



Published in final edited form as:

Cell Rep. 2022 March 29; 38(13): 110574. doi:10.1016/j.celrep.2022.110574.

Dopamine neurons evaluate natural fluctuations in performance quality

Alison Duffy^{1,2}, Kenneth W. Latimer^{1,3}, Jesse H. Goldberg⁴, Adrienne L. Fairhall^{1,2,*}, Vikram Gadagkar^{5,6,*}

¹Department of Physiology and Biophysics, University of Washington, Seattle, WA 98195, USA

²Computational Neuroscience Center, University of Washington, Seattle, WA 98195, USA

³Department of Neurobiology, University of Chicago, Chicago, IL 60637, USA

⁴Department of Neurobiology and Behavior, Cornell University, Ithaca, NY 14853, USA

⁵Department of Neuroscience, Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY 10027, USA

⁶Lead contact

SUMMARY

Many motor skills are learned by comparing ongoing behavior to internal performance benchmarks. Dopamine neurons encode performance error in behavioral paradigms where error is externally induced, but it remains unknown whether dopamine also signals the quality of natural performance fluctuations. Here, we record dopamine neurons in singing birds and examine how spontaneous dopamine spiking activity correlates with natural fluctuations in ongoing song. Antidromically identified basal ganglia-projecting dopamine neurons correlate with recent, and not future, song variations, consistent with a role in evaluation, not production. Furthermore, maximal dopamine spiking occurs at a single vocal target, consistent with either actively maintaining the existing song or shifting the song to a nearby form. These data show that spontaneous dopamine spiking can evaluate natural behavioral fluctuations unperturbed by experimental events such as cues or rewards.

Graphical Abstract

This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

*Correspondence: fairhall@uw.edu (A.L.F.), vikram.gadagkar@columbia.edu (V.G.).

AUTHOR CONTRIBUTIONS

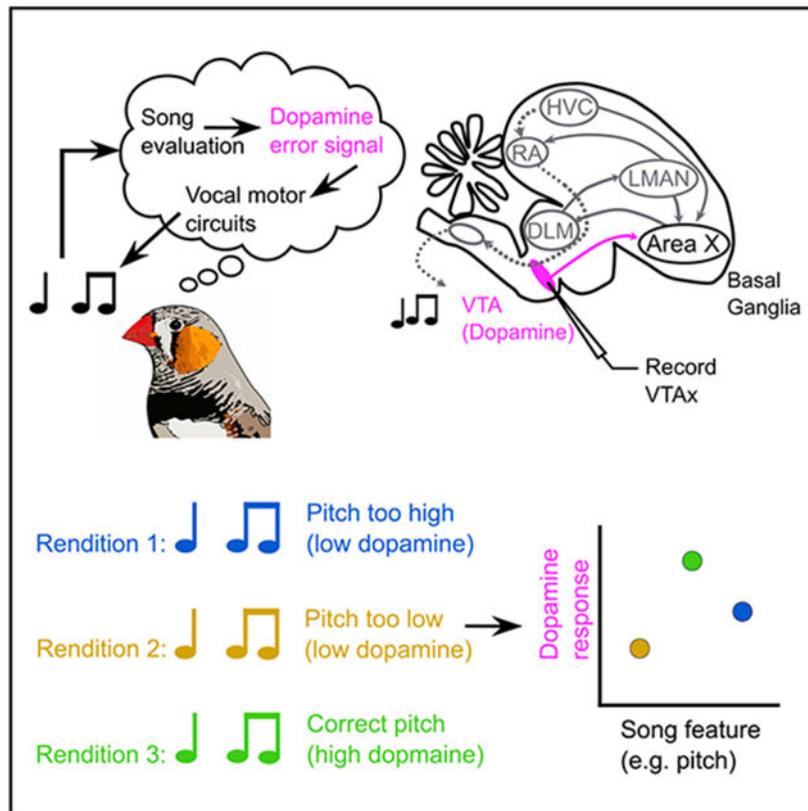
Conceptualization, A.D., K.W.L., J.H.G., A.L.F., and V.G.; experimental methodology and investigation, J.H.G. and V.G.; analysis methodology, A.D., K.W.L., A.L.F., and V.G.; software and formal analysis, A.D., K.W.L., and A.L.F.; writing – original draft, A.D. and V.G.; writing – review and editing, A.D., K.W.L., J.H.G., A.L.F., and V.G.; funding acquisition, J.H.G., A.L.F., and V.G.; supervision, J.H.G., A.L.F., and V.G.

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.celrep.2022.110574>.

DECLARATION OF INTERESTS

The authors declare no competing interests.



In brief

Learning and producing skilled behavior requires an internal measure of performance. Duffy et al. examine dopamine neurons' relationship to natural song in singing birds. Spontaneous dopamine activity correlates with song fluctuations in a manner consistent with evaluation of natural behavioral variations, independent of external perturbations, cues, or rewards.

INTRODUCTION

Dopamine (DA) is associated with fluctuations in future movements as well as the outcomes of past ones. During spontaneous behavior, DA activity can be phasically activated before a movement (da Silva et al., 2018; Hamilos et al., 2020), or can ramp as an animal approaches reward (Hamid et al., 2016; Howe et al., 2013). DA neurons can also signal a reward prediction error (RPE) during reward seeking, where phasic signals represent the value of a current outcome relative to previous outcomes (Schultz et al., 1997). It remains poorly understood how spontaneous DA activity relates to natural fluctuations in behavior that are independent of experimentally induced rewards or perturbations.

Zebra finches provide a tractable model to study the role of DA in natural behavior. First, they sing with a significant amount of trial-to-trial variability, but the overall stereotypy of the song allows renditions to be accurately compared. Second, they have a discrete neural circuit (the song system) that includes a DA-basal ganglia (BG) loop (Figures 1A and 1B) that is necessary for song learning and maintenance (Brainard and Doupe, 2000; Hisey

et al., 2018; Hoffmann et al., 2016; Xiao et al., 2018). Third, BG projecting DA neurons signal performance prediction error (PPE) during singing: they exhibit pauses following worse-than-predicted outcomes caused by distorted auditory feedback (DAF), and they exhibit phasic bursts following better-than-predicted outcomes when predicted distortions do not occur (Figures 1A-1D) (Gadagkar et al., 2016). Yet one limitation of this study was that song quality was controlled with an external sound (DAF; Figure 1C), so it remains unclear if the DA system is simply using song timing to build expectations about an external event (DAF), or if it also evaluates the quality of natural fluctuations (Figure 1E), which would be necessary for natural song learning. Furthermore, this experimental paradigm did not test if DA activity was associated with upcoming syllables, consistent with a premotor signal.

To test DA's role in natural behavior, we recorded from DA neurons in the ventral tegmental area (VTA) (Figure 1B), and examined how spiking activity correlated with natural song fluctuations (Figure 1E). First, if DA activity following externally distorted and undistorted song (Figure 1D) truly reflects a function of the DA system in performance evaluation, then DA activity should correlate with recent song fluctuations (Figure 1E). Second, DAF-associated error signaling was previously only observed in a small subclass of "VTA_{error}" neurons, most of which projected to Area X, the BG nucleus of the song system. "VTA_{other}" neurons were defined by the absence of an error signal during singing. We hypothesize that the VTA_{error} population will carry a performance error signal for natural song (Figure 1E), while the VTA_{other} population will not. Thus, we ask in this analysis: Do VTA neuron activity patterns relate to fluctuations in natural song? If so, what is the structure of these relationships, and do they relate to a performance evaluation framework, a premotor framework, or both?

To answer these questions, we first parameterized natural song into a low-dimensional set of time-varying song features. We then agnostically fit the relationship between rendition-to- rendition variations in song features and spike counts at local time steps in song across a range of song segment-spike window latencies and identified if and when song feature variations predicted spike counts. Finally, we characterized both the timing and the form of these predictive fits. We find that the activity of the VTA_{error}, but not the VTA_{other}, neuronal population encodes fluctuations in natural song in a manner consistent with a performance evaluation signal. These results show that basal ganglia-projecting DA neurons may provide continuous evaluation of natural motor performance independent of external rewards or perturbations.

RESULTS

We performed our analysis on $n = 22$ VTA_{error} neurons and $n = 23$ VTA_{other} neurons during uninterrupted, natural portions of singing. Regions of song selected for DAF were excluded from analysis. Classifications of neuron type were based on responses to the DAF paradigm described above. Neurons were recorded on single days during approximately 20–80 consecutive renditions of song.

A Gaussian process model approach reveals song-spike relationships

We sought to identify how VTA spiking varied with natural fluctuations in song syllables. To identify relationships between natural song fluctuations and VTA spiking, we chose an eight-dimensional, time-varying representation of song based on established song parameterizations, which have been shown in previous studies to relate to neural activity, vary over song development, and drive adult learning in DAF paradigms (Figure 2A; see STAR Methods) (Andalman and Fee, 2009; Deregnaucourt et al., 2004; Kao et al., 2005; Leblois et al., 2010; Ravbar et al., 2012; Sober and Brainard, 2009; Tchernichovski et al., 2000; Tumer and Brainard, 2007; Woolley and Kao, 2014). For each neuron, we identified song syllables and binned both song feature values and spike counts in sliding windows across each syllable to search for relationships between song fluctuations and spike counts at different latencies (song window width = 35 ms; spike window width = 100 ms) (Figure 2B; see STAR Methods). We combined all eight song features and binned spike counts into a single multi-dimensional Gaussian process (GP) regression model (two features shown for illustration) to quantify whether song feature fluctuations predicted spike counts (Figures 2C, S1, and S2; see STAR Methods). This strategy flexibly identifies the most relevant dimensions of song variation within a single model. Specifically, we computed an r^2 value for each model fit using leave-one-out cross-validation to assess how well variations in song features predicted spike counts (Figure 2C). Values of $r^2 > 0$ indicate that song feature variations across renditions can predict spike counts; the larger the r^2 value, the more predictive the song-spike relationship in the model. Finally, we fit the full model to many song-spike latencies and thus built a matrix of r^2 values for each neuron's response to song fluctuations, with each r^2 value in the matrix representing one full model fit (one feature shown for illustration) between a song window-spike window pair (Figure 2D).

