Photogrammetry & Robotics Lab

Bag of Visual Words for Finding Similar Images

Cyrill Stachniss

Slides have been created by Cyrill Stachniss.
Most images by Olga Vysotska and Fei-Fei Li.

Preparation: Watch 5 Min Video

https://www.youtube.com/watch?v=a4cFONdc6nc

What is Bag of Visual Word for?

- Finding images in a database, which are similar to a given query image
- Computing image similarities
- Compact representation of images

?

Analogy to Text Documents

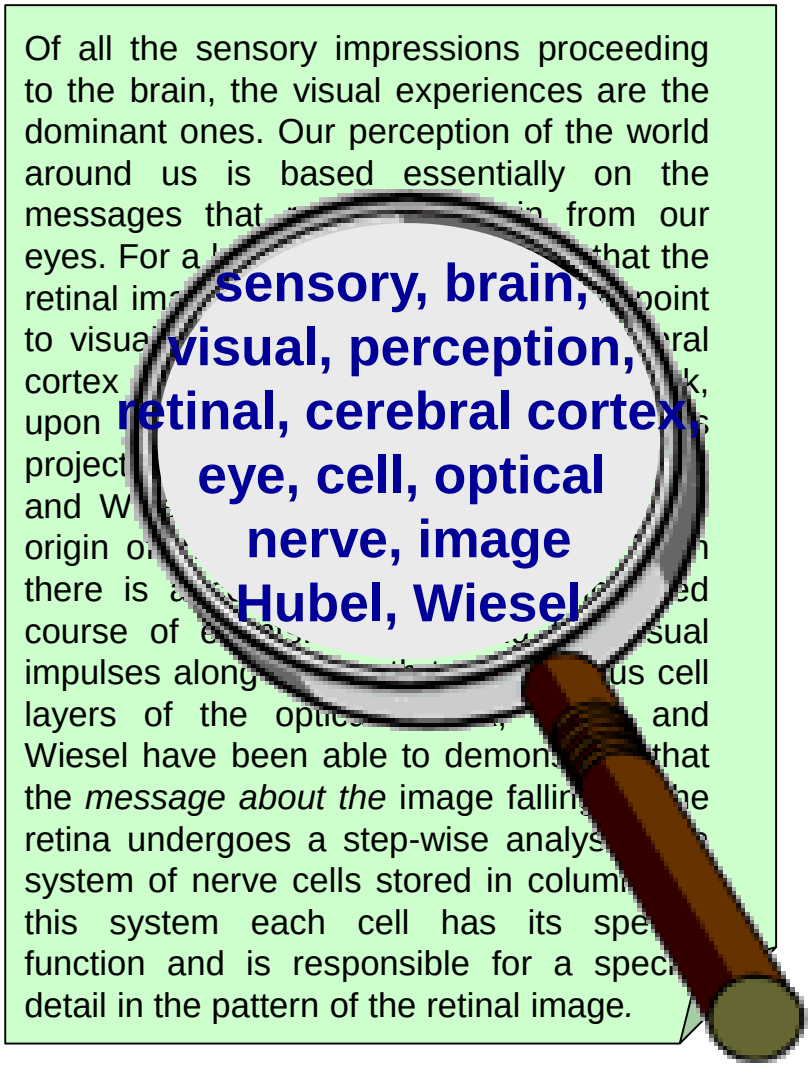

[Image courtesy: Fei-Fei Li]

Looking for Similar Papers

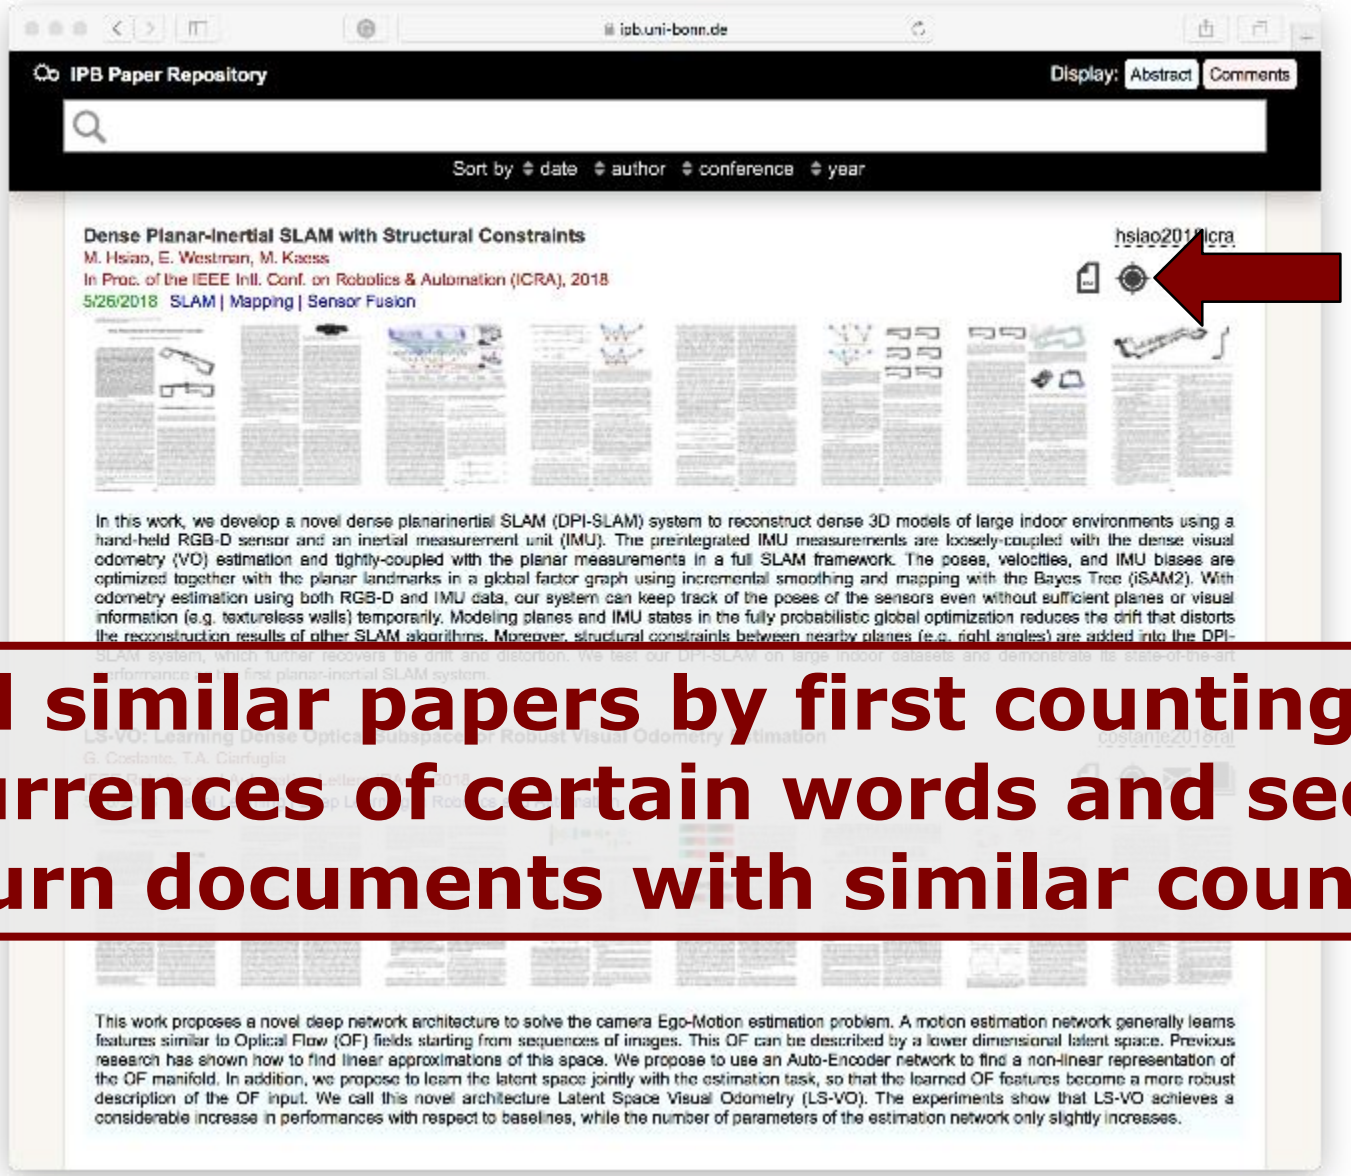

"find similar papers by first counting the occurrences of certain words and second return documents with similar counts."

Bag of (Visual) Words

Analogy to documents: The content of a can be inferred from the frequency of relevant words that occur in a document

object

bag of “visual words”

[Image courtesy: Fei-Fei Li]

Bag of Visual Words

- Visual words = independent features

face

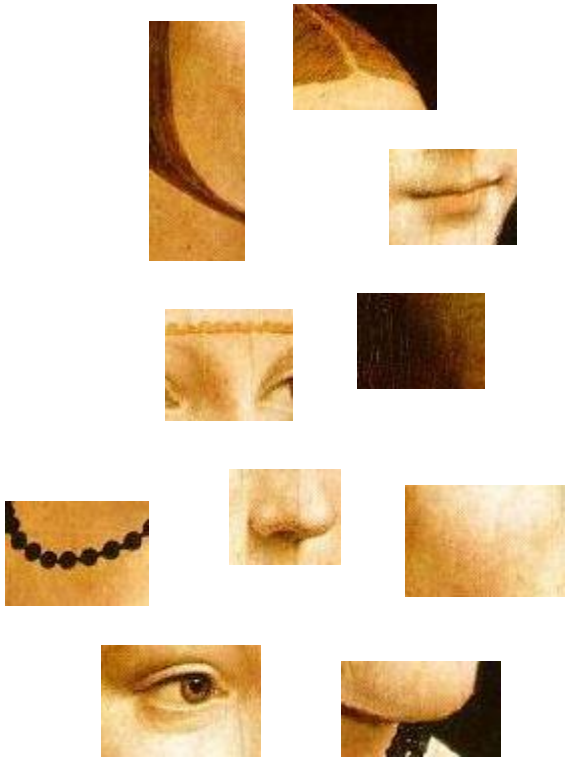

features

[Image courtesy: Fei-Fei Li]

Bag of Visual Words

- Visual words = independent features
- Construct a dictionary of representative words
- Use only words from the dictionary

dictionary ("codebook")

[Image courtesy: Fei-Fei Li]

Bag of Visual Words

- Visual words = independent features
- Words from the dictionary
- Represent the images based on a histogram of word occurrences

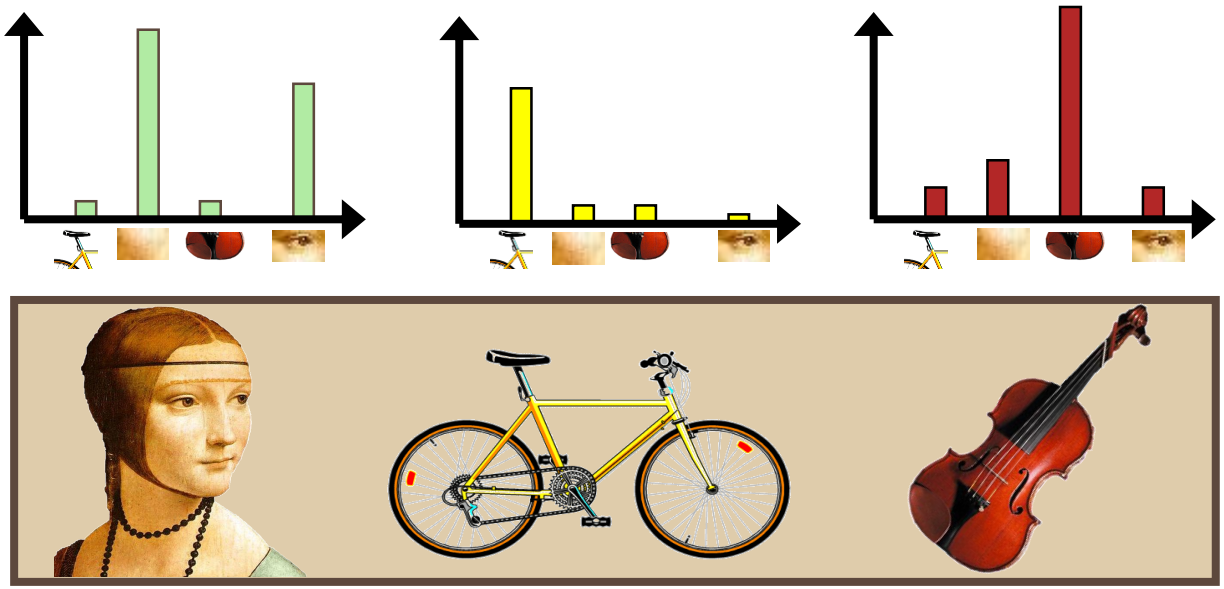

[Image courtesy: Fei-Fei Li]

Bag of Visual Words

- Visual words = independent features
- Words from the dictionary
- Represent the images based on a histogram of word occurrences
- Image comparisons are performed based on such word histograms

[Image courtesy: Fei-Fei Li]

From Images to Histograms

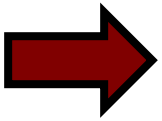

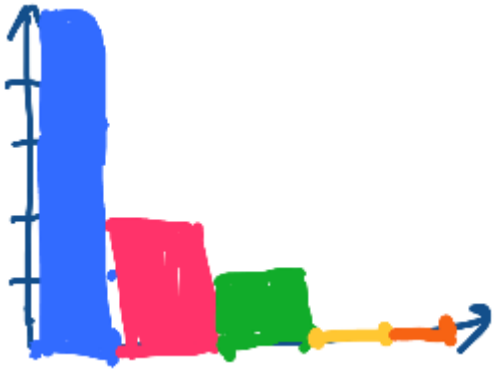

[Image courtesy: Olga Vysotska]

Overview: Input Image

Overview: Extract Features

[Image courtesy: Olga Vysotska]

Overview: Visual Words

[Image courtesy: Olga Vysotska]

