Holistic behavioral modeling

A case study on *Drosophila* larva foraging

Dissertation
for the award of the doctoral degree

Doctor of Philosophy in Natural Sciences
(Doctor rerum naturalium)

Faculty of Mathematics and Natural Sciences
University of Cologne

submitted by

Panagiotis Parthenios Sakagiannis

accepted in the year 2025

University of Cologne
Faculty of Mathematics and Natural Sciences

Doctor of Philosophy in Natural Sciences

Holistic behavioral modeling: A case study on *Drosophila* larva foraging

by Panagiotis Parthenios Sakagiannis

Abstract

This thesis presents the Behavioral Architecture - Dynamic Energy Budget (BA-DEB) framework, a holistic computational approach to behavioral modeling grounded in four core commitments: (i) agent-centered modeling, where the organism is treated as an internally structured and environmentally embedded unit; (ii) comprehensive behavioral scope, instantiated in a modular, hierarchical BA that spans from motor primitives to experience-shaped adaptivity; (iii) cross-timescale integration of neural, behavioral, and metabolic processes; and (iv) homeostatic regulation, achieved through the coupling of the BA with a DEB model that simulates internal energetics and shapes behavior through metabolic feedback. These principles are exemplified in a case study of foraging behavior in the *Drosophila melanogaster* larva – a model organism with a tractable nervous system, rich behavioral repertoire and readily recordable 2D posture – engaged in a well-characterized, structured behavior during a normatively narrow life stage focused on growth and survival.

Standalone mechanistic models form a second core contribution of this thesis. These include a stochastic network model of behavioral intermittency and a coupled-oscillator model that captures the biomechanical interference of crawling on lateral bending. These models formalize biologically plausible mechanistic hypotheses and can be integrated into the BA as modular components, reflecting the framework's capacity to accommodate diverse behavioral mechanisms within a unified control structure.

The BA-DEB framework is implemented in *Larvaworld*, an open-source simulation and analysis platform that generates realistic behavior by combining mechanistic modeling with data-driven fitting. It facilitates exploratory modeling via unbiased empirical validation, and interdisciplinary collaboration through standardized behavioral modules and flexible experimental configurations. Together, the framework and software platform provide a robust and extensible, conceptual and methodological foundation for building, testing, and comparing models of behavior grounded in biological detail and organized in a behavior-based modular logic.

Contents

List of Figures

Part I
Introduction

Modeling behavior as a hierarchically organized, temporally sequenced, adaptive process poses a unique challenge: it requires integrating diverse levels of biological organization – from neural dynamics to ecological context – while preserving the coherence of the individual agent. Traditional modeling approaches often segment this complexity, focusing either on internal mechanisms or environmental interactions. Yet behavior emerges precisely at their interface. This thesis addresses the need for a unified modeling framework that is both holistic and mechanistic: one that treats the behaving organism as an autonomous agent, and behavioral structure as a computationally tractable anchor for integrating currently standalone models. The *Drosophila* larva, with its well-characterized nervous system and rich behavioral repertoire, provides an ideal model system for developing and testing such a framework.

Chapter 1 provides the theoretical and interdisciplinary context for the modeling approach. It examines how systems neuroscience and behavioral ecology, traditionally separate, are converging on the behaving organism as a shared modeling unit. The chapter then introduces architectural perspectives relevant to modeling, focusing on survival circuits, hierarchical control, and behavioral repertoires as key structural principles that support a modular and mechanistic view of behavior.

Chapter 2 establishes the conceptual foundation of the thesis. It articulates the commitments to holism, behavior as a modeling anchor, and mechanistic explanation, framing them as mutually reinforcing strategies for studying purposive biological systems. The chapter then motivates the choice of the *Drosophila* larva and foraging behavior as the empirical focus, and outlines the modeling objectives and design rationale that guide the development of the computational framework.

1 Integrative perspectives on the study of behavior

The study of behavior spans a wide range of scientific disciplines, each offering distinct theoretical insights and methodological traditions, and addressing scientific questions that pertain to different levels of biological organization. To map this interdisciplinary landscape, this chapter provides a structured overview of the main scientific fields contributing to modern behavioral science. Importantly, this overview explicitly aims to provide theoretical grounding for the modeling perspective and methodological choices adopted throughout this thesis.

The chapter is organized into two primary sections. The first part (section 1.1) examines how two traditionally separate traditions – systems neuroscience and behavioral ecology – are increasingly converging toward a unified understanding of behavior cen-

tered on the single behaving organism. The second part (section 1.2) draws from diverse fields of behavioral study to highlight key aspects of living systems that directly support a behavior-based modeling approach. Specifically, it highlights how the behaving agent is organized into parallel homeostatically-driven survival circuits, and how its behavioral repertoire is governed by a hierarchically structured control architecture. Figure 1 provides a visual overview of the behavioral research fields discussed in this chapter – those converging at the level of the individual (such as systems neuroscience and ecology), and those that support the modeling perspective centered on the single behaving agent adopted in this thesis.

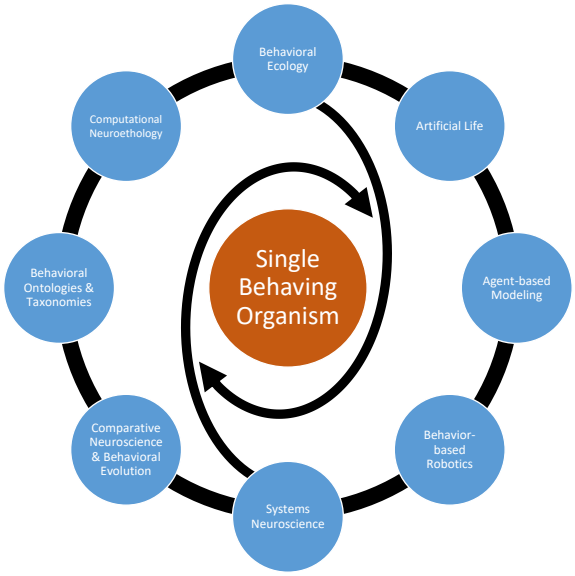

Figure 1: Interdisciplinary landscape around holistic behavioral modeling. At the center lies the individual behaving organism, conceptualized as the interface where subindividual (neuroscientific) and supraindividual (ecological) perspectives converge. Surrounding this focal point are fields that contribute to this convergence by providing theoretical background, scientific insights or methodological tools.

1.1 Convergence between neuroscience and ecology

Despite long-standing methodological and conceptual differences, systems neuroscience and behavioral ecology are increasingly addressing overlapping questions about how organisms generate behavior in response to internal and external conditions. Neuroscience typically approaches this problem by examining subindividual mechanisms – neural circuits, sensory processing, and motor control – under controlled experimental settings. Ecology, by contrast, has focused on behavioral adaptation at the population level, often abstracting away from internal structure.

Recent developments in both fields, however, reflect a growing interest in bridging these perspectives. This convergence is not driven by a unified theoretical framework, but by a shared recognition that understanding behavior requires models that integrate internal dynamics with environmental interaction. This section traces that convergence across three analytical levels: the subindividual (neural architectures), the supraindividual (population dynamics and agent-based models[1] (ABM)), and the individual organism as the interface where these perspectives meet.

1.1.1 The subindividual level: From neural circuits to modular cognitive architectures

The subindividual level refers to modeling approaches that focus on parts of the organism rather than its total behavioral profile. This can be interpreted anatomically, as the study of specific neural substrates (e.g., circuits, neuropiles) in isolation from the entire nervous system, let alone the entire body, or functionally, as the analysis of particular behavioral domains (e.g., foraging, reflexes) in isolation from the full behavioral repertoire. Such decompositions are central to systems neuroscience, which aims to generate mechanistic explanations grounded in neural structure and function.

This level of analysis allows researchers to isolate neural circuits and study how specific components contribute to observable behavior across varying degrees of biological complexity. In organisms with simple nervous systems, such as *C. elegans* – the only organism for which the entire whole-body connectome (302 neurons) has been fully mapped (Cook et al., 2019) – this decomposition can be approached with considerable anatomical precision. Computational modeling can begin from a noisy and biologically grounded connectome and incrementally superimpose behaviorally relevant constraints. This process does not merely reproduce observed behavior, but rather helps formulate plausible explanations of how anatomical structure can support function (Izquierdo and Beer, 2013). These models often focus on specific behaviors, treating them as entry points into larger questions about circuit organization and multifunctionality. As new behavioral tasks are superimposed on the same neural substrate, the constraints on plausible internal models increase, thus refining both their explanatory and predictive scope.

In more complex organisms, such as insects or mammals, complete mapping of behavior onto neural structure is considerably more difficult. Neuroscientific models in these cases typically proceed by identifying correlations between neural activity and behavior, and then testing hypotheses of necessity and sufficiency for specific structures. This process allows tentative functional attributions – for instance, to neuropiles involved in memory, decision-making, or emotional appraisal – hopefully without implying direct one-to-one

[1] A computational modeling paradigm that simulates the actions and interactions of autonomous agents to understand the behavior of complex systems. Agents – each with its own internal structure, goals and capabilities – interact with each other and have only local knowledge of their environment, giving rise to emergent system-level dynamics. ABM is particularly suited for modeling heterogeneous, decentralized, and spatially structured systems (Grimm, 2005).

mappings (Rolls, 2015; Wang, 2008). Increasingly, such attributions are implemented in computational models, which may range from differential-equation-based models of motor control (Daun-Gruhn and Büschges, 2011), local circuit-level simulations (Haenicke et al., 2018) to more ambitious whole-brain constructs (Eliasmith et al., 2012). The latter frequently incorporate optimization methods to account for adaptive behavior under varying task constraints, and reflect a trend towards integrative, multifunctional modeling at various levels of biological realism.

The effort to build models spanning multiple functional domains has led to the development of cognitive architectures – domain-general computational systems that formalize perception, memory, decision-making, and motor control within unified frameworks (Franklin et al., 2016; R. Sun, 2007). While originally rooted in psychology and artificial intelligence (AI), such architectures have been extended to incorporate elements inspired by neurobiological principles, including learning rules and attention dynamics (Verschure, 2012).

These modular constructions are often guided by functionalist assumptions: that cognition consists of separable operations mapping sensory input to appropriate output. However, this view has come under criticism from enactivist[2] perspectives, which challenge the legitimacy of strict modular decomposition (Di Paolo et al., 2017). Nevertheless, modularity remains central to hybrid architectures that aim for both tractability and biological plausibility, supported by empirical research on modular motor control (Büschges and Borgmann, 2013). In this context, explicit assumptions about how low-level sensorimotor modules interface with higher-order processes become crucial – especially when behavioral modeling aspires to be mechanistically explanatory.

1.1.2 The supraindividual level: From population patterns to behavioral ABMs

Whereas systems neuroscience typically dissects behavior by focusing on its subindividual underpinnings, ecological modeling proceeds from the opposite direction: from populations of organisms embedded in dynamic environments. Classical ecological approaches abstract away internal structure, emphasizing population-level dynamics and statistical regularities over individual-level mechanisms. These models operate at the supraindividual level, privileging emergent behavioral patterns over the specific inner workings that give rise to them.

A pivotal transition toward integrating individual-level agency into ecological simulations occurred within the field of Artificial Life[3] (ALife). In many early ALife systems,

[2] For a comprehensive introduction to enactivism and its relevance to behavioral and cognitive modeling, see (Ward et al., 2017).

[3] Artificial Life (ALife) explores "life-as-it-might-be" (Langton, 1992). By abstracting away from specific real-world constraints, it focuses on behavioral patterns and underlying drives that emerge across diverse living systems and subsystems. As a subfield of Artificial Intelligence (AI), AL has pioneered efforts to bridge the nested organizational levels of life – from populations and individuals, to internal functional modules, cells, and genes.

agents were endowed with artificial neural controllers or rule-based systems that mapped sensory input to action. These controllers could evolve under selection pressure, producing adaptive behaviors within dynamic environments (Sims, 1994; Yaeger, 1997). Later models introduced increasingly structured representations of internal states – such as fuzzy cognitive maps – which enabled agents to express emotions, memory, and motivational attitudes, and even to jointly model behavioral evolution and speciation (Gras et al., 2009).

These innovations laid the groundwork for the widespread adoption of agent-based modeling in computational ecology. A key methodological advance was pattern-oriented modeling, which uses observed ecological patterns as criteria for the design and evaluation of models (Grimm, 2005). The ABM paradigm has since developed formalized standards, such as the ODD (Overview, Design Concepts and Details) and ODD+D (which includes decision-making) protocols, to promote transparency, replicability, and biological realism (Grimm et al., 2010; Müller et al., 2013).

Crucially, ABMs offer more than population-level pattern reproduction. They allow researchers to encode intra-individual structure, including morphological traits, internal drives, and decision-making architectures. In evolutionary ABMs, behavioral traits are treated as heritable parameters subject to mutation and selection, a practice known as the phenotypic gambit[4], enabling researchers to link behavioral evolution to population dynamics, now in real-world ecological settings.

Thus, the ABM paradigm has paved the way for a shift in ecological modeling: from a focus on population-level abstraction, increasingly to the internal complexity of individual agents. This methodological shift complements the upward trend in neuroscience toward ecologically valid modeling. The next section explores this convergence explicitly.

1.1.3 The individual as interface between neuroscience and ecology

The separation between neuroscience and ecological traditions, as outlined above, was marked by distinct conceptual frameworks and modeling practices, limiting interdisciplinary integration and a unified understanding of behavior across explanatory levels and timescales. Early attempts to bridge this methodological divide emerged from ecological modeling efforts that sought to embed internal structure within their agents. One early example introduced neural network-based decision-making combined with genetic algorithms into individual-based ecological models (Huse et al., 1999). In this study, neural networks were merely used as computational tools providing agents with internally structured decision-making capabilities, eventually leading to the emergence of survival-maximizing behavioral strategies in dynamic environments. Although the introduction of neural networks does not justify any claim for biological realism, it marks an important

[4]The phenotypic gambit is the assumption that we can model evolutionary outcomes by assuming that natural selection acts directly on phenotypes, treating them as if they were optimally designed, without detailed knowledge of the underlying genetic architecture (Grafen, 1986).

step toward richer models of individual agency within ecological simulations.

More recently, ecological modeling has explicitly incorporated fast homeostatically-driven behavioral mechanisms inspired by neuroscientific and ethological findings. A prominent example is the AHA (Autonomic-Homeostatic-Adaptive) architecture, a framework that embeds "proximate" behavioral mechanisms[5] – such as perception, motivation, decision-making, and attention – into ecological agents (Budaev et al., 2018). By endowing agents with detailed internal structure, the AHA approach captures essential behavioral constraints and dynamics identified through empirical neuroscience (Eliassen et al., 2016). This richer agent design enables not only more plausible simulations of individual behavior but also more powerful ecological models at scale. For instance, in marine computational ecology, AHA-based agents have been employed in large-scale AMBs to simulate and predict the growth, movement, and distribution of millions of fish in oceanic environments, dynamically responding to environmental cues, internal state, and conspecific interactions (Budaev et al., 2018). This integration of neuroethological realism into ecological modeling enhances both explanatory depth and predictive capacity, moving beyond abstract rules toward biologically grounded individual agency.

The shift toward computational agents enriched with neuroscientifically informed internal structures underscores a deeper theoretical convergence: the focal point becomes the single behaving organism, where the fast subindividual mechanisms studied by neuroscience and the slow supraindividual adaptive phenomena investigated by ecology naturally intersect. In this integrative perspective, the behaving individual is recognized as the appropriate locus for computational frameworks that can simultaneously accommodate neural, behavioral, and ecological complexity. By explicitly embedding neural-inspired architectures into ecological models, these computational frameworks represent a promising avenue toward bridging the disciplinary gap that has historically separated systems neuroscience from behavioral ecology. Figure 2 highlights the progressive overlap between previously distinct traditions, ultimately converging at the scale of a single behaving organism.

1.2 Modeling the behaving organism: Architectural perspectives

Understanding behavior as an integrated phenomenon requires characterizing its underlying functional structure. This section draws on diverse fields of behavioral research to identify organizational features of living systems that are especially relevant for behavioral modeling. Rather than attempting to reproduce the full complexity of an organism, the approach adopted here focuses on how internal regulation and overt action can be jointly

[5]Here, "proximate mechanisms" refer specifically to the internal processes that directly influence the immediate control of behavior, grounding ecological simulations in biologically plausible principles derived from neuroscience. The word-choice aims to contrast fast, flexible behavioral reactions typically studied by neuroscience, with the slow adaptive behavioral patterns – e.g., migration or fecundity changes – typically at the focus of ecology.

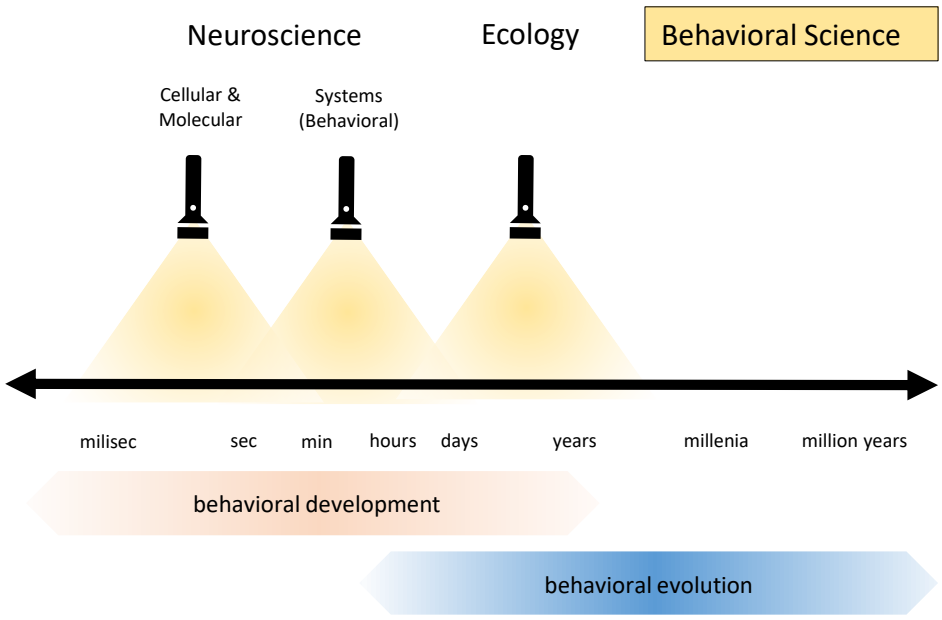

Figure 2: Disciplinary silos in behavioral science. Behavior is studied across a wide range of timescales (logarithmic axis) and at multiple organizational levels – from cellular processes to individual organisms and populations. Traditionally, neuroscience and ecology have focused on distinct, non-overlapping temporal scales, limiting interdisciplinary exchange. However, both fields are increasingly expanding their temporal scope – neuroscience toward longer scales and ecology toward shorter ones – leading to a convergence around the lifespan of the individual behaving organism.

represented through computational architectures. In particular, three interrelated aspects of behaving organisms are examined that provide strategic entry points for modeling.

The first concerns survival circuits – functional subsystems that may diverge structurally across taxa but preserve a core functional role in supporting organismal survival. Each such circuit is a functional unit comprising both homeostatic/metabolic and neural/behavioral components. These circuits operate in parallel to control and regulate behavior according to distinct survival imperatives. Crucially, they do not govern behavior in a top-down manner but are instead expressed through and served by specific behaviors adapted to internal priorities and environmental contingencies.

The second aspect is the hierarchical organization of the neurobehavioral components of a survival circuit. Drawing on both neuroanatomical findings and principles from neurorobotics, layered control architectures are presented as a suitable modeling paradigm to capture the selection, inhibition, and coordination of behavior across nested levels of increasingly complex sensorimotor coupling between agent and environment.

The third involves the extraction and formalization of structured behavioral repertoires, emphasizing their role in populating the layers of the behavioral control architec-

ture. Approaches from computational neuroethology and formal ontology allow for the identification and classification of discrete behavioral units, offering the necessary building blocks for modeling plausible modular behavior-based agents.

Together, these three modeling dimensions reflect a central methodological tension: how to maintain a holistic view of behavior while embracing the necessary simplifications of modeling. One way to address this is by focusing on key organizational nodes – such as a single survival circuit – around which a behavior-based control architecture can be constructed and systematically populated, as will be described in Subsection 2.1.4.

1.2.1 Survival circuits as parallel functional subsystems

Understanding how an individual agent maintains biological viability requires linking internal regulation with context-appropriate behavior. A central concept in this regard is that of survival circuits – evolutionarily conserved neurobehavioral systems that link metabolic priorities to species-typical behavioral strategies. Originally proposed by LeDoux, survival circuits are not anatomically encapsulated modules, but functional subsystems that span slow homeostatic processes (e.g., hunger, thermoregulation) and fast behavioral control loops (e.g., foraging, fleeing) (LeDoux, 2012; Ledoux and Daw, 2018).

Each subsystem consists an autonomous agent-environment feedback loop – sensory appraisal, motivation, decision making and motor competences – anchored to a distinct homeostatic imperative such as energy balance, reproduction, or defense. This combination of internally defined needs and self-deployed behavioral means – an organizing principle for dissecting behavior into biologically grounded units of control – renders them particularly suitable for modeling. Figure 3 provides a schematic illustration of basic survival circuits and their associated behaviors in *Drosophila*, emphasizing the regulatory autonomy of each system, and the nested interaction between slow homeostatic drives and fast sensorimotor control.

Comparative analyses suggest that all evolved nervous systems – from radially organized cnidarians to bilaterally symmetric vertebrates and arthropods – incorporate parallel subsystems dedicated to distinct survival imperatives (Arendt et al., 2016). These subsystems, though anatomically diverse, perform analogous roles: coupling internal homeostatic states to sensorimotor routines that resolve functionally discrete biological challenges. For instance, structures as different as the insect central complex and the vertebrate basal ganglia have been argued to implement similar functions of action selection and spatial coordination, suggesting functional convergence[6] across phyla (Arendt et al., 2016; Cisek, 2021). Similarly the insect mushroom body and the vertebrate hippocampus have both been implicated in spatial navigation and associative learning relevant to foraging contexts

[6]The term "functional convergence" is used here in the sense of analogous – not homologous – traits. Homologous structures share a common evolutionary origin, while analogous ones arise independently but serve similar functions.

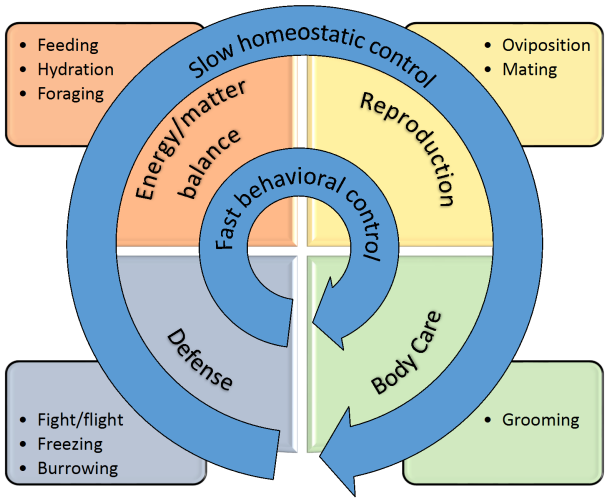

Figure 3: Survival circuits in *Drosophila*. Each functional subsystem enjoys regulatory autonomy by means of nested coordination between slow homeostatic processes and fast sensorimotor control, making them an ideal target for homeostatically-regulated behavioral models. Fast and slow control processes extend to coordination and conflict resolution across circuits.

(Farris, 2011; Xia and Tully, 2007). Bilateral nerve cords and central tracts appear to have evolved independently in multiple lineages, yet serve comparable behavioral control functions. In fact, peripheral motor control circuits across the animal kingdom exhibit common features in terms of sensory feedback and descending input from the brain (Büschges, 2005; Büschges and Ache, 2025; Büschges and Gorostiza, 2023; Mantziaris et al., 2020). This functional alignment across disparate implementations exemplifies how survival circuits can be identified not solely by anatomy, but mainly by their role in resolving core biological imperatives, as fundamental as locomotion.

Survival circuits can be grounded in shared functional roles of behavioral patterns across species. Foraging is a prime example: despite anatomical differences, the core logic – linking energy deficits to exploratory behavior – remains identical. This aligns with Dynamic Energy Budget (DEB) theory, which identifies food acquisition as a universal mechanism for energy assimilation, independent of morphology or phylogeny (Kooijman, 2010). Another paradigmatic case is sleep/quiescence – a behavioral oscillation regulated jointly by circadian and homeostatic drives, and observed across phylogenetically distant species. Despite anatomical differences, core mechanisms such as sleep pressure accumulation and arousal thresholds remain broadly similar (Brown et al., 2012; Fuller et al., 2006; Holland, 2018).

The existence of such analogously organized, evolutionarily converging functional modules supports the view that survival circuits are not idiosyncratic to specific species or brains, but form core constituents of living systems at large and are suitable for modeling behavior across biological taxa. As Klein and Barron (Klein and Barron, 2024) observe,

each such subsystem corresponds to a distinct computational architecture with sufficient autonomy to define a lineage's behavioral repertoire.

From a modeling perspective, the dissection of behavior into parallel survival circuits offers a tractable yet holistic approach. Each such circuit constitutes a teleologically closed unit: it interprets its own deficits, sets its own goals, and recruits its own means of resolution. This autonomy renders it both functionally coherent and normatively self-regulated, allowing researchers to model rich, biologically grounded behavior without reconstructing the entire organismal repertoire.

Moreover, this framework provides a way to reconceptualize composite behavioral categories. Instead of treating terms like "aggression" or "attention" as unitary constructs, one can parse them by the survival circuit in which they operate (Hommel et al., 2019). For example, aggression toward prey, mates, or predators corresponds to distinct motivational and affective regimes – energy acquisition, reproduction, or defense, respectively – and likely recruits different neural substrates although expressed similarly in terms of behavior (Ledoux and Daw, 2018).

In conclusion, the coexistence of multiple survival circuits entails a horizontal modularity in behavioral control: distinct subsystems operate in parallel, each with its intrinsic homeostatically-grounded motivation, regulatory loop and behavioral footprint. Although they may interact or compete for bodily resources and behavioral expression, each retains functional and motivational autonomy. This horizontal structure complements the vertical layering of control described in the next section, which governs prioritization, suppression, and integration of the neural and behavioral components within and across these circuits.

1.2.2 Neurorobotics and hierarchical control architectures

A central feature of biological control is its layered organization – a functional hierarchy that emerges through both evolution and development. Rather than being composed of flat, centralized processors, nervous systems appear to be stratified into interacting levels, where newer structures modulate, override, or subsume the activity of more primitive ones[7].

An influential formulation of this idea in robotics was Brooks' subsumption architecture, which challenged classical cognitivist models[8] by showing that complex behavior can emerge from the hierarchical coordination of simple behavioral modules (Brooks, 1986). In this architecture, higher-order processes inhibit or modulate lower ones to produce adaptive responses without relying on internal world models. Control structures emerge

[7] This idea, often referred to as "Jacksonian dissolution" was first proposed by John Hughlings Jackson in the late 19th century. He argued that higher nervous functions evolved by building upon more basic ones, and that neurological damage reveals this hierarchy through the sequential unmasking of lower reflexes (Jackson, 1884).

[8] Cognitivist models in classical Artificial Intelligence and psychology typically posit that perception, planning, and action occur in sequential stages mediated by symbolic representations of the world. Brooks' approach rejected this representational bottleneck in favor of direct perception – action coupling.

through constraint closure : early-developing substrates constrain the space in which later layers self-organize, while higher cortical systems in turn refine and modulate lower ones (Wilson and Prescott, 2022). This mutual constraint, observed in both phylogeny and ontogeny, accounts for the robust and adaptive nature of layered control. Building on these ideas, Prescott et al. (Prescott et al., 1999) developed biologically grounded control architectures that mirror the stratified organization of the vertebrate nervous system. These insights laid the foundation for behavior-based robotics, where layered control enables situated agents to respond flexibly and robustly to their environment.

Developmental and evolutionary accounts further support the idea that functional hierarchy reflects neuroanatomy. Vertebrate nervous systems exhibit a similar stratification: spinal circuits govern reflexive withdrawal, hindbrain structures mediate startle responses, midbrain and hypothalamus orchestrate species-specific defensive reactions, and cortical areas modulate or suppress these outputs depending on context and internal state. A well-studied case is mammalian defensive behavior, schematized in Figure 4, where behavioral control spans from sensory input to motor and hormonal output, with progressively more abstract or deliberative mediation at higher levels (Prescott et al., 1999).

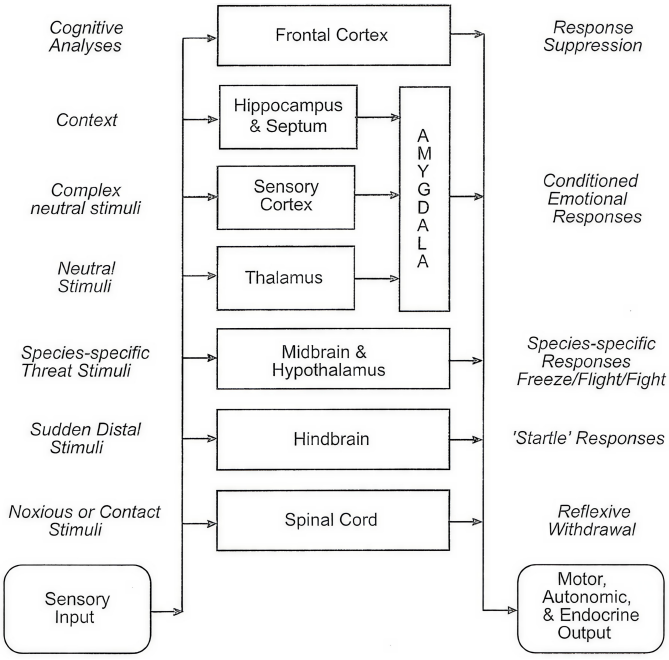

Figure 4: Subsumption architecture of mammalian defensive behaviors. Behavioral control is hierarchically layered in nested sensorimotor loops, involving increasingly complex neural structures and governing progressively more abstract or context-sensitive responses. Reproduced from Prescott et al., 1999.

Beyond descriptive similarity, such hierarchies have been formalized as functional architectures. The Distributed Adaptive Control (DAC) framework, for example, posits three core levels – reactive, adaptive, and contextual – that mirror both neuroanatomical

substrates and computational roles (Verschure, 2012). It has been implemented in neurorobotic agents, enabling integration of low-level perception–action loops with higher-order planning and internal modeling.

These frameworks can be further unified by aligning them with developmental and evolutionary taxonomies. As discussed in Table 1, Dennett's evolutionary classification of minds, LeDoux's taxonomy of defensive behavior, and Verschure's control layers converge on a shared hierarchy of increasing abstraction. Each layer builds upon those below, adding representational depth and temporal foresight – a hierarchy that is reflected in the anatomical layering of the brain and recapitulated in ontogeny.

Cognitive Evolution (D. Dennett)	Behavioral Taxonomy (J. LeDoux)	Cognitive Hierarchy (P. Verschure)
2*Darwinian	Reflexive	Soma
	Reactive	Reactive
2*Skinnerian	Habitual	2*Adaptive
	Action–Outcome contingencies	
Popperian	Subconsciously deliberative	2*Contextual
Gregorian	Consciously deliberative	

Table 1: The cognitive–behavioral hierarchy. Cognitive evolutionary grades by D. Dennett. A behavioral taxonomy of defensive behaviors by J. LeDoux. The layers of a general cognitive architecture by P. Verschure.

Importantly, emotion functions as a modulatory process spanning the control hierarchy, shaping behavior selection across levels by integrating immediate sensorimotor inputs with long-term motivational priorities. This modulation helps resolve competing drives and enables context-sensitive action (Maselli et al., 2023; Ziemke, 2008).

In summary, vertical functional hierarchies – from reflexive to deliberative – are not just modeling conveniences but biological realities. Their layered logic enables integration of behavior across levels of abstraction, while remaining grounded in neuroanatomical structure. This vertical organization complements the horizontal modularity of survival circuits operating in parallel, discussed in the previous section. Together, these dimensions define a dual-axis framework of control: one that is structurally grounded, functionally distinct , and computationally tractable.

1.2.3 Repertoires, patterns, and taxonomies of behavior

To populate the hierarchical control architectures outlined above with behavior-based modules, one must specify how structured behavioral repertoires are extracted, formalized, and integrated into models. This section traces two complementary research programs that pursue this aim from different angles: one grounded in experimental ethology and large-scale behavioral data, the other grounded in theoretical classifications and formal ontology[9].

Recent advances in imaging, recording, and data processing, such as pose-tracking and automated annotation tools (see subsection 2.2.1), have enabled the high-throughput capture of behavior in freely moving organisms in naturalistic or lab conditions. In *Drosophila*, for instance, long-term recordings combined with dimensionality reduction techniques have revealed modular subunits of locomotor behavior – interpreted as stereotyped behavioral building blocks – whose transitions can be modeled using probabilistic and dynamical systems frameworks (Berman et al., 2014, 2016; Katsov et al., 2017). Unsupervised clustering methods have been central in extracting these motifs, bypassing experimenter bias and enabling the construction of state-transition diagrams that uncover latent organizational principles.

This data-driven methodology, increasingly formalized under the label of computational neuroethology, aims to dissect and annotate behavior from time-series data using tools such as unsupervised learning, manifold embedding, and dynamical systems analysis (Datta et al., 2019; Mobbs et al., 2021; Robson and Li, 2022). These approaches provide behavioral decompositions that are not only reproducible but also amenable to cross-modal alignment with neural activity, thereby serving as a bridge between observed action and underlying neural substrate. In this way, behavioral microstructure becomes experimentally tractable, offering a bottom-up route for populating modular control architectures.

Parallel to these empirical efforts, conceptual taxonomies have been proposed that classify behavior based on internal structure, action-outcome dependencies, and degree of flexibility. In particular, taxonomies rooted in theoretical neuroscience and psychology distinguish reflexive, reactive, goal-directed, and habitual behaviors, often framed in terms of the innate-learned spectrum and the nature of stimulus-response coupling (Ledoux and Daw, 2018). These categories serve not only classificatory roles but also inform computational architectures by suggesting distinct layers or modules associated with different levels of control abstraction.

A further line of formalization stems from ontological work aimed at codifying behavioral phenotypes across species. Ontologies such as the Worm Phenotype Ontology, the *Drosophila* Phenotype Ontology, and the Neurobehavior Ontology define standardized

[9] An ontology specified by logical axioms in a formal language. It aims to provide an unbiased foundation for representing reality, guide the development of domain- and application-specific ontologies and avoid conceptual errors in large-scale modeling.

vocabularies and relationships between behavioral terms, enabling computational queries and inter-species comparisons (Gkoutos et al., 2012; Osumi-Sutherland et al., 2013; Schindelman et al., 2011). These ontologies link behavioral descriptions to genetic, cellular, and neural substrates, offering a conceptual infrastructure for unifying behavior with other biological data types. While their integration into mechanistic modeling is still limited, they represent a promising scaffold for cross-disciplinary interoperability.

Despite these advances, challenges remain in capturing temporal organization, intentional binding, and the sequential nature of behavior. Concepts from theoretical biology, such as operational closure[10], have been proposed as tools for defining the initiation and termination conditions of behavioral units – though these require further formalization and are best treated with caution in empirical contexts. Overall, the convergence of empirical and formal approaches offers a promising pathway for modeling behavioral repertoires in a biologically grounded yet computationally tractable manner.

[10] Operational closure refers to a process being defined and maintained through its own internal operations, such that its continuation depends on, and recursively reinforces, its own structure and function.

2 Holistic behavioral modeling

Building on the interdisciplinary landscape outlined in Chapter 1, this chapter shifts focus from the broader context of behavioral research to the specific orientation adopted in this thesis. It first outlines the commitments that motivate a holistic and mechanistic approach centered on behavior, then it justifies the choice of the model organism and behavioral domain as a tractable and illustrative case study, and finally sets out the concrete objectives pursued.

2.1 Conceptual commitments

The title *Holistic Behavioral Modeling* reflects the convergence of three conceptual commitments. "Holistic" marks a perspective that centers the organism not as a bundle of parts or isolated mechanisms, but as a unified, situated, and self-regulating system whose viability is continually at stake. "Behavioral" sets the target of inquiry : normatively regulated sensorimotor processes at the interface of internal needs and environmental demands. "Modeling" reflects mechanistic thinking, not simply as a technical method, but as an epistemological strategy.

This section dives into the philosophical foundation and methodological motivation of these commitments. First, it outlines what is meant by a holistic view of the organism, and why such a view is necessary for modeling purposive biological systems. Second, it frames behavior as a modular, flexible and empirically accessible modeling target. Third, it motivates the use of mechanistic models as a means of generating explanations: formal constructs that aim to capture how observed patterns emerge from organized activity under normative constraints. These commitments shape both the scope and the internal logic of the modeling work that follows.

2.1.1 What does it mean to think holistically?

Holistic thinking in the life sciences arises from the tradition of holism, which opposes reductionism not through a vague appeal to "complexity", but by emphasizing integration, dynamic organization, and context-dependence over part-based decomposition.

This perspective has gained renewed prominence in what's called *the return of the organism* – not as nostalgic metaphor, but as recognition of the organism's central role in explaining biological agency, purposiveness and reciprocal organism-environment causation (Baedke et al., 2021; Fábregas-Tejeda et al., 2024). It marks a shift from explanations rooted solely in genes or networks toward those that treat the organism as a unit of action, constraint, and sense-making. This shift is driven by the growing realization that reductionist methods often overlook the interplay between metabolic, behavioral, and environmental dynamics. A holistic approach views organisms not as passive carriers of genetic information, but as autonomous, self-regulating systems that respond adaptively to their environments.

This epistemic reframing of the organism also has implications for how scientific models are structured. Whole-organism approaches resist the fragmentation of knowledge into disciplinary silos – a condition often lamented in contemporary biological research (Baedke, 2018). Instead of treating molecular biology, neuroscience, and ethology as discrete domains, holistic frameworks seek integrative models that bridge across these levels. This is not merely a methodological ambition, but an ontological claim: that the organism itself is the proper locus of integration. In this light, the idea of starting from the middle – that is, using the organism as the conceptual midpoint for organizing explanation – becomes a principled alternative to reductive decomposition.

Ultimately, to think holistically is to adopt the organism as the primary unit of explanation – not simply because it is complex, but because it anchors the interplay between functional integration, adaptive regulation, and environmental coupling.

2.1.2 Behavior as a modeling anchor

Behavior constitutes a privileged modeling anchor for biological and artificial agents alike. It connects internal regulation with external engagement and serves as a tractable unit of analysis across levels – from sensorimotor control to systemic viability. This section reframes behavior not as an externally triggered output, but as a temporally extended, adaptively modulated, sensorimotor process embedded in the organizational logic of the system.

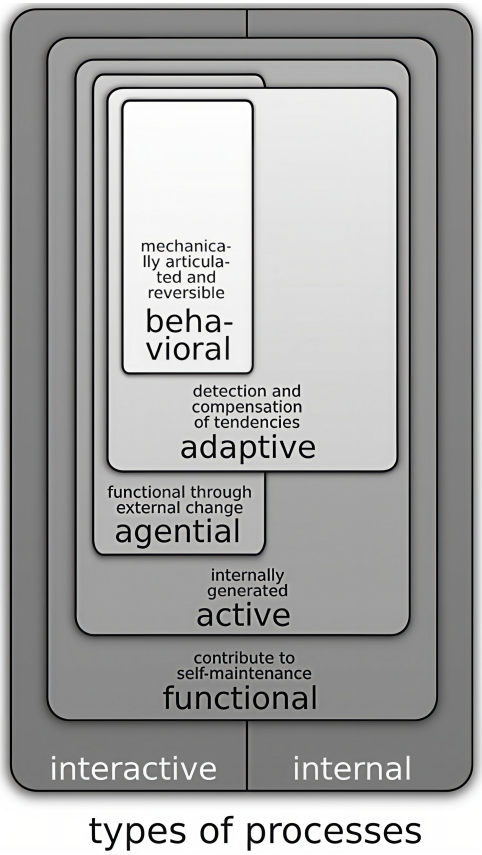

Figure 5: Life processes can be categorized along internal and interactive dimensions, forming a nested hierarchy from self-maintaining processes to those that are active, agential, and adaptive. Behavior is defined as a specific class of processes that are mechanically articulated and reversible interactions with the environment. Reproduced from X. Barandiaran and Moreno, 2008.

Behavior as a process Process-oriented philosophical accounts have attempted to define behavior as a specific kind of process.

Figure 5 summarizes one such proposal, where behavior is situated within a nested hierarchy of life processes, spanning internal and interactive domains (on further distinction within the behavioral set see subsection 1.2.3). Behavioral processes differ from other adaptive processes (e.g., metabolic regulation) in that they involve the articulation of a body in direct interaction with the environment. This makes them especially suitable for modeling in interaction-focused frameworks, due to their selective deployment

and functional flexibility. Furthermore they are internally generated and normatively constrained: they occur under the pressure of maintaining the system's identity and viability (X. E. Barandiaran and Egbert, 2014).

Sensorimotor coupling The primary characteristic of behavior is the dynamic coupling of sensory and motor modalities. From the perspective of enactivist cognitive science, the brain does not process representations of the world but participates in cycles of perception and action that stabilize the organism's relation to its surroundings (Di Paolo et al., 2017). Sensorimotor contingencies emerge through the continuous interplay between neural dynamics, bodily structure, and environmental features (Buhrmann et al., 2013). Sensorimotor coupling is thus the operational layer through which internal state regulation is enacted in real time. Rather than mediating between stimulus and response, it configures the system's mode of interaction with the world. Habits, as stabilized sensorimotor routines, illustrate how repeated interactions shape the topology of the agent's behavioral landscape (X. E. Barandiaran, 2017).

Temporal structure and operational closure Behavior unfolds as a sequence of functionally distinct phases, each defined by its degree of exclusivity, temporal coupling, and interruptibility. For example, walking and standing are mutually exclusive, while chewing can overlap with either. In goal-directed sequences like foraging, transitions from searching to approaching to feeding are marked by cues indicating phase completion or disruption. These transitions reflect a structure of intermediate and final closures – functional endpoints that delimit each interruptible and goal-sensitive subroutine and guide the progression toward goal satisfaction (X. Barandiaran and Moreno, 2008).

Behavior-based robotics These insights align with the modeling tradition of behavior-based robotics (see subsection 1.2.2), where behavior is used as the primary unit of simulation and control. In this paradigm, agents perform modular behaviors, each implemented as a sensorimotor routine with local control logic (X. E. Barandiaran and Chemero, 2009). Instead of representing goals abstractly, such architectures express functionality through the emergent organization of low-level behaviors.

This perspective legitimizes behavior as a modeling anchor, not only because it is empirically accessible, but because it expresses the recursive and embodied nature of autonomous systems. Behavior, in this sense, functions as a modeling pivot between embodiment and normativity, between mechanism and function, supporting both the explanatory and exploratory roles of theory building.

2.1.3 Modeling as mechanistic explanation

In contemporary behavioral science, computational models serve not only predictive or classificatory purposes but increasingly contribute to mechanistic explanation. This shift

reflects a broader epistemological orientation: from forecasting outcomes to understanding the organized structure of the systems that generate them.

Mechanistic models are designed to explain a phenomenon by identifying the relevant entities, their activities, and the organizational structure that produces the target behavior. According to a widely cited formulation, "a mechanism is a structure performing a function in virtue of its component parts, component operations, and their organization" (Bechtel and Abrahamsen, 2005). The explanandum is accounted for by the orchestrated functioning of these components, which may stand in causal or constitutive relation to the phenomenon (Craver, 2006). Mechanistic models thus aim to answer "how" questions by decomposing the system and localizing explanatory responsibility in its internal architecture.

In this sense, explanation is not equated with empirical adequacy or statistical correlation. It is tied to ontological commitments about what exists and how it interacts. For a model to be explanatory, it must map its structural assumptions onto plausible mechanisms in the world. This includes specifying component parts, delineating their activities, and making explicit how their configuration gives rise to the phenomenon under investigation. Scientific explanation, then, is not merely a narrative or summary of observations, but a structured mapping between explanandum and explanans, grounded in the internal logic of the modeled mechanism (Kaiser and Krickel, 2017).

This modeling philosophy supports and requires continuous interaction with experimentation. Empirical findings inform the specification of mechanistic models, which in turn generate new testable claims and help refine experimental design. This iterative loop allows modeling and experimentation to act as mutually corrective and generative processes. A model that fails to reproduce empirical data may reveal gaps in its assumptions or point to overlooked causal pathways. Conversely, a model that generates unexpected behavior may guide researchers toward new observational targets or experimental perturbations.

In addition to their explanatory power, mechanistic models play a crucial hypothesis-generating role. Unlike statistical models that merely describe or interpolate data, mechanistic models open a structured design space. They allow researchers to explore counterfactual configurations – what would happen if certain interactions were removed, modified, or introduced – and to simulate the consequences of perturbations before these are physically tested. This capacity makes them indispensable tools for theoretical exploration, particularly in domains where direct experimental access is limited or where system complexity hinders exhaustive empirical coverage.

To summarize, mechanistic modeling is not a supplement to experimentation but a distinct and complementary epistemic strategy. It enables explanation by revealing how phenomena emerge from the organized activity of parts, generates hypotheses by exploring structural alternatives, and structures scientific discourse by making mechanistic assumptions explicit and testable. The explanatory strength of such models is not a by-product of fit-to-data but a function of their internal coherence, plausibility, and capacity to support

inference about the system's organization.

2.1.4 Grounding the modeling scope

The scope of this modeling approach is defined by four principles that operationalize the preceding commitments and establish the methodological foundation on established modeling paradigms :

(i) Agent-centered modeling : A holistic approach requires centering the organism as an integrated unit – internally structured, environmentally embedded, and behaviorally autonomous. This aligns with ABM, where agents possess internal dynamics and interact locally to produce system-level patterns.

(ii) Comprehensive behavioral scope : Behavior is modeled across its full hierarchy and range, from motor primitives to adaptive action sequences. In line with behavior-based modeling, this repertoire is implemented as a modular, hierarchical control system.

(iii) Cross-timescale integration : Temporally embedded processes – neural, behavioral, and metabolic – levels are jointly modeled. This supports mechanistic explanation and reflects the convergence of neuroscience and ecology around lifespan-level modeling.

(iv) Homeostatic regulation : Behavior is framed as serving the organism's precarious metabolic organization, enabling it to satisfy normative demands, sustain viability and advance through developmental stages.

These principles jointly specify a modeling stance focused on the behaving organism as an embodied, embedded, self-regulated agent.

2.2 Model system and target behavior

This section grounds the modeling approach in a concrete biological context. Rather than operating with abstract agents, the focus is placed on an organism where behavioral complexity, internal structure, and experimental access are sufficiently aligned to support explanatory modeling.

The focus is not only on which organism is modeled, but also on which behavioral domain. The aim is to capture a process that is structured, state-dependent, and ecologically meaningful, while remaining tractable for formalization. What follows is a brief justification for both the model system and the target behavior chosen.

2.2.1 Why choose the *Drosophila* larva?

The *Drosophila melanogaster* larva has emerged as an indispensable model organism at the intersection of developmental biology, behavioral neuroscience, and computational modeling. Its utility stems not from simplicity alone, but from a set of properties that enable an unusually direct mapping between neural circuits, behavior, and physiology.

From a genetic perspective, the larva inherits the powerful molecular toolkit of adult *Drosophila*, but with developmental and anatomical properties that simplify experimen-

tal targeting. Techniques such as the GAL4/UAS binary system (Duffy, 2002) allow for spatially and temporally precise control of gene expression, enabling circuit-level manipulations including activation, silencing, and ablation of selected neuronal populations (Wu et al., 2003). These tools are especially effective in larvae due to their transparent cuticle, facilitating optical access for imaging and optogenetics. This tractability supports causal experimentation on the neural basis of behavior at a resolution rarely attainable in vertebrate models.

A range of specialized tools support the high-resolution tracking and annotation of *Drosophila* larval behavior. FTIR-based imaging platforms such as FIM and its multicolor extension FIM2c enable high-contrast visualization of body-substrate contact (Risse, Berh, et al., 2017, Risse, Otto, et al., 2017). Multi-modal systems like i2Bot (X. Sun et al., 2025) and the modular Ethoscope (Geissmann et al., 2017) allow environmental manipulation and high-throughput recording. Behavioral segmentation and classification can be performed using machine learning pipelines such as JAABA (Kabra et al., 2013) and DeepLabCut (Nath et al., 2019). Additional tools focus on developmental annotation, such as PEDtracker (Schumann and Triphan, 2020) which enables automatic larval staging along with locomotor profiling across instars, or closed-loop experimentation (Tadres and Louis, 2020). Together, these platforms provide a cohesive infrastructure for quantifying behavior and validating computational models, forming a critical empirical interface for computational neuroethology (see subsection 1.2.3).

Perhaps most significantly, the *Drosophila* larva is one of the very few model organisms for which a near-complete connectome at synaptic resolution has been reconstructed across the entire central nervous system (Ohyama et al., 2015, Schneider-Mizell et al., 2016, Winding et al., 2023). This allows for an exceptional level of continuity between structural connectivity and behavioral modeling. Feeding-related circuits have been particularly well mapped, revealing how both monosynaptic and polysynaptic sensory inputs converge onto shared sets of motor and neuroendocrine outputs (Miroschnikow et al., 2018, Hückesfeld et al., 2021). These include neuromodulatory components such as serotonergic and peptidergic neurons, making it possible to dissect sensorimotor transformations underlying value-based behaviors at circuit level. Such completeness supports mechanistic models that integrate sensory topography, circuit motifs, and behavioral function – an opportunity rarely available in more complex nervous systems.

In comparative context, the larva occupies a productive middle ground. Unlike *C. elegans*, which has a fully mapped but minimal neural architecture of around 300 neurons (Cook et al., 2019), the *Drosophila* larva possesses approximately 10,000 neurons – enough to support distributed control, sensory integration, and internal state modulation, yet still manageable for connectomic reconstruction and functional mapping. Compared to vertebrates, its short generation time, optical accessibility, and the absence of strict animal handling regulations – such as those mandated for vertebrate care – make it a more scalable and accessible system for behavioral experimentation.

This convergence of experimental tractability, behavioral richness, and structural com-

pleteness affords an exceptional continuity between structure and function, experiment and simulation – making the *Drosophila* larva an ideal anchor for integrative studies that span from neural circuits to embodied behavior.

2.2.2 Foraging as a structured behavioral process

The *Drosophila* larva exhibits a compact yet functionally rich behavioral repertoire that enables it to adaptively navigate complex, semi-structured environments. Core behaviors include chemotaxis, thermotaxis, phototaxis, anemotaxis, mechanosensory avoidance, and local exploration, all orchestrated through a small number of identifiable sensorimotor primitives, such as intermittent crawling and lateral bending in a head casting or weathervaning mode (Humberg et al., 2018, Jovanic et al., 2019, Hernandez-Nunez et al., 2021).

Foraging behavior in particular involves the interplay of navigational strategies with exploitative feeding behavior. Larvae switch between different exploration modes depending on stimulus reliability and internal state: local or expanded exploration, substrate sampling, or gradient detection and following – strategies that echo gradient-based navigation in flies and bees, and short-range trail tracking as seen in ants. (Tastekin et al., 2018, X. Sun et al., 2020, Paisios et al., 2017). These actions are not mere reflexes but components of a temporally extended decision-making process, shaped by recent experience, metabolic state and structured spatial dynamics.

Recent work has shown that internal states such as satiety, arousal, or developmental stage configure both stimulus sensitivity and motor output (Vogt et al., 2021,Zhu et al., 2021). For example, hungry larvae increase turning frequency and extend search duration in odor-sparse conditions, a pattern that has been reproduced in embodied computational models of foraging (Maffei et al., 2015, Adden et al., 2022).

The ethological richness and algorithmic structure of larval foraging make it particularly well-suited for generative modeling. Unlike stereotyped behaviors, foraging reveals the larva's capacity for adaptive sequencing, contextual switching, and sensory-driven action selection. This makes it an ideal target for studying how embodied agents can operate in dynamic environments while balancing exploration and exploitation according to their metabolic needs.

Crucially foraging will not be treated here as a behavior optimized in evolutionary time, along the lines of optimal foraging theory[11]. Rather it will be regarded as a process under normative homeostatic control that has to be "good enough" to promote viability in precarious conditions.

[11] Optimal foraging theory is a behavioral ecology model that helps predict how an animal behaves when searching for food. It posits that animals adopt foraging strategies that maximize net energy gain by optimizing the trade-off between energy obtained and the time and energy spent obtaining it, within environmental constraints.

2.3 Objectives of the thesis

The overall thesis scope unfolds across three objectives:

(i) Design of the Behavioral Architecture - Dynamic Energy Budget (BA-DEB) framework for *Drosophila* larva foraging

(ii) Implementation of BA-DEB in *Larvaworld*, a simulation and analysis platform

(iii) Development of standalone mechanistic models of *Drosophila* larva behavioral domains

In what follows each of these objectives will be briefly described.

