

Towards teleoperation with human-like dynamics

Human use of elastic tools

Aiple, M.; Schiele, A.

DOI

[10.1109/WHC.2017.7989896](https://doi.org/10.1109/WHC.2017.7989896)

Publication date

2017

Document Version

Accepted author manuscript

Published in

Proceedings of the 2017 IEEE World Haptics Conference - WHC 2017

Citation (APA)

Aiple, M., & Schiele, A. (2017). Towards teleoperation with human-like dynamics: Human use of elastic tools. In M. Harders, M. Otaduy, & E. Ruffaldi (Eds.), *Proceedings of the 2017 IEEE World Haptics Conference - WHC 2017* (pp. 171-176). IEEE. <https://doi.org/10.1109/WHC.2017.7989896>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Towards Teleoperation with Human-like Dynamics: Human Use of Elastic Tools

Manuel Aiple¹, *Student Member, IEEE*, and André Schiele¹

Abstract—Variable stiffness actuators undergo lower peak force in contacts compared to their rigid counterparts, and are thus safer for human-robot interaction. Furthermore, they can store energy in their elastic element and can release it later to achieve human-like dynamic movements. However, it is not clear how to integrate them in teleoperator systems so that they can be controlled intuitively by a human. We performed an experiment to study human use of elastic tools to determine how a teleoperator system with an elastic slave would need to be designed. For this, we had 13 untrained participants hammer with an elastic tool under different stiffness conditions, asking them to try to find the best timing for a backward-forward swing motion in order to achieve the strongest impact. We found that the participants generally executed the task efficiently after a few trials and they converged to very similar solutions. The stiffness influenced the performance slightly, a stiffness between 2.3 Nm/rad and 4.1 Nm/rad showing the best results. We conclude that humans intuitively know how to efficiently use elastic tools for hammering type tasks. This could facilitate the control of teleoperator systems with elastic slave manipulators for tasks requiring explosive movements like hammering.

I. INTRODUCTION

Teleoperation promises to combine the advantages of a human operator in terms of flexibility and planning with the benefits of a robot in terms of robustness to harsh environments (e.g., radioactively contaminated areas, subsea, space). Extensive research has been done to find methods and technologies to efficiently couple the motion of the operator on the master side with the motion of the robot on the slave side, addressing issues like stability [1], transparency, time-delay [2], sharing of control [3]. However, with the robots and input devices of today, the operator is mostly limited to slow motions with low dynamics.

Series Elastic Actuators (SEAs) and variable stiffness actuators (VSAs) have increasingly become a topic of interest over the last years. Vanderborght et al. give a good overview of the existing systems [4]. Unlike active compliance, which emulates a spring effect by force measurement and control, SEAs and VSAs are equipped with a physical spring on the output of the actuator. For SEAs with constant stiffness, this is a mechanical spring, whereas VSAs incorporate a mechanism to change their effective stiffness. This gives SEAs and VSAs true mechanical compliance and higher admittance at high frequencies than actuators with active compliance, resulting in reduced forces in highly dynamic situations. In particular, they are more robust to impacts and safer for humans to interact with.

¹M. Aiple and A. Schiele are affiliated with the Department of BioMechanical Engineering, Delft University of Technology, Delft, The Netherlands. {m.aiple, a.schiele}@tudelft.nl.

Furthermore, SEAs and VSAs can store energy in the spring element to release it again later. Previous studies focused on how to exploit this to make autonomous robots more energy efficient in high dynamic tasks like walking and throwing [5][6]. Our research aims at developing methods and technologies for teleoperation of dynamic real-life human tasks, like hammering, jolting, throwing, etc., using VSAs. But, for teleoperation with elastic tools, there are competing objectives between following the human input accurately to achieve maximum transparency and achieving the best task performance by following an optimal control strategy. In order to resolve this conflict, knowledge about human task performance using flexible tools is required. If humans have an intuitive understanding of how to optimally use elastic tools, teleoperator systems should be optimized for maximum transparency also for teleoperation with flexible tools. On the other hand, if human performance is bad with elastic systems, a shared control approach based on optimal control of the slave might be better.

This paper examines how humans interact with elastic tools for maximum output velocity tasks. Maximum output velocity tasks are tasks like throwing and hammering, aimed at giving the most kinetic energy to an object with as little input energy as possible. We picked this class of tasks for this experiment as it makes use of SEAs' advantage compared to rigid actuators, their capacity to store energy in their spring when switching from a backward motion to forward motion and then releasing it as kinetic energy. Optimal control strategies have been proposed for SEAs and VSAs executing maximum output velocity tasks [7], and output velocities of 272% of the maximum motor speed have been measured with VSAs [8].

