Collaborative Research: NRI: INT: Scalable, Customizable, Robot Learning with Humans

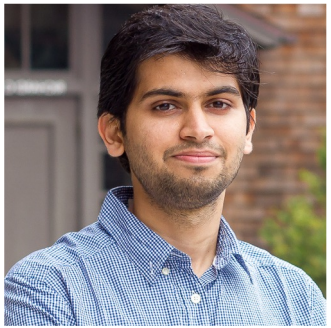

Deepak Pathak
CMU

Pieter Abbeel
UC Berkeley

Anca Dragan
UC Berkeley

Charles Kemp
Georgia Tech

A. Kapusta, Z. Erickson, H. M. Clever, W. Yu, C. K. Liu, G. Turk, C. C. Kemp, “Personalized Collaborative Plans for Robot-Assisted Dressing via Optimization and Simulation,” Autonomous Robots, 2019.

A. Kapusta, P. M. Grice, H. M. Clever, Y. Chitalia, D. Park, and C. C. Kemp, “A system for bedside assistance that integrates a robotic bed and a mobile manipulator,” PloS one, 2019.

Approach

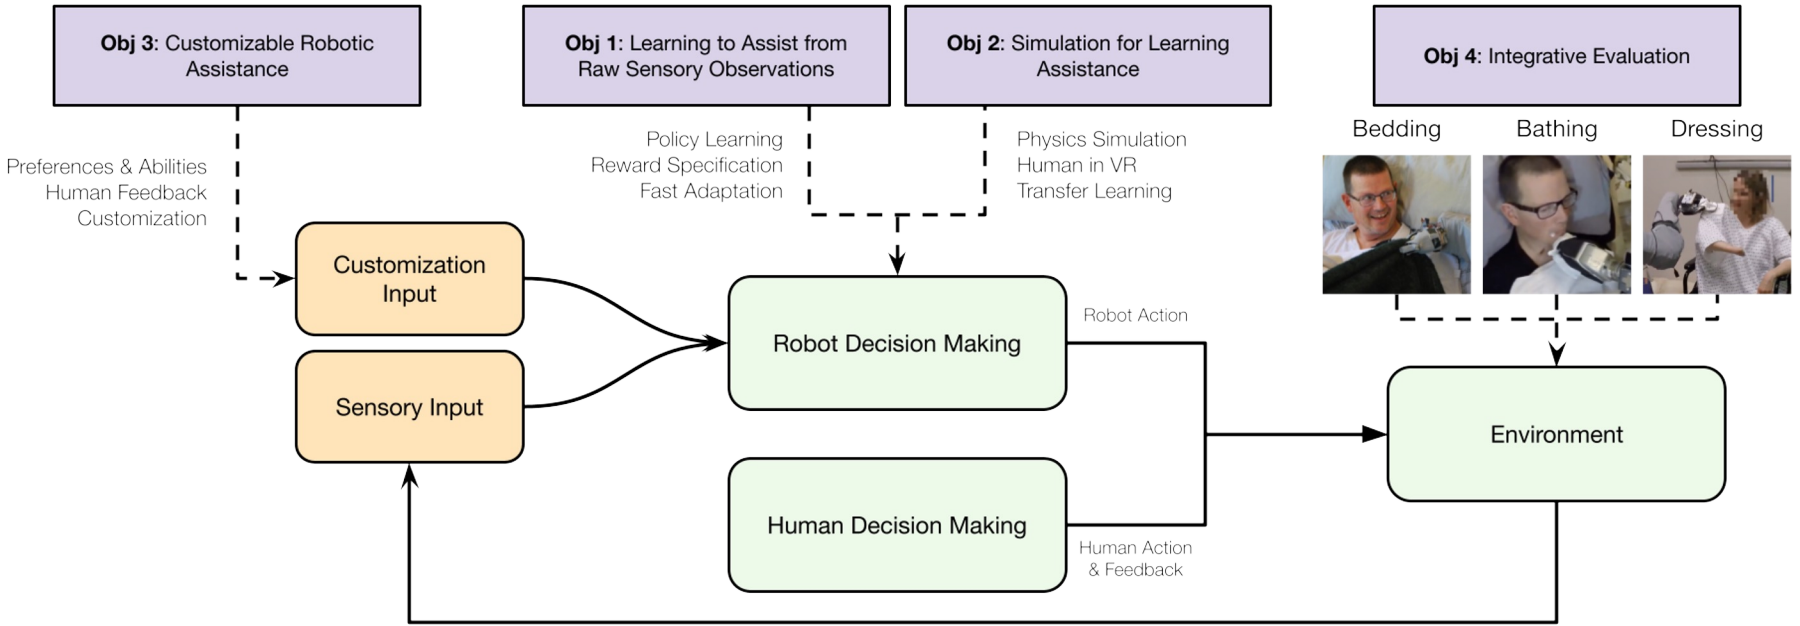
Obj 3: Customizable Robotic Assistance
Obj 1: Learning to Assist from Raw Sensory Observations
Obj 2: Simulation for Learning Assistance
Obj 4: Integrative Evaluation
Preferences & Abilities
Human Feedback
Customization
Policy Learning
Reward Specification
Fast Adaptation
Physics Simulation
Human in VR
Transfer Learning
Bedding
Bathing
Dressing
Customization Input
Sensory Input
Robot Decision Making
Robot Action
Human Decision Making
Human Action & Feedback
Environment

