Saliency for image understanding and manipulation

Speaker: Ming-Ming Cheng

Nankai University

http://mmcheng.net/

My involvement to the problem

- Saliency estimation is key to many applications

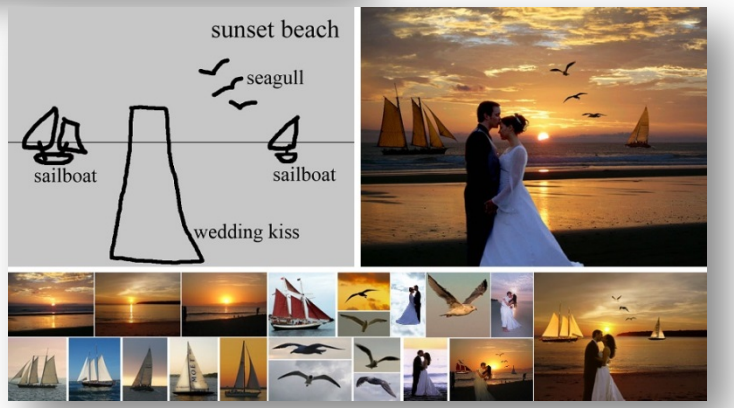

- Shrinkability Maps for Content-Aware Video Resizing, CGF 2008.
- A Shape-Preserving Approach to Image Resizing, CGF 2009.
- Sketch2Photo: internet image montage, ACM TOG 2009.

Predict fixation → detect salient object

- Fixation prediction
  - Predicting saliency points of human eye movement

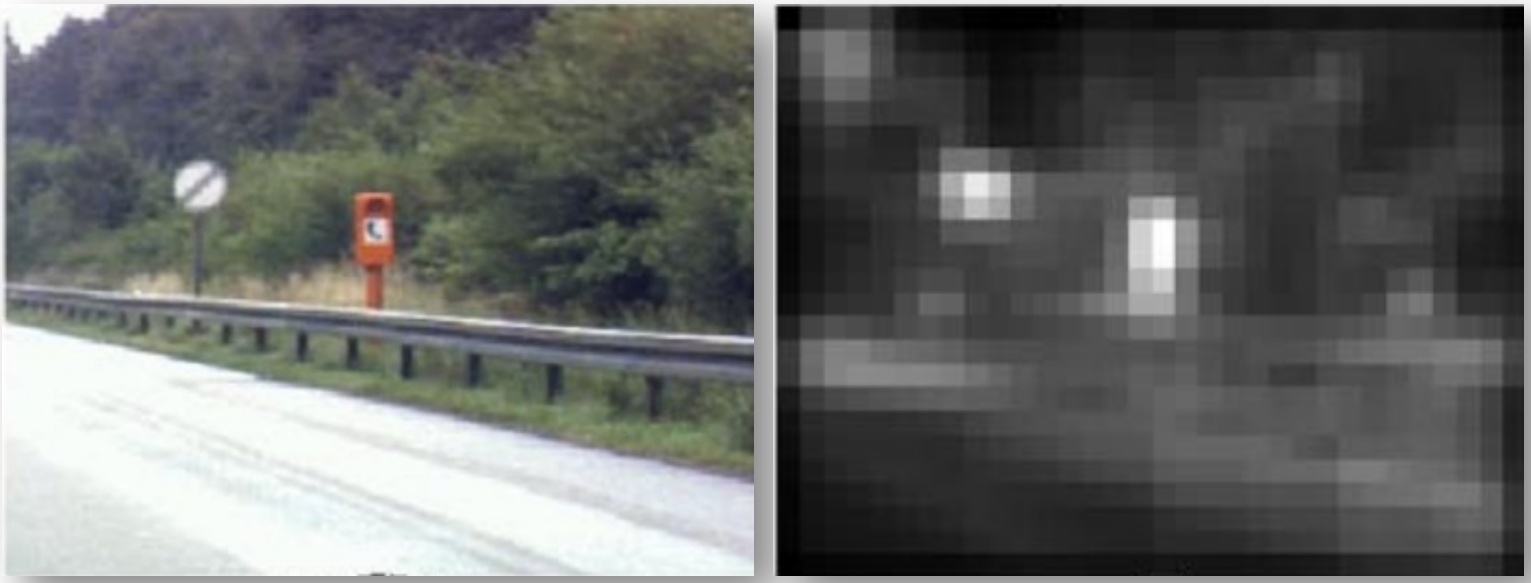

- A model of saliency-based visual attention for rapid scene analysis. PAMI 1998, Itti et al.
- Saliency detection: A spectral residual approach. CVPR 2007, Hou et. al.
- Graph-based visual saliency. NIPS, Harel et. al.
- Quantitative analysis of human-model agreement in visual saliency modeling: A comparative study, IEEE TIP 2012, Borji et. al.

Predict fixation → detect salient object

- Eye tracker
  - Cognitive psychology, neurobiology, etc.

Predict fixation → detect salient object

- Saliency detection as binary segmentation

(a) MSRA10K (b) ECSSD

- Learning to detect a salient object. CVPR 2007, Liu et. al.
- Frequency-tuned salient region detection, CVPR 2009, Achanta et. al.
- Global contrast based salient region detection, CVPR 2011, Cheng et. al.
- Salient object detection: a benchmark, IEEE TIP 2015, Ali et. al.

Salient object detection

- How to define salient objects?
  - Match the human annotators' behavior when they have been asked to pick a salient object in an image.

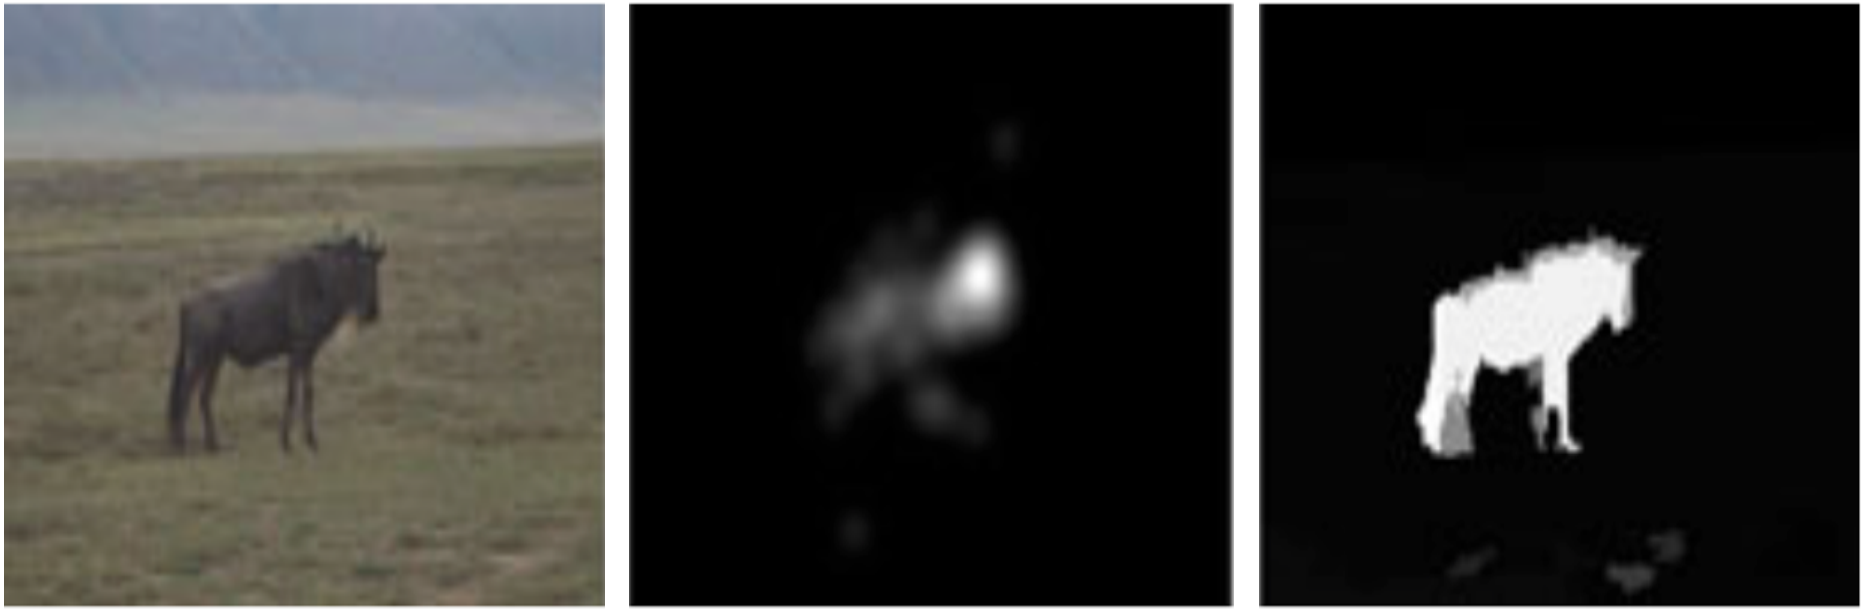

Global contrast based salient region detection, IEEE TPAMI 2015 (CVPR 2011), Cheng et. al.

Salient object detection

- High consistency among labelers.

PASCAL-S dataset

The Secrets of Salient Object Segmentation, CVPR 2014, Li et. al.

Growing interest

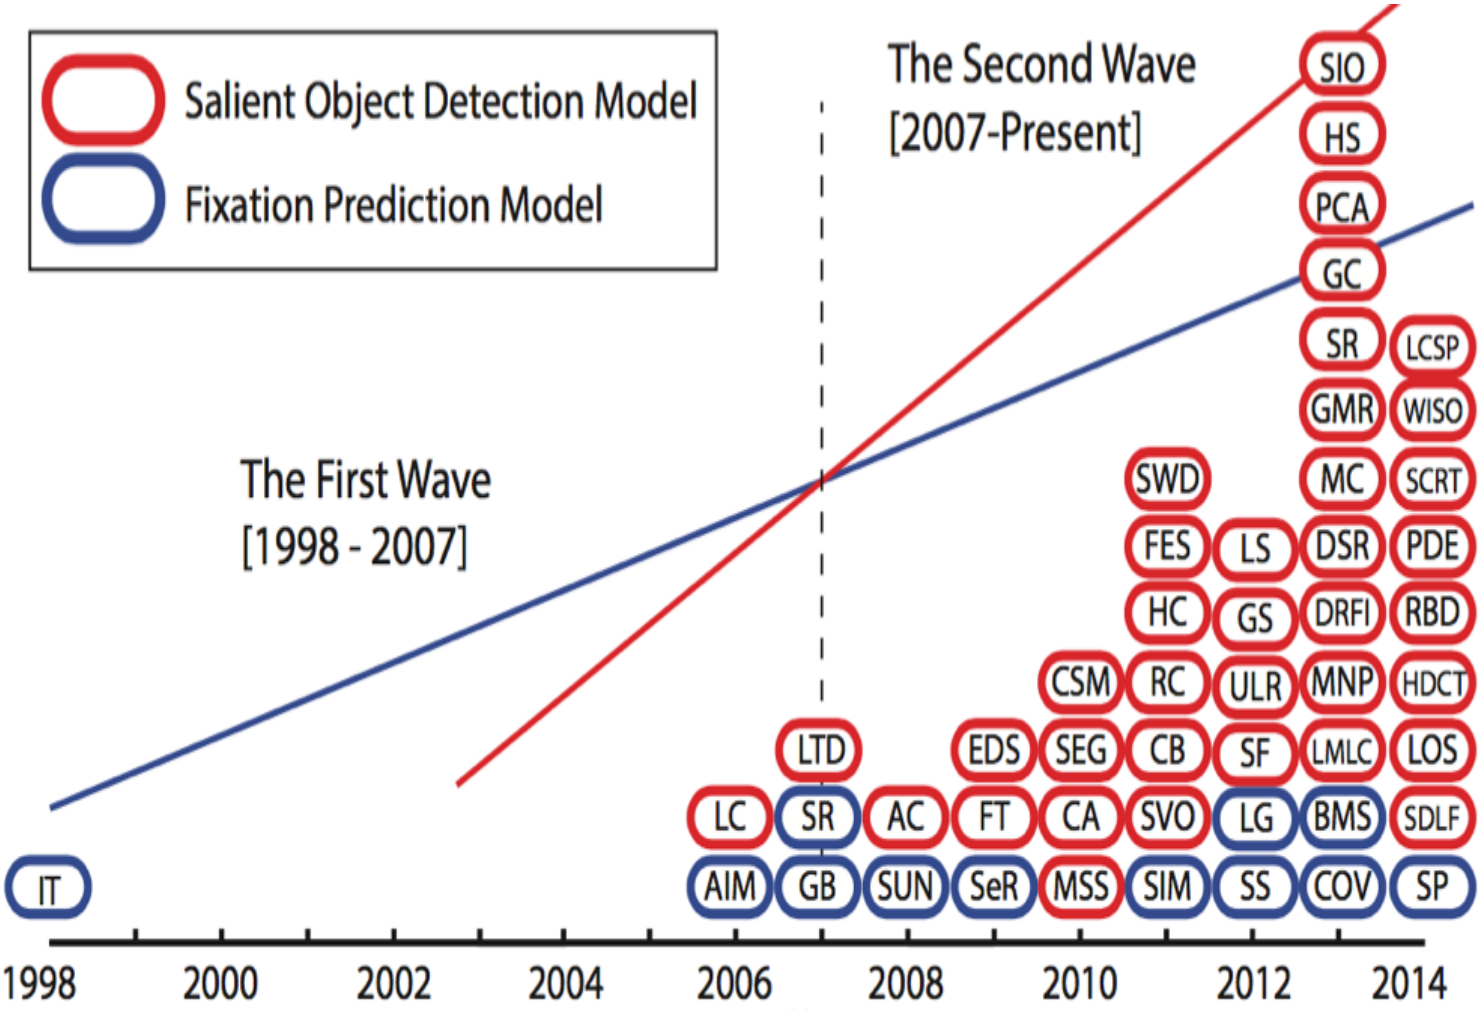

- Salient object detection: a survey, arXiv 2014, Ali et. al.
- Salient object detection: a benchmark, IEEE TIP 2015, Ali et. al.

