Egocentric Direction and the Visual Guidance of Robot Locomotion Background, Theory and Implementation

Simon K. Rushton[1,2], Jia Wen[3], and Robert S. Allison[1,3]

[1]Centre for Vision Research, [2]Department of Psychology, [3]Department of Computer Science, York University,
4700 Keele Street, North York, ON, M3J 1P3, Canada.
rushton@hpl.cvr.yorku.ca; allison@cs.yorku.ca;
http://www.cs.yorku.ca/~percept/robot.html

Abstract and Overview. In this paper we describe the motivation, design and implementation of a system to visually guide a locomoting robot towards a target and around obstacles. The work was inspired by a recent suggestion that walking humans rely on perceived egocentric direction rather than optic flow to guide locomotion to a target. We briefly summarise the human experimental work and then illustrate how direction based heuristics can be used in the visual guidance of locomotion. We also identify perceptual variables that could be used in the detection of obstacles and a control law for the regulation of obstacle avoidance. We describe simulations that demonstrate the utility of the approach and the implementation of these control laws on a Nomad mobile robot. We conclude that our simple biologically inspired solution produces robust behaviour and proves a very promising approach.

1 Theoretical Background: Human Locomotion and Egocentric Direction

For the past 50 years it has been assumed that humans rely on optic flow for the visual guidance of locomotion. This assumption has underpinned psychophysical studies, neurophysiology, imaging and computational modelling (see [1] for a review). Recently this assumption has been challenged.

Rushton et al [2] reported an experimental result seemingly at odds with the use of optic flow. Rushton et al proposed instead a simple heuristic that better described the behaviour they observed. The proposal is that visual guidance of locomotion is achieved by keeping a target at a fixed direction, or *eccentricity*, relative to the body, rather than regulating behaviour so as to maintain a certain pattern of flow on the retina (the optic flow solution). In short, if the current direction of a target object is known, and the observer walks so as to keep the direction constant then they will reach the target. If the target is kept straight-ahead then a straight-line course to the target will result. If the target is maintained at some other direction then the path will be an equi-angular spiral.

The finding of Rushton et al has now been replicated by many others [3-7], and a concise summary of the original study is provided below. In a later section we

H.H. Bülthoff et al. (Eds.): BMCV 2002, LNCS 2525, pp. 576–591, 2002.

illustrate how this simple heuristic can be extended into a general model of the visual guidance of locomotion. We then describe a control law to avoid obstacles.

1.1 The Prism Study, Rushton et al. (1998)

The Rushton *et al.* [2] study involved observers wearing prism glasses. Observers were asked to walk briskly towards a target held out by an experimenter positioned about 10m to 15m away. The glasses contained either paired base-left or base-right wedge prisms. Prisms deflect the image and so shifted the perceived location of objects (relative to the body) approximately 15° to the right or left. Wearing prism glasses had a dramatic effect on the trajectory taken by observers when asked to walk towards the target. Observers veered whilst attempting to walk 'straight towards' the target. A typical veering trajectory is shown in the left panel of Figure 1.

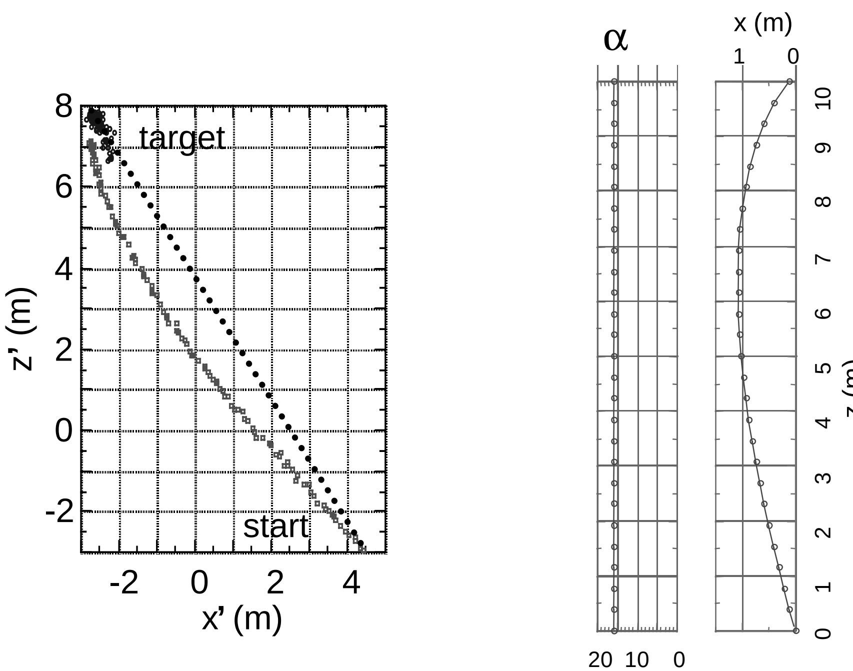

Fig. 1. *Left panel:* A representative trajectory of an observer, wearing a pair of wedge prisms that deflect right, approaching a target. The plot shows raw digitised data with axes x' and z' showing distances in world co-ordinates. *Right panel:* Simulated trajectory and direction error when wearing prisms by a model using target direction. *Right* Plan view of the predicted trajectory of a prism-wearing participant walking in the perceived direction of the target (which is offset from actual position by the simulated 16° angular deflection of the prisms). x and z are distances parallel and perpendicular, respectively, to the starting position of the participant (facing along the z-axis). *left* angle, α, between the instantaneous direction of the target and the direction of locomotion (tangent to the curve), which remains constant.

