Stanford

Chelsea Finn

Assistant Professor of Computer Science and of Electrical Engineering

Bio

BIO

Chelsea Finn is an Assistant Professor in Computer Science and Electrical Engineering at Stanford University, and the William George and Ida Mary Hoover Faculty Fellow. Professor Finn's research interests lie in the ability to enable robots and other agents to develop broadly intelligent behavior through learning and interaction. Her work lies at the intersection of machine learning and robotic control, including topics such as end-to-end learning of visual perception and robotic manipulation skills, deep reinforcement learning of general skills from autonomously collected data, and meta-learning algorithms that can enable fast learning of new concepts and behaviors. Professor Finn received her Bachelors degree in Electrical Engineering and Computer Science at MIT and her PhD in Computer Science at UC Berkeley. Her research has been recognized through the ACM doctoral dissertation award, the Presidential Early Career Award for Scientists and Engineers, and the MIT Technology Review 35 under 35 list, and her work has been covered by various media outlets, including the New York Times, Wired, and Bloomberg. Throughout her career, she has sought to increase the representation of underrepresented minorities within CS and AI by developing an AI outreach camp at Berkeley for underprivileged high school students, a mentoring program for underrepresented undergraduates across three universities, and leading efforts within the WiML and Berkeley WiCSE communities of women researchers.

ACADEMIC APPOINTMENTS

- Assistant Professor, Computer Science
- Assistant Professor, Electrical Engineering
- Faculty Affiliate, Institute for Human-Centered Artificial Intelligence (HAI)
- Member, Wu Tsai Human Performance Alliance

HONORS AND AWARDS

- Presidential Early Career Award for Scientists and Engineers, United States federal government (2025)
- Research Fellowship, Alfred P. Sloan Foundation (2023)
- Early Academic Career Award in Robotics and Automation, IEEE RAS (2022)
- Young Investigator Award, Office of Naval Research (2021)
- Microsoft Faculty Fellowship, Microsoft (2020)
- ACM Doctoral Dissertation Award, ACM (2019)
- 35 Under 35 Innovator, MIT Technology Review (2018)
- C.V. Ramamoorthy Distinguished Research Award, UC Berkeley (2017)

PROGRAM AFFILIATIONS

- Symbolic Systems Program

LINKS

- Academic website: http://ai.stanford.edu/~cbfinn/
- Google Scholar: https://scholar.google.com/citations?user=vfPE6hgAAAAJ
- CV: http://ai.stanford.edu/~cbfinn/_files/cv.pdf

Teaching

COURSES

2025-26

- Deep Reinforcement Learning: CS 224R (Spr)

2024-25

- Deep Reinforcement Learning: CS 224R (Spr)

2023-24

- Deep Multi-task and Meta Learning: CS 330 (Aut)

2022-23

- Deep Multi-task and Meta Learning: CS 330 (Aut)
- Deep Reinforcement Learning: CS 224R (Spr)

STANFORD ADVISEES

Doctoral Dissertation Reader (AC)

Yihuai Gao, Joey Hejna, Garrett Thomas

Postdoctoral Faculty Sponsor

Ji Woong Kim, Yuejiang Liu

Orals Evaluator

Joey Hejna, Stephen Tian

Master's Program Advisor

Emma Beharry, Yifei Cheng, Vanessa Felix, David Fox, Ivan Ge, Devin Gupta, Nidhish Jain, Nick Jiang, James Jin, Alycia Lee, Christine Li, Nesib Muhedin, Isha Sinha, Theo Snoey, Aditya Sood, Andrew Sung, Eric Wang, Michael Yang, Sophie Zhang, Teresa Zhang, Yiwen Zhang

Doctoral Dissertation Co-Advisor (AC)

Henrik Marklund

Doctoral (Program)

Lars Ankile, Tian Gao, Sasha Khazatsky, Moo Kim, Yoonho Lee, Lucy Shi, Anikait Singh, Alexander Swerdlow, Marcel Torne Villasevil, Yonatan Urman, Jonathan Yang

