

Faculty Candidate Seminar – Physical AI**Andrew Wagenmaker**

Postdoc Scholar, Electrical Engineering and Computer Sciences
University of California, Berkeley

Thursday, April 2, 2026

Presentation: 10:30 A.M. – 11:30 A.M.

Q & A: 11:30 A.M. – 12:00 P.M.

MSEE 112 / [Zoom](#)

Physical Agents that Learn from Experience**Abstract**

Humans fundamentally learn through interaction with and experience in the physical world, yet modern approaches applying AI to robotics rely primarily on learning from static, offline sources of data. While this approach has enabled exciting capabilities in some domains, it has proven notoriously difficult to scale to the demands of fully open-world autonomy.

In this talk, I will investigate how we can overcome the limitations of learning with only static data sources, and enable robots to learn from experience as they interact with the physical world. In particular, I will consider how we can collect the experience—explore—that allows for learning and improvement, and how the limited sources of data that are often available to us in the physical world—simulators and human demonstrations—can enable this. I will consider how simulators, even coarse simulators that are insufficient for obtaining effective task-solving policies, can enable efficient exploration, and how the resulting exploration allows for learning performant task-solving robotic behaviors. I will then show how generative robot policies trained on human demonstrations can be utilized to achieve highly focused exploration and enable fast online improvement, and how we can pretrain generative policies on human demonstrations that can themselves collect the experience necessary to learn and improve. Across these examples, I will argue that the insights gained through rigorous analysis are key to uncovering the algorithmic approaches that enable learning from experience, and ultimately bringing AI to the physical world.

Bio

Andrew Wagenmaker is a postdoctoral scholar in Electrical Engineering and Computer Sciences at UC Berkeley working with Sergey Levine. Previously, he completed a PhD in Computer Science at the University of Washington, where he was advised by Kevin Jamieson. Andrew's research focuses on learning in dynamic, interactive settings, spanning fundamental algorithm development to practical approaches for real-world learning and decision-making, particularly toward enabling efficient learning in robotic systems. His work has been recognized by a Best Paper nomination at the Conference on Robot Learning, and he is a recipient of the NSF Graduate Research Fellowship.