



embedded  
**VISION**  
SUMMIT

# Tools for Democratizing Computer Vision

## Automated Performance Characterization



Jayan Eledath

Oct 2 2013

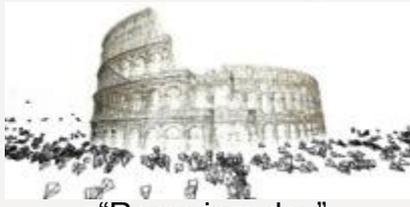
# Computer Vision is becoming ubiquitous

Growing need for support services

- Computer vision is maturing
  - Many CV applications are starting to become widely deployed



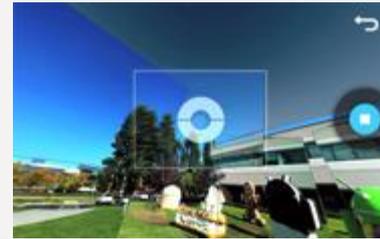
Kinect



“Rome in a day”



SnapTell



Google Pano & Photosphere



Visual odometry for robots

- Libraries like OpenCV spurring rapid growth in algorithm and application developer communities
- Development landscape is vast and cluttered
  - Many algorithms, limited capabilities
- What's Needed
  - An efficient means for assessing algorithm performance
  - Automatic matching of the “best” algorithms to applications

