

Evidence for a causal inverse model in an avian cortico-basal ganglia circuit

Nicolas Giret^{a,b}, Joergen Kornfeld^a, Surya Ganguli^c, and Richard H. R. Hahnloser^{a,b,1}

^aInstitute of Neuroinformatics and ^bNeuroscience Center Zurich, University of Zurich and ETH Zurich, 8057 Zurich, Switzerland; and ^cDepartment of Applied Physics, Stanford University, Stanford, CA 94305-5447

Edited by Thomas C. Südhof, Stanford University School of Medicine, Stanford, CA, and approved March 7, 2014 (received for review September 11, 2013)

Learning by imitation is fundamental to both communication and social behavior and requires the conversion of complex, nonlinear sensory codes for perception into similarly complex motor codes for generating action. To understand the neural substrates underlying this conversion, we study sensorimotor transformations in songbird cortical output neurons of a basal-ganglia pathway involved in song learning. Despite the complexity of sensory and motor codes, we find a simple, temporally specific, causal correspondence between them. Sensory neural responses to song playback mirror motor-related activity recorded during singing, with a temporal offset of roughly 40 ms, in agreement with short feedback loop delays estimated using electrical and auditory stimulation. Such matching of mirroring offsets and loop delays is consistent with a recent Hebbian theory of motor learning and suggests that cortico-basal ganglia pathways could support motor control via causal inverse models that can invert the rich correspondence between motor exploration and sensory feedback.

lateral magnocellular nucleus of the anterior nidopallium | Hebbian learning | mirror neuron

The brain has evolved diverse strategies for combining sensory and motor signals, a prerequisite for many complex behaviors including hunting, communication, and observational learning of motor skills. For example, to accurately sense the external world while simultaneously moving within it, the nervous system must be able to detect changes in its sensory inputs that are not a predictable consequence of self-motion. Indeed, many sensory neurons respond with high sensitivity to unpredictable stimuli during motor behavior despite self-caused sensory feedback (1–5). Such remarkable sensitivity can be achieved by circuit mechanisms that counteract sensory feedback associated with self-generated motor output (6). Such mechanisms are also known as corollary discharges (7) or forward models of the motor system (8), which are synaptic mappings from motor neurons onto sensory neurons that can either predict or suppress sensory feedback from self-generated motor output.

In contrast to our understanding of how the brain cancels predictable, motor-induced sensory feedback, much less is known about neural mechanisms for computing motor codes that produce desired sensory targets. However, the ability to learn to produce desired behaviors by observing others is considered to be a key advantage of sociality and a main driver for the evolution of culture (9, 10). Learning by imitation occurs spontaneously when humans learn to speak, parrots imitate surrounding sounds (11), or songbirds learn to imitate a tutor's song (12). However, insights into the neural implementation of sensory-guided motor learning remain sparse, largely because we lack empirical information about the principles underlying the flow of neural activity through synaptic mappings from sensory to motor areas. Such mappings are known as inverse models and flow in the opposite direction of forward models that direct motor activity to sensory areas.

To learn about principles of sensorimotor integration in the zebra finch, a vocal learner, we focus our attention on cortical premotor nuclei necessary for song production and song learning. The premotor area HVC is involved in generating the stereotyped song motifs of adult birds (13), whereas the lateral magnocellular

nucleus of the anterior nidopallium (LMAN) forms the output of a basal-ganglia pathway involved in generating subtle song variability (14–17). HVC neurons produce highly stereotyped firing patterns during singing (18, 19), whereas LMAN neurons produce highly variable patterns (15, 20, 21).

To gain insights into LMAN's role in motor control, we consider recent theoretical work that establishes a conceptual link between inverse models and vocal-auditory mirror neurons (22–24). We consider three (nonexhaustive) possibilities about the flow of sensory information into motor areas. First, auditory afferents with some feature sensitivity could map onto motor neurons involved in generating those same features (Fig. 1A). This mapping forms a causal inverse, in which a sensory target input generates a motor activity pattern required to cause, or generate that same sensory target. Second, auditory afferents with some feature sensitivity could map onto motor neurons that typically fire after the ones involved in generating those features (Fig. 1B). This mapping forms a predictive inverse, in which a sensory input at some point in a stereotyped acoustic sequence elicits a predictive motor activity pattern required to generate the next acoustic signal in that sequence. Third, the auditory-to-motor connections could be randomly wired (Fig. 1C), in which case there would be no regularity in the relationship between sensory and motor responses.

