Co-GAIL: Learning Diverse Strategies for Human-Robot Collaboration

Chen Wang, Claudia Pérez-D'Arpino, Danfei Xu, Li Fei-Fei, C. Karen Liu, Silvio Savarese

Stanford University

Website and Code: https://sites.google.com/view/cogail/home

Overview

Having a **robot assistant** at home that could seamlessly assist us with daily activities is a long-sought dream.

(a) 2D-Fetch-Quest (b) HR-Handover (c) HR-SeqManip

Challenges for human-robot collaboration:

- It is **unsafe and sample-inefficient** to directly train human-robot collaboration in real-world.
- Human might have **diverse non-stationary** motions.

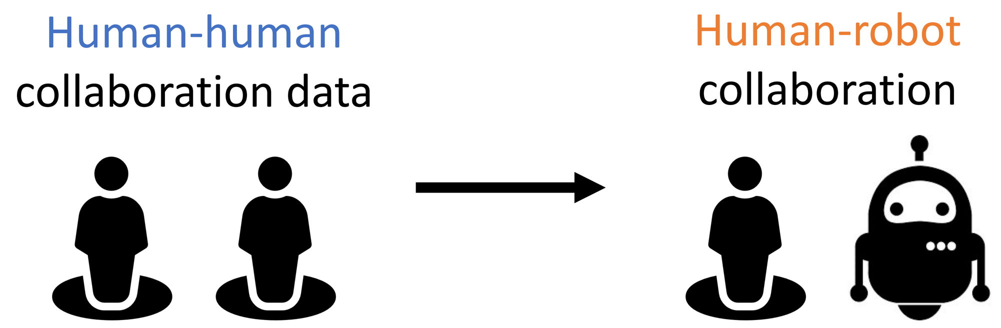

Human collaboration demos are **safe** and **easy to collect** and contains **rich diverse intents and movements**. Introducing **Co-GAIL**, an imitation learning framework to teach diverse human-robot collaboration skills from human-human collaboration demonstrations.

Co-policy Model

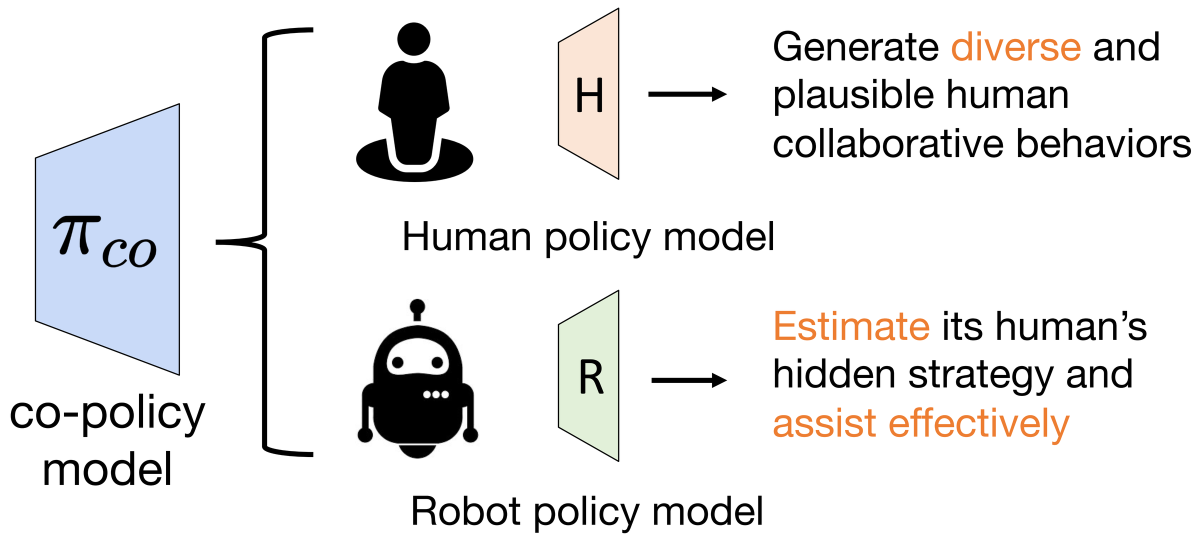

1. **Human motion generator**: Learns to generate diverse human strategies for training the robot.
2. **Robot policy**: Learns to predict hidden human strategy and outputs collaborative actions.