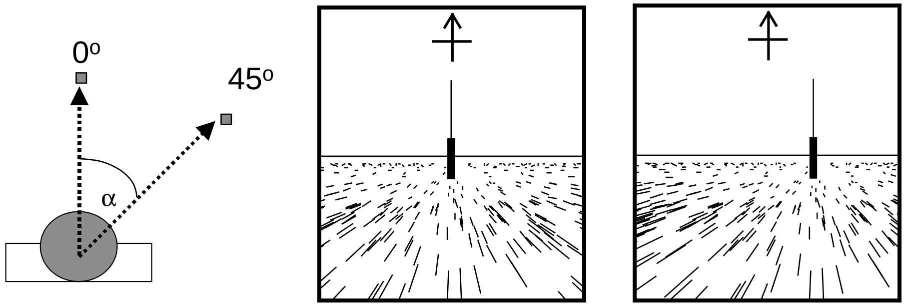

Fig. 2. *Left panel:* Egocentric directions, 'eccentricity', α, measured angle in cardinal plane. *Right panel:* Flow-field during forward translation (magnitude indicates image speed) toward target tower (solid black rectangle) at 16m. Thin vertical line indicates direction of travel. Arrow indicates egocentric straight ahead. *Left:* normal view, the 'focus of expansion' (FoE) is coincident with the tower, which indicates the observer is travelling directly towards tower. Arrow above tower indicates the perceived 'straight-ahead' direction, note it coincides with the tower. *Right:* displacement of whole image by prism. Note FoE is still directly over tower, thus flow indicates the observer is travelling directly towards tower. However, the perceived straight-ahead (denoted by the arrow above) no longer coincides with the tower.

1.1.1 A Flow Explanation?

Can use of optic flow account for such a trajectory? Flow based strategies rely on keeping the flow specified direction of heading (DoH) and the target coincident. More generally, they are concerned with relative positions or patterns within the flow-field. As can be seen from figure 2, although prisms displace the scene and change the perceived location of objects, the critical relations in the flow field are not perturbed. Specifically, the relative positions of the DoH and the target remain unchanged. Therefore, perception of direction of locomotion should remain unchanged and veridical if flow is used, and observers should end up on a straight trajectory towards to the target. The left panel of Figure 1 shows a markedly curved trajectory indicating that the observer did not use a DoH-target strategy. A model based on using the flow field specified DoH is therefore incompatible with the experimental results.

1.1.2 Egocentric Direction Account

A simple model, the *perceived direction* model [2], is compatible with the data. The model predicts that observers take a curved path because they attempt to keep the target perceptually straight-ahead of them. They veer because prisms change the perceived target direction. When wearing prisms, the perceived position of the whole scene, relative to the observer's body, is changed by the angular deflection of the prism – so if the prisms shifts the scene by 15° to the left, an object at 0° relative to the trunk will be seen at approximately 15° to the left. Thus, keeping the target *perceptually* straight-ahead requires the observer to keep the target at a fixed eccentricity (relative to the body) of approximately 15° to the right of the trunk mid-line. If this strategy is used then it should lead to a veering trajectory to the target. The trajectories walked by observers were very similar to those predicted by this simple *perceived-direction* model (compare the left and right panels of figure 1).

1.1.3 Recent Results

The direction of an object relative to the body trunk can be determined from a variety of sources of information. Classically it is assumed that trunk-centric direction is determined by combining non-visual information about the orientation of the eye and head (derived from 'extra-retinal information' – sensors that measure orientation, or copies of motor commands) with retinal target location. However it can be demonstrated theoretically that the eye orientation or head-centric target direction could be determined directly from the binocular disparity field. Visual motion, or slip, of a target as a result of a gaze or body movement could also be used in the determination of trunk-centric direction. Recent findings [3-7] on the visual guidance of locomotion can be interpreted as supporting this less simplistic model of the human perception of egocentric directions (see TICS note [8]). The use of disparity and motion information in refining or calibrating estimation of egocentric direction might usefully be revisited should it be desirable to implement the following algorithms on a robot with a mobile gaze system or stereoscopic vision.

2 The Test Rig

During the next section we complement theoretical predictions with empirical results so we first describe some details of our test rig.

The testing and development proceeded in two parallel stages. In the first stage, locomotion control algorithms were developed through simulations (using Matlab 6). During this stage, evaluation was done by both informal case-based testing and objective measures (derived by analysing the results of batches of simulations). In the second stage, simple image processing and motor output modules were added as front and backends to the locomotion module and experiments were performed on the actual robot. Unfortunately space constraints limited the complexity and length of the robot trajectories.

2.1 Robot

A Nomad Super Scout (Nomadic Technologies Inc.) robot was used for testing. Although the Nomad had a Pentium class CPU, the robot was tele-operated for convenience. A wireless network link was used to send motor commands to the robot. We used only two commands to drive the robot: `rotate()` and `move()`. Therefore, the solutions involved discrete steps. Simulations show that substituting radius of curvature and speed should produce approximately the same trajectories but represent a slightly more elegant solution.

