Optic flow based perception of two-dimensional trajectories and the effects of a single landmark.

R.J.V. Bertin, I. Israël
Collège de France / L.P.P.A.
11, place Marcelin Berthelot
75005 Paris
France
✆ +33 1 44271423
fax: +33 1 44271382
email: rjvbertin AT mailworks.org ; isi AT ccr.jussieu.fr

Optic flow based perception of two-dimensional trajectories and the effects of a single landmark.

R.J.V. Bertin, I. Israël

Abstract

It is well established that human observers can detect their heading direction on a very short time scale on the basis of optic flow. Can they also integrate these perceptions over time to reconstruct a 2D trajectory simulated by the optic flow stimulus? We investigated the visual perception and reconstruction of visually travelled two-dimensional trajectories from optic flow with and without a single landmark. Stimuli in which translation and yaw are unyoked can give rise to illusory percepts; using a structured visual environment instead of only dots can improve perception of these stimuli. Does the additional visual and/or extra-retinal information provided by a *single* landmark have a similar, beneficial effect? Here, seated, stationary subjects wore a head-mounted display showing optic flow stimuli that simulated various manoeuvres: linear or curvilinear 2D trajectories over a horizontal plane. The simulated orientation was either fixed in space, fixed relative to the path, or changed relative to both. Afterwards, subjects reproduced the perceived manoeuvre with a model vehicle, of which we recorded position and orientation. Yaw was perceived correctly. Perception of the travelled path was less accurate, but still good when the simulated orientation was fixed in space or relative to the trajectory. When the amount of yaw was not equal to the rotation of the path, or in the opposite direction, subjects still perceived orientation as fixed relative to the trajectory. This caused trajectory misperception because yaw was wrongly attributed to a rotation of the path. A single landmark could improve perception.

Key words: path reconstruction, ego-motion; optic flow; linear heading, circular heading; landmark; vision.

Introduction

Visual, vestibular and somatosensory systems co-operate to provide coherent information about our ongoing displacements. One of the important immediate types of information available is *where* we are going, often referred to as *heading*. It has been shown that the direction of heading can be detected from an important component of the visual information stream: the *optic flow* (Gibson, 1950; Lee, 1974, 1980), e.g. by judging the direction of their future path after having seen just a short simulated movement through a simple environment. This is typically tested by requiring subjects to indicate to which side they would pass a landmark presented somewhere at the horizon, or to indicate the perceived direction of heading using a pointer. Humans can perform these tasks accurately (Banks *et al.*, 1996; Royden *et al.*, 1992; Royden & Hildreth, 1996; Crowell & Banks, 1993; van den Berg, 1992, 1996; van den Berg & Brenner, 1994a,b; Warren *et al.*, 1988., 1991b; Warren & Saunders, 1995; Grigo & Lappe, 1999; Lappe *et al.*, 1999, 2000), and almost instantaneously (Hooge *et al.*, 2000). This is true for linear trajectories, as well as for circular trajectories (Rieger, 1983; Turano & Wang, 1994; Stone & Perrone, 1997; Warren *et al.*, 1991a,b).

Thus, the optic flow can tell *where* we are going. In navigation, this is not the only interesting or important fact. At some point, it may be necessary to also know *how* we have arrived there, via what route. Is it possible to glean this information from only the optic flow? We addressed the question whether human subjects can reconstruct passively travelled trajectories from the optic flow in a recent study. In this study, we showed optic flow stimuli of 8s duration to 23 subjects, using a head mounted display (HMD). The stimuli simulated horizontal two-dimensional manoeuvres[1] through a virtual environment consisting of dots on a ground surface. The manoeuvres had 3 degrees of freedom: linear and semicircular trajectories, with the simulated (whole-body) orientation fixed in space, fixed relative (yoked) to the trajectory, or changing with respect to both space and trajectory. After each presentation, subjects were required to guide an input device in the form of a model vehicle through the manoeuvre they had perceived: position and orientation of the device were recorded. We showed that under certain conditions this task can indeed be performed, and that the perception of the simulated manoeuvres is similar to the perception of identical, physical manoeuvres travelled passively in the dark ("vestibular perception"; Ivanenko *et al.* 1997a,b).

Problems arise when ambiguous flows are presented, both with heading perception and reconstruction. That is, when the presented flow field resembles the field that another movement would also generate. A well-known example of such a situation is when the stimulus consists of the *retinal* flow that would be generated by making a smooth horizontal eye or head movement during a linear translation. This flow closely resembles the retinal flow generated by a tangential, curvilinear movement[2]. Subjects typically report heading direction estimates that indicate that they perceive a curvilinear movement (Banks *et al.*, 1996; Crowell, 1997; Cutting *et al.*, 1997; Royden *et al.*, 1992, 1994; Royden *et al.*, 1994; van den Berg, 1996; Warren & Hannon, 1990; Warren *et al.*, 1991b; Wann *et al.*, 2000). This is especially true when no extra-retinal or other disambiguating information is available. But if subjects make the appropriate eye movements, or move the head relative to the trunk in the appropriate manner, they are more likely to correctly perceive motion along a straight path (Royden *et al.*, 1994; Crowell *et al.*, 1998, Wann *et al.*, 2000). Adding more structure to the simulated scene (posts, trees, etc.) also helps to disambiguate the optic flow based information (Cutting *et al.*, 1997; Li & Warren, 2000). Our results showed that this illusion also occurs when the stimulus simulates a much larger eye or head rotation, made at higher speed, when the task is reconstruction of the travelled manoeuvre. This suggests that perceived rotation (yaw) is attributed to a rotation of the path (Bertin *et al.*, 2000). Interestingly, the illusion also occurs when using vestibular and somato-kinaesthetic information (Ivanenko *et al.*, 1997b).

In this paper we report on a sequel to the aforementioned study that addressed the effect of rotation (yaw) on the reconstructed manoeuvres (experiment 1). Finally, we also studied the effect of adding a single landmark (that could add additional extra-retinal and/or visual information) to the virtual environment (experiment 2). We compare the results of these two experiments with some of the previously published results to provide a synthesis.

General Methods

Subjects.

16 subjects participated in experiments 1 and 2 and in the previously reported study (these experiments were run in two consecutive sessions with a pause in between): 8 men and 8 women. An additional control experiment (1c) was

[1]We will use the term *manoeuvre* to designate a specific combination of trajectory and orientation along that trajectory. We simulate manoeuvres by presenting the optic flow that would be experienced during the the execution of those manoeuvres in real life.

[2]There is an important difference, though. A tangential, curvilinear movement over a ground plane creates an *ordinal* flow pattern on the retina. That is, all projected image elements travel over non-intersecting retinal paths (*streamlines*): concentric circles in the outer world become "parallel" hyperbolae. This is not so for a linear movement with yaw: the streamlines intersect each other in that case (Kim & Turvey, 1998).

run several months later, in which 16 subjects participated (6 female, 10 male, 8 of which had also participated in experiments 1 and 2). Subjects had varying experience in psychophysical experiments, but all were naive as to the purpose of the experiments. All had normal or corrected-to-normal vision, and were in their middle 20s to early 60s. Stimuli from experiment 1 and experiment 2 were presented intermixed randomly. Subjects gave informed consent for participation; the experimental procedure was approved by the local and national ethics committees, and was conform the Declaration of Helsinki.

Experimental set-up: apparatus.

The experimental set-up has been described in detail in Bertin *et al.*, 2000. We repeat here only the essentials. Optic flow stimuli consisted of dotscapes of white dots distributed (uniform, random; 4 pixel diameter; 4800 dots total) on a (50x50m) ground surface 1m below the virtual eye-level; the horizon was placed at 15m. The flow fields were generated by simulating the movement of a virtual observer along the required manoeuvre, using Performer (Silicon Graphics). Each stimulus consisted of a 2s stationary period followed by 8s of simulated manoeuvre followed by another 2s stationary period. See figure 1a (insets ***I-III***) for time-exposure impressions of stimuli (each with a landmark). Stimulus generation and display were under control of a Silicon Graphics Indigo^2/Extreme workstation with a Virtual Research VR4 head mounted display (HMD; FOV 48° horizontal, 36° vertical, 742x230 pixels, 60Hz refresh), and a Silicon Graphics O2 workstation with an NVision Datavisor LCD HMD (FOV 48° horizontal, 36° vertical, 640x480 pixels, 60Hz refresh). Presentation was binocular but monoscopic.

Subjects were required to reproduce their perception of the simulated manoeuvre after stimulus presentation. That is, they were to guide a model vehicle (a custom-made input device; inset ***V*** in figure 1a) through the manoeuvre they had just perceived, moving it over a graphics tablet. The device's instantaneous position (X,Y; resolution 22860x15240) and orientation (Φo; resolution approx. 4°) were read from the tablet using custom made software running on the Indigo, and saved to disk. During this reproduction (but not during stimulus presentations!), an overhead view of the tablet and the vehicle was presented in the HMD. A cursor showed the device's current position and orientation; previous positions remained visible as dots to indicate the travelled trajectory. Horizontal and vertical lines intersecting in the centre of the image were also shown as a frame of reference (inset ***IV*** in figure 1a). Buttons on the device allowed the subjects to erase unsatisfactory reproductions and accept (save) only those that best represented their percept.

Experimental procedure.

Subjects were seated, unrestrained but under instruction to remain as still as possible during stimulus presentations. The experimenter briefly introduced the experiment to the subjects. It was emphasised in particular that they were going to see "movies" that were intended to give the impression of travelling in a vehicle capable of manoeuvres where its orientation is not yoked to the path, like a hovercraft. The input device was presented as a model of such a hovercraft-like vehicle; we will refer to it hereafter as *the vehicle*. A few possible manoeuvres not used in the experiment were demonstrated to familiarise the subjects with the fact that yaw was not necessarily yoked to the path, in other words, that the vehicle could make skidding and slipping movements in which its orientation could (continuously) change relative to its current direction of movement. Subjects were allowed to to train in the use of the vehicle and the tablet. This also allowed us to ensure they had fully understood the principle of reproducing manoeuvres (2D, 3 degrees of freedom; trajectory and orientation) with the vehicle. Instructions pertaining to the landmarks were given in the pause between the two sessions (see above).

Subjects were instructed to concentrate on reproducing the perceived manoeuvre's geometry but to ignore its scale, making optimal use of the tablet's surface instead (resolution optimisation). After validating their response, they could ask for re-presentations of the same stimulus, until they were sure about the manoeuvre, and satisfied with the way they reproduced it. They were then asked to briefly describe their percept; when this did not match the actual reproduction (e.g. a straight trajectory had been perceived, but the reproduction was curved), they were to either erase and redraw it, or to view another presentation and repeat the reproduction. The experimenter took notes to permit correct off-line interpretation of the reproductions.

Simulated manoeuvres.

A manoeuvre is described at any given instant by a position in space, and three orientations. Here, we ignore scale and use only orientation: *a)* the observer's orientation in space (Φo; independent of the trajectory); *b)* the orientation of the trajectory (Φp: the angle in space of the tangent to the trajectory, *i.e.* the direction of movement) and *c)* the observer's orientation relative to the trajectory, Φr = Φo - Φp. Thus, we can distinguish two types of rotation: Ψo (yaw) and Ψp (the rotation of the trajectory). Angles are expressed in degrees, with negative values indicating clockwise rotation.

The simulated manoeuvres are listed in Table 1 divided in four distinct classes according to the observables described above. Figures 1b through 1d show the stimuli in top view (including stimuli discussed in Bertin *et al.*, 2000 for completeness). Triangular velocity profiles starting from zero velocity (angular and linear) and peaking at t=4s (half the stimulus duration) were used for comparison with earlier vestibular studies (Ivanenko *et al.*, 1997a,b). The figures show the actual scale (in meters) of the simulated manoeuvres.

Condition	Description	Size [m]	P	1	2	Fig
Stimuli with the observer's orientation (yaw) fixed in space:						
semicircle no-yaw	Semicircular trajectory with Φo=0°	R=1.5, 5	☒	⊗	s	C4
Stimuli with the observer's orientation fixed relative to the trajectory:						
semicircle inward	Semicircular trajectory, observer looking inward ("centripetal": Φr=-90°)	R=1.5, 5	☒	-	f	C8
semicircle forward	Semicircular trajectory with tangential orientation (Φr=0°)	R=1.5, 5	☒	⊗	-	C7
Stimuli with the observer's orientation changing in space and relative to the trajectory:						
semicircle 30-yaw	Semicircular trajectory with a quarter rotation (Ψo=30° starting at Φr0=0° and Φr0=15° [C10'])	R=1.5, 5	-	⊗	-	C10
semicircle 90-yaw	Semicircular trajectory with a quarter rotation (Ψo=90° starting at Φr0=0° and Φr0=45° [C5'])	R=1.5, 5	-	☒ ⊗	-	C5
semicircle 120-yaw	Semicircular trajectory with a quarter rotation (Ψo=120° starting at Φr0=0° and Φr0=60° [C11'])	R=1.5, 5	-	⊗	-	C11
semicircle full-turn	Semicircular counterclockwise trajectory with a full rotation (Ψo=360° starting at Φr0=0°)	R=1.5, 5	☒	-	-	C9
linear 90-yaw	Linear translation with rotation Ψo=90° (Φr0=0°)	L=7.8	-	☒	-	L5
linear 140-yaw	Linear translation with rotation Ψo»140° (Φr0=20°)	L=4.7, 7.8	-	☒	f	L6
linear half-turn	Linear translation with rotation Ψo=180° (Φr0=0°)	L=4.7, 7.8	☒	-	s	L7
linear full-turn	Linear translation with rotation Ψo=360° (Φr0=0°)	L=7.8	-	☒	-	L8
rotation in place	A Ψo=-180° clockwise rotation.		☒	-	-	RIP
Idem, with orientation changing in the direction opposite to the trajectory rotation:						
semicircle 64-counter	Semicircular clockwise trajectory with a 64° counterclockwise rotation (Ψo=64°, Φr0= 0°)	R=5	-	☒	s,f	C3
semicircle half-counter	Semicircular clockwise trajectory with 180° counterclockwise rotation (Ψo=180°, Φr0=0°)	R=5	-	☒	-	C2
semicircle full-counter	Semicircular clockwise trajectory with 360° counterclockwise rotation (Ψo=360°, Φr0= 0°)	R=5	-	☒	-	C1

Table 1:

The presented stimulus conditions. The left column lists the condition name, the middle column a description. The other columns lists *Size:* the size of the simulated manoeuvre (L=length, R=radius, in meters). *0,1,2:* in which experiment it was presented (0: our previous study; for experiment 1, a ⊗ indicates that this condition was presented in the later performed control experiment; for experiment 2, this also indicates the type of landmark; s=shifting and f=fixated). *Fig:* how the conditions are labelled in the figures. C= curvilinear; L= linear.