Saliency detection

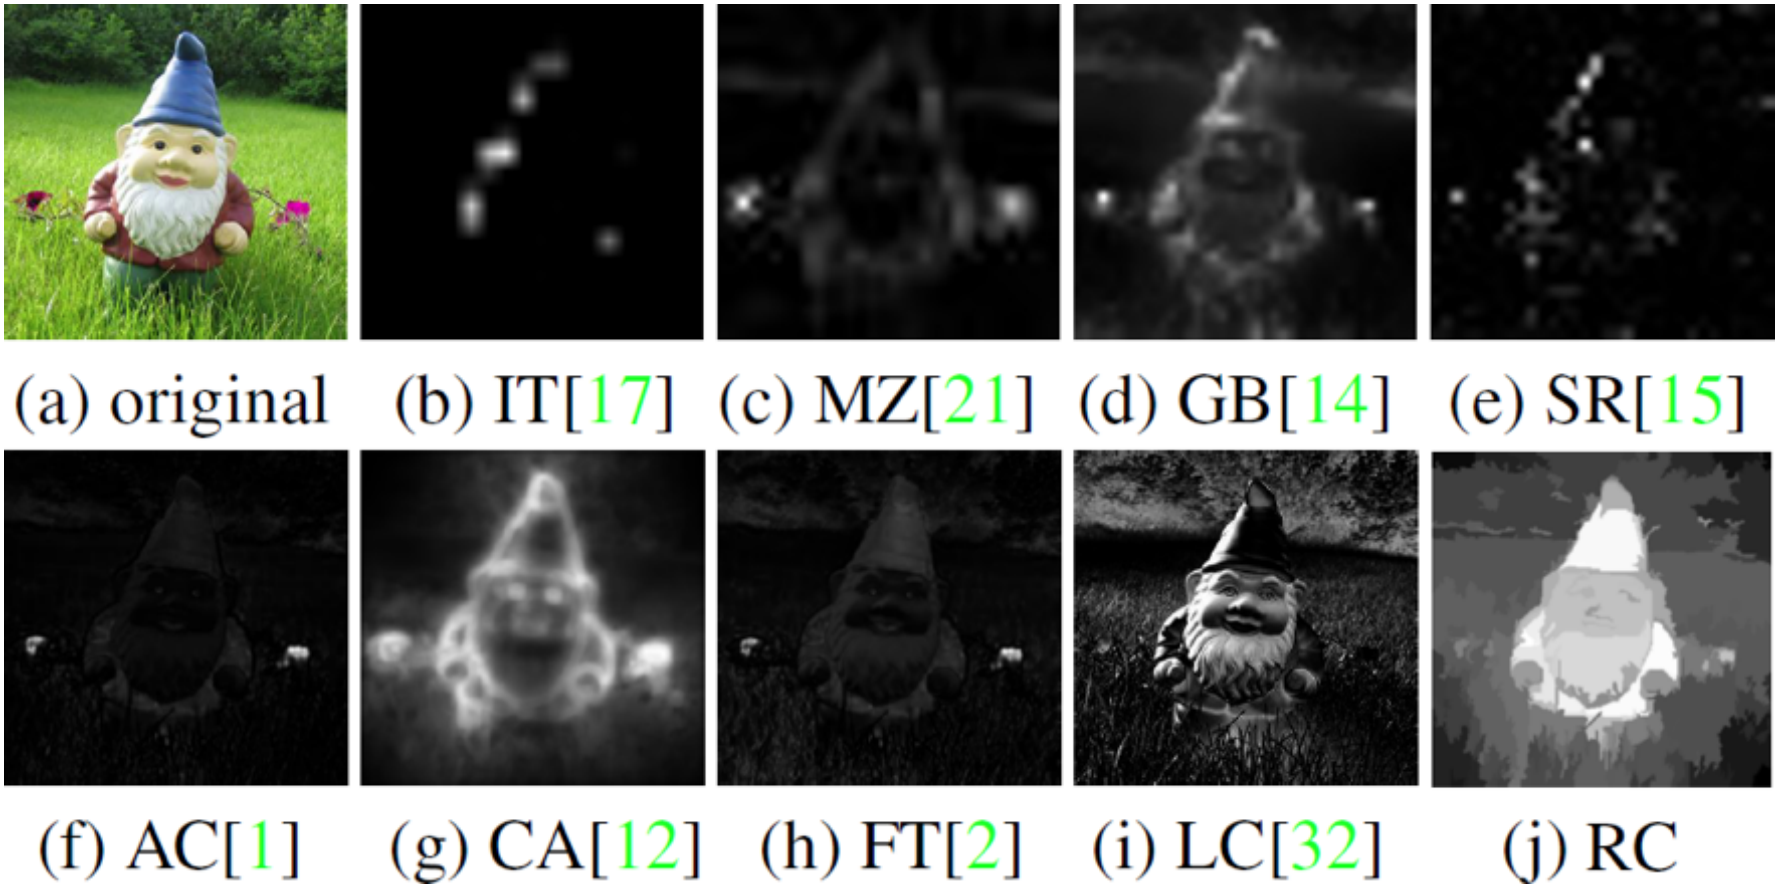

Region contrast (RC)

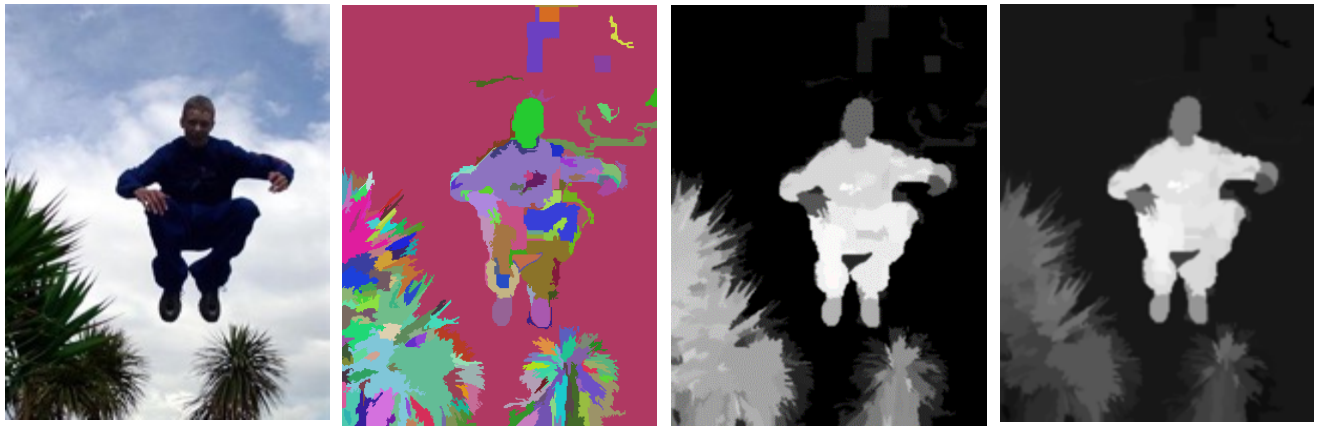

Image | Segmentation | $\sigma_s^2 \to \infty$ | $\sigma_s^2 \to 0.4$

Spatial weighting

Region size

$$S(r_k) = \sum_{r_k \neq r_i} \exp(-D_s(r_k, r_i)/\sigma_s^2)\, \omega(r_i) D_r(r_k, r_i)$$

Region contrast by sparse histogram comparison.

- Global Contrast based Salient Region detection. IEEE TPAMI 2015 (CVPR 2011), Cheng et al.

Experimental results

- Dataset: MSRA1000 [Achanta09]
  - Precision vs. recall

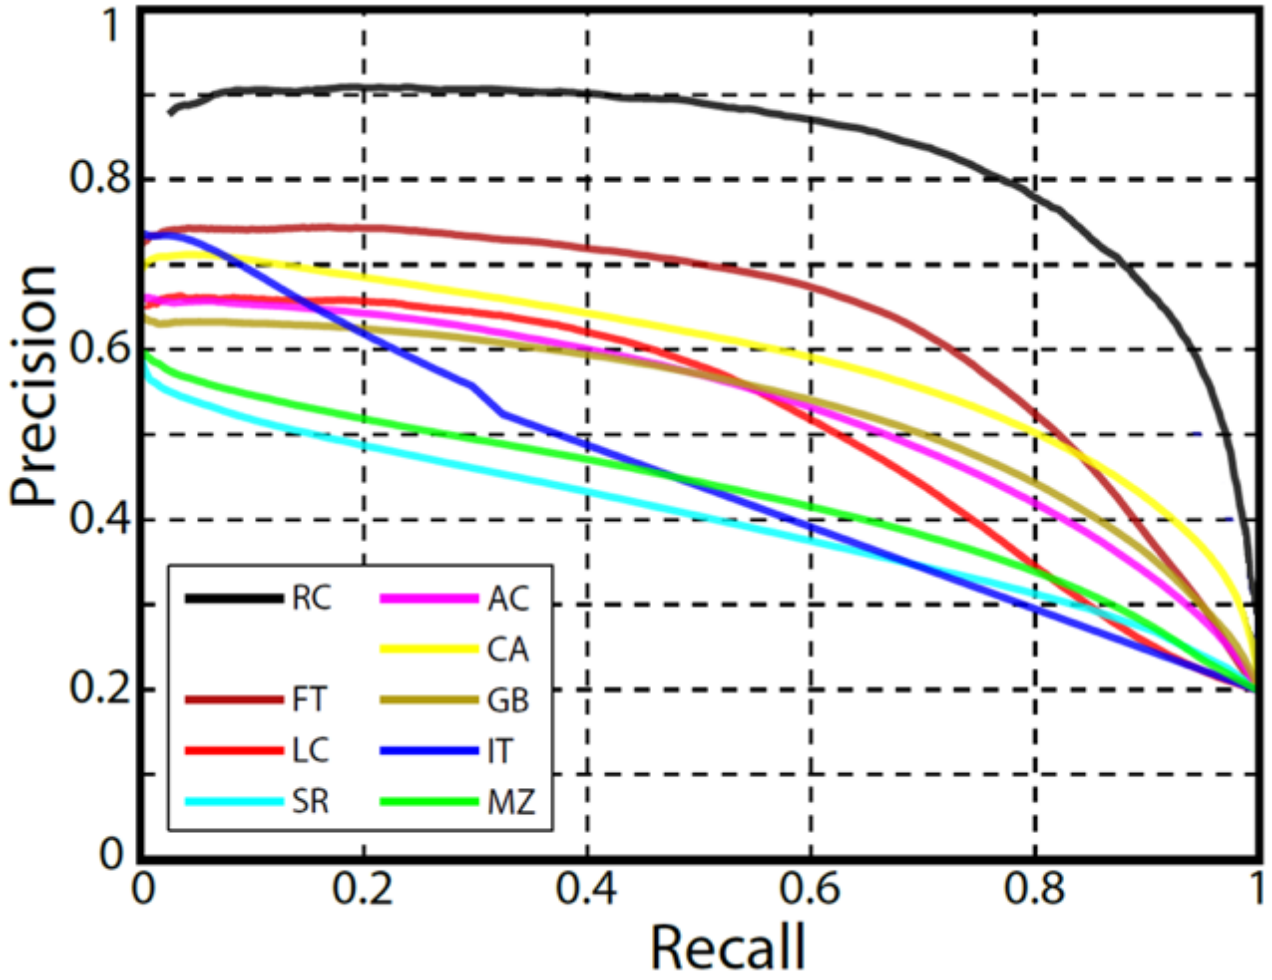

What makes an object salient?

- Contrast
- Spatial distribution
- Focusness
- Backgroundness
- Center bais/prior
- etc.