Framework design At the conceptual level the thesis proposes a structured computational architecture that captures the hierarchical organization of behavior in a behaving organism. This BA formalizes behavior as the output of nested sensorimotor loops, from low-level motor primitives to mid-level reactive behaviors and high-level adaptive modulation. It provides a modular, extensible framework for simulating complex, state-dependent behavior in closed-loop environments. The architecture supports integration of diverse models – neural, statistical, or rule-based – into a coherent whole-organism system. To anchor behavioral control in physiological constraint, the BA is coupled with a DEB model that supplies an internal hunger/satiety signal derived from simulated energy dynamics. This bidirectional coupling enables the study of behavior as a homeostatically regulated process. The resulting BA-DEB framework serves as a concrete, generative platform for embedding and comparing mechanistic hypotheses across levels of organization and timescales.

Software implementation The BA-DEB framework is implemented in *Larvaworld*, an open-source simulation and analysis platform designed to foster collaboration between experimentalists and modelers. *Larvaworld* supports agent-based simulations in closed-loop settings, where virtual larvae – as BA-DEB agents – interact dynamically with their environment. The platform includes standardized pipelines for data analysis and modular modeling components that can be adapted across organisms. It allows for direct comparisons between real and simulated behavior, thereby enabling empirical validation and model refinement. Designed with usability and educational value in mind, it aims to minimize technical barriers and promote methodological convergence across disciplinary boundaries, offering a shared workspace for developing, testing, and communicating behavioral models.

Standalone mechanistic models The thesis also presents some standalone mechanistic models targeting distinct, previously overlooked domains of *Drosophila* larva behavior. Two of them are presented in detail, each advancing a novel mechanistic hypotheses, while additional contributions are briefly summarized in the context of collaborative research efforts. First, a stochastic network model capturing the alternation between activity and

rest observed in freely moving larvae. This model reproduces key statistical regularities and suggests a mechanistic hypothesis for behavioral intermittency. Second, a coupled-oscillator model addressing the interaction between forward crawling and lateral bending. Based on kinematic analysis showing that these motor primitives are not independent, it attributes their interference to biomechanical constraints. These two components are ultimately brought together in the intermittent coupled-oscillator model for larva locomotion.

Part II
Original research

This part of the thesis presents original research conducted during the doctoral project. The core of this work is presented in the form of three first-author studies, which are attached in their original format and constitute Chapters 3, 4 and 5. These chapters correspond to the following papers:

- A plausible mechanism for *Drosophila* larva intermittent behavior, (Sakagiannis et al., 2020) (Chapter 3),
- A behavioral architecture for realistic simulations of *Drosophila* larva locomotion and foraging, (Sakagiannis, Jürgensen, and Nawrot, 2025) (Chapter 4),
- *Larvaworld*: A behavioral simulation and analysis platform for *Drosophila* larva, (Sakagiannis, Rapp, et al., 2025) (Chapter 5).

In addition, selected co-authored studies are summarized in Chapter 6. These collaborative works correspond to the following papers:

- Prediction error drives associative learning and conditioned behavior in a spiking model of *Drosophila* larva, (Jürgensen et al., 2024) (section 6.1),
- Evolution of temperature preference behaviour among *Drosophila* larvae, (Kafle et al., 2025) (section 6.2),
- Feeding-state dependent neuropeptidergic modulation of reciprocally interconnected inhibitory neurons biases sensorimotor decisions in *Drosophila*, (de Tredern et al., 2024) (section 6.3).

The chapter and section titles are identical to the original titles of the corresponding studies. Each study – attached or summarized – is preceded by a page containing bibliographic metadata and a statement of author contributions. For the collaborative works, a brief note of the author's personal involvement is included. Note that since the full papers are attached in their original format, the reader will observe differences in typesetting, section structure, and page numbering between these chapters and the rest of the thesis, as well as reference link malfunctioning. Also, the bibliography of these papers is separate and not included in the thesis' bibliography.
The three papers, presented in full, constitute the core research work of the thesis, forming a coherent progression from mechanistic hypothesis generation to architectural design and software implementation. The contributions to collaborative works reflect the interdisciplinary outreach of the thesis and embody its broader orientation toward integration and collaboration across scientific domains. Collectively, the research presented in this part

will be revisited and contextualized in Part III with respect to the overarching objectives outlined in section 2.3 and integrated into a unified discussion of modeling principles, explanatory strategies, and future directions.

3 A plausible mechanism for *Drosophila* larva intermittent behavior

Authors: Panagiotis Sakagiannis, Miguel Aguilera, and Martin Paul Nawrot
Journal: Biomimetic and Biohybrid Systems, Living Machines 2020
Date: December 2020
DOI: 10.1007/978-3-030-64313-3_28

Author Contributions

Contributions to the publication were as follows: Conceptualization was carried out by **P.S.**, M.A., and M.P.N.; data curation, investigation, visualization, and original draft writing were performed by **P.S.**; formal analysis and methodology were developed by **P.S.** and M.A.; funding acquisition and supervision were provided by M.P.N.; and all authors (**P.S.**, M.A., M.P.N.) contributed to review and editing.

A plausible mechanism for *Drosophila* larva intermittent behavior ⋆

Panagiotis Sakagiannis[1][0000−0002−1033−5387], Miguel Aguilera[2][0000−0002−3366−4706], and Martin Paul Nawrot[1][0000−0003−4133−6419]

[1]Computational Systems Neuroscience, Institute of Zoology, University of Cologne, Germany
[2]IAS-Research Center for Life, Mind, and Society. University of the Basque Country. Donostia, Spain.
p.sakagiannis@uni-koeln.de
http://computational-systems-neuroscience.de/

Abstract. The behavior of many living organisms is not continuous. Rather, activity emerges in bouts that are separated by epochs of rest, a phenomenon known as intermittent behavior. Although intermittency is ubiquitous across phyla, empirical studies are scarce and the underlying neural mechanisms remain unknown. Here we reproduce empirical evidence of intermittency during *Drosophila* larva free exploration. Our findings are in line with previously reported power-law distributed rest-bout durations while we report log-normal distributed activity-bout durations. We show that a simple conceptual stochastic model can transition between power-law and non-power-law distributed states and we suggest a plausible neural mechanism for the alternating rest and activity in the larva. Finally, we discuss possible implementations in behavioral simulations extending spatial Levy-walk or coupled-oscillator models with temporal intermittency.

Keywords: larva crawling · Levy-walks · neuronal avalanches.

1 Introduction

The search for statistical regularities in animal movement is a predominant focus of motion ecology. Random walks form a broad range of models that assume discrete steps of displacement obeying defined statistical rules and acute reorientations. A Levy walk is a random walk where the displacement lengths and the respective displacement durations are drawn from a heavy-tailed, most often

⋆ Supported by the Research Training Group 'Neural Circuit Analysis' (DFG-RTG 1960, grant no. 233886668) and the Research Unit 'Structure, Plasticity and Behavioral Function of the *Drosophila* mushroom body'(DFG-FOR 2705, grant no. 403329959), funded by the German Research Foundation. M.A. was funded by the UPV/EHU post-doctoral training program ESPDOC17/17 and H2020 Marie Skłodowska-Curie grant 892715, and supported in part from the Basque Government (IT1228-19).

a power-law distribution. When considered in a 2D space reorientation angles are drawn from a uniform distribution. This initial basic Levy walk has been extended to encompass distinct behavioral modes bearing different go/turn parameters, thus termed composite Levy walk. Levy walks have been extensively studied in the context of optimal foraging theory. A Levy walk with a power-law exponent between the limit of ballistic ($\alpha = 1$) and brownian motion ($\alpha = 3$) yields higher search efficiency for foragers with an optimum around $\alpha = 2$ when search targets are patchily or scarcely distributed and detection of a target halts displacement (truncated Levy walk) [5].

Nevertheless, the underlying assumption of non-intermittent movement flow of Levy walk models complicates the identification of the underlying generative mechanisms as they focus predominantly on reproducing the observed trajectories, neglecting the temporal dynamics of locomotory behavior. Therefore, Bartumeus (2009) stressed the need for a further extension termed intermittent random walk, emphasizing the integration of behavioral intermittency in the theoretical study of animal movement [2]. Here we aim to contribute to this goal by studying the temporal patterns of intermittency during *Drosophila* larva free exploration in experimental data and conceptual models, bearing in mind that power-law like phenomena can arise from a wide range of mechanisms, possibly involving processes of different timescales [5]. While our study remains agnostic towards whether foragers really perform Levy walks - a claim still disputed [5] - we suggest that intrinsic motion intermittency should be taken into account and the assumption of no pauses and acute reorientations should be dropped in favor of integrative models encompassing both activity and inactivity.

Drosophila larva is a suitable organism for the study of animal exploration patterns and the underlying neural mechanisms. A rich repertoire of available genetic tools allows acute activation, inhibition or even induced death of specific neural components. Crawling in 2D facilitates tracking of unconstrained behavior. Also, fruit flies during this life stage are nearly exclusively concerned with foraging. Therefore a food/odor-deprived environment can be largely considered stimulus-free, devoid of reorientation or pause sensory triggers, while target-detection on contact can be considered certain. Truncated spatial Levy-walk patterns of exploration with exponents ranging from 1.5 to near-optimal 1.96 that hold over at least two orders of magnitude have been previously reported for the *Drosophila* larva. The turning-angle distribution, however, was skewed in favor of small angles and a quasi-uniform distribution was observed only for reorientation events $\geq 50°$ [9,8]. Moreover, it has been shown that these patterns arise from low-level neural circuitry even in the absence of sensory input or brain-lobe function and have therefore been termed 'null movement patterns' [9,8].

Previous studies on larva intermittent locomotory behavior have concluded that the distribution of durations of rest bouts is power-law while that of activity bouts has been reported to be exponential [12] or power-law [8]. Genetic intervention has revealed that dopamine neuron activation affects the activity/rest ratio via modulation of the power-law exponent of the rest bouts, while the dis-

tribution of activity bouts remains unaffected. This observation hints towards a neural mechanism that generates the alternating switches between activity and rest where tonic modulatory input from the brain regulates the activity/rest balance according to environmental conditions and possibly homeostatic state.

Here we analyze intermittency in a large experimental dataset and present a conceptual model that generates alternation between rest and activity, capturing the empirically observed power-law and non-power-law distributions. We discuss a plausible neural mechanism for the alternation between rest and activity and the regulation of the animal's activity/rest ratio via modulation of the rest-bout power-law exponent by top-down modulatory input. Our approach seeks to elaborate on the currently prevailing view that these patterns result from intrinsic neural noise [8].

2 Materials and Methods

2.1 Experimental dataset

We use a larva-tracking dataset available at the DRYAD repository, previously used for spatial Levy-walk pattern detection [9]. The dataset consists of up to one hour long recordings of freely moving larvae tracked as a single point (centroid) in 2D space. We consider three temperature-sensitive shibire^{ts} fly mutants allowing for inhibition of mushroom-body (MB247),brain-lobe/SOG (BL) or brain-lobe/SOG/somatosensory (BLsens) neurons and an rpr/hid mutant line inducing temperature-sensitive neuronal death of brain-lobe/SOG/somatosensory (BLsens) neurons. Each mutant expresses a different behavioral phenotype when activated by 32°-33° C temperature. We compare phenotypic behavior to control behavior in non-activated control groups. A reference control group has been formed consisting of all individuals of the four 32°-33° C control groups (Tab. 1).

T (°C)	treatment	larvae (#)	tracking time (h)	rest bouts (#)	activity bouts (#)	activity ratio (mu±sigma)
22	BLsens > rpr/hid$_{control}$	33	19.9	869	559	0.9 ± 0.21
32	BLsens > rpr/hid	21	17.16	2251	1504	0.6 ± 0.23
33	MB247/+	16	12.28	1067	666	0.72 ± 0.2
33	MB247 > shi^{ts}	16	13.43	1487	1021	0.64 ± 0.23
22	shi^{ts} /+	19	14.82	1370	861	0.77 ± 0.24
33	shi^{ts} /+	21	16.85	1191	768	0.75 ± 0.19
22	BL/+	17	12.5	570	387	0.87 ± 0.22
33	BL/+	17	10.81	391	286	0.87 ± 0.19
33	BL > shi^{ts}	14	12.78	879	629	0.82 ± 0.22
33	BLsens > shi^{ts}	10	10.53	1553	1007	0.44 ± 0.25
32-33	Reference control	87	59.84	3519	2279	0.83 ± 0.21

Table 1. Dataset description and empirical results for rest/activity bout analyses.

For the present study recordings longer than 1024 seconds have been selected. Instances where larvae contacted the arena borders were excluded. The raw

time series of x,y coordinates have been forward-backward filtered with a first-order butterworth low-pass filter of cutoff frequency 0.1 Hz before computing the velocity. The cutoff frequency was selected as to preserve the plateaus of brief stationary periods while suppressing the signal oscillation due to peristaltic-stride cycles. Velocity values ≥ 2.5 mm/sec have been discarded to account for observed jumps in single-larva trajectories that are probably due to technical issues during tracking. This arbitrary threshold was selected as an upper limit for larvae of length up to 5mm, crawling at a speed of up to 2 strides/sec with a scaled displacement per stride of up to 0.25.

2.2 Bout annotation and distribution

In order to designate periods of rest and activity we need to define a suitable threshold V_θ in the velocity distribution as in [12]. We used the density estimation algorithm to locate the first minimum $V_\theta = 0.085mm/sec$ in the velocity histogram of the reference control group. A rest bout is then defined as a period during which velocity does not exceed V_θ. Rest bouts necessarily alternate with periods termed activity bouts. The bout annotation method is exemplified for a single larva track in Fig. 2.

To quantify the duration distribution of the rest and activity bouts we used the maximum likelihood estimation (MLE) method to fit a power-law, an exponential and a log-normal distribution for each group as well as for the reference control group. Given the tracking framerate of 2 Hz and the minimal tracking time of 1024 seconds, we limited our analysis to bouts of duration 2^1 to 2^{10} seconds. The Kolmogorov-Smirnov distance D_{KS} for each candidate distribution was then computed over 64 logarithmic bins covering this range. Findings are summarized in Tab. 2 for the rest bouts and in Tab. 3 for the activity bouts.

3 Results

The results section is organized as follows. Initially we present a simple conceptual two-state model transitioning autonomously between power-law and non-power-law regimes. Next we attempt to reproduce existing findings [12] by studying intermittency during larva free exploration in a different dataset [9]. Finally we compare mutant and control larva phenotypes in the context of intermittency.

3.1 Network model of binary units reproduces larval statistics of intermittent behavior

Previous work on *Drosophila* larva intermittent behavior reported that rest-bout durations are power-law distributed while activity-bout durations are exponentially distributed [12]. Our first contribution is to define a simple model displaying how this dual regime might emerge.

We define a kinetic Ising model with $N = 1000$ binary neurons, with homogeneous all-to-all connectivity (Fig. 1A). Each neuron i is a stochastic variable

with value $s_i(t)$ at time t that can be either 1 or 0 (active or inactive). We assume that this neuron population inhibits behaviour, so that when $\sum_i s_i(t) > 0$ the larva is in the rest phase, and otherwise the larva remains active .

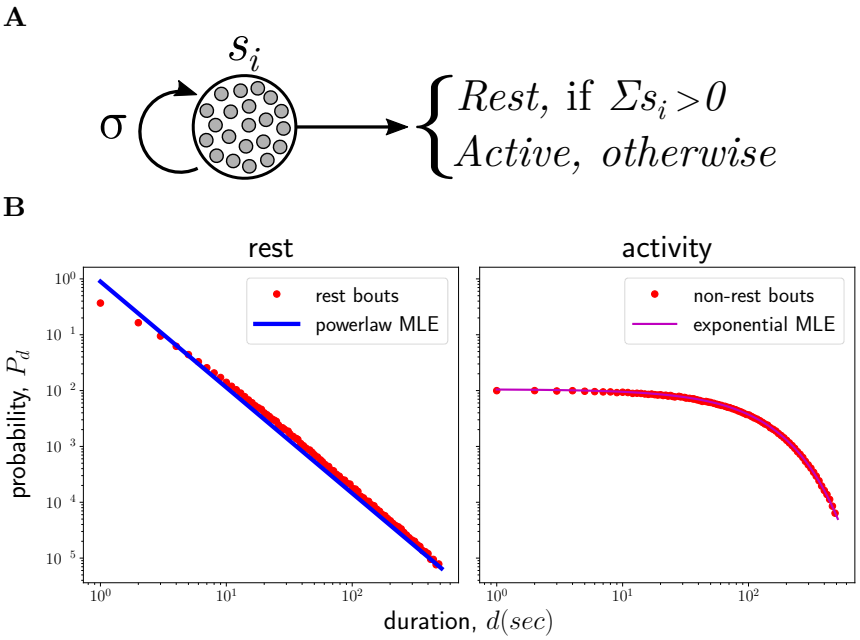

Fig. 1. Probability distribution of the duration d of rest and activity phases in a branching process model of $\sigma = 1$, simulated over 10^5 occurrences of each phase. Duration is measured as the number of updates until a phase is ended. Unit activation $s_i(t)$ propagates to neighbouring units creating self-limiting avalanches. In the rest phase, when $\sum_i s_i(t) > 0$, the system yields a power law distribution with exponent $\alpha \approx 2$. In the activity phase, when $\sum_i s_i(t) = 0$, one unit of the system is activated with probability $\mu = 0.01$, yielding an exponential distribution with coefficient $\lambda = 0.1$.

At time $t + 1$, each neuron's activation rate is proportional to the sum of activities at time t, and will be activated with a linear probability function $p_i(t+1) = \frac{\sigma}{N} \sum_j s_j(t) + \frac{\mu}{N}$. Here, σ is the propagation rate, which indicates that when a node is active at time t, it propagates its activation at time $t+1$ on average to σ other neurons. When one neuron is activated, this model behaves like a branching process [10], with σ as the branching parameter. If $\sigma < 1$, activity tends to decrease rapidly until all units are inactive while, if $\sigma > 1$, activity tends to be amplified until saturation. At the critical point, $\sigma = 1$, activity is propagated in scale-free avalanches, in which duration d of an avalanche once initiated follows a power-law distribution $P(d) \sim d^{-\alpha}$ (Fig. 1B, left), governed by a critical exponent ($\alpha = 2$ at the $N \to \infty$ limit) describing how avalanches at many different scales are generated.

When an avalanche is extinguished, the system returns to quiescence which is only broken by the initiation of a new avalanche. With a residual rate $\mu = 0.01$

the system becomes active by firing one unit and initiating a new avalanche. In this case the duration of quiescence bouts (the interval between two consecutive avalanches) follows an exponential distribution (Fig. 1B, right).

This simple conceptual model alternates autonomously between avalanches of power-law distributed durations and quiescence intervals of exponentially distributed durations. This alternation between power-law and non-power-law regimes can serve as a basic qualitative model of the transition between rest and activity bouts in the larva (cf. Discussion).

3.2 Parameterization of larval intermittent behavior

We analyzed intermittent behavior during larval crawling in a stimulus-free environment (cf. Materials and Methods for dataset description). Each individual larva was video-tracked in space (Fig. 2A). From the time series of spatial coordinates we computed the instantaneous velocity and determined a threshold value (Fig. 2B) that separates plateaus of continued activity (activity bouts) from epochs of inactivity (rest bouts, Fig. 2C-D) following the analyses suggested in [12].

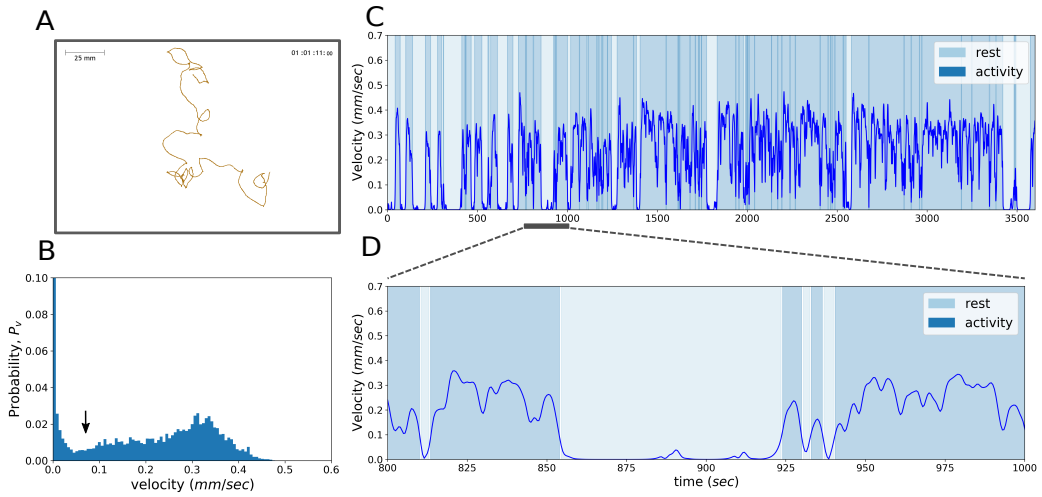

Fig. 2. Bout annotation methodology. A. Individual larva trajectory. Spatial scale and recording duration are noted. B. Velocity distribution for the single larva. The threshold obtained from the reference group, used for rest vs activity bout annotation is denoted by the arrow. C. The entire velocity time series of the larva. Rest and activity bouts are indicated by different background colors. D. Magnification of the velocity time series.

We start out with the analysis of experimental control groups that were not subjected to genetic intervention. As a first step we computed the number of occurrences of rest and activity bouts and the activity ratio, which quantifies the accumulated activity time as fraction of the total time (Tab. 1). For the

reference control group we obtain an activity ratio of 0.83 albeit with a fairly large variance across individuals.

For the duration distribution of rest bouts we find that it is best approximated by a power-law distribution in all six control groups (Tab. 2) confirming previous results on independent datasets [12,8]. The empirical duration distribution of rest-bouts across the reference control group is depicted in Fig. 3 A (red dots). Again, the power law provides the best distribution fit. The exponent α of the power law ranges from 1.514 to 1.938 with $\alpha = 1.598$ for the reference control group.

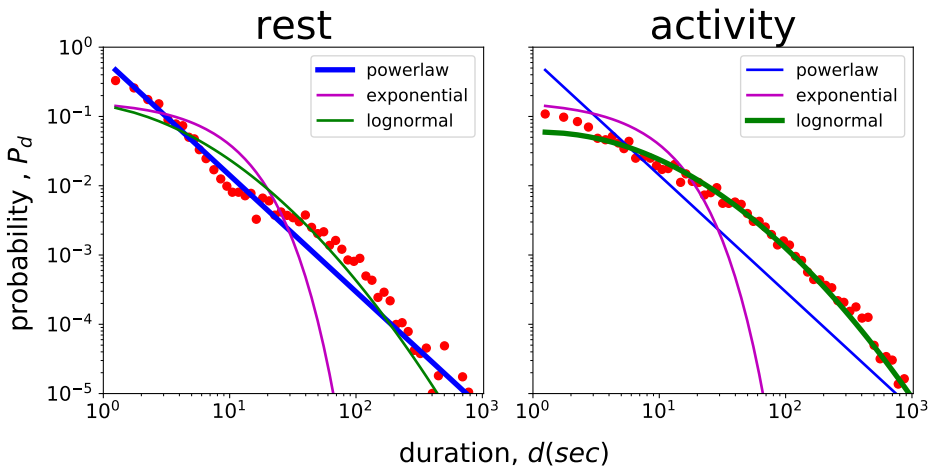

Fig. 3. Probability density of rest and activity bout durations for the reference control group. Dots stand for probability densities over logarithmic bins. Lines are the best fitting power-law, exponential and log-normal distributions. The thick line denotes the distribution having the minimum Kolmogorov-Smirnov distance D_{KS} (Tab. 2 - 3).

When analyzing the durations of activity bouts we found that these are best approximated by a log-normal distribution in all groups (Tab. 3). This result is surprising as previous work suggested the mode of an exponential distribution [12]. For the reference control group Fig. 3 B compares the empirical duration distribution of activity bouts with the fits of the three distribution functions.

3.3 Modification of rest and activity bout durations in mutant flies

Behavioral phenotypes in genetic mutants could help identify brain neuropiles in the nervous system of *Drosophila* larva that are involved in the generation of intermittent behavior, or that have an effect on its modulation. To this end we analyzed 4 experimental groups where genetic intervention was controlled by temperature either via the temperature-sensitive shibire protocol or via temperature-induced neuronal death (rpr/hid genotype).

Interestingly, genetic intervention can have a large effect on the activity ratio. When inactivating sensory neurons and to a lesser extend the mushroom

body the activity ratio is decreased (cf. BLsens $> rpr/hid$, MB247 $> shi^{ts}$ and BLsens $> shi^{ts}$ in Tab. 1). Inspection of the empirical duration distribution of rest bouts in Fig. 4 (first and third columns) shows that while the power-law fit is superior for all control groups, the log-normal fit approximates best the respective mutant distribution in 3 out of 4 cases (cf. MB247 $> shi^{ts}$, BL $> shi^{ts}$ and BLsens $> shi^{ts}$ in Tab. 2. This might hint impairment of the power-law generating processes due to neural dysfunction. In the fourth case of BLsens $> rpr/hid$ the power-law is preserved but shifted to higher values. Regarding activity, the empirical distributions indicate that overall the activity epochs are severely shortened in time for both the BLsens $> rpr/hid$ and the BLsens $> shi^{ts}$ mutants in comparison to the respective control groups (second and fourth columns).

T (˚C)	treatment			rest bouts				
		powerlaw		exponential		lognormal		
		alpha	KS D	lambda	KS D	mu	sigma	KS D
22	BLsens > rpr/hid$_{control}$	**1.938**	**0.089**	0.119	0.411	1.066	1.122	0.171
32	BLsens > rpr/hid	**1.6**	**0.094**	0.04	0.439	1.665	1.476	0.13
33	MB247/+	**1.58**	**0.06**	0.042	0.438	1.724	1.55	0.133
33	MB247 > shi^{ts}	1.41	0.167	0.029	0.246	**2.438**	**1.574**	**0.073**
22	shi^{ts} /+	**1.702**	**0.086**	0.049	0.494	1.425	1.369	0.149
33	shi^{ts} /+	**1.465**	**0.099**	0.027	0.384	2.151	1.699	0.103
22	BL/+	**1.783**	**0.046**	0.048	0.538	1.277	1.399	0.181
33	BL/+	**1.514**	**0.101**	0.039	0.406	1.944	1.623	0.127
33	BL > shi^{ts}	1.666	0.109	0.111	0.246	**1.502**	**1.189**	**0.103**
33	BLsens > shi^{ts}	1.483	0.125	0.033	0.387	**2.072**	**1.554**	**0.105**
32-33	Reference control	**1.598**	**0.061**	0.044	0.448	1.671	1.55	0.14

Table 2. Distribution parameter fits of empirical rest bout duration. The relevant parameters for the best fitting distribution are indicated in bold text.

T (˚C)	treatment			activity bouts				
		powerlaw		exponential		lognormal		
		alpha	KS D	lambda	KS D	mu	sigma	KS D
22	BLsens > rpr/hid$_{control}$	1.351	0.177	0.017	0.249	**2.847**	**1.641**	**0.041**
32	BLsens > rpr/hid	1.591	0.196	0.096	0.189	**1.691**	**1.101**	**0.062**
33	MB247/+	1.428	0.149	0.032	0.248	**2.336**	**1.524**	**0.063**
33	MB247 > shi^{ts}	1.391	0.212	0.027	0.227	**2.557**	**1.412**	**0.035**
22	shi^{ts} /+	1.371	0.175	0.025	0.212	**2.699**	**1.528**	**0.066**
33	shi^{ts} /+	1.359	0.184	0.018	0.287	**2.786**	**1.634**	**0.044**
22	BL/+	1.335	0.191	0.018	0.214	**2.988**	**1.644**	**0.079**
33	BL/+	1.327	0.215	0.015	0.224	**3.058**	**1.566**	**0.042**
33	BL > shi^{ts}	1.483	0.179	0.045	0.252	**2.07**	**1.353**	**0.063**
33	BLsens > shi^{ts}	1.682	0.209	0.147	0.201	**1.466**	**1.011**	**0.096**
32-33	Reference control	1.366	0.173	0.019	0.26	**2.73**	**1.617**	**0.046**

Table 3. Distribution parameter fits of empirical activity bout duration. The relevant parameters for the best fitting distribution are indicated in bold text.

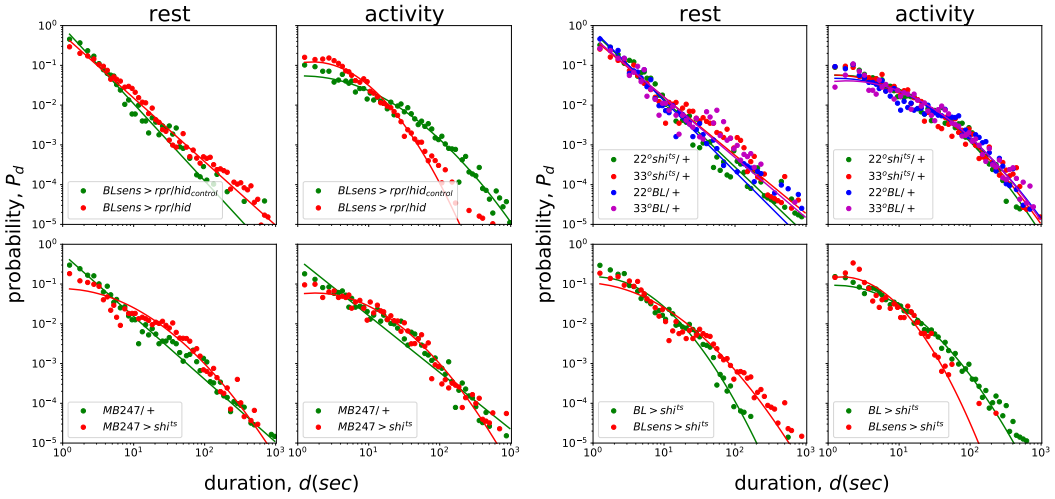

Fig. 4. Probability density of rest and activity bout durations for control and activated mutant genotypes. In the first two diagram pairs mutants are plotted against their single respective controls. In the fourth pair the rest two mutants are plotted with their 4 control groups shown in the third diagram pair. Dots stand for probability densities over logarithmic bins. Lines indicate the distribution with the lowest Kolmogorov-Smirnov distance D_{KS} among the best fitting power-law, exponential and log-normal distributions for each group (Tab. 2 - 3).

4 Discussion

As most neuroscientific research focuses either on static network connectivity or on neural activation/inhibition - behavior correlations, an integrative account of how temporal behavioral statistical patterns arise from unperturbed neural dynamics is still lacking. In this context, we hope to contribute to scientific discovery in a dual way. Firstly by extending existing mechanistic hypothesis for larva intermittent behavior and secondly by promoting the integration of intermittency in functional models of larval behavior. In what follows we elaborate on these goals and finally describe certain limitations of our study.

4.1 Self-limiting inhibitory waves might underlie intermittent crawling and its modulation

The neural mechanisms underlying intermittency in larva behavior remain partly unknown. Displacement runs are intrinsically discretized, comprised of repetitive, stereotypical peristaltic strides. These stem from segmental central pattern generator circuits (CPG) located in the ventral nerve chord, involving both excitatory and inhibitory premotor neurons and oscillating independently of sensory feedback [7]. A 'visceral pistoning' mechanism involving head and tail-segment

synchronous contraction underlies stride initiation [4]. Speed is mainly controlled via stride frequency [4].Crawling is intermittently stopped during both stimulus-free exploratory behavior and chemotaxis, giving rise to non-stereotypical stationary bouts during which reorientation might occur. During the former they are intrinsically generated without need for sensory feedback or brain input [9], while during the latter an olfactory-driven sensorimotor pathway facilitates cessation of runs when navigating down-gradient. Specifically, inhibition of a posterior-segment premotor network by a sub-esophageal zone descending neuron deterministically terminates runs allowing easier reorientation [11].

It is reasonable to assume that this intermittent crawling inhibition is underlying both free exploration and chemotaxis, potentially in the form of transient inhibitory bursts. A neural network controlling the CPG through generation of self-limiting inhibitory waves is well suited for such a role. In the simplest case, during stimulus-free exploration, the durations of the generated inhibitory waves should follow a power-law distribution, behaviorally observed as rest bouts. In contrast, non-power-law distributed quiescent periods of the network would disinhibit locomotion allowing the CPG to generate repetitive peristaltic strides resulting in behaviorally observed runs.

The model we presented (cf. 3.1) alternates autonomously between avalanches of power-law distributed durations and quiescence intervals of exponentially distributed durations without need for external input. Therefore it can serve as a theoretical basis for the development of both generative models that reproduce the intermittent behavior of individual larvae and of the above mechanistic hypothesis for the initiation and cessation of peristaltic locomotion in the larva through disinhibition and inhibition of the crawling CPG respectively. To uncover the underlying neural mechanism and confirm/reject our hypothesis, inhibitory input to the crawling CPG should be sought and correlated to behaviorally observed stride and stride-free bouts during stimulus-free exploration.

Intermittent behavior in the *Drosophila* larva is subject to two modes of modulation, neither of which affects the distribution of the activity bouts. Firstly, high ambient temperature and daylight raise the activity ratio over long timescales by raising the number of activity bouts [12]. This is achieved by lowering the probability of the extremely long rest bouts, without affecting the power-law exponent of the distribution, which coincides with fewer sleep events (> 5 minutes) observed during the day. This modulation is long-lasting and could result from a different constant tonic activation of the system. Secondly, dopamine neuron activation raises the activity ratio acutely by modulation of the power-law exponent upwards [12] skewing locomotion towards the brownian limit. This modulation could be transient in the context of salient phasic stimulation by the environment. As mentioned above, during chemotaxis larvae perform more and sharpest reorientations, terminating runs when navigating down-gradient. A hypothesis integrating both experimental findings could be then that this behavior stems from transient olfactory-driven dopaminergically-modulated inhibition of the crawling CPG. Our conceptual model can be extended to address the above claims by adding tonic and/or phasic input.

4.2 Intermittency can extend functional models of larva locomotion

Traditional random walk models fail to capture the temporal dynamics of animal exploration [5]. Even when time is taken into account in terms of movement speed, reorientations are assumed to occur acutely. Integrating intermittency can address this limitation allowing for more accurate functional models of autonomous behaving agents. Such virtual agents can then be used in simulations of behavioral experiments promoting neuroscientifically informed hypothesis that advance over current knowledge and generate predictions that can stimulate further empirical work [1].

It is widely assumed that *Drosophila* larva exploration can be descibed as a random walk of discrete non-overlapping runs and reorientations/head-casts [9] or alternatively that it is generated by the concurrent combined activity of a crawler and a turner module generating repetitive oscillatory forward peristaltic strides and lateral bending motions respectively and possibly involving energy transfer between the two mechanical modes [3,13,6]. Both models can easily be upgraded by adding crawling intermittency which might or might not be independent of the lateral bending mechanism. In the discrete-mode case, intermittency can simply control the duration and transitions between runs and head-casts or introduce a third mode of immobile pauses resulting in a temporally unfolding random walk. In the overlapping-mode case the two modules are complemented by a controlling intermittency module forming an interacting triplet. Depending on the crawler-turner interaction and the effect of intermittency on the turner module, multiple locomotory patterns emerge including straight runs, curved runs, stationary head-casts and immobile pauses. This simple extension would allow temporal fitting of generative models to experimental observations in addition to the primarily pursued spatial-trajectory fitting, facilitating the use of calibrated virtual larvae in simulations of behavioral experiments.

4.3 Limitations

A limitation of our study is that due to the single-spinepoint tracking, it is impossible to determine whether micro-movements happen during the designated inactivity periods, an issue also unclear in [12]. It follows that in both the analysed dataset and in [12], immobile pauses, feeding motions and stationary head casts are indistinguishable. Therefore, what we define as rest bouts should be considered as periods lacking at least peristaltic strides, not any locomotory activity whatsoever. Our relatively low velocity threshold $V_\theta = 0.085mm/sec$ though allows stricter detection of rest bouts as it is evident from the higher activity ratio (higher than 0.7 in most control groups in comparison to lower than 0.25 in [12]). To tackle this, trackings of higher spatial resolution with more spinepoints tracked per larva are needed, despite the computational challenge of the essentially long recording duration.

Also, our results show that an exponential distribution of activity bouts [12] might not always be the case, as we detected lognormal long-tails in all cases. Still, the exponential-power-law duality in our model illustrates switching

between independent and coupled modes of neural activity. Substituting the exponential regime by other long-tailed distribution such as log-normal might require assuming more complex interactions between the switching regimes and will be pursued in the future so that generative models of the data can be fit.

References

1. Barandiaran, X.E., Chemero, A.: Animats in the Modeling Ecosystem. Adapt. Behav. **17**(4), 287–292 (2009). https://doi.org/10.1177/1059712309340847
2. Bartumeus, F.: Behavioral intermittence , Levy patterns , and randomness in animal movement. Oikos **118:**, 488–494 (2009). https://doi.org/10.1111/j.1600-0706.2008.17313.x
3. Davies, A., Louis, M., Webb, B.: A Model of Drosophila Larva Chemotaxis. PLoS Comput. Biol. **11**(11), 1–24 (2015). https://doi.org/10.1371/journal.pcbi.1004606
4. Heckscher, E.S., Lockery, S.R., Doe, C.Q.: Characterization of Drosophila Larval Crawling at the Level of Organism, Segment, and Somatic Body Wall Musculature. J. Neurosci. **32**(36), 12460–12471 (2012). https://doi.org/10.1523/jneurosci.0222-12.2012
5. James, A., Plank, M.J., Edwards, A.M.: Assessing Lévy walks as models of animal foraging. J. R. Soc. Interface **8**(62), 1233–1247 (2011). https://doi.org/10.1098/rsif.2011.0200
6. Loveless, J., Lagogiannis, K., Webb, B.: Modelling the neuromechanics of exploration and taxis in larval Drosophila. PLoS Comput. Biol. pp. 1–33 (2019). https://doi.org/10.5281/zenodo.1432637
7. Mantziaris, C., Bockemühl, T., Büschges, A.: Central Pattern Generating Networks in Insect Locomotion. Dev. Neurobiol. **00**(October 2019), 1–15 (2020). https://doi.org/10.1002/dneu.22738
8. Reynolds, A.M., Jones, H.B.C., Hill, J.K., Pearson, A.J., Wilson, K., Wolf, S., Lim, K.S., Reynolds, D.R., Chapman, J.W., Reynolds, A.M.: Evidence for a pervasive 'idling-mode' activity template in flying and pedestrian insects. R. Soc. Open Sci. **2** (2015). https://doi.org/10.1098/rsos.150085
9. Sims, D.W., Humphries, N.E., Hu, N., Medan, V., Berni, J.: Optimal searching behaviour generated intrinsically by the central pattern generator for locomotion. Elife **8**, 1–31 (2019). https://doi.org/10.7554/eLife.50316
10. Slade, G.: Probabilistic models of critical phenomena. The Princeton companion to mathematics pp. 343–346 (2008)
11. Tastekin, I., Khandelwal, A., Tadres, D., Fessner, N.D., Truman, J.W., Zlatic, M., Cardona, A., Louis, M.: Sensorimotor pathway controlling stopping behavior during chemotaxis in the Drosophila melanogaster larva. Elife **7**, 1–38 (2018). https://doi.org/10.7554/elife.38740
12. Ueno, T., Masuda, N., Kume, S., Kume, K.: Dopamine Modulates the Rest Period Length without Perturbation of Its Power Law Distribution in Drosophila melanogaster. PLoS One **7**(2) (2012). https://doi.org/10.1371/journal.pone.0032007
13. Wystrach, A., Lagogiannis, K., Webb, B.: Continuous lateral oscillations as a core mechanism for taxis in Drosophila larvae. Elife **5** (2016). https://doi.org/10.7554/elife.15504

4 A behavioral architecture for realistic simulations of *Drosophila* larva locomotion and foraging

Authors: Panagiotis Sakagiannis, Anna-Maria Jürgensen, Martin Paul Nawrot
Journal: eLife
Date: April 2025
DOI: 10.7554/eLife.104262.1

Author Contributions

Contributions to the publication were as follows: Conceptualization was carried out by **P.S.** and M.P.N.; data curation, investigation, software development, methodology, visualization, and original draft writing were performed by **P.S.**; formal analysis and validation were conducted by **P.S.** and A.-M.J.; resources, funding acquisition, and supervision were provided by M.P.N.; and all authors (**P.S.**, A.-M.J., M.P.N.) contributed to review and editing.

Neuroscience

A behavioral architecture for realistic simulations of *Drosophila* larva locomotion and foraging

Panagiotis Sakagiannis, Anna-Maria Jürgensen, Martin Paul Nawrot

Computational Systems Neuroscience, University of Cologne, Cologne, Germany

eLife Assessment

This **useful** study presents a hierarchical computational model that integrates locomotion, navigation, and learning in Drosophila larvae. The evidence supporting the model is **solid**, as it qualitatively replicates empirical behavioral data, but the experimental data is **incomplete**. While some simplifications in neuromechanical representation and sensory-motor integration are limiting factors, the study could be of use to researchers interested in computational modeling of biological movement and adaptive behavior.

https://doi.org/10.7554/eLife.104262.1.sa3

Abstract

The *Drosophila* larva is extensively used as model organism in neuroethological studies where precise behavioral tracking enables the statistical analysis of individual and population-level behavioral metrics that can inform mathematical models of larval behavior. Here, we propose a hierarchical model architecture comprising three layers to facilitate modular model construction, closed-loop simulations, and direct comparisons between empirical and simulated data. At the basic layer, the autonomous locomotory model is capable of performing exploration. Based on novel kinematic analyses our model features intermittent forward crawling that is phasically coupled to lateral bending. At the second layer, navigation is achieved via active sensing in a simulated environment and top-down modulation of locomotion. At the top layer, behavioral adaptation entails associative learning. We evaluate virtual larval behavior across agent-based simulations of autonomous free exploration, chemotaxis, and odor preference testing. Our behavioral architecture is ideally suited for the modular combination of neuromechanical, neural or mere statistical model components, facilitating their evaluation, comparison, extension and integration into multifunctional control architectures.

Introduction

Drosophila larvae express a fairly tractable behavioral repertoire that is consistent across the 4-5 days of the larval life stage (Almeida-Carvalho et al., 2017). Most of the larval lifetime is dedicated to foraging for suitable nutrients while avoiding threats via stereotyped evasive behaviors. Behavioral control is achieved via neural circuits conserved throughout development (Gerhard et al., 2017), that span the entire tripartite CNS consisting of the brain, the subesophageal zone (SEZ) and the ventral nerve cord (VNC), making the larva a formidable system for studying control, execution and adaptation of behavior (Zarin et al., 2019; Jovanic, 2020; Eschbach and Zlatic, 2020; Imambocus et al., 2022). After reaching critical mass for pupation, homeostatic signals switch behavior towards food aversion, hypermobility and collaborative burrowing (Wu et al., 2003), terminating the feeding state and leading to pupation and metamorphosis.

High resolution methods for behavioral tracking (Ohyama et al., 2013; Risse et al., 2017; Schumann and Triphan, 2020; Tadres and Louis, 2020; Thane et al., 2023; Laurent et al., 2024), now routinely used in neuroethological experiments, have revealed detailed insight in the organisation of larval foraging behavior. In the absence of food resources, larvae engage in free exploration to locate food patches (Sims et al., 2019) with a characteristic alternation of locomotory activity and brief pauses (Sakagiannis et al., 2020), a property also reported for adult fly behavior (Ueno et al., 2012; Reynolds et al., 2015). Food consumption involves repetitive feeding motion and digging into the substrate (Kim et al., 2017). Statistical regularities that govern foraging behavior have been un-veiled by analysis at the microscale of body kinematics and at the macroscale of larva trajectories (Denisov et al., 2013; Risse et al., 2017; Karagyozov et al., 2018). Locomotion combines the basic sensorimotor primitives of crawling and turning, and has been in the main focus of recent studies (Heckscher et al., 2012; Wystrach et al., 2016; Sims et al., 2019), whereas studies of feeding behavior remain scarce (Ruiz-Dubreuil et al., 1996). Both, crawling and feeding consist of recurring sensorimotor cycles controlled by central pattern generating circuits (Heckscher et al., 2012; Mantziaris et al., 2020; Miroschnikow et al., 2018a).

Salient olfactory cues can trigger appetitive or aversive chemotaxis, during which larvae navigate up or down a chemical gradient (Gomez-Marin et al., 2011; Slater et al., 2015; Schleyer et al., 2015a; Klein et al., 2017). During appetitive chemotaxis, the detection of minor concentration changes during lateral bending causes a directional bias in the turning movement, establishing a mechanism of active sensing (Gomez-Marin and Louis, 2012; Wystrach et al., 2016; Thane et al., 2019). Encounters with novel odorants in the presence of a food reward or a negative reinforcement such as the bitter taste substance quinine can induce olfactory learning, enabling short- and long-term behavioral adaptations (Schleyer et al., 2011; Schleyer et al., 2015b; Gerber and Stocker, 2007; Diegelmann et al., 2013; Widmann et al., 2018; Weiglein et al., 2019; Jürgensen et al., 2024). For quantification of choice behavior, widely-used standard group assays for behavioral preference testing have been established (Gerber and Stocker, 2007; Gerber et al., 2014; Schleyer et al., 2015a; Schleyer et al., 2015b; Schleyer et al., 2011).

The routine access to detailed behavioral data, and the broad approaches to neuroethological experiments makes the *Drosophila* larva an ideal model system for computational modeling studies on behavioral control principles. Existing generative models typically aim at specific aspects of larval behavior and can vary widely in model type and abstraction level, ranging from basic neuromechanics to the abstract mathematical description. We therefore propose the concept of the behavioral architecture, a three-layered hierarchical and modular framework, nested from behavioral primitives to high-level function, to integrate models of diverse types and different levels of abstraction (**Figure 1*A***). At the basic layer we establish a refined coupled-oscillator

model for larval locomotion capable of autonomous exploratory behavior. To this end we provide a detailed analysis of larval tracking data from experiments of free exploration. Introducing sensory modulation at the second layer allows us to reproduce several experimental findings of appetitive chemotaxis. The combination with a biologically realistic spiking neural network model for associative learning at the adaptive layer reproduces empirical population-level learning scores across a multiple trial training protocol. The architecture allows for agent-based, potentially closed-loop, simulations of virtual larvae in virtual environments and generates simulated data at the level of the experimentally observed behavior facilitating model calibration and evaluation.

Model

Behavioral architecture

Larval behavior is hierarchically structured, i.e. sensorimotor primitives such as crawling, bending and feeding can be integrated into more complex behaviors such as exploration, taxis or exploitation. It has been proposed that the hierarchy of animal behavior is reflected by the underlying neuroanatomy as a hierarchy of nested sensorimotor loops (Prescott et al., 1999; Wilson and Prescott, 2022; Prescott and Wilson, 2023). A functional modeling paradigm that exploits this idea regards the neural system as a layered control architecture (or subsumption architecture) (Brooks, 1986) where low-level stereotyped reflexive and repetitive behaviors are autonomously generated by localized peripheral motor circuitries, while multisynaptic loops involving more centralized neural circuits act as descending modulators on the local circuits in order to coordinate global and complex behavior. The central idea is that energy-efficient decentralized neural control is the rule, while higher centers are recruited only when more extensive integration is needed e.g. in order to react suitably to unexpected sensory stimulation. Furthermore, there are only limited degrees of freedom by which higher layers can influence local sensorimotor loops via descending pathways e.g. by starting or stopping, or by accelerating, decelerating or inverting their autonomous function (Sen et al., 2017; Feng et al., 2020; Bidaye et al., 2020). The complexity of inter-layer top-down modulation is thus predicted to be considerably lower than the complexity of signal integration within layers. Layered control architectures have been used mainly in behavior-based robotics (Bicho, 1999), allowing sequential calibration and modular integration of partial neuroscientific models under a common framework (Brooks, 1986; Prescott et al., 1999).

We here propose a three-layered behavioral architecture for *Drosophila* larva foraging as illustrated in **Figure 1*A***. The bottom layer consists of three basic behaviors: crawling, turning and feeding. Each is realized by a low-level local sensorimotor loop between motor effectors and sensory feedback. For exploration this involves proprioception and mechanoreception, for feeding it additionally involves gustatory input. Integration of these basic behaviors within the layer gives rise to composite behaviors. For example while exploration in stimulus-free conditions only entails crawling and turning, differences in the temporal microstructure between these two yield a spectrum from localized search to remote dispersal. The peripheral ventral nerve cord (VNC) circuits mediating exploration have been shown to be capable of autonomous, decentralised function without the need for any top-down regulation (Sims et al., 2019).

The intermediate layer introduces salient sensory stimulation of different modalities and therefore allows for reactive behavior in the face of presented risks and opportunities. Following the subsumption architecture paradigm, we only assume a single modulatory input from the intermediate to the bottom layer. This assumption is in line with the idea of pre-motor integration of signals from different sensory modalities into a final integrated descending modulatory pathway, which eventually affects behavior (Wystrach et al., 2016; Eschbach and Zlatic, 2020). Modulation of exploratory behavior in the presence of salient stimulation enables coherent navigation along sensory gradients. In the present study we consider odor-driven chemotaxis as a

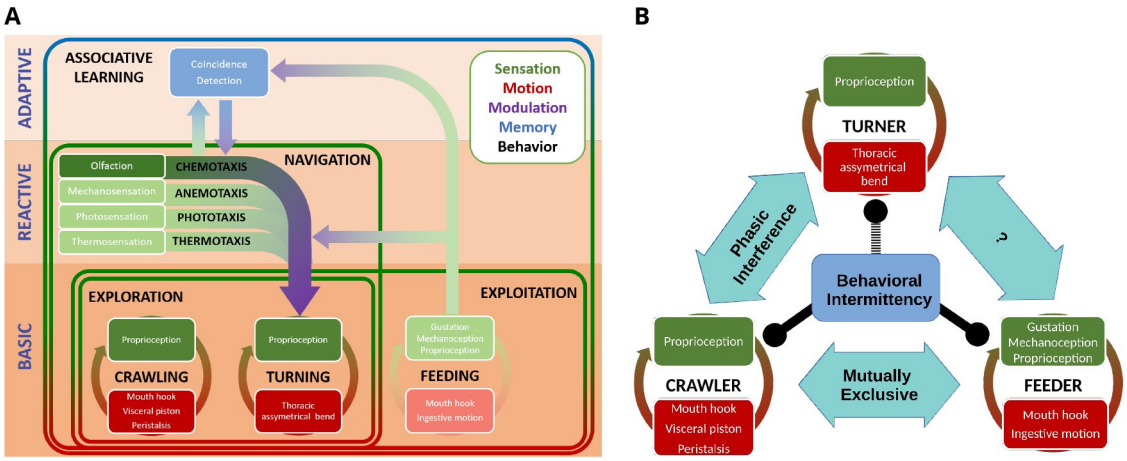

Figure 1.

Behavioral architecture for larva foraging.

A: In the trilayer architecture the bottom layer consists of three basic sensorimotor effectors that constitute the locomotory model. The intermediate layer features innate reactive behavior in response to sensory input that reflects changes in the environment. The top layer allows for behavioral adaptation through experience. Framed areas denote more complex behaviors that require subsumption of subordinate behaviors. **B:** Suggested implementation of basic behavioral modules at the bottom layer. Initiation or cessation of a behavior is controlled by the intermittency module. The turner and crawler module are phasically coupled during forward locomotion, while crawling and feeding are implemented as mutually exclusive sensorimotor primitives.

process of active sensing in which the larva constantly samples the odor concentration during lateral bending motions, enabling odorscape navigation and odor source discovery. In an alternative application, we recently simulated larval behavior during active sensing in thermotaxis experiments across different *Drosophila* species (Kafle et al., 2024).

Finally, at the top layer, associative learning biases the evaluation of recurring sensory stimuli and can modulate the innate valence of odors to allow for experience-dependent adaptation of navigation. Again, the modulatory top-down signal to the underlying layers is limited to the modulation of approach versus avoidance behavior, in line with the research on the mushroom body output signal in *Drosophila* (Gerber and Stocker, 2007; Gerber et al., 2014; Schleyer et al., 2015a; Schleyer et al., 2015b; Schleyer et al., 2011; Owald and Waddell, 2015).

The proposed behavioral architecture is naive to the level of model abstraction of the underlying neural mechanisms. Its role is to create a modular framework where diverse and potentially competing models can be positioned, integrated and tested in behavioral simulations. It is therefore justified to use simple generative models to bridge the gaps that have not yet been studied in detail until they can be substituted by more elaborate future versions. To ensure flexibility and compatibility, every module has a defined basic input and output regardless of the internal structure, complexity, level of abstraction and spatiotemporal scale of its possible instantiations. For the purpose of modeling under the behavioral architecture framework we provide the *larvaworld* python package available at *https://pypi.org/project/larvaworld/*.

Locomotory model

We propose a model for locomotion in the two-dimensional plane that defines the momentary body state by the instantaneous forward velocity v and angular velocity ω as generated through crawling and bending, respectively. Extending on previous models (see Discussion) we incorporate two new features based on our kinematic analysis of larval locomotion: (i) the peristaltic cycle phase dependent attenuation of angular motion, and (ii) the intermittent crawling as transitions between runs and pauses. We briefly describe the modular components of our model. A detailed mathematical description is provided in the Materials and Methods section. The complete pipeline for model calibration is described in *Appendix 2*.

We choose to simplify the larva to a two-segment body (**Figure 2*F***). This abstraction allows to describe the body state at any point in time by only three variables: (i) position of a selected spine midpoint (**Appendix 1-Figure 1**), (ii) absolute orientation of the front segment θ, and (iii) bending angle θ_b between the front and rear segments. This approach is in line with the common practice of quantifying larva bending via a single angle between one anterior and one posterior vector (Gomez-Marin et al., 2011; Lahiri et al., 2011; Paisios et al., 2017).

