

Understanding visual attention beehind bee-inspired UAV navigation

Pranav Rajbhandari¹, Abhi Veda^{1†}, Matthew Garratt¹, Mandayam Srinivasan² and Sridhar Ravi¹

¹ School of Engineering and Technology, University of New South Wales, Canberra, Australia,

² Queensland Brain Institute, University of Queensland, Brisbane, Australia

(E-mail: prajbhan@cs.cmu.edu, {a.veda, m.garratt, sridhar.r}@unsw.edu.au, m.srinivasan@uq.edu.au)

Abstract: Bio-inspired design is often used in autonomous UAV navigation due to the capacity of biological systems for flight and obstacle avoidance despite limited sensory and computational capabilities. In particular, honeybees mainly use the sensory input of optic flow, the apparent motion of objects in their visual field, to navigate cluttered environments. In our work, we train a Reinforcement Learning agent to navigate a tunnel with obstacles using only optic flow as sensory input. We inspect the attention patterns of trained agents to determine the regions of optic flow on which they primarily base their motor decisions. We find that agents trained in this way pay most attention to regions of discontinuity in optic flow, as well as regions with large optic flow magnitude. The trained agents appear to navigate a cluttered tunnel by avoiding the obstacles that produce large optic flow, while maintaining a centered position in their environment, which resembles the behavior seen in flying insects. This pattern persists across independently trained agents, which suggests that this could be a good strategy for developing a simple explicit control law for physical UAVs.

Keywords: Bio-inspired Robotics; Optic Flow; Deep Reinforcement Learning; Explainable AI;

1. INTRODUCTION

1.1. Bio-inspired autonomous navigation

With the growing utilization of Unmanned Aerial Vehicle (UAV) and Micro Aerial Vehicle (MAV) systems for tasks like environment exploration/mapping, surveillance, and payload delivery, autonomous navigation of these systems has gained relevance. These robots often have constraints, such as inability to use GPS and reduction in available sensors, due to their use cases or maximum payload. However, the computational requirements of these systems remain large due to the high fidelity demands of autonomous aerial navigation. To maintain performance under various constraints, navigation methods developed often take inspiration from biological organisms. These methods are designed to be efficient and robust, as they mimic biological systems that have developed their capability through the process of evolution over millions of years. They serve as an attractive alternative to traditional navigation methods, such as [1-4].

In particular, the honeybee has been the subject of various studies for bio-inspired robotic navigation [5, 6]. One of the main sensory inputs that bees utilize during flight is *optic flow*, the relative motion of elements in the visual field of the observer [7-9]. The bee’s reliance on optic flow arises from limited sensory capabilities, including low visual acuity and resolution of honeybee eyes, limited neural capacity, and a lack of stereo vision. Hence, it is unfeasible for them to rely solely on visual features for navigation. [10, 11]. These constraints also make optic flow a useful sensory input for weight-constrained robotic control, as optic flow can be calculated from video captured by a single camera, which is relatively light compared to other sensory payloads. This has led to the adaptation of honeybee capabilities for various robotic tasks, such as autonomous tunnel navigation [12], visual odom-

etry [13-16], landmark-based navigation [17-19] and autonomous landing [20-22]. These results benefit from the honeybee’s capability for associative learning [23], creating greater understanding of navigational principles of the organism through behavioral experiments.

While these projects design systems whose control is inspired from honeybee behavior, our work uses a learning algorithm to train agents with sensory input similar to a bee to fly. Our goal is to inspect the learned behavior to determine what the agent is paying most attention to in its sensory input. Since there is evidence supporting the theory that bees use optic flow from different areas of vision for different control tasks [24-27], we investigate whether similar patterns arise from learning agents that use the same sensory information.

