ARTIFICIAL INTELLIGENCE

PNRR/FAIR - Transferable and efficient learning across task, environment, and embodiment structures

Funded By	Dipartimento DAUIN MINISTERO DELL'UNIVERSITA' E DELLA RICERCA [P.iva/CF:97429780584] Politecnico di TORINO [P.iva/CF:00518460019]

Supervisor	TOMMASI TATIANA - tatiana.tommasi@polito.it

Contact	CAMORIANO RAFFAELLO - raffaello.camoriano@polito.it TOMMASI TATIANA - tatiana.tommasi@polito.it

Context of the research activity	The main project goal is the design of efficient learning methods for high-capacity model training, transfer, and prediction with a focus on safe embodied learning systems (e.g., robots and avatars). Considered approaches will include: i) model compression and approximation methods reducing computational and labeling costs while retaining high accuracy, and ii) constraining the model to low-dimensional structures (e.g., manifolds) arising from the physics of the target problem. Progetto finanziato nell'ambito del PNRR - PNRR M4C2, Investimento 1.3 - Avviso n. 341 del 15/03/2022 - PE0000013 Future Artificial Intelligence Research (FAIR) - CUP E13C22001800001

	Classical learning methods for robotic perception, planning, and control tend to target specific skills and embodiments, due to the difficulties in extracting transferable and actionable representations that are invariant to physical properties of the environment and of the agent (e.g., a physical or simulated robot). However, the performance of such specialized agents can be limited by low model capacity and training on relatively few examples. This can be particularly problematic when tackling complex and long-horizon tasks for which the cost of large-scale data collection on a single robot can be prohibitively high and the complexity of the policy to be learned might benefit from a more expressive function class (i.e., with a larger number of parameters). Conversely, recent high-capacity, highly flexible machine learning models, such as vision transformers and large multimodal models, proved their worth in less constrained domains such as computer vision and natural language processing. In such scenarios, pre-training on large and diverse datasets is

Objectives	possible due to web-scale data availability. This results in rich “generalist” pre-trained models enabling model fine tuning and adaptation to specific target tasks with large savings in terms of target data collection and positive transfer to new tasks and visual appearances. A growing research stream investigates the extension of such high-capacity models to robotic tasks to enable complex skill learning across embodiments and modalities, thanks to the flexibility of architectures such as GATO [1]. RoboCat [2] demonstrates how these models can solve complex robotic manipulation tasks with visually defined goals, while Open X Embodiment [3] demonstrates positive transfer for task goals specified in natural language. Octo further extends this concept by supporting multimodal goal definitions [4], while AutoRT [5] also supports multi-robot coordination. Large language models can also be employed to guide exploration and automate reward design for reinforcement learning [6]. However, these methods rely on very large numbers of parameters (i.e., in the order of billions), rendering model storage and real-time inference a major challenge. This is a relevant roadblock when local execution on limited robotic hardware is required, as is often the case in open-world unstructured environments. Some of the most advanced multi-embodiment models (e.g., RT-2-X [3]) are so extensive that they cannot be stored locally and require communication with cloud environments to perform inference. Even more so when model fine-tuning, open-ended, or continual learning are required for tackling new tasks. Impractical computational and communication costs and catastrophic forgetting of previous tasks indeed represent a major challenge. The objective of this project is the development of efficient methods for training, transfer, and inference of powerful learning models for embodied and robotic tasks. Several approaches will be investigated, including, among others, the use of model compression, recent fine-tuning methods which proved to reduce the cost of execution of robotic policies (i.e., RT-2-X) from quadratic to linear while retaining performance levels [7], and approximation methods to reduce the number of parameters while retaining approximation capacity [8]. Further work will also be devoted to developing novel learning algorithms constrained to low-dimensional structures (e.g., manifolds) arising from the geometric and physical properties of the target tasks and embodiments, with the goal of improving model efficiency and safety [9, 10]. Planned publication venues include the major international AI, ML, robotics, and/or computer vision journals (JMLR, T-RO, RA-L, TPAMI, IJRR, IJCV, TNNLS, etc.) and conferences (e.g. NeurIPS, ICML, ICLR, AISTATS, AAAI, IROS, ICRA, CoRL, CVPR, ICCV, ECCV, etc.). References [1] Reed, Scott, et al. "A generalist agent." TMLR, 2022. [2] Bousmalis, Konstantinos, et al. "RoboCat: A Self-Improving Foundation Agent for Robotic Manipulation." TMLR, 2023. [3] Padalkar, Abhishek, et al. "Open x-embodiment: Robotic learning datasets and rt-x models." arXiv preprint arXiv:2310.08864 (2023). [4] Team, Octo Model, et al. "Octo: An open-source generalist robot policy." 2023. [5] Ahn, Michael, et al. "Autort: Embodied foundation models for large scale orchestration of robotic agents." arXiv, 2024.

	[6] M. Kwon, S. M. Xie, K. Bullard, and D. Sadigh, “Reward design with language models,” ICLR, 2023. [7] Leal, Isabel, et al. "SARA-RT: Scaling up Robotics Transformers with Self-Adaptive Robust Attention." arXiv, 2023. [8] Xiong, Yunyang, et al. "Nyströmformer: A nyström-based algorithm for approximating self-attention." AAAI, 2021. [9] Liu, Puze, et al. "Robot reinforcement learning on the constraint manifold." CoRL, 2022. [10] Duan, Anqing, et al. "A structured prediction approach for robot imitation learning." IJRR, 2023.

Skills and competencies for the development of the activity	We seek candidates highly motivated to conduct methodological research in ML and robotics. An excellent background in ML is required, covering theory and software. Proficiency with Python and ML, robotics, or CV frameworks are a must. Strong communication skills, self-motivation, proven teamwork experience, and independence are necessary. A track record and/or certifications of fluent speaking and technical writing in English are required. Prior research experience is highly appreciated.