

RESEARCH ARTICLE

The roles of visual parallax and edge attraction in the foraging behaviour of the butterfly *Papilio xuthus*

Finlay J. Stewart*, Michiyo Kinoshita and Kentaro Arikawa

ABSTRACT

Several examples of insects using visual motion to measure distance have been documented, from locusts peering to gauge the proximity of prey, to honeybees performing visual odometry en route between the hive and a flower patch. However, whether the use of parallax information is confined to specialised behaviours like these or represents a more general purpose sensory capability, is an open question. We investigate this issue in the foraging swallowtail butterfly *Papilio xuthus*, which we trained to associate a target presented on a monitor with a food reward. We then tracked the animal's flight in real-time, allowing us to manipulate the size and/or position of the target in a closed-loop manner to create the illusion that it is situated either above or below the monitor surface. Butterflies are less attracted to (i.e. slower to approach) targets that appear, based on motion parallax, to be more distant. Furthermore, we found that the number of abortive descent manoeuvres performed prior to the first successful target approach varies according to the depth of the virtual target, with expansion and parallax cues having effects of opposing polarity. However, we found no evidence that *Papilio* modulate the kinematic parameters of their descents according to the apparent distance of the target. Thus, we argue that motion parallax is used to identify a proximal target object, but that the subsequent process of approaching it is based on stabilising its edge in the 2D space of the retina, without estimating its distance.

KEY WORDS: Parallax, Vision, Motion, Flight, Virtual reality, Butterfly

INTRODUCTION

Virtually all animals possessing eyes are faced with the issue of how to reconstruct the three-dimensional layout of their surroundings from the two-dimensional images formed on their retinas. Vertebrates like ourselves solve this problem using, amongst other things, stereopsis (i.e. comparing the images created in the two eyes) and accommodation (i.e. proprioception of the adjustment of the lens required to bring objects into focus). These approaches are generally not suitable for insects because their compound eyes typically exhibit little binocular overlap, poor spatial resolution and fixed-focus optics. Given these constraints, a more viable alternative is to use motion cues (van Breugel et al., 2014): when the animal is performing a translational movement, nearby objects will move more quickly on the retina than distant ones and hence depth can be estimated.

Many insect species measure translational optic flow to regulate their flight speed and/or distance to obstacles (e.g. Fry et al., 2009; Srinivasan et al., 1991). Indeed, the relationship between these

variables means that smooth landings can be achieved using a very simple control scheme (Baird et al., 2013; Chahl et al., 2004; Franceschini et al., 2007). By assuming a roughly constant distance between their eyes and the ground, central-place foragers such as bees and ants can perform odometry by integrating motion parallax over time (Esch et al., 2001; Ronacher and Wehner, 1995). However, in this study we are concerned with the converse process to visual odometry: using visual motion to gauge the distance of objects, rather than to estimate one's own movement.

Honeybees can be trained to select artificial flowers of variable size at particular depths below a Perspex screen, but they fail if the objects offer no green contrast with the background, implying that their colour-blind motion vision channel is essential for this behaviour (Srinivasan et al., 1989). They can also recognise the shape of camouflaged objects (i.e. textured shapes against a similarly textured background) by detecting their edges using discontinuities in optic flow (Zhang et al., 1995) and will approach a camouflaged raised platform differently to a camouflaged hole, attempting to land on the outer edge of the former and pass through the latter (Lehrer and Srinivasan, 1993). There is evidence that bees engage in 'active vision', structuring their flight to optimally extract motion parallax information. Flying honeybees perform lateral 'peering' manoeuvres whilst stabilising their gaze when approaching a feeder (Boeddeker and Hemmi, 2010) and, similarly, bumblebees move in such a way as to maximise 'pivoting parallax' when performing learning flights around the nest (Riabinina et al., 2014).

Although these examples amply illustrate that motion parallax is a valuable cue for motor control and navigation, it is difficult to unambiguously demonstrate that insects use it to measure target distance without specifically manipulating parallax cues. This requires that the target is moved with respect to the changing viewpoint of the animal in a 'closed-loop' manner. Perhaps the earliest insect experiment of this kind was a pioneering study of locusts (Wallace, 1959). These animals perform a stereotyped peering behaviour (whilst standing otherwise still) before pouncing on prey. By manually moving the target with or against these peering movements, the locust could be induced to over- or under-shoot with its subsequent jump, respectively. This finding was later confirmed using more modern techniques (Sobel, 1990; reviewed in Kral and Poteser, 1997). Another closed-loop experiment demonstrated that fruit flies freely walking on a platform are more attracted to targets that give the impression via parallax cues of being closer (Schuster et al., 2002).

Obviously, performing experiments of this kind on flying insects is rather more challenging. In recent years, 'virtual reality' techniques for presenting closed-loop stimuli to 3D-tracked animals have emerged (Fry et al., 2008; Straw et al., 2010). To our knowledge, this method has not yet been used to investigate the role of visual motion in target approach and landing, so, in this study, we adopt this approach to ask whether butterflies use motion cues in addition to colour and brightness information when foraging. Specifically, we consider the

Department of Evolutionary Studies of Biosystems, School of Advanced Sciences, Sokenkai (The Graduate University for Advanced Studies), Shonan Village, Hayama, Kanagawa 240-0193, Japan.

*Author for correspondence (stewart@soken.ac.jp)

List of abbreviations

<i>d</i>	virtual target depth
EO	expansion only
EP	expansion plus parallax
PO	parallax only
S	static (control)
<i>x</i> –	negative apparent depth, i.e. a virtual target above the surface (where <i>x</i> is one of EO, PO, EP)
<i>x</i> +	positive apparent depth, i.e. a virtual target below the surface

Japanese yellow swallowtail *Papilio xuthus*, an animal noted for its elaborate retinal organisation (Arikawa, 2003). We train the butterflies to associate food reward with coloured disks; this paradigm was used a century ago to demonstrate colour vision in bees (von Frisch, 1914), but has been employed more recently to investigate phenomena such as colour constancy, colour contrast and polarisation discrimination in *Papilio* (reviewed in Kinoshita and Arikawa, 2014). Of particular relevance to the current study is the finding that intensity contrast is required for landing. *Papilio* can detect and approach, but not land on, coloured targets that offer only chromatic contrast with their equiluminant grey backgrounds (Koshitaka et al., 2011). This suggests that foraging behaviour is composed of several somewhat distinct processes that utilise different types of visual information, with visual edges playing a crucial role in the final stages of target approach.

We adapt the classical paradigm by using targets displayed on a monitor rather than paper ones. By tracking the animals' flight in real time, we manipulate the stimulus to generate motion cues corresponding to a virtual target at some position other than the monitor surface. A benefit of this approach is that it allows us to independently control different aspects of the motion stimulus, namely expansion and parallax. Three particular questions we address are whether the apparent distance of targets affects (1) their attractiveness, (2) the spatio-temporal structure of the descents the butterfly makes towards them, and (3) their likelihood of eliciting successful versus abortive approaches.

RESULTS

Animals were trained for 5 days to feed from a blue circle displayed in the centre of a monitor placed on the floor of the experimental cage, before being tested on an unrewarded target (Fig. 1A). We define three closed-loop conditions: 'expansion only' (EO), where the target remains in the centre of the monitor but changes its diameter to give the impression from the butterfly's viewpoint of a constant-size object either above or below the monitor; parallax only (PO), where the target's size on the monitor is constant but its position changes to create the illusion of depth; and the more realistic 'expansion+parallax' (EP), where both manipulations occur in tandem (Fig. 1B,C; supplementary material Movie 1). Additionally, we define the open-loop control condition 'static' (S) where the target is fixed, as in the training sessions. Fig. 2 shows a sample flight trajectory recorded in the S condition.

