


The Promise and Peril of Big Data



Alexei (Alyosha) Efros
CMU (school-year), INRIA (summer)

It's hard to be a computer...



How the computer sees the world:



The Guitar Player
Pablo Picasso (1911)

Some early work...

“Data, Data, Data... Watson, I need Data!”
[Sherlock Holmes, 1886]


But If you want to publish a NIPS paper...

Data

Features

Learning Algorithm

Face Detection: Big Success Story



- Rowley, Baluja, and Kanade, 1998
- Schniderman & Kanade, 1999
- Viola & Jones, 2001

Modern Recognition is largely Data-Driven

- In non-linear SVMs:
 - In ML, people report ~10% of data are support vectors
 - In recognition, up to 2/3 of data are support vectors!!!
- In linear SVMs:
 - Typical setup: 4000 dim. HOG, only 300 “chair” examples


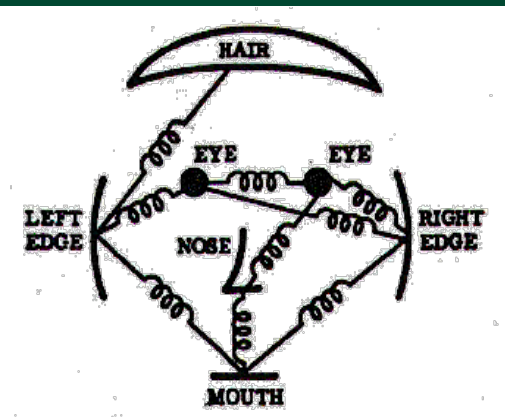



Figure that Francis Bach hates

Recognition Learning Spectrum

Extrapolation problem
Generalization

Interpolation problem
Correspondence



Slide by Antonio Torralba

Everything else being equal...

... the visual world is just much **richer!**

- MNIST Digits

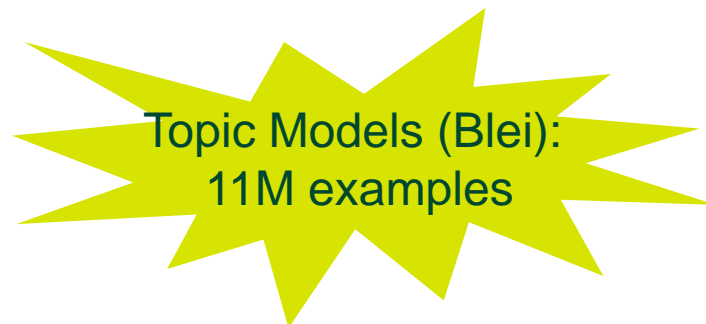
- 10 digits *
- ~1,000 variations = 10,000



MNIST:
60,000 examples

- English words

- ~100,000 words *
- ~5 variations = 500,000



Topic Models (Blei):
11M examples

- Visual world

- ~100,000 objects *
- ~10,000 variations (pose, scale, lighting, intra-category)
- = **1,000,000,000 (1 billion!)**

Yet, we train on 15 examples?!



If you want to start a company...

Data

Features

Algorithm

To make research progress...

Data / Features / Algorithms

Big Message...

Keep the data --

you never know when you will
need it!

Texture

- Texture depicts spatially repeating patterns
- Many natural phenomena are textures



radishes




rocks

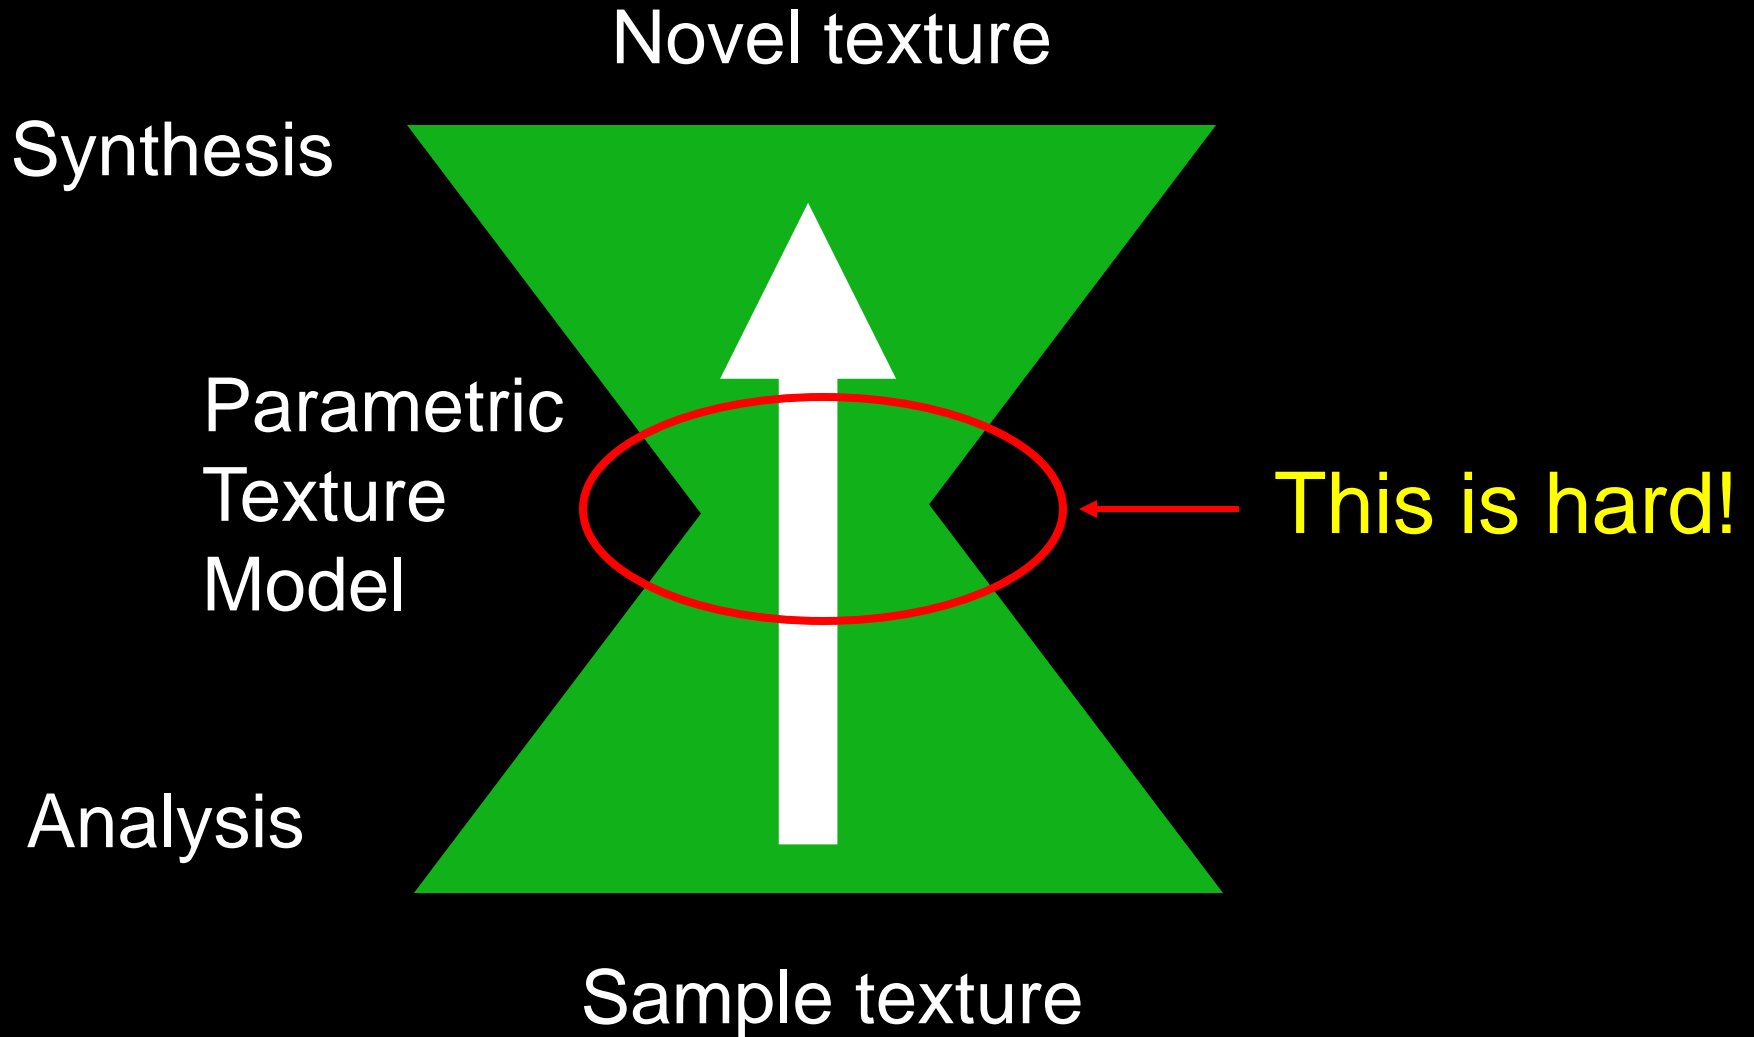


yogurt


Texture Synthesis




Classical Texture Synthesis



Throwing away too much too soon?

