



OPEN ACCESS

EDITED BY

Arij Daou,
The University of Chicago, United States

REVIEWED BY

Luca Patanè,
University of Messina, Italy
Hannes Rapp,
University of Cologne, Germany

*CORRESPONDENCE

Stephan Lochner
✉ lochner@bio.uni-freiburg.de

RECEIVED 05 July 2024

ACCEPTED 20 August 2024

PUBLISHED 09 September 2024

CITATION

Lochner S, Honerkamp D, Valada A and Straw AD (2024) Reinforcement learning as a robotics-inspired framework for insect navigation: from spatial representations to neural implementation. *Front. Comput. Neurosci.* 18:1460006. doi: 10.3389/fncom.2024.1460006

COPYRIGHT

© 2024 Lochner, Honerkamp, Valada and Straw. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Reinforcement learning as a robotics-inspired framework for insect navigation: from spatial representations to neural implementation

Stephan Lochner^{1*}, Daniel Honerkamp², Abhinav Valada² and Andrew D. Straw^{1,3}

¹Institute of Biology I, University of Freiburg, Freiburg, Germany, ²Department of Computer Science, University of Freiburg, Freiburg, Germany, ³Bernstein Center Freiburg, University of Freiburg, Freiburg, Germany

Bees are among the master navigators of the insect world. Despite impressive advances in robot navigation research, the performance of these insects is still unrivaled by any artificial system in terms of training efficiency and generalization capabilities, particularly considering the limited computational capacity. On the other hand, computational principles underlying these extraordinary feats are still only partially understood. The theoretical framework of reinforcement learning (RL) provides an ideal focal point to bring the two fields together for mutual benefit. In particular, we analyze and compare representations of space in robot and insect navigation models through the lens of RL, as the efficiency of insect navigation is likely rooted in an efficient and robust internal representation, linking retinotopic (egocentric) visual input with the geometry of the environment. While RL has long been at the core of robot navigation research, current computational theories of insect navigation are not commonly formulated within this framework, but largely as an associative learning process implemented in the insect brain, especially in the mushroom body (MB). Here we propose specific hypothetical components of the MB circuit that would enable the implementation of a certain class of relatively simple RL algorithms, capable of integrating distinct components of a navigation task, reminiscent of hierarchical RL models used in robot navigation. We discuss how current models of insect and robot navigation are exploring representations beyond classical, complete map-like representations, with spatial information being embedded in the respective latent representations to varying degrees.

KEYWORDS

insect navigation, reinforcement learning, robot navigation, mushroom bodies, spatial representation, cognitive map, world model

1 Introduction

The purpose of this paper is 2-fold: First, we offer a perspective that links our current understanding of spatial navigation by insect navigation researchers together with that of robotics researchers. We do this largely with the help of the theoretical framework of reinforcement learning (RL), which is a central theme in modern robotics research but has so far had relatively little impact in the field of insect navigation or more broadly in

insect learning. We focus on an analysis of spatial representations in current robot and insect navigation models through the lens of RL. Second, we propose neural mechanisms by which the anatomy and physiology of the insect brain may implement RL-like learning, and finally offer a hypothesis on how recent models describing distinct components of insect navigation can be combined into a hierarchical RL model.

Reliably navigating the world in order to acquire essential resources while avoiding potentially catastrophic threats is an existential skill for many animals. In the insect world, central-place foragers like many bee and ant species stake the survival of the entire colony on individuals' ability to return to the nest after extensive foraging trips. Their remarkable navigational capabilities allow them to do so after only a few learning flights or walks under vastly varying environmental conditions. This is so far unrivaled by any artificial autonomous system. Recent years have seen substantial advances in understanding the underlying mechanisms of insect navigation (INav), tentatively converging on what was coined the "insect navigation base model" (INBM) in a comprehensive review by Webb (2019).

This model and its components – rooted in a rich history of behavioral experiments, modeling, and the neuroanatomy of the insect brain – possess substantial explanatory power and offer a mechanistic, bottom-up picture of navigation. Nevertheless, it is only implicitly related to the high-level objective of efficiently exploiting the resources provided by the environment. On the other hand, robot navigation (RNav) research is driven by the practical goal of enabling robots to perform specific spatial tasks, making reinforcement learning a dominant theoretical framework: An RL agent is trained to optimize its interaction with the environment by accumulating positive rewards while avoiding punishment (negative rewards), which it achieves by learning a specific policy: what is the optimal action to take, given the agent's current state? In contrast to other training paradigms, RL requires no additional external supervision. If the task involves a spatial component, this implies learning a navigational strategy which is optimal for achieving the high-level task.

Successful and reliable navigation depends on a robust and efficient choice of the agent's internal representation of its environment (*mapping*) and its own relative pose therein (*localization*), which will be derived from sensory input but otherwise arbitrarily complex. This *spatial representation* of sensory input then serves as the basis to determine a sequence of suitable actions to accomplish a certain objective (*planning*). Until recently, the dominant approach in RNav decoupled the question of finding a suitable spatial representation from the planning phase in modular architectures: a fixed, feed-forward architecture (usually some variety of 'simultaneous localization and mapping', SLAM, see Fuentes-Pacheco et al., 2015 for a review) is used to infer an *explicit* spatial representation from sensory input based on which a policy is optimizing its actions using e.g., classical planners or learned RL methods. In current research, however, end-to-end learning approaches are increasingly gaining traction, where differentiable neural network architectures are trained to learn policies directly from the sensory input. In order to do so efficiently, these networks usually form *latent* spatial representations within hidden layers of the network as an intermediate step. These *latent representations*

are not pre-determined but learned in order to most efficiently solve the navigation task within the constraints of a specific network architecture.¹

Spatial representations can be further characterized by their "geometric content," i.e., how much of the geometric structure of the environment is encoded in the spatial representation: as a biological example, spatial firing fields of hippocampal place-cells (O'Keefe and Dostrovsky, 1971) in mammals tile the entire (accessible) environment of the agent, giving rise to a dense spatial representation akin to grid-like spatial maps used in RNav, although geometry is generally not thought to be preserved accurately. On the other hand, similarity gradients on a retinotopic (pixel-by-pixel) level, as proposed for visual navigation in insects (Zeil et al., 2003), carry no geometric information about the environment at all. This mirrors the long-standing debate among insect navigation researchers whether insects use cognitive maps for navigation (Dhein, 2023). Following a common negative characterization of an animal *without* a cognitive map: "At any one time, the animal knows where to go rather than where it is [...]" (Hoinville and Wehner, 2018), we can restate the question in the language of RL as follows:

What is the geometric content ("where the animal is") of the - latent or explicit - spatial representation ("what the animal knows") of the RL agent?

Free of anatomical and physiological constraints, recent RNav research has produced a plethora of end-to-end learned navigation models with different architectures and policy optimization routines, as discussed in Section 3.2. The resulting pool of "experimentally validated" latent spatial representations can serve as theoretical guidance when thinking about the way space is represented in the insect brain for successful navigation – both in terms of behavioral modeling and in the experimental search for neural correlates of such representations. Conversely, evidence about certain components of spatial representations in insects, like the existence of spatial vectors encoded in the brain, may guide the design of network architectures for artificial agents. To this end, we will analyze what geometric information is represented (Section 2) and how it is represented in recently successful robot navigation models (Section 3) and in the "insect navigation base model" (Section 4).

Going beyond the conceptual considerations outlined above, the question naturally arises whether a link between insect navigation and RL can be established on a more fundamental level. After a brief formal introduction to RL (Section 5.1), we will investigate how the neuroanatomical components involved in the insect navigation base model, the mushroom bodies (MB) and central complex (CX), could support computations similar to certain simple RL algorithms like SARSA or Q-learning. We present current models of MB neural computation (Sections 5.2.1, 5.2.2), to propose specific hypothetical neural connections and their physiological properties by which the models could be augmented

¹ Although other training paradigms with varying degrees of supervision are being used as well, most of the models discussed can be trained in an end-to-end RL fashion, and we will interpret their latent representations within a common RL framework for conceptual clarity.

to support temporal difference learning. In Section 6, we sketch a how such a model could be integrated with a recent MB based visual homing model (Wystrach, 2023) into a full RL-based visual navigation model. Finally, we discuss in Section 7, what kind of spatial representation would result from such a model, its implication for the cognitive map debate and how it aligns with models currently used in robot navigation.

2 Representations of space from a RL perspective

In robot navigation, the problem of navigation has traditionally been partitioned into the subtasks of *localization*, *mapping*, and *planning*. Mapping and localization operations take (potentially multimodal) sensory input to infer a map of the environment and the agent's pose. It has long been acknowledged that the localization problem is most easily solved by reference to locations of salient landmarks in the world—i.e., a map and conversely, constructing a coherent map requires accurate estimates of the agent's pose. This led to the breakthrough of a suite of techniques collectively known as “*simultaneous localization and mapping*” (SLAM) (Mur-Artal and Tardós, 2017; Engel et al., 2014; Endres et al., 2012; Fuentes-Pacheco et al., 2015). Most SLAM techniques combine landmark/feature recognition with odometry to maintain a joint (often probabilistic) representation of the environment and the agent's pose therein, which we will refer to as the *spatial representation* $\phi \in \Phi$ of the sensory input $v \in \mathcal{V}$. We denote sensory input with v , since the paper will focus on visual navigation, for simplicity. Multimodal input spaces are of course possible and highly relevant for a realistic understanding of insect navigation.² Based on ϕ , the planning stage then determines a sequence of actions $a \in \mathcal{A}$ in order to achieve the objective of the navigation task (see Figure 2A). This in turn can be based on planning-based or learned methods. In the following, we analyze the spatial representations Φ found in current robot and insect navigation models. Since these representations differ in many aspects, we first define two dimensions along which our analysis is structured: “*What is represented?*” and “*How is it represented?*”.

2.1 What is represented? The geometric content of the spatial representation ϕ

Confining our discussion to 2D, the simplest, “geometrically perfect” representation of the environment could be imagined as an infinitely extended and infinitesimally spaced grid, filled with binary “occupancy” values. While it may prove useful to enrich the map with layers of meaning (object categories, valuations, etc.) by adding semantic channels, the geometric information is captured fully by this single layer.³ Any practical representation of space, however, must be an abstraction of this ideal to varying degrees,

² Note the definition of “visual input” may differ between INav and RNav and go beyond pure RGB pixel values. For example, insects can visually infer compass cues from celestial polarization patterns, while depth information obtained from RGB-D cameras is often used RNav.

trading off density and geometric accuracy for improved coding efficiency and storage capacity (see Figure 1).

2.1.1 Grid-based maps

The most straightforward simplification is a grid with finite extent and resolution. Grid-based occupancy maps have a long history in SLAM approaches to robot navigation (e.g., Gutmann et al., 2008; Mur-Artal and Tardós, 2017; Engel et al., 2014; Endres et al., 2012) where a probabilistic occupancy grid map is predicted from a time series of observations, as a joint estimate for both the agent's position (localization) and layout of the environment (mapping). More recent methods use hierarchical multi-scale approaches (Zhu et al., 2021), neural radiance fields (Rosinol et al., 2023) or Gaussian splatting (Matsuki et al., 2023) for highly accurate reconstructions. Grid-based maps are both spatially dense and geometrically faithful. The first aspect of such a representation is also reflected in the spatial firing fields of the so-called place-cells in the mammalian hippocampus (O'Keefe and Dostrovsky, 1971). For example, Rich et al. (2014) show that place field density is uniform across a given environment, indicating that all regions of the environment are represented. There is, however, no hard evidence that place-cell representation preserves the geometry of the environment. Furthermore, recent interpretations (Fenton, 2024) of observed “remapping” of spatial fields view place-cell like firing patterns as particular projections of conjoint, collective population activity on a neural attractor manifold, as opposed to the original view, where spatial position was thought to be encoded by dedicated single-neuron activity. It is currently unknown whether insects also possess neuronal populations with similar place-cell like activity. We discuss potential candidate cell types in the insect brain in Section 7.

While more classical approaches focus on binary occupancy or probabilistic occupancy encodings as inputs to motion planners, learning based methods have also encompassed higher-dimensional contexts such as semantics (Wani et al., 2020; Schmalstieg et al., 2022; Younes et al., 2023) or potential functions (Ramakrishnan et al., 2022) as additional channels in these maps.

2.1.2 Vector maps

As a next level of abstraction useful for sparsely populated maps, one could store only the grid indices of occupied cells, instead of an occupancy value for every cell. Increasing the accuracy by replacing grid indices with actual (Cartesian) coordinates with respect to some common origin, we arrive at a vector map, in which geometric relations in the world are represented by relative vectors between salient locations (vector nodes). Stemmler et al. (2015) show how a vector-like spatial representation can be decoded from grid-cell⁴ activity in the mammalian medial entorhinal cortex by

³ Depending on the choice of origin and orientation of the grid axes, this map could be either *allocentric* or *egocentric*. Any sensory (visual) input is by definition *egocentric*.

⁴ Grid-cells have (spatially) periodic firing fields, which can serve as a basis for representing spatial vectors. This is not to be confused with a grid-like spatial representation, which would be more compatible with (non-periodic)

