

The use of computational modeling to link sensory processing with behavior in *Drosophila*

A major goal of systems neuroscience is to understand how the brain represents information, and how those representations are used to drive behavior. Even though genetic model organisms like *Drosophila* grant unprecedented experimental access to the nervous system for manipulating and recording neural activity, the complexity of natural stimuli and natural behaviors still poses a challenge for solving the connections between neural activity and behavior. Here, we advocate for the use of computational modeling to complement (and enhance) the *Drosophila* toolkit. We first lay out a modelling framework for making sense of the relation between natural sensory stimuli, neuronal responses, and natural behavior. We then highlight how this framework can be used to reveal how neural circuits drive behavior, using selected case studies.

The Use of Computational Modeling to Link Sensory Processing with Behavior in *Drosophila*

To appear in: Celik, A., Wernet, M. Decoding Neural Circuit Structure and Function: Cellular Dissection Using Genetic Model Organisms. Springer. 2017.

Jan Clemens and Mala Murthy
Princeton Neuroscience Institute
Princeton University, Princeton, NJ

1. The Challenge

A major goal of systems neuroscience is to understand how the brain represents information, and how those representations are used to drive behavior. Animal brains have evolved to solve particular problems, such as detecting the movements of prey or the features of a suitable mate and changing patterns of locomotion accordingly. Studying these natural behaviors allows systems neuroscientists access to the (potentially conserved) computations and neural mechanisms underlying sensory processing, decision-making, and motor control in these animal model systems.

The genetic model organism *Drosophila melanogaster* exhibits a range of robust and complex behaviors such as acoustic communication during courtship (Coen et al. 2014; Clemens et al. 2015a; Coen et al. 2016), detection and integration of multisensory cues to locate food sources (van Breugel and Dickinson 2014; Bell and Wilson 2016), visually guided flight control (Clark et al. 2011; Aptekar et al. 2012; de Vries and Clandinin 2012; Silies et al. 2014), and avoidance of threats (Reyn et al. 2014). Recording *Drosophila* behaviors can be done in high-throughput, and its numerically compact and largely hardwired nervous system facilitates identifying the underlying neurons. In combination with an unparalleled genetic toolkit for manipulating and monitoring neuronal activity during behavior, flies represent an ideal model system for solving neural mechanisms that link sensory processing with behavior.

Here, we advocate for the use of computational modeling to complement (and enhance) the *Drosophila* toolkit. Modeling, in particular, allows one to make sense of the highly complex relation between natural sensory stimuli, neuronal responses, and natural behavior. For instance, for natural behaviors, both sensory stimuli and motor outputs can vary over multiple temporal and spatial scales and the underlying transformations can be highly nonlinear – this presents a challenge for determining what aspects of the sensory world drive a particular response. This is especially acute for social interactions, in which the behavioral output of one animal constitutes the most relevant sensory stimulus for the other animal.

In addition to illuminating the relation between sensory stimuli and behavior, computational modelling can also reveal how neurons represent behaviorally relevant stimuli and how these neural representations are read out to produce appropriate motor patterns – in short, how sensory information is encoded and decoded by the brain. Ideally, to solve these relationships, neural activity is recorded in behaving animals. Correlated neuronal and behavioral variability can then be exploited for inferring relationships between the two (Britten et al. 1996; Parker and Newsome 1998), and neuronal *decoding* models can identify the dynamical and nonlinear relationship between neuronal codes and behavior (Haefner et al. 2013). Recent advances have now facilitated combining neural and behavioral recordings in head-fixed *Drosophila* (Seelig et al. 2010; Kim et al. 2015), but head-fixing can limit the behavioral repertoire of the animal, particularly for social behaviors. In these cases, modeling is particularly useful since it permits linking data sets from separate experiments – neural recordings from fixed animals with behavioral recordings from freely moving animals. Neural *encoding* models can be derived from recordings in fixed, non-behaving animals – these models can then be used for predicting neural responses to the sensory stimuli from behavioral data sets. Thus, computational models can serve as stand-ins for recording neural activity during behavior and thus facilitate overcoming experimental hurdles when linking neural codes and natural behaviors (Parnas et al. 2013; Schulze et al. 2015; Clemens et

al. 2015a; Badel et al. 2016). Here, we detail this approach using data from *Drosophila* (both adults and larvae), and we discuss selected studies that highlight both the challenges and advantages associated with computational modeling in this model system.

2. The Approach

Understanding how the brain generates behavior in response to sensory stimuli involves bridging three different levels of description and can conceptually be separated into three steps (Figure 1A):

1. Linking stimuli and behavior (“**psychophysics**”) to identify the stimulus features driving behavior.
2. Linking stimuli and neural codes (“**encoding**”) to determine how the behaviorally relevant stimulus features are represented in the brain.
3. Linking neural codes and behavior (“**decoding**”) to reveal how neuronal activity is read out and transformed into behavior.

Interestingly, a single modelling framework has been successfully employed in all three steps, so-called **linear-nonlinear (LN)** models (Figure 1B). This class of models has been largely developed in the context of linking stimuli and neural codes (step 2 above) (e.g. (Schwartz et al. 2006; Pillow et al. 2008; Sharpee 2013; Aljadeff et al. 2016)) but has also recently found applications in predicting behavior from both sensory stimuli and neuronal responses (e.g. (Kato et al. 2014; Coen et al. 2014; Schulze et al. 2015; Clemens et al. 2015a)). LN models treat the brain as a black box, i.e. they reduce the complex action of neural networks, individual neurons, and ion channels to a sequence of relatively simple computations. In this light, it is surprising this framework works as often as it does for characterizing neural mechanisms.

In general, LN models describe the transformation from one or more temporally or spatially varying inputs to an output in two computational steps: a linear filter and a nonlinearity (Figure 1B). In the first step, a linear filter acts as a template that is compared to the input and thus constitutes the stimulus pattern most effective in driving the output. Mathematically, the comparison between the filter and the input is implemented through “convolution”, where the input at each point in time (τ) and/or space (x,y) is multiplied with the template and then summed. By making time explicit in the filter, the LN model can thus capture complex dynamical relationships.

$$s'(t) = \sum_{x,y,\tau} s(x,y,t-\tau)h(x,y,\tau)$$

In the second step of the LN model a static nonlinearity (NL) maps the filtered stimulus to the output – the neuronal or behavioral response (Chichilnisky 2001; Schwartz et al. 2006; Aljadeff et al. 2016). This part mimics properties of the input-output relationship such as thresholds (describing the minimally effective stimulus required to produce a response) or response saturation.

This separation of an input-output transformation into a dynamical linear and a static nonlinear part is a powerful feature of LN models. It facilitates model identification from often limited and noisy experimental data as well as the interpretation of model structure. For example, the linear filter helps identify spatial and dynamical properties of an input driving the response (Figure 2A-C) (Suh and Baccus 2014). A unilobed linear filter indicates that integration or smoothing underlies the response. By contrast, a biphasic filter – with both a positive and a negative lobe – has differentiating properties, indicating that the response relies on transients in the stimulus. More complex filters like Gabor filters, which correspond to short sinusoids whose amplitude is modulated by a Gaussian, have band-pass properties – confining responses only to a narrow range of the input spectrum (Clemens and Hennig 2013).