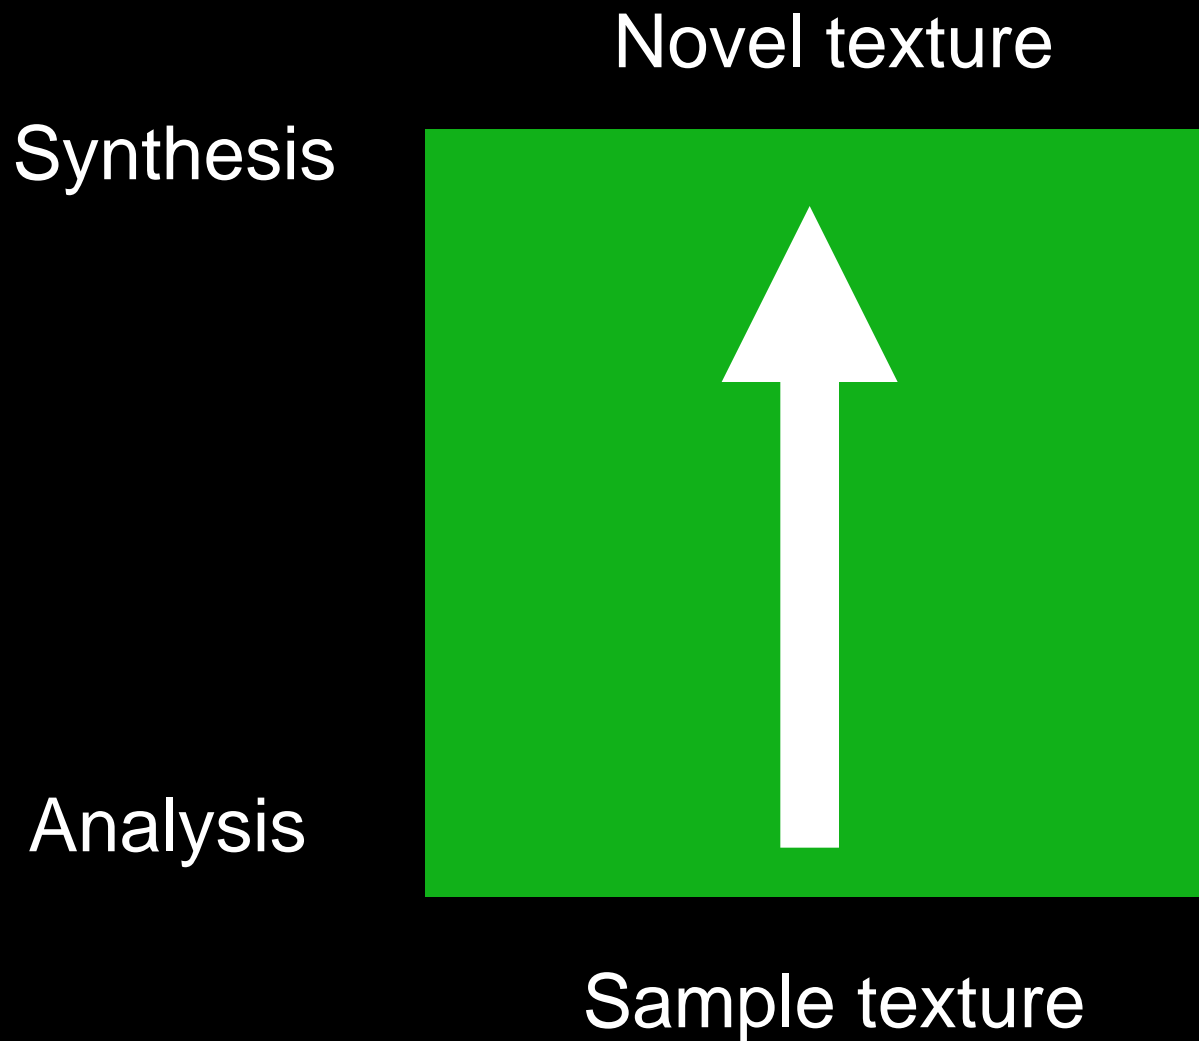


input texture



synthesized texture

Non-parametric Approach



Motivation from Language


- [Shannon, '48] proposed a way to generate English-looking text using N-grams:
 - Assume a generalized Markov model
 - Use a large text to compute prob. distributions of each letter given N-1 previous letters
 - Starting from a seed repeatedly sample this Markov chain to generate new letters
 - Also works for whole words

WE NEED TO EAT CAKE


Mark V. Shaney (Bell Labs)

- Results (using `alt.singles` corpus):
 - *“As I've commented before, really relating to someone involves standing next to impossible.”*
 - *“One morning I shot an elephant in my arms and kissed him.”*
 - *“I spent an interesting evening recently with a grain of salt”*
- Notice how well local structure is preserved!
 - Now, instead of letters let's try pixels...

[Efros & Leung, '99]




Texture Growing




Homage to Shannon

...ing in the unsensational
... Dick Gephardt was fail
...rful riff on the looming
...nly asked, "What's your
...tions?" A heartfelt sigh
...story about the emergen
...es against Clinton. "Boy
...g people about continuin
...ardt began, patiently obs
...s, that the legal system h
...g with this latest tange



...thaim. them . "Whnephartfe lartifelintomimen
...fel ck Clirtioout omaim thartfelins.f out s anento
...the ry onst wartfe lck Gephtoomimeationl sigab
...Cliocoufit Clinut Clil riff on. hat's yo'dn, parut tly
...ons yontonsteht waked, paim t sahe loo riff on
...nskoneploourtfeas leil A nst Clit, "Wheontongal s
...k Cirtioouirtfepe ong pme abegal fartfenstemem
...tiensteneltorydt telemephminsverdt was agemer
...ff ons artientont Cling peme as artfe atich, "Boui s
...nal s fartfelt sig pedr th'dt ske abounutie aboutioo
...tfeonewwas you abounthardt thatins fain, ped, '
...ains. them, pabout wasy arfuit couitly d, ln A h
...ple emthringbooreme agas fa bontinsyst Clinut
...ory about continst Clipeopinst Cloke agatiff out C
...stome zinemen tly ardt beorabol n, thenly as t C
...cons faimeme Diontont wat coutlyohgans as fan
...ien, phrtfaul, "Wbaut cout congagal comininga
...mifmst Clily abon al coountha.emungaint tfoun
...The loocrysta loontieph. intly on, theoplegatick C
...aul fatieeontly atie Dioniomt wal s f thegae ener
...mthahgat's enephhmas fan. "intchthory abons w




input image




Portilla & Simoncelli



Xu, Guo & Shum



Wei & Levoy




Our algorithm

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](#)



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 dnskct



From Lord V



From ConnieFK.....



From e_monk



From RaeA



From mark...



From calbergemin...



From SixRevisions



From GrungeTextur...



From SixRevisions

“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



The Good News

Really stupid algorithms + Lots of Data
= *“Unreasonable Effectiveness”*


The Bad News

Visual Data is much more difficult


- text:
 - clean, segmented, compact, 1D
- Visual data:
 - Noisy, unsegmented, high entropy, 2D/3D

Distance Metrics


CLIME - *CRIME* = hamming distance of 1 letter



= Euclidian distance of 5 units



= Grayvalue distance of 50 values



= ?

L2 norm says these are not similar



?




Make them tiny!



?



Lots of Tiny Images



- 80 million tiny images: a large dataset for non-parametric object and scene recognition
Antonio Torralba, Rob Fergus and William T. Freeman. PAMI 2008.

32x32



office






waiting area



dining room




dining room



c) Segmentation of 32x32 images

Human Scene Recognition



a) Scene recognition

Lots Of Images


Target



7,900



Lots Of Images



Lots Of Images

Target



7,900



790,000



79,000,000



Automatic Colorization

Grayscale input High resolution



Colorization of input using average



Not a pixel lover? No problem!



Let's match gradients



[Hays and Efros. Scene Completion Using Millions of Photographs. SIGGRAPH 2007 and CACM October 2008.]





Efros and Leung result



Scene Matching for Image Completion






Scene Completion Result



The Algorithm





Scene Matching



Scene Descriptor





Scene Descriptor





Scene Gist Descriptor
(Oliva and Torralba 2001)

Scene Descriptor





Scene Gist Descriptor
(Oliva and Torralba 2001)

Scene Descriptor




2 Million Flickr Images

The background of the slide is a dense, colorful mosaic of small image thumbnails, representing 2 million Flickr images. The thumbnails are arranged in a grid-like pattern, creating a complex, multi-colored texture. The colors are varied, including shades of blue, green, red, yellow, and grey, reflecting the diverse content of the images. The overall effect is a rich, textured background that visually represents the vast amount of data being discussed.



... 200 total

Context Matching





Graph cut + Poisson blending





















... 200 scene matches



















Why does it work?





Nearest neighbors from a collection of 20 thousand images



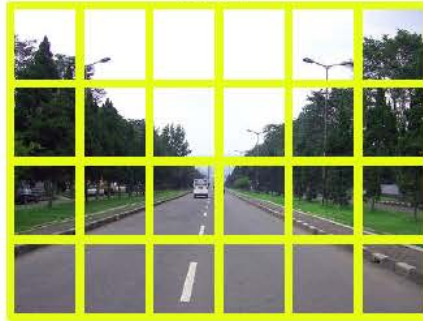
Nearest neighbors from a collection of 2 million images

Scene matching with camera transformations

Query image



GIST



Best match




Top matches




Query image



Camera rotation & GIST



Best match after rotation



Top matches



Sivic, Kaneva, Torralba, Avidan, Freeman, Internet Vision Workshop, 2008
updated version to appear in Proceedings of the IEEE (2010)

Creating and exploring a large, photorealistic virtual space

Josef Sivic (INRIA/ENS), Biliana Kaneva (MIT),
Antonio Torralba (MIT), Shai Avidan (Adobe) and
Bill Freeman (MIT)

IEEE Workshop on Internet Vision, 2008

Cross-Domain Matching

*“Local – bad
Global – good!”*
(after Orwell)



Medici Fountain, Paris



Search by image



[→ Move](#)

Drop image here

[Watch a short video to learn more.](#)

Search

About 2 results (0.29 seconds)

Everything

Images

Maps

Videos

News

Shopping

More



Image size:
1024 × 829

No other sizes of this image found.

Visually similar





Medici Fountain, Paris (winter)



Search

About 2 results (0.29 seconds)

Everything

Images

Maps

Videos

News

Shopping

More



Image size:
713 × 600

No other sizes of this image found.

Visually similar







Search

About 2 results (0.29 seconds)

Everything

Images

Maps

Videos

News

Shopping

More



Image size:
319 x 482

No other sizes of this image found.

Visually similar







Search

About 2 results (0.29 seconds)

Everything

Images

Maps

Videos

News

Shopping




More




Image size:
443 × 482

No other sizes of this image found.

Visually similar




OUR GOAL [SIGGRAPH ASIA'11]




WHY IS THIS SO HARD?




EXAMPLE: SIFT MATCHING





Example: SIFT Matching







Input Query



Top Matches

Input Query



Top Matches


Input Query



Top Matches

IMPORTANT PARTS?