An NTSC resolution camera was connected, via a cable to the image capture card of a Sun Blade 100 workstation. Images were captured on demand, at a resolution of 640x480 pixels. The horizontal field of view of the camera was approximately 60°, and the camera was pitched downwards by approximately 20°.

2.2 Visual Scene

Because the focus of this work was on visual guidance of locomotion and control laws we simplified the image processing requirements. Targets and obstacles were colour-coded, the former being coloured red, the latter blue.

3 From a Single Case to a General Account

The experiments described in section 1 are concerned with a single task, visually guiding locomotion to a static target. Rushton & Harris [9] explored the theoretical sufficiency of the egocentric direction proposal by attempting to extend it to a broader range of locomotion actions. Their ideas are the basis for the algorithms and implementation.

3.1 A List of Fundamental Locomotor Actions

Several authors have attempted to enumerate a list of fundamental locomotor behaviours [10-12]. The most important of these can be summarised as follows: 1) intercepting static and moving targets, 2) following a path, and 3) avoiding obstacles. Below we examine how the first two behaviours could be implemented within an egocentric direction framework. In section 4 we describe in detail our approach to obstacle avoidance.

3.2 Intercepting Static and Moving Targets

Interception of a target is achieved if during locomotion the target is (i) kept at a fixed direction relative to the robot, and (ii) the target gets closer on each step. The direction at which the target is kept will determine the exact trajectory taken. The resultant trajectories are low angle equi-angular spirals. The top left panel of figure 3 illustrates a family of trajectories that intercept a static target. If the target is moving then the same constant direction strategy works. The top middle panel demonstrates a family of constant direction trajectories that intercept a target moving with a constant velocity. The top right panel shows interception with an accelerating target. The lower panels show three constant angle trajectories taken by our robot to a static target.

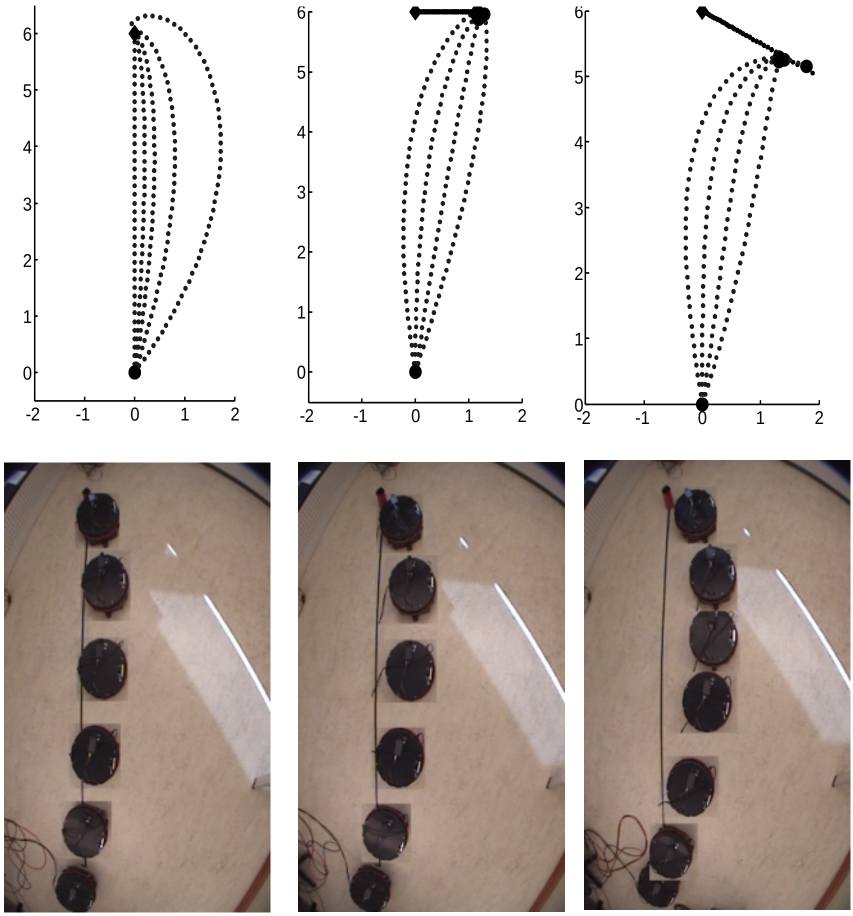

Fig. 3. *Upper Panels:* All panels display a plan view, with the robot starting at (0,0). *Left:* plan view of trajectories that would result from holding a target at a fixed eccentricity, α, of 0°, 5°, 10°, 20° and 40° (from left to right). Robot starts at (0,0), target is at (0, 6). Holding the target 'straight ahead', i.e. at 0° would produce a straight trajectory leading directly to the target. Any other trajectory based upon holding the target at an eccentricity other than zero results in the robot 'veering' to one side before finally reaching the target. *Middle:* Intercepting a moving target. Target starts at (0, 6), and moves rightwards, robot starts at (0,0). Four fixed eccentricity trajectories shown, -10°, 0°, 10°, 20°. *Right:* Intercepting an accelerating target. Target starts at (0, 40), and moves rightwards and downwards with increasing speed (constant acceleration), robot starts at (0,0). Fixed eccentricity trajectories shown are -10°, 0°, 10°, 20°. *Lower Panels:* Fixed eccentricity approaches to a target. Plan view of robot, travelling from bottom to top of the image. *Left*: eccentricity of 6°. *Middle*: eccentricity of 12°. *Right*: eccentricity of 18°.

