
An Optimal Control Model of Zebra Finch Vocalization

Mike Schachter *

Helen Wills Neuroscience Institute
University of California, Berkeley
Berkeley, CA 94720
mike.schachter@gmail.com

Abstract

In this work, a nonlinear oscillator modeling the syringeal folds of the Zebra Finch is controlled by the state of higher level linear dynamical system. We formulate an optimal control cost function and solution for the learning of bird song.

1 Introduction

First we'll discuss the Zebra Finch vocalization system and it's mathematical formulation, and then the control of that system. All code used to generate figures and results in this paper can be found at:

<http://github.com/mschachter/birdy>

1.1 Zebra Finch Song

The Zebra Finch is an *oscine songbird*, it vocalizes complex sounds in order to defend territory and attract mates. A male juvenile Zebra Finch learns a unique song within 90 days of hatching, usually from his father.

The song of a Zebra Finch is organized hierarchically. At the largest time scale is a repeating *motif* that lasts 500-1000ms. The motif is comprised of smaller 20-50ms subunits called *syllables*. The syllables are comprised of sound primitives such as frequency sweeps, noise bursts, and harmonic stacks, collectively called *tones*. Figure 1 illustrates an example of a typical Zebra Finch song. For this project we have focused on the time scale of sound primitives, and more specifically the class of harmonic stacks.

1.2 The Zebra Finch Vocalization System

Like human speech, Zebra Finch song is generated at it's source by oscillations in airflow generated by vibrating vocal cords [3]. Figure 2 shows the Zebra Finch syrinx.

Each syrinx is modeled as a symmetric pair of nonlinear mass-springs. A model for these nonlinear oscillations is taken from [4]:

$$\begin{aligned}\dot{x} &= v \\ \dot{v} &= \gamma^2 \alpha + \gamma^2 \beta x - \gamma^2 x^3 - \gamma x^2 v + \gamma^2 x^2 - \gamma x v\end{aligned}\tag{1}$$

*With a metric ton of help and code from Hédi Soula: hsoula@gmail.com




Figure 1: Zebra Finch song, the sound pressure waveform (top) and associated spectrogram (bottom).




Figure 2: The Zebra Finch syrinx, with two independently controlled vocal cords. Taken from [3]

The control parameters of the model are α and β . $\gamma = 23500$ is a constant. The control space has a rich bifurcation structure, with both Hopf, homoclinic, and saddle-node in a limit cycle bifurcations leading to the birth of oscillations, as illustrated in figure 3. The type of bifurcation that occurs can determine the overall spectral content of the sound produced, as noted in [4].

1.3 Control of the Syrinx Model

The goal is to control the parameters α and β in order to produce an observed vocalization. Let $\phi(t) = [\alpha(t) \beta(t)]^T$ be the state vector.

Assume that the temporal evolution of $\phi(t)$ is defined by a controlled linear dynamical system:

$$\dot{\phi} = A(\phi(t) - \phi_{rest}) + u(t)$$

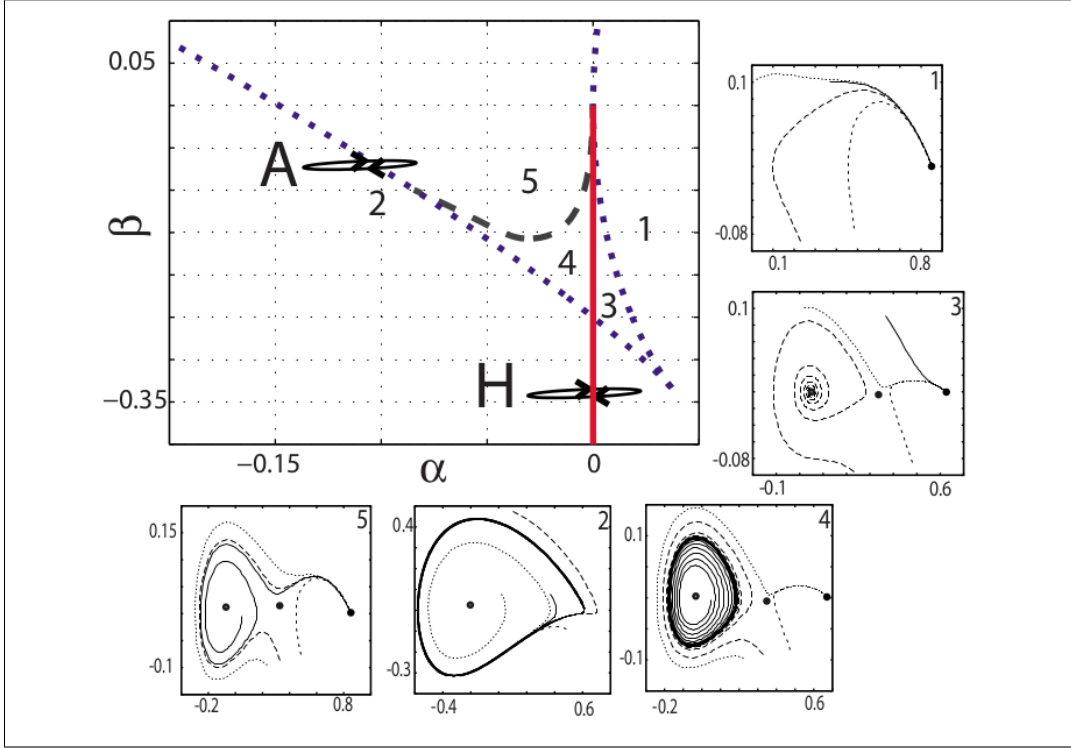


Figure 3: A bifurcation diagram with trajectories, taken from [4]. The red line is a Hopf bifurcation, the dashed black line is homoclinic, and the dotted blue line is a saddle-node bifurcation. Paths A and H represent trajectories that produce Hopf and Saddle-node in a limit cycle bifurcations.

where A is a 2×2 matrix, chosen to make the passive dynamics of the control system decay to some physiologically relevant rest state ϕ_{rest} .

We discretized time to simplify the analysis. Let $\Delta\tau$ be the time step for simulation of the control system, and define $t_k = k\Delta\tau$. Also define $\phi_k = \phi(t_k)$. Note that the time step for the simulation of the control is significantly larger than that of the oscillator. In practice we used the discrete time map representing the control system as given by the forward Euler step:

$$\phi_{k+1} = (A(\phi_k - \phi_{rest}) + \mathbf{u}_k) \Delta\tau + \phi_k \quad (2)$$

In optimal control theory [5], a cost function is specified and is to be minimized over time to produce an optimal control law. But what form should the cost function take? We'll make the following assumptions:

1. The control wants to keep the system's instantaneous energy low: $\phi_k^T \phi_k$.
2. The control wants to keep it's own instantaneous energy low: $\mathbf{u}^T \mathbf{u}$
3. The control wants to produce an instantaneous fundamental frequency that matches that of a stored template.

To elaborate on the last assumption, say we are the given time-varying fundamental frequency of a song syllable that we would like to learn, represented as the function $\hat{F}(t_k)$. Let:

$$\hat{F} = g(\phi_k) \quad (3)$$

be the function that gives the steady state fundamental frequency for a given control ϕ_k . The function g can be empirically estimated through simulation and approximated through interpolation. To follow this frequency means to keep the quantity $(F(t_k) - g(\phi_k))^2$ small.

