


Probabilistic models of appearance for object recognition and pose estimation in 2D images

We developed methods to represent the appearance of objects, and associated inference methods to identify them in images of cluttered scenes. The goal here is to leverage, to a maximum, the information conveyed by 2D images alone, without resorting to stereo or other 3D sensing techniques. We are also interested in recovering the precise pose (3D orientation) of objects, so as to ultimately use such information in the context of robotic interaction and grasping.



Grasp Densities: Dense, Probabilistic Grasp Affordance Representations

Motivated by autonomous robots that need to acquire object manipulation skills on the fly, we are developing methods for representing affordances by nonparametrically by samples from an underlying affordance distribution. We started by representing graspability by distributions of object-relative gripper pose distributions for successful grasps (Detry et al. 2011).




Probabilistic Structural Object Models for Recognition and 3D Pose Estimation

Motivated by autonomous robots that need to acquire object models and manipulation skills on the fly, we developed learnable object models that represent objects as Markov networks, where nodes represent feature types, and arcs represent spatial relations (Detry et al. 2009). These models can handle deformations, occlusion and clutter. Object detection, recognition and pose estimation are solved using classical methods of probabilistic inference.



Reinforcement Learning of Visual Classes - Using learning approaches on visual input is a challenge because of the high dimensionality of the raw pixel data. Introducing concepts from appearance-based computer vision to reinforcement learning, our RLVC algorithm (Jodogne & Piater 2007) initially treats the visual input space as a single, perceptually aliased state, which is then iteratively split on local visual features, forming a decision tree. In this way, perceptual learning and policy learning are interleaved, and the system learns to focus its attention on relevant visual features.



Reinforcement Learning of Perception-Action Categories - Our RLJC algorithm (Jodogne & Piater 2006), extends RLVC to the combined perception-action space. This constitutes a promising new approach to the age-old problem of applying reinforcement learning to high-dimensional and/or continuous action spaces.





Image Classification with Extremely Randomized Trees

- We developed a generic method that achieves highly competitive results on several, very different data sets (Marée et.al. 2005). It is based on three straightforward insights: *randomness* to keep classifier bias down, *local patches* to increase robustness to partial occlusions and global phenomena such as viewpoint changes, and *normalization* to achieve invariance to various transformations. The key contribution was probably the demonstration of how far randomization can take us: Local patches are extracted at random, rotational invariance is obtained by randomly rotating the training patches, and classification is done using Extremely Randomized Trees.



Real-Time Object Tracking in Complex Scenes - While most work at the time was based on background subtraction, we developed new methods for complex scenes by tracking local features for robustness to occlusions and to background changes, taking spatial coherence into account for robustness to overlapping, similar-looking targets. In the context of soccer, we also developed methods for robust, model-based and model-free, absolute and incremental terrain tracking.



From: <https://iis.uibk.ac.at/> - IIS

Permanent link: <https://iis.uibk.ac.at/research?rev=1406813078>

Last update: 2018/09/03 14:57