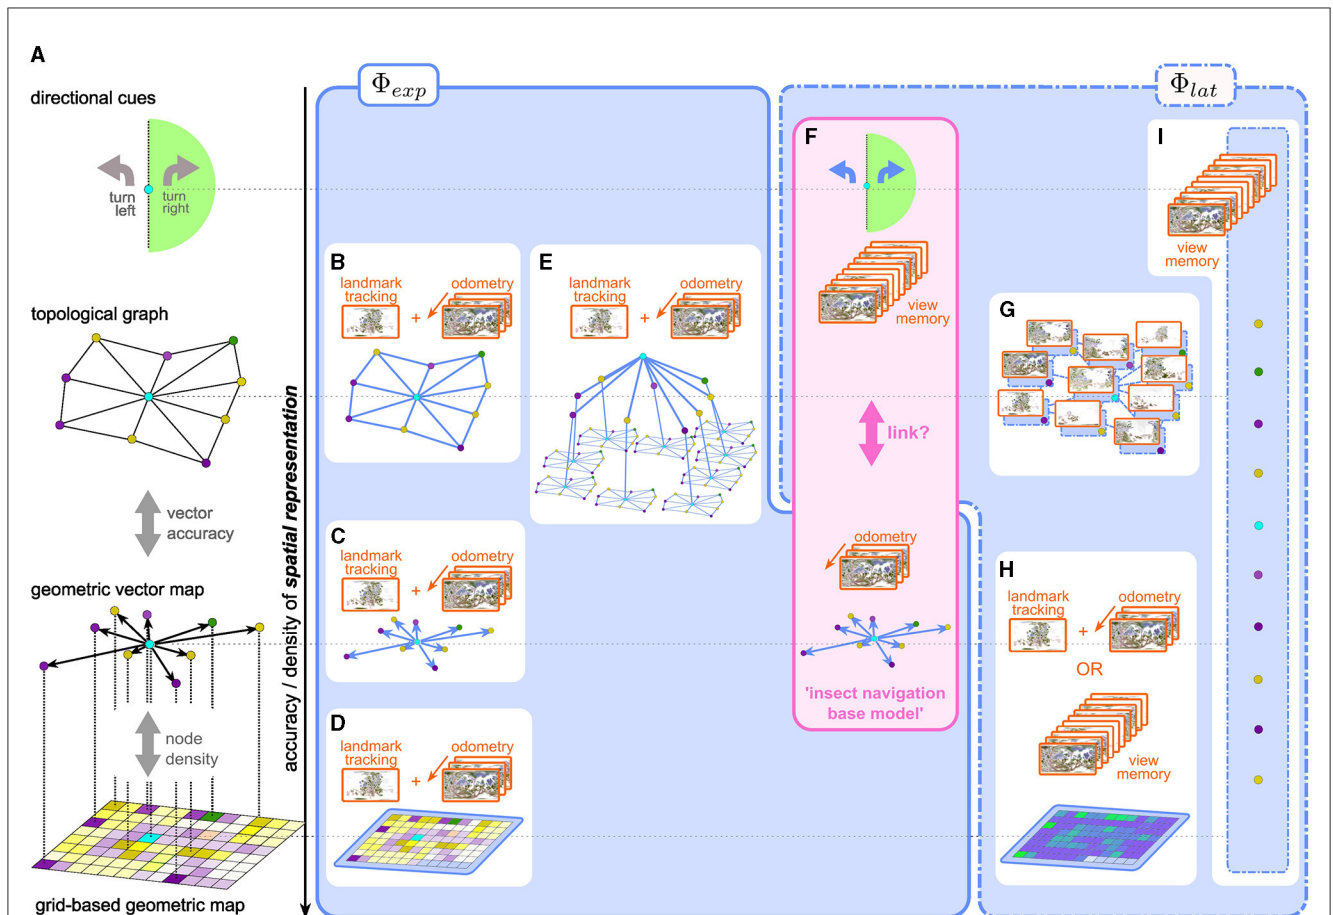


FIGURE 1

(A) Representations of space with varying degrees of geometric accuracy (geometric content) (B–H) explicit and latent spatial representations in robot and insect navigation. (B–D) Visual SLAM methods use a combination of odometry and landmark/feature tracking and matching to construct explicit topological maps (B), vector/landmark (C), or grid-based (D) maps. (E) Scene-graphs (Hughes et al., 2022; Werby et al., 2024; Honerkamp et al., 2024) construct hierarchical, object-centric graph representations to disassemble large scenes into regions, objects, etc., based on SLAM-like mapping approaches. Edges usually encode predefined, i.e., explicit, relational, and semantic attributes. (F) the insect navigation base model uses two distinct mechanisms which result in two spatial representations: an explicit vector map, built from visual odometry alone without the need for landmark recognition and mapping, and latent directional cues learned from a view memory. It is not clear if and how these two are linked to form a unified representation of space. (G–I) more recent approaches in robot navigation use latent representations. (G) topological latent representations (Shah et al., 2023 e.g., RECON) (H) grid-based latent representations can be built from an explicit view-memory based architecture using RNNs (Henriques and Vedaldi, 2018, e.g., MapNet.) or a learnable mapping module as in CMP (Gupta et al., 2019). While the grid structure is fixed, the content of the grid (i.e., the interpretation of the values stored in the grid) is latent, as opposed to an explicit occupancy or semantic grid map (D). (I) Unstructured memory based approaches like SMT (Fang et al., 2019) or unstructured RNNs (like Beeching et al., 2019) learn a completely abstract latent spatial representation, which defies classification along the vertical axis in the figure.

combining populations representing different spatial scales. As we discuss below, the insect navigation base model assumes a vector-based representation of the global geometry of the environment, using a different neural implementation based on phasors (Stone et al., 2017; Lyu et al., 2022).

2.1.3 Topological graphs

If the vector information between connected nodes becomes inaccurate, the vector map gradually loses geometric information and transforms into a topological map, to the extreme case where nodes are connected only by binary “reachability” or

place-cell like activity, where the receptive field corresponds to a particular position within the spatial grid.

“traversability” values. A less extreme case would be a “weighted graph” representation, where edge weights could represent the Euclidean (or temporal) distance between nodes, preserving some geometric information, but not enough to uniquely reconstruct the map. Besides the obvious advantage of memory efficiency, a topological representation may be preferred over geometric maps (as argued for by Warren et al., 2017 in humans) for a different reason: It is more robust to inaccurate or corrupted measurements and therefore a more reliable representation of the coarse structure of the environment, which can then be combined with other mechanisms for local goal finding. Many outdoor navigation approaches in RNav construct topological graphs of the environment (e.g., Shah et al., 2023; Shah and Levine, 2022; Engel et al., 2014). The gradual transition between the map types described above is illustrated in Figure 1A.

2.1.4 Scene-graphs

As another alternative to dense maps, scene graphs (Figure 1E) have arisen as sparse environment representations that disassemble large scenes into objects, regions, etc., and represent them as nodes (Hughes et al., 2022; Gu et al., 2023; Werby et al., 2024). The resulting representation provides a hierarchical and object-centric abstraction that has proven useful in particular in higher-level reasoning and planning (Rana et al., 2023; Honerkamp et al., 2024). In contrast to pure geometric representations, edges mainly focus on semantic or relational attributes, resorting back to grid-based maps for more detailed distance calculations.

All of the above representations establish a relation *between multiple salient locations in the world*, including the agent's own position, and therefore represent knowledge about *where the agent is*.

2.1.5 Directional cues relative to salient location(s)

On the other hand, one could imagine a spatial representation of sensory input that encodes a relation between the agent and salient locations, *without knowledge about how these relate to each other*. For example, the insect navigation base model proposes the use of view memories, which are not attached to any specific location, as discussed in more detail in Section 4. One can interpret visual similarity as a proxy for the distance to the stored view and the similarity gradient as a directional cue (Zeil et al., 2003) toward the location of the snapshot. More recent models based on visual familiarity (Baddeley et al., 2012; Ardin et al., 2016) allow visual homing based on stored view memories regardless of the temporal sequence or locations of the stored views. Wystrach (2023) proposes a visual steering model that categorizes current views into left/right facing with regard to a specific location. These models demonstrate that spatial representations that tell the agent *where to go, rather than where it is*, are sufficient to support surprisingly complex navigation behavior.

2.2 How is it represented? Explicit and latent representations of space: ϕ_{exp} and ϕ_{lat}

We introduce some formal definitions to pose the navigation task as a reinforcement learning problem. Note that while we illustrate the following considerations in the context of RL, they equally apply to other learning paradigms used in the robot navigation literature. RL is usually formalized as a Markov Decision Process (MDP),⁵ which is specified as a 4-tuple $(\mathcal{S}, \mathcal{A}, P, R)$: *State space* \mathcal{S} and *action space* \mathcal{A} characterize the agent, while the

environment⁶ is specified by the (probabilistic) transition function

$$P(s'|s, a), \quad s, s' \in \mathcal{S}, a \in \mathcal{A} \quad (1)$$

between states s and s' , given action a , and a reward function

$$R(s, a), \quad s \in \mathcal{S}, a \in \mathcal{A}. \quad (2)$$

At each timestep, the agent moves across the state space by choosing an action, which determines the next state according to Equation 1, and receives rewards according to Equation 2. The agent's objective is to learn a (probabilistic) policy

$$\pi(a|s), \quad s \in \mathcal{S}, a \in \mathcal{A} \quad (3)$$

over actions given the agent's current state, such that under repeated applications of π , starting from any state s , it maximizes the expected discounted cumulative reward, which we will discuss in more detail in Section 5.1. Note that our notation is meant to implicitly include policies over temporal sequences of states, e.g., eligibility traces (Sutton and Barto, 2018).

In the context of a navigation task, as outlined in Section 2.1, there are now two possible choices for the state space \mathcal{S} of the agent: the conventional approach was to use modular SLAM methods to construct a spatial representation from the sensory input, and then use this *explicit* representation as the state space of the agent: $\mathcal{S} \equiv \Phi_{exp}$ (see Figure 2B). The agent effectively only solves the *planning* sub-problem by either planning or learning a policy over a space of spatial representations whose geometric content is pre-determined by the specific SLAM implementation (e.g., Figures 1B–D).

The other possibility is to model the navigation task as an end-to-end RL (or generally end-to-end learning) problem. This more recent approach takes the raw sensory input as the RL state space: $\mathcal{S} \equiv \mathcal{V}$ (Figure 2C). The policy learned in this case represents a joint implementation of all the sub-problems of the navigation task. In order to achieve this, a sufficiently expressive network architecture needs to be chosen for learning π . For example, Deep Reinforcement Learning (DRL) leverages deep neural networks as trainable function approximators (see Zhu and Zhang, 2021; Zeng et al., 2020 for reviews of DRL for navigation tasks). Crucially, successful end-to-end learning of a navigation task implies the existence of an implicit, or *latent spatial representation* $\Phi_{lat} \neq \mathcal{S}$ within the network architecture for π (Figures 1G–I). Notably, this representation is learned dynamically, ideally converging toward a representation that most efficiently encodes spatial features, not (fully) determined prior to training, relevant for the navigation task. However, the space of possible representations is constrained by the network architecture, which allows imposing certain structural characteristics.

From a biological perspective, the distinction between latent and explicit spatial representation is largely a question of plasticity, recurrent connections, and time scales: In order for a latent (i.e., learned) representation to emerge, sufficient synaptic plasticity is required in the neuronal populations that encode it. Modulation of the synaptic connections based on a training error would require

⁵ Technically, most classical RL navigation problems are formulated as partially observable MDPs (POMDP), which operate on probabilistic belief states about the unobservable true state.

⁶ "Environment" in this context may also include the sensory processing apparatus.

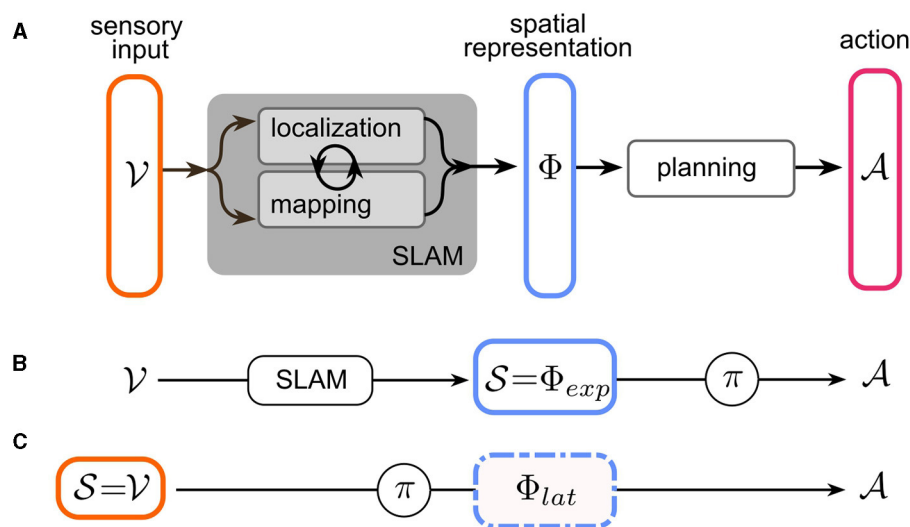


FIGURE 2

Navigation as a Markov Decision Process (MDP). **(A)** The three sub-problems of navigation: Mapping and localization operations take sensory input from an input space \mathcal{V} to infer a map of the environment and the agent's pose, in a space Φ of joint spatial representations. These tasks are usually solved together using simultaneous localization and mapping (SLAM). Based on the spatial representation, the agent plans a sequence of actions from action space \mathcal{A} to achieve the objective of the navigation task. **(B)** In modular robot navigation, only the planning stage is represented as an MDP: existing SLAM methods are used to construct an *explicit* spatial representation Φ_{exp} , which serves as the state space \mathcal{S} of the MDP. The learned policy $\pi(\mathcal{A}|\Phi_{exp})$ then "plans" the action. **(C)** End-to-end RL navigation directly uses the input space \mathcal{V} as the MDP state space \mathcal{S} . Localization, mapping, and planning are jointly solved by learning policy $\pi(\mathcal{A}|\mathcal{V})$. The spatial representation Φ_{lat} is now *latent* in the hidden layers of the learnable (deep) policy network.

some recurrent connectivity to convey that signal (see Section 5). Spatial representations in biological agents likely contain both explicit and latent components. For example, visual processing in the insect optic lobes shows relatively little experience-dependent plasticity, but the mushroom body, which receives high-level sensory input from the visual and other sensory systems, is known as the locus of much insect learning. From a more conceptual point of view, one could argue that the entire neuronal circuit is plastic on an evolutionary timescale. The full spatial representation can then be interpreted as a *latent representation* shaped by ecological constraints over many generations (see Section 4.6). While modeling explicit neuronal implementations of learning mechanisms would be nonsensical in this case, it may be interesting to explore evolutionary computation approaches to RL (see Bai et al., 2023 for a recent review) to serve as normative models for a unified spatial representation in navigating insects. Realistically, the ecological constraints under which such an evolutionary model learns would entail the full behavioral task space of a particular species, potentially giving rise to efficient representations which generalize well to tasks other than navigation. This is generally regarded as a distinguishing feature which sets biological representations apart from most machine learning models.

3 Spatial representations in robot navigation

Having discussed *what* is represented, here we discuss *how* space is represented in robots.

3.1 Explicit spatial representations: variations of SLAM

For completeness, we will very briefly discuss traditional robot navigation models that construct explicit spatial representations using some variation of SLAM. Fuentes-Pacheco et al. (2015) provide a concise review of popular visual SLAM approaches, which operate on different modalities (monocular, stereo, multi-camera, or RGB-D vision) and use different probabilistic [Extended Kalman Filter (EKF), Maximum Likelihood (ML) or Expectation Maximization (EM)] or purely geometric ("Structure from Motion") approaches to maintain a joint representation of the agent's pose within a map of the environment. These map representations come in any of the flavors discussed above. Egocentric occupancy grid maps (Gutmann et al., 2008; Xiao et al., 2022; Schmalstieg et al., 2022 see Figure 1D) are common for dense indoor environments where obstacle avoidance is paramount.

Vector maps (e.g., Klein and Murray, 2007, see Figure 1C), only encode the relative locations of salient features (landmarks) which are tracked in the (retinotopic) camera view across time. Typically, these methods (Mur-Artal and Tardós, 2017) use a loop of (visual) odometry based on these landmarks⁷ and landmark prediction based on self-motion estimates to maintain a joint probabilistic estimate of robot and landmark positions. Loop-closure – the recognition of previously visited locations—allows for the correction of accumulated errors from imperfect odometric information (drift). In contrast to the vector memory in the INBM

⁷ The landmarks encoded in the map and the ones used for odometry need not be identical.

discussed below, this concept of landmark based maps goes beyond our earlier conceptual definition of vector maps: not only are the physical landmark locations stored in the vector map, but they are also linked to their retinotopic locations in the camera frame.

Topological SLAM methods (Konolige et al., 2009; Engel et al., 2014; Greve et al., 2023; Vödisch et al., 2022, see Figure 1B) are particularly useful for mapping larger areas: the world is represented as a graph in which nodes are key-frames (“sensor snapshots”) representing the camera pose. Nodes are connected by edges which represent the relationship between poses obtained from odometry or loop-closure. Global optimization ensures convergence of the topological map. Nevertheless, Fuentes-Pacheco et al. (2015) state that due to “the lack of metric information, [...] it is impossible to use the map for the purpose of guiding a robot”, a limitation which has been overcome by using latent topological representations as in Shah et al. (2023), discussed below.

3.2 Latent representations

Recently, the attention of RNav research has shifted toward end-to-end learning approaches. While these offer the possibility of abstract spatial representations, some implementations choose architectures that constrain the spatial representation to known templates. In the following, we will map out (see Figures 1G–I). the space of possible latent representations along a (non-exhaustive) selection of instructive examples:

The Cognitive Mapper and Planner (CMP) model (Gupta et al., 2019) uses a fully differentiable encoder-decoder architecture to create a grid map of the environment. Instead of occupancy (or pre-defined semantic) values, however, “The model learns to store inside the map whatever information is most useful for generating successful plans”, making the map a latent representation.⁸ Earlier, the same authors (Gupta et al., 2017) suggested a latent representation that combines grid-based with vector (landmark) based maps by synthesizing a global allocentric grid-map from multiple local egocentric grid maps at salient locations. Learning a map from egocentric observations can be viewed as storing encoded egocentric views in a map-like memory. Explicit memory-based models like MapNet (Henriques and Vedaldi, 2018) use a Long Short-Term Memory (LSTM) type Recurrent Neural Network (RNN) with convolutional layers to encode and continually update a grid-map-like state vector by egocentric observations (see Figure 1H).