Locomotion is achieved via forward crawling and lateral body bending. Crawling moves the midpoint along the orientation vector. Bending reorients the front segment by rotation around the midpoint. Forward displacement restores θ_b back to zero by gradually aligning the rear segment to the orientation axis. A demonstration of the bisegmental model simplification is shown in Video 4.

Based on our empirical results (**Figure 2** and **Figure 3**), we propose a coupled oscillator model of locomotion that generates both, weathervaning and headcasts. To this end we implement an oscillatory process (crawler) that generates the forward-velocity v during subsequent stride cycles (**Figure 1*B***) where a tonic input I_C modulates crawling frequency f_C. A second oscillator (turner) generates alternating left-right bending as previously proposed by (Wystrach et al., 2016) with oscillatory amplitude A_T. To incorporate the empirically observed crawl-bend interference, we couple the two oscillators by imposing a crawler-phase dependent suppression on A_T where the phase-dependent modulation is modelled as Gaussian function and fitted to the empirical data (red profile in **Figure 3*C***).

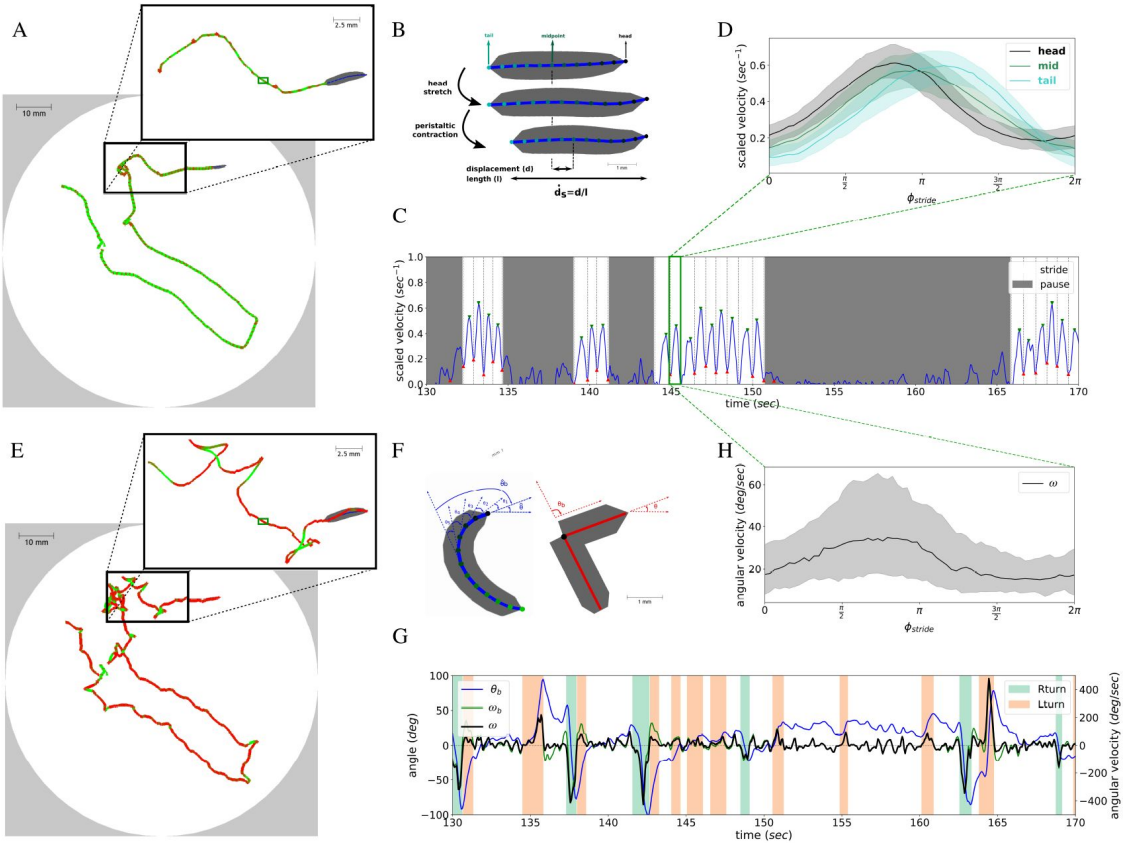

Figure 2.

Kinematic analysis of a single Drosophila larva in locomotion.

A: Individual larva trajectory tracking a posterior point along the midline of the animal . Trajectory color denotes the forward velocity *v* from 0 (red) to maximum (green). Inset focuses on a shorter track epoch analyzed in C and G. The full-length trajectory and the epoch in the inset are shown in Figure 2—video 1 and Figure 2—video 2 respectively. Dark green rectangle denotes the single stride described in B. **B:** Sketch of the single crawling stride indicated in A. The larva first stretches its head forward, anchors it to the substrate and then drags its body forward via peristaltic contraction. Scaled stride displacement $\hat{d_S}$ is defined as the resulting displacement *d* divided by the body-length *l*. **C:** Scaled forward velocity *v* during the 40 s track epoch selected in A (inset). Green and red markers denote the local maxima and minima used for stride annotation. Individual strides are tiled by vertical dashed lines. Successive strides constitute uninterrupted stridechains (white). Epochs that do not show any strides are annotated as crawl-pauses (gray). **D:** Scaled forward velocity *v* of head, midpoint and tail as a function of the stride cycle phase Φ. All detected strides have been interpolated to a stride oscillation cycle of period 2*π*. Solid lines denote the median, shaded areas the lower and upper quartiles across strides. **E:** Same trajectory as in (A) now tracking the head segment. Color denotes the absolute orientation angular velocity *ω* from 0 (red) to maximum (green). The full-length trajectory and the epoch in the inset are shown in Figure 2—video 1 and Figure 2—video 2 respectively. **F:** Definition of bending angle θ_b and orientation angle θ for the original 12-segment (blue) and the simplified 2-segment (red) larvae. **G:** Three angular parameters during the same track epoch shown in (C). Bending angle θ_b, bend and orientation angular velocities ω_b, ω are shown. Background shadings denote left and right turning bouts. For illustration purposes only turns resulting in a change of orientation angle $\Delta\theta > 20_\circ$ are shown. **H:** Absolute orientation angular velocity *ω* during the stride cycle, as shown for *v* in (D).

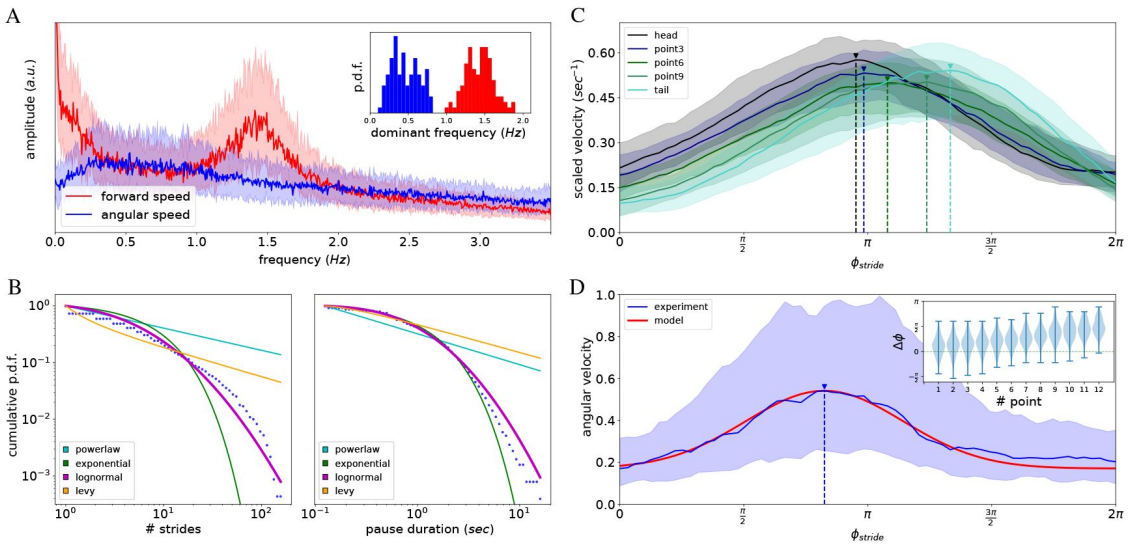

Figure 3.

Population-level analysis.

A: Fourier analysis of the forward *v* (red) and angular *ω* (blue) velocity across 100 larvae. Inset shows the respective dominant frequencies within suitable ranges $1 \leqslant f_C \leqslant 2.5$ and $0.1 \leqslant f_T \leqslant 0.8$ for *v* and *ω*, respectively. Crawling exhibits a dominant frequency of around 1.4 while lateral bending has a slower more variable rhythm of around 0.4. **B:** Epoch-duration distribution. Dots describe the cumulative probability density over logarithmic bins for the length of stridechains and the duration of crawl-pauses pooled across the larva population. Lines indicate the distribution with the lowest Kolmogorov-Smirnov distance among the best-fitting power-law, exponential, log-normal and Levy distributions. Stridechain length and pause duration are best approximated by log-normal distributions. **C-D**: Crawl-bend interference. The stride cycle kinematics are depicted for a single individual. All detected strides have been interpolated into a 64-bin oscillation cycle of period 2π. **C:** Forward velocity of 5 points along the larva midline. Velocity is scaled to the larva body-length. **D:** Absolute angular velocity *ω* (blue) normalized by the average value $\bar{\omega}_P$ computed during the pause epochs. Fitted Gaussian function (red) describes well the phase-dependent attenuation imposed on *ω* and is used for the implementation of the coupled-oscillator locomotory model. Solid lines denote the median, shaded areas the lower and upper quartiles. Vertical dashed lines denote the cycle phase where the respective velocity reaches its maximum value. Inset : Phase offset $\Delta\Phi$ between the peak phase of each midline point's forward velocity $\phi_C^{v_i}$ and the peak phase of angular velocity ϕ_C^{ω} across a dataset of 100 tracked larvae. Notably, *ω* reaches its maximum just before the head forward velocity reaches its maximum.

To accommodate the intermittent behavior with alternating crawling runs and crawling pauses, we chose to implement a stochastic model approach where the number of strides per stridechain and the pause durations each follows a log-normal distribution fitted to the empirical data (**Figure 3*B***). For a demonstration of the locomotory model in different configurations see Video 1.

Under the constraints of the subsumption architecture paradigm the crawler and turner can be influenced by higher-order circuits in a limited number of ways. In the case of the crawler, frequency modulation and the initiation or cessation of the oscillation-cycle are the two available top-down modulations. Likewise for the turner, top-down modulation of the tonic input affects both, the frequency and the amplitude of oscillation.

Results

Kinematic analysis of larva locomotion

We start out with the kinematic analysis of experimental larva trajectories and body postures in order to infer and parameterize several aspects of larva locomotion that will inform our modeling approach. Using diverse metrics that capture spatial and temporal dynamics we specifically assess the oscillation of forward velocity during consecutive peristaltic strides and its influence on lateral bending, the less pronounced oscillation of angular velocity, the intermittent nature of crawling, and the inter-individual variability of a number of locomotion-related parameters across different larvae.

We analyzed the trajectory of each single larva by tracking the forward velocity v of a posterior midline point (**Figure 2*A***) and the head orientation angular velocity ω (**Figure 2*E***). The former illustrates that forward locomotion indeed consists of consecutive steps (strides) that are characterized by the alternation of increased (green) and decreased (red) values of v. The latter confirms that lateral bending occurs both, during crawl-runs (weathervaning) and during stationary crawl-pauses (headcasts). The individual trajectory depicted in **Figure 2** can be seen in full length in Figure 2—video 1, while the short epoch depicted in the inset is shown in Figure 2—video 2.

To characterize both types of oscillations we first perform a Fourier analysis of v and ω on each of 100 larva tracks (**Figure 3*A***). Across larvae we observe a robust peristaltic rhythm with a crawling frequency f_C normally distributed around 1.4 Hz in line with earlier results (Heckscher et al., 2012; Mantziaris et al., 2020; Thane et al., 2023). We also confirm previous reports that lateral bending manifests a slower, more variable rhythm f_T around 0.4 Hz (Wystrach et al., 2016). Then we detect all strides performed by an individual animal and verify their stereotypical structure in terms of stride duration, resulting body displacement, and the phase-dependence of v (**Figure 2*D***).

In order to quantify the phase-dependence of angular motion we interpolate each detected stride into a 64-bin oscillation cycle of period 2π. The progression of forward velocity from head to tail during the peristaltic cycle is illustrated in **Figure 3*C*** (top). To analyze the angular metrics during the peristaltic cycle we consider the absolute angular change ignoring the left or right direction of a turn. This analysis reveals a robust phase dependence of the angular velocity ω and acceleration $\dot{\omega}$, and of the bend angle θ_b (**Appendix 2-Figure 1*A***). Specifically, the angular velocity ω exhibits a smooth unimodal dependency on the stride period peaking during the early stride phase and just before the head forward velocity reaches its maximum (**Figure 3*C***). Nevertheless, ω during crawling is lower than during pauses as shown when we normalize it by dividing with the average $\bar{\omega}_p$ computed during the pause epochs (**Figure 3*C***, bottom). This

implies a phasic attenuating interference of the peristaltic cycle on the lateral bending mechanism (**Figure 1*B***). A plausible mechanistic explanation featuring bodily interference of crawling and bending is suggested in the Discussion.

Larvae intermittently pause crawling before re-assuming it. This results in sequences of concatenated strides (runs or stridechains) that alternate with brief crawl-pauses (**Figure 2*C***). We analyzed the stridechain length with respect to the number of strides and the pause duration in the experimental dataset. To this end we pooled both measures across 100 larvae to obtain the empirical distributions. Testing power-law, exponential, log-normal and Levy distributions revealed the highest quality of fit for the log-normal distribution for both parameters (**Figure 3*B***).

Simulation of behavioral experiments

In this section we simulate increasingly complex behavioral experiments. Starting from stimulus-free exploration we advance to chemotactic navigation and finally to adaptive odor preference experiments. We simulate individual larvae using the model as calibrated in **Appendix 2** and all model parameters are shown in **Appendix 2-Figure 1*A***. Larval populations are constructed by pooling individual virtual larvae that behave independently as they move through the spatial arena and odorscape (Niewalda et al., 2014).

Exploration

In the first virtual experiment, we consider free exploration (**Figure 4**) and evaluate our simulation results against an empirical data set (see Materials and Methods). To this end we simulate a population of 200 virtual larvae during three minutes while exploring a stimulus-free environment on a non-nutritious substrate in a Petri-dish, mimicking the experimental lab conditions (**Figure 4**). Due to the absence of sensory input and food, the proposed intermittent coupled-oscillator locomotory model at the basic layer is sufficient to autonomously generate exploratory movement.

Larvae are initially placed at the center of the dish, each with a random body orientation. Over the course of the experiment, both the simulated and the real larval population disperses in space (**Figure 4*A***, Figure 4—video 1).

The quantitative comparison shows a good agreement between simulated and empirical data with respect to the radial dispersal from the initial position (**Figure 4*B***). Importantly, the variability across individual animals is well captured by the variability between individual model instances as shown for the number of strides, the crawling time, and the total distance travelled by each larva (**Figure 4*C***).

Chemotaxis

Chemotaxis describes the process of exploiting an odor gradient in space to locate an attractive or avoid a repelling odor source. An olfactory sensor (olfactor) placed at the front end of the virtual body enables active sensing during body bending and allows detection of concentration changes that modulate turning behavior accordingly. Sensory-driven behavior is enabled via modulation of the frequency and amplitude of the lateral oscillator (turner) by sensory feedback, while the peristaltic oscillator is not affected, as proposed in (Wystrach et al., 2016).

To assess the chemotactic efficiency of our coupled-oscillator model we reconstruct the arena and odor landscape (odorscape) of two behavioral experiments described in (Gomez-Marin et al., 2011). In the first, larvae are placed on the left side of the arena facing to the right. An appetitive odor source is placed on the right side. The virtual larvae navigate up the odor gradient approaching the source (**Figure 5*A,C,E***), reproducing the experimental observation in Fig.1 C in

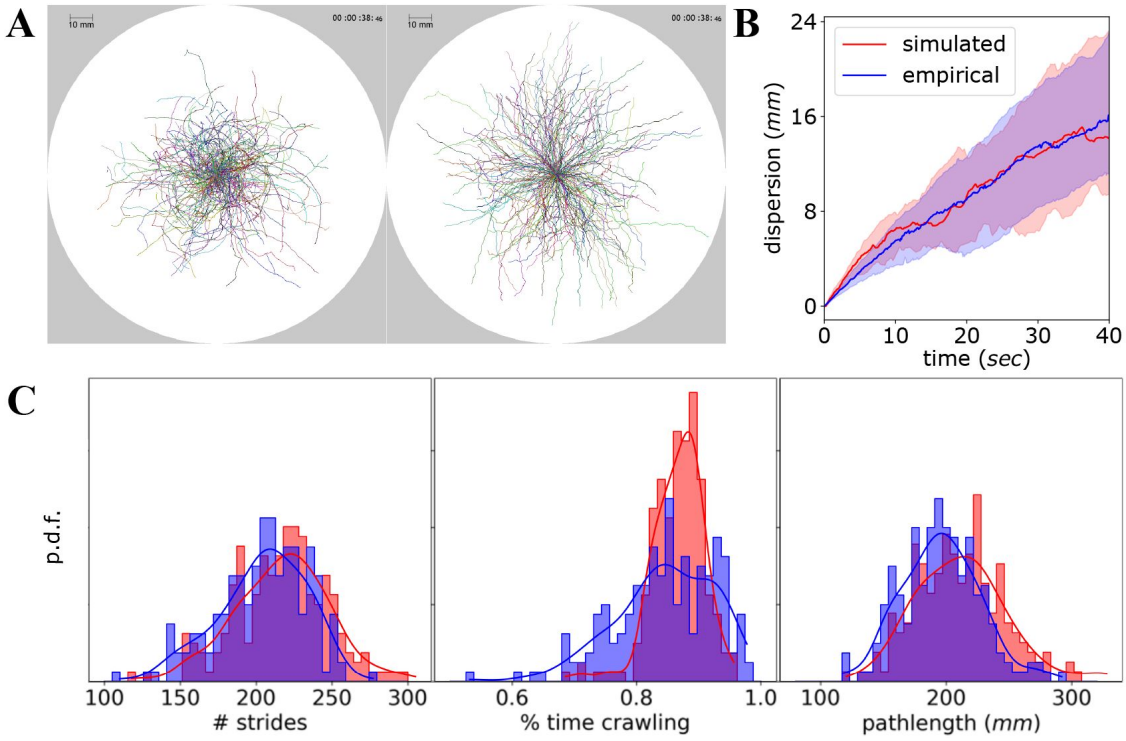

Figure 4.

Free exploration in simulation and experiment.

A: Dispersal of 200 larvae in experiment (left) and simulation (right) during 40 seconds. Individual tracks have been transposed to originate from the center of the arena. The entire temporal course shown in Figure 4—video 1. **B:** Dispersal from origin. Line indicates the group median while shaded area denotes first and third quartiles. **C:** Histograms for total number of strides, time ratio allocated to crawling and pathlength. (arena dimensions = 500×500mm, N = 200 larvae, experiment duration = 3 minutes, simulation timestep =1/16).

(Gomez-Marin et al., 2011). In the second, both the odor source and the virtual larvae are placed at the center of the arena. The larvae perform localized exploration, generating trajectories across and around the odor source. (**Figure 5*B,D,F***), again replicating the observation in **Fig.1 D** in (Gomez-Marin et al., 2011). Two sample simulations can be seen in Figure 5—video 1.

Odor preference test

We simulate the odor preference paradigm as described in the Maggot Learning Manual (Michels et al., 2017). Larvae are placed at the center of a dish containing two odor sources in opposite sides and left to freely explore. The odor concentrations are Gaussian-shaped and overlapping, resulting in an odorscape of appetitive and/or aversive opposing gradients. After 3 minutes the final situation is evaluated. The established population-level metric used is the olfactory preference index (PI), computed for the left odor as $PI_l = \frac{N_l - N_r}{N}$ where N_l and N_r is the number of larvae on the left and right side of the dish while N is the total number of larvae.

The extend of olfactory modulation on the turning behavior is determined by the odor-specific gain G. As this is measured in arbitrary units, we first need to define a realistic value range that correlates with the behaviorally measured PI. We perform a parameter-space search independently varying the gain for left and right odors and measuring the resulting PI in simulations of 30 larvae. The results for a total of 25^2 gain combinations within a suitable range of $G \in [-100, 100]$ are illustrated in **Figure 6*A***. Simulation examples for one appetitive and one aversive odor are shown in Figure 6—video 1.

In order to simulate larval group behavior in response to an associative learning paradigm we interface our behavioral simulation with the spiking mushroom body (MB) model introduced in (Jürgensen et al., 2024) (**Figure 6*C***). It implements a biologically realistic neural network model of the olfactory pathway according to detailed anatomical data using leaky integrate-and-fire neurons (Jürgensen et al., 2021). The MB network undergoes associative plasticity at the synapses between the Kenyon cells and two MB output neurons as a result of concurrent stimulation with an odor and a reward signal. Both, odor and reward is simulated as spike train input to the receptor neurons and the reinforcement signalling dopaminergic neuron, respectively. The model employs two output neurons (MB_+, MB_-), representing a larger number of MB compartments associated with approach/avoidance learning respectively (Saumweber et al., 2018). The initially balanced firing rates between MB_+ and MB_- are skewed after learning and encode the acquired odor valence (Owald et al., 2015; Owald and Waddell, 2015) here defined as:

$$MB_{out} = \frac{MB_+ - MB_-}{MB_+ + MB_-} \quad | \quad MB_{out} \epsilon [-1, 1]$$

We first trained the MB model via a classical conditioning experiment where, in each conditioning trial, it experiences an odor (conditioned stimulus, CS_+) in combination with a sugar reward during 5 min, following the standard training protocol in (Michels et al., 2017). 5 groups of 30 MB models undergo between 1 and 5 sequential conditioning trials (Weiglein et al., 2019). The resulting odor valence MB_{out} from each MB model was converted to an odor gain G via a simple linear transformation and used to generate a virtual larva (**Figure 6*B***). Each population of 30 larvae was then tested in an odor preference simulation. The larvae were placed on a dish in presence of the previously rewarded odor (CS+) and a neutral odor in opposite sides of the dish (Figure 6—video 1), again following standard experimental procedures (Michels et al., 2017). To obtain robust results we replicated the experiment 100 times per population with a different random seed for a total of 600 simulations. The obtained preference indexes (PI) are illustrated in **Figure 6*D***. The PI increase with increasing number of trials as well as its saturation resembles empirical observations (Weiglein et al., 2019). Note that the variability of the PI across the 100 simulations per condition is introduced solely by the behavioral simulation and resembles that seen across real experiments.

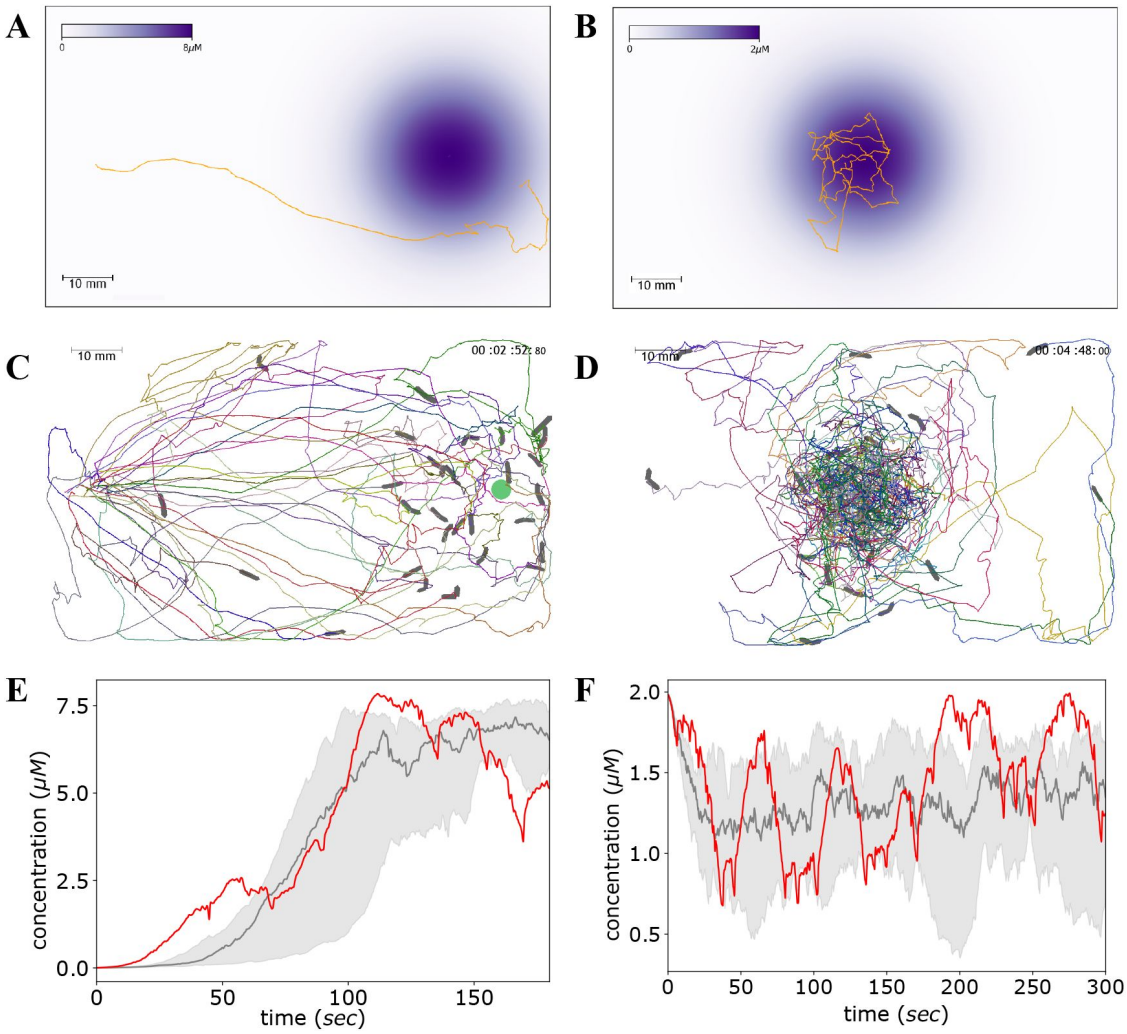

Figure 5.

Simulation of chemotaxis.

A: Experiment 1: A single odor source of 8.9μM peak concentration is placed on the right side of the rectangular arena creating a chemical gradient as indicated by the color scale. Larvae are placed on the left side facing to the right. Larvae are expected to navigate up the gradient approaching the source. A single larva trajectory is shown. This setup mimics the first experiment in (Gomez-Marin et al., 2011). **B:** Experiment 2: A single odor source of 2.0μM peak concentration is placed at the center of the rectangular arena. Larvae are placed in close proximity to the odor source. Larvae are expected to locally explore generating trajectories around and across the source. A single larva trajectory is shown. This setup mimics the second experiment in (Gomez-Marin et al., 2011). **C,D:** The trajectories of 25 virtual larvae during the two experiments. **E,F:** The odor concentration encountered by the virtual larvae as a function of time. Red curves refer to the single larva in A and B. Gray denotes the mean and quartiles of all 25 larvae in C and D. The simulation results fit well to the experimental estimates of concentration sensing during larval chemotaxis in (Gomez-Marin et al., 2011). (arena dimensions = 100×60mm, N = 30 larvae, experiment duration = 3 and 5 minutes respectively, simulation timestep =1/16).

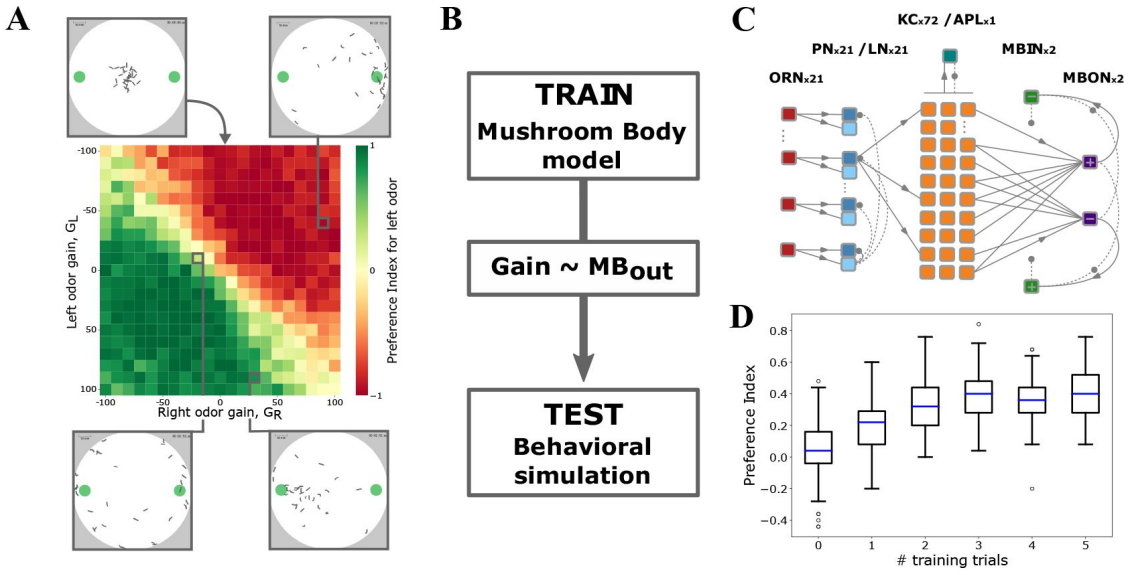

Figure 6.

Simulation of innate and learned odor preference.

A: A total of 25^2 simulations are shown with the resulting Preference Index for different gains of the left and right odor. On the top left the initial state is shown with the larvae randomly generated at the center of the dish. The final state of three additional simulations is depicted on the top right and bottom left and right. See Figure 6—video 1 for videos of two sample simulations. **B:** The pipeline used for coupling the Mushroom Body (MB) model with the behavioral simulation. First a MB model is trained via a classical conditioning experiment where olfactory input is combined with reward. The resulting odor valence MB_{out} is then converted to odor gain G via a simple linear transformation and used to generate a virtual larva. Finally the odor preference of a virtual larva population is evaluated in a behavioral simulation. **C:** The spiking neural network comprising the MB model. The number of neurons comprising each layer is indicated. **D:** The resulting PIs for 100 simulations per number of training trials. In each of the 100 simulations per condition a population of 30 virtual larvae was generated and evaluated using a different random seed, always bearing the exact same 30 odor gains derived from the respective group of 30 trained MB models (arena dimensions = 100*x*100mm, N = 30 larvae, experiment duration = 3 minutes, simulation timestep =0.1.

The current implementation only sequentially couples a trained MB model to be tested in a behavioral simulation. In the discussion we further elaborate on a possible extension featuring their closed-loop integration allowing for full behavioral simulations of both the training and the testing phase of the associative learning paradigm in a virtual environment.

Discussion

Experimental evidence for layered behavioral architecture

The proposed behavioral architecture is based on two underlying principles that justify our modular, hierarchically layered approach. Firstly, we attempt to horizontally structure it into behaviorally functional modules, each generating a well-defined behavior. Secondly, we vertically parse it into semi-independent layers, each under top-down modulation from the ones above but still capable of generating behavior independently of them. We summarize the neuroanatomical observations and behavioral experiments that support this approach.

The neural mechanisms underlying the three basic larval behaviors (**Figure 1*A***, basic layer) have been extensively studied. Crawling is characterized by fairly stereotypical repetitive strides. Initial contraction of the head and tail segments driven by a 'visceral piston' mechanism is followed by a laterally symmetric peristaltic wave traversing neighboring segments longitudinally from back to front (Heckscher et al., 2012) (**Figure 2*B***). Segmental central pattern generators (CPGs) coupled via intersegmental short- and long-range connectivity motifs involving also premotor neurons constitute the underlying neural circuitry (Kohsaka et al., 2019; Zarin et al., 2019; Mantziaris et al., 2020). Lateral bending results from asymmetric contraction of body musculature initiated at the thoracic segments (Berni, 2015). Feeding is generated via a network of mono- and multi-synaptic sensorimotor loops from enteric, pharyngeal and external sensory organs to motor neurons controlling mouth-hook movement, head-tilt and pharyngeal pumping (Miroschnikow et al., 2018b; Schoofs et al., 2024). A noteworthy facet of these behaviors is their autonomous generation with-out the need for descending control. This has been demonstrated by continued exploration even after brain ablation (Sims et al., 2019). Exploitation of a nutritious substrate requires all three basic behaviors as the larva constantly consumes food and re-positions its body. According to the layer-independence principle, we postulate that peripheral consummatory circuits at the sub-esophageal zone (SEZ) are also capable of autonomous exploitation.

Higher brain centers play a pivotal role in behavior modulation, with neurotransmitters like dopamine, serotonin, acetylcholine, and octopamine being key players in this intricate process (Berni et al., 2012; Zhang et al., 2013; Miroschnikow et al., 2018b; Malloy et al., 2019; Eschbach et al., 2020; Vogt et al., 2021). It has been suggested that the transition between exploration and exploitation (feeding) is acutely induced via dopaminergic signaling (Schleyer et al., 2020) while their long-term balance is regulated via hugin-mediated homeostatic neuromodulation (Schoofs et al., 2014). Identification of sensory pathways towards motor effector neuropiles further elucidates the role of interoception in behavioral modulation (Qian et al., 2018).

The neural mechanisms that underlie olfactory modulation of the basic locomotory behavior are also under intense investigation. Chemotactic approach and avoidance to innately valenced odors has been attributed to a predominantly innate pathway involving the antennal lobe (AL) and its direct projection to the lateral horn (LH), both in the larva (Schulze et al., 2015; Vogt et al., 2021) and in the adult fly (de Belle and Heisenberg, 1994; Strutz et al., 2014; Dolan et al., 2018). Modulation of behavior by learned odors strongly involves descending control by the MB in juvenile (Saumweber et al., 2018; Eschbach et al., 2020) and adult (Slater et al., 2015) stages. Both pathways have similar modulating effects on foraging behavior and are likely

integrated in a premotor network downstream of the AL (Schleyer et al., 2015a; Eschbach et al., 2020). In the sensorimotor loop, descending pathways involving the LH control cessation of crawling, possibly triggering sharper re-orientation when navigating down-gradient, facilitating chemotaxis (Tastekin et al., 2018). Finally, the internal homeostatic state (e.g. starvation vs. satiation) is implicated in behavioral regulation via neurotransmitter release at multiple levels, including AL, LH and MB and neuropeptide expression (Vogt et al., 2021; de Tredern et al., 2024).

Previous computational models of larva locomotion

Larva locomotion has been modeled extensively. Generative models of peristaltic crawling feature neural and/or neuromuscular dynamics. Sequential neural activity patterns across the segmental body can be generated by longitudinally repeated CPGs of paired excitatory and inhibitory population rate units, possibly involving proprioceptive feedback (Gjorgjieva et al., 2013; Pehlevan et al., 2016). This relatively abstract CPG model has been elaborated into a connectome-based circuitry of premotor and motor neurons. Under tonic activation the model exhibits two functionally distinct, though structurally overlapping, recurring patterns of neuromuscular activity, responsible for forward and backward peristalsis respectively (Zarin et al., 2019). At the other end of the modeling spectrum the contribution of the visceral-piston mechanism to the peristaltic cycle has been assessed in a biomechanical model (Ross et al., 2015). A more elaborate neuromechanical model, based on segmental localized reflexes and substrate frictional forces and assuming empirically informed axial and transverse oscillatory frequencies, has been shown to generate forward and backward crawling without the need for any neural activation pattern (Loveless et al., 2019).

Lateral bending has also been captured in statistical (Davies et al., 2015) or generative models (Wystrach et al., 2016; Loveless et al., 2019; Loveless and Webb, 2018). In the context of free-exploration it has been shown to be a byproduct of chaotic body neuromechanics underlying peristalsis (Loveless et al., 2019). In the context of chemotaxis it has been attributed to a distinct oscillatory process, autonomous (Wystrach et al., 2016) or semi-autonomous to crawling (Davies et al., 2015). In the former case oscillation is driven by mutual inhibition between excitatory-inhibitory circuits. In the latter, bending behavior is dissected into low-amplitude weathervaning while crawling and high-amplitude headcasting during crawl-pauses, an approach essentially equivalent to an attenuation of the lateral oscillation amplitude due to crawling. Autonomous exploration and chemotaxis can be generated by an integration of the neuromechanical and the independent lateral oscillator models (Loveless and Webb, 2018).

Apart from the rare occasion where such models implement mutually exclusive mechanistic hypotheses and are indubitably incompatible to each other, in most cases they are indeed complementary, overlapping, nested or disconnected in terms of the generative mechanisms they aim to capture and could potentially coexist under a broader control system. In this context our unifying architecture for larval behavior in which any partial model can be positioned can be valuable for modelers and roboticists, interested in a behavior-based synthetic approach.

Intermittent coupled-oscillator model for realistic locomotion

Each of the above models could be adjusted so that at minimum it generates the 2D translational and angular motion of a virtual body and could therefore populate the basic locomotory layer of the behavioral architecture. In this study we propose such a model, assembled in a synthetic approach by distinct modules, which either extend previously proposed locomotory models or integrate novel findings derived from our analysis of kinematic parameters. Given a simplified bisegmental body, a simple oscillator under frequency-regulating tonic activation is adequate for generating recurrent strides, efficiently summarizing the complex underlying neural and

neuromechanical dynamics (Heckscher et al., 2012; Mantziaris et al., 2020). Concerning angular motion, the previously introduced lateral oscillator model (Wystrach et al., 2016) meets the requirements and can therefore be coupled to the forward oscillator.

Our proposed model contributes two novel features. First, the intermittent nature of crawling as transitions between runs and pauses (**Figure 3*B***). And, second, the peristaltic cycle-phase dependent attenuation of angular motion (**Figure 3*D***). By combining these two features, the two bending behavioral modes termed weathervaning and head-casting are naturally generated via a phasic coupling between the two oscillators.

Behavioral intermittency

Larval locomotion is intermittent meaning that crawling runs are transiently intermitted by brief pauses. Transitions between these two behavioral states occur autonomously in stimulus-free conditions as during free exploration. Traditionally, in the context of movement ecology, intermittency has been studied in the spatial regime by characterizing the distributions of run distances and turn amplitudes occurring during brief stationary reorientation events (pauses). Power-law distributed runs in line with Levy-walk theoretical models (Günther et al., 2016; Sims et al., 2019) and diffusion-like kinematics have been reported (Klein et al., 2017). On the other hand, the turn amplitude distribution diverges from the original Levy-walk-predicted uniform distribution, as it is highly skewed towards small amplitudes, even if only significant reorientation events are taken into account (Sims et al., 2019). Moreover, the speed-curvature power-law relationship has been disputed (Zago et al., 2016; Marken and Shaffer, 2017; Zago et al., 2017; Marken and Shaffer, 2018). Regarding the temporal dynamics of intermittency, the duration distribution of these pauses is commonly neglected in traditional Levy-walk literature as they are usually characterised via the amplitude distribution of their resulting turning events. A recent analysis of larval tracks reported that the duration of pauses follows a power-law while that of activity bouts a log-normal distribution (Sakagiannis et al., 2020), partly in line with findings in adult-fly studies (Ueno et al., 2012; Reynolds et al., 2015). In our current study we found log-normal best fits for both pause duration and stridechain length. The disagreement over the pause duration might be attributed to the different timescale assessed in the two studies. Contrary to the high-resolution 180-sec tracks in this study, those analyzed previously lasted at minimum 1024 sec and up to 1 hour, allowing for the detection of longer pauses and imposing the necessity to fit over 4 orders of magnitude bypassing the apparent drop around 10 sec, also seen in Fig.3 of (Sakagiannis et al., 2020).

Concerning computational modeling, the contribution of behavioral intermittency to locomotion has not been adequately appreciated. Candidate generative models can either simply sample statistical distributions (Sims et al., 2019) or feature a generative mechanism that yields state transitions. Stochastic state-transitions have been included in a model of larva exploration yielding exponentially distributed epochs of both runs and stationary headcasts (Davies et al., 2015). At a mechanistic level, a recent study presented a simple binary-neuron model exhibiting state transitions between power-law and non power-law regimes via self-limiting neuronal avalanches and proposed a plausible underlying mechanism that explains initiation/cessation of crawling (Sakagiannis et al., 2020). All these attempts can be considered instantiations of a behavioral intermittency module (**Figure 1*B***) controlling cessation and re-initiation of crawling, central for generating epoch transitions. In the present study, we chose stochastic sampling from the empirically fitted parametric distribution models.

Crawl-bend interference

Crawling includes mouth hook motion. Specifically, the first phase of a crawling stride consists of concurrent forward motion of head and tail segments, aided by a 'visceral pistoning' mechanism that generates forward displacement of the gut. Subsequently, the mouth hooks anchor the head to the substrate so that the second phase of peristaltic motion can drag all other segments forward as

well, completing the stride (Heckscher et al., 2012). With respect to turning, it is still debated whether individual turns should be considered as discrete reorientation events that are temporally non-overlapping with crawling bouts (Sims et al., 2019), or whether lateral bending occurs in an oscillatory fashion generating turns both during crawling (weathervaning) and during pauses (headcasts) (Wystrach et al., 2016; Thane et al., 2019). The latter is supported by detailed eigenshape analysis confirming that larvae rarely crawl straight, rather forward locomotion is always accompanied by continuous small amplitude lateral bending (Szigeti et al., 2015). It follows that crawling does not exclude bending, rather the two strictly co-occur.

Crawling and bending partially recruit the same effector neural circuitry and body musculature at least at the level of the thorax. Peristaltic motion during crawling includes sequential symmetric bilateral contraction of all segments while bending occurs due to asymmetric unilateral contraction of the thoracic segments. This partial effector overlap could result in interference between the two processes. Indeed here we report a phase-dependent attenuation of angular velocity (**Figure 3*C***). Attenuation reaches a minimum at a specific phase of the cycle, closely preceding the increase of head forward velocity (**Figure 3***C:inset*). This coincides with the stride phase when the head stops being anchored to the substrate and is therefore free to move laterally. When applying a phase-dependent Gaussian attenuation kernel on angular velocity we managed to accurately reproduce the empirical relation (**Appendix 2-Figure 1*C***).

A reasonable hypothesis would then be that the asymmetric thoracic contraction generating lateral bending becomes easier when the head is not anchored to the substrate therefore during a specific phase interval of the stride cycle. We propose that crawling phasically interferes with lateral bending because of these bodily constraints. A consequence of this hypothesis is that the amplitude of turns generated during crawl-pauses (headcasts) is larger in comparison to those generated during crawling (weathervaning) because during pauses the crawling interference to lateral bending is lifted. It is this phenomenon that dominates the description of larva exploration as a Levy-walk with non-overlapping straight runs and reorientation events, where weathervaning is neglected (Günther et al., 2016; Sims et al., 2019). Nevertheless, it has been included in a previous stochastic model of larva exploration, where it has been treated as entirely distinct to headcasts, occurring during crawl-pauses, via the application of differential constraints on both the angular velocity and the resulting turn amplitude (Davies et al., 2015). By implementing a coupled-oscillator locomotory model we avoid such dual treatment of headcasts and weathervaning.

Olfactory learning in closed loop behavioral simulations

We have reproduced the results of a basic associative learning experiment in the fruit fly larva (Schleyer et al., 2018; Jürgensen et al., 2024) by the open-loop simulation of classical conditioning trials and subsequent closed-loop behavioral simulation during the memory retention test for individual virtual larvae (**Figure 6***B-D*). This modeling approach can be extended in multiple ways. First, the larva demonstrates a number of additional learning capabilities such as differential conditioning (Schleyer et al., 2011; Schleyer et al., 2015b; Schleyer et al., 2018), extinction learning (Felsenberg and Waddell, 2019; Lesar et al., 2021), and relief learning (Saumweber et al., 2018; Weiglein et al., 2019; Gerber et al., 2014; König et al., 2018). Interfacing neural network simulations with candidate circuit and synaptic mechanisms of plasticity with our behavioral model allows to directly compare virtual and empirical behavioral experiments, both at the level of individual and group assays. Second, while information about odor concentration is provided through active sensing, simultaneous input from a feeder module could activate the dopaminergic reward pathway required for synaptic plasticity in the mushroom body. At the same time, simulation of food uptake and energy expenditure will regulate the agent's energy homeostasis. This would further allow realistic foraging scenarios with food depletion and competition among animals.

Closing the loop from active sensing to associative memory formation and behavioral control requires to synchronize a (spiking) neural network at the adaptive layer with the sensory module (reactive layer), and the locomotory and feeding modules at the basic layer (**Figure 1*A***). This will enable the simulation of virtual larvae experiencing spatial and temporal dynamics in a virtual environment or on a robotic platform (Helgadottir et al., 2013), and it will allow to test model hypotheses on sensory-motor integration and to infer predictions for experimental interventions such as optogenetic stimulation (Saumweber et al., 2018) or genetic manipulation (Saumweber et al., 2011; Michels et al., 2011; Widmann et al., 2016; Springer and Nawrot, 2021).

Materials and Methods

We first describe the experimental dataset and the software package used in this study. Then we explicitly describe each of the computational modules that comprise the proposed behavioral architecture. The metrics used throughout the analyses of both real and simulated datasets are described in Appendix 1.

Dataset description

The larva-tracking dataset was obtained by M.Schleyer and J. Thoener at the Leipzig Institute of Neurobiology. It consists of 31 experimental groups. Each group of ~ 30 third-instar larvae (Canton S) was placed on agarose-filled Petri dishes of 15 cm diameter with no particular stimuli and video-filmed from above at a framerate of 16 *Hz* for 3 minutes. During video-tracking, 12 points are detected along the longitudinal axis of each larva. For the present study groups were pooled together in a single population and timepoints of detected collisions have been excluded. A subset consisting of the 200 most complete, uninterrupted tracks was selected. The x-y coordinate timeseries have been filtered with a first-order butterworth low-pass filter with a cutoff frequency of 2 *Hz* in order to decrease tracking-related noise but retain the behaviorally relevant crawling frequency of $f_C \simeq 1.4\ Hz$. The effect of inadequate and excessive filtering is illustrated in Video 3.

Software package and code availability

All data processing, data analyses and model simulations were performed using our freely available python package *Larvaworld* (*https://pypi.org/project/larvaworld/*), a behavioral analysis and simulation platform for *Drosophila* larva. In *Larvaworld* simulated and empirical data are treated indistinguishably, meaning that the exact same analysis pipeline and behavioral metrics are applied to both. Concerning modeling, the behavioral architecture we propose here comprises the backbone for the construction, extension, configuration and fitting of behavioral models in *Larvaworld*. Moreover, the intermittent coupled-oscillator model as introduced in the present manuscript is available for simulations alongside other preconfigured models. Extensive documentation can be found at *https://larvaworld.readthedocs.io*.

Inter-individual variability in virtual populations

The dominant methodological approach in animal behavioral modeling yields generative models that capture the average behavior of a group of animals. Each model therefore strives to generate behavior resembling that of an idealized average animal, neglecting inter-individual variability. To challenge this, we contrast it with a group-level modeling approach that preserves individuality within the population. Here, the generative model is instantiated by a group of non-identical animats, with parameter combinations drawn from a fitted multivariate normal distribution, preserving their pairwise correlations. We first measured a number of endpoint parameters across a population of 200 larvae and fitted a multivariate Gaussian distribution. A subset of three

of these parameters, body-length l, crawling frequency f_C, and scaled stride displacement $\mathring{d}_S$ is shown in **Figure 7**. The univariate and bivariate empirical distributions are illustrated in blue, while the bivariate projections are shown in red.

Module definition

The building blocks comprising the behavioral architecture are described as separate modules of defined input and output. The modular architecture only specifies the required placeholders and remains agnostic to the specific module implementations. Nevertheless, in what follows, the specific modular implementation for the proposed coupled-oscillator locomotory model will be described alongside the general module-placeholder definition. Calibration of each module on the empirical dataset and parameter specification is described in detail in **Appendix 2**. The parameters of the final model configuration are shown in **Table 1**.

Bisegmental body

The virtual body is the architecture's interface with the simulated environment. It moves through space under the control of the motor effectors comprising the basic locomotory layer of the architecture, and thereby repositions the sensors introduced in the intermediate reactive layer. For the simplest case of 2D locomotion it is therefore imperative to define at least the forward velocity v and the angular orientation velocity ω. Any locomotory model dynamically generating these two parameters is adequate for an abstract point-body or a single-segment body implementation.

For the proposed coupled-oscillator model we choose a bisegmental body implementation, additionally featuring the bending angle (posture) θ_b between the front and rear segment. Following (Wystrach et al., 2016), the body is modeled as a torsional analog of the mass-spring damper model. Torque angularly accelerates an inertia I against angular damping that resists motion and viscoelastic forces that resist deformation (lateral bending). The original mass-spring damper model and its torsional analogue are defined by the equations:

$$m\dot{v} = F_{ext} - zv - kx$$

$$I\dot{\omega} = T_{ext} - Z \cdot \omega - K \cdot \theta_b.$$

We introduce simplified coefficients for angular damping $z = {}^{Z}/_{I}$ in sec^{-1} and bend deformation resistance $k = {}^{K}/_{I}$ in sec^{-2} as the dimensions of inertia ML^2 cancel out. Similarly, the external torque-per-unit-of-inertia is now in angular acceleration units $\dot{\omega}_{ext} = {}^{T_{ext}}/_{I}$ This drive is generated by the turner module and will be described below. Additionally we introduce a crawl-phase dependent suppression of angular motion during crawling described by the coefficient $c_{CT}(t) = c_{CT}(\Phi_C)$, which will also be described below.

$$\dot{\omega} = \dot{\omega}_{ext} - k \cdot \theta_b - z \cdot \omega - (1 - c_{CT}) \cdot \omega \qquad (1)$$

In its original implementation (Wystrach et al., 2016) the torsional body model deliberately neglects two aspects of the real turning behavior of the larva. First, there is no distinction between bending ω_b and orientation ω angular velocities. Second, there is no correction of the bending angle θ_b due to forward motion. We tackle the first via the bisegmental body so that ω_b between

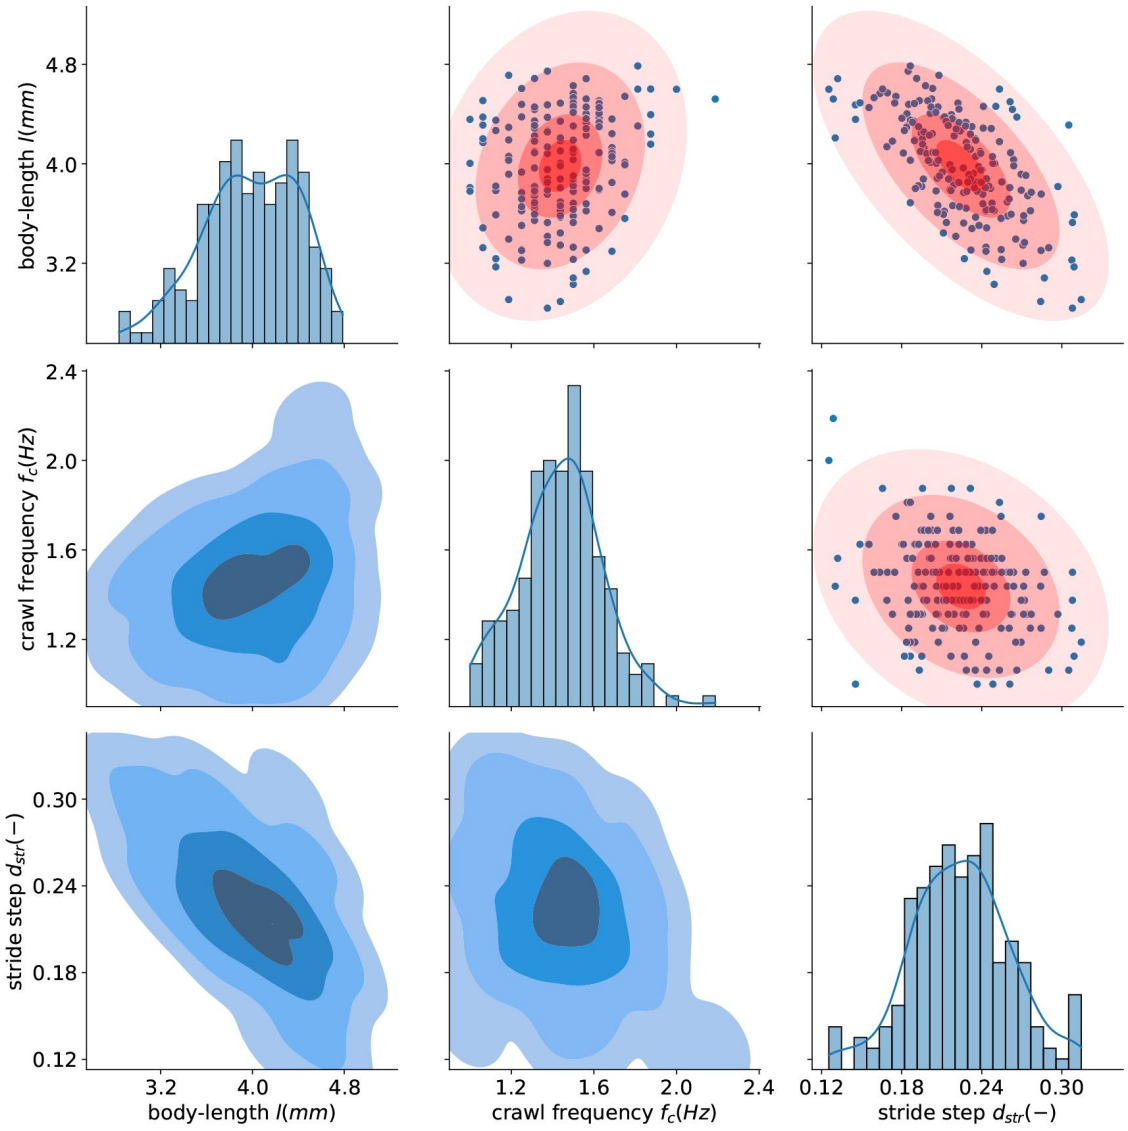

Figure 7.