Obj 1: Learning to Assist from Raw Sensory Observations

Robot need not only operate from **raw data** but do so in a **dynamic** way to be safe

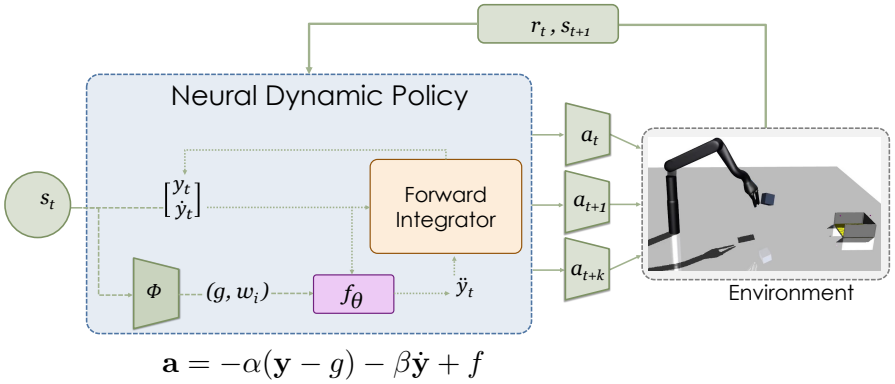

Neural Dynamic Policy (NDP)

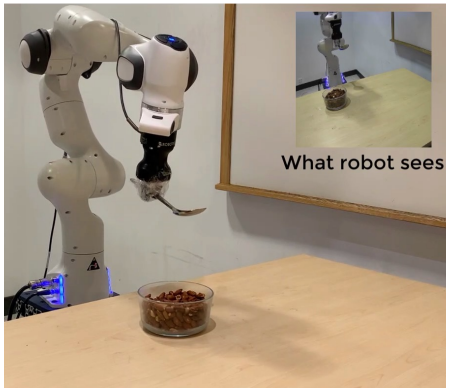

Scooping

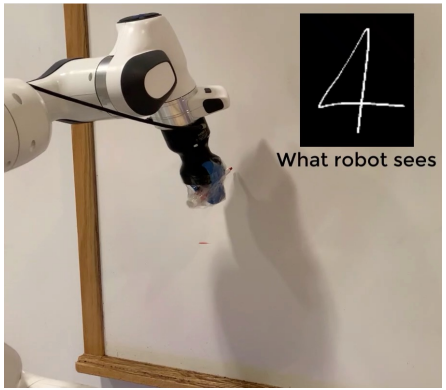

Writing

Obj 1: Learning to Assist from Raw Sensory Observations

More efficient RL algorithms from pixels

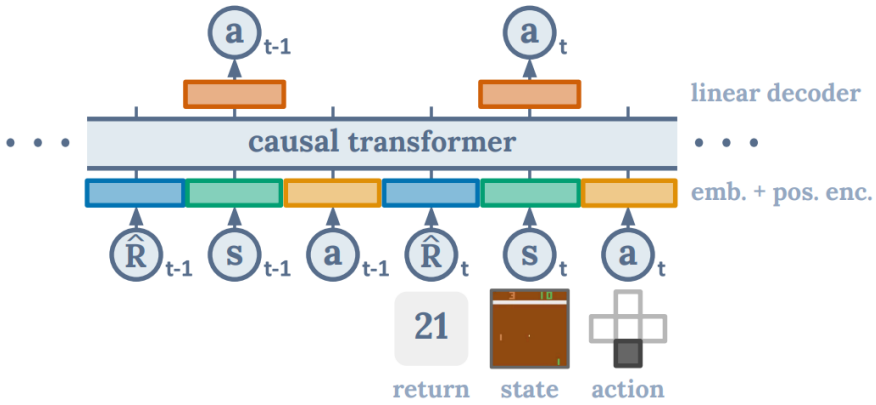

Decision Transformers
[NeurIPS 2021]

