


VISUAL SIGNAL PROCESS & PIPELINE

Chih-Chung Hsu (許志仲)
Assistant Professor
ACVLab, Institute of Data Science
National Cheng Kung University
<https://cchsui.info>



Supervised learning

$$y = f(x)$$

output prediction function Image feature

▪ Training:

- given a *training set* of labeled examples $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)\}$, estimate the prediction function f by minimizing the prediction error on the training set

▪ Testing:

- apply f to a never before seen *test example* \mathbf{x} and output the predicted value $y = f(\mathbf{x})$

Image Categorization

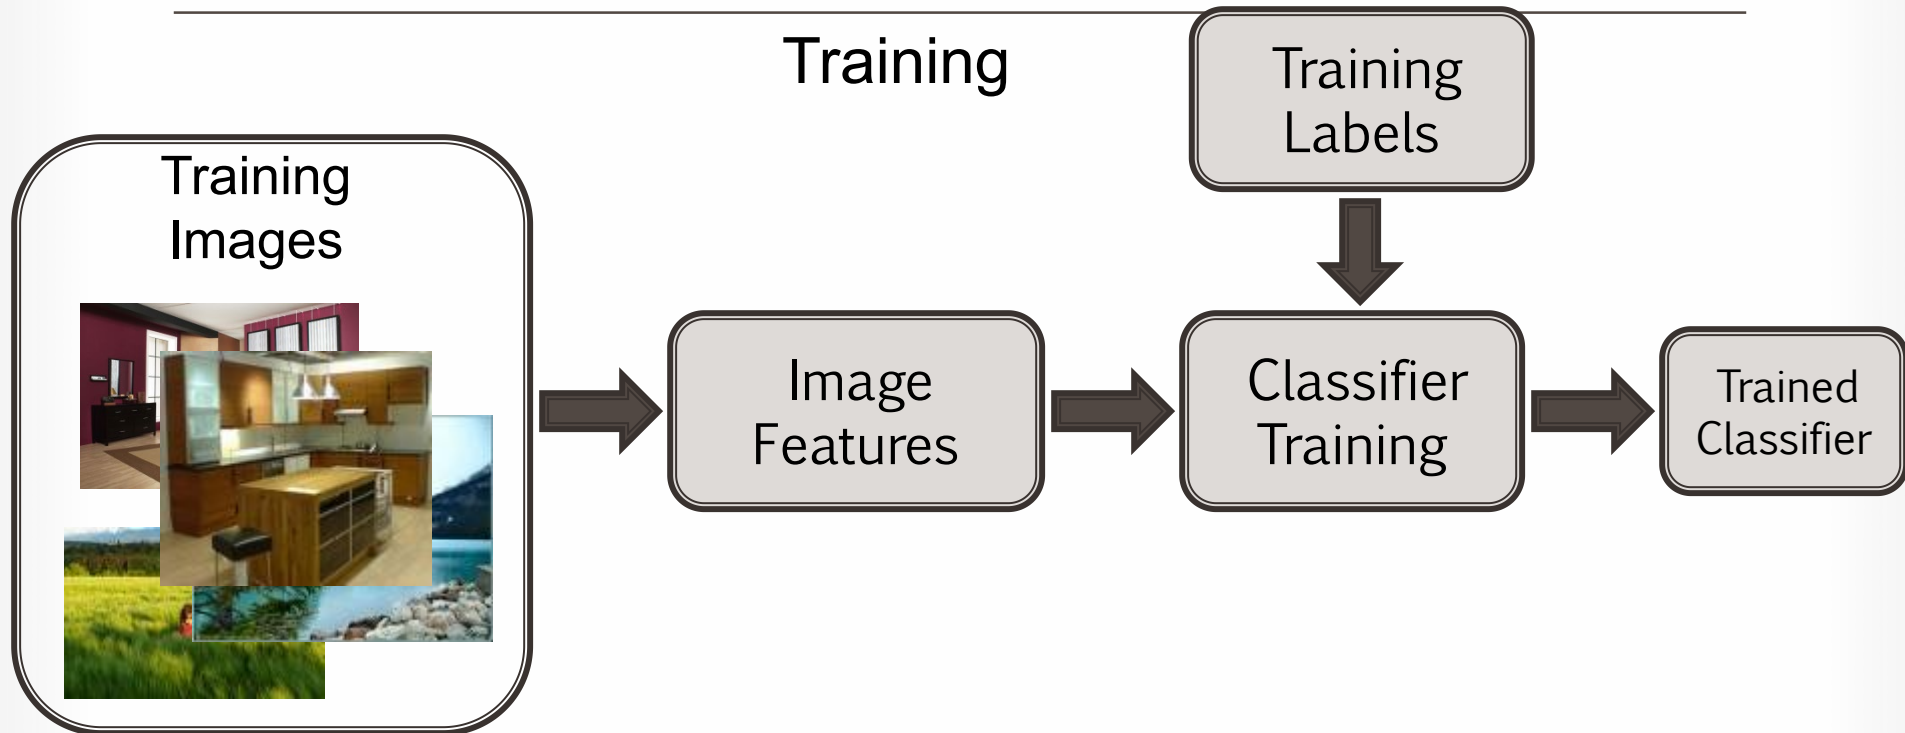
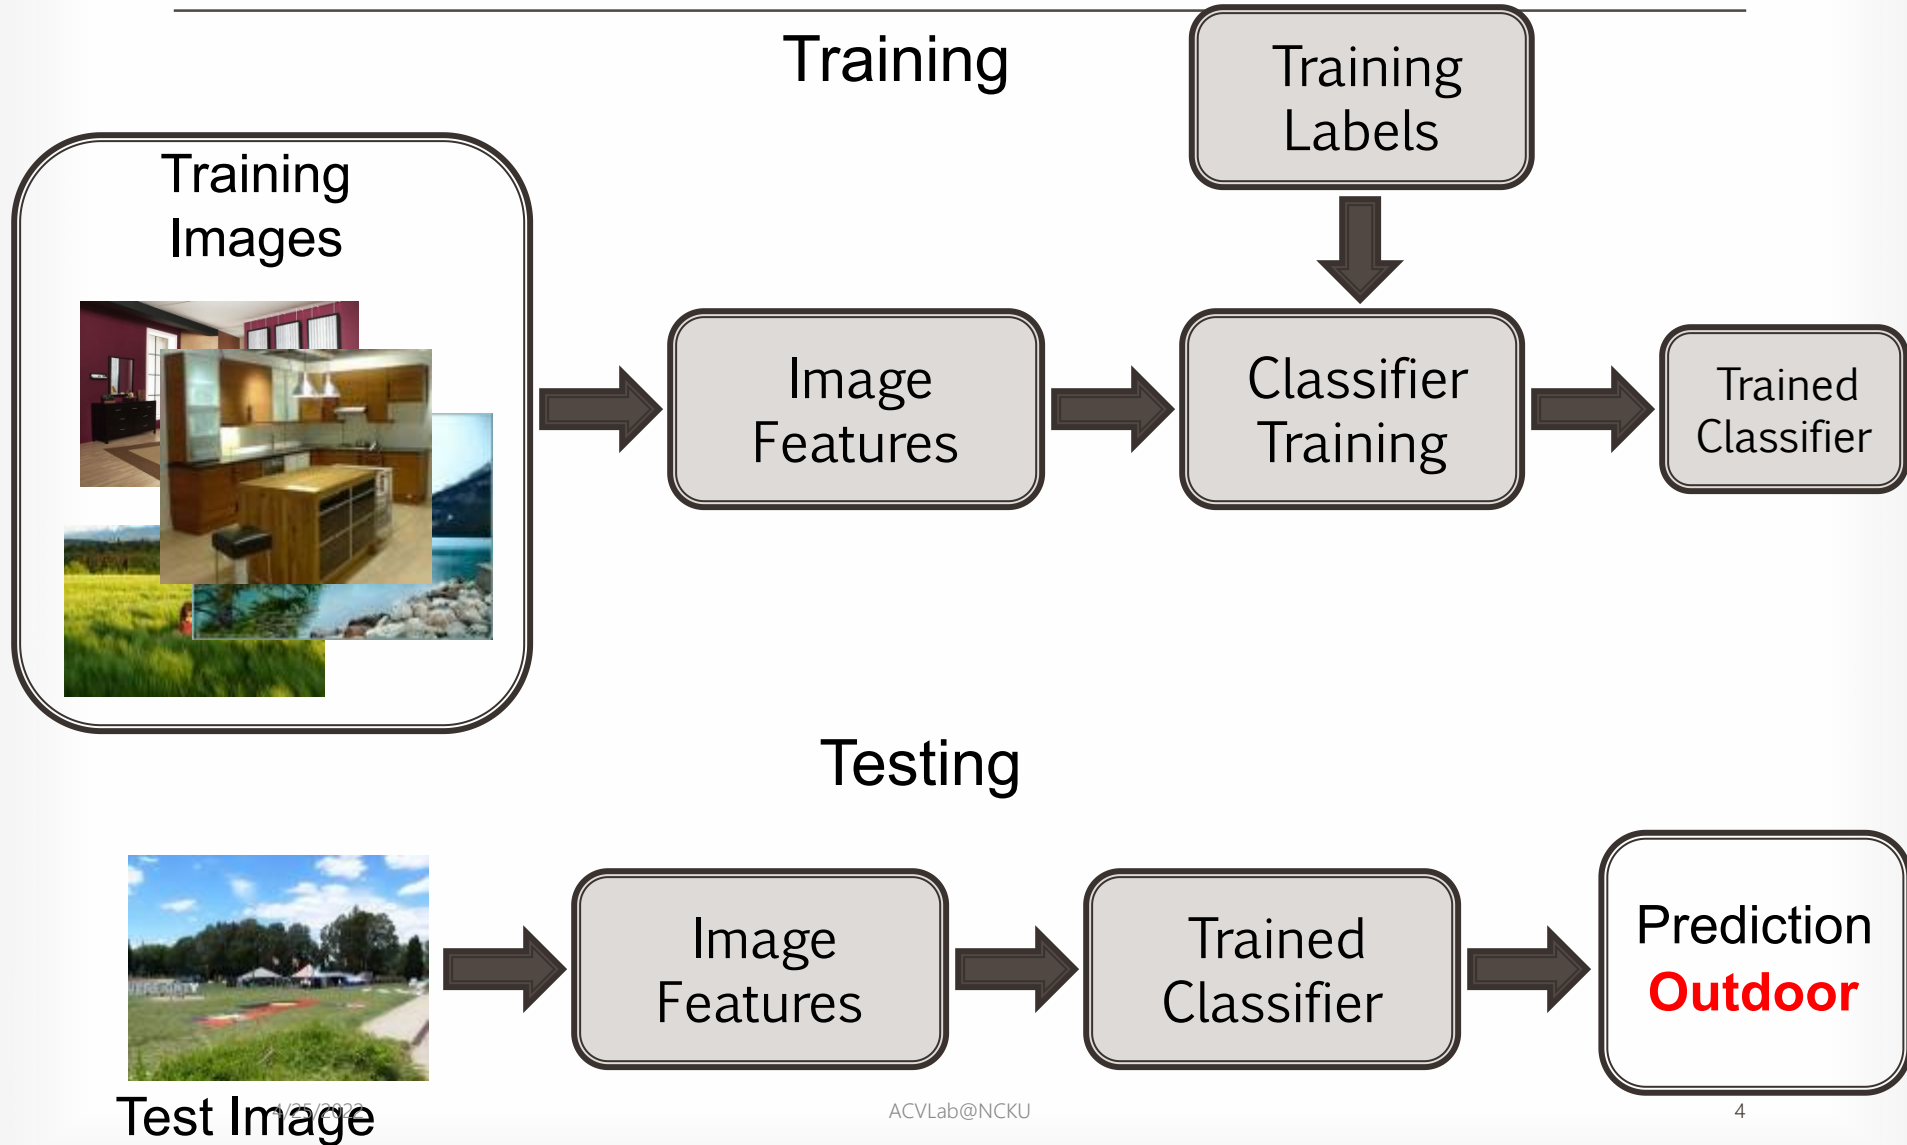


Image Categorization

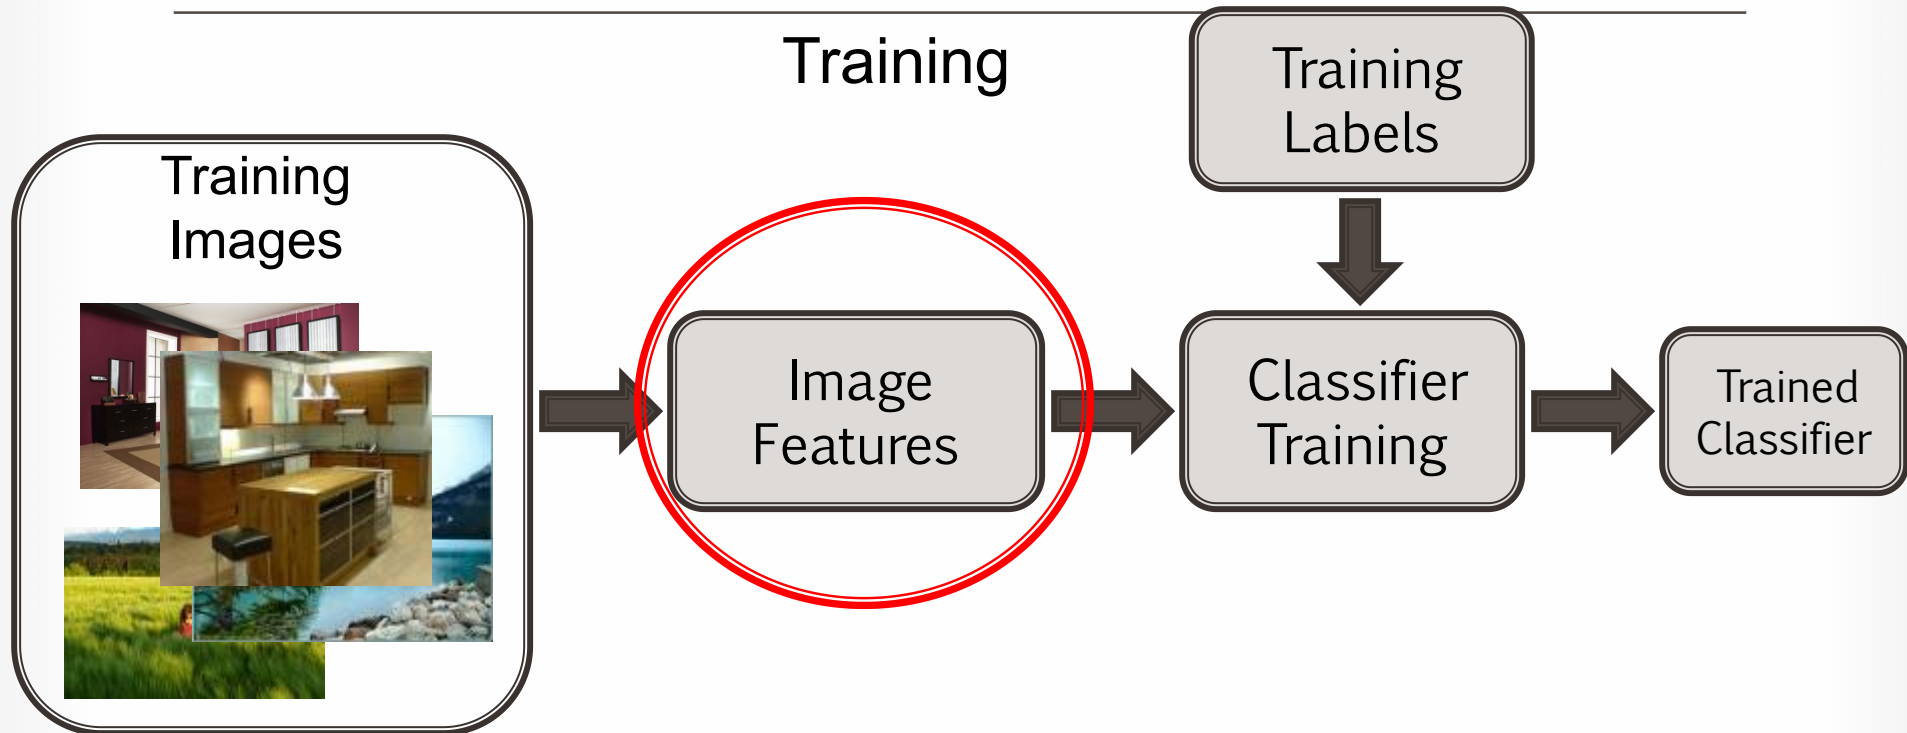


Example: Scene Categorization

- Is this a kitchen?



