Chelsea Finn

353 Jane Stanford Way
Stanford, CA 94305
cbfinn@cs.stanford.edu
`http://ai.stanford.edu/~cbfinn`

Current Positions

Stanford University, Computer Science Department and Electrical Engineering Department, Assistant Professor — 2019 – present

Google, Inc., Brain Team, Research Scientist — 2018 – present

Education

University of California, Berkeley, PhD — 2014 – 2018
Thesis: *"Learning to Learn with Gradients"*.
Department of Electrical Engineering and Computer Science

Massachusetts Institute of Technology, Bachelor of Science — 2010 – 2014
Electrical Engineering and Computer Science

Honors and Awards

Okawa Foundation Research Grant — 2023
Awarded to 7 US faculty members in the fields of information and telecommunications.

Stanford Tau Beta Pi Teaching Honor Role — 2023
Awarded to 12 Stanford faculty members who promote excellent teaching in the School of Engineering.

NSF CAREER Award — 2023
Supports early-career faculty who have the potential to serve as academic role models in research and education.

Alfred P. Sloan Research Fellowship — 2023
Awarded to 22 early-career scholars in computer science in the US and Canada

IEEE RAS Early Academic Career Award in Robotics and Automation — 2022
Awarded to two early-career academics for major impact on robotics & automation
For pioneering contributions in deep robotic learning, and their application to vision-based robotic manipulation

ONR Young Investigator Award — 2021
Awarded to 38 early-career faculty

Samsung AI Researcher of the Year — 2020
Awarded to five early-career researchers in AI worldwide

CoRL Best Paper Award — 2020
For the paper *"Learning Latent Representations to Influence Multi-Agent Interaction"*

Intel Rising Star Faculty Award — 2020
Awarded to ten early-career professors worldwide

Microsoft Faculty Fellowship Award — 2020
Awarded to five early-career professors in North America

ACM Doctoral Dissertation Award — 2019
Awarded to one doctoral dissertation in computer science and engineering, worldwide

MIT TR35 Innovator Award 2018
Awarded to 35 innovators under 35 worldwide

Rising Stars in EECS 2017
Awarded to 70 EECS graduate and postdoctoral women

C.V. Ramamoorthy Distinguished Research Award 2017
For outstanding contributions to a new research area in computer science and engineering

ICRA Best Cognitive Robotics Paper Finalist 2017
For the paper *"Deep Visual Foresight for Planning Robot Motion"*

Tong Leong Lim Pre-Doctoral Prize 2016
For achieving the highest distinction in the pre-doctoral examination

Computing Community Consortium (CCC) Blue Sky Ideas Award 2015
For the paper *"End-to-End Training of Deep Visuomotor Policies"*

National Science Foundation Graduate Research Fellowship 2015-2018

National Defense Science and Engineering Graduate Fellowship *(declined)* 2015

IEEE-HKN Alton B. Zerby and Carl T. Koerner Outstanding Student Award 2015
Awarded annually to one undergraduate student in the United States

SanDisk Fellowship 2015

UC Berkeley EECS Department Fellowship 2014

MIT SuperUROP Outstanding Research Presentation Award 2014
"Real-time Text Detection in Human Scenes"

Teaching

Instructor

Stanford CS224R: Deep Reinforcement Learning Spring 2023

Stanford CS330: Deep Multi-Task and Meta Learning Fall 2019, Fall 2020, Fall 2021, Fall 2022

Stanford CS221: Artificial Intelligence: Principles and Techniques Spring 2020, Spring 2021

Berkeley CS294-112: Deep Reinforcement Learning Spring 2017

Teaching Assistant

Berkeley CS188 Introduction to Artificial Intelligence Spring 2015

MIT 6.008 Introduction to Inference Spring 2014

MIT 6.141 Robotics: Science and Systems I Spring 2013

MIT 6.02 Digital Communication Systems Spring 2012

Invited Tutorials

Tutorial on Meta-Learning for Bridging Labeled and Unlabeled Data in Biomedicine Fall 2021
in Intelligent Systems for Molecular Biology/European Conference on Computational Biology.

Meta Reinforcement Learning Summer 2020
at the CIFAR Deep Learning & Reinforcement Learning Summer School.

Tutorial on Meta-Learning: from Few-Shot Learning to Rapid Reinforcement Learning Summer 2019
at the International Conference on Machine Learning (ICML).
at the Conference on Computer Vision and Pattern Recognition (CVPR).

Tutorial on Deep Visuomotor Learning Summer 2019
in Computational Vision Summer School, Freudenstadt.

Tutorial on Deep Visuomotor Learning Summer 2018
in International Computer Vision Summer School, Sicily.

Tutorial on Deep Reinforcement Learning, Decision Making, and Control Summer 2017
at the International Conference on Machine Learning (ICML).

Selected Invited Talks

Detecting and Adapting to Distribution Shift.
ICML Workshop on Machine Learning for Astrophysics. June 2023.

History and Future of Artificial Intelligence and Computer Vision.
IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Plenary Panel. June 2023.

Benchmarking, Detecting, and Adapting to Distribution Shift.
CVPR Workshop on Autonomous Driving. June 2023.

Why Robots Should Learn in the Real World.
ICRA Conference Keynote. May 2023.

How to Generalize Your Robot with Diverse Offline Data.
ICRA Workshop on Learning from Diverse, Offline Data (LDOD). May 2023. *CVPR Workshop on Visual Pre-Training for Robotics*. June 2023.

Neural Networks Make Stuff Up. What Should We Do About It?.
Harvard Machine Learning Foundations Seminar. February 2023. *Women in Data Science Worldwide Virtual Conference*. September 2023.

How to Train Your Robot from Demonstrations.
University of Utah Robotics Seminar. February 2023.
Wallenberg AI, Autonomous Systems and Software Program. March 2023.
ICRA Workshop on Effective Representations, Abstractions, and Priors for Robot Learning. May 2023.

Robots Need Initiative and Autonomy.
CoRL Workshop on Learning to Adapt and Improve in the Real World. December 2022.

Meta-Reinforcement Learning: Algorithms and Applications.
NeurIPS Workshop on Meta-Learning. December 2022.

Extrapolation via Interpolation.
NeurIPS Workshop on Interpolation and Beyond. December 2022.

How to Genearlize Your Robot.
CMU Robotics Institute Seminar. October 2022.
CoRL Workshop on Pre-Training Robot Learning. December 2022.

Practical Robustness Despite Underspecification.
ECCV Robust Vision Challenge Workshop. October 2022.

Ingredients for Developing Intelligent Robots.
Cognitive Computational Neuroscience (CCN) Conference Keynote. August 2022.

Robust Deep Networks through Invariance and Adaptation.
Korean Conference on Computer Vision (KCCV) Keynote. August 2022.

Distribution Shift as Underspecification: And What We Might Do About It.
Simons Institute Deep Learning Theory Workshop. August 2022.

Meta-Learning for Education.
AutoML Conference Keynote. July 2022.

Adapting Neural Networks to Distribution Shift with Minimal Assumptions.
ICML Workshop on Updatable Machine Learning. July 2022.

Robots that Learn to Be Safe.
ICML Workshop on Safe Learning for Autonomous Driving. July 2022.

Benchmarks and Methods Beyond I.I.D. Data.
ICML Shift Happens Workshop. July 2022.

Revisiting the Ins and Outs of Imitation Learning for Robotic Manipulation.
RSS Workshop on Overlooked Aspects of Imitation Learning. June 2022.

Flexible Machine Learning for Mitigating Distribution Shift.
Machine Learning at Berkeley. April 2022.
Apple, Inc.. May 2022
Princeton AI Club. May 2022

Robots Need to Reduce, Reuse, and Recycle.
Allen Institute for AI, Embodied AI Lecture Series. April 2022.

Scaling Data-Driven Robot Learning.
NVIDIA GTC Conference. March 2022.

Robustness through the Lens of Invariance.
NeurIPS Workshop on Distribution Shifts: Connecting Methods and Applications. December 2021.

Compositional Task Generalization in Vision-Based Robotic Manipulation.
ICCV BEHAVIOR Workshop. October 2021.

Inductive Biases for Robust Deep Learning.
ICCV Workshop on Visual Inductive Priors for Data-Efficient Deep Learning. October 2021.

Understanding and Controlling Transfer in Multi-Task Learning.
ICCV Workshop on Multi-Task Learning in Computer Vision. October 2021

Meta-Learning Unsupervised Adaptation Strategies.
IJCAI Workshop on Continual Semi-Supervised Learning. August 2021.

Few-Shot Learning in the Real World: Meta-Learning for Giving Feedback to Students.
Beneficial AI Seminar, UC Berkeley Center for Human-Compatible AI (CHAI). August 2021.
ICML Workshop on Self-Supervised Learning for Reasoning and Perception. July 2021.
ACL Workshop on Meta Learning and Its Applications to Natural Language Processing. August 2021.