Overview: No Pixel Values

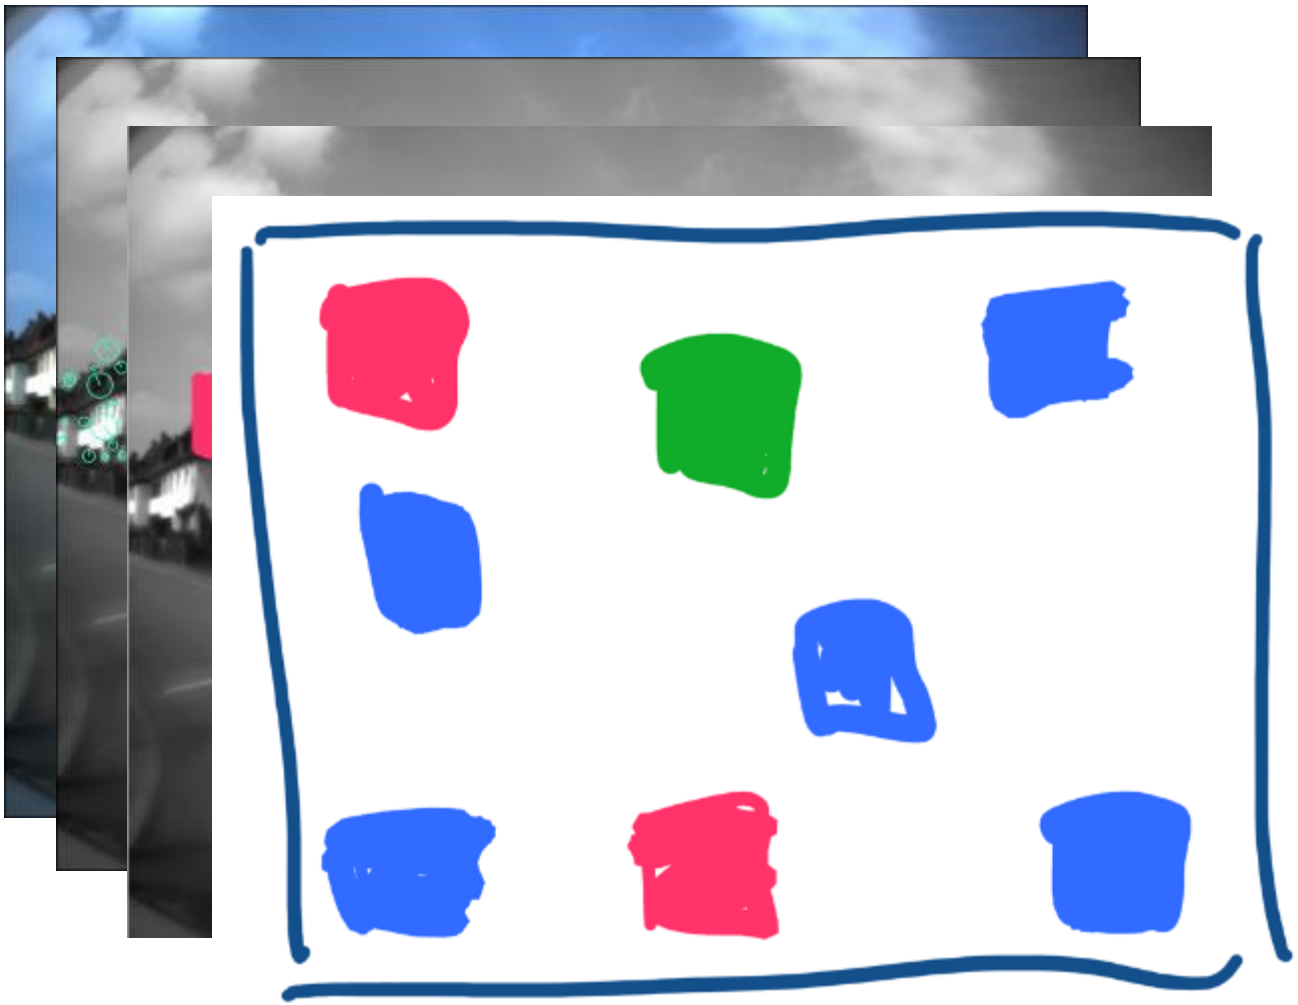

[Image courtesy: Olga Vysotska]

Overview: Word Occurrences

[Image courtesy: Olga Vysotska]

Images to Histograms

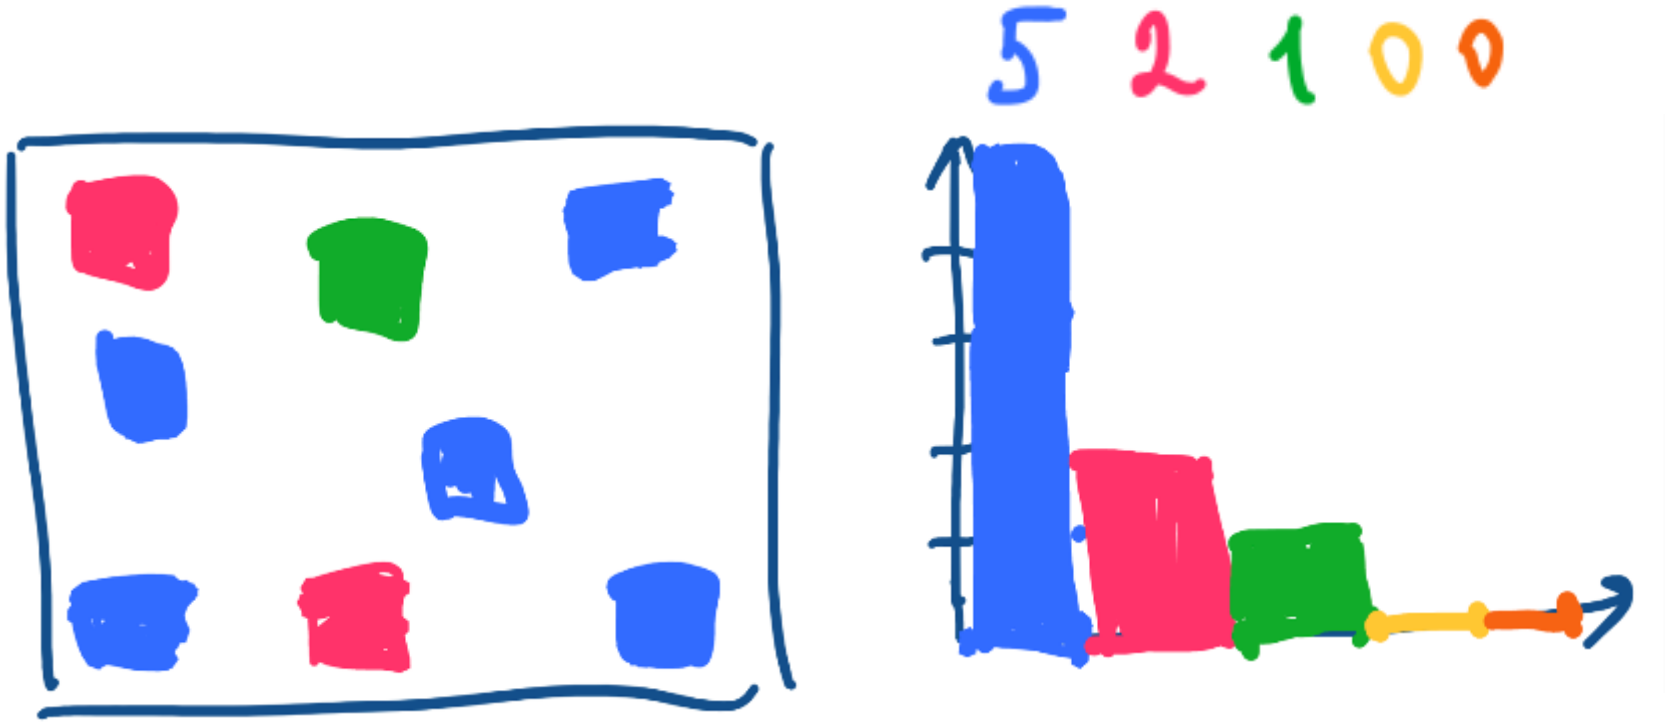

[Image courtesy: Olga Vysotska]

Where Do the Visual Words Come Form?

Dictionary

- A dictionary defines the list of words that are considered
- The dictionary defines the x-axes of all the word occurrence histograms

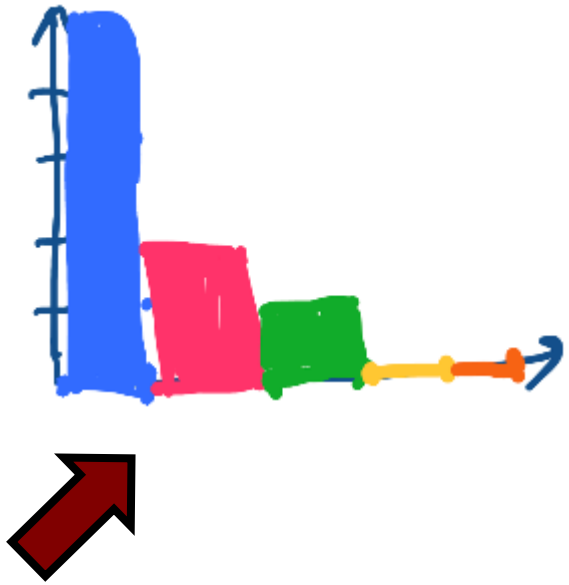

[Image courtesy: Olga Vysotska]

Dictionary

- A dictionary defines the list of words that are considered
- The dictionary defines the x-axes of all the word occurrence histograms
- The dictionary must remain fixed

The dictionary is typically learned from data. How can we do that?

Extract Feature Descriptors from a Training Dataset

Visual feature descriptor vectors (e.g., SIFT)

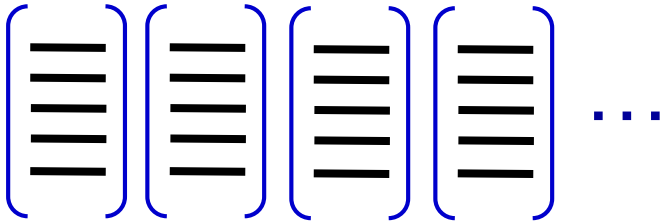

[Partial image courtesy: Fei-Fei Li]

Feature Descriptors are Points in a High-Dimensional Space

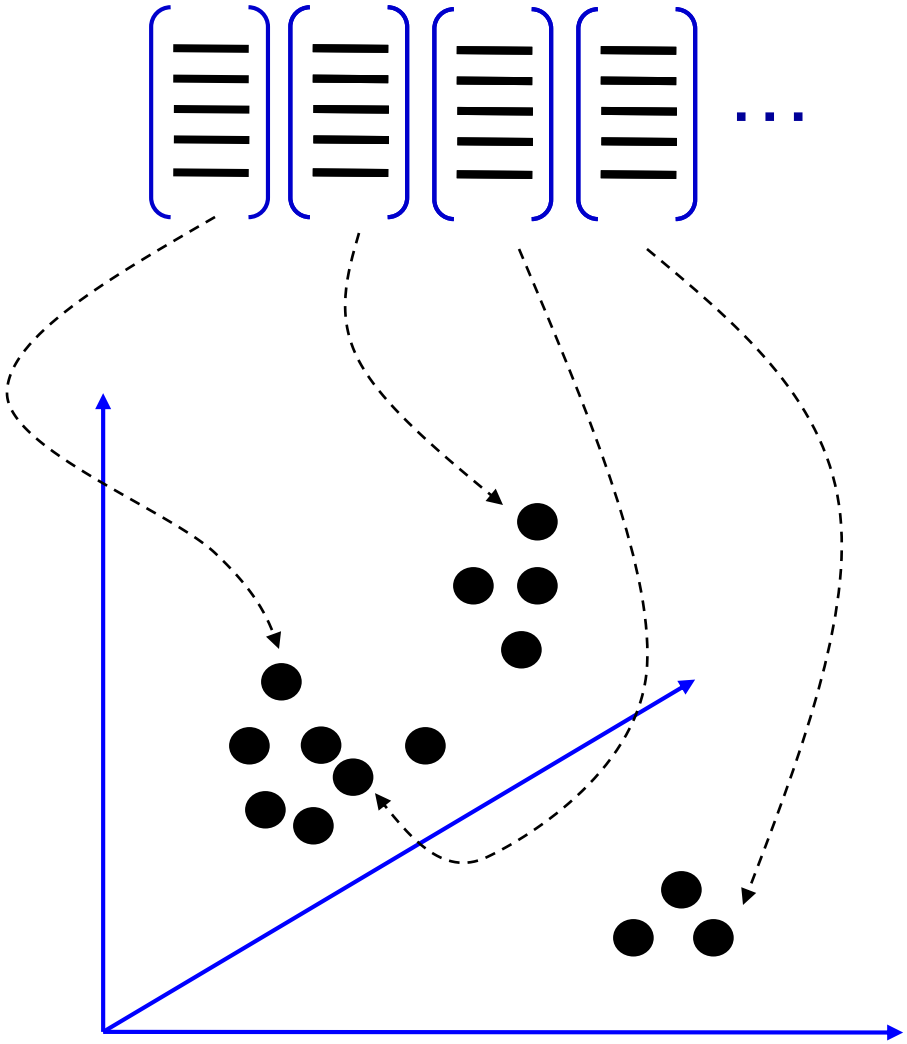

[Image courtesy: Fei-Fei Li]

