

AFRL-AFOSR-UK-TR-2015-0006



**Air-to-air and air-to-ground attack strategies in trained birds
of prey**

**Prof. Graham K. Taylor, Prof. Adrian L. R. Thomas,
and Ms. Caroline H. Brighton**

**University of Oxford
Department of Zoology
South Parks Road
Oxford, OX1 3PS United Kingdom**

EOARD GRANT #FA8655-11-1-3065

Report Date: September 2014

Final Report from 1 October 2011 to 30 September 2014

Distribution Statement A: Approved for public release distribution is unlimited.

**Air Force Research Laboratory
Air Force Office of Scientific Research
European Office of Aerospace Research and Development
Unit 4515, APO AE 0942-4515**

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 30 September 2014	2. REPORT TYPE Final	3. DATES COVERED (From – To) 1 October 2011 – 30 September 2014
---	--------------------------------	---

4. TITLE AND SUBTITLE Air-to-air and air-to-ground attack strategies in trained birds of prey	5a. CONTRACT NUMBER
	5b. GRANT NUMBER FA8655-11-1-3065
	5c. PROGRAM ELEMENT NUMBER 61101F

6. AUTHOR(S) Prof. Graham K. Taylor, Prof. Adrian L. R. Thomas, and Ms. Caroline H. Brighton	5d. PROJECT NUMBER
	5d. TASK NUMBER
	5e. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Oxford Department of Zoology South Parks Road Oxford, OX1 3PS United Kingdom	8. PERFORMING ORGANIZATION REPORT NUMBER N/A
--	--

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) EOARD Unit 4515 APO AE 09421-4515	10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/AFOSR/IOE (EOARD)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-AFOSR-UK-TR-2015-0006

12. DISTRIBUTION/AVAILABILITY STATEMENT
Distribution A: Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The aim of this research was to analyse the strategies used by two species of birds of prey (peregrine falcons *Falco peregrinus* and Harris' hawks *Parabuteo unicinctus*) during air-to-air and air-to-ground attacks, from the perspective of guidance and control. This was done experimentally for three test cases: 1. Peregrine attacks on targets thrown to the ground by a falconer; 2. Harris' hawk attacks on targets towed over the ground at high speed along a zigzagging course of pulleys; 3. Peregrine attacks on targets towed through the air by a maneuvering remotely piloted air vehicle. We used onboard GPS loggers or high-speed video cameras to provide accurate measurements of the instantaneous position of the bird and its target (typically <0.05m error), and used ground-based video cameras and onboard video cameras to provide qualitative recordings of head, wing, or tail movements.

15. SUBJECT TERMS
EOARD, avian ground attack, avian attack strategies

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 50	19a. NAME OF RESPONSIBLE PERSON Gregg L. Abate
a. REPORT UNCLAS	b. ABSTRACT UNCLAS	c. THIS PAGE UNCLAS			19b. TELEPHONE NUMBER (Include area code) +44 (0)1895 616021

Final Report on Award no. FA8655-11-1-3065

Air-to-air and air-to-ground attack strategies in trained birds of prey

Period of Performance

1 October 2011 to 30 September 2014

Principal Investigator: Prof. Graham K. Taylor

Co-Principal Investigator: Prof. Adrian L. R. Thomas

Doctoral Research Student: Ms Caroline H. Brighton

*Department of Zoology, University of Oxford, South Parks Road,
Oxford, OX1 3PS, United Kingdom*

email: graham.taylor@zoo.ox.ac.uk

AFOSR/EOARD Programme Manager: Dr Gregg Abate

Date report submitted: 30 September 2014

Table of Contents

List of Figures	1
List of Tables	2
Summary	3
Introduction	4
Geometric rules of pursuit in nature	4
Guidance laws in biology and engineering	8
Rationale	9
Methods, Assumptions, and Procedures	10
Air-to-ground attacks in Peregrine Falcons	10
Animals	10
Equipment	10
Experimental protocol	10
Error analysis	12
Data selection and synchronization	12
Air-to-ground attacks in Harris' Hawks	13
Animals	13
Equipment	13
Experimental protocol	15
Experimental design	15
Marker tracking software	17
Photogrammetry	17
Error analysis	20
Data selection	20
Air-to-air attacks in Peregrine Falcons	22
Animals	22
Equipment	22
Experimental protocol	22
Data synchronization	22
Results and Discussion	23
Air-to-ground attacks in Peregrine Falcons	23
Data processing	23
Analysis of results	24
Air-to-ground attacks in Harris' Hawks	27
Data processing	27
Analysis of results	31
Air-to-air attacks in Peregrine Falcons	36
Preliminary analysis of results	36
Conclusions	41

Acknowledgments	42
References	43
List of Symbols, Abbreviations, and Acronyms	45

List of Figures

- Figure 1 Theoretical pursuit trajectories under three simple geometric rules.
- Figure 2 Anatomy and field of view of the visual system in hawks and falcons.
- Figure 3 GPS trajectory and onboard video of an incidental pursuit of a live prey item by a peregrine falcon.
- Figure 4 Schematic of the camera setup and lure-pulley setup used to elicit air-to-ground attacks in Harris' hawks.
- Figure 5 Schematic of the 16 possible courses the lure could take in air-to-ground attack experiments in Harris' hawks.
- Figure 6 Screenshot of the color-tracking software used to track markers automatically for photogrammetry.
- Figure 7 Reconstruction of a calibrated camera setup and of the position and pose of the calibration object used.
- Figure 8 Histogram quantifying the error in measuring the calibration object length using photogrammetry.
- Figure 9 Horizontal projection of peregrine trajectories during attack passes against stationary ground targets.
- Figure 10 Estimated turn rate versus estimated line-of-sight rate during peregrine attacks on stationary ground targets.
- Figure 11 Comparison of actual trajectories of peregrine attacks with those simulated assuming proportional navigation.
- Figure 12 Comparison of actual trajectories of peregrine attacks with those simulated assuming proportional navigation.
- Figure 13 Comparison of actual trajectories of peregrine attacks with those simulated assuming proportional navigation.
- Figure 14 Example trajectory for a Harris' hawk chasing a maneuvering lure in an air-to-ground attack.
- Figure 15 Method of calculating track angle and line-of-sight angle from photogrammetric position data.
- Figure 16 Track angle and line-of-sight angle plotted against time during Harris' hawk attacks on a maneuvering lure.
- Figure 17 Correlation analysis of time lag associated with Harris' hawk pursuit of a maneuvering lure.
- Figure 18 Lagged track angle and line-of-sight angle plotted against time during Harris' hawk attacks on a maneuvering lure.
- Figure 19 Turn rate and line-of-sight rate plotted against time during Harris' hawk attacks on a maneuvering lure.
- Figure 20 Lagged turn rate and line-of-sight rate plotted against time during Harris' hawk attacks on a maneuvering lure.
- Figure 21 Raw GPS position data from peregrine, lure, and remotely piloted air vehicle during air-to-air attacks.

List of Tables

None.

Summary

This final report describes research undertaken on AFOSR Grant FA8655-13-1-3077 during the period of performance 1 October 2011 to 30 September 2014. The aim of this research was to analyse the strategies used by two species of birds of prey (peregrine falcons *Falco peregrinus* and Harris' hawks *Parabuteo unicinctus*) during air-to-air and air-to-ground attacks, from the perspective of guidance and control. This was done experimentally for three test cases: 1. Peregrine attacks on targets thrown to the ground by a falconer; 2. Harris' hawk attacks on targets towed over the ground at high speed along a zigzagging course of pulleys; 3. Peregrine attacks on targets towed through the air by a maneuvering remotely piloted air vehicle. We used onboard GPS loggers or high-speed video cameras to provide accurate measurements of the instantaneous position of the bird and its target (typically <0.05m error), and used ground-based video cameras and onboard video cameras to provide qualitative recordings of head, wing, or tail movements. Analysis of the trajectory data showed that both the peregrine falcons and the Harris' hawks turned at a rate proportional to the rate of change of the direction of their line of sight to target. This is consistent with their use of a proportional navigation guidance law, as commonly used by modern guided missiles, indicating that birds of prey have converged upon broadly the same solution as missile engineers. In the case of the Harris' hawks, the turn rate was typically approximately equal to the line-of-sight rate, albeit with a short (c. 0.19s) time delay, with the result that the birds followed a pure or deviated pursuit course (i.e. they flew directly at, or less commonly at a constant angle ahead of, their target). In the case of the peregrines, the constant of proportionality (i.e. the navigation constant of pure proportional navigation) was typically greater than one, varying between – but not within – separate attack passes at the target. This variation is likely to be adaptive, but our results to date do not yet allow us to conclude how. We discuss how the guidance of the recorded attacks is likely to have been implemented in relation to the known visual physiology and anatomy of birds of prey, and in relation to our video recordings of head movements during the attacks. We conclude that the way in which guidance is implemented in a bird is fundamentally different to the way in which guidance is implemented in most missiles. Specifically, birds have a much wider field of view than most guided missiles, and are most closely comparable in this sense to missiles with strapdown seekers. On the other hand, in contrast to a missile with a strapdown seeker, birds use head movements to keep the target within the field of view of their visual sensors, and more especially within the most acute parts of that field of view. Furthermore, they stabilize their gaze during the periods between these saccadic head movements in a manner that is reminiscent of an image-stabilized camera system rather than a gimballed missile seeker. Further ongoing analysis of the data collected during this research effort now seeks to elucidate in detail the role of head movements in implementing effective guidance strategies in birds.

Introduction

Aerial attack and pursuit behaviors are essential to predation and courtship in many flying animals. Previous research on target-directed flight behaviors in insects (Land & Collett, 1974; Collett & Land, 1975, 1978; Land, 1993; Olberg *et al.*, 2000, 2007; Mizutani *et al.*, 2003; Olberg, 2012), birds (Tucker, 2000a,b; Tucker *et al.*, 2000; Kane & Zamani, 2014), and bats (Ghose *et al.*, 2006) has aimed at identifying simple geometric rules characterizing their observed flight trajectories. Such geometric rules can only be a part of the picture, however, because they merely characterize the integrated outcome of the commanded changes in flight velocity. Those commands must themselves be generated by an underlying guidance law, according to which sensory feedback is used to determine the required acceleration. Completing the picture therefore requires identification of the guidance law underpinning a given behavior, but this has not previously been achieved for any aerial predator. The aim of the research described in this report was to identify the guidance laws used by birds of prey to attack targets in the air and on the ground, and to place this in the context of their use of high-acuity vision and highly agile flight.

Geometric rules of pursuit in nature

Previous studies have identified three simple geometric rules used by flying animals, each defined by the constancy of one of two angles characterizing the local geometry of the pursuit. These are the line-of-sight angle (λ), defined as the geographic direction of the line of sight drawn from the pursuer to its target, and the deviation angle (δ), defined as the angle between the pursuer's velocity vector and its line of sight to target. Under the simplest geometric rule, observed in chasing flies (Land & Collett, 1974; Land, 1993), the pursuer heads directly towards its target at all times t during the pursuit, thereby holding its deviation angle δ at zero. This results in a pure pursuit course (Fig. 1A), defined by the geometric rule $\delta(t) = \delta(0) = 0$, where $t = 0$ defines the start of the pursuit. Pure pursuit behaviors have also been observed in bees landing on a target (Zhang *et al.*, 1990), beetles running after prey (Gilbert, 1997), and fish swimming after sinking food (Lanchester & Mark, 1975). In the second geometric rule, observed in hoverflies (Collett & Land, 1975, 1978), the pursuer aims its velocity vector at a constant non-zero deviation angle δ ahead of its line of sight to target. This results in a deviated pursuit course (Fig. 1B), defined by the geometric rule $\delta(t) = \delta(0) \neq 0$. Deviated pursuit behaviors have also been observed in dogs (Shaffer *et al.*, 2004) and humans (Regan & Gray, 2000; Shaffer & McBeath, 2002) running to catch thrown targets. In a third geometric rule, observed in bats (Ghose *et al.*, 2006), dragonflies (Olberg *et al.*, 2000, 2007; Mizutani *et al.*, 2003; Olberg, 2012), and falcons (Kane & Zamani, 2014), the line-of-sight angle λ is held constant. This results in a parallel navigation course (Fig. 1C), defined by the geometric rule $\lambda(t) = \lambda(0)$, and leads incidentally to a form of motion camouflage, because from the target's perspective, the pursuer appears stationary relative to suitably distant landmarks (Srinivasan & Davey, 1995; Justh & Krishnaprasad, 2006).

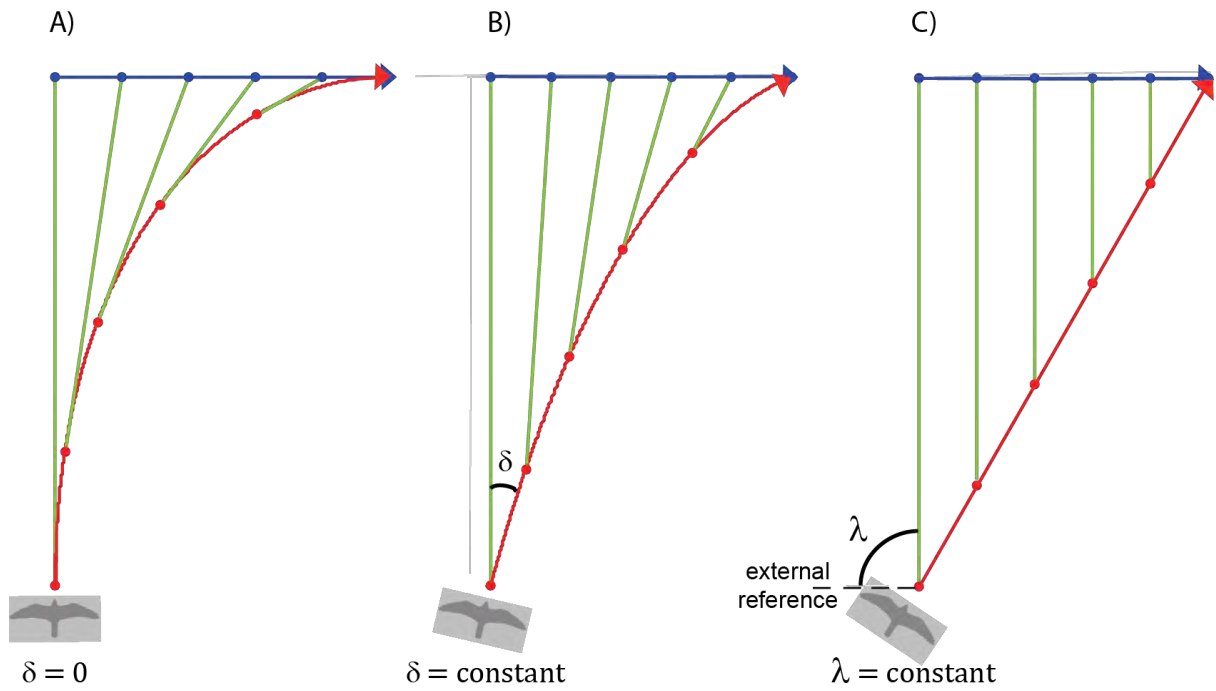


Figure 1 (A-C): Theoretical pursuit trajectories for a pursuer (red) capturing a non-maneuvring target (blue) showing the line of sight (green) between corresponding pursuer and target positions for the three geometric rules discussed in the text: A) pure pursuit course with zero deviation angle δ between the pursuer's velocity vector and the line of sight to target, i.e. $\delta = \delta(0) = 0$ such that the line of sight is always tangent to the pursuer's trajectory; B) deviated pursuit course with a constant non-zero deviation angle δ , i.e. $\delta = \delta(0) \neq 0$; C) parallel navigation course with a constant line-of-sight angle λ measuring the geographic direction of the line of sight relative to an external reference, i.e. $\lambda = \lambda(0)$.

Missile engineers first designed guidance laws to implement these same geometric rules over sixty years ago (Shneydor, 1998; Siouris, 2004), but the question of how they are implemented in nature has gone unanswered. Previous research on falcons has attempted to relate their attack behaviors to their visual anatomy (Tucker, 2000a,b; Tucker *et al.*, 2000; Kane & Zamani, 2014). Diurnal birds of prey vary considerably in the extent of their eye movements (O'Rourke *et al.*, 2010a), but the eyes appear to have very little mobility in falcons ($\pm 2^\circ$ azimuth in the American kestrel *Falco sparverius*, O'Rourke *et al.*, 2010a; but see Frost *et al.*, 1990) and quite limited mobility in hawks ($\pm 10^\circ$ azimuth in the red-tailed hawk *Buteo jamaicensis*; $\pm 12^\circ$ azimuth in the Cooper's hawk *Accipiter cooperi*; O'Rourke *et al.*, 2010a), so that head orientation effectively determines gaze direction (O'Rourke *et al.*, 2010b). Falcons and hawks are bifoveate (Fig. 2), each retina having one forward-facing acute zone (the temporal, or shallow, fovea) and one laterally directed acute zone (the central, or deep) fovea. Bifoveate birds use their forward-facing temporal fovea preferentially when inspecting nearby targets, but typically look sideways at more distant targets (Frost *et al.*, 1990; Tucker, 2000b; O'Rourke *et al.*, 2010b), which led Tucker (2000b) to propose that the laterally directed central fovea would be used to target distant prey items. Tucker (2000a) further proposed that the head would have to be held straight for streamlining during the high-speed stoops for which peregrine falcons are renowned. The conjunction of these two propositions led Tucker (2000b) to hypothesize that stooping falcons would fly a deviated pursuit course, with their velocity vector aimed at a constant deviation angle $\delta(t) = \phi$ ahead of their line of sight to target, where ϕ is the angle between the central fovea and the sagittal plane (Fig. 2A). Thus, under Tucker's hypothesis, the geometry of a falcon's attack trajectory is supposed to emerge as an adaptation to the constraints imposed by its anatomy and aerodynamics (Tucker, 2000b; Tucker *et al.*, 2000).

Unfortunately, no detailed measurements of eye geometry exist for our study species the peregrine falcon *Falco peregrinus* and the Harris' hawk *Parabuteo unicinctus*, but measurements from two closely related species (see Fig. 2) indicate that $\phi \approx 48^\circ$ in the American kestrel *Falco sparverius* (Frost *et al.*, 1990) and $\phi \approx 40^\circ$ in the broad-winged hawk *Buteo platypterus* (Slonaker, 1897). Assuming that peregrines and Harris' hawks have a similar foveal angle, then flying with the head straight and the target fixated on either central fovea would naturally produce a deviated pursuit course with deviation angle $\delta = \phi \approx 48^\circ$. Tucker *et al.* (2000) claimed that the curved trajectories that they observed in wild peregrines were qualitatively consistent with this hypothesis, but more recent work using head-mounted video cameras (Kane & Zamani, 2014) has cast doubt upon the supposition that the laterally directed central fovea is used ubiquitously – or often – when targeting aerial prey. In fact, although falcons appear to fixate targets during aerial pursuits, the target's inferred position on the retina is typically nearer to the forward-facing temporal fovea (Kane & Zamani, 2014). It therefore appears that the curved attack trajectories of falcons cannot be explained away as an emergent property of the geometry of their visual system. On the contrary, it may only become possible to understand the visual anatomy and physiology of hawks and falcons after we have first understood the guidance laws that feedback from their visual system enables.

