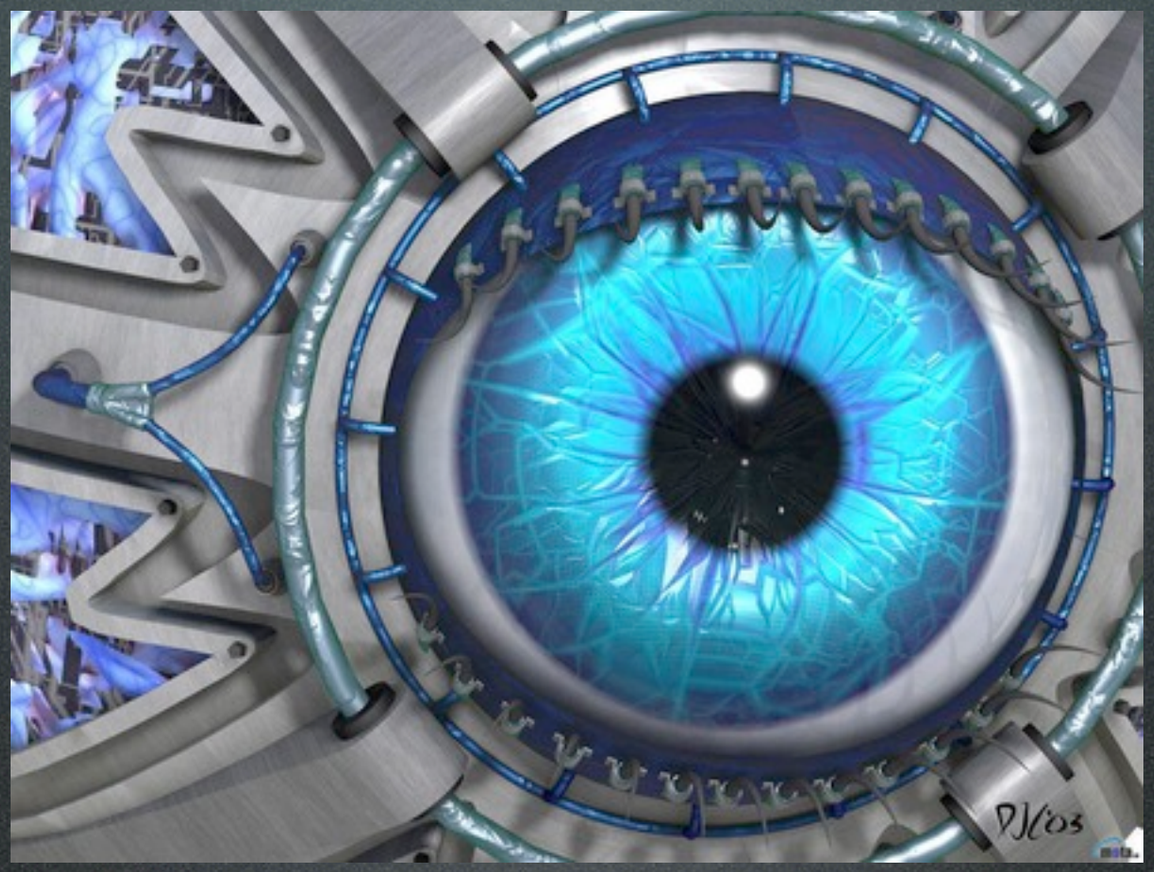

Computer Vision

Working to help machines see...

What do you do in computer vision?

- In computer vision, we try to make machines/robots "see" like humans

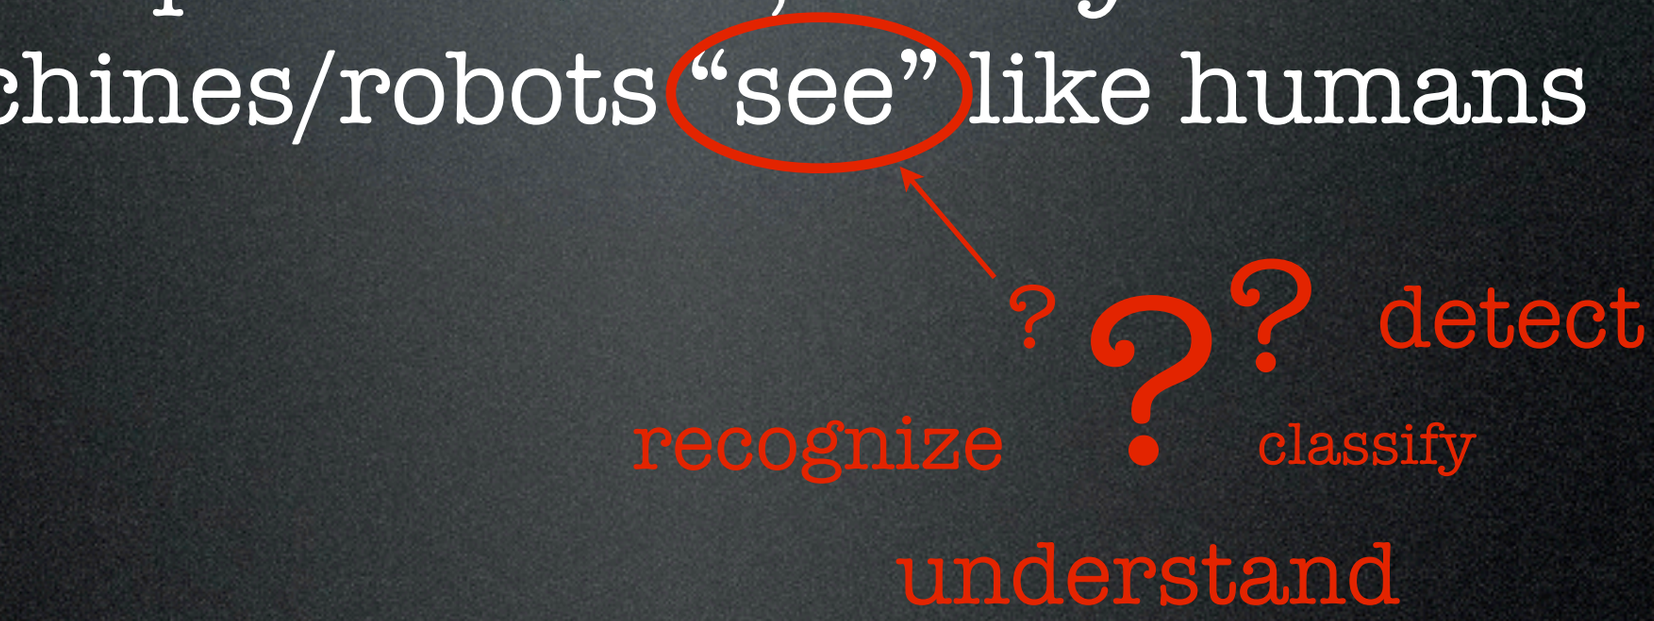

You know what one great vision system is?

YOU

Ready to use your
visual wisdom and worldly knowledge
to help a robot “see” like you
?

Challenge 1: help a robot recognize a chair

Materials? Colors?

Parts? Shape?

What's in common?

What's in common?

Stand on ground?


Hopeless!

Forget the properties,
describe the affordances instead

what can a chair
be used for?

Surface, at a certain height, used for sitting...

Surface, at a certain height, used for sitting...

Surface, at a certain height, used for sitting...

Hopeless!

