Detecting object boundaries using focus, zoom, shadows, and active probing

Paul Fitzpatrick and Eduardo Torres-Jara

What: We are building an active vision head with controllable zoom and focus, to be tightly integrated with a robotic manipulation system capable of actively probing the structure of its environment. We are seeking to address the perceptual difficulties associated with robot manipulation.

Why: This work is part of a larger project to build a robot capable of manipulating unfamiliar objects in an unstructured environment. This is a very challenging task since perceptual uncertainty translates directly into clumsy motion. For example, if the robot estimates the boundary of an object incorrectly, then there is a limit to what can be salvaged by clever control algorithms.

Our goal is to give our robot the perceptual abilities it needs to actively resolve ambiguity in its environment, when the passive vision algorithms it uses fail. A key example of this is figure/ground separation [2]. By equipping our robot with a camera with a narrow, controllable depth of field, the robot can bring a questionable edge into focus and blur out distracting background. By giving it the ability to zoom, the robot can look closely around the edge for hints of texture (see Figure 1). And finally, by bringing its manipulator into view, it can physically probe regions and cast its own controlled shadows.

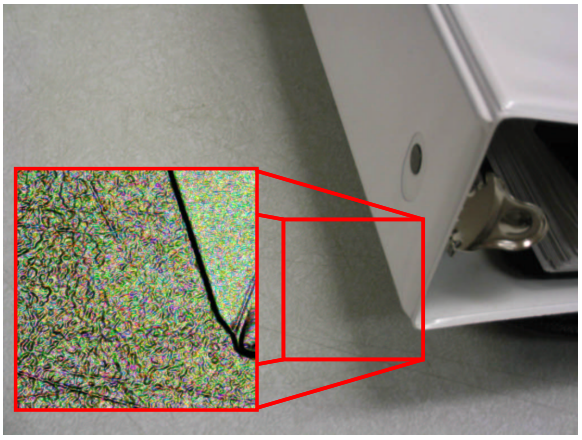

Figure 1: Shadows tend to be low-frequency illumination changes. With sufficient zoom and good focus, higher-frequency texture can be found and used to show the shadow is not a physical edge. Comparing cast edges with true edges can suggest the presence of physical gaps – very useful for the purposes of manipulation.

How: Our approach is to fixate a possible object edge using pan, tilt, and zoom, and then hunt with focus to find high-frequency texture [4]. If this is present (and with sufficient zoom most surfaces do have some texture), shadows cast on the surface are easy to discount as low-frequency illumination changes. This is particularly true for moving shadows, where the high-frequency texture pattern is fixed as a low-frequency illumination edge passes across it. In difficult cases, we plan to have the robot intervene with its own arm, either directly or by casting its own shadow.

Progress: Previous work on Cog showed manipulation for detecting object boundaries experimentally [1]. In that work, shadows were dealt with as a nuisance. Now we want to make them a positive benefit. We are currently working with off-line pictures while the active head is being built. At the moment we are working on the interface to control zoom and focus in the active head.

Future: We will integrate our active vision system with our manipulator. The manipulator is comprised of an arm and a hand, both with force control [3]. The hand is mechanically compliant so that the robot can actually touch the environment without damaging itself or the objects. These features together with the advantages of the vision system will allow us to operate the robot in somewhat unstructured environments.

Research Support: Funds for this project were provided by DARPA as part of the "Natural Tasking of Robots Based on Human Interaction Cues" project under contract number DABT 63-00-C-10102, and by the Nippon Telegraph and Telephone Corporation as part of the NTT/MIT Collaboration Agreement.

References:

[1] P. Fitzpatrick and G. Metta. Grounding vision through experimental manipulation. *Philosophical Transactions of the Royal Society: Mathematical, Physical, and Engineering Sciences*, 361(1811):2165–2185, 2003.

[2] I. Nourbakhsh, D. Andre, C. Tomasi, and M. Genesereth. Mobile robot obstacle avoidance via depth from focus. *Robotics and Autonomous Systems*, 22:151–158, June 1997.

[3] E. Torres-Jara and J. Banks. A simple and scalable force actuator. *35th International Symposioum on Robotics*, March 2004.

[4] J. Y. Zheng, T. Sakai, and N. Abe. Guiding robot motion using zooming and focusing. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 3:1076–1082, June 1996.