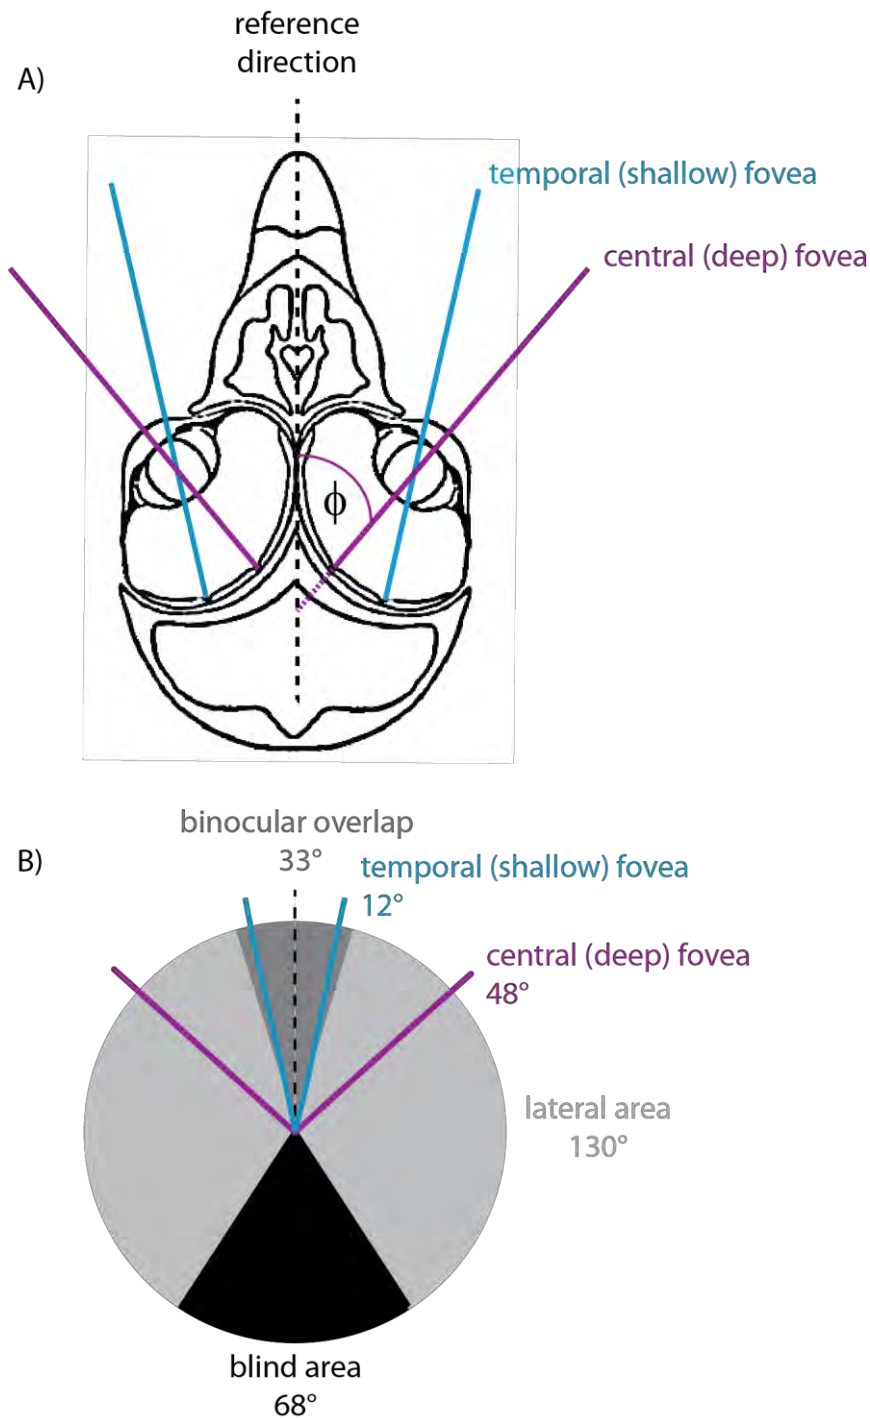


Figure 2: A) Diagram of a broad-winged hawk (*Buteo platypterus*) skull showing the eye anatomy and locations of the two foveae on the retina and their respective gaze directions. The foveal angle ϕ denotes the angle between the gaze direction of the lateral fovea and the sagittal plane. Note that the eyes of this specimen may be more diverged in death than they would have been in life. Modified from Slonaker (1897). B) Gaze directions of the foveae in a small falcon (American Kestrel, *Falco sparverius*). The dashed line denotes the sagittal plane. Data from O'Rourke *et al.* (2010a) and Frost *et al.* (1990).

Guidance laws in biology and engineering

Most guided missiles use proportional navigation, or some variant thereof, as their guidance law. In what is conceptually the simplest case – of pure proportional navigation – turning is commanded at a rate proportional to the rate of change in the line-of-sight angle, according to the guidance law:

$$\frac{d\gamma}{dt} = N \frac{d\lambda}{dt} \quad (1)$$

where $\gamma(t)$ is the geographic direction of the missile's velocity vector, and where the parameter N is called the navigation constant (Shneydor, 1998). Integrating Eq. 1 between limits of $t = 0$ and $t = \tau$, and solving for $\lambda(\tau)$ yields:

$$\lambda(\tau) = \lambda(0) + \frac{\gamma(\tau) - \gamma(0)}{N} \quad (2)$$

It is clear by inspection of Eq. 2 that $\lambda(t)$ approaches $\lambda(0)$ in the limit as $N \rightarrow \infty$, and hence that proportional navigation effectively implements the parallel navigation rule $\lambda(t) = \lambda(0)$ when N is very large. At the other extreme, setting $N = 1$ in Eq. 2 and making use of the identity $\delta(t) \stackrel{\text{def}}{=} \gamma(t) - \lambda(t)$, we find that $\delta(t) = \gamma(0) - \lambda(0) = \delta(0)$. Hence, when N is equal to one, proportional navigation implements either the pure pursuit rule $\delta(t) = \delta(0) = 0$ or the deviated pursuit rule $\delta(t) = \delta(0) \neq 0$, according to the initial conditions defined by $\delta(0)$. It follows that the proportional navigation guidance law in Eq. 1 is capable of implementing any of the three geometric rules that have been identified in nature (see above). Naturally, proportional navigation can also produce a continuum of intermediate trajectories, depending upon the numerical value of the navigation constant. This remarkable property – coupled with its exquisite simplicity as a guidance law – makes proportional navigation an obvious candidate for explaining how natural systems intercept moving targets.

Proportional navigation has been mooted previously as a candidate guidance law for aerial predators (Anderson, 1982; Shneydor, 1998; Olberg, 2012), but the hypothesis has not been tested explicitly. Proportional navigation is especially useful against aerial targets, but is equally effective against ground targets. This is because the proportional navigation guidance law (Eq. 1) is fundamentally a statement about how an attacker should respond to changes in line-of-sight angle (λ) – regardless of whether these changes arise through the motion of the target or the self-motion of the attacker. Proportional navigation therefore serves to correct accumulated steering errors on the part of the attacker in just the same way as it corrects for target motion, which explains why some guided missiles designed for use against stationary ground targets employ proportional navigation (Shneydor, 1998). Thus, it is reasonable to suppose that an aerial predator adapted to use proportional navigation against aerial targets might also use proportional navigation against a ground target that could take flight at any time.

Rationale

Although some more advanced guided weapons systems use optimal control, we think it unlikely that flying animals perform the computations necessary to implement this. Our approach in this study has therefore been to test whether hawks and falcons use a) some variant of proportional navigation or b) some simpler form of pursuit, whilst remaining open to the possibility that neither hypothesis may be correct. Within this broader framework, we have also sought to test the narrower hypothesis that birds of prey adapt their guidance to different situations, and have sought to identify how this is implemented in relation to the physiology of their visual system (see above).

In order to accomplish our aim of identifying the guidance laws used by birds of prey when attacking targets in the air and on the ground, we have chosen to analyze three experimental situations of increasing complexity:

- I. In the first test case, captive peregrine falcons *Falco peregrinus* were flown at targets thrown to the ground by their handler, thereby simulating a low-level swoop against a perched or otherwise stationary ground target. This experimentally accessible behaviour simulates an ecologically important foraging strategy in this species, since a large proportion – or even a majority – of peregrine attacks in open habitats involve low-level swoops at perched, sitting, or swimming prey (flooded fields: Dekker, 1987, 1988, 1995; estuaries: Cresswell, 1996; sea's surface: Dekker & Bogaert, 1997). More importantly for our present purposes, this simplest test case is an attractive one analytically, because the equations of motion governing guidance under proportional navigation have a closed-form solution in the case of a stationary target.
- II. In the second test case, captive Harris' hawks *Parabuteo unicinctus* were flown against a target towed along the ground at speed in an unpredictable (because randomly varied) zigzag trajectory. This behaviour simulates the pursuit of a fleeting target on the ground, which is a foraging strategy typical of this species.
- III. In the third test case, captive peregrine falcons *Falco peregrinus* were flown against a target that was towed at speed through the air by a remotely piloted air vehicle. In this case, the pilot attempted to evade capture of the lure during the early stages of the pursuit, before eventually allowing a successful capture. This behaviour simulates the aerial pursuit of a manoeuvring target.

Analysis of the first two test cases is complete, whilst analysis of the third test case remains ongoing. The present report summarizes our methods and findings in respect of each.

Methods, Assumptions, and Procedures

The following experimental protocols were approved by the United States Air Force, Surgeon General's Human and Animal Research Panel, and by the Local Ethical Review Committee of the University of Oxford's Department of Zoology, and were considered not to pose any significant risk of causing pain, suffering, damage or lasting harm to the animals involved. We are pleased to confirm that no pain, suffering, damage or lasting harm was caused to any of the animals used during the course of this research. One of the birds died of natural causes before the end of the project, for reasons unrelated to experimental work and husbandry.

Air-to-ground attacks in Peregrine Falcons

Animals

Three trained peregrine falcons *Falco peregrinus* ("Bella": ♀, 0.96kg; "Nina": ♀, 0.90kg; "Weirdo": ♂, 0.65kg) were flown at targets – a dead pheasant *Phasianus colchicus*, or another dead prey item – thrown to the ground by their handler. The birds usually made more than one pass at the target, resulting in a total of $n = 57$ passes over the course of 26 flights. The peregrines attacked two live prey encountered incidentally during testing: neither prey item was harmed, but we were able to include 4 ground-attack passes made at a mallard *Anas platyrhynchos*, in the analysis (Fig. 3).

Equipment

Each of the three peregrines carried a GPS receiver logging position and groundspeed at 5Hz (Qstarz BT-Q1300) and a high-definition video camera looking forward over the head (HD720P Mini DV, 30fps). The equipment weighed 0.031kg in total, and was carried dorsally on a commercially available harness (Marshall Radio Telemetry, Trackpack) comprising a curved plastic mounting plate held in place between the shoulders by a pair of flat tubular Teflon ribbons drawn once around the bird's body in a figure-of-eight pattern. Each bird also carried a tail-mounted radio transmitter for tracking purposes (Marshall Radio Telemetry Micro Transmitter).

Experimental protocol

The GPS unit was allowed to settle for at least 5 minutes before beginning the experiment, after which time the peregrine was released by its falconer and allowed to gain height. The falconer then shouted a command that the birds associated with food, and threw the target a short distance up in the air. Although the birds occasionally caught their targets mid-air, the targets were never thrown more than a few metres from the falconer, and were usually only taken after they had hit the ground. We therefore treat the targets as being

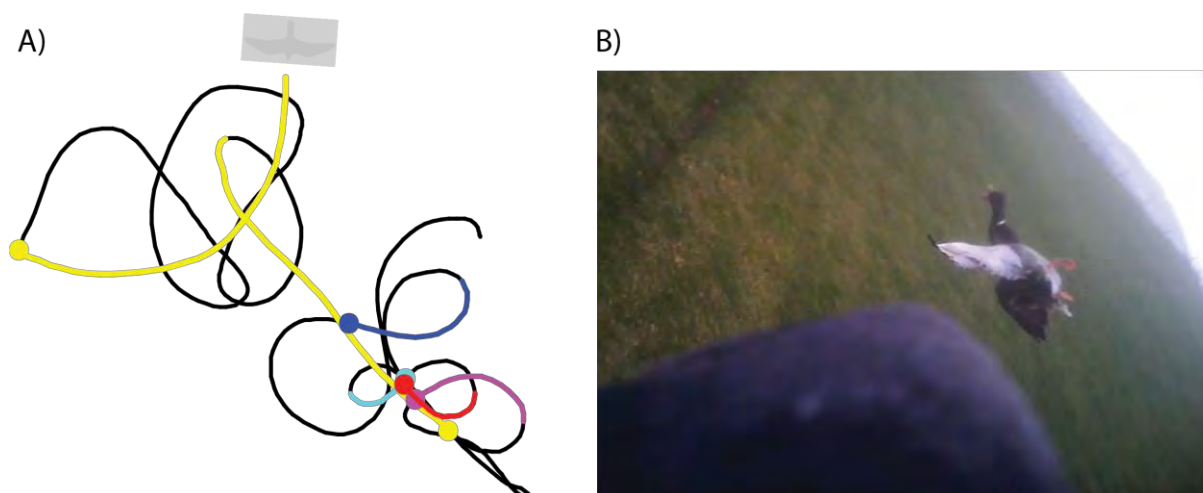


Figure 3: A) GPS trajectory and B) onboard camera image of a flight in which the peregrine chased a mallard duck in flight (yellow) before making several further passes at it on the ground. A) The different colours represent the separate attack passes with target position at the corresponding coloured circle (yellow trajectories denote pursuits which were not included in the analysis because the duck was flying, so that its position, until the point of interaction, could not be determined). B) The image shows the duck with its body inverted as it attempts (successfully) to evade the peregrine, and the tilt of the horizon indicates the bank angle of the peregrine during the terminal phase of this attempted interception.

stationary for the purposes of the analysis: an assumption which is satisfied exactly in most of the attack passes that we have analysed, and approximately in the remainder of the passes. This being so, it is possible – hypothetically at least – that the peregrines might have computed or otherwise determined an intercept trajectory that they then followed in open loop. We think this hypothesis improbable on two grounds: first, because trajectory planning in open loop is much more computationally demanding than using a guidance law in closed loop; and second, because the birds were subject to buffeting by winds, and will in any case have been unable to implement forward commands without error, so that some form of closed-loop guidance would always be required.

Error analysis

According to the current GPS Standard Positioning Service Performance Standard (Department of Defense Positioning, Navigation, and Timing Executive Committee, 2008), civilian GPS units typically provide an absolute positioning accuracy of $\leq 7.8\text{m}$ (95% of data, as a global average). This figure is not, however, representative of the much greater accuracy with which such units can measure changes in position over short periods of time. To assess the relative positioning accuracy of our GPS units directly, we characterized the noise in two of the units whilst stationary over sampling intervals of comparable duration to the longest of the attack passes that we measured (approximately 10s). After allowing each GPS unit to settle for 5 minutes, we sampled the next 200 consecutive 10s blocks of position data from each unit. Then, for each 10s block of data, we calculated the three-dimensional distance of each of the position fixes from the mean of all of the 50 position fixes within that block. Since there was no appreciable difference in the noise level in the two units, we pooled the data. The median distance deviation within these 10s blocks was 0.014m, and 95% of the 20,000 position fixes fell within 0.050m of the mean position fix for their respective block. The noise in our GPS position fixes is therefore typically $\leq 0.05\text{m}$ on the timescale of the longest recorded attack passes – at least under stationary test conditions.

Data selection and synchronization

We synchronized the GPS data and onboard video data by identifying the start and end time of each flight in both of the data streams. We then used the onboard video data to identify the precise time(s) at which the peregrine contacted or otherwise passed immediately over its target, and used this time to identify the GPS position of the target on each pass.

Air-to-ground attacks in Harris' Hawks

Animals

We used five captive bred Harris Hawks *Parabuteo unicinctus* ("Ruby": ♀, 0.93kg; "Findo": ♂, 0.70kg; "Aggy": ♀, 0.92kg; "Jake": ♂, 0.62kg; "Spitfire": ♂, 0.66kg), all trained to capture an artificial lure pulled along the ground to simulate a maneuvering terrestrial prey item. After more than 4 months in the field, we achieved a total of 20 good flights for each bird, defined as a flight in which the bird actively chased and intercepted the lure within view of all four of the high-speed cameras that we used to film the chase. We captured data from many more flights in which the bird failed to intercept the lure, recording well in excess of 100 chases in total.

Equipment

Each bird was fitted with a Trackpack harness (Marshall Radio Telemetry) before each flight (see above for detailed description). The birds were then fitted with two 4cm diameter polystyrene balls coloured with fluorescent paint for ease of visualization: a pink ball attached to the crossover point of the harness on the breast, and a green ball attached to the mounting plate of the harness on the back, each using Teflon ribbon. A measurement of wind speed and direction was taken at the end of each flight day, using a Kestrel 4500 Pocket Weather Tracker anemometer and Kestrel Wind Vane (Nielsen-Kellerman, Boothwyn, PA, USA).

An artificial lure was attached via 100m of kite line to a 12V DC motorised winch ('Expert Winch', Gliders Distribution, Newark, Notts, UK.) powered by a car battery. The lure was made from a fluorescent yellow mitt, to facilitate visualization and to provide a good gripping surface for the birds. On days when the grass was very wet, this was substituted for a waterproof version to prevent the lure being saturated and thus too heavy for the winch to pull at speed. A food reward was attached to the lure to motivate the birds. The lure was placed beneath a cover in one of three starting positions, and was run quickly around a series of pulleys to generate abrupt changes in direction simulating those of an evasively moving prey item (Fig. 4). We used ten pulleys in total, six of which were built into tunnels for the purpose of guiding the lure around the pulley, and to provide the birds with additional motivation to chase the target. (In pilot trials without the tunnels, the birds appeared less motivated to chase the target, but adding the tunnels seemed to provoke them to attempt to capture the lure before it disappeared.)

The birds were filmed using four synchronised S-PRI high-speed cameras (Lake Image Systems Ltd, Tring, Herts, UK) giving 8 seconds of record time (resolution: 1280 x 1024 pixels; frame rate: 250 Hz; exposure time: 0.002s). The cameras were fitted with Nikon AF Nikkor 28mm f/2.8D wide-angle lenses (Nikon UK Ltd, Kingston upon Thames, Surrey, UK) to give a field of view adequate to cover the volume in which the lure moved. The cameras were arranged in a rectangle measuring approximately 12m by 25m, with the long axis of the rectangle perpendicular to the slope (gradient approximately 1:10), and were pointed uphill with overlapping fields of view (Fig. 4).

