Following the Human Thread in Social Navigation

Luca Scofano[1]*, Alessio Sampieri[1,4]*, Tommaso Campari[2]*, Valentino Sacco[1]*,
Indro Spinelli[1], Lamberto Ballan[3], Fabio Galasso[1]
[1]Sapienza University of Rome, [2]Fondazione Bruno Kessler,
[3]University of Padova, [4]ItalAI

Code available!

Objective: Build a model allowing an agent to adapt to human movements in shared environments, approaching to assist when needed and stepping back to avoid collisions.

Core Idea: Infer social dynamics from the agent's past actions and states, without relying on hard-to-obtain privileged information.

Social Dynamics Adaptation

We introduce Social Dynamics Adaptation (SDA): a two-stage reinforcement learning framework enabling robots to adapt to human movement.

- **Stage 1: t**he model learns a latent representation of social dynamics from fully observable human trajectories (privileged information) to condition robot navigation.
- **Stage 2:** SDA infers social dynamics solely from the robot's past actions and states, allowing real-time adaptation without external (privileged) trajectory data.

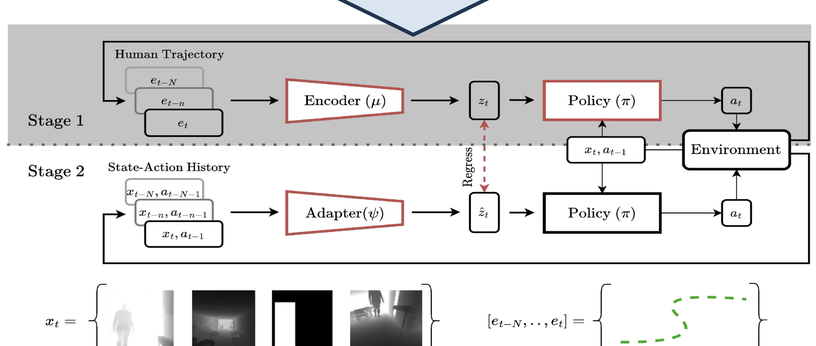

Results

The SDA model was evaluated on the Habitat 3.0 simulator, showing state-of-the-art performance in finding and following humans:

- **Success (S):** Increased from 76% to 91%. **(+15%)**
- **Standardized Path Success (SPS):** Improved from 0.34 to 0.45. **(+11%)**
- **Following Success (F):** Increased from 0.29 to 0.39. **(+10%)**

Habitat 3.0

Models	hGPS	traj.	S ↑	SPS ↑	F ↑	CR ↓	ES ↑
Heuristic Expert	-	-	1.00	0.97	0.51	0.52	-
Baseline	GT		$0.97^{\pm 0.00}$	$0.65^{\pm 0.00}$	$0.44^{\pm 0.01}$	$0.51^{\pm 0.03}$	$0.55^{\pm 0.01}$*
Baseline+Proximity	GT		$0.97^{\pm 0.01}$	$0.64^{\pm 0.00}$	$0.57^{\pm 0.01}$	$0.58^{\pm 0.03}$	$0.63^{\pm 0.02}$
SDA - S1		GT	$0.92^{\pm 0.00}$	$0.46^{\pm 0.01}$	$0.44^{\pm 0.02}$	$0.61^{\pm 0.02}$	$0.50^{\pm 0.01}$
Baseline			$0.76^{\pm 0.02}$	$0.34^{\pm 0.01}$	$0.29^{\pm 0.01}$	**0.48**$^{\pm 0.03}$	$0.40^{\pm 0.02}$*
Baseline+Proximity			$0.85^{\pm 0.02}$	$0.41^{\pm 0.02}$	$0.37^{\pm 0.01}$	$0.58^{\pm 0.02}$	$0.41^{\pm 0.01}$
SDA - S2			**0.91**$^{\pm 0.01}$	**0.45**$^{\pm 0.01}$	**0.39**$^{\pm 0.01}$	$0.57^{\pm 0.02}$	**0.43**$^{\pm 0.02}$

Towards Real World Scenarios

Update Rate	S ↑	SPS ↑	F ↑	CR ↓	ES ↑
1 *(Proposed)*	**0.91**$^{\pm 0.01}$	**0.45**$^{\pm 0.01}$	$0.39^{\pm 0.01}$	**0.57**$^{\pm 0.02}$	$0.43^{\pm 0.02}$
1/2	$0.87^{\pm 0.01}$	$0.39^{\pm 0.01}$	**0.44**$^{\pm 0.01}$	$0.63^{\pm 0.02}$	**0.48**$^{\pm 0.02}$
1/100	$0.85^{\pm 0.01}$	$0.38^{\pm 0.01}$	$0.43^{\pm 0.01}$	$0.64^{\pm 0.03}$	$0.46^{\pm 0.01}$

SDA performance considering missing readers

SDA	S ↑	SPS ↑	F ↑	CR ↓	ES ↑
Habitat 3.0	**0.91**$^{\pm 0.01}$	**0.45**$^{\pm 0.01}$	**0.39**$^{\pm 0.01}$	$0.57^{\pm 0.02}$	$0.43^{\pm 0.02}$
Habitat 3.0 + ORCA	$0.90^{\pm 0.01}$	$0.43^{\pm 0.02}$	$0.38^{\pm 0.01}$	**0.37**$^{\pm 0.01}$	**0.48**$^{\pm 0.01}$

Comparison of SDA versus the variant with ORCA

Key Contributions

Innovative Social Inference: Introduces the novel SDA framework, enabling robots to infer social cues from their own past actions and states.

Dual-Phase Learning: Employs a two-stage RL approach: first building a latent representation of human movement, then adapting it in real time.

Enhanced Navigation: Evaluated on Habitat 3.0, SDA significantly improves the robot's ability to detect and follow humans.

Latent Analysis

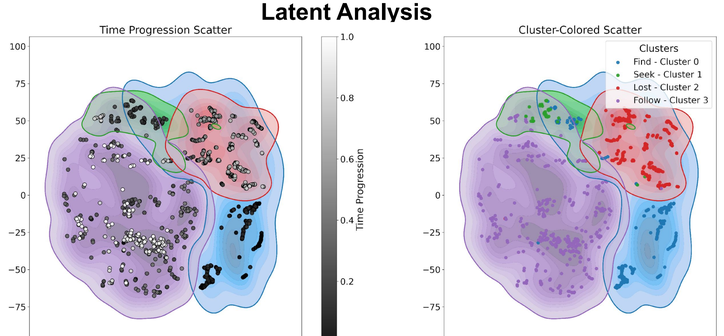

Acknowledgements. We are grateful to Panasonic for partially supporting this work. We acknowledge the financial support from the PNRR MUR project PE0000013-FAIR and from the Sapienza grant RG123188B3EF6A80 (CENTS). Thanks CINECA and the ISCRA initiative for high-performance computing resources and support.