Displays 34 (2013) 105–113

Contents lists available at SciVerse ScienceDirect

Displays

journal homepage: www.elsevier.com/locate/displa

Biologically-inspired heuristics for human-like walking trajectories toward targets and around obstacles

Simon K. Rushton [a,b,*], Robert S. Allison [a]

[a] Centre for Vision Research, York University, 4700 Keele Street, Toronto, ON, Canada M3J 1P3
[b] School of Psychology, Cardiff University, Tower Building, Park Place, Cardiff CF10 3AT, UK

ARTICLE INFO

Article history:
Available online 27 November 2012

Keywords:
Locomotion
Walking
Guidance
Egocentric direction
Target drift
Obstacle avoidance

ABSTRACT

We describe simple heuristics, based on perceptual variables, that produce human-like trajectories towards moving and stationary targets, and around moving and stationary obstacles. Interception of moving and stationary objects can be achieved through regulation of self-movement to maintain a target at a constant eccentricity, or by cancelling the change (drift) in the eccentricity of the target. We first show how a constant eccentricity strategy can be extended to home in on optimal paths and avoid obstacles. We then identify a simple visual speed ratio that signals a future collision, and the change in path needed for avoidance. The combination of heuristics based on eccentricity and the speed-ratio produces human-like behaviour. The heuristics can be used to animate avatars in virtual environments or to guide mobile robots. Combined with higher-level goal setting and way-finding behaviours, such navigation heuristics could provide the foundation for generative models of natural human locomotion.

1. Introduction

Human beings locomote through an environment with ease. Whether travelling on foot, or by wheel, humans typically weave around obstacles to reach target objects without incident, or apparent effort. Recently, interest has developed [1–31] in the algorithms or heuristics that underpin these behaviours (mobile robot researchers began investigation into the same problem somewhat earlier). This paper illustrates the power and potential utility of simple heuristics in the guidance of locomotion. The heuristics are based upon the pick up or regulation of two perceptual variables, object direction [32] and object drift [33]. The heuristics generate trajectories that resemble those produced by humans. The heuristics have potential application in the generation of natural looking behaviour in simulations, virtual environments and games; the guidance of mobile robots; and as reference models in the study of human or animal locomotion.

When modelling a complicated control system such as a walking human it is necessary to choose between two approaches. The first approach begins with data (i.e. empirical locomotion trajectories) and the other begins with the system (its known structure or variables). Most researchers choose the former approach: they attempt to fit empirical data with standard dynamical models and hope to end up with variables and parameters that are biologically plausible. In this paper we adopt the second approach. We start with perceptual variables to which humans have a known sensitivity and then use them to construct a model that generates human behaviour.

A second decision in the development of such models is to decide whether to attempt to simulate the whole system (taking into account processing latencies, kinematic constraints, physical inertia and so on) or to concentrate instead on the perceptual-motor control laws. Here we take the latter option. We note that it is relatively easy to extend the heuristics we describe to include processing delays, perceptual noise and mechanical damping, and we have demonstrated the robustness of the heuristics in a robotic implementation elsewhere [34].

We begin by outlining how an observer can intercept a static or moving target by maintaining the target at a constant egocentric direction (‘eccentricity’), α, before discussing the trajectories that result from implementing such a rule. In the next section we describe an alternative (but related) heuristic, cancelling target drift, $\dot{\alpha}$. We then go on to describe how target drift, together with other visual parameters to which an observer has access, can be used to detect obstacles (i.e. determine whether maintaining the current course will result in collision). Finally we describe how the identified variables can be used to solve the general problem of locomoting through the world while avoiding stationary or moving obstacles.

For brevity we will use the terms “observer” and “step”. The former term can be read as human or animal, simulated agent or

* Corresponding author at: School of Psychology, Cardiff University, Tower Building, Park Place, Cardiff CF10 3AT, UK. Tel.: +44 29 208 70086; fax: +44 29 208 74858.

E-mail address: rushtonsk@cardiff.ac.uk (S.K. Rushton).

0141-9382/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.displa.2012.10.006

robot. The latter should be considered to include not just the physical steps of a human, animal or legged robot, but also the discrete perception–action cycles of a wheeled robot.

2. Intercepting static and moving targets

2.1. Maintaining direction

During locomotion, interception of a target is achieved if the target is (i) kept at a fixed direction, α, relative to the observer and (ii) the target gets closer on each step. The direction, α, at which the observer maintains the target will determine the exact trajectory taken. The trajectories that result from use of this heuristic are equi-angular spirals.

The left panel of Fig. 1 illustrates a family of trajectories that intercept a static target (see [34] for illustrations of a mobile robot producing such trajectories). The same constant direction strategy works if the target is moving. The middle panel illustrates a family of constant direction trajectories that intercept a target moving with a constant velocity. The right panel illustrates interception of an accelerating target. Thus, a simple heuristic is able to guide locomotion to intercept a static or moving target without modification. Under this control strategy interception relies on the observer perceiving and regulating a single degree of freedom, the egocentric direction, α, of the target. A locomoting observer could select any trajectory from the constant direction family and switch between trajectories at will. A comparison of the form of trajectory generated by the heuristic described here and a real world trajectory is shown in Fig. 2.

Unless an obstacle must be avoided there is an obvious advantage to selecting a straight-line ($\alpha = 0°$) trajectory. A straight-line trajectory minimises the distance to be travelled, and hence the time to reach the goal and the energy expended. Further during a non-straight approach the curvature increases throughout the course, requiring the modification of the trajectory on every step

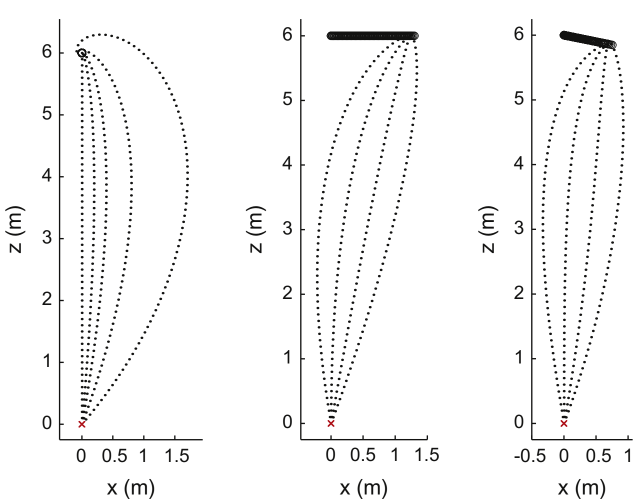

Fig. 1. Family of equi-angular trajectories to a target. All panels display a plan view, with the observer starting at (0,0). Left panel, trajectories that would result from closing on a target while holding it at a fixed egocentric eccentricity, α, of 0°, 5°, 10°, 20° and 40° (from left to right). Observer 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 trajectory based upon holding the target at an eccentricity other than zero results in the observer 'veering' to one side before finally reaching the target. Middle panel, intercepting a moving target. Target starts at (0,6), and moves rightwards, observer starts at (0,0). Four fixed eccentricity trajectories shown, −10°, 0°, 10°, 20° (from left to right). Right panel, intercepting an accelerating target. Target starts at (0,6), and moves rightwards and downwards with increasing speed (constant acceleration), observer starts at (0,0). Fixed eccentricity trajectories shown are −10°, 0°, 10°, 20° (from left to right).