1.2. Reinforcement Learning

A Markov Decision Process (MDP) is a framework useful for formalizing robotic optimization tasks. MDPs consist of a state space S , action space A , transition function $\mathcal{T}(S | S, A)$, reward function $R: S \times A \times S \rightarrow \mathbb{R}$, and discount factor γ . An agent in an MDP is described by its policy $\pi(A | S)$, which chooses a distribution over actions at each state. Given an MDP and an agent, we may generate a trajectory $(s_0, a_0, r_0), (s_1, a_1, r_1), \dots$, where $s_i \sim \mathcal{T}(\cdot | s_{i-1}, a_{i-1})$, $a_i \sim \pi(\cdot | s_i)$, and $r_i = R(s_i, a_i, s_{i+1})$. An agent’s goal is to maximize the expected *return* $\mathbb{E}[\sum_i \gamma^i r_i]$ over trajectories. A Partially Observable MDP (POMDP) is an extension of this framework that includes an observation space Ω and an observation function $O(\Omega | S)$. An agent in a POMDP must choose action a_i based on its *belief*, a probability distribution over states. We simplify this to a policy $\pi(A | \Omega)$ depending only on the current observation $o_i \sim O(\cdot | s_i)$.

Reinforcement Learning (RL) is a method to solve these tasks by repeatedly sampling trajectories to optimize an agent’s policy. Deep RL uses deep neural networks to define the agent’s policy (the *policy network*),

[†] Abhi Veda is the presenter of this paper.

as well as for functions like the expected return from a state (the *value network*).

In our work, we aim to train robotic agents to perform the task of avoiding obstacles while navigating a tunnel using bee-like optic flow observations. To do this, we structure the task as a POMDP, where agents are given reward for navigating successfully through the tunnel, and a penalty for crashing. Our goal is to inspect the policy of an RL agent that learned this task, to find out to which portion of each observation the agent pays most attention.

1.3. Explainable AI/Shapley Values

To understand the attention patterns of a trained agent, we use methods from Explainable AI, a field concerned with exactly this problem. Specifically, SHapley Additive exPlanations (SHAP) is a method rooted in game theory that assigns values to a set of input features that represent how much each feature contributed to a decision [28]. In our case, the features will be specific regions of observed optic flow. We report the absolute SHAP value since we care about the magnitude of the contribution of each feature, as opposed to whether it had a positive or negative effect. We use the attention patterns of trained neural networks to hypothesize how real bees may utilize different regions of optic flow for a particular task.

2. PREVIOUS WORK

2.1. Buggy robotics

Insects, albeit limited in computational capacity ($\approx 10^4$ times less neurons in the brain than humans [29]), solve similar navigation problems as MAVs with remarkable performance, even with a miniature form factor and no complex sensory system. This has resulted in insect-inspired systems for autonomous navigation tasks [30-34]. Roboticists mimic insects both via physical robot hardware that imitates morphological design and vision capabilities [35-38], as well as via robot behavior that imitates insect-based navigational strategies [39-43]. These low-resource vision and computation capabilities of insects have also influenced neuromorphic systems and techniques for navigation [44, 45]. These methods allow for rapid visual information to be utilized to drive navigational systems at lower power and computational cost compared to traditional hardware (such as SLAM, proximity sensors, or depth cameras).

2.2. Explainable UAV navigation

Previous studies have developed sophisticated UAV navigation methods utilizing sensory input from LiDAR, depth sensors, RGB cameras, or a combination of these [46-48]. These works analyze the learned navigation behavior through the use of SHAP values or Grad-CAM [49]. This allows for an interpretation of how the input features are utilized, giving insight into the working of the control algorithms. However, none of these methods are directly bio-inspired, and none utilize optical flow for sensory input. Through our work, we aim to address this gap by utilizing optic flow in an autonomous navigation




Fig. 1: Screenshot of the AirSim simulation environment

system, and exploring how an autonomous system interprets this sensory input. This method of explaining a bio-inspired autonomous system stems from ideas presented in [50, 51]. We use our results, which reflect how an autonomous agent formulates a control behavior from optic flow, to hypothesize how real bees navigate based on the same sensory input.

3. SIMULATION ENVIRONMENT

We conduct our experiments in a simulated UAV environment, using the AirSim platform [52] for the Unreal physics simulation engine [53]. The agent is a quadrotor drone that is traveling through a constrained tunnel containing cylindrical obstacles (displayed in Fig. 1). There are six tunnels, each with different widths and obstacle configurations. The agent's goal is to reach the end of the tunnel without colliding with the walls or obstacles. We structure this as a POMDP with the OpenAI Gymnasium environment¹ [55]. We define elements of the state space S as states of the physics simulation, and we define sampling from the transition function $\mathcal{T}$ as running the physics simulation for a duration $\Delta t = 0.05$ s. We set discount factor $\gamma = 0.99$.