Timing of song-spike relationships for VTA_{error} neurons is consistent with an evaluative process

Using the GP model approach described above, we asked if significant relationships between natural song fluctuations and VTA neuron spiking exist and, if so, at what song-spike latencies they occur. If VTA spiking is predictive of upcoming syllable fluctuations in a premotor fashion, then significant relationships would be observed at negative lags. Alternatively, if VTA spiking is playing an evaluative function, then variations in spike counts should follow variations in syllable acoustic structure, and relationships should be observed at positive lags. Based on past work (Gadagkar et al., 2016), an evaluation signal is predicted to occur at a positive lag of ~50 ms with a duration range of 0–150 ms. Figure 3A shows an example VTA_{error} neuron's song-spike relationship (the r^2 matrix) for a single syllable. The y axis is the midpoint of each song window aligned to syllable onset ($t = 0$). The x axis is the latency, defined as the time between the song window midpoint and the spike window midpoint. Colored pixels in the r^2 matrix indicate that song feature fluctuations predict spike counts ($r^2 > 0$); grayscale pixels indicate that song feature fluctuations do not predict spike counts ($r^2 \approx 0$). The pink box indicates the song-spike latencies (0–150 ms) where we expect to see evaluation-like relationships based on the DAF experimental results (Figure 1D). We assessed the significance of finding predictive fits by shuffling entire spike trains relative to song renditions and refitting our model across all latencies and song windows (Figure 3A, bottom and S3; see STAR Methods).

This population method of shuffling the data preserves the underlying temporal correlation structure of song and spiking while randomizing the song-spike relationship and allowed us to assess the significance of the entire set of model fits and account for multiple comparisons (see STAR Methods). The bottom matrix in Figure 3A shows the r^2 values from one such randomized shuffle of the same neuron's activity patterns. Positive r^2 values were found to be less frequent in the shuffled data (one-sided bootstrap test: $p < 0.01$; see STAR Methods).

We then analyzed the temporal relationship between song features and spiking by plotting the latency distribution for all song window-spike window pairs within the r^2 matrix in which song features predicted spike counts ($r^2 > 0$) (Figure 3B). The latencies of the predictive fits were clustered within the expected error evaluation range (0–150 ms) (Figure 3B). In Figure 3B, the blue line indicates the song-spike latency distribution for which song fluctuations predict spike counts ($r^2 > 0$) while the black line and gray shading are the mean and standard deviation of the same latency distributions across all randomized population shuffles. Figure 3C shows the result of the same analysis performed across all the VTA_{error} neurons in our dataset ($n = 22$). The true data showed a large peak within the expected PPE latency range (3.74 standard deviations from the mean, one-sided bootstrap test: $p < 0.01$; see STAR Methods). We found that song fluctuations are most predictive of spike counts 0–100 ms after the song fluctuation occurs, consistent with a PPE-like signal based on our previous DAF experiments (Figure 1D). In addition, across the population of VTA_{error} neurons there were significantly more predictive fits within the PPE latency window than expected by chance (one-sided bootstrap test: $p < 0.01$; see STAR Methods). Thus, the timing and frequency of the predictive song-spike relationships was remarkably consistent with a PPE-like response to natural song variations.

We next performed the identical analysis on a population of VTA_{other} neurons ($n = 23$), which did not show an error-like response in previous DAF experiments (Figure 1D). The number of predictive song-spike relationships from this population was also significantly larger than expected by chance (one-sided bootstrap test, $p < 0.01$; see STAR Methods). However, unlike the VTA_{error} neurons, the predictive relationships from this population did not cluster within the PPE latency range, nor did the variance of the distribution significantly deviate from the randomized latency distributions (one-sided bootstrap test: $p = 0.21$; Figure 3D; see STAR Methods). Thus, consistent with results from the DAF experiments (Figure 1D), only the responses of VTA_{error} , and not the non-error responsive VTA_{other} , neuron populations were predicted by natural song fluctuations within the expected PPE latency range with significantly increased frequency. The same neurons that exhibited error responses to the DAF sound exhibited significant relationships with natural syllable fluctuations. Remarkably, both the DAF-induced error and the natural fluctuations were at a similar latency with respect to song.

The form of the predominant song-spike relationship for VTA_{error} neurons is consistent with song maintenance

The hypothesis that VTA_{error} neurons evaluate natural song fluctuations led to further predictions about the forms of these song-spike relationships. If a bird is trying to maintain the acoustic structure of a syllable, then typical variations should be followed by more spikes

and rare, outlying syllable variations should be followed by fewer spikes (Figure 4A, top left). Alternatively, if the bird is trying to modify a syllable, e.g., increase its pitch, then the relationship between syllable acoustic structure and spike counts should be directional: spike counts should peak at whatever shifted variant to which the bird aspires but is not yet consistently producing (Figure 4A, right panels). We did not expect PPE-like signals to have multiple maxima in a *disruptive* fit: we assumed there is a single “best” version of the song at each time step (Figure 4A, bottom left).

To test these possible outcomes, we characterized tuning curve shapes of song-spike relationships. For this analysis, we focused our attention on the subset of GP model fits that were predictive ($r^2 > 0$). Within this subset, we further selected single-feature fits that were also predictive ($r^2 > 0$ for the 1D feature fit). We selected this subset of song-spike relationships because we are interested only in the tuning curves that might carry information about song. We re-fit all such song-spike relationships with a generalized linear model (GLM) using both linear (l-GLM) and linear and quadratic (q-GLM) features (Figure 4B; see STAR Methods). We chose these models because the parameters can be used to directly quantify aspects of the tuning curve shapes. If the song-spike fit has a single peak, the spiking response is stabilizing, and the quadratic coefficient of the q-GLM is negative (Figure 4A, top left). If the song-spike fit has two peaks, the spiking response is disruptive, and the quadratic coefficient of the q-GLM is positive (Figure 4A, bottom left). If the song-spike fit is monotonic, the spiking response is directional, and the l-GLM (with only linear features) is the more appropriate model (Figure 4A, right panels). Figure 4C shows examples of predictive relationships between individual song features and spike counts along with all model fits (GP, q-GLM, and l-GLM). Each point on these plots represents the song feature value and the spike count for a single rendition. Specific models were better fits for some distributions than others. For example, in panel 4 (Figure 4C, fourth panel from left), the q-GLM produced the same model fit as the l-GLM because the quadratic term added no improvement to the fit, whereas in panel 3 (Figure 4C, third panel from left), the quadratic term was necessary to accurately follow the spiking response and thus the q-GLM resulted in a better fit than the l-GLM.

When the spiking response is either stabilizing or disruptive, the quadratic coefficient of the q-GLM distinguishes between these two response types by indicating the direction of curvature. We used the Akaike information criterion (ΔAIC) to compare the relative success of the q-GLM and the l-GLM and, from this, identified potential curvature in our data (see STAR Methods). ΔAIC introduces a penalty for added parameters in a model to account for possible overfitting. $\Delta AIC > 0$ indicates that the quadratic model outperforms the linear model, considering both the likelihood of the model fit (the goodness of the fit) and the complexity of the model used (the number of parameters in the model). Because of this, we use $\Delta AIC > 0$ as the threshold for potential curvature in a song feature-spike count fit. The larger the ΔAIC , the better the q-GLM fit relative to the l-GLM. In Figure 4D, each point represents a fit to a single song feature with an $r^2 > 0$ within a multi-dimensional model fit with an $r^2 > 0$ for the population of VTA_{error} neurons within the expected PPE latency range. When we identified potential curvature in the data ($\Delta AIC > 0$), we found a significantly greater fraction (0.78) of q-GLM fits with negative quadratic coefficients, which indicates more stabilizing tuning curves than disruptive tuning curves (Figures 4D

and 4E; two-tailed bootstrap test, $p < 0.02$; see STAR Methods). Furthermore, when $AIC > 0$ the fraction of stabilizing fits also increased (Figure 4E, inset). Thus, the predictive fits from the GP model had significantly more stabilizing tuning curves than disruptive when their shape had identified curvature, as we expect for a PPE signal with a single best outcome (Figure 4E). This finding is consistent with our hypothesis that a PPE signal should respond most strongly to a single best performance of song. The AIC measure also allowed us to examine the fraction of tuning curves that are better fit by a linear versus quadratic model. The VTA_{error} population did not differ from chance in this fraction (fraction fits with $AIC > 0 = 0.43$; two-tailed bootstrap test, $p = 0.34$; see STAR Methods), consistent with a PPE signal with both directional and stabilizing responses depending on the current level of song error.

DISCUSSION

Value judgments in the brain are necessary to drive appropriate changes in behavior during learning. Using experimentally constrained tasks with external rewards, previous studies have found that DA neurons in VTA can encode a key component of value judgment: the mismatch between expected and actual reward outcomes, the RPE (Schultz et al., 1997). However, extending these findings to natural behavior and intrinsic reward has been a challenge. Here, we made use of a novel opportunity to use an experimental context to partition songbird VTA neurons into error and non-error classes and analyze their spiking responses in the context of a natural behavior (Gadagkar et al., 2016). We compared natural song fluctuations at a local, within-syllable scale with variations in spike counts of VTA neurons. We developed a GP regression analysis to quantify the non-stationary spiking response to variations in performance at different points in song and with different temporal relationships to song. We found evidence that VTA DA neurons' activity patterns correlate with variations in natural song in a manner consistent with performance evaluation: both the timing and tuning properties of the DA response was consistent with a PPE-like response. This finding corroborates and extends complementary discoveries of RPE signals emerging from mammalian DA neurons in VTA in experimentally imposed tasks. Previous studies have shown the significance of midbrain DA in song learning and maintenance in artificial manipulations of the circuit and behavior: 6-OHDA lesions disrupt DAF-induced shifts in pitch (Hoffmann et al., 2016). Optogenetic manipulations of VTA inputs to Area X induce changes in song consistent with a reinforcement paradigm (Hisey et al., 2018; Xiao et al., 2018). While these experiments implicated DA in song evaluation, our work directly shows that VTA inputs to Area X are active in a manner consistent with song evaluation during natural behavior. We did not find significant temporal relationships between DA and song fluctuations consistent with a premotor signal as has been observed in previous studies of DA (Barter et al., 2015; Engelhard et al., 2019). Our results provide direct evidence that DA neurons in VTA respond to fluctuations in natural behavior in a manner consistent with evaluation.

