Behavioural principles underlying navigational decision-making in *Drosophila melanogaster* larvae

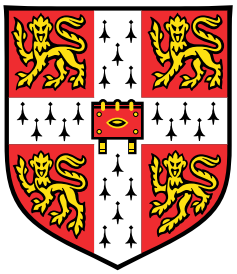

Elise Colette Croteau-Chonka

Department of Zoology
University of Cambridge

This thesis is submitted for the degree of
Doctor of Philosophy

Trinity College August 2021

Preface

This thesis is the result of work carried out within the Howard Hughes Medical Institute Janelia Research Campus, the University of Cambridge Department of Zoology, and the Medical Research Council Laboratory of Molecular Biology. My research was funded by a C T Taylor Cambridge International Scholarship, with overrun funding support from the Cambridge Trust and Trinity College.

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in this Preface and specified in the text. It is not substantially the same as any work that has already been submitted before for any degree or other qualification except as declared in this Preface and specified in the text. It does not exceed the prescribed 60,000 word limit for the Biology Degree Committee.

Some of the work I present in this thesis is the result of collaboration with Kristina Klein, Dr Ann Hermundstad, Dr Carey Priebe, Dr Jean-Baptiste Masson, Dr Lakshmi Narayan, Dr Peter Polidoro, Dr Christopher McRaven, Dr Michael Winding, Monti Mercer, and Dr Brandi Sharp as specified in the relevant passages of the text. In Chapter 6, I enclose a modified version of the preprint manuscript 'Serotonergic Neurons Mediate Operant Conditioning in *Drosophila* Larvae' (Klein et al., 2021), for which I am co-first author alongside K Klein. For readability, I have moved the associated methods and materials to Section 2.2.

Elise Colette Croteau-Chonka
August 2021

Summary

Title: Behavioural principles underlying navigational decision-making in *Drosophila melanogaster* larvae
Author: Elise Colette Croteau-Chonka

An animal's survival depends on timely decisions informed by sensory information. Studies in humans and large model organisms have elucidated auxiliary roles of large brain regions in the evolution of such perceptual decisions. What remains challenging is acquiring a detailed understanding of the underlying neural mechanisms at a synaptic level and across entire brain circuits.

The *Drosophila melanogaster* larva is an apt model system for probing the mechanisms of decision-making given its rich behavioural repertoire, small nervous system, genetic tractability, and available neuronal wiring diagrams. Taking inspiration from the application of two-alternative forced choice (TAFC) tasks to study perceptual decision-making in other model systems, I employed a closed-loop system to optogenetically activate larval nociceptive neurons based on the direction of precisely detected lateral head sweeps (i. e. casts). I sought to uncover the behavioural computations driving the stereotyped larval navigation sequence comprising repeated head casts followed by crawling in a new direction.

I found that in control conditions where stimulus intensity is identical between left and right casts, the percentage of larvae that stop exploration and crawl in the direction favourable for survival (i.e. toward the first stimulated direction) significantly increases with number of casts. However, in experimental conditions where the aversive stimulus differs between sides, the percentage that accept the correct side (i. e. lower intensity) increases more significantly with cast number. When controlling for integrated intensity across casts, I observe a higher fraction of larvae accepting the lower intensity stimulus in experimental conditions compared to controls. These results suggest a mechanism of side-to-side comparison and possible sensory evidence accumulation that facilitates improved decision-making.

In this thesis, I introduce the construction and implementation of two computational models for comparison to the larval behaviour trajectories. Both models reflect features of the experiment paradigm, though they differ in their assumptions about how the larva uses information from its environment to guide the acceptance or rejection of a given cast. The resulting predictions I generated about larval behaviour capture some, but not all, qualitative signatures within both the experimental and control datasets. I explore avenues for future model investigation and collection of additional behavioural data in order to draw more definitive mechanistic conclusions.

While powerful, the closed-loop system I employed tracks only a single larva at a time. Transitioning my sensory discrimination task to a high-throughput system would be advantageous not only to expand the investigation of other stimulus levels but also to screen stimuli of different valences or from other sensory modalities. In this thesis, I detail my contributions to the development, validation, testing, and experimental application of a new tracking system that is capable of behaviour detection and closed-loop optogenetic and thermogenetic stimulation of 16 larvae simultaneously. This facilitated the first observations of operant conditioning in the *Drosophila* larva in which the animal successfully adapted its casting behaviour following repeated coupling with reward presentation. Although operant learning occurs over a longer time scale than perhaps what is required for perceptual decision-making, the two tasks are related in creating an association between the animal's body posture and available sensory information. Together, my work on the sensory discrimination task, behavioural modeling, tool development, and analysis of the operant learning results lays a foundation for future investigation of decision-making behaviour in *Drosophila* larvae, with implications for further understanding the circuit mechanisms underlying larval taxis, learning, and memory.

To all the graduate students who did not get the chance to succeed as they could have.

Acknowledgements

My path in higher education leading to this moment has taken many turns, some planned and others wholly unexpected. These twelve years have spanned six institutions, four states, and two countries. To name every single individual or group who gave me their support, mentorship, guidance, friendship, or love would be a monumental task. I instead take this opportunity to highlight some special ones.

First, I extend many thanks to my supervisor, Dr Marta Zlatic, for her helpful discussions, valuable writing and presentation feedback, and generous laboratory resources. Thank you for opening the door to Janelia and Cambridge. Thank you to Dr Matthias Landgraf and my advisors, Dr Greg Jefferis and Dr Berthold Hedwig, for continued professional feedback and personal guidance. Thank you, also, to Dr Albert Cardona, whose response to my inquiry about graduate school opportunities helped ease the most challenging academic transition I've ever faced. You not only vouched for me to Marta but also continued to see my potential when I wasn't always convinced of it myself. I also want to thank Kristina Klein, Dr Ann Hermundstad, Dr Carey Priebe, Dr Jean-Baptiste Masson, Dr Lakshmi Narayan, Dr Peter Polidoro, Dr Christopher McRaven, Dr Michael Winding, Monti Mercer, and Dr Brandi Sharp who contributed their scientific knowledge to various parts of this project. I'm also grateful for the many members of the Zlatic and Cardona labs who were a part of this unforgettable experience not only as comrades in science but also as good friends.

Graduate school is challenging, particularly when you're more than 3000 miles from home. I am immensely grateful to the Trinity College BA Society for helping us international postgraduate students find home in Cambridge. I never thought it possible to have so much fun in such a short period of time. I'm reminded of these lyrics from a favourite song of mine: "There's not a word yet; for old friends who've just met." I will forever be lovingly tied to Trinity because of you.

Like many other graduate students I know, I never thought I'd be finishing my PhD during a pandemic. I took solace in frequent Zoom check-ins with close friends and loved ones, but I must highlight those with Margo Heston in particular. Our long-distance

co-working started out of mutual need for accountability and structure while stuck at home, but our relationship grew and you became a trusted friend and confidant. Thanks for sharing your screen with me.

When I auditioned for Ensemble Companio in the fall of 2015, my path had just taken one of those unexpected turns, and I was seeking a space to heal and find my voice again. I am honoured that Joe Gregorio chose me to join this choral family. Making music with EC over the years has been undeniably joyous and cathartic. Thank you.

And most importantly, I am forever indebted to my parents, my two older siblings, and my partner, Brian, who stood by me every step of the way. Thank you for your boundless love, incredible patience, and constant encouragement. I know this hasn't always been easy on you. But I also know I wouldn't have wanted to do it (or, honestly, been capable of doing it) without you.

Table of contents

List of figures

List of tables

Glossary

C. elegans *Caenorhabditis elegans*.

E. coli *Escherichia coli*.

CNS central nervous system.

CS conditioned stimulus.

CSD contralaterally projecting serotonin-immunoreactive deutocerebral.

DDM drift diffusion model.

DMD digital micromirror device.

EM electron microscopy.

FPGA field-programmable gate array.

GFP green fluorescent protein.

GUI graphical user interface.

IR infrared.

KC Kenyon cell.

LNM linear-nonlinear model.

MB mushroom body.

MBIN MB input neuron.

MBON MB output neuron.

mdIV multidendritic class IV sensory neuron.

MWT Multi-Worm Tracker.

PI performance index.

PKC protein kinase C.

SEZ subesophageal zone.

TAFC two-alternative forced choice.

UCB Upper-Confidence-Bound.

US unconditioned stimulus.

VNC ventral nerve cord.

Chapter 1

Introduction

An animal's survival depends on making informed decisions from sensory information. One behavioural strategy for this perceptual decision-making could be strict reinforcement of actions based on outcomes, but such highly strengthened associations with stimuli can be maladaptive for learning in dynamic environments (Cohen et al., 2007; Reitich-Stolero et al., 2019). Animals must be flexible enough to manage inherently ambiguous sensory information that itself may change over time. Both vertebrates (Brunton et al., 2013; Feng et al., 2009; Kira et al., 2015; Shadlen and Shohamy, 2016) and invertebrates (DasGupta et al., 2014; Groschner and Miesenböck, 2019; Tanimoto and Kimura, 2019) repeatedly sample stimuli to reduce uncertainty surrounding features of the environment. Top-down research facilitates translation of these observations into testable predictions about the cognitive algorithms guiding perceptual decisions and the neural mechanisms through which those computations occur (Gerstner et al., 2012; Kriegeskorte and Douglas, 2018).

Alongside egg-laying and mate selection, foraging remains a prime context for studying perceptual decision-making in naturalistic environments. Foraging studies span animal taxa, from humans (Kolling et al., 2012) and various bird species (McNamara et al., 2006; Valone, 2006), to bees (Foley and Marjoram, 2017) and the nematode *Caenorhabditis elegans* (Calhoun et al., 2014; Flavell et al., 2013). Perceptual decision-making is also probed through psychophysics, with a common assay being the two-alternative forced choice (TAFC) task. Perhaps the most well-known formulation is one in which a monkey must discriminate coherent motion within a visual field comprising other randomly moving dots. The monkey reports its belief about which direction the coherent group is moving by saccading to one of two targets (Feng et al., 2009; Shadlen and Newsome, 1996). Variations include stimulus discrimination in visual (Kira et al.,

2015) and vibrotactile (Romo et al., 2002) sequences. Rodent studies highlight TAFC versatility across sensory modalities (Shadlen and Kiani, 2013) including olfaction (Gold and Shadlen, 2007; Kepecs et al., 2008; Uchida and Mainen, 2003) and audition (Brunton et al., 2013; Erlich et al., 2011; Marbach and Zador, 2016; Otazu et al., 2009).

Controlled laboratory experiments across model organisms have been integral to advancing investigation of perceptual decision-making toward comparative study of underlying neural mechanisms. In general, findings across large model organisms suggest an auxiliary role of large brain regions in forming these decisions. In monkeys performing TAFC, graded neuronal firing reflects decision formation dynamics for visual (Smith and Ratcliff, 2004) and vibrotactile (Romo et al., 2002) stimuli. In both monkeys and rodents, cortical activity in motor planning areas during TAFC is consistent with evidence accumulation to a threshold (de Lafuente et al., 2015; Eckhoff et al., 2008; Feng et al., 2009; Gold and Shadlen, 2001; Kira et al., 2015; Li et al., 2016; Li and Krishnamurthy, 2015; Shadlen and Newsome, 1996). Relationships between smaller neuron groups are, however, poorly understood at this scale (Hanks and Summerfield, 2017). A full mechanistic understanding of perceptual decision-making within a single organism, from behaviour down to the synapse, has yet to be described (Tanimoto and Kimura, 2019).

1.1 *D melanogaster* larval navigation

The *Drosophila melanogaster* larva is an apt model system for study given its rich behavioural repertoire (Green et al., 1983; Lahiri et al., 2011; Ohyama et al., 2013; Takagi et al., 2017), small nervous system (approximately 10,000 neurons), genetic tractability for selective manipulation of individual neuron types (Brand and Perrimon, 1993; Duffy, 2002; Jenett et al., 2012; Klapoetke et al., 2014; Luo et al., 2008; Simpson and Looger, 2018), and available brain synaptic-level wiring diagram (Berck et al., 2016; Eichler et al., 2017; Fushiki et al., 2016; Ohyama et al., 2015; Schneider-Mizell et al., 2016). These larvae also exhibit a stereotyped behavioural sequence as they navigate their environment: periods of forward crawling (runs) are interrupted by reorientation manoeuvres in which head sweeps (casts) once or more to either side of the body precede crawling in a new direction (turns). Such changes in larval heading are considered navigational decision points (Fig. 1.1; Lahiri et al., 2011; Luo et al., 2010; Riedl and Louis, 2012), making reorientation manoeuvres ideally suited for investigating perceptual decision-making. In homogeneous environments, the larval ventral nerve

cord (VNC) is sufficient for executing and alternating between runs and reorientation manoeuvres. In dynamic environments, the brain, subesophageal zone (SEZ), and VNC modulate these behaviours to facilitate movement away from potential threats toward more favourable living conditions (Berni et al., 2012).

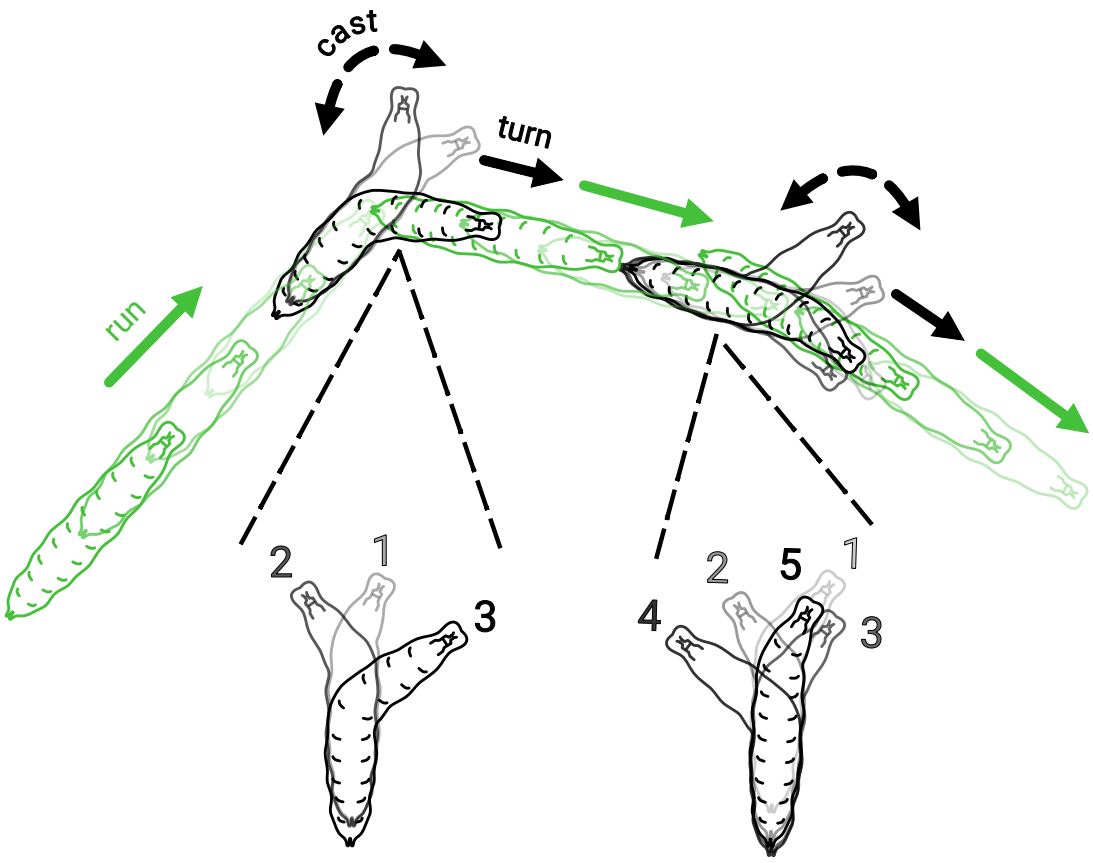

Fig. 1.1 Time series of larval navigation. Periods of forward peristaltic motion (runs; green) are interrupted by reorientation manoeuvres (casts followed by a turn; black). Arrows indicate direction of motion. Reorientation manoeuvres comprise one or several head casts (annotated below the trajectory, with numbers corresponding to the first, second, etc. cast in the sequence), followed by crawling in a new direction. Although both reorientation manoeuvres show the larva alternating sequential head casts between the left and right sides of its body, larvae are also known to cast repeatedly to one side before sweeping their head to the other (behaviour not shown). The heading direction set by the final cast in a reorientation manoeuvre dictates the direction of the subsequent run, making this stereotyped behavioural sequence a form of navigational decision-making. Larva schematic inspired by Gjorgjieva et al. (2013).

One such dynamic environment that has been carefully studied in *Drosophila* is the smoothly varying sensory gradient. The success with which larval *Drosophila* traverse these environments has encouraged further investigation into underlying trajectory-guiding strategies. Consider three features of the stereotyped reorientation behaviour sequence: the average duration of the run preceding a reorientation manoeuvre (manoeuvre frequency), the magnitude of change in heading direction between the start and end of a reorientation manoeuvre (manoeuvre size), and the

acceptance of a head cast when oriented orthogonal to the gradient (manoeuvre direction). Larvae have been shown to bias the frequency, size, and direction of reorientation manoeuvres to navigate favourably in gradients of odour (chemo; Gershow et al., 2012), temperature (thermo; Klein et al., 2015; Luo et al., 2010), light (photo; Humberg et al., 2018; Kane et al., 2013), and wind (anemo; Jovanic et al., 2019). Although *Escherichia coli* and *C. elegans* also modulate run length based on stimulus intensity, this method alone results in indirect orientation toward attractants (kinesis). Through the addition of head casting, *Drosophila* larvae take advantage of local sensory information to orient themselves directly toward appetitive stimuli (taxis).

But what function, if any, do separated sensory organs serve in directed navigation? Adult *Drosophila* appear to rely solely on simultaneous spatial comparison of bilateral inputs (tropotaxis; Gomez-Marin and Louis, 2012). In the larva, this reliance seems to be modality-dependent. In both chemotaxis and thermotaxis, larvae neither significantly change their heading direction during runs nor bias the first head cast of a reorientation manoeuvre to the favourable direction (Gershow et al., 2012; Luo et al., 2010). The absence of steering or biased cast initiation suggests that larvae do not use tropotaxis to spatially compare odour or temperature stimuli while crawling. Larvae instead move their sensory organs through sequential head casts during reorientation manoeuvres, enabling temporal stimulus comparison to drive navigational decisions (klinotaxis). Successful larval chemotaxis and thermotaxis using unilateral sensory input further supports this conclusion (Gershow et al., 2012; Gomez-Marin et al., 2011; Klein et al., 2015; Louis et al., 2008). In contrast, the dynamics of larval phototaxis and anemotaxis seem dependent on both temporal and spatial comparisons between stimuli. In wind gradients, larvae do not steer but they do bias the first head cast of a reorientation manoeuvre (Jovanic et al., 2019). In light gradients, binocularity provides spatial information for steering (Humberg et al., 2018), though biased cast initiation appears absent (Kane et al., 2013).

Despite detailed behavioural investigation of larval navigation strategies through taxis behaviour, gaps remain in our knowledge of precisely how temporal information acquisition leads to successful navigation. Notably, the function of repeated head casts within a reorientation manoeuvre has not been well characterised. Previous work in taxis has observed larvae casting as little as once and as often as seven times per manoeuvre (Gomez-Marin and Louis, 2012; Luo et al., 2010). Despite this variability in manoeuvre length, statistics describing the cast-to-run transition have, to my knowledge, only focused on cast acceptance as a function of direction (to the favourable or

unfavourable side) and not a cast's numerical placement within the sequence (Gershow et al., 2012; Kane et al., 2013; Luo et al., 2010). What behavioural algorithm mediates transitions between repeated cast rejections leading to the final cast's acceptance? It has been proposed that sensory systems in insects may accumulate evidence for different alternatives and use coupled inhibition to improve decision-making performance (Barron et al., 2015). I hypothesise that larvae accumulate sensory information with each successive head cast, reducing their uncertainty about stimulus intensities and improving the likelihood that the direction in which they turn is associated with favourable conditions. Such a strategy could also underlie the increased head casting observed when stimulus contrast is low (Gomez-Marin et al., 2011). Such knowledge would have implications for our understanding of the circuit mechanisms underlying not only larval taxis but also decision-making associated with learning and memory.

1.2 Explanation of chapters

The aim of my thesis is to detail my investigation of the behavioural computations underlying repeated head casts during stereotyped reorientation manoeuvres in *Drosophila melanogaster* larvae. In Chapter 2, I detail all methodology I followed and materials I required to conduct this research. In Chapter 3, I describe my design for a sensory discrimination task to probe features of larval perceptual decision-making within a single-larva closed-loop optogenetic activation system. In that same chapter, I interrogate the data I acquired from larvae engaging with this task and propose an underlying behavioural mechanism. In Chapters 4 and 5, I outline the conceptual construction and mathematical implementation of two different computational models and evaluate their ability to explain features of the larval behavioural data. Chapter 6 is the preprint manuscript I co-authored with Kristina Klein (methods and materials moved to latter half of Chapter 2) detailing how we built a multi-larva closed-loop system and observed, for the first time, operant learning in the *Drosophila* larva. I begin the chapter with a brief preamble outlining how this work complements the investigation I describe in Chapters 3, 4, and 5. The final chapter of my thesis (Chapter 7) is a discussion of my results and their contribution to our understanding of larval behaviour as a whole. I outline limitations of my approach and suggest various avenues for future investigation.

Chapter 2

Methods and materials

2.1 Larval sensory discrimination task

2.1.1 Single-larva closed-loop stimulation system

Hardware and software frameworks

The single-larva tracker electronics and hardware were nearly identical to the tracker described in Schulze et al. (2015), with notable exceptions to the camera and backlight which I detail here. Engineers from Janelia Research Campus' Instrument Design and Fabrication team assembled this system. The behaviour arena comprised a layer of agarose sitting atop a fixed 45 cm x 45 cm glass plate. Above and below the arena sat a pair of motorised linear slides (#T-LSR450B, Zaber Technologies), with single slides mounted perpendicular to one another. Units oriented along the same axis were daisy-chained via a 6-pin mini din cable and shared the same power supply. This arrangement enabled movement of mounted components to any (x, y) location within both slides' travel range specifications (total coverage area was 34 cm x 38 cm) (Fig. 2.1a). The linear slides interfaced with the host computer via a USB connection and received (x, y) position updates at a rate of 4 Hz (Fig. 2.1b).

Mounted to the slides above the behaviour arena was a 2048 x 2048 resolution camera (Grasshopper3 #GS3-U3-41C6NIR-C, Point Grey Research) (Fig. 2.1a) that captured images at 20 Hz (Fig. 2.1b). A long-distance microscope (Model KC/S VideoMax with IF2 objective, Edmund Optics) was c-mounted to the camera. The camera-to-arena coordinate relationships were -1132 pixels (x) / mm (x), 70.56 pixels (x) / mm (y), 71.45 pixels (y) / mm (x), and 2831 pixels (y) / mm (y). The camera

interfaced with the host computer via a USB 3.0 connection and the aerial images were displayed in real time on a web-based graphical user interface (GUI) (Fig. 2.1b).

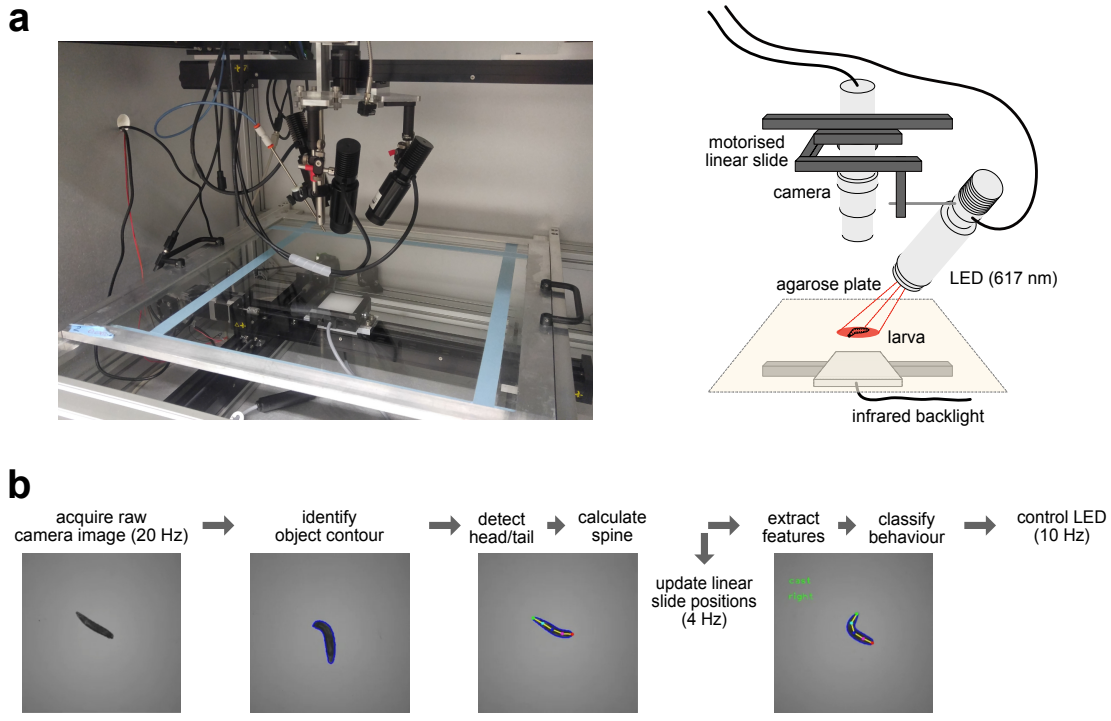

Fig. 2.1 Single-larva closed-loop stimulation system. A camera and an infrared backlight, both mounted on motorised linear slides, were used to track the real-time behaviours of a single larva moving along a stationary agarose plate. When a chosen behaviour was detected, the red LED turned on, stimulating the larva. **a.** System hardware was nearly identical to that described in Schulze et al. (2015) (see text for details). **b.** Software framework. Still images from the graphical user interface show aerial views of a single *Drosophila* larva in the behaviour arena. Processing raw camera images included identifying and drawing a 200-point contour (blue outline) around the larva's body. This enabled tracking throughout the experiment. The larval spine was defined as equally distributed landmarks along the larva's midline between the head (green dot) and tail (red dot). Using the `neck` landmark (dark blue dot) as a guide, the linear slides performed position updates to keep the larva centred in the camera's field of view. The motion of landmarks over time (both individually and with respect to one another) was the basis for extracting motion, shape, and velocity features. Features were combined to train behaviour classifiers. Software protocols combined detected behaviours (displayed in green words) with user-defined experiment parameters to control the timing and intensity of LED stimulation.

Also mounted to the slides above the behaviour arena was a 617 nm LED (#PLS-0617-030-S, Mightex Systems) which I equipped with a neutral density 3.0 filter (Fig. 2.1a). The LED was connected to a universal driver with a serial data connection to the host computer and a 12 bit output driving current of up to 1000 mA (#SLC-xx04-US, Mightex Systems) at an update rate of 10 Hz (Fig. 2.1b). The centre of the LED projection was aligned with the centre of the camera image to facilitate targeted optogenetic stimulus presentation. To maintain the integrity of behaviour detection, stimulus presentation, and optogenetic experimentation, all hardware was housed inside a light-tight enclosure. A 2 in x 2 in 880 nm LED backlight (#BL0202-880IC, Advanced

Illumination) was mounted to the linear slides below the behaviour arena to allow the camera to continuously track the larva in the dark (Fig. 2.1a).

Dr Peter Polidoro wrote the hardware communication and operation software of the single-larva tracker using a Robot Operating System (ROS) framework. Running in an Ubuntu Linux environment, the individual software modules communicated with one another through a TCP connection. Custom experiment protocols written by myself and Kristina Klein combined classified behaviour with user-defined stimulus parameters from the web-based GUI to control LED timing, intensity, and duration. The tracker output files contained frame-by-frame data associated with feature calculation, behaviour classification, and stimulus presentation. I modified the operation software to save experiment parameters and other metadata (e. g. larval genotype) as a header within each output file. I also debugged and further developed the default user interface to increase flexibility in choosing experiment protocols and recording metadata without repeatedly needing to modify the software.

Object contouring, landmark detection, and behaviour classification

Dr Jean-Baptiste Masson wrote the software for object contouring, landmark detection, and behaviour classification. Here I describe these components, emphasising the modifications that I made to the software either by myself or in collaboration with Dr J-B Masson and/or K Klein. The software processed the raw camera images with an inverse binary threshold and then identified the largest visible contour (i. e. the larva) (Fig. 2.1b). With the infrared backlight positioned beneath the behaviour arena, K Klein and I established that the transparency of the glass plate together with the concentration of the agarose substrate influenced contour detection stability. Together, we iteratively tested different camera shutter speeds to determine which yielded contours that not only accurately represented the larval outline and but also did not flicker between frames. Dr J-B Masson's software kept the number of contour points identical between frames using online Fourier decomposition and reconstruction (Masson et al., 2020). Larval head and tail positions were extracted from the reconstructed contour's sharpest and second-sharpest internal angles (Fig. 2.1b). These calculations alone were prone to error, with false detection during the first frame or as a consequence of large-angle head casts. Dr J-B Masson, K Klein, and I designed preventive proximity measures and a corrective vote system to improve online head and tail detection accuracy between image frames (for algorithm details, see Section 2.2.1 and Fig. 2.3).

Dr J-B Masson's software used head and tail positions to divide the contour in half and calculate the larval midline (i. e. spine) as in Swierczek et al. (2011). Three additional landmarks were isolated from equally spaced locations along the spine. The `neck` landmark was used to update the linear slide positions, keeping the moving larva centred within the camera's field of view (Fig. 2.1b and Fig. 2.4). As in the high-throughput tracker described in Section 2.2.1, the contour, spine, and landmarks were used to calculate real-time larval motion, shape, and velocity features on each camera image. Temporal smoothing and differentiation by convolution removed high-frequency noise.

Together, Dr J-B Masson, K Klein, and I developed behaviour classifiers for this single-larva tracker that consisted of linear thresholds on the feature space and/or supervised machine learning with a neural network algorithm (Bishop, 2006; Hagan et al., 2014) (Fig. 2.1b). I contributed to training classifiers and manually assessing their performance using a custom MATLAB (MathWorks, Natick, Massachusetts) GUI equipped with software functions from the Neural Network, Deep Learning, and Statistics and Machine Learning toolboxes. This GUI was developed by Dr J-B Masson with necessary modifications implemented by K Klein. Three classifiers integral to my sensory discrimination task included `cast`, `ball`, and `forward`.

Dr J-B Masson's existing `cast` classifier comprised a two-layer neural network trained on two features that characterised larval body axis alignment and overall shape. I empirically determined a strict threshold on the former to ensure sufficiently large head cast amplitudes ($\mathtt{s} < 0.82$; Section 2.2.1; Masson et al., 2020). Larval direction (`left` or `right`) was computed using a single threshold on the angle between the head and lower body (`asymmetry`; Section 2.2.1). If exponential smoothing of the resulting boolean `left` and `right` values exceeded a predetermined threshold, the respective classifier was set to `true`. This separate `left` or `right` classification was combined with the `cast` classifier to detect either `left cast` or `right cast`. I manually assessed `left cast` and `right cast` performance based on 358 events from 30 minutes of video data across 46 larvae. Precision was 76.0% and recall was 97.2%. Online false positives were largely attributed to `cast` flickering off after initial detection (3.0 flickers per minute). These were manually aligned offline with the closest true positive cast event (see Section 2.1.4).

A larva would occasionally make a high amplitude cast, touching its head to its tail. This "ball" shape transiently warped the calculated contour and caused head and tail detection to rapidly flicker (Fig. 2.3). This flickering often led to incorrect head tail

detection as the larva straightened. K Klein and I trained a two-layer neural network classifier for `ball` behaviour based on features that defined body axis alignment, overall shape, and ratios of the contour's perimeter and enclosed area to those of its convex hull. We established that when a `ball` was detected, the previous frame's `left cast` or `right cast` behaviour classifications must be maintained. We set these to update only 1.5 s after `ball` detection stopped. False head and tail detection tallies were also reset during the ball to reduce flickering once real-time `left` and `right` detection resumed (for details, see Section 2.2.1). I manually assessed the performance of this `ball` classifier (produced 0.03 left/right cast errors per minute).

K Klein developed a `forward` crawl classifier with thresholds on larval body axis alignment, overall shape, and tail movement. To reduce detection errors across larvae that moved at different speeds, I added to this classifier a unique tail movement threshold for each larva based on it's average tail movement during the 30 s waiting period preceding each stimulation protocol. I manually assessed this new `forward` classifier's performance based on 809 events from 30 minutes of video data across 46 larvae. Precision was 95.3% and recall was 90.7%.

2.1.2 Fly strains and larval rearing

All larvae I used for experimentation were early third instar (72-84 hours after egg laying) *Drosophila melanogaster* with the *CsChrimson* transgene (Klapoetke et al., 2014) expressed in nociceptive multidendritic class IV neurons (ppk1.9; Ainsley et al., 2003; Hwang et al., 2007; Xiang et al., 2010) under the Gal4-UAS system (Brand and Perrimon, 1993). Virgin adult females of the genotype *20xUAS-CsChrimson-mVenus trafficked in attp18* (Stock Number 55134, Bloomington Drosophila Stock Center) were crossed in a 3:1 ratio with healthy males of the genotype w^{1118}*;ppk1.9-Gal4* (Ainsley et al., 2003). Monti Mercer and Dr Brandi Sharp occasionally helped set up crosses in behavioural cages. I collected eggs in the dark for four hours at 25℃ and 50% humidity on fly food with dry yeast added to the surface to encourage egg laying. The food was prepared by the Janelia Research Campus' Media Prep Facility following the standard cornmeal recipe detailed in Section 2.2.3. Following egg collection, I incubated plates in the dark for three days at 25℃ and 50% humidity. The short collection window ensured that all larvae were stage-matched to early third instar.

2.1.3 Experiment conditions and procedures

I conducted experiments at 25℃ and 50% humidity. Prior to each day of experiments, I poured 1% agarose (Fisher Scientific) (Apostolopoulou et al., 2014b) on the glass plate and allowed it to harden . I left all larvae in the dark on the fly food until they were needed for an experiment. To easily extract larvae for experiments, I immersed small amounts of fly food in a separate petri dish of water. Using a brush, I lifted each larva out of the water using a brush, then gently dried and placed them on a room-temperature 4% agar petri dish. I kept this dish in the dark to allow the larva to recover from extraction for a minimum of five minutes. Just prior to the experiment, I gently picked up the larva from the petri dish and placed in the centre of the behaviour arena on top of the cooled and set agarose. I then closed the rig door.

I wrote the Python software code for my larval sensory discrimination task which required triggering the presentation of different stimulus intensities as a function of detected `cast` direction (Fig. 3.1a). 30 s without light preceded each experiment, allowing for behaviour detection stability and larval acclimation to the substrate. Immediately following this waiting period, light stimuli triggering was withheld until the larva completed four consecutive `forward` crawls, each within 2 s of one another and uninterrupted by a `cast`. I established this criterion to prevent partial stimulation of casting events that started during the waiting period and continued into the experiment window. This also distinguished larvae that may have been more motivated to solve the task: if larvae were too sluggish and failed to meet the crawling criterion within the experiment's first minute, I discarded them.

Individual larvae were tested under one of five predetermined stimuli conditions, each constituting a pair of 617 nm red light intensities. I measured the following intensities with a Thorlabs digital optical power meter equipped with a compatible photodiode sensor:

- $E_{L,M}$ (18.72 µW/cm^2, 19.78 µW/cm^2)
- $E_{L,H}$ (18.72 µW/cm^2, 20.90 µW/cm^2)
- $C_{L,L}$ (18.72 µW/cm^2, 18.72 µW/cm^2)
- $C_{M,M}$ (19.78 µW/cm^2, 19.78 µW/cm^2)
- $C_{H,H}$ (20.90 µW/cm^2, 20.90 µW/cm^2)

$E_{L,M}$ and $E_{L,H}$ represent *experimental* conditions under which a different stimulus intensity was triggered when the larva cast in one direction versus the other. Importantly, the two conditions shared the same lower value. $C_{L,L}$, $C_{M,M}$, and $C_{H,H}$ represent *control* conditions under which the same stimulus intensity was triggered regardless of the direction in which the larva cast. These controls accounted for the three unique stimulus values present in the experimental conditions (L: low, M: medium, H: high).

Irrespective of direction (`left` or `right`), I set each larva's first head cast after meeting the crawling criterion to trigger the presentation of the lower intensity stimulus in the condition (Fig. 3.1a). By remembering the first cast's direction, the software code ensured that the larva casting to the alternate side would trigger presentation of the higher intensity stimulus in the condition. Regardless of intensity, the red LED remained on for the entirety of each detected `left` or `right` `cast`, including during `ball` behaviour.

My criterion for a larva ending a reorientation manoeuvre was completion of two forward peristaltic waves in the direction set by the last head cast, with neither wave requiring tail movement nor motion along a straight body axis (for details, see Chapter 3). I visually evaluated this criterion in real time and stopped the experiment manually after the larva completed a single reorientation manoeuvre (Fig. 3.1a).

2.1.4 Assessing larval performance

All tracker output files were post-processed with a custom MATLAB pipeline built by Dr J-B Masson that regenerated contour and spine data for subsequent visualisation in a custom MATLAB user interface. This interface, initially designed by Dr J-B Masson and modified by K Klein, supported playback of both the rendered larval images and the associated time-series graphs of feature values and detected behaviours. For each larva, the full experiment was visually validated and the number of casts the larva performed in the reorientation manoeuvre, the direction (`left` or `right`) of the manoeuvre's last cast, and the image frame numbers corresponding to the beginning and end of the manoeuvre were manually recorded in Excel. Because occasional flickering of the `cast` classifier translated to ON/OFF flickering of the LED, visual validation and manual recording were necessary to merge all LED signal events associated with the same larval head cast. Image frame numbers corresponding to the beginning and end of the reorientation manoeuvre (see criterion above) were also manually recorded for each larva. Larval data were excluded from further analysis if they met any of the following criteria that I defined:

- larva had a potential motor impairment exhibited while crawling (e. g. the larva maintained an "S" body shape while crawling)
- larva exhibited rearing behaviour by lifting its head off the substrate
- larva did not appear to react to the first stimulus presentation (i. e. larva moved its head such that a cast was detected, but the motion appeared associated with a run rather than the beginning of a reorientation manoeuvre)
- direction of the reorientation manoeuvre's last cast was ambiguous
- larva crawled straight at the end of the reorientation manoeuvre rather than accepting a left or right cast
- LED stayed off or triggered improperly during a reorientation manoeuvre
- delays or unresolved flips in head and tail detection

The tracker output and manual review files were then post-processed through a custom R pipeline that I wrote. The pipeline first selected larvae that alternated sides with each successive head cast, excluding those that cast repeatedly to one side before either casting again or ending the reorientation manoeuvre. The last head cast of a reorientation manoeuvre determined the larva's turn direction and I considered this to be the larva's decision (Fig. 3.1a).

For the larval populations within each stimuli condition, I calculated multiple performance metrics which I detail in Section 3.2.2. In brief, the probability of accepting a given head cast ($P(\mathrm{accept})$) equals the number of larvae that accepted that cast conditioned on the number of larvae that performed that cast (whether or not they accepted it). The probability of making the correct decision on a given head cast ($P(\mathrm{correct})$) was calculated similarly, except a correct decision was defined as accepting the less noxious stimulus or rejecting the more noxious stimulus. Because I designed the experiment to start with the less noxious stimulus and selected larvae that performed consecutive casts by alternating sides, a correct decision for each larva translated to accepting odd-numbered casts and rejecting even-numbered casts. I applied the same definition to control conditions for which the stimulus intensity was equivalent on both sides.

I used the Clopper-Pearson Exact method to calculate 95% confidence intervals for $P(\mathrm{accept})$ and $P(\mathrm{correct})$. This method is preferred over the normally distributed error

approximation because it addresses the need for a more accurate binomial proportion confidence interval. Although Clopper-Pearson is more conservative than other approaches, it computes asymmetric intervals around the proportion, avoiding issues of overshoot when the sample probability lies near zero or one. I used binary logistic regression to quantitatively evaluate whether larvae are more likely to choose the less noxious side following more head casts and compared the resulting predictor variable coefficients between conditions (for details, see Section 3.2.2).

In my R pipeline, the duration of each cast within each larval manoeuvre was calculated as the total time for which the LED stimulus was on during the cast. Using this data and scaled values of the corresponding stimulus intensities on each cast, I worked alongside Dr Ann Hermundstad to calculate three separate acceptance probabilities for larvae that made more than one cast in their reorientation manoeuvre (for details, see Fig. 3.4 and Section 3.2.2).

2.1.5 Bayesian inference model

In Chapter 4, I explicitly detail mathematics descriptions of the generative model, inference calculation, decision process, and simplifying assumptions of a Bayesian inference model for decision-making in my sensory discrimination task. The following text outlines how the resulting model was implemented in MATLAB software code. Together, Dr A Hermundstad and I iteratively implemented various aspects of the model into a custom built MATLAB analysis pipeline. I performed all software refactoring, validation, testing, and finalising by myself. The *in silico* $\langle P(\mathrm{correct})\rangle$ trajectory calculated in the modeling code followed the closed-form solution described by Eqn. 4.13, with variables defined by Eqn.'s 4.5, 4.9, and 4.10. The modeling code employed MATLAB's constrained minimisation solver, `fmincon` (Optimization Toolbox), to find values for the parameters μ^L, μ^M, μ^H, and σ^2 that generated the best fit of this solution to user-supplied behavioural data. Each parameter's lower bound was set to 0 and was unbounded above. The relationships between parameters were defined through the inequalities shown in Section 4.7. μ^L was initialised by drawing a pseudorandom number from a uniform distribution over the open interval $(0, 1)$. μ^M was initialised as μ^L plus another pseudorandom number. μ^H was initialised as μ^M plus another pseudorandom number. This ensured that μ^L, μ^M, and μ^H were initialised in accordance with the numerical relationship between them. σ^2 was initialised by multiplying another pseudorandom number by a factor of two.

Generating the best fit of the model to the behavioural data required minimising the weighted objective function I defined in Eqn. 4.20. The modeling code ran 1000 different constrained optimisations following `fmincon`'s default interior-point algorithm. Each returned optimised values for each parameter and the solution to the objective function when evaluated at those values. Of the 1000 iterations, the optimised parameters that yielded the smallest fit error were used to build $\langle P(\mathrm{correct})\rangle$ trajectories for qualitative comparison to the behavioural data.

2.1.6 Acceptance pressure model

I explicitly detail mathematics descriptions of this model in Chapter 5. The following text outlines how the resulting model was implemented in MATLAB software code.

Together, Dr A Hermundstad and I implemented various aspects of the model into custom built MATLAB code. I performed all software refactoring, validation, testing, and finalising by myself. The *in silico* $P(\mathrm{accept})$ trajectory calculated in the modeling code followed the form of either Eqn. 5.2, 5.3, or 5.4 depending on the user's input preferences. The modeling code employed MATLAB's unconstrained minimisation solver, `fminunc` (Optimization Toolbox), to find values for the parameters r^M, r^H, m, and b that generated the best fit of the chosen equation to user-supplied behavioural data. Although none of the parameter values were constrained, each was initialised by drawing a pseudorandom number from a uniform distribution over the open interval $(0, 1)$.

Generating the best fit of the model to the behavioural data required minimising the weighted objective function I defined in Eqn. 5.5. The modeling code ran 1000 different unconstrained optimisations following `fminunc`'s default interior-point algorithm. Each returned optimised values for each parameter and the solution to the objective function when evaluated at those values. Of the 1000 iterations, the optimised parameters that yielded the smallest fit error were used to build *in silico* $P(\mathrm{accept})$ trajectories for qualitative comparison to the behavioural data.

2.1.7 Software availability

All software code is available upon request.

2.2 High-throughput operant conditioning

This section is the methods and materials from the preprint manuscript 'Serotonergic Neurons Mediate Operant Conditioning in *Drosophila* Larvae' (Klein et al., 2021), for which I am co-first author. I have modified the text to explicitly detail my contributions to method design, development, and testing as well as data collection and analysis. For the purposes of readability in relation to my sensory discrimination task (see Chapter 3), I have changed written instances of 'bend' to 'cast'. The remainder of the manuscript resides in Chapter 6.

2.2.1 High-throughput closed-loop tracker

Hardware set-up

A high-resolution camera (3072 x 3200 pixels) (#TEL-G3-CM10-M5105, Teledyne DALSA, Ontario, Canada) positioned above a 23 cm x 23 cm 4% agarose plate captured 8-bit greyscale images at 20 Hz. The agarose plate was illuminated from below by a 30 cm x 30 cm 850 nm LED backlight (#SOBL-300x300-850, Smart Vision Lights, Norton Shores, Michigan) equipped with intensity control (#IVP-C1, Smart Vision Lights, Norton Shores, Michigan). An 800 nm longpass filter (#LP800-40.5, Midwest Optical Systems, Palatine, Illinois) mounted on the camera blocked all visible wavelengths, including those used for optogenetics. When the agarose plate comprised most of the camera image, each pixel corresponded to a 72.92 µm diameter section of the plate. Dr Lakshmi Narayan and K Klein chose all hardware and performed the necessary functional testing of each component.

Each camera image was processed in parallel on both the host computer (#T7920, running Windows 10, Dell Technologies Inc, Round Rock, Texas) and an field-programmable gate array (FPGA) device (#PCIe-1473R-LX110, National Instruments, Austin, Texas), itself programmed by Dr L Narayan. LabVIEW 2017 (National Instruments, Austin, Texas) software extracted larval contours and interfaced with C++ software that performed real-time behaviour detection. The LabVIEW software controlled closed-loop optogenetic and thermogenetic stimulation in response to these detected behaviours. K Klein wrote the C++ software and Dr L Narayan coded the LabVIEW software.

Multi-animal detection and tracking

K Klein designed and validated the following tracking algorithms. Dr L Narayan implemented these algorithms in the LabVIEW code. Raw camera images were read by the FPGA at 20 Hz and then sent to the host computer. The LabVIEW process on the host computer then filtered out non-larval objects by combining background subtraction and binary thresholding. The remaining objects were each enclosed in a rectangular box of minimal size, with edges parallel to the camera image axes. The following criteria were defined to detect third-instar larvae within these boxes:

- Pixel intensity range (default 25–170): the minimum and maximum brightness values for pixels selected by binary thresholding (between 0 and 255 for an 8-bit image).
- Box side length (pixels) (default 6–100): the range of eligible values for width and height of each box.
- Box width + height (pixels) (default 12–200): the range of eligible values for the sum of each box's width and height.
- Box area (pixels) (default 300–900): the range of eligible values for the area of each box.

To track larvae over time, the host computer assigned a numerical identifier to each eligible object. Distance-based tracking with a hard threshold of 40 pixels maintained larval ID based on `centroid` position. Although identity was lost when larvae touched or reached the plate's edge, new IDs were generated when larvae matched detection criteria. For each of the largest 16 objects, the host computer sent a binary pixel pattern and location (defined as the centre of the box) to the FPGA. Since the host computer required more than 50 ms of run time for object detection, this process was not executed in every frame. On average, the FPGA received updated objects and their locations every three frames.

The FPGA extracted object contours in three steps. Within a 2 cm^2 region of interest around the object's centre, the FPGA first applied a user-defined binary threshold, then applied both vertical and a horizontal convolution with a 2 x 1 XOR kernel, and finally generated edge pixels by combining the results of the two convolutions using an OR operation. Contours were extracted from edge images using the Moore boundary tracing algorithm (Gonzalez and Woods, 2018) with three added error capture procedures. First,

if the algorithm yielded a contour that ended prematurely or contained small loops, the construction process could be reversed by up to 16 contour points to find an alternative contour. Second, 10,000 FPGA clock cycles ($\approx$ 100 us) was the maximum allotted execution time, with each pixel comparison occurring within one clock cycle. In the rare event that this window was exceeded, the algorithm returned the already constructed contour points. Third, a contour containing fewer than 63 points was rejected and the FPGA returned the last valid contour detected for a given larva ID. The algorithm stopped when none of the remaining neighbours were edge pixels (Fig. 2.2).

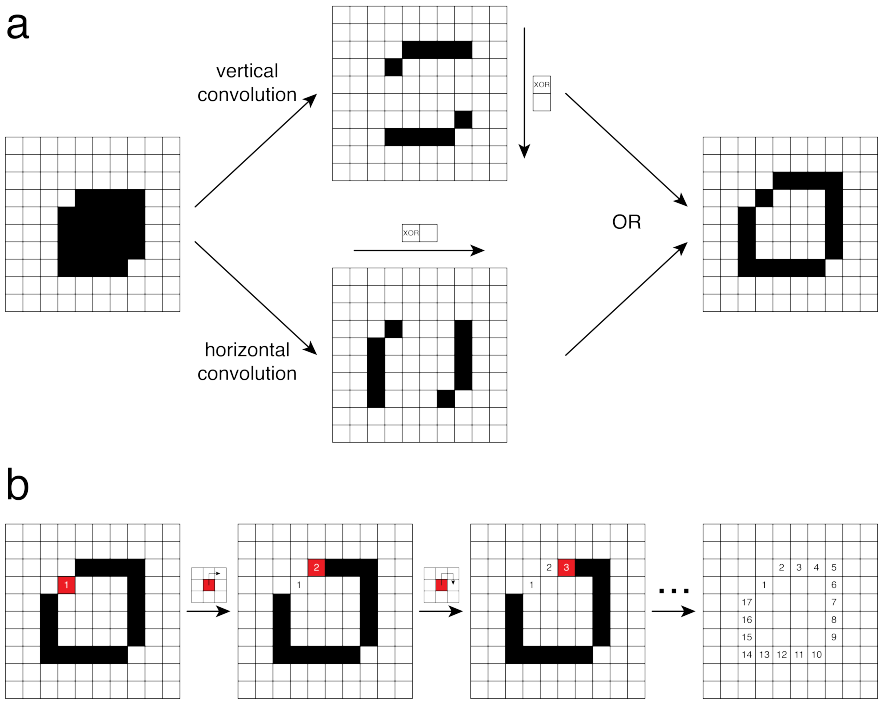

Fig. 2.2 Contour calculation on field-programmable gate array (FPGA). A simplified example is shown using a 10 x 10 pixel box containing a small object. **a.** The object (black) was detected against the background (white) using binary thresholding. Edge pixels were detected by combining the results of vertical and horizontal image convolution with a 2 x 1 XOR kernel using an OR operator. **b.** The contour points were reconstructed in an iterative process, starting with the edge pixel closest to the centre of the box. The next contour point was defined as the first neighbouring pixel that was found to be an edge pixel. Neighbouring pixels were assessed clockwise from the pixel directly above the contour point. The process ended when no eligible edge pixels could be found.