3.1. Observation Space

To imitate sensory input of real bees during flight, we allow the agent to observe optic flow at each timestep from a camera with field of view 120° mounted to the front of the drone. The optic flow is calculated geometrically, using knowledge of the drone's true velocity and distance from each obstacle. In the case of a real vehicle, there exist algorithms to calculate optic flow from single camera input [56-59]. We calculate optic flow from the formulation presented in [60]:

$$\begin{bmatrix} \dot{u} \\ \dot{v} \end{bmatrix} = \begin{bmatrix} \frac{f_x}{d_{u,v}} & 0 & -\frac{u}{d_{u,v}} & -\frac{uv}{f_x} & -\frac{f_x^2+u^2}{f_x} & -v \\ 0 & \frac{f_y}{d_{u,v}} & -\frac{v}{d_{u,v}} & -\frac{f_y^2+v^2}{f_y} & \frac{uv}{f_y} & u \end{bmatrix} \begin{bmatrix} V \\ W \end{bmatrix} \quad (1)$$

Where f_x and f_y are the focal lengths of the camera in x and y direction, u and v are the horizontal and vertical pixel coordinates, $V = [\dot{x}, \dot{y}, \dot{z}]^T$ is the linear velocity of the vehicle in cartesian coordinates, $W = [W_x, W_y, W_z]^T$ is the angular velocity, and $d_{u,v}$ is the distance of the vehicle from the object at pixel (u, v) .

¹Our experiment implementation is available at [54].

Explicitly, an element of the observation space Ω is an array containing the optic flow magnitude and orientation at each pixel in the camera’s field of view. We represent this as an ‘image’ with three channels: the optic flow magnitude, and the horizontal and vertical components of the normalized optic flow vector. We set angular velocity to zero during optic flow calculation ($W = \mathbf{0}$), since we found that the contribution from the drone’s angular velocity is often much larger than the contribution from its linear motion.³ We visualize an observation in Fig. 3(b), where the agent is traveling forward and right. Here, we display (unnormalized) optic flow vectors as black arrows over a heatmap of the optic flow magnitude.

3.2. Action Space

The agent controls the acceleration of the drone in the basis where the $+x$ and $+y$ directions correspond to the drone’s forward and starboard directions respectively. We project the $+x$ and $+y$ direction onto the global xy plane, so that both of these are normal to the global z axis. The acceleration is added to a target velocity, which is then sent as a velocity command to the AirSim platform. Both the acceleration and the target velocity have a limited magnitude on each dimension, which we denote by $a_{\max} = 3 \text{ m/s}^2$ and $v_{\max} = 2 \text{ m/s}$ respectively. We represent the action space A as $\{\mathbf{a} \in \mathbb{R}^2 : \|\mathbf{a}\|_{\infty} \leq a_{\max}\}$. We fix the drone’s height above the ground in our experiments, as our environment contains only column obstacles and ‘flowers’ below the drone’s flying height.

3.3. Reward Function/Termination

Let m_i be the maximum global x coordinate of the agent throughout timesteps $[0, \dots, i]$. We set the reward $r_i := m_{i+1} - m_i$. The reward is positive if the agent increases the furthest it has traveled, otherwise it is zero. The episode terminates with a penalty ($r_i = -1$) if the agent crashes at timestep i . The episode terminates with no penalty if the agent successfully navigates the tunnel or runs out of time (after 30 s). The motivation behind our reward function is to encourage forward travel while not directly penalizing a path that slows to avoid obstacles.

3.4. Agent Policy

To train an agent in this environment we use Proximal Policy Optimization (PPO), a Deep RL algorithm that defines policy and value networks [61]. For the policy network, we use a network architecture whose layers with learnable weights are a Convolutional Neural Network (CNN) layer with ReLU activation, a fully connected linear layer with Tanh activation, and a linear output layer (Fig. 2). The architecture of the value network is identical to that of the policy network, with the exception of the output layer, which maps to a one-dimensional output instead. We chose this architecture since it is relatively simple and performed well in initial experiments.




Fig. 2: Policy Network Architecture. Yellow boxes represent transitions with learnable weights.

4. RESULTS

We train our agent in six different tunnel configurations, which vary in width and in number of obstacles. At the start of each trial, the agent’s position is randomized at the beginning of a tunnel randomly chosen from the six. We train for 100 epochs, with each epoch consisting of 1024 timesteps ($\approx 50 \text{ s}$ of simulation).