Co-optimization for learning collaborative skills

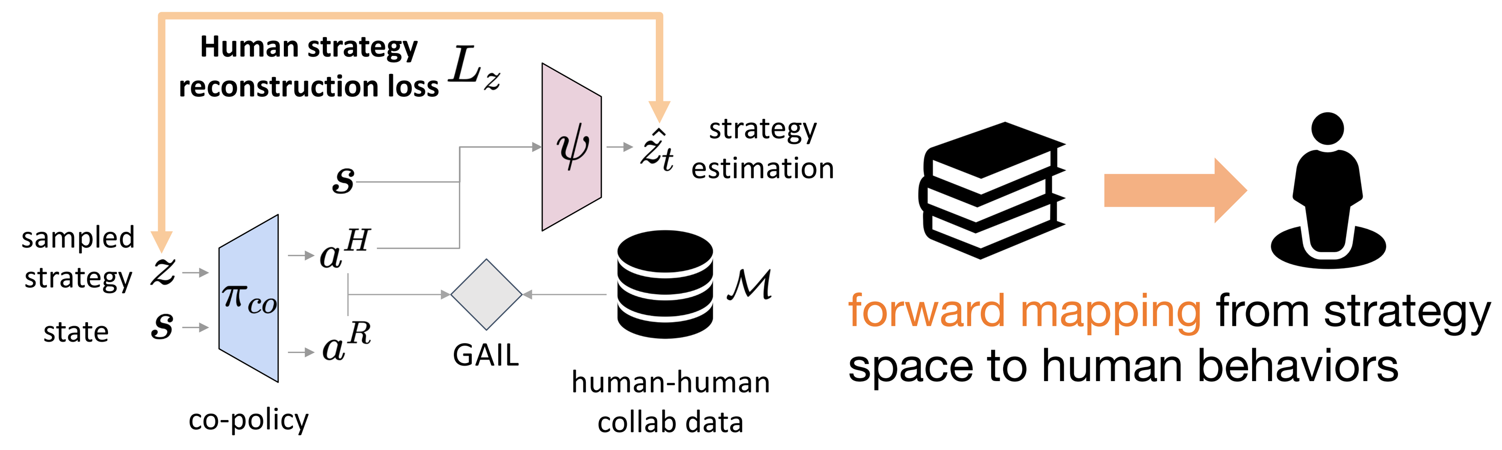

Forward mapping guarantees that each strategy code could lead to a unique type of human behavior.

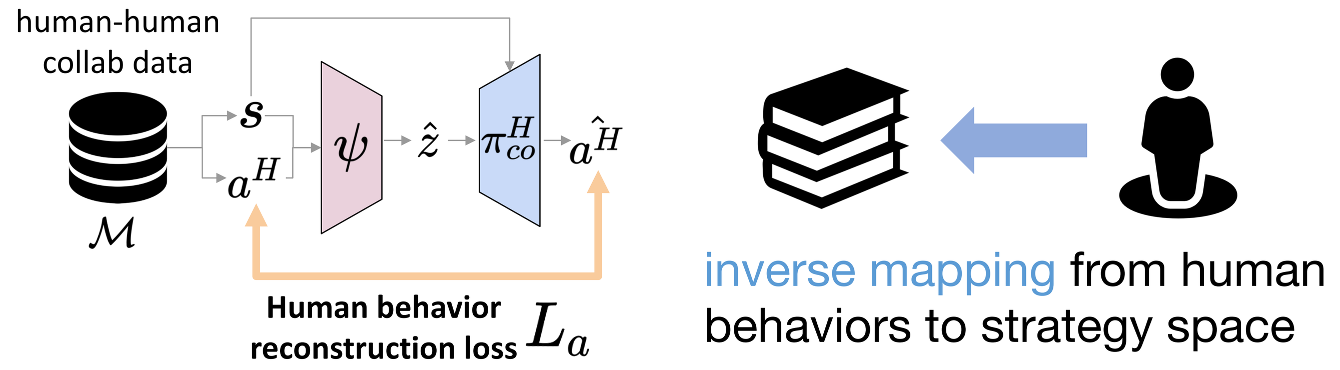

Inverse mapping objective encourages the model to explore diverse human behavior from the data.

