

Antagonistic Muscle based Robot Control for Physical Interactions

Tapomayukh Bhattacharjee¹ and Günter Niemeyer²

Abstract—Robots are ever more present in human environments and effective physical human-robot interactions are essential to many applications. But to a person, these interactions rarely feel biological or equivalent to a human-human interactions. It is our goal to make robots feel more human-like, in the hopes of allowing more natural human-robot interactions. In this paper, we examine a novel biologically-inspired control method, emulating antagonistic muscle pairs based on a nonlinear Hill model. The controller captures the muscle properties and dynamics and is driven solely by muscle activation levels. A human-robot experiment compares this approach to PD and PID controllers with equivalent impedances as well as to direct human-human interactions. The results show the promise of driving motors like muscles and allowing users to experience robots much like humans.

I. INTRODUCTION

Robots are increasingly being used in human environments for applications ranging from entertainment to assistance. Physical human-robot interactions become part of the normal tasks and often have to replace traditional human-human interactions [1], [2]. In addition, cognitive interactions depend on and are influenced by physical interactions. People are more likely to accept and work well with robot partners if they appear human-like [3]. We therefore strive to make physical human-robot interactions truly feel akin to human-human interactions.

We believe this is best achieved by a control scheme that directly emulates muscles. Biological muscles exhibit numerous properties from auto-stiffening to limited activation speeds that combine into significant nonlinear dynamics. We utilize a detailed model of antagonistic muscles presented by Winters and Stark [4], in turn based on a nonlinear Hill model. Both flexor and extensor muscle contain nonlinear series elastic, active, and viscous elements. The resultant controller is commanded only by the two muscle activation levels.

To explore a user’s physical experience of this controller, we conduct a human-subjects experiment comparing against traditional PD and PID controllers with equivalent impedances. We also compare against direct manual interactions. To focus on the basic haptic interaction, we target a simplified stationary task. Depicted in Fig. 1, the performer (controlled robot or human) is commanded to hold their position while the test’s subject is allowed to perturb and interact with the performer. We see this as a precursor to handshaking, without the need of implicit communication or

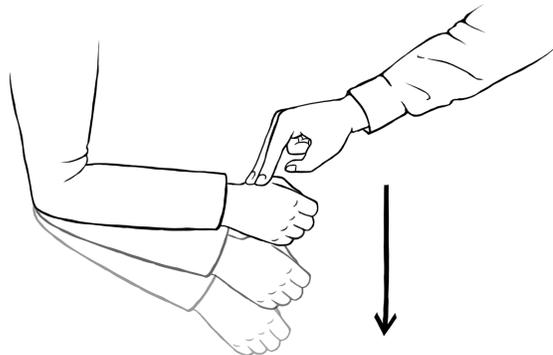


Fig. 1: The performer (robot or human) holds their arm out with elbow bent and eyes closed. They are charged to maintain this posture regardless of external disturbances. The tester injects stimuli and observes the response haptically and visually.

negotiation that occur during the handshake. The Turing-like test seeks to answer whether the experience appears human-like. Results expose limitations, but also confirm that muscle-based control has the potential to convincingly convey a human-like feel.

II. RELATED WORK

A. Human Arm Impedance and Muscle Models

There have been many studies focusing on identifying human arm impedance characteristics such as stiffness [5], [6], damping [7], and mass [8]. There have also been many studies on the modeling of human arm muscles [9], [10]. Winters and Stark [4] identified a non-linear eighth order agonist-antagonist muscle model using engineering analysis and obtained parameter values for human arm muscle constitutive equations.

B. Impedance Control and Human-Robot Interactions

Robotics has long known that antagonistic muscles set the interaction impedance in humans [11]. And it has embraced Stiffness Control [12] and Impedance Control [13], which specify and maintain impedance over precise position or force tracking. Most recently, actuators are even being developed that can directly and physically vary their impedance [14].

