
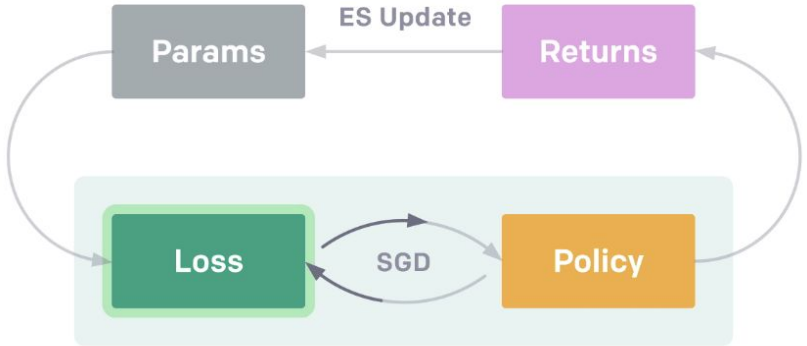


# Evolved Policy Gradients

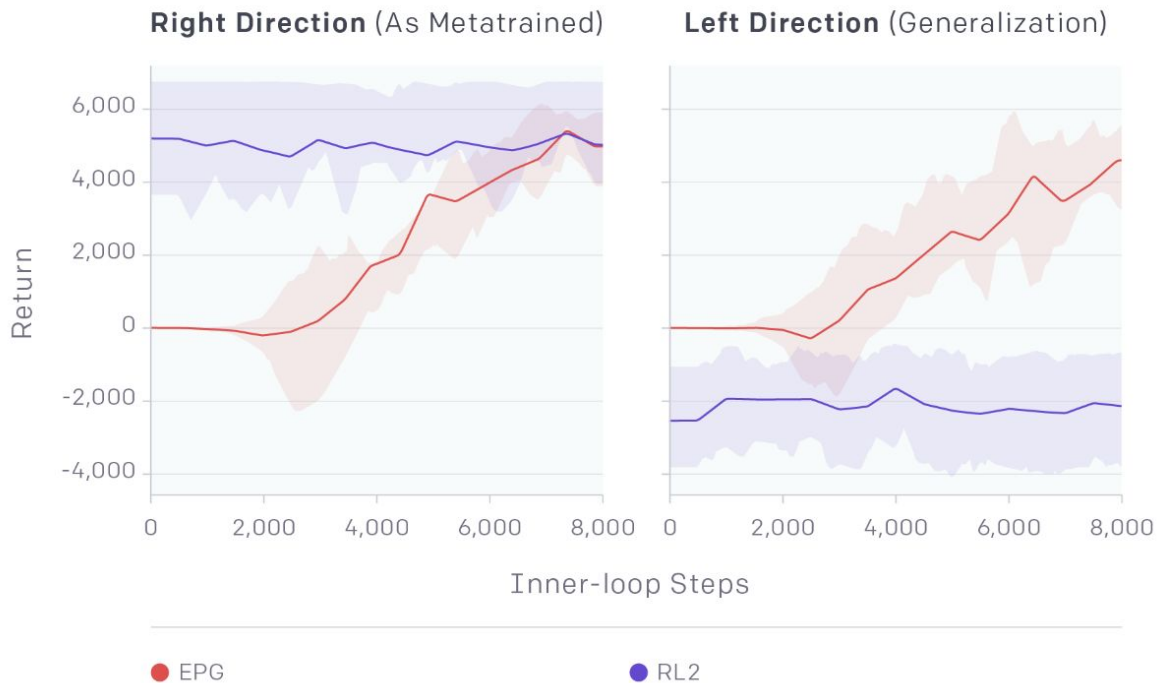

Rein Houthoofd, Richard Y. Chen, Phillip Isola, Bradly C. Stadie,  
Filip Wolski, Jonathan Ho, Pieter Abbeel



# Evolution Strategies + Stochastic Gradient Descent




# Generalization Beyond Training Task Distribution



# Generalization Beyond Training Task Distribution



EPG



RL<sup>2</sup>

# More Information

- Poster: tonight 05:00 -- 07:00 PM @ Room 210 & 230 AB **#16**
- Blog and videos: [blog.openai.com/evolved-policy-gradients](https://blog.openai.com/evolved-policy-gradients)
- Open-source code: [github.com/openai/EPG](https://github.com/openai/EPG)