

Computer Vision Research at the University of Massachusetts

Edward M. Riseman and Allen R. Hanson
with J. Inigo Thomas and Members of the Computer Vision Laboratory

Computer Science Department
University of Massachusetts
Amherst, MA 01003

Abstract

This video first summarizes current research at the University of Massachusetts on mobile vehicle navigation using landmark recognition and a partial 3D world model. We then show how landmarks and world models might be automatically acquired and updated over time.

A fundamental goal in robot navigation is to determine the "pose" of the robot - that is, the position and orientation of the robot with respect to a 3D world model, such as a hallway. In order to determine its pose, the robot identifies modeled 3D landmarks such as doors and baseboards in a 2D image of the hallway. Identifying landmarks involves determining correspondences of extracted image line segments with predicted landmark lines projected into the image. Model matching is achieved by a combinatorial optimization technique (local search) which minimizes the error in the model to data fit. From the model- data feature correspondences thus obtained, the 3D pose of the robot is computed via a non-linear optimization procedure. The best pose requires that lines in the 2D image lie on the planes formed by the corresponding 3D landmark lines and the camera center. Robust statistical methods are employed to detect outliers. Extension of the initial partial model (over time) is achieved by determining the camera pose over a sequence of images while simultaneously tracking new unmodeled features; triangulation is then used to determine the depth of these new features, allowing them to be incorporated into the 3D model.

One of our goals is to automatically construct a 3D model of the environment, without assuming the availability of an initial partial model. One technique identifies objects in the scene that are shallow in depth and which therefore can be represented by a frontal planar surface at some recovered depth. Using this model we automatically generate a path for the robot that avoids the obstacles. This shallow model has also been successfully used as the initial partial model for model extension from

pose. Another 3D inference technique uses motion analysis over a sequence of images. This technique isolates and successfully corrects for a common major source of error—caused by error in the robot's motion—that has been traditionally neglected in other structure-from-motion algorithms. Model extension has also been achieved through a semi-automated method involving invariants from projective geometry. Given four points or lines on a plane, this technique is able to accurately project from the image to the model any number of such additional features on the same plane.

A test site for the UMass component of the DARPA Unmanned Ground Vehicle program, and a partial 3D model of a portion of the UMass campus has been built. Experiments are underway to use these models for landmark-based autonomous navigation, and automatic 3D model acquisition at this test site.

Acknowledgment

This work has been supported in part by the Advanced Research Projects Agency under contract numbers DAAE07-91-C-RO35 (via TACOM), and DACA76-92-C-0041, by the Defense Advanced Research Laboratories (via Harry Diamond Labs), and by the National Science Foundation under grant numbers IRI-9208920 and CDA-8922572.