While, as a population, the non-error VTA_{other} neuron activity did not relate to song fluctuations in a manner consistent with a PPE signal, many neurons exhibited correlated relationships to song (Chen et al., 2021). These findings are consistent with previous studies in mammals, which found that both DA and non-DA neurons in VTA contribute to an RPE

calculation and that elements of the RPE signal are computed, in part, locally within VTA (Cohen et al., 2012; Dobi et al., 2010; Tian et al., 2016; Wood et al., 2017). Correlations with song variations in this population could represent components of the PPE calculation.

This project uses the structure of an experimentally grounded characterization of individual neurons' response to song-triggered DAF to analyze the same neurons during natural behavior. The connection to an existing experiment (Gadagkar et al., 2016) as well as to a reinforcement learning framework (Sutton and Barto, 1998) anchors our interpretations of natural behavior in a constrained laboratory paradigm and theory. The unusually high stereotypy of the natural behavior we consider, zebra finch song, allows reasonable inferences to be made both in the experimental and natural context about the behavior of the bird and a reasonable way to characterize and align a complex, natural behavior. We found a parallel relationship, including a striking temporal correspondence, between the VTA_{error} neuron activity in experimental and natural contexts that corroborates the experimental finding that VTA_{error} neurons encode time step-specific PPEs in song. Our analysis of natural song addresses the critique that the DAF experimental paradigm is aversive rather than perturbative and thus qualitatively different from natural song evaluation. A frequent debate in neuroscience is whether artificial behavioral paradigms serve as true building blocks for understanding neural activity in complex, freely behaving contexts, or whether they represent a different, overly simplified context that will not extrapolate to natural behavior. This experimentally guided study of natural behavior is a fruitful direction that permits the control of experimental contexts and the complexity of natural contexts to interact and build upon one another.

Limitations of the study

Two important predictions from these PPE-like signals we find in DA VTA neurons are that (1) future song renditions will shift toward song variations that correlate with the peak response in the DA neurons, and that (2) this shift will be accompanied by a decrease in the PPE peak response. We could not address these predictions here because of the limited duration of our single recording sessions. This will be an important direction for future work and will help disambiguate the role of the VTA DA responses from other possible relationships to song. We chose a pre-defined set of song features ($n = 8$) that have been shown to represent biologically relevant song variations in previous studies, because we focused on single sessions with limited data. Future work could apply more flexible, non-parametric dimensionality reduction methods using more song renditions to better identify VTA's relationship to song features that are most modulated by the bird at different points in song (Goffinet et al., 2021; Kollmorgen et al., 2020). Because of the size of our dataset and the subtlety of the natural behavior, our study draws conclusions at the population level. Populations of neurons were defined by their responses to the DAF experiments reported in Gadagkar et al. (2016). If, in future work, recordings of single cells could be carried out for longer periods of time, this would increase the signal-to-noise ratio and allow single-cell analyses of response significance. Finally, although we have only focused on natural song in this analysis and excluded distorted song regions, our dataset did not include a control in which songs were not targeted for distortion. Future work could apply our analysis techniques to recordings made with no distortions to any parts of song.

STAR★METHODS

RESOURCE AVAILABILITY

Lead contact—Further information and requests for resources and reagents should be directed to and will be fulfilled by Vikram Gadagkar (vikram.gadagkar@columbia.edu).

Materials availability—This study did not generate new unique reagents.

Data and code availability

- All data reported in this paper will be shared by the lead contact upon request.
- All original code has been deposited at Figshare and is publicly available as of the date of publication. The DOIs are listed in the key resources table.
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Subjects were 19 adult male zebra finches, 74–300 days old, singing undirected song. Male zebra finches (*Taeniopygia guttata*) used in this study were obtained from Magnolia Bird Farm in Anaheim, CA and kept on a 12h:12h light-dark cycle with ad-libitum access to food. All experiments were conducted in accordance with NIH guidelines and were approved by the Cornell Institutional Animal Care and Use Committee. Before implant surgeries, each bird was anesthetized with isoflurane. A bipolar stimulation electrode was implanted into Area X at previously established coordinates (+5.6A, +1.5L relative to lambda and 2.65 ventral relative to pial surface; head angle 20 degrees). Intraoperatively in each bird, antidromic methods were applied to locate the precise part of VTA containing VTax neurons. Next, custom made, plastic printed microdrives carrying an accelerometer, linear actuator, and homemade electrode arrays (5 electrodes, 3–5 MOhms, microprobes.com) were implanted into the region of VTA containing VTax neurons.

METHOD DETAILS

Syllable-targeted distorted auditory feedback—Detailed description of all aspects of the distorted auditory feedback (DAF) experiments is described elsewhere (Gadagkar et al., 2016). Descriptions of experimental details relevant to this study are presented here. Postoperative birds were put in a sound isolation chamber. The chamber was equipped with a microphone and two speakers which provided DAF. To carry out targeted DAF, the microphone signal was analyzed every 2.5 ms using custom Labview software. Specific syllables were targeted either by detecting a characteristic spectral feature in the previous syllable (using Butterworth band-pass filters) or by identifying an inter-onset interval (onset time of previous syllable to onset time of target syllable) using the sound amplitude (this procedure has been previously described (Ali et al., 2013; Hamaguchi et al., 2014; Tumer and Brainard, 2007)). A delay between 10–200 ms was applied between the detected song interval and the target time.

To help ensure that DAF would not be perceived as an aversive stimulus, the DAF sound was generated with the same amplitude and spectral content as normal zebra finch song. For broadband DAF ($n = 16$ birds), DAF was a broadband sound, band passed at 1.5-8kHz in order to be in the same spectral range of zebra finch song (Andalman and Fee, 2009). For displaced syllable DAF ($n = 10$ birds), DAF was a segment of one of the bird's own syllables displaced in the song. For both types of DAF, the amplitude was measured with a decibel meter (CEM DT-85A) and kept at less than 90 dB, which is the average maximum loudness of zebra finch song (Mandelblat-Cerf et al., 2014). This ensured that DAF was not an unusually loud sound for the bird; the distorted part of the song was always softer than the loudest parts of the song.

Target time in the song was defined as the median DAF onset time across renditions of target syllables; jitter of the target time in each bird was defined as the standard deviation of the distribution of DAF onset times relative to the target syllable onset. Syllable truncations following DAF events were rare and were excluded from analysis.

Electrophysiology—Neural signals were band pass filtered (0.25-15 kHz) in homemade analog circuits and obtained at 40 kHz using custom Matlab software. Single neurons were identified as Area X-projecting (VTax) by antidromic identification (stimulation intensities 50–400 μ A, 200 μ s on the bipolar stimulation electrode in Area X). All neurons identified as VTax were additionally validated by antidromic collision testing. Spike widths were calculated as the trough-to-peak interval in the mean spike waveform.

Histological verification of electrode location was performed after each experiment by making small, electrolytic lesions (30uA for 60 s) with the recording electrodes. Histological confirmation of reference lesions among dopamine neurons was then made by first fixing the brains, then cutting the brains into 100 μ m thick sagittal sections and immuno-staining them with tyrosine hydroxylase. Anatomical location was used to classify neurons as VTA, and antidromic identification and collision testing was used to classify VTax neurons. Note that antidromic testing can produce false negatives, so it is possible that non-antidromically identified VTAerror neurons in fact project to Area X.

Spike sorting and analyzing responses to distorted auditory feedback—Offline spike sorting was performed using custom Matlab software. Firing rate histograms were constructed with 25 ms bins and smoothed with a 3-bin moving average. To calculate whether error responses were significant (Figure 1D), spiking activity within ± 1 second relative to target onset was binned in a 30 ms, moving window in 2 ms steps. Each bin after the target was tested against the bins in the entire previous 1 second using a z-test (Mandelblat-Cerf et al., 2014). Response onset (latency) was defined as the first bin for which the next 3 consecutive bins (6 ms) were significantly different from the previous 1 second of activity (z-test, $p < 0.05$); response offset was defined as the first bin after response onset for which the next 7 consecutive bins (14 ms) did not differ from the prior activity ($p > 0.05$, z-test); the response onset and offset were needed to set the maximum (undistorted) or minimum (distorted) response after target time.

Parameterizing song—We used open-source MATLAB software, Sound Analysis Pro 2011 (SAP 2011), to assemble spectrograms and to define and extract song features. Given the limited size of our data, we were unable to use an unsupervised method to determine relevant song features. We therefore used an existing feature set in SAP 2011, a customized software package for analysis of animal communication developed to study bird song (Tchernichovski et al., 2000) for our parameterization. These features have been previously used to connect song variations to spiking activity or neuromodulator concentrations (Kao et al., 2005; Leblois et al., 2010; Woolley and Kao, 2014), to study variation in song over development (Deregnacourt et al., 2004; Lipkind and Tchernichovski, 2011; Ravbar et al., 2012) and to drive adult learning in DAF paradigms (Andalman and Fee, 2009; Sober and Brainard, 2009; Tumer and Brainard, 2007). Therefore, we can use this form of dimensionality reduction knowing that these features are behaviorally relevant to song variation in other contexts. The song features used were Wiener entropy, pitch, goodness of pitch, amplitude, amplitude modulation (AM), frequency modulation (FM), mean frequency, and aperiodicity. These features create an eight-dimensional representation of song at each time-step. We further applied a moving-average filter (35 ms) to smooth the feature signals in time and sampled the smoothed value every 5 ms across song.

Aligning syllables across renditions—To compare song across renditions, syllables were classified using custom Matlab code (Gadagkar et al., 2016). Groups of unique syllables were labelled alphabetically as ‘a’, ‘b’, ‘c’ etc. depending on their order within a rendition. The number of unique syllables each bird sings differs bird-to-bird from 3-7 syllables. We identified syllable onsets and offsets across renditions for every syllable set for which there were at least 15 renditions of that syllable. We used an amplitude threshold chosen to match the amplitude variance of that syllable. All alignments were then checked by eye. When alignment was ambiguous by eye, renditions were excluded from analysis.

All syllable types (i.e. ‘a’ or ‘b’ etc.) were isolated and aligned across renditions by syllable onset times. Individual syllable types have a stereotyped, consistent duration; however, there is slight variation of this duration from rendition-to-rendition. In order to make sure that small differences in syllable durations were not misaligning local syllable features at the later portions of the syllable, we linearly time-warped the feature wave forms of each syllable so that they all lasted the median duration of that syllable type (Kao et al., 2008).

