Time and Output Warping of Control Systems: Comparing and Imitating Motions

Peter Kingston, Magnus Egerstedt

Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332, USA

Abstract

The problem of making one system "mimic" a motion generated by another is studied. To address this, an optimal tracking problem is introduced that, in addition to tracking, optimizes over functions that deform or "warp" the time axis and the output space. Parametric and nonparametric versions of the time-warped tracking problem are presented and reduced to standard Bolza problems. The output warping problem is treated for affine and piecewise affine output warping functions. Example applications are given, including that of controlling a marionette to mimic a human dancer.

Key words: Time warping; output warping; optimal control; tracking; robotics.

1 Introduction

The ineffable quality of "subjective similarity" is difficult to formalize, especially when it is items of different kinds or classes that are compared. Nevertheless, the rather ill-posed question of how to define it lies at the heart of a number of areas of human-robot interaction, particularly Programming by Demonstration, (e.g. Billard et al. (2007),Breazeal and Scassellati (2002)), where humans ask robots to behave "like" them. In these areas, and particularly when goals are aesthetic or artistic, the question of what constitutes similarity is inescapable.

In this paper we fix a particular, control-aware definition for similarity a priori. Namely, imitation is defined in terms of a novel optimal tracking problem, in which the robot is given the flexibility to optimize over homeomorphisms that "warp" or "deform" both the output space and the time axis. An interesting aspect of this problem is that the choice of warping functions is inextricably coupled with the design of the control input that achieves the tracking.

The first portion of this paper is concerned with the deformation of the time axis, and is motivated by the recognition that the controlled system either simply may not be able to move as quickly as a human, or may have natural modes of oscillation (imagine a marionette swinging from its strings) that can be exploited by interpreting time liberally. For instance, a marionette actuated by weak servos whose arm swings with a natural frequency of 1 Hz may be asked to mimic a human waving her hand at 2 Hz. By choosing to interpret the human's wave at a slower speed, the oscillatory dynamics of the marionette can be used, rather than fought against, to mimic the wave. To likewise illustrate output warping, suppose that the human waves at head level yet the puppet cannot raise its hand above shoulder height. By translating the human's "wave" down to shoulder level, the puppet can work around this limitation, the motion can be tracked, and the human's gesture can be convincingly imitated. Combining time and output warping, in this example one obtains a motion that is lower and slower than the human's, but still recognizably the same.

For the case when one simply wants to compare signals rather than control systems, the *time warping* problem has been studied extensively in the automatic speech recognition literature, where this related problem is solved with a dynamic programming algorithm known as Dynamic Time Warping (DTW) (for example, see Sakoe and Chiba (1971), C. Myers (1980), Waibel and Yegnanarayana (1978)). The first portion of this paper is inspired by that work, but differs in that we are now interested in the *control* of systems, and since *the control problem and the "warping" problems are bidirectionally*

Email addresses: kingston@gatech.edu (Peter Kingston), magnus@ece.gatech.edu (Magnus Egerstedt).

coupled, one cannot now simply perform dynamic time warping on a reference signal and then control the system to track it; rather, the two problems must be solved simultaneously.

A number of different but related problems have also been studied in the controls context, as *time reparametrization* (e.g., Pappas (1996)) or *path following* (e.g. Hauser and Hindman (1997),Skjetne et al. (2005),Aguiar and Hespanha (2007)). In Pappas (1996), for instance, infeasible reference trajectories for feedback-linearizable systems with bounded control inputs are made feasible by reparametrization. Certain minimum-time optimal control problems are introduced, and necessary conditions for the existence of time warping functions are given for this large class of systems. Exact and asymptotic tracking are considered, but finite-time approximate tracking – which we consider in this paper – is not.

Of the previous work, that most similar in approach to our own appears in Raptis et al. (2007), in which the authors discuss a number of time warping problems in which dynamical constraints need to be enforced. The primary concern of Raptis et al. (2007) is the comparison of motions for computer vision purposes rather than the control of systems, and this does however motivate different problem formulations and solutions. Linear time-invariant systems are studied; a combination of dynamic time warping and deconvolution is used to solve one formulation efficiently, and another formulation results in an optimal control problem. Key differences from our work include that in Raptis et al. (2007) time warping is applied to the input to the systems rather than the output as in our case, and in Raptis et al. (2007) the control effort required to effect the motion is not penalized in the optimization problem. These differences follow naturally from the different goals of the two papers.

Conceptually, output warping is inspired in large part by work in "elastic" image deformation, which has applications to image registration (e.g., Kybic (2001), Uchida and Sakoe (2001)), video compression, and image processing (e.g., Akleman (1997)). As we will be "warping" not images but output trajectories, there are however very few technical similarities to our work in this literature (besides the use of simplicial complexes to define warping functions in Akleman (1997)).

Finally, before proceeding, we note that a more detailed exposition of what follows, including additional examples and figures, can be found in the technical report Kingston and Egerstedt (2011).

2 Time Warping

A *time warping function* $w : [0, T] \to \mathbb{R}$, $T \in \mathbb{R}_+$ is a strictly increasing, continuously differentiable function satisfying $w(0) = 0$; i.e., it is an increasing homeomorphism. We will denote the set of all time warping functions by Ω, and the set of derivatives of time warping functions by Ω'; i.e., Ω' is the set of continuous non-negative functions which are zero at at most a finite number of points.

Given a reference signal $r : [0, T] \to \mathbb{R}^p$, $T \in \mathbb{R}_+$, and another signal $y : \mathbb{R} \to \mathbb{R}^p$, the goal of time warping is to find a function w satisfying the above (i.e., $w \in \Omega$) that minimizes the functional $J : \Omega \to \mathbb{R}_+ \cup \{0\}$ defined,

$$J(w) = \int_0^T ||r(\tau) - (y \circ w)(\tau)||_Q^2 d\tau \quad \forall w \in \Omega \,. \tag{1}$$

For instance, consider the following example:

Example 2.1 *Suppose* $T = 2\pi$, $r(\tau) = \sin(\frac{\tau^2}{2\pi})$, *and* $y(t) = \sin(t)$ *for all* $\tau \in [0, T]$ *and* $t \in \mathbb{R}$. *Then the optimal warping function is given by,*

$$w^*(\tau) = \frac{\tau^2}{2\pi} \quad \forall \tau \in [0, T]$$

since $(y \circ w)(\tau) = \sin(\frac{t^2}{2\pi}) = r(\tau)$, *so* $J(w^*) = 0 \leq J(w)\ \forall w \neq w^*$. *Also note that although in this particular example, exact matching is possible (i.e., there exists a* $w \in \Omega$ *such that* $y \circ w = r$*), this is not generally the case.*

We can use the time warping idea to formulate a "time-warped" output-tracking problem as well. We will do this in two ways: First, we will solve the problem when "any" time warping function is allowed; we call this the *non-parametric* time warping problem. We follow this with a look at techniques that fix a particular parametrized form for the time warping function – we address polynomial functions – and which optimize over those parameters.

2.1 Tracking with Nonparametric Time Warping

Given a (possibly) nonlinear system of the form,

$$\begin{aligned} \frac{dx_t}{dt}(t) &= f(x_t(t), u_t(t), t) \\ y_t(t) &= h(x_t(t)) \end{aligned} \tag{2}$$

(with $x_t(t) \in \mathbb{R}^n$, $u_t(t) \in \mathbb{R}^m$, $y_t(t) \in \mathbb{R}^p$, and compatible dimensions for the domain and codomain of f and h) [1] we consider, to begin, the optimal control problem,

$$\min_{u,w} \int_0^T \left[||y_t(w(\tau)) - r(\tau)||_Q^2 + w'(\tau)||u_t(w(\tau))||_{R_u}^2 \right] d\tau$$

[1] Note that the subscript t is simply part of the function names x_t, y_t, etc, and is used to distinguish these functions from others to be introduced later.

where $r : [0,T] \to \mathbb{R}^p$ is a reference signal.[2] In other words, we would like the output y to "track" the reference signal r as closely as possible. However, note that this differs from the usual tracking problem by the introduction of the time warping function $w : [0,T] \to [0,\infty)$, and also in that we integrate tracking error over *reference time* ("τ time") instead of *system time* ("t time").