The RECON (Rapid Exploration Controllers for Outcome-driven Navigation) model by Shah et al. (2023) uses a network architecture whose latent representations capture the topology of a large-scale environment. The map is represented as a graph with egocentric views (“goal images”) at specific locations as nodes (see Figure 1G), which are determined by a goal-directed exploration algorithm.⁹ The model employs a variation of the information bottleneck architecture (Alemi et al., 2019): an encoder-decoder

pair, conditioned on the current egocentric view—learns to compress the goal image into a latent representation (conditional encoder), which is predictive of both the (temporal) distance to the goal, and the best action to reach it (conditional decoder). The encoder and decoder are trained together in a self-supervised manner to learn the optimal (most predictive) latent representation, with the actual time to reach the goal as ground truth. Crucially, the resulting conditional latent representation now encodes the relative distance to the nodes, and thus the topology of the environment. In contrast to topological SLAM models, goal-directed actions are learned alongside the topology, enabling successful robot navigation.

Memory-based approaches like MapNet are based on the insight that all spatial representation is inherently contained in the history of previous observations. However, these need not necessarily resort to fixed grid-like spatial representations. Unstructured end-to-end RL approaches using RNNs (e.g., Beeching et al., 2019) build completely abstract latent representations. Similarly, the Scene Memory Transformer (SMT) architecture (Fang et al., 2019) learns an abstract representation free of inductive biases about the memory structure. Instead of updating an RNN state vector with each observation, an efficient embedding of every observation is stored in an unstructured scene memory. This serves as the state space for an attention-based policy network based on the Transformer architecture (Vaswani et al., 2017), which enables the model to transform the embedding of each memory item according to a specific context. In a nutshell, the transformer blocks are used to “[...] first encode the memory by transforming each memory element in the context of all other elements. This step has the potential to capture the spatio-temporal dependencies in the environment” (Fang et al., 2019). Thus, the encoded scene memory contains a completely abstract latent spatial representation without any preimposed structure (Figure 1I). A second attention block is then used to decode the current observation in the context of the transformed (encoded) scene memory into a distribution over actions. The lack of prior assumptions about spatial representation makes this model very versatile and allows applications in a variety of navigation domains. Wani et al. (2020) compare models using map-based and map-less spatial representations on a multi-object navigation task.

It is important to note that in most of the above approaches, latent representations are built based on a large store of memories of past experiences, which are then transformed into representations most suitable for navigation. One might argue that in this case, learning the navigation task end-to-end effectively ‘reduces’ to the optimization of a memory storage and retrieval process. This makes sense in silicon-based agents, where storage capacity comes cheap. However, for biological brains—especially in insects—storage capacity is very limited, and solving complex tasks must rely more on computational efficiency rather than memory. One possible exception is RECON, where the memory is explicitly compressed into an efficient topological representation, bearing some similarity to how insects might solve navigation: As we will summarize in the next chapter, the insect brain has evolved to compute very effective vector representations of space, and we will sketch in Section 6 how a RL mechanism could be built on top of this innate representation

⁸ This is in contrast to (e.g., neural SLAM Chaplot et al., 2020), which is explicitly trained to produce an occupancy map against a ground truth.

⁹ Akin to the “key-frames” in pose graphs (Engel et al., 2014).

to combine the flexibility of memory based learning with computational efficiency.

4 Spatial representations of the insect navigation base model

In this section, we discuss *how* space is represented in insects. We describe constituent components of the proposed insect navigation base model INBM (Webb, 2019) and analyze inherent spatial representations in the light of the previous discussion for artificial agents (see Figure 1F). Current INav research has identified three main mechanisms as the minimal set of assumptions that may be sufficient to explain observed navigation behavior.

4.1 Path integration

Central place foraging insects are able to maintain a reasonably accurate estimate of their position with respect to a central nest location as a vector-like representation, known as the path integration (PI) home vector. Stone et al. (2017) propose an anatomically constrained model for path integration in the central complex (CX) region of the bee brain: a self-stabilizing representation of the current heading direction is maintained in the ring-attractor architecture of the protocerebral bridge (PB): Neuronal activity of TB1 neurons in eight (per hemisphere) columnar compartments of the PB encodes heading direction relative to the sky compass, projected onto eight axes shifted by $360/8 = 45^\circ$, leading to a periodic, sinusoidal activity pattern.¹⁰ Another population of CX neurons (CPU4) accumulates a speed signal derived from optic flow, modulated by the current heading direction signal from the PB neurons. As a result, the PI home vector is again (redundantly) encoded by its projection along eight axes. This representation essentially amounts to a (discrete) phasor representation of the home vector, with the amplitude and phase of the periodic signal representing its length and angle, respectively. The home vector can then be used to drive the animal back toward the nest. Note that this ring attractor structure can be viewed as an example of a broader class of attractor networks used to represent spatial information across species, such as head-direction or grid-cell activity in the mammalian entorhinal cortex (for a comparative review see Khona and Fiete, 2022). In the context of this work, we want to stress two important aspects of PI in flying navigators: First, it must rely to a large extent on vision alone, since proprioceptive modalities used by walking insects like desert ants are highly unreliable due to wind drift and other atmospheric parameters. We can therefore interpret the PI home vector as a predominantly visual representation of space.¹¹ Secondly, for the same reason, basing PI on *heading direction* is an oversimplification since heading and traveling

direction will often differ. Lyu et al. (2022) proposed a circuit model of the fly (*Drosophila melanogaster*) CX, demonstrating how a representation of the *allocentric traveling direction* can be computed from heading direction and *egocentric directional* optic flow cues by a phasor-based neural implementation of vector addition. However, this important implementation detail does not invalidate the general PI mechanism outlined above.

4.2 Vector memory

Efficient navigation entails more than returning to the nest: Foragers need to be able to reliably revisit known food sources. The INBM posits that whenever insects visit a salient location, they store the current state of the home vector in a vector memory. Le Moél et al. (2019) suggest a mechanism where an individual vector memory is stored in the synaptic weights of a memory neuron, which forms tangential inhibitory synapses onto all directional compartments of the CPU4 population. Activation of this neuron when the home vector is zero would leave a negative imprint of the memorized PI vector (i.e., the vector from the nest to the remembered location) in CPU4 activity, driving the animal to recover a CPU4 activity corresponding to a zero vector (which is now the case at the remembered location, where vector memory and home vector are equal). If the initial home vector is not zero, this mechanism for vector addition effectively computes the direct shortcut to the remembered location, an ability frequently cited as strong evidence for the existence of a cognitive map.¹² The neuronal mechanisms for storing, retrieving, and choosing between multiple vector memories remain speculative. For the former, the authors suggest dopaminergic synaptic modulation directly at the CPU4 dendrites, providing direct reinforcement from extrinsic rewards (food).

This combination of accurate path integration and vector memories constitutes a vector map, i.e., a geometrically accurate (within limits of PI accuracy) representation of the world, which the insect can access for navigation, as long as the PI vector is not corrupted or manipulated. Unlike landmark-based maps in vSLAM, the vector locations are not associated with any visual landmarks or features in retinotopic space.

4.3 View memory

A large body of INav research has been concerned with the ability to return to the nest when an accurate PI vector is not accessible to the animal, making it reliant on visual homing and route-following mechanisms. Originating from the snapshot model (Cartwright and Collett, 1983) which matched the retinal positions of landmarks between the current view and stored snapshots, more recent models suggest retinotopic representations of a low-resolution panoramic view, with only elementary processing like edge filters, and without the need for explicit landmark recognition.

¹⁰ This can be interpreted as a redundant - and thus more robust - generalization of simple Cartesian encoding (along two axes).

¹¹ The role of vestibular-like inertial sensory information in PI is largely unknown.

¹² Another implementation of a vector memory may be sustained neuronal "phasor type" activity of (unknown) cell types in the CBU, and performing vector arithmetic as in Lyu et al. (2022).

Zeil et al. (2003) showed that similarity gradients based on pixel-by-pixel intensity differences are sufficient for successful visual homing. Combining multiple view memories along frequently traveled routes allows for complex routes following toward the nest. Webb (2019) emphasizes that no information about the location or temporal sequence of the stored views is necessary: Baddeley et al. (2012) proposed a computational familiarity model, which encodes the entire view memory in an InfoMax (Lee et al., 1999; Lulham et al., 2011) neural network architecture. From a scan of the environment, the agent can then infer the most familiar viewing direction over all stored memories. If the view memories are acquired during inbound routes (i.e., linked to a homing motivational state), this will guide the agent toward the nest. Ardin et al. (2016) proposed a biological implementation of a familiarity model based on the insect mushroom body (MB), a learning-associated region of the insect brain discussed in more detail in Section 5. In RL, different neural network approaches to assess view familiarity have been used e.g., in Random Network Distillation (Burda et al., 2018). Instead of a homing signal, it serves as motivation for exploring unknown states (cf. discussion Section 7).

4.4 Discussion: spatial representation in the insect navigation base model

As presented thus far, the base model entails two independent spatial representations: A vector map, which is not linked to specific egocentric views, and directional cues based on view memories, which are not linked to any geometric information from the vector memory (see Figure 1F). According to our previous classification, the vector map is an *explicit spatial representation*: It is evolutionarily pre-determined by the path integration circuitry, just like classical SLAM architecture by a pre-defined inference algorithm. The construction of the vector representation differs from SLAM methods in that exact localization is assumed, based on PI, and the map is constructed based on that ground truth, obviating the need to maintain correspondences between retinotopic and geometric locations of features and landmarks. On the other hand, a crude *latent spatial representation* in terms of directional cues is implicit in the view memory.

For example, the visual features used by the InfoMax architecture for familiarity discrimination are latent in the learned network weights. Figure 1F illustrates how the INBM aligns with our classification of spatial representations used for robot navigation. As mentioned, the base model explicitly does not link view memories to vector memories.

4.5 Beyond the base model

How these two distinct representations are linked is an active research question, covering two major aspects: First, how do insects balance conflicting information from the two systems? Sun et al. (2020) proposed a unified model inspired by joint MB/CX neuroanatomy, combining PI, visual homing, and visual route following. The model balances off-route (PI and visual homing)

with on-route (visual route following) steering outputs based on visual novelty and uncertainty of the PI signal. Goulard et al. (2023) recently proposed a different mechanism to integrate view based and vector based navigation, based on an extended concept of vector memories: Positioning vector computations as the fundamental scaffold for integrating navigational strategies, they exploit the fact that both PI home vector and the current heading direction are represented in ring attractor networks within the CX (c.f. Stone et al., 2017; Le Moël et al., 2019). In addition to *location specific* vector memories [i.e., imprints of the home vector, as in Le Moël et al. (2019)] they introduce *direction specific* memories, which store the current heading direction via inhibition by a second set of “memory neurons.” View familiarity is again assumed to be learned in the MB, producing a binary signal which is used to store the instantaneous heading direction as a vector memory, whenever a familiar view is encountered. View familiarity is thus translated directly into a directional vector memory, which is anchored in an allocentric coordinate frame using the PI based home vector (or other location specific vector memories). This makes the resulting steering output more robust to overshooting and temporary disappearance of the visual cue. In theory, this mechanism could also be used to stabilize directional cues obtained from other, possibly more volatile, sensory modalities such as odor gradients.

Conceptually more interesting is a second aspect: is the view memory truly independent of the geometry of the vector memory? This is closely related to a question not thoroughly addressed in the work cited above: When and where does an animal form a view memory? Most models just assume that views are stored regularly along a homeward-bound route. Ardin et al. (2016) suggest that “the home reinforcement signal could [...] be generated by decreases in home vector length.” Note that this already associates the stored views with a specific node in the vector map. Wystrach (2023) recently proposed a neuroanatomically constrained model for visual homing which obviates the need for storing individual view memories: during learning, views are continuously associated with facing left or right with respect to the nest, using the difference between PI vector and current heading direction for reinforcement. We will discuss this model in detail in Section 6.1. The spatial representation of egocentric views is now decidedly conditioned on a specific vector. One could easily imagine an extension of this model by vector memories, enabling learnable visual guidance along arbitrary vectors encoded in the vector map, essentially using each vector as a motivational state. Note that such a mechanism would be different from ‘reloading’ a PI state from the view memory, although the expected behavior is similar: Insects would be able to recover previously known ‘shortcuts’ based on visual guidance alone. The joint spatial representation would be a topological latent representation similar to Shah et al. (2023), see the discussion in Section 7.

4.6 An evolutionary perspective: insect inspired RL as a normative model

Conceptually, such a unified spatial representation could itself be viewed as a single, latent embedding of visual input,

learned over evolutionary time to best adapt to ecological constraints, i.e., reap the largest long-term reward from the environment. The dichotomy between static, explicit versus plastic, latent components of the representation would then be relaxed to a continuum of plasticity for different model components, realized via differential learning rates. We propose to design an end-to-end RL-learnable navigation model constrained by the insect navigation base model, in the sense that the resulting spatial representation is compatible with its basic assumptions. This will be instructive for both the field of insect and robot navigation: For the former, it can serve as a normative model for a possible unified spatial representation that goes beyond the base model, providing theoretical guidance for how vector and view-based representations may interact to support efficient navigation.

On the level of sensory processing, having the network learn representations that match, e.g., PI based vector maps may yield valuable insights into which visual features are useful in intermediate processing steps to reliably support such computations in a variety of visual conditions and environments. These model predictions would yield testable hypotheses for further neuroanatomical, physiological, and behavioral experiments, as discussed in Section 7. Given the superior performance of insect navigators in terms of training efficiency, robustness, and generalization capability, robot navigation may profit from this biologically inspired and constrained spatial representation. Pre-training such a network extensively under varying conditions and then freezing the slow components may yield a highly robust, adaptive spatial representation for applications similar to natural insect task spaces, e.g., visual outdoor navigation for ground or aerial autonomous agents. Implementing network architectures that support phasor representations may be a useful avenue for robotic navigation research.

5 Reinforcement learning with an insect brain

Given the success of reinforcement learning as a framework for robot navigation, it seems reasonable to ask if and how navigation could be implemented based on actual RL-type computations in the insect brain. Furthermore, extensive literature involving dopamine, learning, and reward prediction errors exists in the mammalian neuroscience community but despite these topics being relevant in insect learning and navigation, the discussion of potential connections is limited. To explore this line of thought, we will first continue our formal treatment of RL (Section 5.1). In Section 5.2 we will discuss how current computational models of the MB - the prominent learning associated region of the insect brain - could be augmented to support simple RL algorithms. This will allow us to discuss the recent MB/CX based visual homing model by Wystrach (2023) in the context of RL and extrapolate it to roughly outline a neuroanatomically inspired end-to-end RL model for insect navigation in Section 6.

