

Perceiving and Acting Out of the Box

Fabien Lagriffoul¹[0000-0002-8631-7863]

Marjan Alirezaie¹[0000-0002-4001-2087]

Örebro University, 70182 Örebro, Sweden

`fabien.lagriffoul@oru.se`

`marjan.alirezaie@oru.se`

Abstract. This paper discusses potential limitations in learning in autonomous robotic systems that integrate several specialized subsystems working at different levels of abstraction. If the designers have anticipated what the system may have to learn, then adding new knowledge boils down to adding new entries in a database and/or tuning parameters of some subsystem(s). But if this new knowledge does not fit in predefined structures, the system can simply not acquire it, hence it cannot “think out of the box” designed by its creators. We show why learning out of the box may be difficult in integrated systems, hint at some existing potential approaches, and finally suggest that a better approach may come by looking at constructivist epistemology, with focus on Piaget’s schemas theory.

Keywords: Autonomous Robot · Learning · Piaget’s constructivist theory of knowledge.

1 Introduction

The typical approach for designing intelligent robots is *Divide and Conquer*: A team of experts with different domains of competence is formed, each of which shall develop the individual components (perception, actuation, reasoning) required for the system to achieve global intelligent behavior. Then, these subsystems need to be *integrated*, which requires some sort of interface between subsystems and global coordination mechanisms.

As compared to disembodied AI, there are several reasons why robotic systems need to integrate several subsystems. For functional reasons (e.g., the robot needs vision, path planning, dialogue), for engineering reasons (reusing existing software modules), or for reasoning upon different types of knowledge (causal, spatial, temporal) for which specific representations and reasoners have been developed. One may for instance represent causal relations by some action language and reason upon it with a satisfiability solver, while spatial matters may be represented by transformation matrices and reasoned upon with graph search.

These specialized subsystems perform well in their own domains, and for most of them, learning “variants” have been devised: vision systems can learn new categories, motion planners can learn from previous queries or by imitation,

task planners can learn heuristics, rule-based systems can learn by chunking, etc. This paper points out the issue of drawing meaningful relations between what is individually learned by the different subsystems of integrated systems and, furthermore, questions the capacity of current learning methods for robots to develop new representations and skills.

2 Learning in Integrated Systems

Let us consider the following example: a robot that should learn to manipulate various objects. The robot is standing in front of a table and its task is to clear the table, i.e, picking up any object from the table and releasing it in a nearby trash bin. We assume a classical *sense-plan-act* architecture [1] with three subsystems: deliberation, perception, and actuation, each of which being capable of learning. Initially, the subsystems have initial knowledge about bottles and glasses:

- the deliberative subsystem knows that the task can be solved using first the *pick_bottle* or *pick_glass* operators, and then the *release_trash* operator;
- the perception subsystem uses an Artificial Neural Network (ANN) which can label images from a camera as *glass* or *bottle*;
- the actuation subsystem has a database of motion primitives for *pick_bottle*, *pick_glass*, and *release*.

Then a novel object is introduced, e.g., a credit card, and the system should learn how to complete the task. The credit card cannot be grasped directly from the table, i.e., it has to be slid to the edge of the table before it can be grasped (see Fig. 1). We do not assume any particular learning methods, and suppose that after sufficient training, the subsystems have learned as follows:

- the deliberative subsystem has learned a new operator *grasp_34* (the system does not know it is a slide-and-grasp) for objects of type *type_23* (the system does not know it is a credit card);
- the perception subsystem has been trained to recognize a new class of objects (*type_23*);
- the actuation subsystem has added a new motion primitive to its database for *grasp_34*.

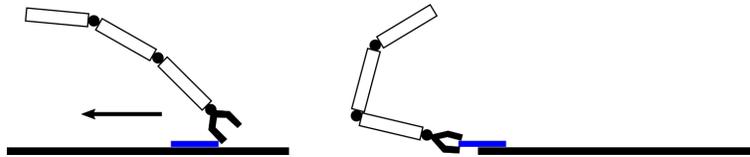


Fig. 1. Motion primitive for grasping a flat object.