Parameterizing and aligning spiking activity—Spike sorting was performed offline using custom MATLAB software (Gadagkar et al., 2016). For each syllable, we included the spike train \pm 500 ms around the syllable onset. To align spiking activity to the song features, we first applied the same linear time-warping map to the spike times that we used to align syllables for each rendition (Kao et al., 2008). In all cases, we applied this map to the time window in which the syllable occurred. When possible, we built a piece-wise linear time warping map using syllable boundaries in surrounding syllables. In the time windows where there was no song with which to generate a time warping map, we left the spike train un-warped.

We binned spike counts within a sliding window (100 ms) across the 1000 ms length of spike train we considered for each syllable. We selected this window based on the firing rate of the VTAerror neurons (mean firing rate = 13 ± 5 Hz).

Fitting spikes to song with a Gaussian process regression—The goal of our analysis is to quantify non-stationary spiking responses to a time-varying sensory signal, with the following characteristics. First, if VTA neuron activity encodes prediction error responses to song fluctuations, these responses would be specific to the time in song; an identical vocalization occurring at the beginning of the song might elicit a very different response than at the middle. Second, the relevant dimensions of the signal space could vary throughout the song; thus, different parameterizations of the song might provide better, low-dimensional representations of error-relevant song variation at different song time-steps. Third, the form of a PPE-like tuning curve could also vary across the song.

Towards this goal, we used a regression approach to determine if spike counts are related to the variation in song (Aljadeff et al., 2016). The relationship between spike counts and song is likely non-linear and related to a variable number of features depending on the point in song. To address this, we used a non-parametric Gaussian process (GP) regression to fit the relationship between our eight song features and spike counts within single time windows (Williams and Rasmussen, 2006).

There are multiple sources of model uncertainty in this task: it is unclear which and how many features to use at a given point in song. Furthermore, the prediction of the model depends heavily on which features are used. To address this uncertainty, we used a Bayesian model averaging approach to determine the predicted spike count wherein we integrated over all possible feature combinations and weighted their predictions according to their posterior probability given the observed spike counts (Hoeting et al., 1998).

For each neuron, in every non-target (distorted or undistorted) syllable for which there were $N = 15$ renditions, we sampled the smoothed song features every 5 ms across the syllable and sampled spike counts in 100 ms windows every 10 ms across 1s of spike train centered around syllable onset. We fit the multi-dimensional GP model across all song segment-spike bin pairs and generated song-spike relationships at many time latencies. We additionally fit a GP model using each feature individually. For all of these fits we computed the r^2 value (coefficient of determination).

Construction of the Gaussian process regression model—We modeled the relationship between the set of N , z-scored song features on a single rendition i , x_i , and the spike counts on that given rendition, y_i , in single time windows (e.g. the song feature values 20 ms after syllable onset and the spike count in a 100 ms window, 75 ms after syllable onset). We used a non-parametric Gaussian process (GP) regression to fit the relationship between the eight song features and spike counts across song and spike window pairs (Williams and Rasmussen, 2006). We used a Bayesian model averaging approach to combine a weighted average of GP regressions using all subsets of song features into a single model prediction. Note that the full GP model can result in worse predictions than individual features since the full model takes a weighted average across all feature

combinations. While the model learns to shrink away unrelated features, the model won't completely shrink features (exactly zero weights) with finite data due to remaining posterior uncertainty. If only one feature is truly related to the spiking activity, this will impact the final weighted average. This extra stringency enables us to produce predictive fits across each fold in our cross-validation procedure using a single model (the GP model), without requiring a complicated cross-validation procedure to coherently select a single best feature model and correctly evaluate withheld performance. We therefore expect that occasionally a single feature model will do better than the full-feature model in the limited data regime, but without knowing which features to select a priori the single-feature fits may be overly optimistic due to the problem of multiple comparisons.

We selected a subset of features, \mathcal{M} , for a single GP regression, where feature is indexed by $n = 1, 2, \dots, N$ such that $\mathcal{M} \subseteq \{1, 2, \dots, N\}$, $\mathcal{M} \neq \emptyset$, $N = 8$. The GP regression model for a single set of \mathcal{M} is:

$$y_i | f_{\mathcal{M}} \sim \mathcal{N}(f_{\mathcal{M}}(\mathbf{x}_i, \mathcal{M}), \sigma^2) \quad (\text{Equation 1})$$

$$f_{\mathcal{M}}(\cdot) \sim \mathcal{GP}(\mu, \omega^2 \kappa^{(\mathcal{M})}(\mathbf{x}, \mathbf{x}')) \quad (\text{Equation 2})$$

where f is a function that relates song features to spike rate and κ is the covariance function and defines how spike counts will correlate with one another in feature space. We used the commonly selected kernel function for κ ,

$$\kappa^{(\mathcal{M})}(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x}_{\mathcal{M}} - \mathbf{x}'_{\mathcal{M}}\|^2}{2l^2}\right). \quad (\text{Equation 3})$$

ω^2 is the GP variance term and specifies how strongly the spike counts vary as a function of the song features. σ^2 is the variance term that describes noise in the spike count (i.e. the variation of spike counts at a single point in song feature space). The length scale, l , sets how close points must be in feature space to have correlated spike counts. In order to reduce computational complexity, we set $l = 0.5$, for all model fits and z-scored individual song features at each time step.

Although we do not observe f directly, the GP framework allows us to compute the marginal likelihood of the data with respect to the model parameters. The likelihood of all T renditions of spike count-song segment pairs is:

$$p(y_{1:T} | \mathbf{x}_{1:T}, \mathcal{M}, l, \sigma^2, \omega^2, \mu) = \mathcal{N}([y_1, y_2, \dots, y_T]^\top; [\mu, \mu, \dots, \mu]^\top, \omega^2 \mathbf{K}^{(\mathcal{M})} + \sigma^2 \mathbf{I}_T); \quad (\text{Equation 4})$$

$$\mathbf{K}^{(\mathcal{M})} = \begin{bmatrix} \kappa^{(\mathcal{M})}(\mathbf{x}_1, \mathbf{x}_1) & \dots & \kappa^{(\mathcal{M})}(\mathbf{x}_1, \mathbf{x}_T) \\ \vdots & \ddots & \vdots \\ \kappa^{(\mathcal{M})}(\mathbf{x}_T, \mathbf{x}_1) & \dots & \kappa^{(\mathcal{M})}(\mathbf{x}_T, \mathbf{x}_T) \end{bmatrix}, \quad (\text{Equation 5})$$

where I_T is the identity matrix of dimension T .

The prediction mean-squared error for the GP model is:

$$MSE_{loo}^{(GP)} = \frac{1}{T} \sum_{i=1}^T (y_i - \hat{y}_i)^2 \quad (\text{Equation 6})$$

$$\hat{y}_i = \mathbb{E}[y_i | \mathbf{x}_{/i}, \mathbf{y}_{/i}, \mathbf{x}_i] \quad (\text{Equation 7})$$

where \hat{y}_i is the predicted spike count from the model for rendition i , and $\mathbf{x}_{/i}$ and $\mathbf{y}_{/i}$ are the song features and spike counts for all renditions except for the i^{th} rendition.

We determined the predicted spike count by applying a Bayesian model averaging approach. We integrated over all possible values of \mathcal{M} and then weighted their predictions based on their posterior probability given observed spike counts (Hoeting et al., 1998):

$$\mathbb{E}[y_i | \mathbf{x}_{/i}, \mathbf{y}_{/i}, \mathbf{x}_i] = \sum_{\mathcal{M} \subseteq \{1, 2, \dots, N\}, \mathcal{M} \neq \emptyset} \int_0^\infty \mathbb{E}[y_i | \mathcal{M}, r, \mathbf{x}_{/i}, \mathbf{y}_{/i}, \mathbf{x}_i] p(\mathcal{M}, r | \mathbf{x}_{/i}, \mathbf{y}_{/i}) dr, \quad (\text{Equation 8})$$

where $r = \alpha^2/\omega^2$ is the ratio of the GP variance to the observation noise. We then integrated over all possible values of \mathcal{M} and weighted their predictions according to their posterior probability given the observed spike counts.

$$\mathbb{E}[y_i | \mu, \mathcal{M}, r, \mathbf{x}_{/i}, \mathbf{y}_{/i}, \mathbf{x}_i] = \mathbf{K}_{/i,i}^{(\mathcal{M})\top} \left(\mathbf{K}_{/i,/i}^{(\mathcal{M})} + r\mathbf{I}_{(T-1)} \right)^{-1} (\mathbf{y}_{/i} - \mu) + \mu. \quad (\text{Equation 9})$$

In this way, we incorporated all possible combinations of song features into a single model prediction for each song-spike count pair. We re-parameterized (σ^2, ω^2) to (ψ^2, r^2) where ψ^2 is the total variance:

$$\psi^2 = \sigma^2 + \omega^2, \quad (\text{Equation 10})$$

$$\sigma^2 = \frac{r}{r+1} \psi^2, \quad \omega^2 = \frac{\psi^2}{r+1} \quad (\text{Equation 11})$$

And evaluated the posterior over model parameters using Bayes' rule:

$$p(\mathcal{M}, r | \mathbf{x}_{/i}, \mathbf{y}_{/i}) = \frac{p(\mathbf{y}_{/i} | r, \mathcal{M}, \mathbf{x}_{/i}) p(\mathcal{M}, r)}{\sum_{\mathcal{M}^* \subseteq \{1, 2, \dots, N\}} p(\mathcal{M}^*) \int_0^\infty p(\mathbf{y}_{/i} | r^*, \mathcal{M}^*, \mathbf{x}_{/i}) p(\mathcal{M}^*, r^*) dP(r^*)}. \quad (\text{Equation 12})$$

We again applied Bayes' rule to compute the likelihood term in Equation 12:

$$p(\mathbf{y}_{/i} | r, \mathcal{M}, \mathbf{x}_{/i}) = \frac{p(\mathbf{y}_{/i} | \mu, \psi^2, r, \mathcal{M}, \mathbf{x}_{/i})p(\mu, \psi^2)}{p(\mu, \psi^2 | r, \mu, \mathbf{x}_{/i}, \psi^2, \mathcal{M})}. \quad (\text{Equation 13})$$

The likelihood term is computed as in Equation 4. We again used Bayes' rule to compute the posterior over μ and ψ^2 :

$$p(\mu, \psi^2 | r, \mu, \mathbf{x}_{/i}, \psi^2, \mathcal{M}) \propto p(\mathbf{y}_{/i} | \mu, \psi^2, r, \mathcal{M}, \mathbf{x}_{/i})p(\mu, \psi^2). \quad (\text{Equation 14})$$