Given these assumptions, the cost at time t_k of applying control \mathbf{u} is:

$$\ell_k(\phi_k, \mathbf{u}) = \phi_k^T \phi_k + \mathbf{u}^T \mathbf{u} + (F(t_{k+1}) - g(\phi_{k+1}))^2$$

Let N be the number of time points we want to control, and let $\pi_j = \{\mathbf{u}_j(\phi_j), \dots, \mathbf{u}_N(\phi_N)\}$ be the control law applied from time j to time N . Note that the control law is a sequence of functions! Each control uses feedback information about most recently observed state. The total cost of applying a control law π_0 given initial state ϕ_0 is:

$$v(\phi_0, \pi_0) = \sum_{k=1}^N \ell_k(\phi_k, \mathbf{u}_k)$$

1.4 Dynamic Programming Solution to Optimal Control

The optimal cost-to-go is a function that represents the total cost from a time j to time N given that the optimal control function is applied:

$$v_j^*(\phi_j) = \min_{\pi_j} \sum_{k=j}^N \ell_k(\phi_k, \mathbf{u}_k)$$

Dynamic programming is typically used to solve for the optimal cost-to-go [1]. When applying DP, we work backwards in time, creating a recursive algorithm that determines the functional form of the entire optimal control policy π_0^* . To illustrate this, we are going to take $N = 2$ with a specified initial condition ϕ_0 .

First the solution must be found for $N = 2$:

$$v_2^*(\phi_2) = \min_{\mathbf{u}_2} \ell_2(\phi_2, \mathbf{u}_2)$$

A straightforward but expensive way to solve would be to perform a search over a grid of values for ϕ_2 . For each value on the grid, we then will use gradient descent to minimize $\ell_2(\phi_2, \mathbf{u}_2)$ with respect to \mathbf{u}_2 . The end result is that we have a lookup table of values, giving us a function $\mathbf{u}_2(\phi_2)$, and a function $v_2^*(\phi_2)$.

The optimal cost-to-go at time $N - 1 = 1$ is

$$v_1^*(\phi_1) = \min_{\mathbf{u}_1} [\ell_1(\phi_1, \mathbf{u}_1) + v_2^*(\phi_2)] \quad (4)$$

We know the functional form of $\ell_1(\phi_1, \mathbf{u}_1)$ and the numerical form of $v_2^*(\phi_2)$. We also know that:

$$\phi_2 = (A\phi_1 + \mathbf{u}_1) \Delta\tau + \phi_1$$

so we can again use gradient descent for a grid on ϕ_1 to create a lookup table for $\mathbf{u}_1(\phi_1)$ and $v_1^*(\phi_1)$. We now have a functional form for the optimal control policy: $\pi_0^* = \{\mathbf{u}_1(\phi_1), \mathbf{u}_2(\phi_2)\}$

Now going forward in time, with the knowledge of a start state ϕ_0 , we can compute ϕ_1 , and then the optimal control $\mathbf{u}_1(\phi_1)$. Given \mathbf{u}_1 , we can then compute ϕ_2 and then \mathbf{u}_2 . The technique is generally applicable for any N .

1.4.1 A More Efficient Implementation

Solving the dynamic programming problem at a time t_k requires finding a function $\mathbf{u}_k(\phi_k)$ that minimizes $v_k(\phi_k)$ for any $\phi_k \in \mathcal{P}$, where \mathcal{P} is a pre-specified set of realistic values. This requires us to uniformly sample from \mathcal{P} and solve an optimization problem for each sample, a very computationally expensive prospect!

To reduce the computational burden, we first model the function $\mathbf{u}_k(\phi_k)$ as a linear combination of M radial basis functions $\{\psi_1, \dots, \psi_M\}$. Each basis function takes the form:

$$\psi_i(\phi) = \exp\left(-\frac{\|\phi - \theta_i\|^2}{\sigma_i^2}\right) \quad (5)$$

where θ_i is the center of the basis function, σ_i is the bandwidth of the basis function. Each basis function is multiplied by a coefficient $\mathbf{c}_i = [c_{i1} \ c_{i2}]^T$. Let $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_M]$ be the $2 \times M$ matrix of coefficients, the control is computed for a given \mathbf{C} as:

$$\mathbf{u}_k(\phi_k, \mathbf{C}) = \sum_{i=1}^M \mathbf{c}_i \psi_i(\phi_k) \quad (6)$$

We then formulate a single optimization problem to minimize the sum of cost across all points in \mathcal{P} for a given \mathbf{C} :

$$\mathbf{C}^* = \underset{\mathbf{C}}{\operatorname{argmin}} \sum_{\phi \in \mathcal{P}} v_k(\phi, \mathbf{C}) \quad (7)$$

The optimal control at time t_k is then given as $\mathbf{u}_k(\phi_k, \mathbf{C}^*)$.

