

# Are Categories Necessary?



Alexei (Alyosha) Efros  
CMU

# Before We Begin...


- Not an AI/learning person
- My work is in Computer Vision and Computer Graphics
- But I want to utilize lots of data

# A Confession

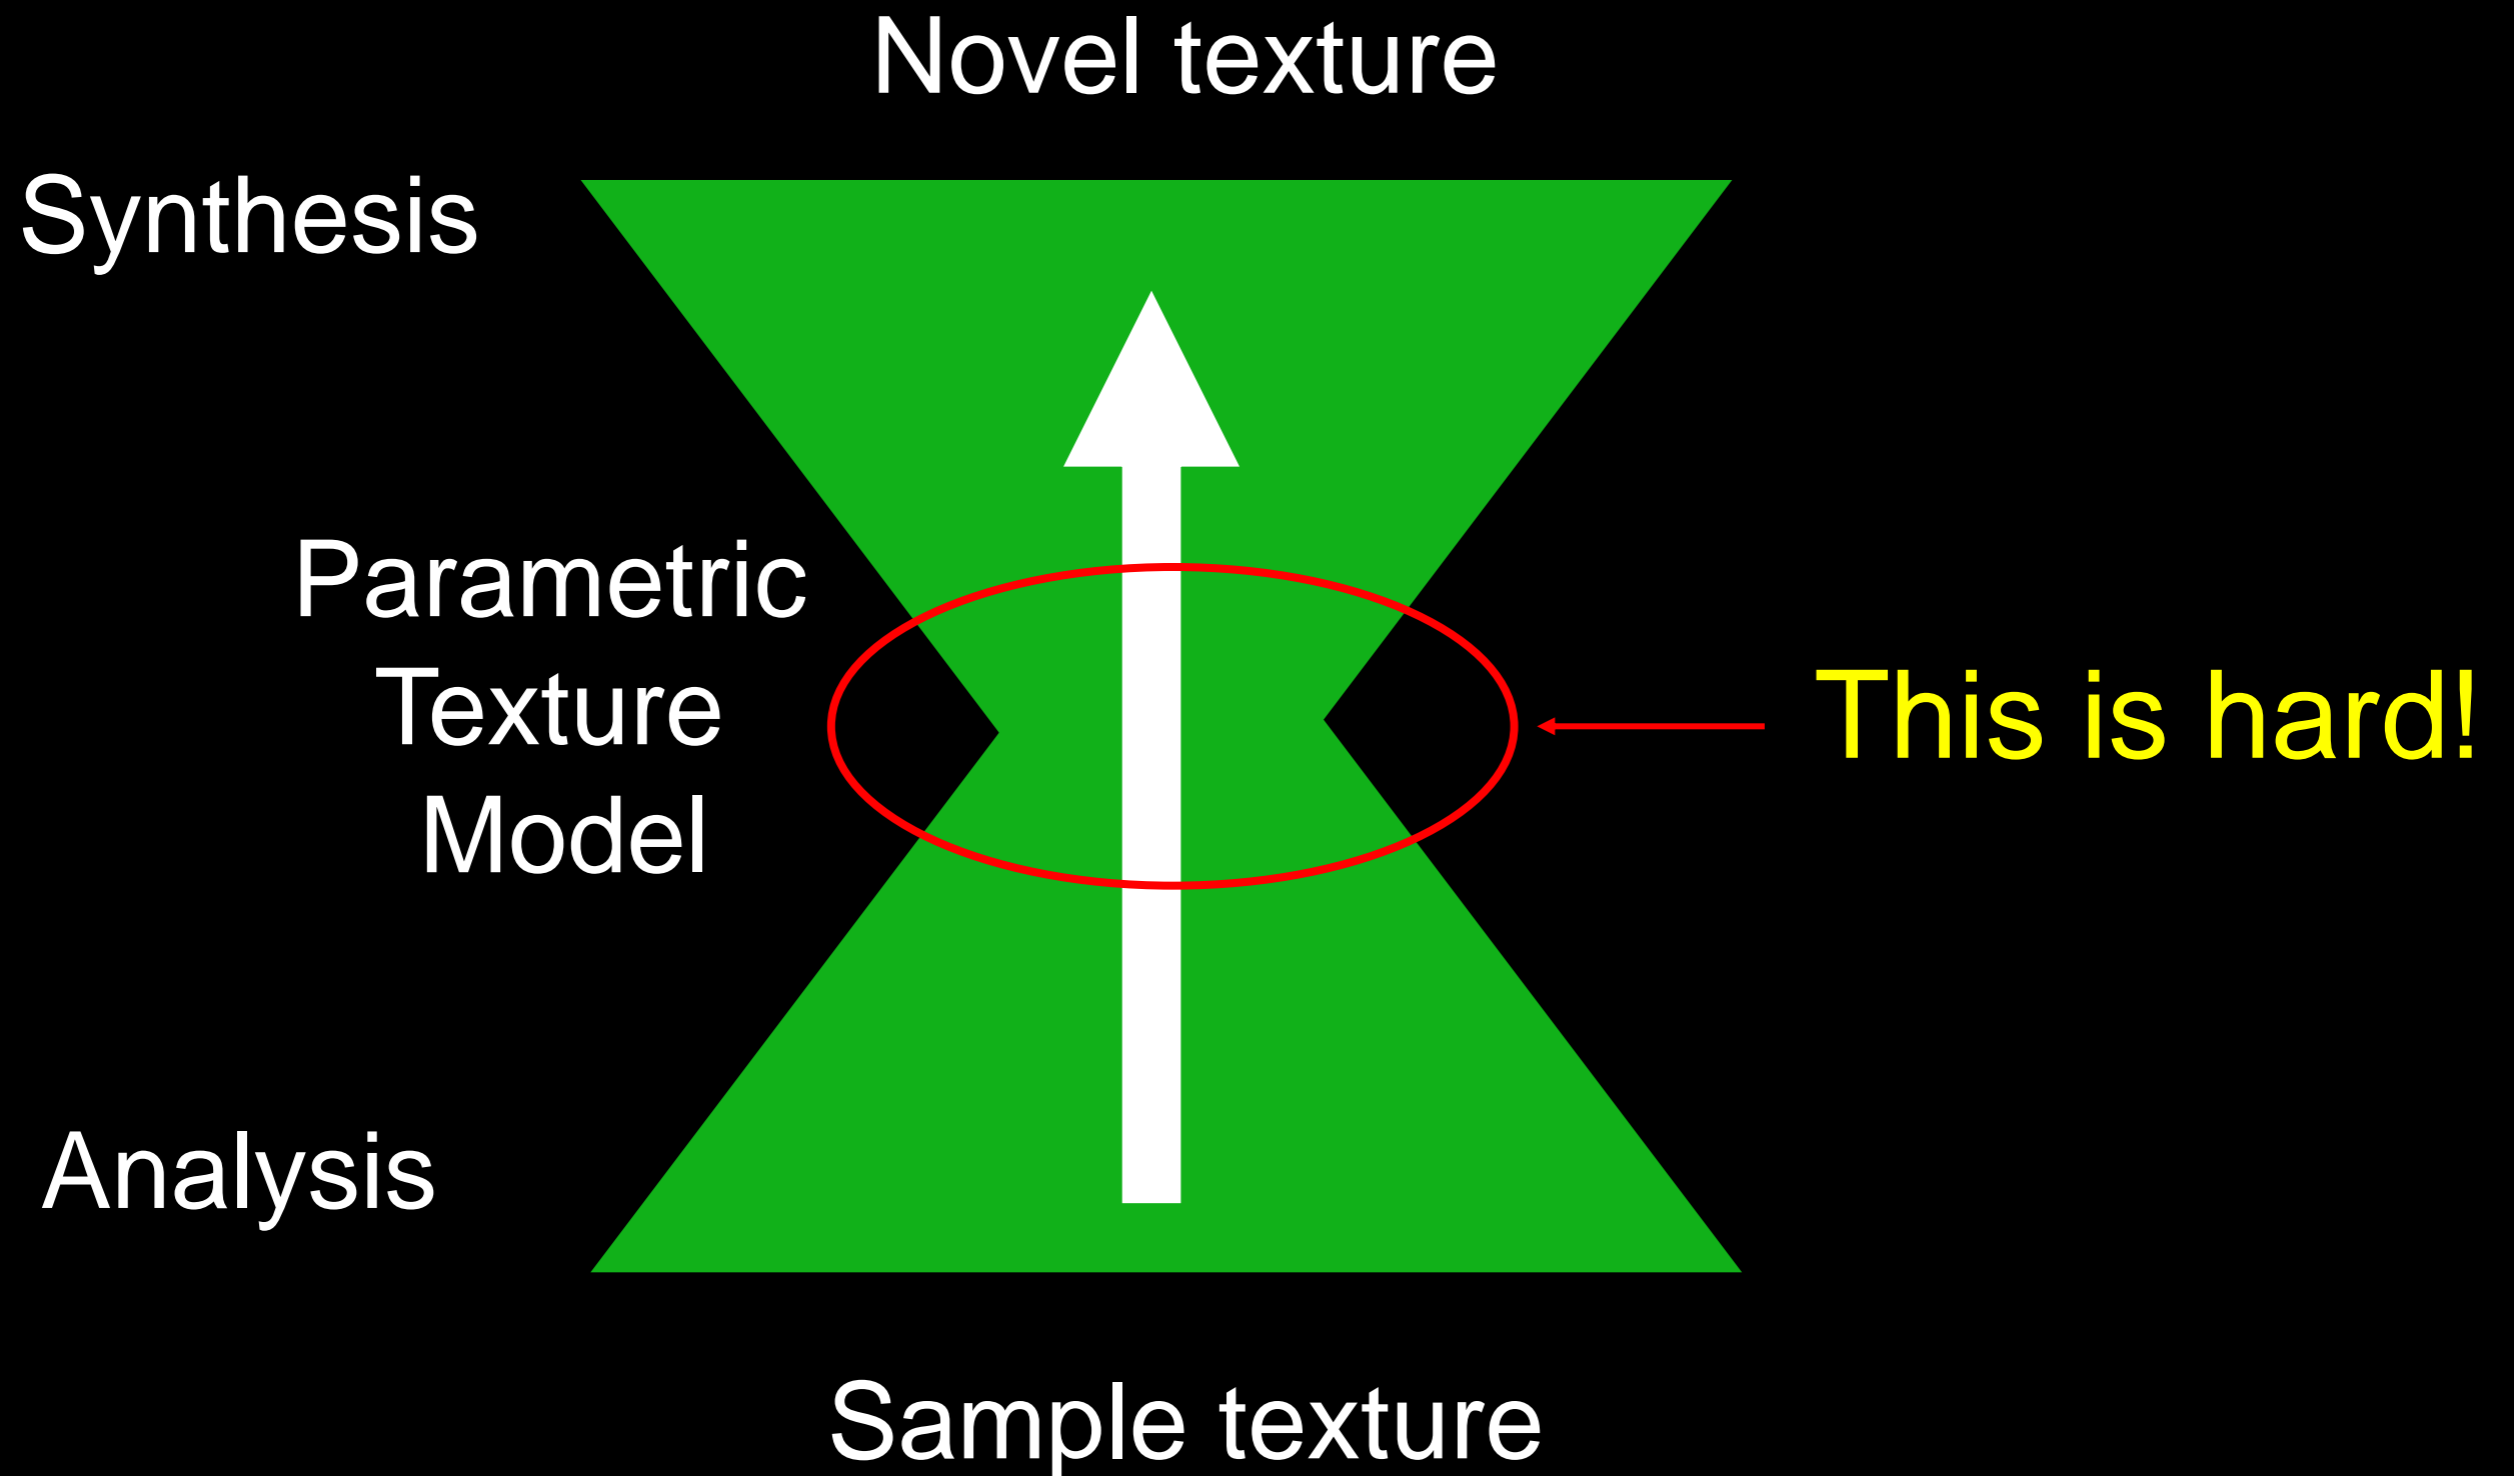
I have a problem...

...I am a nearest-neighbor addict!

