laboratory. An acousto-inertial receiver station was placed on the robot so that its position and orientation could be recorded continuously at a sampling rate of 100 Hz. This allowed us to compare the trajectories taken by the models under noisy real-world conditions when all information available for steering was retrieved from visual information received via the camera, as well as from non-visual information from the potentiometer signalling camera orientation. For evaluation, we compared real robot performance to Matlab simulations that provided 'ideal' noise-free perceptual inputs (Figure 4, left panel) or explicitly biased information sources (Figure 6, left panel).

In order for the robot to identify the steering goal the scene contained colour-coded steering targets. In the trials for Experiments 1 and 2, a single red target was placed at a distance of 2.5m offset from the robot midline by 14° (equivalent to a lateral offset of .62m). To recover target position the Image Processing Toolbox from Matlab was used to filter all colours except red from the camera image. The corners of the red object were identified and the centroid of the object was taken as its position in the image, this was converted into an angular direction. The resolution of the camera limited detection of the angular direction to within 0.2°, however we wanted to mimic human localisation which may be less precise than this. It has already been demonstrated that Equation 1 can effect steering based on a small number of discrete values (Wilkie and Wann [2003], p378) so here we limited the localisation accuracy to be within $\pm 1°$, which meant that non-zero $\dot{\alpha}$ values were limited to a minimum rate of 1° per step. When gaze was mobile the 'retinal' location was combined with the angle of the camera and then a motor command was sent to fixate the target. The centre of rotation of the mobile camera did not perfectly align with the optical centre, however, we envisaged the steering model to be robust to such sources of noise.

For trials where we wished the robot to steer via multiple targets (Experiment 3), a range of colours and filters (Red, Green, Blue, Magenta, Orange and Cyan) were used to uniquely identify targets. The estimated direction of the target was passed through one of the steering models (either Equation 1 or 2), and the output was used to specify the amount that the robot turned. There were limitations in how quickly the robot could acquire and process a single camera image. To identify the position of targets a well balanced image with distinct coloration of targets was required. Even small increases in illumination required adaptive changes to the lens aperture as well as reductions in contrast and gain of the image (equivalent to pupil contraction of the human eye with accompanying neural adaptation). Because of the time taken to acquire and process the camera image the movement of the robot had to be carried out in a series of steps. A larger step would result in fewer attempts

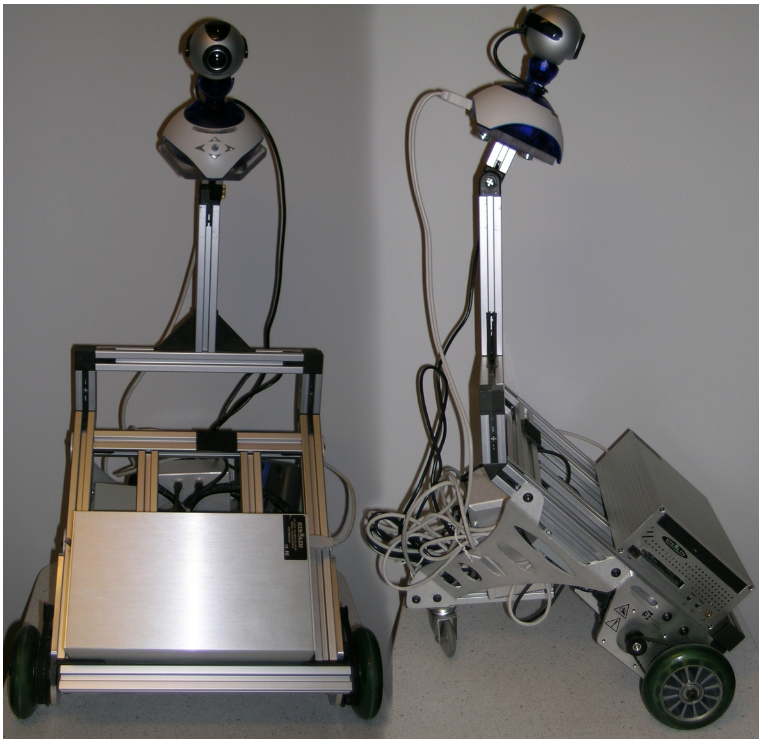

Fig. 3. The robot used was an Evolution Robotics model (ER1) which has two 4-inch diameter wheels (wheel base = 38cm) driven by stepper motors that were positioned at the front of the robot with a passive stabilising wheel trailing at the rear. The ER1 has a mount for a single static camera that we replaced with Logitech Quickcam (320×240 pixels that limited resolution to $\sim \pm 0.2°$) on a TrackerPod (Eagletron) mobile platform that could be rotated at speeds of up to $100°$/sec within a range of $160°$ pan and $110°$ tilt. The camera position (effective eye-height) was 75cm high and 19cm behind the midpoint of the front wheels, which was the centre of rotation of the robot.

to localise the target over a trial, and longer straight trajectories, potentially off target. A variety of step sizes were piloted both in simulation and with the robot and 10 steps, each step traversing 30cm, were found to produce reasonably accurate trajectories without making trials too time consuming. After each step a new snapshot of the scene was taken and the target was refixated (in mobile gaze trials). If the target could not be found within the camera image then a search routine was used to scan the camera over its full range, vertically and horizontally, until the target was re-acquired, at which point steering recommenced. Note that the speed of locomotion does not directly influence steering in this implementation, but speed information is not used by either steering model.

