

Toward Learning to Press Doorbell Buttons

Liping Wu, Vladimir Sukhoy and Alexander Stoytchev

Developmental Robotics Laboratory
Department of Electrical and Computer Engineering
Iowa State University
{lipingwu, sukhoy, alexs}@iastate.edu

Introduction

To function in human-inhabited environments a robot must be able to press buttons. There are literally thousands of different buttons, which produce various types of feedback when pressed. This work focuses on doorbell buttons, which provide auditory feedback. Our robot learned to predict if a specific pushing movement would press a doorbell button and produce a sound. The robot explored different buttons with random pushing behaviors and perceived the proprioceptive, tactile, and acoustic outcomes of these behaviors.

Previous related work in robotics has focused mainly on visual feedback. Thomaz (2006) proposed using social guidance to teach a robot which button it has to press. Miura et al.'s (2005) robot searched for an image template of a button in the region pointed to by a human and pressed it. Nguyen et al. (2009) used PPS-Tags to make switches more visible so that a robot can press them more easily. In contrast, the robot in our study relied exclusively on multi-modal feedback, i.e., auditory, proprioceptive and tactile feedback.

In Psychology, Hauf and Ascherleben (2008) reported that 9 months old human infants already anticipate that pressing a button would produce interesting sounds or lights. E.J. Gibson (1988) suggested that humans actively apply exploratory behaviors on objects to learn their affordances.

Experimental Setup

The experiments were performed using the upper-torso humanoid robot shown in Fig. 1(a), which has two 7-dof Barrett Whole Arm Manipulators (WAMs) for arms and two three-finger hands also made by Barrett Technology. The robot experimented with the 3 doorbell buttons shown in Fig. 1(b) and performed 400 trials with each, for a total of 1200 trials. The experiments were performed in 4 rounds. In each round the robot performed 100 trials with each of the buttons. Each trial took approximately 10 seconds to complete. During each trial the robot first pushed the area around the button and then slid its finger five times over the area. The start position of pushing behaviors was fixed above the button; the end position was randomly sampled around the button. The sliding behaviors were implemented by moving the robot's finger in a random direction from the stop position of the previous behavior (either a push or a slide). For

Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Experimental setup: (a) The upper-torso humanoid robot; (b) The three buttons explored by the robot.

each trial the robot recorded auditory, proprioceptive, and tactile data and the temporal intervals for all behaviors. The Audio-Technica U853AW Hanging Microphone mounted in the robot's head recorded sound at 44.1 kHz. The Analog Devices ADXL345 3-axis accelerometer attached to the robot's finger recorded the tactile data at 400 Hz. Proprioceptive data, in the form of joint angles and torques, was recorded from the WAM arm at 500 Hz.

Methodology

For each of the three sensory modalities the robot detected 'interesting' events and timestamped them. If the robot heard a buzzing sound, it recorded an auditory event. If the robot noticed that the torque on any joint exceeded its predefined limit then it recorded a proprioceptive event, interrupted the current behavior, and started the next one. Spikes in the magnitude of the accelerometer's readings were recorded as tactile events. The goal of the experiments was three-fold: 1) to learn which behaviors trigger different events; 2) to localize events in space for each button; 3) to investigate if and how events from different modalities co-occur in time.

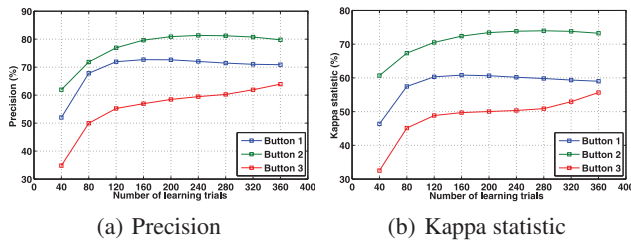


Figure 2: Learning curves for predicting if a pushing behavior triggers an auditory event as a function of the number of learning trials: (a) precision — number of true positives as a percentage of all positives; (b) kappa statistic — a more robust measure of how good the robot’s predictions are, which takes into account the probability of chance agreement. For each data point, the values were averaged over 1000 runs of 10-fold cross-validation.