We took one step back from the teleoperation scenario to exclude effects coming from the teleoperator system (e.g., time-delay, missing fidelity in force-feedback rendering, dynamic constraints of the slave) and performed an experiment with direct use of an elastic tool, asking participants to hit a target as hard as possible with an elastic hammer (section II). We analyzed the task performance as a function of the stiffness of the elastic hammer, the number of trials and the motion profile (section III). Our results suggest that humans can intuitively use an elastic hammer exploiting the spring effect (section IV). We conclude that this might be beneficial for future teleoperation systems with VSAs (section VI).

II. METHOD

For this experiment, we modified the one degree of freedom teleoperator system designed by Rebelo et al. [2] for




Fig. 1. Experimental setup. The motion is constrained to rotation around one axis. The extension can be exchanged to allow for different stiffness of the hammer. The movement of the suspended mass shows to the participants how hard they hit the target.

direct interaction with an elastic tool. Fig. 1 shows a diagram of the setup. We used only one unit and decoupled the motor from the output axis by removing the wire from the capstan gear. We replaced the original end-effector by a handle and a hammer. The hammer had two parts: the hammer head and an interchangeable extension. The latter could be completely rigid, in which case the tool was a rigid hammer, or a flexible extension realized with a leaf-spring, in which case the tool was an elastic hammer.

We mounted the hammer hanging down, so that the spring was stable in neutral position with the handle in vertical position. Participants could execute hammering motions by moving the handle with their right hand. The setup constrained the motion to one degree of freedom. Therefore, the motion was quite different from the hammering participants might have been accustomed to. In addition, the backswing was a pushing motion for the participants and the forward swing a pulling motion. Although one might consider this as making the experimental setup unrealistic, we thought that this was beneficial as it reduced possible effects from previously acquired skills of participants. We assumed that participants had a general understanding of elastic tools, but no knowledge about the setup in particular.

We asked 13 participants (8 male, 5 female, 12 right-handed, 1 ambidextrous, age 21-41) to perform a hammering task on the setup. The participants were told that the goal of this experiment was to study how intuitive it is for humans to interact with an elastic tool. The participants were instructed to hit the target with the hammer by executing one backward swing followed by a forward swing. Before every hit, the hammer should be in the resting position with the hammer head slightly touching the target in order to ensure the hammer head was at rest before the motion. The hammer should hit the target only once per movement and not repeatedly. Comfortable and similar effort should be used for all hits, focusing on the timing for switching

from backward swing to forward swing (related to the input frequency explained below) to achieve the maximum impact. The experiment was approved by the human research ethics committee of TU Delft and all participants gave written informed consent before participation.

The mass of the hammer head and its distance to the center of rotation were constant at 115 g and 150 mm. We used five different extensions: the rigid extension and four leaf-springs of different thickness and stiffness: 0.4 mm (0.62 Nm/rad), 0.6 mm (2.3 Nm/rad), 0.8 mm (4.1 Nm/rad), and 1.0 mm (11 Nm/rad). In the following we will use “Rigid” or the spring thickness as names for the experiment conditions. As a first approximation, the combination of leaf-spring and hammer head can be modeled as a mass-spring system. Its resonance frequency f_0 depends only on the rotational spring constant κ and the moment of inertia of the hammer head I according to the following equation: $f_0 = \frac{1}{2\pi} \sqrt{\frac{\kappa}{I}}$. If the input excitation through the handle is close to the resonance frequency, mechanical resonance occurs. In this case, the hammer head velocity amplitude is greater than the handle velocity amplitude. Thus, for each condition, the participants needed to adapt to the resonance frequency of the elastic tool to achieve a high peak output velocity with a low peak input velocity corresponding to a high impact with little effort. However, besides the instructions explained above, the participants were not taught a specific input profile, but had to adjust in the way they deemed the best for moving the suspended mass compared to the effort they put into the strike.

For every condition, the participants had a learning phase of approximately 100 trials and a performance phase of approximately 20 trials. During the learning phase, the participants were asked to try to improve their performance by experimenting with different timings (i.e., executing the motion at different speeds). During the performance phase, the participants were asked to use the timing which they had found to be the best during the learning phase. We chose 100 trials, as experience from the pilot experiments showed that this was a trade-off allowing enough time to get used to a condition without getting tired. After every phase there was a short interruption to download the measurement data, and between every condition there was a short break to change the hammer extension. We chose a different sequence of conditions per participant in order to compensate for learning effects from one condition to the next.

