The Basal Ganglia Is Necessary for Learning Spectral, but Not Temporal, Features of Birdsong

Farhan Ali,[1,2] Timothy M. Otchy,[2,3,7] Cengiz Pehlevan,[2,4,7] Antoniu L. Fantana,[1,2,5] Yoram Burak,[2,4,6] and Bence P. Ölveczky[1,2,*]

[1]Department of Organismic and Evolutionary Biology
[2]Center for Brain Science
[3]Program in Neuroscience
[4]Swartz Program in Theoretical Neuroscience
Harvard University, Cambridge, MA 02138, USA
[5]Present address: Department of Psychology, Bowling Green State University, Bowling Green, OH 43403, USA
[6]Present address: Edmond and Lily Safra Center for Brain Sciences and Racah Institute of Physics, The Hebrew University, Jerusalem 91904, Israel
[7]These authors contributed equally to this work
*Correspondence: olveczky@fas.harvard.edu
http://dx.doi.org/10.1016/j.neuron.2013.07.049

SUMMARY

Executing a motor skill requires the brain to control which muscles to activate at what times. How these aspects of control—motor implementation and timing—are acquired, and whether the learning processes underlying them differ, is not well understood. To address this, we used a reinforcement learning paradigm to independently manipulate both spectral and temporal features of birdsong, a complex learned motor sequence, while recording and perturbing activity in underlying circuits. Our results uncovered a striking dissociation in how neural circuits underlie learning in the two domains. The basal ganglia was required for modifying spectral, but not temporal, structure. This functional dissociation extended to the descending motor pathway, where recordings from a premotor cortex analog nucleus reflected changes to temporal, but not spectral, structure. Our results reveal a strategy in which the nervous system employs different and largely independent circuits to learn distinct aspects of a motor skill.

INTRODUCTION

To master a motor skill, both its timing and specific motor implementation must be learned and adaptively refined. Increasing the power of your tennis serve, for example, might mean speeding up certain parts of the service motion (modifying timing), while adding top spin might require changing the angle of your elbow (modifying motor implementation). Both improvements will require changes to the motor program underlying your serve, but the nature of these changes can be construed as different.

Modifying timing equates to changing the temporal progression of the muscle activity patterns to slow down or speed up certain parts of the action, whereas changing motor implementation means modifying specific muscle commands while maintaining the temporal dynamics of the action (Figures 1A–1C). Whether this conceptual distinction reflects a dissociation in how the motor system learns and refines motor skills has not been explored.

The zebra finch, a songbird, provides a unique model system for addressing this question. Through a process that resembles human speech learning (Doupe and Kuhl, 1999), juvenile zebra finches gradually improve both temporal (Glaze and Troyer, 2012; Lipkind and Tchernichovski, 2011) and spectral (Tchernichovski et al., 2001) aspects of their songs (Figures 1D–1F) until they resemble those of their tutors (Immelmann, 1969). Spectral features of song are largely determined by the activity of vocal muscles (Goller and Suthers, 1996) and thus serve as a proxy for "motor implementation."

The neural circuit architecture underlying song production is well delineated (Figure 1G) and suggests a hierarchical organization (Yu and Margoliash, 1996) with a descending motor cortical pathway that encompasses premotor nucleus HVC (proper name) (Vu et al., 1994) and motor cortex analog robust nucleus of the arcopallium (RA) (Nottebohm et al., 1982). RA projection neurons synapse onto brainstem motor neurons involved in singing (Wild, 1993). HVC and RA are also indirectly connected through the anterior forebrain pathway (AFP), a basal ganglia-thalamo-cortical circuit that is critical for song learning but not essential for producing learned song (Figure 1G) (Bottjer et al., 1984; Scharff and Nottebohm, 1991). A separate basal ganglia circuit, medial to the AFP, receives input from and provides output to HVC (Foster et al., 1997; Kubikova et al., 2007; Williams et al., 2012) (Figure 6A), but the role of this circuit in song learning, if any, remains to be elucidated (Foster and Bottjer, 2001).

The analogies and homologies between the AFP and basal ganglia circuits in mammals (Farries and Perkel, 2002; Reiner et al., 2004) have made the songbird a tractable model for
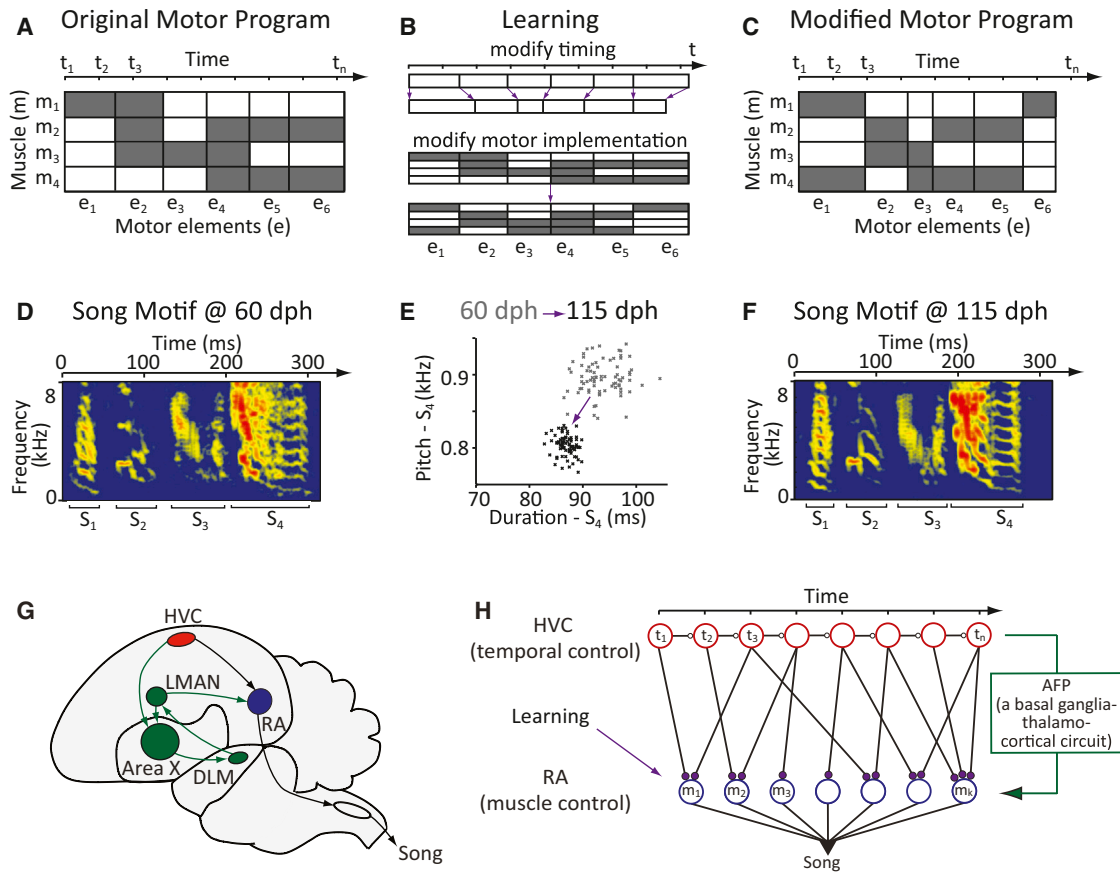
Figure 1. Using Songbirds to Test whether the Nervous System Distinguishes Learning in the Temporal and Motor Implementation Domains

(A–C) Conceptual schematic that parses motor skill learning into separate processes for timing and motor implementation. (A) Muscle activity patterns underlying a hypothetical six-element motor sequence. Each element is defined by its duration (timing) and the set of recruited muscles (motor implementation). Gray, muscle is "active." (B) Learning can be conceptualized as the process of changing timing (e.g., duration; top) and motor implementation (e.g., which muscles are active; bottom) of the individual motor elements. (C) Modified motor program resulting from changes to both aspects.

(D–F) Birdsong learning as an example of the process outlined in (A)–(C). (D) Spectrogram of a juvenile zebra finch song. (E) Learning modifies both temporal and spectral (i.e., motor implementation) aspects of song, as exemplified by changes to the duration and pitch of syllable "S_4." Pitch and duration estimates for 80 consecutive renditions of the syllable recorded at 60 (gray) and 115 (black) days posthatch (dph), respectively; pitch calculated from the harmonic stack part of the syllable. (F) Spectrogram of a song from the bird in (D) at 115 dph.

(G) Schematic of the song circuit underlying vocal learning and production. HVC and RA constitute the cortical part of the descending motor pathway. These motor regions are also indirectly connected through the anterior forebrain pathway (AFP), a basal ganglia (Area X)-thalamo (DLM)-cortical (LMAN) circuit.

(H) Presumed functional organization of the motor pathway in which HVC represents time (t) in the form of a synaptic chain network and RA neurons control specific muscles or muscle groups (m). Learning in the motor pathway is thought to be driven by plasticity in RA, which is facilitated and guided by input from the AFP (Doya and Sejnowski, 1995; Fiete et al., 2004, 2007; Troyer and Doupe, 2000). Adapted from Fee et al. (2004).

exploring how the basal ganglia (used as singular noun, as we refer to it as a functional entity) contributes to motor learning (Doupe et al., 2005; Fee and Goldberg, 2011). Recent models have the AFP implement aspects of a reinforcement learning process that shapes connectivity in motor cortex analog RA (Doya and Sejnowski, 1995; Fee and Goldberg, 2011; Fiete et al., 2007; Troyer and Doupe, 2000). Besides being the direct target of the AFP, the focus on RA as the nexus for song learning is also motivated by the finding that neurons in premotor nucleus HVC that project to RA encode time in the song (Hahnloser et al., 2002). This "clock code" in HVC has been hypothesized to provide a stable temporal input to RA during learning and production of song (Fee and Goldberg, 2011; Fee et al., 2004).