Blending Learning and Planning for Flexible, Generalizable Robotic Manipulation.
RSS Workshop on Integrating Planning and Learning. July 2021.

Data Scalability in Reinforcement Learning.
ICML Workshop on Unsupervised Reinforcement Learning. July 2021.

Learning Transferable Exploration Strategies via Meta Reinforcement Learning.
ICML Workshop on Theory and Foundation of Continual Learning. July 2021.

Broad Robot Generalization by Reusing Broad Data.
ETH Zurich Distinguished Seminar in Robotics, Systems and Control. November 2021.
IROS Conference Keynote. October 2021.
CVPR Embodied AI Workshop. June 2021

Few Shot Learning in the Real World: Meta-Learning for Giving Feedback to Students.
CVPR Workshop on Learning with Limited and Imperfect Data. June 2021

Learning Exploration Strategies with Meta Reinforcement Learning.
ICRA Workshop on Learning to Learn for Robotics. June 2021

Robots that Anticipate and Adapt to Change.
ICRA Workshop on Perception and Action in Dynamic Environments. June 2021

Broad Data for Broad Robot Generalization.
Technical University of Munich (TUM) AI Guest Lecture Series. April 2021.

Principles for Tackling Distribution Shift: Pessimism, Adaptation, and Anticipation.
DeepMind/Ellis Seminar, Computational Statistics and Machine Learning Centre, University College London. February 2021.
Vector Institute & The Fields Institute for Research in Mathematical Sciences Seminar, University of Toronto. February 2021.
Google Workshop on the Conceptual Understanding of Deep Learning. May 2021.
CVPR Workshop on Bridging the Gap Between Computational Photography and Visual Recognition. June 2021

Reinforcement Learning for Real Robots.
AAAI New Faculty Highlight. January 2021.
NeurIPS Workshop on Real World Reinforcement Learning. December 2020.

Underfitting and Uncertainty in Self-Supervised Predictive Models.
NeurIPS Workshop on Self-Supervised Learning – Theory and Practice. December 2020.
NeurIPS Workshop on Self-Supervised Learning for Speech and Audio Processing. December 2020.

Data Scalability for Robot Learning.
CMU Robotics Institute Seminar. November 2020.
RSS Self-Supervised Learning Workshop. July 2020.

Meta-Learning: From Few-Shot Adaptation to Uncovering Symmetries.
Samsung AI Forum Keynote. November 2020.

Meta-Learning for Robustness to our Changing World.
BayLearn: Bay Area Machine Learning Symposium Keynote. October 2020.

How Not to Create a Robot's Mind.
Stanford Human-Centered Artificial Intelligence Conference Keynote. October 2020.

From Neuroscience to Artificially Intelligent Systems (NAISys) Conference. November 2020.

Learning Exploration Strategies with Meta-Reinforcement Learning.
Simons Institute Workshop on Deep Reinforcement Learning. September 2020.

Learning Structured Exploration Strategies via Language and Simple Supervision.
ECCV Workshop on Embodied Vision, Actions & Language. August 2020.

How Can Robots Get the Most out of People?
ICML Workshop on Human-in-the-Loop Learning. July 2020.

Beyond the Training Distribution: Embodiment, Adaptation, and Symmetry.
MIT Embodied Intelligence Seminar. June 2020.

Extrapolation via Adaptation.
L4DC Conference Keynote. June 2020.
CVPR Workshop on Continual Learning in Computer Vision. June 2020.

Meta-Learning Beyond Few-Shot Classification.
CVPR Workshop on Deep Declarative Networks. June 2020.

Meta-Learning Symmetries and Distributions.
CVPR Workshop on Compositionality. June 2020.

Peculiar Optimization and Regularization Challenges in Multi-Task Learning and Meta-Learning.
Workshop on New Directions in Optimization, Statistics and Machine Learning, The Institute for Advanced Study. April 2020
CVPR Workshop on Efficient Deep Learning. June 2020.

Meta-Learning and Memorization.
CIFAR Learning in Machines and Brains Program Meeting. December 2019
NeurIPS Workshop on Bayesian Deep Learning. December 2019

The Next Generation of Robot Learning.
Stanford Robotics Seminar. December 2019.

Flexible Neural Networks and the Frontiers of Meta-Learning.
Simons Institute Workshop on Emerging Challenges in Deep Learning. August 2019.

Reinforcement Learning for Robots.
The Multi-Disciplinary Conference on Reinforcement Learning and Decision Making (RLDM). July 2019.

Learning to Adapt to Dynamic, Real-World Environments.
RSS Workshop on Simulation to Real-World Transfer. June 2019.

Learning Compound Tasks through Interaction and Observation.
RSS Workshop on Task-Informed Graping. June 2019.

Learning Models of the World and its Intentions.
CVPR Workshop on Vision Meets Cognition. June 2019.

A Practical View on Generalization and Autonomy in the Real World.
ICML Workshop on Understanding and Improving Generalization in Deep Learning. June 2019.
ICML Workshop on AI for Autonomous Driving. June 2019.

Complexity without Losing Generality: The Role of Supervision and Composition.
ICML Workshop on Generative Modeling and Model-Based Reasoning for Robotics and AI. June 2019.

Agents that Set Measurable Goals for Themselves.
ICML Workshop on Self-Supervised Learning. June 2019.

Meta-Learning: Challenges and Frontiers.
ICLR Workshop on Learning from Limited Data. May 2019.
CIFAR Learning in Machines and Brains Program Meeting. May 2019.
ICML Workshop on Multi-Task and Adaptive Learning. June 2019.

What can we learn from unlabeled interaction?
ICLR Workshop on Task-Agnostic Reinforcement Learning. May 2019

Versatility and Self-Supervision in Deep Robotic Learning.
University of Pennsylvania, GRASP Lab. May 2019

Meta-Learning Deep Networks. *Re-work Deep Learning Summit, San Francisco*. January 2019.

Meta-Learning across Time. *NeurIPS Workshop on Continual Learning*. December 2018.

An agent that can do many things (by modeling the world). *NeurIPS Workshop on Modeling the Physical World*. December 2018.

Learning Generalizable Models through Unsupervised Interaction. *NeurIPS Workshop on Modeling and Decision-Making in the Spatiotemporal Domain*. December 2018.

Model-Based Deep Reinforcement Learning Tutorial. *CIFAR Learning in Machines and Brains Program Meeting*. December 2018

Building Versatile Agents through Unsupervised Interaction.
Stanford Minds, Brains, and Computation (MBC) Colloquium. November 2018.
Stanford DAWN Seminar. November 2018

Robots that Excel in Diverse Environments. *Bay Area Robotics Symposium*. November 2018

Building Unsupervised, Versatile Agents with Meta Learning.
University of Washington Robotics Colloquium. October 2018.
Allen Institute for Artificial Intelligence. October 2018.
OpenAI. November 2018.

Meta-Learning Frontiers: Universal, Uncertain, and Unsupervised. *Google DeepMind*. July 2018.

Properties of Good Meta-Learning Algorithms (And How to Achieve Them). *ICML AutoML Workshop*. July 2018.

Meta-Learning for Goal Inference, Imitation, and Planning. *RSS Workshop on Learning from Demonstrations for High-Level Tasks*. June 2018.

Efficiency through Learning to Learn. *Clarifai*. April 2018.

Generalization and Self-Supervision in Deep Robotic Learning.
Toyota Technical Institute in Chicago (TTIC). February 2018.
Stanford University. March 2018.
MIT. March 2018.

Google. April 2018.

Learning Versatile Behavior and Reusable Models through Real-World Interaction. *Caltech Young Investigator Lecture*. February 2018.

Model-Agnostic Meta-Learning: Universality, Inductive Bias, and Weak Supervision. *NIPS Workshop on Meta-Learning*. December 2017.

Deep Predictive Learning for Acquiring Vision-Based Skills. *ICML Workshop on Reinforcement Learning*. August 2017.

Learning Representations for Versatile Behavior. *RSS Workshop on New Fronteirs for Deep Learning in Robotics*. July 2017.

Learning through Interaction: Generalization in Robot Reinforcement Learning.
Symposium on Robot Learning, Berkeley, CA. May 2017.
MIT. April 2017.
Stanford University. March 2017.

End-to-End Deep Robotic Learning. *Re-work Deep Learning Summit, San Francisco*. January 2017.

Guided Cost Learning and Connections to Generative Adversarial Modeling. *NIPS Deep Learning Symposium*. December 2016.

Large Scale Self-Supervised Robotic Learning.
NIPS Deep Reinforcement Learning Workshop. December 2016.
NIPS Neurorobotics Workshop. December 2016.

Robotic Visuomotor Learning. *3DV Tutorial: Workshop on Understanding 3D and Visuomotor Learning*. October 2016.

Learning Visuomotor Skills.
OpenAI. March 2016.
Google DeepMind. May 2016.

Robotic Visuomotor Learning. *Redwood Center for Theoretical Neuroscience*. November 2015.