Taking inspiration from biology, impedance variations have been learned [15]. Ganesh *et. al.* [16] developed a human-like automatic motor behavior which adapts torque

¹Tapomayukh Bhattacharjee is with Healthcare Robotics Lab, Georgia Institute of Technology, Atlanta, USA. He did this work at Disney Research, Los Angeles, as an R&D Lab Associate tapomayukh@gatech.edu

²Günter Niemeyer is with Disney Research, Los-Angeles, USA gunter.niemeyer@disneyresearch.com

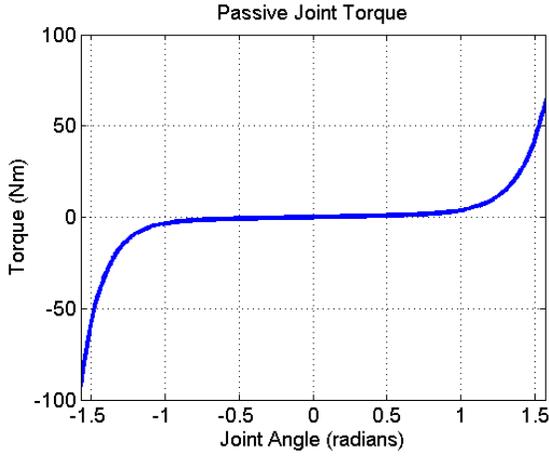


Fig. 2: *Passive Joint Torque*.

in response to slow disturbances but increases stiffness for high-frequency disturbances.

Inspiration has also been found in handshaking. Giannopoulos *et al.* [17] compared real and virtual handshakes. They found human handshakes were consistently rated above robot handshakes, even with their more advanced controller. Avraham *et al.* [18] performed Turing-like tests to compare three virtual handshake models.

III. ANTAGONISTIC MUSCLE BASED CONTROL

We control the robot by emulating a pair of antagonistic muscles to drive each joint. The robot, acting akin to the human skeleton, is driven by the muscle torque. The activation levels are set, possibly with some delay, by a highlevel behavior.

We follow the muscle model developed by Winters and Starks [4] and illustrated in Fig. 3. Each joint is driven by an extensor and flexor muscle as well as passive joint mechanics. We define the joint angle q as positive when flexed, and negative when extended. The muscle elongations ϕ are positive when lengthened, negative when contracted, and measured relative to their nominal length.

$$\phi^{\text{flex}} = -q, \phi^{\text{ext}} = +q \quad (1)$$

We assume each muscle independently generates a positive contracting torque $T(\phi) > 0$, so the total joint torque τ is

$$\tau = T^{\text{flex}}(\phi^{\text{flex}}) - T^{\text{ext}}(\phi^{\text{ext}}) - \tau_p \quad (2)$$

The passive joint torque τ_p stems from a nonlinear exponential stiffening spring [4], computed as

$$\tau_p = k_p q + \begin{cases} M_{pflex} \left(e^{q/\theta_{pflex}} - 1 \right) & \text{if } q \geq 0 \\ -M_{pext} \left(e^{q/\theta_{pext}} - 1 \right) & \text{if } q < 0 \end{cases} \quad (3)$$

where k_p is the spring constant, M_p and θ_p are torque and angle parameters. We obtain parameter values summarized

TABLE I: Muscle Model Parameters for a Typical Human Elbow obtained from [4].

Muscle Component	Parameter	Flexor	Extensor
Passive Joint Torque	k_p	1.5 Nm/rad	
	M_{pflex}	0.0074 Nm	
	M_{pext}	0.0023 Nm	
	θ_{pflex}	0.1745 rad	
	θ_{pext}	0.1484 rad	
Series Element	M_s	3.6486 Nm	4.5359 Nm
	θ_s	0.1122 rad	0.2327 rad
	M_{max}	60 Nm	50 Nm
Active Element	θ_o	0 rad	-0.6981 rad
	θ_{sh}	1.7 rad	1.6 rad
	c_s	-0.01 1/rad	0.2 1/rad
Viscous Element	v_{max}	22 rad/s	28 rad/s
	c_v	0.34	0.4

in Table I, for a typical human elbow from [4] and show the passive torque in Fig. 2.