Input Query



Important Parts




Input Query




Top Matches





“Data-driven Uniqueness”









Possible Explanation

- If the space of images was uniform, nearest neighbor would work perfectly well
- But the space is very non-linear, non-Euclidian
- The Exemplar-SVM is trying to make a small, linear rescaling of visual space, near the query point
 - i.e. capturing the natural image statistics near the exemplar
 - Or maybe global is good enough (see Deva)


FEATURE REPRESENTATION

HISTOGRAM OF ORIENTED GRADIENTS (HOG)






Input Query



HOG



Top Match




Learnt Weights



Top Match

SEARCH USING IMAGES


Input Query



Top Matches

SEARCH USING IMAGES


Input Query



Top Matches

SEARCH USING IMAGES

Input Query




Top Matches

SEARCH USING PAINTINGS




Input Painting



GIST




Bag-of-Words



Our Approach



Tiny Images



HOG

SEARCH USING PAINTINGS



Input Painting



Top Matches

SEARCH USING PAINTINGS




Input Painting



Top Matches

SEARCH USING SKETCHES



Input Sketch




Tiny Images



GIST



Bag-of-
Words





HOG



Our Approach

SEARCH USING SKETCHES



SALIENCY




PROXY FOR SALIENCY?





PREDICTING SALIENCY

SALIENCY DATASET [Judd et al., 2009]



WHERE DOES IT FAIL?



Top Matches

APPLICATIONS

WHERE WAS THE PAINTER STANDING?

Input Painting



PAINTING2GPS

Input Painting



Retrieval set

10,000 Geo-tagged Flickr Images

100 top matches used to estimation

PAINTING2GPS

Input Painting




Estimated Geo-location



Estimated using 100 top matches

PAINTING2GPS

Input Painting



Sydney Opera House

VISUAL SCENE EXPLORATION

VISUAL SCENE EXPLORATION

flickr

Sign up or Admin: Show options

Home You **Organize & Create** Contacts Groups Explore Upload

Search

Search

Photos Groups People

Feyyaz's Uploads

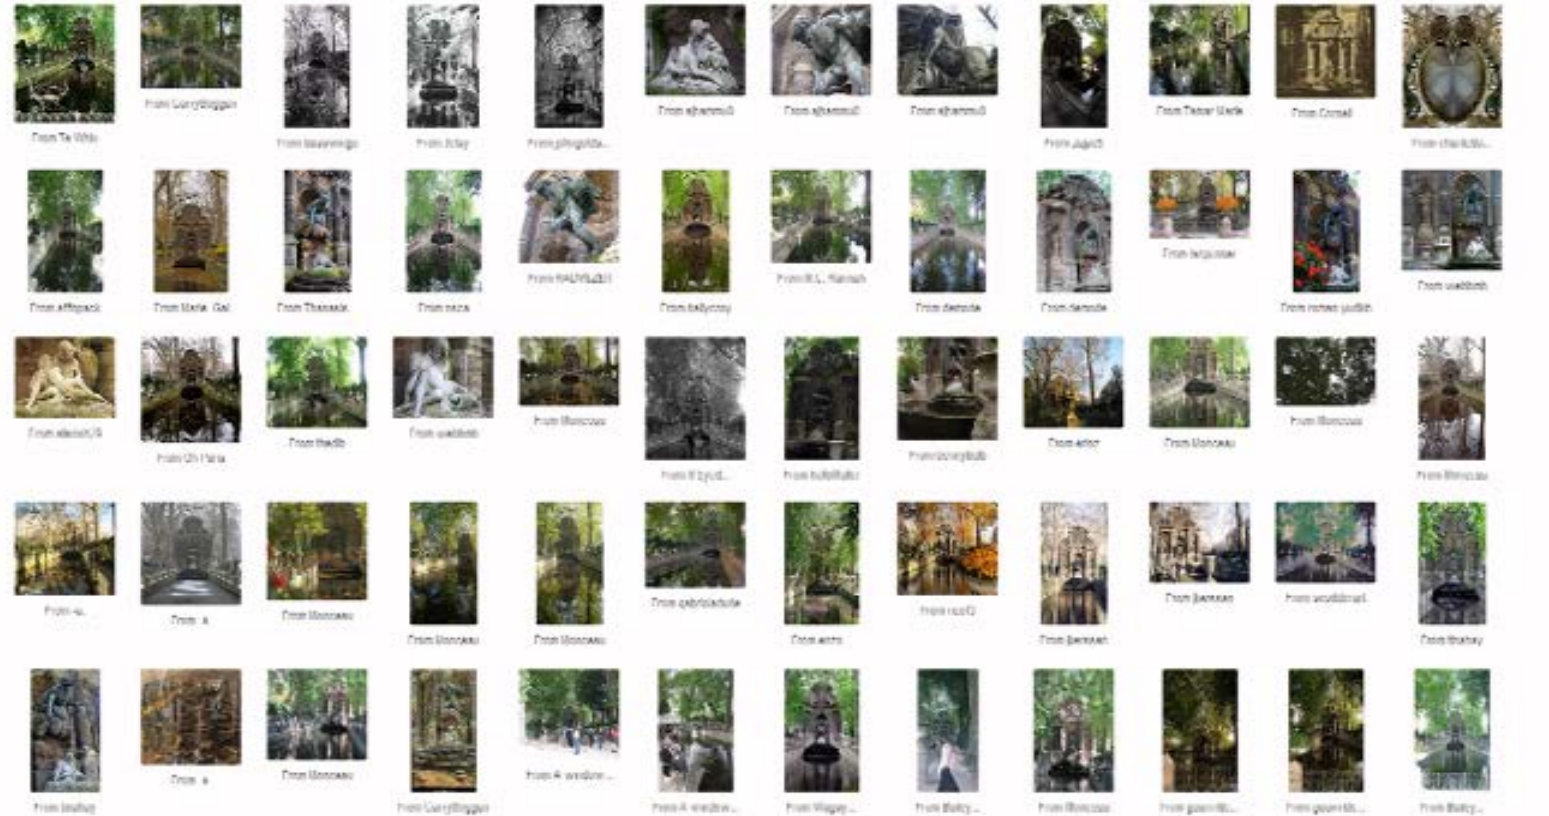
medici fountain

SEARCH

Full Text | Text Only
Advanced Search


Sort: **Relevant** Recent Interesting

View: **Small** Medium Detail Slideshow



1 2 3 4 5 6 7

(177 results)



What to go with? Get that!

Groups

- Florence - Florence**
1,313 members | 16,370 photos
- Paris**
25,034 members | 244,343 photos

Photographers

- Emilio Secchi** - Curcio, United States
See related photos
- wally** - Jersey City, NJ, US
See related photos

Tags

- Florence** - Photos with tags like florence, italy and fountain
- Fountain** - Photos with tags like park, fountain and fountain

Looking to license a photo?

Looking to license an image? You can search through the licensed image of Flickr members enrolled in the Getty Images Program.

Photos

more...

Organizing Visual “Garbage Heap”



“It irritated him that the “dog” of 3:14 in the afternoon, seen in profile, should be indicated by the same noun as the dog of 3:15, seen frontally...”

“My memory, sir, is like a garbage heap.”

*-- Jorge Luis Borges, *Funes the Memorious**

PHOTOSYNTH

[Snavely et al., 2006]



Synth 1



Synth 2



Synth 3

...






Dataset size: 136 photos (from flickr)