Group Similar Descriptors

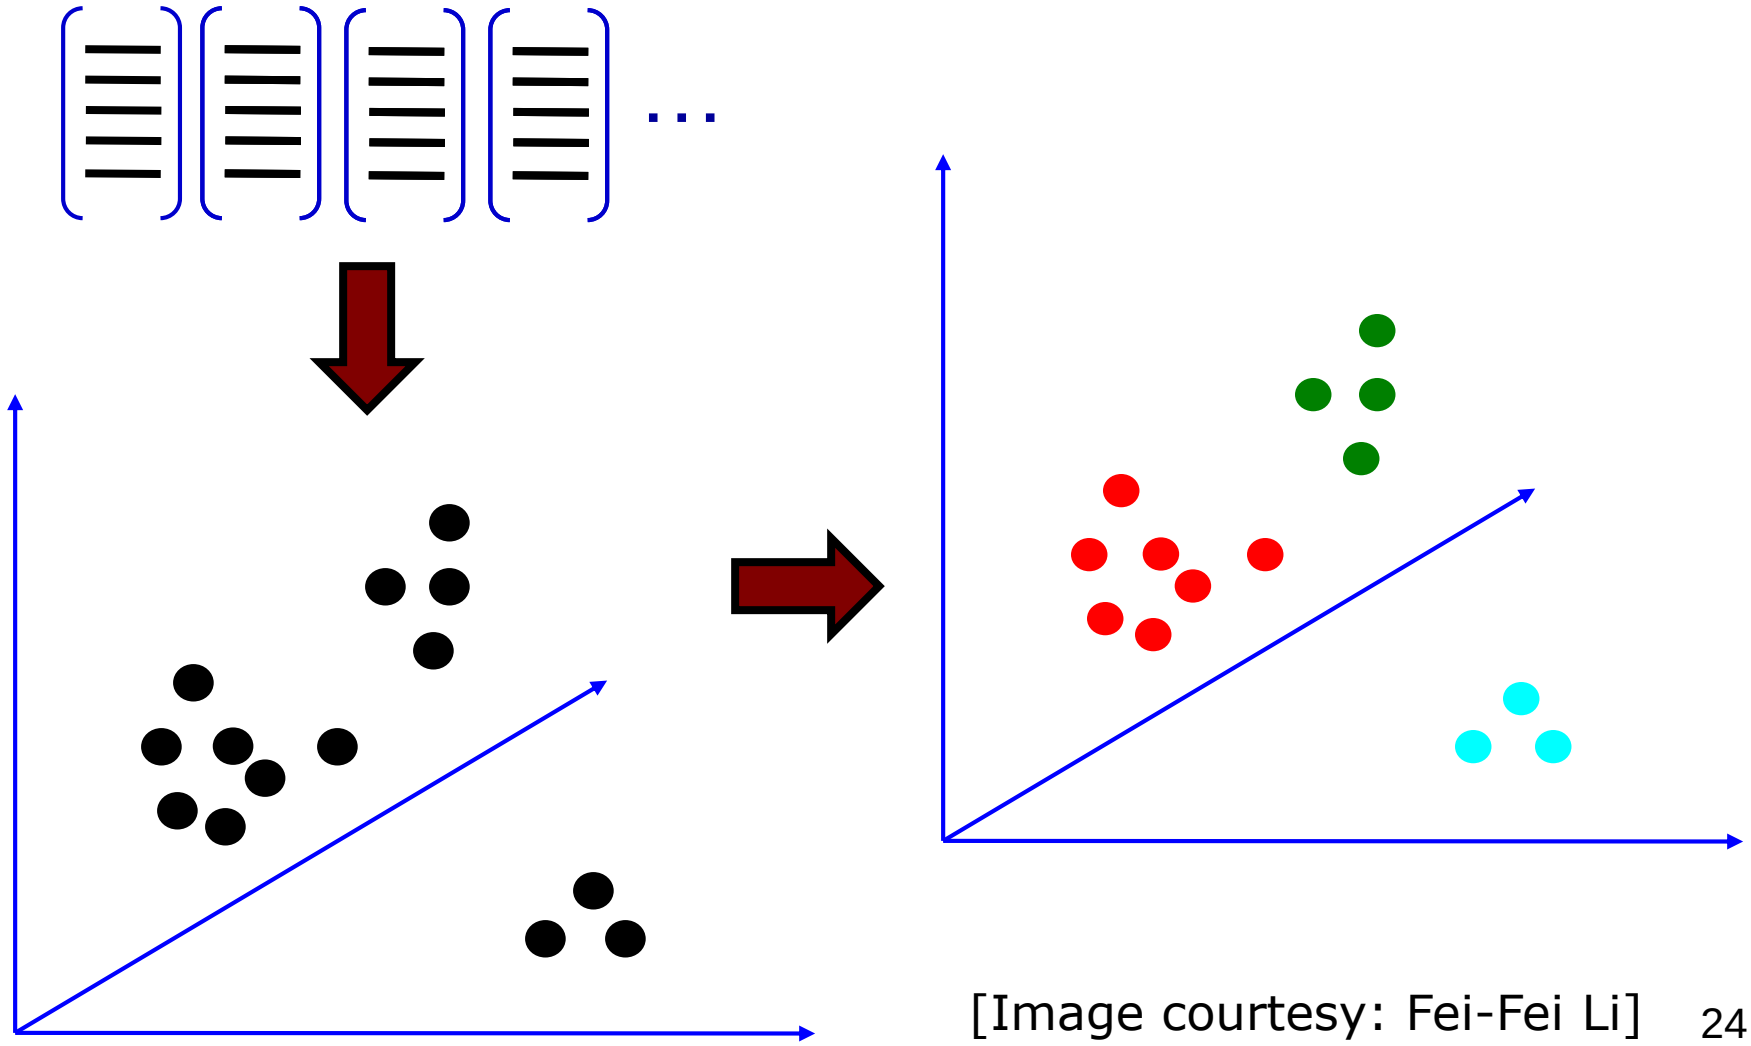

[Image courtesy: Fei-Fei Li]

Clusters of Descriptors from Data Forms the Dictionary

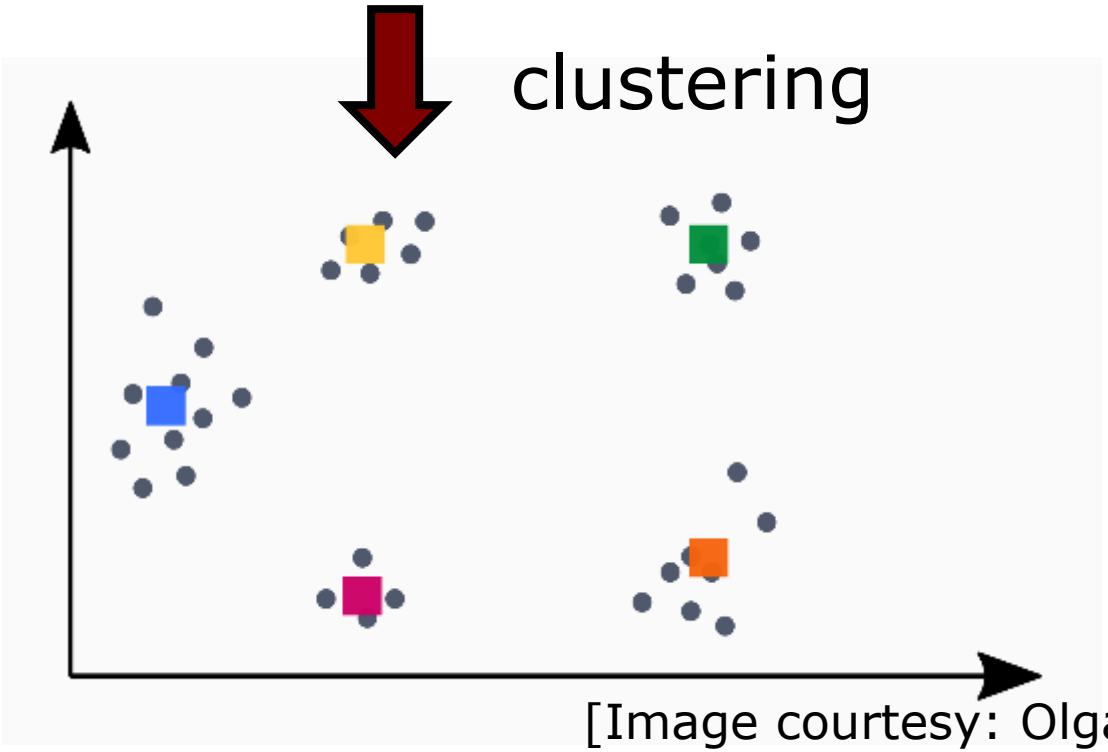

[Image courtesy: Olga Vysotska]

K-Means Clustering

K-Means Clustering

- Partitions the data into k clusters
- Clusters are represented by centroids
- A centroid is the mean of data points

Objective:

- Find the k cluster centers and assign the data points to the nearest one, such that the squared distances to the cluster centroids are minimized

K-Means Clustering for Learning the BoVW Dictionary

- Partitions the features into k groups
- The centroids form the dictionary
- Features will be assigned to the closest centroid (visual word)

Approach:

- Find k word and assign the features to the nearest word, such that the squared distances are minimized

K-Means Clustering (Informally)

- Initialization: Choose k arbitrary centroids as cluster representatives

- Repeat until convergence
  - Assign each data point to the closest centroid
  - Re-compute the centroids of the clusters based on the assigned data points

K-Means Algorithm

Initialize $\boldsymbol{m}_i, i = 1, \ldots, k$, for example, to k random $\boldsymbol{x}^t$

Repeat

For all $\boldsymbol{x}^t \in \mathcal{X}$

$$b_i^t \leftarrow \begin{cases} 1 & \text{if } \|\boldsymbol{x}^t - \boldsymbol{m}_i\| = \min_j \|\boldsymbol{x}^t - \boldsymbol{m}_j\| \\ 0 & \text{otherwise} \end{cases}$$

For all $\boldsymbol{m}_i, i = 1, \ldots, k$

$$\boldsymbol{m}_i \leftarrow \sum_t b_i^t \boldsymbol{x}^t / \sum_t b_i^t$$

Until $\boldsymbol{m}_i$ converge

Re-compute the cluster means using the current cluster memberships

Assign each data point to the closest cluster

K-Means Example

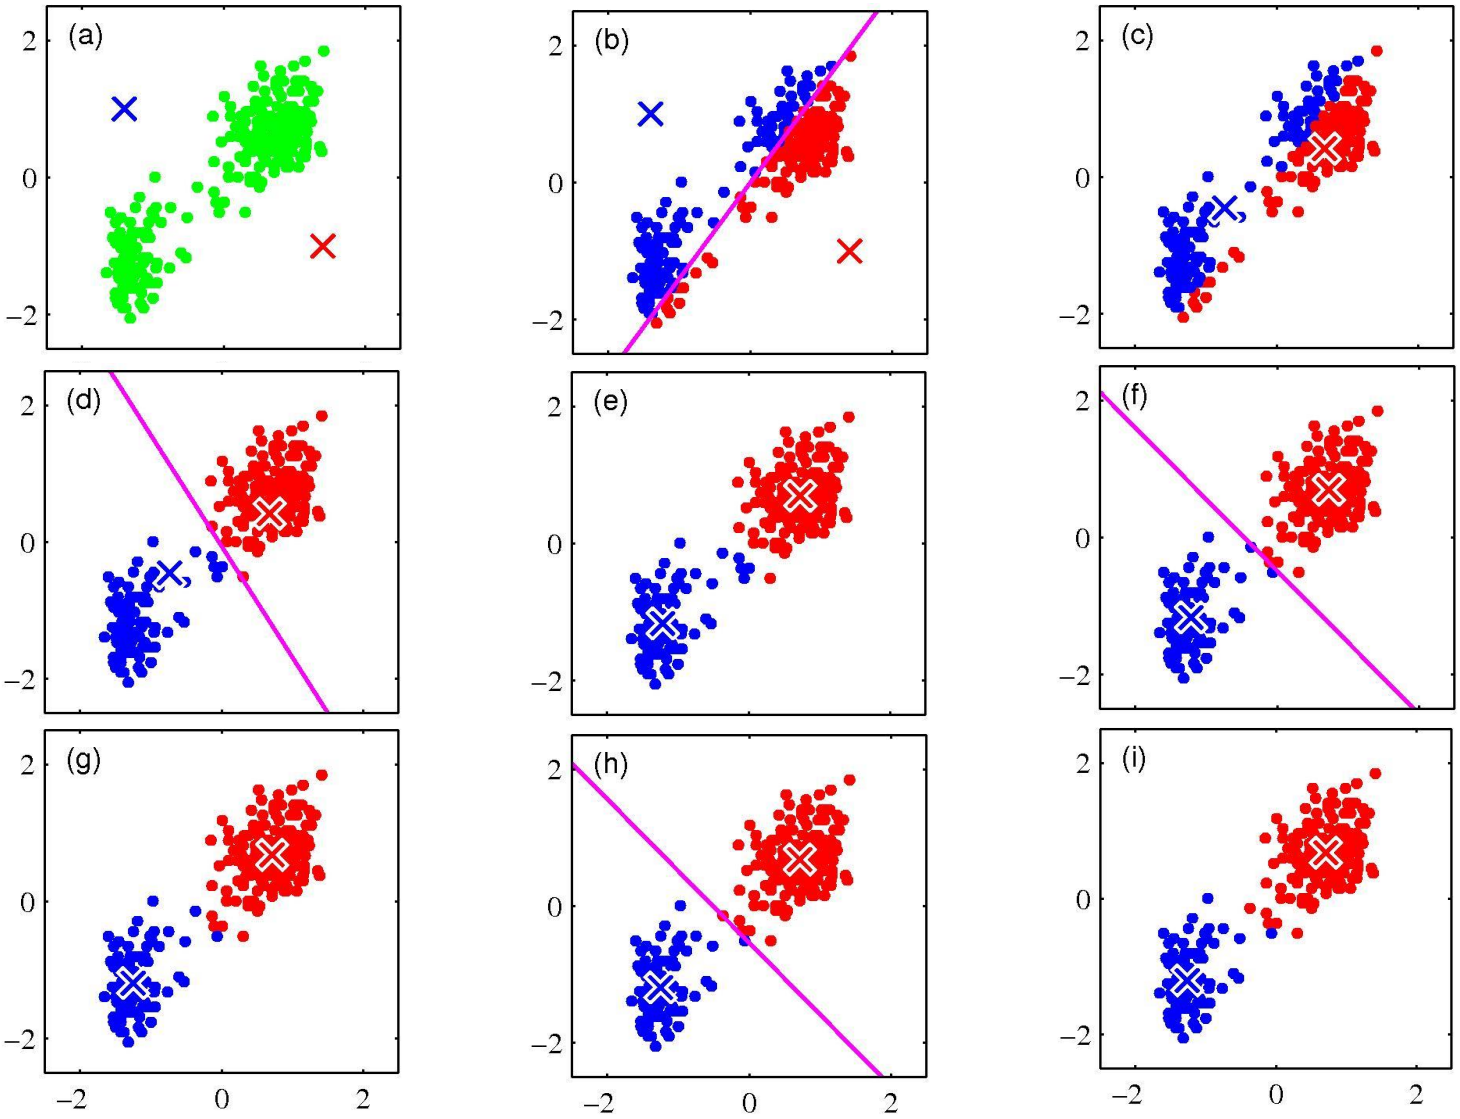

[Image courtesy: Bishop]

Summary K-Means

- Standard approach to clustering
- Simple to implement
- Number of clusters k must be chosen
- Depends on the initialization
- Sensitive to outliers
- Prone to local minima

We use k-means to compute the dictionary of visual words

K-Means for Building the Dictionary from Training Data

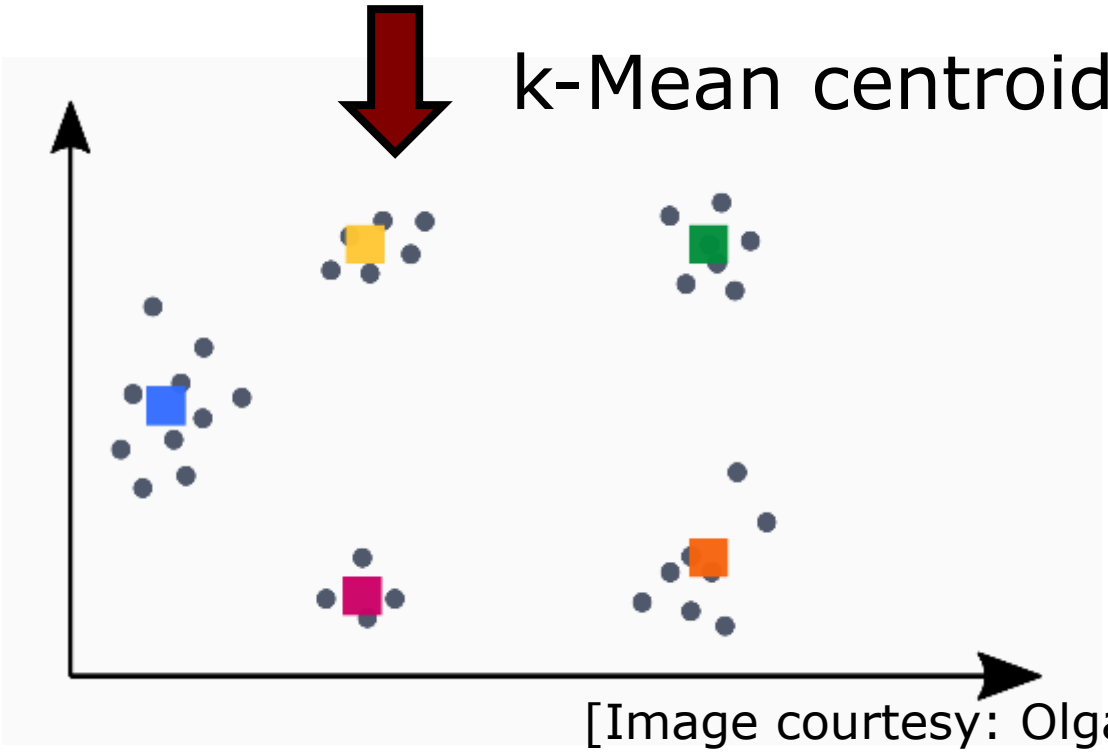

[Image courtesy: Olga Vysotska]

