


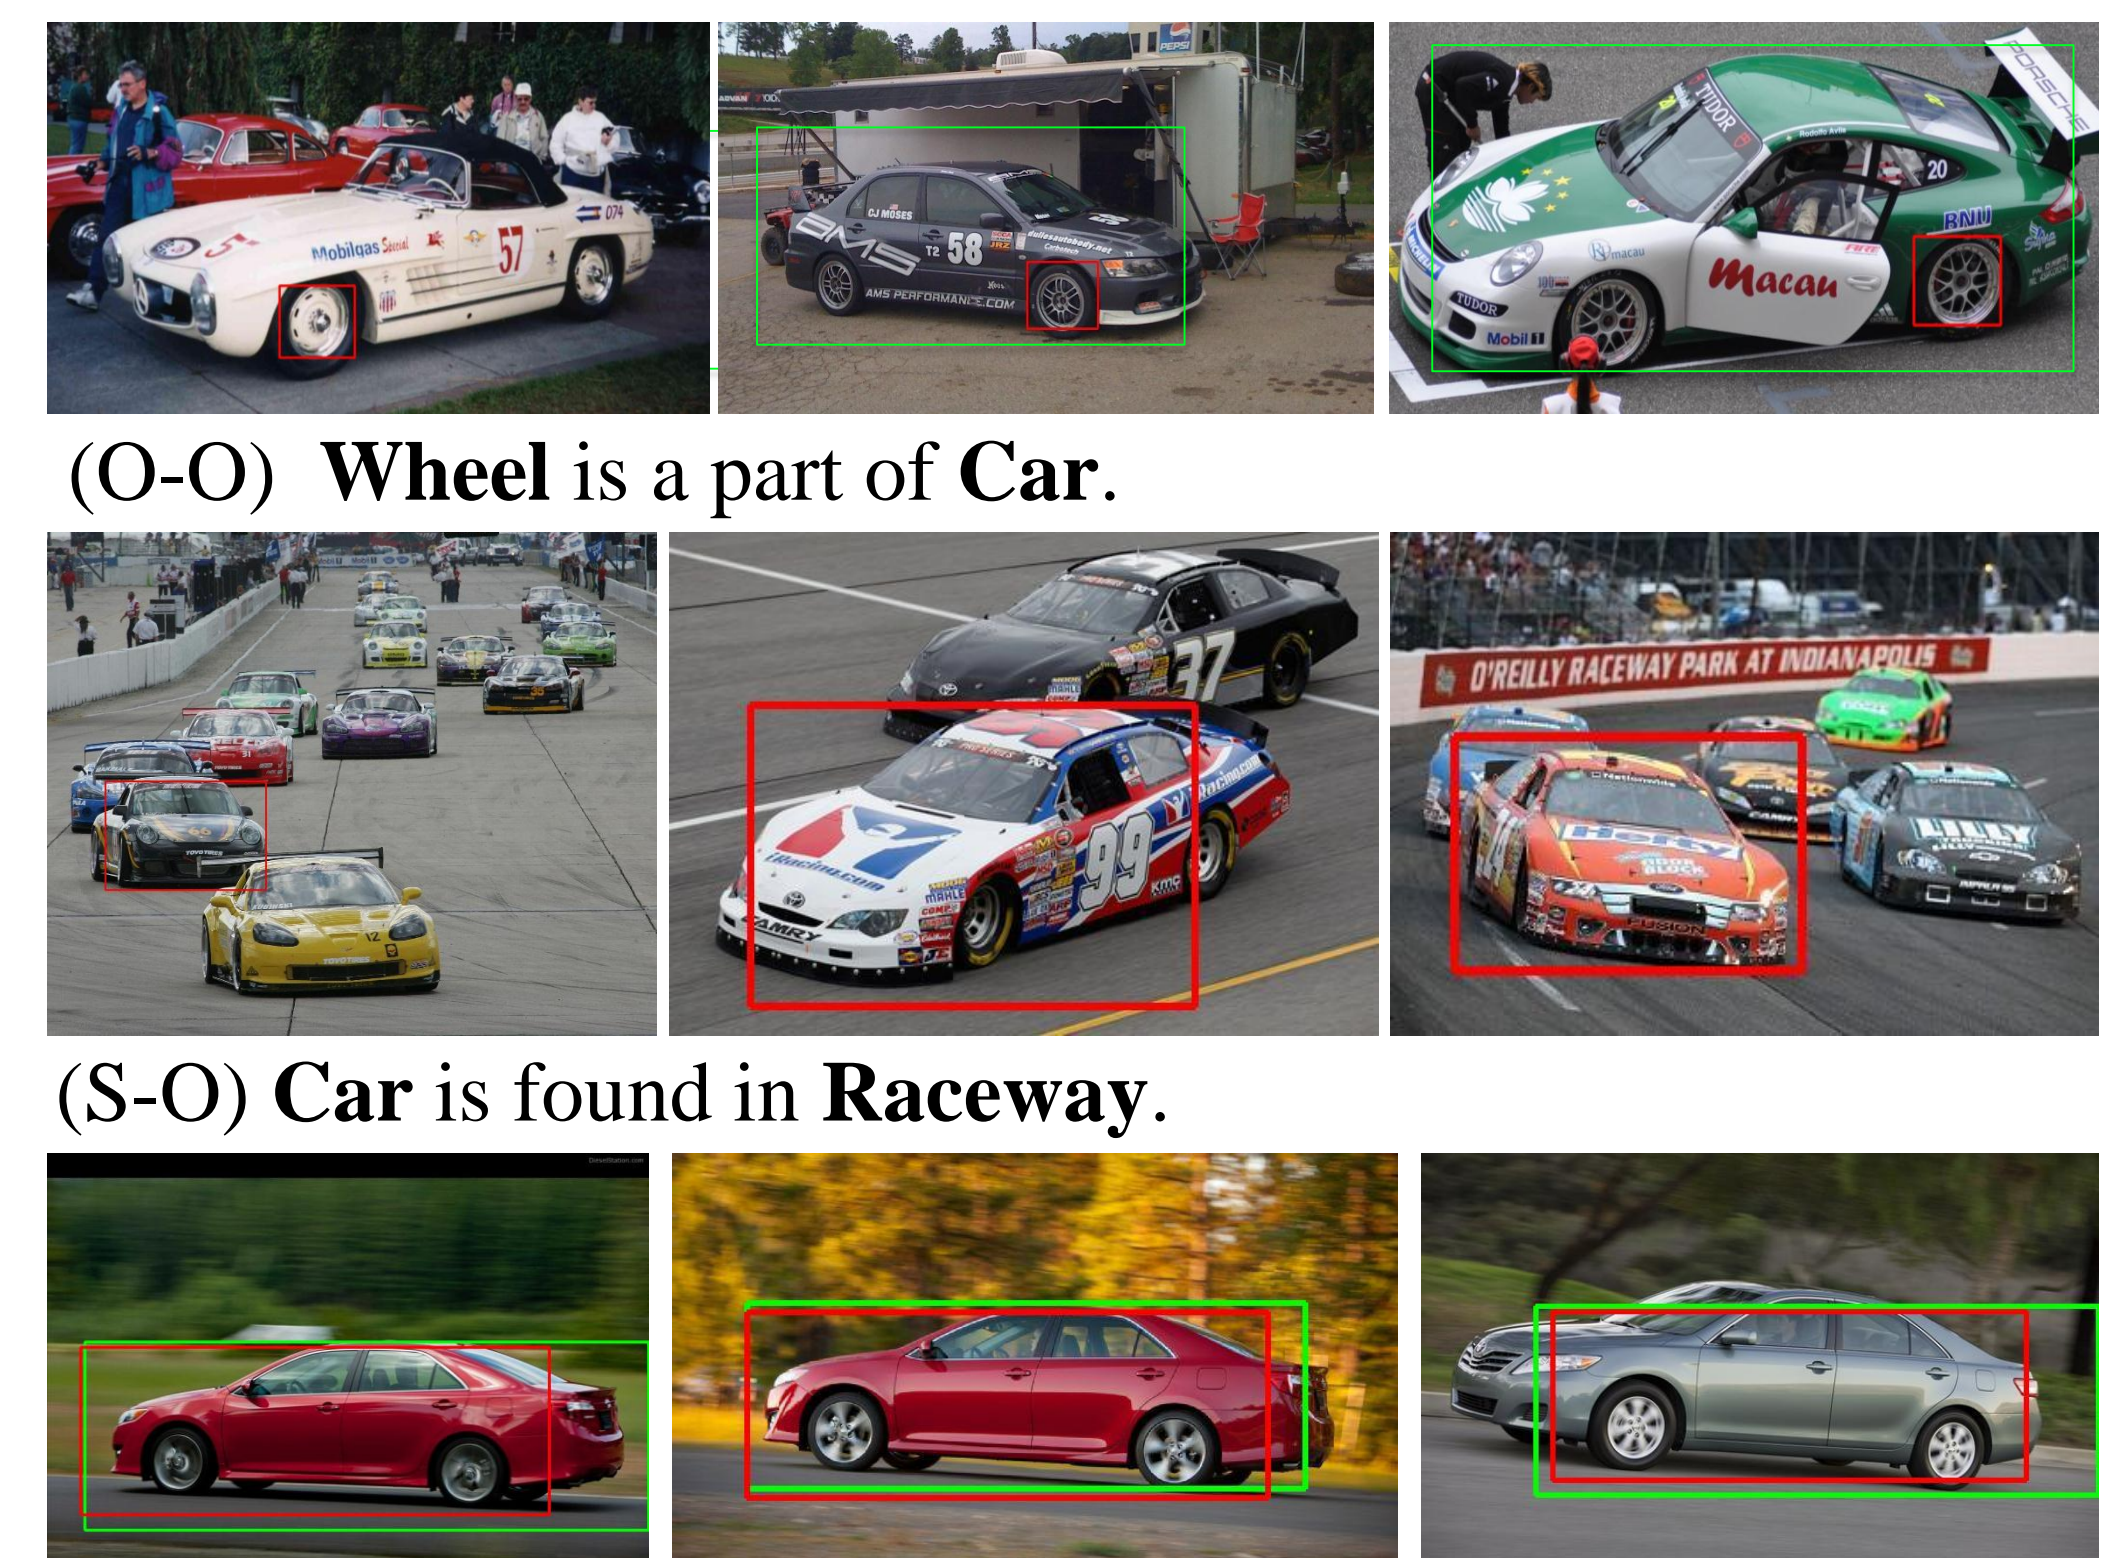