of discovered synths: 14

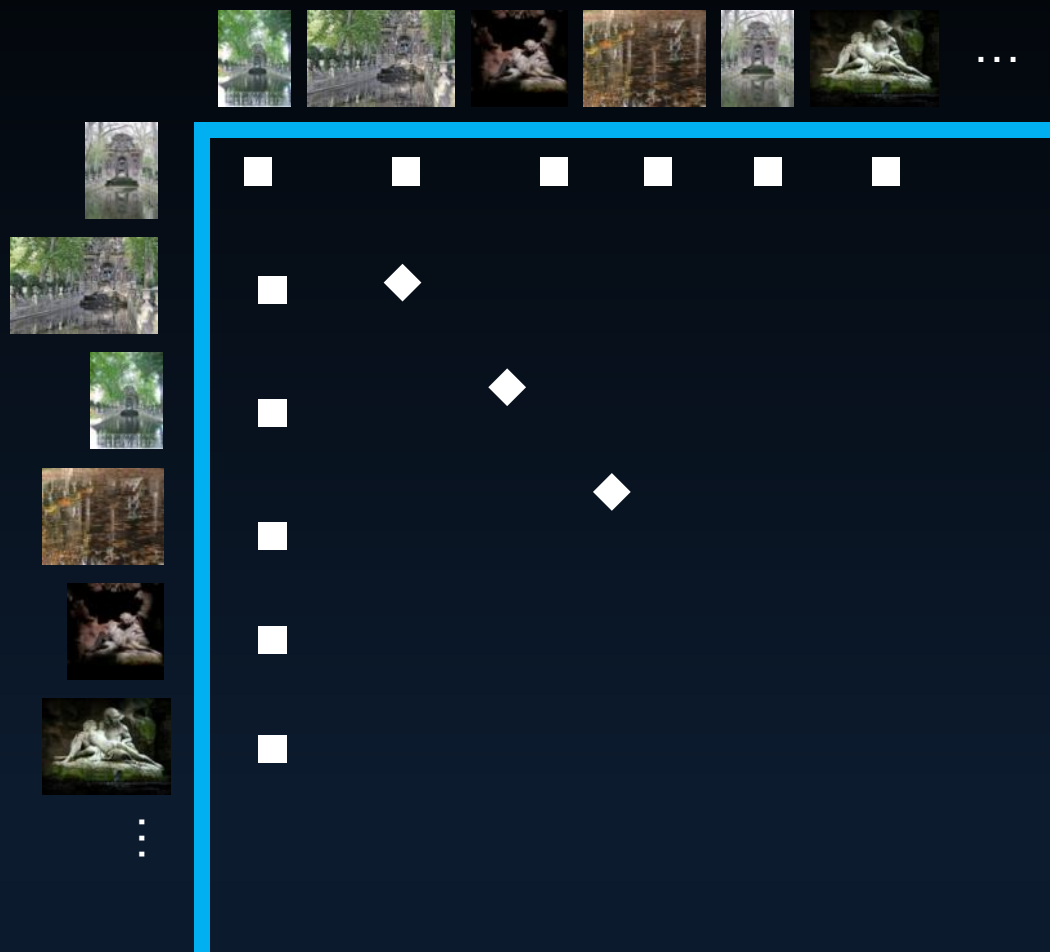
82 photos not part of any synth

FINDING SIMILAR IMAGES


Query image



PAIRWISE SIMILARITY MATRIX

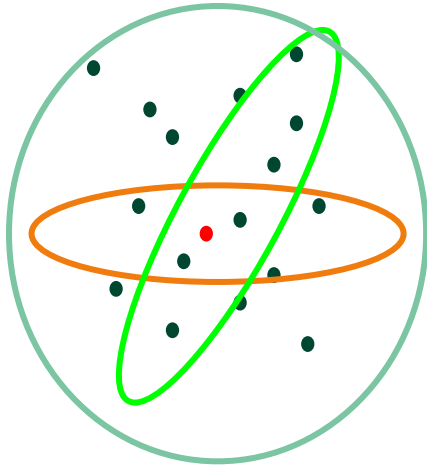


TRAVERSING THE GRAPH







Ways to use Big Data




1. See what different subsets of data think of you



2. Use kNN to make a smaller, cozier sub-problem



3. Visual Data Mining: find needles in a haystack



4. Ditch Categories – keep all instances and connect them

Priors for Large Photo Collections & What they Reveal about Cameras

Sujit Kuthirummal

Columbia University

Aseem Agarwala

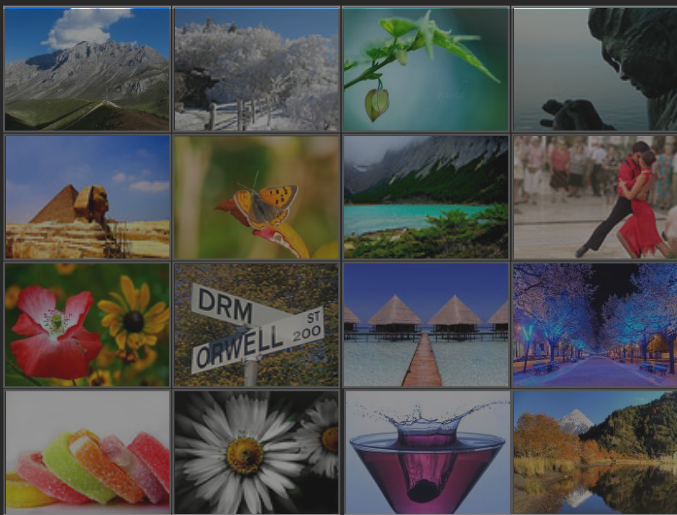
Adobe Systems, Inc.

Dan B Goldman

Adobe Systems, Inc.


Shree K. Nayar

Columbia University



Camera Distortion Free

— Compute
Aggregate
Statistic —>




Independent of
Scenes, Photographers
& Cameras



One Camera's Distortion

— Compute
Aggregate
Statistic —>




Independent of
Scenes & Photographers
Dependent on
Camera


Recover
Camera Properties




Ways to use Big Data



1. See what different subsets of data think of you



2. Use kNN to make a smaller, cozier sub-problem



3. Visual Data Mining: find needles in a haystack



4. Ditch Categories – keep all instances and connect them

1. kNN + Label Transfer



Sky, Water, Hills, Beach,
Sunny, mid-day




80 Million Tiny Images [PAMI'08]




a) Input image



b) Neighbors

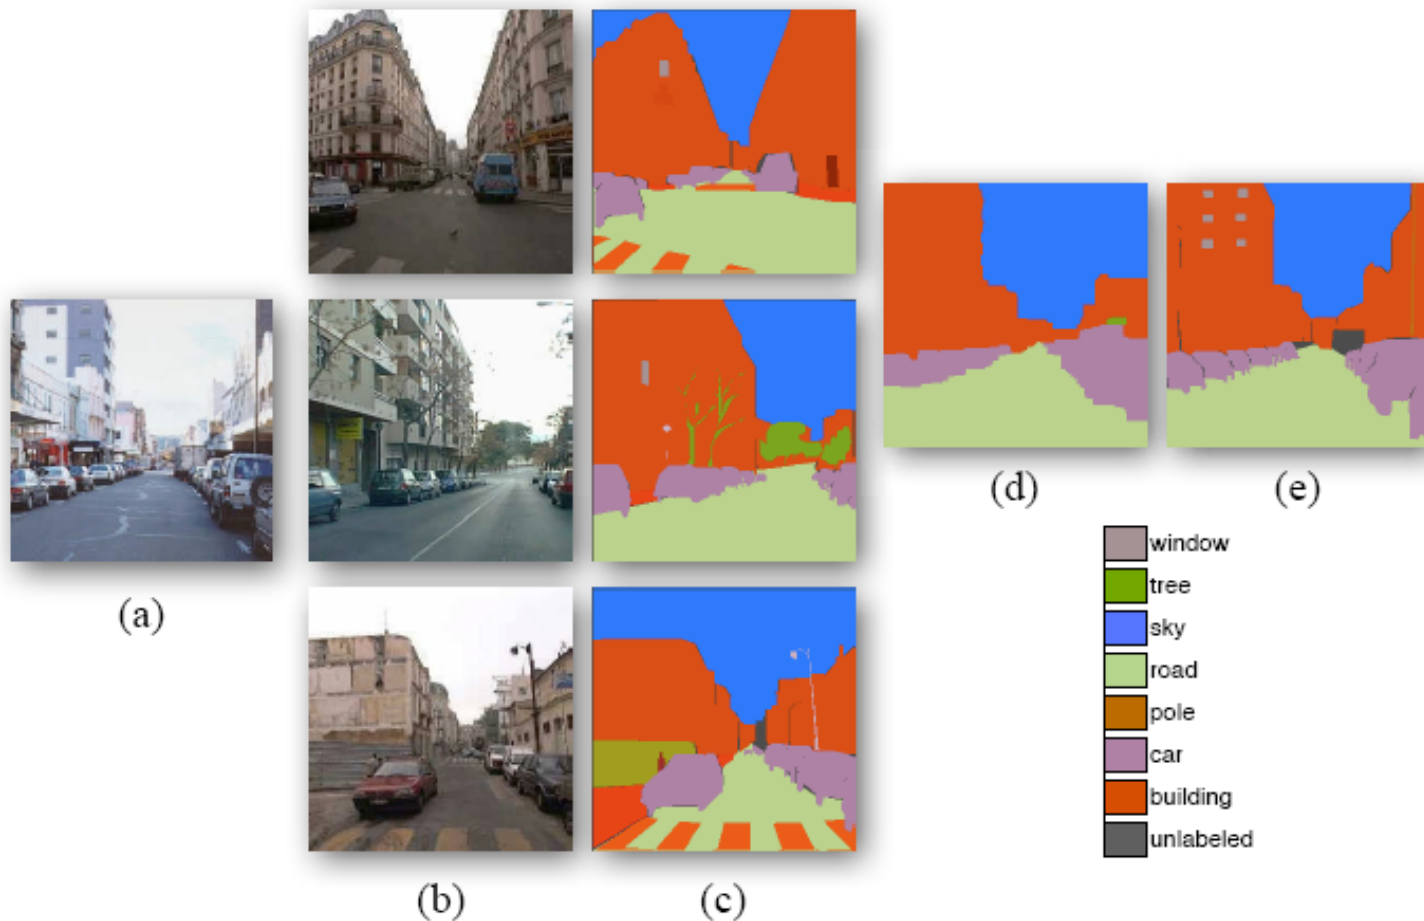


c) Ground truth



d) Wordnet voted branches

Non-parametric Scene Parsing [CVPR'09]

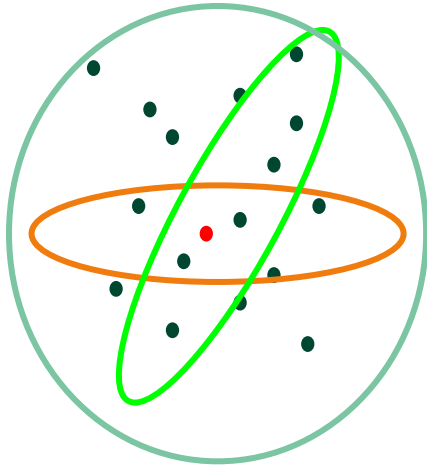


im2gps [CVPR'08]




Query Photograph


Ways to use Big Data




1. See what different subsets of data think of you



2. Use kNN to make a smaller, cozier sub-problem




3. Visual Data Mining: find needles in a haystack



4. Ditch Categories – keep all instances and connect them

Visual Words or Letters?



Spectrum of Visual Features

Low-Level

High-Level



Pixel

Filter-Banks

Sparse-SIFT

Parts,
Segments

Objects

Image



Visual Words



Our Approach


Discriminative Patches

Two key requirements


1. Need to occur frequently (representative)
2. ...but not **too** frequently

Discriminative: Need to be different enough from the rest of the visual world.


First some examples




Finding needles in a haystack




K-Means Clusters



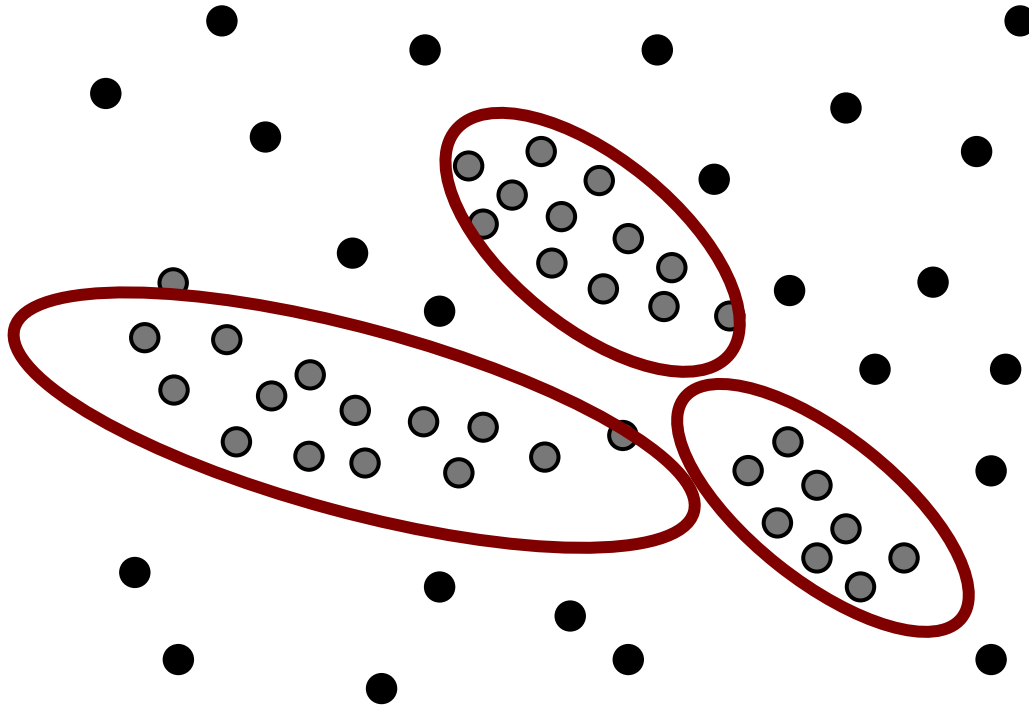
K-Means



Discriminative K-means



Discriminative Clustering+



Discriminative Clustering+

Initial



Final



Initial




Final



Discriminative Clustering++

KMeans



Iter 1



Iter 2



Iter 3



Iter 4



Discriminative Clustering++

KMeans



Iter 1



Iter 2




Iter 3



Iter 4




More Discovered Patches



Church




Meeting Room



Bathroom




Bakery




Bookstore




Bowling




Closet



Dining Room




Grocery Store




Auditorium



Staircase



Hair Salon



Laundromat



Office



Shoe Shop




Game Room



Subway



Wine Cellar



What makes Paris look like Paris?



