

PAPER

From biokinematics to a robotic active vision system

To cite this article: Ouriel Barzilay *et al* 2017 *Bioinspir. Biomim.* **12** 056004

View the [article online](#) for updates and enhancements.

Related content

- [A small-scale hyperacute compound eye featuring active eye tremor: application to visual stabilization, target tracking, and short-range odometry](#)
Fabien Colonnier, Augustin Manecy, Raphaël Juston *et al.*
- [A Kinect™ camera based navigation system for percutaneous abdominal puncture](#)
Deqiang Xiao, Huoling Luo, Fucang Jia *et al.*
- [Toward autonomous avian-inspired grasping for micro aerial vehicles](#)
Justin Thomas, Giuseppe Loianno, Joseph Polin *et al.*

Bioinspiration & Biomimetics



PAPER

From biokinematics to a robotic active vision system

RECEIVED
26 December 2016

REVISED
16 May 2017

ACCEPTED FOR PUBLICATION
5 June 2017

PUBLISHED
20 September 2017

Ouriel Barzilay¹, Lihi Zelnik-Manor², Yoram Gutfreund³, Hermann Wagner⁴ and Alon Wolf¹

¹ Faculty of Mechanical Engineering Technion, Israel Institute of Technology, Haifa, Israel

² Faculty of Electrical Engineering Technion, Israel Institute of Technology, Haifa, Israel

³ Faculty of Medicine, Technion, Israel Institute of Technology, Haifa, Israel

⁴ RWTH Aachen University, Aachen, Germany

E-mail: alonw@technion.ac.il (A Wolf)

Keywords: biokinematics, biorobotics, active vision, barn owls

Supplementary material for this article is available [online](#)

Abstract

Barn owls move their heads in very particular motions, compensating for the quasi-immovability of their eyes. These efficient predators often perform peering side-to-side head motions when scanning their surroundings and seeking prey. In this work, we use the head movements of barn owls as a model to bridge between biological active vision and machine vision. The biomotions are measured and used to actuate a specially built robot equipped with a depth camera for scanning. We hypothesize that the biomotions improve scan accuracy of static objects. Our experiments show that barn owl biomotion-based trajectories consistently improve scan accuracy when compared to intuitive scanning motions. This constitutes proof-of-concept evidence that the vision of robotic systems can be enhanced by bio-inspired viewpoint manipulation. Such biomimetic scanning systems can have many applications, e.g. manufacturing inspection or in autonomous robots.

Introduction

Active vision designates the controlled manipulation of visual sensors for enhancing visual acquisition and perception. Active vision mechanisms can be categorized into macro behaviors, associated with selective attention and visual search [1, 2], and micro behaviors, related to the analysis of a single object of interest, such as micro eye saccades or motion parallax [3]. This process is stereotypically achieved by eye, head and body movements [4–6]. Similarly, the vision of autonomous robots may be enhanced via active viewpoint manipulation strategies. In this study, we focused on the investigation of active vision micro behavior in the barn owl: the peering motion. The barn owl (*Tyto alba*) constitutes an excellent model system for exploring active vision mechanisms, since these birds have very limited eye movements [7, 8]; thus, gaze changes may be studied by only recording head motions [9]. Compensating for the restricted eye mobility, the long and flexible necks of barn owls allow them to perform large and elaborate head movements while focusing on objects of interest [10–12]. When introduced to new environments or before preying, barn owls typically exhibit conspicuous side-to-side head movements, called *peering* motions [13, 14]. Peering motions occur in various species [15–17], and

have been shown in some species to mediate distance estimation via motion parallax [13, 17–20].

Barn owls have frontally oriented eyes, similar to humans but in contrast to most other avian species. The resulting binocular overlap allows the birds to use binocular disparity for depth vision [21]. Barn owls have also been shown to exploit motion parallax and to be able to transfer depth information from stereo to parallax [13, 20]. We used these findings in a biomimetic approach to develop a barn-owl-inspired robotic platform equipped with a depth camera and aimed at the exploration of active vision strategies to test whether the inclusion of barn-owl like scanning strategies would improve scanning the performance of a robot.

This study aims at corroborating three research hypotheses, stating that peering motions (1) can be identified from kinematic analysis, (2) are linked with the enhancement of visual perception in barn owls, and (3) can improve the scan accuracy of static scenes in artificial robotic systems.

Materials and methods

Head tracking

Experiments were made with three adult barn owls (*Tyto alba*) of both sexes. All owls were hatched in captivity, and kept in large flying cages. Owls were hand-raised

from hatching and accustomed to human handling. All experiments were approved by the Technion Institutional Care and Use of Animals Committee and performed in adherence to the NIH Guide for the Care and Use of Laboratory Animals. Several weeks before the beginning of the recording sessions, owls were prepared for experiments in a single surgical procedure, in which a holding bolt was cemented to the dorsal part of the skull. The owls were allowed to recover for about a week after surgery and then acclimatized for perching on a small rod, 1.5 m above the ground, with the motion capture device attached to the head bolt (figure 1(A)). Experiments began once the owls showed no behavioral abnormalities carrying the head device. During the experiment, each owl was positioned on a perch at the center of the laboratory room ($6 \times 10 \times 3 \text{ m}^3$, with lights off and faint daylight coming through closed shutters). The owl was left alone in the room to spontaneously scan the surroundings (figure 1(A)).

The bird's head kinematics was captured by means of a Vicon MX-13 motion capture system [22] by reconstructing the locations of five markers situated on the head-mounted device, at a frequency of 120 Hz and with an accuracy of 0.2 mm. The device was screwed into the head bolt, so that it formed a rigid body with the owl's head. In addition to the motion capture markers, the device also supported a miniature wireless camera for an estimation of the bird's point of interest. The device's weight was minimized to 12 g (including markers, camera and battery pack) and the birds were accustomed to wearing the device on a daily basis. Previous studies indicated that similar head-mounted devices did not seem to significantly alter the owl behavior [23–25]. A remote controlled car with reflective markers attached (identified by the Vicon system) was activated at the beginning of the experiment for synchronization of the kinematic recordings and of the video footages. For the rest of the experiment, the scene was entirely static. From the recorded marker coordinates G in the laboratory coordinate system, the owl's head position C and orientation R , hence the gaze direction (up to limited eye movements), could be estimated from solution of the equation $G = RL + C$, where L are the known marker coordinates in the device's local coordinate system. After extraction of the side-to-side peering motions, each trajectory was translated and rotated such that its average local system coincided with the laboratory coordinate system, for easy replication on the robotic platform. While three non-aligned markers sufficed to determine the instantaneous location and orientation of the device, using five markers (figure 1(A)) allowed overcoming the occasional tracking failure of the motion capture system. The markers were placed on the device in an asymmetric manner, increasing the variance of the distances separating each pair of markers, so that the marker labeling could be validated and corrected after the experiments, and small gaps could be completed by interpolation. To compensate for reconstruction error and to filter unnecessary high frequencies, the location of each marker was smoothed over 20 time samples.