Defining $v \triangleq w'$ as the derivative of w; applying the Chain Rule to (2); augmenting the state with the time $t = w(\tau)$; defining $x \triangleq x_t \circ w$, $u \triangleq u_t \circ w$, and $y \triangleq y_t \circ w$; and moreover treating t as a function of time (so in fact $t = w$), we can state the following standard (Bolza) optimal control problem: Given the system,

$$\frac{d}{d\tau}\begin{bmatrix} x \\ t \end{bmatrix}(\tau) = \begin{bmatrix} v(\tau) f(x(\tau), u(\tau), t(\tau)) \\ v(\tau) \end{bmatrix} \qquad (3)$$
$$y(\tau) = h(x(\tau))$$

with known initial conditions $(x,t)(0) = (x_0, 0)$, minimize the cost functional (dropping τ arguments)

$$J_{\text{track}}(u,v) = \int_0^T \left[||y - r||_Q^2 + v||u||_{R_u}^2 \right] d\tau \qquad (4)$$

over the functions u and v, which can be viewed as control inputs to the system.

In fact, however, we will actually use a slightly modified cost that also penalizes large deviations of $w'(\tau)$ from one, both to regularize the problem and to capture the intuition that signals that must be "warped" by a great deal are more dissimilar than those that do not need to be warped as much; this is given below:

$$\begin{aligned} J &: L^2([0,T], \mathbb{R}^m) \times \Omega' \to \mathbb{R} \\ J(u,v) &= J_{\text{track}}(u,v) + J_{\text{timewarp}}(v) \\ &= \int_0^T \left[||y(\tau) - r(\tau)||_Q^2 + v(\tau)||u(\tau)||_{R_u}^2 \right. \\ &\qquad \left. + R_v (v(\tau) - 1)^2 \right] d\tau \end{aligned} \qquad (5)$$

This, together with (3), is the problem with which we will be interested in this section. Since it is also a Bolza problem, the solution can be found using standard optimal control theory, which we apply in the next section. First, however, we will present a brief discussion of the significance of dynamics to the time warping problem in order to build some intuition:

Example 2.2 *Consider the underdamped simple harmonic oscillator with state* $x_t(t) = (x_{t,1}, x_{t,2})(t) \in \mathbb{R}^2$ *and dynamics*

$$\frac{d}{dt}\begin{bmatrix} x_{t,1} \\ x_{t,2} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -\omega_0^2 & -2\zeta\omega_0 \end{bmatrix}\begin{bmatrix} x_{t,1} \\ x_{t,2} \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u_t \qquad (6)$$
$$y_t = [\,1\ 0\,] x_t$$

with $\zeta \in (0,1) \subset \mathbb{R}$, $\omega_0 \in \mathbb{R}_+$, *and suppose that we would like to solve the infinite-time problem (a modified version of (4)),*

$$\min_{u,w} \lim_{T\to\infty} \int_0^T \left[\frac{1}{T} ||(y_t \circ w)(\tau) - r(\tau)||_Q^2 + \frac{1}{w(T)} w'(\tau) ||(u_t \circ w)(\tau)||_{R_u}^2 \right] d\tau \qquad (7)$$

where the reference signal to be tracked is the sinusoid,

$$r(t) = \cos(\omega_r t)\ \forall t \in [0,\infty)\ . \qquad (8)$$

with $\omega_r \in \mathbb{R}_+$. *Moreover for clarity of exposition we will limit our attention to time warping functions of the form*

$$w(\tau) = \xi\tau \quad \forall \tau \in [0,T] \qquad (9)$$

for some $\xi \in \mathbb{R}_+$ *(Such* parametric time warping functions *are discussed in more detail in section 2.2).*

Before proceeding, we point out that in this example the control input which globally minimizes the cost (7) will not be unique. This is, informally, because components of signals which are bounded and occur for finite time "disappear" in the infinite-time average of (7). In fact, if u^* *is a minimizer for (7) (with* ξ *fixed), then so is* $u^* + \tilde{u}$ *for any bounded* $\tilde{u}$ *such that* $\lim_{t\to\infty} \tilde{u}(t) = 0$. *The consequence for this example is that it is only the behavior of the various signals as* $t \to \infty$ *which is of interest.*

We note that the presence of frequencies other than ω_r *in* $u_t \circ w$ *(and hence in* $y_t \circ w$*) increases both terms of (7),*[3] *so* $u_t \circ w$ *and* $y_t \circ w$ *must approach sinusoids with angular frequency* ω_r *as* $t \to \infty$*; without loss of generality (by the previous paragraph) we will assume that they are in fact sinusoids. Also observing that the phase of* $u_t \circ w$ *has no effect on the second term of (7), it must be that* $y \circ w = a\cos(\omega_r t)$ *for some* $a \in \mathbb{R}_+$, *and*

$$u_t(t) = \text{Re}\left(\frac{a}{h(i\omega_r/\xi)} e^{i\frac{\omega_r}{\xi}t} \right) \qquad (10)$$

where h *is the transfer function defined by* $h(s) = 1/(s^2 + 2\zeta\omega_0 s + \omega_0^2)$.

[2] We denote $||\cdot||_M^2 \triangleq (\cdot)^T M(\cdot)$ and assume $M = M^T \succ 0$, for whichever matrix M is used in the subscript.

[3] These arguments can be made rigorous using Plancherel's identity for Fourier series and considering a sequence of values for T that are multiples of $\frac{2\pi}{\omega_r}$; we have omitted this lengthier development for the purposes of our informal discussion.

It follows that the minimization problem (7) then reduces to,

$$\min_{a,\xi} \lim_{T\to\infty} \left[\frac{Q}{T} \int_0^T (a-1)^2 \cos^2(\omega_r \tau) d\tau + \frac{R_u}{\xi T} \int_0^{\xi T} \mathrm{Re}\left(\frac{a}{h(i\omega_r/\xi)} e^{i\frac{\omega_r}{\xi}t} \right)^2 dt \right] = \min_{a,\xi} \left[\frac{Q}{2}(a-1)^2 + \frac{R_u}{2} \left| \frac{a}{h(i\omega_r/\xi)} \right|^2 \right]. \quad (11)$$

For any fixed a, this is minimized with respect to ξ when the magnitude of the transfer function $|h(i\omega_r/\xi)|$ is maximized. This occurs when $\xi = \xi^* \triangleq \omega_r/(\omega_0\sqrt{1-\zeta^2})$ – that is, when ξ is chosen so that the frequency of the reference signal coincides with the resonant frequency of the "reference time" system matrix

The key point demonstrated by this example is that the steady-state effect of time warping is to scale the frequency axis (by the Fourier Dilation Theorem) so that passbands of the system coincide with concentrations of energy in the reference signal. (In fact, this effect is sufficient to overcome certain performance limitations, caused by unstable zero dynamics, that are inherent to the unwarped reference tracking problem; we refer the interested reader to Kwakernaak and Sivan (1972) and Aguiar et al. (2008).)