The nonlinearity in the LN model corresponds to a tuning curve for the filter – or the selectivity of the response for filter outputs. A monotonically increasing NL indicates that response magnitude is determined by stimulus magnitude, while a unimodal NL indicates that only a limited range of filter outputs drives the response (Figure

2D). Separating input-output transformations thus helps to discriminate changes in overall response gain from changes in stimulus selectivity. For example, motivation, context, or attention can change the overall gain of responses – and therefore affect the nonlinearity – without changing the feature selectivity of the response represented by the filter (Baccus and Meister 2002; Rabinowitz et al. 2011; Clemens et al. 2015b).

How can one identify the two components of an LN model from experimental data? Usually, linear filters and nonlinearities are estimated in separate steps. The standard approach for estimating the first stage of the model – the filter – is by averaging all stimuli leading to a response (Figure 3A-B). This is most straightforward for LN models of neuronal responses: the so-called spike-triggered average (STA) is computed by aligning the stimuli preceding all spikes and averaging them.

$$c_{STA}(\tau) = \sum_{\{t_{event}\}} r(t_{event}) s(t_{event} - \tau)$$

This can also be applied to behavioral events (in lieu of spikes) like the onset of an escape behavior (Reyn et al. 2014) or the individual pulses of fly song (Coen et al. 2016). A simple generalization enables applying the STA approach to continuous outputs like neuronal calcium levels or the amplitude of the animal's communication signal (Kato et al. 2014; Coen et al. 2016). In these cases, the filter is estimated by weighting the stimulus by the response value.

$$c(\tau) = \sum_t r(t) s(t - \tau)$$

The STA is the special case in which the response r is either 0 (no event occurred) or 1 (event occurred). The advantage of this approach is that the filter can be directly computed from stimulus response pairs, without the need for computationally expensive optimization procedures. While the STA has intuitive appeal and is useful for visualizing the relation between input and output, it provides an accurate estimate of the stimulus feature driving the response only for inputs with uncorrelated structure (such as white noise) (Paninski 2003; Sharpee 2013). However, natural stimuli are highly structured and characterized by higher-order correlations associated with edges in visual scenes or transients in acoustic signals (Simoncelli and Olshausen 2001; Machens and Zador 2003).

Alternatives to the STA method have been developed that permit identifying the filter without the strong requirement of uncorrelatedness of input statistics (Truccolo et al. 2005; Pillow et al. 2008; Sharpee 2013; Williamson et al. 2015; Aljadeff et al. 2016). These methods find the filter by directly minimizing the mismatch between the prediction of the model and the experimentally measured response using numerical optimization techniques in which the filter shape is modified gradually over many iterations to improve model fit. Efficient implementations of the optimization procedure are standard in all common statistical analysis packages and often include a procedure termed *regularization*, which involves adding a penalty term to the error function that is minimized to find the optimal filter. The penalty biases the filter to have desired or biologically “reasonable” properties. For instance, stimulus correlations lead to filter structure that does not contribute to model fit, but rather impede interpretation. By penalizing filters with many non-zero or large values, the impact of stimulus correlations can be minimized, greatly facilitating the interpretation of filter structure and improving model performance for stimuli with different correlation structure (Mineault et al. 2009; Park and Pillow 2011).

Generalized linear models (GLMs) are a powerful implementation of this approach, in which filter shape is free to vary, but the nonlinearity is constrained to a specific parametric form (Aljadeff et al. 2016) – most commonly the NLS used are logistic, exponential or linear (Figure 3C). Pre-defining the type of nonlinearity in the second stage of the LN model reduces the number of free model parameters and hence the amount of data needed for fitting the model, making GLMs especially suitable when data is scarce or noisy. However, in some cases it may be necessary to fit a model with an unconstrained nonlinearity. In these cases, the filtered stimulus is

usually plotted against the actual responses (Figure 3D) and the relationship between the two is estimated either by binning and averaging or by fitting a suitable parametric function (Chichilnisky 2001).

Model evaluation and selection

Model performance is usually quantified as the match between actual data and model predictions, using a quantity such as the mean-square error or the correlation coefficient. For GLMs that predict probabilities, models should be evaluated based on the model likelihood (see (Aljadeff et al. 2016) for details). Alternatively, when the output is the probability of a binary event – e.g. whether a spike will occur or not – one can predict events from the model output by setting a decision threshold (e.g. if $p > 0.5$ generate the event, otherwise generate no event) and use methods from signal detection theory (e.g., the percentage of correct predictions) to assess model performance (Coen et al. 2014). However, care needs to be taken when the event to be predicted is extremely rare – e.g. when only 1% of the inputs induce an event. A model that always predicts non-events would be correct in 99% of the cases. To improve model fit, such data sets should be “balanced”, by selecting random subsets of data for fitting and evaluation such that events and non-events occur with equal frequency (Coen et al. 2014).

If a model has many free parameters it can reflect properties of the data set used for fitting, not the general input-output relation of the system under study. The model performance calculated from the same data used to train the model can thus greatly overestimate the model quality. The impact of this so-called “overfitting” is usually assessed via cross-validation – computing model performance with data not used for fitting (e.g. (Schönfelder and Wichmann 2012; Aljadeff et al. 2016)). Practically, this is done by splitting the data into a training set for estimating model parameters (usually 80% of data) and a test set for quantifying model performance (usually 20%). This split can be randomized and repeated – parameters and performance values from different runs can be averaged and their variance estimated to statistically compare the parameters or performances of models obtained for different experimental conditions.

The ability to statistically compare models is essential for model selection – a method of selecting the simplest model that can explain the response. Generally, there exists a trade-off between model complexity and performance – a model with many parameters usually performs better than a simpler one. Several measures of model fit – such as the Akaike or the Bayesian information criterion (Zucchini 2000) – approach this problem by penalizing models with many parameters. In simple cases, one can select models based on a simple cutoff criterion – e.g. that a more complex model needs to outperform the simple one by at least 10% (Coen et al. 2014).

Extensions of the standard linear-nonlinear model

The relatively simple, standard LN model with a single linear filter and a static nonlinearity can be extended to account for more complex input-output transformations. Several extensions explicitly take into account possible nonlinear interactions between multiple filters operating on a single stimulus, like spike-triggered covariance analysis (Schwartz et al. 2006) or generalized quadratic models (Rajan et al. 2013). However, with greater model power comes the need for more data to identify model parameters, rendering these methods impractical when data are limited. Often, neuronal or behavioral responses are non-stationary or dependent on internal factors (e.g. behavioral states). While the standard GLM framework can still accommodate these cases (Pillow et al. 2008; Fründ et al. 2014), more complex types of non-stationarities affecting the stimulus response mapping can be modelled by combining LNs with hidden Markov models (HMM), in which the HMM models the switching between states and one LN model per state captures state-specific input-output transformations ((Escola et al. 2011), see (Wiltschko et al. 2015) for a related approach).

