

## ARTIFICIAL INTELLIGENCE

## Co-designing hardware and control for robot hands

Tianjian Chen<sup>†</sup>, Zhanpeng He<sup>†</sup>, Matei Ciocarlie\*

Policy gradient methods can be used for mechanical and computational co-design of robot manipulators.

Copyright © 2021  
The Authors, some  
rights reserved;  
exclusive licensee  
American Association  
for the Advancement  
of Science. No claim  
to original U.S.  
Government Works

Roboticians have long argued that embodied agents' intelligence can extend beyond a purely computational brain: Hardware plays a key role in determining behavior and can embody intelligence just like computational components do. Design optimization thus goes hand in hand with deriving effective control algorithms or policies. Here, we focus specifically on the concept of co-optimization (or co-design), which aims to achieve desired behaviors by optimizing mechanical and computational components simultaneously, with and for each other.

A number of foundational studies on co-optimization in robotics used genetic algorithms (1) or evolutionary computation (2) to optimize both hardware and software. The choice of biologically inspired optimization algorithms seems natural given the well-known evidence of co-optimization in nature, including the evolution of dexterous manipulators (3). However, when designing artificial agents, many other optimization strategies have also been used, such as efficient gradient-based methods (4). As more powerful optimization algorithms are introduced, they could potentially enable new directions in co-optimization. In particular, the last decade has seen the advent of GPU-accelerated optimization for computational components such as universal function approximators. Could such methods be used to simultaneously advance mechanical intelligence in robot hands and beyond?

Recent computational results have only strengthened the evidence favoring co-design by showing that end-to-end optimization often outperforms sequential alternatives. We now know that filters for low-level feature detection in images work best when optimized end to end based on performance of the desired task, which could range from image classification to autonomous game play (5). Through this lens, mechanical/computational co-design in robotics is simply


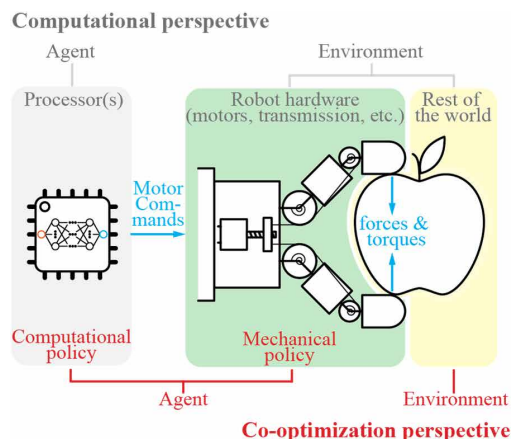
end-to-end optimization taken to its logical conclusion. Aligned with such renewed incentives, we bring attention to three developments, none of them motivated primarily by co-optimization but all bearing notable promise in this direction:

**1) Continued progress in rigid- and soft-body dynamics simulators.** Co-optimization in robotics can rarely afford to wait for the physical prototyping of each hardware iteration; thus, simulation has always been central to the effort. Over the past 5 years, multipurpose physics engines (such as MuJoCo or Bullet) able to handle complex kinematic chains and transient frictional contact have finally become general and robust enough to be considered commodities. Concurrently, general access to cloud computing has lowered the computational cost barrier to widespread use of large-scale simulation.

**2) Sim-to-real transfer via domain randomization.** Building accurate simulators is still an active and compelling research

problem. However, domain randomization has emerged as a complementary approach: By introducing variation in the behavior of an imperfect simulator, we can obtain a policy or controller that is general enough to discount simulation inaccuracies and thus transfer to the real world. Domain randomization inherently requires optimization algorithms that can formally reason about the large ensuing variance in the behavior of the agent, which is particularly difficult because the simulated agent itself can rarely be given access to observe the parameters of the simulation.

**3) Advances in policy optimization for complex motor skills.** The last half decade has seen impressive advances in the optimization of the software side of a motor ability. Although examples also include results in evolutionary algorithms or optimal control, we focus here specifically on policies optimized through deep reinforcement learning using models and stochastic algorithms that



**Fig. 1. Hardware as policy.** (Left) We contrast the co-optimization perspective, where part of the robot's hardware is considered as a mechanical policy and optimized alongside the computational policy, against the computational-only optimization perspective, where the robot's hardware is typically considered part of the environment. (Right) A robot hand with a compliant underactuated transmission designed using the co-optimization paradigm.

Department of Mechanical Engineering, School of Engineering and Applied Science, Columbia University, 228 Mudd Building, 500 W 120th St., New York, NY 10027, USA.

<sup>†</sup>These authors contributed equally to this work.

