Reduced Dimensionality Control for the ACT Hand

Mark Malhotra, Eric Rombokas, Evangelos Theodorou, Emanuel Todorov, and Yoky Matsuoka

Abstract— Redundant tendon-driven systems such as the human hand or the ACT robotic hand are high-dimensional and nonlinear systems that make traditional control strategies ineffective. The synergy hypothesis from neuroscience suggests that employing dimensionality reduction techniques can simplify the system without a major loss in function. We define a dimensionality reduction framework consisting of separate observation and activation synergies, a first-order model, and an optimal controller. The framework is implemented for two example tasks: adaptive control of thumb posture and hybrid position/force control to enable dynamic handwriting.

I. INTRODUCTION

The amazing versatility and dexterity of the human hand remain unmatched by robotic hands in production today. The superiority of the hand is a result of its biomechanical properties, the neuromuscular control system, and the richness of sensory feedback available, but comes at a cost in complexity. The 15 joints of the human fingers are controlled by more than 30 muscles located in either the forearm or the palm acting through a network of tendons that eventually terminate at various connection points on the finger bones, resulting in a nonlinear system with a variable moment arm. A contentious subject in neuromuscular control is the extent to which models of all these intracacies are used for different levels of control and planning [1].

The hypothesis of "muscle synergies" offers a way for the nervous system to reduce the degrees of freedom (DOF) and thus simplify the control, learning, or planning problem: rather than independently activating muscles, groups of muscles are used in a coordinated manner corresponding to the kinematic [2], [3] or dynamic [4] regularities of the system itself or a task in the environment [5], [6]. Many studies have shown that this method could explain different aspects of the neuromusclar control system: for instance, frog leg wiping trajectories can be composed from premotor drive pulses [7] and a small subspace of possible movements is used by people when manipulating an object [2]. A number of different matrix factorization algorithms have been applied to such analyses [8].

By mimicking the biomechanical features of the hand, the Anatomically Correct Testbed (ACT) robotic hand [9] enjoys many of same benefits as the human neuromusclar control system and provides a platform for studying how synergies may be used to control the hand. Most biomimetic robotic

Fig. 1. The Anatomically Correct Testbed robotic hand uses a biomimetic tendon hood to transmit tendon forces to the fingers. Complex and dynamic tasks such as handwriting are enabled by a dimensionality reduction framework that simplifies the system model and control.

hands try to replicate the DOFs of the joints and the geometry of the fingers and palm. However, the ACT hand is unique in its actuation; rather than controlling each joint DOF with a pair of tendons [10], tendons pull through a crocheted tendon hood and along a surface defined by accurately reproduced bones. The ACT hand does not currently have embedded joint angle sensors, instead estimating the robot's pose from tendon lengths, as does a human from muscle spindles. These features result in system nonlinearities similar to those of the human hand that make traditional robotic control schemes either difficult to implement or ineffective. In particular, when controlled in tendon space, the system is underactuated because of the tendon force non-negativity constraint.

Task space control of underactuated systems has been studied for torque-driven systems [11] but must be adapted for tendon-driven systems; synergies provide a framework for controlling systems that are underactuated in tendon space by instead doing control in task space. Previous work has shown that employing linear observation synergies with PID control can allow multi-DOF manipulation tasks such as handwriting [12] or accurate single-DOF dynamic movements such as striking a piano key [13]. In this paper, we extend that work and show that a framework using synergies for both observation and control of the tendon-driven system enables dynamic movements, adaptive control, and hybrid position/force control much improved over using observation synergies and PID control alone.

