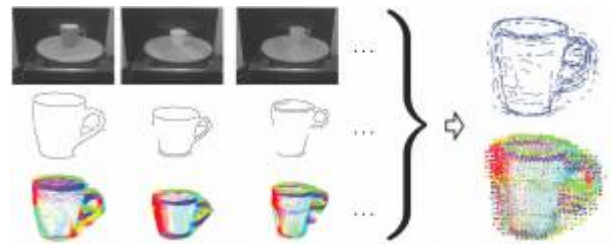
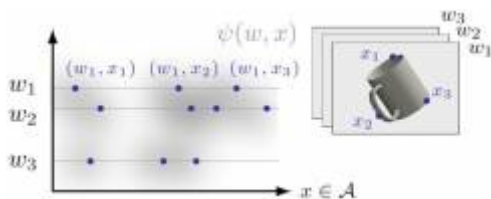


Probabilistic Models of Appearance for Object Recognition and Pose Estimation


We developed probabilistic models to encode the appearance of objects, and inference methods to perform **detection (localization) and pose estimation** of those object in **2D images** of cluttered scenes. Some of our early work used 3D, CAD-style models (Teney et al. 2011), but we then solely focused on **appearance-based models** (Teney et al. 2012).




Those are trained using 2D example images alone, the goal being here to leverage, to a maximum, the information conveyed by 2D images, without resorting to stereo or other 3D sensing techniques. Our models are identically applicable to either specific object instances, or to object categories/classes (Teney et al. 2013). The appearance is modeled as a **distribution of low-level, fine-grained image features**. The strength of the approach is its straightforward formulation, applicable to virtually any type of image feature. We have applied the method to different types of such low-level features: points along image edges, and intensity gradients extracted densely over the image.



Such models of



appearance have been applied to the tasks of object detection/localization, object recognition, and pose classification (by matching the test view with one of several trained viewpoints of the object). A notable advantage of the proposed model is its **ability to use dense gradients directly** (extracted over entire images), versus relying on typical hand-crafted image descriptors. Using gradients extracted at a coarse scale over the images allows us to **use shading and homogeneous regions** to recognize untextured objects, when edges alone would be ambiguous.



We also proposed extensions of



this generative model to perform continuous pose estimation, by explicitly interpolating appearance between trained viewpoints. This makes it one of the rare methods capable of doing **appearance-based continuous pose estimation at category level**, this capability being usually reserved to methods based on 3D CAD models of objects, and limited to specific object instances.

Figure 1 consists of three panels. The left panel shows a physical robotic arm with orange joints and a yellow gripper, positioned over a workbench with various tools. The top right panel shows a 3D model of the robotic arm, rendered in orange and yellow, against a dark background. The bottom right panel shows the skeleton of the robotic arm, represented by white lines and joints on a dark background.