Publications

PUBLICATIONS

- **SRT-H: A hierarchical framework for autonomous surgery via language-conditioned imitation learning.** *Science robotics*
  Kim, J. W., Chen, J. T., Hansen, P., Shi, L. X., Goldenberg, A., Schmidgall, S., Scheikl, P. M., Deguet, A., White, B. M., Tsai, D. R., Cha, R. J., Jopling, J., Finn, et al
  2025; 10 (104): eadt5254

- **A Tutorial on Meta-Reinforcement Learning** *FOUNDATIONS AND TRENDS IN MACHINE LEARNING*
  Beck, J., Vuorio, R., Liu, E., Xiong, Z., Zintgraf, L., Finn, C., Whiteson, S.
  2025; 18 (2-3)

- **CoT-VLA: Visual Chain-of-Thought Reasoning for Vision-Language-Action Models**
  Zhao, Q., Lu, Y., Kim, M., Fu, Z., Zhang, Z., Wu, Y., Li, Z., Ma, Q., Han, S., Finn, C., Handa, A., Lin, T., Wetzstein, et al
  IEEE COMPUTER SOC.2025: 1702-1713

- **Commonsense Reasoning for Legged Robot Adaptation with Vision-Language Models**
  Chen, A. S., Lessing, A. M., Tang, A., Chada, G., Smith, L., Levine, S., Finn, C.
  edited by Ott, C.
  IEEE.2025: 12826-12833

- **SPEEDTUNING: Speeding Up Policy Execution with Lightweight Reinforcement Learning**
  Yuan, D. D., Zhao, T. Z., Burns, K., Finn, C.
  edited by Ott, C.
  IEEE.2025: 1184-1192

- **RoboCrowd: Scaling Robot Data Collection through Crowdsourcing**
  Mirchandani, S., Yuan, D. D., Burns, K., Islam, M., Zhao, T. Z., Finn, C., Sadigh, D.
  edited by Ott, C.
  IEEE.2025: 1392-1399

- **Bayesian Embeddings for Few-Shot Open World Recognition.** *IEEE transactions on pattern analysis and machine intelligence*
  Willes, J., Harrison, J., Harakeh, A., Finn, C., Pavone, M., Waslander, S. L.
  2024; 46 (3): 1513-1529

- **A Fast and Accurate Machine Learning Autograder for the Breakout Assignment**
  Liu, E., Yuan, D., Ahmed, A., Cornwall, E., Woodrow, J., Burns, K., Nie, A., Brunskill, E., Piech, C., Assoc Computing Machinery
  ASSOC COMPUTING MACHINERY.2024: 736-742

- **Surgical Robot Transformer (SRT): Imitation Learning for Surgical Tasks**
  Kim, J., Zhao, T. Z., Schmidgall, S., Deguet, A., Kobilarov, M., Finn, C., Krieger, A.
  edited by Kroemer, O., Agrawal, P., Burgard, W.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2024

- **Evaluating Real-World Robot Manipulation Policies in Simulation**
  Li, X., Hsu, K., Gu, J., Pertsch, K., Mees, O., Walke, H., Fu, C., Lunawat, I., Sieh, I., Kirmani, S., Levine, S., Wu, J., Finn, et al
  edited by Kroemer, O., Agrawal, P., Burgard, W.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2024

- **HumanPlus: Humanoid Shadowing and Imitation from Humans**
  Fu, Z., Zhao, Q., Wu, Q., Wetzstein, G., Finn, C.
  edited by Kroemer, O., Agrawal, P., Burgard, W.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2024

- **OpenVLA: An Open-Source Vision-Language-Action Model**
  Kim, M., Pertsch, K., Karamcheti, S., Xiao, T., Balakrishna, A., Nair, S., Rafailov, R., Foster, E., Sanketi, P., Vuong, Q., Kollar, T., Burchfiel, B., Tedrake, et al
  edited by Kroemer, O., Agrawal, P., Burgard, W.