We train four agents independently in this way, so that we may inspect general trends in attention patterns rather than the learned attention of any one specific trained agent. We sample a few trajectories from one of these trained agents navigating a tunnel. We use the recorded optic flows at each timestep in these trajectories to obtain SHAP values for each of the four agents on these observations. We take the absolute value of these SHAP values, apply Gaussian smoothing, and take the average across all four agents. The resulting attention patterns for a single timestep are displayed in Figures 3 and 5.

We find that the trained agents are successful in navigating the tunnels using optic flow observations. Their trajectories (displayed in Fig. 4) appear to center the agent in the non-obstructed region of the tunnel.

We find that regions of discontinuity in optic flow magnitude make a disproportionately larger contribution to the agent’s decisions. This is seen in Fig. 3, where the agents pay most attention to the nearest edges of an obstacle. Aside from this, the agents seem to pay attention to regions of high optic flow, such as the center of the nearest obstacle and the closest parts of the tunnel walls in Fig. 3. We observe this attention pattern throughout the trajectories, in addition to the example shown.

These trajectories and attention patterns could be explained by the learned behavior being to move away from regions with high optic flow, while placing the most value on regions with discontinuities in the optic flow. The reason for the larger attention towards discontinuities could be that discontinuities are a good visual indicator of an obstacle, which create large optic flow values in a small portion of the visual field. If the agent did not handle these cases specially, the thinness of the obstacle may outbalance the magnitude of the optic flow values, resulting in not avoiding a close obstacle due to its small size.

These attention patterns seem to be learned indepen-

³This can be simulated on physical systems either through measuring angular velocity, or through estimating it from observed optic flow.




Fig. 3: The attention pattern of trained agents flying through a tunnel, averaged over four independently trained agents.




Fig. 4: Trajectories from a trained agent flying through different tunnels: (a) Easier tunnel with two obstacles placed off-center, (b) Difficult tunnel with three obstacles, one placed centrally and two placed off-center.




Fig. 5: Comparison of attention patterns for four independently trained agents. The average attention is largest at nearby obstacle edges; this pattern exists in each of the four agent policy networks as well.

dently by all four of our agents (Fig. 5), which indicates this arises as a result of the task and observation space as opposed to arising from a quirk of a particular agent policy network.

5. DISCUSSION

In our experiments, our trained agents appear to balance the optic flow perceived in the left and right regions of their visual field to maintain a centered flight path. We find that the agents pay most attention to nearby edges of obstacles, where discontinuities in optic flow are most significant (Fig. 5). This results in the well documented centering behavior of honeybees [62, 63] appearing in the agents' trajectories (Fig. 4), as well as detection and obstacle avoidance evident from lateral deviation and reduction in forward speed. The similarities of the trajectories to navigation patterns found in bees [6, 12] indicate that real bees also respond strongly to discontinuities in optic flow, such as those induced by the edges

of obstacles. This also suggests that an effective computational navigation strategy should involve discontinuity detection, as opposed to regulating raw optic flow values.

6. CONCLUSION AND FUTURE WORK

We demonstrate that behavior learned through deep RL with the same sensory input as a bee induces similar navigation capability as the organism. The underlying attention pattern for this learned behavior is a novel contribution and can inspire a formalized control scheme for an autonomous UAV. In future work, we plan to use similar methods in a Goal Conditioned RL environment, and compare the trained agent's attention patterns for various goals such as moving forward, landing, or passing through gaps. We also plan to define control laws inspired by the demonstrated agent behavior that induce similar SHAP attention patterns to the trained agents.