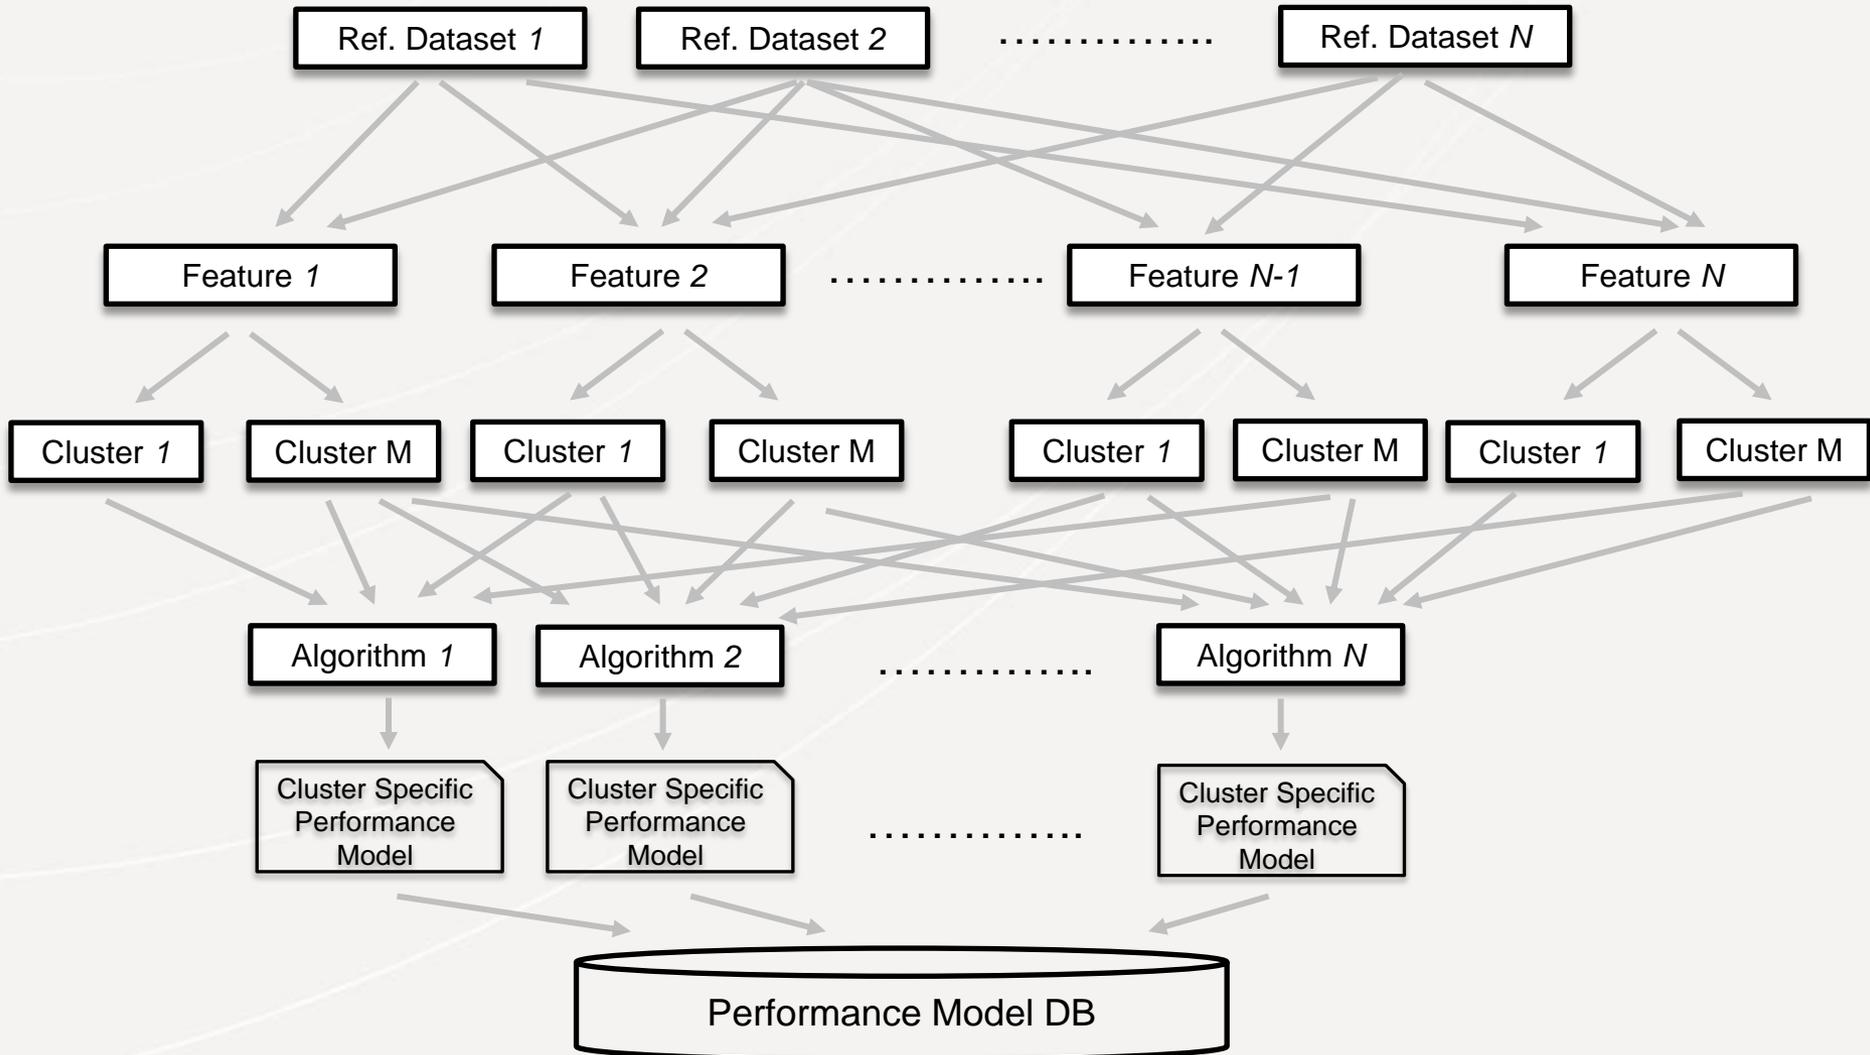
- Principled solution to the problem: “Every CV researcher is eager to create their own broken algorithm rather than improving on others’ broken algorithms”
- Automated performance characterization enables rigorous testing and a level playing field on which to compare various algorithms
- Faster development of working applications
  - Automatically select from multiple algorithms for a (sub)problem
  - Automatically identify the best parameters for an algorithm and a particular image
  - Minimize barrier to getting started with state-of-the-art algorithms, their operating parameters and their use under various operating conditions