All Images are Reduced to Visual Words

[Image courtesy: Olga Vysotska]

All Images are Represented by Visual Word Occurrences

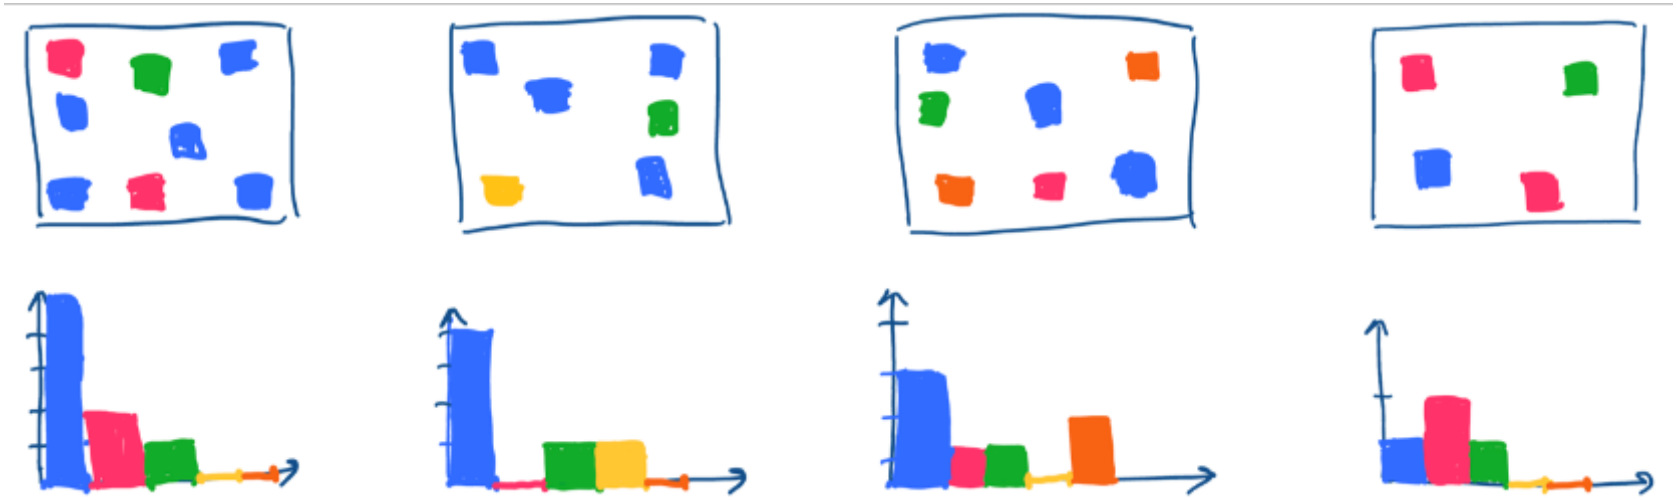

Every image turns into a histogram

[Image courtesy: Olga Vysotska]

Bag of Visual Words Model

- Compact summary of the image content
- Largely invariant to viewpoint changes and deformations
- Ignores the spatial arrangement
- Unclear how to choose optimal size of the vocabulary
  - Too small: Words not representative of all image regions
  - Too large: Over-fitting

How to Find Similar Images?

Task Description

- **Task:** Find similar looking images
- **Input:**
  - Database of images
  - Dictionary
  - Query image(s)
- **Output:**
  - The N most similar database images to the query image

?

Image Similarity by Comparing Word Occurrence Histograms

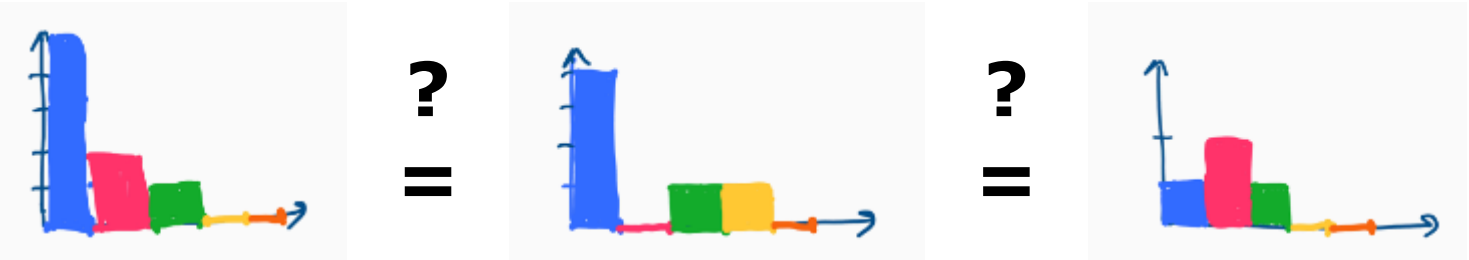

[Image courtesy: Olga Vysotska]

How to Compare Histograms?

- Euclidean distance of two points?
- Angle between two vectors?
- Kullback Leibler divergence (KLD)?
- Something else?

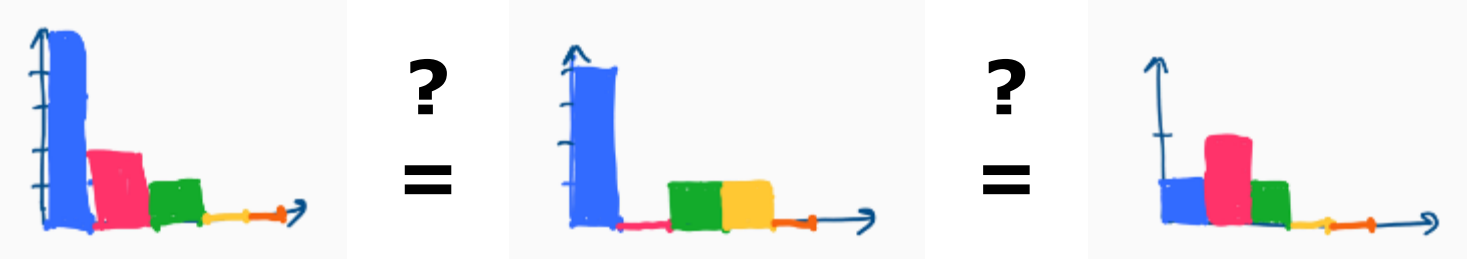

[Image courtesy: Olga Vysotska]

Are All Words Expressive for Comparing Histograms?

- Should all visual words be treated in the same way?
- Text analogy: What about articles?

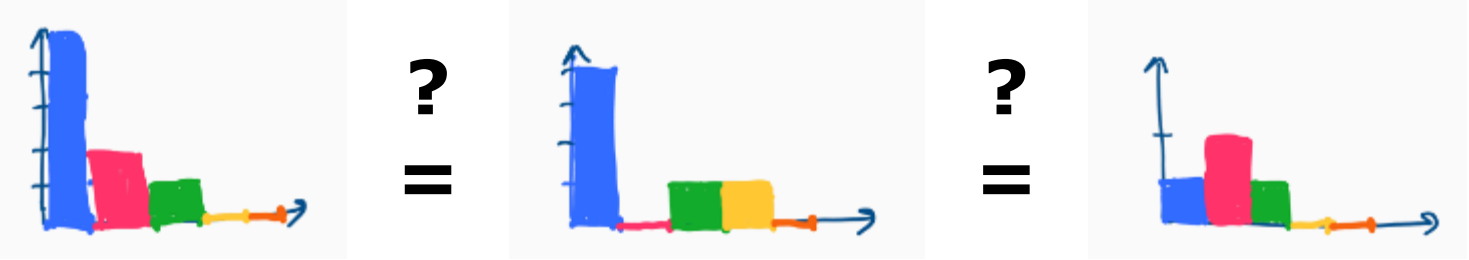

[Image courtesy: Olga Vysotska]

Some Word are Less Expressive Than Others!

- Words that occur in every image do not help a lot for comparisons

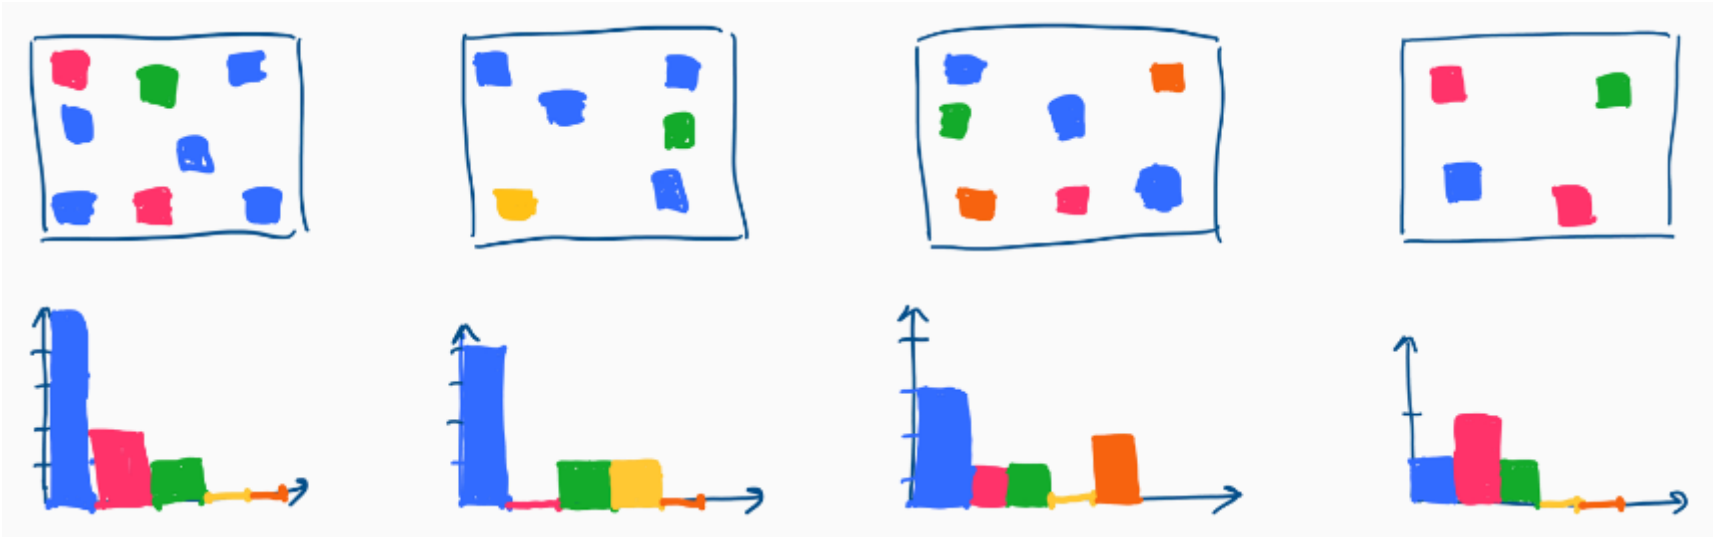

- Example: the “green word” is useless

[Image courtesy: Olga Vysotska]

TF-IDF Reweighting

- Weight words considering the probability that they appear
- TF-IDF = term frequency – inverse document frequency
- Every bin is reweighted

$$t_{id} = \frac{n_{id}}{n_d} \log \frac{N}{n_i}$$

bin **normalize** **weight**

TF-IDF

term frequency

$$t_{id} = \frac{n_{id}}{n_d} \log \frac{N}{n_i}$$

inverse document frequency

bin of word i in image d

- t_{id}: histogram bin of word i for image d
- n_{id}: occurances of word i in image d
- n_d: number of word occurances in image d
- n_i: number of images that contain word i
- N: number of images

Computing the TF-IDF (1)

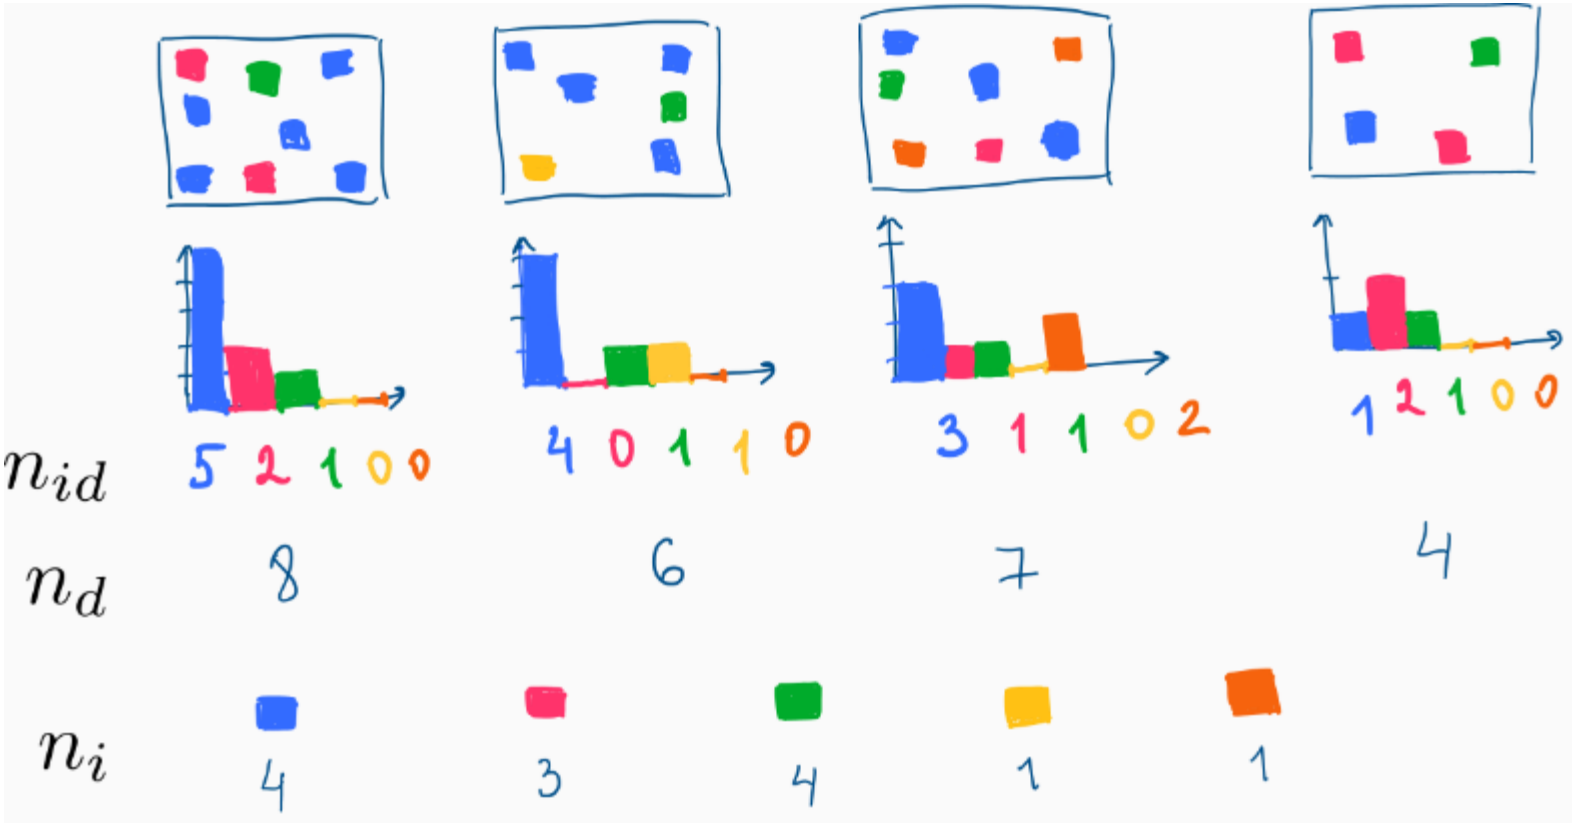

$$t_{id} = \frac{n_{id}}{n_d} \log \frac{N}{n_i}$$

[Image courtesy: Olga Vysotska]

