

# A Compositional Paradigm for Foundation Models: Towards Smarter Robotic Agents

Luigi Quarantiello\*<sup>§</sup>, Elia Piccoli\*, Jack Bell\*, Malio Li\*, Giacomo Carfi\*, Eric Nuertey Coleman\*  
Gerlando Gramaglia\*, Lanpei Li\*, Mauro Madeddu\*, Irene Testa\*, Vincenzo Lomonaco<sup>†</sup>

\*University of Pisa, Italy

<sup>†</sup>LUISS University, Rome, Italy

<sup>§</sup>Corresponding author: luigi.quarantiello@phd.unipi.it

**Abstract**—The birth of Foundation Models brought unprecedented results in a wide range of tasks, from language to vision, to robotic control. These models are able to process huge quantities of data, and can extract and develop rich representations, which can be employed across different domains and modalities. However, they still have issues in adapting to dynamic, real-world scenarios without retraining the entire model from scratch. In this work, we propose the application of Continual Learning and Compositionality principles to foster the development of more flexible, efficient and smart AI solutions.

## I. INTRODUCTION

Recent advancements in Artificial Intelligence have led to the emergence of Foundation Models (FMs). Such architectures [1], leveraging billions of parameters, terabytes of training data and the most powerful (and energy-hungry) machines in the world, are able to encode broad knowledge and skills in their hidden spaces. Ultimately, these features are employed to solve the most diverse tasks, encompassing different modalities and objectives. The most impressive results, in particular from the user standpoint, comes from the large language models powering different chatbots, which excel at web search, code/text generation and brainstorming. They show an incredible fluency (at least for widely spoken languages) and a great capability in elaborating knowledge and information from multiple sources. Beyond language, FMs have demonstrated superior capabilities in other settings too, such as Computer Vision (*e.g.* ViT for image classification), and multi-model learning (*e.g.* CLIP for joint image-text understanding). This led to networks with richer representations, which are more flexible and applicable to multiple scenarios, ultimately showing better performance.

However, recent releases show that noticeable parameters and complexity increases lead only to marginal performance improvements. In this work, we propose a alternative vision on AI that encompasses Continual Learning and Compositionality, to enhance FMs with an increased adaptability and flexibility to address more complex, dynamic scenarios.

## II. LARGE MODELS FOR ROBOTICS

Given their strong performance across different tasks, Transformer-based models have been applied to robotics too. Early efforts, such as Robotics Transformer (RT), proved the feasibility of high-capacity, data-driven control policies. RT-1 [4] demonstrated scalable generalization from large, diverse

real-robot datasets, while RT-2 [16] injected internet-scale vision-language priors directly into action policies, establishing the “*foundation model paradigm*” for robotic control. In order to further scale these agents, the researchers focused on different aspects. A first key contribution was the release of Open X-Embodiment dataset [11], that unifies demonstrations across 22 robot types to enable cross-embodiment transfer. On top of that, a great effort has been put to in the development of benchmarks, *e.g.* Meta-World [15], and realistic yet highly efficient, GPU accelerated simulators, like ManiSkill3 [14]. These tools enable efficient training and systematic evaluation of generalization and multi-task performance across different robotics tasks. More recently, given the increased computational power and huge breakthroughs of language models, a parallel—or perhaps orthogonal—line of research has focused on the models architecture. Some pioneering work [6], [13] tried to leverage CLIP’s embeddings to align natural language instruction in the robotics domain. As pre-trained models evolved, research aimed at modifying architectures to have specific components that handle different data modalities. While scaling architectures to billions of parameters helped in achieving better agents performance, a huge effort was spent to adapt general pre-trained models to specific robotics tasks. This effort lead to the development of Vision Language Action (VLA) models, in particular OpenVLA [9]. By fine-tuning a 7B-Llama backbone, the agent obtains a 16.5% increase in absolute task success rate compared to RT-2-X.

## III. THE NEED FOR CONTINUAL LEARNING AND COMPOSITIONALITY

In the previous sections, we examined Transformer-based models and their most recent applications, particularly in robotics. As it often happens in Computer Science, many of these improvements has been driven by scaling, increasing model size, training data, or both, at the cost of growing computational and energy demands. This alone has yielded impressive results: language models are able to handle complex requests and process longer inputs, offering a more human-like experience and demonstrating reduced biases and improved safety, while robotic agents have gained more complex, multi-skill capabilities. However, we are at a point of diminishing returns, where enlarging models, datasets or computational resources is not enough to obtain noticeable results [10].

TABLE I  
IMAGE CLASSIFICATION TASK. OUR COMPOSITION OF ADAPTERS GREATLY OUTPERFORMS COMPETITORS, WHILE BEING MORE EFFICIENT

Method	Accuracy (%)	Training Time (s)
SD-LoRA	47.56 $\pm$ 3.77	318.22
InfLoRA	36.02 $\pm$ 3.26	196.89
HAM ( <i>Ours</i> )	<b>55.17</b> $\pm$ 1.04	<b>170.61</b>

TABLE II  
ROBOTIC MANIPULATION TASK. OUR METHOD OUTPERFORMS MORE COMPLEX BASELINES, WITH REDUCED TRAINING TIMES

Method	Reward/Step	Success Rate	Training Time
InstructRL	0.14	0.0	40 h
OpenVLA	0.10	0.0	92 h
WSA ( <i>Ours</i> )	<b>0.60</b>	<b>0.91</b>	<b>14 h</b>

Furthermore, despite being trained on massive datasets, FMs still struggle to generalize across tasks over time, especially when data distributions shift. Their performance relies on statistical correlations rather than deeper reasoning, so scaling data [2] or using Parameter-Efficient Fine-Tuning methods only yields limited improvements. Applying Continual Learning principles and methods to FMs can represent a viable solution to give such models a higher degree of flexibility, allowing for updates on the pre-training distribution and efficient adaptation to domain-specific tasks [3].

