Attention patterns of autonomously learned bee-inspired flight

Pranav Rajbhandari[1*], Abhi Veda[1*], Matthew Garratt[1], Mandyam Srinivasan[2], Sridhar Ravi[1]

[1] School of Engineering and Technology, University of New South Wales, Canberra, Australia,
[2] Queensland Brain Institute, University of Queensland, Brisbane, Australia
[*] equal contribution

Bio-inspiration for UAV navigation

- UAVs must efficiently navigate complex environments for various applications. Due to various limitations, these systems use monocular camera setups for navigation [1].
- To solve this, robotic systems have been inspired by flying insects, such as honeybees, which, despite simple vision and neural systems, perform similar vision-based navigation successfully [2]. These animals utilize optic flow, a visuo-sensory cue for the motion of objects in their environment [3].

Research Goals

In this work, we train agents for navigation with similar sensory input to a honeybee. Our goal is to determine the attention that agents pay to different regions of sensory input [4]. The findings from this study can lead to a better understanding of the flight behavior of real biological systems, as in [5].

Methodology

Reinforcement Learning & Simulation

Using Deep RL, an agent is trained to fly a quadrotor drone, rewarding it for navigating through a corridor and penalizing crashing.

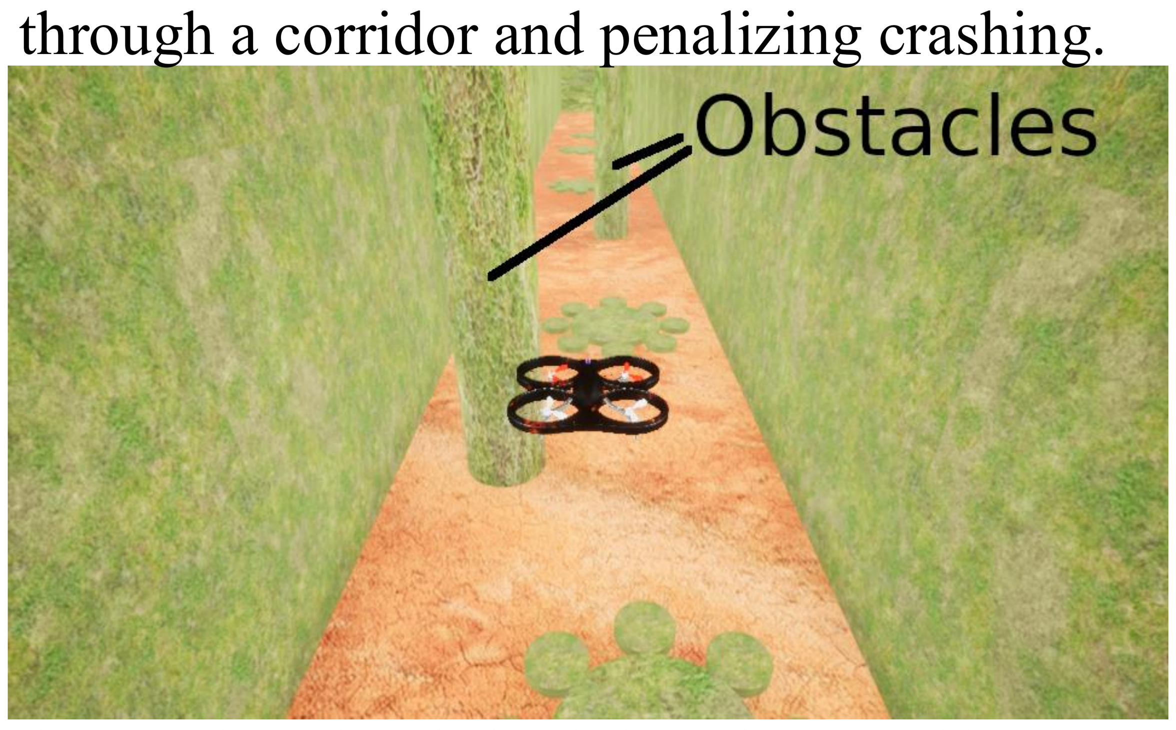

Fig. 1. Quadrotor drone in the AirSim environment with obstacles.