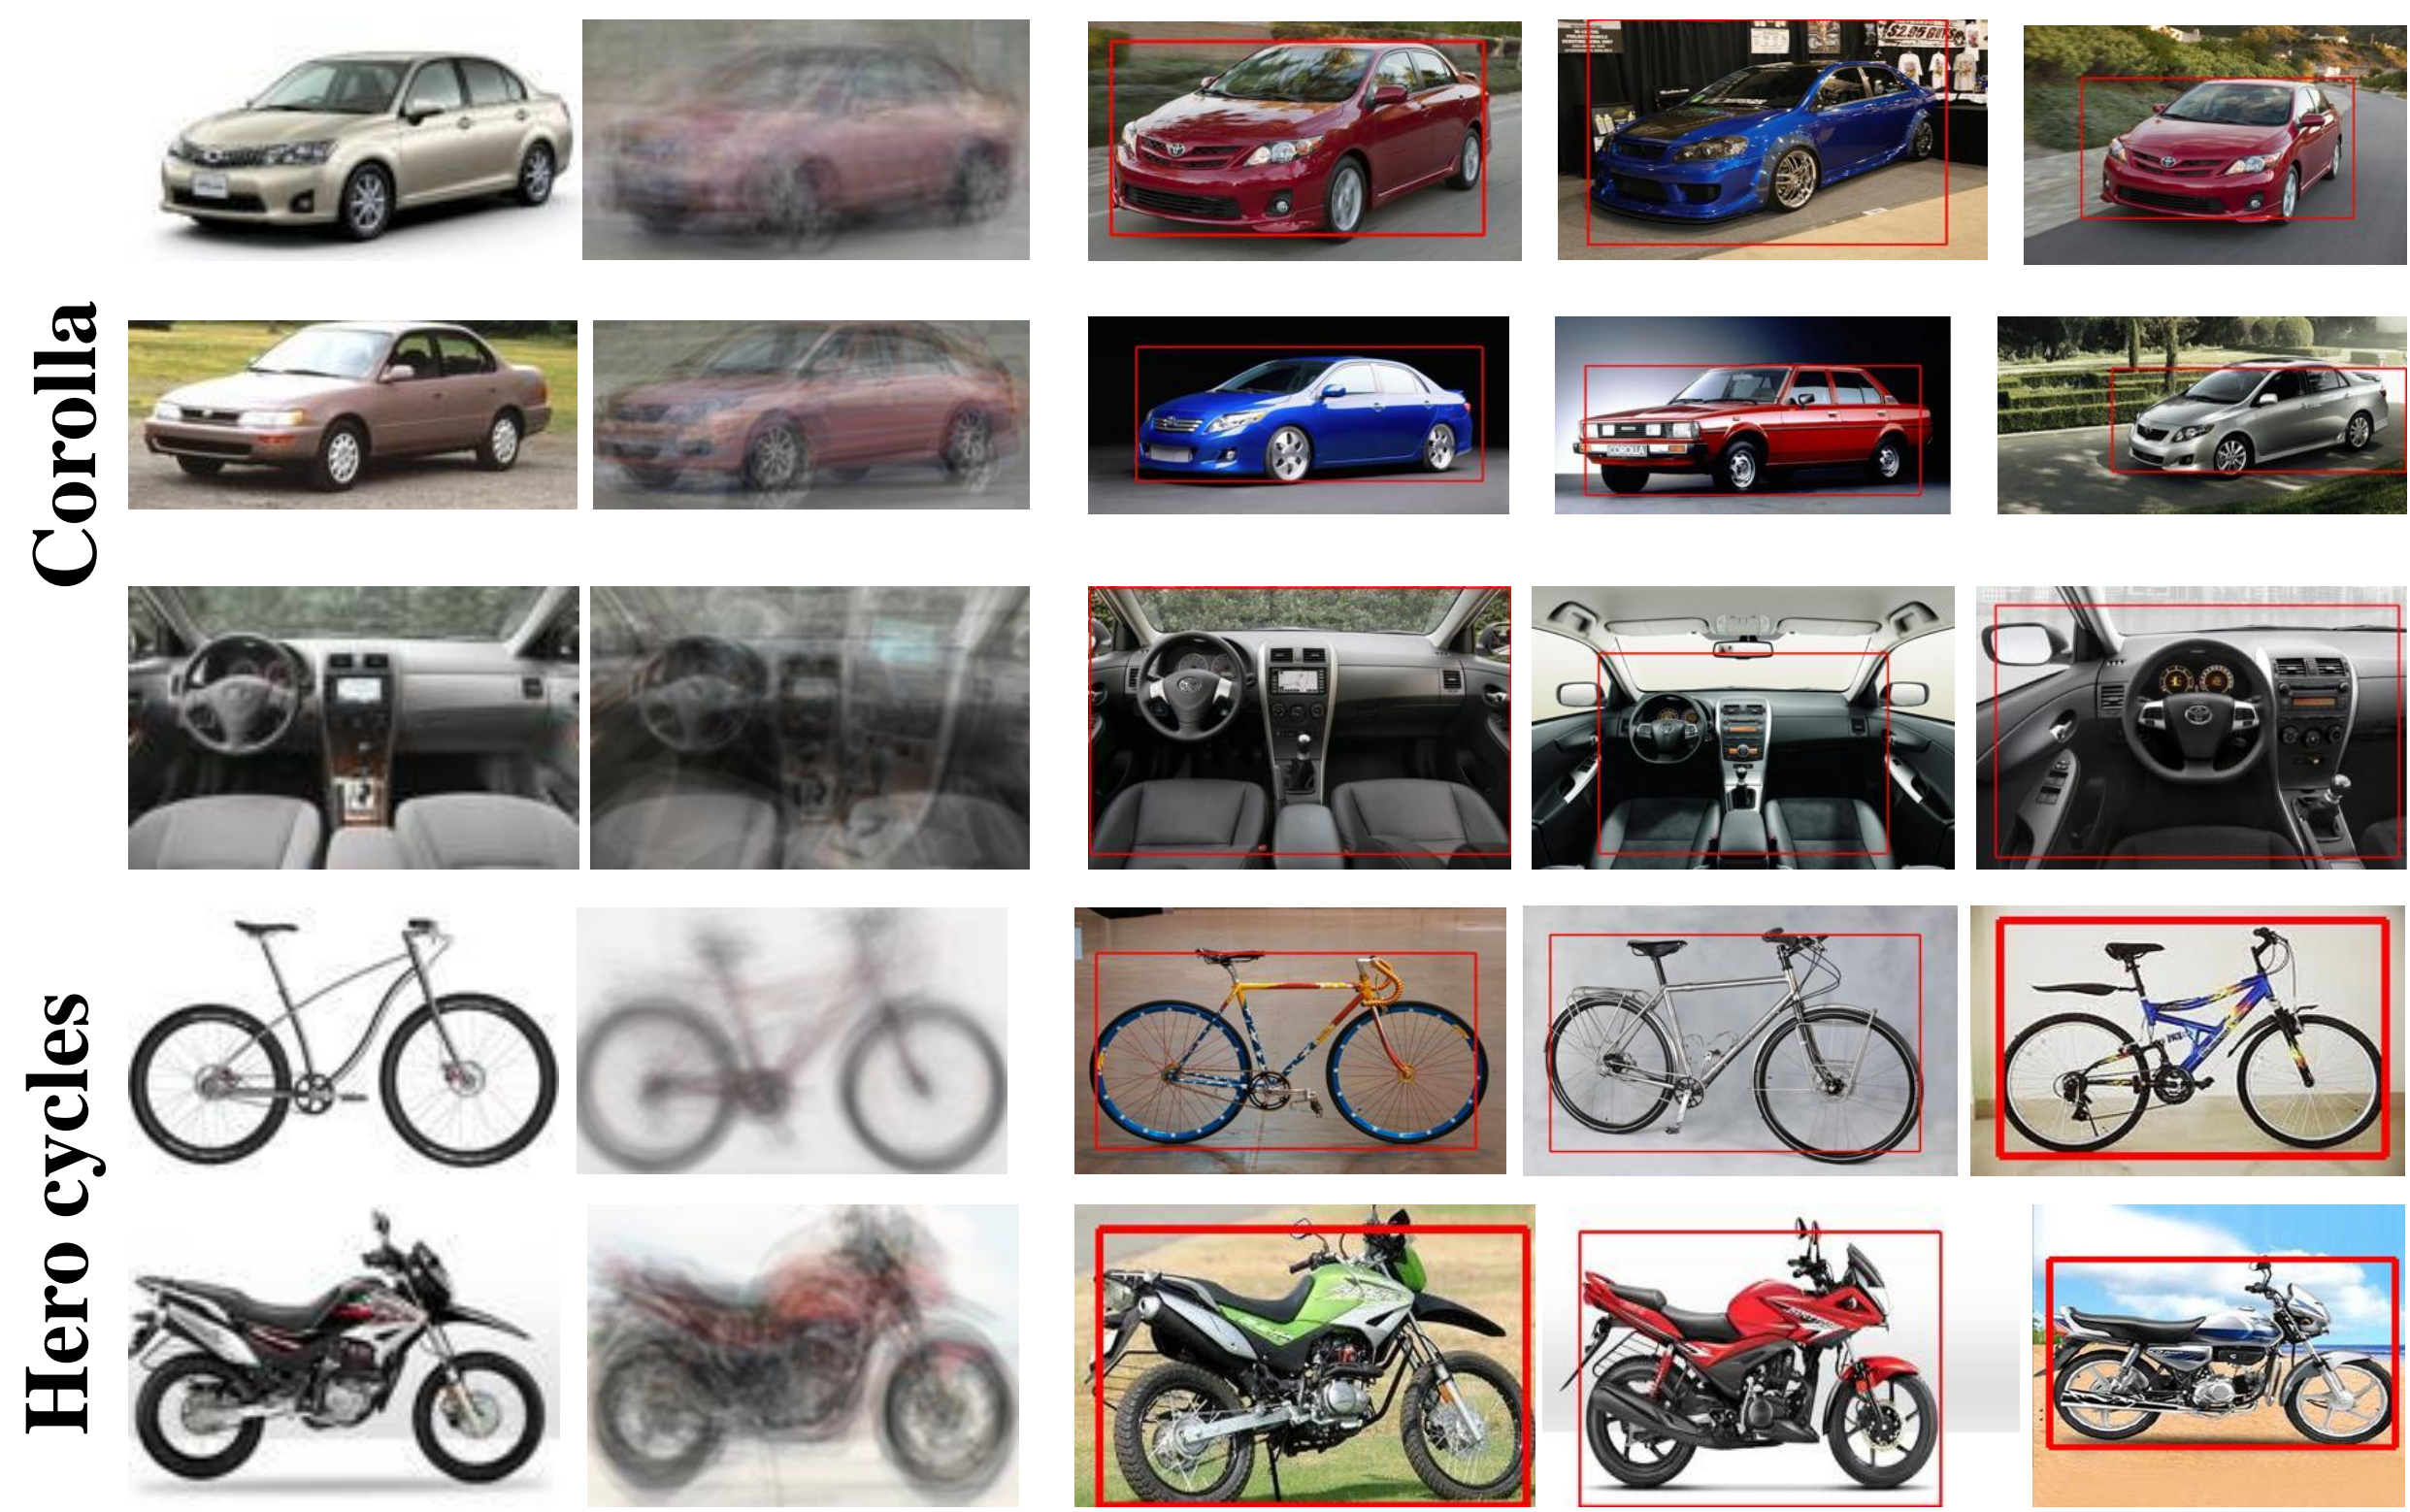
NEIL: Extracting Visual Knowledge from Web Data