End-to-End Training of Deep Visuomotor Policies. *Google, Inc.*. March 2015.

Peer-Reviewed Publications (Journals and Conferences)

[157] Lucy Xiaoyang Shi, Archit Sharma, Tony Z. Zhao, **Chelsea Finn**. Waypoint-Based Imitation Learning for Robotic Manipulation. *Conference on Robot Learning (CoRL)*. 2023.

[156] Rafael Rafailov, Kyle Beltran Hatch, Victor Kolev, John D Martin, Mariano Phielipp, **Chelsea Finn**. MOTO: Offline Pre-training to Online Fine-tuning for Model-based Robot Learning. *Conference on Robot Learning (CoRL)*. 2023.

[155] Archit Sharma, Ahmed M Ahmed, Rehaan Ahmad, **Chelsea Finn**. Self-Improving Robots: End-to-End Autonomous Visuomotor Reinforcement Learning. *Conference on Robot Learning (CoRL)*. 2023.

[154] Ziwen Zhuang, Zipeng Fu, Jianren Wang, Christopher G Atkeson, Soren Schwertfeger, **Chelsea Finn**, Hang Zhao. Robot Parkour Learning. *Conference on Robot Learning (CoRL)*. 2023.

[153] Jonathan Heewon Yang, Dorsa Sadigh, **Chelsea Finn**. Polybot: Training One Policy Across Robots While Embracing Variability. *Conference on Robot Learning (CoRL)*. 2023.

[152] Homer Walke, Kevin Black, Abraham Lee, Moo Jin Kim, Max Du, Chongyi Zheng, Tony Zhao, Philippe Hansen-Estruch, Quan Vuong, Andre He, Vivek Myers, Kuan Fang, **Chelsea Finn**, Sergey Levine. BridgeData V2: A Dataset for Robot Learning at Scale. *Conference on Robot Learning (CoRL)*. 2023.

[151] Austin Stone, Ted Xiao, Yao Lu, Keerthana Gopalakrishnan, Kuang-Huei Lee, quan vuong, Paul Wohlhart, Sean Kirmani, Brianna Zitkovich, Fei Xia, **Chelsea Finn**, Karol Hausman. Open-World Object Manipulation using Pre-Trained Vision-Language Models. *Conference on Robot Learning (CoRL)*. 2023.

[150] Yevgen Chebotar, quan vuong, Karol Hausman, Fei Xia, Yao Lu, Alex Irpan, Aviral Kumar, Tianhe Yu, Alexander Herzog, Karl Pertsch, Keerthana Gopalakrishnan, Julian Ibarz, Ofir Nachum, Grecia Salazar, Huong T Tran, Jodilyn Peralta, Clayton Tan, Deeksha Manjunath, Jaspiar Singh, Brianna Zitkovich, Tomas Jackson, Kanishka Rao, **Chelsea Finn**, Sergey Levine. Q-Transformer: Scalable Offline Reinforcement Learning via Autoregressive Q-Functions. *Conference on Robot Learning (CoRL)*. 2023.

[149] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, **Chelsea Finn**, Pete Florence, Chuyuan Fu, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alexander Herzog, Jasmine Hsu, Brian Ichter, Alex Irpan, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Lisa Lee, Tsang-Wei Edward Lee, Sergey Levine, Yao Lu, Henryk Michalewski, Igor Mordatch, Karl Pertsch, Kanishka Rao, Krista Reymann, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Pierre Sermanet, Jaspiar Singh, Anikait Singh, Radu Soricut, Huong Tran, Vincent Vanhoucke, Quan Vuong, Ayzaan Wahid, Stefan Welker, Paul Wohlhart, Jialin Wu, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, Brianna Zitkovich. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control. *Conference on Robot Learning (CoRL)*. 2023.

[148] Peter Henderson, Eric Mitchell, Christopher Manning, Dan Jurafsky, **Chelsea Finn**. Self-Destructing Models: Increasing the Costs of Harmful Dual Uses in Foundation Models. *AAAI/ACM Conference on AI, Ethics, and Society (AIES)*. 2023.

[147] Evan Zheran Liu, Sahaana Suri, Tong Mu, Allan Zhou, **Chelsea Finn**. Simple Embodied Language Learning as a Byproduct of Meta-Reinforcement Learning. *International Conference on Machine Learning (ICML)*. 2023.

[146] Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D Manning, **Chelsea Finn**. DetectGPT: Zero-Shot Machine-Generated Text Detection using Probability Curvature. *International Conference on Machine Learning (ICML)*. 2023.

[145] Tony Zhao, Vikash Kumar, Sergey Levine, **Chelsea Finn**. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. *Robotics: Science and Systems (RSS)*. 2023.

[144] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, **Chelsea Finn**, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil J Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Kuang-Huei Lee, Sergey Levine, Yao Lu, Utsav Malla, Deeksha Manjunath, Igor Mordatch, Ofir Nachum, Carolina Parada, Jodilyn Peralta, Emily Perez, Karl Pertsch, Jornell Quiambao, Kanishka Rao, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Kevin Sayed, Jaspiar Singh, Sumedh Sontakke, Austin Stone, Clayton Tan, Huong Tran, Vincent Vanhoucke, Steve Vega, Quan Vuong, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, Brianna Zitkovich. RT-1: Robotics

Transformer for Real-World Control at Scale. *Robotics: Science and Systems (RSS)*. 2023.

[143] Siddharth Karamcheti, Suraj Nair, Annie Chen, Thomas Kollar, **Chelsea Finn**, Dorsa Sadigh, Percy Liang. Language-Driven Representation Learning for Robotics. *Robotics: Science and Systems (RSS)*. 2023.

[142] Aviral Kumar, Anikait Singh, Frederik Ebert, Mitsuhiko Nakamoto, Yanlai Yang, **Chelsea Finn**, Sergey Levine. Pre-Training for Robots: Offline RL Enables Learning New Tasks in a Handful of Trials. *Robotics: Science and Systems (RSS)*. 2023.

[141] Maximilian Du, Suraj Nair, Dorsa Sadigh, **Chelsea Finn**. Behavior Retrieval: Few-Shot Imitation Learning by Querying Unlabeled Datasets. *Robotics: Science and Systems (RSS)*. 2023.

[140] Gaoyue Zhou*, Victoria Dean*, Mohan Kumar Srirama, Aravind Rajeswaran, Jyothish Pari, Kyle Hatch, Aryan Jain, Tianhe Yu, Pieter Abbeel, Lerrel Pinto, **Chelsea Finn**, Abhinav Gupta. Train Offline, Test Online: A Real Robot Learning Benchmark. *International Conference on Robotics and Automation (ICRA)*. 2023.

[139] Allan Zhou*, Moo Jin Kim*, Lirui Wang, Pete Florence, **Chelsea Finn**. NeRF in the Palm of Your Hand: Corrective Augmentation for Robotics via Novel-View Synthesis. *Computer Vision and Pattern Recognition (CVPR)*. 2023.

[138] Yoonho Lee, Annie S Chen, Fahim Tajwar, Ananya Kumar, Huaxiu Yao, Percy Liang, **Chelsea Finn**. Surgical Fine-Tuning Improves Adaptation to Distribution Shifts. *International Conference on Learning Representations (ICLR)*. 2023.

[137] Yoonho Lee, Huaxiu Yao, **Chelsea Finn**. Diversify and Disambiguate: Out-of-Distribution Robustness via Disagreement. *International Conference on Learning Representations (ICLR)*. 2023.

[136] Stephen Tian, **Chelsea Finn**, Jiajun Wu. A Control-Centric Benchmark for Video Prediction. *International Conference on Learning Representations (ICLR)*. 2023.

[135] Amrith Setlur, Don Dennis, Benjamin Eysenbach, Aditi Raghunathan, **Chelsea Finn**, Virginia Smith, Sergey Levine. Bitrate-Constrained DRO: Beyond Worst Case Robustness To Unknown Group Shifts. *International Conference on Learning Representations (ICLR)*. 2023.

[134] Gaoyue Zhou, Victoria Dean, Mohan Kumar Srirama, Aravind Rajeswaran, Jyothish Pari, Kyle Beltran Hatch, Aryan Jain, Tianhe Yu, Pieter Abbeel, Lerrel Pinto, **Chelsea Finn**, Abhinav Gupta. Train Offline, Test Online: A Real Robot Learning Benchmark. *International Conference on Robotics and Automation (ICRA)*. 2023.

[133] Eric A. Mitchell, Joseph Noh, Siyan Li, Will Armstrong, Ananth Agarwal, Patrick Liu, **Chelsea Finn**, Christopher D. Manning. Enhancing Self-Consistency and Performance of Pretrained Language Models with NLI. *Empirical Methods in Natural Language Processing (EMNLP)*. 2022.

[132] Huaxiu Yao*, Caroline Choi*, Yoonho Lee, Pang Wei Koh, **Chelsea Finn**. Wild-Time: A Benchmark of in-the-Wild Distribution Shift over Time. *Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*. 2022.