II. SYNERGISTIC CONTROL FRAMEWORK

A. System Definition

Consider a tendon-driven system composed of n torque-driven motors, each with an encoder to measure the displacements, $\mathbf{L}$, of all n tendons. First-order system dynamics can

be expressed in discrete state space formulation as:

$$\mathbf{L}_{k+1} = A\mathbf{L}_k + B\mathbf{T}_k, \tag{1}$$

such that $\mathbf{L}_k \in \mathbb{R}^n$ is the state vector at timestep k, $\mathbf{T}_k \in \mathbb{R}^n \geq 0$ is the control vector, A is the $n \times n$ state transition matrix, and B the $n \times n$ control matrix. Note that transmitting force through tendons results in a non-negative control vector, so that this system is not always controllable in the n-DOF space of tendon lengths. This problem is often combatted by controlling in joint space, using tendons antagonistically [14] and possibly redundantly to ensure that the system is tendon-controllable [15]. However, the configuration of tendons in the human hand does not adhere to such design constraints. The following synergistic control framework 1) reduces the dimensionality of observation and control and 2) enables controllability for a tendon-driven robot without enforcing an antagonistic joint-space control.

B. Observation

Let the observation synergy S be an $s \times n (s < n)$ orthonormal matrix that is used to project $\mathbf{L}$ into a s-dimensional virtual space. Define $\mathbf{x}$ as the "synergy position":

$$\mathbf{x} = S\mathbf{L}. \tag{2}$$

The synergy position, $\mathbf{x} \in \mathbb{R}^s$, is a low-dimensional respresentation of the measured tendon lengths and provides a virtual state on which to control. S must be orthonormal so that all the different bases in the virtual space are separately observable, an important property for preserving system linearity. The selection of S alters the extent to which each tendon's displacement is observed in the virtual space and also defines the dimensionality of the virtual space. Two different methods for choosing S are presented in sections III-B and III-C .

Since S is orthonormal, the inverse operation is easily computed. For instance, a target in tendon length space can be reconstructed from the target synergy position by the equation:

$$\hat{\mathbf{L}} = S^T \hat{\mathbf{x}}. \tag{3}$$

This definition differs slightly from standard postural synergies [2]: tendon lengths are the input rather than joint angles. This is analogous to human's proprioceptive feedback from muscle spindles, which is considered the primary input for awareness of joint position in the absence of vision [16], [17]. Consequently, tendon-space synergies may account for both inter-joint coupling and for the coupled movement of multiple tendons that results from a single joint movement.

C. Control

The dimensionality of the control input can be reduced to simplify modelling and optimal control selection. Let the activation synergy W be a $n \times w (s < w < n)$ nonnegative matrix used to project the w-dimensional "synergy activation," $\mathbf{u}$, into torque commands for each of the n motors.

$$\mathbf{T} = W\mathbf{u}. \tag{4}$$

Thus, activating a single component of $\mathbf{u}$ produces a weighted activation of a set of motors. This is a common form of linear synchronous muscle synergies [1] usually used to interpret neural motor control. Although this strategy decreases the number of free parameters available to a controller, a proper W will encourage smoother control signals and reduce the complexity of modelling and feedback control problems. Two ways to determine W are dicussed in Section III.

D. Virtual Space System

Applying both the observation and control reduction models significantly simplifies the system. Now, the first order system dynamics are expressed in discrete state space as:

$$\mathbf{x}_{k+1} = \tilde{A}\mathbf{x}_k + \tilde{B}\mathbf{u}_k. \tag{5}$$

The state transition and control matrices, $\tilde{A}$ and $\tilde{B}$, are no longer $n \times n$, but instead $\tilde{A}$ is $s \times s$ and $\tilde{B}$ is $s \times w$. For instance a system of 6 tendons would require $(6 \times 6) + (6 \times 6) = 72$ parameters to represent the tendon-space system dynamics, but reduced to $s = 2$ and $w = 3$ instead has only $(2 \times 2) + (2 \times 3) = 10$ parameters. In terms of the original system dynamics,

$$\tilde{A} = SAS^T \tag{6}$$

$$\tilde{B} = SBW \tag{7}$$

E. Modelling and Optimal Control

Joint-space tendon-driven control methods can be applied to the virtual space system with the synergy position substituted for joint angle and the synergy activation substituted for tendon force. Analogously to joint torque, the "synergy torque" can be defined as

$$\tau_S = J^T \mathbf{u} \tag{8}$$

where J is the Jacobian. Making a simplifying assumption of first-order dynamics in the virtual space, $\tilde{B} = J^T$, $\tilde{A}$ is the identity matrix, and τ_S can be estimated from the change in synergy position: $\delta\mathbf{x} = J^T\mathbf{u}$. The system is then virtual-space-controllable if the non-negative rank of $\tilde{B} > s$.

