Learning for Tactile Manipulation

Tucker Hermans, Filipe Veiga, Herke van Hoof,
Janine Hoelscher, and Jan Peters

November 8, 2014

Tactile sensing affords robots the opportunity to dexterously manipulate objects in-hand without the need of strong object models and planning. Our demonstration focuses on learning for tactile, in-hand manipulation by robots. We address learning problems related to the control of objects in-hand, as well as perception problems encountered by a robot exploring its environment with a tactile sensor. We demonstrate applications for three specific learning problems: learning to detect slip for grasp stability, learning to reposition objects in-hand, and learning to identify objects and object properties through tactile exploration.

We address the problem of learning to detect slip of grasped objects. We show that the robot can learn a detector for slip events which generalizes to novel objects. We leverage this slip detector to produce a feedback controller that can stabilize objects during grasping and manipulation. Our work compares a number of supervised learning approaches and feature representations in order to achieve reliable slip detection.

Tactile sensors provide observations of high enough dimension to cause problems for traditional reinforcement learning methods. As such, we introduce a novel reinforcement learning (RL) algorithm which learns transition functions embedded in a reproducing kernel Hilbert space (RKHS). The resulting policy search algorithm provides robust policy updates which can efficiently deal with high-dimensional sensory input. We demonstrate the method on the problem of repositioning a grasped object in the hand.

Finally, we present a method for learning to classify objects through tactile exploration. The robot collects data from a number of objects through various exploratory motions. The robot learns a classifier for each object to be used during exploration of its environment to detect objects in cluttered environments. Here again we compare a number of learning methods and features present in the literature and synthesize a method to best work in human environments the robot is likely to encounter.

Users will be able to interact with a robot hand by giving it objects to grasp and attempting to remove these objects from the robot. The hand will also perform some basic in-hand manipulation tasks such as rolling the object between the fingers and rotating the object about a fixed grasp point. Users

will also be able to interact with a touch sensor capable of classifying objects as well as semantic events such as slipping from a stable contact location.