Carl Doersch, Saurabh Singh, Abhinav Gupta,
Josef Sivic and Alexei Efros, [SIGGRAPH '12]

How well can people do?

- <http://baikal.graphics.cs.cmu.edu/cdoersch/im2gps2/corr/test2.html>

How well can people do?

- <http://baikal.graphics.cs.cmu.edu/cdoersch/im2gps2/corr/test2.html>
- Mean performance = 79%
 - Subjects who have been to Paris
(up to 90% if allowed to scrutinize the images)

What makes Paris look like Paris?



What makes Paris look like Paris?



Goal

To automatically discover geo-informative visual elements, that (hopefully) capture the “look and feel” of a place



Goal

To automatically discover geo-informative visual elements, that are:

- **Representative:** frequently occur in Paris.
- **Discriminative:** occur in Paris but not at other places.

Need both conditions

Discriminative only:



Need both conditions

Frequently occurring only:







Extracted Visual Elements from Prague




Extracted Elements from London



Extracted Elements from Barcelona



Extracted Elements from SF



Extracted Elements from Boston

Correspondence






Paris, France

Prague, Czech Republic

London, England

Many Elements Capture Context




Mapping architectural features




Figure 6: Examples of geographic patterns in Paris (shown as red dots on the maps) for three discovered visual elements (shown below each map). Balconies with cast-iron railings are concentrated on the main boulevards (left). Windows with railings mostly occur on smaller streets (middle). Arch supporting columns are concentrated on Place des Vosges and the St. Germain market (right).


Ways to use Big Data




1. See what different subsets of data think of you



2. Use kNN to make a smaller, cozier sub-problem



3. Visual Data Mining: find needles in a haystack



4. Ditch Categories – keep all instances and connect them


Down with Categories!!!



Alexei (Alyosha) Efros
CMU


Joint work with
Tomasz Malisiewicz

Acknowledgements



Murphy

Big Book of Concepts



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



Object naming -> Object categorization



Object categorization

sky

building

flag

face

banner

wall

street lamp

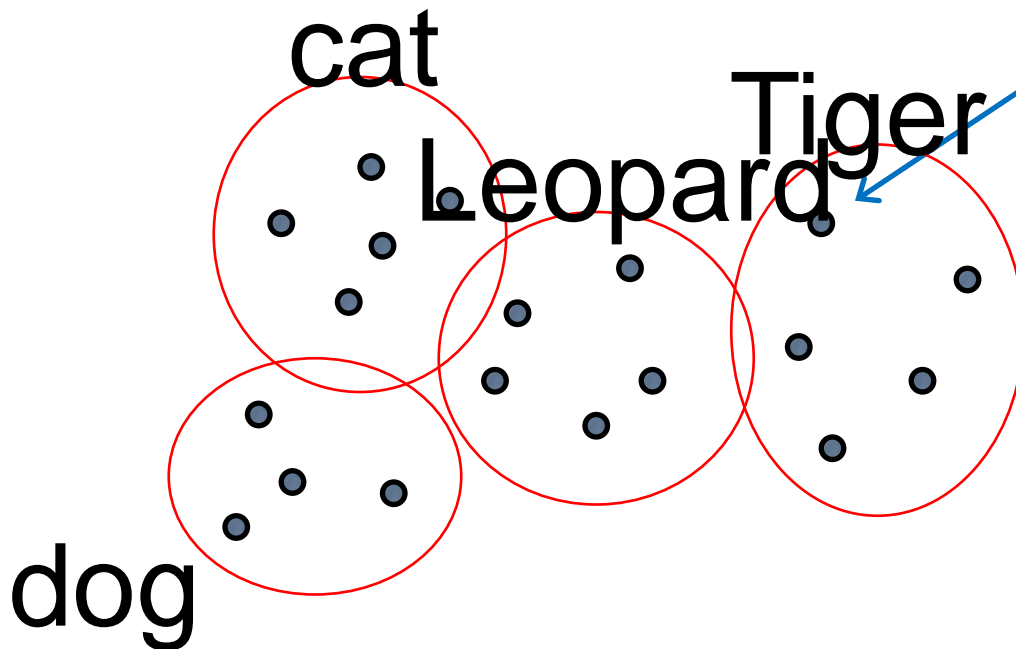
bus

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 2

Why not?



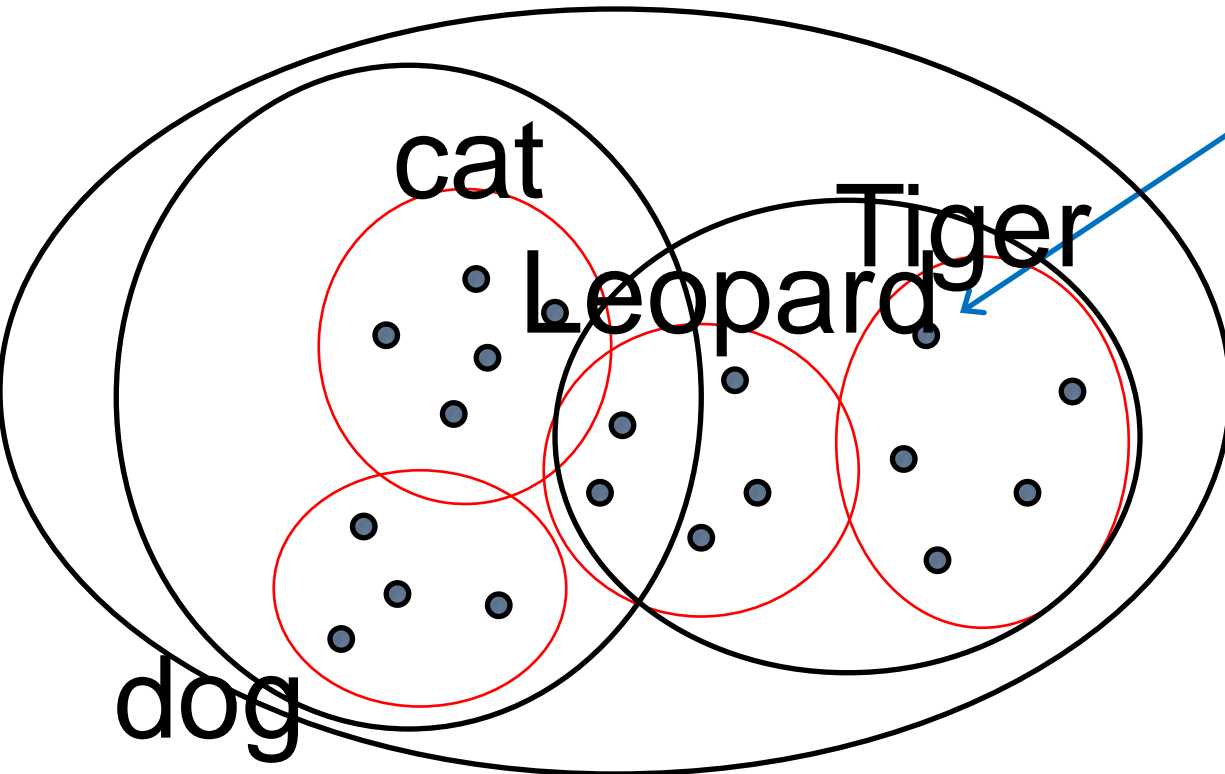
+



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!



Clay Shirky, "Ontologies are Overrated"

Fundamental Problem with Categorization



Making decisions too early!

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

The Dictatorship of Librarians



categories are losing...


The logo for Yahoo!, featuring the word "YAHOO!" in a bold, red, serif font with a slight shadow effect.

vs.

The logo for Google, featuring the word "Google" in its signature multi-colored font (blue, red, yellow, blue, green, red) with a slight shadow effect.

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)

