Amplified Control for Robotic Teleoperation

Barton L. Paulhamus, Edward W. Staley, Corban G. Rivera, Kapil D. Katyal, and Chien-Ming Huang

***Abstract*— Assistive robotics holds the promise of bettering the lives of countless people throughout the world. As robots become more complex, the degrees-of-freedom for controlling robotic systems is rapidly outpacing the degrees-of-control that can be supplied by humans via conventional interfaces. In this paper, we describe a novel paradigm, *amplified control*, that strives to capture the adaptability of teleoperation while also leveraging the reduced user burden offered by shared control approaches. The novelty of this approach is that machine intelligence amplifies human intelligence for robotic control as opposed to replacing it, supplementing it, or augmenting it. If successful, our novel control paradigm will lower the barrier of entry (e.g., overcoming physical limitations and lessening cognitive load) for people to operate complex robotic systems for assistive robotics as well as other domains.**

I. Introduction

Assistive robotics holds the promise of bettering the lives of countless people throughout the world including individuals with debilitating conditions, individuals recovering from injuries, and the elderly. As robots become more complex, the degrees-of-freedom for controlling robotic systems is rapidly outpacing the degrees-of-control that can be supplied by humans via conventional interfaces. Machine intelligence offers the potential of lowering user burden by making robots more self-sufficient. However, the current state of machine intelligence lacks the ability to infer and reason, leading to robots that cannot adapt to novel circumstances (i.e., novel tasks and/or novel environments). This introduces a tradeoff in the design space between increasing the system's adaptability to novel circumstances and reducing the user burden to operate the system. This tradeoff is particularly relevant to assistive robotics. As an assistive device, robots need to be robust to the novel circumstances that accompany humans throughout their activities of daily living. At the same time, assistive robotics strives to make robotics available to anyone. In fact, often with assistive robotics, the users' physical and/or cognitive abilities are limited.

In this paper, we describe a novel paradigm, ***amplified control***, that strives to capture the adaptability of teleoperation while also leveraging the reduced user burden offered by shared control approaches (Fig 1). We call this new paradigm amplified control in reference to its ability to amplify a small number of input channels in order to control a larger number of robotic command signals. Amplified control is an extension of teleoperation that expands the user's degrees-of-control beyond a 1-to-1 mapping to the robot's degrees-of-freedom. Instead, machine intelligence is

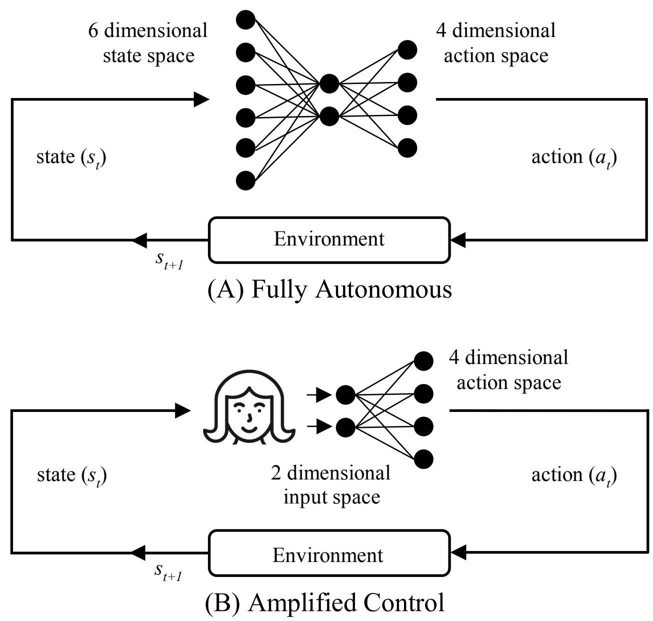

Fig. 1. Our approach to amplified control uses deep neural networks and consists of two steps: (A) training a deep neural network, and (B) replacing the front half of the network with a human user.

used to amplify a lower degree-of-control in order to operate a higher degree-of-freedom robot. However, unlike shared control approaches, the user remains in complete control of the robot. This makes it possible for a user to "teleoperate" a high degree-of-freedom robotic system with lower degrees-of-freedom of input. The novelty of this approach is that machine intelligence amplifies human intelligence for robotic control as opposed to replacing it, supplementing it, or augmenting it. If successful, our novel control paradigm will lower the barrier of entry (e.g., overcoming physical limitations and lessening cognitive load) for people to operate complex robotic systems for assistive robotics as well as other domains.

With amplified control, the user generates a small number of control channels and machine intelligence amplifies these inputs in order to control a higher degree-of-freedom robot. Amplified control is most similar to teleoperation [1] and supervised control [2]. Like teleoperation, the user maintains control of the robot throughout the lifecycle of the task. This enables the robot to be used under novel circumstances (with the caveat that a level of machine intelligence is contained in the amplification function). At the same time, amplified control shares similarities with supervised control by decreasing user burden. However, unlike supervisory control, control of the robot is never handed over to the machine intelligence.

In addition to assistive devices, robotic arms have shown benefits to a range of domains including manufacturing,

B.L. Paulhamus, E. S. Staley, C. G. Rivera, and K. D. Kapil are with the Applied Physics Laboratory, The Johns Hopkins University, Baltimore, MD, USA. (email: Bart.Paulhamus@jhuapl.edu).

C.-M. Huang is with the Department of Computer Science, The Johns Hopkins University, Baltimore, MD, USA. (e-mail: cmhuang@cs.jhu.edu).

robotic surgery, and extreme environments such as space, undersea, and bomb disposal. The goals of amplified control make it particularly attractive to the unique priorities of assistive robotics.