3. EXPERIMENT 1: STEERING TO A SINGLE TARGET WITH FIXED OR ACTIVE GAZE

The first experiment was designed to examine how well Equation 1 & 2 could steer a robot towards offset targets and how active gaze influenced steering. Firstly we

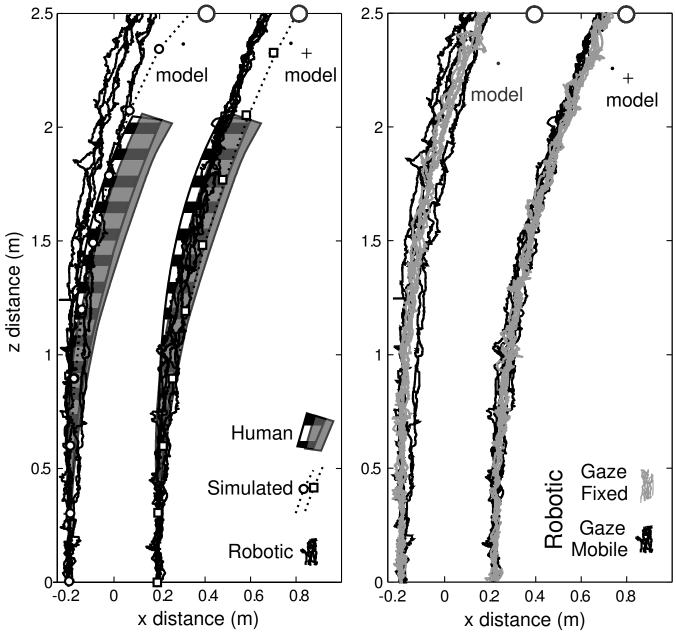

Fig. 4. Steering paths modelled using the $\dot{\alpha}$ or $\dot{\alpha} + \alpha$ steering strategies (these are offset in the figure for clarity of presentation). The left panel shows robotic paths with mobile gaze (solid black lines) as compared with software simulated paths (dashed lines). To indicate the variability of human performance, the data from Figure 2 has been represented as zones of steering that are superimposed behind the robot and model paths (solid grey region = Wilkie and Wann [2002]; black and white striped region = Wilkie and Wann [2003a]). These trajectories are truncated because trials were halted before the target was reached. It can be seen that the robotic paths tend to turn less quickly than the modelled trajectories. The $\dot{\alpha}+\alpha$ model fits well with the human data (falling centrally within the distribution of trajectories), and the robotic data falls at the edge of the distribution that tends towards understeer. The right panel shows 6 robotic paths when gaze (the camera) was mobile (black) or fixed (grey). The same parameter values were used for both fixed and mobile gaze. Both models had the damping parameter $b = 1$; in the $\dot{\alpha} + \alpha$ model $k_1 = 4.7$ and $k_2 = 1$; in the α model $k_1 = 17.7$.

compared how the robot performed in relation to an 'ideal' noise-free simulation with an equivalent number and size of steps as were used by the robot. Secondly we compared the paths taken by the robot when steering was based on Equation 1 [Wilkie and Wann 2002] versus Equation 2 [Wilkie et al. 2008]. Thirdly, we evaluated the robotic and simulated paths by comparing them to human data collected under similar circumstances. Finally, to examine whether the use of a mobile camera interfered with steering control, we compared steering when the camera on the robot was fixed or free.

3.1 Results and Discussion

It can be seen from Figure 4 (left panel) that all robotic paths fail to turn as rapidly as the ideal simulations, especially over the last 1m of trajectories. The closest passing distance of the robot was calculated for each trial (see Figure 5 for end-point

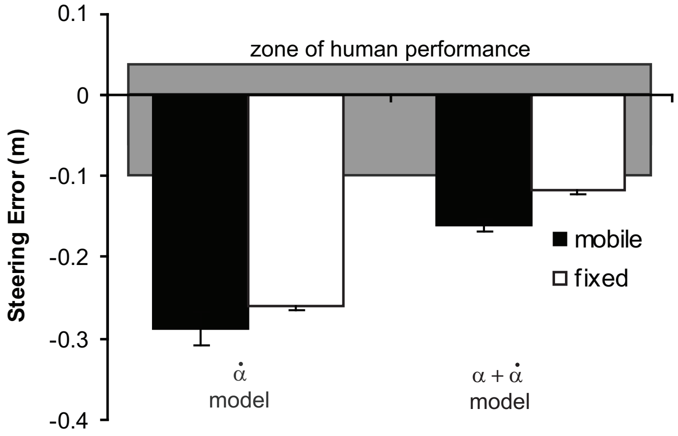

Fig. 5. The final distance of the robot from the target when steering using a model based on $\dot{\alpha}$ or $\dot{\alpha}+\alpha$ when the camera was either mobile (black) or fixed (white). A negative steering error shows understeer resulting from insufficient steering. The grey box indicates the range and maximum and minimum end errors of human observers steering as shown in Figure 2, which matches the combined striped and grey zones displayed in the left panel of Figure 4 as originally tested both within Wilkie and Wann [2002] and Wilkie and Wann [2003a].

errors), and the grand mean of errors (across all robotic conditions) indicates that understeer caused the target to be missed by $\sim .2$m. This understeer did, however, vary significantly based on which steering model was used ($F(1,3) = 240.9; p < .001$), or whether the camera was fixed or mobile ($F(1,3) = 20.6; p < .05$), but there was no interaction between these two effects ($F(1,3) = .07; p = .81$; ns). The amount of understeer for both simulations and real robotic paths reduced when using the $\dot{\alpha} + \alpha$ model in comparison to the $\dot{\alpha}$ model (Figure 4, left panel). This supports the suggestion of Wilkie et al. [2008] that α is a crucial variable for the successful control of steering.