- Pd: = Probability of True Detections
  - Number of True Positives / Number of Ground Truth Detections
- Pfa: = Probability of False Alarms
  - (Num. of Detections—Num. of True Positives) / Num. of Detections
- Algorithms parameter space: Input parameters and their ranges
  - Example algorithm: The Zhu and Ramanan (CVPR 2012) Face Pose and Landmark detection
  - Input parameters: Classifier Threshold, Non maximal suppression Threshold, Classifier Model
- Probability distributions
  - Probability (detector\_score | positive)
  - Probability (detector\_score | negative)
  - Probability (detector\_score | +/-, operating condition)

- Performance Metrics
  - Prob. of Detections vs. Prob. False Alarms plots (with distributions)
- Algorithm Selection:
  - Determine what algorithm to use on a given instance (image) or instances (collection of images) to maximize performance
- Predicting Performance:
  - Predict how an algorithm will behave (in terms of some metric) on a new instance (image) or datasets (collection of images) without executing the algorithm

# Large-scale APC Framework

## Conceptual Design



## Live APC Tool Demo

---



# Performance Characterization—Datasets

Reference Dataset	Object Class	Number of Ground-Truthed Images	Source
LFW	Faces	2845	Public Domain
Pascal VOC 2011	Vehicle, Person, Bike	14961	Public Domain
Google Image Faces	Faces	4384	Internal

- Examples of algorithms characterized
  - PittPatt (face detection)
  - Face and Pose Detector (Zhu and Ramanan, 2012)
  - Parts-based Object Detector (Felszenswalb; models for vehicles, person, bicycle)

# APC with Image-Derived Attributes

- Algorithm performance changes based on image content
- Attributes can be used to group (cluster) images into semantic categories (e.g., water scene, indoors)
- Goal: Characterize performance based on image content
- Example attributes
  - GIST descriptor: Holistic scene representation in terms of dominant image gradients
  - Color descriptor: Bag of color histograms that are computed from image patches.
  - Scene Attributes\*: 102 attribute classifiers for various semantic classes

agedworn asphalt att bathing biking brick camping  
 carpet cleaning climbing cloth clouds clutteredspace  
 cold competing concrete conductingbusiness  
 congregating constructingbuilding digging  
 directsunsunny dirtsoil dirty diving driving dry eating  
 electricindoorlighting enclosedarea exercise far-  
 awayhorizon farming fencing fire flowers foliage  
 gaming glass glossy grass hiking ice leaves man-made  
 marble matte medicalactivity metal moistdamp  
 mostlyhorizontalcomponents natural naturallight  
 mostlyverticalcomponents ocean paper pavement  
 nohorizon notpartofatree playing ppenarea praying railing reading research  
 rockstone rubberplastic ruggedscene runningwater  
 rusty sand scary semi-enclosedarea shingles shopping  
 shrubbery smoke snow socializing  
 spectatingbeinginanaudience sports sterile stillwater  
 stressful studyinglearning sunbathing swimming  
 symmetrical teachingtraining tiles  
 transportingthingsorpeople trees txt usingtools  
 vacationingtouring vegetation vinylinoleum  
 waitinginlinequeueing warm wavessurf wire wood  
 working

