

What Makes a Great Picture?



© Robert Doisneau, 1955

CS194: Image Manipulation & Computational Photography

*With many slides from Yan Ke,
as annotated by Tamara Berg*

Alexei Efros, UC Berkeley, Fall 2015


Photography 101: the where and when

- Composition
 - Framing
 - Rule of Thirds
 - Leading Lines
 - Textures and Patterns
 - Simplicity
- Lighting
 - Light Direction
 - Color coordination / balance
 - “Golden Hour”
 - (sur) realism

Framing

“Photography is all about framing. We see a subject -- and we put a frame around it. Essentially, that is photography when all is said and done.”


-- from photo.blorge.com



Frame serves several purposes:

1. It gives the image depth
2. Used correctly, framing can draw the eye of the viewer of an interest to a particular part of the scene.
3. Framing can bring a sense of organization or containment to an image.
4. Framing can add context to a shot.

Examples of nice framing



<http://flickr.com/photos/paulosacramento/226545698/>


<http://flickr.com/photos/chrisbeach/13868545/>

<http://flickr.com/photos/74531485@N00/929270814/>

<http://flickr.com/photos/freakdog/223117229/>

<http://flickr.com/photos/cdm/253805482/>

Rules of Thirds




Other examples



Leading Lines



More examples



Textures and Patterns



Color Coordination



Complementary colors (of opposite hue on color wheel)


Front Lighting



Side Lighting



Back Lighting



“Golden Hour”



Simplicity




“Look Into” by Josh Brown @ Flickr



Prof - Obvious what one should be looking at ie easy to separate subject from the background. Snap – unstructured, busy, filled with clutter.

Simplicity



Simplicity




(Sur) Realism



“Golden Gate Bridge at Sunset” by Buzz Andersen @ Flickr

“Golden Gate 3” by Justin Burns @ Flickr

(Sur) Realism




Aki Jinn 2015

“Somewhere Only We Know Prt2 (sic)” by Aki Jinn @ Flickr


Anyone can take great pictures...



I am a sucky photographer...



...but I am a pretty good photo critic!



<http://flickr.com/photos/aaefros/>

of my Paris photos on Flickr: 32

Total # of my Paris photos: ~1250

~2%

The Postmodern Photographer

The Old Days: a pre-process

- Load film
- Find subject
- Position camera
- Set all the settings “just right”
- Take a deep breath...
- ...Press button!

The New Digital Days: a post-process

- Get a 16 GB memory cartridge
- Take pictures like there is no tomorrow!!!
- ...
- Back home, spend hours of agony trying to find 1-2 good ones

How to recognize the good photos?



Applications

Image search for improved quality along with relevance.

Automatically select the best photos from a set of vacation pictures to choose the best ones to show.


See if computer can perform well on a traditionally human task.

Outline

- Photography 101
- Recognition (CVPR '06)
 - What makes one photo better than another?
 - What features can we extract?
 - How can we measure our performance?

[Y. Ke, X. Tang, and F. Jing. *The Design of High-Level Features for Photo Quality Assessment*. CVPR 2006.](#)

Not Critiquing Art



Piet Modrian



Lothar Wolleh

Not considering semantic measures of what makes a photo good (subject matter, humor, etc).
Professional = those you would frame, snapshot = those that would stay in photo album.


Features – Spatial Distribution of Edges



More edges
near border
due to
background
clutter



More edges
near center
of img




Spatial Distribution of Edges

Mean Laplacian of snapshots




M_s

More uniformly distributed




Low quality photos

Mean Laplacian of professional



M_p

More concentrated



High quality photos

Expect high quality photos to have high spatial frequency edges nearer to center than snapshots

Edge width

Calculate area that edges occupy – width of bounding box covering 96% of edge energy

Cluttered regions should tend to produce a larger bounding box, and well defined subjects should produce a smaller one.