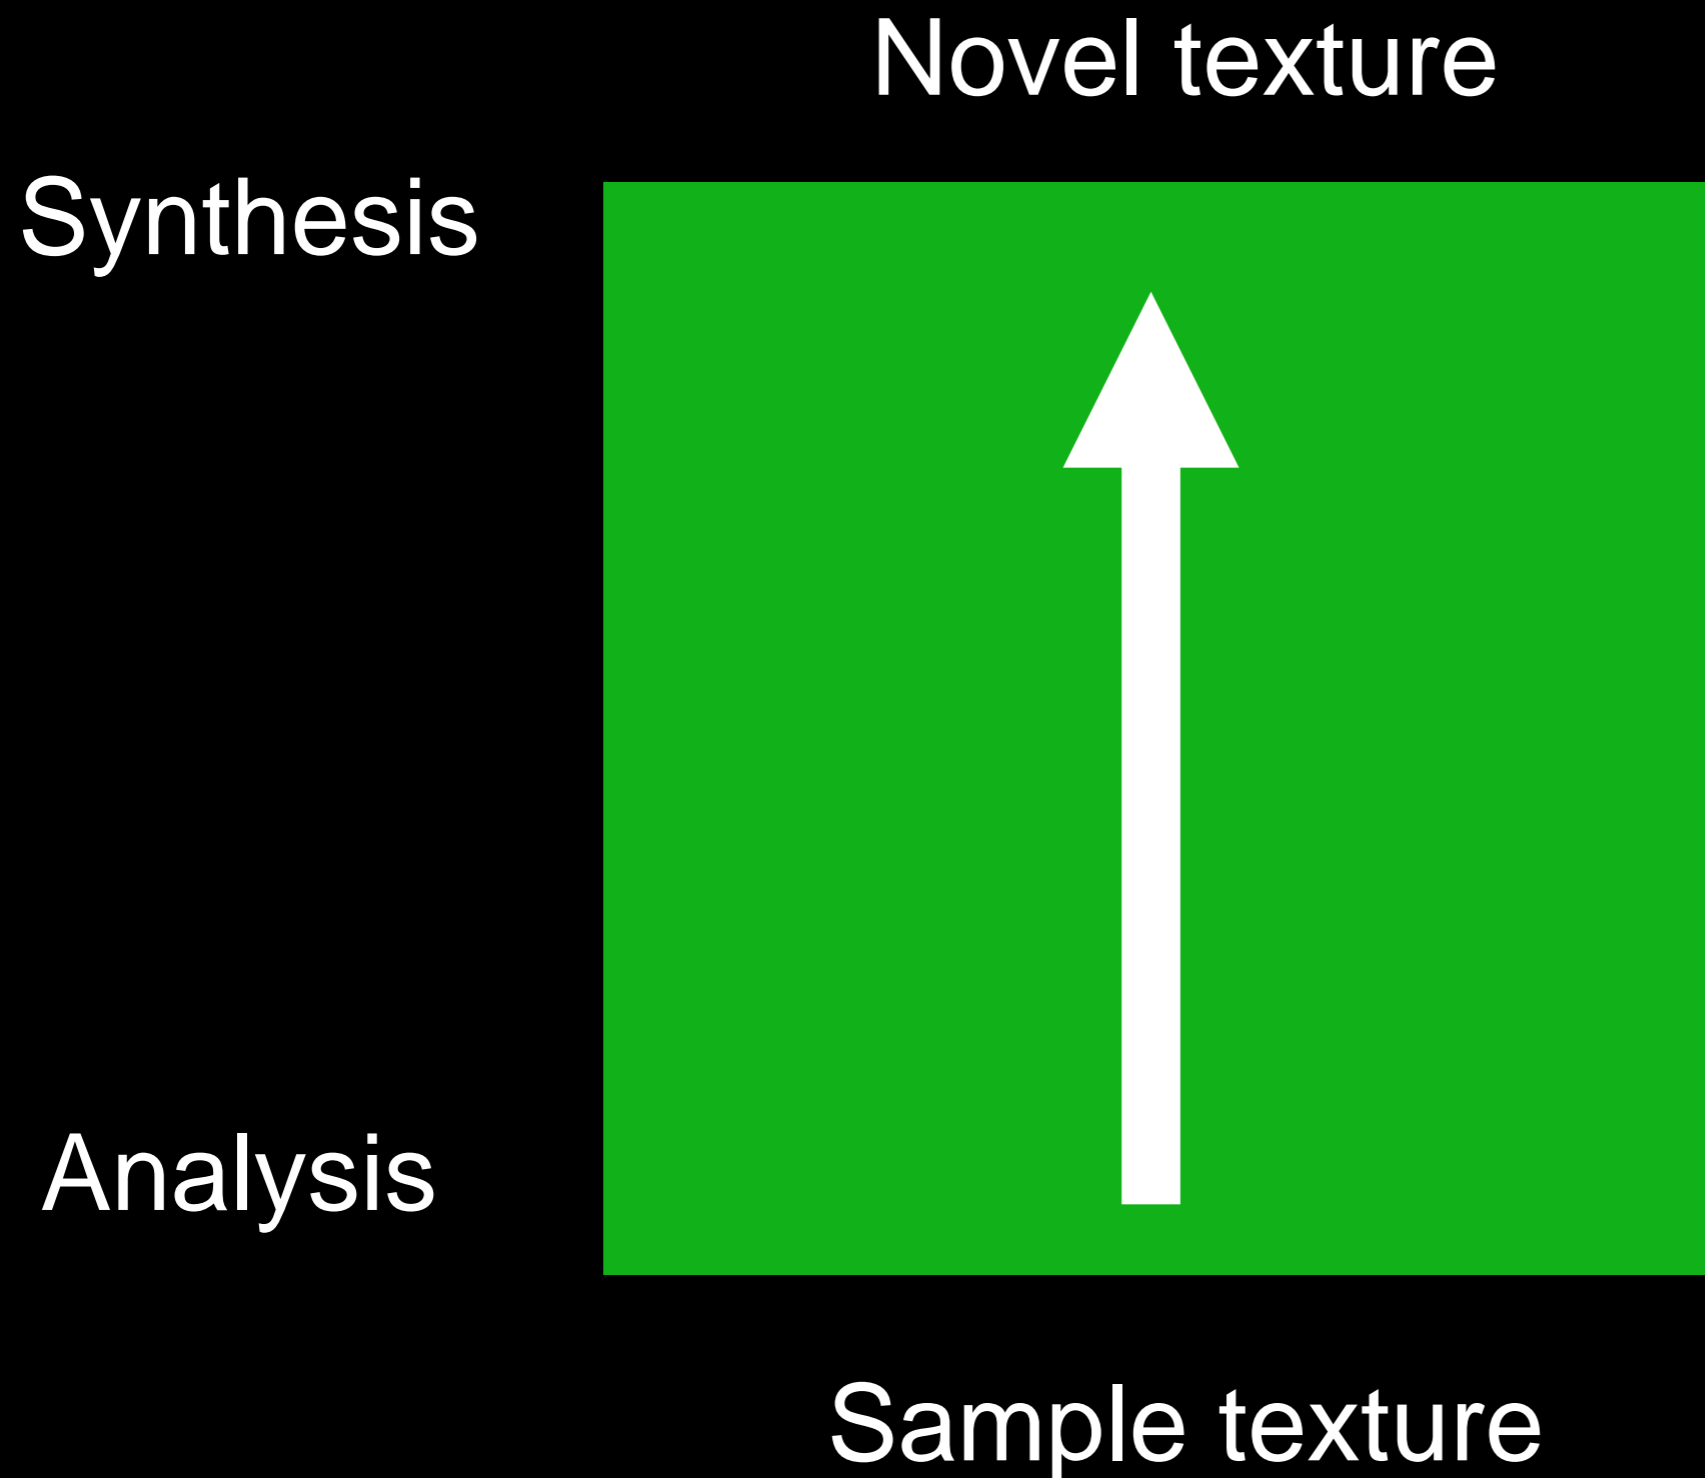
# Texture Synthesis




# Classical Texture Synthesis



# Non-parametric Approach




# Non-parametric Synthesis [Efros&Leung'99]




- Assuming Markov property, compute  $P(\mathbf{p}|\mathbf{N}(\mathbf{p}))$ 
  - Building explicit probability tables infeasible
  - Instead, we *search the input image* for all similar neighbourhoods — that's our distribution for  $\mathbf{p}$
  - To sample from distribution, just pick one match at random

# Hole Filling








**input image**




**Portilla & Simoncelli**



**Our algorithm**

... of a visual cortical neuron—their  
... describing the response of that neuro  
... as a function of position—is perhaps  
... functional description of that neuron.  
... seek a single conceptual and mathemat  
... describe the wealth of simple-cell recep  
... and neurophysiologically<sup>1-3</sup> and inferred  
... especially if such a framework has the  
... it helps us to understand the function  
... deeper way. Whereas no generic model  
... sians (DOG), difference of offset G  
... rivative of a Gaussian, higher derivati  
... function, and so on—can be expected  
... simple-cell receptive field, we noneth

**input image**




**Portilla & Simoncelli**

... sition—is perhaps a single conceptual and  
... of that neuron describe the wealth of simple-  
... al and neurophysiologically<sup>1-3</sup> and  
... simple-cell especially if such a framework  
... y<sup>1-3</sup> and infer helps us to understand the  
... amework has perhaps. Whereas no ge  
... and the neuron(DOG), difference o  
... s no generic a single conceptual and m  
... rence of offset the wealth of simple-ce  
... , higher deriescribing the response of t  
... —can be expected a function of position-  
... helps us to understand theription of the  
... per way. Whereas no conceptual an  
... sians (DOG), differencealth of simple

**Our algorithm**



**[Hays & Efros, SIGGRAPH'07]**





# 2 Million Flickr Images

The background of the slide is a dense, multi-colored mosaic composed of millions of tiny, square images. The colors are varied, including shades of blue, green, red, yellow, and grey, creating a complex, textured pattern. The overall effect is that of a vast collection of small visual elements.



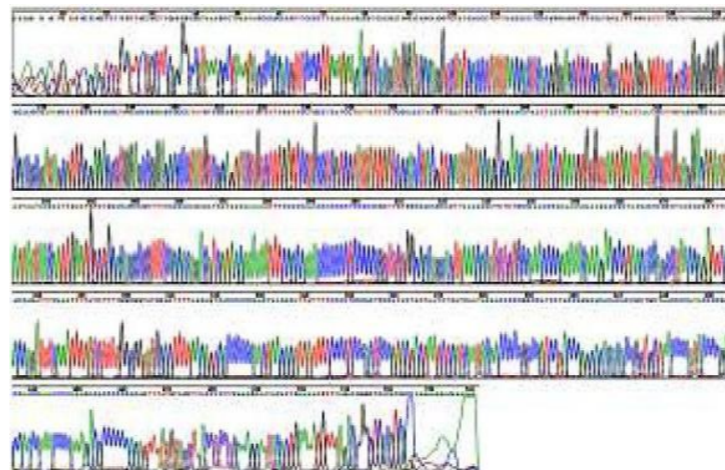


# Two Kinds of Things in the World



Navier-Stokes Equation

$$\frac{\partial \mathbf{u}}{\partial t} = - (\mathbf{u} \cdot \nabla) \mathbf{u} + \nu \nabla^2 \mathbf{u} - \frac{1}{\rho} \nabla p + \mathbf{f}$$



+ weather  
+ location  
+ ...



# Lots of data available

flickr® from YAHOO!