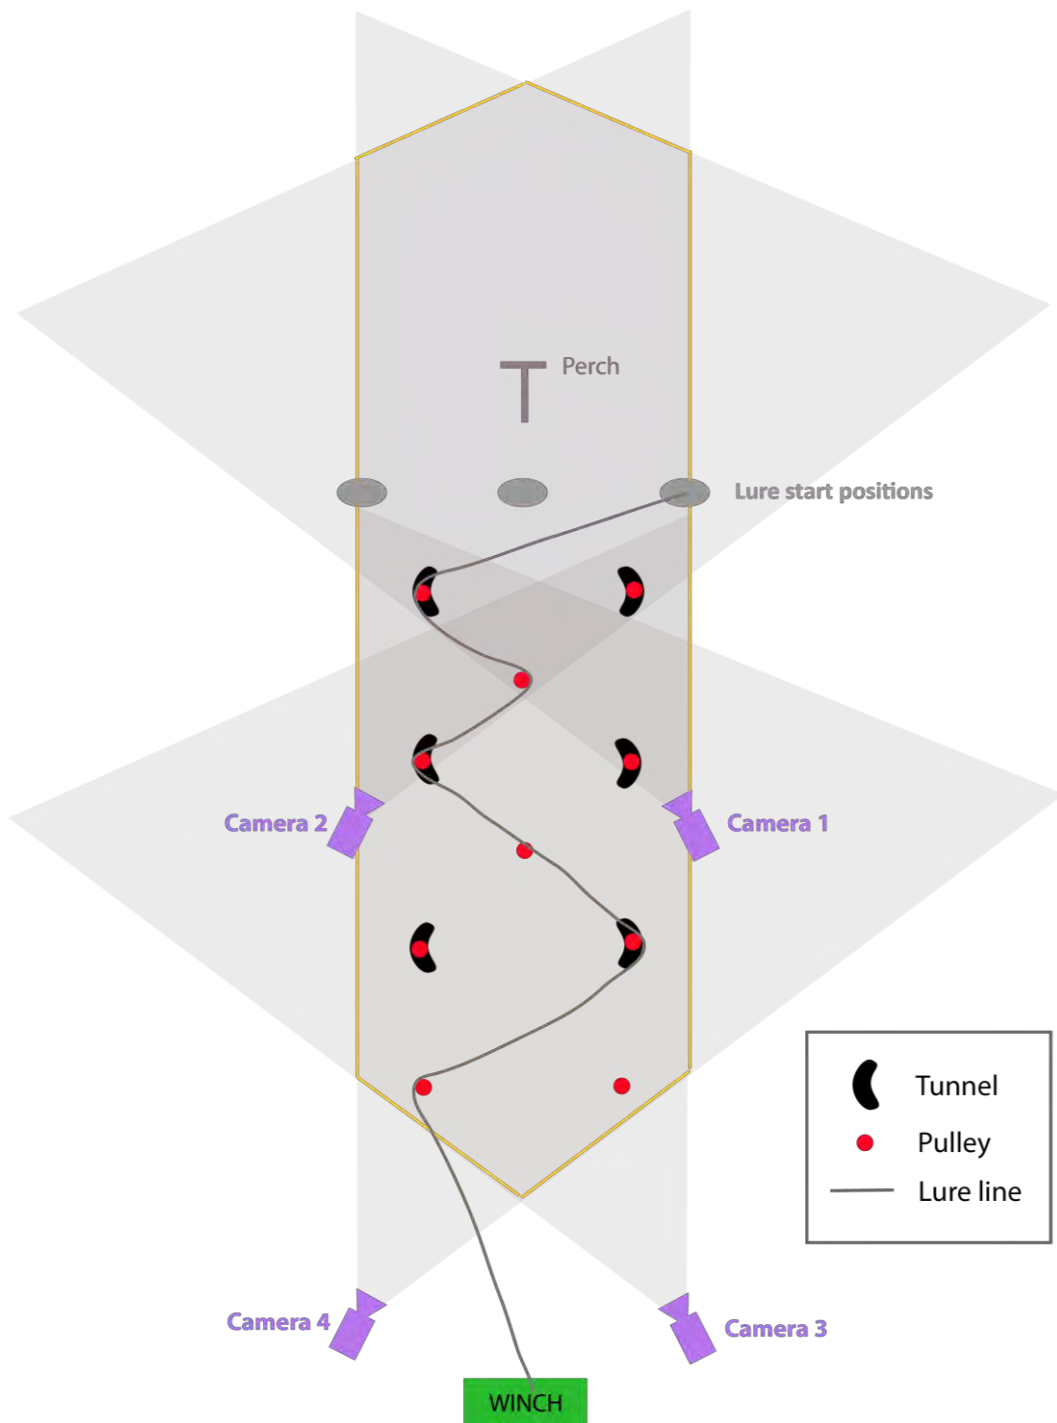


Figure 4: Diagrammatic representation of the pulley set-up showing camera positions and pose (purple), camera field of views (opaque triangular areas), winch (green), pulleys (red), tunnels (black), perch (brown), lure starting positions (grey circles), and one example lure-course (grey line).

Experimental protocol

The tests with two of the birds (Ruby and Findo) were conducted on an open grassy field at Great Llywgy Farm (Lower Cymyoy, Abergavenny, UK) between October and December 2012. The tests with the other three birds (Aggy, Jake and Spit) were conducted on an open grassy field near Ffawyddden Cottage (Llandewi Rhydderch, Abergavenny, UK) between July and September 2013. The birds were flown individually, either from the falconer's fist, or from a perch positioned at head-height at the top of the test area, or from a nearby tree. The winch was activated at the start of each flight, which prompted the bird to take off and chase the target down the hill. The speed of the winch was controlled manually according to the bird's proximity to the lure, so as to maintain the bird's motivation. (Pilot trials indicated that the birds were not motivated to continue chasing the lure if it was allowed to run too far ahead of them.) Interception occurred within the measurement volume, and the winch was stopped immediately upon capture to avoid harming the bird. The cameras were set to post-trigger, and were triggered manually after the bird had caught the lure. The video data from each camera pair was immediately saved onto a laptop in the field (Toshiba Tecra R850-140, Dell Latitude E5510). Together, the cameras captured the whole of the attack sequence, typically involving an initial wing beat acceleration, followed by fast twists, turns and wingover maneuvers, terminating with a rapid leg extension at the point of capture.

We found that it was possible to undertake 2 to 4 flights on a single day with each bird. If the bird proved not to be highly motivated on a given flight (e.g. because it spent a long time gliding after the lure, rather than in active flight), then it would be made to experience a 'miss' by stopping the lure out of reach to prevent interception – either by stopping the lure in a tunnel, or by winding the lure in fully and covering it. This approach was found to improve the bird's motivation on subsequent trials, and was intended to create more realistic hunting conditions. In any case, the bird always received a full crop of food after its last flight to keep it motivated for subsequent test days.

Experimental design

With the pulley setup that we used (Fig. 4), it was possible to pull the lure along 16 different courses beginning at one of 3 different starting points (Fig. 5). We randomly selected which course to use on each flight, and laid decoy lure lines to control for the possibility that the birds might simply follow the line, as well as decoy covers over the starting positions to prevent the bird from anticipating from where the lure would appear and in what initial direction of travel. The lure switched its direction of travel 4, 5, or 6 times on each complete run, depending upon the course that was set (Fig. 5). The lure always switched its direction 4 times at the periphery of the course, which always resulted in the lure being pulled towards the middle of the course, so it was possible in principle for the bird to anticipate the direction of these switches if it had learned this aspect of the task. On the other hand, it should not have been possible for the bird to anticipate whether the lure would change direction as it passed around

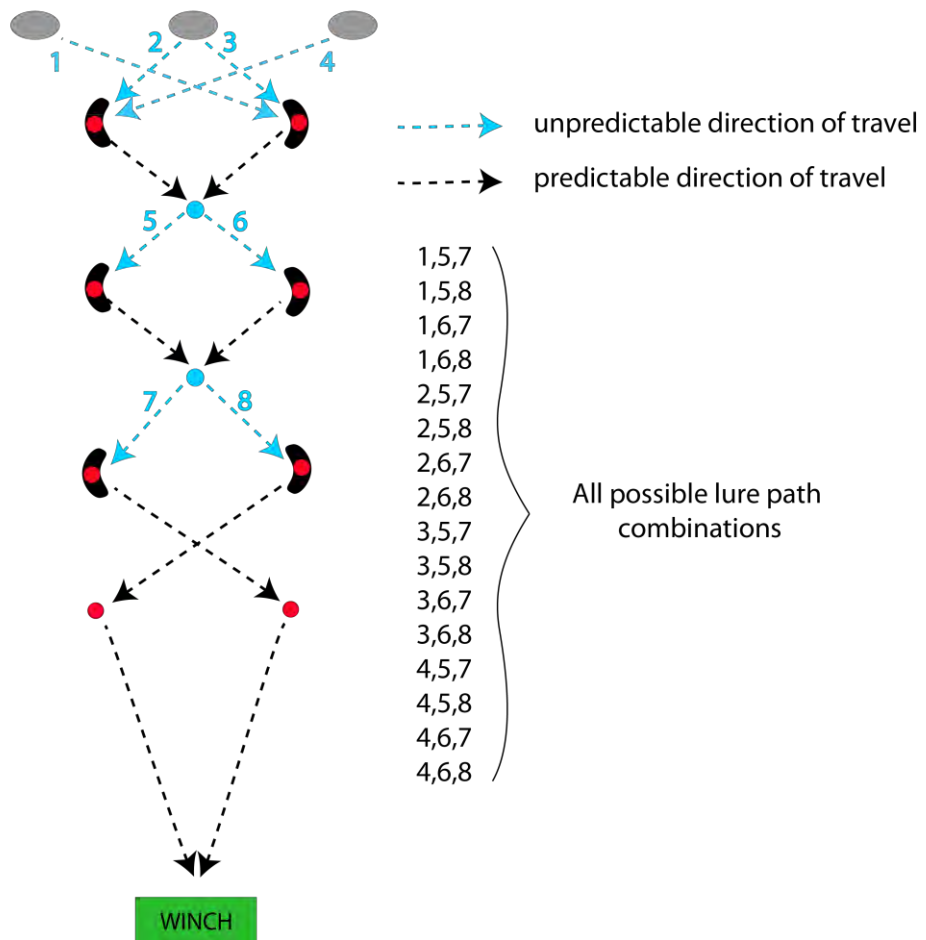


Figure 5: Diagram showing the 16 possible lure courses in the pulley set-up that we used, highlighting those pulleys which create an unpredictable change in lure direction and those which do not.

the two pulleys in the middle of the course (excepting the possibility that the bird could discriminate which of the lines was under tension on a given run).

Marker tracking software

The pixel coordinates of the lure and of the two marker balls on the bird were identified for each camera view using in-house software in Matlab v.7.4 (courtesy of Shane Windsor; Fig. 6). In brief, the user first creates a color template by outlining the pixels corresponding to the object to be tracked in a series of specified video frames. A local search area around the marker is also defined, and the software then uses the color template to identify matching pixels in each frame of the video sequence, updating the search area by centering the search area on the last successfully tracked point. The color of each pixel within the search area is then compared to the distribution of the color template by computing its Mahalanobis distance. Any pixels having a Mahalanobis distance less than a set threshold are assumed to be candidate marker pixels, and the centroid of the single largest contiguous group of candidate marker pixels is assumed to correspond to the centroid of the marker itself. The software uses a Kalman filter to estimate the position of the marker if no marker was found, or in the event that the estimated position of the marker jumps by more than a given amount between frames. The markers were obscured when the lure entered a tunnel or if the bird flew outside of the field of view, and in these cases the data were treated as missing. All estimated marker positions were checked manually and corrected if necessary using in-house software in Matlab v.7.4 (courtesy of Simon Walker).

Photogrammetry

The cameras were set up and taken down at the beginning and end of each test day, so had to be calibrated separately each day. This was done by filming a dumbbell-shaped calibration object, comprising a 1m long clear acrylic tube with a ball secured at each end, being moved through the test volume in a range of different orientations. The position of the two marker balls was tracked automatically offline using the color-tracking software described above. The camera collinearity equations were then solved by the method of nonlinear least squares in a bundle adjustment that identified jointly optimal estimates of the position and pose of the calibration object and the cameras (Fig. 7).

Lens distortions were found to be minimal and we therefore assumed a central perspective projection, assuming no lens distortion and no offset of the principal point with respect to the center of the camera sensor. Field estimates of camera position and pose were used to provide initial values for these parameters in the optimization, in order to guarantee convergence. Having first optimized the focal length separately for each flight, we ultimately fixed the principal distance of the cameras at 28.42mm for all flights, calculated as the mean for all cameras over all flights. This was done on the grounds that there was no reason to think that the small variations in the estimated principal

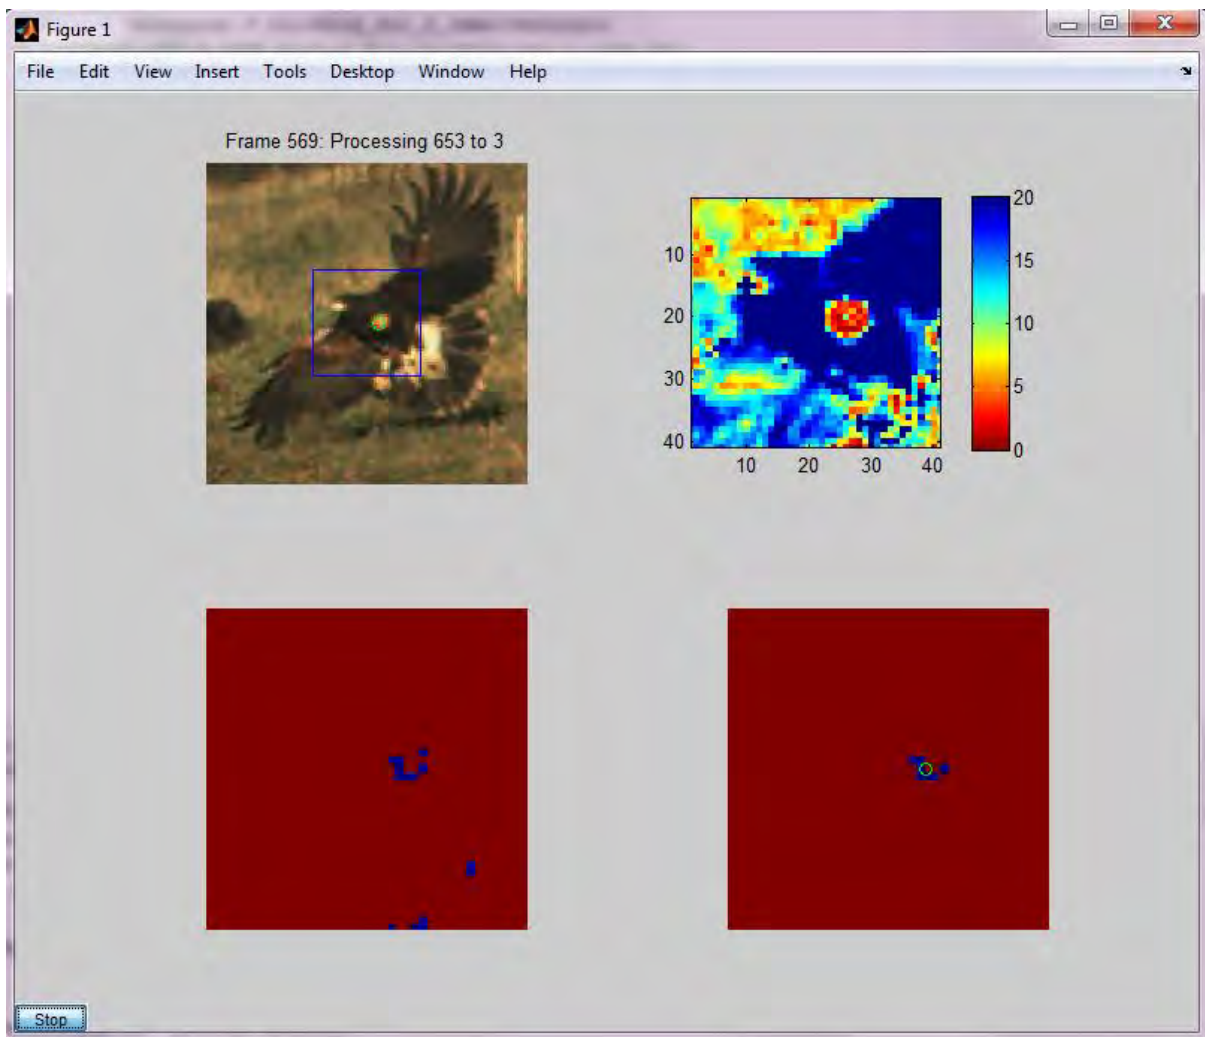


Figure 6: Screenshot of color tracking software graphical user interface. The top left panel shows the search area defined by the user (blue square) and the estimated centroid of the marker (green circle). The top right panel shows the Mahalanobis distance of the pixels within the search area, measured with respect to the color distribution of the user-defined color template. The bottom left panel shows all of the candidate marker pixels, identified as those pixels whose Mahalanobis distance falls below a set threshold. The bottom right panel shows the largest contiguous group of candidate marker pixels, the centroid of which is taken as an estimate of the marker position

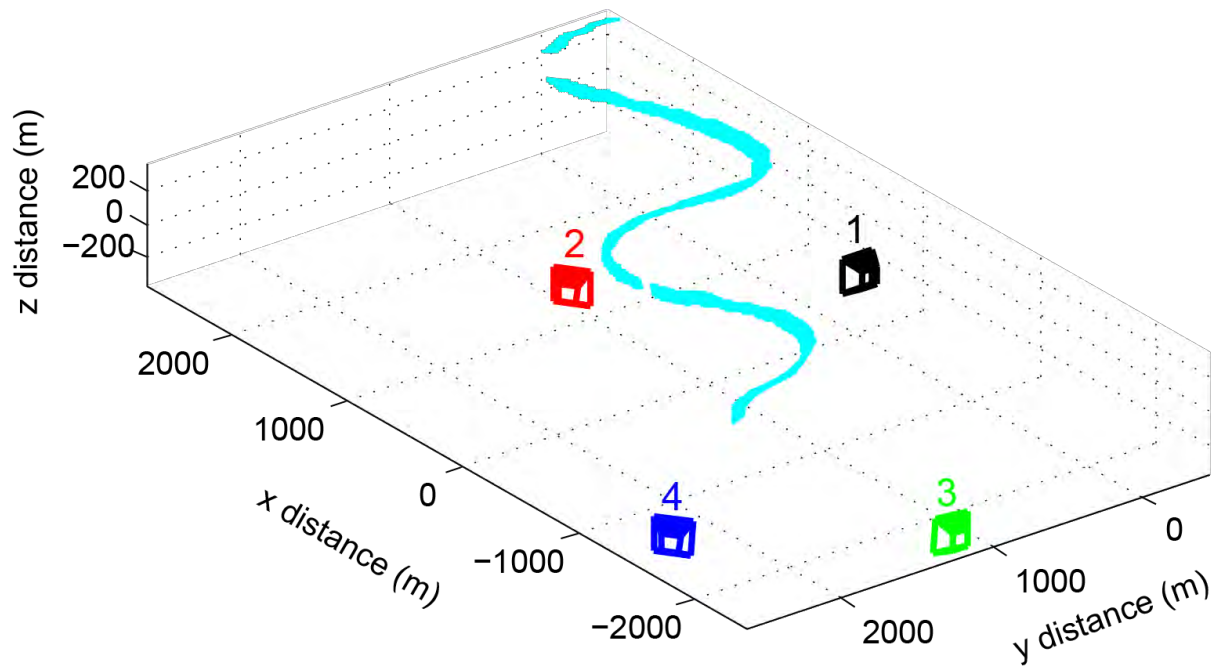


Figure 7: Example 3D reconstruction of the camera calibration for one test day, showing the estimated position and pose of the calibration object (cyan line) and estimated position and pose of the four cameras. The snaking path of the calibration object shows how it was moved through the test volume.

distances between flights or between cameras were real, so that the grand mean represents the best estimate of the true value of this parameter.

The resulting camera calibration parameters were used to solve, in a least squares sense, for the position of the lure mitt and of the two marker balls on the bird for each flight.

Error analysis

We established the measurement error of our photogrammetric technique by calculating the standard deviation of the estimated distance between the two calibration markers over the 14 camera calibrations that we ultimately used (see below). In the camera calibration software (see above), the distance between the calibration markers is treated as a known (i.e. fixed) parameter, so having run the calibration software already we then estimated the positions of these markers separately using the same software that we used to estimate the position of the lure mitt and of the two marker balls on the bird. Because the calibration object itself is a rigid body that had been moved through the test volume in a range of different positions and orientations, the standard deviation of our estimate of its length (i.e. the distance between the calibration markers) provides an isotropic estimate of the measurement error.

A fixed-effects ANOVA testing for differences in mean estimated calibration object length with respect to calibration date revealed the presence of small but statistically significant differences between days (one-way ANOVA: $F_{13,21174} = 48.4$, $p < 0.0005$). These differences could have been due to systematic error in the calibrations, or due to real variation in the length of the calibration object between days, but the effect size was very small, and we therefore considered it reasonable to pool our estimates of calibration object length for all days (Fig. 8). The grand mean and standard deviation of the estimated calibration object length were 1.054m and 0.023m, respectively, and 95% of the measurements fell within the interval [-0.032m, 0.036m] of the mean. It follows that the isotropic positioning error of our photogrammetric measurements is typically <0.04m.

Data selection

In order to deal practically with the volume of data that we had collected, whilst preserving the deliberate balance of the experimental design, we determined to analyse only the last 10 flights from each of the 5 birds. Given that we recorded at least 20 good flights from each bird in total, it follows that each of the birds had already experienced a training period comprising at least 10 good flights before the block of 10 good flights that we actually analysed. Future work will test whether there is evidence of learning during the training phase, but at present the data remain unanalyzed.