The closed-loop manipulation of the target's appearance has no obvious disruptive effect on gross flight behaviour; repeated-measures ANOVAs detect neither significant effects of closed-loop condition (i.e. S, EO, PO or EP) nor of the target's apparent depth on either mean horizontal flight speed (condition: $P=0.64$, d.f.=3,123; depth: $P=0.28$, d.f.=1,41; supplementary material Fig. S1A) or median altitude (condition: $P=0.45$, d.f.=3,123; depth: $P=0.37$, d.f.=1,41; supplementary material Fig. S1C). There are, however, pronounced differences between individuals in terms of these

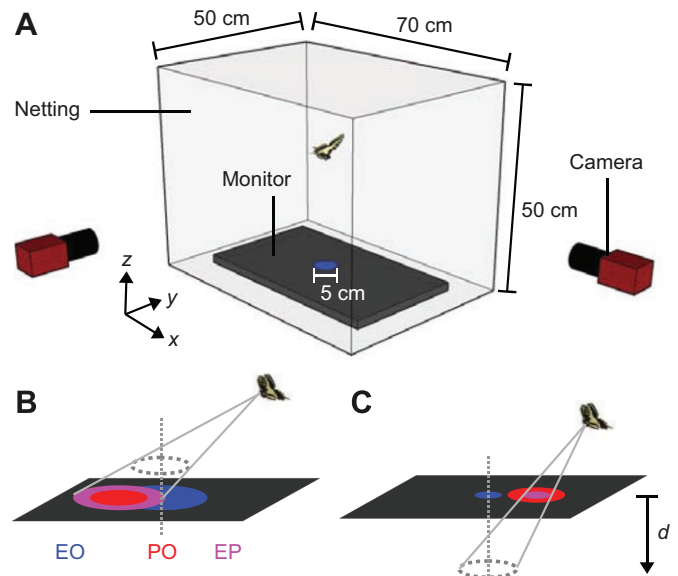


Fig. 1. Experimental set-up. (A) Schematic diagram of the experimental arena. (B,C) Illustration of the closed-loop stimuli generated for a virtual target above (B) or below (C) the monitor. Colours denote the closed-loop condition; the actual target is always blue. The expansion-plus-parallax (EP) stimulus is the projection of the virtual target to the monitor surface with respect to the butterfly's position; the expansion-only (EO) and parallax-only (PO) stimuli hold the position and size of the target on the monitor constant, respectively.

metrics (main effect of individual on flight speed: $P<0.0001$; altitude: $P<0.0001$; d.f.=41,531).

Attraction to apparently closer objects

The animals tended to fly near the ceiling of the cage when no target was displayed on the monitor. In six 30 s control tests (in different individuals) with a blank screen, only one instance of flight below 15 cm altitude was recorded (median minimum altitude: 29.8 cm, supplementary material Fig. S2). Therefore, we consider descents to below 15 cm to represent attempts to approach the target. To quantify the attractiveness of the stimulus, we measured the time between its onset and the animal's first foray below this altitude (Fig. 3A). A significant main effect of apparent depth exists ($P=0.0029$, d.f.=1,41), such that 'closer' targets are approached sooner. This analysis includes the S condition (apparent depth=0 cm), but a similar result is obtained if we consider only those cases where the appearance of the target is manipulated ($P=0.0064$, d.f.=1,41). Continuing to exclude S, neither a main effect of closed-loop condition ($P=0.71$, d.f.=2,82) nor an interaction between closed-loop condition and apparent depth ($P=0.38$, d.f.=2,82) was found. However, performing separate ANOVAs for each closed-loop condition revealed significant relationships between apparent depth and time to approach in those involving parallax manipulation (PO: $P=0.026$; EP: $P=0.025$; d.f.=2,82), but not in EO ($P=0.56$).

To confirm that this effect could be attributed to the motion of the target from the butterfly's perspective, we repeated the experiment using a 'yoked-loop' set-up, where the animal was played a recording of the stimulus generated by a previous individual's closed-loop flight, i.e. a target whose size and/or position fluctuates with no relation to the current animal's behaviour. No main effect of 'apparent depth' (of course, the illusion was then broken) and time to approach existed ($P=0.27$, d.f.=1,22; Fig. 3B). More importantly, analysing both data sets together (excluding the S condition)

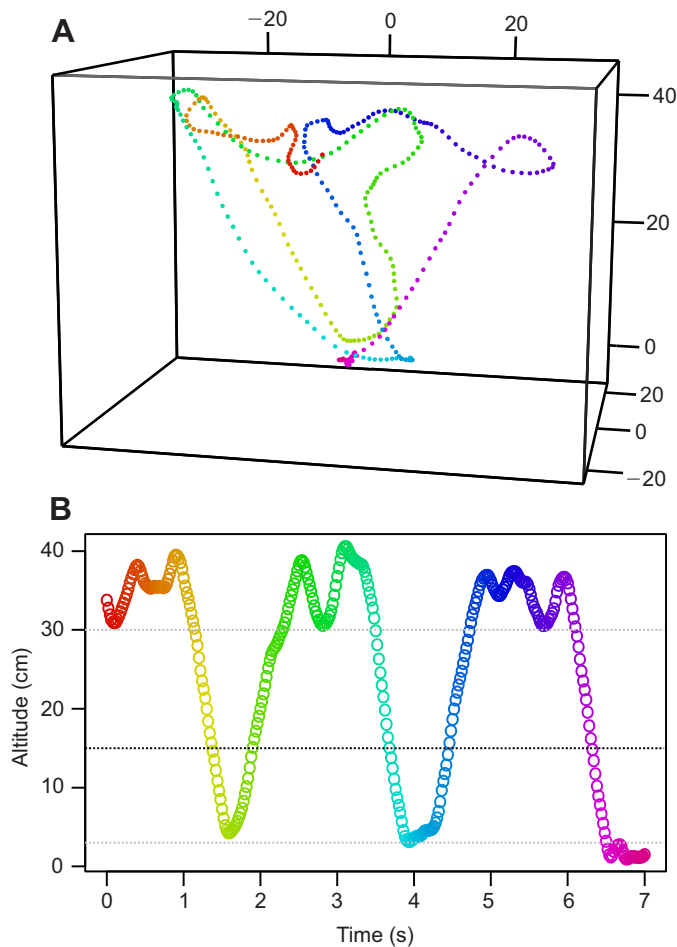


Fig. 2. Sample flight trajectory of a *Papilio xuthus* butterfly. (A) 3D view of 7 s of flight in the experimental arena, from the static control (S) condition. Data is smoothed (Gaussian average, s.d.=30 ms) and plotted at 20 ms intervals; colour spectrum indicates sequence (starting with red). Axes show measurements in cm. (B) Altitude plotted against time. Black dotted line shows the 15 cm threshold for identifying approaches; the time to approach is 1.3 s in this instance. Grey dotted lines are the thresholds for classifying descents. The first two descents (starting at $t \approx 1, 3.5$) fail to reach the 3 cm level (minima of 4.3 and 3.1 cm respectively), so are deemed abortive; the third manoeuvre ($t \approx 6$), ending with a landing, is identified as the first descent for analysis.

revealed a significant interaction between paradigm (i.e. closed-versus yoked-loop) and apparent depth ($P=0.013$, d.f.=1,774).

To establish that flight behaviour was broadly similar between paradigms, we measured flight speed and altitude in the yoked-loop conditions, as we did previously for closed-loop conditions (supplementary material Fig. S1B,D). Because these experiments were performed on different individuals, and results from one individual in different trials cannot be considered statistically independent, we calculated means for each animal across all 12 EO, PO and EP trials. There was no difference in mean horizontal flight speed ($P=0.54$, d.f.=63, Mann–Whitney *U*-test), but a significantly higher median altitude was observed in the yoked-loop paradigm ($P=0.039$), which probably reflects the lower overall attractiveness of the yoked-loop targets (Fig. 3). The S condition was ostensibly identical in both paradigms, yet we observed a significant difference in time to approach ($P=0.0007$, d.f.=63, Mann–Whitney *U*-test). In addition to the aforementioned large individual differences, this may be caused by the differing experience of the previous trials in the session. The possibility of such a learning effect is a troublesome

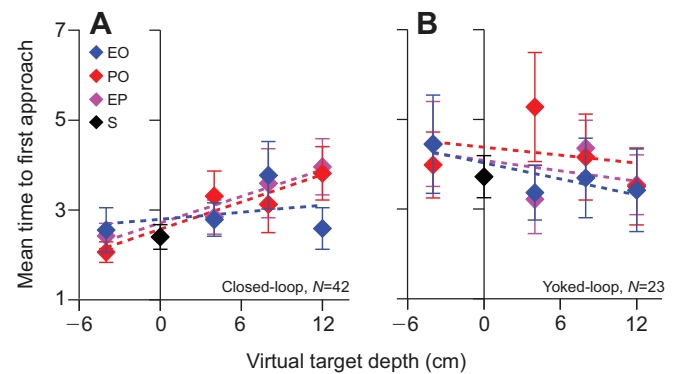


Fig. 3. Time taken to approach the target in the experimental arena by *Papilio xuthus* butterflies. (A) The y-axis shows the time from stimulus onset to the animal first reaching an altitude <15 cm, in the three closed loop conditions, plus the static control (depth=0). (B) Equivalent data for individuals in the yoked-loop paradigm, i.e. being played back stimuli generated by the individuals in A. Dashed lines are linear regressions; error bars are ± 1 s.e. EO, expansion only; PO, parallax only; EP, expansion and parallax; S, static control.

confounding factor that must be borne in mind when interpreting our results. However, as the order of trials within a session was randomised, it should introduce no systematic bias.

