Learning How to Plan for Multi-Step Manipulation in Collaborative Robotics

Claudia Pérez-D'Arpino

I. Introduction

The use of robots for complex manipulation tasks in collaboration with humans is currently challenged by the limited ability of robots to construct a rich representation of the activity at both the motion and task levels in ways that are both functional and apt for human-supervised execution. Traditionally, the need for algorithms to build this representation is avoided by either resorting to teleoperation or by having expert programmers code the actions for a robot. For instance, the operator of a remote robot (Figure 1 a), which currently uses joint-by-joint teleoperation, would benefit from planning assistance from the robot. This assistance requires the robot to autonomously generate manipulation plans by leveraging the current context and its existing skills. In manufacturing, robots have been successfully deployed in structured environments on tasks where all actions can be coded by domain experts. This paradigm needs to change in order to enable robots to execute manipulation tasks in shared workspace with humans (Figure 1 b), where the structure of the tasks is subject to uncertainty about the human actions and the planner can no longer execute a predesigned strategy, but instead needs to handle predictions of human actions while still accomplishing the task. In both cases, it is beneficial to advance the state of the art with systems that are capable of:

1) Learning constrained multi-step manipulation tasks from observed demonstrations, as this would enable the application of robotics by a larger set of users in increasingly complex scenarios;
2) Implementing efficient workflows for human-in-the-loop execution of the learned multi-step manipulation tasks;
3) Using this knowledge to plan for a broader set of tasks under uncertainty, so that the current context and online predictions can be incorporated.

The technical approach is to develop models and algorithms to learn tasks in the form of a knowledge base (KB) that serves as information for the planner, and to devise strategies for these planners to use this information effectively in quasi-static settings. This representation is learned from observing human demonstrations that are taken as the initial information seed needed to reason about the functionality of a manipulation task. This demonstration seed is further exploited through computation to simulate self-experience and improve upon the learned strategies.

I motivate and evaluate the work in the context of two main applications that involve collaboration with humans: remote robot operation and shared workspace collaborative robotics in manufacturing.

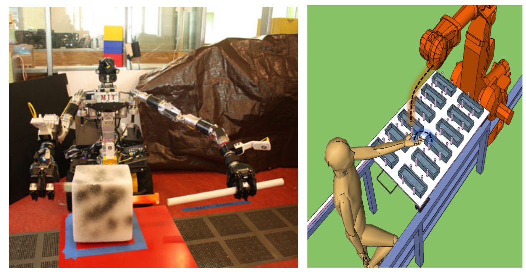

Fig. 1: Domains of application: (a) Remote robot operation in shared autonomy, (b) Shared workspace collaboration.

I focus my thesis work on the aforementioned three aspects of the manipulation problem. Section II summarizes my work to date on the first two aspects: (1) learning and planning, and (2) supervised execution; whereas Section III introduces the proposed work for (3) planning new strategies under uncertainty.

II. Work To Date

Figure 2 presents a systems-level diagram, where the following components are highlighted:

1) Learning and Planning: Human demonstrations are used to learn a representation of the activity, which is encoded in the KB. I have proposed C-LEARN [1], a method for learning geometric constraints from demonstrations for multi-step functional manipulation tasks with multiple end effectors in quasi-static settings. It contributes the ability to (1) learn geometric constraints, including CAD constraints (parallel, perpendicular, move in a line), (2) transfer a skill learned with a source robot to a target robot without requiring new demonstrations, and (3) use the planner and the multi-step representation to formulate a series of motion suggestions to be presented to an operator in shared autonomy.

