

The Promise and Perils of Big Data





Some Slides from A. Efros and A. Torralba

Why do we need data?

Most problems in vision are ambiguous and hard.

- 2D -> 3D
- Segmentation/Edges



So, how do we solve these problems?



- Magic of data !
- Use data to learn better likelihoods: how things look like.
- Use data to learn priors of what is more likely than others.

But how much data do we need?

The extremes of learning

Extrapolation problem
Generalization
Transfer learning

Interpolation problem
Correspondence
Finding the differences



So how much data does humans use?

What's the Capacity of Visual Long Term Memory?

What we know...

Standing (1973)

10,000 images

83% Recognition

*... people can
remember thousands
of images*

High Fidelity Visual
Memory is possible
(Hollingworth 2004)

What we don't know...

*... what people are remembering for
each item?*



According to Standing

“Basically, my recollection is that we just separated the pictures into **distinct thematic categories**: e.g. cars, animals, single-person, 2-people, plants, etc.) Only a few slides were selected which fell into each category, and they were visually distinct.”



“Gist” Only




Sparse Details



Highly Detailed

Slide by Aude Oliva

Massive Memory I: Methods



Showed 14 observers 2500 **categorically unique objects**

1 at a time, 3 seconds each

800 ms blank between items

Study session lasted about 5.5 hours

Repeat Detection task to maintain focus

Followed by 300 2-alternative forced choice tests



how far can we push the fidelity of visual LTM representation ?

Same object category, different instance

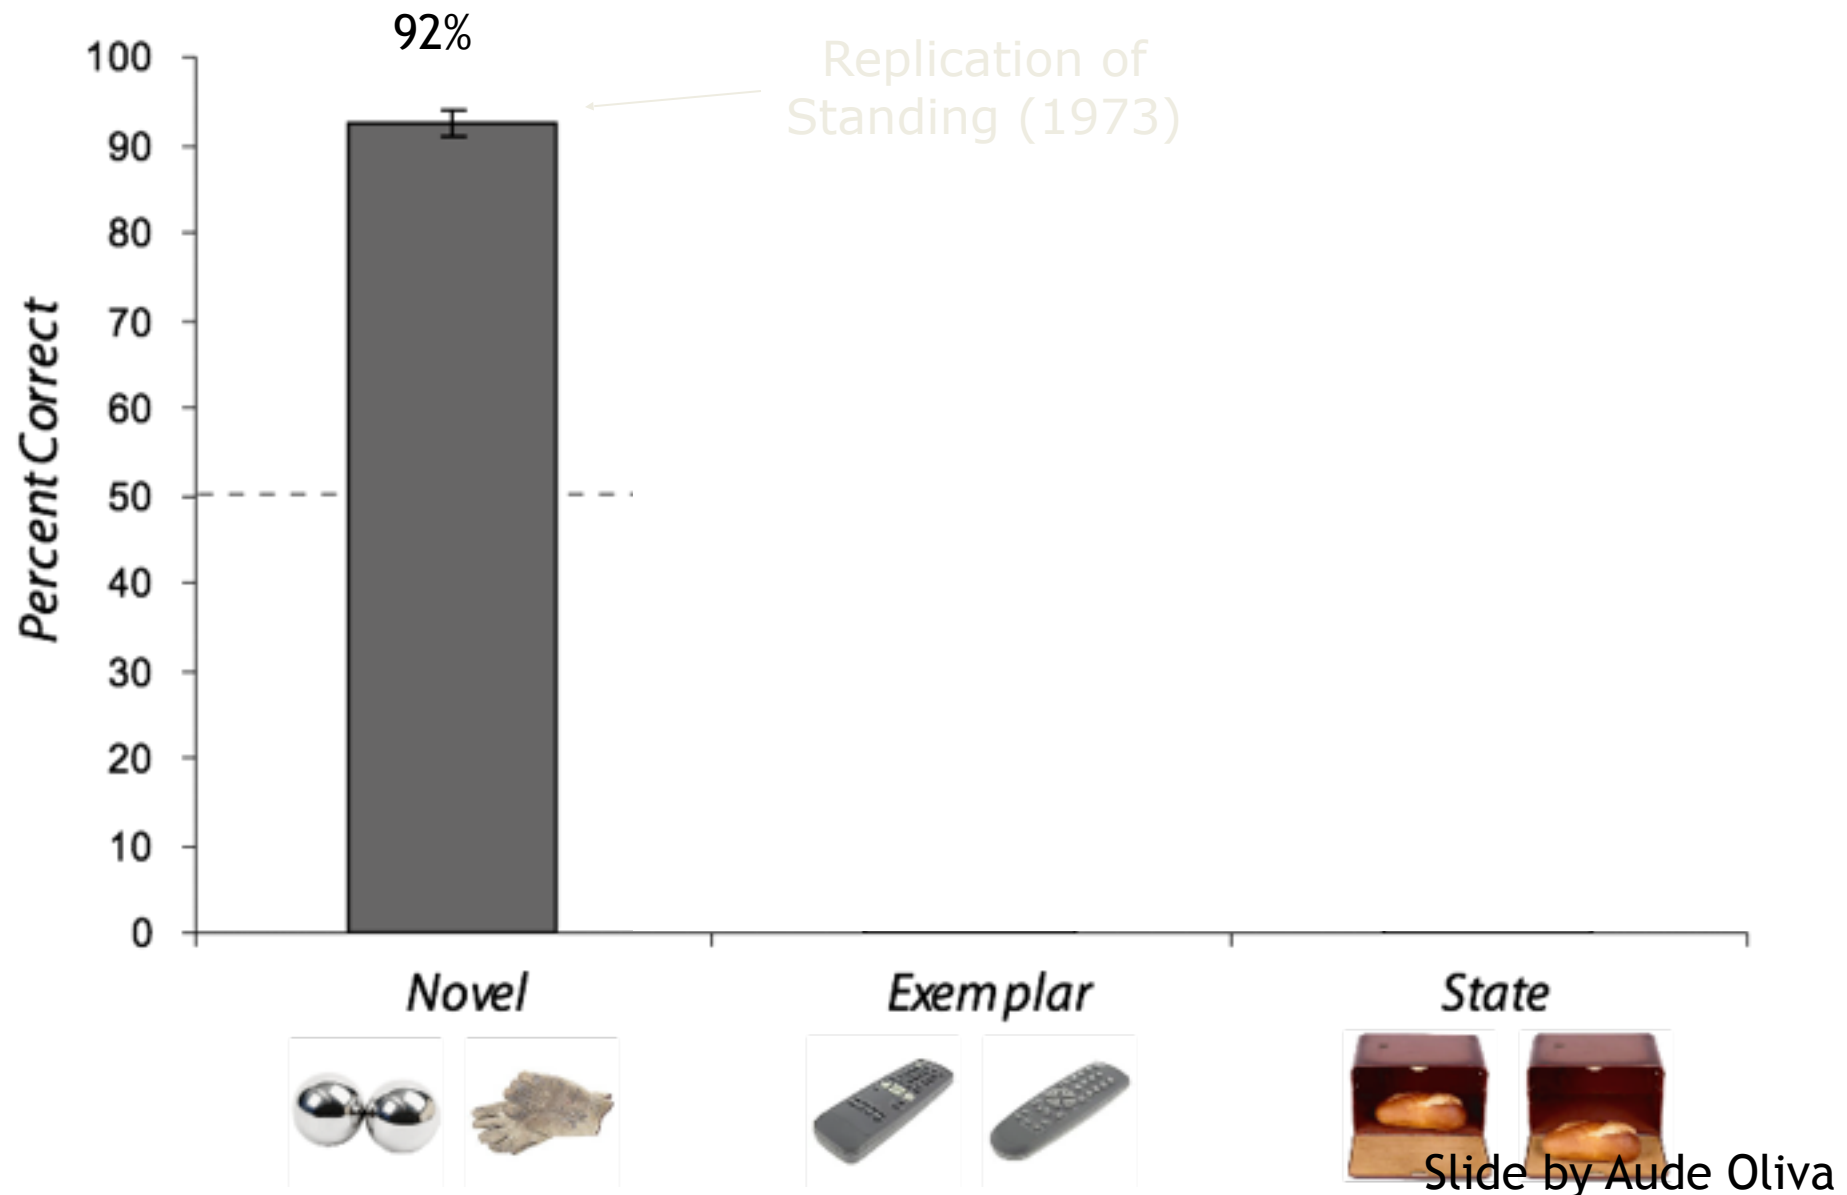


how far can we push the fidelity of visual LTM representation ?


Same object, different states




Massive Memory I: Recognition Memory Results



Massive Memory I: Recognition Memory Results



Extrapolation of Repeat Detection Data




how much data does computer vision
researchers use?



1972

10¹
images



10^1
images





FIGURE 4. This figure provides a high quality reproduction of the six images discussed in the text. (a) and (b) were taken with a considerably modified Information International Incorporated Vidisector, and the rest were taken with a Telemation TMC-2100 vidicon camera attached to a Spatial Data Systems digitizer (Camera Eye 108). The full dynamic range from black to white is represented by 256 grey-levels. The images reproduced here were created by an Optronics P1500h Photowriter from intensity arrays that measured 128 elements square. This size of intensity array corresponds to viewing a 1 in square at 5 ft with the human retina. The image of the period at the end of this sentence probably covers more than 40 retinal receptors. The reader should view the images from a distance of about 5 ft when assessing the performance of the programs.

10^{2-4}
images




The faces and cars scale

10^{2-4}
images



In 1996 DARPA released 14000 images, from over 1000 individuals.



The PASCAL Visual Object Classes

In 2007, the twenty object classes that have been selected are:

Person: person

Animal: bird, cat, cow, dog, horse, sheep

Vehicle: aeroplane, bicycle, boat, bus, car, motorbike, train


Indoor: bottle, chair, dining table, potted plant, sofa, tv/monitor



10^{2-4}
images



10^5
images



Caltech 101 and 256

10^5
images



Fei-Fei, Fergus, Perona, 2004



Griffin, Holub, Perona, 2007

Lotus Hill Research Institute image corpus

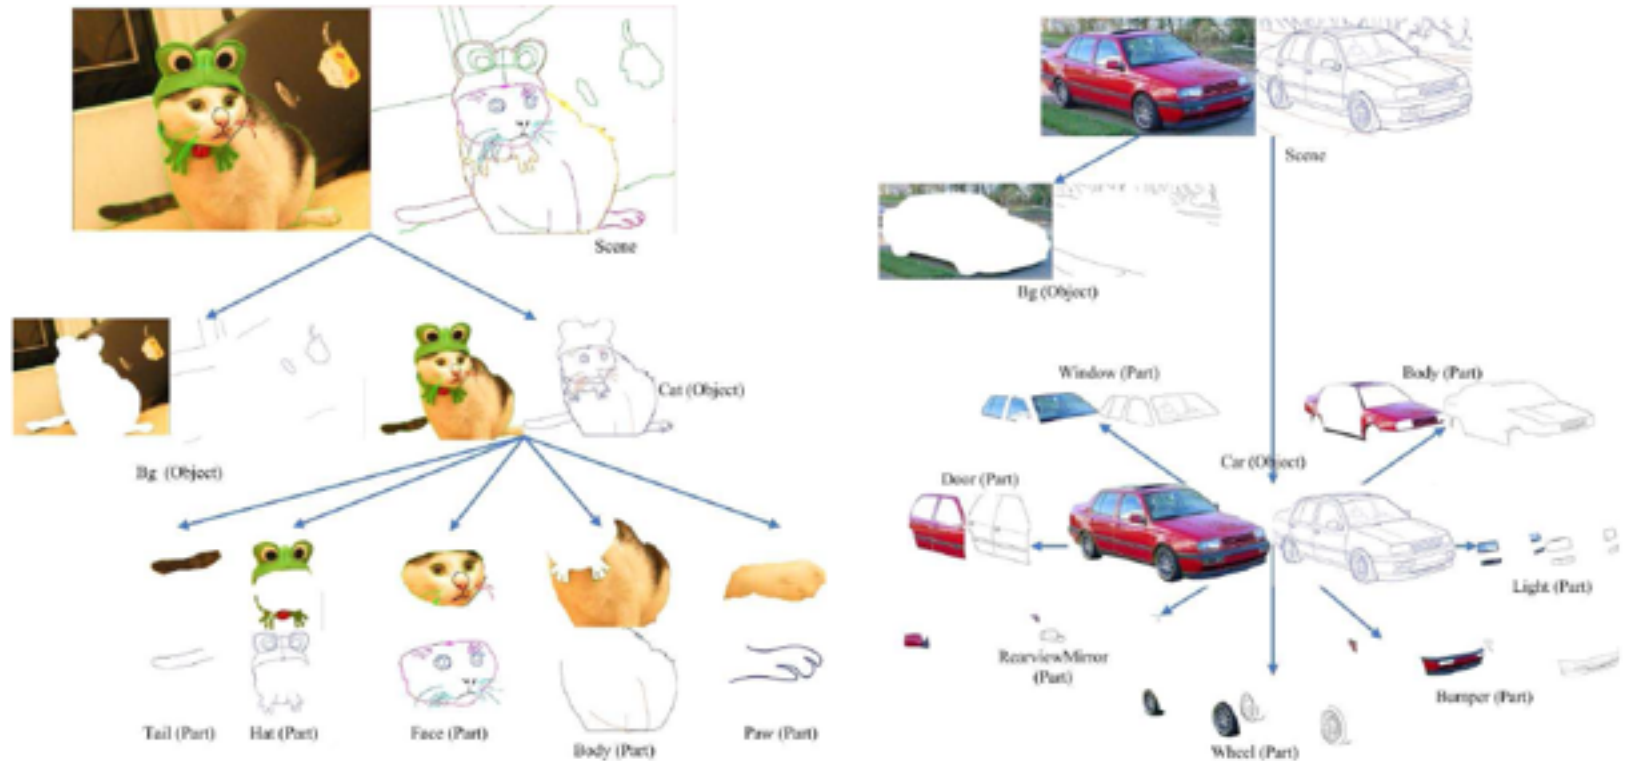
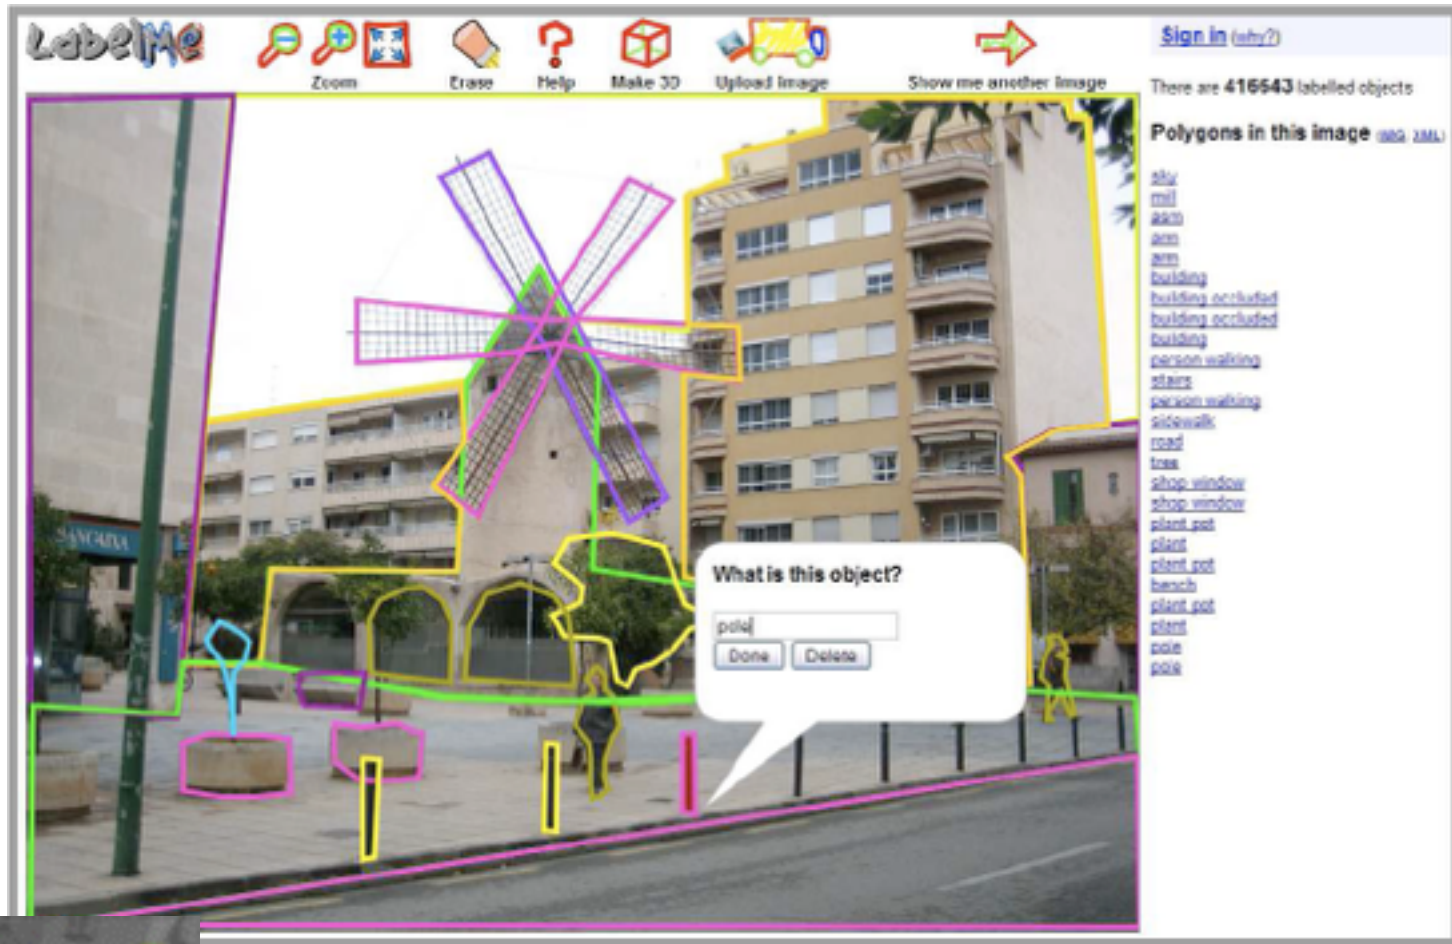


Figure 5: Two examples of the parse trees (cat and car) in the Lotus Hill Research Institute image corpus. From [87].

LabelMe

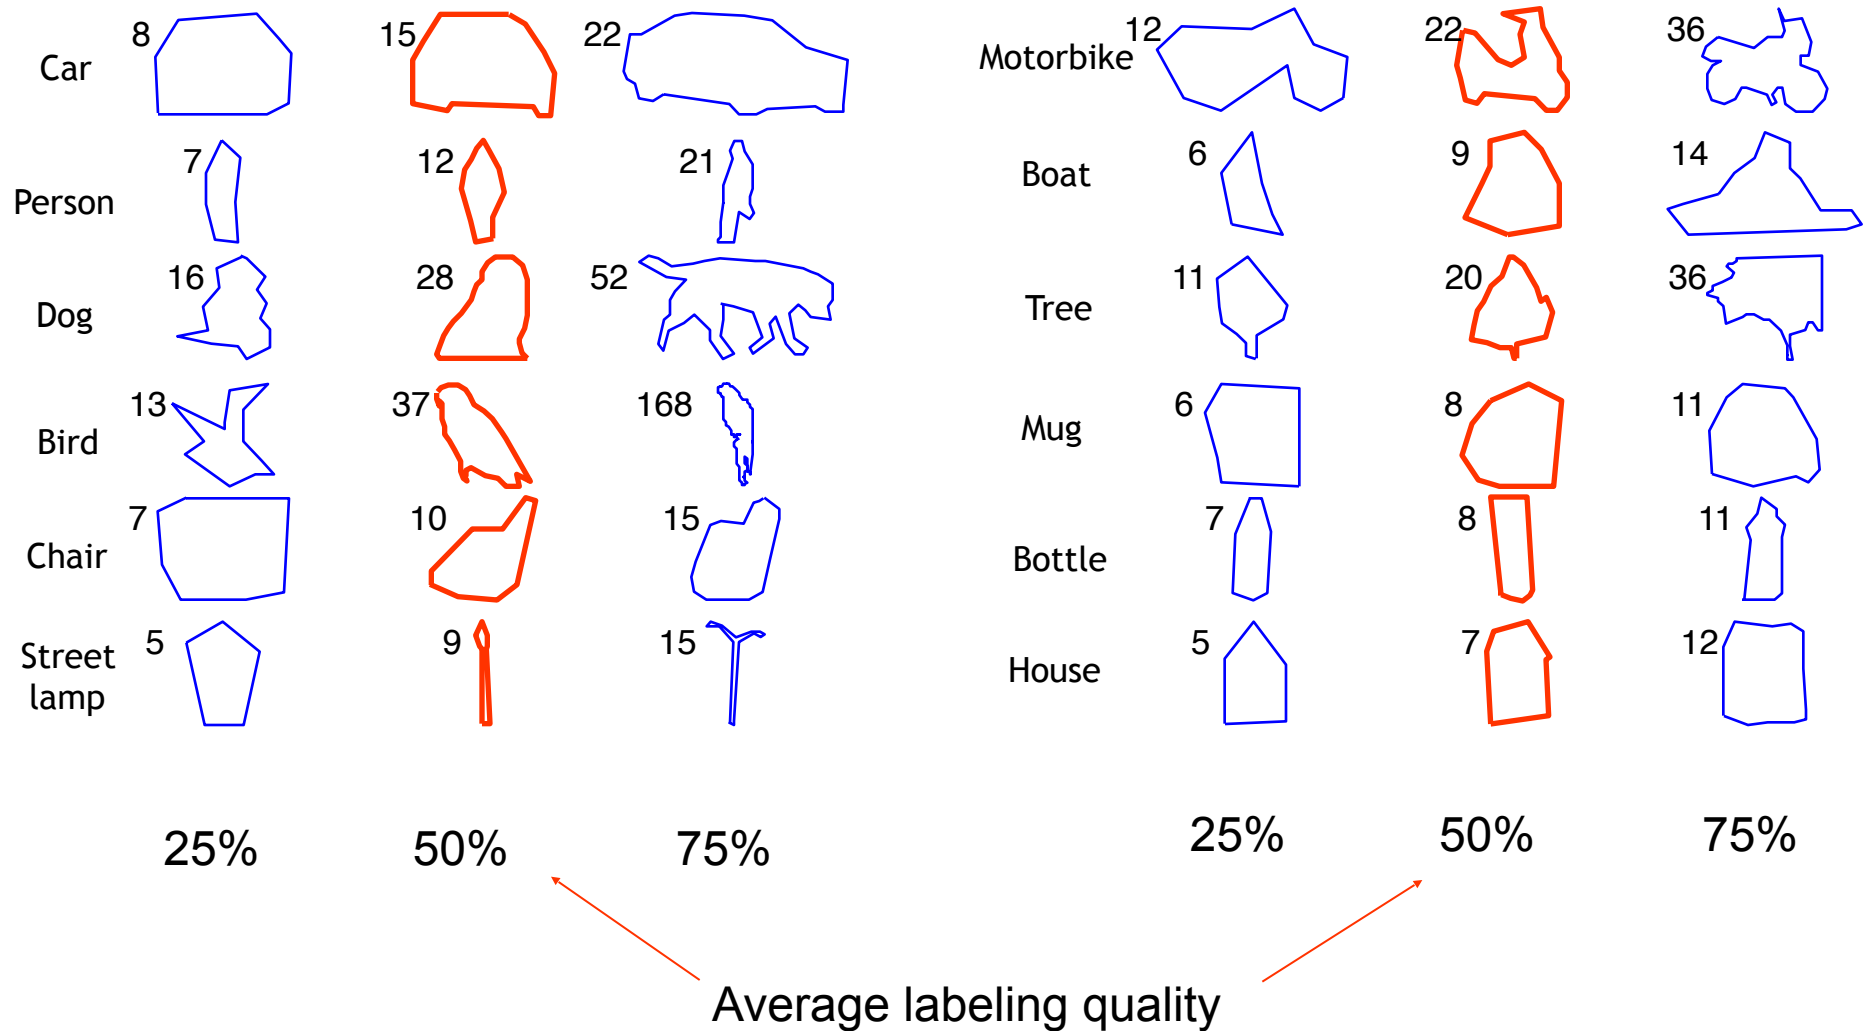
10^5
images



Tool went online July 1st, 2005
530,000 object annotations collected

Labelme.csail.mit.edu

Quality of labeling



Extreme labeling

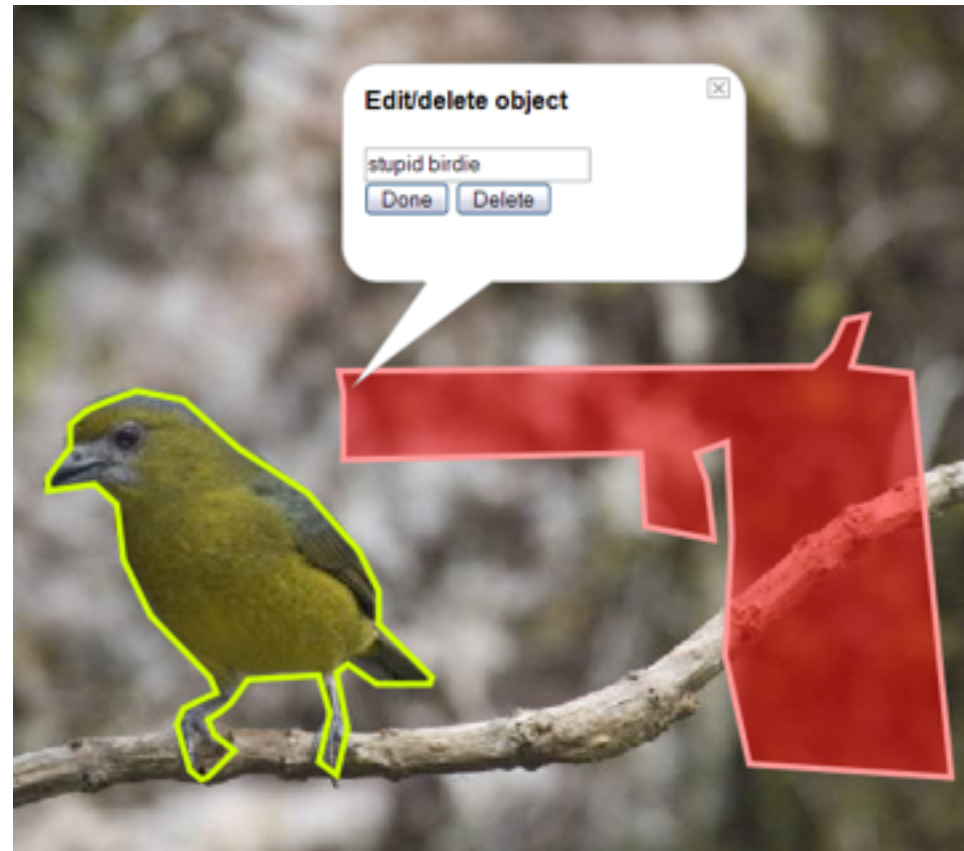
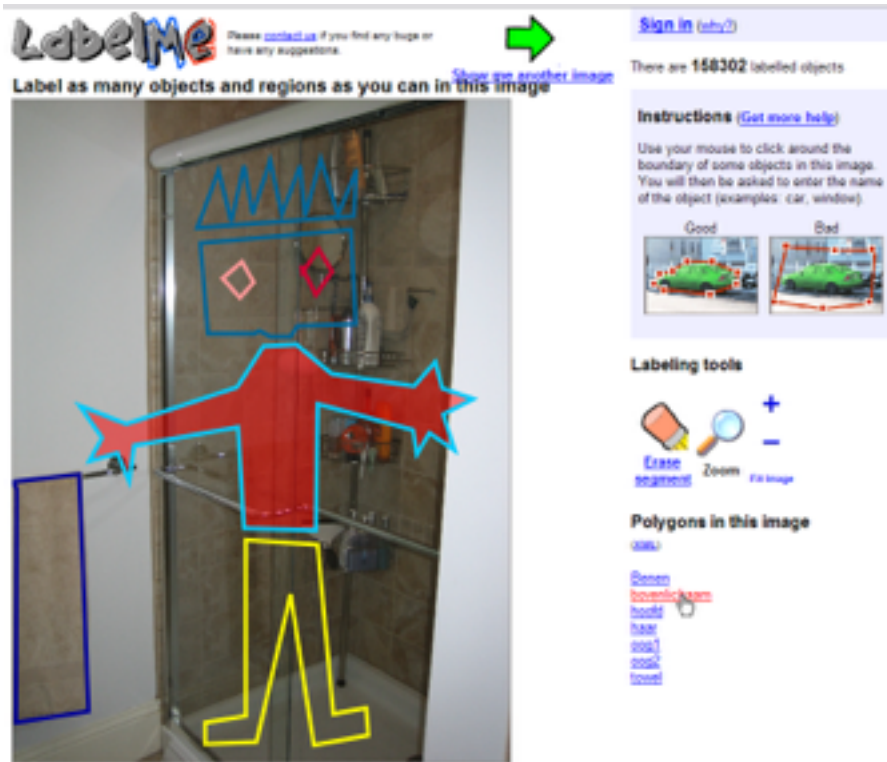