Challenge 2: help a robot understand a scene

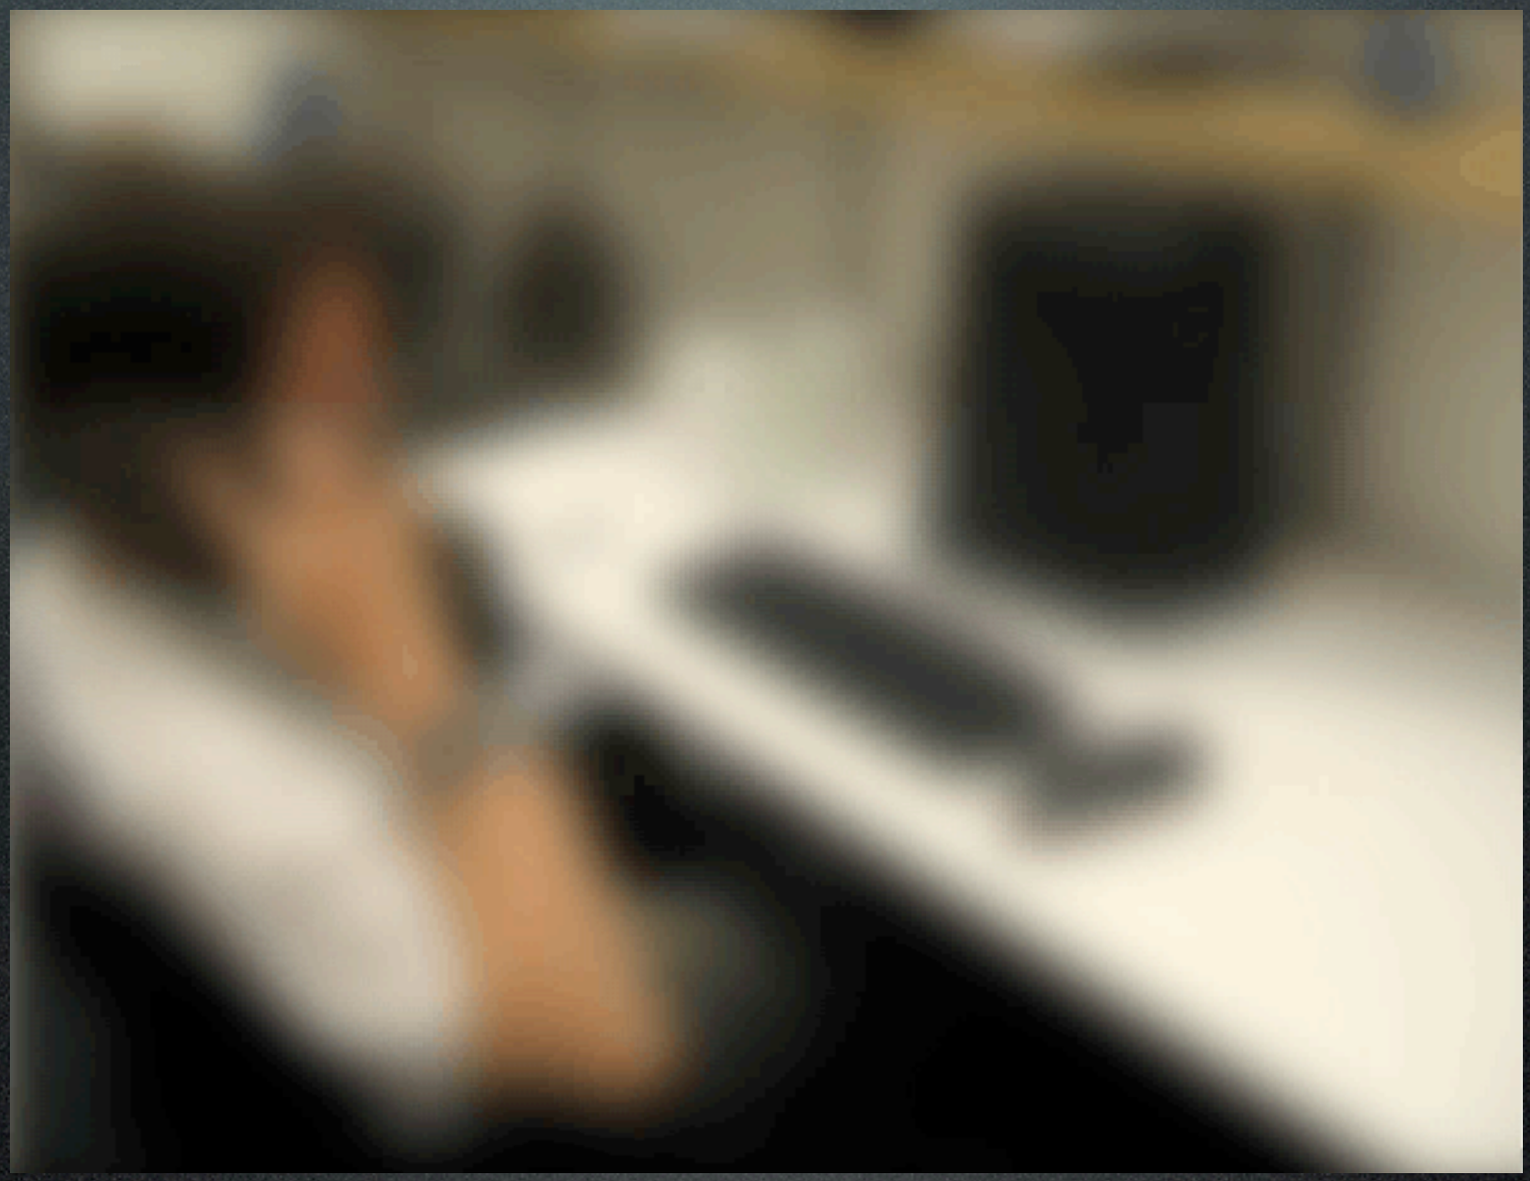

Joan Steiner

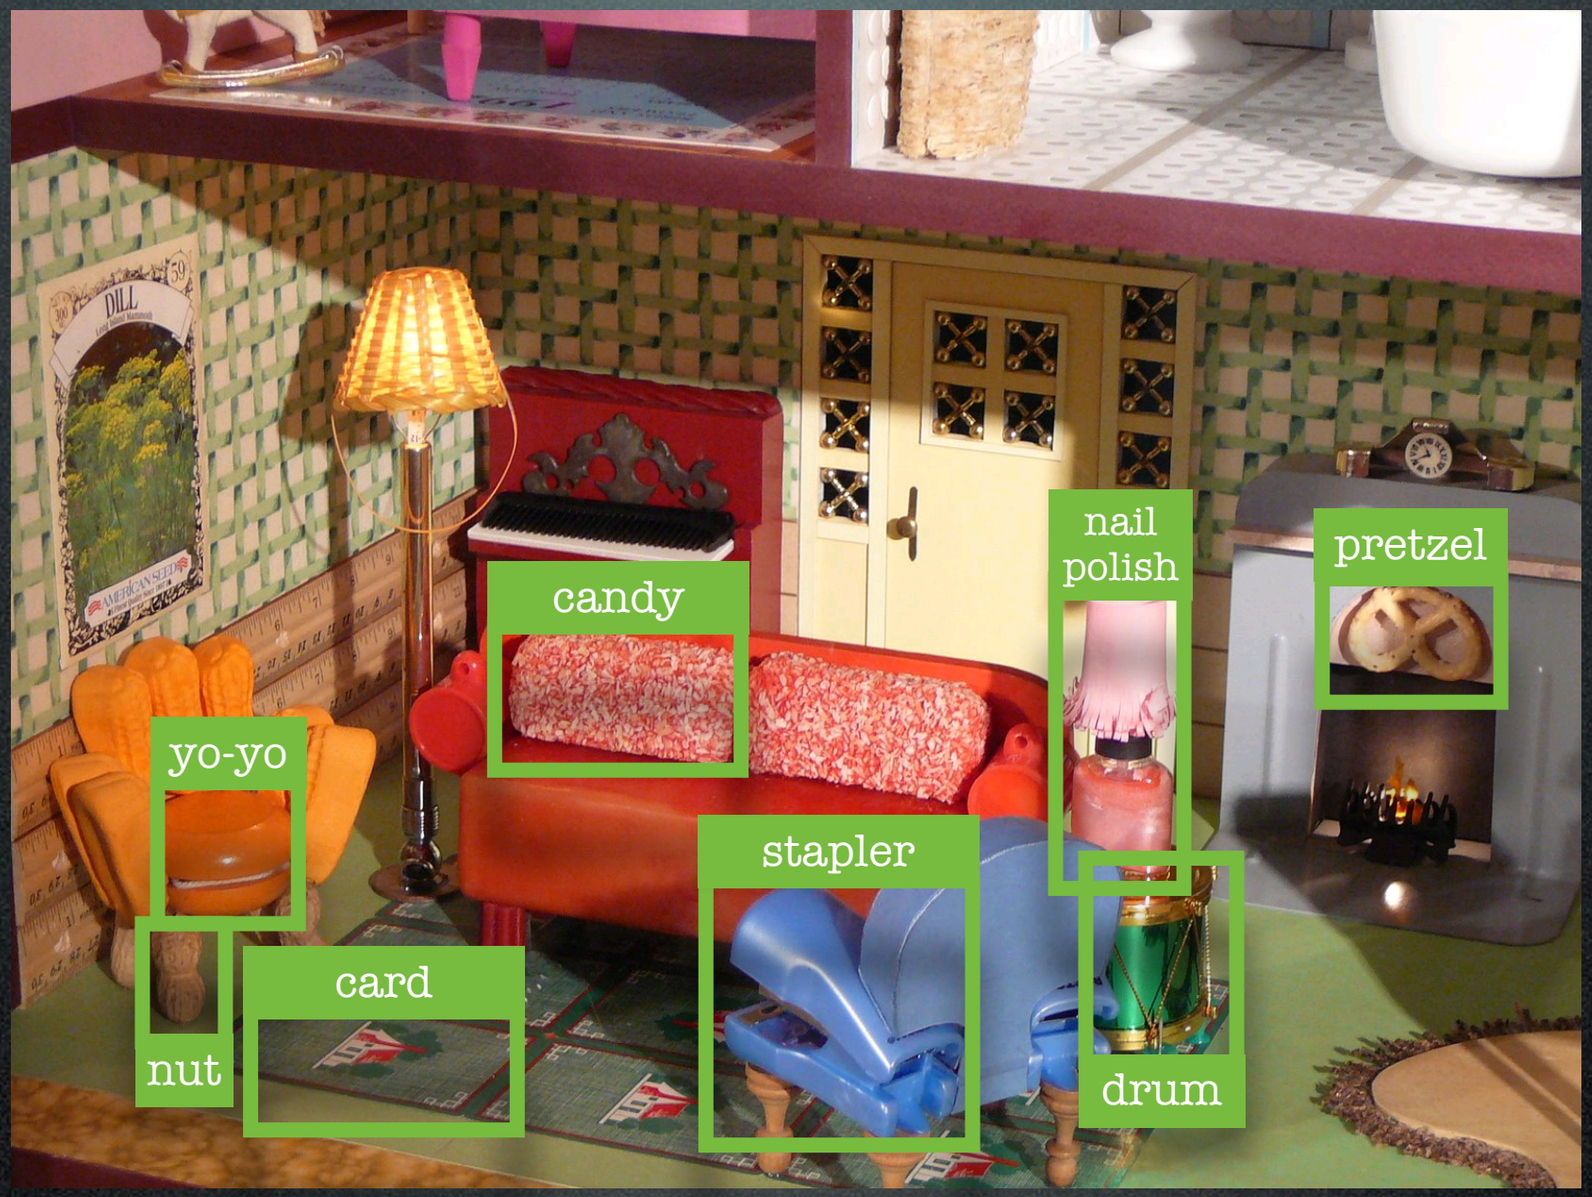

Joan Steiner

Is the identity of objects context-dependent?

Joan Steiner

Challenge 3: help a robot find animals

Find Wald...animals

Some of the challenges in computer vision

- classification/categorization

"what is it?"
(e.g. recognize a chair)

- detection

"find it"
(e.g. find animals in picture)

- scene/image understanding

"what is happening here?"
(e.g. understand a scene, describe the elements, understand the relationships and actions)

- ...