JMLR-JOURNAL MACHINE LEARNING RESEARCH.2024

- **What Makes Pre-Trained Visual Representations Successful for Robust Manipulation?**
  Burns, K., Witzel, Z., Ibn Hamid, J., Yu, T., Finn, C., Hausman, K.
  edited by Kroemer, O., Agrawal, P., Burgard, W.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2024

- **Mobile ALOHA: Learning Bimanual Mobile Manipulation using Low-Cost Whole-Body Teleoperation**
  Fu, Z., Zhao, T. Z., Finn, C.
  edited by Kroemer, O., Agrawal, P., Burgard, W.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2024

- **Robotic Control via Embodied Chain-of-Thought Reasoning**
  Zawalski, M., Chen, W., Pertsch, K., Mees, O., Finn, C., Levine, S.
  edited by Kroemer, O., Agrawal, P., Burgard, W.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2024

- **Decomposing the Generalization Gap in Imitation Learning for Visual Robotic Manipulation**
  Xie, A., Lee, L., Xiao, T., Finn, C., IEEE
  IEEE.2024: 3153-3160

- **Efficient Imitation Learning with Conservative World Models**
  Kolev, V., Rafailov, R., Hatch, K., Wu, J., Finn, C.
  edited by Abate, A., Cannon, M., Margellos, K., Papachristodoulou, A.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2024: 1776-1789

- **Robot Fine-Tuning Made Easy: Pre-Training Rewards and Policies for Autonomous Real-World Reinforcement Learning**
  Yang, J., Mark, M., Vu, B., Sharma, A., Bohg, J., Finn, C., IEEE
  IEEE.2024: 4804-4811

- **PIGEON: Predicting Image Geolocations**
  Haas, L., Skreta, M., Alberti, S., Finn, C., IEEE
  IEEE COMPUTER SOC.2024: 12893-12902

- **Open X-Embodiment: Robotic Learning Datasets and RT-X Models**
  O'Neill, A., Rehman, A., Gupta, A., Maddukuri, A., Gupta, A., Padalkar, A., Lee, A., Pooley, A., Gupta, A., Mandlekar, A., Jain, A., Tung, A., Bewley, et al
  IEEE.2024: 6892-6903

- **SERL: A Software Suite for Sample-Efficient Robotic Reinforcement Learning**
  Luo, J., Hu, Z., Xu, C., Tan, Y., Berg, J., Sharma, A., Schaal, S., Finn, C., Gupta, A., Levine, S., IEEE
  IEEE.2024: 16961-16969

- **Disentangling Length from Quality in Direct Preference Optimization**
  Park, R., Rafailov, R., Ermon, S., Finn, C.
  edited by Martins, A., Srikumar, Ku, L. W.
  ASSOC COMPUTATIONAL LINGUISTICS-ACL.2024: 4998-5017

- **Clarify: Improving Model Robustness With Natural Language Corrections**
  Lee, Y., Lam, M. S., Vasconcelos, H., Bernstein, M. S., Finn, C., ACM
  ASSOC COMPUTING MACHINERY.2024

- **NeRF in the Palm of Your Hand: Corrective Augmentation for Robotics via Novel-View Synthesis**
  Zhou, A., Kim, M., Wang, L., Florence, P., Finn, C., IEEE
  IEEE COMPUTER SOC.2023: 17907-17917

- **Supervised Pretraining Can Learn In-Context Reinforcement Learning**
  Lee, J. N., Xie, A., Pacchiano, A., Chandak, Y., Finn, C., Nachum, O., Brunskill, E.
  edited by Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S.
  NEURAL INFORMATION PROCESSING SYSTEMS (NIPS).2023

- **Contrastive Example-Based Control**
  Hatch, K., Eysenbach, B., Rafailov, R., Yu, T., Salakhutdinov, R., Levine, S., Finn, C.
  edited by Pappas, G. J., Matni, N., Morari, M.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2023