Computing the TF-IDF (2)

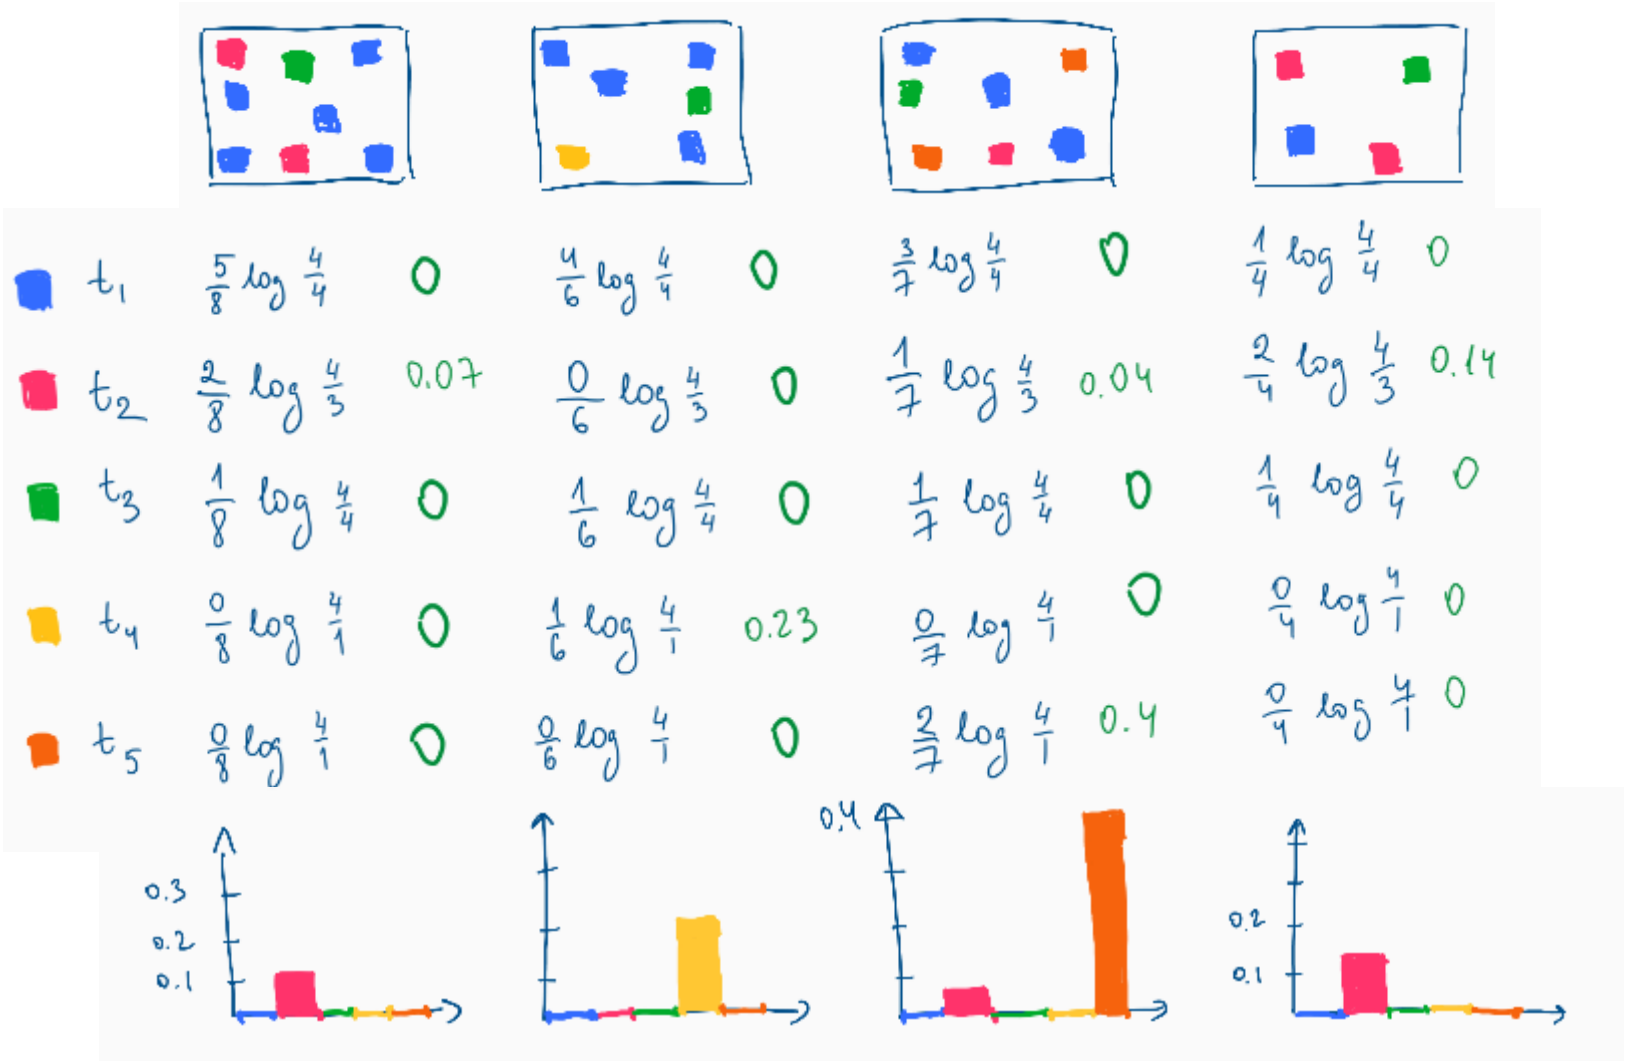

[Image courtesy: Olga Vysotska]

Reweighted Histograms

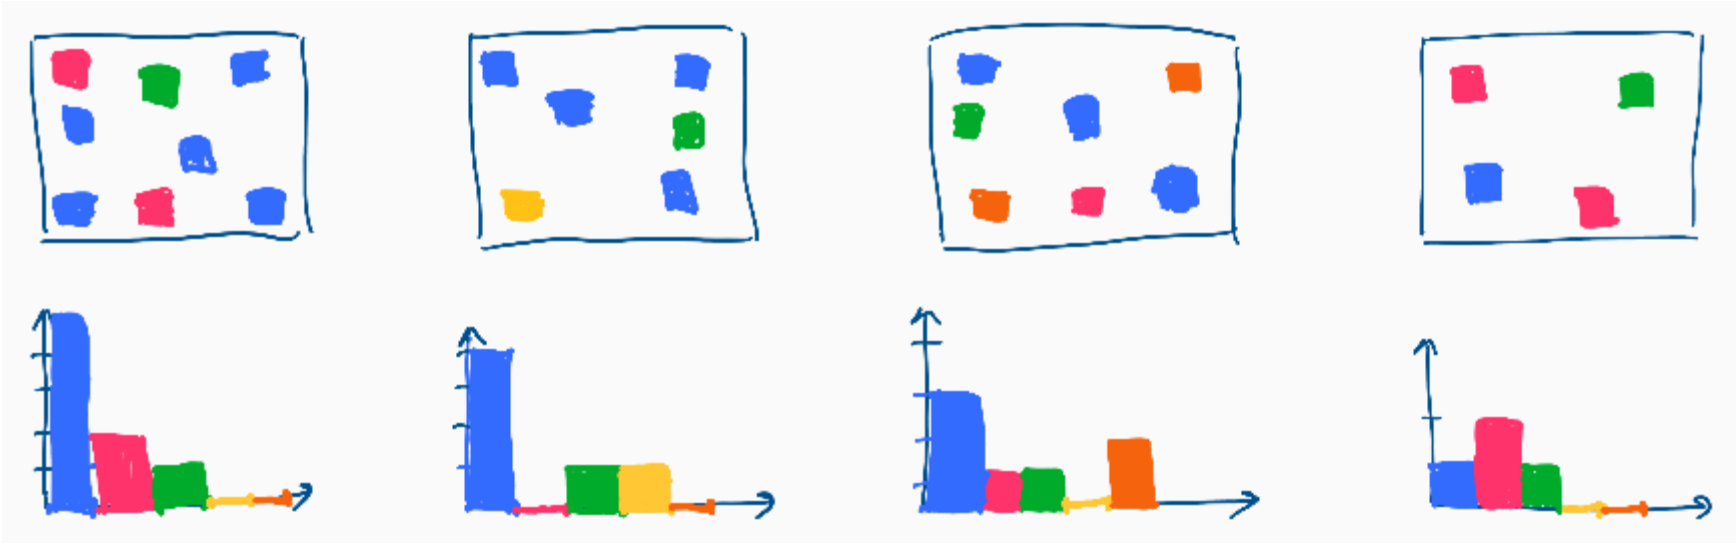

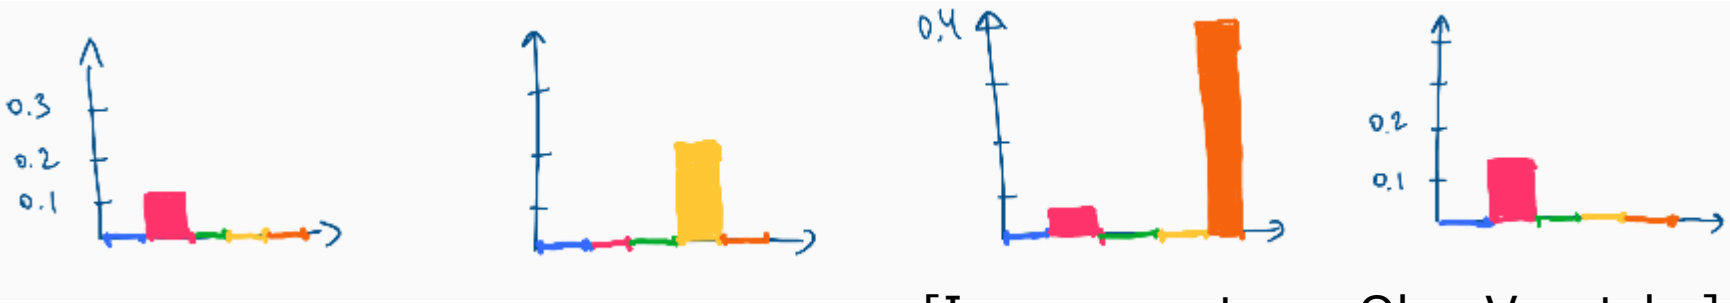

[Image courtesy: Olga Vysotska]

Reweighted Histograms

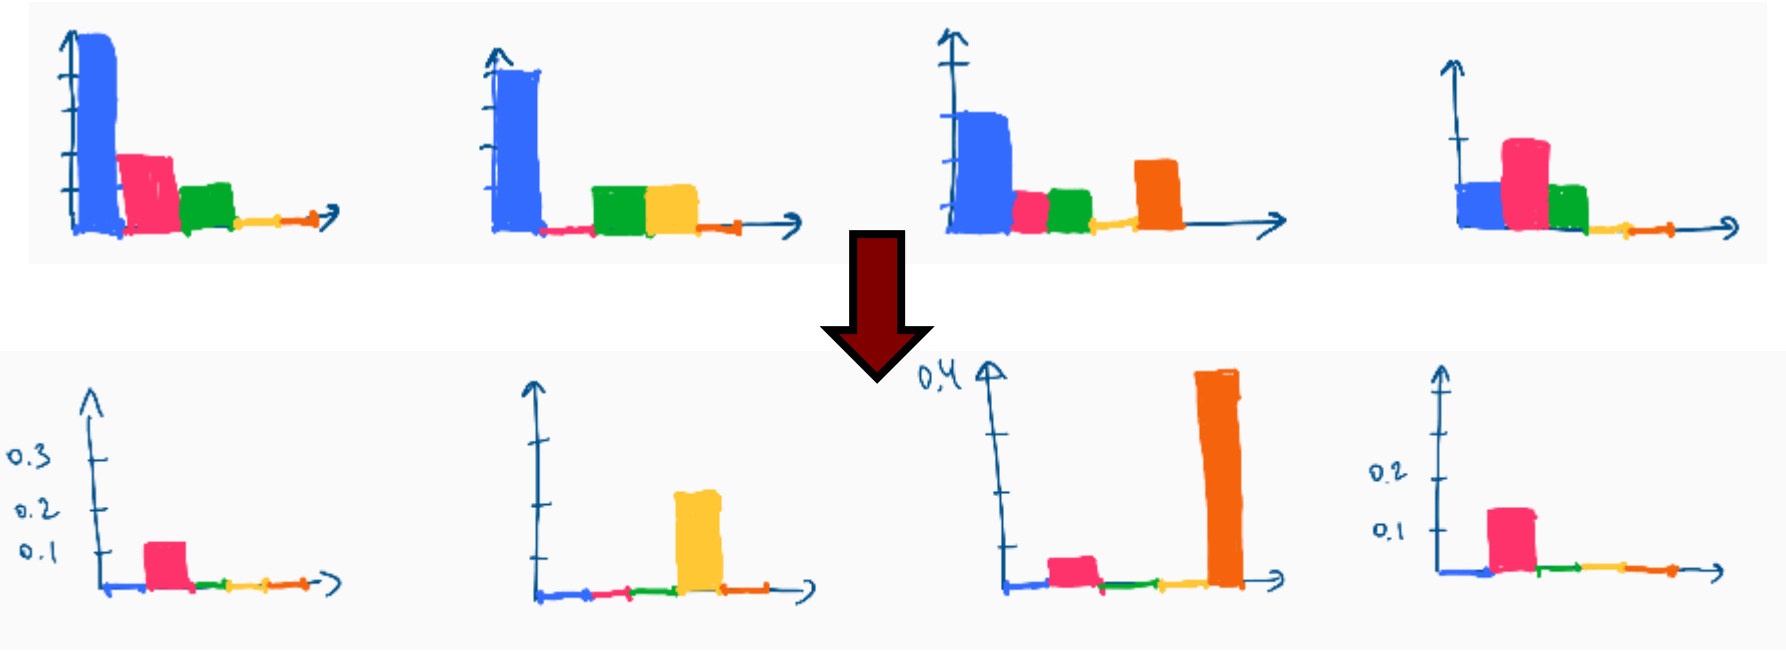

- Relevant words get higher weights
- Others are weighted down to zero (those occurring in every image)