Xinlei Chen, Abhinav Shrivastava, Abhinav Gupta
Carnegie Mellon University



All models and relationships can be found at www.neil-kb.com



(a) Objects (w/Bounding Boxes and Visual Subcategories)

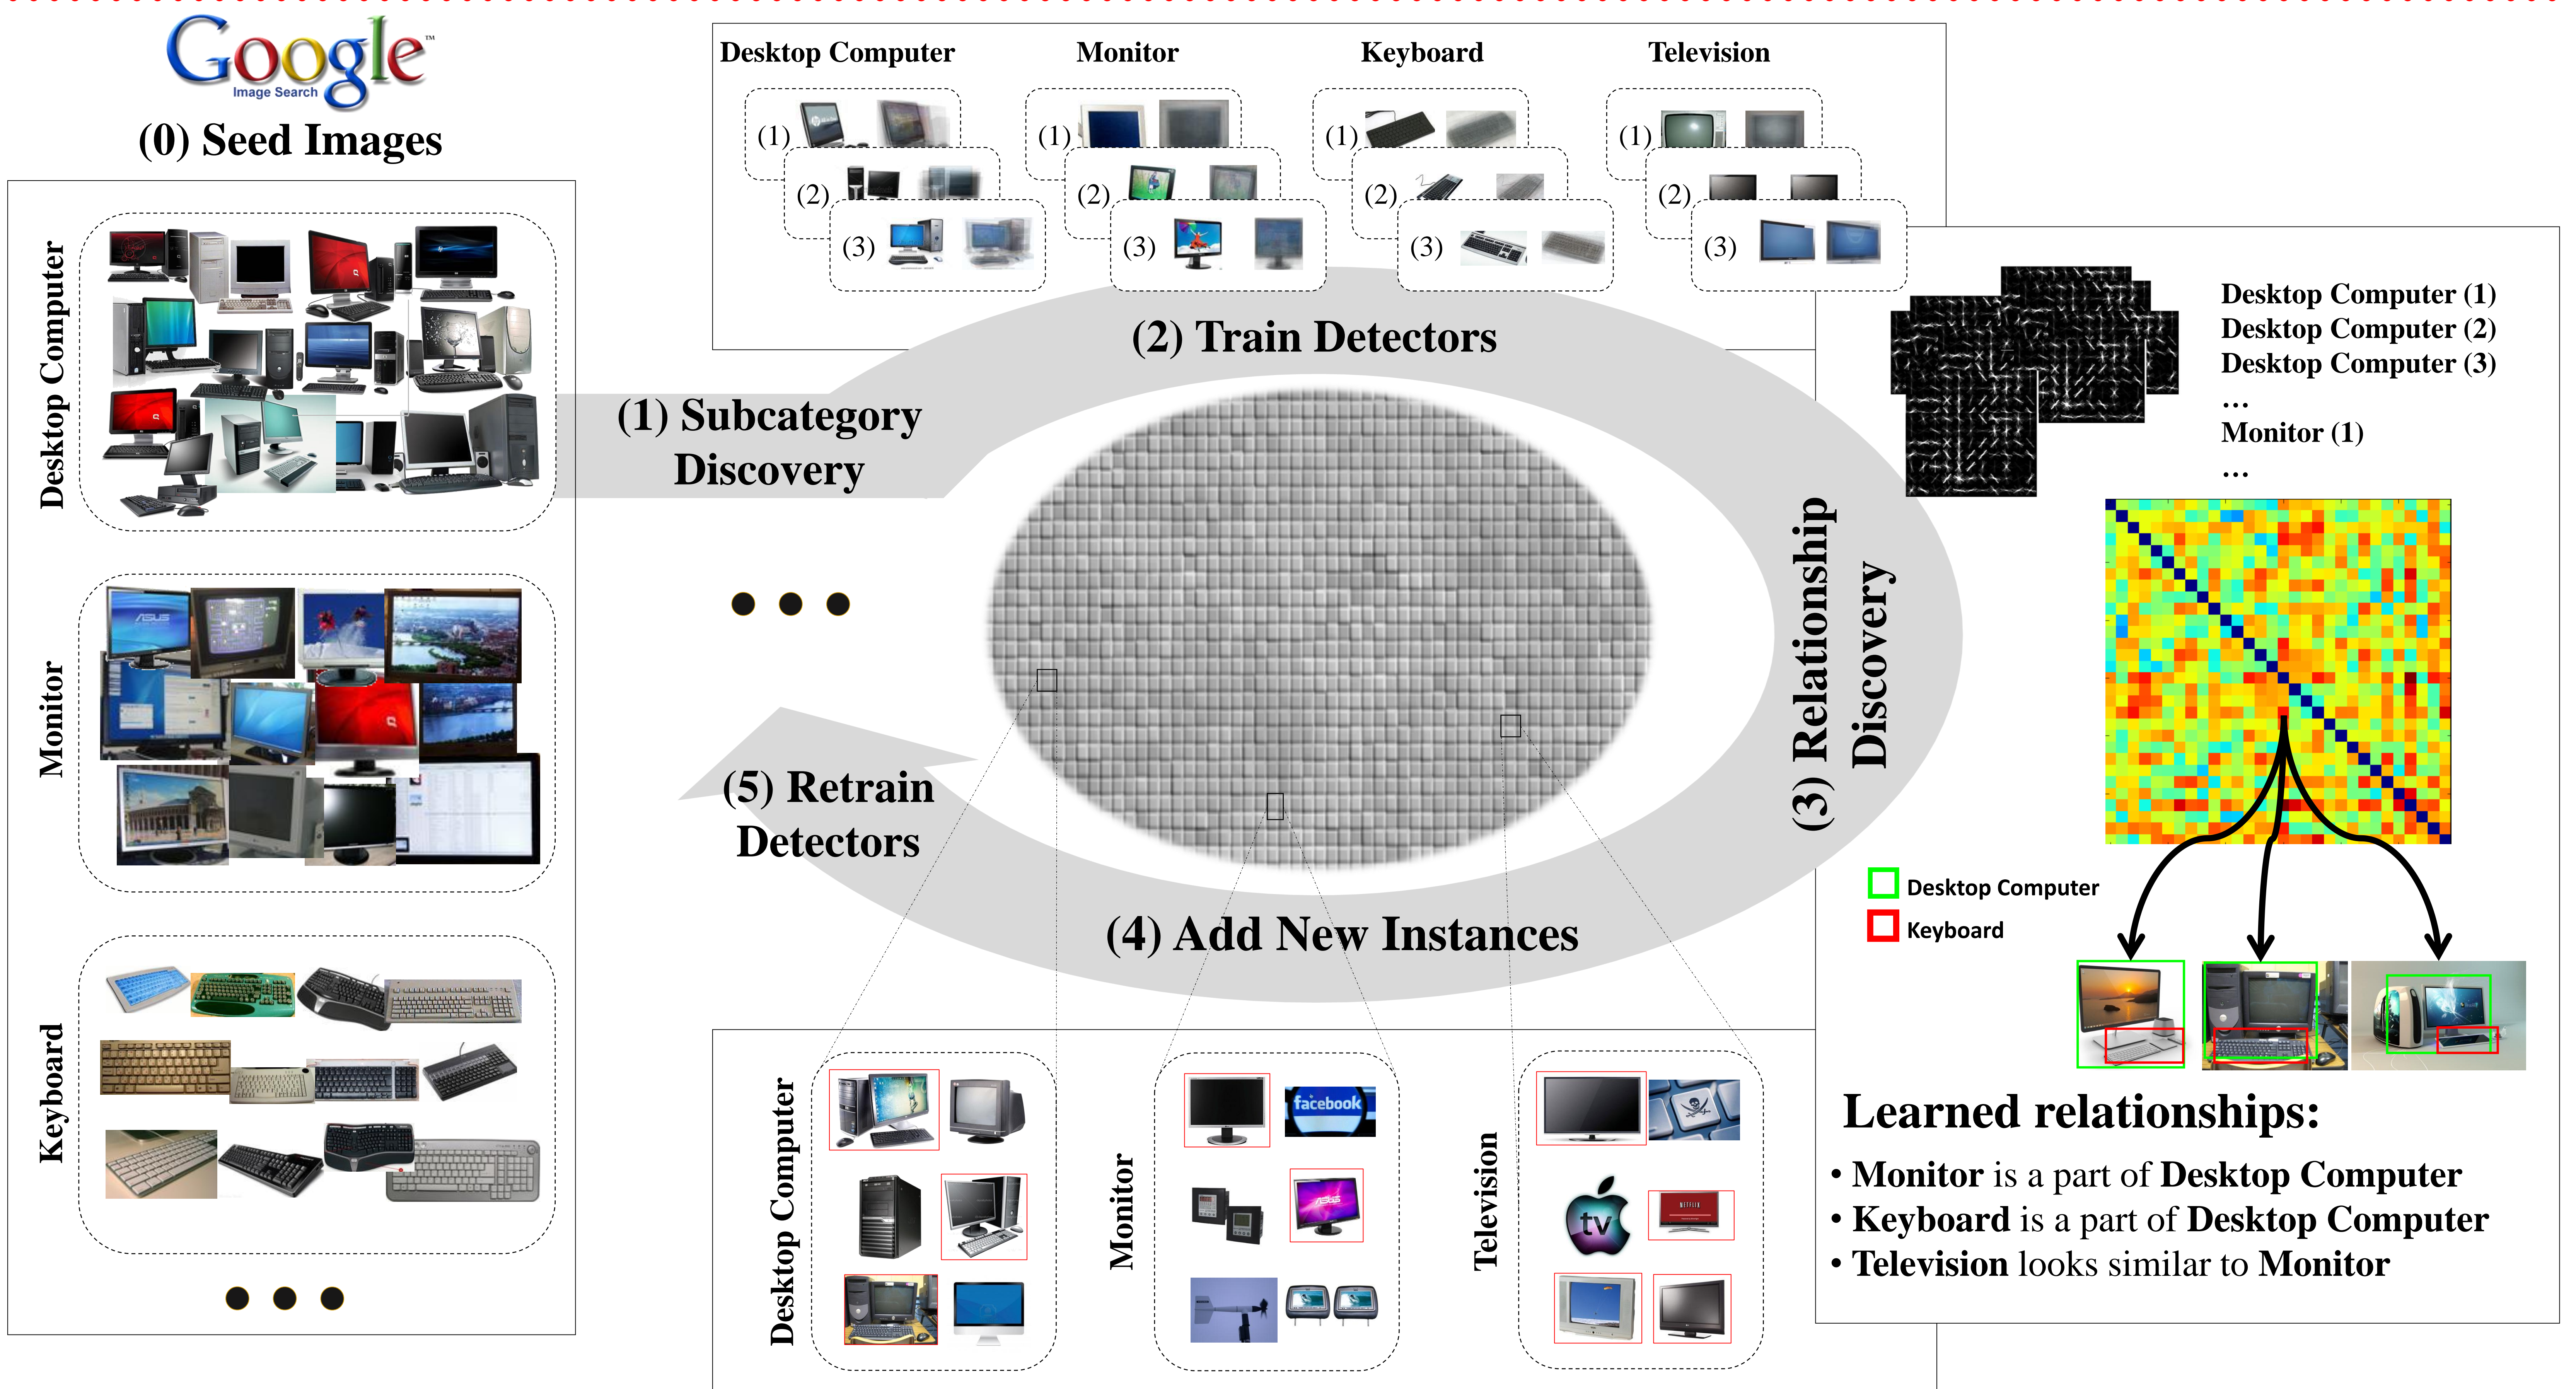


(b) Scenes

(c) Attributes

Visual Instances Labeled by NEIL

Relationships Extracted by NEIL



Why NEIL Works?

1. Micro-Vision → Macro-Vision
2. Structure in Visual World
3. Semantically-Driven Acquisition


Statistics

- Running for ~4 months on ~200 cores
- >2000 concepts, including >1000 objects, >1000 scenes, ~90 attributes