Experiments results

Replay evaluation results with unseen human motions

	Replay 2D-FetchQ.	Replay HR-H.	Replay HR-S.
BC-single	21.1 ± 1.9	13.3 ± 3.3	15.0 ± 2.3
BC-GAIL	27.2 ± 5.0	14.1 ± 4.6	16.6 ± 7.1
MA-GAIL[26]	30.0 ± 3.4	23.7 ± 5.3	21.0 ± 3.7
MA-InfoGAIL[25]	40.0 ± 1.7	27.8 ± 4.4	26.6 ± 8.8
DIAYN[32]	44.4 ± 8.4	22.7 ± 1.5	18.3 ± 7.1
Co-GAIL (ours)	$\mathbf{53.3 \pm 4.4}$	$\mathbf{43.9 \pm 6.3}$	$\mathbf{40.0 \pm 8.2}$

Qualitative results on diverse human behavior exploration

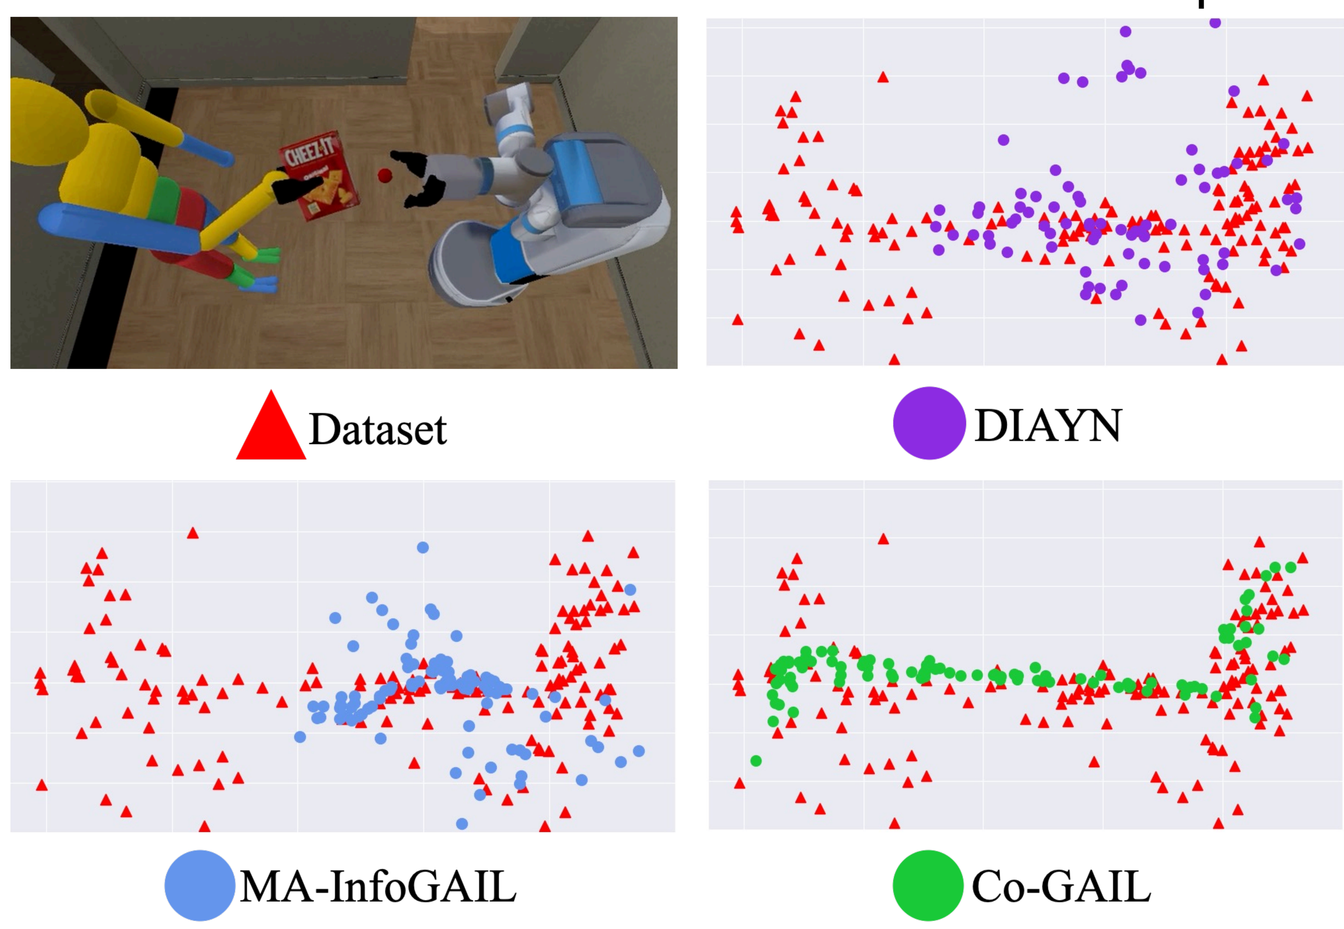

Visualization of the learned latent strategy space

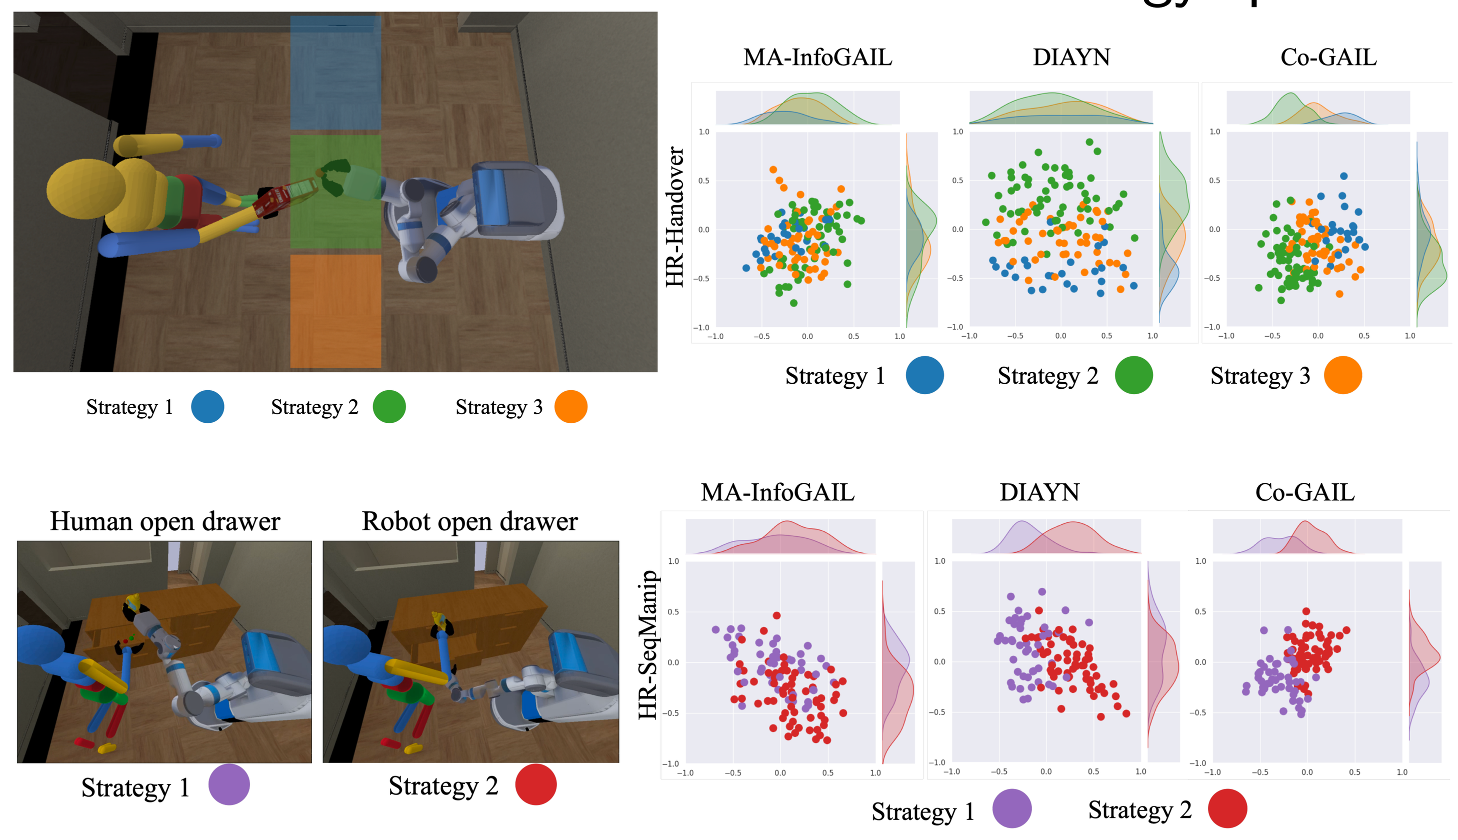

We observe a better separation of latent strategy learned using Co-GAIL than the other two methods.

