

Jacob Krantz

Email: jkrantz989@gmail.com

Web: jacobkrantz.github.io

LinkedIn: [linkedin.com/in/jacob-krantz](https://www.linkedin.com/in/jacob-krantz)

SUMMARY

I build embodied artificial intelligence. I work on **perception** (scene understanding, long-horizon tracking), **task planning** (LLMs, imitation, reinforcement learning, exploration), and the full **data lifecycle** thereof (curation, pipelining, benchmarking). I publish in vision and learning venues such as CVPR and ICLR. My works have been reported on by Tech Crunch and demoed at the White House Correspondents' Dinner.

EDUCATION

- Ph.D. in Computer Science, Oregon State University** 2019–2023
Advised by [Stefan Lee](#)
- B.Sc. in Computer Science, Gonzaga University** 2015–2019
Magna Cum Laude. Minor in Physics.

PROFESSIONAL EXPERIENCE

- Applied Scientist**, Amazon Frontier AI & Robotics (FAR) Apr 2026 –
San Francisco, CA
- Member of the Technical Staff** (short-term), [Vmax AI](#) Mar 2026
San Francisco, CA
- Established the startup's long-term research agenda: sample-efficient training of SWE coding agents via code world models.
 - Developed a prototype demonstrating improved sample efficiency of analytic policy gradients over model-free RL methods.
 - Analyzed the central task generation pipeline; quantified impacts of multi-task synthesis and identified 1 critical bug.
- Senior Research Scientist**, FAIR Labs @ Meta Platforms Jan 2024–Dec 2025
Menlo Park, CA
Under Roozbeh Mottaghi and Jitendra Malik
- Led VLM training for embodied perception + memory using egocentric videos.
 - Trained LLMs to solve robot task planning from language (“*Tidy up the living room*”) after collecting 100,000 human demonstrations.
 - Unlocked scaling of task generation to enable [PARTNR](#) to become the largest human-robot collaboration benchmark. I designed the eval engine and owned the analysis and delivery of all datasets.
 - Managed interns and collaborated XFN with product design and UX research.
- Research Intern**, FAIR Labs @ Meta Platforms Jun–Dec 2022
Menlo Park, CA
Advised by Devendra Singh Chaplot
- Research Intern**, FAIR Labs @ Meta Platforms Aug–Feb 2021
Menlo Park, CA
Advised by Oleksandr Maksymets
- NSF REU Fellow**, University of Colorado, Colorado Springs May–Aug 2018
Colorado Springs, CO
Advised by Jugal Kalita
- Data Analytics Intern**, Faithlife Corporation May–Aug 2017
Bellingham, WA

Peer-Reviewed Conference Papers

Planning with an Embodied Learnable Memory.P. Parashar, **J. Krantz**, M. Chang, K. Shah, X. Puig, R. MottaghiTo appear: *International Conference on Learning Representations (ICLR)*, 2026.[OpenReview:pre-print](#) [28%]**Do Visual Imaginations Improve Vision-and-Language Navigation Agents?**A. Perincherry, **J. Krantz**, S. Lee.In *Computer Vision and Pattern Recognition (CVPR)*, 2025.[arXiv:2503.16394](#) [22%]**PARTNR: A Benchmark for Planning and Reasoning in Embodied Multi-agent Tasks.**

M. Chang, G. Chhablani, A. Clegg, M.D. Cote, R. Desai, . Hlavac, V. Karashchuk,

J. Krantz, R. Mottaghi, P. Parashar, S. Patki, I. Prasad, X. Puig, A. Rai, R. Ramrakhya,

D. Tran, J. Truong, J. Turner, E. Undersander, T. Yang.

In *International Conference on Learning Representations (ICLR)*, 2025.[arXiv:2411.00081](#) [32%]**Navigating to Objects Specified by Images.****J. Krantz**, T. Gervet, K. Yadav, A. Wang, C. Paxton,