The understeer was reduced with the $\dot{\alpha}+\alpha$ model but the bias did not completely disappear (mean error = $-.16$m). To determine whether the mobile camera was the cause of the understeer mobile camera trials were compared with trials when the camera was fixed (Figure 4, right panel). When using the $\dot{\alpha}$ model, mobile gaze (black lines, $SD = .044$) caused paths three times as variable compared to fixed gaze conditions (grey lines, $SD = .015$). This indicates that fixation of the steering target added noise to the estimate of $\dot{\alpha}$ returned from the mobile camera. Interestingly, the variability of the paths taken by the robot running under the $\dot{\alpha}+\alpha$ model were no different when gaze was mobile ($SD = .016$) or fixed ($SD = .012$) suggesting that Equation 2 is robust to the noise introduced to $\dot{\alpha}$ during gaze fixation because it contributes 3.8 times less to steering than Equation 1. The fixed gaze conditions did still result in an average error of $-.12$m. Whilst this was a small but significant ($t(3) = 5.13; p < .05$) increase in accuracy compared to free gaze, it does suggest that the predominant cause of the understeer was unrelated to mobile gaze.

We re-examined data for human participants steering to an eccentric target when only extra-retinal information was available [Wilkie and Wann 2002]. When human trajectories are scaled to map onto the target distance used in the robotic trials the end errors are small (root mean squared error = .03m). The variability of human paths, however, was relatively high, as shown by the spread of light and dark grey paths in Figure 2, which are also represented by the grey and barred zones in the left panel of Figure 4. The range of steering errors are represented in Figure 5 by the grey box (errors: -.1m to +.04m). The only robotic trials that approached the accuracy of the human data were those using the $\dot{\alpha} + \alpha$ model (Figure 5). In Figure 4 (left panel) it can be seen that the trajectories of the robot under the control of the $\dot{\alpha} + \alpha$ model seems to 'hug' the edge of the zone of human performance up until $\sim$ .5m from the target. At this point the human trajectories would continue to curve towards the target (this final section of human trajectories were not recorded but extrapolated across the last .5m) whereas the robotic paths do not turn a great deal over the final .5m. This highlights the importance of steering sufficiently at an appropriate time-point, since the stage of steering at which certain information maximally contributes does seem to vary [Wilkie and Wann 2003b]. Human trajectories tend to be more direct when retinal flow information is degraded (compare light and dark grey human trajectories in Figure 2 and see Wilkie and Wann [2003a]. The 'tighter' more direct paths result from increased steering rates when far from the target (when α is large and $\dot{\alpha}$ tends to be small). To model more accurately this aspect of human steering would require a dynamic change in the scaling of α (the k_2 term in Equation 2) and link this with the quality of retinal flow information that provides an estimate of $\dot{\alpha}$. In this case the model should increase the weighting of α when retinal flow is degraded and when the value of $\dot{\alpha}$ is small. This weighting is subtly different from previous proposals (e.g. Wilkie and Wann [2002]) that suggested the individual sources of information that contribute to an estimate of $\dot{\alpha}$ are weighted based on their variability (the β values in Equation 4, Wilkie et al. [2008]). In addition to weighting individual estimates, here we propose that human steering may well involve a flexible and dynamic combination of α and $\dot{\alpha}$.

In conclusion the results of Experiment 1 suggest that performance of a robot running under the $\dot{\alpha}+\alpha$ model with a mobile camera compares favourably with trials when the camera is fixed. Even in a noisy real-world steering situation a mobile gaze need not be a major source of error and variability and it seems, therefore, that the benefits of the mobile camera (such as increased field of view for wider sampling from the scene) may outweigh the problems.