II. Approach

Our approach to amplified control uses deep neural networks and consists of two steps as described in Fig. 1. Initially, a deep neural network is trained as a fully autonomous agent to perform a specific task within our environment. In Fig. 1.A, this network is a control policy that takes a six-dimensional state space as input and outputs a four-dimensional action space. The critical aspect of this concept is that we purposely design a bottleneck layer into the network architecture. The number of nodes in the bottleneck layer must be less than the number of nodes in the output layer (i.e., the dimensionality of the action space). In Fig. 1 the bottleneck layer has two nodes and represents two latent variables of the policy function.

In the second step of our approach we divide the trained neural network at the bottlenecked layer and discard the portion of the network that precedes the bottleneck layer. The bottleneck layer becomes the new input layer to our network. We maintain the back half of the trained network that converts the bottleneck layer to the output layer. In Fig. 1.B, a user supplies two inputs (i.e., the two latent variables in the bottleneck layer) and the network translates those inputs into a four-dimensional action. It is important to highlight that in our initial designs these latent variables are uninterpretable. That is, the latent variables are simply real numbers between -1.0 and 1.0 without any meaning. We speculate that humans will be able to quickly learn how to manipulate these variables to successfully perform actions with a robot in the environment.

III. Preliminary Results

Our preliminary results focus on validating the underlying intuition behind our approach in the context of wheelchair mounted robotic arms [3]. First, we focused on the feasibility of introducing bottleneck layers to neural networks for robot control. In simulation (see center of Fig. 2), we manually generated 100 trajectories of a UR5 robotic arm performing three separate motion primitives: reach, grasp, and lift (300 total trajectories). With this data, we used supervised learning to learn a policy for mapping an 18D state space to a 7D action space for each primitive. We repeated this process and varied the number of latent variables in the bottleneck layer of the neural network. This allowed us to analyze how the number of latent variables affected our performance (see right side of Fig. 2). Performance was measured by the mean squared error (MSE) of the validation data (10% of the total dataset). Our results reveal two interesting findings: 1) we were able to bottleneck all three networks down to four latent variables without sacrificing performance, and 2) with less than four latent variables, performance degraded at different rates for each of the primitives.

Additionally, our internal engineering tests confirmed our intuition that amplified control is a plausible solution for high DOF control with a lower DOF of inputs. In these tests, we

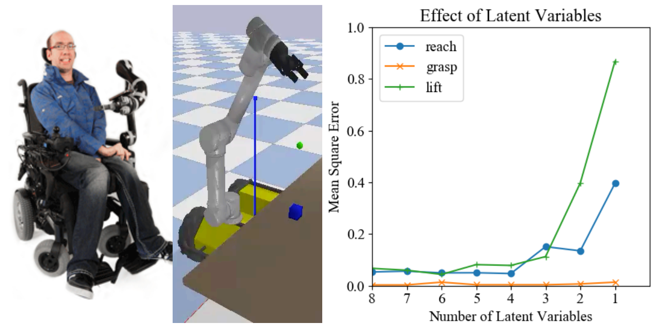

Fig. 2. JACO robotic arm mounted on a power wheelchair [3] (left), our simulation environment (center), and preliminary results (right)

constructed a simple 5DOF robotic arm in simulation. We used deep reinforcement learning to learn a control policy using a network architecture that contained a bottleneck layer with 3 latent variables. Then we chopped the network in half, exposing the 3 latent variables as 3 inputs to the human operator. The inputs were uninterpretable. Nevertheless, we mapped those inputs to 3 inputs on a standard gaming joystick to see if we could successfully control the 5 DOF arm.

Our anecdotal engineering tests generated a few observations: 1) we successfully used amplified control to complete simple tasks, and 2) we noticed a degree of familiarity; the longer we operated the system, the more comfortable it was to use it. These were by no means rigorous experiments or analyses. However, it did provide proof-of-concept evidence to further pursue this research.

IV. Conclusion and Future Work

Amplified control presents new opportunities for direct, physical human-robot interactions for assistive applications. Our near-term future work has two foci. First is understanding the best way to design amplified control algorithms. Our preliminary results have inspired a number of design-based questions: What is the relationship between the action space and the number of required latent variables? How generalizable are the learned networks? Can we combine latent variables (or subsets of latent variables) to create networks capable of performing more complex and novel tasks? Can we increase the interpretability, transparency, and semantic applicability of latent variables? If so, what is the effect on performance? The second focus of our near-term future work is to conduct a human-robot interaction study to investigate how the task performance and user experience of amplified control compares to traditional and state-of-the-art robotic control paradigms.

References

[1] G. S. Gupta, S. C. Mukhopadhyay, C. H. Messom, S. N. Demidenko, "MasterSlave Control of a Teleoperated Anthropomorphic Robotic Arm With Gripping Force Sensing", *IEEE Transactions on Instrumentation and Measurement*, vol. 55, pp. 2136-2145, 2006.

[2] D. Gopinath, S. Jain, and B.D. Argall, "Human-in-the-Loop Optimization of Shared Autonomy in Assistive Robotics," *IEEE Robotics and Automation Letters*, July 2016.

[3] A. Campeau-Lecours, V. Maheu, S. Lepage, S. Latour, L. Paquet, N. Hardie, "JACO assistive robotic device: Empowering people with disabilities through innovative algorithms", *Proc. 2016 Rehabil. Eng. Assistive Technol. Soc. North Amer. Annu. Conf.*, Jul. 2016.