Rather than directly characterize the auditory-to-motor mapping (a daunting task), we probe this mapping indirectly by studying the neural responses it causes in experiments in which we compare auditory responses elicited by playback of the bird's own song (BOS) to motor responses recorded during production of these songs. Indeed, the three possibilities in Fig. 1 make specific, testable predictions at the level of single neurons. Consider for example a motor neuron downstream of a causal inverse model (i.e., neuron 2 in Fig. 1A). This motor neuron generates song feature B after a motor latency (τ_m). Now, when the bird is not singing, this motor neuron also has a sensory selectivity for song feature B (because it is downstream of a causal inverse model). Therefore, during song playback, this same motor neuron will fire after playback of song feature B with an auditory latency τ_a . Thus, if one temporally aligns both the playback-evoked spike train and

Significance

Auditory neural responses mirror motor activity in a songbird cortical area. The average temporal offset of mirrored responses is roughly equal to short sensorimotor loop delays. This correspondence between mirroring offsets and loop delays constitutes evidence for a causal inverse model. Causal inverse models can map a desired sensation into the required action.

Author contributions: N.G., J.K., S.G., and R.H.R.H. designed research; N.G. and J.K. performed research; N.G., J.K., and R.H.R.H. analyzed data; and N.G., S.G., and R.H.R.H. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹To whom correspondence should be addressed. E-mail: rich@ini.phys.ethz.ch.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1317087111/-DCSupplemental.

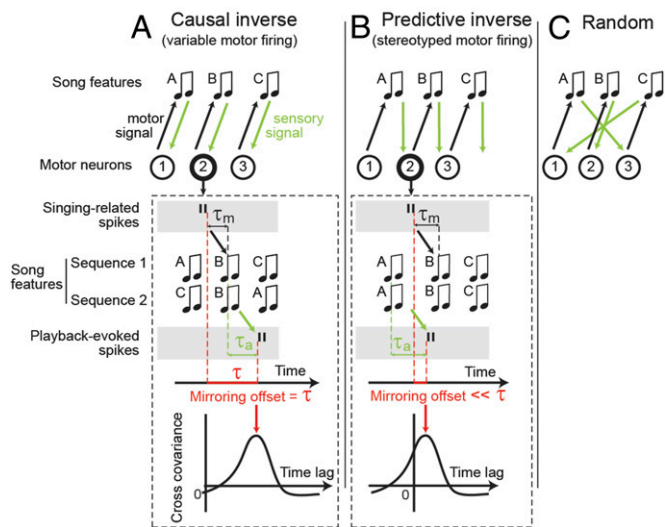


Fig. 1. Three hypothetical sensorimotor mappings and associated mirroring offsets. Sensory-to-motor mappings could implement a causal inverse of the motor plan (A), a predictive inverse (B), or be random (C). Under a causal inverse, generated by variable sequences of song features (ABC-CBA), a spike burst in a motor neuron (neuron 2) triggers the production (black arrow) of a song feature (feature B) after latency τ_m , and the neuron receives sensory feedback (thick green arrow) from that same feature after an additional latency τ_a . In such a neuron, we expect to see a cross-covariance (CC) peak (red arrow) between singing-related and playback-evoked spike bursts (black vertical bars) at a time lag (the so-called mirroring offset, red horizontal bar) given by the delay of the sensorimotor loop $\tau = \tau_m + \tau_a$. Under a predictive inverse (B), generated by stereotyped sequences of song features (ABC-ABC), the motor neuron 2 again triggers song feature B, but at the same time receives reliable feedback from the previous song feature A (thick green arrow). Thus, we expect to see a CC peak at a time lag much smaller than the sensorimotor loop delay τ . Finally, under a random sensory-to-motor mapping (C), we expected no CC between the motor- and sensory-evoked firing.

the singing-related spike train with the onset of song, the playback train will lag, or mirror the singing related train by a mirroring offset equal to the sensorimotor loop delay $\tau = \tau_a + \tau_m$. In more subtle scenarios in which the motor neuron has selectivity for multiple acoustic features, or is subject to greater degrees of noise, this temporal alignment between the motor and sensory responses of the neuron can still be detected through the position of a peak in the cross-covariance function (CC function) between the playback and singing related spike trains both time-aligned to song (Fig. 1A, Lower), even when the offset may not be visually apparent by simply looking at spike trains (see *SI Methods* and Fig. S1 for a theory of the CC function in the case of multiple latency auditory responses).

Conversely, consider a motor neuron downstream of a predictive inverse model (i.e., neuron 2 in Fig. 1B). Just as before, because it is downstream of a predictive inverse model, when the bird is not singing, this neuron now has a sensory selectivity for the previous song feature, A. This selectivity occurs because a predictive inverse model takes a sensory stimulus, in this case song feature A, and generates a motor command for the next feature, in this case song feature B. Thus, when aligned to song, the playback spike occurs τ_a after syllable A while the singing related spike occurs τ_m before song feature B. The result is that the mirroring offset will be much smaller than the total sensorimotor loop delay τ , and the peak in the CC function will be much closer to zero time lag. Finally, in the random scenario in Fig. 1C, we expect no pronounced peak in the CC function.