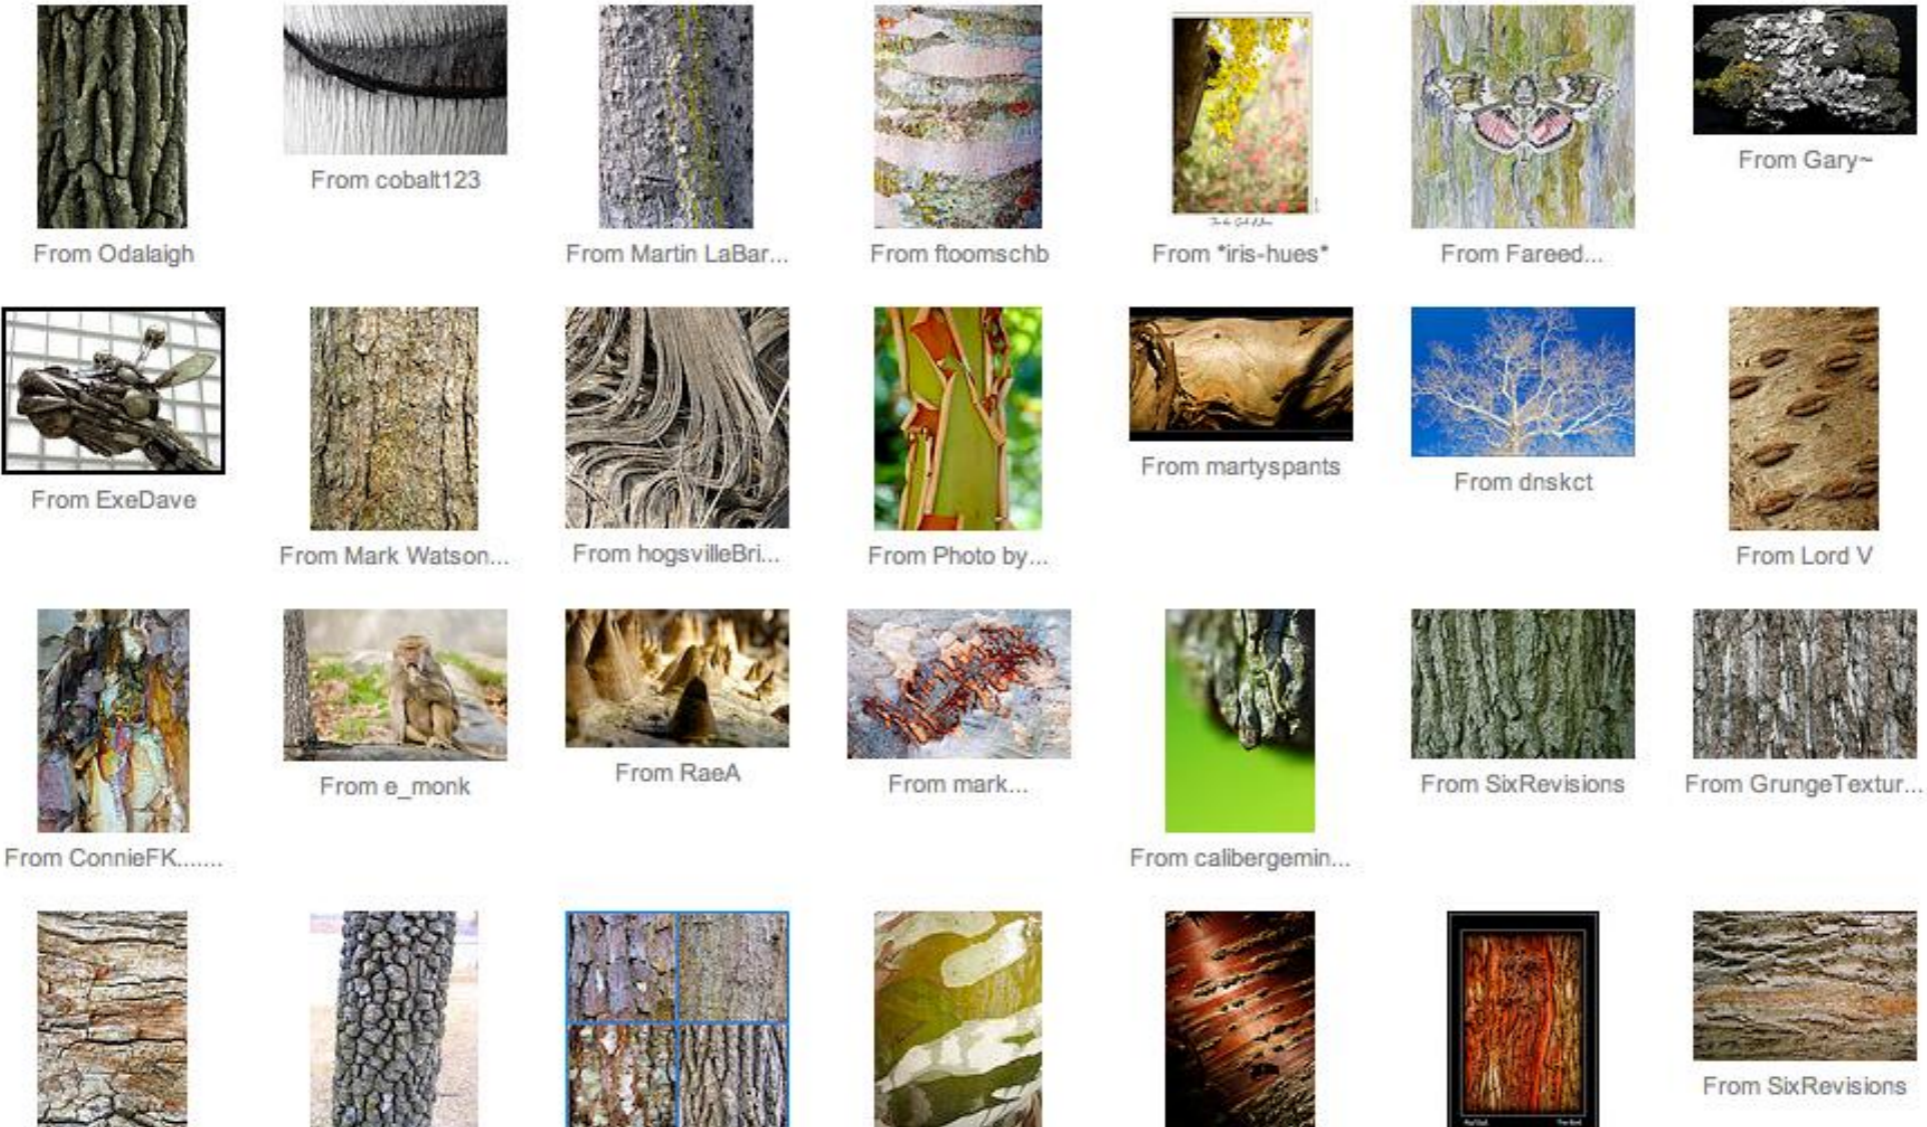
Signed in as [swatjarial](#)

[Home](#) [You](#) [Organize & Create](#) [Contacts](#) [Groups](#) [Explore](#)

**Search** | [Photos](#) | [Groups](#) | [People](#)

Everyone's Uploads | tree bark | **SEARCH** | [Full Text](#) | [Tags Only](#) | [Advanced Search](#)

Sort: [Relevant](#) | [Recent](#) | [Interesting](#) | View: [Small](#) | [Medium](#) | [Detail](#) | [Slideshow](#)



The image displays a grid of 28 Flickr search results for the query 'tree bark'. Each result consists of a small thumbnail image and a caption below it. The thumbnails show a variety of tree bark textures, including smooth, rough, peeling, and colorful. The captions identify the source user for each image, such as 'From Odalaigh', 'From cobalt123', 'From Martin LaBar...', 'From foomschb', 'From \*iris-hues\*', 'From Fareed...', 'From Gary~', 'From ExeDave', 'From Mark Watson...', 'From hogsvilleBri...', 'From Photo by...', 'From martyspants', 'From dnsckt', 'From Lord V', 'From ConnieFK.....', 'From e\_monk', 'From RaeA', 'From mark...', 'From calibergemin...', 'From SixRevisions', and 'From GrungeTextur...'. The grid is organized into four rows of seven images each.

# “Unreasonable Effectiveness of Data”

[Halevy, Norvig, Pereira 2009]

- Parts of our world can be explained by elegant mathematics:
  - physics, chemistry, astronomy, etc.
- But much cannot:
  - psychology, genetics, economics, etc.
- Enter: The Magic of **Big Data**
  - Great advances in several fields:
    - e.g. speech recognition, machine translation, Google



- A.I. for the postmodern world:
  - all questions have already been answered...many times, in many ways
  - Google is dumb, the “intelligence” is in the data




# Are Categories Necessary?



Alexei (Alyosha) Efros  
CMU

Joint work with  
Tomasz Malisiewicz

# Acknowledgements




Murphy

*Big Book of Concepts*

Everything Is  
Miscellaneous

THE POWER OF THE  
NEW DIGITAL DISORDER

"Perfectly placed to tell us what's really new  
about the second-generation Web."  
—Los Angeles Times



Weinberger

*Everything is Miscellaneous*

*Talks by Moshe Bar; writings of Shimon Edelman*

Many great discussions with many colleagues, especially  
Tomasz Malisiewicz, James Hays, and Derek Hoiem

# Understanding an Image



slide by Fei Fei, Fergus & Torralba

# Object naming -> Object categorization



sky

building

flag

face

banner

wall

street lamp

bus

bus

cars

slide by Fei Fei, Fergus & Torralba

# Object categorization

sky

building

flag

face

banner

wall

street lamp

bus

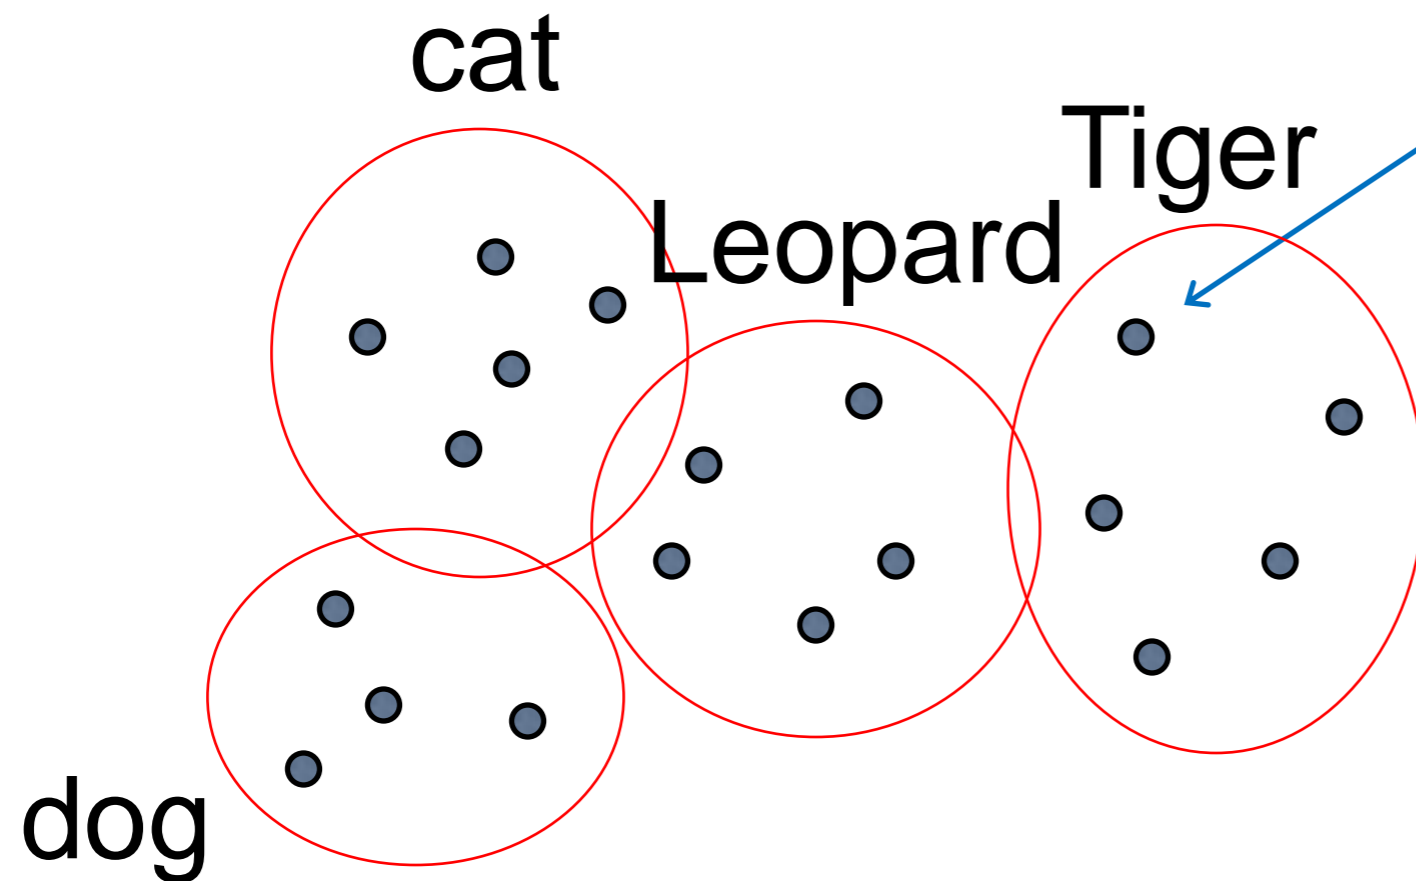
bus

cars



# Why Categorize?

1. Knowledge Transfer
2. Communication



# Classical View of Categories

- Dates back to Plato & Aristotle
  1. Categories are defined by a list of properties shared by all elements in a category
  2. Category membership is binary
  3. Every member in the category is equal




# Problems with Classical View

- Humans don't do this!
  - People don't rely on abstract definitions / lists of shared properties (Wittgenstein 1953, Rosch 1973)
    - e.g. define the properties shared by all "games"
    - e.g. are curtains furniture? Are olives fruit?
  - Typicality
    - e.g. Chicken -> bird, but bird -> eagle, pigeon, etc.
  - Language-dependent
    - e.g. "Women, Fire, and Dangerous Things" category is Australian aboriginal language (Lakoff 1987)
  - Doesn't work even in human-defined domains
    - e.g. Is Pluto a planet?

