Self-Adaptive Perception Systems for Varying Environments

David Hershberger

Problem: Most robot systems have parameters which require tuning in order to work properly. Tuning these parameters autonomously is important for adapting to changes within a robot system or in its environment. Changes affecting robots might include degradation of equipment such as cameras and light sensors gradually losing intensity, antennae being damaged, or gyroscopes losing accuracy. Another type of change is the move from a laboratory environment to the field: parameters tuned in the lab normally need adjusting once a robot is deployed. Recently Nomad was deployed in Antarctica, and the meteorite-finding vision software which worked perfectly in the lab took several weeks of adjustment before working in the field. Change also happens when robot software is ported from one robot to another: different physical configurations, sensors, and speeds must all be adjusted for.

Handling varying environments is important because good parameter settings learned for one situation may not work at all in another. While working on a visual servoing system at JSC last summer, a frequent source of frustration was the need to tune parameters to keep it working when the environment changed: lighting, background objects, and angles of fiducials. Environmental variations can affect perception in other robots as well. Stereo vision algorithms rely on visual texture to produce accurate depth measurements, so in a sandy landscape with few rocks a vision algorithm will need different adjustments than are appropriate for a rocky environment. Likewise camera settings should be adjusted differently when looking nearly into the sun than when looking away from the sun, or into a region of shadow. Perhaps most importantly, there may be completely unanticipated relationships such as temperature variations affecting sensor performance. A learning system which does not take environmental changes into account and assumes that there is one unique set of best parameters will be forced to continually re-learn good settings as it moves from one situation to another and back again.

Impact: This project will advance the state of the art in automatic parameter optimization primarily in its use of sensed environmental features to learn which situations require different parameter settings for best performance. Existing parameter optimization systems have to continously experiment and re-learn as they move from one environment to another. The proposed system will discover in which regions of the environment the system works best with which parameter settings, allowing a robot to move from one area to another and use the best settings in each. For robust robot perception, this project has the potential to greatly improve performance in a wide variety of systems and applications.

State of the Art: Our search for parameter optimization algorithms thus far has revealed a large body of work addressing many different aspects of the problem. “Q2”, by Andrew Moore and Jeff Schnieder [3], has been applied successfully to plant optimization problems. Q2 addresses the need for active experimentation, but not the need for safety assurances. Other algorithms include those discussed in [2], such as Kaelbling’s IE algorithm [1], and the safety-oriented COMAX algorithm. These algorithms focus on the idea that some experiments are more expensive than others and try to minimize the cost of finding the best parameters.

Approach: The proposed system will adapt parameter values automatically to continually improve performance. Varying environments will be handled by learning which different situations require different parameter settings, and learning the best settings for each situation. This approach avoids the problem of re-learning settings for situations the robot has experienced before. In order to adjust parameters automatically, the system will select experiments (which parameters to change, and by how much), run them, then adjust its estimate of which settings are best based on the results. An important aspect in experiment selection is safety: ensuring that previously untested parameter values will not damage the equipment. To learn which situations work best in what settings, aspects of the environment must be measured, for example image brightness, dominant colors, sun angle, terrain slope, and terrain texture. Correlations between changes in these features and changes in performance must be discovered in order to either divide possible environments into different types. Once a division is made, parameters can be learned separately for each type of environment. The process can continue dividing and learning until the system works well in all environments encountered.

An alternative approach to learning which settings work in which environments currently under investigation is a modification of Q2 enabling it to hold certain parameters fixed while optimizing the other parameters. These fixed parameters would actually be measured features of the environment.

Future Work: This research is just beginning, so future work includes implementing and evaluating one or more optimization techniques on synthetic data, then implementing and evaluating performance improvement on actual robot systems. The initial robot system targetted is that of the visual perception system in the Distributed Visual Servoing project.

Selected References:

[1] Kaelbling. *Learning in Embedded Systems*. PhD thesis, Stanford University, 1990.

[2] Andrew W. Moore and Jeff Schneider. Memory-based stochastic optimization. In *NIPS-95*, 1995.

[3] Andrew W. Moore, Jeff Schneider, Justin Boyan, and Mary Soon Lee. Q2: Memory-based active learning for optimizing noisy continuous functions. In *Proceedings of the International Conference of Machine Learning*, 1998.