We placed a conjugate normal-inverse gamma prior over μ and ψ^2 :

$$(\mu, \psi^2) \sim N_{-}\Gamma^{-1}(\mu_0, \lambda_0, \alpha_0, \beta_0) \quad (\text{Equation 15})$$

where,

$$\mu_0 = 0; \lambda_0 = 1; \alpha_0 = 10; \beta_0 = \alpha_0 + 1. \quad (\text{Equation 16})$$

Thus,

$$(\mu, \psi^2 | r, \mu, \psi^2, \mathcal{M}, \mathbf{x}_{/i}) \sim N_{-}\Gamma^{-1}(\mu_{post}^{(i)}, \lambda_{post}^{(i)}, \alpha_{post}^{(i)}, \beta_{post}^{(i)}), \quad (\text{Equation 17})$$

where,

$$\mu_{post}^{(i)} = \frac{b^{(i)}}{a^{(i)}}, \quad (\text{Equation 18})$$

$$\lambda_{post}^{(i)} = a^{(i)}, \quad (\text{Equation 19})$$

$$\alpha_{post}^{(i)} = \alpha_0 + \frac{T-1}{2}, \quad (\text{Equation 20})$$

$$\beta_{post}^{(i)} = \frac{1}{2} \left(c^{(i)} - \frac{b^{(i)}}{a^{(i)}} \right), \quad (\text{Equation 21})$$

$$a^{(i)} = (r+1) \mathbf{1}^T \left(K_{/i, /i}^{(\mathcal{M})} + rI \right)^{-1} \mathbf{1} + \lambda_0, \quad (\text{Equation 22})$$

$$b^{(i)} = (r+1) \mathbf{1}^T \left(K_{/i, /i}^{(\mathcal{M})} + rI \right)^{-1} \mathbf{y}_{/i}, \quad (\text{Equation 23})$$

$$c^{(i)} = (r + 1) \mathbf{1}^\top \left(\mathbf{K}_{/i, /i}^{(\mathcal{M})} + r \mathbf{I} \right)^{-1} \mathbf{y}_{/i} + 2\beta_0, \quad (\text{Equation 24})$$

where $\mathbf{1}$ is a vector of ones. With this, we compute the terms in Equation 13.

The integral over Equation 12 is over one dimension and therefore tractable to compute. We selected a discrete distribution for the prior $P(r)$ to improve computation speed:

$$P(r) = \text{Uniform}(\{3, 4, 5.67, 9\}), \quad (\text{Equation 25})$$

such that the GP model could account for 25%, 20%, 15% or 10% of the total variance.

We set a truncated binomial prior over the number of included features that privileged models with fewer features (i.e., sparse models):

$$p(\mathcal{M}) = \frac{1}{1 - (1 - p)^N} \binom{N}{|\mathcal{M}|} p^{|\mathcal{M}|} (1 - p)^{N - |\mathcal{M}|}. \quad (\text{Equation 26})$$

We set $p = 0.1$ such that approximately 2/3 of the prior probability mass rests on single-feature models. We could integrate over a sparse prior in our model, rather than a shrinkage prior such as the Lasso, because we considered only a small ($N = 8$) set of features (Park and Casella, 2012). Using this normal inverse-gamma description of the posterior, we can then compute the prediction of y_i given \mathcal{M} and r :

$$\mathbb{E}[y_i \mid \mathcal{M}, r, \mathbf{x}_{/i}, \mathbf{y}_{/i}, \mathbf{x}_i] = \mathbb{E}[\mathbb{E}[y_i \mid \mu, \mathcal{M}, r, \mathbf{x}_{/i}, \mathbf{y}_{/i}, \mathbf{x}_i] \mid \mathcal{M}, r, \mathbf{x}_{/i}, \mathbf{y}_{/i}, \mathbf{x}_i], \quad (\text{Equation 27})$$

$$= \mathbf{K}_{/i, i}^{(\mathcal{M}) \top} \left(\mathbf{K}_{/i, /i}^{(\mathcal{M})} + r \mathbf{I}_{(T-1)} \right)^{-1} (\mathbf{y}_{/i} - \mu_{post}^{(i)}) + \mu_{post}^{(i)}. \quad (\text{Equation 28})$$

We then insert Equation 28 and Equation 12 into Equation 8 to obtain the prediction of y_i .

CONSTRUCTION OF THE LATENCY DISTRIBUTION

We defined the latency distribution as the set of all latencies between spike bins and song feature windows in which there was a predictive relationship ($r^2 > 0$) within the GP model.

CHARACTERIZING TUNING CURVES OF CELL RESPONSES

The GP model is flexible: it will fit any relationship between the independent and dependent variables and additionally is computationally efficient. However, the GP model provides no easily interpretable means of characterizing the shape of the fit. To characterize the form of the spike-count to song relationships across many fits, we needed an automated method to categorize the shapes of the tuning curves.

To do this, we applied a generalized linear model (GLM) with (q-GLM) and without (l-GLM) a quadratic transformation of the song features (Park et al., 2013). A GLM is

comprised of a stimulus filter, an invertible non-linearity (the link function) and a stochastic exponential non-linearity, such as a Poisson process:

$$y | \mathbf{x} \sim \text{Pois}(f(Q(\mathbf{x}))) \quad (\text{Equation 29})$$

$$Q(x) = ax^2 + bx + c, \quad (\text{Equation 30})$$

where y is the spike count, x is the stimulus, f is the inverse link function and Q is a quadratic function of x with coefficients a , b , c . We considered song features separately in this tuning curve analysis, so $\dim(\mathbf{x}) = 1$, and the quadratic coefficients were scalars (a , b , c). We chose the link function to be an exponential and the noise process to be Poisson. We fit the quadratic coefficients by maximizing the log-likelihood:

$$\log P(Y | X, a, b, c) = \sum_{i=1}^N [-\exp(Q(x_i)) + a * y_i x_i^2 + b * y_i x_i + c - \log(y_i!)], \quad (\text{Equation 31})$$

We maximized the log-likelihood numerically using conjugate gradient methods. The sign of the quadratic coefficient, a , of this model specifies whether the data is better fit by an upwards-facing, quadratic basis in which the data is double-peaked, or a downwards-facing quadratic basis in which the data is single-peaked. We compared this model to a nested model fit where the quadratic term is set to zero (l-GLM).

We compared the performance of the two models using the Akaike information criterion (AIC) (Akaike, 1974). The AIC metric is defined as:

$$AIC = 2k - 2 \ln \widehat{\mathcal{L}}, \quad (\text{Equation 32})$$

$$\widehat{\mathcal{L}} = \text{argmax}_{a, b, c} \log P(Y | X, a, b, c), \quad (\text{Equation 33})$$

where $\widehat{\mathcal{L}}$ is the maximum of the log-likelihood function for a given model and k is the number of estimated parameters in the model. This metric incorporates both goodness of fit and model complexity. A lower AIC metric indicates better performance. Therefore, the difference in the AIC metrics of two models denotes the relative success of one model over another, adjusting for differences in model complexity (Akaike, 1974; Raftery, 1995). We can then ask, when the quadratic model (q-GLM) is a better fit than the linear model (l-GLM), does the tuning curve form of spike counts to song features have a positive or negative curvature?

We compared the q-GLM and l-GLM models on all GP model fits with $r^2 > 0$ for all song features which individually had predictive fits within the multi-dimensional model. In this model comparison, $AIC = -2$ is the lower bound of the metric and indicates that the quadratic term of the q-GLM = 0. Thus, no improvement in fit was gained by the quadratic complexity. $AIC > 0$ indicates that the q-GLM outperforms the l-GLM, taking into account both the goodness of fit and the added complexity of the extra model parameter. We used

this comparison to identify potential curvature in the spike-song feature relationship. Note, that when the AIC metric is used in model selection, $AIC > 10$ is the standard threshold for significance. However, we are using this metric not to make a model selection, but rather to identify instances of possible curvature in the data. Thus, we use $AIC > 0$ as our threshold. In our data at $AIC > 10$, 90% of our fits have a negative, quadratic coefficient. However, because our data is noisy, this accounts for only 2% of our data. We calculated the fraction of fits with the quadratic coefficient, a < 0 , as a function of the AIC metric:

$$\Delta AIC \equiv 2k_{linear} - 2 \ln \widehat{\mathcal{L}}_{linear} - 2k_{quad} + 2 \ln \widehat{\mathcal{L}}_{quad}. \quad (\text{Equation 34})$$

QUANTIFICATION AND STATISTICAL ANALYSIS

Evaluating the Gaussian process model performance—To evaluate the GP model performance we estimated the mean-squared prediction error for new observations using a leave-one-out cross validation method as in (Vehtari et al., 2017):

$$MSE_{loo}^{(GP)} = \frac{1}{T} \sum_{i=1}^T (y_i - \hat{y}_i)^2, \quad (\text{Equation 35})$$

$$\hat{y}_i = \mathbb{E} [y_i | \mathbf{x}_{/i}, \mathbf{y}_{/i}, \mathbf{x}_i], \quad (\text{Equation 36})$$

where \hat{y}_i is the model prediction spike count for rendition ‘i’, and $\mathbf{x}^{/i}$ and $\mathbf{y}^{/i}$ are the song features and spike counts of all renditions excluding the i^{th} rendition. We then compared the GP model to a model with constant mean firing rate equal to the mean spike count over all renditions excluding the i^{th} rendition:

$$y_{(i)} \sim \mathcal{N}(\alpha, \tau^2). \quad (\text{Equation 37})$$

The predictive mean-squared error of this model is:

$$MSE_{loo}^{(null)} = \frac{1}{T} \sum_{i=1}^T (y_i - \overline{y}_{/i})^2 \quad (\text{Equation 38})$$

where,

$$\overline{y}_{/i} = \frac{1}{T-1} \sum_{j \neq i} y_j. \quad (\text{Equation 39})$$

The cross-validated r^2 value is:

$$r^2 = 1 - \frac{MSE_{loo}^{(GP)}}{MSE_{loo}^{(null)}}. \quad (\text{Equation 40})$$

An $r^2 > 0$ indicates that the model predicts new observations better than simply using the mean to predict new observations. The maximum value the r^2 can take is one—this indicates perfect model prediction and, in practice, is never reached. The r^2 value is our measure of model performance.