<u>Figure 1 about here.</u>

At the beginning of each session, before the experiment proper, subjects were presented with 1) a simple forward translation and 2) a lateral translation (Φo=90°: condition **linear lateral**). For these two stimuli, feedback was given to help the subjects arrive at the correct interpretation. This also served as a final check whether they completely understood the task, and to help them get used to the optic flow and its presentation in the HMD[3].

Data analysis.

The responses were filtered, smoothed and then resampled to 20 equidistant points per trace as described in Bertin *et al.* (2000). Resampling to 20 points was more than adequate for the large majority of the responses. The stimulus traces in figures 1b-d were obtained by the same process from 60Hz position and orientation data generated by the stimulus programme. Figure 1e shows 2 examples of raw reproduction (1 more, 1 less typical), and the resampled versions: both retain the essential characteristics: angular turns, number of loops (and hence total path rotation) and the total yaw.

The resulting resampled reproductions were used to evaluate the subjects' performance. Performance was classified using two criteria: overall correctness (if a reproduction retains certain key properties of the presented manoeuvre; *globally correct*) and the correctness of the trajectory (*trajectory correct*). The trajectory correct index gives a measure of how often subjects reproduced trajectories of the correct type (i.e. linear, "circular" or "in place"). A response is trajectory-correct when the path-rotation is within a stimulus-specific interval. The globally correct index also takes reported yaw (Ψo) into account and requires stimulus-specific combinations of yaw and trajectory. A response is globally-correct when it is trajectory-correct and in addition *either* the average path-relative orientation *or* the total yaw is within stimulus-specific intervals, *as well as* possibly the initial path-relative orientation. For instance, for a lateral (or oblique) translation stimulus, a reproduced manoeuvre is globally correct when it is linear with the observer facing in the correct direction at a perpendicular (or oblique) orientation to the path. These qualitative criteria were determined automatically, after which a manual verification was performed to remove false positives and an occasional false negative.

Our protocol only allowed us to analyse only scale-independent quantitative properties of the reproductions: reproduced orientation (Φ) and change in orientation (rotation; Ψ). The three orientations and the two types of rotation that can be distinguished have been introduced above. Of these, we quantified our results with Ψp, Ψo and the average orientation relative to the path <Φr>; cf. figure 1f and Bertin *et al.* (2000). For rotations in place, only Ψo is defined. All the other conditions (manoeuvres) could be described by Ψp, Ψo and <Φr>[4].

For statistics (ANOVA) tests, we used relative errors in these measures. For the rotation of the path, the Ψp error was normalised to (divided by) the presented Ψo: this corrects for the attribution of yaw to path rotation that the subjects generally made.

We also analysed the responses with a *figural distance* measure, modified from Conditt *et al.* (1997). This measure quantifies the difference between stimulus and response, with zero indicating a perfect overlap. It is a measure that can be applied to all responses (whether or not a rotation in place), and that captures other aspects as well (e.g. three orthogonal linear segments can correspond to a 180° path rotation, while not being equal to a semicircle). Our measure (Df) consists of two independent components (Dfs and Dfa) that were calculated after all stimulus and reproduced manoeuvres had been normalised to unit length. A spatial figural distance (Dfs) was calculated as the average of the pairwise, squared distances between all points on the two trajectories (point 1 to point 1, point 2 to point 2, etc.[5]). An angular figural distance (Dfa) was calculated as the average squared difference (maintaining the sign) between the orientations

[3] The VR4 has some cushion distortion in the corners of its view. This is an overly common problem with HMDs, probably due to the projected size of the field of view and the physical size and closeness of the screens. The distortion in the Datavisor LCD is less pronounced, and of a slightly different type.

[4]The <Φr> index is meaningful only for conditions with the orientation (approximately) fixed relative to the trajectory.

[5]Conditt *et al.* used a "two-pass average minimal distance" algorithm in their figural distance calculation, determining for each point on trace 1 the nearest point on trace 2, and afterwards the same for the points on trace 2. Since we assured a fixed number of samples per trace through resampling, we could use the simpler method of point-wise calculation, that in addition preserves the traces' inherent directionality. For traces with identical sample counts and reasonable overlap, the 2 algorithms give identical results.

Ψo in radians at the corresponding points (Dfa was normalised by (1+Dfs)/π to put Dfs and Dfa on a comparable scale). The *combined* figural distance was then calculated as the length of the vector spanned by the spatial and angular components: $Df = \ln(\sqrt{Dfs^2 + Dfa^2})$ (the logarithm was taken to obtain a distribution closer to the normal). Prior to this analysis, the maximum overlap was determined between all stimulus/response pairs. This was necessary because the position and overall orientation of the reproduced manoeuvre had not been constrained. It was done by shifting and rotating the response with respect to the corresponding stimulus, under control of a Simplex downhill method that minimised Df.

General Results

Introduction.

Our results indicate that it is possible to reconstruct passively travelled manoeuvres from optic flow. However, almost all subjects indicated that this was a difficult task. Most indicated that they had experienced impressions of ego-motion (vection). When the stimulus simulated a relatively large amount of yaw relative to the simulated translation, subjects often reproduced a rotation in place.

Not all subjects systematically requested re-presentations. Very few subjects asked systematically for many (more than 3) re-presentations. In almost all cases, responses were consistent between re-presentations, with the final re-presentation serving as a check, notably for the perceived direction of rotation (that appeared to be easily forgotten). The results we present hereafter are thus based on the final response only.

Influence of yaw (experiment 1, 1c).

Introduction and Methods

Our previous study had shown that in our reconstruction-from-optic-flow paradigm, subjects assume a fixed orientation relative to the path they travel, and attribute perceived rotation to a rotation of that path (Bertin *et al.*, 2000). This was shown for manoeuvres consisting of a linear and a semicircular trajectory with 180° more yaw than path rotation. In both cases, it led to strongly overshot curvature of the reproduced manoeuvres. Similarly, stimuli with the same amount of yaw (yoked to the path or not) but with different orientation with respect to the direction of travel gave rise to different perceptual errors. Condition semicircle forward was quite well perceived, condition semicircle outward was often perceived as a linear, lateral manoeuvre, and condition semicircle inward was not perceived in any consistent manner. All these conditions simulated a path rotation and yaw of -180°, but a (fixed) path-relative orientation <Φr> of 0°, 90° and -90° respectively.

In experiment 1 complemented by the later control experiment 1c, we studied this influence of the amount of yaw and the (initial) path-relative orientation on the perceived manoeuvre in more detail. Manoeuvres along semicircular and linear trajectories were presented (table 1, figure 1c), as described in the General Methods.

The quantitative measures (yaw, path rotation and figural distance) were averaged per condition across subjects, and compared to the presented values with Student's t-tests.

Results

Figures 2 shows the perceived yaw as a function of presented yaw (Ψo) and figure 3 shows the perception rotation of the path (Ψp) per condition. These observables are averaged over subjects, per condition, and include only those responses that were not rotations in place and in which there was no change in direction of rotation (either Ψo or Ψp). This excludes responses that are clearly uncorrelated with the stimulus and for which Ψp is undefined and/or ambiguous.

Figures 2 & 3 about here.

Subjects appear to be capable to approximate the angle over which they are rotated, even during complex manoeuvres (figure 2a). A small range effect occurs for the linear+yaw stimuli: overshoot for the smaller rotations, and undershoot for the large rotation ($\Psi o=360°$). Difficulties in reproducing may partly account for this, as they may also (partly) for the large standard deviations in the more complex conditions. Does the good performance mean that on average subjects correctly perceive their simulated rotation, fully "specified" in the optic flow? Probably not, as shown by the cases in which rotations in place are (erroneously) reported ("RIP Responses" in the figure). It would appear that subjects incorrectly perceive the simulated rotation when they are not capable of detecting the simulated displacement in the stimulus. Yaw is never incorrectly perceived when the stimulus does not simulate any observer rotation (also cf. Bertin *et al.* 2000).

The perception of path rotation is shown in figure 3a (left part). Perception is close to veridical when orientation remains approximately fixed in the environment ($\Psi o \approx 0°$; slighty undershot for $\Psi o=0°$; p=0.044), and when it is tangential to the trajectory. In experiment 1, a single condition with $0 < \Psi o < \Psi p$ was presented. Path rotation perception in this condition **semicircle 90-yaw** was not significantly different from condition **semicircle forward**. This suggested that perception is close to correct when $0° \leq |\Psi o| \leq |\Psi p|$. The control experiment 1c (data included in the figures) showed that path rotation is, or tends to be, undershot (p< 0.002) when $|\Psi o| \lesseqgtr |\Psi p|$, although the subjects who participated in both experiments did not respond differently in the control. Recall that yaw perception is veridical in almost all conditions. This is not the case in conditions **semicircle 90-yaw** (p=5e-6), and possibly not in condition **semicircle 30-yaw** for which many subjects verbally reported the perception of no yaw (but they reproduced on average 30° of yaw).

When the orientation changes relative to the trajectory, such that $\Psi o > \Psi p$ (this includes the linear+yaw conditions, where $\Psi p=0°$; middle and right parts in figure 3a), subjects seem to assume that the detected rotation results from a rotation of the trajectory, or mostly so. They respond in a similar manner when the simulated yaw is in the direction opposite to the direction in which the trajectory rotates. In this case, the number of reports of straight paths, or even paths curving in the wrong direction, increase with the presented Ψo, highly significant in all cases (p<3e-5) except for **linear full-turn** (only 5 non-RIP responses). This effect is so strong that in condition **semicircle half-counter** (a clockwise semicircle with a counterclockwise half-turn of the observer), the *average* perceived path is more or less a straight line (p=0.007). In condition **semicircle full-counter**, the average response is a three-quarter counterclockwise arc (p=3.9e-7)! The effect is already significant in condition **semicircle 64-counter** (p=0.046; t-test), but not when a landmark is present (see the next section).

During the time course of these ambiguous stimuli, the perceived manoeuvre becomes increasingly incompatible with the simulated manoeuvre. It should be evident from the reproduction if this conflict is perceived, e.g. in the form of corrections. We split all responses in two equal parts, and compared the errors in the two halves (using a condition x halve ANOVA design) to test if such corrections occurred. We only found significant differences for the ambiguous stimuli. However, these differences were not systematic. Typically, the sign of the error changed, without a significant decrease or increase in absolute error. There is thus no support (neither in our current nor in our previously published data) for the hypothesis that the subjects correctly interpret a perceived mismatch between their initial percept and the simulated manoeuvre at some later time. Thus, it seems that most subjects adhere to the manoeuvre they perceive in the first moments of the stimulus presentation, regardless of whether a mismatch is perceived later on. There are subjects who correctly perceive these stimuli, however.

Effects of initial/average path-relative orientation

The good perception of yaw and even of rotation in general breaks down when manoeuvres are simulated in which the observer's (initial) orientation is no longer tangential relative to the followed path: figures 2b and 3b (large radii only; large dots between the conditions indicate significant differences between tangential and non-tangential orientation). This applies to both the perception of yaw (figure 2b) and (by extension?) the perception of the path rotation (figure 3b). When the orientation is outward relative to the path (i.e. the observer "looks over the left shoulder" for rightward curving paths), the perceived rotations undershoot the presented values, in most cases significantly so. In the cases where the orientation is (almost) 90° outward, this can lead to perceived manoeuvres that are (almost) linear trajectories with lateral orientation (yaw: $p=0.006$; path-rotation: $p=4.7e-7$; difference with **semicircle forward**: $p<2.1e-4$). When the orientation is inward (a single condition), rotations are perceived correctly — but only in the 9 (out of 16) responses that are not rotations in place.

Influence of a single landmark (experiment 2).

Introduction and Methods

Experiment 1 confirmed a result from our previous study (Bertin *et al.*, 2000), namely that illusions occur when the optic flow simulates manoeuvres in which the observer doesn't maintain a fixed orientation (or gaze) relative to the travelled path. This confirms that what others have called the 'rotation problem' also occurs when the task is to reconstruct a manoeuvre from optic flow, rather than indicate current heading. There were indications that a structured virtual environment could help disambiguate the optic flow (Cutting *et al.*, 1997), and recently, Li and Warren (2000) showed that heading perception along linear+yaw manoeuvres was more veridical when *multiple* landmarks are present.

In experiment 2, we studied whether minimal additional structuring of the virtual environment — a single landmark — has an effect on reconstruction of a travelled manoeuvre from optic flow. A landmark, when present, was always fixed in the virtual environment, and visible during the whole stimulus presentation so that the additional information provided was available the whole time. This constraint made it impossible to present all stimuli both with and without a landmark. Figure 2d shows the manoeuvres that were presented, once without and once with a landmark (indicated by the large dot).

Landmarks were constructed of dots identical to the dots making up the rest of the stimulus, except for their colour (they could thus be identified individually). Every landmark had a single red dot on top that had to be kept fixated at all times. It was attached to the ground surface via a string of equally spaced blue dots (cf. figure 1a). Some landmarks moved through the observer's field of view because of the simulated manoeuvre and their (fixed!) location in the simulated environment. These landmarks are referred to as *shifting* landmarks (indicated with the suffix **sLM**). Thus, eye movements were induced, that provided extra-retinal information. Motion in depth relative to such a landmark generated parallax of the landmark's dots with respect to each other (retinal information; see inset ***II*** in figure 1a: the blue dots are smeared). Other, *fixated*, landmarks (indicated with the suffix **LM**) remained stationary in (*fixated to*) the middle of the screen and could provide only parallax. Fixated landmarks were also used to test for the effect of the presence of a reference point alone in conditions without movement in depth relative to the landmark. Note that the point to be fixated is almost always at eye-level; when eye-movements are to be made, they are thus always horizontal (with respect to the virtual environment), and the flow contains only radial and horizontal fronto-parallel components[6] (cf. Wann *et al.*, 2000). Exceptions are 4 **semicircle inward** conditions in which the red dot was placed directly on the ground surface; in 2 of these, the (simulated) gaze was tilted downward so that the landmark dot was imaged in the middle of the subjects' field of view (HMD image). A similar **linear 140-yaw** condition was also presented, with a fixated landmark of half height, requiring (simulated) gaze tilt. Only qualitative, aggregate results from these conditions are given (Table 2a,b).