# Problems with Visual Categories

- A lot of categories are functional

**Chair**



- World is too varied




- Categories are 3D, but images are 2D

**car**



# Typical HOG car detector



Felzenszwalb et al, PASCAL 2007

# Why not?




+



# Solution: hierarchy?


Ontologies, hierarchies, levels of categories (Rosch), etc.

WordNet, ImageNet, etc etc



# Still Problematic!

- Intransitivity
  - e.g. car seat is chair, chair is furniture, but ...
- Multiple category membership
  - it's not a tree, it's a forest!





# Fundamental Problem with Categorization



**Making decisions too early!**

**We should only categorize at run-time, once we know the task!**

# The Dictatorship of Librarians



Arts and recreation

Literature

Technology

Religion

amazon.com Hello, Tomasz Malisiewicz. We have [recommendations](#) for you. (Not Tomasz?)  
Tomasz's Amazon.com | Today's Deals | [Gifts & Wish Lists](#) | [Gift Cards](#)

Shop All Departments Search

Books | [Advanced Search](#) | [Browse Subjects](#) | [New Releases](#) | [Bestsellers](#) | [The New York Times® Bestsellers](#)

Department  
< Any Department  
**Books**  
Professional & Technical (7)  
Science (5)  
Computers & Internet (6)

Format  
**Any Format**  
Printed Books (7)  
HTML (1)

Binding  
**Any Binding**  
Hardcover (4)  
Paperback (3)

Shipping Option (What's this?)

Books > "autonomous driving"

Showing 10 Results

- Autonomous driving in traffic: boss and the Urban AI Magazine** by Chris Urmson, Chris Baker, John Dolan (2009) - HTML  
[Buy: \\$9.95](#)  
Available for download now
- The DARPA Urban Challenge: Autonomous Vehicles and Advanced Robotics** by Martin Buehler, Karl Iagnemma (2009)  
[Buy new: \\$199.00 \\$143.20](#)  
[10 new](#) from \$143.20 [5 used](#) from \$142.17  
Get it by **Thursday, Feb. 18** if you order in the next **3 hours** and are eligible for **FREE** Super Saver Shipping.

categories are losing...


The image shows the classic Yahoo! logo, which consists of the word "YAHOO!" in a bold, red, serif font with a slight shadow effect. The logo is centered within a white rectangular area.

vs.

The image shows the Google logo, which is the word "Google" in its signature multi-colored font (blue, red, yellow, blue, green, red) with a trademark symbol (TM) to the upper right. The logo is centered within a white rectangular area.

# On-the-fly Categorization?

1. Knowledge Transfer
2. ~~Communication~~



# Association instead of categorization

*Ask not “what is this?”, ask “what is this like”*



– Moshe Bar

- Exemplar Theory (Medin & Schaffer 1978, Nosofsky 1986, Krushke 1992)
  - categories represented in terms of remembered objects (exemplars)
  - Similarity is measured between input and all exemplars
  - *think* non-parametric density estimation
- Vanevar Bush (1945), Memex (MEMory EXtender)
  - Inspired hypertext, WWW, Google...