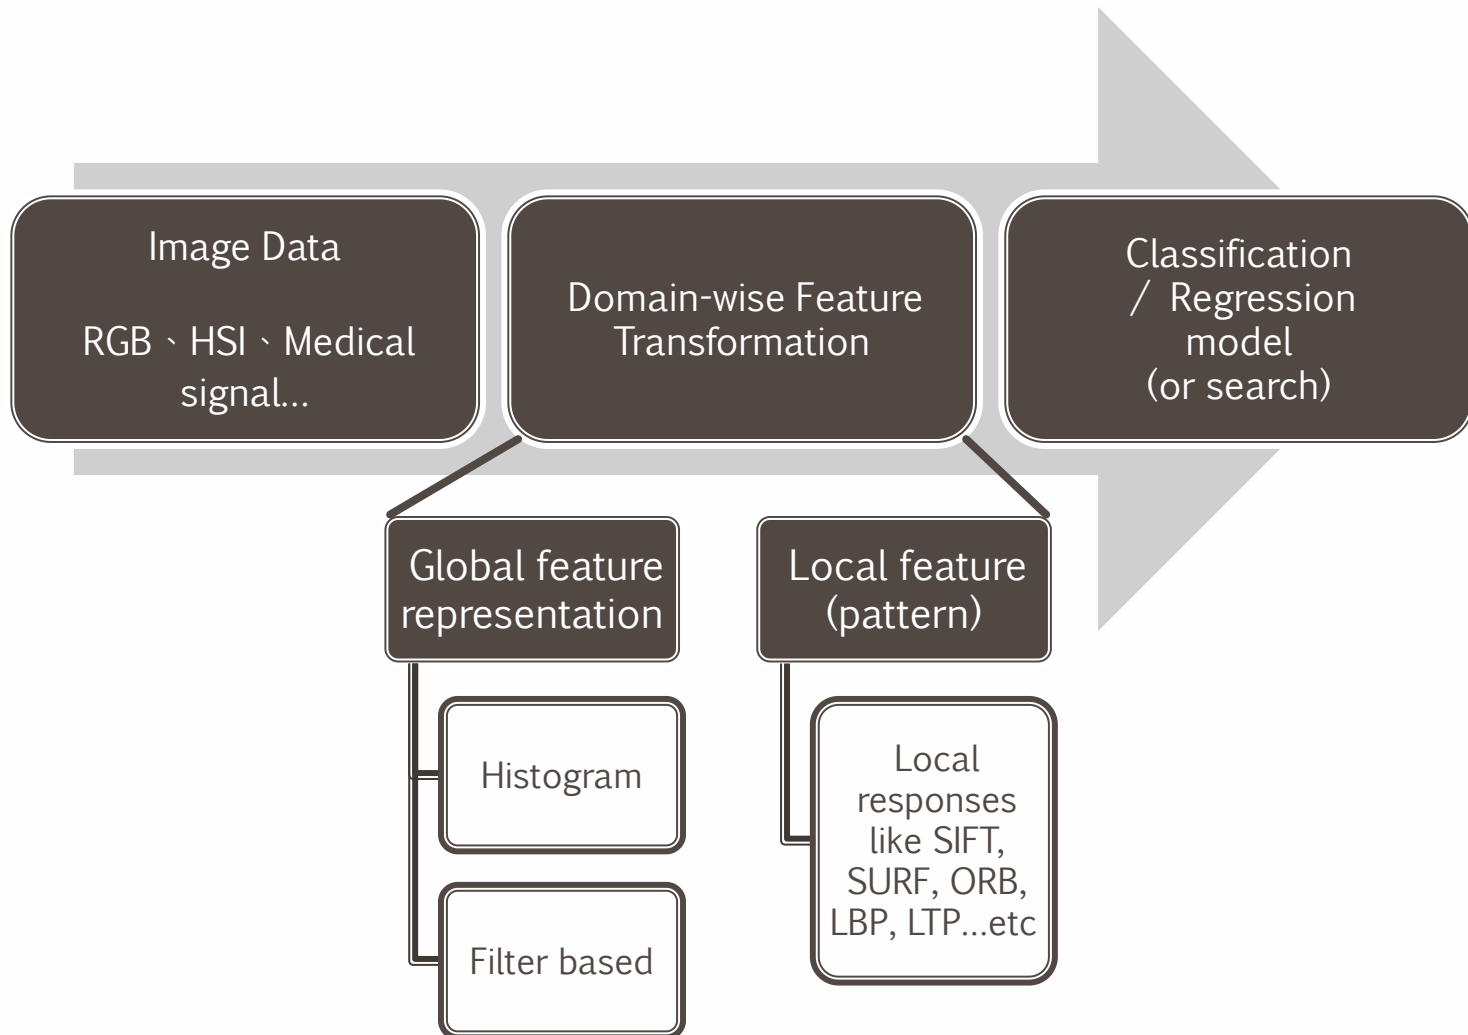
Image features





BASELINE: IMAGE RECOGNITION WITH FEATURES

Data Pipeline for Image Recognition





Feature-based Image Classification

- Global Feature
- It is nothing new.
 - Just use the feature extracted from images and followed by feeding the feature to the classifier.

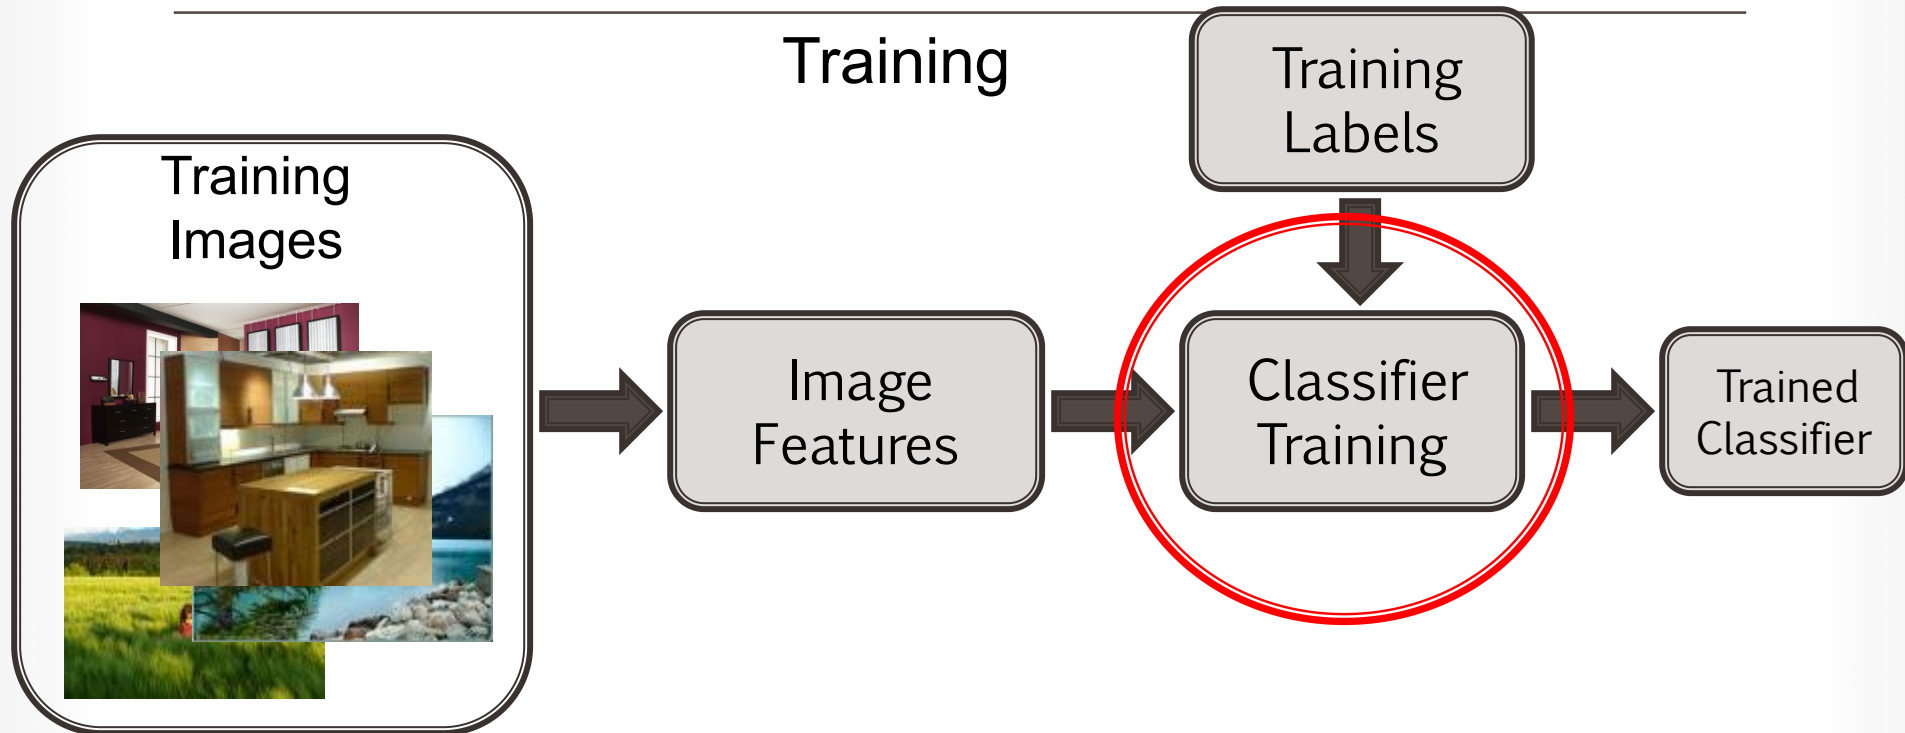
- Local Feature
 - #features per image fixed.
 - Similar to global one
 - #features variant
 - Matching is necessary....

Features

- Raw pixels
- Histograms
- GIST descriptors
- SIFT/SURF/LBP/HOG...
- Learning features from data

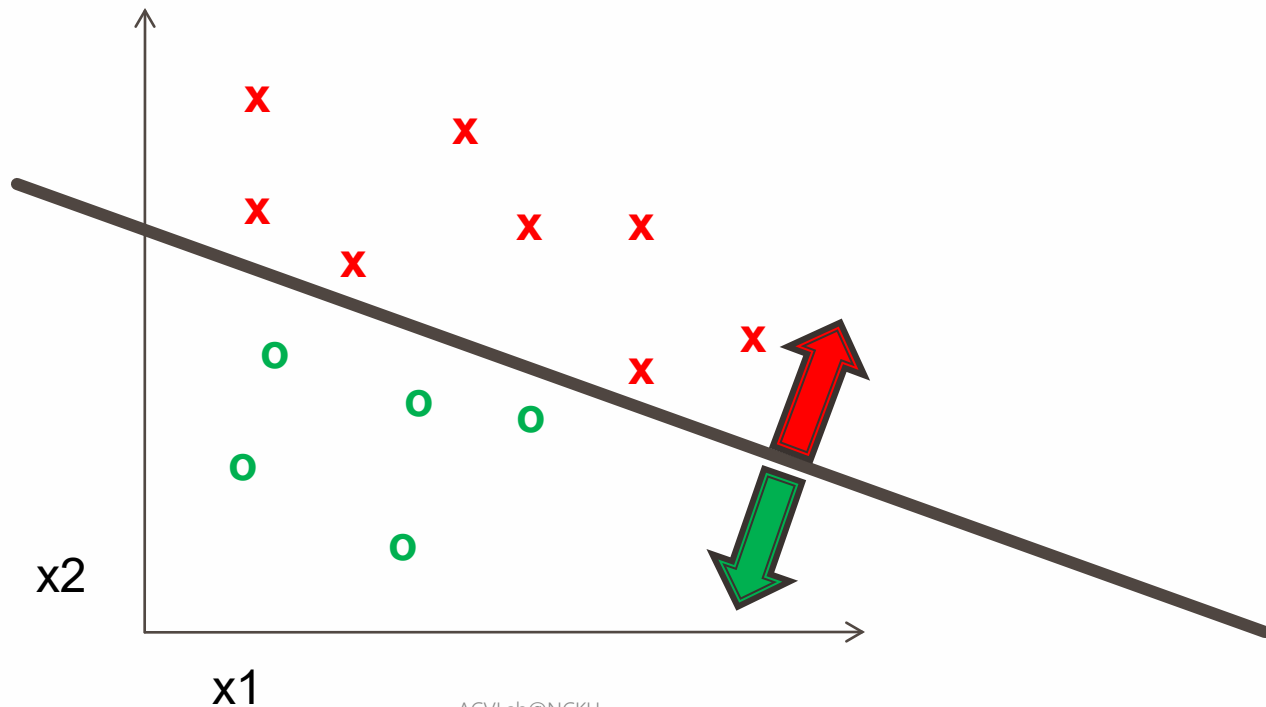


Classifiers




Learning a classifier

Given some set of features with corresponding labels, learn a function to predict the labels from the features



How many visual object categories are there?