The other extreme of extreme labeling

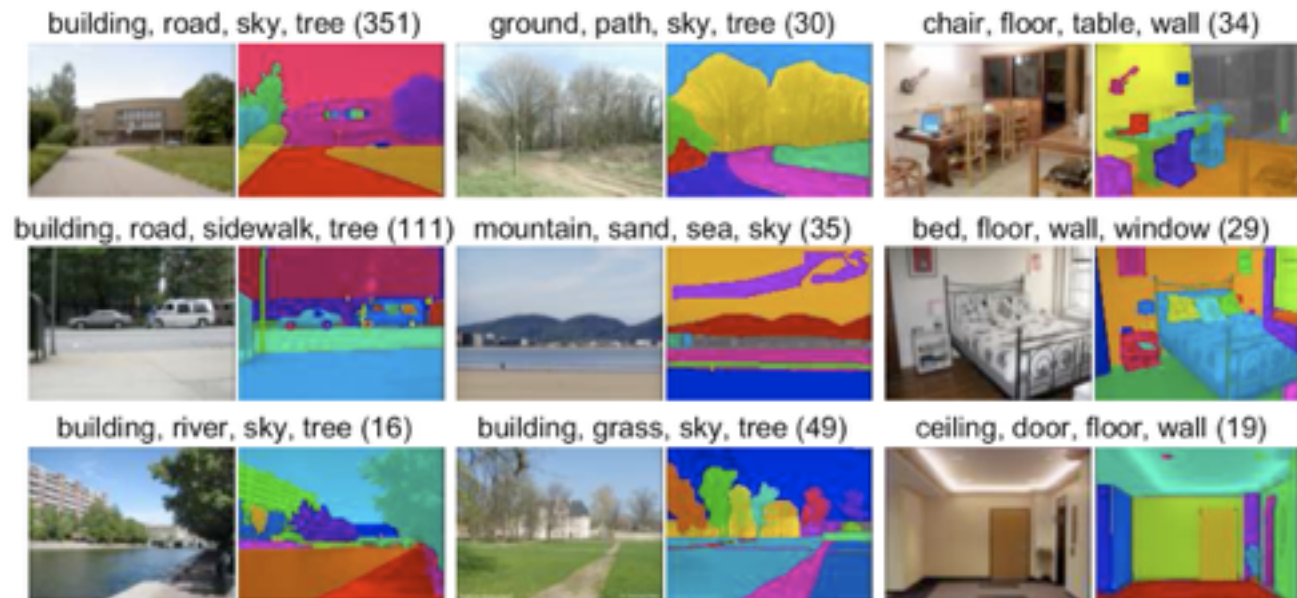
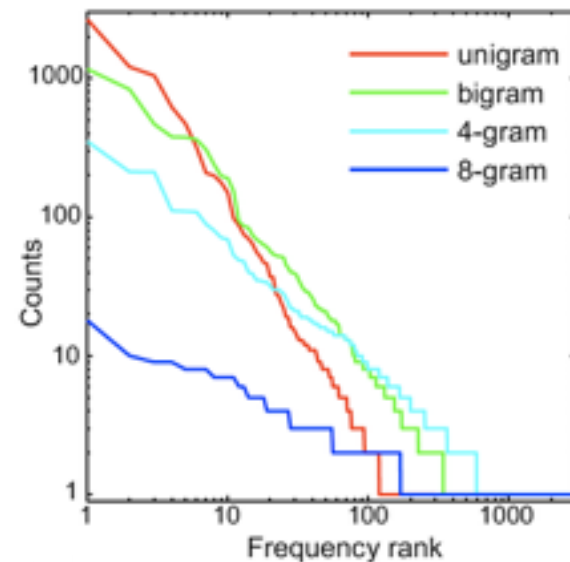
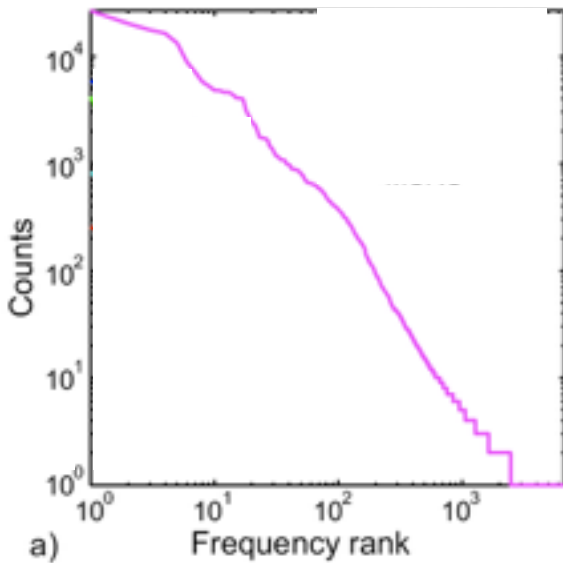
... things do not always look good...



Creative testing




Scene and object biases



10^5
images



10^{6-7}
images

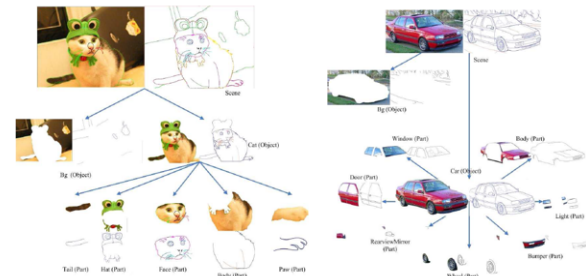


Things start getting out of hand

Collecting big datasets

10⁶⁻⁷
images

- **ESP game (CMU)**
Luis Von Ahn and Laura Dabbish 2004
- **LabelMe (MIT)**
Russell, Torralba, Freeman, 2005
- **StreetScenes (CBCL-MIT)**
Bileschi, Poggio, 2006
- **WhatWhere (Caltech)**
Perona et al, 2007
- **PASCAL challenge**
2006, 2007
- **Lotus Hill Institute**
Song-Chun Zhu et al, 2007
- **80 million images**
Torralba, Fergus, Freeman, 2007



80.000.000 images

10^{6-7}

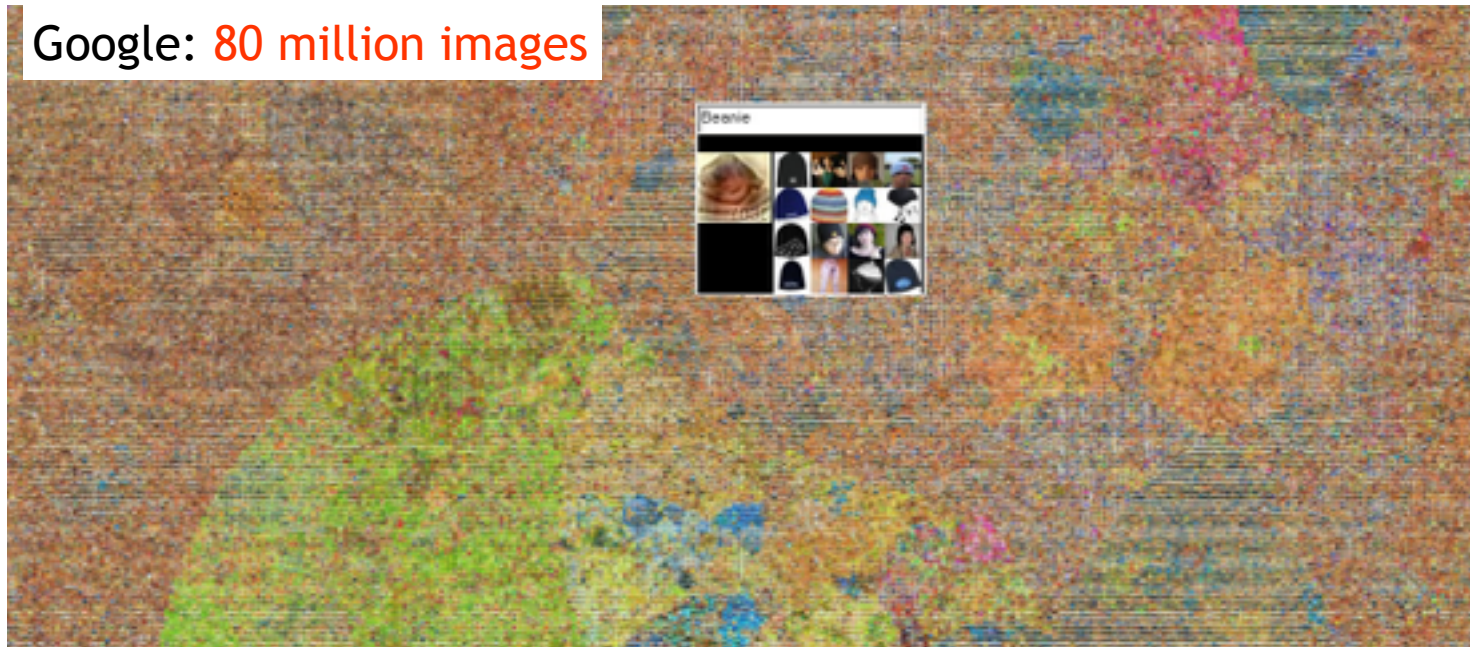
75.000 non-abstract nouns from WordNet

7 Online image search engines

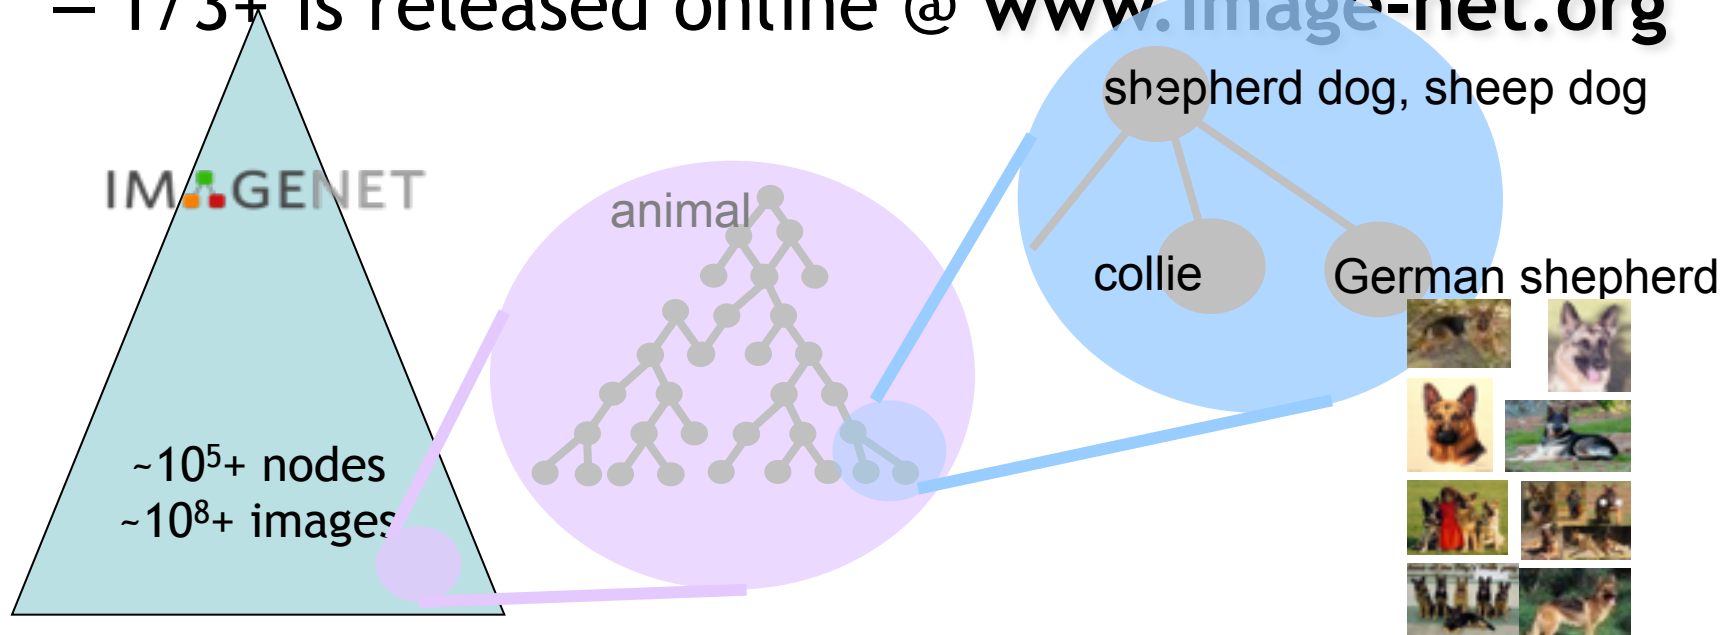


And after 1 year downloading images

Google: 80 million images



- An **ontology of images** based on WordNet
- ImageNet currently has
 - 22,000+ categories of visual concepts
 - 15 million human-cleaned images (~700im/categ)
 - 1/3+ is released online @ **www.image-net.org**



Mechanical Turk is a marketplace for work.