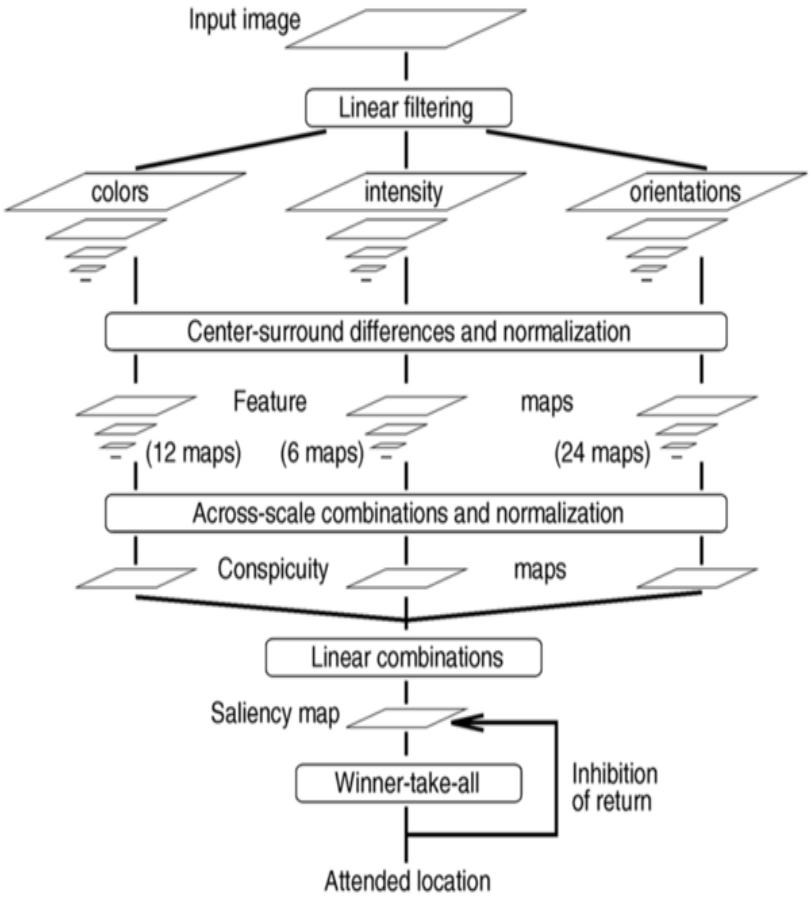

- A Model of Saliency-Based Visual Attention for Rapid Scene Analysis, IEEE TPAMI 1998, Itti et al.
- Salient object detection: a survey, arXiv 2014, Ali et. al.

Content aware image resizing

A Shape-Preserving Approach to Image Resizing, CGF 2009.

Image mosaic

Saliency for image Manipulation, The Visual Computer 2013, Margolin et al.

SaliencyCut

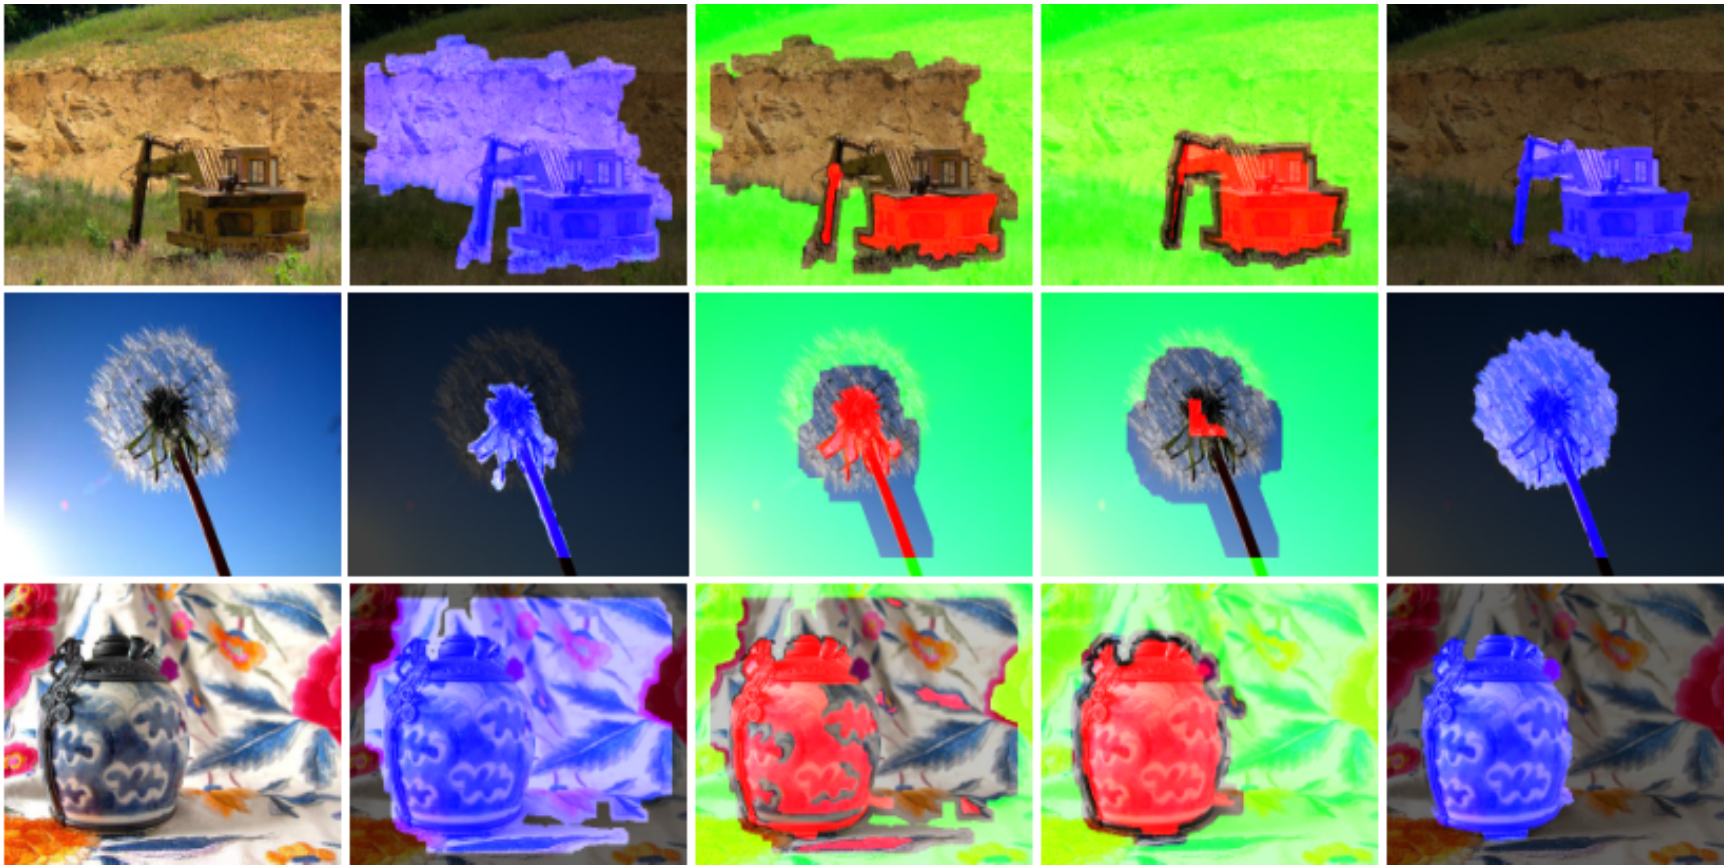

Enables automatic initialization provided by salient object detection.

Application requirements

- Is salient object detection for ‘simple’ images useful?

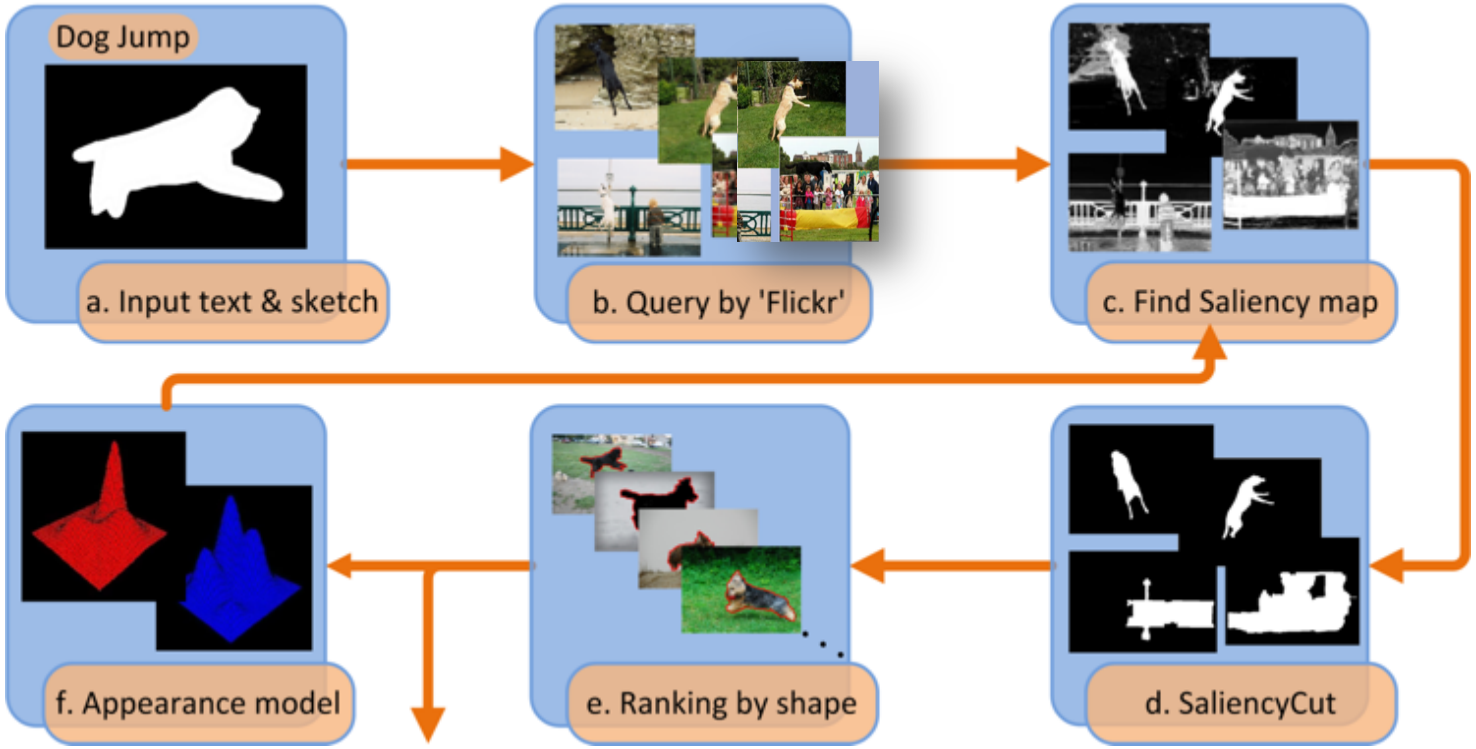

SalientShape: Group Saliency in Image Collections, The Visual Computer 2014. Cheng et. al.

Application requirements

- Illustration of learned appearance models
  - Accords with our understanding of these categories