[Image courtesy: Olga Vysotska]

Comparing Two Histograms

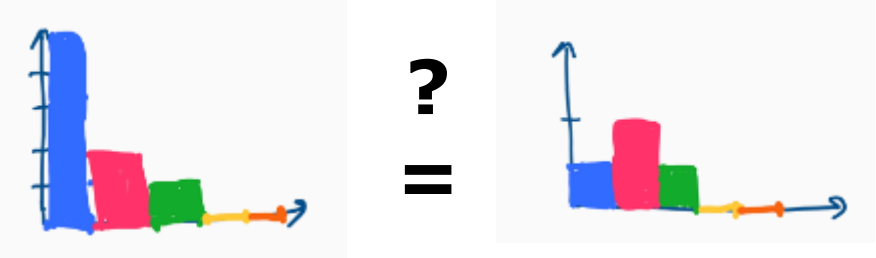

Options

- Euclidean distance of two points
- Angle between two vectors
- ~~Kullback Leibler divergence (KLD)~~

[Image courtesy: Olga Vysotska]

Comparing Two Histograms

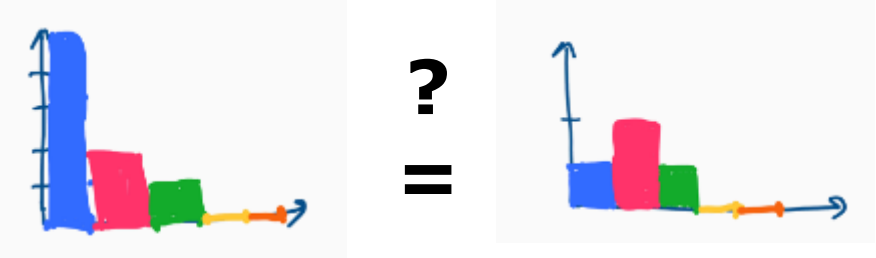

Options

- Euclidean distance of two vectors
- **Angle between two vectors**
- ~~Kullback Leibler divergence (KLD)~~

BoVW approaches often use the cosine distance for comparisons

[Image courtesy: Olga Vysotska]

Cosine Similarity and Distance

- Cosine similarity considers the cosine of the angle between vectors:

$$\text{cossim}(\mathbf{x}, \mathbf{y}) = \cos(\theta) = \frac{\mathbf{x}^\top \mathbf{y}}{||\mathbf{x}|| \, ||\mathbf{y}||}$$

- We use the cosine distance

$$\text{d}_{\cos}(\mathbf{x}, \mathbf{y}) = 1 - \text{cossim}(\mathbf{x}, \mathbf{y}) = 1 - \frac{\mathbf{x}^\top \mathbf{y}}{||\mathbf{x}|| \, ||\mathbf{y}||}$$

- Takes values between 0 and 1 (for vectors in the 1st quadrant)

Example Comparing Histograms

- 4 images
- Image 0 and image 3 are similar

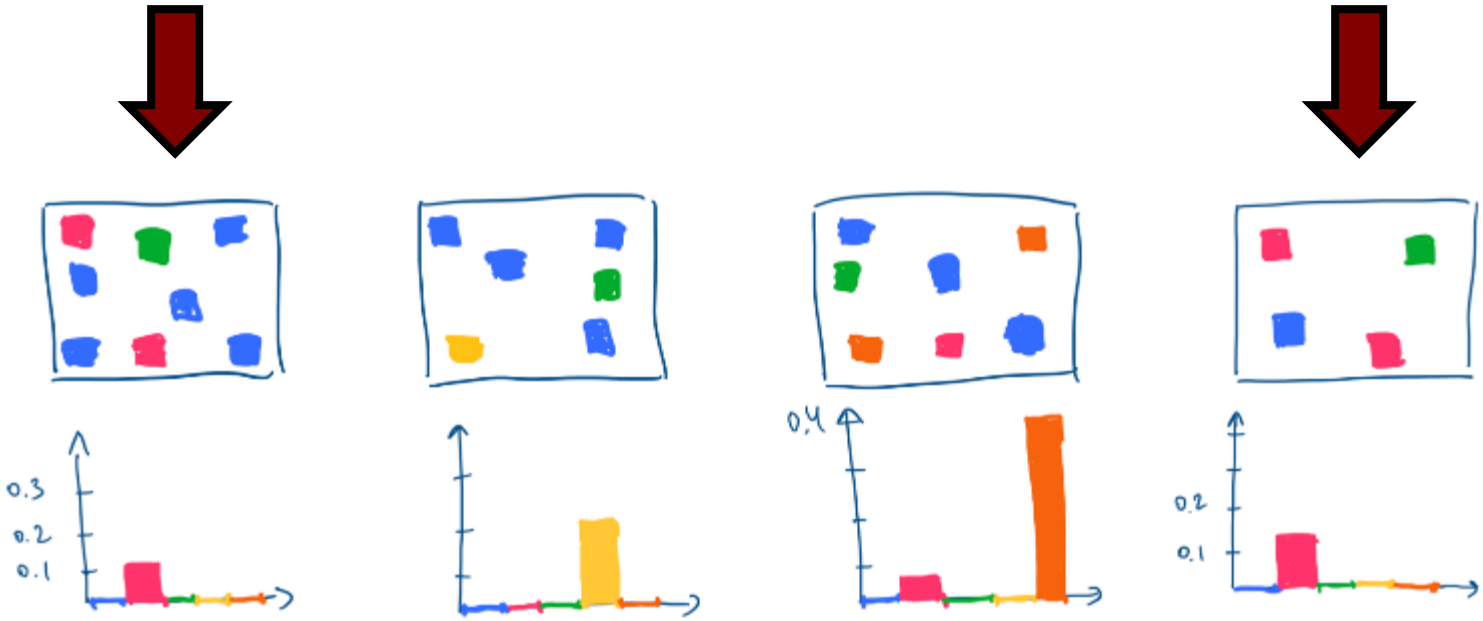

[Image courtesy: Olga Vysotska]

Example Comparing Histograms

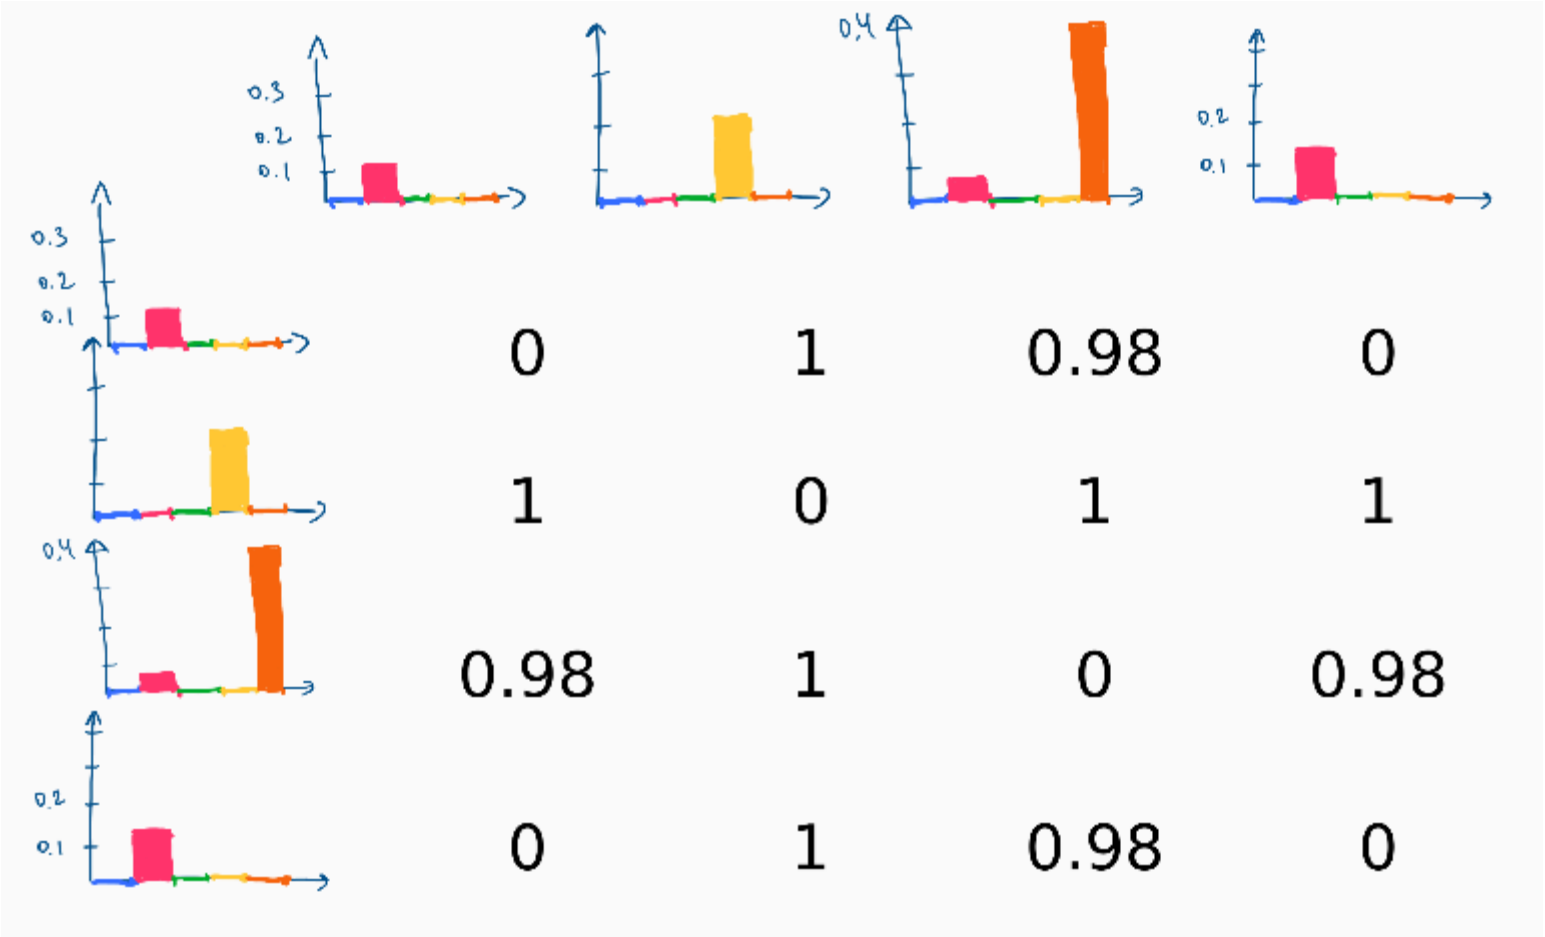

[Image courtesy: Olga Vysotska]

Example Comparing Histograms

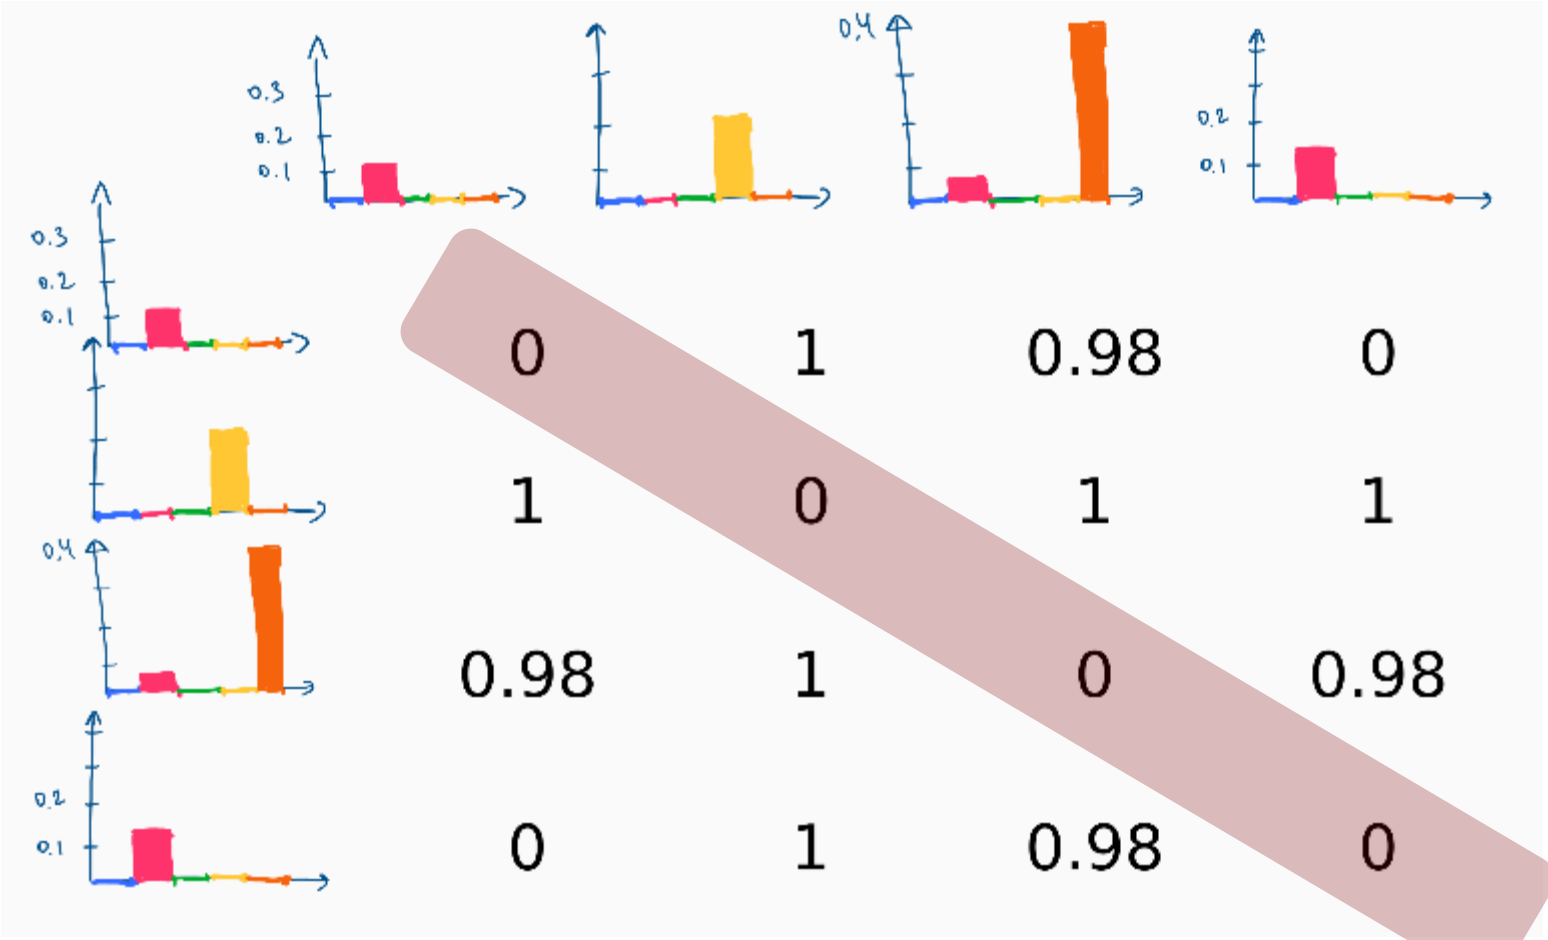

Images have a zero distance to themselves

[Image courtesy: Olga Vysotska]