4. EXPERIMENT 2: STEERING TO A SINGLE TARGET UNDER BIASED CONDITIONS

Experiment 1 suggests that the $\dot{\alpha} + \alpha$ steering strategy supports reasonably robust steering even in the presence of noise, although there was a consistent tendency to understeer. To determine what caused the understeer we looked at each of the processing stages within the steering model. Under ideal conditions the $\dot{\alpha} + \alpha$ steering model shows no understeer (Figure 2), so there must be some aspect of the real-world inputs to (or outputs from) the model to cause this error. The fixed gaze condition still showed significant understeer which meant the error must be caused either by a misperception of the target angle, or the production of motor output that does not match that required by the model. The centre of rotation of the camera was positioned 19cm behind the centre of rotation of the robot, which would have caused the target to be visibly offset by only 13° at the start of trials. In Experiment 2, we investigated the role of the misperception further by introducing systematic bias to the perceptual inputs to explore how this effects the accuracy of robotic steering to a fixed target. We introduced a bias into α while leaving the rate of rotation signal $\dot{\alpha}$ veridical, which in the human would be the equivalent of having an imprecise estimate of the absolute gaze-position, but a good estimate of gaze rotation. We do not know how accurately the absolute gaze position (eye in head + head on shoulders) is specified in humans. Because of the properties of the eye motor system the rate of rotation may be more accurately estimated than the absolute position [Brindley and Merton 1960]. In a human participant α and $\dot{\alpha}$ are intrinsically linked, so it is very difficult to manipulate one source without changing the other. Behavioural experiments have shown that manipulating potential sources of $\dot{\alpha}$ at a rate of $\pm 1°$/s is sufficient to directionally bias steering [Wilkie & Wann, 2002, 2005). Because altering $\dot{\alpha}$ also changes the time history of α these experiments could not distinguish between the relative contributions of these inputs during online steering. Using the robotic platform we can selectively and independently interfere with either input and examine the sensitivity of the system to noise in either source. First we simulated α bias in software to see how the system would respond with noise-free estimates (Figure 6, Left Panel). We then ran robotic trials where the registration of the visual direction of the target (α) from the camera orientation was biased (Figure 6, Right Panel).

4.1 Results and Discussion

Simulated paths using an $\dot{\alpha} + \alpha$ steering strategy were not particularly sensitive to addition of a constant 5° bias to the estimate of target direction when veridical $\dot{\alpha}$ was present (Figure 7). Purely nulling α (as shown in Figure 2) would have resulted in errors of $\pm$.24m, whereas end point errors were only in the order of .03m. We can compare the magnitude of bias with that of studies with human participants. Wilkie and Wann [2002] caused steering errors by systematically biasing the visual direction of the target by $\pm 1°$/s, and the resulting bias (after being appropriately scaled as described in Experiment 1) was .06m. While the magnitude of bias is comparable between the simulation and the human data, we are unable to determine what proportion of the human error was due to the misperception of α, and what the additional contribution $\dot{\alpha}$ played in the biased performance.

In the robotic system it was possible to add bias to the potentiometer signal that indicated the camera orientation. Adding bias to α at this point is broadly equivalent to biasing the afferent eye-position signal in the human. This bias caused larger errors than when simulating bias in α and caused errors in the order of $\sim$ .18m. For simulated trajectories the pattern of oversteer vs. understeer bias was perfectly symmetrical (Figure 6, left panel) however this was not the case for the robotic system (Figure 6, right panel). Robotic paths overshot by only .11m for +5° bias, in contrast to undershoot of .25m during the $-5°$ bias conditions (relative to unbiased trials; Figure 7). This asymmetry suggests a cause for the understeer observed in Experiment 1. The understeer could have been due to an underestimate of α (due to lack of calibration or an offset between the centre of rotation of the camera and of

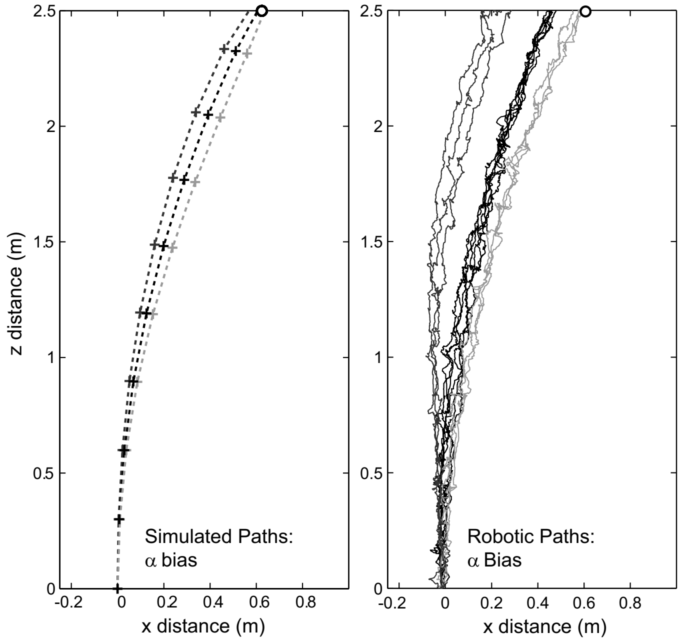

Fig. 6. Steering paths modelled using an $\dot{\alpha} + \alpha$ steering strategy when bias is added to α. Black traces indicate paths taken when α was unbiased. Dark grey paths are those paths followed when α was misperceived by 5° to the left of the actual target, and light grey paths 5° to the right. *Left Panel:* Paths simulated in software shows an effect of α bias, despite the veridical $\dot{\alpha}$ information that is available. *Right Panel:* Robotic paths when steering towards a target, with or without α bias. The bias caused small shifts in steering behaviour similar in size to simulated trials, however bias caused asymmetric shifts in the robotic paths.

the robot). If this was the explanation, however, then artificially increasing α would be expected to counteract the understeer. Furthermore if the sole explanation for robot understeer was a constant error in perceiving α (for instance an inappropriate scaling of the image), then we would still predict a symmetric bias to robotic paths when a directional bias was added to α. One possibility is that understeer was caused by erroneous motor output from the robotic wheels despite appropriate commands being transmitted from the steering model (e.g. due to factors such as wheel slippage). It may be that in the real-world an increase in gain (controlled by response rate) is required to offset situations where motor output is attenuated. The main problem with an increase in gain is the increased likelihood of steering oscillations, however the low-pass dynamics of the motor action in the real situation would dampen oscillations somewhat (though with a sequence of discrete steps the robot may be more prone to this problem).