The system can now deal with credit cards or (similar objects), but it has not learned anything about why *grasp_34* is appropriate for objects of type *type_23*.

Therefore if it is presented with a novel flat but different object, for instance a coin, it will not be able to reuse what it has learned about credit cards. Assuming that the perceptual subsystem has learned a feature related to “flatness”, and given previous experiences with flat and non-flat objects, the system could infer through statistical methods a correlation between flatness and a particular grasping strategy, hence using this knowledge to grasp unforeseen flat objects. But since the subsystems work –by construction– in different domains, this may simply not happen. What is learned by one subsystem is not necessarily relevant for other subsystems. Unless a human designer has anticipated which features may be of interest and built them in.

3 Horizontal and Vertical Learning

Learning is a general function that may take a variety of forms. For subsequent discussion, we introduce an informal distinction between two types of learning processes: *horizontal* and *vertical* learning.

We denote by *horizontal* learning the type of learning commonly found in artificial systems (supervised/unsupervised learning, Reinforcement Learning). Horizontal learning takes place in predefined structures, which have been set up for that end. For instance, rules in a logic program, weights in an ANN, spline parameters of a motion primitive. During the learning process, new knowledge is created by tuning existing knowledge or appending the existing one with new instances. The knowledge acquired through horizontal learning can be subsequently used by the system without modifying its core algorithm, since the data structures and/or semantics are the same as for previous knowledge.

Vertical learning is a more fundamental type of learning which involves modification of the system itself as new knowledge is acquired. This is what humans (and probably other evolved species) do as they grow up. As they develop, infants gradually acquire representations about causality, space, time, quantity, and other concepts [2]. Vertical learning goes beyond acquiring new data: it requires to “update” the reasoning process itself. Consider for instance a system capable of causal reasoning using, e.g., task planning methods [3]. Causality is represented by means of operators with preconditions and effects. Such system can learn new causal relations by augmenting its domain with new operators (horizontal learning), but if the system is to learn something about duration of actions, then it needs both new representations (i.e., operators with duration) and to update the planning algorithm for reasoning upon time intervals. The same applies to perception and actuation: robots can only see or act what their representations and algorithms allow for.

We believe that both types of learning are necessary as a basis for intelligent robots: horizontal learning for adapting to new objects/environments, and vertical learning for being able to solve problems that have not been anticipated by their designers. Next, we examine some approaches to address the vertical learning problem.

4 Vertical Learning

To our knowledge, there exists no automated system capable of updating its core reasoning process through learning.

One way to circumvent the problem is to learn new representations. Learning new representations allows to see the world from a new perspective, therefore it is a key ability for solving unforeseen problems [4]. This approach has been used in Reinforcement Learning for learning new representations of the action space [5], in computer vision for image attributes [6], or in some cognitive architectures, e.g., SOAR, for learning macro-operators [7]. Learning new representations speeds up learning and improves generalization by better exploiting structure in the training data, but it does not modify the system’s core reasoning method. A system can for instance learn macro-operators, but the semantics of these macro-operators and the algorithm that reason upon them are predefined and remain unchanged through learning, which inherently bounds the scope of such systems.

Another approach to tackle vertical learning is to come up with a form of knowledge representation which can represent everything. If causal, perceptual and motor knowledge could be represented seamlessly with the same language, learning could take place in a single system, thereby avoiding the issue of learning in integrated systems addressed in Section 2. Ontologies are good candidates to this end. Some systems have been developed both for perception, e.g., SceneNet [8] and physical actions and processes [9]. The first issue with this approach is completeness. Manually modeling knowledge about, e.g., all existing physical objects, in the form of hierarchical subsumption relations is intractable [10], and has to be done manually (i.e., by human), which shifts the problem of vertical learning to ontology design. The second issue comes with reasoning upon this knowledge, which may be computationally intensive when it requires to merge knowledge across different domains [11].

Deep *end-to-end* learning allows to learn perceptual features, deliberation rules, and motor control parameters within a single process. But the training process is data and computationally intensive, even for narrow tasks such as object grasping [12] or driving [13]. Therefore it is not clear how this approach could scale up for robots learning to solve a wide range of problems.