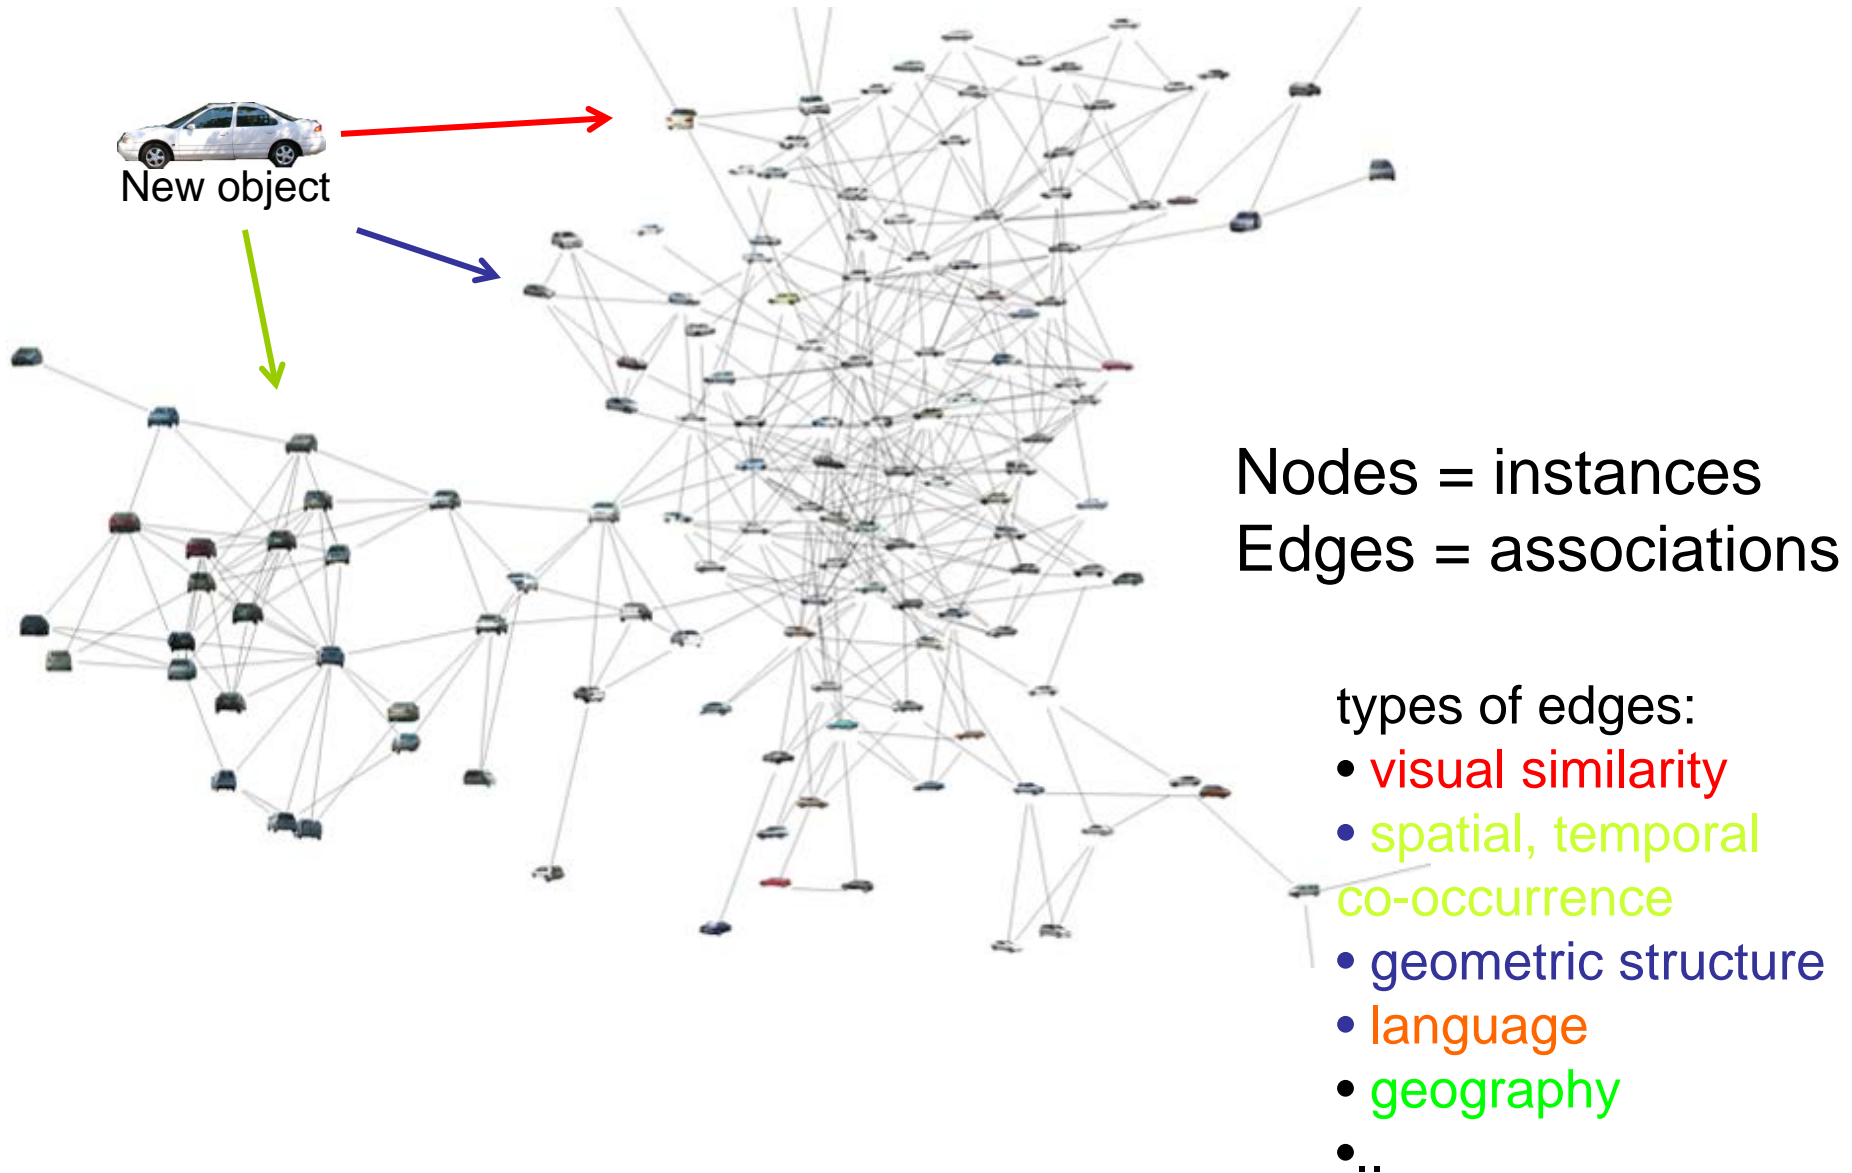


- Store publications, correspondence, personal work, on microfilm
- Items retrieved rapidly using index codes
 - Builds on “rapid selector”
- Can annotate text with margin notes, comments
- Can construct a *trail* through the material and save it
 - Roots of hypertext
- Acts as an external memory




Visual Memex, a proposal

[Malisiewicz & Efros]