Contour processing and landmark detection

An undesired result of the FPGA contouring algorithm was the variable number of contour points across larvae and frames. Detected behaviour was therefore based on a smooth contour with a fixed number of 100 contour points. This contour regularization was achieved inside the Behaviour Programme using Fourier decomposition and reconstruction as in Masson et al. (2020).

Dr L Narayan implemented the initial detection of head and tail on the FPGA. The fundamental algorithm was developed and improved for the single-larva tracker by Dr J-B Masson, K Klein, and myself (see Section 2.1.1). K Klein applied the algorithm to this high-throughput system, iteratively adapting and testing the software to establish robust detection across multiple larvae. This algorithm defined the larva's head and tail as the contour points with the sharpest and second-sharpest curvature, respectively (Fig. 2.3). While correct in most cases, this calculation sometimes led to flipped detection of the two body ends. The Behaviour Programme flagged and corrected these false detection events at run time by calculating the distance head and tail traveled between frames and tracking the number of correct versus flipped detection events. The vote system correction commonly failed when the larva made large angle casts. The resulting contour was nearly-circular and exhibited similar curvature across all points. The solution required resetting the vote tallies when detecting these `ball` events (Fig. 2.3).

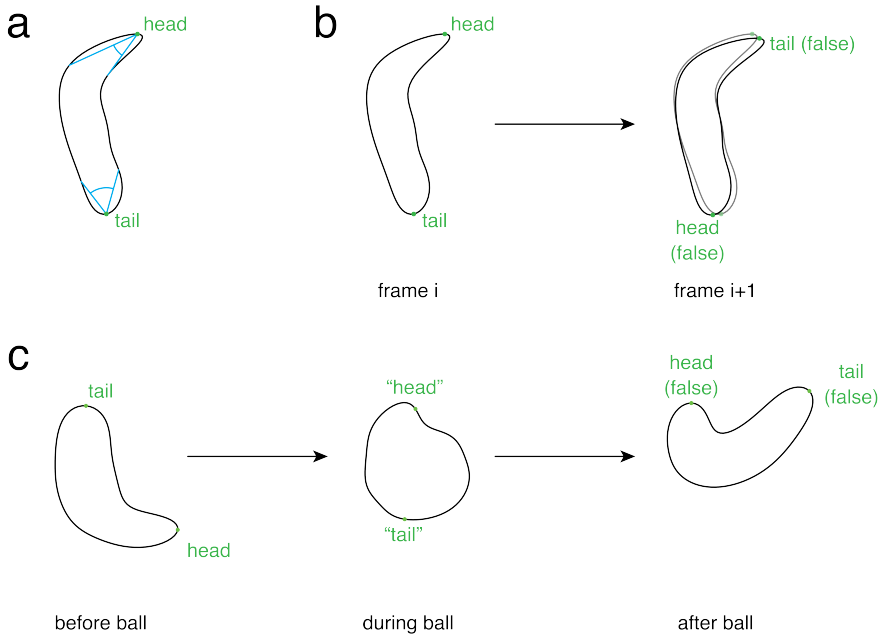

Fig. 2.3 Detecting head and tail. The larval contour (black outline) and head and tail (green) are shown. **a.** Initial detection of head and tail. The head was the contour point with the sharpest curvature. The tail was the contour point with the next-sharpest curvature which did not lie in close proximity to the head. **b.** The initial detection of head and tail was incorrect in some cases. False detection could be corrected by swapping head and tail, thereby minimising the distances from head and tail in the current frame (solid contour) to head and tail in the previous frame (transparent contour). **c.** The correction described in **b** failed if larvae curled up such that the contour appeared circular ("ball"). To eliminate this source of false head and tail detection, these events were detected using a `ball` classifier.

The larval spine was defined as 11 points running along the central body axis from `head` to `tail` (Fig. 2.4; Swierczek et al., 2011). In addition to `head` and `tail`, the Behaviour Programme calculated three equally distributed landmark points along the

spine (`neck_top`, `neck`, and `neck_down`). A fourth landmark, the `centroid`, defined the larva's location. The six landmarks were collectively used to extract features for training behaviour classifiers (Fig. 2.4).

The Behaviour Programme transformed the raw contour and spine from camera coordinates (in pixels) to world coordinates (in mm). If stable larval detection criteria were met, all spine points were temporally smoothed using exponential smoothing (Fig. 2.4).

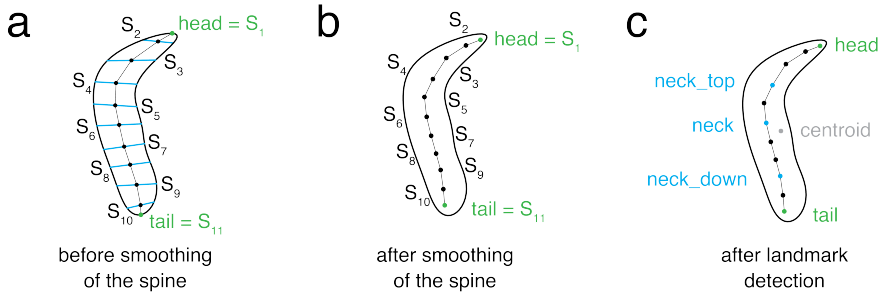

Fig. 2.4 Calculating a smooth spine and landmark points. The larval contour is shown (black outline). The spine S was comprised of eleven points (black), including head and tail (green). **a.** The raw spine points were obtained by finding the centres between equally spaced contour points on either half of the contour as defined by head and tail. The first spine point was the head, the last spine point was the tail. **b.** The smooth spine was obtained by exponentially smoothing the raw spine. **c.** Four additional landmark points, `neck_top`, `neck`, and `neck_down` (blue), and the contour `centroid` (grey), were calculated.

Feature extraction

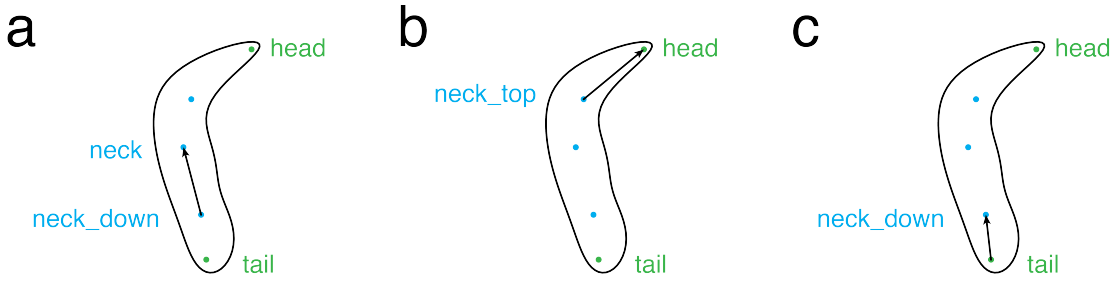

Fig. 2.5 Calculating direction vectors. Three direction vectors were calculated based on `head`, `tail`, and the landmark points. **a.** `direction_vector` was the normalised vector from `neck_down` to `neck`. **b.** `direction_head_vector` was the normalised vector from `neck_top` to `head`. **c.** `direction_tail_vector` was the normalised vector from `tail` to `neck_down`.

A machine learning approach was developed to address the high deformability of the larva shape, ensure live execution, reduce overfitting, and limit the volume of data tagging. What follows is a brief summary of larval features describing motion direction, body shape, and velocity that were calculated from the contour and spine data inside the Behaviour Programme. Features were designed as in Masson et al. (2020), with notable

modifications implemented on the high-throughout system by K Klein. These modifications were based on work completed on the single-larva tracker by Dr J-B Masson, K Klein, and myself, and were necessary to run the inference live:

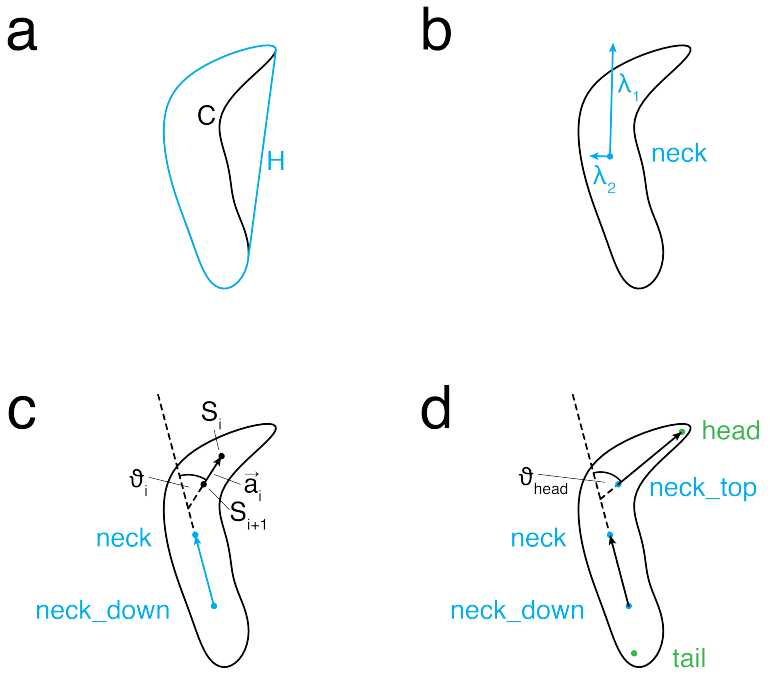

Fig. 2.6 Features describing body shape. a. Outline of a larva with contour C (black) and its convex hull H (blue). **b.** Shown here are the eigenvectors (blue) of the larval contour (black) structure tensor with respect to `neck` and their corresponding eigenvalues λ_1 and λ_2. **c.** ϑ_i was defined as the angle between `direction_vector` (blue) and the vector $\vec{a_i}$ that passed through spine points S_i and S_{i+1} (black). **d.** $\vartheta_{\texttt{head}}$ was defined as the angle between `direction_vector` and `direction_head_vector`. `head` and `tail` are shown in green.

1. Motion Direction (Fig. 2.5)
    - `direction_vector`: normalised vector describing the main body axis
    - `direction_head_vector`: normalised vector describing the head axis
    - `direction_tail_vector`: normalised vector describing the tail axis
2. Body Shape (Fig. 2.6)
    - `skeleton_length`: summed distances between consecutive spine points
    - `perimeter`: summed distances between neighbouring contour points
    - `larva_arc_ratio`: ratio of contour perimeter to convex hull perimeter (`larva_arc_ratio` ≥ 1 and was close to 1 when larva was in either straight or ball-like shape)

- `larva_area_ratio`: ratio of the areas enclosed by the contour and its convex hull ($0 \leq$ `larva_area_ratio` ≤ 1 and was close to 1 when the larva was in either straight, heavily curved, or ball-like shape)
- `eig_reduced`: `eig_reduced` $= \frac{|\lambda_1 - \lambda_2|}{\lambda_1 + \lambda_2}$ where λ_1, λ_2 were the eigenvalues of the structure tensor of the larval contour with respect to the `neck` ($0 \leq$ `eig_reduced` ≤ 1 and `eig_reduced` decreased as the cast amplitude of the larva increased)
- `s`: normalised angle along the body ($-0.5 \leq$ `s` ≤ 1, was close to 1 when larva was straight, and decreased with increasing cast amplitude)
- `asymmetry`: sine of the angle between `direction_vector` and `direction_head_vector` (`asymmetry` > 0 when larva bent left and `asymmetry` < 0 when larva bent right)
- `angle_upper_lower`: absolute angle between `direction_vector` and `direction_head_vector` (despite similarity to `asymmetry`, this develops different dynamics following temporal smoothing, which are valuable for stable left and right cast detection)

3. Velocity (Fig. 2.7)

   - Velocity of all six landmark points (`head_speed`, `neck_top_speed`, `neck_speed`, `neck_down_speed`, `tail_speed`, and `v_centroid`) in mm/s over interval $dt = 0.2\,s$ (four frames)
   - `v_norm`: arithmetic mean of `neck_top_speed`, `neck_speed`, and `neck_down_speed`, passed through a hyperbolic tangent activation function to suppress excessively large values
   - `speed_reduced`: relative contribution of `neck_top_speed` to `v_norm`, passed through a hyperbolic tangent activation function to suppress excessively large values (`speed_reduced` increased when the anterior larval body moved quickly compared to the posterior, e. g. when a cast was initiated)
   - `damped_distance`: distance (mm) travelled by `neck`, giving greater weight to recent over past events
   - `crab_speed`: lateral velocity (mm/s), defined as the component of `neck_speed` orthogonal to `direction_vector_filtered`

- `parallel_speed`: forward velocity (mm/s), defined as the component of `neck_speed_filtered` parallel to `direction_vector_filtered`
- `parallel_speed_tail_raw`: tail's forward velocity (mm/s), defined as the component of `tail_speed_filtered` parallel to `direction_tail_vector_filtered`
- `parallel_speed_tail`: similar to `parallel_speed_tail_raw`, with the difference that `tail_speed_filtered` was normalised prior to calculating the dot product (i. e. a measure of tail movement direction which took values between -1 (backward) and +1 (forward))

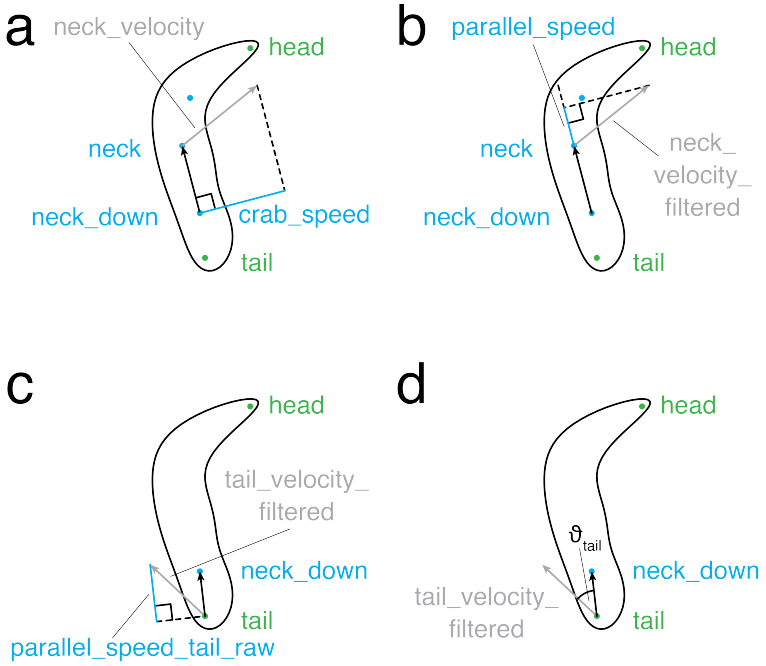

Fig. 2.7 Velocity features. The larval contour is shown in black while `head` and `tail` are shown in green. **a.** `crab_speed` (blue) was defined as the component of `neck_speed` (grey) that was orthogonal to `direction_vector_filtered` (black). **b.** `parallel_speed` (blue) was defined as the component of `neck_speed_filtered` (grey) that was parallel to `direction_vector_filtered` (black). **c.** `parallel_speed_tail_raw` (blue) was defined as the component of `tail_speed_filtered` (grey) that was parallel to `direction_tail_vector_filtered` (black). **d.** $\vartheta_{\texttt{tail}}$ was defined as the angle between `tail_speed_filtered` (grey) and `direction_tail_vector_filtered` (black).

Exponential smoothing was implemented to extract features in real time and address various sources of noise. Smoothing is here defined for a given feature `f` (Fig. 2.8):

$$\texttt{f\_filtered}_t = (1-\alpha)\cdot \texttt{f\_filtered}_{t-\Delta t} + \alpha \cdot \texttt{f}_t$$

where t is unitless, but derived from the experiment time in seconds, $\alpha = \frac{\Delta t}{\tau}$ with $\Delta t = 0.05\,s$ and $\tau = 0.25\,s$. Features that had the potential to exhibit large value

deviations (e. g. `v_norm`) were instead bounded using a hyperbolic tangent function. Additionally, some features were exponentially smoothed over a longer time window (where $\alpha_{long} = \frac{\Delta t}{\tau_{long}}$ with $\Delta t = 0.05\,s$ and $\tau_{long} = 5\,s$) (Fig. 2.8).

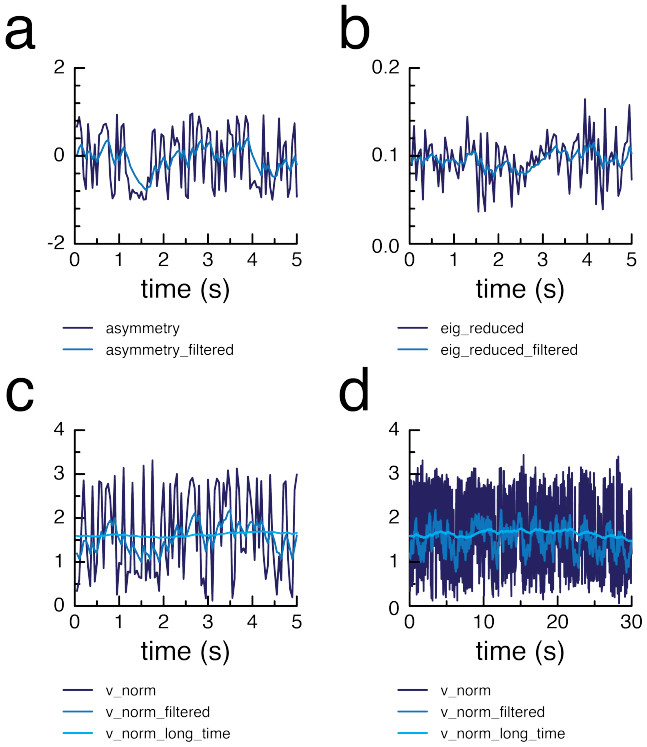

Fig. 2.8 Temporal smoothing of features. a-b. Example graphs of raw (dark blue) and filtered (mid blue) `asymmetry` (**a**) and `eig_reduced` (**b**) values over time. **c–d.** Example graphs of raw (dark blue), filtered (mid blue), and long-time filtered (light blue) `v_norm` values over a short (**c**) and a long (**d**) period of time.

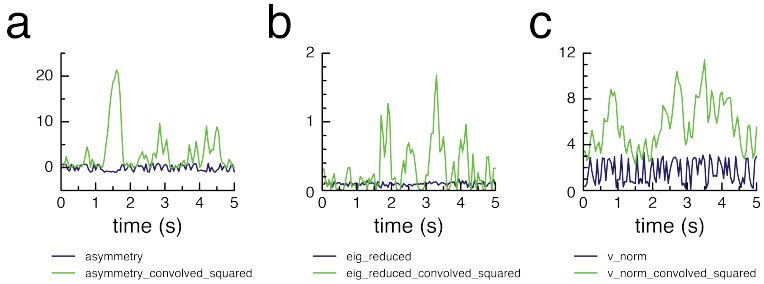

Fig. 2.9 Differentiation by convolution. Example graphs of raw (dark blue) and convolved squared (green) `asymmetry` (**a**), `eig_reduced` (**b**) and `v_norm` (**c**) values over time.

Convolution was used to approximate a smoothed squared derivative for each feature (Fig. 2.9); useful for integrating information over time without needing to further expand the feature space. The underlying mathematical concepts were motivated by Masson et al. (2012). For a given feature `f` at time t, `f_convolved_squared` was

calculated as follows:

$$\begin{aligned}
\texttt{f1}_t &= (1-\lambda\Delta t)\cdot \texttt{f1}_{t-\Delta t} + \frac{1}{2}\Delta t\cdot(\texttt{f}_{t-\Delta t}+\texttt{f}_t)\\
\texttt{f2}_t &= \lambda\Delta t\cdot \texttt{f1}_{t-\Delta t} + (1-\lambda\Delta t)\cdot \texttt{f2}_{t-n\Delta t}\\
\texttt{f_convolved_squared}_t &= k\cdot(\texttt{f1}_t - \texttt{f2}_t)^2,
\end{aligned}$$

where $\Delta t = 0.05\,s$, $\lambda = \frac{1}{\tau}$, $\tau = 0.25\,s$, and $n = 5\,s$. k values were empirically chosen for each feature.

Behaviour classifiers

Behaviour classifiers were developed using a user interface similar to JAABA (Kabra et al., 2013). The underlying algorithms were initially developed by Dr J-B Masson, K Klein, and myself for use on the single-larva tracker, combining trained neural networks and empirically determined linear thresholds. K Klein performed the necessary updates to utilise these classifiers on the high-throughput system. She also trained each classifier and manually validated its performance on multi-larva data (Table 2.1). K Klein also developed a MATLAB user interface with functions for data visualisation, manual annotation, and machine learning using the Neural Network Toolbox, the Deep Learning Toolbox, and the Statistics and Machine Learning Toolbox. Below is a brief description of the behaviour classifiers and associated performance results.

The `cast` classifier was based on predefined thresholds for temporally smoothed body shape features and was itself exponentially smoothed over time. Independent `left` and `right` classifiers were used to initially detect cast direction. To detect left and right casts, these classifiers were combined with the smoothed `cast` classifier using an AND conjunction. The raw time series of left and right casts was further smoothed post-acquisition using a custom MATLAB script: two casts to the same side separated by less than 200 ms were combined into a single long cast, and short casts of less than 200 ms were removed from analysis.

To improve `left` and `right` detection performance, a classifier was developed for circular larval contours. This `ball` classifier used a feed-forward neural network with a single fully connected hidden layer whose inputs were normalised values of `eig_reduced`, `larva_arc_ratio`, and `larva_area_ratio`. The hidden layer consisted of five neurons with a hyperbolic tangent activation function. The output layer contained a single neuron and used a sigmoid activation function. The neural network was trained

Table 2.1 Manual quantification of behaviour detection performance

`back` (268 events from 24 larvae in 60 minutes of video data)	
Precision	86.5%
Recall	88.4%
`cast` (714 events from 24 larvae in 60 minutes of video data)	
Precision	95.6%
Recall	96.4%
Accuracy of `left` and `right` detection (true-positive casts)	97.3%
`forward` (425 events from 24 larvae in 60 minutes of video data)	
Precision	97.8%
Recall	94.1%
`forward_peristaltic` (2954 events from 24 larvae in 60 minutes of video data)	
Precision	99.5%
Recall	93.6%
Events which are falsely combined with another event	10.7%
Events which are detected as more than one event	1.2%
`roll` (240 events from 24 larvae in 60 minutes of video data)	
Precision (rolls and roll-like events)	96.6%
Recall (rolls)	86.7%
Recall (roll-like events)	25.8%

in MATLAB on a manually annotated data set for 500 epochs using a cross-entropy loss function and scaled conjugate gradient backpropagation. If a `ball` was detected within the previous 1.5 s, `left` and `right` classifiers were overwritten to match the last detected cast direction prior to the beginning of the `ball`.

The `back` classifier detected individual backward peristaltic waves based on thresholds for smoothed tail velocity features combined with no `ball` detection within the previous 1.5 s.

Two different classifiers were used to detect crawling. `forward` detected longer forward crawl periods based on thresholds for smoothed tail velocity features combined with no `ball` detection within the previous 1.5 s. `forward_peristaltic` detected individual forward peristaltic waves based on the `forward` classifier and a threshold on forward tail velocity.

The `roll` classifier was based on thresholds for body shape and velocity combined with no `ball` detection and was exponentially smoothed over time. If a `roll` was detected within the previous 1.5 s, `forward`, `forward_peristaltic`, and `back` classifier values were reset to reduce false-positive detection for these classifiers. Unusual behaviour patterns such as rapid casting or twitching could be observed in addition to true larval rolling. These behaviours were considered "roll-like" events during manual validation of the `roll` classifier's performance.

Optogenetic stimulation

Optogenetic stimulation was achieved using two digital micromirror devices (DMDs) to project light patterns onto larvae on the agarose plate. During the hardware design process, Dr L Narayan and K Klein tested two different DMD models. One contained an integrated 613 nm LED (#CEL-5500-LED, Digital Light Innovations, Austin, Texas) and the other (#CEL-5500-FIBER, Digital Light Innovations, Austin, Texas) received input from an external 625 nm LED (#BLS-GCS-0625-38-A0710, Mightex Systems, Ontario, Canada) controlled by a BioLED light source control module (#BLS-13000-1, Mightex Systems, Ontario, Canada) and fed through an optic fibre (#LLG-05-59-420-2000-1, Mightex Systems, Ontario, Canada). Both DMDs operated like a 768 x 1024 pixel monochrome red light projector with numerous rotatable micromirrors used to modulate the intensity of individual pixels. Each DMD pixel corresponded to a 291.7 µm diameter section of the plate. Although both achieved similar light intensities, each DMD on its own was insufficient for optogenetic stimulation of larvae. Dr L Narayan therefore installed both devices on the system such that their projections each covered the entire agarose plate. In this way, the summed light intensities of the two DMDs could be achieved at all locations. Accurately aiming light at crawling larvae required spatial calibration of each DMD. Dr L Narayan and K Klein developed a spatial calibration procedure in which square spots were projected at fixed DMD pixel locations and the corresponding camera coordinates were linearly fit. I performed validation and testing of the spatial calibration output on live animals. Dr L Narayan updated the camera coordinate mapping to address initial errors I identified in both the overlap of both DMD projections and their centring over individual larvae.

K Klein determined that DMD illumination using the default light output was not uniform at plate level, which could have resulted in variable optogenetic stimulation depending on larval location. The maximum achievable light intensity at the plate's edge was approximately 40% of the peak value at its centre. The pixel intensity of the DMD

image was therefore normalised to the highest intensity uniformly achievable at all plate locations. A look-up table containing the normalisation factor for each DMD pixel was then calculated using bi-linear interpolation with approximately 100 light intensity values measured across the plate. To accommodate for possible differences in non-uniformity between the two DMDs, this intensity calibration (developed by Dr L Narayan and K Klein) was performed for both DMDs simultaneously following spatial calibration. When fully calibrated, the system could achieve a uniform light intensity of 285 µW/cm^2. I performed all spatial and intensity calibrations for both DMDs preceding the collection of behavioural data for larval operant learning experiments.

A user-defined Behaviour Programme protocol operated on the behaviour detection output and sent 8-bit optogenetic stimulation instructions to the LabVIEW application. In its final configuration, the LabVIEW application updated DMD projections at 20 Hz, therefore the delay between behaviour detection and closed-loop optogenetic stimulation of individual larvae did not exceed 50 ms. Furthermore, K Klein and I decided that if two or more larvae were close enough such that their corresponding stimulation areas overlapped, the light intensity in the overlapping region should be set to the smallest of those values to avoid undesired stimulation.

I contributed to validating the high-throughput system's capability for delivering spatially accurate optogenetic stimulation with minimal temporal delay. I gathered open-loop and closed-loop data on live larvae. I helped K Klein manually analyse these data for preliminary readouts on larval responses to red light. K Klein also used the data I collected to iteratively test calculations within the Behaviour Programme. Alongside K Klein, I also debugged system crashes, data processing delays, data file writing errors, and broken links between the GUI and the Behaviour Programme. While not detailed here, it is important to emphasise that K Klein and I continued to improve the end user's experience of the high-throughput system. With such enhanced system capabilities comes complexity in maintenance and daily operation. Our awareness of potential sources of human error guided calibration procedures, user interface design, and data output structure.

Thermogenetic stimulation

Thermogenetic stimulation was achieved by heating up larvae with a custom infrared (IR) laser set-up built by Dr Chris McRaven and Dr Michael Winding. A 1490 nm laser diode beam (#2CM-101, SemiNex, Peabody, Massachusetts) was fed into a two-axis galvanometer system (#GVSM002, Thorlabs, Newton, New Jersey), both controlled by

an analogue output device (#PCIe-6738, National Instruments, Austin, Texas). Two mirrors inside the galvanometer were rotated around orthogonal axes to target the beam spot to any user-defined location on the agarose plate. The beam spot measured approximately 5 mm in diameter, depending on the beam's angle of incidence to the plate. Mirror positions were controlled by two integrated motors that received voltage inputs. Each voltage pair clearly defined the laser beam's position.

Spatially calibrating the galvanometer was necessary to obtain a map between larval locations in world coordinates and the mirror motor input voltages. The procedure for spatial calibration of the galvanometers was developed by Dr L Narayan, K Klein, and myself. A visible aiming beam was scanned across the agarose plate using a fixed set of voltage pair inputs to the galvanometer. With the optical filter removed from the camera, the aiming beam's location in camera coordinates was automatically extracted from the image using binary thresholding. Two voltage-to-camera look-up tables were generated through bi-linear interpolation of these measured coordinates. For accurately targeted thermogenetic stimulation, the location of the larval `centroid` was first converted to camera coordinates using the existing world-to-camera transform and was then mapped to a pair of galvanometer input voltages using the look-up tables. I tested the spatial accuracy of galvanometer positioning on live animals.

Laser intensity calibration was also necessary to ensure that all larvae received the same stimulation regardless of their position on the agarose plate. The procedure for laser intensity calibration was developed by Dr L Narayan, K Klein, and myself. A larva's location changed the laser beam's angle of incidence, causing the illuminated spot at plate level to take an elliptical shape with variable size. Although laser beam power was constant, the changing spot area generated inconsistencies in the amount of IR light covering each larva. Calibration was used to normalise the desired laser intensity to achieve constant power per unit area. A visible aiming beam was scanned across the plate and the camera image automatically measured the beam's spot size at various locations. Bi-linear interpolation was then used to generate a pixel-wise look-up table containing the necessary scaling factors for the laser power. At the location where the laser spot area was smallest, the maximum power was reduced to 67.3%. To account for a nonlinear relationship between the laser source input voltage and the laser's total power output, a voltage-to-power map was generated from manual measurements I acquired. With these transformations, the system could calculate the laser source input voltage necessary to produce uniform, 5.26 W stimulation at any location. I performed

spatial and intensity calibration for the laser preceding the collection of behavioural data for proof-of-principle experiments.

A user-defined Behaviour Programme protocol operated on the 20 Hz behaviour detection output and sent thermogenetic stimulation instructions to the LabVIEW application which controlled the galvanometer and laser. Four `centroid` locations were specified on every frame, enabling a single galvanometer to cycle the laser beam between four individual larvae at 20 Hz. Although 80 Hz position updates are well within the galvanometer's mechanical capabilities, K Klein and I worked together to determine a consistent intra-cycle stimulus exposure time for each larva. The primary constraint to consider was the maximum likely travel time between detected objects. Using a thermal sensor, K Klein and I observed temporal changes in power as the laser cycled between four locations. We considered different path orders, including those over the maximum possible travel distance when the locations resided at the four plate corners. Within the available 50 ms time window, we therefore set the Behaviour Programme to heat each larva for 11 ms. Switching off the laser input for 1.5 ms between larvae accounted for small time fluctuations surrounding each new galvanometer position update and helped avoid undesired stimulation of other plate areas (Fig. 6.2d). If fewer than four objects were detected in a given frame, the remaining galvanometer target locations were set to the plate's centre and the corresponding laser intensity was set to zero. This temporal pattern of galvanometer position updates yielded no more than 100 ms delay between behaviour detection and closed-loop thermogenetic stimulation.

Three parameters influenced larval temperature increase following thermogenetic stimulation with the IR beam: i) the laser power, ii) the total duration of the stimulus, and iii) the order in which the galvanometer cycles between locations in its 80 Hz movement. Preliminary experiments performed by myself and K Klein suggested that these parameters could be adjusted to simultaneously stimulate eight or twelve larvae using a single galvanometer. This could potentially eliminate the need to install three additional laser sources to target all 16 larvae.

2.2.2 Software availability

All software code is available upon request.

2.2.3 Fly strains and larval rearing

We used the following fly strains: *58E02-Gal4* (Bloomington stock 41347), *69F06-Gal4* (Bloomington stock 39497), *72F11-Gal4* (Bloomington stock 39786), *attP2* (Pfeiffer et al., 2008), *Ddc-Gal4* (Li et al., 2000), *SS01989* (own stock), *TH-Gal4* (Friggi-Grelin et al., 2003), *Tph-Gal4* (Park et al., 2006), *Trh-Gal4* (Alekseyenko et al., 2010), *UAS-CsChrimson* (Bloomington stock 55134), *UAS-CsChrimson; tsh-LexA, LexAop-Gal80* (Dr Stefan Pulver, Dr Yoshinori Aso), *UAS-dTrpA1* (Dr Paul Garrity), *UAS-GFP* (Nern et al., 2015), and w^{1118} (Hazelrigg et al., 1984).

Fly stocks were maintained in vials filled with standard cornmeal food (Wirtz and Semey, 1982; 49.2 ml of molasses, 19.9 g of yeast, 82.2 g of cornmeal, 7.4 g of agarose, 9.8 ml of 20% Tegosept solution in 95% ethanol and 5.2 ml of propionic acid in 1 litre of water). For proof-of-principle and operant and classical learning experiments, eggs were collected overnight for approximately 12–18 hours on standard cornmeal food plates with additional dry yeast to increase laying. These experiments were performed using foraging-stage third-instar larvae (72–96 hours after egg laying) reared at 25 ℃ and 65% humidity (Eschbach et al., 2020b; Jovanic et al., 2016, 2019; Ohyama et al., 2013, 2015). Specifically for optogenetics experiments, larvae were raised in the dark and a 1:200 retinal solution (diluting 1 g of powdered all-*trans*-retinal (#R240000, Toronto Research Chemicals, Ontario, Canada) in 35.2 ml of 95% ethanol) was added to the food unless indicated otherwise. For immunohistochemistry, eggs were collected during daytime for approximately four hours on standard cornmeal food plates with added yeast. Dissections were performed using wandering-stage third-instar larvae (118–122 hours after egg laying).

2.2.4 Immunohistochemistry and confocal imaging

Janelia Research Campus' FlyLight team performed all dissections, immunohistochemical stainings, and confocal imaging following a procedure adapted from Jenett et al. (2012) and Li et al. (2014). Larval central nervous systems (CNSs) were dissected in cold 1x phosphate buffer saline (PBS, Corning Cellgro, #21-040) and transferred to tubes filled with cold 4% paraformaldehyde (Electron Microscopy Sciences, #15713-S) in 1x PBS. Tubes were incubated for one hour at room temperature. The tissue was then washed four times in 1x PBS with 1% Triton X-100 (#X100, Sigma Aldrich St. Louis, Missouri) (PBT) and incubated in 1:20 donkey serum (#017-000-121,

Jackson Immuno Research, West Grove, Pennsylvania) in PBT for two hours at room temperature.

The tissue was then incubated in the primary antibody solution, first for four hours at room temperature and then for two nights at 4℃. This solution contained mouse anti-Neuroglian (1:50, #BP104 anti-Neuroglian, Developmental Studies Hybridoma Bank, Iowa City, Iowa), rabbit anti-green fluorescent protein (GFP) (1:500, #A11122, Life Technologies, Waltham, Massachusetts) and rat anti-N-Cadherin (1:50, #DN-Ex #8, Developmental Studies Hybridoma Bank, Iowa City, Iowa) in PBT. This solution was then removed and the tissue washed four times in PBT. The tissue was then incubated in the secondary antibody solution, first for four hours at room temperature and then for two nights at 4℃. This solution contained Alexa Fluor 568 donkey anti-mouse (1:500, #A10037, Invitrogen, Waltham, Massachusetts), FITC donkey anti-rabbit (1:500, #711-095-152, Jackson Immuno Research West Grove, Pennsylvania) and Alexa Fluor 647 donkey anti-rat (1:500, #712-605-153, Jackson Immuno Research West Grove, Pennsylvania) in PBT. After removal of the secondary solution, the tissue was washed in PBT four times and mounted on a coverslip coated with poly-L-lysine (#P1524-25MG, Sigma Aldrich, St. Louis, Missouri).

The coverslip with the CNSs was dehydrated by moving it through a series of jars containing ethanol at increasing concentrations (30%, 50%, 75%, 95%, 100%, 100%, 100%) for ten minutes each. The tissue was then cleared by soaking the coverslip with xylene (#X5-500, Fisher Scientific, Waltham, Massachusetts) three times for five minutes each. Finally, the coverslips were mounted in dibutyl phthalate in xylene (DPX, #13512, Electron Microscopy Sciences, Hatfield, Pennsylvania) with the tissue facing down on a microscope slide with spacers. The DPX was allowed to dry for at least two nights prior to confocal imaging with an LSM 710 microscope (Zeiss).

Details on the confocal imaging settings are provided in the respective figure captions. Confocal images were analysed using Fiji (ImageJ). Neurons were counted by specifying regions of interest around the cell bodies using raw image stacks.

2.2.5 Verification of optogenetic and thermogenetic stimulation efficiency

K Klein assessed the multi-larva tracker's optogenetic and thermogenetic stimulation efficiency through open-loop experiments. The behavioural readout was rolling upon exposure to stimulation. All larval handling and experiments were performed in the dark

to avoid unintended optogenetic stimulation. The one-minute experiment protocol began with a 15 s initialisation period in which larvae acclimated to the agarose plate and the `roll` behaviour classifier stabilised. In three subsequent 15 s stimulation cycles, larvae received 5 s of open-loop stimulation followed by 10 s without stimulation (Fig. 6.2b, Fig. 6.2e). Optogenetics were performed with the maximum available red light intensity of 285 µW/cm^2. Thermogenetics were performed with 40% of the maximum available laser intensity.

K Klein analysed both optogenetic and thermogenetic experiment data using identical assessment and exclusion criteria. For each larva, the criterion for a single `roll` was detection of the behaviour for at least 300 ms during a given 15 s stimulation cycle. This threshold ensured true rolls were counted, as opposed to rapid larval casts characteristic of aversion to light.

2.2.6 Operant conditioning

Experiment procedures

K Klein and I performed high-throughput operant conditioning experiments using our multi-larva closed-loop tracker. All larval handling and experiments were performed in the dark to avoid unintended optogenetic stimulation. We used water to wash approximately 10–12 larvae out of their food. Using a brush, we immediately placed these larvae into the centre of the agarose plate in such a way that they were not touching each other. We placed the agarose plate inside the tracker on top of the backlight and then shut the tracker door. Larvae were given at least 30 s to accustom to their new environment before we started the experiment.

The experiment protocol (designed by K Klein) began and ended with a one-minute test period without optogenetic stimulation. Between these test periods were four, three-minute training sessions during which larvae received red light stimulation of 285 µW/cm^2 for the entire duration of the detected cast. Which side received stimulation was randomised across trials such that approximately 50% of larvae were trained to develop a right cast preference and 50% a left cast preference. No stimulus was triggered when the larva was casting right or when its body was straight. The test periods were each separated by three-minute periods without stimulation. After the first minute of this period, we used a brush to gently move all larvae back to the centre of the plate and larvae were given time to recover before the beginning of the next training session. This recentring addresses problems encountered when performing prolonged

experiments with freely behaving larvae on a small agarose plate. The longer larvae are left undisturbed, the more likely they are to touch the plate's edge, causing tracking disruption and temporary loss of valid objects. This shrinks sample size and reduces training efficiency by decreasing the proportion of animals which are receiving the stimulus.

K Klein designed control experiments so that valid objects received optogenetic stimulation uncorrelated with behaviour. These control experiments were split into 60 s time bins, during which each valid object was randomly assigned a stimulus train from this same time bin, pulled from a prior experiment where stimulation correlated with behaviour.

Data analysis

Data analysis was conducted using custom MATLAB software. To ensure high quality data, it was necessary to remove invalid objects from the data set prior to behavioural analysis. These included corrupted objects (e. g. scratches on the plate or residual food) that the software mistook for larvae. They also included larvae that lost their object identity and were consequently detected for only part of the experiment (e. g. after temporarily reaching the plate's edge or touching other larvae). After equally splitting each experiment into 60 s time bins, we retained objects for analysis that fulfilled strict criteria: i) the object must have been detected in every frame of the bin; ii) the object's initial detection must have occurred at least 20 s prior to the start of the bin; iii) at no point during the bin did the smoothed velocity of the larval centroid exceed 1.5 mm/s; and iv) the mean of the smoothed centroid velocity across the object's detection period in the bin was at least 0.5 mm/s. To quantify the accuracy of this method, K Klein manually reviewed 350 videos of objects flagged as valid for a given 60 s bin. In this group, K Klein observed no severely corrupted objects. In one case (0.3%), a larva briefly touched another larva. In another case (0.3%), head and tail of a larva were falsely detected the majority of the time, leading to flipped detection of left and right casts.

When analysing valid bin data for operant conditioning of cast direction preference, we counted, for each larva, the numbers of left and right casts initiated within the bin. This was defined as the cast rate towards the respective direction. The difference in cast rate within the bin was defined, for each larva, as the number of casts towards the side paired with the optogenetic stimulus minus the number of casts toward the unstimulated side. We pooled together all larval data within each bin because casts to the left and right were each paired with the optogenetic stimulus for approximately half of the larvae.

Mean and standard error were calculated for cast rate to the stimulated side, cast rate to the unstimulated side, and difference in cast rate. For the control condition in which larvae received random stimulation during 50% of casts regardless of direction, we calculated mean and standard error for cast rates to the left and right and the difference in cast rate between left and right. Within a bin, cast rates to either side were compared to each other using a paired, two-sided Wilcoxon signed-rank test. This statistical choice was driven by the known pairing of these spatial observations for individual larvae and the non-normality exhibited across distributions of differences between these paired cast rates. Similar reasoning guided the usage of a two-sided Wilcoxon signed-rank test to compare the difference in cast rate against 0. The behaviour characteristics of experimental animals were compared to each control group using a non-parametric analogue to the two-sample t-test, a two-sided Mann-Whitney *U* test. The Mann-Whitney *U* test was also used to test for statistically significant differences in cast rate within a group, from before training to after training. Loss of object identity contributes to differences in number of larvae over time. Because of the resulting inability to pair observations between bins, Mann-Whitney *U* was performed assuming that the observations were independent from one another.

For all statistical comparisons described above, I performed a non-hierarchical bootstrap analysis. The foundation of this analysis was sampling (with replacement) from the observed behavioural data to build a bootstrap resample comprising the same number of larvae as was in the observed data set. When sampling cast rate data within a bin, I maintained the number of casts an individual larva made to the stimulated and unstimulated sides to preserve the correlation between these spatial observations. Once I constructed a bootstrap resample, I calculated the mean of the bootstrapped cast direction preference measure (either #casts/min or difference in #casts/min) for this resample. I wrote Matlab code to repeat this process of resampling and calculating the mean 1000 times for both groups within a given statistical comparison. A one-sided threshold test was then conducted across each of the 1000x1000 possible combinations of the bootstrap resamples from both groups. The notable exception to this was comparing the difference in cast rate against 0, for which there was only one group. The bootstrap result was computed as the proportion of random resample combinations that satisfied the threshold test.

2.2.7 Classical conditioning

Experiment procedures

CsChrimson (Klapoetke et al., 2014) was expressed under the control of driver lines targeting candidate valence-conveying neurons. These driver lines were classified based on expression pattern and previous functional data and are known to drive expression in larvae. K Klein paired optogenetic activation of neurons (unconditioned stimulus (US)) with odour presentation (conditioned stimulus (CS)) to induce olfactory memory (Fig. 6.7a). For each driver line, data was acquired from at least two separate crosses.

K Klein followed a procedure for classical conditioning similar to those described in Gerber and Hendel (2006), Saumweber et al. (2011) and Eschbach et al. (2020b). Approximately 40 third-instar larvae were transferred onto a 4% agarose petri dish. Larvae were presented with an odour (1:10^4 ethyl acetate in ddH_2O) pipetted onto two pieces of filter paper attached to the lid of the dish. This enclosed dish was exposed to red light (630 nm, 350 µW/cm^2) for three minutes. Larvae were then transferred to a new agarose-filled petri dish with no odour on its lid ("air") and placed in the dark for three minutes. This training procedure was repeated three times, with alternating presentation of odour/light and air/dark (paired group). An unpaired group receiving reciprocal stimulus presentation (odour/dark, air/light) was trained simultaneously. This ensured that any observed effects were attributable to learning rather than innate odour preference or avoidance. The training trial order was reversed in half of the experiments, starting with air instead of odour presentation.

After training, larvae of both paired and unpaired groups were immediately transferred to a 1 cm middle zone in the centre of fresh agarose-filled petri dishes. A lid was placed on each dish, with odour presented on one side (odour side) but not the other (air side). After a three-minute test period in the dark, the number of larvae on the odour side, the air side, and in the middle zone were manually counted and entered into an Excel spreadsheet (Microsoft Corporation, Remond, Washington).

Data analysis

K Klein performed all data analysis for classical conditioning experiments. All data was manually entered into MATLAB and analysed using custom software. For each

experiment, a performance index (PI) was calculated as follows:

$$\text{Pref}_{\text{paired}} = \frac{\#(\text{larvae on odour side}) - \#(\text{larvae on air side})}{\#(\text{larvae on plate})} \quad \text{(paired dish)}$$

$$\text{Pref}_{\text{unpaired}} = \frac{\#(\text{larvae on odour side}) - \#(\text{larvae on air side})}{\#(\text{larvae on plate})} \quad \text{(unpaired dish)}$$

$$\text{PI} = \frac{\text{Pref}_{\text{paired}} - \text{Pref}_{\text{unpaired}}}{2} \quad \text{(combined)}$$

PIs take values between -1 and +1, where a positive PI reflects appetitive learning and a negative PI reflects aversive learning. Mean and standard error were calculated for each condition. Statistical differences between two groups were tested using a two-sided Mann-Whitney *U* test with Bonferroni correction. Significance compared to zero was tested with a two-sided Wilcoxon signed-rank test with Bonferroni correction.

Chapter 3

Larval decision-making in a sensory discrimination task

3.1 Introduction

A *Drosophila* larva's innate navigation behaviour includes reorientation manoeuvres during which a series of lateral head sweeps (casts) precedes turning in a new direction (Fig. 1.1). Although previous research has shown that head casts can inform the turn decision process through temporal comparison of sensory stimuli, the function of increasing the number of head casts within a reorientation manoeuvre has not been investigated. Larvae can form associative memories between punishment or reward and previously neutral environmental stimuli (Eschbach et al., 2011; von Essen et al., 2011), but whether they employ working memory to select turn direction during reorientation manoeuvres is unknown. I hypothesise that larvae use repeated casts to accumulate additional sensory information when stimuli are difficult to resolve, improving the accuracy of their turn direction decision. To test this hypothesis, I investigated whether larvae are more likely to turn toward the "better direction" following more head casts in a reorientation manoeuvre. In this chapter, I describe the larval decision-making task I designed, briefly summarise necessary improvements I made to experimental software design, outline how I assessed larval performance in the task, and propose a mechanism underlying the behaviour I observed.

3.2 Results

3.2.1 Developing sensory discrimination task structure

Animal behaviour within a two-alternative forced choice (TAFC) task has long been considered a readout for likely cognitive algorithms supporting perceptual decision-making. The simple structure of TAFC is experimentally advantageous in constraining both the time and information available to an animal without sacrificing ecological relevance (Bogacz et al., 2006). Inspired by TAFC, I designed a sensory discrimination task in which larval head cast direction (either left or right) during a reorientation manoeuvre triggered differential presentation of 617 nm red light stimuli. Within this closed-loop framework, I optogenetically activated nociceptive (pain-sensing) neurons in transgenic *ppk1.9-Gal4 x UAS-CsChrimson* larvae for the entirety of each detected head cast (Fig. 3.1a). Much research has concluded that the multidendritic class IV sensory neurons targeted by this genetic driver are the primary nociceptive neurons in the larva. They not only are necessary and sufficient for sensing noxious thermal and mechanical stimuli (Hwang et al., 2007; Tracey et al., 2003), but also mediate photoavoidance of noxious, high intensity light (Xiang et al., 2010). Given nociception's importance in warning animals of possible tissue-damaging stimuli, I anticipated that differential activation of these polymodal nociceptors in my task would encourage larvae to repeatedly head cast to determine which side is the less noxious, more favourable direction in which to crawl (Hwang et al., 2007). In this way, larvae in this task could be viewed as engaging in perceptual decision-making.

In identifying optimal intensity values for my sensory discrimination task, I aimed to observe active larval exploration rather than reflexive responses or stereotypical nocifensive escape behaviour. In this way, larval behaviour served as a proxy for neural activity and was valuable in establishing a lower perceptual bound. What resulted from my investigation were five possible fictive noxious stimuli conditions, of which individual larvae were randomly assigned one (for measured intensity values, see Section 2.1.3). In two of these conditions, $E_{L,M}$ and $E_{L,H}$, I established a difference in light intensity when the larva cast in one direction versus the other (as indicated by the subscripts). My goal in having experimental conditions with differing contrast between stimulus alternatives was to create behavioural tasks of varying difficulty. Results from olfactory discrimination tasks in adult (DasGupta et al., 2014) and larval (Gomez-Marin et al., 2011) *Drosophila* support the intuition that low stimulus contrast prompts increased environment sampling. By assigning the same intensity to the less noxious side, I aimed

to isolate whether any observed behavioural differences between the $E_{L,M}$ and $E_{L,H}$ conditions could be attributed to attempts at resolving stimuli differences as opposed to reflexive responses to overall higher intensities. The three remaining stimuli conditions were controls, each defined by equal light intensity irrespective of cast direction ($C_{L,L}$, $C_{M,M}$, and $C_{H,H}$). Importantly, their intensities matched the three unique values comprising the experimental conditions' stimulus alternatives (see Section 2.1.3). These control conditions provided baseline behavioral dynamics without sacrificing numerous casts per reorientation manoeuvre.