Extraction of peering motions

From the long recording sessions, several cues were extracted for identification and segmentation of the peering motions. First, the peering motions were easily recognized by an expert from direct observation of the barn owl's head on the videos. An initial selection of peering motions was performed from observation of the footages of the head-mounted camera and of an additional video camera directed at the bird (see supplementary movie 1, available at stacks.iop.org/BIB/12/056004/mmedia). Second, we used the criteria proposed by Ohayon *et al* [18] for the identification of peering motion and initially regarded peering motions as successions of peering samples, with angular velocity below 10° s^{-1} and linear velocity above 30 mm s^{-1} (figure 1(C)). However, in our view, the constraints used by Ohayon *et al* were too stringent, because they failed to identify many sequences that seemed to be peering motions. For example, in the sequence shown in figure 2(C) angular velocity is often higher than 10° s^{-1} , although the sequence looks like a peering motion upon inspection of the video. Thus, we concluded that a further cue was necessary for the segmentation and that, after using this cue, the threshold for the angular velocity could be relaxed to 20° s^{-1} .

The third cue used for isolation of peering motions was the change in gaze direction (figures 2(A) and (B)). This criterion allowed identifying successions of peering motions performed over different objects of interest, separated by quick rotations. This parameter was computed as the sinus of the angle (in absolute value) between two consecutive gaze direction vectors. Prominent peaks were observed every time the owl turned its head. At the same time, no peering motion was found entirely embedded in the peering quadrant (e.g. figure 2(C)). This criterion thus helped in checking the consistency of peering motions. In our hands, the combination of these three peering indications provided an accurate and exhaustive extraction of peering segments from the long recording sessions.

Recording sessions were conducted with four different barn owls while perching and freely observing the motion capture laboratory. A summary of the extraction of head capture and peering motions of three owls is presented in table 1.

The owls exhibited recurrent and consistent patterns of movement during the entire recording sessions as indicated by the observation of the translational and rotational components of the head velocity (figure 3). These movements correspond to what has already been described by Ohayon *et al* [18] and, therefore, are only described briefly for the sake of completeness here. Each owl spent more than half the time (53–63% of the samples) resting or fixating. Mixed rotations and translations constituted 21–25% of the samples. This demonstrated a marked proportional relationship between the rotational and translational components. As can be seen in figure 3, the highest peaks in velocity were reached only in motions simultaneously combining great angular and translational velocity components,

