

Position Talk: Building a Curious Agent that Learns to Plan with Images Schemas

Henri-Jacques Geiss^{1,*}, Alejandro Agostini¹ and Justus Piater¹

¹Department of Computer Science, University of Innsbruck, Austria

Keywords

Image Schemas, Abstraction Learning, Reinforcement Learning, Robotics, Intrinsic Motivation

1. Presentation Abstract

As children we seem to effortlessly pick up everyday physical concepts such as the effect and direction of gravity or the ability of a glass to contain a liquid. Once internalized, we can use these abstractions to intuitively plan in new scenarios and by that achieve zero-shot task solving. Naturally, as *robotic* agents are increasingly deployed in complex everyday environments such qualities are also desired for their control.

While recently-presented, end-to-end Vision-Language-Action (VLA) models are performant on a variety of tasks and even show hints of emergent capabilities [1], they fail to include a conceptual, mechanistic understanding of the physical regularities of their environment, highlighting the need for models with explicit inductive biases for *abstraction*. Here, Image Schemas represent a promising candidate for serving as a semantic structure for these abstractions. Given their nature of being pre-verbal, spatially structured representations of the relations between entities and their movement through space [2], they indicate to be at the same time expressive enough to reason about the effect of actions and interactions, while also exhibiting enough generalizability to be used in symbolic planning and transfer learning.

Additionally, said VLA-models for sequential decision making are learned with internet-scale datasets [1], and thus turn the cognitive processes with which Image Schemas are theorized to ground language upside-down. Instead of acquiring basic, meaningful concepts through *intrinsically-motivated*, low-level sensorimotor exploration in their surroundings [2], they infer which actions to take based on the enormous language corpus they are (pre-)trained with. In contrast, in the field of Reinforcement Learning (RL), computational modelling of intrinsic drives like the search for novelty [3] or regularity [4] has seen abundant attention. With such inductive biases, RL agents have been shown to “curiously” explore intricate video game worlds or learn about object interactions [4].

To date, no approach has been presented in which a robot, without a pre-built action repertoire, can be placed in a previously-unknown environment, explore it through intrinsic motivation, and distill an abstract, image-schematic understanding of physical concepts. Additionally, forming such symbols is not an end in itself; rather we aim to use it for high-level reasoning and planning based on knowledge that is transferred from prior experience, from other agents, or that pre-exists.

Grounded on insights from cognitive science and computational methods for long-horizon task solving [5], we opt for a *hierarchical* framework with increasing level of abstraction at which sensorimotor information is processed to equip an agent with these capabilities (see Figure 1). Provided that the framework should be embodied in a physical robot platform, both the upwards path from low-level signals to an image-schematic description of the current state and the top-down flow of symbolic commands to continuous actions is included in our architecture. We present how existing approaches from the literature could be leveraged to solve the sub-problems

The 8th Image Schema Day, ISD8@AI*IA, November 27-29, 2024, Bolzano, Italy

*Corresponding and presenting author.

✉ henri.geiss@uibk.ac.at (H. Geiss); alejandro.agostini@uibk.ac.at (A. Agostini); justus.piater@uibk.ac.at (J. Piater)

ORCID 0000-0002-1898-3362 (J. Piater)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

that emerge in the respective layers, namely “Sensorimotor Exploration”, “Abstraction Formation” and “Symbolic Planning”. In the sensorimotor domain, this includes the discussion of recent advances in regards to exploratory priors that enable self-supervised learning, be it intrinsic rewards in model-based RL [4] or processes for *autotelic* goal generation in goal-conditioned RL [6].

For abstraction formation we will compare approaches that entail inductive biases for symbolic learning on different time-scales; while “Relational DeepSym” [7] learns symbolic attention-based relationships given pre-built action primitives. Many strategies exist that learn *mechanisms* of object interactions in world models [8]. We will also explore how to characterize dynamic concepts like sliding, rolling, and falling, along with the associated stable dynamics in a non-parametric fashion using nonlinear dynamic analysis [9]. Lastly, we will present a set of candidates for high-level reasoning given these acquired abstractions. While Task and Motion Planning (TAMP) represents a classical approach to leveraging physical abstractions for robot task execution [10], given the *semantics* that are carried in the connotation of the learned symbols with Image Schemas, more elaborate methods promise to pay dividends. For example, Image Schema Logic [11] could be used to reason about the effect of actions given abstract relations between objects.