Another feature of neural computations is that they often bridge several time scales. For instance, the decision to engage in a behavior is often based on the detection of a stimulus on short time scales (milliseconds) while the resulting motor output evolves on much longer time scales (seconds or minutes). Models that exploit this

separation of time scales usually contain an explicit integration stage. This approach has been applied successfully to study the behavioral evaluation of courtship songs (Clemens and Hennig 2013; Clemens et al. 2015a). When extended with a behavioral threshold, such models resemble so-called “drift diffusion models”, which integrate sensory information – in this case the output of the LN model – to a decision boundary, the crossing of which corresponds to decision making (Brunton et al. 2013; DasGupta et al. 2014; Clemens et al. 2014).

3. Linking sensory stimuli with behavior

Fitting a model to link sensory cues and behaviors requires measuring 1) the putatively relevant sensory stimuli and 2) the behavioral output to be predicted. A number of studies using *Drosophila* have examined the impact of visual, chemosensory, or thermal cues on navigational decisions (Gomez-Marin et al. 2011; Clark et al. 2011; Censi et al. 2013; Klein et al. 2014; Gepner et al. 2015; Schulze et al. 2015). In these studies, the detailed recording of fly movement parameters with a quantitative, model-supported analysis of the behavior revealed the navigational strategies with which *Drosophila* negotiates its environment to find food and avoid noxious chemical, heat, or light stimuli. For example, (Censi et al. 2013) examined the visual features driving navigation in adult *Drosophila* during free-flight in a circular arena. They tracked the position and angle of flies in 3D and estimated the stimulus features – given by the distance and angle from the arena’s wall – that triggered rapid changes in direction (or saccades). They found that flies turn away from the wall when approaching it, e.g. they tended to turn right when the wall was to their left. Given that their model could predict more than 90% of the saccades, the authors concluded that the majority of flight decisions in their arena was driven by visual cues. (Gomez-Marin et al. 2011) used behavioral manipulation and computational modeling to show how *Drosophila* larvae find the source of an attractive odorant. The first step in understanding the computations underlying this behavior was to reconstruct the larva’s sensory environment – in this case the concentration of odorants at the chemosensory organ in the head. This was done by tracking the larval head position and determining – either through measurement or physical modeling – the local concentration of odorants the larva receives at each moment in time. This highlights the importance of reconstructing the stimulus relative to the animal’s frame of reference. An animal’s sensory experience depends on its position relative to the stimulus source, on how the stimulus propagates through the medium, and on the properties of their sense organs (e.g. their directionality). *Drosophila* larval behaviors can easily be classified into three relatively stereotyped motor programs during chemotaxis: phases of straight locomotion (termed “runs”), intense side-to-side movement of the head (“head casting”), and fast reorientation events (“turns”). Models that explain the relation between chemosensory cues and navigation behaviors usually use the rate or probability of performing one of these behaviors as a function of the stimulus (Gomez-Marin et al. 2011; Gepner et al. 2015; Schulze et al. 2015; Hernandez-Nunez et al. 2015). Based on an accurate description of the larva’s sensory experience as well as of its complete body posture during navigation, the authors determined the features of the odor concentration gradient that triggered turning and head casting using STA-like analyses. This revealed the behavioral program controlling navigation: decreases in concentration stop a run and induce head casting. Upon realizing that it has moved away from the source, the larva actively samples the concentration gradient by casting its head, which translates weak spatial differences in concentration into strong temporal differences. When perceiving an increase in concentration during the head cast, the larva initiates the turn, therefore reorienting it towards higher concentration. This study exemplifies the power of modeling for revealing how different behavioral modules – running, head-casting, and turning – are engaged in a stimulus-dependent manner to support successful navigation behavior.

Modeling is also useful to characterize sensorimotor relationships during more complex behaviors, like social interactions. During courtship, *Drosophila* males chase females and produce a song by vibrating their wings. The song’s pattern is evaluated by the female and informs her mating decision (Coen et al. 2014; Clemens et al. 2015a). (Coen et al. 2014) investigated the origin of the male’s highly variable song pattern. The song consists of two, relatively stereotyped modes – “pulse” and “sine” – which are sequenced into variable bouts, stemming from the fact that the duration of each song mode is not fixed and that males rapidly transition

between the two modes. (Coen et al. 2014) asked whether the variable bout structure is purely stochastic – e.g. due to noise in the song-producing circuits – or whether it is the outcome of the male dynamically and adaptively shaping song structure in response to sensory cues. The authors tracked the movement of the male and female, along with recording the song produced by the male – the movements and interactions of the pair served as inputs to a GLM with regularization. The regularization during model fitting (Mineault et al. 2009) helped the authors reduce the impact of strong correlations among different inputs, in addition to minimizing temporal correlations within a single input. This model-based analysis of the behavior revealed that almost every aspect of song patterning – song starts, the choice of song mode, or song ends – are predicted by the movements and interactions of the flies. For instance, the distance between male and female is most predictive of song starts: the filter that predicts whether or not a song start will occur reveals that the male only begins song in sine mode when near the female. In a follow-up study, using the same GLM method, the authors showed that not only do males bias toward pulse mode (the louder song mode) when far away from the female, but they also modulate the amplitude of each pulse within this mode relative to distance to the female (Coen et al. 2016). Overall, the model-based analysis of courtship interactions revealed that the diversity present in song structure results from the male dynamically and adaptively controlling when and what to sing in response to changing feedback from the female.

These three case studies showcase that modelling can facilitate extracting information from large and complex behavioral data sets. In providing a detailed and quantitative description of the mapping between the animal's sensory experience and its behavior, these studies 1) reveal the computations that govern behavior and 2) lay the groundwork for identifying the neural circuits and neural codes underlying the behavior.

4. Linking sensory stimuli with neural responses – neuronal encoding models

How are behaviorally relevant stimulus features represented by the brain? Modeling – especially the use of LN models – has a long tradition in characterizing neural responses to sensory stimuli for many model systems and sensory modalities (see e.g. (Schwartz et al. 2006; Aljadeff et al. 2016)). In the context of *Drosophila* systems neuroscience, LN models have provided insight into the dynamical stimulus features driving neuronal responses in the olfactory system (Nagel and Wilson 2011; Martelli et al. 2013), visual system (de Vries and Clandinin 2012; Seelig and Jayaraman 2013; Behnia et al. 2014)), and auditory system (Clemens et al. 2015a). These LN models provide information about what sensory information is available to the brain for driving behaviors and in which format it is represented. Moreover, the LN model of a neuron – obtained from responses recorded in fixed animals, for example – can serve as a surrogate for neural recordings when it is not feasible to record neural activity during behavior. This is particularly useful in organisms like *Drosophila* with genetically and morphologically identifiable cell types, whose response properties are stereotyped in most, but not all (e.g. (Murthy et al. 2008)), cases.