3.3 Calibration of Straight-Ahead (0^o) through Target Drift, or Visually Guiding Uncalibrated Systems

The algorithm described so far relies on a calibrated system. If straight-ahead (0^o) is not known, or has drifted then an observer or robot could not take a straight-line course to a target if they wished to. How might the system be calibrated? Held [13] proposed that the focus of expansion of the optic flow field sampled by an observer could be used for calibration (humans normally walk either straight, or on a curving trajectory, seldom do they walk diagonally, therefore if the position of the focus of expansion was averaged over several minutes of locomotion it would provide a good estimate of straight-ahead or 0^o). A faster, on-line alternative, more in keeping with the proposal outlined so far would be to use target drift.

Llewellyn [14] proposed a heuristic for visual guidance of locomotion that is related to the constant eccentricity heuristic described above. Llewellyn suggested that an observer could reach a target by simply cancelling target drift, that is the visual movement of the target. So if a target drifts 1^o to the left after one step then if the observer rotates left by 1^o (so returning the target to its original eccentricity) and takes another step they will eventually reach their target. It should be obvious that the course will be the same equi-angular spirals produced by the constant eccentricity strategy. The use of a motion signal instead of a direction signal has one disadvantage and one related advantage. First off it is not possible to explicitly choose a trajectory. A sharply curving 50^o equi-angular trajectory cannot be selected in advance or distinguished from a 0^o trajectory. However the problem of selecting a 0^o trajectory can be avoided. During a non-zero approach, the target will drift on each step. By "overcompensating" for this drift the trajectory can be straightened into a 0^o trajectory. So if the target drifts 1^o left on a given step, if instead of rotating 1^o left to compensate (100% compensation) the observer rotates 2^o left (200% compensation) then they will end up reducing the eccentricity of their trajectory, and thus straightening the trajectory, until it reaches zero. This is illustrated in the left panel of figure 4 below.

The right and middle panels of figure 4 illustrate robot trajectories. It can be seen that in the set up here, the system calibrates rapidly. If target drift is to be used for calibration then once the target drift has settled below a preset limit, straight-ahead can be taken from the windowed average of target image position.

3.4 Path Following

Many models have been proposed to account for steering a car round a bend. Land & Lee [15] proposed that the curvature of a bend can be determined using a function of the direction of the 'tangent' (or 'reversal') point and its distance. The curvature can then be used to set the steering angle. Murray et al [16] proposed a similar rule for determining curvature and demonstrated that it could be used to guide a robot around a track.

Here we propose a simpler (but related) solution. Rather than estimate the curvature of the bend it is sufficient simply to keep a portion of the road a fixed distance ahead (and distance can be determined simply from either relative or absolute height in the visual field), or the tangent point, at a constant direction.

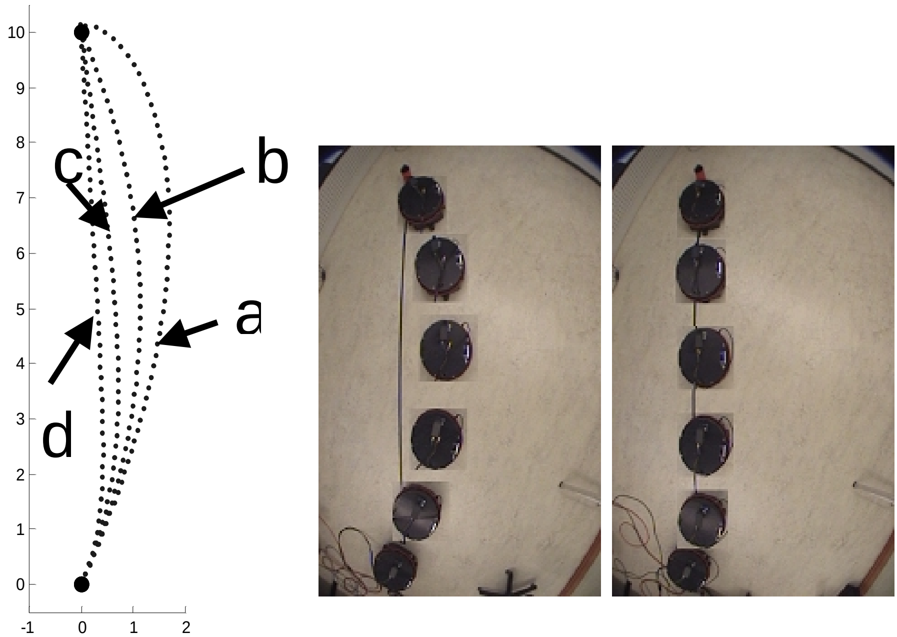

Fig. 4. *Left Panel:* Robot heads towards the target (0, 10). Initial target-heading direction, or eccentricity, is 25°. Trajectory **a** shows the course taken by cancelling target drift on each step (100% compensation) resulting in a constant direction trajectory. Trajectory **b** shows the course taken when the observer "over-compensates" for the target drift by a factor of 2 (200% compensation). Trajectory **c** "over-compensation" is 400%, trajectory **d** is 800%. *Middle and Right Panels:* Overhead view of robot travels from right to left. Initial heading angle is approximately 18 °. *Middle Panel:* Compensation of 200%. *Right Panel:* Compensation of 600%