We give businesses and developers access to an on-demand, scalable workforce.
Workers select from thousands of tasks and work whenever it's convenient.

216,070 HITs available. [View them now.](#)

Make Money by working on HITs

HITs - *Human Intelligence Tasks* - are individual tasks that you work on. [Find HITs now.](#)

As a Mechanical Turk Worker you:

- Can work from home
- Choose your own work hours
- Get paid for doing good work



or [learn more about being a Worker](#)

Get Results from Mechanical Turk Workers

Ask workers to complete HITs - *Human Intelligence Tasks* - and get results using Mechanical Turk. [Register Now](#)

As a Mechanical Turk Requester you:

- Have access to a global, on-demand, 24 x 7 workforce
- Get thousands of HITs completed in minutes
- Pay only when you're satisfied with the results



Labeling for money

amazonmechanical turk
Artificial Artificial Intelligence

beta

Your Account | **HITs** | Qualifications | **56,035 HITs** available now

Bryan C Russell | Account Settings | Sign Out | Help

All HITs | HITs Available To You | HITs Assigned To You

Search for **HITs** containing _____ that pay at least \$ **0.00** for which you are qualified **GO**

Timer: 00:00:13 of 60 minutes

Finished with this HIT? **Submit HIT** | Let someone else do it? **Return HIT**


☐ Automatically accept the next HIT

Total Earned: \$0.01
Total HITs Submitted: 12

LabelMe: Label objects in this image
Requester: Bryan C Russell
Qualifications Required: None


Reward: \$0.01 per HIT | HITs Available: 269 | Duration: 60 minutes

Please label as many objects as you want in this image. Scroll down to see the entire image. **Submit HIT**



Alexander Sorokin, David Forsyth, "Utility data annotation with Amazon Mechanical Turk", First IEEE Workshop on Internet Vision at CVPR 08.

10^{6-7}
images



10^{8-11}
images



Datasets in perspective

Number of images on my hard drive: 10^4

PASCAL

LabelMe



Number of images seen during my first 10 years: 10^8
(3 images/second * 60 * 60 * 16 * 365 * 10 = 630720000)

flickr
Google



Number of images seen by all humanity: 10^{20}
 $106,456,367,669 \text{ humans}^1 * 100 \text{ years} * 3 \text{ images/second} * 60 * 60 * 16 * 365 =$
1 from <http://www.prb.org/Articles/2002/HowManyPeopleHaveEverLivedonEarth.aspx>



Number of all 32x32 images: 10^{7373}
 $256^{32*32*3} \sim 10^{7373}$



Number of
samples


When do we need big data?

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 Odalagh



From cobalt123



From Martin LaBar...



From floomschb



From "iris-hues"



From Fareed...



From Gary~



From EkeDave



From Mark Watson...



From hogsvilleBri...



From Photo by...



From martyspants



From dmskct



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,... visual understanding?
- Enter: The **Magic of Data**
 - Great advances in several fields:
 - e.g. speech recognition, machine translation, Google

Unreasonable Effectiveness of Data

Simple Algorithms (Dumb) + Lot of Data
are better than Complicated algorithms

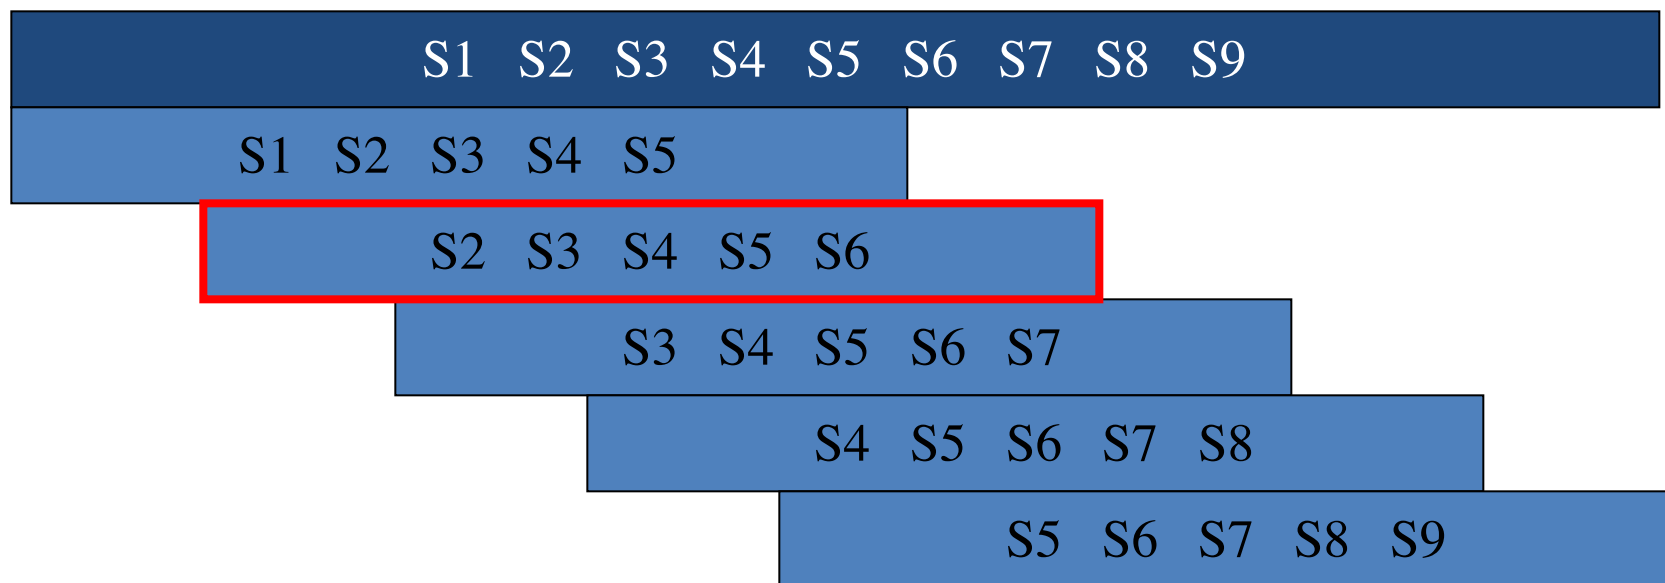
Example: Machine Translation

Example: Texture Generation

Machine Translation

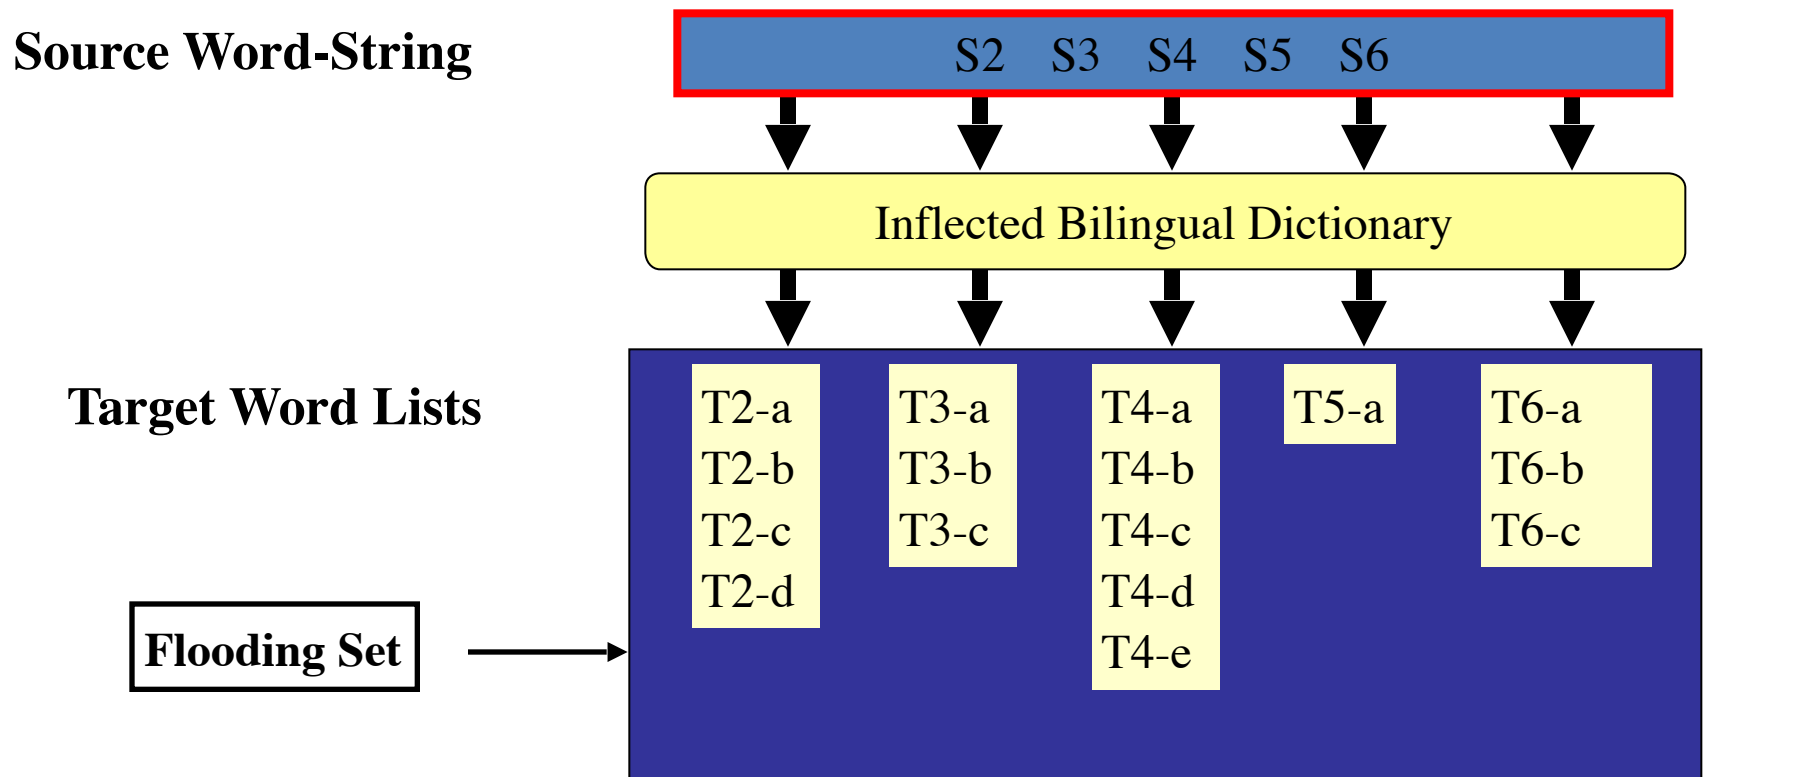
Step 1: Source Sentence Chunking

- Segment source sentence into overlapping n-grams via sliding window
- Typical n-gram length 4 to 9 terms
- Each term is a word or a known phrase
- Any sentence length



Step 2: Dictionary Lookup


- Using bilingual dictionary, list all possible target translations for each source word or phrase