Integrated systems have issues for relating knowledge learned across different subsystems, while monolithic systems have computational issues or heavily rely on designer’s knowledge. In the next section, we question (and wish to foster discussions on that theme during the workshop) the possibility of drawing inspiration from a constructivist psychology for addressing our problem from a different perspective.

5 Towards a Constructivist Approach

AI/Robotics essentially tries to reproduce cognitive and sensorimotor skills of humans *adults*. This approach has been successful for solving variety of problems,

even outperforming humans in narrow domains. In the constructivist paradigm, the question of interest is not “How do humans grasp different objects?” but rather how a system who is initially barely aware of itself –the infant– acquires knowledge and skills which allow him to grasp different objects.

The logician and psychologist Jean Piaget has long studied how knowledge is constructed, particularly in infants. His main contribution is the discovery of universal developmental stages in cognitive development, which may occur at different times, but always in the same sequence, regardless of cultural or social environment [14]. In other words intelligence is not innate, but constructed through necessary steps. The first stage is the *sensorimotor stage*, in which the infant progressively builds knowledge about the world through interactions within it (mainly trial and error at that stage). Piaget theorizes *schemas* as abstract elementary building blocks of knowledge. In a nutshell, schemas can represent objects, actions, or more abstract concepts. Knowledge builds up through acquisition of new and more abstract schemas. Piaget’s theory also provides two basic general mechanisms for developing schemas:

- *assimilation* is the process by which an existing schema is used on a novel object, e.g., a kid sees a bold man and shouts “clown!”;
- *accomodation* is the process of modifying existing schemas when assimilation failed, e.g., the father tells his kid that the bold man his not a clown because he does not have red hair. The kid then modifies his “clown” schema accordingly [15].

Piaget’s ideas have been implemented and tested in micro-world simulations or simple systems [16][17][18][19]. As argued by Guerin et al., Piaget’s theory is incomplete in different aspects [20] and requires more research to fill in the gaps, which makes it a potentially rich field of investigation. To our knowledge, no work as been done on applying Piaget’s ideas to robotic systems. In theory, the assimilation/accomodation learning mechanism proposed by Piaget allows for bottom-up hierarchical knowledge creation, from basic sensorimotor skills, know-hows, up to more abstract cognitive operations.

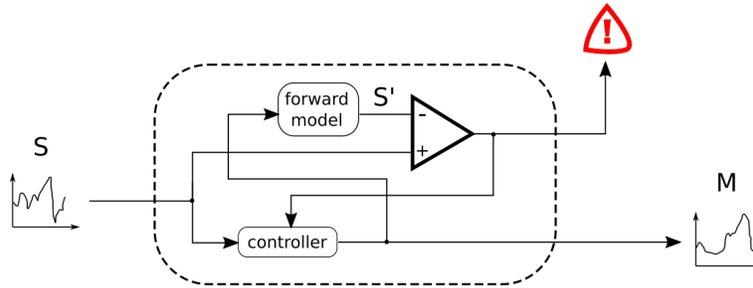


Fig. 2. Conceptual view of the proposed sensorimotor schema.

In order to investigate the application of Piaget’s theory to robotics, we propose a model which could be used as a basic building block for a schema-based learning robotic controller (see Fig. 2). Unlike previous attempts, this schema model operates in the continuous time domain, i.e., inputs and outputs are multidimensional time-dependent signals. The schema continuously learns a forward model, which predicts sensory signals (S') as a function of motor control signals (M). The difference between predictions and actual sensory input is used by the controller for adjusting control parameters in face of disturbances, which corresponds to the assimilation mechanism. When prediction (S') and actual sensory input (S) diverge beyond a certain threshold, a warning signal is issued to trigger the accommodation mechanism at a higher level.

In the initial stage, the system creates sensorimotor schemas through random exploration and motor babbling. These schemas are reinforced as they are re-enacted. When a sufficient number of sensorimotor schemas has been reached, they produce patterns of activation which can be assimilated by higher level schemas. Higher-level schemas follow the same principles as sensorimotor schemas do, except that their input and output come from other schemas instead of sensors and actuators. Hence, assimilation and accommodation take place with the same mechanism, but at higher level of abstraction.