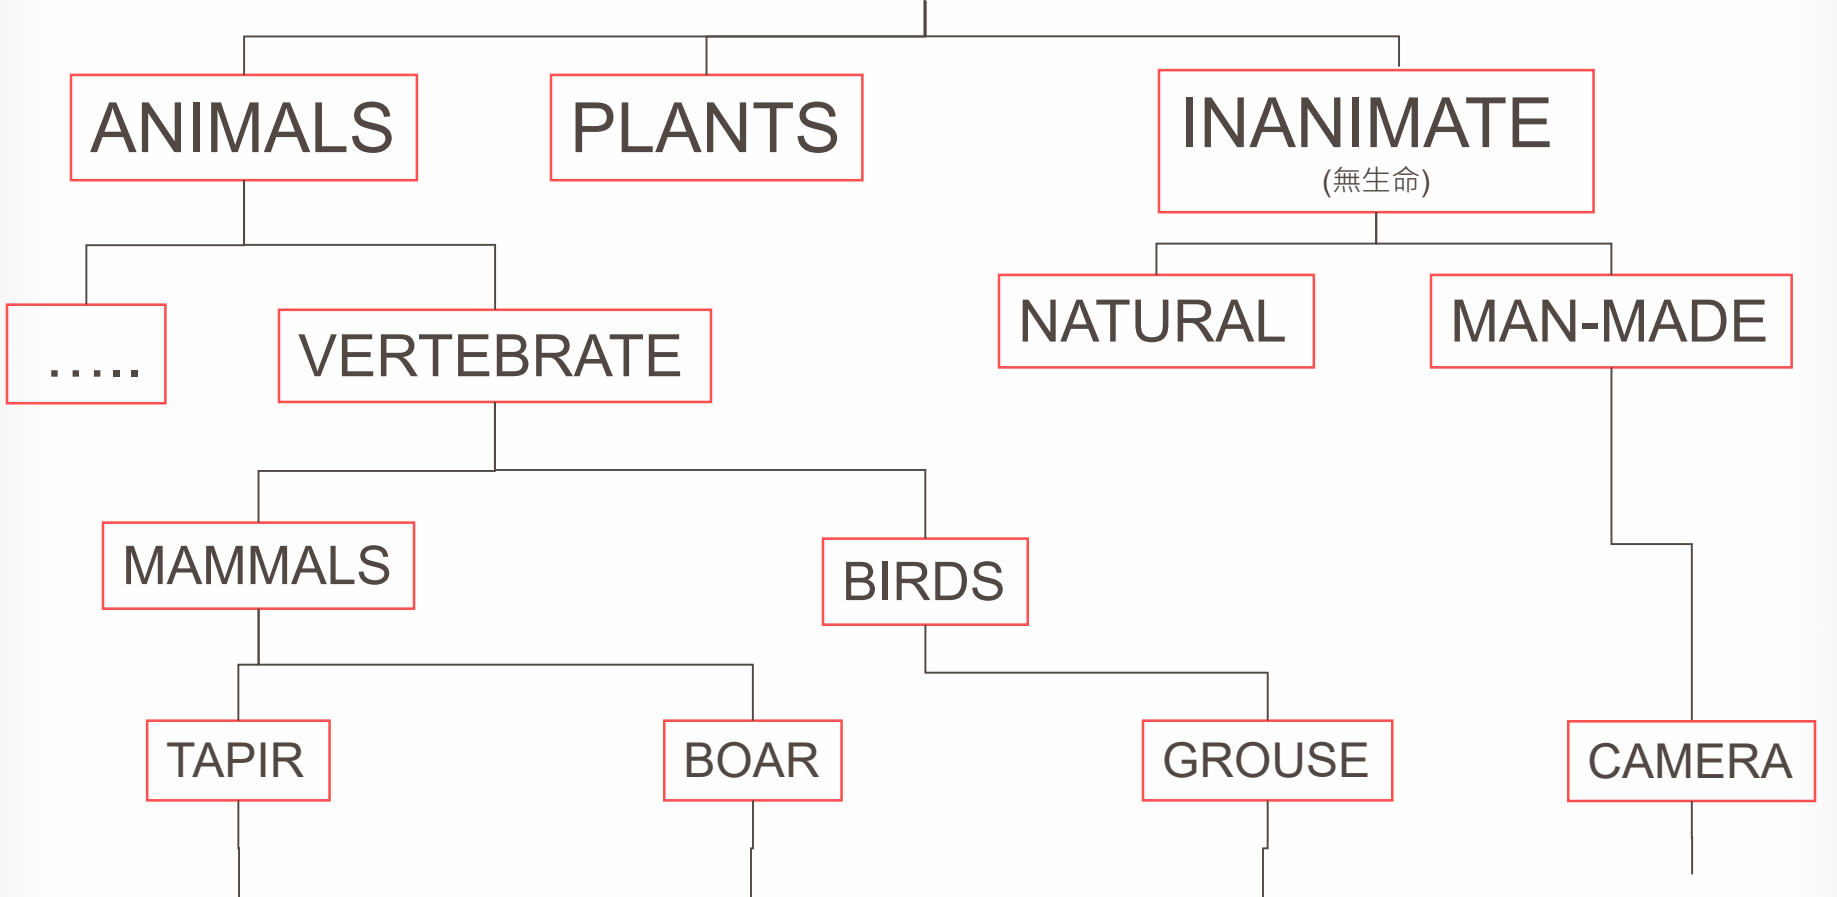




~10,000 to 30,000



OBJECTS



Specific recognition tasks



Scene categorization or classification



Image annotation / tagging / attributes




- street
- people
- building
- mountain
- tourism
- cloudy
- brick
- ...

Object detection

- find pedestrians




Image parsing / semantic segmentation



Scene understanding?


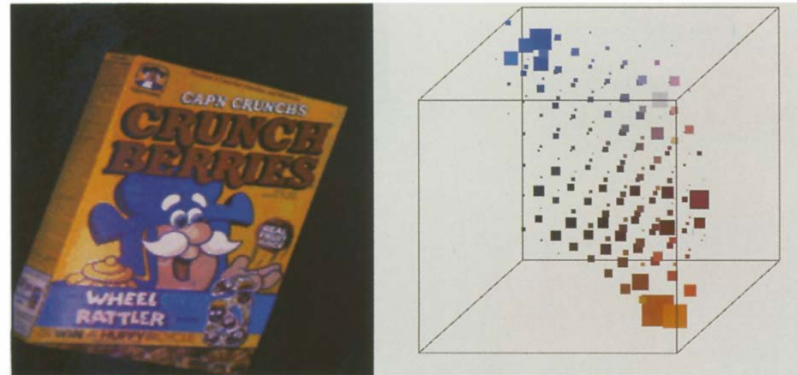


Eigenfaces (Turk & Pentland, 1991)



Experimental Condition	Correct/Unknown Recognition Percentage		
	Lighting	Orientation	Scale
Forced classification	96/0	85/0	64/0
Forced 100% accuracy	100/19	100/39	100/60
Forced 20% unknown rate	100/20	94/20	74/20

Color Histograms



4/25/2022 Swain and Ballard, Color Indexing, IJCV 1991.

History of ideas in recognition

- 1960s – early 1990s: the geometric era
- 1990s: appearance-based models
- 1990s – present: sliding window approaches

Sliding window approaches



History of ideas in recognition

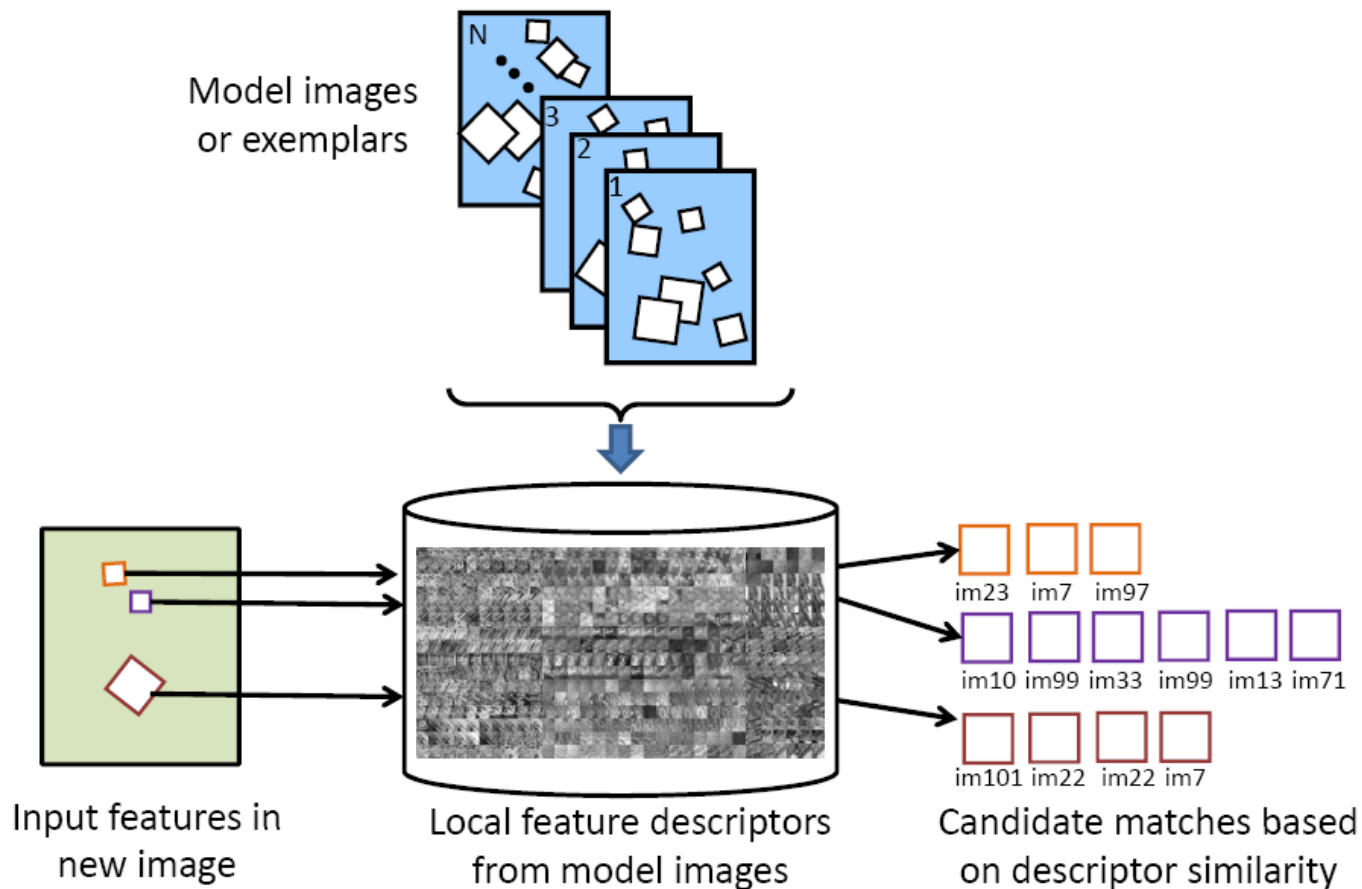
- 1960s – early 1990s: the geometric era
- 1990s: appearance-based models
- Mid-1990s: sliding window approaches
- Late 1990s: local features

Local features for object instance recognition




Large-scale image search

Combining local features, indexing, and spatial constraints



Large-scale image search

Combining local features, indexing, and spatial constraints





Large-scale image search

Combining local features, indexing, and spatial constraints

Google Goggles in Action

Click the icons below to see the different ways Google Goggles can be used.




History of ideas in recognition

- 1960s – early 1990s: the geometric era
- 1990s: appearance-based models
- Mid-1990s: sliding window approaches
- Late 1990s: local features
- Early 2000s: parts-and-shape models

Parts-and-shape models

- Model:
 - Object as a set of parts
 - Relative locations between parts
 - Appearance of part



Representing people


Pictorial structure model

Fischler and Elschlager(73), Felzenszwalb and Huttenlocher(00)

$$\Pr(P_{tor}, P_{arm}, \dots | \text{Im}) \propto \prod_{i,j} \Pr(P_i | P_j) \prod_i \Pr(\text{Im}(P_i))$$



\uparrow
 part geometry

\uparrow
 part appearance



Discriminatively trained part-based models



- P. Felzenszwalb, R. Girshick, D. McAllester, D. Ramanan, "Object Detection with Discriminatively Trained Part-Based Models," PAMI 2009



History of ideas in recognition

- 1960s – early 1990s: the geometric era
- 1990s: appearance-based models
- Mid-1990s: sliding window approaches
- Late 1990s: local features
- Early 2000s: parts-and-shape models
- Mid-2000s: bags of features

Bag-of-features models




Bag-of-features models



Objects as texture




- All of these are treated as being the same










- No distinction between foreground and background: scene recognition?