Given the functional organization of the song circuit (Figure 1H), learning can be understood as the process of establishing and refining connections between time-keeper neurons in HVC and muscle-related neurons in RA and further between RA collaterals (Sizemore and Perkel, 2011), such that the "right" muscles get activated at the appropriate times (Fee and Goldberg, 2011; Fee et al., 2004; Fiete et al., 2004, 2007). The AFP is thought to contribute to this process by inducing variability in RA neurons and thus song (Kao et al., 2005; Ölveczky et al., 2005, 2011) and by providing an instructive signal that biases the motor program toward improved performance (Andalman and Fee, 2009; Charlesworth et al., 2012; Fee and Goldberg, 2011; Warren et al., 2011).
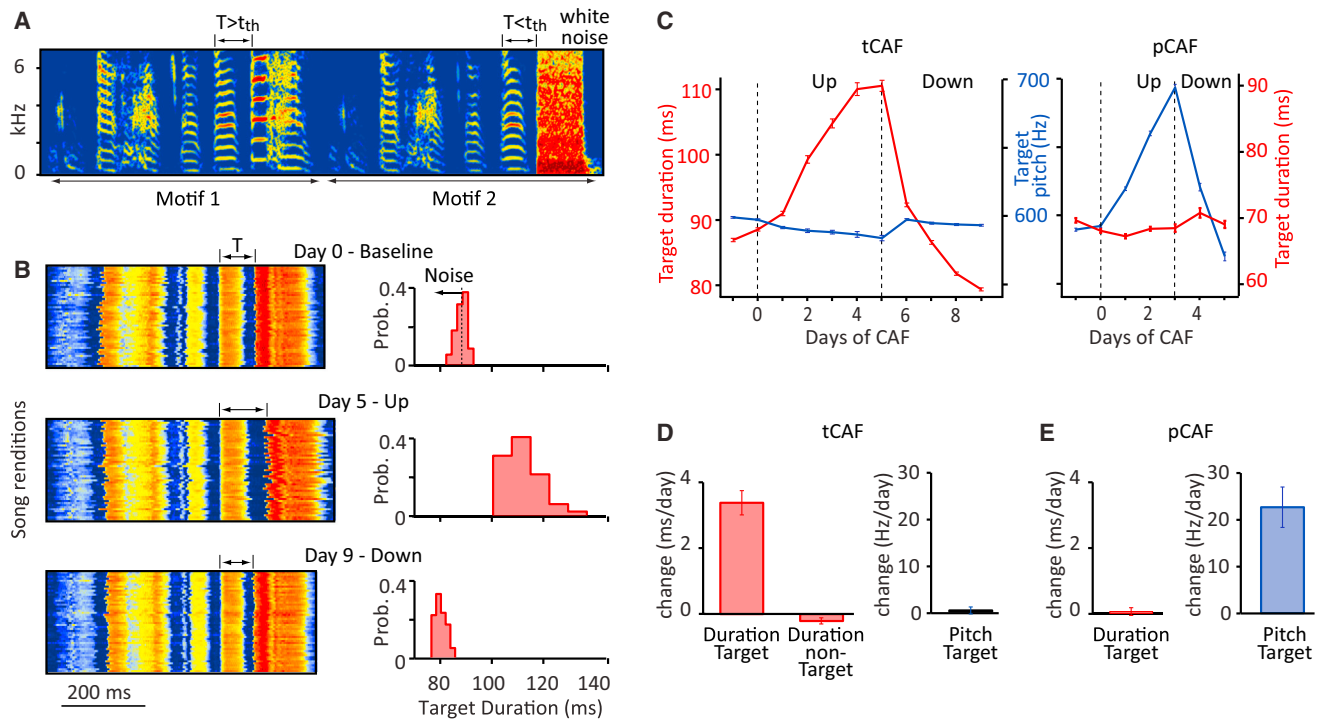
Figure 2. Independent Modification of Temporal and Spectral Song Features Using an Aversive Reinforcement Learning Paradigm
(A) Spectrogram of the song for the bird in (B) and (C). The duration of the target segment (T) is measured online and aversive white noise presented contingent on its duration being above (to shorten) or below (to lengthen) a threshold (t_{th}) (see Experimental Procedures).
(B) Song power (red, high; blue, low) for 50 consecutive motifs aligned on target onset at baseline (top) and after driving the target duration up for 5 days (middle) and down for 4 days (bottom); (right) associated target duration distributions.
(C) Learning trajectories (mean ± SEM) for duration (tCAF) and pitch (pCAF) for the target (see A). For pCAF, the target was the pitch of the syllable.
(D and E) Summary statistics (mean ± SEM) for changes in pitch and duration during tCAF (n = 24 birds) (D) and pCAF (n = 14 birds) (E). Duration values in (C), (D), and (E) refer to the targeted syllable for pCAF and the targeted segment (mostly "syllable + gap" as in C) for tCAF.

While this framework for song learning, i.e., plasticity in RA, can plausibly account for both temporal and spectral changes in song (Figure S1A available online), the extent to which other circuits are involved, and whether motor cortical and basal ganglia circuits distinguish learning in the temporal and spectral domains, has not been explored. To address this, we developed a reinforcement learning paradigm to independently modify both temporal and spectral features of zebra finch song. We perturb activity in different parts of the AFP, including its basal ganglia component Area X (Person et al., 2008) and cortical output (lateral magnocellular nucleus of the anterior nidopallium [LMAN]) and quantify how these circuit manipulations affect the capacity for learning temporal and spectral aspects of song. To probe whether the descending motor pathway encodes learned changes in the two domains differently, we record from neurons in HVC during modification to both temporal and spectral structure.

RESULTS

Independent Modification of Temporal and Spectral Song Structure
Testing whether the song system (Figures 1G and 1H) differentiates between learning in the temporal and spectral domains

requires experimentally modifying both aspects of song. A paradigm in which disruptive auditory feedback is delivered to the bird contingent on the pitch of one of its syllables has proven effective in adaptively altering spectral structure of song (pitch-conditional auditory feedback [pCAF]) (Tumer and Brainard, 2007). To probe whether temporal structure of adult zebra finch song is similarly plastic, we adapted this method to the temporal domain. This involved delivering aversive loud noise bursts every time the duration of a targeted song segment was below (to lengthen) or above (to shorten) a given threshold value (timing-conditional auditory feedback [tCAF], see Experimental Procedures and Figure 2A). To get precise and reliable online estimates of target duration, we targeted segments bounded by large and abrupt changes in sound amplitude, which in practice mostly meant intervals between ensuing syllable starts, i.e., "syllable + gap" segments (see Figure 2A and Experimental Procedures).

This paradigm induced rapid and predictable changes in the duration of targeted segments (Figures 2B–2D), demonstrating a remarkable capacity for changing the temporal structure of zebra finch song even well past song crystallization. Across the population of birds (n = 24), the duration of targeted segments changed by, on average, 3.4 ± 1.7 ms/day (mean ± SD) across 4–10 days of tCAF (Figure 2D; range: 0.9–6.4 ms/day,

p = 1.8 × 10^{-9}). Changes to temporal structure were specific to the targeted segments (Figure 2D), with minimal changes to the duration of nontargeted elements (−0.21 ± 0.43 ms/day). When targeting "syllable + gap" segments, both syllables and gaps changed in duration (syllables: 0.7 ± 0.6 ms/day, p = 4.6 × 10^{-5}; gaps: 2.8 ± 1.6, p = 7.7 × 10^{-8}; Figures S2 and S3C), though gaps changed significantly more than syllables (p = 1.3 × 10^{-5}). This difference was largely explained by the reinforcer being further removed in time from the syllables (by on average 47.2 ± 13.6 ms). When we experimentally delayed the noise burst by 50 ms relative to the end of the gap, the rate at which gaps changed decreased dramatically (79.7% ± 4.1%, n = 3 birds; Figures S2C and S2D). The effect was consistent with the difference in syllable and gap learning rate in our experiments being due to the differential delay in reinforcement (Figure S2E), though contribution from other factors cannot be discounted (Glaze and Troyer, 2012).

Learning was restricted not only in time but also to the feature being targeted. Changing the duration of a syllable did not alter its pitch (Figure 2D; pitch change during tCAF = 0.2 ± 2.6 Hz/day, p = 0.72). Similarly, modifying the pitch of a syllable using pCAF (Andalman and Fee, 2009; Warren et al., 2011) (Figure 2E; 22.6 ± 16.2 Hz/day; range: 7.3–62.8 Hz/day, n = 14 birds, p = 1.60 × 10^{-4}) did not affect its duration (Figures 2C and 2E; duration change during pCAF = 0.05 ± 0.43 ms/day, p = 0.65), suggesting that the two features, duration and pitch, may be independently learned and controlled (Figure S3). Having a method (CAF) for inducing rapid and reproducible changes to both spectral and temporal aspects of song allowed us to address the neural underpinnings of learning in the two domains and gauge the extent to which they are distinct.

Dissecting the Role of the AFP, a Basal Ganglia-Thalamo-Cortical Circuit

In our paradigm, adaptive changes to both pitch and duration rely on differential reinforcement of variable actions and as such are examples of reinforcement learning (Sutton and Barto, 1998). In the context of motor learning, this process requires two main ingredients: (1) motor variability producing exploratory actions and (2) a process converting information from this exploration into improved motor performance. LMAN, the output of the AFP, has been implicated in both aspects. Activity in this nucleus induces variability in vocal output (Kao et al., 2005; Ölveczky et al., 2005) and, in the spectral domain at least, drives an error-correcting premotor bias through its action on RA (Andalman and Fee, 2009; Charlesworth et al., 2012; Warren et al., 2011).

While LMAN has been a convenient proxy for understanding the role of the song-specialized basal ganglia-thalamo-cortical circuit (AFP), questions of how the basal ganglia itself (Area X) contributes to song learning (Kojima et al., 2013; Scharff and Nottebohm, 1991) and whether its role—and the role of LMAN—differs for learning in the temporal and spectral domains, have yet to be explored. To address this, we lesioned Area X and LMAN in separate experiments and compared variability and learning rates in the spectral and temporal domains before and after lesions.

## Area X Is Required for Learning in the Spectral, but Not Temporal, Domain

Bilateral lesions of Area X (Figure 3A, Tables S1 and S2, and Figure S5A) revealed a striking dissociation as to its role in learning. In the spectral domain (pCAF), learning was largely abolished following lesions (Figures 3B and 3E; pitch change 4.52 ± 4.05 Hz/day versus 32.42 ± 18.97 Hz/day before lesions, n = 6 birds; p = 2.03 × 10^{-5}). In fact, pCAF-induced changes to pitch after Area X lesions were not significantly different from normal baseline drift (Figure 3E; p = 0.48). In contrast, the capacity for modifying temporal structure remained unchanged. Average learning rates in tCAF experiments before and after lesions were similar with daily changes to target duration of 3.90 ± 2.03 ms before versus 3.30 ± 1.72 ms after lesion (Figures 3C and 3F; p = 0.63; n = 7 birds, 5 of which were also tested in pCAF).