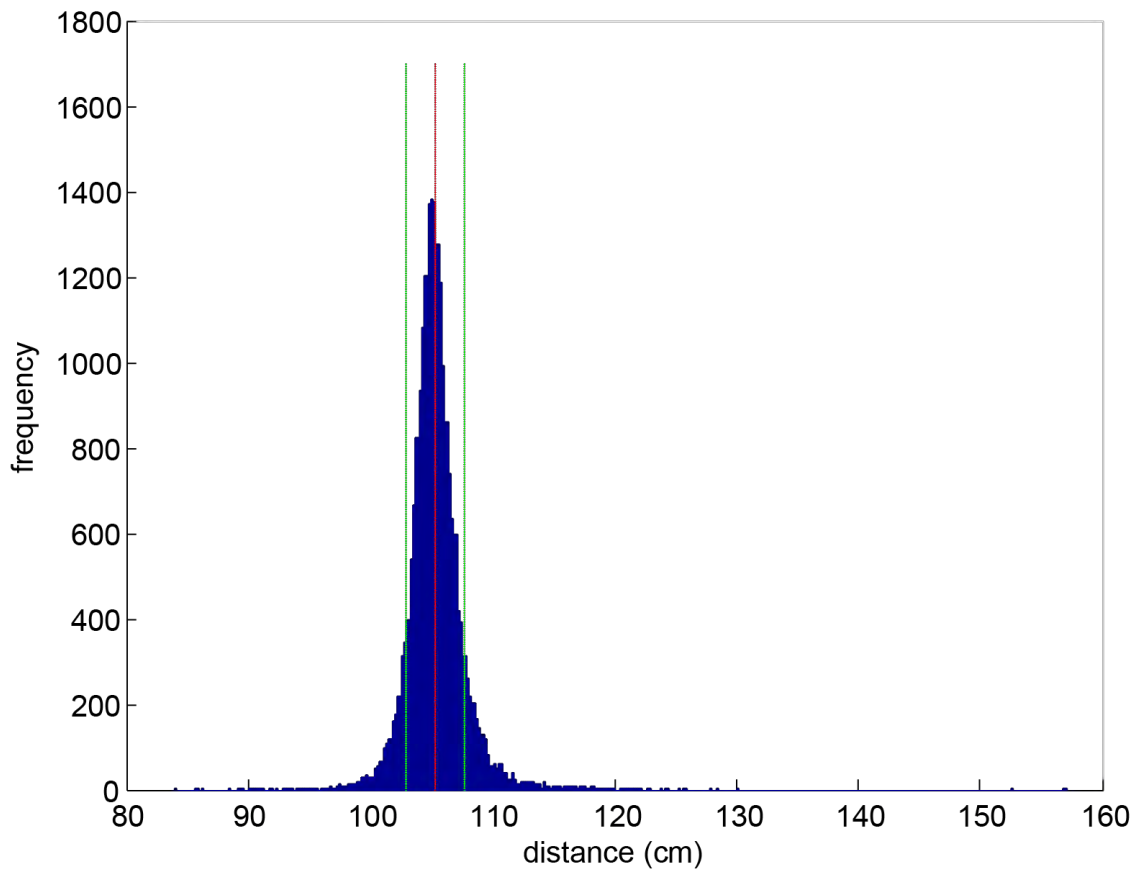


Figure 8: Distribution of estimated calibration object length for all of the 14 calibrations that we used, showing the mean estimated length (1.054m, red) and standard deviation (± 0.023 m, green). The measured distance assumed in the calibrations was 1.05m, and the true length of the calibration object should not have varied by more than a few millimetres between calibrations. The dispersion of this distribution therefore provides an estimate of the isotropic measurement error.

Air-to-air attacks in Peregrine Falcons

Animals

Two trained peregrine falcons *Falco peregrinus* (“Lizzie”: ♀; “Izzy”: ♀) were flown at a target towed behind a remotely piloted aircraft, giving 14 test flights in total.

Equipment

Each of the peregrines carried a GPS receiver logging position and groundspeed at 5Hz (Qstarz BT-Q1300) and a high-definition video camera looking forward over the head (HD720P Mini DV, 30fps). The equipment weighed 0.031kg in total, and was carried dorsally on a commercially available harness (Marshall Radio Telemetry, Trackpack) comprising a curved plastic mounting plate held in place between the shoulders by a pair of flat tubular Teflon ribbons drawn once around the bird’s body in a figure-of-eight pattern. Each bird also carried a tail-mounted radio transmitter for tracking purposes (Marshall Radio Telemetry Micro Transmitter). The bird was filmed from the ground by two JVC GC-PX100 HD video cameras mounted on tripods spaced approximately 20m apart. The cameras were zoomed out so as to cover as much of the test volume as possible, and were set in a convergent configuration so that their fields of view overlapped. The positions of the two cameras were recorded by leaving a GPS unit stationary on the camera, and all of the cameras were synchronised using a digital stopwatch display.

Experimental protocol

The target – a leather falconer’s lure with food reward attached – was towed by a small remotely piloted aircraft on c. 10m of kite line attached to the aircraft’s tail boom. A Qstarz BT-Q1300 GPS logger was attached to the airplane close to the center of mass, and an HD720P Mini DV camera was fixed beneath the aircraft so that it faced backwards towards the towed lure. A third Qstarz BT-Q1300 GPS logger was attached directly to the lure. The bird was released by the falconer at the start of each flight test, and was allowed to gain height. The aircraft was then launched and operated by a pilot on the ground. On windy days, the plane and lure were launched while the bird remained hooded, to encourage an attack upon first sight of the target and to reduce the risk of the bird flying off. The birds chased the target as it swung behind the aircraft, which was flown through a series of maneuvers by the pilot. As a safety precaution, the lure was released from the aircraft as soon as the bird had made contact with it.

Data synchronization

We plotted the GPS groundspeed of the lure and bird against time, and used this to establish the start time of the pursuit, which we defined as the time

at which the bird or plane first began moving – whichever was the later of the two. The start of the pursuit was also identified separately in the ground-based video recordings, and the GPS time at the start of the pursuit was used as a reference to convert the timestamp of the ground-based video recordings to absolute time. The absolute time of the end of the pursuit was identified by using the ground-based video cameras to establish the point of intercept, which allowed us to identify the point of intercept in the three GPS data streams recorded from the bird, aircraft, and the lure. For the few flights where the intercept occurred out of view of the ground-based video cameras, the onboard video data from the back of the bird or the plane were instead used to establish the time of the intercept.

Results and Discussion

Air-to-ground attacks in Peregrine Falcons

Data processing

Because there was no way of knowing from first principles when the peregrines first switched from loitering to target-directed flight behavior, we used an arbitrary but objective empirical criterion to define the time t_0 at which to begin analysing the flight trajectory on a given pass. For each pass, we defined t_0 as the beginning of the time interval during which the peregrine was both closing range on its target and turning monotonically towards it. We attribute no particular biological meaning to this definition of t_0 , which is merely intended to provide a consistent and objective criterion for selecting sections of flight for analysis. Thus, there is no reason to think, for example, that t_0 must coincide with the time at which the bird initiated target-directed flight behaviour. The peregrines usually followed a shallow, curving trajectory from time t_0 to the point of intercept.

The observed flight trajectories were almost planar, and we treat them as planar in all of the following analyses. We define the plane of each pass in two distinct ways. The first two principal components of the bird's 3D flight trajectory measure its 2D position within the plane of best fit, and in one set of analyses we project therefore the 3D position data for each attack pass onto these first two principal components. This approach has the advantage of capturing as much of the positional variation as possible within a single analysis, but has the disadvantage of mixing accurate latitudinal and longitudinal data with less accurate altitudinal data. It also prevents us from exploring the effects of varying the assumed starting point of the trajectory about t_0 without simultaneously affecting the plane of best fit. For this reason, we begin by adopting the simpler approach of analyzing only the horizontal components of the attack trajectory, by taking only the latitudinal and longitudinal components of the GPS position data. Because all of the attack trajectories were shallow, this approach captures most of the positional variation, whilst avoiding the

disadvantages of considering the first two principal components of the trajectory instead.

We used forwards differencing to estimate the bird's horizontal ground velocity from the GPS latitude and longitude data, and estimated the bird's instantaneous track angle as the azimuth of its ground velocity vector. We calculated the instantaneous line-of-sight angle to the target as the azimuth of the target with respect to the bird, and we used central differencing to calculate the turn rate and the line-of-sight rate from the track angle and the line-of-sight angle data, respectively.

Analysis of results

It is clear from the latitudinal and longitudinal GPS position data that there is no simple geometric rule that is capable of describing all of the observed trajectories. In particular, the deviation angle between the velocity vector of the bird and its line of sight to target is not held constant during the most highly curved flights (Figs. 9), as would have to be the case if the birds were following a deviated pursuit course as hypothesized by Tucker (2000b). In contrast, the GPS data do reveal a positive correlation between horizontal turn rate and horizontal line-of-sight rate (Fig. 10), as required under the hypothesis that turn rate is commanded through some form of proportional navigation guidance law (i.e. in proportion to the line-of-sight rate). Our estimates of line-of-sight rate and turn rate are inherently noisy, because they are both calculated by differencing GPS position data. Hence, to provide a more thorough test of the hypothesis that turn rate is commanded through proportional navigation, we next asked whether the detailed shapes of the measured flight trajectories could be explained under the hypothesis of proportional navigation.

Given GPS estimates of a peregrine's initial position and flight direction relative to its target at t_0 , it is straightforward to simulate the trajectory that the bird would have followed under pure proportional navigation (Eq. 1) with a given navigation constant N . This involves solving the equations of motion for guidance under proportional navigation, which have a closed-form solution against a stationary target (Shneydor, 1998). For simulation purposes, it proved easier to solve these equations numerically, but we verified our numerical trajectories against the closed-form solution, as a check upon the accuracy of the integration scheme. We simulated the trajectory that would have been followed under different values of N given the initial conditions at the start of each pass, and identified the best-fitting value of N separately for each pass. We did this by minimizing the L1-norm of the distance between the measured and simulated flight trajectories in the horizontal plane, using golden section search and parabolic interpolation subject to the constraint $N \geq 1$. The L1-norm measures the mean absolute distance between the measured and simulated flight trajectories, and was $<0.55\text{m}$ (median: 0.15m) at the best-fitting values of N . Expressed as a percentage of the distance to target at t_0 , this represents a mean absolute difference of $<2.3\%$ (median: 0.6%). It follows that the trajectories are well explained under the hypothesis that peregrines use proportional navigation, provided that the value of the navigation constant is allowed to vary between passes.

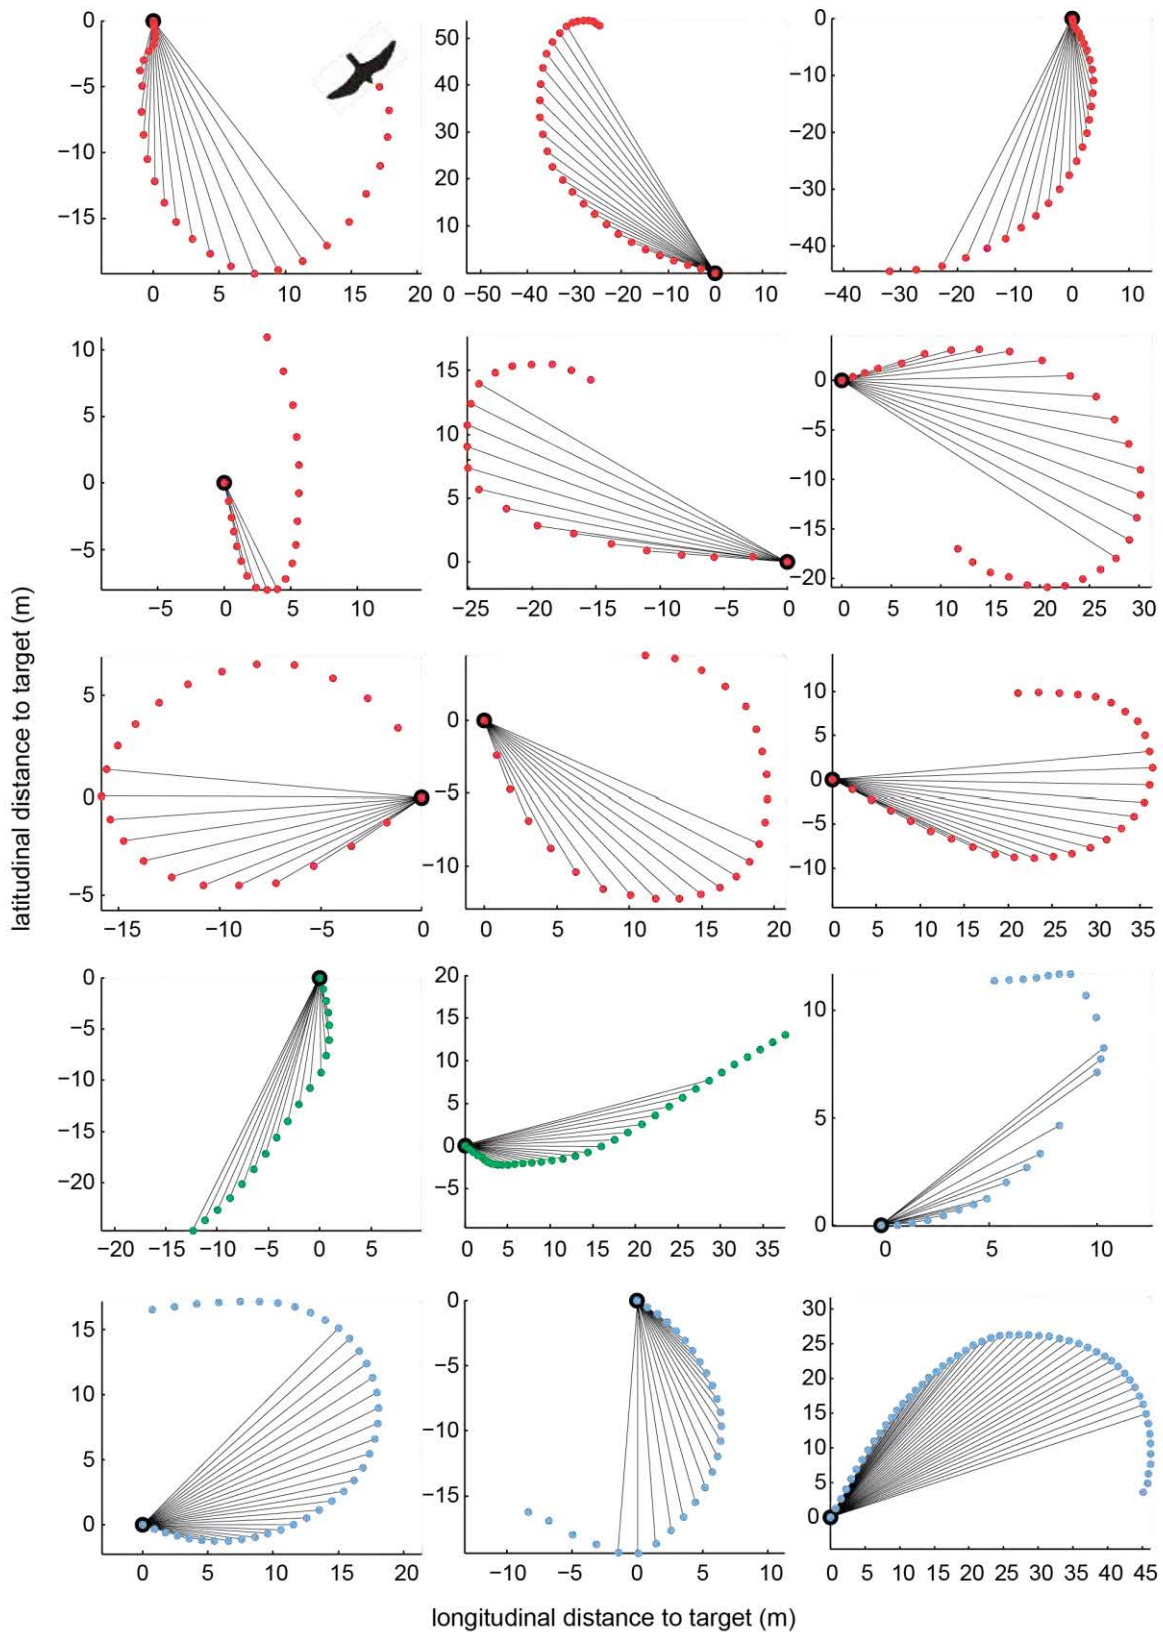


Figure 9. Horizontal GPS position (coloured points) plotted for the 15 most highly curved peregrine air-to-ground attack trajectories. Colours identify individuals: Bella (red), Nina (green), Weirdo (blue). Black lines show the instantaneous line of sight to target (black dot), from t_0 to end of pass. Note that the deviation angle between the velocity vector and line-of-sight to target is not constant, demonstrating that the birds did not use deviated pursuit.

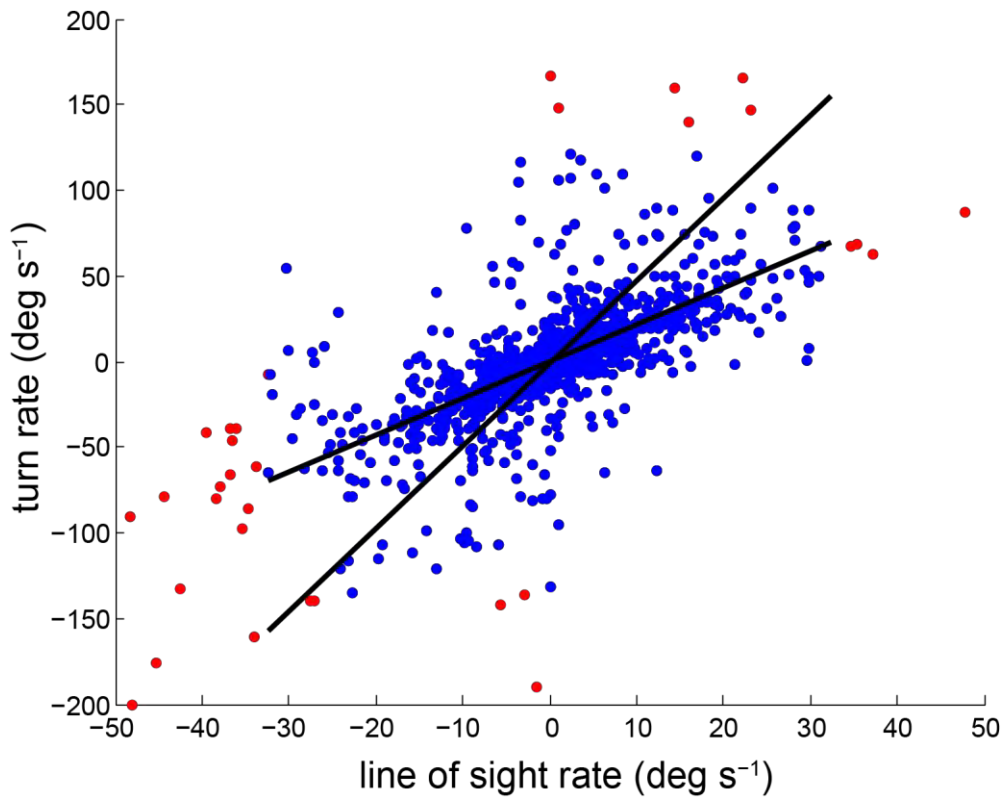


Figure 10: Estimated turn rate versus estimated line of sight rate for all measured instants during all measured peregrine air-to-ground attack passes. Points in red are outliers, defined as points falling within the upper and low 2.5% percentiles for turn rate and line of sight rate, and are likely to reflect measurement error rather than real data. The two black lines denote the linear regression of turn rate on line of sight rate, and line of sight rate on turn rate, respectively. Although the data are noisy, it is clear that turn rate is positively correlated with line of sight rate, as required under the hypothesis that the birds use proportional navigation.