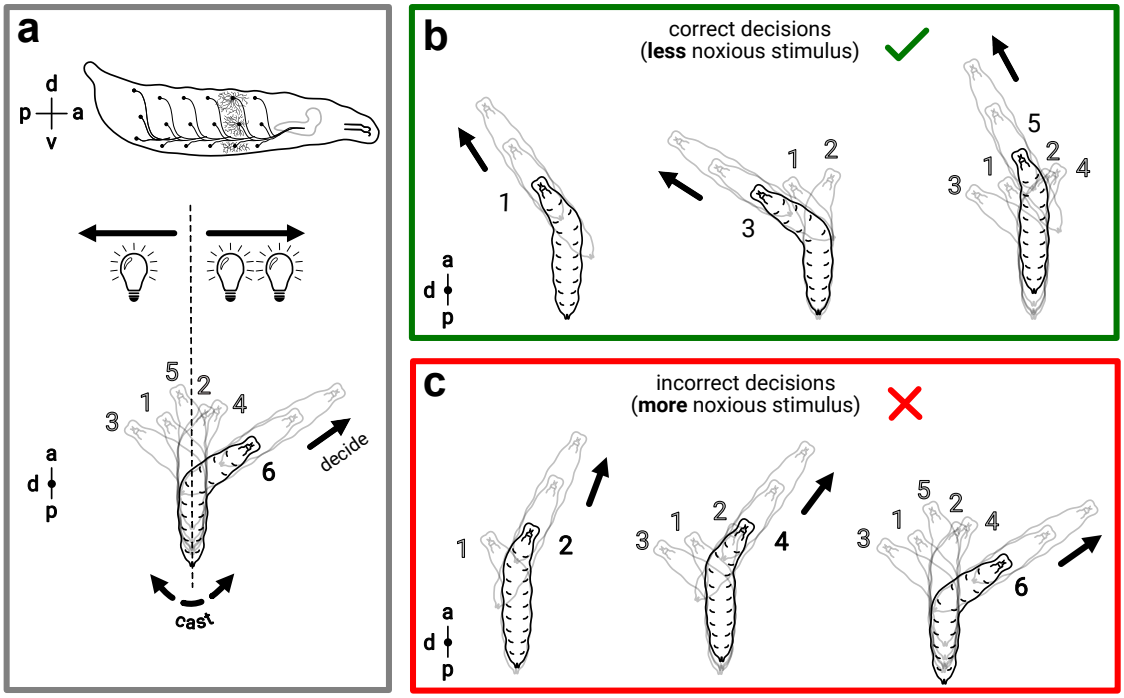

Fig. 3.1 Larval sensory discrimination task. Inspired by two-alternative forced choice (TAFC) task designs, my sensory discrimination task necessitated using a closed-loop tracker to optogenetically activate larval nociceptive neurons based on the direction of precisely detected lateral head sweeps (casts). **a.** (top) Side view schematic of a transgenic *Drosophila* larva with expression in multidendritic class IV neurons driven by *ppk-Gal4* (adapted from Yoshino et al. (2017)). Although only one hemisegment is shown, these neurons' naked dendrites completely cover the larval epidermis without overlap (Grueber et al., 2002; Xiang et al., 2010). (middle) In my sensory discrimination task, larval head casting triggered a 617 nm red light stimulus (light bulbs), the amplitude of which was predetermined and, in the case of the experimental stimuli conditions, direction-dependent (one light bulb versus two). (bottom) Schematic depicting a series of head casts within a reorientation manoeuvre, with the larva viewed from above. Although it is shown here pointed left, the larva's first cast could be toward either the left or right side and it always triggered the lower intensity stimulus in the condition. Larvae I selected for analysis alternated sides with each consecutive cast, as shown. I defined the last cast of the reorientation manoeuvre (here, cast six) as the larva's decision in the task. **b.** The correct decision on a given cast was to crawl in the direction of the less noxious stimulus. Because the larva's first cast decided the direction that always triggered the lower intensity stimulus, a correct decision was for the larva to accept any one of the odd-numbered casts in the sequence, rejecting all the evens. **c.** An incorrect decision was the reverse: accepting any one of the even-numbered casts to the high intensity side, rejecting all the odds. d: dorsal, v: ventral, a: anterior, p: posterior.

Irrespective of direction (left or right), the first head cast of the larva's reorientation manoeuvre triggered presentation of the stimuli condition's lower light intensity. By programming the tracker's stimulation protocol to remember the direction of this first cast,

I ensured that subsequent casts to the alternate side triggered presentation of the stimuli condition's higher light intensity (Fig. 3.1a). Larvae rarely finish reorientation manoeuvres by crawling straight, leaving two available decisions: crawl to the left or right (Fig. 3.1a). I defined the last cast of the larva's reorientation manoeuvre as the cast immediately preceding the completion of two forward peristaltic waves. I observed that larvae often initiated peristalsis from the bend in their body formed by this cast; a feature of the larval motor program that has previously been documented and analysed (Lahiri et al., 2011). This informed two additional criteria I established for visually evaluating peristalsis: peristalsis neither had to originate at the larva's tail nor had to occur along a straight body axis. I stopped the experiment manually once I observed the larva complete a reorientation manoeuvre. In a previous iteration of the experiment design, I afforded each larva three minutes to perform several manoeuvres but was uncertain whether larvae were transferring knowledge about the stimulus-cast direction relationship between manoeuvres. The results I describe in this thesis are from experiments in which I restricted the task to one reorientation manoeuvre per larva. This ensured all larvae were naïve to the task and eliminated potential operant learning confounds.

I performed all experiments using a single-larva closed-loop tracker to precisely couple real-time detection of larval behaviour with stimulus presentation (Fig. 2.1, Section 2.1.1, Schulze et al., 2015). For further details about my experiment procedures, see Section 2.1.3. In collaboration with Dr Jean-Baptiste Masson and Kristina Klein, I made marked improvements to the tracker's existing software framework. I helped stabilise both larval contouring and head and tail detection using preventive proximity measures and a corrective vote system. I also helped reduce high false positive and false negative rates for behaviour classifiers by redefining their feature descriptors, retraining them on newly acquired live data, and visually validating their performance myself (for further details, see Section 2.1.1). K Klein and I worked to establish baseline operations across all custom experiment protocols, but I wrote the software code specific to my sensory discrimination task. I also debugged and developed the user interface and data output structure beyond their default frameworks. My purpose in doing so was to provide the end user with greater options when customising experiment protocols and establish human-readable documentation of crucial experiment metadata.

3.2.2 Analysing larval performance

Following experimentation, I post-processed the single-larva closed-loop tracker output files in a custom MATLAB pipeline. This pipeline was written by Dr J-B Masson, with necessary updates to accommodate experiment metadata and additional online feature data by myself and K Klein. Following post-processing, I visually reviewed the resulting reconstructed behavioural data in a custom MATLAB GUI, itself designed by K Klein. K Klein kindly assisted me in reviewing a portion of the larval data I acquired. The review process required assessing against exclusion criteria (see Section 2.1.4), validating ON/OFF light signals against behaviour classification, and manually recording the following larval metrics:

- number of casts in the reorientation manoeuvre
- direction (left or right) of the reorientation manoeuvre's last cast (see criteria above)
- image frame numbers for both the beginning and end of the reorientation manoeuvre

I wrote a custom R pipeline to analyse this manual review data alongside each output file's metadata (e. g. stimulation protocol and light intensities defined via the GUI), and time-series data (e. g. frame-by-frame feature values, classified behaviours, and ON/OFF light signals). The first stage of my analysis pipeline filtered out larvae that cast more than once to the same side before either casting to the alternate side or finishing the reorientation manoeuvre. This enabled more direct comparison between larvae, as those remaining strictly switched direction with each consecutive head cast (schematised in Fig. 3.1a). Because the direction of the first cast always triggered the stimuli condition's lower light intensity, these larvae cast from the less noxious to the more noxious side, back and forth until they accepted either of the two (i. e. crawled in the direction set by this cast). Said another way, if a larva completed a reorientation manoeuvre following an odd number of casts, it decided to accept the less noxious stimulus (Fig. 3.1b). In contrast, if a larva completed a reorientation manoeuvre following an even number of casts, it decided to accept the more noxious stimulus (Fig. 3.1c).

To assess larval decision-making in my sensory discrimination task, I considered behavioural dynamics of the sample population within each stimuli condition. I describe here a mathematical formulation to clearly illustrate these ideas. A stimuli condition

contains a sample population of n larvae. A single larva j ($j \in [1, .., n]$) performed a single reorientation manoeuvre under this stimuli condition. The result of this manoeuvre was acceptance of one of k mutually exclusive cast outcomes. Said another way, k represents the number of the last cast in the manoeuvre. I let the discrete random variable X_j represent this number for the reorientation manoeuvre performed by larva j. For this larva, each outcome k occurs with probability $P(X_j = k) = r_k$ where $r_i \geq 0$ and $\sum_{i=1}^{k} r_i = 1$. Based on my own empirical knowledge, I can claim that X_j has no bearing on X_{j+1}. Furthermore, I assume that the probability distribution $r_1, ..., r_k$ over the cast accepted to end the reorientation manoeuvre does not differ from larva j to larva $j + 1$. The behaviours of the n larvae in the sample population are therefore independent and identically distributed. I let the discrete random variable Y_i represent the number of occurrences, in n larvae, of a reorientation manoeuvre comprising i casts ($i \in [1, ..., k]$). In this way, $Y \sim Multinomial(n; r_1, ..., r_k)$.

Probability of acceptance

Within a stimuli condition, I defined the total number of larvae that accepted cast k as:

$$n_k^{acc} = \sum_{j=1}^{n} I(X_j = k) \tag{3.1}$$

with I representing the indicator function:

$$I = \begin{cases} 1, & if X_j = k \\ 0, & \text{otherwise} \end{cases}$$

Consequently, $\sum_{i=1}^{k} n_k^{acc} = n$. I also defined the total number of larvae that performed cast k (whether or not they accepted it) as:

$$n_k^{tot} = \sum_{j=1}^{n} I(X_j \geq k) \tag{3.2}$$

with I representing the indicator function:

$$I = \begin{cases} 1, & if X_j \geq k \\ 0, & \text{otherwise} \end{cases}$$

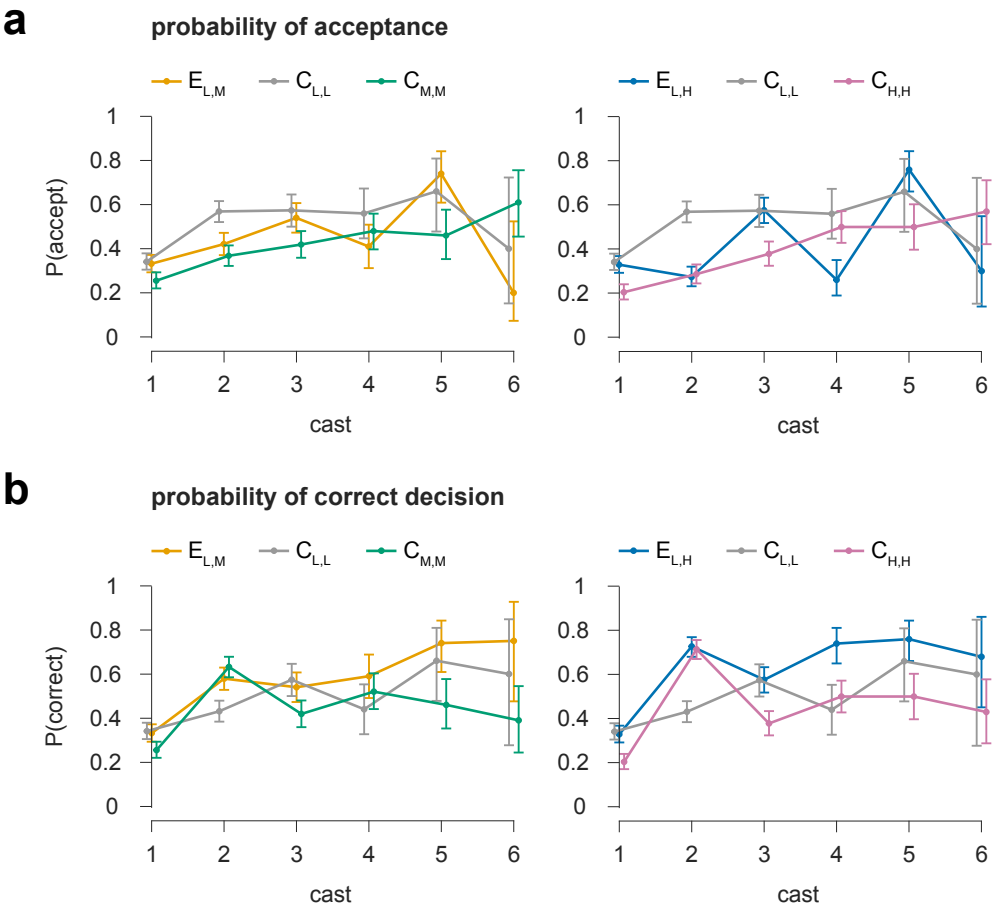

c

Condition	cast 1	2	3	4	5	6
$E_{L,M}$	579	387	224	103	61	16
$E_{L,H}$	605	406	295	125	92	22
$C_{L,L}$	662	436	188	80	35	12
$C_{M,M}$	573	427	270	157	82	44
$C_{H,H}$	554	441	315	196	98	49

d

Condition	correct	cast 1	2	3	4	5	6
$E_{L,M}$	0	387	163	103	42	16	4
	1	192	224	121	61	45	12
$E_{L,H}$	0	406	111	125	33	22	7
	1	199	295	170	92	70	15
$C_{L,L}$	0	436	248	80	45	12	5
	1	226	188	108	35	23	7
$C_{M,M}$	0	427	157	157	75	44	27
	1	146	270	113	82	38	17
$C_{H,H}$	0	441	126	196	98	49	28
	1	113	315	119	98	49	21

Fig. 3.2 Sensory discrimination task cast-by-cast decisions. Experiments were performed as described in Chapter 2 through optogenetic stimulation of transgenic *ppk1.9-Gal4 x UAS-CsChrimson* larvae. Data are shown for experimental conditions in which there is an intensity differential between the two stimulus alternatives ($E_{L,M}$ and $E_{L,H}$). Data are also shown for the associated control conditions in which intensities are equivalent regardless of cast direction ($C_{L,L}$, $C_{M,M}$, and $C_{H,H}$). Individual larvae belong to only one stimuli condition. By nature of the task design, the less noxious stimulus was presented on odd-numbered casts and the more noxious stimulus was presented on even-numbered casts. **a.** Both graphs show the probability of accepting a cast (ending the reorientation manoeuvre; $P(\text{accept})$) after performing the number of casts indicated by the x-axis. **b.** Both graphs show the probability of making the correct decision (acceptance of less noxious stimulus or rejection of more noxious stimulus; $P(\text{correct})$) after performing the number of casts indicated by the x-axis. Error bars in **b** and **c** represent the Clopper-Pearson Exact 95% confidence interval. Data is jittered for easier visualisation. **c.** This table shows, for each condition, the number of larvae that performed each cast (n_k^{tot}, $k \in [1, ...6]$). Note that the greater the proportion of larvae that end a reorientation manoeuvre following a given head cast, the fewer larvae remain to continue casting. **d.** This table shows, for each condition, the number of larvae that made the incorrect (0) or correct (1) decision after performing the indicated cast. Within each condition, the column totals match the values in panel **c**.

I was interested in the probability of larvae accepting a head cast given how many head casts they already performed. Combining Eqn.'s 3.1 and 3.2, this equates to:

$$\begin{aligned} P(\mathrm{accept})_\mathrm{k} &= n_k^{acc}/n_k^{tot} \\ &\sim Bernoulli \equiv P[X_j = k | X_j \geq k] \\ &= \frac{r_k}{\sum_{k' \geq k} r_{k'}} \end{aligned}$$

Fig. 3.2a shows this metric plotted across casts one through six for both experimental conditions and their associated controls. Although some larvae performed more than six head casts, I restricted the data to this window to ensure that, within a condition, at least ten larvae performed each cast (Fig. 3.2c).

Probability of correct decision

I developed a success probability metric, $P(\mathrm{correct})$, to assess whether performing more head casts in a reorientation manoeuvre improves the decision about the manoeuvre's final direction. In my sensory discrimination task, a correct decision is to accept a cast and finish the reorientation manoeuvre on the side associated with less noxious stimulus. By definition, this also requires rejecting casts to the side with the more noxious stimulus by casting to the other side (Fig. 3.1b). Mathematically,

$$P(\mathrm{correct})_\mathrm{k} = \begin{cases} P(\mathrm{accept})_\mathrm{k}, & \text{if } k \text{ is odd} \\ 1 - P(\mathrm{accept})_\mathrm{k}, & \text{if } k \text{ is even} \end{cases} \tag{3.3}$$

While $P(\mathrm{accept})$ measures when larvae ended their reorientation manoeuvre, $P(\mathrm{correct})$ tracks the action larvae made on each cast within the manoeuvre. Fig. 3.2b shows the $P(\mathrm{correct})$ metric plotted across casts one through six for both experimental conditions and their associated controls. Although some larvae performed more than six head casts, I restricted the data to this window to ensure that, within a condition, at least ten larvae performed each cast (Fig. 3.2d).

Evolution of acceptance and correct decision

The errors surrounding the $P(\mathrm{accept})$ and $P(\mathrm{correct})$ data in Fig. 3.2 represent Clopper-Pearson Exact 95% confidence intervals (for details on method selection, see

Section 2.1.4). Visible overlap of these confidence intervals within and between conditions raises uncertainty about whether a true, meaningful difference exists between observed probabilities. This did not, however, preclude me from qualitatively analysing the behaviour of these sample populations. It is also important to emphasise that $P(\mathrm{accept})$ and $P(\mathrm{correct})$ results from my sensory discrimination task exhibit an inherent dependence from one cast to the next, functioning like conditional probabilities based on previous behavioural choices. An appropriate statistical test for quantitative comparison within and between these conditions may therefore be one that considers trends based on the totality of each condition's data trajectory, rather than individual casts. In the paragraphs that follow, I outline my qualitative observations and the quantitative results I obtained in accordance with this thinking. I make specific reference to Fig. 3.3 which shows the same $P(\mathrm{accept})$ and $P(\mathrm{correct})$ data as Fig. 3.2 but omits the Clopper-Pearson confidence intervals so as not to confuse those error estimates with my logistic regression analysis.

Notably, the $P(\mathrm{accept})$ data show an inverse relationship between the probability of accepting the first cast and the intensity of the associated noxious stimulus. For conditions in which cast one triggered intensity L, acceptance probabilities lie around 0.335 ($E_{L,M} = 0.332; E_{L,H} = 0.329; C_{L,L} = 0.341$). The acceptance probability on cast one is smaller for $C_{M,M}$ (0.255) and smaller still for $C_{H,H}$ (0.204) (Fig. 3.3a). This graded behavioural response is an important indicator that larvae can perceive stimuli at these intensities. The low acceptance probabilities are also unsurprising, considering that cast one is each larva's first exposure to a noxious environmental stimulus.

Under the $P(\mathrm{accept})$ metric, I hypothesised that experimental larvae would show increased probability of acceptance of the less noxious side (k odd) with increased cast number. I expected the converse over casts that triggered the more noxious stimulus (k even): decreased probability of acceptance with increased reorientation manoeuvre length. The $E_{L,M}$ condition exhibits this behaviour but the probability of acceptance on cast two is greater than that of cast one despite the increase in noxious intensity. The $E_{L,H}$ condition is almost identical to $E_{L,M}$ in its magnitude of $P(\mathrm{accept})$ increase over odd-numbered casts. $E_{L,H}$, however, shows stable values of $P(\mathrm{accept})$ over even-numbered casts. What results in both experimental conditions is an oscillating trajectory with an envelope that widens with increasing cast number (Fig. 3.3a). In each control condition, I expected the probability of acceptance to lie close to 0.5 from cast two onward. The $P(\mathrm{accept})$ values in condition $C_{L,L}$ are nearly identical on casts two, three, and four (close to 0.57) before peaking at 0.66 on cast five and dropping to 0.40 on

cast six. Notably, the $C_{L,L}$ trajectory jumps much higher than $C_{M,M}$ and $C_{H,H}$ on cast two and stays consistently above these other control conditions through cast five. Conditions $C_{M,M}$ and $C_{H,H}$ show monotonic increases in $P(\mathrm{accept})$ with the exception of the value at cast five (Fig. 3.3a).

Under the $P(\mathrm{correct})$ metric, I hypothesised that experimental larvae faced with a differential between noxious stimulus alternatives would show increased probability of correct decision with increased cast number. I further hypothesised that task difficulty would alter the rate of this ascent. Specifically, the greater the separation between stimulus intensities, the quicker the larva may resolve this difference and end the reorientation manoeuvre with a cast toward the less noxious side. In agreement with my hypothesis, condition $E_{L,M}$ shows an upward trend in $P(\mathrm{correct})$ with increasing cast number. Condition $E_{L,H}$ shows a faster plateau to nearly the same values of $P(\mathrm{correct})$ as $E_{L,M}$, suggesting that larvae exposed to a larger stimulus differential need less sensory information to reach the same level of performance (Fig. 3.3b).

Because in the control conditions neither stimulus alternative was better than the other, I retained the $P(\mathrm{correct})$ definition for the control conditions as in Eqn. 3.3. In each control condition, I expected a flat $P(\mathrm{correct})$ trajectory from cast two onward, with values close to 0.5. Although both $C_{M,M}$ and $C_{H,H}$ show a spike in $P(\mathrm{correct})$ values above 0.5 at cast two, both hover at or below 0.5 from casts three through six. Condition $C_{L,L}$ does not show a large spike in $P(\mathrm{correct})$ at cast two and oscillates more widely around 0.5 from casts three through six (Fig. 3.3b). This difference in behaviour between $C_{L,L}$ and the remaining two controls may be a consequence of the smaller n_k^{tot} values at higher cast numbers (Fig. 3.3c), which itself may be due to limited perception of the L noxious stimulus intensity. I confirmed that the exact $P(\mathrm{correct})$ definition was irrelevant for the control conditions by calculating, for each condition, an average $P(\mathrm{correct})$ value for each cast over 1000 separate analyses. In each analysis, I assigned a random half of the larvae to have the correct side as the odd-numbered side and the other half to have the correct side as the even-numbered side. The result (not shown) was all three control conditions tightly oscillating around 0.5, beginning at cast one, with greater deviations away from 0.5 as cast number increased (reflecting the decrease in n_k^{tot}).

The qualitative differences in $P(\mathrm{correct})$ trajectories between the experimental conditions and their respective controls suggest that the more larvae sample the environment, the better their decision about the direction in which to crawl to end the reorientation manoeuvre. I aimed to determine whether this was supported statistically by a more significant rise in experimental $P(\mathrm{correct})$ trajectories than those of the

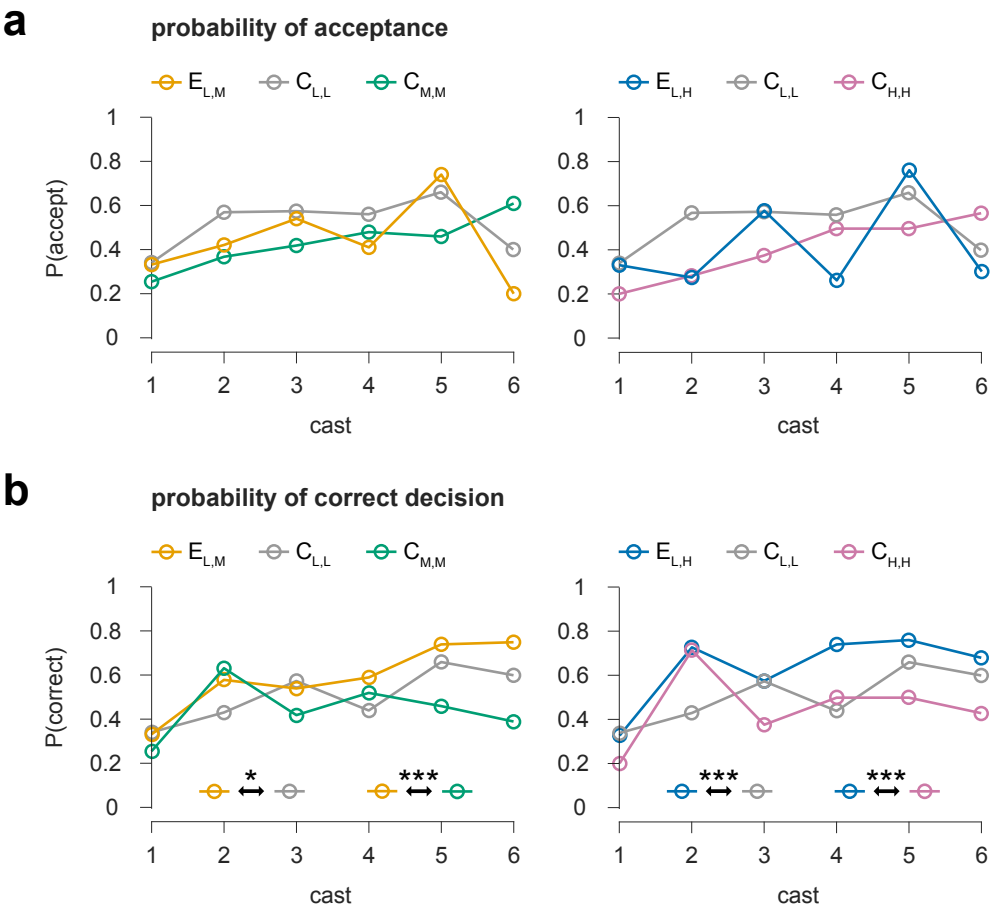

c

Condition	cast 1	2	3	4	5	6
$E_{L,M}$	579	387	224	103	61	16
$E_{L,H}$	605	406	295	125	92	22
$C_{L,L}$	662	436	188	80	35	12
$C_{M,M}$	573	427	270	157	82	44
$C_{H,H}$	554	441	315	196	98	49

d

Condition	correct	cast 1	2	3	4	5	6
$E_{L,M}$	0	387	163	103	42	16	4
	1	192	224	121	61	45	12
$E_{L,H}$	0	406	111	125	33	22	7
	1	199	295	170	92	70	15
$C_{L,L}$	0	436	248	80	45	12	5
	1	226	188	108	35	23	7
$C_{M,M}$	0	427	157	157	75	44	27
	1	146	270	113	82	38	17
$C_{H,H}$	0	441	126	196	98	49	28
	1	113	315	119	98	49	21

Fig. 3.3 Evolution of sensory discrimination task cast-by-cast decisions. Data in all panels is identical to that of Fig. 3.2, but is shown here without confidence intervals. Instead, asterisks below the trajectories in panels **a** and **b** show whether the slope of the logit regression is significantly different between the experimental condition and each of its respective controls (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$; for exact values of the parameter $\hat{\beta}_{cast}$, see Table 3.1). **a.** Both graphs show the probability of accepting a cast (ending the reorientation manoeuvre; $P(\text{accept})$) after performing the number of casts indicated by the x-axis. **b.** Both graphs show the probability of making the correct decision (acceptance of less noxious stimulus or rejection of more noxious stimulus; $P(\text{correct})$) after performing the number of casts indicated by the x-axis. **c.** This table shows, for each condition, the number of larvae that performed each cast (n_k^{tot}, $k \in [1, ...6]$). Note that the greater the proportion of larvae that end a reorientation manoeuvre following a given head cast, the fewer larvae remain to continue casting. **d.** This table shows, for each condition, the number of larvae that made the incorrect (0) or correct (1) decision after performing the indicated cast. Within each condition, the column totals match the values in panel **c**.

control conditions. In quantitatively comparing these behavioural trajectories, I considered correctness as a categorical response variable with two possible values, incorrect and correct ($Y \in 0, 1$; $P(Y = 1) = p$). I considered cast number as an independent variable with a set upper limit ($X_{cast} \in [1, 6]$). In accordance with this view, coordinate pairs (x_{cast}, y) represent every decision made within a condition's data set (i. e. one for each rejected or accepted cast; for counts by condition and cast number, see Fig. 3.3d). I sought advice from expert statistician Dr Carey Priebe, as I was unfamiliar with available tests for trends in binary response variables. I directed our conversations, combining my existing statistical knowledge with my experience in designing my sensory discrimination task and $P(\mathrm{correct})$ performance criterion to effectively relay my analysis considerations and interpret Dr C Priebe's suggestions.

Dr C Priebe's first suggestion was the Cochran-Armitage trend test, which tests for trends in proportions across an ordinal variable. In researching how Cochran-Armitage could be applied to comparing trajectories of proportions to one another, I learned about its similarities to hypothesis testing in binary logistic regression. Notably, Cochran-Armitage provides similar results to the Wald test statistic for $H_0 : \beta = 0$ in the linear logit model, with β representing the coefficient of the independent variable (Agresti, 2002). This discovery led me to utilise binary logistic regression for direct, quantitative comparison of my observed $P(\mathrm{correct})$ behavioural trajectories. Dr C Priebe guided me on how to implement this statistical approach using the R programming language, which was necessary for compatibility with my existing R analysis pipeline.

My first step in analysing the data was transforming $P(\mathrm{correct})$. Because this dependent variable (here denoted as Y) is categorical, a log-odds (logit) transformation makes it easier to interpret Y as a linear function of the predictor variable x_{cast}:

$$logit(p) = ln(\frac{p}{1-p}) = \beta_{int} + \beta_{cast} x_{cast} + \epsilon \tag{3.4}$$

I used logistic regression to estimate the β_{cast} parameter from the data. The Wald z-score is the ratio of this estimate, $\hat{\beta}_{cast}$, to its standard error and I used this ratio test $H_0 : \beta_{cast} = 0$. Table 3.1 shows, for each condition, $\hat{\beta}_{cast}$, the standard error representing the uncertainty surrounding $\hat{\beta}_{cast}$, and the two-sided p-value based on the Wald test. I have ordered the table rows based on increasing $\hat{\beta}_{cast}$ value. I visually assessed Q-Q plots in R and observed approximate normality for all $\hat{\beta}_{cast}$ values. All of the standard errors are small, indicating that there is little variation in the estimate of the relationship between cast number and correct decision. A notable disadvantage of employing logistic

Table 3.1 Logistic regression results

Condition	$\hat{\beta}_{cast}$	s.e.($\hat{\beta}_{cast}$)	P(>|z|)
$C_{M,M}$	0.186	0.0385	1.37×10^{-6}
$C_{H,H}$	0.206	0.0369	2.26×10^{-8}
$C_{L,L}$	0.288	0.0508	1.42×10^{-8}
$E_{L,M}$	0.392	0.0484	5.34×10^{-16}
$E_{L,H}$	0.460	0.0460	1.52×10^{-23}

regression here is the assumption of a uniform relationship between the predictor and response variables over a range of values. While this makes it difficult to capture a potentially more complex relationship, I consider it to be a biologically feasible starting point.

In all five stimuli conditions, p-values suggest rejection of the null hypothesis as cast number appears to be a significant positive predictor of correct decisions. Knowing that the standard errors are approximately the same across conditions, I investigated significance between groups fit with different logit regressions by measuring the distance between $\hat{\beta}_{cast}$ coefficients based on standard error. Taking this approach, $E_{L,H}$ shows a bigger effect than $E_{L,M}$, supporting the qualitative observation of a faster plateau to equivalently high levels of $P(\text{correct})$ (Fig. 3.3b). Together, these two experimental conditions appear significantly larger than the three control conditions by, at minimum, approximately two standard errors. Given the qualitative differences already observed in the $C_{L,L}$ condition, it is unsurprising that $C_{L,L}$ shows a comparatively higher $\hat{\beta}_{cast}$ value than $C_{M,M}$ and $C_{H,H}$.

Although the $\hat{\beta}_{cast}$ coefficients neatly stratified the stimuli conditions into controls (lower values) and experimentals (higher values), it was important to formally compare the most ambiguous pair of conditions, $E_{L,M}$ and $C_{L,L}$. I achieved this with a logistic regression on the combined data from both conditions. This regression featured a second covariate, x_{cond}. This categorical variable represented the condition to which each data point belongs. The corresponding $\hat{\beta}_{cond}$ coefficient demonstrates significant predictive power (z-score: 2.55; p = 0.0108), reinforcing the previously described informal evaluation of a statistically significant difference in $\hat{\beta}_{cast}$ values between conditions $E_{L,M}$ and $C_{L,L}$.

The quantitative assessment I have outlined has shown both that cast number is a significant predictor of making the correct decision on a cast and that this predictive power is significantly more pronounced in the experimental conditions than the controls.

The improvement in larval performance with increasing head casts per reorientation manoeuvre seems to suggest that that larvae accumulate evidence for each stimulus alternative over repeated space samplings. In this way, repeated environmental sampling may function to reinforce the larva's association between its body position and the intensity of sensory information it receives. This result appears inconsistent with previous mechanistic explanations of navigational decisions during larval chemo- and thermotaxis. Those results suggest that head casts serve to compute temporal, derivative changes in sensory stimuli that, once having exceeded a threshold, trigger the end of a reorientation manoeuvre (Gomez-Marin et al., 2011; Kane et al., 2013; Klein et al., 2015). In Chapter 4 and Chapter 5, I explore two different computational models I built in an effort to elucidate the mechanistic underpinning the sequential sampling behaviour I observed in my sensory discrimination task.

3.2.3 Exploring cast duration data

A valuable metric I calculated for each larval reorientation manoeuvre was the duration of individual casts. Although all analysis I have described thus far focused on casts as the unit of sensory information, I thought it important to also consider whether the length of time larvae spent performing a cast influenced their decision on which direction to end the reorientation manoeuvre. I observed cast duration variability both within a single manoeuvre and between larvae; the latter perhaps due to varying nociceptive sensitivity across animals. This variability unfortunately precluded me from reasonably stratifying larvae into groups, making it challenging to formulate informative duration data analyses for individual casts. I discussed these challenges at length with Dr Ann Hermundstad, emphasising the importance of leveraging my available cast duration data to explore a continuous measure of head cast acceptance. Together, we settled on an approach that considered the total duration of larval reorientation manoeuvres as a function of stimulus intensity experienced over each cast. I wrangled all available duration data into the appropriate input format for a prototype MATLAB script that Dr A Hermundstad wrote to calculate integrated intensity for each larva. This calculation involved multiplying each cast's duration (in seconds) by the corresponding stimulus intensity for that cast (linearly scaled relative to least noxious intensity L), and summing this value over all casts the larva performed in its reorientation manoeuvre. Larvae that made a single cast were excluded from this analysis. I meticulously validated this prototype code, auditing the calculations for mathematical accuracy and biological relevance, and modifying various

code elements to ensure reproducibility of function outputs and compatibility with my existing analysis workflow.

Dr A Hermundstad and I calculated three separate metrics as a function of integrated intensity over a sliding data window. The window was used to reduce high frequency noise that otherwise may have occurred with non-overlapping data bins. Fig. 3.4a shows the probability of accepting a reorientation manoeuvre within the window centred at the indicated integrated intensity value. Because the probability is based on the total number of larvae that made more than one cast, this metric functions similarly to a probability density function of the data, though the sliding window means the probabilities do not sum to one. All five conditions exhibit a similarly shaped trajectory, with a peak occurring at approximately the same integrated intensity. This peak could suggest the existence of an internal drive to accept a head cast following exploration up to a specific "level" of integrated intensity. An opposing drive to find the "correct side" may explain why some animals continue exploring and the distribution tails extend toward higher integrated intensity values. If integrated intensity explained all the variability between conditions, then the data from all five conditions would overlap. The trajectories are, however, vertically translated from each other across much of the integrated intensity values. This vertical translation exhibits a rank ordering based on the summation of stimuli values within each condition, with the least noxious condition at the top of the graph and the most noxious at the bottom. Together, these qualitative features seem to suggest that larvae maintain a consistent baseline relationship between nociceptive stimuli and the rate of behavioural response across different stimuli conditions.

Two other duration data metrics further dissect the probability of acceptance across integrated intensity. Conditional on having accepted the reorientation manoeuvre within the integrated intensity window, Dr A Hermundstad and I calculated the probability that that acceptance was of the less noxious stimulus (odd-numbered cast; Fig. 3.4b) or, conversely, of the more noxious stimulus (even-numbered cast; Fig. 3.4c). Recall that because the choice is binary, $P(\text{less noxious}|\text{accept}) = 1 - P(\text{more noxious}|\text{accept})$. For both conditional probabilities, each stimuli condition exhibits a similarly shaped trajectory across integrated intensity. As before, this seems to suggest an element of behavioural consistency irrespective of the noxious environment. The rank ordering of $P(\text{accept})$ trajectories from Fig. 3.4a is not, however, preserved here. At any given integrated intensity, the $E_{L,H}$ trajectory is positioned above that of $E_{L,M}$ which itself lies above the three controls that eventually plateau near 0.5 (Fig. 3.4b). This pattern is understandably

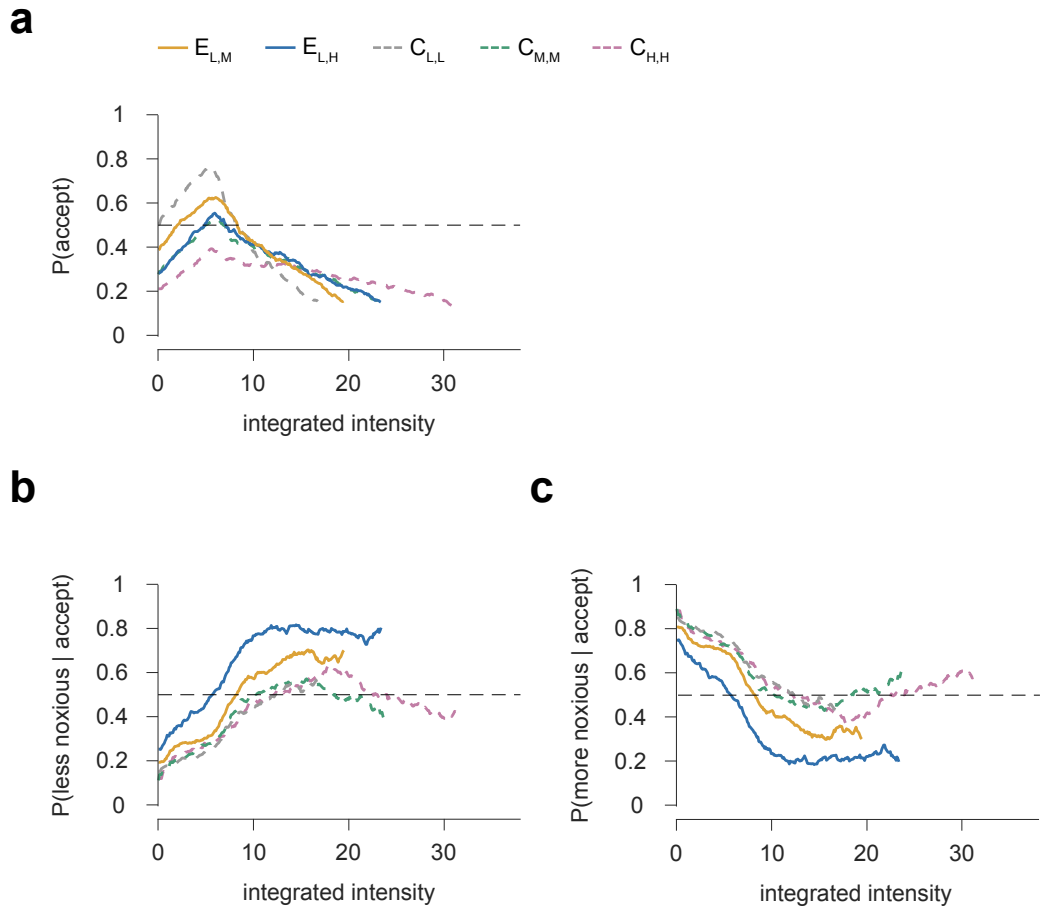

Fig. 3.4 Sensory discrimination task duration results. Experiments were performed as described in Chapter 2 through optogenetic stimulation of transgenic *ppk1.9-Gal4 x UAS-CsChrimson* larvae. Data are shown for experimental conditions in which there is an intensity differential between the two alternatives ($E_{L,M}$ and $E_{L,H}$). Data are also shown for the associated control conditions in which intensities are equivalent regardless of cast direction ($C_{L,L}$, $C_{M,M}$, and $C_{H,H}$). Individual larvae belong to only one condition. For each larva, integrated intensity was calculated by multiplying the duration of each cast with its corresponding stimulus intensity (linearly scaled relative to least noxious intensity L) and summing this value over all casts in the reorientation manoeuvre. By nature of the task design, the less noxious stimulus in a condition was presented on odd-numbered casts and the more noxious stimulus was presented on even-numbered casts. **a.** For larvae that made more than one cast, this is the probability of accepting a reorientation manoeuvre within the sliding window centred at the indicated integrated intensity value. **b.** For larvae that made more than one cast and accepted their reorientation manoeuvre within the indicated sliding window, this is the probability of that manoeuvre ending on the side corresponding to the less noxious stimulus. Conditions as in the legend shown in **a**. **c.** For larvae that made more than one cast and accepted their reorientation manoeuvre within the indicated sliding window, this is the probability of that manoeuvre ending on the side corresponding to the more noxious stimulus. Conditions as in the legend shown in **a**.

reversed in Fig. 3.4c. The near superposition of the three control conditions on one another compared to the separation of the experimental groups suggests that an intensity difference between the stimulus alternatives influences performance. Specifically, the more difficult the task, the lower the rate of acceptance of the less noxious stimulus. Simulated duration data sets could help predict what may be driving not only the vertical translation of stimuli conditions across acceptance probabilities in this larval data, but also the rate of change in acceptance over the sliding window. Gathering more behavioural data could also aid in investigating whether correlations exist between performing a particular cast within the reorientation manoeuvre and that cast's duration or duration of an individual cast and the continuous probability of its acceptance.

These duration data metrics are an avenue for which different drivers of the decision-making process can be explored. Sensory adaptation is one such driver. In the most fundamental sense, adaptation is a manifestation of sensory fatigue. Repeated stimulus exposure can cause neurons to adapt, calibrating their responses to pre-exposure baseline levels. Although an animal may then be unresponsive to the original stimulus, this gain modulation of the sensory system heightens the animal's sensitivity to novel stimuli (Rahnev and Denison, 2018). A notable behavioural consequence of adaptation is a reduction in the time required to make a decision (Theodoni et al., 2011). Any observed adaptation to noxious stimuli in my larval sensory discrimination task would therefore call into question my proposed mechanism of sensory evidence accumulation. Such an explanation is not without precedent, as *Drosophila* larvae have been shown to exhibit adaptation to both visual and olfactory stimuli (Cobb and Domain, 2000; Gepner et al., 2018). The results shown in Fig. 3.4b begin to address this concern. The difference between experimental and control conditions in the probability of accepting the less noxious stimulus seems to suggest that adaptation is not occurring. Said another way, if larvae engaged with my task adapted to the stimulus intensity they already experienced, then the probability of acceptance of the less noxious stimulus would be equivalent between experimental and control conditions across the range of integrated intensities. Further study is needed, however, to definitively rule out an adaptation mechanism underlying the observed decision-making process. My existing task design does not expose larvae in all conditions to persistent, unwavering stimulation before presenting a novel stimulus intensity. In the first instance, I would consider an alternative, adaptation-specific experiment paradigm that

necessitates intentional prolonged exposure to an adaptor stimulus to determine that stimulus' impact on future perceptual decisions (Theodoni et al., 2011).

Chapter 4

Modeling with Bayesian inference

4.1 Introduction

The sensory discrimination task results show a significant positive correlation between the number of head casts performed by larvae and the probability that those larvae end their exploration on the side with the less noxious stimulus. I hypothesise that continually gathering sensory evidence with each cast improves the larva's understanding of each side's true nociceptive level and facilitates better decisions with time. Computational models are one tool used to investigate such predictions about the decision rules that govern a behavioural readout. In the case of my evidence accumulation hypothesis, larvae may utilise the inherently noisy sensory information to update their internal beliefs about the true external environment.

It is common to conceptualise such perceptual computations in a Bayesian inference framework (Foley and Marjoram, 2017; Knill and Pouget, 2004; Ma et al., 2008). Bayesian decision theory was formulated from the application of Bayesian statistics to decision-making paradigms. A Bayesian observer infers the true value of environmental parameters by updating their prior beliefs about the world (either empirically derived or genetically encoded) based on newly acquired knowledge (Körding, 2007; McNamara et al., 2006; Trimmer et al., 2011; Valone, 2006). Although the parameters' true value will always remain unknown to the observer, they can use their inference results to inform a variety of task-dependent decisions.

A plethora of psychophysical experiments across various perceptual and motor tasks have shown that humans and monkeys perform Bayesian behavioural computations (Körding, 2007; Körding and Wolpert, 2004; Ma et al., 2008; Pouget et al., 2013; Stengård and van den Berg, 2019). Foraging decisions and mate selection are the

predominant ecological contexts in which other animals including various bird species (McNamara et al., 2006; Valone, 2006), honey bees (Naug and Arathi, 2007) and bumblebees (Foley and Marjoram, 2017), and tephritid fruit fly larvae (Morimoto, 2019) have been shown to exhibit Bayesian behaviour. Such results have encouraged further investigation into the underlying neuronal mechanisms (Ma et al., 2008). Although Bayesian behaviour does not necessitate that neurons themselves implement Bayesian algorithms (Ma and Jazayeri, 2014; McNamara et al., 2006; Trimmer et al., 2011), neuronal populations in several animal taxa have been shown to encode features of Bayesian inference calculations. Examples include sensory stimulus uncertainty in cue localisation (Rich et al., 2015), belief distributions over stimulus values in perceptual decision-making (Beck et al., 2008), and both the establishment and updating of predictions about the external environment in goal-directed behaviour (Funamizu et al., 2016). Computational modeling has also shed light on how single neurons can perform inference computations (Deneve, 2008; Pouget et al., 2013).

Beginning my investigation of the sensory discrimination task decisions with Bayesian inference is a reasoned approach given the vast space of models from which I could choose. Bayesian inference models are conceptually straight forward, specifying an optimal computation given assumptions about easily interpretable and separable model parameters. They also serve as a broad computational tool that generalises well across domains from perceptual decision-making to value-based decision-making, sensorimotor tasks, learning, and cognition (Beck et al., 2008; Ma and Jazayeri, 2014; Pouget et al., 2013). Arguments have been made, however, that classical Bayesian inference models of the brain do not account for constraints on the observer's time, memory capacity, computational power (Körding, 2007; Pouget et al., 2013; Trimmer et al., 2011), or fidelity in neurally encoding stimulus properties (Młynarski and Hermundstad, 2018). Indeed it is true that real biological systems must balance inference accuracy with the cost of such constraints (Tavoni et al., 2019), so neuronal computations often deviate from optimality (Rahnev, 2019; Rahnev and Denison, 2018). Even so, Bayesian inference models can provide insight into suboptimal (Młynarski and Hermundstad, 2018; Stengård and van den Berg, 2019) and even non-Bayesian algorithms. Furthermore, the availability of advanced genetic and imaging tools for studying *Drosophila* larvae make it more feasible to investigate the neuronal implementation of sensory inference and any subsequent behavioural decision rule.

I modeled the larval behaviour I observed as a Bayesian inference process, adapted to accommodate the sensory discrimination task structure. In this chapter, I

mathematically describe the model scenario (Section 4.2), the model observer's inferences about the true nociceptive stimulus levels (Section 4.3), and how the resulting posterior distributions map to the decision process (Section 4.4). I then extend these calculations through simplification of parameter values (Section 4.5) and consideration of special cases dictated by the larval sensory discrimination task (Section 4.6). Throughout the text, I describe how different model parameters, individually or in combination, could drive the observed difference in larval-averaged $P(\mathrm{correct})$ trajectories between stimuli conditions. I end the chapter with a qualitative comparison of the model predictions to my experiment data (Section 4.7).

I conceptualised and articulated the hypotheses and experiment quantities likely relevant for the decision process and collaborated with Dr Ann Hermundstad whose expertise was integral to mathematically translating those ideas into a Bayesian inference framework. Together, we formalised the design and iteratively implemented various aspects of the framework into software code. I validated all mathematical derivations, derived model predictions, and tested and finalised all code implementations.

4.2 The generative model

A generative model of Bayesian inference characterises how sensory observations are generated from probabilistic relationships between stimuli and associated world states. Features of my generative model mimic the larval sensory discrimination task in which cast direction triggers differential stimulus presentation. We model the environment as consisting of two sides ($s \in \{a, b\}$), each distinguished by a fixed nociceptive level, μ_s. For the purposes of this mathematical formulation, side a has a low nociceptive level ($\mu_a = \mu^L$) and side b has a high nociceptive level ($\mu_b = \mu^H$). Each nociceptive level is corrupted by Gaussian noise of variance σ_s^2, attributable to noise in the presentation of the nociceptive stimulus and/or the sensing and processing of such information by the larval central nervous system. The model observer's goal in interacting with this environment is to accept the less noxious side based on evolving beliefs about each side's true nociceptive level (Fig. 4.1a). To aid in its decision process, the observer alternates sampling stimulus evidence, x_s, from each side of the environment. The existence of sensory measurement noise on both sides precludes the observer from directly mapping x_s to the true nociceptive level from which it came. The observer must therefore estimate each side's true nociceptive level before deciding which direction it

believes to be less noxious. It performs this inference by combining incoming sensory evidence with its existing beliefs about the world, first building and then iteratively updating a posterior distribution over the possible stimulus levels.

We assume the observer knows each side's stimulus distribution follows a Gaussian. Additionally, we assume the observer knows the magnitude of the corresponding noise σ_s^2. The observer does not know the true nociceptive level for either side, which here is equivalent to the mean of that side's corresponding stimulus distribution (i. e. the specific values μ^L and μ^H). As is standard in other Bayesian inference constructions, we assume that the model observer enters the environment with a preestablished prior belief about the true nociceptive level. We also assume that this prior distribution takes the form $\mathcal{N}(\mu_0, \sigma_0^2)$. The observer maintains this prior for both sides of the environment until it acquires stimulus evidence to update its beliefs (Fig. 4.1b).

4.3 Inferring the true nociceptive levels

Analogous to larvae that alternate cast direction in the sensory discrimination task, the model observer alternates sampling the two sides of the environment. This is a distinguishing feature of this framework, as the observer updates their belief distribution over nociceptive levels for each side individually and sequentially (Fig. 4.1b). Because the observer follows the same inference process on both sides, I generalise the following formulation by using the subscript s where $s \in \{a, b\}$.

Let the observer acquire n_s samples of stimulus evidence, $\mathcal{E}_s = \{x_{s_1}, ..., x_{s_{n_s}}\}$, from side s of the environment. The observer calculates a posterior belief distribution about all possible levels of nociception on side s, $p(\mu_s|\mathcal{E}_s)$, by combining two quantities:

- The likelihood of observing a specific series of stimulus evidence, $\mathcal{E}_s$, conditioned on the belief that the true nociceptive level on that side equals μ_s. This takes the form

$$p(\mathcal{E}_s|\mu_s) \propto \exp\left(-\frac{n_s}{2\sigma_s^2}\left(\bar{\mathcal{E}}_s - \mu_s\right)^2\right) \propto \mathcal{N}\left(\bar{\mathcal{E}}_s; \mu_s, \frac{\sigma_s^2}{n_s}\right), \tag{4.1}$$

  where $\bar{\mathcal{E}}_s$ is the mean of n_s samples of acquired evidence (i. e. $\bar{\mathcal{E}}_s = \sum_{i=1}^{n_s} x_{s_i}/n_s$).