Origin 1: Texture recognition

- Texture is characterized by the repetition of basic elements or *textons*
- For stochastic textures, it is the identity of the textons, not their spatial arrangement, that matters



Origin 1: Texture recognition



Origin 2: Bag-of-words models

- Orderless document representation: frequencies of words from a dictionary Salton & McGill (1983)

Origin 2: Bag-of-words models

- Orderless document representation: frequencies of words from a

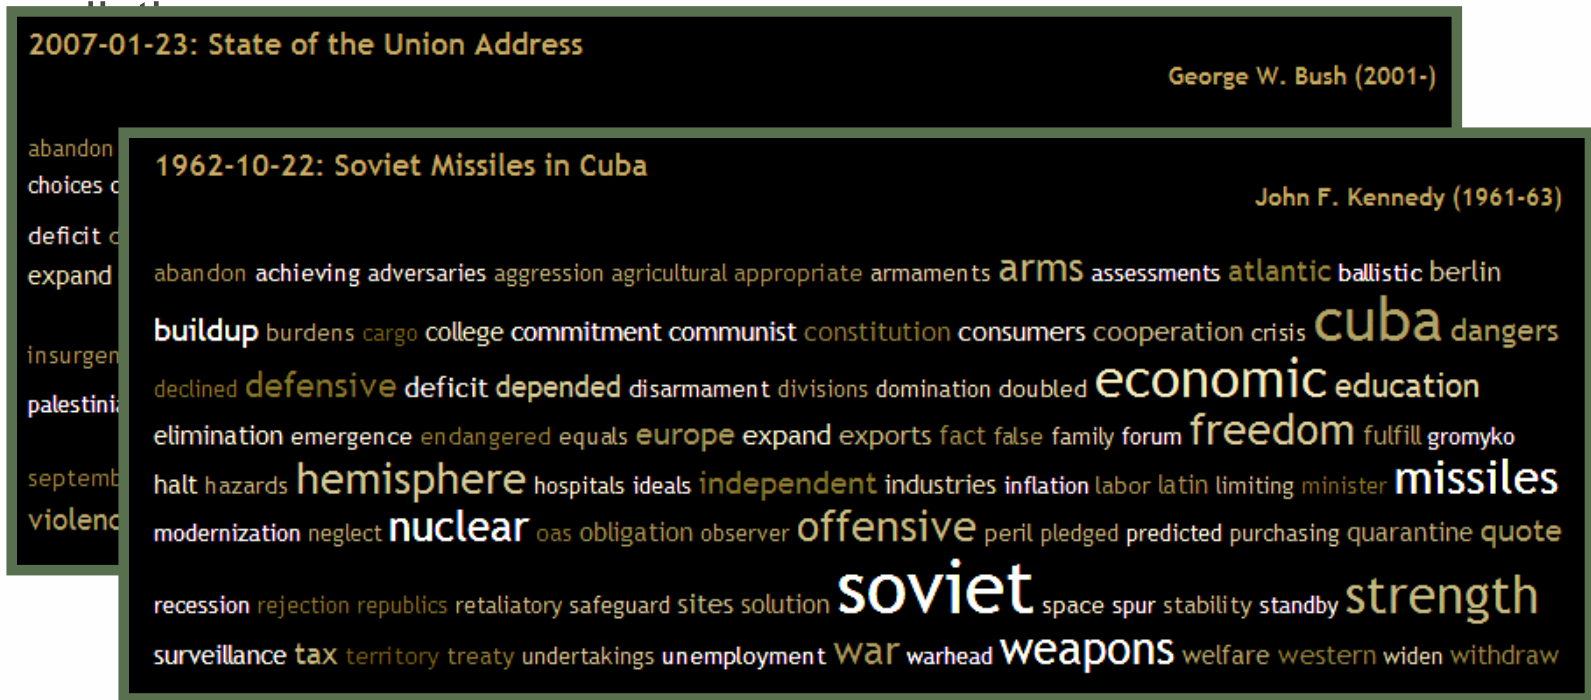
2007-01-23: State of the Union Address

George W. Bush (2001-)

abandon accountable **affordable** **afghanistan** africa aided ally **anbar** armed army **baghdad** bless **challenges** chamber chaos
 choices civilians coalition **commanders** **commitment** confident confront congressman constitution corps debates deduction
 deficit deliver **democratic** deploy dikembe diplomacy disruptions earmarks **economy** einstein **elections** eliminates
 expand **extremists** failing faithful families **freedom** fuel **funding** god haven ideology immigration impose
 insurgents **iran** **iraq** islam julie lebanon love madam marine math **medicare** moderation neighborhoods nuclear offensive
 palestinian payroll province pursuing **qaeda** radical regimes resolve retreat **rieman** sacrifices science sectarian senate
 september **shia** stays strength students succeed sunni **tax** territories **terrorists** threats uphold victory
 violence violent **war** washington weapons wesley

Origin 2: Bag-of-words models

- Orderless document representation: frequencies of words from a



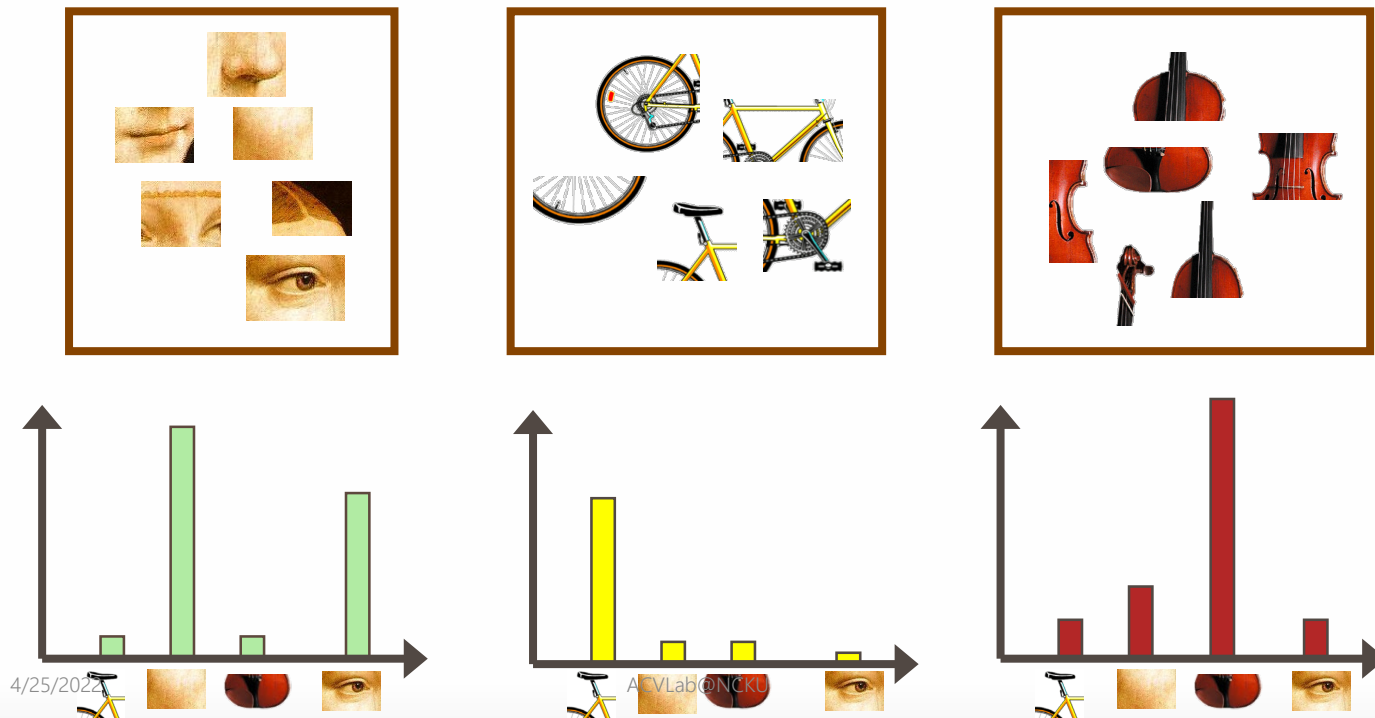
Origin 2: Bag-of-words models

- Orderless document representation: frequencies of words from a



Bag-of-features steps

- Extract features
- Learn “visual vocabulary”
- Quantize features using visual vocabulary
- Represent images by frequencies of “visual words”