# Bush's Memex (1945)

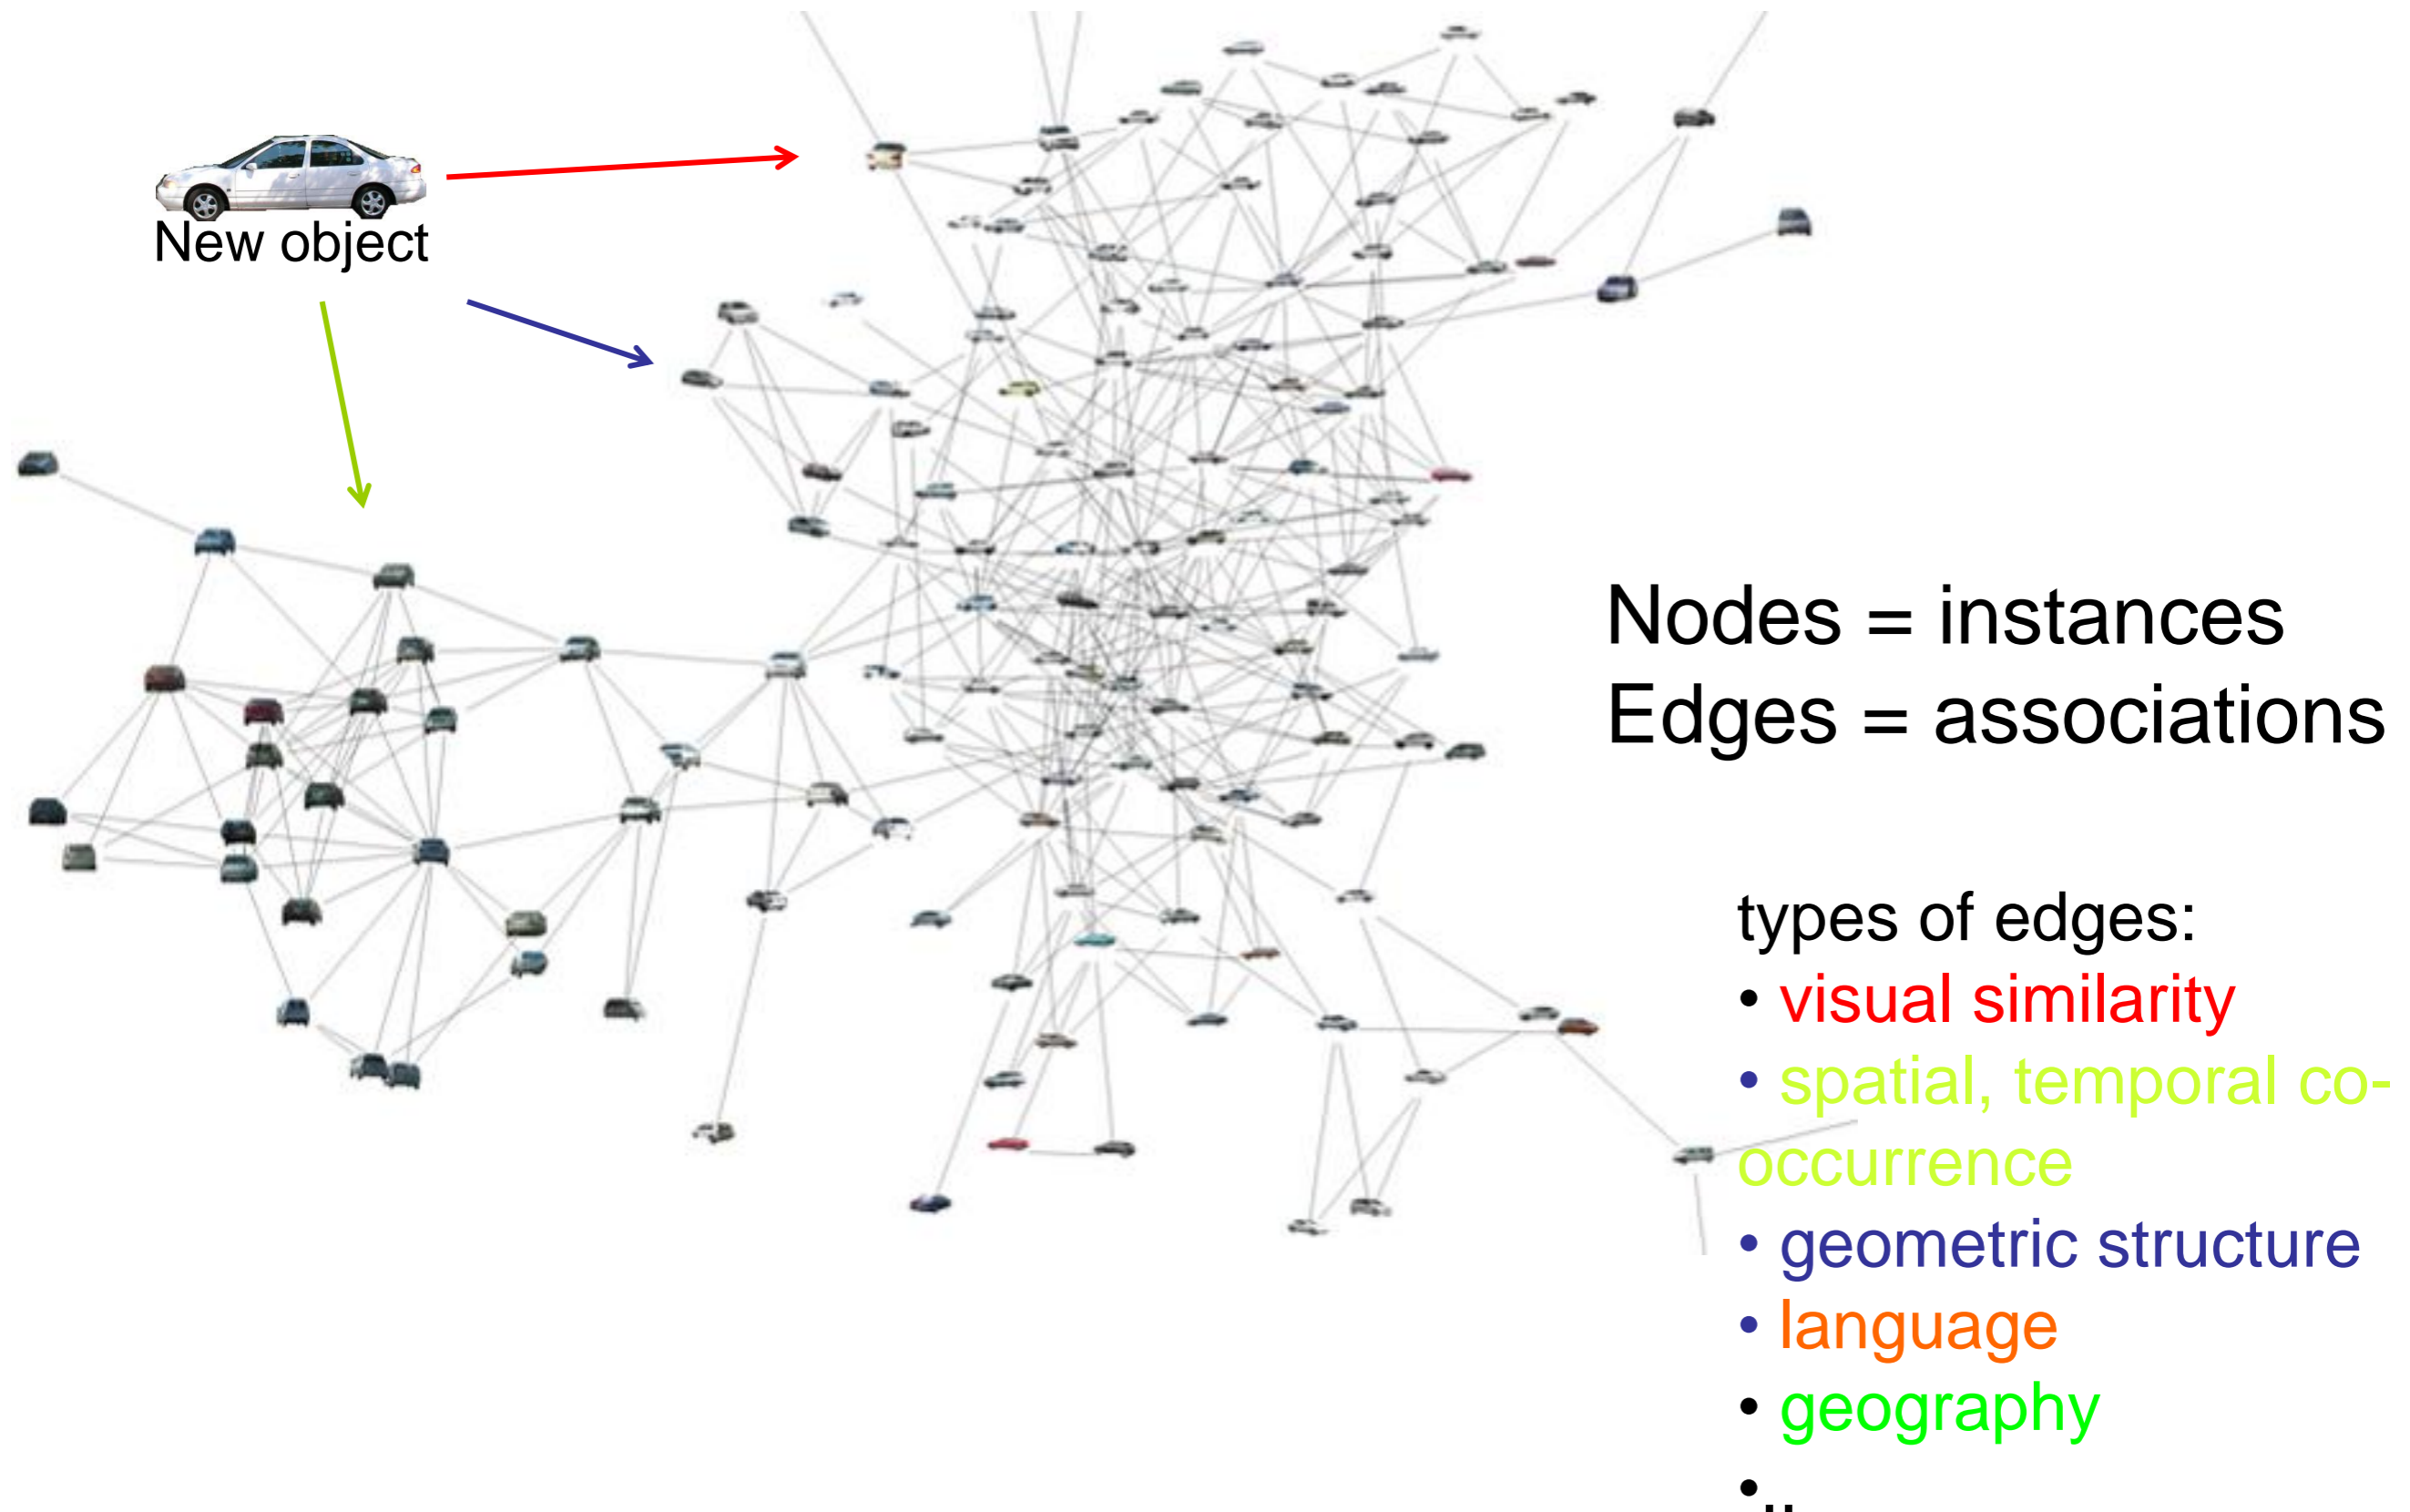


- Store publications, correspondence, personal work, on microfilm
- Items retrieved rapidly using index codes
  - Builds on “rapid selector”
- Can annotate text with margin notes, comments
- Can construct a *trail* through the material and save it
  - Roots of hypertext
- Acts as an external memory



# Visual Memex, a proposal

[Malisiewicz & Efros]



# “What is this?”


Input Image



He 2004, Tu 2004, Shotton 2006, Galleguillos 2008,  
Fei-Fei 2009, Gould 2009, etc.



“What is this *like*?”



Malisiewicz & Efros, CVPR'08

# Visual Associations

- How are objects similar?



