# Incrementally Learning a Library of Full-Pose Via-Point Movement Primitives for a Humanoid Robot

Tilman Daab[1], Noémie Jaquier[1,2], Christian Dreher[1], Andre Meixner[1], Franziska Krebs[1], Annemarie Schaub[1], David Vigneron[1], and Tamim Asfour[1]

***Abstract*— Robots should be able to continuously learn and enhance their skills and abilities over time. Such skills can be represented as sequences of movement primitives, which are known for their generalization abilities and can be re-used across tasks when being collected in a library. Incrementally learning such a library efficiently allows movement primitives to be updated without the need to permanently store all demonstrations. This extended abstract builds on our previous work on incrementally learning Full-Pose Via-Point Movement Primitives based on 7 fundamental operations. Here, we integrate this movement primitive library into a cognitive architecture and conceptually show how to incrementally learn such a library for a humanoid robot from demonstrations of multiple modalities.**

## I. Introduction

Robots at home, in care settings or at work need to adapt their skills to new situations. For this, movement primitives (MPs) [1], [2] are a commonly used representation, particularly in the field of learning from demonstrations [3] and imitation learning [2]. Storing MPs in a library [2], [4] enables their re-use across tasks.

Robots may need to extend and improve their skills over time [5], [4], mandating to update the MP library. Using all previous demonstrations for such updates results in unnecessary storage and computation requirements. This is avoided by incrementally learning from a new demonstration alone. Akin to Gepperth and Hammer [6], we consider *online learning* as learning from sequentially-arriving data, and *incremental learning* as online learning with limited amount of memory. Specifically, we require the memory for learning $m$ models from $n$ samples with $m < n$ to be bound by $\mathcal{O}(f(m))$, i. e., independently of $n$.

Incrementally learning a MP library is not just about improving an existing MP, but requires further operations, such as merging or splitting MPs. Many previous works enable a single operation [7], [8], [9]. Others enable more operations, but are only applied to 2D motions [10], achieve bounded memory consumption by discarding old demonstrations [5], or are purely on-line instead of incremental, due to storing all demonstrations permanently [11], [12]. In contrast, we aim at providing multiple operations on a MP library that is capable to describe full-pose motions. The operations shall be performed in an incremental fashion without the need to discard some of the demonstrations.

The research leading to these results has received funding from the European Union's Horizon 2021 Research and Innovation programme under grant agreement No 101070292 (HARIA) and from the Carl Zeiss Foundation through the JuBot project, and was partially supported by the German Federal Ministry of Education and Research (BMBF) under the Robotics Institute Germany (RIG).

[1]Institute for Anthropomatics and Robotics, High Performance Humanoid Technologies Lab (H2T), Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany. {tilman.daab, asfour}@kit.edu

[2]Division of Robotics, Perception and Learning (RPL), KTH Royal Institute of Technology, Stockholm, Sweden.

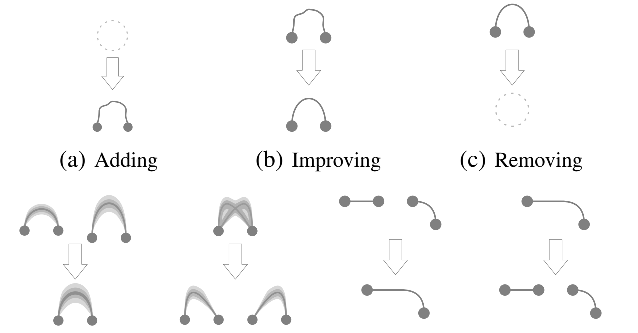

(a) Adding (b) Improving (c) Removing (d) Merging modes (e) Splitting a mode (f) Merging temporally (g) Splitting temporally

Fig. 1: Fundamental operations to incrementally learn MP libraries.

In the following, we outline our previous work [13] on incrementally learning MP libraries (Section II and Section III), show its embedding into a cognitive architecture (Section IV), and an outlook on future directions (Section V).