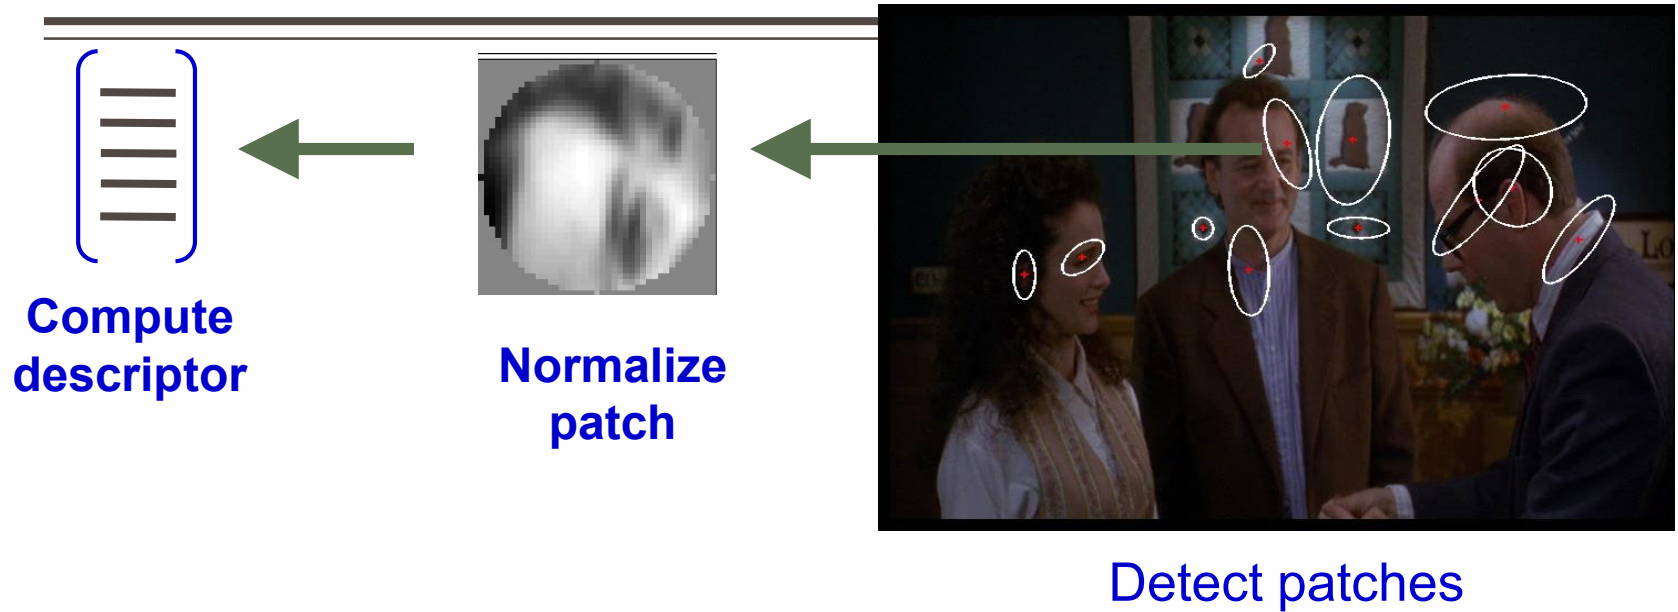


1. Feature extraction



- Regular grid or interest regions



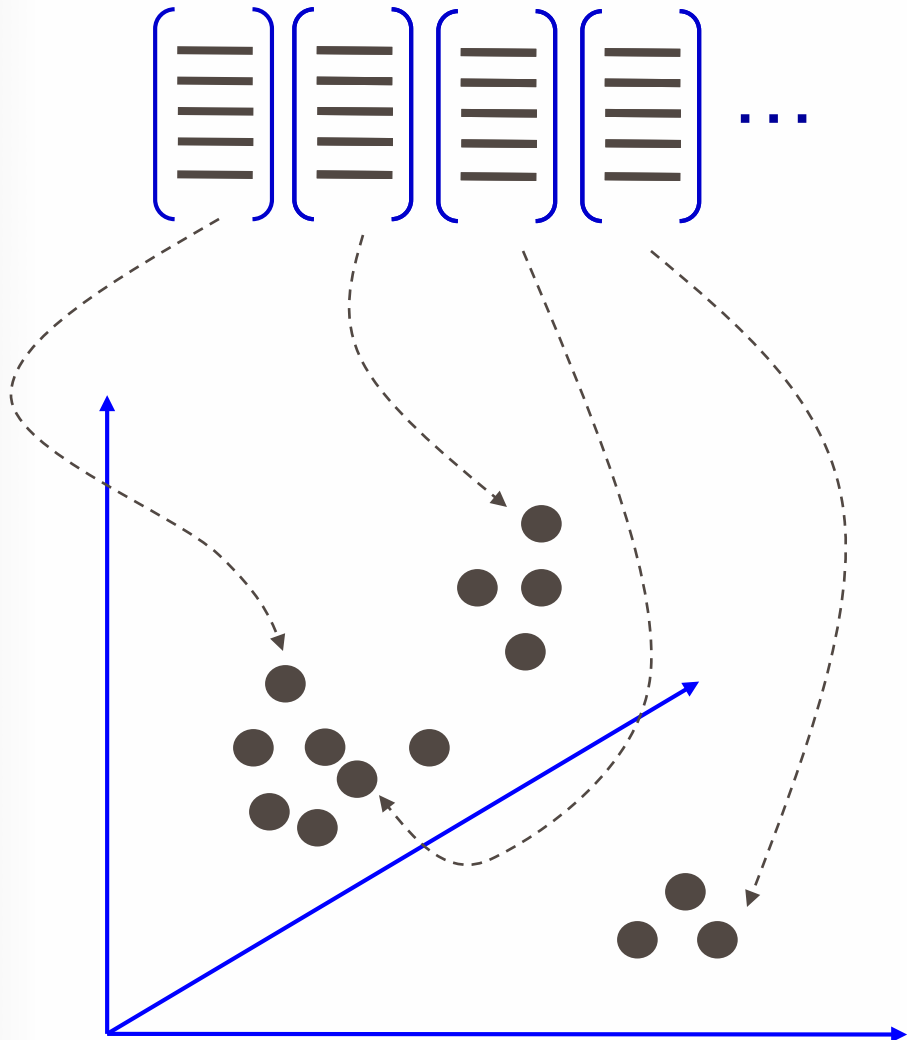
1. Feature extraction




1. Feature extraction



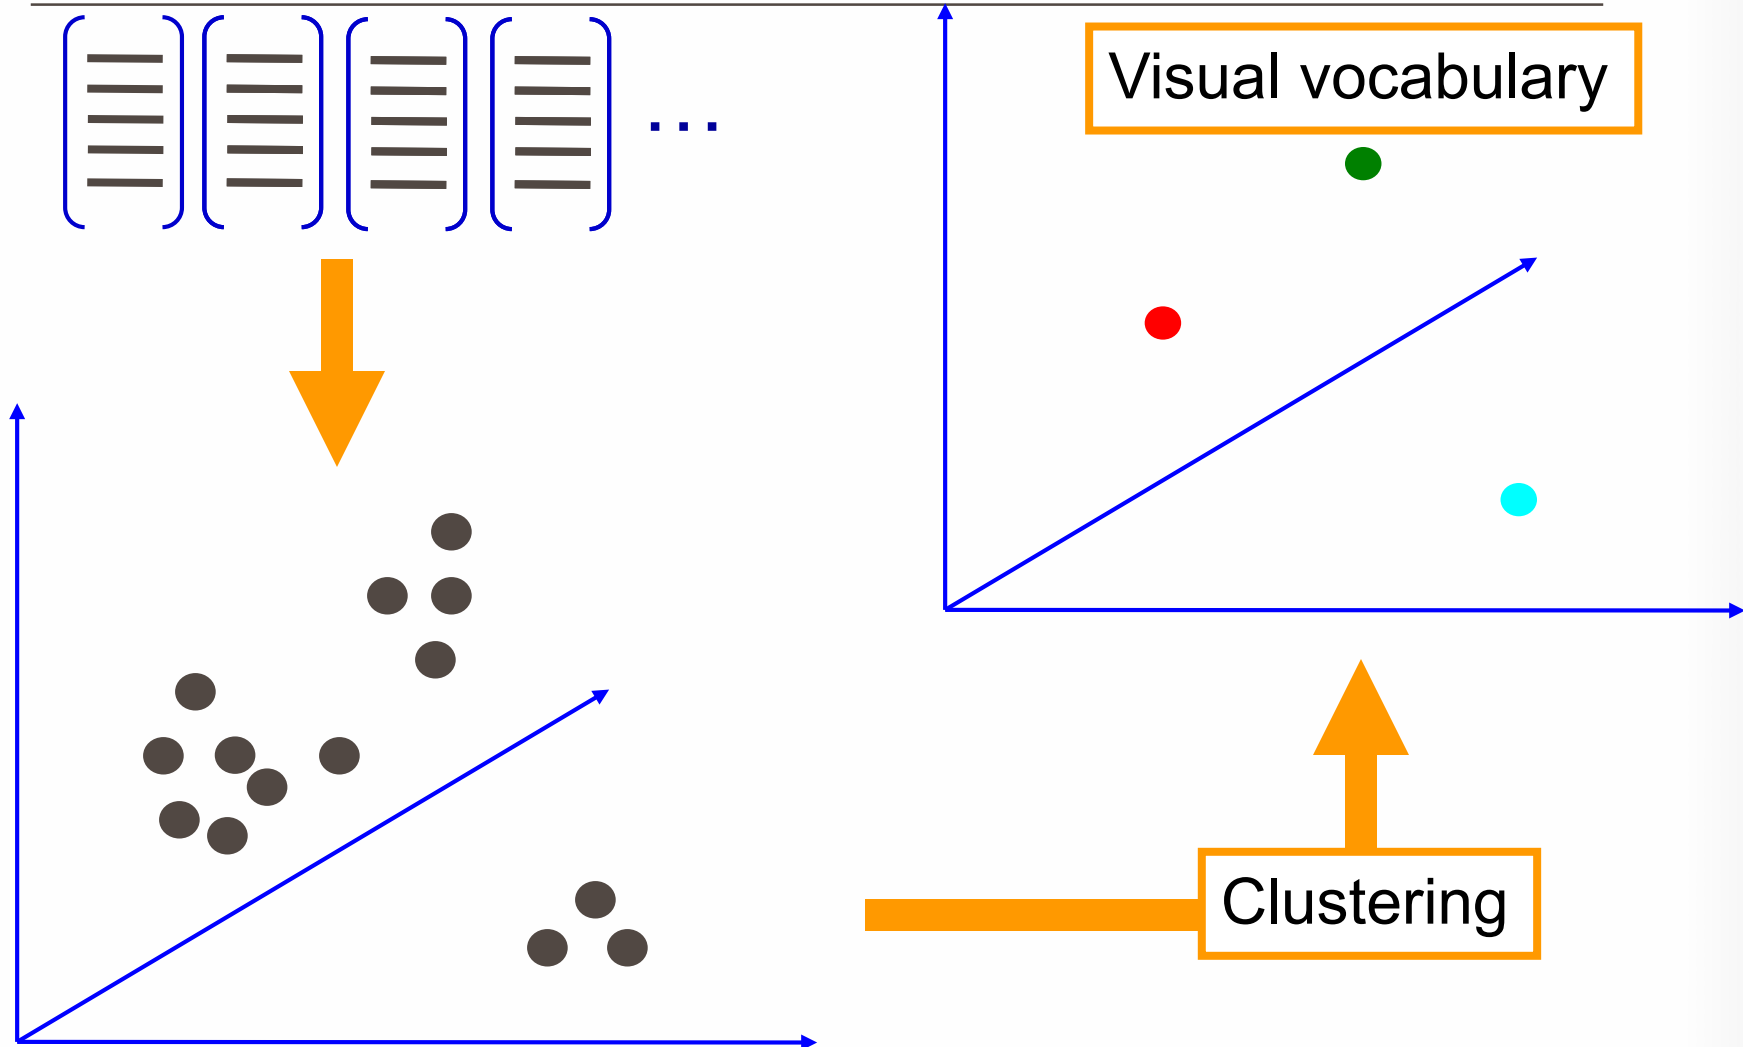
2. Learning the visual vocabulary



2. Learning the visual vocabulary



2. Learning the visual vocabulary



K-means clustering

- Want to minimize sum of squared Euclidean distances between points x_i and their nearest cluster centers m_k



$$D(X, M) = \sum_{\text{cluster } k} \sum_{\substack{\text{point } i \text{ in} \\ \text{cluster } k}} (x_i - m_k)^2$$

- Algorithm:
 - Randomly initialize K cluster centers
 - Iterate until convergence:
 - Assign each data point to the nearest center
 - Recompute each cluster center as the mean of all points assigned to it

Clustering and vector quantization

- Clustering is a common method for learning a visual vocabulary or codebook
 - Unsupervised learning process
 - Each cluster center produced by k-means becomes a codevector
 - Codebook can be learned on separate training set
 - Provided the training set is sufficiently representative, the codebook will be “universal”
- The codebook is used for quantizing features
 - A vector quantizer takes a feature vector and maps it to the index of the nearest codevector in a codebook
 - Codebook = visual vocabulary
 - Codevector = visual word


Example codebook



...

Appearance codebook


Another codebook





Appearance codebook

Visual vocabularies: Issues

- How to choose vocabulary size?
 - Too small: visual words not representative of all patches
 - Too large: quantization artifacts, overfitting
- Computational efficiency
 - Vocabulary trees (Nister & Stewenius, 2006)



But what about layout?



All of these images have the same color histogram


Spatial pyramid



Compute histogram in each spatial bin

Spatial pyramid representation


- Extension of a bag of features
- Locally orderless representation at several levels of resolution




level 0

Spatial pyramid representation

- Extension of a bag of features
- Locally orderless representation at several levels of resolution





level 0




level 1

Spatial pyramid representation


- Extension of a bag of features
- Locally orderless representation at several levels of resolution



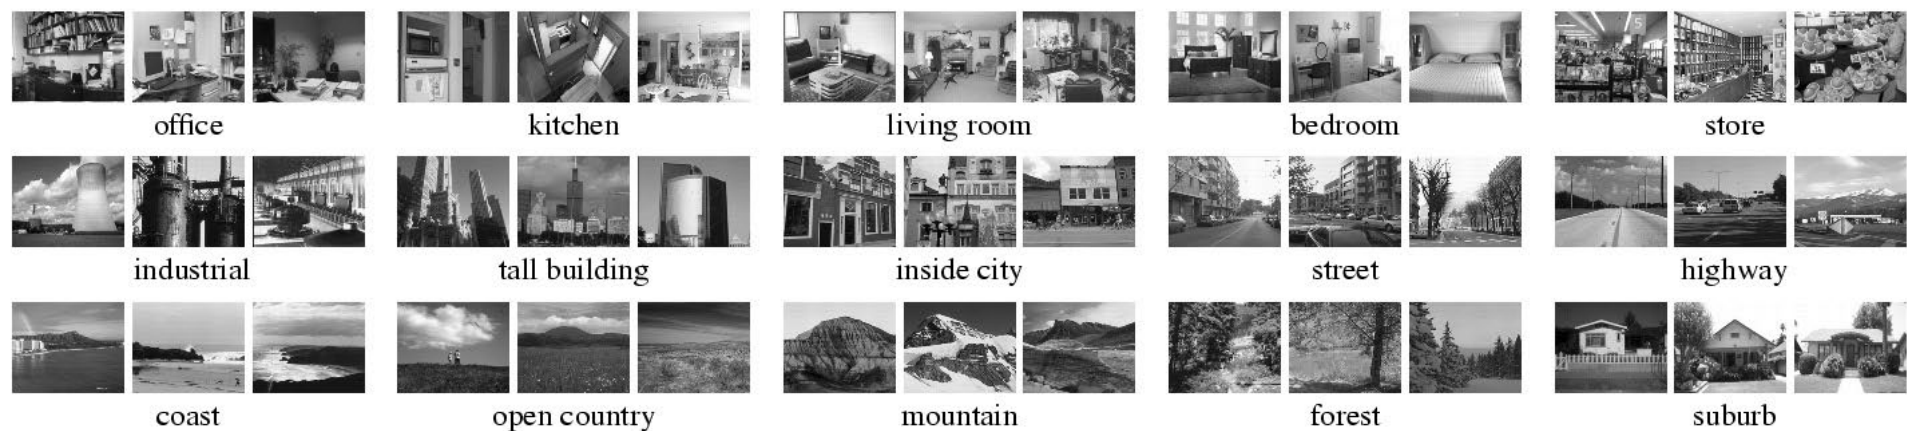
level 0



level 1




level 2



Multi-class classification results (100 training images per class)

Level	Weak features (vocabulary size: 16)		Strong features (vocabulary size: 200)	
	Single-level	Pyramid	Single-level	Pyramid
0 (1 × 1)	45.3 ±0.5		72.2 ±0.6	
1 (2 × 2)	53.6 ±0.3	56.2 ±0.6	77.9 ±0.6	79.0 ±0.5
2 (4 × 4)	61.7 ±0.6	64.7 ±0.7	79.4 ±0.3	81.1 ±0.3
3 (8 × 8)	63.3 ±0.8	66.8 ±0.6	77.2 ±0.4	80.7 ±0.3



Caltech101 dataset



Multi-class classification results (30 training images per class)


	Weak features (16)		Strong features (200)	
Level	Single-level	Pyramid	Single-level	Pyramid
0	15.5 ±0.9		41.2 ±1.2	
1	31.4 ±1.2	32.8 ±1.3	55.9 ±0.9	57.0 ±0.8
2	47.2 ±1.1	49.3 ±1.4	63.6 ±0.9	64.6 ±0.8
3	52.2 ±0.8	54.0 ±1.1	60.3 ±0.9	64.6 ±0.7