Shape



~~Color~~



~~Shape~~



Color



# Distance “Similarity” Functions


- Positive Linear Combinations of Elementary Distances Computed Over 14 Features

$$D_e(z) = \mathbf{w}_e \cdot \mathbf{d}_{ez}$$

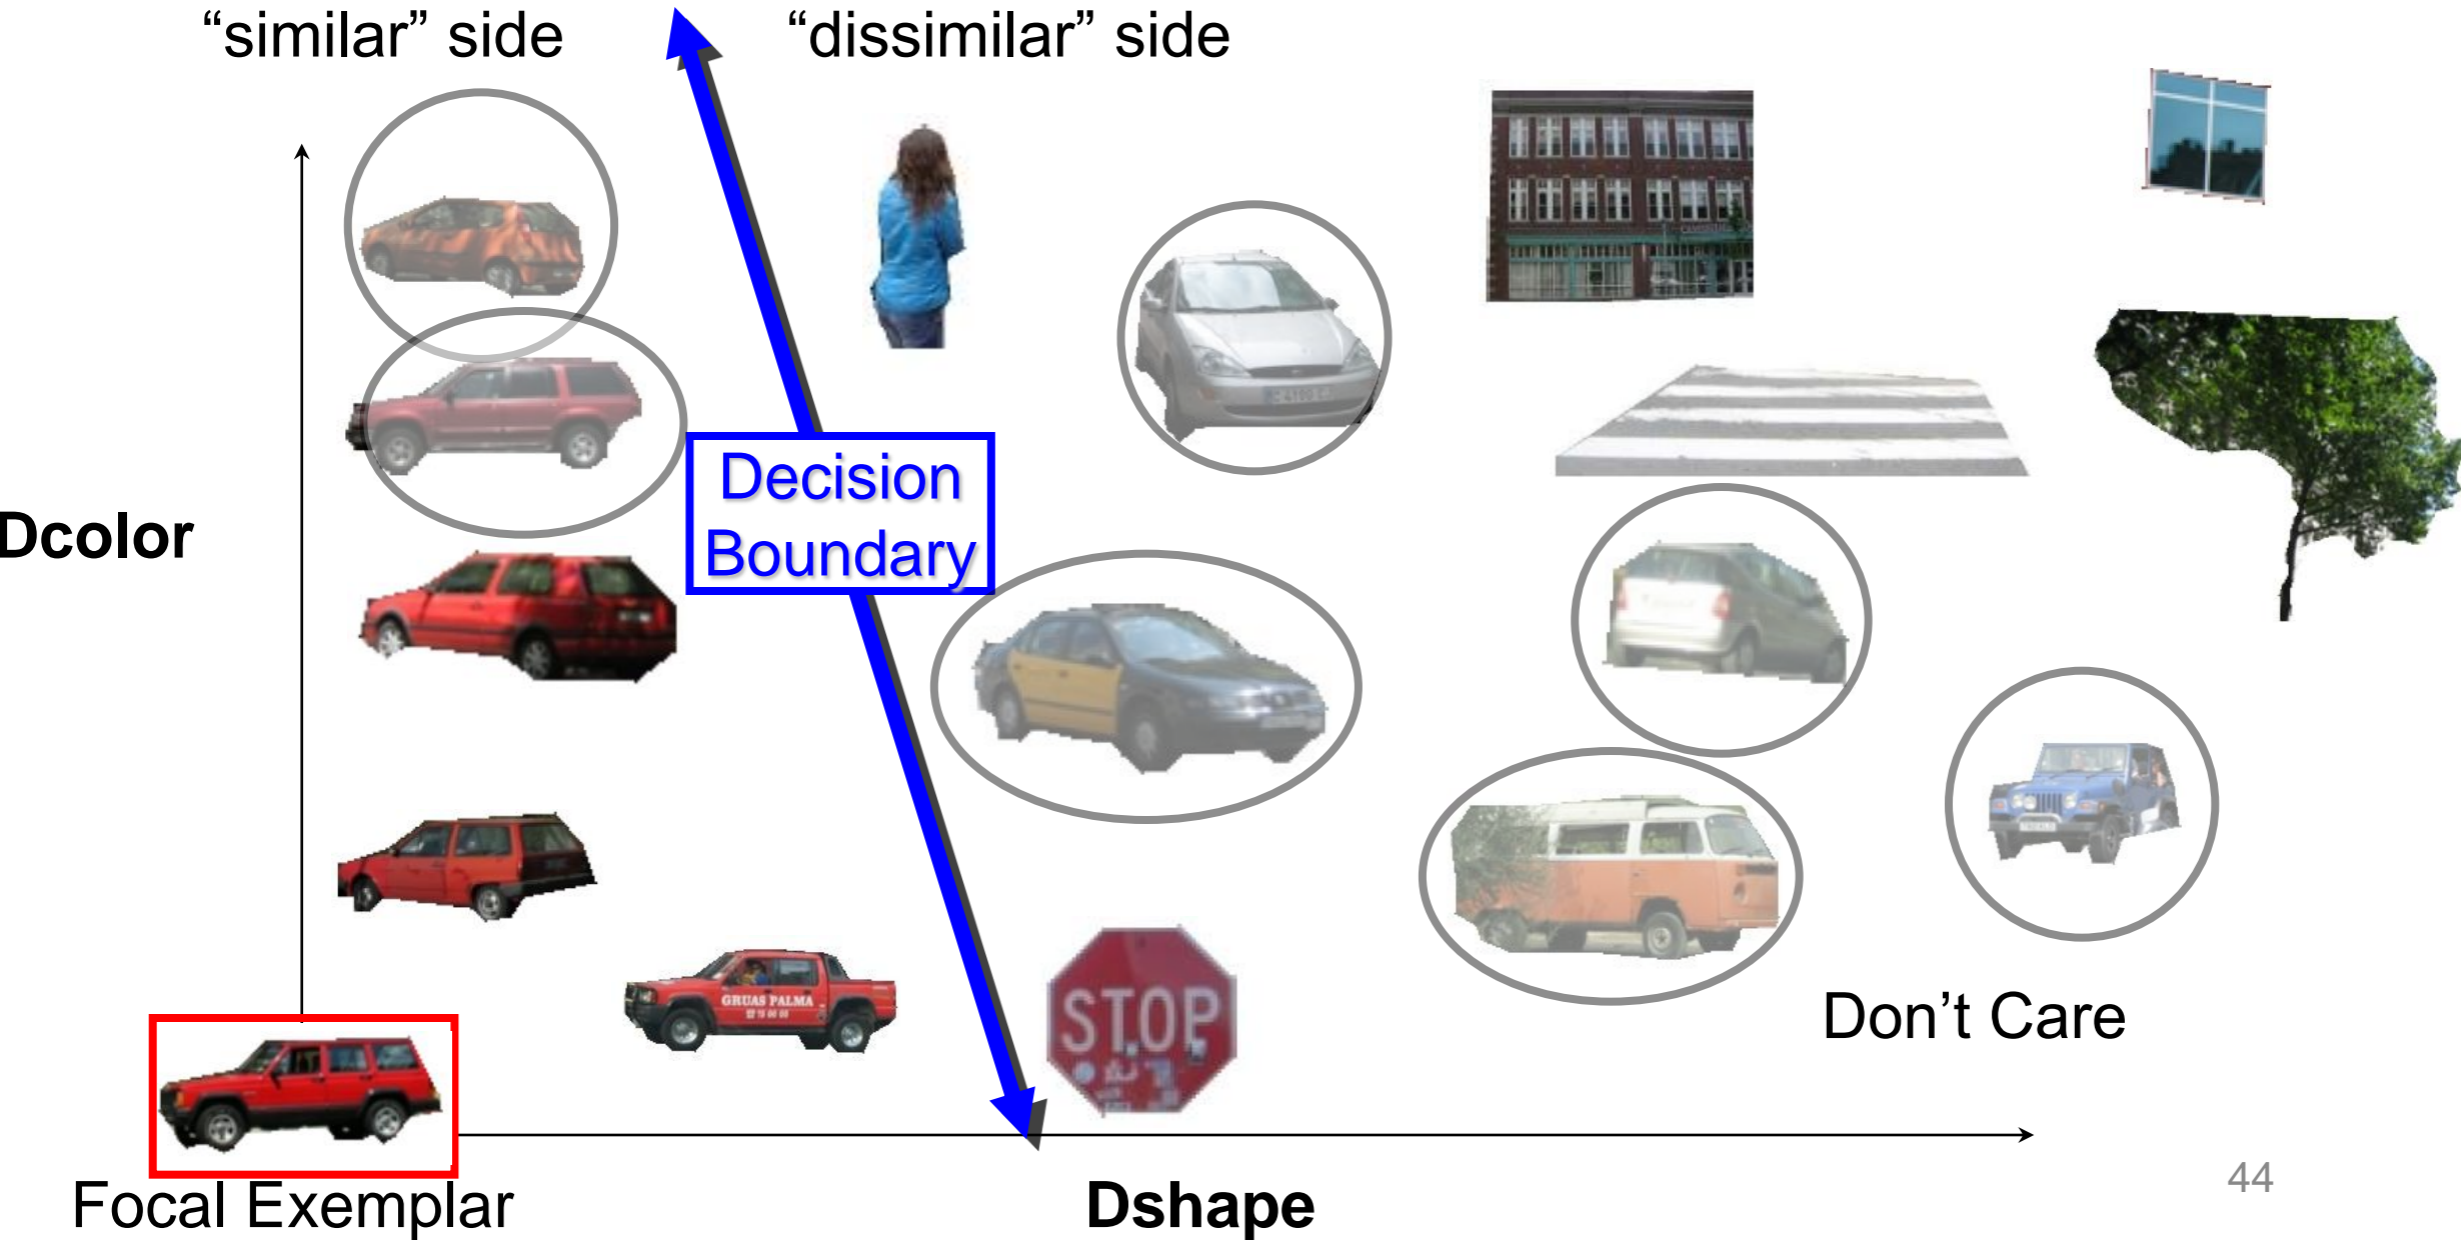
Building e



Building e Distance Function



# Learning Distance Functions



# Visualizing Distance Functions (Training Set)

Query



Top Neighbors with Tex-Hist Dist




Distance Function

Query




Top Neighbors with Learned Dist



# Visualizing Distance Functions (Training Set)



Distance Function



# Labels Crossing Boundary

stop sign	sign	7.8%
pole	streetlight	6.7%
motorcycle	motorbike	6.2%
mountains	mountain	6.2%
ground grass	sidewalk	3.7%
grass	lawn	3.6%
road highway	road	3.4%
painting	picture	3.4%
sidewalk	road	3.2%
cloud	sky	3.1%
grass	ground grass	3.1%
mountain	mountains	2.7%

Table 2: Top dozen label confusions discovered after distance function learning.

# Image Parsing with Context

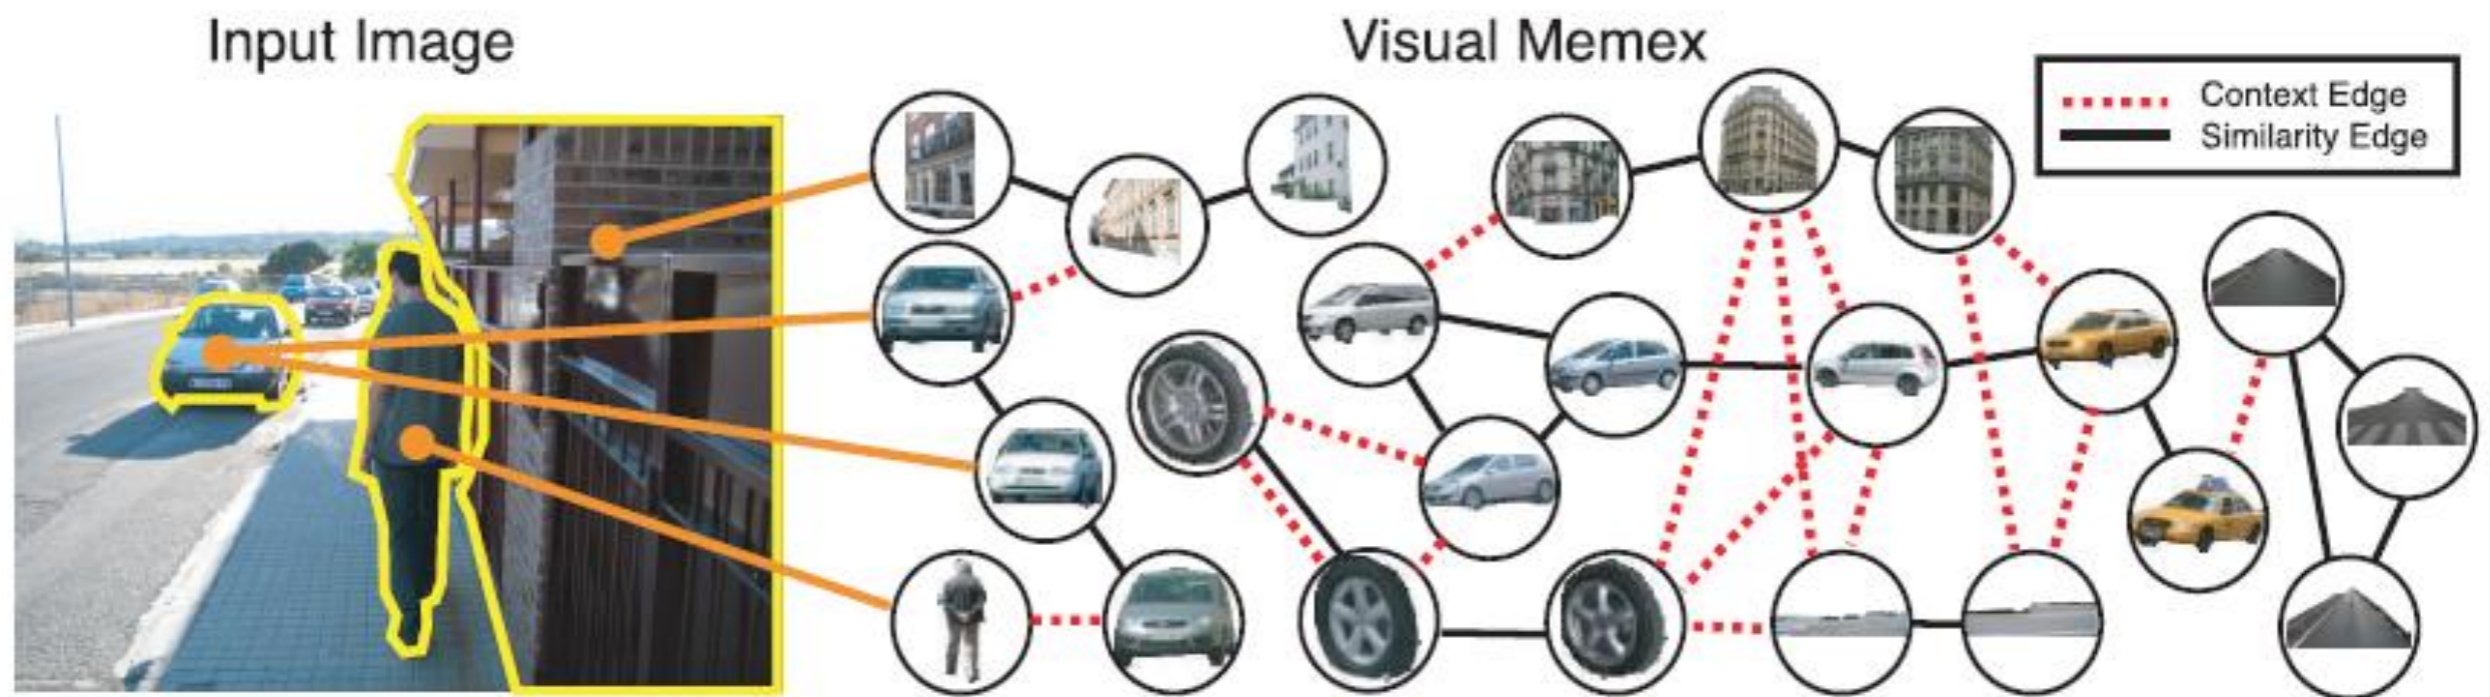


Figure 1: The **Visual Memex** graph encodes object similarity (solid black edge) and spatial context (dotted red edge) between pairs of object exemplars. A spatial context feature is stored for each context edge. The Memex graph can be used to interpret a new image (left) by associating image segments with exemplars in the graph (orange edges) and propagating the information.