Agent Control Policy

The RL agent is trained using the PPO algorithm. The agent controls the drone's x and y acceleration, with a policy network with architecture as in Fig. 2:

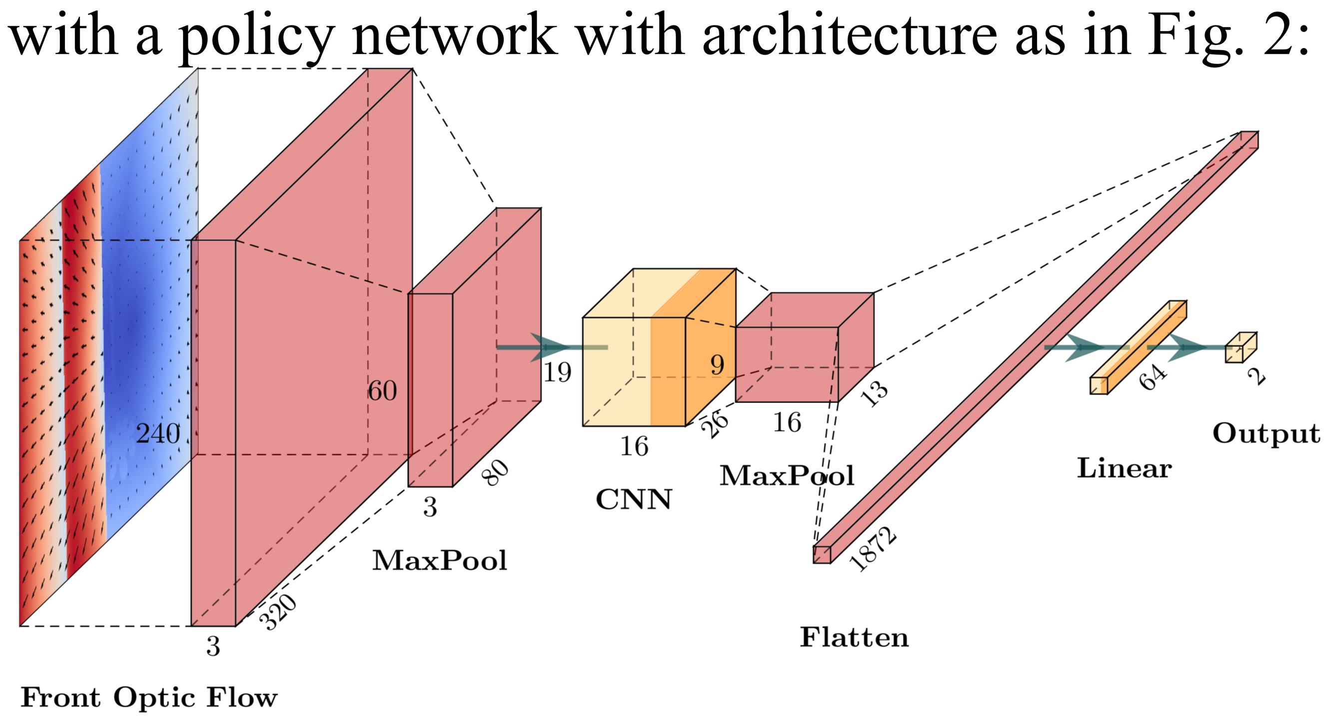

Fig. 2. Policy network architecture, yellow layers have learnable weights.

Optic Flow Perception

- To imitate the input of real bees during their flight, the agent observes optic flow from a monocular camera with a 120° FOV.
- The observation for one time-step is the flow magnitude and orientation at each pixel in the camera's FOV. This is calculated as below:

$$\begin{bmatrix}\dot{u}\\ \dot{v}\end{bmatrix} = \begin{bmatrix}\frac{f_x}{d_{u,v}} & 0 & -\frac{u}{d_{u,v}} & -\frac{uv}{f_x} & -\frac{f_x^2+u^2}{f_x} & -v\\ 0 & \frac{f_y}{d_{u,v}} & -\frac{v}{d_{u,v}} & -\frac{f_y^2+u^2}{f_y} & \frac{uv}{f_y} & u\end{bmatrix}\begin{bmatrix}V\\ W\end{bmatrix}$$

Results

- The trained agents successfully navigate the corridors using only optic flow observations.
- Their trajectories (Fig. 4) appear to center the agent in the non-obstructed region of the corridor and avoid obstacles.
- SHAP values are obtained by sampling a few trajectories from trained agents navigating a corridor and recording the optic flow from each timestep in these trajectories.
- We take the absolute value of these SHAP values and apply Gaussian smoothing.
- The attention patterns for a single timestep are displayed in Fig. 3C.