[131] Annie Xie, Fahim Tajwar, Archit Sharma, **Chelsea Finn**. When to Ask for Help: Proactive Interventions in Autonomous Reinforcement Learning. *Neural Information Processing Systems (NeurIPS)*. 2022.

[130] Evan Zheran Liu*, Moritz Pascal Stephan*, Allen Nie, Christopher J Piech, Emma Brunskill, **Chelsea Finn**. Giving Feedback on Interactive Student Programs with Meta-Exploration. *Neural*

Information Processing Systems (NeurIPS). 2022.

[129] Marvin Mengxin Zhang, Sergey Levine, **Chelsea Finn**. MEMO: Test Time Robustness via Adaptation and Augmentation. *Neural Information Processing Systems (NeurIPS)*. 2022.

[128] Huaxiu Yao, Yiping Wang, Linjun Zhang, James Zou, **Chelsea Finn**. C-Mixup: Improving Generalization in Regression. *Neural Information Processing Systems (NeurIPS)*. 2022.

[127] Yiding Jiang, Evan Zheran Liu, Benjamin Eysenbach, J Zico Kolter, **Chelsea Finn**. Learning Options via Compression. *Neural Information Processing Systems (NeurIPS)*. 2022.

[126] Annie S Chen, Archit Sharma, Sergey Levine, **Chelsea Finn**. You Only Live Once: Single-Life Reinforcement Learning via Learned Reward Shaping. *Neural Information Processing Systems (NeurIPS)*. 2022.

[125] Xi Chen, Ali Ghadirzadeh, Tianhe Yu, Yuan Gao, Jianhao Wang, Wenzhe Li, Liang Bin, **Chelsea Finn**, Chongjie Zhang. Latent-Variable Advantage-Weighted Policy Optimization for Offline Reinforcement Learning. *Neural Information Processing Systems (NeurIPS)*. 2022.

[124] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, **Chelsea Finn**, Abhinav Gupta. R3M: A Universal Visual Representation for Robot Manipulation. *Conference on Robot Learning (CoRL)*. 2022.

[123] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, **Chelsea Finn**, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario Jauregui Ruano, Kyle Jeffrey, Sally Jesmonth, Nikhil J Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee, Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander Toshev, Vincent Vanhoucke, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Mengyuan Yan, Andy Zeng. Do As I Can, Not As I Say: Grounding Language in Robotic Affordances. *Conference on Robot Learning (CoRL)*. 2022.

[122] Kaylee Burns, Tianhe Yu, **Chelsea Finn**, Karol Hausman. Offline Reinforcement Learning at Multiple Frequencies. *Conference on Robot Learning (CoRL)*. 2022.

[121] John Willes, James Harrison, Ali Harakeh, **Chelsea Finn**, Marco Pavone, Steven Waslander. Bayesian Embeddings for Few-Shot Open World Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*. 2022.

[120] Ali Ghadirzadeh, Petra Poklukar, Karol Arndt, **Chelsea Finn**, Ville Kyrki, Danica Kragic, Marten Bjorkman. Training and Evaluation of Deep Policies Using Reinforcement Learning and Generative Models. *Journal of Machine Learning Research (JMLR)*. 2022.

[119] Annie Xie, **Chelsea Finn**. Lifelong Robotic Reinforcement Learning by Retaining Experiences. *Conference on Lifelong Learning Agents (CoLLAs)*. 2022.

[118] Huaxiu Yao, Yu Wang, Sai Li, Linjun Zhang, Weixin Liang, James Zou, **Chelsea Finn**. Improving Out-of-Distribution Robustness via Selective Augmentation. *International Conference on Machine Learning (ICML)*. 2022.

[117] Tianhe Yu*, Aviral Kumar*, Yevgen Chebotar, Karol Hausman, **Chelsea Finn**, Sergey Levine. How to Leverage Unlabeled Data in Offline Reinforcement Learning? *International Conference on Machine Learning (ICML)*. 2022.

[116] Eric Mitchell, Charles Lin, Antoine Bosselut, Christopher Manning, **Chelsea Finn**. Memory-Based Model Editing at Scale. *International Conference on Machine Learning (ICML)*. 2022.

[115] Annie Xie, Shagun Sodhani, **Chelsea Finn**, Joelle Pineau, Amy Zhang. Robust Policy Learning over Multiple Uncertainty Sets. *International Conference on Machine Learning (ICML)*. 2022.

[114] Archit Sharma*, Rehaan Ahmad*, **Chelsea Finn**. A State-Distribution Matching Approach to Non-Episodic Reinforcement Learning. *International Conference on Machine Learning (ICML)*. 2022.

[113] Michael Zhang, Nimit Sohoni, Hongyang Zhang, **Chelsea Finn**, Christopher Re. Correct-N-Contrast: a Contrastive Approach for Improving Robustness to Spurious Correlations. *International Conference on Machine Learning (ICML)*. 2022.

[112] Maximilian Du, Olivia Y Lee, Suraj Nair, **Chelsea Finn**. Play it by Ear: Learning Skills amidst Occlusion through Audio-Visual Imitation Learning. *Robotics: Science and Systems (RSS)*. 2022.

[111] Frederik Ebert, Yanlai Yang, Karl Schmeckpeper, Bernadette Bucher, Georgios Georgakis, Kostas Daniilidis, **Chelsea Finn**, Sergey Levine. Bridge Data: Boosting Generalization of Robotic Skills with Cross-Domain Datasets. *Robotics: Science and Systems (RSS)*. 2022.

[110] Kyle Hsu*, Moo Jin Kim*, Rafael Rafailov, Jiajun Wu, **Chelsea Finn**. Vision-Based Manipulators Need to Also See from Their Hands. *International Conference on Learning Representations (ICLR)*. 2022.

[109] Eric Mitchell, Charles Lin, Antoine Bosselut, **Chelsea Finn**, Christopher Manning. Fast Model Editing at Scale. *International Conference on Learning Representations (ICLR)*. 2022.

[108] Huaxiu Yao, Linjun Zhang, **Chelsea Finn**. Meta-Learning with Fewer Tasks through Task Interpolation. *International Conference on Learning Representations (ICLR)*. 2022.

[107] Archit Sharma, Kelvin Xu, Nikhil Sardana, Abhishek Gupta, Karol Hausman, Sergey Levine, **Chelsea Finn**. Autonomous Reinforcement Learning: Formalism and Benchmarking. *International Conference on Learning Representations (ICLR)*. 2022.

[106] Allan Zhou*, Fahim Tajwar*, Alexander Robey, Tom Knowles, George J. Pappas, Hamed Hassani, **Chelsea Finn**. Do deep networks transfer invariances across classes? *International Conference on Learning Representations (ICLR)*. 2022.

[105] Shiori Sagawa, Pang Wei Koh, Tony Lee, Irena Gao, Sang Michael Xie, Kendrick Shen, Ananya Kumar, Weihua Hu, Michihiro Yasunaga, Henrik Marklund, Sara Beery, Etienne David, Ian Stavness, Wei Guo, Jure Leskovec, Kate Saenko, Tatsunori Hashimoto, Sergey Levine, **Chelsea Finn**, Percy Liang. Extending the WILDS Benchmark for Unsupervised Adaptation. *International Conference on Learning Representations (ICLR)*. 2022.

[104] Glen Berseth, Zhiwei Zhang, Grace Zhang, **Chelsea Finn**, Sergey Levine. CoMPS: Continual Meta Policy Search. *International Conference on Learning Representations (ICLR)*. 2022.

[103] Ferran Alet*, Dylan Doblar*, Allan Zhou, Joshua B. Tenenbaum, Kenji Kawaguchi, **Chelsea Finn**. Noether Networks: Meta-Learning Useful Conserved Quantities. *Neural Information Processing Systems (NeurIPS)*. 2021.

[102] Nicholas Rhinehart, Jenny Wang, Glen Berseth, John D Co-Reyes, Danijar Hafner, **Chelsea Finn**, Sergey Levine. Information is Power: Intrinsic Control via Information Capture. *Neural Information Processing Systems (NeurIPS)*. 2021.

[101] Huaxiu Yao*, Yu Wang*, Ying Wei, Peilin Zhao, Mehrdad Mahdavi, Defu Lian, **Chelsea Finn**. Meta-Learning with an Adaptive Task Scheduler. *Neural Information Processing Systems (NeurIPS)*. 2021.

[100] Tianhe Yu*, Aviral Kumar*, Yevgen Chebotar, Karol Hausman, Sergey Levine, **Chelsea Finn**. Conservative Data Sharing for Multi-Task Offline Reinforcement Learning. *Neural Information Processing Systems (NeurIPS)*. 2021.

[99] Tianhe Yu*, Aviral Kumar*, Rafael Rafailov, Aravind Rajeswaran, Sergey Levine, **Chelsea Finn**. COMBO: Conservative Offline Model-Based Policy Optimization. *Neural Information Processing Systems (NeurIPS)*. 2021.