To create an initial guess for the optimization problem, each \mathbf{c}_i is initialized independently by finding the minimum cost control at θ_i , the center of basis function i :

$$\mathbf{c}_i^0 = \underset{\mathbf{c}_i}{\operatorname{argmin}} v_k(\theta_i, \mathbf{c}_i) \quad (8)$$

2 Methods and Results

2.1 Simulation

Equation (1) was implemented in C++ and utilized the GNU Scientific Library. It was accessed and simulated using Cython/Python. The integration method used was an 8th order Runge-Kutta method [2] with a step size of $1 \mu\text{s}$. A GUI was created using PyQt in order to examine the dependence of oscillations with the control parameters, as shown in figure 4.

2.2 Mapping Controls to Fundamental Frequency

Equation (3) maps a control ϕ to a steady-state fundamental frequency \hat{F} , and required simulation to determine. First, a grid of points was constructed that spanned $\alpha \in [-1.25, 0.05]$ and $\beta \in [-1.10, 1.10]$, with a spacing of 0.05. For each point in the grid, the oscillator was simulated with a zero initial condition and the control specified by the grid point, for 15 ms.

The power spectrum of the resulting oscillation was taken and the fundamental frequency \hat{F} was identified as it's maximum. This produced samples for a real-valued function of two variables. These samples were fit with a radial basis function implementation from SciPy with a bandwidth 0.8 and Gaussian kernel. A RBF kernel was placed at each of the sample points, and allowed for a continuously-varying representation of equation (3). The sampling and it's RBF interpolation are shown in figure 5.

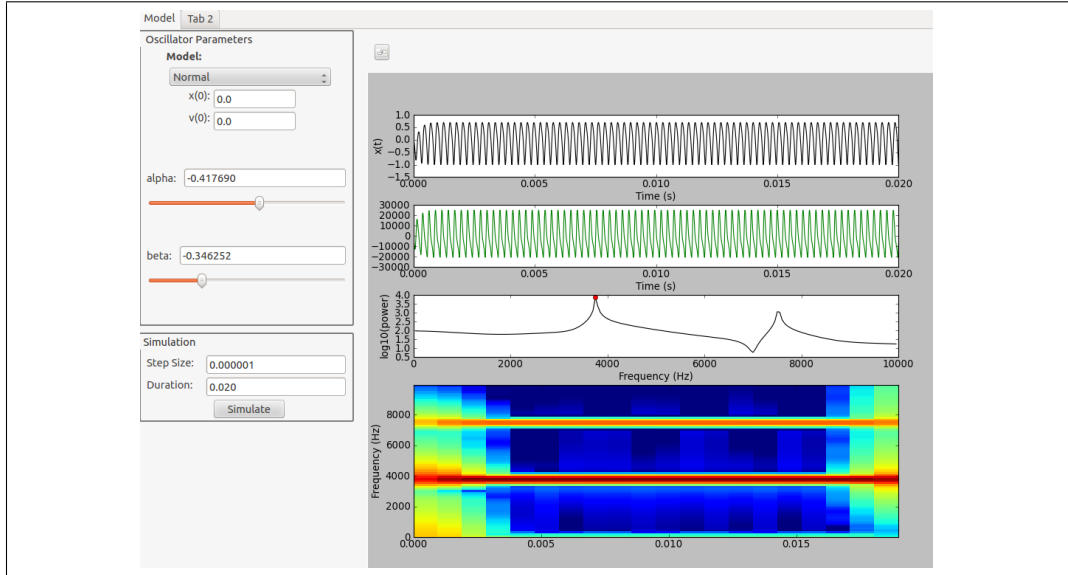


Figure 4: The UI for the simulation environment.