5.1 RL formalism

Starting from the definitions from Section 2.2, RL methods find the optimal policy (Equation 3) which may maximize the expectation of the temporally discounted sum of instantaneous rewards (Equation 2) over time:

$$V^\pi(s) = \mathbb{E}_\pi \left[\sum_{k=0}^T \gamma^k R_{t+k+1} \mid S_t = s \right] \quad (4)$$

$$= \mathbb{E}_\pi [R_{t+1} + \gamma V^\pi(S_{t+1}) \mid S_t = s] \quad (5)$$

$$= \sum_a \pi(a|s) \left[R(s, a) + \gamma \sum_{s'} p(s'|s, a) V^\pi(s') \right] \quad (6)$$

for all initial states s . This function is therefore called the *value function* of s under policy π , with a temporal discount factor $\gamma \in [0, 1]$. Equations 5, 6, versions of the *Bellman equation*, illustrate the recursive nature of the value function: it can be decomposed into the average immediate reward from the current state s under policy π , plus the discounted value of the subsequent state, averaged over all possible successor states s' . Maximization of V with respect to π can now be understood intuitively: For a single time step, the optimal policy π^* for the recursion (Equation 6) would be simply choosing the action which maximizes the term in square brackets. Iterating through the recursion then leads to the *Bellman optimality equation* for the optimal state value function

$$V^*(s) = \max_a \left[R(s, a) + \gamma \sum_{s'} p(s'|s, a) V^*(s') \right] \quad (7)$$

$$= \max_a \mathbb{E} [R_{t+1} + \gamma V^*(S_{t+1}) \mid S_t = s, A_t = a] \quad (8)$$

Another way to interpret Equation 6 would be as a policy average over a *state-action value function* $Q^\pi(s, a)$:

$$V^\pi(s) \equiv \sum_a \pi(a|s) Q^\pi(s, a). \quad (9)$$

By the same logic, recursive (optimality) relations can be derived for Q :

$$Q^\pi(s, a) = \mathbb{E}_\pi [R_{t+1} + \gamma Q^\pi(S_{t+1}, A_{t+1}) \mid S_t = s, A_t = a] \quad (10)$$

$$Q^*(s, a) = \mathbb{E} \left[R_{t+1} + \gamma \max_{a'} Q^*(S_{t+1}, a') \mid S_t = s, A_t = a \right] \quad (11)$$

5.1.1 Value-based, policy-based, and actor-critic methods

One way to find the optimal policy π^* is by trying to solve it directly. This is most commonly done using policy gradient methods which parameterize the policy $\pi(a|s; \theta)$ and then perform gradient ascent on a suitable performance metric, like the average reward per timestep: $\Delta\theta \propto \nabla_\theta \mathbb{E}[R_t]$.

We will not dwell on pure policy-based methods further, for more detail see Sutton et al. (1999) and Williams (1992). Alternatively, one can estimate the value function $V(s)$ or $Q(s, a)$ and infer an optimal policy indirectly. For example, assuming an optimal Q^* is found, the optimal policy is simply

$$\pi^* = \operatorname{argmax}_a Q^*(s, a) \quad (12)$$

These value-based approaches have the strongest connection to insect neuroscience and will therefore feature prominently in the rest of this paper.

Last, so-called actor-critic methods constitute a hybrid approach, involving both a policy (actor) and value (critic) estimation. The policy gradient is then computed to maximize the *advantage* of an action derived from the policy over the estimated baseline value. Actor-critic methods play an important part in robot navigation.

5.1.2 Temporal difference methods: SARSA and Q-learning

TD methods approach the problem of value estimation by deriving single-timestep update rules from the recursive relations (Equations 10, 11) for Q or Equations 5, 8 for V : for each timestep, the squared difference between the LHS and RHS is treated as a prediction error to be minimized. For Equation 10, this leads to the update rule

$$Q^{\pi}(S_t, A_t) \leftarrow Q^{\pi}(S_t, A_t) + \alpha \cdot \Delta_t, \quad (13a)$$

$$\Delta_t = [R_{t+1} + \gamma Q^{\pi}(S_{t+1}, A_{t+1}) - Q^{\pi}(S_t, A_t)] \quad (13b)$$

with learning rate $\alpha \in [0, 1]$. This update rule is known as SARSA due to the tuple $(S_t, A_t, R_{t+1}, S_{t+1}, A_{t+1})$ required to compute the update. In order to accommodate exploration, the agent follows a non-deterministic policy based on the current estimate of Q , making SARSA an *on-policy* algorithm. The most common choices are ϵ -greedy (choose the $\text{argmax}_a(Q)$ with probability $1 - \epsilon$, random action with probability ϵ) and softmax policies (choose an action from a Boltzmann distribution based on Q , with inverse temperature β). These policies converge to the optimal policy π^* , if the stochasticity is reduced systematically toward a deterministic (greedy) policy ($\epsilon \rightarrow 0$ or $\beta \rightarrow \infty$) during learning (see Sutton and Barto, 2018).

Finally, a TD error without reference to a specific policy can be derived to estimate Q^* directly from Equation 11:

$$\Delta_t = [R_{t+1} + \gamma \max_a Q^*(S_{t+1}, a) - Q^*(S_t, A_t)] \quad (14)$$

This is known as Q-learning. Note that in contrast to Equation 13b, this *off-policy* update is now independent of the consecutive action A_{t+1} prescribed by the policy. However, the policy still determines which action-value pair receives the update. This effectively decouples the learned policy from the policy employed during learning. In particular, this also allows for the agent to perform *off-line* updates, i.e. updates which are not based on the current state transition, but e.g., sampled from a replay buffer of previous experiences. This property was key to the success of the Deep Q-Network (DQN) by Mnih et al. (2015), an early milestone of Deep Reinforcement Learning, which has since found numerous applications in robot navigation.

As we will show, the neural substrate of the insect mushroom body has the potential to support Q-based TD computations like SARSA or Q-learning to solve navigational tasks. Conversely, (Deep) Reinforcement Learning as a key toolset for robot navigation provides a useful framework to think about the neural computations underlying insect navigation.

5.2 The mushroom bodies as a neural substrate for RL

5.2.1 The canonical MB learning model for classical conditioning

The mushroom bodies are bilateral neuropils in the insect brain, with homologous structures largely conserved across different species, whose crucial role in learning and memory has long been established. Extensively studied in the context of olfactory learning in *Drosophila melanogaster* (reviewed in Cognigni et al., 2018), inputs from other sensory modalities, in particular vision (e.g., Ehmer and Gronenberg, 2002; Strube-Bloss and Rössler, 2018; Vogt et al., 2014) likely support general behavioral learning tasks, including visual-spatial navigation. The main intrinsic anatomical components of the MB are *Kenyon cells* (KC), whose dendrites form the calyx, while the axons constitute the lobes of the MB. They receive sensory inputs via *projection neurons* (PN), which are thought to form non-plastic, sparse and random (Caron et al., 2013) synapses onto the KCs at the calyx. KC activity is transmitted to *mushroom body output neurons* (MBONs) at the MB lobes. *Dopaminergic neurons* (DANs), which also target the MB lobes, induce (usually depressive) modulation of the KC-MBON synapses and thus enable an adaptive response to the sensory stimulus. MBON activity is integrated downstream by (pre)motor neurons (MN) to produce an action. Conventionally, the reinforcement signal mediated by the DANs is assumed to encode a direct extrinsic reward. Figure 3A shows the “canonical” MB circuitry (without the dashed MBON \rightarrow DAN synapses).

5.2.1.1 Prediction targets and reinforcement signals: RL vs. associative learning

At first glance, all the ingredients for reinforcement learning seem to be there: KC activity defines the state space \mathcal{S} , based on which MBON activity encodes *some* value prediction over an action space \mathcal{A} . DAN activity encodes a reward function $R(s, a)$. However, there is a crucial difference in the *kind of value prediction* which is computed: So far, MB-based learning has been studied in the context of trial-by-trial associative learning (AL) paradigms like classical conditioning, where the agent is presented with isolated stimulus-reward pairs. This is *not* a full MDP, since neither future states nor rewards are contingent upon the current state and action, i.e., the transition function $P(s'|s, a)$ is not specified. The prediction target of RL, the long-term cumulative reward, is therefore ill-defined, since it is determined by the experimenter's choice of stimulus/reward pairs, and not (only) by the agent's action. Instead, the agent is trained to predict the immediate external reward following an action. This highlights a fundamental difference in the interpretation of rewards in AL vs. reinforcement learning: In the former, it serves as an immediate feedback signal used to evaluate individual actions. The agent does not learn to maximize future rewards, but merely to react to a stimulus according to the associated reward. In the latter, rewards define a long-term objective which the agent learns to achieve by a series of optimal actions.

Furthermore, the models differ in *how the value prediction is learned*: in the canonical MB model, the DAN reinforcement signal *directly* encodes the absolute value of the external reward R (direct

reinforcement), while for TD-RL methods, Δ_t in Equations 13, 14 encode a prediction error of $Q(s, a)$ (which is a proxy for the prediction target $V(s)$). Crucially, an MDP cannot be learned using direct reinforcement, since there is no directly provided ground truth for the prediction target. A neural mechanism for computing prediction errors is therefore a prerequisite to reconcile RL with MB based computations.

5.2.2 Prediction errors in the MB

Recent computational studies proposed rate-based (Bennett et al., 2021) and spiking (Jürgensen et al., 2024) models of the MB which employ recurrent MBON→DAN connections to compute a reward prediction error (RPE). They show that the resulting behavior of the agents in a classic conditioning paradigm aligns with experimental evidence. The postulated recurrent connections are supported by anatomical evidence (Bennett et al., 2021) and are indicated by the dashed arrows in Figure 3A. It illustrates the simplest iteration of the models investigated by Bennett et al. (2021): Agent behavior is the net result of approach and avoid opponent processes within the MB. The two antagonistic behaviors are encoded by the activity of two distinct sets of MBONs (M_{\pm}), whose outputs are integrated by downstream descending neurons. KC→MBON synapses of these distinct sets are targeted by two sets of valence specific, i.e., appetitive and aversive, DANs (D_{\pm}). The major innovation of this model lies in the interpretation of M_{\pm} activity as predictions \tilde{R}_{\pm} of the positive and negative external reward R_{\pm} , respectively.¹³ The prediction error ΔR of the total reward $R = R_+ - R_-$ is then computed indirectly by recurrent excitatory MBON → DAN connections of opposite valence, i.e., the prediction of negative reward is added to the direct positive reward, and vice versa: $D_{\pm} \leftrightarrow R_{\pm} + \tilde{R}_{\mp}$. The difference between D_+ and D_- activity then encodes the full RPE:

$$\begin{aligned} D &= D_+ - D_- \leftrightarrow (R_+ + \tilde{R}_-) - (R_- + \tilde{R}_+) = \Delta R_+ - \Delta R_- \\ &= \Delta R \end{aligned} \quad (15)$$

Since the synaptic modulation by DANs is depressive, reinforcement is achieved by inhibiting MBONs of the opposite valence, i.e., appetitive DANs inhibit aversive MBONs and vice versa. In contrast to the temporal difference errors from SARSA (Equation 13) and Q-learning (Equation 14), here the RPE reflects the prediction error of the single-timestep (or “timeless”) total external reward of the Rescorla-Wager type (Rescorla, 1972). We propose that only a few additional assumption can turn this MB circuit into a neural implementation of a TD RL agent.

5.2.3 Temporal dynamics of recurrent connections in the MB can support temporal difference learning

Once we change the experimental paradigm to an interaction task that can be modeled as a full MDP, optimizing cumulative rewards becomes a meaningful objective. The agent’s learning objective is now no longer the immediate reward $R(s, a)$ following

an action, but the long-term value of a state-action pair $Q(s, a)$. The recursive TD update rules (Equations 13, 14) are prediction errors of the current estimate of $Q(s, a)$, with a more complex prediction target: $R(s, a)$ plus the agent’s estimate for $Q(s', a')$ in the next timestep. This temporal link is the core principle allowing TD methods to compute estimates of cumulative rewards over time. As before, we can interpret the activity of distinct sets of MBONs as predictions, now for Q instead of R .

To illustrate the key differences to the RPE model discussed above, let us consider the simple case of SARSA learning on a discrete n -dimensional action space $\{a_i\}$ beyond binary approach/avoid behavior—representing for example the choice of a navigational sub-goal (“vector memory,” see Section 6.2). Unlike in standard “computational” SARSA, positive and negative rewards in the MB model are encoded in valence-specific pathways, leading to a valence-specific decomposition $Q = Q_+ - Q_-$, represented by the activity of two distinct sets of n MBONs: $M_{\pm}^{(i)} \leftrightarrow Q_{\pm}(s, a_i)$. In order to implement the SARSA update rule via dopaminergic modulation, the following differences to the RPE model are crucial: (i) To convey estimates of Q at subsequent timesteps, recurrent MBON→DAN connections would have to carry activity traces, for example via multiple pathways with different temporal dynamics,¹⁴ synaptic strength and potentially different effective signs. Eschbach et al. (2020) have demonstrated the existence of these different kinds of pathways in larval *Drosophila* in the context of classical conditioning. It would be interesting to investigate whether they converge in a way that enables the computation of TD errors. (ii) The TD-error (Equation 13b) only involves predictions of $Q(\bar{s}, \bar{a})$ corresponding to the actually experienced state-action pair (\bar{s}, \bar{a}) , chosen according to a value-based (probabilistic) policy: $\bar{a} \sim \pi(Q(\bar{s}, a_i))$ (see Section 5.1). Since MBON population activity $\{M_{\pm}^{(i)}\}$ encodes the entire function $Q(\bar{s}, a_i)$, this would require selective MBON→DAN feedback, reflecting the implementation of π . (iii) Finally, synaptic modulation must selectively affect only KC→MBON synapses corresponding to $Q(\bar{s}, \bar{a})$, according to Equation 13a.

Action selection and policy implementation The simplest way to implement a policy-dependent action selection is directly at the level of MBON activity. For example, a winner-take-all type lateral inhibition mechanism between MBONs would correspond to a (greedy) argmax policy (cf. Equation 12). Accounting for noise in the circuit, this could be interpreted as a non-deterministic policy,¹⁵ potentially giving rise to an ϵ -greedy or softmax-like selection mechanism. As a result, only the ‘conjugate’ MBON pair \tilde{M}_{\pm} would be active, reflecting the value $\tilde{Q}_{\pm} \equiv Q_{\pm}(\bar{s}, \bar{a})$ of the chosen (experienced) action in the given state. This ensures that unspecific recurrent MBON→DAN connections would still convey only the

¹⁴ Technically, Equation 13 looks forward in time, while an activity trace only gives access to previous states. This could be reconciled by simply shifting the time index such that $Q_t(S_{t-1}, A_{t-1})$ is updated at time t . In the neural model, this could be captured by the temporal dynamics of synaptic modulation.

¹⁵ This could even account for adaptive exploration, with an increasingly exploitative policy as the signal-to-noise ratio increases, i.e. the agent becomes more evaluation of a state-action pair becomes more informative (larger in absolute value).

¹³ Somewhat confusingly, rewards of opposite valence are often termed “reward” and “punishment” in the AL literature.