The effect of landmark presence (and type) on the quantitative measures was tested for significance using analyses of variance (ANOVA).

Results

The presence of a landmark was perceived as helpful, but only in the cases where it was stationary in the field of view ("fixated"); otherwise it was considered distracting. The effect of the presence of such a landmark on the subjects' performance (response classification) is shown in Table 2a-d. The number of rotation-in-place responses observed are also

[6]The optic flow is a vector field describing the movements (speed and direction) of individual elements of a visual scene through the field of view of a moving observer, or on a projection screen. This vector field can be decomposed into several orthogonal components. In the present study, we are concerned only with the radial component (pure expansion of contraction of the projected image; corresponds to forward/backward ego-motion), and the horizontal fronto-parallel component (whole-field image motion to the left or right, corresponding to clockwise or counterclockwise ego-rotation).

listed, and the criteria used for determining globally correctness and trajectory correctness. (All responses were included in the following analyses.)

Table 2 and Figure 4 about here.

Positive effects

There are significant improvements in yaw and path perception for the **linear+yaw** conditions, when a landmark is present (figure 4). There is a clear, positive effect on the perception of the travelled path: the (yaw-relative) error in Ψp decreases when a landmark is present (figure 4a; $F(1,14)=8.274$; $p=0.012$), as does the spatial component of the figural distance, Dfs (figure 4b; $F(1,14)=9.043$; $p=0.00942$).

Figures 4c and 4d compare the three different landmark conditions: no landmark, shifting through the field of view, and fixed at the centre of the HMD's field of view (fixated). The path rotation is significantly better reproduced with a shifting landmark (4c; $F(2,30)=3.804$; $p=0.034$). Perception of yaw is also (slightly) better when a shifting landmark is present (4d; $F(2,30)=2.91$; $p=0.07$). The results suggest that the errors are smaller for a shifting landmark than for a fixated one; in fact, only the effect of a shifting landmark is significant.

This improvement can also be seen in the qualitative measures: table 2a. The number of RIP responses decreases sharply ($p=0.022$; Kruskal-Wallis), supporting the fact that the translation component of the simulated manoeuvre was better perceived. The number of globally correct responses increases, although this is significant only for the short version of the manoeuvre ($p<0.0001$).

There are also significant improvements in the perception of the **semicircle inward** conditions. This can most easily be seen in table 2b: the number of RIP responses decreases, and the number of trajectory correct and globally correct responses increase when a landmark is present. This effect was also visible in the figural distance (Df=-1.468 without, Df=-2.287 with; $F(1,14)=34.63$; $p<4e-5$), but not clear in the other quantitative measures.

Finally, there are significant improvements in the perception of the **semicircle 64-counter** conditions. The figural distance Df decreases (from -1.172 to -1.99; $F(1,13)=42.17$; $p=2e-5$; mostly due to the spatial component Dfs), and the relative error in Ψo decreases too (from 0.3018 to -0.0926; $F(1,13)=4.862$; $p=0.046$). Both the spatial and the angular (yaw) aspects of the manoeuvres are thus better perceived when a landmark is present. But curiously, Df improves for the shifting landmark type only, whereas the yaw perception improves only with fixated landmarks. The effect of landmark type is highly significant: $F(2,28)=13.19$; $p=9e-5$. There are no significant effects on the qualitative performance indices (table 2c).

Negative effects.

Subjects thought that only fixated landmarks were helpful, whereas shifting ones were generally judged distracting. Such landmarks that move through the field of view during the stimulus presentation can indeed have a negative effect on the reconstruction of the travelled manoeuvre. This occurred for the **semicircle no-yaw** conditions. Without landmark, these were quite well perceived, but when a landmark was present, the (yaw) relative error in the path perception increased from -0.039 to -0.187 ($F(1,14)=4.847$; $p=0.045$) and the qualitative performance measures reflect this (table 2d).

General Discussion

We studied the perception of ego-movement from optic flow along passive 2D manoeuvres in the horizontal plane, and how subjects reconstructed such manoeuvres. Our question was whether subjects can reconstruct such passively travelled linear and curvilinear trajectories together with the orientation along them and, if so, how well. Or, in other words: can human subjects integrate subsequent instantaneous perceptions of heading (that can be perceived) into a percept of a complex, 2D manoeuvre? In the light of this question we studied the influence of simulated yaw, and also asked whether a single landmark can have a similar, positive effect in our reconstruction task as multiple landmarks can have on heading perception from ambiguous flow.

Perception of simulated 2-D manoeuvres and the influence of yaw

Our results show that subjects could perceive the qualitative nature of a simulated manoeuvre; correct directions, the correct form of trajectory and sometimes also the correct average orientation relative to the trajectory ($<\Phi r>$). How much of these qualitative aspects they perceived, and how well, depended on the complexity of the simulated manoeuvre. Human observers can thus indeed integrate their instantaneous perceptions of heading directions, at least to some extent. This does of course not give any information about the distances travelled in these directions, since optic flow alone does not provide information on absolute linear ego-motion speed. An absolute judgement of the distance travelled can therefore not be made. But human observers are quite capable of making relative distance judgements from optic flow (Bremmer & Lappe, 1999; Bremmer *et al.* 1999). They can also estimate distances when the flow is generated by an accurate rendering of an environment they know well (Redlick, Jenkin *et al.*, 2001). Finally, they can compare the radii of presented curvilinear manoeuvres (our preliminary results).

It is possible to quantitatively estimate rotation, *i.e.* turned angles, from the optic flow, at least in theory. More specifically, it can be retrieved from the instantaneous flow field; no integration of information over time is necessary. Our results suggest that subjects could indeed correctly detect the angle (Ψo) over which they were turned.

The reconstruction of the travelled trajectory was correct when the simulated trajectory was linear, and no yaw was simulated. In this case, the perceived path-relative orientation was also veridical (Bertin *et al.*, 2000); this corresponds to the "classical" heading tasks. The angle (Ψp) over which the trajectory rotates can be determined from the optic flow just as well as any other angle, again in theory. This should at least be so when the optic flow corresponds to a manoeuvre maintaining a fixed orientation relative to the travelled path. Then all rotation is caused by a rotation of the path. Subjects were indeed capable of correctly judging the rotation of the path when orientation was yoked to the path, but only when it was tangential to the trajectory. Perception was much less veridical when the orientation was either inward (towards the centre of rotation, causing unreliable perception; perceived as rotations in place) or outward (away from the centre of rotation; perceived as rotations in place or as lateral translations when the radius was large).

When rotations are present and yaw is not necessarily yoked to the path (the "generic" case), correct perception requires that one can determine whether (and if so, how much) one-self rotates and whether the path (also) rotates. A manoeuvre in which the observer moves with zero yaw (orientation fixed in space) along a circular trajectory is a special case: there is only path rotation. The flow generated by such a manoeuvre in which "the path rotates underneath the non-rotating observer" is very different from when "the observer rotates on top of a straight path". For these manoeuvres, perceived path rotation was also close to (although significantly different from) veridical (it was not good at all when subjects were physically displaced along such manoeuvres: cf. Ivanenko *et al.*, 1997b). Perception was probably good because the rotation of the displacement vector was easily detectable in these flow fields; it was not masked by any other rotation. In other words, no pure frontoparallel flow masked the transitions between radial flow and laminar flow. When yaw (in the same direction as the path rotation) increased, perception became less veridical (path rotation was undershot), until there was more than half as much yaw as path rotation.

The average reported path rotation was always less than the presented value for the stimuli discussed above, even if the difference was not always significant. This may be in line with the results from a recent study by van den Berg *et al.*

(2001) who found that subjects tend to undershoot the curvature of the simulated path. These authors measured heading and path-curvature perception from stimuli with less than 1s duration, not reconstruction as we did. Furthermore, curvature is inversely proportional to radius: a semicircle with a 10m radius will have half the curvature of a semicircle with a 5m radius, and yet they both have 180° path rotation. We did not measure scale (radius), so our results are not directly comparable to those of van den Berg *et al.* However, a number of our subjects reproduced less-than-180° circular arcs with a "flat" appearance (as if part of a big radius circle), and some described them as elliptical rather than circular. This may indeed correspond to an underestimation of the path's curvature.

We also presented stimuli in which there was yaw but no path rotation (linear trajectories), or more yaw than path rotation, or yaw in the direction opposite to the rotation of the path. These stimuli are ambiguous, and give rise to illusory perception — that is they induce the perception of being on a trajectory other than the one actually simulated. A well known example is a "linear+yaw" stimulus that simulates a linear path travelled in combination with yaw (e.g. making a horizontal pursuit eye or head movement while walking straight forward). For such a stimulus, subjects often perceive curvilinear, tangential manoeuvres: they attribute the perceived rotation to a rotation of the path. We also found this behaviour in response to our version of this stimulus. We also found it when the path rotated (a semicircle), but there was more yaw than path-rotation (a full turn; some subjects then reported spiral manoeuvres) and when there was less yaw (120° or 90°; see above). Surprisingly, the attribution also occurred for the stimuli in which yaw was in the direction opposite to the path rotation. Thus, a semicircle could be perceived as a straight line, or as a curvilinear manoeuvre in the opposite direction, depending on the amount of yaw. For condition **semicircle 64-counter** (64° counter clockwise yaw and a clockwise semicircular path, no landmark), subjects on average perceived approximately 130° of path rotation.

Many authors have argued why this kind of illusion occurs. The primary reasons put forward are that 1) the optic flow resembles the flow that would be generated by the illusory manoeuvre, and that 2) it simply does not provide enough information to choose between the alternative interpretations. Thus, the simplest (or most familiar) of the possibly matching manoeuvres is perceived. This is probably a sufficient explanation for short stimulus presentations, in which no big conflicts can occur between the illusory perceived manoeuvre (and its corresponding flowfield), and the manoeuvre specified by the optic flow. All of our linear+yaw stimuli evolved from a forward to a lateral manoeuvre, or beyond (*i.e.* backward), due to the amount of yaw simulated. Thus, there was an increasing incompatibility with the initially perceived, illusory manoeuvre, that should create conflict. Yet, subjects in many cases reproduced manoeuvres that suggest that they more or less extrapolated the initially perceived manoeuvre. Comparison of the 1st and 2nd halves of their reproductions did not reveal any indication that they might have detected a conflict during the course of the stimulus, and adjusted their perception/response accordingly. Instead, the reproductions were in most cases based on the illusory initial percept. Maybe not all subjects monitored the optic flow continuously for gradual changes in heading direction. They could instead have relied on qualitative changes in the simulated manoeuvre (or in the presented optic flow) to update their perceived movement, extrapolating the previous (initial) percept as long as no such "events" occurred. A few subjects were experienced observers; they did detect the qualitatively different phases, and arrived at (more) veridical interpretations. It is relevant to note here that if one casually observes these stimuli (but also the "counter yaw" stimuli), one gets a strong impression of setting off on a trajectory that curves (in the wrong direction), even if one knows that the trajectory is straight, or in which direction it curves.

The effect of initial orientation

The perception of yaw and even of rotation breaks down when manoeuvres are simulated in which the observer's orientation is no longer tangential relative to the followed path, even if the manoeuvres were correctly perceived with tangential orientation. When the orientation is outward relative to the path (i.e. the observer "looks over the left shoulder" for rightward curving paths), the perceived rotations undershoot the presented values, in most cases significantly so. Thus, curvilinear, outward manoeuvres can be perceived as (almost) linear trajectories with lateral orientation. When the orientation is inward, rotations are perceived correctly — but only in the relatively small number of responses that are not rotations in place (which is highest in the small radius conditions).

These effects can probably be explained from a decomposition of the optic flow. A purely linear forward movement gives rise to a radial flow: expansion, with a focus aligned with the direction of heading. This remans true with non-tangential gaze (path-relative orientation): changing gaze only shifts this focus of expansion (FOE) from the field's centre to an off-centre location (for gaze angles smaller than ± 90°). Adding observer rotation to a forward linear movement, so that it becomes a tangential, curvilinear manoeuvre adds a frontoparallel component to the radial flow, which removes the FOE altogether. Heading perception on this sort of manoeuvres remains possible (Rieger, 1983; Turano & Wang, 1994; Stone & Perrone, 1997; Warren *et al.*, 1991a,b), and results from our comparison of different radii suggest that this is true to the extent that the future path is contained within the available field of view (FOV; even if the path is not visible or indicated; in our small radius conditions that were so often were perceived as RIPs, the future path fell outside the FOV). A pure *outward*, curvilinear manoeuvre (e.g. **semicircle outward**) can be obtained by adding observer and path rotation to a linear, lateral manoeuvre. This consists of adding pure frontoparallel flow to a flow that is purely laminar (and frontoparallel; there is no FOE in this case). The balance between these 2 frontoparallel types of flow (and the depth of the scene) will determine how well (and how much of) the rotational and translational components of the simulated manoeuvre will be perceived. Manoeuvres with less than 90° outward orientation are less extreme cases, and it can thus be expected that perception of rotation should be relatively better. This indeed seems to be the case. A pure *inward* curvilinear manoeuvre is obtained in a similar manner, except that here the flows are in different directions. In our stimuli, the laminar flow due to the translational component would be rightward on the screen, whereas the (clockwise) rotational component gives rise to a leftward frontoparallel flow. These components combine to a shear (or vortex) like flow that again depends on the proportion rotation/translation in the simulate manoeuvre. In this light, it is not amazing that so many subjects had trouble interpreting this stimulus! Again, manoeuvres with less than 90° inward orientation should give relatively better performance, but unfortunately we did not test this.

