




Collecting a Large-Scale Dataset of Fine-Grained Cars

Jonathan Krause¹, Jia Deng¹, Michael Stark^{1,2}, and Li Fei-Fei¹

¹Stanford University, ²Max Planck Institute for Informatics



Why Cars?

One of the most important object categories

Difficult:

Quiz: How Many Classes Below?



New avenues of research:

- Simplifies pose normalization
- Easier 3D reasoning

Finding Classes

1. Scrape initial class list and example images from popular car website
2. Identify similar example images via perceptual hashing and merge categories
3. Merge categories with similar example images
4. Subsample $N=197$ categories




Image Collection

Collect candidate images for each class from Google, flickr, and bing.

Filter carefully on Amazon Mechanical Turk



Get bounding boxes, deduplicate images.



Modeling Annotator Quality


Estimate worker quality and probability of class membership via *Get-Another-Label* system of Ipeirotis *et al.*[1]

EM procedure on worker confusion matrices and label probabilities

[1] P.G. Ipeirotis, F. Provost, and J. Wang. Quality Management on Amazon Mechanical Turk. *WS ACM SIGKDD*, 2010

Dataset Statistics

197 classes
16,185 images



Training Annotators

Detailed examples of discriminative parts



2010 BMW 3-Series

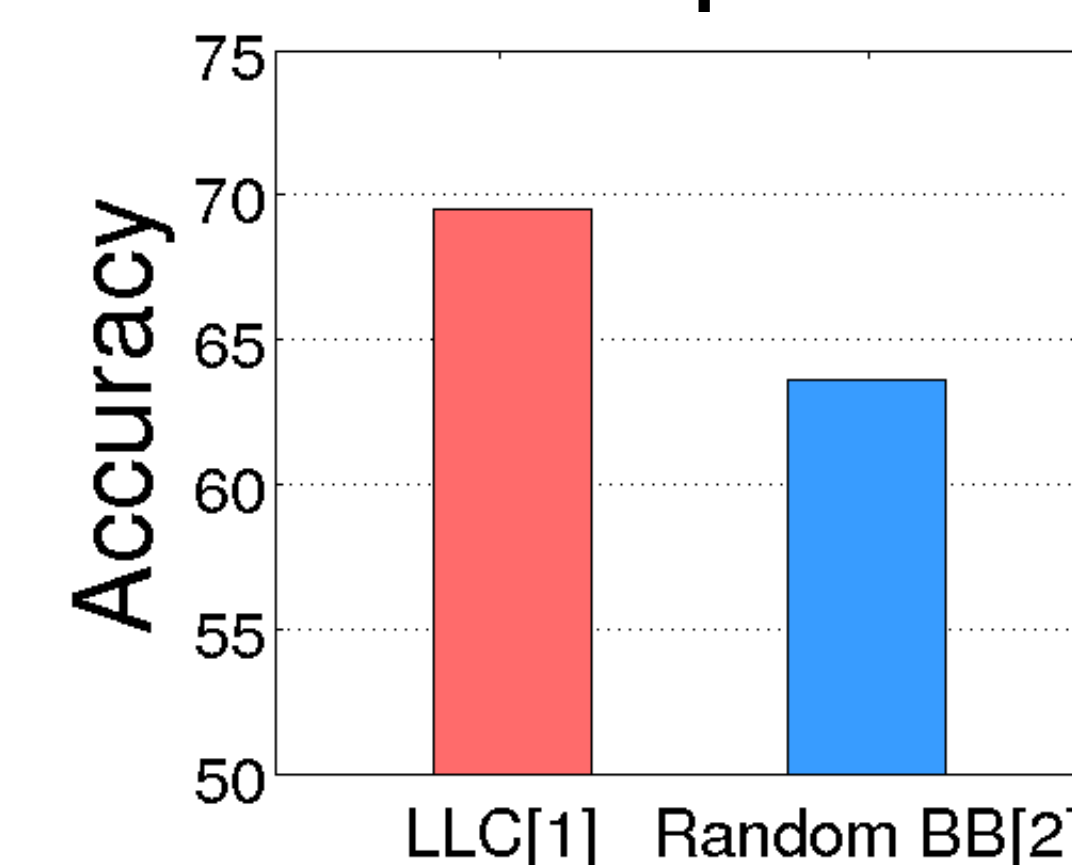
2010 BMW 5-Series

Provide images of both positive and negative example classes in annotation task



Click on the photos that contains a 2012 Tesla Model S Sedan. It should look identical to the GOOD examples above EXCEPT background, viewing angle and color. There can be multiple body styles for this model and you must identify each of the styles given in the examples. AVOID CARS THAT LOOK IDENTICAL TO THE BAD EXAMPLES. PREVIEW MODE. TO WORK ON THIS HIT, PLEASE ACCEPT IT FIRST.

Baseline Experiment



50/50 training/test split

[1] J. Wang, J. Yang, K. Yu, F. Lv, T. Huang, and Y. Gong. Locality-Constrained Linear Coding for Image Classification. *CVPR* 2010

[2] J. Deng, J. Krause, and L. Fei-Fei. Fine-Grained Crowdsourcing for Fine-Grained Recognition. *CVPR* 2013

Future Directions

Good performance:

- Go even larger-scale!
- Ready for application?

Car attributes:

- Color, size, year, attributes
- Large-scale viewpoint/pose prediction
- 3D