Descent kinematics and apparent distance

One might expect the animals to modulate the speed and/or steepness of their approach according to the perceived distance of the target, such that objects apparently further below the monitor would elicit more violent contacts with the surface. To assess whether this is the case, we identified the first descent in each trial that reached an altitude of less than 3 cm (Fig. 2B). This rather strict criterion was employed to exclude abortive approaches (or ‘fly-bys’; van Breugel and Dickinson, 2012) from our analysis. For each of these initial descents we measured the instantaneous downward velocity at 5 cm altitude; the downward velocity over the altitude interval 25 to 5 cm; the altitude of the animal at the temporal midpoint of this spatial range, to characterise deceleration; and the total horizontal trajectory length over that vertical interval, to characterise steepness (Fig. 4). Neither any main effects of apparent depth, nor any interactions between condition and apparent depth, were found for any of these metrics (supplementary material Table S1).

Frequency of abortive approaches

As Fig. 2B illustrates, not all descents resulted in successful landings. Fig. 5A shows the number of forays the animal made below 15 cm altitude prior to its first successful descent (i.e. reaching an altitude of less than 3 cm). A significant interaction between apparent depth and closed-loop condition was found ($P=0.034$, d.f.=2,82). As there was little difference in elicited behaviour between virtual targets 4, 8 or 12 cm below the surface, Fig. 5C pools the data according to the polarity of the apparent depth. In the EO condition, virtual targets above the surface (EO–) elicited significantly fewer abortive approaches than those below it (EO+) ($P=0.014$, d.f.=42, Wilcoxon signed-rank test), while the opposite pattern was observed in PO ($P=0.0090$). When both cues are present (EP), no significant difference was found ($P=0.18$). Furthermore, the number of abortive approaches in both EO– and PO+ was significantly lower than in S ($P=0.018$, 0.0086, respectively). This phenomenon was never observed in the yoked-

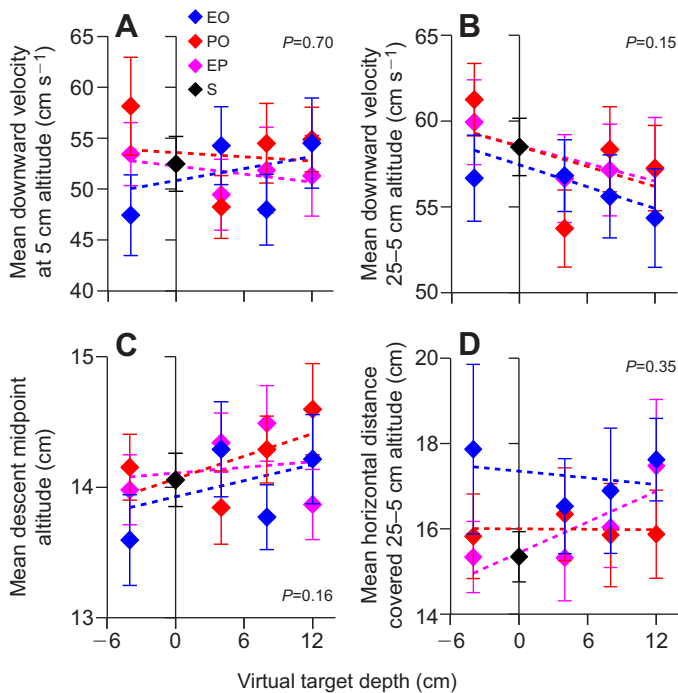


Fig. 4. Descent kinematics of *Papilio xuthus* butterflies in closed-loop conditions. In each trial, the first descent from >30 to <3 cm is identified and the following metrics are calculated: (A) Instantaneous downward velocity at 5 cm altitude. (B) Downward velocity averaged over the spatial interval 25 to 5 cm altitude. (C) Altitude of the animal at the temporal midpoint of its trajectory between 25 and 5 cm altitude; lower values mean greater deceleration, i.e. gentler landings. (D) Total horizontal trajectory length over the same vertical interval. Dashed lines show linear regressions; P -values are for main effects of virtual depth on the metric in question (repeated-measures ANOVA, d.f.=1,41); error bars are ± 1 s.e.; $n=42$. See supplementary material Table S1 for full statistical results.

loop paradigm (Fig. 5B,D; S versus PO+: $P=0.11$, d.f.=23), indicating that successful approaches are not facilitated by dynamic stimuli per se, although certain types of yoked-loop stimulus motion may hinder them: PO– elicited significantly more abortive approaches than PO+ ($P=0.015$, d.f.=23). Unlike the time to approach data, these results may therefore be attributed, at least partially, to the general nature of the target’s movement rather than the illusion created from the animal’s perspective by the closed-loop system.

Returning to the closed-loop paradigm, the absence of a significant difference in number of abortive approaches between EP– and EP+ suggests that the opposing effects of expansion and parallax may combine additively and thus cancel each other out. To further explore this possibility, we performed a separate experiment using ‘chimeric’ stimuli: E–P+, where the target gives expansion cues as though it were situated 4 cm above the monitor, but parallax cues as though it were 12 cm below; and E+P–, the converse case (Fig. 5E). The former target required significantly fewer attempts than the static control to approach successfully ($P=0.024$, d.f.=54, Wilcoxon signed-rank test), while the latter required more ($P=0.00085$, Fig. 5F).

Edge attraction

Although it is known that *Papilio* require intensity contrast to land on targets (Koshitaka et al., 2011) and in the present study, we observe that they frequently land close to and extend their proboscis towards the perimeter rather than the centre of the target, it has not

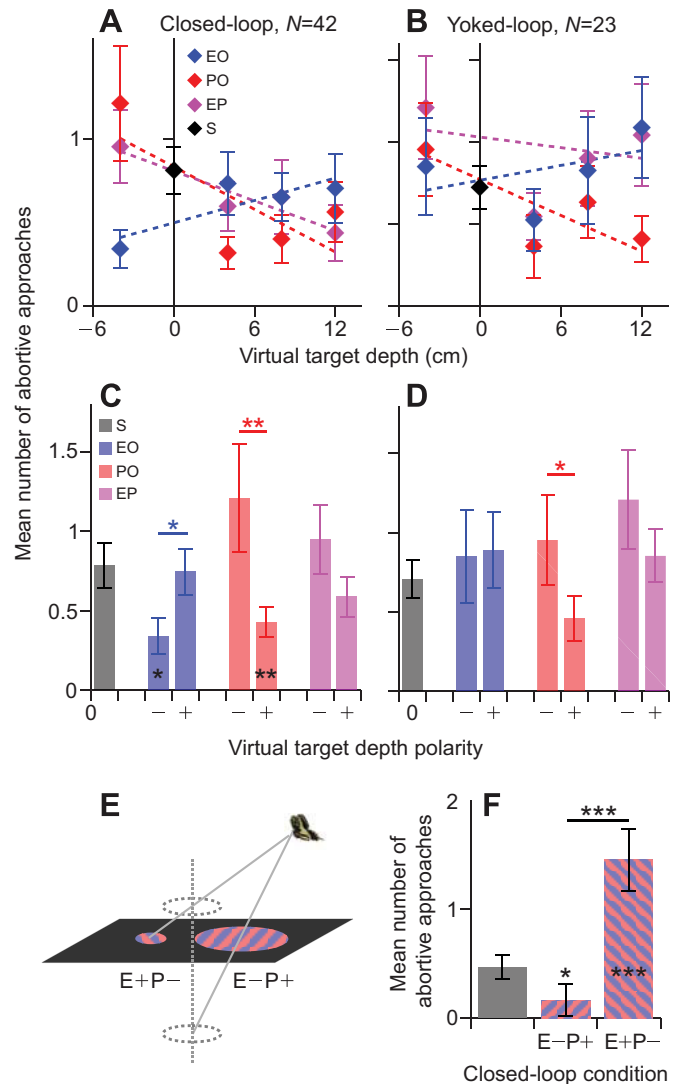


Fig. 5. Abortive approaches made towards target by *Papilio xuthus* butterflies. (A,B) The mean number of times the butterfly descends below 15 cm altitude prior to making its first successful approach. Dashed lines are linear regressions. (C,D) The same data as in A,B, but pooled according to whether the virtual target is above (–) or below (+) the monitor surface. Asterisks within bars denote differences from the S condition, while those over horizontal lines show comparisons between polarities of the same closed-loop condition. (E) Chimeric stimuli: E–P+ is a target that gives expansion cues as though it were at a depth of –4 cm, but parallax cues corresponding to a depth of 12 cm; E+P– is the converse. (F) Abortive approaches in the chimeric conditions, $N=54$. * $P<0.05$, ** $P<0.01$, *** $P<0.001$; all values calculated using Wilcoxon signed-rank test.