The effect of a single landmark

Can path perception from our longer, more complex stimuli that give rise to illusory perception also benefit from additional, visual or extra-retinal information, as it can in shorter stimuli? To study this, we placed a *single* landmark in a number of our stimuli. This landmark provided additional extra-retinal information when it moved through the field of view; it provided additional visual information (parallax) when there was movement in depth relative to the landmark. Even when it did not provide these two types of information, it could still serve as a reference for the movement of the other points (conditions **semicircle inward**).

In general, the subjects reported that the presence of a landmark facilitated the task, but only when it remained stationary in the field of view. This was confirmed by the figural distance measure: the reproductions indeed resembled the actually presented manoeuvres more when a landmark was present than when it was absent.

For the linear+yaw conditions, perception was more veridical when a landmark was present. This is in agreement with the study of Li and Warren (2000) who also found more veridical perception in shorter presentations of similar stimuli, but using many more landmarks. Our result can also be compared to the results of van den Berg (1996) who finds that heading perception is "much more accurate" when subjects were asked to judge their heading relative to the perceived motion in depth of the fixation point that was stationary on the screen and that corresponded to a fixed point in the virtual environment.

Perception also improved in a number of the curvilinear conditions. Much less RIP responses were found to the conditions **semicircle inward** when the stimulus contained a landmark, compared to when there was no landmark. In the conditions **semicircle 64-counter**, the reproduced manoeuvres were more veridical with than without a landmark's presence, although the effects are harder to interpret in this case. In the conditions **semicircle no-yaw**, however, the reproduced manoeuvres were slightly but significantly less veridical when a landmark was present. That means either that the additional information was of no use to the subject and that the attention needed to interpret that information actually interfered with the ego-motion perception (this agrees with the subjects' dislike of shifting landmarks). But it is also possible that the information provided by the landmark was picked up and used, but changed the perceived manoeuvre, i.e. gave rise to illusions. It is likely that the structuring of an environment (the elements in addition to the "basic" elements, dots, that provide the optic flow) has to fulfil certain conditions that depend on the simulated manoeuvre, in order to be useful. The more complex the manoeuvre, the more additional elements (landmarks) that may be needed, for example. Adding more landmarks relaxes the big constraint faced with our single landmark, namely that at least one should be visible throughout the stimulus duration. This constraint was particularly big in the linear+yaw conditions. In these conditions, the shifting landmark moved across a large part of the field of view — in a zig-zag manner. Yet we have some evidence for these conditions that the improvement of path rotation perception was bigger for a shifting landmark than for a non-shifting landmark. This evidence is tentative and not very consistent. In the linear+yaw conditions, the improvement in perception of both path-rotation and yaw seems to be best for the shifting landmarks. For the **semicircle 64-counter** conditions, path-rotation perception improves with the presence of a shifting landmark, whereas yaw perception improves with the presence of a fixated landmark. For a pure simulated RIP, this would be understandable: a fixated landmark suppresses the OKN that otherwise might interfere with retinal image motion (suppressing the OKN improves vection too). But since path-rotation and yaw occur together in all these stimuli, this can not be the sole explication.

Two critical notes should be made about our response task. First, the good perception of yaw may in part be a result of an observational and/or response strategy. Our stimuli simulated manoeuvres with multiples of 90° and 180° of rotation. Subjects seemed to categorise in steps of at least 90°; estimating a large rotation that did not seem to be a full turn to be a half turn, or three quarters of a turn, and so forth. Similar observations were made in vestibular experiments by one of us (I. Israël). Sometimes subjects also indicated that they had little idea about the true amount of rotation, and their estimation was often considerably off when they erroneously perceived a rotation in place. But for stimuli simulating 64° yaw against 180° path-rotation in the other direction, the reported path rotation was around 130°. This was not significantly different from 180°-64°=116°, which suggests that the accuracy of rotation perception that we found was not solely due to a rough 90° categorisation. Second, the recorded reproductions that subjects made do not always exactly match how they verbally described them. For instance, in the **semicircle 30-yaw** condition, many subjects declared having seen a **semicircle no-yaw** manoeuvre, and yet the average reproduced yaw was approximately -32° — almost perfectly correct. It is impossible to tell whether the reproduction corresponded to the actual perception, and the verbal description was a categorisation; or whether the verbal description was closest to the actual percept, and the reproduction was just sloppy; or whether the truth was somewhere in the middle. But similar, inherent problems are likely to exist in any psychophysical method, quantitative or qualitative.

Conclusion

To summarise, we have shown that it is possible in certain circumstances to reconstruct virtually travelled manoeuvres from optic flow. Reconstructions were veridical when yaw was zero (observer orientation fixed in space; but the trajectory may be curved!); performance dropped but could still be good when yaw was non-zero but yoked to the path. When yaw was unyoked to the path and non-zero, it was perceived as stemming from a rotation of the path. The reconstructions tended to be governed by the initial percept. Perceptual veracity could be increased by structuring the simulated environment to add additional information, which can be as simple as adding a single landmark — that can also provide additional extra-retinal information.

References

Banks, M.S., Ehrlich, S.M., Backus, B.T., & Crowell, J.A. (1996). Estimating heading during real and simulated eye movements. *Vision Res 36*, 431-443.

Bertin, R.J.V., Israël, I., & Lappe, M. (2000). Perception of two-dimensional, simulated ego-motion trajectories from optic flow. *Vision Res 40*, 2951-2971.

Bremmer, F., Kubischik, M., Pekel, M., Lappe, M., & Hoffmann, K.P. (1999). Linear vestibular self-motion signals in monkey medial superior temporal area. *Ann N Y Acad Sci 871*, 272-281.

Conditt, M.A., Gandolfo, F. & Mussa-Ivaldi, F.A. (1997). The motor system does not learn the dynamics of the arm by rote memorisation of past experience. *J. Neurophysiol 78*, 554-560.

Crowell, J.A. (1997). Testing the Perrone and Stone (1994) model of heading estimation. *Vision Res 37*, 1653-1671.

Crowell, J.A. & Banks, M.S. (1993). Perceiving heading with different retinal regions and types of optic flow. *Percept Psychophys 53*, 325-337.

Crowell, J.A., Banks, M.S., Shenoy, K.V., & Andersen, R.A. (1998). Visual self-motion perception during head turns. *Nature Neuroscience 1*, 732-737.

Cutting, J.E., Vishton, P.M., Flückiger, M., Baumberger, B., & Gerndt, J.D. (1997). Heading and path information from retinal flow in naturalistic environments. *Percept Psychophys 59*, 426-441.

Gibson, J.J. (1950). The perception of the visual world. Boston: Houghton Mifflin.

Grigo, A. & Lappe, M. (1999). Dynamical use of different sources of information in heading judgments from retinal flow. *J Opt Soc Am A Opt Image Sci Vis 16*, 2079-2091.

Hooge, I.T.C., Beintema, J.A., & van den Berg, A.V. (2000). Visual search of heading direction. *Perception 29*, 10.

Ivanenko, Y.P., Grasso, R., Israël, I., & Berthoz, A. (1997a). Spatial orientation in humans: perception of angular whole-body displacements in two-dimensional trajectories. *Experimental Brain Research 117*, 419-427.

Ivanenko, Y.P., Grasso, R., Israël, I., & Berthoz, A. (1997b). The contribution of otoliths and semicircular canals to the perception of two-dimensional passive whole-body motion in humans. *J Physiol (Lond) 502 (Pt 1)*, 223-233.

Kim,N.G. & Turvey,M.T. (1998). Visually perceiving heading on circular and elliptical paths. *J Exp Psychol Hum Percept Perform* 24, 1690-1704.

Lappe, M., Bremmer, F., & van den Berg, A.V. (1999). Perception of self-motion from visual flow [Review]. *Trends in Gognitive Sciences 3*, 329-336.

Lee, D.N. (1974). Visual information during locomotion. In: R.B. MacLeod & H.L. Pick (Eds), *Perception. Essays in honor of J.J. Gibson.* (Ch. 14, pp. 250-267), Cornell University Press.

Lee, D.N. (1980). The optic flow field: the foundation of vision. *Philos Trans R Soc Lond B Biol Sci 290*, 169-179.

Li, L. & Warren, W.H.Jr. (2000). Perception of heading during rotation: Sufficiency of dense motion parallax and reference objects (in press). *Vision Res*

Redlick FP, Jenkin M, Harris LR (2001). Humans can use optic flow to estimate distance of travel. *Vision Res.* 41: 213-219.

Rieger, J.H. (1983). Information in optical flows induced by curved paths of observation. *J Opt Soc Am 73*, 339-344.

Royden, C.S. (1994). Analysis of misperceived observer motion during simulated eye rotations. *Vision Res 34*, 3215-3222.

Royden, C.S., Banks, M.S., & Crowell, J.A. (1992). The perception of heading during eye movements [see comments]. *Nature 360*, 583-585.

Royden, C.S., Crowell, J.A., & Banks, M.S. (1994). Estimating heading during eye movements. *Vision Res 34*, 3197-3214.

Royden, C.S. & Hildreth, E.C. (1996). Human heading judgments in the presence of moving objects. *Percept Psychophys 58*, 836-856.

Stone, L.S. & Perrone, J.A. (1997). Human heading estimation during visually simulated curvilinear motion. *Vision Res 37*, 573-590.

Turano, K.A. & Wang, X. (1994). Visual discrimination between a curved and straight path of self motion: effects of forward speed. *Vision Res 34*, 107-114.

van den Berg, A.V. (1992). Robustness of perception of heading from optic flow. *Vision Res 32*, 1285-1296.

van den Berg, A.V. (1996). Judgements of heading. *Vision Res 36*, 2337-2350.

van den Berg, A.V. & Brenner, E. (1994a). Humans combine the optic flow with static depth cues for robust perception of heading. *Vision Res 34*, 2153-2167.

van den Berg, A.V. & Brenner, E. (1994b). Why two eyes are better than one for judgements of heading. *Nature 371*, 700-702.

van den Berg, A.V., Beintema, J.A. & Frens, M.A. (2001). Heading and path percepts from visual flow and eye pursuit signals. *Vision Res 41*, 3467-3486

Wann, J.P., Swapp, D., & Rushton, S.K. (2000). Heading perception and the allocation of attention. *Vision Research 40*, 2533-2543.

Warren, W.H.Jr. & Saunders, J.A. (1995). Perceiving heading in the presence of moving objects. *Perception 24*, 315-331.

Warren, W.H.Jr, Blackwell, A.W., Kurtz, K.J., Hatsopoulos, N.G., & Kalish, M.L. (1991a). On the sufficiency of the velocity field for perception of heading. *Biol Cybern 65*, 311-320.

Warren, W.H.Jr. & Hannon, D.J. (1990). Eye movements and optical flow. *J Opt Soc Am [A] 7*, 160-169.

Warren, W.H.Jr., Mestre, D.R., Blackwell, A.W., & Morris, M.W. (1991b). Perception of circular heading from optical flow. *J Exp Psychol Hum Percept Perform 17*, 28-43.

Warren, W.H.Jr., Morris, M.W., & Kalish, M. (1988). Perception of translational heading from optical flow. *J Exp Psychol Hum Percept Perform 14*, 646-660.

Acknowledgements

This research was carried out as part of the Human Frontiers Science Project RG71/96B. RJVB was supported by the Fondation pour la Recherche Médicale and TMR grant ERBFMDICT 97 2358. The authors wish to thank Michel Ehrette (C.N.R.S., L.P.P.A.) who created the "vehicle"; Dr. S. Wiener for critical reading of the manuscript, France Maloumian for help with some of the figures; the technical staff at CalComp (Scottsdale, CA) and Virtual Research for valuable support.

Tables:

Legend to Table 2

Performance classification; summary of landmark effects. Separate tables list, in the "score" line, the effect of landmark presence on the number of rotation in place responses (RIP), the number of trajectory correct responses (TC) and the number of globally correct responses (GC). Values in parentheses list the number of observations available. Where appropriate, the significance of the with/without landmark differences (Kruskal-Wallis tests) are shown too. The "remarks" line lists for the TC and GC columns the criteria that were used (Ψp: path rotation; <Φr> average path-relative orientation).

2a: results for the **linear 140-yaw** and the **linear half-yaw** conditions (including the half height landmark fixated by tilting).

2b: results for all linear+yaw conditions, for all landmark types (including the landmark dots on the ground fixated by tilting).

2c: results for the **semicircle no-yaw** conditions.

2d: results for the **semicircle 64-counter** conditions.

linear+yaw	**RIP**		**TC**		**GC**	
Landmark	without	with	without	with	without	with
score	11 (63)	3 (77)	12 (63)	19 (77)	6 (63)	11 (77)
	X^2=7.033, p=0.0080		N.S.		N.S.	
remarks			Yp ∈ <-∞°,-20°]		<Φr> ∈ [-110°,70°]	
					short: 1	7
					X^2=15.30, p<1e-4	

Table 2a

semicircle inward	**RIP**		**TC**		**GC**	
Landmark	without	with	without	with	without	with
score	16 (31)	7 (62)	15 (31)	52 (62)	3 (31)	36 (62)
	X^2=17.86, p<2.4e-5		X^2=12.78, p=0.00035		X^2=19.66, p<1e-5	
remarks	due to large radius		Yp ∈ <-∞°,-20°]		<Φr> ∈ [-110°,70°]	

Table 2b

semicircle no-yaw	**RIP**		**TC**		**GC**	
Landmark	without	with	without	with	without	with
score	0 (32)	0 (31)	25 (32)	18 (31)	30 (32)	24 (31)
			X^2=2.878, p=0.089		X^2=3.375, p=0.066	
remarks			Ψp ∈ <-205°,-20°]		<Φr> ∈ [80°,105°]	

Table 2c

semicircle 64c	**RIP**		**TC**		**GC**	
Landmark	without	with	without	with	without	with
score	0 (47)	0 (45)	38 (47)	36 (45)	24 (47)	30 (45)
			N.S.		N.S.	
remarks			Yp ∈ <-∞°,-20°]		<Φr> ∈ [65°,115°]	

Table 2d

Legends to the figures:

Legend to figure 1:

1a: Information displays in the head mounted display: impressions of the optic flow stimuli and the feedback during reproduction. In the experimental conditions, only single dots were seen to be moving, with a slightly higher density and otherwise identical geometry and field of view. See Table 1 for detailed descriptions of the stimuli. ***I***): condition **linear half-turn** with a shifting landmark (after 1.87s); ***II***) condition **semicircle 64-counter** with a fixated landmark (after 2.4s); ***III***) condition **semicircle no-yaw** with a shifting landmark (after 4.68s). ***IV***) an example of the reproduction feedback the subjects saw in the HMD: here, the input device was guided through a tangential, curvilinear manoeuvre (from left to right). ***V***) an exploded view of the "vehicle", the input device manipulated by the subjects.