Step 3: Search Target Text

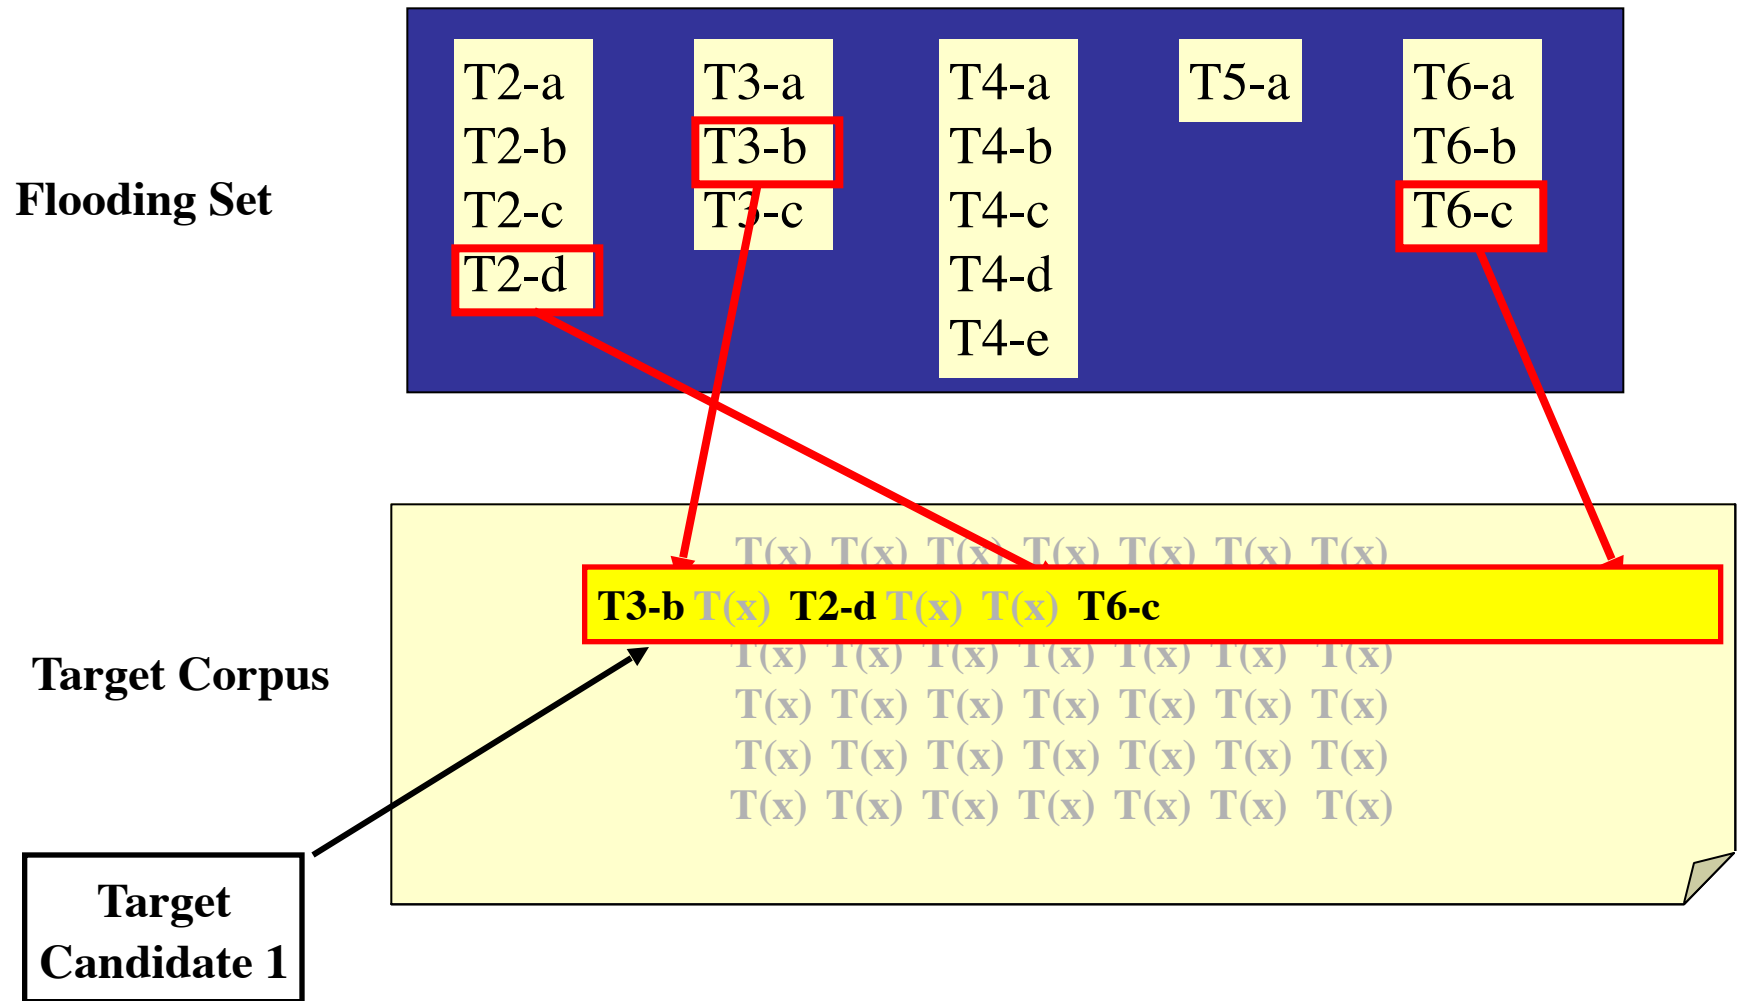
- Using the Flooding Set, search target text for word-strings containing one word from
 - each group

Flooding Set



- Find maximum number of words from Flooding Set in minimum length word-string
 - *Words or phrases can be in any order*
 - *Ignore function words in initial step (T5 is a function word in this example)*

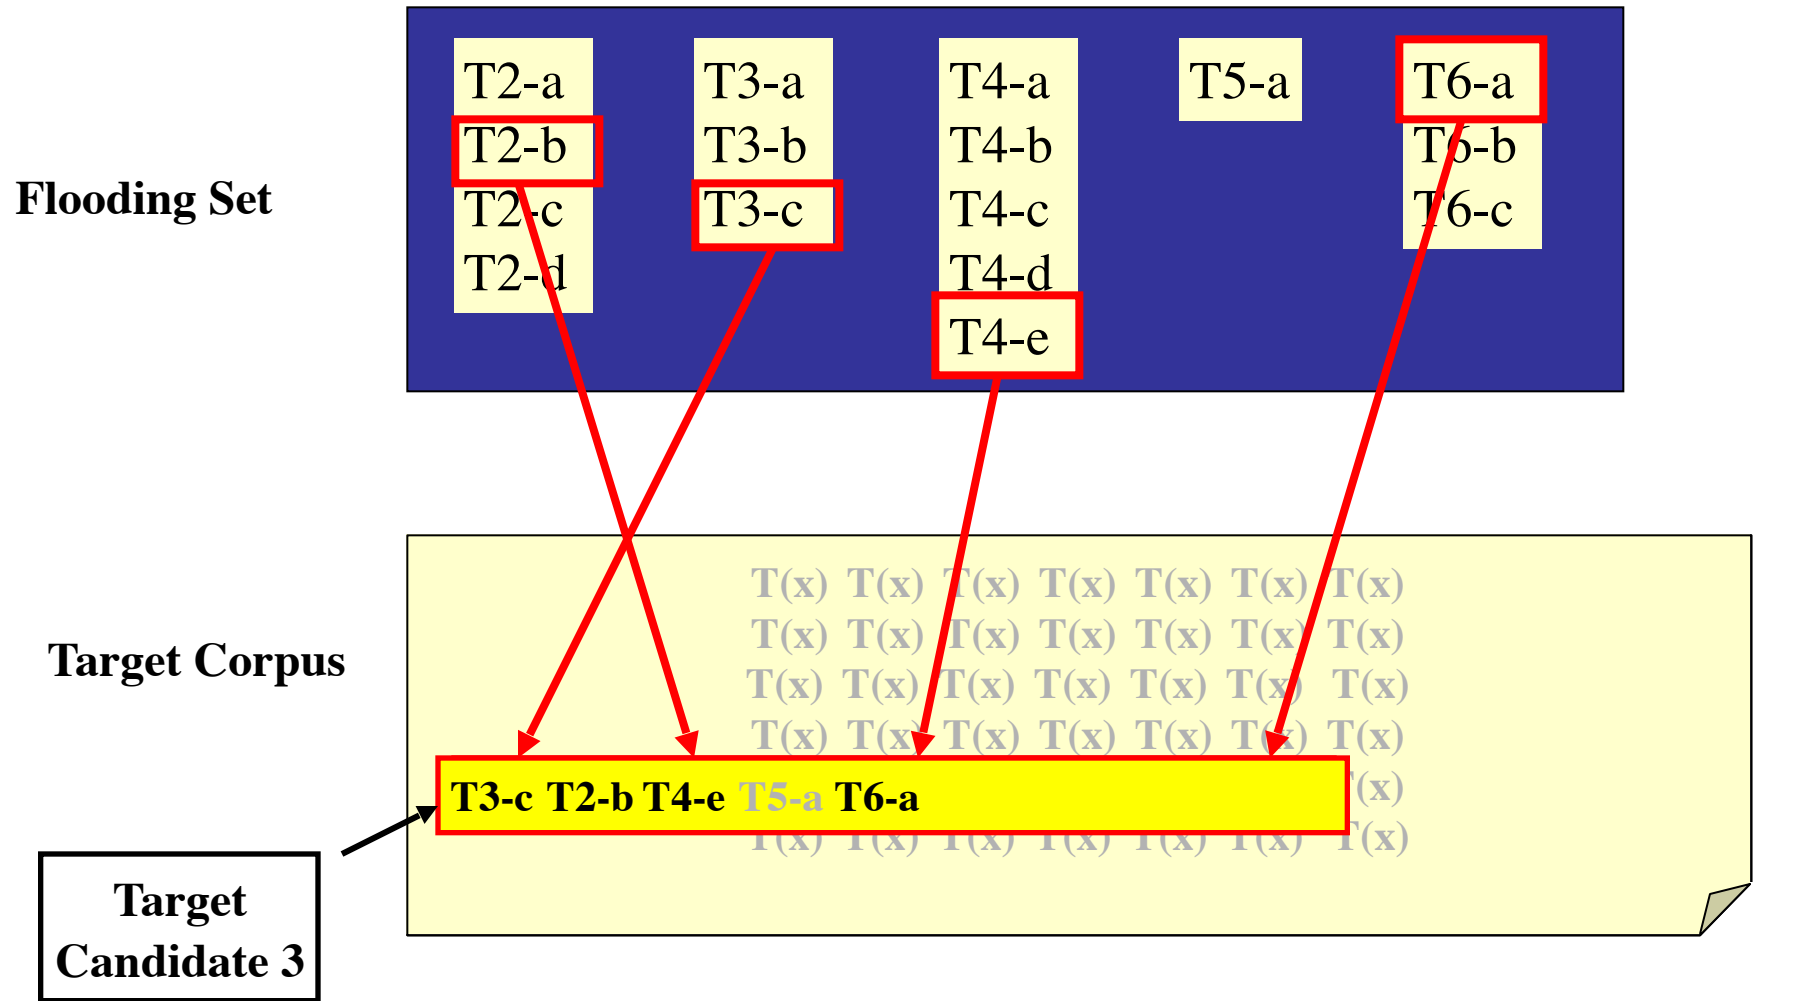
Step 3: Search Target Text (Example)



Flooding Set



Step 3: Search Target Text (Example)




Reintroduce function words after initial match (e.g. T5)

Slide by Jaime Carbonell

Step 4: Score Word-String Candidates

- Scoring of candidates based on:
 - Proximity (minimize extraneous words in target n-gram \approx precision)
 - Number of word matches (maximize coverage \approx recall)
 - Regular words given more weight than function words
 - Combine results (e.g., optimize F_1 or p-norm or ...)

Target Word-String Candidates



Step 5: Select Candidates Using Overlap (Propagate context over entire sentence)

**Word-String 1
Candidates**

T(x1)	T2-d	T3-c	T(x2)	T4-b
T(x1)	T3-c	T2-b	T4-e	
T(x2)	T4-a	T6-b	T(x3)	T2-c

**Word-String 2
Candidates**

T3-b	T(x3)	T2-d	T(x5)	T(x6)	T6-c
T4-a	T6-b	T(x3)	T2-c	T3-a	
T3-c	T2-b	T4-e	T5-a	T6-a	

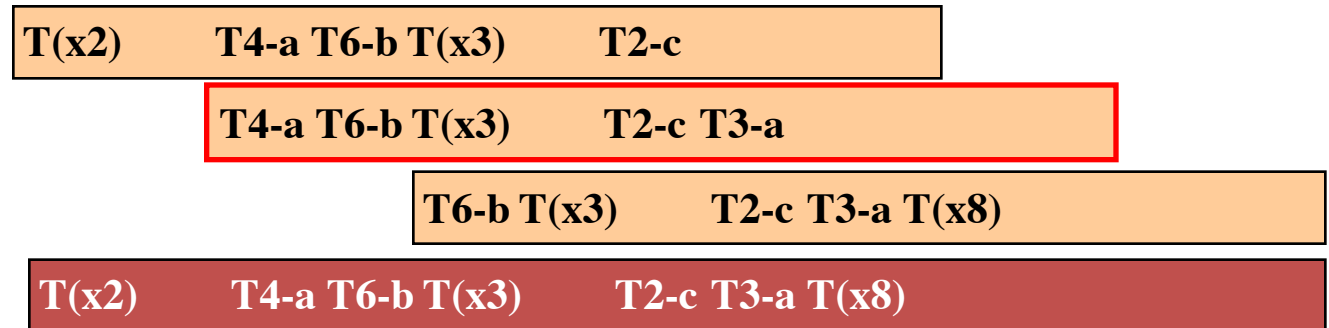
**Word-String 3
Candidates**

T2-b	T4-e	T5-a	T6-a	T(x8)
T6-b	T(x11)	T2-c	T3-a	T(x9)
T6-b	T(x3)	T2-c	T3-a	T(x8)