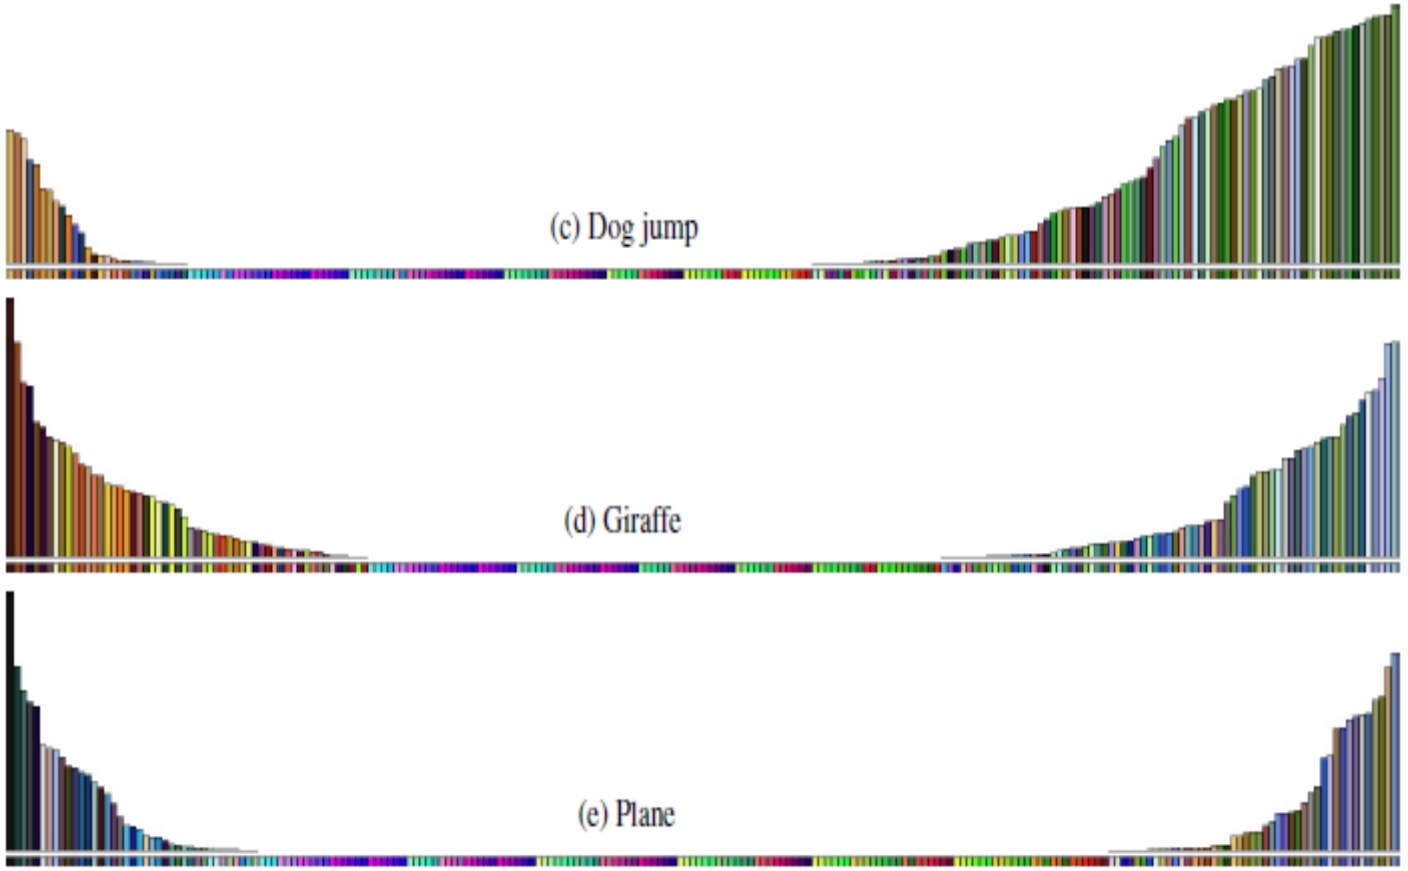

Applications

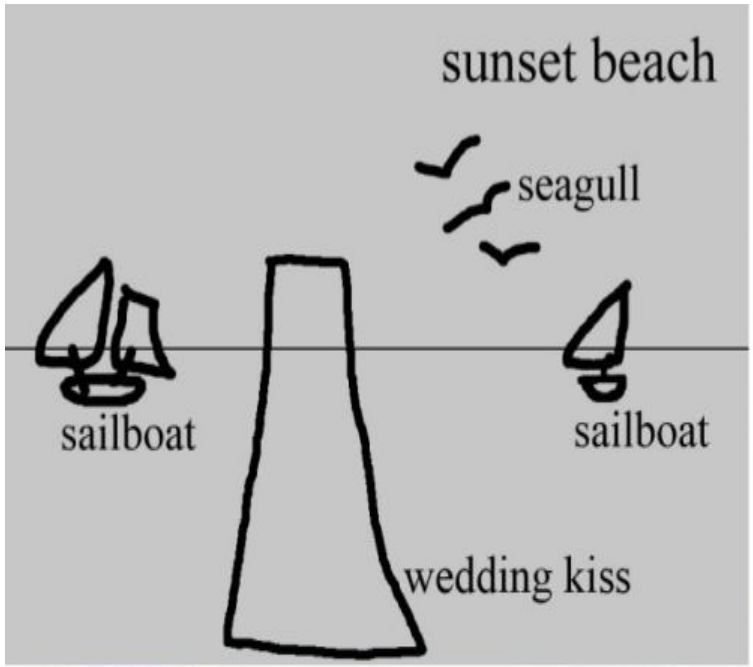

- Sketch2Photo: internet image montage, ACM TOG 2009. Chen et. al.

Image Collage

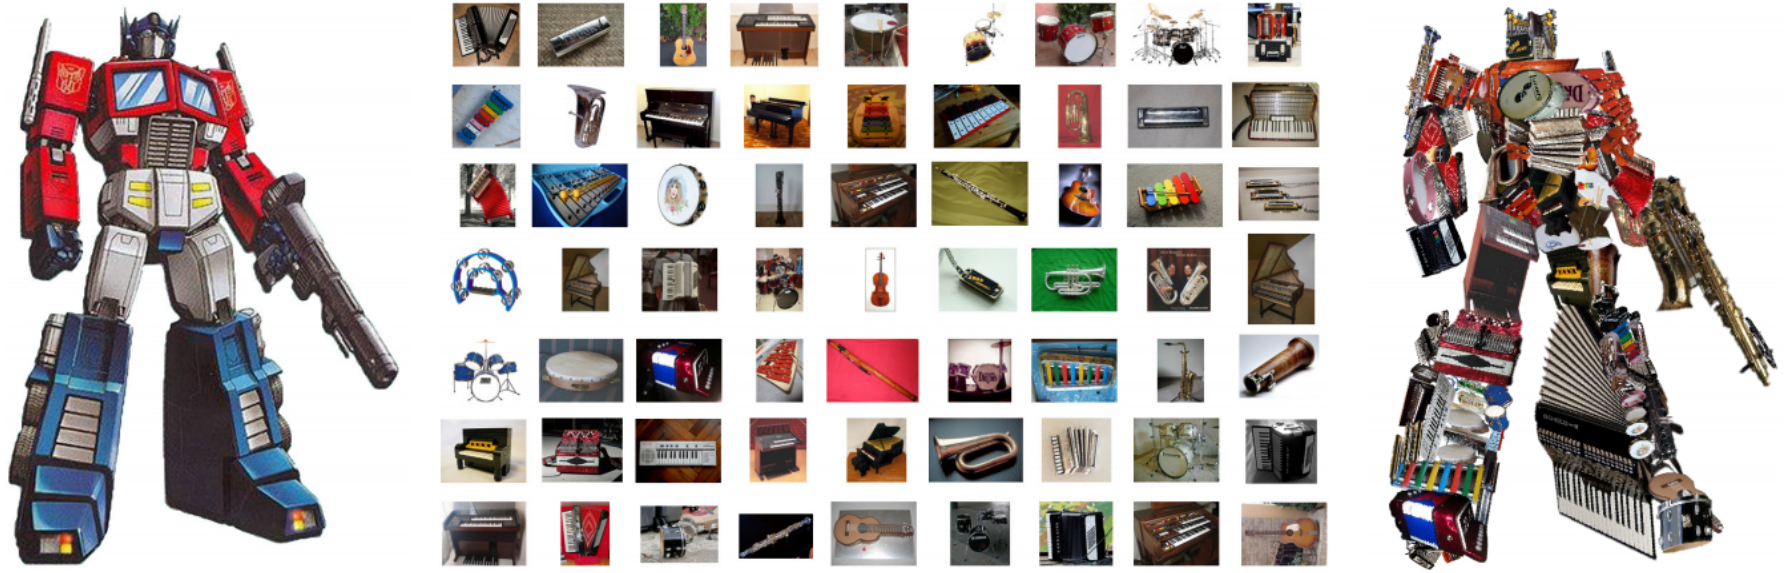

- Arcimboldo-like Collage Using Internet Images, ACM TOG 2011. Huang et. al.

View selection

- Web-image driven best views of 3d shapes, The Visual Computer 2012. Liu et al.

Colorization

- Semantic Colorization with Internet Images, ACM TOG 2011. Chia et al.

Unsupervised object discovery

- Unsupervised Object Discovery via Saliency-Guided Multiple Class Learning, IEEE CVPR 2012. Zhu et al.