\* <http://cs.brown.edu/~gen/sunattributes.html>

# Sample Cluster Visualization

PASCAL VOC 2011 | GIST Descriptor



# Sample Cluster Visualization

PASCAL VOC 2011 | Color Descriptor



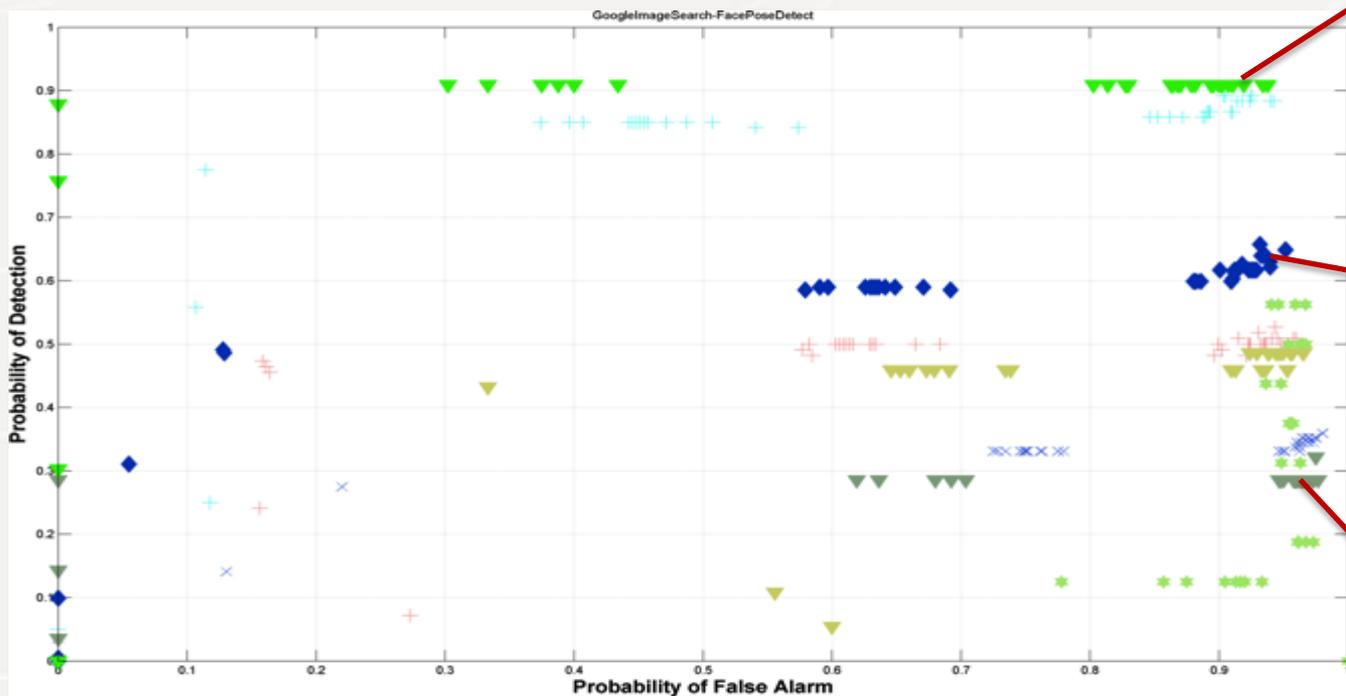
# Sample Cluster Visualization

PASCAL VOC 2011 | Scene Attributes



# Cluster-specific Pd vs. Pfa

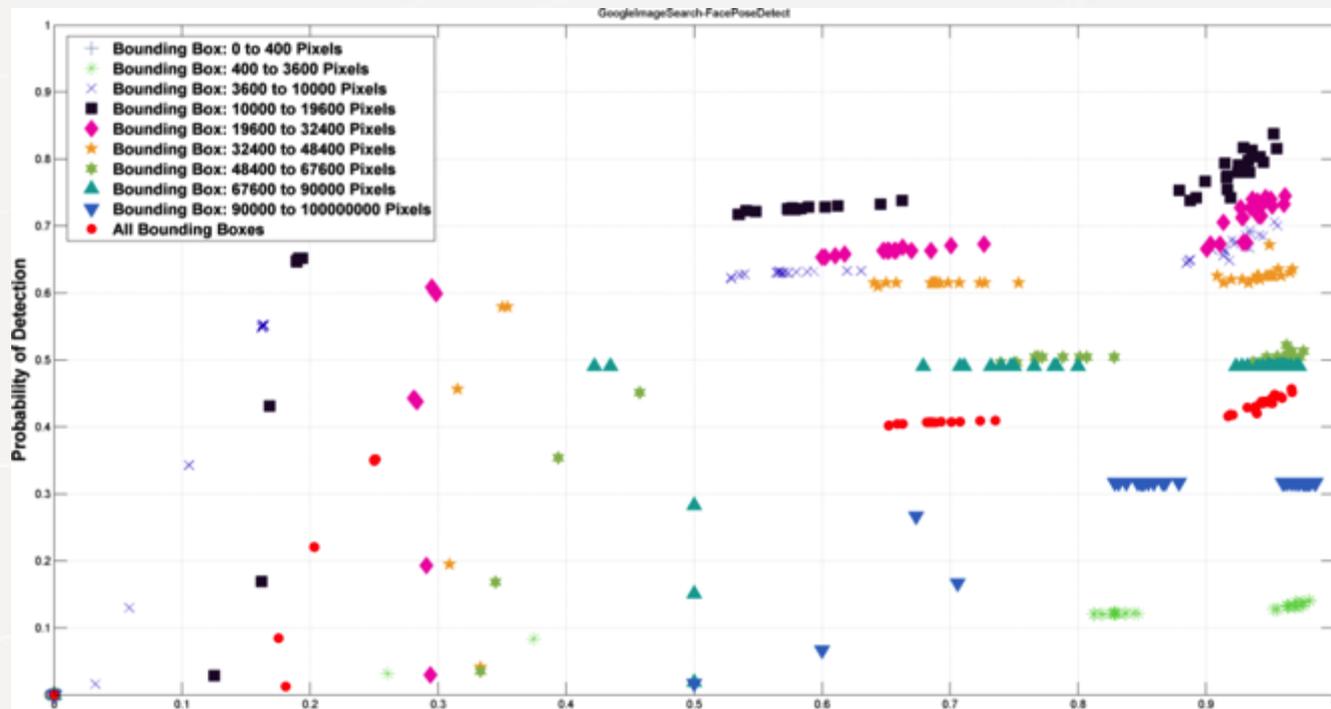
## Scene Attribute-based clusters



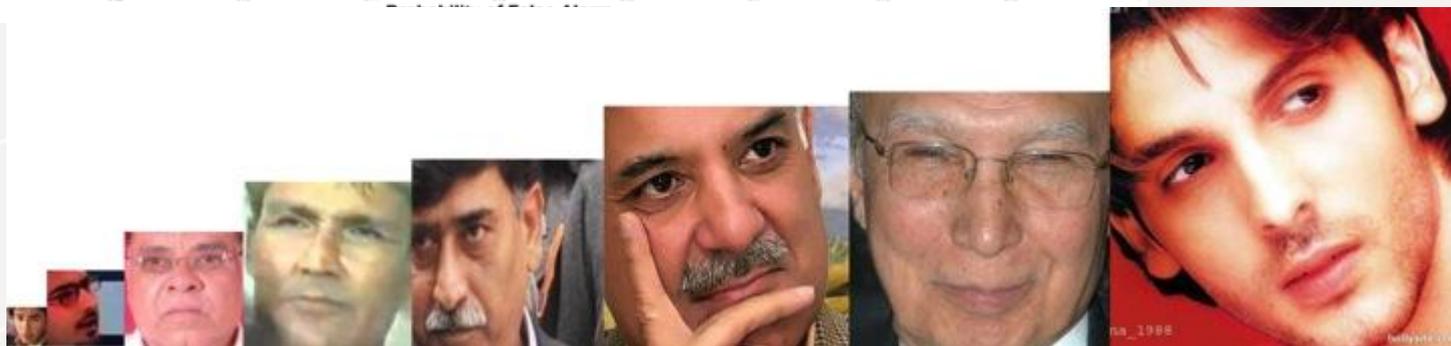
Algorithm: Face & Pose Detect  
Clustering Feature: Scene Attribute  
Number of Clusters: 100  
Dataset: "Google Faces" (internal)

# Cluster-specific Pd vs. Pfa

Face size-based clusters



Algorithm: Face & Pose Detect  
Number of Clusters: 100  
Dataset: "Google Faces" (internal)



- Automated performance characterization enables rapid development of robust CV applications by providing a means to:
  - Efficiently determine how an algorithm performs on large reference datasets
  - Automatically determine which algorithm to use for a particular image (or set of images)
- APC implemented as a web-based software service that allows easy evaluation of algorithm performance against various reference datasets
- Demonstrated how algorithms perform as a function of image content