Each behavior was parameterized by its start and desired destination positions in joint space. The joint space coordinates of the destination position of a pushing behavior were used as features to learn to predict whether the robot will hear the buzzer if it pushes in that direction. The robot trained a separate k -NN classifier for each button to facilitate the learning task. Other classifiers were tested as well, but k -NN was typically among the best.

Results

The robot learned to predict if a pushing behavior triggers an auditory event. By varying the number of training trials, the robot found that approximately 160 learning trials were required to get the best possible performance with both rectangular buttons (see Fig. 2). For the round button, which is smaller and thus harder to press, the robot continuously gained new insights over all available trials.

To localize events in space the robot used forward kinematics to find the Cartesian coordinates of its fingertip when the event was detected. Fig. 3 shows the spatial distribution of auditory events around each of the three buttons. The events were mapped to the pictures of the buttons by applying an affine transform. The affine transform was computed separately using several anchor points for which the pixel coordinates and the Cartesian coordinates were known. Fig. 3 shows that the auditory events are localized around the functional components of the buttons. The apparent scattering of the points is due to the fact that in some cases the robot pressed the button with the side of its finger and not with the fingertip, which was the only thing that was tracked.

For the trials with auditory events the robot measured the temporal intervals between events in different modalities (see Fig. 4). The results indicate that auditory events co-occur with proprioceptive and tactile ones. This suggests a possibility for future work: using these cross-modal temporal dependencies and their co-occurrence in space, detect if the functional components of a button are working properly.

Conclusion and Future Work

This paper showed that a robot can learn to predict if pushing doorbell buttons in a certain way triggers an auditory event. Depending on the button, a robot may continuously gain new insights from unfocused random exploration or it may reach a limit and make no further learning progress.




Figure 3: Auditory events localized in space for each button for both pushing and sliding behaviors.

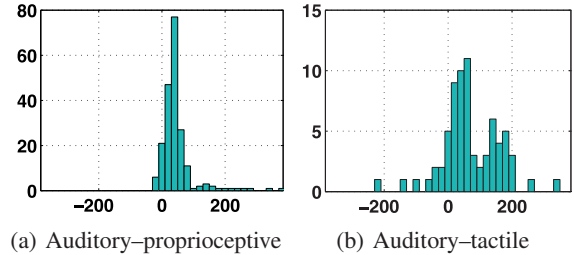


Figure 4: Histograms for the temporal intervals between auditory events and proprioceptive events (a), and auditory and tactile events (b) for both pushing and sliding behaviors. The bin size of each histogram is 20 ms.

In future work, we will investigate the effect of active learning on the speed of learning. Preliminary results indicate that a robot learns faster when it selects behaviors that are more likely to trigger an auditory event. Given the ability to tell in advance if a specific pushing behavior will result in an auditory event, a robot can generate movements that press buttons more efficiently. Different active learning policies may be evaluated in this context.

The spatial co-occurrence patterns of events from different modalities that a robot observes while exploring, may provide the data necessary for segmenting objects into functional parts and identifying these parts in novel objects. The temporal co-occurrence patterns may provide the data necessary for the robot to optimize its behaviors for triggering specific types of events in different modalities. For additional details and results, see <http://www.ece.iastate.edu/~lipingwu/AAAI10.html>.

References

- Gibson, E. 1988. Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge. *Annual review of psychology* 39(1):1–42.
- Hauf, P., and Aschersleben, G. 2008. Action-effect anticipation in infant action control. *Psychological Research* 72(2):203–210.
- Miura, J.; Iwase, K.; and Shirai, Y. 2005. Interactive teaching of a mobile robot. In *IEEE ICRA*, 3378–3383.
- Nguyen, H.; Deyle, T.; Reynolds, M.; and Kemp, C. 2009. PPS-tags: Physical, Perceptual and Semantic tags for autonomous mobile manipulation. In *Proceedings of the IROS 2009 Workshop on Semantic Perception for Mobile Manipulation*.
- Thomaz, A. 2006. *Socially guided machine learning*. Ph.D. Dissertation, Massachusetts Institute of Technology.