[98] Rafael Rafailov, Tianhe Yu, Aravind Rajeswaran, **Chelsea Finn**. Visual Adversarial Imitation Learning using Variational Models. *Neural Information Processing Systems (NeurIPS)*. 2021.

[97] Archit Sharma, Abhishek Gupta, Sergey Levine, Karol Hausman, **Chelsea Finn**. Autonomous Reinforcement Learning via Subgoal Curricula. *Neural Information Processing Systems (NeurIPS)*. 2021.

[96] Christopher Fifty, Ehsan Amid, Zhe Zhao, Tianhe Yu, Rohan Anil, **Chelsea Finn**. Efficiently Identifying Task Groupings for Multi-Task Learning. *Neural Information Processing Systems (NeurIPS)*. 2021.

[95] Marvin Zhang*, Henrik Marklund*, Nikita Dhawan*, Abhishek Gupta, Sergey Levine, **Chelsea Finn**. Adaptive Risk Minimization: A Meta-Learning Approach for Tackling Group Shift. *Neural Information Processing Systems (NeurIPS)*. 2021.

[94] Guodong Zhang, Kyle Hsu, Jianing Li, **Chelsea Finn**, Roger Grosse. Differentiable Annealed Importance Sampling and the Perils of Gradient Noise. *Neural Information Processing Systems (NeurIPS)*. 2021.

[93] Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, **Chelsea Finn**. BC-Z: Zero-Shot Task Generalization with Robotic Imitation Learning. *Conference on Robot Learning (CoRL)*. 2021.

[92] Bohan Wu, Suraj Nair, Li Fei-Fei, **Chelsea Finn**. Example-Driven Model-Based Reinforcement Learning for Solving Long-Horizon Visuomotor Tasks. *Conference on Robot Learning (CoRL)*. 2021.

[91] Aviral Kumar*, Anikait Singh*, Stephen Tian, **Chelsea Finn**, Sergey Levine. A Workflow for Offline Model-Free Robotic Reinforcement Learning. *Conference on Robot Learning (CoRL)*. 2021.

[90] Suraj Nair, Eric Mitchell, Kevin Chen, Brian Ichter, Silvio Savarese, **Chelsea Finn**. Learning Language-Conditioned Robot Behavior from Offline Data and Crowd-Sourced Annotation. *Conference on Robot Learning (CoRL)*. 2021.

[89] Dmitry Kalashnikov*, Jake Varley*, Yevgen Chebotar, Benjamin Swanson, Rico Jonschkowski, **Chelsea Finn**, Sergey Levine, Karol Hausman. MT-Opt: Continuous Multi-Task Robotic Reinforcement Learning at Scale. *Conference on Robot Learning (CoRL)*. 2021.

[88] Ali Ghadirzadeh, Xi Chen, Petra Poklukar, **Chelsea Finn**, Marten Bjorkman, Danica Kragic. Bayesian Meta-Learning for Few-Shot Policy Adaptation Across Robotic Platforms. *International Conference on Intelligent Robots and Systems (IROS)*. 2021.

[87] Annie S. Chen, Suraj Nair, **Chelsea Finn**. Learning Generalizable Robotic Reward Functions from "In-The-Wild" Human Videos. *Robotics: Science and Systems (RSS)*. 2021.

[86] Evan Z. Liu*, Behzad Haghgoo*, Annie S. Chen*, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa, Percy Liang, **Chelsea Finn**. Just Train Twice: Improving Group Robustness without Training Group Information. *International Conference on Machine Learning (ICML)*. 2021.

[85] Evan Z. Liu, Aditi Raghunathan, Percy Liang, **Chelsea Finn**. Decoupling Exploration and Exploitation for Meta-Reinforcement Learning without Sacrifices. *International Conference on Machine Learning (ICML)*. 2021.

[84] Annie Xie, James Harrison, **Chelsea Finn**. Deep Reinforcement Learning amidst Lifelong Non-Stationarity. *International Conference on Machine Learning (ICML)*. 2021.

[83] Eric Mitchell, Rafael Rafailov, Xue Bin (Jason) Peng, Sergey Levine, **Chelsea Finn**. Offline Meta-Reinforcement Learning with Advantage Weighting. *International Conference on Machine Learning (ICML)*. 2021.

[82] Yevgen Chebotar, Karol Hausman, Yao Lu, Ted Xiao, Dmitry Kalashnikov, Jake Varley, Alex Irpan, Ryan Julian, **Chelsea Finn**, Sergey Levine. Actionable Models: Unsupervised Offline Reinforcement Learning of Robotic Skills. *International Conference on Machine Learning (ICML)*. 2021.

[81] Pang Wei Koh*, Shiori Sagawa*, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Sara Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, **Chelsea Finn**, Percy Liang. WILDS: A Benchmark of in-the-Wild Distribution Shifts. *International Conference on Machine Learning (ICML)*. 2021.

[80] Jared Davis, Albert Gu, Tri Dao, Krzysztof Choromanski, Christopher Re, **Chelsea Finn**, Percy Liang. Catformer: Designing Stable Transformers via Sensitivity Analysis. *International Conference on Machine Learning (ICML)*. 2021.

[79] Bohan Wu, Suraj Nair, Roberto Martin-Martin, Li Fei-Fei, **Chelsea Finn**. Greedy Hierarchical Variational Autoencoders for Large-Scale Video Prediction. *Conference on Computer Vision and Pattern Recognition (CVPR)*. 2021.

[78] Julian Ibarz, Jie Tan, **Chelsea Finn**, Mrinal Kalakrishnan, Peter Pastor, Sergey Levine. How to Train Your Robot with Deep Reinforcement Learning; Lessons We've Learned. *International Journal of Robotics Research (IJRR))*. 2021.

[77] Rafael Rafailov*, Tianhe Yu*, Aravind Rajeswaran, **Chelsea Finn**. Offline Reinforcement Learning from Images with Latent Space Models. *Conference on Learning for Decision Making and Control (L4DC)*. 2021.

[76] Annie Chen*, HyunJi Nam*, Suraj Nair*, **Chelsea Finn**. Batch Exploration with Examples for Scalable Robotic Reinforcement Learning. *IEEE Robotics and Automation Letters (RA-L)* and *International Conference on Robotics and Automation (ICRA)*. 2021.

[75] Brijen Thananjeyan, Ashwin Balakrishna, Suraj Nair, Michael Luo, Krishnan Srinivasan, Minho Hwang, Joseph E. Gonzalez, Julian Ibarz, **Chelsea Finn**, Ken Goldberg. Recovery RL: Safe Reinforcement Learning with Learned Recovery Zones. *IEEE Robotics and Automation Letters (RA-L)* and *International Conference on Robotics and Automation (ICRA)*. 2021.

[74] Allan Zhou, Tom Knowles, **Chelsea Finn**. Meta-Learning Symmetries by Reparametrization. *International Conference on Learning Representations (ICLR)*. 2021.

[73] Stephen Tian, Suraj Nair, Frederik Ebert, Sudeep Dasari, Ben Eysenbach, **Chelsea Finn**, Sergey Levine. Model-Based Visual Planning with Self-Supervised Functional Distances. *International Conference on Learning Representations (ICLR)*. 2021.

[72] Glen Berseth, Daniel Geng, Coline Devin, **Chelsea Finn**, Dinesh Jayaraman, Sergey Levine. SMiRL: Surprise Minimizing RL in Dynamic Environments. *International Conference on Learning*

Representations (ICLR). 2021.

[71] Karl Schmeckpeper, Oleh Rybkin, Kostas Daniilidis, Sergey Levine, **Chelsea Finn**. Reinforcement Learning with Videos: Combining Offline Observations with Interaction. *Conference on Robot Learning (CoRL)*. 2020.

[70] Ryan Julian, Benjamin Swanson, Gaurav Sukhatme, Sergey Levine, **Chelsea Finn**, Karol Hausman. Never Stop Learning: The Effectiveness of Fine-Tuning in Robotic Reinforcement Learning. *Conference on Robot Learning (CoRL)*. 2020.

[69] Annie Xie, Dylan Losey, Ryan Tolsma, **Chelsea Finn**, Dorsa Sadigh. Learning Latent Representations to Influence Multi-Agent Interaction. *Conference on Robot Learning (CoRL)*. 2020.

[68] Anusha Nagabandi, Zihao Zhao, Kate Rakelly, **Chelsea Finn**, Sergey Levine. Latent State Models for Meta-Reinforcement Learning from Images. *Conference on Robot Learning (CoRL)*. 2020.

[67] Saurabh Kumar, Aviral Kumar, Sergey Levine, **Chelsea Finn**. One Solution is Not All You Need: Few-Shot Extrapolation via Structured MaxEnt RL. *Neural Information Processing Systems (NeurIPS)*. 2020.