The results of the biased conditions revealed that the implementation of robotic steering using the $\dot{\alpha} + \alpha$ model is susceptible to understeer. The asymmetry of our results shows that it is easier to cause understeer through a reduction in α than it is to promote oversteer through an increase in α. Increasing the perceived offset of the target by +5° did improve the accuracy of trajectories which suggests that

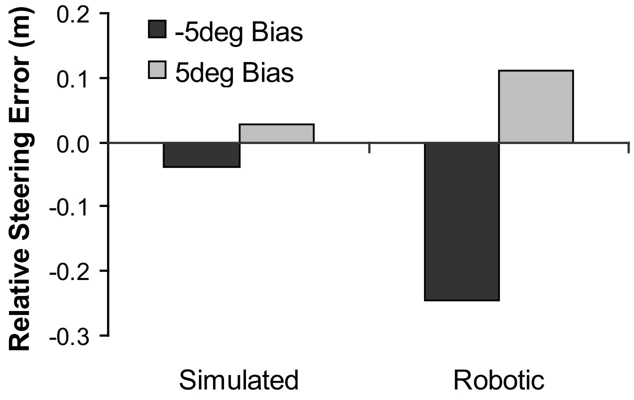

Fig. 7. The amount and direction of steering error for biased trials relative to unbiased trials. When the target is misperceived to the left of the actual position (-5deg) steering errors indicate understeer (negative values).

calibration of the camera-orientation could be an important component for ensuring accurate steering results from using this model.

5. EXPERIMENT 3: STEERING VIA MULTIPLE TARGETS

The experiments so far have examined the case of steering to a single target, but real-world trajectories are usually constrained to follow paths that are made up of many waypoints [Wilkie et al. 2008]. We can learn more about the underlying mechanisms of steering through a cluttered environment by simplifying the representation of the scene to comprise a number of distinct targets or waypoints. The models outlined in Equation 1 & 2 can be used to steer to multiple waypoints if a mechanism for shifting gaze from one target to the next is specified [Wilkie et al. 2008]. For these models the key to splining together waypoints into a single path is to smoothly shift gaze from one waypoint to the next at the appropriate time. The decision to shift gaze can be based upon the concept of immediacy, where the most immediate target captures attention up until a threshold point, at which point the next region of interest is fixated upon [Wilkie et al. 2008]. Wilkie et al. [2008] showed that Equation 2 could successfully model steering to multiple targets when gaze shifted up to 2 seconds before passing the most immediate target. Here we test this model using the robotic platform. Having the camera at a fixed height above a flat ground plane means that when a steering target is fixated, the angle of declination becomes a useful estimate of target immediacy [Rushton et al. 2002]. The vertical angle of a point on the camera image is directly related to the distance of that point in the world. In order to steer via multiple waypoints the robot merely needs to identify the most immediate target to fixate and then steer towards this. When a target is sufficiently close (in this case the immediacy specified by vertical angle exceeds a threshold) the next most immediate target is identified, and fixated, which shapes subsequent steering. In these slalom conditions the camera fixation switches before the most immediate target is reached in a manner analogous to the behaviour observed by Wilkie et al. [2008]. But because the robot moves in a series of steps the temporal aspects of gaze switching are not directly matched with human behaviour. The step-wise progression also means there is no momentum contained within the system and so we expect the robot to turn sooner than would be predicted from the simulations and the observed human trajectories presented in Wilkie et al. [2008].

5.1 Results and Discussion

The basic task for the robot was to pass over the waypoints in sequence. The actual path taken for the robot steering via multiple waypoints shows that it was successful in following the course (Figure 8). The curvature of the path matches the offset position of the targets well, and the steering is sufficiently accurate to ensure the targets stay within the wheelbase of the robot. Ideal steering would result in the target being centred under the wheel base when passed over, but we did not expect perfect performance, because of the understeer and momentum issues already noted. Not only was the vertical angle found to be useful and highly effective in this robotic situation, but under high-speed locomotor conditions the rate of change of the vertical angle should usefully specify the speed at which a fixated target is approaching, and so this could also be used to generate a richer term of immediacy. In the human system an unmediated optical variable such as optic expansion may be used to judge immediacy but this was beyond the scope of our robotic implementation.

A comparison can be made of the robot steering around multiple targets and with the task given to human participants in Wilkie et al. [2008]. This paper reported experimental work that examined steering via a series of slalom gates, laid out to be ~ 32m (or ~ 4 sec) apart. The human experiments showed that enforced early switching (looking to the next gate, when the most immediate gate was still more than 2 seconds away) resulted in greatly elevated steering errors. Wilkie et al. [2008] also ran simulated trials using Equation 2 to examine how the model coped with these fixation requirements, and it was also found that gaze switching > 2 seconds away from the target resulted in elevated errors in modelled paths. In a manner analogous to the scaled comparison in Experiment 1 & 2, we wished to determine whether the imprecision in the steering of the robot could be explained by premature gaze switching. The robot used a vertical gaze angle threshold ($\sim 74°$) which meant that gaze switching occurred $\sim .22$m from the target (approximately half-way towards each target). Whilst a direct comparison with the gaze-switch time cannot be made, when scaled to be proportional to the distances used in the human trajectories, this switch zone approximates the region where human errors start to be made.