.94



.56

Color Distribution

- K-NN on color histogram

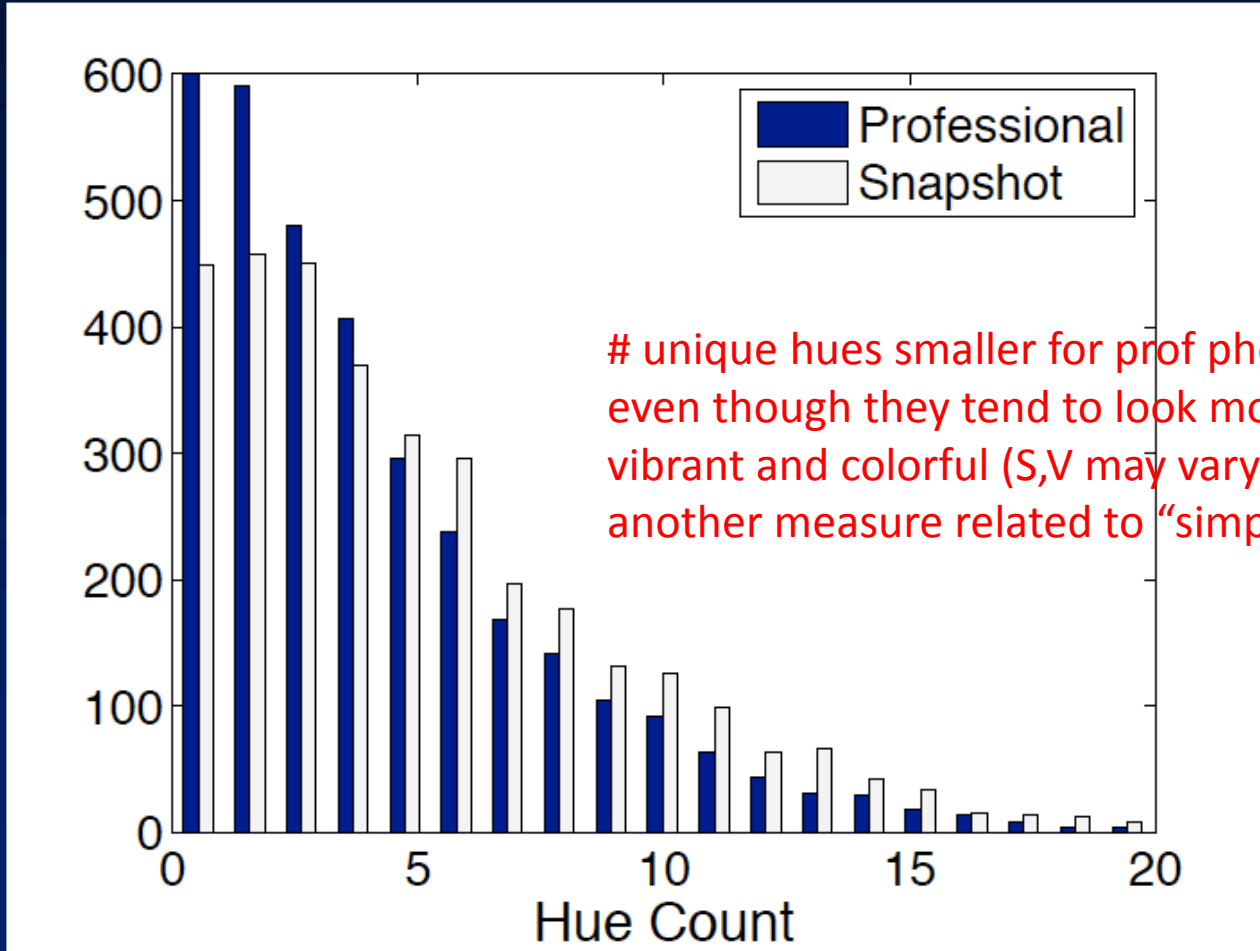


For query image find k nearest neighbors in training set.
Quality = number of prof neighbors in top 5.



$$q_{cd} = \# \textit{professional_neighbors}$$

20 bin histogram defining possible unique hues

Hue Count




unique hues smaller for prof photos even though they tend to look more vibrant and colorful (S,V may vary more) – another measure related to “simplicity”



$$q_h = 20 - (\# \text{ hues } > \text{ threshold})$$



Most unlikely colors...