2.1.1 Optimality Conditions

Theorem 2.1 *The first order necessary optimality conditions for the minimization of (5) are (dropping τ arguments),*

$$2vu^T R_u + v\lambda^T \frac{\partial f}{\partial u}(x,u,t) = \mathbf{0}^T \quad (12)$$

$$||u||^2_{R_u} + 2R_v v(\tau-1) + \lambda^T f(x,u,t) + \mu = 0 \quad (13)$$

for all $\tau \in [0,T]$, where λ is the solution to the backwards differential equations (16), (17).

Proof : Taking (x,t) to be the state, (u,v) the control input, and (λ,μ) the costate, the Hamiltonian for this problem is,

$$H((x,t),(u,v),(\lambda,\mu),\tau) = L(x,(u,v),\tau) + v\lambda^T f(x,u,t) + \mu v \quad (14)$$

where

$$L(x,(u,v),\tau) = ||h(x) - r(\tau)||^2_Q + v(\tau)||u||^2_{R_u} + R_v(v(\tau)-1)^2 \quad (15)$$

and the costate equations are,

$$\begin{aligned} -\frac{d\lambda}{d\tau} &= \frac{\partial H}{\partial x}^T ((x,t),(u,v),(\lambda,\mu),\tau) \\ &= \frac{\partial L}{\partial x}^T (x,(u,v),\tau) + v\frac{\partial f}{\partial x}^T (x,u,t(\tau))\lambda \\ &= 2h'(x)^T Q(h(x)-r) + v\frac{\partial f}{\partial x}^T (x,u,t)\lambda \quad (16) \\ -\frac{d\mu}{d\tau} &= \frac{\partial H}{\partial t}((x,t),(u,v),(\lambda,\mu),\tau) \\ &= v\frac{\partial f}{\partial t}^T (x,u,t)\lambda \quad (17) \end{aligned}$$

with $(\lambda,\mu)(T) = 0$. Note that when f is not time-varying, $\frac{\partial f}{\partial t} = 0$, which gives the simplification that $\mu(\tau) = 0\ \forall \tau \in [0,T]$.

In any case, the first order necessary optimality conditions (FONCs) are,

$$\begin{aligned} &\frac{\partial H}{\partial u}((x,t),(u,v),(\lambda,\mu),\tau) = \\ &\qquad 2vu^T R_u + v\lambda^T \frac{\partial f}{\partial u}(x,u,t) = \mathbf{0}^T \quad (18) \\ &\frac{\partial H}{\partial v}((x,t),(u,v),(\lambda,\mu),\tau) = \\ &\quad ||u(\tau)||^2_{R_u} + 2R_v v(\tau-1) + \lambda^T f(x,u,t) + \mu = 0 \end{aligned}$$

for all $\tau \in [0,T]$. ∎

In fact, these equations (18) can be given the stronger interpretation of stating that the gradient of the functional (5) in the function space of which (u,v) is an element must be zero; we will leverage this interpretation in the later section 2.2.2 which describes an algorithm for computing the optimal (u,v).

Example 2.3 *The preceding, more general formulation can be used to solve the usual, more specific time-warping problem (1). Given signals y and r to be compared, we define the "signal generator" system,*

$$\dot{x}_t(t) = y'(t) \ , \ x_t(0) = y(0) \quad (19)$$

in which case the corresponding augmented system in reference time is

$$\frac{d}{d\tau}\begin{bmatrix} x \\ t \end{bmatrix}(\tau) = \begin{bmatrix} v(\tau)y'(t) \\ v(\tau) \end{bmatrix}. \quad (20)$$

The Hamiltonian is (letting $Q = 1$, $R_v = k \in \mathbb{R}$)

$$H((x,t),v,(\lambda,\mu),\tau) = (x - r(\tau))^2 + k(v-1)^2 + \lambda v y'(t) + \mu v, \quad (21)$$

and we obtain the costate equations

$$-\frac{d\lambda}{d\tau}(\tau) = 2(x - r(\tau)) \tag{22}$$

$$-\frac{d\mu}{d\tau}(\tau) = \lambda(\tau)v(\tau)y''(t(\tau)) \tag{23}$$

with $(\lambda, \mu)(T) = 0$. *Finally, the FONC is,*

$$\frac{\partial H}{\partial v}((x,t), v, (\lambda,\mu), \tau) = \mu + \lambda y'(t) = 0 \tag{24}$$

which converts the problem into a two point boundary value problem.

2.2 Tracking with Parametric Time Warping

In some situations, we may be interested only in time warping functions with a particular parametric form. One example is linear time-warping functions, which are of special interest since they represent a uniform scaling of the time axis. Another motive for investigating time warping functions with given parametric forms is the discretization of the problem for numerical solution.

To express these ideas, we introduce a parameter vector ξ in some parameter set $\Xi \subset \mathbb{R}^q$, and a *parametrization function* $\phi_v : \Xi \to \Omega'$ which, given a parameter vector, returns the derivative of a time warping function. Then, we are in fact considering the problem,

$$\min_{\xi,u} J(u, \phi_v(\xi)) \ . \tag{25}$$

Parametrization functions that return linear and polynomial time warping functions yield problems that can be treated nicely under the Bolza framework – and, more generally, under the assumption that Fréchet derivatives exist, we may apply the Chain Rule to (25) to compute the gradient of J with respect to ξ. This is explored in the following sections; more detail can be found in Kingston and Egerstedt (2011).

2.2.1 Polynomial Time Warping

For instance, polynomial time warping functions are of the form,

$$w(\tau) = \phi_v(\xi)(\tau) = \sum_{i=1}^{N_v} \xi_i \tau^i \tag{26}$$

for some integer $N \geq 1$, and with discrete parameter vector $\xi = [\xi_1, \ldots, xi_N] \in \mathbb{R}^{N_v}$ (Note that the requirement that $w(0) = 0$ implies that there is no constant term in the polynomial). These can be handled nicely under the standard Bolza framework by treating w' as the output of an autonomous system (a chain of integrators).

Specifically, we define the augmented system,

$$\frac{d}{d\tau}\begin{bmatrix} x \\ t \\ v_1 \\ \vdots \\ v_{N_v-1} \\ v_N \end{bmatrix}(\tau) = \begin{bmatrix} v_1(\tau)(f(x(\tau), u(\tau), t(\tau)) \\ v_1(\tau) \\ v_2(\tau) \\ \vdots \\ v_N(\tau) \\ 0 \end{bmatrix} \tag{27}$$

$$y(\tau) = h(x(\tau))$$

with partially-known initial conditions $(x,t)(0) = (x_0, 0)$, and seek to minimize (5) with respect to the initial conditions $v_1(0), \ldots, v_N(0)$, which define the polynomial. In order for t to be a valid time warping function, it is also necessary that this minimization be performed such that the constraint $v_1(\tau) > 0 \ \forall \tau \in [0, T]$ is satisfied. The FONCs are given in the subsequent theorem:

Theorem 2.2 *The FONCs for the polynomial time warping problem are given by (31) and*

$$\frac{\partial J}{\partial \xi}^T(\xi) = \text{diag}(1, \tfrac{1}{2}, \tfrac{1}{3!}, ..., \tfrac{1}{N_v!})\nu(0) = 0 \tag{28}$$

where ν *is the solution to (29).*

Proof : For the polynomial time warping problem, the Hamiltonian is (here we identify the costate by (λ, μ, ν), where $\nu(\tau) \in \mathbb{R}^{N_v} \ \forall \tau \in [0, T]$),

$$\begin{aligned} &H((x,t,v), u, (\lambda,\mu,\nu), \tau) \\ &\quad = L(x, (u, v_1), \tau) + \lambda^T v_1 f(x, u, t) + \mu v_1 \\ &\qquad + \nu^T [v_2, \ldots, v_N, 0]^T \\ &\quad = ||h(x) - r(\tau)||_Q^2 + v_1||u||_{R_u}^2 + R_v(v_1 - 1)^2 \\ &\qquad + \lambda^T v_1 f(x, u, t) + \mu v_1 + \nu^T [v_2, \ldots, v_N, 0]^T \end{aligned}$$

and we obtain the costate equation for ν, (those for λ and μ are the same as given in (16, 17)):

$$\begin{aligned} -\frac{d\nu}{d\tau}(\tau) &= \frac{\partial H}{\partial v}^T((x,t,v)(\tau), u(\tau), (\lambda,\mu,\nu)(\tau), \tau) \\ &= \begin{bmatrix} ||u||_{R_u}^2 + 2R_v(v_1 - 1) + \lambda^T f(x,u,t) + \mu \\ \nu_1 \\ \vdots \\ \nu_{N_v-1} \end{bmatrix} \\ &= \mathbf{0} \end{aligned} \tag{29}$$

for all $\tau \in [0, T]$, and with $\nu(T) = \mathbf{0}$. This gives us a

convenient method for minimizing over $\xi = v(0)$ since,

$$\frac{\partial J}{\partial \xi}^T(\xi) = \text{diag}(1, \tfrac{1}{2}, \tfrac{1}{3!}, ..., \tfrac{1}{N_v!}) \frac{\partial J}{\partial v(0)}^T(v(0)) = \text{diag}(1, \tfrac{1}{2}, \tfrac{1}{3!}, ..., \tfrac{1}{N_v!})\nu(0) \ . \quad (30)$$

In other words the gradient of the cost with respect to the polynomial coefficients is simply a scaled version of the initial value of the costate ν.

Additionally, the optimality conditions for the control input are,

$$\frac{\partial H}{\partial u}((x,t,v)(\tau), u(\tau), (\lambda,\mu,\nu)(\tau), \tau) = 2v_1 u^T R_u + v_1 \lambda^T \frac{\partial f}{\partial u}(x,u,t) = \mathbf{0}^T \quad (31)$$

for all $\tau \in [0,T]$. ■

Finally, we note that, in principle, since Theorem 2.2 is a restriction of Theorem 2.1, it can also be proven as a corollary following the approach of the next section (2.2.2), in which case (28) would be seen as the projection of the second equation of (18) onto the subspace of polynomial functions. That said, the introduction of the exosystem is particularly convenient, in that it reduces the problem to one that can be addressed within the standard Bolza framework.

2.2.2 The Chain Rule for Parametrization Functions

If the Fréchet derivatives of both the parametrization function ϕ and the cost J [defined in (5)] exist, then we may in fact apply the solution given in section 2.1 directly to the discretized problem (25) through a simple application of the chain rule.

Since we will be interested in discretizing not just the time warping function v but also the control input u, at this point we will introduce a second parametrization function $\phi_u : \Sigma \to L^2([0,T], \mathbb{R}^m)$ for some $N_u \in \mathbb{N}$ and parameter space $\Sigma \subset \mathbb{R}^{N_u}$. In other words, given some finite-dimensional $\sigma \in \Sigma$, the function ϕ_u returns a control input function $u \in L^2([0,T], \mathbb{R}^m)$. This yields the fully discretized problem,

$$\min_{\sigma,\xi} J(\phi_u(\sigma), \phi_v(\xi)) \triangleq \min_{\sigma,\xi} J_{\phi_u,\phi_\sigma}(\sigma,\xi) \ . \quad (32)$$

With the interpretation that the partial derivative of the Hamiltonian with respect to the control input (as a function of time) is the gradient of J with respect to the control input (projected onto the dynamical constraint),

$$\nabla_u J(u,v)(\tau) = \frac{\partial H}{\partial u}^T((x(\tau),t(\tau)),(u(\tau),v(\tau)),(\lambda,\mu)(\tau),\tau)$$
$$\nabla_v J(u,v)(\tau) = \frac{\partial H}{\partial v}^T((x(\tau),t(\tau)),(u(\tau),v(\tau)),(\lambda,\mu)(\tau),\tau)$$

for all $\tau \in [0,T]$. Then, applying the chain rule to (32), the partial gradients of J_{ϕ_u,ϕ_σ} with respect to σ and ξ are given by the inner products,

$$\nabla_\sigma J_{\phi_u,\phi_v}(\sigma,\phi) = \begin{bmatrix} \langle \nabla_{\sigma_1}\phi_u(\sigma), \nabla_u J(\phi_u(\sigma),\phi_v(\xi))\rangle_{L^2([0,T],\mathbb{R}^m)} \\ \vdots \\ \left\langle \nabla_{\sigma_{N_u}}\phi_u(\sigma), \nabla_u J(\phi_u(\sigma),\phi_v(\xi))\right\rangle_{L^2([0,T],\mathbb{R}^m)} \end{bmatrix}$$

$$\nabla_\xi J_{\phi_u,\phi_v}(\sigma,\phi) = \begin{bmatrix} \langle \nabla_{\xi_1}\phi_v(\xi), \nabla_v J(\phi_u(\sigma),\phi_v(\xi))\rangle_{\Omega'} \\ \vdots \\ \left\langle \nabla_{\xi_{N_v}}\phi_v(\xi), \nabla_v J(\phi_u(\sigma),\phi_v(\xi))\right\rangle_{\Omega'} \end{bmatrix}$$

where the inner products are defined in the L^2 sense. By solving the state and costate ODEs and evaluating the inner-product integrals numerically we thus obtain a principled and general way to discretize the problem.

Example 2.4 *Suppose ϕ_u and ϕ_v parametrize control inputs by uniform linear b-splines. In other words, the control inputs u and v are represented by uniformly spaced samples, and linear interpolation is used for the intermediate values. Then, $\phi_u(\sigma)$ and $\phi_v(\xi)$ can both be expressed as sums of triangular basis functions, as given below,*

$$\phi_u(\sigma)(\tau) = \sum_{i=1}^{K} \begin{bmatrix} \sigma_{1+m(i-1)} \\ \vdots \\ \sigma_{mi} \end{bmatrix} \text{tri}\left(\frac{K-1}{T}(\tau - i - 1)\right)$$
$$\phi_v(\xi)(\tau) = \sum_{i=1}^{K} \xi_i \, \text{tri}\left(\frac{K-1}{T}(\tau - i - 1)\right)$$

where $K \in \mathbb{N}$ is the number of samples used. Then, the functional gradients are given by,

$$(\nabla_{\sigma_i}\phi_u(\sigma,\xi))(\tau) = e_{\text{mod}(i-1,m)+1} \, \text{tri}\left(\frac{K-1}{T}(\tau - \lceil \tfrac{i}{m} \rceil - 1)\right)$$

for all $i \in \{1, \ldots, mK\}$, *and*

$$(\nabla_{\xi_i} \phi_v(\sigma, \xi))(\tau) = \mathrm{tri}\left(\frac{K-1}{T}(\tau - i - 1)\right)$$

for all $i \in \{1, \ldots, K\}$, *where* e_i *denotes the* i*-th element of the natural basis, and*

$$\mathrm{tri}(\tau) = \begin{cases} 1 - |\tau| & \text{if } |\tau| \leq 1 \\ 0 & \text{if } |\tau| > 1 \end{cases}. \tag{33}$$

This gives us the finite-dimensional partial gradient for J_{ϕ_u, ϕ_v} *with respect to* ξ_i,

$$\begin{aligned}(\nabla_\xi J_{\phi_u,\phi_v}(\sigma,\xi))_i = \\ \int_0^T \left[\frac{\partial H}{\partial v}((x,t,v)(\tau), u(\tau), (\lambda,\mu,\nu)(\tau), \tau)\right. \\ \left. \mathrm{tri}\left(\frac{K-1}{T}(\tau - i - 1)\right)\right] d\tau\end{aligned}$$

for all $i \in \{1, \ldots, K\}$. *The expression for* $\nabla_\sigma J_{\phi_u,\phi_v}$ *is similar.*

3 Output Warping

In the previous sections, we assumed that the spatial correspondence between values of the output signal and those of the reference signal was known, and that only the temporal correspondence needed to be determined. In this section, we will additionally assume that the spatial or output-space correspondence is unknown. This explicitly addresses the fact that it is not just the dynamics of the "mimicking" system that may differ from those of the system that generated the reference motion, but also spatial constraints and scales – a problem evident even in the prototypical example of a large industrial robot arm asked to imitate a human operator.

