May 22, 2008

Learning to Predict Affordances Through Visual Object Categorization

Prof. Jim Rehg

School of Interactive Computing
Georgia Tech.

ABSTRACT OF THE TALK

A fundamental requirement of any autonomous robot system is the ability to predict the affordances of its environment. Affordances are latent "action possibilities" jointly determined by the robot and its environment. For example, if a robot can drive over a patch of grass, then we say that the grass possesses the traversability affordance. In Gibson's theory of direct perception, affordances are perceived directly from sensor data. In the case of computer vision sensing, this consists of learning a direct mapping between image features and affordance properties, such as traversability. We have developed a technique for autonomously learning such affordance mappings from the robot's experience. I will present results of on-line traversability learning for several autonomous outdoor navigation tasks.

A major disadvantage of the direct perception approach is its limited scalability, since the mapping from appearance to affordance must be learned separately for each new object instance and each new affordance type. We explore the hypothesis that visual object categorization can significantly improve the scalability of the affordance learning task. For example, if the object category for a new object instance can be predicted, than its affordance properties can be identified immediately. I will describe

a probabilistic graphical model and associated learning algorithm that supports the joint learning of object categories and their associated affordance labels. In comparison to the direct perception approach, this new model makes more effective use of scarce training data and supports the incremental learning of new affordance concepts. I will present experimental results for category-based affordance learning using a mobile robot platform operating in an office environment.

BIOGRAPHY

Jim Rehg is an Associate Professor in the School of Interactive Computing at Georgia Tech. He received his Ph.D. from CMU in 1995 and worked at the Cambridge Research Lab of DEC (and then Compaq) from 1995-2001, where he managed the computer vision research group. His research interests include computer vision, robotics, machine learning, and computer graphics. He is currently serving as a general co-chair of CVPR 2009.