[66] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, **Chelsea Finn**. Gradient Surgery for Multi-Task Learning. *Neural Information Processing Systems (NeurIPS)*. 2020.

[65] Kelvin Xu, Siddharth Verma, **Chelsea Finn**, Sergey Levine. Learning Skillful Resets: Acquisition of Behavior via Reset-Free Games. *Neural Information Processing Systems (NeurIPS)*. 2020.

[64] Lisa Lee, Ben Eysenbach, Ruslan Salakhutdinov, Shixiang Gu, **Chelsea Finn**. Weakly-Supervised Reinforcement Learning for Controllable Behavior. *Neural Information Processing Systems (NeurIPS)*. 2020.

[63] Tianhe Yu*, Garrett Thomas*, Lantao Yu, Stefano Ermon, James Zou, Sergey Levine, **Chelsea Finn**, Tengyu Ma. MOPO: Model-based Offline Policy Optimization. *Neural Information Processing Systems (NeurIPS)*. 2020.

[62] Karl Pertsch*, Oleh Rybkin*, Frederik Ebert, **Chelsea Finn**, Dinesh Jayaraman, Sergey Levine. Long-Horizon Visual Planning with Goal-Conditioned Hierarchical Predictors. *Neural Information Processing Systems (NeurIPS)*. 2020.

[61] James Harrison, Apoorva Sharma, **Chelsea Finn**, Marco Pavone. Continuous Meta-Learning without Tasks. *Neural Information Processing Systems (NeurIPS)*. 2020.

[60] Xingyou Song, Yuxiang Yang, Krzysztof Choromanski, Ken Caluwaerts, Wenbo Gao, **Chelsea Finn**, Jie Tan. Rapidly Adaptable Legged Robots via Evolutionary Meta-Learning. *International Conference on Intelligent Robots and Systems (IROS)*. 2020.

[59] Karl Schmeckpeper, Annie Xie, Oleh Rybkin, Stephen Tian, Kostas Daniilidis, Sergey Levine, **Chelsea Finn**. Learning Predictive Models from Observation and Interaction. *European Conference on Computer Vision (ECCV)*. 2020.

[58] Suraj Nair, Silvio Savarese, **Chelsea Finn**. Goal-Aware Prediction: Learning to Model What Matters. *International Conference on Machine Learning (ICML)*. 2020.

[57] Jesse Zhang, Brian Cheung, **Chelsea Finn**, Sergey Levine, Dinesh Jayaraman. Cautious Adaptation For Reinforcement Learning in Safety-Critical Settings. *International Conference on Machine Learning (ICML)*. 2020.

[56] Akhil Padmanabha, Frederik Ebert, Stephen Tian, Roberto Calandra, **Chelsea Finn**, Sergey Levine. OmniTact: Compact Multi-Directional Optical Tactile Sensor for Robotic Manipulation. *International Conference on Robotics and Automation (ICRA)*. 2020.

[55] Suraj Nair, Mohammad Babaeizadeh, **Chelsea Finn**, Sergey Levine, Vikash Kumar. Time Reversal as Self-Supervision. *International Conference on Robotics and Automation (ICRA)*. 2020.

[54] Avi Singh, Eric Jang, Daniel Kappler, Mohi Khansari, Murtaza Dalal, Alex Irpan, Sergey Levine, Mohi Khansari, **Chelsea Finn**. Scalable Multi-Task Imitation Learning with Autonomous Improvement. *International Conference on Robotics and Automation (ICRA)*. 2020.

[53] Mingzhang Yin, George Tucker, Mingyuan Zhou, Sergey Levine, **Chelsea Finn**. Meta-Learning without Memorization. *International Conference on Learning Representations (ICLR)*. 2020.

[52] Suraj Nair, **Chelsea Finn**. Hierarchical Foresight: Self-Supervised Learning of Long-Horizon Tasks via Visual Subgoal Generation. *International Conference on Learning Representations (ICLR)*. 2020.

[51] Allan Zhou, Eric Jang, Daniel Kappler, Alex Herzog, Mohi Khansari, Paul Wohlhart, Yunfei Bai, Mrinal Kalakrishnan, Sergey Levine, **Chelsea Finn**. Watch, Try, Learn: Meta-Learning from Demonstrations and Rewards. *International Conference on Learning Representations (ICLR)*. 2020.

[50] Manoj Kumar, Mohammad Babaeizadeh, Dumitru Erhan, **Chelsea Finn**, Sergey Levine, Laurent Dinh, Durk Kingma. VideoFlow: A Conditional Flow-Based Model for Stochastic Video Generation. *International Conference on Learning Representations (ICLR)*. 2020.

[49] Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H Campbell, Konrad Czechowski, Dumitru Erhan, **Chelsea Finn**, Piotr Kozakowski, Sergey Levine, Afroz Mohiuddin, Ryan Sepassi, George Tucker, Henryk Michalewski. Model-Based Reinforcement Learning for Atari. *International Conference on Learning Representations (ICLR)*. 2020.

[48] Mark Woodward, **Chelsea Finn**, Karol Hausman. Learning to Interactively Learn and Assist. *AAAI Conference on Artificial Intelligence (AAAI)*. 2020.

[47] Sudeep Dasari, Frederik Ebert, Stephen Tian, Suraj Nair, Bernadette Bucher, Karl Schmeckpeper, Siddharth Singh, Sergey Levine, **Chelsea Finn**. RoboNet: Large-Scale Multi-Robot Learning. *Conference on Robot Learning (CoRL)*. 2019.

[46] Tianhe Yu*, Deirdre Quillen*, Zhanpeng He, Ryan Julian, Karol Hausman, **Chelsea Finn**, Sergey Levine. Meta-World: A Benchmark and Evaluation for Multi-Task and Meta Reinforcement Learning. *Conference on Robot Learning (CoRL)*. 2019.

[45] John Co-Reyes, Rishi Veerapaneni, Michael Chang, Michael Janner, **Chelsea Finn**, Jiajun Wu, Josh Tenenbaum, Sergey Levine. Entity Abstraction in Visual Model-Based Reinforcement Learning. *Conference on Robot Learning (CoRL)*. 2019.

[44] Allan Jabri, Kyle Hsu, Ben Eysenbach, Abhishek Gupta, Sergey Levine, **Chelsea Finn**. Unsupervised Curricula for Visual Meta-Reinforcement Learning. *Neural Information Processing Systems (NeurIPS)*. 2019.

[43] Russell Mendonca, Abhishek Gupta, Rosen Kralev, Pieter Abbeel, Sergey Levine, **Chelsea Finn**. Guided Meta Policy Search. *Neural Information Processing Systems (NeurIPS)*. 2019.

[42] Yiding Jiang, Shixiang Gu, Kevin Murphy, **Chelsea Finn**. Language as an Abstraction for Hierarchical Reinforcement Learning. *Neural Information Processing Systems (NeurIPS)*. 2019.

[41] Aravind Rajeswaran*, **Chelsea Finn***, Sham Kakade, Sergey Levine. Meta-Learning with Implicit Gradients. *Neural Information Processing Systems (NeurIPS)*. 2019.

[40] Lantao Yu, Tianhe Yu, **Chelsea Finn**, Stefano Ermon. Meta-Inverse Reinforcement Learning with Probabilistic Context Variables. *Neural Information Processing Systems (NeurIPS)*. 2019.

[39] Tianhe Yu, Pieter Abbeel, Sergey Levine, **Chelsea Finn**. One-Shot Hierarchical Imitation Learning of Compound Visuomotor Tasks. *International Conference on Intelligent Robots and Systems (IROS)*. 2019.

[38] Tianhe Yu, Gleb Shevchuk, Dorsa Sadigh, **Chelsea Finn**. Unsupervised Visuomotor Control via Distributional Planning Networks. *Robotics: Science and Systems (RSS)*. 2019.

[37] Annie Xie, Frederik Ebert, Sergey Levine, **Chelsea Finn**. Improvisation through Physical Understanding: Using Novel Objects as Tools with Visual Foresight. *Robotics: Science and Systems (RSS)*. 2019.

[36] Avi Singh, Larry Yang, Kristian Hartikainen, **Chelsea Finn**, Sergey Levine. End-to-End Robotic Reinforcement Learning without Reward Engineering. *Robotics: Science and Systems (RSS)*. 2019.

[35] **Chelsea Finn***, Aravind Rajeswaran*, Sham Kakade, Sergey Levine. Online Meta-Learning. *International Conference on Machine Learning (ICML)*. 2019.

[34] Kate Rakelly*, Aurick Zhou*, Deirdre Quillen, **Chelsea Finn**, Sergey Levine. Efficient Off-Policy Meta-Reinforcement Learning via Probabilistic Context Variables. *International Conference on Machine Learning (ICML)*. 2019.

[33] Kelvin Xu, Ellis Ratner, Anca Dragan, Sergey Levine, **Chelsea Finn**. Learning a Prior over Intent via Meta-Inverse Reinforcement Learning. *International Conference on Machine Learning (ICML)*. 2019.