More details about the envisioned system will be presented at the workshop and, given the preliminary status of our proposal, rather than presenting results, we hope to foster discussions and to get inspiring ideas and suggestions from the community.

References

1. Robin R. Murphy. *Introduction to AI Robotics*. MIT Press, Cambridge, MA, USA, 1st edition, 2000.
2. Jean Piaget. *The construction of reality in the child*. Basic Books, New York, 1954.
3. Dana Nau, Malik Ghallab, and Paolo Traverso. *Automated Planning: Theory & Practice*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2004.
4. Ana-Maria Olteanu, Mikkel Schttner, and Arpit Bahety. Towards a multi-level exploration of human and computational re-representation in unified cognitive frameworks. *Frontiers in Psychology*, 10:940, 2019.
5. Yash Chandak, Georgios Theocharous, James Kostas, Scott Jordan, and Philip S. Thomas. Learning action representations for reinforcement learning. *CoRR*, abs/1902.00183, 2019.
6. Zeynep Akata, Florent Perronnin, Zaïd Harchaoui, and Cordelia Schmid. Label-embedding for image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38:1425–1438, 2016.
7. John E. Laird, Paul S. Rosenbloom, and Allen Newell. Chunking in soar: The anatomy of a general learning mechanism. *Mach. Learn.*, 1(1):11–46, March 1986.
8. Ilan Kadar and Ohad Ben-Shahar. Scenenet: A perceptual ontology for scene understanding. In Lourdes Agapito, Michael M. Bronstein, and Carsten Rother, editors, *Computer Vision - ECCV 2014 Workshops*, pages 385–400, Cham, 2015. Springer International Publishing.

9. Moritz Tenorth and Michael Beetz. A unified representation for reasoning about robot actions, processes, and their effects on objects. *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1351–1358, 2012.
10. Viviana Mascardi, Valentina Cordi, and Paolo Rosso. A comparison of upper ontologies. Technical Report DISI-TR-06-21, Dipartimento di Informatica e Scienze dell'Informazione (DISI), Universit'a degli Studi di Genova, Via Dodecaneso 35, 16146, Genova, Italy, 2006.
11. Kathrin Dentler, Ronald Cornet, Annette ten Teije, and Nicolette de Keizer. Comparison of reasoners for large ontologies in the owl 2 el profile. *Semant. web*, 2(2):71–87, April 2011.
12. Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *CoRR*, abs/1504.00702, 2015.
13. Huazhe Xu, Yang Gao, Fisher Yu, and Trevor Darrell. End-to-end learning of driving models from large-scale video datasets. *CoRR*, abs/1612.01079, 2016.
14. Piaget Jean. *The origins of intelligence in children*. 1952.
15. R. S. Siegler, J. S. DeLoache, N. Eisenberg, J. Saffran, and C Leaper. *How children develop*. Worth Publishers, New York, 4th edition, 2004a.
16. Ezequiel Alejandro Di Paolo, Xabier E. Barandiaran, Michael Beaton, and Thomas Buhrmann. Learning to perceive in the sensorimotor approach: Piagets theory of equilibration interpreted dynamically. *Frontiers in Human Neuroscience*, 8:551, 2014.
17. Gary L. Drescher. *Made-Up Minds: A Constructivist Approach to Artificial Intelligence*. Cambridge: MIT Press, 1991.
18. Harold H. Chaput. *The Constructivist Learning Architecture: A Model of Cognitive Development for Robust Autonomous Robots*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin, August 2004. Also Technical Report TR-04-34.
19. Olivier L. Georgeon and Frank E. Ritter. An intrinsically-motivated schema mechanism to model and simulate emergent cognition. *Cogn. Syst. Res.*, 15-16:73–92, May 2012.
20. Frank Guerin and D. McKenzie. A piagetian model of early sensorimotor development. In *Proceedings of the Eighth International Conference on Epigenetic Robotics*, number - in Lund University Cognitive Studies, pages 29–36. Kognitionsforskning, Lunds universitet, 2008.