Robot learning in factories of the future: adaptability, redundancy and exploitation of movement options

João Silvério and Sylvain Calinon

***Abstract*— In the factories of the future, robots are expected to operate alongside humans in seamless ways, raising a multitude of challenges for robot learning. Particularly, learning algorithms should be able to exploit variability information contained in task demonstrations to adapt robot motions in ways that are both safe for the human partners and efficient for the tasks at hand, with certain behaviors prioritized over others. One way to achieve this is by leveraging the redundancy of robots both at the level of their kinematics and trajectory planning. In this abstract, we discuss the challenges in learning priorities by imitation, the connections to trajectory redundancy and identify the main problems to be solved for bringing learning from demonstration to industrial case studies.**

I. Introduction

The challenges of modern and future factories make imitation learning a promising avenue to cope with the need for rapid robot re-programming. However, learning algorithms should be re-formulated so that controllers can adapt to the conditions that robots will find in such environments. The possible presence of human partners in the surroundings of the robot raises the bar when it comes to safety, with the robot having to perform its task, as long as it does not harm its partner, prioritizing the safety of the human. Other levels of priority exist, where the robot may try to explore alternative motions for its degrees of freedom, as long as it does not affect the task at hand. To add to the complexity, the way in which these different priorities are organized should flexibly be modified depending on inputs from the environment (e.g. if sensors give a high certainty to the robot that there are no humans close to its workspace).

In previous work [1], we leveraged the linear structure of task-parameterized Gaussian Mixture Models (TP-GMM) to learn task priorities from demonstrations, where candidate task hierarchies with null space matrices were used to automatically extract the relevant hierarchies based on their variability. In [2], a null space formulation for the classical model predictive control problem is proposed where the spatiotemporal redundancy in the trajectory plan is exploited to pass through different via-points before reaching the final goal state, without compromising the task success. These research lines show that redundancy is common to different robotics problems, opening various possibilities while learning new skills by imitation.

In the following section we propose research challenges relevant to the factories of the future that can benefit from the aforementioned insights and discuss how we plan to tackle them.

The authors are with the Idiap Research Institute, CH-1920 Martigny, Switzerland (e-mail: name.surname@idiap.ch). This work was supported by Innosuisse (Cob'Hook project) and by the European Commission's Horizon 2020 Programme (CoLLaboratE project, https://collaborate-project.eu/, grant agreement 820767).

II. Research challenges for robot learning in the factories of the future

Here we identify research questions connected to modern industrial scenarios and how we propose to address them using insights from previous work and the state-of-the-art.

A. Combining safety, task fulfillment and refinement into a learning framework

Safety for the human partners should always be the highest priority for robot controllers, with manipulation tasks being performed only when they do not compromise safety. At the same time, the refinement of movement primitives required to succeed at a task should be done without compromising the task itself. Interestingly, these common assumptions may be allowed to change when the task conditions change. For instance, if the workspace of the robot becomes free, the robot can perform its task with a joint posture that it would not use in the presence of users, or follow a different path to fulfill the task. It is therefore crucial to answer the question of *how to teach priority behaviors in a way that they can change in response to variations of different inputs (potentially high-dimensional)?*

Our previous approach [1] is well-suited for this but so far our results have been applied at the level of the robot kinematic redundancy. A promising way to extend the same concept to the redundancy of the task itself is by merging the approach with [2]. In this case, using demonstrations of the robot passing through different via-points while fulfiling the main task should, in principle, generate enough variability for TP-GMM to extract the required priorities and how they change with respect to the inputs.

Another situation that arises frequently is that, for certain inputs, the robot may not have enough information about how to perform its task. We propose that, in these cases, the robot is allowed to perform exploratory motions outside of the null space of the task. In [3] we introduced an approach based on kernelized movement primitives (KMP) [4] to render the robot safer (through high compliance) when facing input data that it did not see during demonstrations of collaborative tasks. We are convinced that the uncertainty information provided by KMP is very well suited to guide exploratory behaviors, where the robot would be allowed to explore outside its null space when the uncertainty is high and vice versa when it is low.

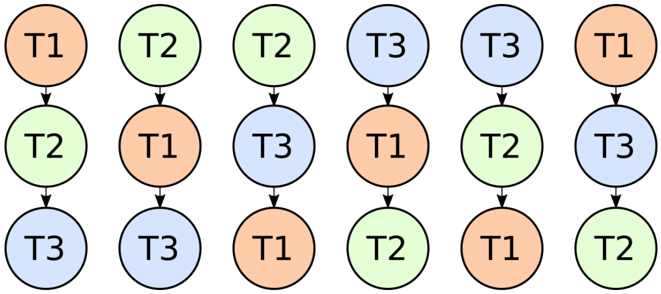

Fig. 1: 3-task illustration of candidate hierarchy definition, as used in [1]. Arrows point in the direction of priority, i.e. the task below the arrow is performed in the null space of the task above. Even though the task context may provide prior information that allows for removing hierarchies (e.g. if one know that T3 cannot be the last task) the complexity still scales significantly with the number of tasks.