- The prior belief that the true nociception is μ_s:

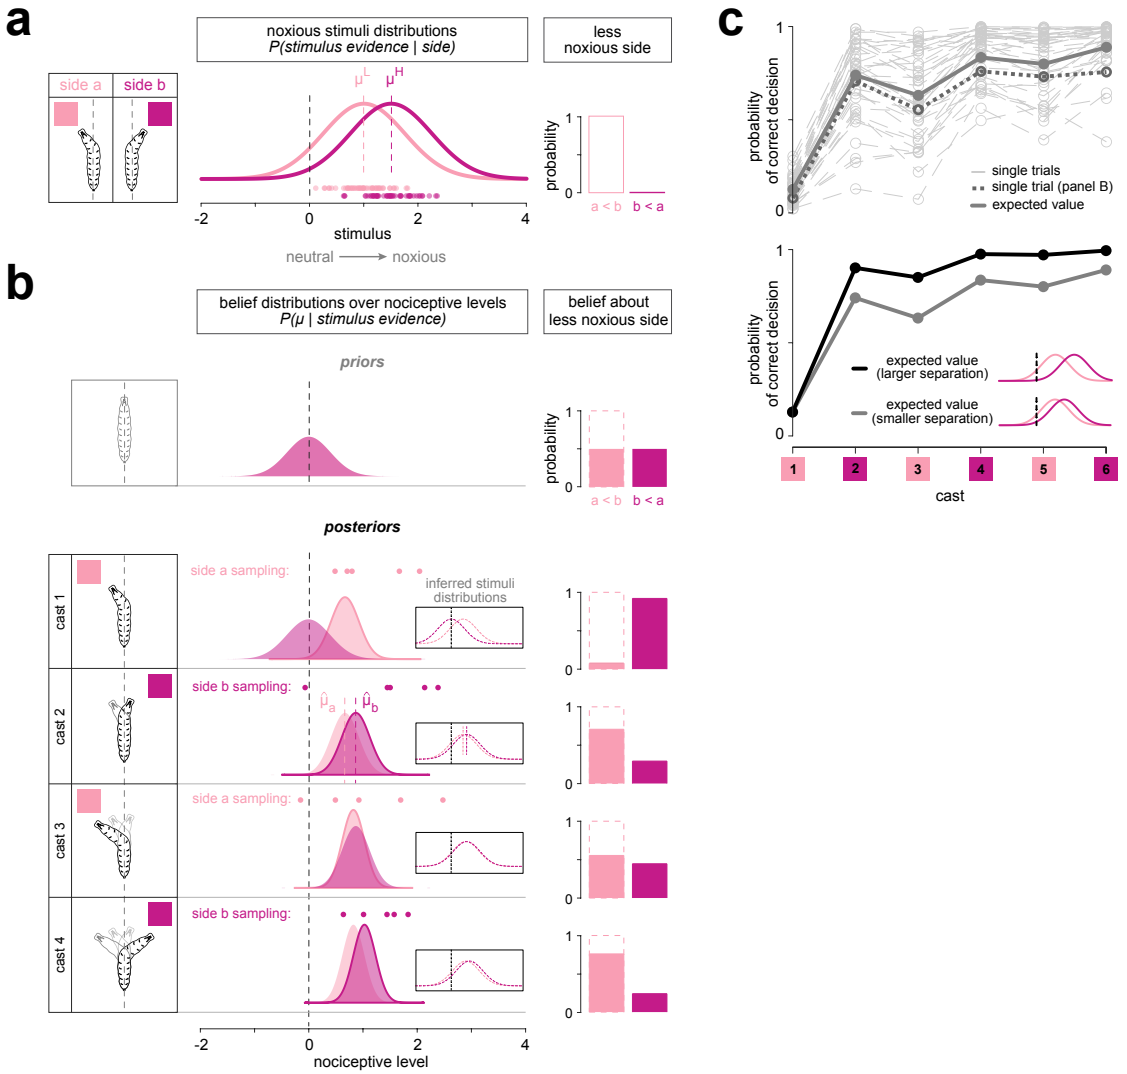

Fig. 4.1 Bayesian inference model construction. a. The generative model. (left) The model environment is divided into two sides, a and b, for comparison to the larval sensory discrimination task. (middle) Noxious stimuli on each side are Gaussian distributed with equal variance. More noxious stimuli are numerically represented by increasingly positive values. The mean values of the two distributions, μ^L and μ^H, are unknown to the observer and represent the true nociceptive level on each side. The close proximity of the distributions makes it difficult for the observer to accurately attribute stimulus evidence samples (colored circles immediately below the stimulus distributions) to the true nociceptive level that generated them. (right) In this framework, side a is defined as the less noxious side. **b.** The inference and decision processes. (left) Before "casting", the model observer holds identical belief distributions over the true nociceptive levels for each side of the environment (middle). (right) The means of both priors equal the neutral value 0, and the observer believes they are equally noxious. The observer then alternates acquiring stimulus evidence (colored circles) from one side at a time, beginning with a. The observer uses this information to derive a posterior probability distribution over nociceptive levels for that side. Only one posterior is updated with each "cast" (solid colored outline). The mean values of both posteriors ($\hat{\mu}_a$ and $\hat{\mu}_b$, shown only in cast 2 for visual simplicity) represent the most likely nociceptive level for that side. (middle inset) In this way, they also specify the observer's inferred stimuli distributions with fixed variance. It is assumed that the larva accepts a given side with the probability that that side is less noxious than the other. (right) The belief about the less noxious side is calculated using the posterior distributions. The light pink dashed box around the $a < b$ probability reflects the truth that side a is less noxious than side b. **c.** Illustrating the model readout. The probability of correct decision equals the probability of accepting side a. In the top graph, single trials represent the evolution of this probability as an individual observer casts, with the observer from panel **b** highlighted specifically. The expected value of the probability of correct decision represents the behaviour of a population of observers in the environment. A greater separation between μ^L and μ^H causes the probability of correct decision to rise more quickly than if the separation were smaller. Smaller separation as in the solid gray line in the top graph.

$$p(\mu_s) \propto \exp\left(-\frac{1}{2\sigma_0^2}(\mu_s - \mu_0)^2\right) \propto \mathcal{N}\left(\mu_s; \mu_0, \sigma_0^2\right) \tag{4.2}$$

Choosing a conjugate prior conveniently yields a posterior distribution of the same Gaussian form when applying Bayes' Rule (see Murphy (2007) for a detailed derivation of this closed-form solution):

$$p\left(\mu_s|\mathcal{E}_s\right) \sim \mathcal{N}\left(\mu_s; \hat{\mu}_s, \hat{\sigma}_s^2\right) \tag{4.3}$$

with mean $\hat{\mu}_s$ and variance $\hat{\sigma}_s^2$ given by

$$\hat{\mu}_s = \frac{\sigma_s^2\mu_0/n_s + \sigma_0^2\bar{\mathcal{E}}_s}{\sigma_s^2/n_s + \sigma_0^2} \tag{4.4}$$

$$\hat{\sigma}_s^2 = \frac{\sigma_s^2\sigma_0^2}{n_s\sigma_0^2 + \sigma_s^2} \tag{4.5}$$

Because this posterior has a Gaussian form, the most likely value equals the mean. In this way, $\hat{\mu}_s$ is the observer's estimate of the true nociceptive level on side s. Importantly, $\hat{\sigma}_s^2$ is not an estimate of the true variance but instead represents the observer's uncertainty over $\hat{\mu}_s$. As the observer acquires more evidence, the posterior mean approaches the sample mean (i. e. $\lim_{n_s\to\infty}\hat{\mu}_s = \bar{\mathcal{E}}_s$) and the uncertainty surrounding this nociceptive level estimate decreases (i. e. $\lim_{n_s\to\infty}\hat{\sigma}_s^2 = 0$) (Fig. 4.1b).

4.4 The decision process

The model observer relies on its inference of the nociceptive levels on each side to assess which side of the environment is less noxious. In this framework, a nociceptive level of 0 is neutral and increasingly positive values represent more nociception (Fig. 4.1a). Without loss of generality, I detail calculations for the probability that side a is less noxious than side b (i. e. $P(\mu_a < \mu_b)$) for all possible values of μ_b (Fig. 4.1b). I assume that the observer makes an optimal decision by accepting side a with this probability (Beck et al., 2008). Because I define side a as the less noxious side, the probability of accepting side a equals the probability of a correct decision (Fig. 4.1c):

$$
\begin{aligned}
&P(\text{accept state a}) = P(\text{correct}) \\
&= \int P(\mu_a < \mu_b) p\left(\mu_b | \mathcal{E}_b\right)\, d\mu_b \\
&= \int \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{\mu_b - \hat{\mu}_a}{\hat{\sigma}_a\sqrt{2}}\right)\right] \mathcal{N}\left(\mu_b; \hat{\mu}_b, \hat{\sigma}_b^2\right)\, d\mu_b \\
&= \int \frac{1}{2}\left[\frac{1}{\sqrt{2\pi\hat{\sigma}_b^2}} \exp\left(-\frac{(\mu_b - \hat{\mu}_b)^2}{2\hat{\sigma}_b^2}\right)\right] d\mu_b + \\
&\quad \int \frac{1}{2}\left[\operatorname{erf}\left(\frac{\mu_b - \hat{\mu}_a}{\hat{\sigma}_a\sqrt{2}}\right)\left(\frac{1}{\sqrt{2\pi\hat{\sigma}_b^2}} \exp\left(-\frac{(\mu_b - \hat{\mu}_b)^2}{2\hat{\sigma}_b^2}\right)\right)\right] d\mu_b \\
&= \frac{1}{2} + \int \frac{1}{2}\left[\operatorname{erf}\left(\frac{\mu_b - \hat{\mu}_a}{\hat{\sigma}_a\sqrt{2}}\right)\left(\frac{1}{\sqrt{2\pi\hat{\sigma}_b^2}} \exp\left(-\frac{(\mu_b - \hat{\mu}_b)^2}{2\hat{\sigma}_b^2}\right)\right)\right] d\mu_b
\end{aligned}
\tag{4.6}
$$

Eqn. 4.6 is simplified by applying the following identity for the indefinite integral of an error function multiplied by a Gaussian:

$$
\int_{-\infty}^{\infty} \operatorname{erf}\left(ax + b\right) \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) dx = \operatorname{erf}\left(\frac{a\mu + b}{\sqrt{1 + 2a^2\sigma^2}}\right) \tag{4.7}
$$

The resulting closed-form solution reveals each parameter's contribution to a single observer's decision (Fig. 4.1c):

$$
\begin{aligned}
P(\text{correct}) &= \frac{1}{2}\left(1 + \operatorname{erf}\left(\frac{(\hat{\mu}_b - \hat{\mu}_a)/\hat{\sigma}_a\sqrt{2}}{\sqrt{1 + \hat{\sigma}_b^2/\hat{\sigma}_a^2}}\right)\right) \\
&= \frac{1}{2}\left(1 + \operatorname{erf}\left(\frac{\hat{\mu}_b - \hat{\mu}_a}{\sqrt{2\left(\hat{\sigma}_a^2 + \hat{\sigma}_b^2\right)}}\right)\right) \\
&= \frac{1}{2}\left(1 + \operatorname{erf}\left(\frac{\gamma_b\bar{\mathcal{E}}_b - \gamma_a\bar{\mathcal{E}}_a}{\beta}\right)\right)
\end{aligned}
\tag{4.8}
$$

with γ and β defined, for convenience, as follows:

$$
\gamma_s = \frac{\sigma_0^2}{\sigma_s^2/n_s + \sigma_0^2} \tag{4.9}
$$

$$
\beta = \sqrt{2(\hat{\sigma}_b^2 + \hat{\sigma}_a^2)} \tag{4.10}
$$

To describe $P(\text{correct})$ dynamics across a population of observers, we average over the expected sampling variability in $\bar{\mathcal{E}}_a$ and $\bar{\mathcal{E}}_b$. Notably γ and β depend only on the known parameters σ_0^2, σ_a^2, σ_b^2, n_a, and n_b, so they do not contribute to variability between observers. Within this model framework, I know that the evidence samples, x_s, acquired by the observer are Gaussian distributed with mean μ_s (hereafter representing the fixed, true nociceptive level on side s) and variance σ_s^2. The central limit theorem states that the sampling distribution over the sample mean, $\bar{\mathcal{E}}_s$, is Gaussian distributed with the same mean μ_s and a scaled variance σ_s^2/n_s:

$$p(\bar{\mathcal{E}}_s) = \mathcal{N}\left(\bar{\mathcal{E}}_s; \mu_s, \frac{\sigma_s^2}{n_s}\right) \tag{4.11}$$

The expected value of $P(\text{correct})$, $\langle P(\text{correct})\rangle$, equals the summation of $P(\text{correct})$ over all possible series of evidence drawn from n_a and n_b samples, weighted by the probability of drawing each evidence series, $p(\bar{\mathcal{E}}_a)$ and $p(\bar{\mathcal{E}}_b)$:

$$\begin{aligned}
&\langle P(\text{correct})\rangle \\
=&\frac{1}{2}\iint\left(1+\text{erf}\left(\frac{\gamma_b\bar{\mathcal{E}}_b-\gamma_a\bar{\mathcal{E}}_a}{\beta}\right)\right)\mathcal{N}\left(\bar{\mathcal{E}}_a;\mu_a,\frac{\sigma_a^2}{n_a}\right)\mathcal{N}\left(\bar{\mathcal{E}}_b;\mu_b,\frac{\sigma_b^2}{n_b}\right)d\bar{\mathcal{E}}_a d\bar{\mathcal{E}}_b \\
=&\frac{1}{2}\int\left(1+\text{erf}\left(\frac{\gamma_b\bar{\mathcal{E}}_b-\gamma_a\bar{\mathcal{E}}_a}{\beta}\right)\right)\mathcal{N}\left(\bar{\mathcal{E}}_a;\mu_a,\frac{\sigma_a^2}{n_a}\right)d\bar{\mathcal{E}}_a\int\mathcal{N}\left(\bar{\mathcal{E}}_b;\mu_b,\frac{\sigma_b^2}{n_b}\right)d\bar{\mathcal{E}}_b
\end{aligned} \tag{4.12}$$

Applying the identity from Eqn. 4.7 to the integral over $\bar{\mathcal{E}}_a$, with $a = -\gamma_a/\beta$ and $b = \gamma_b\bar{\mathcal{E}}_b/\beta$, yields:

$$\begin{aligned}
&\int\left(1+\text{erf}\left(\frac{\gamma_b\bar{\mathcal{E}}_b-\gamma_a\bar{\mathcal{E}}_a}{\beta}\right)\right)\mathcal{N}\left(\bar{\mathcal{E}}_a;\mu_a,\frac{\sigma_a^2}{n_a}\right)d\bar{\mathcal{E}}_a \\
&= 1+\text{erf}\left(\frac{\gamma_b\bar{\mathcal{E}}_b/\beta-\gamma_a\mu_a/\beta}{\sqrt{1+2\gamma_a^2\sigma_a^2/\beta^2 n_a}}\right) \\
&= 1+\text{erf}\left(\frac{\gamma_b\bar{\mathcal{E}}_b-\gamma_a\mu_a}{\sqrt{\beta^2+2\gamma_a^2\sigma_a^2/n_a}}\right)
\end{aligned}$$

Eqn. 4.12 then simplifies to:

$$\langle P(\text{correct})\rangle = \frac{1}{2}\int\left(1+\text{erf}\left(\frac{\gamma_b\bar{\mathcal{E}}_b-\gamma_a\mu_a}{\sqrt{\beta^2+2\gamma_a^2\sigma_a^2/n_a}}\right)\right)\mathcal{N}\left(\bar{\mathcal{E}}_b;\mu_b,\frac{\sigma_b^2}{n_b}\right)d\bar{\mathcal{E}}_b$$

Reapplying Eqn. 4.7, with $a = \gamma_b / \sqrt{\beta^2 + 2\gamma_a^2 \sigma_a^2 / n_a}$ and $b = -\gamma_a \mu_a / \sqrt{\beta^2 + 2\gamma_a^2 \sigma_a^2 / n_a}$ yields the following:

$$\begin{aligned}
\langle P(\text{correct}) \rangle &= \frac{1}{2}\left(1 + \text{erf}\left(\frac{\gamma_b \mu_b - \gamma_a \mu_a}{\sqrt{\beta^2 + 2\gamma_a^2 \sigma_a^2 / n_a}} \frac{1}{\sqrt{1 + \frac{2\gamma_b^2 \sigma_b^2 / n_b}{\beta^2 + 2\gamma_a^2 \sigma_a^2 / n_a}}}\right)\right) \\
&= \frac{1}{2}\left(1 + \text{erf}\left(\frac{\gamma_b \mu_b - \gamma_a \mu_a}{\sqrt{\beta^2 + 2\gamma_a^2 \sigma_a^2 / n_a}} \frac{\sqrt{\beta^2 + 2\gamma_a^2 \sigma_a^2 / n_a}}{\sqrt{\beta^2 + 2\gamma_a^2 \sigma_a^2 / n_a + 2\gamma_b^2 \sigma_b^2 / n_b}}\right)\right) \\
&= \frac{1}{2}\left(1 + \text{erf}\left(\frac{\gamma_b \mu_b - \gamma_a \mu_a}{\sqrt{\beta^2 + 2\gamma_a^2 \sigma_a^2 / n_a + 2\gamma_b^2 \sigma_b^2 / n_b}}\right)\right)
\end{aligned} \tag{4.13}$$

This analytic result reveals that the probability of a correct decision directly depends on the separation between nociceptive levels μ_b and μ_a (Fig. 4.1c). Consider that any overlap in the two stimulus distributions means that a single sample x_s could have originated from either nociceptive level, increasing the difficulty of the inference process and lowering the probability of making the correct decision. This mirrors my hypothesis that the magnitude of the stimulus differential impacts the probability of larvae making a correct decision on a given head cast. Plugging in values for n_a and n_b samples of stimulus evidence into this closed-form solution yields an *in silico* $\langle P(\text{correct}) \rangle$ trajectory that can be qualitatively compared to the larval behavioural data (Section 4.7).

4.5 Simplifying assumptions

Establishing reasonable assumptions about the numeric value of several model parameters reduces the number of free parameters that an optimisation algorithm must fit to available data. Here I detail these assumptions and show how they simplify the $\langle P(\text{correct}) \rangle$ result from Eqn. 4.13.

The first assumption is that each model observer holds a neutral expectation about the nociceptive level prior to sampling (i. e. $\mu_0 = 0$) (Fig. 4.1b). Because the magnitude of the surrounding variance, σ_0^2, is encapsulated by both γ and β in Eqn. 4.13, I can rearrange parameters within the related Eqn.'s 4.5 and 4.9 to obtain

$$\hat{\sigma}_s^2 = \frac{\sigma_s^2 / n_s}{1 + \sigma_s^2 / (n_s \sigma_0^2)}$$

$$\gamma_s = \frac{1}{1 + \sigma_s^2/(n_s\sigma_0^2)},$$

This reveals a consistent relationship between σ_0^2, σ_s, and n_s, with σ_0^2 never appearing in isolation. Therefore, without loss of generality, I can set $\sigma_0^2 = 1$.

I constructed the model such that the number of stimulus evidence samples the observer acquires on each side of the environment, n_s, is equal to the number of "casts" the observer makes to that side during the reorientation manoeuvre, N_s, each comprising T samples:

$$n_s = N_s * T \tag{4.14}$$

As in the larval sensory discrimination task, the global number of casts, N, is the summation of casts to both sides:

$$N = N_a + N_b \tag{4.15}$$

In this way, T controls the sampling rate per cast, as it is independent of both cast number and side. Although it is possible to make T a fittable parameter, I do not have enough information about the timescale of sampling to explicitly constrain its value. For numerical simplicity, I set $T = 1$. In this way, the resulting model approximates a continuous stream of evidence acquisition as discrete.

Finally, I assume that the nociceptive stimulus distributions from which the observer alternates sampling have equivalent Gaussian variances ($\sigma_a^2 = \sigma_b^2 = \sigma^2$).

Taken together, these four assumptions ($\mu_0 = 0;\ \sigma_0^2 = 1;\ T = 1;\ \sigma_a^2 = \sigma_b^2 = \sigma^2$) allow me to further simplify Eqn. 4.13 to the following:

$$\langle P(\text{correct})\rangle = \frac{1}{2}\left(1 + \text{erf}\left(\frac{\gamma_b\mu_b - \gamma_a\mu_a}{\sqrt{2\sigma^2\left(\left(\gamma_b + \gamma_b^2\right)/n_b + \left(\gamma_a + \gamma_a^2\right)/n_a\right)}}\right)\right) \tag{4.16}$$

$$\text{with } \gamma_s = \frac{1}{1 + \sigma^2/n_s}$$

This not only reemphasises that the difference in μ_b and μ_a drives $\langle P(\text{correct})\rangle$, but also reveals that σ^2, the sensory measurement noise surrounding both nociceptive levels, directly influences those trajectory dynamics.

4.6 Model predictions and subtypes

Here, I analytically outline special cases of $\langle P(\text{correct})\rangle$ (Eqn. 4.16) to illustrate important model predictions at defined stages of the larval sensory discrimination task. I also describe different model subtypes based on an observer's evolving belief over a set number of previous casts.

Consider the inference and decision processes immediately following the observer's first "cast" (i. e. $N = 1$). If this cast is to side a, the observer is naïve to side b. Mathematically, this means that the observer's posterior distribution for side b equals the prior (i. e. setting $n_b = 0$ in Eqn.'s 4.4 and 4.5 yields $\hat{\mu}_b = \mu_0$ and $\hat{\sigma}_b^2 = \sigma_0^2$). Finding $\langle P(\text{correct})\rangle$ following this first cast requires recalculating $P(\text{correct})$ from Eqn. 4.8 by setting $n_b = 0$:

$$P(\text{correct}) = \frac{1}{2}\left(1 - \text{erf}\left(\frac{\gamma_a \bar{\mathcal{E}}_a}{\sqrt{2\left(\sigma_0^2 + \hat{\sigma}_a^2\right)}}\right)\right)$$

then summing $P(\text{correct})$ over all possible series' of evidence drawn from n_a samples in this first cast, weighted by the probability of drawing that evidence series, $p(\bar{\mathcal{E}}_a)$:

$$\langle P(\text{correct})\rangle = \frac{1}{2}\int \left(1 - \text{erf}\left(\frac{\gamma_a \bar{\mathcal{E}}_a}{\sqrt{2\left(\sigma_0^2 + \hat{\sigma}_a^2\right)}}\right)\right) \mathcal{N}\left(\bar{\mathcal{E}}_a; \mu_a, \frac{\sigma_a^2}{n_a}\right) d\bar{\mathcal{E}}_a$$

Applying the identity in Eqn. 4.7, with $a = \gamma_a/\sqrt{2(\sigma_0^2 + \hat{\sigma}_a^2)}$ and $b = 0$ and further simplifying with the assumptions I outlined in Section 4.5 yields:

$$\begin{aligned}
\langle P(\text{correct})\rangle &= \frac{1}{2}\left(1 - \text{erf}\left(\frac{\gamma_a \mu_a}{\sqrt{2(\sigma_0^2 + \hat{\sigma}_a^2)}} \frac{1}{\sqrt{1 + \frac{2\gamma_a^2\sigma_a^2}{2(\sigma_0^2 + \hat{\sigma}_a^2)n_a}}}\right)\right) \\
&= \frac{1}{2}\left(1 - \text{erf}\left(\frac{\gamma_a \mu_a}{\sqrt{2(\sigma_0^2 + \hat{\sigma}_a^2 + \gamma_a^2\sigma_a^2/n_a)}}\right)\right) \\
&= \frac{1}{2}\left(1 - \text{erf}\left(\frac{\gamma \mu_a}{\sqrt{2\left(1 + \gamma\sigma^2 + \gamma^2\sigma^2\right)}}\right)\right)
\end{aligned} \tag{4.17}$$

$$\text{with } \gamma = \frac{1}{1 + \sigma^2}$$

Notably, this shows a direct influence of μ_a on $\langle P(\text{correct})\rangle$ following the first cast. A lower value of μ_a generates a greater probability of correct decision (i. e. a lower level of nociception leads to a higher probability of accepting the first cast).

Now consider the dynamics that arise from casts beyond the first one, during which the observer samples each side individually. If the observer's first cast is to side a, $n_a = n_b + T_a$ when N is odd. During these odd-numbered casts where $n_a \neq n_b$, $\langle P(\text{correct})\rangle$ follows Eqn. 4.16. I observe a different special solution to $\langle P(\text{correct})\rangle$ when $n_a = n_b = n$ (i. e. for all even-numbered casts). Substituting n for both n_1 and n_2 in Eqn. 4.16, I obtain:

$$\langle P(\text{correct})\rangle = \frac{1}{2}\left(1 + \text{erf}\left(\frac{\gamma\,(\mu_b - \mu_a)}{\sqrt{(4\sigma^2/n)\,(\gamma + \gamma^2)}}\right)\right) \tag{4.18}$$

$$\text{with } \gamma = \frac{1}{1 + \sigma^2/n}$$

If the model observer accumulated evidence from both states simultaneously, one would observe the trajectory mapped out by computing Eqn. 4.18 over all values of n. Because the observer instead alternates sampling from each side individually, Eqn. 4.18 provides a qualitative envelope for the $\langle P(\text{correct})\rangle$ trajectory, with deviations caused by the inherently asynchronous updates of each side's posterior distribution.

I also consider the effect of the two nociceptive stimulus distributions being identical (i. e. $\mu_a = \mu_b = \mu$). In this control condition case, Eqn. 4.16 simplifies to:

$$\langle P(\text{correct})\rangle = \frac{1}{2}\left(1 + \text{erf}\left(\frac{(\gamma_b - \gamma_a)\,\mu}{\sqrt{2\sigma^2\left(\left(\gamma_b + \gamma_b^2\right)/n_b + \left(\gamma_a + \gamma_a^2\right)/n_a\right)}}\right)\right) \tag{4.19}$$

$$\text{with } \gamma_s = \frac{1}{1 + \sigma^2/n_s}$$

This analytic form predicts that when $n_a = n_b = n$, $\langle P(\text{correct})\rangle = 0.5$.

The aforementioned special cases are derived from the default model framework (i. e. *all casts*) in which each observer continues to update its posterior belief distributions for each side with each newly acquired piece of evidence. In Section 4.7, I also show results from other model subtypes in which we implemented an element of forgetting. In these cases, the observer iteratively updates its posterior beliefs over a select number of sequential casts before reverting to updating the original prior; as if it

were experiencing the environment for the first time. In the *two casts* subtype, the sequential casts include the current cast and the cast just before, occurring on the opposite side. The two posterior distributions (one for each side) therefore revert back to the prior with each cast. Under the *three casts* subtype, the three sequential casts include the current cast and the two immediately preceding it. The resulting imbalance in the number of updates to the posterior distribution before reverting to the prior persists from $N = 3$ onward, favouring the side the observer is currently sampling.

4.7 Fitting the model

I have outlined a Bayesian inference model that closely mimics the larval sensory discrimination task design. Although I have explored the predictions this model makes about key features of the decision readout, these analytic solutions are not always easily interpretable when viewed by themselves. An important next step is qualitative evaluation of what the model can capture when fit to the behavioural data.

Fitting model-generated $\langle P(\text{correct})\rangle$ trajectories to larval $P(\text{correct})$ trajectories requires finding the best possible values for the unknowns, μ_s and σ^2 (Eqn. 4.16). The parameters μ^L, μ^M, and μ^H account for the three unique nociceptive levels (Section 2.1.3) across all five possible sensory discrimination task stimuli conditions and σ^2 accounts for their equivalent Gaussian variance (Section 4.5). Informed by the task design, the Bayesian inference framework dictates the following relationships between these parameters:

$$\begin{aligned} \mu^L, \mu^M, \mu^H &< 0 \\ \mu^H - \mu^M &< 0 \\ \mu^M - \mu^L &< 0 \\ \sigma^2 &> 0 \end{aligned}$$

My metric for assessing whether a set of parameter values produces a good model fit to the behavioural data is the following calculated error:

$$\begin{aligned}&\sum_{N=1}^{6}\sqrt{\mathrm{nLarvae}_{E_{L,M}}\big|_N}\left(\langle P(\mathrm{correct})\rangle\big|_{N,\,\mu^L,\,\mu^M,\,\sigma}-P(\mathrm{correct})_{E_{L,M}}\big|_N\right)^2\\+&\sum_{N=1}^{6}\sqrt{\mathrm{nLarvae}_{E_{L,H}}\big|_N}\left(\langle P(\mathrm{correct})\rangle\big|_{N,\,\mu^L,\,\mu^H,\,\sigma}-P(\mathrm{correct})_{E_{L,H}}\big|_N\right)^2\end{aligned}\tag{4.20}$$

where $E_{L,M}$ and $E_{L,H}$ are the two experimental stimuli conditions, N is the cast number, and $\mathrm{nLarvae}$ is the number of larvae that performed that cast (normalised by the total number of larvae in the stimuli condition). Including the weighted sum of squared errors for both experimental $P(\mathrm{correct})$ trajectories increases the ratio of empirical data points to free parameters and reduces over-fitting. Unlike a standard objective function, this weighted alternative accounts for the decrease in total larvae with increasing cast number and better balances the model fit between both conditions.

I evaluated the fit of Bayesian inference trajectories to the experimental conditions' behavioural data under each of the three model subtypes outlined in Section 4.6. Using MATLAB's `fmincon` solver, I found numerical values for each of the four free parameters that, together, minimise Eqn. 4.20, subject to the constraints outlined above. I randomly initialised parameter values for 1000 separate optimisations to discourage trapping the solution set in a single local minimum. For each model subtype, at least 89% of the 1000 separate optimisations yielded sum of squared error values within 10% of the smallest calculated (Table 4.1). Because these results reveal strong consistency of the model output across instantiations, I selected the solution parameters corresponding to that with the smallest error. The differences in magnitude between these μ^L, μ^M, and μ^H solutions (when normalised to μ^L for comparison) are neither equivalent across model subtypes nor do they show a discernible monotonic relationship between model subtypes (Table 4.1). Furthermore, none of the subtypes show an exact match between these normalised μ^L, μ^M, and μ^H parameters and the true stimulus power values (Section 2.1.3). These results, however, do not preclude the model from capturing notable qualitative features of the behavioural data.

All three model subtypes predict the equivalent $P(\mathrm{correct})$ value observed across both experimental conditions, $E_{L,M}$ and $E_{L,H}$, on the first cast and mimic the increased probability of correct decision on cast two (Fig. 4.2a). Adjusting the amount that individual observers forget does, however, impact the model output trajectories at later casts. Notably, the probability of correct decision in the *all casts* subtype oscillates between consistently lower values on odd-numbered casts compared to subsequent

Table 4.1 Bayesian inference optimisation results

Model Subtype	Number within 10% of Min	Error over $E_{L,M}$, $E_{L,H}$	Solution Parameters			
			μ^L	μ^M	μ^H	σ
all casts	966	0.0141	1.1119	1.7256	2.0951	0.9533
three casts	979	0.0249	0.5884	0.8136	0.9904	0.4277
two casts	891	0.0261	2.3624	3.7781	4.9792	1.6954

even-numbered casts. A function of the asynchronous updates of each side's posterior distribution, this model feature yields a close match the observed decrease in $E_{L,M}$ and $E_{L,H}$ trajectories at cast three. The envelope of the model's oscillation also decreases with cast number (Fig. 4.2a), reflecting a decrease in uncertainty over the posterior beliefs. Overall, the construction of the *all casts* model is guaranteed to show an increase in the probability of correct decision over head casts when faced with a stimulus differential between sides. Although this yields reasonable qualitative fits to the larval behaviour trajectories, there is notable deviation between the model and the larval behaviour at casts four and five in $E_{L,M}$ and cast six in $E_{L,H}$. In contrast, the *two casts* and *three casts* subtypes show a performance ceiling, where the probability of correct decision does not increase beyond a set threshold. Under the *three casts* subtype, $\langle P(\text{correct})\rangle$ consistently oscillates between the same two values from cast three onward, with odd-numbered casts exhibiting lower probabilities than even-numbered. In the model fit to $E_{L,H}$, the bounds of the oscillation envelope are vertically shifted to greater values of $\langle P(\text{correct})\rangle$ than when fit to $E_{L,M}$. Finally, the *two casts* subtype reliably plateaus to the same probability value from cast two onward, with the plateau reaching a higher value when fit to $E_{L,H}$ versus $E_{L,M}$ (Fig. 4.2a).

An additional point of comparison between the three model subtypes is the weighted sum of squared error values (a goodness-of-fit measure) between the behavioural data trajectories and their corresponding model trajectories. For $E_{L,M}$ and $E_{L,H}$, *all casts* has the lowest summed error of the model subtypes (Fig. 4.2a and Table 4.1).

Using the optimised parameters acquired from the model fits to the experimental conditions, I generated $\langle P(\text{correct})\rangle$ trajectories for comparison to the control data. In this way, the control data is a test for whether the model under these parameters can predict previously unseen data. As predicted by the inverse relationship between $\langle P(\text{correct})\rangle$ and μ_a in Eqn. 4.17, all three model subtypes capture the decrease in probability of correct decision on cast one from $C_{L,L}$ to $C_{M,M}$ to $C_{H,H}$ (Fig. 4.2b).

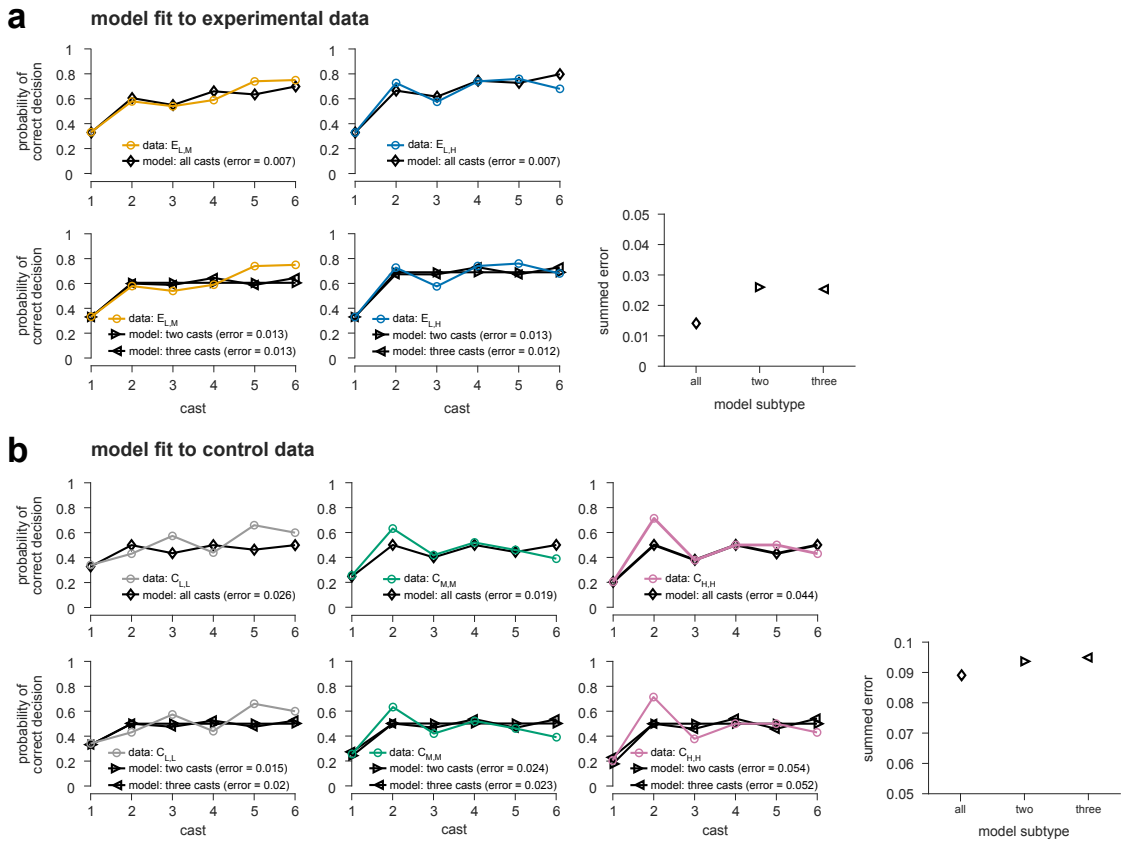

Fig. 4.2 Bayesian inference model fits to behavioural data. The *in silico* $\langle P(\text{correct})\rangle$ population readouts (black traces), evaluated at condition-specific solution parameters, are plotted against corresponding behavioural data (coloured traces, as in Fig. 3.3b). The results from three different model subtypes are shown. Each subtype differs in the number of casts over which the observer updates its posterior belief before it forgets the acquired evidence and reverts to the prior distribution (for details, see Section 4.6). For each subtype, the sum of squared errors (shown in each plot legend) are themselves summed across conditions and plotted on the bottom right of each panel. **a.** The solution parameters corresponding to these *in silico* $\langle P(\text{correct})\rangle$ population readouts were found by running 1000 optimisations to both experimental data traces and selecting the instantiation with the lowest summed error (see Eqn. 4.20). **b.** Control data was used to test the solution parameters on previously unseen data.

Irrespective of the control condition to which it is fit, the *all casts* subtype shows oscillation in the probability of correct decision, with a maximum value of 0.5 on all even-numbered casts (analytically predicted in Eqn. 4.19). While this model subtype captures some of the oscillations in the behavioural data signatures, it misses peak probability values that occur above 0.5 in all three control conditions. The *three casts* subtype also oscillates when fit to each control, but does not max out at 0.5. The envelope of this oscillation shows a small increase in size from conditions $C_{L,L}$ to $C_{M,M}$ to $C_{H,H}$. For all three control conditions, the *two casts* subtype plateaus at 0.5 from cast two onward.

The weighted sum of squared error values between the control condition trajectories and their corresponding model trajectories were higher than that of the experimental conditions. I expected this result, given that I did not use the controls themselves to optimise the model parameters. Notably, optimising the Bayesian inference parameters to the controls (results not shown) yields nearly identical quantitative and qualitative fits to that which I observe for all three control conditions under the experimental-optimised *all casts* framework.

The *all casts* subtype qualitatively and quantitatively outperforms the *three casts* and *two casts* subtypes when fit to all conditions, with the exception of $C_{L,L}$. This Bayesian inference framework well captures features of the experimental conditions consistent with the significant increase I observe in $P(\mathrm{correct})$ (see Chapter 3). It does not, however, fully explain the control conditions; a shortcoming most apparent at low cast numbers (Fig. 4.2b). If larvae rely solely on an optimal Bayesian strategy to guide their performance in this sensory discrimination task, then this model would be capable of better fits to previously unseen data, irrespective of the stimuli conditions under which those data were generated. These results prompted me to consider whether behavioural trajectory features could be fully captured by a computationally simpler model construction in which the animal does not rely on comparison between the two sides but still relies on information accumulation over time.

Chapter 5

Modeling with acceptance pressure

5.1 Introduction

While Bayesian inference was a principled framework to explore the behaviour strategy in my larval sensory discrimination task, it did not explain notable features of the probability of correct decision in the control conditions (Section 4.7). This encouraged me to reevaluate whether the larval decision computation results from a comparison of the estimated noxious intensity between both sides of the environment; a key mechanistic assumption of my Bayesian inference construction. In this chapter, I propose a simpler algorithmic process to explain the larval behaviour data. This process still relies on information integration with time, but does not rely on the observer making any comparison calculations. Features of this alternate model were derived from the relationship between experimental and control conditions' $P(\mathrm{accept})$ values when matched by both cast number and noxious stimulus intensity. For example, I compared $E_{L,M}$ to $C_{L,L}$ on odd-numbered casts and $E_{L,M}$ to $C_{M,M}$ on even-numbered casts. Similar behavioural readouts occur between these matched groups on early casts (but not on later ones) despite larvae in the separate stimuli conditions being exposed to different sequences of alternating noxious intensities (Fig. 3.3b). If the control data explains the experimental data in this way, then perhaps a comparison between sides does not influence the larva's evaluation of the current cast. Following consultation with Dr Ann Hermundstad, I explored whether larval behaviour in the sensory discrimination task is driven by a combination of conflicting pressures: aversion to the current nociceptive level and either constant or increasing pressure to end the reorientation manoeuvre. Dr A Hermundstad's expertise was integral to implementing this framework into software code, which I tested and finalised.

5.2 Model construction and predictions

The readout of this "acceptance pressure" model is the overall probability of acceptance on a given cast, calculated as a combination of opposing pressures to accept and reject:

$$\begin{aligned} P(\text{accept})_{\text{in silico}} &= (a_x) - (r_x) \\ &= (mx + b) + (r_x) \end{aligned} \tag{5.1}$$

In this general form, a_x represents a pressure to accept a given cast, $x \in [1, 6]$, and takes the form of a linear polynomial with slope m and intercept b. Here, r_x represents an aversion to cast x based on its relative stimulus intensity ($r_x \in \{r^L, r^M, r^H\}$ where the superscripts denote low, medium, and high intensities like those across the five stimuli conditions outlined in Section 2.1.3).

Here, I outline three subtypes of this model. The defining feature for each subtype is the form of the dependent variable within a_x. The *instant noci* subtype is defined by a constant value of a_x, irrespective of cast number. Consequently, r_x is the primary driver of change in the calculated $P(\text{accept})_{\text{in silico}}$ value:

$$P(\text{accept})_{\text{instant noci}} = (m + b) - (r_x) \tag{5.2}$$

In the *cast only* subtype, a_x linearly increases as a function of cast number, x. $P(\text{accept})_{\text{in silico}}$ is therefore influenced by both the noxious stimulus intensity and the cast on which it is experienced:

$$P(\text{accept})_{\text{cast only}} = (mx + b) - (r_x) \tag{5.3}$$

What is not captured by the use of cast number as a dependent variable is the notion that any observed differences in stimulus acceptance may also be driven by the sequence of intensities experienced up until the accepted cast. We formulated the third model subtype in an attempt to disentangle the history of prior experience from cast number itself. In this *integ noci* subtype, the dependent variable is the total summed nociceptive stimulus intensity from the beginning of the reorientation manoeuvre through the current cast, x. As with the *cast only* subtype, the pressure to end the manoeuvre, a_x, monotonically increases with time:

$$P(\text{accept})_{\text{integ noci}} = \left(m \sum_{i=1}^{x} r_i + b \right) - (r_x) \tag{5.4}$$

5.3 Fitting the model

Because we conceptualised this model based on whether the control conditions' data explain those of the experimental, I fit model-generated $P(\text{accept})_{\text{in silico}}$ trajectories to larval $P(\text{accept})$ trajectories across all three control stimuli conditions, $C_{L,L}$, $C_{M,M}$, and $C_{H,H}$. Using MATLAB's `fminunc` solver, I found numerical values for each of the four unknown model parameters m, b, r^M, and r^H that minimise the following weighted sum of squared errors:

$$\begin{aligned}
&\sum_{x=1}^{6} \sqrt{\text{nLarvae}_{C_{L,L}}\big|_x} \left(P(\text{accept})_{\text{in silico}}\big|_{x;\, m;\, b;\, r^L} - P(\text{accept})_{C_{L,L}}\big|_x \right)^2 \\
+&\sum_{x=1}^{6} \sqrt{\text{nLarvae}_{C_{M,M}}\big|_x} \left(P(\text{accept})_{\text{in silico}}\big|_{x;\, m;\, b;\, r^M} - P(\text{accept})_{C_{M,M}}\big|_x \right)^2 \\
+&\sum_{x=1}^{6} \sqrt{\text{nLarvae}_{C_{H,H}}\big|_x} \left(P(\text{accept})_{\text{in silico}}\big|_{x;\, m;\, b;\, r^H} - P(\text{accept})_{C_{H,H}}\big|_x \right)^2
\end{aligned} \tag{5.5}$$

where x is the cast number and $nLarvae$ is the number of larvae that performed that cast (normalised by the total number of larvae in the stimuli condition). By nature of the decrease in total larvae with increasing cast number, this weighted objective function places greater emphasis on fitting earlier casts within each trajectory. To reduce the number of unknown parameters, we set r^L equal to 1 and assume it to be a baseline against which the magnitudes of r^M and r^H are scaled. By fitting m and b across the three control conditions, we assume that the mechanism of the pressure to accept is the same regardless of the environment in which the larva is casting.

For each model subtype (Section 5.2), I performed 1000 separate optimisations following Eqn. 5.5. At no point did I constrain parameters by upper or lower bounds or linear relationships with other parameters. Instead, I randomly initialised m, b, r^M, and r^H at the beginning of each optimisation to discourage trapping the solution set in a single local minimum (see Chapter 2 for additional model implementation details). For each subtype, all 1000 separate optimisations yielded error values within 10% of the smallest calculated (Table 5.1). Because of this consistency in model output across

Table 5.1 Acceptance pressure optimisation results

Model Subtype	Number within 10% of Min	Error over Controls	Solution Parameters			
			r^M	r^H	m	b
instant noci	1000	0.1417	1.1036	1.1345	0.8529	0.6408
cast only	1000	0.0354	1.1243	1.1593	0.0646	1.3337
integ noci	1000	0.0334	1.1463	1.1912	0.0579	1.3515

optimisations, I selected for subsequent analysis the solution parameters corresponding to that with the smallest error (Table 5.1). Although the scaling of the r_x parameters to each other does not precisely match that of the stimulus intensities I set in my sensory discrimination task, their rank ordering ($r^L < r^M < r^H$) matches that of the true intensities (see Section 2.1.3). This occurs across model subtypes and is especially notable since the model parameters were unconstrained during optimisation.

The $P(\text{accept})_{\text{in silico}}$ trajectories resulting from these optimised parameters are plotted against each control conditions' behavioural data in the bottom row of Fig. 5.1b. As was predicted mathematically, the *instant noci* model subtype produces an unchanging probability of acceptance when fit to any of the control data trajectories. The observed difference in the y-intercept of this model fit between $C_{L,L}$, $C_{M,M}$, and $C_{H,H}$ is solely caused by the r^L, r^M, and r^H values, respectively. The high error values (see plot legends in Fig. 5.1b) further emphasise this model subtype's poor goodness-of-fit to the larval control data. The *cast only* and *integ noci* model subtypes exhibit near-identical performance against the control data (Fig. 5.1b, top row). This is not surprising when one considers that the optimised r_x values over which *integ noci* is integrating are close to 1. The fit of their model trajectories to condition $C_{L,L}$ is qualitatively poor, with the larval behavioural data deviating from the model's linear increase at casts two, three, and (most notably) six. This translates to a nearly ten-fold increase in computed error compared to these models' fits against $C_{M,M}$ and $C_{H,H}$. In these latter cases, both model subtypes yield near matches to the behavioural data on the probability of accepting cast one. The linear increase in $P(\text{accept})_{\text{in silico}}$ characteristic of both model subtypes explains almost all of the data in these two control conditions, with the exception of cast five in $C_{M,M}$ and cast four in $C_{H,H}$. In the case of the *integ noci* subtype, this suggests larvae follow a mechanism of "giving up" in which they accept a cast after exceeding a threshold of integrated stimulus intensity.

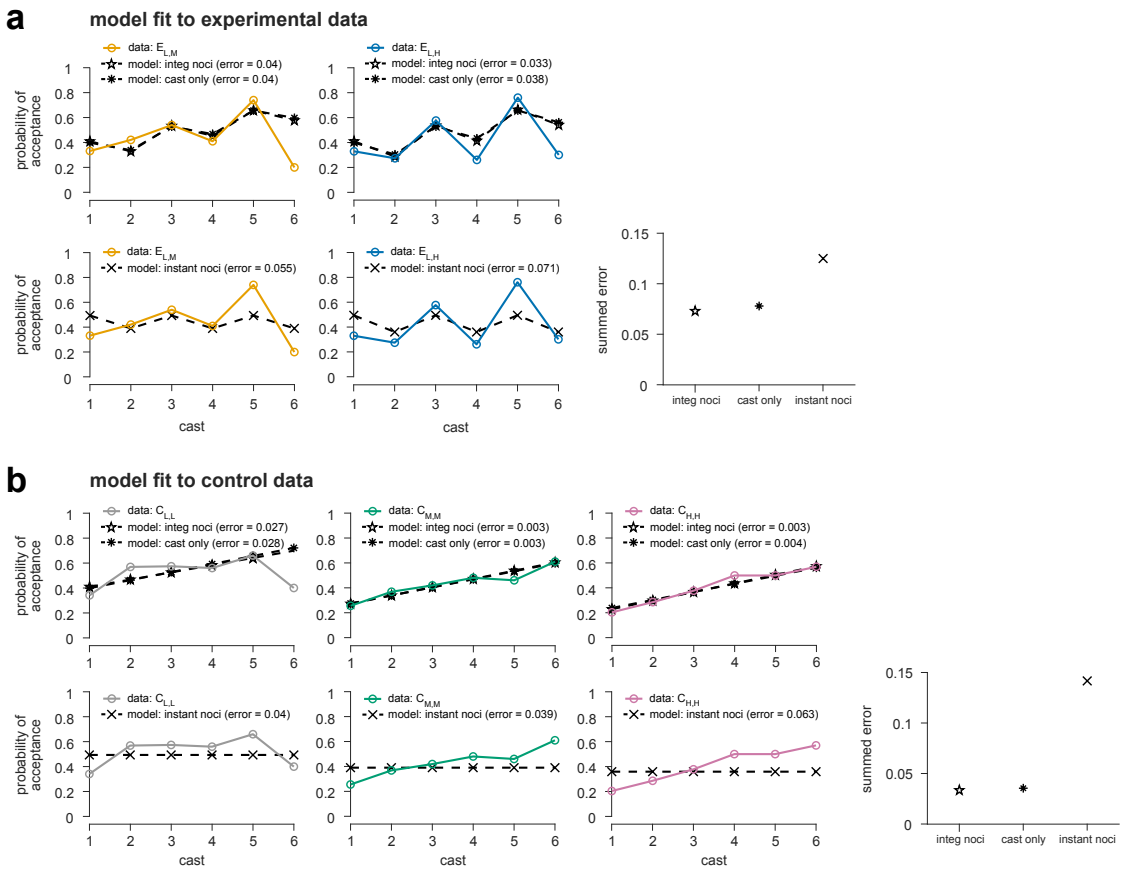

Fig. 5.1 Acceptance pressure model fits to behavioural data. The *in silico* $P(\text{accept})$ population readouts (black traces), evaluated at condition-specific solution parameters (Table 5.1), are plotted against corresponding behavioural data (coloured traces, as in Fig. 3.3a). The results from three different model subtypes are shown, each differing by the dependent variable driving the pressure to accept a given cast (for details, see Section 5.2). Each plot legend details the normalised sum of squared errors between the model-generated data and that of the displayed stimuli condition. The plot in the bottom right of each panel displays the summation of these errors across either both experimental conditions (as in **a**) or the three control conditions (as in **b**). **a.** Data from the experimental conditions was used to test the solution parameters on previously unseen data. **b.** The solution parameters were found by running 1000 optimisations over all three control condition data traces and selecting the run with the lowest summed error (see Eqn. 5.5).

Using the optimised parameters from the best model fits to the control conditions, I generated condition-specific $P(\text{accept})_{\text{in silico}}$ trajectories for comparison to the experimental data (Fig. 5.1a). These served as test cases, as the experimental data were unseen by the model subtypes during optimisation. In each model subtype, the probability of acceptance predictably oscillates up and down as a function of the alternating value of r_x from one cast to the next. As was the case when fit to the control data, the *instant noci* subtype does not exhibit a global increase in the probability of acceptance with casts, making it a poor fit to the experimental conditions' data (Fig. 5.1a, bottom row). The trajectories formed by the *cast only* and *integ noci* subtypes are nearly identical to one another, both globally increasing with cast number. When fit to the

experimental data, these subtypes explain the oscillatory behaviour in the larval population but fall short of capturing some of the highest and lowest probability values (e. g. casts two and six in $E_{L,M}$ and casts four and six in $E_{L,H}$). Both the *integ noci* and *cast only* model subtypes qualitatively and quantitatively outperform *instant noci* when fit to any of the five sensory discrimination task conditions.