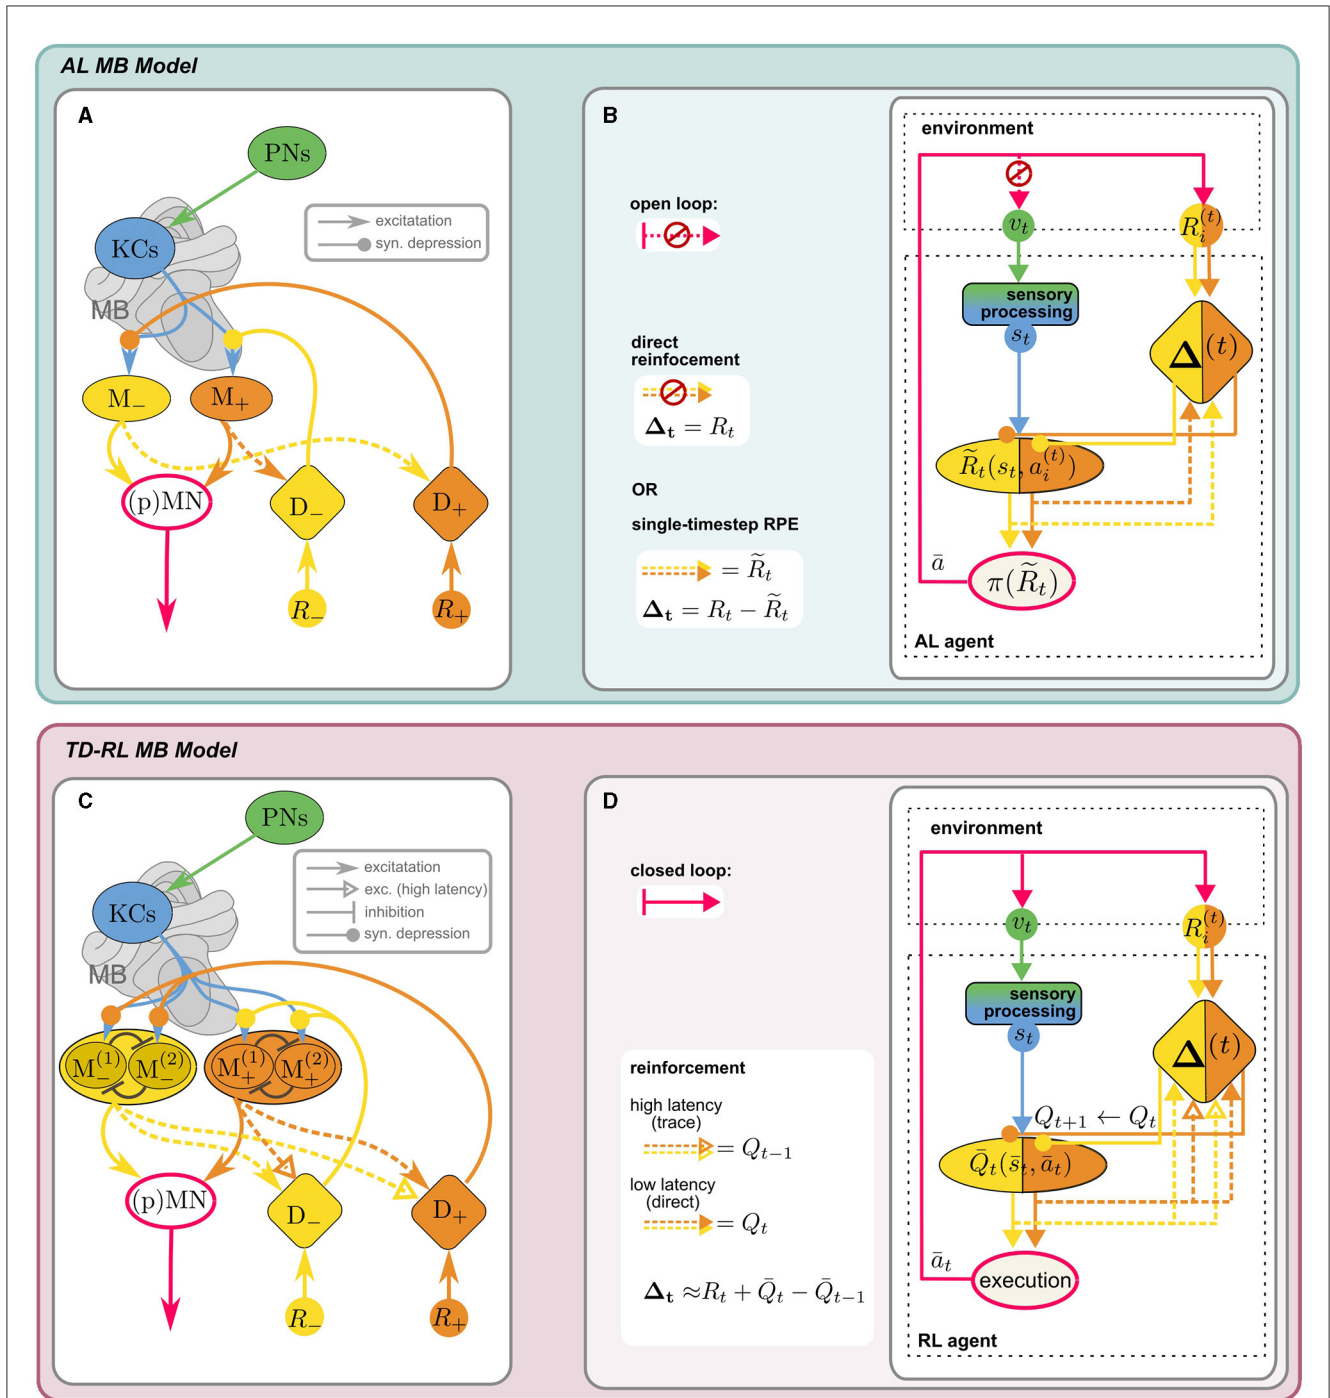


FIGURE 3

Anatomical [(A), adapted from Bennett et al., 2021 licensed under CC-BY 4.0] and computational (B) components of MB model for AL: KCs receive sensory input v_t from the PN, encoding a decorrelated and sparse representation of the sensory environment s_t . KCs synapse onto distinct sets of approach/avoid MBONs (M_{\pm}) driving opposing responses, which are integrated in downstream (pre)motor neurons, (p)MN, to produce an effective action, according to a “policy” π . MBON activity can be interpreted as a prediction \tilde{R} of external reward R , stored in KC→MBON synaptic weights. These are depressively modulated by reinforcement signals Δ_{\pm} from distinct sets of aversive/appetitive (D_{\pm}) DANs of opposite valence. In the canonical model, DANs encode direct (external) positive or negative reinforcement $\Delta = R$. Including recurrent MBON → DAN connections (dashed) enables DANs to compute an RPE signal for reinforcement: $\Delta = R - \tilde{R}$ (see Section 5.2). (C, D) MB as a neural substrate for RL (Section 5.2.3): When the state-action loop is closed, s_t are valid states of an MDP. A discrete action space (here: $n = 2$) is represented by corresponding sets of MBONs (in “conjugate” valence pairs). A value-based policy could be implemented by lateral inhibition between MBONs (schematic illustration only). The MB circuit can support TD-RL with the additional assumption that high-latency recurrent MBON→DAN connections can carry an activity trace. DAN activity can then encode a TD error (Equation 16) and MBON activity reflects the agent’s current estimate of \bar{Q} , the value prediction for the current state-action pair (\bar{s}_t, \bar{a}_t).

relevant Q-prediction such that joint DAN activity represents the TD error (Equation 13b): Extrapolating from the RPE computation (Equation 15), this could be implemented by assuming *low-latency* (t), *excitatory* recurrent connections between DANs and MBONs of *equal valence* and *high-latency* (t'), *excitatory* connections between *opposite valence* compartments:

$$\begin{aligned} D &= D_+ - D_- \leftrightarrow [R_+(t) + \bar{Q}_+(t) + \bar{Q}_-(t')] - [+ \leftrightarrow -] \\ &= [R_+(t) + \bar{Q}_+(t) - \bar{Q}_+(t')] - [+ \leftrightarrow -] \\ &= \Delta_{t'}^+ - \Delta_{t'}^- = \Delta_{t'} \end{aligned} \quad (16)$$

If we additionally assume a DAN gated 3-factor plasticity rule at the KC→MBON synapse, depending on both presynaptic KC and postsynaptic MBON activity, updates to the value prediction would apply selectively only to $Q(\bar{s}, \bar{a})$, see (ii) above. Both lateral MBON inhibition (e.g., Huerta et al., 2004) and effective 3-factor plasticity rules (e.g., Huerta and Nowotny, 2009; Faghihi et al., 2017) have been employed in MB learning models, summarized nicely in a recent review by Webb (2024). Other combinations of mechanisms presented therein may be equally viable to effectively perform RL with a MB inspired circuit. Alternatively, action selection may take place downstream of the MB, for example in the CX which has been suggested as a potential substrate for action selection in *Drosophila* (Hulse et al., 2021). The recurrent MBON→DAN connections described by Eschbach et al. (2020) also include multi-synaptic pathways, which would be necessary for such a more indirect action-selection and reinforcement mechanism. In any case, the agent's current estimate of the *function* Q is fully captured by the KC→MBON synaptic weights, learned by synaptic modulation through the DANs. Figure 3 illustrates how a full TD model can be obtained by augmenting the canonical classical conditioning model of the MB, and how anatomical and algorithmic components map onto each other. If rewards are given only externally at sparse locations, successful learning in such an RL agent will require training over many episodes.

Note that while we will continue to discuss the proposed circuit in the context of navigation, it may well be applicable to other MB based learning functions. For example, the spiking MB model for motor skill learning in Arena et al. (2017) interprets MBON output as modulation to the parameters of a central pattern generator. Similar to our model, learning is achieved by optimizing cumulative rewards (c.f. Equation 4), essentially optimizing a MDP. However, the learning rule is based on direct comparison of cumulative reward between episodes. It would be interesting to modify this model to instead implement a continuous, value-based TD learning rule (Equations 13b, 14) via the recurrent connections postulated above.

5.2.4 Experimental assessment of MB based RL

So far, we have treated the premise of RL—the agent trying to optimize cumulative rewards in a MDP over time—as given and asked how the MB architecture could provide the neural substrate for the necessary computations. However, experimental validation of the underlying assumption, as well as conjectures about possible neural implementations, requires a paradigm that goes beyond simple associative learning (which has been the

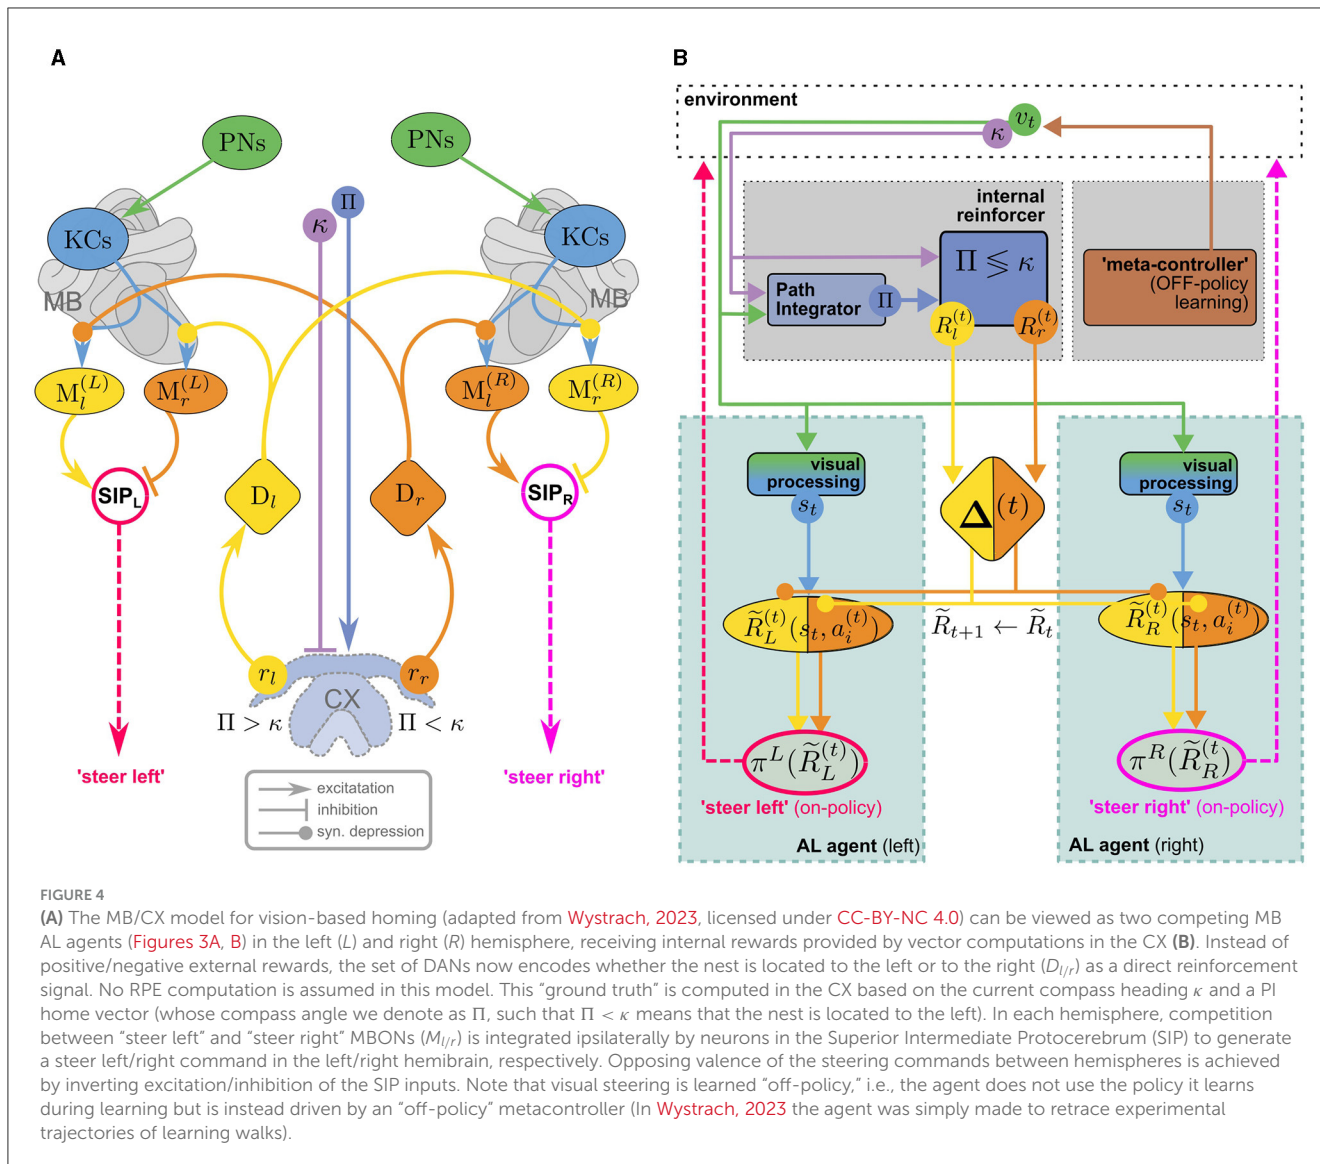
dominant approach in MB learning experiments, e.g., Liu et al., 2012; Aso et al., 2014; Felsenberg et al., 2017) and is instead capable of capturing the full action-state feedback (transition function) of a MDP. Ideally, the experimental setup would grant direct control over states, transition and reward functions. For example, one could imagine a closed-loop virtual reality setup with a tethered insect on a treadmill, similar to the setup used by Lafon et al. (2021) or Geng et al. (2022). The simplest way to encode (purely visual) states would be as uniform visual input of different color/intensity/polarization which changes along predefined gradients (the transition function) according to the walking direction of the animal.¹⁶ Reward could be delivered optogenetically at specific locations in this artificial state space. This setup would allow a direct comparison between observed and modeled behavior (also using different RL algorithms). Computational techniques like inverse reinforcement learning (Kalweit et al., 2020, 2022) would allow a direct comparison of the inferred Q-function of the agent with the one learned by the model. In combination with electrophysiology or calcium imaging, DAN activity in the behaving and learning animal could be compared to the TD error of SARSA or Q-learning models. In particular, one would expect increased DAN activity in states close to (but distinct from) a change in the reward landscape, which could not be explained by associative learning mechanisms alone. Finally, assuming that a healthy animal exhibits behavior consistent with a RL model, targeted disruption of MBON→DAN connectivity could be used to probe whether the proposed connections are indeed necessary for the computations underlying this behavior.