In figure 5, in the panel A, the inside edge of the road a fixed distance ahead is kept at a constant direction. In panel B the outside edge is used. Logically if there was a centre line then this could be used instead. One disadvantage of using a fixed distance ahead strategy is that it only works if the observer does not use a portion of the road too far ahead. The maximum distance ahead is proportional to the radius of curvature of the bend. A strategy that automatically compensates for the curvature of the bend is to use the tangent point. The result of such a strategy is shown in fig. 5C.

An intuitive solution would be to scale the distance of the road edge that is used to control steering (control point) as a function of speed – to look a fixed time ahead. The distance of the control point could then be bound by the distance of the tangent point. If the control point would lie beyond the tangent point the observer could either use the tangent point (and so look less time ahead making visual control more difficult), or slow down so as to bring the control point back to the tangent point and the look-ahead distance back to an optimal value.

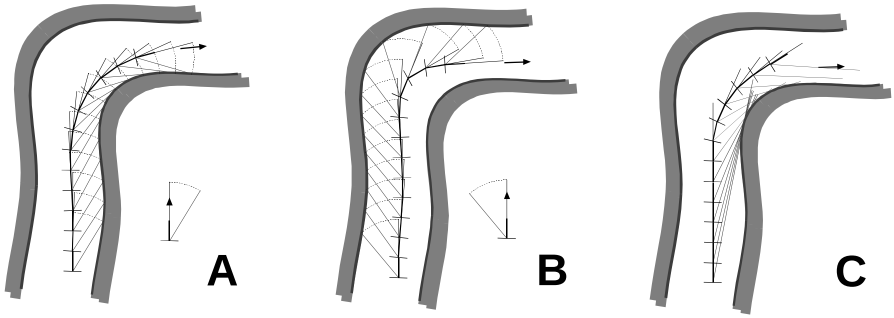

Fig. 5. A. Fixed distance, fixed direction inside of bend. **B**. Fixed distance, fixed direction outside of bend. **C**. Tangent point, fixed direction (30° threshold)

The maintain-eccentricity solution described above is unlikely to be used in isolation. If we consider a classic model of steering by Donges [17] we note that it relies on two control variables, a far point and lateral lane position. The lateral position is monitored to ensure that the car has not drifted. It is likely that walking observers would, and moving robots should, also respond to changes in lateral position. However it might be sufficient to monitor lateral position only intermittently and to correct any drift with an approximate ballistic/open-loop change in lateral position. Land [18] and Wann & Land [19] provide useful reviews of behavioural data and models and provide some alternative perceptual control solutions.

4 Detecting and Avoiding Obstacles

4.1 Detecting Obstacles

If during approach an object remains at a constant direction as you move then you are on a collision course. If the obstacles and the robot had zero horizontal extent then this would be sufficient for obstacle avoidance.

An observer could search for any imminent collisions, and if a collision is detected, change the eccentricity of their approach to the target. So they might go from a 0° eccentricity (straight) trajectory to a 10° trajectory. Note, this does not require that the observer change their immediate goal and navigate around an obstacle, but rather that they simply change the parameters of their current target approach trajectory. So even if an observer ended up changing direction by a large angle (eg 40°) to avoid a target they would still be on a course to their target. However, both the observer or robot, and obstacle have some non-zero horizontal extent, so identifying only objects that remain at a constant direction as obstacles is not sufficient.

So how might an obstacle be detected? It would be possible to use a ratio of the x (lateral) and z (depth) distances of an obstacle to generate a change of trajectory response.

Another solution would be to use other variables to which the human visual system is sensitive. One such variable is crossing distance [20-21]. Crossing distance is the lateral distance measured in the Cyclopean plane that passes through the eyes, at which a projectile will pass. It was first proposed by Bootsma [20], who showed that:

$$\frac{XDIST}{2R} = \frac{\dot{\alpha}}{\dot{\theta}} \tag{1}$$

where *XDIST* is the crossing distance, *R* the object radius, $\dot{\alpha}$ is the rate of changing direction, and $\dot{\theta}$ is the rate of changing size.

4.1.1 Calculation of XDIST

A variant of equation 1 based upon changing binocular disparity instead of changing size is:

$$\frac{XDIST}{I} = \frac{\dot{\alpha}}{\dot{\phi}} \tag{2}$$

where I is the inter-camera separation, and $\dot{\phi}$ is changing disparity [22].

We use neither of these variants, the problem with the first being that it returns *XDIST* as a multiple of obstacle width and therefore requires prior knowledge of obstacle dimensions. The second returns a more useful measure, *XDIST,* as a multiple of inter-ocular or inter-camera distance. However, use of this formulation would require a binocular viewing system and associated stereo matching algorithms.

