Learning Task- and Touch-based Grasping

Yasemin Bekiroglu, Dan Song, Lu Wang and Danica Kragic

In order to equip robots with goal-directed grasping ability, the integration of high-level task information with low-level sensory data is needed. For example, if a robot is given a task, e.g., *pour me a cup of coffee*, it needs to 1) make decision on which object to use, 2) how the hand should be placed around the object, and 3) how much gripping force should be applied so that the subsequent manipulation is feasible and stable for the pouring action. Several sensory streams (visual, proprioceptive and haptic) are relevant for these three steps. The problem domain and hence the state space becomes high-dimensional involving both continuous and discrete variables with complex relations. We study how these can be encoded in a suitable manner using probabilistic generative models so that robots can achieve stable and robust goal-directed grasps by exploiting feedback loops from multisensory data.

To enable goal-directed grasp planning, grounding symbolic representation of goal state (e.g. to pour) to continuous representation of low-level sensory feedback (e.g. grasping pose) is the main challenge. We resolve this by adopting a probabilistic graphical model, a Bayesian network (BN), which encode relations between variables using conditional probabilistic distributions. Such distributions do not require the variables to comply to the same underlying representations. Therefore, a symbolic task goal can be grounded to a continuous grasp pose through inferring the conditional $p(pose|task)$.

In this work, the robot (a Schunk Dextrous hand attached on a Kuka arm) learns such a goal-directed grasp BNs by a combination of self-exploration and human-supervision. Exploration enables the robot to learn about its own sensorimotor ability (how to grasp an object to stably lift and manipulate it), while human tutoring helps the robot to associate its sensorimotor ability to high-level goals. During exploration, the robot collects visual, proprioceptive (joint sensor) and haptic (tactile sensing array) data by executing a set of grasps on a set of objects. During human supervision, each grasp is labeled with tasks it affords by a human tutor. The database collected is the instantiation of a set of variables $\{O, A, H, T\}$, where T denotes task variables, O, A, H denote object (vision), action (proprioception) and haptic (or tactile) feature sets. A learned BN using this data encodes the joint distribution $p(T, O, A, H)$. This BN allows inference of conditional probability of task success $p(T|X)$, where $X \subseteqq \{O, A, H\}$ can be a full or partial observation of all sensory data. The BN also allows inference of the class conditional distributions such as $p(O, A|T)$ which is the basis for goal-directed object selection and grasp planning.

The table presents evaluation on task inference performance using the area under the ROC curves (mean (std)). We observe high task classification performances with both full and partial observations. With this as a basis, we can design a 2-loop grasp adaptation to allow goal-directed grasp selection in an efficient manner. In loop 1, the robot predicts grasp success in simulation environment by inferring $p(T|O, A)$ before execution, and in loop 2, the robots executes a grasp selected in loop 1, obtains the H reading from tactile sensor, then checks grasp success by inferring $p(T|O, A, H)$ before lifting the object. Fig. 1 demonstrates such a process. From left to right, the grasps (1-3) are adapted because the inferred task likelihoods are very low (see the location of each grasp on the likelihood map below). Once a good grasp (4) is selected in first loop, it is executed on real platform and H data is available. The task probability with full observation predicts this as a bad grasp (see the location on the likelihood map of the tactile feature), which is confirmed by the failing subsequent manipulation (a 90° rotation). A replan is triggered until a good grasp (5) is found.

Task (Nsamples)	Full	partial
Hand-over (1026)	0.90 (0.04)	0.86 (0.01)
Pouring (1143)	0.88 (0.02)	0.86 (0.02)
Dishwashing (831)	0.92 (0.01)	0.86 (0.02)

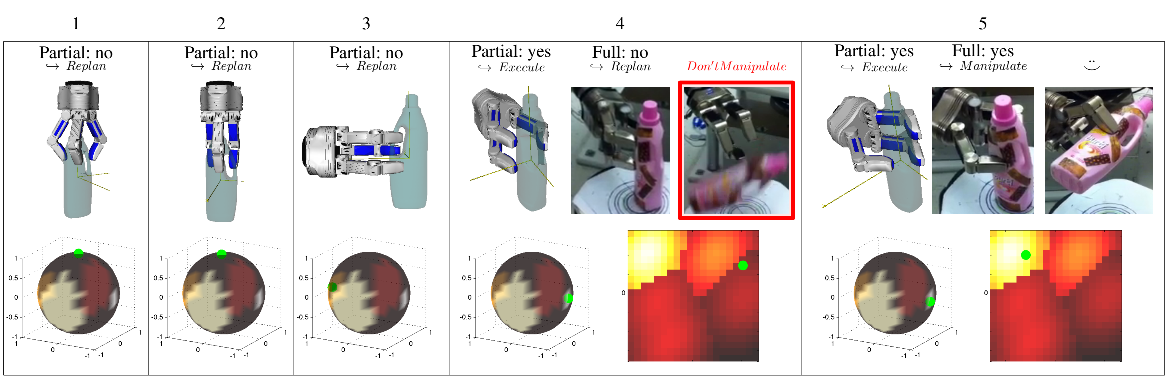

Fig. 1. Two-loop grasp adaptation for pouring. Top is grasp hypotheses sequentially produced by a grasp planner. Bottom is the likelihood map of grasp position (sphere) and tactile image (square) conditioned on the task. Dark color means low likelihood. Green dot represents the grasp hypothesis above.