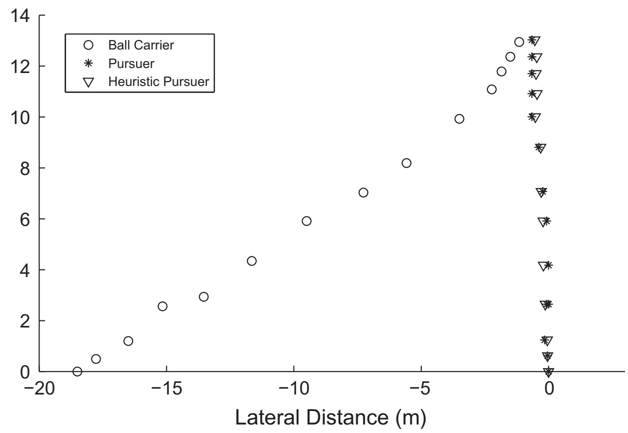

Fig. 2. Shown in plan view is the trajectory taken by an American Football player running to intercept an opposing player carrying the ball [49]. The starting position of both players is towards the bottom of the figure. Interception is shown towards the top of the figure. The intercepting player starts from a position to the right. Overlaid is the trajectory generated by a program that is based on the heuristics described here. Figure is provided for illustrative purposes only, to show that our heuristics can generate "human-like" trajectories.

and challenging the observer's stability (due to lateral acceleration). However under some circumstances it might be preferable to take a non-straight trajectory. Common cases might be when the observer wishes to bypass an obstacle located between the observer and the target (see below) and when the observer wishes to avoid an abrupt change of course.

2.2. Target drift and overcompensation

As noted by Llewellyn [33]: from simple geometry, if the target is to the left of straight-ahead (which is assumed to be coincident with the locomotor axis) then as the observer moves forward the target will drift left. If the target is to the right of straight-ahead then the target will drift right. If the target is only a little away from straight-ahead then the target will drift slowly, if it is a long way away then it will drift rapidly. The rate of drift is the time derivative, $\dot{\alpha}$, of the quantity α, defined above.

2.2.1. Overcompensation

If on each step, when the target drifts, the observer rotates so as to maintain the target at the same direction ('compensates' for the drift) then the observer will follow an equi-angular spiral to the target (as in Fig. 1). However, the observer could 'over-compensate' – for example, if on one step the target drifts $\Delta\alpha$ degrees left, instead of rotating $\Delta\alpha$ degrees left (compensation) the observer rotates twice that amount, $2\Delta\alpha$ degrees left (200% compensation or over-compensation). If the observer uses this strategy she will end up reducing the eccentricity of her trajectory, and thus straightening the trajectory, until the eccentricity reaches zero or perhaps oscillates by a small amount around zero. This is illustrated in Fig. 3.

It can be seen that with overcompensation the trajectory can straighten rapidly (Fig. 3, left panel). The higher the 'gain' or the magnitude of the over-compensation, the more rapidly the trajectory straightens. In the example shown, with the 200% and 400% approaches, the egocentric target direction, α, can be seen to decrease during the course of the approach trajectory (Fig. 3, right panel). The precision (and latency) of the drift estimates and the turning responses will place an upper limit on the magnitude of the gain if the system is to remain stable. This is basically a visual servo mechanism that uses the drift as an error measure and then applies proportional control to correct for the change in direction

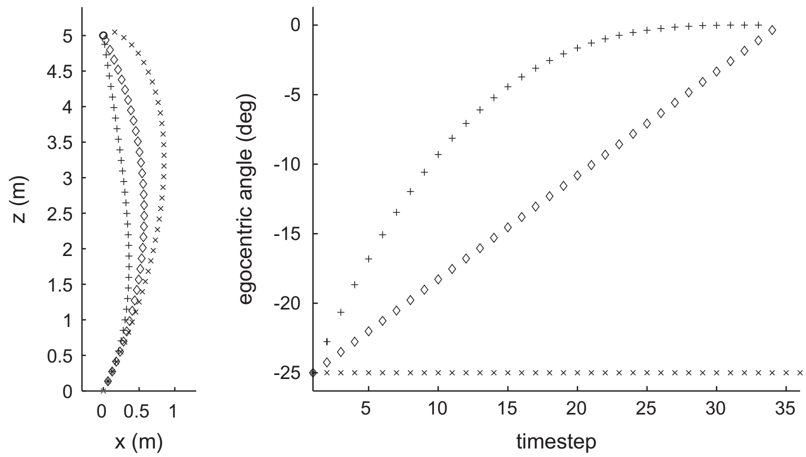

Fig. 3. Over-compensation heuristic with a static target. Left panel, plan view, observer heads towards the target (0,5). Right panel, egocentric direction of the target, α, during the course of the approach. Initial target direction, α, is 25°. Crosses indicate 100% compensation (a fixed angle, equi-angular spiral trajectory) for target drift. Diamonds indicate 200% (over-) compensation, plus symbols indicate 400% compensation.

of the target. Note that it is important for convergence of the path to ensure that the corrective turns are not included in the drift signal. In other words on each step the observer should calculate the correction, turn and then move forward monitoring $\dot{\alpha}$ during the forward motion.

A similar solution works for interception of a moving target. In the left panels of Fig. 4 it can be seen that the interception trajectory straightens with increasing gain. The right panels plots the egocentric direction of the target, α, as a function of time. Note with a moving target the optimal target direction may not be $\alpha = 0$; keeping the target straight-ahead will not necessarily generate a straight line trajectory to the target (see Fig. 4). The over-compensation heuristic converges not on $\alpha = 0$ but rather on the optimal value, the value that generates a straight-line trajectory toward the predicted future location of the target (see Fig. 4).