Weakly supervised sematic segmentation

- Example of simple images in Flickr and saliency maps

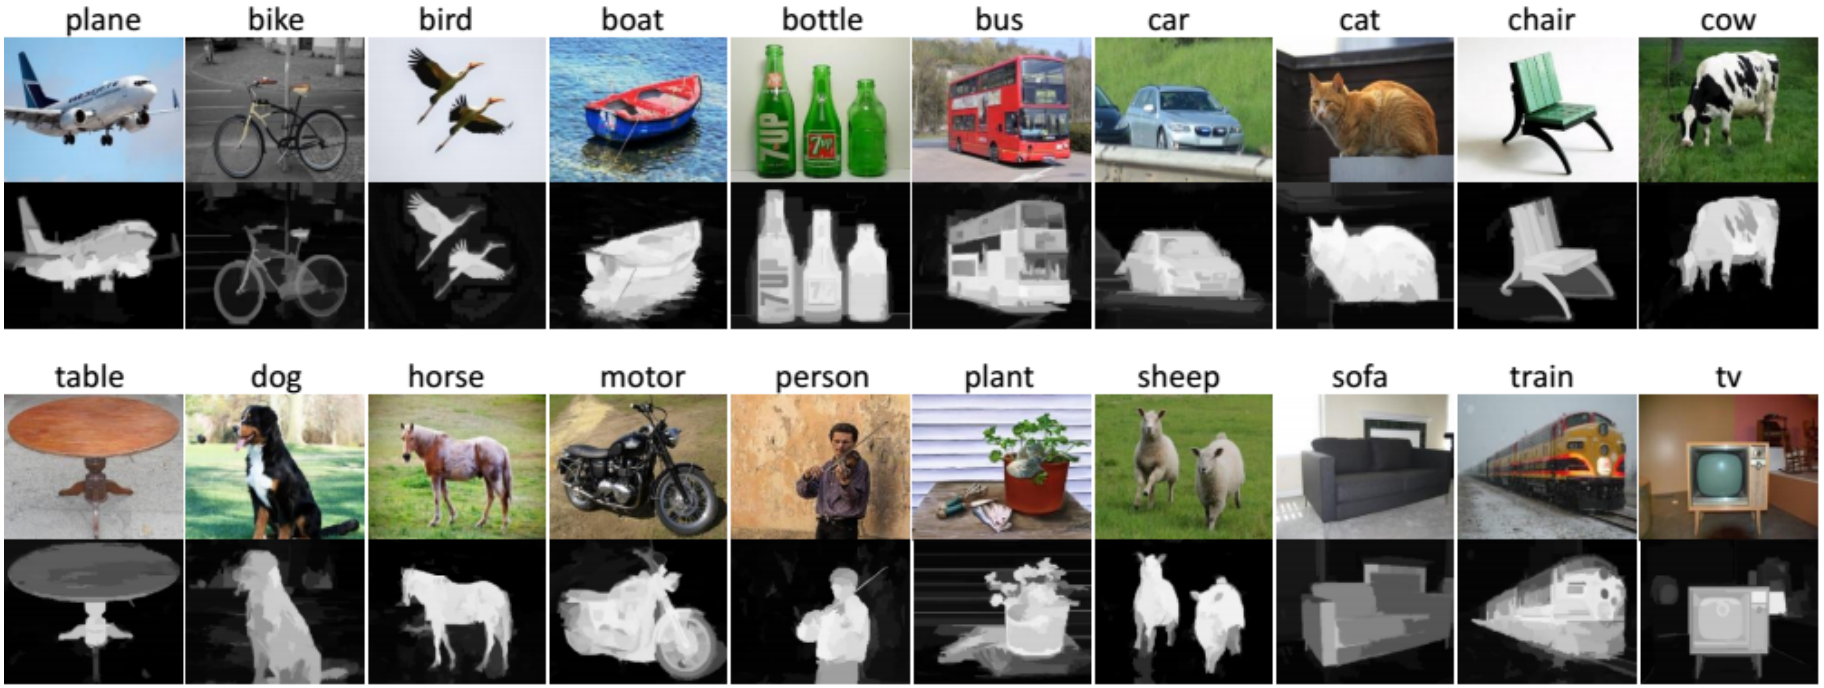

Weakly supervised sematic segmentation

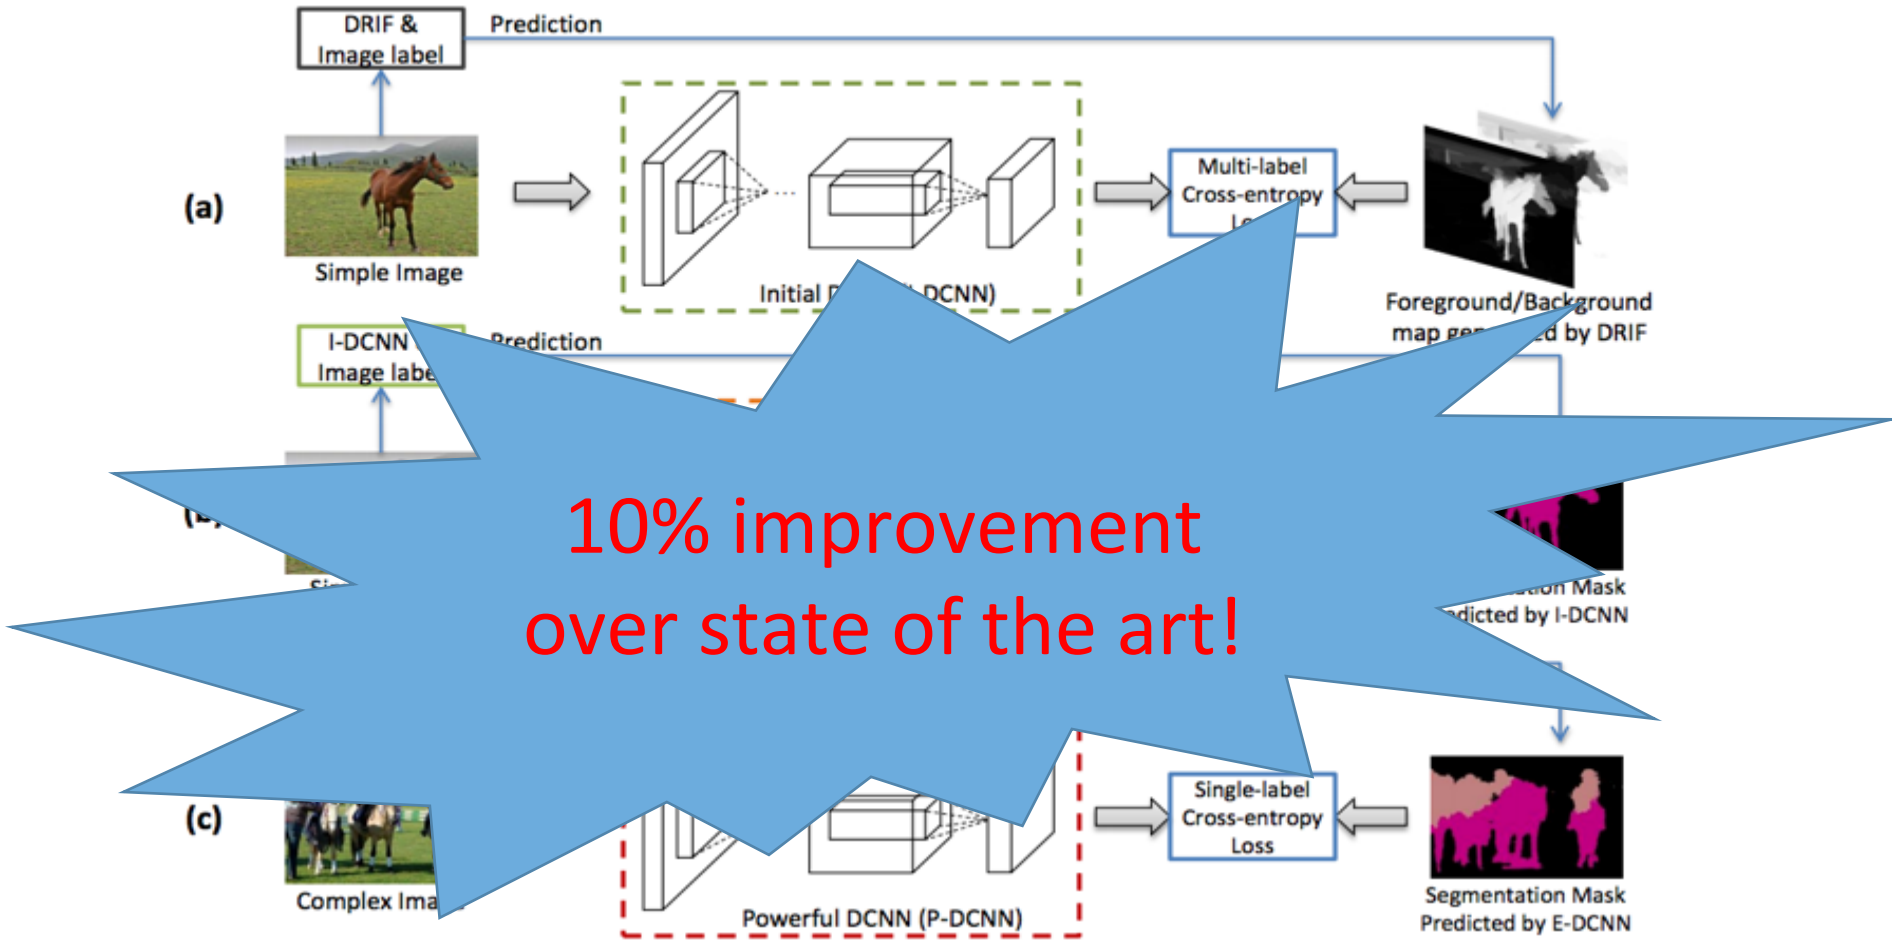

- STC: A Simple to Complex Framework for Weakly-supervised Semantic Segmentation, Wei et al., arXiv 2015.

Applications: what we learnt?

Don't ask what segments can do for you, ask what you can do for the segments.

— Jitendra Malik

Salient Object Subitizing

- Predicting the number of salient objects
  - Is it possible without any object localization process?
  - How fast can you tell the number of prominent objects in each of these images?
  - Are we able to do it using holistic cues only?



Salient Object Subitizing, IEEE CVPR 2015, Zhang et al.

Joint Detection and Existence Prediction

- Joint Salient Object Detection and Existence Prediction, Computational Visual Media, Jiang et al.

How about complicated images?

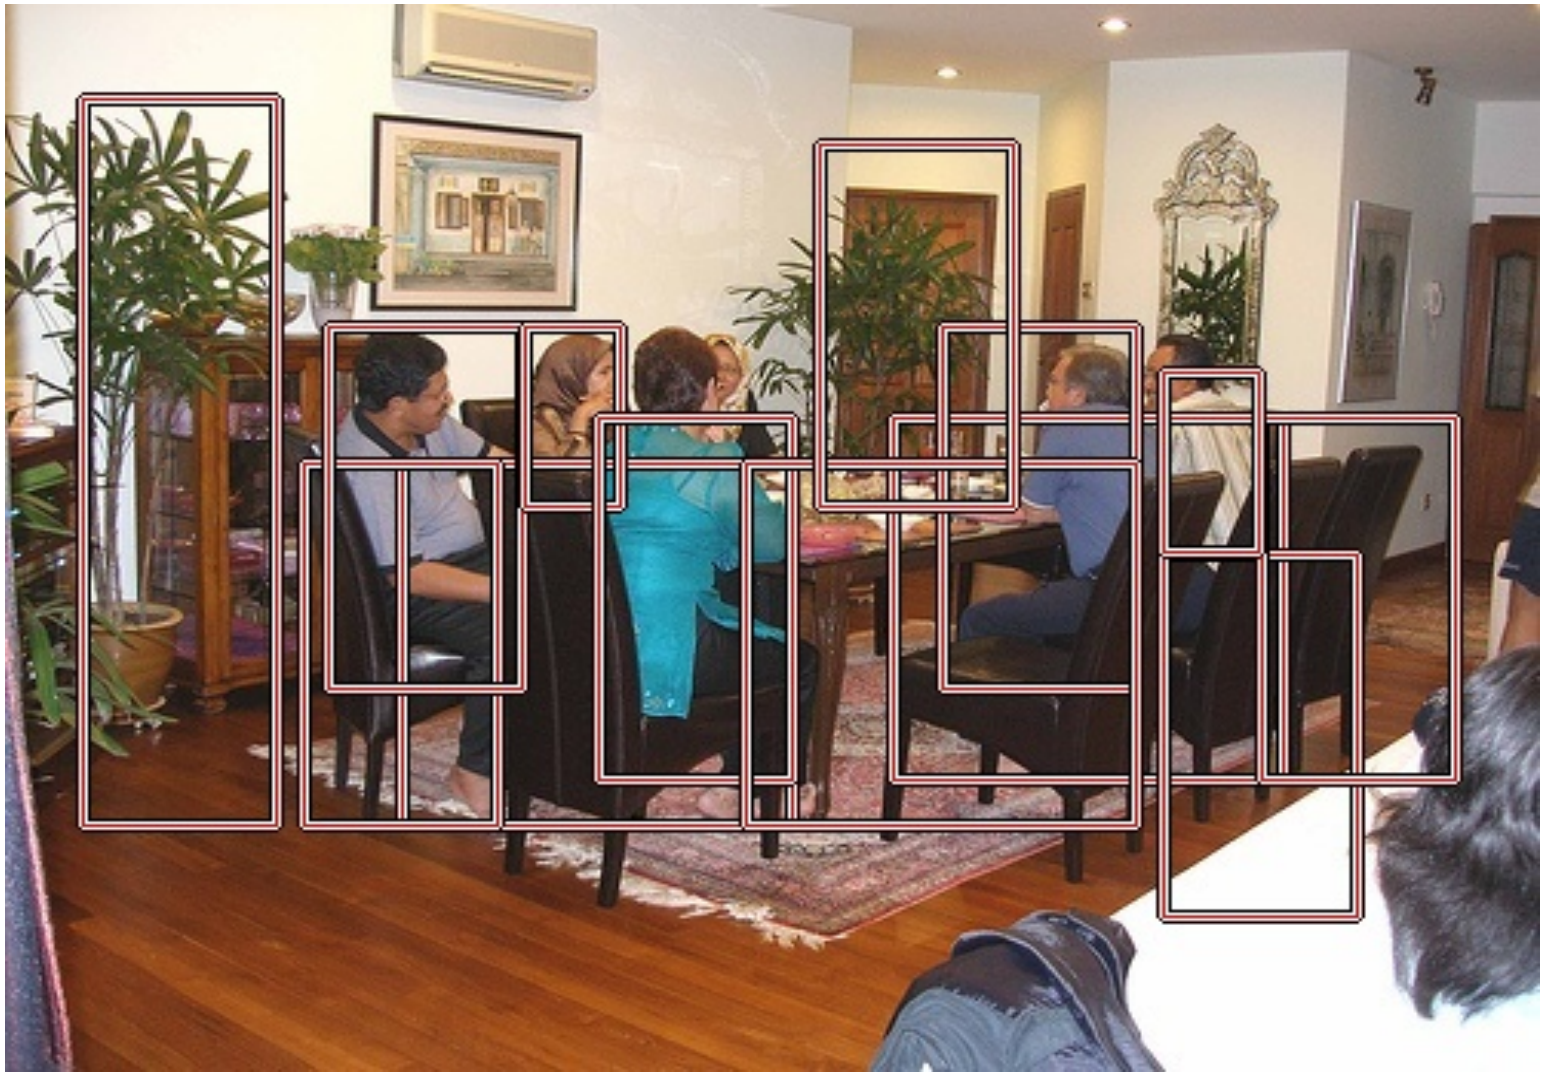

Objectness proposals

Motivation: What is an object?

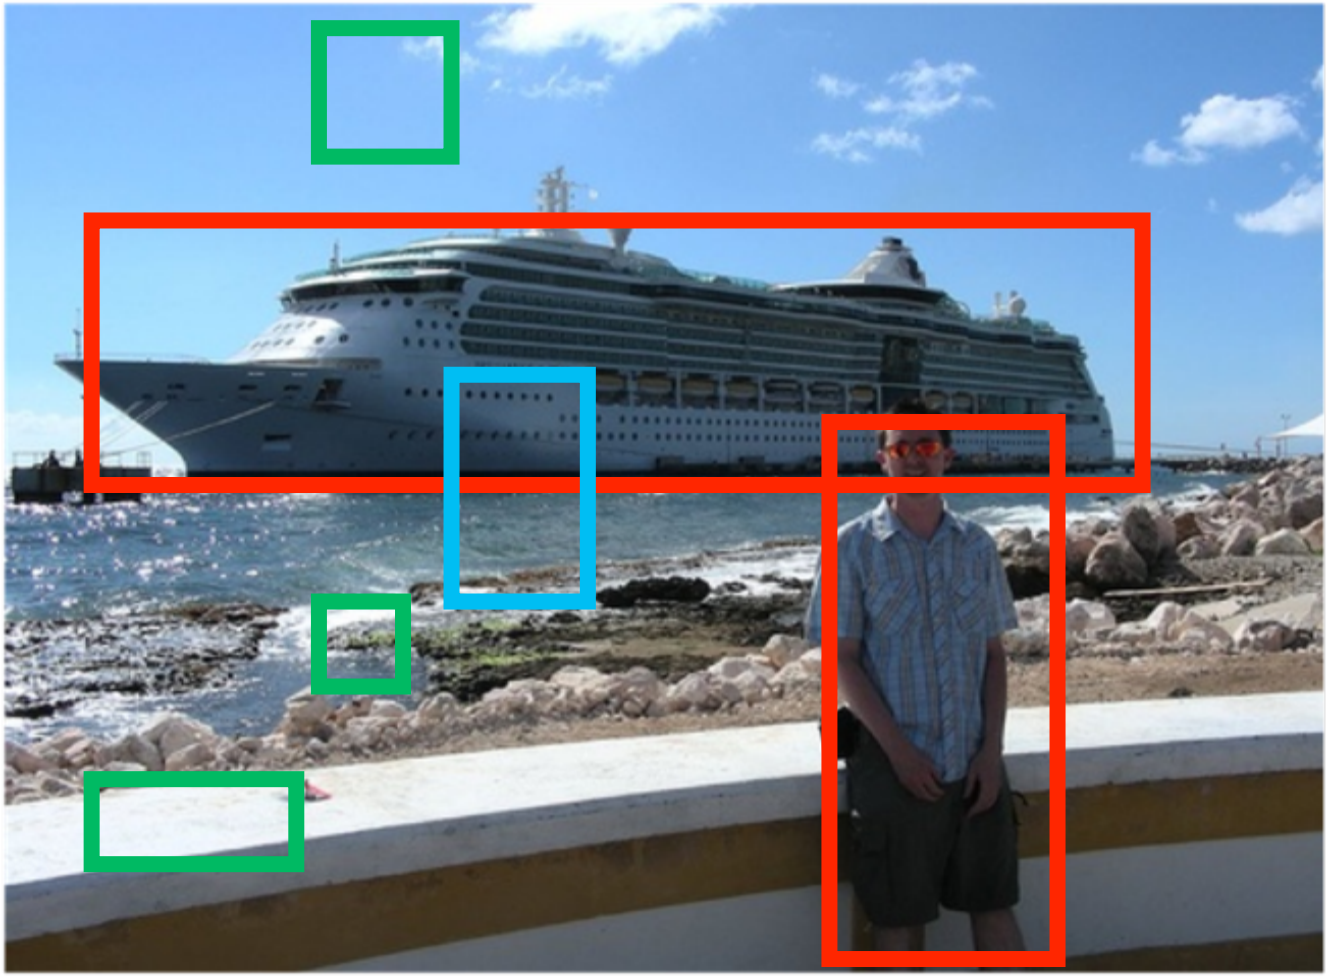

Motivation: What is an object?