1b: Top view of the stimuli; stimuli used in our previous study (Bertin *et al.*, 2000). Each curve represents a trajectory (X,Y), the arrows point in the direction of the observer's orientation (Φo) at the indicated locations. These were obtained at 60Hz from the stimulus generator, and afterwards smoothed and resampled to 20 equidistant samples (like the responses; cf. figure 1e). The figure shows only the large conditions, from left to right, top to bottom:
(left): L1: **linear lateral**, L2: **linear oblique 30°**, L3: **linear oblique 120°**; L4: **linear oblique 135°**.
(middle): C6: **semicircle outward** (Φr=90°); C7: **semicircle forward** (Φr=0°); RIP: the **rotation in place;** C8: **semicircle inward** (Φr=-90°).
(right): C4: **semicircle no-yaw;** C9: **semicircle full-**turn; L7: **linear half-turn.**

1c: Top view of the stimuli; experiment 1. Presentation as described for figure1b. For simplicity, those stimuli from set **A** that will directly be compared to stimuli from set **B** are also shown, whereas those that will be compared to stimuli from set **C** are shown in figure 1d. C1: **semicircle full-counter**; C2: **semicircle half-counter**; C4: **semicircle no-yaw**; C5: **semicircle 90-yaw**; C7: **semicircle forward**; C9 **semicircle full-turn**; L5: **linear 90-yaw**; L7: **linear half-turn**; L8: **linear full-turn**.

1d: Top view of the stimuli; experiment 2. Presentation as described for figure 1b; landmarks are shown as large filled dots. All stimuli were presented at least once with and once without the landmark at the indicated position. No distinction is made in this figure between the different types of landmark (sLM: shifting landmark; LM: fixated landmark). C3-s: **semicircle 64-counter, sLM**; C3-f: **semicircle 64-counter, LM**; C4-s: **semicircle no-yaw, LM**; C8-f: **semicircle inward, LM**; L6-s: **linear half-turn, LM**; L6-f: **linear 140-yaw, LM**. Note the small differences between C3-s and C3-f, and L7-s and L6-f.

1e: Examples of raw reproductions (traces labelled 'raw') and the same responses resampled to 20 approximately equidistant points (traces labelled 'cooked'). Reproductions are dimensionless, hence no scale is shown. Both responses are *not* typical. The left-upper pair shows a response with 2 sharp angles that were faithfully preserved by our technique (also cf. the semicircles in figures 1b-d). The right-bottom pair shows a response that is much more different from a semicircle, yet the key features are preserved.

1f: Explanation of the indices used in the quantitative analyses. Ψp, the average rotation of the path is calculated from the average difference between the tangents to the trajectory in 2 consecutive (resampled) points, multiplied by the number of segments per curve (19). The total yaw Ψo is calculated by (non circular) summation over Φo, discounting the initial orientation; thus, 2 full observer turns give Ψo=720°. The average orientation relative to the path, <Φr>, is calculated as the average difference between Φo and Φp in the 20 resampled points. All these measures are expressed in degrees and averaged over subjects. In this example (**semicircle half-counter**), Ψo=180°, Ψp=-180° and <Φr>=179.7°±109.8°. NB: points at which orientation is shown are arbitrarily chosen in this figure!

Legend to figure 2:

Reproduced yaw for conditions without landmark.

2a: Reproduced Ψo (yaw) against the presented Ψo for all conditions (excluding the landmark conditions from exp. 3) with Ψo≠0° and initial orientation Φo=0°. The leftmost collection also shows the responses to the curvilinear stimuli with congruent Ψo (Ψo and Ψp having the same sign; this includes data from the control experiment 1c). The middle collection shows the responses to the curvilinear stimuli with Ψo opposite to Ψp, and the rightmost collection the responses to the linear+yaw stimuli. Ideally, the responses should all fall on the corresponding shaded lines. Points with errorbars (standard deviation in the mean) each represent the average of all responses to one particular manoeuvre, lumped over all sizes at which that manoeuvre was presented. The leftmost collection also shows the Ψo for all cases were the subjects reported rotations in place. These are *not* lumped over size, and, for clarity, the standard deviations are not shown. Asterisks indicate significant deviation from the expected values, determined by Student's t-tests using mean and average standard deviation; * ≡ $p < 0.05$, ** ≡ $p < 0.005$, *** ≡ $p < 0.001$. The counter-yaw conditions are labelled opp(osite) in the figures.

2b: Ψo responses as a function of initial path-relative orientation, per condition. Light grey bars show the average reported yaw, with the standard deviation of the mean. Behind them, dark grey bars show the presented Ψp, and open bars the presented Ψo. Numbers above the horizontal axis give the number of responses per condition. Asterisks indicate significant difference between reported and presented yaw (bars above), large dots indicate significant effect of the initial path-relative orientation (bars left and right below) (t-tests). NB: C6, C7 and C8 are all semicircles with 180° yaw.

Legend to figure 3

Reproduced path rotation for conditions without landmark.

3a: Rotation of the path (Ψp) responses (light shaded bars) for all conditions with initial orientation Φo≈0°, averaged over all sizes in which each condition was presented. Dark shaded bars show the expected (i.e. stimulus) Ψp values; open "hooks" show the presented Ψo. The conditions are ordered according to decreasing counter-clockwise rotation/increasing clockwise rotation. Left part: less yaw than path rotation; middle part: more yaw than path rotation; right part: yaw, no path rotation. Numbers above the horizontal axis show the number of analysed responses. The response value for condition **semicircle 64-counter** is the mean over all responses to the 3 presentations without landmark. Condition **semicircle 120-yaw** (C11) and **semicircle 30-yaw** (C10) were presented in the control experiment 1c. The other results from this experiment are incorporated in the averages shown for the corresponding conditions: **semicircle no-yaw**, **semicircle 90-yaw** and **semicircle forward**. Errorbars show standard deviation of the mean. Asterisks indicate significant differences from the expected values, determined by Student's t-tests; * ≡ $p < 0.05$, ** ≡ $p < 0.005$, *** ≡ $p < 0.001$.

3b: Idem, showing the effect of the initial path-relative orientation. Presentation as in figures 3a and 2b.

Legend to figure 4:

Effects of landmark presence on the perception of yaw (Ψo), path rotation (Ψp) and the spatial figural distance for the linear+yaw conditions. The graphs show the results from ANOVAs (R 1.5.1: http://www.r-project.org) on the difference with the expected values, using box-and-whisker plots. For Ψo, relative error is shown. For Ψp, we show the error relative to the presented Ψo (yaw-relative error), because of the observed tendency to attribute perceived yaw to path-rotation. Big dots indicate the median, thick lines connect the means also shown numerically, the boxes show the interquartile range, the whiskers 1.5 times that range and open dots indicate outliers.

4a: Plot of means for landmark presence on the perception of Ψp. There is a significant increase in performance when a landmark is present.

4b: Plot of means for landmark presence on the perception of the spatial component Dfs of the figural distance. There is a significant increase in performance when a single (shifting or fixated) landmark is present.

4c: Influence of landmark presence and type on Ψp perception. There is a significant effect of landmark presence/type; the benefit of a *shifting* landmark presence is significant.

4d: Influence of landmark presence and type on Ψo perception. A shifting landmark causes marginally significantly better performance.