A simple optimal controller can be achieved using the efficient non-negative least squares algorithm [18] at every timestep to select synergy activations to achieve target synergy position displacements:

$$\hat{\mathbf{u}} = \underset{\mathbf{u}}{\operatorname{argmin}} \|J^T\mathbf{u} - \widehat{\delta\mathbf{x}}\| \text{ subject to } \hat{\mathbf{u}} \geq 0 \tag{9}$$

This algorithm minimizes the norm of the synergy activation and produces smooth control.

F. Adaptive Control

Adaptive control is a well-known framework for system identification and control of dynamical systems with parametric uncertainty [19]. In this study, the parameter under identification is the Jacobian, as the observation and activation synergies are fixed. The Jacobian can be estimated offline by linear regression from a set of identification data:

$$J_0 = (UU^T)^{-1}UX^T, \tag{10}$$

where U is the $w \times N$ matrix of input synergy activations and X is the $s \times N$ matrix of measured synergy torques over N timesteps. If the estimated Jacobian has full non-negative rank this indicates that the virtual-space system should be controllable.

A model-adaptive controller can then revise the estimated Jacobian as a trajectory is tracked through synergistic control. The revision is done using with a batched, damped recursive least squares [20] formulation of the form:

$$J_{k+1} = J_k - c((UU^T)^{-1}U(J_k^T U - X)^T) \quad (11)$$

where U and X are recorded over time horizon N and the damping coefficient c controls adaptation speed. The time horizon can be lengthened to combat noise and avoid non-singular matrix inversion (this condition is caught and J is not altered).

III. EXPERIMENTS

A. Hardware

We demonstrate the usefulness of the synergistic control framework on the ACT robotic hand's thumb, index, and middle fingers. Each finger mimics the human's degrees of freedom in its joints, number of muscle/tendon actuator units, and complex tendon routing. Motor torques are controlled at 200Hz via CAN-bus communicating to 20 Barrett Pucks [21] that run the commutation for the brushless motors. Tendon lengths are measured at 200Hz by optical encoders with a 0.68 μm resolution. To show the simplicity of synergistic control, no other sensors were used within the control loops for these experiments. Experiment 1 is an example usage of the framework to perform tracking of thumb posture in 2 DOF with a recursively estimated Jacobian. Experiment 2 shows feedback control of coordinated three finger movement with an additional grip force target, which together produce handwriting on a surface.

B. Experiment 1: Model Adaptive Control

Fast and accurate movement of the thumb is easily achieved with the synergistic control framework by implementing a model-adaptive feedforward controller. An un-moderated process of system identification quickly builds the observation synergy S, activation synergy W, and an initial estimate of the Jacobian J. The thumb is placed in the neutral position where it is far from any of its joint limits. Each of the 8 tendon/motor actuators are commanded a 5-ms impulse at full torque and all of the resulting tendon displacements are logged. After each pulse, the thumb is controlled back to its initial posture. This process is repeated 10 times for each tendon to observe the average movement in tendon space caused by pulling a single tendon. For instance, pulling one of the flexors on the thumb may result in a negative displacement (contraction) of multiple flexors and a positive displacement (extension) of multiple extensors.

Such regularities in the output are captured by Principal Component Analysis (PCA). The matrix M of all tendon displacements caused by the motor pulses is run through PCA and the first two ($s = 2$) components are extracted


Fig. 2. Observation synergies found through self-exploration of thumb movement. Bars depict weighting of each tendon's displacement for the two components of synergy position, the rows of observation synergy matrix S. Muscle names are abbreviated as in [22].