- reconstruct a 3D scene from 2D images
- recognize the emotion on the face
- split the images into its components
- etc.

But it's not so hopeless after all!

How do you do it?

Let's recognize a chair:

input: image

feature extraction

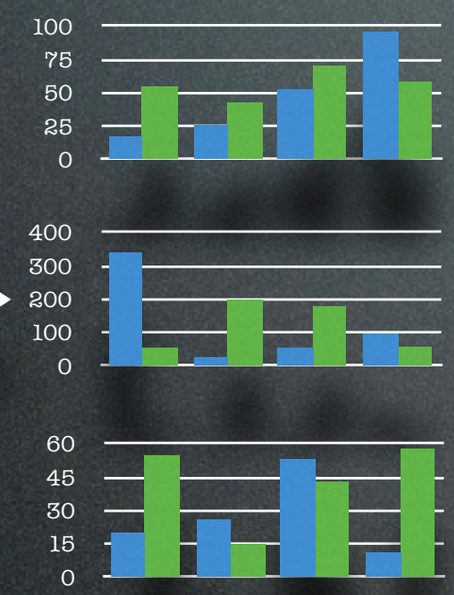

a bit of magic

output: label

a typical computer vision pipeline

Let's recognize a chair:

input: image

feature extraction

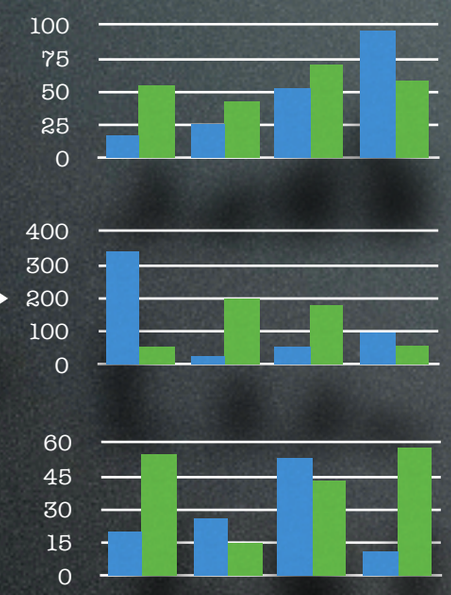

a bit of magic

output: label

+ ASSUMPTIONS

As you saw before, the general computer vision task is too hard

So, let's assume our task is not to recognize chairs in the general sense, but to tell apart wooden chairs from glass wine bottles

Class discrimination

can we find a simple way
to tell them apart?