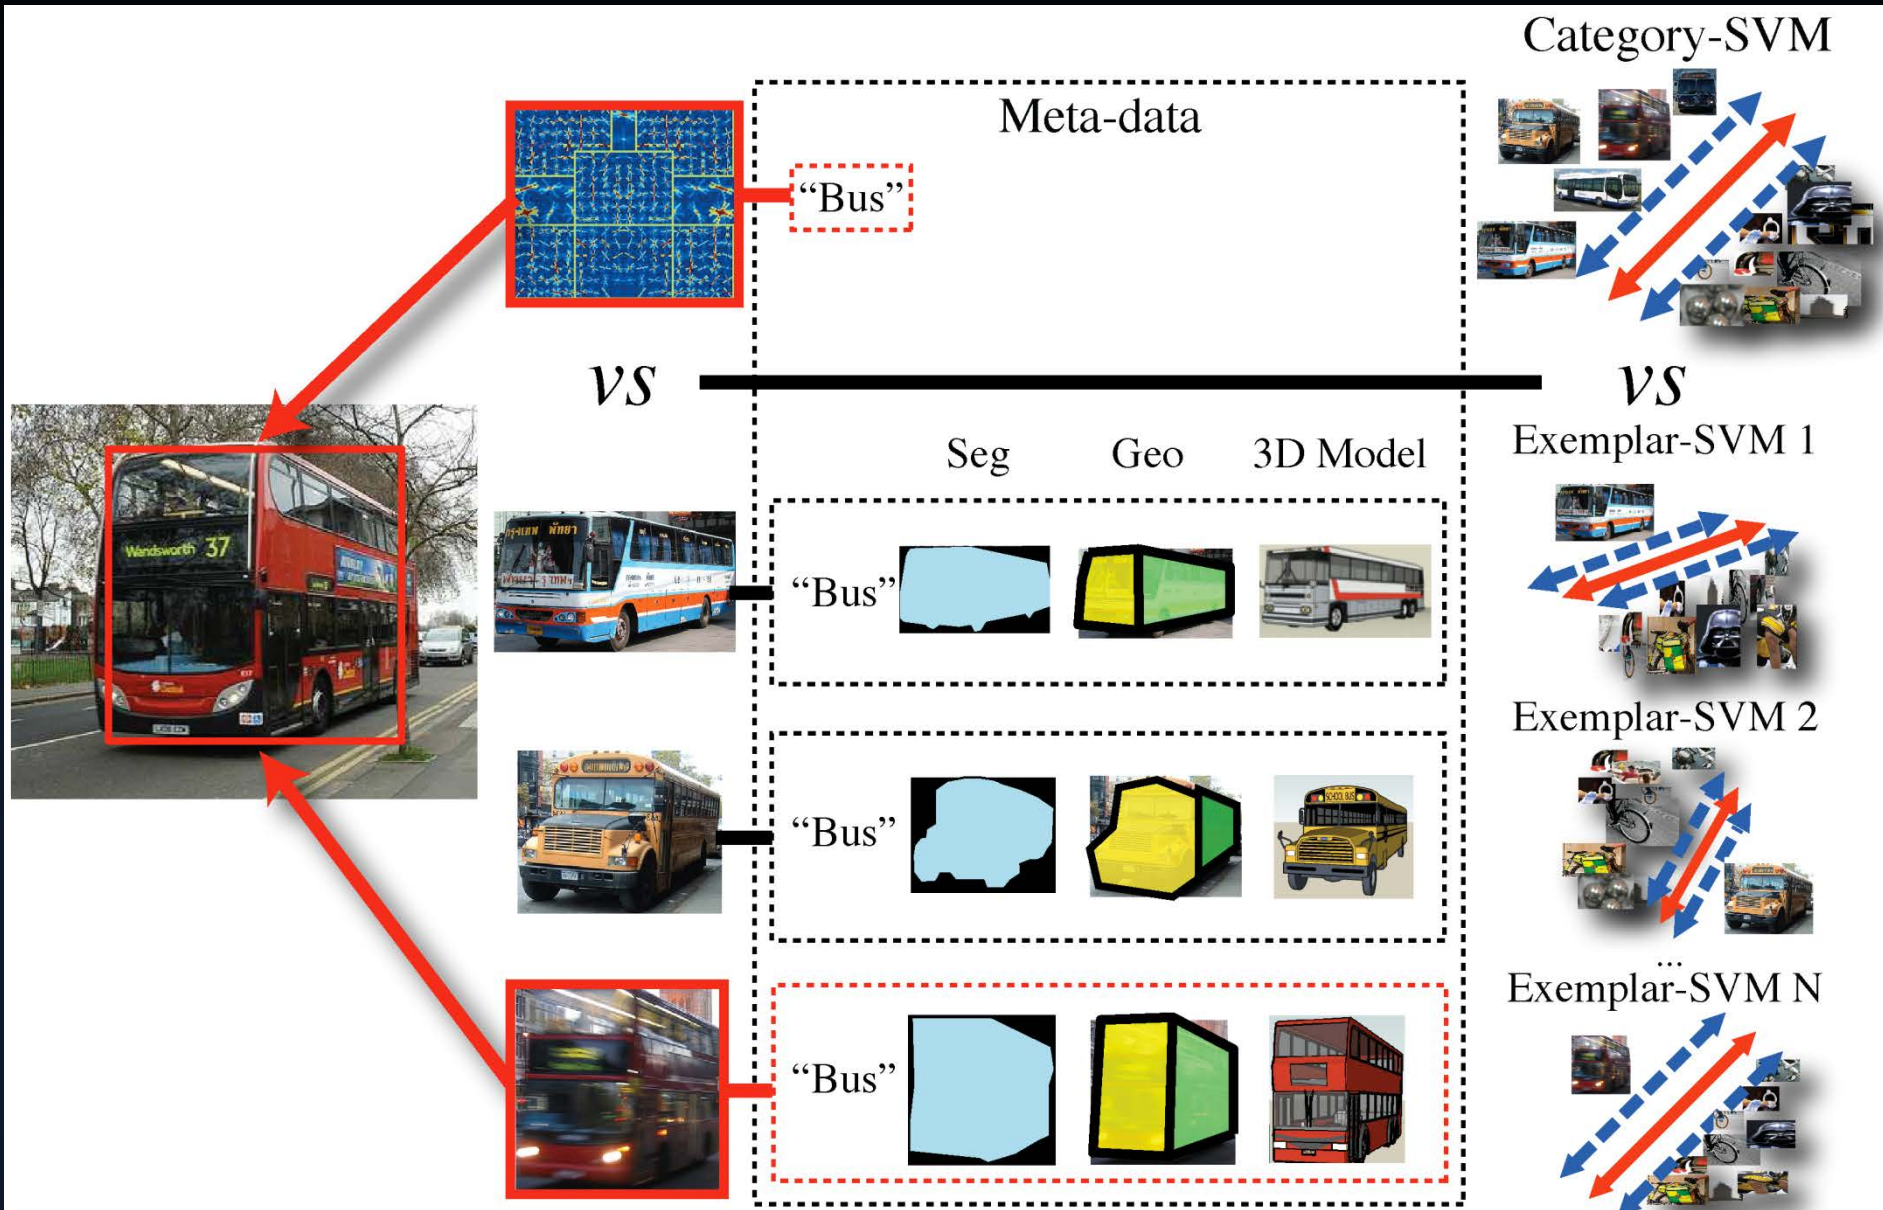


Object Detection




“bus”

Exemplar Object Detection




Ensemble of Exemplar-SVMs

Exemplar-SVM 1




Exemplar-SVM 2




...

Exemplar-SVM N




“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

Image Parsing with Context

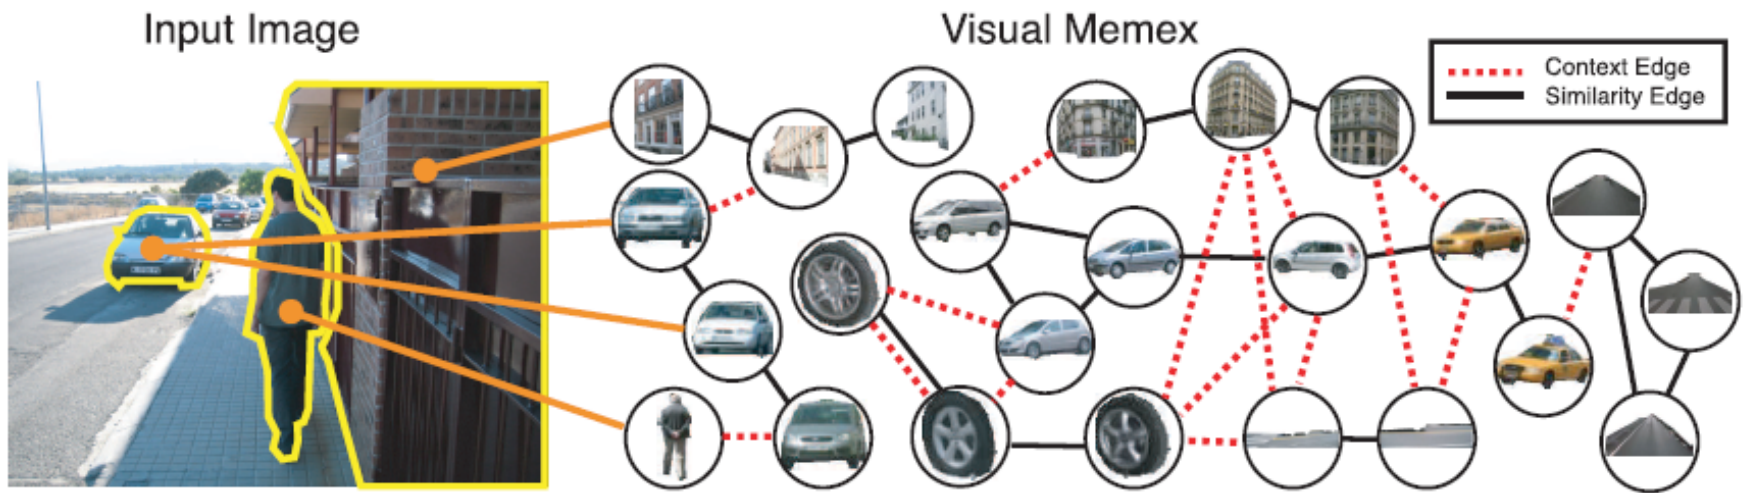


Figure 1: The **Visual Memex** graph encodes object similarity (solid black edge) and spatial context (dotted red edge) between pairs of object exemplars. A spatial context feature is stored for each context edge. The Memex graph can be used to interpret a new image (left) by associating image segments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

Torralba's Context Challenge

Torralba's Context Challenge



Torralba's Context Challenge



Slide by Antonio Torralba

Our Challenge Setup

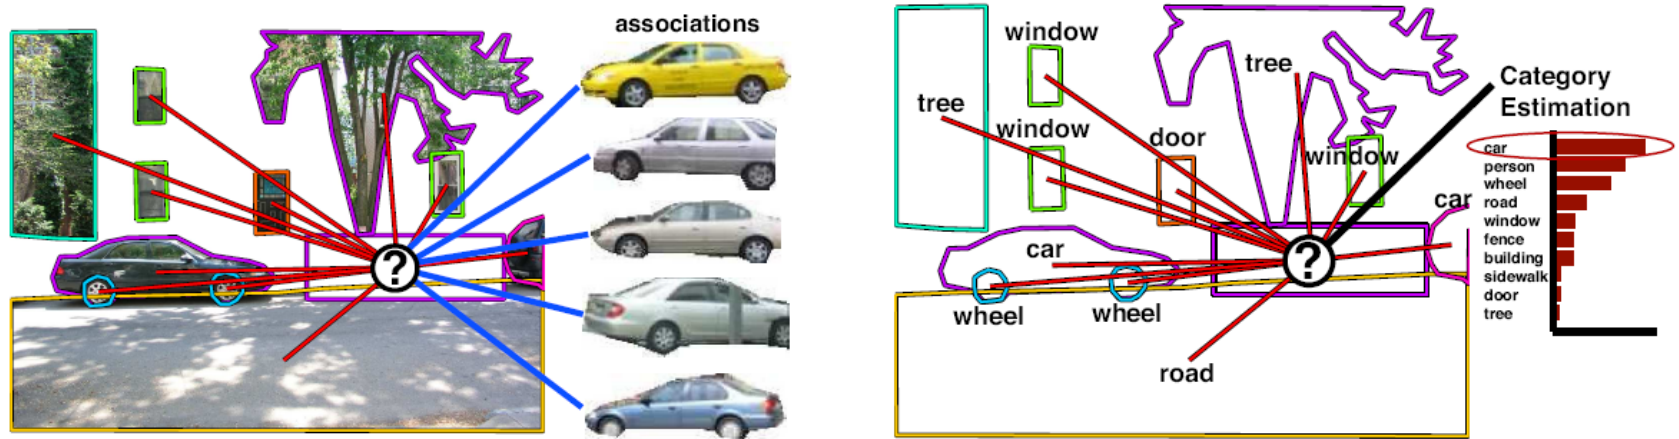


Figure 2: Torralba's Context Challenge: "How far can you go without running a local object detector?" The task is to reason about the identity of the hidden object (denoted by a "?") without local information. In our category-free Visual Memex model, object predictions are generated in the form of exemplar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


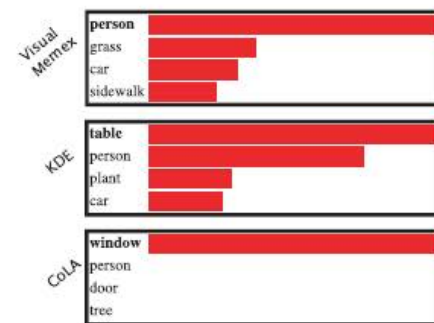
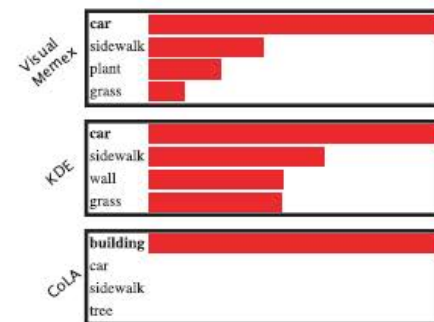