As an example, (Clemens et al. 2015a) used an LN model extension to characterize the encoding of sound in the fly brain. The authors recorded intracellularly from a set of 10, morphologically diverse neurons from early stages of the fly's auditory pathway: the antennal mechanosensory and motor center (AMMC), which is the projection area of the auditory receptor neurons, and the wedge and the ventrolateral protocerebrum, which receives its main auditory inputs from the AMMC. Sound is largely encoded via graded membrane voltage (Vm) responses, not spikes. The authors hence fit LN models to the neuronal membrane voltage changes. A standard LN model was not sufficient to reproduce the Vm traces of these neurons due to prominent adaptation, and hence the authors used an extension – a so-called “adaptive LN model” – in which the stimulus is first processed by an adaptation stage and the adapted stimulus is then fed into a standard LN model with a filter and a nonlinearity. Crucially, the model which was fitted to neuronal responses for artificial song pulse trains, generalized well to predict Vm responses to natural courtship song, suggesting that these computational steps are sufficient to describe the neural code for song processing in these auditory neurons. While adaptation parameters were relatively diverse, reflecting the diversity of adaptation time scales and strengths for the neurons in the data set, the shapes of the linear filters and nonlinearities were similar across

the models of all 10 cell types, which is surprising given the large morphological diversity of the neurons in the data set. The linear filters were biphasic – with a dominant positive lobe followed by a shallow but long negative lobe. The positive lobe conveys low-pass filter properties – smoothing the stimulus envelope – while the negative lobe has a weakly differentiating effect – it accentuates transients like the onsets or offsets of song bouts (Figure 2A-C). Analysis of female behavioral responses to male song during natural courtship revealed that female slowing – a proxy of her willingness to mate – was strongly correlated with the duration of song bouts over timescales of 10s of seconds. The temporal properties of the neuronal code for song in the auditory system are well-suited for extracting information about bout structure. Bout starts are encoded in positive Vm transients due to adaptation, negative Vm transients encode bout ends due to the negative filter lobe, and the amount of song is represented in sustained responses during bouts due to the positive filter lobe.

(Schulze et al. 2015) used models to show how responses of larval olfactory sensory neurons (OSN) encode the temporal pattern of odorant concentration as encountered during chemotaxis. Instead of using a linear-nonlinear model to represent the dynamical codes underlying concentration coding in these OSN, they employed a model of the signal transduction cascade involved in the transformation of odorant concentration to OSN firing rate. Similar models have been used to model neural codes in other chemosensory systems (Kato et al. 2014) or in the vertebrate retina (Ozuysal and Baccus 2012). This encoding model revealed the temporal stimulus features represented in the time-varying OSN firing rate: for positive concentration gradients, OSNs encode the slope; for negative gradients, OSNs act as offset detectors. These temporal features of olfactory stimuli are known to guide chemotaxis behaviors (see above, (Gomez-Marin et al. 2011)): the duration of runs as well as the initiation and direction of turns are guided by positive slopes in concentration, while stopping and head casting occur upon abrupt decreases of concentration. Again, the properties of the dynamical codes in OSNs directly support navigation behavior. While these two studies highlight the power of careful quantification of behavior for interpreting neural codes, these interpretations must ultimately be tested by linking neural codes directly with behavior via decoding models.

5. Linking neural responses with behavior – neuronal decoding models

Having identified both the sensory cues driving behavior and how these cues are encoded by the brain, the next task is to determine how neuronal codes are transformed into behavior. Modeling helps 1) to infer the brain's readout strategies and 2) to identify the properties of neural codes most relevant for driving a particular behavior. As laid out in the Introduction, the neural responses for fitting decoding models can be directly recorded during the behavior, or, if not feasible, obtained either by replay of the animal's sensory experience while recording from a fixed preparation or by generating surrogate neuronal responses through virtual recordings from neural encoding models.

Several studies of *Drosophila* have employed decoding models to investigate the relationship between chemosensory codes and behavior (Parnas et al. 2013; Gepner et al. 2015; Hernandez-Nunez et al. 2015; Badel et al. 2016; Bell and Wilson 2016). For example, (Hernandez-Nunez et al. 2015) asked how the activity of gustatory receptor neurons (GRN) affects chemotaxis behavior. Optogenetic activation with temporal white noise patterns and STA analyses revealed the filters underlying the transition from run to turn behavior (and from turn to run behavior). Comparing the STA filters obtained from single-GRN activation experiments to those for multi-GRN activation experiments revealed that the magnitude as well as the shape of the filters for multi-GRN activation could not be predicted from a simple, linear combination of single-GRN STA filters. This demonstrated that gustatory information is integrated nonlinearly between individual GRNs in the context of behavior.

Modeling was also used to infer algorithms underlying the integration of sensory cues across modalities. (Gepner et al. 2015) studied how visual and chemosensory cues are integrated during larval orientation behavior. Visual stimuli were provided using blue light and chemosensory neurons were activated optogenetically (via csChrimson, (Klapoetke et al. 2014)) using red light, which is outside the sensitivity range

of the photoreceptors. Stimulation with temporal white noise patterns allowed the authors to infer the features driving behavioral responses (turn rates) for light and fictive chemosensory stimulation via STA analyses. This revealed that visual cues increased turn rates, while attractive chemosensory cues decreased turn rates. Using a decoding approach, they then determined how these two opposing cues that feed into the same motor output (turn rate) are combined. Specifically, they tested whether the two modalities were combined using “early integration” versus “independent pathways”. That is, whether the filtered signals from each modality were linearly combined to yield a single “turn rate” signal, or whether each modality generated its own, independent turn rate, which was then summed. Their modeling approach lent strong support to the early integration model.

Other studies have examined odortaxis in adult *Drosophila* and used decoding models to determine how the neural population code in the fly's antennal lobe (AL) is read out by downstream circuits. Each glomerulus in the AL receives exclusive input from OSNs expressing a single odorant receptor. Second-order projection neurons (PN) typically innervate single glomeruli and link the AL to higher-order olfactory centers in the lateral horn (for innate behaviors) and the mushroom body (for learned behaviors) (Jefferis et al. 2007; Aso et al. 2014).

Similar to (Hernandez-Nunez et al. 2015), (Bell and Wilson 2016) asked how glomerular activation is translated into walking behavior to odors (odortaxis). Specifically, they tested whether the behavioral effects of multiple OSN activation sum linearly or nonlinearly. Interestingly, responses to pairwise activation frequently did not match the sum of the responses to the activation of the individual OSN in a pair. Instead, the responses to pairwise OSN activation often equaled the responses to the stronger OSN in the pair, suggesting a “max” pooling of OSN activity by downstream circuits. In a similar study, (Badel et al. 2016) used odorant stimulation – not optogenetic activation of individual OSNs – to directly link naturalistic glomerular codes for odors to odortaxis in flying *Drosophila*. They showed that the glomerular code is relatively linear, since the neuronal responses to mixtures can be predicted from the responses to their constituents. This was consistent with the observation of linearity in the behavioral responses to mixtures. However, by examining behavioral responses to odors presented in sequences, they found that odortaxis over longer timescales depends on the olfactory “context” – the other odors in the sequence. This context effect can be strong and even switch the valence of an individual odor from attractive to aversive and is inconsistent with a linear glomerular code. The authors then used an extension of their linear decoding model that includes canonical computations – mean subtraction and divisive normalization – to reproduce the context effect. These computations are likely implemented in higher-order olfactory areas downstream of the AL glomeruli.