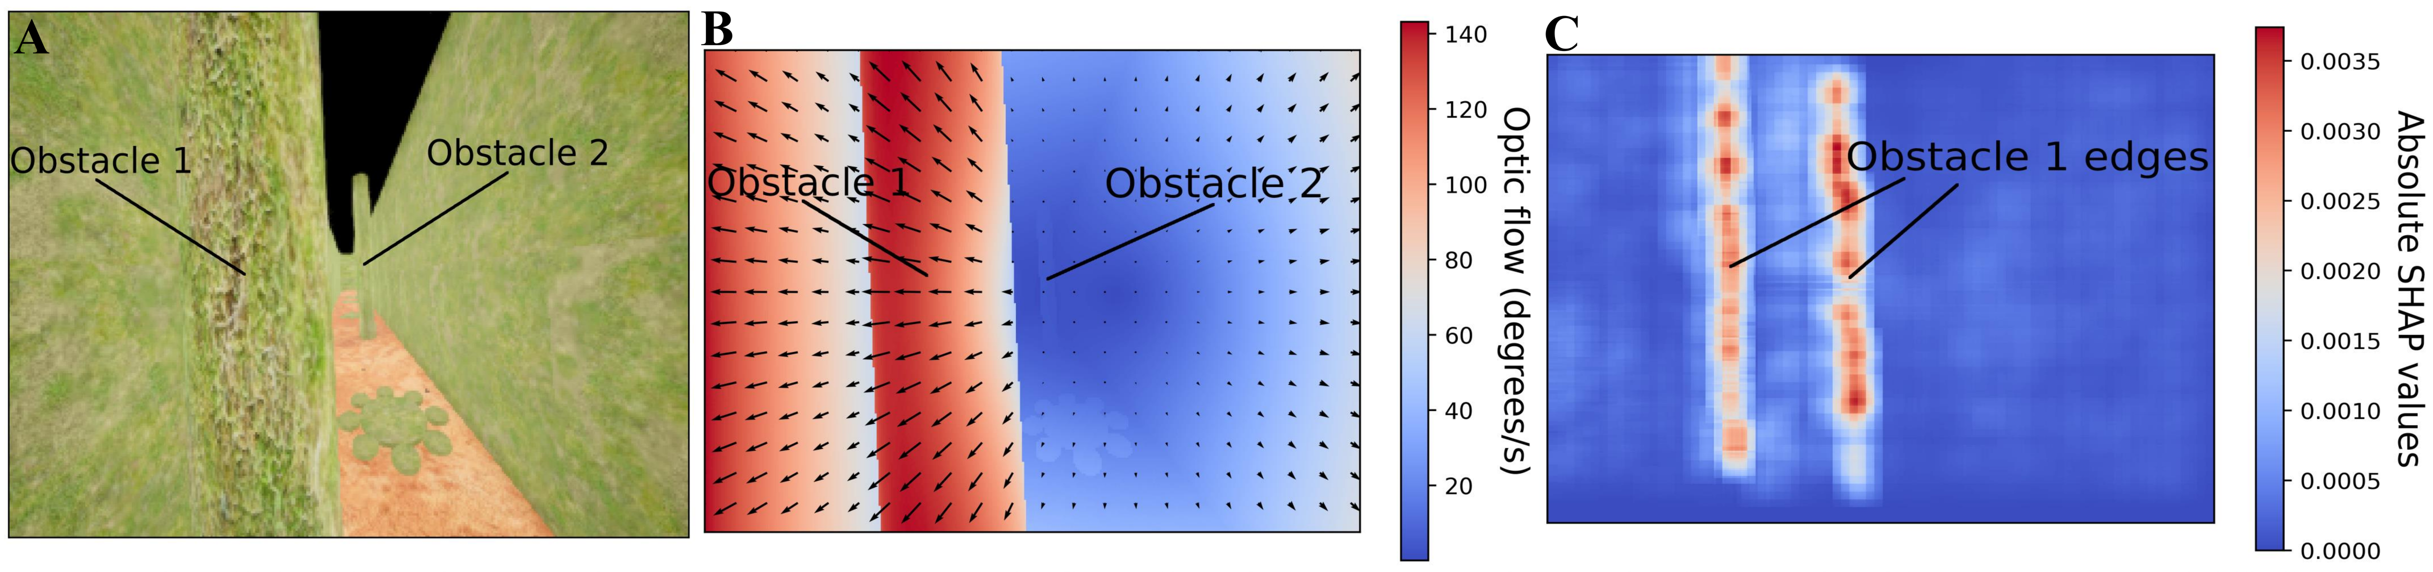

Fig. 3. Visualizing (A) The raw visual scene, (B) The optic flow direction vectors and magnitude perceived by the agent and (C) Sum of the absolute SHAP values denoting the averaged attention pattern of four independently trained agents flying through a corridor.