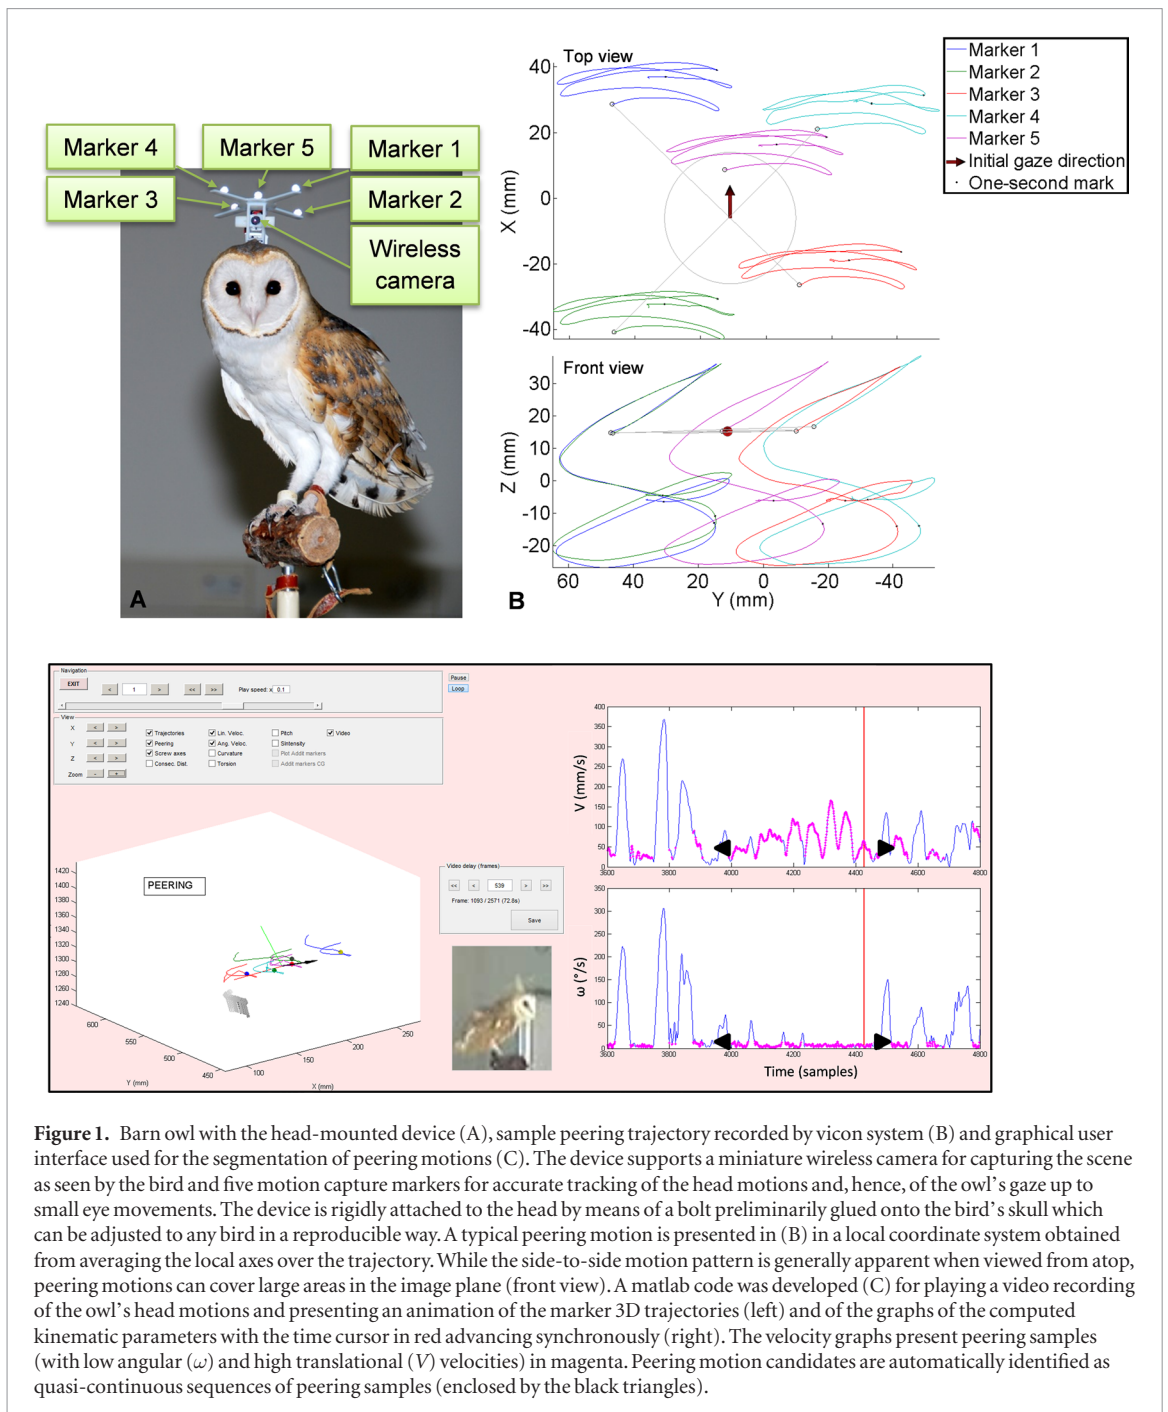


Figure 1. Barn owl with the head-mounted device (A), sample peering trajectory recorded by vicon system (B) and graphical user interface used for the segmentation of peering motions (C). The device supports a miniature wireless camera for capturing the scene as seen by the bird and five motion capture markers for accurate tracking of the head motions and, hence, of the owl's gaze up to small eye movements. The device is rigidly attached to the head by means of a bolt preliminarily glued onto the bird's skull which can be adjusted to any bird in a reproducible way. A typical peering motion is presented in (B) in a local coordinate system obtained from averaging the local axes over the trajectory. While the side-to-side motion pattern is generally apparent when viewed from atop, peering motions can cover large areas in the image plane (front view). A matlab code was developed (C) for playing a video recording of the owl's head motions and presenting an animation of the marker 3D trajectories (left) and of the graphs of the computed kinematic parameters with the time cursor in red advancing synchronously (right). The velocity graphs present peering samples (with low angular (ω) and high translational (V) velocities) in magenta. Peering motion candidates are automatically identified as quasi-continuous sequences of peering samples (enclosed by the black triangles).

in line with the composite nature of the ballistic head saccades in barn owls [18]. Pure rotations were rarely observed (4–5% of the samples). In contrast, in 12–18% of the cases, the owls executed high translational speeds with low rotational components. These segments corresponded to peering movements. The peering movements often consisted not only of a single side-to-side movement, but of a sequence of such movements with often complex trajectories (see figure 1(B)). Moreover, as already described by Ohayon *et al* [18], translations did not only consist of horizontal side-to-side movements, but had additional vertical translational components (see example in figure 1(B)).

Active vision robot

The active vision robotic platform (figure 4 and supplementary movie 2 (available at stacks.iop.org/BB/12/056004/mmedia)) was designed to include five degrees of freedom. The three translational degrees allowed the robot to move in a $25 \times 25 \times 25 \text{ cm}^3$ workspace. Although the largest observed peering motions (figure 5) were bounded in a 10 cm-edged cube, the robot's extent was defined wider so as to allow additional types of active vision motions. The next two links in the robot's serial kinematic chain enabled yaw and pitch rotations, so that the depth camera was always maintained oriented towards the

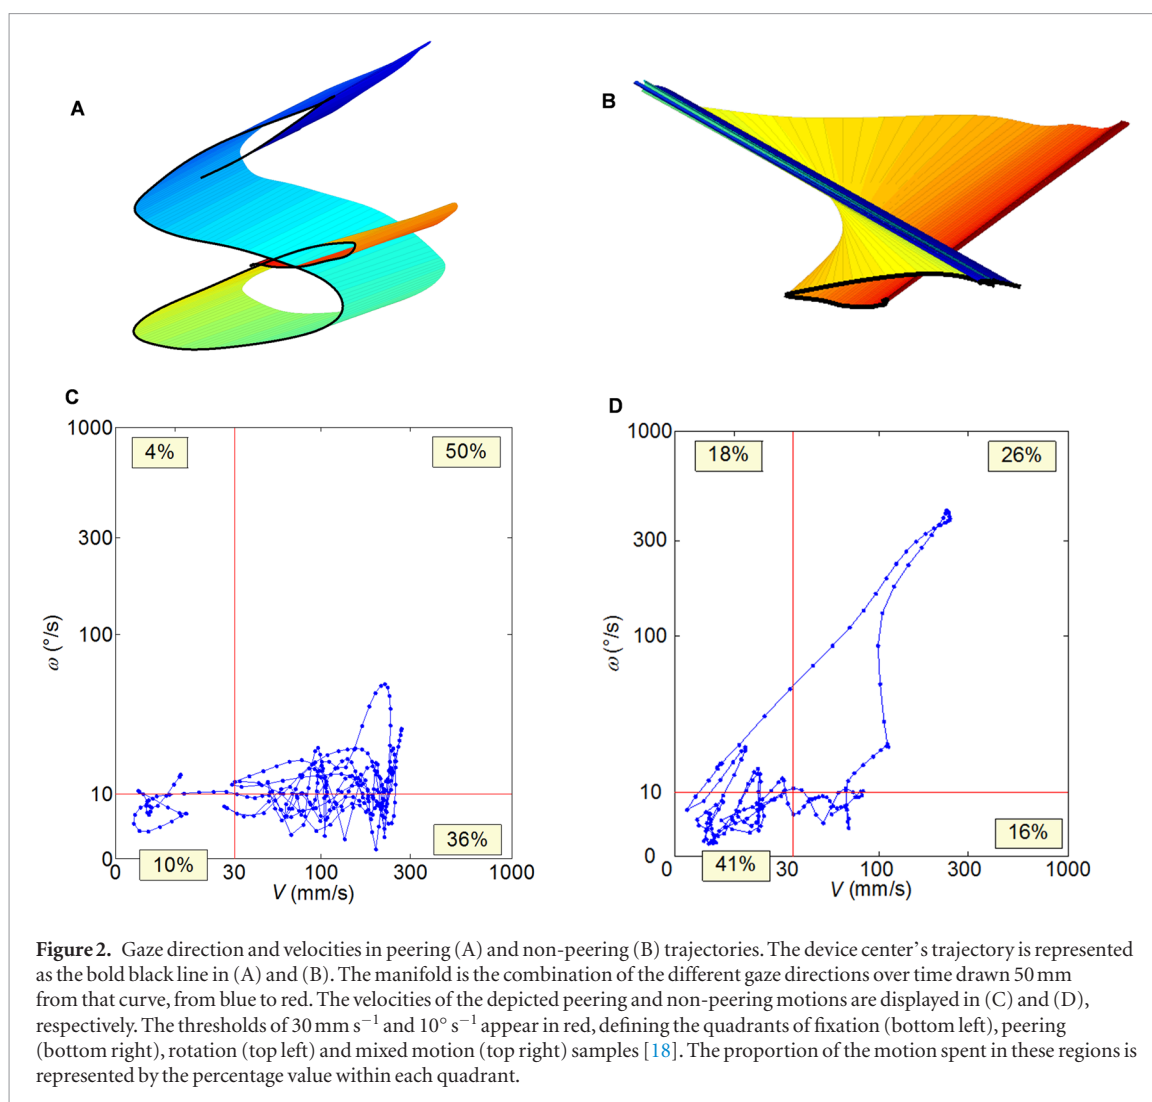


Table 1. Extracted peering motions—summary.