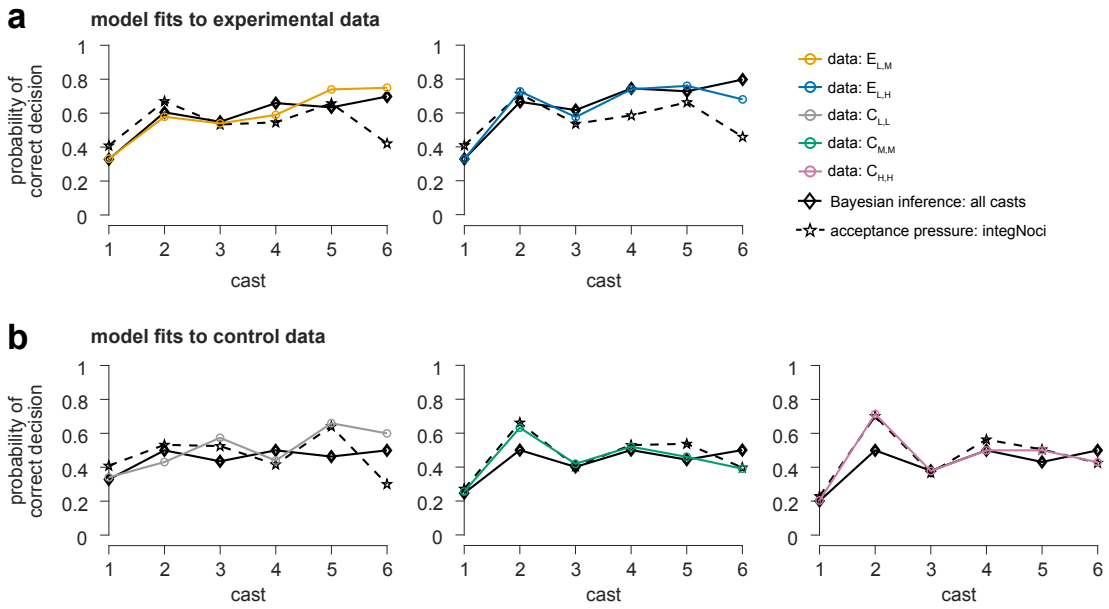

Fig. 5.2 Best Bayesian inference and acceptance pressure model fits. The *in silico* population readouts (black traces) for probability of correct decision from the *all casts* Bayesian inference model subtype and the *integ noci* acceptance pressure model subtype plotted against corresponding behavioural data (coloured traces, as in Fig. 3.3b). Bayesian inference model data as in Fig. 4.2. Acceptance pressure model data translated from that shown in Fig. 5.1 using Eqn. 3.3.

I pursued the formulation and investigation of this acceptance pressure model in response to the shortcomings of my Bayesian inference mechanism in fully explaining the larval control data (for details, see Chapter 4). I have included Fig. 5.2 for visual comparison of the best performing Bayesian inference fits to those of the acceptance pressure model. As per the relationship defined in Eqn. 3.3, I translated the *integ noci* model results from probability of acceptance to probability of correct decision. While the mechanism of Bayesian inference shows qualitatively better fits to the experimental data trajectories than acceptance pressure (Fig. 5.2a), acceptance pressure is a better fit to the control stimuli conditions than Bayesian inference (Fig. 5.2b). To favour one model framework over the other would necessitate that all five experimental and control conditions be well explained through the associated mechanism. Although I cannot accept either my Bayesian inference or acceptance pressure model based on this criterion, I take the beginning of Chapter 7 to speculate on what their contrasting model

results may suggest about the larval behavioural strategy employed during this sensory discrimination task.

Chapter 6

High-throughput operant conditioning

6.1 Preamble

The closed-loop system I employed to study larval decision-making behaviour in my sensory discrimination task was low-throughput, tracking only a single larva at a time. Expanding the functionality of this system design to simultaneously track and stimulate several larvae would facilitate efficient and thorough investigation of other nociceptive stimulus intensities, genetic lines, and sensory modalities. Such a high-throughput system would also enable development of other behavioural tasks requiring closed-loop stimulation in response to larval action, furthering investigation of working memory in larvae. In close collaboration with Kristina Klein, I developed a novel tracking system capable of closed-loop optogenetic and thermogenetic stimulation of up to 16 freely behaving larvae. Proof-of-principle experiments highlighted the capabilities of this high-throughput system and we observed, for the first time, reward-driven operant conditioning in the *Drosophila* larva. Foundational to operant learning is the learned association between an animal's self-directed behaviour and an external stimulus. A similar action-outcome principle may apply to larval decision-making behaviour in my sensory discrimination task, albeit on a much shorter time scale. The neural circuits supporting behaviour in these separate paradigms may themselves be shared, making the results from operant conditioning experiments valuable for guiding future exploration of the sensory discrimination task mechanism.

My role in this collaboration was shepherding the development of the high-throughput system from the last stages of proof-of-principle prototyping through to the final product and subsequent behavioural experiments. The larval features and behaviour algorithms employed on this multi-larva closed-loop system were originally

developed by Dr Jean-Baptiste Masson, K Klein, and myself for use on the single-larva closed-loop tracker (Fig. 2.1). K Klein updated these techniques and developed the multi-larva Behaviour Programme software (Fig. 6.1e), with the intention of performing high-throughput larval operant conditioning through a paradigm of her own design (Fig. 6.3a). Given my expertise in operating and maintaining the single-larva tracker hardware and software, I was best suited to efficiently and effectively manifest developments on the multi-larva system. My role was especially critical given K Klein's physical absence from the Janelia Research Campus and the timeline of project deadlines prior to physically moving the system to our new lab space at the University of Cambridge. I worked alongside K Klein to debug and test software code components, validate the stability and reproducibility of optogenetic and thermogenetic stimulation procedures on live animals, develop the laser stimulus calibration pipeline, run operant conditioning experiments (Fig. 6.3), and analyse data for manuscript preparation. I was solely responsible for performing spatial and intensity calibration of both the optogenetic and thermogenetic stimulation systems, validating their accuracy prior to all data collection. I also led efforts to address underlying hardware processing delays and improve the integrity of data backup and storage. For more granular detail on my contributions, see Section 2.2.

Achieving project deliverables necessitated regular, in-depth, long-distance communication between myself and K Klein. In these discussions, I provided observational feedback on design implementations, often suggesting additional software or hardware modifications based on observed system operation with live animals. I also drew upon my extensive knowledge of closed-loop larval experiment considerations during in-person conversations that I facilitated with our engineering project collaborator, Dr Lakshmi Narayan. Dr L Narayan worked alongside K Klein throughout the prototyping phase to choose and quality control test all hardware components for image processing, behaviour detection, and optogenetic stimulation. He had also programmed the FPGA to work together with the host computer to read the raw camera images, detect eligible objects, and extract and process object features (Fig. 6.1e). When I joined the project, I was responsible for relaying many of my and K Klein's design goals to Dr L Narayan and confirming that his resulting software modifications met our collective experimental needs. I also led in-person consultations with Dr Michael Winding on thermogenetic stimulation dynamics, given his work with Dr Chris McRaven in designing the housing and control mechanisms for the laser light source.

K Klein and I deposited a version of this work on bioRxiv as co-first authors ('Serotonergic Neurons Mediate Operant Conditioning in *Drosophila* Larvae', Klein et al. (2021)), with Dr L Narayan, Dr M Winding, Dr J-B Masson, and Dr Marta Zlatic listed as contributing authors. I contributed to the writing, editing, and revision of this preprint text. Within this thesis Chapter 6, I enclose a modified version of the preprint, with revised conclusions I have written in accordance with additional tests I performed including bootstrap statistical sensitivity analyses and analysis of larger data sets reflecting fewer larval exclusion criteria. For the purposes of readability in Chapter 6, I have changed the instances of 'bend' in the preprint to 'cast' and moved the associated methods and materials text to Chapter 2 of this thesis.

6.2 Introduction

Animals must rapidly alter their behaviour in response to environmental changes. An important adaptation strategy is associative learning (Dickinson, 1981; Rescorla, 1988), in which an animal learns to predict an US by the occurrence of a CS. The US is often a punishing or rewarding event such as pain or the discovery of a new food source (Pavlov, 1927). The nature of the CS distinguishes two major associative learning types: classical conditioning (Pavlov, 1927) and operant conditioning (Skinner, 1938; Thorndike, 1911).

In classical conditioning, the CS is an inherently neutral environmental stimulus such as a sound, odour, or visual cue. Pairing with an appetitive or aversive US leads to learned approach or avoidance of the CS in the future. Many vertebrates (Andreatta and Pauli, 2015; Braubach et al., 2009; Brown et al., 1951; Jones et al., 2005) and invertebrates (Alexander et al., 1984; Cognigni et al., 2018; Davis, 2005; Scherer et al., 2003; Takeda, 1961; Vinauger et al., 2014; Vogt et al., 2014; Wen et al., 1997) can make these associations. Across the animal kingdom, neural circuits have been identified as convergence sites for the external CS and the rewarding or punishing US (Caroni, 2015; Gründemann and Lüthi, 2015; Hawkins and Byrne, 2015; Heisenberg et al., 1985; Owald and Waddell, 2015; Tonegawa et al., 2015). In classical conditioning of both larval and adult *Drosophila*, the mushroom body (MB) brain area serves this purpose (Cognigni et al., 2018; Heisenberg, 2003; Heisenberg et al., 1985; Owald and Waddell, 2015; Rohwedder et al., 2016; Saumweber et al., 2018; Vogt et al., 2014). In each larval brain hemisphere, the CS is encoded by a subset of the approximately 110 Kenyon cells (KCs) (Aso et al., 2014a; Berck et al., 2016; Campbell et al., 2013; Eichler et al., 2017; Honegger et al., 2011; Lin et al., 2014; Owald and Waddell, 2015; Turner et al., 2008),

which synapse onto 24 MB output neurons (MBONs) driving approach or avoidance (Aso et al., 2014b; Eichler et al., 2017; Owald et al., 2015; Perisse et al., 2016; Plaçais et al., 2013; Saumweber et al., 2018; Séjourné et al., 2011; Shyu et al., 2017). KC to MBON connection strength is modulated by dopaminergic and octopaminergic neurons, which represent the rewarding or punishing US (Honjo and Furukubo-Tokunaga, 2009; Saumweber et al., 2018; Schroll et al., 2006; Schwaerzel et al., 2003; Vogt et al., 2014; Waddell, 2013). Activation of the MB-innervating PAM cluster dopaminergic neurons serves as both a necessary and sufficient reward signal in classical conditioning (Cognigni et al., 2018; Liu et al., 2012; Rohwedder et al., 2016; Vogt et al., 2014; Waddell, 2013).

In operant conditioning, the CS is an animal's own action (Skinner, 1938; Thorndike, 1911). After memory formation, the animal can predict the outcome of its behaviour and bias future action selection accordingly, usually to maximise reward and avoid punishment (Skinner, 1938). This behavioural adaptation can facilitate novel action sequences (Fee and Goldberg, 2011; Nottebohm, 1991; Topál et al., 2006) and, in some cases, repetitive, high-frequency motor activity (Corbett and Wise, 1980; Jin and Costa, 2010; Lovell et al., 2015; Olds and Milner, 1954). Such observations have wider implications for understanding diseases including obsessive-compulsive disorder and addiction (Balleine et al., 2015; Everitt et al., 2018; Joel, 2006). Invertebrates are also capable of operant conditioning (Abramson et al., 2016; Booker and Quinn, 1981; Brembs, 2003; Hoyle, 1979; Nuwal et al., 2012). Despite countless operant conditioning experiments across species using various CS–US combinations, the underlying neural mechanisms remain poorly understood. For an animal to associate an action with its outcome, behavioural information must converge with circuits encoding positive or negative valence. Although vertebrate basal ganglia-like structures exemplify this (Balleine et al., 2009; Fee and Goldberg, 2011; Redgrave et al., 2011), some learned action-outcome associations do not require the brain (Booker and Quinn, 1981; Grau et al., 1998; Horridge, 1962). Operant conditioning may hence occur in more than one area of the CNS. It is also unclear to what extent learning at these sites is mediated by synaptic plasticity (Gómez-Pinilla et al., 2007; Joynes et al., 2004; Lovinger, 2010; Reynolds and Wickens, 2002; Surmeier et al., 2007) versus changes in the intrinsic excitability of individual neurons (Brembs et al., 2002; Dong et al., 2006; Nargeot et al., 1997, 2009; Shen et al., 2005). We aim to establish the *Drosophila* larva as a tractable model system for studying the neural circuit mechanisms underlying operant conditioning.

Drosophila melanogaster larvae perform various different actions. Typically, when exploring an environment, a larva alternates between crawling via forward peristalsis (Heckscher et al., 2012) and casting its head once or more to the left or right (Gomez-Marin et al., 2011; Kane et al., 2013; Luo et al., 2010; Fig. 6.1a). In the presence of nociceptive stimuli, larvae exhibit escape behaviour. While the most common response is an increase in head casting away from undesirable conditions, including extreme temperature (Lahiri et al., 2011; Luo et al., 2010), light (Kane et al., 2013), or wind (Jovanic et al., 2019), larvae also retreat from aversive sources using backward peristalsis (Heckscher et al., 2012; Kernan et al., 1994; Masson et al., 2020; Vogelstein et al., 2014; Fig. 6.1a). The fastest escape response is rolling, where the larva moves laterally by curling into a C-shape and quickly turning around its own body axis (Hwang et al., 2007; Ohyama et al., 2013; Robertson et al., 2013; Fig. 6.1a). In nature, rolling is only observed after exposure to a strong noxious stimulus, such as heat or a predator attack (Ohyama et al., 2015; Robertson et al., 2013; Tracey et al., 2003).

Powerful genetic toolkits have advanced the observation and manipulation of larval behaviour at the cellular level, making *Drosophila* larvae particularly well-suited for studying the neural mechanisms underlying learning. In *Drosophila*, individual neurons are uniquely identifiable, with morphology and function preserved across animals (Jefferis et al., 2007; Marin et al., 2002; Skeath and Thor, 2003; Wong et al., 2002). Together with tissue-localised protein expression afforded by the GAL4-UAS binary expression system (Brand and Perrimon, 1993; Fischer et al., 1988), this has yielded neuron-specific GAL4 drivers (Jenett et al., 2012; Luan et al., 2006; Pfeiffer et al., 2010) that reproducibly target the same group of cells in each individual. Adding fluorescent markers helps pinpoint a neuron's location and reveal its anatomical features (Lee and Luo, 1999), while producing light-sensitive channelrhodopsins and temperature-sensitive ion channels facilitates optogenetic (Lima and Miesenböck, 2005; Zemelman et al., 2002) or thermogenetic (Hamada et al., 2008; Kitamoto, 2001) modulation of neural activity. Furthermore, the larva's compact CNS has made it feasible to manually reconstruct neurons and their synaptic partners from a larval electron microscopy (EM) volume (Berck et al., 2016; Eichler et al., 2017; Fushiki et al., 2016; Jovanic et al., 2016, 2019; Larderet et al., 2017; Ohyama et al., 2015; Schlegel et al., 2016), giving rise to a full wiring diagram of the MB (Eichler et al., 2017; Eschbach et al., 2020a,b).

There is overwhelming evidence that larvae are capable of classical conditioning. They can be trained to approach an odour paired with a gustatory reward (Hendel et al., 2005; Kudow et al., 2017; Niewalda et al., 2008; Schleyer et al., 2011), or avoid an odour

paired with light (von Essen et al., 2011), electric shock (Aceves-Piña and Quinn, 1979; Tully et al., 1994), heat (Khurana et al., 2012), vibration (Eschbach et al., 2011), or the bitter compound quinine (Apostolopoulou et al., 2014b; Gerber and Hendel, 2006). Light can also be a CS: innate avoidance of light and preference for darkness (Sawin-McCormack et al., 1995) can be modulated when paired with reward or punishment (Gerber et al., 2004; von Essen et al., 2011). It has remained an open question, however, whether *Drosophila* larvae can form action–outcome associations and where in the CNS these memories are formed.

Conducting operant conditioning with larvae requires real-time behaviour detection such that reward or punishment can be administered with minimal delay (Fig. 6.1b). Single-animal closed-loop trackers have recently been developed (Schulze et al., 2015; Tadres and Louis, 2020). However, the efficiency of training paradigms would improve with automated US delivery and simultaneous conditioning of multiple animals. Therefore, we here introduce a high-throughput tracker for *Drosophila* larvae with real-time behaviour detection and closed-loop stimulation. Efficiency of the setup stems from the simultaneous, real-time, behaviour detection for up to 16 freely moving larvae, and targeted closed-loop optogenetic and thermogenetic stimulus delivery with full intensity control and minimal delay.

6.3 Results

6.3.1 High-throughput closed-loop tracker

Hardware design

Designing an automated operant conditioning protocol for the *Drosophila* larva was challenging due to the larva's physical characteristics. We excluded partial immobilisation protocols similar to the ones used to condition adult *Drosophila* navigation through virtual environments (Brembs, 2011; Nuwal et al., 2012; Wolf and Heisenberg, 1991; Wolf et al., 1998). We instead built a high-throughput multi-larva tracker combining live computer vision behaviour detection with closed-loop control of US delivery in response to unrestricted larval behaviour.

All hardware resided within an optically opaque enclosure to ensure experiments were performed without environmental light. Larvae moved freely on an agarose plate, backlit from below by an infrared LED and observed from above through a

high-resolution camera (Fig. 6.1c). A Camera Link communication protocol interfaced with a high-performance FPGA, which itself interacted with the host computer. The FPGA and the host computer performed image processing, behaviour detection, and stimulus calculation (Fig. 6.1d).

Our operant conditioning paradigm targeted individual larvae performing specific behaviours. Optogenetic stimulation was achieved by directing red light through two DMDs which were programmed to project small 1 cm^2 squares at the location of individual larvae. Both DMDs, which were positioned to project over the entire plate area, were operated simultaneously (Fig. 6.1c).

Thermogenetic stimulation of individual larvae was achieved by directing a 1490 nm IR laser beam through a two-axis scanning galvanometer mirror positioning system (Fig. 6.1c), a technique previously used to stimulate single adult flies (Bath et al., 2014; Wu et al., 2014). Because the 1490 nm wavelength is well-absorbed by water (Curcio and Petty, 1951), larvae exposed to the IR beam were rapidly heated. We took advantage of the galvanometer's high scanning velocity to rapidly cycle the beam between four larvae (Fig. 6.1d).

Software architecture

Several computer vision algorithms exist for real-time tracking of freely behaving animals. Stowers et al. (2017) and Krynitsky et al. (2020) developed software for tracking mice, and Mischiati et al. (2015) developed high-speed tracking of single dragonflies in three-dimensional space. There are numerous tracking frameworks for adult *Drosophila*, some requiring the flies to move within a two-dimensional plane (Donelson et al., 2012; Straw and Dickinson, 2009) while others detect the three-dimensional position of single (Fry et al., 2008) or multiple (Grover et al., 2008; Straw et al., 2011) flies. The Multi-Worm Tracker (MWT) software developed by Swierczek et al. (2011) is suitable for simultaneously tracking a large number of *C. elegans* and has been adapted to analyse *Drosophila* larvae reactions in response to various stimuli (Jovanic et al., 2019; Masson et al., 2020; Ohyama et al., 2013; Vogelstein et al., 2014).

Operant conditioning requires live behaviour detection to trigger delivery of reward or punishment. Numerous algorithms have been developed to analyse offline behavioural recordings of animals such as *C. elegans* (Gupta and Gomez-Marin, 2019; Huang et al., 2006; Stephens et al., 2008), zebrafish larvae (Mirat et al., 2013; Reddy et al., 2020), adult *Drosophila* (Berman et al., 2014; Branson et al., 2009; Dankert et al., 2009; Katsov and Clandinin, 2008; Klibaite et al., 2017; Robie et al., 2017), bees

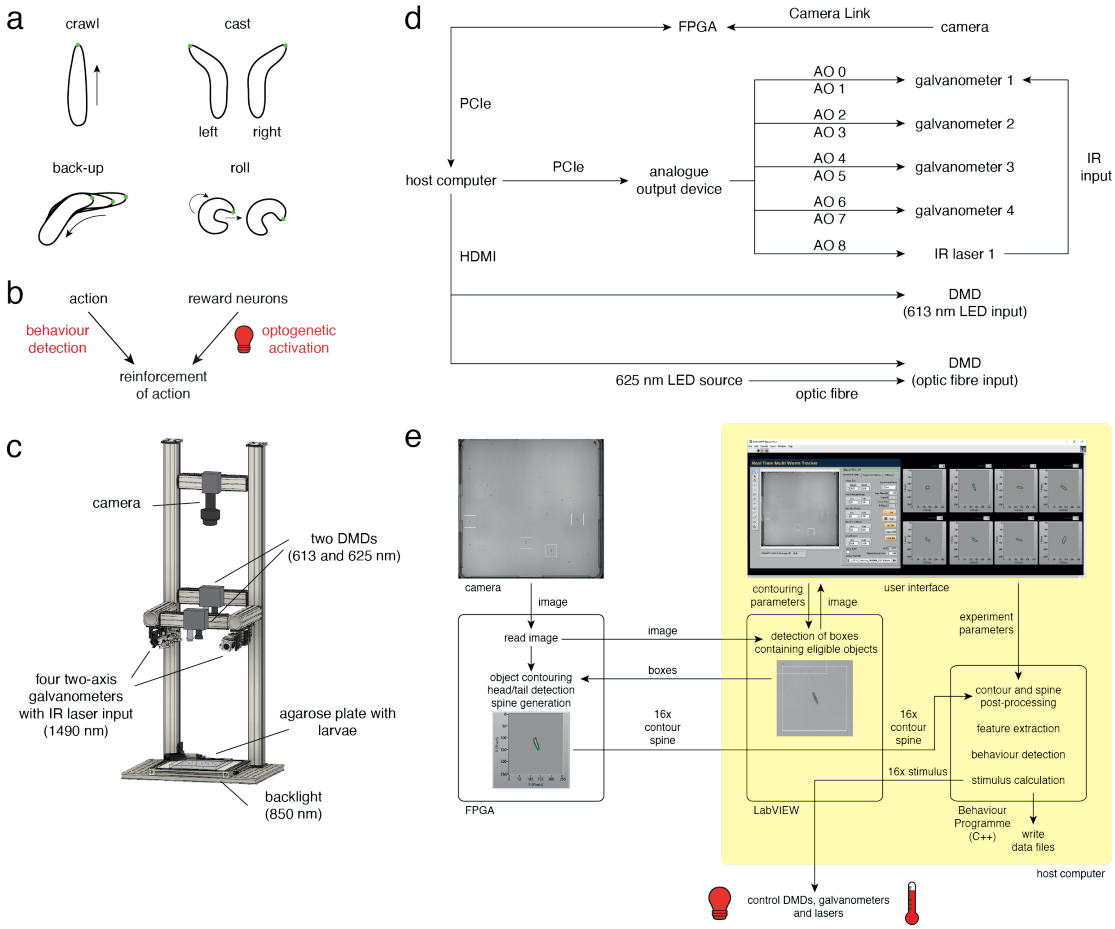

Fig. 6.1 High-throughput operant conditioning in *Drosophila* larvae. a. Behavioural repertoire of *Drosophila* larvae. Schematics show the four most prominent actions displayed by *Drosophila* larvae (crawl, left and right cast, back-up and roll). The larval contour is displayed as a black outline with a green dot marking the head. **b.** In fully automated operant conditioning, an action of interest was reinforced by coupling real-time behaviour detection with optogenetic activation of reward circuits. **c.** High-throughput tracker schematic showing the relative positions of the agarose plate, backlight, camera, digital micromirror devices (DMDs), and galvanometers. IR: infrared. **d.** Block diagram of hardware components. AO: analogue output, FPGA: field-programmable gate array. **e.** Data flow between software elements.

(Veeraraghavan et al., 2008), and mice (Luxem et al., 2020; Mathis et al., 2018; van Dam et al., 2020). The *Drosophila* larva has also attracted attention due to analytical challenges surrounding its deformable body and limited set of distinguishing features (Denisov et al., 2013; Gershow et al., 2012; Gomez-Marin et al., 2011; Luo et al., 2010; Masson et al., 2020; Ohyama et al., 2013, 2015; Vogelstein et al., 2014). Most of these approaches are not ideal to run in real time or require a mix of past and future information to provide reliable behaviour detection (Gomez-Marin et al., 2011; Masson et al., 2020). More generally, machine learning based methods have gained momentum in providing both supervised and unsupervised approaches to behaviour analysis. It is worth noting a recent trend in developing unsupervised learning methods (e. g. Graving and Couzin, 2020; Luxem et al., 2020).

While real-time behaviour detection of casts and runs has been developed for a single animal (Schulze et al., 2015), our study of operant conditioning in freely behaving *Drosophila* larvae required efficient, real-time behaviour detection of multiple animals. We built a system to simultaneously track up to 16 larvae in real time, using LabVIEW for the user interface and algorithm implementation (Fig. 6.1e). Instrumental to this software architecture was the fast image processing speed afforded by FPGA-based parallelisation (Li et al., 2011; Soares dos Santos and Ferreira, 2014; Zhang et al., 2017). Neuroscientists have adapted FPGA's real-time analysis capabilities (Chiuchisan, 2013; Shirvaikar and Bushnaq, 2009; Uzun et al., 2005; Yasukawa et al., 2016) to track rats (Chen et al., 2005), zebrafish larvae (Cong et al., 2017), and fluorescently labelled neurons in freely behaving *Drosophila* larvae (Karagyozov et al., 2018). In our system, the FPGA and host computer worked together to read the raw camera images, detect eligible objects, and extract and process object features (i. e. contour, head and tail position, and body axis) (Fig. 6.1e). Larval body shape, velocity, and direction of motion facilitated robust behaviour detection which, in turn, drove closed-loop optogenetic and thermogenetic stimulation. All relevant experiment parameters and time-series data were output for offline analysis through a custom MATLAB framework (see Chapter 2).

Optogenetic and thermogenetic stimulation efficiency verified by behavioural readout

We conducted proof-of-principle experiments to ensure that our set-up could be successfully used for optogenetic stimulation (Fig. 6.2a). Ohyama et al. (2015) have identified two GAL4 lines expressed in neurons whose activation triggers strong rolling behaviour. *69F06-Gal4* drives expression in command neurons for rolling, whereas

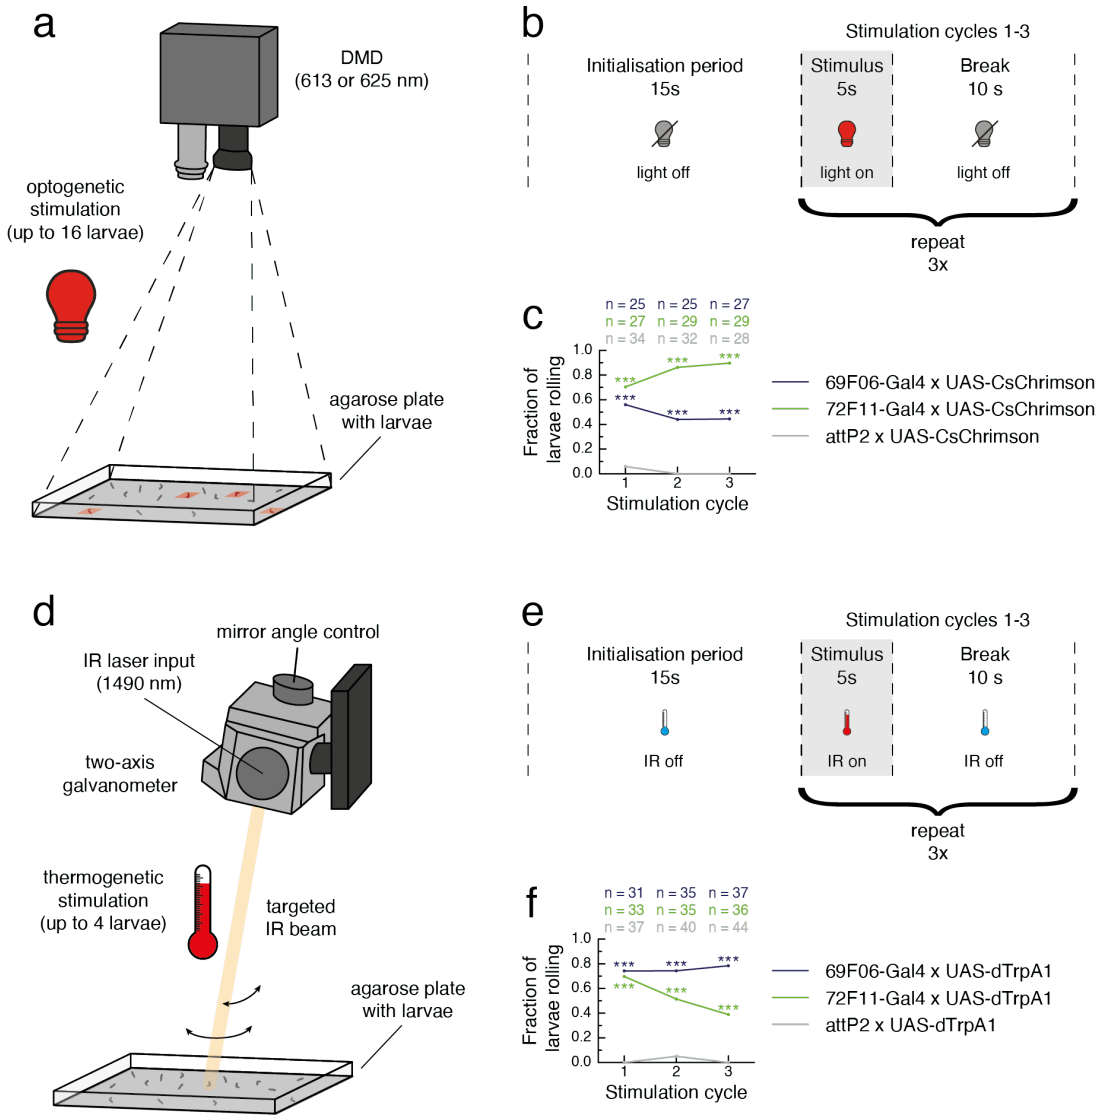

Fig. 6.2 Optogenetic and thermogenetic stimulation with the high-throughput tracker. **a.** Hardware design schematic for optogenetic stimulation. Although the high-throughput tracker included two digital micromirror devices (DMDs), only one is shown for simplicity. **b.** Proof-of-principal experiment protocol for optogenetic stimulation. **c.** The fraction of larvae for which a `roll` was detected in each stimulation cycle. *69F06-Gal4 x UAS-CsChrimson* and *72F11-Gal4 x UAS-CsChrimson* larvae (*CsChrimson* expressed in neurons triggering roll behaviour; experiment groups) were compared to *attP2 x UAS-CsChrimson* larvae (no *CsChrimson* expression; control group). Fisher's exact test was used to calculate statistical differences between the experiment and control groups (*** $p < 0.001$). **d.** Hardware design schematic for thermogenetic stimulation. Although the high-throughput tracker included four two-axis galvanometers, only one is shown for simplicity. IR: infrared. **e.** Proof-of-principal experiment protocol for thermogenetic stimulation. **f.** The fraction of larvae for which a `roll` was detected in each stimulation cycle. *69F06-Gal4 x UAS-dTrpA1* and *72F11-Gal4 x UAS-dTrpA1* larvae (*dTrpA1* expressed in neurons triggering roll behaviour; experiment groups) were compared to *attP2 x UAS-dTrpA1* larvae (no *dTrpA1* expression; control group). Fisher's exact test was used to calculate statistical differences between the experiment and control groups (*** $p < 0.001$).

72F11-Gal4 drives expression in the Basin neurons, which integrate mechanosensory and nociceptive stimuli. Klapoetke et al. (2014) have developed the red-shifted channelrhodopsin *CsChrimson*, which can be expressed under GAL4 control. We tested whether *69F06-Gal4 x UAS-CsChrimson* and *72F11-Gal4 x UAS-CsChrimson* larvae rolled upon exposure to red light (Fig. 6.2b, see also Chapter 2). In each stimulation cycle, we observed above-threshold rolls in over 40% of *69F06-Gal4 x UAS-CsChrimson* larvae and over 70% of *72F11-Gal4 x UAS-CsChrimson* larvae. This behaviour significantly contrasted with that of *attP2 x UAS-CsChrimson* control larvae (Fig. 6.2c), suggesting that the DMDs could be used for optogenetic stimulation without activating the animals' photoreceptors.

We also verified the efficacy of the galvanometer set-up for thermogenetic stimulation (Fig. 6.2d). We tested whether *69F06-Gal4 x UAS-dTrpA1* and *72F11-Gal4 x UAS-dTrpA1* larvae rolled upon exposure to the IR laser (Fig. 6.2e, see also Chapter 2). In each stimulation cycle, we observed above-threshold rolls in over 70% of *69F06-Gal4 x UAS-dTrpA1* larvae and over 35% of *72F11-Gal4 x UAS-dTrpA1* larvae; a significant contrast to the *attP2 x UAS-dTrpA1* control larvae whose roll rate was close to zero. We concluded that these heating conditions were effective for targeted *Trp* channel activation without larvae perceiving strong pain (Fig. 6.2f).

6.3.2 Operant conditioning of larval cast direction

We chose optogenetic activation of reward circuits as a US for automated operant conditioning. The main challenge was determining which neurons could convey a sufficient reinforcement signal, especially as the capacity for *Drosophila* larvae to exhibit operant learning was not yet demonstrated. Across the animal kingdom, it has been observed that biogenic amine neurotransmitters can provide such a signal (Fee and Goldberg, 2011; Giurfa, 2006; Hawkins and Byrne, 2015; Meneses and Liy-Salmeron, 2012). It is also conceivable that the *Drosophila* PAM cluster dopaminergic neurons that can signal reward in classical conditioning (Cognigni et al., 2018; Liu et al., 2012; Rohwedder et al., 2016; Vogt et al., 2014; Waddell, 2013) may perform similarly in operant conditioning. We therefore aimed to induce operant conditioning by stimulating a broad set of dopaminergic and serotonergic neurons. If valence signalling relevant for operant conditioning is mediated by one of these two neurotransmitters, activation of this large set of neurons paired with behaviour should be sufficient to induce learning.

We expressed *UAS-CsChrimson* under the control of the *Ddc-Gal4* driver, which covers a large set of dopaminergic and serotonergic neurons in the CNS (Li et al., 2000; Lundell and Hirsh, 1994; Sitaraman et al., 2008), including the PAM cluster (Aso et al., 2012; Liu et al., 2012). Although the function of most *Ddc* neurons is unknown, their collective activation can substitute for an olfactory conditioning reward in adult flies (Aso et al., 2012; Liu et al., 2012; Shyu et al., 2017). The goal of our paradigm was to establish a learned direction preference for casting, conditioning *Ddc-Gal4 x UAS-CsChrimson* larvae to cast more often to one side than the other. Although stimulation side was randomized across trials, we describe (for simplicity) the experiment procedure where this predefined side was the left. Each experiment began with a one-minute test period where no light was presented. What followed were four training sessions, each three-minutes long, in which larvae received optogenetic stimulation when casting to the left. Between training sessions, larvae experienced three minutes without stimulation. Larvae were periodically brushed back to the centre of the agarose plate to mitigate the experimental side effects of reaching the plate's edge (see Chapter 2 for more details). Following the fourth training session was a one-minute test period without stimulation (Fig. 6.3a).

For each larva, cast rate, measured as the number of casts per minute performed towards a given side, served as a read-out for cast direction preference. Within a given one-minute time bin, the statistical test comparing cast rates to each other is mathematically equivalent to the difference in cast rates compared to 0. However, computing the difference in cast rate facilitated comparison across time and between genotypes that themselves may differ in basal cast rate. This difference was calculated as the number of casts per minute to the stimulated side minus the number of casts per minute to the unstimulated side. In the one-minute test prior to the first training session, we observed no significant difference in *Ddc-Gal4 x UAS-CsChrimson* larval cast rate to either side (Fig. 6.3b, Fig. 6.3c). Close investigation of larval cast rate during training revealed a significantly greater number of casts to the stimulated side throughout each of the four training sessions, with the exception of the final minute of session three. Notably, the observed difference in cast rate between stimulated and unstimulated sides tended to increase over the course of each training session (Fig. 6.4a). Together, these findings suggested that activation of *Ddc* neurons can serve as a rewarding stimulus that larvae increasingly seek with time. A direct comparison of individual cast rates from before training to after training showed a significant decrease in casts to both the stimulated and unstimulated sides (Fig. 6.3b). Because the experiment protocol exceeded 17

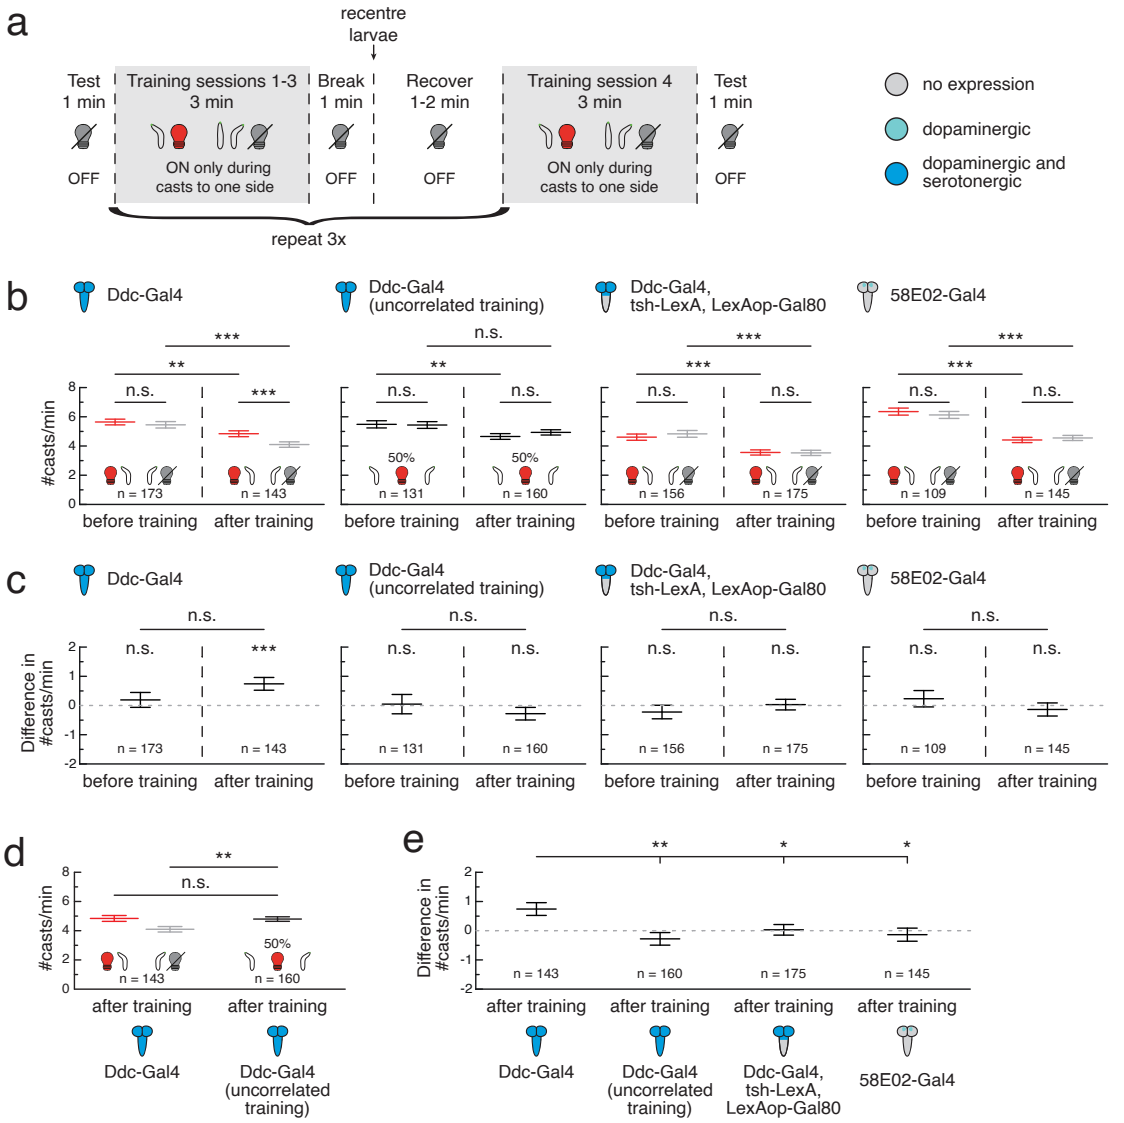

Fig. 6.3 Operant conditioning of cast direction in *Drosophila* larvae requires the ventral nerve cord. a. Experiment protocol using the high-throughput closed-loop tracker. Behaviours are depicted as larval contours (black) with head (green). During training, the larva received an optogenetic stimulus (red light bulb) whenever it cast to one predefined side (here depicted as the left for simplicity), and light was switched off during all other behaviours (grey light bulb). **b–e.** Gal4 expression depicted as color-coded CNS. The effector for all genetic lines, *UAS-CsChrimson*, is omitted from the figure for visual clarity. All data is shown as (mean $\pm$ s. e. m.). **b, d.** Larval cast rate shown as the number of casts per minute, grouped by cast direction. The cast rate to the stimulated side (depicted as a left cast with a red light bulb for simplicity) is shown in red and the cast rate to the unstimulated side (depicted as a right cast with a grey light bulb for simplicity) is shown in grey. For larvae that received random, uncorrelated stimulation during 50% of casts, the cast rates to the left and right are shown in black. Data in **b** is shown from the test period before the first training session and the test period after the fourth training session. Data in **d** is the same after training data as in **b**, but cast rate for uncorrelated training group was calculated without stratification by cast direction. Statistical differences within groups (i. e. within before training time bin or within after training time bin) were tested with a two-sided Wilcoxon signed-rank test. Statistical differences between two groups (e. g. stimulated side before training versus stimulated side after training or stimulated side after training between genotype conditions) were tested with a two-sided Mann-Whitney *U* test. n. s. $p \geq 0.05$ (not significant), * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. **c, e.** Difference in cast rate between stimulated and unstimulated sides. Grey dashed line at 0 indicates equal cast rates to either side. **c.** Within-group statistics calculated from a two-sided Wilcoxon signed-rank test. Between-group statistics calculated from a two-sided Mann-Whitney *U* test. n. s. $p \geq 0.05$ (not significant), *** $p < 0.001$. **e.** Same after training data as in **c**, with statistical comparisons to *Ddc-Gal4 x UAS-CsChrimson* calculated using a two-sided Mann-Whitney *U* test with Bonferroni correction. * $p < 0.05/3$, ** $p < 0.01/3$.

minutes (including recovery periods following recentring larvae), this may reflect general larval fatigue. In the one-minute test following the fourth training session, larvae showed a preference for casts towards the side paired with red light stimulation during training (Fig. 6.3b, Fig. 6.3c), though directly comparing this difference to that observed before training did not reveal a statistically significant increase (Fig. 6.3c).

To confirm that the cast preference we observed in the after training period was attributable to pairing light with casts solely in one direction, we conducted another control experiment in which larvae received random, uncorrelated stimulation during 50% of casts regardless of direction. Before, during, and after training, larvae showed no difference in absolute left and right cast rates (with the exception of the first minute of the fourth training session, see Fig. 6.3b, Fig. 6.3c, and Fig. 6.4c). These larvae did, however, show a significantly lower cast rate to the left side after training compared to before; a trend that appears to reflect larval fatigue as also observed with the experimental larvae (Fig. 6.3b). These larvae also showed a significantly lower difference between cast rates after training compared to pair-trained larvae (Fig. 6.3e). Further dissection of cast rates to each side showed that the after training cast rates averaged together for larvae that received uncorrelated training were indistinguishable from the rate of pair-trained larvae casting to the previously stimulated side. However, larvae that received uncorrelated training showed a significantly higher cast rate overall compared to pair-trained larvae casting to the previously unstimulated side (Fig. 6.3d). This raised the question whether pair-trained *Ddc-Gal4 x UAS-CsChrimson* larvae were learning to prefer the side paired with the rewarding US, or rather to avoid the side without the stimulus.

A bootstrap sensitivity analysis was also performed on these data to assess the robustness of the observed statistical inference results. There may exist systematic, non-randomly distributed measurement error not modeled within the existing experiment design that could improperly suggest a statistically significant difference between experimental conditions of interest. A primary benefit of this approach is the ability to model such measurement error without needing to perform further laboratory experiments. Table 6.1 displays the bootstrap results for all statistical comparisons within Fig. 6.3 (for details on bootstrap methodology, see Section 2.2.1). In short, we asked what proportion of bootstrapped data sets reproduce the hypothesised relationship between the original, non-bootstrapped groups? Consider taking two groups of comparative interest and introducing principled variability in their measurement distributions by sampling observations (with replacement) from each group. Repeating

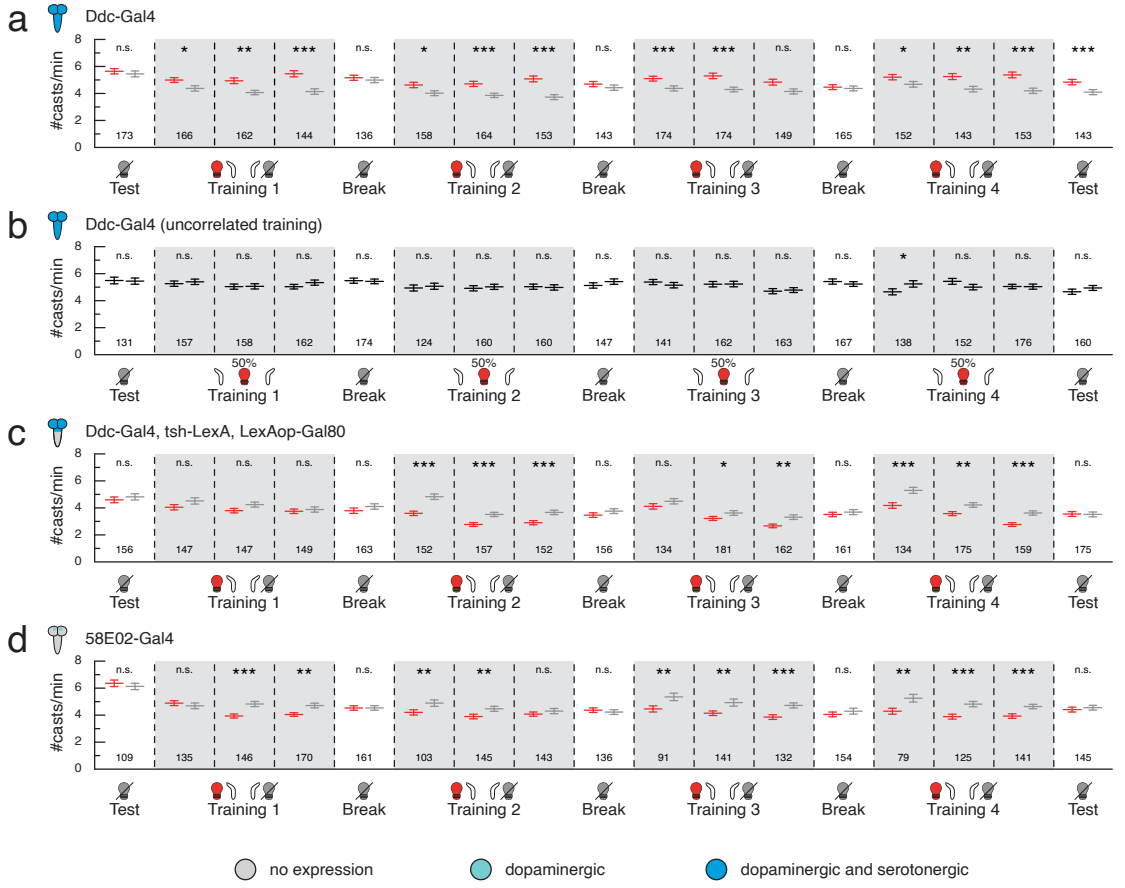

Fig. 6.4 *Drosophila* larvae show cast direction preference during operant paradigm training sessions. Gal4 expression depicted as color-coded CNS. All fly lines contained the *UAS-CsChrimson* effector, which is omitted from the figure for visual clarity. Experiments followed the protocol depicted in Fig. 6.3a. Data is shown in one-minute time bins (separated by vertical dashed lines) across the entire experiment, beginning with the test period before the first training session and ending with the test period after the fourth training session. Larval cast rate is shown as the number of casts per minute, grouped by cast direction. The cast rate to the stimulated side (depicted as a left cast with a red light bulb for simplicity) is shown in red and the cast rate to the unstimulated side (depicted as a right cast with a grey light bulb for simplicity) is shown in grey. For larvae that received random, uncorrelated stimulation during 50% of casts (panel **c**), the cast rates to the left and right are shown in black. All data is shown as (mean $\pm$ s. e. m.). Number of larvae in each time bin is written just above the x-axis. Statistical differences within bins were tested with a two-sided Wilcoxon signed-rank test; n. s. $p \geq 0.05$ (not significant), * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

this process 1000 times for both groups yields 1000x1000 "experiments", representing all possible combinations of bootstrapped observations across the comparison.

After calculating the mean of each group within these 1 million experiments, our goal was to calculate the proportion of these bootstrapped data sets that satisfy a specific threshold criterion on the relationship between means. This criterion matched the original biological comparison. This proportion was then interpreted in the context of whether it supports the original experimental finding. In other words, for each experimental comparison, we defined an arbitrary threshold for the bootstrap proportion to assess its consistency with the observed result. This required assessing whether the direction (above or below 0.5) and magnitude (greater than 0.95 or less than 0.05) of a bootstrap proportion was consistent with the effect direction and significance of the original statistical test.

Bootstrap analysis revealed that the original statistical testing framework was broadly robust for the comparisons described thus far. Bootstrap proportions were consistent with effect direction and significance in all *Ddc-Gal4 x UAS-CsChrimson* and *Ddc-Gal4 x UAS-CsChrimson* (uncorrelated training) tests with the exception of comparing *Ddc-Gal4 x UAS-CsChrimson* (uncorrelated training) cast rates to the right side after training to those on the right side before training (Table 6.1). Bootstrap analysis of this comparison indicates that a significant proportion of the time ($probability = 0.9593$), the cast rate to the right side is lower after training than before. This is inconsistent with the inference conducted on the original, non-bootstrapped data set that showed no significant difference in cast rates to the right side from before training to after (Fig. 6.3b). In other words, the original statistical inference is not robust to measurement error and perhaps less weight should be given to the resulting conclusion. This inconsistency is informative because it lends further biological consideration to the observed significant decreases in cast rates on both sides following training for pair-trained larvae (Fig 6.3b).