[32] Stephen Tian*, Frederik Ebert*, Dinesh Jayaraman, Mayur Mudigonda, **Chelsea Finn**, Roberto Calandra, Sergey Levine. Manipulation by Feel: Touch-Based Control with Deep Predictive Models. *International Conference on Robotics and Automation (ICRA)*. 2019.

[31] Yuxiang Yang, Ken Caluwaerts, Atil Iscen, Jie Tan, **Chelsea Finn**. NoRML: No-Reward Meta Learning. *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2019.

[30] Michael Janner, Sergey Levine, Bill Freeman, Josh Tenenbaum, **Chelsea Finn**, Jiajun Wu. Reasoning About Physical Interactions with Object-Oriented Prediction and Planning. *International Conference on Learning Representations (ICLR)*. 2019.

[29] Anusha Nagabandi, **Chelsea Finn**, Sergey Levine. Deep Online Learning Via Meta-Learning: Continual Adaptation for Model-Based RL. *International Conference on Learning Representations (ICLR)*. 2019.

[28] Kyle Hsu, Sergey Levine, **Chelsea Finn**. Unsupervised Learning via Meta-Learning. *International Conference on Learning Representations (ICLR)*. 2019.

[27] Anusha Nagabandi*, Ignasi Clavera*, Simin Liu Ronald S. Fearing, Pieter Abbeel, Sergey Levine, **Chelsea Finn**. Learning to Adapt in Dynamic, Real-World Environments Through Meta-Reinforcement Learning. *International Conference on Learning Representations (ICLR)*. 2019.

[26] **Chelsea Finn***, Kelvin Xu*, Sergey Levine. Probabilistic Model-Agnostic Meta-Learning. *Neural Information Processing Systems (NIPS)*. 2018.

[25] Annie Xie, Avi Singh, Sergey Levine, **Chelsea Finn**. Few-shot Goal Inference for Visuomotor

Learning and Planning. *Conference on Robot Learning (CoRL)*. 2018.

[24] Frederik Ebert, Sudeep Dasari, Alex Lee, Sergey Levine, **Chelsea Finn**. Robustness via Retrying: Closed-Loop Robotic Manipulation with Self-Supervised Learning. *Conference on Robot Learning (CoRL)*. 2018.

[23] Aravind Srinivas, Allan Jabri, Pieter Abbeel, Sergey Levine, **Chelsea Finn**. Universal Planning Networks. *International Conference on Machine Learning (ICML)*. 2018.

[22] Tianhe Yu*, **Chelsea Finn***, Annie Xie, Sudeep Dasari, Pieter Abbeel, Sergey Levine. One-Shot Imitation from Observing Humans via Domain-Adaptive Meta-Learning. *Robotics: Science and Systems (RSS)*. 2018.

[21] Deirdre Quillen, Eric Jang, Ofir Nachum, **Chelsea Finn**, Julian Ibarz, Sergey Levine. Deep Reinforcement Learning for Vision-Based Robotic Grasping: A Simulated Comparative Evaluation of Off-Policy Methods. *International Conference on Robotics and Automation (ICRA)*. 2018.

[20] **Chelsea Finn**, Sergey Levine. Meta-Learning and Universality: Deep Representations and Gradient Descent can Approximate any Learning Algorithm. *International Conference on Learning Representations (ICLR)*. 2018.

[19] Erin Grant, **Chelsea Finn**, Sergey Levine, Trevor Darrell, Tom Griffiths. Recasting Gradient-Based Meta-Learning as Hierarchical Bayes. *International Conference on Learning Representations (ICLR)*. 2018.

[18] Mohammad Babaeizadeh, **Chelsea Finn**, Dumitru Erhan, Roy H. Campbell, Sergey Levine. Stochastic Variational Video Prediction. *International Conference on Learning Representations (ICLR)*. 2018.

[17] **Chelsea Finn***, Tianhe Yu*, Tianhao Zhang, Pieter Abbeel, Sergey Levine. One-Shot Visual Imitation Learning via Meta-Learning. *Conference on Robot Learning (CoRL)*. 2017.

[16] Frederik Ebert, **Chelsea Finn**, Alex Lee, Sergey Levine. Self-Supervised Visual Planning with Temporal Skip-Connections. *Conference on Robot Learning (CoRL)*. 2017.

[15] **Chelsea Finn**, Pieter Abbeel, Sergey Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. *International Conference on Machine Learning (ICML)*. 2017.

[14] **Chelsea Finn**, Tianhe Yu, Justin Fu, Pieter Abbeel, Sergey Levine. Generalizing Skills with Semi-Supervised Reinforcement Learning. *International Conference on Learning Representations (ICLR)*. 2017.

[13] **Chelsea Finn**, Sergey Levine. Deep Visual Foresight for Planning Robot Motion. *International Conference on Robotics and Automation (ICRA)*. 2017.

[12] William Montgomery*, Anurag Ajay*, **Chelsea Finn**, Pieter Abbeel, Sergey Levine. Reset-Free Guided Policy Search: Efficient Deep Reinforcement Learning with Stochastic Initial States. *International Conference on Robotics and Automation (ICRA)*. 2017.

[11] **Chelsea Finn**, Ian Goodfellow, Sergey Levine. Unsupervised Learning for Physical Interaction through Video Prediction. *Neural Information Processing Systems (NIPS)*. 2016.

[10] Eric Tzeng, Coline Devin, Judy Hoffman, **Chelsea Finn**, Pieter Abbeel, Sergey Levine, Kate Saenko and Trevor Darrell. Adapting Deep Visuomotor Representations with Weak Pairwise Constraints. *Workshop on the Algorithmic Foundations of Robotics (WAFR)*. 2016.

[9] **Chelsea Finn**, Sergey Levine, Pieter Abbeel. Guided Cost Learning: Deep Inverse Optimal Control via Policy Optimization. *International Conference on Machine Learning (ICML)*. 2016.

[8] **Chelsea Finn**, Xin Yu Tan, Yan Duan, Trevor Darrell, Sergey Levine, Pieter Abbeel. Deep Spatial Autoencoders for Visuomotor Learning. *International Conference on Robotics and Automation (ICRA)*. 2016.

[7] Marvin Zhang, Zoe McCarthy, **Chelsea Finn**, Sergey Levine, Pieter Abbeel. Learning Deep Neural Network Policies with Continuous Memory States. *International Conference on Robotics and Automation (ICRA)*. 2016.

[6] Sergey Levine*, **Chelsea Finn***, Trevor Darrell, Pieter Abbeel. End-to-End Training of Deep Visuomotor Policies. *Journal of Machine Learning Research (JMLR)*. 2016.

[5] Hsueh-Cheng Wang, **Chelsea Finn**, Liam Paull, Michael Kaess, Ruth Rosenholtz, Seth Teller, John Leonard. Bridging text spotting and SLAM with junction features. *International Conference on Intelligent Robots and Systems (IROS)*. 2015.

[4] Dylan Hadfield-Menell, Alex Xavier Lee, **Chelsea Finn**, Eric Tzeng, Sandy Huang, Pieter Abbeel. Beyond Lowest-Warping Cost Action Selection in Trajectory Transfer. *International Conference on Robotics and Automation (ICRA)*. 2015.

[3] James Duyck, **Chelsea Finn**, Andy Hutcheon, Pablo Vera, Joaquin Salas, Sai Ravela. Sloop: A pattern retrieval engine for individual animal identification. *Pattern Recognition*. 2014.

[2] **Chelsea Finn**, James Duyck, Andy Hutcheon, Pablo Vera, Joaquin Salas, Sai Ravela. Relevance feedback in biometric retrieval of animal photographs. *Mexican Conference on Pattern Recognition (MCPR)*. 2014.

[1] Sai Ravela, James Duyck, **Chelsea Finn**. Vision-Based Biometrics for Conservation. *Mexican Conference on Pattern Recognition (MCPR)*. 2013.

Workshop Papers and Abstracts

Chelsea Finn*, Paul Christiano*, Pieter Abbeel, Sergey Levine. A Connection between Generative Adversarial Networks, Inverse Reinforcement Learning, and Energy-based Models. *NIPS Workshop on Adversarial Training*. 2016.

Mark Woodward, **Chelsea Finn**. Active One-Shot Learning. *NIPS Deep Reinforcement Learning Workshop*. 2016.

Chelsea Finn, Lisa Anne Hendricks, Trevor Darrell Learning Compact Convolutional Neural Networks with Nested Dropout. *International Conference on Learning Representations (ICLR) – Workshop Contribution*. 2015.