To treat the problem of spatial correspondence, we will assume that the reference signal r that we have been considering so far is in fact the composition of two functions: the "actual" reference signal $\bar{r} : [0, T] \to \mathbb{R}^p$, and an "output warping function" $s : \mathbb{R}^p \to \mathbb{R}^p$ of our choosing which transforms values of $\bar{r}$ before they are compared to those of the output signal y. In other words, $r = s \circ \bar{r}$.

More precisely, an *output warping function* $s : \mathbb{R}^p \to \mathbb{R}^p$ is a continuous bijective map with continuous inverse (That is, s is a homeomorphism from $\mathbb{R}^p$ to $\mathbb{R}^p$). We will denote the set of all such functions by $\mathcal{S}$.

We will additionally assume that s has a particular parametric form. This is expressed by saying that s is returned by a parametrization function $\phi_s : \mathcal{C} \to \mathcal{S}$, where $\mathcal{C}$ is some finite-dimensional parameter space. Specifically, $\mathcal{C}$ is a space diffeomorphic to $\mathbb{R}^{N_s}$ for some $N_s \in \mathbb{N}$.

With these definitions, we can extend the original cost functional (5) to obtain the new cost functional to be minimized,

$$\begin{aligned}\bar{J} &: L^2([0,T], \mathbb{R}^m) \times \Omega' \times \mathcal{C} \to \mathbb{R} \\ \bar{J}(u, v, c) &= \bar{J}_{\text{track}}(u, v, c) + \bar{J}_{\text{timewarp}}(v) + \bar{J}_{\text{outwarp}}(c) \\ &= \int_0^T \left[||y(\tau) - (\phi_s(c) \circ \bar{r})(\tau)||_Q^2 \right. \\ &\qquad \left. + v(\tau)||u(\tau)||_{R_u}^2 + R_v(v(\tau) - 1)^2\right] d\tau \\ &\quad + \bar{J}_{\text{outwarp}}(c)\end{aligned} \tag{34}$$

where $\bar{J}_{\text{outwarp}}$ is some cost used to penalize "large" output warpings, regularize the problem, and in certain cases enforce constraints; its form will be determined by the choice of ϕ_s and is discussed in more detail later.

Differentiating (34) to find the partial gradient with respect to c, we obtain the FONC,

$$\begin{aligned}\nabla_c \bar{J} = &- \int_0^T (\nabla_c \phi_s(c))(\bar{r}(\tau))^T Q\left[y(\tau) - \phi_s(c)(\bar{r}(\tau))\right] d\tau \\ &+ \nabla_c \bar{J}_{\text{outwarp}}(c) = \mathbf{0}\end{aligned} \tag{35}$$

which must be satisfied in addition to those given in section 2.1. Note that in this equation we assume for notational simplicity that c is a column vector, but as we have stipulated that $\mathcal{C}$ be diffeomorphic to $\mathbb{R}^{N_s}$ this is no restriction.

3.1 Affine Output Warping

Affine output warping functions are of the form,

$$\phi_s((M, z))(r) = s(r) = Mr + z \tag{36}$$

where $M \in \mathbb{R}^{p \times p}$ is an invertible matrix, and $z \in \mathbb{R}^p$. Hence, for affine output warping functions, $c = (M, z)$, $\mathcal{C} = \mathbb{R}^{p \times p} \times \mathbb{R}^p$, and $N_s = p^2 + p$.

In selecting an appropriate $\bar{J}_{\text{outwarp}}$ for this parametrization, we have the goals (1) of rewarding "smaller" transformations – that is, those "close" to the identity transformation in some sense – over "larger" ones; and (2) of ensuring that s remain a bijection – which means penalizing values of (M, z) for which M is singular or nearly-singular. The following function achieves these goals:

$$\begin{aligned}&\bar{J}_{\text{outwarp}}((M, z)) = \\ &\alpha \frac{\mathrm{tr}(M^T M)^p}{\det(M^T M)} + \beta \frac{1}{p} \mathrm{tr}[(M - I)^T (M - I)] + \gamma z^T z\end{aligned} \tag{37}$$

where $\alpha, \beta, \gamma \in \mathbb{R}_+ \cup \{0\}$ are scalar coefficients used to weight the relative "importance" of the various parts of the cost.

The first term is inversely proportional to the Gram determinant of the columns of M and so gives a measure of "how singular" M is. The presence of the Frobenius norm $\mathrm{tr}(M^T M)$ in the numerator makes this term somewhat independent of the absolute scaling of the matrix, since $\frac{\mathrm{tr}((\alpha M)^T(\alpha M))^p}{\det((\alpha M)^T(\alpha M))} = \frac{\mathrm{tr}(M^T M)^p}{\det(M^T M))}$ for all $\alpha \in \mathbb{R}$; that is, it is constant on lines passing through the origin in $\mathbb{R}^{p\times p}$. Furthermore, if $\sigma_1, \ldots, \sigma_p$ are the singular values of M, then,

$$p^p \left(\frac{\sigma_{\min}^2}{\sigma_{\max}^2}\right)^p \leq \frac{\mathrm{tr}(M^T M)^p}{\det(M^T M)} \leq p^p \left(\frac{\sigma_{\max}^2}{\sigma_{\min}^2}\right)^p \qquad (38)$$

so the $p-$th root of this term is always between the (squared) l_2 condition number for M and its reciprocal, up to a constant factor p.

The second term is the squared Frobenius norm of $M-I$, and gives the expected value of $||Mx - x||^2/||x||^2$ if x is a random variable drawn from a uniform distribution on a ball of (any) fixed radius centered at the origin in $\mathbb{R}^p$. Together with the third term, it penalizes "large" transformations – i.e., those that differ substantially from the identity map.

Hence for this problem we have,

$$\begin{aligned}\bar{J}(u, v, (M, z)) = \int_0^T \big[||y(\tau) - M\bar{r}(\tau) - z||_Q^2 + \\ ||u(\tau)||_{R_u}^2 + R_v(v(\tau) - 1)^2\big]\, d\tau \\ + \bar{J}_{\text{outwarp}}(c)\end{aligned} \qquad (39)$$

Combining (39) with (37) and taking partial gradients with respect to M and z we obtain the FONCs expressed by Theorem 3.1.