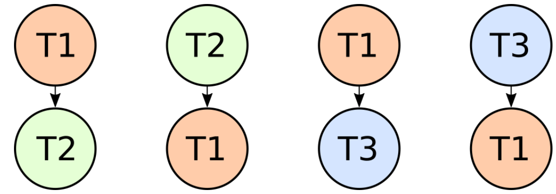

Fig. 2: Simplified version of candidate hierarchy definition, where the hierarchies are not exhaustively defined but only couples of hierarchies are given. The goal is then to find the optimal combination of these.

It should be noted that the refinement of movement primitives can, and should, also occur in joint space. Control aspects such as manipulability and energy efficiency are essential for adequate robot manipulation and they are often optimized at the joint level. Exploiting null space structures allows robots to first learn the task to be fulfilled while subsequently exploiting the null space to find the optimal configuration to perform it. We propose to enhance this aspect by allowing the robot to autonomously explore its joint space to a degree that is proportional to how much the main task can be disrupted (i.e. explore more when the uncertainty of the main task is high).

B. Reducing amount of prior priority knowledge

One of the limitations of [1] is that our approach requires an extensive definition of candidate hierarchies, with the limit case being all the possible combinations of tasks (although this can be strongly reduced via domain knowledge about the application), see Fig. 1. We propose to study alternative approaches to reduce the required prior knowledge. A promising direction is to consider hierarchies of fewer tasks (e.g. only pairs of tasks), see Fig. 2, and either estimate their optimal combination from the data or let the robot discover them through exploration. The optimal solution is, in this case, the result of a *product of experts* [5], which for normally-distributed prioritization primitives, corresponds to a product of Gaussians. As explained in [1], representing demonstration data in the subspaces associated with each candidate hierarchy provides enough information (through variability) for the robot to autonomously extract the used hierarchies in a seamless way, including for varying inputs. This property is compatible with the simplified hierarchy definition depicted in Fig. 2.

C. Learning distributions of null space structures from demonstrations

Approaches like [1], [2] rely on null space structures that are parameterized by the kinematics of the robot and its configuration. However, the environment typically constrains the robot motion (e.g. planar motions on a table), allowing for movements in some directions without compromising the actual task. This allows for the encoding of the task constraints as null space projection matrices. Lin *et al.* [6] propose an approach for learning such matrices from data. However, it cannot learn distributions of null space matrices, which hinders the approach from being used in cases where the null space needs to be modulated by external inputs. Null space projection matrices can be represented as points in Grassmann manifolds. Thus, we propose to exploit the Riemannian geometry of this space in combination with statistical tools [7], [8], to obtain probabilistic representations of task constraints that can be externally modulated and generalized. Particularly, operations like Gaussian mixture regression or information fusion through the product of Riemannian Gaussians are formulated in [7], [8], allowing the extension of classical imitation learning concepts to data with unique geometric properties.

III. Conclusion

In this abstract we highlighted some of the current challenges in learning new skills by imitation in modern industrial contexts. In particular, we made the case for adequately exploiting the variability and priority found in the data to learn movement options that are leveraged to optimize the robot behavior in the presence of humans. We also reviewed work showing that redundancy is present in various robot control problems and proposed ways to exploit it in combination with variability/priority information.

References

[1] J. Silvério, S. Calinon, L. Rozo, and D. G. Caldwell, "Learning task priorities from demonstrations," *IEEE Transactions on Robotics*, vol. 35, no. 1, pp. 78–94, 2019.

[2] H. Girgin and S. Calinon, "Nullspace structure in model predictive control," *CoRR*, vol. abs/1905.09679, 2019. [Online]. Available: http://arxiv.org/abs/1905.09679

[3] J. Silvério, Y. Huang, L. Rozo, F. Abu-Dakka, and D. G. Caldwell, "Uncertainty-aware imitation learning using kernelized movement primitives," in *Proc. IEEE/RSJ Intl Conf. on Intelligent Robots and Systems (IROS)*, Macau, China, November 2019, p. (to appear).

[4] Y. Huang, L. Rozo, J. Silvério, and D. G. Caldwell, "Kernelized movement primitives," *International Journal of Robotics Research*, vol. 38, no. 7, pp. 833–852, 2019.

[5] G. Hinton, "Products of experts," in *Proc. Intl Conf. on Artificial Neural Networks (ICANN)*, 1999, pp. 1–6.

[6] H.-C. Lin, M. Howard, and S. Vijayakumar, "Learning null space projections," in *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA)*, May 2015, pp. 2613–2619.

[7] M. J. A. Zeestraten, I. Havoutis, J. Silvério, S. Calinon, and D. G. Caldwell, "An approach for imitation learning on Riemannian manifolds," *IEEE Robotics and Automation Letters (RA-L)*, vol. 2, no. 3, pp. 1240–1247, June 2017.

[8] S. Calinon and N. Jaquier, "Gaussians on Riemannian manifolds for robot learning and adaptive control," *arXiv*, pp. 1–10, 2019.