6 Reinforcement learning for insect navigation

6.1 Visual homing with vector-based internal rewards

Wystrach (2023) proposed an MB/CX based visual homing model that uses an internal reward signal to alleviate this problem. Figure 4A illustrates how the model consists of two antagonistic copies (one in each hemisphere) of the “canonical” MB circuit in Figure 3 which receive internal reward signals computed in the CX: comparing the agent's current compass heading κ and PI home vector (with compass angle P_i) by a mechanism similar to the one described by Lyu et al. (2022), a set of two reward signals $r_{l/r}$ is computed, encoding whether the nest is located to the left/right of the agent's current heading direction. These provide input for two sets of dopaminergic neurons $D_{l/r}$, assuming the role of external rewards in the canonical model (Figure 3A). However, they don't encode a rewarding experience coming from the environment, like the agent reaching a food source, but relate to an internal state of the agent, the home vector. Since the latter is continually updated, the rewards are no longer sparse which makes learning considerably more efficient. DANs convey copies of the respective reward signal to MBs in both hemispheres, giving rise to double opponent processes: In each hemisphere, $r_{r/l}$ serve as direct reinforcement

¹⁶ More naturalistic visual state spaces could be used, as long as they are easily parameterized along the action dimensions.



to associate representations of the current view (encoded in KC activity) with populations of MBONs, $M_{l/r}$, corresponding to “steer left/steer right” responses respectively. In either hemisphere, $M_{l/r}$ activity is integrated by downstream neurons, but with inverted signs, leading to competing “steer left” and “steer right” premotor commands from the left/right hemisphere, respectively.¹⁷ This double opponent architecture increases performance robustness as there are now two sets of MBONs that independently encode the appropriate behavioral response.

Evidently, the model implements an associative learning algorithm: It does not involve computation of a TD error as proposed above, or even an RPE for the internal reward. While it may be interesting to extend the model to use RPEs for reinforcement, a computation of TD errors would serve to achieve an erroneous objective: The formulation of the underlying

MDP would imply that the agent’s goal is to maximize the cumulative *internal* reward. Since the internal reward is higher (e.g., for steering right) the further off-target the agent is heading, maximizing it over time would lead to the opposite of the desired behavior. It would be interesting to investigate if a model with inverted reward valences could be extended to an RL model for visual homing. In the following, however, we will explore a different line of thought, sketching an MB/CX inspired RL model that integrates all three components of the insect navigation base model from Section 4—path integration, vector memories, and view memories—and links them to the behavioral objective of optimizing external rewards. Finally, note that—in a liberal interpretation of RL terminology—we can classify the visual homing model as an *off-policy* learning algorithm: During the learning phase, the agent’s actions are assumed to follow an exploration strategy in agreement with observations (using data from Jayatilaka et al., 2018; Wystrach et al., 2014) of learning flights/walks performed by insects after emerging from their nest for the first time (see, for a recent review).

17 The implementation of the actual steering mechanism is again located in the CX, but we will not discuss this further here. For more detail refer to the original paper (Wystrach, 2023).




FIGURE 5

The visual homing model of Figure 4 can conceptually be turned into a hierarchical MB/CX based RL full navigation model which roughly fits to the insect navigation base model. This can be done by replacing the off-policy controller in Figure 4B with an MB RL agent as an “RL meta-controller” (Figure 3D). As one additional component, it relies on a store of vector memories m , “snapshots” of the PI home vector Π , learned via direct associative reinforcement with an external reward at salient locations (see Section 4 and Section 6.2 for possible anatomical implementations) which now form the (discrete) action space of the TD-RL agent. A specific vector memory $\bar{m} \sim \pi(m|s)$ is selected as a navigational subgoal based on a policy learned from the same external reward as the vector memories on a state space of processed visual input s_t . The internal reward now encodes whether the agent is facing left/right with respect to the relative vector to the currently active vector memory selected by the meta-controller. This vector computation based on current Π , current heading κ , and active m is assumed to be performed in the CX. The rest of the visual homing circuit is identical to Figure 4. Note that s_t and Π could be viewed as components of a joint representation of the visual input, with different components of the algorithm acting only on subspaces of this representation.

6.2 Toward an MB/CX based RL model for (Insect) navigation

The visual homing model discussed above will serve as a representative of models for view memory. It obviates the need to decide when to store a view memory by learning continuous

associations, as described above. It makes a very explicit connection between view memories and the home vector, while other models only implicitly associate view memories with a motivational state (Webb, 2019). However, taking into account the proposed mechanism for storing and recalling vector memories (Le Moëll et al., 2019), one could broadly interpret the presently loaded vector

memory as a motivational state (feeder vector loaded = “foraging motivation,” no vector loaded = “homing motivation”). Loading the vector memory of a specific location effectively replaces the home vector in the visual homing model with a relative vector toward that location, theoretically allowing the agent to associatively learn visual homing relative to *every location in the vector memory*, using internal rewards. Vector memories on the other hand, are learned by one-shot associative learning from external reward, e.g., the presence of food. As suggested by Le Moël et al. (2019), this could be achieved by direct dopaminergic modulation of synapses between “vector-memory neurons” and the CPU4 integrator neurons in the CX. Alternatively, specific populations of neurons, each of which conveys a specific vector in a phasor-like representation, might be activated. If these hypothetical “vector-memory neurons” are, or receive excitatory input from, a subpopulation of MBONs, the pool of vector memories could in turn serve as the action space for an MB-based implementation of a TD-RL algorithm discussed above (Figure 4B). This RL meta-controller learns a policy over sub-goals from the current view and *external* reward, which may serve two purposes: Steering the agent toward the selected sub-goal using PI, and providing sub-goal-directed *internal* reward for the low-level AL steering/homing controller. Interestingly, however, due to the off-policy architecture of the AL controller, view association with respect to any (inactive) vector memory could be learned while homing toward another (active) one (using either views or PI), assuming a dedicated set of visual homing MBONs associated with each vector memory. This would be consistent with our interpretation of vector memories as motivational states, which are often modeled by distinct MBON populations in the MB literature [see Webb and Wystrach (2016)]. Theoretically, this enables the agent to continually learn view associations with respect to all vectors stored in memory while using a specific one, or PI, for steering. If we further assume an anatomical link between “vector memory MBONs” and their corresponding set of “visual homing MBONs,” the RL-meta controller could select previously visited vector locations as sub-goals for visual homing. Figure 5 illustrates how the proposed model is built from the components discussed above. We argue that a unified MB/CX RL model along these lines would give rise to a remarkably versatile spatial representation for navigation, while being relatively memory efficient: Each memorized location would require at most three MBONs per hemisphere to encode it (one for selection plus two for steering), which would be reasonably within the capacity of e.g., honeybees with ~200 MBONs per hemisphere (Rybak and Menzel, 1993; Mobbs and Young, 1997). This computational efficiency, rooted in the combination of an innate vector representation with an adaptive learning and optimization mechanism, could serve as a blueprint for robotic applications where compute power and energy are severely limited, like autonomously navigating aerial vehicles.

7 Discussion

In this paper, we presented reinforcement learning as a common framework to compare representations of space (Section 2) in current models for robot (Section 3) and insect navigation (Section 4). We argue that the insect navigation based

model combines an *explicit* vector map with a *latent* spatial representation based on view memories. This combination of explicit and latent spatial representations, evidently very successful for insects, may provide inspiration for future approaches for robot navigation. We further proposed reinforcement learning as a novel framework for interpreting and expanding existing models for insect navigation. In Section 5, we investigated how existing models of the insect MB circuit could implement temporal-difference (TD) RL algorithms, and proposed a circuit model with specific recurrent MBON→DAN connections which would support TD-like reinforcement at the MBON→KC synapse via differential temporal dynamics (Section 5.2.3) In Section 6, we sketched how this could be used to build a full MB/CX inspired TD-RL navigation model, which has the potential to link the disjoint spatial representations of the insect navigation-based model. To conclude, we will discuss aspects of such a hypothetical model with respect to spatial representations and RL-based robot navigation models more generally.

7.1 Task hierarchies, intrinsic rewards, and exploration

The model we outline here is, in a sense, a hierarchical RL (HRL) model: A metacontroller selects a sub-goal for a low-level controller, to provide the latter with dense internal rewards instead of operating on the sparse external reward landscape. For similar reasons, hierarchical RL architectures also play an important role in robot navigation (Schmalstieg et al., 2023; Nachum et al., 2018; Haarnoja et al., 2018) to avoid “dimensional disaster.” There are, however, crucial differences: In RL, the meta-controller usually provides *intrinsic* rewards for the low-level agent, designed to facilitate the exploration of complex environments with sparse external rewards. Oudeyer and Kaplan (2008) define a situation as “*intrinsically motivating* [...] if its interest depends primarily on the collation or comparison of information from different stimuli and independently of their semantics, [...] understood in an information theoretic perspective, in which what is considered is the intrinsic mathematical structure of the values of stimuli, independently of their meaning”, as opposed to extrinsic motivation. Intrinsic rewards can be viewed as a formalization of curiosity, motivating the agent to explore unfamiliar terrain, and various approaches to model it exist in the RL literature (Pathak et al., 2017; Burda et al., 2018; Savinov et al., 2019). In the proposed model, the internal reward – facing left/right of a location related to external rewards – is inherently *extrinsic* and does not motivate the agent to explore, and the low-level controller is not an RL agent. However, there are biologically plausible ways to include intrinsic reward in the model. For example, Sun et al. (2020) proposes a mechanism for switching between on-route and off-route navigation strategies based on visual novelty. Such a signal could theoretically serve as an intrinsic reward in our model.

Finally, it is important to note that so far we have treated vision in isolation. While it is clear that moving to multi-sensory state spaces is necessary for a realistic model of insect navigation,

including other sensory modalities—especially olfaction—also has great potential to alleviate the problem of sparse rewards in our model. It may be very useful to consider the “map” defined by spatially varying odor concentrations (c.f. spatial olfactory hypothesis in [Jacobs, 2012](#)) as a potential source for a spatially and temporally more dense reward signal. The simplest and least realistic implementation would be using large-scale gradients of rewarding odors as a reward signal. For more realistic odor landscapes, e.g., odor filaments caused by turbulent flow, the frequency/probability of binary odor detections (encoded for example by temporally sparse firing of tuned MBONs, as in [Rapp and Nawrot, 2020](#)) could be interpreted as a proxy for odor density. [Matheson et al. \(2022\)](#) further demonstrated the existence of an odor-gated representation of the wind direction in the fanshaped body. Interpreting this as a rough estimate of the direction of a rewarding location, its deviation from the current heading could be used as a direct reward signal via recurrent connections from CX to MB. Such a mechanism could contribute both to increased learning performance due to denser rewards, as well as improve exploration by introducing an inherently stochastic element to the reward landscape.

7.2 A unified spatial representation

The proposed CX/MB inspired RL navigation model would fuse the disjoint spatial representations of the insect navigation base model – an explicit vector map and latent directional cues ([Figure 1](#)) – into a unified latent graph-like representation which could be used in the absence of accurate PI information (while preserving the vector map for pure PI-based navigation). The agent learns to select a high-level subgoal—a vector memory—based on the current view, from the high-level objective of external reward optimization using an MB/CX RL circuit. This association defines a latent topological relation between the location of the current view and nodes represented by the vector memories: the learned policy will likely reflect distance since learning to visit distant goals would be less rewarding in the long run. The low-level MB/CX AL homing circuit learns rough directional cues to steer the agent toward the selected goal. Although both directional and distance information are therefore available to the agent, both are indirect and inaccurate, making the representation topological rather than geometric.

In light of the cognitive map debate, we can characterize an agent with such a spatial representation as “*knowing where to go, on different spatial scales*”: It can infer from the current view both where it wants to go, i.e., which vector memory to load, and how to get there. This would enable behavior commonly associated with a cognitive map. For example, the agent could visually infer the direction of a novel shortcut from the present location A to a previously visited location B stored in the vector memory, since view memories in the vicinity of A have also been associated with left/right steering commands with respect to B. By a change in the policy over vector memories, for example, because a previously closer food source C has been depleted, the agent could be incentivized to choose to steer along a novel route (“shortcut”) toward B.

7.3 Predictions, memory traces, and internal world models

The previous example shows that our model could support surprisingly flexible navigation, adapting behavior locally depending on distant changes in the environment. In RL, this kind of flexibility is traditionally associated with model-based algorithms: Unlike the model-free algorithms discussed above, they learn a model of the environment – i.e., the transition ([Equation 1](#)) and reward functions ([Equation 2](#)) – directly and then infer the optimal policy from them via ([Equation 5](#)). Explicit knowledge about the environment allows them to plan out actions virtually in order to optimize the policy (see [Hafner et al., 2022](#) for a current application using Deep Hierarchical Planning). This is particularly useful to adapt flexibly to changes in the environment, like changes in reward magnitudes (e.g., an empty feeder) or new navigational obstacles. Model-based algorithms don’t need to slowly and incrementally update their state evaluation by repeated exposure to the change in the environment, but can simply update their model of the environment and virtually plan a new route based on that update. A neural implementation of such a *predictive* model of the environment presupposes a state representation in the form of recurrent neural activity, like the (p)replay of spatial sequences observed in the mammalian hippocampus ([Dragoi and Tonegawa, 2011](#)). The likely insect analog, spontaneous KC activity in the MB, would be challenging to access experimentally and has not yet been observed. On the modeling side, [Wei et al. \(2024\)](#) recently proposed the intriguing idea that adaptive axo-axonal gap junctions between KCs in the MB calyx could represent the state transitions of a learned transition function and thus support model-based RL computations in the MB.

On the other hand, a suitably predictive (latent) representation can produce similar flexible behavior. E.g., [Russek et al. \(2017\)](#) show that the successor representation ([Dayan, 1993](#)) can link model-free TD methods to model-based behavior. Abstract latent representations like in the SMT model ([Fang et al., 2019](#), see Section 3) can also adapt flexibly to environmental changes without an explicit model: Adding an observation of a change in the environment to the scene memory would globally change the embedding of all other scene memories, and the latent representation could quickly adapt to reflect the new environment. Another memory mechanism linking model-free TD-RL algorithms to seemingly *predictive* behavior are *eligibility traces* ([Sutton, 1988](#)). Instead of updating only values of the current (single timestep) state-action pair with the current reward, a trace of previous experiences is updated as well. A current negative reward, e.g., related to an obstacle, would for example affect the agent’s evaluation of an earlier state, causing it to adapt behavior early. However, while the behavior looks predictive, it is in fact reactive: In order to learn to avoid a new obstacle, one would still have to experience the novel situation a couple of times. [Wystrach et al. \(2020\)](#) used a very similar concept of memory trace learning,¹⁸ mediated by KC activity traces in the MB, to show how desert ants learn to avoid new obstacles. Note that memory traces could potentially also account for other learning

¹⁸ This is essentially the AL analog to eligibility traces.

behavior reported in honeybees, like the delayed match-to-sample task (Giurfa et al., 2001): traces containing similar (“matching”) and dissimilar (“non-matching”) stimuli would constitute distinct states for which different actions are learned, eliciting different behaviors depending on whether a match is ‘recognized’ or not.

Our MB/CX RL model has the potential to enable fast adaptation to changes: As discussed in the previous paragraph, a small change in global high-level policy due to environmental changes (empty feeder) can lead to sudden change in the local, low-level behavior (steering toward a different goal): the question would now be, how fast can the agent adapt the high-level policy to the environmental change? Since local steering is handled by the low-level controller, and the agent learns a policy over vector memories on a relatively large spatial scale, the agent doesn’t need to constantly reassess said policy, but can do so in larger intervals. This would effectively increase the temporal scale of the MDP, and therefore reduce the number of intermediate states between rewarding experiences. In this “sped-up” MDP with much denser rewards, policy adaptation due to a local change in the environment is propagated faster to distant states, through fewer intermediate states. This effect could be further enhanced by a memory trace mechanism. Taken together, this would enable the agent to quickly adapt local behavior based on distant changes in the environment, by updating its hierarchical, topological spatial representation of the world, without the need for a predictive internal world model.