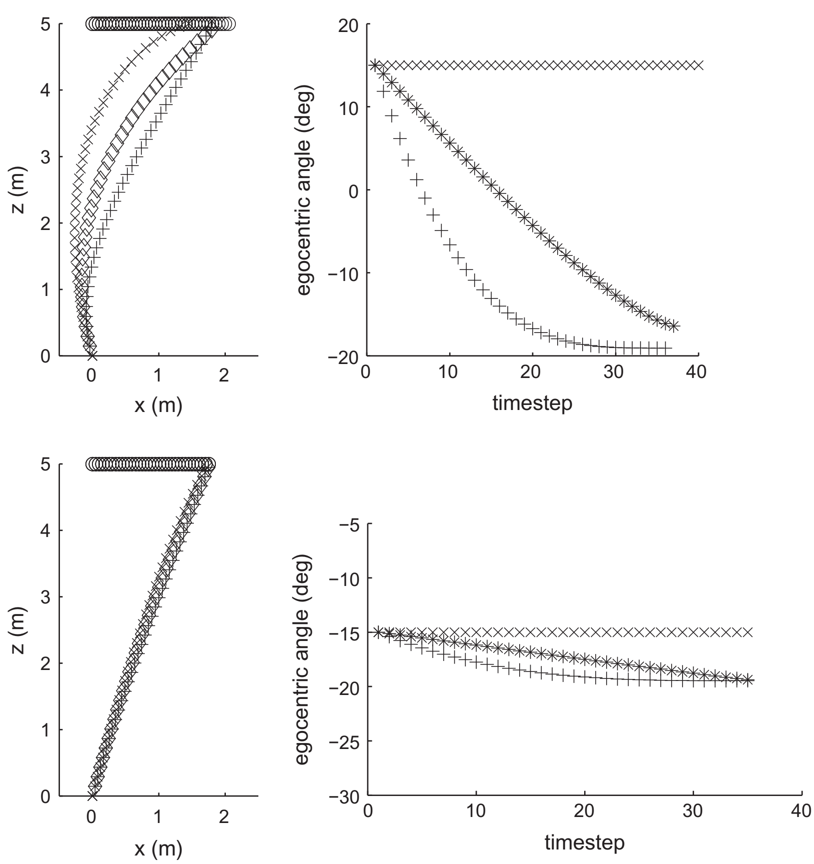

Fig. 4. Over-compensation heuristic with a moving target. Left panels, target starts from (0,5) and moves laterally with a speed of 0.5 m/s. Observer starts at (0,0) and moves at 1.5 m/s. The initial direction of observer travel for the top panels is 15° to the left of the starting position of the target and in the bottom panels, 15° to the right. Crosses indicate 100% compensation, diamonds indicate 200% compensation, plus symbols indicate 400% compensation. Right panels show how the egocentric angle of the target changes during the course of the trajectory.

2.2.2. Target drift in the calibration and prediction

An alternative to using $\dot{\alpha}$ to directly guide locomotion is to use it in calibration or prediction. Let us first consider the role of $\dot{\alpha}$ in calibration.

The ability to rotate and place a target straight-ahead is dependent upon a calibrated system for perception of direction. In the case of humans, laboratory studies indicate that there is drift in the signal for eye-orientation used in the perception of direction [35]. A robot with a mobile head is likely to also suffer calibration problems. We can reverse the logic of the drift-cancellation strategy described above: if a static object is not drifting during forward locomotion it is straight-ahead (aligned with the locomotor axis). Therefore an observer could notice when this occurs and recalibrate straight-ahead accordingly. Alternatively the observer could use the over-compensation strategy outlined above and when the target stops drifting simply reset the direction system with that direction as zero.

Another way to use $\dot{\alpha}$ is in the prediction of α. Rather than maintain $\alpha = k$, the observer could maintain, $\alpha + \dot{\alpha} \cdot \Delta t = k$.[1] Prediction would reduce the amount of drift that would occur between steps. The observer would follow successive chords of the equiangular spiral rather than the tangents to the spiral curve that they follow in the non-predictive approach. As a result they would tend to travel a path interior to the continuous-time spiral rather than outside and, as Δt was increased, the path would straighten.

2.3. Target direction and target drift

To summarise the last two sub-sections, we have two simple heuristics for interception of targets (either moving or stationary): If an observer wishes to reach a target she can regulate her locomotor direction so as to keep the current (or predicted) egocentric direction of the target, α, constant. From the observer's current position to the target there is a family of possible trajectories. An observer can switch from one trajectory within the family to another (for example to avoid an obstacle) by simply changing the direction at which the target is maintained. If the target is stationary and observer wishes to take the optimal straight-line path to the target then she need simply rotate so that the target is straight-ahead ($\alpha = 0$) and then maintain that target direction. If the target is moving no such simple solution exists.

Alternatively the observer can regulate the temporal derivative, $\dot{\alpha}$, or drift-rate. By cancelling target drift she will reach the target by an equi-angular spiral (as in the constant direction strategy above). If the observer over-compensates for target drift she will straighten the trajectory she takes towards the target.

Both the constant direction and the cancel drift heuristics are very simple to implement and "computationally cheap", requiring regulation of only a single visual variable in each case. Furthermore, it is an attractive property of the heuristics that they generalise to the moving target case without additional control principles or heuristics. In contrast, models of behavioural dynamics which use a differential equation approach to fit parameters to empirical trajectory data (e.g. [7]) become increasingly complicated with additional task complexity.

3. Obstacles – detection and avoidance

To successfully move around in an environment it is also necessary to avoid obstacles. In this and the following sections we show that the heuristic suggested above for target interception can also be used to avoid obstacles.

As already noted, if an obstacle is detected then the observer can bypass it by simply changing the egocentric direction, α, at which the target is maintained. Examine the left panel of Fig. 1. If there is an obstacle at (0,3) and the observer is currently on an $\alpha = 0°$ approach, she could switch to a $\alpha = 30°$ approach to the target and pass clear of the obstacle. An important point to note here is that the observer need not switch to an *avoid-obstacle* procedure. At an implementational level a *reach-target* process could run continuously, the obstacle-avoidance process would just intermittently change the target value of α (maintain direction strategy), or inject a rotation of the observer (cancel drift strategy), thus "pushing" it onto a different trajectory (see Fig. 5).

To avoid an obstacle it must first be detected. As noted, if the observer regulates her locomotor direction to maintain the target at a fixed direction, α, she will arrive at the target. This heuristic can be reversed to avoid obstacles – if the observer regulates her locomotor direction so that the obstacle does not remain at a fixed