We measured the handle position and the hammer position using a Vicon motion capture system running at a frame rate of 1 kHz. For this we attached passive reflective markers on an extension to the handle and at the hammer head (cf. Fig. 2).

III. ANALYSIS

The markers fixed to the handle and the hammer head defined two positions in the motion capture system: position E , coupled to the handle, and position H , coupled to the hammer head. We converted the Cartesian coordinates x , y , z measured by the motion capture system into polar




Fig. 2. Participant holding the handle of the setup. The reflective markers on the handle and on the hammer head allow position tracking with the motion capture system.




Fig. 3. Mapping from motion capture system coordinates x, y, z in the reference frame R to rotation angles of the handle θ_e and the hammer head θ_h around the rotation axis center C . The points E and H are the positions of the objects defined by the markers on the handle extension and the hammer head, respectively.

coordinates θ_e and θ_h by mapping them to two circles around a common center C and lying in the same plane (cf. Fig. 3). This produced two one-dimensional signals for analysis. We segmented the time sequences to isolate the individual hammer strikes and synchronized the segments to have the impact with the suspended mass at time 0 s. For the following calculations, we cropped the segments to the time -425 ms up to 0 s. For every segment, we calculated the gain G defined by the peak hammer head velocity $\dot{\theta}_{h,max}$ divided by the peak handle velocity $\dot{\theta}_{e,max}$, and estimated the input frequency f from the half period between minimum and maximum handle velocity (cf. Fig. 4).

As explained above, we modeled the elastic hammer as a mass-spring system with the handle velocity as input and the hammer head velocity as output. We extracted the second-order transfer function for the different extensions through system identification to determine the resonance frequency f_0 and damping ratio ζ (table I). For the system identification we used a separate measurement with manual input at different frequencies with the hammer head able to move freely (suspended mass removed). Based on this transfer function, we calculated the reference gain G_{ref} with a theoretical reference velocity profile input, which was a bang-bang type profile with half a period of back-swing

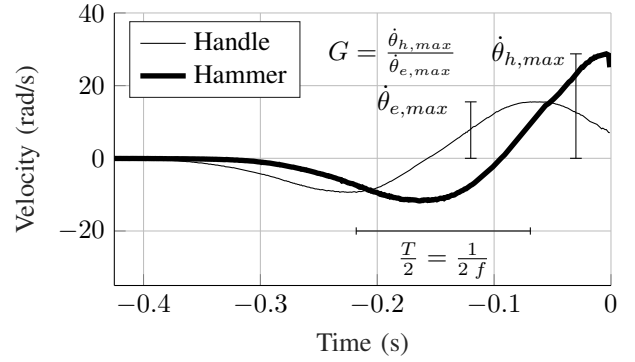


Fig. 4. Typical velocity profile for one hammer strike with the metrics used for analysis, showing the handle velocity $\dot{\theta}_e$ (thin line), the hammer velocity $\dot{\theta}_h$ (thick line), the peak velocities $\dot{\theta}_{e,max}$ and $\dot{\theta}_{h,max}$, and the half-period of the input velocity $\frac{T}{2}$ from minimum to maximum of the handle velocity.

TABLE I
SYSTEM IDENTIFICATION RESULTS
FOR THE TRANSFER FUNCTION

$$H(s) = \frac{\dot{\theta}_h}{\dot{\theta}_e} = \frac{1}{1 + \frac{2\zeta}{\omega_0}s + \frac{1}{\omega_0^2}s^2}$$

Condition	$f_0 = \frac{\omega_0}{2\pi}$	ζ
0.4 mm	2.98 Hz	$23 \cdot 10^{-3}$
0.6 mm	4.82 Hz	$32 \cdot 10^{-3}$
0.8 mm	6.93 Hz	$40 \cdot 10^{-3}$
1.0 mm	9.89 Hz	$17 \cdot 10^{-3}$

followed by half a period of forward swing according to Haddadin et al. [7].

