Autonomous Exploration: An Integrated Systems Approach

Marc Bolduc, Eric Bourque, Gregory Dudek, Nicholas Roy, and Robert Sim
Centre for Intelligent Machines
McGill University
Montreal, Quebec, Canada

Introduction[1]

Mobile robotic systems offer an ideal platform for testing and implementing many of the concepts developed in more abstract artificial intelligence. Robotic systems embody a complex interaction of computation, perception and actuation that depend upon such familiar tasks as recognition and reaction. In order for robots to perform real-world tasks such as navigation, localization and exploration, the subsystems of motion, sensing and computation must be merged into a single, realizable unit that uses the different techniques together.

Our group is investigating problems in the domain of computational perception, in the context of mobile robotics. In particular, we are concerned with environment exploration, and map construction. We are using the AAAI 1997 Mobile Robot competition as an opportunity to test a number of implementations of systems in navigation, spatial reasoning and perception.

Methodology

We are using a modular, distributed software architecture for the implementation of the various components of the exploration and object recognition tasks. This implementation is distributed across a network, with the critical processes running on the robotic hardware and off-board planning processes allowed to run on remote hardware, thus taking advantage of processor power available over a network.

To support this work, we have developed a layered software architecture that facilitates a modular approach to problems in addition to building an abstraction of a robotic system (?). This abstraction allows external software to interact with either a simulated robot and environment or a real robot complete with sensors. In appreciation of the necessity of simulation in addition to real robot control, we have developed a graphical environment for the development of algorithms and software for mobile robotics. Our system permits multiple robots to be controlled from a number of standpoints, and (most usefully) from planning processes. The simulation environment can then be replaced by a pipeline to the robot, or can be used as a high-level control mechanism.

The modular aspect of the software architecture allows us to swap different solutions to various aspects of the problem in and out of the architecture. We are using a weighted combination of multiple approaches to the problems of robot localisation and object localisation in the image. In particular, it is this use of multiple techniques that allows us to inspect the success of a number of mobile robot methodologies, and to determine which aspects of the parameter space are most relevant to our systems.

Approach

Object Detection, Localization and Recognition

A solution to the object recognition problem is inevitably tailored to suit the nature of the objects under consideration. Given that the objects encountered in this exercise will be both distinctively colored and relatively small, we are employing a variety of techniques for the steps of object detection, localization and classification. All three problems are solved by integrating a variety of methods derived from laser range data, image color-space analysis, template matching, principle componenents analysis and statistical measures of image features.

The first such method is a color-space analysis of images obtained from an on-board camera. By identifying pixels that match desired hues, and clustering these pixels into blobs, we can locate possible objects in the image. The second process makes use of local maxima in edge density in the image to detect regions

The support of NSERC and the Federal Centres of Excellence program is gratefully acknowledged.

of interest. Thus, the detection stage relies on image properties to generate regions in the image that may be of interest.

The second step of the recognition task is to localize objects in the world and servo the pan-tilt unit with the camera to center the object in the image. The geometry of the camera pose is then used to obtain an estimate of the world coordinates (x, y) of the object. The process of localization is crucial if we want to identify and move the robot close to objects.

The third step in the object recognition task is to classify the objects. We combine the data obtained from three techniques to perform the identification: visible surface estimation using laser range data, template matching to detected blobs, and principal components analysis of the image itself. Laser striping is used to estimate the visible surface of the object. A simple surface classification of curved or flat is sufficient to distinguish between ellipsoids and cuboids. A measure of how informative the data is also yields the relative uncertainty of the classification.

Template matching is performed on blobs detected by the color space analysis. Given that the objects detected will project either to ellipses or parallelograms in the image, we perform a simple minimization that finds the best possible elliptical and quadrilateral template matches to the data. The template with the best match is used to identify the object.

Finally, principle components analysis offers a third method for object classification. (**?**) Using a set of exemplar images obtained prior to actual competition, a database can be built that projects the examplars into a low-dimensional linear (and hence separable) subspace. The desired result is that when images are obtained on the fly, they can be projected into this subspace, and the objects depicted by the nearest neighbours in the database will be of the same shape as the object under consideration.

By integrating data from the concurrent classification processes, and assuming a good estimate of the position of an object, we can then move on to the task of navigating amongst and interacting with the objects in the world.

Navigation

The issue of navigating within the environment presents two contrasting problems. The robot must avoid the obstacles within the environment, but must approach the obstacles close enough to classify, and touch the "interesting" ones. These are essentially the problems of obstacle avoidance and environmental exploration.

We combine our solution to the problems of exploration and obstacle avoidance through the use of a potential field representation of the environment. (**?**) In this representation, the (x, y) position of the obstacles detected in the vision process are assigned negative (repelling) forces. Attractive (positive) forces are applied to areas which are unexplored, or areas which may be deemed "interesting" from previous vantage points. As the robot moves through the environment, it essentially moves from areas of high potential to lower potential. During the exploration phase, we build an internal metric map of the environment which aids in the detection of unexplored areas. If we come to a local minimum in the potential field manifold, we use an A* breadth-first search to find the closest attractive area, and to generate the best obstacle-free path to that region. The region of exploration is bounded, and therefore our metric map must converge on completeness.

During the exploration phase, odometry error, or error in the position of the robot estimated via dead-reckoning, will accumulate with the motion of the robot. We will be using a combination of a number of localization methods to compensate for this accumulation of positional error. One method is vision-based, and uses principle component analysis and interpolation between basis images of the environment (**?**). This method has the advantage of being relatively robust under a wide variety of conditions, but requires some training time for the acquisition of the basis images. Another localization method is the sonar-based line-fitting to an existing metric map (**?**). The combination of multiple localization methods will allow us accurate localization under a wider set of conditions than any of the localization methods alone.

References

Dudek, G., and Jenkin, M. 1993. A multi-layer distributed development environment for mobile robotics. In *Proc. Conf. on Intelligent Autonomous Systems (IAS-3)*, 542–550. Pittsburgh, PA: IOS Press.

Dudek, G., and Zhang, C. 1994. Pose estimation from image data without explicit object models. In *Proceedings of Vision Interface 1994*.

Hwang, Y. K., and Ahuja, H. 1992. A potential field approach to path planning. *IEEE Transactions on Robotics and Automoation* 8(1):23–32.

MacKenzie, P., and Dudek, G. 1994. Precise positioning using model-based maps. In *Proc. IEEE/RSJ ICRA.*. San Diego, CA: IEEE Press.

Nayar, S. K.; Murase, H.; and Nene, S. A. 1994. Learning, positioning, and tracking visual appearance. In *Proc. IEEE/RSJ ICRA.*, 3237–3244. San Diego, CA: IEEE Press.