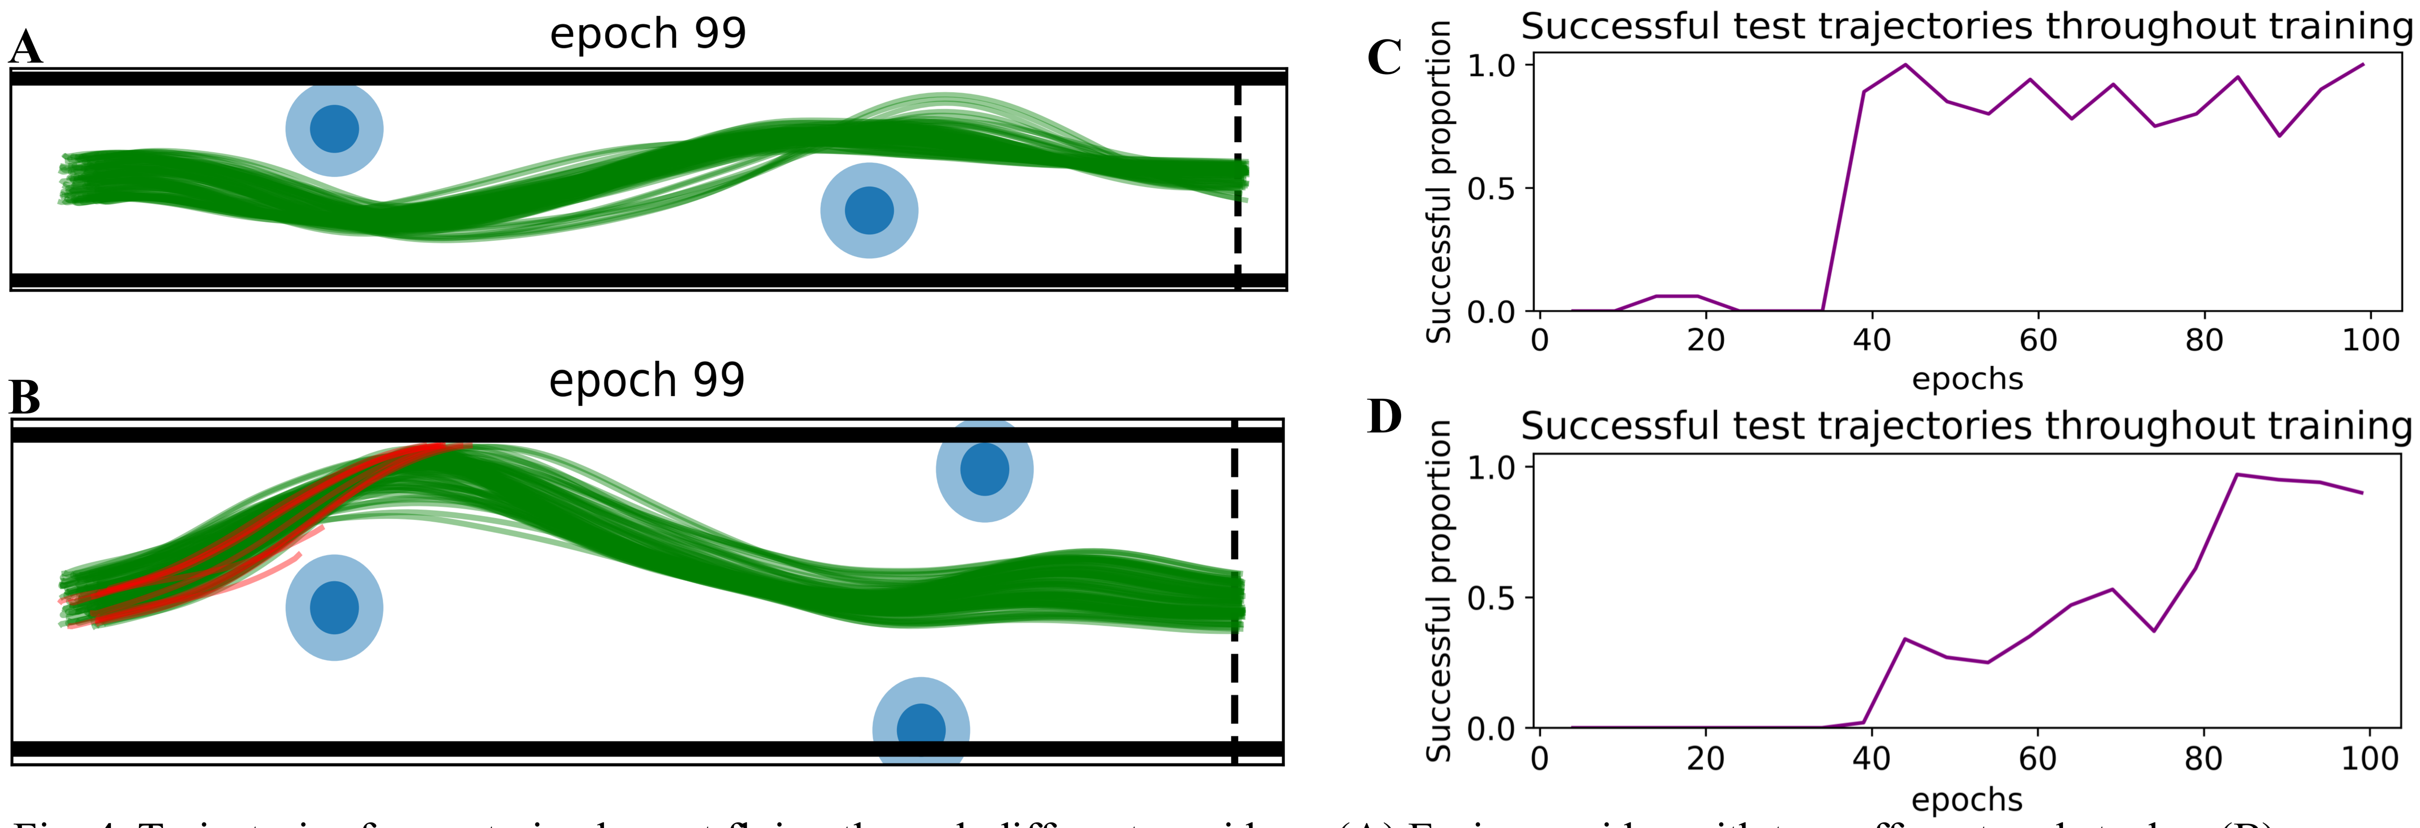

Fig. 4. Trajectories from a trained agent flying through different corridors: (A) Easier corridor with two off-center obstacles. (B) Difficult corridor with three obstacles, one placed centrally and two off-center. Success rate in the (C) easier and (D) difficult corridor.

Discussion

- Regions of high optic flow and regions of discontinuities in optic flow contribute disproportionally to the agents' decision (Fig. 5).
- Trained agents pay most attention to the edges of nearby obstacles. The attention to discontinuities can be explained by agents learning to detect and avoid obstacles, which produce large optic flow in a small visual region.