yet been established that they ‘aim for’ visual edges when landing. To confirm that this is the case, we presented trained individuals with a static asymmetrically blurred target (supplementary material Fig. S3). Landings were significantly biased towards the side with the steeper intensity gradient (Wilcoxon signed-rank test, $P=0.021$, d.f.=38).

Fig. 6 offers a dynamic systems analysis of approaches under various conditions, in terms of the elevation of the nearest edge of the target from the animal’s viewpoint and the instantaneous gradient of its descent (Fig. 6A). Successful trajectories in the S condition include many segments where both parameters are rather well matched at around 60 deg (Fig. 6B). Note that, in general, as the animal’s altitude decreases, the retinal position of the target edge

will tend to become increasingly unstable, corresponding to a leftward drift on the diagram. Fig. 6D shows a similar analysis of abortive approaches reaching a minimum altitude between 5 and 10 cm. Here, the bimodal distribution of trajectory segments reveals two qualitatively different failure modes, of which one example of each type is shown in Fig. 6C: undershooting, i.e. descending too steeply (cyan) and overshooting, i.e. failing to descend sufficiently steeply (blue). In Fig. 6E,F we consider two closed-loop conditions where few abortive approaches were observed (Fig. 5): EO(-4) and PO(+12). In both cases, the distributions are shifted rightward; the nature of these closed-loop manipulations means that the target edge will tend to remain in a ventral position rather than drifting towards the horizon as the animal descends.

A serious issue with this analysis is that we presented the stimuli on an LCD monitor, so contrast was attenuated (and colours distorted) at shallow viewing angles. To quantify this effect, we measured the contrast between the target and background from various angles, and found that it was relatively constant up to at least 40 deg, but then drops off rapidly (supplementary material Fig. S4). Thus, an elevation of ~45 deg may represent a ‘point of no return’ where the target becomes invisible and the approach cannot be recovered. This is reflected in the leftward directionality of almost all the arrows towards the extreme left of Fig. 6B,D.

DISCUSSION

Using a virtual reality system, we presented stimuli to free-flying *Papilio* that mimicked a (previously rewarded) target situated at some vertical displacement above or below the arena floor. The butterflies were slower to initiate approaches to targets the deeper they appeared to be, and this effect appears to be elicited primarily by parallax cues. We interpret this to mean that objects that move quickly across the retina (and thus are presumably close) are more attractive. The animal can perform this computation rather quickly; for targets of apparent depth -4 cm in the PO condition, the first quartile of time to approach is just 1.02 s from stimulus onset, which includes the time taken to descend to 15 cm altitude.

It is important to stress that this effect cannot be trivially explained by target size, as it occurs even when this parameter is (unrealistically) held constant, in the PO condition. Furthermore, if the butterfly were simply attracted to targets based on their proximity in the horizontal plane (either because closer targets subtend larger visual angles, or because the steeper viewing angle to the monitor maximises contrast) we would expect to see the opposite pattern (see Fig. 1B,C).

Contrary to our expectations, we find no clear evidence that the insects approach apparently shallower targets more slowly, with sharper deceleration, or at shallower angles; though of course, we cannot exclude the possibility that subtle effects on these kinematic metrics exist (see particularly Fig. 4B,C). However, we do find that the frequency of abortive approaches depends on the apparent depth of the target, and interestingly, the polarity of this effect is opposite for expansion and parallax cues. This effect is demonstrated particularly clearly using ‘chimeric’ stimuli, which give expansion and parallax cues corresponding to two different depths. The observed facilitation of successful approaches can be understood as the effect of the closed-loop system helping the animal to stabilise the edge of the target in a ventral retinal position.

Abortive approaches

On the one hand, it is rather remarkable that significantly fewer abortive descents are observed in the EO- and PO+ closed-loop conditions than the S control (Fig. 5C); why would a stationary

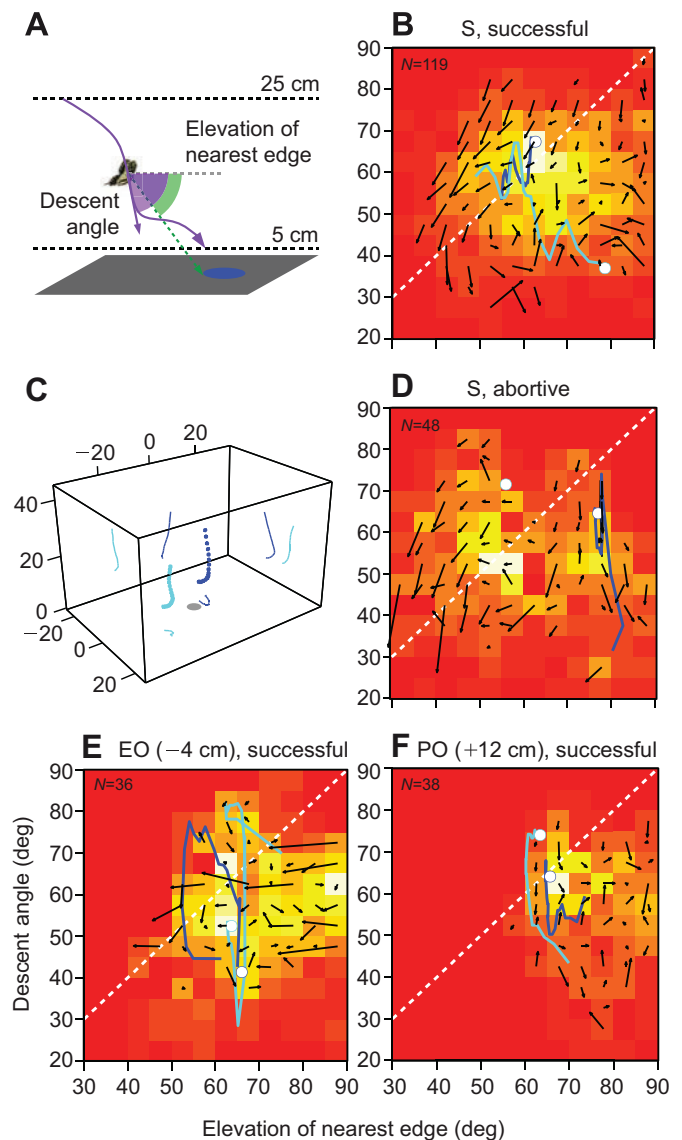


Fig. 6. Edge targeting of *Papilio xuthus* butterflies. (A) Approaches are analysed in terms of the vertical visual angle (i.e. elevation) of the nearest edge of the target, and the instantaneous gradient of the animal's descent; an angle of 90 deg corresponds to directly downwards. Trajectories are sampled at 1 cm altitude intervals from 25 to 5 cm. (B) Successful approaches (i.e. reaching an altitude <3 cm) in the S condition. Heat maps denote the relative frequency of trajectory segments falling into each bin. Arrows indicate the mean next trajectory segment, scaled by the square root of the number of segments in that bin; thus, long arrows correspond to large, repeatedly observed transitions. No arrow is displayed for bins containing fewer than five segments. Two sample trajectories are shown in blue and cyan in each plot, with the circle denoting the start (i.e. 25 cm altitude). The white dashed line indicates where the two parameters are equal. All points above this line correspond to undershooting trajectories if extrapolated, but those below it do not necessarily lead to overshoot; consider a spiral staircase-shaped descent and/or the possibility of landing on the far side of the target. Note that *n*-values refer to approaches, not individuals. (C) 3D plot of two sample abortive approaches in S. Points are plotted at 1 cm altitude intervals, with lines showing 2D projections to the back walls and floor. The target is indicated in grey; axis measurements in cm. (D) Abortive approaches in S reaching a minimum altitude between 5 and 10 cm, before any successful approaches have occurred. As these descents terminate before reaching 5 cm, they are composed of fewer segments. The sample trajectories shown in C are marked. (E) Successful approaches in EO with an apparent depth of -4 cm. (F) Successful approaches in PO with apparent depth 12 cm.