Feature extraction example: color

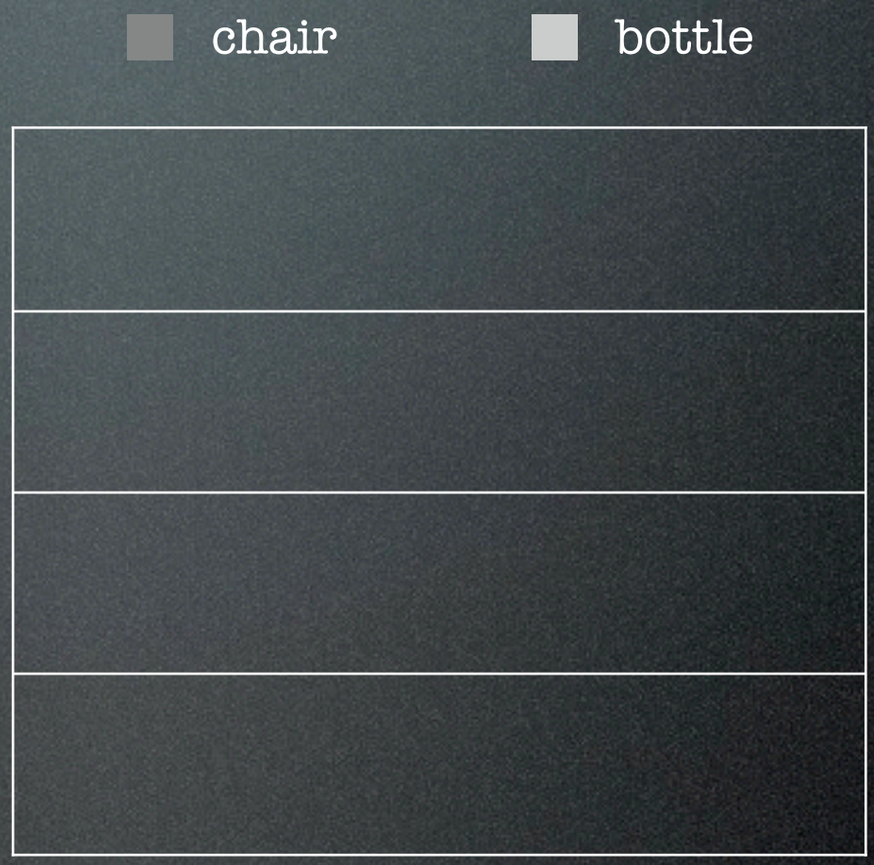

Feature extraction example: color

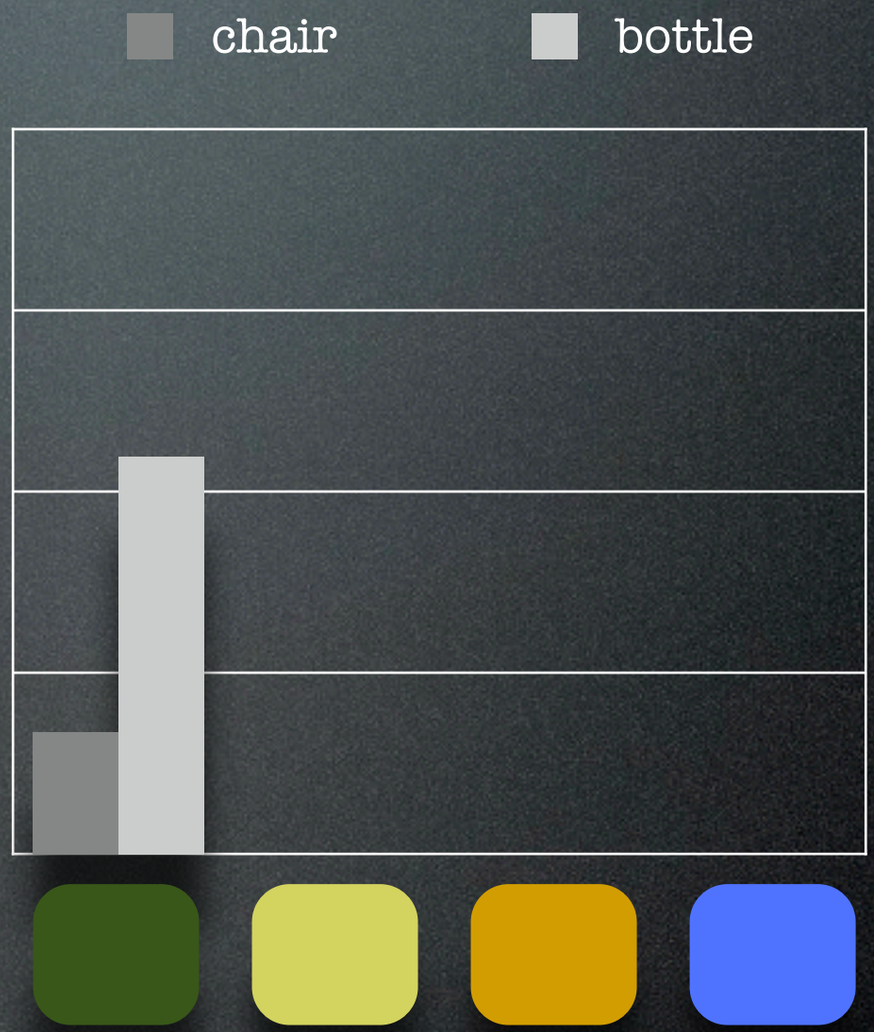

Feature extraction example: color

Feature extraction example: color

Feature extraction example: color

Collect more data

wooden chairs

glass wine bottles

Compute features

wooden chairs



glass wine bottles



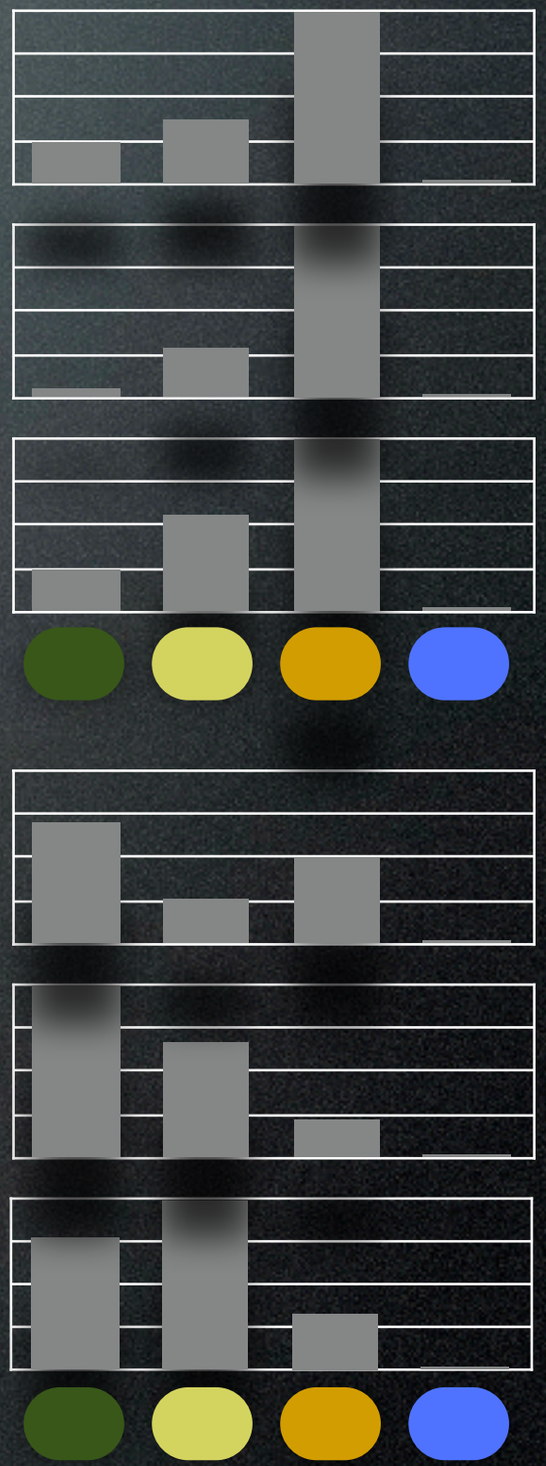

Discover discriminative features

wooden chairs

glass wine bottles

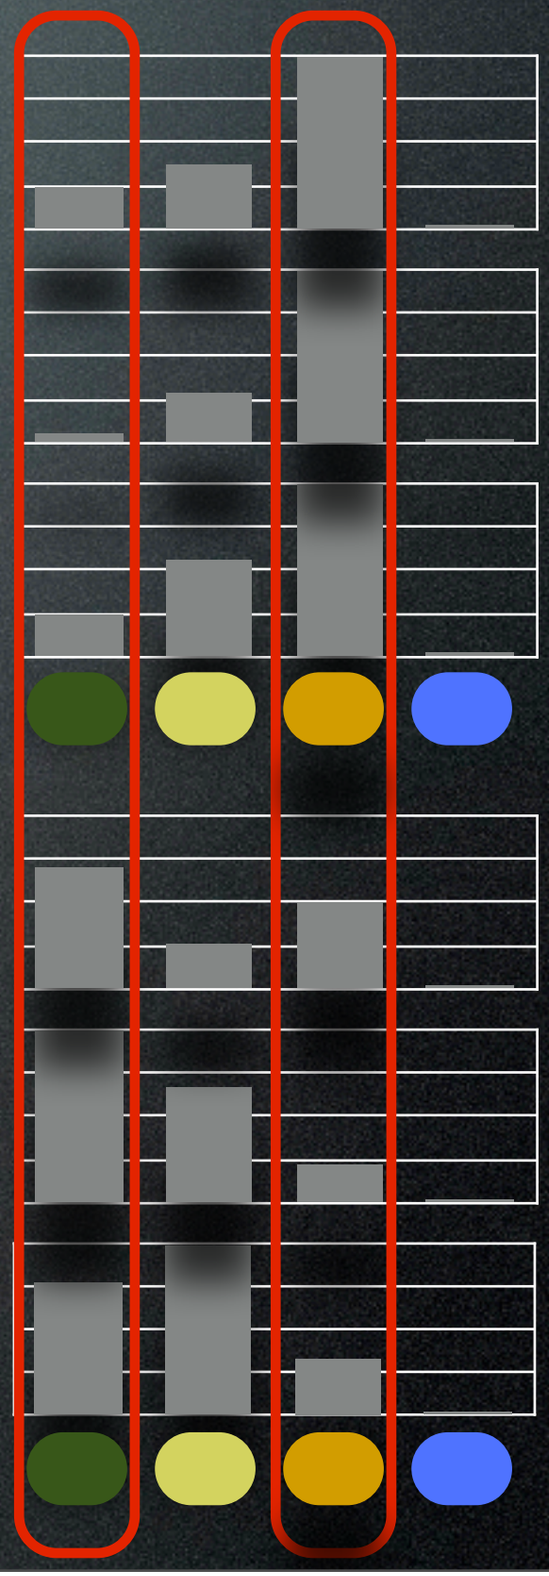

Peek inside the black box:

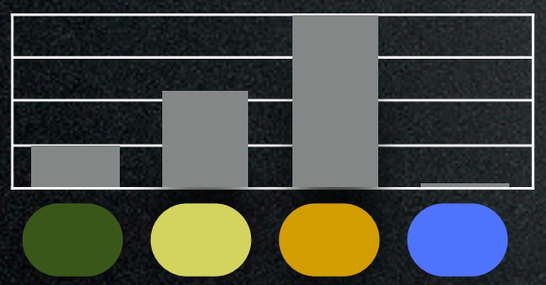

IF more → then RETURN bottle

ELSE more → then RETURN chair

Let's test our system...

Let's test our system...

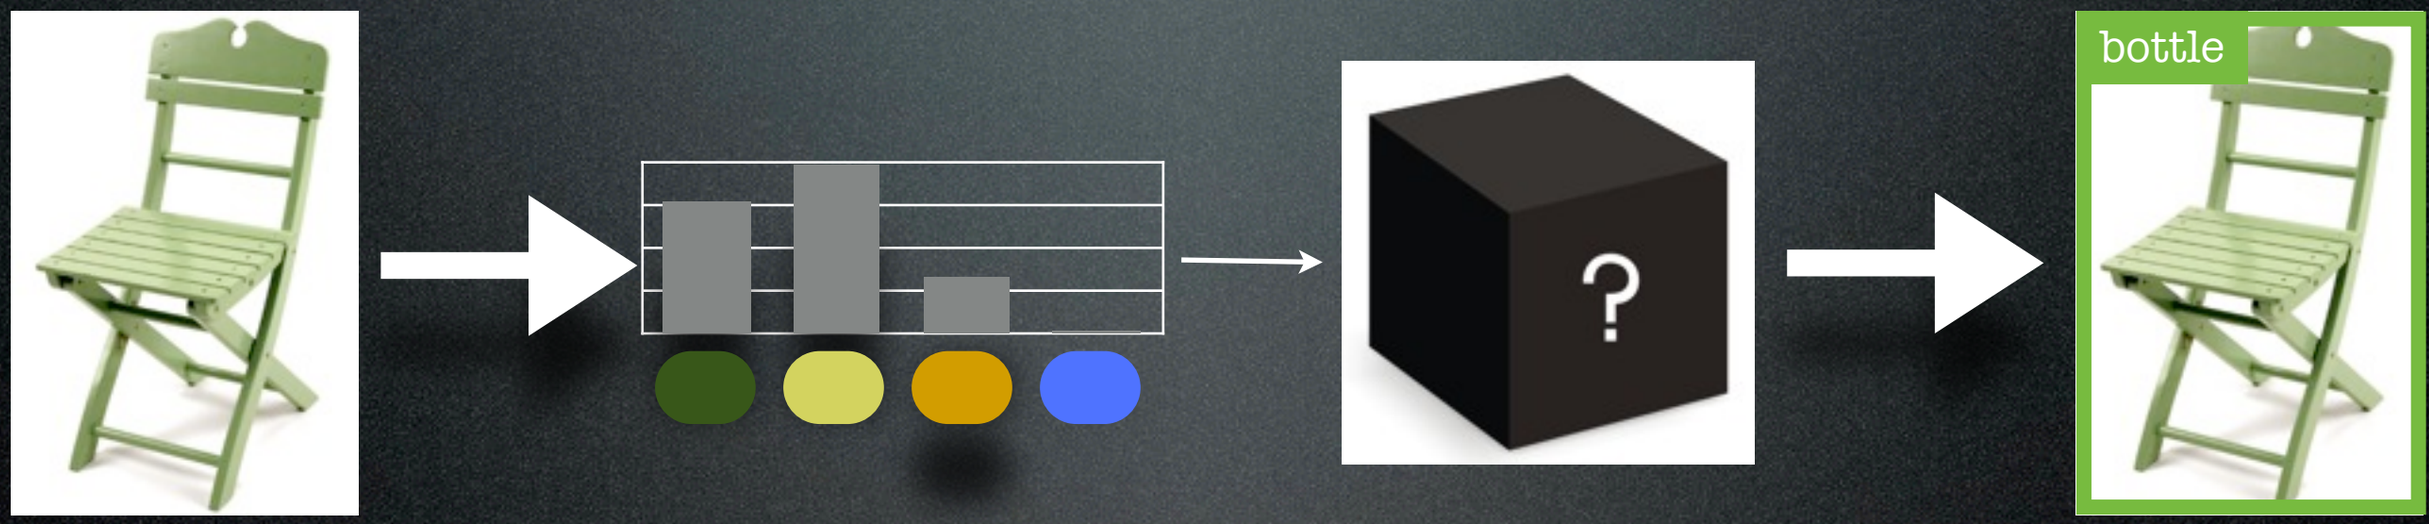

Back to the drawing board...

More features!

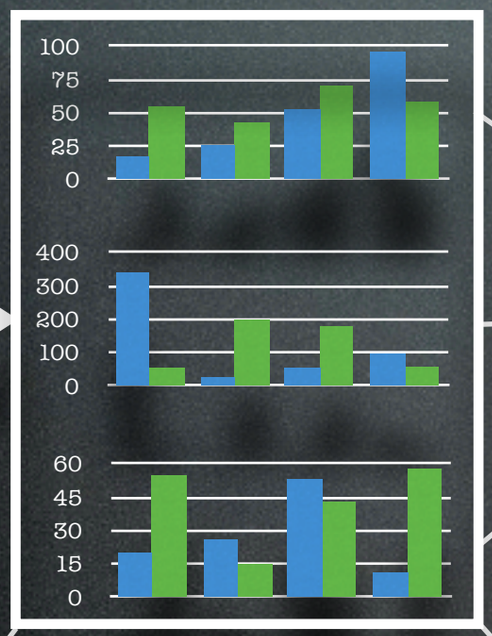

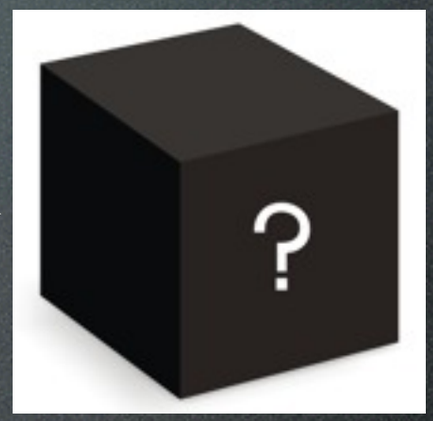

materials
textures
gradients
corners
surfaces
...
...