Instead while acknowledging the utility of these sources of *XDIST* information, we elected to keep our implementation (hardware and software) as simple as possible and instead take advantage of one of the constraints in our setup. The robot always moves over a flat ground plane, and obstacles rest on the ground plane. Therefore we can modify the second *XDIST* equation and use change in height-in-the-image, $\dot{\rho}$, in place of change of disparity. We elected to use the changing direction of the inside or closest edge, $\dot{\beta}$:

$$\frac{XDIST}{H} = \frac{\dot{\beta}}{\dot{\rho}} \tag{3}$$

where H is the known height of the camera. In our system the camera is pitched downwards so we must multiply $\dot{\rho}$ by sec(pitch).

The inside edge of the obstacle is the edge that is closest to the straight-ahead direction of the camera, or the edge that is moving most slowly. The use of the nearest edge simplifies the algorithm as it avoids the need to explicitly calculate and take into account object width. Use of object edge also fits well with a recent suggestion that the human perceptuo-motor system works on the position of object edges rather than the object's centroid and width [23].

4.1.2 Detecting Collision

We can define a safe crossing distance, SAFEDIST, which is 'body-scaled' [24] to the robot. If the lateral distance at which an obstacle will pass the robot, XDIST is less than a minimum distance (SAFEDIST) then it is on a collision course. Therefore an observer or robot can continuously examine the obstacles in the scene and look for any on a collision course.

4.1.3 Knowing How Quickly to Turn

Once an obstacle is detected, how quickly does a change in trajectory need to occur? We can define "temporal distance" as the amount of time remaining before we will collide with the robot. The temporal distance can then be used to assess how urgent a change of course is.

4.1.4 Calculation of Temporal Distance

TTC (time to contact), the time remaining before an obstacle collides with the eye or camera can be determined (to a first order approximation) indirectly by the ratio of the obstacle distance to the obstacle's closing speed. It can be determined "directly" [25] from $\theta / \dot{\theta}$ where θ is the size of the image of the obstacle at the eye or camera. It is also given by $\phi / \dot{\phi}$ where ϕ is the binocular subtense of the obstacle viewed from eyes or a stereo-head. Rushton & Wann [26] recently proposed a computational solution that combines both these estimates and demonstrates robust behaviour in the case of cue-drop out, cue-conflict and optimal weighting of information from size and disparity as a function of object size.

$$TTC = (\theta + \phi) / (\dot{\theta} + \dot{\phi}) \tag{4}$$

Our implementation involved only a single camera so eq 4. cannot be used, however, the principle behind it can be used to optimise estimation of TTC from the monocular information:

$$TTC = (\theta_h + \theta_v) / (\dot{\theta}_h + \dot{\theta}_v) \tag{5}$$

where θ_h is the horizontal extent of the image of the obstacle, θ_v is the vertical extent.

The above equation will lead to the expansion of horizontal image size having the most influence with a short and wide obstacle, and vertical extent with a thin and tall obstacle. Therefore, it makes optimal use of the information available without the additional computational cost or the difficulty of determining priors or variance on the fly associated with Bayesian approaches.

4.2 Change of Path Equation

Our first constraint in deriving an obstacle avoidance algorithm is that we shouldn't lose track of our target, therefore any avoidance behaviour taken will simply change the parameters of the approach to the target rather than spawn an obstacle avoidance sub-goal. Therefore we only change the *eccentricity* parameter in our target approach algorithm.

If we set a safe crossing distance, SAFEDIST, then when XDIST < SAFEDIST, we can calculate a change in eccentricity or a turnrate, ϖ , to avoid the obstacle while proceeding to the target:

$$\varpi = k.\frac{(SAFEDIST - XDIST)}{TTC} \tag{6}$$

This equation is of the same form as that proposed by Peper et al [27] to describe projectile interception by human observers. Recent work indicates that human catching behaviour may be better described by models that include a prediction of lateral position [28], however work on human locomotion suggests that prediction is not used [2]. Therefore for now, we do not include a predictive term for future XDIST.

We change the eccentricity of the approach to the target as follows:

$$\alpha_{t+1} = \alpha_t + \varpi \tag{7}$$

Where α_t is the eccentricity of approach to the target at time t. Equation 6 leads to a response to obstacles as shown in figure 6.

On the basis of our interpretation of behavioural data [29] we only modify the eccentricity of approach on the basis of the closest obstacle. This decision contrasts with decisions made by others to include other obstacles too [30-31]. Closest obstacle could be decided on the basis of distance, TTC (time before collision with the observation point) or TTP (time before the obstacle will pass the robot). Distance is less useful when the environment is not static, therefore we used TTC.

4.3 Left vs. Right Decision Rules

Consider a robot approaching an obstacle. The robot could change the eccentricity of its target approach so as to pass to the left or the right of the obstacle. How should it decide?

- Our first rule says that it should take the route that requires the smallest change in eccentricity of approach.
- Our second rule says it should take the route that reduces the eccentricity of the current approach (gets it closest to zero or straight-ahead).