Figure 5: The simulated steady-state fundamental frequency as a function of control (left), and it's interpolation by radial basis functions (right).

2.3 Optimal Control

The first step in solving the optimal control problem was to choose the passive dynamics of the system defined by (2). We chose the following matrix:

$$A = \begin{bmatrix} -1000 & 0 \\ 0 & -900 \end{bmatrix}$$

and specified $\phi_{rest} = [-0.30 \ 0.30]^T$, a value where only one stable fixed point exists and no oscillations can occur. The passive dynamics had the phase plot specified in figure 6. The values of A were specified such that both α and β decayed to their rest points within 10ms.

Once the passive dynamics were selected, we applied the algorithm described in section 1.4.1. First, we created a grid of points for u that spanned $[-1, 1] \times [-1, 1] \subset \mathbb{R}^2$, with a spacing of 0.25. Each point in the grid was a center θ_i as listed in equation (5). The bandwidths were chosen uniformly as $\sigma_i = 0.30$.

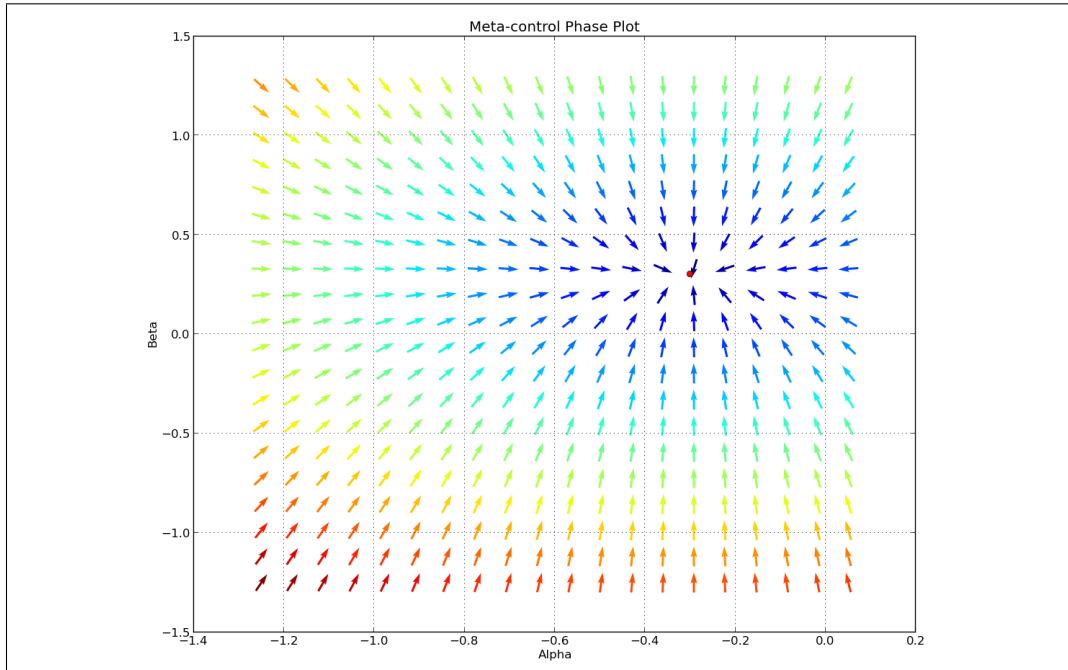


Figure 6: A phase plot for the passive dynamics of the control system. Color indicates magnitude, red is large, blue is small.

In the next step, the desired sequence of fundamental frequencies was specified with a sample spacing of $1ms$. The last frequency in the sequence was selected, and the dynamic programming optimization began.