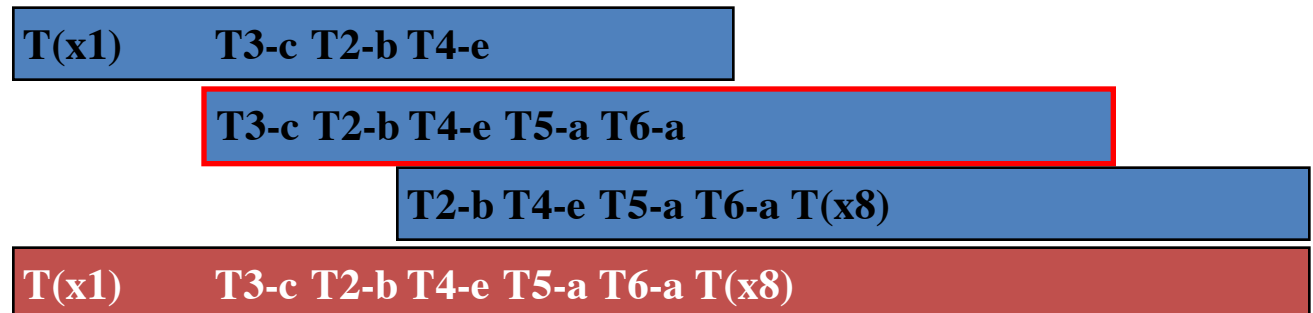
Step 5: Select Candidates Using Overlap

Best translations selected via maximal overlap

Alternative 1



Alternative 2



A (Simple) Real Example of Overlap

Flooding → N-gram fidelity

Overlap → Long range fidelity

**N-grams
generated
from
Flooding**

a United States soldier

United States soldier died

soldier died and two others

died and two others were injured

two others were injured Monday

**N-grams connected via
Overlap**

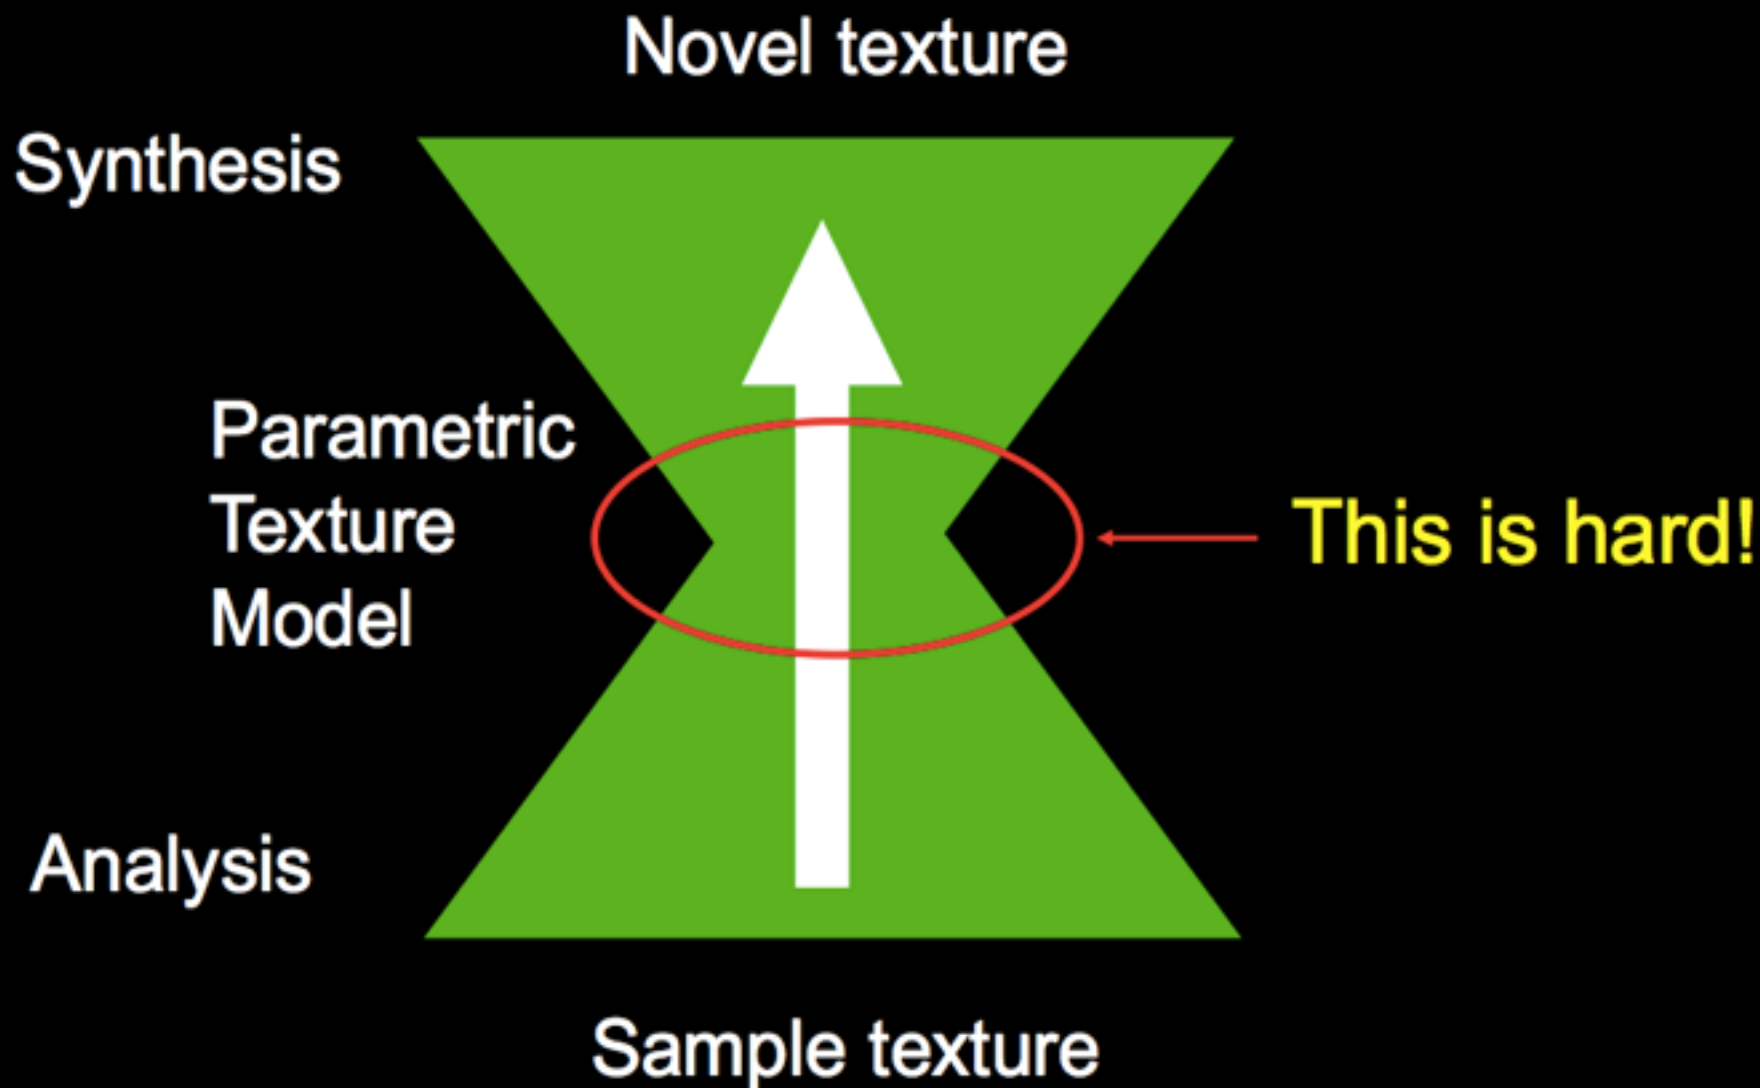
a United States soldier died and two others were injured Monday

Texture Synthesis


Texture Synthesis



Classical Texture Synthesis



Throwing away too much too soon?

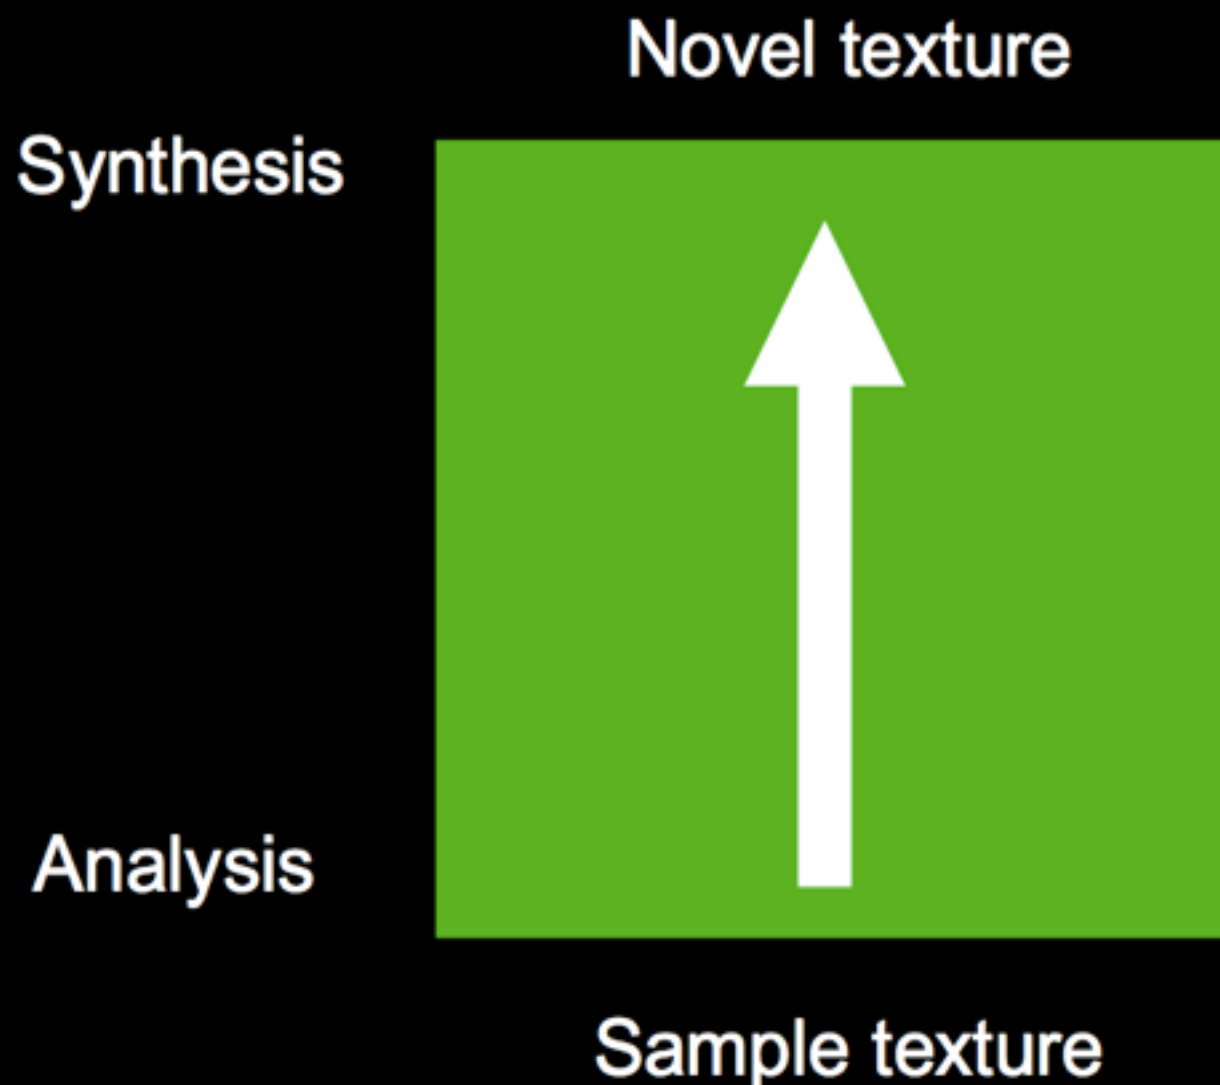


input texture




synthesized texture


Non-parametric Approach



[Efros & Leung, '99, Efros & Freeman '01]




non-parametric
sampling




Input image

Texture Growing




So, how do we use big data?

Two ways to use Lots of Data



Brute Force Vision: Find that needle in the haystack and disregard the rest (a.k.a. kNN)



See what different subsets of data think of you



kNN matching is great...

- because we live in a (mostly) boring world!



Lots Of Images


Target



7,900




Lots Of Images



Lots Of Images

Target



7,900



790,000



79,000,000




Automatic Colorization Result

Grayscale input High resolution



Colorization of input using average



im2gps

Instead of using objects labels, the web provides other kinds of metadata associate to large collections of images



Figure 2. The distribution of photos in our database. Photo locations are cyan. Density is overlaid with the jet colormap (log scale).

20 million geotagged and geographic text-labeled images




Figure 5. *Geolocation performance across features.* Percentage of test cases geolocated to within 200km for each feature. We compare geolocation by 1-NN vs. largest mean-shift mode.



Image completion



Original Image

Input


Criminisi et al.

MS *Smart Erase*

Instead, generate proposals using millions of images



Input



16 nearest neighbors
(gist+color matching)



output

Hays, Efros, 2007

With a good image similarity
and a lot of data...

Input image




Nearest neighbors






22,000 LabelMe scenes

With a good image similarity and a lot of data...






With a good image similarity
and a lot of data...








Russell, Liu, Torralba, Fergus, Freeman. NIPS 2007

With a good image similarity and a lot of data...





Russell, Liu, Torralba, Fergus, Freeman. NIPS 2007

With a good image similarity and a lot of data...



Russell, Liu, Torralba, Fergus, Freeman. NIPS 2007

Outputs



While many scenes are boring...