The models mentioned above infer behavioral readout strategies from a *static* description of neuronal population activity – but modeling can also facilitate reading out *dynamical* neuronal responses. For example, (Schulze et al. 2015) used a nonlinear dynamic encoding model of *Drosophila* larval OSNs (see above) to relate neural responses in OSNs to behavioral dynamics (time-varying turn rates). The decoding model used was a GLM and consisted of a single linear weight followed by a sigmoidal transform. To determine whether OSN dynamics were essential for driving behavioral dynamics, the authors predicted behavior either directly from the stimulus or from model OSN responses. Taking into account OSN dynamics greatly improved behavioral predictions, suggesting the neuronal dynamics in the sensory periphery strongly shape behavioral dynamics.

In the (Clemens et al. 2015a) study (see above), the authors discovered that female slowing during courtship, a proxy of her willingness to mate, was strongly associated with long timescale features of the male's courtship song – in particular, the average duration of song bouts over timescales of 10s of seconds. The authors then used a decoding model to link the neuronal responses to song in the auditory system with the female slowing response. The highly interactive nature of the courtship chase prevented recording neural activity in the auditory system during behavior and the authors instead relied on an encoding model to reconstruct the neural representation of song recorded during courtship. The encoding model (see above) consisted of two computations: an adaptation stage, which produced positive transients at the start of each song bout, and a

linear filter with a positive and a negative lobe, which produced sustained neural responses during a bout and negative neural responses at the end of each bout. Using a decoder, the authors asked i) How is this neuronal representation of bout structure read out? and ii) What encoder computations are crucial for producing the behavior? They started with a decoder that transformed the reconstructed neuronal response to female speed in two computational steps. First, a sigmoidal nonlinearity transformed the neuronal responses and mimicked common nonlinearities in putative downstream circuits like thresholds or saturation. Second, an integration stage linked the timescales of neuronal computations in the auditory system (a few hundred milliseconds) to the long behavioral timescales (tens of seconds). By manipulating the encoding model, the authors determined the computations in the auditory system essential for reproducing the behavior. Although both adaptation and biphasic filtering were necessary to reproduce the neuronal responses to song in the auditory system, the behavior mainly relied on aspects of the neuronal response associated with biphasic filtering – sustained neural responses during a bout and negative neural responses at the end of each bout. While the integral of the sustained response corresponded to the amount of song, the integral of the negative offset response encoded the number of bouts (since each bout in the integration time window produced the same, stereotypical offset response). Thus, the decoder predicted behavior by combining song amount and bout number. Because bout duration, which the behavioral analysis had identified as being most strongly associated with the female's slowing behavior, is given by the ratio of these two quantities, the authors modified their decoder to directly compute bout duration from the neuronal responses. This decoder almost perfectly reproduced the behavioral relation between bout duration and female speed during natural courtship. Overall, the decoding analysis revealed 1) the aspects of the neural representation of song in the auditory system likely to be essential for generating behavior (in this case, biphasic filtering) and 2) probable computations (such as integration and division) that transform neural codes in the auditory system into behavior. The approach highlights how computational modeling can both help to overcome the experimental difficulties associated with recording neuronal activity during interactive, social behaviors and to generate hypotheses regarding the computations underlying sensorimotor transformations.

6. Experimental tests of computational models

The above examples reveal how modeling can provide insight into the neural computations that transform natural stimuli into behavior. However, to establish causality, models should be tested experimentally. Experimental tests can involve manipulations of sensory stimuli, of known sensory pathways, or of the activity of subsets of neurons. This is greatly facilitated in *Drosophila* via genetic tools for activating or inactivating genes or neuronal activity during behavior. For instance, based on behavioral analyses (Ramdya et al. 2015) posited that collective behavior in *Drosophila* relies on mechanosensory cues. While individual flies only poorly avoid aversive odorants, groups of flies, by bumping into each other, “push” each other out of the aversive odor zone. The authors genetically activated and inactivated leg mechanosensory neurons to demonstrate the sufficiency and necessity of mechanosensory cues for this collective behavior. That is, inactivation abolished the behavior while activation was sufficient to induce flies to move. Similarly, the decoding model by (Badel et al. 2016) (see above) linked the olfactory population code in the fly's antennal lobe to odortaxis and predicted that individual glomeruli have only a small impact on navigational decisions. Consistent with this notion, they found that inactivating single glomeruli negligibly changed behavioral responses.

Another strategy for testing models fitted to natural behavior is to use artificial, controlled stimuli. This approach has a long tradition for testing models of visual motion-processing in *Drosophila*, where models are tested by monitoring the behavioral responses of tethered flies to artificial visual stimuli (Eichner et al. 2011; Fitzgerald et al. 2015; Leonhardt et al. 2016). This strategy was also employed by (Coen et al. 2016) who used GLMs to identify the distance between males and females during courtship as the most significant predictor of the male's song amplitude. The authors tested their model using a virtual reality setup in which the visual cues available to the male could be precisely manipulated. The fly was tethered and allowed to walk on an air-suspended ball in front of a computer screen. Artificial visual stimuli – a black square moving on a white background with movement statistics similar to those encountered during courtship by the male - elicited song

amplitude modulation as a function of the size of the square, a visual feature strongly correlated with distance. Similarly, (van Breugel and Dickinson 2014) used stimulus manipulations to test the behavioral observation that olfactory and visual stimuli interact during odortaxis in free flight. Specifically, they placed high-contrast visual cues within a flight tunnel to show that attractive olfactory stimuli increase the visual saliency of objects.


In the light of correlations between sensory cues, models derived from natural behaviors often have difficulties discriminating between internally and externally generated cues (Censi et al. 2013). In (Coen et al. 2014), GLMs predicted an association between male forward velocity and song mode - faster males were more likely to produce the so-called “pulse song” versus “sine song”. However, the nature of the sensory cue associated with male forward velocity was unclear. Was the cue the optic flow generated from the male’s own motion? Was it a signal internally generated, like an efference copy of the motor drive or proprioceptive feedback from the muscles or joints? Using flies carrying mutations that ablate all photoreceptor cells, the authors found that models built on blind male data performed similarly to models built on wild type male data – in both cases, the male’s own motion effectively predicted song mode choice. To more explicitly test for a link between male locomotor circuits and song patterning, the authors fixed the male flies in place and induced singing by optogenetically activating song command neurons. Comparing the optogenetically induced song of fixed versus freely moving males, they found that males that cannot move produced more “pulse song”, demonstrating that interfering with locomotion alters singing. These examples highlight how experimental tests reduce ambiguity regarding which sensory cues and neural circuits drive behavior.

