

Are Categories Necessary for Recognition?

Alexei Efros
Carnegie Mellon University

The use of categories to represent concepts (e.g. visual objects) is so prevalent in computer vision and machine learning that most researchers don't give it a second thought. Faced with a new task, one simply carves up the solution space into classes (e.g. cars, people, buildings), assigns class labels to training examples and applies one of the many popular classifiers to arrive at a solution.

In this talk, I will discuss a different way of thinking about object recognition -- not as object naming, but rather as object association. Instead than asking "What is it?", a better question might be "What is it like?"[M. Bar]. The etymology of the very word "re-cognize" (to know again) supports the view that association plays a key role in recognition. Under this model, when faced with a novel object, the task is to associate it with the most similar objects in one's memory which can then be used directly for knowledge transfer, bypassing the categorization step all-together. I will present some very preliminary results on our new model, termed "The Visual Memex", which aims to use object associations (in terms of visual similarity and spatial context) to reason about and parse visual scenes. We show that our model offers better performance at certain tasks than standard category-driven approaches.