

Visual Learning

A core interest lies in visual perception as part of closed-loop interactive tasks, and in particular, on systems that improve their performance with experience. Examples of our work include reinforcement learning within perception-action loops, image classification that drives machine learning to the extreme, and [visuomotor learning](#) for various purposes including object detection, recognition and manipulation.

Reinforcement Learning on Visual Perception

Using learning approaches on visual input is a challenge because of the high dimensionality of the raw pixel data. In this work, we bring concepts from appearance-based computer vision to reinforcement learning. Our RLVC algorithm (Reinforcement Learning of Visual Classes, [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.

Our RLJC algorithm (Reinforcement Learning of Joint Classes, [Jodogne & Piater, 2006](#)), extends this idea 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 using Extra-Trees and Random Patches

Image classification remains a difficult problem in general, and the best results on specific problems are usually obtained using specifically tailored methods.

We came up with a generic method that turns this principle upside-down and nevertheless 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,
- **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.

From:

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

Permanent link:

<https://iis.uibk.ac.at/research/visual-learning?rev=1329341960>

Last update: **2018/09/03 14:57**