REFERENCES

- [1] Yuncheng Lu, Zhucun Xue, Gui Song Xia, and Liangpei Zhang. A survey on vision-based UAV navigation. *Geo-Spatial Information Science*, 21:21–32, 1 2018.
- [2] Haiyang Chao, Yu Gu, Jason Gross, Guodong Guo, Mario L. Fravolini, and Marcello R. Napolitano. A comparative study of optical flow and traditional sensors in UAV navigation. *Proceedings of the American Control Conference*, pages 3858–3863, 2013.
- [3] 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.
- [4] Nasser Gyagenda, Jasper V. Hatilima, Hubert Roth, and Vadim Zhmud. A review of GNSS-independent UAV navigation techniques. *Robotics and Autonomous Systems*, 152:104069, 6 2022.
- [5] Mandyam V. Srinivasan. Honey bees as a model for vision, perception, and cognition. *Annual Review of Entomology*, 55:267–284, 1 2010.
- [6] 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.
- [7] Martin Egelhaaf and Jens P. Lindemann. Path integration and optic flow in flying insects: a review of current evidence. *Journal of Comparative Physiology A* 2025, pages 1–27, 3 2025.
- [8] Martin Egelhaaf. Optic flow based spatial vision in insects. *Journal of Comparative Physiology A* 2023 209:4, 209:541–561, 1 2023.
- [9] Emily Baird, Norbert Boeddeker, and Mandyam V. Srinivasan. The effect of optic flow cues on honeybee flight control in wind. *Proceedings of the Royal Society B: Biological Sciences*, 288(1943):20203051, 2021.
- [10] John Paul Curraea, Yash Sondhi, Akito Y. Kawahara, and Jamie Theobald. Measuring compound eye optics with microscope and microct images. *Communications Biology* 2023 6:1, 6:1–12, 3 2023.
- [11] M. Asmi Jezeera, Pierre Tichit, G. S. Balamurali, Emily Baird, Almut Kelber, and Hema Somanathan. Spatial resolution and sensitivity of the eyes of the stingless bee, *Tetragonula iridipennis*. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, 208:225–238, 3 2022.
- [12] Julien R. Serres and Franck Ruffier. Optic flow-based collision-free strategies: From insects to robots. *Arthropod Structure & Development*, 46:703–717, 9 2017.
- [13] Fumiya Iida and Dimitrios Lambrinos. Navigation in an autonomous flying robot by using a biologically inspired visual odometer. <https://doi.org/10.1117/12.403708>, 4196:86–97, 10 2000.
- [14] Fumiya Iida. Biologically inspired visual odometer for navigation of a flying robot. *Robotics and Autonomous Systems*, 44:201–208, 9 2003.
- [15] Daniel Ribeiro dos Santos. Robotic models for the honeybee visual odometer. page 255, 3 2023.
- [16] Mandyan V. Srinivasan, Shaowu Zhang, Monika Altwein, and Jürgen Tautz. Honeybee navigation: Nature and calibration of the ‘odometer’. *Science*, 287:851–853, 2 2000.
- [17] S. N. Fry and R. Wehner. Look and turn: landmark-based goal navigation in honey bees. *Journal of Experimental Biology*, 208:3945–3955, 10 2005.
- [18] Giovanni M. Bianco. Getting inspired from Bees to perform large scale visual precise navigation. *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1:619–624, 2004.
- [19] Kodi C.A. Cumbo, Samantha Heck, Ian Tanimoto, Travis Devault, Robert Heckendorn, and Terence Soule. Bee-inspired landmark recognition in robotic navigation. *GECCO 2016 Companion - Proceedings of the 2016 Genetic and Evolutionary Computation Conference*, pages 1039–1042, 7 2016.
- [20] Ahmad Hammad and Sophie F. Armanini. Landing and take-off capabilities of bioinspired aerial vehicles: a review. *Bioinspiration & Biomimetics*, 19:031001, 3 2024.
- [21] Pulkit Goyal, Emily Baird, Mandyam V. Srinivasan, and Florian T. Muijres. Visual guidance of honeybees approaching a vertical landing surface. *Journal of Experimental Biology*, 226:jeb245956, 9 2023.
- [22] Saul Thurrowgood, Richard J. D. Moore, Dean Soccol, Michael Knight, and Mandyam V. Srinivasan. A Biologically Inspired, Vision-based Guidance System for Automatic Landing of a Fixed-wing Aircraft. *Journal of Field Robotics*, 31(4):699–727, 2014.
- [23] Karl Von Frisch. Decoding the language of the bee. *Science*, 185:663–668, 8 1974.
- [24] Thomas Labhart. Specialized photoreceptors at the dorsal rim of the honeybee’s compound eye: Polarizational and angular sensitivity. *Journal of Comparative Physiology A*, 141:19–30, 3 1980.
- [25] Miriam Lehrer. Looking all Around: Honeybees use Different Cues in Different Eye Regions. *Journal of Experimental Biology*, 201:3275–3292, 12 1998.
- [26] M. Giurfa, G. Zaccardi, and M. Vorobyev. How bees detect coloured targets using different regions of their compound eyes. *Journal of Comparative Physiology - A Sensory, Neural, and Behavioral Physiology*, 185:591–600, 12 1999.
- [27] Reinhard Seidl and Walter Kaiser. Visual field size, binocular domain and the ommatidial array of the compound eyes in worker honey bees. *Journal of Comparative Physiology A*, 143:17–26, 3 1981.
- [28] Scott M Lundberg and Su-In Lee. A Unified Approach to Interpreting Model Predictions. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach,