Individuality: empirical (blue) and fitted (red) parameter distributions.

Diagonal: Histogram and kernel density estimates (KDE) for body-length *l*, crawling frequency f_C and mean scaled displacement per stride $\mathring{d}_S$ across a population of 200 larvae in the experimental dataset. Below: Bivariate projection of 3-dim. KDE outlined contours for each parameter pair. Above: Red ellipses represent the bivariate projections of the 3-dim. fitted Gaussian distributions at 0.5, 1, 2 and 3 standard deviations. In our model this Gaussian is used to sample a parameter set for each individual larva. The blue dots denote the empirically measured parameters.

	Parameter	Symbol	Value	Unit
PHYSICS	torque coefficient	c_T	0.5	sec^{-2}
	angular damping	z	1.0	sec^{-1}
	body spring constant	k	1.0	sec^{-2}
	bend correction coefficient	c_b	1.0	-
BODY	length	l	4.0	*mm*
	number of body segments	-	2	*# segments*
TURNER	tonic input	I_T^0	21.43	-
	input range	$[I_T^{min}, I_T^{max}]$	[10. 40.]	-
	time constant	tau_T	0.1	sec
	spike response steepness	n_T	2.37	-
CRAWLER	crawling frequency	f_C	1.42	*Hz*
	maximum scaled velocity	$\dot{v}_{max}$	0.51	$body-lengths/sec$
	stride distance mean	$\bar{\bar{d}}_S$	0.24	$body-lengths$
	stride distance std	$\bar{\bar{d}}_S$	0.04	$body-lengths$
	max velocity phase	ϕ_C^v	3.49	*rads*
INTERFERENCE	suppression coefficient	c_{CT}^0	0.46	-
	suppression relief coefficient.	c_{CT}^1	0.54	-
	max relief phase	ϕ_C^ω	2.05	*rads*
INTERMITTER	run length distribution	N_R	Lognormal(m=1.4, s=1.15)	-
	run length range	$[N_R^{min}, N_R^{max}]$	[1, 142]	*# strides*
	pause duration distribution	t_P	Exp(b=1.0)	-
	pause duration range	$[t_P^{min}, t_P^{max}]$	[0.12, 16.0]	*sec*

Table 1.

Locomotory model configuration.

The parameters of the calibrated average locomotory model, organized per module.

the front and rear vector is distinct from the front vector's ω. Regarding the second, we introduce a simple linear bending-angle correction as the rear vector is aligned to the front vector's orientation during forward motion, according to the equation:

$$\theta_b = \begin{cases} \check{\theta}_b \cdot (1 - c_b^{max}) & if \quad 0 \le c_b^{max} < 1 \\ 0 & if \quad c_b^{max} \ge 1 \end{cases}$$

$$c_b^{max} = \frac{2dc_b}{l}$$

where d is the linear displacement during a timestep, $\check{\theta}_b$ and θ_b is the bending angle before and after correction, l is the body length and $c_b = 1$ is the bend correction coefficient.

Concerning the forward motion, the simple kinematic implementation of the crawler module directly generates forward velocity v without taking into account any biomechanical dynamics.

Crawler module

The crawler module generates forward velocity v, under a continuous (tonic) activation signal I_C and an additional intermittent initiation/cessation signal generated by the intermittency module. Velocity generation involves three parameters : the larval body-length l, the scaled displacement per stride $\mathring{d}_S$ and the crawling frequency f_C. The latter linearly reflects the tonic input I_C, which is kept constant during the short-duration simulations. For an individual virtual larva $l, \bar{\mathring{d}}_S, \tilde{\mathring{d}}_S$ and f_C are set during initialization, where $\bar{\mathring{d}}_S, \tilde{\mathring{d}}_S$ are the mean and standard deviation of $\mathring{d}_S$.

In the present study crawling behavior is modeled as an oscillatory process. Each oscillation generates a cycle of forward velocity v increase-decrease resulting in displacement $l \cdot \mathring{d}_S$ of the larva along the axis of its front-segment, simplistically modeling the result of exactly one peristaltic stride. After termination of a stride a new $\mathring{d}_S$ is sampled. An analytically tractable curve is fitted to the average empirical velocity curve measured during strides by setting all 5 crawler-related parameters to the median empirically measured value (**Appendix 2-Figure 1*B***):

$$\mathring{d}_S = N\left(\bar{\mathring{d}}_S, \tilde{\mathring{d}}_S\right) \tag{2}$$

$$v = l \cdot \mathring{d}_S \cdot f_C \cdot \left(\mathring{v}_{max} \cos\left(\phi_C - \phi_C^v\right) + 1\right) \tag{3}$$

where Φ_C is the instantaneous phase of the crawler oscillation iterating from 0 to 2π during an oscillatory cycle, $\mathring{v}_{max}$ is the maximum scaled velocity during the cycle and ϕ_C^v is the phase where the maximum velocity occurs.

Turner module

The turner module is defined as the generative model of a torque-like output $A_T(t)$ under a continuous activation signal $I_T(t)$. The (considered non-dimensional) output is scaled to angular acceleration by a coefficient c_T in sec^{-2} and applied to the body as external bending drive (**Equation 1**):

$$\dot{\omega}_{ext}(t) = c_T \cdot A_T(t), \tag{4}$$

which is then applied to the body yielding the instantaneous angular velocity $\omega(t)$ and angular acceleration $\dot{\omega}(t)$ bending the virtual body laterally. It has been previously suggested that an underlying oscillatory process drives alternating bending to the left and right side (Wystrach et al., 2016). We empirically confirmed this by detecting a slow,variable rhythm $f_T \approx 0.4$ on the angular velocity timeseries of individual larvae. Based on this observation we adopt an oscillatory approach for the alternating lateral bending and extend the lateral oscillator model described in

(Wystrach et al., 2016). It consists of two mutually inhibiting components each having an excitatory and inhibitory neuron ($E_{L\&}C_L$ vs $E_{R\&}C_R$). The spike-response of a neuron to a cumulative input x is given by $R(x, h)$ where h is the half-response threshold. The neuromodulatory activation $I_T(t)$ excites both components but also affects the gain $g(I_T)$ and time constant $\tau_H(I_T)$ of an additional adaptation component H_x of each neuron. The left and right components quickly settle in antiphase, while adaptation ensures that periodic transitions occur. Under the baseline activation I_T^0 the oscillations occur at a dominant frequency f_T. Perturbations of this external drive cause transient changes in both amplitude and frequency because both, up to transient loss of oscillation. This feature is exploited during olfactory modulation. The resulting oscillator activity is defined as the instantaneous difference in the firing rates $A_T = (E_L - E_R)$.

$$R(x,h) = \begin{cases} \frac{mx^n}{h^n + x^n} & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \tag{5}$$

$$g(I_T) = 6 + (0.09 \cdot I_T)^2$$

$$\tau_H(I_T) = \frac{3}{1 + (0.04 \cdot I_T^2)}$$

$$\tau_T \dot{E}_L = -E_L + R(I_T + W_{ee}E_L - W_{ec}E_R, 64 + g(I_T)H_{EL})$$

$$\tau_T \dot{C}_L = -C_L + R(I_T + W_{ce}E_L - W_{cc}C_R, 64 + g(I_T)H_{CL})$$

$$\dot{H}_{EL} = \frac{E_L - H_{EL}}{\tau_H(I_T)}$$

$$\dot{H}_{CL} = \frac{E_L - H_{CL}}{\tau_H(I_T)}.$$

Calibration of the turner module parameters in order to achieve a resulting angular motion that is in agreement to the empirical observations requires concurrent specification of the angularmotion related angular damping z, bend deformation resistance k and torque scaling coefficient c_T. The calibration process is described in detail in Appendix 2.

Crawler-turner coupling

We couple the crawler and turner oscillators by imposing a crawler-phase dependent attenuation $c_{CT}(\Phi_C)$ on the angular velocity ω (Equation 1). During the entire stride-cycle ω is scaled by a baseline suppression coefficient c_{CT}^0. Our kinematic analysis revealed that this baseline attenuation is partially lifted during a phase interval of the stride-cycle. The maximum additional relief is defined by the relief coefficient c_{CT}^1. Concerning the mode of transition from c_{CT}^0 to $c_{CT}^0 + c_{CT}^1$ and back to c_{CT}^0, we study two implementations, each defined by an additional parameter. In the “square” mode there is an acute transition to maximum relief and back to baseline for a specific interval of the cycle $[\phi_C^{\omega_0}, \phi_C^{\omega_1}]$. In the “phase” mode the transition is

described as a Gaussian kernel reaching maximum relief at phase ϕ_C^{ω} when crawl-induced angular attenuation is minimum and the anterior body is maximally free to bend laterally, as shown in the equation:

$$c_{CT} = \begin{cases} c_{CT}^1 \cdot A_{CT} + c_{CT}^0, & \text{during runs} \\ 1, & \text{during pauses} \end{cases} \tag{6}$$

$$A_{CT} = \begin{cases} \exp\left(-\frac{(\phi_C - \phi_C^{\omega})^2}{2}\right) & \text{phase mode} \\ 1 \quad \text{if } \phi_C^{\omega_0} \le \phi_C \le \phi_C^{\omega_1} \\ 0 \quad \text{otherwise} & \text{square mode} \end{cases} \tag{7}$$

The parameter set is optimized to match the average empirically measured angular motion observed during the stride cycle. The optimization process used is described in Appendix 2. The optimal Gaussian kernel is shown in **Figure 3**.

Intermittency module

We define the intermittency module as a placeholder for any model capable of generating transitions between runs and pauses (or equivalently runs and headcasts) (**Figure 1*B***). During pauses, the input to the crawler module $I_C = 0$ and therefore forward velocity $v = 0$ while during runs this is unaffected. For the specific study we implement a statistical model where the duration for both the run and pause states is sampled at each initiation from a distribution that has been fitted to the experimental data. In the case of pauses the duration is normally measured in time units. For runs we employ an equivalent metric and therefore measure them as number of consecutive crawling strides (stridechains). The fitted distributions are shown in **Table 1**.

Olfactory sensor

Olfaction is introduced in the intermediate layer of the behavioral architecture allowing chemotactic behavior. The olfactory sensor is located at the front end of the virtual larva therefore any reorientation and/or displacement influences sensory input. As in (Gomez-Marin et al., 2011) we assume that olfactory perception A_O relates to changes in odor concentration C according to the Weber-Fechner law, meaning that $\Delta A_O \sim \ln \Delta C$. We further add a decay term that slowly resets A_O back to 0. The rate of change is given in Equation 8 where $c_O = 1$ is the olfactory decay coefficient, G_i is the gain for odor i and C_i the respective odor concentration:

$$\dot{A}_O = -c_O A_O + \sum_i G_i \cdot \frac{\dot{C}_i}{C_i} \qquad with \qquad -1 \le A_O \le 1. \tag{8}$$

Concerning how the perceived olfactory stimulation A_O modulates exploratory behavior we adopt the mechanism proposed in (Wystrach et al., 2016) (**Figure 5**). According to this model, the turner activation I is perturbed from its baseline value I_T^0 within a suitable range $I_T^{range} = [I_T^{min}, I_T^{max}] = [10, 40]$:

$$I_T = I_T^0 + A_O\left(I_T^{lim} - I_T^0\right) \qquad where \qquad I_T^{lim} = \begin{cases} I_T^{max} & if \quad A_O \ge 0 \\ I_T^{min} & if \quad A_O \le 0. \end{cases} \tag{9}$$

Mushroom body module

The underlying spiking network model of the *Drosophila* larva olfactory pathway and mushroom body (*MB*) implements the model published by (Jürgensen et al., 2024). It consists of conductance-based leaky integrate-and-fire neurons. It encompasses 21 olfactory receptor neurons (*ORNs*) and the same number of projection neurons (*PNs*) and local interneurons (*LNs*) that receive input from the *ORNs* in a one-to-one manner (**Figure 6*C***). *LNs* form inhibitory connections with the *PNs*. In the *MB*, each of the 72 Kenyon cells (*KC*) receive excitatory input from 2-6 random *PNs*. Feedback inhibition to the *KCs* is provided via the single GABAergic anterior paired lateral neuron (*APL*). All *KC* provide excitatory input to two *MB* output neurons (*MBONs*). The two *MBONs* represent two different types of output neurons that exist in the larval *MB* and either represent approach or avoidance behavior. Plasticity at the *KC* :: *MBON* synapses facilitates the association of odors with reward or punishment (associative learning). *KC* :: *MBON* synapses employ a two factor learning rule: Pre-synaptic activation by an odor stimulus (action potential of the KC) triggers an exponentially decaying eligibility trace $e(t)$, which determines the window of opportunity for synaptic change. Additional neuromodulatory input $r(t)$ from one of the two *DANs* will lead to a reduction in synaptic strength (**Figure 6*C***) proportional to $e(t) \cdot r(t)$. The acquired imbalance between the outputs of the two *MBONs* (behavioral bias) represents the association of odors with rewards/punishments. An initially balanced interaction of excitatory and inhibitory feedback components from *MBONs* onto both *DANs* encodes the models' learning history in form of a prediction error. The concept of prediction error suggests that associative learning about stimulus A is proportional to the difference between the reinforcement currently received with stimulus A minus the reinforcement predicted by stimulus A (previously learned association) (Rescorla, 1972). In vertebrates an implementation of this error signal has been demonstrated in dopaminergic neurons (Schultz, 2015, 2016) and a similar mechanism is proposed for learning in the insect *MB* (Riemensperger et al., 2005; Springer and Nawrot, 2021; Bennett et al., 2021; Jürgensen et al., 2024). The time-resolved behavioral bias is used to compute gains for odors, which bias motor output of the virtual larvae.

Videos

Here we provide the videos cited in the text. All videos are derived from real or simulated experiments as visualized via the *Larvaworld* software package, after some additional editing.

Acknowledgements

This project received funding from the Ministry of Culture and Science of the State of Northrhine Westphalia through the science network 'iBehave' (*https://ibehave.nrw/*) within the program "Netzwerke 2021", and from the German Research Foundation through the Research Unit 'Structure, Plasticity and Behavioral Function of the *Drosophila* mushroom body' (DFG-FOR 2705, grant no. 365082554, *https://www.unigoettingen.de/en/601524.html*). PS received a stipendship within the Research Training Group 'Neural Circuit Analysis' (DFG-RTG 1960, grant no. 233886668). We would like to thank Michael Schleyer and Juliane Thoener for providing the tracking dataset and for valuable discussions on the data analyses. We thank Bertram Gerber for valuable discussion and Anna Morozova for substantial support with the generation of the video and figure material.

Additional files

Video 1. *Locomotory model for Drosophila larva. The function of the locomotory model in Figure 1B is illustrated by gradually integrating its 4 modules (crawler, turner, oscillator-coupling, crawling-intermittency). In each of the 6 videos the implemented modules are shown in the inset.*

Video 2. *Free exploration simulation. A population of 25 real (left) or virtual (right) larvae is placed on a dish and left to freely explore. The body of the real larvae has been simplified into 2 segments as described in Video 4.*

Video 3. *Filter selection. The effect of inadequate or excessive filtering of the empirical larva recordings is illustrated. The left video shows the jittery original recording while the effect of lowpass filtering at cut-off frequencies of 4 Hz, 2 Hz and 0.5 Hz is shown on the rest. Selection of an intermediate 2 Hz cut-off frequency eliminates the unrealistic jitter while preserving the behaviorally relevant crawling frequency.*

Video 4. *Bisegmental larva-body simplification. The first video shows the original larva body as recorded by the tracker. 12 points are tracked along its longitudinal axis defining 11 segments while 22 points constitute the body contour. In the second video the body contour is dropped. In the third video an artificial rectangular contour is added for each body segment. In the last video the body-midline is segmented into 2 segments. The absolute head orientation angle θ is preserved while the single bending angle between the 2 segments is defined as* $\theta_b = \sum_{i=1}^{5} \theta_i$

Figure 2—video 1. *The full-length trajectory (A, E) colored according to angular and forward velocity.*

Figure 2—video 2. *The short track epoch depicted in the insets (A, E) colored according to forward and angular velocity.*

Figure 4—video 1. *The temporal course of dispersal for real (left) and virtual (right) larvae shown in (A).*

Figure 5—video 1. *Time course of the two simulated chemotaxis experiments.*

Figure 6—video 1. *Two odor preference simulations, one with an appetitive and one with an aversive odor source placed on the left side of the dish. A non-valenced odor source is placed on the right side.*

Appendix 1

Metric definition

Body length

The instantaneous body-length of an individual larva fluctuates during crawling due to subsequent stretching and contraction. Its histogram is well fitted by a Gaussian distribution (data not shown). Therefore individual larva length l is defined as the median of the midline length across time (total length of the line connecting all 12 midline points). All spatial parameters including

displacement and velocity can be scaled to this body-length, converting spatial units *m* or *mm* to dimensionless body-length units. Scaled spatial metrics are denoted by an additional° over the metric symbol.

Segmentation and angular metrics

To specify the body segmentation providing the most suitable contact/rotation point for the definition of the bending ω_b and orientation ω angular velocities we analyse their relationship in a subset of 40 larvae. Tracking of 12 midline points allows computation of the absolute orientation of 11 body-segments and the respective 10 angles θ_{1-10} between successive body segments (**Figure 2*F***). We define θ as the head-segment absolute orientation in reference to the x-axis because this defines the movement orientation of the animal. We ask how ω results from the bending of the body as this is captured by the 10 angular velocities ω_{1-10}. The regression analysis depicted in **Figure 1*B*** shows as expected that ω depends primarily on the front angular velocities while this dependence decays as we move towards the rear segments, in line with previous studies (Lahiri et al., 2011). Timeshift analysis also shows that the front 3 angles change concurrently while angles further down the midline are increasingly lagging behind (data not shown). The correlation analysis depicted in **Figure 1*C*** shows that the sum of the front 5 angular velocities best correlates to ω. In other words the cumulative body bend of the front 5 segments best predicts head reorientation. Therefore we define the reorientation-relevant bending angle as $\theta_b = \sum_{i=1}^{5} \theta_i$ (**Figure 2*F***). The remaining 5 angles between the rear body-segments can safely be neglected as they do not contribute to reorientation. This analysis results in a segmentation of the body in a front and a rear segments of length ratio 5 : 6. The segmentation process is demonstrated in Video 4.

Forward velocity

To define forward velocity we need to choose which midline-point is most suitable to track and which velocity metric to use for defining the start and end of a stride. To this end we perform stride annotation of 3-minute tracks of a population of 20 larvae using each of 24 candidate instantaneous velocity metrics, namely the velocities of the 12 points, the component velocities of the rearest 11 points parallel to their front segment's absolute orientation and finally the centroid velocity. To compare the candidate metrics we compute the spatial cv_s and temporal cv_t coefficient of variation of the annotated strides for each larva to assess how variant their time duration and displacement is. We finally compute the mean $\overline{cv}_s$ and $\overline{cv}_t$ across individuals. In **Figure 1*A*** the spatiotemporal stride variance is shown for each candidate metric. We choose the metric that provides the minimal spatial and temporal stride variance, assuming that strides of an individual larva are more or less stereotypical in both duration and displacement (Heckscher et al., 2012). Our study reveals that the centroid velocity is the most suitable metric for stride annotation. All spatial metrics are therefore computed via this point's displacement.

Track tortuosity

Track tortuosity is quantified by the straightness index (S.I), a metric previously used in larva-track analysis (Sims et al., 2019), computed by advancing a fixed time window along the track and calculating at each point the ratio of the straight line distance to the actual distance travelled. This index, which varies from 0 (no movement) to 1 (straight line movement), can capture very well the complexity of the movement at various scales (set by the window time frame) throughout the track. As the time window decreases the smoothing effect is also reduced revealing increasing track details. Different window sizes have been used ranging from 2 to 20 seconds in order to capture both large-scale spatial trajectories and small-scale local movements. Tortuosity was computed for each larva through time and revealed changes in movement as larva alternated between straight line relocation, changes of direction and different degrees of tortuosity. Hence, a

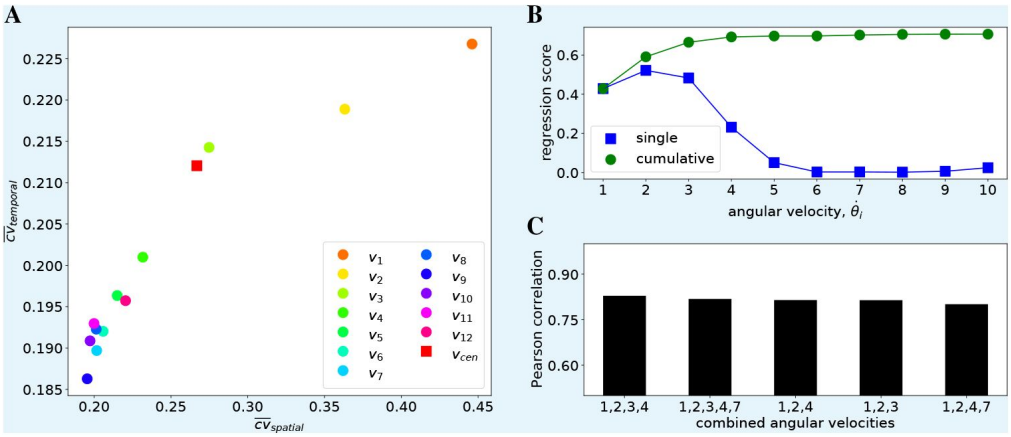

Appendix 1—figure 1.

Segmentation and velocity definition.

A: Forward velocity definition. 13 candidate velocity metrics are compared for use in stride annotation of 3-minute tracks of a population of 30 larvae. For each candidate the mean coefficient of variation of temporal duration $\overline{cv_t}$ and spatial displacement $\overline{cv_s}$ of the annotated strides is shown. Midline point 9 velocity provides the most temporally and spatially stereotypical strides, therefore it is selected as the reference forward velocity for stride annotation and model fitting. v_{cen} : centroid velocity, $v_1 - v_{12}$: 1^{st}-12^{th} point's velocity. **B:** Regression analysis of individual and cumulative angular velocities $\omega_{i=1-10}$ to orientation angular velocity ω. When considered individually, ω_2 best predicts reorientation with the ω_3 and ω_1 following. When considered cumulatively the anterior 4 ω_i allow optimal prediction of reorientation velocity. **C:** Correlation analysis of the sum of all possible ω_i combinations to ω. The sum $\sum_{i=1}^{4}\omega_i$ shows the highest correlation therefore we define $\theta_b = \sum_{i=1}^{4}\theta_i$ as shown in A. For illustration purposes only the 5 highest correlations are shown.

change in the S.I. along a larva's track captures the magnitude of the change in movement pattern from intensive, area-restricted searching movements (higher tortuosity) to extensive, straighter line movements (lower tortuosity), and vice versa, across a wide range of spatial scales.

Epoch annotation

Strides (*S*) are annotated using the scaled forward velocity $\mathring{v}$ First we apply fourier analysis to detect the dominant crawling frequency within a suitable range $1 \leqslant f_C \leqslant 2.5$ (**Figure 3*A***). From this we derive the reference stride duration $t_S = f_C^{-1}$ Then epochs are annotated under a number of constraints (**Figure 2*C***):

- Each stride is contained between two $\mathring{v}$ local minima.
- The $\mathring{v}$ local maxima contained in the epoch needs to exceed a threshold $\mathring{v}_{max} \geqslant 0.3$.
- The duration of the epoch *t* needs to range within $0.75 \leqslant {}^{t}S/_{\leqslant}2$. This allows individual strides to temporally vary without overlapping so that adjacent strides can be concatenated in stridechains.

After stride annotation the displacement due to each individual stride is computed for each larva and divided by the larva's body-length (**Figure 2***A-C*). The individual distributions are well fitted by Gaussians (data not shown). Therefore $\bar{d}_S, \bar{d}_S$ are defined as the average and standard deviation of the scaled displacement per stride for each larva.

Crawling runs (*R*) are defined as uninterrupted sequences of successive strides, also termed stridechains (**Figure 2*C***). Stridechain length is the number of concatenated strides in a run and is a discrete metric equivalent to the continuous crawling run duration. If the locomotory model under evaluation does not generate *v* oscillations and therefore strides can not be detected, runs are annotated using the plain forward velocity *v* timeseries as in a previous study (Sakagiannis et al., 2020). We first define a suitable threshold v_{thr} by detecting the minima of the pooled *v* bimodal distribution. Then we define runs as epochs where constantly $v \geqslant v_{thr}$. Pauses are then defined as epochs containing no strides (or equivalently not overlapping with runs) and during which $\mathring{v} \leq 0.3$ (or $v \leq v_{thr}$).

Turn epochs (*T*) are contained between pairs of successive sign changes of orientation angular velocity *ω*. Annotated epochs of left (T_L) and right (T_R) turns yield the respective turn angle amplitudes θ_{T_L} and θ_{T_R} as the absolute total change of orientation angle *θ* (**Figure 2*G***), which can then be pooled into the overall absolute turn-angle amplitudes θ_T. Notably by this definition turns include both headcasts (*H*) occurring during crawl-pauses and weathervaning (*W*) occurring during runs. Turns happening exclusively during runs and pauses yield the turn-angle amplitudes of weathervaning θ_W and headcasts θ_H respectively.

Appendix 2

Model calibration

Here we describe in detail the calibration procedure for a locomotory model of an average idealized individual based on a reference dataset of 150 larvae tracked while freely exploring a Petri-dish over 3 minutes. The pipeline consists of 3 initial steps, which can be performed independently of each other and set the configuration parameters of the crawler, intermittency and turner modules respectively. In a final step that builds upon the results of the initial three, the parameters of the crawl-bend phase-locked interference are set. The first two steps are directly defined by the kinematic analysis results. The third and the final steps involve an optimization process to reach the parameter set best fitting the empirical data. Notably, along with the intrinsic

parameters of a module, a number of additional module-related parameters might need to be defined, such as the body-length for the crawler module and the angular-motion relevant physical parameters for the turner module.

For the hereby described average-larva model no noise is introduced to any module input or output and no parameter variability is allowed across the population of the generated virtual larvae. The sole source of stochasticity then is due to the random distribution sampling processes operated during the behavioral simulations. For the analysis of the generated datasets the exact same pipeline is used as for the empirical dataset. The configuration of the final locomotory model is illustrated in **Table 1**.

Forward motion

For the Crawler module a set of 5 parameters need to be defined, along with a default body-length. This parameter-set is adequate to dynamically generate realistic forward velocity v oscillations as defined in **Equation 3**. All parameters are set to the median value across the empirical larva group. In **Figure 1*A*** the empirical distributions of the 5 crawler-module parameters are shown along with the median value selected for the average module.

Crawl intermittency

For the Intermittency module the sampling distributions for the run and pause epochs need to be defined. For both a continuous time-duration distribution is define while for the former an additional discrete distribution of number of concatenated strides per stridechain (run) is computed. In all cases, the pooled distribution of the given epoch across the entire dataset is approximated by power-law, exponential, Levy, log-normal and combined log-normal/power-law distributions. All distributions are truncated within an empirically observed range and are fully defined by their usual arguments. The best-fitting candidates from each class are sorted according to their Kolmogorov-Smirnov distance D_{KS} from the empirical distribution and the one exhibiting the minimum D_{KS} is selected for use in the model's intermittency module (**Table 1**). A validation is shown in **Figure 1*D*** where the distributions of the generated epochs (blue) match the empirical ones (red).

Angular motion

Realistically calibrating the locomotory model in the angular domain is more demanding. The Turner module only generates a torque-equivalent output while the torsional-spring body model applies both angular damping and restorative force due to the body-bend angle (**Equation 1**). We need to calibrate the 3 involved physics parameters, namely the angular damping z in sec^{-1}, the bend deformation resistance k in sec^{-2} and the torque coefficient c_T in sec^{-2}. Additionally the parameters of the turner module itself need to be specified. Here we study two implementations. The first is the above described neural oscillator, for which we need to define the baseline tonic input I_T^0, the spike-response steepness coefficient n_T and the time constant τ_T. The second is a simple sinusoidal oscillator, for which we need to define the baseline amplitude A_T^0 and frequency f_T. In order to have an empirical reference for isolated angular-only motion while avoiding the effect of crawling on the angular motion and we select the detected pause epochs exceeding 3 seconds and derive the pooled distributions for 3 angular metrics, namely the bending angle θ_b, the angular velocity ω and the angular acceleration $\dot{\omega}$ We run an optimization algorithm aiming to minimize the Kolmogorov-Smirnov distance between the simulated and empirical distributions for these 3 angular metrics. In each iteration a Turner module of a given configuration is simulated for 3 minutes while **Equation 1** is applied to derive θ_b and ω. The obtained optimal turner and physics parameters are shown in **Table 1**.

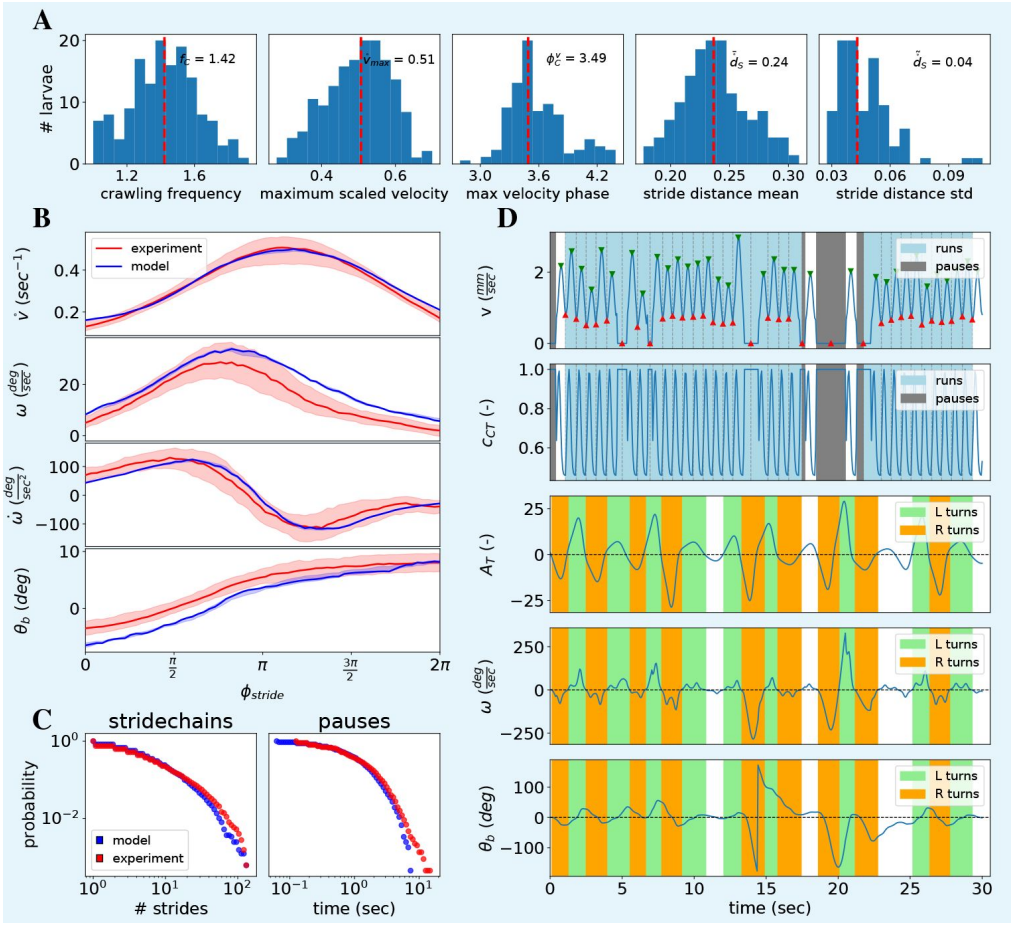

Appendix 2—figure 1.

Average locomotory model summary.

A: Distribution of stride-cycle related parameters over the empirical dataset. Red dotted line denotes the median value used in the Crawler module of the average locomotory model. **B:** Normalized average curves of angular metrics during the stride cycle for each individual larva. Red line denotes the group median. **C:** Pooled distribution of runs and pauses over the entire dataset (blue). Runs are detected as chains of concatenated strides (stridechains). Stridechains and pauses generated by the fitted distributions are shown in red. These distributions are used in the Intermitter module of the average locomotory model. **D:** Sample simulation of the model with all modules active. The model features additionally bend correction due to forward motion and crawling phase-coupled suppression of angular motion.

Crawl-bend interference

For the crawl-phase dependent attenuation of angular motion we study two implementations, each having three parameters that define the attenuation coefficient $c_{CT}(\Phi_C)$ during the stride-cycle (**Equation 6**). To achieve this we use a genetic algorithm (GA) optimization process to reach a best-fitting parameter-set for this module. Although optimization is only allowed to search the parameter space for these three parameters, each configuration is evaluated using a complete virtual larva model, combined with the optimal configurations of all other modules defined in the previous three calibration steps. During each iteration of the GA process a group of 30 model configurations is simulated for 3 minutes in free exploration conditions. The best 6 configurations are then mutated and/or combined in a novel generation. The evaluation function minimized is two-fold.Firstly the distributions of the three angular metrics derived from the simulated track, namely the bending angle θ_b, the angular velocity ω and the angular acceleration $\dot{\omega}$ are compared to the empirically measured target distributions via the Kolmogorov-Smirnov distance as described above for the angular motion. Secondly, stride-cycle analysis is carried out for each simulated larva. Detected strides are interpolated in over $n = 64$ bins in a 2π cycle and direction-normalized by inverting the right-turning strides. The average stride-cycle curve $\hat{x}$ is then computed for each of the three metrics and compared to the average empirically measured curve x via the evaluation metric defined below. The metric has been selected so that the different scales of angles, angular velocities and accelerations are normalized before being summed. The final evaluation metric to be minimized is defined as the cumulative error across the three angular metrics for both the distributions and the average stride-cycle curves.

$$Error(x, \hat{x}) = \sqrt{\frac{1}{n}\sum_{i=1}^{n}\frac{(x_i - \hat{x}_i)^2}{|\max(x) - \min(x)|^2}} \tag{10}$$

Here we study two configurations for the crawl phase-dependent suppression of the angular motion, namely a smooth gaussian relief curve and an acute square transition as described in **Equation 6**.The target empirical and optimal model's average curve for the angular velocity ω are shown in **Figure 1*B***.

A sample simulation of the complete model is shown in **Figure 1*D***. The Crawler-generated forward velocity oscillations (1st row) attenuate the angular motion in a phase-locked manner (2nd row) resulting in low angular velocity during runs and acute headcasts during pauses (3rd row). The bending angle is additionally restored to 0 during forward motion (4th row).

References

Almeida-Carvalho MJ, Berh D, Braun A, Chen YC, Eichler K, Eschbach C, Fritsch PMJ, Gerber B, Hoyer N, Jiang X, Kleber J, Klämbt C, König C, Louis M, Michels B, Miroschnikow A, Mirth C, Miura D, Niewalda T, Otto N, et al. (2017) **The Ol1mpiad: Concordance of behavioural faculties of stage 1 and stage 3 Drosophila larvae** *J Exp Biol* **220**:2452–2475 https://doi.org/10.1242/jeb.156646

de Belle JS, Heisenberg M. (1994) **Associative Odor Learning in Drosophila Abolished by Chemical Ablation of Mushroom Bodies** *Science* **263**:692–695 https://doi.org/10.1126/science.8303280

Bennett JE, Philippides A, Nowotny T. (2021) **Learning with reinforcement prediction errors in a model of the Drosophila mushroom body** *Nature communications* **12**:2569 https://doi.org/10.1038/s41467-021-22592-4

Berni J. (2015) **Genetic dissection of a regionally differentiated network for exploratory behavior in drosophila larvae** *Curr Biol* **25**:1319–1326 https://doi.org/10.1016/j.cub.2015.03.023

Berni J, Pulver SR, Griffith LC, Bate M. (2012) **Autonomous circuitry for substrate exploration in freely moving drosophila larvae** *Curr Biol* **22**:1861–1870 https://doi.org/10.1016/j.cub.2012.07.048

Bicho E. (1999) **Dynamic approach to behavior-based robotics : design, specification, analysis, simulation and implementation** University of Minho

Bidaye SS, Laturney M, Chang AK, Liu Y, Bockemühl T, Büschges A, Scott K. (2020) **Two brain pathways initiate distinct forward walking programs in Drosophila** *Neuron* **108**:469–485 https://doi.org/10.1016/j.neuron.2020.07.032

Brooks R. (1986) **A Robust Layered Control System For A Mobile Robot** *IEEE J Robot Autom* **RA-2**:14–23

Davies A, Louis M, Webb B. (2015) **A Model of Drosophila Larva Chemotaxis** *PLoS Comput Biol* **11**:1–24 https://doi.org/10.1371/journal.pcbi.1004606

Denisov G, Ohyama T, Jovanic T, Zlatic M. (2013) **Model-based detection and analysis of animal behaviors using signals extracted by automated tracking** In: *BIOSIGNALS 2013 - Proc Int Conf Bio-Inspired Syst Signal Process* pp. 175–181 https://doi.org/10.5220/0004235101750181

Diegelmann S, Klagges B, Michels B, Schleyer M, Gerber B. (2013) **Maggot learning and synapsin function** *J Exp Biol* **216**:939–951 https://doi.org/10.1242/jeb.076208

Dolan MJ, Belliart-Guérin G, Bates AS, Frechter S, Lampin-Saint-Amaux A, Aso Y, Roberts RJV, Schlegel P, Wong A, Hammad A, Bock D, Rubin GM, Preat T, Plaçais PY, Jefferis GSXE. (2018) **Communication from Learned to Innate Olfactory Processing Centers Is Required for Memory Retrieval in Drosophila** *Neuron* **100**:651–668 https://doi.org/10.1016/j.neuron.2018.08.037

Eschbach C, Fushiki A, Winding M, Schneider-Mizell CM, Shao M, Arruda R, Eichler K, Valdes-Aleman J, Ohyama T, Thum AS, Gerber B, Fetter RD, Truman JW, Litwin-Kumar A, Cardona A, Zlatic M. (2020) **Recurrent architecture for adaptive regulation of learning in the insect brain** *Nat Neurosci* **23**:544–555 https://doi.org/10.1038/s41593-020-0607-9

Eschbach C, Zlatic M. (2020) **Useful road maps: studying Drosophila larva's central nervous system with the help of connectomics** *Curr Opin Neurobiol* **65**:129–137 https://doi.org/10.1016/j.conb.2020.09.008

Felsenberg J, Waddell S. (2019) **Switching Gears, Structuring the Right Search Strategy** *Neuron* **102**:273–275 https://doi.org/10.1016/j.neuron.2019.03.040

Feng K, Sen R, Minegishi R, Dübbert M, Bockemühl T, Büschges A, Dickson BJ (2020) **Distributed control of motor circuits for backward walking in Drosophila** *Nature Communications* **11**:6166 https://doi.org/10.1038/s41467-020-19936-x

Gerber B, Stocker RF (2007) **The drosophila larva as a model for studying chemosensation and chemosensory learning: A review** *Chem Senses* **32**:65–89 https://doi.org/10.1093/chemse/bjl030

Gerber B, Yarali A, Diegelmann S, Wotjak CT, Pauli P, Fendt M. (2014) **Pain-relief learning in flies, rats, and man: Basic research and applied perspectives** *Learn Mem* **21**:232–252 https://doi.org/10.1101/lm.032995.113

Gerhard S, Andrade I, Fetter RD, Cardona A, Schneider-Mizell CM (2017) **Conserved neural circuit structure across drosophila larval development revealed by comparative connectomics** *eLife* **6**:1–17 https://doi.org/10.7554/eLife.29089

Gjorgjieva J, Berni J, Evers JF, Eglen SJ (2013) **Neural circuits for peristaltic wave propagation in crawling drosophila larvae: Analysis and modeling** *Front Comput Neurosci* **7**:1–19 https://doi.org/10.3389/fncom.2013.00024

Gomez-Marin A, Louis M. (2012) **Active sensation during orientation behavior in the Drosophila larva: More sense than luck** *Curr Opin Neurobiol* **22**:208–215 https://doi.org/10.1016/j.conb.2011.11.008

Gomez-Marin A, Stephens GJ, Louis M. (2011) **Active sampling and decision making in Drosophila chemotaxis** *Nat Commun* **2**:410–441 https://doi.org/10.1038/ncomms1455

Günther MN, Nettesheim G, Shubeita GT (2016) **Quantifying and predicting Drosophila larvae crawling phenotypes** *Sci Rep* **6**:1–10 https://doi.org/10.1038/srep27972

Heckscher ES, Lockery SR, Doe CQ (2012) **Characterization of Drosophila Larval Crawling at the Level of Organism, Segment, and Somatic Body Wall Musculature** *J Neurosci* **32**:12460–12471 https://doi.org/10.1523/jneurosci.0222-12.2012

Helgadottir LI, Haenicke J, Landgraf T, Rojas R, Nawrot MP (2013) **Conditioned behavior in a robot controlled by a spiking neural network** In: *Int. IEEE/EMBS Conf. Neural Eng. NER* https://doi.org/10.1109/NER.2013.6696078

Imambocus BN, Zhou F, Formozov A, Wittich A, Tenedini FM, Hu C, Sauter K, Macarenhas Varela E, Herédia F, Casimiro AP, Macedo A, Schlegel P, Yang CH, Miguel-Aliaga I, Wiegert JS, Pankratz MJ, Gontijo AM, Cardona A, Soba P. (2022) **A neuropeptidergic circuit gates selective escape**

behavior of Drosophila larvae *Current Biology* **32**:149–163 https://doi.org/10.1016/j.cub.2021.10.069

Jovanic T. (2020) **Studying neural circuits of decision-making in Drosophila larva** *J Neurogenet* **34**:162–170 https://doi.org/10.1080/01677063.2020.1719407

Jürgensen AM, Khalili A, Chicca E, Indiveri G, Nawrot MP (2021) **A neuromorphic model of olfactory processing and sparse coding in the Drosophila larva brain** *Neuromorphic Comput Eng* https://doi.org/10.1088/2634-4386/ac3ba6

Jürgensen AM, Sakagiannis P, Schleyer M, Gerber B, Nawrot MP (2024) **Prediction error drives associative learning and conditioned behavior in a spiking model of Drosophila larva** *iScience* **27**:108640 https://doi.org/10.1016/j.isci.2023.108640

Kafle T, Grub M, Sakagiannis P, Nawrot MP, Arguello R. (2024) **Fast and recurrent evolution of temperature preference among drosophilids** *bioRxiv* https://doi.org/10.1101/2024.03.15.585210

Karagyozov D, Mihovilovic Skanata M, Lesar A, Gershow M. (2018) **Recording Neural Activity in Unrestrained Animals with Three-Dimensional Tracking Two-Photon Microscopy** *Cell Reports* **25**:1371–1383 https://doi.org/10.1016/j.celrep.2018.10.013

Kim D, Alvarez M, Lechuga LM, Louis M. (2017) **Species-specific modulation of food-search behavior by respiration and chemosensation in Drosophila larvae** *eLife* **6**:1–23 https://doi.org/10.7554/eLife.27057

Klein M, Krivov SV, Ferrer AJ, Luo L, Samuel AD, Karplus M. (2017) **Exploratory search during directed navigation in C. elegans and Drosophila larva** *Elife* **6**:1–14 https://doi.org/10.7554/elife.30503

Kohsaka H, Zwart MF, Fushiki A, Fetter RD, Truman JW, Cardona A, Nose A. (2019) **Regulation of forward and backward locomotion through intersegmental feedback circuits in Drosophila larvae** *Nat Commun* **10**:2654 https://doi.org/10.1038/s41467-019-10695-y

König C, Khalili A, Ganesan M, Nishu AP, Garza AP, Niewalda T, Gerber B, Aso Y, Yarali A. (2018) **Reinforcement signaling of punishment versus relief in fruit flies** *Learn Mem* **25**:247–257 https://doi.org/10.1101/lm.047308.118

Lahiri S, Shen K, Klein M, Tang A, Kane E, Gershow M, Garrity P, Samuel ADT (2011) **Two alternating motor programs drive navigation in Drosophila larva** *PLoS One* **6** https://doi.org/10.1371/journal.pone.0023180

Laurent F, Blanc A, May L, Gándara L, Cocanougher BT, Jones BMW, Hague P, Barré C, Vestergaard CL, Crocker J, Zlatic M, Jovanic T, Masson JB (2024) **LarvaTagger: manual and automatic tagging of Drosophila larval behaviour** *Bioinformatics* **40**:btae441 https://doi.org/10.1093/bioinformatics/btae441

Lesar A, Tahir J, Wolk J, Gershow M. (2021) **Switch-like and persistent memory formation in individual Drosophila larvae** *eLife* **10**:e70317

Loveless J, Lagogiannis K, Webb B. (2019) **Modelling the neuromechanics of exploration in larval Drosophila** *PLoS Comput Biol* **15**:7 https://doi.org/10.5281/zenodo.1432637

Loveless J, Webb B. (2018) **A Neuromechanical Model of Larval Chemotaxis** *Integr Comp Biol* **58**:906–914 https://doi.org/10.1093/icb/icy094

Malloy CA, Somasundaram E, Omar A, Bhutto U, Medley M, Dzubuk N, Cooper RL (2019) **Pharmacological identification of cholinergic receptor subtypes: modulation of locomotion and neural circuit excitability in Drosophila larvae** *Neuroscience* **411**:47–64 https://doi.org/10.1016/j.neuroscience.2019.05.016

Mantziaris C, Bockemühl T, Büschges A. (2020) **Central Pattern Generating Networks in Insect Locomotion** *Dev Neurobiol* **00**:1–15 https://doi.org/10.1002/dneu.22738

Marken RS, Shaffer DM (2017) **The power law of movement : an example of a behavioral illusion** *Exp Brain Res* **235**:1835–1842 https://doi.org/10.1007/s00221-017-4939-y

Marken RS, Shaffer DM (2018) **The power law as behavioral illusion: reappraising the reappraisals** *Exp Brain Res* **236**:1537–1544 https://doi.org/10.1007/s00221-018-5208-4

Michels B, Chen YC, Saumweber T, Mishra D, Tanimoto H, Schmid B, Engmann O, Gerber B. (2011) **Cellular site and molecular mode of synapsin action in associative learning** *Learn Mem* **18**:332–344 https://doi.org/10.1101/lm.2101411

Michels B, Saumweber T, Biernacki R, Thum J, Glasgow RDV, Schleyer M, Chen YC, Eschbach C, Stocker RF, Toshima N, Tanimura T, Louis M, Arias-Gil G, Marescotti M, Benfenati F, Gerber B. (2017) **Pavlovian conditioning of larval Drosophila: An illustrated, multilingual, hands-on manual for odor-taste associative learning in maggots** *Front Behav Neurosci* **11**:1–6 https://doi.org/10.3389/fnbeh.2017.00045

Miroschnikow A, Schlegel P, Schoofs A, Hueckesfeld S, Li F, Schneider-Mizell CM, Fetter RD, Truman JW, Cardona A, Pankratz MJ (2018) **Convergence of monosynaptic and polysynaptic sensory paths onto common motor outputs in a Drosophila feeding connectome** *eLife* **7**:e40247 https://doi.org/10.7554/eLife.40247

Miroschnikow A, Schlegel P, Schoofs A, Hueckesfeld S, Li F, Schneider-Mizell CM, Fetter RD, Truman JW, Cardona A, Pankratz MJ (2018) **The Feeding Connectome: Convergence of Monosynaptic and Polysynaptic Sensory Paths onto Common Motor Outputs** *eLife* **7**:1–23 https://doi.org/10.7554/elife.40247

Niewalda T, Jeske I, Michels B, Gerber B. (2014) **Peer pressure' in larval Drosophila** *Biol Open* **3**:575–582 https://doi.org/10.1242/bio.20148458

Ohyama T, Jovanic T, Denisov G, Dang TC, Hoffmann D, Kerr RA, Zlatic M. (2013) **High-Throughput Analysis of Stimulus-Evoked Behaviors in Drosophila Larva Reveals Multiple Modality-Specific Escape Strategies** *PLoS One* **8** https://doi.org/10.1371/journal.pone.0071706

Owald D, Felsenberg J, Talbot CB, Das G, Perisse E, Huetteroth W, Waddell S. (2015) **Activity of defined mushroom body output neurons underlies learned olfactory behavior in Drosophila** *Neuron* **86**:417–427 https://doi.org/10.1016/j.neuron.2015.03.025

Owald D, Waddell S. (2015) **Olfactory learning skews mushroom body output pathways to steer behavioral choice in Drosophila** *Curr Opin Neurobiol* **35**:178–184 https://doi.org/10.1016/j.conb.2015.10.002

Paisios E, Rjosk A, Pamir E, Schleyer M. (2017) **Common microbehavioral "footprint" of two distinct classes of conditioned aversion** *Learn Mem* **24**:191–198 https://doi.org/10.1101/lm.045062.117

Pehlevan C, Paoletti P, Mahadevan L. (2016) **Integrative neuromechanics of crawling in D. melanogaster larvae** *Elife* **5**:1–23 https://doi.org/10.7554/elife.11031

Prescott TJ, Redgrave P, Gurney K. (1999) **Layered control architectures in robots and vertebrates** *Adapt Behav* **7**:99–127 https://doi.org/10.1177/105971239900700105

Prescott TJ, Wilson SP (2023) **Understanding brain functional architecture through robotics** *Science Robotics* American Association for the Advancement of Science (AAAS) **8** https://doi.org/10.1126/scirobotics.adg6014

Qian CS, Kaplow M, Lee JK, Grueber WB (2018) **Diversity of internal sensory neuron axon projection patterns is controlled by the POU-domain protein pdm3 in drosophila larvae** *J Neurosci* **38**:2081–2093 https://doi.org/10.1523/JNEUROSCI.2125-17.2018

Rescorla RA (1972) **Informational Variables in Pavlovian Conditioning** *Psychol Learn Motiv* **6**:1–64

Reynolds AM, Jones HBC, Hill JK, Pearson AJ, Wilson K, Wolf S, Lim KS, Reynolds DR, Chapman JW, Reynolds AM (2015) **Evidence for a pervasive 'idling-mode' activity template in flying and pedestrian insects** *R Soc Open Sci* **2** https://doi.org/10.1098/rsos.150085

Riemensperger T, Völler T, Stock P, Buchner E, Fiala A. (2005) **Punishment prediction by dopaminergic neurons in Drosophila** *Curr Biol* **15**:1953–1960 https://doi.org/10.1016/j.cub.2005.09.042

Risse B, Berh D, Otto N, Klämbt C, Jiang X. (2017) **FIMTrack: An open source tracking and locomotion analysis software for small animals** *PLoS Comput Biol* **13**:1–15 https://doi.org/10.1371/journal.pcbi.1005530

Ross D, Lagogiannis K, Webb B. (2015) **A Model of Larval Biomechanics Reveals Exploitable Passive Properties for Efficient Locomotion** In: Wilson SP, Verschure PFMJ, Mura A, Prescott TJ, editors. *Biomimetic and Biohybrid Systems* Cham: Sprimger International Publishing pp. 1–12

Ruiz-Dubreuil G, Burnet B, Connolly K, Furness P. (1996) **Larval foraging behaviour and competition in Drosophila melanogaster** *Heredity (Edinb)* **76**:55–64 https://doi.org/10.1038/hdy.1996.7

Sakagiannis P, Aguilera M, Nawrot MP (2020) **A Plausible Mechanism for Drosophila Larva Intermittent Behavior** In: Vouloutsi V, Mura A, Tauber F, Speck T, Prescott TJ, Verschure PFMJ, editors. *Biomimetic and Biohybrid Systems* Cham: Sprimger International Publishing pp. 288–299

Saumweber T, Rohwedder A, Schleyer M, Eichler K, Chen Yc, Aso Y, Cardona A, Eschbach C, Kobler O, Voigt A, Durairaja A, Mancini N, Zlatic M, Truman JW, Thum AS, Gerber B. (2018) **Functional architecture of reward learning in mushroom body extrinsic neurons of larval Drosophila** *Nat Commun* :1–19 https://doi.org/10.1038/s41467-018-03130-1

Saumweber T, Weyhersmüller A, Hallermann S, Diegelmann S, Michels B, Bucher D, Funk N, Reisch D, Krohne G, Wegener S, Buchner E, Gerber B. (2011) **Behavioral and synaptic plasticity are impaired upon lack of the synaptic protein SAP47** *J Neurosci* **31**:3508–3518 https://doi.org/10.1523/JNEUROSCI.2646-10.2011

