

Computer Vision

Instructor: David Forsyth

This will be an introductory tutorial in computer vision. The two core problems of computer vision are reconstruction, where an agent builds a model of the world from an image or a set of images and recognition, where an agent draws distinctions among the objects it encounters based on visual and other information. Both problems should be interpreted very broadly.

Building a geometric model from images is obviously reconstruction (and solutions are very valuable), but sometimes we need to build a map of the different textures on a surface, and this is reconstruction, too. Attaching names to objects that appear in an image is clearly recognition. Sometimes we need to answer questions like: is it asleep? does it eat meat? which end has teeth?, and answering these questions is recognition, too.

We will discuss major ideas in reconstruction and recognition, focusing on recent highly successful local feature representations. The main topics will be: camera models; smoothing, filtering and edge detection; constructing local point representations (SIFT features); searching for near-duplicate images using information retrieval methods; 3D reconstruction using matching methods; object detection and recognition using classifiers; and computer vision methods applied to finding people and reasoning about what they're doing.