From Lalonde and Efros, ICCV'2007


Blur

- Look at frequency distribution.
- Measure the amount of blur in the sharpest object, instead of the *average* blur.




Prof photos
should
have some
part of
photo in
sharp focus

Low Level Features - Contrast

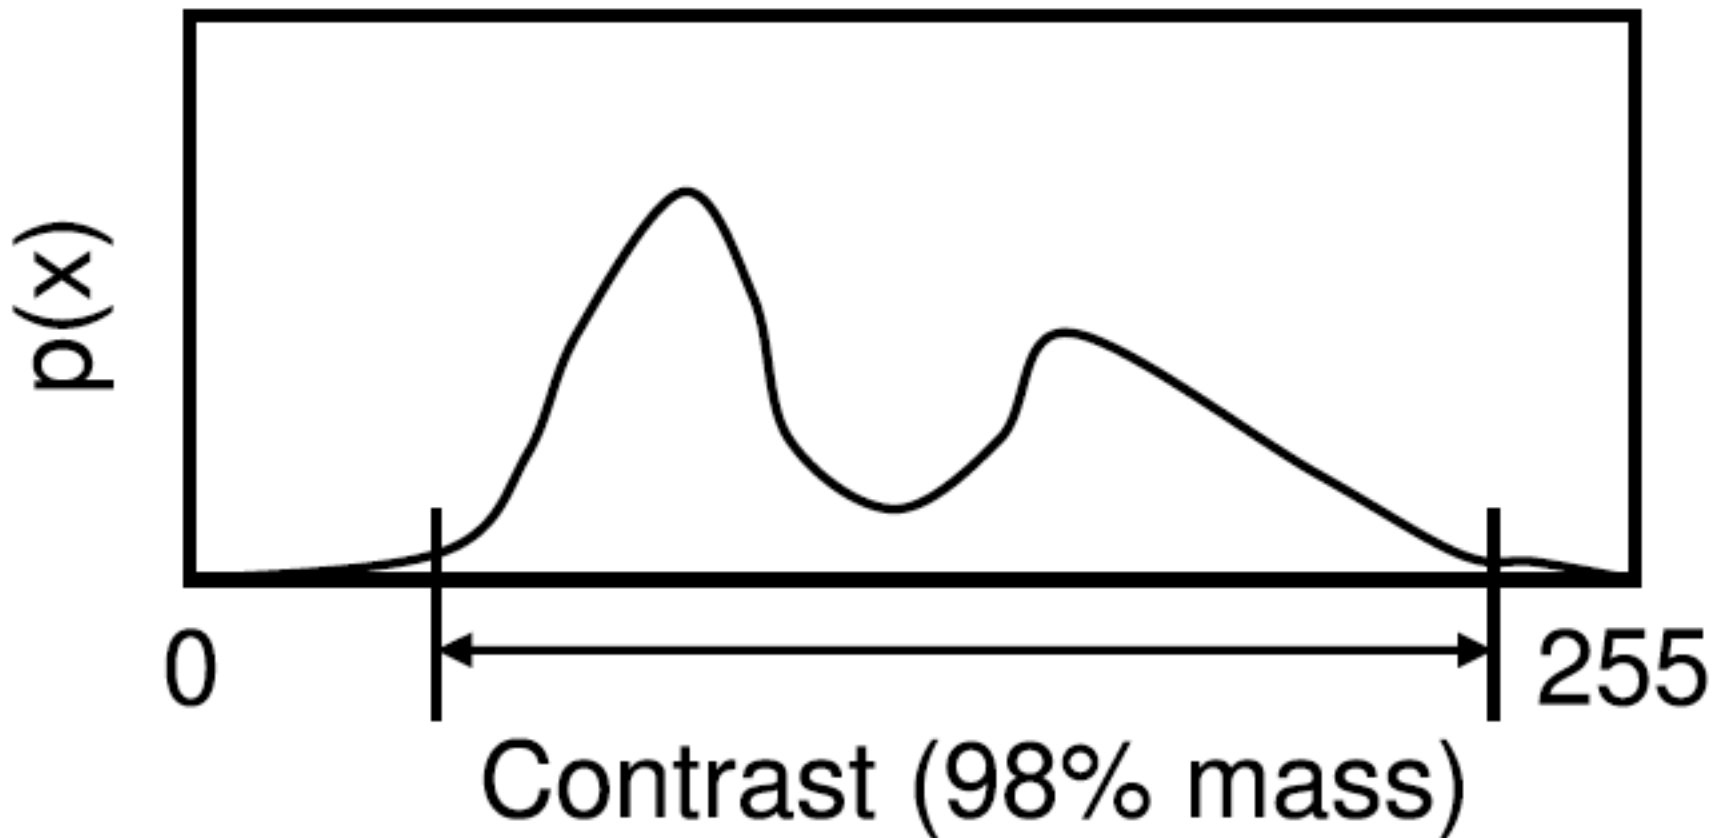


Prof photos usually have higher contrast



Contrast = width of middle 98% mass of hist



Contrast