Variability in both temporal and spectral features was unchanged from prelesion levels when measured 6 ± 2.5 days postlesion (range: 3–12 days; see also Figure S4 for acute but transient effects immediately following lesions), consistent with previous studies (Goldberg and Fee, 2011; Scharff and Nottebohm, 1991). The coefficient of variation (CV) in the duration of syllables and intersyllable gaps (Glaze and Troyer, 2013) was 2.9% ± 0.9% and 2.8% ± 0.6% before and after lesions, respectively (Figure 3D; n = 9 birds, p = 0.89), whereas the CV of pitch was 1.9% ± 1.3% and 1.9% ± 1.5% (Figure 3D; n = 9 birds, p = 0.79). This suggests that Area X is instrumental for learning spectral features not because it produces variability in this domain, but because it is required for generating the instructive signal expressed at the level of LMAN (Fee and Goldberg, 2011).

Reduction in LMAN Activity Reveals an Error-Correcting Motor Bias in the Spectral, but Not Temporal, Domain

In pCAF experiments, the learning-related instructive signal produced by the AFP manifests as an LMAN-dependent motor bias that shifts the pitch in the direction of learning (Andalman and Fee, 2009; Charlesworth et al., 2012; Warren et al., 2011). This bias can be estimated from the reversion in learned changes upon silencing of LMAN. If, however, learning temporal structure does not require the AFP, as our Area X lesion experiments suggest, then LMAN should also not contribute an error-correcting bias in this domain. To test this, we exposed our experimental subjects to female birds (see Experimental Procedures), a social manipulation known to dramatically reduce the variability and rate of LMAN firing (Kao et al., 2008) and thus decrease song variability in a way that mirrors the effect of pharmacological inactivations or lesions of LMAN (Kao et al., 2005; Ölveczky et al., 2005). Suppressing LMAN activity this way after 4–7 hr of pCAF exposure resulted in a 40.1% ± 20.3% mean reversion of that day's learned pitch changes (Figures 4A and 4B; n = 11 birds, 22 experiments, p = 6.5 × 10^{-5}), an effect very similar to what is seen after LMAN inactivations (Andalman and Fee, 2009; Warren et al., 2011). This reversion was seen both when the pitch was driven away from baseline (reversion toward baseline, 49.1% ± 41.3%) and toward it (reversion away from baseline, 35.2% ± 17.9%). After tCAF, however, there was no significant reversion in learned duration changes, consistent with LMAN not contributing an instructive bias in the temporal domain
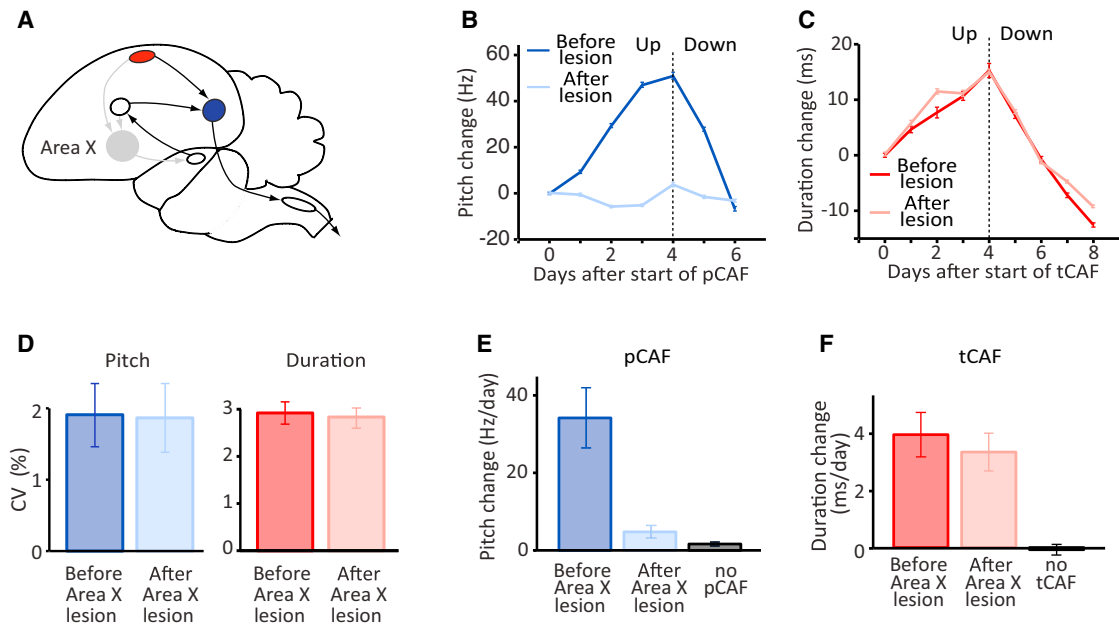
Figure 3. Area X Lesions Reveal a Dissociation in How This Basal Ganglia Structure Contributes to Adaptive Modification of Spectral and Temporal Song Features

(A) Schematic of the song circuit following Area X lesions. Gray denotes disrupted pathways/circuits.

(B) pCAF-induced changes to the pitch (mean ± SEM) of a targeted syllable in an example bird before (blue) and after (light blue) bilateral Area X lesions.

(C) tCAF-induced shifts in the duration (mean ± SEM) of the target interval before (red) and after (light red) Area X lesions in the same bird as (B).

(D) Variability in pitch and duration (mean ± SEM) before and after Area X lesions.

(E and F) Effects of Area X lesions on learning rates (mean ± SEM) in pCAF (n = 6 birds) (E) and tCAF (n = 7 birds) (F).

(Figures 4A and 4B; n = 5 birds, 12 experiments, 10.0% ± 11.2% reversion of the day's learned duration change, p = 0.12; see Experimental Procedures).

## LMAN Lesions Affect Variability and Learning in the Temporal Domain

If the AFP is not guiding adaptive changes to temporal structure, we reasoned that the capacity for learning in this domain should be robust to LMAN lesions. To test this, we ablated LMAN bilaterally in a separate group of birds (Figures 4C and S5B, Tables S1 and S2). A prior study, using pharmacological inactivation of LMAN in the context of pCAF (Charlesworth et al., 2012), had shown that LMAN is necessary for adaptively modifying pitch in pCAF. We confirmed this result in LMAN-lesioned birds, with learning rates in pCAF going from 13.3 ± 5.9 Hz/day before lesions to 0.7 ± 1.1 Hz/day after lesions (p = 6.7 × 10^{-4}; n = 4 birds, 3 of which also tested for tCAF; p = 0.11 when comparing LMAN-lesioned birds in pCAF to normal drift; Figure 4D). In the temporal domain, however, LMAN-lesioned birds retained the ability to learn, albeit at a reduced rate compared to prelesion (Figure 4E; prelesion: 2.8 ± 1.6 ms/day, postlesion: 0.9 ± 0.6 ms/day, p = 0.003 when comparing LMAN-lesioned birds in tCAF to normal drift). Mean reduction in the learning rate within a bird was 60.7% ± 29.4% (n = 8 birds, p = 6.3 × 10^{-4}).

Since LMAN is known to induce vocal exploration in both the temporal and spectral domains (Thompson et al., 2011), we wondered whether the decreased learning rates in tCAF following lesions could be explained by a reduction in temporal variability. Consistent with this, we found that variability in the duration of song elements (CV of syllable and intersyllable gaps [Glaze and Troyer, 2013], see Experimental Procedures) decreased within a bird by, on average, 38%, from 3.3% ± 1.2% to 2.1% ± 1.1% (Figure 4F; p = 4.8 × 10^{-4}). These results suggest that LMAN contributes to temporal learning by inducing variability in song timing. The process of converting information derived from this variability into improved motor timing, however, is probably implemented outside the AFP, as this process does not require an intact Area X or LMAN.

A Basal Ganglia Circuit Projecting to HVC Is Not Required for Learning Temporal Song Structure

Given the architecture of the song circuit, and the assumed role of the basal ganglia in reinforcement learning, an obvious candidate for driving temporal learning is the only other known song-related basal ganglia-thalamo-cortical circuit—a parallel circuit to the AFP that includes a basal ganglia-like structure medial to traditionally defined Area X (mArea X) (Kubikova et al., 2007), the thalamic nucleus DMP and the medial part of MAN (MMAN) (Figure 5A). Whereas the AFP projects directly to RA, which encodes spectral features (Sober et al., 2008), MMAN outputs directly to HVC and could, in analogy to its lateral counterpart (LMAN), provide the instructive signal for altering neural dynamics in HVC and thus temporal structure of song.
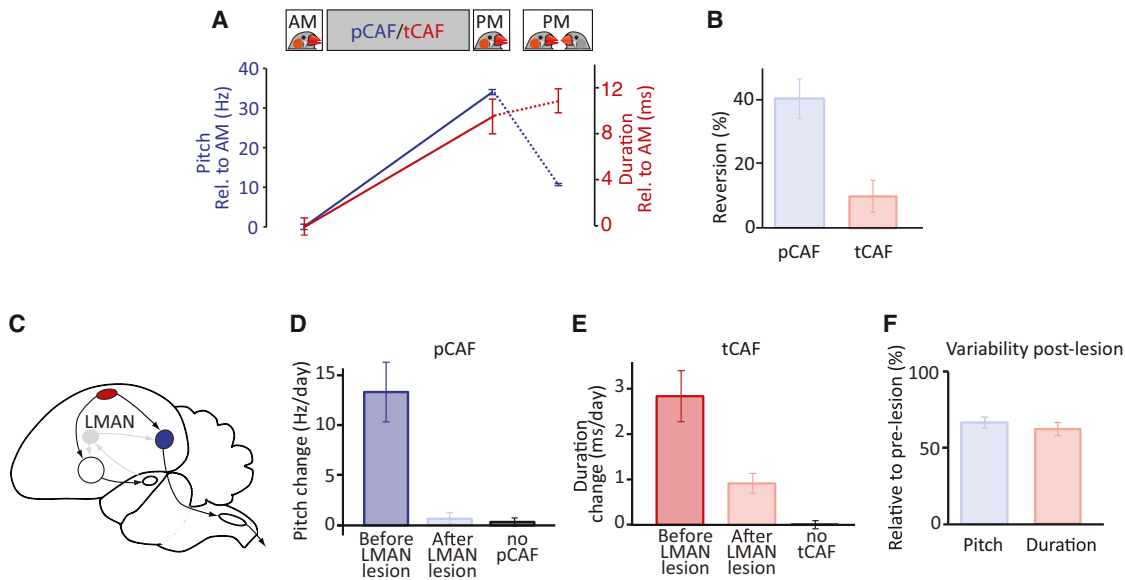
Figure 4. Distinct Roles for LMAN in Adaptive Modification of Temporal and Spectral Structure

(A) Female-directed singing, which reduces LMAN activity, caused a reversion in learned changes to pitch, but not interval duration. Example data (mean ± SEM) from the same bird shows the pitch and duration of the target before and after a day of pCAF (solid blue line) and tCAF (solid red line), respectively; last data point shows the corresponding values after presentation of a female (dashed lines).