IV. NONLINEAR MUSCLE IMPEDANCE

Inspired by Hill [9], each antagonistic muscle combines an active and viscous element in parallel, connected to the joint via a series elastic element. The total elongation can thus be separated into the deflection of the active/viscous elements ϕ_a and the series element ϕ_s

$$\phi = \phi_a + \phi_s \quad (4)$$

The torques must balance such that

$$T(\phi) = T_s(\phi_s) = T_a(\phi_a) + T_v(\dot{\phi}_a) \quad (5)$$

A. Series Element

The series element consists of a nonlinear spring, modulating how quickly and with what impedance the active/viscous torques are transmitted to joint. The series element torque T_s is

$$T_s = \begin{cases} M_s \left(e^{\frac{\phi_s}{\theta_s}} - 1 \right) & \text{if } \phi_s \geq 0 \\ 0 & \text{if } \phi_s < 0 \end{cases} \quad (6)$$

where M_s and θ_s are torque and deflection constants. Fig. 4 shows the torque for flexor and extensor muscles.

B. Active Element

The active element produces torque proportional to a normalized activation level A ($0 \leq A \leq 1$). We assume that the activation level of each muscle is independent. The active torque

$$T_a = AM_{\max} S(\phi_a) \quad (7)$$

is further shaped by the elongation according to

$$S(\phi_a) = e^{-((\phi_a - \theta_o)/\theta_{sh})^2} + c_s \phi_a \quad (8)$$

where θ_o , θ_{sh} and c_s are angle and slope parameters. As shown in Fig. 5, full torque can only be produced in the center of the muscle's range. Near its limits, the muscle becomes weaker.

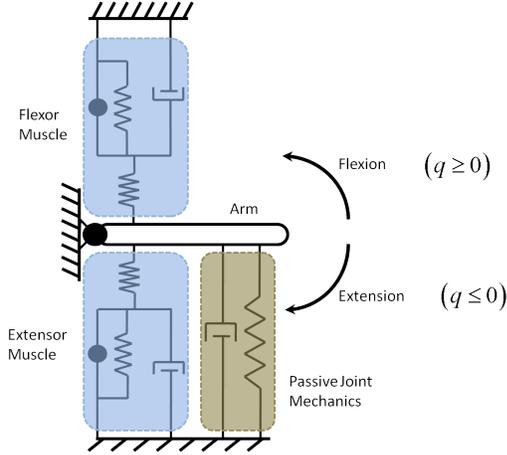


Fig. 3: An antagonistic muscle model where each muscle is modeled using non-linear high-order model.

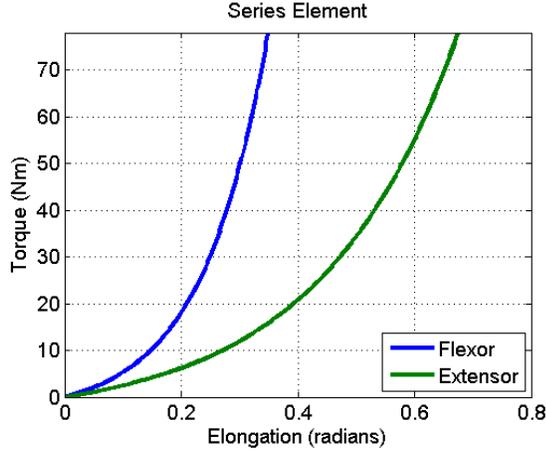


Fig. 4: Series Element Torque.