Figs. 5 to 7 show the qualitative difference between the velocity profiles and system responses. Fig. 5 shows the measured input (handle) and output (hammer head) velocity profiles after normalization and taking the median of all participant performance phase trials. The median velocity profile is suitable as a representation of a typical velocity profile as the participants adopted very similar profiles at the end of the learning phase. Fig. 6 visualizes the simulated system response to the human input velocity profile, showing a good match between simulation and measurement data compared to Fig. 5. Fig. 7 shows the results of the simulated system response to the reference input velocity profiles. We can see that the difference in timing is more prominent for the reference profiles than for the human input profiles. The reference profiles have much shorter back swing and forward swing phases for stiffer springs than for softer springs whereas the human input profiles are fairly similar (this is detailed quantitatively in the results section).

We used the following metrics to evaluate the performance of the participants:

- 1) The peak hammer velocity $\dot{\theta}_{h,max}$ shows which condition gives the best results overall.
- 2) The gain G shows which condition is the most efficient in terms of output per input.
- 3) The relative gain $\eta = \frac{G}{G_{ref}}$ indicates how well the participants performed compared to the bang-bang type input profile.




Fig. 5. Median normalized velocity profiles of input (handle) and output (hammer head) for the performance phase trials of the different conditions (1 = peak input velocity of the respective trial). Thin lines represent the respective input velocity and thick lines the output velocity. For the rigid condition, input and output velocity are identical. The input velocity profile is similar for the different conditions but executed slower for softer springs and faster for stiffer springs.




Fig. 6. Simulated system response to a human input profile. Thin lines are used for the input and thick lines for the output velocity (1 = peak input velocity of the respective trial).




Fig. 7. Simulated system response to a square input profile of an ideal velocity source for reference gain calculation. Thin lines are used for the input and thick lines for the output velocity.

- 4) The estimated input frequency f shows how the participants adapted to the stiffness of the hammer for the different conditions (values closer to the resonance frequency meaning better adaptation).
- 5) The relative frequency error $e = \frac{f-f_0}{f_0}$ indicates how close to the resonance frequency the participants excited the mass-spring system for the different conditions.




Fig. 8. Peak hammer velocity by condition. The thin lines connect the median of the performance trials of each subject. The thick line connects the medians of the medians of all subjects. Superscript numbers identify pairs of conditions with statistically significant difference at the 0.5%-level.

IV. RESULTS

For the following results, we only considered the trials of the performance phase. We calculated the median per participant for each condition and used Friedman's test to check for column effects after adjusting for possible row effects ($\alpha = 5\%$, row effects: subject-dependent variations, column effects: condition-dependent variations). We then used Wilcoxon's signed-rank test for pairwise testing of the conditions ($\alpha = \frac{2}{k(k+1)} \cdot 5\%$, with k the number of conditions). With two participants, not all conditions could be performed for organizational reasons (participant 3 missing 1.0 mm condition and rigid condition, participant 6 missing 0.4 mm condition and 0.6 mm condition). Therefore, the results of these participants were excluded from Friedman's test ($n = 11$) and the concerned conditions were excluded from the pairwise testing ($n = 12$ for the tests including one of the conditions mentioned above, $n = 11$ for the tests including two of the conditions mentioned above). Fig. 8 shows the achieved peak output velocity results, ranging from 10.69 rad/s to 46.06 rad/s median per participant per condition. The column effects after adjusting for possible row effects are different at the 5%-level (Friedman's test: $\chi^2 = 25.60$, $n = 11$). The peak output velocity is significantly higher at the 0.5%-level with the 0.6 mm, 0.8 mm and 1.0 mm leaf-springs compared to the rigid extension. The increase of median peak output velocity is more than 63% from the rigid extension to the 0.6 mm and 0.8 mm leaf-springs (rank-biserial correlation for 0.6 mm vs rigid: 0.97, rank-biserial correlation for 0.8 mm vs rigid: 1.00, rank-biserial correlation for 1.0 mm vs rigid: 1.00).

Fig. 9 shows the achieved gain results, ranging from 1.04 to 3.76 median per participant per condition (the gain for the rigid extension is always 1). After adjusting for possible row effects, the column effects are different at the 5%-level (Friedman's test: $\chi^2 = 30.84$, $n = 11$). The gain is significantly higher at the 0.5%-level with the flexible extension compared to the rigid extension. The increase in median gain is more than 100% from the rigid extension to the 0.6 mm leaf-spring (rank-biserial correlation of all




Fig. 9. Gain by condition. The thin lines connect the median of the performance trials of each subject. The thick line connects the medians of the medians of all subjects. Superscript numbers identify pairs of conditions with statistically significant difference at the 0.5%-level.