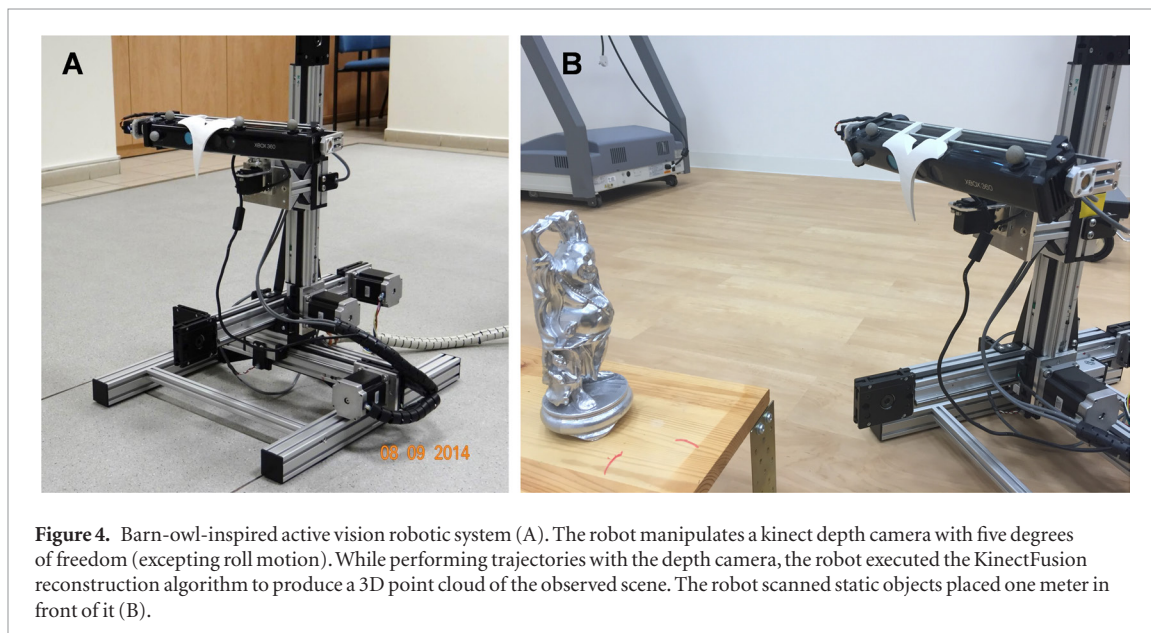
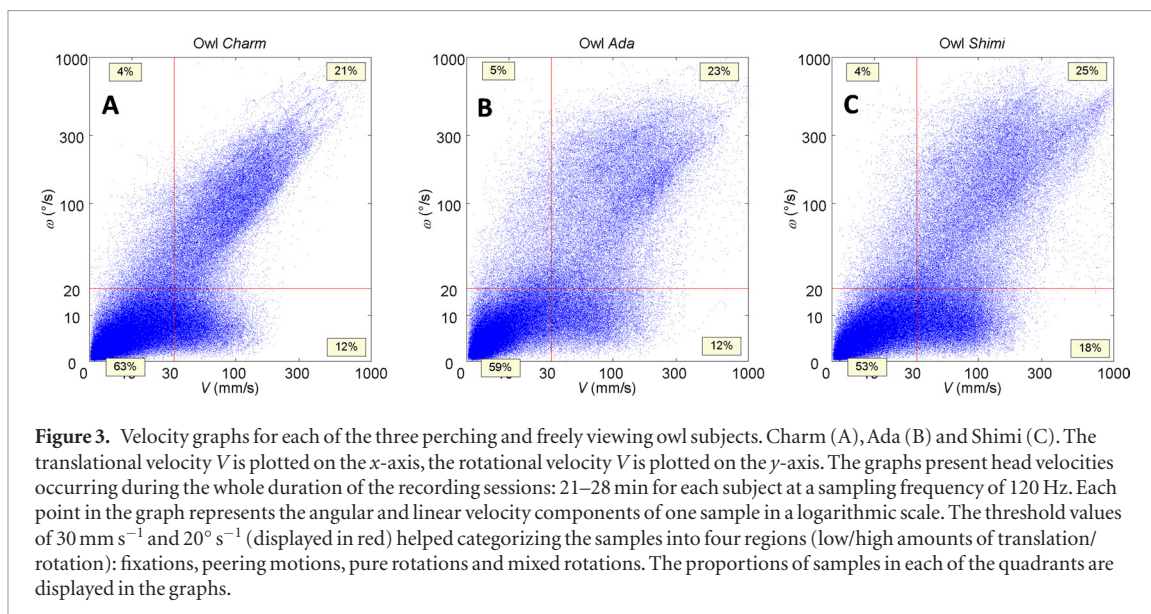
Owl	Recording duration (min)	Number of peering motions	Total peering duration (min:s)	Average peering duration (s)	Average peering distance (mm)	Peering motions replicated on the robot
Ada	24	20	0:30	1.51 ± 0.85	81.3 ± 50.9	2
Shimi	26.5	69	1:28	1.27 ± 0.85	65.2 ± 60.9	4
Charm	28.5	118	2:28	1.25 ± 0.84	55.9 ± 58.1	13
Overall	79	207	4:26	1.28 ± 0.85	61.5 ± 58.6	19

target object. Although the Kinect has a motorized tilt angle of $\pm 27^\circ$, this degree of freedom was replaced for centralized and customized control. Rotations around the roll axis were considered unlikely to improve the quality of reconstructions. The camera holding support was designed to maintain the camera horizontally level at all times. The precision was set as 0.25 mm for the linear motion and 2° for the angular motion.

An Arduino MEGA 2560 microcontroller transmitted the commands received from the workstation by USB serial connection to the servo motors via a similar Arduino board by inter-integrated circuit (I2C) communication, in a master/slave configuration, allowing a task-oriented partition of the microcontroller boards.