7. Prospects

The case studies cited above highlight how the use of computational modeling powerfully complements the *Drosophila* genetic toolkit for solving open questions in systems neuroscience. We expect that methodological innovations – both experimental and theoretical – will further increase the utility of computational modeling in this model system. Imaging technology now permits recording from larger numbers of *Drosophila* neurons simultaneously (Bouchard et al. 2015; Harris et al. 2015; Lu et al. 2016; Aimon et al. 2016) or from subsets of neurons in freely behaving flies (Grover et al. 2016). Computational models applied to such data will facilitate both interpreting population neural activity and connecting neural activity with behavior. In parallel, the use of unsupervised classification methods has revealed stereotyped structure in animal behavior – an animal’s movements over time can be described as sequences of discrete behavioral modules (Vogelstein et al. 2014; Berman et al. 2014; Wiltischko et al. 2015). Computational modeling’s task will now be to determine how sensory cues and internal states affect behavioral sequencing, and how neural codes underlie the choice of behavioral modules. In conclusion, combining the wealth of genetic tools to dissect the neural circuits underlying behavior in *Drosophila* with advances in machine learning and computational modeling now makes it more feasible than ever to link sensory processing, neural representations, and behavior in this system.

Figures & Legends


Figure 1



A Linking sensory stimuli, neuronal responses and behavior can conceptually be separated into three steps: First, linking sensory stimuli and behavior to identify the stimulus features driving behavior (psychophysics). Second, linking sensory stimuli and neuronal responses to describe how behaviorally relevant stimulus features are represented in the brain (encoding). Third, linking neuronal responses with behavior to understand how neuronal representations are read out and transformed into behavior (decoding).

B Linear nonlinear (LN) models are useful in all three steps, since they describe the input-output transformation with two computational steps: First, a linear filter (L) acts as a template that is compared to the input. In a second step the filtered input is transformed to the output using a fixed nonlinearity (N), which can mimic a threshold or saturation in the input-output transformation.

Figure 2




A The action of linear filters can be best understood as lying between two extremes: a uni-lobed filter with integrating properties (blue) and a bi-lobed filter with differentiating properties (orange). Typically, experimentally derived filters constitute a mixture between these purely integrating and purely differentiating filters.

B The frequency domain representation, obtained by applying a Fourier transform to the filters in **A**, reveals that the two filters have fundamentally different frequency transfer properties: integrating filters (blue) mainly let pass low-frequencies, while differentiators (red) are selective for a limited range of frequencies, rejecting both very low and very high frequencies.

C The response properties of integrating and differentiating filters become obvious when looking at responses to a step-like input (black): integrators sum up the input over their duration and reject high-frequencies associated with sharp edges in the stimulus, effectively smoothing the stimulus. By contrast, differentiators only encode stimulus changes but not sustained stimulus epochs and hence only respond transiently during the step's on- and offset.

D The nonlinearity (NL) in an LN model transforms the filtered input (bottom) to the output. An accelerating NL, e.g. an exponential function (orange), amplifies strong inputs and accentuates stimulus differences. By contrast, a saturating nonlinearity, e.g. a sigmoidal function (yellow), compresses large values and thereby reduces differences in the input. An optimum-like, bell-shaped NL (purple) produces outputs only for intermediate input values.

Figure 3



A To illustrate the process of identifying an LN model from experimental data, we generated a trial-averaged firing rate (orange) and a sparse spike train (black lines, top) from a white noise stimulus (gray) using an LN model with a bandpass filter and a simple, threshold-linear nonlinearity.

B Extracting all stimuli preceding a spike yields the so-called spike-triggered stimulus ensemble (STE, grey, shown for the 4 spikes in A, spike is to right – at the end – of the stimulus snippet). The average of this STE is the spike-triggered average (STA, black, average over 200 spikes) which resembles the true filter (blue) used for generating the spikes.

C The nonlinearity (NL) in the LN model can be obtained using two approaches. In the GLM approach, the nonlinearity is fixed and only the filter can vary. The specific form of the NL is usually dictated by the type of output to be predicted. For binary events (e.g. spike/no spike), the NL is a logistic function. For event rates (e.g. firing rate), the NL is an exponential function. For continuous outputs with positive and negative values the standard NL is usually a linear function.

D LN models can also be estimated using a free-form NL. One way of identifying the NL in these cases is by filtering the stimulus with the STA and plotting the filtered stimulus vs. the actual response (gray). The relationship can then be binned and averaged or fitted using a suitable parametric function. For the example shown here, the resulting plot recovers the thresholding NL (black) used for generating the response in A.

References

- Aimon S, Katsuki T, Grosenick L, Broxton M, Deisseroth K, Sejnowski TJ, Greenspan RJ (2016) Linking stimuli and behavior with fast near-whole brain recordings in adult *Drosophila*. *bioRxiv* 033803. doi: 10.1101/033803
- Aljadeff J, Lansdell BJ, Fairhall AL, Kleinfeld D (2016) Analysis of Neuronal Spike Trains, Deconstructed. *Neuron* 91:221–259. doi: 10.1016/j.neuron.2016.05.039
- Aptekar JW, Shoemaker PA, Frye MA (2012) Figure Tracking by Flies Is Supported by Parallel Visual Streams. *Current Biology* 22:482–487. doi: 10.1016/j.cub.2012.01.044
- Aso Y, Hattori D, Yu Y, Johnston RM, Iyer NA, Ngo T-T, Dionne H, Abbott L, Axel R, Tanimoto H, Rubin GM (2014) The neuronal architecture of the mushroom body provides a logic for associative learning. *eLife*. doi: 10.7554/eLife.04577
- Baccus SA, Meister M (2002) Fast and Slow Contrast Adaptation in Retinal Circuitry. *Neuron* 36:909–919. doi: 10.1016/S0896-6273(02)01050-4
- Badel L, Ohta K, Tsuchimoto Y, Kazama H (2016) Decoding of Context-Dependent Olfactory Behavior in *Drosophila*. *Neuron*. doi: 10.1016/j.neuron.2016.05.022
- Behnia R, Clark DA, Carter AG, Clandinin TR, Desplan C (2014) Processing properties of ON and OFF pathways for *Drosophila* motion detection. *Nature* 512:427–430. doi: 10.1038/nature13427
- Bell JS, Wilson RI (2016) Behavior Reveals Selective Summation and Max Pooling among Olfactory Processing Channels. *Neuron*. doi: 10.1016/j.neuron.2016.06.011
- Berman GJ, Bialek W, Shaevitz JW (2016) Predictability and hierarchy in *Drosophila* behavior.
- Berman GJ, Choi DM, Bialek W, Shaevitz JW (2014) Mapping the stereotyped behaviour of freely moving fruit flies. *J R Soc Interface* 11:20140672–20140672. doi: 10.1098/rsif.2014.0672
- Bouchard MB, Voleti V, Mendes CS, Lacefield C, Grueber WB, Mann RS, Bruno RM, Hillman EMC (2015) Swept confocally-aligned planar excitation (SCAPE) microscopy for high-speed volumetric imaging of behaving organisms. *Nature Photon* 9:113–119. doi: 10.1038/nphoton.2014.323
- Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13:87–100.
- Brunton BW, Botvinick MM, Brody CD (2013) Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. *Science* 340:95–98. doi: 10.1126/science.1233912
- Censi A, Straw AD, Sayaman RW, Murray RM, Dickinson MH (2013) Discriminating External and Internal Causes for Heading Changes in Freely Flying *Drosophila*. *PLoS Comput Biol* 9:e1002891. doi: 10.1371/journal.pcbi.1002891.s002
- Chichilnisky EJ (2001) A simple white noise analysis of neuronal light responses. *Network: Computation in Neural Systems* 12:199–213.
- Clark DA, Bursztyn L, Horowitz MA, Schnitzer MJ, Clandinin TR (2011) Defining the Computational Structure