The goal of the DP algorithm, for a given time step, was to determine the matrix C^* of equation (7). C^* was then used in equation (6) to determine the optimal control.

Because the optimization problem of (7) is (probably) difficult and (possibly) nonconvex, and definitely expensive, the first step was to find a good initial guess for each element of C^* . So we evaluated a locally-optimal value at each center, solving the optimization problem specified by equation (8). This was accomplished using gradient descent, with an initial guess of $c_i^0 = [0\ 0]^T$. All gradients were computed using a finite-difference approximation.

Once C was initialized, we then ran the full optimization of equation (7). We only let it run for 10 iterations, with the intention of actually getting at least one result before the due date of this project!

The technique above was applied to the last time point. Using dynamic programming, we then solved for the second-to-last time point, applying a recursion to the cost similar to that shown in equation (4). The time for optimization increased linearly as a function of how far back in time we went, as shown in figure 7.

Given time constraints, we only had the opportunity to apply the equation to a single example, that of a constant desired fundamental frequency of 3000Hz for 10ms. The total time to solve this problem, computed as the cumulative sum of the times in figure 7, was 314 minutes.

For each time point, we determined the form of the function specified in (6). One such example of this function is shown in figure 8. Some explanation is in order. The upper left is the optimal cost-to-go function $v_{10}^*(\phi)$. It shows that the most expensive states are those that produce a mismatch of the fundamental frequency (see figure 5).

The lower left panel shows the value of u_0 for a given value of ϕ . Strongly negative values of u_0 will push α to be significantly more negative. The control at this time point seems to have the ability to push α to two different regions of the space, to the extreme left and to the extreme right. Given that time has run out to check for bugs, we can only hope that this makes sense...




Figure 7: The computation time for optimization at each time point.

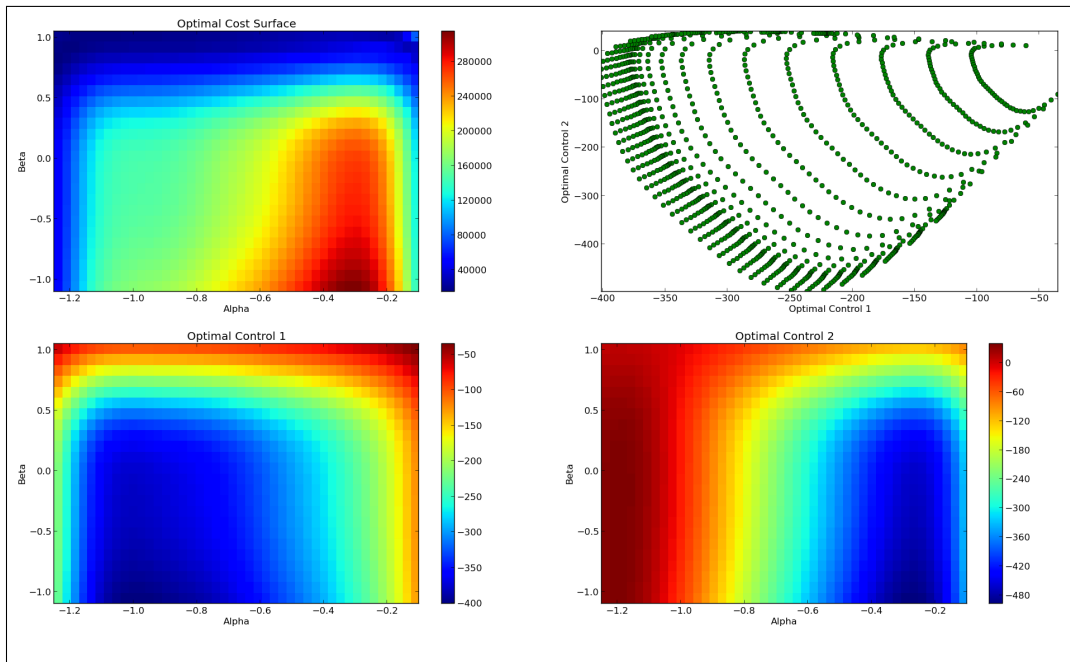


Figure 8: Upper Left: the optimal cost-to-go function at time point $N = 10$. Upper Right: the optimal control vector elements plotted against each other. Lower Left: optimal control 1 as a function of ϕ . Lower Right: optimal control 1 as a function of ϕ .