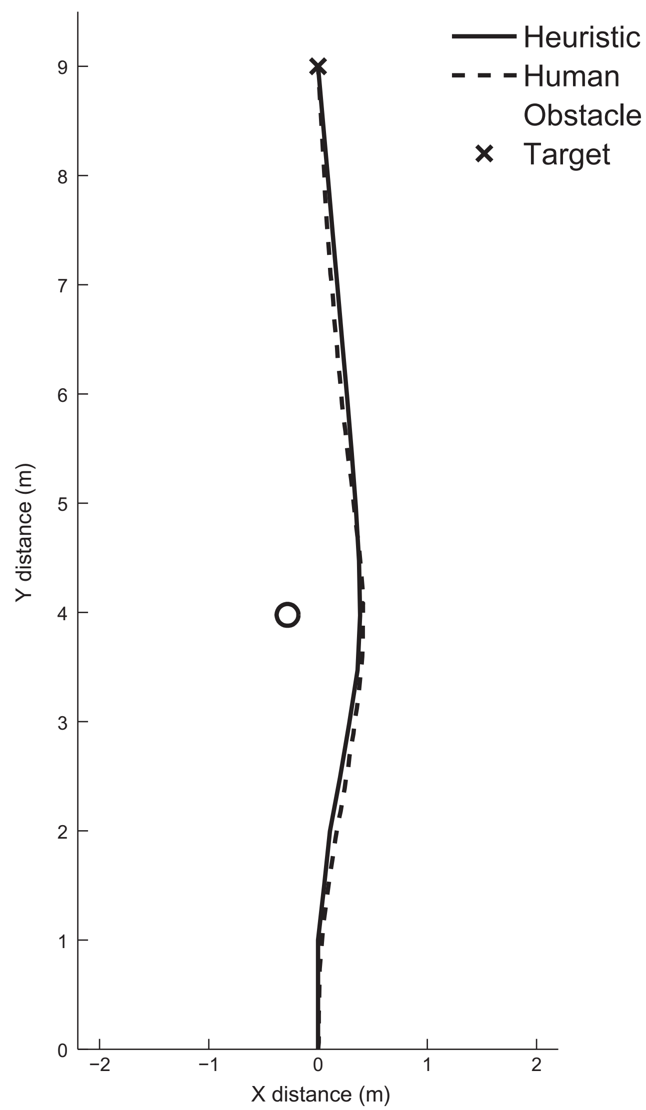

Fig. 5. Adapted with permission from Ref. [7], published by the APA. Shows the mean trajectory taken by an observer walking towards a target and avoiding an obstacle. Trajectory, obstacle and target are shown in plan view. Starting position is towards the bottom of the figure, the target towards the top and the obstacle towards the middle. Overlaid is the trajectory generated by a program that is based on the heuristics described in Section 3. Figure is provided for illustrative purposes only, to show that our heuristics can generate "human-like" trajectories.

[1] A more complex prediction that estimates the distance and direction is possible. This would require use of additional information [36].

direction, α, then she will avoid the obstacle. This rule would be sufficient to avoid collisions if both the observer and the obstacle were simple points. However, in the natural world objects (and the observer) are not points, they have a non-zero width.

First consider obstacles of differing extents and shapes. A very simple solution to avoiding objects of non-zero extent is to steer with respect to the object edges: Rather than monitoring the centre of the obstacle and then scaling the turning response based upon the width of the object a heuristic can work on the boundary edges of the object. If the observer avoids visible obstacle edges (by visible edge we refer to the bounding or 'silhouette' edge in the image rather than an edges defined by the object shape) then the solution works irrespective of the shape and size of the obstacle: silhouette edges constrain the maximal extent of the object or its 'visual hull' [37]. This solution is also in-line with a recent proposal regarding human perceptor-motor control during grasping movements [38].

Next we turn to the more difficult problem, the problem that observers do not have a non-zero width. The solution we propose is based upon a speed ratio, a ratio of two visual parameters.

3.1. Useful speed ratios

If $\dot{\alpha}$ equals zero then the observer is on a head-on collision course; if $\dot{\alpha}$ does not equal zero then the observer is not on a head-on collision course. However, presuming the observer has a non-zero width, if the value of $|\dot{\alpha}|$ is only slightly greater than zero, the observer might be on course for a brushing collision. We can calculate a critical value, such that if $|\dot{\alpha}|$ is higher, the observer will travel safely past the object, if it is less than the observer will collide. Regan [39] pointed out that the ratio $\dot{\alpha}/\dot{\phi}$, where $\dot{\phi}$ is the rate of change of binocular disparity, is approximately proportional to the future passing distance of the approaching projectile, where passing distance is defined as the distance to the left or right at which the object passes the observer (see [40] for an equivalent equation based upon optical looming). The multiplier needed to turn the ratio into a distance (the constant of proportionality) is the interocular separation, I.

$$X_c = I \cdot \dot{\alpha}/\dot{\phi} \tag{1}$$

Therefore an animal or robot with a binocular-vision system could use this ratio to directly determine the distance X_c at which an approaching object will pass.[2]

When the approaching object is travelling over a flat ground plane towards the observer we can substitute $\dot{\rho}$, the change in vertical direction of the bottom edge of the object, for $\dot{\phi}$ [34]. Note that now the constant of proportionality ("multiplier") is the height of the eye or camera, h:

$$X_c = h \cdot \dot{\alpha}/\dot{\rho} \tag{2}$$

Rearranging Eq. (2):

$$R^{X_c} = \frac{\dot{\alpha}}{\dot{\rho}} = \frac{X_c}{h} \tag{3}$$

where R^{X_c} is the value of the ratio for a passing distance of X_c (of course the same quantity can be derived using $\dot{\phi}$ and I). If the absolute value of the speed ratio (Eq. (3)) for an object is greater than a given threshold R^T – perhaps corresponding to a safe crossing distance threshold, X_c^T, of 1.5 times the radius of the observer – then it can be assumed that the observer will safely pass the object if the current course is maintained. If the speed ratio is less than the threshold then the object is on a collision course and so avoidance action needs to be taken.

In the surface plots in Fig. 6 we use the speed ratio to determine *collision areas* – the areas in the environment in which the observer will collide with stationary objects if the current course is maintained. In the shaded areas the speed-ratio is below a critical value, R^T. The observer will hit objects within the shaded area unless a course change occurs.

Moving obstacles are not identified by falling within the collision areas. Moving obstacles on a collision path are indicated by their speed ratio. If an object is currently away from the intended path but is on a collision course then its speed ratio will be below the critical threshold. Similarly, an object moving away from the danger zone will have a speed ratio higher than the critical threshold. This is illustrated in Fig. 7. Fig. 7 is a flow-field representation of the scene with the moving eye or camera as the origin. The speed ratio is indicated by the angle of the motion vector. The collision area for static objects is the area bounded by the dashed lines (compare to left panel of Fig. 6). Moving objects outside of the area may collide with object (see dotted vector to the left of the figure), and moving objects within the area may miss the observer (see dotted vector towards centre of figure).

Moving objects such as those illustrated in Fig. 7 should be detected effortlessly by a human observer. Recent experimental work has shown that objects moving within a scene "pop-out" [41]; it appears that the visual system is able to parse retinal motion into components due to self-movement and components due to object movement [41–48]. A robot system could similarly use optic flow processing to isolate and identify objects moving within the scene. Once a moving object has been detected its speed ratio can be assessed for indication of a future collision.

Previously in the paper we referred to observers moving on equi-angular spirals towards a goal. This begs the question, does the visual speed ratio identify potential obstacles when the observer is on such a trajectory? The simple answer is yes. A slightly modified version of the speed ratio is:

$$R = \frac{(\dot{\alpha} + \omega/2)}{\dot{\rho}} \tag{4}$$

where ω is the angle through which the observer turned during the last step to maintain one's locomotor axis tangent to the curved path. A derivation of this equation is provided in the Appendix. Fig. 8 shows the use of this ratio during movement along an equi-angular spiral trajectory towards a target.