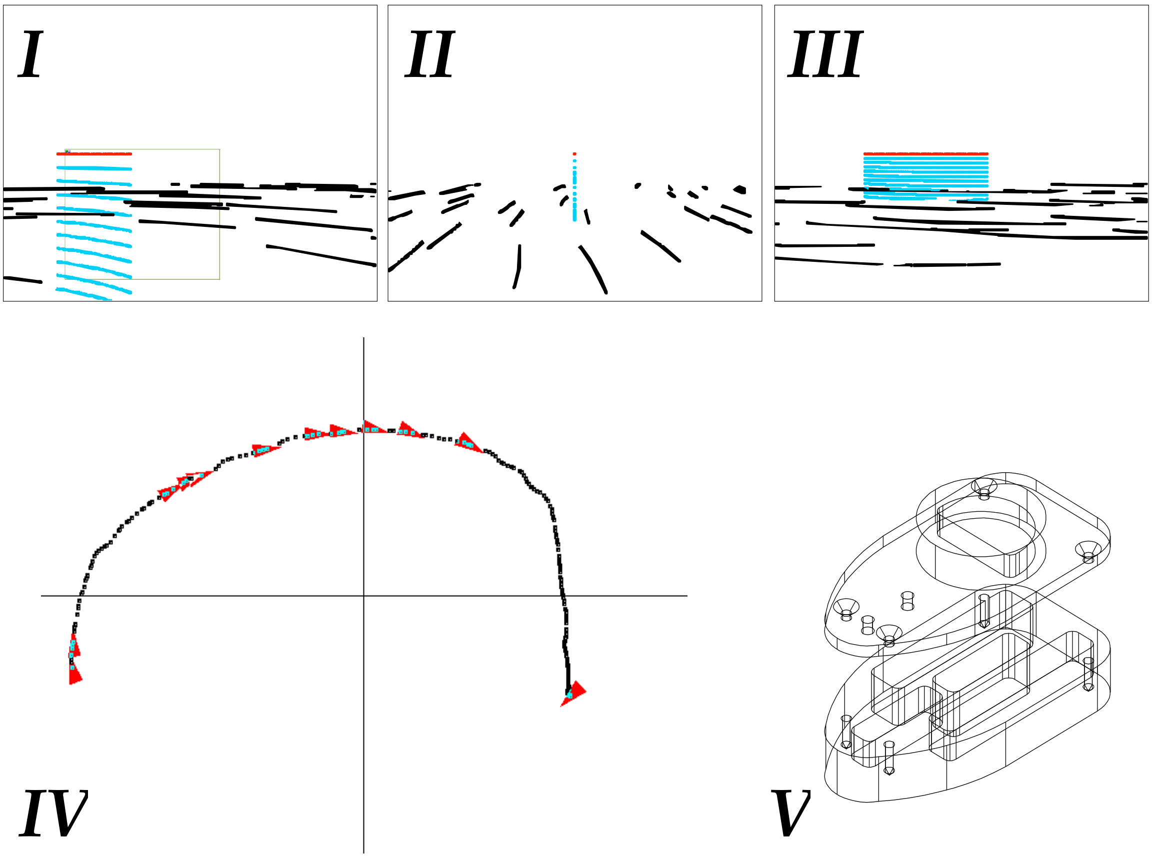

Figure 1a

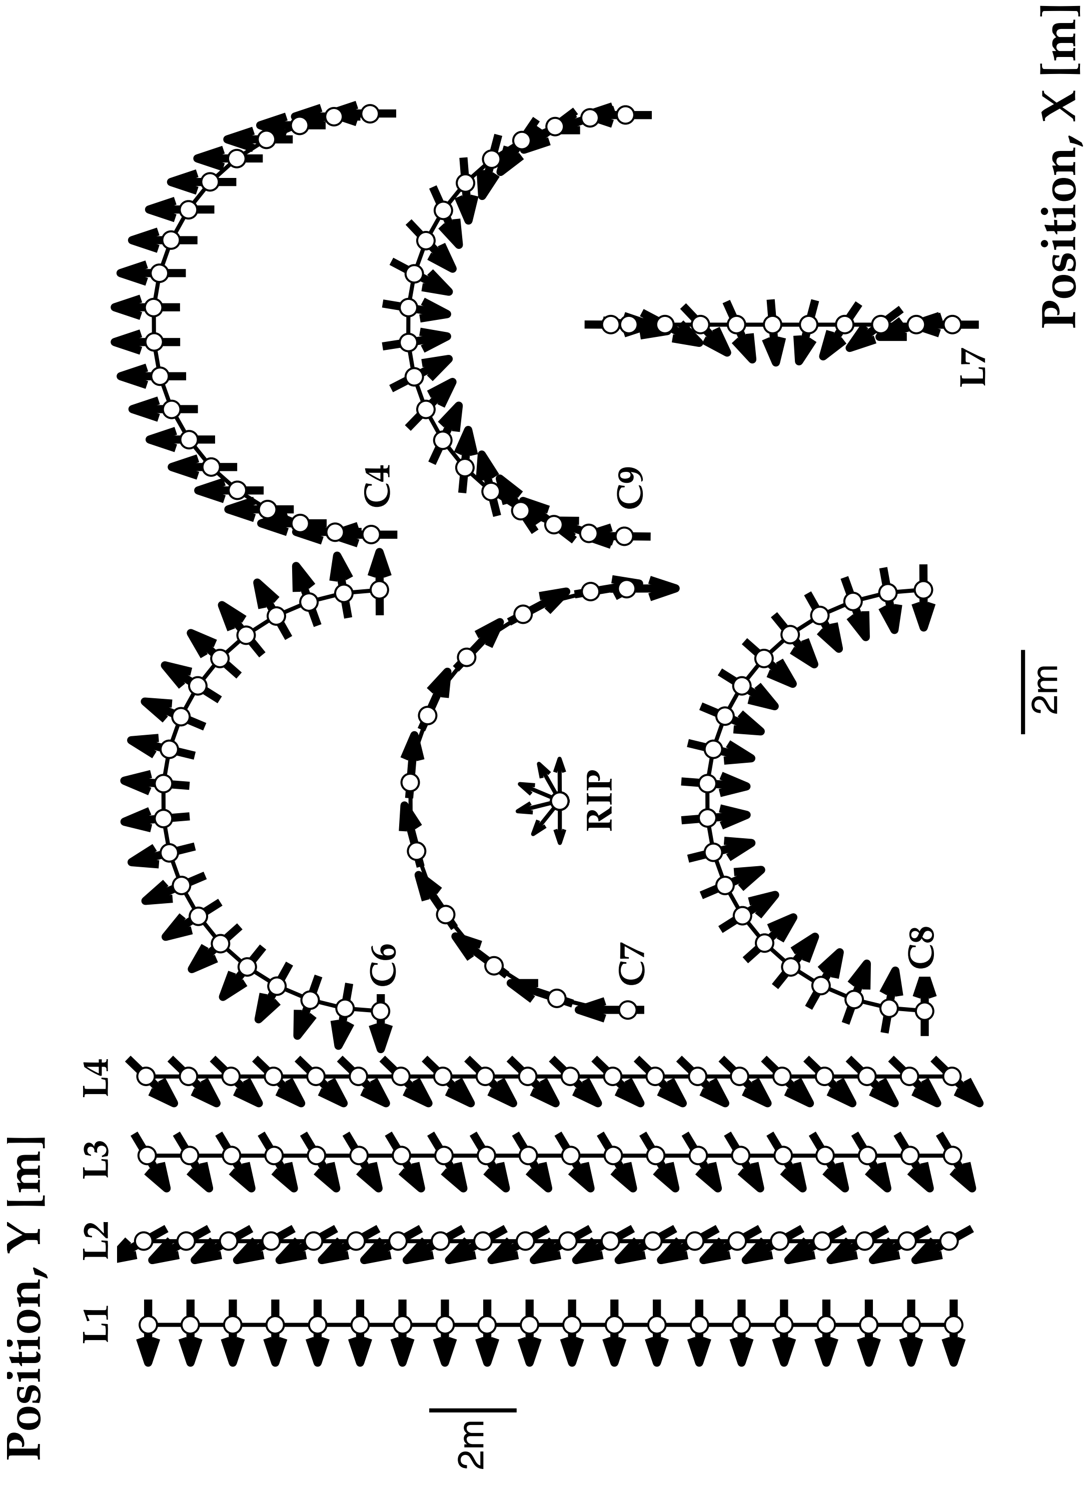

Figure 1b

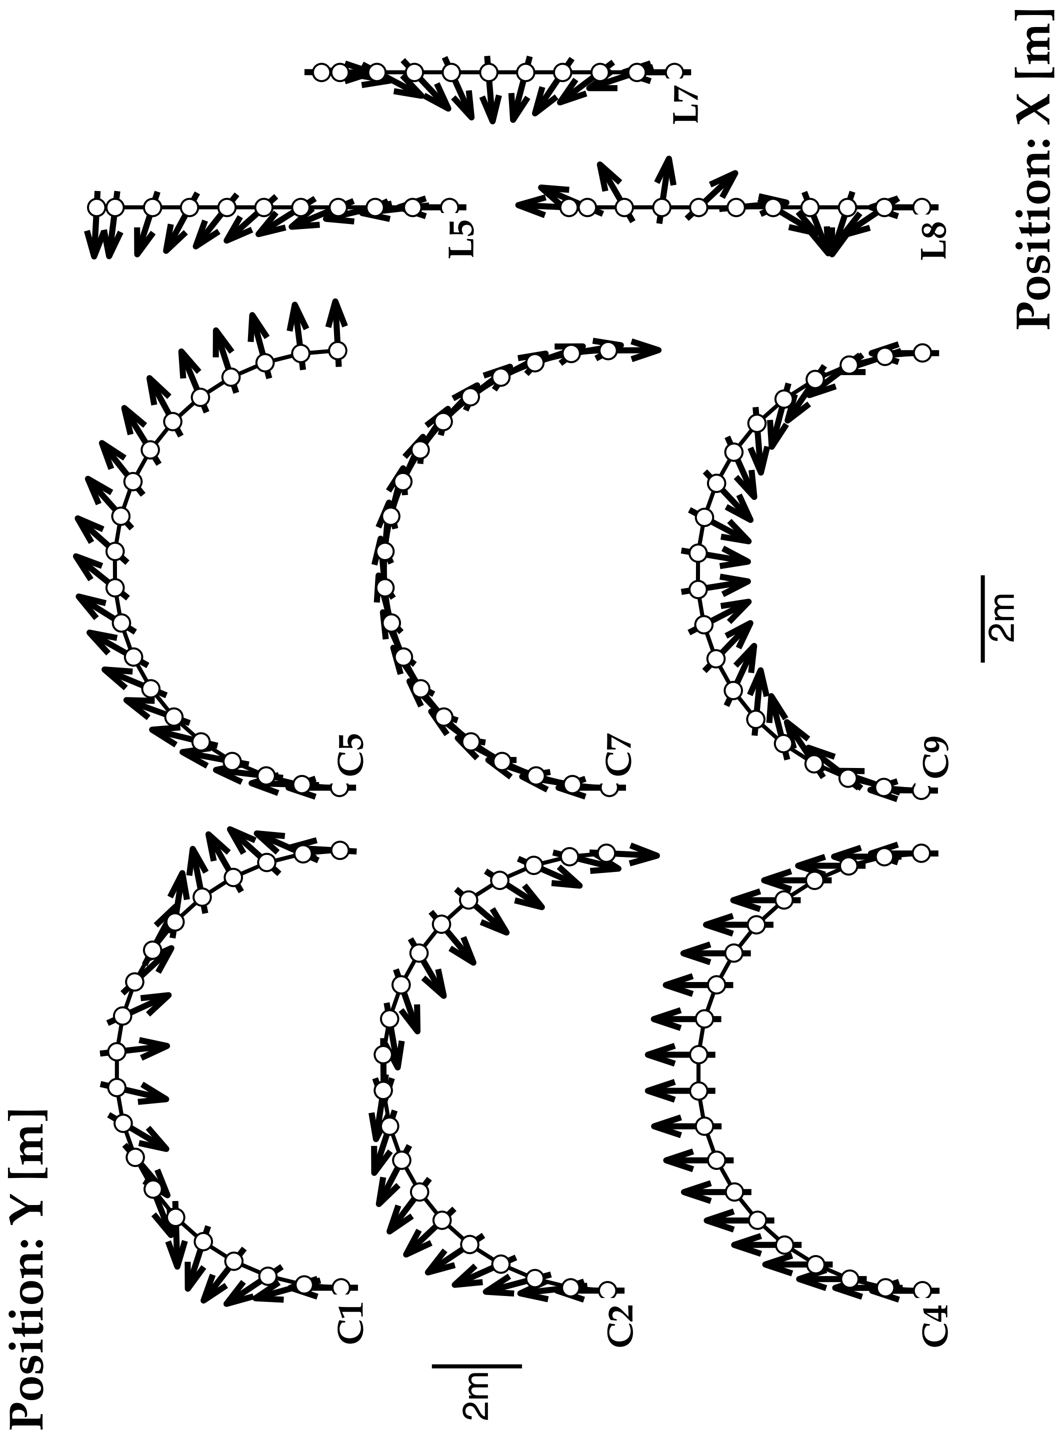

Figure 1c

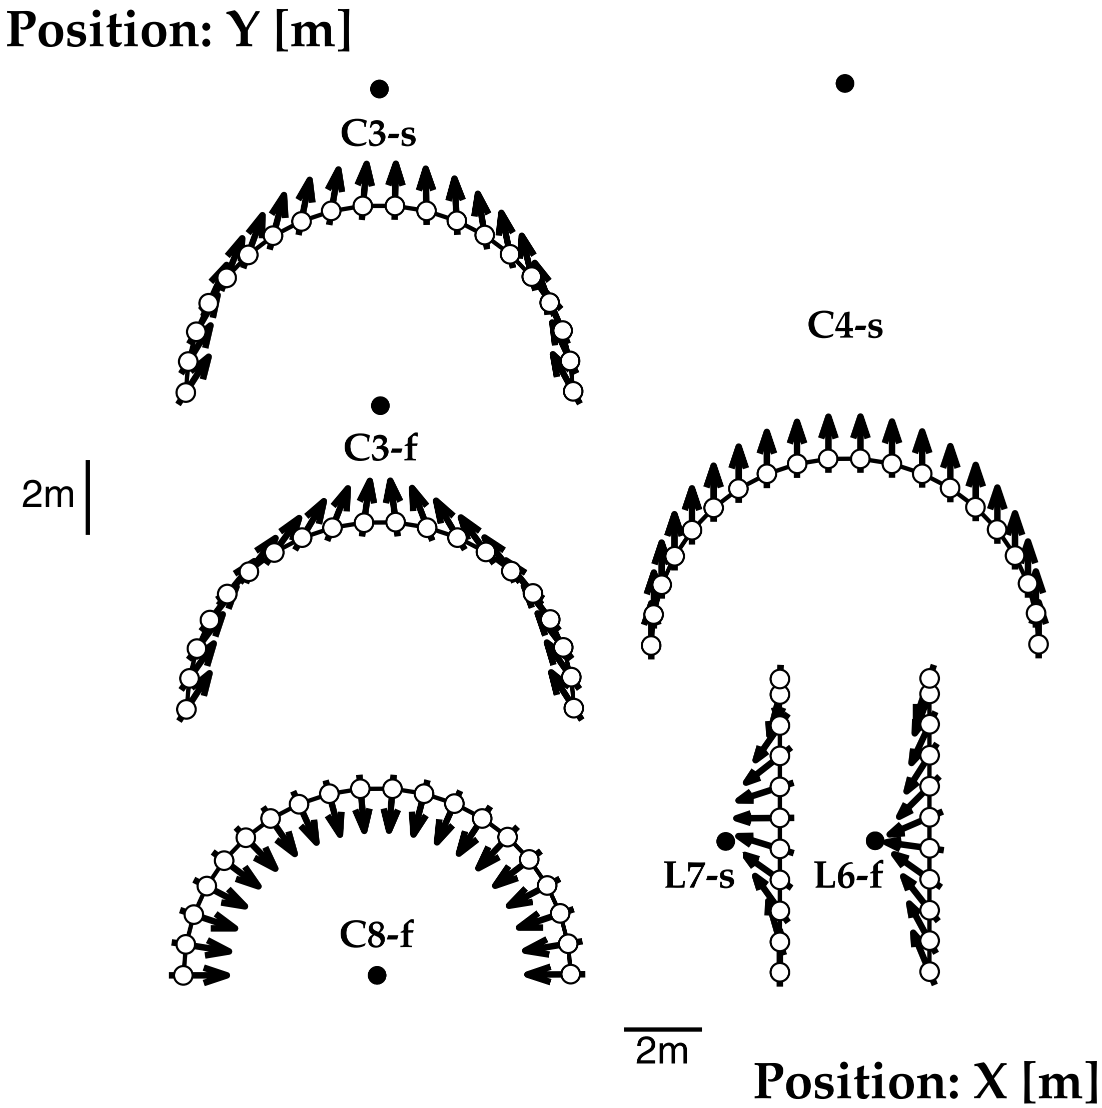

Figure 1d

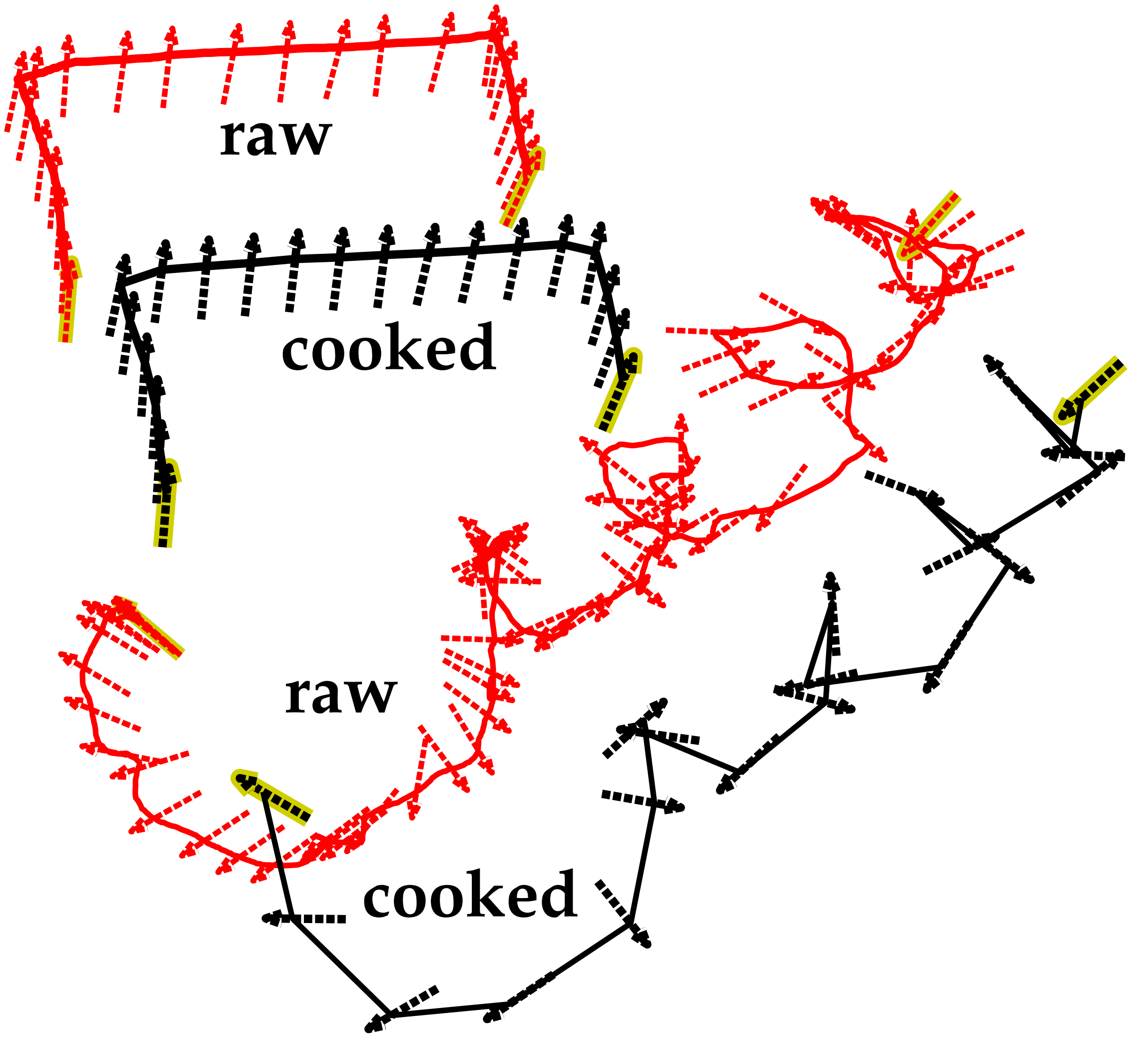

Figure 1e

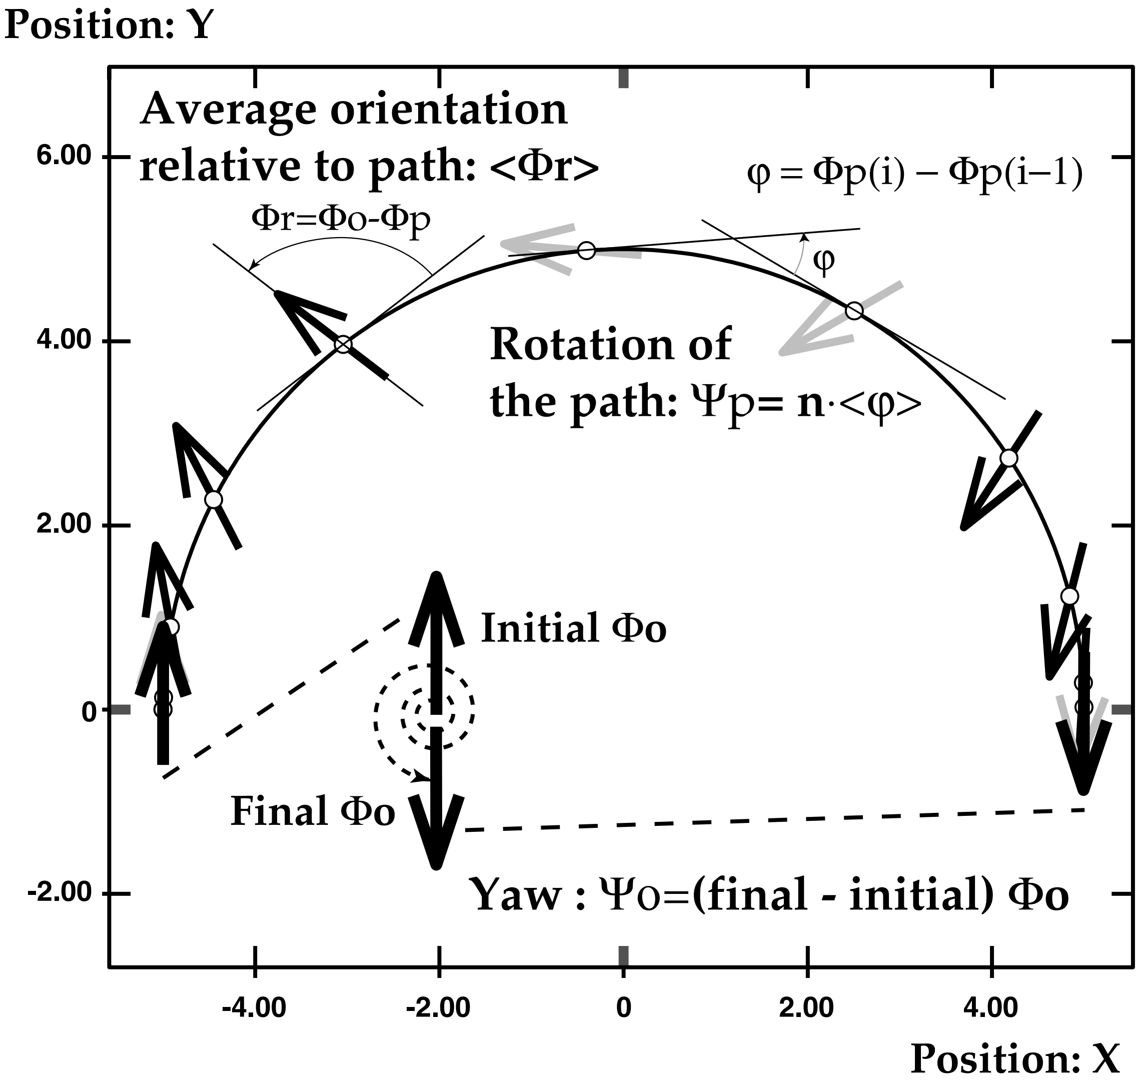

Position: Y
Average orientation
relative to path: <Φr>
φ = Φp(i) – Φp(i–1)
Φr=Φo-Φp
φ
Rotation of
the path: Ψp= n·<φ>
Initial Φo
Final Φo
Yaw : Ψo=(final - initial) Φo
6.00
4.00
2.00
0
-2.00
-4.00
-2.00
0
2.00
4.00
Position: X

Figure 1f

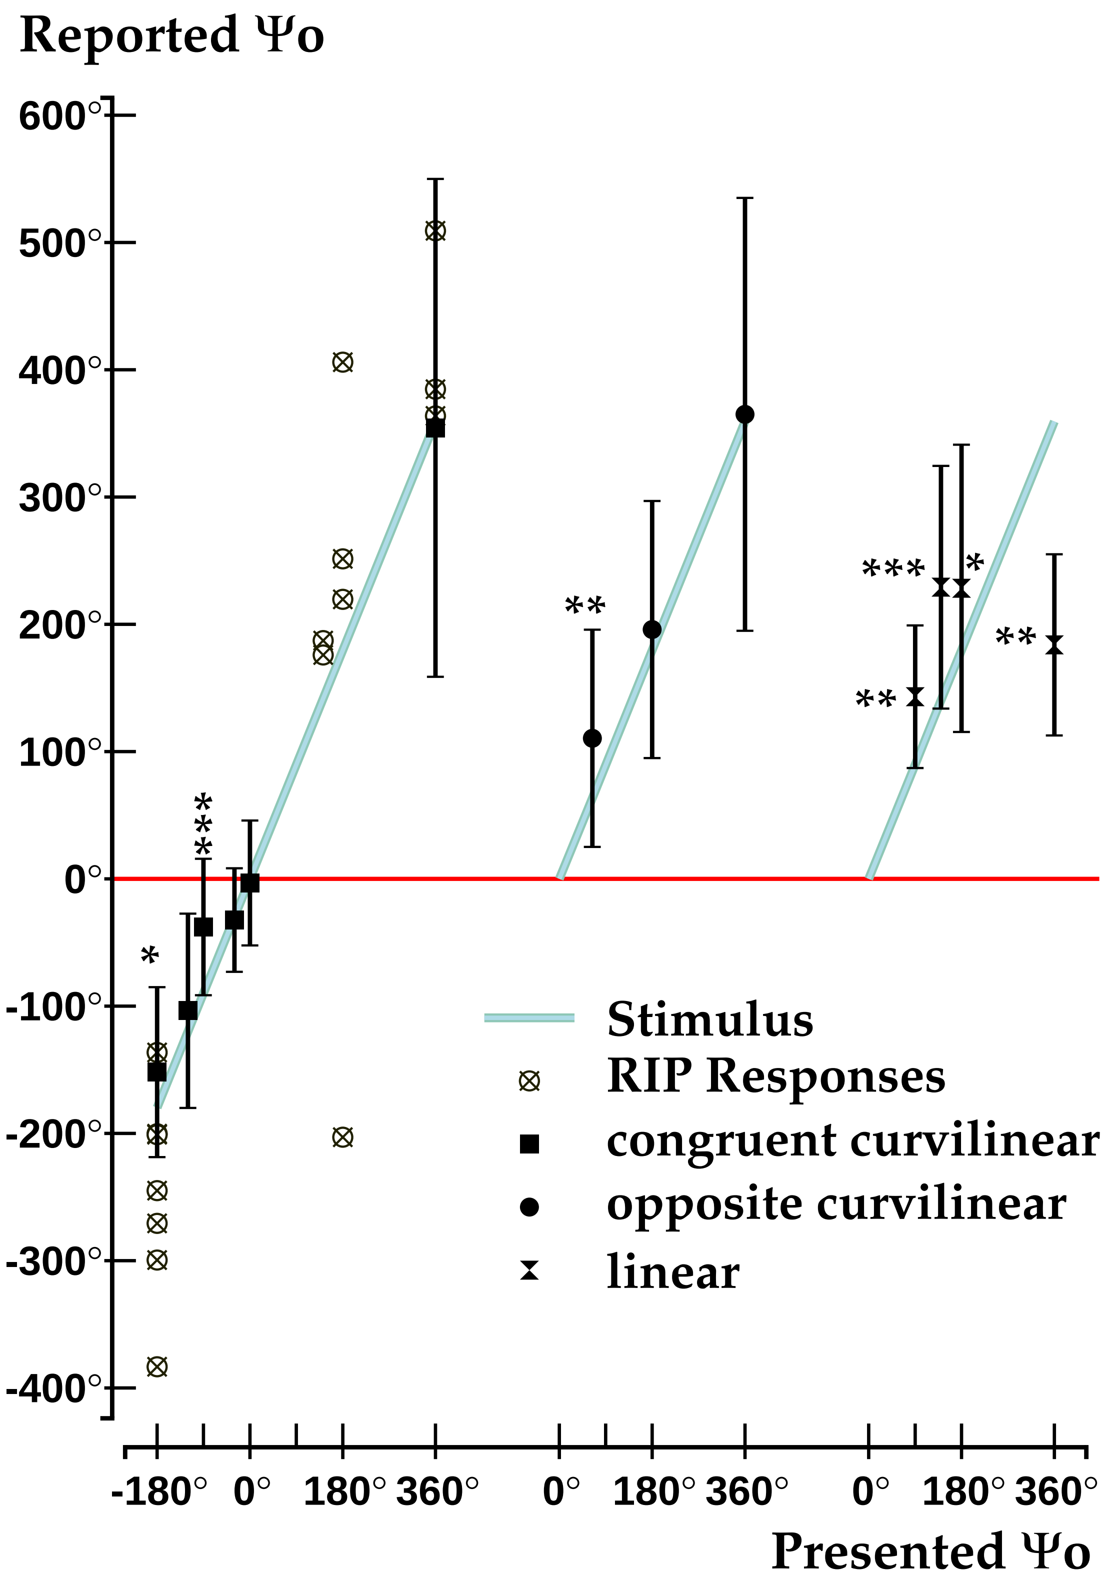
Reported Ψo
600°
500°
400°
300°
200°
100°
0°
-100°
-200°
-300°