For these reasons, we argue that advancing to the next generation of AI requires a paradigm shift towards systems that are more flexible, adaptable and computationally efficient. Compositionality—*divide and conquer*—is a key enabling factor of the next revolution. It allows multiple specialized models to coexist, share knowledge and collectively solve complex tasks more effectively than single models. Such model compositions are inherently dynamic; by changing the way in which the models are orchestrated, such systems could easily adapt to ever-changing scenarios, being more flexible than a monolithic solution. Lastly, such combinations may be performed with little to no training phases, exploiting FMs full potential while being more efficient than the static counterparts.

#### IV. EXPERIMENTAL RESULTS

We explored the idea of Continual Learning and Compositionality in different domains, performing some preliminary experiments.

For image classification tasks, we studied the adaptation of the ViT model to a stream of tasks with the composition of multiple LoRA adapters [7]. For each experience, we train a separate module, which is then merged with previously learned LoRAs to obtain a single final model. We employ a hierarchical merging process: a first combination is performed within LoRAs grouped by similarity, while group adapters are later merged into a single module, so to have a unified model capable of solving every seen task [5]. We found that such method is effective in reducing the merging process interference. Table I shows the results obtained on CUB200 dataset with 50 tasks: our methods achieves an higher accuracy, while being faster in the training process.

In the context of robotics applications, we propose a simple architecture that using small adapters and attention mechanism can dynamically scale agents’ pre-trained components while adapting to different contexts [12]. The results show that agents using small off-the-shelf pre-trained models and a small computational budget are able to learn to effectively solve robotics manipulation task while more complex counterparts as OpenVLA fail (given the same computational constraints). The results obtained are presented in Table II, where its clear that our method, while being much more lightweight, obtains an higher reward per step on average, and achieves a far superior success rate than the competitors.

#### V. CONCLUSION

For FMs to reach their full potential and approach agentic AI, they must continuously adapt to evolving scenarios and generalize skills in open world contexts. In this work, we show that composing multiple models enables strong performance with limited resources. Robotic agents must learn low level skills from both data and their interactions with the world in order to perform precise and specific behaviors. At the same time, agents need to be able to abstract from these policies and reason over their capabilities in order to come up with plans [8]. Learning a precise model of the world is key to understand how to interact with objects and humans, to continuously evolve agents abilities and adapt to sudden changes in the world without catastrophic forgetting and dangerous behaviors from safety misalignment.

#### REFERENCES

- [1] Yang M et al. Learning interactive real-world simulators[J]. arXiv preprint arXiv:2310.06114, 2023.
- [2] Chen Zhe et al. "Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling." arXiv preprint arXiv:2412.05271, 2024.
- [3] Jack Bell et al. The future of continual learning in the era of foundation models: Three key directions.arXiv preprint arXiv:2506.03320, 2025.
- [4] Anthony Brohan et al. RT-1: robotics transformer for real-world control at scale. Robotics: Science and Systems XIX, 2023.
- [5] Eric Nuertey Coleman et al. Ham: Hierarchical adapter merging for scalable continual learning. arXiv preprint arXiv:2509.13211, 2025.
- [6] Pierre-Louis Guhur et al. Instruction-driven history-aware policies for robotic manipulations. Conference on Robot Learning, 2022.
- [7] Edward J Hu et al. Lora: Low-rank adaptation of large language models. ICLR, 2022.
- [8] Physical Intelligence et al.  $\pi_{0.5}$ : a vision-language-action model with open-world generalization. abs/2504.16054, 2025.
- [9] Moo Jin Kim et al. Openvla: An open-source vision-language-action model. Conference on Robot Learning, 2024.
- [10] Charles Luo. Has llm reached the scaling ceiling yet? unified insights into llm regularities and constraints. arXiv preprint arXiv:2412.16443, 2024.
- [11] Abby O’Neill et al. Open x-embodiment: Robotic learning datasets and RT-X models : Open x-embodiment collaboration. ICRA, 2024.
- [12] Elia Piccoli et al. Combining pre-trained models for enhanced feature representation in reinforcement learning. CoLLAs, 2025.
- [13] Mohit Shridhar et al. Cliport: What and where pathways for robotic manipulation. Conference on Robot Learning, 2021.
- [14] Stone Tao et al. Maniskill3: GPU parallelized robotics simulation and rendering for generalizable embodied AI. abs/2410.00425, 2024.
- [15] Tianhe Yu et al. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. Conference on Robot Learning, 2019.
- [16] Brianna Zitkovich et al. RT-2: vision-language-action models transfer web knowledge to robotic control. Conference on Robot Learning, 2023.