In theory, whether to expect a causal or predictive inverse depends on the sequence stereotypy of produced song features

(23, 24). What are those features? A song feature could be a song syllable, in which case the sequence of features is stereotyped because adult zebra finches sing stereotyped syllable sequences ABC-ABC (Fig. 1B). Stereotyped song sequences are generated mainly by stereotyped firing patterns in HVC, and according to previous theory (23, 24), we expect a predictive inverse and its associated signature of small mirroring offsets, to arise in HVC.

Alternatively, a song feature could be a brief pitch increase or decrease, in which case the features and their variable sequences (ABC, DBA; Fig. 1A) are mainly generated by variable LMAN firing patterns. Accordingly, we expect to find a causal inverse upstream of LMAN (23, 24). Indeed, in HVC neurons the mirroring offset between motor-related spiking and song-playback evoked spiking is less than 10 ms (25), much less than the roughly 40-ms loop delay of HVC estimated using electrical stimulation of HVC and using auditory stimulation of the ear (26–31). By contrast, mirroring offsets in LMAN have not been quantified yet, leaving it open as to whether they provide evidence for causal inverses.

Results

We first estimated LMAN’s motor latency τ_m by electrically stimulating LMAN neurons during singing using chronically implanted electrode pairs (Fig. 2A). Brief single or paired current pulses delivered at a random or a fixed time during a harmonic song syllable induced transient increases in frequency modulation (FM) of song (Fig. 2B and Fig. S2A). Transient song distortions (also of nonharmonic syllables) started in the range of

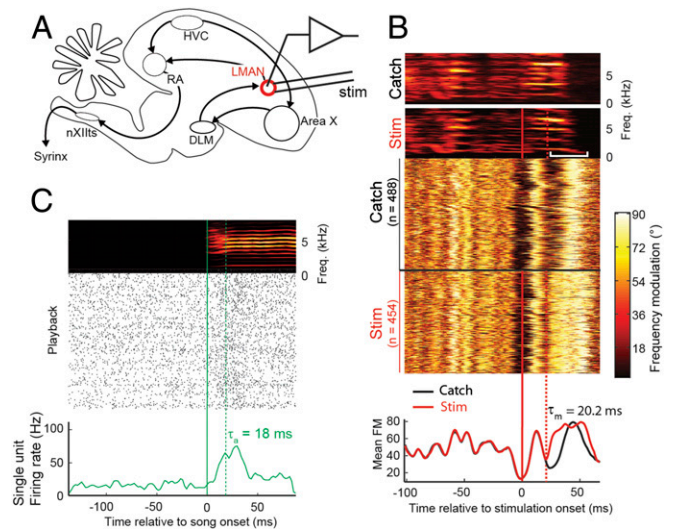


Fig. 2. LMAN sensorimotor loop delay. (A) Sagittal schematic of the songbird brain. Both HVC and the LMAN project to the premotor RA. DLM, dorsal lateral nucleus of the medial thalamus; nXIIIts, hypoglossal nucleus. (B) Electrical stimulation in LMAN using paired 0.2-ms current pulses of 500 μ A (separated by 1 ms) during song leads to transient distortions of song syllables (brief pitch decrease, red square bracket) compared with catch trials. (Top) Log-power sound spectrograms (high and low power shown in yellow and black, respectively) of a nonstimulated (catch) syllable and a stimulated (stim) syllable. A stack plot of frequency modulation (FM; Middle) and the mean FM (Bottom) across 488 nonstimulated syllables (catch trials) and 454 stimulated syllables (Stim) reveals a transient FM increase corresponding to a brief pitch decrease (white square bracket in the sound spectrogram) roughly 20 ms (dashed red line) after stimulation onset (time origin, thick red line). (C) Log-power sound spectrogram (Top), raster plot (Middle), and mean firing rate (Bottom) of a LMAN single unit with short auditory latency of 18 ms to playback onset of the bird’s own song ($n = 307$ playbacks).

function (Fig. 3B) peaked at a time lag in the range of 40–60 ms (suggesting a lead of motor activity on sensory responses).

To investigate mirroring also beyond single song motifs, we computed bout CC functions over entire song bouts (including song motifs and introductory notes). The resulting population CC function, normalized and averaged over all recording sites ($n = 48$ LMAN sites, $n = 7$ birds), also peaked near 40 ms (Fig. 3C), in agreement with the peak of the population CC function over motifs (Fig. 3B). In summary, at individual recording sites, we found a tendency for nonzero mirroring offsets. This tendency was amplified by averaging over recording sites to result in a significant average mirroring offset near 40 ms, suggesting a correspondence between sensory and motor responses.