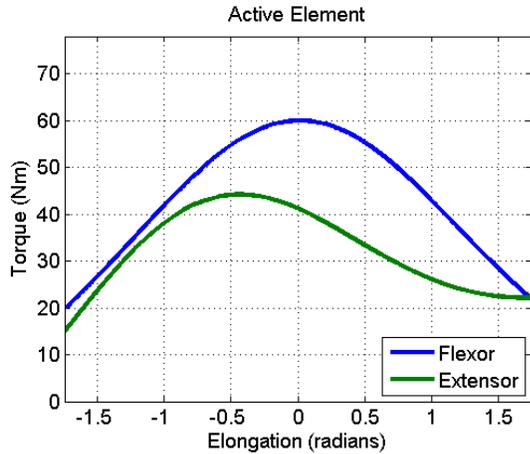


Fig. 5: Active Element Torque at full activation.

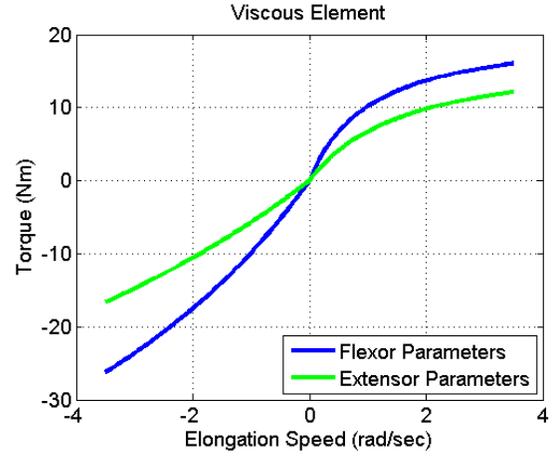


Fig. 6: Viscous Component Torque. The flexor and extensor activation levels are fixed at 1.0 and $v_m = v_{max}$.

C. Viscous Element

The viscous element produces the torque

$$T_v = T_a \begin{cases} (-1) & \text{if } \dot{\phi}_a = -v_m \\ \left(-\frac{|\dot{\phi}_a| + c_v |\dot{\phi}_a|}{|\dot{\phi}_a| + c_v v_m} \right) & \text{if } -v_m \leq \dot{\phi}_a \leq 0 \\ \eta \left(\frac{|\dot{\phi}_a| + 0.5c_v |\dot{\phi}_a|}{|\dot{\phi}_a| + 0.5c_v \eta v_m} \right) & \text{if } 0 \leq \dot{\phi}_a \leq \eta v_m \\ \eta & \text{if } \eta v_m \leq \dot{\phi}_a \end{cases} \quad (9)$$

where,

$$v_m = v_{max} (0.5 + 0.5 A S(\phi_a)) S(\phi_a) \quad (10)$$

specifies a maximum contracting velocity. We note T_v is proportional to and effectively modulates T_a . From (5) we see $-T_a \leq T_v \leq \eta T_a$ implies $0 \leq T \leq (1 + \eta) T_a$. Typically, $\eta = 0.3$ [4], so the total muscle torque is non-negative and limited to 130% of the active torque.

The four cases alter the behavior according to the elongation speed. First, if the muscle is externally forced to contract beyond its maximum speed, the series element completely collapses to $T_s = 0$, the viscous element cancels all active torque, and the active element contracts with $\dot{\phi}_a = -v_m$. Only when the active element catches up, will any output torque be restored. Second, in normal contraction, the viscous element reduces the total torque $T < T_a$. Third, in normal lengthening, the viscous element resists and increases the output torque $T > T_a$. Fourth, if the muscle is externally lengthened faster than ηv_m , the active element lengthens equally to maintain the series element deflection and torque at $T = (1 + \eta) T_a$. Figure 6 shows the viscous torque at normal elongation speeds for $A = 1$ and $S(\phi_a) = 1$. Please note that because the viscous torque is proportional to the active torque, the effective damping will drop to zero with no activation.

D. Implementation

Combining a nonlinear stiffness, a shaped active element, and a nonlinear viscosity, each muscle presents its own internal dynamics. Computationally, we use the active element's elongation ϕ_a as the muscle's state. Using the torque balance (5), series and active elements (6), (7), we can invert the viscous element (9) to determine the state derivative $\dot{\phi}_a$, mindful of the current case. However, we note that the damping gain may drop to zero making the dynamics infinitely fast. As such, a discretized implementation needs to compute the steady state solution of ϕ_a to the current inputs, eg. using the Newton-Raphson method. Knowing the steady state solution, we can ensure that the discretized state does not overshoot and the implementation remains stable, reaching steady state in a single time step if appropriate.