It was also explored whether relaxing object inclusion criteria impacted the observed behavioural outcomes in the conservative data set. The data set shown in Fig. 6.5 included all larvae even if they lost their object identity and were therefore only detected for part of the experiment. This can be viewed as an alternative method to assess the robustness of an observed statistical inference result following the introduction of additional measurement error. For larvae that were tracked for less than the full one-minute duration of a given time bin, the number of casts that they performed on both sides in the period they were tracked was scaled up to reflect cast rate

Table 6.1 Bootstrap sensitivity analyses of data in Fig. 6.3. Each statistical inference is denoted by genotype and Comparison. The associated Bootstrap Test is defined by a threshold criterion on the relationship between means. The Bootstrap Proportion reflects the proportion of all bootstrapped data sets that satisfied the Bootstrap Test. Consistency with Observed P-value indicates whether the direction and magnitude of the Bootstrap Proportion supports the original experimental finding. Inconsistencies are highlighted in grey.

	Comparison	Bootstrap Test	Bootstrap Proportion	Consistency with Observed P-value
Fig. 6.3b – #casts/min				
Ddc-Gal4	before training	p(stimulated > unstimulated)	0.7469	consistent
	after training	p(stimulated > unstimulated)	0.9966	consistent
	stimulated side	p(after < before)	0.9986	consistent
	unstimulated side	p(after < before)	1.000	consistent
Ddc-Gal4 (uncorrelated training)	before training	p(left > right)	0.5298	consistent
	after training	p(left > right)	0.1402	consistent
	left side	p(after < before)	0.9956	consistent
	right side	p(after < before)	0.9593	inconsistent
Ddc-Gal4, tsh-LexA, LexAop-Gal80	before training	p(stimulated > unstimulated)	0.2227	consistent
	after training	p(stimulated > unstimulated)	0.5379	consistent
	stimulated side	p(after < before)	0.9997	consistent
	unstimulated side	p(after < before)	1.000	consistent
58E02-Gal4	before training	p(stimulated > unstimulated)	0.7500	consistent
	after training	p(stimulated > unstimulated)	0.2769	consistent
	stimulated side	p(after < before)	1.000	consistent
	unstimulated side	p(after < before)	1.000	consistent
Fig. 6.3c – Difference in #casts/min				
Ddc-Gal4	before training	p(before > 0)	0.7710	consistent
	after training	p(after > 0)	1.000	consistent
	after *vs* before training	p(after > before)	0.9447	consistent
Ddc-Gal4 (uncorrelated training)	before training	p(before > 0)	0.5310	consistent
	after training	p(after > 0)	0.0930	consistent
	after *vs* before training	p(after > before)	0.2114	consistent
Ddc-Gal4, tsh-LexA, LexAop-Gal80	before training	p(before > 0)	0.1600	consistent
	after training	p(after > 0)	0.5570	consistent
	after *vs* before training	p(after > before)	0.8151	consistent
58E02-Gal4	before training	p(before > 0)	0.7800	consistent
	after training	p(after > 0)	0.2350	consistent
	after *vs* before training	p(after > before)	0.1528	consistent
Fig. 6.3d – #casts/min after training				
Ddc-Gal4 *vs* Ddc-Gal4 (uncorrelated training) stimulated side		p(Ddc > Ddc uncorrelated)	0.5576	consistent
Ddc-Gal4 *vs* Ddc-Gal4 (uncorrelated training) unstimulated side		p(Ddc < Ddc uncorrelated)	0.0021	consistent
Fig. 6.3e – Difference in #casts/min after training				
Ddc-Gal4 *vs* Ddc-Gal4 (uncorrelated training)		p(Ddc > Ddc uncorrelated)	0.9997	consistent
Ddc-Gal4 *vs* Ddc-Gal4, tsh-LexA, LexAop-Gal80		p(Ddc > Ddc tsh)	0.9939	consistent
Ddc-Gal4 *vs* 58E02-Gal4		p(Ddc > 58E02)	0.9763	consistent

per-minute. Notably, this scaling assumed that larval cast rate does not change with time. All statistical results show the same direction of effect, but noticeably the operant learning effect previously observed after pair-trained *Ddc* activation is less pronounced (Fig. 6.5a, Fig. 6.5b). Furthermore, all statistically significant relationships previously observed between *Ddc-Gal4 x UAS-CsChrimson* larvae and the controls or *58E02-Gal4 x UAS-CsChrimson* larvae are abolished (Fig. 6.5d). While absence of the observed effect does not serve as direct evidence that larvae are not capable of operant learning following *Ddc* neuron activation, these results suggest that the statistical inference and its underlying assumptions may not be robust when additional noise is added to the measurement. Further analysis of the behavioural data and perhaps modeling of

underlying larval dynamics and biases may elucidate how best to interpret results following inclusion of all data in this way. An important consideration is that, by including larvae that have not been tracked for the full duration of a given one-minute time bin, there is a risk that larvae may be detected more than once within the bin and therefore contribute multiple cast rate values to the data set. There is currently no framework to determine whether a given larva was previously detected and therefore larva IDs cannot be associated across time with individual larvae. Future software development could aim to fill this knowledge gap.

6.3.3 The mushroom body is not sufficient to mediate operant conditioning in larvae

Our experiments showed that activation of *Ddc* neurons is a sufficient US for operant conditioning. While we did not identify which individual neurons mediate the observed effect, we hypothesised that not all *Ddc* neurons are involved. Some prior work in adult flies suggests that the MB is involved in operant conditioning (Sun et al., 2020), while other studies in the adult suggest that operant conditioning does not require the MB (Booker and Quinn, 1981; Colomb and Brembs, 2010, 2016; Wolf et al., 1998) and may instead involve motor neuron plasticity (Colomb and Brembs, 2016). The extent to which the MB is dispensable in larval operant conditioning is unknown. We investigated whether smaller subsets of *Ddc* neurons in the brain and SEZ could support memory formation in our cast direction paradigm.

GAL80 under control of the *tsh* promoter suppresses GAL4 expression in the VNC, but not in the brain or SEZ (Clyne and Miesenböck, 2008; Fig. 6.6). Prior to training under our operant conditioning protocol (Fig. 6.3a), *Ddc-Gal4 x UAS-CsChrimson; tsh-LexA, LexAop-Gal80* larvae showed no directional bias in cast rate (Fig. 6.3b, Fig. 6.3c). Observations of larval cast rates during training revealed a persistent, and in some cases statistically significant, direction preference to the unstimulated side, suggesting perhaps a mild aversion to optogenetic stimulation in the paradigm (Fig. 6.4d). Following training, these larvae were equally likely to cast towards the side where they had previously received the optogenetic stimulus as they were to cast towards the unstimulated side (Fig. 6.3b, Fig. 6.3c). As we observed in all genotype conditions described thus far, these larvae showed significantly lower cast rates to both sides after training compared to before training (Fig. 6.3b). Together, these results suggested that activating these neurons was a seemingly insufficient rewarding US in

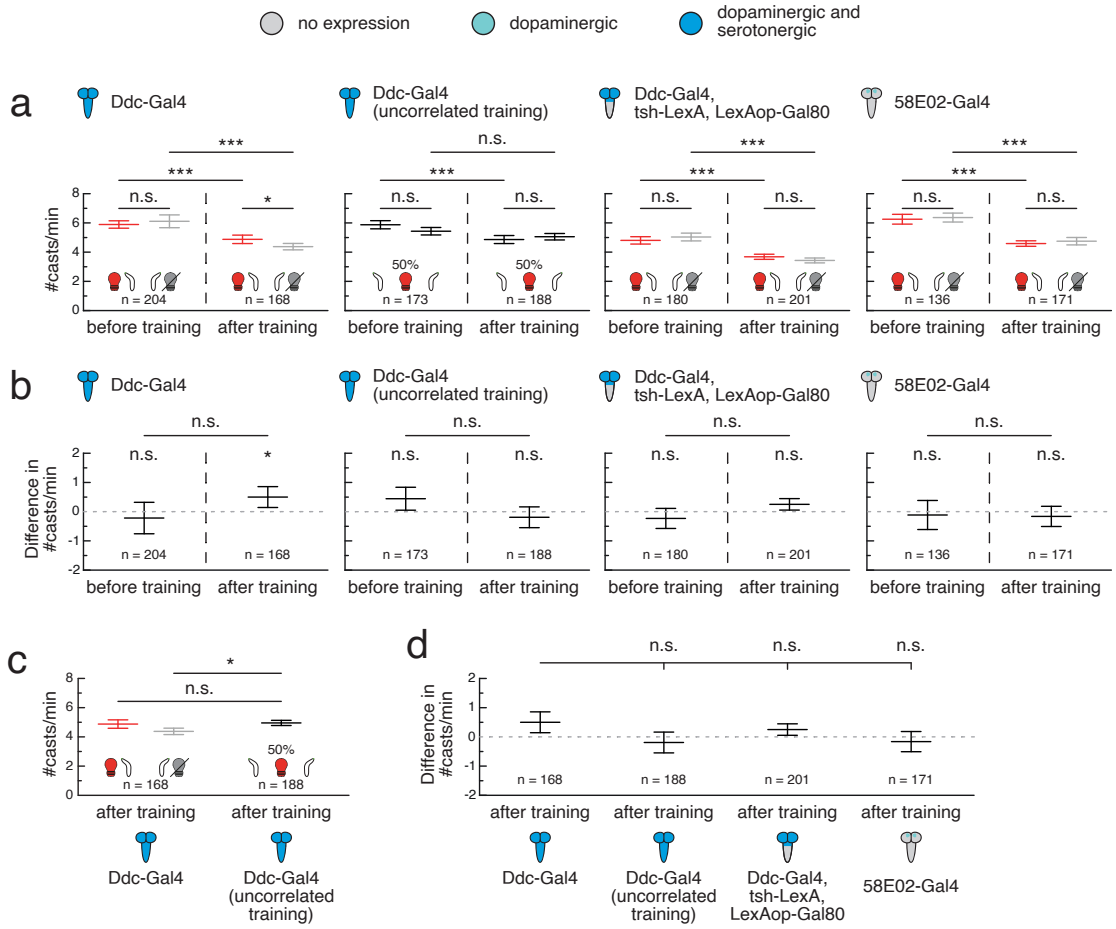

Fig. 6.5 *Ddc* and *58E02* data sets including all larvae regardless of tracking. Experiments were conducted following the protocol outlined in Fig. 6.3a. Gal4 expression depicted as color-coded CNS. The effector for all genetic lines, *UAS-CsChrimson*, is omitted from the figure for visual clarity. Criteria to retain objects for analysis is identical to that described in Section 2.2.1 with the exception that larvae detected for only part of each bin were also included. All data is shown as (mean $\pm$ s. e. m.). **a, c.** Larval cast rate shown as the number of casts per minute, grouped by cast direction. The cast rate to the stimulated side (depicted as a left cast with a red light bulb for simplicity) is shown in red and the cast rate to the unstimulated side (depicted as a right cast with a grey light bulb for simplicity) is shown in grey. For larvae that received random, uncorrelated stimulation during 50% of casts, the cast rates to the left and right are shown in black. Data in **a** is shown from the test period before the first training session and the test period after the fourth training session. Data in **c** is the same after training data as in **a**, but cast rate for uncorrelated training group was calculated without stratification by cast direction. Statistical differences within groups (i. e. within before training time bin or within after training time bin) were tested with a two-sided Wilcoxon signed-rank test. Statistical differences between two groups (e. g. stimulated side before training versus stimulated side after training or stimulated side after training between genotype conditions) were tested with a two-sided Mann-Whitney *U* test. n. s. $p \geq 0.05$ (not significant), * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. **b, d.** Difference in cast rate between stimulated and unstimulated sides. Grey dashed line at 0 indicates equal cast rates to either side. **b.** Within-group statistics calculated from a two-sided Wilcoxon signed-rank test. Between-group statistics calculated from a two-sided Mann-Whitney *U* test. n. s. $p \geq 0.05$ (not significant), *** $p < 0.001$. **d.** Same after training data as in **b**, with statistical comparisons to *Ddc-Gal4 x UAS-CsChrimson* calculated using a two-sided Mann-Whitney *U* test with Bonferroni correction. n. s. $p \geq 0.05/3$ (not significant).

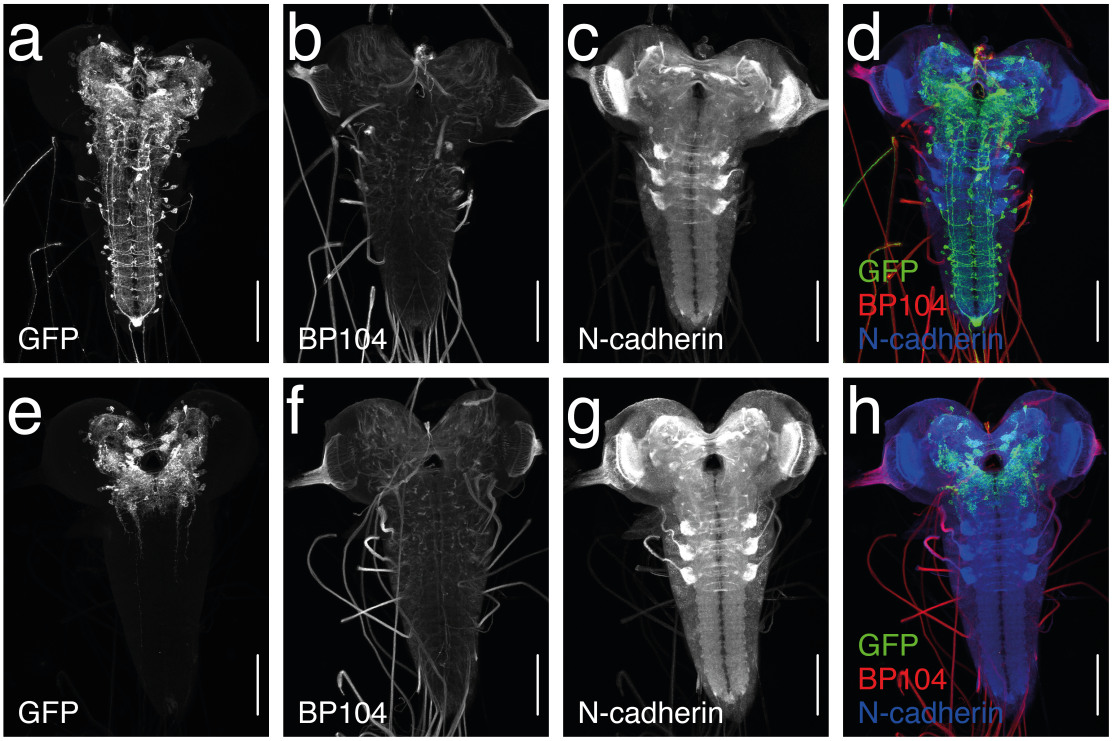

Fig. 6.6 *Ddc-Gal4* expression pattern without and with *tsh-Gal80* restriction. Maximum intensity projections of confocal images obtained after immunohistochemical staining. **a**, **e**; green in **d** and **h.** Targeting a green fluorescent protein (GFP) antibody to the *mVenus* tag of *CsChrimson*. **b**, **f**; red in **d** and **h**. Staining against BP104. **c**, **g**; blue in **d** and **h.** Staining against N-cadherin. **a–d.** *Ddc-Gal4 x UAS-CsChrimson* larvae. Manually counting the cell bodies in the image stacks revealed more than 200 GFP-positive neurons located in the brain, subesophageal zone (SEZ), and ventral nerve cord (VNC), including the PAM cluster dopaminergic neurons innervating the mushroom body ($n = 2$). This confirmed that *Ddc-Gal4* drives broad expression across the central nervous system (CNS) (Li et al., 2000; Lundell and Hirsh, 1994). **e–h.** *Ddc-Gal4 x UAS-CsChrimson; tsh-LexA, LexAop-Gal80* larvae. As expected, no GFP-positive neurons were found in the VNC ($n = 6$). *Ddc-Gal4* brain and SEZ expression remained largely unaffected by GAL80, as the GFP-positive neurons in both areas that could be consistently identified in *Ddc-Gal4 x UAS-CsChrimson* larvae ($n = 3$) were also present in *Ddc-Gal4 x UAS-CsChrimson; tsh-LexA, LexAop-Gal80* larvae ($n = 3$). **a–h.** Plan-Apochromat 20x objective, resolution: 592 x 800 pixels, scale bar: 100 µm. Images courtesy of the HHMI Janelia FlyLight team.

this paradigm. A direct comparison of the after training difference in cast rate between these larvae and the *Ddc-Gal4 x UAS-CsChrimson* group showed a significantly lower difference in cast rate for *Ddc-Gal4 x UAS-CsChrimson; tsh-LexA, LexAop-Gal80* larvae (Fig. 6.3e). The loss of the operant conditioning effect we otherwise observed in after training cast rates of *Ddc-Gal4 x UAS-CsChrimson* larvae highlighted that dopaminergic or serotonergic neurons in the VNC are necessary for the formation of a learned cast direction preference. Their sufficiency was inconclusive, however, since perhaps two or more distinct groups of *Ddc* neurons needed collective activation in order to form a memory. Bootstrap proportions were consistent with each of these results (Table 6.1).

We then assessed whether exclusively activating the PAM cluster dopaminergic neurons innervating the MB could induce operant conditioning, as is the case for classical conditioning. *58E02-Gal4* drives expression in the majority of these neurons (Rohwedder et al., 2016). In the test period before training, *58E02-Gal4 x UAS-CsChrimson* larvae did not exhibit a cast direction preference (Fig. 6.3b, Fig. 6.3c). During training, however, these larvae showed a significant cast direction preference to the unstimulated side (Fig. 6.4d). These preference results contrast the appetitive behaviour observed when activating *58E02* neurons in classical conditioning. They may be partly a consequence of these neurons' role in motor control, as *58E02* neurons are known to form synaptic connections to aversive MBONs that promote turning (Eichler et al., 2017). Following training, *58E02-Gal4 x UAS-CsChrimson* larvae did not exhibit a learned direction preference for casts to either side (Fig. 6.3b, Fig. 6.3c), though they did show a significant decrease in cast rates to both sides after training compared to before (Fig. 6.3b). These larvae also exhibited a significantly lower difference in cast rate between sides after training compared to *Ddc-Gal4 x UAS-CsChrimson* larvae (Fig. 6.3e). Bootstrap proportions were consistent with each of these results (Table 6.1). It is unsurprising that activation of *58E02* neurons alone could not act as a rewarding US in this paradigm, given our finding that *Ddc* neurons in the brain and SEZ are insufficient. It is remarkable, however, because it suggests that the neural circuits signalling reward in operant conditioning differ from those of classical conditioning. Although it remains to be seen whether these PAM cluster neurons contribute to memory formation by interacting with other *Ddc* neurons, these results further supported the idea that operant conditioning in *Drosophila* may not be mediated by the MB.

6.3.4 Serotonergic neurons in brain and SEZ are a sufficient reward signal in classical conditioning

Pairing an action with activation of numerous dopaminergic and serotonergic neurons across the CNS was sufficient to induce operant conditioning of cast direction preference. Furthermore, our results indicated that the VNC subset of these neurons was essential for memory formation in the paradigm. It was an open question, however, whether this learning was mediated by dopamine, serotonin, or both. Dopamine and serotonin receptors are necessary for different classical conditioning tasks in honeybees, suggesting that the two neurotransmitters may carry out separate functions (Wright et al., 2010). We conducted a high-throughput classical conditioning screen of sparser dopaminergic and serotonergic driver lines to identify US candidates for comparison with our operant conditioning paradigm.

We expressed *CsChrimson* under the control of different GAL4 driver lines and tested whether pairing optogenetic activation of these neurons (US) with odour presentation (CS) could induce olfactory memory. Conditioning was performed using a similar procedure to those described in Gerber and Hendel (2006), Saumweber et al. (2011) and Eschbach et al. (2020b). In the paired group, larvae were exposed to alternating three-minute presentations of ethyl acetate with red light and air with no light. To ensure that any observed effects were a result of learning rather than innate odour preference or avoidance, an unpaired group was trained simultaneously with reciprocal stimulus presentation (odour/dark, air/light). Following training, larvae in both groups were tested on their preference for the odour in the absence of light (Fig. 6.7a). All learning scores were compared to a negative control containing no GAL4 driver, *w^{1118} x UAS-CsChrimson*, which did not exhibit a learning phenotype (Fig. 6.7b). Consistent with prior study results (Almeida-Carvalho et al., 2017; Eichler et al., 2017; Rohwedder et al., 2016), *58E02-Gal4 x UAS-CsChrimson* larvae showed appetitive olfactory learning with a significantly higher performance index than *w^{1118} x UAS-CsChrimson* larvae and so were used as a positive control (Fig. 6.7b). *Ddc-Gal4 x UAS-CsChrimson* larvae exhibited appetitive memory comparable to *58E02-Gal4* ($p = 0.1304$, two-sided Mann-Whitney *U* test); an unsurprising result since the *Ddc-Gal4* expression pattern includes the PAM cluster neurons. Consistent with previous studies in the larva (Schroll et al., 2006) and adult (Aso et al., 2012; Claridge-Chang et al., 2009; Liu et al., 2012), *TH-Gal4 x UAS-CsChrimson* larvae exhibited significant aversive olfactory learning. *TH-Gal4* covers most dopaminergic neurons, excluding the PAM cluster (Rohwedder

et al., 2016). The effect we observed may be mediated by punishment-signalling dopaminergic neurons that project to the MB vertical lobes (Eschbach et al., 2020b; Selcho et al., 2009). Isolating the locus of this effect may prove challenging, given the dearth of larval driver lines targeting dopaminergic neurons without MB innervation.

Serotonergic signalling is required for associative learning in both larval (Huser et al., 2017) and adult (Johnson et al., 2011; Sitaraman et al., 2012) *Drosophila*. We tested *Trh-Gal4* and *Tph-Gal4*, two driver lines that target the majority of serotonergic neurons and no dopaminergic neurons across the CNS of third-instar larvae (Huser et al., 2012). Consistent with previous reports (Ganguly et al., 2020), larvae expressing *CsChrimson* under either driver line formed strong appetitive olfactory memory, highlighting the sufficiency of serotonin as a US in associative learning. *Tph-Gal4* targets fewer serotonergic neurons than *Trh-Gal4*, making it valuable for narrowing down which serotonergic neurons serve as a relevant reward signal. We eliminated all *Tph-Gal4* expression in the VNC using *tsh-Gal80* (Fig. 6.8). Activating the remaining *Tph* neurons in the brain and SEZ was sufficient to induce strong appetitive memory (Fig. 6.7b). This result was notable and raised further questions: are serotonergic neurons in the brain and SEZ indirectly connected to MB-innervating dopaminergic neurons or do alternative learning circuits exist that altogether bypass the MB?

The contralaterally projecting serotonin-immunoreactive deutocerebral (CSD) neuron (Roy et al., 2007) is one previously described serotonergic brain neuron within the *Tph-Gal4* expression pattern (Huser et al., 2012) that innervates the antennal lobe and only has a few indirect connections to the MB (Berck et al., 2016). Combining anatomical features from existing EM reconstruction (Berck et al., 2016) with available lineage information facilitated identification of a split-GAL4 line (*SS01989*) that drives expression exclusively in the CSD neuron (Fig. 6.9). Pairing activation of *SS01989* with ethyl acetate was insufficient for inducing olfactory memory (Fig. 6.7b), suggesting that the classical conditioning phenotype we observed under *Tph-Gal4 x UAS-CsChrimson; tsh-LexA, LexAop-Gal80* was mediated by at least one other group of serotonergic neurons in the brain or SEZ.

6.3.5 Serotonergic neurons in VNC may play a role in operant conditioning of cast direction

Given their strong associative learning phenotypes, we used the *TH-Gal4* and *Tph-Gal4* drivers to investigate whether operant conditioning of cast direction could be induced

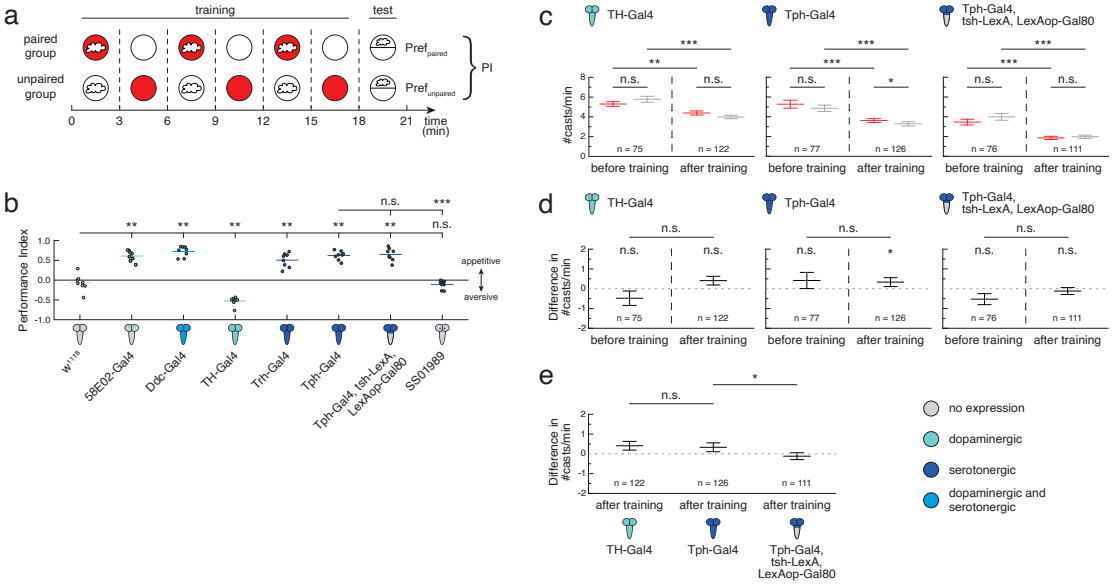

Fig. 6.7 Different serotonergic neurons may mediate classical and operant conditioning. All fly lines contained the *UAS-CsChrimson* effector, which is omitted from the figure for visual clarity. Gal4 expression depicted as color-coded CNS. **a.** Olfactory conditioning experiment protocol. During training, larvae in the paired group received three minutes of optogenetic red light stimulation (solid red circles) paired with the odour (white cloud) followed by three minutes of darkness (solid white circles) paired with air (no cloud). The unpaired group received reciprocal stimulus presentation (dark paired with odour, light paired with air). This procedure was repeated three times. In half of the experiments, the order of training trials was reversed, starting with air presentation instead of odour presentation. Both groups were then tested for learned odour preference in the dark with odour presented on one side of the plate and no odour on the other (PI = performance index). **b.** Performance indices following olfactory conditioning, plotted as raw data points and mean. *w^{1118} x UAS-CsChrimson* was the negative control (grey, $n = 8$), *58E02-Gal4 x UAS-CsChrimson* was the positive control (blue, $n = 8$). Statistical comparisons to *w^{1118} x UAS-CsChrimson* were calculated using a two-sided Mann-Whitney *U* test with Bonferroni correction; n. s. $p \geq 0.05/7$ (not significant), ** $p < 0.01/7$. Statistical comparisons to *Tph-Gal4 x UAS-CsChrimson* were calculated using a two-sided Mann-Whitney *U* test with Bonferroni correction; n. s. $p \geq 0.05/2$ (not significant), *** $p < 0.001/2$. **c–e.** All data is shown as (mean $\pm$ s. e. m.). Experiments followed the protocol depicted in Fig. 6.3a. **c.** Larval cast rate shown as the number of casts per minute, grouped by cast direction. The cast rate to the stimulated side is shown in red and the cast rate to the unstimulated side is shown in grey. Data is shown from the test period before the first training session and the test period immediately following the fourth training session. Statistical differences within groups (i. e. within before training time bin or within after training time bin) were tested with a two-sided Wilcoxon signed-rank test. Statistical differences between two groups (e. g. stimulated side before training versus stimulated side after training) were tested with a two-sided Mann-Whitney *U* test. n. s. $p \geq 0.05$ (not significant), * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. **d.** Difference in cast rate between stimulated and unstimulated sides. Grey dashed line at 0 indicates equal cast rates to either side. Data is shown from the test period before the first training session and the test period immediately following the fourth training session. Within-group statistics calculated from a two-sided Wilcoxon signed-rank test. Between-group statistics calculated from a two-sided Mann-Whitney *U* test. n. s. $p \geq 0.05$ (not significant), * $p < 0.05$. **e.** Same after training data as in **d**, with statistical comparisons to *Tph-Gal4 x UAS-CsChrimson* calculated using a two-sided Mann-Whitney *U* test with Bonferroni correction. n. s. $p \geq 0.05/2$ (not significant), * $p < 0.05/2$.

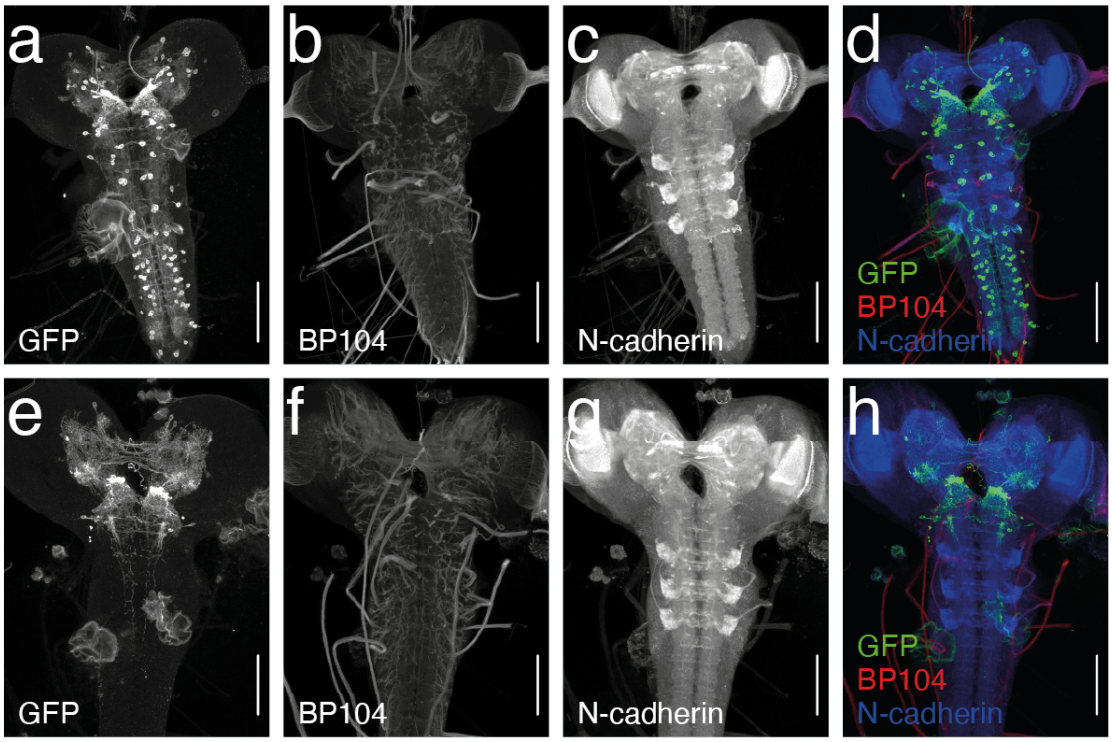

Fig. 6.8 *Tph-Gal4* expression pattern without and with *tsh-Gal80* restriction. Maximum intensity projections of confocal images obtained after immunohistochemistry. **a–d.** *Tph-Gal4 x UAS-CsChrimson* larvae, **e–h.** *Tph-Gal4 x UAS-CsChrimson; tsh-LexA, LexAop-Gal80* larvae. **a**, **e**; green in **d** and **h.** Staining against green fluorescent protein (GFP) antibody targeting the *mVenus* tag of *CsChrimson*. **b**, **f**; red in **d** and **h.** Staining against BP104. **c**, **g**; blue in **d** and **h.** Staining against N-cadherin. **a–h.** Plan-Apochromat 20x objective, resolution: 592 x 800 pixels, scale bar: 100 µm. Image courtesy of the HHMI Janelia FlyLight team.

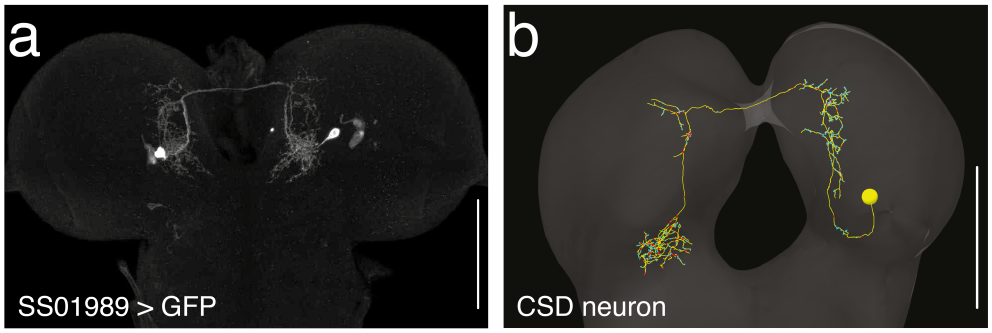

Fig. 6.9 SS01989 exclusively drives expression in the CSD neuron. a. Confocal image of a third-instar *SS01989 x UAS-GFP* larva CNS, derived from maximum intensity projections, obtained after immunohistochemical staining against *GFP*. C-Apochromat 40x objective, resolution: 975 x 651 pixels, scale bar: 100 µm. Image courtesy of the HHMI Janelia FlyLight team. **b.** Electron microscopy reconstruction of the contralaterally projecting serotonin-immunoreactive deutocerebral (CSD) neuron from the central nervous system of a first-instar larva (Berck et al., 2016), scale bar: 50 µm.

exclusively by dopaminergic or serotonergic neurons, respectively. Under our high-throughput training protocol (Fig. 6.4a), *TH-Gal4 x UAS-CsChrimson* larvae showed no difference in cast rate between the previously stimulated and unstimulated sides in the one-minute test periods either before training or after training (Fig. 6.7c, Fig. 6.7d). Larval fatigue also seemed to occur in these larvae, suggested by the significant decrease in cast rate to both stimulated and unstimulated sides after training (Fig. 6.7c). While the cast rate after training was distinctly higher than before training, this increase was not statistically significant (Fig. 6.7d). Based on these results, activating these dopaminergic neurons appeared to be an insufficient substitute for reward or punishment in operant conditioning. Bootstrap analyses, however, were inconsistent with the statistical inference results associated with the difference in cast rate (Table 6.2). Bootstrap analysis showed a significant proportion of bootstrap comparisons in which the mean difference in cast rate after training is greater than 0 ($probability = 0.9610$). Furthermore, bootstrap analysis showed a significant proportion of bootstrap comparisons in which the mean difference in cast rate after training is higher than that before training ($probability = 0.9827$). Inclusion of larvae tracked for less than a minute also revealed a statistically significant difference in cast rates to either side after training (Fig. 6.10a, Fig. 6.10b). Together these inconsistencies suggest that the original statistical inference is not robust to measurement error. Consequently, we cannot exclude the possibility that activation of these dopamine neurons can serve as a US in operant conditioning.

Before training, *Tph-Gal4 x UAS-CsChrimson* larvae showed no difference in cast rate between sides (Fig. 6.7c, Fig. 6.7d). Paired activation of *Tph-Gal4* neurons during casts to one side resulted in a significantly higher cast rate to the stimulated side relative to the unstimulated side during the test period after training (Fig. 6.7c, Fig. 6.7d), though the elevated difference in cast rate after training was not significantly higher than that observed before training (Fig. 6.7d). In accordance with all other genotypes described so far, cast rate comparison between the before training and after training periods showed a significant decrease in cast rate to both sides following training (Fig. 6.7c), suggesting a common behavioural outcome within this operant conditioning paradigm. No statistically significant difference was observed when directly comparing the difference in cast rate between *TH-Gal4 x UAS-CsChrimson* and *Tph-Gal4 x UAS-CsChrimson* larvae after training (Fig. 6.7e). Bootstrap analysis was, however, inconsistent with both statistical results in the period after training (Table 6.2). This analysis showed a large, but not significantly large, proportion of bootstrap comparisons

Table 6.2 Bootstrap sensitivity analyses of data in Fig. 6.7. Each statistical inference is denoted by genotype and Comparison. The associated Bootstrap Test is defined by a threshold criterion on the relationship between means. The Bootstrap Proportion reflects the proportion of all bootstrapped data sets that satisfied the Bootstrap Test. Consistency with Observed P-value indicates whether the direction and magnitude of the Bootstrap Proportion supports the original experimental finding. Inconsistencies are highlighted in grey.

	Comparison	Bootstrap Test	Bootstrap Proportion	Consistency with Observed P-value
Fig. 6.7c – #casts/min				
TH-Gal4	before training	p(stimulated > unstimulated)	0.1061	consistent
	after training	p(stimulated > unstimulated)	0.9338	consistent
	stimulated side	p(after < before)	0.9985	consistent
	unstimulated side	p(after < before)	1.000	consistent
Tph-Gal4	before training	p(stimulated > unstimulated)	0.7930	consistent
	after training	p(stimulated > unstimulated)	0.8703	inconsistent
	stimulated side	p(after < before)	0.9999	consistent
	unstimulated side	p(after < before)	1.000	consistent
Tph-Gal4, tsh-LexA, LexAop-Gal80	before training	p(stimulated > unstimulated)	0.1159	consistent
	after training	p(stimulated > unstimulated)	0.3022	consistent
	stimulated side	p(after < before)	1.000	consistent
	unstimulated side	p(after < before)	1.000	consistent
Fig. 6.7d – Difference in #casts/min				
TH-Gal4	before training	p(before > 0)	0.0870	consistent
	after training	p(after > 0)	0.9610	inconsistent
	after *vs* before training	p(after > before)	0.9827	inconsistent
Tph-Gal4	before training	p(before > 0)	0.8570	consistent
	after training	p(after > 0)	0.9320	inconsistent
	after *vs* before training	p(after > before)	0.4375	consistent
Tph-Gal4, tsh-LexA, LexAop-Gal80	before training	p(before > 0)	0.0270	inconsistent
	after training	p(after > 0)	0.2230	consistent
	after *vs* before training	p(after > before)	0.9040	consistent
Fig. 6.7e – Difference in #casts/min after training				
Tph-Gal4 *vs* TH-Gal4		p(Tph > TH)	0.4097	consistent
Tph-Gal4 *vs* Tph-Gal4, tsh-LexA, LexAop-Gal80		p(Tph > Tph tsh)	0.9430	inconsistent

in which the mean cast rate to the stimulated side was greater than that to the unstimulated side after training ($probability = 0.8703$). Bootstrap analysis of the difference in *Tph-Gal4 x UAS-CsChrimson* larval cast rate after training also showed a large, and nearly significant, proportion of bootstrap comparisons in which the mean difference in cast rate was greater than 0 ($probability = 0.9320$). Inclusion of larvae tracked for less than a minute also showed no significant difference in cast rates to either side after training (Fig. 6.10a, Fig. 6.10b). Together, these results suggest caution when interpreting the *Tph-Gal4 x UAS-CsChrimson* cast rate results. Activation of *Tph*-positive serotonergic neurons paired with casts to one side may contribute to the formation of a learned direction preference, though further data acquisition and analysis of the underlying measurement distributions is necessary to have more confidence in this conclusion. Knowing that operant conditioning was impaired following restriction of *Ddc-Gal4 x UAS-CsChrimson* expression to the brain and SEZ suggests that the serotonergic neurons of the VNC are perhaps contributors to memory formation in this paradigm. Because *Tph-Gal4* is a broad driver line, it is possible that its expression pattern contains brain or SEZ neurons outside of those in *Ddc-Gal4*. The existence of

these neurons may have induced the observed learning in the non-bootstrapped data set through an alternate mechanism independent from that which drove memory formation following *Ddc* neuron activation.

To assess whether VNC serotonergic neurons were necessary for the observed operant conditioning effect, we used *tsh-Gal80* to restrict the *Tph-Gal4* expression pattern to the brain and SEZ. Although these larvae appeared to show no direction preference prior to training (Fig. 6.7c, Fig. 6.7d), bootstrap analysis of the difference in cast rates shows a significant proportion of bootstrap comparisons for which this difference was less than 0 (proportion above 0 was 0.0270; Table 6.2). Further investigation either through additional data collection or analysis is required to determine whether this inconsistency is of biological relevance. Paired optogenetic activation of *Tph-Gal4 x UAS-CsChrimson; tsh-LexA, LexAop-Gal80* with larval casts to one side was insufficient to induce a learned direction preference consistent with operant conditioning (Fig. 6.7c, Fig. 6.7d). Consistent with all other tested genotypes, cast rate to both the stimulated and unstimulated sides was significantly lower after training compared to before (Fig. 6.7c). A direct comparison of the difference in cast rate after training between *Tph-Gal4 x UAS-CsChrimson; tsh-LexA, LexAop-Gal80* and *Tph-Gal4 x UAS-CsChrimson* larvae showed a significant difference between groups (Fig. 6.7e), though bootstrap analysis was inconsistent with this finding (Table 6.2). The proportion of bootstrap comparisons in which the difference in cast rate for *Tph-Gal4 x UAS-CsChrimson* larvae exceeded that of *Tph-Gal4 x UAS-CsChrimson; tsh-LexA, LexAop-Gal80* larvae was 0.9430. Because this falls just short of the 0.95 threshold for significance, the original statistical inference is not robust enough to measurement error and caution should be used in interpreting the associated comparison. The *Tph-Gal4* expression pattern contains two neurons per VNC hemisegment (with the exception of a single neuron in each A8 abdominal hemisegment), all of which are serotonergic (Huser et al., 2012). Future experimentation exclusively targeting VNC serotonergic neurons could be valuable in not only determining whether these neurons are sufficient to establish operant learning like that observed following *Ddc* activation, but also give insight into whether synergistic activity of serotonergic neurons from both the VNC and the brain or the SEZ may mediate the observed effect.

Under a classical conditioning paradigm, we have confirmed that there exist learning pathways in *Drosophila* that rely on serotonergic neurons. We have also shown that serotonergic neurons may serve as a sufficient US for operant conditioning, though we cannot exclude the possibility that dopaminergic neurons may themselves be

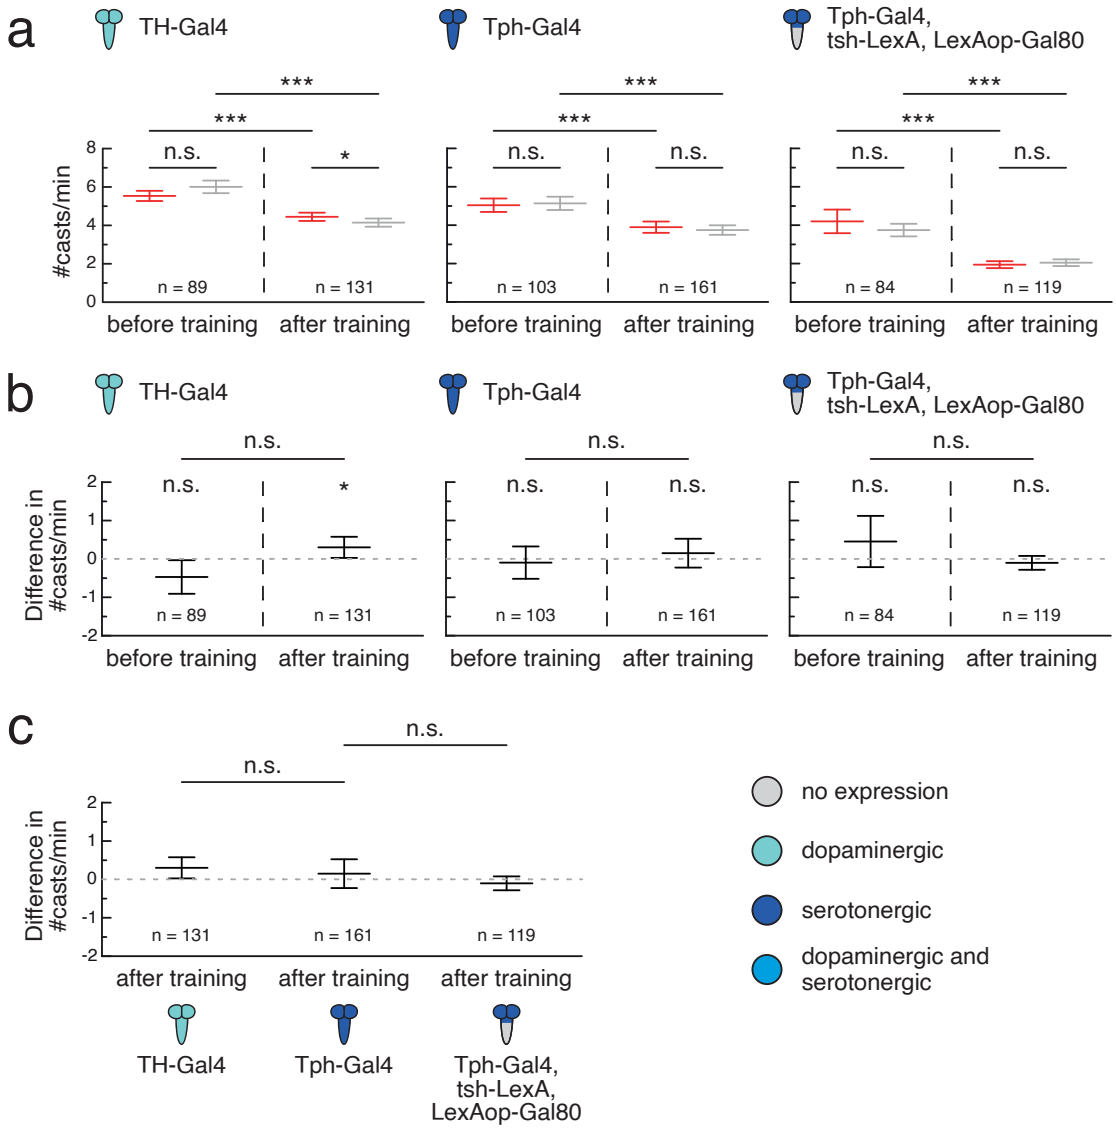

Fig. 6.10 *TH* and *Tph* operant conditioning data sets including all larvae regardless of tracking. Experiments were conducted following the protocol outlined in Fig. 6.3. Gal4 expression depicted as color-coded CNS. All fly lines contained the *UAS-CsChrimson* effector, which is omitted from the figure for visual clarity. Criteria to retain objects for analysis is identical to that described in Section 2.2.1 with the exception that larvae detected for only part of each bin were also included. All data is shown as (mean $\pm$ s. e. m.). **a.** Larval cast rate shown as the number of casts per minute, grouped by cast direction. The cast rate to the stimulated side is shown in red and the cast rate to the unstimulated side is shown in grey. Data is shown from the test period before the first training session and the test period immediately following the fourth training session. Statistical differences within groups (i. e. within before training time bin or within after training time bin) were tested with a two-sided Wilcoxon signed-rank test. Statistical differences between two groups (e. g. stimulated side before training versus stimulated side after training) were tested with a two-sided Mann-Whitney *U* test. n. s. $p \geq 0.05$ (not significant), * $p < 0.05$, *** $p < 0.001$. **b.** Difference in cast rate between stimulated and unstimulated sides. Grey dashed line at 0 indicates equal cast rates to either side. Data is shown from the test period before the first training session and the test period immediately following the fourth training session. Within-group statistics calculated from a two-sided Wilcoxon signed-rank test. Between-group statistics calculated from a two-sided Mann-Whitney *U* test. n. s. $p \geq 0.05$ (not significant), * $p < 0.05$. **c.** Same after training data as in **b**, with statistical comparisons to *Tph-Gal4 x UAS-CsChrimson* calculated using a two-sided Mann-Whitney *U* test with Bonferroni correction. n. s. $p \geq 0.05/2$ (not significant).

sufficient. Future experimentation is necessary to form more robust conclusions, which themselves may have implications for whether classical and operant conditioning are mediated by different circuit mechanisms.

6.4 Discussion

Due to available genetic tools and the emerging connectome, the *Drosophila* larva is a uniquely advantageous model organism for neuroscience. We explored whether classical and operant associative learning share neurotransmitters or circuit components. Future experimentation will help uncover whether there exist fundamental differences in connectivity, function, and underlying neural mechanisms between these learning paradigms. Notably, the experimental system we built was instrumental in investigating the neural circuits of operant conditioning, as it combined FPGA-based real-time tracking of multiple larvae with robust online behaviour detection and closed-loop stimulus presentation. This system could facilitate further research in taxis (Gomez-Marin et al., 2011; Jovanic et al., 2019; Kane et al., 2013; Luo et al., 2010), decision-making (DasGupta et al., 2014; Krajbich, 2019), and spatial navigation and memory (Haberkern et al., 2019; Neuser et al., 2008). While further work is necessary, our cast direction paradigm provides a strong foundation for continued study of the circuit and cellular mechanisms underlying operant conditioning.

6.4.1 High-throughput operant conditioning in *Drosophila* larvae

We have shown that *Drosophila* larvae are capable of operant conditioning and that optogenetic activation of *Ddc* neurons serves as a rewarding US during this learning process. With training, larvae formed an association between the presence or absence of this US and the direction in which they were casting. During testing, in the absence of any stimulation, larvae showed a significant learned preference for casting towards the previously stimulated side. This learned modification of future behaviour was only observed when the CS and US were paired during training; a hallmark of operant conditioning. Because *Ddc-Gal4* drives expression in dopaminergic and serotonergic neurons (Li et al., 2000; Sitaraman et al., 2008), we concluded that one or both neurotransmitters are involved in memory formation under these experiment conditions.

Strong parallels exist between our operant learning paradigm and those employed for conditioning direction preference in adult *Drosophila*. Consider the work of Nuwal

et al. (2012), in which tethered flies walked on a rotating ball and were rewarded with optogenetic activation of sugar-sensing neurons upon turns to one direction. As a consequence of this pairing, the flies learned to increase the number of turns to this side. Notably, the nature of the US remains an important difference between these paradigms. Our initial attempts to operantly condition larvae using activation of sugar-sensing neurons as a rewarding US were unsuccessful. These neurons, defined by two different *Gr43a-Gal4* drivers, were also insufficient for memory formation when activated with a paired odour in an olfactory conditioning screen. This was surprising, considering extensive evidence that natural sugar can serve as a rewarding US for classical conditioning in larvae (Apostolopoulou et al., 2013; Honjo and Furukubo-Tokunaga, 2005; Neuser et al., 2005; Rohwedder et al., 2012; Scherer et al., 2003; Schipanski et al., 2008; Schleyer et al., 2015; Weiglein et al., 2019). One possible explanation for these discrepancies is that multiple groups of sensory neurons must be co-activated in order to relay a meaningful reward signal. Alternatively, it may be necessary to adjust the temporal pattern or intensity of optogenetic stimulation.

It remains to be seen whether operant learning can occur by pairing roll or back-up behaviour with reward or punishment. Conditioning these actions is challenging given their infrequency in naïve, freely behaving animals. Rolls only occur in response to noxious stimuli (Ohyama et al., 2013, 2015; Robertson et al., 2013; Tracey et al., 2003). Back-ups also occur at very low rates (Masson et al., 2020). Consequently, the amount of US which larvae would receive during paired training would be very small, making observable memory formation more difficult. Our high-throughput tracker could potentially address this challenge with probabilistic, thermogenetic activation of roll or back-up command neurons (Carreira-Rosario et al., 2018; Ohyama et al., 2015) and optogenetic reward when performing the desired action.