R. Mottaghi, D. Batra, J. Malik, S. Lee, D.S. Chaplot.

In *International Conference on Computer Vision (ICCV)*, 2023.[arXiv:2304.01192](#) [26%]**Iterative Vision-and-Language Navigation.****J. Krantz***, S. Banerjee*, W. Zhu, J. Corso, P. Anderson, S. Lee, J. Thomason.In *Computer Vision and Pattern Recognition (CVPR)*, 2023.[arXiv:2210.03087](#) [26%]**Sim-2-Sim Transfer for Vision-and-Language Navigation in Continuous Environments.****J. Krantz**, S. Lee.In *European Conference on Computer Vision (ECCV)*, 2022.[arXiv:2110.02207](#)**Waypoint Models for Instruction-guided Navigation in Continuous Environments.****J. Krantz**, A. Gokaslan, D. Batra, S. Lee, O. Maksymets.In *International Conference on Computer Vision (ICCV)*, 2021.[arXiv:2110.02207](#) [Oral: 210/6236 = top 3.3%]**Where Are You? Localization from Embodied Dialog.**M. Hahn, **J. Krantz**, D. Batra, D. Parikh, M. Rehg, S. Lee, P. Anderson.In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.[arXiv:2011.08277](#) [25%]**Beyond the Nav-Graph: Vision-and-Language Navigation in Continuous Environments.****J. Krantz**, E. Wijmans, A. Majumdar, D. Batra, S. Lee.In: *European Conference on Computer Vision (ECCV)*, 2020.[arXiv:2004.02857](#) [27%]**Language-Agnostic Syllabification with Neural Sequence Labeling.****J. Krantz**, M. Dulin, P. De Palma.In *IEEE International Conference on Machine Learning and Applications (ICMLA)*, 2019.[arXiv:1909.13362](#) [Oral: 14%]**Abstractive Summarization Using Attentive Neural Techniques.****J. Krantz**, J. Kalita.In *International Conference on Natural Language Processing (ICON)*, 2018.[arXiv:1810.08838](#) [Oral: 9%]**Syllabification by Phone Categorization.****J. Krantz**, M. Dulin, P. De Palma, M. VanDam.In *ACM Genetic and Evolutionary Computation Conference Companion (GECCO)*, 2018.[arXiv:1807.05518](#)

In Review

Asymmetric On-Policy Distillation can solve POMDPs.

S. Poddar, **J. Krantz**, M. Chang, X. Puig, R. Mottaghi, N. Jaques, A. Gupta

Technical Reports

Instance-Specific Image Goal Navigation: Training Embodied Agents to Find Object Instances.

J. Krantz, S. Lee, J. Malik, D. Batra, D.S. Chaplot.

arXiv Preprint, 2022.

[arXiv:2211.15876](https://arxiv.org/abs/2211.15876)

Retrospectives on the Embodied AI Workshop.

M. Deitke, D. Batra, Y. Bisk, T. Campari, A. Chang, D. Chaplot, C. Chen, C. D’Arpino, K. Ehsani, A. Farhadi, F. Li, A. Francis, C. Gan, K. Grauman, D. Hall, W. Han, U. Jain, A. Kembhavi, **J. Krantz**, S. Lee, C. Li, S. Majumder, O. Maksymets, R. Martín-Martín, R. Mottaghi, S. Raychaudhuri, M. Roberts, S. Savarese, M. Savva, M. Shridhar, N. Sünderhauf, A. Szot, B. Talbot, J. Tenenbaum, J. Thomason, A. Toshev, J. Truong, L. Weihs, J. Wu

arXiv Preprint, 2022.

[arXiv:2210.06849](https://arxiv.org/abs/2210.06849)

Ph.D. Dissertation

Semantic Embodied Navigation: Developing Agents That Navigate From Language and Vision.

J. Krantz.

In *Oregon State University*, 2023.

[ScholarsArchive:n583z3248](https://scholarsarchive.library.oregonstate.edu/handle/1957/63248)

INVITED TALKS

Embodied Agents for the Open World: Pursuing Generalization through Home Robotics.

Apple Machine Learning Research (MLR)

Jan 2026

Amazon Frontier AI and Robotics (FAR)

Mar 2026

Human-Aware Embodied Intelligence: Performing Household Tasks For and With Humans.