More data!

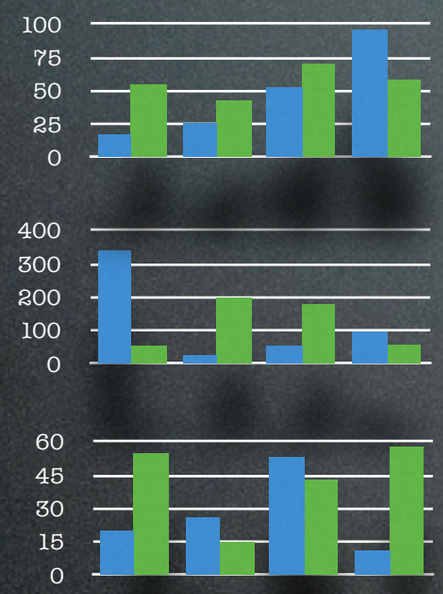

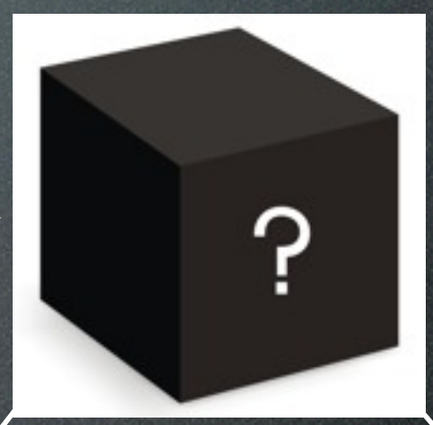

Computer vision is hard but possible!

Let's look at some successes...

Face detection

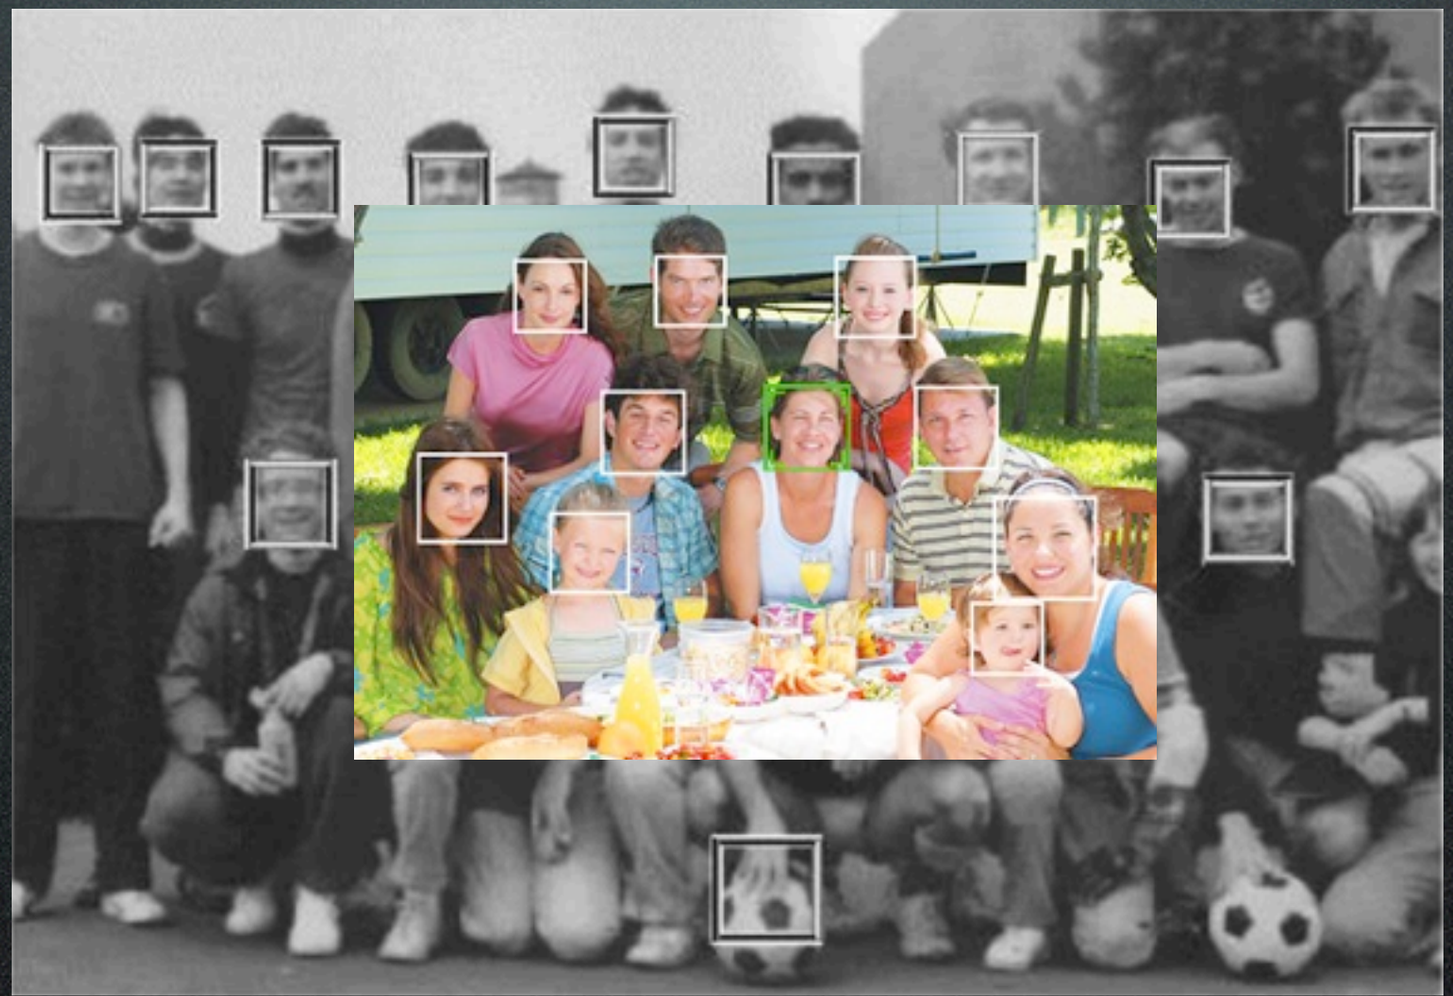

Viola & Jones (2001)