Theorem 3.1 *The additional FONCs for the problem (39) are,*

$$\begin{aligned}&\nabla_M \bar{J}(u, v, (M, z)) = \\ &\int_0^T 2Q\left(M\bar{r}(\tau) + z - y(\tau)\right) r^T(\tau) d\tau \\ &+ \alpha \frac{2}{\det(M^T M)} \big(-\mathrm{tr}(M^T M)^p M^{-T} + \\ &\qquad p\,\mathrm{tr}(M^T M)^{p-1} M\big) \\ &+ 2\beta \frac{1}{p}(M - I) = \mathbf{0}\end{aligned} \qquad (40)$$

and

$$\begin{aligned}&\nabla_z \bar{J}(u, v, (M, z)) = \\ &\int_0^T 2Q(M\bar{r}(\tau) + z - y(\tau)) d\tau + 2\gamma z = \mathbf{0}\end{aligned} \qquad (41)$$

which must be satisfied in addition to (18).

Before proving this, we present a few preliminaries with regard to notation. In what follows, $(x \mapsto [\text{expression}])$ denotes functions that take x as an argument and return [expression]; e.g., $(x \mapsto x^2)$ is the function that squares its argument. The notation $D_c(f)$ is used for the differential of a function f at a point c in the domain of f. Note that $D_c(f)$ is itself a (linear) function which can be evaluated; we denote its evaluation at h (the "direction" of variation of the argument to f) by $D_c(f)(h)$. For instance, $D_c(x \mapsto x^3)(z) = (h \mapsto 3c^2 h)(z) = 3c^2 z$. We denote the *gradient* of f at c by $\nabla f(c)$. The derivative is the vector *dual* to the gradient; that is $D_c(f)(h) = \langle \nabla f(c), h\rangle$ for all variations h of the argument to f, where $\langle \cdot, \cdot \rangle$ is an appropriate inner product. Continuing the preceding example, $\nabla(x \mapsto x^3)(c) = (x \mapsto 3x^2)(c) = 3c^2$.

We now give the proof of Theorem 3.1.

Proof : Considering the first term of (39),

$$\begin{aligned}&D_A\left(M \mapsto \frac{\mathrm{tr}(M^T M)^p}{\det(M^T M)}\right) \\ &= \mathrm{tr}(A^T A)^p D_A\left(M \mapsto \frac{1}{\det(M^T M)}\right) + \\ &\quad \frac{1}{\det(A^T A)} D_A\left(M \mapsto \mathrm{tr}(M^T M)^p\right)\end{aligned} \qquad (42)$$

where, since

$$\begin{aligned}&D_A\left(M \mapsto \det(M^T M)\right) \\ &= \left(dA \mapsto 2\det(A^T A)(\mathrm{tr}(A^{-1} dA)\right)\end{aligned} \qquad (43)$$

we have

$$\begin{aligned}&D_A\left(M \mapsto \frac{1}{\det(M^T M)}\right) \\ &= \left(dA \mapsto \frac{-2\,\mathrm{tr}(A^{-1} dA)}{\det(A^T A)}\right)\end{aligned} \qquad (44)$$

and

$$\begin{aligned}&D_A\left(M \mapsto \mathrm{tr}(M^T M)^p\right) \\ &= \left(dA \mapsto 2p\,\mathrm{tr}(A^T A)^{p-1}\,\mathrm{tr}(A^T dA)\right)\end{aligned} \qquad (45)$$

Substituting (44) and (45) into (42),

$$D_A\left(M \mapsto \frac{\operatorname{tr}(M^TM)^p}{\det(M^TM)}\right) = \left(dA \mapsto \frac{-2}{\det(A^TA)}\left\langle \operatorname{tr}(A^TA)^pA^{-T} - p\operatorname{tr}(A^TA)^{p-1}A, dA\right\rangle\right) \quad (46)$$

where the angle brackets denote the Frobenius inner product. This (46) is the dual to the gradient, which is given by (and replacing the symbol A by M),

$$-\frac{2}{\det(M^TM)}\left(\operatorname{tr}(M^TM)^pM^{-T} - p\operatorname{tr}(M^TM)^{p-1}M\right) .$$

Likewise differentiating the other terms we obtain (40) and (41). ∎

3.2 Piecewise Affine Output Warping

The essential idea of piecewise affine output warping will be that we divide the space $\mathcal{R}$ into some number of p-simplices, and use an affine warping function within each of these, chosen in such a way that the resulting piecewise function is continuous. In order to enforce that s remain a bijection, this will require both that the individual affine warping functions be full rank, and that the images of their domains remain disjoint.

To begin, let S (the "input simplices") be a finite pure simplicial p-complex covering $\mathcal{R}$, whose 1-skeleton is rigid; and R ("the output simplices") be another finite pure simplicial p-complex, which is isomorphic to S.[4] Basically, we will optimize over the positions of vertices in R, and use the induced simplicial map (which interpolates vertex positions barycentrically) as our output warping function.

In more detail: We denote the p-simplices contained in S by $S^1, \ldots, S^{|S|}$, and those contained in R by $R^1, \ldots, R^{|S|}$; that is, the p-simplices in S and R are indexed. We also denote the vertices (0-simplices) of S and R by $V_1^R, V_2^R, \ldots$ and $V_1^S, V_2^S, \ldots$, respectively, and the vertices of a given simplex R^i by $R_1^i, \ldots, R_p^i$. Then, we define the output warping function by,

$$s(\bar{r}) = \sum_{i=1}^{p} R_i^{\pi_S(\bar{r})}\beta_i(\bar{r}, S^{\pi_S(\bar{r})}) \quad (47)$$

[4] *Pure* means that the only simplices of dimension less than p are the faces of higher-dimensional simplices. *p-complex* means that the highest dimensional simplices are p-simplices. *Isomorphic* means that there is a bijection between elements of R and S that preserves topology.

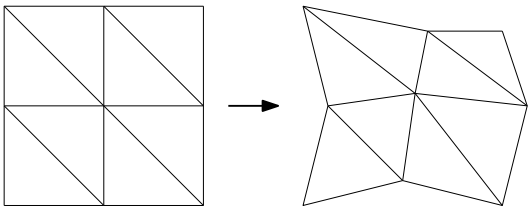

Fig. 1. Given the set S of input simplices, the output warping function s is determined by the positions of the vertices of the corresponding output simplices R. This example uses a Coxeter-Kuhn-Freudenthal tessellation of a regular grid of cubes in $\mathbb{R}^2$.

where $\pi_S(r) : \mathcal{R} \to \mathbb{N}$ is the function that, given a point $\bar{r} \in \mathcal{R}$, returns the index of the p-simplex in S containing $\bar{r}$, and $\beta : \mathcal{R} \times S \to \mathbb{R}^{p+1}$ is the function that, given a point $\bar{r} \in \mathcal{R}$ and a simplex $S^i \in S$, returns the barycentric coordinates of r in S^i if $r \in S^i$ and $\mathbf{0}$ otherwise. The map s is called the *simplicial map* induced by the vertex map from S to R.

Defining for each simplicial complex $\mathcal{K} \in \{S, R\}$ a graph $\mathcal{G}^{\mathcal{K}}$ whose vertices are the 0-simplices in $\mathcal{K}$, and in which an edge exists between two vertices iff they are both contained within the same 1-simplex, then a cost which tends to maintain the bijectivity of s is given by,

$$\bar{J}_{\text{outwarp}}(c) = 1/2\sum_{(V_i^R, V_j^R)\in\text{edges}(\mathcal{G}^R)}\left(||V_i^R - V_j^R||_K - ||V_i^S - V_j^S||_K\right)^2. \quad (48)$$

The idea here is that $\mathcal{G}^R$ is a rigid graph, and that by maintaining edge distances we ensure that simplices can neither "collapse" nor "collide." If $\mathcal{G}^R$ is visualized as a network of springs, then (48) gives their overall potential energy.