(B) Reversion (mean ± SEM) of the day's learned changes upon presentation of a female bird (i.e., the directed singing-induced change in pitch or duration relative to the day's total change; n = 11 birds for pCAF, n = 5 birds for tCAF). Duration values corrected for global tempo changes observed during directed singing (Stepanek and Doupe, 2010) (see Experimental Procedures).

(C) Schematic of the song circuit following LMAN lesions. Gray denotes disrupted pathways/circuits.

(D and E) Learning rates (mean ± SEM) in pCAF (n = 4 birds) (D) and tCAF (n = 8 birds) (E) before and after LMAN lesions.

(F) Effects of LMAN lesions on variability (mean ± SEM) in pitch and interval duration.

To test this, we lesioned MMAN bilaterally (Tables S1 and S2 and Figure S5C), comparing learning rates in our tCAF paradigm before and after lesions. We saw no significant change in the capacity of birds to shift the duration of targeted song segments after MMAN lesions (Figure 5B; prelesion: 3.4 ± 1.8 ms/day, postlesion: 3.0 ± 1.3 ms/day, n = 3 birds, p = 0.34). Neither did MMAN lesions influence variability (CV) in temporal (prelesion: 2.6% ± 0.6%, postlesion: 2.6% ± 0.4%) or spectral (prelesion: 3.0% ± 0.9%, postlesion: 2.8% ± 0.5%) features of song (Figure 5C; p = 0.94 and 0.53, respectively), leaving its role in song learning, if any, to be elucidated.

Song Recovery in the Temporal Domain Does Not Require Area X

While the CAF paradigm allows us to address how the song system implements reinforcement learning in the spectral and temporal domains, the extent to which the same circuits and neural processes underlie "normal" song learning is unclear. Song learning is thought to be driven by an evaluation of the bird's vocalizations relative to an auditory template acquired from listening to a tutor early in life (Konishi, 2010). This auditory feedback-dependent learning process maintains stable adult song and restores it after experimental manipulations drive it away from the presumed template (Leonardo and Konishi, 1999; Sober and Brainard, 2009). Such a song recovery process would be expected to interact with CAF-based learning, working against it when the targeted feature (duration or pitch) is driven

away from baseline ("a" in Figure 6A) and in conjunction with it when driven toward it ("b" in Figure 6A). Consistent with this, learning rates in our CAF experiments were significantly higher when the targeted feature was driven toward baseline than when it was driven away (Figure 6B; pCAF: 42.0 ± 25.1 versus 26.0 ± 15.4 Hz/day, p = 0.03; Figure 6C; tCAF: 5.4 ± 2.2 versus 3.4 ± 1.9 ms/day, respectively, p = 0.04). To compare learning rates in the CAF paradigm to "normal" (i.e., CAF-free) song recovery, we drove the pitch or duration of targeted segments away from baseline by exposing birds to 3–5 days of CAF and then measured the rate at which the feature returned with and without CAF (Warren et al., 2011). Though both pitch and duration returned toward baseline, the rate of return was much lower without CAF (Figure 6B; 12.1 ± 12.4 Hz/day after cessation of pCAF; Figure 6C; 0.8 ± 0.3 ms/day after cessation of tCAF).

To test whether the dissociation in basal ganglia function uncovered with the CAF-paradigm (Figure 3) extends also to normal song learning, we lesioned Area X in a subset of birds and compared spontaneous (i.e., CAF-free) returns toward baseline before and after lesions. Because birds could not predictably alter the spectral structure of their vocal output (pCAF) after Area X lesions (Figures 3B and 3E), we drove targeted syllables away from their baseline pitch for 4 days (average drive away from baseline: 100.2 ± 76.0 Hz, n = 3 birds) before lesioning Area X bilaterally. Consistent with our CAF experiments, we saw no significant return to baseline even after 7 postlesion days of singing (Figure 6D; p = 0.17). The spontaneous change in pitch
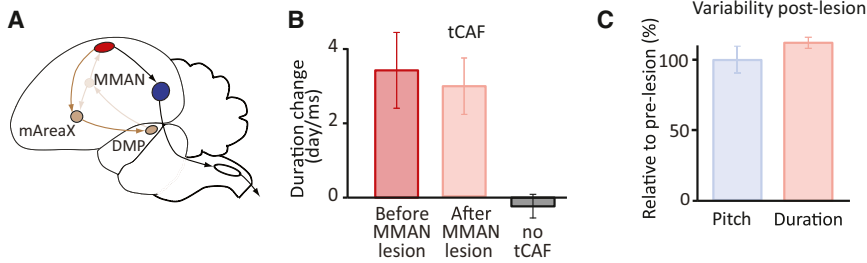
Figure 5. A Basal Ganglia-Thalamo-Cortical Circuit Parallel to the AFP with Projections to HVC Is Not Required for Learning Temporal Structure

(A) Schematic outline of the basal ganglia loop originating from and projecting back to HVC (brown). mArea X is a basal ganglia region medial to Area X; DMP is the dorsomedial nucleus of the posterior thalamus; MMAN is the medial magnocellular nucleus of the anterior nidopallium. Faint lines denote pathways/circuit disrupted by lesions to MMAN.

(B and C) Effect of MMAN lesions on learning rates (mean ± SEM) in tCAF (B) and spectral and temporal variability (mean ± SEM) of song (C) across three birds.

went from 15.3 ± 13.3 Hz/day before lesion to 1.6 ± 1.4 Hz after, suggesting that Area X is required for maintaining the spectral identity of song (Kojima et al., 2013). In the temporal domain, however, lesioning Area X did not affect the spontaneous recovery toward baseline (0.63 ± 0.13 ms/day before lesion versus 0.71 ± 0.71 ms/day after lesion, n = 4 birds, p = 0.76). Area X lesions also did not affect the difference in learning rates for tCAF drives toward and away from baseline ("b − a" in Figure 6D; p = 0.57), a difference we hypothesize being due, in part at least, to the template-based learning process working in opposite directions in the two cases. These results suggest that the dissociation in how the basal ganglia contributes to learning in the spectral and temporal domains extends to normal CAF-free song learning.

Premotor Cortical Region HVC Encodes Changes to Temporal, but Not Spectral, Structure

Given the difference in how the AFP contributes to learning in the temporal and spectral domains, we wondered whether learning-related changes in the motor pathway show a similar dissociation. While changes to both temporal and spectral structure can be understood within the existing framework for song learning (i.e., plasticity in RA), significant modifications to the duration of song segments, like those induced by our tCAF paradigm, would require an extensive reorganization of HVC-RA connectivity (Figure S1A). An alternative, which confers more flexibility on the learning process by capitalizing on the functional organization of the song control circuits (Figure 1H), would be for temporal changes to be encoded at the level of HVC (Figure S1B). Though white-noise feedback does not acutely affect song-related HVC activity (Kozhevnikov and Fee, 2007), we speculated that chronic exposure to the tCAF protocol could alter its dynamics to reflect adaptive changes to temporal structure. This would extend the current framework for song learning (Doya and Sejnowski, 1995; Fiete et al., 2004, 2007; Troyer and Doupe, 2000) to include changes in HVC activity, while also expanding the role of HVC beyond that of a generic "clock" (Fee et al., 2004; Fiete et al., 2004, 2007).

Describing the relationship between HVC dynamics and adaptive changes to temporal structure (Figure 2C) requires tracking the activity of HVC neurons over the course of learning. Given the difficulty in recording single units in HVC of freely behaving songbirds for extended periods (i.e., more than a few hours [Kozhev-

nikov and Fee, 2007; Sakata and Brainard, 2006; Yu and Margoliash, 1996]), we recorded multiunit activity (Crandall et al., 2007; Schmidt, 2003) while exposing birds to the CAF protocols (see Experimental Procedures). Song-aligned neural signals thus acquired were stable over many days (see Figures 7A and 7D for examples), allowing us to explore how HVC dynamics change with significant modifications to the song's temporal structure.

Relating HVC dynamics to vocal output requires taking into account the temporal lag between premotor activity in HVC and the sound produced. We estimated this lag by cross-correlating the HVC signal with sound amplitude and by computing the covariance in the temporal variability of the two signals (see Experimental Procedures). Both analyses showed HVC activity leading sound by, on average, 35 ms (Figure S6), consistent with the anticipatory premotor nature of HVC reported in previous studies (Fee et al., 2004; Schmidt, 2003; Vu et al., 1994).