The master board also managed the commands to the stepper motors via motor drivers, and the incoming signals from the limit switches enabling the homing process and ensuring legitimate motion of the robot in the delimited workspace, and from the user interface, such as the stop button.

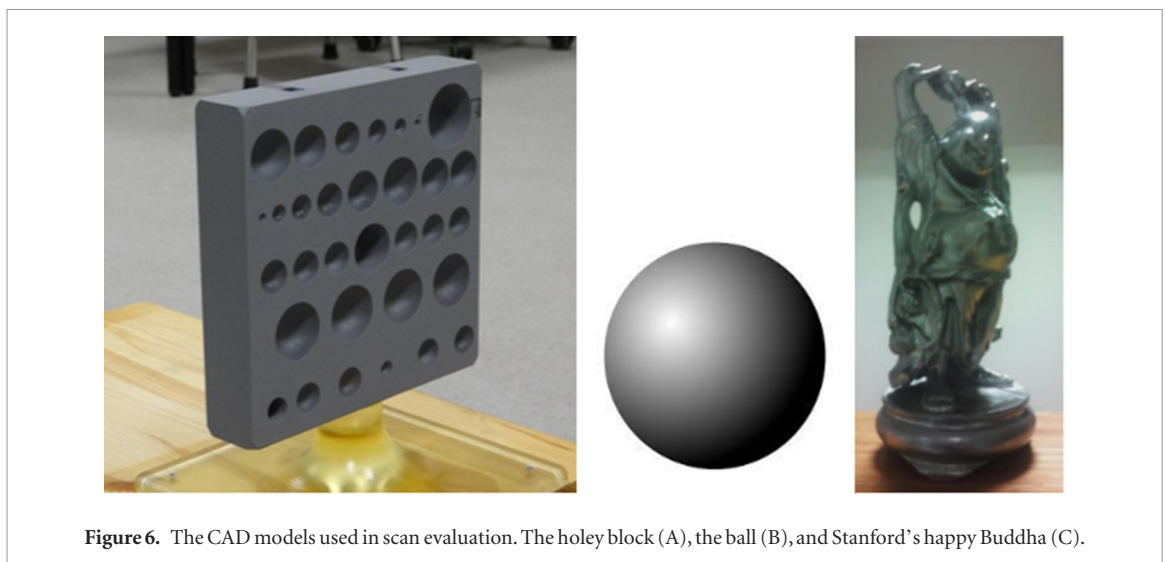
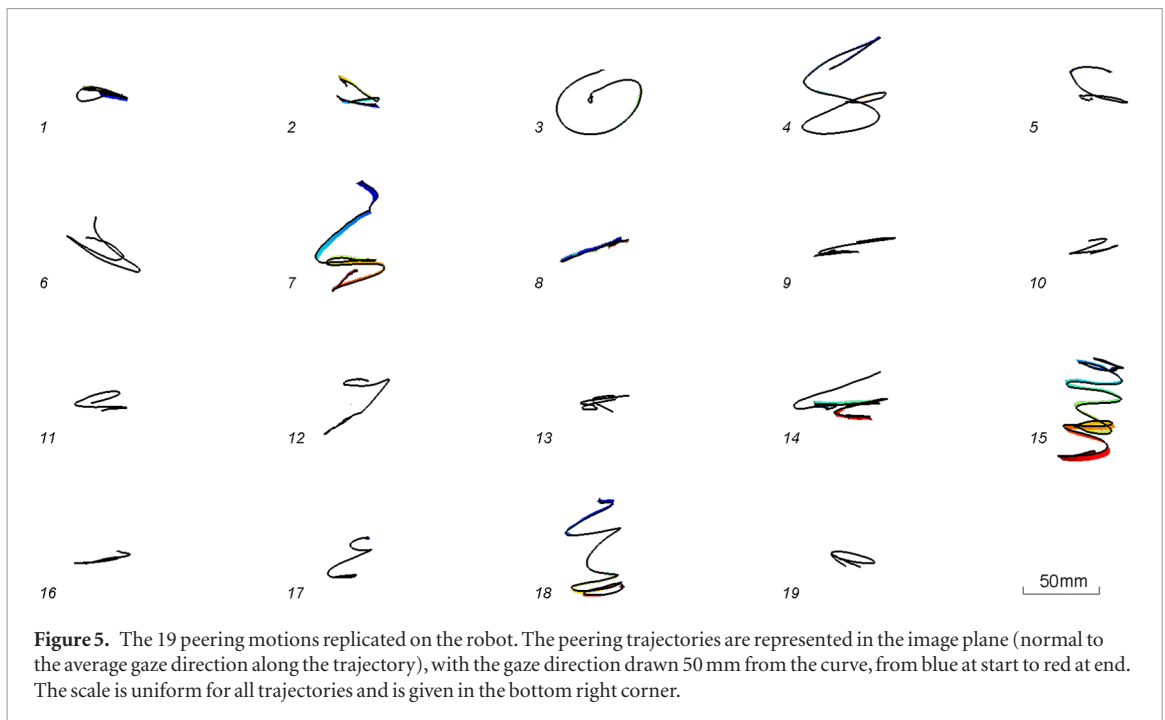
At initialization, the trajectory extracted in Matlab was loaded as an ASCII list of spatial points. After some level of preprocessing aimed at formatting the data for the robot (downsampling, conversion from millimeters to microsteps, speed computation, etc), the trajectory was sent as a whole to the microcontroller once the robot signaled completion of the homing procedure. The computer agent then waited for the robot to get into position to start the main loop.



As the robot started performing the transmitted trajectory, the 3D reconstruction procedure was initiated. No communication was required at that stage between the microcontroller and the PC, until completion of the trajectory. For that reason, a multi-threaded approach was adopted on the PC software agent. The first thread was responsible for regularly fetching images from the Kinect and updating the reconstruction with *KinectFusion* [26, 27]. The second thread continuously checked the communication stream in search of the trajectory completion signal. Once this signal was perceived on the PC, the communication thread terminated the reconstruction thread and continued with output processing. The PC software agent was implemented in C++, with use of the Point Cloud Library providing *KinFu*: an open-source GPU implementation of KinectFusion in CUDA. The code was executed on a Toshiba Qosmio laptop (i7-3630QM @2.40 GHz, 8Gb DDR3 RAM, NVIDIA GeForce GTX 670 M), at a rate of 8fps.