It can be seen that potential obstacles in the immediate vicinity are clearly identified. To understand intuitively why the collision area surrounds the future path, decompose a curving path into its two components: Considered instantaneously, a curving course can be thought of as a linear translation plus a rotation about a vertical line passing through the body. The rotation adds a constant motion component vector across the field. As the magnitude of $\dot{\alpha}$ and $\dot{\rho}$ vary as a function of distance, the addition of a constant to $\dot{\alpha}$ has a differing effect on the ratio, $\dot{\alpha}/\dot{\rho}$, as a function of distance. In Eq. (4), ω can be estimated from either retinal or extra-retinal information, or a combination of the two. Unfortunately the addition of the $\omega/2$ term means that R is not *locally* specified in the retinal flow field.

3.2. The turning function

Now we consider how we can generate an appropriate turning response from the speed ratios described above. We will work through the problem of making a robot turn. A biological

[2] In all the equations and simulations we use a longitude–latitude coordinate system to express directional coordinates in optic arrays (although the analyses generalise to other coordinate systems after suitable modification of the equations). In this system, lines of constant horizontal angle or azimuth correspond to vertically oriented lines of longitude on a spherical projection surface (such as the retina) and lines of constant vertical angle or elevation correspond to horizontally oriented lines of latitude.

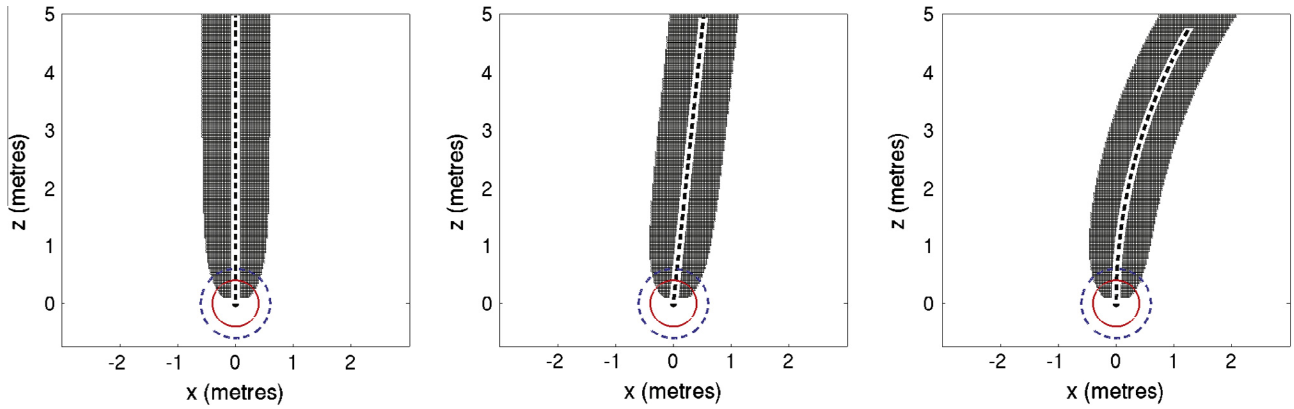

Fig. 6. Plan views of observer on straight, diagonal and constant curvature paths. Dotted line shows future path, shaded area is 'collision area', the portion of space in which an object would have a speed ratio below the "safe ratio". Simulations assume a binocular observer with an IPD of 6.4 cm, a radius of 40 cm and a "safe ratio" of 1.5 times the radius (60 cm). In each panel inner disc indicates radius of the observer, outer disc indicates "safe distance". In all cases it can be seen that the space defined by the safe ratio anticipates the future path.

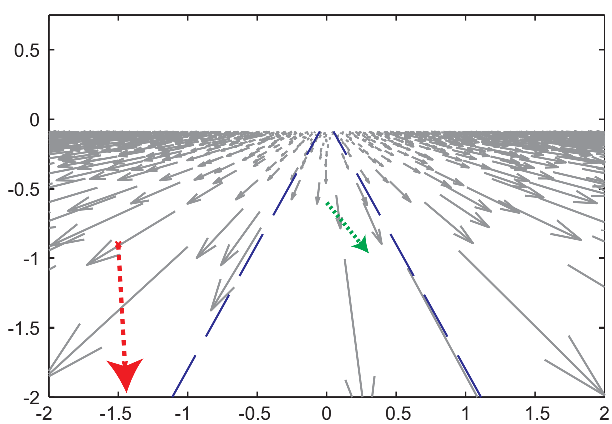

Fig. 7. Flow field representation. Motion vectors seen from eye of moving observer. Vector angle in the flow field indicates the path of the object relative to the observer. The observer would collide with any static objects within the area marked by the dashed lines which define the collision zone. The moving object directly ahead of the observer (dotted vector originating at approximately $(0, -0.6)$) would pass by the observer. The moving object to the left (dotted vector originating at approximately $(-1.5, -0.8)$) is on a collision course.

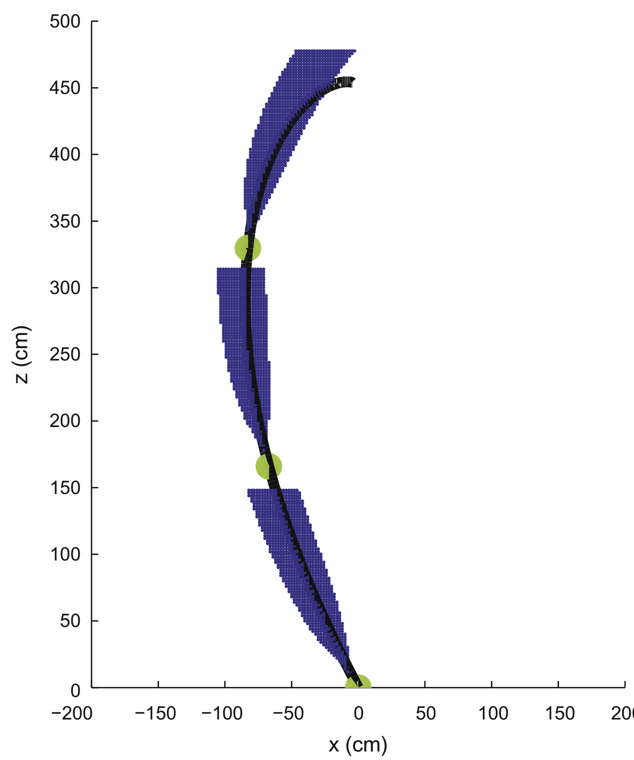

Fig. 8. Plan views of observer on an equiangular spiral path. Dark line shows future path, shaded area is 'collision area', the portion of space in which an object would have a speed ratio below the "safe ratio" for three positions along the trajectory. Simulations assume an observer (green disc) with eye height of 40 cm, a radius of 10 cm and a "safe ratio" of 2 times the radius (20 cm). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

system may pick up $\dot{\alpha}$, $\dot{\rho}$, and $\dot{\phi}$ directly. Therefore it could base its turning on $\dot{\alpha}/\dot{\rho}$ or one of the equivalent ratios. A robot with a single camera would not have access to binocular disparity, ϕ, and so we will use ρ. Also, unlike biological systems, robotic systems do not normally sense speeds – $\dot{\alpha}$, $\dot{\rho}$ – directly so we will use $\Delta\alpha$ (change in lateral direction), $\Delta\rho$ (change in vertical direction) and $\Delta\alpha/\Delta\rho$ instead.