- **Robot Parkour Learning**
  Zhuang, Z., Fu, Z., Wang, J., Atkeson, C., Schwertfeger, S., Finn, C., Zhao, H.
  edited by Tan, J., Toussaint, M., Darvish, K.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2023

- **RoboCLIP: One Demonstration is Enough to Learn Robot Policies**
  Sontakke, S. A., Zhang, J., Arnold, S. M. R., Pertsch, K., Biyik, E., Sadigh, D., Finn, C., Itti, L.
  edited by Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S.
  NEURAL INFORMATION PROCESSING SYSTEMS (NIPS).2023

- **Waypoint-Based Imitation Learning for Robotic Manipulation**
  Shi, L., Sharma, A., Zhao, T. Z., Finn, C.
  edited by Tan, J., Toussaint, M., Darvish, K.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2023

- **Self-Improving Robots: End-to-End Autonomous Visuomotor Reinforcement Learning**
  Sharma, A., Ahmed, A. M., Ahmad, R., Finn, C.
  edited by Tan, J., Toussaint, M., Darvish, K.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2023

- **MOTO: Offline Pre-training to Online Fine-tuning for Model-based Robot Learning**
  Rafailov, R., Hatch, K., Kolev, V., Martin, J. D., Phielipp, M., Finn, C.
  edited by Tan, J., Toussaint, M., Darvish, K.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2023

- **Polybot: Training One Policy Across Robots While Embracing Variability**
  Yang, J., Sadigh, D., Finn, C.
  edited by Tan, J., Toussaint, M., Darvish, K.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2023

- **BridgeData V2: A Dataset for Robot Learning at Scale**
  Walke, H., Black, K., Lee, A., Kim, M., Du, M., Zheng, C., Zhao, T., Hansen-Estruch, P., Vuong, Q., He, A., Myers, V., Fang, K., Finn, et al
  edited by Tan, J., Toussaint, M., Darvish, K.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2023

- **Disentanglement via Latent Quantization**
  Hsu, K., Dorrell, W., Whittington, J. C. R., Wu, J., Finn, C.
  edited by Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S.
  NEURAL INFORMATION PROCESSING SYSTEMS (NIPS).2023

- **Cal-QL: Calibrated Offline RL Pre-Training for Efficient Online Fine-Tuning**
  Nakamoto, M., Zhai, Y., Singh, A., Mark, M., Ma, Y., Finn, C., Kumar, A., Levine, S.
  edited by Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S.
  NEURAL INFORMATION PROCESSING SYSTEMS (NIPS).2023

- **Neural Functional Transformers**
  Zhou, A., Yang, K., Jiang, Y., Burns, K., Xu, W., Sokota, S., Kolter, J., Finn, C.
  edited by Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S.
  NEURAL INFORMATION PROCESSING SYSTEMS (NIPS).2023

- **Permutation Equivariant Neural Functionals**
  Zhou, A., Yang, K., Burns, K., Cardace, A., Jiang, Y., Sokota, S., Kolter, J., Finn, C.
  edited by Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S.

NEURAL INFORMATION PROCESSING SYSTEMS (NIPS).2023

- **Direct Preference Optimization: Your Language Model is Secretly a Reward Model**
  Rafailov, R., Sharma, A., Mitchell, E., Ermon, S., Manning, C. D., Finn, C.
  edited by Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S.
  NEURAL INFORMATION PROCESSING SYSTEMS (NIPS).2023

- **Train Offline, Test Online: A Real Robot Learning Benchmark**
  Zhou, G., Dean, V., Srirama, M., Rajeswaran, A., Pari, J., Hatch, K., Jain, A., Yu, T., Abbeel, P., Pinto, L., Finn, C., Gupta, A., IEEE
  IEEE.2023: 9197-9203