Schleyer M, Fendt M, Schuller S, Gerber B. (2018) **Associative learning of stimuli paired and unpaired with reinforcement: Evaluating evidence from maggots, flies, bees, and rats** *Front Psychol* **9**:1–15 https://doi.org/10.3389/fpsyg.2018.01494

Schleyer M, Miura D, Tanimura T, Gerber B. (2015) **Learning the specific quality of taste reinforcement in larval Drosophila** *eLife* **4** https://doi.org/10.7554/eLife.04711

Schleyer M, Reid SF, Pamir E, Saumweber T, Paisios E, Davies A, Gerber B, Louis M. (2015) **The impact of odor-reward memory on chemotaxis in larval Drosophila** *Learn Mem* **22**:267–277 https://doi.org/10.1101/lm.037978.114

Schleyer M, Saumweber T, Nahrendorf W, Fischer B, von Alpen D, Pauls D, Thum A, Gerber B. (2011) **A behavior-based circuit model of how outcome expectations organize learned behavior in larval Drosophila** *Learn Mem* **18**:639–653 https://doi.org/10.1101/lm.2163411

Schleyer M, Weiglein A, Thoener J, Strauch M, Hartenstein V, Weigelt MK, Schuller S, Saumweber T, Eichler K, Rohwedder A, Merhof D, Zlatic M, Thum AS, Gerber B. (2020) **Identification of dopaminergic neurons that can both establish associative memory and acutely terminate its behavioral expression** *J Neurosci* **40**:5990–6006 https://doi.org/10.1523/JNEUROSCI.0290-20.2020

Schoofs A, Hückesfeld S, Schlegel P, Miroschnikow A, Peters M, Zeymer M, Spieß R, Chiang AS, Pankratz MJ (2014) **Selection of Motor Programs for Suppressing Food Intake and Inducing Locomotion in the Drosophila Brain** *PLoS Biol* **12** https://doi.org/10.1371/journal.pbio.1001893

Schoofs A, Miroschnikow A, Schlegel P, Zinke I, Schneider-Mizell CM, Cardona A, Pankratz MJ (2024) **Serotonergic modulation of swallowing in a complete fly vagus nerve connectome** *Current Biology* https://doi.org/10.1016/j.cub.2024.08.025

Schultz W. (2015) **Neuronal reward and decision signals: From theories to data** *Physiol Rev* **95**:853–951 https://doi.org/10.1152/physrev.00023.2014

Schultz W. (2016) **Dopamine reward prediction error coding** *Dialogues Clin Neurosci* **18**:23

Schulze A, Gomez-Marin A, Rajendran VG, Lott G, Musy M, Ahammad P, Deogade A, Sharpe J, Riedl J, Jarriault D, Trautman ET, Werner C, Venkadesan M, Druckmann S, Jayaraman V, Louis M. (2015) **Dynamical feature extraction at the sensory periphery guides chemotaxis** *eLife* **4**:1–52 https://doi.org/10.7554/eLife.06694

Schumann I, Triphan T. (2020) **The PEDtracker: An Automatic Staging Approach for Drosophila melanogaster Larvae** *Front Behav Neurosci* **14** https://doi.org/10.3389/fnbeh.2020.612313

Sen R, Wu M, Branson K, Robie A, Rubin GM, Dickson BJ (2017) **Moonwalker Descending Neurons Mediate Visually Evoked Retreat in Drosophila** *Current Biology* Elsevier Ltd **27**:766–771 https://doi.org/10.1016/j.cub.2017.02.008

Sims DW, Humphries NE, Hu N, Medan V, Berni J. (2019) **Optimal searching behaviour generated intrinsically by the central pattern generator for locomotion** *eLife* **8**:1–31 https://doi.org/10.7554/eLife.50316

Slater G, Levy P, Andrew Chan KL, Larsen C. (2015) **A central neural pathway controlling odor tracking in drosophila** *J Neurosci* **35**:1831–1848 https://doi.org/10.1523/JNEUROSCI.2331-14

.2015

Springer MA, Nawrot MP (2021) **A mechanistic model for reward prediction and extinction learning in the fruit fly** *eNeuro - Press* https://doi.org/10.1523/ENEURO.0549-20.2021

Strutz A, Soelter J, Baschwitz A, Farhan A, Grabe V, Rybak J, Knaden M, Schmuker M, Hansson BS, Sachse S. (2014) **Decoding odor quality and intensity in the Drosophila brain** *eLife* **3**:1–21 https://doi.org/10.7554/elife.04147

Szigeti B, Deogade A, Webb B. (2015) **Searching for motifs in the behaviour of larval Drosophila melanogaster and Caenorhabditis elegans reveals continuity between behavioural states** *J R Soc Interface* **12** https://doi.org/10.1098/rsif.2015.0899

Tadres D, Louis M. (2020) **PiVR: An affordable and versatile closed-loop platform to study unrestrained sensorimotor behavior** *PLoS Biol* **18**:1–25 https://doi.org/10.1371/jour-nal.pbio.3000712

Tastekin I, Khandelwal A, Tadres D, Fessner ND, Truman JW, Zlatic M, Cardona A, Louis M. (2018) **Sensorimotor pathway controlling stopping behavior during chemotaxis in the Drosophila melanogaster larva** *eLife* **7**:1–38 https://doi.org/10.7554/elife.38740

Thane M, Paisios E, Stöter T, Krüger AR, Gläß S, Dahse AK, Scholz N, Gerber B, Lehmann DJ, Schleyer M. (2023) **High-resolution analysis of individual Drosophila melanogaster larvae uncovers individual variability in locomotion and its neurogenetic modulation** *Open Biology* **13**:220308 https://doi.org/10.1098/rsob.220308

Thane M, Viswanathan V, Meyer TC, Paisios E, Schleyer M. (2019) **Modulations of microbehaviour by associative memory strength in Drosophila larvae** *PLoS One* **14**:e0224154 https://doi.org/10.1371/journal.pone.0224154

de Tredern E, Manceau D, Blanc A, Sakagiannis P, Barre C, Sus V, Viscido F, Hasan MA, Autran S, Nawrot M, Mason JB, Jovanic T. (2024) **Feeding-state dependent neuropeptidergic modulation of reciprocally interconnected inhibitory neurons biases sensorimotor decisions in Drosophila** *bioRxiv* https://doi.org/10.1101/2023.12.26.573306

Ueno T, Masuda N, Kume S, Kume K. (2012) **Dopamine Modulates the Rest Period Length without Perturbation of Its Power Law Distribution in Drosophila melanogaster** *PLoS One* **7** https://doi.org/10.1371/journal.pone.0032007

Vogt K, Zimmerman DM, Schlichting M, Hernandez-Nunez L, Qin S, Malacon K, Rosbash M, Pehlevan C, Cardona A, Samuel ADT (2021) **Internal state configures olfactory behavior and early sensory processing in Drosophila larvae** *Sci Adv* **7**:eabd6900 https://doi.org/10.1126/sciadv.abd6900

Weiglein A, Gerstner F, Mancini N, Schleyer M, Gerber B. (2019) **One-trial learning in larval Drosophila** *Learn Mem* **26**:109–120 https://doi.org/10.1101/lm.049106.118

Widmann A, Artinger M, Biesinger L, Boepple K, Peters C, Schlechter J, Selcho M, Thum AS (2016) **Genetic Dissection of Aversive Associative Olfactory Learning and Memory in Drosophila Larvae** *PLoS Genet* **12**:1–32 https://doi.org/10.1371/journal.pgen.1006378

Widmann A, Eichler K, Selcho M, Thum AS, Pauls D. (2018) **Odor-taste learning in Drosophila larvae** *J Insect Physiol* **106**:47–54 https://doi.org/10.1016/j.jinsphys.2017.08.004

Wilson SP, Prescott TJ (2022) **Scaffolding layered control architectures through constraint closure: insights into brain evolution and development** *Philosophical Transactions of the Royal Society B: Biological Sciences* The Royal Society **377** https://doi.org/10.1098/rstb.2020.0519

Wu Q, Wen T, Lee G, Park JH, Cai HN, Shen P. (2003) **Developmental control of foraging and social behavior by the Drosophila neuropeptide Y-like system** *Neuron* **39**:147–161 https://doi.org/10.1016/S0896-6273(03)00396-9

Wystrach A, Lagogiannis K, Webb B. (2016) **Continuous lateral oscillations as a core mechanism for taxis in Drosophila larvae** *eLife* **5** https://doi.org/10.7554/elife.15504

Zago M, Lacquaniti F, Gomez-marin A, Gomez-marin A. (2016) **The speed - curvature power law in Drosophila larval locomotion** *Biol Lett* **12** https://doi.org/10.6084/m9.fig-share.c.3517350

Zago M, Matic A, Flash T, Gomez-Marin A, Lacquaniti F. (2017) **The speed - curvature power law of movements: a reappraisal** *Exp Brain Res* **236**:69–82 https://doi.org/10.1007/s00221-017-5108-z

Zarin AA, Mark B, Cardona A, Litwin-Kumar A, Doe CQ (2019) **A multilayer circuit architecture for the generation of distinct locomotor behaviors in Drosophila** *eLife* **8**:1–34 https://doi.org/10.7554/eLife.51781

Zhang T, Branch A, Shen P. (2013) **Octopamine-mediated circuit mechanism underlying controlled appetite for palatable food in Drosophila** *Proc Natl Acad Sci* **110**:15431–15436 https://doi.org/10.1073/pnas.1308816110

Author information

Panagiotis Sakagiannis[†]
Computational Systems Neuroscience, University of Cologne, Cologne, Germany
ORCID iD: 0000-0002-1033-5387

For correspondence: p.sakagiannis@uni-koeln.de

[†]Present address: Computational Systems Neuroscience, Institute of Zoology, University of Cologne, Cologne, Germany

Anna-Maria Jürgensen[†]
Computational Systems Neuroscience, University of Cologne, Cologne, Germany
ORCID iD: 0000-0002-7871-1887

[†]Present address: Computational Systems Neuroscience, Institute of Zoology, University of Cologne, Cologne, Germany

Martin Paul Nawrot[†]
Computational Systems Neuroscience, University of Cologne, Cologne, Germany
ORCID iD: 0000-0003-4133-6419

For correspondence: mnawrot@uni-koeln.de
[†]Present address: Computational Systems Neuroscience, Institute of Zoology, University of Cologne, Cologne, Germany

5 *Larvaworld*: A behavioral simulation and analysis platform for *Drosophila* larva

Authors: Panagiotis Sakagiannis, Hannes Rapp, Tihana Jovanic, Martin Paul Nawrot
Archive: bioRxiv
Date: June 18, 2025
DOI: 10.1101/2025.06.15.659765
Links Python package, Documentation, Source Code

Author Contributions

Contributions to the publication were as follows: Conceptualization was carried out by P.S. and M.P.N.; data curation was performed by T.J.; methodology, visualization, and original draft writing were performed by P.S.; software development was conducted by P.S. and H.R.; funding acquisition and supervision were provided by M.P.N.; and all authors (P.S., H.R., T.J., M.P.N.) contributed to review and editing.

Larvaworld : A behavioral simulation and analysis platform for *Drosophila* larva

Panagiotis Sakagiannis[1*], Hannes Rapp[1], Tihana Jovanic[2], Martin Paul Nawrot[1]

1 Computational Systems Neuroscience, Institute of Zoology, 50674, University of Cologne, Germany
2 Université Paris-Saclay, CNRS, Institut des neurosciences Paris-Saclay, 91400, Saclay, France

* p.sakagiannis@uni-koeln.de

Abstract

Behavioral modeling supports theory building and evaluation across disciplines. Leveraging advances in motion-tracking and computational tools, we present a virtual laboratory for Drosophila larvae that integrates agent-based modeling with multiscale neural control and supports analysis of both simulated and experimental data. Virtual larvae are implemented as 2D agents capable of realistic locomotion, guided by multimodal sensory input and constrained by a dynamic energy-budget model that balances exploration and exploitation. Each agent is organized as a hierarchical, behavior-based control system comprising three layers: low-level locomotion, optionally incorporating neuromechanical models; mid-level sensory processing; and high-level behavioral adaptation. Neural control models can range from simple linear transfer models to rate-based or spiking neural network models, e.g. to accomodate associative learning. Simulations operate across sub-millisecond neuronal dynamics, sub-second closed-loop behavior, and circadian-scale metabolic regulation. Users can configure both larval models and virtual environments, including sensory landscapes, nutrient sources, and physical arenas. Real-time visualization is integrated into the simulation and analysis pipeline, which also allows for standardized processing of motion-tracking data from real experiments. Distributed as an open-source Python package, the platform includes tutorial experiments to support accessibility, customization, and use in both research and education.

Author summary

Larvaworld was developed to address two key challenges in behavioral neuroscience and computational modeling. First, it responds to the growing call for closer collaboration between experimentalists and modelers by providing a shared platform -a virtual laboratory- where experimental data analysis and behavioral modeling can be seamlessly

integrated. By standardizing dataset formats and ensuring identical, unbiased analysis pipelines for experimental and simulated data, *Larvaworld* facilitates methodological consistency and enables rigorous model evaluation.

Second, it aims to bridge a long-standing gap in theory building and computational modeling at the level of the individual behaving organism. Historically, neuroscience has focused on sub-individual processes, while ecology has concentrated on supra-individual dynamics, resulting in discontinuities among the respective modeling approaches. Recent advances, however, have begun to align these fields, with neuroscience incorporating slower homeostatic processes and ecology integrating faster neurally-mediated mechanisms. *Larvaworld* boosts this convergence by adopting a nested, multi-timescale modeling approach, thus achieving behavioral regulation within the normative homeostatic constraints as these dynamically unfold during larval development. By combining established modeling paradigms from neuroscience and ecology, it provides a novel and flexible platform for studying behavior at the level of the individual organism, promoting cross-disciplinary insights and advancing computational neuroethology [1].

Introduction

Drosophila is a widely studied model organism in neuroscience, alongside lamprey, zebrafish, mouse, rat, and monkey, listed in order of increasing nervous system complexity. Researchers working with mammals are acutely aware of the scarcity and high cost of experimental animals, making computational modeling a valuable alternative to real-life experiments. Insects, on the contrary, are more abundant, affordable, and easier to rear, and their use is subject to fewer ethical restrictions. Nevertheless, growing concerns about animal welfare have recently fueled support for alternative research and educational tools aimed at reducing reliance on live animal experiments. Behavioral studies, in particular, often require far larger sample sizes than neurophysiological experiments, further highlighting the need for computational approaches. In this context, virtual laboratories and simulation platforms are emerging as indispensable tools for both scientific research and education [1, 2]. Many academic institutions have already developed and integrated such tools into their curricula, with notable applications in fields like Mendelian genetics and beyond.

Not all scientific fields readily translate to the virtual domain, and behavioral science is particularly challenging in this regard. The widespread use of arbitrary behavioral metrics across different laboratories further complicates experiment replication. Despite significant advances in experimental setups, genetic manipulations, and behavioral recordings, tools for behavioral modeling remain limited, often lacking key functionalities. In particular, tools that support direct comparisons between simulated and experimental datasets are still underdeveloped.

A recent perspective paper, authored by pioneering researchers in AI introduced the concept of the *embodied Turing test* as a potential goal for computational modeling of animal behavior [3]. This approach seeks to bridge neuroscience and AI by generating robotic or virtual animats, whose embodied behavior is indistinguishable from that of real animals. Achieving this goal requires a platform that facilitates continuous and unbiased comparisons between experimental and simulated datasets to more effectively construct, calibrate, and evaluate models.

Here we propose one such tool for *Drosophila* larvae. The core concept is to create and simulate virtual larvae whose behavior can be analyzed using the exact same pipelines applied to real animals. From the platform's perspective, experimental and simulated

datasets are indistinguishable. Experimental data can be imported and automatically converted into a standardized format identical to that produced by the simulations.

To enhance replicability and accommodate variability in behavioral metrics, the platform relies solely on the originally tracked x-y coordinate time series. All additional behavioral metrics are transparently defined and consistently derived from these coordinates for both experimental and simulated datasets. Virtual larvae are modeled with a simplified 2D body structure that mirrors key features of the real animal. Their movements are tracked during simulations as if observed by a virtual tracker, navigating a detailed virtual replica of the real experimental arena.

Once a trackable body is established, behavioral control can be implemented through computational modeling. To maximize flexibility and support integration and expansion, *Larvaworld* adopts a hybrid, modular, and hierarchical control architecture [4]. Its hierarchical design follows the layered behavioral control paradigm [5], in which successive control layers are stacked according to two key principles: decentralization and subsumption. Decentralization ensures that each layer can operate autonomously, independent of top-down input-consistent with evidence of decentralized neural circuits capable of generating simple behavioral primitives even when isolated from higher-order control [6]. Subsumption, on the other hand, allows top-down modulation to influence only a few key parameters, reflecting subtle adjustments by higher neural centers.

Each control layer is composed of interconnected modules, specialized for processing specific sensory and modulatory inputs, motor outputs, or sensorimotor integration. This toolkit-like, modular design has been already adopted in some studies of larva behavioral neuroscience because it offers a high degree of configurability, enabling researchers to compare models by adding, removing, or replacing modules [7]. It also facilitates expansion through the seamless integration of new modules.

The hybrid nature of the framework imposes minimal constraints on the modeling detail within individual modules. Whether deterministic, stochastic, rule-based, rate-coded neural models, or neuron-level spiking models, modules can be combined, replaced, and compared, provided they conform to standardized input-output formats. Once a modular larva model has been assembled, it can undergo a genetic algorithm optimization process prior to its use in evaluation studies.

The following sections elaborate on the implementation, usage, and accessibility of the platform in greater detail. Section Design and Implementation describes the architecture and implementation of the Larvaworld platform, including its modeling principles, simulation modes, configuration options, and analysis tools. Section Results presents a range of scientific applications, demonstrating how Larvaworld has been used for behavioral modeling, data-driven analysis, and model evaluation in various settings. Section Code availability provides information on code availability and installation. A schematic of the main components is shown in Figure 1 and a summary of preconfigured virtual experiments is included in Table 2, offering quick-start examples for further exploration.

Design and Implementation

Modeling principles

The design of the Larvaworld platform addresses four aims:

- Integration of established theoretical and modeling principles from distinct fields across behavioral sciences

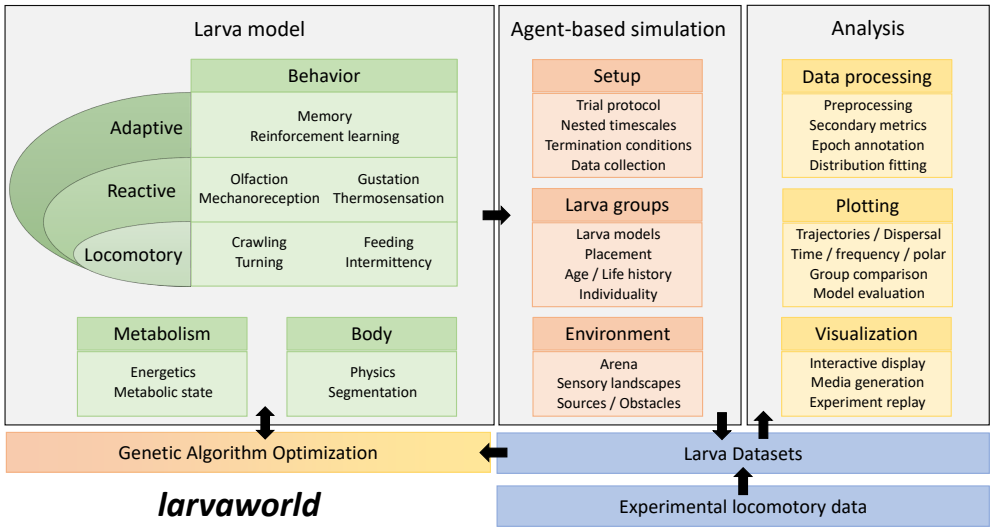

Fig 1. *Larvaworld* architecture. A schematic of the main components and functionalities of *Larvaworld*.

- User-friendly interface for both behavioral modeling and analysis purposes
- Modularity and extendability
- Computational efficiency and storage management

In the following sections the motivation behind each developer decision will be articulated along with the specific implementation choices.

Agent-based modeling (ABM)

The platform is designed to integrate established computational paradigms from diverse research fields, with a primary focus on the agent-based modeling (ABM) approach, which has been extensively used in computational ecology [8, 9]. This choice is motivated by theoretical and practical considerations. ABM provides a powerful framework for simulating complex, dynamic systems where individual agents -representing organisms, behavior, or components- interact with each other and with their environment, making it particularly well-suited for modeling the behavior of *Drosophila* larvae in a virtual space.

The core simulation and agent classes in *Larvaworld* are built upon the *agentpy* package, a Python-based ABM framework and a key dependency of the platform [10]. *Agentpy*'s basic *Model*, *Space* and *Object* classes are adjusted to meet the specific needs of *Larvaworld* especially the nested-dictionary implementation of agent parametrization, ensuring that the platform can accommodate the unique requirements of simulating complex modular biological systems. ABM provides a flexible and efficient structure for the creation and management of agents, including the ordering of their actions, simulation setup, and workflow, and turn-based data retrieval.

The Dynamic Energy Budget (DEB) theory

To bridge the gap between the neuroscientific and homeostatic timescales of behavioral modeling, *Larvaworld* can place fast neural control of behavior under the regulation of a slower energetics model which addresses energy allocation according to the metabolic needs of the virtual animal as these develop during the larval life stage. The energetics simulation is formulated as a dynamic energy budget (DEB) model following the established literature [11].

The DEB model takes into account the post-hatch age of the larva and the rearing conditions including nutritious substates and starvation or partial food-deprivation periods. It enables the introduction of discrete foraging phenotypes, such as rovers and sitters [12], by differentially configuring the nutrient absorption rate during feeding, eventually regulating the behaviorally expressed exploration-exploitation balance via a dynamic energy-reserve dependent hunger drive. The energetics model runs in the background of the behavioral simulation at an adjusted circadian timescale achieving realistic growth curves for virtual larvae.

Parametrization

All classes in *Larvaworld* are parameterized using the *param* Python library, a member of the *holoviz* ecosystem, developed to enhance transparent subclassing, interactive visualization through dynamic widgets (see Fig 4), and intuitive real-time display of their parameters within Jupyter notebooks, as demonstrated in the tutorials available on the documentation page (see Figures 6, 7, 5).

These interactive features are illustrated in a number of browser-based applications (see Fig 8) that can later be integrated into a comprehensive Graphical User Interface (GUI) to enhance the accessibility of the virtual laboratory, making it easier to use for researchers and educators without extensive coding experience.

Visualization

Larvaworld uses the *pygame* Python library to support visualization of behavioral simulations as well as replays of previously stored simulations and real-world experiments reconstructed from imported data. These can be run at real-time, slower or faster up to the device's computational limits. If visualization is enabled for a simulation, a pop-up screen shows the larvae and the environment objects (odor/food sources, impassable borders) on a 2D arena with a realistic spatial scale and timer. Keyboard shortcuts and mouse clicks allow real-time interaction with the visualization screen. Zooming in and out is allowed as well as selecting and locking the screen on specific individuals. Larva unique IDs as well as the screen scale and timer can be toggled on and off. Additionally, larva midline and contour can be toggled, and their head or centroid can be pointed out. Larva trajectories can be shown and their duration can be adjusted. Larvae can be colored with default or random colors or dynamically according to their instantaneous behavior. The arena color can be changed and the distribution of specific odors (odorscape) dynamically visualized. Larvae, sources and borders can be added or deleted instantly. Finally videos of the simulations and replays can be stored, arena or odorscape snapshots captured and videos collapsed on a single image file comprised of all video frames overlaid. A number of snapshots illustrating the real-time visualization of both reconstructed experiments and simulations are shown in Fig 2. A summary of the available online controls is shown in Table 1.

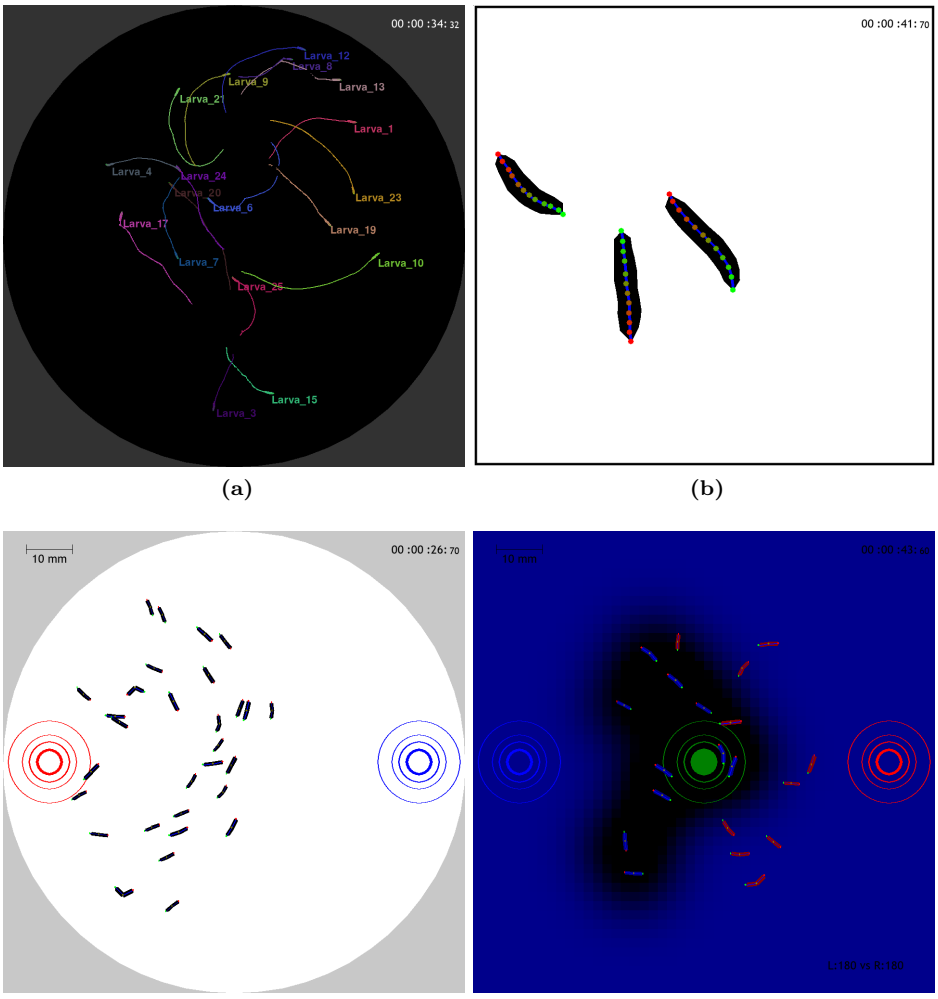

Fig 2. Real-time visualization of reconstructed real-animal experiments and agent simulations. Snapshots from the experimental arena illustrate: (a) real-animal behavior displayed on a black background with visible larval IDs and reconstructed locomotion trajectories [13]; (b) a close-up showing the 12-point midline tracking of individual larvae [13]; (c) a simulated odor preference experiment: A group of larvae is randomly placed at the centre of the dish with an appetitive odor source placed on the left (red) and a non-valenced odor on the right (blue). The larvae are gradually attracted to the side of the appetitive odor; and (d) a capture-the-flag mini-game: Two larva groups (blue on the left VS red on the right) compete to capture a highly-valenced centrally-placed odor source (green) and carry it back to their base (low-valenced odor source of the respective color). The simulated olfactory sensory landscape ("odorscape") emanating from the green odor source is shown in black over a blue background.

Screen		Drawing		Color		Interaction		Simulation/Storage	
Scale	s	midline	m	random	r	select	L*	snapshot	i
Timer	t	contour	c	behavior	b	lock screen	f	odorscape	o
IDs	TAB	head	h	background	g	delete	del	pause	p
Zoom in/out	M*	centroid	e	odorscape	0-9	add	L*		
Screen Motion	↕↔	trail	p			inspect	R*		
		trail duration	+/-			dynamic graph	g		

Table 1. Visualization default keyboard/mouse controls. L*, R*, M* stand for mouse left, right, and center buttons respectively

Replays of simulated or imported experimental datasets allow additional configuration. Inclusion of specific individuals and time range can be defined. Replays can be run on arenas different than their original one as long as all tracks remain within their boundaries. To this end, larva x-y coordinates can be transposed to the arena center. Additionally all tracks can be aligned to start from the same origin at the arena center, a mode favoring inspection of larva dispersion over time. Larva trails can be colored according to the instantaneous forward or reorientation angular velocity. For closely inspecting single-larva replays, the screen center can be locked to a specific midline point or segment and a background hue can be enabled to easily assess body movement in space. Finally when replaying imported experiments, larva bodies can be reconstructed as segmented virtual bodies of a given number of segments in order to make them visually comparable to virtual larvae.

Simulation modes

Several simulation modes are available, each serving a distinct purpose and featuring its own set of configurable parameters. Such a mode-specific parameter set fully configures a simulation and can be stored as a nested Python dictionary under a unique ID for future usage. A rich repertoire of preconfigured parameter sets is already available upon installation consisting of several simulation examples for each mode.

The Command Line Interface (CLI) can be summoned directly from the terminal using the `larvaworld` command. This has to be followed by an essential argument that defines the desired simulation mode by its respective shortcut. In some modes, a second argument is needed to retrieve an existing, stored configuration by its unique ID. Additional arguments can be used to overwrite some basic parameters of the predefined configuration. The available simulation modes are shown in Fig 3 while their usage via CLI commands is showcased separately.

Single experiment

This is the basic simulation mode which runs a single virtual experiment once, based on a detailed configuration of the virtual environment, involved larva groups, temporal simulation features and subsequent analysis pipeline. Such a configuration might aim to enable the detailed, standardized, realistic, in silico replication of a real-world experiment, in which case its setup and analysis is drawn from published behavioral studies or

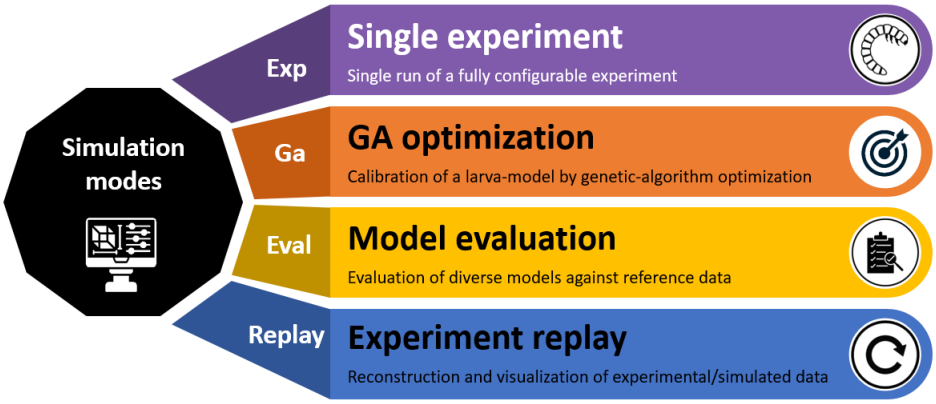

Fig 3. CLI simulation modes. The simulation modes available in *Larvaworld* along with the respective argument to launch them via the command-line interface.

established behavioral protocols. Such examples can be found among the multiple preconfigured experiments listed below (Table 2). Alternatively, it might define a difficult-to-implement or entirely fictitious experiment for model-testing, visualization or proof-of-concept purposes. As a distinct simulation mode, an experiment launched individually supports a complete pipeline to analyze the results, generate the respective plots and store the simulated datasets for future access. In contrast, the single run forms the core simulation unit, the building block for more complex simulation modes described below.

Each of these lines runs a dish simulation (30 larvae, 3 minutes) without analysis:

```
larvaworld Exp dish -N 30 -duration 3.0 -vis_mode video
larvaworld Exp patch_grid -N 30 -duration 3.0 -vis_mode video
```

This line runs a dispersion simulation and compares the results to the existing reference dataset. We choose to only produce a final image of the simulation.

```
larvaworld Exp dispersion -N 30 -duration 3.0 -vis_mode image -a
```

Genetic Algorithm for model optimization

As explained above, the *behavioral architecture* applying neural control over an agent's behavior is hierarchically layered and modular. The latter implies that diverse mutually exclusive models can compete for each module defined within the architecture. The nature of these competing module-specific models is not constrained apart from the need to adhere to the input/output specification of the module. In other words each of these candidate models might entail an intrinsic set of configuration parameters, not shared among other candidate models. An issue that surfaces when comparing across them, is that each model should be represented by the best possible parameter configuration, so that the comparative simulation study indeed selects between optimal model instances. It follows that an optimization process for each candidate model that will yield its optimal representative, must precede the comparative study. This need has motivated the integration of a genetic algorithm (GA) optimization process into *Larvaworld*.

The GA simulation mode accepts three sets of parameters as illustrated in Fig 4 :

Type/Behavior	Experiment	Description	Literature source
exploration	close-view	Single larva closely inspected in a tiny arena	-
	dish	Exploration of a non-nutritious Petri-dish	-
	dispersion	Larva dispersion from the arena center	-
chemotaxis	navigation	Navigation up an odor gradient	Gomez-Marin et al. (2012) [14]
	local search	Exploration in the vicinity of an odor source	Gomez-Marin et al. (2012) [14]
odor preference	train & test	Olfactory associative learning (train & test phase)	The Maggot Learning Manual
	test on/off food	Test in the presence/absence of nutritious substrate	The Maggot Learning Manual
foraging	patchy food	Foraging in arena with one/two/multiple food patches	-
	uniform food	Foraging in uniformly distributed nutritious substrate	-
growth	rearing	Larva rearing in ad-libitum conditions	-
	rovers VS sitters	Foraging phenotypes compared in diverse conditions	Kaun et al. (2007) [12]
imitation	realistic bodies	Multisegment larvae in Box2D physics engine	-
	dataset imitation	Experimental dataset imitation	-
games	maze	Navigation in a maze towards an odor source	-
	capture/keep the flag	Larva teams competing for a portable nutritious object	-

Table 2. Summary of preconfigured behavioral experiments

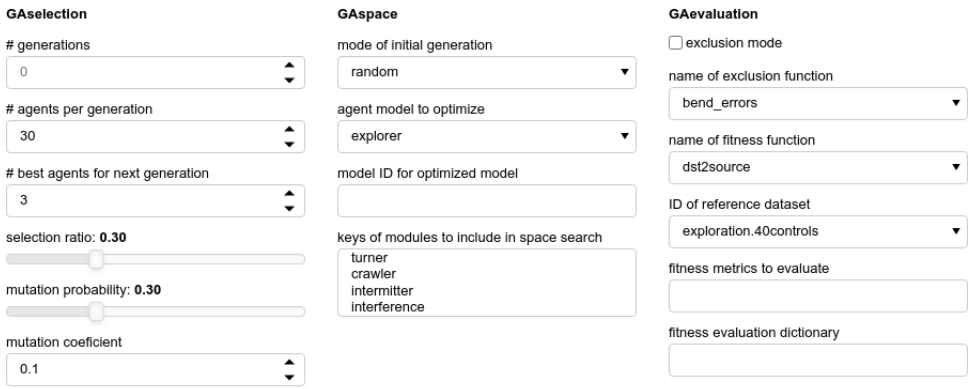

Fig 4. GA configuration panel

- **Selection algorithm** : Here we define the number of generations to run and the number of agents per generation. Additionally, the criteria to select and optionally mutate agent configurations to generate the next generation.
- **Parameter space** : A basic model configuration needs to be provided along with the specification of the parameters that will be optimized. Each of the parameters-to-be optimized is part of an existing module of the agent's behavioral architecture and the allowed range of the space search is predefined.
- **Performance evaluation** : Parameters that ultimately define a fitness function that quantifies the perfomance of each agent in each generation. In the simplest case the fitness function is defined externally and provided as an argument. Alternatively, it may be constructed by the GA engine based on provided parameters. For example when the agents are evaluated against an experimental reference dataset.

This line optimizes a model for kinematic realism against a reference experimental dataset :

```
larvaworld Ga realism -refID exploration.30controls -N 20 -duration 0.5
    ↪ -mID1 GA_test_loco -init_mode model
```

Model evaluation - comparison with real data

Sometimes we focus on evaluating our agent models against some reference experimental dataset. The model evaluation mode does exactly that. As shown in Fig 5, it requires defining :

- Reference dataset, designated via ID or directory
- Larva models retrieved via ID and the size and IDs of the respective larvagroups
- Evaluation metrics and configuration
- Simulation experiment and its setup

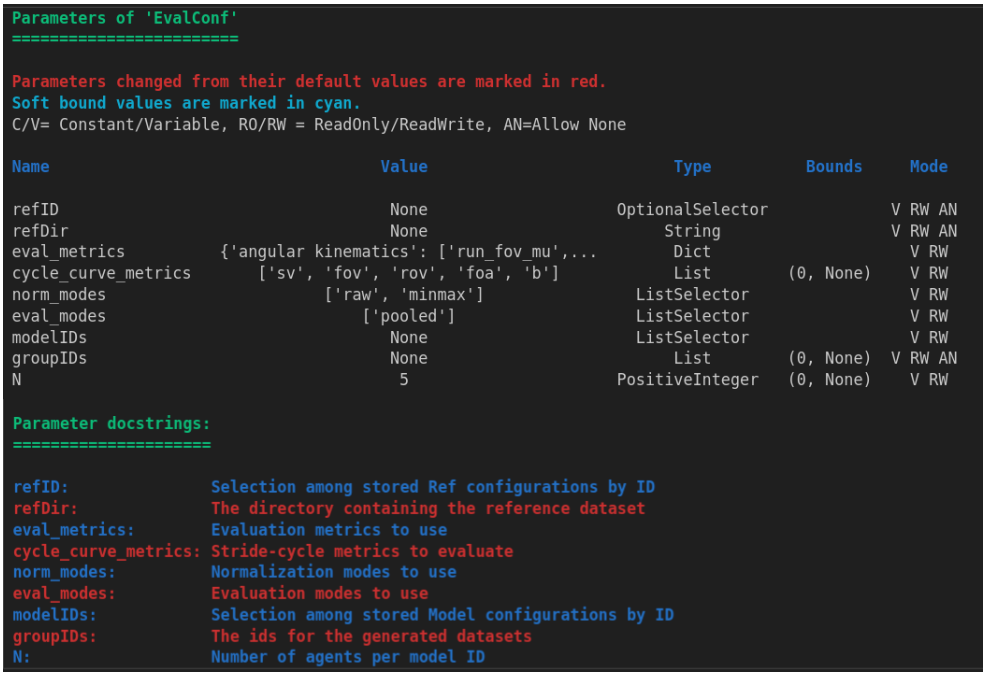

Fig 5. Model Evaluation configuration parameters

Illustrative examples of model evaluation simulations are presented in the Results section. This mode is available via CLI as well. The following line evaluates two models against a reference experimental dataset :

```
larvaworld Eval -refID exploration.30controls -mIDs RE_NEU_PHI_DEF
    ↪ RE_SIN_PHI_DEF -N 3
```

Experiment replay

Replay a real-world experiment. This line replays a reference experimental dataset :

```
larvaworld Replay -refID exploration.30controls -vis_mode video
larvaworld Replay -refDir SchleyerGroup/processed/exploration/30controls
    ↪ -vis_mode video
```

Simulation configuration

Environment setup

The simulation environment in *Larvaworld* is designed as a spatial arena with configurable shape and dimensions. Within this arena, single or grouped odor/food sources and impassable borders can be placed at specified locations. A key feature of the environment is its customizable sensory landscape, which includes gradients of various sensory modalities such as olfactory (odorscape) emanating from odor sources placed inside the arena, thermal (thermoscape), and wind (windscape). The parameters of the respective class are shown in Fig 6. Mini videos of windscape and odorscape are available online [15–19].

To facilitate diverse experimental setups, multiple predefined environmental configurations are available, each of which can be adjusted to suit specific research needs.

```
Parameters of 'EnvConf'
=======================

Parameters changed from their default values are marked in red.
Soft bound values are marked in cyan.
C/V= Constant/Variable, RO/RW = ReadOnly/ReadWrite, AN=Allow None

Name                          Value                        Type      Mode

arena         AreaUnit(dims=(0.1, 0.1), geometry='c...  ClassAttr    V RW
food_params   FoodConf(food_grid=None, name='FoodCo...  ClassAttr    V RW
border_list                     {}                      ClassDict    V RW
odorscape                      None                     ClassAttr  V RW AN
windscape                      None                     ClassAttr  V RW AN
thermoscape                    None                     ClassAttr  V RW AN

Parameter docstrings:
=====================

arena:       The arena configuration
food_params: The food sources in the arena
border_list: The obstacles in the arena
odorscape:   The sensory odor landscape in the arena
windscape:   The wind landscape in the arena
thermoscape: The thermal landscape in the arena
```

Fig 6. Virtual environment configuration parameters

Researchers can modify attributes such as the position and intensity of sensory sources, the spatial distribution of sensory gradients, and the boundaries of the arena. This flexibility enables precise control over the simulation environment, allowing for the design of experiments that replicate real-world conditions or explore novel scenarios.

Substrate

Larva rearing and behavioral lab experiments are carried out on diverse nutritious or non-nutritious substrates. Several substrate types have been established in terms of their nutrient compound composition. Some of these have been implemented in *Larvaworld* with their characteristic composition. Virtual larvae can be grown, starved or tested on such substrates (Table 3). Any food source in the arena is characterized by a substrate of specific type, nutritional quality and available amount of food. Alternatively, the substrate can be placed locally as patches in a grid that covers the entire arena as a grid where each cell can independently hold a certain amount of food. This food can be detected and consumed until depleted.

Larva groups

In *Larvaworld*, virtual larvae are generated in groups, with members of each group sharing specific traits that distinguish them from individuals in other groups. These traits include not only the configuration of the larva model itself but also elements such as a shared life history, a spatial distribution that defines their initial pose (Table 4), a group-specific color for easy visualization and a distinct odor signature (Fig 7).

Substrate	Compound density ($\mu g/ml$)						Literature source
	glucose	dextrose	saccharose	yeast	agar	cornmeal	
standard-medium	100	-	-	50	16	-	Kaun et al. (2007) [12]
PED-tracker	-	-	10	187.5	5000	-	Schumann et al (2020) [20]
cornmeal	-	70.3	-	14.1	6.6	65.6	Wosniack et al. (2021) [21]
sucrose	17.1	-	-	-	4	-	Wosniack et al. (2021) [21]

Table 3. Compound composition of established nutritious arena substrates

Parameter	Description
N	Number of virtual larvae in the group
location	Centre of the spatial distribution in the arena
scale	Spatial extent of the distribution
shape	Shape of the distribution (circular/rectangular/oval)
placement	Placement within the distribution's shape (uniform/normal/periphery)
orientation	Range of initial spatial body orientations

Table 4. Larva group initial spatial placement parameters

Virtual larvae can have their own life history at the moment they enter the experimental arena. To simulate that, a *Drosophila*-specific dynamic energy budget (DEB) model simulates larva growth up to a specific age post-hatch, on a predefined rearing substrate (Table 3) of specified quality so that the behavioral simulation can be subsequently initiated. Periods of food deprivation or complete starvation during rearing or during the actual behavioral simulation can also be defined.

Groups can be assigned a linked reference dataset from which parameters are sampled. The sampling mode is highly configurable, offering three main options: optimizing for an average individual, preserving interindividual variability, or replicating the reference dataset on an individual-by-individual basis for specified parameters. For a detailed comparative evaluation of these group-generation modes, see *Individuality and Variability* in the *Results* section.

This functionality provides researchers with the flexibility to design experiments featuring multiple, distinguishable groups of virtual larvae. By enabling group-specific traits and parameter configurations, *Larvaworld* supports a broad spectrum of experimental scenarios, ranging from comparative analyses to simulations that emulate the diversity and dynamics observed in real-world populations.

Web-based applications

A number of web-based *Larvaworld* applications are available to facilitate inspection, configuration and real-time visualization of various elements. The panel showing the

```
Parameters of 'LarvaGroup'
==========================

Parameters changed from their default values are marked in red.
Soft bound values are marked in cyan.
C/V= Constant/Variable, RO/RW = ReadOnly/ReadWrite, AN=Allow None

Name                          Value                          Type          Mode

group_id                    'LarvaGroup'                    String         V RW
model                          None                    OptionalSelector  V RW AN
color                        'black'                        Color          V RW
odor          Odor(id=None, intensity=None, name='O...     ClassAttr        V RW
distribution  Larva_Distro(N=30, loc=(0.0, 0.0), mo...     ClassAttr        V RW
life_history  Life(age=0.0, epochs=ItemList (0 obje...     ClassAttr        V RW
sample                         None                    OptionalSelector  V RW AN
imitation                      False                       Boolean         V RW

Parameter docstrings:
=====================

group_id:     The distinct ID of the group
model:        Selection among stored Model configurations by ID
color:        The default color of the group
odor:         The odor of the agent
distribution: The spatial distribution of the group agents
life_history: The life history of the group agents
sample:       Selection among stored Ref configurations by ID
imitation:    Whether to imitate the reference dataset.
```

Fig 7. Virtual larva group parameters

Application	Description
Experiment Viewer	Inspect/launch preconfigured experiments
Larva Models	Inspect/visualize modular larva-models
Locomotory Modules	Inspect/test behavioral modules
Track Viewer	Visualize stored datasets

Table 5. Web-based applications

available applications can be launched via the CLI command `larvaworld-app`. Some of the available applications are listed in Table 5 and a snapshot is shown in Fig 8.

Data management and analysis

Standardized format for both experimental and simulated datasets

Larvaworld is set up primarily to analyze and simulate *Drosophila* larva motion-tracking experiments. In such studies recorded data are usually collected per trial and per animal group. Datasets are then comprised of a number of (identical) trials for each of a number of animal groups (for example different genotypes, satiation states, rearing conditions) and their control counterparts. Furthermore it is common to perform the same behavioral experiment under a number of different environmental conditions or even to compile an essay of diverse behavioral experiments across animal groups. To make the importation of such experimental datasets straightforward to the user, the platform's data organization builds on such a *LarvaDataset* as a core class. Crucially

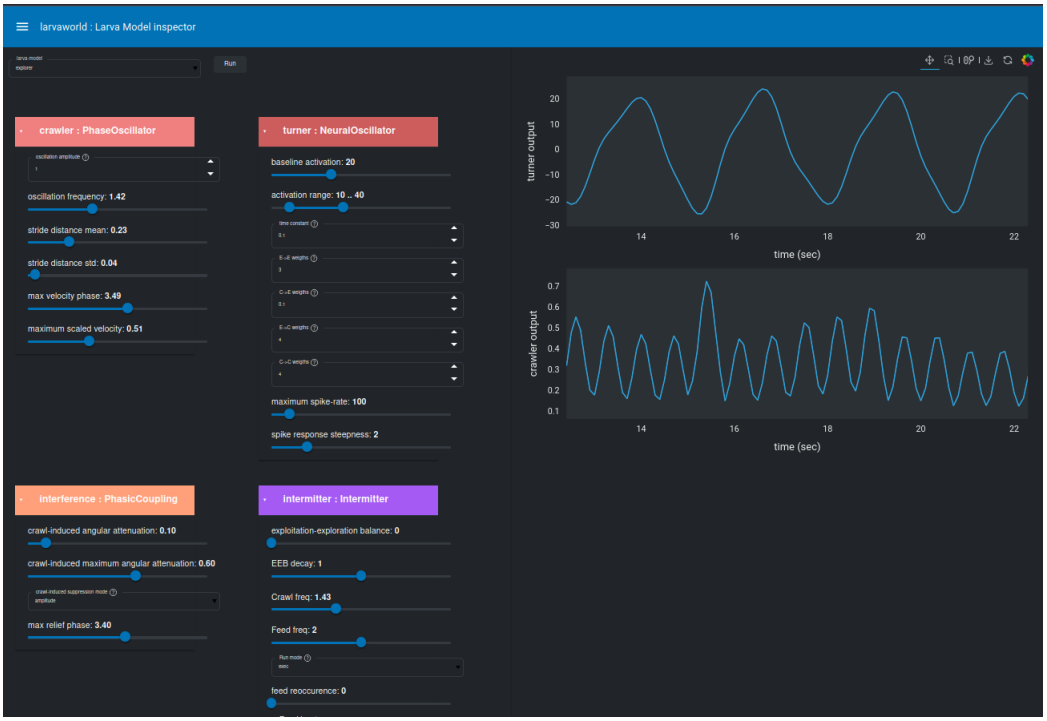

Fig 8. Web-based *Larvaworld* application to inspect the modular composition of any preconfigured locomotory larva-model. This can be selected by its unique ID from a drop-down list of all available models. The configuration of all 4 basic modules comprising the locomotory layer of the model is available on the left sidebar. Real-time simulation variables (e.g. input/output) are dynamically plotted in the middle.

this class is common for both simulated and experimental data in order to facilitate comparison and model fitting. The basic elements of a *LarvaDataset* instance are :

- Timeseries data computed at every timestep for every larva are stored as a double-indexed *Pandas* dataframe where the index defines the timestep and the unique ID of the larva, while the values of each parameter are stored in a separate column. Initially the dataframe contains only the primary tracked parameters (x-y coordinates of at least the centroid, often a number of points along the larva midline and optionally the contour of the larva body). The dataframe is enriched with derived parameters during data-processing as described in the section below.
- Endpoint metrics computed once for every larva at the end of the simulation are also stored as a *Pandas* dataframe indexed by the agent's unique ID. This is also enriched during data-processing with summary measurements of the timeseries data.
- Metadata about experimental conditions, tracking parameters, animal groups etc are structured as a nested dictionary. Storage paths are also stored there for easy access during analysis.

To facilitate efficient data storage, retrieval, and analysis the dataframes are stored in an *HDF* file under different keys (e.g. midline, contour, trajectory) for easy access. The metadata dictionary is stored as a configuration text file alongside the data. The dataset can optionally be registered as a reference dataset, by providing a unique ID as key, a

	Function	Description
Preprocessing	Scaling	x-y scaling by a scalar
	Transposition	x-y transposition
	Alignment	Trajectory alignment e.g. to common origin
	Interpolation	Missing data interpolation
	Exclusion	Data exclusion on condition
	Filtering	Low-pass filtering at a cut-off frequency
Processing	Angular analysis	Bend/orientation angle, angular velocity/acceleration
	Spatial analysis	Spatial distance/velocity/acceleration
	Forward components	Spatial metric components along orientation axis
	Dispersal	Larva spatial dispersal during time ranges
	Tortuosity	Trajectory tortuosity for sliding temporal windows
	Odor preference	Olfactory preference index
	Odor concentration	Absolute and perceived odor concentration along trajectory
Annotation	Strides/Crawl-runs	Individual strides and uninterrupted chains of concatenated strides
	Crawl-pauses	Immobility epochs without peristaltic strides
	Turns	Turning events based on reorientation amplitude or angular velocity
	Bout analysis	Spatial/angular metric change during bouts
	Bout distribution	Distribution fitting for bout duration/length

Table 6. Data processing methods

functionality often used to easily access real experimental data during model evaluation and optimization but also for quickly visualizing comparative plots of simulated and experimental data. Additional files are created during data-processing.

Unbiased parameter computation and data analysis from primary 2D time-series

In *Larvaworld*, data processing methods are identical for both simulated and imported experimental data, ensuring consistency and unbiased analysis. The platform supports three sequential, configurable pipelines that can be applied to any dataset containing at least the x-y coordinates of a single point per larva (Table 6):

During preprocessing the dataset x-y series can be rescaled by a scalar (e.g. to convert datasets stored in millimeters to the default *Larvaworld* space unit of meters) or transposed to align with the arena center. Datasets can also be aligned to a common origin, facilitating visualization of dispersal patterns. Noise from tracking artifacts can be reduced using low-pass filtering at a configurable cut-off frequency. Additionally, missing data points can be interpolated, or specific data slices can be excluded based on predefined conditions (e.g. instances of larva collisions).

In the main processing stage, secondary parameters are derived from the primary x-y time series:

- Angular analysis: Bending angles and absolute orientation in 2D space are computed for individual larval segments. Instantaneous orientation and bend can be defined either directly from these angles or via front/rear body vectors.
- Spatial analysis: Key metrics such as distance, velocity, and acceleration (including their components along the larva's forward orientation axis) are calculated.
- Dispersal patterns over specified time intervals

- Trajectory tortuosity using sliding windows of configurable durations
- Odorscape Navigation: Metrics such as instantaneous odor concentration and perceived concentration changes along the larva's trajectory are available. Additionally it is possible to track the instantaneous distance and bearing to odor or food sources.
- Preference index calculations for olfactory preference experiments.

Finally the bout annotation pipeline identifies discrete behavioral events, including:

- Strides: Individual locomotor steps.
- Crawl-Runs: Continuous sequences of strides.
- Crawl-Pauses: Periods of immobility.
- Turns: Defined either by reorientation amplitude or changes in orientation/bending angular velocity sign.

A fitting algorithm determines optimal distributions for the temporal durations or spatial lengths of bouts, selecting from options such as power-law, exponential, and log-normal distributions. Additional metrics include distance, orientation change, and bending angle change during individual bouts, as well as crawling frequency estimates. For virtual larvae, the platform enables tracking of a broad range of model-specific, simulation-derived metrics, offering flexibility for detailed behavioral analysis and model evaluation.