When the change in eccentricity associated with turning left vs. right is approximately the same we defer to the second rule. By varying the definition of “approximately the same” we can trade off the two rules.

5 Absolute Performance

We performed extensive testing through simulation. The system demonstrated robust performance in a range of environments (size, shape, number of obstacles, moving or stationary). It is difficult to capture the absolute performance of a system in a few

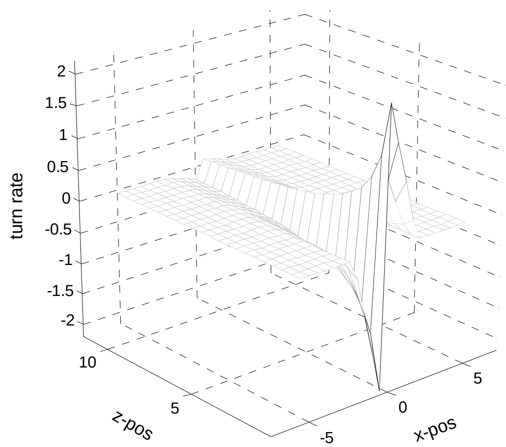

Fig. 6. Graphical representation of equation 6. X is lateral position, Z is distance in depth. Turn rate as a function of (x, z) position of an obstacle relative to the observer/robot.

statistics, as it is necessary to qualify results with an extensive description of the testing conditions, and to rigorously document any implementation specific functions, defaults or assumptions. We intend to report such details and results elsewhere.

5.1 Limitations of the Algorithm

The algorithm was designed to steer around obstacles, thus it does not deal with situations such as dead-ends or obstacles placed too close at beginning. We envisage that other processes would recognise and deal with these situations.

One implementation specific cause of failure was an inability to correctly segment obstacles when they overlapped in the projective view. This is not due to a shortcoming of our approach, but rather simply results from implementation decisions made (all obstacles were rendered flat-shaded in pure blue), and would not normally be a problem when features such as distance, texture, colour and so on could be used in the segmentation process.

One solution to this problem is to add predictive object tracking. Such an addition proves very useful as it removes other errors associated with mistaking object identity (and thus incorrectly calculating TTC etc). Every different implementation will bring its own specific problems, but it appears that predictive object tracking might be solve a whole raft of implementation specific problems and so it might be judicious to include an object tracking system by default.

A general problem that applies to any implementation of the algorithms described in this paper or any others is the problem of spatial scale. Our system can deal well with obstacles of varying width, but would run into problems when obstacles become so tall that the top of the obstacle falls out of the field of view. It is instructive consider what a human might do under such circumstances. They would be likely to switch to determining TTC, XDIST etc from local texture patches. In other words they use a different spatial scale. Trying to reformulate algorithms so that they are spatial scale invariant is an important and interesting problem.

6 Relative Performance

Although we have not done any formal testing of relative performance we can make the following comparisons to alternative approaches. Compared to potential field [31] and dynamical system [30] approaches, we have the following advantages: (i) simplicity; (ii) ability to deal with obstacles of varying size and shape (unlike the current formulation of the dynamical model [30]); (iii) based upon human perceptual variables.

Duchon et al [32] review the literature on the use of optic flow for obstacle avoidance and describe their own approach. From our understanding, the solution we propose is markedly simpler than the optic flow solutions, not least because we do not need to do optic flow processing over the visual field.

7 Some Examples

The following figures show some sample trajectories from our current model.

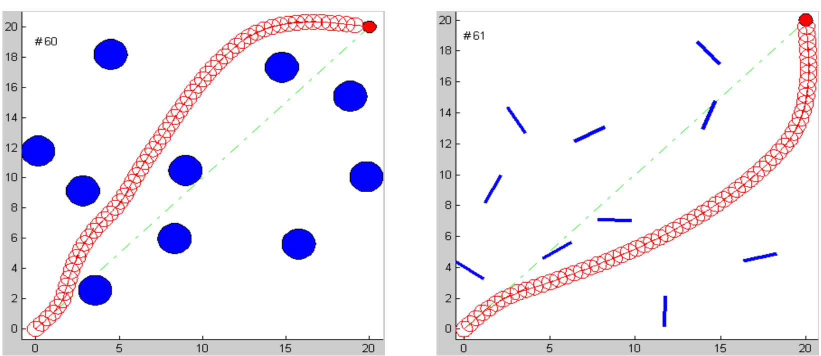

Fig. 7. Two example simulations of robot proceeding to target and avoiding obstacles along the way. Plan view with robot travelling from (0,0) to (20,20).

8 Summary and Conclusion

We have implemented a robot guidance system built on the regulation of object direction. This solution avoids the complexities associated with optic flow (computational cost, and the problems of a non-static environments). Our solution is inspired by research that suggests that humans navigate through the use of regulation of object direction. Fajen, et al [30] have recently modelled the locomotion of humans around obstacles and towards targets using a system of attractors (target) and repellors (obstacles). The work of Fajen et al is similar to the artificial potential field method and related methods of robot navigation [30-32]. These methods have the disadvantage of being computationally demanding. The work was constrained to only use simple visual variables to which humans have a documented sensitivity. Our approach leads to a system that is very simple, produces robust behaviour and utilises a biologically plausible strategy.