V. EXPERIMENTS

A. Activation Functions

Emulating muscle pairs on a zero-friction skeleton, the system now has appropriate 'intrinsic' dynamics. It is stable and well damped with steady state location and torque levels as well as stiffness determined by the extensor and flexor activation levels. Of course, humans constantly vary their muscle activations. As such, we also seek to vary activation levels to make our system appear more life-like.

Based on observed human behaviors from preliminary tests, we believe humans tends to build up force when held away from a goal position, similar to integral action. Also, humans show stiffening and relaxing actions.

To create an equivalent behavior, we dynamically vary the activation levels within the allowable region shown in Fig. 7. Away from boundaries, we modify the activations differentially based on a measured error. This effectively raises or lowers the joint torque in an integral fashion and is depicted as direction '1'. The error is a linear combination of position, velocity, and force errors.

At a boundary, we raise the activation levels together to increase torque and stiffness. The changes are scaled to produce the same differential action as before and shown as direction '2'. When the error changes sign and the torque wants to be updated in the reverse direction, the differential action pulls the activations into the allowable region, marked '3'. And finally, when no external torque is observed, the activations drop with a time constant of 4s to relax the system.

B. Experimental Design

Our experiment was designed similar to a Turing test. Human subjects were asked to interact with the robot which was trying to maintain a fixed position. Unbeknownst to the subjects, the robot was controlled either by a muscle-based controller with an automatically varying activation, or by a traditional PD or PID controller, or manually. In the last case, a human operator would disconnect the motor and manually hold the robot arm. As such, the experiment was designed to evaluate how muscle-based control compares to traditional controllers and truly biological regulation. To

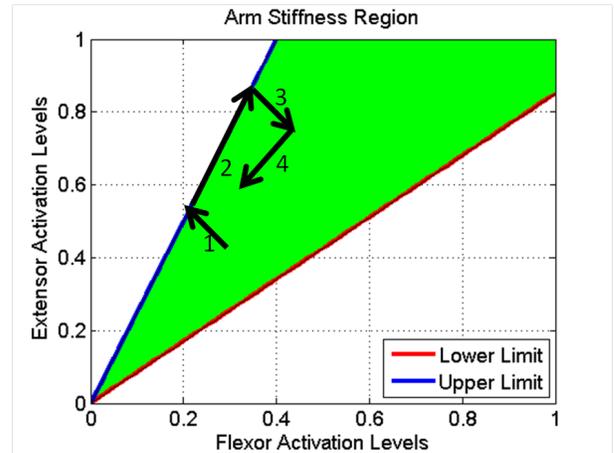


Fig. 7: The flexor and extensor activation signals are limited to the green region. The arrows indicate directions of change for (1) torque increase, (2) torque and stiffness increase, (3) torque decrease, and (4) relaxing.

minimize variation, the robot was always held in a bent-elbow configuration and the task involved only the elbow joint control. For the PD and PID controllers, the position gain was 30 Nm/rad, derivative gain was 3 Nm-s/rad, and the integral gain, if present, was 200 Nm/rad-s.

Our experimental setup is shown in Fig. 8. We used one of the 3 degrees of freedom and locked the wrist and shoulder joint to obtain the posture as shown in the middle of Fig. 8. We attached a stick at the end-effector of the robot-arm to provide an interface with which a human can interact. The joints of the robot-arm are powered by Maxon motors with 100:1 gear-ratio, capable of producing 19.8Nm of torque and sufficient to achieve most required muscle torques. Elmo servo drives provide current control and measure the joint position. They are controlled via a CAN bus at 500 Hz with the control code executing in MATLAB Realtime Workshop. The robot is also equipped with a 6-axis ATI Mini-45 F/T sensor at the end-effector to measure environment forces.