In support of HVC encoding temporal changes, modifications to the duration of discrete song segments (mean shift per tCAF drive: 11.3 ± 4.2 ms; 3–5 days per drive; n = 13 tCAF drives in 6 birds) were associated with significant and target-specific changes in the underlying HVC signal (Figure 7A). Indeed, the correlation between the average song-aligned neural activity pattern before and after tCAF training was 0.50 ± 0.26 and 0.86 ± 0.18 for target and nontarget segments, respectively (Figure 7B, p = 0.002; see Experimental Procedures). Learning-related changes in HVC activity manifested predominantly as a temporal rescaling of the baseline signal, stretching or shrinking it in segments where the song had experienced lengthening or shortening, respectively. Accounting for the temporal changes in song by time warping the neural traces accordingly yielded a marked increase in the correlation between the neural signals before and after tCAF for the targeted segment (0.83 ± 0.09, see Experimental Procedures), making it not significantly different from the correlation values for time-warped nontargeted segments (0.88 ± 0.07, p = 0.24; Figure 7B). Time warping the average neural trace recorded at the end of a tCAF drive to best fit the pre-CAF recordings (see Experimental Procedures) yielded warping estimates that were very similar to those derived from warping the corresponding average song spectrograms to each other (R = 0.95 for targeted segments, n = 23 segments; Figure 7C), suggesting a strong mechanistic link between temporal restructuring of behavior and HVC dynamics.
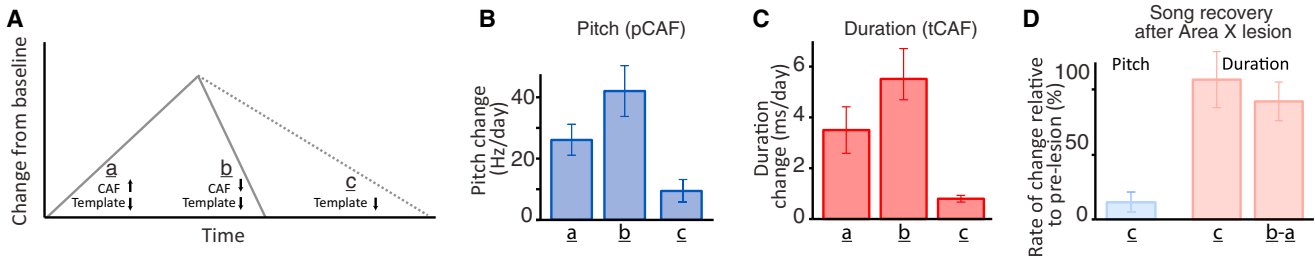
Figure 6. The Dissociation in How Area X Contributes to Learning in the Spectral and Temporal Domains Extends to "Normal" CAF-free Song Recovery

(A) Schematic of how CAF and the normal song recovery process ("template") may contribute and interact during various forms of learning: a, CAF away from baseline; b, CAF toward baseline; c, spontaneous return to baseline.

(B and C) Rate of change (mean ± SEM) in pitch (B) and duration (C) for the scenarios in (A).

(D) Effect of Area X lesions on the spontaneous return to baseline (mean ± SEM) for pitch and duration (n = 3 and 4 birds, respectively).

Inducing shifts in the pitch of targeted syllables (pCAF), on the other hand, yielded no target-specific change in HVC activity (Figure 7D; mean total shift per pCAF drive: 52.9 ± 31.3 Hz; 3–5 days per drive; n = 8 pCAF drives in 4 birds). Correlations in the neural traces before and after pCAF for target and nontarget segments were 0.89 ± 0.13 and 0.87 ± 0.13, respectively (Figure 7E; p = 0.76). These observations are consistent with the idea that changes to spectral structure are implemented downstream of HVC (Doya and Sejnowski, 1995; Fiete et al., 2007; Sober et al., 2008; Troyer and Doupe, 2000).

DISCUSSION

By making reinforcement contingent on variability in either temporal or spectral features of birdsong, we demonstrate the capacity of the nervous system to independently modify timing and motor implementation aspects of a motor skill (Figures 1 and 2). In dissecting the underlying circuits, we discovered a surprising dissociation in how learning is implemented in the two domains, with the basal ganglia essential for modifying spectral, but not temporal, features of song (Figure 3) and a premotor cortex analog area (HVC) encoding changes to temporal, but not spectral, features (Figure 7). The dissociation in how the different aspects of vocal output are learned extended to the normal song maintenance process (Figure 6), suggesting that "template-based" song learning (Konishi, 2010) may be an instantiation of reinforcement learning (Doya and Sejnowski, 1995; Fee and Goldberg, 2011). This also further validates the CAF paradigm (Tumer and Brainard, 2007) as a proxy for normal song learning, though the extent to which the two are similar need to be further explored.

Our results show that reinforcement learning in the spectral and temporal domains is implemented by distinct but partially overlapping circuits. Much of the exploratory variability in both aspects of vocal output is driven by the same thalamo-cortical circuit (DLM-LMAN [Goldberg and Fee, 2011]), which outputs directly to RA and indirectly to HVC (Hamaguchi and Mooney, 2012; Schmidt et al., 2004) (Figure 4). However, the circuits that convert the information gained from vocal exploration into a learning signal capable of driving changes in motor circuitry differ. For pitch, our results point to Area X as a key locus of rein-

forcement learning (Fee and Goldberg, 2011; Kojima et al., 2013). This basal ganglia homolog can affect the RA motor program by modulating activity in the downstream thalamo-cortical circuit to produce an error-correcting motor bias at the level of LMAN (Andalman and Fee, 2009; Warren et al., 2011; Charlesworth et al., 2012) (Figures 4A and 4B). For learning in the temporal domain, however, the circuits that translate the consequences of exploration into improved performance do not seem to involve the AFP or, more generally, the song-related basal ganglia circuits (Figures 3 and 5).

The anatomy of the song circuit together with our results showing learning-related changes in HVC activity points to this time-keeper circuit as a possible nexus for reinforcement learning of temporal features. This would require variability in motor timing to be expressed within HVC and for a performance-based evaluation signal to reach it—both plausible scenarios: LMAN, which drives much of the temporal variability underlying learning (Figure 4F), can influence HVC network dynamics through indirect connections (Hamaguchi and Mooney, 2012; Roberts et al., 2008; Schmidt et al., 2004), while midbrain dopaminergic projection neurons, a common source of reinforcement in vertebrate circuits (Fields et al., 2007), project directly to HVC (Appellants et al., 2000; Hamaguchi and Mooney, 2012) and, interestingly, also to Area X (Person et al., 2008). Thus, the same source of variability (LMAN) and reinforcement (midbrain dopamine neurons) could, in principle, underlie two distinct reinforcement learning processes. While follow-up studies are needed to conclusively establish where and how temporal learning happens within the song system, our result showing basal-ganglia-independent changes to HVC activity (Figure 7) makes this premotor nucleus a plausible candidate.

The basal ganglia is generally thought to be involved in the acquisition of learned motor behaviors (Doyon et al., 2009; Graybiel, 2005; Turner and Desmurget, 2010), yet the specifics of how it contributes to the learning process remain poorly understood. Our results, showing that the basal ganglia in songbirds is necessary for learning spectral, but not temporal, aspects of vocal output add important nuance to this question. Whether this reflects a general difference in how the basal ganglia contributes to motor skill learning remains to be explored, but our current study strongly suggests that the distinction between timing
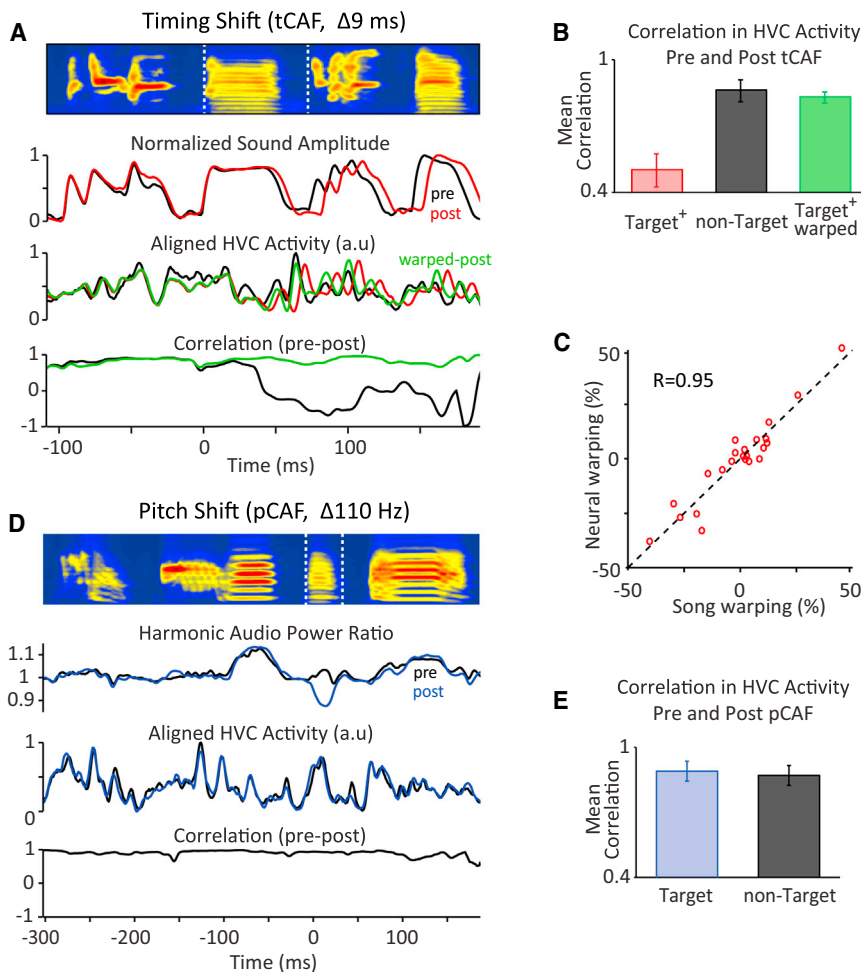
Figure 7. HVC Network Activity Reflects Learned Changes to Temporal, but Not Spectral, Structure

(A) Comparing HVC recordings before and after 3 days of tCAF, during which the duration of the target segment (bracketed by dashed white lines) increased by 9 ms. Top row: average song spectrogram prior to tCAF. Second row: mean audio power envelope ("Sound Amplitude") for the song motif before (black) and after (red) tCAF. Third row: mean neural power ("HVC Activity") before and after tCAF. Overlaid (green) is the HVC activity post-CAF time warped to account for temporal changes in the song (see Experimental Procedures). Bottom row: local Pearson's correlation (50 ms sliding window) between the song-aligned neural traces before and after tCAF, using original (black) and warped (green) post-tCAF traces.

(B) Summary statistics (mean ± SEM) for n = 13 tCAF experiments in 6 birds, showing correlations between HVC activity at the start and end of tCAF for conditions and song segments as indicated (target$^+$ = target + 100 ms).