Scan accuracy evaluation

Different objects were placed at a distance of one meter in front of the robot and several trajectories were executed for each of the objects. The models were scaled to reach a height of 30 cm. At that distance, the Kinect for Xbox 360, with an angular field of view of 57° horizontally and 43° vertically, was able to see the models entirely at all times. The first model was specifically designed to test the accuracy of the selected reconstruction method KinectFusion, and included holes of diameters ranging from 0.5 mm to 4 cm (figure 6(A)). Both front and back faces presented different semi-spherical cavities and conical holes to provide a rich specimen of geometric primitives on which the robot's scan accuracy could be checked. In the produced point clouds, holes with diameters smaller than 2 cm were hardly observable. A spherical model (figure 6(B)) was also produced by rapid prototyping to produce scan results on a uniform model. The

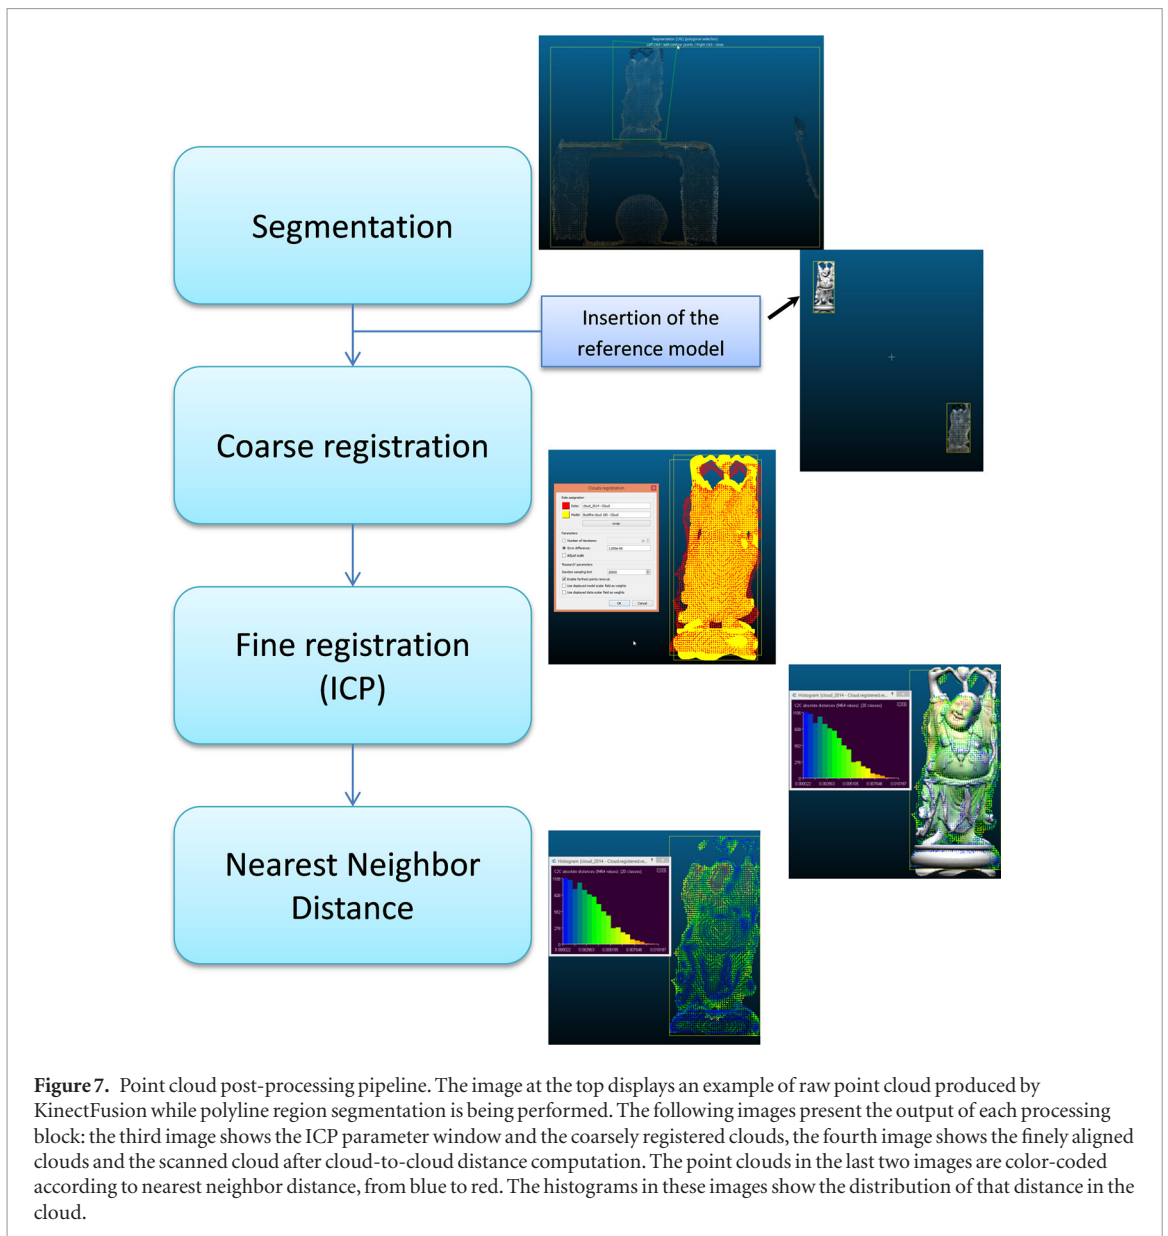


principal model on which the final accuracy of the robot was tested was the Happy Buddha statuette from the Stanford model repository (figure 6(C)). This model is very popular in the fields of computational geometry and computer graphics, as it is composed of free-form surfaces with a wide range of levels of details. After scaling it to a 30 cm height, the model was printed by rapid prototyping and painted. The paint coating greatly improved scan accuracy as the rapid prototyping material used was translucent and scattered the infrared light emitted by the Kinect camera, making the object virtually invisible to the camera. This property of the material was made useful by leaving the pedestal supporting the models unpainted, rendering the scanned target model easier to segment from the rest of the scenery (figure 6(A)).

For execution on the robot, the trajectories recorded in the Vicon motion capture system at 120 Hz needed to be downsampled to a rate more suitable to the hard-

ware limitations but with minimal loss of information. Each dozen samples were averaged to reach a rate of 10 Hz, providing an acceptable trade-off between hardware capabilities and trajectory accuracy. Additionally, the trajectories were stretched to durations of 5.5–6.0s to allow the capture of 26 depth frames exactly, allowing KinectFusion to produce a dense point cloud. As KinectFusion relies on the iterative closest point (ICP) method for fitting the data arriving from the last frame to the previously scanned model and for estimating the camera pose, tracking occasionally failed during scans. Only scans where ICP had not failed once were considered.

After scanning, the produced point clouds required post-processing for isolating the object of interest from the scan of the scene and comparing it to the reference 3D model. The processing of the point clouds obtained from the different scans was performed by means of CloudCompare, a software package allowing



the segmentation of an object and its alignment and comparison with another reference model.

