Achieving Robustness in a Drawer Manipulation Task by using High-level Feedback instead of Planning

Manuel Baum and Oliver Brock

***Abstract*— Robotic manipulation behavior should be robust to disturbances that violate high-level task-structure. Such robustness can be achieved by constantly monitoring the environment to observe the discrete high-level state of the task. This is possible because different phases of a task are characterized by different sensor patterns and by monitoring these patterns a robot can decide which controllers to execute. This eliminates the need to plan a temporal sequence of those controllers and makes the behavior robust to unforeseen disturbances. We implement this idea as a probabilistic filter over discrete states where each state directly activates a controller. Based on this framework we present a robotic system that is able to robustly open a drawer and grasp tennis balls from it.**

Preferred typ of presentation: Oral

I. Motivation and Problem Definition

Planning makes assumptions about the temporal evolution of the world when robots act in it. It creates a *model* of the world over future time-spans, but such models may be invalidated by the numerous unforeseen real-world contingencies. Instead of relying on such model predictions, we argue it is better to follow the idea that "the world is its own best model" [1]. Why should we try to predict the world state and make decisions based on anticipated contingencies when we can instead simply react to the state of the world in the moment the contingencies actually occur?

II. Related Work

Classical planning approaches assume a determinstic world and explicitly predict its state after a linear sequence of robot actions. But linear plans likely fail when the world does not develop as expected. Contingent plans are more flexible, tree-like structures [2] that allow some robustness to unforeseen disturbances, but still only consider a limited set of temporal evolutions and fail if actual contingencies were not predicted. The idea of reactive planning is to avoid such predictions and instead act based on the environment's state [3]. Recently Robust Logical-Dynamical Systems[4] showed behavior with remarkable robustness to disturbances by constantly estimating the world state. But do we really need logical-geometric state-estimation?

III. Own Approach and Contribution

The purpose of state estimation is to extract information from sensor input that is sufficient to choose appropriate actions in the current context. If we use feedback controllers that can directly solve all sub-problems of a task, then for certain manipulation tasks the only state estimate we need is to identify which controller to activate. For such problems we can avoid complex logical-geometric state-estimation.

Both authors are with the Robotics and Biology Laboratory, Technische Universität Berlin, Germany and the Cluster of Excellence "Science of Intelligence" (SCIoI). Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2002/1 "Science of Intelligence" - project number 390523135.

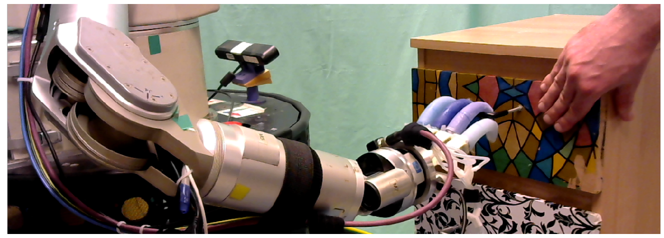

Fig. 1. A person holds a drawer shut while a robot tries to open it. Irrespective of such disturbances the robot can reliably open the drawer.

We present a method that can be used in complex manipulation tasks like opening a drawer and grasping tennis balls from it. It is based on an HMM which filters the discrete high-level state of the task and activates feedback controllers based on that state. It does not rely on a pre-defined sequence of controllers and does not implement prior knowledge about such a sequence as the HMM's state transition matrix is almost a diagonal matrix with a small uniform off-diagonal term. Yet the system is able to create complex interactions with the environment such as loops to re-establish lost grasps on the drawer handle or to re-grasp tennis balls when they fell out of the hand. By use of high-level feedback the system so robust that it can detect [5] and recover from significant interferences such as forcefully removing the end-effector from the handle during a successful grasp, holding the drawer shut during attempts to open it and randomly poking at the end-effector. The system recovers from such interferences even without explicitly programmed recovery behaviors.

References

[1] R. A. Brooks, "Elephants don't play chess," *Robotics and Autonomous Systems*, vol. 6, pp. 3–15, 1990.

[2] L. Pryor and G. Collins, "Planning for Contingencies: A Decision-based Approach," *Journal of Artificial Intelligence Research*, vol. 4, pp. 287–339, May 1996.

[3] R. J. Firby, "An Investigation into Reactive Planning in Complex Domains," in *Proceedings of the Sixth National Conference on Artificial Intelligence - Volume 1*, ser. AAAI'87. AAAI Press, pp. 202–206.

[4] C. Paxton, N. Ratliff, C. Eppner, and D. Fox, "Representing robot task plans as robust logical-dynamical systems," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 5588–5595.

[5] M. Baum and O. Brock, "Achieving robustness by optimizing failure behavior," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 5806–5811.