Body part recognition

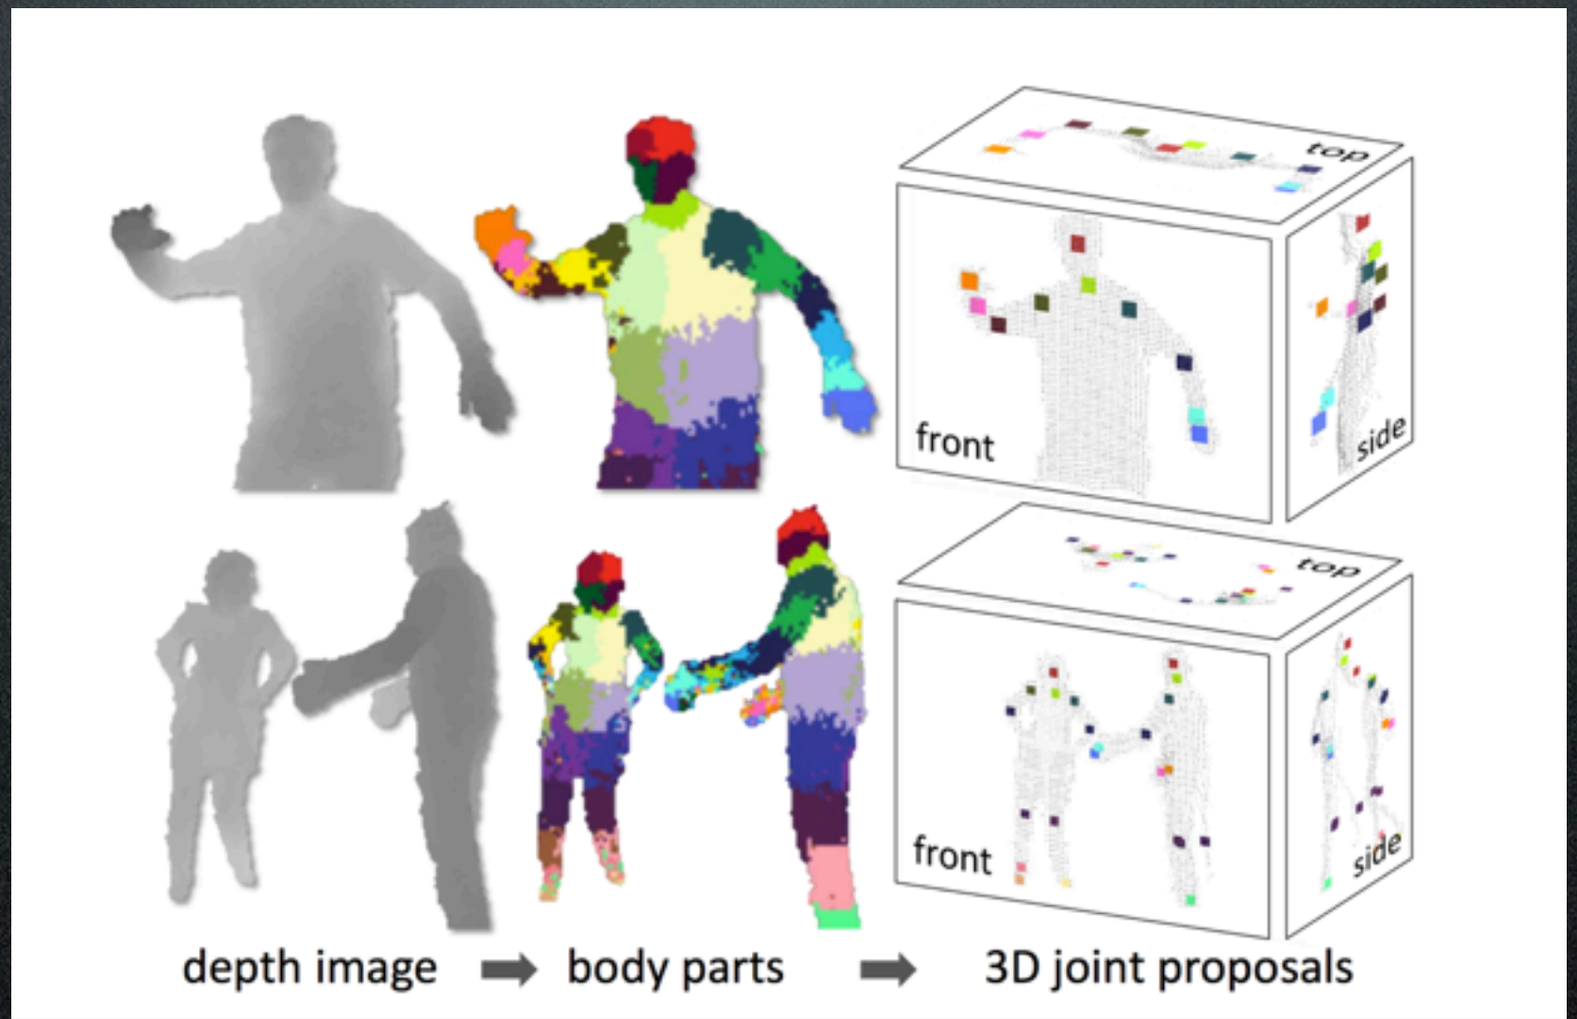

Microsoft's Kinect (2011)

Pedestrian Detection

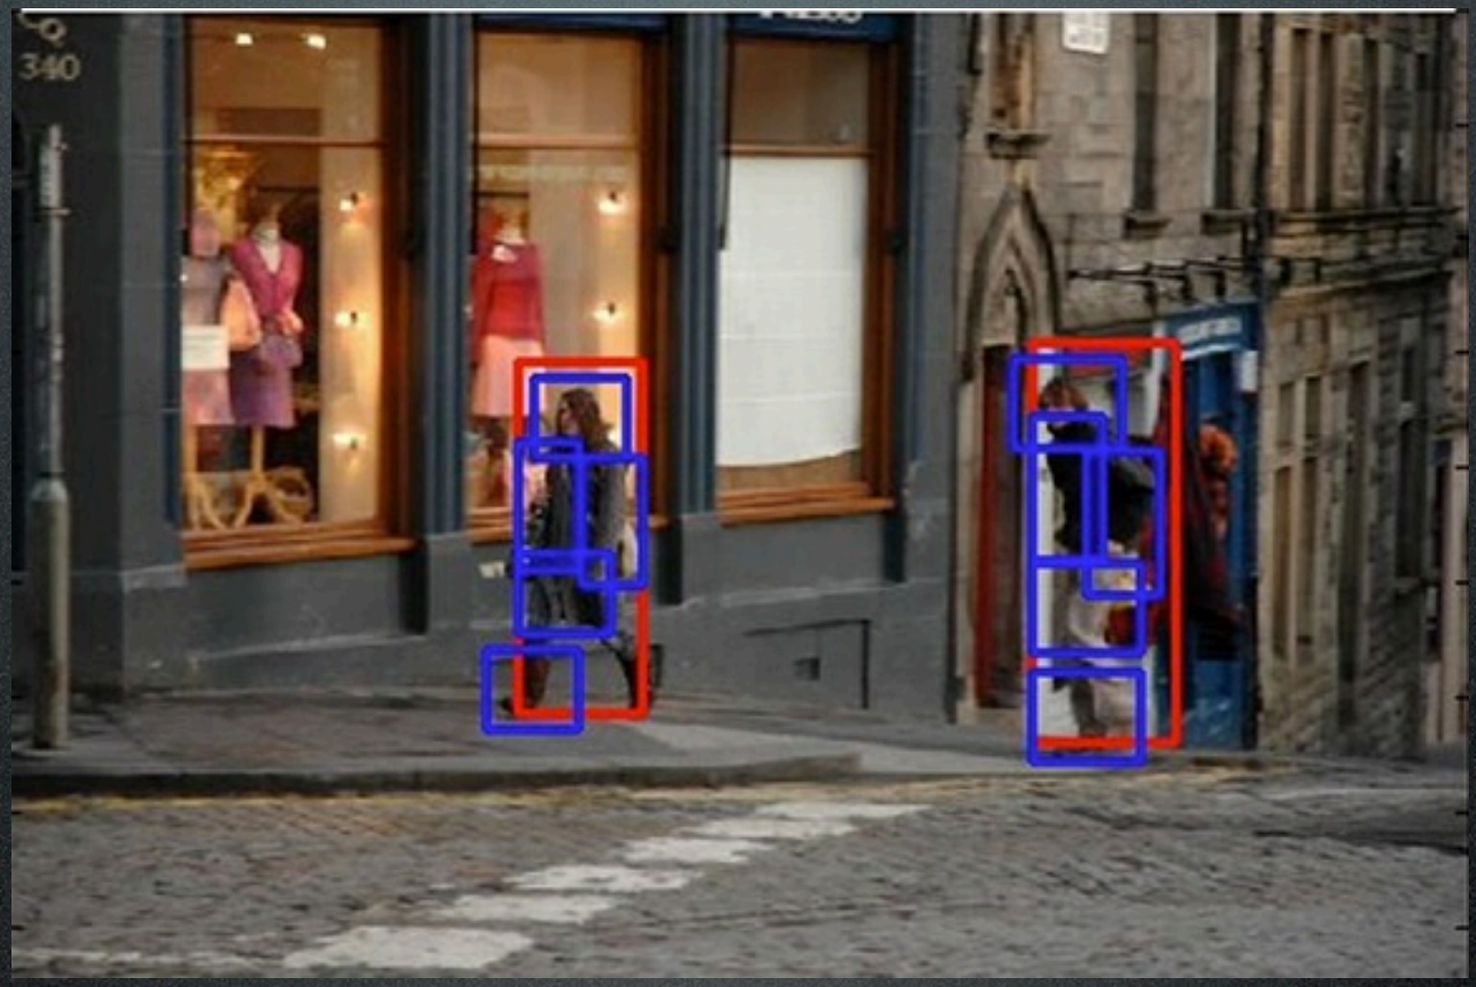

Felzenszwalb, McAllester, & Ramanan (2008)