The point cloud post-processing workflow is depicted in figure 7. As the produced point cloud represents a 3D model of the entire scene, the object of interest first needed to be segmented out. CloudCompare allows visualization and editing of point clouds, and particularly extracting the model of interest by selecting polyline regions on different views and discarding irrelevant points. After insertion of the 3D model used to print the scanned object, the two models were aligned one to the other. This was done in two steps, a coarse registration followed by a fine registration. The fine registration is based on ICP, a state-of-the-art alignment method known to strongly rely on the preliminary coarse alignment. The coarse registration was performed by aligning the principal axes and then by making the centroids of the clouds' bounding boxes coincide. ICP registration was performed several times thereafter for improved alignment accuracy, until the error stabilized.

Once the models were finely aligned, the produced point cloud could be evaluated. The reference model was sampled to produce a dense point cloud with one million points. The evaluation of the scanning trajectory was then performed by computation of the cloud-to-cloud distance between the output cloud and the model. The cloud-to-cloud distance was computed based on the evaluation, for each point of the scanned cloud, of the distance to its nearest neighbor in the reference model (figure 7).

Results

Although scans obtained from the peering trajectories were globally more accurate by 4.11% (from 1.85 ± 0.13 mm for peering motions to 1.925 ± 0.03 mm for basic trajectories), statistical significance was not achieved for the Ball model scans ($p = 0.12$). However, if two outliers among the 19 peering motions were discarded (trajectories #6 and #7, yielding an average error of 1.82 ± 0.11 mm for peering motions), the

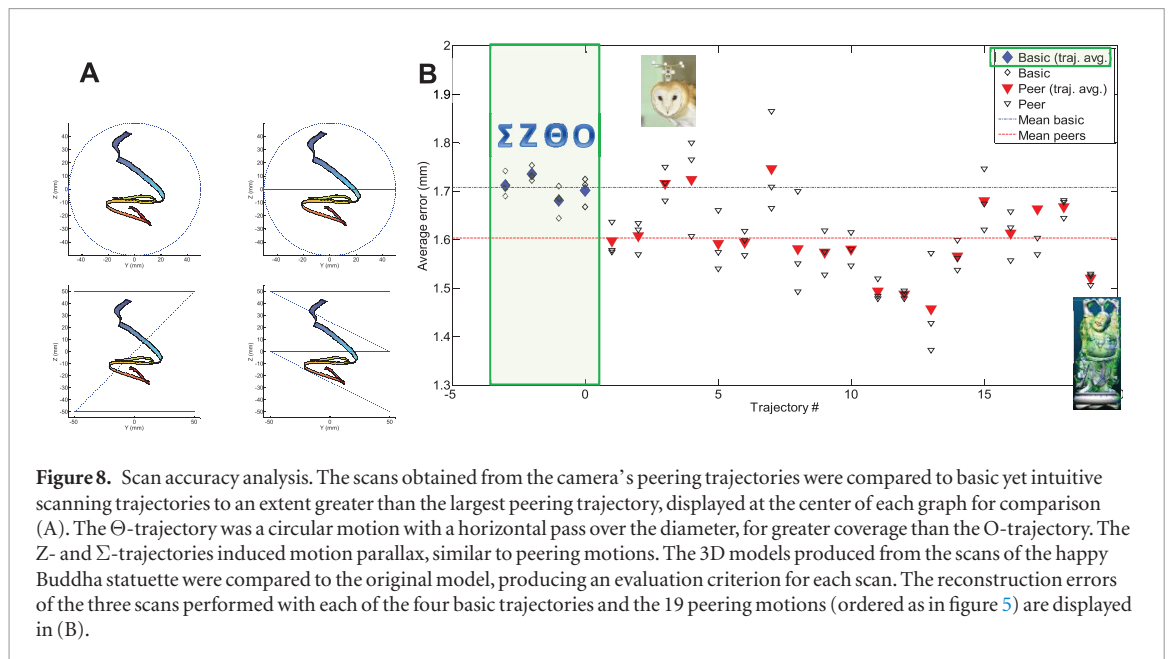


Figure 8. Scan accuracy analysis. The scans obtained from the camera's peering trajectories were compared to basic yet intuitive scanning trajectories to an extent greater than the largest peering trajectory, displayed at the center of each graph for comparison (A). The Θ -trajectory was a circular motion with a horizontal pass over the diameter, for greater coverage than the O-trajectory. The Z- and Σ -trajectories induced motion parallax, similar to peering motions. The 3D models produced from the scans of the happy Buddha statuette were compared to the original model, producing an evaluation criterion for each scan. The reconstruction errors of the three scans performed with each of the four basic trajectories and the 19 peering motions (ordered as in figure 5) are displayed in (B).

accuracy improvement of barn owls peering-inspired motions over the control group increased to 5.25% ($p = 0.04$). The omitted trajectories were the ones presenting the highest discrepancies over the three scan repetitions, in addition to relatively high deviations from the reference ball model. In a general manner, the scan accuracy achieved in this study with the Kinect for Xbox 360 and KinectFusion compares fairly with the results found in the literature for similar scenes [28, 29]. Nonetheless, our focus was set on accurately reproducing the same methods for the biomimetic and basic trajectories for the sake of comparison.

The results of the scan accuracy analysis on the Ball model were less conclusive than the results on the Happy Buddha model. For the Happy Buddha model, the average scan accuracy of the intuitive scanning trajectories ($n = 4$, three repetitions each, see blue curves in figure 8(A)) was 1.71 ± 0.03 mm (mean \pm standard deviation). The peering motions ($n = 19$, three repetitions each) presented an average accuracy of 1.60 ± 0.10 mm (figure 8). In other words, moving the camera along peering motions recorded from barn owls produced an average improvement in scan accuracy of 6.1% compared to the control group of intuitive scanning trajectories. A Wilcoxon rank-sum test [30] was performed to evaluate the statistical significance of the difference in scan accuracy observed between peering motions and reference motions. The improvement of scan accuracy from execution of peering trajectories with the camera was statistically significant ($p = 0.0001$).

Discussions

While the peering motions provided greater accuracy, they generally provided less coverage than the reference

motions. Scans with motions in the control group provided, after segmentation, point clouds including 8304 ± 243 points while peering scans produced clouds of 7778 ± 379 points. It is reasonable to assume that the peering motions could not improve both scan accuracy and coverage of the object of interest, especially when the motions of the control group had extents larger than the extent of the largest peering motion.

The peering trajectories were 3D, while the reference motions were planar trajectories. Peering motions were generally restricted to an extent of 4 cm within the image plane, but often included motion along the gaze direction with an average extent of 1.5 cm, (and 2.4 cm at most, but the gaze direction was never the principal axis of motion). Although motions toward the line of sight are not believed to improve depth estimation, their influence on scan accuracy could be assessed by scanning with peering motions projected onto the average image plane.