Slide by Antonio Torralba

Some scenes are unique



Slide by Antonio Torralba

Dealing with sparse data (rare scenes)

- better similarity



Medici Fountain, Paris



Search by image



[→ Move](#)

Drop image here

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



medici_summer.jpg

luxembourg gardens

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)



medici_winter.png x luxembourg gardens



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







painting.png



describe image here



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







medici_sketch.bmp x

describe image here



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






Input Query




Top Matches

Input Query



Top Matches


Input Query



Top Matches

IMPORTANT PARTS?

Input Query





Important Parts






Top Matches

Input Query





“Data-driven Uniqueness”



Search using Images

Input Query



Top Matches

Search using Sketches



Search using Paintings



Input Painting



Top Matches

Search using Paintings



Input Painting



Top Matches

Dealing with sparse data (rare scenes)

- better similarity
- better alignment
 - e.g. reduce resolution, sifting, warping, etc.

Matching scenes

Two images taken from the same scene category, but different instances

- Contain different objects with different scales, perspectives and spatial location



Image representation









Image gradients

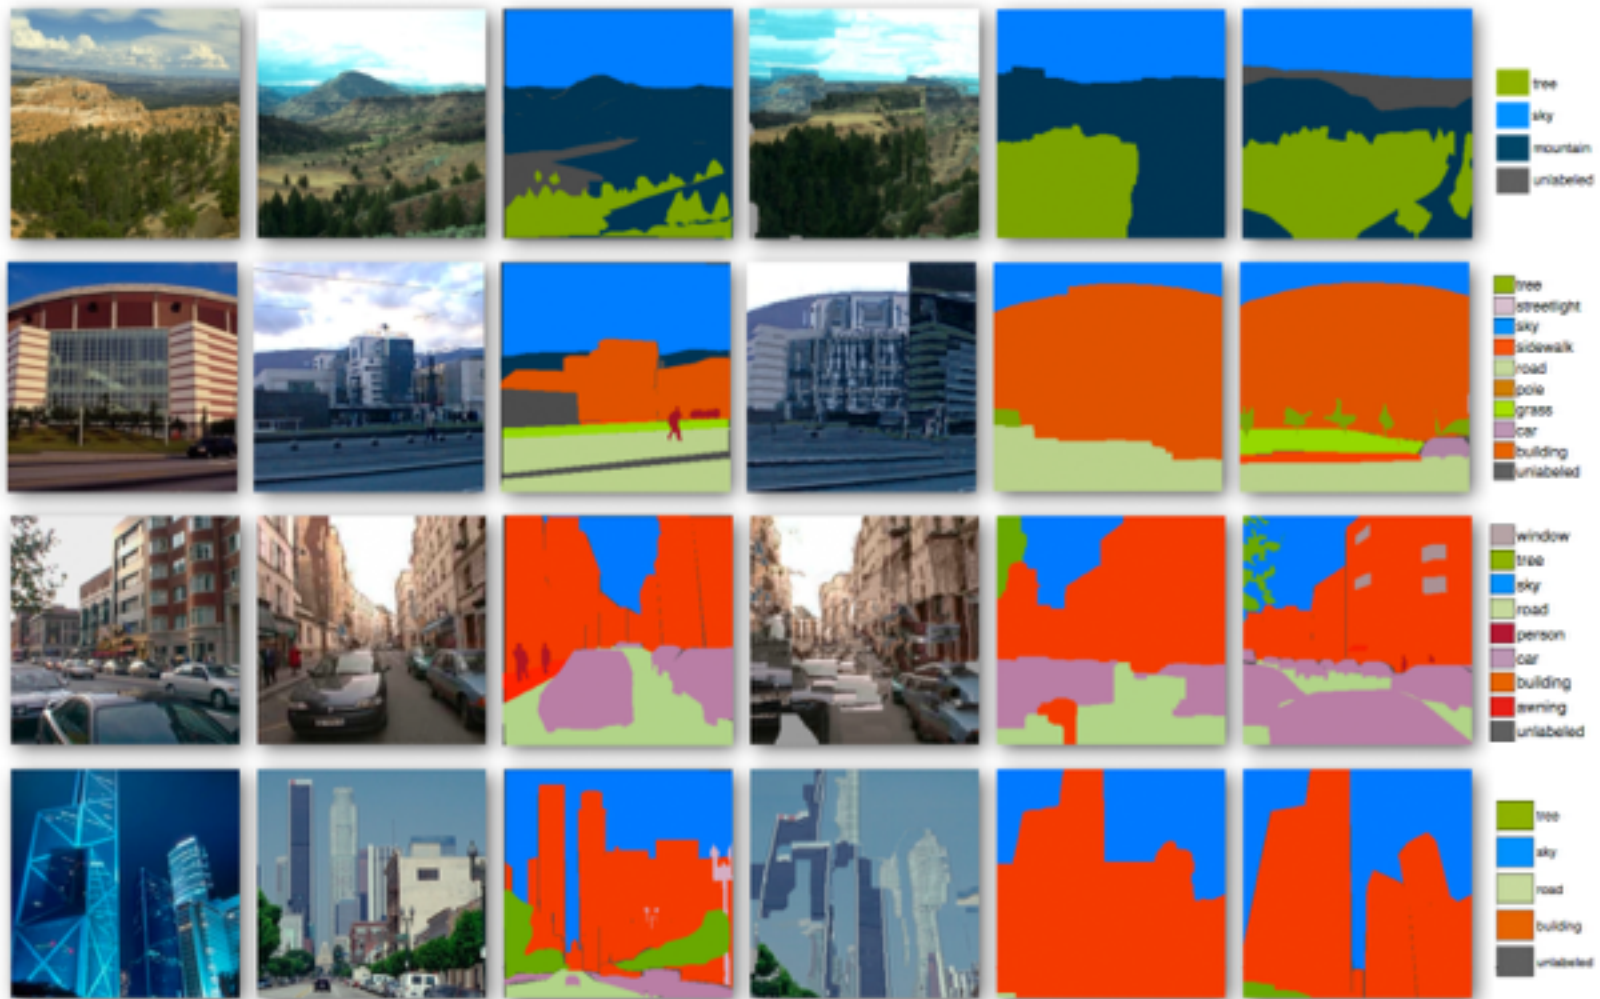


Keypoint descriptor





Scene parsing results



Query

Best match

Annotation of
best match

Warped best
match to query

Parsing result

Ground truth



(1)



(2)



(3)

(a)

(b)


(c)

Prediction


Dealing with sparse data (rare scenes)

- better similarity
- better alignment
 - e.g. reduce resolution, sifting, warping, etc.
- Use sub-images (primitives) to match
 - Allows matching from multiple images



Predicting Surface Normals





Matching Parts



Matching Parts



Matching Parts



Dealing with sparse data (rare scenes)

- better similarity
- better alignment
- Use sub-images (primitives) to match
- Understand the simple stuff first
 - e.g. tracking via recognition, background subtraction, “object pop-out”, etc.

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

Guess structure



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


Guess structure



Subtracting away structure



Structure



Objects




Wall appearance modeling

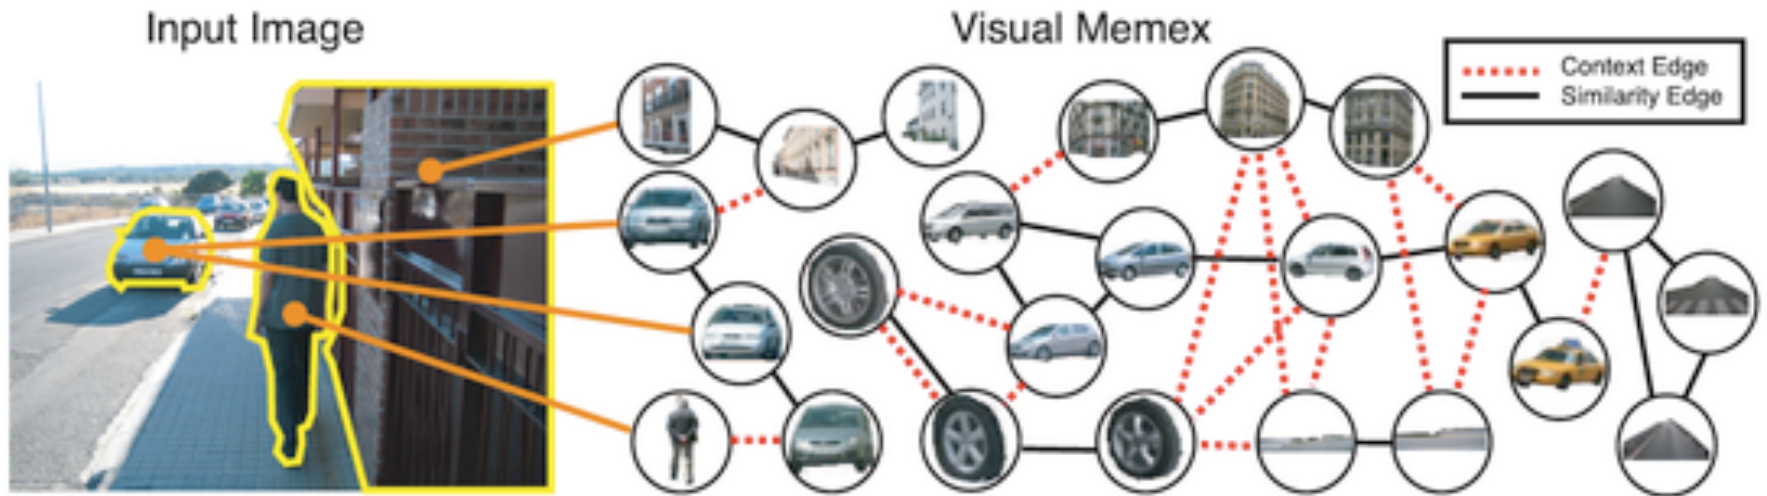
David C. Lee, Martial Hebert, Takeo Kanade, CVPR'09

Dealing with sparse data (rare scenes)

- better similarity
- better alignment
 - e.g. reduce resolution, sifting, warping, etc.
- segment into chunks
 - e.g. segmentation for recognition approaches
- get rid of simple stuff first
 - e.g. background subtraction, “object pop-out”, etc.
- Moving away from kNN methodology...
- use data to make connections
 - e.g. The Memex, manifold learning, data association, subpopulation means, etc.



Memex - Knowledge Graph




Manifolds

- Images are high dimensional: A 64x64 image is 4096 dimensional vector.
- But the possible images are much less!
- Is there a subspace where the set of images lie?