Bootstrapping to assess population-level significance—Evaluating the significance of the model predictions must be done on a population level for this analysis. We computed model fits to hundreds of song segment-spike count pairs for each syllable. Simply by chance, some of these fits would generate a predictive $r^2 > 0$ value.

In addition, spike-song pairs are correlated, not just because spike counts and song segment windows overlap but also because of possible underlying correlations in the song and spike fluctuations across the song. To address this, we randomized the relationship between full spike trains and song renditions and then re-fit our model on the randomized, spike count-song segment pairs. By leaving the timing of the song and spiking activity intact and only randomizing the relationship between them, we constructed a randomized population of fits for each cell-syllable pair, which retained whatever underlying temporal structure was present in the spike trains and song (Tusher et al., 2001). We then repeated this process 100 times for the VTA_{error} cell population and 100 times for the VTA_{other} cell population.

From this distribution of coherently randomized cell sets, we computed bootstrap test, p value assessments of the r^2 values of the individual spike count-song segment model fits as well as on population measures of significance in the VTA_{error} and VTA_{other} cell populations independently. We assessed four population measures:

The number of predictive signals across the whole cell population: An $r^2 > 0$ indicates that the model predicts the data better than an estimate based solely on the mean spike count, and we label this a ‘predictive signal’. We therefore computed the significance of the total number of $r^2 > 0$ song segment-spike count fits within the PPE latency window for both the VTA_{error} ($p < 0.01$) population and VTA_{other} ($p < 0.01$) population of syllable-cell pairs with a one-sided bootstrap test.

The distribution of the predictive signal across the cell population: We asked whether a small number of cell-syllable pairs contained the majority of the positive r^2 values or if the signal appeared across multiple cell-syllable pairs in the population. To answer this, we first labeled each cell-syllable pair as ‘significant’ if the number of fits with positive r^2 values within the PPE latency window (0–150 ms) had a single-tailed $p < 0.05$. We then computed the single-tailed p value for the number of significant cell-syllable pairs across each cell population (VTA_{error} population: one-sided bootstrap test: $p < 0.01$; VTA_{other} population: one-sided bootstrap test: $p < 0.01$).

The magnitude of the peak in signal occurrence within the PPE latency window across the cell population: We computed latency distributions as the latencies of the full set of spike-count song feature pairs that resulted in GP model fits with $r^2 > 0$. We compared the variance within the latency distributions of the randomized populations to the magnitude of the peak we found in the actual data. We computed the single-tailed p value for the

maximum fluctuation of a latency distribution at any point in the latency domain. Thus, we tested the significance not only of finding a peak of that size in the data at the PPE window but of finding a peak of that size anywhere in the latency distribution. The VTA_{error} population latency distribution had a peak within the expected PPE latency region (Figure 3C). This peak was 3.74 standard deviations from the mean. The variance in relation to the maximum variance in randomized latency distributions was significant (one-sided bootstrap test: $p < 0.01$). The time-bin of this latency distribution was 25 ms. The VTA_{other} population latency distribution had a peak 2.30 standard deviations from the mean (Figure 3D). This variance was not significant (one-sided bootstrap test: $p = 0.21$). The time-bin of this latency distribution was 25 ms.

The curvature of tuning curves we find via our GLM parameterization technique: We computed the two-sided bootstrap test p value of the fraction of tuning curves with a negative quadratic coefficient in the population data relative to the randomized populations (Figure 4E legend and main text).

Note that the goal of this significance strategy allows us to assess the VTA_{error} cell activity as a population, not the significance of particular song segment-spike count pairs. In this framework, we do not require that a single feature or even song-spike pattern be significant. More data are needed for this level of specificity.

Assessing natural song relationship to distorted auditory feedback—We have been unable to find any impact of the presence or absence of a distortion event in other portions of the song. As reported in Gadagkar et al., (2016), there is no significant correlation between VTA spiking activity and the distortion event outside of a 0–150 ms lag. We additionally looked for correlations between the distortion event and natural song fluctuations in the syllable immediately following the distortion event (when the greatest impact might be expected) and found no significant correlations there either.

As an additional control, we re-performed the analyses presented in Figures 3 and 4 on only the sections of song which occurred before the distorted auditory feedback across all error cells. During this portion of the behavior, it is unknown whether a distortion will occur later in the song, and thus, no hidden correlations can exist between song fluctuations and distortion events. In this subset of the data, we find a robust recapitulation of our original results: there is a highly significant peak in the latency distribution of predictive relationships between spike counts and song with the expected RPE timing latency (0–150 ms) (p value < 0.01) and a highly significant number of predictive relationships (p value < 0.01), repeating results in Figure 3. Additionally, there is a significant over-representation of stabilizing tuning curve relationships between DA neurons and song fluctuations (p value < 0.02), again recapitulating the results of Figure 4.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

This work was supported by NIH/NINDS (grant no. R01NS094667), Pew Charitable Trusts, and the Klingenstein Neuroscience Foundation (to J.H.G.); a Simons Collaboration on the Global Brain (SCGB) grant (grant no. 542975SPI) and NSF (grant no. IIS 1421245) (to A.L.F. and A.D.); an SCGB Postdoctoral Fellowship and NIH/NINDS Pathway to Independence award (grant no. K99NS102520 and R00NS102520) (to V.G.); and facilitated through the use of the Hyak supercomputing system and funded by the STF at the University of Washington. This work is based on Chapter 5 of author A.D.'s PhD thesis at the University of Washington (Duffy, 2018).

REFERENCES

- Akaike H (1974). A new look at the statistical model identification. *IEEE Trans. Automatic Control* 19, 716–723.
- Ali F, Otchy TM, Pehlevan C, Fantana AL, Burak Y, and Olveczky BP (2013). The basal ganglia is necessary for learning spectral, but not temporal, features of birdsong. *Neuron* 80, 494–506. 10.1016/j.neuron.2013.07.049. [PubMed: 24075977]
- Aljadeff J, Lansdell BJ, Fairhall AL, and Kleinfeld D (2016). Analysis of neuronal spike trains, deconstructed. *Neuron* 91, 221–259. 10.1016/j.neuron.2016.05.039. [PubMed: 27477016]
- Andalman AS, and Fee MS (2009). A basal ganglia-forebrain circuit in the songbird biases motor output to avoid vocal errors. *Proc. Natl. Acad. Sci. U S A* 106, 12518–12523. 10.1073/pnas.0903214106. [PubMed: 19597157]
- Barter JW, Li S, Lu D, Bartholomew RA, Rossi MA, Shoemaker CT, Salas-Meza D, Gaidis E, and Yin HH (2015). Beyond reward prediction errors: the role of dopamine in movement kinematics. *Front. Integr. Neurosci* 9, 39. 10.3389/fnint.2015.00039. [PubMed: 26074791]
- Brainard MS, and Doupe AJ (2000). Interruption of a basal ganglia-fore-brain circuit prevents plasticity of learned vocalizations. *Nature* 404, 762–766. 10.1038/35008083. [PubMed: 10783889]
- Chen R, Gadagkar V, Roeser AC, Puzerey PA, and Goldberg JH (2021). Movement signaling in ventral pallidum and dopaminergic midbrain is gated by behavioral state in singing birds. *J. Neurophysiol* 125, 2219–2227. 10.1152/jn.00110.2021. [PubMed: 33949888]
- Cohen JY, Haesler S, Vong L, Lowell BB, and Uchida N (2012). Neuron-type-specific signals for reward and punishment in the ventral tegmental area. *Nature* 482, 85–88. 10.1038/nature10754. [PubMed: 22258508]
- da Silva JA, Tecuapetla F, Paixão V, and Costa RM (2018). Dopamine neuron activity before action initiation gates and invigorates future movements. *Nature* 554, 244–248. 10.1038/nature25457. [PubMed: 29420469]
- Deregnaucourt S, Mitra PP, Feher O, Maul KK, Lints TJ, and Tchernichovski O (2004). Song development: in search of the error-signal. *Ann. N. Y. Acad. Sci* 1016, 364–376. 10.1196/annals.1298.036. [PubMed: 15313785]
- Dobi A, Margolis EB, Wang HL, Harvey BK, and Morales M (2010). Glutamatergic and nonglutamatergic neurons of the ventral tegmental area establish local synaptic contacts with dopaminergic and nondopaminergic neurons. *J. Neurosci* 30, 218–229. 10.1523/JNEUROSCI.3884-09.2010. [PubMed: 20053904]
- Duffy A (2018). *Neural Mechanisms for Trial and Error Learning: A Study from Bird Song*, PhD (University of Washington).
- Engelhard B, Finkelstein J, Cox J, Fleming W, Jang HJ, Ornelas S, Koay SA, Thiberge SY, Daw ND, Tank DW, and Witten IB (2019). Specialized coding of sensory, motor and cognitive variables in VTA dopamine neurons. *Nature* 570, 509–513. 10.1038/s41586-019-1261-9. [PubMed: 31142844]
- Gadagkar V, Puzerey PA, Chen R, Baird-Daniel E, Farhang AR, and Goldberg JH (2016). Dopamine neurons encode performance error in singing birds. *Science* 354, 1278–1282. 10.1126/science.aah6837. [PubMed: 27940871]
- Goffinet J, Brudner S, Mooney R, and Pearson J (2021). Low-dimensional learned feature spaces quantify individual and group differences in vocal repertoires. Preprint at bioRxiv. 811661. 10.1101/811661.