- An objectness measure
  - A value to reflects how likely an image window covers an object of any category.
- What's the benefits?
  - Improve computational efficiency, reduce the search space
  - Allowing the usage of strong classifiers during testing, improve accuracy

Measuring the objectness of image window, IEEE TPAMI 2012, Alexe et. al.

A feature integration approach

- Objectness proposal generation
  - A small number (e.g. 1K) of category-independent proposals
  - Expected to cover all objects in an image

Measuring the objectness of image windows. PAMI 2012, Alexe, et. Al.

Region merging & Diversification

- Region merging
  - Merge two most similar regions based on region similarity.
  - Update similarities between the new region and its neighbors.
- Diversification

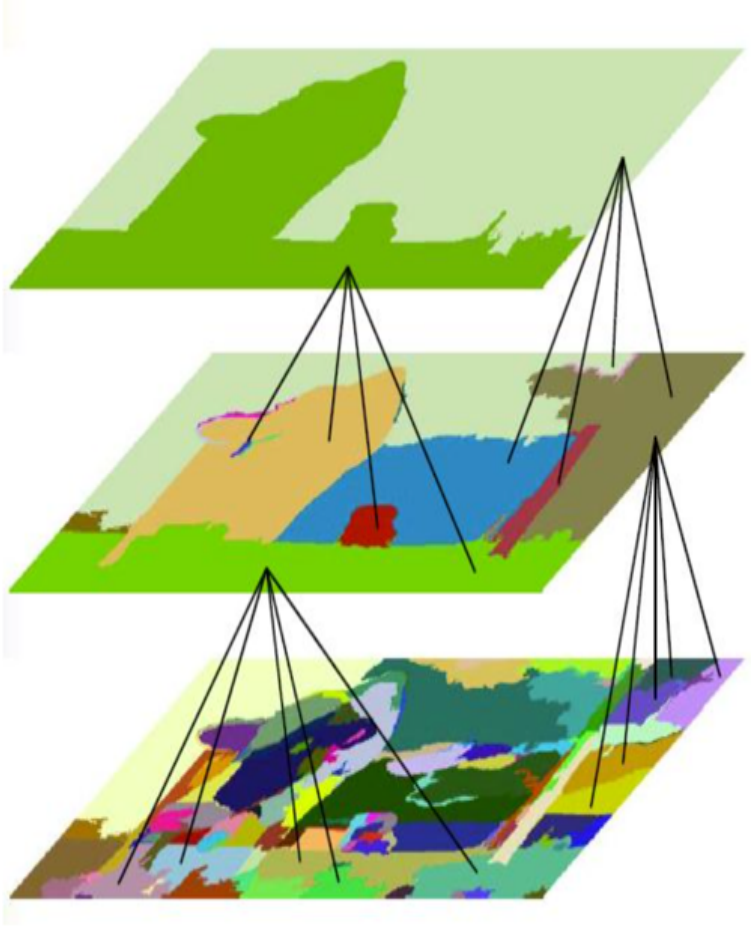

Selective Search for Object Recognition, IJCV 2013, Uijlings et. Al.

Local & global search

- Local search
  - Unsuitable for object with distinct parts
- Global search
  - Initialize with foreground/background seeds
  - A global optimization function for each parameter set

Generating object segmentation proposals using global and local search, CVPR 2014, Rantalankila et al.

BING method

- Our observation: a small interactive demo
  - Take you pen and paper and draw an object which is current in your mind.
  - What if we resize it to a tiny fixed size?
    - E.g. 8x8. Not only changing the scale, but also aspect ratio.

BING: Binarized Normed Gradients for Objectness Estimation at 300fp, IEEE CVPR 2014 (Oral), M.M. Cheng, et. al.

BING method

- Objects are stand-alone things with well defined **closed boundaries** and centers.

- Finding pictures of objects in large collections of images. Springer Berlin Heidelberg, 1996, Forsyth et. al.
- Using stuff to find things. ECCV 2008, Heitz et. al.
- Measuring the objectness of image window, IEEE TPAMI 2012, Alexe et. al.

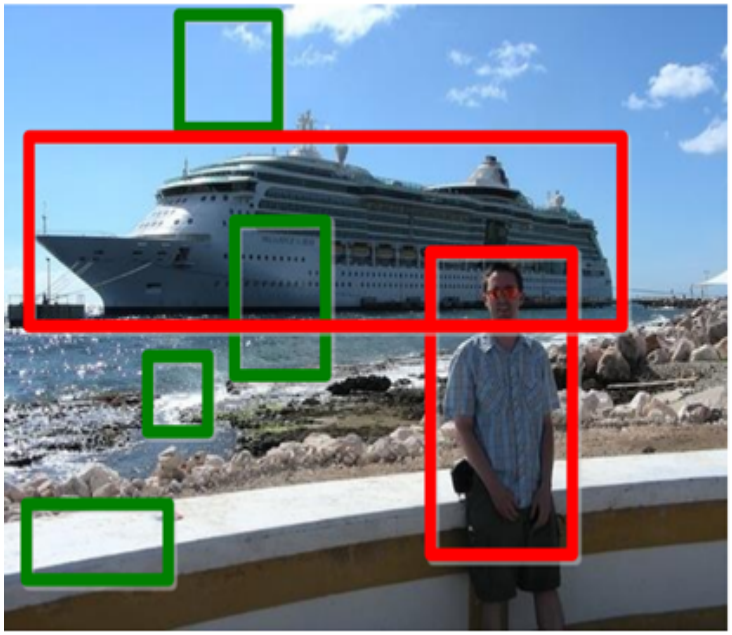

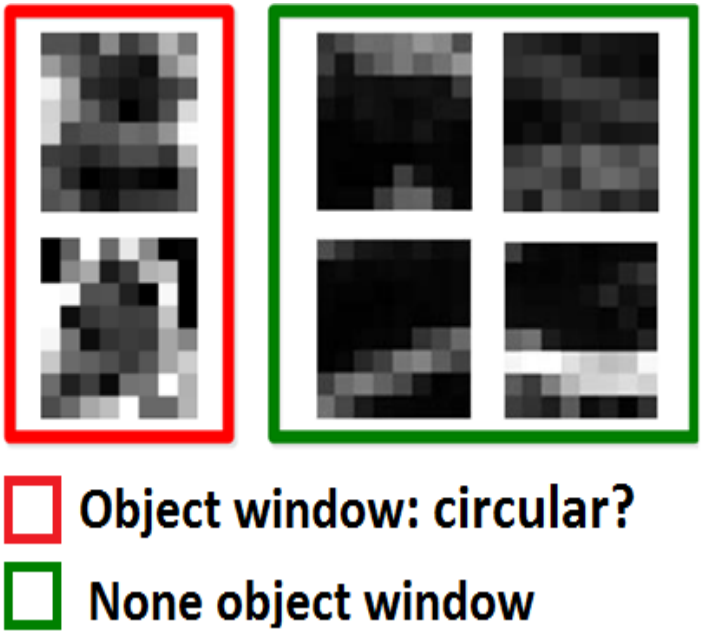

Experimental results of BING method

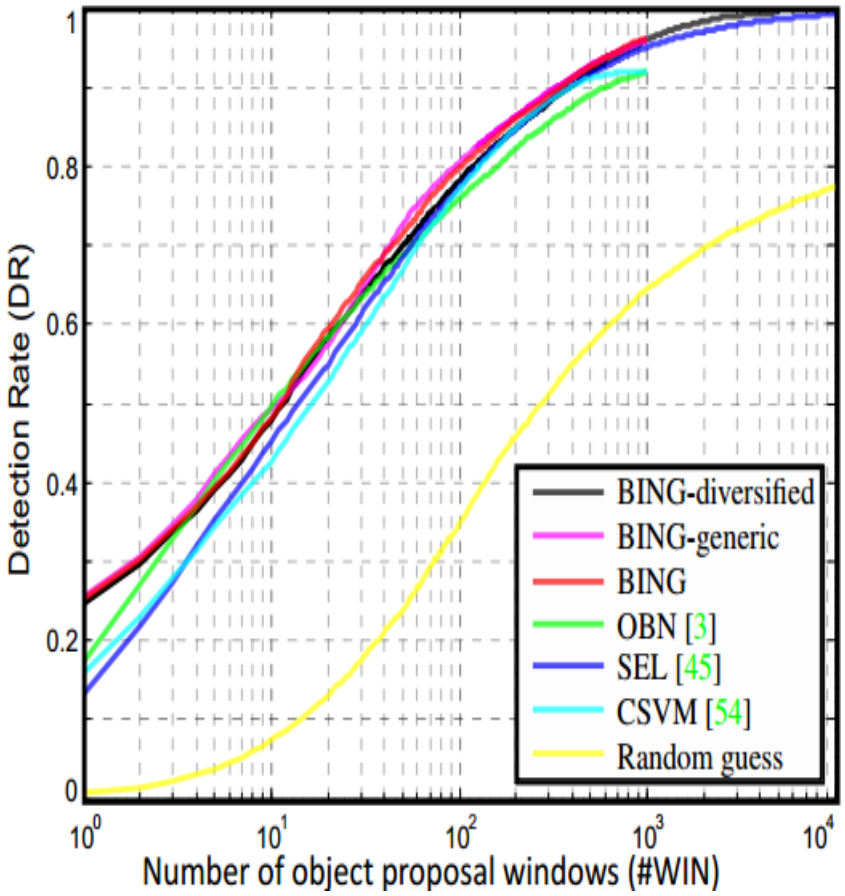

Method	[22]	OBN [3]	CSVM [57]	SEL [48]	Our BING
Time (seconds)	89.2	3.14	1.32	11.2	**0.003**

Table 1. Average computational time on VOC2007.