\*Corresponding author. Email: matei.ciocarlie@columbia.edu

are inherently able to consider different outcomes from similar observations (due to partially observable states). As a result, these algorithms are a natural match for domain randomization, allowing the physical realization of complex motor skills trained in simulation.

These developments all point to previously unexplored opportunities for mechanical/computational co-optimization. We have better simulation tools, better methods for transferring simulation results to real robots, and better methods to derive controllers or policies for complex motor skills. Indeed, recent years have seen many promising new results. Reinforcement learning algorithms have optimized both design parameters and policies for simulated robots performing difficult motor skills (6–8). Bayesian optimization has also been used to similar effect (9). Looking beyond simulation, studies are also explicitly considering the fact that a number of real physical hardware iterations might still be expected when doing simulated co-optimization (8,9).

In our own recent work (10), we noticed that hardware components can be deeply similar to policies. A computational policy takes as input sensor observations, processes them (e.g., through a deep neural network), and outputs actions as motor commands. Hardware takes the problem from here: It accepts motor commands as input and processes them through the robot's body to finally output actions such as forces acting on the world. The behavior of a computational policy is governed by tunable parameters (e.g., weights and biases), much like hardware

behavior is determined by tunable design parameters.

We applied this concept to the problem of underactuated manipulation, where a computational policy can move the hand and provide setpoints for the motor controlling the fingers. The hand transmission mechanism then acts as a mechanical policy, converting motor forces to joint torques. We combine models of the computational and mechanical policies into a joint policy, which we optimize via off-the-shelf policy gradient algorithms. During training, we simulate the rest of the world (beyond the transmission mechanism) in an off-the-shelf dynamics simulator, with extensive domain randomization. Once training is complete, the optimized parameters of the mechanical policy are used to build physical prototypes, which are in turn controlled by the computational policy to operate in the real world. Importantly, simultaneous optimization of hardware and software is shown to outperform approaches that alternate between the two.

The results of the studies we cited above highlight the fact that current optimization methods originally developed for the computational engines of intelligent systems can also be applied, in the case of embodied agents, to hardware. They suggest that, for complex motor skills such as manipulation, intelligence is not confined to computation: It can reside in actuators, sensors, or other hardware components just as well. Mechanical and computational co-design can take advantage of such methods to jointly optimize hardware and software, aiming for

intelligent behavior by physical agents in complex environments.

## REFERENCES

1. K. Sims, Evolving virtual creatures, in *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques* (ACM, 1994), pp. 15–22.
2. H. Lipson, J. B. Pollack, Automatic design and manufacture of robotic lifeforms. *Nature* **406**, 974–978 (2000).
3. R. W. Young, Evolution of the human hand: The role of throwing and clubbing. *J. Anat.* **202**, 165–174 (2003).
4. J.-H. Park, H. Asada, Concurrent design optimization of mechanical structure and control for high-speed robots. *J. Dyn. Sys. Meas. Control.* **116**, 344–356 (1994).
5. V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, M. Riedmiller, Playing atari with deep reinforcement learning. arXiv:1312.5602 (2013).
6. D. Ha, Reinforcement learning for improving agent design. arXiv:1810.03779 (2018).
7. C. Schaff, D. Yunis, A. Chakrabarti, M. R. Walter, Jointly learning to construct and control agents using deep reinforcement learning, in *IEEE International Conference on Robotics and Automation* (IEEE, 2019), pp. 9798–9805.
8. K. S. Luck, H. Ben Amor, R. Calandra, Data-efficient co-adaptation of morphology and behavior with deep reinforcement learning, in *Conference on Robot Learning* (2019).
9. T. Liao, G. Wang, B. Yang, R. Lee, K. Pister, S. Levine, R. Calandra, Data-efficient learning of morphology and controller for a microrobot, in *2019 International Conference on Robotics and Automation (ICRA)* (IEEE, 2019), pp. 2488–2494.
10. T. Chen, Z. He, M. Ciocarlie, Hardware as policy: Mechanical and computational co-optimization using deep reinforcement learning, in *Conference on Robot Learning* (2020).

10.1126/scirobotics.abg2133

**Citation:** T. Chen, Z. He, M. Ciocarlie, Co-designing hardware and control for robot hands. *Sci. Robot.* **6**, eabg2133 (2021).

## Co-designing hardware and control for robot hands

Tianjian Chen, Zhanpeng He, and Matei Ciocarlie

*Sci. Robot.* **6** (54), eabg2133. DOI: 10.1126/scirobotics.abg2133

### View the article online

<https://www.science.org/doi/10.1126/scirobotics.abg2133>

### Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

---

*Science Robotics* (ISSN 2470-9476) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.

Copyright © 2021 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works