Because the peregrines may well have initiated their target-directed flight behaviour before or after the time that we have designated as t_0 , we next iterated the start time of the simulations so as to identify the earliest start time for which the L1-norm of every trajectory was $<0.17\text{m}$ at the best-fitting value of N . This arbitrary but objective criterion was chosen to keep the median of the L1-norm for all passes the same as it had been when starting the simulations at t_0 . The resulting simulations match the measured trajectories very closely indeed (Fig. 11), and show some potentially interesting variation in the best-fitting value of N . We therefore conducted a graphical sensitivity analysis to test whether this variation was genuine, plotting the envelope of simulated trajectories for navigation constants $\pm 10\%$ of the best-fitting value of N (Fig. 11). The best-fitting simulation always matched the curvature of the measured trajectory closely, but only for the six most highly curved trajectories did the $\pm 10\%$ bounds have curvatures that were clearly distinguishable from the measured trajectory (Fig. 11). It follows that we are only able to estimate N with sensitivity better than 10% for those six trajectories. For the less highly curved trajectories, the $\pm 10\%$ bounds are not clearly distinguishable from the measured trajectory, so although the simulations still match the measured trajectory very closely, the curvature of the measured trajectory is consistent with a wide range of values of N . Nevertheless, because the best-fitting value of N for the six most highly curved passes varied between 2.5 and 4.0, which is an interval covering a range $\pm 23\%$ of its midpoint, we can be confident that at least some of the variation in the best-fitting values of N reflects genuine variability in its true value.

Finally, to deal with the fact that the measured trajectories were not quite horizontal, we projected the three-dimensional GPS coordinates for each pass onto the inclined plane defined by their first two principal components, and reran the simulations. These inclined projections (Fig. 13) were very similar in shape to the horizontal projections (Fig. 13), and we found little change in the estimate of N for each pass. We therefore conclude that our peregrines turned with respect to a stationary ground target in a manner consistent with the use of proportional navigation, and that they varied the value of the navigation constant N between, though not within, passes. Many advanced missile guidance systems use a form of augmented proportional navigation, in which the value of N is effectively updated in light of information about the target to give best performance against some criterion such as time to intercept (Shneydor, 1998; Siouris, 2004). We hypothesise that variation in the navigation constant N may be similarly adaptive in peregrines, and are now in the process of testing this hypothesis formally using the trajectory data that we have collected during air-to-air attacks against maneuvering targets (see below).

Air-to-ground attacks in Harris' Hawks

Data processing

We took the estimated position of the marker carried on the breast of the bird as our estimate of the bird's instantaneous position. This marker was not fully visible in a small number of cases, and in those cases we instead used the

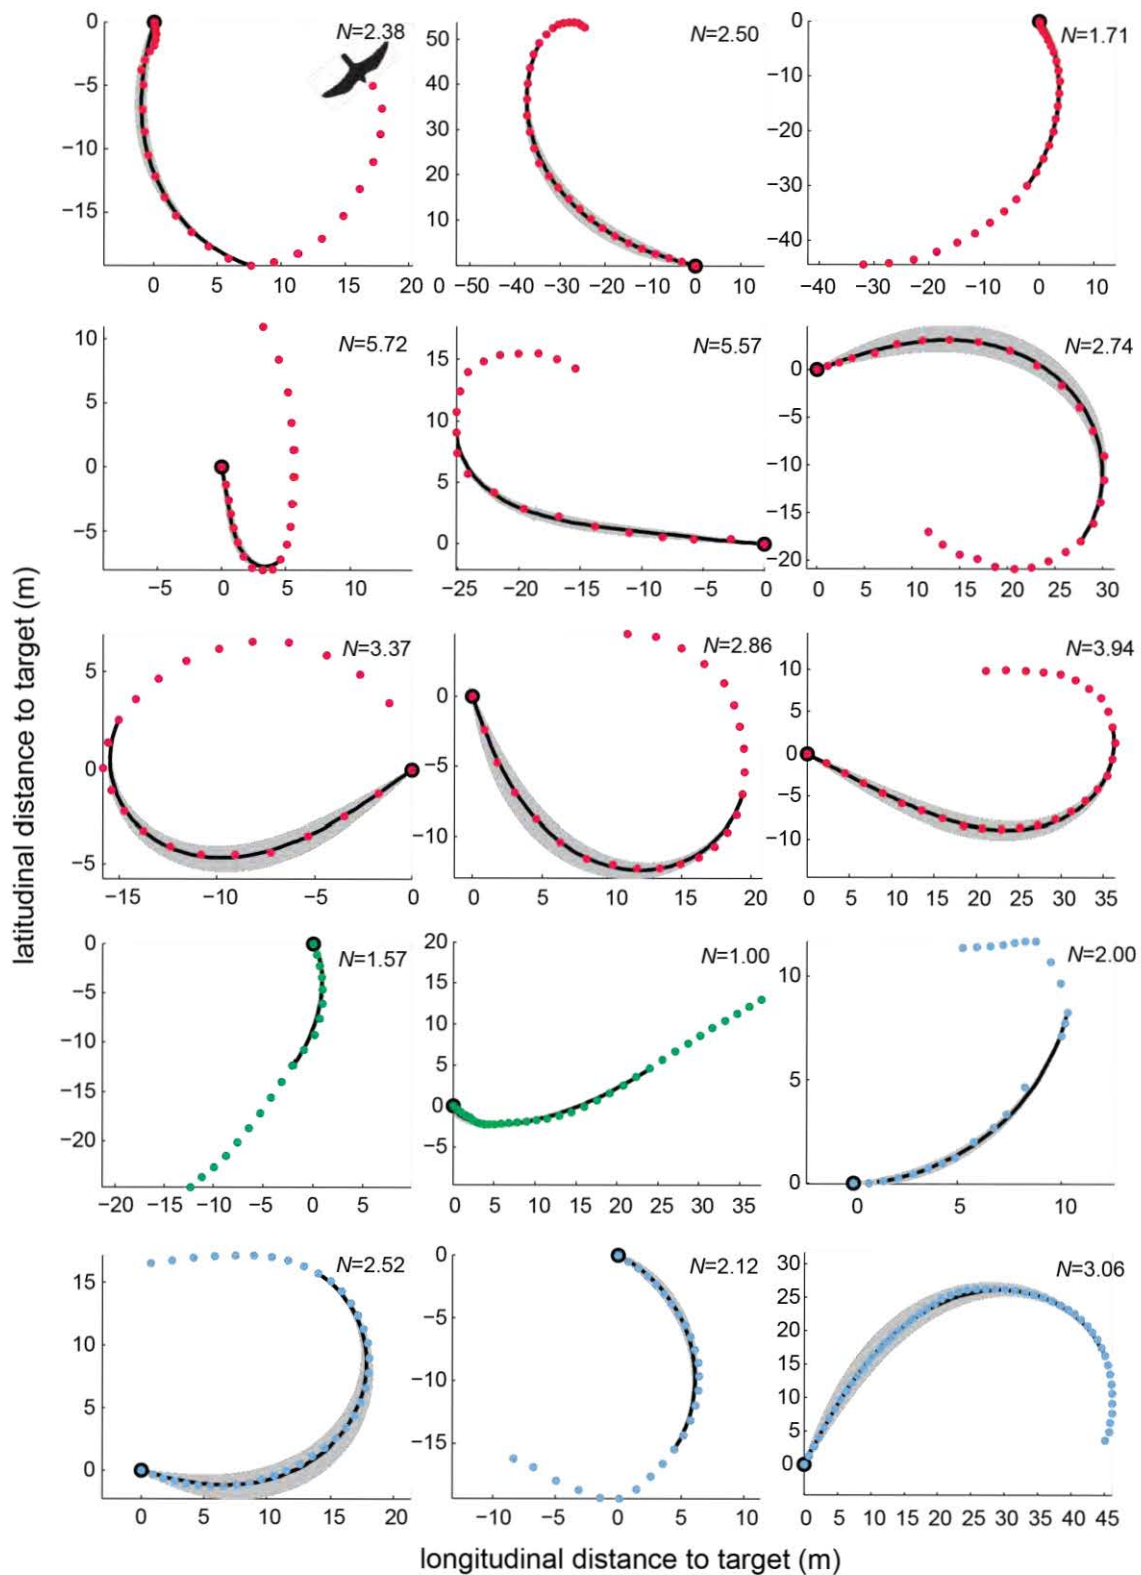


Figure 11. Measured GPS positions (coloured points) overlay on trajectories simulated under the hypothesis of proportional navigation (black lines), for the 15 most highly curved peregrine air-to-ground attack trajectories (see Fig. 12 for all recorded attack trajectories). Grey shading (not always visible) indicates the envelope of simulated trajectories for navigation constants $\pm 10\%$ of the best-fitting value (N) shown at the top right of each plot. Colours identify individuals, as in Fig. 2. Note the very close match between the measured and simulated trajectories, consistent with the use of proportional navigation, and the variation in N among passes.

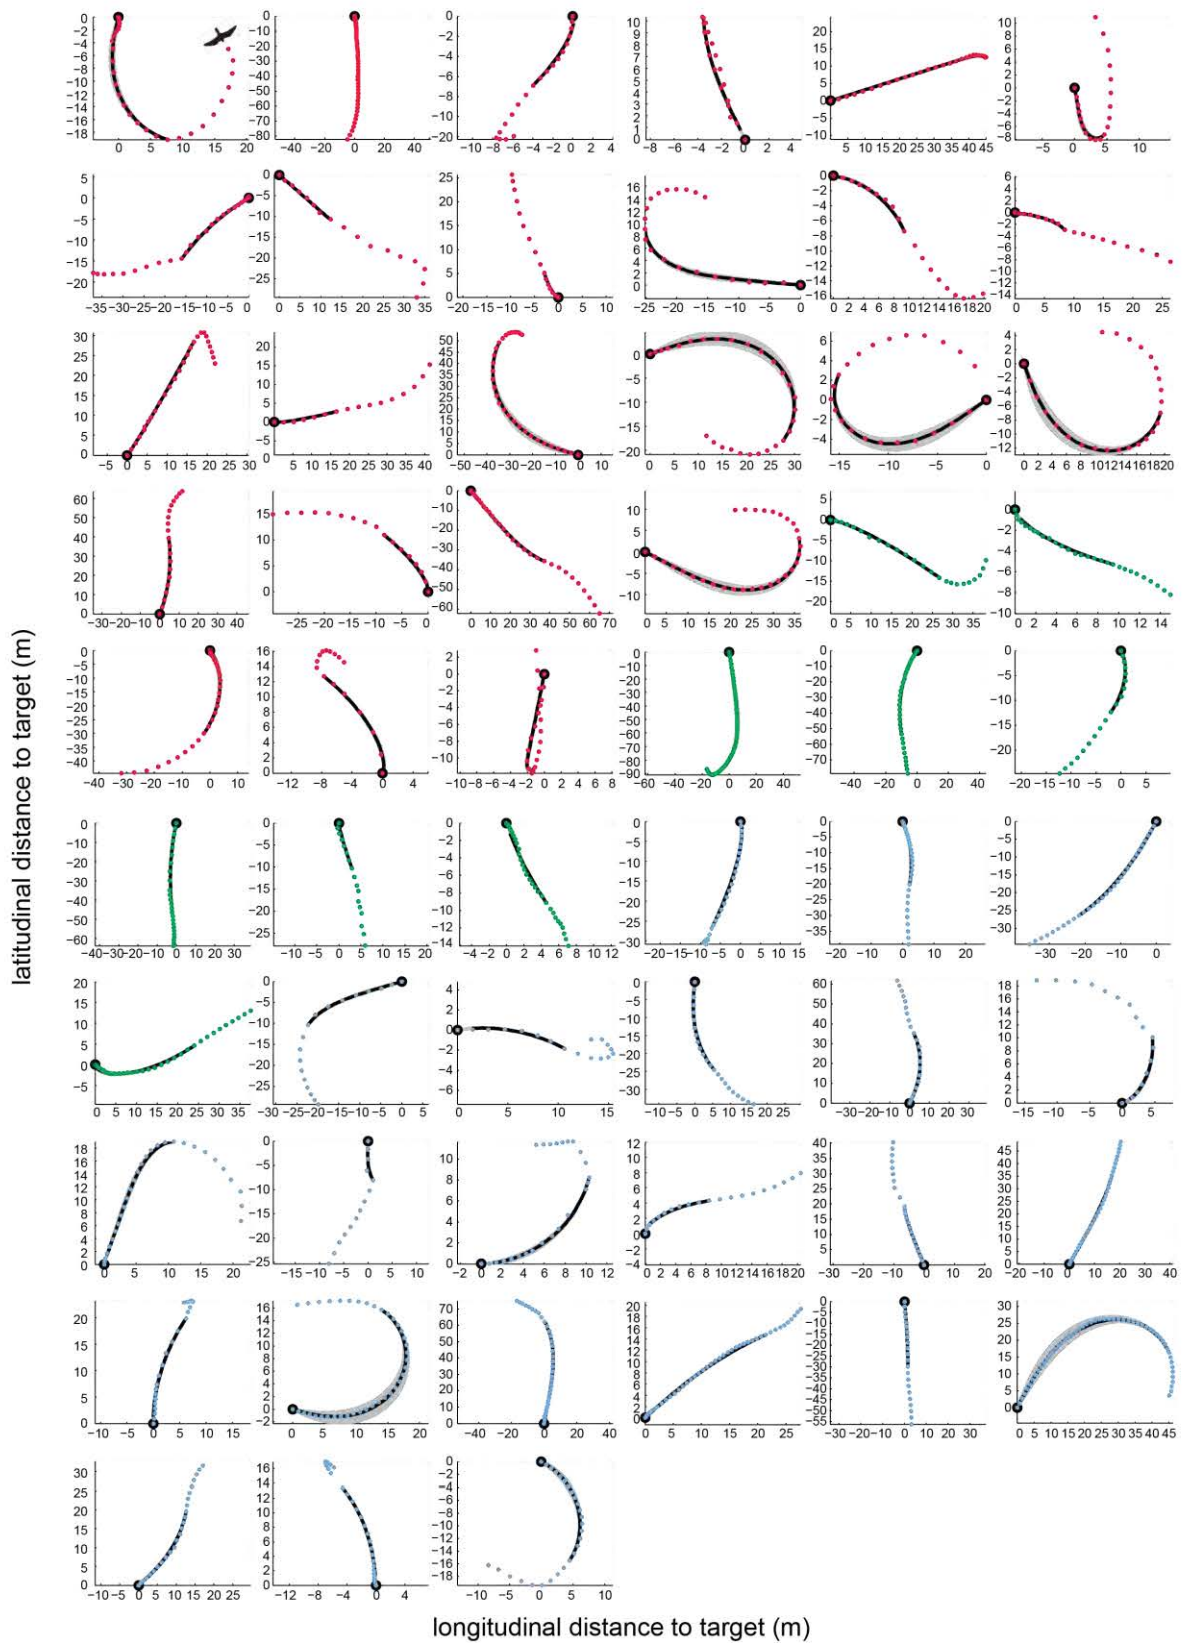


Figure 12. Measured GPS positions (coloured points) overlain on trajectories simulated under the hypothesis of proportional navigation (black lines), for all peregrine air-to-ground attack trajectories. See Fig. 11 legend for explanation.

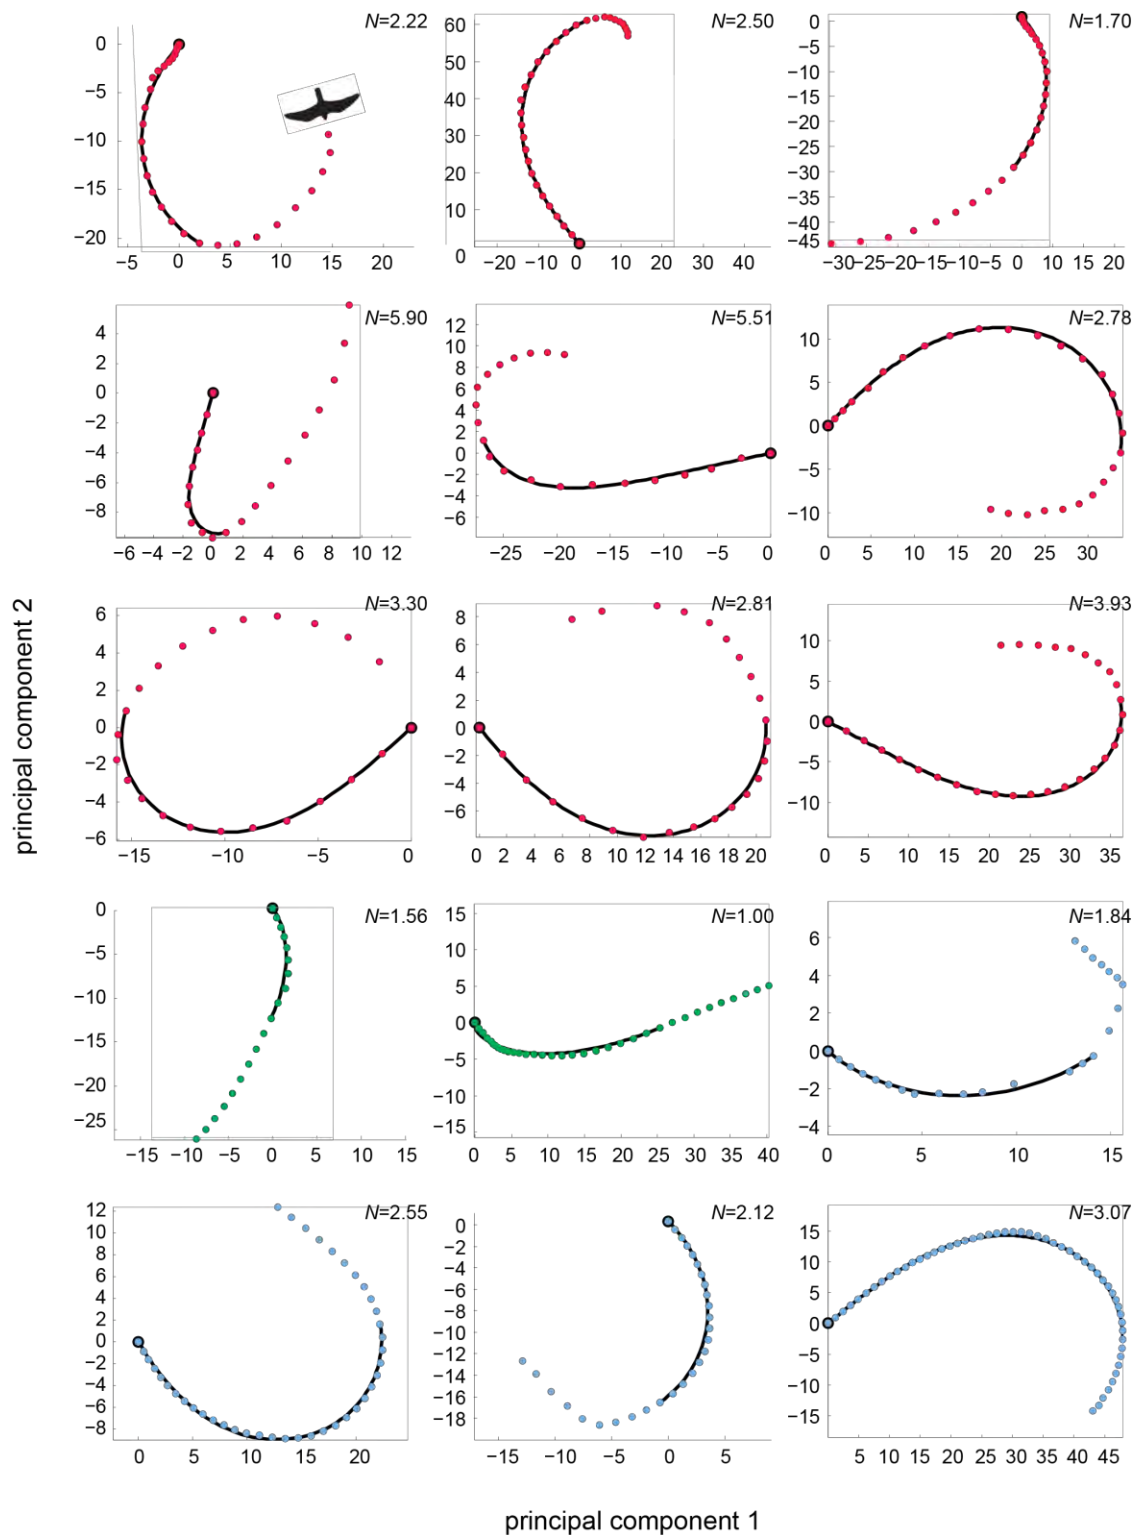


Figure 13: First and second principal components of three-dimensional GPS position (coloured points) overlain on trajectories simulated under the hypothesis of proportional navigation (black lines), for the 15 most highly curved attack trajectories. Colours identify individuals. The first and second principal components represent a projection of the three-dimensional position data onto the plane that best fits the measured trajectory for each pass. Note the very close match between the measured and simulated trajectories, confirming the use of proportional navigation, and the variation in N between passes. Compare with Fig. 11, showing the horizontal components of the same trajectories.