Proof-of-concept analysis with real users

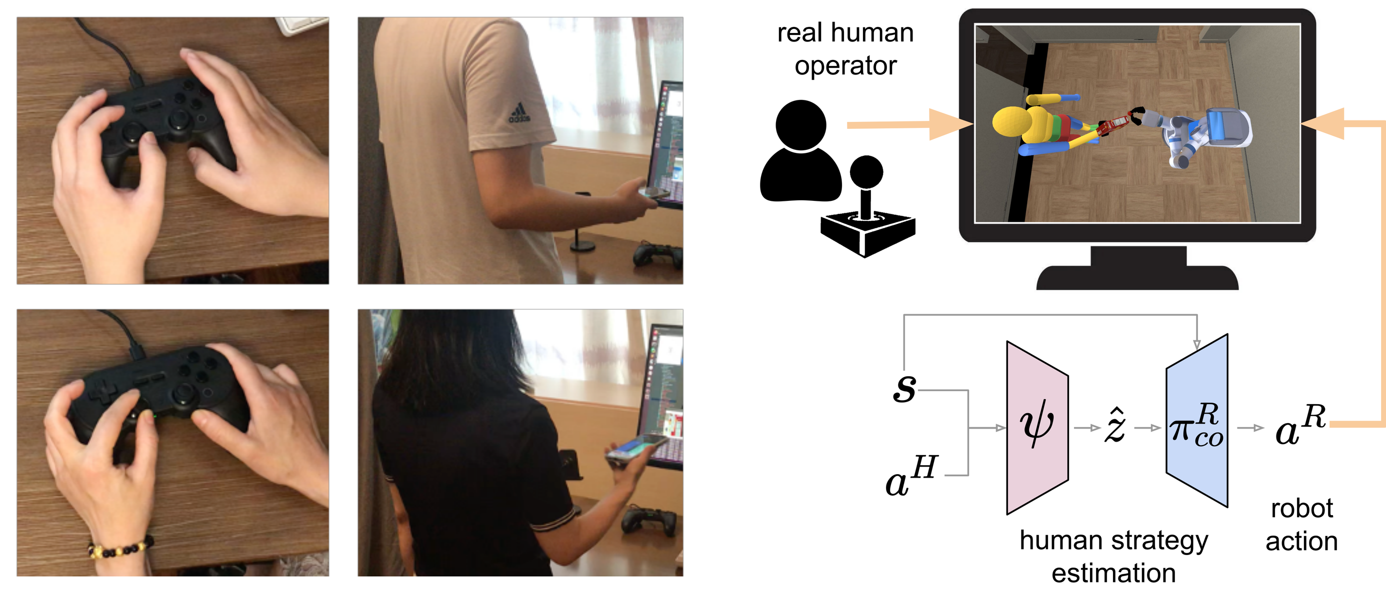

Setup: We invite four human users to test the robot trained with different algorithms. Users interacted with all methods in randomized order and the success rate over 20 trials is reported.

User id	2D-FetchQ. 1	2	3	4	HR-H. 1	2	3	4	HR-S. 1	2	3	4
BC-single	70	65	50	65	20	65	10	20	10	5	10	15
MA-GAIL[24]	65	50	45	50	50	35	25	20	30	25	35	25
MA-InfoGAIL[25]	85	80	40	50	55	**90**	65	25	35	40	40	50
Co-GAIL	**100**	**90**	**85**	**80**	**75**	70	**70**	**75**	**60**	**50**	**65**	**70**

Results: Co-GAIL consistently outperforms the other methods with different users and environments.

References:

[1] Shen B, Xia F, Li C, et al. iGibson, a Simulation Environment for Interactive Tasks in Large Realistic Scenes, IROS-21
[2] Tung A, Wong J, Mandlekar A, et al. Learning multi-arm manipulation through collaborative teleoperation, ICRA-21
[3] Ho J, Ermon S. Generative adversarial imitation learning, NeurIPS, 2016
[4] Li Y, Song J, Ermon S. Infogail: Interpretable imitation learning from visual demonstrations, NeurIPS, 2017
[5] Eysenbach B, Gupta A, Ibarz J, et al. Diversity is all you need: Learning skills without a reward function, ICLR, 2019.