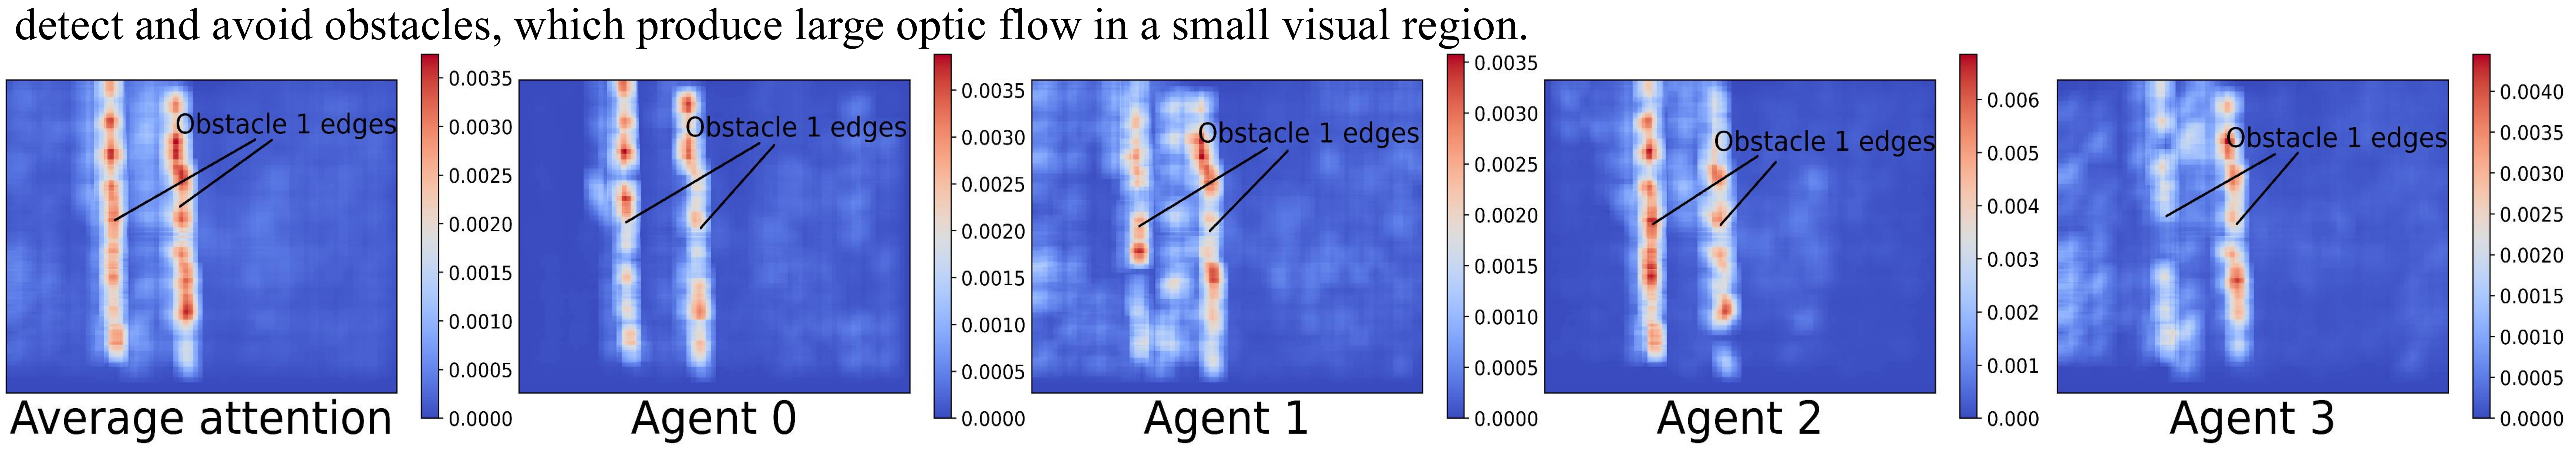

Fig. 5. Comparison of attention patterns for four independently trained agents. The attention across the agents is consistently larger towards the obstacle edges, around which the largest optic flow discontinuity exists. This is reinforced in the average attention plot and suggests that the policy is basing its decision on the largest detected optic flow imbalance.

Conclusions and future work

We show that behavior learned through Deep RL with the same sensory input as a bee induces a similar navigation capability as the organism. The underlying attention pattern of this behavior is a novel contribution and can inspire a control scheme for a UAV. In future work, we aim to compare an agent's attention patterns for various goals, such as moving forward, landing, or passing through gaps.

References

[1] Yingxiu Chang, Yongqiang Cheng, Umar Manzoor, and John Murray. A review of UAV autonomous navigation in GPS-denied environments. Robotics and Autonomous Systems, 170:104533, 12 2023.
[2] Mandyam V. Srinivasan. Honeybees as a model for the study of visually guided flight, navigation, and biologically inspired robotics. Physiological Reviews, 91:413–460, 4 2011.
[3] Martin Egelhaaf. Optic flow based spatial vision in insects. Journal of Comparative Physiology A 2023209:4, 209:541–561, 1 2023.
[4] Miriam Lehrer. Looking all Around: Honeybees Use Different Cues in Different Eye Regions. Journal of Experimental Biology, 201:3275–3292, 121998.
[5] Nick Gravish and George V. Lauder. Robotics-inspired biology. Journal of Experimental Biology,221(7):jeb138438, 2018.