Towards the Emergence of Procedural Memories from Lifelong Multi-Modal Streaming Memories for Cognitive Robots

Maxime Petit* Tobias Fischer* Yiannis Demiris

***Abstract*— Various research topics are emerging as the demand for intelligent lifelong interactions between robot and humans increases. Among them, we can find the examination of persistent storage, the continuous unsupervised annotation of memories and the usage of data at high-frequency over long periods of time. We recently proposed a lifelong autobiographical memory architecture tackling some of these challenges, allowing the iCub humanoid robot to 1) create new memories for both actions that are self-executed and observed from humans, 2) continuously annotate these actions in an unsupervised manner, and 3) use reasoning modules to augment these memories *a-posteriori*. In this paper, we present a reasoning algorithm which generalises the robots' understanding of actions by finding the point of commonalities with the former ones. In particular, we generated and labelled templates of pointing actions in different directions. This represents a first step towards the emergence of a procedural memory within a long-term autobiographical memory framework for robots.**

I. Introduction

Various studies have shown that cognitive robots can significantly benefit from a multi-modal long-term memory; for example in the context of navigation planning [1], adaptive assistance [2], and cooperative task learning [3]. According to [4], robots can benefit from taking their procedural and episodic memories into consideration when choosing appropriate actions according to previous experiences, in the same way as humans do [5]. Furthermore, we subscribe to the view that memory should be considered as an active cognitive component [6]: agents constantly use their own memories to find patterns among them, leading to the generalisation of *e.g.* actions concepts.

Previously, we have proposed an autobiographical multi-modal memory framework for cognitive robots to store, augment, and recall streaming episodes [7]. We have shown that this enables autonomous robots to 1) learn from self exploration and social interaction using motor babbling and imitation, and 2) to aid computer vision algorithms which are not currently able to run in real-time to contribute to the analysis and augmentation of streaming data (*e.g.* [8]). Other examples for the usage of the framework are the hierarchical learning of actions, which requires new knowledge to be built upon previously acquired information [9]; as well as perspective taking in order to reconstruct occluded views from memory [10].

The authors are with the Personal Robotics Lab, Department of Electrical and Electronic Engineering, Imperial College London, UK. Email: *{m.petit, t.fischer, y.demiris}@imperial.ac.uk*

*M. Petit and T. Fischer contributed equally to this work.

This work was supported in part by the EU FP7 project WYSIWYD (Grant 612139).

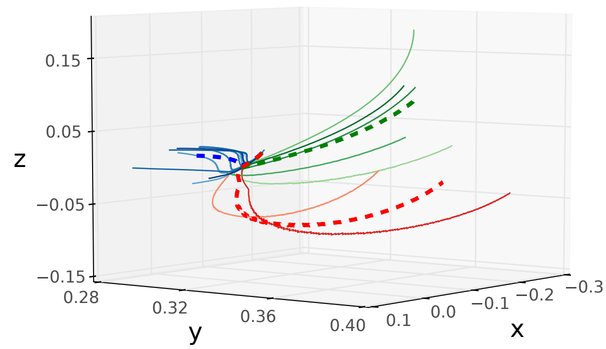

Fig. 1. Trajectories of 19 distinct pointing actions of the iCub's left hand end-effector retrieved from the autobiographical memory. One can see three clusters, whereas the blue-coloured cluster originates from pointing actions to the left, the green-coloured from pointing actions to the right upwards and the red-coloured right downwards. The computed mean trajectories are plotted dashed and bold. Figure is best viewed in colour.

In this paper, we present steps towards the emergence of a procedural memory by clustering annotated self-actions. Thus, our method first learns templates of actions, then acquires corresponding labels in a human-robot interaction, and finally stores the labelled actions in the long-term memory.

II. Related Works

Our proposed framework is based on the theoretical paradigm of lifelong machine learning as proposed by Silver *et al.* [11]. The authors argued in favour of lifelong memories to adapt to new situations as these allow employing both universal and domain specific knowledge. They proposed to consolidate memories to abstract knowledge from specific episodes, which results in concepts that can be used as prior knowledge when learning new tasks.

Dubba *et al.* [12] proposed a cognitive architecture that allows learning of higher level concepts after an initial learning phase where the robot learns to recognise objects. Their framework allows multiple input sources and the interaction can take place with non-expert users in relatively open environments. However, they are only providing a graphical web-based interface to allow annotations from a human, despite highlighting that natural language is the most preferred interface to robots. Our architecture allows such natural human-robot interactions.

III. Architecture and Analysis

In this section, we briefly review the architecture we use to store and recall autobiographical memories [7]. We then provide initial evidence that the data contained in the memories can be clustered, which subsequently allows on-the-fly classification of new template actions.

A. Architecture Overview

The central component of the employed framework [7] is a SQL database, which was designed to store data originating from various sources in a general manner. The data can cover multiple modalities (proprioception, vision, language, *etc.*), as well as multiple levels of abstraction (*e.g.* from raw sound signals over sentences to extracted meanings).

Several interfaces are in place to store, augment and recall the memories. Action generation modules trigger episodic events and provide basic annotations. During an episode, continuous data from specified sensors are acquired through the input interface. For example, when triggered, the module responsible for pointing actions indicates the beginning and end of an action and provides the name (*i.e.* "pointing") of the performed action along with the used parameters. While the action is performed, proprioceptive and visual information is stored. The annotations can be used by reasoning modules (*e.g.* machine learning algorithms) to retrieve related episodes, and add augmented data to these original memories.