## II. Fundamental Operations to Incrementally Learn a Movement Primitive Library

As presented previously [13], we consider *adding*, *improving*, *removing*, *merging modes*, *splitting a mode*, *merging temporally*, and *splitting temporally* as the fundamental operations of learning a MP library incrementally, which are depicted in Fig. 1. The first three operations act on a single MP, while the others involve merging or splitting of two MPs. Thereby, the latter operations re-organize the MP library. Moreover, the first five operations are spatial ones, affecting MP amplitudes, while the others are temporal operations, addressing the concatenation of MPs.

## III. Full-Pose Via-Point Movement Primitives

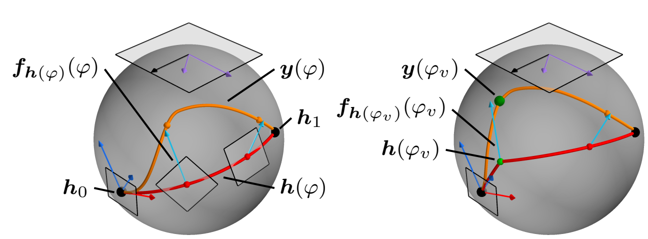


(a) Decomposition of the trajectory (b) With via-point at $\boldsymbol{y}(\varphi_v)$

Fig. 2: Illustration of a Riemannian VMP on $\mathcal{S}^2$.

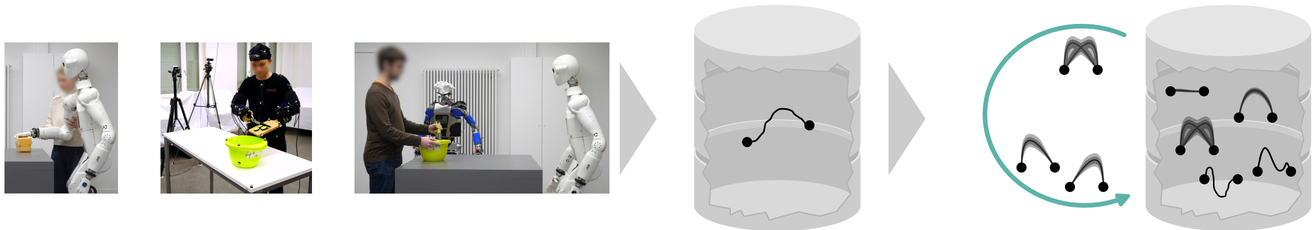

(a) Illustrations of demonstration modalities (left to right): Kinesthetic teaching of a *wiping* motion, motion capture for a *transfer* motion [14], visual observation of a *stirring* motion.

(b) Consolidation as demonstration.

(c) Movement primitive library with mode splitting shown as an example.

Fig. 3: Illustration of processing demonstrations as part of the memory-centric cognitive architecture.

Via-Point Movement Primitives [15] (VMPs) are highly flexible and combine the extrapolation capabilities of Dynamic Movement Primitives [16] with the ability of Probabilistic Movement Primitives [17] to handle via-points. We exemplarily use an extension of VMPs to full poses [13] to represent the movement primitives of the library. This enables to accurately handle orientations by leveraging Riemannian geometry.

VMPs decompose the trajectory $\boldsymbol{y}$ into an elementary trajectory $\boldsymbol{h}(\varphi)$ and a shape modulation $\boldsymbol{f}(\varphi)$ (Fig. 2),

$$\boldsymbol{y}(\varphi) = \mathrm{Exp}_{\boldsymbol{h}(\varphi)}\left(\boldsymbol{f}_{\boldsymbol{h}(\varphi)}(\varphi)\right) \tag{1a}$$

$$\boldsymbol{y}(\varphi) = \boldsymbol{h}(\varphi) + \boldsymbol{f}(\varphi), \tag{1b}$$