Low Level Features – Avg. Brightness



Professional photographers may adjust exposure to be correct on subject only so subj pops from bkd. Cameras tend to adjust brightness to average at 50% gray, but prof photos might deviate significantly. Use ave brightness as feature.

Dataset – DPChallenge.com

Use photos average rating as ground truth quality measure

Use only top 10%, bottom 10% as dataset.

Use half for training/half for testing.



Statistics	Voting Breakdown
Place: 1 out of 829	1 0
Avg (all users): 7.987	2 1
Avg (commenters): 8.805	3 2
Avg (camera): 7.998	4 5
Avg (no camera): 6.333	5 24
Views since voting: 6597	6 54
Views during voting: 1003	7 89
Votes: 478	8 109
Comments: 190	9 89
Favorites: 133 (view)	10 105


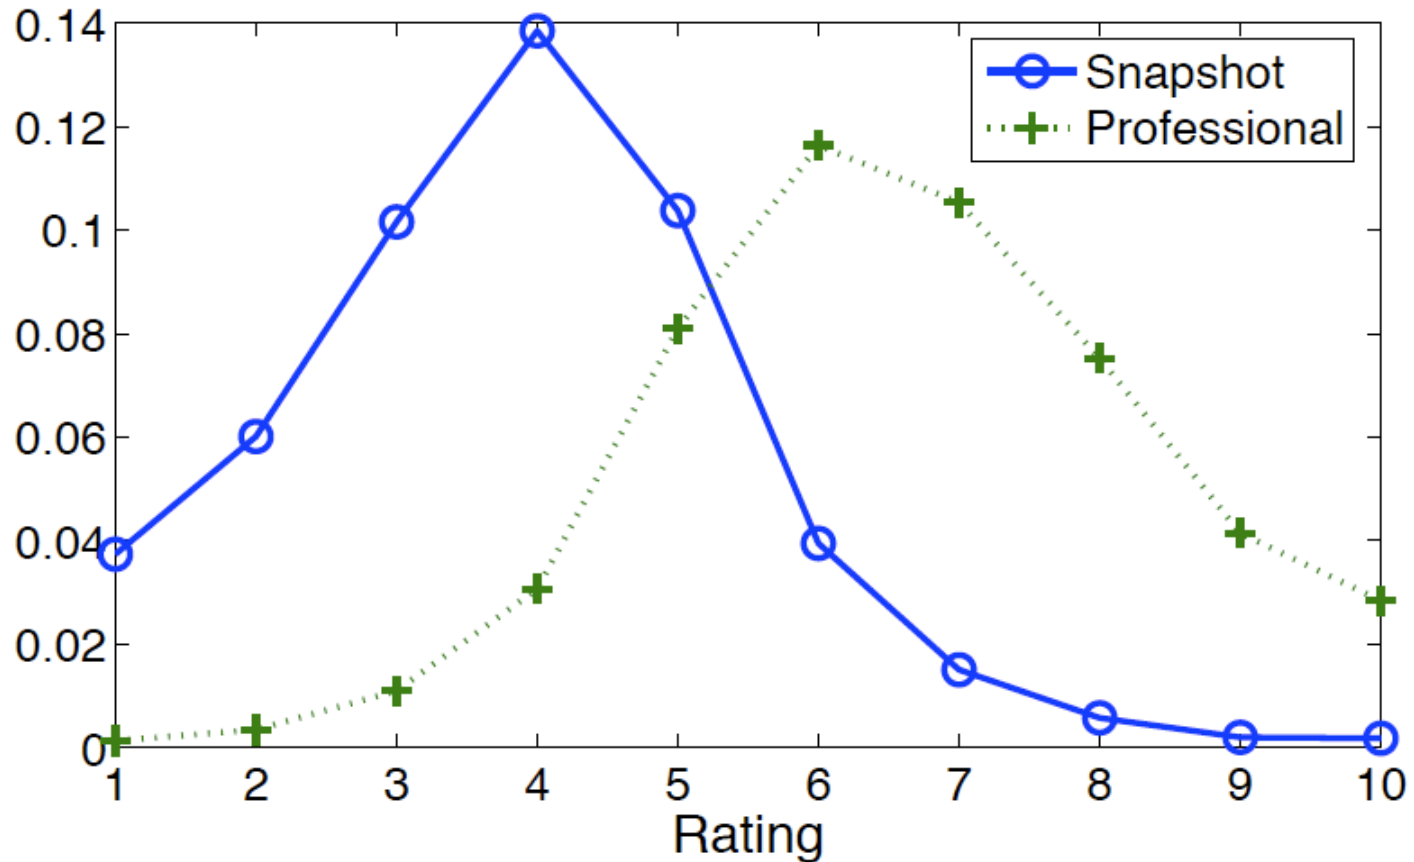
 [Add this photograph to your favorites!](#)