One possible side effect of the multi-larva system design is unintended optogenetic stimulation of nearby larvae. Visual inspection of the behaviour arena during DMD calibration and stimulation protocol testing revealed scattering of visible light on the agarose plate around each projected 1 cm^2 square. It is therefore possible that red light projected at target larvae during closed-loop experiments may activate neurons in nearby larvae that themselves are not performing the behaviour of interest. Although the intensity of scattered red light is markedly lower than that within a given target square, further investigation is needed to measure and mitigate the behavioural impact of any stimulus cross-talk. We programmed the existing multi-larva software to account for any overlap of 1 cm^2 stimulation areas between proximate larvae by projecting the lowest

stimulus intensity to the overlapping region (see Section 2.2.1). The scattered light of greatest concern may therefore exist outside of any possible overlapped region (i. e. more than 1 cm away from a target larva). One approach for measuring the effect of ambient light is to project a DMD stimulation square at least 1 cm away from a larva's location and observe any coincident larval behaviour. Behavioural responses can also be recorded as a function of distance from the projected light source. In an attempt to quantify any stimulus cross-talk, it would be fitting to consider the rolling phenotypes from *69F06-Gal4 x UAS-CsChrimson* and *72F11-Gal4 x UAS-CsChrimson* larvae. Their behaviour was noticeably robust in the proof-of-principle optogenetic stimulation experiments (Fig. 6.2c). To address any concerns that high intensity basal stimulation is required to evoke energy-intensive rolling behaviour, a more sensitive effector could also be explored. Optogenetic activation of *OK6-Gal4 x UAS-CsChrimson* larvae at lower light intensity than the maximum available on our system (see Chapter 2) has been shown to temporarily paralyse larvae via simultaneous muscle contraction (Hernandez-Nunez et al., 2015). Such a distinct behavioural phenotype could help in determining whether scattered light from DMD projections is sufficient to activate neurons in nearby non-targeted larvae. It may be possible to mitigate any observed cross-talk by adjusting the agarose plate concentration so the surface is less reflective and more absorptive. The effects of changing agarose concentration on baseline larval behaviour must, however, be considered (Apostolopoulou et al., 2014a). Other mitigation strategies include reducing the size of the stimulation square or shutting off all DMDs projections for larvae within activation distance of one another.

6.4.2 Neural circuits of operant conditioning

From the available data, it cannot be concluded that the brain and SEZ are dispensable for operant conditioning in *Drosophila* larvae. Examples from both vertebrates (Grau et al., 1998) and invertebrates (Booker and Quinn, 1981; Horridge, 1962) show the spinal cord or VNC as sufficient for learning, suggesting that conserved mechanisms exist for brain-independent operant conditioning across species. This does not, however, exclude the possibility that there exist alternative learning pathways using the brain. In mammals (Balleine et al., 2009; Redgrave et al., 2011) and birds (Fee and Goldberg, 2011), brain correlates of operant conditioning have been identified. It is unclear where such pathways would be located in the insect brain. Previous adult fly studies (Booker and Quinn, 1981; Colomb and Brembs, 2010, 2016; Wolf et al., 1998) support the idea

that operant conditioning can occur independently of the MB, such that other learning centres might exist. To determine whether larval operant conditioning can be fully mediated by the VNC or whether the brain or SEZ are necessary, new driver lines must be created. A collection of sparse split-GAL4 lines could help identify the minimum subset of neurons necessary for conveying a US in our cast direction paradigm.

Even if the learning signal for operant conditioning can be mapped to a few neurons, there remain several open questions regarding the neuronal mechanisms underlying this learning. Locally, neurons could drive synaptic plasticity or modulate the intrinsic excitability of their postsynaptic partners. Alternatively, the learning signal could propagate further downstream, yielding learning correlates elsewhere in the network. Furthermore, memory formation requires integrating the US with information about the occurrence of the reinforced action. Motor feedback (e. g. efference copy, Fee, 2014; Webb, 2004) or proprioceptive input could be used to transmit this movement signal to higher-level circuits for convergence with the valence-encoding US. However, if memory formation occurred at a lower level, the action-specific signal and associated valence could be locally integrated inside the motor or premotor neuron without the need for feedback loops.

Lorenzetti et al. (2008) proposed intracellular mechanisms for modulating the intrinsic excitability of the *Aplysia* premotor neuron B51 during operant conditioning, mediated by the highly conserved protein kinase C (PKC) gene. PKC signalling is also essential for operant conditioning in *Lymnaea* (Rosenegger and Lukowiak, 2010) and adult *Drosophila* (Brembs and Plendl, 2008; Colomb and Brembs, 2016). If the *Drosophila* larva showed evidence of PKC-induced motor neuron plasticity, EM reconstruction of the pathways between neurons mediating operant learning and the PKC-positive motor neurons could further elucidate the mechanisms of memory formation and retrieval. A similar investigation of the *Drosophila* gene *FoxP* may also be informative, as its mutation in the adult results in impaired operant self-learning (Mendoza et al., 2014). The vertebrate homologue *FOXP2* is associated with deficits in human speech acquisition (Lai et al., 2001), song learning in birds (Haesler et al., 2007), and motor learning in mice (Groszer et al., 2008).

It remains unclear to what extent dopaminergic and/or serotonergic neurons are involved in the operant conditioning effect we observed. Further investigation is necessary to better understand the function of these neurotransmitters in operant learning memory formation. It is possible that even a single instance of learning leads to a variety of changes across the nervous system. In the case of operant conditioning,

higher brain centres, motor command neurons, premotor circuits and motor neurons would all qualify as potential learning sites. In addition to thoroughly analysing the expression patterns of driver lines used in our classical conditioning screen, developing new, sparse driver lines targeting dopaminergic and serotonergic neurons would be valuable for identifying the minimal subset of neurons which provide the learning signal. The larval connectome could be used to subsequently trace the paths from these neurons to the MB. One could then test whether learning remains intact when these connections are silenced. The expression pattern of dopamine and serotonin receptors could also provide clues about whether and how these signals trigger learning. One should certainly consider the possibility that learning is not induced by dopamine and/or serotonin themselves, but by other coexpressed neurotransmitters.

Chapter 7

Discussion

7.1 Larval sensory discrimination task

In this thesis, I have introduced a sensory discrimination task to investigate navigational decision-making behaviour in *Drosophila* larvae. Given the necessity of larval head casts to local sensory information acquisition, I aimed to closely investigate the influence of increased head casting on larval performance in reorientation manoeuvres. I asked for the first time whether and how larvae improve their decision about which direction to turn the longer they spend casting (see Chapter 3).

It is known that the larval nervous system is highly sensitive to changes in stimuli as opposed to static stimulus values. Research in CO_2, odour, temperature, and light gradients has revealed how the derivative of stimulus intensity can drive transitions between navigation behaviours (Gershow et al., 2012; Gomez-Marin et al., 2011; Kane et al., 2013; Klein et al., 2015; Schulze et al., 2015). In each of these environments, larvae bias the initiation of a reorientation manoeuvre based on derivative changes that increasingly signal movement toward unfavourable conditions. The reverse occurs to end a manoeuvre: larvae are more likely to turn in the direction set by a head cast if that head cast coincided with a derivative change signaling movement toward favourable conditions. These observations emphasise that larvae can make accurate decisions about where to end a reorientation manoeuvre based on sensory information from their environment. Nearly all of this investigation, however, focused on either the first or last cast individually or all casts irrespective of their numerical placement in the manoeuvre.

In my analysis of larval behaviour, I found a significant positive correlation between head cast number and the likelihood that larvae choose to end the reorientation manoeuvre on the less noxious side. This increase in probability of correct decision is

significantly more pronounced for conditions in which a stimulus differential exists between sides compared to control conditions in which the stimulus is the same on both sides. This difference is qualitatively apparent in the divergence of $P(\mathrm{correct})$ values between experimental and control conditions at higher cast numbers. My findings from larval behaviour within the experimental conditions are consistent with increased environmental sampling possibly facilitating evidence accumulation that improves decision-making.

These results appear to contradict a prevailing view in the field that larval navigation is based on the derivative change computed immediately prior to a behaviour transition (Wystrach et al., 2016). Some studies of larval taxis have sought to better characterise the sensory history preceding the end of a reorientation manoeuvre. A common analytic approach for relating sensory cell activity to observed behaviour is reverse-correlation statistics. Experiments involve supplying random trains of stimuli to sensory cells. Reverse-correlation performed on larval behaviour traces reveals the stimulus history that triggered the behaviour transition of interest. These results are then used to build parameters for linear-nonlinear models (LNMs) that predict navigation decision dynamics in response to novel sensory information. LNMs have been shown to successfully predict these larval behavioural transitions in response to fictive olfactory (Gepner et al., 2015; Hernandez-Nunez et al., 2015) and gustatory (Hernandez-Nunez et al., 2015) stimuli in addition to naturally aversive blue light (Gepner et al., 2015). These studies have been inspired by the observation that stimulus derivatives monotonically increase to a peak just prior to the acceptance of a cast (Gomez-Marin et al., 2011; Kane et al., 2013; Klein et al., 2015). Researchers have interpreted this finding as evidence for a sensorimotor algorithm in which sensory integration dictates the transition into or out of a reorientation manoeuvre. Wystrach et al. (2016) observed similar sensorimotor interactions in their own computational model of larval chemotaxis. At the same time, however, their model recapitulates larval behaviour in an odour gradient based on the perceived stimulus change from only the immediately preceding time step. My sensory discrimination task results are unique in suggesting that larvae rely on sensory information acquired over each successive cast to better inform their decision about which direction to end the reorientation manoeuvre.

It has been speculated that larvae rely on short term memory to successfully compare computed changes in sensory stimuli during a reorientation manoeuvre (Gomez-Marin et al., 2011). I take this a step further by speculating that short term memory aids in the improvement of decisions over multiple casts in my sensory

discrimination task. In other words, larvae may use short term memory to relate proprioceptive information to perceived sensory stimuli, continuing to improve their understanding of that relationship over several casts. Despite originating from an appetitive conditioning paradigm and over a longer time scale, we have shown that larvae are capable of relating body position to sensory stimuli through operant conditioning of cast direction (see Chapter 6). Further investigation into the larval sensory discrimination task can help elucidate whether larvae have the capacity for short term memory.

7.1.1 Varying task design

Deciphering larval behaviour in my sensory discrimination task remains a complex problem. Additional experimentation with other task designs will not only broaden our understanding of larval decision-making but also inform the extent to which my current interpretation of evidence accumulation holds across different environments.

A logical first approach will be to consider other light intensity combinations for the noxious stimulus alternatives, the most notable being the experimental condition $E_{M,H}$ that I was unable to test. Another modification could include depriving larvae of food for a longer period. Research in both adult *Drosophila* (Corrales-Carvajal et al., 2016) and honeybees (Katz and Naug, 2015) emphasises the importance of an animal's internal state on foraging decisions. Nutrition is perhaps the most relevant internal state to consider for *Drosophila* larvae. Evidence from foraging adult *Drosophila* shows that starved individuals exhibit greater local exploration of single food patches compared to satiated flies that perform a wider search between patches (Corrales-Carvajal et al., 2016). In a similar way, greater food deprivation for larvae may encourage greater numbers of head casts per reorientation manoeuvre in my sensory discrimination task. Future investigation could also include more precise timing of larval acclimation to the agar petri dish (see Section 2.1.3) so that individuals are more precisely matched according to satiety.

Transitioning to high-throughput experimentation will be advantageous for future exploration of larval behaviour within the various aforementioned sensory discrimination task designs. The existing hardware framework and software design of our multi-larva closed-loop system facilitate fast tracking of larval behaviour and precise targeting of optogenetic stimuli (see Chapter 6). Before this system is ready to run simultaneous sensory discrimination tasks across several larvae, a few software elements need to be

developed. The first is a custom stimulation protocol following each task design. Specifying the relationship between detected behaviour and presented stimulus in the written software code is standard for all experiment protocols implemented by this system. The second necessary software element involves robust detection of both the start and end of a reorientation manoeuvre. The former can be easily implemented based on automated detection of forward crawls in accordance with my design on the single-larva system (see Section 2.1.3). The latter will require a combination of larval motion, shape, and velocity features that capture two forward peristaltic waves regardless of the larva's tail movement or alignment along a straight body axis (see Section 2.1.3). As with other forms of behaviour detection, the system's analysis of the start and end of reorientation manoeuvres will require human validation prior to widespread use. Once finalised, this addition to the system's behaviour detection software will eliminate any need for the manual online validation and offline analysis that I performed on the existing larval behavioural data (see Sections 2.1.3 and 2.1.4). The increased efficiency of data analysis will encourage greater and more varied testing of different experimental hypotheses.

7.1.2 Considering other larval genotypes

An important caveat to my sensory discrimination task results is the relationship of the fictive noxious stimulus to the larva's body position. In the task I designed, larval head casting triggers differential presentation of red light stimuli that physically envelop the whole larva. Because *ppk1.9-Gal4* dendrites span the entire larval body wall, such optogenetic stimulation during head casting also differentially excites all *ppk1.9* neurons in the abdominal and tail segments despite those areas of the larval body remaining stationary. This mismatch between noxious information and postural feedback posterior to the head and thoracic segments is a key limitation to effective analysis of larval decision-making behaviour in my existing closed-loop sensory discrimination task. Future experiments using the multi-larva closed-loop system can help address this concern. One approach could be to reduce the size of the projected stimulus area and reposition the projection to target *ppk1.9-Gal4* activation solely within the larva's head. Any head casts would thus be strictly correlated with changes in sensory information detected via the head. The multi-larva closed-loop system is well-equipped to implement this approach. Not only does the software algorithm already localise the head landmark (see Section 2.1.1), but the DMDs operate with a spatial resolution of 291.7 µm/pixel

(see Chapter 2). This resolution is small enough to pinpoint optogenetic stimuli to anterior thoracic segments T1-T3 of third instar *Drosophila* larvae, which together measure approximately 600 µm in length (Berni et al., 2012).

Future experiments could also involve investigating other genetic lines in the context of this larval sensory discrimination framework. It is unknown whether larval decision-making in response to nociception manifests differently compared to other sensory information. Investigating other information processes within the larva's head (e. g. taste or odour discrimination) may elucidate whether the larval behaviour strategy I observed occurs regardless of the sensory landscape. It is also worth considering whether changing the valence of the incoming sensory information influences head casting dynamics. We have shown that larvae are capable of operant learning following reward-based conditioning of cast direction (see Chapter 6). Perhaps larvae will be motivated to continue exploring their environment following either differential activation of the same dopaminergic *Ddc-Gal4* neurons or fictive presentation of innately rewarding gustatory or olfactory stimuli. Other candidate genetic lines include those that consistently evoke appetitive memory in larval associative learning tasks. Larval MBONs are notably strong candidates for study, given their role in instructing approach or avoidance behaviour as a function of previously learnt stimuli valences. MBON-m1 and MBON-i1 are two MBONs whose distinct structural and functional relationships with a shared dopaminergic MB input neurons (MBINs) have been characterised (Eichler et al., 2017; Eschbach et al., 2020b). It would be informative to explore whether the activity of these MBONs biases head casting dynamics in accordance with evidence accumulation in my sensory discrimination task.

7.2 Larval behaviour models

7.2.1 Bayesian inference and acceptance pressure

My goal in formulating the Bayesian inference and acceptance pressure models was to computationally elucidate the underlying behavioural process that leads larvae to choose the less noxious of two stimuli. I conclude based on their fits to the larval data that I cannot accept either model as the governing behavioural mechanism for my sensory discrimination task. Neither model explains notable qualitative features of *both* the experimental and control conditions. Instead, Bayesian inference better fits the experimental conditions while acceptance pressure better fits the controls (see

Chapters 4 and 5). This difference prompts me to speculate whether a combination of processes described by each model underlies the larval behavioural readout.

Recall that in the case of *all casts* Bayesian inference, the decision surrounding final cast direction results from direct comparison of the most likely nociceptive level based on continually evolving posterior beliefs. In the *integ noci* subtype of the acceptance pressure model, ending a reorientation manoeuvre could be interpreted as a larva "giving up" after exceeding a threshold of accumulated stimulus intensity. Now consider the possibility that larvae detect differences in noxious stimulus intensity, regardless of their magnitude of separation, within the first two casts of a reorientation manoeuvre. In this way, context may drive larvae to select between the Bayesian inference and acceptance pressure strategies beginning on cast three. This scenario is supported by the similarity in $P(\mathrm{correct})$ between each experimental condition and its high intensity control ($E_{L,M}$ and $C_{M,M}$; $E_{L,H}$ and $C_{H,H}$) on casts one and two, followed by divergence between these experimentals and their controls beginning on cast three and continuing over higher cast numbers (Fig. 3.3b). Perhaps *Drosophila* larvae in my sensory discrimination task detect a stimulus difference (or lack thereof) within the first two casts of a reorientation manoeuvre that biases them to adopt a particular behavioural strategy across the remaining casts? While I am not aware of evidence for whether or how the fly brain chooses between competing strategies, there is literature describing how vertebrates engage in strategy switches based on environmental context. A recently described example is that of larval zebrafish switching from active to passive swimming behaviour. Notably, this switch occurs after rigorous attempts to move through their environment fail to resolve a mismatch between their motor pattern and external visual cues (Mu et al., 2019).

7.2.2 Related formulations based on evidence accumulation

In developing and analysing the Bayesian inference and acceptance pressure models, I established a valuable reference point for future modeling pursuits. The best performing subtype of each model construction is contingent on evidence accumulation with time. In the case of *all casts* Bayesian inference, this manifests as two posterior belief distributions, one for each side's true nociceptive level, updated by newly acquired stimulus samples on each cast. The *integ noci* acceptance pressure model equates accumulated information to the summed total noxious stimulus experienced by the larva. I prioritised the evidence accumulation principle in accordance with my observation that

repeated head casts significantly improved decision-making in both experimental conditions. This remains a viable hypothesis for the observed larval behaviour, supported, in part, by the existence of sensory evidence accumulation mechanisms in adult *Drosophila*. Biophysical properties of the fly's $\alpha\beta$ core ($\alpha\beta_c$) Kenyon cell (KC) dendrites enable the integration of sequential samples of olfactory information during perceptual decision-making tasks. These samples take the form of subthreshold changes in membrane potential. Membrane depolarisations contribute to an evolving decision variable, changing magnitude and speed in accordance with the difference in contrast between odourants. Any evoked action potential thereby represents the decision bound; a conclusion further supported by its closely timed relationship to the fly's behavioural decision (Groschner et al., 2018; Groschner and Miesenböck, 2019).

My larval evidence accumulation hypothesis also helps parse the vast space of models that could be explored in the context of my sensory discrimination task. Considering that sensory discrimination is itself a form of perceptual decision-making, it may be advantageous to direct future investigation toward computational models derived from known perceptual decision-making tasks. One such task is two-alternative forced choice (TAFC) (see Chapter 1). Of available computational frameworks, the drift diffusion model (DDM) represents the optimal solution for reaction time and response accuracy in TAFC (Bitzer et al., 2014; Bogacz et al., 2006; Eckhoff et al., 2008; Tajima et al., 2016). The DDM is an integrator model where, at each time step, the animal accumulates evidence for one of the two choice alternatives. A phenomenon known as the speed-accuracy trade-off is reflected in the magnitude of evidence required to favor one choice over the other. While noise disturbance contributes to the random walk of accumulated evidence, a decision is made either when one of the two choice thresholds is reached or evidence runs out (Bitzer et al., 2014; Bogacz et al., 2006; de Lafuente et al., 2015; Gold and Shadlen, 2007; Li and Krishnamurthy, 2015; Shadlen and Kiani, 2013). Recent work in adult *Drosophila* has shown successful application of the DDM to a small, neurally tractable model organism. Specifically, the behavioural choices displayed by adult flies in an odour discrimination task follow drift diffusion dynamics (DasGupta et al., 2014). Consider, also, that DDM has been applied to studies of foraging behaviour (Calhoun et al., 2014). It is therefore logical to question whether the DDM can also explain the larval behavioural signatures in my sensory discrimination task. Results from such investigation could reveal a common decision-making mechanism in *Drosophila*, shared across distinct developmental stages and/or sensory modalities.

In light of ongoing debate over the generalisability of drift diffusion models to other decision-making tasks, research has been directed toward model constructions with broader applicability. Bitzer et al. (2014) developed a Bayesian inference formulation of TAFC and mathematically demonstrated how the inference results are equivalent to those of the DDM. In doing so, the researchers emphasise inherent advantages of Bayesian inference over the DDM. These include explicitly defining the model observer's prior knowledge and its perceptual uncertainty (Bitzer et al., 2014). Research on foraging behaviour has simultaneously considered the performance advantages of the Upper-Confidence-Bound (UCB) algorithm over Bayesian inference (Audibert et al., 2009; Morimoto, 2019). An observer following the UCB algorithm chooses an action that has the highest upper bound over its estimated value. The observer is optimistic in this choice, following through regardless of whether the magnitude of uncertainty surrounding the action's value estimate is large. Even if the resulting choice outcome is not equal in value to the upper confidence bound, the observer can still use the experience to reduce its uncertainty over the action's value (Morimoto, 2019). Whether the DDM or UCB algorithm could elucidate the mechanism underlying my observations of larval evidence accumulation remains an open question. It is indeed possible that either could yield a better fit to the larval sensory discrimination data than my Bayesian inference and acceptance pressure model constructions. Importantly, the applicability of both DDM and UCB to evidence accumulation processes across taxa makes them appealing and logical alternatives to consider in future modeling of my larval sensory discrimination task.

Future experimentation and modeling of larval decision-making behaviour in my sensory discrimination task could also include consideration of a foraging phenomenon known as the exploration-exploitation trade-off. Repeatedly sampling stimuli in the environment comes at the cost of energy resources and, in some circumstances, leaves animals vulnerable to predation. This conflict results in a trade-off between exploring unknown alternatives that may reap greater benefits or exploiting those for which some information is already known. Perhaps *Drosophila melanogaster* larval head casts can be viewed as a behavioural mechanism for assessing outcome value, estimating outcome uncertainty, and selecting a subsequent action. If so, then variable head casting sequences like those in my sensory discrimination task may reflect a decision-making process akin to the exploration-exploitation trade-off (Cohen et al., 2007). Animal behaviour researchers have found the well-known "n-armed bandit" problem especially conducive to investigating explore-exploit behaviour dynamics. With

origins in statistical decision theory and practical applications to machine learning, economics, and medicine (Steyvers et al., 2009), the bandit problem is a decision-making task with a structure analogous to a gambler playing casino slot machines (i. e. bandits). The player's goal is to maximise the total financial payout across a sequence of arm pulls from several slot machines. When they pull an arm, a payout is drawn from the machine's own reward distribution. While each machine's reward is often fixed, its value is always initially unknown to the player. This uncertainty forces the player to search among the machines, pulling arms and learning reward distributions to make the most informed decision. But continued search comes at the cost of losing additional money. What results is a trade-off between continuing to play the current best choice arm based on existing knowledge of the reward distributions (exploitation) or testing other unknown arms to potentially improve the long-term payout (exploration) (Audibert et al., 2009; Reid et al., 2016). In primates and rodents, recording neural activity and manipulating neurotransmitter concentration during bandit tasks has implicated large subcortical structures in regulating the balance between exploration and exploitation (Cinotti et al., 2019; Costa et al., 2019). Investigation in insects has been largely centred around *in silico* modeling (Morimoto, 2019) or testing learning rules or search algorithms against observed behavioural data (Keasar et al., 2002; Naug and Arathi, 2007). It will be interesting to develop a bandit task for *Drosophila* larvae under which different, yet fixed, probabilistic stimulus patterns are linked to cast direction. This would contrast my all-or-nothing stimulus protocol by closely mimicking reward distributions characteristic of multi-armed bandit problems. This addition of uncertainty in stimulus presentation may have interesting ramifications for larval decision-making behaviour. If, for example, larvae in a nociceptive bandit paradigm show continued casting, perhaps the benefit of greater certainty about which direction to choose outweighs the environmental risks and energetic costs associated with repeated exposure to noxious stimuli.

7.3 Investigating underlying neuronal circuitry

7.3.1 Identifying candidate regions or neurons

My sensory discrimination task comprised low-level, cast direction-dependent stimulation of larval nociceptive neurons during a reorientation manoeuvre. The observed improvement in navigational decisions with increasing cast number suggests an underlying sensory evidence accumulation mechanism. My existing data set does

not, however, purport a specific region of the *Drosophila* larval CNS as the locus for such computations.

Future investigation of underlying neuronal circuitry could be guided by existing knowledge of the sensorimotor transformations underlying nocifensive rolling behaviour. Although rapid, energy-intensive rolling is biomechanically distinct from head casting, it is partly triggered by activating multidendritic class IV sensory neurons (mdIVs); the same nociceptive neuronal population I stimulated in my sensory discrimination task. Research on rolling has been extensive, with at least three neuronal circuits found to relay primary sensory information from polymodal mdIVs. In the first of these circuits, mdIVs in each larval hemisegment target two of four Basin neurons. These Basins also receive mechanosensory information from upstream chordotonal neurons. From this first-order multisensory layer stems a local VNC pathway via the A05q neuron and a global ascending brain pathway via A00c. Both paths end by targeting Goro, a command neuron for rolling. This circuit also comprises SEZ feedback neurons that target Basins, A05q, and A00c (Ohyama et al., 2015). Yoshino et al. (2017) identified a second circuit comprising medial clusters of mdIV second-order interneurons (mCSIs) sufficient to trigger rolling independently of Basins and Goro. Hu et al. (2017) identified additional candidate interneurons downstream of mdIV: DP-ilp7 triggered the C-shape bend formed by larvae just prior to a roll, while A08n was sufficient to induce rolling (Hu et al., 2017). By virtue of their motor output, these circuits are unlikely to fully overlap those that drive continued head casting in my sensory discrimination task. But such work remains foundational to locating sites of convergence between mild nociceptive sensation and either proprioceptive information or motor commands.

Candidate regions or neurons implicated in larval taxis may also be a reasonable starting point to identify the sensory discrimination task's neuronal substrate. In homogeneous environments, transitions between runs and reorientation manoeuvres occur independently of the brain and SEZ. In some heterogeneous environments, sensory information processing localised to the VNC modulates larval navigation. In anemotaxis, silencing the thoracic ladder projection neuron significantly decreases the probability of ending a reorientation manoeuvre with a head cast directed downwind (Jovanic et al., 2019). Larval phototaxis can proceed unimpeded by the absence of brain or SEZ activity (Berni et al., 2012). Despite these findings, higher order centres also contribute to the modulation of larval behaviour that improves navigation (Berni et al., 2012). The SEZ has been implicated as a critical premotor centre in which odour, light, and temperature information converge to trigger the run-to-reorientation manoeuvre

transition (Tastekin et al., 2015). As it relates to chemotaxis, identification of the upstream PDM descending neuron has elucidated a functional sensorimotor connection between lateral horn interneurons, SEZ interneurons, and central pattern generators in the VNC (Tastekin et al., 2018). Neuronal circuits involved in navigating down wind gradients have also been discovered in the brain. Silencing the MB-innervating PAM cluster dopaminergic neurons causes larvae to adversely bias the first cast of a reorientation manoeuvre upwind (Jovanic et al., 2019).

The aforementioned circuits implicated in both larval rolling and taxis span the entirety of the CNS, with local loops in the nerve cord and/or additional sensory convergence centres in higher order regions. A rudimentary first approach may be necessary to narrow the field of candidates potentially responsible for improved perceptual decision-making over noxious stimulus alternatives. For example, inactivating proprioceptive projection neurons and their downstream targets during my sensory discrimination task could elucidate the role of VNC circuitry in improved decision-making. It will be crucial to confirm that these candidate neurons are not required to maintain baseline locomotor coordination. If the targeted ascending pathways are deemed necessary for sensory discrimination task behaviour, then subsequent inactivation experiments could involve higher order brain centres. *FoxP-GAL4*-positive $\alpha\beta_c$ KCs have been implicated in adult flies' ability to accumulate odourant information to a response threshold (DasGupta et al., 2014). Given their role in perceptual decision-making, MB KCs may serve a sensory integration role in my task. Results from future operant learning experiments could also suggest possible candidate circuits that relate valence signals with larval body position. A second approach to further narrow the field of candidates could be the application of reverse-correlation analysis to data from an unbiased optogenetic behaviour screen of GAL4 driver lines. Computational modeling work in *Drosophila* larval taxis has emphasised the utility of this statistical approach in identifying neurons responsible for behaviour state transitions and attributing valence to their activity (Gepner et al., 2015; Hernandez-Nunez et al., 2015; Klein et al., 2015).

7.3.2 Exploring connectomic data sets

Serial section electron microscopy (EM) volumes of the larval nervous system remain a valuable tool for elucidating how observed behaviour is supported by neuronal circuit architecture. Reconstructing neurons from these images can reveal important details about cellular morphology and synaptic connectivity. Our understanding of neuronal

functionality can be attributed, in part, to knowledge of such features. EM data have revealed highly stereotyped synaptic connectivity in sensory systems across individual larvae (Gerhard et al., 2017). Exploring available connectomic data downstream of mdIVs will undoubtedly complement experiment-based understandings of how and where noxious information in my sensory discrimination task combines with proprioceptive and/or motor signals.

7.3.3 Quantifying neuronal activity

Future investigations will also benefit from quantifying neuronal activity. In establishing stimuli conditions for the sensory discrimination task, I relied on larval behaviour readouts as a proxy for larval sensory perception. Measuring downstream intracellular calcium transients upon activation of mdIVs would not only confirm that larvae are sensing each stimulus but also quantify the magnitude of perceived stimulus contrast in the experimental conditions. This knowledge may then inform future model constructions by improving fundamental assumptions about the observer and its environment. Neuronal activity traces can also supplement larval connectomic data by elucidating whether sites of structural connectivity between candidate neurons are themselves functional. Such measurements will also reveal whether neuronal circuit activity reflects observed behavioural decision dynamics in my sensory discrimination task. Recording neuronal activity in unrestrained larvae is preferable to fictive behaviour preps that preclude proprioceptive feedback. Karagyozov et al. (2018) developed a system that combines closed-loop stimulus presentation with calcium imaging in freely behaving larvae. Usage of this system will enable observation of real-time calcium dynamics as larvae perform the sensory discrimination task. This simultaneous collection of behavioural and functional data will permit verification of candidate regions responsible for improvement of decisions with head casts.

Functional data has the added benefit of informing better assumptions within computational models of both behaviour and neural processing. Work by Schulze et al. (2015) is exemplary of this experimental approach. Their model relied on results from a combination of electrophysiology and behaviour experiments which revealed that a single olfactory sensory neuron, Or42a, responds to both stimulus intensity as well as concentration changes computed via its first derivative. Models at both the behavioural and circuit levels have the potential to yield experimentally testable predictions about larval decision-making under different conditions.

7.4 Concluding remarks

My work designing the larval sensory discrimination task, building behavioural models, and developing high-throughput tools lays the foundation for understanding navigational decision-making during *Drosophila* larval reorientation manoeuvres. My efforts also carry wide implications for how sensory processing by the larval nervous system supports learning and memory formation.

References

Abramson, C. I., Dinges, C. W., and Wells, H. (2016). Operant Conditioning in Honey Bees (Apis mellifera L.): The Cap Pushing Response. *PLoS ONE*, 11(9):e0162347.

Aceves-Piña, E. O. and Quinn, W. G. (1979). Learning in Normal and Mutant Drosophila Larvae. *Science*, 206(4414):93–96.

Agresti, A. (2002). *Categorical Data Analysis*. Wiley-Interscience, Hoboken, N.J., 2nd edition.

Ainsley, J. A., Pettus, J. M., Bosenko, D., Gerstein, C. E., Zinkevich, N., Anderson, M. G., Adams, C. M., Welsh, M. J., and Johnson, W. A. (2003). Enhanced Locomotion Caused by Loss of the Drosophila DEG/ENaC Protein Pickpocket1. *Current Biology*, 13(7):1557–1563.

Alekseyenko, O. V., Lee, C., and Kravitz, E. A. (2010). Targeted Manipulation of Serotonergic Neurotransmission Affects the Escalation of Aggression in Adult Male Drosophila melanogaster. *PLoS ONE*, 5(5):e10806.

Alexander, J., Audesirk, T. E., and Audesirk, G. J. (1984). One-trial reward learning in the snail Lymnea stagnalis. *Journal of Neurobiology*, 15(1):67–72.

Almeida-Carvalho, M. J., Berh, D., Braun, A., Chen, Y.-C., Eichler, K., Eschbach, C., Fritsch, P. M. J., 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., Paisios, E., Pankratz, M. J., Petersen, M., Ramsperger, N., Randel, N., Risse, B., Saumweber, T., Schlegel, P., Schleyer, M., Soba, P., Sprecher, S. G., Tanimura, T., Thum, A. S., Toshima, N., Truman, J. W., Yarali, A., and Zlatic, M. (2017). The Ol1mpiad: concordance of behavioural faculties of stage 1 and stage 3 Drosophila larvae. *Journal of Experimental Biology*, 220:2452–2475.

Andreatta, M. and Pauli, P. (2015). Appetitive vs. Aversive conditioning in humans. *Frontiers in Behavioral Neuroscience*, 9:128.

Apostolopoulou, A. A., Hersperger, F., Mazija, L., Widmann, A., Wüst, A., and Thum, A. S. (2014a). Composition of agarose substrate affects behavioral output of Drosophila larvae. *Frontiers in behavioral neuroscience*, 8(January):11.

Apostolopoulou, A. A., Mazija, L., Wüst, A., and Thum, A. S. (2014b). The neuronal and molecular basis of quinine-dependent bitter taste signaling in Drosophila larvae. *Frontiers in Behavioral Neuroscience*, 8:6.

Apostolopoulou, A. A., Widmann, A., Rohwedder, A., Pfitzenmaier, J. E., and Thum, A. S. (2013). Appetitive Associative Olfactory Learning in Drosophila Larvae. *Journal of Visualized Experiments*, 72:e4334.

Aso, Y., Hattori, D., Yu, Y., Johnston, R. M., Iyer, N. A., Ngo, T.-T. B., Dionne, H., Abbott, L. F., Axel, R., Tanimoto, H., and Rubin, G. M. (2014a). The neuronal architecture of the mushroom body provides a logic for associative learning. *eLife*, 3:e04577.

Aso, Y., Herb, A., Ogueta, M., Siwanowicz, I., Templier, T., Friedrich, A. B., Ito, K., Scholz, H., and Tanimoto, H. (2012). Three Dopamine Pathways Induce Aversive Odor Memories with Different Stability. *PLoS Genetics*, 8(7):e1002768.

Aso, Y., Sitaraman, D., Ichinose, T., Kaun, K. R., Vogt, K., Belliart-Guérin, G., Plaçais, P.-Y., Robie, A. A., Yamagata, N., Schnaitmann, C., Rowell, W. J., Johnston, R. M., Ngo, T.-T. B., Chen, N., Korff, W., Nitabach, M. N., Heberlein, U., Preat, T., Branson, K. M., Tanimoto, H., and Rubin, G. M. (2014b). Mushroom body output neurons encode valence and guide memory-based action selection in Drosophila. *eLife*, 3:e04580.

Audibert, J.-Y., Munos, R., and Szepesvári, C. (2009). Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876–1902.

Balleine, B. W., Liljeholm, M., and Ostlund, S. B. (2009). The integrative function of the basal ganglia in instrumental conditioning. *Behavioural Brain Research*, 199:43–52.

Balleine, B. W., Morris, R. W., and Leung, B. K. (2015). Thalamocortical integration of instrumental learning and performance and their disintegration in addiction. *Brain Research*, 1628(Pt A):104–116.

Barron, A. B., Gurney, K. N., Meah, L. F. S., Vasilaki, E., and Marshall, J. A. R. (2015). Decision-making and action selection in insects: inspiration from vertebrate-based theories. *Frontiers in Behavioral Neuroscience*, 9:216.

Bath, D. E., Stowers, J. R., Hörmann, D., Poehlmann, A., Dickson, B. J., and Straw, A. D. (2014). FlyMAD: rapid thermogenetic control of neuronal activity in freely walking Drosophila. *Nature Methods*, 11(7):756–762.

Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., Shadlen, M. N., Latham, P. E., and Pouget, A. (2008). Probabilistic Population Codes for Bayesian Decision Making. *Neuron*, 60(6):1142–1152.

Berck, M. E., Khandelwal, A., Claus, L., Hernandez-Nunez, L., Si, G., Tabone, C. J., Li, F., Truman, J. W., Fetter, R. D., Louis, M., Samuel, A. D., and Cardona, A. (2016). The wiring diagram of a glomerular olfactory system. *eLife*, 5:e14859.

Berman, G. J., Choi, D. M., Bialek, W., and Shaevitz, J. W. (2014). Mapping the stereotyped behaviour of freely moving fruit flies. *Journal of The Royal Society Interface*, 11(99):20140672.

Berni, J., Pulver, S. R., Griffith, L. C., and Bate, M. (2012). Autonomous Circuitry for Substrate Exploration in Freely Moving Drosophila Larvae. *Current Biology*, 22(20):1861–1870.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.

Bitzer, S., Park, H., Blankenburg, F., and Kiebel, S. J. (2014). Perceptual decision making: drift-diffusion model is equivalent to a Bayesian model. *Frontiers in Human Neuroscience*, 8:102.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4):700–765.

Booker, R. and Quinn, W. G. (1981). Conditioning of leg position in normal and mutant Drosophila. *Proceedings of the National Academy of Sciences*, 78(6):3940–3944.

Brand, A. H. and Perrimon, N. (1993). Targeted gene expression as a means of altering cell fates and generating dominant phenotypes. *Development (Cambridge, England)*, 118(2):401–15.

Branson, K., Robie, A. A., Bender, J., Perona, P., and Dickinson, M. H. (2009). High-throughput ethomics in large groups of Drosophila. *Nature Methods*, 6(6):451–457.

Braubach, O. R., Wood, H.-D., Gadbois, S., Fine, A., and Croll, R. P. (2009). Olfactory conditioning in the zebrafish (Danio rerio). *Behavioural Brain Research*, 198(1):190–198.

Brembs, B. (2003). Operant conditioning in invertebrates. *Current Opinion in Neurobiology*, 13(6):710–717.

Brembs, B. (2011). Spontaneous decisions and operant conditioning in fruit flies. *Behavioural Processes*, 87(1):157–164.

Brembs, B., Lorenzetti, F. D., Reyes, F. D., Baxter, D. A., and Byrne, J. H. (2002). Operant Reward Learning in Aplysia: Neuronal Correlates and Mechanisms. *Science*, 296(5573):1706–1709.

Brembs, B. and Plendl, W. (2008). Double Dissociation of PKC and AC Manipulations on Operant and Classical Learning in Drosophila. *Current Biology*, 18(15):1168–1171.

Brown, J. S., Kalish, H. I., and Farber, I. E. (1951). Conditioned fear as revealed by magnitude of startle response to an auditory stimulus. *Journal of Experimental Psychology*, 41(5):317–328.

Brunton, B. W., Botvinick, M. M., and Brody, C. D. (2013). Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. *Science*, 340(6128):95–98.

Calhoun, A. J., Chalasani, S. H., and Sharpee, T. O. (2014). Maximally informative foraging by Caenorhabditis elegans. *eLife*, 3:e04220.

Campbell, R. A. A., Honegger, K. S., Qin, H., Li, W., Demir, E., and Turner, G. C. (2013). Imaging a Population Code for Odor Identity in the Drosophila Mushroom Body. *Journal of Neuroscience*, 33(25):10568–10581.

Caroni, P. (2015). Inhibitory microcircuit modules in hippocampal learning. *Current Opinion in Neurobiology*, 35:66–73.

Carreira-Rosario, A., Zarin, A. A., Clark, M. Q., Manning, L., Fetter, R. D., Cardona, A., and Doe, C. Q. (2018). MDN brain descending neurons coordinately activate backward and inhibit forward locomotion. *eLife*, 7:e38554.

Chen, Y.-J., Li, Y.-C., Huang, K.-N., and Young, M.-S. (2005). The Implementation of a Stand-alone Video Tracking and Analysis System for Animal Behavior Measurement in Morris Water Maze. In *IEEE Engineering in Medicine and Biology 27th Annual Conference*, pages 1766–1768. IEEE.

Chiuchisan, I. (2013). A New FPGA-based Real-Time Configurable System for Medical Image Processing. In *E-Health and Bioengineering Conference (EHB)*, pages 1–4. IEEE.

Cinotti, F., Fresno, V., Aklil, N., Coutureau, E., Girard, B., Marchand, A. R., and Khamassi, M. (2019). Dopamine blockade impairs the exploration-exploitation trade-off in rats. *Scientific Reports*, 9:6770.

Claridge-Chang, A., Roorda, R. D., Vrontou, E., Sjulson, L., Li, H., Hirsh, J., and Miesenböck, G. (2009). Writing Memories with Light-Addressable Reinforcement Circuitry. *Cell*, 139(2):405–415.

Clyne, J. D. and Miesenböck, G. (2008). Sex-Specific Control and Tuning of the Pattern Generator for Courtship Song in Drosophila. *Cell*, 133(2):354–363.

Cobb, M. and Domain, I. (2000). Olfactory coding in a simple system: adaptation in Drosophila larvae. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 267(1457):2119–2125.

Cognigni, P., Felsenberg, J., and Waddell, S. (2018). Do the right thing: neural network mechanisms of memory formation, expression and update in Drosophila. *Current Opinion in Neurobiology*, 49:51–58.

Cohen, J. D., McClure, S. M., and Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481):933–942.

Colomb, J. and Brembs, B. (2010). The biology of psychology. *Communicative & Integrative Biology*, 3(2):142–145.

Colomb, J. and Brembs, B. (2016). PKC in motorneurons underlies self-learning, a form of motor learning in Drosophila. *PeerJ*, 4:e1971.

Cong, L., Wang, Z., Chai, Y., Hang, W., Shang, C., Yang, W., Bai, L., Du, J., Wang, K., and Wen, Q. (2017). Rapid whole brain imaging of neural activity in freely behaving larval zebrafish (Danio rerio). *eLife*, 6:e28158.

Corbett, D. and Wise, R. A. (1980). Intracranial self-stimulation in relation to the ascending dopaminergic systems of the midbrain: A moveable electrode mapping study. *Brain Research*, 185(1):1–15.

Corrales-Carvajal, V. M., Faisal, A. A., and Ribeiro, C. (2016). Internal states drive nutrient homeostasis by modulating exploration-exploitation trade-off. *eLife*, 5:e19920.

Costa, V. D., Mitz, A. R., and Averbeck, B. B. (2019). Subcortical Substrates of Explore-Exploit Decisions in Primates. *Neuron*, 103(3):533–545.

Curcio, J. A. and Petty, C. C. (1951). The Near Infrared Absorption Spectrum of Liquid Water. *Journal of the Optical Society of America*, 41(5):302–304.

Dankert, H., Wang, L., Hoopfer, E. D., Anderson, D. J., and Perona, P. (2009). Automated monitoring and analysis of social behavior in Drosophila. *Nature Methods*, 6(4):297–303.

DasGupta, S., Ferreira, C. H., and Miesenböck, G. (2014). FoxP influences the speed and accuracy of a perceptual decision in Drosophila. *Science*, 344(6186):901–4.

Davis, R. L. (2005). Olfactory Memory Formation in Drosophila: From Molecular to Systems Neuroscience. *Annual Review of Neuroscience*, 28:275–302.

de Lafuente, V., Jazayeri, M., and Shadlen, M. N. (2015). Representation of Accumulating Evidence for a Decision in Two Parietal Areas. *Journal of Neuroscience*, 35(10):4306–4318.

Deneve, S. (2008). Bayesian Spiking Neurons I: Inference. *Neural Computation*, 20(1):91–117.

Denisov, G., Ohyama, T., Jovanic, T., and Zlatic, M. (2013). Model-based Detection and Analysis of Animal Behaviors using Signals Extracted by Automated Tracking. In *BIOSIGNALS*, pages 175–181.

Dickinson, A. (1981). Conditioning and Associative Learning. *British Medical Bulletin*, 37(2):165–168.

Donelson, N., Kim, E. Z., Slawson, J. B., Vecsey, C. G., Huber, R., and Griffith, L. C. (2012). High-Resolution Positional Tracking for Long-Term Analysis of Drosophila Sleep and Locomotion Using the “Tracker” Program. *PLoS ONE*, 7(5):e37250.

Dong, Y., Green, T., Saal, D., Marie, H., Neve, R., Nestler, E. J., and Malenka, R. C. (2006). CREB modulates excitability of nucleus accumbens neurons. *Nature Neuroscience*, 9(4):475–477.

Duffy, J. B. (2002). GAL4 System in Drosophila: A Fly Geneticist’s Swiss Army Knife. *genesis*, 34(1-2):1–15.

Eckhoff, P., Holmes, P., Law, C., Connolly, P. M., and Gold, J. I. (2008). On diffusion processes with variable drift rates as models for decision making during learning. *New Journal of Physics*, 10:015006.

Eichler, K., Li, F., Litwin-Kumar, A., Park, Y., Andrade, I., Schneider-Mizell, C. M., Saumweber, T., Huser, A., Eschbach, C., Gerber, B., Fetter, R. D., Truman, J. W., Priebe, C. E., Abbott, L. F., Thum, A. S., Zlatic, M., and Cardona, A. (2017). The complete connectome of a learning and memory centre in an insect brain. *Nature*, 548(7666):175–182.

Erlich, J. C., Bialek, M., and Brody, C. D. (2011). A Cortical Substrate for Memory-Guided Orienting in the Rat. *Neuron*, 72(2):330–343.

Eschbach, C., Cano, C., Haberkern, H., Schraut, K., Guan, C., Triphan, T., and Gerber, B. (2011). Associative learning between odorants and mechanosensory punishment in larval Drosophila. *The Journal of Experimental Biology*, 214(23):3897–3905.

Eschbach, C., Fushiki, A., Winding, M., Afonso, B., Andrade, I. V., Cocanougher, B. T., Eichler, K., Gepner, R., Si, G., Valdes-Aleman, J., Gershow, M., Sxe Jefferis, G., Truman, J. W., Fetter, R. D., Samuel, A., Cardona, A., and Zlatic, M. (2020a). Circuits for integrating learnt and innate valences in the fly brain. *bioRxiv*.

Eschbach, C., Fushiki, A., Winding, M., Schneider-Mizell, C. M., Shao, M., Arruda, R., Eichler, K., Valdes-Aleman, J., Ohyama, T., Thum, A. S., Gerber, B., Fetter, R. D., Truman, J. W., Litwin-Kumar, A., Cardona, A., and Zlatic, M. (2020b). Recurrent architecture for adaptive regulation of learning in the insect brain. *Nature Neuroscience*, 23:544–555.

Everitt, B. J., Giuliano, C., and Belin, D. (2018). Addictive behaviour in experimental animals: prospects for translation. *Philosophical Transactions of the Royal Society B*, 373(1742):20170027.

Fee, M. S. (2014). The role of efference copy in striatal learning. *Current Opinion in Neurobiology*, 25:194–200.

Fee, M. S. and Goldberg, J. H. (2011). A hypothesis for basal ganglia-dependent reinforcement learning in the songbird. *Neuroscience*, 198:152–170.

Feng, S., Holmes, P., Rorie, A., and Newsome, W. T. (2009). Can Monkeys Choose Optimally When Faced with Noisy Stimuli and Unequal Rewards? *PLoS Computational Biology*, 5(2):e1000284.

Fischer, J. A., Giniger, E., Maniatis, T., and Ptashne, M. (1988). GAL4 activates transcription in Drosophila. *Nature*, 332(6167):853–856.

Flavell, S. W., Pokala, N., Macosko, E. Z., Albrecht, D. R., Larsch, J., and Bargmann, C. I. (2013). Serotonin and the Neuropeptide PDF Initiate and Extend Opposing Behavioral States in C. elegans. *Cell*, 154(5):1023–1035.

Foley, B. R. and Marjoram, P. (2017). Sure enough: efficient Bayesian learning and choice. *Animal Cognition*, 20(5):867–880.

Friggi-Grelin, F., Coulom, H., Meller, M., Gomez, D., Hirsh, J., and Birman, S. (2003). Targeted Gene Expression in Drosophila Dopaminergic Cells Using Regulatory Sequences from Tyrosine Hydroxylase. *Journal of Neurobiology*, 54(4):618–627.

Fry, S. N., Rohrseitz, N., Straw, A. D., and Dickinson, M. H. (2008). TrackFly: Virtual reality for a behavioral system analysis in free-flying fruit flies. *Journal of Neuroscience Methods*, 171(1):110–117.

Funamizu, A., Kuhn, B., and Doya, K. (2016). Neural substrate of dynamic Bayesian inference in the cerebral cortex. *Nature Neuroscience*, 19(12):1682–1689.

Fushiki, A., Zwart, M. F., Kohsaka, H., Fetter, R. D., Cardona, A., and Nose, A. (2016). A circuit mechanism for the propagation of waves of muscle contraction in Drosophila. *eLife*, 5:e13253.

Ganguly, A., Qi, C., Bajaj, J., and Lee, D. (2020). Serotonin receptor 5-HT7 in Drosophila mushroom body neurons mediates larval appetitive olfactory learning. *Scientific Reports*, 10(1):21267.

Gepner, R., Mihovilovic Skanata, M., Bernat, N. M., Kaplow, M., and Gershow, M. (2015). Computations underlying Drosophila photo-taxis, odor-taxis, and multi-sensory integration. *eLife*, 4:e06229.

Gepner, R., Wolk, J., Wadekar, D. S., Dvali, S., and Gershow, M. (2018). Variance adaptation in navigational decision making. *eLife*, 7:e37945.

Gerber, B. and Hendel, T. (2006). Outcome expectations drive learned behaviour in larval Drosophila. *Proceedings of the Royal Society B*, 273(1604):2965–2968.

Gerber, B., Scherer, S., Neuser, K., Michels, B., Hendel, T., Stocker, R. F., and Heisenberg, M. (2004). Visual learning in individually assayed Drosophila larvae. *The Journal of Experimental Biology*, 207(1):179–188.

Gerhard, S., Andrade, I., Fetter, R. D., Cardona, A., and Schneider-Mizell, C. M. (2017). Conserved neural circuit structure across Drosophila larval development revealed by comparative connectomics. *eLife*, 6:e29089.

Gershow, M., Berck, M., Mathew, D., Luo, L., Kane, E. A., Carlson, J. R., and Samuel, A. D. T. (2012). Controlling airborne cues to study small animal navigation. *Nature Methods*, 9(3):290–296.

Gerstner, W., Sprekeler, H., and Deco, G. (2012). Theory and Simulation in Neuroscience. *Science*, 338(6103):60–65.

Giurfa, M. (2006). Associative Learning: The Instructive Function of Biogenic Amines. *Current Biology*, 16(20):R892–R895.