The model outlined in Equation 2 does not explicitly represent multiple waypoints, nor does it attempt to differentially represent obstacles. Instead the model relies upon gaze fixation on unobstructed parts of the environment to guide successful steering. This method is in line with expert motorcycle instruction [Motorcycle Safety Foundation 1992] that suggests that to avoid hitting an obstacle you look to the side of it, rather than at it. To steer via a series of slalom gates or in a cluttered environment the model should 'look to the gap'. For obstacles (and vehicles) with a larger extent, it may be necessary to ensure that gaze is far enough away from the

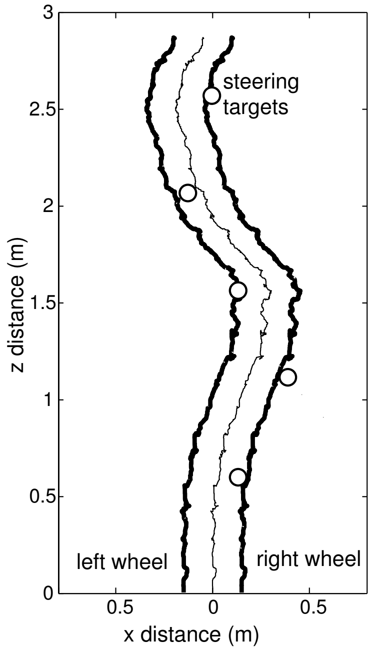

Fig. 8. Robotic steering through a course of five consecutive targets using Equation 2 (the $\dot{\alpha} + \alpha$ steering strategy). The robot used the vertical gaze angle to decide which target was the most immediate and steered towards it. When the angle of declination of the target exceeded 74° ($\sim$ .22m away) a search for the next nearest target was initiated. This routine of steering and searching continued until the robot was unable to find a target. A mobile gaze system was essential to sample extensively from the scene, as well as to provide the visual direction of targets.

obstacle to allow safe passage of the vehicle past the edge of the obstacle without collision. For certain fine manoeuvres (for example squeezing through a tight gap), this may not square with everyday experience, however these cases tend to occur at low speeds where a different pattern of eye-movements become invoked.

6. GENERAL DISCUSSION

The experiments outlined in this paper show that it is possible for a robot to use simple perceptual inputs such as α and $\dot{\alpha}$ to steer effectively to single and multiple targets. Estimates of the noise in the detection of these variables did not need to be factored into the contribution of each perceptual input in order to steer effectively, though some improvement in performance may have resulted from weighting each term based on the variability. Furthermore a mobile camera, mimicking the mobile head and eye of the human, has been shown to be an effective way of increasing field-of-view as well as supplying useful perceptual inputs (both yaw and pitch) without excessive compromise on steering accuracy. Although use of a single extra-retinal information source for estimating α and $\dot{\alpha}$ may have limited precision (compared to multiple unbiased sources available to the human) the level of understeer observed was not out of line with human performance when steering under similar impoverished conditions [Wilkie and Wann 2002]. One notable characteristic of paths taken by humans when steering to a single target is that the trajectories become more direct under limited cue conditions (see Figure 3, Middle panel of Wilkie and Wann [2002] and Figure 9 of Wilkie and Wann [2003a]). This 'tightening' of behaviour could be modelled by dynamically changing the weighting of α to increase the influence of this input when the target is far away and/or information from other cues is limited. When travelling at higher speeds the relative influence of $\dot{\alpha}$ would naturally increase, however, the human gaze fixation system usually compensates for locomotor speed and directs gaze to a zone 1-2s ahead [Land and Horwood 2005; Wilkie et al. 2008]. We did not examine how locomotor speed influences these steering models implemented within the robotic platform, but this would be a natural next step.

Though understeer was exhibited by both the steering models, this was reduced markedly for the $\alpha + \dot{\alpha}$ model compared to using $\dot{\alpha}$ alone. Understeer would most likely have been reduced by recalibrating the camera before each trial, however it is not clear how often and at what level humans perform equivalent calibrations. In terms of identifying the locomotor midline, flow field would provide highly useful information for continuously recalibrating a mobile system [Held and Bossom 1961]. For the robot, the camera was positioned so that returning the camera to zero aimed the camera approximately straight-ahead, however motor slippage may have caused this to change and may well have contributed to understeer. Calibration and learning are essential aspects of locomotor control that have not been examined in this work.

The Wilkie et al. [2008] model uses active gaze to shape a steering response through a cluttered or high risk environment. Although the steering system does not require any explicit internal representations it does provide a route for cognitive learning to influence steering performance via priority scheduling of gaze fixations (see Figure 10, Wilkie et al. [2008]). The steering model outlined in Equation 2 was effective with a relatively naïve robot, with no compensatory mechanisms or internal representations. This suggests that decoupling gaze from steering, for example when a human driver looks away from their desired path to a road sign, need not be problematic for the control model. Under high-speed high-risk conditions, where changes to the scene occur very rapidly, we would expect a very tight coupling between gaze and steering, whereas at slower speeds steering can be carried out in a series of steps with gaze uncoupled from steering for the majority of the time. One aspect that this work highlights is the need to determine when the gaze fixation becomes critical for the safe control of locomotion.