Learning from demonstrations (LfD) has been shown to be a successful method for non-experts to teach manipulation tasks to robots. These methods typically build generative models from demonstrations and then use regression to reproduce skills. However, this approach has limitations in capturing hard geometric constraints imposed by the task. On the other hand, while sampling and optimization-based motion planners exist that reason about geometric constraints, these are typically carefully hand-crafted by an expert. C-LEARN addresses this technical gap. The system builds a knowledge base for reaching and grasping objects,

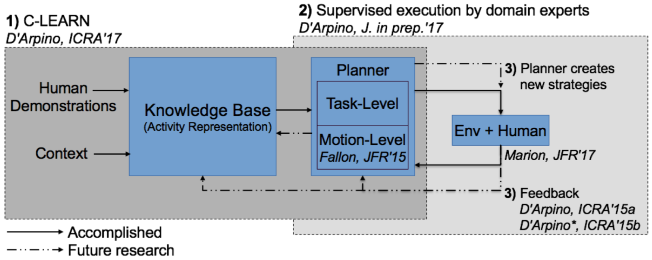

Fig. 2: High-level diagram of the system

which is then leveraged to learn multi-step tasks from a single demonstration.

2) Human-in-the-loop execution: Motivated by the target applications, the execution of the manipulation tasks is realized in collaboration with humans, either in close physical proximity or by remote operation though an user interface. I leverage the shared autonomy framework designed at MIT for the DARPA Robotics Challenge (DRC) – in particular, an optimization-based motion planner [2] and a user interface [3] – to implement a two-step workflow of planning and execution, where motion plans are shown to the operator and executed upon approval.

I have integrated the motion suggestions produced by C-LEARN [1] with this execution workflow [2][3] to evaluate the advantages of increased levels of autonomy for remote robot control, and conducted a within-subjects user study with an expert population to evaluate this method. Detailed results will be presented in a journal paper currently in preparation.

III. Proposed Work

3) Planning new strategies under uncertainty: The previous steps implement an end-to-end framework for learning and planning multi-step manipulation tasks. This step involves generalizing C-LEARN to new tasks that are to be executed in the presence of uncertainty. Uncertainty is due to the lack of knowledge about the effects and feasibility of new actions that were not present in the demonstrations, and that are now required to handle variations of the task configuration or of human motions in close proximity. While methods to model and bound the uncertainty have been explored, such as prediction of human motions [4][5][6], still to be investigated is how to give feedback to the planner and to the knowledge base learned from demonstrations, for the case where a pre-designed manipulation strategy is not feasible.

The next step involves creating a method to use the knowledge base in generating new strategies that are variations of the originally learned multi-step manipulation task. As illustrated in Figure 2 , this third phase involves expanding the KB and the planner to now support this source of uncertainty, while still accomplishing the manipulation task. From the point of view of the KB, this demands a representation that learns the geometric constraints in a flexible fashion that allows the planner to make use of them for performing modifications to the plan according to the new task configuration.

To summarize, C-LEARN achieved generalization across different robots and different positions and orientations of the objects involved, but it does not support generalization across tasks. This means that it is only able to execute a multi-step manipulation task using the same order of steps that was learned. Creating new strategies for the same task or generating new possible tasks with similar objects remains a challenge in the field of artificial intelligence and robotics.

IV. Survey of Related Work

The following list is a representative sample of related work on various topics. Learning and LfD: [7] [8] [9] [10] [11] [12] [13] [14]; learning from one demonstration: [15] [16] [17] [18]; planning with constraints and TAMP: [19] [20] [21] [22] [23] [24] [25] [26]; multi-step constrained manipulation: [27] [28] [29] [30] [31] [32] [33]; shared autonomy systems: [34] [35] [36] [37] .

References

[1] **Claudia Pérez-D'Arpino** and J. A. Shah, "C-LEARN: Learning geometric constraints from demonstrations for multi-step manipulation in shared autonomy," in *IEEE ICRA 2017*, 2017.

[2] M. Fallon, S. Kuindersma, S. Karumanchi, M. Antone, T. Schneider, H. Dai, **C. Pérez-D'Arpino**, R. Deits, M. DiCicco, D. Fourie, T. Koolen, P. Marion, M. Posa, A. Valenzuela, K.-T. Yu, J. Shah, K. Iagnemma, R. Tedrake, and S. Teller, "An architecture for online affordance-based perception and whole-body planning," *Journal of Field Robotics*, vol. 32, no. 2, pp. 229–254, 2015.