- R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [29] Anastasia A. Makarova, Alexey A. Polilov, and Dmitri B. Chklovskii. Small brains for big science. *Current Opinion in Neurobiology*, 71:77–83, 12 2021.
- [30] Reuben Strydom, Aymeric Denuelle, and Mandyam V. Srinivasan. Bio-Inspired Principles Applied to the Guidance, Navigation and Control of UAS. *Aerospace 2016, Vol. 3, Page 21*, 3:21, 7 2016.
- [31] G. C.H.E. de Croon, J. J.G. Dupeyroux, S. B. Fuller, and J. A.R. Marshall. Insect-inspired AI for autonomous robots. *Science Robotics*, 7:6334, 6 2022.
- [32] Jochen Zeil. Visual navigation: properties, acquisition and use of views. *Journal of Comparative Physiology A* 2022 209:4, 209:499–514, 12 2022.
- [33] Yu Hu, Yang Zhang, Yunlong Song, Yang Deng, Feng Yu, Linzuo Zhang, Weiyao Lin, Danping Zou, and Wenxian Yu. Seeing Through Pixel Motion: Learning Obstacle Avoidance From Optical Flow With One Camera. *IEEE Robotics and Automation Letters*, 2025.
- [34] Haibin Duan and Pei Li. Bio-inspired computation in unmanned aerial vehicles. *Bio-Inspired Computation in Unmanned Aerial Vehicles*, 9783642411960:1–269, 11 2014.
- [35] Jayant Ratti and George Vachtsevanos. Inventing a biologically inspired, energy efficient micro aerial vehicle. *Journal of Intelligent and Robotic Systems: Theory and Applications*, 65:437–455, 1 2012.
- [36] Sejal Budholiya, Aayush Bhat, S. Aravind Raj, Mohamed Thariq Hameed Sultan, Ain Umaira Md Shah, and Adi A. Basri. State of the Art Review about Bio-Inspired Design and Applications: An Aerospace Perspective. *Applied Sciences* 2021, Vol. 11, Page 5054, 11:5054, 5 2021.
- [37] Zhe Hui, Yang Zhang, and Gang Chen. Aerodynamic performance investigation on a morphing unmanned aerial vehicle with bio-inspired discrete wing structures. *Aerospace Science and Technology*, 95:105419, 12 2019.
- [38] Renan S. Geronel and Douglas D. Bueno. An overview on bio-inspired UAVs: State-of-the art, challenges and opportunities. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 2025.
- [39] Faten Aljalaud, Heba Kurdi, and Kamal Youcef-Toumi. Bio-Inspired Multi-UAV Path Planning Heuristics: A Review. *Mathematics* 2023, Vol. 11, Page 2356, 11:2356, 5 2023.
- [40] Hans Lehnert, Mauricio Araya, Rodrigo Carrasco-Davis, and Maria Jose Escobar. Bio-Inspired Deep Reinforcement Learning for Autonomous Navigation of Artificial Agents. *IEEE Latin America Transactions*, 17:2037–2044, 12 2019.
- [41] Tingjun Lei, Chaomin Luo, Simon X. Yang, Daniel W. Carruth, and Zhuming Bi. Bio-Inspired Intelligence-Based Multiagent Navigation with Safety-Aware Considerations. *IEEE Transactions on Artificial Intelligence*, 5:2946–2961, 6 2024.
- [42] Aymeric Denuelle and Mandyam V. Srinivasan. Bio-inspired visual guidance: From insect homing to UAS navigation. *2015 IEEE International Conference on Robotics and Biomimetics, IEEE-ROBIO 2015*, pages 326–332, 2015.
- [43] Chelsea Sabo, Alex Cope, Kevin Gurny, Eleni Vasilaki, and James A.R. Marshall. Bio-inspired visual navigation for a quadcopter using optic flow. *AIAA Infotech @ Aerospace Conference*, 2016.
- [44] Thorben Schoepe, Ella Janotte, Moritz B. Milde, Olivier J.N. Bertrand, Martin Egelhaaf, and Elisabetha Chicca. Finding the gap: neuromorphic motion-vision in dense environments. *Nature Communications* 2024 15:1, 15:1–14, 1 2024.
- [45] F. Paredes-Vallés, J. J. Hagenaars, J. Dupeyroux, S. Stroobants, Y. Xu, and G. C.H.E. de Croon. Fully neuromorphic vision and control for autonomous drone flight. *Science Robotics*, 9:591, 5 2024.
- [46] Yunfan Ren, Fangcheng Zhu, Guozheng Lu, Yixi Cai, Longji Yin, Fanze Kong, Jiarong Lin, Nan Chen, and Fu Zhang. Safety-assured high-speed navigation for MAVs. *Science Robotics*, 10, 1 2025.
- [47] Didula Dissanayaka, Thumeera R. Wanasinghe, and Raymond G. Gosine. Explainable Artificial intelligence for Autonomous UAV Navigation. *IEEE International Conference on Intelligent Robots and Systems*, pages 10439–10446, 2024.
- [48] Lei He, Nabil Aouf, and Bifeng Song. Explainable Deep Reinforcement Learning for UAV autonomous path planning. *Aerospace Science and Technology*, 118:107052, 11 2021.
- [49] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *Proceedings of the IEEE International Conference on Computer Vision*, 2017-October:618–626, 12 2017.
- [50] Nick Gravish and George V. Lauder. Robotics-inspired biology. *Journal of Experimental Biology*, 221(7):jeb138438, 2018.
- [51] F Iida and A J Ijspeert. Biologically Inspired Robotics. In *Springer Handbook of Robotics*, pages 2015–2034. 2016.
- [52] Shital Shah, Debadeepta Dey, Chris Lovett, and Ashish Kapoor. AirSim: High-Fidelity Visual and Physical Simulation for Autonomous Vehicles. In *Field and Service Robotics*, 2017.
- [53] Epic Games. Unreal Engine.
- [54] Pranav Rajbhandari and Abhi Veda. Attention patterns in Beebehavior. <https://github.com/pranavraj575/beebehavior>, 2025.
- [55] Greg Brockman et al. OpenAI Gym, 2016.