manifolds in vision


appearance variation




manifolds in vision



reasonable distance metrics



reasonable distance metrics




linear interpolation

reasonable distance metrics




manifold interpolation

reasonable distance metrics




reasonable distance metrics




Some observations about data collection

Object distributions




LabelMe



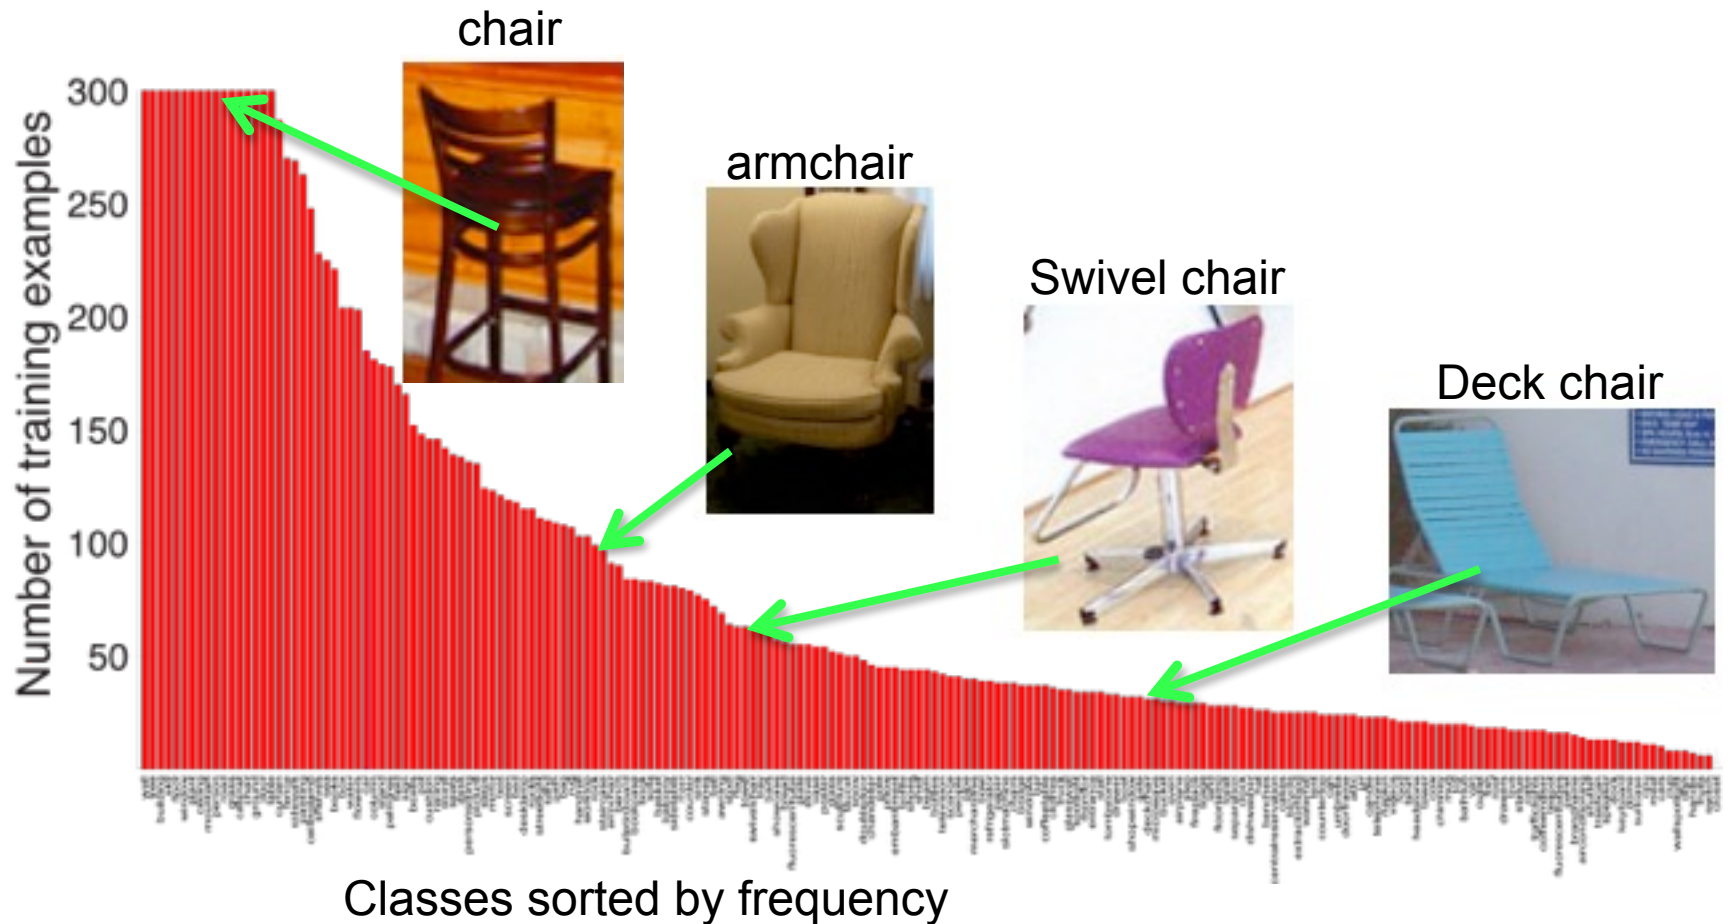
IMAGENET

SUN database




The first 9 objects account for 50% of all training examples
17 classes with more than 300 examples
109 classes with less than 50 examples

Rare objects are similar to frequent objects



Rare objects are similar to frequent objects



Some bias comes from the way the data is collected

mug Search SafeSearch moderate ▼


About 10,100,000 results (0.09 seconds) Advanced search

59¢ Logo Coffee Mugs
www.DiscountMugs.com Lead Free & Dishwasher Safe. Save 40-50%. No Catch. Factory Direct!


Custom Mugs On Sale
www.Vistaprint.com Order Now & Save 50% On Custom Mugs No Minimums. Upload Photos & Logos.

Promotional Mugs from 69¢
www.4imprint.com/Mugs Huge Selection of Styles Colors- Buy 72 Mugs @ \$1.35 ea-24hr Service


Related searches: [white mug](#) [coffee mug](#) [mug root beer](#) [mug shot](#)




Representational
500 × 429 - 91k - jpg
eagereyes.org
[Find similar images](#)




Ceramic Happy Face
300 × 300 - 77k - jpg
larose.com
[Find similar images](#)




Here I go then, trying
600 × 600 - 35k - jpg
beeper.wordpress.com
[Find similar images](#)




The Chalk Mug »
304 × 314 - 17k - jpg
coolest-gadgets.com
[Find similar images](#)




mug




Bring your own
500 × 451 - 15k - jpg
cookstounited.ca
[Find similar images](#)




ceramic mug
980 × 1024 - 30k - jpg
diytrade.com




Dual Purpose Drinking
490 × 428 - 16k - jpg
freshome.com
[Find similar images](#)




This coffee mug,
300 × 300 - 22k - jpg
gizmodo.com
[Find similar images](#)




Back to Ceramic
400 × 400 - 8k - jpg
freshpromotions.com.au
[Find similar images](#)




Coffee Mug as a
303 × 301 - 10k - jpg
dustbowl.wordpress.com
[Find similar images](#)




SASS Life Member
300 × 302 - 6k - jpg
sassnet.com



personalized coffee
400 × 343 - 15k - jpg
walyou.com
[Find similar images](#)



We like our mugs
290 × 290 - 6k - jpg
kitchencontraptions.com
[Find similar images](#)



mug

Search

SafeSearch moderate

About 10,100,000 results (0.09 seconds)

Advanced search

59¢ Logo Coffee Mugs

www.DiscountMugs.com

Lead Free & Dishwasher Safe. Save 40-50%. No Catch. Factory Direct !

Custom Mugs On Sale

www.Vistaprint.com

Order Now & Save 50% On Custom Mugs No Minimums. Upload Photos & Logos.


Promotional Mugs from 69¢

www.4imprint.com/Mugs


Huge Selection of Styles Colors- Buy 72 Mugs @ \$1.35 ea-24hr Service

Sponsored


Related searches: [white mug](#) [coffee mug](#) [mug root beer](#) [mug shot](#)




Representational
500 × 429 - 91k - jpg
[eagereyes.org](#)
[Find similar images](#)




Ceramic Happy Face
300 × 300 - 77k - jpg
[larose.com](#)
[Find similar images](#)




Here I go then, trying
600 × 600 - 35k - jpg
[beeper.wordpress.com](#)
[Find similar images](#)




The Chalk Mug »
304 × 314 - 17k - jpg
[coolest-gadgets.com](#)
[Find similar images](#)




mug
300 × 279 - 54k - jpg
[reynosawatch.org](#)




Bring your own
500 × 451 - 15k - jpg
[cookstownunited.ca](#)
[Find similar images](#)




ceramic mug
980 × 1024 - 30k - jpg
[diytrade.com](#)




Dual Purpose Drinking
490 × 428 - 16k - jpg
[freshome.com](#)
[Find similar images](#)




This coffee mug.
300 × 300 - 22k - jpg
[gizmodo.com](#)
[Find similar images](#)




Back to Ceramic
400 × 400 - 8k - jpg
[freshpromotions.com.au](#)
[Find similar images](#)



Coffee Mug as a
303 × 301 - 10k - jpg
[dustbowl.wordpress.com](#)
[Find similar images](#)




SASS Life Member
300 × 302 - 6k - jpg
[sassnet.com](#)



personalized coffee
400 × 343 - 15k - jpg
[walyou.com](#)
[Find similar images](#)

Google mugs



Mugs from LabelMe



WHAT BIAS? I AM SURE THAT MY MSRC CLASSIFIER WILL WORK ON ANY DATA!

1. Denial



RECOGNITION IS HOPELESS., IT WILL NEVER WORK. WE WILL JUST KEEP OVERFITTING TO THE NEXT DATASET...

3. Despair

OF COURSE THERE IS BIAS! THAT'S WHY YOU MUST ALWAYS TRAIN AND TEST ON THE SAME DATASET.



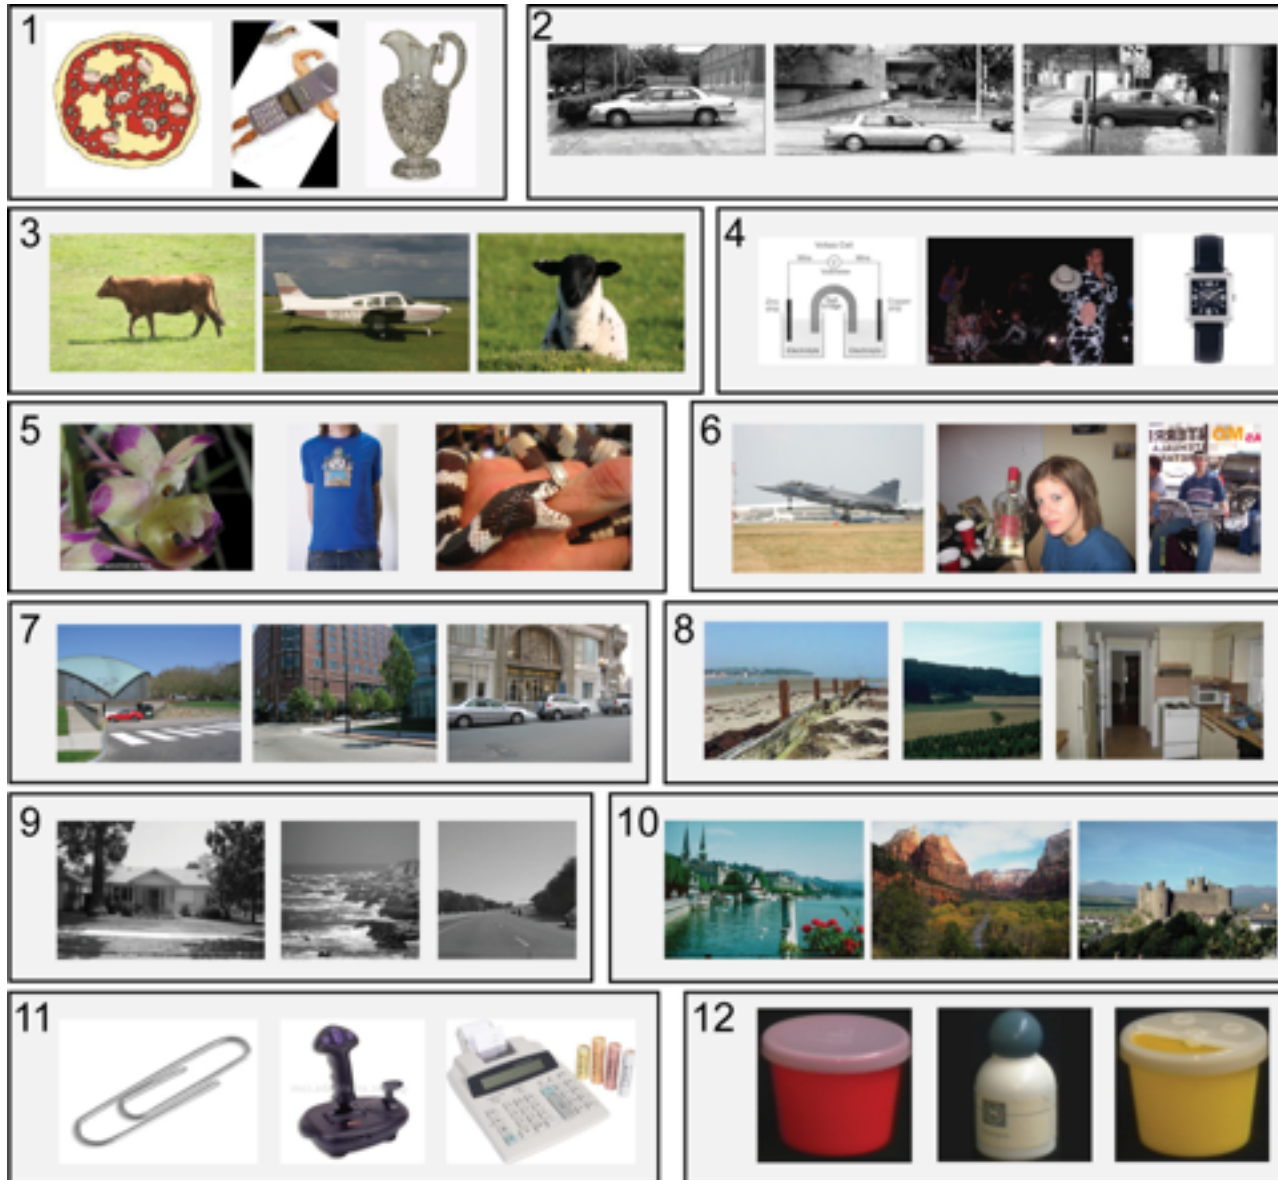
2. Machine Learning

BIAS IS HERE TO STAY, SO WE MUST BE VIGILANT THAT OUR ALGORITHMS DON'T GET DISTRACTED BY IT.



4. Acceptance

“Name That Dataset!” game



- Caltech 101
- Caltech 256
- MSRC
- UIUC cars
- Tiny Images
- Corel
- PASCAL 2007
- LabelMe
- COIL-100
- ImageNet
- 15 Scenes
- SUN'09


SVM plays “*Name that dataset!*”

SVM plays “*Name that dataset!*”

UIUC	0	29	8	21	5	10	2	17	6	3	2	0
LabelMe Spain	0	54		7	8	6		2	2		8	0
PASCAL 2007	0	10	29	10	10		7		7	7	11	1
MSRC	0	3	7	60		3			2		7	0
SUN09	0	14	9	9	24	17	11	3	3	3		0
15 Scenes	0	8	3		13	51	11	2	2	2	2	0
Corel	1	2	6		8	11	35	10	7	7	9	0
Caltech101	1	2	9	9	2		7	38	14	7	6	1
Caltech256	1	2	8				10	18	20	11	12	1
Tiny	1	2	8	9			11	12	13	24	12	1
ImageNet	1	3	11	9			11	8	12	13	21	1
COIL-100	0	0	0	0	0	0	0	0	0	0	0	99
	UIUC	LabelMe	PASCAL07	MSRC	SUN09	15 Scenes	Corel	Caltech101	Caltech256	Tiny	ImageNet	COIL-100

- 12 1-vs-all classifiers
- Standard full-image features
- 39% performance (chance is 8%)

SVM plays “*Name that dataset!*”



Datasets have different goals...

- Some are object-centric (e.g. Caltech, ImageNet)
- Otherwise are scene-centric (e.g. LabelMe, SUN'09)
- What about playing “*name that dataset*” on bounding boxes?

Cross-Dataset Generalization

MSRC



Classifier trained on MSRC cars

Mixing datasets

Test on PASCAL