where Eqn. (a) describes the Riemannian [18] and (b) the Euclidean case. The phase $\varphi$ describes the progress, in our case using $\varphi \in [0; 1]$. The elementary trajectory $\boldsymbol{h}(\varphi)$ interpolates between start pose $\boldsymbol{h}_0$ and end pose $\boldsymbol{h}_1$,

$$\boldsymbol{h}(\varphi) = \mathrm{Exp}_{\boldsymbol{h}_0}\left(\varphi \cdot \mathrm{Log}_{\boldsymbol{h}_0}(\boldsymbol{h}_1)\right) \tag{2a}$$

$$\boldsymbol{h}(\varphi) = \boldsymbol{h}_0 + \varphi \cdot (\boldsymbol{h}_1 - \boldsymbol{h}_0). \tag{2b}$$

The shape modulation describes the deviation of the actual trajectory from the point of the elementary trajectory as a sum of kernels $\Psi$ activated by phase, multiplied by coefficients stored as the weight vector $\boldsymbol{w}$,

$$\boldsymbol{f}_{\boldsymbol{h}(\varphi)}(\varphi) = \Gamma_{\boldsymbol{h}_0 \to \boldsymbol{h}(\varphi)}\left(\Psi(\varphi)\boldsymbol{w}_{\boldsymbol{h}_0}\right) \tag{3a}$$

$$\boldsymbol{f}(\varphi) = \Psi(\varphi)\boldsymbol{w}. \tag{3b}$$

For detailed descriptions, we refer to Daab et al. [13]. With the spatial operations having been realized therein, we are currently working on the temporal operations.

## IV. Integration into a Cognitive Architecture

To apply the incremental learning methods to a robot, an integration into its cognitive architecture is required. We leverage a memory-centric approach [19], allowing multi-modal data to be organized in an episodic manner. Depending on the demonstration method (Fig. 3a), relevant modalities could be raw joint angles or processed human poses. Objects in the scene or force measurements could be added as well. The memory allows to access these different modalities of a demonstration via a unified interface (Fig. 3b). The MP library is represented in the robot's memory as well (Fig. 3c), and can be learned over prolonged periods by using the long term memory. It not only allows to retrieve the stored MPs, but to apply the fundamental operations to it. To conserve the acquired knowledge, the long term memory is used. In the future, sequences of MPs and their task parameters could be represented and learned in a similar manner.

Learning from demonstrations of different modalities is not only a matter of a unified interface, but also of the demonstrations adhering to different kinematic structures – here, following the robot's or the human's embodiment. For a general solution, motion retargeting is needed to learn across these embodiments. When just learning the end effector motion in task space, one could also attempt to neglect the remaining differences and directly learn across such demonstrations.

The proposed integration of incrementally learning a MP library allows the following: Initially, the MP library can either be empty, or equipped with prior knowledge. The data of a demonstration are spread over the robot's distributed memory system. They are consolidated and provided to the memory of the MP library for incremental learning. For example, if an MP already existed that was to be split into two MPs representing different modes, that MP would be retrieved from the MP library. The demonstration would be used to identify one of the modes, which is separated into a first MP, and the remaining estimation forms a second MP. Finally, both MPs would be stored in the MP library, and the demonstration could be discarded. Not only can MPs learned from different modalities be stored in the same MP library, but even the same MP could be incrementally learned from different modalities.

## V. Outlook

On the long term, we aim for the robot to learn with less manual interaction, so that the decision of when to segment demonstrations and which incremental learning operation to apply are perform automatically. This also facilitates to learn from larger data sets, or even by observing the environment. When including the temporal operations, the robustness of such learning is expected to increase, as erroneous segmentation decisions can be corrected when learning from further observations.

## REFERENCES


[1] S. F. Giszter, F. A. Mussa-Ivaldi, and E. Bizzi, "Convergent Force Fields Organized in the Frog's Spinal Cord," *Journal of Neuroscience*, vol. 13, pp. 467–491, 1993.