# Torralba's Context Challenge

# Torralba's Context Challenge

---



# Torralba's Context Challenge

---



# Our Challenge Setup

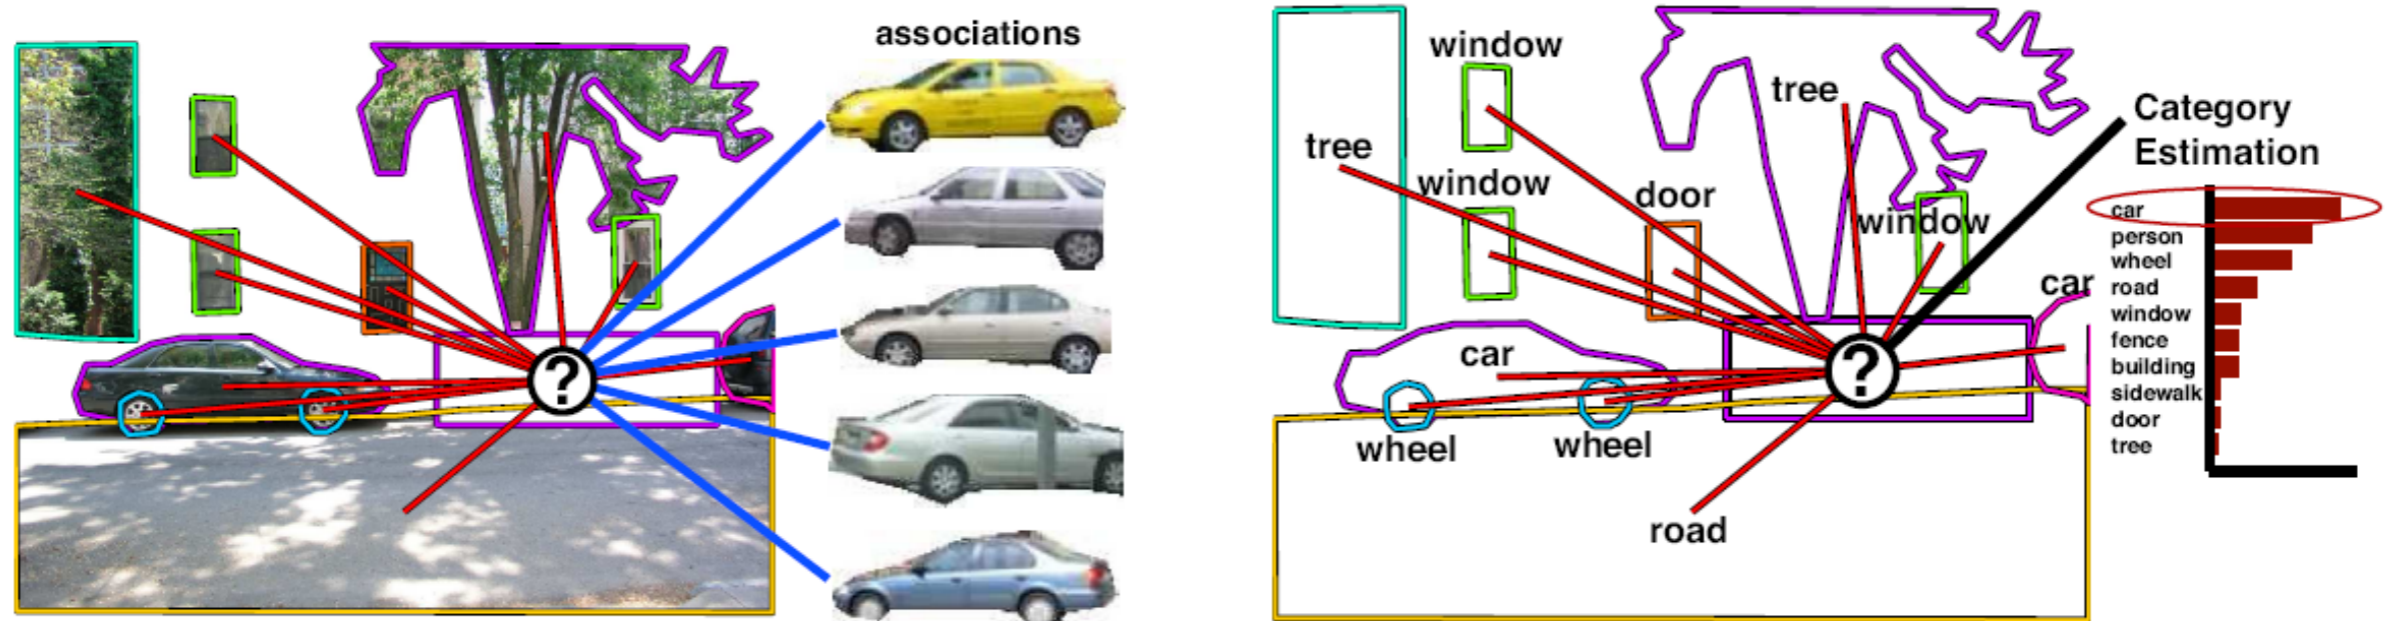


Figure 2: Torralba's Context Challenge: "How far can you go without running a local object detector?" The task is to reason about the identity of the hidden object (denoted by a "?") without local information. In our category-free Visual Memex model, object predictions are generated in the form of exemplar associations for the hidden object. In a category-based model, the category of the hidden object is directly estimated.

# 3 models

---

Visual Memex: exemplars, non-parametric  
object-object relationships

- Recurse through the graph

Baseline: CoLA: categories, parametric object-  
object relationships

Reduced Memex: categories, non-parametric  
relationships

# Qual. results

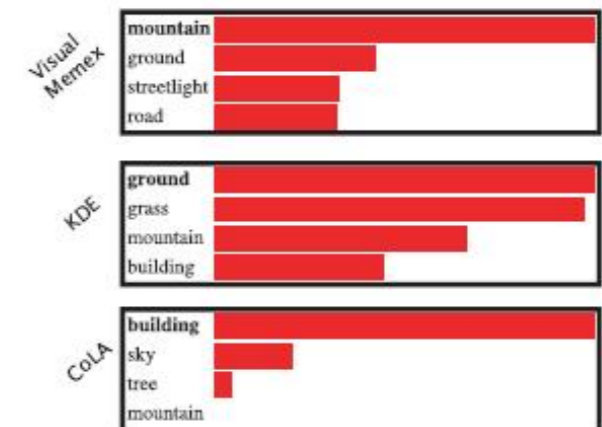
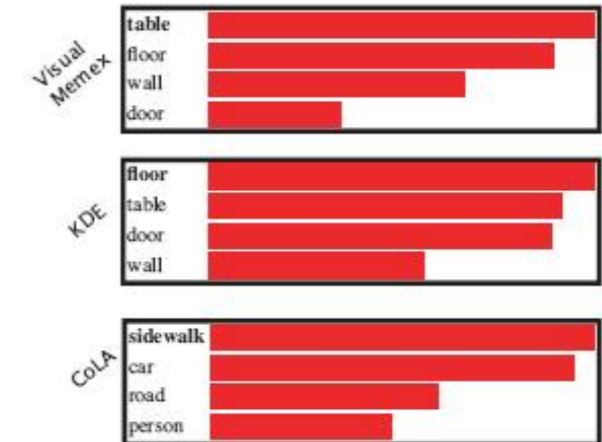
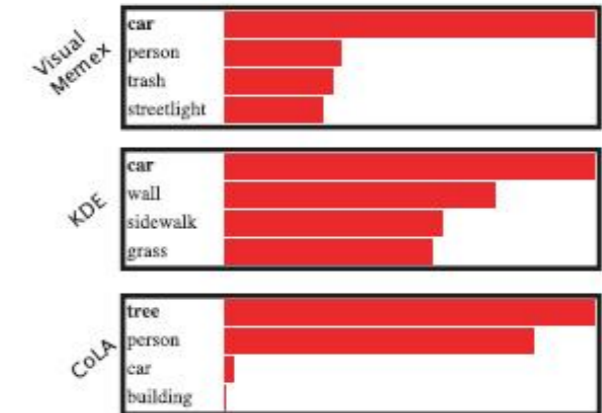
Input Image + Hidden Region