- Analyzed 5m images
 - Labeled 600K images
 - Discovered 3K relationships

Subcategories



Relationships



More Relationships

- Monitor is found in Control room
- Washing machine is found in Utility room
- Siberian tiger is found in Zoo
- Baseball is found in Butters box
- Bullet train is found in Train station platform
- Cougar looks similar to Cat
- Urn looks similar to Goblet
- Samsung galaxy is a kind of Cellphone
- Computer room is/has Modern
- Hallway is/has Narrow
- Building facade is/has Check texture
- Trading floor is/has Crowded
- Umbrella looks similar to Ferris wheel
- Bonfire is found in Volcano

Scene Classification

Classifier	mAP
Seed Classifier (15 Google Images)	0.52
Bootstrapping (without relationships)	0.54
NEIL Scene Classifiers	0.57
NEIL (Classifiers + Relationships)	0.62

Object Detection

Classifier	mAP
Latent SVM (50 Google Images)	0.34
Latent SVM (450 Google Images)	0.28
Latent SVM (450, Aspect Ratio Clustering)	0.30
Latent SVM (450, HOG-based Clustering)	0.33
Seed Detector (NEIL Clustering)	0.44
Bootstrapping (without relationships)	0.45
NEIL Detector	0.49
NEIL Detector + Relationships	0.51