

About

I'm a Stanford CS BS/MS student, concentrating in AI. My research is at the intersection of reinforcement learning, large language models, and robot learning. More specifically, I am currently interested the following topics:

- **RL post-training techniques:** self-distillation, continual learning, RLVR, reasoning
- **Robot learning:** on-policy RL, VLA's, video models, world models

RL sits in the middle of both of these interests: how can we develop powerful methods using reinforcement learning to (1) best leverage pre-trained representations and (2) adapt to new tasks in a robust and general manner? The core modeling work is the same across the two domains: just that the backbone is different.

For a more detailed writeup, see my [Research Interests Statement](#).

This notebook serves as a public-facing log of experiments, paper notes, and technical deep dives from my ongoing work.

Contact

Please feel free to reach out if you have any thoughts or are open to collaborating on these topics! I've always thought that the best medium for learning is through conversation—though perhaps we are entering an age where the conversations are with non-human intelligence it seems.

- **GitHub:** [roikand](#)
- **Email:** rsikand at stanford.edu

Note: a good bit of the content on this webpage is AI-generated. I often use AI tools to help me understand papers, concepts, and ideas and as a result, some of the content here and there is just a dump from a chat thread for safe keeping, with annotations.