Fig. 3. Activation synergies found through self-exploration of thumb movement. Bars depict (non-negative) weighting of each tendon's force for the three components of synergy activation, the columns of activation synergy W.


Fig. 4. Estimated Jacobian shows the movement, in the observation synergy frame represented by $S1$ and $S2$, that is expected when each activation synergy component ($W1$, $W2$ and $W3$) is applied.

t=0 sec   t=300 sec   t=400 sec   t=500 sec

Fig. 5. Observed open-loop tracking of thumb posture in virtual space as the initial model ($t = 0$) adapts over time. Circles were traversed in 1 second and cover approximately 4cm^2 workspace at the fingertip.


Fig. 6. Mean squared error (MSE) of actual versus target positions for circles in virtual space, normalized. The open-loop synergistic controller began with a very poor initial model and converged to a plateau at about 5% error within 500 seconds. Also shown is baseline MSE for the tendon-space PID controller (see Section III-C) and the 1σ level. The open-loop synergistic controller adapts to outperform the closed-loop tendon-space controller and reaches the estimated lower bound of open-loop tracking error.

as the strongest modes of thumb movement (see Figure 2). These components are set as the rows of the observation synergy S. PCA satisfies the requirement of orthonormality.

Similarly, regularities in the activation are found from the same identification data set. However, PCA is not appropriate in this case as the activations are non-negative. Instead, the well known non-negative matrix factorization (NMF) algorithm of alternating least squares [23] is used to extract which tendons contract together. NMF is used to find three ($w = 3$) components that well explain the variance of the matrix $N = max(-M, 0)$ of coupled tendon contractions. We chose w to be the lowest possible while maintaining virtual-space controllability as defined in Section II-E; w must be at least $s + 1$ and should result in $rank(J) = s$. As shown in Figure 3, thumb muscle groups are intuitively separated.

Finally, the initial estimate of the Jacobian J_0 is made through the linear regression described in 10 and is full non-negative rank (see Figure 4). Although the initial estimate is poor, as shown by the initial ($t = 0$) feedforward tracking results in Figure 5, a model-adaptive controller as described in Section II-F revises the estimated Jacobian online.

Using the identified S, W, and J_0, a fast 1Hz circle in the 2-dimensional synergy position space was tracked by a feedforward controller as the model was adapted in real-time. The desired trajectory covers an approximately 4cm^2 workspace at the fingertip. Using adaptation parameters $N = 200, c = .01$, the model was changed by more than an order of magnitude to more accurately track the goal. Within ten minutes, the performance plateaued to about 5% tracking error as compared to nearly 100% using the initial estimate of the Jacobian (see Figure 6). This result shows that dynamic movements are possible with the synergistic controller and demonstrates the utility of a reduced number of parameters for fast model identification and adaptation.

Optimal open-loop tracking performance is limited by the noise in the system. To verify that the adaptive controller converges to the level of minimum error, the system noise was identified by repeatedly applying a fixed sequence of motor torques. The final estimate of J at $t = 600sec$ was used to track the circle 100 times; the standard deviation of the normalized tracking error was calculated to be $\sigma = 0.041$. As shown in Figure 6, the adaptive controller converged to this 1σ level.

C. Experiment 2: Hybrid Force/Position Control

Whereas thumb posture tracking was achieved using synergies built from self-experience of the system dynamics, application of the synergistic control framework to specific tasks may require a human demonstration to provide the initial identification data. For instance, a more complex task such as handwriting involves coordination of multiple fingers and is not well described by kinematic trajectories alone. In order to control the motion of the pen endtip on a piece a paper, the hand must apply forces not only to induce movement but also to maintain grip on the pen. Observation and activation synergies were built from separate demonstrations of each 1) pen endtip horizontal displacement (East-West), 2) pen endtip vertical displacement (North-South), and 3) squeezing on the pen to maintain grip. These movements utilize all 20 motors of the ACT hand's thumb, index, and middle finger.