Importing experimental datasets from diverse tracker setups

Experimental laboratories employ a wide range of tracking software to accurately capture 2D larval locomotion. Typically, primary data consist of coordinates for multiple points along the body midline and around the body contour, tracked at either fixed or variable frame rates. These raw data are exported in tracker-specific formats and subsequently processed to derive secondary parameters for analysis. *Larvaworld* supports the direct import of tracked datasets from a variety of tracker-specific formats (Table 7). For each supported format, the platform defines key parameters such as the tracking frame rate and the number of midline or contour points, along with a format-specific conversion function to transform the raw data into the standardized LarvaDataset format. Users can specify which tracks to include using a set of import arguments, such as track duration, start or termination time, and the option to limit the dataset size by defining a maximum number of animals. The imported datasets are standardized, making them directly comparable both to each other and to simulated data. Importantly, to promote reproducibility, only primary parameters are imported, ensuring that all subsequent derived metrics are transparently calculated within the platform. This approach facilitates consistent and unbiased comparisons across experimental and simulated datasets while maintaining a high degree of accessibility and flexibility for data analysis.

Results

Scientific applications

The *Larvaworld* package has already been successfully used in several scientific studies. These will be briefly described here, focusing on the way each of them used *Larvaworld*.

Lab	Framerate (Hz)	Midline (#)	Contour (#)	Source
Schleyer	16	12	22	Paisios et al. (2017) [22]
Jovanic	11.27*	11	30**	de Tredern et al. (2024) [23]
Berni	2	1	0	Sims et al. (2019) [6]
Arguello	10	5	0	Kafle et al. (2025) [24]

* variable, average framerate used instead

** variable, convex hull of defined size used instead

Table 7. Lab-specific experimental data-formats

As a whole, these studies illustrate the spectrum of functionalities supported by the package, both as a data analysis and a behavioral modeling tool.

The simplest and most straightforward case is usage of Larvaworld as a data analysis tool. A study of feeding-state dependent aversive behavioral responses of *Drosophila* larvae to mechanosensory stimulation aimed to elucidate neuropeptidergic modulation of reciprocally interconnected inhibitory neurons [23]. The tracker-specific datasets of recorded larva locomotion across all experimental groups were successfully imported into *Larvaworld*. Using the platform's comparative analysis methods, differences in several key kinematic locomotory parameters were identified among fed, starved, and sucrose only fed groups. The analysis was further extended to examine locomotory differences in dehydrated (sucrose only fed) and rehydrated, as well as food-deprived and refed animals, across multiple genetically distinct strains. This comprehensive approach allowed for a detailed assessment of how nutritional and hydration states influence locomotory behavior.

The *behavioral architecture* framework that forms the backbone of the modular modeling functionality in Larvaworld has been presented in [4]. In this study both data analysis and modeling is done via *Larvaworld*. Novel insights derived from the analysis of experimental datasets are utilized in order to extend a previously published locomotory model of the larva [25] with two novel features, namely crawl-bend interference and behavioral intermittency, the latter also based on a previous modeling study [26]. The resulting intermittent coupled-oscillator model of larva locomotion is the default model in *Larvaworld* for performing behavioral simulations in stimulus free or sensory-rich environments, in the latter case augmenting the locomotory model with sensors to detect e.g. odor or food sources.

The modular approach to behavioral modeling featured in *Larvaworld* is showcased in a study of associative learning and conditioned behavior of the *Drosophila* larva [27]. The study features a spiking model of the larva's mushroom body (MB), a neuropile that has been implicated in memory and learning. In order to evaluate the MB model in realistic odor-preference simulations, and to compare the results to lab experiments, it has been plugged in as a memory module. This highlights how the modular architecture of *Larvaworld* contributes to answering relevant research questions across domains.

Notably the resulting modular larva model was a hybrid also in the sense that its comprising modules operated in nested timescales, namely the spiking MB module at a 0.1 msec timestep whereas the behavioral locomotory simulation at the default 0.1 sec timestep. This is an illustration of the capability of the behavioral architecture to

integrate disparate computational modules regardless of their level of detail or temporal resolution.

Moreover the modular modeling approach facilitates extendability. In a comparative evolutionary study of thermotaxis across drosophilids, larvae from 8 distinct species have been tested in a suitably configured thermal-gradient arena (thermoscape) in order to quantify and compare their preferred temperatures [24]. Simulations were used to evaluate competing hypotheses about the neural underpinnings of the observed temperature preferences across species. In order to couple real animals to their virtual counterparts, species-specific models were fitted. A virtual thermoscape arena was configured and the default locomotory model was extended with an additional thermosensation modality. The performed temperature-preference simulations closely resembled the experiments in their spatiotemporal configuration, larva groups and behavioral results, while additionally allowing the evaluation of competing circuitries of thermotaxis models.

Import and analysis of experimental datasets

This example illustrates the process of importing datasets of *Drosophila* larva locomotion. The goal is to compare the locomotion patterns of larvae subjected to different diets over 5 hours : normally fed, fed only with sucrose (therefore protein-deprived), and completely starved. The distinct metabolic states of the larva groups might have an impact on their locomotion. We specifically focus on the temporal evolution of their dispersal in space.

The import process involves converting the raw tracker-specific data format into the *Larvaworld* format. This conversion is performed only once, and the processed data is stored for future use without the need for repeated conversions.

Once the datasets are imported, various visualizations can be generated to analyze the locomotion patterns of the larvae. The trajectories of the larvae are plotted, showing their movement over time. Boxplots of endpoint metrics such as mean, final, and maximum dispersal during the first 60 seconds are generated to provide a statistical comparison of the groups. In addition, the dispersal of larvae from their starting point is plotted over time, capturing the mean and variance of their movement (see Fig 9). Replay simulations are run to visualize the trajectories of the larvae aligned at the origin, and videos of these simulations are generated and combined into a single video.

The combined video output from the replay simulations provides a visual comparison of the locomotion patterns of the three larva groups, showcasing the impact of different metabolic states on larval locomotion. The complete tutorial notebook, including the code and detailed steps, is available in the software's documentation.

Demonstration of a chemotaxis essay

To illustrate the usability of the virtual lab, we present a demonstration of a chemotaxis assay, replicating two well-established experimental paradigms [14]. In the first experiment, both the odor source and the larvae are positioned at the center of the arena, where the larvae are expected to exhibit zigzagging or orbiting trajectories around the source, remaining in its vicinity. In the second experiment, the larvae and the odor source are placed on opposite sides of the arena, prompting the larvae to approach the source.

We focus on two distinct chemotaxis algorithms that have been proposed in the literature. The first, as described in [25], suggests that lateral bending is driven by continuous

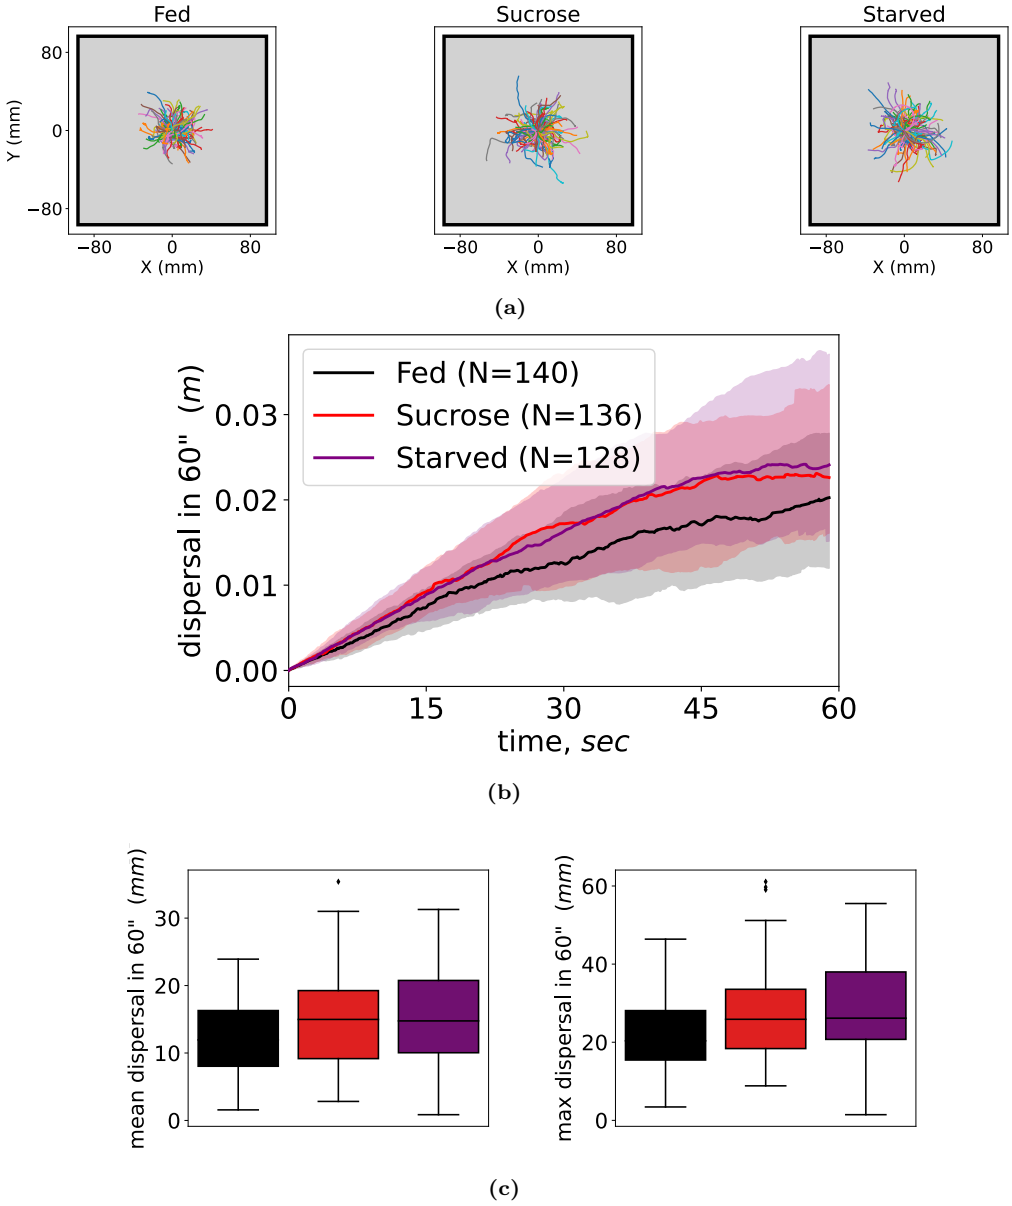

Fig 9. Impact of metabolic state on spatial dispersal analyzed in lab experiments. Three larva groups subjected to different dietary conditions are compared in terms of their spatial dispersal in a quadratic arena during the first minute of the experiment: (a) The trajectories of the larvae, aligned to start from the center of the arena. (b) Temporal course of the larva spatial dispersal. Line indicates the group median while shaded area denotes first and third quartiles. (c) Boxplots of average and maximum dispersal.

Video: The temporal course of dispersal for the three groups shown in (a) [28].

Model configurations	Tastekin (log)	Tastekin (lin)	Wystrach (log)	Wystrach (lin)	controls
olfaction sensory transduction mode	logarithmic	linear	logarithmic	linear	logarithmic
ability to interrupt locomotion	True	True	False	False	False
appetitive odor valence	True	True	True	True	False

Table 8. Chemotaxis essay. Five locomotory models are tested in a chemotaxis essay. The two basic models employ different algorithms for turning behavior, namely "Tastekin" is a stochastic Levy-walker with the ability to interrupt locomotion when navigating down-gradient , while "Wystrach" features a lateral oscillator generating continuous lateral bending, modulated by the incoming olfactory stimulation. The models are tested using either a logarithmic or a linear function of instantaneous concentration change for sensory transduction. An olfactory-blind model is used as control.

oscillations, which are dynamically modulated by incoming olfactory stimuli. This model lacks behavioral intermittency, resulting in uninterrupted crawling along curved trajectories. In contrast, an alternative model proposes that larvae interrupt their crawling when navigating down an odor gradient to reorient themselves toward the appetitive odor source [29]. This finding can be integrated as an extension of the Lévy walk model, adjusting turning probabilities based on incoming sensory stimulation, as acknowledged in previous studies [30]. The model configurations are summarized in Tab 8.

In *Larvaworld* simulations, multiple models can be evaluated simultaneously, with each model represented by a dedicated group of virtual larvae. Utilizing this functionality, we implemented two variations of each chemotaxis model by modifying the sensory transduction algorithm, which translates changes in odor concentration into olfactory stimulation. Additionally, we included an odor-blind control group to serve as a baseline for comparison.

The results of the chemotaxis assay are presented in Fig 10-11. Analysis of the simulated trajectories and corresponding absolute concentration time plots indicates that the intermittent Lévy walk model demonstrates superior performance in both experiments. Meanwhile, the lateral oscillator model performs better than chance in the approach experiment. It is important to note that this demonstration does not incorporate any model optimization step (see *Genetic Algorithm Optimization* in *Simulation Modes*), and thus, further parameter adjustments could potentially improve model performance.

Model evaluation

Evaluation of a model configuration is performed by comparing the dataset derived from the simulation of the respective virtual larva group to the empirical dataset of real larva recordings. Given that the analysis pipeline is exactly the same in both cases, the empirical and simulated datasets have exactly the same shape and format and are therefore directly comparable. As there is no established methodological framework for evaluating behavioral similarity to the real animals, we use a broad array of metrics and refrain from defining a unique global error score per model.

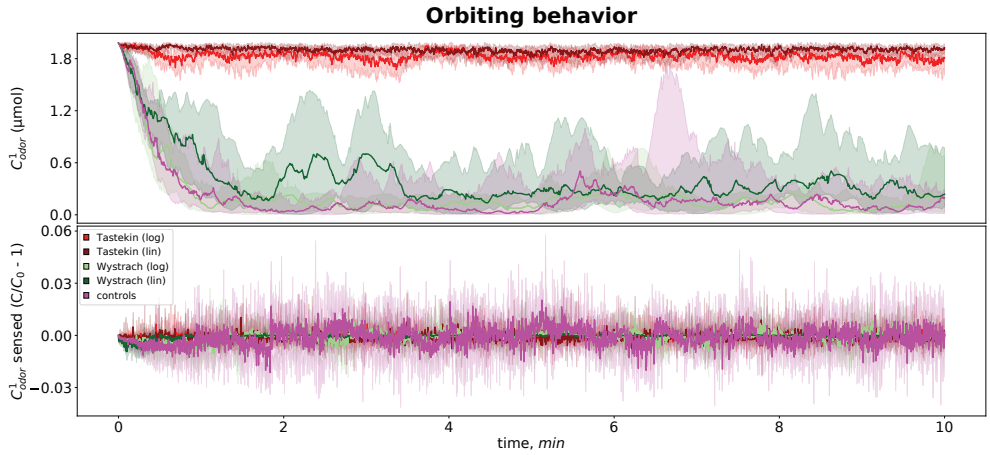

(a)

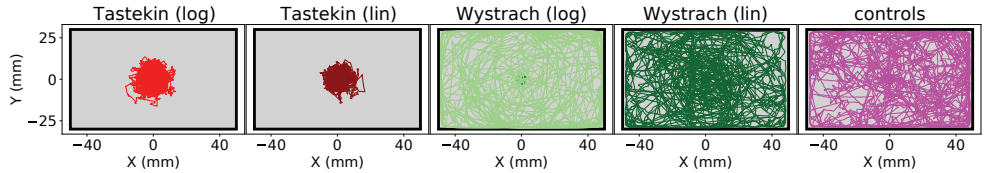

(b)

Fig 10. Chemotaxis around the odor source. (a) Time plots of the absolute odor concentration (top) and perceived odor concentration change (bottom) encountered by larvae during the experiments, along with (b) the larva trajectories around the odor source.

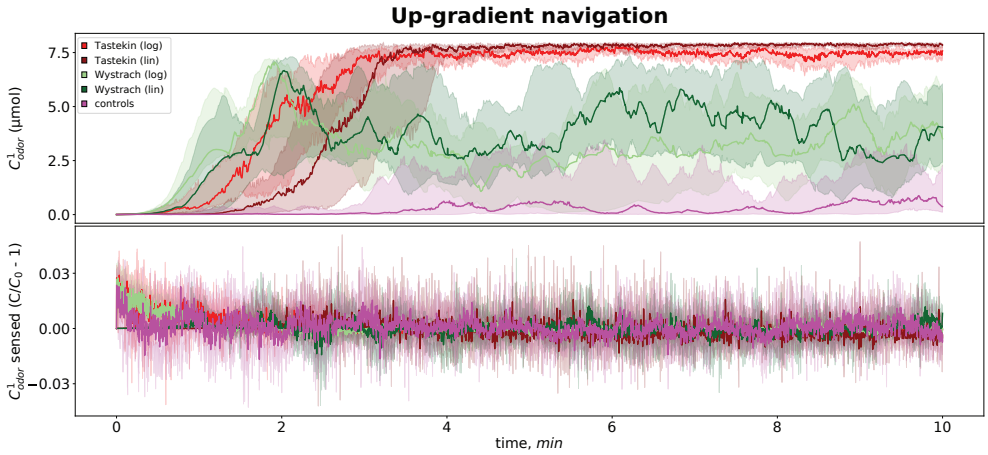

(a)

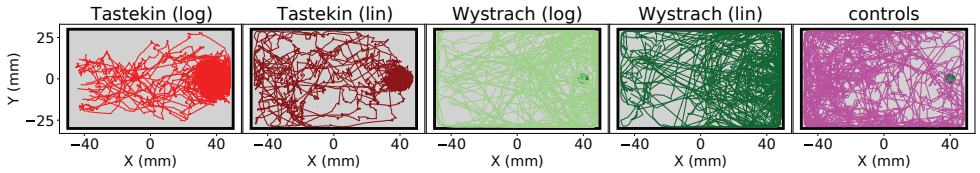

(b)

Fig 11. Chemotaxis towards the odor source. (a) Time plots of the absolute odor concentration (top) and perceived odor concentration change (bottom) encountered by larvae during the experiments, along with (b) the larva trajectories towards the odor source.

Metric	PHI→NEU	PHI→SIN	SQ→NEU	SQ→SIN
mean angular velocity during runs	**0.693**	0.9	0.9	0.967
mean angular velocity during pauses	**0.14**	0.633	0.767	0.47
total pathlength	0.51	**0.46**	0.467	0.517
mean velocity	0.373	**0.33**	0.347	0.38
mean velocity during pauses	**0.993**	1	**0.993**	1
mean velocity during runs	0.553	0.553	**0.513**	0.54
# strides	0.46	**0.39**	0.45	0.463
dominant crawling frequency	0.427	**0.407**	0.447	0.44
dominant bending frequency	0.667	0.693	**0.467**	0.693
% time crawling	0.507	0.43	**0.387**	0.46
% time pausing	0.527	**0.457**	0.46	0.507

Table 9. Locomotory model evaluation. Exact KS_D values for all endpoint-based evaluation metrics presented in Fig 12. Each row corresponds to a specific metric, and each column to one of the four evaluated model configurations. For each metric the best fitting model is highlighted in bold.

The evaluation metrics used are of two kinds. First, endpoint metrics where a single measurement is performed per larva, useful to summarize observations as a sum, average, variance, extrema, dominant frequency etc. Second, the timeseries of a continuous metric derived by the step-by-step simulations of individual animats, equal in shape to the number of simulation timesteps (3' duration at 16 Hz framerate yields 2880 timesteps). In both cases the values are pooled together across the group and the error is computed as the Kolmogorov-Smyrnov distance D_{KS} between the pooled real and simulated distributions.

A summary of the evaluation process is illustrated in Fig 12. Four different model configurations have been calibrated through the pipeline described above. Specifically the models are combinations of the two turner module implementations ("NEU" for the neural oscillator, "SIN" for the sinusoidal oscillator) and the two oscillator-coupling modes ("PHI" for a continuous gaussian $c_{CT}(\phi_C)$, "SQ" for a step-change of the angular suppression during a stride-cycle phase interval) The evaluation metrics are grouped into angular, spatial and temporal domains with distinct color coding for visualization purposes. The error scores of the various parameters for each of the 4 cases are min-max normalized before visualization. Exact values used for plotting are provided in Tab 9 and Tab 10.

Individuality and variability

To test the effect of inter-individual variability in simulations, we first measured a number of endpoint parameters across a population of 200 larvae and fitted a multivariate Gaussian distribution. We select the five crawler-related parameters described in the crawler calibration section. A generated set of these parameters is adequate to completely define the crawler module. We run an evaluation simulation to compare the above-described average model to the group-level variability method. The former is therefore represented by a group of 50 identical animats while the latter by a group of 50 non-identical animats each with a parameter-set sampled from the fitted multivariate Gaussian. The results are shown in Fig 13 and Tab 11-12.

Metric	PHI→NEU	PHI→SIN	SQ→NEU	SQ→SIN
bending angle	**0.079**	0.109	0.112	0.14
angular velocity	**0.031**	0.067	0.082	0.102
angular acceleration	0.181	**0.166**	0.217	0.181
rear angular velocity	**0.089**	0.142	0.145	0.167
rear angular acceleration	**0.143**	0.201	0.201	0.209
turn-angle amplitude	0.267	0.291	**0.151**	0.288
run distance	0.358	**0.318**	0.354	0.33
# strides per run (stridechain)	0.317	0.3	0.328	**0.287**
velocity	0.159	0.158	0.159	**0.157**
acceleration	0.086	**0.085**	0.086	**0.085**
tortuosity over 5”	**0.241**	0.649	0.303	0.678
tortuosity over 20”	**0.224**	0.336	0.351	0.389
run duration	0.346	**0.318**	0.358	**0.318**
pause duration	0.1	0.096	**0.087**	0.105

Table 10. Locomotory model evaluation. Exact KS_D values for all timeseries-derived metrics presented in Fig 12. Each row corresponds to a specific metric, and each column to one of the four evaluated model configurations. For each metric the best fitting model is highlighted in bold.

Metric	NEU mean	NEU var	SIN mean	SIN var
mean angular velocity during runs	**0.63**	0.64	0.907	0.92
mean angular velocity during pauses	**0.147**	0.243	0.393	0.193
total pathlength	0.443	0.187	0.433	**0.163**
mean velocity	0.297	**0.097**	0.283	0.187
mean velocity during pauses	0.987	**0.717**	0.983	0.8
mean velocity during runs	0.51	**0.337**	0.51	0.36
maximum dispersal in 40”	0.713	**0.653**	0.713	0.697
maximum dispersal in 60”	0.713	**0.653**	0.713	0.697
# strides	0.303	**0.093**	0.283	0.117
dominant crawling frequency	0.427	**0.057**	0.437	0.067
dominant bending frequency	**0.667**	**0.667**	0.693	0.693
% time crawling	0.27	**0.217**	0.283	0.257
% time pausing	0.383	**0.18**	0.357	0.183

Table 11. Average vs group-variability model. Exact KS_D values for all endpoint-based evaluation metrics presented in Fig 13. Each row corresponds to a specific metric, and each column to one of the four evaluated model configurations. For each metric the best fitting model is highlighted in bold.

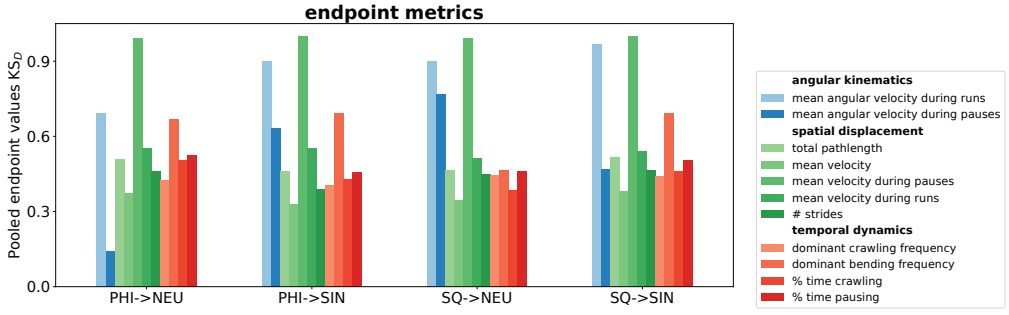

(a)

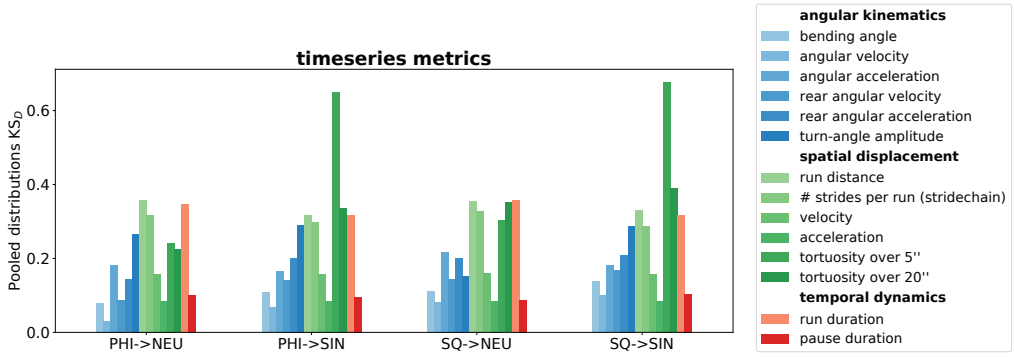

(b)

Fig 12. Locomotory model evaluation. Four locomotory models are evaluated against an empirical dataset of 150 larvae freely exploring a stimulus-free Petri-dish for 3 minutes. Two turner-oscillator implementations ("NEU" : neural, "SIN" : sinusoidal) are combined with two crawler-turner coupling modes ("PHI": smooth gaussian suppression relief, "SQ" : acute step-wise suppression relief). For each model configuration 150 identical animats are placed at the exact same initial positions and with the same initial orientations as the real larvae and are simulated with the same time step as the tracked dataset over the duration of the experiment. Model evaluation is carried out over two sets of metrics: (a) single-valued endpoint measurements such as sums, averages, counts or dominant frequencies and (b) step-by-step timeseries spanning the entire experiment duration. In both cases data is pooled across the entire group and the resulting distribution's Kolmogorov-Smyrnov distance KS_D to the empirical one is computed. For visualization purposes the metrics are grouped into angular, spatial and temporal domains with distinct color coding.

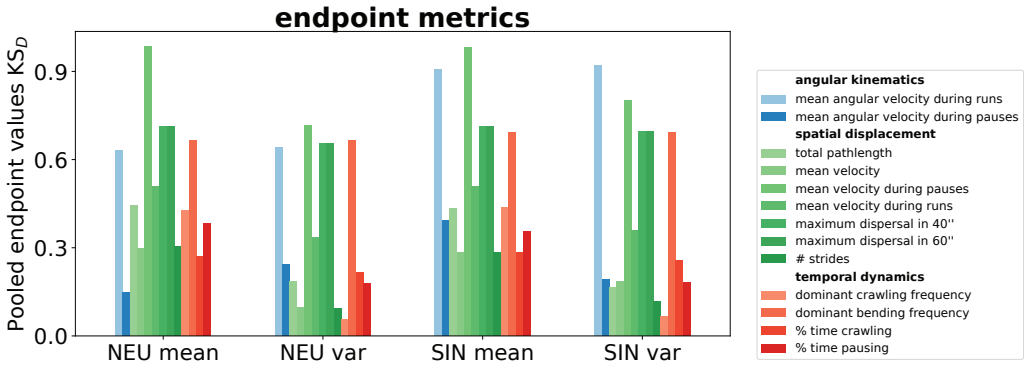

(a)

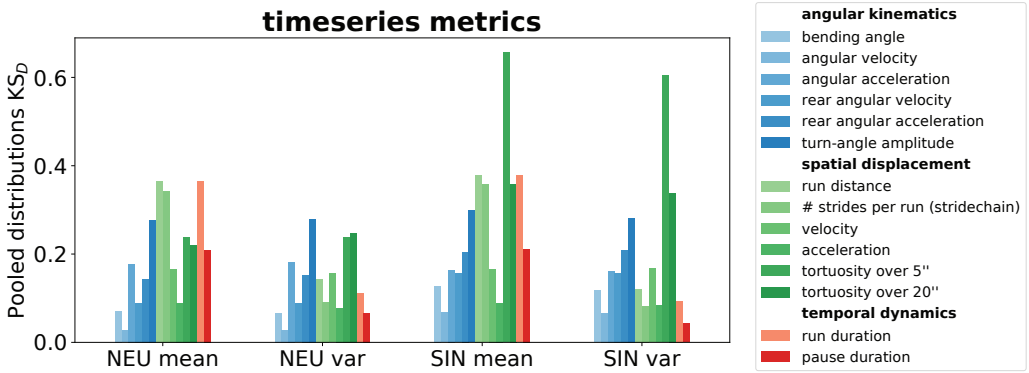

(b)

Fig 13. Average vs group-variability model. Two methods of generating virtual populations are compared and contrasted over each of two locomotory models already calibrated to optimally fit the average individual of the empirical dataset. A group of 50 identical animats is generated for each average model and another group of 50 non-identical animats is generated by sampling the five crawler-related parameters from a multivariate Gaussian distribution. All groups are evaluated against an empirical dataset of 150 larvae freely exploring a stimulus-free Petri-dish for 3 minutes. Model evaluation is carried out over two sets of metrics: (a) single-valued endpoint measurements such as sums, averages, counts or dominant frequencies and (b) step-by-step timeseries spanning the entire experiment duration. In both cases, data is pooled across the entire group and the resulting distribution's Kolmogorov-Smyrnov distance KS_D to the empirical one is computed. For visualization purposes, the metrics are grouped into angular, spatial and temporal domains with distinct color coding.

Metric	NEU mean	NEU var	SIN mean	SIN var
bending angle	0.071	**0.067**	0.128	0.117
angular velocity	**0.027**	0.029	0.067	0.066
angular acceleration	0.178	0.181	0.163	**0.161**
rear angular velocity	0.089	**0.088**	0.156	0.157
rear angular acceleration	**0.143**	0.153	0.205	0.21
turn-angle amplitude	**0.278**	0.279	0.299	0.282
run distance	0.365	0.143	0.38	**0.121**
# strides per run (stridechain)	0.343	0.092	0.357	**0.082**
velocity	0.166	**0.156**	0.166	0.167
acceleration	0.089	**0.078**	0.089	0.084
tortuosity over 5”	0.239	**0.238**	0.657	0.604
tortuosity over 20”	**0.22**	0.248	0.358	0.339
run duration	0.364	0.111	0.378	**0.094**
pause duration	0.209	0.066	0.211	**0.043**

Table 12. Average vs group-variability model. Exact KS_D values for all timeseries-derived metrics presented in Fig 13. Each row corresponds to a specific metric, and each column to one of the four evaluated model configurations. For each metric the best fitting model is highlighted in bold.

Code availability

Larvaworld is available as a python package, freely distributed under the GNU General Public License 3.0. The latest version can be found at https://pypi.org/project/larvaworld/ and can be installed easily using the pip Python Installer. Code development, new features,reported issues and contributions to the project are hosted at the respective github repository (https://github.com/nawrotlab/larvaworld). Extensive documentation can be found at https://larvaworld.readthedocs.io

Author contributions

- Conceptualization: Panagiotis Sakagiannis, Martin Paul Nawrot
- Data Curation: Tihana Jovanic
- Funding acquisition: Martin Paul Nawrot
- Methodology: Panagiotis Sakagiannis
- Software: Panagiotis Sakagiannis, Hannes Rapp
- Supervision: Martin Paul Nawrot
- Visualization: Panagiotis Sakagiannis
- Writing - Original draft: Panagiotis Sakagiannis
- Writing - Review, Editing: Panagiotis Sakagiannis, Hannes Rapp, Tihana Jovanic, Martin Paul Nawrot

Acknowledgments

This project received funding from the Ministry of Culture and Science of the State of Northrhine Westphalia through the science network 'iBehave' (https://ibehave.nrw/) within the program "Netzwerke 2021". Within the Collaborative Research in Computational Neuroscience (CRCNS) program (project 'DrosoExpect') funding has been received from the German Federal Ministry of Education and Research (BMBF, grant no. 01GQ2103A to M.N.), the French National Research Agency (ANR, grant no. ANR-21-NEUC-0002 to T.J.), and the United States Department of Energy (DOE, grant no. SC0021922 to B.H.S.). Additional funding has been received by T.J. from ANR-NEUROMOD (grant no. ANR-22-CE37-0027), Fédération pour la recherche sur le cerveau (FRC), and Fondation pour la Recherche Médicale: Équipe FRM EQU202303016317. We would like to thank Nick del Grosso of the iBOTS: the iBehave Open Technology Support Platform for his technical support.

References

1. Datta SR, Anderson DJ, Branson K, Perona P, Leifer A. Computational Neuroethology : A Call to Action. Neuron. 2019;104(1):11-24. Publisher: Elsevier Inc. Available from: https://doi.org/10.1016/j.neuron.2019.09.038.

2. Arreguit J, Ramalingasetty ST, Ijspeert A. FARMS: Framework for Animal and Robot Modeling and Simulation. bioRxiv. 2023 Sep. Available from: http://biorxiv.org/lookup/doi/10.1101/2023.09.25.559130.

3. Zador A, Escola S, Richards B, Ölveczky B, Bengio Y, Boahen K, et al. Catalyzing next-generation Artificial Intelligence through NeuroAI. Nature Communications. 2023 Mar;14(1):1597. Available from: https://www.nature.com/articles/s41467-023-37180-x.

4. Sakagiannis P, Jürgensen AM, Nawrot MP. A behavioral architecture for realistic simulations of *Drosophila* larva locomotion and foraging. eLife. 2025 Apr. Available from: https://doi.org/10.7554/eLife.104262.1.sa3.

5. Wilson SP, Prescott TJ. Scaffolding layered control architectures through constraint closure: insights into brain evolution and development. Philosophical Transactions of the Royal Society B: Biological Sciences. 2022 Feb;377(1844). Publisher: The Royal Society. Available from: https://royalsocietypublishing.org/doi/10.1098/rstb.2020.0519.

6. Sims DW, Humphries NE, Hu N, Medan V, Berni J. Optimal searching behaviour generated intrinsically by the central pattern generator for locomotion. eLife. 2019;8:1-31. Available from: https://elifesciences.org/articles/50316.

7. Sun X, Yue S, Mangan M. A decentralised neural model explaining optimal integration of navigational strategies in insects. eLife. 2020;9:1-30.

8. Cardoso RC, Ferrando A. A Review of Agent-Based Programming for Multi-Agent Systems. Computers. 2021 Jan;10(2):16. Available from: https://www.mdpi.com/2073-431X/10/2/16.

9. An L, Grimm V, Sullivan A, Turner II BL, Malleson N, Heppenstall A, et al. Challenges, tasks, and opportunities in modeling agent-based complex systems. Ecological Modelling. 2021 Oct;457:109685. Available from: https://linkinghub.elsevier.com/retrieve/pii/S030438002100243X.

10. Foramitti J. AgentPy: A package for agent-based modeling in Python. Journal of Open Source Software. 2021 Jun;6(62):3065. Available from: https://joss.theoj.org/papers/10.21105/joss.03065.

11. Kooijman SaLM. Dynamic Energy Budget theory for metabolic organisation : Summary of concepts of the third edition. Water. 2010;365:68. ISBN: 9780521131919. Available from: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2981979&tool=pmcentrez&rendertype=abstract.

12. Kaun KR, Riedl CAL, Chakaborty-Chatterjee M, Belay AT, Douglas SJ, Gibbs AG, et al. Natural variation in food acquisition mediated via a Drosophila cGMP-dependent protein kinase. Journal of Experimental Biology. 2007;210(20):3547-58.

13. Thoener J, Schleyer M. Locomotion of naive Drosophila larvae. G-Node; 2021. Available from: https://doi.gin.g-node.org/10.12751/g-node.5e1ifd.

14. Gomez-Marin A, Louis M. Active sensation during orientation behavior in the Drosophila larva: More sense than luck. Current Opinion in Neurobiology. 2012;22(2):208-15. Publisher: Elsevier Ltd. Available from: http://dx.doi.org/10.1016/j.conb.2011.11.008.

15. Sakagiannis PP. *Larvaworld* platform generated video: Wind of variable speed; 2025. Available from: https://computational-systems-neuroscience.de/wp-content/uploads/2025/03/wind_speed.mp4.

16. Sakagiannis PP. *Larvaworld* platform generated video: Wind-affected odorscape; 2025. Available from: https://computational-systems-neuroscience.de/wp-content/uploads/2025/03/windNodorscape.mp4.

17. Sakagiannis PP. *Larvaworld* platform generated video: Wind of variable direction; 2025. Available from: https://computational-systems-neuroscience.de/wp-content/uploads/2025/03/wind_direction.mp4.

18. Sakagiannis PP. *Larvaworld* platform generated video: Single air-puff of variable direction; 2025. Available from: https://computational-systems-neuroscience.de/wp-content/uploads/2025/03/single_air-puffs.mp4.

19. Sakagiannis PP. *Larvaworld* platform generated video: Repetitive air-puffs; 2025. Available from: https://computational-systems-neuroscience.de/wp-content/uploads/2025/03/repetitive_air-puffs.mp4.

20. Schumann I, Triphan T. The PEDtracker: An Automatic Staging Approach for Drosophila melanogaster Larvae. Frontiers in Behavioral Neuroscience. 2020;14.

21. Wosniack ME, Festa D, Hu N, Gjorgjieva J, Berni J. Adaptation of Drosophila larva foraging in response to changes in food resources. eLife. 2022 Dec;11:e75826. Available from: https://elifesciences.org/articles/75826.

22. Paisios E, Rjosk A, Pamir E, Schleyer M. Common microbehavioral "footprint" of two distinct classes of conditioned aversion. Learning and Memory. 2017;24(5):191-8.

23. de Tredern E, Manceau D, Blanc A, Sakagiannis P, Barre C, Sus V, et al. Feeding-state dependent neuropeptidergic modulation of reciprocally interconnected inhibitory neurons biases sensorimotor decisions in Drosophila. bioRxiv. 2024 Jan:2023.12.26.573306. Available from: http://biorxiv.org/content/early/2024/04/18/2023.12.26.573306.abstract.

24. Kafle T, Grub M, Sakagiannis P, Nawrot MP, Arguello JR. Evolution of temperature preference behaviour among Drosophila larvae. iScience. 2025 May:112809. Available from: https://linkinghub.elsevier.com/retrieve/pii/S2589004225010703.

25. Wystrach A, Lagogiannis K, Webb B. Continuous lateral oscillations as a core mechanism for taxis in Drosophila larvae. eLife. 2016;5.

26. Sakagiannis P, Aguilera M, Nawrot MP. A Plausible Mechanism for Drosophila Larva Intermittent Behavior. In: Vouloutsi V, Mura A, Tauber F, Speck T, Prescott TJ, Verschure PFMJ, editors. Biomimetic and Biohybrid Systems. Cham: Springer International Publishing; 2020. p. 288-99.

27. Jürgensen AM, Sakagiannis P, Schleyer M, Gerber B, Nawrot MP. Prediction error drives associative learning and conditioned behavior in a spiking model of Drosophila larva. iScience. 2024 Jan;27(1):108640. Available from: https://linkinghub.elsevier.com/retrieve/pii/S2589004223027177.

28. Sakagiannis PP. *Larvaworld* platform generated video: Feeding-state dependent dispersal of Drosophila larvae; 2025. Available from: https://computational-systems-neuroscience.de/wp-content/uploads/2025/03/3conditions.mp4.

29. Tastekin I, Khandelwal A, Tadres D, Fessner ND, Truman JW, Zlatic M, et al. Sensorimotor pathway controlling stopping behavior during chemotaxis in the Drosophila melanogaster larva. eLife. 2018;7:1-38.

30. Schulze A, Gomez-Marin A, Rajendran VG, Lott G, Musy M, Ahammad P, et al. Dynamical feature extraction at the sensory periphery guides chemotaxis. eLife. 2015;4(JUNE):1-52.

6 Applications in collaborative studies

Beyond the core studies presented in the main chapters of this thesis, additional scientific contributions were made to several international collaborative research projects. These interdisciplinary efforts explored a range of topics – including the neural mechanisms of associative learning, the evolution of thermotactic behavior, and the influence of internal physiological states on behavioral decisions – using approaches that combined behavioral data analysis with computational modeling and agent-based simulations.

The nature of these contributions varied across projects and included diverse combinations of modeling, simulation, data processing, code development, and scientific writing. Collectively, they underscore two key points: first, how the Behavioral Architecture can incorporate additional neuroscientific models as modular components or extend into new sensory modalities and their respective behavioral domains; and second, how *Larvaworld* serves as a versatile tool, supporting both modelers in simulation-based work and experimentalists in behavioral data analysis.

The following sections briefly present each collaborative project, highlighting its scientific context and the specific contributions made. Each section is titled after the corresponding scientific publication and includes its metadata, a description of the personal contribution, and a summary of the study – since the full papers are not included here as in previous chapters. The broader relevance of these projects and their place within the thesis framework are discussed in Part III.

6.1 Prediction error drives associative learning and conditioned behavior in a spiking model of *Drosophila* larva

Authors: Anna-Maria Jürgensen, Panagiotis Sakagiannis, Michael Schleyer, Bertram Gerber, Martin Paul Nawrot
Journal: iScience
Date: December 23, 2023
DOI: 10.1016/j.isci.2023.108640

Author Contributions

Conceptualization was carried out by A.-M.J. and M.P.N.; computational experiments and data analyses were performed by A.-M.J. and **P.S.**; animal data were provided by M.S. and B.G.; the original draft was written by A.-M.J., **P.S.**, and M.P.N.; review and editing were carried out by A.-M.J., **P.S.**, M.S., B.G., and M.P.N.; funding acquisition was provided by M.P.N. and B.G.; and supervision was conducted by M.P.N.

Personal involvement In this project, contributions were made to computational modeling, simulation and data analysis, alongside manuscript preparation, review and editing. Regarding the former, agent-based odor preference simulations were designed, conducted and analyzed using the *Larvaworld* platform, by interfacing the mushroom body model to the larva chemotaxis model.

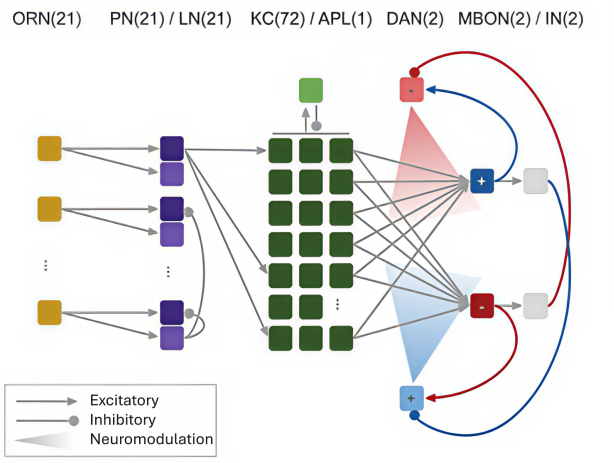

Figure 6: A mechanistic model of the *Drosophila* larva mushroom body including sensory input from olfactory receptor neurons (ORNs), lateral and feedback inhibition via local interneurons and the APL neuron, and dopaminergic neuromodulation by reward- and punishment-associated DANs. Learned odor associations modulate the activity of approach- or avoidance-mediating MBONs that feedback to DANs. Reproduced from Jürgensen et al., 2024.

Summary

This study presents a biologically grounded, spiking neural network model of the *Drosophila* larva mushroom body (MB) that elucidates how prediction error (PE) drives associative learning and conditioned behavior. Anchored in the principles of classical conditioning and inspired by the Rescorla-Wagner (RW) model, the research advances our understanding of how reinforcement learning is mechanistically implemented in the insect brain, particularly through dopaminergic neuron (DAN) activity and synaptic plasticity in the MB circuitry.

The authors model a circuit where DANs receive both external reinforcement signals and feedback from MB output neurons (MBONs), allowing them to compute a PE defined as the difference between received and expected reinforcement. This is achieved via an anatomically plausible feedback loop in which excitatory and inhibitory inputs from MBONs to DANs encode the animal's reinforcement expectations. Synaptic plasticity at the Kenyon cell (KC) to MBON synapses is governed by a two-factor rule: the eligibility trace from KC spiking and the DAN-mediated reward signal. A third component, homeostatic synaptic regulation, is introduced to account for memory decay in the absence of reinforcement, enabling the model to reproduce extinction phenomena (Figure 6).

The model simulates odor-reward associative learning, demonstrating key properties of PE-driven learning: saturating acquisition curves, dependence on reward intensity and timing, and extinction through omission of the unconditioned stimulus (US). Importantly, when feedback from MBONs to DANs is disabled, learning curves fail to saturate, underscoring the necessity of feedback for accurate PE computation. Additionally, the model reproduces behavioral generalization to similar odors and time-dependent learning in trace conditioning paradigms, where the conditioned stimulus (CS) precedes the US by varying intervals. The presence of an odor-evoked eligibility trace enables effective learning even when CS and US are not temporally overlapping, consistent with larval behavioral data.

To validate the model against empirical data, paired and unpaired odor-reward conditioning protocols are simulated, paralleling behavioral experiments in real larvae. In paired conditions, the model shows robust odor preference that saturates with training duration, while unpaired conditions result in a decay of initial preference – a pattern also observed in biological counterparts. When pre-training is included to reflect innate or previously learned odor preferences, the model accurately reproduces the decline in preference during unpaired trials, likely due to violation of reward expectations.

Crucially, the model is integrated with a realistic simulation of larval locomotion, enabling direct comparison between computational outputs and group-level behavioral metrics such as preference and performance indices. The behavioral simulations align closely with experimental data, demonstrating that the modeled PE mechanism accounts for observed learning dynamics in larval *Drosophila* (Figure 7).

Compared to prior rate-based models and non-PE models of insect learning, this is

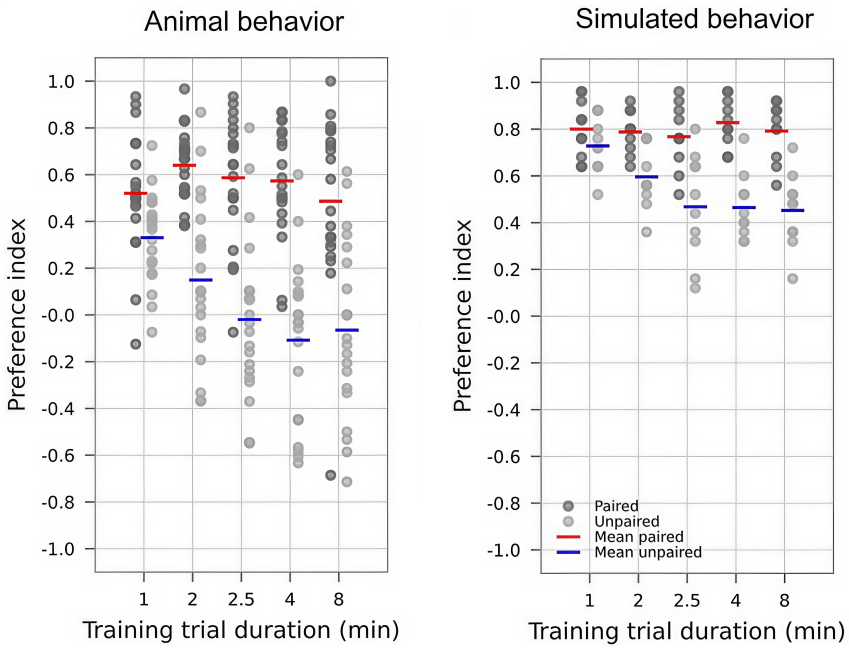

Figure 7: Replicating odor preference experiments with paired and unpaired training. (a) Experimental preference indices for real animal groups each for paired and unpaired experiments with randomized order of odor and reward. (b) Simulated preference indices based on the protocol in (a). Reproduced from Jürgensen et al., 2024.

the first spiking neural network model to implement PE coding in the MB. It supports a mechanistic instantiation of the RW model using biophysically plausible dynamics, where DAN spikes act as modulatory signals for learning and extinction. The model also accommodates individual variability via randomized PN-KC connectivity and stochastic sensory inputs, capturing the diversity in larval learning behavior.

In summary, the study proposes and validates a spiking circuit model in which PE, computed via MBON>DAN feedback, underlies associative learning in *Drosophila* larvae. This model not only reproduces a wide range of experimental findings but also offers testable predictions about circuit mechanisms and learning dynamics, providing a robust computational framework for studying reward-driven learning in compact neural systems.

6.2 Evolution of temperature preference behaviour among *Drosophila* larvae

Authors: Tane Kafle, Manuel Grub, Panagiotis Sakagiannis, Martin Paul Nawrot, J. Roman Arguello
Journal: iScience
Date: May 31, 2025
DOI: 10.1016/j.isci.2025.112809

Author Contributions

Conceptualization was carried out by T.K. and J.R.A.; methodology was developed by T.K., **P.S.**, and J.R.A.; software was implemented by T.K. and **P.S.**; formal analysis was conducted by T.K. and J.R.A.; investigation was performed by T.K., M.G., and J.R.A.; data curation was carried out by T.K. and J.R.A.; the original draft was written by T.K.; funding acquisition was provided by J.R.A.; and review and editing were carried out by T.K., **P.S.**, M.P.N., and J.R.A.

Personal involvement This project involved the design and implementation of agent-based thermotaxis simulations within the *Larvaworld* platform, alongside participation to methodological planning and manuscript preparation. Software development included extending the default locomotory model to accomodate thermosensory input, fitting of species-specific larval models, and designing a virtual thermoscape arena.

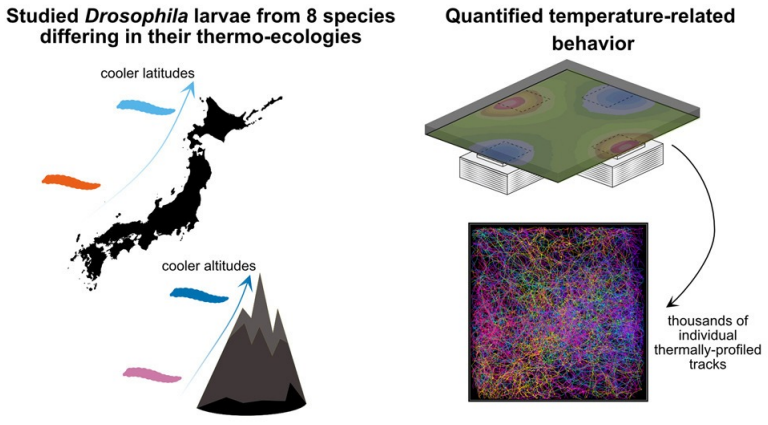

Figure 8: Thermo-related behavioral evolution. Left: Variation in latitude and altitude distribution correlates with temperature preference across *Drosophila* species. Right: Schematic of the behavioral arena and the overlaid thermal gradient. Adapted from Kafle et al., 2025.

Summary

The study investigates the evolution of temperature preference behavior among larvae of eight closely related *Drosophila* species from ecologically diverse habitats. Using a controlled thermal gradient assay and high-resolution behavioral tracking, the researchers quantified species-specific thermotactic behaviors, identifying both coarse and fine-scale interspecific differences (Figure 8). The results show that temperature preference traits – namely the optimal temperature (T_{opt}) and the breadth of preferred temperatures ($T_{breadth}$) – have evolved recurrently and rapidly across the *Drosophila* phylogeny, reflecting adaptation to local climatic conditions.

Species from the Oriental subgroup, including *D. lutescens, D. takahashii, D. pseudotakahashii,* and *D. suzukii*, consistently displayed a stronger preference for cooler temperatures compared to species from the *D. melanogaster* subgroup. Within both clades, species pairs occupying different thermal habitats exhibited significant divergence in thermotactic behavior. For example, *D. lutescens* and *D. santomea*, both native to cooler environments, showed significantly lower T_{opt} values than their respective sister species, *D. takahashii* and *D. yakuba*. These behavioral differences are consistent with previously documented thermotolerance disparities and support the hypothesis of adaptive divergence in thermal preference.

In addition to aggregate temperature occupancy, the study analyzed fine-scale navigational metrics – velocity, tortuosity, and head sweep size – which provide mechanistic insight into larval thermotaxis. Larvae from cooler-preferring species not only preferred lower temperature zones but also altered their movement dynamics accordingly: they exhibited increased linearity and velocity in preferred (cooler) zones and more tortuous, exploratory movements in warmer ones. These differences were especially pronounced in *D. lutescens* and *D. santomea*, further reinforcing the interpretation of behavioral specialization to local thermal environments.

To explore the underlying neural mechanisms of these behaviors, the study used agent-based simulations parameterized by empirical data. These simulations were based on a cross-inhibition model of the larval thermotaxis circuit involving distinct cooling- and warming-responsive neurons (CCs and WCs) (Figure 9). By fitting model parameters – specifically, homeostatic set point and slope (which reflects sensitivity) – to each species' behavioral data, the analysis revealed that species differences in thermal preference are primarily driven by shifts in the balance between CC and WC circuit outputs rather than changes in overall thermosensitivity. This suggests that the evolutionary changes observed are not due to heightened detection thresholds but to altered central processing or weighting of thermal signals.

Overall, the study demonstrates that temperature preference behaviors in *Drosophila* larvae evolve rapidly and repeatedly across short evolutionary timescales. These shifts are accompanied by changes in navigational strategies and are mechanistically linked to alterations in the neural circuitry governing thermotaxis. The results offer a detailed

picture of how behavioral thermoregulation can evolve in small ectotherms and highlight the utility of combining behavioral assays, phylogenetic comparisons, and neural modeling to dissect the evolutionary dynamics of complex behaviors.