Meta Reality Labs

Dec 2025

Human-Robot Collaboration for Everyday Household Tasks.

Oregon State University AI Seminar

Mar 2025

Semantic Embodied Navigation: Developing Agents That Navigate From Language and Vision.

Boston Dynamics AI Institute

Sep 2023

FAIR Labs, Meta Platforms

Sep 2023

ADVISED STUDENTS

[Sriyash Poddar](#) (UW PhD Student, FAIR Research Intern)

Jun– Dec 2025

Project: Asymmetric On-Policy Distillation can solve POMDPs.

[Gunjan Chhablani](#) (Georgia Tech MS Student, FAIR) Next: Waymo

Jun– Dec 2024

Project: Visualizing and verifying task planning datasets with [PrediViz](#).

PROFESSIONAL ACTIVITIES

Workshop Organization

Human-aware Embodied AI Workshop ([HEAI](#)) at IROS

2025

FAIR Habitat [Challenge](#) at the CVPR [Embodied AI Workshop](#)

2023

RxR-Habitat [Challenge](#) at the CVPR [Embodied AI Workshop](#)

2021,2022,2023

Visually Grounded Interaction and Language Workshop ([ViGIL](#)) at NAACL

2021

Service

Design Advisory Board Member Center for Engineering Design & Entrepreneurship (CEDE), Gonzaga University	2021–2026
AI Application Support Program (AIASP) Mentor Guidance to historically underrepresented groups interested in graduate programs	2023

Reviewing

International Conference on Intelligent Robots and Systems (IROS)	2026
Computer Vision and Pattern Recognition (CVPR)	2024, 2025, 2026
IEEE Robotics and Automation Letters (RA-L)	2023, 2026
Winter Conference on Applications of Computer Vision (WACV)	2026
Pattern Recognition Letters	2025
Workshop on Space in Vision, Language, and Embodied AI (SpaVLE; NeurIPS)	2025
Robotics: Science and Systems (RSS)	2025
IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)	2023, 2024, 2025
International Journal of Computer Vision (IJCV)	2024
Workshop on Spatial Language Understanding and Grounded Communication for Robotics (SPLU-RoboNLP; ACL)	2021, 2023, 2024
Neural Information Processing Systems (NeurIPS)	2021, 2023
Journal of Artificial Intelligence Research (JAIR)	2023
International Conference on Robotics and Automation (ICRA)	2022, 2023
Association for Computational Linguistics Rolling Review (ARR)	2022

OPEN SOURCE

[PARTNR-Planner](#) (0.3k stars)

A repository accompanying the PARTNR benchmark for using Large planning models (LPMs) to solve robot instruction following tasks and human-robot collaboration in the Habitat simulator.

[Home Robot](#) (1.2k stars)

Contributed demo code for an agent that can perform Image Goal Navigation on real-world robotic hardware (Hello Robot Stretch; 2023). Home Robot is built and maintained by FAIR at Meta.

[AI Habitat](#) (2.8k stars)

Contributed Vision-and-Language Navigation (2020) and Instance ImageGoal Navigation (2022) to Habitat Lab. Habitat is built and maintained by FAIR at Meta.

[VLN-CE](#) (0.7k stars)

Developer and maintainer of the Vision-and-Language Navigation in Continuous Environments task, challenge, and baselines.

ABSTRACTS AND PRESENTATIONS

Beyond the Nav-Graph: Vision-and-Language Navigation in Continuous Environments **Extended Abstract.**

J. Krantz, E. Wijmans, A. Majumdar, D. Batra, S. Lee.

In: *Language in Reinforcement Learning (LaReL) Workshop at ICML*, 2020.

Probabilistic Syllabification of English Words.

J. Krantz, M. Dulin.

In: *Spokane Intercollegiate Research Conference (SIRC)*, 2017.

Best Presentation Award. [1/75=1.3%]

TEACHING

Theory of Computation (CS 321). Teaching Assistant.	Fall 2019
Artificial Intelligence (CS 427). Teaching Assistant.	Spring 2018, Spring 2019
Natural Language Processing (CS 475). Teaching Assistant.	Fall 2018