Again, PCA was the primary tool for extracting the observation synergy S from the kinematic demonstrations (1 and 2). As PCA maximizes the variance captured by each successive component, simultaneously reducing both of the datasets is likely to cause an undesired mixing of the demonstrated modes of motion. For instance, the first component may capture Northeast-Southwest endtip motion and the second Northwest-Southeast. To avoid this mixing and enforce the required orthogonality of components, PCA was used to sequentially extract the strongest component from each demonstration and Hotelling's deflation[24] was used between the datasets. This process is amenable to selecting any number of components from many demonstrations and also gives a priority to earlier demonstrations.

An alternative NMF method was used to find the activation synergy W. S is separated into its negative and positive parts and then rectified to give $w = 2s$ non-negative activation components. That is, $W = \begin{bmatrix} max(S, 0) \\ -1 * min(S, 0) \end{bmatrix}$. Intuitively,

this method weights tendon tension during a movement in the virtual space by the amount that tendon contracts. This method was used because the resulting Jacobian was found to be better-conditioned than the corresponding factorization found by alternating least squares.

The Jacobian J is found by repeating the identification process from Section III-B for all 20 motors with the hand in a pen-holding posture but without the pen in place. J is again estimated by the linear regression from Equation 10 and has full non-negative rank.

A prior demonstration of synergistic control for handwriting used PID control on tendon lengths to meet targets specified in the virtual space [12]. As mentioned in Section II-A, this system may not be controllable and showed limited dynamic response accordingly. In this experiment we show that incorporating non-negative activation synergies, the Jacobian, and optimal controller allows much better dynamic tracking of trajectories in the virtual space.

Controller 1 uses a PID controller acting on tendon lengths to follow trajectories in the virtual space. Control torques are clamped to be non-negative and integrators were saturated at 75% of max torque. Equation 3 is used to find the tendon length setpoints for each synergy position setpoint. Selecting feedback gains is difficult: 20 tendons $\times$ 3 gains (P, I, and D) = 60 variables to tune, tendon coupling means that tuning cannot be done on each tendon independently.

Even careful selection of all the gains provides little intuition into task-relevant controller performance. Essentially, there exist tendons for which good tracking performance may not be critical to the overall task. On the hand tracking performance plays very important role in tendons with big excursion changes. Given these limitations, we chose identical PID gains for all tendons and used a Kalman filter on each tendon to smooth noisy velocity measurements and give the best achievable tracking in the virtual space. To track a particular motion in virtual space, it is more intuitive to work in the virtual space, tuning gains that reflect the movement of groups of tendons in synchrony; the synergistic controller gives access to the system at this level.

Controller 2 uses a PID controller in the virtual space that acts directly on synergy position, the same space in which trajectories are defined. The gains were tuned to enable dynamic movements while maintaining good disturbance rejection and stability. To allow increased stiffness through the proportional term, the derivative term acts on an estimated synergy velocity calculated by a Kalman filter that smooths out noisy measurements. The Kalman filter considers each component of the synergy velocity to be independent and does not operate on the tendon lengths, but on the synergy velocity measurement directly. Future extensions of the filter to include automatic parameterization of covariance matrices, outlier rejection, and coupling of each component may allow higher PID gains and improved trajectory tracking. Figure 8 shows the much improved performance of the virtual-space controller versus the tendon length controller in tracking a virtual-space trajectory.

In addition to improved tracking, the synergistic controller

Fig. 7. A Kalman filter acting on measurements of virtual-space velocity produces a smoothed estimate that allows stiffer PID gains. The estimate shows minimal lag while attenuating the high frequency noise.