- Hamaguchi K, Tschida KA, Yoon I, Donald BR, and Mooney R (2014). Auditory synapses to song premotor neurons are gated off during vocalization in zebra finches. *Elife* 3, e01833. 10.7554/eLife.01833. [PubMed: 24550254]
- Hamid AA, Pettibone JR, Mabrouk OS, Hetrick VL, Schmidt R, Vander Weele CM, Kennedy RT, Aragona BJ, and Berke JD (2016). Mesolimbic dopamine signals the value of work. *Nat. Neurosci* 19, 117–126. 10.1038/nn.4173. [PubMed: 26595651]
- Hamilos AE, Spedicato G, Hong Y, Sun F, Li Y, and Assad JA (2020). Dynamic dopaminergic activity controls the timing of self-timed movement. Preprint at bioRxiv. 10.1101/2020.05.13.094904.
- Hisey E, Kearney MG, and Mooney R (2018). A common neural circuit mechanism for internally guided and externally reinforced forms of motor learning. *Nat. Neurosci* 21, 589–597. 10.1038/s41593-018-0092-6. [PubMed: 29483664]
- Hoeting JA, Madigan D, Raftery AE, and Volinsky CT (1998). Bayesian model averaging. *Proc. AAAI Workshop Integrating Mult. Learned Models* 335, 77–83.
- Hoffmann LA, Saravanan V, Wood AN, He L, and Sober SJ (2016). Dopaminergic contributions to vocal learning. *J. Neurosci* 36, 2176–2189. 10.1523/JNEUROSCI.3883-15.2016. [PubMed: 26888928]
- Howe MW, Tierney PL, Sandberg SG, Phillips PE, and Graybiel AM (2013). Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. *Nature* 500, 575–579. 10.1038/nature12475. [PubMed: 23913271]
- Kao MH, Doupe AJ, and Brainard MS (2005). Contributions of an avian basal ganglia-forebrain circuit to real-time modulation of song. *Nature* 433, 638–643. 10.1038/nature03127. [PubMed: 15703748]
- Kao MH, Wright BD, and Doupe AJ (2008). Neurons in a forebrain nucleus required for vocal plasticity rapidly switch between precise firing and variable bursting depending on social context. *J. Neurosci* 28, 13232–13247. 10.1523/JNEUROSCI.2250-08.2008. [PubMed: 19052215]
- Kollmorgen S, Hahnloser RHR, and Mante V (2020). Nearest neighbours reveal fast and slow components of motor learning. *Nature* 577, 526–530. 10.1038/s41586-019-1892-x. [PubMed: 31915383]
- Leblois A, Wendel BJ, and Perkel DJ (2010). Striatal dopamine modulates basal ganglia output and regulates social context-dependent behavioral variability through D1 receptors. *J. Neurosci* 30, 5730–5743. 10.1523/JNEUROSCI.5974-09.2010. [PubMed: 20410125]
- Lipkind D, and Tchernichovski O (2011). Quantification of developmental birdsong learning from the subsyllabic scale to cultural evolution. *Proc. Natl. Acad. Sci. U S A* 108, 15572–15579. 10.1073/pnas.1012941108. [PubMed: 21436035]
- Mandelblat-Cerf Y, Las L, Denisenko N, and Fee MS (2014). A role for descending auditory cortical projections in songbird vocal learning. *Elife* 3, e02152. 10.7554/eLife.02152.
- Park T, and Casella G (2012). The Bayesian Lasso. *J. Am. Stat. Assoc* 103, 681–686. 10.1198/016214508000000337.
- Park IM, Priebe N, and Pillow J (2013). Spectral methods for neural characterization using generalized quadratic models. *Adv. Neural Inf. Process. Syst* 26, 2454–2462.
- Raftery RE (1995). Bayes factors. *J. Am. Statistical Assoc* 90, 773–795.
- Ravbar P, Lipkind D, Parra LC, and Tchernichovski O (2012). Vocal exploration is locally regulated during song learning. *J. Neurosci* 32, 3422–3432. 10.1523/JNEUROSCI.3740-11.2012. [PubMed: 22399765]
- Schultz W, Dayan P, and Montague PR (1997). A neural substrate of prediction and reward. *Science* 275, 1593–1599. 10.1126/science.275.5306.1593. [PubMed: 9054347]
- Sober SJ, and Brainard MS (2009). Adult birdsong is actively maintained by error correction. *Nat. Neurosci* 12, 927–931. 10.1038/nn.2336. [PubMed: 19525945]
- Sutton RS, and Barto AG (1998). *Reinforcement Learning: An Introduction* (MIT Press).
- Tchernichovski O, Nottebohm F, Ho CE, Pesaran B, and Mitra PP (2000). A procedure for an automated measurement of song similarity. *Anim. Behav* 59, 1167–1176. 10.1006/anbe.1999.1416. [PubMed: 10877896]

- Tian J, Cohen JY, Osakada F, Kobak D, Machens CK, Callaway EM, Uchida N, and Watabe-Uchida M (2016). Distributed and mixed information in monosynaptic inputs to dopamine neurons. *Neuron* 91, 1374–1389. [PubMed: 27618675]
- Tumer EC, and Brainard MS (2007). Performance variability enables adaptive plasticity of ‘crystallized’ adult birdsong. *Nature* 450, 1240–1244. 10.1038/nature06390. [PubMed: 18097411]
- Tusher VG, Tibshirani R, and Chu G (2001). Significance analysis of micro-arrays applied to the ionizing radiation response. *Proc. Natl. Acad. Sci. U S A* 98, 5116–5121. 10.1073/pnas.091062498. [PubMed: 11309499]
- Vehtari A, Gelman A, and Gabry J (2017). Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Stat. Comput* 27, 1413–1432.
- Williams CK, and Rasmussen CE (2006). *Gaussian Processes for Machine Learning* (The MIT Press).
- Wood J, Simon NW, Koerner FS, Kass RE, and Moghaddam B (2017). Networks of VTA neurons encode real-time information about uncertain numbers of actions executed to earn a reward. *Front. Behav. Neurosci* 11, 140. 10.3389/fnbeh.2017.00140. [PubMed: 28848408]
- Woolley SC, and Kao MH (2014). Variability in action: contributions of a songbird cortical-basal ganglia circuit to vocal motor learning and control. *Neuroscience* 296, 39–47. 10.1016/j.neuroscience.2014.10.010. [PubMed: 25445191]
- Xiao L, Chattree G, Oscos FG, Cao M, Wanat MJ, and Roberts TF (2018). A basal ganglia circuit sufficient to guide birdsong learning. *Neuron* 98, 208–221.e205. 10.1016/j.neuron.2018.02.020. [PubMed: 29551492]

Highlights

- Songbird dopamine activity correlates with natural song fluctuations
- The form and timing of the dopamine activity is consistent with song evaluation

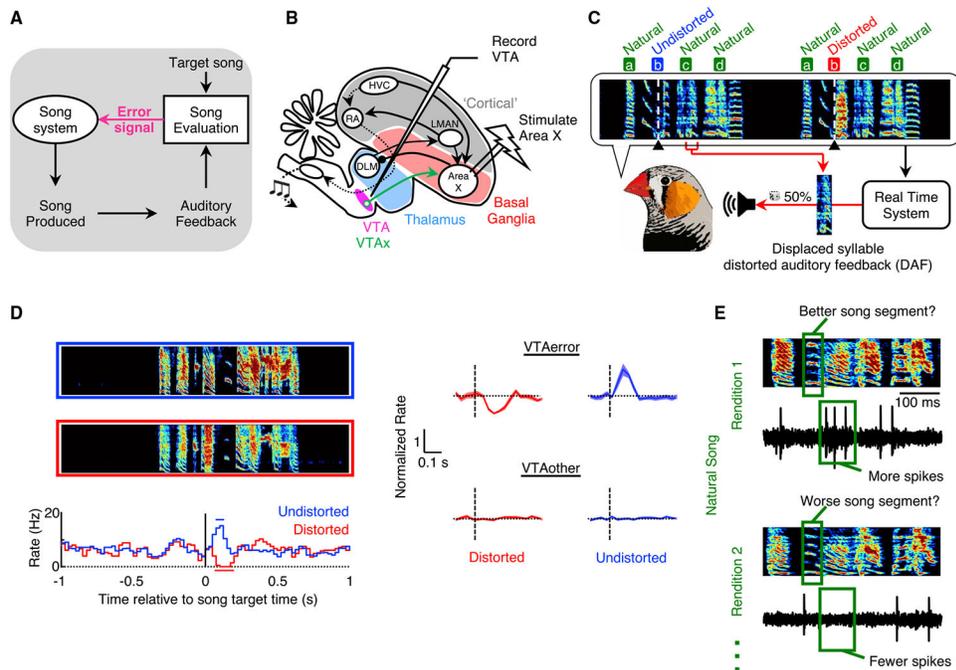


Figure 1. Experimental identification of performance error in VTA DA neurons in singing birds
 (A) Evaluation of auditory feedback during singing is thought to produce an error signal for song learning (reproduced from Gadagkar et al., 2016, with the permission of AAAS).
 (B) Basal ganglia (Area X)-projecting DA neurons from VTA were antidromically identified (reproduced from Gadagkar et al., 2016, with the permission of AAAS).
 (C) Example of DAF. The target syllable was randomly distorted across motifs. All other syllables (labeled “Natural”) were left undisturbed (reproduced from Gadagkar et al., 2016, with the permission of AAAS).
 (D) (Left, top to bottom) Example spectrograms of renditions with the target syllable undistorted (enclosed in blue box) and distorted (enclosed in red box); rate histogram of distorted and undistorted renditions (the horizontal bar indicates significant deviations from baseline [$p < 0.05$, z test; see STAR Methods]); (Right) Normalized response to target syllable in VTA_{error} and VTA_{other} neurons (mean \pm SEM; see STAR Methods) (reproduced from Gadagkar et al., 2016, with the permission of AAAS).
 (E) The experimental results suggest a hypothesis that fluctuations in natural song should also result in VTA_{error} responses.

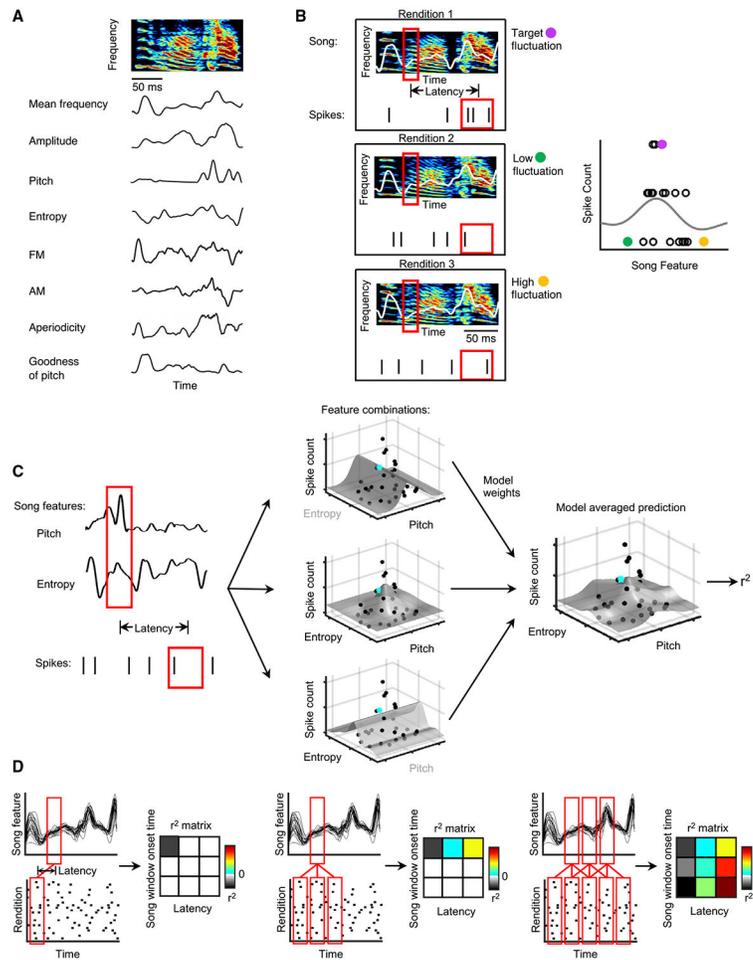


Figure 2. A Gaussian process model approach reveals song-spike relationships

(A) Natural song was parameterized into eight time-varying song features.