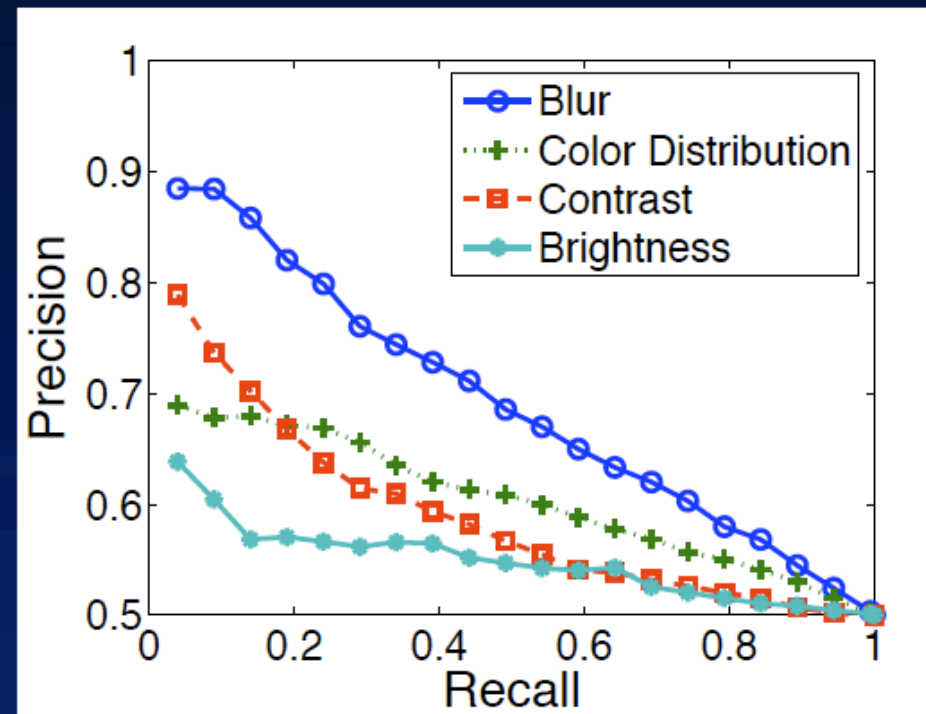
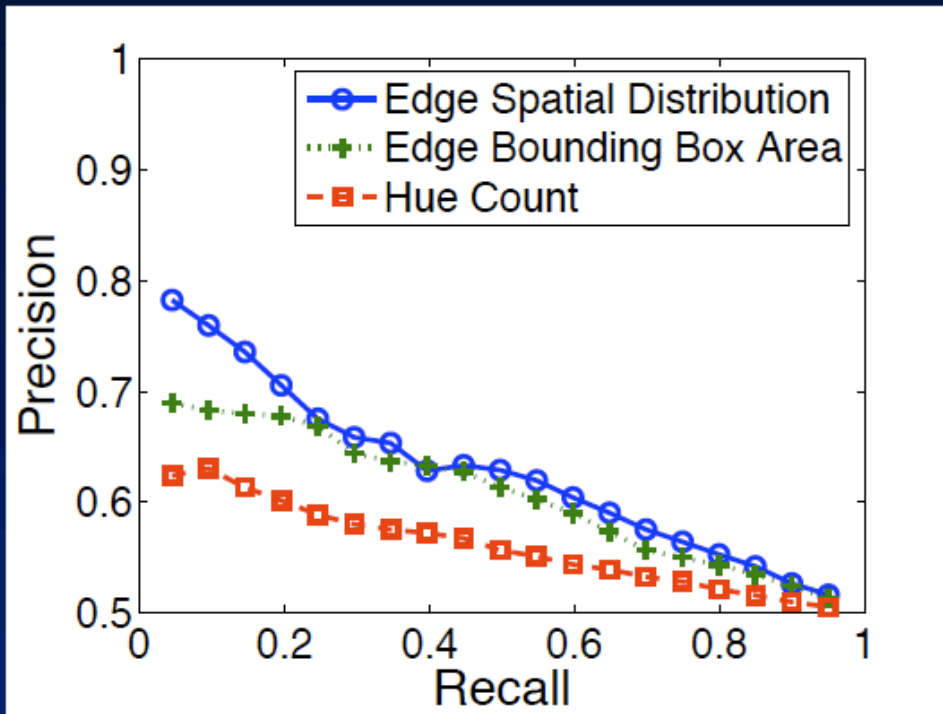
Photo contest website,
user rated