B. Understanding of Actions

In this section, we provide preliminary results towards unsupervised understanding of self-actions. The work presented in this paper allows the robot to form procedural memories by autonomously discovering higher-level representations of actions from previous episodes. The labels of the concepts stored in the procedural memory (*i.e.* of the action templates) can then be acquired in human-robot interactions.

In [7], we have shown that a reasoning module can be used to extract the kinematic structure of the human hand given an image sequence of a specific episode provided by the autobiographical memory. In this paper, there are two key differences. Firstly, we are interested in using data which was recorded without any specific intention, and thus lacking annotation by the human. Secondly, rather than augmenting a single, specific episode, the framework finds patterns in groups of related episodes. As an example, we use pointing actions executed with the left hand within this paper.

We analyse these pointing actions, and plot the trajectories in Fig. 1. The framework uses the Mean Shift algorithm [13] to find clusters in the data, allowing to detect three clusters among the 19 pointing gestures of the iCub with the left hand. The different types of pointing actions that emerge are: pointing to the left and right (blue and red trajectories respectively), and pointing right+upwards (green trajectories). However, the autobiographical memory does not yet contain this semantic information.

Thus, we extend the human-robot interaction abilities of [7][1] such that the iCub is able to ask a human to provide labels of the newly found clusters. After executing motor commands which correspond to the mean of the motor commands found in the clusters, the iCub then asks for a label of this action, which is also used as additional annotation for the original memories.

[1]So far, the iCub can a) greet a human, b) remember a unique event with and without augmented memories, c) remember a subset of events including active recalling with its body, and d) acquire feedback about the quality of reasoning results (for more details, see [7]).

Additionally to the in-build action primitives which allow pointing towards specified coordinates, the iCub has then discovered commonalities within these pointing actions, which can be seen as higher-level primitives. Now, the iCub can "point upwards with your left hand", "point to the right with your left hand" and "point to the left with your left hand" in human-robot scenarios (note that the pointing is irrelevant of a physical target). We can use the new pointing actions as building components for hierarchical action learning as described in [9].

IV. Discussion and Conclusion

We have proposed using the Mean Shift clustering algorithm to gain further insights to already annotated actions. In our example, an annotation module provided only the label "pointing" along with the arm used. After the iCub has found three clusters automatically, it asked a human to name the template actions, resulting in three distinct labels discriminating the three pointing actions. Thus, we have shown how a long term autobiographical memory allows *a-posteriori* reasoning, which can be used to learn new action concepts without conducting specific experiments. This can be seen as the first step towards the emergence of a procedural memory. Investigating other types of actions, as well as non-action related concepts will show how well our method scales beyond the presented pointing actions. For our future works, we aim to use the action clusters found on the iCub, and transfer this knowledge to other robots such as the Baxter.

References

[1] J. P. Fentanes *et al.*, "Now or later? Predicting and Maximising Success of Navigation Actions from Long-Term Experience," in *IEEE Int. Conf. Robot. Autom.*, 2015, pp. 1112–1117.

[2] Y. Demiris, "Knowing when to assist: Developmental issues in lifelong assistive robotics," in *IEEE Int. Conf. Eng. Medicine and Biology Soc.*, 2009, pp. 3357–3360.

[3] M. Petit *et al.*, "The Coordinating Role of Language in Real-Time Multimodal Learning of Cooperative Tasks," *IEEE Trans. Auton. Ment. Dev.*, vol. 5, no. 1, pp. 3–17, 2013.

[4] D. Vernon, M. Beetz, and G. Sandini, "Prospection in Cognition: The Case for Joint Episodic-Procedural Memory in Cognitive Robotics," *Front. Robot. AI*, vol. 2, pp. 1–14, 2015.

[5] C. M. Atance and D. K. O'Neill, "Episodic future thinking," *Trends Cogn. Sci.*, vol. 5, no. 12, pp. 533–539, 2001.

[6] R. Wood, P. Baxter, and T. Belpaeme, "A review of long-term memory in natural and synthetic systems," *Adapt. Behav.*, vol. 20, no. 2, pp. 81–103, 2012.

[7] M. Petit, T. Fischer, and Y. Demiris, "Lifelong Augmentation of Multi-Modal Streaming Autobiographical Memories," *IEEE Trans. Cogn. Develop. Syst.*, vol. 8, no. 3, pp. 201–213, 2016.

[8] H. J. Chang, T. Fischer, M. Petit, M. Zambelli, and Y. Demiris, "Kinematic Structure Correspondences via Hypergraph Matching," in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 4216–4425.

[9] M. Petit and Y. Demiris, "Hierarchical Action Learning by Instruction Through Interactive Body Part and Proto-Action Grounding," in *IEEE Int. Conf. Robot. Autom.*, 2016, pp. 3375–3382.

[10] T. Fischer and Y. Demiris, "Markerless Perspective Taking for Humanoid Robots in Unconstrained Environments," in *IEEE Int. Conf. Robot. Autom.*, 2016, pp. 3309–3316.

[11] D. L. Silver, Q. Yang, and L. Li, "Lifelong Machine Learning Systems: Beyond Learning Algorithms," in *AAAI Symp.*, 2013, pp. 49–55.

[12] K. S. Dubba *et al.*, "Grounding Language in Perception for Scene Conceptualization in Autonomous Robots," in *AAAI Symp.*, 2014, pp. 26–33.

[13] D. Comaniciu and P. Meer, "Mean Shift: A Robust Approach Toward Feature Space Analysis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 5, pp. 603–619, 2002.