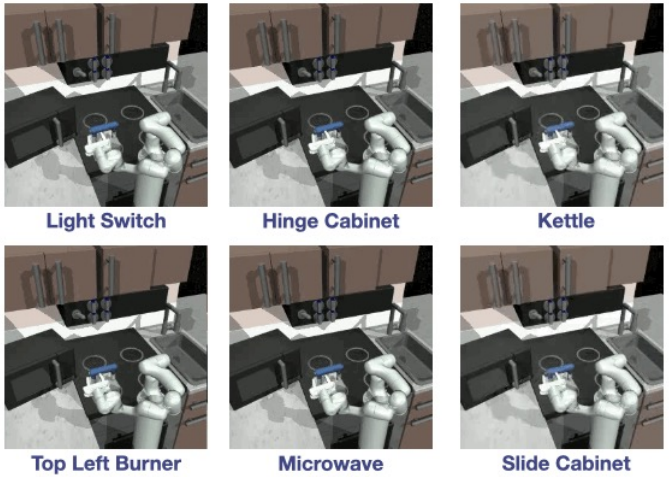

Robot Action Primitives
[NeurIPS 2021]

Obj 2: Simulation for Learning Assistance

Sim2Real Transfer from “Assistive Gym” – our physics simulation framework

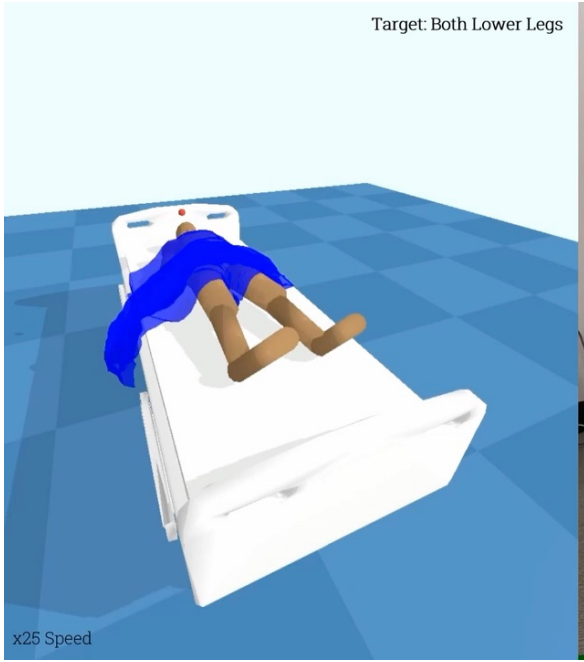

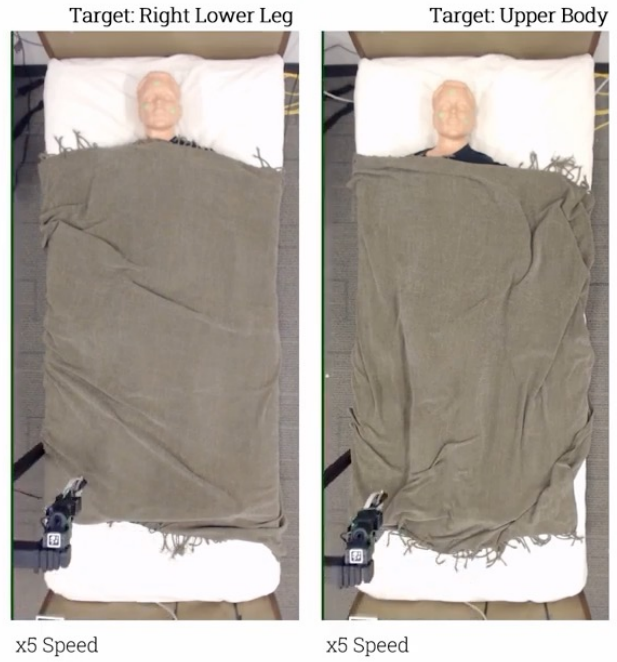

Bedding Manipulation

Obj 3: Customizable Robotic Assistance

Learning Rewards for Robot's Policy from Human Preferences

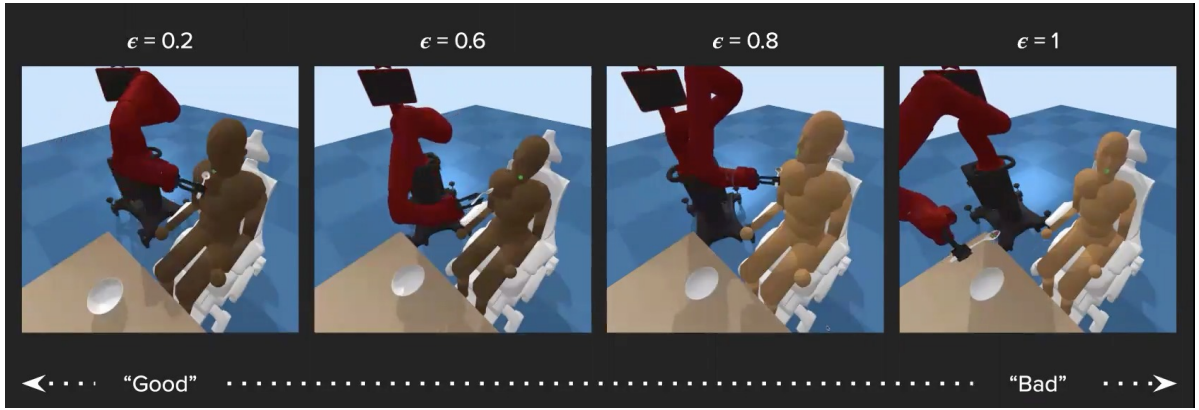

Preference variations in Feeding

Itch Scratching Assistance

Thank you! (reach out: dpathak@cs.cmu.edu)

- Towards more Generalizable One-shot Visual Imitation Learning. Zhao Mandi, Fangchen Liu, Kimin Lee, Pieter Abbeel. ICRA 2022.
- Decision Transformer: Reinforcement Learning via Sequence Modeling. Lili Chen*, Kevin Lu*, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Michael Laskin, Pieter Abbeel, Aravind Srinivas, Igor Mordatch. Neural Information Processing Systems (NeurIPS), Virtual, December 2021.
- Improving Computational Efficiency in Visual Reinforcement Learning via Stored Embeddings. Lili Chen, Kimin Lee, Aravind Srinivas, Pieter Abbeel. NeurIPS 2021.
- Skill Preferences: Learning to Extract and Execute Robotic Skills from Human Feedback. Xiaofei Wang, Kimin Lee, Kourosh Hakhamaneshi, Pieter Abbeel, Michael Laskin. CoRL 2021.
- PEBBLE: Feedback-Efficient Interactive Reinforcement Learning via Relabeling Experience and Unsupervised Pre-training. Kimin Lee*, Laura Smith*, Pieter Abbeel. International Conference on Machine Learning (ICML), Virtual, July 2021.
- State Entropy Maximization with Random Encoders for Efficient Exploration. Younggyo Seo, Lili Chen, Jinwoo Shin, Honglak Lee, Pieter Abbeel, Kimin Lee. ICML 2021.