Example Comparing Histograms

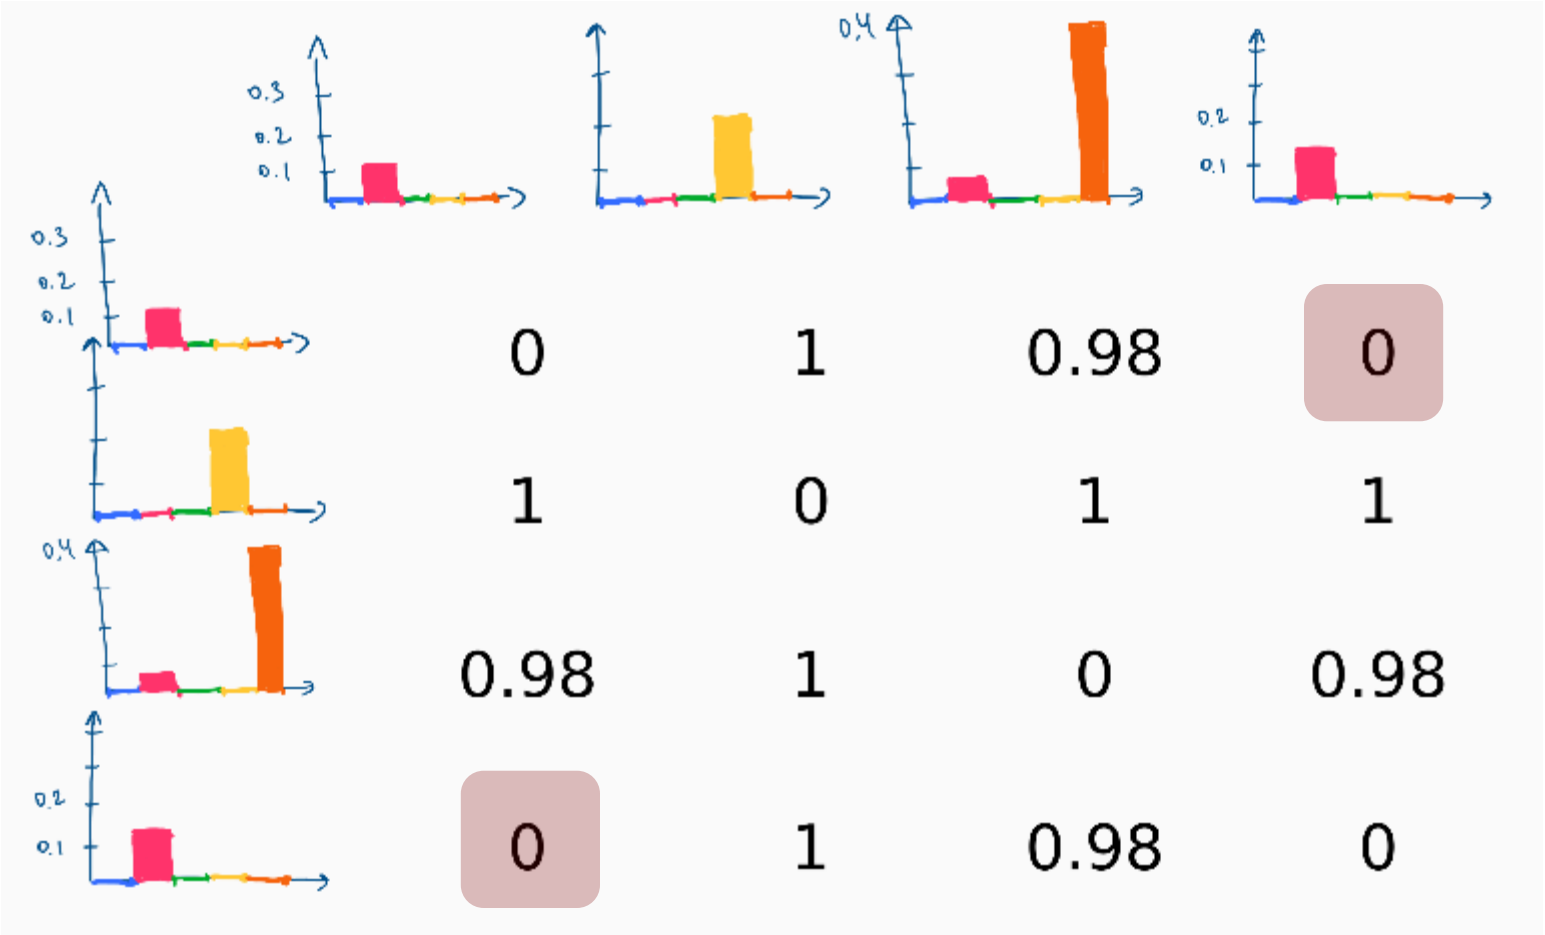

Images 0 and 3 are highly similar

[Image courtesy: Olga Vysotska]

Cost Matrix

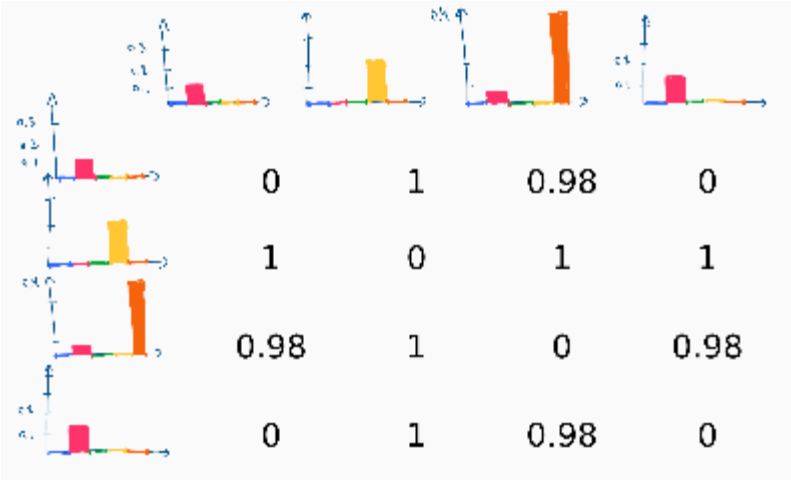

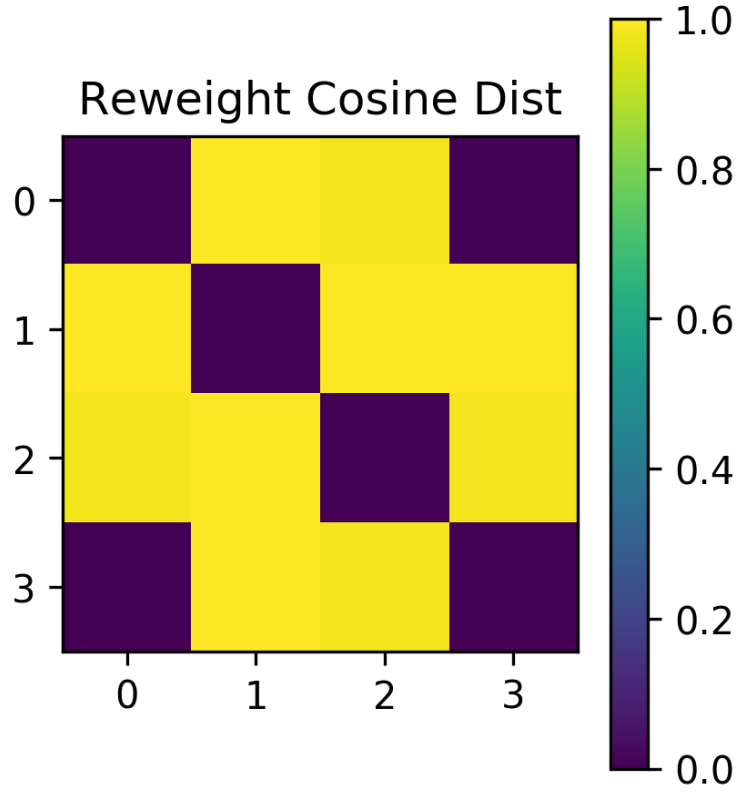

[Image courtesy: Olga Vysotska]

IF-IDF Actually Helps

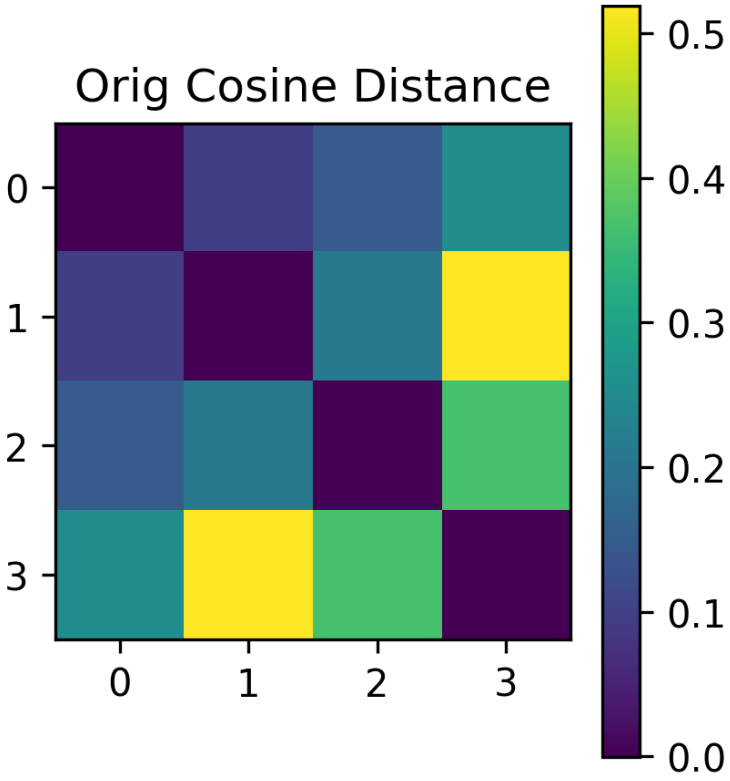

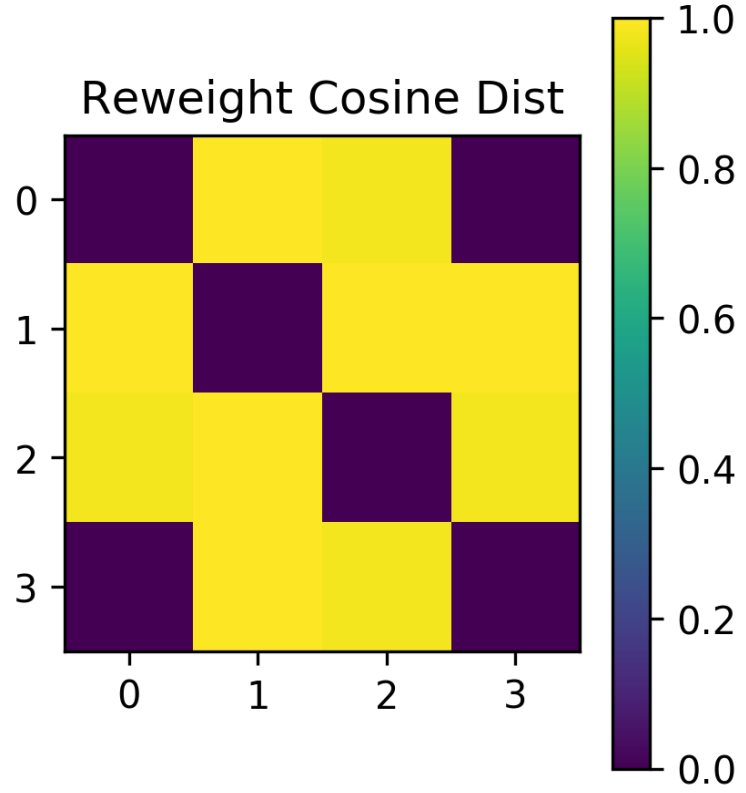

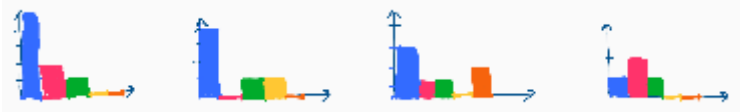

original histograms

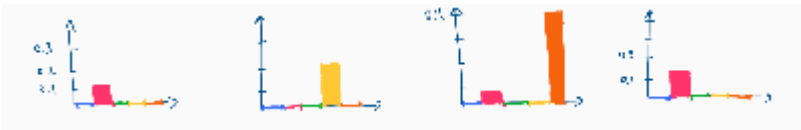

TF-IDF histograms

[Image courtesy: Olga Vysotska]

Euclidean vs. Cosine Distance

- Cosine distance ignores the length of the vectors
- **For vectors of length 1**, the squared Euclidean and the cosine distance only differ by a factor of 2:

$$\begin{aligned} ||\mathbf{x}-\mathbf{y}||^2 &= (\mathbf{x}-\mathbf{y})^\top(\mathbf{x}-\mathbf{y}) \\ &= \mathbf{x}^\top\mathbf{x} - 2\mathbf{x}^\top\mathbf{y} + \mathbf{y}^\top\mathbf{y} \end{aligned}$$

as $||\mathbf{x}|| = ||\mathbf{y}|| = 1$

$$\begin{aligned} ||\mathbf{x}-\mathbf{y}||^2 &= 2 - 2\mathbf{x}^\top\mathbf{y} = 2 - 2\cos\theta \\ &= 2\,\mathrm{d}_{\cos}(\mathbf{x},\mathbf{y}) \end{aligned}$$

