A.8. Details on local classifiers from Experiment 4

The raw LabelMe annotations for our data included several typos, misspellings, synonyms and subordinant-level category terms (such as “red sand” or “man bending over”). The 16 semantic concepts were: *sky, water, foliage, mountains, ice/snow, rock, sand, animals, hills, fog/mist, clouds, grass, dirt, manmade objects, canyon* and *road.* These include the nine semantic concepts used by Vogel and Schiele (2007) as well as others that were needed to fully explain our database.

In order to ensure that the 16 item list was not too general, we compared this model to a model that represented the basic-level names of regions. The following are the basic-level object and region names decided upon for this data: *bird, branch, bridge, buildings, bush, cactus, canyon, car, cliff, cloud, cow, dirt, dock, fence, flowers, fog, glacier, grass, ground, hill, horse, house, ice, leaves, moon, moss, mountain, mud, path, pebbles, person, plateau, raft, railroad tracks, rainbow, road, rock, root, sand, sheep, sign, sky, snow, telephone pole, tree, trees, trunk, valley, water, weeds.* We found no significant differences between the performance of the 16 region semantic concept model and the 50 item semantic concept model. Furthermore, to ensure that our abstractions from the raw data were not leaving out critical evidence that could be used for classification, we also compared these to the raw annotation data, and again found no significant differences in the overall performance of these models ($t(398) < 1$).

The percent of images correctly classified for the local semantic concept model were desert: 68%, field: 84%, forest: 80%, lake: 16%, mountain: 84%, ocean: 48%, river: 16% and waterfall: 84%. For the prominent object model, these were desert: 60%, field: 84%, forest: 92%, lake: 24%, mountain: 80%, ocean: 12%, river: 20% and waterfall: 44%.