For experiments, the robot was placed behind a curtain. It was setup such that the human subject cannot visually ascertain whether the robot is automatically controlled or whether its manually controlled. Only the stick attached to the end effector would be visible to human subjects. The human subjects would then push on the stick with one finger. For some interactions, a human will be manually controlling the robot-arm whereas for others it will be controlled using the muscle-based behavior / PD / PID behavior.

The experiment presented the four behaviors to each subject three times. The order of the twelve trials was randomized. Also, since the motor decoupling / coupling procedure took around 30 seconds time, a uniform break of 30 seconds was given in between all the trials so that the subject could not guess the behaviors just based on the time-gap between the trials.

Subjects were instructed to compare a particular interaction behavior to a human-held stick using a seven point Likert

TABLE II: Programmed Behaviors and Quantitative Behavior Trends.

Behavior Number	Controlled by	Behavior Type	Average Likert Score	Minimum Likert Score	Maximum Likert Score
1	Robot	Muscle-based Behavior	4.17	1	7
2		PID Control	3.57	1	7
3		PD Control	2.57	1	6
4	Human	Human Behavior	4.87	1	7

TABLE III: Average Likert Scores.

Behavior Number	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6	Subject 7	Subject 8	Subject 9	Subject 10
1	4.67	3	5	4	3.33	3.33	5.67	3.67	3.33	5.67
2	4.33	4.33	3.33	5.67	3.33	1.33	4	3.67	2.33	3.33
3	5	1.33	4.33	1.33	1.33	1.0	2.67	3.33	3	2.33
4	5.67	7	4.67	6	4.33	3.67	1.33	3	6.67	6.33

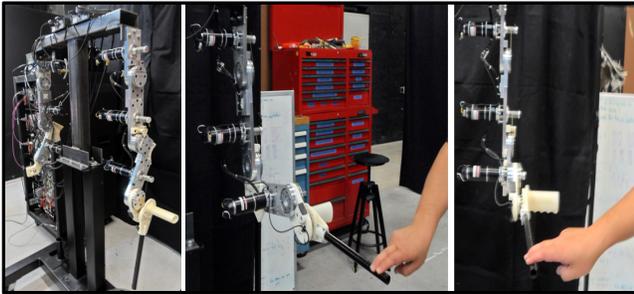


Fig. 8: Experimental setup showing the robot-arm and a human-subject physically interacting with it.

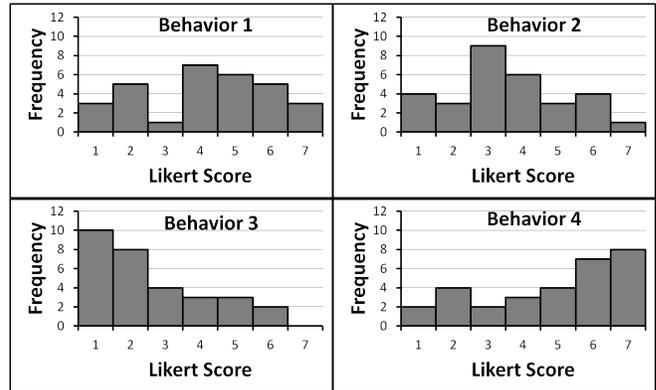


Fig. 9: Histograms of responses to all the 4 behaviors.

scale. As reference, before any trial, subjects were shown the PD and the manual behavior. They were granted sufficient time to familiarize themselves with the interactions.

For each trial, the subjects wore headphones and white noise was played to cancel any auditory cues from the motors. Also, for each trial, they were instructed to interact with the stick using a finger as shown in Fig. 8 for around 10 seconds in anyway they wanted. They could push or pull it in any direction or hold it in any position.