- **Self-Destructing Models: Increasing the Costs of Harmful Dual Uses of Foundation Models**
  Henderson, P., Mitchell, E., Manning, C. D., Jurafsky, D., Finn, C., ACM
  ASSOC COMPUTING MACHINERY.2023: 287-296

- **Play it by Ear: Learning Skills amidst Occlusion through Audio-Visual Imitation Learning**
  Du, M., Lee, O. Y., Nair, S., Finn, C.
  edited by Hauser, K., Shell, D., Huang, S.
  RSS FOUNDATION-ROBOTICS SCIENCE & SYSTEMS FOUNDATION.2022

- **Memory-Based Model Editing at Scale**
  Mitchell, E., Lin, C., Bosselut, A., Manning, C. D., Finn, C.
  edited by Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2022

- **A State-Distribution Matching Approach to Non-Episodic Reinforcement Learning**
  Sharma, A., Ahmad, R., Finn, C.
  edited by Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2022: 19645-19657

- **Robust Policy Learning over Multiple Uncertainty Sets**
  Xie, A., Sodhani, S., Finn, C., Pineau, J., Zhang, A.
  edited by Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2022

- **How to Leverage Unlabeled Data in Offline Reinforcement Learning**
  Yu, T., Kumar, A., Chebotar, Y., Hausman, K., Finn, C., Levine, S.
  edited by Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2022

- **Improving Out-of-Distribution Robustness via Selective Augmentation**
  Yao, H., Wang, Y., Li, S., Zhang, L., Liang, W., Zou, J., Finn, C.
  edited by Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2022

- **Correct-N-Contrast: A Contrastive Approach for Improving Robustness to Spurious Correlations**
  Zhang, M., Sohoni, N. S., Zhang, H. R., Finn, C., Re, C.
  edited by Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2022

- **Training and Evaluation of Deep Policies Using Reinforcement Learning and Generative Models** *JOURNAL OF MACHINE LEARNING RESEARCH*
  Ghadirzadeh, A., Poklukar, P., Arndt, K., Finn, C., Kyrki, V., Kragic, D., Bjorkman, M.
  2022; 23

- **Bridge Data: Boosting Generalization of Robotic Skills with Cross-Domain Datasets**
  Ebert, F., Yang, Y., Schmeckpeper, K., Bucher, B., Georgakis, G., Daniilidis, K., Finn, C., Levine, S.
  edited by Hauser, K., Shell, D., Huang, S.
  RSS FOUNDATION-ROBOTICS SCIENCE & SYSTEMS FOUNDATION.2022