In summary, in this position talk we will discuss how existing approaches for intrinsic motivation, symbol extraction and symbolic planning can be combined and extended to synthesize an agent that learns to plan with an image-schematic understanding of its environment.

References

- [1] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Choromanski, T. Ding, D. Driess, A. Dubey, C. Finn, et al., Rt-2: Vision-language-action models transfer web knowledge to robotic control, arXiv preprint arXiv:2307.15818 (2023).
- [2] J. M. Mandler, How to build a baby: Ii. conceptual primitives., Psychological review 99 (1992) 587.
- [3] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, R. Munos, Unifying count-based exploration and intrinsic motivation, Advances in neural information processing systems 29 (2016).
- [4] C. Sancaktar, J. Piater, G. Martius, Regularity as intrinsic reward for free play, Advances in Neural Information Processing Systems 36 (2024).
- [5] M. Eppe, C. Gumbsch, M. Kerzel, P. D. Nguyen, M. V. Butz, S. Wermter, Intelligent problem-solving as integrated hierarchical reinforcement learning, Nature Machine Intelligence 4 (2022) 11–20.
- [6] C. Colas, T. Karch, O. Sigaud, P.-Y. Oudeyer, Autotelic agents with intrinsically motivated goal-conditioned reinforcement learning: a short survey, Journal of Artificial Intelligence Research 74 (2022) 1159–1199.
- [7] A. Ahmetoglu, B. Celik, E. Oztop, E. Ugur, Discovering predictive relational object symbols with symbolic attentive layers, IEEE Robotics and Automation Letters (2024).

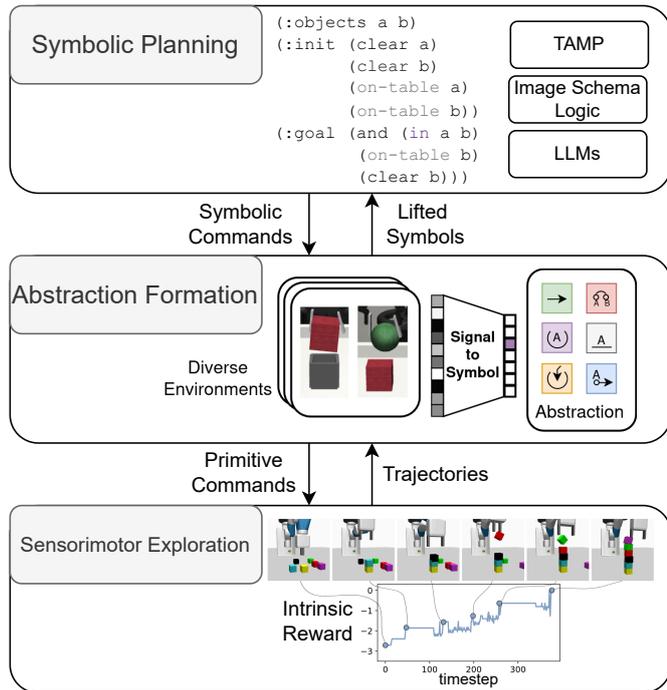


Figure 1: Overview over the proposed architecture. Intrinsic reward graphic is adapted from [4].

- [8] A. Lei, F. Nolte, B. Schölkopf, I. Posner, Compete and compose: Learning independent mechanisms for modular world models, arXiv preprint arXiv:2404.15109 (2024).
- [9] S. H. Strogatz, Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering, CRC press, 2018.
- [10] A. Agostini, J. Piater, Unified task and motion planning using object-centric abstractions of motion constraints, arXiv preprint arXiv:2312.17605 (2023).
- [11] M. M. Hedblom, M. Pomarlan, R. Porzel, R. Malaka, M. Beetz, Dynamic action selection using image schema-based reasoning for robots., in: JOWO, 2021.