-400°
Stimulus
RIP Responses
congruent curvilinear
opposite curvilinear
linear
-180° 0° 180° 360° 0° 180° 360° 0° 180° 360°
Presented Ψo

Figure 2a

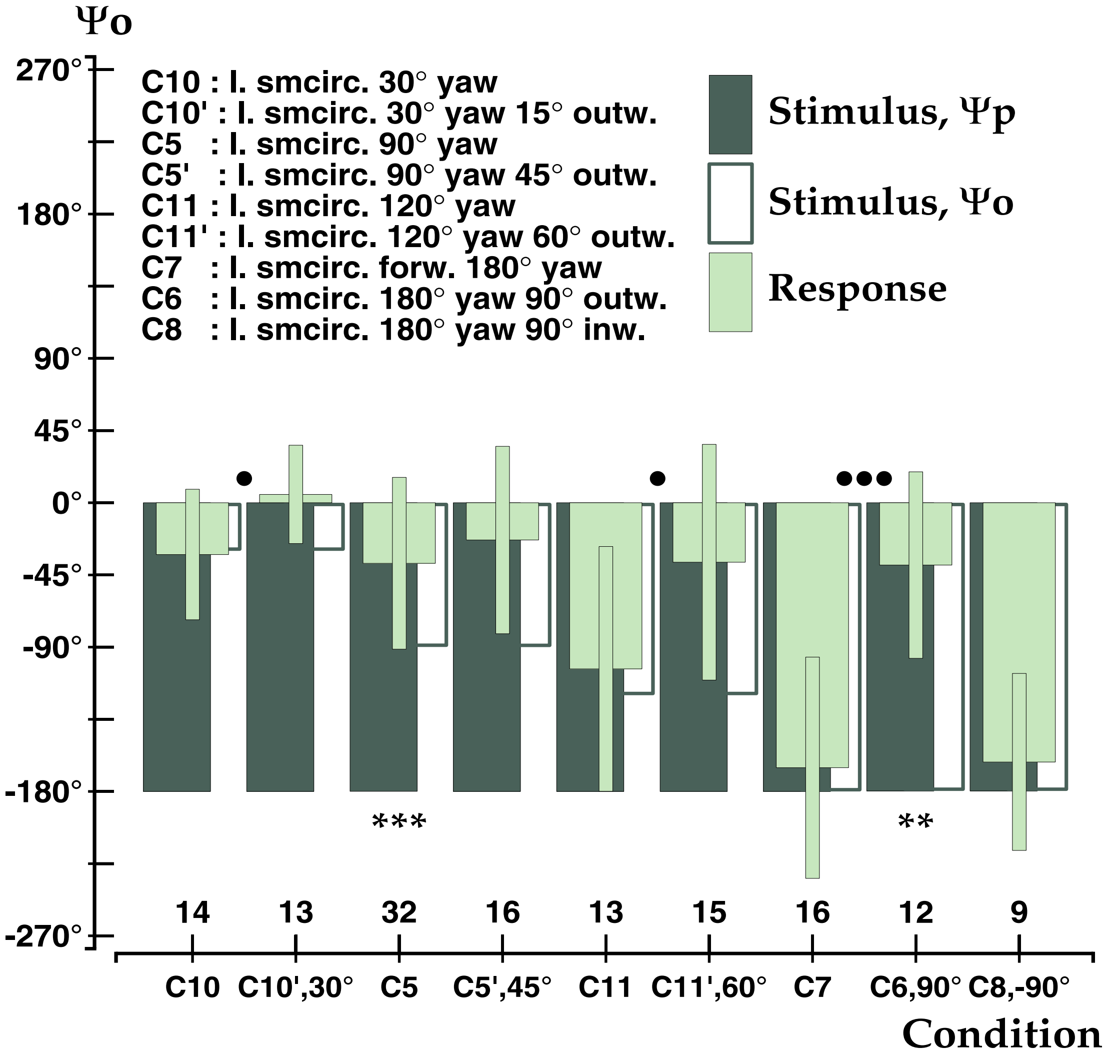
Ψo
270°
180°
90°
45°
0°
-45°
-90°
-180°
-270°
C10 : l. smcirc. 30° yaw
C10' : l. smcirc. 30° yaw 15° outw.
C5 : l. smcirc. 90° yaw
C5' : l. smcirc. 90° yaw 45° outw.
C11 : l. smcirc. 120° yaw
C11' : l. smcirc. 120° yaw 60° outw.
C7 : l. smcirc. forw. 180° yaw
C6 : l. smcirc. 180° yaw 90° outw.
C8 : l. smcirc. 180° yaw 90° inw.
Stimulus, Ψp
Stimulus, Ψo
Response
***
**
14 13 32 16 13 15 16 12 9
C10 C10',30° C5 C5',45° C11 C11',60° C7 C6,90° C8,-90°
Condition