(C) Time warping the average neural traces recorded in HVC after tCAF to those recorded before tCAF yielded estimates of temporal rescaling (percent stretching/shrinking) in the target intervals that were highly correlated with those derived from the respective sound recordings. Data points correspond to warping estimates for individual song elements (syllables and gaps) making up the target (n = 23 song elements from 13 tCAF drives in 6 birds).

(D) Comparing HVC recordings before and after a pCAF drive. Top, third, and bottom rows: same as in (A). Second row: ratio of power in the frequency bands corresponding to the first ten harmonics of the target syllable at baseline (pitch = 530 Hz, harmonic bandwidth = 10 Hz) to the total power in bands offset by half-pitch.

(E) Similar to (B) but for n = 8 pCAF drives in four birds.

and motor implementation (Figures 1A–1C) is a crucial one to make when considering basal ganglia function in the context of motor learning.

Control of motor timing in humans is thought to involve prefrontal regions (Halsband et al., 1993; Harrington and Haaland, 1999), yet little is known about how these circuits represent the temporal structure of motor output, and whether they are involved in learning. HVC, the equivalent structure in songbirds, has been studied in far greater detail. It is thought to control song timing in the form of a synaptically connected chain of neurons, where each node represents a specific time point in the song (Li and Greenside, 2006; Long et al., 2010) (Figure 1H). Our HVC recordings during temporal learning, however, show HVC to be more than an immutable time keeper. We observed activity patterns in this premotor nucleus stretch and shrink with the song (Figure 7), suggesting that temporal structure is modified by locally tuning the propagation speed within the network. Thus, rather than representing time, our result suggests that neurons in HVC encode specific parts of the song, e.g., the starts and ends of syllabic or subsyllabic elements, the relative timings of

which can be adjusted independently from other features of the song.

Modulating dynamics in HVC by means of temperature has previously been shown to uniformly alter song tempo without interfering with spectral content (Aronov and Fee, 2012; Long and Fee, 2008). Our results show that similar changes to HVC dynamics and song can be induced and consolidated through reinforcement learning. Moreover, we show that the temporal changes to song structure can be specific to certain parts of the song. The ability to shape the temporal structure of birdsong in such a specific manner is likely to be ethologically relevant: temporal features, such as syllable duration, distinguish song dialects (Wonke and Wallschläger, 2009) and can be shaped by exposure to different habitats (Kopuchian et al., 2004).

The ability to adaptively modify timing without interfering with other aspects of behavior may be critical to the acquisition and refinement of many motor skills also in humans (Gentner, 1987). Subtle changes to the temporal structure of syllables in human speech, for example, do not unduly change spectral aspects of vocal output (Cai et al., 2011). Furthermore, when a

targeted syllable segment is experimentally lengthened (Cai et al., 2011), subsequent speech patterns are similarly delayed to account for the increase in target duration, i.e., a phenomenology similar to what we see in songbirds (Figures 2B and 2D). Our results suggest a powerful and potentially very general solution for how this and other processes that alter temporal structure of learned motor output could be instantiated in neural circuitry (Figure S1B).

Having separate learning processes shape distinct aspects of a motor skill can have several advantages, chief among them the flexibility to modify them independently (Figures 1 and 2). The success of "slow practice," a method for training complex motor sequences championed by many music and dance teachers, is one of many examples attesting to this flexibility. Students are first taught proper motor implementation (i.e., which fingers/limbs to move in what sequence and to what extent) before refining the temporal structure of their performance. The underlying premise is that learning in the time domain does not interfere with other learned aspects of motor output. Our results show that this intuition is codified in the organization of the nervous system, which divides up the task of learning precise motor skills into functional modules for timing and motor implementation (Figure 1B), each with its distinct circuitry. This modularity may also be necessary to overcome the inherent limitations of reinforcement learning, basic implementations of which do not cope well with large task domains (Botvinick et al., 2009). Indeed, parsing up complex learning tasks into hierarchically connected, but largely independent, modules (Diuk et al., 2013) may have enabled increasingly complex behaviors to evolve by using (and reusing) the same rudimentary learning algorithms.

EXPERIMENTAL PROCEDURES

Adult male zebra finches (90+ days after hatch, n = 40) were obtained from the Harvard breeding facility and housed on a 13:11 hr light/dark cycle in individual sound-attenuating chambers with food and water provided ad libitum. The care and experimental manipulation of the animals were carried out in accordance with the guidelines of the National Institutes of Health and were reviewed and approved by the Harvard Institutional Animal Care and Use Committee.

Conditional Auditory Feedback Protocol
Custom software (LabVIEW) was used to implement the conditional auditory feedback (CAF) protocol used to manipulate pitch and duration of targeted song segments. The target was detected based on the correlation between the bird's song and a template spectrogram of the preceding 100–500 ms in the bird's song motif. Average detection rates as quantified by manually examining at least 80 songs both early and late in the CAF drive were generally high (>80%) and did not differ after any of the lesions (98% ± 3% prelesion versus 97% ± 4% postlesion).

Once a target was detected, its feature (pitch or duration) was computed. If it did not meet the escape threshold, white-noise feedback (lasting between 25–100 ms, but constant for a given bird) was played back through a loudspeaker with short latency (~1–3 ms). We calibrated the feedback volume to be marginally higher than the bird's loudest syllable, effectively setting it to ~80–95 dB (A-weighting), 10 cm away from the speaker. The threshold to escape white-noise feedback was dynamically updated based on the bird's performance over the last 200 renditions of the target. If the fraction of escapes exceeded 80%, the threshold was automatically adjusted to the bird's mean in those last 200 renditions, but the adjustment was only made in the direction of learning.

Target Estimates: pCAF
We chose target syllables with well-defined pitch (i.e., harmonic stacks) that were reliably (>80%) detected. Pitch was computed on a 5 ms sound segment of the target syllable using an algorithm fitting different sets of harmonics (see Supplemental Experimental Procedures). We computed pitch either at the very start of the syllable or 15–50 ms into it (varied between birds but constant within a bird).

Target Estimates: tCAF
Online estimates of targeted segment durations used threshold crossings of the smoothed (5 ms boxcar filter with 1 ms advancement) amplitude envelope. The threshold was set to ~2×–10× the background noise levels and kept constant throughout an experiment. Syllable onsets are associated with rapid increases in amplitude, which makes the estimates of their timing more robust to noise. Thus, we mostly targeted "syllable + gap" segments and estimated the target duration from the onset of the target syllable to the onset of the following syllable. However, in one bird, we made white noise conditional on the duration of a syllable, with the additional contingency that the subsequent gap duration did not change significantly. In four additional birds, we targeted intersyllable gaps (offset of last syllable to onset of next syllable). These five birds were pooled with the rest because they produced similar effects in response to experimental manipulations (e.g., lesions).

Experimental Design
The design for birds that underwent pCAF and tCAF both before and after lesions was as follows: one group did a continuous block of pCAF for at least 6 days, followed by at least a week of no CAF. This was followed by a continuous block of tCAF for at least 6 days. The birds then underwent surgery for lesions and were given at least 1 week to recover before repeating the pCAF and tCAF blocks in the same order. Another group of birds experienced the same protocol but with the order reversed (tCAF followed by pCAF). Because pCAF was impaired after Area X lesions, we wanted to rule out potential short-term effects of lesions on learning. We thus ran pCAF for two birds more than 4 weeks after lesion to confirm abolished learning. We typically exposed birds to CAF for the same number of days before and after lesion and targeted the same song segment. Some birds experienced either tCAF or pCAF only, in which cases we did at least one round of CAF (in both directions). See main text for details of sample sizes for the various experiments. In a subset of birds, we conducted spontaneous return-to-baseline experiments before and after Area X lesions (Figure 6). For tCAF experiments, we drove the targeted segment duration away from baseline for 3–5 days before removing white-noise feedback. The same protocol was repeated after Area X lesions (as birds can still shift duration). However, since birds cannot shift pitch after Area X lesions (Figure 3), for pCAF we drove targeted syllables away from their baseline pitch for 4 days and then turned CAF off, allowing birds to spontaneously recover to baseline for up to 7 days. The same birds were then driven up again for 4 days before lesioning Area X. The pitch was subsequently monitored for up to 7 days postlesion to assess any recovery to baseline.

Lesions
Birds were anesthetized under 1%–3% isoflurane in carbogen and placed in a stereotaxic apparatus. Targeted brain areas were lesioned by injecting 4% (w/v) of N-methyl-DL-aspartic acid (NMA; Sigma) at stereotactically defined locations (see Table S1). Lesions were confirmed histologically using cresyl violet staining. We identified Area X and LMAN based on regions of stronger staining and/or higher density of cells than surrounding areas and were additionally guided by anatomical landmarks (e.g., lamina palliosubpstatus lialis and lamina mesopallialis) (Karten et al., http://www.zebrafinch.org/neuroanatomy.html). MMAN was identified based on landmarks and presence of LMAN. Remaining Area X, LMAN, or MMAN volumes were quantified and compared to volumes from adult control birds (n = 4) with intact brains. Between 80%–100% of LMAN, 72%–98% of Area X, and 75%–100% of MMAN were lesioned (see Figure S5 and Table S2).

Directed Singing
To test LMAN-mediated premotor bias, we presented a female bird to the experimental subject after 4–7 hr of CAF. Each female was presented for

2–3 min, after which it was replaced with a different female. This sequence of single-female presentations continued for 15–30 min. All directed songs as well as catch trials just before presentation of females were uncontaminated by white noise (i.e., CAF was turned off).

Electrophysiology
Surgery
Birds (n = 7) were anesthetized with 1%–3% isoflurane in carbogen and placed in a stereotaxic apparatus. The location of Area X was estimated (see above) and confirmed by electrophysiological criteria (Kojima and Doupe, 2009). A bipolar electrode was acutely placed in Area X and used to identify the boundaries of HVC through antidromic stimulation. A custom recording array (four channels, ~250 μm spacing) of 100 kΩ tungsten or platinum electrodes (Microprobes) was implanted within the boundaries of HVC and a silver ground reference placed outside of HVC between the dura and the surface of the brain. Implanted components were secured to the skull with dental cement. All birds exhibited normal song output within 3 days of surgery; pre- and postsurgery song spectrograms were similar by visual inspection, suggesting minimal disruption of the targeted tissue. After completion of the experiment, the animals were sacrificed, their brains harvested, and the placement of recording and stimulating electrodes confirmed by histology.