Combining boxes and regions

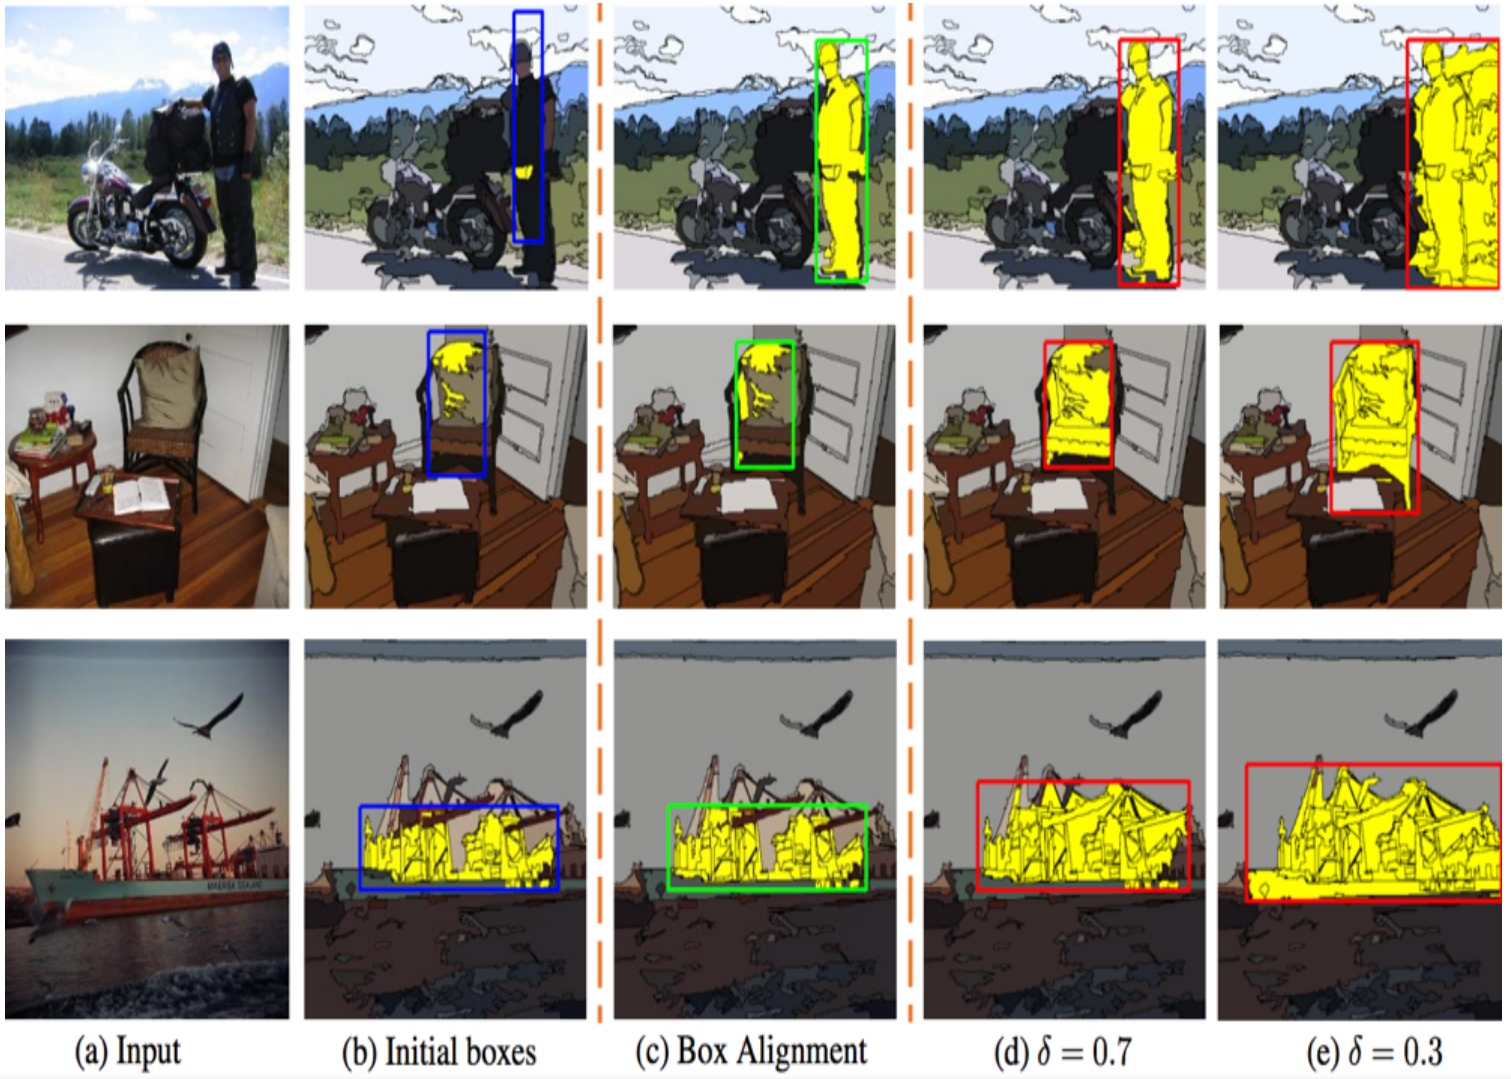

Improving Object Proposals with Multi-Thresholding Straddling Expansion, IEEE CVPR 2015, Chen, et. al.

Combining boxes and regions

- Experimental results

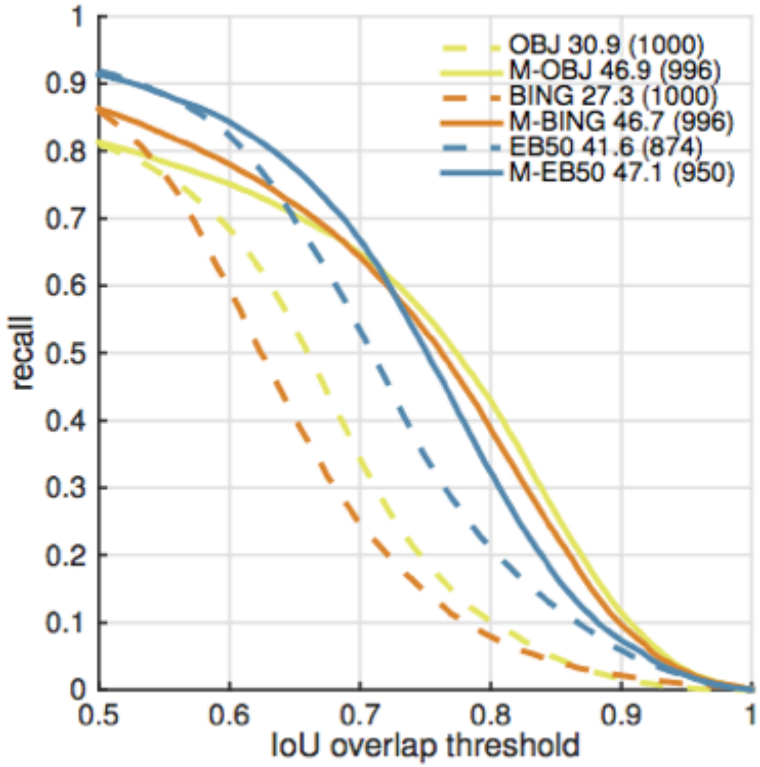

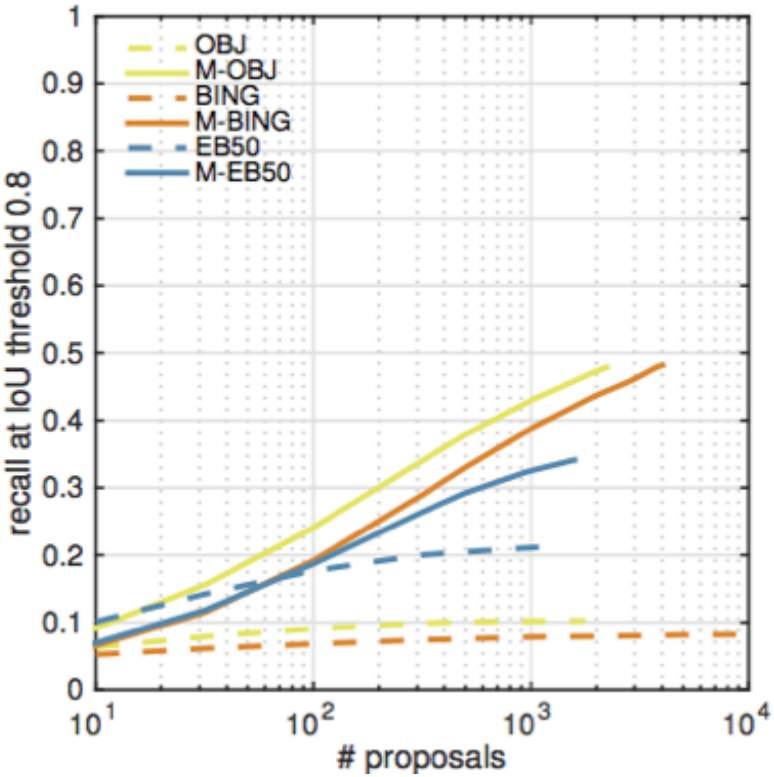

Improving Object Proposals with Multi-Thresholding Straddling Expansion, IEEE CVPR 2015, Chen, et. al.

Applications

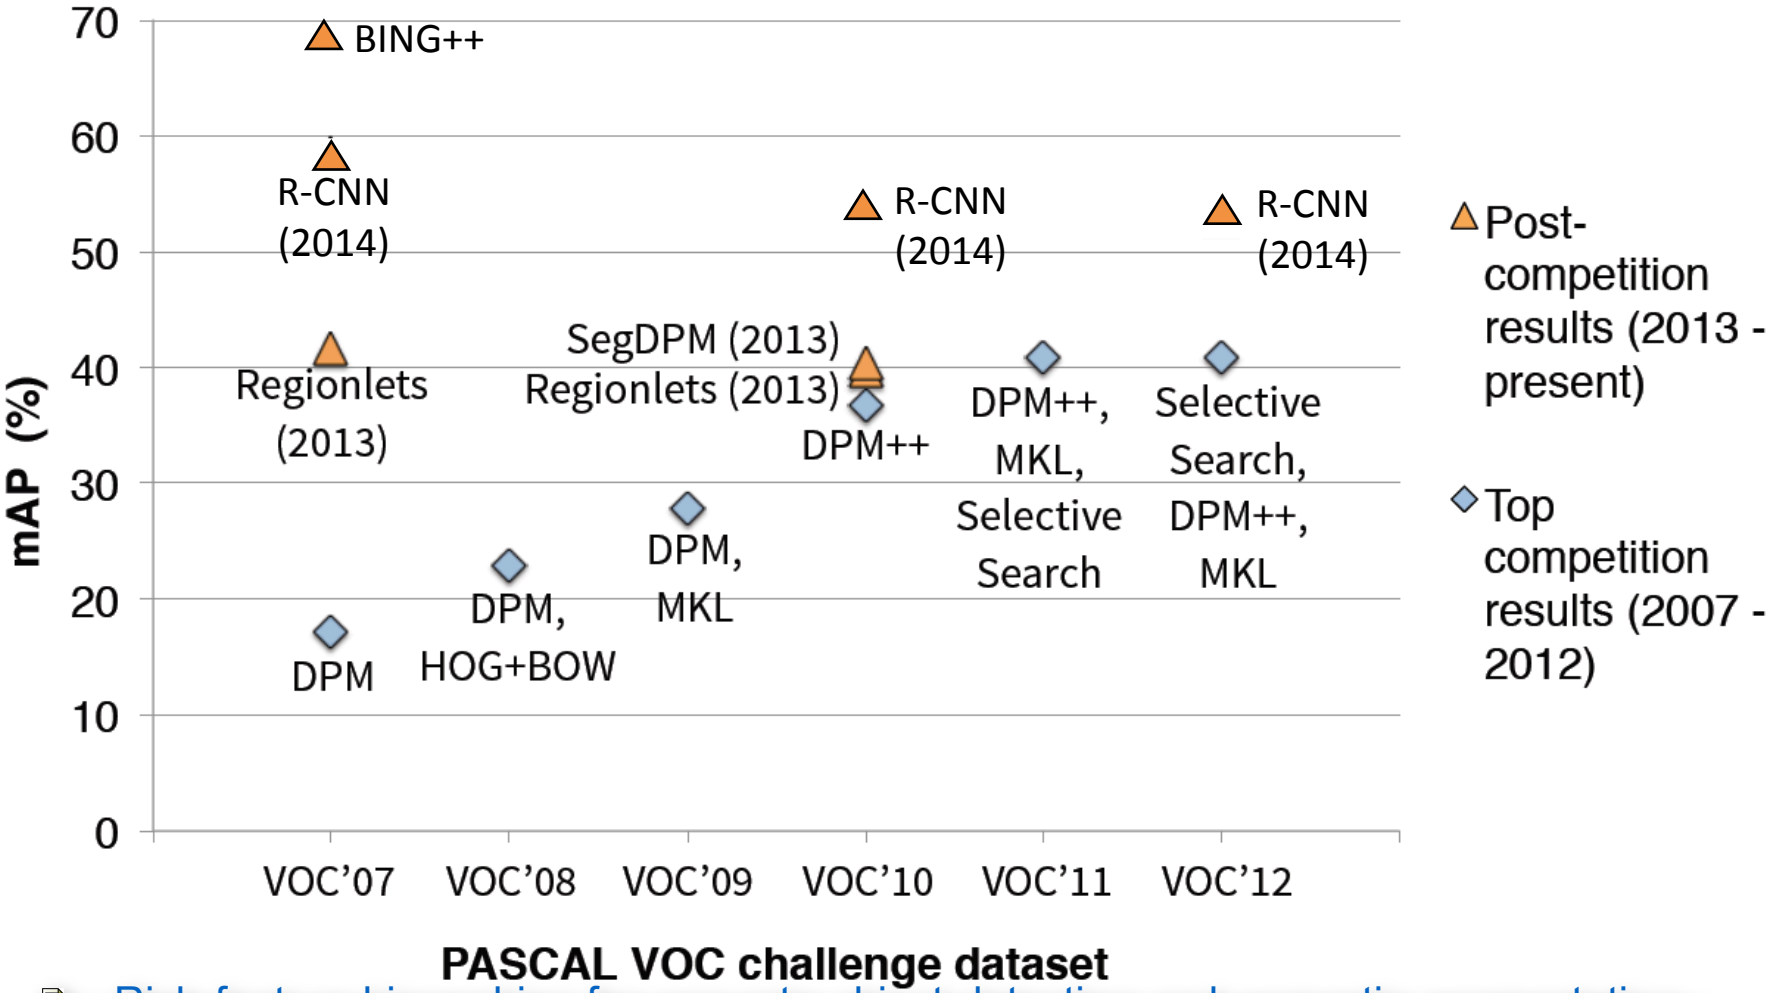

- Rich feature hierarchies for accurate object detection and semantic segmentation, CVPR 2014 (Oral), Girshick et al.
- BING++: A Fast High Quality Object Proposal Generator at 100fps, arXiv, Zhang et al.