To place the observed CC functions into perspective, we compared them to theoretical upper and lower bounds. We formed a lower bound by circularly shifting all playback-evoked spike trains by a random shift uniformly chosen within the duration of the song motif, thus eliminating any hypothetical temporal relationship between playback and motor responses. The resulting random-shift predictor (obtained after normalizing and averaging CC functions associated with randomly shifted spike trains) did not exceed the previously defined significance threshold of 3 Jackknife SDs (Fig. 3C), demonstrating that the nonzero mirroring offset we found could not simply occur by chance. We also estimated an upper bound of mirroring strength that corresponds to perfect mirroring limited only by intrinsic variability in playback-evoked responses. That is, we replaced all song-related spike-trains by one trial of a corresponding playback-evoked spike train, shifting it by -40 ms, and cross-correlated it with all other spikes evoked by the same song stimulus. If intrinsic variability was absent, all playback-evoked responses would be identical and the population CC function (which by virtue of division by the SD is a signal-to-noise estimate) would exhibit an infinitely high peak. Instead, we found an upper bound of the population CC function of 1.2, about six times above the observed peak of 0.2 in Fig. 3C. In summary, the LMAN mirroring strength of 0.2 (in units of signal to noise) at the 40-ms offset was more than five times higher than that of a random code (SD of random shift predictor) and about 17% of the upper bound corresponding to perfect mirroring.

We also estimated the temporal offset between LMAN sensory- and motor-evoked activity by inspection of the spike-triggered average (STA) sound amplitude, a curve that reports the average sound amplitude preceding and following a spike (Fig. S4 B and C). Sound amplitudes peaked on average 29 ms after LMAN motor spikes, and during playback, they peaked on average 8 ms before LMAN spikes ($n = 50$ LMAN sites), suggesting that LMAN neurons respond preferentially to increases in sound amplitudes and evidencing a combined motor-playback temporal offset of about 37 ms, which is within the range of estimated loop delays and mirroring offsets.

Discussion

Taken together, our mirroring offset analysis supports the hypothesis that LMAN neurons are part of, or downstream of, a causal inverse model (Fig. 1A). The roughly 40-ms mirroring offset we found was expectedly near the short end of estimated LMAN loop delays and much larger than mirroring offsets reported in HVC (25), all in agreement with a Hebbian view of sensorimotor integration in premotor areas (23, 24).

We found LMAN mirroring to be much weaker than the reported HVC mirroring. Among the sources contributing to weak mirroring in LMAN, we identified intrinsic noise in LMAN BOS responses. Intrinsic noise by itself limited mirroring strength at the level of individual neurons to a signal-to-noise ratio of less than 1.2. Obviously such noisy LMAN auditory responses hinder our ability to precisely measure sensorimotor correspondences. To mitigate the effects of intrinsic noise and response gating as much as possible, we performed the playback sessions while birds

were quietly resting in the dark. Possibly the sleep state was not very deep throughout the experiment, which may have reduced auditory responses and their correlation with singing-related activity (21, 32). Although noisy LMAN responses to BOS playback may limit our ability to perform mirroring analyses, they need not limit the bird's ability to learn inverse models through synaptic learning rules simply because birds learn naturally using auditory feedback during singing and not BOS playback in a quiescent state. In the auditory forebrain, signal-to-noise ratios of more than 10 have been observed (4), revealing that highly sensitive song-related auditory signals are present in the forebrain and that such sensitive signals could underlie the formation of causal inverses.

A shortcoming of our experiments is that we have tested the mirroring hypothesis only on a population level, although in theory it could be tested on a single-neuron level. Namely, if the motor latency of an individual neuron is τ_m and its auditory latency is τ_a , then the expected mirroring offset is $\tau_m + \tau_a$. A major obstacle to testing the single-neuron hypothesis is the difficulty of measuring motor latencies in single neurons: To do so, we would have to stimulate neurons during singing either alone or together with their cofiring neurons, which we were unable to do. Also, LMAN neurons could exhibit diverse auditory response latencies depending on both the time point in the motif at which they are stimulated and the feature composition of the auditory stimulus, as is the case for auditory responses in the primary auditory cortex analog field L (4). Thus, we imagine that a single LMAN neuron may exhibit a wide range of auditory latencies depending on the presynaptic neurons that drive its spiking at a given time. In the wake of this complexity associated with testing the mirroring hypothesis in single neurons, we tested for the existence of inverse models at a population level under the simplifying theoretical assumption that the eligibility trace for synaptic learning is a monotonically decaying function of time (nearby spike pairs lead to stronger potentiation than more widely separated spike pairs, which agrees with nearly all known spike-time dependent plasticity rules). This assumption implies that the population averaged mirroring offsets would lie within the earlier range of loop delays because synapses that mediate short-latency responses have higher eligibility and experience more potentiation than synapses that mediate long-latency responses (see *SI Methods* for a theoretical analysis).