of the Motion Detector in *Drosophila*. *Neuron* 70:1165–1177. doi: 10.1016/j.neuron.2011.05.023

Clemens J, Girardin CC, Coen P, Guan X-J, Dickson BJ, Murthy M (2015a) Connecting Neural Codes with Behavior in the Auditory System of *Drosophila*. *Neuron* 87:1332–1343. doi: 10.1016/j.neuron.2015.08.014

Clemens J, Hennig RM (2013) Computational principles underlying the recognition of acoustic signals in insects. *Journal of Computational Neuroscience* 35:75–85. doi: 10.1007/s10827-013-0441-0

Clemens J, Krämer S, Ronacher B (2014) Asymmetrical integration of sensory information during mating decisions in grasshoppers. *Proc Natl Acad Sci U S A* 111:16562–16567. doi: 10.1073/pnas.1412741111

Clemens J, Rau F, Hennig RM, Hildebrandt KJ (2015b) Context-dependent coding and gain control in the auditory system of crickets. *The European journal of neuroscience* 42:2390–2406. doi: 10.1111/ejn.13019

Coen P, Clemens J, Weinstein AJ, Pacheco DA, Deng Y, Murthy M (2014) Dynamic sensory cues shape song structure in *Drosophila*. *Nature* 507:233–237. doi: 10.1038/nature13131

Coen P, Xie M, Clemens J, Murthy M (2016) Sensorimotor Transformations Underlying Variability in Song Intensity during *Drosophila* Courtship. *Neuron* 89:629–644. doi: 10.1016/j.neuron.2015.12.035

DasGupta S, Ferreira CH, Miesenböck G (2014) FoxP influences the speed and accuracy of a perceptual decision in *Drosophila*. *Science* 344:901–904. doi: 10.1126/science.1252114

de Vries SEJ, Clandinin TR (2012) Loom-Sensitive Neurons Link Computation to Action in the *Drosophila* Visual System. *Current Biology* 22:353–362. doi: 10.1016/j.cub.2012.01.007

Eichner H, Joesch M, Schnell B, Reiff DF, Borst A (2011) Internal Structure of the Fly Elementary Motion Detector. *Neuron* 70:1155–1164. doi: 10.1016/j.neuron.2011.03.028

Escola S, Fontanini A, Katz D, Paninski L (2011) Hidden Markov Models for the Stimulus-Response Relationships of Multistate Neural Systems. *Neural computation* 23:1071–1132. doi: 10.1162/neco_a_00118

Fitzgerald JE, Clark DA, Carandini M (2015) Nonlinear circuits for naturalistic visual motion estimation. *eLife* 4:e09123. doi: 10.7554/eLife.09123

Fründ I, Wichmann FA, Macke JH (2014) Quantifying the effect of intertrial dependence on perceptual decisions. *Journal of Vision* 14:9–9. doi: 10.1167/14.7.9

Gepner R, Mihovilovic Skanata M, Bernat NM, Kaplow M, Gershow M (2015) Computations underlying *Drosophila* photo-taxis, odor-taxis, and multi-sensory integration. *eLife* 4:599. doi: 10.7554/eLife.06229

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

Grover D, Katsuki T, Greenspan RJ (2016) Flyception: imaging brain activity in freely walking fruit flies. *Nat Methods*. doi: 10.1038/nmeth.3866

Haefner RM, Gerwinn S, Macke JH, Bethge M (2013) Inferring decoding strategies from choice probabilities in the presence of correlated variability. *Nature neuroscience* 16:235–242. doi: 10.1038/nn.3309

Harris DT, Kallman BR, Mullaney BC, Scott K (2015) Representations of Taste Modality in the *Drosophila* Brain. *Neuron* 86:1449–1460. doi: 10.1016/j.neuron.2015.05.026

- Hernandez-Nunez L, Belina J, Klein M, Si G, Claus L, Carlson JR, Samuel AD, Calabrese RL (2015) Reverse-correlation analysis of navigation dynamics in *Drosophila* larva using optogenetics. *eLife* 4:e06225. doi: 10.7554/eLife.06225
- Jefferis GSXE, Potter CJ, Chan AM, Marin EC, Rohlifing T, Maurer CR, Luo L (2007) Comprehensive maps of *Drosophila* higher olfactory centers: spatially segregated fruit and pheromone representation. *Cell* 128:1187–1203. doi: 10.1016/j.cell.2007.01.040
- Kato S, Xu Y, Cho CE, Abbott LF, Bargmann CI (2014) Temporal Responses of *C. elegans* Chemosensory Neurons Are Preserved in Behavioral Dynamics. *Neuron* 81:616–628. doi: 10.1016/j.neuron.2013.11.020
- Kim AJ, Fitzgerald JK, Maimon G (2015) Cellular evidence for efference copy in *Drosophila* visuomotor processing. *Nature neuroscience* 18:1247–1255. doi: 10.1038/nn.4083
- Klapoetke NC, Murata Y, Kim SS, Pulver SR, Birdsey-Benson A, Cho YK, Morimoto TK, Chuong AS, Carpenter EJ, Tian Z, Wang J, Xie Y, Yan Z, Zhang Y, Chow BY, Surek B, Melkonian M, Jayaraman V, Constantine-Paton M, Wong GK-S, Boyden ES (2014) Independent optical excitation of distinct neural populations. *Nat Methods* 11:338–346. doi: 10.1038/nmeth.2836
- Klein M, Afonso B, Vonner AJ, Hernandez-Nunez L, Berck M, Tabone CJ, Kane EA, Pieribone VA, Nitabach MN, Cardona A, Zlatic M, Sprecher SG, Gershow M, Garrity PA, Samuel ADT (2014) Sensory determinants of behavioral dynamics in *Drosophila* thermotaxis. *Proc Natl Acad Sci U S A* 112:E220–E229. doi: 10.1073/pnas.1416212112
- Leonhardt A, Ammer G, Meier M, Serbe E, Bahl A, Borst A (2016) Asymmetry of *Drosophila* ON and OFF motion detectors enhances real-world velocity estimation. *Nature neuroscience* 19:706–715. doi: 10.1038/nn.4262
- Lu R, Sun W, Liang Y, Kerlin A, Bierfeld J, Seelig J, Wilson DE, Scholl B, Mohar B, Tanimoto M, Koyama M, Fitzpatrick D, Orger MB, Ji N (2016) Video-rate volumetric functional imaging of the brain at synaptic resolution. *bioRxiv* 058495. doi: 10.1101/058495
- Machens CK, Zador AM (2003) Auditory modeling gets an edge. *Journal of Neurophysiology* 90:3581–3582. doi: 10.1152/jn.00832.2003
- Martelli C, Carlson JR, Emonet T (2013) Intensity invariant dynamics and odor-specific latencies in olfactory receptor neuron response. *Journal of Neuroscience* 33:6285–6297. doi: 10.1523/JNEUROSCI.0426-12.2013
- Mineault PJ, Barthelmé S, Pack CC (2009) Improved classification images with sparse priors in a smooth basis. *Journal of Vision* 9:1–24.
- Murthy M, Fiete I, Laurent G (2008) Testing odor response stereotypy in the *Drosophila* mushroom body. *Neuron* 59:1009–1023. doi: 10.1016/j.neuron.2008.07.040
- Nagel KI, Wilson RI (2011) Biophysical mechanisms underlying olfactory receptor neuron dynamics. *Nature neuroscience* 14:208–216. doi: 10.1038/nn.2725
- Ozuysal Y, Baccus SA (2012) Linking the computational structure of variance adaptation to biophysical mechanisms. *Neuron* 73:1002–1015. doi: 10.1016/j.neuron.2011.12.029
- Paninski L (2003) Convergence properties of three spike-triggered analysis techniques. *Network: Computation* 16