1/6 Hz   1 Hz

Fig. 8. Performance of virtual-space synergistic controller (black) versus tendon-space PID controller (gray) while tracking a handwriting trajectory in virtual space. Shown on the left is a slower 1/6 Hz circle, on the right a 1 Hz circle. In both cases the synergistic controller tracks much better, because of simpler gain tuning and the application of a grip force.


Fig. 9. Bode plot comparing synergistic vs. tendon-space PID controller tracking virtual-space circles at speeds ranging from 1/8 Hz to 1 Hz. Although both controllers are subject to the same hardware limitations, such as noise and latency, the synergistic controller is much more successful at following these dynamic handwriting movements.

is simultaneously maintaining a grip force to keep the pen from slipping out of the hand, stabilizing the pen while the virtual-space trajectory is tracked. With the addition of force-sensitive skin on the fingers or instrumentation on the pen, the grip force could further be be controlled through feedback. Notably, because the tendon-space PID controller does not account for the grip force, the controller is not appropriately designed to manipulate the pen; as shown in 8 there is a large offset resulting from the stiffness of the pen. Stiffer PID gains that would allow force development on the pen would also result in instability when moving out of the pen. Such considerations are easily handled by the synergistic controller.

Performance of the two controllers was compared over a dynamic range. For each test, a small circular motion at constant frequency was tracked and both magnitude and phase of the response were calculated using the FFT. As shown in Figure 9 both controllers show decreased magnitude and increased phase lag at higher speeds, as expected, but the synergistic virtual-space controller is substantially better.

IV. DISCUSSION

Control of the tendon-driven biomimetic ACT robotic hand is complicated by the system's high dimensionality, underactuation, redundancy, and nonlinearities. System identification, model-predictive control, and planning are greatly simplified by dimensionality reduction through the synergistic control framework. Applying a first-order virtual-space model quickly enabled adaptive controller performance surpassing the model-free full dimensionality controller. The complex task of dynamic handwriting is also more easily described through the use of task-relevant observation and activation synergies. Handwriting movements in the virtual space are well tracked while simultaneously maintaining a grip force on the pen, providing a foundation for future exploration of closed-loop control of the pen position on the tablet.

In this work we have demonstrated the challenges when PID control in tendon space is considered. In a system with many tendons, gain tuning is not intuitive and is task dependent; each tendon has different tracking requirements. With the synergistic control framework we can reduce the high-dimensional control space into a low-dimensional space where gain tuning, identification, and trajectory definition are more intuitive.

Further work within the structure of the synergistic control framework will include refinement of the system model, the optimal controller, and the Kalman filter for state estimation. We are planning to improve high-speed tracking performance by incorporating gain scheduling methods and using filtering techniques that automatically reject outliers.

Future research will also consider the use of the proposed framework as the underlying controller for learning new tasks based on a reward feedback specified in task-relevant space. We are aiming to take advantage of policy gradient methods and the iterative path integral control framework [25]. Our goal is to investigate whether or not reinforcement learning with synergistic control has better performance and is faster than learning in tendon space.

REFERENCES

[1] M. Tresch and A. Jarc, "The case for and against muscle synergies," *Current Opinion in Neurobiology*, vol. 19, no. 6, pp. 601–607, 2009.

[2] M. Santello, M. Flanders, and J. Soechting, "Postural hand synergies for tool use," *Journal of Neuroscience*, vol. 18, no. 23, 1998.

[3] E. Weiss and M. Flanders, "Muscular and postural synergies of the human hand," *Journal of Neurophysiology*, vol. 92, no. 1, p. 523, 2004.

[4] L. Ting and J. Macpherson, "A limited set of muscle synergies for force control during a postural task," *Journal of Neurophysiology*, vol. 93, no. 1, p. 609, 2005.

[5] E. Todorov, W. Li, and X. Pan, "From task parameters to motor synergies: A hierarchical framework for approximately optimal control of redundant manipulators," *Journal of Robotic Systems*, vol. 22, no. 11, pp. 691–710, 2005.