C. Data Collection

We collected experimental data for 10 human subjects. First, the subjects answered a 'Pre-Task Background Questionnaire'. This questionnaire collected general data about the subjects' age, gender, their background / experience in robotics and haptics, engineering / programming experience, and identified their dominant hand. After each trial subjects were asked *How close was this interaction compared to a human-held stick?*. The subjects were asked to answer using a 7-point Likert scale where '1' was *Robot-Held* and '7' was *Human-Held*. At the end of the last trial, we asked them some open-ended questions to understand their methodology to judge the performance of the system. The questions were: 1. *What approach did you use to test the interaction?* 2. *What criteria did you use to distinguish a human-held stick from a robot-held stick?*, and 3. *Other Comments / Suggestions ?*.

VI. RESULTS AND DISCUSSION

We performed experiments with 10 human subjects. 5 of them were male, and 5 were female. The average age of the participants was approx. 30 years. 5 out of the 10 participants had interacted with a robot / haptic device previously, 4 of them were from non-engineering backgrounds, and 1 participant noted his/her left hand as the dominant hand.

Table III shows the results, averaging each subject's response over the three repetitions of each behavior. A summary is also included in Table II. We see that 3 of the 10 subjects found the muscle-based behavior more human-like than the manual behavior. In general, however, the likert scores were quite varied. The histograms collected in Fig. 9 show the distributions.

We notice that subjects consistently scored behavior 3, traditional PD control, lowest and very different from human behavior. Note that Behavior 3 does not include integral control. This supports our belief that integral action is important.

We also observe that the human behavior got very low scores once in a while. We are very pleased that on numerous occasions, muscle-based control received the highest score, supporting the value of muscle controllers.

To select appropriate significance tests, we first performed *Shapiro-Wilk test* to check the normality of the data. As we might suspect from Fig. 9, the data is not normally

distributed.

We used the *Mann-Whitney test* to compare the muscle-based behavior to all other behaviors. Considering a 95% confidence-interval, we found significant difference from behavior 3 (p - value = 0.0008). If we consider a 90% confidence-interval, we also find difference with behavior 2 (p - value = 0.086). However, interestingly, we could not find any significant difference between the muscle-based behavior and the human behavior. It would seem that using muscles, we can approach the illusion of interacting with a human-arm.

In addition to the quantitative results discussed above, it is interesting to note some of the open-ended responses of the human-subjects. Most participants used a combination of force and visual techniques to identify the behaviors. The force techniques included pushing / pulling with varied amounts of force repeatedly, impulse response, holding the stick with increasing strength, and watching for subtle robot related vibrations. The visual techniques included watching for overshoot response.

VII. CONCLUSIONS

In this paper, we utilized a novel bio-inspired control methodology suitable for physical interactions. To evaluate the system we programmed a basic activation behavior and created a haptic Turing-like test on a simple stationary task. We compared muscle-based control to tradition PD and PID control, as well as to manual interactions.

We are delighted that muscle-based control performed well and was at times able to create the illusion of a biological interaction. We feel the intrinsic dynamics and properties will be invaluable in achieving natural human-robot interactions. In future work, we hope to explore the sensitivity of individual parameters and aspects of the muscle model, potentially to simplify the necessary complexity. Furthermore, we hope to extend the illusion beyond reflexive to voluntary and rhythmic movements, in particular in combination with high-level strategies, so that we may achieve a truly convincing hand-shake. Ultimately, we hope this will allow robots to function better in human contact.

ACKNOWLEDGMENT

We gratefully acknowledge the support from Disney Research, Los Angeles for this project. We would like to thank our subjects for participating and the entire team of Disney Research, Los Angeles for providing us with valuable feedback. We extend our special thanks to Dr. Quinn Smithwick and Dr. Rick Cory for their help in conducting our study.

REFERENCES

- [1] R. Alami, A. Albu-Schaeffer, A. Bicchi, R. Bischoff, R. Chatila, A. De Luca, A. De Santis, G. Giralt, J. Guiochet, G. Hirzinger, *et al.*, "Safe and dependable physical human-robot interaction in anthropic domains: State of the art and challenges," in *Proc. IROS*, vol. 6, no. 1. Citeseer, 2006.
- [2] A. De Santis, B. Siciliano, A. De Luca, and A. Bicchi, "An atlas of physical human-robot interaction," *Mechanism and Machine Theory*, vol. 43, no. 3, pp. 253–270, 2008.