target with which the animal has had at least 5 sessions' experience be 'harder' to land on than a jittery, unfamiliar, dynamic one? However, when one considers the nature of these closed-loop manipulations, perhaps this result is not so surprising. In EO–, the target expands as the butterfly approaches; and in PO+, it moves to track the animal's flight. In both cases, the animal cannot miss; as long as it descends towards some point on the monitor, the target will come to it (albeit briefly in the case of EO–, as it disappears when the butterfly reaches 4 cm altitude). Viewed in these terms, it is somewhat impressive that the animal is able to approach the target at all in the PO– and especially E+P– conditions, as these represent the opposite case where the butterfly repels rather than attracts the target. Furthermore, the static target is more challenging to land on than an equivalent physical object, as the nature of the experimental apparatus means that its visibility will be attenuated if it is approached from shallow angles (supplementary material Fig. S4).

At the very least, these findings strongly imply that target approach is not the result of an open-loop, feed-forward motor program, but rather that the insect tracks the target in a continuous fashion. Prior to initiating an approach, when the animal's altitude is >30 cm, it experiences little difference between any of the closed-loop conditions, so approaches are presumably aborted as a result of sensory input received once the manoeuvre is under way. This issue could be investigated more directly by having the target abruptly vanish at a certain point during the approach (for an analogous experiment on walking flies, see Strauss and Pichler, 1998).

Edge fixation

We have based our analyses of the butterfly's flight control on the retinal position of the nearest edge of the target (Fig. 6). There are three reasons why we consider this to be a more salient feature than its centre. The first argument is logical: if one has the goal of contacting an object, then its edge – not its centre – defines the boundary between success and failure. Second, we found that when given a choice between sharp and gradual contrasts, butterflies display a clear bias towards landing on the former (supplementary material Fig. S3). The third line of evidence comes from the EO condition; as the position of the target is fixed in these trials, we would expect to see no difference in behaviour across apparent depths if the animal were targeting its centre. On the contrary, the frequency of abortive approaches is significantly higher in EO+ than EO– (Fig. 5C), which we attribute to the edge receding away from the animal during its descent. Note that the abortive approaches in EO+ cannot be explained simply by the target becoming too small to be clearly visible, as it still expands on the retina during the descent, albeit at a slower rate than a real object on the monitor surface would.

Descent velocity control

An obvious question is why *Papilio* do not use the apparent distance to the target to control their approach speed. One possible answer is that during training they have learnt the dimensions of the cage and thus 'know' how far they are from contact with the surface at all times. However, this account is difficult to reconcile with their increased attraction to apparently closer targets. We think a more plausible explanation is that the butterflies are physically robust enough that they simply have no need to land softly, at least over the range of flight velocities observed in this small arena. This may be a rather common feature of insect flight; when fruit flies land on a vertical post, over a third of landings are 'crashes' where the head or wing contacts the surface before the legs (van Breugel and Dickinson, 2012). This question could be clarified by investigating behaviours such as leg and

proboscis extension, which are performed in anticipation of imminent contact. Unfortunately, our present camera set-up lacks the necessary resolution for such an analysis.

Given this apparent disregard for the perceived distance of the target, together with the evidence for closed-loop edge tracking discussed above, it seems unlikely that *Papilio* truly plan their trajectories in three-dimensional space. Rather, our results are more consistent with them simply adjusting the steepness of their descent in order to stabilise the position of the target on the retina. The tendency of flying insects to fixate and land upon visual edges is well documented (e.g. van Breugel and Dickinson, 2012; Maimon et al., 2008). However, most studies of this kind have been concerned with the problem of stabilising vertical contrasts in the yaw dimension (but see Lehrer and Srinivasan, 1993). Our study provides evidence that a similar targeting process occurs towards edges on ventrally positioned horizontal surfaces. Like male blowflies pursuing potential mates by holding their retinal position and size constant (Boeddeker et al., 2003) or human baseball players catching fly balls by moving to keep their retinal trajectory linear (McBeath et al., 1995), this may represent a computational short-cut to reduce a complex but underconstrained three-dimensional problem to a simpler one in two-dimensional visual space.

Expansion versus parallax cues

If we consider only the EO condition (where parallax is not manipulated), there is no significant relationship between the attractiveness of the target and its apparent depth. However, it would be hasty to conclude that object expansion is not used to estimate distance; we quantify attractiveness by measuring the time taken to initiate an approach, but by definition, little expansion can be observed by the animal until it starts approaching the target. Having said this, it is reasonable to suppose that in general, expansion is a less-reliable cue than parallax in terms of signal-to-noise ratio. This is because the magnitude of the optic flow generated by an object at a certain distance is proportional to the sine of the angle between the object and the direction of translation, and probably explains why peering locusts and mantids sweep their heads laterally rather than back and forth.

The importance of distance perception for foraging

At first blush, foraging for flowers or other static, highly visible food sources seems to be a somewhat trivial task: as soon as a target is detected, approach it. Having acute vision to extend the range at which targets can be detected is presumably advantageous, but introduces a complication: some means of favouring closer targets is required if foraging efficiency is to be maintained. While retinal size could be used as a proxy for proximity, this may be unreliable as a result of the effects of orientation and partial occlusion. Worse still, such a system could be fooled by large distant objects that shared the same visual properties (e.g. colour, polarisation, brightness, etc.) as target flowers. However, as long as the targets are approximately stationary, motion cues provide a way to unambiguously estimate their distance. Despite finding themselves in a restricted, artificial situation where the target can never be more than a metre away, the butterflies in this study do appear to be performing a computation of this kind, suggesting that it represents a fundamental foraging strategy. It should be stressed, however, that without performing experiments using physical targets at various distances, it remains somewhat unclear how well our behavioural metrics such as time to approach correspond to real-world distance perception, or indeed whether *Papilio* can truly be said to possess such a capability.

From this and previous behavioural studies on *Papilio xuthus*, a detailed picture is emerging of the sensorimotor processes involved in foraging. First, the butterfly must acquire a target, which it presumably does based on spectral (and possibly polarisation) information; *Papilio* can discriminate colours of small targets observed at distances where they would be imaged by a single ommatidium (Takeuchi et al., 2006). Then, it gauges the distance using motion cues, and based on this information, decides whether to approach. Finally, once the target is close, it aims to land on an edge, requiring a contrast in intensity (Koshitaka et al., 2011). Far from being trivial, foraging flight in fact relies on the delicate interplay of a number of visual modalities and visuomotor control schemes.

MATERIALS AND METHODS

Animals

Laboratory-reared spring-form adult Japanese yellow swallowtail butterflies (*Papilio xuthus* Linnaeus) were used for all experiments. The laboratory culture was derived from eggs laid by females caught in Kanagawa, Japan. Larvae were fed fresh citrus leaves under a 10 h:14 h light:dark cycle. Pupae were stored at 4°C for at least 3 months before being allowed to emerge at room temperature. Between training or test sessions, animals were kept in polystyrene boxes at room temperature, with no more than three individuals to a box, all of the same sex. They were not fed other than during training or testing.

Experimental set-up

The experimental arena was a cuboid cage measuring 70×50×50 cm (length, width, height), consisting of a frame of aluminium tubes covered with plastic netting with 9 mm square holes (Fig. 1A). It was lit with six overhead xenon lamps, and care was taken to light the arena evenly at approximately 4000 lx at floor level. A 24" (61 cm) LCD monitor (Benq XL2420T) was placed flat on the arena floor, with a transparent acrylic sheet put directly on its surface to protect it. We define the centre of the monitor surface as the origin of our co-ordinate system, meaning that the ceiling's height was approximately 44.5 cm. The animals were tracked using two orthogonally oriented cameras (AVT Prosilica GE680, monochrome), which were synchronised by triggering them programmatically using a USB interface board (PhidgetInterfaceKit 8/8/8). The cameras each recorded 60 frames s⁻¹ with a 2.0 ms exposure at a resolution of 640×480 pixels (cropped to 500×480 for the camera oriented along the *x*-axis).