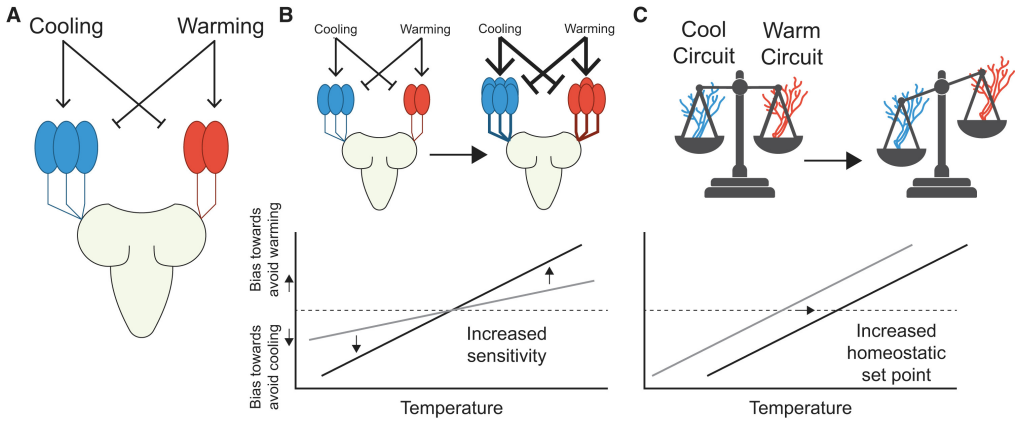

Figure 9: Modeling parameters of larval cooling and warming circuits. (A) Simplified schematic of three CCs and two WCs in the dorsal organ: warming activates WCs and inhibits CCs; cooling does the reverse. Mutual inhibition shapes outputs to brain centers guiding temperature preference. (B) The slope parameter controls temperature sensitivity: steeper slopes produce stronger avoidance responses to smaller temperature changes. This may reflect peripheral factors like cell number (as illustrated) or differences in thermosensor function. Graphically, slope determines how strongly behavior deviates from the dashed line representing circuit balance. (C) The homeostatic set point is where cool- and warm-avoidance circuits are balanced – marked by the intersection of the response curve and the dashed line. Shifting this balance alters the larva's preferred temperature. Reproduced from Kafle et al., 2025.

6.3 Feeding-state dependent neuropeptidergic modulation of reciprocally interconnected inhibitory neurons biases sensorimotor decisions in *Drosophila*

Authors: Eloïse de Tredern, Dylan Manceau, Alexandre Blanc, Panagiotis Sakagiannis, Chloe Barre, Victoria Sus, Francesca Viscido, Md Amit Hasan, Sandra Autran, Martin Nawrot, Jean-Baptiste Masson, Tihana Jovanic
Journal: Nature Communications (in press)
Date: 27 June, 2025 (accepted)
DOI: 10.21203/rs.3.rs-4018128/v1 (preprint in bioRxiv)

Author Contributions

Contributions to the publication were as follows: Behavioral and physiology experiments, including calcium imaging, were performed by E.d.T. and D.M.; immunohistochemistry was carried out by M.A.H. and S.A.; modeling was conducted by A.B.; data analysis was performed by **P.S.**; behavioral classification and statistical analysis were conducted by C.B.; the original draft was written by D.M.; figures, edits, and revisions were contributed by E.d.T., D.M., J.-B.M., and T.J.; methodology and supervision were provided by J.-B.M.; funding acquisition was secured by J.-B.M., M.P.N., and T.J.; and conceptualization, supervision, and project administration were led by T.J.

Personal involvement For this project, behavioral tracking data of larval locomotion were analyzed using the *Larvaworld* platform. The analysis enabled extraction of kinematic features and revealed locomotory differences between nutritional states and hydration conditions across genetically distinct strains, as shown in Figure 10.

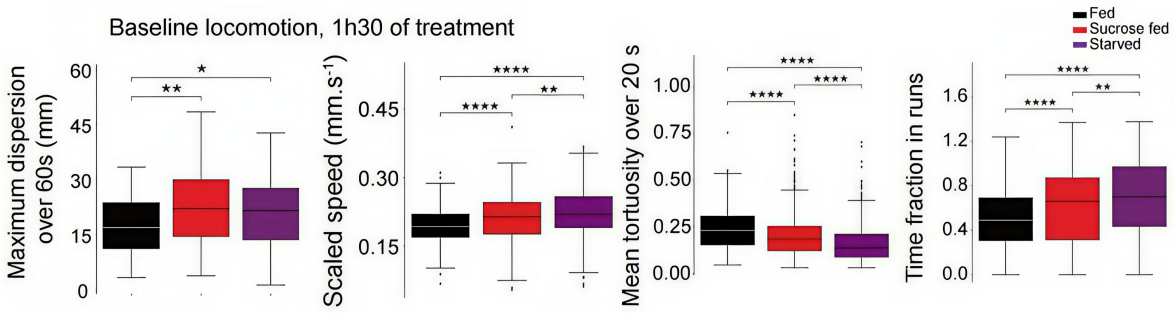

Figure 10: Impact of nutritional state on locomotion. After 90 min of sucrose feeding or complete starvation, larval dispersion, time spent crawling and speed are significantly increased, while the tortuosity of the trajectory is decreased compared to larvae fed on a standard diet. Reproduced from de Tredern et al., 2024.

Summary

This study investigates how the internal physiological state – specifically short-term alterations in feeding condition – modulates sensorimotor decision-making in *Drosophila melanogaster* larvae in response to non-feeding-related stimuli. Focusing on the behavioral responses to an aversive mechanical cue (air-puff), the research uncovers how competing actions – protective hunching versus exploratory head casting – are dynamically biased by the feeding state via neuropeptidergic modulation of reciprocally inhibitory neural circuits.

A 90-minute feeding manipulation was sufficient to shift locomotor behavior and stimulus-driven decisions. Both sucrose-fed and starved larvae displayed increased locomotion and decreased tortuosity, consistent with an elevated exploratory drive (Figure 10). These changes reversed upon brief refeeding or rehydration, demonstrating the flexibility of the motivational feedback. Importantly, sucrose-fed larvae reduced their consumption of sucrose in subsequent assays, implicating nutrient sensing over taste aversion in behavioral modulation.

Behavioral responses to the air-puff included five mutually exclusive actions, with hunching and head casting being the primary ones studied. Under normal conditions, a balance exists between these responses. However, larvae subjected to starvation or sucrose-only diets showed a shift towards head casting, indicating a preference for exploratory behavior. This behavioral bias emerged without significant changes in the calcium responses of the mechanosensory chordotonal neurons, suggesting that modulation occurs downstream in the circuit.

Calcium imaging and modeling revealed that the activity of two reciprocally inhibitory interneurons, Griddle-2 (promoting hunching) and Handle-b (promoting bending/head casting), is differentially modulated by feeding state. Sucrose-fed and starved larvae exhibited reduced Griddle-2 activity and enhanced Handle-b responses, thereby biasing circuit output toward head casting. These interneurons are not only reciprocally connected but also influence the activity of downstream projection neurons (Basin-1 and Basin-2), which gate the behavioral outcome. Specifically, Basin-2 activity – associated with bending – was increased in sucrose-fed larvae in a manner dependent on the activity of Handle-b (Figure 11).

The study identified neuromodulatory inputs that mediate this feeding state-dependent tuning. One descending neuron expressing neuropeptide F (NPF), the fly homolog of vertebrate NPY, synapses directly onto Handle-b and increases its activity in sucrose-fed and starved larvae. Silencing this NPF neuron or knocking down its receptor NPFR1 in Handle-b abolished the feeding-state-induced bias, confirming its functional role. In contrast, Griddle-2 did not receive direct NPF input or express NPFR1.

Instead, both Griddle-2 and Handle-b express the receptor for short Neuropeptide F (sNPF), another NPY homolog. Knockdown of the sNPF receptor sNPFR1 in either neuron mimicked the sucrose-induced behavioral phenotype: reduced hunching and increased head casting. Functional imaging revealed that sNPF has an excitatory effect on Griddle-2

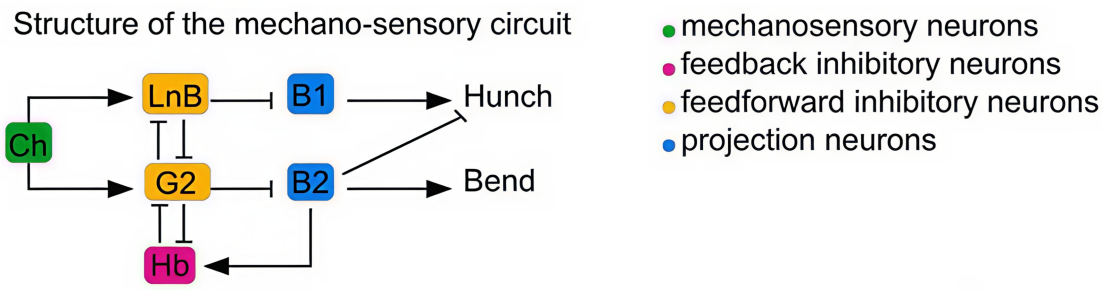

Figure 11: Schematic of the reconstructed Basin mechanosensory circuit. Edge width increases with number of synapses. Sharp arrowheads are excitatory, squared inhibitory. Reproduced from de Tredern et al., 2024.

and an inhibitory effect on Handle-b, further supporting a role for sNPF in the observed bidirectional modulation.

This work elucidates a dual neuromodulatory mechanism – mediated by NPF and sNPF – that enables state-dependent reconfiguration of behavioral circuits through targeted modulation of inhibitory neurons. The modulation does not alter sensory neuron responsiveness but instead biases the competition between motor programs downstream, providing a mechanistic basis for flexible action selection under varying internal conditions. The identified motifs – reciprocal inhibition and feedback disinhibition – allow for context-sensitive control of escape behavior, illustrating how internal state signals like hunger or protein deficiency can reshape decision-making even in non-feeding contexts.

Part III
Discussion

This final part of the thesis reflects on the research contributions presented in Part II and situates them within a broader conceptual and methodological context. As a first step in this discussion, we revisit the three objectives outlined in section 2.3, synthesizing the findings from individual studies and connecting them to the overarching aim of developing a holistic, mechanistic framework for behavioral modeling.

The first objective is the design of a framework for behavioral modeling that can accommodate partial models and integrate neurobehavioral and metabolic components. To set the stage, Figure 12 offers a synthetic view of this framework. It brings together two key components: a layered behavioral architecture (BA), introduced in Chapter 4, and briefly revisited below in section 7.2, and a Dynamic Energy Budget (DEB) model of metabolic regulation, which will be presented in Chapter 8 as the main future direction of the thesis. Although these components are developed separately, their integration here adheres to a synthetic agent design capable of reconciling behavior and physiology within a unified control scheme.

The second objective calls for a software implementation of the BA-DEB framework, as the backbone of a virtual agent's internal structure. Indeed Chapter 5 introduces *Larvaworld*, an open-source behavioral modeling and analysis platform, along with its functionalities – agent-based modeling and simulations, hypothesis testing, data analysis, model evaluation and optimization – and applications in scientific studies. Chapter 9 complements this discussion by focusing on core principles prioritized in its design that highlight its potential as a tool for behavioral research.

The third objective concerns the generation and validation of plausible mechanistic hypotheses on distinct domains of *Drosophila* larval behavior by analyzing empirical data and formalizing observations into computational models. These studies were presented in full or in summary in Chapters 3, 4 and 6. Chapter 7 will briefly revisit them under a mechanistic lens and situate them within the broader framework.

To avoid redundancy, this part will revisit previous material selectively, adopting an integrative perspective on the overall thesis. Notably, pending projects and future directions of this work will not be assembled in a dedicated section but rather be referenced in context where deemed appropriate. The concluding Chapter 10 adopts a reflective Q&A format, addressing questions about the scientific, methodological, and broader implications of the thesis.

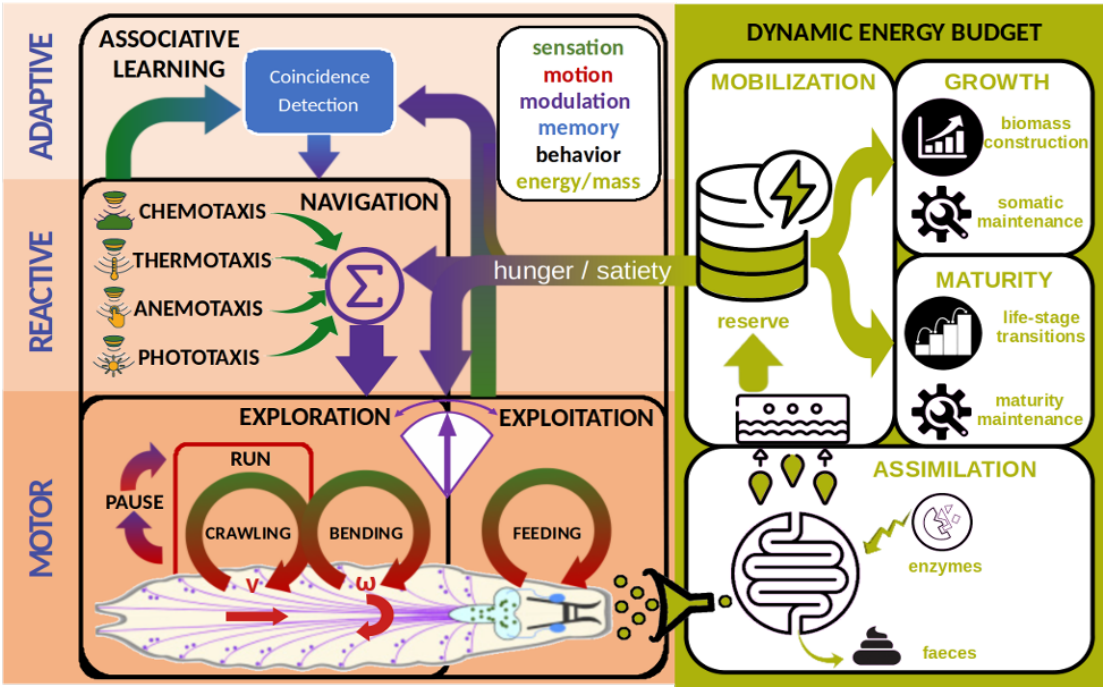

Figure 12: . The BA-DEB framework interfaces a layered behavioral control architecture (BA) with a Dynamic Energy Budget (DEB) metabolic model. **Behavioral Architecture (BA)**: The left portion of the diagram depicts a layered control architecture comprising motor, reactive, and adaptive layers. At the motor layer, the larval body is represented along with its translational velocity v and angular velocity ω in red. Circular arrows denote embodied sensorimotor couplings of oscillatory nature. Run–pause intermittency is illustrated as an alternation between activity and inactivity states. Interference between crawling and bending is indicated by the intersection of the corresponding behavioral modules. An arrow-shaped pendulum symbolizes the dynamic balance between exploration and exploitation modes (EEB). In the reactive layer, multimodal sensory inputs converge into a common modulation pathway ($\sum$), guiding navigational responses to sensory landscapes. The adaptive layer incorporates a coincidence detection module, recurrently connected to the reactive layer, which supports associative learning. Colored arrows indicate pathways for sensation (green), locomotion (red), modulation (purple), and learning (blue). Behavioral modes are shown in black, enclosed in nested rectangles, reflecting their hierarchical organization. **Dynamic Energy Budget (DEB)**: The right portion of the diagram depicts a DEB model : consumed food is converted into energy, which is then allocated to growth, maintenance, and maturation. Rectangular frames delineate the four interconnected energy fluxes defined in DEB theory: ingestion and assimilation, reserve mobilization, biomass construction and maintenance, and maturation. The assimilation process, catalyzed by gut enzymes, transforms food into energy, stored as reserve. Mobilized reserve fuels both growth and maturation, where growth entails biomass construction and its metabolic upkeep, and maturation governs developmental transitions across life stages. **BA↔DEB**: A hunger/satiety signal bridges the metabolic and behavioral domains, linking energy reserve to behavioral decision-making. This coupling modulates the EEB pendulum and also affects sensory responsiveness and memory processes. The architecture thus enables bidirectional interactions between physiological state and behavioral control, supporting context-sensitive, energy-aware behavior.

7 Placing mechanisms within a unified framework

This chapter brings into focus the concrete scientific contributions of the thesis by examining the standalone mechanistic models developed throughout. Initially treating them in isolation, it then situates each model in relation to the overarching BA introduced in Chapter 4. The goal is to show how specific hypotheses about larval behavior, informed by empirical observations and implemented through computational modeling, can be framed as modular components that both draw from and inform the architecture as a whole. In doing so, the chapter highlights the interplay between mechanism-specific insight and the broader challenge of building integrative models of behavior.

Initially, the individual studies are presented in a homogeneous format, each framed as a standalone instance of the scientific method with emphasis on the proposed mechanisms themselves. This is followed by briefly revisiting the BA, considering how these mechanisms can be positioned on its layers.

7.1 Generating and validating mechanistic hypotheses

Scientific understanding advances through the formulation and validation of mechanistic hypotheses that link observed behavior to underlying processes. In the case of larval behavior, these mechanisms may span multiple substrates and timescales, from neural dynamics and body biomechanics to state-dependent modulation. The studies presented in this section follow a common pathway to scientific discovery: observations derived from detailed behavioral data give rise to mechanistic hypotheses informed by neuroscientific insight; these are then formalized into computational models, tested against experimental evidence, and iteratively refined to enhance their explanatory and predictive power. While differing in modeling approaches and degree of supportive empirical evidence, all contributions adhere to the same scientific pipeline and converge on a shared goal: to identify and validate plausible mechanisms underlying larval behavior.

In what follows, we briefly review each study, focusing on the mechanistic hypothesis it proposes—how it emerges, how it is modeled and supported through simulations, and how it might ultimately be tested empirically.

Behavioral intermittency as emergent avalanche dynamics (Chapter 3)

Freely moving larvae exhibit a spontaneous alternation between activity and quiescence. This intermittency displays statistical signatures such as heavy-tailed duration distributions for both behavioral states, prompting the hypothesis that it may originate from internal network dynamics. Drawing inspiration from avalanche[12] phenomena in neural

[12]Neural avalanches refer to cascades of neural activity that follow a power-law distribution in size and duration, indicating that the system operates near a critical state between order and disorder. They are typically observed in spontaneous activity across cortical and subcortical networks and are thought to reflect optimal conditions for information processing and network adaptability (Beggs and Plenz, 2003).

systems, the proposed model implements a minimal stochastic binary-neuron network in which activity propagates and self-terminates through threshold-bounded interactions, resulting in self-limiting cascades.

This simple network architecture reproduces salient temporal features of larval behavioral transitions, notably the asymmetric distribution of activity and pause durations. The model, although it does not claim to replicate any empirical circuit anatomy, points to a testable mechanistic hypothesis : autonomous critical-like neural activity as a potential driver of behavioral transitions. This hypothesis draws deeper justification from known features of larval neurobiology. Larval crawling is driven by central pattern generators (CPGs) in the ventral nerve cord, involving both excitatory and inhibitory premotor neurons and oscillating independently of sensory feedback (Mantziaris et al., 2020; Pulver et al., 2015). The model posits that transient self-limiting bursts of inhibitory activity affecting these CPGs may underlie the observed behavioral pauses, naturally giving rise to power-law distributed pause durations. Conversely, the cessation of inhibitory input – akin to quiescent periods in the model – enables the reactivation of the CPGs and results in activity bouts – uninterrupted chains of crawling strides (runs) – which are empirically found to follow log-normal distributions. To confirm or reject this hypothesis, inhibitory activity affecting the crawling CPGs should be sought and correlated to behaviorally observed runs and pauses during free exploration.

Crawl–bend interference as biomechanical constraint(Chapter 4)

Kinematic analysis of recorded larval locomotion revealed that the lateral bending amplitude (turns) is systematically reduced during particular phases of the crawling stride cycle. This pointed to the mechanistic proposal that this phenomenon does not depend on neural coordination but on bodily interference: the posterior-to-anterior push phase imposes a biomechanical constraint that limits head bending. Specifically, lateral bending is attenuated during the stride phase interval that the head is anchored to the substrate, and facilitated when the head is lifted, a pattern confirmed by phase-dependent angular velocity attenuation in empirical data.

To test this, a coupled-oscillator model was developed in which oscillatory angular motion is under interference by the peristaltic crawling oscillation. The resulting motion patterns match those seen in experimental trajectories and reproduce the smaller amplitude of weathervaning turns during crawling compared to larger head casts during pauses. The model offers a unified generative account that avoids ad hoc distinctions between turn events, instead attributing them to a shared underlying mechanism modulated by embodied constraints – an example of how embodiment can structure behavior independently of central control.

Thermotaxis evolution through circuit rebalancing (section 6.2)

Behavioral assays across eight *Drosophila* species revealed pronounced differences in larval temperature preferences. Phylogenetic analysis suggested that these preferences evolved recurrently, including a notable "cool shift" in the ancestor of the *D. melanogaster* subgroup and Oriental clade. These patterns point to repeated evolutionary tuning of thermotactic behavior – raising the question: how might thermotaxis evolve through changes in underlying neural mechanisms?

Two mechanistic hypotheses were proposed: divergence driven by altered sensitivity of first-order Cooling Cell (CC) and Warming Cell (WC) receptors, or by reweighting of their antagonistic downstream circuits. Agent-based simulations supported the latter—modulating circuit weights alone reproduced species-specific thermotactic patterns, implying that changes in circuit balance, not receptor sensitivity, underlie behavioral evolution.

This mechanistic hypothesis offers a parsimonious explanation: adaptive shifts in thermotaxis can arise from rebalancing existing pathways. It predicts that species with different thermal preferences will show corresponding differences in CC and WC pathway influence on motor output. Future studies should test this by linking behavioral variation to physiological measures of circuit balance, using tools like calcium imaging or targeted perturbations.

Feedback loops in expectation-driven associative learning (section 6.1)

The mushroom body (MB) has long been established as the neuropile mediating associative learning in insects. Behavioral observations – saturating learning curves when across increasing training durations and intensities of reinforcement – favored the hypothesis that feedback from MB output neurons (MBONs) to reinforcement-signaling dopaminergic neurons (DANs) may implement a form of prediction error[13] coding. It was tested in a spiking MB model incorporating MBON→DAN feedback which reproduced learning dynamics consistent with experimental data, both at its output level and when interfaced with virtual larva odor preference simulations.

Positioning the MB as a site of prediction error computation reframes it as more than an associative relay, aligning with theoretical models of reward learning. The model predicts that disrupting MBON→DAN feedback should impair learning or its sensitivity to reinforcement timing. Targeted manipulations – e.g., optogenetics or silencing – could empirically test this mechanism and refine our understanding of learning in compact neural circuits.

[13] Defined as the difference between expected and received reinforcement, prediction error serves as a central learning signal in both theoretical and empirical accounts of associative learning. See (Schultz et al., 1997) for a foundational neurobiological perspective.

Behavioral modulation via neuropeptidergic feedback (section 6.3)

Larvae exhibit feeding-state dependent changes in locomotory patterns and defensive reactions to aversive mechanical stimuli. This was attributed to a mechanism of differential neuropeptide-mediated (NPF and sNPF) modulation of a set of inhibitory interneurons within the decision-making circuit and evidenced via calcium imaging, neuronal manipulations, behavioral analysis and computational modeling.

Regarding the latter, a previously established model of hunch-bend action selection was extended by incorporating feeding state modulation of the activity of feedforward and feedback inhibitory neurons. Simulated reweighting of these pathways reproduced the state-dependent behavioral shift – reduced hunching and increased head casting – demonstrating how internal states reshape behavior through targeted circuit-level changes.

Each of the mechanisms hypothesized here is implicated in a distinct aspect of behavioral control and can be formulated in a way that supports computational modeling. Beyond capturing each in a standalone model, they collectively lend themselves to integration as modular components within a unified architecture. Their behavior-based classification enables a synthetic approach, allowing interactions across control layers such as motor execution, sensory modulation, internal state regulation, and learning. Together, they provide a coherent basis for constructing a layered behavioral system grounded in mechanistic insight.

7.2 The Behavioral Architecture in a nutshell

In the BA-DEB framework behavior arises from the BA, a layered, modular system designed to explain how organized behavior can emerge from the interaction of structurally distinct but relatively simple subsystems. It follows the subsumption principle from behavior-based robotics, wherein each control layer functions autonomously of higher layers. This preserves decentralized control and allows for simulation-based ablation or lesion experiments by selectively disabling layers to assess functional dependencies.

We briefly describe the BA here to illustrate how it serves as a conceptual canvas for situating the mechanisms discussed in the previous section.

Embodied agent The BA operates on a minimal embodied agent. In this case study, the agent is a bisegmental larval body model (Chapter 4) comprising two interacting segments that generate motion. This abstraction supports kinematic-level control—forward crawling and lateral bending—serving as the structural foundation for motor-level behaviors.

Motor layer The motor layer generates stereotyped behavioral primitives – crawling, bending, and feeding – via distinct, potentially concurrent or exclusive modules. Two additional functionalities enhance its generative capacity, namely behavioral intermittency

(Chapter 3), – rest–activity transitions producing realistic pause–run cycles – and intermodular interference (Chapter 4), capturing the stride cycle's phasic suppression of bending and accounting for headcasts during crawl-pauses and weathervaning during sustained runs. Together, these mechanisms form the intermittent coupled oscillator model, enabling the motor layer to autonomously produce realistic exploratory behavior.

Reactive layer The reactive layer modulates motor output based on sensory input, enabling navigation behaviors such as chemotaxis (Chapter 4), thermotaxis (section 6.2), anemotaxis. These second-order sensorimotor couplings build upon the motor layer's primitives but function independently of the adaptive layer. Sensory modalities – olfaction, thermo-, and mechanosensation – are processed by dedicated transduction modules and integrated into a unified behavior-modulating pathway (Eschbach and Zlatic, 2020; Wystrach et al., 2016). Sensors are spatially distributed along the 2D body contour, providing localized input. The mechanosensory circuit engaged in Hunch-Bend action selection (Figure 11) could be positioned here and act as target of the neuropeptidergic state-dependent modulation studied in section 6.3.

Adaptive layer The adaptive layer enables associative learning via a coincidence detection mechanism, allowing behavior to be shaped by prior experience. In section 6.1 this is implemented through a spiking model of the mushroom body (MB), capable of learning odor–reinforcement associations and modulating downstream reactive outputs. Its recurrent interaction with the reactive layer supports the acquisition of sensory-guided behaviors.

The next chapter will introduce the DEB model of metabolism coupled recursively to the BA and bringing it under homeostatic regulation. The behavior-modulating neuropeptidergic feedback mechanism proposed in section 6.3 could be integrated exactly as such a regulatory signal, targeting distinct loci across all BA layers.

8 Energetically-regulated behavior

Individuals persist in time by dynamically regulating behavior throughout life according to both current and future homeostatic needs. From the perspective of metabolism, foraging behavior can be viewed as a fast-responding, though not always reliable, means of energy gain to be used for self-maintenance, growth and reserve build-up. From the behaving agent's perspective homeostatic regulation imposes slowly-adjusting normative constrains that induce goal-directed behavior in order to allow the individual to advance along its predetermined life-history stages. It follows that models of energetics can help bridge the neuroscientifically and ecologically relevant timescales. DEB theory is the most prominent formal theory of homeostasis and energy utilization capturing the energy state and fluxes of organisms across subsequent life stages (Kooijman, 2010). Although it has already been integrated into ABMs of computational ecology (Martin et al., 2012), its integration within neuroscience is limited.

This chapter completes the previously outlined computational framework by linking a DEB model with the BA as shown in Figure 12. The goal is to enable reciprocal coupling between internal metabolic state and overt foraging behavior. In short the interoceptively accessible level of energy reserve regulates, among others, the exploration–exploitation balance (EEB), while feeding behavior determines nutrient intake and assimilation.

We first explain how behavior is treated in standard DEB theory drawing motivation for establishing the aforementioned feedback loop. We then describe the EEB and the implementation of both the DEB model and feeding behavior, before turning to the BA↔DEB feedback loop. We then outline the remaining steps toward a species-calibrated, metabolically-informed agent and suggest a suitable implementation to showcase its scientific value.

8.1 Why integrate behavior into DEB?

The standard formulation of DEB theory provides a quantitative framework for tracking energy assimilation, allocation, and usage across the life stages of an organism. It describes how assimilated energy is divided into reserves and subsequently mobilized for growth, maintenance, and maturation, assuming a continuous flux-based system governed by ordinary differential equations. However, DEB models do not incorporate explicit mechanisms for behavioral modulation. The processes of food search, ingestion, and handling are treated as an undifferentiated contribution to assimilation, subsumed under the single variable f, the scaled functional response.

In DEB theory, the scaled functional response is defined by the relation $f = X/(K+X)$, where X denotes the substrate (food) density and K is the half-saturation constant. The variable f determines the proportion of the maximum assimilation rate that is realized at a given substrate density. Crucially, this formulation treats f as a static function of environmental availability, without modeling how behavioral strategies may dynamically

adjust the actual ingestion rate or energy expenditure associated with foraging.

As a result, conditions of reduced substrate quality ($f < 1$) are not associated with compensatory changes in behavior. There is no mechanism for increasing exploratory effort, extending feeding duration, or reallocating time between competing behavioral modes. The absence of behavioral feedback in standard DEB implies that the organism's interaction with its energetic environment is one-directional and passive: assimilation is simply reduced in suboptimal conditions, and the consequences propagate through the system without modulation at the behavioral level.

The integration of a behavioral controller into the DEB framework enables the recovery of this missing loop. By embedding a foraging mechanism that alternates between exploratory and exploitative modes – each with distinct energy costs and payoffs – it becomes possible to simulate how an agent might actively adjust its behavior in response to internal metabolic signals. A hunger–satiety signal, derived from reserve density, can modulate the agent's state along the EEB, thus restoring a degree of reactivity and control.

This behavioral–metabolic feedback loop not only enables a more realistic simulation of foraging behavior, but also introduces a minimal form of normativity. Rather than executing fixed action sequences, the agent becomes capable of regulating its own behavioral allocation in order to sustain metabolic viability. The target of this regulation is an ontogenetically specified developmental trajectory – such as completing the larval stage and initiating pupation – which places constraints on growth rates and reserve dynamics. In this view, behavioral norms are not externally imposed but emerge from the organism's ongoing attempt to remain within viable energetic bounds under fluctuating environmental conditions. This stands in sharp contrast to optimal foraging theory, where normativity is effectively reduced to optimality and divorced from the internal regulatory dynamics of the organism.

8.2 The exploration–exploitation balance

Larva foraging consists, at minimum, of an exploratory and a consummatory behavioral mode. The temporal microstructure of feeding behavior has been studied predominantly for the adult fly (Itskov et al., 2014). Single reccurent motions, termed sips in the adult, are organized in uninterrupted sequences (bursts or meals) much like single peristaltic waves are concatenated into runs. Meals are interspersed by non-consummatory locomotion intervals. These behavioral modes are mutually exclusive as they partially recruit the same neuromuscular machinery (head segment and mouth-hooks). The larva therefore alternates constantly between the two while exploiting a nutritious substrate.

In the absence of food, all available time and energy are allocated to exploration. Starvation modulates this behavior, inducing straighter trajectories and increased crawling frequency, which enhances dispersal and the likelihood of encountering new food sources (de Tredern et al., 2024). Previous encounters with food bias exploration toward local search, in a manner dependent on starvation state. In this context, idiothetic cues appear

sufficient to guide recurrent visits to previously rewarding locations, suggesting an underlying path integration mechanism that operates independently of visual or olfactory input (Kim and Dickinson, 2017).

In the presence of food, the mutually exclusive modes of consumption (exploitation) and local search (exploration) are dynamically balanced according to the larva's homeostatic state. Starvation induces a transient phase of increased consumption (hyperphagia) before returning to a baseline EEB (Kaun et al., 2008). Under ad libitum feeding conditions, this baseline EEB is normatively imposed: from hatching through the end of the third-instar feeding stage, the larva consumes precisely the amount of food required to reach a critical body mass after a highly predictable developmental interval. This milestone marks the onset of a non-feeding exploratory phase that culminates in pupation.

8.3 Implementation of feeding behavior

Feeding behavior is implemented in *Larvaworld* as a distinct oscillatory process generating repetitive feeding motions, analogous to crawling strides but occurring at a higher frequency. During a feeding motion the larva remains stationary but consumes food if any is present within a radius – proportional to its body length – around its head. The amount of food ingested per motion is proportional to the larva's volume, following the assumptions of V1-morphy[14]. Feeding behavior is activated only when the agent is located on edible substrate. To determine this, a probabilistic sampling mechanism evaluates substrate properties at the agent's position. Feeding is initiated only if the substrate is identified as ingestible, ensuring that consummatory behavior is contingent on local environmental affordances rather than being hardcoded into the behavioral cycle.

In scenarios where sensory-driven chemotaxis is active, the model supports a sequential organization of navigation and ingestion. The agent follows odor gradients toward nutrient sources, and upon arrival, transitions into feeding behavior. This modular chaining enables the simulation of context-sensitive foraging sequences in which distal sensory cues guide exploration, while local substrate detection triggers exploitation. This mechanism reproduces characteristic patterns of larval foraging under spatially structured nutrient distributions, as shown in Figure 13.

8.4 DEB model

An implementation of the DEB model has been developed to simulate the entire larval stage of *Drosophila melanogaster* under predefined or dynamically changing nutritional conditions. The parameterization shown in Table 2 was manually derived and anticipates

[14] In DEB theory, a V1-morph is an organism whose surface area scales linearly with its volume, typically due to one-dimensional growth (e.g., filamentous or sheet-like organisms). This implies that feeding and maintenance costs both scale with volume rather than surface area, simplifying energetic dynamics. Despite morphological deviations, such scaling often approximates real organisms well in narrow size ranges (Kooijman, 2010).

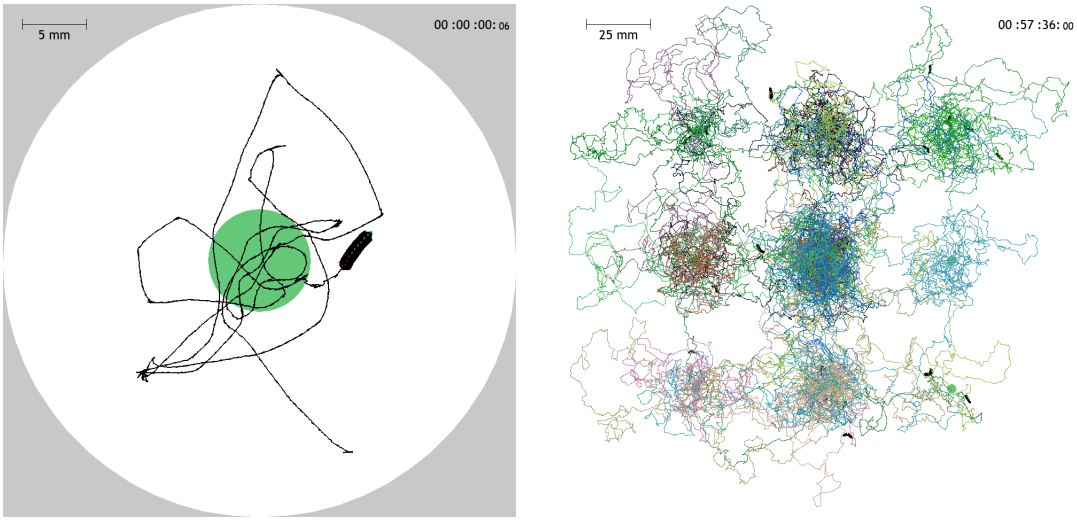

Figure 13: Empirical and simulated foraging on substrates with spatially localized, odor-emitting nutrient sources. Left: Trajectory of a real larva navigating and feeding on a food blob atop an agar dish. The zoomed view centers on the blob, with the larva's body reconstructed as a 12-segment contour for visualization. Right: Simulated trajectories of virtual larvae in a nutrient-poor arena containing a 3×3 grid of food blobs. The agents combine chemotactic navigation with feeding behavior upon detecting edible substrate.

a more formal fitting procedure (see section 8.6). The model tracks all core DEB state variables – energy reserve, structural biomass, and maturity – and derives secondary metrics such as wet mass and body length.

To support integration within an agent-based modeling framework, the original DEB formulation – differential equations solved analytically – was reformulated as a stepwise simulation. This temporal discretization allows the organism's state to be updated incrementally at each simulation tick, enabling synchronization with the behavioral timescale, which typically operates at subsecond resolution (notably the default time unit in DEB is a day, see subsection 9.2.1). Coarser update intervals (e.g., minute-level integration) can also be configured to improve computational efficiency in scenarios where behavioral detail is not critical.

The DEB model can be run in an offline mode, meaning it operates independently of any behavioral control loop. In this mode, food intake is modeled implicitly: the larva is assumed to continuously assimilate nutrients from a substrate of fixed quality, represented as a scalar between 0 and 1. No explicit ingestion or feeding decisions are simulated. The assimilation rate is calculated as if the larva were feeding constantly on a substrate of the specified quality. Under optimal-quality conditions, the model yields growth curves consistent with empirical ad libitum feeding scenarios, capturing stage-specific changes in size, biomass, and energetic state over time. When starvation periods are introduced, the model tracks the depletion of energy reserve, delayed developmental timing, and reduction

Parameter	Description	Value	Units
E_G	volume-specific cost of structure	13022	J/cm^3
$[\dot{p}_M]$	volume-specific somatic maintenance rate	97.1	$J/cm^3 d$
$\dot{v}$	energy conductance	0.018	cm/d
κ	fraction of mobilized reserve allocated to soma	0.456	-
$\dot{k}_J$	maturity maintenance rate coefficient	0.002	d^{-1}
$E_H{}^b$	maturity threshold for hatching	0.032	J
$E_R{}^p$	threshold for the onset of pupation	1.1	J
z	zoom factor	96	-
Δ_M	shape (morph) coefficient	6.4	-
d_W	density of wet mass	0.75	g/cm^3

Table 2: DEB parameter values for *Drosophila melanogaster*. Note that this parameter set is preliminary, as fitting is still ongoing and includes only values relevant up to pupation.

in mass and length while prolonged starvation induces death when the reserve can not support biomass maintenance (Figure 14).

This offline implementation serves as both a baseline for validating closed-loop BA↔DEB simulations and as a testbed for exploring the energetic dynamics of the larval stage under controlled conditions. The corresponding online implementation should be calibrated to yield equivalent growth trajectories when virtual larvae are simulated foraging freely on non-depletable edible substrate from hatching through pupation.

8.5 The BA↔DEB feedback loop

The integration of behavior and metabolism is illustrated in Figure 12 as a composite architecture combining the BA (left) with a DEB model (right). The diagram captures the bidirectional coupling between these two domains: behavior influences metabolism via ingestion (bottom), while metabolic state regulates behavior through a homeostatic signal (top).

This feedback loop comprises two functional arcs. The first arc (BA→DEB) involves a dedicated oscillatory feeding module, added to the BA's motor layer and mutually exclusive with crawling. When active, this consummatory behavior suspends exploration and initiates ingestion, increasing input to the DEB model's assimilation flux.

The second arc (DEB→BA) introduces a hunger/satiety signal, sensitive to fluctuations of the scaled reserve density[15]. The choice of the latter as an interoceptively accessible, metabolic state variable is justified both conceptually – it remains constant during growth

[15]In DEB theory the dimensionless scaled reserve density is defined as $e = [E]/[E_m]$ where $[E]$ is the reserve density – reserve per unit of structural volume – and $[E_m] = \{\dot{p}_{Am}\}/\dot{v}$ is the reserve density capacity – a species-specific parameter that scales with maximum structural length (Kooijman, 2010).

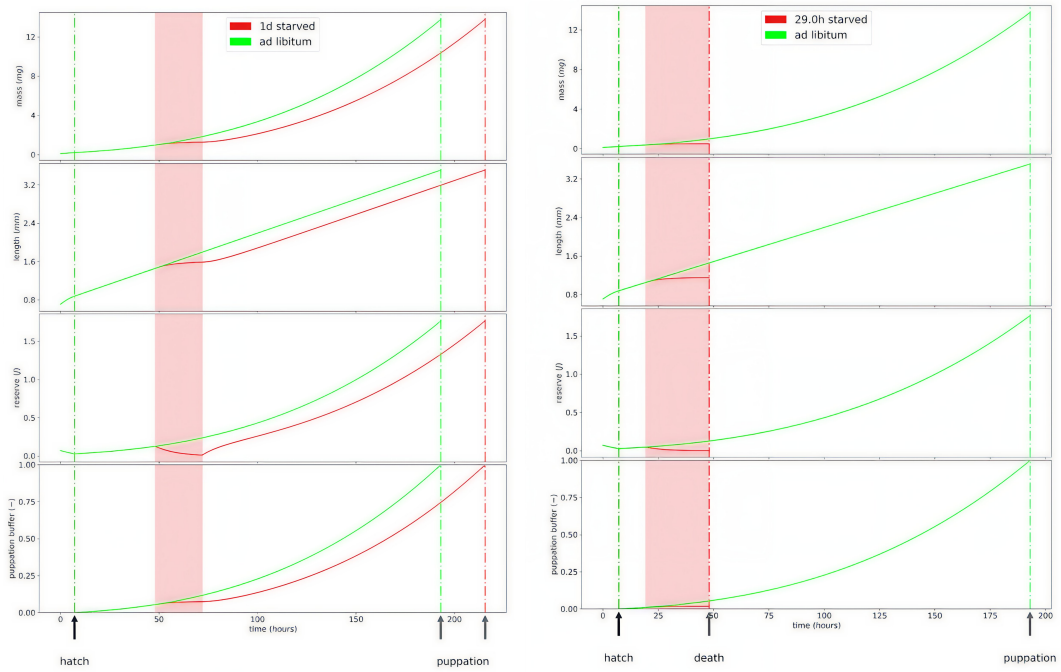

(a) (b)

Figure 14: Offline simulation of the *Drosophila* DEB model for the larval stage. Green and red lines represent development under ad libitum feeding and transient starvation conditions, respectively (with starvation periods indicated by red shading). (a): A one-day starvation episode impacts energy reserves, delays pupation, and extends developmental time. (b): Prolonged starvation results in death once energy reserves are fully depleted.

in constant food availability conditions, according to the weak homeostasis assumption[16] – and computationally – it is a scalar ranging from 0 (death) to 1 (ad libitum feeding). This signal biases the EEB toward exploitation in conditions of food scarcity and also modulates responsiveness to stimuli and memory processes. Through this slow, interoceptive regulation, the DEB model shapes behavioral priorities to maintain homeostasis and unobstructed development.

Together, these mechanisms enable simulation of foraging sequences that adapt to internal energetic demands. By coupling fast-timescale behavioral control with slower metabolic regulation, the architecture captures the reciprocal relationship between action and physiological state throughout the larval life stage.

[16]The weak homeostasis assumption posits that *if food density does not change reserve density*, i.e. the ratio between the amounts of reserve and structure, the latter remains constant even when growth continues; reserve and structure grow in harmony and biomass no longer changes in composition (Kooijman, 2010).

8.6 Steps toward project completion

While the behavioral and metabolic components of the model have been integrated and tested in various configurations, several essential steps remain in order to fulfill the project's goals.

First, a proper species-specific parameterization of the DEB model for *Drosophila melanogaster* is still underway. The necessary empirical observations have been assembled in Table 3 but the optimization process using the DEBtool software package has yet to be concluded. This step is expected to result in an entry for *Drosophila* in the Add-my-Pet database. The process initiates from the nearest available entry – *Aedes aegypti* – and proceeds by parameter fitting for the hex type DEB model, appropriate for holometabolous insects (Llandres et al., 2015). Once complete, the offline simulations should replicate the full growth trajectory of the organism across all developmental stages.

Symbol	Description	Data	Units	Source
a_b	age at birth	0.7	d	Schumann and Triphan, 2020
a_j	time since birth at pupation	7.8	d	Schumann and Triphan, 2020
t_1	duration of instar 1	1.9	d	Schumann and Triphan, 2020
t_2	duration of instar 2	1.6	d	Schumann and Triphan, 2020
t_3	duration of instar 3	4.3	d	Schumann and Triphan, 2020
a_m	life span as imago	27.1	d	Oxenkrug et al., 2011
L_1	length of instar 1	1.8	mm	Schumann and Triphan, 2020
L_2	length of instar 2	2.5	mm	Schumann and Triphan, 2020
L_3	length of instar 3	3.8	mm	Schumann and Triphan, 2020
W_j	wet weight of pupa	15	mg	Ormerod et al., 2017
W_e	wet weight of imago	9.3	mg	Ormerod et al., 2017
E_0	initial energy content	10	mJ	Song et al., 2019

Table 3: Empirical observations for fitting the *Drosophila* DEB model.

Second, further work is needed to calibrate the coupling from metabolism to behavior, specifically the mapping from reserve density to hunger drive and, ultimately, to the EEB. Although the qualitative effect has been demonstrated in simulation – where declining reserves lead to more persistent feeding behavior – this arc of the feedback loop requires quantitative tuning to ensure that behavioral adaptations match ecologically plausible survival strategies under varying environmental conditions.

8.7 Application to a naturalistic scenario

An illustrative application of the BA↔DEB framework involves modeling a well-documented case of natural behavioral variation in *Drosophila* larvae. Wild-type populations exhibit distinct foraging phenotypes, most notably categorized as rovers and sitters, a difference

attributed to allelic variants of the foraging (*for*) gene, which encodes a cGMP-dependent protein kinase (PKG) (Kaun et al., 2007).

When reared under nutrient-rich conditions, rovers (for^r,70%) ingest less food and engage more in exploratory locomotion, while sitters (for^s, 30%) spend more time feeding and exhibit reduced movement. Despite these behavioral differences, both phenotypes grow to similar sizes and display comparable metabolic rates. The key distinguishing factor lies in food absorption efficiency: rovers reportedly absorb approximately 50% of ingested glucose, compared to only 15% in sitters (Kaun et al. (2007)).

Under poor food conditions, both phenotypes increase food intake to a shared maximum, yet rovers maintain their superior absorption efficiency. This divergence is reflected in survival outcomes: rovers show enhanced survivorship under chronic food scarcity, whereas sitters are more resilient in crowded highly-competitive conditions and during short food-deprivation periods (Fitzpatrick et al., 2007; Kaun et al., 2008).

The challenge lies in mapping this interindividual variability in foraging strategy onto the metabolic model. Preliminary simulations have explored parameter variations reflecting differential feeding dynamics and absorption efficiencies. However, a formal identification of the specific DEB parameters that best capture the energetic profiles of rovers and sitters remains an open modeling task. Establishing this mapping would enable comparative simulations under both ad libitum and restricted feeding conditions, with results testable against existing experimental data.

9 Design principles in *Larvaworld*

This chapter addresses core design principles and implemented functionalities of *Larvaworld*, presenting both as expressions of a broader modeling perspective. While the focus is on the software infrastructure and its components, the underlying priorities and conceptual commitments apply more generally to the BA-DEB framework it instantiates. Throughout, the distinction between implementation and framework is maintained only for clarity; methodologically, they are treated as two aspects of the same modeling approach.

In what follows, we first examine two complementary strategies for constructing realistic models of behavior. One emphasizes black-box fitting of observed behavior; the other targets biological plausibility based on modular mechanisms. Rather than treating these approaches as mutually exclusive, it is shown how *Larvaworld* is built to support and integrate both. The next section turns to the methodological choices that structure the platform – including its multiscale temporal logic, modular architecture, and support for closed-loop dynamics.

9.1 A word on realism: Behavioral fidelity vs mechanistic plausibility

In behavioral modeling the term "realism" can be understood in at least two complementary senses. One refers to the external similarity between simulated and biological behavior, often evaluated in terms of indistinguishability by human observers or standardized analytical procedures. The other concerns the internal structure of the model – to what degree its mechanisms reflect known biological organization and function.

Larvaworld is designed to accommodate both perspectives. It supports model optimization to fit observable behavior without presupposing internal structure, and at the same time enables the construction of modular, mechanistically grounded architectures. What follows is a detailed account of how each of these modeling strategies is pursued, and how they are combined to enable gradual, evidence-driven development of increasingly interpretable models.

9.1.1 Realistic behavior: Fitting animats to animals

The first sense of realism emphasizes the capacity of a model to generate behavior that closely resembles that of a real organism. This includes not only its locomotor trajectories but also broader patterns such as responsiveness to stimuli, spatial movement patterns, or temporal organization of activity. A model is considered realistic in this sense if its behavior cannot be distinguished from that of its biological counterpart when assessed under equivalent experimental conditions. This perspective aligns with the notion of an "embodied Turing test" proposed in AI and robotics, where the aim is to generate agents whose behavior appears natural and integrated within their environment (Zador et al.,

2023).

Achieving such realistic behavior does not require that the model reflect the biological system's internal structure, as long as the behavior matches. This approach favors flexibility and context-specific calibration, treating the model as a black-box that maps inputs to outputs without explicit assumptions about internal mechanisms. Behavioral fit is obtained through parameter tuning, typically using iterative procedures guided by empirical data.

Larvaworld accommodates this approach by providing a standardized format for both simulated and experimental data. Trajectories from virtual agents are exported in the same x–y coordinate format as those obtained through motion tracking of real larvae. Behavioral metrics are derived using identical pipelines, allowing for direct comparison. Model parameters may be tuned using optimization techniques (e.g., genetic algorithms) with realism assessed at the level of output similarity.

9.1.2 Mechanistic realism: Constraining internal structure

A second perspective on realism focuses on the internal organization of the model. Here, a model is considered realistic not primarily if it reproduces observed behavior, but if it does so through mechanisms consistent with known anatomical, physiological, or computational principles of the modeled biological system. This view aligns with mechanistic explanation in neuroscience, where explanation requires identifying the parts, interactions, and organization responsible for a phenomenon (Craver, 2006).

In this context, breaking down a behavioral model into smaller, interconnected modules based on neuroscientific insights (anatomical connectivity, functional specialization, or known biophysical properties) imposes internal constraints on the mapping from input to output. Unlike a monolithic black-box that can flexibly fit any behavioral data, a modular architecture must satisfy the constraints imposed by each component and their interconnections.

Larvaworld supports such modular decomposition through its hierarchical control architecture. The platform allows users to implement novel or extend existing modules, provided they conform to defined input-output interfaces. This modularity permits researchers to explicitly embed mechanistic assumptions within any partial module and to study their behavioral consequences at the agent's level. Constraints imposed at the module level shape system-wide dynamics, narrowing the solution space in favor of biologically plausible mechanisms.

9.1.3 Combining approaches: From fitting to explaining

Both notions of realism – behavioral fidelity and mechanistic plausibility – are supported within the present framework, and their integration provides a practical path for model development. Models that prioritize empirical fit may begin as black-box systems, optimized to reproduce behavioral observations across experimental conditions. This allows rapid

prototyping and benchmarking against data when internal mechanisms remain unknown or are too complex to include in early stages.

Conversely, models grounded in mechanistic realism aim to explain behavior through biologically meaningful processes. These rely on modular decomposition informed by anatomical and physiological evidence, enabling targeted hypotheses about internal organization and control. As empirical knowledge increases, modular black-box components can be progressively replaced with structured, mechanism-driven modules, facilitating a smooth transition from behavioral matching to biological explanation.

Larvaworld's architecture enables this dual-track strategy by supporting both output-driven parameter tuning and modular internal design. Models can thus evolve flexibly, incorporating detail where evidence allows, while preserving the capacity to replicate behavior at the whole-organism level. This balance between generality and specificity helps accommodate evolving biological insight without sacrificing empirical grounding.

9.2 Methodological priorities

The computational framework presented in this thesis has been implemented in the *Larvaworld* software package, with the aim of supporting modeling at the interface of behavior, neural control, and physiology. Its design reflects a set of methodological priorities aimed at balancing biological relevance with computational flexibility. Rather than prescribing a fixed modeling approach, the framework is designed to accommodate multiple levels of abstraction and timescales, enabling researchers to build and test models suited to a range of experimental and theoretical contexts.

This section outlines three core methodological trends prioritized in *Larvaworld*: (i) the use of nested timescales to capture dynamics at neural, behavioral, and metabolic levels; (ii) a commitment to modularity and extensibility, enabling the flexible composition, replacement, and reuse of model components; and (iii) the emphasis on closed-loop modeling, supporting real-time interactions within the agent and between the agent and its environment. Together, these trends reflect a broader commitment to flexible, transparent, and biologically grounded modeling across multiple organizational levels. Each is elaborated in the subsections below.

9.2.1 Nested temporal dynamics

A central feature of the framework is its ability to accommodate and coordinate processes operating across multiple timescales, reflecting the temporal stratification observed in biological systems. This nested multiscale structure, illustrated in Figure 15, underpins both the framework's flexibility and its biological relevance.

At the core is the behavioral timescale, which governs the main control loop in *Larvaworld*. Its simulation timestep defaults to 0.1 seconds, matching the 10 Hz framerate at which larval locomotion is typically recorded. This choice is warranted since these datasets

form the empirical testbed for evaluating simulated locomotion. Since all simulated behavior – locomotion, taxis, feeding etc – operates at this timescale, it sets the rhythm of agent-environment interaction.