Conclusions

The results obtained in this study support the evidence that peering motions can improve depth estimation in artificial systems, not only at depth discontinuities but also in the observation of fine details. Similarly, it is reasonable to assume that, with these slight head movements, barn owls achieve a general finer scan resolution, rather than solely observe depth anomalies between the prey's contour and the background.

In this study, we showed that the peering-inspired manipulation of a depth camera within a restricted space could improve scan accuracy on static objects. Additionally to supporting the evidence that peering motions enrich the scanning capabilities of barn owls, this finding suggests that manipulation of the visual sensor of a robot with miniature peering-inspired

trajectories, in combination with models of selective attention [31] and navigation capabilities, may enhance the vision of autonomous robots.

Acknowledgments

We thank the Israel Science Foundation (ISF) (grants 1274/11 and 1086/16), the German Research Foundation (DFG) (grant WA 606/17-1) and the Ollendorf Foundation for their financial support; Dr S Netser and Y Hazan for their help with the owl experiments; A Bar-Yehuda and M Souffir for the construction of the robot.

References

- [1] Treisman A 1985 Preattentive processing in vision *Comput. Vis. Grap. Image Process.* **31** 156–77
- [2] Wolfe J M 1994 Visual search in continuous, naturalistic stimuli *Vis. Res.* **34** 1187–95
- [3] Engbert R and Kliegl R 2003 Microsaccades uncover the orientation of covert attention *Vis. Res.* **43** 1035–45
- [4] Findlay J 1998 Active vision: visual activity in everyday life *Curr. Biol.* **8** R640–2
- [5] Mokeichev A, Segev R and Ben-Shahar O 2010 Orientation saliency without visual cortex and target selection in archer fish *Natl Acad. Sci.* **107** 16726–31
- [6] Lehrer M and Srinivasan M V 1994 Active vision in honeybees: task-oriented suppression of an innate behaviour *Vis. Res.* **34** 511–6
- [7] Steinbach M J and Money K E 1973 Eye movements of the owl *Vis. Res.* **13** 889–91
- [8] du Lac S and Knudsen E I 1990 Neural maps of head movement vector and speed in the optic tectum of the barn owl *J. Neurophysiol.* **63** 131–46
- [9] Harmening W M, Orłowski J, Ben-Shahar O and Wagner H 2011 Overt attention toward oriented objects in free-viewing barn owls *Natl Acad. Sci.* **108** 8461–6
- [10] Krings M, Nyakatura J A, Fischer M S and Wagner H 2014 The cervical spine of the American barn owl (*Tyto furcata pratincola*): I. Anatomy of the vertebrae and regionalization in their S-shaped arrangement *PLoS One* **9** e91653
- [11] Boumans M L, Krings M and Wagner H 2015 Muscular arrangement and muscle attachment sites in the cervical region of the American barn owl (*Tyto furcata pratincola*) *PLoS One* **10** e0134272
- [12] de Kok-Mercado F, Habib M, Phelps T, Gregg L and Gailloud P 2013 Adaptations of the owl's cervical & cephalic arteries in relation to extreme neck rotation *Science* **339** 514
- [13] van der Willigen R F, Frost B J and Wagner H 2002 Depth generalization from stereo to motion parallax in the owl *J. Comp. Physiol. A* **187** 997–1007
- [14] Wagner H 1989 Peering in barn owls *Neural Mechanisms of Behavior* ed J Erber et al (New York: Thieme Verlag) pp 238–9
- [15] Bruckstein A, Holt R J, Katsman I and Rivlin E 2005 Head movements for depth perception: praying mantis versus pigeon *Auton. Robots* **18** 21–42
- [16] Wallace H R 1959 The movement of eelworms in water films *Ann. Appl. Biol.* **47** 366–70
- [17] Collett T S 1978 Short communication: peering—a locust behaviour pattern for obtaining motion parallax information *J. Exp. Biol.* **76** 237–41
- [18] Ohayon S, van der Willigen R F, Wagner H, Katsman I and Rivlin E 2006 On the barn owl's visual pre-attack behavior: I. Structure of head movements and motion patterns *J. Comp. Physiol. A* **192** 927–40
- [19] Fux M and Eilam D 2009 How barn owls (*Tyto alba*) visually follow moving voles (*Microtus socialis*) before attacking them *Physiol. Behav.* **98** 359–66
- [20] van der Willigen R F 2011 Owls see in stereo much like humans do *J. Vis.* **11** 10
- [21] van der Willigen R F, Frost B J and Wagner H 1998 Stereoscopic depth perception in the owl *Neuroreport* **9** 1233–7
- [22] Peak V 2005 *Vicon Motion Capture System* www.vicon.com/
- [23] Hazan Y, Kra Y, Yarín I, Wagner H and Gutfreund Y 2015 Visual-auditory integration for visual search: a behavioral study in barn owls *Front. Integr. Neurosci.* **9** 11
- [24] Ohayon S, Harmening W, Wagner H and Rivlin E 2008 Through a barn owl's eyes: interactions between scene content and visual attention *Biol. Cybern.* **98** 115–32
- [25] Barzilay O, Gutfreund Y and Wolf A 2012 Biokinematic study of barn owl head movements for the development of a bio-inspired active vision robotic system *Latest Advances in Robot Kinematics* (Dordrecht: Springer) pp 139–46
- [26] Newcombe RA, Izadi S, Hilliges O, Molyneaux D, Kim D, Davison A J, Kohli P, Shotton J, Hodges S and Fitzgibbon A 2011 KinectFusion: real-time dense surface mapping and tracking *2011 10th IEEE Int. Symp. on Mixed and Augmented Reality (ISMAR)* (IEEE) pp 127–36
- [27] Izadi S, Kim D, Hilliges O, Molyneaux D, Newcombe RA, Kohli P, Shotton J, Hodges S, Davison A J and Fitzgibbon A 2011 KinectFusion: real-time 3D reconstruction and interaction using a moving depth camera *Proc. of the 24th Annual ACM Symp. on User Interface Software and Technology (ACM)* pp 559–68
- [28] Khoshelham K and Elberink S O 2012 Accuracy and resolution of kinect depth data for indoor mapping applications *Sensors* **12** 1437–54
- [29] Meister S, Izadi S, Kohli P, Hämmerle M, Rother C and Kondermann D 2012 When can we use kinectfusion for ground truth acquisition *Proc. Workshop on Color-Depth Camera Fusion in Robotics* vol 2
- [30] Mann H B and Whitney D R 1947 On a test of whether one of two random variables is stochastically larger than the other *Ann. Math. Stat.* **18** 50–60
- [31] Rucci M, Bullock D and Santini F 2007 Integrating robotics and neuroscience: brains for robots, bodies for brains *Adv. Robot.* **21** 1115–29