If the ratio $\Delta\alpha/\Delta\rho$ is less than a safety threshold, R^T, this indicates that the obstacle will collide with the observer. To get the robot onto a trajectory that will avoid the obstacle all that is required is to change the eccentricity of approach towards the target. The change in eccentricity ($\Delta\eta$) should be sufficiently large to push the ratio, $\Delta\alpha/\Delta\rho$, to the threshold safe ratio, R^T, so:

$$\left|\frac{(\Delta\alpha + \Delta\eta)}{\Delta\rho}\right| = R^T \tag{5}$$

The above equation can be re-arranged as:

$$\Delta\eta = \begin{cases} R^T\Delta\rho - \Delta\alpha & 0 < \frac{\Delta\alpha}{\Delta\rho} \le R^T \\ 0 & \left|\frac{\Delta\alpha}{\Delta\rho}\right| > R^T \\ -R^T\Delta\rho - \Delta\alpha & 0 \ge \frac{\Delta\alpha}{\Delta\rho} \ge -R^T \end{cases} \tag{6}$$

This equation can be re-written in terms of the quantity, $\Delta\alpha/\Delta\rho$, used above:

$$\Delta\eta = \begin{cases} (R^T - \frac{\Delta\alpha}{\Delta\rho})\Delta\rho & 0 < \frac{\Delta\alpha}{\Delta\rho} \le R^T \\ 0 & \left|\frac{\Delta\alpha}{\Delta\rho}\right| > R^T \\ -(R^T + \frac{\Delta\alpha}{\Delta\rho})\Delta\rho & 0 \ge \frac{\Delta\alpha}{\Delta\rho} \ge -R^T \end{cases} \tag{7}$$

This equation states that the change in eccentricity necessary, $\Delta\eta$, is a function of the difference between the desired value (R^T) and the current value ($\Delta\alpha/\Delta\rho$). This value is large if the object is

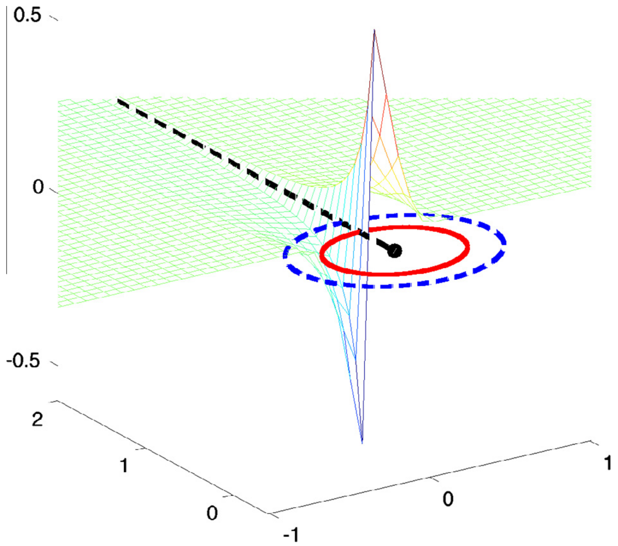

Fig. 9. Surface plot showing turning response as a function of object position. As with Fig. 4, inner disc indicates radius of the observer (40 cm radius), outer disc indicates safe distance (60 cm radius). Simulation assumes observer with binocular vision and inter-ocular distance of 6.4 cm. Vertical axis indicates turning rate as a function of the position of a stationary virtual object relative to the observer. The magnitude of the response decreases with the distance in both the lateral and forward (coincident with the direction of travel) directions.

close to the future path. The change in eccentricity also depends on the quantity $\Delta\rho$. $\Delta\rho$ is proportional to the speed of approach and is inversely proportional to the distance away. This quantity is large if the object is nearby. As a result of these considerations we should expect the largest turning rates to occur if the obstacle is on or close to the future path trajectory and collision is imminent.

Fig. 9 shows a surface plot for the turning function. R^T is set at 60 cm which is 1.5 times the radius of the observer (40 cm). The surface plot is based upon the ratio, $\dot{\alpha}/\dot{\phi}$, which involves binocular disparity, but we remind that a similar plot can be generated based on the ratio $\dot{\alpha}/\dot{\phi}$. As is expected, static objects that are close to the observer and that lie directly on the future path generate the most pronounced response. The turning function attempts to swing to the right of objects that will pass near and to the left and attempts to swing to the left of objects that will pass near and to the right. This minimizes the extent of the avoidance manoeuvre and provides the best chance of avoiding the obstacle.

The surface plot shown in Fig. 9 can be compared to those produced by other locomotion guidance algorithms (e.g. [7]). The general shape is very similar, but the equations from which the surfaces have been computed differ considerably in complexity. We also stress that the variables in Eq. (7) are directly available in the visual flow field, and the constants are simple body scaled parameters.

4. Discussion

The description of heuristics based upon egocentric direction and target drift may inevitably raise some questions regarding the specifics of implementation. We discuss a few of the more interesting or important issues below, and for others we direct the reader to [34] that described the implementation of an earlier variant of the heuristics described here.

4.1. Multiple obstacles

Most scenes contain more than a single potential obstacle. There are two ways to deal with multiple obstacles: either all obstacles could be taken into consideration in the selection of trajectory, or alternatively obstacles can be avoided one at a time. Previous simulations suggest that the second approach is sufficient [34]. The decision that remains is how to select which obstacle to attend and respond to. Obstacles can be prioritised on the basis of current distance from the robot, anticipated distance (taking into account approach speed), temporal distance, or closeness to the robot's path. Subtle differences in performance will result dependent upon the choice, but in many situations the results will be equivalent (for example, with static obstacles the object which is closest in distance is also closest in time).

4.2. Left vs. right decision rules

Consider an observer 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? The simple solution described above suggests that it should swing to the left of objects that will pass near and to the right, and swing right of objects that will pass left. However, when the observer is already on a curved path there are additional considerations of time and efficiency. The solution we used in our robotic implementation [34] was as follows:

- Our first rule says that it should take the route that requires the smallest change in eccentricity of approach.
- If there is little difference between going left or right then 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).

By varying what constitutes "little difference" the two rules can be traded-off.

4.3. Loss of view of the target

If the view of a target is transiently lost then it is of no great consequence, if the robot continues with the same velocity then it will remain on approximately the path to the target and any (small) corrections can be made once the target is back in view.

4.4. Dead-ends

The heuristics proposed here will not function if the robot finds itself in a dead-end or faced with a gap between two obstacles that is too small to get through. We do not see this as a shortcoming. If a human finds themselves if such a situation they normally stop and then reason about the situation.

5. Conclusions

In this paper we have described simple heuristics that can be used to guide an observer around a scene: a heuristic for intercepting targets (constant eccentricity approach); straightening paths (overcompensation); and detection and avoidance of obstacles (speed ratio). Using the avoidance algorithms with visible object boundaries makes the solution independent of object size and shape.