In a sense, our theoretical prediction that mirroring offsets correspond to shorter loop delays is analogous to the immediacy effect in operant conditioning: In this effect, immediate reinforcement is more effective for modifying a response than delayed reinforcement. Indeed, birds can adapt their songs to escape from negatively reinforcing acoustic stimuli that are delivered during low-pitch renditions of their songs (33). In these operant conditioning experiments, birds were found to be unable to adapt their songs to escape negative reinforcement when the latency between pitch measurements and acoustic stimuli was 100 ms (33), revealing that birds are able to detect correlations between their songs and auditory stimuli when latencies of the latter are short but not when they are long.

The lag of 8 ms in our population analysis of STA sound amplitudes during playback (Fig. S4C) is short, even compared with the mean peak latency of 14 ms observed in spectro-temporal receptive fields (STRFs) in field L, the analog of primary auditory cortex (34). Although STRFs are not directly comparable to STAs, we conclude from the short 8-ms lag that presumably LMAN neurons are tuned to increases in sound amplitudes (such as occurring during syllable onsets) and not to peak amplitudes (such as occurring near the middle of syllables). Indeed, the average STA curve in Fig. S4C reaches its maximal slope about 25 ms before spikes, in agreement with the median lag of auditory response onsets.

In essence, our population analysis allowed us to conduct a first test of the inverse model hypothesis, but it leaves open several directions. Clearly, the range of validity of the causal inverse hypothesis remains unresolved; it remains to be investigated whether causal inverses apply to single neurons and even down to specific sound features generated by these neurons. In future work, to obtain improved latency estimates and to test the inverse model hypothesis on the single neuron level, it will be necessary to hold the signal of a single cell for a much longer time than the few minutes we were able to. There are promising new recording techniques that may provide the order of magnitude improvement required (35).

In the postulated Hebbian learning rule leading to formation of inverse models, synaptic strengthening occurs when presynaptic activity follows postsynaptic activity and not vice versa (24). This order allows sensory feedback arriving at motor neurons to be associated with past postsynaptic patterns of motor activity that could have caused this sensory feedback. Such rules have first been described in mormyrid electric fish (36) where they explain the formation of negative mirror images (collary discharges) of reafferent (self-generated) sensory input (2). In mammals, a similar rule has been found to describe synaptic connections from cortex to the basal ganglia (37). Thus, a causal inverse may reside in the efferent synapses of a cortical area upstream of the basal ganglia homolog; or, based on anatomy, a causal inverse could be connected to a dopaminergic ventral tegmental area (38, 39), thereby establishing a possible link of our findings with reinforcement learning theories (see below). Other candidate pathways for the causal inverse could lead to LMAN through HVC or through the thalamo-cortical projections from the dorsolateral nucleus of the thalamus (DLM). HVC projects to the basal ganglia loop and so indirectly to LMAN. However, inactivation of HVC in anesthetized birds had almost no effect on auditory responses of LMAN neurons (Fig. S5). Further experiments will be needed to explore these and possibly other pathways leading to LMAN, like the recently discovered LMAN shell (40) (note that we cannot rule out that some of the antidromically activated cells in our study were located in the LMAN shell rather than in LMAN because of the close proximity of LMAN and its shell, their combined efference to the arcopallium, and the uncertainty of recording sites in our experiments).

A causal inverse upstream of LMAN could be formed during a sensorimotor phase of song development and be beneficial to that development. That is, the causal inverse could be formed when young birds engage in motor explorations during the sensorimotor learning phase, in agreement with LMAN's known role in producing song variability in this critical phase (41, 42). Once the inverse mapping is established, it could be used to select downstream motor patterns in agreement with a sensory target and therefore steer the developing song in the right direction. By contrast, it is less clear how predictive inverses leading to HVC could be involved in motor learning. However, predictive inverses may aid vocal communication, especially when fast processing is required (43), e.g., during vocal exchanges with millisecond timescale precision in pair duetting birds (44) or during counter singing episodes (25).

Causal inverse models can in theory support some forms of single-trial imitation, but there are few reports on fast learning in zebra finches [although changes in the juvenile songs are identified after 2 d of exposure to the tutor's song (45), which motivates further investigations into this possibility]. A promising avenue into LMAN-mediated fast learning are recent experiments on negative song reinforcement in which LMAN, after having been transiently prevented from contributing to vocal output, was shown after recovery to instantaneously mediate vocal escape behaviors (41), which was suggested to provide support for a motor efference copy to LMAN (41). Alternatively, we propose that such instantaneous escape behaviors arise from

a causal inverse model upstream of LMAN. Our alternative explanation has the advantage that it invokes no postulated efference copy between a hypothetical premotor area and LMAN (but rather a causal inverse for which we find support). Moreover, in our scenario, LMAN can implement escape behaviors regardless of the source of variability (e.g., be it of neural or muscular origin), whereas in the efference copy scenario, escape would be confined to explorations mediated by the source of the efference copy (and by LMAN).