Future work in objectness proposals

- High detection rate under large IoU
- Running speed
- Small number of proposals
- Exploring more applications

Other research problems

- New research topics lies in the requirements of applications

Primary object auto-segmentation

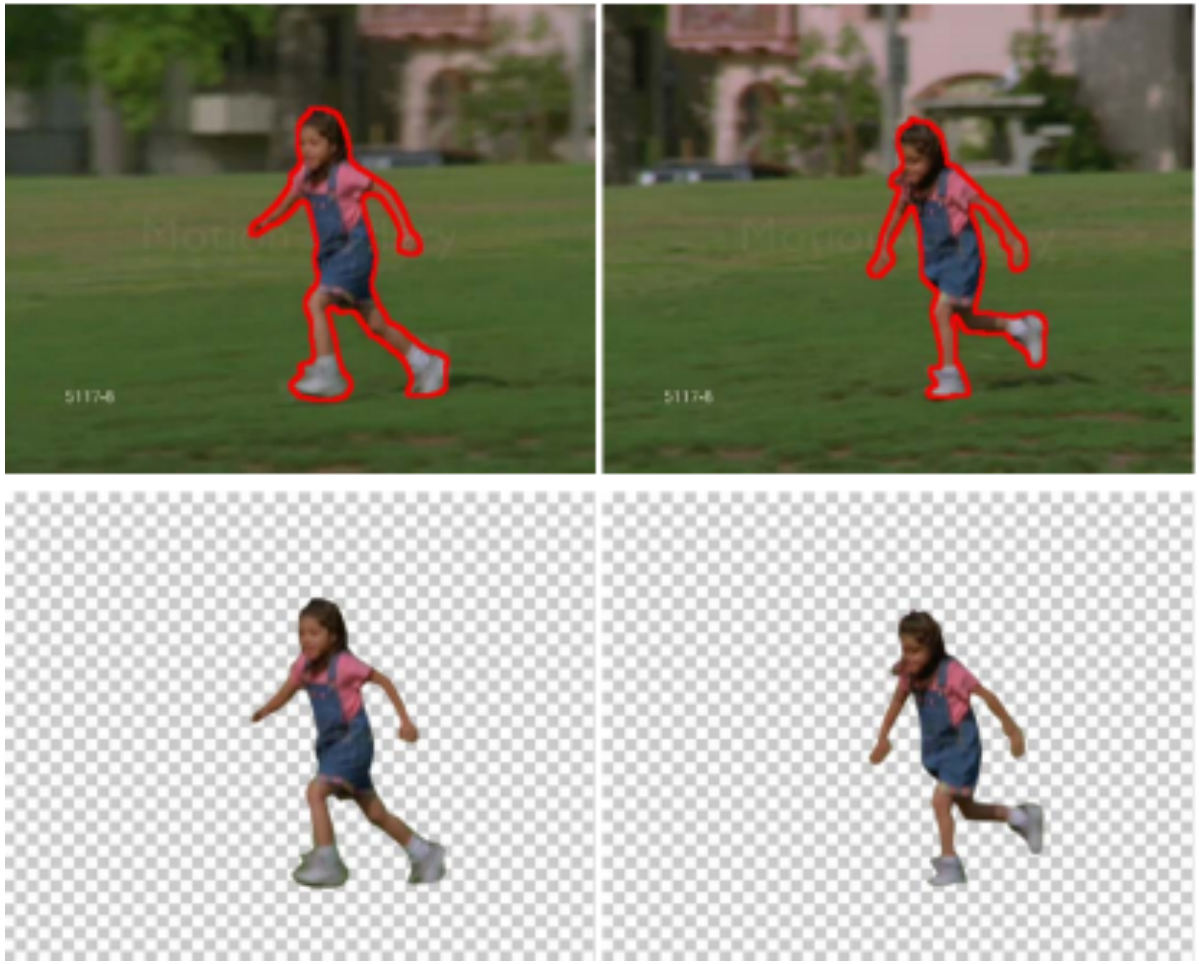

- Fast Appearance Modeling for Automatic Primary Video Object Segmentation, IEEE TIP 2016, Yang et al.

Finding Distractors

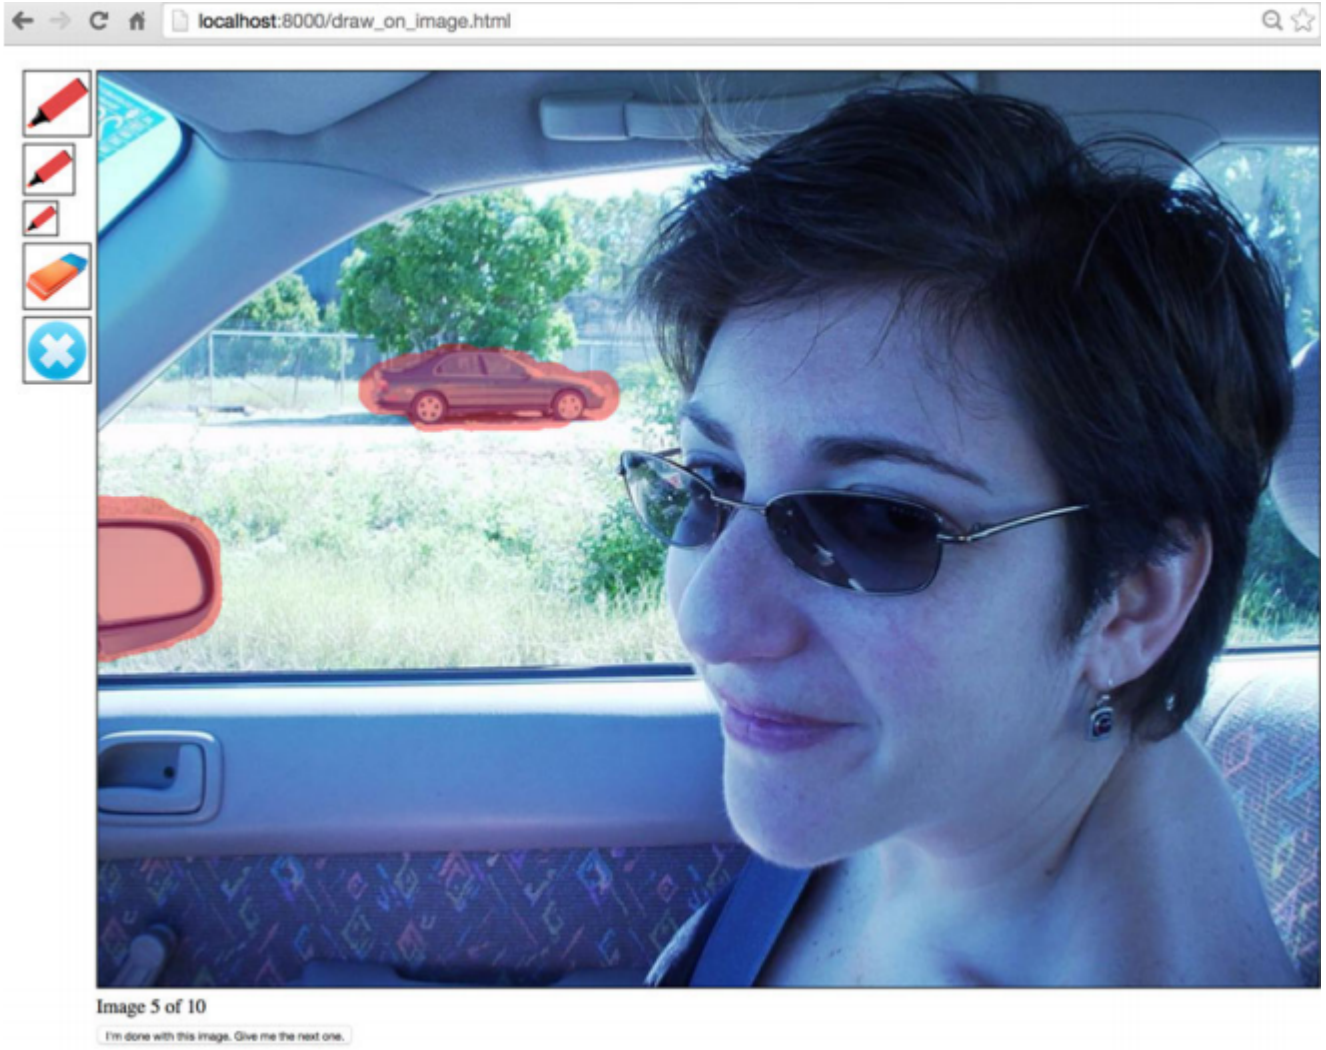

Finding Distractors In Images, IEEE CVPR 2015, Fried et al.

Finding Distractors

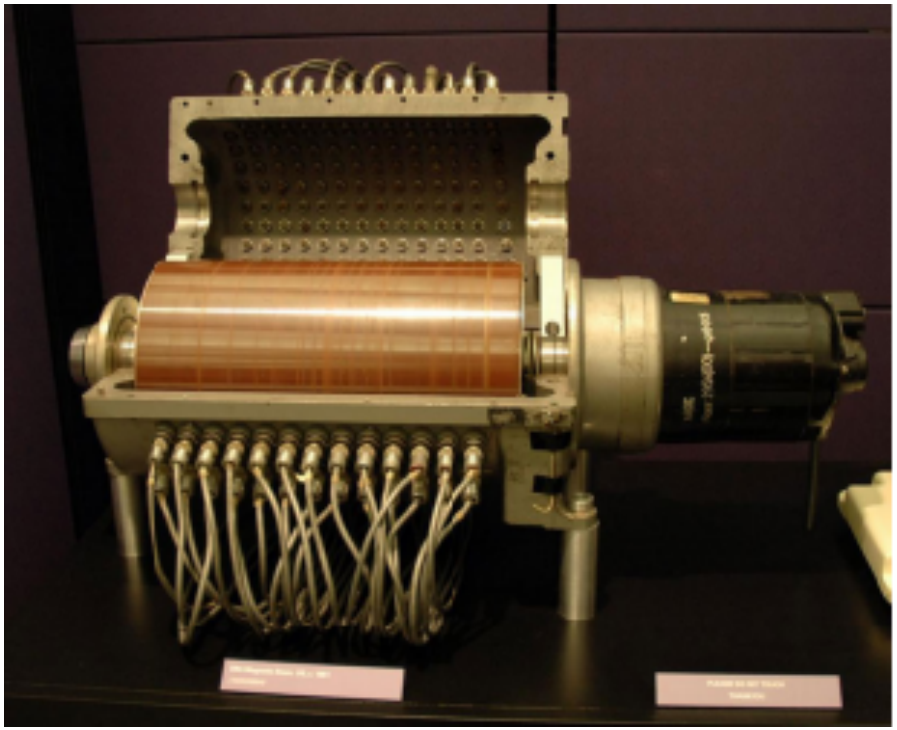

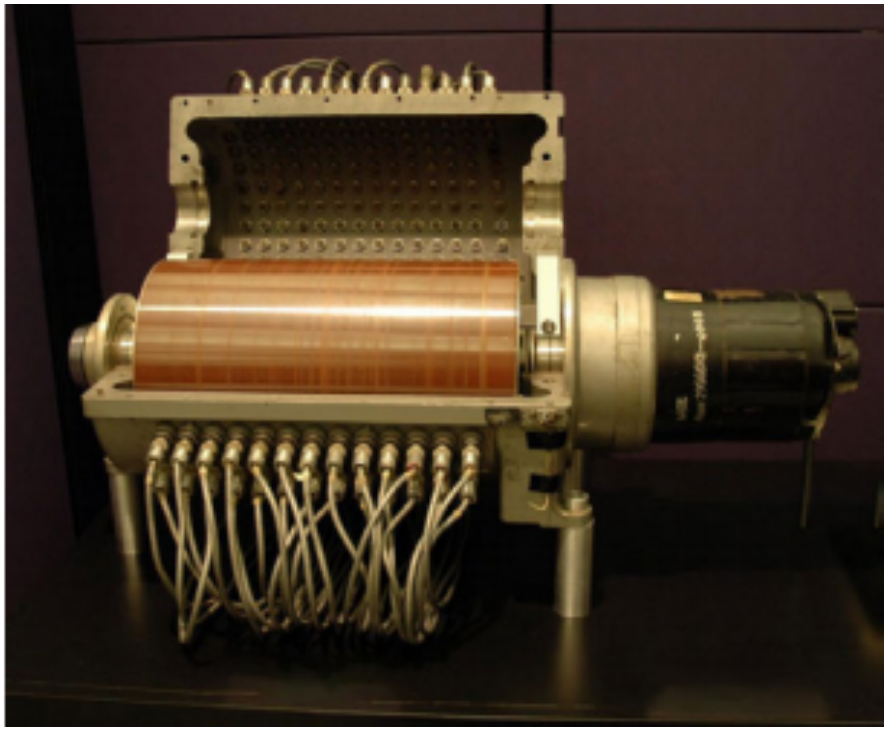

Finding Distractors In Images, IEEE CVPR 2015, Fried et al.

Thanks!

Tutorial webpage:
http://saliency.mit.edu/ECCVTutorial/ECCV_saliency.htm