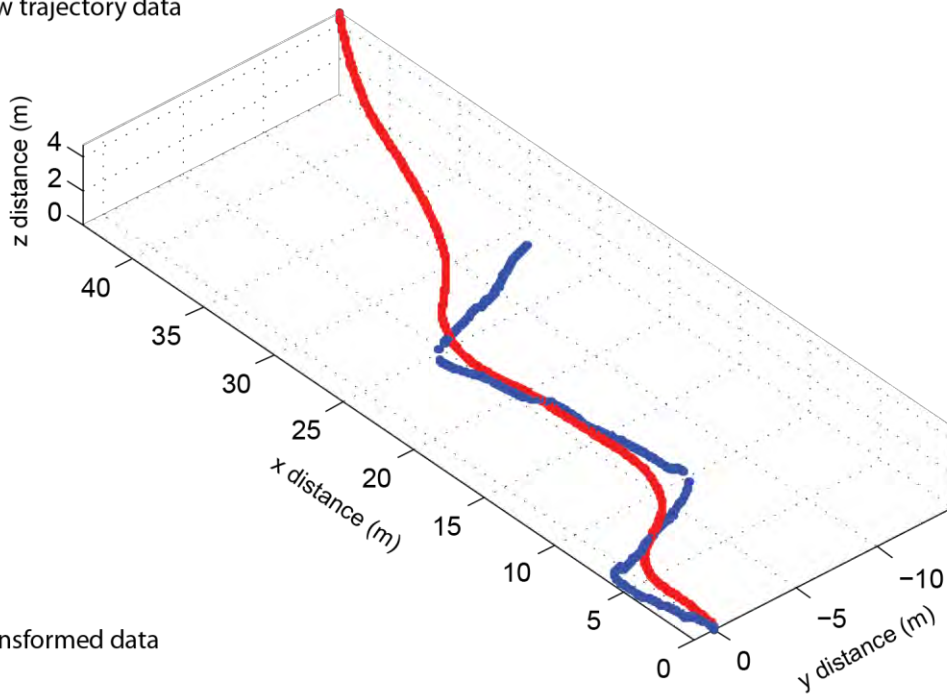
estimated position of the marker carried on the back of the bird. Because we had selected flatly sloping fields for the flight tests, the chase trajectories that we recorded were all close to planar, but the planes of the trajectories did not in general coincide with the Cartesian coordinate system generated by the camera calibration software (Fig. 14A). The first two principal components of the bird's 3D flight trajectory measure its 2D position within the plane of best fit. We therefore computed the principal components of the bird's 3-dimensional flight trajectory, and used these to rotate the Cartesian coordinate system generated by the camera calibration software so that the $x'y'$ -coordinates of the rotated coordinate system corresponded to the plane of best fit (Fig. 14B). All of the subsequent analysis is performed on the $x'y'$ -components of the position data in this rotated coordinate system.

Determining the bird's flight velocity involves differencing its position, but because we had used high-speed video data at 250Hz, the bird only travelled a short distance between consecutive video frames. We were therefore obliged to calculate the bird's flight velocity by differencing over a longer time interval than the separation between the video frames. Since the relative positioning error of our photogrammetric method was comparable to the relative positioning error of the GPS units used in our work with peregrines (<0.05m in both cases; see above) we chose to calculate the bird's flight velocity by central differencing the position estimated on every 50th frame (Fig. 15A,B), which resulted in the same 0.2s temporal separation between the position fixes used to estimate flight velocity as was already present in the 5Hz GPS data that we had collected with the peregrines. We then calculated the track angle of the bird as the azimuth of the resulting velocity data (Fig. 15A,B), and determined the line-of-sight angle to target at the corresponding point in time (Fig. 15A,B). Finally, we computed the turn rate of the bird and the rate of change in its line-of-sight angle by central differencing the track angle and the line-of-sight angle, respectively, using the same 0.2s temporal separation as before (Fig. 15C). Hence, because the relative positional error (<0.05m) and temporal separation (0.2s) used in differencing the position data are similar in both the GPS-based and photogrammetric analyses, it follows that the measurement error in all of the measured quantities should be similar in all of the various analyses presented in this report.

Analysis of results

The track angle of the bird closely follows the line-of-sight angle to the target in all of the 50 flight trajectories that we analysed, albeit with a small time delay (Fig. 16). It follows immediately that the Harris' hawks were following a pursuit course in which the deviation angle was held approximately constant: in most cases, the deviation angle was close to zero (Fig. 16), implying that the birds were normally using a pure pursuit strategy. To provide a quantitative estimate of the delay in the pursuit dynamics, we calculated the correlation coefficient between the track angle and the line-of-sight angle for lags within the interval ± 0.5 s (Fig. 17). The correlation coefficient was low on a handful of the flights, which may indicate that the bird was not always locked-on to the lure in these cases. In most cases, however, the correlation coefficient had a maximum >0.8

A) Raw trajectory data



B) Transformed data

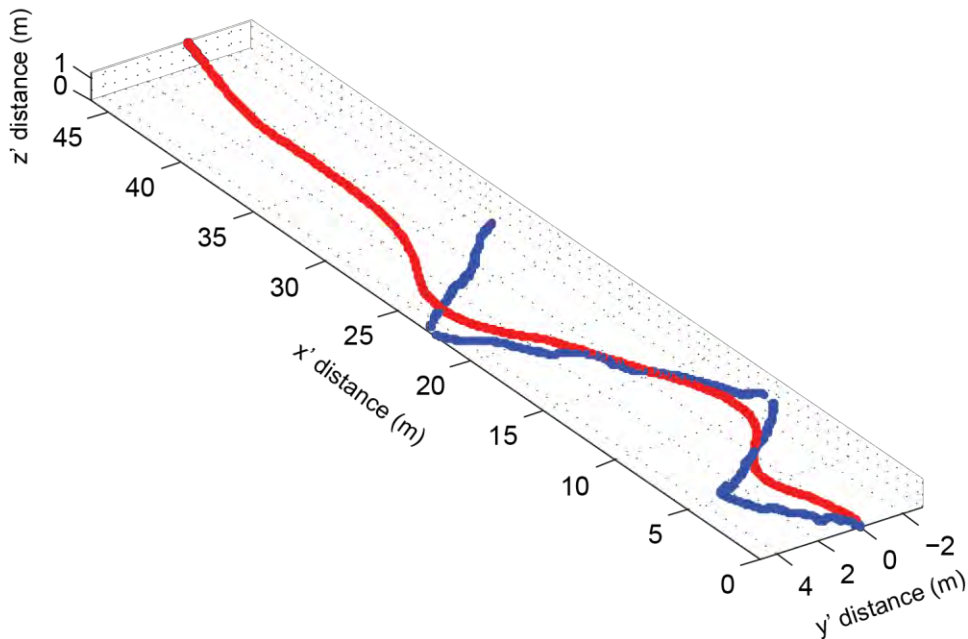


Figure 14: Example trajectory of a Harris' hawk (red) chasing the towed lure (blue) from one flight test, with interception occurring where the two trajectories meet at the origin. A) Data plotted in the Cartesian coordinate system generated by the camera calibration software (see also Fig. 7); B) Data plotted in the rotated Cartesian coordinate system in which x' and y' denote the first and second principal components of the bird's flight trajectory. The $x'y'$ plane is the plane which best fits the measured trajectory (i.e. the plane which captures most of the measured variation in position).

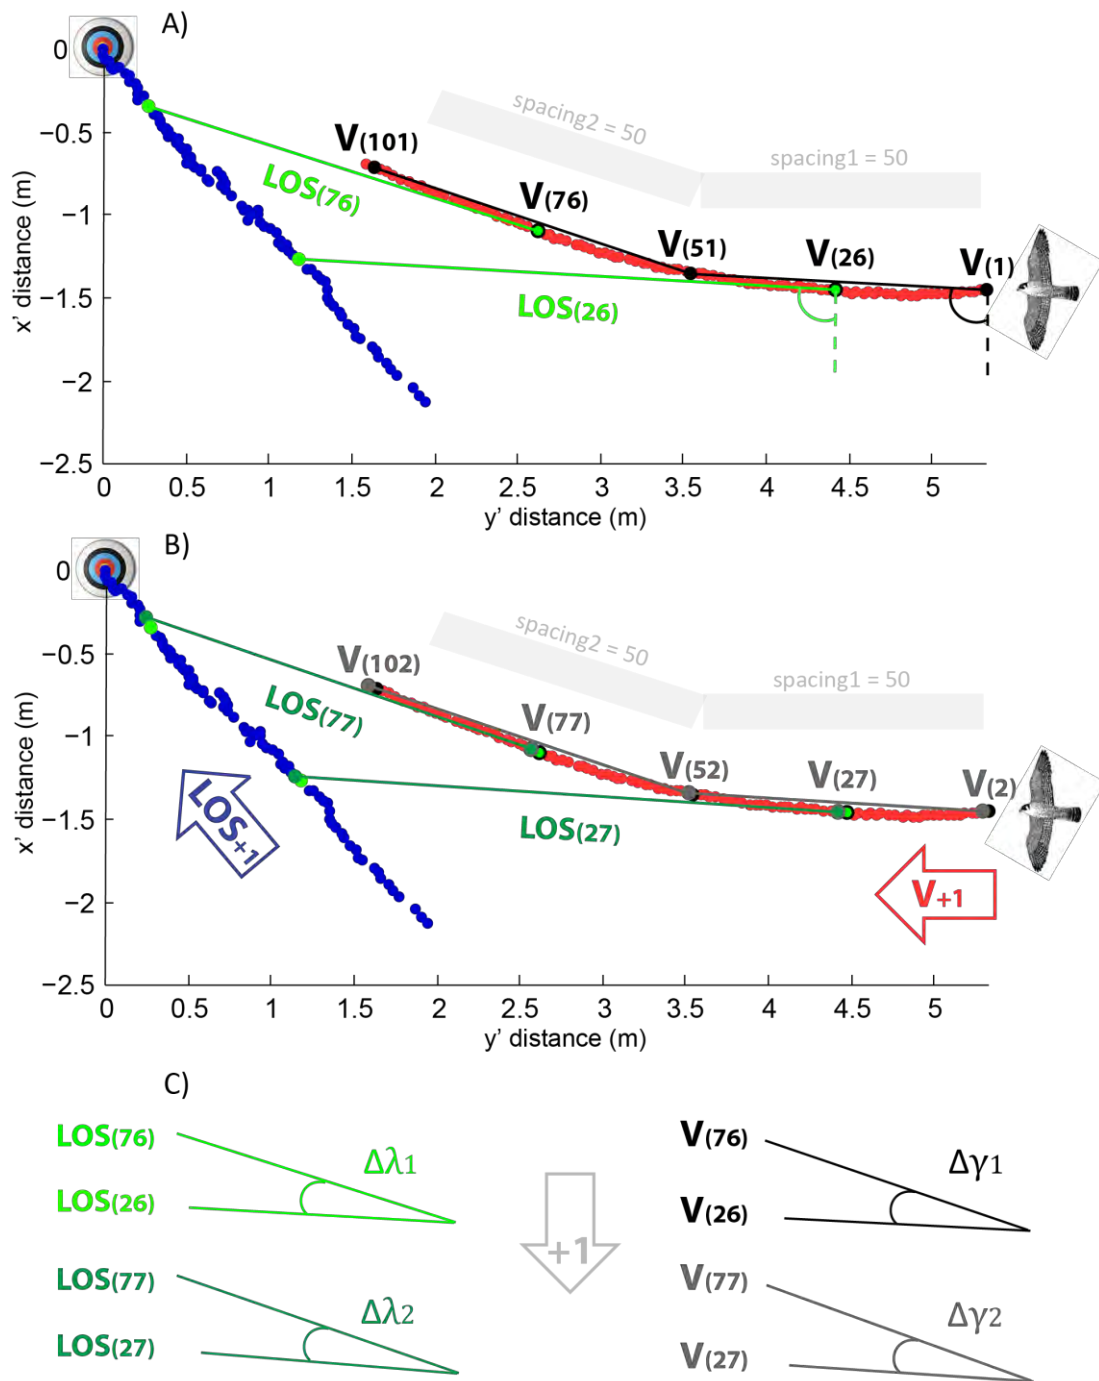


Figure 15: Calculation of line-of-sight angle, track angle, line-of-sight rate and turn rate from our photogrammetric measurements of bird position (red points) and lure position (blue points). LOS: line of sight; V: velocity vector; numerical subscripts denote camera frame numbers.). A) Track angle at Frame 26 is calculated as the bearing of the velocity vector V_{26} , calculated from the difference in position between Frames 1 and 51. Line-of-sight angle at Frame 26 is calculated as the bearing of the line of sight LOS_{26} drawn from bird to target at Frame 26. The track angle and line-of-sight angle at Frame 76 are calculated similarly as shown. B) Calculation of the same quantities at Frames 27 and 77, to show how the calculation preserves the 250Hz sampling rate of the data, despite using a temporal separation of 0.2s for the calculations. C) Turn rate and line-of-sight rate at Frames 51 and 52 are calculated from the difference in track angle $\Delta\gamma$ and line-of-sight angle $\Delta\lambda$ at Frames 26 and 76, and 27 and 77, respectively.

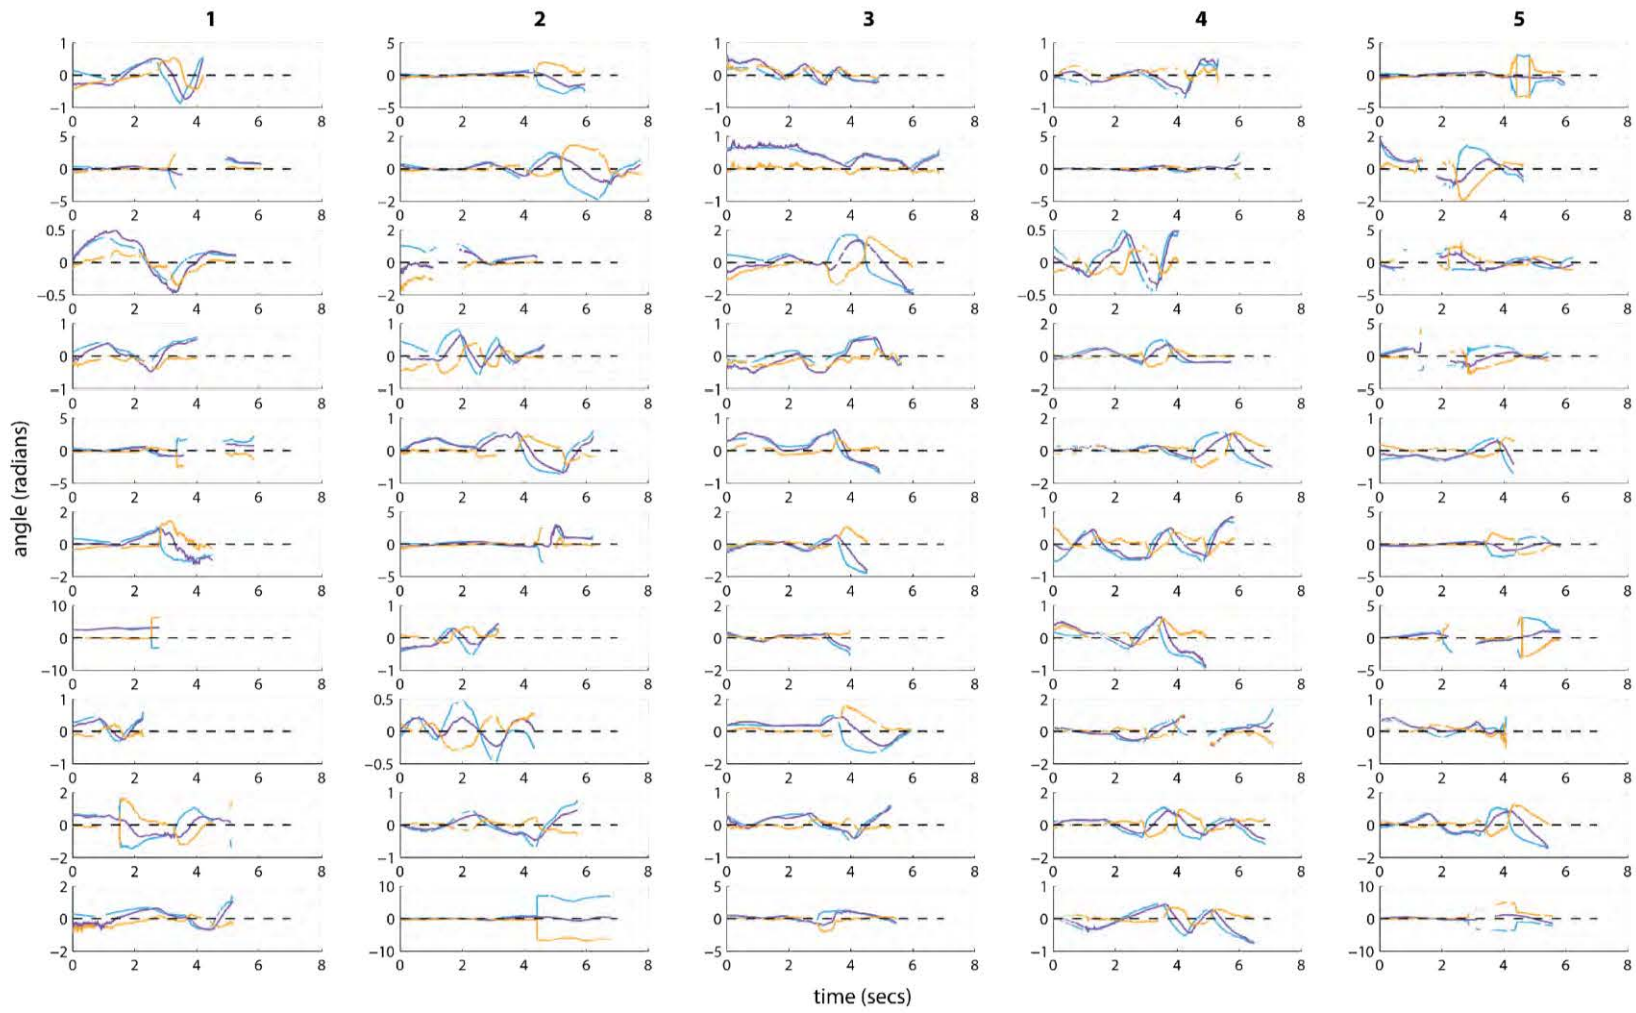


Figure 16. Track angle γ (blue), line-of-sight angle λ (purple), and deviation angle $\delta = \lambda - \gamma$ (orange) plotted against time for all of the air-to-ground attack passes by Harris' hawks that we analysed. Each column (1-5) corresponds to a different bird.