If reinforcement signals (e.g., dopamine) were mediated via causal inverse models instead of being released nonspecifically as assumed in many computational models (46, 47), then motor learning could be more efficient, in accordance with model-based reinforcement strategies (48). Indeed, simple reinforcement learning strategies can be enhanced with inverse models as a means to solve the structural credit assignment problem inherent in reinforcement learning (49). Our findings can thus be seen as providing neural support for efficient learning capabilities in basal ganglia pathways. Indeed many physiological and theoretical studies of learning in the basal ganglia have focused on the role of dopaminergic afferents onto striatal neurons in carrying reward prediction errors that are useful for reinforcement learning. On the other hand, the results of ref. 37 reveal that afferent cortical synapses onto striatal neurons also exhibit plasticity in which postsynaptic before presynaptic firing strengthens the synapse, a rule useful for inverse model learning, according to the theory of refs. 23 and 24. These results, when combined with evidence presented here for a causal inverse model in the cortical output area of a basal ganglia pathway, suggest that cortico-basal ganglia pathways could contain rich physiological mechanisms capable of combining inverse model and reinforcement based learning.

In general, learning to imitate any complex action poses a challenging problem for sensorimotor circuits: a sensory representation of the action must be converted into motor activity patterns capable of reproducing the action. Our approach of focusing on mirroring offsets and to relate these offsets to the existence and type of inverse models may be useful not only for uncovering basal ganglia circuit mechanisms for vocal learning, but also for explaining a wide range of imitation behaviors beyond vocal learning. We speculate that the careful measurement of mirroring offsets may also provide insights into other mirror neuron systems including visual-tactile neurons in mammalian premotor cortex (50).

Overall, our work suggests we may need to enlarge our hypothesis space for cortico-basal ganglia function by developing and testing theoretical models in which both associative Hebbian plasticity and dopamine dependent reward driven plasticity could interact to mediate sophisticated model based reinforcement learning strategies. At a computational level, these multiple, interacting forms of plasticity could allow organisms to combine classical reinforcement and control-theoretic inverse models to efficiently learn sensorimotor behaviors.

Methods

We chronically recorded from LMAN neurons in $n=9$ freely moving adult male zebra finches (>90 d after hatch) using miniature motorized microdrives. Each song bout was played back a random number of times in the dark because auditory responses in premotor areas of birds are state dependent—they tend to be gated off in the awake and aroused bird but gated on during sleep (51, 52).

We analyzed song (S) and playback (P) spike trains $\rho_S(t)$ and $\rho_P(t)$ (mean-subtracted) that we aligned to the common sonogram and that were either restricted to stereotyped song motifs or to song bouts composed of several motifs and introductory notes. We averaged CC functions $C_{SP}(\tau) = \frac{1}{T} \int_0^T \rho_S(t) \rho_P(t+\tau) dt$ (with T representing motif or bout duration) over all playbacks of a given song motif/bout. We then computed the mean CC function (Fig. 3A, iv) by averaging over all song motifs/bouts. In the mirroring analysis, we excluded the data from three

LMAN single units because they were suppressed by playback and did not produce sufficient spikes to be correlated with motor activity.

Note that the key signature of inverse models is the CC between spike trains in motor and precisely corresponding sensory states. In principle, one could observe a peak in CC evidencing an inverse model even if spiking during song and playback was uniformly distributed in time, i.e., even if song and playback-evoked firing rate curves in Fig. 3*B*, *ii* and *iii* and Fig. 5*4A*, *ii* were perfectly flat. In other words, the shapes of motif-aligned firing-rate curves may be indicative of mirrored responses but are useless to demonstrate their absence.

When computing population-averaged CC functions, to discount firing-rate differences among recording sites, we first normalized individual CC functions by their Jackknife estimate of SD before averaging (average in the interval [−150, 150] ms; Fig. 3 *B* and *C*). We assessed significance of peaks in population-averaged CC functions (Fig. 3 *B* and *C*) by their exceeding a threshold of 3 SDs estimated using Jackknifing over recording sites (3 Jackknife SDs σ correspond roughly to $P = 0.01$). We also explored

normalizing CC functions by the Jackknife estimate of SD in the interval [30, 50] ms, yielding qualitatively similar results.