- SUNRISE: A Simple Unified Framework for Ensemble Learning in Deep Reinforcement Learning. In the proceedings of the International Conference on Machine Learning (ICML), Virtual, July 2021. Kimin Lee, Michael Laskin, Aravind Srinivas, Pieter Abbeel.
- A study of causal confusion in preference-based reward learning. IROS (in submission), 2022. J. Tien, J.Z. He, Z. Erickson, A.D. Dragan, and D. Brown.
- B-Pref: Benchmarking Preference-Based Reinforcement Learning. Neural Information Processing Systems (NeurIPS), 2021. K. Lee, L. Smith, A.D. Dragan, and P. Abbeel.
- Assisted Robust Reward Design. Conference on Robot Learning (CoRL), 2021. J.Z. He, A.D. Dragan.
- Unsupervised Learning of Visual 3D Keypoints for Control. International Conference on Machine Learning (ICML) 2021. Boyuan Chen, Pieter Abbeel, Deepak Pathak
- Hierarchical Neural Dynamic Policies. Robotics: Science and Systems (RSS) 2021. Shikhar Bahl, Abhinav Gupta, Deepak Pathak
- Accelerating Robotic Reinforcement Learning via Parameterized Action Primitives. NeurIPS 2021. Murtaza Dalal, Deepak Pathak*, Ruslan Salakhutdinov*
- Functional Regularization for Reinforcement Learning via Learned Fourier Features. Neural Information Processing Systems (NeurIPS) 2021. Alexander C. Li, Deepak Pathak
- RB2: Robotic Manipulation Benchmarking with a Twist. Neural Information Processing Systems Datasets and Benchmark Track (NeurIPS) 2021. Sudeep Dasari, Jianren Wang, Joyce Hong, Shikhar Bahl, Abitha Thankaraj, Karanbir Chahal, Berk Calli, Saurabh Gupta, David Held, Lerrel Pinto, Deepak Pathak, Vikash Kumar, Abhinav Gupta
- Generalization in Dexterous Manipulation via Geometry-Aware Multi-Task Learning. In submission, 2022. Huang, W., Mordatch, I., Abbeel, P. and Pathak, D.