As well as inform models of human locomotor guidance, experiments such as those presented here could provide a foundation for biologically-motivated algorithms for autonomous robots in the field. Currently, many remotely piloted vehicles are teleoperated based on visual displays transmitted to the operator. This allows the operator to make navigation decisions but autonomous operation is often preferred. In some cases, such as Mars rovers, the communications delay is several minutes and the link is only available at certain times in the Martian day, precluding teleoperation. Planners must upload an entire days worth of motion plans to the rover. The ability to act autonomously is thus a key factor in the amount and type of exploration that can be accomplished [Bajracharya et al. 2008]. Similarly autonomous operation of vehicles on earth has been touted as a means to operate in hazardous environments, reduce labour costs or improve high-safety. The challenges in achieving such visions come to the fore in events such as the Darpa Grand Challenge/Urban Challenge [Darms et al. 2009]. State-of-the-art engineering has led to notable successes in these contests including autonomous navigation of difficult desert terrain and urban streets. While impressive, these systems relied on suites of sensors such as RADAR and LIDAR and moved relatively slowly and deliberately. For instance, the Boss system from Carnegie Mellon University that won the 2007 Urban Challenge detected obstacles near an intersection with a 10 s delay and negotiated them at reduced speed [Darms et al. 2009]. In contrast, humans usually avoid such obstacles easily almost entirely using vision. LIDAR and RADAR are very useful in such situations but are active (emit radiation that can be detected) and can be expensive, physically large and consume power making passive vision sensing preferable for small vehicles or stealth operation [Bansal et al. 2008]. Human vision provides a model for such systems and simple algorithms such as those tested in the present paper could allow for low-complexity implementation in software or hardware. This could prove useful for small robotic systems constrained in power and computational budget. The Mars Exploration Rovers use their mobile cameras to track and guide the robot to the target in a closed-loop fashion [Kim et al. 2009]. Interestingly computational delays require the rover to move in discrete steps and perform the tracking at the stops much as we do in the present experiments. In one mode the rover uses position control of the target position tracking (essentially controlling α). The models evaluated in the present paper could be easily adapted for algorithms on such a platform.

REFERENCES

Bajracharya, M., Maimone, M. W., and Helmick, D. 2008. The absence of position sense in the human eye. *Computer 41*, 44–50.

Bansal, M., Das, A., Kreutzer, G., Eledath, J., Kumar, R., and Sawhney, H. 2008. Obstacle detection and tracking for the urban challenge. *Proceedings of the 11th International IEEE Conference on Intelligent Transportation Systems*, 434–440.

Bicho, E. and Schoner, G. 1997. The dynamic approach to autonomous robotics demonstrated on a low-level vehicle platform. *Robotics and Autonomous Systems 21,* 1, 23–35.

Brindley, G. S. and Merton, P. A. 1960. The absence of position sense in the human eye. *Journal of Physiology 153,* 1, 127–130.

Darms, M. S., Rybski, P. E., Baker, C., and Urmson, C. 2009. Obstacle detection and tracking for the urban challenge. *IEEE Transactions on Intelligent Transportation Systems 10,* 3, 475–485.

Fajen, B. R. and Warren, W. H. 2003. Behavioral dynamics of steering, obstacle avoidance, and route selection. *Journal of Experimental Psychology: Human Perception and Performance 29,* 2, 343–62.

Fajen, B. R. and Warren, W. H. 2004. Visual guidance of intercepting a moving target on foot. *Perception 33,* 6, 689–715.

Ghose, K. and Moss, C. 2006. Steering by hearing: A bat's acoustic gaze is linked to its flight motor output by a delayed, adaptive linear law. *JOURNAL OF NEUROSCIENCE 26,* 6, 1704–1710.

Gibson, J. J. 1958. Visually controlled locomotion and visual orientation in animals. *Br J Psychol 49,* 3, 182–94.

Gnadt, J. W., Bracewell, R. M., and Andersen, R. A. 1991. Sensorimotor transformation during eye movements to remembered visual targets. *Vision Research 31,* 4, 693.

Held, R. and Bossom, J. 1961. Neonatal deprivation and adult rearrangement: complementary techniques for analyzing plastic sensory-motor coordinations. *J Comp Physiol Psychol 54*, 33–7.

Jagacinski, R. and Flach, J. 2003. *Control theory for humans: Quantitative approaches to modeling performance*. Erlbaum, Mahwah, New Jersey.

Kim, N. and Turvey, M. 1999. Eye movements and a rule for perceiving direction of heading. *Ecological Psychology 11*, 233–248.

Kim, W. S., Nesnas, I. A., Bajracharya, M., Madison, R., Ansar, A. I., Steele, R. D., Biesiadecki, J. J., and Ali, K. S. 2009. Targeted driving using visual tracking on mars: From research to flight. *Journal of Field Robotics 26*, 243–263.

Land, M. 1969. Movements of retinae of jumping spiders (salticidae - dendryphantinae) in response to visual stimuli. *Journal of Experimental Biology 51,* 2, 471–493.

Land, M. 1971. Orientation by jumping spiders in absence of visual feedback. *Journal of Experimental Biology 54,* 1, 119–139.