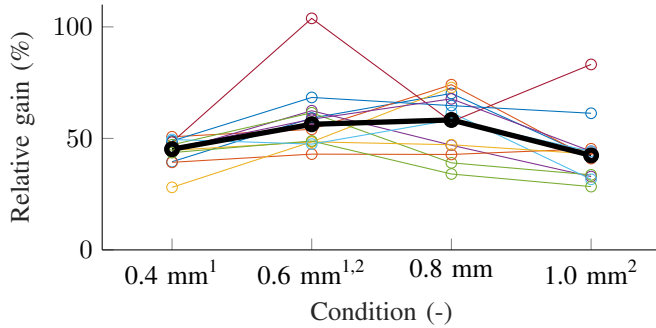


Fig. 10. Relative gain by condition. The thin lines connect the median of the performance trials of each subject. The thick line connects the medians of the medians of all subjects. Superscript numbers identify pairs of conditions with statistically significant difference at the 0.83%-level.

flexible conditions vs the rigid condition: 1.00).

Fig. 10 shows the achieved relative gain results, ranging from 28.05% to 103.82% median per participant per condition (the latter being an unusual subject as the others achieved at most 68.36% for this condition). After adjusting for possible row effects, the column effects are different at the 5%-level (Friedman's test: $\chi^2 = 14.67$, $n = 11$). The relative gain is significantly higher at the 0.83%-level for the 0.6 mm leaf-spring compared to the 0.4 mm and the 1.0 mm leaf-spring (rank-biserial correlation for 0.6 mm vs 0.4 mm: 0.97, rank-biserial correlation for 0.6 mm vs 1.0 mm: 0.97). There is no significant difference between the 0.6 mm and the 0.8 mm leaf-spring. Fig. 11 shows the estimated input frequency results, ranging from 1.57 Hz to 7.27 Hz median per participant per condition. After adjusting for possible row effects, the column effects are different at the 5%-level (Friedman's test: $\chi^2 = 24.87$, $n = 11$). The estimated input frequency is significantly lower at the 0.83%-level for the 0.4 mm leaf-spring compared to all other conditions (rank-biserial correlation for 0.4 mm vs 0.6 mm: -1.00, rank-biserial correlation for 0.4 mm vs 0.8 mm: -0.97, rank-biserial correlation for 0.4 mm vs 1.0 mm: -1.00, rank-biserial correlation for 0.4 mm vs rigid: -1.00).

Fig. 12 shows the relative frequency error results, ranging from -64.51% to 26.49% median per participant per con-




Fig. 11. Input frequency by condition. The thin lines connect the median of the performance trials of each subject. The thick line connects the medians of the medians of all subjects. Superscript numbers identify pairs of conditions with statistically significant difference at the 0.5%-level.




Fig. 12. Relative frequency error by condition. The thin lines connect the median of the performance trials of each subject. The thick line connects the medians of the medians of all subjects. Superscript numbers identify pairs of conditions with significant difference at the 0.83%-level.

dition. After adjusting for possible row effects, the column effects are different at the 5%-level (Friedman's test: $\chi^2 = 27.76$, $n = 11$). The relative frequency error is significantly better at the 0.83%-level for the 0.4 mm and 0.6 mm leaf-spring compared to the 0.8 mm and 1.0 mm conditions (rank-biserial correlation for 0.4 mm vs 0.8 mm: 0.97, rank-biserial correlation for 0.4 mm vs 1.0 mm: 1.00, rank-biserial correlation for 0.6 mm vs 0.8 mm: 1.00, rank-biserial correlation for 0.6 mm vs 1.0 mm: 1.00).

Fig. 13 shows the gain over the cumulative trial count, independent of the condition, combining learning and performance trials. We normalized the gains so that a value of 1 corresponds to the median over all trials of the same condition, and we grouped per ten trials, calculating the medians per participant and the medians of the medians of all participants. No significant learning effect can be seen for most learning phases. In some cases, a small learning effect can be observed over the first 20 to 30 trials of the learning phase, and in others, the performance decreases after 60 to 70 trials, perhaps due to participants getting bored.

V. DISCUSSION

The participants were able to use the elastic tool to their benefit, consistently reaching higher peak output velocities in




Fig. 13. Normalized gain over cumulated trial count for all flexible condition trials. A value of 1 corresponds to the median over all trials of the same condition. Thin lines connect the median gains over ten trials per participant. The thick line connects the median of the medians over ten trials of all participants. Learning phases, labeled L1 to L4, consist of 100 trials and performance phases, labeled P1 to P4, consist of 20 trials. Vertical bars show the limits of the phases. After every learning phase (L1, L2, L3, L4) there was a short interruption to download the measurement data; after every performance phase (P1, P2, P3, P4) there was a short break to change the hammer extension.

the hammering motion in the flexible conditions than in the rigid condition, exploiting the spring gain. They achieved this performance after a short learning phase of less than 20 to 30 trials for new stiffness conditions, not showing significant improvements over the longer learning phase of 100 trials. Therefore, the participants seem to have an intuitive understanding of the dynamic system properties.