- [3] B. R. Duffy, "Anthropomorphism and the social robot," *Robotics and autonomous systems*, vol. 42, no. 3, pp. 177–190, 2003.
- [4] J. M. Winters and L. Stark, "Analysis of fundamental human movement patterns through the use of in-depth antagonistic muscle models," *Biomedical Engineering, IEEE Transactions on*, no. 10, pp. 826–839, 1985.
- [5] F. A. Mussa-Ivaldi, N. Hogan, and E. Bizzi, "Neural, mechanical, and geometric factors subserving arm posture in humans," *The Journal of Neuroscience*, vol. 5, no. 10, pp. 2732–2743, 1985.
- [6] R. Shadmehr, "Control of equilibrium position and stiffness through postural modules," *Journal of motor behavior*, vol. 25, no. 3, pp. 228–241, 1993.
- [7] T. Tsuji, Y. Takeda, and Y. Tanaka, "Analysis of mechanical impedance in human arm movements using a virtual tennis system," *Biological cybernetics*, vol. 91, no. 5, pp. 295–305, 2004.
- [8] J. M. Dolan, M. B. Friedman, and M. L. Nagurka, "Dynamic and loaded impedance components in the maintenance of human arm posture," *Systems, Man and Cybernetics, IEEE Transactions on*, vol. 23, no. 3, pp. 698–709, 1993.
- [9] A. Hill, "The heat of shortening and the dynamic constants of muscle," *Proceedings of the Royal Society of London. Series B, Biological Sciences*, vol. 126, no. 843, pp. 136–195, 1938.
- [10] R. Shadmehr and S. P. Wise. (2005) A simple muscle model. [Online]. Available: <http://www.shadmehrlab.org/book/musclemodel.pdf>
- [11] N. Hogan, "Adaptive control of mechanical impedance by coactivation of antagonist muscles," *Automatic Control, IEEE Transactions on*, vol. 29, no. 8, pp. 681–690, 1984.
- [12] J. Salisbury, "Active stiffness control of a manipulator in cartesian coordinates," in *Decision and Control including the Symposium on Adaptive Processes, 1980 19th IEEE Conference on*, vol. 19, Dec 1980, pp. 95–100.
- [13] N. Hogan, "Impedance control: An approach to manipulation: Part i-theory," *Journal of dynamic systems, measurement, and control*, vol. 107, no. 1, pp. 1–7, 1985.
- [14] B. Vanderborght, A. Albu-Schäffer, A. Bicchi, E. Burdet, D. G. Caldwell, R. Carloni, M. Catalano, O. Eiberger, W. Friedl, G. Ganesh, *et al.*, "Variable impedance actuators: A review," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1601–1614, 2013.
- [15] J. Buchli, F. Stulp, E. Theodorou, and S. Schaal, "Learning variable impedance control," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 820–833, 2011.
- [16] G. Ganesh, A. Albu-Schaffer, M. Haruno, M. Kawato, and E. Burdet, "Biomimetic motor behavior for simultaneous adaptation of force, impedance and trajectory in interaction tasks," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*. IEEE, 2010, pp. 2705–2711.
- [17] E. Giannopoulos, Z. Wang, A. Peer, M. Buss, and M. Slater, "Comparison of people's responses to real and virtual handshakes within a virtual environment," *Brain research bulletin*, vol. 85, no. 5, pp. 276–282, 2011.
- [18] G. Avraham, I. Nisky, H. L. Fernandes, D. E. Acuna, K. P. Kording, G. E. Loeb, and A. Karniel, "Toward perceiving robots as humans: Three handshake models face the turing-like handshake test," *Haptics, IEEE Transactions on*, vol. 5, no. 3, pp. 196–207, 2012.