[2] S. Schaal, "Is Imitation Learning the Route to Humanoid Robots?," *Trends in Cognitive Sciences*, vol. 3, no. 6, pp. 233–242, 1999.

[3] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, "Robot Programming by Demonstration," in *Handbook of Robotics*, pp. 1371–1394, Springer, 2008.

[4] P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal, "Learning and Generalization of Motor Skills by Learning From Demonstration," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 763–768, 2009.

[5] D. Kulić, W. Takano, and Y. Nakamura, "Incremental Learning, Clustering and Hierarchy Formation of Whole Body Motion Patterns using Adaptive Hidden Markov Chains," *Intl. Journal of Robotics Research*, vol. 27, pp. 761–784, July 2008.

[6] A. Gepperth and B. Hammer, "Incremental Learning Algorithms and Applications," in *European Symposium on Artificial Neural Networks (ESANN)*, 2016.

[7] S. Calinon and A. Billard, "Incremental Learning of Gestures by Imitation in a Humanoid Robot," in *ACM/IEEE Intl. Conf. on Human-Robot Interaction (HRI)*, p. 255, 2007.

[8] F. Meier, E. Theodorou, and S. Schaal, "Movement Segmentation and Recognition for Imitation Learning," in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 761–769, 2012.

[9] A. Gams, T. Petrič, M. Do, B. Nemec, J. Morimoto, T. Asfour, and A. Ude, "Adaptation and Coaching of Periodic Motion Primitives through Physical and Visual Interaction," *Robotics and Autonomous Systems*, vol. 75, pp. 340–351, 2016.

[10] A. Lemme, R. F. Reinhart, and J. J. Steil, "Self-Supervised Bootstrapping of a Movement Primitive Library From Complex Trajectories," in *IEEE/RAS Intl. Conf. on Humanoid Robots (Humanoids)*, pp. 726–732, 2014.

[11] D. Kulić, C. Ott, D. Lee, J. Ishikawa, and Y. Nakamura, "Incremental Learning of Full Body Motion Primitives and Their Sequencing Through Human Motion Observation," *Intl. Journal of Robotics Research*, vol. 31, no. 3, pp. 330–345, 2012.

[12] S. Niekum, S. Osentoski, G. Konidaris, S. Chitta, B. Marthi, and A. G. Barto, "Learning Grounded Finite-State Representations From Unstructured Demonstrations," *Intl. Journal of Robotics Research*, vol. 34, no. 2, pp. 131–157, 2015.

[13] T. Daab, N. Jaquier, C. Dreher, A. Meixner, F. Krebs, and T. Asfour, "Incremental learning of full-pose via-point movement primitives on riemannian manifolds," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pp. 2317–2323, 2024.

[14] F. Krebs, A. Meixner, I. Patzer, and T. Asfour, "The KIT Bimanual Manipulation Dataset," in *IEEE/RAS Intl. Conf. on Humanoid Robots (Humanoids)*, pp. 499–506, 2021.

[15] Y. Zhou, J. Gao, and T. Asfour, "Learning Via-Point Movement Primitives With Inter- and Extrapolation Capabilities," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 4301–4308, 2019.

[16] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, "Dynamical Movement Primitives: Learning Attractor Models for Motor Behaviors," *Neural Computation*, vol. 25, no. 2, pp. 328–373, 2013.

[17] A. Paraschos, C. Daniel, J. R. Peters, and G. Neumann, "Probabilistic Movement Primitives," in *Neural Information Processing Systems (NeurIPS)*, vol. 26, pp. 2616–2624, 2013.

[18] J. M. Lee, *Introduction to Riemannian Manifolds*. Springer, 2018.

[19] F. Peller-Konrad, R. Kartmann, C. R. Dreher, A. Meixner, F. Reister, M. Grotz, and T. Asfour, "A memory system of a robot cognitive architecture and its implementation in armarx," *Robotics and Autonomous Systems*, vol. 164, p. 104415, 2023.