[6] V. Krishnamoorthy, M. Latash, J. Scholz, and V. Zatsiorsky, "Muscle synergies during shifts of the center of pressure by standing persons," *Experimental Brain Research*, vol. 152, no. 3, pp. 281–292, 2003.

[7] W. Kargo and S. Giszter, "Individual premotor drive pulses, not time-varying synergies, are the units of adjustment for limb trajectories constructed in spinal cord," *Journal of Neuroscience*, vol. 28, no. 10, p. 2409, 2008.

[8] M. Tresch, V. Cheung, and A. d'Avilla, "Matrix factorization algorithms for the identification of muscle synergies: evaluation on simulated and experimental data sets," *Journal of Neurophysiology*, vol. 95, no. 4, pp. 2199–2212, 2006.

[9] A. Deshpande, Z. Xu, M. J. V. Weghe, L. Y. Chang, B. H. Brown, D. D. Wilkinson, S. M. Bidic, and Y. Matsuoka, "Mechanisms of anatomically correct testbed (ACT) hand," *IEEE Trans. Mechatronics*, to appear, 2012.

[10] P. Tuffield and H. Elias, "The shadow robot mimics human actions," *Industrial Robot: An International Journal*, vol. 30, no. 1, pp. 56–60, 2003.

[11] A. Shkolnik and R. Tedrake, "High-dimensional underactued motion planning via task space control," in *Int. Conf. on Intelligent Robots and Systems*, 2008.

[12] E. Rombokas, M. Malhotra, and Y. Matsuoka, "Task-specific demonstration and practiced synergies for writing with the ACT hand," in *IEEE International Conference on Robotics & Automation*, 2011.

[13] A. Zhang, M. Malhotra, and Y. Matsuoka, "Musical piano performance by the ACT hand," 2011.

[14] S. Jacobsen, H. Ko, E. Iversen, and C. Davis, "Antagonistic control of a tendon driven manipulator," in *IEEE Int. Conf. Robotics and Automation*, 1998, pp. 1334–1339.

[15] H. Kobayashi, K. Hyodo, and D. Ogane, "On tendon-driven robotic mechanisms with redundant tendons," *Int. J. Robotic Research*, vol. 17, no. 5, pp. 561–571, 1998.

[16] P. Burgess, J. Wei, F. Clark, and J. Simon, "Signaling of kinesthetic information by peripheral sensory receptors," *Annual Review of Neuroscience*, vol. 5, no. 1, pp. 171–188, 1982.

[17] D. Burke, S. Gandevia, and G. Macefield, "Responses to passive movement of receptors in joint, skin and muscle of the human hand." *The Journal of physiology*, vol. 402, no. 1, p. 347, 1988.

[18] C. Lawson and R. Hanson, *Solving least squares problems*. Society for Industrial Mathematics, 1995, vol. 15, ch. 23.

[19] P. Ioannou and J. Sun, *Robust adaptive control*. Prentice-Hall, 1996.

[20] L. Ljung, *System Identification: Theory for the User, 2nd ed.* Prentice-Hall, 1999, ch. 11.4.

[21] W. Townsend, A. Crowell, G. Pratt, and T. Hauptman, "Ultra-compact, high-performance motor controller and method of using same," Feb. 22 2011, uS Patent 7,893,644.

[22] H. Schmidt, *Surgical Anatomy of the Hand*. New York: Thieme, 2003, ch. 5.

[23] M. Berry, M. Browne, A. Langville, V. Pauca, and R. Plemmons, "Algorithms and applications for approximate nonnegative matrix factorization," *Computational Statistics & Data Analysis*, vol. 52, no. 1, pp. 155–173, 2007.

[24] L. Mackey, "Deflation methods for sparse pca," in *Advances in Neural Information Processing Systems*, vol. 21, 2009, pp. 1017–1024.

[25] E. Theodorou, J. Buchli, and S. Schaal, "A generalized path integral control approach to reinforcement learning," *Journal of Machine Learning Research*, no. 11, pp. 3137–3181, 2010.