Multi-view matching



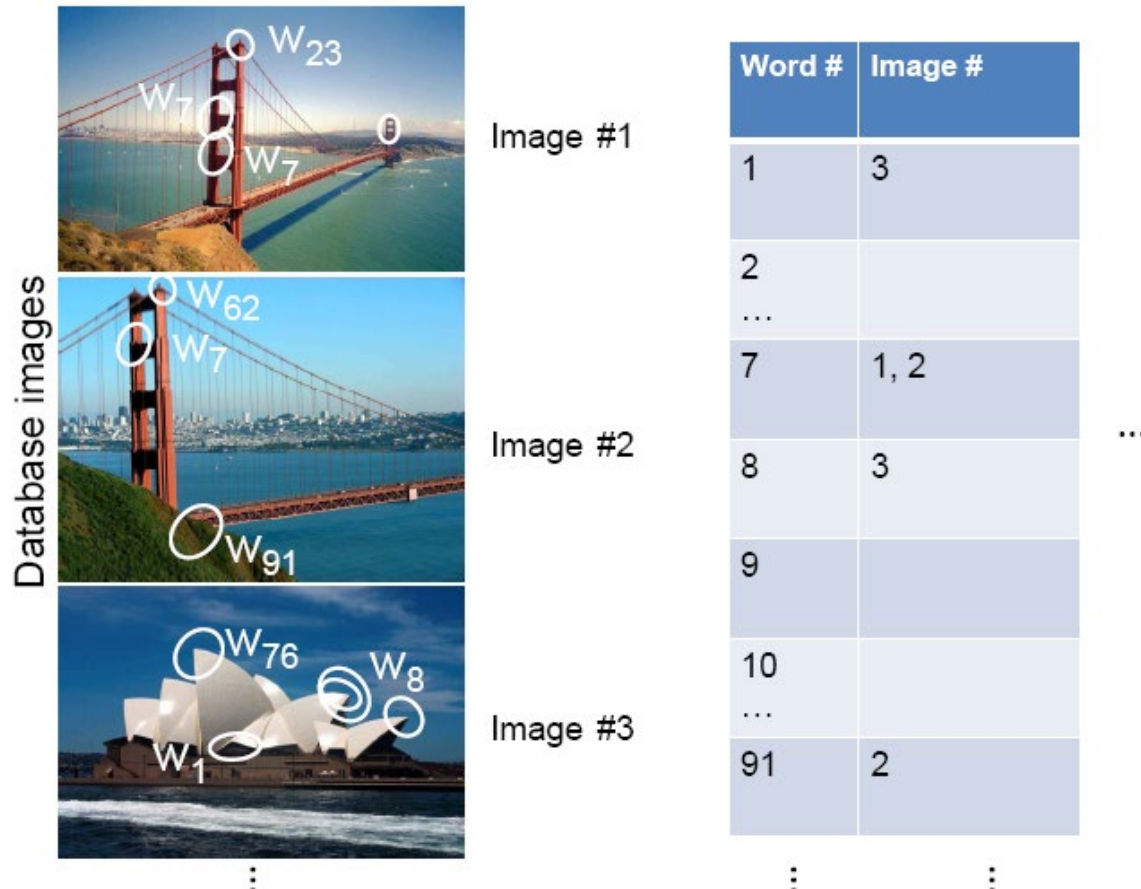
Matching two given views for depth

vs



Search for a matching view for recognition

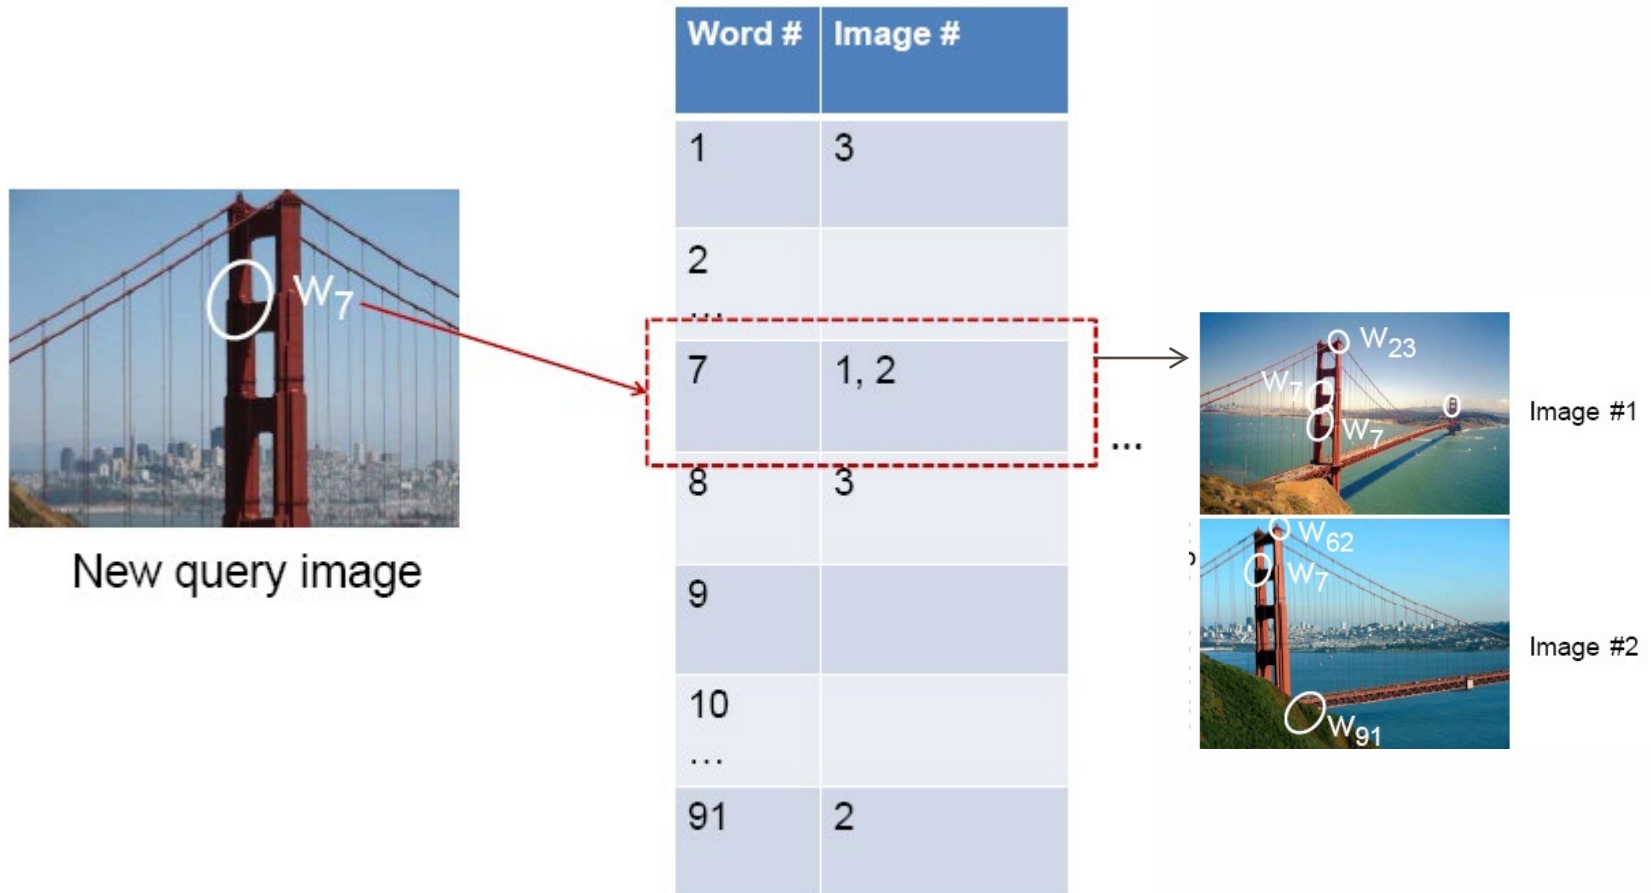
Inverted file index



Database images are loaded into the index mapping words to image numbers

Inverted file index

- New query image is mapped to indices of database images that share a word.




Inverted file index

- Key requirement for inverted file index to be efficient: sparsity
- If most pages/images contain most words then you're no better off than exhaustive search.
 - Exhaustive search would mean comparing the word distribution of a query versus every page.

Comparing bags of words

- Rank frames by normalized scalar product between their (possibly weighted) occurrence counts---nearest neighbor search for similar images.



[1 8 1 4]



\vec{d}_j

4/25/2022

[5 1 1 0]



\vec{q}

$$\text{sim}(d_j, q) = \frac{\langle d_j, q \rangle}{\|d_j\| \|q\|}$$


$$= \frac{\sum_{i=1}^V d_j(i) * q(i)}{\sqrt{\sum_{i=1}^V d_j(i)^2} * \sqrt{\sum_{i=1}^V q(i)^2}}$$

for vocabulary of V words


Comparing bags of words

- Other common histogram comparisons:

[1 8 1 4]



[5 1 1 0]



\vec{d}_j



\vec{q}

Histogram intersection

$$\frac{\sum_{j=1}^n \min(I_j, M_j)}{\sum_{j=1}^n M_j}$$


Chi squared

$$\sum_{i=1}^n \frac{(x_i - y_i)^2}{(x_i + y_i)}$$


Instance recognition: remaining issues

- How to summarize the content of an entire image? And gauge overall similarity?
- How large should the vocabulary be? How to perform quantization efficiently?
- Is having the same set of visual words enough to identify the object/scene? How to verify spatial agreement?
- How to score the retrieval results?

Vocabulary size



Results for recognition task with 6347 images




Influence on performance, sparsity


Nister & Stewenius, CVPR 2006





BOOSTED IMAGE MATCHING


Following slides by David Nister (CVPR 2006)











