Procedure

During training, a 100% brightness blue circle of diameter 5 cm was displayed in the centre of the monitor, with several drops of 6% sucrose solution placed on the acrylic sheet over it. One box of butterflies at a time were released into the arena and allowed to feed *ad libitum*. Any individuals not spontaneously feeding were caught and had their proboscis manually extended onto the sucrose drops. This procedure was repeated for 5 days, starting 24–48 h post-emergence.

For testing, a single animal was released at a time. As before, it was allowed to feed from the target, but would be quickly shooed away in order to prevent satiety and thus maintain high foraging motivation. After three successful landings, the monitor was blanked, the sucrose solution wiped off, and the testing phase begun. Any animal failing to land would not be tested, but rather given a further training session with manual feeding. If failure continued for five attempted test sessions, the animal was discarded.

In the main experiment, each testing session consisted of fourteen 30 s trials presented in a random order: three closed-loop conditions (EO, PO and EP) at four apparent depths each [−4 (i.e. above the monitor), 4, 8 and 12 cm], plus two control (S) tests. If on any particular trial the animal failed to descend below 5 cm altitude, that trial was repeated. If this happened more than five times in the whole session, or if the animal stopped flying, the test was abandoned and a further training session given. A similar procedure was followed for the chimeric experiment, but with just five trials per session (S, E−P+, E+P−, and two other chimeric conditions not discussed in this paper).

Flight tracking

The cameras were calibrated by placing vertical posts in each corner of the monitor, giving eight points (top and bottom, ×4) in monitor-centric 3D space from which the projection matrices of the cameras could be fitted using a least-squares method.

The closed-loop system was controlled by a custom-written Java program. Immediately prior to the start of each trial, each camera captured 15 frames from which the pixel-wise median was taken to effectively erase the flying animal and obtain a background image of the empty arena. Once the trial was under way, each incoming frame was subtracted from the background to identify the dark region corresponding to the butterfly. This difference image was thresholded, dilated and the largest connected region found. The centroid of this was taken as the animal's position in the image plane; note that this could be ~2 cm from the head in the worst case. Together with the projection matrix, this defined a ray in 3D space along which the animal must have lain. The point at which the rays of the two cameras come closest to intersecting was calculated; if they did not pass within 2.5 cm of each other, a tracking error was deemed to have occurred and the animal's position was unknown, otherwise the 3D position was recorded and made available to the stimulus presentation process.

Stimulus display

The monitor was set to a resolution of 1024×768 pixels and a refresh rate of 60 Hz, to match the frequency of image acquisition. This is lower than the flicker fusion frequencies of *Papilio xuthus* photoreceptors, which are around 100 Hz (Nakagawa and Eguchi, 1994). However, filming the monitor at 1000 frames s⁻¹ revealed negligible flicker.

Every refresh cycle, the display was updated based on the last logged 3D position (*x*,*y*,*z*) of the butterfly. In all closed-loop conditions, the 'true' 3D position of the centre of the virtual target was (0,0,−*d*), where *d* was the apparent depth (Fig. 1B,C).

In expansion conditions, the 'true' radius *r*₀ of the hypothetical object was adjusted such that it would appear identical to the training image from an altitude of 30 cm:

$$r_0 = 2.5 \times \frac{30 + d}{30}. \quad (1)$$

Its on-screen radius as a function of the butterfly's altitude *z* was then:

$$r(z) = r_0 \times \frac{z}{z + d}. \quad (2)$$

In parallax conditions, the on-screen centre of the target (*p*_{*x*},*p*_{*y*}) was:

$$(p_x, p_y) = (x, y) \times \frac{d}{d + z}. \quad (3)$$

If the virtual target was above the monitor (*d*<0), then its appearance was undefined when the animal was below its altitude (*z*<−*d*), so no target was displayed. For this reason, virtual targets >4 cm above the monitor were not used, because the target would disappear long before the animal reached the surface, making comparisons between conditions difficult. In the case of chimeric stimuli (Fig. 5E), one value of *d* was used for Eqns 1 and 2, and a different one for Eqn 3.

In the yoked-loop paradigm, everything was done identically, except that instead of using the last logged 3D position of the animal, the corresponding time-point of a previously recorded closed-loop trajectory was used as the basis for generating the stimulus. Each yoked-loop animal was paired with a different closed-loop counterpart in a one-to-one fashion.

To quantify the latency of the closed-loop system, an LED was placed in the arena and one of the cameras was repositioned such that it imaged the monitor surface. The program was modified to display an indication on the monitor as soon the value of a pixel corresponding to the LED exceeded a threshold. Whilst the system was tracking a butterfly, the LED was lit. In the resulting recording, the delay between the LED being activated and the pattern appearing on the monitor was rather repeatably found to be three frames, corresponding to a latency of approximately 50 ms.

Asymmetric blur

To test edge attraction, a special static stimulus was prepared, based on an 8-cm-diameter blue circle. This was convolved with a 1D Gaussian blur in the x -dimension, where the standard deviation varied linearly according to the position on x -axis (supplementary material Fig. S3). The resulting image was then blurred again with a uniform 1D Gaussian in the y -dimension. This stimulus was designed so that the centre of mass remained at (0,0). For testing, the stimulus was presented twice, in mirror-image orientations about the y -axis, and the animal's trajectory was tracked in real-time as before. In each trial the landing location was taken as the (x,y) co-ordinates of the lowest point of the first descent that reached <2 cm altitude within a 9 cm radius of the centre of the monitor.

Analysis

The trajectories obtained from real-time tracking are rather jittery; in particular, an oscillation in the z -dimension due to the wing-stroke cycle is often apparent. Therefore, we smoothed the data by resampling it at 20 ms intervals using Gaussian averaging with s.d.=30 ms. Note that this is for analysis purposes only; the closed-loop stimulus presentation used the raw signal as averaging would increase latency.

When calculating time-to-approach, in the rare case that the butterfly was at an altitude <25 cm at the onset of the stimulus, the trial was excluded. Otherwise, we simply recorded the time at which its altitude first drops below 15 cm. For the analysis of descents, we identified the first manoeuvre in each trial where the butterfly moved from an altitude >30 to <3 cm, without climbing more than 2 cm, i.e. its altitude could never exceed the lowest point reached so far in the descent +2 cm. If no descents satisfying these criteria were found, that trial was excluded from the analysis.

A 'successful' approach was defined as one reaching a minimum altitude <3 cm. We term these manoeuvres 'approaches' rather than 'landings', because the butterflies frequently made only fleeting contact with the surface, particularly (but not exclusively) in those closed-loop conditions where the target disappears in the final stages of the descent. To ensure consistency across closed-loop conditions, the position of the animal in the horizontal plane was disregarded. Thus, in rare cases, a trajectory where the animal approached some part of the monitor distant from the target could be deemed successful. 'Abortive' approaches were those reaching a minimum altitude <15 cm but >3 cm, unless otherwise specified.

When performing ANOVAs, virtual target depth was treated as an interval type independent variable, whereas closed-loop condition was a categorical variable with four categories: EO, PO, EP and S (though S was often excluded).

Acknowledgements

We thank the anonymous reviewers for their very helpful comments.

Competing interests

The authors declare no competing or financial interests.

Author contributions

F.J.S.: experimental design, system development, performing of experiments, analysis, manuscript writing. M.K.: experimental design, manuscript editing. K.A.: experimental design, manuscript editing.

Funding

This work was supported in part by the JSPS (Japanese Society for Promotion of Science) Grants-in-Aid for Scientific Research to F.J.S. [26650116], M.K. [24570084] and K.A. [26251036].