The heuristics we propose are simple and general purpose. Additionally, they have the advantage of being based on variables that are directly available to the biological or machine visual system. The heuristics produce human-like behaviour and can be used to guide robots [34] or simulated agents. They may also have some explanatory power when applied to biological systems.

Acknowledgement

This work was supported in part by funding from Nissan Technical Center, North America.

Appendix A. derivation of Eq. (4)

Assume for convenience that we align the coordinate system so that the object of interest is at $\alpha = 0$. Let D be the radial distance of the object and h be the eye height

$$TTP = \frac{D(t)}{D'(t)} = \frac{\frac{h}{\tan(\rho(t))}}{h \cdot \frac{-\rho'(t).\sec^2(\rho(t))}{\tan^2(\rho(t))}} = \frac{\sin(\rho(t)) \cdot \cos(\rho(t))}{-\rho'(t)}$$

where TTP is the time to passage for the object across a line through the eye, perpendicular to the direction of the object (and ρ is the vertical direction of the bottom edge of the object – see main text). The negative sign arises since distance should be decreasing for a positive TTP.

$$x(t) = \frac{h \cdot \tan(\alpha(t))}{\sin(\rho(t))} \cong \frac{h \cdot \alpha(t)}{\sin(\rho(t))}$$

the small angle approximation for $\tan(\alpha)$ is justified by our choice of coordinates which ensures α is near zero

$$x'(t) = \frac{h \cdot \alpha'(t) \cdot \sin(\rho(t)) - h \cdot \alpha(t) \cdot \cos(\rho(t))}{\sin^2(\rho(t))}$$

$$\cong \frac{h \cdot \alpha'(t)}{\sin(\rho(t))}$$

Substituting

$$X_c \cong x'(t) \cdot TTP = \frac{h \cdot \alpha'(t)}{\sin(\rho(t))} \cdot \frac{\sin(\rho(t)) \cdot \cos(\rho(t))}{-\rho'(t)}$$

$$= -h \cdot \frac{\alpha'(t) \cdot \cos(\rho(t))}{\rho'(t)}$$

$$\cong -h \cdot \frac{\alpha'(t)}{\rho'(t)}$$

The last approximation ignores the $\cos(\rho)$ which leads to little error for modest ρ but overestimates the crossing distance somewhat when the object is very near (by about 40% if the object is 1 eye height away).

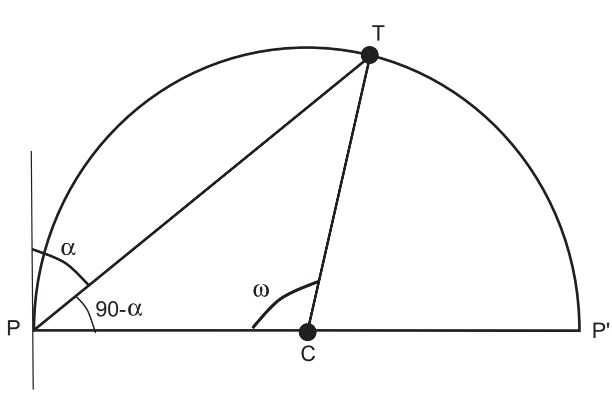

Fig. A1. Geometry when taking a curved path.

A.1. On a curved path

Instantaneous direction is tangent to the curve specifying the future path along a constant radius of curvature path (by definition).

Form a circle tangent to the current direction with current radius of curvature. Circle is centred at C. Target is at T and current position is P. Let P' be opposite (antipodal) to P on the circle (see Fig. A1).

By fundamental theorem of geometry (central angle theorem)

$$\angle P'CT = 2 \cdot \angle P'PT$$

The visual direction of T is

$$\alpha = 90 - \angle P'PT = 90 - \frac{\angle P'CT}{2}$$

Let ω be the rotation around the curved path that will take us from P to T. Assume the observer now makes a small movement along the path (i.e. along the arc). This will change the path length between the target and the observer by a small amount which can be expressed in terms of the change in rotational angle $d\omega$

$$\alpha = 90 - \frac{(180 - \omega)}{2} = \omega/2$$

$$\frac{d\alpha}{d\omega} = \frac{1}{2}$$

or on each time step

$$\Delta\alpha = \frac{1}{2} \cdot \Delta\omega$$

that is, the target drifts at ½ the rate we traverse the circle. All points on future path will have a rate of change of visual direction of ½ the change in direction of travel on each step. We want targets on the future path not to drift so to compensate we need to compensate by subtracting ½ the turning rate. This is the controlled parameter and is therefore available to the observer or could be estimated from the optic flow.

References

[1] J.P. Wann, M.F. Land, Steering with or without the flow: is the retrieval of heading necessary?, Trends Cogn Sci. 4 (2000) 319–324.
[2] B.R. Fajen, Steering toward a goal by equalizing taus, J. Exp. Psychol. Hum. Percept. Perform. 27 (2001) 953–968.
[3] W.H. Warren, B.A. Kay, A.P. Duchon, W. Zosh, S. Sahuc, Optic flow is used to control human walking, Nat. Neurosci. 4 (2001) 213–216.
[4] A. Chardenon, G. Montagne, M. Buekers, M. Laurent, The visual control of ball interception during human locomotion, Neurosci. Lett. 334 (2002) 13–16.
[5] M. Lenoir, E. Musch, E. Thiery, G.J.P. Savelsbergh, Rate of change of angular bearing as the relevant property in a horizontal interception task during locomotion, J. Mot. Behav. 34 (4) (2002) 385–401.
[6] R.M. Wilkie, J.P. Wann, Driving as night falls. The contribution of retinal flow and visual direction to the control of steering, Curr. Biol. 12 (23) (2002) 2014–2017.
[7] B.R. Fajen, W.H. Warren, The behavioral dynamics of steering, obstacle avoidance, and route selection, J. Exp. Psychol. Hum. Percept. Perform. 29 (2) (2003) 343–362.
[8] B.R. Fajen, W.H. Warren, S. Temizer, L. Kaelbling, A dynamical model of steering, obstacle avoidance, and route selection, Int. J. Comput. Vision 54 (1/2) (2003) 13–34.
[9] G. Montagne, M. Buekers, C. Camachon, A. De Rugy, M. Laurent, The learning of goal-directed locomotion: a perception-action coupling perspective, Quart. J. Exp. Psychol. 56 (2003) 551–567.
[10] R.M. Wilkie, J.P. Wann, Controlling steering and judging heading: retinal flow, visual direction and extra-retinal information, J. Exp. Psychol. Hum. Percept. Perform. 29 (2) (2003) 363–378.
[11] A. Chardenon, G. Montagne, M. Laurent, R. Bootsma, The perceptual control of goal-directed locomotion: a common control architecture for interception and navigation?, Exp Brain Res. 158 (2004) 100–108.
[12] B.R. Fajen, W.H. Warren, Visual guidance of intercepting a moving target on foot, Perception 33 (2004) 689–715.