Faster dynamics are captured at the neural timescale, invoked when modules of the BA are implemented as neuronal circuits in a neural simulator interfacing with *Larvaworld*. For example, the spiking MB model (Chapter 6) operates at a 0.1 millisecond timestep, allowing precise representation of spike timing and synaptic plasticity. This neural activity is embedded within the slower behavioral control loop, enabling mechanisms such as reward-driven learning to modulate behavior in real time. The result is a coherent integration of fast neural dynamics within slower behavioral ones.

At the opposite end of the hierarchy lies the metabolic timescale, introduced through coupling of the BA with the DEB model (Chapter 8). Here, the behavioral loop operates within a slower cycle in which variables like energy reserve, growth, and maturity evolve over minutes to hours of simulated time. These internal states influence behavior via homeostatic feedback, linking long-term physiological change to short-term behavioral expression. Notably, the DEB formalism itself is grounded in circadian-scale dynamics, abstracting away faster fluctuations to model developmental and energetic trends.

By supporting tightly coupled processes across milliseconds to hours, the framework enables simulations that are both mechanistically rich and biologically grounded. It also provides a platform for exploring how dynamics at one timescale constrain, regulate, or enable those at another. Importantly, the modular structure allows these multiscale interactions to be selectively activated or omitted, depending on the aims of a given modeling scenario.

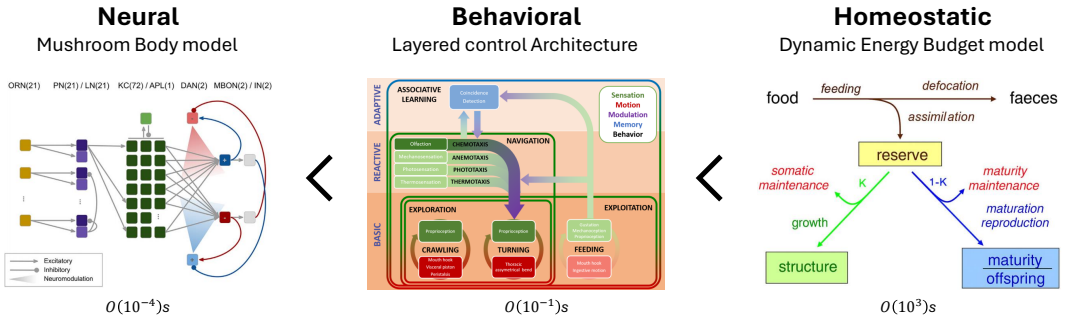

Figure 15: Nested temporal processes across agent components. Neural dynamics (e.g., mushroom body model) unfold at the sub-millisecond scale. Behavioral control (BA) operates at the default timescale of 0.1 seconds. Metabolic processes (DEB model) evolve over minutes to hours ($\sim 10^3$ s). Arrows denote the hierarchical nesting of faster processes within slower ones within a unified simulation environment.

9.2.2 Modularity and extensibility

One of the foundational design principles of the framework developed in this thesis is modularity. From both a conceptual and implementational standpoint, the architecture favors a separation of concerns: distinct behavioral functions are assigned to separate control modules, each of which can be modified, replaced, or omitted independently of the others. This reflects a biologically inspired commitment to distributed control and facilitates the iterative development of more complex behavioral models.

In practical terms, this modular structure is realized in the *Larvaworld* software package (Chapter 5). Each layer of the BA – motor, reactive, and adaptive – consists of discrete, interconnected software modules with well-defined interfaces. These modules can be run independently, configured, or extended with minimal changes to the rest of the codebase, as long as they adhere to the input/output interface specifications. For instance, the reactive layer's sensory transduction mechanisms can be replaced or extended without interfering with motor control, and new learning mechanisms can be plugged into the adaptive layer seamlessly. In fact for each of the modules several alternative implementations are available, mirroring the different mechanistic models that have been proposed for a given behavioral mode. Importantly, individual modules can differ substantially in their level of abstraction – from simple rule-based action selection or algebraic equations to complex systems of differential equations or spiking neural networks – or target timescale as described in the previous section. This heterogeneity is supported natively by the framework and reflects the framework's emphasis on functional encapsulation over algorithmic uniformity.

Extensibility is not limited to behavioral modules. The platform also allows users to configure new environmental designs and data analysis pipelines. This flexibility enables researchers to tailor the framework to diverse experimental paradigms, while maintaining consistency in agent structure and behavioral control. Moreover, the compatibility with third-party libraries and simulation backends ensures that *Larvaworld* can evolve alongside external advances in computational neuroscience and behavioral modeling.

Crucially, this modular design supports both top-down and bottom-up workflows. Researchers may start with a high-level behavioral objective and gradually introduce biological detail using increasingly sophisticated modules. Conversely, detailed but partial neuroscientific models can be plugged into preconfigured control architectures, enabling investigation of their behavioral implications. This bidirectional flexibility makes the framework well-suited both for abstract coarse-grained and for focused, neuroanatomically informed approaches.

9.2.3 Domains of closed-loop modeling

Closed-loop modeling is a central objective of this thesis, referring to simulations in which distinct processes interact dynamically with each other, in real time, through feedback loops. These loops may occur within the agent or extend to its interactions with the

environment. This section outlines four domains of closed-loop processes as modeling targets, spanning sensory, motor, physiological, and neural dimensions.

One domain involves active sensing during taxis behaviors, such as chemotaxis and thermotaxis, where the larva's movement alters its sensory input. As the agent navigates an odor or temperature gradient, its displacement modulates the stimulus it perceives, which in turn influences its motor output. This feedback loop – from sensor to motor to environment and back – engages the reactive and motor layers of the BA and enables successful navigation in multimodal sensory landscapes.

A second domain links foraging behavior and metabolic state through coupling of the BA with the DEB model. In this loop, food ingestion updates the internal energy reserve via assimilation. The interoceptively accessible level of reserve then feeds back into behavior through a hunger/satiety signal, modulating the balance between exploratory and exploitative actions. Thus, physiological state and behavioral decisions are linked in a recurrent feedback loop, establishing an internal closed-loop pathway between metabolism and action selection.

A third domain concerns sensorimotor interaction between the body and the environment, specifically in the context of neuromechanical feedback. While the current motor layer does not incorporate proprioception or explicit physics-based modeling of body–substrate interaction, the *Larvaworld* platform includes a validated interface to the Box2D physics engine. This engine simulates the physical dynamics of segmented larval motion and provides a foundation for integrating biomechanical feedback into future locomotory models.

Finally, a fourth domain is real-time interfacing with neural simulation engines, such as NENGO (Bekolay et al., 2014) and BRIAN2 (Stimberg et al., 2019). *Larvaworld* includes communication pipelines allowing spiking or rate-coded neural circuits to be integrated as modules within the BA. While the technical infrastructure is in place, applying these interfaces in closed-loop neural-behavioral experiments remains a future research direction. A particularly promising application would be to extend the integration of the MB model within the BA, as presented in section 6.1, which so far has operated in an offline manner. The goal is to enable both odor concentration information, acquired through active sensing, and feeding-dependent reward or punishment signals to feed back into the MB circuit in real time, serving as input and modulatory cues respectively. Together, these signals would shape synaptic plasticity and establish a fully closed-loop, experience-driven learning process.

Together, these domains reflect a broader commitment to integrating neural, behavioral, environmental, and physiological processes into unified, biologically grounded simulations. Of the four domains described, the first two have already been implemented and tested in this thesis, while the latter two are currently under development and represent key targets for future work, as discussed in Chapter 10.

10 Questions answered, questions raised

This final chapter adopts a dialogical format to address critical questions that arose during the development and presentation of this work. Rather than summarizing results, it clarifies assumptions, justifies modeling choices, and reflects on the framework's scope, applicability, limitations, and future directions.

The questions derive from interdisciplinary discussions, conceptual challenges in integrative modeling, and anticipated objections from readers across fields such as computational neuroscience, behavioral ecology, and robotics.

Framing these reflections as a Q&A serves two goals: it disambiguates key conceptual and methodological choices, and it highlights opportunities for future extension by confronting unresolved challenges. Each set of questions is thematically grouped and addressed with conceptual and technical precision. This chapter is not a supplement, but a critical reflection on the modeling philosophy, inviting deeper engagement with the framework.

10.1 Beyond the case study

While this thesis centers on the modeling of *Drosophila* larva foraging behavior, its aim is not to remain limited to this case study. Rather, it seeks to develop a modeling framework that is both mechanistically grounded and generative – one that can be adapted to a broader class of organisms and behavioral domains. From this perspective, questions naturally arise:

Why use ABM at all, since larva foraging does not include interaction among conspecifics?

The *Drosophila* larva is not a solitary organism in ecological terms, rather it exhibits a rich array of social behaviors modulated by environmental context, developmental stage, and internal neuromodulation. Larvae aggregate via chemical cues deposited by conspecifics, promoting shared exploitation of resources and visually mediated cooperative digging (Dombrovski et al., 2017, 2019; Louis and de Polavieja, 2017) but also age-related, competitive adaptive shifts in foraging and locomotion (Ruiz-Dubreuil et al., 1996; Sarangi et al., 2016; Wu et al., 2003). Under starvation or crowding, they display predatory cannibalism – young instars attacking larger conspecifics – driven by chemical cues from injured victims and enhanced by nutritional deprivation; this behavior can evolve rapidly under selection and supports full development on a cannibalistic diet (Vijendravarma et al., 2013).

Larvae are implemented as fully embodied agents, and the simulation environment supports direct and indirect inter-larva interactions. Currently implemented scenarios include indirect competition for scarce food resources and direct collision-induced freezing (Otto

et al., 2016). By adopting an ABM framework from the outset, the platform developed here remains structurally open to extensions toward more elaborate social behaviors.

Are there conflicting needs/urgencies? How are they resolved?

The current study focuses exclusively on the behavioral and metabolic components of energy and matter regulation, with a particular emphasis on homeostatically driven foraging behavior. As only a single survival circuit is modeled, homeostatic normativity is expressed in a unidimensional manner – as a hunger/satiety drive – and does not give rise to norm-based conflicts between goals such as feeding, threat avoidance, or reproduction.

Nevertheless, two forms of internal tension are observed within the present framework. First, sensory conflicts may arise, for instance when an appetitive and an aversive odor coincide spatially. These are resolved reactively, through differential activation of sensory channels, without any higher-order arbitration across competing goals – the resolution lies entirely in the sensorimotor output. Second, a dynamic balance is maintained between exploration and exploitation, two behaviors are governed by distinct controllers that compete for motor output, giving rise to stochastic behavioral transitions. Rather than being treated as antagonistic drives, balance between these modes emerges as context-dependent allocation of time and energy, modulated by metabolic state and food availability (section 8.2).

Properly conflicting needs would only emerge with the inclusion of additional survival circuits – such as those related to defense or reproduction. The latter is particularly tractable – although irrelevant for the larval stage – as it introduces its own metabolic demands and constraints, which are already formalized within the DEB theory (Kooijman, 2010) .

How would your framework be adapted for modeling other organisms?

The computational framework has been designed at a sufficiently abstract level to allow its adaptation to organisms beyond the *Drosophila* larva. Implementing it for a species with a richer behavioral repertoire or more complex motivational landscape requires a minimal set of species-specific adjustments, which nonetheless preserve the general architectural principles. These include:

1. A 2D simplification of the body and environment, identifying the degrees of freedom relevant for locomotion, orientation, and sensory access in the target organism. Anatomical details are abstracted unless they bear behavioral or energetic significance.

2. A modular decomposition of the behavioral repertoire into discrete, possibly nested components. Each module corresponds to a behavioral mode – from low-level primitives to higher-level routines – ideally linked to specific sensory inputs, internal states, and motor outputs. This structure supports behavioral switching, mutual inhibition, and state-dependent modulation.

3. The integration or development of computational models for each behavioral module. These may draw from existing work on the target species and may be implemented using diverse paradigms (e.g. dynamical systems, spiking or rate-based networks, or rule-based systems). The modular design allows for their coherent integration into a unified control architecture.

4. A DEB model tailored to the species and developmental stage, parameterized via the Add-my-Pet database procedure. Since DEB theory provides a formal structure for energy acquisition, allocation, and trade-offs, it enables principled modeling of physiological constraints and life-history dynamics (Kooijman, 2010).

These requirements define the minimal basis for modeling a single behavioral domain, as exemplified in the present thesis through foraging. Extending the framework to capture richer behavioral repertoires and conflicting drives entails incorporating additional survival circuits and expanding the control architecture to accommodate multi-dimensional motivational dynamics.

10.2 Deeper into neuroscience

The BA presented in this thesis is grounded in neuroscience, but it does not claim to be a neuroscientific model in the narrow sense. The focus lies in modeling behavior as an emergent outcome of interconnected behavior-based modules that ultimately stand for increasingly nested sensorimotor loops, modulated by internal states and energy dynamics. Nevertheless, questions arise about the neuroscientific detail of the framework and the relation between its current level of abstraction and its extensibility in terms of biological, mechanistic realism:

Can the framework integrate more elaborate, even connectome-informed, neural circuits?

It is true that the modeling framework developed in this thesis primarily operates at the level of functional modules rather than detailed neuronal circuitry. However, this choice is not due to lack of concern for neuroanatomy, connectomics or neural plausibility, but rather reflects a deliberate abstraction aimed at capturing behaviorally relevant dynamics with minimal computational overhead.

That said, one of the collaborative studies does make use of a spiking neuronal model of the MB, with explicit sensory input and dopaminergic modulation (section 6.1). This showcases that the *Larvaworld* software package supports interfacing with dedicated neural simulators, such as BRIAN2 and NENGO, enabling future integration of connectome-informed neural circuits, biologically grounded plasticity rules and neural dynamics. Moreover, support for nested simulations at neural and behavioral timescales has been one of the technical goals already achieved (section 9.2)

More broadly, the entire framework is designed with modularity in mind: each func-

tional component – whether sensory, modulatory or motor – can be replaced or extended by more detailed neuronal models provided that the appropriate interface to the rest of the architecture is defined. In that sense, the current level of abstraction should be seen not as a fixed commitment, but as a practical instantiation within a flexible and extensible architecture.

Does the framework support advanced neuromechanics, realistic body–substrate dynamics and environment physics?

This is a valid concern. The use of a simplified locomotory model – a 2D, bisegmental body – is not meant to dismiss the value of detailed neuromechanical models involving multisegment bodies, CPGs, or intersegmental coordination Bidaye et al., 2018; Büschges, 2005; Mantziaris et al., 2020. Rather, it reflects a deliberate modeling trade-off in which the goal is to implement and test high-level control architectures under minimal embodiment constraints. That said, since the framework's implementation is modular, individual components – including the behavior-based modules of the motor layer, body dynamics, and environmental physics – can be replaced or extended as needed.

Notably, *Larvaworld* interfaces with the Box2D physics engine, allowing the creation of multisegment larval bodies composed of interconnected rigid segments. Unlike the default kinematic approach, motion here is governed by physical forces, with the physics engine handling the computational load of simulating torques, friction, and dynamics. However, such models require careful calibration and validation, therefore higher neuromechanical realism remains an open direction for future development.

On the environmental side, some basic physical processes have already been implemented to support dynamic sensory landscapes. These include an odor diffusion algorithm for simulating the odor gradients emitted by localized odor sources, combination of such odorscapes with wind-based advection of odor plumes, and spatiotemporal evolution of environmental traits such as expansion or depletion of localized substrate nutrients. Together, these features allow for nontrivial agent–environment interactions, even within the currently simplified 2D setting.

10.3 Scientific usability beyond modeling

The development of *Larvaworld* is primarily motivated by modeling needs, but its design favors possible uses beyond simulation. This section reflects on how the platform might intersect with experimental workflows and educational settings. These perspectives have gradually emerged during interdisciplinary collaboration and community engagement and are included here as plausible directions worth exploring further.

Can *Larvaworld* also benefit experimentalists, not just modelers?

Yes, *Larvaworld* is explicitly designed to bridge the gap between computational modeling and experimental behavioral research. For experimentalists, it functions as a tool for importing, standardizing, and analyzing motion-tracking data from real *Drosophila* larvae. Datasets from various tracking systems can be converted to a standardized format, allowing for consistent analysis across experiments and laboratories. This helps reduce methodological variability and facilitates reproducible data processing pipelines.

Importantly, experimental datasets and simulation outputs are treated identically within the platform. This enables direct, unbiased comparisons between real and simulated behavior using the same quantitative metrics and visualization tools. Experimentalists can explore how specific behavioral features vary across genotypes, feeding conditions, or environmental manipulations, and test how well computational models reproduce their findings. A fruitful future direction would be to interface it with PiVR – a tool for closed-loop experimentation – with which it shares multiple features (Tadres and Louis, 2020). In short, *Larvaworld* helps establish a tight feedback loop between empirical findings and mechanistic hypotheses, a valuable tool not only for model development but also for hypothesis generation and testing in experimental behavioral neuroscience.

How can *Larvaworld* be used in scientific education?

While *Larvaworld* was originally developed for research, its design makes it suitable as a virtual laboratory in educational contexts, particularly where access to live animal experimentation is limited. It offers interactive tutorial notebooks, browser-based tools, and preconfigured simulations that allow students to explore core concepts in behavioral modeling through direct engagement. Its modular architecture enables users to build synthetic agents step-by-step by combining functional components – such as locomotory, sensory or learning modules – and watch increasingly complex behaviors unfold. Or they can construct and modify virtual arenas, configure multimodal sensory landscapes, and observe how environmental features influence behavior, gaining insight into behavioral control in stimulus-rich environments.

To broaden its accessibility in educational settings, several directions are being considered. These include the development of a graphical interface that would allow users to configure and run experiments without writing code, making the platform accessible to non-programmers. Integration into open online courses and the creation of student-oriented project modules could further support its use in digital education settings. Although these applications are still under development, they suggest that *Larvaworld* has the potential to serve not only as a research toolkit but also as a flexible platform for interactive and model-driven scientific education.

References

Adden, A., Stewart, T. C., Webb, B., & Heinze, S. (2022). A neural model for insect steering applied to olfaction and path integration. *Neural Computation*, *34*(11), 2205–2231. https://doi.org/10.1162/neco_a_01540 (cit. on p. 21).

Arendt, D., Tosches, M. A., & Marlow, H. (2016). From nerve net to nerve ring, nerve cord and brain-evolution of the nervous system. *Nature Reviews Neuroscience*, *17*(1), 61–72. https://doi.org/10.1038/nrn.2015.15 (cit. on p. 8).

Baedke, J. (2018). O organism, where art thou? old and new challenges for organism-centered biology. *Journal of the History of Biology*, *52*(2), 293–324. https://doi.org/10.1007/s10739-018-9549-4 (cit. on p. 16).

Baedke, J., Fábregas-Tejeda, A., & Prieto, G. I. (2021). Unknotting reciprocal causation between organism and environment. *Biology and Philosophy*, *36*(5). https://doi.org/10.1007/s10539-021-09815-0 (cit. on p. 15).

Barandiaran, X., & Moreno, A. (2008). Adaptivity: From Metabolism to Behavior. https://doi.org/10.1177/1059712308093868 (cit. on pp. 16, 17).

Barandiaran, X. E. (2017). Autonomy and Enactivism: Towards a Theory of Sensorimotor Autonomous Agency. *Topoi*, *36*(3), 409–430. https://doi.org/10.1007/s11245-016-9365-4 (cit. on p. 17).

Barandiaran, X. E., & Chemero, A. (2009). Animats in the Modeling Ecosystem. *Adaptive Behavior*, *17*(4), 287–292. https://doi.org/10.1177/1059712309340847 (cit. on p. 17).

Barandiaran, X. E., & Egbert, M. D. (2014). Norm-Following in Autonomous Agency. *28*, 5–28. https://doi.org/10.1162/ARTL (cit. on p. 17).

Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *36*(2), 421–441. https://doi.org/10.1016/j.shpsc.2005.03.010 (cit. on p. 18).

Beggs, J. M., & Plenz, D. (2003). Neuronal avalanches in neocortical circuits. *Journal of Neuroscience*, *23*(35), 11167–11177. https://doi.org/10.1523/JNEUROSCI.23-35-11167.2003 (cit. on p. 125).

Bekolay, T., Bergstra, J., Hunsberger, E., DeWolf, T., Stewart, T., Rasmussen, D., Choo, X., Voelker, A., & Eliasmith, C. (2014). Nengo: A Python tool for building large-scale functional brain models. *Frontiers in Neuroinformatics*, *7*(48), 1–13. https://doi.org/10.3389/fninf.2013.00048 (cit. on p. 143).

Berman, G. J., Bialek, W., & Shaevitz, J. W. (2016). Predictability and hierarchy in drosophila behavior. *Proceedings of the National Academy of Sciences*, *113*(42), 11943–11948. https://doi.org/10.1073/pnas.1607601113 (cit. on p. 13).

Berman, G. J., Choi, D. M., Bialek, W., & Shaevitz, J. W. (2014). Mapping the stereotyped behaviour of freely moving fruit flies. *Journal of the Royal Society Interface*, *11*(99). https://doi.org/10.1098/rsif.2014.0672 (cit. on p. 13).

Bidaye, S. S., Bockemühl, T., & Büschges, A. (2018). Six-legged walking in insects: How cpgs, peripheral feedback, and descending signals generate coordinated and adaptive motor rhythms. *Journal of Neurophysiology*, *119*(2), 459–475. https://doi.org/10.1152/jn.00658.2017 (cit. on p. 147).

Brooks, R. (1986). A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation*, *2*(1), 14–23. https://doi.org/10.1109/jra.1986.1087032 (cit. on p. 10).

Brown, R. E., Basheer, R., McKenna, J. T., Strecker, R. E., & McCarley, R. W. (2012). Control of Sleep and Wakefulness. *Physiological Reviews*, *92*(3), 1087–1187. https://doi.org/10.1152/physrev.00032.2011 (cit. on p. 9).

Budaev, S., Giske, J., & Eliassen, S. (2018). AHA: A general cognitive architecture for Darwinian agents. *Biologically Inspired Cognitive Architectures*, *25*(July), 51–57. https://doi.org/10.1016/j.bica.2018.07.009 (cit. on p. 6).

Buhrmann, T., Di Paolo, E. A., & Barandiaran, X. (2013). A dynamical systems account of sensorimotor contingencies. *Frontiers in Psychology*, *4*(MAY), 1–19. https://doi.org/10.3389/fpsyg.2013.00285 (cit. on p. 17).

Büschges, A. (2005). Sensory control and organization of neural networks mediating coordination of multisegmental organs for locomotion. *Journal of Neurophysiology*, *93*(3), 1127–1135. https://doi.org/10.1152/jn.00615.2004 (cit. on pp. 9, 147).

Büschges, A., & Ache, J. M. (2025). Motor control on the move: From insights in insects to general mechanisms. *Physiological Reviews*, *105*(3), 975–1031. https://doi.org/10.1152/physrev.00009.2024 (cit. on p. 9).

Büschges, A., & Borgmann, A. (2013). Network modularity: Back to the future in motor control. *Current Biology*, *23*(20), R936–R938. https://doi.org/10.1016/j.cub.2013.09.021 (cit. on p. 4).

Büschges, A., & Gorostiza, E. A. (2023). Neurons with names: Descending control and sensorimotor processing in insect motor control. *Current Opinion in Neurobiology*, *83*, 102766. https://doi.org/10.1016/j.conb.2023.102766 (cit. on p. 9).

Cisek, P. (2021). Evolution of behavioural control from chordates to primates. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *377*(1844). https://doi.org/10.1098/rstb.2020.0522 (cit. on p. 8).

Cook, S. J., Jarrell, T. A., Brittin, C. A., Wang, Y., Bloniarz, A. E., Yakovlev, M. A., Nguyen, K. C. Q., Tang, L. T.-H., Bayer, E. A., Duerr, J. S., Bülow, H. E., Hobert, O., Hall, D. H., & Emmons, S. W. (2019). Whole-animal connectomes of both Caenorhabditis elegans sexes. *Nature*, *571*(7763), 63–71. https://doi.org/10.1038/s41586-019-1352-7 (cit. on pp. 3, 20).

Craver, C. F. (2006). When mechanistic models explain. *Synthese*, *153*(3), 355–376. https://doi.org/10.1007/s11229-006-9097-x (cit. on pp. 18, 139).

Datta, S. R., Anderson, D. J., Branson, K., Perona, P., & Leifer, A. (2019). Computational Neuroethology : A Call to Action. *Neuron*, *104*(1), 11–24. https://doi.org/10.1016/j.neuron.2019.09.038 (cit. on p. 13).

Daun-Gruhn, S., & Büschges, A. (2011). From neuron to behavior: Dynamic equation-based prediction of biological processes in motor control. *Biological Cybernetics*, *105*(1), 71–88. https://doi.org/10.1007/s00422-011-0446-6 (cit. on p. 4).

de Tredern, E., Manceau, D., Blanc, A., Sakagiannis, P., Barre, C., Sus, V., Viscido, F., Hasan, M. A., Autran, S., Nawrot, M., Mason, J.-B., & Jovanic, T. (2024). Feeding-state dependent neuropeptidergic modulation of reciprocally interconnected inhibitory neurons biases sensorimotor decisions in Drosophila. *bioRxiv*. https://doi.org/https://doi.org/10.1101/2023.12.26.573306 (cit. on pp. 24, 120, 122, 131).

Di Paolo, E. A., Buhrmann, T., & Barandiaran, X. E. (2017). *Sensorimotor life: An enactive proposal*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198786849.001.0001 (cit. on pp. 4, 17).

Dombrovski, M., Kim, A., Poussard, L., Vaccari, A., Acton, S., Spillman, E., Condron, B., & Yuan, Q. (2019). A Plastic Visual Pathway Regulates Cooperative Behavior in Drosophila Larvae. *Current Biology*, *29*(11), 1866–1876.e5. https://doi.org/10.1016/j.cub.2019.04.060 (cit. on p. 144).

Dombrovski, M., Poussard, L., Moalem, K., Kmecova, L., Hogan, N., Schott, E., Vaccari, A., Acton, S., & Condron, B. (2017). Cooperative Behavior Emerges among Drosophila Larvae. *Current Biology*, *27*(18), 2821–2826.e2. https://doi.org/10.1016/j.cub.2017.07.054 (cit. on p. 144).

Duffy, J. B. (2002). GAL4 system in *drosophila* : A fly geneticist's swiss army knife. *genesis*, *34*(1-2), 1–15. https://doi.org/10.1002/gene.10150 (cit. on p. 20).

Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., & Rasmussen, D. (2012). A large-scale model of the functioning brain. *Science*, *338*(6111), 1202–1205. https://doi.org/10.1126/science.1225266 (cit. on p. 4).

Eliassen, S., Andersen, B. S., Jørgensen, C., & Giske, J. (2016). From sensing to emergent adaptations: Modelling the proximate architecture for decision-making. *Ecological Modelling*, *326*, 90–100. https://doi.org/10.1016/j.ecolmodel.2015.09.001 (cit. on p. 6).

Eschbach, C., & Zlatic, M. (2020). Useful road maps: Studying Drosophila larva's central nervous system with the help of connectomics. *Current Opinion in Neurobiology*, *65*, 129–137. https://doi.org/10.1016/j.conb.2020.09.008 (cit. on p. 129).

Fábregas-Tejeda, A., Baedke, J., Prieto, G. I., & Radick, G. (2024, June). *The riddle of organismal agency: New historical and philosophical reflections*. Routledge. https://doi.org/10.4324/9781003413318 (cit. on p. 15).

Farris, S. M. (2011). Are mushroom bodies cerebellum-like structures? *Arthropod Structure and Development*, *40*(4), 368–379. https://doi.org/10.1016/j.asd.2011.02.004 (cit. on p. 9).

Fitzpatrick, M. J., Feder, E., Rowe, L., & Sokolowski, M. B. (2007). Maintaining a behaviour polymorphism by frequency-dependent selection on a single gene. *Nature*, *447*(7141), 210–212. https://doi.org/10.1038/nature05764 (cit. on p. 137).

Franklin, S., Madl, T., Strain, S., Faghihi, U., Dong, D., Kugele, S., Snaider, J., Agrawal, P., & Chen, S. (2016). A LIDA cognitive model tutorial. *Biologically Inspired Cognitive Architectures*. https://doi.org/10.1016/j.bica.2016.04.003 (cit. on p. 4).

Fuller, P. M., Gooley, J. J., & Saper, C. B. (2006). Neurobiology of the sleep-wake cycle: Sleep architecture, circadian regulation, and regulatory feedback. *Journal of biological rhythms*, *21*(6), 482–493. https://doi.org/10.1177/0748730406294627 (cit. on p. 9).

Geissmann, Q., Garcia Rodriguez, L., Beckwith, E. J., French, A. S., Jamasb, A. R., & Gilestro, G. F. (2017). Ethoscopes: An open platform for high-throughput ethomics. *PLOS Biology*, *15*(10), e2003026. https://doi.org/10.1371/journal.pbio.2003026 (cit. on p. 20).

Gkoutos, G. V., Schofield, P. N., & Hoehndorf, R. (2012). *The Neurobehavior Ontology : An Ontology for Annotation and Integration of Behavior and Behavioral Phenotypes* (Vol. 103). https://doi.org/10.1016/B978-0-12-388408-4.00004-6 (cit. on p. 14).

Grafen, A. (1986). Natural selection, kin selection and group selection. *Behavioural Ecology : An Evolutionary Approach*. https://cir.nii.ac.jp/crid/1571417125212123776 (cit. on p. 5).

Gras, R., Devaurs, D., Wozniak, A., & Aspinall, A. (2009). An individual-based evolving predator-prey ecosystem simulation using a fuzzy cognitive map as the behavior model. *Artificial Life*, *15*(4), 423–463. https://doi.org/10.1162/artl.2009.Gras.012 (cit. on p. 5).

Grimm, V. (2005). Pattern-Oriented Modeling of Agent-Based Complex Systems: Lessons from Ecology. *Science*, *310*(5750), 987–991. https://doi.org/10.1126/science.1116681 (cit. on pp. 3, 5).

Grimm, V., Berger, U., DeAngelis, D. L., Polhill, J. G., Giske, J., & Railsback, S. F. (2010). The ODD protocol: A review and first update. *Ecological Modelling*, *221*(23), 2760–2768. https://doi.org/10.1016/j.ecolmodel.2010.08.019 (cit. on p. 5).

Haenicke, J., Zwaka, H., & Nawrot, M. (2018). Neural Correlates of Odor Learning in the Presynaptic Microglomerular Circuitry in the Honeybee Mushroom Body Calyx. *eNeuro*, *5*(3), 1–29. https://doi.org/10.1523/ENEURO.0128-18.2018 (cit. on p. 4).

Hernandez-Nunez, L., Chen, A., Budelli, G., Berck, M. E., Richter, V., Rist, A., Thum, A. S., Cardona, A., Klein, M., Garrity, P., & Samuel, A. D. T. (2021). Synchronous and opponent thermosensors use flexible cross-inhibition to orchestrate thermal homeostasis. *Science Advances*, *7*(35). https://doi.org/10.1126/sciadv.abg6707 (cit. on p. 21).

Holland, C. (2018). Sleep homeostasis indrosophila: A window on the vital function of sleep. *Bioscience Horizons: The International Journal of Student Research*, *11*. https://doi.org/10.1093/biohorizons/hzy009 (cit. on p. 9).

Hommel, B., Chapman, C. S., Cisek, P., Neyedli, H. F., Song, J. H., & Welsh, T. N. (2019). No one knows what attention is. *Attention, Perception, and Psychophysics*, *81*(7), 2288–2303. https://doi.org/10.3758/s13414-019-01846-w (cit. on p. 10).

Hückesfeld, S., Schlegel, P., Miroschnikow, A., Schoofs, A., Zinke, I., Haubrich, A. N., Schneider-Mizell, C. M., Truman, J. W., Fetter, R. D., Cardona, A., & Pankratz, M. J. (2021). Unveiling the sensory and interneuronal pathways of the neuroendocrine connectome in drosophila. *eLife*, *10*, 1–25. https://doi.org/10.7554/eLife.65745 (cit. on p. 20).

Humberg, T. H., Bruegger, P., Afonso, B., Zlatic, M., Truman, J. W., Gershow, M., Samuel, A., & Sprecher, S. G. (2018). Dedicated photoreceptor pathways in Drosophila larvae mediate navigation by processing either spatial or temporal cues. *Nature Communications*, *9*(1), 1–16. https://doi.org/10.1038/s41467-018-03520-5 (cit. on p. 21).

Huse, G., Strand, E., & Giske, J. (1999). Implementing behaviour in individual-based models using neural networks and genetic algorithms. *Evolutionary Ecology*, *13*(5), 469–483. https://doi.org/10.1023/A:1006746727151 (cit. on p. 5).

Itskov, P. M., Moreira, J. M., Vinnik, E., Lopes, G., Safarik, S., Dickinson, M. H., & Ribeiro, C. (2014). Automated monitoring and quantitative analysis of feeding behaviour in Drosophila. *Nature Communications*, *5*. https://doi.org/10.1038/ncomms5560 (cit. on p. 131).

Izquierdo, E. J., & Beer, R. D. (2013). Connecting a connectome to behavior: An ensemble of neuroanatomical models of c. elegans klinotaxis. *PLoS Computational Biology*, *9*(2), e1002890. https://doi.org/10.1371/journal.pcbi.1002890 (cit. on p. 3).

Jackson, J. H. (1884). The croonian lectures on evolution and dissolution of the nervous system. *BMJ*, *1*(1215), 703–707. https://doi.org/10.1136/bmj.1.1215.703 (cit. on p. 10).

Jovanic, T., Winding, M., Cardona, A., Truman, J. W., Gershow, M., & Zlatic, M. (2019). Neural Substrates of Drosophila Larval Anemotaxis. *Current Biology*, *29*(4), 554–566.e4. https://doi.org/10.1016/j.cub.2019.01.009 (cit. on p. 21).

Jürgensen, A.-M., Sakagiannis, P., Schleyer, M., Gerber, B., & Nawrot, M. P. (2024). Prediction error drives associative learning and conditioned behavior in a spiking model of Drosophila larva. *iScience*, *27*(1), 108640. https://doi.org/10.1016/j.isci.2023.108640 (cit. on pp. 24, 114, 116).

Kabra, M., Robie, A. A., Rivera-Alba, M., Branson, S., & Branson, K. (2013). Jaaba: Interactive machine learning for automatic annotation of animal behavior. *Nature Methods*, *10*(1), 64–67. https://doi.org/10.1038/nmeth.2281 (cit. on p. 20).

Kafle, T., Grub, M., Sakagiannis, P., Nawrot, M. P., & Arguello, J. R. (2025). Evolution of temperature preference behaviour among drosophila larvae. *iScience*, *28*(7), 112809. https://doi.org/10.1016/j.isci.2025.112809 (cit. on pp. 24, 117, 119).

Kaiser, M. I., & Krickel, B. (2017). The metaphysics of constitutive mechanistic phenomena. *British Journal for the Philosophy of Science*, *68*(3), 745–779. https://doi.org/10.1093/bjps/axv058 (cit. on p. 18).

Katsov, A. Y., Freifeld, L., Horowitz, M., Kuehn, S., & Clandinin, T. R. (2017). Dynamic structure of locomotor behavior in walking fruit flies. *eLife*, *6*, 1–32. https://doi.org/10.7554/eLife.26410 (cit. on p. 13).

Kaun, K. R., Chakaborty-Chatterjee, M., & Sokolowski, M. B. (2008). Natural variation in plasticity of glucose homeostasis and food intake. *Journal of Experimental Biology*, *211*(19), 3160–3166. https://doi.org/10.1242/jeb.010124 (cit. on pp. 132, 137).

Kaun, K. R., Riedl, C. A., Chakaborty-Chatterjee, M., Belay, A. T., Douglas, S. J., Gibbs, A. G., & Sokolowski, M. B. (2007). Natural variation in food acquisition mediated via a Drosophila cGMP-dependent protein kinase. *Journal of Experimental Biology*, *210*(20), 3547–3558. https://doi.org/10.1242/jeb.006924 (cit. on p. 137).

Kim, I. S., & Dickinson, M. H. (2017). Idiothetic Path Integration in the Fruit Fly Drosophila melanogaster. *Current Biology*, *27*(15), 2227–2238.e3. https://doi.org/10.1016/j.cub.2017.06.026 (cit. on p. 132).

Klein, C., & Barron, A. B. (2024). Comparing cognition across major transitions using the hierarchy of formal automata. *WIREs Cognitive Science*, *15*(4). https://doi.org/10.1002/wcs.1680 (cit. on p. 9).

Kooijman, S. a. L. M. (2010). Dynamic Energy Budget theory for metabolic organisation : Summary of concepts of the third edition. *Water*, *365*, 68. https://doi.org/10.1098/rstb.2010.0167 (cit. on pp. 9, 130, 132, 134, 135, 145, 146).

Langton, C. (1992). Artificial Life II. In *Volume X of SFI Studies in the Sciences of Complexity* (pp. xiii–xviii). (Cit. on p. 4).

LeDoux, J. (2012). Rethinking the Emotional Brain. *Neuron*, *73*(4), 653–676. https://doi.org/10.1016/j.neuron.2012.02.004 (cit. on p. 8).

Ledoux, J., & Daw, N. D. (2018). Surviving threats: Neural circuit and computational implications of a new taxonomy of defensive behaviour. *Nature Reviews Neuroscience*, *19*(5), 269–282. https://doi.org/10.1038/nrn.2018.22 (cit. on pp. 8, 10, 13).

Llandres, A. L., Marques, G. M., Maino, J. L., Kooijman, S. A. L. M., Kearney, M. R., & Casas, J. (2015). A dynamic energy budget for the whole life-cycle of holometabolous insects. *Ecological Monographs*, *85*(3), 353–371. https://doi.org/10.1890/14-0976.1 (cit. on p. 136).

Louis, M., & de Polavieja, G. (2017). Collective Behavior: Social Digging in Drosophila Larvae. *Current Biology*, *27*(18), R1010–R1012. https://doi.org/10.1016/j.cub.2017.08.023 (cit. on p. 144).

Maffei, G., Santos-Pata, D., Marcos, E., Sánchez-Fibla, M., & Verschure, P. F. (2015). An embodied biologically constrained model of foraging: From classical and operant conditioning to adaptive real-world behavior in DAC-X. *Neural Networks*, *72*, 88–108. https://doi.org/10.1016/j.neunet.2015.10.004 (cit. on p. 21).

Mantziaris, C., Bockemühl, T., & Büschges, A. (2020). Central Pattern Generating Networks in Insect Locomotion. *Developmental Neurobiology*, *00*, 1–15. https://doi.org/10.1002/dneu.22738 (cit. on pp. 9, 126, 147).

Martin, B. T., Zimmer, E. I., Grimm, V., & Jager, T. (2012). Dynamic Energy Budget theory meets individual-based modelling: A generic and accessible implementation. *Methods in Ecology and Evolution*, *3*(2), 445–449. https://doi.org/10.1111/j.2041-210X.2011.00168.x (cit. on p. 130).

Maselli, A., Gordon, J., Eluchans, M., Lancia, G. L., Thiery, T., Moretti, R., Cisek, P., & Pezzulo, G. (2023). Beyond simple laboratory studies: Developing sophisticated models to study rich behavior. *Physics of Life Reviews*, *46*, 220–244. https://doi.org/10.1016/j.plrev.2023.07.006 (cit. on p. 12).

Miroschnikow, A., Schlegel, P., Schoofs, A., Hueckesfeld, S., Li, F., Schneider-Mizell, C. M., Fetter, R. D., Truman, J. W., Cardona, A., & Pankratz, M. J. (2018). The Feeding Connectome: Convergence of Monosynaptic and Polysynaptic Sensory Paths onto Common Motor Outputs. *eLife*, *7*, 1–23. https://doi.org/10.7554/elife.40247 (cit. on p. 20).

Mobbs, D., Wise, T., Suthana, N., Guzmán, N., Kriegeskorte, N., & Leibo, J. Z. (2021). Promises and challenges of human computational ethology. *Neuron*, *109*(14), 2224–2238. https://doi.org/10.1016/j.neuron.2021.05.021 (cit. on p. 13).

Müller, B., Bohn, F., Dreßler, G., Groeneveld, J., Klassert, C., Martin, R., Schlüter, M., Schulze, J., Weise, H., & Schwarz, N. (2013). Describing human decisions in agent-based models - ODD+D, an extension of the ODD protocol. *Environmental Modelling and Software*, *48*, 37–48. https://doi.org/10.1016/j.envsoft.2013.06.003 (cit. on p. 5).

Nath, T., Mathis, A., Chen, A. C., Patel, A., Bethge, M., & Mathis, M. W. (2019). Using deeplabcut for 3d markerless pose estimation across species and behaviors. *Nature Protocols*, *14*(7), 2152–2176. https://doi.org/10.1038/s41596-019-0176-0 (cit. on p. 20).

Ohyama, T., Schneider-Mizell, C. M., Fetter, R. D., Aleman, J. V., Franconville, R., Rivera-Alba, M., Mensh, B. D., Branson, K. M., Simpson, J. H., Truman, J. W., Cardona, A., & Zlatic, M. (2015). A multilevel multimodal circuit enhances action selection in drosophila. *Nature*, *520*(7549), 633–639. https://doi.org/10.1038/nature14297 (cit. on p. 20).

Ormerod, K. G., LePine, O. K., Abbineni, P. S., Bridgeman, J. M., Coorssen, J. R., Mercier, A. J., & Tattersall, G. J. (2017). Drosophila development, physiology, behavior, and lifespan are influenced by altered dietary composition. *Fly*, *11*(3), 153–170. https://doi.org/10.1080/19336934.2017.1304331 (cit. on p. 136).

Osumi-Sutherland, D., Marygold, S. J., Millburn, G. H., McQuilton, P. A., Ponting, L., Stefancsik, R., Falls, K., Brown, N. H., & Gkoutos, G. V. (2013). The Drosophila phenotype ontology. *Journal of Biomedical Semantics*, *4*(1), 1–10. https://doi.org/10.1186/2041-1480-4-30 (cit. on p. 14).

Otto, N., Risse, B., Berh, D., Bittern, J., Jiang, X., & Klämbt, C. (2016). Interactions among Drosophila larvae before and during collision. *Scientific Reports*, *6*(July), 1–11. https://doi.org/10.1038/srep31564 (cit. on p. 144).

Oxenkrug, G., Navrotskaya, V., Vorobyova, L., & Summergrad, P. (2011). Extension of life span of drosophila Melanogaster by the inhibitors of tryptophan-kynurenine metabolism. *Fly*, *5*(4), 307–309. https://doi.org/10.4161/fly.5.4.18414 (cit. on p. 136).

Paisios, E., Rjosk, A., Pamir, E., & Schleyer, M. (2017). Common microbehavioral "footprint" of two distinct classes of conditioned aversion. *Learning and Memory*, *24*(5), 191–198. https://doi.org/10.1101/lm.045062.117 (cit. on p. 21).

Prescott, T. J., Redgrave, P., & Gurney, K. (1999). Layered control architectures in robots and vertebrates. *Adaptive Behavior*, *7*(1), 99–127. https://doi.org/10.1177/105971239900700105 (cit. on p. 11).

Pulver, S. R., Bayley, T. G., Taylor, A. L., Berni, J., Bate, M., & Hedwig, B. (2015). Imaging fictive locomotor patterns in larval drosophila. *Journal of Neurophysiology*, *114*(5), 2564–2577. https://doi.org/10.1152/jn.00731.2015 (cit. on p. 126).

Risse, B., Berh, D., Otto, N., Klämbt, C., & Jiang, X. (2017). FIMTrack: An open source tracking and locomotion analysis software for small animals. *PLoS Computational Biology*, *13*(5), 1–15. https://doi.org/10.1371/journal.pcbi.1005530 (cit. on p. 20).

Risse, B., Otto, N., Berh, D., Jiang, X., Kiel, M., & Klambt, C. (2017). Fim^{2c} : Multicolor, multipurpose imaging system to manipulate and analyze animal behavior. *IEEE Transactions on Biomedical Engineering*, *64*(3), 610–620. https://doi.org/10.1109/tbme.2016.2570598 (cit. on p. 20).

Robson, D. N., & Li, J. M. (2022). A dynamical systems view of neuroethology: Uncovering stateful computation in natural behaviors. *Current Opinion in Neurobiology*, *73*, 102517. https://doi.org/10.1016/j.conb.2022.01.002 (cit. on p. 13).

Rolls, E. T. (2015). Limbic systems for emotion and for memory, but no single limbic system. *Cortex*, *62*, 119–157. https://doi.org/10.1016/j.cortex.2013.12.005 (cit. on p. 4).

Ruiz-Dubreuil, G., Burnet, B., Connolly, K., & Furness, P. (1996). Larval foraging behaviour and competition in Drosophila melanogaster. *Heredity*, *76*(1), 55–64. https://doi.org/10.1038/hdy.1996.7 (cit. on p. 144).

Sakagiannis, P., Aguilera, M., & Nawrot, M. P. (2020). A Plausible Mechanism for Drosophila Larva Intermittent Behavior. In *Biomimetic and Biohybrid Systems* (pp. 288–299, Vol. 12413). Springer International Publishing. https://doi.org/10.1007/978-3-030-64313-3_28 (cit. on p. 24).

Sakagiannis, P., Jürgensen, A.-M., & Nawrot, M. P. (2025). A behavioral architecture for realistic simulations of drosophila larva locomotion and foraging. https://doi.org/10.7554/elife.104262.1 (cit. on p. 24).

Sakagiannis, P., Rapp, H., Jovanic, T., & Nawrot, M. P. (2025). Larvaworld : A behavioral simulation and analysis platform for drosophila larva. *bioRxiv*. https://doi.org/10.1101/2025.06.15.659765 (cit. on p. 24).

Sarangi, M., Nagarajan, A., Dey, S., Bose, J., & Joshi, A. (2016). Evolution of increased larval competitive ability in Drosophila melanogaster without increased larval feed-

ing rate. *Journal of Genetics*, *95*(3), 491–503. https://doi.org/10.1007/s12041-016-0656-8 (cit. on p. 144).

Schindelman, G., Fernandes, J. S., Bastiani, C. A., Yook, K., & Sternberg, P. W. (2011). Worm Phenotype Ontology: Integrating phenotype data within and beyond the C. elegans community. *BMC Bioinformatics*, *12*. https://doi.org/10.1186/1471-2105-12-32 (cit. on p. 14).

Schneider-Mizell, C. M., Gerhard, S., Longair, M., Kazimiers, T., Li, F., Zwart, M. F., Champion, A., Midgley, F. M., Fetter, R. D., Saalfeld, S., & Cardona, A. (2016). Quantitative neuroanatomy for connectomics in drosophila. *eLife*, *5*. https://doi.org/10.7554/elife.12059 (cit. on p. 20).

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593–1599. https://doi.org/10.1126/science.275.5306.1593 (cit. on p. 127).

Schumann, I., & Triphan, T. (2020). The PEDtracker: An Automatic Staging Approach for Drosophila melanogaster Larvae. *Frontiers in Behavioral Neuroscience*, *14*. https://doi.org/10.3389/fnbeh.2020.612313 (cit. on pp. 20, 136).

Sims, K. (1994). Evolving 3d morphology and behavior by competition. *Artificial Life*, *1*(4), 353–372. https://doi.org/10.1162/artl.1994.1.4.353 (cit. on p. 5).

Song, Y., Park, J. O., Tanner, L., Nagano, Y., Rabinowitz, J. D., & Shvartsman, S. Y. (2019). Energy budget of Drosophila embryogenesis. *Current Biology*, *29*(12), R566–R567. https://doi.org/10.1016/j.cub.2019.05.025 (cit. on p. 136).

Stimberg, M., Brette, R., & Goodman, D. F. (2019). Brian 2, an intuitive and efficient neural simulator. *eLife*, *8*. https://doi.org/10.7554/elife.47314 (cit. on p. 143).

Sun, R. (2007). The importance of cognitive architectures: An analysis based on clarion. *Journal of Experimental & Theoretical Artificial Intelligence*, *19*(2), 159–193. https://doi.org/10.1080/09528130701191560 (cit. on p. 4).

Sun, X., Mangan, M., Peng, J., & Yue, S. (2025). I2bot: An open-source tool for multimodal and embodied simulation of insect navigation. *Journal of The Royal Society Interface*, *22*(222). https://doi.org/10.1098/rsif.2024.0586 (cit. on p. 20).

Sun, X., Yue, S., & Mangan, M. (2020). A decentralised neural model explaining optimal integration of navigational strategies in insects. *eLife*, *9*, 1–30. https://doi.org/10.7554/eLife.54026 (cit. on p. 21).

Tadres, D., & Louis, M. (2020). PiVR: An affordable and versatile closed-loop platform to study unrestrained sensorimotor behavior. *PLOS Biology*, *18*(7), e3000712. https://doi.org/10.1371/journal.pbio.3000712 (cit. on pp. 20, 148).

Tastekin, I., Khandelwal, A., Tadres, D., Fessner, N. D., Truman, J. W., Zlatic, M., Cardona, A., & Louis, M. (2018). Sensorimotor pathway controlling stopping behavior during chemotaxis in the Drosophila melanogaster larva. *eLife*, *7*, 1–38. https://doi.org/10.7554/elife.38740 (cit. on p. 21).

Verschure, P. F. (2012). Distributed adaptive control: A theory of the mind, brain, body nexus. *Biologically Inspired Cognitive Architectures*, *1*, 55–72. https://doi.org/https://doi.org/10.1016/j.bica.2012.04.005 (cit. on pp. 4, 12).

Vijendravarma, R. K., Narasimha, S., & Kawecki, T. J. (2013). Predatory cannibalism in Drosophila melanogaster larvae. *Nature Communications*, *4*, 1788–1789. https://doi.org/10.1038/ncomms2744 (cit. on p. 144).

Vogt, K., Zimmerman, D. M., Schlichting, M., Hernandez-Nunez, L., Qin, S., Malacon, K., Rosbash, M., Pehlevan, C., Cardona, A., & Samuel, A. D. (2021). Internal state configures olfactory behavior and early sensory processing in Drosophila larvae. *Science Advances*, *7*, 1, eabd6900. https://doi.org/10.1126/sciadv.abd6900 (cit. on p. 21).

Wang, X.-J. (2008). Decision Making in Recurrent Neuronal Circuits. *Neuron*, *60*(2), 215–234. https://doi.org/10.1016/j.neuron.2008.09.034 (cit. on p. 4).

Ward, D., Silverman, D., & Villalobos, M. (2017). Introduction: The varieties of enactivism. *Topoi*, *36*(3), 365–375. https://doi.org/10.1007/s11245-017-9484-6 (cit. on p. 4).

Wilson, S. P., & Prescott, T. J. (2022). Scaffolding layered control architectures through constraint closure: Insights into brain evolution and development. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *377*(1844). https://doi.org/10.1098/rstb.2020.0519 (cit. on p. 11).

Winding, M., Pedigo, B. D., Barnes, C. L., Patsolic, H. G., Park, Y., Kazimiers, T., Fushiki, A., Andrade, I. V., Khandelwal, A., Valdes-Aleman, J., Li, F., Randel, N., Barsotti, E., Correia, A., Fetter, R. D., Hartenstein, V., Priebe, C. E., Vogelstein, J. T., Cardona, A., & Zlatic, M. (2023). The connectome of an insect brain. *Science*, *379*(6636), eadd9330. https://doi.org/10.1126/science.add9330 (cit. on p. 20).

Wu, Q., Wen, T., Lee, G., Park, J. H., Cai, H. N., & Shen, P. (2003). Developmental control of foraging and social behavior by the Drosophila neuropeptide Y-like system. *Neuron*, *39*(1), 147–161. https://doi.org/10.1016/S0896-6273(03)00396-9 (cit. on pp. 20, 144).

Wystrach, A., Lagogiannis, K., & Webb, B. (2016). Continuous lateral oscillations as a core mechanism for taxis in Drosophila larvae. *eLife*, *5*. https://doi.org/10.7554/elife.15504 (cit. on p. 129).

Xia, S., & Tully, T. (2007). Segregation of odor identity and intensity during odor discrimination in Drosophila mushroom body. *PLoS Biology*, *5*(10), 2398–2407. https://doi.org/10.1371/journal.pbio.0050264 (cit. on p. 9).

Yaeger, L. S. (1997). Computational genetics, physiology, metabolism, neural systems, learning, vision, and behavior or polyworld: Life in a new context. https://api.semanticscholar.org/CorpusID:14599677 (cit. on p. 5).

Zador, A., Escola, S., Richards, B., Ölveczky, B., Bengio, Y., Boahen, K., Botvinick, M., Chklovskii, D., Churchland, A., Clopath, C., DiCarlo, J., Ganguli, S., Hawkins, J., Körding, K., Koulakov, A., LeCun, Y., Lillicrap, T., Marblestone, A., Olshausen,

B., ... Tsao, D. (2023). Catalyzing next-generation Artificial Intelligence through NeuroAI. *Nature Communications*, *14*(1), 1597. https://doi.org/10.1038/s41467-023-37180-x (cit. on p. 138).

Zhu, M. L., Herrera, K. J., Vogt, K., & Bahl, A. (2021). Navigational strategies underlying temporal phototaxis in Drosophila larvae. *Journal of Experimental Biology*, *224*(11). https://doi.org/10.1242/JEB.242428 (cit. on p. 21).

Ziemke, T. (2008). On the role of emotion in biological and robotic autonomy. *BioSystems*, *91*(2), 401–408. https://doi.org/10.1016/j.biosystems.2007.05.015 (cit. on p. 12).