- [56] G Farnebäck. Two-Frame Motion Estimation Based on Polynomial Expansion. In J Bigun and T Gustavsson, editors, *Image Analysis. SCIA 2003. Lecture Notes in Computer Science*, volume 2749. Springer, Berlin, Heidelberg, 2003.
- [57] Mingliang Zhai, Xuezhi Xiang, Ning Lv, and Xiandong Kong. Optical flow and scene flow estimation: A survey. *Pattern Recognition*, 114:107861, 6 2021.
- [58] M V Srinivasan. An image-interpolation technique for the computation of optic flow and egomotion. *Biological Cybernetics*, 71(5):401–415, 9 1994.
- [59] R Benosman, S.-H. Ieng, C Clercq, C Bartolozzi, and M Srinivasan. Asynchronous frameless event-based optical flow. *Neural Networks*, 27:32–37, 3 2012.
- [60] Ping Li, Matthew Garratt, Andrew Lambert, and Shanggang Lin. Metric sensing and control of a quadrotor using a homography-based visual inertial fusion method. *Robotics and Autonomous Systems*, 76:1–14, 2 2016.
- [61] John Schulman et al. Proximal Policy Optimization Algorithms, 2017.
- [62] Julien R. Serres, Guillaume P. Masson, Franck Ruffier, and Nicolas Franceschini. A bee in the corridor: centering and wall-following. *The Science of Nature Naturwissenschaften*, 95(12):1181–1187, 12 2008.
- [63] Emily Baird, Mandyam V. Srinivasan, Shaowu Zhang, and Ann Cowling. Visual control of flight speed in honeybees. *J Exp Biol*, 208(20):3895–3905, 10 2005.