[13] J.P. Wann, R.M. Wilkie, How do we control high speed steering?, in: L.M. Vaina, S.A. Beardsley, S.K. Rushton (Eds.), Optic Flow and Beyond, vol. 324, Kluwer Academic Publishers, Synthese Library, Dordrecht, 2004.
[14] W.H. Warren, B.R. Fajen, From optic flow to laws of control, in: L.M. Vaina, S.A. Beardsley, S. Rushton (Eds.), Optic Flow and Beyond, Kluwer, 2004, pp. 307–337.
[15] W.H. Warren, B.R. Fajen, Behavioral dynamics of human locomotion, Ecol. Psychol. 16 (1) (2004) 61–66.
[16] J. Bastin, G. Montagne, The perceptual support of goal-directed displacements is context dependent, Neurosci. Lett. 376 (2005) 121–126.
[17] R.M. Wilkie, J.P. Wann, The role of visual and non-visual information in the control of locomotion, J. Exp. Psychol. Hum. Percept. Perform. 31 (5) (2005) 901–911.
[18] J. Bastin, C. Craig, G. Montagne, Prospective strategies undelie the control of interceptive action, Hum. Mov. Sci. 25 (2006) 718–732.
[19] J. Bastin, S. Calvin, G. Montagne, Muscular proprioception contributes to the control of interceptive actions, J. Exp. Psychol. Hum. Percept. Perform. 32 (2006) 964–972.
[20] W. Huang, B.R. Fajen, J. Fink, W.H. Warren, Visual navigation and obstacle avoidance using a steering potential function, J. Robot. Auton. Syst. 54 (4) (2006) 288–299.
[21] L. Li, B.T. Sweet, L.S. Stone, Humans can perceive heading without visual path information, J. Vis. 6 (2006) 874–881.
[22] B.R. Fajen, W.H. Warren, Behavioral dynamics of intercepting a moving target, Exp. Brain Res. 180 (2007) 303–319.
[23] Q.-C. Pham, H. Hicheur, G. Arechavaleta, J.-P. Laumond, A. Berthoz, The formation of trajectories during goal-oriented locomotion in humans. II. A maximum smoothness model, Eur. J. Neurosci. 26 (2007) 2391–2403.
[24] R.M. Wilkie, J.P. Wann, R.S. Allison, Active gaze, visual look-ahead, and locomotor control, J. Exp. Psychol. Human Percept. Perform. 34 (2008) 1150–1164.
[25] H. Hicheur, Q.-C. Pham, G. Arechavaleta, J.-P. Laumond, A. Berthoz, The formation of trajectories during goal-oriented locomotion in humans. I. A stereotyped behavior, Eur. J. Neurosci. 26 (2007) 2376–2390.
[26] Q.-C. Pham, H. Hicheur, On the open-loop and feedback processes that underlie the formation of trajectories during visual and nonvisual locomotion in humans, J. Neurophysiol. 102 (2009) 2800–2815.
[27] L. Li, J.C. Cheng, Heading but not path or the tau-equalization strategy is used in the visual control of steering toward a goal, J. Vis. 11 (2011).
[28] Q.-C. Pham, A. Berthoz, H. Hicheur, Invariance of locomotor trajectories across visual and gait direction conditions, Exp. Brain Res. 210 (2011) 207–215.
[29] M. Cinelli, W. Warren, Do walkers follow their heads? Investigating the role of head rotation in locomotor control, Exp. Brain Res. (2012) 1–16.
[30] T.A. Herlihey, S.K. Rushton, The role of discrepant retinal motion during walking in the realignment of egocentric space, J. Vis. (2012) 12.
[31] S.K. Rushton, J.M. Harris, M.R. Lloyd, J.P. Wann, Guidance of locomotion on foot uses perceived target location rather than optic flow, Curr. Biol. 8 (1998) 1191–1194.
[32] K.R. Llewellyn, Visual guidance of locomotion, J. Exp. Psychol. 91 (1971) 245–261.
[33] S.K. Rushton, J. Wen, R.S. Allison, Egocentric Direction and the Visual Guidance of Robot Locomotion: Background, Theory and Implementation. in: H.H. Bülthoff, S-W. Lee, T.A. Poggio, C. Wallraven (Eds.), Biologically Motivated Computer Vision: Lecture Notes in Computer Science, Springer, 2002. pp. 576–591.
[34] K.R. Paap, S.M. Ebenholtz, Perceptual consequences of potentiation in the extra-ocular muscles: an alternative explanation for adaptation to wedge prisms, J. Exp. Psychol. Hum. Percept. Perform. 2 (1976) 457–468.
[35] S.K. Rushton, Projectile interception, from where & when to where once, in: H. Hecht, G.J.P. Savelsbergh (Eds.), Theories of Time-to Contact, Elsevier, 2004, pp. 327–354.
[36] A. Laurentini, The visual hull concept for silhouette-based image understanding, IEEE Trans. Pattern Anal. Mach. Int. 16 (1994) 150–162.
[37] J.B.J. Smeets, E. Brenner, A new view on grasping, Mot. Control 3 (1999) 237–271.
[38] D. Regan, Binocular correlates of the direction of motion in depth, Vision. Res. 33 (1993) 2359–2360.
[39] R.J. Bootsma, Predictive information and the control of action: what you see if what you get, Int. J. Sport Psychol. 22 (1991) 271–278.
[40] S.K. Rushton, M.F. Bradshaw, P.A. Warren, The pop out of scene relative object movement against retinal motion due to self-movement, Cognition 105 (2007) 237–245.
[41] S.K. Rushton, P.A. Warren, Moving observers, relative retinal motion and the perception of object movement, Curr. Biol. 15 (2005) R542–R543.
[42] P.A. Warren, S.K. Rushton, Perception of object trajectory: parsing retinal motion into self and object movement components, J. Vis. 7 (2007) 1–11.
[43] P.A. Warren, S.K. Rushton, Evidence for flow-parsing in radial flow displays, Vision. Res. 48 (2008) 655–663.
[44] P.A. Warren, S.K. Rushton, Perception of scene-relative object movement: optic flow parsing and the contribution of monocular depth cues, Vision. Res. 11 (2009) 1406–1419.
[45] P.A. Warren, S.K. Rushton, Optic flow processing for the assessment of object movement during ego movement, Curr. Biol. 19 (2009) 1555–1560.
[46] C.S. Royden, E.M. Connors, The detection of moving objects by moving observers, Vision. Res. 50 (2010) 1014–1024.
[47] F.J. Calabro, S. Soto-Faraco, L.M. Vaina, Acoustic facilitation of object movement detection during self-motion, Proc. Biol. Sci. 278 (2011) 2840–2847.
[48] D.N. Lee, R. Lishman, Visual control of locomotion, Scand. J. Psychol. 18 (1977) 224–230.
[49] D.M. Shaffer, T.B. Gregory, How football players determine where to run to tackle other players: a mathematical and psychological description and analysis, The Open Sports Sciences Journal 2 (2009) 29–36.