7.4 Experimental validation of RL based insect navigation

In any RL based navigation model, including the hierarchical MB/CX model proposed above, the agent will learn policies, and suitable underlying representations of the sensory input, which optimize long term reward over extended multi-location foraging trips. While it has been demonstrated in flight cages (Lihoreau et al., 2010) and in the field (Lihoreau et al., 2012) that bumblebees are able to optimize their route between multiple artificial feeders (“traplining”) in an approximate solution to the “traveling salesman problem” (TSP), the underlying optimization mechanisms are not yet fully understood. Overly simplistic heuristics like nearest-neighbor rules have been ruled out based on observed behavior (Lihoreau et al., 2011). In Lihoreau et al. (2012), the authors propose an iterative optimization heuristic—reinforcing route segments when the total length of the foraging bout is shorter than previous ones—which seems to be in line with observations, but remain vague on neural mechanisms allowing the bee to compare total distance traveled between bouts. On the other hand, (deep) reinforcement learning methods have a long history of application to variations of the TRP both theoretically (Gambardella and Dorigo, 1995; Zhang et al., 2023) and in hardware applications like a drone (Bogyrbayeva et al., 2023). The model we propose may therefore provide an alternative explanation to the observed optimization, based on a neuro-anatomically plausible learning mechanism rather than innate heuristics. As a first step, it will be useful to replicate the relevant experiments in simulations based on the proposed architecture and compare observed with RL based and heuristic behavior. We expect differences mostly

in the way optimal routes are consolidated over time, with a more gradual transition if the optimal policy is learned continually as in the proposed RL model. More recent analyses (Kembro et al., 2019) also revealed that trained bumblebees switch between exploitative (optimized) and more exploratory trajectories, which may easily be integrated into the RL model by a variable “determinism parameter” (c.f. Section 5.1.2). The crucial difference between the heuristic and the RL model, however, is in the internal representation of the optimized route: In the former, it is represented as a probability distribution over sequences of relative vectors between feeder locations, in the latter it is directly coupled to sensory input, based on which the optimal next feeder (vector memory) is learned and chosen. As mentioned earlier, this is a crucial property of the proposed model, setting it apart from most current models of insect navigation (Webb, 2019; Goulard et al., 2023 e.g.). One way to assess this would be a displacement experiment with bees that have already learned an optimized route. Starting from a visually familiar location, but with a zero home vector, an agent using visual RL would be able to immediately follow the remainder of the learned route from that location, while an agent relying on vector sequences alone would not. Novel outdoor tracking technology using a drone platform (Vo-Doan et al., 2024) could enable these kinds of experiments, tracking individual foraging bouts between *natural* food sources in unprecedented spatial and temporal resolution.

7.5 Limitations of the proposed model

Besides being a mere model sketch which has not been implemented and tested yet, some important limitations can already be identified: A navigation model based purely on vision and excluding other sensory modalities is a useful simplification for now, but ultimately a comprehensive, realistic model must take these into account—both as the basis for multi-sensory state spaces and a potential source for a reward signal (c.f. Section 7.1). The model can also not account for different behavioral modes like learning, foraging and exploration flights observed in insects. This may be possible by using parallel instances of the model operating on different reward signals associated with those modes, but the feasibility and biological plausibility of this approach has yet to be tested. Lastly, our model can currently not account for a biological mechanism for the eviction and re-assignment of vector memories to MBONs, which is essential for the stability of the learning algorithm. We suspect that this could be implemented by activity-dependent plasticity in combination with temporal decay of the memory, but a concrete mechanism has yet to be formulated.

7.6 RL as a normative framework: place-cell like activity in KCs?

So far, we have treated the visual system as a static component of the model. While this assumption is largely consistent with current knowledge about plasticity in the insect visual system, we can again expand the temporal horizon to evolutionary timescales and view the anatomical components of the insect MB, CX, and visual system

in the light of end-to-end RL, as outlined in the discussion of Section 4. Without going into anatomical details of the insect visual system here, a sensor-motor signal transduction pathway: retina → lamina → medulla → lobula/lobula plate → visual PNs → KCs → MBONs → (pre)motor neurons, could be modeled as a deep neural network implementing a DRL architecture, whose latent spatial representation would then correspond to KC activity (the input layer to the actual policy network, represented by the MB). The learned representations could be indicative of actual spatial representations used by navigating insects, and provide guidance for experimental work. For example, it would be interesting to see if such a model learns a grid-like spatial representation which would correspond to place-cell like KC activity. This hypothesis is compatible with known sparse firing patterns of KCs, but otherwise speculative given current neurophysiological data. It is also unclear how such a representation would fit into the existing model for insect navigation. Since CX based PI is firmly established as a key component for insect navigation, it seems imperative to eventually include the CX in such an end-to-end RL model. This would enable representations that embed PI-like components in a more complex latent space. However, current insect navigation models do not include CX→KC connections and it is not straightforward how the CX would be integrated into a MB based end-to-end RL model. It will also be challenging to strike a balance between expressive power of the network architecture – essential for gaining new insights about possible representations – and necessary constraints to match empirically known components of these representations, like the PI home vector.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

Author contributions

SL: Conceptualization, Visualization, Writing – original draft, Writing – review & editing. DH: Conceptualization, Writing –

review & editing. AV: Conceptualization, Funding acquisition, Supervision, Writing – review & editing. AS: Conceptualization, Funding acquisition, Supervision, Writing – review & editing.

Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article. We gratefully acknowledge financial support from the Brain Worlds Initiative at the University of Freiburg and the Volkswagen Foundation Momentum Program (AZ 98692 to AS). We further acknowledge support by the Open Access Publication Fund of the University of Freiburg.

Acknowledgments

We thank Joschka Boedecker, Paulina Friemann, Michael J. M. Harrap and Christian Leibold, as well as two reviewers for helpful comments and discussions.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Alemi, A. A., Fischer, I., Dillon, J. V., and Murphy, K. (2019). Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*.
- Ardin, P., Peng, F., Mangan, M., Lagogiannis, K., and Webb, B. (2016). Using an insect mushroom body circuit to encode route memory in complex natural environments. *PLoS Comput. Biol.* 12:e1004683. doi: 10.1371/journal.pcbi.1004683
- Arena, E., Arena, P., Strauss, R., and Patané, L. (2017). Motor-skill learning in an insect inspired neuro-computational control system. *Front. Neurobot.* 11. doi: 10.3389/fnbot.2017.00012
- Aso, Y., Sitaraman, D., Ichinose, T., Kaun, K. R., Vogt, K., Belliard-Guérin, G., et al. (2014). Mushroom body output neurons encode valence and guide memory-based action selection in *Drosophila*. *Elife* 3:e04580. doi: 10.7554/eLife.04580
- Baddeley, B., Graham, P., Husbands, P., and Philippides, A. (2012). A model of ant route navigation driven by scene familiarity. *PLoS Comput. Biol.* 8:e1002336. doi: 10.1371/journal.pcbi.1002336
- Bai, H., Cheng, R., and Jin, Y. (2023). Evolutionary reinforcement learning: a survey. *Intell. Comput.* 2:0025. doi: 10.34133/icomputing.0025
- Beeching, E., Wolf, C., Dibangoye, J., and Simonin, O. (2019). Deep reinforcement learning on a budget: 3D control and reasoning without a supercomputer. *arXiv preprint arXiv:1904.01806*.
- Bennett, J. E. M., Philippides, A., and Nowotny, T. (2021). Learning with reinforcement prediction errors in a model of the *Drosophila* mushroom body. *Nat. Commun.* 12:2569. doi: 10.1038/s41467-021-22592-4
- Bogyrbayeva, A., Yoon, T., Ko, H., Lim, S., Yun, H., and Kwon, C. (2023). A deep reinforcement learning approach for solving the Traveling Salesman Problem with Drone. *Transport. Res. Part C* 148:103981. doi: 10.1016/j.trc.2022.103981
- Burda, Y., Edwards, H., Storkey, A., and Klimov, O. (2018). Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*.

- Caron, S. J. C., Ruta, V., Abbott, L. F., and Axel, R. (2013). Random convergence of olfactory inputs in the drosophila mushroom body. *Nature* 497, 113–117. doi: 10.1038/nature12063
- Cartwright, B. A., and Collett, T. S. (1983). Landmark learning in bees. *J. Compar. Physiol.* 151, 521–543. doi: 10.1007/BF00605469
- Chaplot, D. S., Gandhi, D., Gupta, S., Gupta, A., and Salakhutdinov, R. (2020). Learning to explore using active neural slam. *arXiv preprint arXiv:2004.05155*.
- Cognigni, P., Felsenberg, J., and Waddell, S. (2018). Do the right thing: neural network mechanisms of memory formation, expression and update in drosophila. *Curr. Opin. Neurobiol.* 49, 51–58. doi: 10.1016/j.conb.2017.12.002
- Collett, T. S., and Hempel de Ibarra, N. (2023). An 'instinct for learning': the learning flights and walks of bees, wasps and ants from the 1850s to now. *J. Exper. Biol.* 226:jeb245278. doi: 10.1242/jeb.245278
- Dayan, P. (1993). Improving generalization for temporal difference learning: the successor representation. *Neural Comput.* 5, 613–624. doi: 10.1162/neco.1993.5.4.613
- Dhein, K. (2023). The cognitive map debate in insects: a historical perspective on what is at stake. *Stud. Hist. Philos. Sci.* 98, 62–79. doi: 10.1016/j.shpsa.2022.12.008
- Dragoi, G., and Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature* 469, 397–401. doi: 10.1038/nature09633
- Ehmer, B., and Gronenberg, W. (2002). Segregation of visual input to the mushroom bodies in the honeybee (*Apis mellifera*). *J. Compar. Neurol.* 451, 362–373. doi: 10.1002/cne.10355
- Endres, F., Hess, J., Engelhard, N., Sturm, J., Cremers, D., and Burgard, W. (2012). "An evaluation of the rgb-d slam system," in *2012 IEEE International Conference on Robotics and Automation (IEEE)*, 1691–1696. doi: 10.1109/ICRA.2012.6225199
- Engel, J., Schöps, T., and Cremers, D. (2014). "LSD-SLAM: large-scale direct monocular SLAM," in *Computer Vision-ECCV 2014*, eds. D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars (Cham: Springer International Publishing), 834–849. doi: 10.1007/978-3-319-10605-2_54
- Eschbach, C., Fushiki, A., Winding, M., Schneider-Mizell, C. M., Shao, M., Arruda, R., et al. (2020). Recurrent architecture for adaptive regulation of learning in the insect brain. *Nat. Neurosci.* 23, 544–555. doi: 10.1038/s41593-020-0607-9
- Faghihi, F., Moustafa, A. A., Heinrich, R., and Wörgötter, F. (2017). A computational model of conditioning inspired by Drosophila olfactory system. *Neural Netw.* 87, 96–108. doi: 10.1016/j.neunet.2016.11.002
- Fang, K., Toshev, A., Fei-Fei, L., and Savarese, S. (2019). "Scene memory transformer for embodied agents in long-horizon tasks," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 538–547. doi: 10.1109/CVPR.2019.00063
- Felsenberg, J., Barnstedt, O., Cognigni, P., Lin, S., and Waddell, S. (2017). Re-evaluation of learned information in Drosophila. *Nature* 544, 240–244. doi: 10.1038/nature21716
- Fenton, A. A. (2024). Remapping revisited: How the hippocampus represents different spaces. *Nat. Rev. Neurosci.* 25, 428–448. doi: 10.1038/s41583-024-00817-x
- Fuentes-Pacheco, J., Ruiz-Ascencio, J., and Rendón-Mancha, J. M. (2015). Visual simultaneous localization and mapping: a survey. *Artif. Intell. Rev.* 43, 55–81. doi: 10.1007/s10462-012-9365-8
- Gambardella, L. M., and Dorigo, M. (1995). "Ant-Q: a reinforcement learning approach to the traveling salesman problem," in *Machine Learning Proceedings*, eds. A. Prieditis, and S. Russell (San Francisco, CA: Morgan Kaufmann), 252–260. doi: 10.1016/B978-1-55860-377-6.50039-6
- Geng, H., Lafon, G., Avargués-Weber, A., Buatois, A., Massou, I., and Giurfa, M. (2022). Visual learning in a virtual reality environment upregulates immediate early gene expression in the mushroom bodies of honey bees. *Commun. Biol.* 5, 1–11. doi: 10.1038/s42003-022-03075-8
- Giurfa, M., Zhang, S., Jenett, A., Menzel, R., and Srinivasan, M. V. (2001). The concepts of 'sameness' and 'difference' in an insect. *Nature* 410, 930–933. doi: 10.1038/35073582
- Goulard, R., Heinze, S., and Webb, B. (2023). Emergent spatial goals in an integrative model of the insect central complex. *PLoS Comput. Biol.* 19:e1011480. doi: 10.1371/journal.pcbi.1011480
- Greve, E., Büchner, M., Vödisch, N., Burgard, W., and Valada, A. (2023). Collaborative dynamic 3d scene graphs for automated driving. *arXiv preprint arXiv:2309.06635*. doi: 10.1109/ICRA57147.2024.10610112
- Gu, Q., Kuwajerwala, A., Morin, S., Jatavallabhula, K. M., Sen, B., Agarwal, A., et al. (2023). ConceptGraphs: open-vocabulary 3D scene graphs for perception and planning. *arXiv preprint arXiv:2309.16650*.
- Gupta, S., Fouhey, D., Levine, S., and Malik, J. (2017). Unifying map and landmark based representations for visual navigation. *arXiv preprint arXiv:1712.08125*.
- Gupta, S., Tolani, V., Davidson, J., Levine, S., Sukthankar, R., and Malik, J. (2019). Cognitive mapping and planning for visual navigation. *arXiv preprint arXiv:1702.03920*.
- Gutmann, J.-S., Fukuchi, M., and Fujita, M. (2008). 3D perception and environment map generation for humanoid robot navigation. *Int. J. Rob. Res.* 27, 1117–1134. doi: 10.1177/0278364908096316
- Haarnoja, T., Hartikainen, K., Abbeel, P., and Levine, S. (2018). Latent space policies for hierarchical reinforcement learning. *arXiv preprint arXiv:1804.02808*.
- Hafner, D., Lee, K.-H., Fischer, I., and Abbeel, P. (2022). Deep hierarchical planning from pixels. *arXiv preprint arXiv:2206.04114*.
- Henriques, J. F., and Vedaldi, A. (2018). "MapNet: an allocentric spatial memory for mapping environments," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8476–8484. doi: 10.1109/CVPR.2018.00884
- Hoinville, T., and Wehner, R. (2018). Optimal multiguide integration in insect navigation. *Proc. Nat. Acad. Sci.* 115, 2824–2829. doi: 10.1073/pnas.1721668115
- Honerkamp, D., Büchner, M., Despinoy, F., Welschhold, T., and Valada, A. (2024). Language-grounded dynamic scene graphs for interactive object search with mobile manipulation. *arXiv preprint arXiv:2403.08605*.
- Huerta, R., and Nowotny, T. (2009). Fast and robust learning by reinforcement signals: explorations in the insect brain. *Neural Comput.* 21, 2123–2151. doi: 10.1162/neco.2009.03-08-733
- Huerta, R., Nowotny, T., García-Sánchez, M., Abarbanel, H. D. I., and Rabinovich, M. I. (2004). Learning classification in the olfactory system of insects. *Neural Comput.* 16, 1601–1640. doi: 10.1162/08997660474201613
- Hughes, N., Chang, Y., and Carlone, L. (2022). Hydra: a real-time spatial perception system for 3D scene graph construction and optimization. *arXiv preprint arXiv:2201.13360*. doi: 10.15607/RSS.2022.XLVIII.050
- Hulse, B. K., Haberkern, H., Franconville, R., Turner-Evans, D., Takemura, S.-y., Wolff, T., et al. (2021). A connectome of the Drosophila central complex reveals network motifs suitable for flexible navigation and context-dependent action selection. *Elife* 10:e66039. doi: 10.7554/eLife.66039
- Jacobs, L. F. (2012). From chemotaxis to the cognitive map: the function of olfaction. *Proc. Nat. Acad. Sci.* 109, 10693–10700. doi: 10.1073/pnas.1201880109
- Jayatilaka, P., Murray, T., Narendra, A., and Zeil, J. (2018). The choreography of learning walks in the Australian jack jumper ant *Myrmecia croslandi*. *J. Exper. Biol.* 221:jeb185306. doi: 10.1242/jeb.185306
- Jürgensen, A.-M., Sakagiannis, P., Schleyer, M., Gerber, B., and Nawrot, M. P. (2024). Prediction error drives associative learning and conditioned behavior in a spiking model of Drosophila larva. *iScience* 27:108640. doi: 10.1016/j.isci.2023.108640
- Kalweit, G., Huegle, M., Werling, M., and Boedecker, J. (2020). "Deep inverse q-learning with constraints," in *Advances in Neural Information Processing Systems* (Curran Associates, Inc.), 14291–14302.
- Kalweit, G., Kalweit, M., Alyahyay, M., Jaeckel, Z., Steenbergen, F., Hardung, S., et al. (2022). NeuRL: closed-form inverse reinforcement learning for neural decoding. *arXiv preprint arXiv:2204.04733*.
- Kembro, J. M., Lihoreau, M., Garriga, J., Raposo, E. P., and Bartumeus, F. (2019). Bumblebees learn foraging routes through exploitation-exploration cycles. *J. R. Soc. Interface* 16:20190103. doi: 10.1098/rsif.2019.0103
- Khona, M., and Fiete, I. R. (2022). Attractor and integrator networks in the brain. *Nat. Rev. Neurosci.* 23, 744–766. doi: 10.1038/s41583-022-00642-0
- Klein, G., and Murray, D. (2007). "Parallel tracking and mapping for small AR workspaces," in *2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality* (Nara, Japan: IEEE), 1–10. doi: 10.1109/ISMAR.2007.4538852
- Konolige, K., Bowman, J., Chen, J. D., Mihelich, P., Calonder, M., Lepetit, V., et al. (2009). "View-based maps," in *Robotics: Science and Systems V*, 05. doi: 10.15607/RSS.2009.V.020
- Lafon, G., Howard, S. R., Paffhausen, B. H., Avargués-Weber, A., and Giurfa, M. (2021). Motion cues from the background influence associative color learning of honey bees in a virtual-reality scenario. *Sci. Rep.* 11:21127. doi: 10.1038/s41598-021-00630-x
- Le Moël, F., Stone, T., Lihoreau, M., Wystrach, A., and Webb, B. (2019). The central complex as a potential substrate for vector based navigation. *Front. Psychol.* 10:690. doi: 10.3389/fpsyg.2019.00690
- Lee, T. W., Girolami, M., and Sejnowski, T. J. (1999). Independent component analysis using an extended infomax algorithm for mixed subgaussian and supergaussian sources. *Neural Comput.* 11, 417–441. doi: 10.1162/089976699300016719
- Lihoreau, M., Chittka, L., Le Comber, S. C., and Raine, N. E. (2011). Bees do not use nearest-neighbour rules for optimization of multi-location routes. *Biol. Lett.* 8, 13–16. doi: 10.1098/rsbl.2011.0661
- Lihoreau, M., Chittka, L., and Raine, N. E. (2010). Travel optimization by foraging bumblebees through readjustments of traplines after discovery of new feeding locations. *Am. Nat.* 176, 744–757. doi: 10.1086/657042
- Lihoreau, M., Raine, N. E., Reynolds, A. M., Stelzer, R. J., Lim, K. S., Smith, A. D., et al. (2012). Radar tracking and motion-sensitive cameras on flowers reveal the development of pollinator multi-destination routes over large spatial scales. *PLoS Biol.* 10:e1001392. doi: 10.1371/journal.pbio.1001392
- Liu, C., Plaçais, P.-Y., Yamagata, N., Pfeiffer, B. D., Aso, Y., Friedrich, A. B., et al. (2012). A subset of dopamine neurons signals reward for odour memory in Drosophila. *Nature* 488, 512–516. doi: 10.1038/nature11304