Scene Parsing

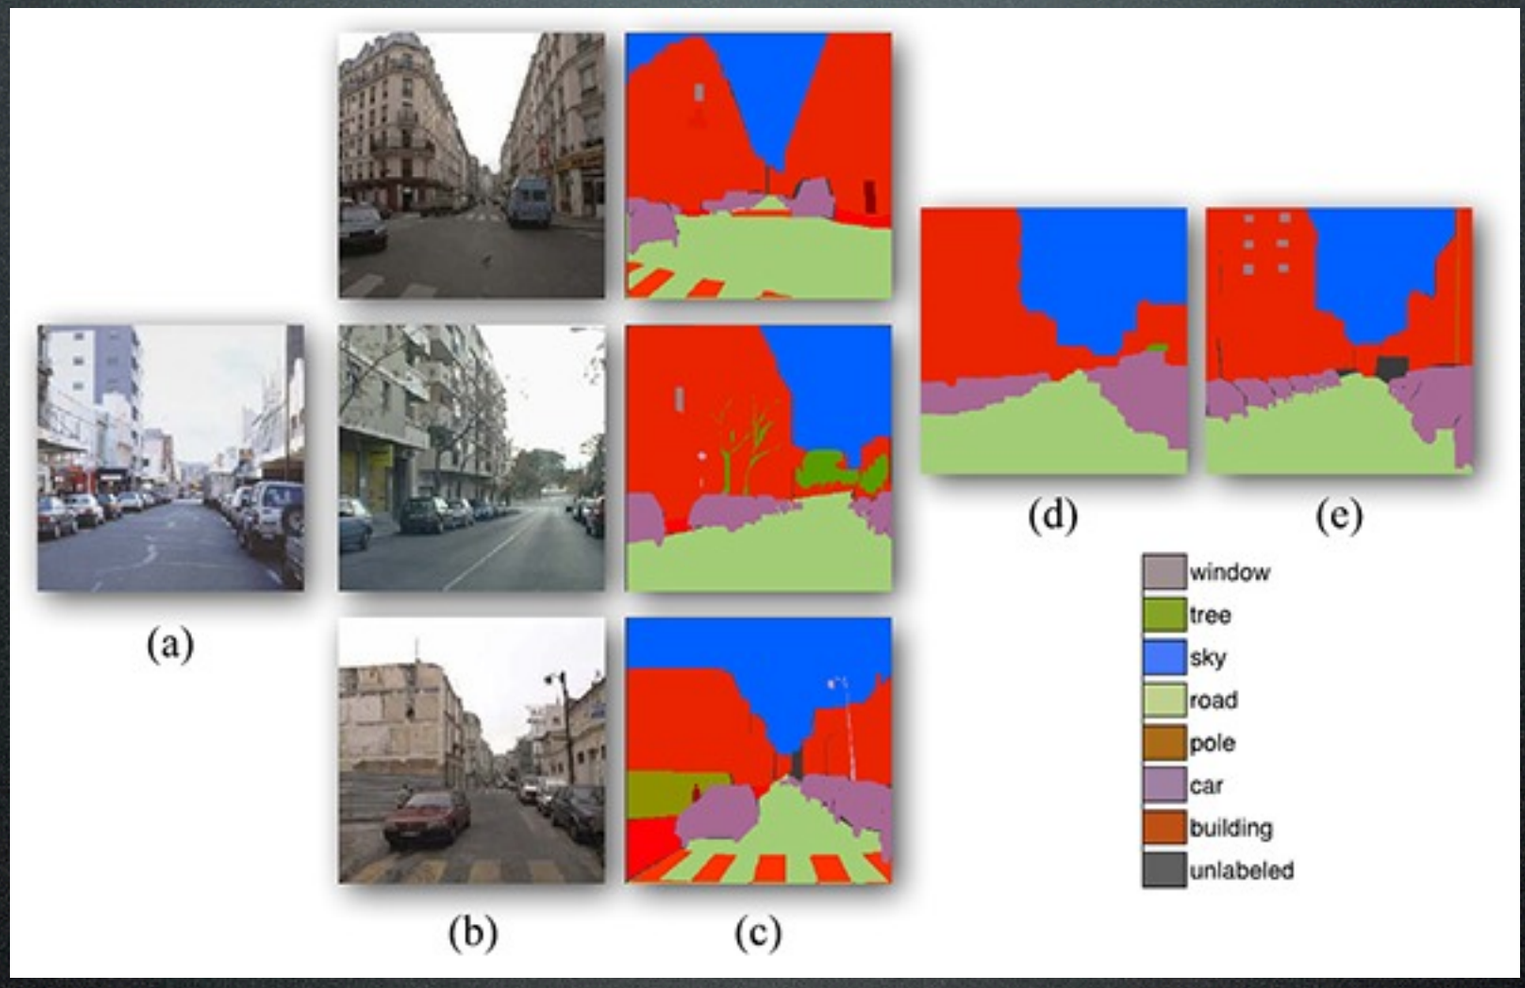

Liu, Yuen, & Torralba (2011)

Image Segmentation

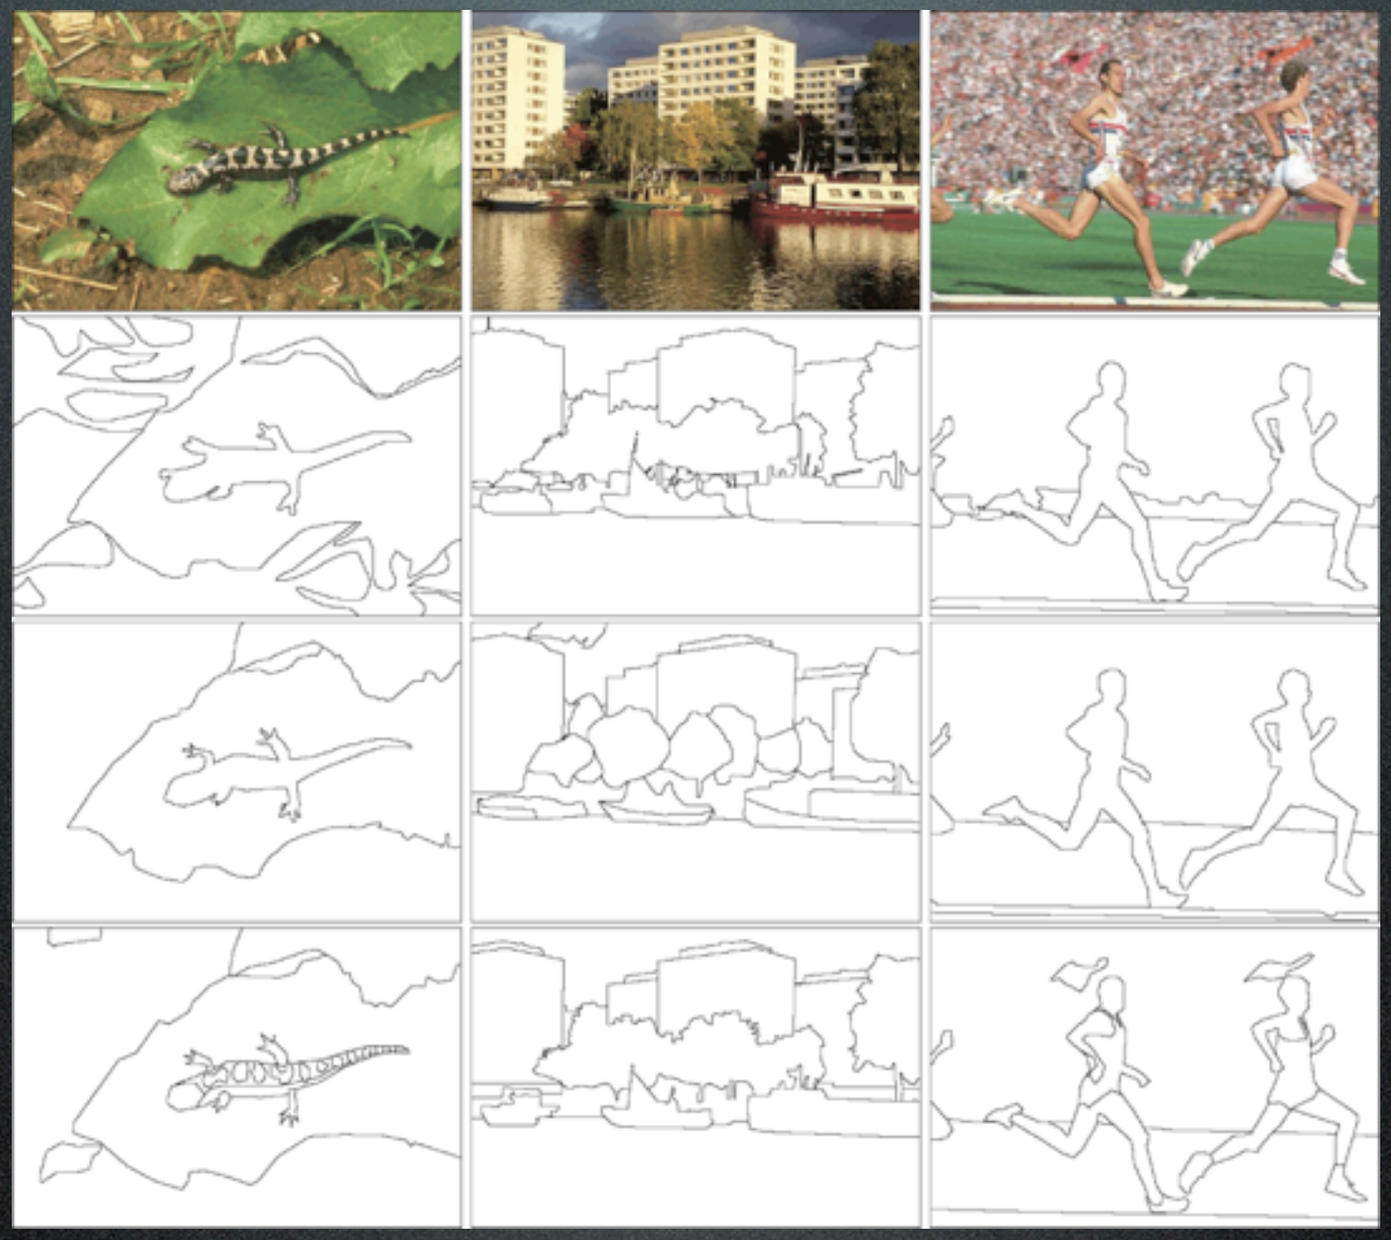

Berkeley Segmentation Dataset (2001)

