

Object identification (OID) is a specific type of recognition problem in which the category is known (e.g. cars) and the algorithm recognizes an object's exact identity (e.g. Bob's BMW). Two special challenges characterize object identification. First, inter-class variation is often small (many cars look alike) and may be dwarfed by illumination or pose changes. Second, there may be many classes but few or just one positive "training" examples per class. Because of the low inter-class variation, a solution must locate possibly subtle object-specific salient features (like a door handle) while avoiding distracting ones (such as specular highlights). However, the lack of abundant training data rules out direct techniques of feature selection. We describe an on-line algorithm that takes one query image from a known category and builds an efficient "same" vs. "different" classification cascade by predicting the most discriminative feature set for that object. Our method not only estimates the saliency and scoring function for each candidate feature, but also models the dependency between features, building an ordered feature sequence unique to a specific query image, maximizing cumulative information content.