The participants actively adapted to the system resonance by adopting higher excitation frequencies for stiffer springs, consistent with Hatsopoulos et al. [9]. However, the participants stayed below approximately 6 Hz excitation frequency (one outlier for the 0.8 mm condition), even for the 0.8 mm and 1.0 mm conditions that would require higher excitation frequencies for best performance (resonance frequency of 6.93 Hz and 9.89 Hz). We conclude that human dynamic limitations only allow input frequencies of up to approximately 6 Hz for the task used in our experiment. Overall, the participants achieve around half to two-thirds of the gain possible with a perfect velocity source (unexplained outliers achieved a higher gain), another effect of the human dynamic limitations.

The participants achieved the best results with the 0.6 mm leaf-spring, obtaining a median gain of 2.04. Although some participants performed better with the 0.8 mm leaf-spring, others had more difficulties with the faster motion, resulting in a bigger spread of results between participants. Therefore, we observed no statistically significant improvement in the gain for the 0.8 mm leaf-spring compared to the other flexible conditions, and the 0.6 mm leaf-spring with a resonance frequency of 4.82 Hz seems to be the one the participants could best adjust to.

Although we could show that humans can use elastic tools intuitively, we cannot yet explain the feedback mechanisms important for humans to be able to find and excite the system resonance frequency. We intend to do follow-up experiments in a teleoperation setup to separate the effects of force and visual feedback on human performance.

VI. CONCLUSION

This paper studied how humans interact with elastic tools. We performed an experiment with untrained participants

hammering with an elastic hammer. The participants were able to intuitively exploit the elastic properties to their advantage, achieving significantly higher peak output velocities with an elastic hammer than with a rigid hammer. The participants quickly adapted to the resonance frequency for different stiffnesses. We conclude that humans should be able to use series elastic actuators also in teleoperator systems, if the teleoperator is sufficiently transparent. Our future research will study the feedback mechanisms used by humans for dynamic tasks and how to develop dynamic teleoperator systems that are intuitive to use. With this basis, a novel class of teleoperator systems can be developed that allow force amplification for dynamic tasks requiring high peak forces on the slave side.

REFERENCES

- [1] D. A. Lawrence, "Stability and transparency in bilateral teleoperation," *Robotics and Automation, IEEE Transactions on*, vol. 9, pp. 624–637, 1993.
- [2] J. L. Pinto Rebelo, *Robust and Transparent Multi-Degree-of-Freedom Bilateral Teleoperation with Time-Delay*. TU Delft, Delft University of Technology, 2015.
- [3] R. Groten, D. Feth, A. Peer, and M. Buss, "Shared decision making in a collaborative task with reciprocal haptic feedback - an efficiency-analysis," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, 2010, pp. 1834–1839.
- [4] R. Van Ham, T. G. Sugar, B. Vanderborght, K. W. Hollander, and D. Lefeber, "Compliant actuator designs," *Robotics & Automation Magazine, IEEE*, vol. 16, pp. 81–94, 2009.
- [5] D. Braun, M. Howard, and S. Vijayakumar, "Optimal variable stiffness control: formulation and application to explosive movement tasks," *Autonomous Robots*, vol. 33, pp. 237–253, 2012.
- [6] M. Garabini, A. Passaglia, F. Belo, P. Salaris, and A. Bicchi, "Optimality principles in variable stiffness control: The VSA hammer," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, 2011, pp. 3770–3775.
- [7] S. Haddadin, M. Weis, S. Wolf, and A. Albu-Schäffer, "Optimal control for maximizing link velocity of robotic variable stiffness joints," in *IFAC World Congress*, 2011, pp. 6863–6871.
- [8] S. Wolf and G. Hirzinger, "A new variable stiffness design: Matching requirements of the next robot generation," in *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, 2008, pp. 1741–1746.
- [9] N. G. Hatsopoulos and W. H. Warren Jr, "Resonance Tuning in Rhythmic Arm Movements," *Journal of Motor Behavior*, vol. 28, pp. 3–14, 1996.