Figure 2b

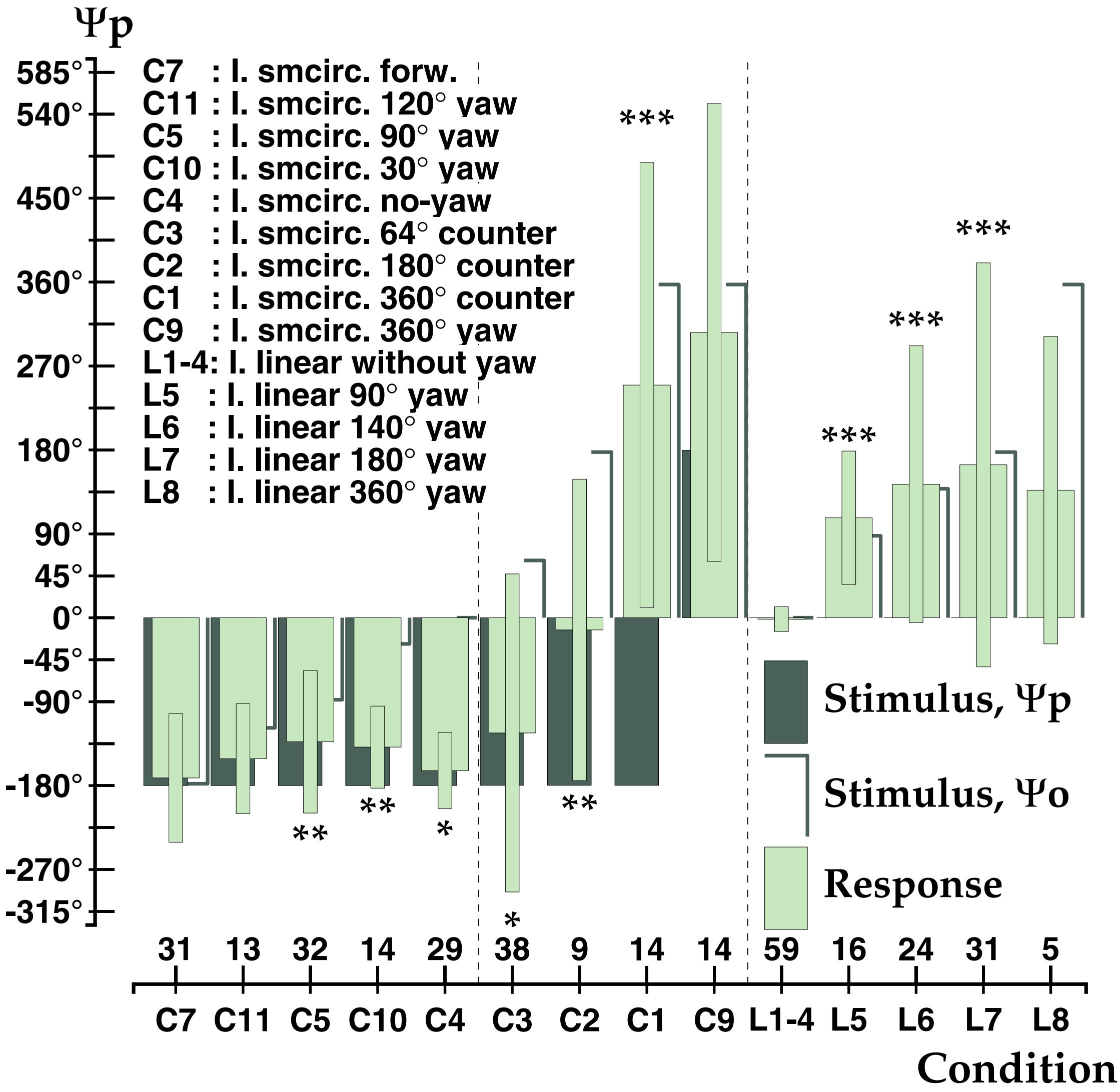

Ψp
585°
540°
450°
360°
270°
180°
90°
45°
0°
-45°
-90°
-180°
-270°
-315°
C7 : l. smcirc. forw.
C11 : l. smcirc. 120° yaw
C5 : l. smcirc. 90° yaw
C10 : l. smcirc. 30° yaw
C4 : l. smcirc. no-yaw
C3 : l. smcirc. 64° counter
C2 : l. smcirc. 180° counter
C1 : l. smcirc. 360° counter
C9 : l. smcirc. 360° yaw
L1-4: l. linear without yaw
L5 : l. linear 90° yaw
L6 : l. linear 140° yaw
L7 : l. linear 180° yaw
L8 : l. linear 360° yaw
***
***
***
***
**
**
*
*
**
Stimulus, Ψp
Stimulus, Ψo
Response
31 13 32 14 29 38 9 14 14 59 16 24 31 5
C7 C11 C5 C10 C4 C3 C2 C1 C9 L1-4 L5 L6 L7 L8
Condition

Figure 3a

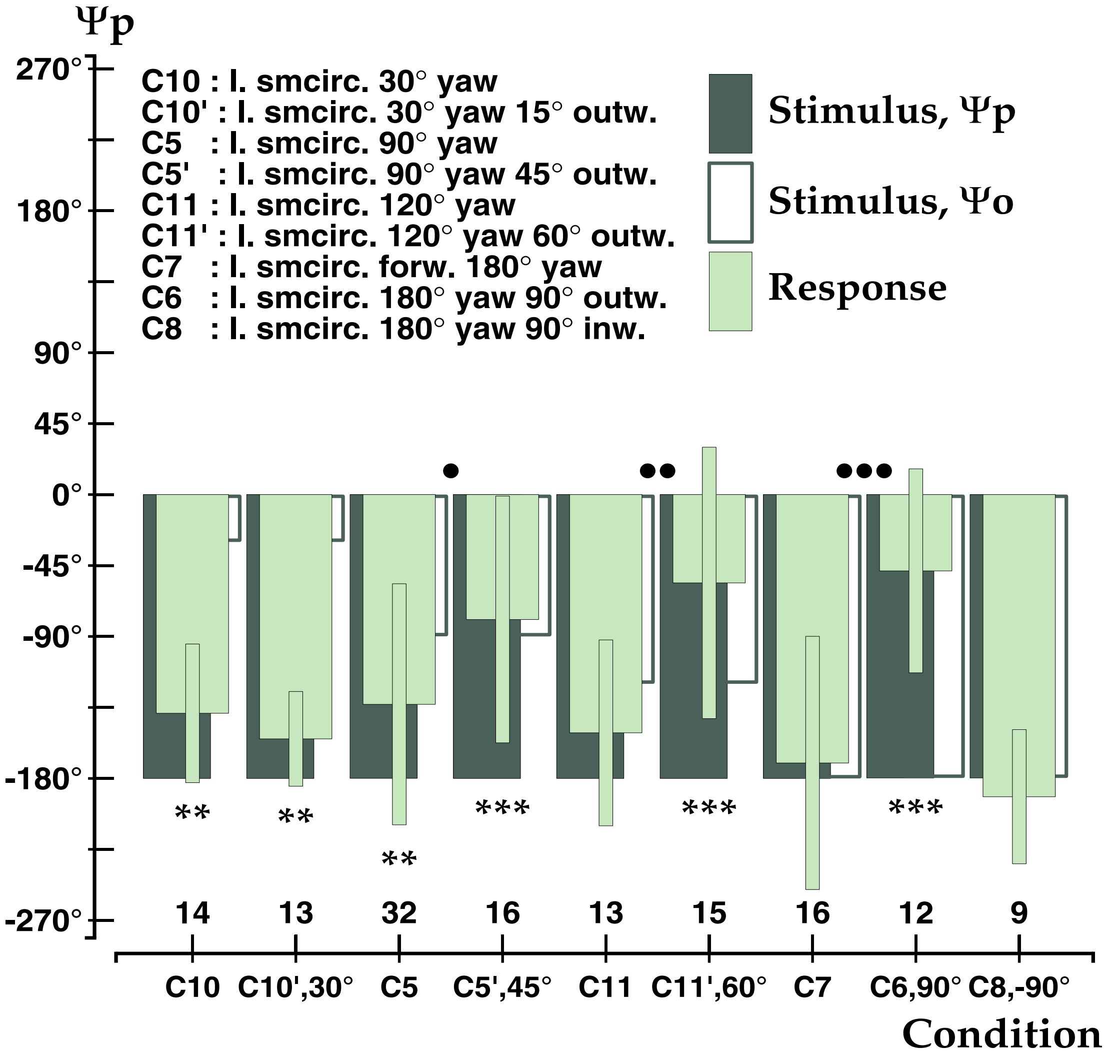
Ψp
270°
180°
90°
45°
0°
-45°
-90°
-180°
-270°
C10 : l. smcirc. 30° yaw
C10' : l. smcirc. 30° yaw 15° outw.
C5 : l. smcirc. 90° yaw
C5' : l. smcirc. 90° yaw 45° outw.
C11 : l. smcirc. 120° yaw
C11' : l. smcirc. 120° yaw 60° outw.
C7 : l. smcirc. forw. 180° yaw
C6 : l. smcirc. 180° yaw 90° outw.
C8 : l. smcirc. 180° yaw 90° inw.
Stimulus, Ψp
Stimulus, Ψo
Response
** ** ** *** *** ***
14 13 32 16 13 15 16 12 9
C10 C10',30° C5 C5',45° C11 C11',60° C7 C6,90° C8,-90°
Condition

Figure 3b

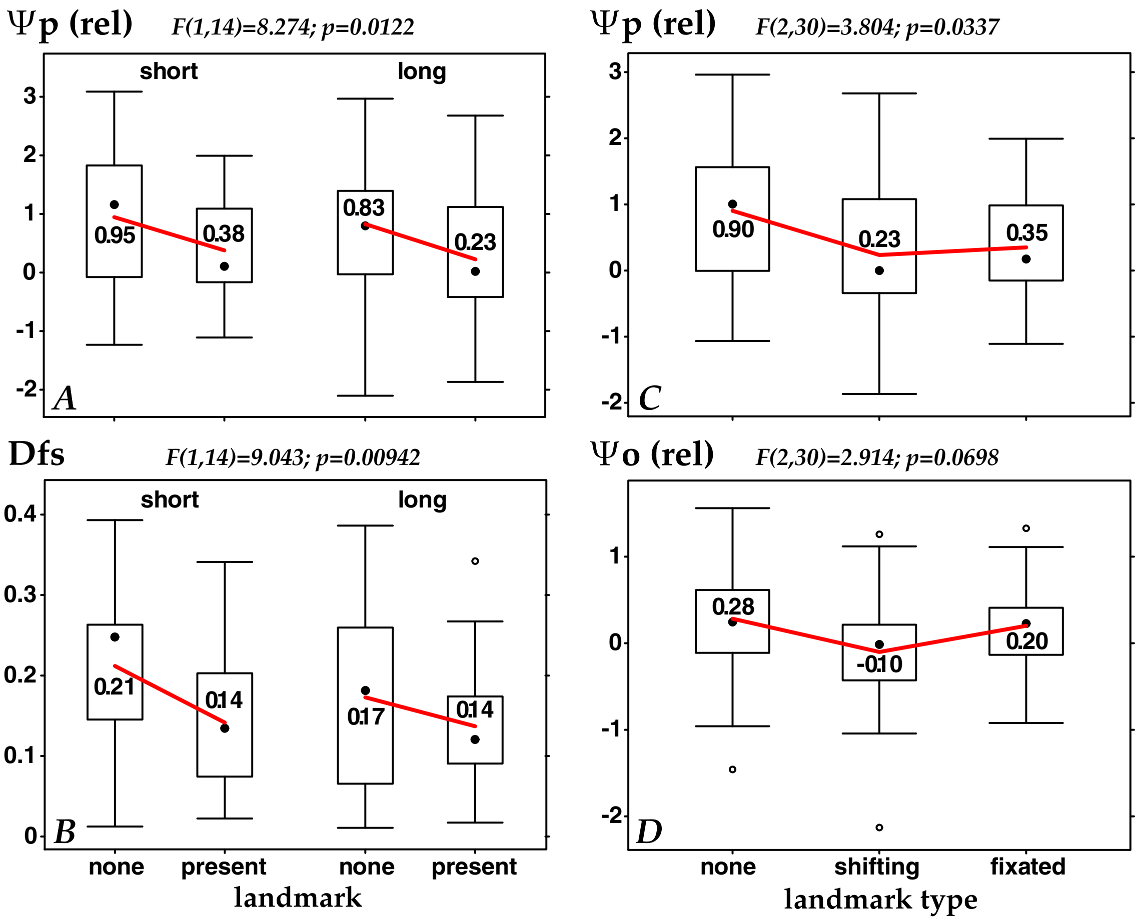
Ψp (rel)
F(1,14)=8.274; p=0.0122
short
long
0.95
0.38
0.83
0.23
A
Ψp (rel)
F(2,30)=3.804; p=0.0337
0.90
0.23
0.35
C
Dfs
F(1,14)=9.043; p=0.00942
short
long
0.21
0.14
0.17
0.14
B
none
present
none
present
landmark
Ψo (rel)
F(2,30)=2.914; p=0.0698
0.28
-0.10
0.20
D
none
shifting
fixated
landmark type

Figure 4