Land, M. F. 1992. Predictable eye-head coordination during driving. *Nature 359,* 6393, 318–20. 0028-0836 (Print) Journal Article.

Land, M. F. 2004. The coordination of rotations of the eyes, head and trunk in saccadic turns produced in natural situations. *Exp Brain Res 159,* 2, 151–60. 0014-4819 (Print) Journal Article.

Land, M. F. and Furneaux, S. 1997. The knowledge base of the oculomotor system. *Philos Trans R Soc Lond B Biol Sci 352,* 1358, 1231–9.

Land, M. F. and Hayhoe, M. 2001. In what ways do eye movements contribute to everyday activities? *Vision Research 41,* 25-26, 3559–65.

Land, M. F. and Horwood, J. 2005. Different retina-lamina projections in mosquitoes with fused and open rhabdoms. *J Comp Physiol A Neuroethol Sens Neural Behav Physiol 191,* 7, 639–47. 0340-7594 (Print) Journal Article.

Land, M. F. and Lee, D. N. 1994. Where we look when we steer. *Nature 369,* 6483, 742–4.

Land, M. F. and Tatler, B. W. 2001. Steering with the head. the visual strategy of a racing driver. *Current Biology 11,* 15, 1215–20.

Lee, D. 1998. Guiding movement by coupling taus. *Ecological Psychology 10,* 3-4, 221–250.

Llewewllyn, K. R. 1971. Visual Guidance of Locomotion. *Journal of Experimental Psychology 91,* 2, 245–261.

Motorcycle Safety Foundation, I. 1992. *Motorcycling excellence: Skills, knowledge and strategies for riding right*. Whitehorse Press, North Conway NH.

Murray, D. W., Reid, I. D., and Davison, A. J. 1997. Steering without representation with the use of active fixation. *Perception 26,* 12, 1519–28.

Robertshaw, K. D. and Wilkie, R. M. 2008. Does gaze influence steering around a bend? *Journal of Vision 8,* 4, 18.1–13.

Royden, C. S., Banks, M. S., and Crowell, J. A. 1992. The perception of heading during eye movements. *Nature 360,* 6404, 583–5.

Rushton, S., Wen, J., and Allison, R. 2002. Egocentric direction and the visual guidance of robot locomotion background, theory and implementation. *Biologically Motivated Computer Vision, Proceedings 2525*, 576–591.

Rushton, S. K., Harris, J. M., Lloyd, M. R., and Wann, J. P. 1998. Guidance of locomotion on foot uses perceived target location rather than optic flow. *Curr Biol 8,* 21, 1191–4.

Salvucci, D. D. and Gray, R. 2004. A two-point visual control model of steering. *Perception 33,* 10, 1233–1248.

Schoner, G., Dose, M., and Engels, C. 1995. Dynamics of behavior: Theory and applications for autonomous robot architectures. *Robotics and Autonomous Systems 16,* 2-4, 213–245.

Tucker, V. 2000. The deep fovea, sideways vision and spiral flight paths in raptors. *Journal of Experimental Biology 203,* 24, 3745–3754.

Wann, J. P. and Ibrahim, S. F. 1992. Does limb proprioception drift? *Exp Brain Res 91,* 1, 162–6.

WANN, J. P. AND LAND, M. 2000. Steering with or without the flow: is the retrieval of heading necessary? *Trends in Cognitive Sciences 4,* 8, 319–324.

WANN, J. P. AND SWAPP, D. K. 2000. Why you should look where you are going. *Nature Neuroscience 3,* 7, 647–648.

WARREN, W. H., J., KAY, B. A., ZOSH, W. D., DUCHON, A. P., AND SAHUC, S. 2001. Optic flow is used to control human walking. *Nat Neurosci 4,* 2, 213–6.

WHITE, J. M., SPARKS, D. L., AND STANFORD, T. R. 1994. Saccades to remembered target locations: an analysis of systematic and variable errors. *Vision Research 34,* 1, 79.

WILKIE, R. AND WANN, J. 2003a. Controlling steering and judging heading: retinal flow, visual direction, and extraretinal information. *Journal of Experimental Psychology: Human Perception and Performance 29,* 2, 363–78.

WILKIE, R. AND WANN, J. 2003b. The stages of steering. *Journal of Vision 3,* 9, 551a.

WILKIE, R. AND WANN, J. 2006. Judgements of path not heading guide steering. *Journal of Experimental Psychology: Human Perception and Performance 32,* 1, 88–96.

WILKIE, R. M. AND WANN, J. P. 2002. Driving as night falls: the contribution of retinal flow and visual direction to the control of steering. *Current Biology 12,* 23, 2014–7.

WILKIE, R. M. AND WANN, J. P. 2003c. Eye-movements aid the control of locomotion. *Journal of Vision 3,* 11, 677–84.

WILKIE, R. M. AND WANN, J. P. 2005. The role of visual and nonvisual information in the control of locomotion. *Journal of Experimental Psychology: Human Perception and Performance 31,* 5, 901–11.

WILKIE, R. M., WANN, J. P., AND ALLISON, R. S. 2008. Active gaze, visual look-ahead and locomotor control. *Journal of Experimental Psychology: Human Perception and Performance 34*, 1150–1164.

...