Advising

PhD research:
Frederik Ebert (now founder & CEO of Emancro)
Tianhe Yu (now research scientist at Google Brain)
Suraj Nair (now research scientist at TRI)
Evan Z. Liu (now research scientist at Generally Intelligent)
Allan Zhou
Annie Xie

Eric Mitchell
Archit Sharma
Kyle Hsu
Annie S. Chen
Alexander Khazatsky
Yoonho Lee
Zhihao (Tony) Zhao
Zipeng Fu
Jonathan Yang
Rafael Rafailov

Masters research:
Frederik Ebert (PhD at UC Berkeley)
Henrik Marklund (PhD student at Stanford)
Rafael Rafailov (PhD student at Stanford)
Ahmed Ahmed (PhD student at Stanford)
Moo Jin Kim (PhD student at Stanford)

Undergraduate research:
Nopphon Sirinart (MS at Stanford)
Justin Fu (PhD student at UC Berkeley)
Marvin Zhang (PhD student at UC Berkeley)
Anurag Ajay (PhD student at MIT)
Tianhe Yu (PhD at Stanford)
Xin Yu Tan
Annie Xie (PhD student at Stanford)
Sudeep Dasari (PhD student at CMU)
Russell Mendonca (PhD student at CMU)
Kyle Hsu (PhD student at Stanford)
Tom Knowles
HyunJi (Alex) Nam (PhD student at Stanford)
Annie Chen (PhD student at Stanford)
Fahim Tajwar (PhD student at CMU)
Behzad Haghgoo
Kyle Hatch
Max Sobol Mark
Max Du
Olivia Lee
Takao Yagatai
Leo Dong
Caroline Choi

Independent research:
Mark Woodward (next: Google AI resident)
Rosen Kralev

Outreach

LINXS Summer Research Program, Faculty Host 2022
Hosted an HBCU undergraduate researcher for one summer.

AI Research Mentoring Program, Co-Organizer 2017-present
Coordinating a research mentoring program for underrepresented undergraduates.
Grew the program to UC Berkeley, Stanford, and CMU

Berkeley AI & AI4ALL Camp, Co-Organizer 2018
Organized 5-day camp for underprivileged high-school students
Free camp with hands-on introduction to CS and AI, aiming to increase diversity in AI.

Berkeley AI & AI4ALL Camp, Co-Organizer 2017
Organized inaugural 2-day camp for 24 underprivileged high-school students
Free camp with hands-on introduction to CS and AI, aiming to increase diversity in AI.

Women in Machine Learning (WiML) 2017-present
Invited speaker or panelist, CoRL 2019, 2021
Lunch mentor, ICML 2017, NeurIPS 2018, 2019, 2020, 2021
Co-organized WiML evening event, CoRL 2017

UC Berkeley Women in EECS, Outreach Co-coordinator 2016-2017
Organized events for minorities, with advice on pursuing research & grad school
Organized day-long STEM exploration workshop for Girl Scouts.

UC Berkeley Women in EECS, Co-President 2015-2016

Career Panels and Talks at Minorities in STEM events 2015-present
Stanford WITE and BASES Women in AI Panel, panelist, 2023
Stanford-Berkeley Women in EECS Meet Up, speaker & panelist, 2015, 2019, 2022, 2023
ProjectCS Girls hackathon, speaker, 2023
Stanford AI4ALL Summer Camp, speaker, 2020, 2021, 2022, 2023
RSS Pioneers Workshop, panelist, 2022
Stanford Engineering Research Introductions (SERIS), faculty speaker, 2022
VEX Robotics Girl Powered Workshop, keynote speaker, 2021
MIT Graduate Women in Robotics Community, lunch speaker, 2021
REsearch Exposure in Socially Relevant Computing, panelist, 2021
Harker School Research Symposium, keynote, 2021
Stanford Society of Women Engineers (SWE), mentor, 2021
Stanford Women in Electrical Engineering (WEE), lunch panelist, 2021
Stanford Women in Computer Science (WiCS), speaker, 2020
Inclusion@RSS, panelist, 2020
ICML NewInML Workshop, panelist, 2020
CVPR Women in Computer Vision Workshop, keynote, panelist, mentor, 2020
RSS Women in Robotics Workshop, speaker, 2020
CISCO Women Rock IT Live Broadcast, featured speaker, 2019
Khipu: Latin American Meeting in AI, Women in AI event, panelist 2019
Girls Programming League (GPL), keynote, 2019
Pioneers in Engineering (PiE) Kick-Off, keynote, 2018
Graduate Pathways to STEM, panelist, 2016
SWE Parent Education Outreach Program, panelist, 2017
NASA When I Grow Up Career Exploration Event, panelist, 2016

Professional Activities

Board Member:
International Conference on Learning Representations (ICLR) (2023-present)

Program Chair:
International Conference on Learning Representations (ICLR) 2022

Workshops Chair:
International Conference on Learning Representations (ICLR) 2021

Tutorials Chair:
Reinforcement Learning and Decision Making (RLDM) 2022

Senior Area Chair:
International Conference on Machine Learning (ICML) 2023
Neural Information Processing Systems (NeurIPS) 2023

Area Chair:
Neural Information Processing Systems (NeurIPS) 2019, 2020, 2021, 2022
Robotics: Science and Systems (RSS) 2020, 2021
International Conference on Machine Learning (ICML) 2019, 2020, 2021
International Conference on Learning Representations (ICLR) 2019, 2020, 2021, 2023
Conference on Robot Learning (CoRL) 2018, 2019, 2021, 2022, 2023

Reviewing:
Proceedings of the Royal Society A, 2022
CRA Computing Innovation Fellows, Reviewer 2020
IEEE Robotics and Automation Letters (RA-L) 2016, 2017, 2018, 2019, 2020, 2021
IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2016, 2017, 2019
Robotics: Science and Systems (RSS) 2016, 2019, 2022
IEEE International Conference on Robotics and Automation (ICRA) 2016, 2017, 2018, 2019
Foundations and Trends in Machine Learning 2018
ACM Siggraph 2018
Neural Information Processing Systems (NIPS) 2016, 2017, 2018
International Conference on Machine Learning (ICML) 2017, 2018
International Conference on Learning Representations (ICLR) 2017, 2018
Conference on Robot Learning (CoRL) 2017, 2020
International Journal of Robotics Research (IJRR) 2016, 2017
Communications of the ACM 2016

Workshop Organization:
Workshop on Distribution Shifts: Connecting Methods and Applications, NeurIPS 2022
Workshop on Pre-Training: Perspectives, Pitfalls, and Paths Forward, ICML 2022
Workshop on Learning from Diverse Offline Data, RSS 2022
Workshop on Robot Learning in the Cloud, RSS 2022
Deep Reinforcement Learning Workshop, NeurIPS 2021
Robotics for People (R4P): Perspectives on Interaction, Learning and Safety, RSS 2021
Deep Reinforcement Learning Workshop, NeurIPS 2020
Beyond "Tabula Rasa" in Reinforcement Learning Workshop, ICLR 2020
Deep Reinforcement Learning Workshop, NeurIPS 2019
Workshop on Learning with Rich Experience, NeurIPS 2019
Workshop on Multi-Task and Lifelong Reinforcement Learning, ICML 2019
Workshop on Imitation, Intent, and Interaction, ICML 2019
Workshop on Structures and Priors in Reinforcement Learning, ICLR 2019
Workshop on Deep Learning for Action and Interaction, NIPS 2016

Selected Press Coverage

"Robot can find keys in a bag just by listening as it rummages around," by Jeremy Hsu. New Scientist. 13 June 2022.

"Can A.I. Grade Your Next Test?," by Cade Metz. The New York Times. 20 July 2021.

"The key to smarter robot collaborators may be more simplicity," by Karen Hao. MIT Technology Review. 13 November 2020.

"Artificial Imagination: How machines could learn creativity and common sense, among other human qualities," by George Musser. Scientific American. May 2019.

"A Robot has Figured Out How to Use Tools," by Will Knight. MIT Technology Review. 11 April 2019.

"The Robots are Here: All they need is a brain," by Daniel Cossins. New Scientist. 2 February 2019.

"Don't Just Lecture Robots – Make Them Learn," by Matt Simon. Wired. 9 July 2018.

"Robot learns by playing and imagines its own future," by Jonathan Bloom. ABC 7 News. 18 December 2017.

"Researchers train robots to see into the future," by John Biggs. TechCrunch. 8 December 2017.

"Building A.I. That Can Build A.I.," by Cade Metz. The New York Times. 5 November 2017.

"The Education of Brett the Robot," by Matt Simon. Wired. 21 September 2017.

"Google Builds a Robotic Hive-Mind Kindergarten," by Will Knight. MIT Technology Review. 3 October 2016.

"This Preschool is for Robots," by Jack Clark. Bloomberg Business. 2 September 2015.

"Robot Demonstrates Human-Like Learning Abilities," by Jonathan Bloom. ABC 7 News. 22 May 2015.

"Deep Learning Robots, DRC Practice, and Drone Pilot Competition," by Evan Ackerman. IEEE Spectrum. 22 May 2015.

"New approach trains robots to match human dexterity and speed," by John Markoff. The New York Times. 21 May 2015.