- Lulham, A., Bogacz, R., Vogt, S., and Brown, M. W. (2011). An Infomax algorithm can perform both familiarity discrimination and feature extraction in a single network. *Neural Comput.* 23, 909–926. doi: 10.1162/NECO_a_00097
- Lyu, C., Abbott, L. F., and Maimon, G. (2022). Building an allocentric travelling direction signal via vector computation. *Nature* 601, 92–97. doi: 10.1038/s41586-021-04067-0
- Matheson, A. M. M., Lanz, A. J., Medina, A. M., Licata, A. M., Currier, T. A., Syed, M. H., et al. (2022). A neural circuit for wind-guided olfactory navigation. *Nat. Commun.* 13:4613. doi: 10.1038/s41467-022-32247-7
- Matsuki, H., Murai, R., Kelly, P. H., and Davison, A. J. (2023). Gaussian splatting slam. *arXiv preprint arXiv:2312.06741*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., et al. (2015). Human-level control through deep reinforcement learning. *Nature* 518, 529–533. doi: 10.1038/nature14236
- Mobbs, P. G., and Young, J. Z. (1997). The brain of the honeybee *Apis mellifera*. I. The connections and spatial organization of the mushroom bodies. *Philos. Trans. R. Soc. London* 298, 309–354. doi: 10.1098/rstb.1982.0086
- Mur-Artal, R., and Tardós, J. D. (2017). ORB-SLAM2: an open-source SLAM system for monocular, stereo, and RGB-D cameras. *IEEE Trans. Robot.* 33, 1255–1262. doi: 10.1109/TRO.2017.2705103
- Nachum, O., Gu, S., Lee, H., and Levine, S. (2018). Data-efficient hierarchical reinforcement learning. *arXiv preprint arXiv:1805.08296*.
- O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* 34, 171–175. doi: 10.1016/0006-8993(71)90358-1
- Oudeyer, P.-Y., and Kaplan, F. (2008). “How can we define intrinsic motivation?” in *Proceedings of the Eight International Conference on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, eds. M. Schlesinger, L. Berthouze, and C. Balkenius (Lund University Cognitive Science).
- Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction. *arXiv preprint arXiv:1705.05363*. doi: 10.1109/CVPRW.2017.70
- Ramakrishnan, S. K., Chaplot, D. S., Al-Halah, Z., Malik, J., and Grauman, K. (2022). “Poni: Potential functions for objectgoal navigation with interaction-free learning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18890–18900. doi: 10.1109/CVPR52688.2022.01832
- Rana, K., Haviland, J., Garg, S., Abou-Chakra, J., Reid, I., and Suenderhauf, N. (2023). Sayplan: grounding large language models using 3D scene graphs for scalable task planning. *arXiv preprint arXiv:2307.06135*.
- Rapp, H., and Nawrot, M. P. (2020). A spiking neural program for sensorimotor control during foraging in flying insects. *Proc. Nat. Acad. Sci.* 117, 28412–28421. doi: 10.1073/pnas.2009821117
- Rescorla, R. A. (1972). A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and non-reinforcement. *Class. Condit. Curr. Res. Theory* 2, 64–69.
- Rich, P. D., Liaw, H.-P., and Lee, A. K. (2014). Large environments reveal the statistical structure governing hippocampal representations. *Science* 345, 814–817. doi: 10.1126/science.1255635
- Rosinol, A., Leonard, J. J., and Carlone, L. (2023). “Nerf-slam: real-time dense monocular slam with neural radiance fields,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE), 3437–3444. doi: 10.1109/IROS55552.2023.10341922
- Russek, E. M., Momennejad, I., Botvinick, M. M., Gershman, S. J., and Daw, N. D. (2017). Predictive representations can link model-based reinforcement learning to model-free mechanisms. *PLoS Comput. Biol.* 13:e1005768. doi: 10.1371/journal.pcbi.1005768
- Rybak, J., and Menzel, R. (1993). Anatomy of the mushroom bodies in the honey bee brain: the neuronal connections of the alpha-lobe. *J. Compar. Neurol.* 334, 444–465. doi: 10.1002/cne.903340309
- Savinov, N., Raichuk, A., Marinier, R., Vincent, D., Pollefeys, M., Lillicip, T., et al. (2019). Episodic curiosity through reachability. *arXiv preprint arXiv:1810.02274*.
- Schmalstieg, F., Honerkamp, D., Welschehold, T., and Valada, A. (2022). “Learning long-horizon robot exploration strategies for multi-object search in continuous action spaces,” in *The International Symposium of Robotics Research* (Springer), 52–66. doi: 10.1007/978-3-031-25555-7_5
- Schmalstieg, F., Honerkamp, D., Welschehold, T., and Valada, A. (2023). Learning hierarchical interactive multi-object search for mobile manipulation. *IEEE Robot. Autom. Lett.* 8, 8549–8556. doi: 10.1109/LRA.2023.3329619
- Shah, D., Eysenbach, B., Kahn, G., Rhinehart, N., and Levine, S. (2023). Rapid exploration for open-world navigation with latent goal models. *arXiv preprint arXiv:2104.05859*.
- Shah, D., and Levine, S. (2022). “ViKiNG: vision-based kilometer-scale navigation with geographic hints,” in *Robotics: Science and Systems XVIII*. doi: 10.15607/RSS.2022.XVIII.019
- Stemmler, M., Mathis, A., and Herz, A. V. M. (2015). Connecting multiple spatial scales to decode the population activity of grid cells. *Sci. Adv.* 1:e1500816. doi: 10.1126/science.1500816
- Stone, T., Webb, B., Adden, A., Weddig, N. B., Honkanen, A., Templin, R., et al. (2017). An anatomically constrained model for path integration in the bee brain. *Curr. Biol.* 27, 3069–3085.e11. doi: 10.1016/j.cub.2017.08.052
- Strube-Bloss, M. F., and Rössler, W. (2018). Multimodal integration and stimulus categorization in putative mushroom body output neurons of the honeybee. *R. Soc. Open Sci.* 5:171785. doi: 10.1098/rsos.171785
- Sun, X., Yue, S., and Mangan, M. (2020). A decentralised neural model explaining optimal integration of navigational strategies in insects. *Elife* 9:e54026. doi: 10.7554/eLife.54026
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Mach. Learn.* 3, 9–44. doi: 10.1007/BF00115009
- Sutton, R. S., and Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. Adaptive Computation and Machine Learning Series. Cambridge, MA: The MIT Press.
- Sutton, R. S., McAllester, D., Singh, S., and Mansour, Y. (1999). “Policy gradient methods for reinforcement learning with function approximation,” in *Advances in Neural Information Processing Systems* (MIT Press).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). “Attention is all you need,” in *Advances in neural information processing systems* (Curran Associates, Inc.).
- Vödisch, N., Cattaneo, D., Burgard, W., and Valada, A. (2022). “Continual SLAM: beyond lifelong simultaneous localization and mapping through continual learning,” in *The International Symposium of Robotics Research* (Springer), 19–35. doi: 10.1007/978-3-031-25555-7_3
- Vo-Doan, T. T., Titov, V. V., Harrap, M. J. M., Lochner, S., and Straw, A. D. (2024). High resolution outdoor videography of insects using fast lock-on tracking. *bioRxiv preprint*, 2023.12.20.572558. doi: 10.1101/2023.12.20.572558
- Vogt, K., Schnaitmann, C., Dylla, K. V., Knapke, S., Aso, Y., Rubin, G. M., et al. (2014). Shared mushroom body circuits underlie visual and olfactory memories in *Drosophila*. *Elife* 3:e02395. doi: 10.7554/eLife.02395
- Wani, S., Patel, S., Jain, U., Chang, A. X., and Savva, M. (2020). MultiON: benchmarking semantic map memory using multi-object navigation. *arXiv preprint arXiv:2012.03912*.
- Warren, W. H., Rothman, D. B., Schnapp, B. H., and Ericson, J. D. (2017). Wormholes in virtual space: from cognitive maps to cognitive graphs. *Cognition* 166, 152–163. doi: 10.1016/j.cognition.2017.05.020
- Webb, B. (2019). The internal maps of insects. *J. Exper. Biol.* 222:jeb188094. doi: 10.1242/jeb.188094
- Webb, B. (2024). Beyond prediction error: 25 years of modeling the associations formed in the insect mushroom body. *Lear. Memory* 31:a053824. doi: 10.1101/lm.053824.123
- Webb, B., and Wystrach, A. (2016). Neural mechanisms of insect navigation. *Curr. Opin. Insect Sci.* 15, 27–39. doi: 10.1016/j.cois.2016.02.011
- Wei, T., Guo, Q., and Webb, B. (2024). Learning with sparse reward in a gap junction network inspired by the insect mushroom body. *PLoS Comput. Biol.* 20:e1012086. doi: 10.1371/journal.pcbi.1012086
- Werby, A., Huang, C., Büchner, M., Valada, A., and Burgard, W. (2024). Hierarchical open-vocabulary 3D scene graphs for language-grounded robot navigation. *arXiv preprint arXiv:2403.17846*.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.* 8, 229–256. doi: 10.1007/BF00992696
- Wystrach, A. (2023). Neurons from pre-motor areas to the Mushroom bodies can orchestrate latent visual learning in navigating insects. *bioRxiv preprint* 2023-03. doi: 10.1101/2023.03.09.531867
- Wystrach, A., Buehlmann, C., Schwarz, S., Cheng, K., and Graham, P. (2020). Rapid aversive and memory trace learning during route navigation in desert ants. *Curr. Biol.* 30, 1927–1933.e2. doi: 10.1016/j.cub.2020.02.082
- Wystrach, A., Philippides, A., Aurejac, A., Cheng, K., and Graham, P. (2014). Visual scanning behaviours and their role in the navigation of the Australian desert ant *Melophorus bagoti*. *J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* 200, 615–626. doi: 10.1007/s00359-014-0900-8
- Xiao, X., Xu, Z., Wang, Z., Song, Y., Warnell, G., Stone, P., et al. (2022). Autonomous ground navigation in highly constrained spaces: lessons learned from the benchmark autonomous robot navigation challenge at icra 2022 [competitions]. *IEEE Robot. Autom. Mag.* 29, 148–156. doi: 10.1109/MRA.2022.3213466
- Younes, A., Honerkamp, D., Welschehold, T., and Valada, A. (2023). Catch me if you hear me: audio-visual navigation in complex unmapped environments with moving sounds. *IEEE Robot. Autom. Lett.* 8, 928–935. doi: 10.1109/LRA.2023.3237466
- Zeil, J., Hofmann, M. I., and Chahl, J. S. (2003). Catchment areas of panoramic snapshots in outdoor scenes. *J. Opt. Soc. Am. A* 20:450. doi: 10.1364/JOSA.20.000450

Zeng, F., Wang, C., and Ge, S. S. (2020). A survey on visual navigation for artificial agents with deep reinforcement learning. *IEEE Access* 8, 135426–135442. doi: 10.1109/ACCESS.2020.3011438

Zhang, Z., Liu, H., Zhou, M., and Wang, J. (2023). Solving dynamic traveling salesman problems with deep reinforcement learning. *IEEE Trans. Neural Netw. Lear. Syst.* 34, 2119–2132. doi: 10.1109/TNNLS.2021.3105905

Zhu, K., and Zhang, T. (2021). Deep reinforcement learning based mobile robot navigation: a review. *Tsinghua Sci. Technol.* 26, 674–691. doi: 10.26599/TST.2021.9010012

Zhu, Z., Peng, S., Larsson, V., Xu, W., Bao, H., Cui, Z., et al. (2021). “Nice-slam: neural implicit scalable encoding for slam. 2022 IEEE,” in *CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12776–12786. doi: 10.1109/CVPR52688.2022.01245