60K photos
40K photographers
10/90 percentile

Difficulty of Dataset



Results



$$\text{recall} = \frac{\# \text{ professional photos above threshold}}{\text{total } \# \text{ professional photos}}$$



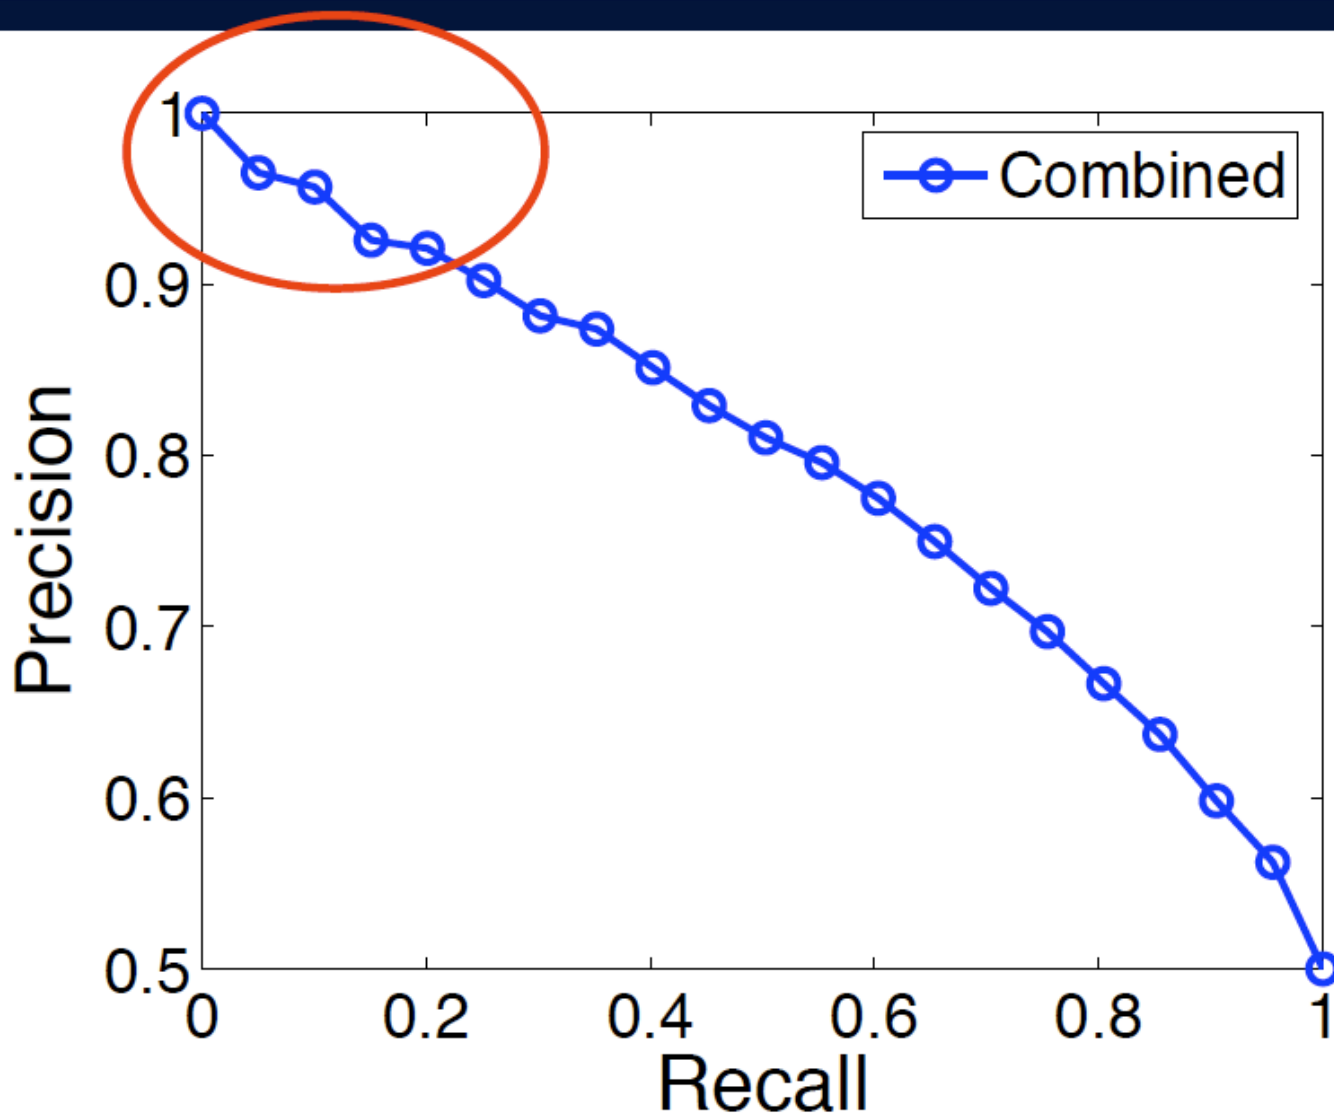
$$\text{precision} = \frac{\# \text{ professional photos above threshold}}{\# \text{ photos above threshold}}$$

Most Distinctive Feature: Blur

- A *badness* metric, rather than a *goodness* metric.



Results



72% classification rate

Web Retrieval Results



...



Web Retrieval Results





...



Web Retrieval Results



Lucky Imaging



Juliet Bernstein, **Aseem Agarwala**, Brian Curless. "Candid Portrait Selection From Video," *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia 2011)*

http://grail.cs.washington.edu/projects/candid_video_portraits/

THANK YOU!