Chronic Recordings in HVC
Sound and neural activity were recorded using a custom LabVIEW application. The raw neural signal was amplified (1,000×–10,000×) and band-pass filtered (1 Hz–15 kHz). Multiunit activity was recorded from up to four sites from each bird over 4–6 weeks. Because multiday stability of the recordings was crucial for our analysis, all subsequent analysis was done on data collected from the most stable recording site in each bird.

Data Analysis
Song Segmentation
All song and HVC recording analysis was performed offline using custom-written software (LabVIEW and MATLAB). Songs were sampled at 44.15 kHz and band-pass filtered (0.3–7 kHz). The dominant song motif for each bird was determined by visual inspection. Once a motif was chosen, it was identified in the sound recordings using a semiautomated routine, which included visual inspection of the segmented songs to verify that they indeed matched the chosen motif. These segmented motifs constituted the data for subsequent analysis.

Catch Trials
Song analysis was done on catch trials, i.e., songs recorded with the CAF protocol turned off, in the early morning (a.m. session) and evening (p.m. session). Approximately 100–200 songs/day were analyzed for each bird. Baseline data were analyzed for ~200 songs recorded 1–2 days before the start of CAF at comparable times to the CAF catch trials.

Pitch Estimates
Pitch estimates for the catch trials were calculated as described in Supplemental Experimental Procedures. Since pitch can be defined robustly only for harmonic stacks, we computed pitch variability for harmonic stack syllables in birds that had them. If a bird did not have any harmonic stack syllable, we analyzed pitch variability in a subsyllabic harmonic stack (see the latter half of syllable S_4 in Figure 1F for an example).

Interval Duration Estimates
Offline duration estimates from the catch trials were obtained by dynamically time warping (DTW) the songs to an average template (Glaze and Troyer, 2006). We implemented our DTW algorithm on spectrograms, using the L^2-norm of the difference in the log-transformed spectrogram at each time point as the local distance metric. Slopes of the warping paths were constrained to be between 0.5 and 2. Template start and end points were not constrained to align to the start and end points in the rendition. For details on how interval durations were estimated using DTW, see Supplemental Experimental Procedures.

Temporal Variability
Temporal variability in interval (i.e., syllable and gap) durations was estimated as described previously (Glaze and Troyer, 2012). Briefly, rendition-to-rendition variability of interval durations in the song was parsed into local, global, and jitter components by factor analysis. Local variability refers to independent variations in interval lengths, global variability captures correlated variability across intervals (due to e.g., temperature [Aronov and Fee, 2012; Long and Fee, 2008] or circadian [Glaze and Troyer, 2006] effects), and jitter is the variance in determining an interval's boundary. Given that we are inducing temporal shifts in the duration of individual segments, we report on the "independent" variability component (Glaze and Troyer, 2013) but note that the other components showed similar trends after Area X and LMAN lesions. Coefficient of variation was calculated for each interval and averaged over the intervals in the bird's song.

Pre- and Postlesion Variability Comparison
We compared both temporal and pitch variability before and after lesion. For prelesion, we used songs produced in the mornings up to 2 days preceding the surgery, grouped into a single catch trial block to increase our sample size. For postlesion, we analyzed morning songs for up to 2 days, at different times after surgery to parse acute (1–3 days postlesion) and persistent (3+ days postlesion) effects (Figures 3D and S4).

Learning Rates
For tCAF and pCAF, respectively, we computed learning rates as the difference in the average pitch or duration in a.m. catch trials on the first and last day of CAF, divided by the number of intervening days. We did the same for p.m. catch trials and the overall learning rate was then averaged across a.m. and p.m. catch trials for the whole drive up and down (sign inverted) for each bird to obtain a more robust estimate of the learning. For a small number of birds that did not sing during either a.m. or p.m. catch trial blocks, we computed learning rate from the remaining block only (e.g., a.m. only). Comparing the same time periods in the day allowed us to rule out circadian effects.

Directed Singing
Estimates of pitch and duration were computed as described above. In addition, we corrected duration estimates for global tempo changes during directed singing (Stepanek and Doupe, 2010), estimated as the average change in the duration of nontarget intervals during directed songs compared to undirected songs immediately before presentation of females. Reversion was calculated as the difference between the pitch or duration estimate just prior to presentation of the female (undirected p.m., see Figure 4A) and during directed singing p.m. and normalized to the total change in pitch or duration during the 4–7 hr of CAF.

Song and Neural Alignment within a Block
Songs during catch trial blocks were segmented and a song template created as described in Supplemental Experimental Procedures. Starts and ends of intervals (syllables and gaps) were extracted for each rendition and linearly warped to the template. The warping path was time shifted by 35 ms to account for the lag between HVC and sound output (Figure S6) and then applied to the band-pass filtered HVC voltage trace (0.3–6 kHz, zero-phase, 2-pole Butterworth). The squared voltage was averaged across all renditions in the block and smoothed with a 5 ms boxcar window to generate the mean neural power trace. Spectrograms warped to the common template were similarly averaged to generate a mean spectrogram for the block. The average warping paths across the renditions were then applied to the mean spectrogram and neural trace to remove any template specific effects.

Song and Neural Alignment across Blocks
The mean neural traces and spectrograms were calculated as described above for the start and end of a CAF drive. To account for CAF-induced changes in temporal song structure, the post-CAF spectrogram was warped to the baseline spectrogram, using the same DTW warping routine as described above. Warping estimates for each interval were calculated as the ratio of post-CAF to pre-CAF interval duration. The warping paths thus derived were applied to the average post-CAF neural trace, yielding the green traces in Figure 7A. The same DTW routine was also applied to the neural traces to compare the warping in the underlying neural signal to warping in the song (Figure 7C). To make the warping estimates for the neural data more reliable, we flagged salient points in the neural trace (i.e., well-defined peaks and troughs) and calculated the time shifts in these points over the course of the CAF drive. Since these points did not always line up with the interval boundaries in the song, we took the weighted average of the time shifts in the points within 10 ms of the interval boundary, each point being weighted inversely to its distance from the boundary. The estimate for the neural warping in a given

interval was then derived from the difference in the estimated time shifts corresponding to the start and end points of the interval.

Correlations in Neural Power

To quantify the degree and temporal specificity of the changes in neural power induced by CAF, we calculated running Pearson's correlations (50 ms boxcar window, 1 ms advance) between the neural power in baseline and post-CAF conditions. For each analyzed CAF drive, we compared the mean correlation of nontargeted song intervals (motif onset to 50–100 ms prior to CAF target) with those in the targeted interval (pCAF) or targeted interval plus 100 ms (tCAF).

Statistical Testing

All statistics presented in the main text refer to mean ± SD, while error bars in the figures all represent SEM. All statistical tests assessing significance across manipulations in the same birds were done using paired-samples t tests or one-sample t tests against mean zero unless otherwise noted.

REFERENCES

Andalman, A.S., and Fee, M.S. (2009). A basal ganglia-forebrain circuit in the songbird biases motor output to avoid vocal errors. Proc. Natl. Acad. Sci. USA *106*, 12518–12523.

Appeltants, D., Absil, P., Balthazart, J., and Ball, G.F. (2000). Identification of the origin of catecholaminergic inputs to HVc in canaries by retrograde tract tracing combined with tyrosine hydroxylase immunocytochemistry. J. Chem. Neuroanat. *18*, 117–133.

Aronov, D., and Fee, M.S. (2012). Natural changes in brain temperature underlie variations in song tempo during a mating behavior. PLoS ONE *7*, e47856.

Bottjer, S.W., Miesner, E.A., and Arnold, A.P. (1984). Forebrain lesions disrupt development but not maintenance of song in passerine birds. Science *224*, 901–903.

Botvinick, M.M., Niv, Y., and Barto, A.C. (2009). Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective. Cognition *113*, 262–280.

Cai, S., Ghosh, S.S., Guenther, F.H., and Perkell, J.S. (2011). Focal manipulations of formant trajectories reveal a role of auditory feedback in the online control of both within-syllable and between-syllable speech timing. J. Neurosci. *31*, 16483–16490.

Charlesworth, J.D., Warren, T.L., and Brainard, M.S. (2012). Covert skill learning in a cortical-basal ganglia circuit. Nature *486*, 251–255.

Crandall, S.R., Aoki, N., and Nick, T.A. (2007). Developmental modulation of the temporal relationship between brain and behavior. J. Neurophysiol. *97*, 806–816.

Diuk, C., Tsai, K., Wallis, J., Botvinick, M., and Niv, Y. (2013). Hierarchical learning induces two simultaneous, but separable, prediction errors in human basal ganglia. J. Neurosci. *33*, 5797–5805.

Doupe, A.J., and Kuhl, P.K. (1999). Birdsong and human speech: common themes and mechanisms. Annu. Rev. Neurosci. *22*, 567–631.

Doupe, A.J., Perkel, D.J., Reiner, A., and Stern, E.A. (2005). Birdbrains could teach basal ganglia research a new song. Trends Neurosci. *28*, 353–363.

Doya, K., and Sejnowski, T. (1995). A novel reinforcement model of birdsong vocalization learning. Adv. Neural Inf. Process. Syst. *7*, 101–108.

Doyon, J., Bellec, P., Amsel, R., Penhune, V., Monchi, O., Carrier, J., Lehéricy, S., and Benali, H. (2009). Contributions of the basal ganglia and functionally related brain structures to motor learning. Behav. Brain Res. *199*, 61–75.

Farries, M.A., and Perkel, D.J. (2002). A telencephalic nucleus essential for song learning contains neurons with physiological characteristics of both striatum and globus pallidus. J. Neurosci. *22*, 3776–3787.

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

Fee, M.S., Kozhevnikov, A.A., and Hahnloser, R.H.R. (2004). Neural mechanisms of vocal sequence generation in the songbird. Ann. N Y Acad. Sci. *1016*, 153–170.

Fields, H.L., Hjelmstad, G.O., Margolis, E.B., and Nicola, S.M. (2007). Ventral tegmental area neurons in. learned appetitive behavior and positive reinforcement. Annu. Rev. Neurosci. *30*, 289–316.

Fiete, I.R., Hahnloser, R.H., Fee, M.S., and Seung, H.S. (2004). Temporal sparseness of the premotor drive is important for rapid learning in a neural network model of birdsong. J. Neurophysiol. *92*, 2274–2282.