- **Batch Exploration With Examples for Scalable Robotic Reinforcement Learning** *IEEE ROBOTICS AND AUTOMATION LETTERS*
  Chen, A. S., Nam, H., Nair, S., Finn, C.
  2021; 6 (3): 4401–8
- **Recovery RL: Safe Reinforcement Learning With Learned Recovery Zones** *IEEE ROBOTICS AND AUTOMATION LETTERS*
  Thananjeyan, B., Balakrishna, A., Nair, S., Luo, M., Srinivasan, K., Hwang, M., Gonzalez, J. E., Ibarz, J., Finn, C., Goldberg, K.
  2021; 6 (3): 4915-4922
- **How to train your robot with deep reinforcement learning: lessons we have learned** *INTERNATIONAL JOURNAL OF ROBOTICS RESEARCH*
  Ibarz, J., Tan, J., Finn, C., Kalakrishnan, M., Pastor, P., Levine, S.
  2021; 40 (4-5): 698-721
- **WILDS: A Benchmark of in-the-Wild Distribution Shifts**
  Koh, P., Sagawa, S., Marklund, H., Xie, S., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R., Gao, I., Lee, T., David, E., Stavness, et al
  edited by Meila, M., Zhang, T.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2021
- **Bayesian Meta-Learning for Few-Shot Policy Adaptation Across Robotic Platforms**
  Ghadirzadeh, A., Chen, X., Poklukar, P., Finn, C., Bjorkman, M., Kragic, D., IEEE
  IEEE.2021: 1274-1280
- **Offline Meta-Reinforcement Learning with Advantage Weighting**
  Mitchell, E., Rafailov, R., Peng, X., Levine, S., Finn, C.
  edited by Meila, M., Zhang, T.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2021
- **Deep Reinforcement Learning amidst Continual Structured Non-Stationarity**
  Xie, A., Harrison, J., Finn, C.
  edited by Meila, M., Zhang, T.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2021
- **Learning Generalizable Robotic Reward Functions from "In-The-Wild" Human Videos**
  Chen, A. S., Nair, S., Finn, C.
  edited by Shell, D. A., Toussaint, M., Hsieh, M. A.
  RSS FOUNDATION-ROBOTICS SCIENCE & SYSTEMS FOUNDATION.2021
- **Decoupling Exploration and Exploitation for Meta-Reinforcement Learning without Sacrifices**
  Liu, E., Raghunathan, A., Liang, P., Finn, C.
  edited by Meila, M., Zhang, T.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2021
- **Just Train Twice: Improving Group Robustness without Training Group Information**
  Liu, E., Haghgoo, B., Chen, A. S., Raghunathan, A., Koh, P., Sagawa, S., Liang, P., Finn, C.
  edited by Meila, M., Zhang, T.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2021
- **Catformer: Designing Stable Transformers via Sensitivity Analysis**
  Davis, J., Gu, A., Choromanski, K., Dao, T., Re, C., Finn, C., Liang, P.
  edited by Meila, M., Zhang, T.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2021
- **Actionable Models: Unsupervised Offline Reinforcement Learning of Robotic Skills**
  Chebotar, Y., Hausman, K., Lu, Y., Xiao, T., Kalashnikov, D., Varley, J., Irpan, A., Eysenbach, B., Julian, R., Finn, C., Levine, S.
  edited by Meila, M., Zhang, T.
  JMLR-JOURNAL MACHINE LEARNING RESEARCH.2021
- **Greedy Hierarchical Variational Autoencoders for Large-Scale Video Prediction**
  Wu, B., Nair, S., Martin-Martin, R., Li Fei-Fei, Finn, C., IEEE COMP SOC

IEEE COMPUTER SOC.2021: 2318-2328

- **Scalable Multi-Task Imitation Learning with Autonomous Improvement**
  Singh, A., Jang, E., Irpan, A., Kappler, D., Dalal, M., Levinev, S., Khansari, M., Finn, C., IEEE
  IEEE.2020: 2167-2173
- **OmniTact: A Multi-Directional High-Resolution Touch Sensor**
  Padmanabha, A., Ebert, F., Tian, S., Calandra, R., Finn, C., Levine, S., IEEE
  IEEE.2020: 618-624
- **Meta-Inverse Reinforcement Learning with Probabilistic Context Variables**
  Yu, L., Yu, T., Finn, C., Ermon, S.
  edited by Wallach, H., Larochelle, H., Beygelzimer, A., d'Alche-Buc, F., Fox, E., Garnett, R.
  NEURAL INFORMATION PROCESSING SYSTEMS (NIPS).2019
- **Unsupervised Curricula for Visual Meta-Reinforcement Learning**
  Jabri, A., Hsu, K., Eysenbach, B., Gupta, A., Levine, S., Finn, C.
  edited by Wallach, H., Larochelle, H., Beygelzimer, A., d'Alche-Buc, F., Fox, E., Garnett, R.
  NEURAL INFORMATION PROCESSING SYSTEMS (NIPS).2019
- **Unsupervised Visuomotor Control through Distributional Planning Networks**
  Yu, T., Shevchuk, G., Sadigh, D., Finn, C.
  edited by Bicchi, A., KressGazit, H., Hutchinson, S.
  MIT PRESS.2019
- **One-Shot Composition of Vision-Based Skills from Demonstration**
  Yu, T., Abbeel, P., Levine, S., Finn, C., IEEE
  IEEE.2019: 2643–50