Visual Memex Exemplar Predictions



Categorization Results





Quant. results

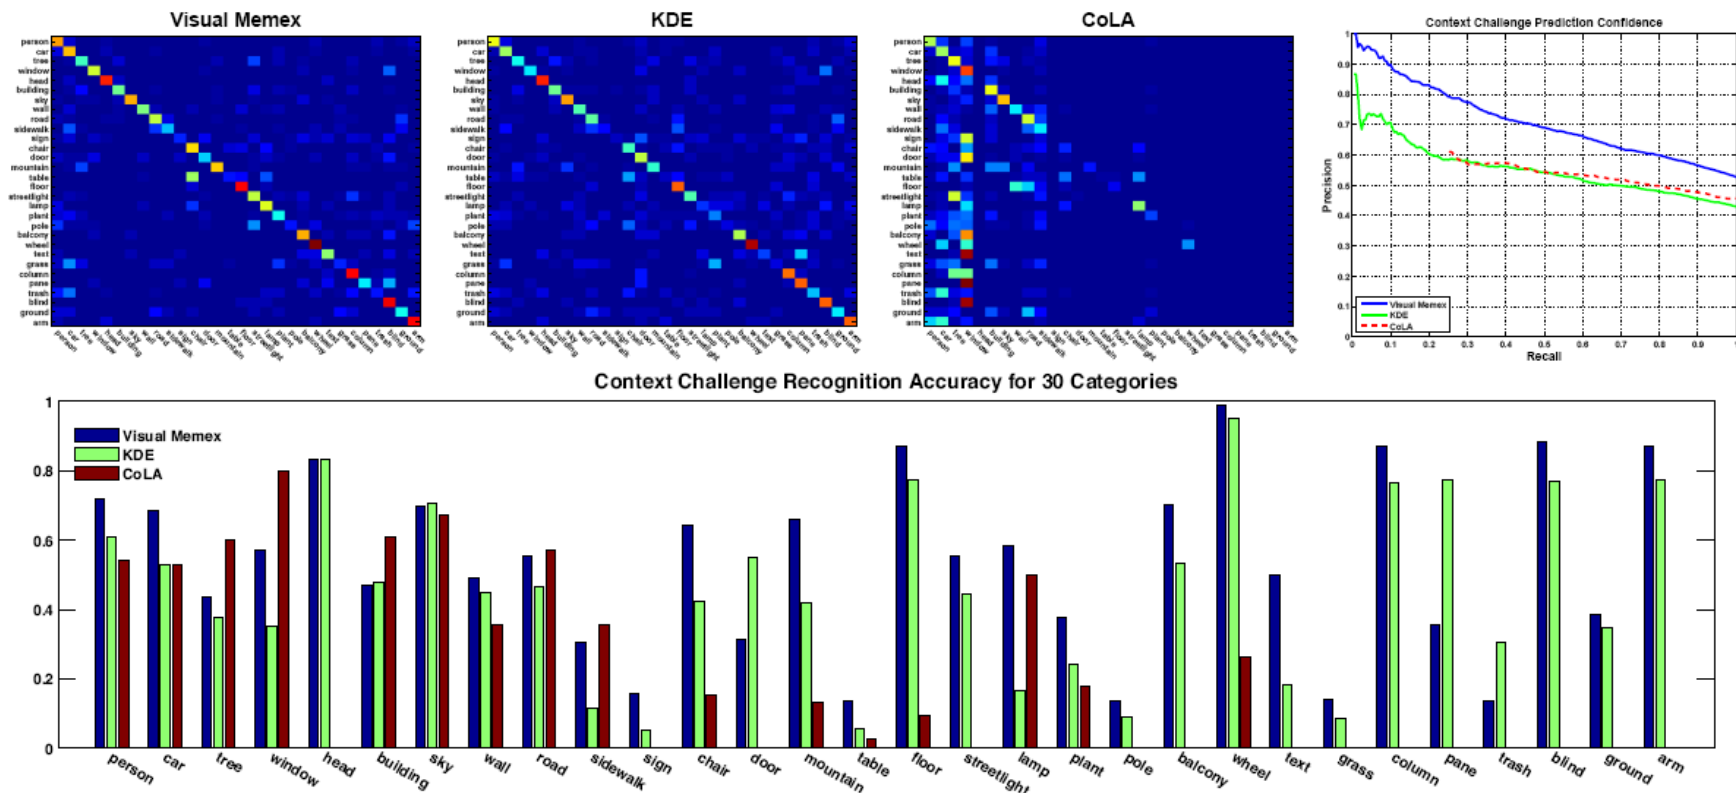


Figure 3: a.) Context Challenge confusion matrices for the 3 methods: Visual Memex, KDE, and CoLA. b.) Recognition Precision versus Recall when thresholding output based on confidence. c.) Side by side comparison of the 3 methods' accuracies for 30 categories.


**will Big Data solve all your
problems?**

1. Data is Biased


- Internet is a tremendous repository of visual data (Flickr, YouTube, Picassa, etc)
- But it's not random samples of visual world



Flickr Paris





Google StreetView Paris



My Paris



Real Notre Dame




Sampling Bias

- People like to take pictures on vacation



Photographer Bias

- People want their pictures to be recognizable and/or interesting



vs.



Social Bias




Little Leaguer



Kids with Santa



The Graduate



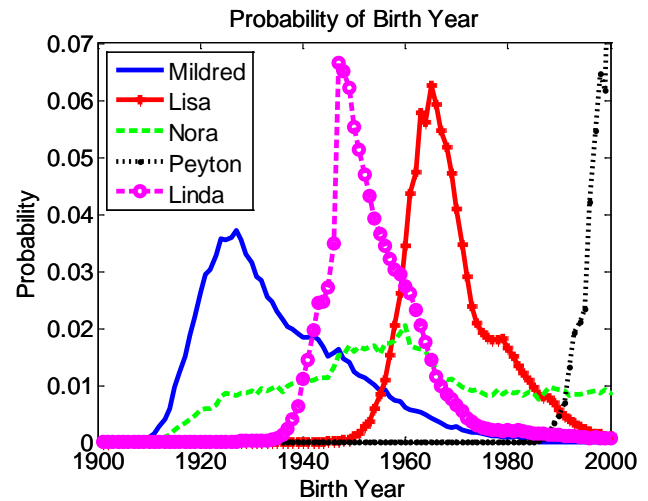
Newlyweds

“100 Special Moments” by Jason Salavon

Social Bias



Mildred and Lisa



Source: U.S. Social Security Administration

Social Bias



Gallagher et al CVPR 2008

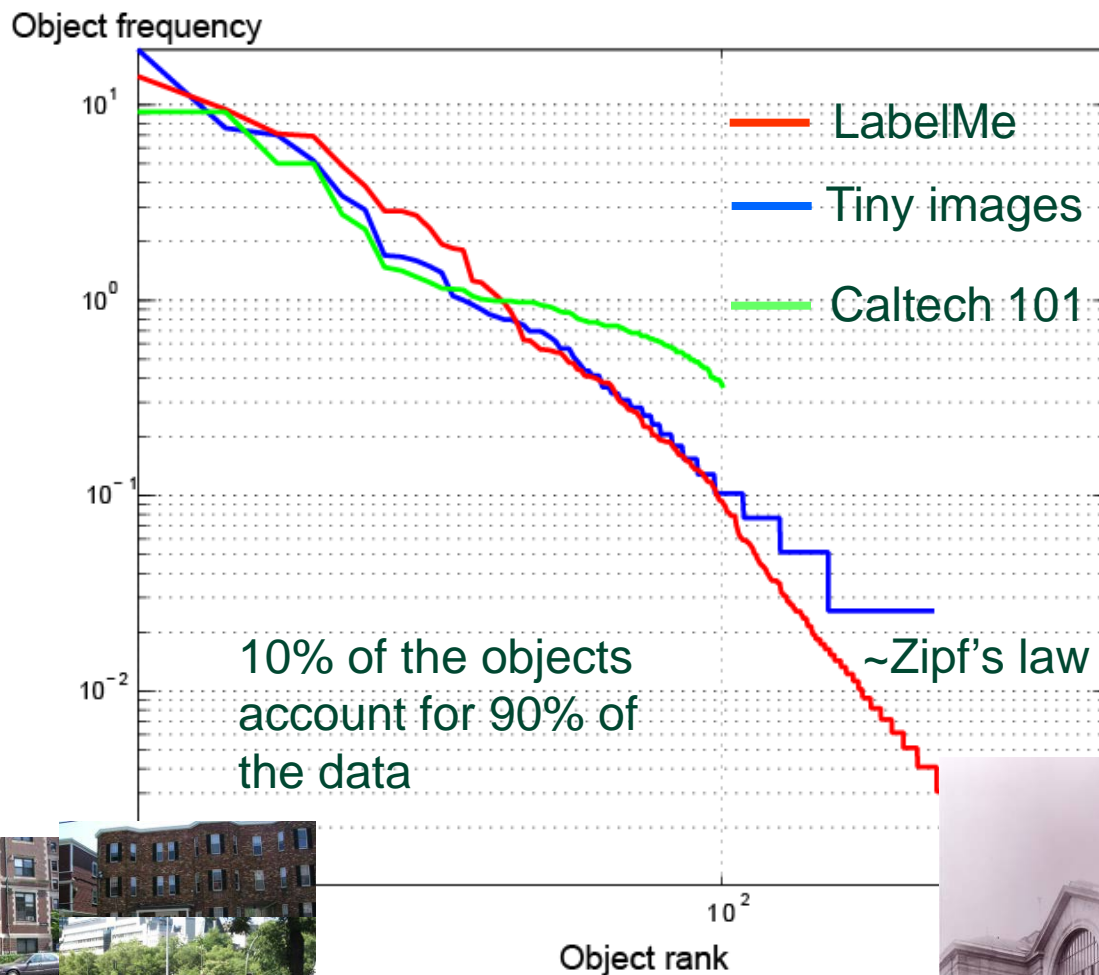


Gallagher et al, CVPR 2009

2. We will never have enough data



Long Tails -- Unfamiliar is Common



Dealing with sparse data (rare scenes)

Quick Fixes:

better alignment

- e.g. reduce resolution, sifting, warping, etc.

Understand the simple stuff first

Recognize when it's easy!

People take on a variety of **poses**, aspects, scales



self-occlusion

rare pose

motion blur



non-distinctive pose

too small

just right
detect this ²⁷

“Popping out” foreground objects

Hoiem et al, ICCV 2007







Figure 10. Object popout. We show five out of the fifteen most “solid” regions in the Geometric Context dataset. Our algorithm often finds foreground objects, which would be helpful for unsupervised object discovery [21].

Guess structure



David C. Lee, Martial Hebert, Takeo Kanade, CVPR'09

Guess structure




David C. Lee, Martial Hebert, Takeo Kanade, CVPR'09


Subtracting away structure



Structure



Objects



Wall appearance modeling

Dealing with sparse data (rare scenes)

Long-term Fixes:

segment into chunks

- e.g. segmentation for recognition approaches

Attributes – densifying the labels

From categorization to association

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

Conclusion...

*“If you torture data long enough,
it might confess”*

- Ronald Coase