in *Neural Systems* 14:437–464.

- Park M, Pillow JW (2011) Receptive Field Inference with Localized Priors. *PLoS Comput Biol* 7:e1002219. doi: 10.1371/journal.pcbi.1002219
- Parker AJ, Newsome WT (1998) Sense and the single neuron: probing the physiology of perception. *Annu Rev Neurosci* 21:227–277. doi: 10.1146/annurev.neuro.21.1.227
- Parnas M, Lin AC, Huetteroth W, Miesenböck G (2013) Odor Discrimination in *Drosophila*: From Neural Population Codes to Behavior. *Neuron* 79:932–944. doi: 10.1016/j.neuron.2013.08.006
- Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454:995–999. doi: 10.1038/nature07140
- Rabinowitz NC, Willmore BDB, Schnupp JW, King AJ (2011) Contrast gain control in auditory cortex. *Neuron* 70:1178–1191. doi: 10.1016/j.neuron.2011.04.030
- Rajan K, Marre O, Tkačik G (2013) Learning Quadratic Receptive Fields from Neural Responses to Natural Stimuli. *Neural computation*. doi: 10.1162/NECO_a_00463
- Ramdyia P, Lichocki P, Cruchet S, Frisch L, Tse W, Floreano D, Benton R (2015) Mechanosensory interactions drive collective behaviour in *Drosophila*. *Nature* 519:233–236. doi: 10.1038/nature14024
- Reyn von CR, Breads P, Peek MY, Zheng GZ, Williamson WR, Yee AL, Leonardo A, Card GM (2014) A spike-timing mechanism for action selection. *Nature neuroscience* 17:962–970. doi: 10.1038/nn.3741
- Schönfelder VH, Wichmann FA (2012) Sparse regularized regression identifies behaviorally-relevant stimulus features from psychophysical data. *The Journal of the Acoustical Society of America* 131:3953–3969. doi: 10.1121/1.3701832
- Schulze A, Gomez-Marin A, Rajendran VG, Lott G, Musy M, Ahammad P, Deogade A, Sharpe J, Riedl J, Jariault D, Trautman ET, Werner C, Venkadesan M, Druckmann S, Jayaraman V, Louis M (2015) Dynamical feature extraction at the sensory periphery guides chemotaxis. *eLife*. doi: 10.7554/eLife.06694
- Schwartz O, Pillow JW, Rust NC, Simoncelli EP (2006) Spike-triggered neural characterization. *Journal of Vision* 6:484–507. doi: 10.1167/6.4.13
- Seelig JD, Chiappe ME, Lott GK, Dutta A, Osborne JE, Reiser MB, Jayaraman V (2010) Two-photon calcium imaging from head-fixed *Drosophila* during optomotor walking behavior. *Nat Methods* 7:535–540. doi: 10.1038/nmeth.1468
- Seelig JD, Jayaraman V (2013) Feature detection and orientation tuning in the *Drosophila* central complex. *Nature*. doi: 10.1038/nature12601
- Sharpee TO (2013) Computational identification of receptive fields. *Annu Rev Neurosci* 36:103–120. doi: 10.1146/annurev-neuro-062012-170253
- Silies M, Gohl DM, Clandinin TR (2014) Motion-Detecting Circuits in Flies: Coming into View. *Annu Rev Neurosci* 37:307–327. doi: 10.1146/annurev-neuro-071013-013931
- Simoncelli EP, Olshausen BA (2001) Natural image statistics and neural representation. *Annu Rev Neurosci*

24:1193–1216. doi: 10.1146/annurev.neuro.24.1.1193

Suh B, Baccus SA (2014) Building blocks of temporal filters in retinal synapses. *PLoS Biology* 12:e1001973. doi: 10.1371/journal.pbio.1001973

Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN (2005) A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology* 93:1074–1089. doi: 10.1152/jn.00697.2004

van Breugel F, Dickinson MH (2014) Plume-Tracking Behavior of Flying *Drosophila* Emerges from a Set of Distinct Sensory-Motor Reflexes. *Current Biology* 24:274–286. doi: 10.1016/j.cub.2013.12.023

Vogelstein JT, Park Y, Ohyama T, Kerr R, Truman JW, Priebe CE, Zlatic M (2014) Discovery of Brainwide Neural-Behavioral Maps via Multiscale Unsupervised Structure Learning. *Science*. doi: 10.1126/science.1250298

Williamson RS, Sahani M, Pillow JW (2015) The equivalence of information-theoretic and likelihood-based methods for neural dimensionality reduction. *PLoS Comput Biol* 11:e1004141. doi: 10.1371/journal.pcbi.1004141

Wiltschko AB, Johnson MJ, Iurilli G, Peterson RE, Katon JM, Pashkovski SL, Abaira VE, Adams RP, Datta SR (2015) Mapping Sub-Second Structure in Mouse Behavior. *Neuron* 88:1121–1135. doi: 10.1016/j.neuron.2015.11.031

Zucchini W (2000) An Introduction to Model Selection. *Journal of Mathematical Psychology* 44:41–61. doi: 10.1006/jmps.1999.1276