(B) Schematic of fitting song fluctuations to spike counts within specific time windows.

Local feature averages (one feature shown for illustration) were used to predict local spike counts using a GP model.

(C) Schematic of fitting a single, multivariate model using multiple song features. The multi-dimensional model takes a weighted average of the model predictions from every combination of eight song features (two shown here for illustration). The middle column shows the three feature combinations for two example features (t-b): pitch only; pitch and entropy; entropy only. The model's goodness of fit was quantified by the cross-validated r^2 calculated from the final, weighted average model (see STAR Methods). The cyan dot indicates an example held out data point in the cross-validation procedure.

(D) Schematic of modeling technique shown in (B and C) now extended across a range of song windows and song-spike latencies, thus building a matrix of r^2 values. The top panels show a sliding window along the song (single feature shown for illustration). The bottom panels show the time-aligned spiking activity across renditions in a raster plot. Each entry in the r^2 matrix (middle panels) represents the fit between one song window and one spike window, shown here connected with red lines.

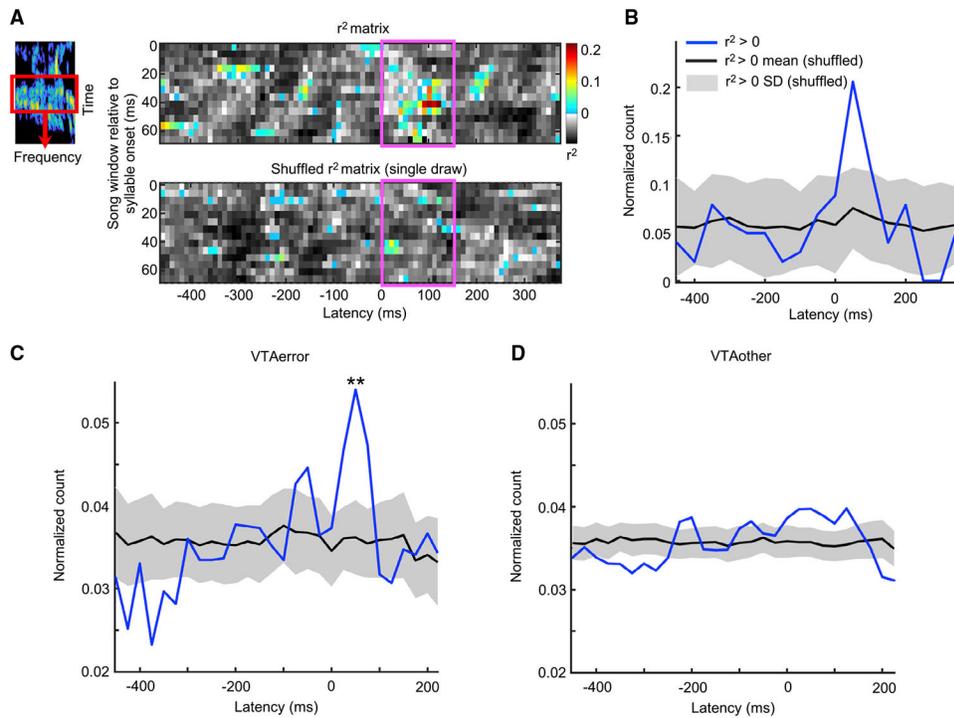


Figure 3. Timing of song-spike relationships for VTA_{error} neurons suggests an evaluative process
 (A) Spectrogram of example syllable (top left). Heatmap of r^2 values for fitted relationships between local song feature averages and binned spike counts (top right). $r^2 > 0$ indicates a predictive relationship. The pink box indicates the region where the latency matches the hypothesized response for a PPE, 0–150 ms. The lower heatmap shows an r^2 matrix for a shuffled version of the data (see STAR Methods).
 (B) Histogram of latencies for predictive fits shown in (A).
 (C) Latency distribution of predictive fits over all VTA_{error} neurons ($n = 22$) showed a significant peak in the number of responses in the expected PPE time window (** $p < 0.01$; see STAR Methods).
 (D) Same as in (C), but for the VTA_{other} neuron population ($n = 23$).

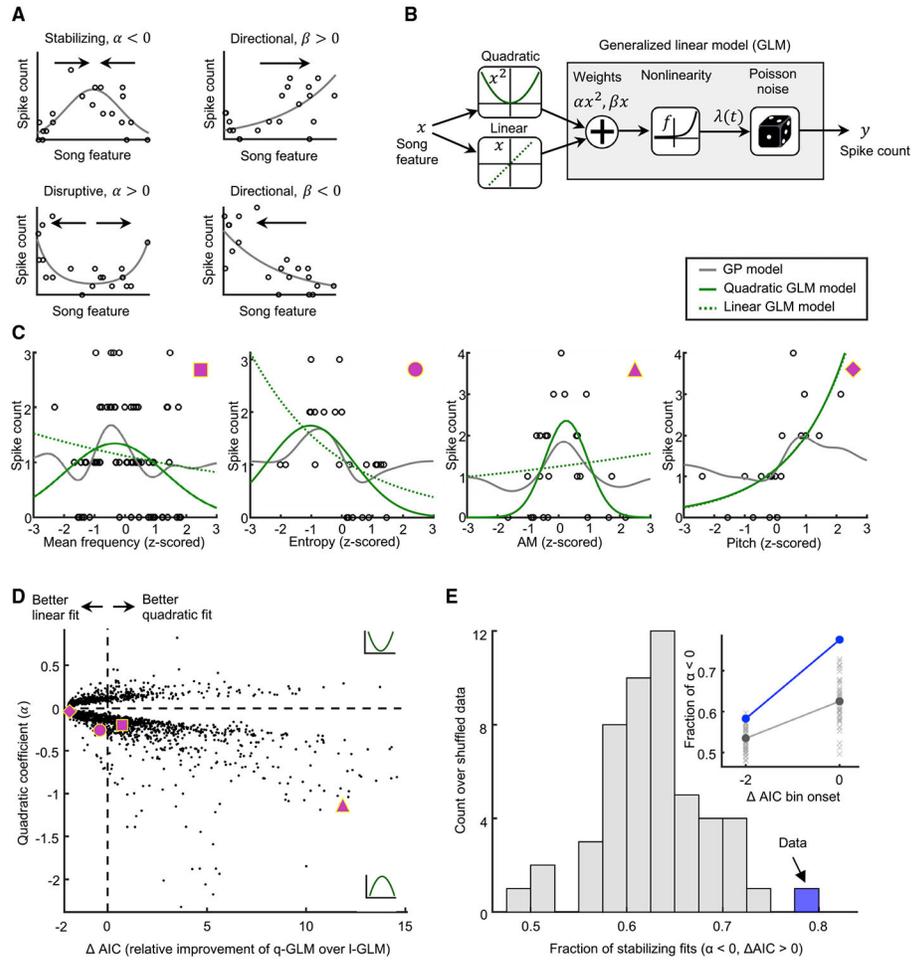


Figure 4. The form of the predominant song-spike relationship for VTA_{error} neurons is consistent with song maintenance

(A) The form of a song-spike relationship determines how the song is being reinforced. α and β correspond to the quadratic and linear coefficients in the GLM shown in (B).

(B) Schematic of the nested GLM fitting process to quantify tuning curve shape for VTA_{error} neuron activity to natural song fluctuations.

(C) Example tuning curves obtained with the GP model, l-GLM, and q-GLM between single song features and spike counts for a selection of song-spike model fits. Each point on each plot represents a single rendition. Pink shapes denote fit locations marked in (D). The r^2 values for each example fit, l-r are: 0.13, 0.27, 0.13, 0.24. Additional, single-feature examples are given in Figure S4.

(D) The quadratic coefficient for q-GLM model fits to predictive song-spike relationships (defined within the GP model) as a function of ΔAIC values in the GLM model comparison within the PPE latency range. Each point represents one q-GLM fit to a significant song feature-spike count pair. Pink shapes denote fits shown in (C). The fraction of total data points in each quadrant about the [0,0] origin, clockwise from top left is: 0.24, 0.09, 0.32, 0.34. This plot zooms in on 99.5% of data. Outlier points follow the same trend but increase scale and obscure visualization. All data are used in analysis.

(E) Fraction of stabilizing fits (negative quadratic coefficient) for all fits better described as quadratic than linear ($AIC > 0$) compared with shuffled population fractions. The blue point is the data and each value in the gray histogram is a single fraction from an independent population shuffle (see STAR Methods). The data showed a greater fraction of stabilizing fits than expected by chance (two-sided bootstrap test: $p < 0.02$; see STAR Methods). Inset: same distribution but now shown for both the binned $AIC > 0$ group and $AIC \in [-2, 0]$. The blue point is the true fraction and gray points are fractions from shuffled populations (see STAR Methods).

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Antibodies		
Anti-Tyrosine Hydroxylase antibody	Millipore	Cat#AB152; RRID:AB_390204
Dextran, Alexa Fluor™ 488	Invitrogen	Cat#D22910
Experimental models: Organisms/strains		
Zebra Finch (<i>Taeniopygia guttata</i>)	Magnolia Bird Farm, Anaheim CA	N/A
Software and algorithms		
MATLAB	MathWorks	https://www.mathworks.com/products/matlab.html
Sound Analysis Pro	Tchernichovski et al., 2000	http://soundanalysispro.com
Original analysis scripts	This paper	https://doi.org/10.6084/m9.figshare.15019380