[3] P. Marion, M. Fallon, R. Deits, A. Valenzuela, **C. Pérez-D'Arpino**, G. Izatt, L. Manuelli, M. Antone, H. Dai, T. Koolen, J. Carter, S. Kuindersma, and R. Tedrake, "Director: A user interface designed for robot operation with shared autonomy," *Journal of Field Robotics*, vol. 34, no. 2, pp. 262–280, 2017.

[4] **Claudia Pérez-D'Arpino** and J. A. Shah, "Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification," in *IEEE ICRA 2015*, pp. 6175–6182, 2015.

[5] V. V. Unhelkar*, **Claudia Pérez-D'Arpino***, L. Stirling, and J. A. Shah, "Human-robot co-navigation using anticipatory indicators of human walking motion," in *IEEE ICRA 2015*, 2015.

[6] **Claudia Pérez-D'Arpino** and J. A. Shah, "Fast motion prediction for collaborative robotics," *Joint Conference on Artificial Intelligence (IJCAI-16), Doctoral Consortium*, 2016.

[7] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469 – 483, 2009.

[8] S. Calinon and A. Billard, "A probabilistic programming by demonstration framework handling constraints in joint space and task space," in *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pp. 367–372, IEEE, 2008.

[9] G. Y. C. Bowen and R. Alterovitz, "Asymptotically optimal motion planning for learned tasks using time-dependent cost maps," *IEEE Transactions on Automation Science and Engineering*, vol. 12, pp. 171–182, Jan 2015.

[10] D. Berenson, S. S. Srinivasa, and J. Kuffner, "Task Space Regions: A framework for pose-constrained manipulation planning," *The International Journal of Robotics Research*, p. 14351460, 2011.

[11] C. Li and D. Berenson, "Learning object orientation constraints and guiding constraints for narrow passages from one demonstration," *ISER 2016*, 2016.

[12] B. Akgun, M. Cakmak, K. Jiang, and A. L. Thomaz, "Keyframe-based learning from demonstration," *International Journal of Social Robotics*, vol. 4, no. 4, pp. 343–355, 2012.

[13] A. Kurenkov, B. Akgun, and A. L. Thomaz, "An evaluation of GUI and kinesthetic teaching methods for constrained-keyframe skills," in *IEEE IROS 2015*, pp. 3608–3613, 2015.

[14] M. Laskey, C. Chuck, J. Lee, J. Mahler, S. Krishnan, K. Jamieson, A. Dragan, and K. Goldberg, "Comparing human-centric and robot-centric sampling for robot deep learning from demonstrations," in *IEEE ICRA 2017*, 2017.

[15] S. Alexandrova, M. Cakmak, K. Hsiao, and L. Takayama, "Robot programming by demonstration with interactive action visualizations," *RSS 2014*, 2014.

[16] L. Fei-Fei, R. Fergus, and P. Perona, "One-shot learning of object categories," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, pp. 594–611, April 2006.

[17] C. Groth and D. Henrich, "One-shot robot programming by demonstration by adapting motion segments," in *IEEE ROBIO 2014*, pp. 1068–1075, Dec 2014.

[18] C. Finn, T. Yu, J. Fu, P. Abbeel, and S. Levine, "Generalizing skills with semi-supervised reinforcement learning," in *ICLR*, 2017.

[19] L. Pais, K. Umezawa, Y. Nakamura, and A. Billard, "Learning robot skills through motion segmentation and constraints extraction," in *HRI Workshop on Collaborative Manipulation*, 2013.

[20] J. Schulman, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, and P. Abbeel, "Finding locally optimal, collision-free trajectories with sequential convex optimization.," in *RSS 2013*, 2013.