The partial gradient of $\bar{J}_{\text{outwarp}}$ with respect to each V_i^R is then,

$$\nabla_{V_i^R}\bar{J}_{\text{outwarp}}(c) = \sum_{V_j^R \in \mathcal{N}_{\mathcal{G}^R}(V_i^R)} \alpha_{ij}K(V_j^R - V_i^R) \quad (49)$$

where

$$\alpha_{ij} \triangleq \frac{||V_i^R - V_j^R||_K - ||V_i^S - V_j^S||_K}{||V_j^R - V_i^R||_K} \quad (50)$$

and $\mathcal{N}_{\mathcal{G}^R}(V_i^R)$ is the neighborhood of V_i^R in $\mathcal{G}^R$.

Admittedly, this cost does leave something to be desired, since simplices can collapse with finite energy. Nevertheless, we believe it is useful for its simplicity. One may wish to also apply (37) for each simplex in cases where (48) is not sufficient.

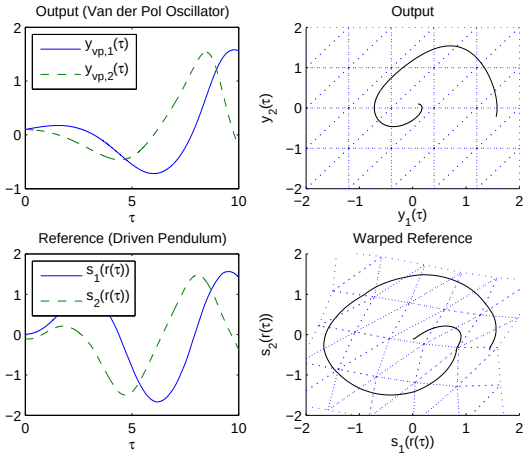

Fig. 2. Van der Pol oscillator vs. driven pendulum, after warping. Time warping matches the first part of the Van der Pol oscillator's transient to that of the pendulum (left top, bottom), and output warping rotates and deforms the reference output space to better match the output (right top, bottom).

Now, define $c = \left[(V_1^R)^T \ \ldots \ (|V^R_{\text{verts}(R)|})^T \right]^T$. Letting $i_1(r), \ldots i_{p+1}(r)$ be the indexes into verts(S) corresponding to the vertices of the simplex in S containing r, letting $\pi_S(r)$ be the simplex in S containing r, and defining the $p|\,\text{verts}(S)| \times p$ matrix $Z(r) = \left[I\alpha_1, \cdots, I\alpha_{|\,\text{verts}(S)|}\right]^T$ where $\alpha_{i_1(r)} = \beta_1(r, \pi_S(r)), \ldots, \alpha_{i_{p+1}(r)} = \beta_{p+1}(r, \pi_S(r))$ and $\alpha_i = 0 \, \forall i \notin \{i_1(r), \ldots, i_{p+1}(r)\}$, then the the partial gradient of (34) without the last term $\bar{J}_{\text{outwarp}}$ is given by,

$$\nabla_c(\bar{J} - \bar{J}_{\text{outwarp}})(R) = \\ -2\int_0^T Z(r(\tau))Q\left[y(\tau) - (\phi_s(c) \circ r)(\tau)\right] d\tau \ . \quad (51)$$

Hence the partial gradient of $\bar{J}$ with respect to c is simply the sum of (49) and (51).

We apply piecewise affine output warping together with linear time warping in the example of Figure 2.

3.3 The Puppet

To demonstrate the application of these ideas, we computed optimal controls for a simplified model of a marionette, which we wished to "mimic" the movements of a human as recorded by a motion capture system (Figure 4). The marionette (Figure 3) is modeled as a collection of ten point masses $m_1, \cdots, m_{10}$ joined by massless rods and suspended from four massless strings in two dimensions, as illustrated by Figure 5. The free endpoints $p_1, \cdots, p_4$ of the strings can be moved kinematically, and the remainder of the model is fully dynamic. The resulting mechanical system has 11 dynamic and 8 kinematic degrees of freedom, for a state space dimension of 30. The

Fig. 3. The marionette

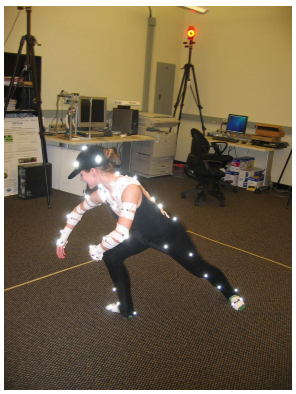

Fig. 4. Subject, dancing

strings are in principle represented by inequality constraints on the distance between the points at which their ends are anchored; in our model these constraints are for simplicity relaxed to exponentially-stiffening springs (for a complete treatment of marionette dynamics including inequality constraints, see Johnson and Murphey (2007)).

The dynamics of the state $x(t) = (q, \nu, p)(t)$ are then summarized,

$$\dot{q} = \nu \quad (52)$$

$$\dot{\nu} = M(q)^{-1}\left(-\sum_{i,j}\Gamma_{ij}(q)\nu_i\nu_j - \nabla U_{\text{grav}}(q) \right. \\ \left. -\nabla U_{\text{spring}}(q,p) - \mu\nu\right) \quad (53)$$

$$\dot{p} = u \quad (54)$$

where $M(q)$ is the mechanical system's mass matrix, each $\Gamma_{ij}(q) \in \mathbb{R}^{11}$ is a vector of Christoffel symbols of the first kind, U_{grav} is the gravitational potential, $\mu \in \mathbb{R}_+$ is a damping coefficient, and $U_{\text{spring}}(q,p)$ is the sum of spring potentials; specifically,

$$U_{\text{spring}}(q,p) = \sum_{(i,j)\in I_{\text{spring}}} k_1 \exp\left[k_2\left(||\rho_i(q) - p_j||^2 - l_{ij}^2\right)\right]$$

where $I_{\text{spring}} = \{(2,1),(7,3),(4,2),(8,4)\}$ is a set of index pairs summarizing the spring connections, each $\rho_i : \mathbb{R}^{11} \to \mathbb{R}^2$ is the forward kinematic map from the configuration q to the position of the i-th point, l_{ij} is the nominal length of the string connecting m_i and p_j, and $k_1, k_2 \in \mathbb{R}_+$ are chosen constants.

The output map h in this case depends only on the configuration component q of the state, and returns the positions $\rho_1(q), \cdots, \rho_{10}(q)$ of the ten masses $m_1, \cdots, m_{10}$. Our goal is that the output signal y track, in the time- and output- warped sense, the corresponding positions

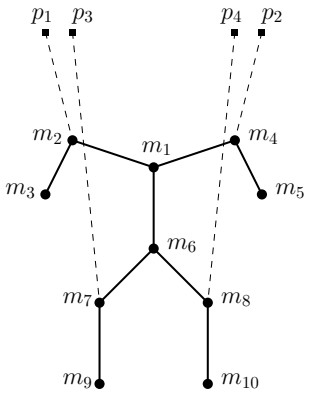

Fig. 5. The marionette is modeled as a system of point masses $(m_1, \cdots, m_{10})$ interconnected by massless rods (solid lines) and suspended by strings (dashed lines) from four kinematically-controllable points $(p_1, \cdots, p_4)$.

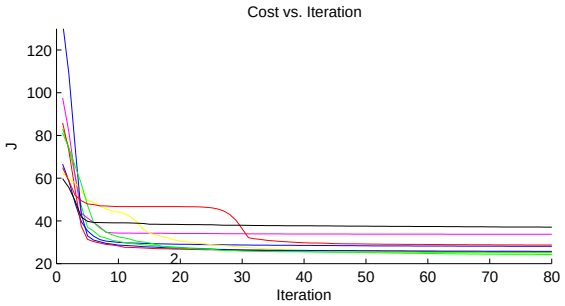

Fig. 6. The gradient descent procedure is characterized by a rapid initial descent followed by slower final descent. Each curve corresponds to a different initial guess for the control trajectory, drawn randomly from a space of sinusoids.