Gjorgjieva, J., Berni, J., Evers, J. F., and Eglen, S. J. (2013). Neural circuits for peristaltic wave propagation in crawling Drosophila larvae: analysis and modeling. *Frontiers in Computational Neuroscience*, 7:24.

Gold, J. I. and Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, 5(1):10–16.

Gold, J. I. and Shadlen, M. N. (2007). The Neural Basis of Decision Making. *Annual Review of Neuroscience*, 30(1):535–574.

Gomez-Marin, A. and Louis, M. (2012). Active sensation during orientation behavior in the Drosophila larva: more sense than luck. *Current Opinion in Neurobiology*, 22(2):208–215.

Gomez-Marin, A., Stephens, G. J., and Louis, M. (2011). Active sampling and decision making in Drosophila chemotaxis. *Nature Communications*, 2:441.

Gómez-Pinilla, F., Huie, J., Ying, Z., Ferguson, A., Crown, E., Baumbauer, K., Edgerton, V., and Grau, J. (2007). BDNF and learning: Evidence that instrumental training promotes learning within the spinal cord by up-regulating BDNF expression. *Neuroscience*, 148(4):893–906.

Gonzalez, R. C. and Woods, R. E. (2018). *Digital image processing, Ebook*. Pearson Education, Limited, fourth, gl edition.

Grau, J. W., Barstow, D. G., and Joynes, R. L. (1998). Instrumental Learning Within the Spinal Cord: I. Behavioral Properties. *Behavioral Neuroscience*, 112(6):1366–1386.

Graving, J. M. and Couzin, I. D. (2020). VAE-SNE: a deep generative model for simultaneous dimensionality reduction and clustering. *bioRxiv*.

Green, C. H., Burnet, B., and Connolly, K. J. (1983). Organization and patterns of inter- and intraspecific variation in the behaviour of Drosophila larvae. *Animal Behaviour*, 31(1):282–291.

Groschner, L. N., Chan Wah Hak, L., Bogacz, R., DasGupta, S., and Miesenböck, G. (2018). Dendritic Integration of Sensory Evidence in Perceptual Decision-Making. *Cell*, 173(4):894–905.

Groschner, L. N. and Miesenböck, G. (2019). Mechanisms of Sensory Discrimination: Insights from Drosophila Olfaction. *Annual Review of Biophysics*, 48(1):209–229.

Groszer, M., Keays, D. A., Deacon, R. M. J., de Bono, J. P., Prasad-Mulcare, S., Gaub, S., Baum, M. G., French, C. A., Nicod, J., Coventry, J. A., Enard, W., Fray, M., Brown, S. D. M., Nolan, P. M., Pääbo, S., Channon, K. M., Costa, R. M., Eilers, J., Ehret, G., Rawlins, J. N. P., and Fisher, S. E. (2008). Impaired Synaptic Plasticity and Motor Learning in Mice with a Point Mutation Implicated in Human Speech Deficits. *Current Biology*, 18(5):354–362.

Grover, D., Tower, J., and Tavaré, S. (2008). O fly, where art thou? *Journal of the Royal Society Interface*, 5(27):1181–1191.

Grueber, W. B., Jan, L. Y., and Jan, Y. N. (2002). Tiling of the Drosophila epidermis by multidendritic sensory neurons. *Development*, 129(12):2867–2878.

Gründemann, J. and Lüthi, A. (2015). Ensemble coding in amygdala circuits for associative learning. *Current Opinion in Neurobiology*, 35:200–206.

Gupta, S. and Gomez-Marin, A. (2019). A context-free grammar for Caenorhabditis elegans behavior. *bioRxiv*.

Haberkern, H., Basnak, M. A., Ahanonu, B., Schauder, D., Cohen, J. D., Bolstad, M., Bruns, C., and Jayaraman, V. (2019). Visually Guided Behavior and Optogenetically Induced Learning in Head-Fixed Flies Exploring a Virtual Landscape. *Current Biology*, 29(10):1647–1659.e8.

Haesler, S., Rochefort, C., Georgi, B., Licznerski, P., Osten, P., and Scharff, C. (2007). Incomplete and Inaccurate Vocal Imitation after Knockdown of FoxP2 in Songbird Basal Ganglia Nucleus Area X. *PLoS Biology*, 5(12):e321.

Hagan, M. T., Demuth, H. B., Beale, M. H., and De Jesús, O. (2014). *Neural Network Design*. Martin Hagan, 2nd edition.

Hamada, F. N., Rosenzweig, M., Kang, K., Pulver, S. R., Ghezzi, A., Jegla, T. J., and Garrity, P. A. (2008). An internal thermal sensor controlling temperature preference in Drosophila. *Nature*, 454(7201):217–220.

Hanks, T. D. and Summerfield, C. (2017). Perceptual Decision Making in Rodents, Monkeys, and Humans. *Neuron*, 93(1):15–31.

Hawkins, R. D. and Byrne, J. H. (2015). Associative Learning in Invertebrates. *Cold Spring Harbor Perspectives in Biology*, 7(5):a021709.

Hazelrigg, T., Levis, R., and Rubin, G. M. (1984). Transformation of white locus DNA in Drosophila: Dosage compensation, zeste interaction, and position effects. *Cell*, 36(2):469–481.

Heckscher, E. S., Lockery, S. R., and Doe, C. Q. (2012). Characterization of Drosophila Larval Crawling at the Level of Organism, Segment, and Somatic Body Wall Musculature. *Journal of Neuroscience*, 32(36):12460–12471.

Heisenberg, M. (2003). Mushroom body memoir: from maps to models. *Nature Reviews Neuroscience*, 4(4):266–75.

Heisenberg, M., Borst, A., Wagner, S., and Byers, D. (1985). Drosophila Mushroom Body Mutants are Deficient in Olfactory Learning. *Journal of Neurogenetics*, 2(1):1–30.

Hendel, T., Michels, B., Neuser, K., Schipanski, A., Kaun, K., Sokolowski, M. B., Marohn, F., Michel, R., Heisenberg, M., and Gerber, B. (2005). The carrot, not the stick: Appetitive rather than aversive gustatory stimuli support associative olfactory learning in individually assayed Drosophila larvae. *Journal of Comparative Physiology A*, 191(3):265–279.

Hernandez-Nunez, L., Belina, J., Klein, M., Si, G., Claus, L., Carlson, J. R., and Samuel, A. D. (2015). Reverse-correlation analysis of navigation dynamics in Drosophila larva using optogenetics. *eLife*, 4:e06225.

Honegger, K. S., Campbell, R. A. A., and Turner, G. C. (2011). Cellular-Resolution Population Imaging Reveals Robust Sparse Coding in the Drosophila Mushroom Body. *Journal of Neuroscience*, 31(33):11772–11785.

Honjo, K. and Furukubo-Tokunaga, K. (2005). Induction of cAMP Response Element-Binding Protein-Dependent Medium-Term Memory by Appetitive Gustatory Reinforcement in Drosophila Larvae. *Journal of Neuroscience*, 25(35):7905–7913.

Honjo, K. and Furukubo-Tokunaga, K. (2009). Distinctive Neuronal Networks and Biochemical Pathways for Appetitive and Aversive Memory in Drosophila Larvae. *Journal of Neuroscience*, 29(3):852–862.

Horridge, G. A. (1962). Learning of leg position by the ventral nerve cord in headless insects. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 157(966):33–52.

Hoyle, G. (1979). Mechanisms of simple motor learning. *Trends in Neurosciences*, 2:153–155.

Hu, C., Petersen, M., Hoyer, N., Spitzweck, B., Tenedini, F., Wang, D., Gruschka, A., Burchardt, L. S., Szpotowicz, E., Schweizer, M., Guntur, A. R., Yang, C.-H. H., and Soba, P. (2017). Sensory integration and neuromodulatory feedback facilitate Drosophila mechanonociceptive behavior. *Nature Neuroscience*, 20(8):1085–1095.

Huang, K.-M., Cosman, P., and Schafer, W. R. (2006). Machine vision based detection of omega bends and reversals in C. elegans. *Journal of Neuroscience Methods*, 158(2):323–336.

Humberg, T.-H., Bruegger, P., Afonso, B., Zlatic, M., Truman, J. W., Gershow, M., Samuel, A., and Sprecher, S. G. (2018). Dedicated photoreceptor pathways in Drosophila larvae mediate navigation by processing either spatial or temporal cues. *Nature Communications*, 9(1):1260.

Huser, A., Eschment, M., Güllü, N., Collins, K. A. N., Böpple, K., Pankevych, L., Rolsing, E., and Thum, A. S. (2017). Anatomy and behavioral function of serotonin receptors in Drosophila melanogaster larvae. *PLoS ONE*, 12(8):e0181865.

Huser, A., Rohwedder, A., Apostolopoulou, A. A., Widmann, A., Pfitzenmaier, J. E., Maiolo, E. M., Selcho, M., Pauls, D., von Essen, A., Gupta, T., Sprecher, S. G., Birman, S., Riemensperger, T., Stocker, R. F., and Thum, A. S. (2012). The Serotonergic Central Nervous System of the Drosophila Larva: Anatomy and Behavioral Function. *PLoS ONE*, 7(10):e47518.

Hwang, R. Y., Zhong, L., Xu, Y., Johnson, T., Zhang, F., Deisseroth, K., and Tracey, W. D. (2007). Nociceptive Neurons Protect Drosophila Larvae from Parasitoid Wasps. *Current Biology*, 17(24):2105–2116.

Jefferis, G. S. X. E., Potter, C. J., Chan, A. M., Marin, E. C., Rohlfing, T., Maurer, Jr., C. R., and Luo, L. (2007). Comprehensive Maps of Drosophila Higher Olfactory Centers: Spatially Segregated Fruit and Pheromone Representation. *Cell*, 128(6):1187–1203.

Jenett, A., Rubin, G. M., Ngo, T.-T., Shepherd, D., Murphy, C., Dionne, H., Pfeiffer, B. D., Cavallaro, A., Hall, D., Jeter, J., Iyer, N., Fetter, D., Hausenfluck, J. H., Peng, H., Trautman, E. T., Svirskas, R. R., Myers, E. W., Iwinski, Z. R., Aso, Y., DePasquale, G. M., Enos, A., Hulamm, P., Lam, S. C. B., Li, H.-H., Laverty, T. R., Long, F., Qu, L., Murphy, S. D., Rokicki, K., Safford, T., Shaw, K., Simpson, J. H., Sowell, A., Tae, S., Yu, Y., and Zugates, C. T. (2012). A GAL4-Driver Line Resource for Drosophila Neurobiology. *Cell Reports*, 2(4):991–1001.

Jin, X. and Costa, R. M. (2010). Start/stop signals emerge in nigrostriatal circuits during sequence learning. *Nature*, 466(7305):457–462.

Joel, D. (2006). The signal attenuation rat model of obsessive-compulsive disorder: a review. *Psychopharmacology*, 186(4):487–503.

Johnson, O., Becnel, J., and Nichols, C. (2011). Serotonin receptor activity is necessary for olfactory learning and memory in Drosophila melanogaster. *Neuroscience*, 192:372–381.

Jones, S. V., Heldt, S. A., Davis, M., and Ressler, K. J. (2005). Olfactory-Mediated Fear Conditioning in Mice: Simultaneous Measurements of Fear-Potentiated Startle and Freezing. *Behavioral Neuroscience*, 119(1):329–335.

Jovanic, T., Schneider-Mizell, C. M., Shao, M., Masson, J.-B., Denisov, G., Fetter, R. D., Mensh, B. D., Truman, J. W., Cardona, A., and Zlatic, M. (2016). Competitive Disinhibition Mediates Behavioral Choice and Sequences in Drosophila. *Cell*, 167(3):858–870.e19.

Jovanic, T., Winding, M., Cardona, A., Truman, J. W., Gershow, M., and Zlatic, M. (2019). Neural Substrates of Drosophila Larval Anemotaxis. *Current Biology*, 29(4):554–566.

Joynes, R. L., Janjua, K., and Grau, J. W. (2004). Instrumental learning within the spinal cord: VI: The NMDA receptor antagonist, AP5, disrupts the acquisition and maintenance of an acquired flexion response. *Behavioural Brain Research*, 154(2):431–438.

Kabra, M., Robie, A. A., Rivera-Alba, M., Branson, S., and Branson, K. (2013). JAABA: interactive machine learning for automatic annotation of animal behavior. *Nature Methods*, 10(1):64–67.

Kane, E. A., Gershow, M., Afonso, B., Larderet, I., Klein, M., Carter, A. R., de Bivort, B. L., Sprecher, S. G., and Samuel, A. D. T. (2013). Sensorimotor structure of Drosophila larva phototaxis. *Proceedings of the National Academy of Sciences*, 110(40):E3868–E3877.

Karagyozov, D., Mihovilovic Skanata, M., Lesar, A., and Gershow, M. (2018). Recording Neural Activity in Unrestrained Animals with Three-Dimensional Tracking Two-Photon Microscopy. *Cell Reports*, 25(5):1371–1383.e10.

Katsov, A. Y. and Clandinin, T. R. (2008). Motion Processing Streams in Drosophila Are Behaviorally Specialized. *Neuron*, 59(2):322–335.

Katz, K. and Naug, D. (2015). Energetic state regulates the exploration–exploitation trade-off in honeybees. *Behavioral Ecology*, 26(4):1045–1050.

Keasar, T., Rashkovich, E., Cohen, D., and Shmida, A. (2002). Bees in two-armed bandit situations: foraging choices and possible decision mechanisms. *Behavioral Ecology*, 13(6):757–765.

Kepecs, A., Uchida, N., Zariwala, H. A., and Mainen, Z. F. (2008). Neural correlates, computation and behavioural impact of decision confidence. *Nature*, 455(7210):227–231.

Kernan, M., Cowan, D., and Zuker, C. (1994). Genetic Dissection of Mechanosensory Transduction: Mechanoreception-Defective Mutations of Drosophila. *Neuron*, 12(6):1195–1206.

Khurana, S., Robinson, B. G., Wang, Z., Shropshire, W. C., Zhong, A. C., Garcia, L. E., Corpuz, J., Chow, J., Hatch, M. M., Precise, E. F., Cady, A., Godinez, R. M., Pulpanyawong, T., Nguyen, A. T., Li, W.-k., Seiter, M., Jahanian, K., Sun, J. C., Shah, R., Rajani, S., Chen, W. Y., Ray, S., Ryazanova, N. V., Wakou, D., Prabhu, R. K., and Atkinson, N. S. (2012). Olfactory Conditioning in the Third Instar Larvae of Drosophila melanogaster Using Heat Shock Reinforcement. *Behavior Genetics*, 42(1):151–161.

Kira, S., Yang, T., and Shadlen, M. N. (2015). A Neural Implementation of Wald's Sequential Probability Ratio Test. *Neuron*, 85(4):861–873.

Kitamoto, T. (2001). Conditional Modification of Behavior in Drosophila by Targeted Expression of a Temperature-Sensitive shibire Allele in Defined Neurons. *Journal of Neurobiology*, 47(2):81–92.

Klapoetke, N. C., Murata, Y., Kim, S. S., Pulver, S. R., Birdsey-Benson, A., Cho, Y. K., Morimoto, T. K., Chuong, A. S., Carpenter, E. J., Tian, Z., Wang, J., Xie, Y., Yan, Z., Zhang, Y., Chow, B. Y., Surek, B., Melkonian, M., Jayaraman, V., Constantine-Paton, M., Wong, G. K.-S., and Boyden, E. S. (2014). Independent optical excitation of distinct neural populations. *Nature Methods*, 11(3):338–346.

Klein, K. T., Croteau-Chonka, E. C., Narayan, L., Winding, M., Masson, J.-B., and Zlatic, M. (2021). Serotonergic Neurons Mediate Operant Conditioning in Drosophila Larvae. *bioRxiv*.

Klein, M., Afonso, B., Vonner, A. J., Hernandez-Nunez, L., Berck, M., Tabone, C. J., Kane, E. A., Pieribone, V. A., Nitabach, M. N., Cardona, A., Zlatic, M., Sprecher, S. G., Gershow, M., Garrity, P. A., and Samuel, A. D. T. (2015). Sensory determinants of behavioral dynamics in Drosophila thermotaxis. *Proceedings of the National Academy of Sciences*, 112(2):E220–E229.

Klibaite, U., Berman, G. J., Cande, J., Stern, D. L., and Shaevitz, J. W. (2017). An unsupervised method for quantifying the behavior of paired animals. *Physical Biology*, 14(1):015006.

Knill, D. C. and Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12):712–719.

Kolling, N., Behrens, T. E. J., Mars, R. B., and Rushworth, M. F. S. (2012). Neural Mechanisms of Foraging. *Science*, 336(6077):95–98.

Körding, K. (2007). Decision Theory: What "Should" the Nervous System Do? *Science*, 318(5850):606–610.

Körding, K. P. and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971):244–247.

Krajbich, I. (2019). Accounting for attention in sequential sampling models of decision making. *Current Opinion in Psychology*, 29:6–11.

Kriegeskorte, N. and Douglas, P. K. (2018). Cognitive computational neuroscience. *Nature Neuroscience*, 21(9):1148–1160.

Krynitsky, J., Legaria, A. A., Pai, J. J., Garmendia-Cedillos, M., Salem, G., Pohida, T., and Kravitz, A. V. (2020). Rodent Arena Tracker (Rat): A Machine Vision Rodent Tracking Camera and Closed Loop Control System. *eNeuro*, 7(3):ENEURO.0485–19.2020.

Kudow, N., Miura, D., Schleyer, M., Toshima, N., Gerber, B., and Tanimura, T. (2017). Preference for and learning of amino acids in larval Drosophila. *Biology Open*, 6(3):365–369.

Lahiri, S., Shen, K., Klein, M., Tang, A., Kane, E., Gershow, M., Garrity, P., and Samuel, A. D. T. (2011). Two Alternating Motor Programs Drive Navigation in Drosophila Larva. *PLoS ONE*, 6(8):e23180.

Lai, C. S. L., Fisher, S. E., Hurst, J. A., Vargha-Khadem, F., and Monaco, A. P. (2001). A forkhead-domain gene is mutated in a severe speech and language disorder. *Nature*, 413(6855):519–523.

Larderet, I., Fritsch, P. M. J., Gendre, N., Neagu-Maier, G. L., Fetter, R. D., Schneider-Mizell, C. M., Truman, J. W., Zlatic, M., Cardona, A., and Sprecher, S. G. (2017). Organization of the Drosophila larval visual circuit. *eLife*, 6:e28387.

Lee, T. and Luo, L. (1999). Mosaic analysis with a repressible cell marker for studies of gene function in neuronal morphogenesis. *Neuron*, 22(3):451–461.

Li, H., Chaney, S., Forte, M., and Hirsh, J. (2000). Ectopic G-protein expression in dopamine and serotonin neurons blocks cocaine sensitization in Drosophila melanogaster. *Current Biology*, 10(4):211–214.

Li, H.-H., Kroll, J. R., Lennox, S. M., Ogundeyi, O., Jeter, J., Depasquale, G., and Truman, J. W. (2014). A GAL4 driver resource for developmental and behavioral studies on the larval CNS of Drosophila. *Cell Reports*, 8(3):897–908.

Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. *Nature*, 532(7600):459–464.

Li, Y. and Krishnamurthy, K. (2015). Is There a General Role for the Monkey Oculomotor System in Perceptual Decision-Making? *The Journal of Neuroscience*, 35(27):9783–9785.

Li, Y., Yao, Q., Tian, B., and Xu, W. (2011). Fast Double-Parallel image processing based on FPGA. In *IEEE International Conference on Vehicular Electronics and Safety*, pages 97–102. IEEE.

Lima, S. Q. and Miesenböck, G. (2005). Remote Control of Behavior through Genetically Targeted Photostimulation of Neurons. *Cell*, 121(1):141–152.

Lin, S., Owald, D., Chandra, V., Talbot, C., Huetteroth, W., and Waddell, S. (2014). Neural correlates of water reward in thirsty Drosophila. *Nature Neuroscience*, 17(11):1536–1542.

Liu, C., Plaçais, P.-Y., Yamagata, N., Pfeiffer, B. D., Aso, Y., Friedrich, A. B., Siwanowicz, I., Rubin, G. M., Preat, T., and Tanimoto, H. (2012). A subset of dopamine neurons signals reward for odour memory in Drosophila. *Nature*, 488(7412):512–516.

Lorenzetti, F. D., Baxter, D. A., and Byrne, J. H. (2008). Molecular Mechanisms Underlying a Cellular Analog of Operant Reward Learning. *Neuron*, 59(5):815–828.

Louis, M., Huber, T., Benton, R., Sakmar, T. P., and Vosshall, L. B. (2008). Bilateral olfactory sensory input enhances chemotaxis behavior. *Nature Neuroscience*, 11(2):187–199.

Lovell, J. M., Mylius, J., Scheich, H., and Brosch, M. (2015). Stimulation of the Dopaminergic Midbrain as a Behavioral Reward in Instrumentally Conditioned Monkeys. *Brain Stimulation*, 8(5):868–874.

Lovinger, D. M. (2010). Neurotransmitter roles in synaptic modulation, plasticity and learning in the dorsal striatum. *Neuropharmacology*, 58(7):951–961.

Luan, H., Peabody, N. C., Vinson, C. R., and White, B. H. (2006). Refined Spatial Manipulation of Neuronal Function by Combinatorial Restriction of Transgene Expression. *Neuron*, 52(3):425–436.

Lundell, M. J. and Hirsh, J. (1994). Temporal and Spatial Development of Serotonin and Dopamine Neurons in the Drosophila CNS. *Developmental Biology*, 165(2):385–396.

Luo, L., Callaway, E. M., and Svoboda, K. (2008). Genetic Dissection of Neural Circuits. *Neuron*, 57(5):634–60.

Luo, L., Gershow, M., Rosenzweig, M., Kang, K., Fang-Yen, C., Garrity, P. A., and Samuel, A. D. T. (2010). Navigational Decision Making in Drosophila Thermotaxis. *Journal of Neuroscience*, 30(12):4261–4272.

Luxem, K., Fuhrmann, F., Kürsch, J., Remy, S., and Bauer, P. (2020). Identifying behavioral structure from deep variational embeddings of animal motion. *bioRxiv*.

Ma, W. J., Beck, J. M., and Pouget, A. (2008). Spiking networks for Bayesian inference and choice. *Current Opinion in Neurobiology*, 18(2):217–222.

Ma, W. J. and Jazayeri, M. (2014). Neural Coding of Uncertainty and Probability. *Annual Review of Neuroscience*, 37(1):205–220.

Marbach, F. and Zador, A. M. (2016). A self-initiated two-alternative forced choice paradigm for head-fixed mice. *bioRxiv*.

Marin, E. C., Jefferis, G. S. X. E., Komiyama, T., Zhu, H., and Luo, L. (2002). Representation of the Glomerular Olfactory Map in the Drosophila Brain. *Cell*, 109(2):243–255.

Masson, J.-B., Laurent, F., Cardona, A., Barré, C., Skatchkovsky, N., Zlatic, M., and Jovanic, T. (2020). Identifying neural substrates of competitive interactions and sequence transitions during mechanosensory responses in Drosophila. *PLoS Genetics*, 16(2):e1008589.

Masson, J.-B., Voisinne, G., Wong-Ng, J., Celani, A., and Vergassola, M. (2012). Noninvasive inference of the molecular chemotactic response using bacterial trajectories. *Proceedings of the National Academy of Sciences*, 109(5):1802–1807.

Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience*, 21(9):1281–1289.

McNamara, J. M., Green, R. F., and Olsson, O. (2006). Bayes' theorem and its applications in animal behaviour. *Oikos*, 112(2):243–251.

Mendoza, E., Colomb, J., Rybak, J., Pflüger, H. J., Zars, T., Scharff, C., and Brembs, B. (2014). Drosophila FoxP Mutants are Deficient in Operant Self-Learning. *PLoS ONE*, 9(6):e100648.

Meneses, A. and Liy-Salmeron, G. (2012). Serotonin and emotion, learning and memory. *Reviews in the Neurosciences*, 23(5-6):543–553.

Mirat, O., Sternberg, J. R., Severi, K. E., and Wyart, C. (2013). ZebraZoom: an automated program for high-throughput behavioral analysis and categorization. *Frontiers in Neural Circuits*, 7:107.

Mischiati, M., Lin, H.-T., Herold, P., Imler, E., Olberg, R., and Leonardo, A. (2015). Internal models direct dragonfly interception steering. *Nature*, 517(7534):333–338.

Młynarski, W. F. and Hermundstad, A. M. (2018). Adaptive coding for dynamic sensory inference. *eLife*, 7:e32055.

Morimoto, J. (2019). Foraging decisions as multi-armed bandit problems: Applying reinforcement learning algorithms to foraging data. *Journal of Theoretical Biology*, 467:48–56.

Mu, Y., Bennett, D. V., Rubinov, M., Narayan, S., Yang, C.-T., Tanimoto, M., Mensh, B. D., Looger, L. L., and Ahrens, M. B. (2019). Glia Accumulate Evidence that Actions Are Futile and Suppress Unsuccessful Behavior. *Cell*, 178(1):27–43.

Murphy, K. P. (2007). Conjugate Bayesian Analysis of the Gaussian Distribution. *Def*, 1(7):1–29.

Nargeot, R., Baxter, D. A., and Byrne, J. H. (1997). Contingent-Dependent Enhancement of Rhythmic Motor Patterns: An In Vitro Analog of Operant Conditioning. *The Journal of Neuroscience*, 17(21):8093–8105.

Nargeot, R., Le Bon-Jego, M., and Simmers, J. (2009). Cellular and Network Mechanisms of Operant Learning-Induced Compulsive Behavior in Aplysia. *Current Biology*, 19(12):975–984.

Naug, D. and Arathi, H. S. (2007). Sampling and decision rules used by honey bees in a foraging arena. *Animal Cognition*, 10(2):117–124.

Nern, A., Pfeiffer, B. D., and Rubin, G. M. (2015). Optimized tools for multicolor stochastic labeling reveal diverse stereotyped cell arrangements in the fly visual system. *Proceedings of the National Academy of Sciences*, 112(22):E2967–E2976.

Neuser, K., Husse, J., Stock, P., and Gerber, B. (2005). Appetitive olfactory learning in Drosophila larvae: effects of repetition, reward strength, age, gender, assay type and memory span. *Animal Behaviour*, 69(4):891–898.

Neuser, K., Triphan, T., Mronz, M., Poeck, B., and Strauss, R. (2008). Analysis of a spatial orientation memory in Drosophila. *Nature*, 453(7199):1244–1247.

Niewalda, T., Singhal, N., Fiala, A., Saumweber, T., Wegener, S., and Gerber, B. (2008). Salt Processing in Larval Drosophila: Choice, Feeding, and Learning Shift from Appetitive to Aversive in a Concentration-Dependent Way. *Chemical Senses*, 33(8):685–692.

Nottebohm, F. (1991). Reassessing the mechanisms and origins of vocal learning in birds. *Trends in Neurosciences*, 14(5):206–211.

Nuwal, N., Stock, P., Hiemeyer, J., Schmid, B., Fiala, A., and Buchner, E. (2012). Avoidance of Heat and Attraction to Optogenetically Induced Sugar Sensation as Operant Behavior in Adult Drosophila. *Journal of Neurogenetics*, 26(3-4):298–305.

Ohyama, T., Jovanic, T., Denisov, G., Dang, T. C., Hoffmann, D., Kerr, R. A., and Zlatic, M. (2013). High-Throughput Analysis of Stimulus-Evoked Behaviors in Drosophila Larva Reveals Multiple Modality-Specific Escape Strategies. *PLoS ONE*, 8(8):e71706.

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., and Zlatic, M. (2015). A multilevel multimodal circuit enhances action selection in Drosophila. *Nature*, 520(7549):633–639.

Olds, J. and Milner, P. (1954). Positive reinforcement produced by electrical stimulation of septal area and other regions of rat brain. *Journal of Comparative and Physiological Psychology*, 47(6):419–427.

Otazu, G. H., Tai, L. H., Yang, Y., and Zador, A. M. (2009). Engaging in an auditory task suppresses responses in auditory cortex. *Nature Neuroscience*, 12(5):646–654.

Owald, D., Felsenberg, J., Talbot, C. B., Das, G., Perisse, E., Huetteroth, W., and Waddell, S. (2015). Activity of Defined Mushroom Body Output Neurons Underlies Learned Olfactory Behavior in Drosophila. *Neuron*, 86(2):417–427.

Owald, D. and Waddell, S. (2015). Olfactory learning skews mushroom body output pathways to steer behavioral choice in Drosophila. *Current Opinion in Neurobiology*, 35:178–184.

Park, J., Lee, S. B., Lee, S., Kim, Y., Song, S., Kim, S., Bae, E., Kim, J., Shong, M., Kim, J.-M., and Chung, J. (2006). Mitochondrial dysfunction in Drosophila PINK1 mutants is complemented by parkin. *Nature*, 441(7097):1157–1161.

Pavlov, I. P. (1927). *Conditioned reflexes: an investigation of the physiological activity of the cerebral cortex*. Oxford University Press, Oxford.

Perisse, E., Owald, D., Barnstedt, O., Talbot, C. B., Huetteroth, W., and Waddell, S. (2016). Aversive Learning and Appetitive Motivation Toggle Feed-Forward Inhibition in the Drosophila Mushroom Body. *Neuron*, 90(5):1086–1099.

Pfeiffer, B. D., Jenett, A., Hammonds, A. S., Ngo, T.-T. B., Misra, S., Murphy, C., Scully, A., Carlson, J. W., Wan, K. H., Laverty, T. R., Mungall, C., Svirskas, R., Kadonaga, J. T., Doe, C. Q., Eisen, M. B., Celniker, S. E., and Rubin, G. M. (2008). Tools for neuroanatomy and neurogenetics in Drosophila. *Proceedings of the National Academy of Sciences*, 105(28):9715–9720.

Pfeiffer, B. D., Ngo, T.-T. B., Hibbard, K. L., Murphy, C., Jenett, A., Truman, J. W., and Rubin, G. M. (2010). Refinement of Tools for Targeted Gene Expression in Drosophila. *Genetics*, 186(2):735–755.

Plaçais, P. Y., Trannoy, S., Friedrich, A. B., Tanimoto, H., and Preat, T. (2013). Two Pairs of Mushroom Body Efferent Neurons are Required for Appetitive Long-Term Memory Retrieval in Drosophila. *Cell Reports*, 5(3):769–780.

Pouget, A., Beck, J. M., Ma, W. J., and Latham, P. E. (2013). Probabilistic brains: knowns and unknowns. *Nature Neuroscience*, 16(9):1170–1178.

Rahnev, D. (2019). The Bayesian brain: What is it and do humans have it? *Behavioral and Brain Sciences*, 42:e238.

Rahnev, D. and Denison, R. N. (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences*, 41:e223.

Reddy, G., Desban, L., Tanaka, H., Roussel, J., Mirat, O., and Wyart, C. (2020). A lexical approach for identifying behavioral action sequences. *bioRxiv*.

Redgrave, P., Vautrelle, N., and Reynolds, J. (2011). Functional properties of the basal ganglia's re-entrant loop architecture: selection and reinforcement. *Neuroscience*, 198:138–151.

Reid, C. R., MacDonald, H., Mann, R. P., Marshall, J. A. R., Latty, T., and Garnier, S. (2016). Decision-making without a brain: how an amoeboid organism solves the two-armed bandit. *Journal of The Royal Society Interface*, 13(119):20160030.

Reitich-Stolero, T., Aberg, K. C., and Paz, R. (2019). Re-exploring Mechanisms of Exploration. *Neuron*, 103(3):360–363.

Rescorla, R. A. (1988). Behavioral Studies of Pavlovian Conditioning. *Annual Review of Neuroscience*, 11:329–352.

Reynolds, J. N. and Wickens, J. R. (2002). Dopamine-dependent plasticity of corticostriatal synapses. *Neural Networks*, 15(4-6):507–521.

Rich, D., Cazettes, F., Wang, Y., Peña, J. L., and Fischer, B. J. (2015). Neural representation of probabilities for Bayesian inference. *Journal of Computational Neuroscience*, 38(2):315–323.

Riedl, J. and Louis, M. (2012). Behavioral Neuroscience: Crawling Is a No-Brainer for Fruit Fly Larvae. *Current Biology*, 22(20):R867–R869.

Robertson, J. L., Tsubouchi, A., and Tracey, W. D. (2013). Larval Defense against Attack from Parasitoid Wasps Requires Nociceptive Neurons. *PLoS ONE*, 8(10):e78704.

Robie, A. A., Hirokawa, J., Edwards, A. W., Umayam, L. A., Lee, A., Phillips, M. L., Card, G. M., Korff, W., Rubin, G. M., Simpson, J. H., Reiser, M. B., and Branson, K. (2017). Mapping the Neural Substrates of Behavior. *Cell*, 170(2):393–406.

Rohwedder, A., Pfitzenmaier, J. E., Ramsperger, N., Apostolopoulou, A. A., Widmann, A., and Thum, A. S. (2012). Nutritional Value-Dependent and Nutritional Value-Independent Effects on Drosophila melanogaster Larval Behavior. *Chemical Senses*, 37(8):711–721.

Rohwedder, A., Wenz, N. L., Stehle, B., Huser, A., Yamagata, N., Zlatic, M., Truman, J. W., Tanimoto, H., Saumweber, T., Gerber, B., and Thum, A. S. (2016). Four Individually Identified Paired Dopamine Neurons Signal Reward in Larval Drosophila. *Current Biology*, 26(5):661–669.

Romo, R., Hernández, A., Zainos, A., Lemus, L., and Brody, C. D. (2002). Neuronal correlates of decision-making in secondary somatosensory cortex. *Nature Neuroscience*, 5(11):1217–1225.

Rosenegger, D. and Lukowiak, K. (2010). The participation of NMDA receptors, PKC, and MAPK in the formation of memory following operant conditioning in Lymnaea. *Molecular Brain*, 3:24.

Roy, B., Singh, A. P., Shetty, C., Chaudhary, V., North, A., Landgraf, M., VijayRaghavan, K., and Rodrigues, V. (2007). Metamorphosis of an identified serotonergic neuron in the Drosophila olfactory system. *Neural Development*, 2:20.

Saumweber, T., Husse, J., and Gerber, B. (2011). Innate Attractiveness and Associative Learnability of Odors Can Be Dissociated in Larval Drosophila. *Chemical Senses*, 36(3):223–235.

Saumweber, T., Rohwedder, A., Schleyer, M., Eichler, K., Chen, Y.-c., Aso, Y., Cardona, A., Eschbach, C., Kobler, O., Voigt, A., Durairaja, A., Mancini, N., Zlatic, M., Truman, J. W., Thum, A. S., and Gerber, B. (2018). Functional architecture of reward learning in mushroom body extrinsic neurons of larval Drosophila. *Nature Communications*, 9(1):1104.

Sawin-McCormack, E. P., Sokolowski, M. B., and Campos, A. R. (1995). Characterization and Genetic Analysis of Drosophila Melanogaster Photobehavior During Larval Development. *Journal of Neurogenetics*, 10(2):119–135.

Scherer, S., Stocker, R. F., and Gerber, B. (2003). Olfactory Learning in Individually Assayed Drosophila Larvae. *Learning & Memory*, 10(3):217–225.

Schipanski, A., Yarali, A., Niewalda, T., and Gerber, B. (2008). Behavioral Analyses of Sugar Processing in Choice, Feeding, and Learning in Larval Drosophila. *Chemical Senses*, 33(6):563–573.

Schlegel, P., Texada, M. J., Miroschnikow, A., Schoofs, A., Hückesfeld, S., Peters, M., Schneider-Mizell, C. M., Lacin, H., Li, F., Fetter, R. D., Truman, J. W., Cardona, A., and Pankratz, M. J. (2016). Synaptic transmission parallels neuromodulation in a central food-intake circuit. *eLife*, 5:e16799.

Schleyer, M., Miura, D., Tanimura, T., and Gerber, B. (2015). Learning the specific quality of taste reinforcement in larval Drosophila. *eLife*, 4:e04711.

Schleyer, M., Saumweber, T., Nahrendorf, W., Fischer, B., von Alpen, D., Pauls, D., Thum, A., and Gerber, B. (2011). A behavior-based circuit model of how outcome expectations organize learned behavior in larval Drosophila. *Learning & Memory*, 18(10):639–653.

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., and Cardona, A. (2016). Quantitative neuroanatomy for connectomics in Drosophila. *eLife*, 5:e12059.

Schroll, C., Riemensperger, T., Bucher, D., Ehmer, J., Völler, T., Erbguth, K., Gerber, B., Hendel, T., Nagel, G., Buchner, E., and Fiala, A. (2006). Light-Induced Activation of Distinct Modulatory Neurons Triggers Appetitive or Aversive Learning in Drosophila Larvae. *Current Biology*, 16(17):1741–1747.

Schulze, A., Gomez-Marin, A., Rajendran, V. G., Lott, G., Musy, M., Ahammad, P., Deogade, A., Sharpe, J., Riedl, J., Jarriault, D., Trautman, E. T., Werner, C., Venkadesan, M., Druckmann, S., Jayaraman, V., and Louis, M. (2015). Dynamical feature extraction at the sensory periphery guides chemotaxis. *eLife*, 4:e06694.

Schwaerzel, M., Monastirioti, M., Scholz, H., Friggi-Grelin, F., Birman, S., and Heisenberg, M. (2003). Dopamine and Octopamine Differentiate between Aversive and Appetitive Olfactory Memories in Drosophila. *The Journal of Neuroscience*, 23(33):10495–10502.

Séjourné, J., Plaçais, P.-Y., Aso, Y., Siwanowicz, I., Trannoy, S., Thoma, V., Tedjakumala, S. R., Rubin, G. M., Tchénio, P., Ito, K., Isabel, G., Tanimoto, H., and Preat, T. (2011). Mushroom body efferent neurons responsible for aversive olfactory memory retrieval in Drosophila. *Nature Neuroscience*, 14(7):903–910.

Selcho, M., Pauls, D., Han, K.-A., Stocker, R. F., and Thum, A. S. (2009). The Role of Dopamine in Drosophila Larval Classical Olfactory Conditioning. *PLoS ONE*, 4(6):e5897.

Shadlen, M. N. and Kiani, R. (2013). Decision Making as a Window on Cognition. *Neuron*, 80(3):791–806.

Shadlen, M. N. and Newsome, W. T. (1996). Motion perception: Seeing and deciding. *Proceedings of the National Academy of Sciences of the United States of America*, 93(2):628–33.

Shadlen, M. N. and Shohamy, D. (2016). Decision Making and Sequential Sampling from Memory. *Neuron*, 90(5):927–939.

Shen, W., Hamilton, S. E., Nathanson, N. M., and Surmeier, D. J. (2005). Cholinergic Suppression of KCNQ Channel Currents Enhances Excitability of Striatal Medium Spiny Neurons. *Journal of Neuroscience*, 25(32):7449–7458.

Shirvaikar, M. and Bushnaq, T. (2009). A comparison between DSP and FPGA platforms for real-time imaging applications. In Kehtarnavaz, N. and Carlsohn, M. F., editors, *Real-Time Image and Video Processing*, volume 7244, page 724406. International Society for Optics and Photonics.

Shyu, W.-H., Chiu, T.-H., Chiang, M.-H., Cheng, Y.-C., Tsai, Y.-L., Fu, T.-F., Wu, T., and Wu, C.-L. (2017). Neural circuits for long-term water-reward memory processing in thirsty Drosophila. *Nature Communications*, 8:15230.

Simpson, J. H. and Looger, L. L. (2018). Functional Imaging and Optogenetics in Drosophila. *Genetics*, 208(4):1291–1309.

Sitaraman, D., LaFerriere, H., Birman, S., and Zars, T. (2012). Serotonin is Critical for Rewarded Olfactory Short-Term Memory in Drosophila. *Journal of Neurogenetics*, 26(2):238–244.

Sitaraman, D., Zars, M., LaFerriere, H., Chen, Y.-C., Sable-Smith, A., Kitamoto, T., Rottinghaus, G. E., and Zars, T. (2008). Serotonin is necessary for place memory in Drosophila. *Proceedings of the National Academy of Sciences*, 105(14):5579–5584.

Skeath, J. B. and Thor, S. (2003). Genetic control of Drosophila nerve cord development. *Current Opinion in Neurobiology*, 13(1):8–15.

Skinner, B. F. (1938). *The Behavior of Organisms: An Experimental Analysis*. NY: Appleton-Century-Crofts, New York.

Smith, P. L. and Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, 27(3):161–168.

Soares dos Santos, M. P. and Ferreira, J. (2014). Novel intelligent real-time position tracking system using FPGA and fuzzy logic. *ISA Transactions*, 53(2):402–414.

Stengård, E. and van den Berg, R. (2019). Imperfect Bayesian inference in visual perception. *PLoS Computational Biology*, 15(4):e1006465.

Stephens, G. J., Johnson-Kerner, B., Bialek, W., and Ryu, W. S. (2008). Dimensionality and Dynamics in the Behavior of C. elegans. *PLoS Computational Biology*, 4(4):e1000028.

Steyvers, M., Lee, M. D., and Wagenmakers, E.-J. (2009). A Bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53(3):168–179.

Stowers, J. R., Hofbauer, M., Bastien, R., Griessner, J., Higgins, P., Farooqui, S., Fischer, R. M., Nowikovsky, K., Haubensak, W., Couzin, I. D., Tessmar-Raible, K., and Straw, A. D. (2017). Virtual reality for freely moving animals. *Nature Methods*, 14(10):995–1002.

Straw, A. D., Branson, K., Neumann, T. R., and Dickinson, M. H. (2011). Multi-camera real-time three-dimensional tracking of multiple flying animals. *Journal of the Royal Society Interface*, 8(56):395–409.

Straw, A. D. and Dickinson, M. H. (2009). Motmot, an open-source toolkit for realtime video acquisition and analysis. *Source Code for Biology and Medicine*, 4:5.

Sun, R., Delly, J., Sereno, E., Wong, S., Chen, X., Wang, Y., Huang, Y., and Greenspan, R. J. (2020). Anti-instinctive Learning Behavior Revealed by Locomotion-Triggered Mild Heat Stress in Drosophila. *Frontiers in Behavioral Neuroscience*, 14:41.

Surmeier, D. J., Ding, J., Day, M., Wang, Z., and Shen, W. (2007). D1 and D2 dopamine-receptor modulation of striatal glutamatergic signaling in striatal medium spiny neurons. *Trends in Neurosciences*, 30(5):228–235.

Swierczek, N. A., Giles, A. C., Rankin, C. H., and Kerr, R. A. (2011). High-throughput behavioral analysis in C. elegans. *Nature Methods*, 8(7):592–598.

Tadres, D. and Louis, M. (2020). PiVR: An affordable and versatile closed-loop platform to study unrestrained sensorimotor behavior. *PLoS Biology*, 18(7):e3000712.

Tajima, S., Drugowitsch, J., and Pouget, A. (2016). Optimal policy for value-based decision-making. *Nature Communications*, 7(1):12400.

Takagi, S., Cocanougher, B. T., Niki, S., Miyamoto, D., Kohsaka, H., Kazama, H., Fetter, R. D., Truman, J. W., Zlatic, M., Cardona, A., and Nose, A. (2017). Divergent Connectivity of Homologous Command-like Neurons Mediates Segment-Specific Touch Responses in Drosophila. *Neuron*, 96:1–15.

Takeda, K. (1961). Classical conditioned response in the honey bee. *Journal of Insect Physiology*, 6(3):168–179.

Tanimoto, Y. and Kimura, K. D. (2019). Neuronal, mathematical, and molecular bases of perceptual decision-making in C. elegans. *Neuroscience Research*, 140:3–13.

Tastekin, I., Khandelwal, A., Tadres, D., Fessner, N. D., Truman, J. W., Zlatic, M., Cardona, A., and Louis, M. (2018). Sensorimotor pathway controlling stopping behavior during chemotaxis in the Drosophila melanogaster larva. *eLife*, 7:e38740.

Tastekin, I., Riedl, J., Schilling-Kurz, V., Gomez-Marin, A., Truman, J. W., and Louis, M. (2015). Role of the Subesophageal Zone in Sensorimotor Control of Orientation in Drosophila Larva. *Current Biology*, 25(11):1448–1460.

Tavoni, G., Balasubramanian, V., and Gold, J. I. (2019). What is optimal in optimal inference? *Current Opinion in Behavioral Sciences*, 29:117–126.

Theodoni, P., Kovacs, G., Greenlee, M. W., and Deco, G. (2011). Neuronal Adaptation Effects in Decision Making. *Journal of Neuroscience*, 31(1):234–246.

Thorndike, E. L. (1911). *Animal intelligence; experimental studies*. The Macmillan company, New York.

Tonegawa, S., Pignatelli, M., Roy, D. S., and Ryan, T. J. (2015). Memory engram storage and retrieval. *Current Opinion in Neurobiology*, 35:101–109.

Topál, J., Byrne, R. W., Miklósi, Á., and Csányi, V. (2006). Reproducing human actions and action sequences: "Do as I do!" in a dog. *Animal Cognition*, 9(4):355–367.

Tracey, W. D., Wilson, R. I., Laurent, G., and Benzer, S. (2003). painless, a Drosophila Gene Essential for Nociception. *Cell*, 113(2):261–73.

Trimmer, P. C., Houston, A. I., Marshall, J. A. R., Mendl, M. T., Paul, E. S., and McNamara, J. M. (2011). Decision-making under uncertainty: biases and Bayesians. *Animal Cognition*, 14(4):465–476.

Tully, T., Cambiazo, V., and Kruse, L. (1994). Memory through Metamorphosis. *Journal of Neuroscience*, 14(1):68–74.

Turner, G. C., Bazhenov, M., and Laurent, G. (2008). Olfactory Representations by Drosophila Mushroom Body Neurons. *Journal of Neurophysiology*, 99(2):734–746.

Uchida, N. and Mainen, Z. F. (2003). Speed and accuracy of olfactory discrimination in the rat. *Nature Neuroscience*, 6(11):1224–1229.

Uzun, I. S., Amira, A., and Bouridane, A. (2005). FPGA implementations of fast Fourier transforms for real-time signal and image processing. *IEE Proceedings - Vision, Image, and Signal Processing*, 152(3):283–296.

Valone, T. J. (2006). Are animals capable of Bayesian updating? An empirical review. *Oikos*, 112(2):252–259.

van Dam, E. A., Noldus, L. P., and van Gerven, M. A. (2020). Deep learning improves automated rodent behavior recognition within a specific experimental setup. *Journal of Neuroscience Methods*, 332:108536.

Veeraraghavan, A., Chellappa, R., and Srinivasan, M. (2008). Shape-and-Behavior Encoded Tracking of Bee Dances. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(3):463–476.

Vinauger, C., Lutz, E. K., and Riffell, J. A. (2014). Olfactory learning and memory in the disease vector mosquito Aedes aegypti. *Journal of Experimental Biology*, 217(Pt 13):2321–2330.

Vogelstein, J. T., Park, Y., Ohyama, T., Kerr, R., Truman, J. W., Priebe, C. E., and Zlatic, M. (2014). Discovery of Brainwide Neural-Behavioral Maps via Multiscale Unsupervised Structure Learning. *Science*, 344(6182):386–392.

Vogt, K., Schnaitmann, C., Dylla, K. V., Knapek, S., Aso, Y., Rubin, G. M., and Tanimoto, H. (2014). Shared mushroom body circuits underlie visual and olfactory memories in Drosophila. *eLife*, 3:e02395.

von Essen, A. M. H. J., Pauls, D., Thum, A. S., and Sprecher, S. G. (2011). Capacity of Visual Classical Conditioning in Drosophila larvae. *Behavioral Neuroscience*, 125(6):921–929.

Waddell, S. (2013). Reinforcement signalling in Drosophila; dopamine does it all after all. *Current Opinion in Neurobiology*, 23(3):324–329.

Webb, B. (2004). Neural mechanisms for prediction: do insects have forward models? *Trends in Neurosciences*, 27(5):278–282.

Weiglein, A., Gerstner, F., Mancini, N., Schleyer, M., and Gerber, B. (2019). One-trial learning in larval Drosophila. *Learning & Memory*, 26(4):109–120.

Wen, J. Y. M., Kumar, N., Morrison, G., Rambaldini, G., Runciman, S., Rousseau, J., and van der Kooy, D. (1997). Mutations That Prevent Associative Learning in C. elegans. *Behavioral Neuroscience*, 111(2):354–368.

Wirtz, R. A. and Semey, H. G. (1982). The Drosophila kitchen-equipment, media preparation, and supplies. *Drosophila Information Service*, 58:176–180.

Wolf, R. and Heisenberg, M. (1991). Basic organization of operant behavior as revealed in Drosophila flight orientation. *Journal of Comparative Physiology A*, 169(6):699–705.

Wolf, R., Wittig, T., Liu, L., Wustmann, G., Eyding, D., and Heisenberg, M. (1998). Drosophila Mushroom Bodies Are Dispensable for Visual, Tactile, and Motor Learning. *Learning & Memory*, 5(1-2):166–78.

Wong, A. M., Wang, J. W., and Axel, R. (2002). Spatial Representation of the Glomerular Map in the Drosophila Protocerebrum. *Cell*, 109(2):229–241.

Wright, G. A., Mustard, J. A., Simcock, N. K., Ross-Taylor, A. A. R., McNicholas, L. D., Popescu, A., and Marion-Poll, F. (2010). Parallel Reinforcement Pathways for Conditioned Food Aversions in the Honeybee. *Current Biology*, 20(24):2234–2240.

Wu, M.-C., Chu, L.-A., Hsiao, P.-Y., Lin, Y.-Y., Chi, C.-C., Liu, T.-H., Fu, C.-C., and Chiang, A.-S. (2014). Optogenetic control of selective neural activity in multiple freely moving Drosophila adults. *Proceedings of the National Academy of Sciences*, 111(14):5367–5372.

Wystrach, A., Lagogiannis, K., and Webb, B. (2016). Continuous lateral oscillations as a core mechanism for taxis in drosophila larvae. *eLife*, 5:e15504.

Xiang, Y., Yuan, Q., Vogt, N., Looger, L. L., Jan, L. Y., and Jan, Y. N. (2010). Light-avoidance-mediating photoreceptors tile the Drosophila larval body wall. *Nature*, 468(7326):921–926.

Yasukawa, S., Okuno, H., Ishii, K., and Yagi, T. (2016). Real-time object tracking based on scale-invariant features employing bio-inspired hardware. *Neural Networks*, 81:29–38.

Yoshino, J., Morikawa, R. K., Hasegawa, E., and Emoto, K. (2017). Neural Circuitry that Evokes Escape Behavior upon Activation of Nociceptive Sensory Neurons in Drosophila Larvae. *Current Biology*, 27(16):2499–2504.

Zemelman, B. V., Lee, G. A., Ng, M., and Miesenböck, G. (2002). Selective Photostimulation of Genetically ChARGed Neurons. *Neuron*, 33(1):15–22.

Zhang, C., Liang, T., Mok, P. K. T., and Yu, W. (2017). FPGA Implementation of the Coupled Filtering Method and the Affine Warping Method. *IEEE Transactions on Nanobioscience*, 16(5):314–325.