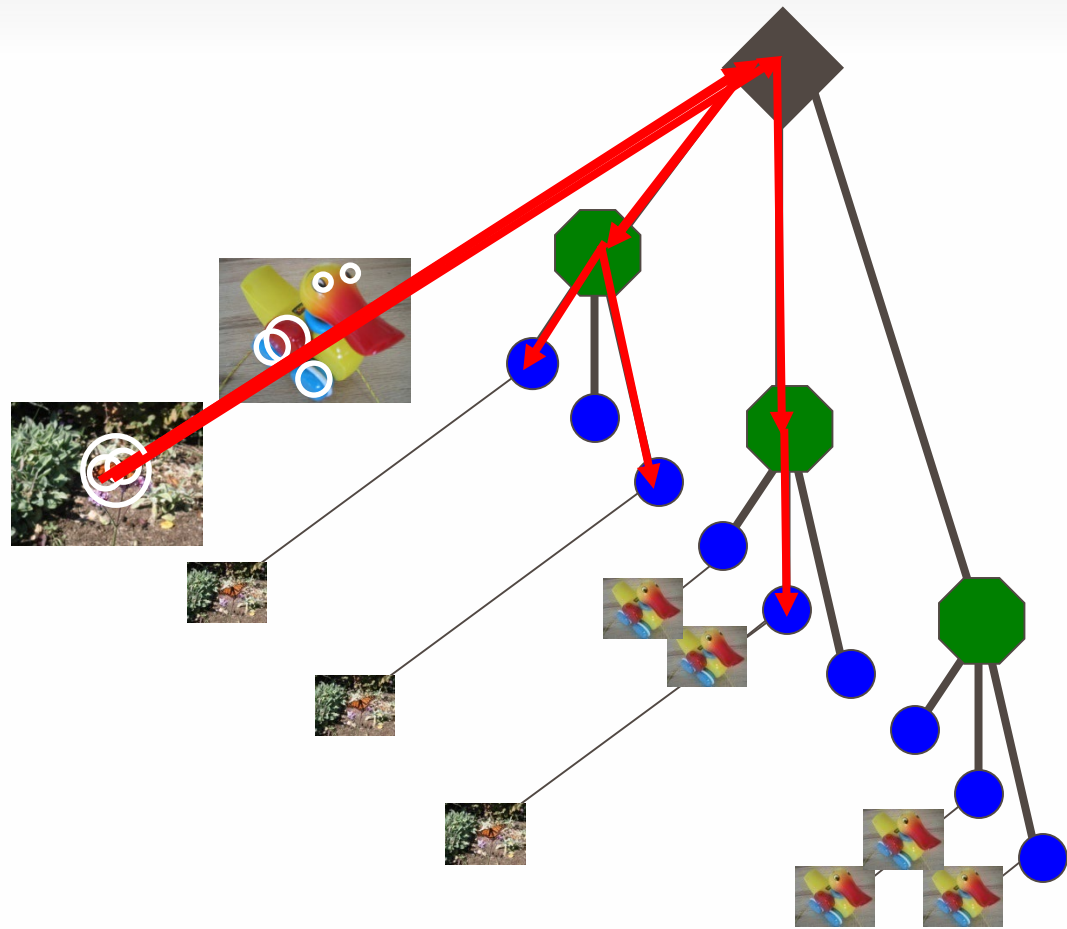





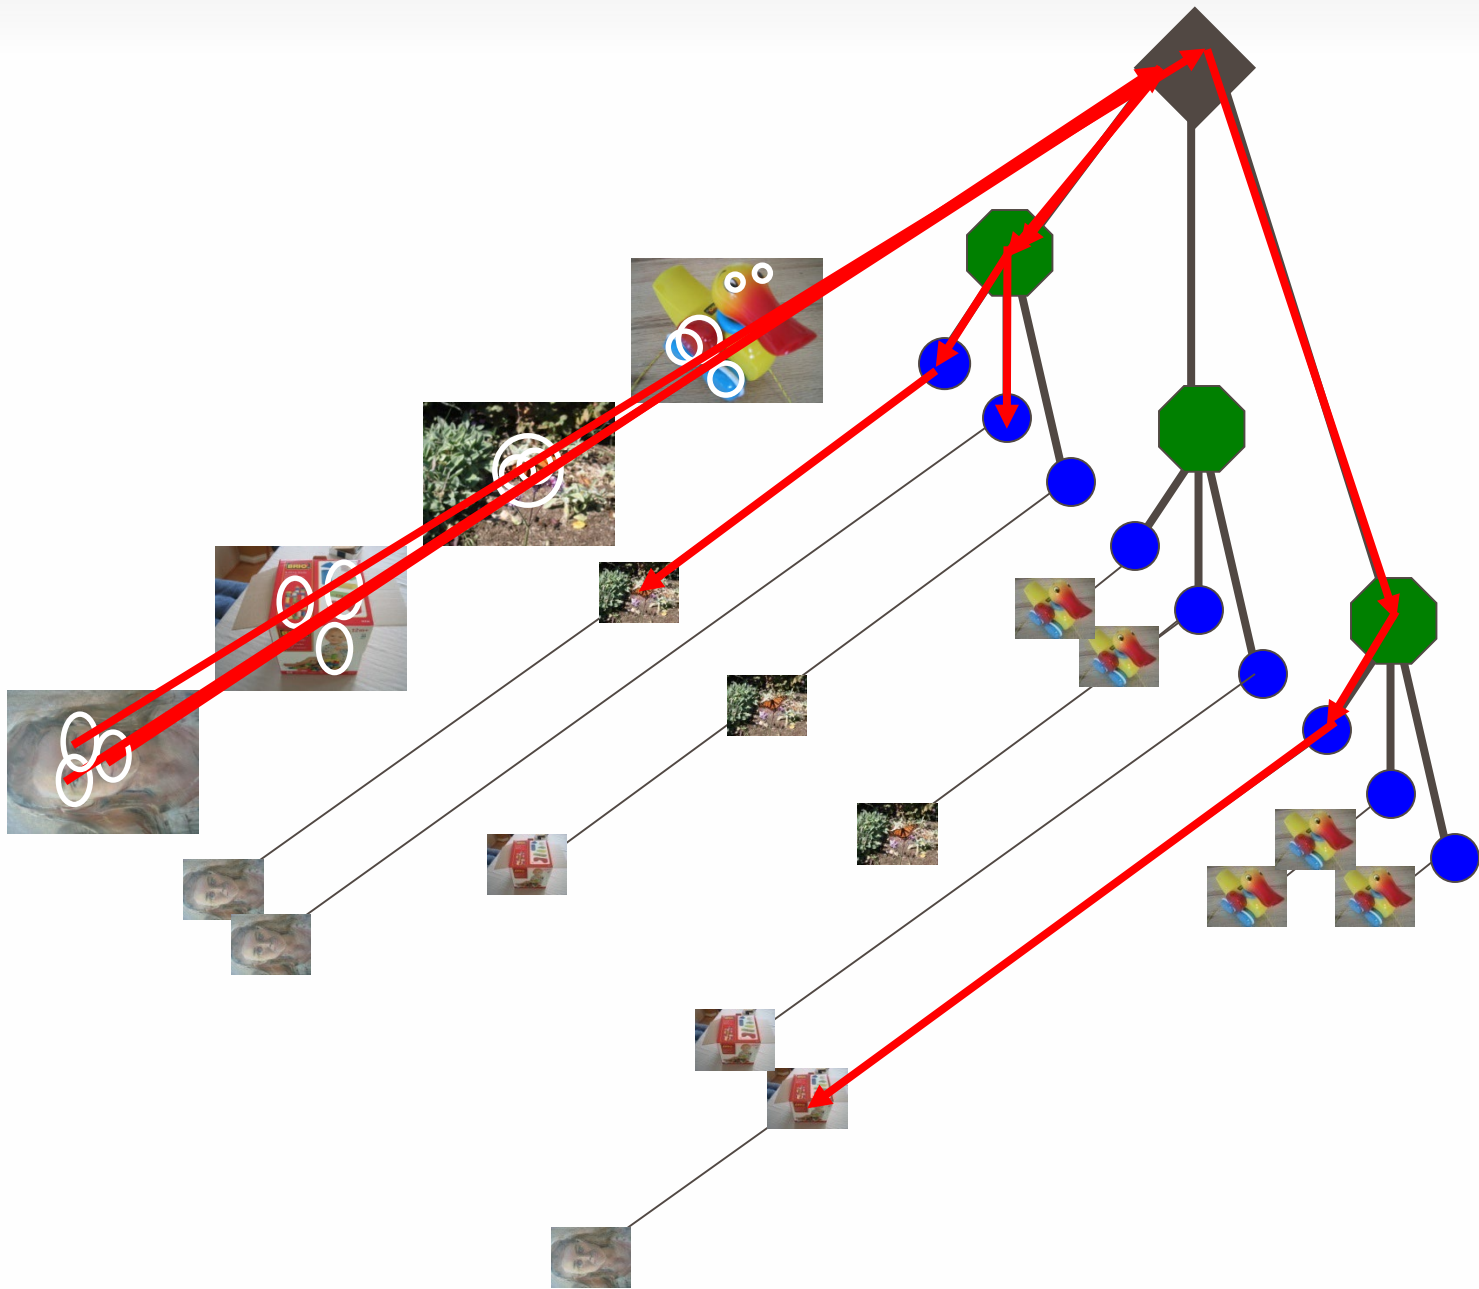


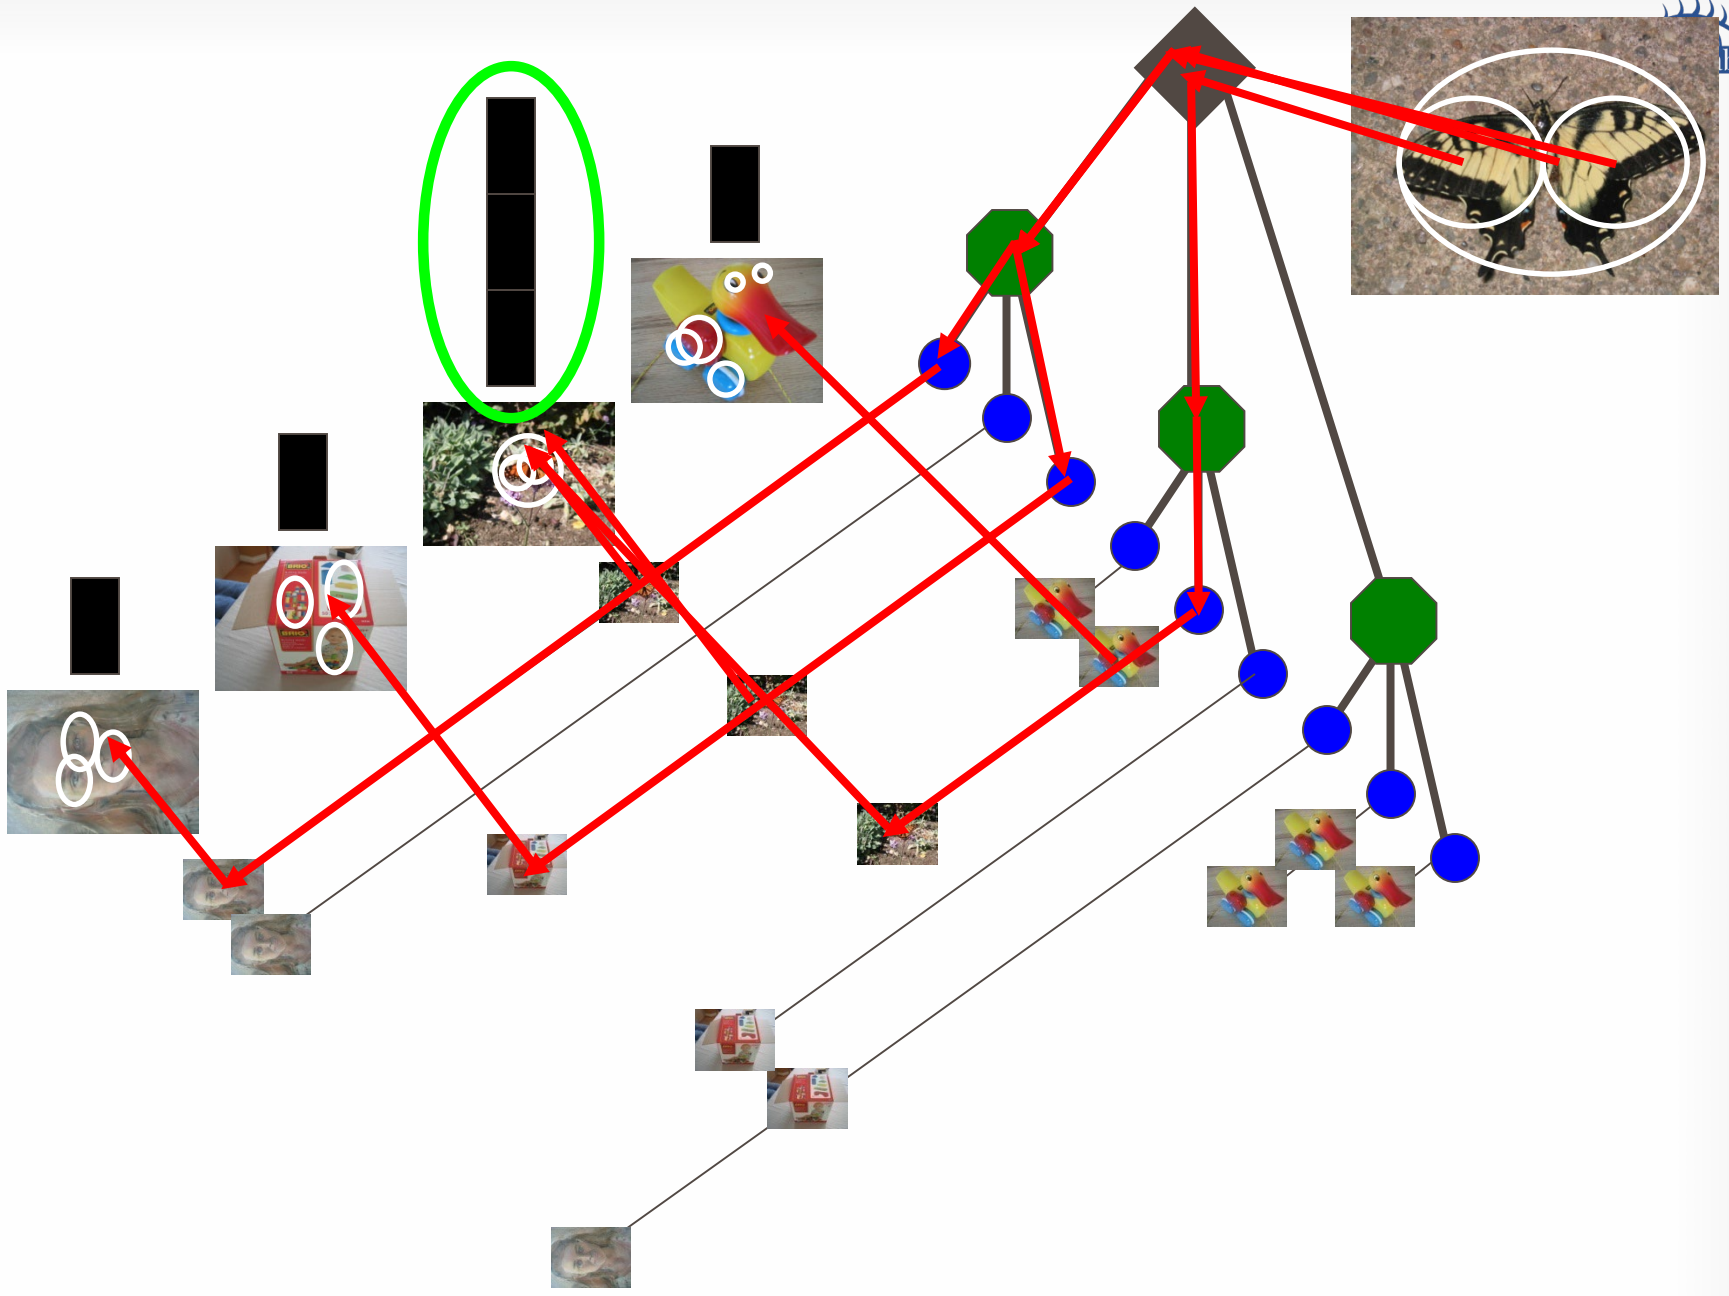












Vocabulary trees: complexity

Number of words given tree parameters: branching factor and number of levels

$$\text{branching_factor}^{\text{number_of_levels}}$$

Word assignment cost vs. flat vocabulary

$O(k)$ for flat


$$O(\log_{\text{branching_factor}}(k) * \text{branching_factor})$$

Is this like a kd-tree?

Yes, but with better partitioning and defeatist search.

This hierarchical data structure is lossy – you might not find your true nearest cluster.

Higher branch factor works better (but slower)



Visual words/bags of words

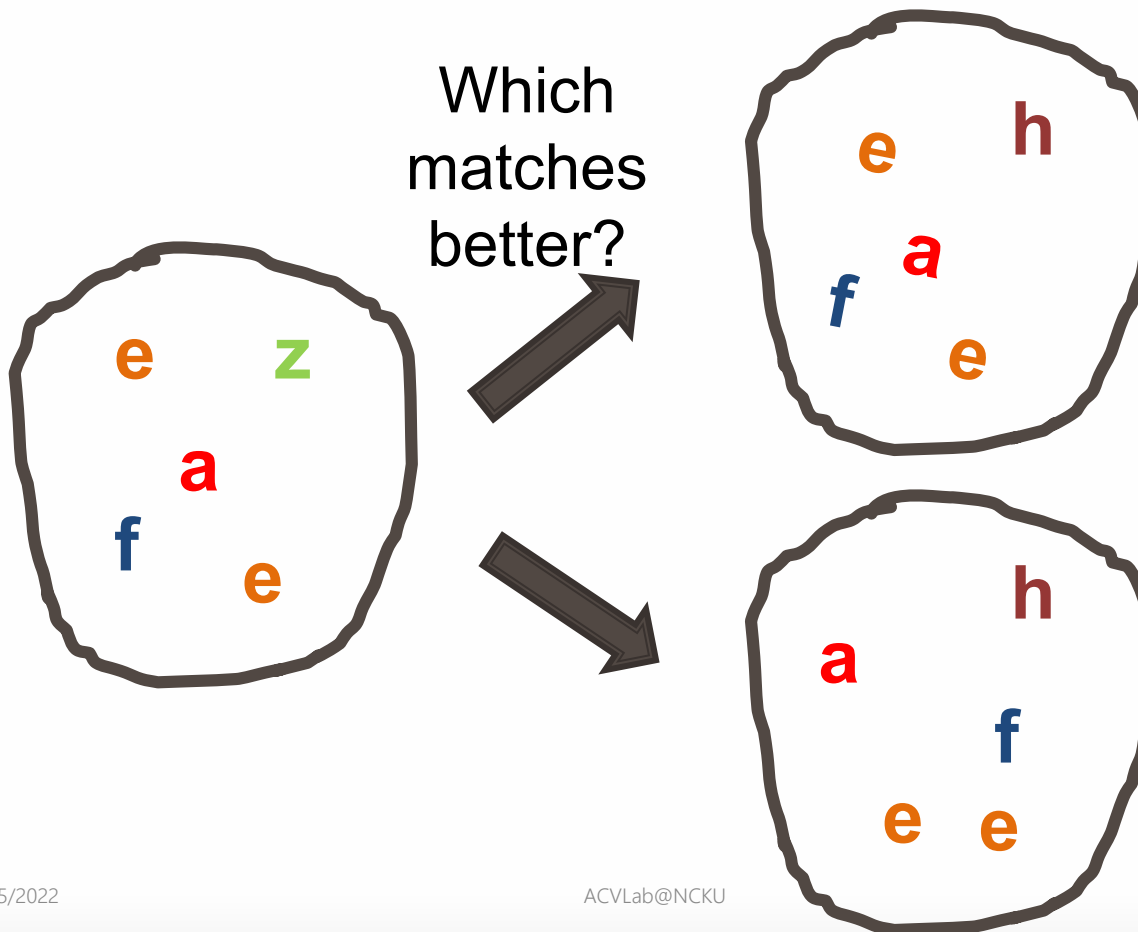
- + flexible to geometry / deformations / viewpoint
 - + compact summary of image content
 - + provides fixed dimensional vector representation for sets
 - + very good results in practice
-
- background and foreground mixed when bag covers whole image
 - optimal vocabulary formation remains unclear
 - basic model ignores geometry – must verify afterwards, or encode via features

Instance recognition: remaining issues

- How to summarize the content of an entire image? And gauge overall similarity?
- How large should the vocabulary be? How to perform quantization efficiently?
- Is having the same set of visual words enough to identify the object/scene? How to verify spatial agreement?
- How to score the retrieval results?



Can we be more accurate?

- So far, we treat each image as containing a “bag of words”, with no spatial information



Can we be more accurate?


- So far, we treat each image as containing a “bag of words”, with no spatial information



Real objects have consistent geometry



Spatial Verification

Query



DB image with high BoW
similarity

Query



DB image with high BoW
similarity

Both image pairs have many visual words in common.


Spatial Verification

Query



DB image with high BoW similarity

Query



DB image with high BoW similarity

Only some of the matches are mutually consistent

Spatial Verification: two basic strategies

■ RANSAC

- Typically sort by BoW similarity as initial filter
- Verify by checking support (inliers) for possible transformations
 - e.g., “success” if find a transformation with $> N$ inlier correspondences

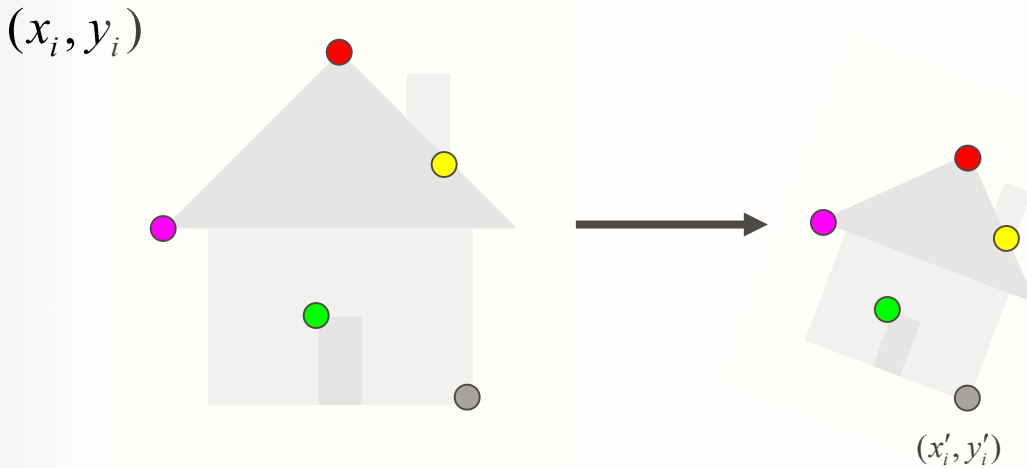
■ Generalized Hough Transform

- Let each matched feature cast a vote on location, scale, orientation of the model object
- Verify parameters with enough votes

RANSAC verification



Recall: Fitting an affine transformation






Approximates viewpoint changes for roughly planar objects and roughly orthographic cameras.

$$\begin{bmatrix} x'_i \\ y'_i \end{bmatrix} = \begin{bmatrix} m_1 & m_2 \\ m_3 & m_4 \end{bmatrix} \begin{bmatrix} x_i \\ y_i \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \end{bmatrix}$$

$$\begin{bmatrix} \dots & \dots & \dots & \dots & \dots & \dots \\ x_i & y_i & 0 & 0 & 1 & 0 \\ 0 & 0 & x_i & y_i & 0 & 1 \\ \dots & \dots & \dots & \dots & \dots & \dots \end{bmatrix} \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} \dots \\ x'_i \\ y'_i \\ \dots \end{bmatrix}$$

RANSAC verification



Instance recognition: remaining issues

- How to summarize the content of an entire image? And gauge overall similarity?
- How large should the vocabulary be? How to perform quantization efficiently?
- Is having the same set of visual words enough to identify the object/scene? How to verify spatial agreement?
- How to score the retrieval results?