$\bar{r} = [\bar{\rho}_1, \cdots, \bar{\rho}_{10}]^T$ of points on a human subject. More specifically, the reference signal was created by a human dancer who performed the *bhangra* in a motion-capture environment. It consists of the coordinates of the subject's joints as computed by standard motion-capture software,[5] and projected onto a coronal (or *frontal*) plane. The optimization takes place with both time- and output- warping. The time warping is nonparametric, and the output warping optimizes over the six scalar parameters defining a two-dimensional affine transformation that is applied to all of the reference points.

The results are shown in Figures 6, 7, and 8, where it can be seen that marked improvements in similarity are achieved, both in numerical cost (Figure 6) and in subjective appearance (Figure 7). Both Figure 6 and Figure 8 also illustrate the point that, since the dynamical constraints are nonconvex, the local optimum to which the gradient descent procedure converges depends on the initial condition used to initialize the algorithm. Nevertheless, even for fairly widely-separated initial guesses in the puppet example, similar final costs are obtained.

[5] Vicon *VisionIQ* was used.

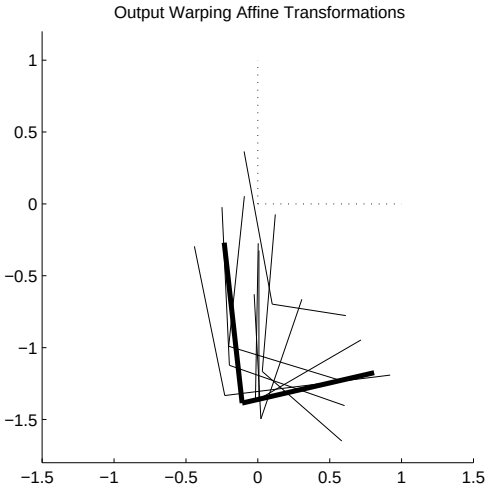

Fig. 8. The affine transformations arrived at by the optimization algorithm for different initial guesses (as described in Figure 6) are illustrated as frames in $\mathbb{R}^2$ (solid), along with the identity transformation (dotted). The lowest-cost transformation is drawn in bold.

Acknowledgments

This work was supported by the U.S. National Science Foundation through grant number 0820004. We would also like to acknowledge the help of Akhil Bahl in producing the bhangra data used in the dancing puppet example, and Amy LaViers (Figure 4) for ballet data.

References

Aguiar, A., Hespanha, J., 2007. Trajectory-tracking and path-following of underactuated autonomous vehicles with parametric modeling uncertainty. In: Automatic Control, IEEE Transactions on. Vol. 52. pp. 1362–1379.

Aguiar, A. P., Hespanha, J. P., Kokotović, P. V., Mar. 2008. Performance limitations in reference-tracking and path-following for nonlinear systems. Automatica 44 (3), 598–610.

Akleman, E., 1997. Making caricatures with morphing. In: SIGGRAPH: ACM Special Interest Group on Computer Graphics and Interactive Techniques. p. 145.

Billard, A., Calinon, S., Dillman, R., Schaal, S., 2007. Robot Programming by Demonstration. MIT Press, Ch. 59.

Breazeal, C., Scassellati, B., 2002. Robots that imitate humans. Trends in Cognitive Sciences 6 (1), 481–487.

C. Myers, L.R. Rabiner, A. R., Dec 1980. Performance tradeoffs in dynamic time warping algorithms for isolated word recognition. In: IEEE Transactions on Acoustics, Speech and Signal Processing. Vol. ASSP-6. pp. 623–635.

Hauser, J., Hindman, R., 1997. Aggressive flight maneu-

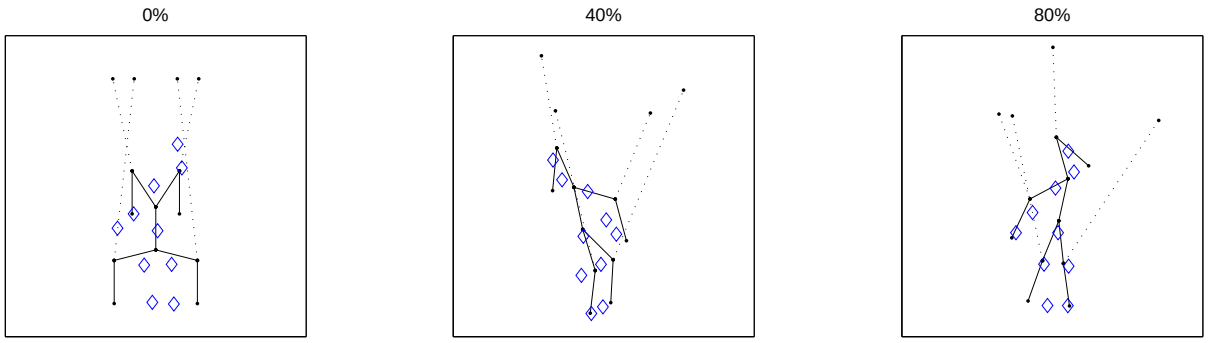

Fig. 7. Animation frames showing the lowest-cost imitation of the human subject's bhangra performance by the puppet; percentages are of the total playback time elapsed. The puppet begins hanging in an equilibrium state (at time "0%").

vers. In: Decision and Control, 1997., Proceedings of the 36th IEEE Conference on. Vol. 5. pp. 4186–4191.

Johnson, E., Murphey, T., 2007. Dynamic modeling and motion planning for marionettes: Rigid bodies articulated by massless strings. In: Robotics and Automation, 2007 IEEE International Conference on. pp. 330–335.

Kingston, P., Egerstedt, M., 2011. Time and output warping of control systems: Comparing and imitating motions. Tech. rep., Georgia Institute of Technology, `http://www.prism.gatech.edu/~pkingston3/papers/time_out_warp_2.pdf`.

Kwakernaak, H., Sivan, R., 1972. The maximal achievable accuracy of linear optimal regulators and linear optimal filters. IEEE Trans. on Automat. Contr. 17 (1), 79–86.

Kybic, J., July 2001. Elastic image registration using parametric deformation models. Phd in biomedical image processing, Ecole Polytechnique Fédérale de Lausanne, `http://cmp.felk.cvut.cz/~kybic/thesis/`.

Pappas, G., Dec 1996. Avoiding saturation by trajectory reparameterization. In: Proceedings of the 35th IEEE Decision and Control. Vol. 1. pp. 76–81.

Raptis, M., Bustreo, M., Soatto, S., 2007. Time warping under dynamic constraints. In: Eleventh IEEE International Conference on Computer Vision, Workshop on Dynamical Vision.

Sakoe, H., Chiba, S., 1971. A dynamic programming approach to continuous speech recognition. In: Proc. Int'l. Cong. Acoust.

Skjetne, R., Fossen, T. I., Kokotovic, P. V., 2005. Adaptive maneuvering, with experiments, for a model ship in a marine control laboratory. In: Automatica. Vol. 41. pp. 289–298.

Uchida, S., Sakoe, H., 2001. Piecewise linear two-dimensional warping. In: Systems and Computers in Japan. Vol. 32. pp. 1–9.

Waibel, A., Yegnanarayana, Feb 1978. Comparative study of nonlinear time warping techniques in isolated word speech recognition. In: IEEE Transactions on Acoustics, Speech and Signal Processing.