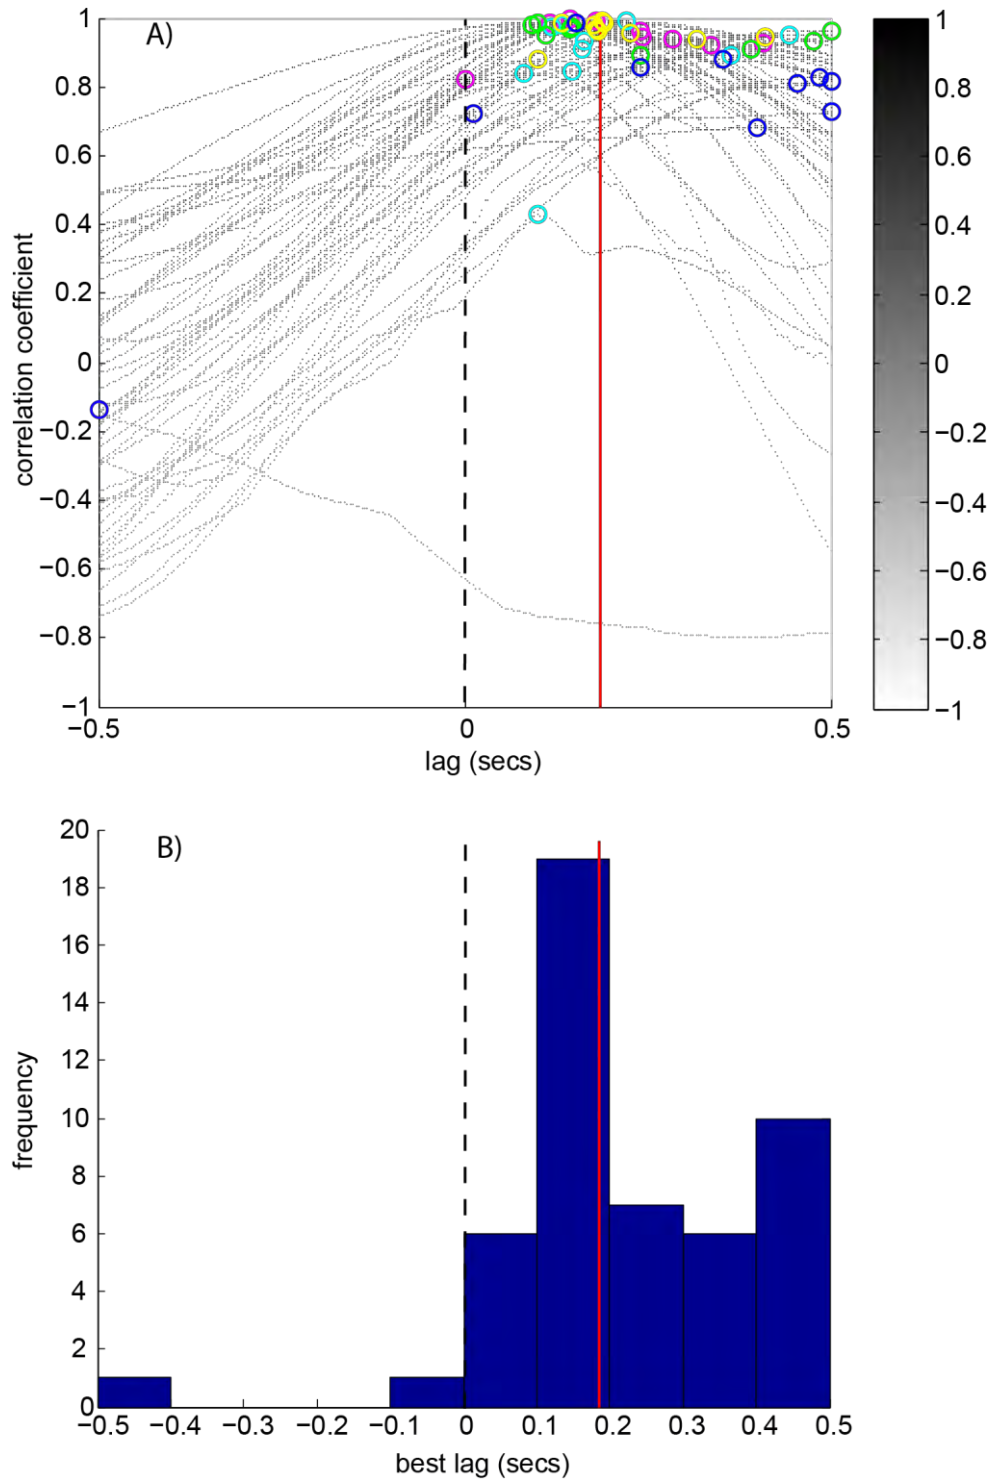


Figure 17: A) Correlation coefficients of track angle γ against line-of-sight angle λ for assumed lags in the interval ± 0.5 s for all analysed air-to-ground attacks in Harris' hawks. The lines plotting how the correlation coefficient for each flight varies with time lag are colored gray according to their value, to draw attention to the regions of highest correlation coefficient. The point of maximum correlation coefficient for each line is identified by a colored circle (color identifies individual birds), and the median of the lags corresponding to the peak correlation coefficient for each flight (0.19s) is shown as a vertical red line. B) Histogram of best lags, defined as the time lag corresponding to the point of maximum correlation coefficient for each of the flights in (A). The red line denotes the median of the lags corresponding to the peak correlation coefficient for each flight (0.19s).

and was maximal having lagged the track angle by approximately 0.19s (median of the lags corresponding to the peak correlation coefficient for each flight; Fig. 17). This is confirmed by Fig. 18, which plots track angle and line-of-sight angle against time for all 50 flight trajectories, with a lag of 0.19s applied to the track angle data.

As discussed in the Introduction, it is possible in principle to implement pure or deviated pursuit by using pure proportional navigation with a navigation constant of $N = 1$. Conversely, it is easily seen that if a bird is engaged in a pure or deviated pursuit, then its turn rate must necessarily equal its line-of-sight rate (albeit with a possible lag in the dynamics). It follows that any predator engaging in pure pursuit will look as if it is engaging in pure proportional navigation with $N = 1$. This is shown directly in Fig. 19, which plots turn rate and line-of-sight rate against time. Although the plots are noisier than they are for the track angle and line-of-sight angle (Figs. 16, 18), on account of the numerical differencing, it is obvious that the turn rate indeed follows the line-of-sight rate. Fig. 20 plots the same data again, but with the turn rate lagged by 0.19s.

As has already been mentioned, the Harris' hawks were normally using a pure pursuit strategy, which can be seen from the fact that the lagged track angle is similar to the line-of-sight angle at most times (Fig. 18). Nevertheless, there are some sections of flight during which the track angle differs from the line-of-sight angle by some near-constant amount. Thus, there are some sections of flight during which the lagged track angle is substantially different to the line-of-sight angle, but during which the lagged turn rate and line-of-sight rate are approximately equal. These sections of flight correspond to a deviated pursuit, and although this observation is not in itself sufficient to demonstrate that the birds were using proportional navigation with $N = 1$ to implement their pursuits, it does at least show that the hypothesis that the birds use proportional navigation with $N = 1$ has the potential to explain most of the data. In summary, the Harris' hawks were using pure pursuit for most of the time, but occasionally made use of a deviated pursuit. In each case, the dynamics of the pursuit involve a time delay of approximately 0.19s.

Visual inspection of the raw trajectory data confirms that the bird's trajectory does indeed only begin to turn approximately 0.19s after the lure's trajectory turns. This presumably reflects not only the latency of the bird's physiological response to the lure, but also the finite time that is required physically to develop the unsteady aerodynamic force needed to initiate a turn. Ongoing analysis of the high-speed video data is now being conducted to measure the latency with which control inputs are first visible in the wings and tail.

Air-to-air attacks in Peregrine Falcons

Preliminary analysis of results

A total of 16 successful intercepts were achieved between the two birds (6 flights from "Lizzie"; 8 flights from "Izzy"). Fig. 21 plots the latitude and longitude measurements from the GPS units carried by the bird, aircraft, and the lure for all 16 flights. Although the GPS units are accurate to within <0.05m when

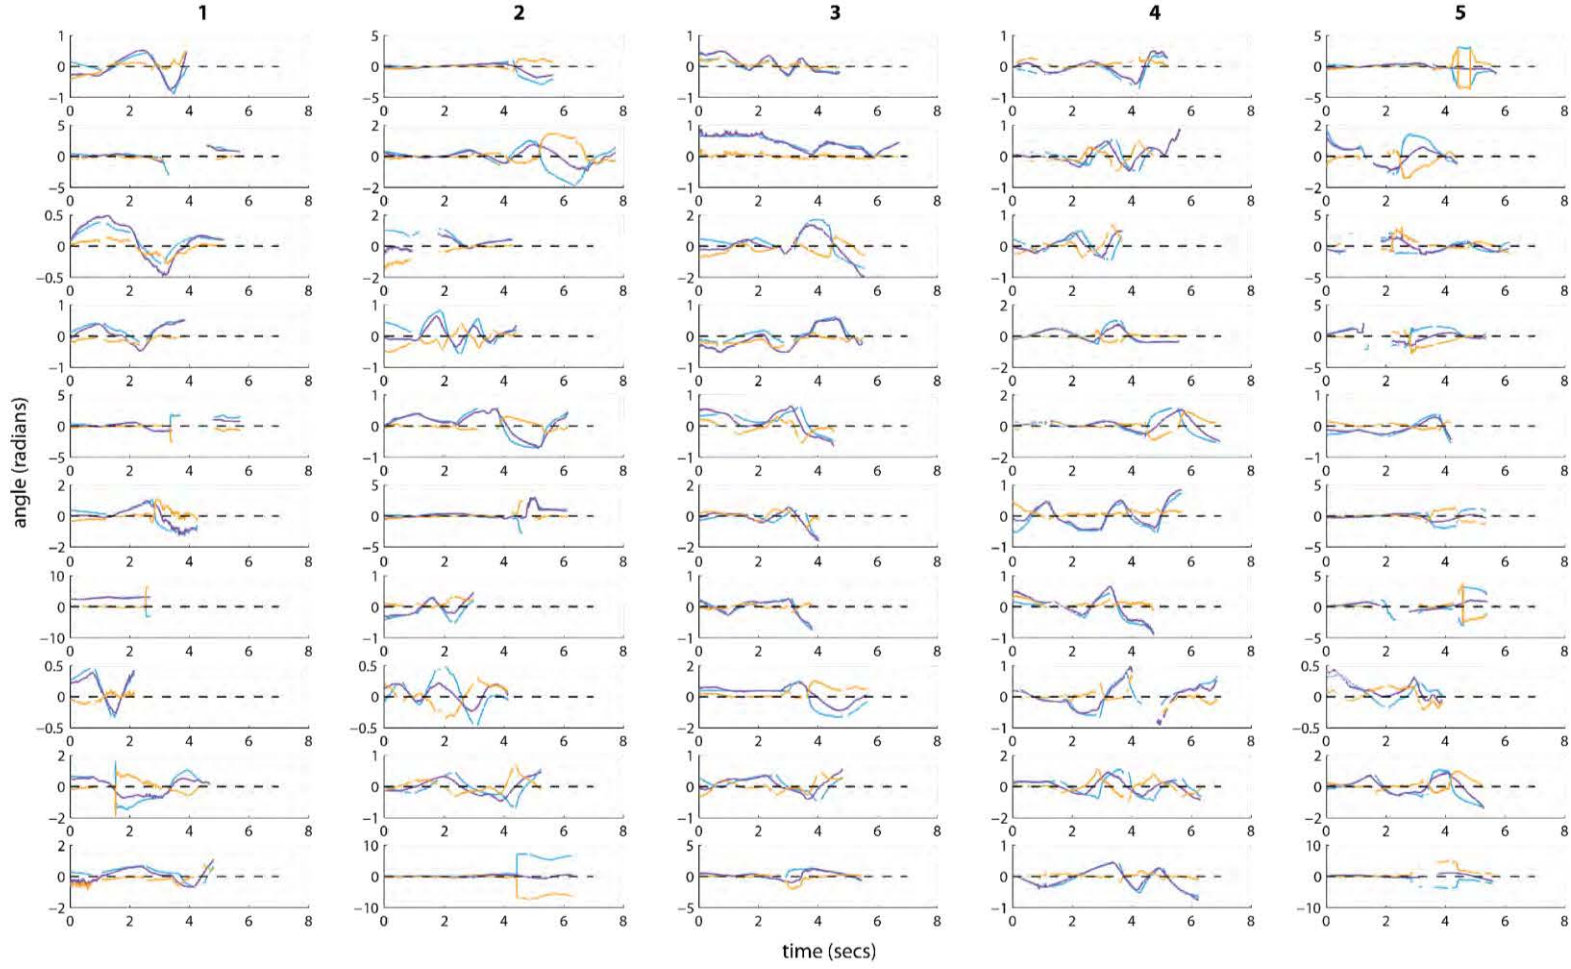


Figure 18. Lagged track angle $\gamma(t + 0.19)$ (blue), line-of-sight angle $\lambda(t)$ (purple), and lagged deviation angle $\delta(t + 0.19) = \lambda(t) - \gamma(t + 0.19)$ (orange) plotted against time for all of the air-to-ground attack passes by Harris' hawks that we analysed. Each column (1-5) corresponds to a different bird. The time lag of 0.19s is our best estimate of the average lag from the correlation coefficient analysis in Fig. 17.

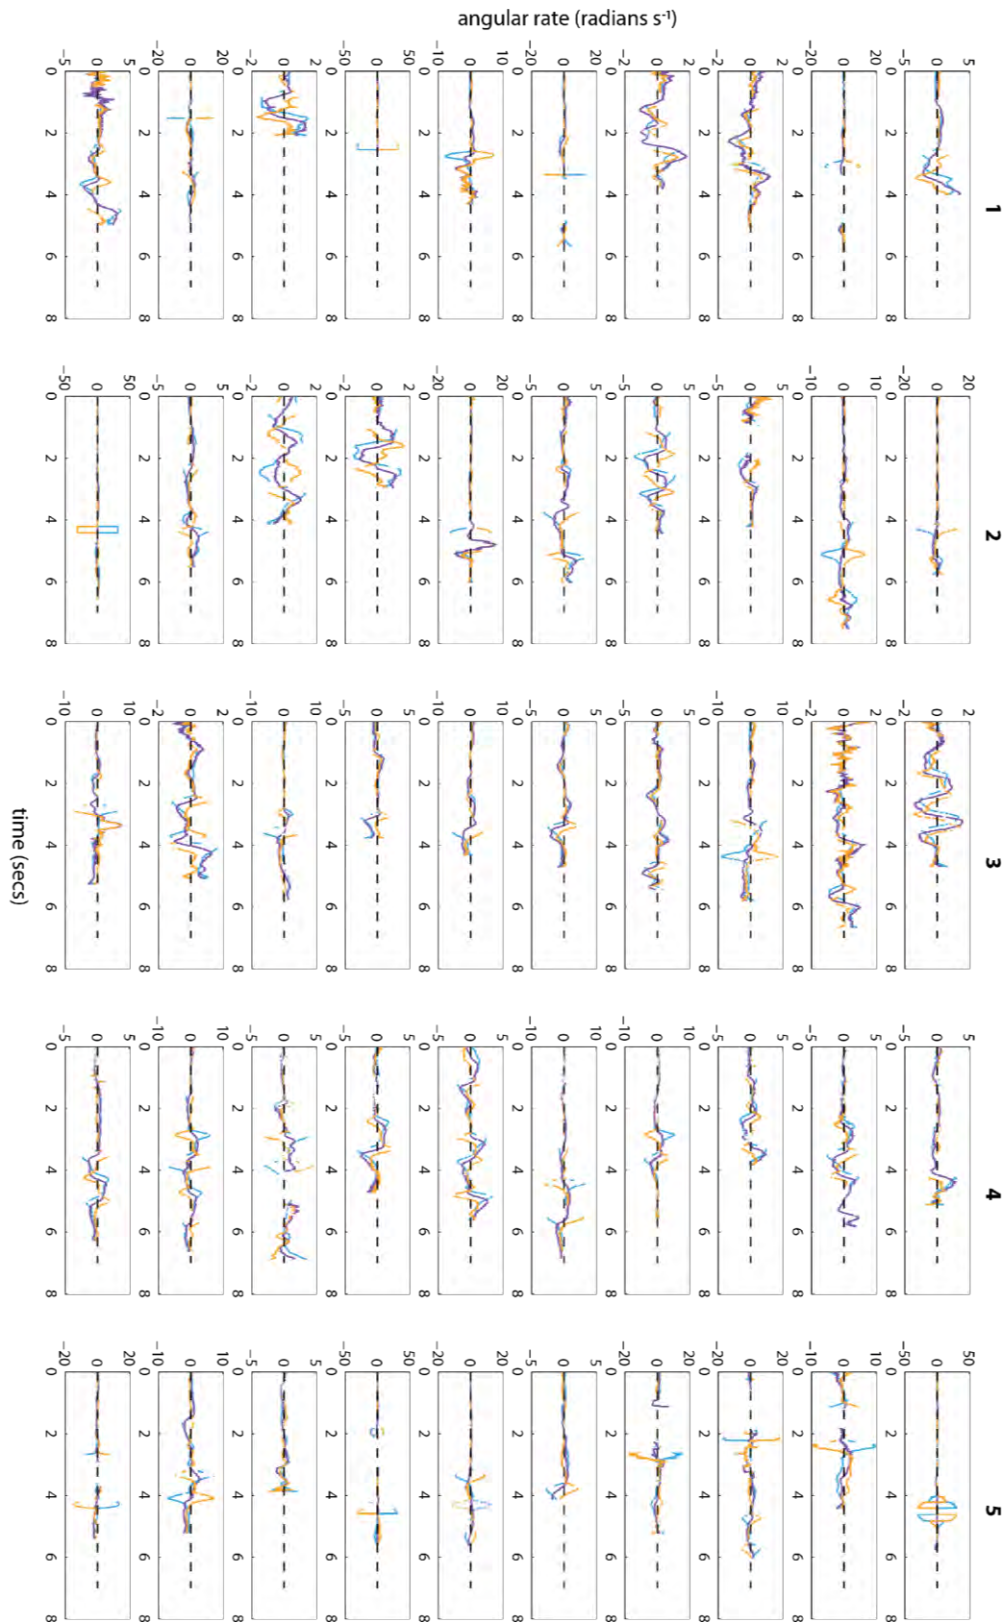


Figure 19. Turn rate $\dot{\gamma}$ (blue), line-of-sight rate $\dot{\lambda}$ (purple), and deviation rate $\dot{\delta} = \dot{\lambda} - \dot{\gamma}$ (orange) plotted against time for all of the air-to-ground attack passes by Harris' hawks that we analysed. Each column (1-5) corresponds to a different bird.