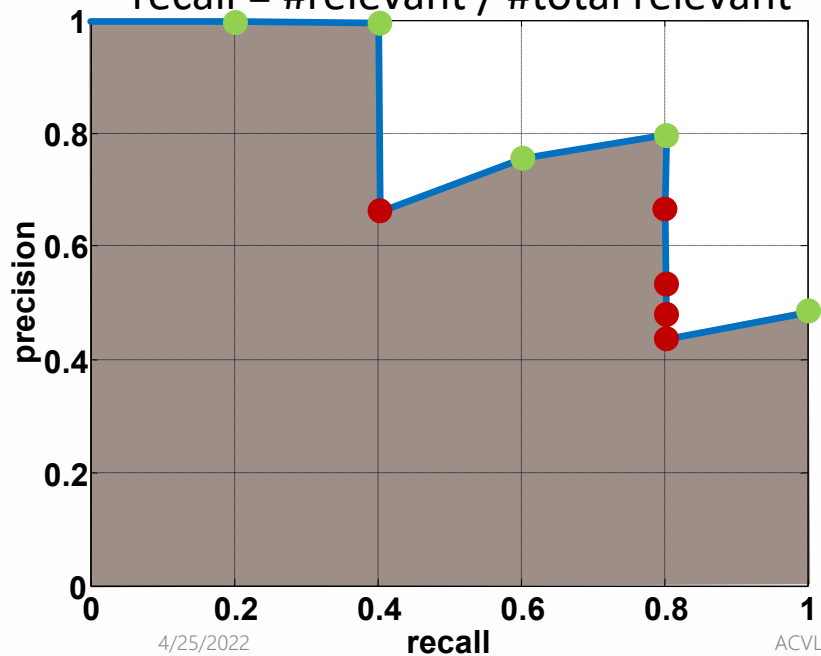
Scoring retrieval quality





Database size: 10 images
Relevant (total): 5 images

Query

precision = #relevant / #returned
recall = #relevant / #total relevant



Results (ordered):



Query Expansion

Results



Query image




Spatial verification



New results



New query



Chum, Philbin, Sivic, Isard, Zisserman: Total Recall..., ICCV 2007

Slide credit: Ondrej Chum

Recognition via alignment

- Pros:
 - Effective when we are able to find reliable features within clutter
 - Great results for matching specific instances
- Cons:
 - Scaling with number of models
 - Spatial verification as post-processing – not seamless, expensive for large-scale problems
 - Not suited for category recognition.

Summary

- Matching local invariant features
 - Useful not only to provide matches for multi-view geometry, but also to find objects and scenes.
- Bag of words representation: quantize feature space to make discrete set of visual words
 - Summarize image by distribution of words
 - Index individual words
- Inverted index: pre-compute index to enable faster search at query time
- Recognition of instances via alignment: matching local features followed by spatial verification
 - Robust fitting : RANSAC, GHT

INTRODUCTION TO IMAGE PROCESSING

Chih-Chung Hsu (許志仲)
Assistant Professor
ACVLab, Institute of Data Science
National Cheng Kung University
<https://cchsu.info>



Applications- Why/ Who

- Useful in a variety of fields from science to design
- Adjust images and run analysis
 - Add filters and text
 - Make images easier to view
 - Count cells or items in an image
 - Track movement
 - Facial Recognition
 - Create movies/gifs

Basics- What is an Image?

- Image Type
- Vector
 - Made of independent and editable lines/shapes
- Pixel
 - Also called bitmap or raster
 - Made of uniform grid of colored dots (pixels)
- Compression
 - JPEG → Common
 - Loss vs Lossless

Format

Extension	Colour	Compression	Common Uses
JPG, JPEG	24-bit	Lossy	Photos, web pics
GIF	8-bit	Lossless	Web graphics – buttons, icons, etc
PNG	up to 24-bit	Lossless	Web – replacement for GIF
TIF, TIFF	24-bit	Lossless	Professional Photos etc

<https://www.slideshare.net/bobwatson/image-file-formats>

Opening an image in Python

- Need Numpy and Scipy:
- Numpy: basic array manipulation
 - Pixel images are stored as arrays
 - Each pixel has (x,y,rgb)
 - Means you can loop through images and has array functions
- Scipy: dedicated to image processing

Opening an image in Python

Here are the imports for this first section:

```
import os
import numpy as np
from scipy import ndimage, misc
import matplotlib.pyplot as plt
from PIL import Image, ImageEnhance
```

This is how I import an image using numpy and PIL.
Make sure you have navigated to the correct file folder.

```
28 ##Open and view an image #####
29 im = misc.imread('cricket.jpg')
30 type(im) #should show as np.array
31 im.shape, im.dtype #shows current shape and type
32 plt.imshow(im) #view image
33
34 imPIL = Image.open('cricket.jpg')
35 imPIL.show()
36 np.array(imPIL).shape
--
```

Image Properties

Resolution

- Size of each pixel expressed as number of pixels per unit
 - Dots per inch (DPI)
 - Pixels per inch (PPI)
- Dimensions
- Number of pixels along X and Y axes
- Can be:
 - Pixels (800 X 600 pixels)
 - Physical (89 mm X 66 mm)

Color

- Models:
 - RGB
 - CMYK
- Format:
 - RGB Values
 - Hex Values

Code

- The following sections reference python code
 - The title is a description of the section
 - Information about the images needed is in notes
 - Most of it works with your image of choice

Basics of Numpy/Scipy and PIL

```

41 ##crop dog image
42 im = im[100: 500, 0: 360] #YAxis, XAxis
43 plt.imshow(im) #view image
44
45 imPIL = imPIL.crop((0, 100, 360, 500)) #ULX, ULY, LRX, LRY
46 imPIL.show()
47
48 ##Rotate and Flip dog
49 flipped_im = np.flipud(im)
50 plt.imshow(flipped_im)
51
52 rotate_im = ndimage.rotate(im,30)
53 plt.imshow(rotate_im)
54
55 rotate_imPIL = imPIL.rotate(30)
56 rotate_imPIL.show()
57
58 ##Change to GreyScale
59 from skimage.color import rgb2gray
60 grayIm = rgb2gray(im)
61 plt.imshow(grayIm, cmap = plt.get_cmap('gray'))
62 plt.imshow(grayIm) #EXAMPLE: Does not work
63
64 grayImPIL = imPIL.convert('LA') #using PIL
65 grayImPIL.show()
66

```

```

67 ##Change Contrast of Image
68 from skimage import util
69 inverted_im = util.invert(im)
70 plt.imshow(inverted_im)
71
72 from skimage import exposure
73 vmin, vmax = np.percentile(im, (0.2,99.8))
74 vmin, vmax
75 contrast_im = exposure.rescale_intensity(im, in_range=(vmin,vmax))
76 plt.imshow(contrast_im) #FIX: Not working right now
77
78 imPIL2 = ImageEnhance.Contrast(imPIL)
79 imPIL2.enhance(2).show()

```


Converting between np.array and PIL

```
81 #Convert between numpy array and PIL  
82 Array2PIL = Image.fromarray(im) #converts to PIL from array  
83 Array2PIL.show()  
84  
85 PIL2Array = np.array(imPIL) #change PIL to a numpy array  
86 plt.imshow(PIL2Array)
```

Manually Working with np.array

```


3 from PIL import Image
4 from numpy import *
5 import matplotlib.pyplot as plt
6 import sys
7
8 #read in an image
9 pil_im = Image.open('picToFix.jpg') #change to sys.argv[1] if calling from cmd
10 pil_im.show()
11 imArray = array(pil_im)
12
13 #find shape of array and get number of rows and columns
14 arrayShape = imArray.shape
15 rows = arrayShape[0]
16 rows
17 cols = arrayShape[1]
18 cols
19
20 #create a negative image
21 negIm = []
22 negIm = 255-imArray
23 plt.imshow(negIm)
24
25 #create a greyscale image
26 greyIm = 0 + imArray
27 for row in range(rows):
28     for col in range(cols):
29         grey = (float(imArray[row,col,0]) + float(imArray[row,col,1]) + float(imArray[row,col,2]))/3
30         for color in range(3):
31             grey = int(grey)
32             greyIm[row][col][color] = grey
33 plt.imshow(greyIm)

```

Using Numpy allows you to iterate through pixels and make very specific changes. Most things are best done using PIL or a pre-made function, but all things can be done “manually”.

Manually Working with np.array

- You can get very (overly?) specific like in the following case where I have removed the green check from the image. I used PixelMath to find the exact range of green in the image.



```

35 #create an image without the green check
36 no_g = 0 + imArray
37 for row in range(rows):
38     for col in range(cols):
39         if imArray[row,col,0] < 240 and imArray[row,col,1] > 200 and imArray[row,col,1] < 250:
40             for color in range(3):
41                 no_g[row][col][color] = 255

```

Creating a Gif or Movie from a Series of Images

- First, I need to navigate to an empty folder and fill it with numbered images that I want to turn into a movie/gif.

```
89 ## Create a gif from images in file #####
90 import glob
91 import moviepy.editor as mpy
92
93 os.chdir(r'C:\Users\Stefanie\Documents\WSU\CougPy\gif')
94
95 angle = 0
96 while(angle<360): #creates images that rotate in a circle
97     rotate_imPIL = imPIL.rotate(angle)
98     #you will need to change your filepath as needed
99     fileNameTemplate = r'C:\Users\Stefanie\Documents\WSU\CougPy\gif\dog_{0:02d}.png'
100     rotate_imPIL.save(fileNameTemplate.format(angle), format='png')
101     angle = angle + 10
102
103 gif_name = 'godog'
104 fps = 50 #set the frame rate
105 file_list = glob.glob('*.png') # Get all the pngs in the current directory
106 list.sort(file_list, key=lambda x: int(x.split('_')[1].split('.png')[0])) # Sort the images by #,
107 clip = mpy.ImageSequenceClip(file_list, fps=fps) #creates the image sequence
108 clip.write_gif('{}.gif'.format(gif_name), fps=fps) #this will create a gif
109 #clip.write_videofile('{}.mp4'.format(gif_name), fps=fps) #this will create a video
```



Denoising and Feature Extraction

- This code is from the SciKit tutorial on plot boundaries listed in references

```
114 ##Removing noise with filters
115 from skimage.morphology import disk
116 from skimage import data
117 from skimage import filters
118 from skimage import restoration
119 coins = data.coins()
120 coins_zoom = coins[10:80, 300:370]
121 plt.imshow(coins_zoom, cmap='gray', interpolation='nearest')
122 median_coins = filters.median(coins_zoom, disk(1))
123 plt.imshow(median_coins, cmap='gray',
124            interpolation='nearest')
125 tv_coins = restoration.denoise_tv_chambolle(coins_zoom, weight=0.1)
126 plt.imshow(tv_coins, cmap='gray',
127            interpolation='nearest')
128 gaussian_coins = filters.gaussian(coins_zoom, sigma=2)
129 plt.imshow(gaussian_coins, cmap='gray',
130            interpolation='nearest')
```

Denoising and Feature Extraction

- This code is from the SciPy tutorial listed in references

```
132 ##Fixing Mathematical Morphology
133 square = np.zeros((32, 32))
134 square[10:-10, 10:-10] = 1
135 np.random.seed(2)
136 x, y = (32*np.random.random((2, 20))).astype(np.int)
137 square[x, y] = 1
138 plt.imshow(square, cmap='gray')
139
140 open_square = ndimage.binary_opening(square)
141 plt.imshow(open_square, cmap='gray')
142
143 eroded_square = ndimage.binary_erosion(square)
144 plt.imshow(eroded_square, cmap='gray')
145
146 reconstruction = ndimage.binary_propagation(eroded_square, mask=square)
147 plt.imshow(reconstruction, cmap='gray')
148
```

Denoising and Feature Extraction

```

151 ##Feature Extraction: Counting "Red Blood Cells" ##### 177 ##Remove anything that is too small for us to care about
152 ##Create synthetic data
153 n = 10
154 l = 256
155 im = np.zeros((l, l))
156 points = l*np.random.random((2, n**2))
157 im[(points[0]).astype(np.int), (points[1]).astype(np.int)] = 1
158 im = ndimage.gaussian_filter(im, sigma=l/(4.*n))
159 plt.imshow(im)
160
161 mask = im > im.mean() #threshold the image into binary colors by mean
162 plt.imshow(mask, cmap='gray')
163
164 label_im, nb_labels = ndimage.label(mask) #count the 'red blood cells'
165 nb_labels #will give int count
166 plt.imshow(label_im) #colors our individual counts
167 #clearly some are touching and it is counting that as one large cell
168 #can be adjusted with threshold
169
170 ##compute the size and mean volume of each region
171 sizes = ndimage.sum(mask, label_im, range(nb_labels + 1))
172 sizes
173 mean_vals = ndimage.sum(im, label_im, range(1, nb_labels + 1))
174 mean_vals

```

This code is from the SciPy tutorial listed in references

More Techniques

- Plenty of other things possible:
- Segmentation:
 - Mark edges of features in image
- More Feature Extraction
 - Computer vision can detect features (ie corners)
- Facial Recognition
 - Using OpenCV or others

Resources

- Practical Computing for Biologists by Haddock & Dunn
- <https://realpython.com/blog/python/face-recognition-with-python/>
- <https://pillow.readthedocs.io/en/3.1.x/reference/Image.html>
- http://www.scipy-lectures.org/advanced/image_processing/#geometrical-transformations
- http://scikit-image.org/docs/dev/user_guide/transforming_image_data.html
- http://www.scipy-lectures.org/packages/scikit-image/auto_examples/plot_boundaries.html
- <https://www.safaribooksonline.com/library/view/programming-computer-vision/9781449341916/ch01.html>