[21] Y. Wu and Y. Demiris, "Towards one shot learning by imitation for humanoid robots," in *IEEE ICRA 2010*, pp. 2889–2894, IEEE, 2010.

[22] N. Mehr, R. Horowitz, and A. Dragan, "Inferring and assisting with constraints in shared autonomy," in *Conference on Decision and Control (CDC)*, 2016.

[23] J. Silvrio, L. Rozo, S. Calinon, and D. G. Caldwell, "Learning bimanual end-effector poses from demonstrations using task-parameterized dynamical systems," in *IEEE IROS 2015*, pp. 464–470, 2015.

[24] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "CHOMP: Gradient optimization techniques for efficient motion planning," in *IEEE ICRA 2009*, pp. 489–494, 2009.

[25] L. P. Kaelbling and T. Lozano-Perez, "Integrated task and motion planning in belief space," in *International Journal of Robotics Research*, 2013.

[26] C. Devin*, A. Gupta*, T. Darrell, P. Abbeel, and S. Levine, "Learning modular neural network policies for multi-task and multi-robot transfer," in *IEEE ICRA 2017*, 2017.

[27] S. Niekum, G. K. S. Osentoski, and A. G. Barto, "Learning and generalization of complex tasks from unstructured demonstrations," in *IEEE IROS 2012*, pp. 5239–5246, Oct 2012.

[28] M. Phillips, V. Hwang, S. Chitta, and M. Likhachev, "Learning to plan for constrained manipulation from demonstrations," *Autonomous Robots*, vol. 40, no. 1, pp. 109–124, 2016.

[29] N. Somani, A. Gaschler, M. R. A. Perzylo, and A. Knoll, "Constraint-based task programming with CAD semantics: From intuitive specification to real-time control," in *IEEE IROS 2015*, pp. 2854–2859, 2015.

[30] M. Dogar, R. A. Knepper, A. Spielberg, C. Choi, H. I. Christensen, and D. Rus, "Towards coordinated precision assembly with robot teams," in *ISER 2014*, 2014.

[31] M. Dogar, A. Spielberg, S. Baker, and D. Rus, "Multi-robot grasp planning for sequential assembly operations," in *IEEE ICRA 2015*, 2015.

[32] S. Pillai, M. R. Walter, and S. Teller, "Learning articulated motions from visual demonstrations," in *RSS 2014*, 2014.

[33] T. Lozano-Perez and L. P. Kaelbling, "A constraint-based method for solving sequential manipulation planning problems," in *IEEE/RSJ IROS*, 2014.

[34] T. B. Sheridan, *Telerobotics, Automation, and Human Supervisory Control*. Cambridge, MA, USA: MIT Press, 1992.

[35] M. Johnson, B. Shrewsbury, S. Bertrand, T. Wu, D. Duran, M. Floyd, P. Abeles, D. Stephen, N. Mertins, A. Lesman, J. Carff, W. Rifenburgh, P. Kaveti, W. Straatman, J. Smith, M. Griffioen, B. Layton, T. de Boer, T. Koolen, P. Neuhaus, and J. Pratt, "Team IHMC's Lessons Learned from the DARPA Robotics Challenge Trials," *Journal of Field Robotics*, vol. 32, no. 2, pp. 192–208, 2015.

[36] O. Khatib, X. Yeh, G. Brantner, B. Soe, B. Kim, S. Ganguly, H. Stuart, S. Wang, M. Cutkosky, A. Edsinger, P. Mullins, M. Barham, C. R. Voolstra, K. N. Salama, M. L'Hour, and V. Creuze, "Ocean one: A robotic avatar for oceanic discovery," vol. 23, pp. 20–29, 2016.

[37] A. Romay, S. Kohlbrecher, A. Stumpf, O. von Stryk, S. Maniatopoulos, H. Kress-Gazit, P. Schillinger, and D. C. Conner, "Collaborative autonomy between high-level behaviors and human operators for remote manipulation tasks using different humanoid robots," *Journal of Field Robotics*, 2016.