Acknowledgements. This research was supported in part by funds from Nissan Technical Center North America Inc, and the National Science and Engineering Research Council of Canada.

References

1. Lappe, M. Bremmer, F.; & van den Berg, A.V.: Perception of self-motion from visual flow *Trends in Cognitive Sciences* 3 (1999) 329–336.
2. Rushton, S.K., Harris, J.M., Lloyd, M.L. & Wann, J.P.: Guidance of locomotion on foot uses perceived target location rather than optic flow. *Current Biology* 8 (1998) 1191-1194
3. Rogers, B.J. & Allison, R.S.: When do we use optic flow and when do we use perceived direction to control locomotion? *Perception* 28 (1999) S2
4. Wood, R.M et al.: Weighting to go with the flow *Current Biology* 10 (2000) 545–546
5. Harris, M.G. & Carré, G.: Is optic flow used to guide walking while wearing a displacing prism? *Perception* 30 (2001) 811–818
6. Harris, J.M. & Bonas W.: Optic flow and scene structure do not always contribute to the control of human walking. *Vision Research* 42 (2002), 1619–1626
7. Warren, W.H. et al: Optic flow is used to control human walking. *Nature Neuroscience* 4 (2001) 213–216
8. Rushton, S.K. & Salvucci, D.D.: An egocentric account of the visual guidance of locomotion. *Trends in Cognitive Science* 5 (2001) 6–7
9. Rushton, S.K. & Harris, J.M.: The utility of not changing direction and the visual guidance of locomotion (submitted).
10. Gibson J.J.: Visually controlled locomotion and visual orientation in animals. *British Journal of Psychology* 19 (1958) 182–194
11. Loomis, J. M. & Beall, A. C.: Visually controlled locomotion: Its dependence on optic flow, three - dimensional space perception, and cognition. *Ecological Psychology* 10 (1998) 271–286
12. Lee, D. N. Guiding movement by coupling taus. *Ecological Psychology* 10 (1998) 221–250
13. Held, R., & Bossom, J.: Neonatal deprivation and adult rearrangement: complementary techniques for analyzing plastic sensory–motor coordinations. *J. Comp. Physiol. Psychol.* (1961) 33–37
14. Llewellyn KR.: Visual guidance of locomotion. *Journal of Experimental Psychology* 91 (1971) 245–261.
15. Land, M.F. & Lee, D.N.: Where we look when we steer. *Nature* 369 (1994) 742–744.
16. Murray, D.W., Reid, I.D. & Davison, A.J.: Steering without representation with the use of active fixation *Perception* 26 (1997) 1519–1528

17. Donges, E. A two-level model of driver steering behavior *Human Factors* 20 (1978) 691–707.
18. Land, M.F..: The visual control of steering IN Harris, L.R. & Jenkin M. (eds): *Vision and Action.* Cambridge University Press.
19. Wann, J.P. & Land, M.F.: Steering with or without the flow: is the retrieval of heading necessary? *Trends in Cognitive Science* 4 (2000) 319–324
20. Bootsma, R.J.: Predictive information and the control of action: what you see is what you get. *International Journal of Sports Psychology* 22 (1991) 271–278
21. Regan, D. & Kaushall, S.: Monocular judgements of the direction of motion in depth. *Vision. Research* 34 (1994) 163–177
22. Laurent, M., Montagne G. & Durey A.: Binocular invariants in interceptive tasks: a directed perception approach. *Perception* 25 (1996) 1437–1450
23. Smeets, J.B.J. & Brenner, E.: A new view on grasping. *Motor Control* 3 (1999) 237–271
24. Warren WH & Whang S.: Visual guidance of walking through apertures: body-scaled information for affordances. *Journal Experimental Psychology Human Perception and Performance* 13 (1987) 371–83
25. Lee, D.N.: A theory of visual control of braking based on information about time-to-collision. *Perception* 5 (1976) 437–459.
26. Rushton, S.K. & Wann, J.P.: Weighted combination of size and disparity: a computational model for timing a ball catch. *Nature Neuroscience* 2 (1999) 186–190
27. Peper, L., Bootsma, R,J,, Mestre, D.R. & Bakker, F.C.: Catching Balls - How To Get The Hand To The Right Place At The Right Time. *Journal Of Experimental Psychology-Human Perception And Performance* 20 (1994) 591–612
28. Girshick, A. R., Rushton, S.K., and Bradshaw, M.F.: The use of predictive visual information in projectile interception. *Investigative Ophthalmology and Visual Science,* 42 (2001) S 3335
29. Duchon, A.P. & Warren,W.H.: Path Planning vs. on-line control in visually guided locomotion. *Investigate Ophthalmology & Vision Science* 38 (1997) s 384.
30. Fajen, B.R., Warren, W.H., Temizer, S & Kaelbling, L.P.: A dynamical model of visually-guided steering, obstacle avoidance, and route selection (in press) *International Journal of Computer Vision*
31. Khatib, O.: Real-time obstacle avoidance for manipulators and mobile robots. *International Journal of Robotics Research* 5 (1986) 90–99.
32. Duchon, A.P, Warren, W.H. & Kaelbling, L.P.: Ecological Robots. *Adaptive Behavior* 6 (1998) 473–507