All playback trials in which the bird produced a vocalization or other sound were discarded. Given the weak auditory responses in LMAN neurons, we recorded their spike responses during long periods of time to obtain sufficient statistics (order of hundreds of song playbacks within 10–20 min). Because it was difficult to maintain either good single-unit isolation or stable multiunit activity across these long periods, in the population plots (Fig. 3 *B* and *C*), we averaged clear single-unit responses with stable multiunit responses.

ACKNOWLEDGMENTS. We thank Klaus Hepp, Alexander Hanuschkin, and Walter Senn for helpful discussions and critical reading. This work was funded by Swiss National Science Foundation Grant 31003A_127024, the European Research Council (ERC) under the European Community's Seventh Framework Programme (FP7/2007-2013/ERC Grant AdG 268911), Burroughs Wellcome Foundation, Sloan Foundation, and Defense Advanced Research Planning Agency.

- Ahrens MB, et al. (2012) Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature* 485(7399):471–477.
- Bell CC (1981) An efference copy which is modified by reafferent input. *Science* 214(4519):450–453.
- Eliades SJ, Wang X (2008) Neural substrates of vocalization feedback monitoring in primate auditory cortex. *Nature* 453(7198):1102–1106.
- Keller GB, Hahnloser RHR (2009) Neural processing of auditory feedback during vocal practice in a songbird. *Nature* 457(7226):187–190.
- Poulet JFA, Hedwig B (2002) A corollary discharge maintains auditory sensitivity during sound production. *Nature* 418(6900):872–876.
- Poulet JFA, Hedwig B (2006) The cellular basis of a corollary discharge. *Science* 311(5760):518–522.
- Crapse TB, Sommer MA (2008) Corollary discharge circuits in the primate brain. *Curr Opin Neurobiol* 18(6):552–557.
- Webb B (2004) Neural mechanisms for prediction: Do insects have forward models? *Trends Neurosci* 27(5):278–282.
- Gergely G, Csibra G (2006) Sylvia's recipe: The role of imitation and pedagogy in the transmission of cultural knowledge. *Roots of Human Sociality: Culture, Cognition, and Human Interaction*, eds Enfield NJ, Levenson SC (Berg Publishers, Oxford, UK), pp 229–255.
- Meltzoff AN, Decety J (2003) What imitation tells us about social cognition: A rapprochement between developmental psychology and cognitive neuroscience. *Philos Trans R Soc Lond B Biol Sci* 358(1431):491–500.
- Giret N, Albert A, Nagle L, Kreutzer M, Bovet D (2012) Context-related vocalizations in African grey parrots (*Psittacus erithacus*). *Acta Ethol* 15(1):39–46.
- Funabiki Y, Konishi M (2003) Long memory in song learning by zebra finches. *J Neurosci* 23(17):6928–6935.
- Nottebohm F, Stokes TM, Leonard CM (1976) Central control of song in the canary, *Serinus canarius*. *J Comp Neurol* 165(4):457–486.
- Kao MH, Wright BD, 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(49):13232–13247.
- Olviczky BP, Andalman AS, Fee MS (2005) Vocal experimentation in the juvenile songbird requires a basal ganglia circuit. *PLoS Biol* 3(5):e153.
- Kao MH, Brainard MS (2006) Lesions of an avian basal ganglia circuit prevent context-dependent changes to song variability. *J Neurophysiol* 96(3):1441–1455.
- Stepanek L, Doupe AJ (2010) Activity in a cortical-basal ganglia circuit for song is required for social context-dependent vocal variability. *J Neurophysiol* 104(5):2474–2486.
- Hahnloser RHR, Kozhevnikov AA, Fee MS (2002) An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* 419(6902):65–70.
- Long MA, Jin DZ, Fee MS (2010) Support for a synaptic chain model of neuronal sequence generation. *Nature* 468(7322):394–399.
- Hessler NA, Doupe AJ (1999) Social context modulates singing-related neural activity in the songbird forebrain. *Nat Neurosci* 2(3):209–211.
- Hessler NA, Doupe AJ (1999) Singing-related neural activity in a dorsal forebrain-basal ganglia circuit of adult zebra finches. *J Neurosci* 19(23):10461–10481.
- Oztop E, Kawato M, Arbib MA (2013) Mirror neurons: Functions, mechanisms and models. *Neurosci Lett* 540:43–55.
- Hahnloser R, Ganguli S (2013) Vocal learning with inverse models. *Principles of Neural Coding*, eds Panzeri S, Quiroga P (CRC Taylor and Francis, Boca Raton, FL).
- Hanuschkin A, Ganguli S, Hahnloser RHR (2013) A Hebbian learning rule gives rise to mirror neurons and links them to control theoretic inverse models. *Front Neural Circuits* 7:106.