Supplementary material

Supplementary material available online at <http://jeb.biologists.org/lookup/suppl/doi:10.1242/jeb.115063/-/DC1>

References

- Arikawa, K. (2003). Spectral organization of the eye of a butterfly, *Papilio*. *J. Comp. Physiol. A Sens. Neural Behav. Physiol.* **189**, 791–800.
- Baird, E., Boeddeker, N., Ibbotson, M. R. and Srinivasan, M. V. (2013). A universal strategy for visually guided landing. *Proc. Natl. Acad. Sci. USA* **110**, 18686–18691.
- Boeddeker, N. and Hemmi, J. M. (2010). Visual gaze control during peering flight manoeuvres in honeybees. *Proc. R. Soc. B Biol. Sci.* **277**, 1209–1217.
- Boeddeker, N., Kern, R. and Egelhaaf, M. (2003). Chasing a dummy target: smooth pursuit and velocity control in male blowflies. *Proc. R. Soc. B Biol. Sci.* **270**, 393–399.
- Chahl, J. S., Srinivasan, M. V. and Zhang, S. W. (2004). Landing strategies in honeybees and applications to uninhabited airborne vehicles. *Int. J. Robot. Res.* **23**, 101–110.
- Esch, H. E., Zhang, S., Srinivasan, M. V. and Tautz, J. (2001). Honeybee dances communicate distances measured by optic flow. *Nature* **411**, 581–583.
- Franceschini, N., Ruffier, F. and Serres, J. (2007). A bio-inspired flying robot sheds light on insect piloting abilities. *Curr. Biol.* **17**, 329–335.
- Fry, S. N., Rohrseitz, N., Straw, A. D. and Dickinson, M. H. (2008). TrackFly: virtual reality for a behavioral system analysis in free-flying fruit flies. *J. Neurosci. Methods* **171**, 110–117.
- Fry, S. N., Rohrseitz, N., Straw, A. D. and Dickinson, M. H. (2009). Visual control of flight speed in *Drosophila melanogaster*. *J. Exp. Biol.* **212**, 1120–1130.
- Kinoshita, M. and Arikawa, K. (2014). Color and polarization vision in foraging *Papilio*. *J. Comp. Physiol. A* **200**, 513–526.
- Koshitaka, H., Arikawa, K. and Kinoshita, M. (2011). Intensity contrast as a crucial cue for butterfly landing. *J. Comp. Physiol. A* **197**, 1105–1112.
- Kral, K. and Poteser, M. (1997). Motion parallax as a source of distance information in locusts and mantids. *J. Insect Behav.* **10**, 145–163.
- Lehrer, M. and Srinivasan, M. V. (1993). Object detection by honeybees: why do they land on edges? *J. Comp. Physiol. A* **173**, 23–32.
- Maimon, G., Straw, A. D. and Dickinson, M. H. (2008). A simple vision-based algorithm for decision making in flying *Drosophila*. *Curr. Biol.* **18**, 464–470.
- McBeath, M. K., Shaffer, D. M. and Kaiser, M. K. (1995). How baseball outfielders determine where to run to catch fly balls. *Science* **268**, 569–573.
- Nakagawa, T. and Eguchi, E. (1994). Differences in flicker fusion frequencies of the five spectral photoreceptor types in the swallowtail butterfly's compound eye. *Zool. Sci.* **11**, 759–762.
- Riabinina, O., de Ibarra, N. H., Philippides, A. and Collett, T. S. (2014). Head movements and the optic flow generated during the learning flights of bumblebees. *J. Exp. Biol.* **217**, 2633–2642.
- Ronacher, B. and Wehner, R. (1995). Desert ants *Cataglyphis fortis* use self-induced optic flow to measure distances travelled. *J. Comp. Physiol. A* **177**, 21–27.
- Schuster, S., Strauss, R. and Götz, K. G. (2002). Virtual-reality techniques resolve the visual cues used by fruit flies to evaluate object distances. *Curr. Biol.* **12**, 1591–1594.
- Sobel, E. C. (1990). The locust's use of motion parallax to measure distance. *J. Comp. Physiol. A* **167**, 579–588.
- Srinivasan, M. V., Lehrer, M., Zhang, S. W. and Horridge, G. A. (1989). How honeybees measure their distance from objects of unknown size. *J. Comp. Physiol. A* **165**, 605–613.
- Srinivasan, M. V., Lehrer, M., Kirchner, W. H. and Zhang, S. W. (1991). Range perception through apparent image speed in freely flying honeybees. *Vis. Neurosci.* **6**, 519–535.
- Strauss, R. and Pichler, J. (1998). Persistence of orientation toward a temporarily invisible landmark in *Drosophila melanogaster*. *J. Comp. Physiol. A Sens. Neural Behav.* **182**, 411–423.
- Straw, A. D., Lee, S. and Dickinson, M. H. (2010). Visual control of altitude in flying *Drosophila*. *Curr. Biol.* **20**, 1550–1556.
- Takeuchi, Y., Arikawa, K. and Kinoshita, M. (2006). Color discrimination at the spatial resolution limit in a swallowtail butterfly, *Papilio xuthus*. *J. Exp. Biol.* **209**, 2873–2879.
- van Breugel, F. and Dickinson, M. H. (2012). The visual control of landing and obstacle avoidance in the fruit fly *Drosophila melanogaster*. *J. Exp. Biol.* **215**, 1783–1798.
- van Breugel, F., Morgansen, K. and Dickinson, M. H. (2014). Monocular distance estimation from optic flow during active landing maneuvers. *Bioinspir. Biomim.* **9**, 025002.
- von Frisch, K. (1914). *Der farbennsinn und Formensinn der Biene*. Jena: Fischer.
- Wallace, G. K. (1959). Visual scanning in the desert locust *Schistocerca gregaria* Forskål. *J. Exp. Biol.* **36**, 512–525.
- Zhang, S. W., Srinivasan, M. V. and Collett, T. (1995). Convergent processing in honeybee vision: multiple channels for the recognition of shape. *Proc. Natl. Acad. Sci. USA* **92**, 3029–3031.