The lower right panel shows a tendency for the control u_1 to push β into a negative range, dependent on the value of α . It's probably best not to over-interpret this figure in the absence of a more rigorous analysis of bugs. The upper right panel shows u_0 vs u_1 , showing a very much nonlinear relationship between the two.

You may ask yourself at this point, "does it work"? The answer for now is "not yet". Figure 9 is a spectrogram of a 10ms simulation for the optimal control law applied to a given initial condition. There's plenty of power in the 3000Hz range, but clearly the fundamental frequency is closer to 1000Hz. More debugging is necessary!




Figure 9: A spectrogram for an optimally-controlled vocalization.

3 Conclusion

This work was an exciting foray into the world of nonlinear oscillators and optimal control theory. It's a work-in-progress, setting the stage for my thesis work, which will require the generation of synthetic Zebra Finch syllables.

The basic idea for future work will be to identify the "motor primitives" of Zebra Finch song. These motor primitives are the motor commands, i.e. the feedback control functions $u_k(\phi_k)$, that dictate the spectral content of a song for durations of 5-10ms.

The major bottleneck right now is computation. The dynamic programming algorithm takes a long time. There are many ways to speed it up. Some ways to do so may include coming up with a better parameterization for the control functions than RBFs, faster generation of initial guesses, implementation of an analytic gradient (if possible), and an optimization algorithm such as scaled conjugate gradient descent that should converge faster.

Note that there was no actual feedback used in this work. The Zebra Finch auditory system cannot provide sensory feedback in less than 20-50ms, which is often longer than the duration of the syllable. Instead of using feedback, we predicted the fundamental frequency through an empirically estimated function (equation (3)). We used what is called a *corrolary discharge*, a copy of the motor command, and ran it through a *forward model*, a model that takes a motor command and produces a prediction of the output state. By doing so we formulated the error in the predicted fundamental frequency, and tried to minimize that error.

The use of forward models and corrolary discharge is very general and has been used to implement a Zebra Finch learning model in [6]. We hope to extend this methodology to sound features that span longer periods of time. The idea would be to start building associations between the properties of a syllable at several time scales, and in doing so actually simplify the function form of the feedback control laws. We can begin to explore the various timescales that make learning easier, and explore their representation by neurons in the Zebra Finch auditory cortex.

Acknowledgments

Thank you to Hédi Soula, who created a C++/Python implementation of this model in our lab and has provided ideas, expertise, and feedback on this project. A thousand thanks to Jim Crutchfield, who facilitated an excellent NCASO course and offered lots of advice on this project.

References

- [1] BERTSEKAS, D. P. *Dynamic Programming and Optimal Control*. Athena Scientific, 2000.
- [2] DORMAND, J., AND PRINCE, P. A family of embedded runge-kutta formulae. *Journal of Computational and Applied Mathematics* 6 (1980), 19–26.
- [3] MINDLIN, G. B., AND LAJE, R. *The Physics of Birdsong*. Springer, 2005.
- [4] SITT, J. D., M, E., ARNEODO, GOLLER, F., AND MINDLIN, G. B. Physiologically driven avian vocal synthesizer. *Physical Review E* 81, 031927 (2010).
- [5] TODOROV, E., AND JORDAN, M. I. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience* 5 no 11 (2002).
- [6] TROYER, T. W., AND DOUPE, A. J. An association model of birdsong sensorimotor learning i. efference copy and the learning of song syllables. *J. Neurophysiol* 84 (2000), 1204–1223.