3D Reconstruction

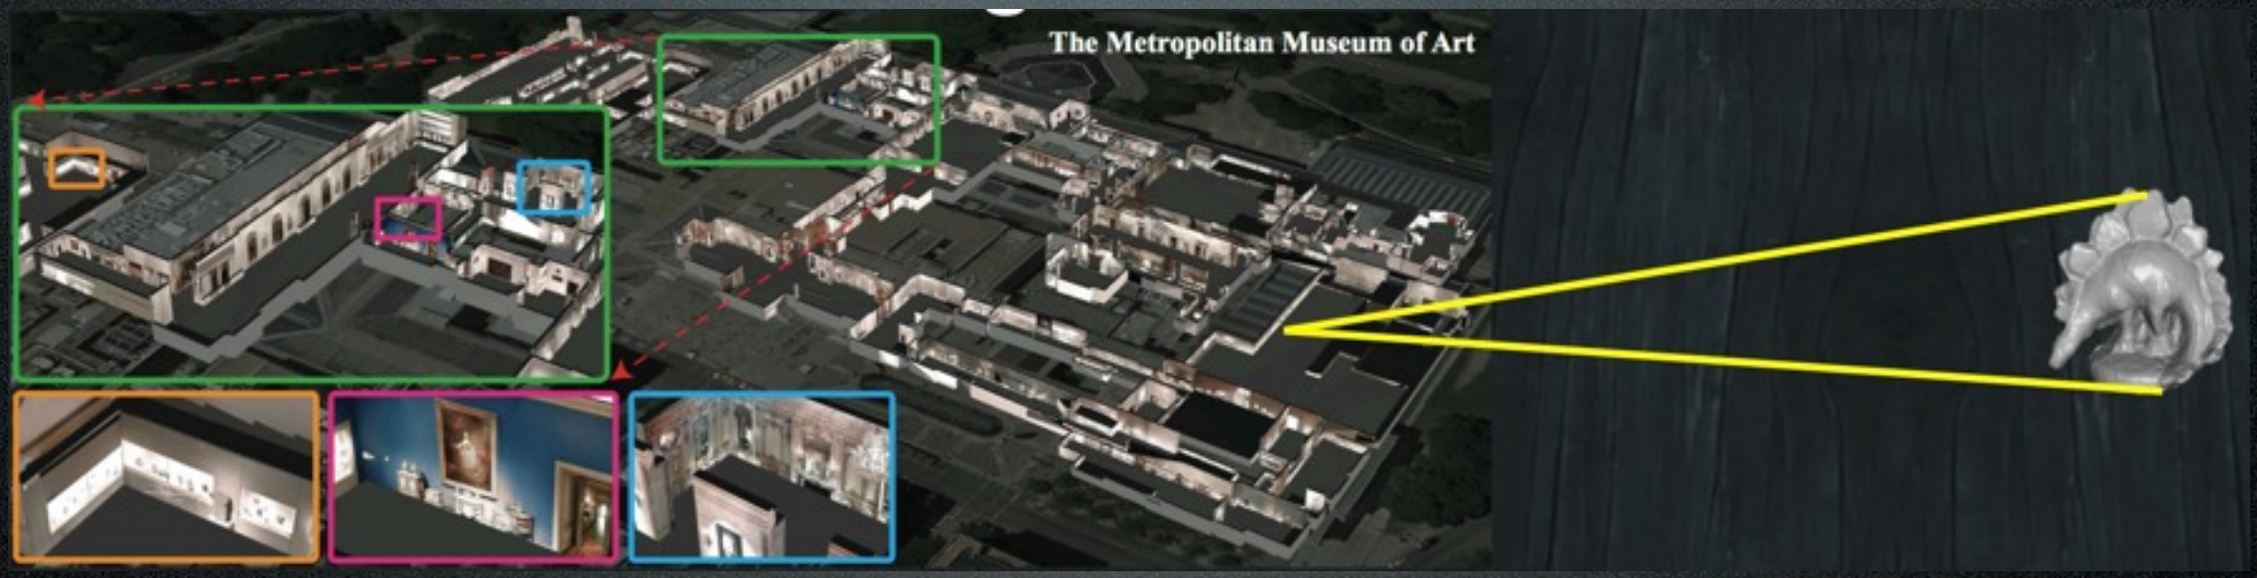

Xiao & Furukawa (2012)

Predicting where people look

Judd, Ehinger, Torralba, & Durand (2009)

Predicting Image Memorability

Isola, Xiao, Torralba, & Oliva (2011)

Coming soon!

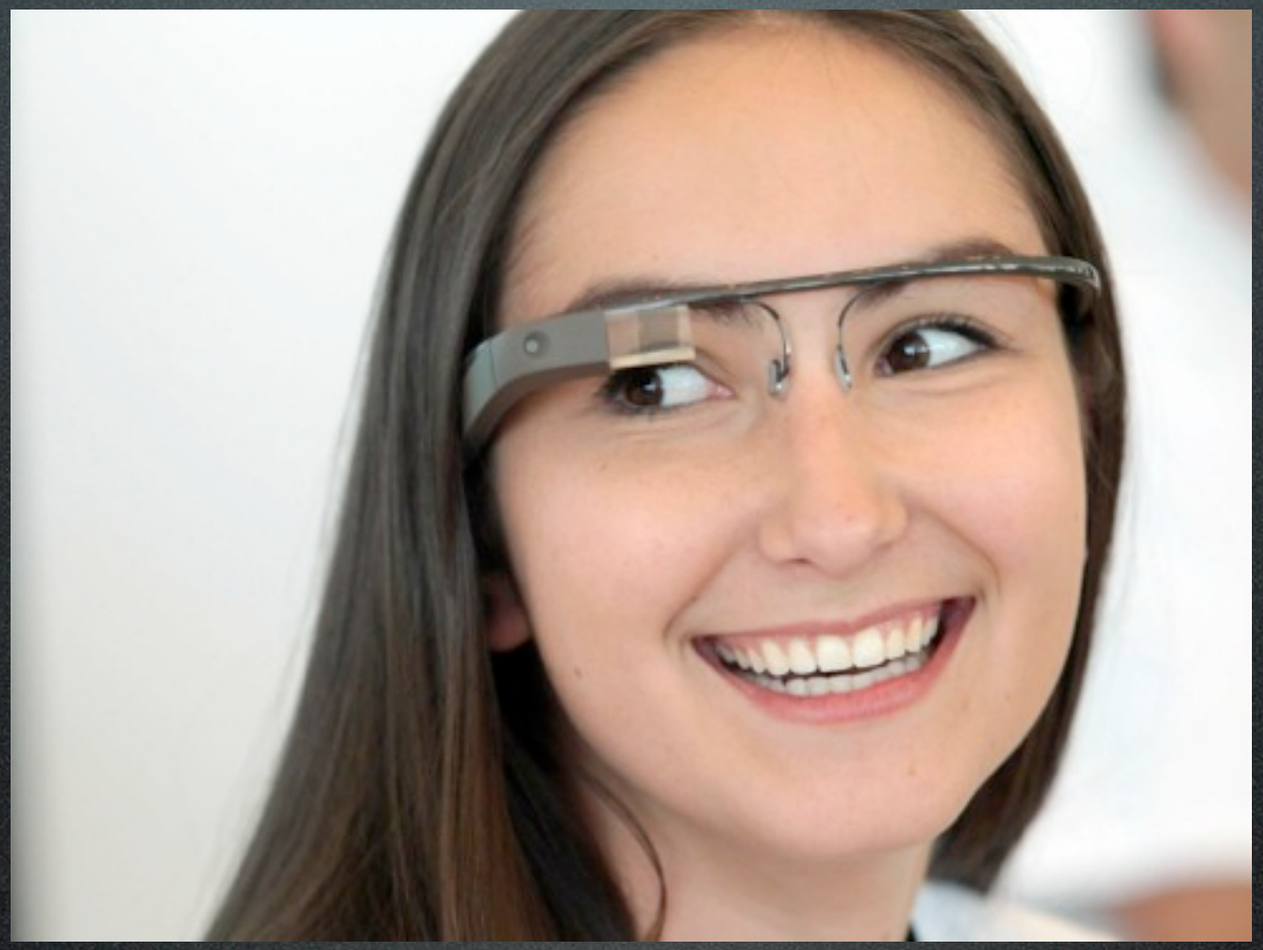

Google Glass

The next great black box:

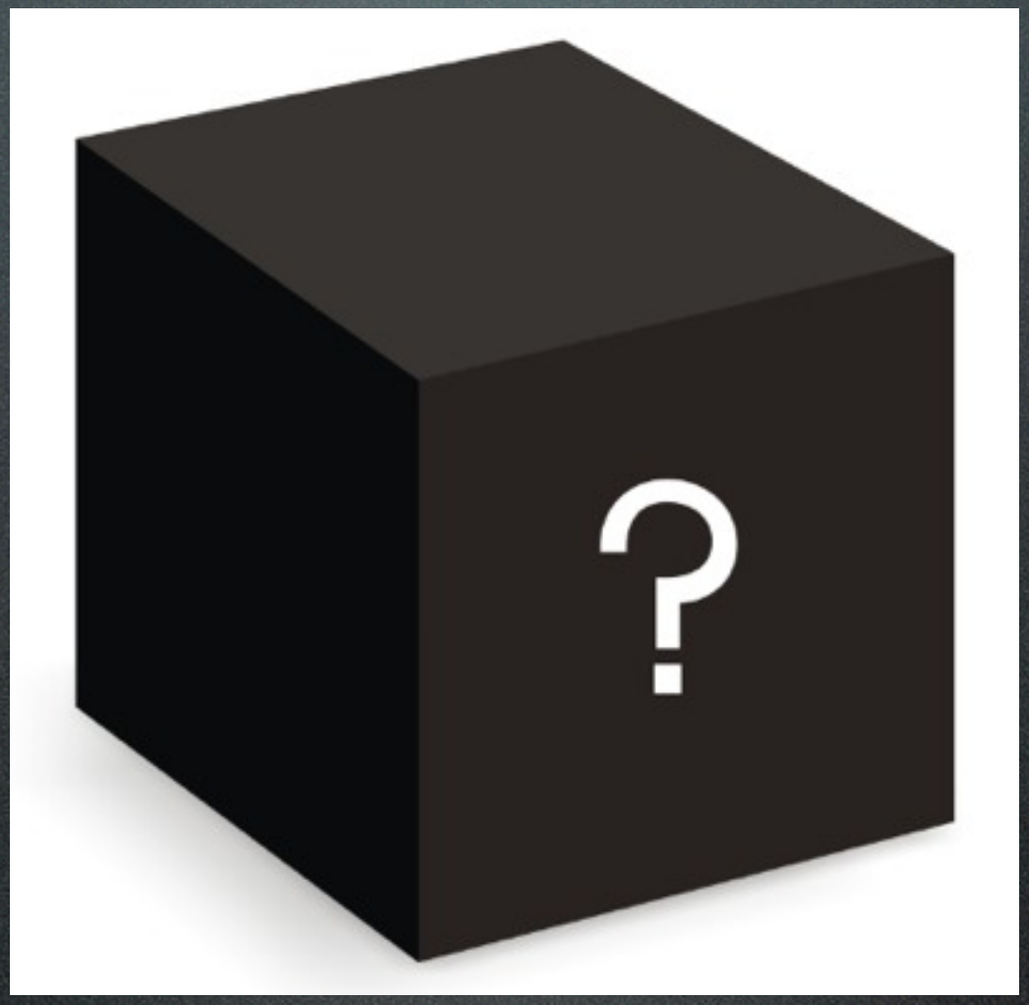

Designed by YOU

THANK YOU :)

THA :)

Questions?