Fiete, I.R., Fee, M.S., and Seung, H.S. (2007). Model of birdsong learning based on gradient estimation by dynamic perturbation of neural conductances. J. Neurophysiol. *98*, 2038–2057.

Foster, E.F., and Bottjer, S.W. (2001). Lesions of a telencephalic nucleus in male zebra finches: Influences on vocal behavior in juveniles and adults. J. Neurobiol. *46*, 142–165.

Foster, E.F., Mehta, R.P., and Bottjer, S.W. (1997). Axonal connections of the medial magnocellular nucleus of the anterior neostriatum in zebra finches. J. Comp. Neurol. *382*, 364–381.

Gentner, D.R. (1987). Timing of skilled motor performance: Tests of the proportional duration model. Psychol. Rev. *94*, 255–276.

Glaze, C.M., and Troyer, T.W. (2006). Temporal structure in zebra finch song: implications for motor coding. J. Neurosci. *26*, 991–1005.

Glaze, C.M., and Troyer, T.W. (2012). A generative model for measuring latent timing structure in motor sequences. PLoS ONE *7*, e37616.

Glaze, C.M., and Troyer, T.W. (2013). The Development of Temporal Structure in Zebra Finch Song. J. Neurophysiol. *109*, 1025–1035.

Goldberg, J.H., and Fee, M.S. (2011). Vocal babbling in songbirds requires the basal ganglia-recipient motor thalamus but not the basal ganglia. J. Neurophysiol. *105*, 2729–2739.

Goller, F., and Suthers, R.A. (1996). Role of syringeal muscles in controlling the phonology of bird song. J. Neurophysiol. *76*, 287–300.

Graybiel, A.M. (2005). The basal ganglia: learning new tricks and loving it. Curr. Opin. Neurobiol. *15*, 638–644.

Hahnloser, R.H., Kozhevnikov, A.A., and Fee, M.S. (2002). An ultra-sparse code underlies the generation of neural sequences in a songbird. Nature *419*, 65–70.

Halsband, U., Ito, N., Tanji, J., and Freund, H.J. (1993). The role of premotor cortex and the supplementary motor area in the temporal control of movement in man. Brain *116*, 243–266.

Hamaguchi, K., and Mooney, R. (2012). Recurrent interactions between the input and output of a songbird cortico-basal ganglia pathway are implicated in vocal sequence variability. J. Neurosci. *32*, 11671–11687.

Harrington, D.L., and Haaland, K.Y. (1999). Neural underpinnings of temporal processing: a review of focal lesion, pharmacological, and functional imaging research. Rev. Neurosci. *10*, 91–116.

Immelmann, K. (1969). Song development in the zebra finch and other estrildid finches. In Bird Vocalizations, R.A. Hinde, ed. (Cambridge: Cambridge University Press), pp. 61–74.

Kao, M.H., Doupe, A.J., and Brainard, M.S. (2005). Contributions of an avian basal ganglia-forebrain circuit to real-time modulation of song. Nature *433*, 638–643.

Kao, M.H., Wright, B.D., and Doupe, A.J. (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.

Kojima, S., and Doupe, A.J. (2009). Activity propagation in an avian basal ganglia-thalamocortical circuit essential for vocal learning. J. Neurosci. *29*, 4782–4793.

Kojima, S., Kao, M.H., and Doupe, A.J. (2013). Task-related "cortical" bursting depends critically on basal ganglia input and is linked to vocal plasticity. Proc. Natl. Acad. Sci. USA *110*, 4756–4761.

Konishi, M. (2010). From central pattern generator to sensory template in the evolution of birdsong. Brain Lang. *115*, 18–20.

Kopuchian, C., Alejandro Lijtmaer, D., Luis Tubaro, P., and Handford, P. (2004). Temporal stability and change in a microgeographical pattern of song variation in the rufous-collared sparrow. Anim. Behav. *68*, 551–559.

Kozhevnikov, A.A., and Fee, M.S. (2007). Singing-related activity of identified HVC neurons in the zebra finch. J. Neurophysiol. *97*, 4271–4283.

Kubikova, L., Turner, E.A., and Jarvis, E.D. (2007). The pallial basal ganglia pathway modulates the behaviorally driven gene expression of the motor pathway. Eur. J. Neurosci. *25*, 2145–2160.

Leonardo, A., and Konishi, M. (1999). Decrystallization of adult birdsong by perturbation of auditory feedback. Nature *399*, 466–470.

Li, M., and Greenside, H. (2006). Stable propagation of a burst through a one-dimensional homogeneous excitatory chain model of songbird nucleus HVC. Phys. Rev. E Stat. Nonlin. Soft Matter Phys. *74*, 011918.

Lipkind, D., and Tchernichovski, O. (2011). Quantification of developmental birdsong learning from the subsyllabic scale to cultural evolution. Proc. Natl. Acad. Sci. USA *108*(*Suppl 3*), 15572–15579.

Long, M.A., and Fee, M.S. (2008). Using temperature to analyse temporal dynamics in the songbird motor pathway. Nature *456*, 189–194.

Long, M.A., Jin, D.Z., and Fee, M.S. (2010). Support for a synaptic chain model of neuronal sequence generation. Nature *468*, 394–399.

Nottebohm, F., Kelley, D.B., and Paton, J.A. (1982). Connections of vocal control nuclei in the canary telencephalon. J. Comp. Neurol. *207*, 344–357.

Ölveczky, B.P., Andalman, A.S., and Fee, M.S. (2005). Vocal experimentation in the juvenile songbird requires a basal ganglia circuit. PLoS Biol. *3*, e153.

Ölveczky, B.P., Otchy, T.M., Goldberg, J.H., Aronov, D., and Fee, M.S. (2011). Changes in the neural control of a complex motor sequence during learning. J. Neurophysiol. *106*, 386–397.

Person, A.L., Gale, S.D., Farries, M.A., and Perkel, D.J. (2008). Organization of the songbird basal ganglia, including area X. J. Comp. Neurol. *508*, 840–866.

Reiner, A., Perkel, D.J., Bruce, L.L., Butler, A.B., Csillag, A., Kuenzel, W., Medina, L., Paxinos, G., Shimizu, T., Striedter, G., et al.; Avian Brain Nomenclature Forum. (2004). Revised nomenclature for avian telencephalon and some related brainstem nuclei. J. Comp. Neurol. *473*, 377–414.

Roberts, T.F., Klein, M.E., Kubke, M.F., Wild, J.M., and Mooney, R. (2008). Telencephalic neurons monosynaptically link brainstem and forebrain premotor networks necessary for song. J. Neurosci. *28*, 3479–3489.

Sakata, J.T., and Brainard, M.S. (2006). Real-time contributions of auditory feedback to avian vocal motor control. J. Neurosci. *26*, 9619–9628.

Scharff, C., and Nottebohm, F. (1991). A comparative study of the behavioral deficits following lesions of various parts of the zebra finch song system: implications for vocal learning. J. Neurosci. *11*, 2896–2913.

Schmidt, M.F. (2003). Pattern of interhemispheric synchronization in HVc during singing correlates with key transitions in the song pattern. J. Neurophysiol. *90*, 3931–3949.

Schmidt, M.F., Ashmore, R.C., and Vu, E.T. (2004). Bilateral control and interhemispheric coordination in the avian song motor system. Ann. N Y Acad. Sci. *1016*, 171–186.

Sizemore, M., and Perkel, D.J. (2011). Premotor synaptic plasticity limited to the critical period for song learning. Proc. Natl. Acad. Sci. USA *108*, 17492–17497.

Sober, S.J., and Brainard, M.S. (2009). Adult birdsong is actively maintained by error correction. Nat. Neurosci. *12*, 927–931.

Sober, S.J., Wohlgemuth, M.J., and Brainard, M.S. (2008). Central contributions to acoustic variation in birdsong. J. Neurosci. *28*, 10370–10379.

Stepanek, L., and Doupe, A.J. (2010). Activity in a cortical-basal ganglia circuit for song is required for social context-dependent vocal variability. J. Neurophysiol. *104*, 2474–2486.

Sutton, R.S., and Barto, A.G. (1998). Reinforcement Learning: An Introduction. (Cambridge: MIT Press).

Tchernichovski, O., Mitra, P.P., Lints, T., and Nottebohm, F. (2001). Dynamics of the vocal imitation process: how a zebra finch learns its song. Science *291*, 2564–2569.

Thompson, J.A., Basista, M.J., Wu, W., Bertram, R., and Johnson, F. (2011). Dual pre-motor contribution to songbird syllable variation. J. Neurosci. *31*, 322–330.

Troyer, T.W., and Doupe, A.J. (2000). An associational model of birdsong sensorimotor learning I. Efference copy and the learning of song syllables. J. Neurophysiol. *84*, 1204–1223.

Tumer, E.C., and Brainard, M.S. (2007). Performance variability enables adaptive plasticity of 'crystallized' adult birdsong. Nature *450*, 1240–1244.

Turner, R.S., and Desmurget, M. (2010). Basal ganglia contributions to motor control: a vigorous tutor. Curr. Opin. Neurobiol. *20*, 704–716.

Vu, E.T., Mazurek, M.E., and Kuo, Y.C. (1994). Identification of a forebrain motor programming network for the learned song of zebra finches. J. Neurosci. *14*, 6924–6934.

Warren, T.L., Tumer, E.C., Charlesworth, J.D., and Brainard, M.S. (2011). Mechanisms and time course of vocal learning and consolidation in the adult songbird. J. Neurophysiol. *106*, 1806–1821.

Wild, J.M. (1993). Descending projections of the songbird nucleus robustus archistriatalis. J. Comp. Neurol. *338*, 225–241.

Williams, S.M., Nast, A., and Coleman, M.J. (2012). Characterization of synaptically connected nuclei in a potential sensorimotor feedback pathway in the zebra finch song system. PLoS ONE *7*, e32178.

Wonke, G., and Wallschläger, D. (2009). Song dialects in the yellowhammer Emberiza citrinella: bioacoustic variation between and within dialects. J. Ornithol. *150*, 117–126.

Yu, A.C., and Margoliash, D. (1996). Temporal hierarchical control of singing in birds. Science *273*, 1871–1875.