- Prather JF, Peters S, Nowicki S, Mooney R (2008) Precise auditory-vocal mirroring in neurons for learned vocal communication. *Nature* 451(7176):305–310.
- Day NF, Kinnischtzke AK, Adam M, Nick TA (2008) Top-down regulation of plasticity in the birdsong system: "Premotor" activity in the nucleus HVC predicts song variability better than it predicts song features. *J Neurophysiol* 100(5):2956–2965.
- Margoliash D, Fortune ES (1992) Temporal and harmonic combination-sensitive neurons in the zebra finch's HVC. *J Neurosci* 12(11):4309–4326.
- McCasland JS (1987) Neuronal control of bird song production. *J Neurosci* 7(1):23–39.
- McCasland JS, Konishi M (1981) Interaction between auditory and motor activities in an avian song control nucleus. *Proc Natl Acad Sci USA* 78(12):7815–7819.
- Troyer TW, Doupe AJ (2000) An associational model of birdsong sensorimotor learning II. Temporal hierarchies and the learning of song sequence. *J Neurophysiol* 84(3):1224–1239.
- Wang CZH, Herbst JA, Keller GB, Hahnloser RHR (2008) Rapid interhemispheric switching during vocal production in a songbird. *PLoS Biol* 6(10):e250.
- Doupe AJ, Konishi M (1991) Song-selective auditory circuits in the vocal control system of the zebra finch. *Proc Natl Acad Sci USA* 88(24):11339–11343.
- Tumer EC, Brainard MS (2007) Performance variability enables adaptive plasticity of "crystallized" adult birdsong. *Nature* 450(7173):1240–1244.
- Sen K, Theunissen FE, Doupe AJ (2001) Feature analysis of natural sounds in the songbird auditory forebrain. *J Neurophysiol* 86(3):1445–1458.
- Guitchounts G, Markowitz JE, Liberti WA, Gardner TJ (2013) A carbon-fiber electrode array for long-term neural recording. *J Neural Eng* 10(4):046016.
- Bell CC, Han VZ, Sugawara Y, Grant K (1997) Synaptic plasticity in a cerebellum-like structure depends on temporal order. *Nature* 387(6630):278–281.
- Fino E, Glowinski J, Venance L (2005) Bidirectional activity-dependent plasticity at corticostriatal synapses. *J Neurosci* 25(49):11279–11287.
- Person AL, Gale SD, Farries MA, Perkel DJ (2008) Organization of the songbird basal ganglia, including area X. *J Comp Neurol* 508(5):840–866.
- Gale SD, Person AL, Perkel DJ (2008) A novel basal ganglia pathway forms a loop linking a vocal learning circuit with its dopaminergic input. *J Comp Neurol* 508(5):824–839.
- Bottjer SW, Altenau B (2010) Parallel pathways for vocal learning in basal ganglia of songbirds. *Nat Neurosci* 13(2):153–155.
- Charlesworth JD, Warren TL, Brainard MS (2012) Covert skill learning in a cortical-basal ganglia circuit. *Nature* 486(7402):251–255.
- Bottjer SW, Miesner EA, Arnold AP (1984) Forebrain lesions disrupt development but not maintenance of song in passerine birds. *Science* 224(4651):901–903.
- Bonini L, Ferrari PF (2011) Evolution of mirror systems: A simple mechanism for complex cognitive functions. *Ann N Y Acad Sci* 1225:166–175.
- Fortune ES, Rodriguez C, Li D, Ball GF, Coleman MJ (2011) Neural mechanisms for the coordination of duet singing in wrens. *Science* 334(6056):666–670.
- Tchernichovski O, Mitra PP, Lints T, Nottebohm F (2001) Dynamics of the vocal imitation process: How a zebra finch learns its song. *Science* 291(5513):2564–2569.
- Frémaux N, Sprekeler H, Gerstner W (2010) Functional requirements for reward-modulated spike-timing-dependent plasticity. *J Neurosci* 30(40):13326–13337.
- Fiete IR, Fee MS, Seung HS (2007) Model of birdsong learning based on gradient estimation by dynamic perturbation of neural conductances. *J Neurophysiol* 98(4):2038–2057.
- Atkeson CG, Santamaria JC (1997) A comparison of direct and model-based reinforcement learning. *Proceedings of the IEEE International Conference on Robotics and Automation*, pp 3557–3564.
- O'Reilly RC, Frank MJ (2006) Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Comput* 18(2):283–328.
- Graziano MSA, Hu XT, Gross CG (1997) Visuospatial properties of ventral premotor cortex. *J Neurophysiol* 77(5):2268–2292.
- Dave AS, Yu AC, Margoliash D (1998) Behavioral state modulation of auditory activity in a vocal motor system. *Science* 282(5397):2250–2254.
- Schmidt MF, Konishi M (1998) Gating of auditory responses in the vocal control system of awake songbirds. *Nat Neurosci* 1(6):513–518.