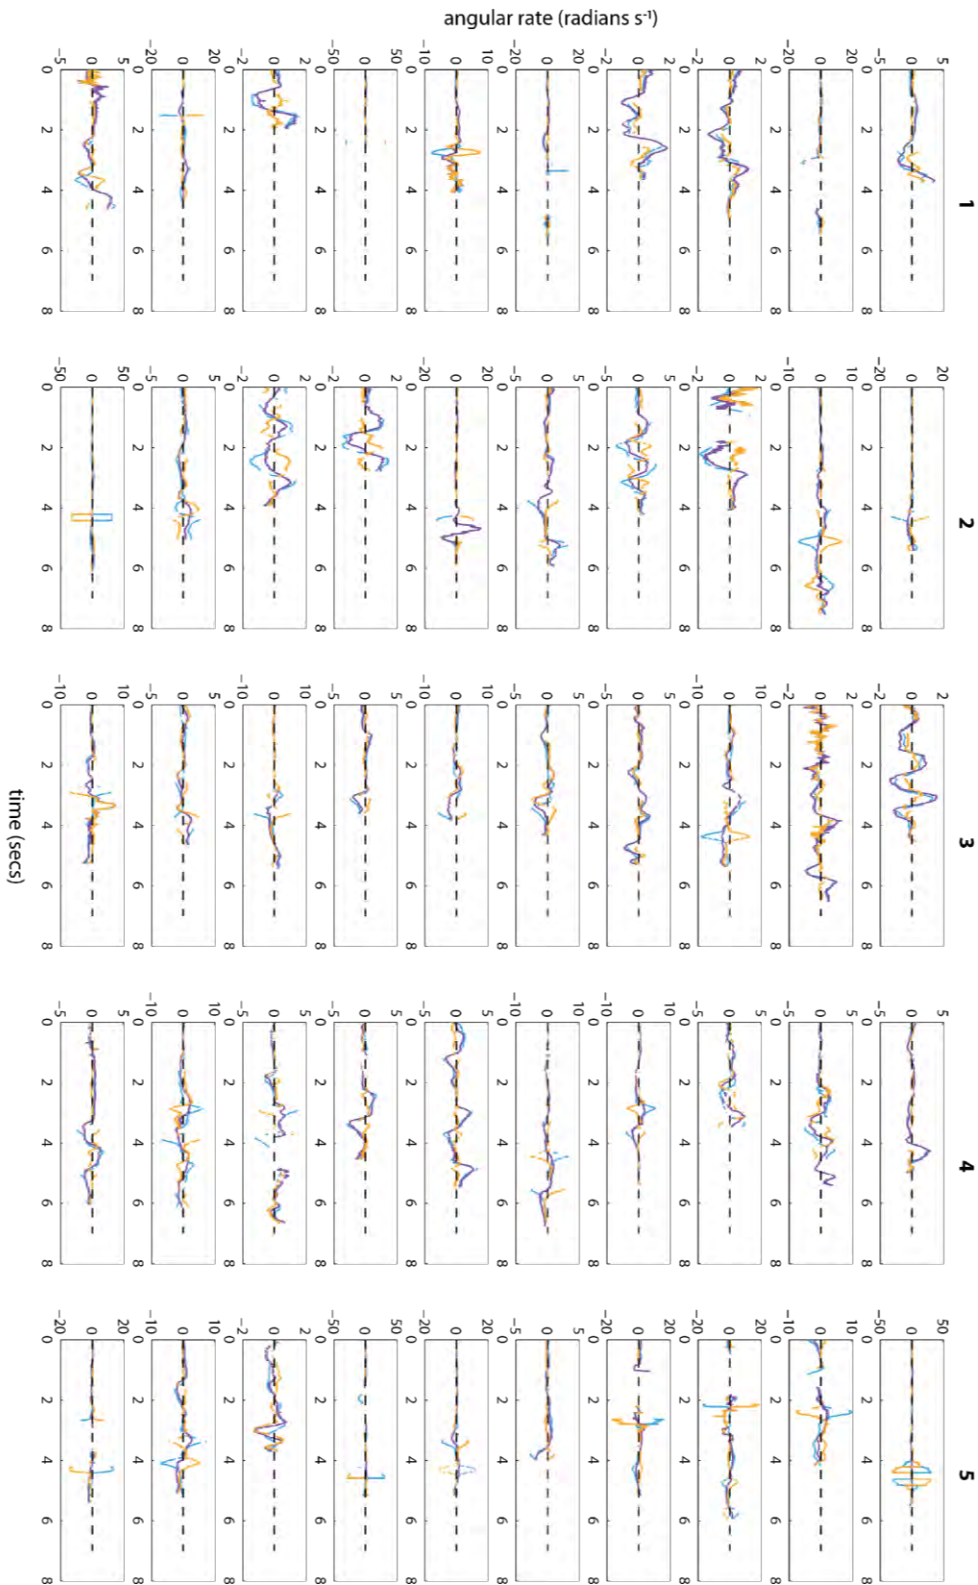


Figure 20. Lagged turn rate $\dot{\gamma}(t + 0.19)$ (blue), line-of-sight rate $\dot{\lambda}(t)$ (purple), and lagged deviation rate $\dot{\delta}(t + 0.19) = \dot{\lambda}(t) - \dot{\gamma}(t + 0.19)$ (orange) plotted against time for all of the air-to-ground attack passes by Harris' hawks that we analysed. Each column (1-5) corresponds to a different bird. The time lag of 0.19s is our best estimate of the average lag from the correlation coefficient analysis in Fig. 17.

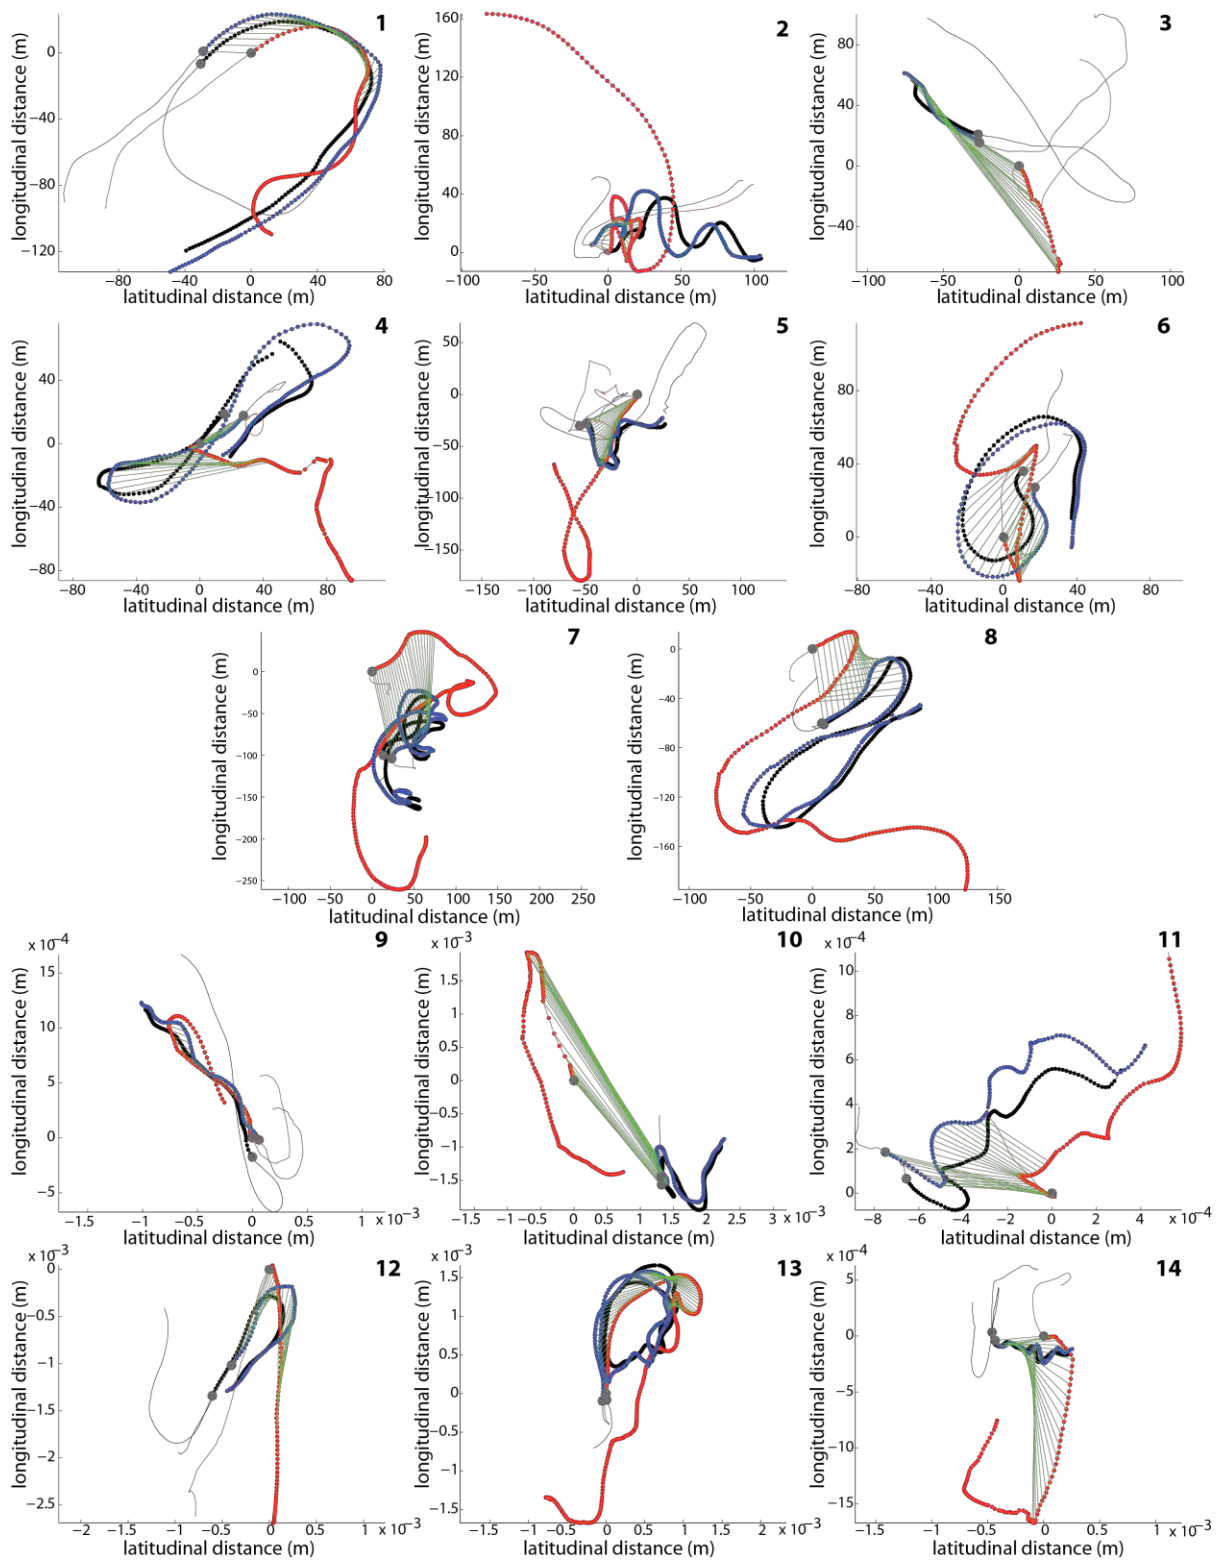


Figure 21: Raw GPS trajectory plots of air-to-air attacks by peregrines (red) against a lure (black) towed behind a remotely piloted air vehicle (blue). Flights 1-8: “Izzy”; Flights 9-14: Lizzie. Note that the relative position of these trajectories has not been adjusted to ensure that the bird and lure meet at the established point of intercept. Any deviation between the red and black trajectories at the point of intercept (grey circle) is therefore due to the absolute positional error in the GPS, but can be corrected at a later stage.

measuring changes in position over short time scales (see Methods above), their absolute positional accuracy is only expected to be on the order of 7.8m (Department of Defense Positioning, Navigation, and Timing Executive Committee, 2008). It is therefore necessary to adjust the relative position of the GPS tracks separately so that the trajectories of the bird and lure meet at the established point of intercept. Analysis of these trajectories is ongoing, in conjunction with analysis of the accompanying ground-based video data, which were collected so as to enable photogrammetric reconstruction if required. Ongoing analysis of the onboard video data will be used to identify head movements where possible, and to corroborate our findings from the onboard GPS units.

Conclusions

Our results on air-to-ground attacks on stationary targets by peregrine falcons and on air-to-ground attacks on maneuvering targets by Harris' hawks clearly demonstrate that both species turn at a rate approximately proportional to the rate of change in their line-of-sight to target. In the case of the Harris' hawks, the turn rate is typically approximately equal to the line-of-sight rate, albeit with a short time delay (*c.* 0.19s) in the dynamics. In the case of the peregrines, the constant of proportionality varies between, though not within, attack passes. It follows that both species effectively achieve proportional navigation, although it does not necessarily follow that either species actually implements this guidance law explicitly anywhere within their control system. In the particular case of the Harris' hawks, the effective navigation constant is $N = 1$, which corresponds to a pure or deviated pursuit course with constant deviation angle δ between the bird's velocity vector and its line of sight to target. It follows that guidance could in this case be implemented by acting directly upon the deviation angle itself, rather than acting upon the rate of change in the line of sight. In the case of the peregrines, the variation in navigation constant between passes makes it difficult to see alternative ways in which proportional navigation could emerge unless implemented directly by the birds' guidance system, but we must remain open to the possibility that such alternatives could exist. Nevertheless, the most parsimonious explanation of our data is that the birds do indeed implement pure proportional navigation directly within their guidance system, so we proceed on this assumption for the remainder of the discussion.

A necessary condition for implementing proportional navigation is that a non-rotating reference direction is established by the system, against which the line-of-sight rate can be measured. Birds go to great lengths to stabilize gaze against the visual scene (see e.g. Fig. 3B), using fast saccadic head movements to shift gaze between longer periods of gaze fixation (Warrick *et al.*, 2002; Eckmeier *et al.*, 2008). We hypothesise that a non-rotating reference direction is established visually in birds of prey, by stabilizing the head with respect to the landscape – probably with assistance from the vestibular system. In contrast, most guided missiles have a gimballed seeker that tracks the target continuously to measure its line-of-sight rate relative to an inertial reference established by gyros in the seeker or missile body. Thus, even if birds of prey and guided missiles both intercept targets using proportional navigation, they may very well

implement the same guidance law through fundamentally different mechanisms. This difference, if real, presumably reflects the fact that birds of prey have a much wider field of view than most guided missiles. In effect, this means that they have a wide-angle sensor array that is more comparable to the wide-angle sensor of a strapdown seeker than it is to the narrow-angle sensor of a gimbaled seeker, which they deploy on a rotating platform – the head – that is more comparable to a gimbaled image-stabilized camera mount than it is to a gimbaled seeker. Under this paradigm, the target is kept within the field of view, and perhaps more especially within the acute foveal regions of that field of view, by intermittent gaze shifts interspersed with longer periods of gaze stabilization. Our ongoing analysis of the data will focus upon the detailed role of head movements in such an implementation.

Acknowledgments

We thank Martin Cray for his assistance as falconer during all of the experiments, and Jo Binns for his assistance in husbandry of the birds. We thank all of the landowners who have permitted us to use their land for the purposes of this research. We thank Gregg Abate, Pat Bradshaw, Johnny Evers, Willard Larkin, and Ric Wehling for their support of this work.

References

- Collett T. S., and Land M. F.** (1978). How hoverflies compute interception courses. *J. Comp. Physiol.* 125, 191-204.
- Collett T. S., and Land M. F.** (1975). Visual control of flight behavior in the hoverfly, *Syritta pipiens* L. *J. Comp. Physiol.* 99, 1-66.
- Cresswell, W.** (1996). Surprise as a winter hunting strategy in Sparrowhawks, Peregrine Falcons, and Merlins. *Ibis* 138, 684-692.
- Dekker, D.** (1987). Peregrine Falcon predation on ducks in Alberta and British Columbia. *J. Wildl. Manag.* 51, 156-159.
- Dekker, D.** (1988). Peregrine Falcon and Merlin predation on small shorebirds and passerines in Alberta. *Can. J. Zool.* 66, 925-928.
- Dekker, D.** (1995). Prey capture by Peregrine Falcons wintering on southern Vancouver Island, British Columbia. *J. Raptor Res.* 29, 26-29.
- Dekker, D. and L. Bogaert.** (1997). Over-ocean hunting by Peregrine Falcons in British Columbia. *J. Raptor Res.* 31, 381-383.
- Department of Defense Positioning, Navigation, and Timing Executive Committee** (2008). *Global Positioning System Standard Positioning Service Performance Standard*. 4th edn. (Washington: Department of Defense Positioning, Navigation, and Timing Executive Committee).
- Eckmeier D., Geurten B. R. H., Kress, D., Mertes, M., Kern, R., Egelhaaf, M., and Bischof, H.-J.** (2008). Gaze strategy in the free flying zebra finch (*Taeniopygia guttata*). *PLoS ONE* 3, e3956.
- Frost, B. J., Wise, L. Z., Morgan. B. and Bird, D.** (1990). Retinotopic representation of the bifoveate eye of the kestrel (*Falco sparverius*) on the optic tectum. *Vis. Neurosci.* 5, 231-239.
- Ghose K., Horiuchi T. K., Krishnaprasad P. S., and Moss C. F.** (2006). Echolocating bats use a nearly time-optimal strategy to intercept prey. *PLoS Biol.* 4(5), e108.
- Gilbert C.** (1997). Visual control of cursorial prey pursuit by tiger beetles (Cicindelidae). *J. Comp. Physiol. A* 181, 217-230.
- Justh E. W., and Krishnaprasad P. S.** (2006). Steering laws for motion camouflage. *Proc. R. Soc. A* 462, 3629-3643.
- Kane, S. A. and Zamani, M.** (2014). Falcons pursue prey using visual motion cues: new perspectives from animal-borne cameras. *J. Exp. Biol.* 217, 225-234.
- Lanchester B. S., and Mark R. F.** (1975). Pursuit and prediction in the tracking of moving food by a teleost fish (*Acanthaluteres spilomelanurus*). *J. Exp. Biol.* 63, 627-645.
- Land M. F.** (1993). Chasing and pursuit in the dolichopodid fly *Poecilobothrus nobiletatus*. *J. Comp. Physiol.* 173, 605-613.
- Land M. F., and Collett T. S.** (1974). Chasing behavior of houseflies (*Fannia canicularis*). A description and analysis. *J. Comp. Physiol.* 89, 331-357.
- Mizutani A., Chahl J. S., and Srinivasan M. V.** (2003). Motion camouflage in dragonflies. *Nature* 423, 604.
- Olberg R. M., Worthington A. H., and Venator K. R.** (2000). Prey pursuit and interception in dragonflies. *J. Comp. Physiol. A.* 186, 155-162.
- Olberg R. M., Seaman R. C., Coats M. I., and Henry A. F.** (2007). Eye movements and target fixation during dragonfly prey-interception flights. *J. Comp. Physiol. A.* 193, 685-693.

- Olberg R. M.** (2012). Visual control of prey-capture flight in dragonflies. *Curr. Opin. Neurobiol.*, 22, 267–271.
- O'Rourke, C. T., Hall, M. I., Pitlik, T., and Fernández-Juricic, E.** (2010a) Hawk eyes I: Diurnal raptors differ in visual fields and degree of eye movement. *PLoS ONE* 5(9): e12802. doi:10.1371/journal.pone.0012802
- O'Rourke CT, Pitlik T, Hoover M, and Fernández-Juricic E** (2010b) Hawk eyes II: diurnal raptors differ in head movement strategies when scanning from perches. *PLoS ONE* 5(9). doi:10.1371/journal.pone.0012169.
- Regan D., and Gray R.** (2000). Visually guided collision avoidance and collision achievement. *Trends Cogn. Sci.* 4, 99-107.
- Shaffer D. M., Krauchunas S. M., Eddy M., and McBeath M. K.** (2004). How dogs navigate to catch Frisbees. *Psych. Sci.*, 15, 437-441.
- Shaffer D. M., and McBeath M. K.** (2002). Baseball outfielders maintain a linear optical trajectory when tracking uncatchable fly balls. *J. Exp. Psych.* 28, 335–348.
- Shneydor N. A.** (1998). *Missile Guidance and Pursuit: Kinematics, Dynamics and Control* (Cambridge: Woodhead).
- Siouris G. M.** (2004). *Missile Guidance and Control Systems* (New York: Springer-Verlag).
- Slonaker, J. R.** (1897). A comparative study of the area of acute vision in vertebrates. *J. Morphol.* 13, 445-502.
- Srinivasan M. V., and Davey M.** (1995). Strategies for active camouflage of motion. *Proc. R. Soc. B* 259, 19-25.
- Tucker V. A.** (2000). Gliding flight: drag and torque of a hawk and a falcon with straight and turned heads, and a lower value for the parasite drag coefficient. *J. Exp. Biol.* 203, 3733–3744.
- Tucker, V. A.** (2000). The deep fovea, sideways vision and spiral flight paths in raptors. *J. Exp. Biol.* 203, 3745–3754.
- Tucker V. A., Tucker A. E., Akers K., and Enderson J. H.** (2000). Curved flight paths and sideways vision in peregrine falcons (*Falco peregrinus*). *J. Exp. Biol.* 203, 3755–3763.
- Warrick D. R., Bundle M. W., and Dial K. P.** (2002). Bird maneuvering flight: blurred bodies, clear heads. *Integr. Comp. Biol.* 42, 141-148.
- Zhang S. W., Xiang W., and Srinivasan M. V.** (1990). Visual tracking of moving targets by freely flying honeybees. *Visual Neurosci.* 4, 379-386.

List of Symbols, Abbreviations, and Acronyms

γ	track angle <i>geographic direction of the flight velocity vector</i>
δ	deviation angle <i>angle between flight velocity vector and line of sight to target</i>
θ	foveal angle <i>angle between gaze of central fovea and sagittal plane</i>
λ	line-of-sight angle <i>geographic direction of the line of sight to target</i>
t	time
t_0	beginning of the time interval during which a bird both closes range on its target and turns monotonically towards it before interception
N	navigation constant of pure proportional navigation guidance law
GPS	Global Positioning System
LOS	line of sight