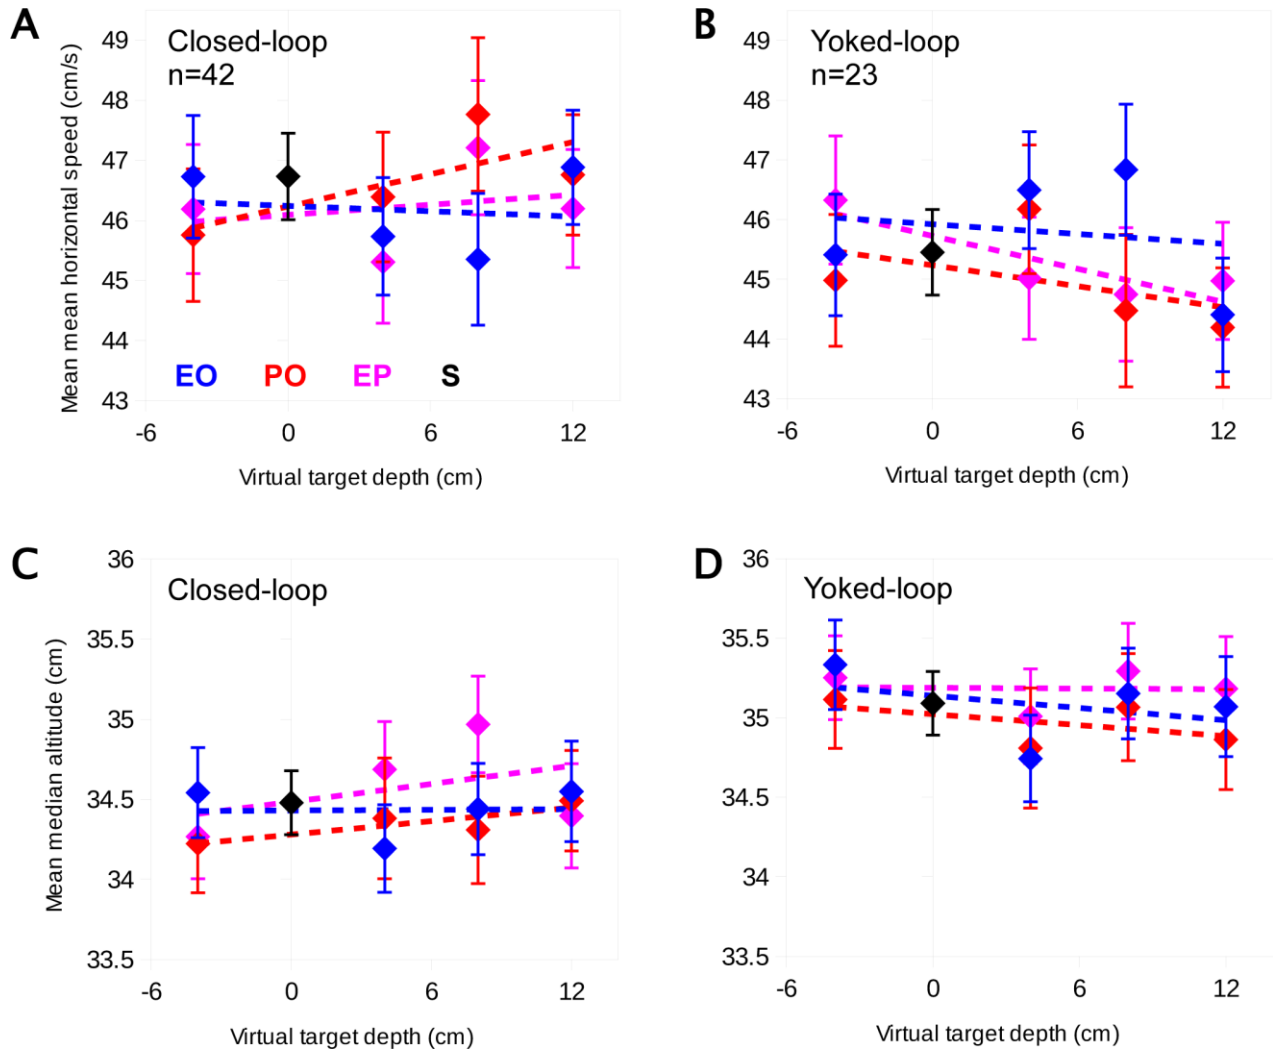


Fig. S1. Gross flight behaviour. (A,B) Mean horizontal speed in the closed-loop (A, n = 42) and yoked-loop (B, n = 23) paradigms. (C,D) Median altitude. In all cases, portions of the trajectory below 25cm altitude are ignored, to exclude approaches to the target. Rare occasions when the animal is stationary (e.g. hanging on the netting) are also excluded. Dashed lines show linear regressions; bars are ± 1 SE.

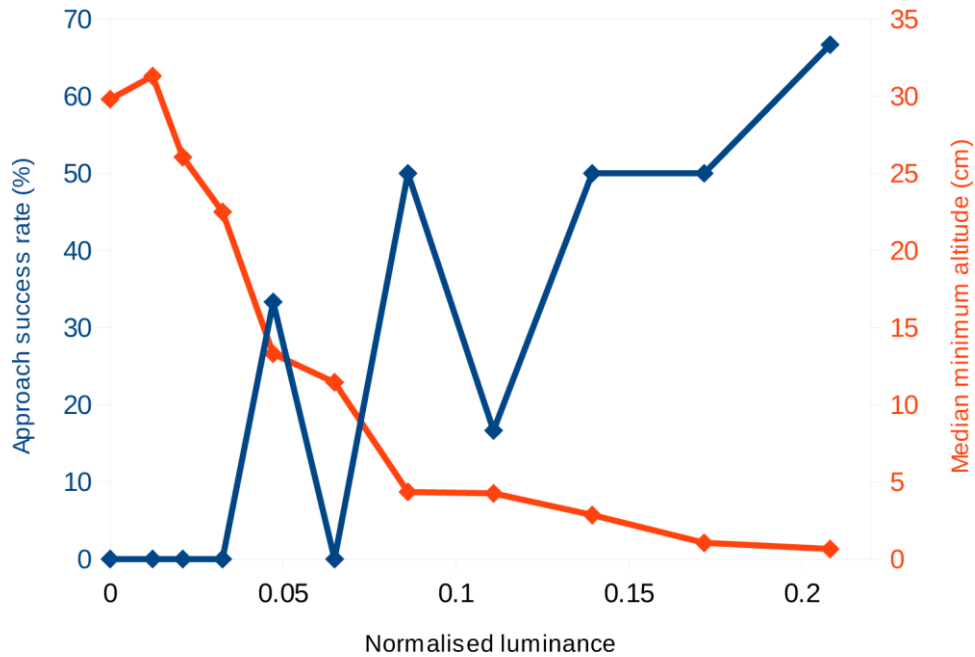


Fig. S2. Low intensity static targets. Luminance is expressed as a proportion of the full brightness target used in the main experiments, based on the gamma setting of the monitor (= 2.4). Note that luminance = 0 is a control condition with no target displayed. Blue line indicates the percentage of individuals that successfully approached the target. Orange line shows the median lowest altitude reached over the course of the 30 s trial. Unlike in the main experiments, trials were not repeated if the animal failed to approach the target first time. $n = 6$.




Fig. S3. Asymmetric blur experiment. Landing positions on the (static) asymmetric target, normalised for left/right polarity. $n = 38$ landings, 19 individuals.

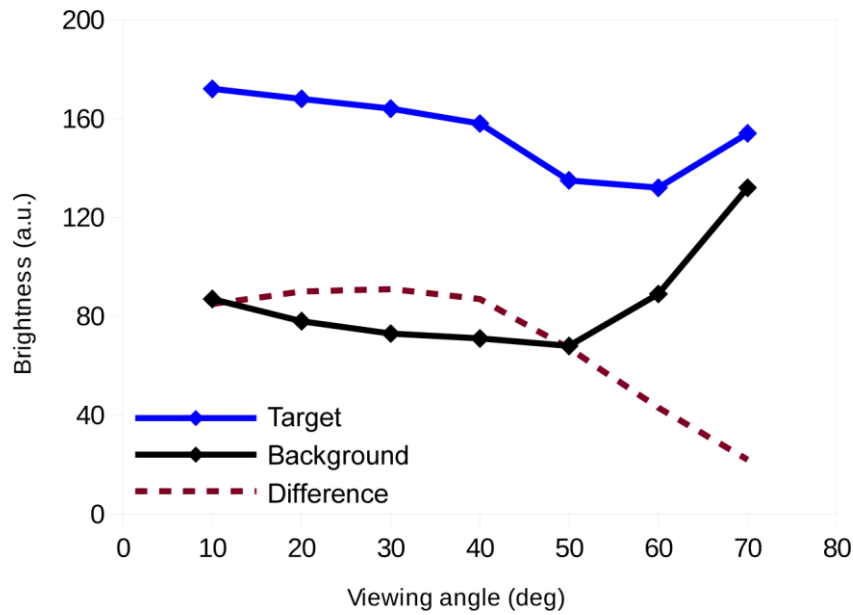



Fig. S4. Viewing angle effects on contrast. Using a camera with fixed shutter, iris, and digital gain settings, we measured the brightness of the target and adjacent black background from various angles. The cage was removed but otherwise conditions were identical to those used for experiments, i.e. the xenon floodlights were on and the acrylic sheet present on the monitor. Measurements were taken along the long axis of the arena (i.e. horizontal w.r.t. the monitor), as this is the direction from which the animals will tend to experience the shallowest angles. The sharp increase in background brightness for angles $> 50^\circ$ is largely due to reflection from the acrylic sheet. No measurement is reported for 0° due to the unavoidable problem of the camera's shadow and reflection contaminating the image.



Movie 1. Closed-loop stimulus. Overhead view of a trained butterfly in the EP condition, virtual depth = 12 cm. Video begins 1 s prior to stimulus onset. Reflected in the acrylic sheet (smaller rectangle within the monitor), the smartphone used to record the footage can be seen resting on the top of the cage; obviously this is not present during experiments.

Table S1. Closed-loop descent kinematics statistics			
Metric (dependent variable)	p		
	Apparent depth (main effect) df = (1, 41)		Apparent depth x closed-loop condition (interaction) df = (2, 82)
	incl S	excl S	excl S
A Downward velocity, 5 cm	.70	.61	.47
B Downward velocity, 25 ~ 5 cm	.15	.19	.72
C Descent midpoint altitude	.16	.13	.77
D Horizontal distance	.35	.61	.41

Full results of the ANOVAs performed for the data shown in Fig. 4; refer to it for details of the metrics used (bold letters correspond to the figure). When testing for main effects of apparent depth on the metrics in question, the control condition S can either be included with depth = 0, or excluded; for interactions between apparent depth and closed-loop condition, it must be excluded.