Visual Memex Exemplar Predictions



Categorization Results





# Quant. results

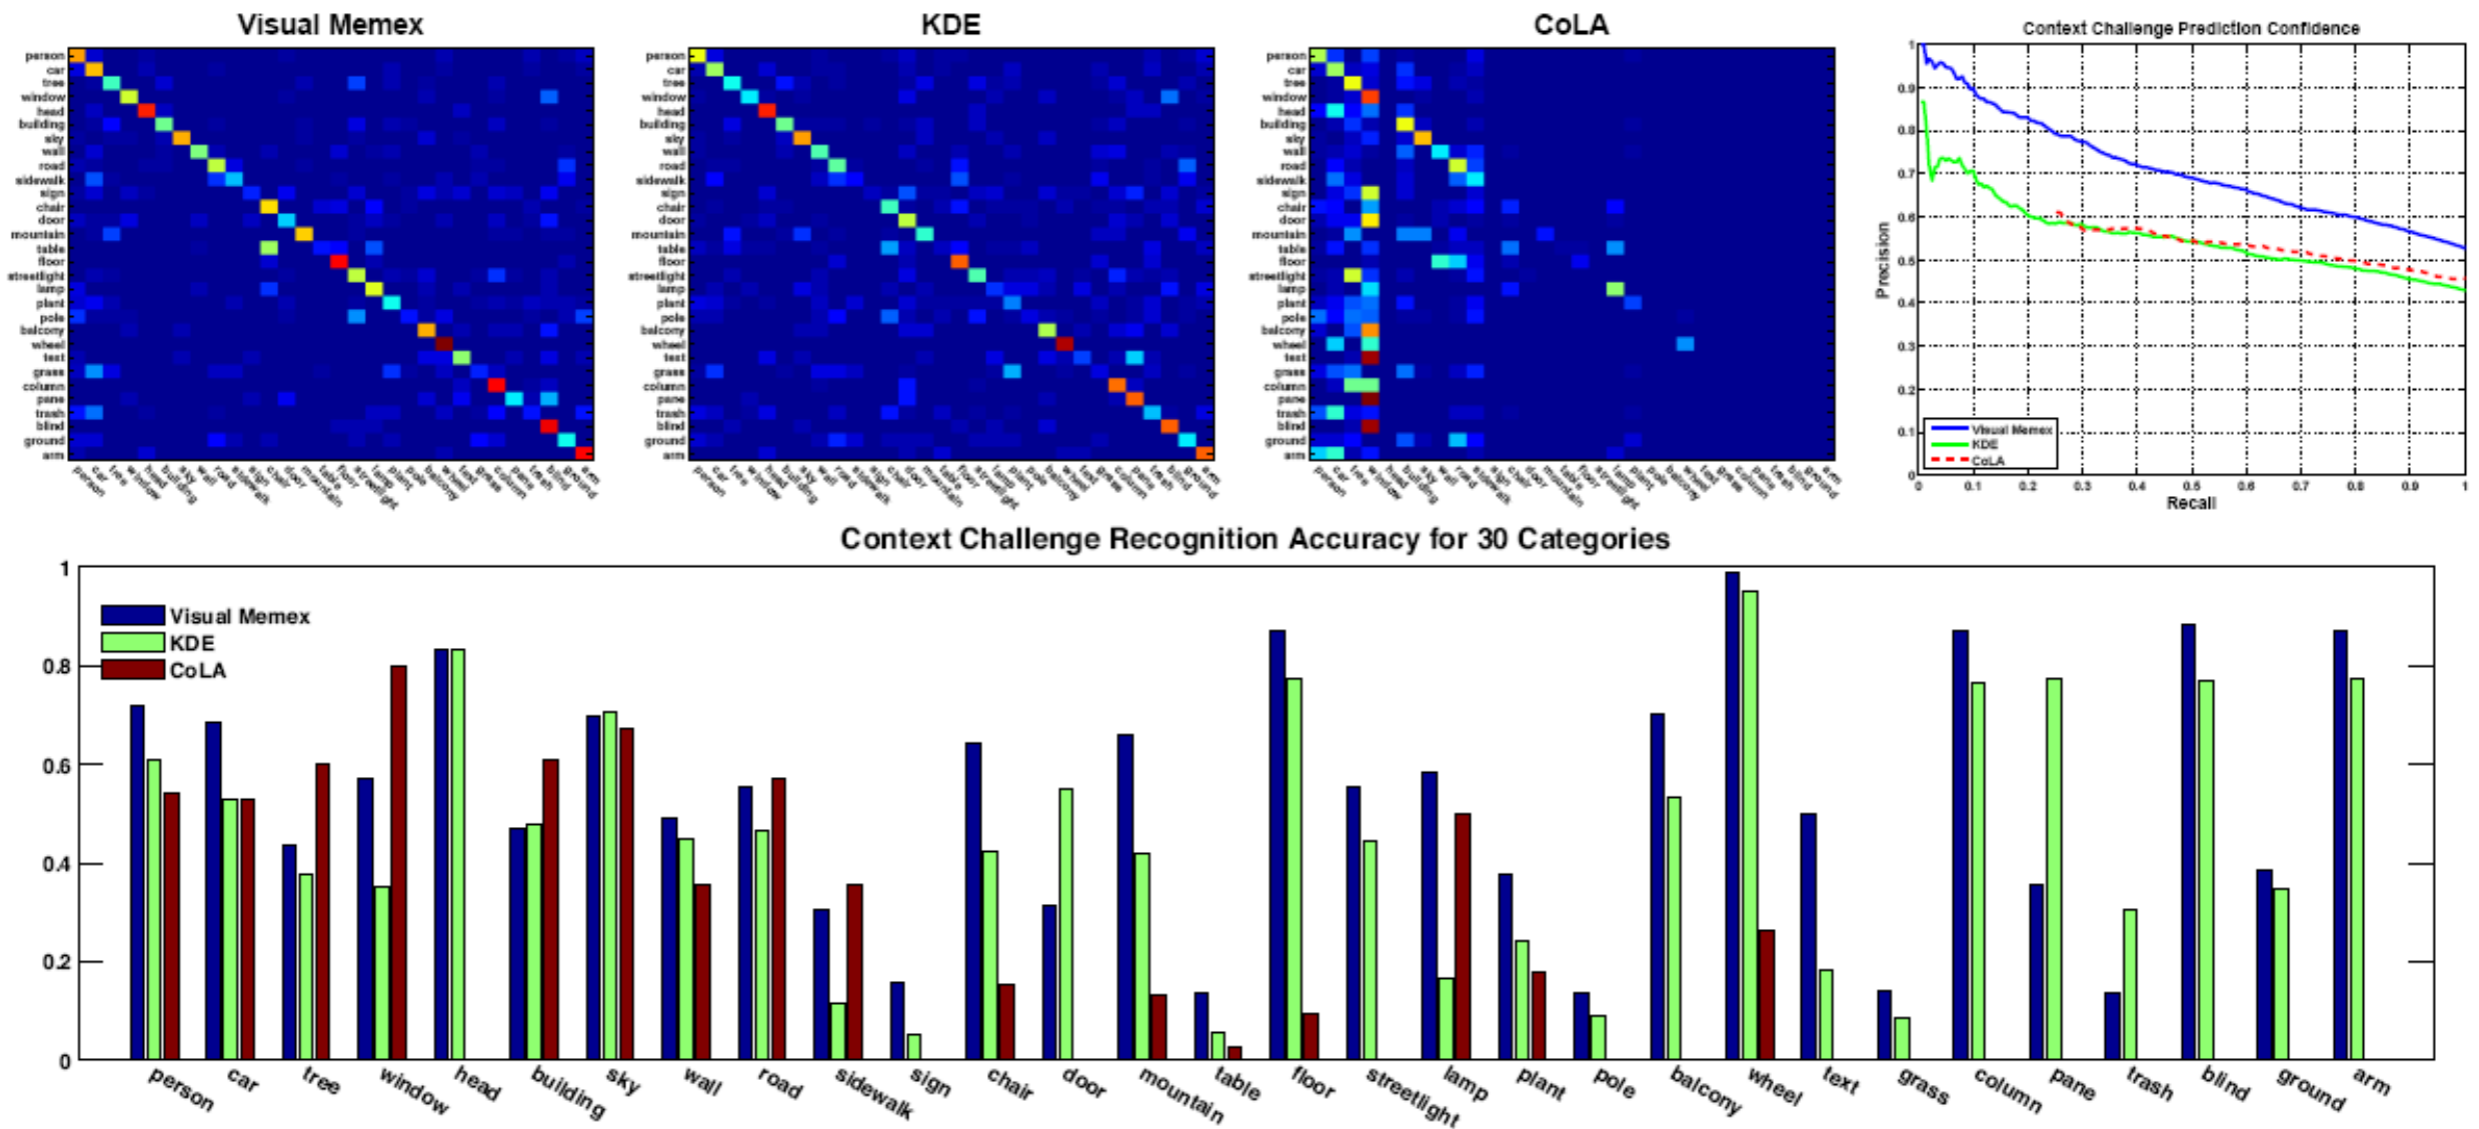
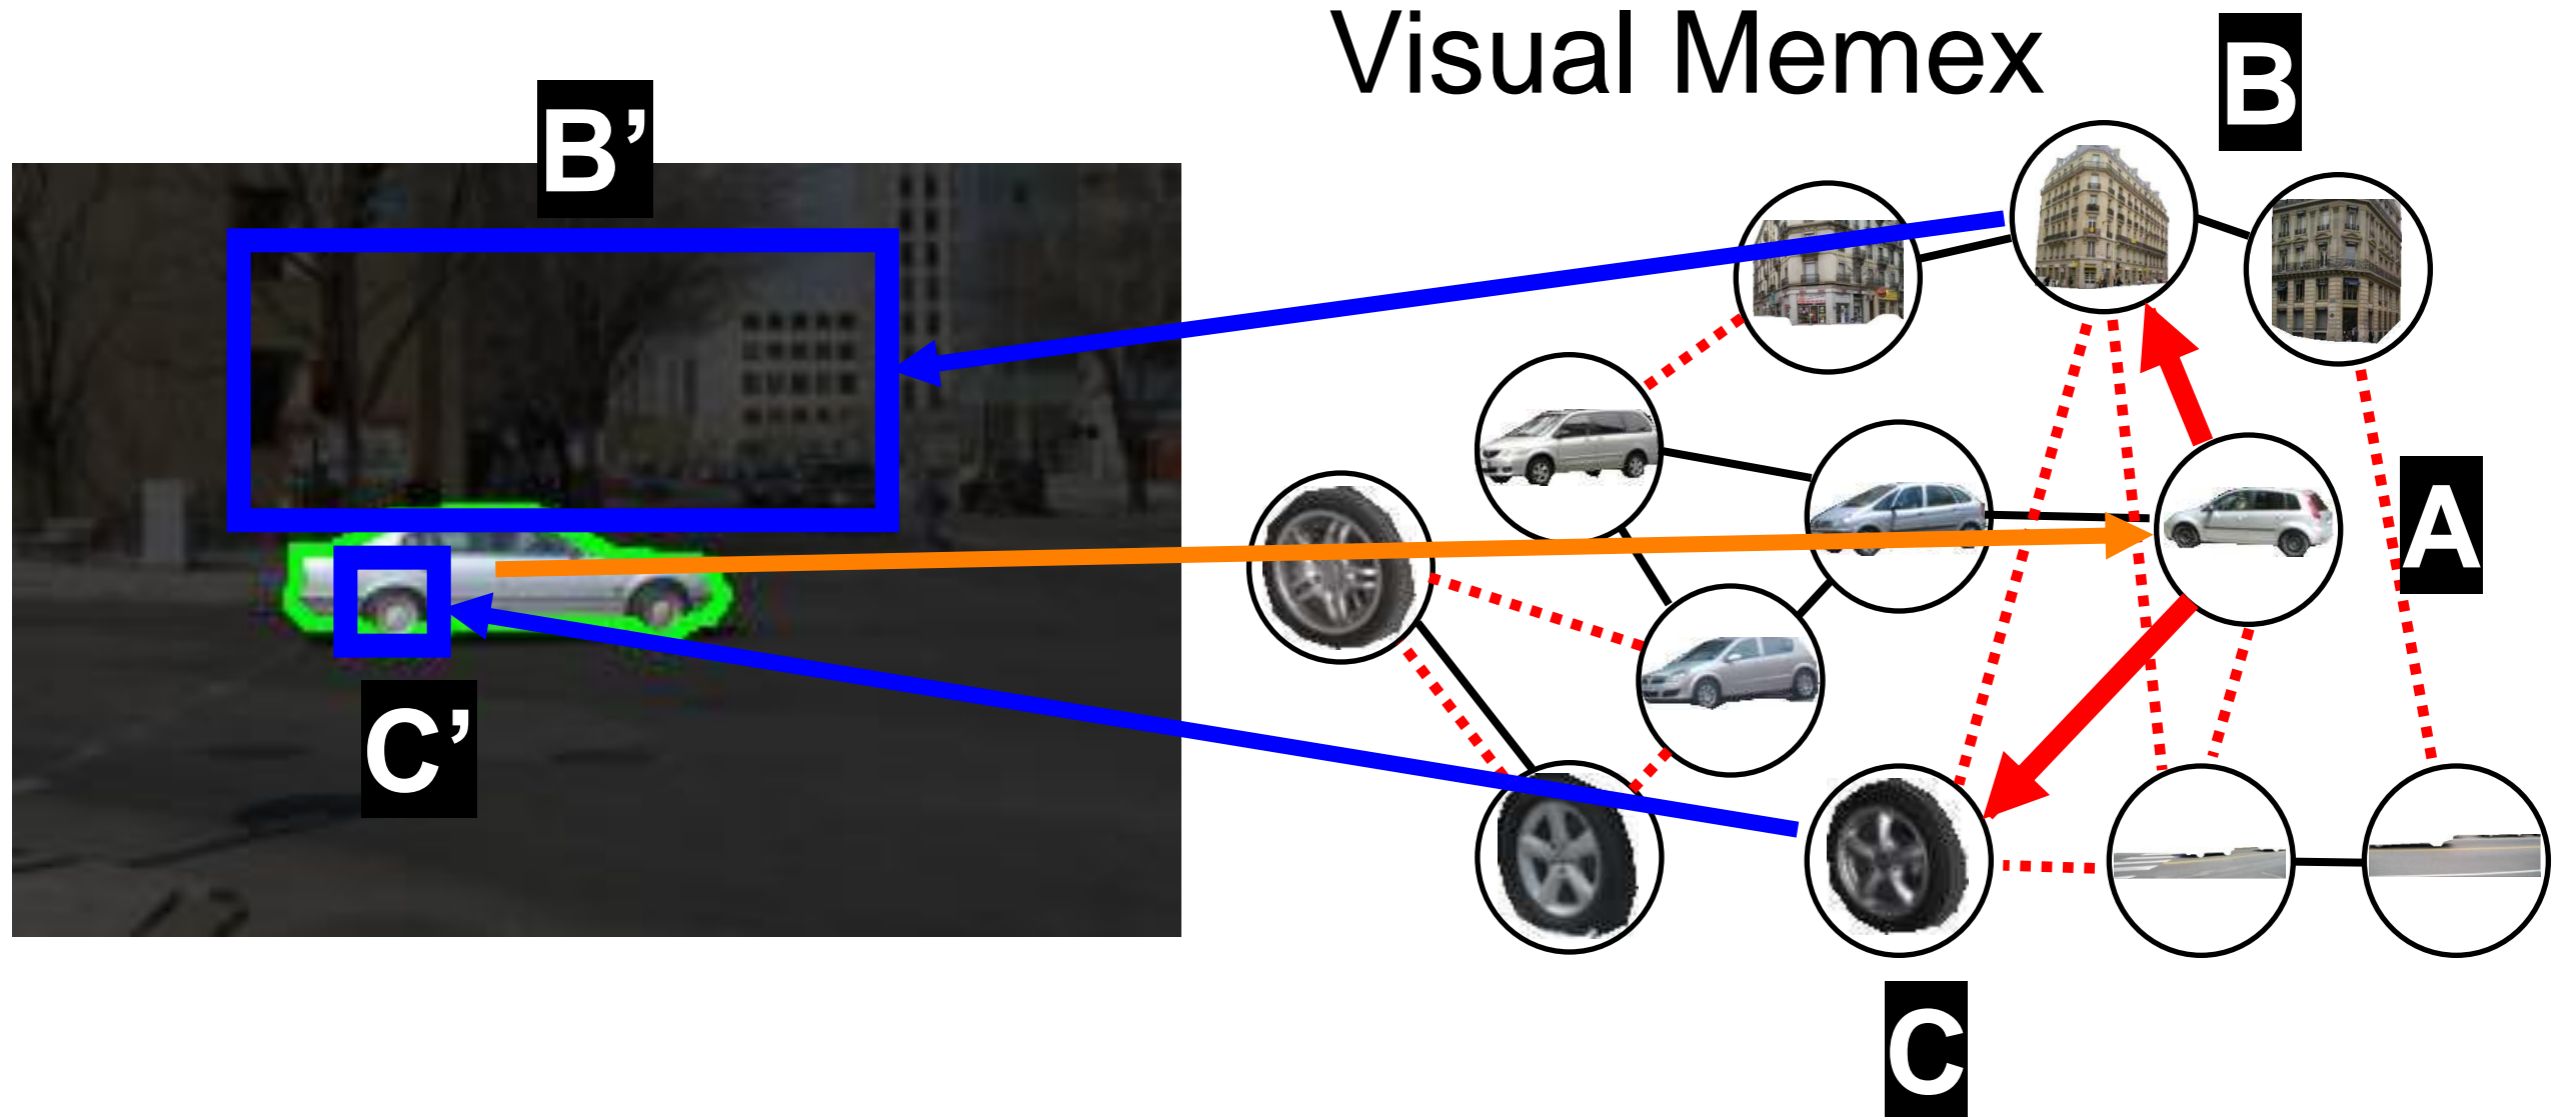


Figure 3: a.) Context Challenge confusion matrices for the 3 methods: Visual Memex, KDE, and CoLA. b.) Recognition Precision versus Recall when thresholding output based on confidence. c) Side by side comparison of the 3 methods' accuracies for 30 categories.



# Next Step: top-down segmentation

Visual Memex



# Take Home Message

- Categorization is not a goal in itself
  - Rather, it is a means for transferring knowledge onto a new instance
- Skipping explicit categorization might make things easier, not harder
  - The “harder intermediate problem” syndrome
- Keeping around all your data isn't so bad...
  - you never know when you will need it

# Questions?



+