Comparison of Distance Metrics

Euclidean

cosine distance

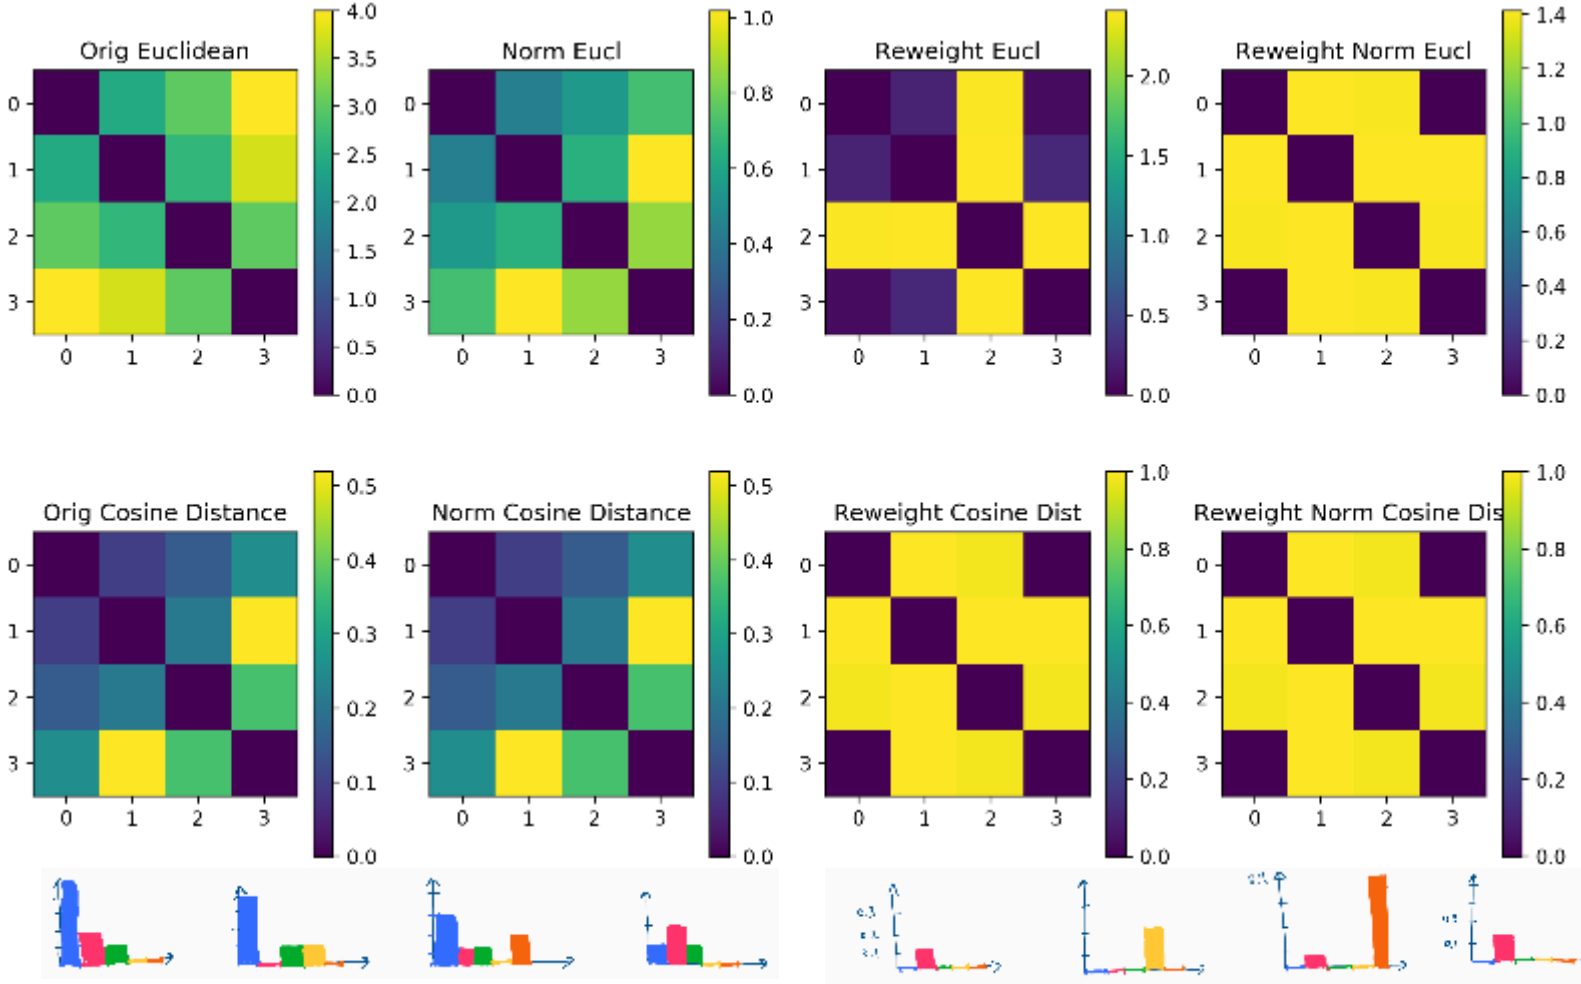

original histograms

TF-IDF histograms

[Image courtesy: Olga Vysotska]

Comparison of Distance Metrics

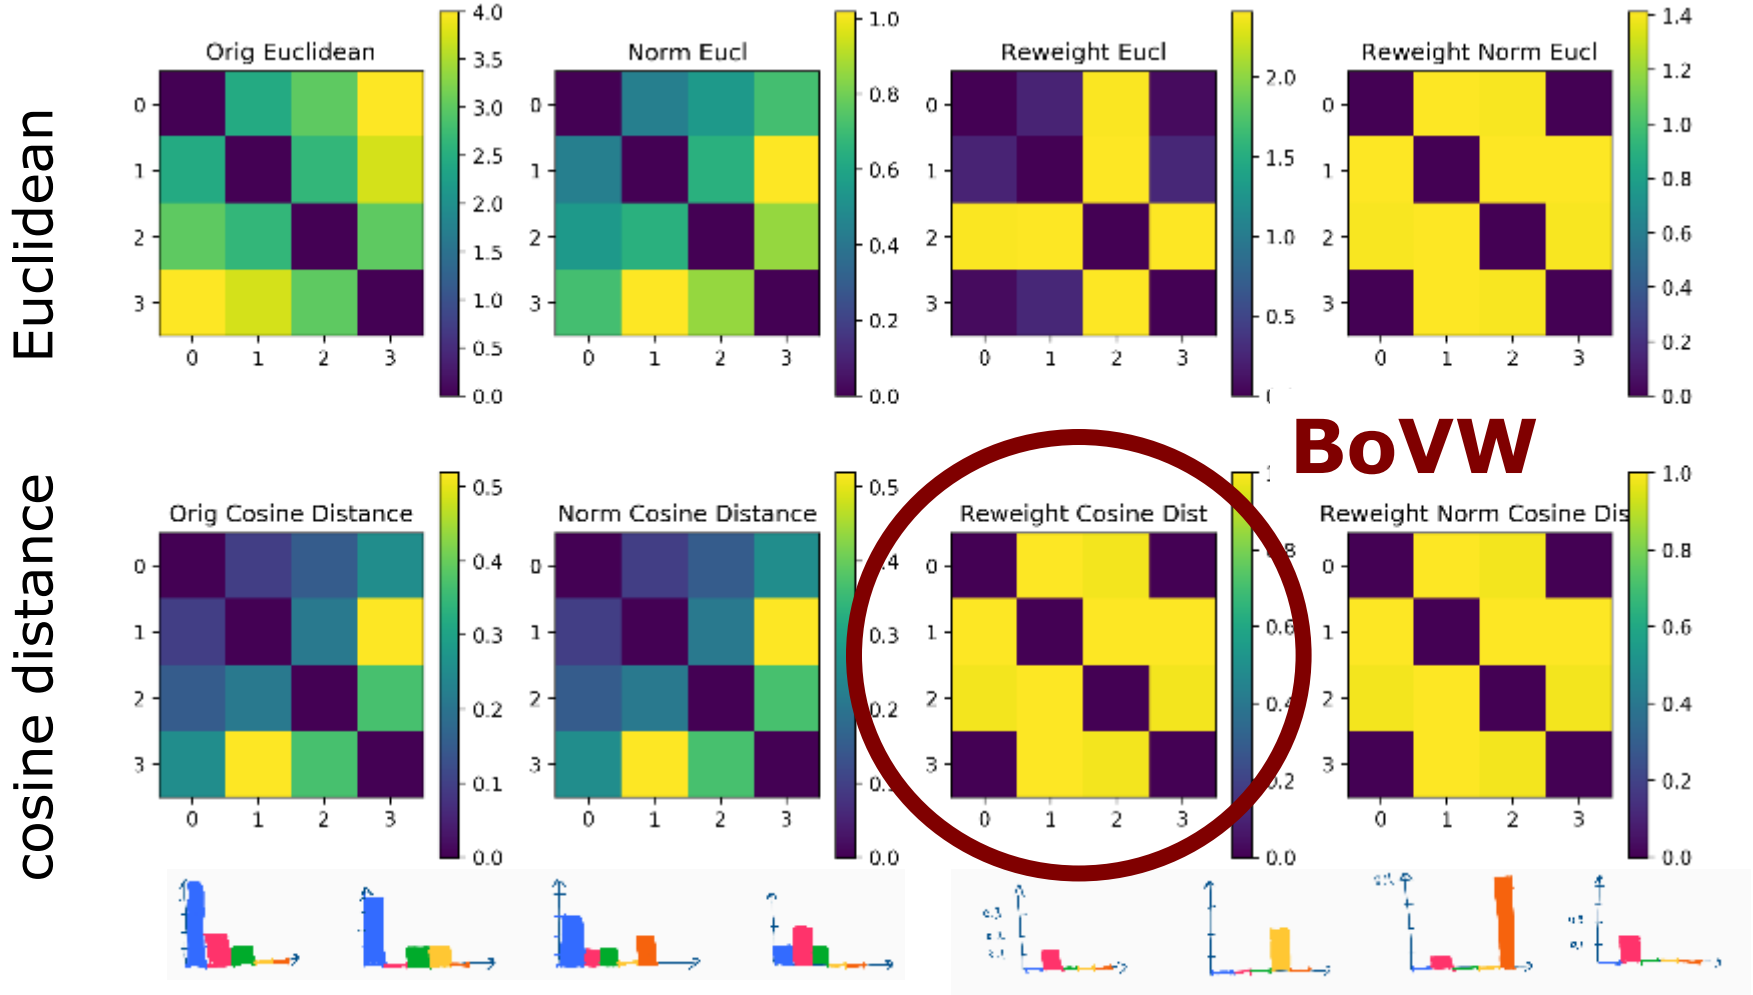

[Image courtesy: Olga Vysotska]

Similarity Queries

- Database stores TF-IDF weighted histograms for all database images

Find similar images by

- Extract features from query image
- Assign features to visual words
- Build TF-IDF histogram for query image
- Return N most similar histograms from database under cosine distance

Further Material

- Bag of Visual Words in 5 Minutes:
  https://www.youtube.com/watch?v=a4cFONdc6nc

Further Material

- Jupyter notebook by Olga Vysotska: https://github.com/ovysotska/in_simple_english/blob/master/bag_of_visual_words.ipynb

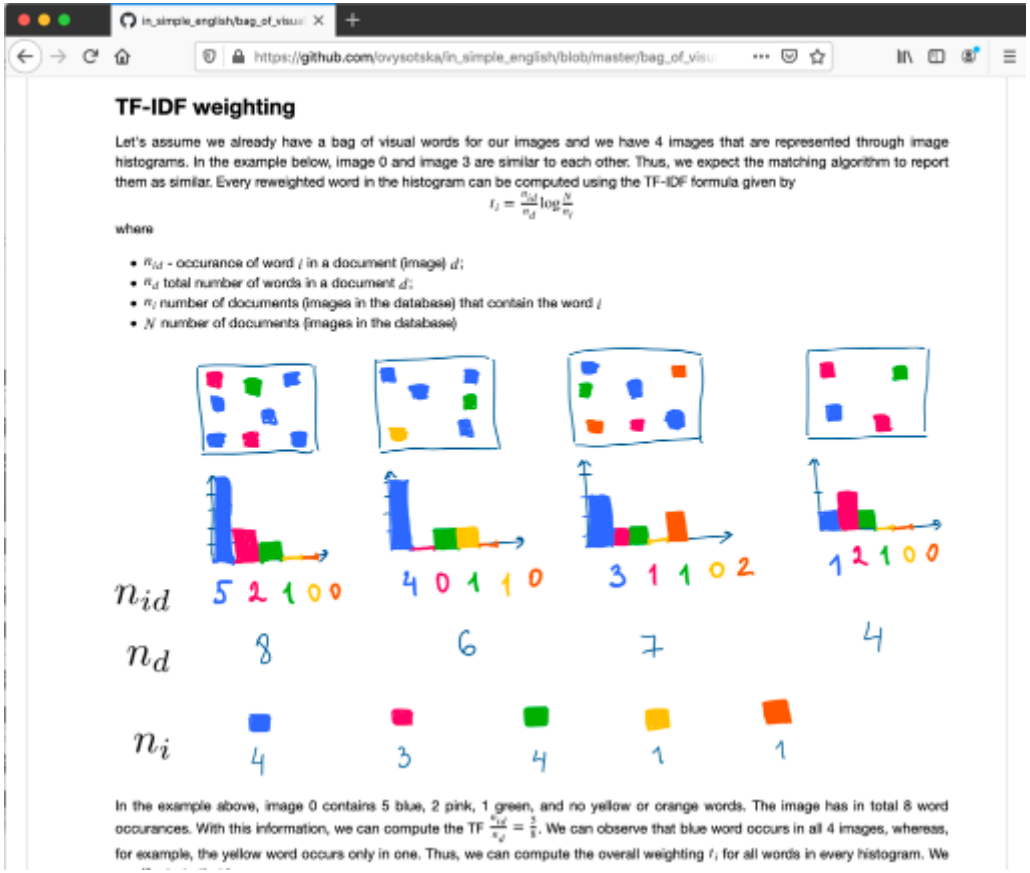

TF-IDF weighting

Let's assume we already have a bag of visual words for our images and we have 4 images that are represented through image histograms. In the example below, image 0 and image 3 are similar to each other. Thus, we expect the matching algorithm to report them as similar. Every reweighted word in the histogram can be computed using the TF-IDF formula given by

$$t_i = \frac{n_{id}}{n_d} \log \frac{N}{n_i}$$

where

- n_{id} - occurance of word i in a document (image) d;
- n_d total number of words in a document d;
- n_i number of documents (images in the database) that contain the word i
- N number of documents (images in the database)

In the example above, image 0 contains 5 blue, 2 pink, 1 green, and no yellow or orange words. The image has in total 8 word occurances. With this information, we can compute the TF $\frac{n_{id}}{n_d} = \frac{5}{8}$. We can observe that blue word occurs in all 4 images, whereas, for example, the yellow word occurs only in one. Thus, we can compute the overall weighting t_i for all words in every histogram. We

Further Material

- Bag of Visual Words in 5 Minutes: https://www.youtube.com/watch?v=a4cFONdc6nc
- Jupyter notebook by Olga Vysotska: https://github.com/ovysotska/in_simple_english/blob/master/bag_of_visual_words.ipynb
- Sivic and Zisserman. Video Google: A Text Retrieval Approach to Object Matching in Videos, 2003: http://www.robots.ox.ac.uk/~vgg/publications/papers/sivic03.pdf
- TF-IDF information: https://en.wikipedia.org/wiki/Tf%E2%80%93idf

Further Material

- Bag of Visual Words in 5 Minutes: https://www.youtube.com/watch?v=a4cFONdc6nc
- Jupyter notebook by Olga Vysotska: https://github.com/ovysotska/in_simple_english/blob/master/bag_of_visual_words.ipynb
- Sivic and Zisserman. Video Google: A Text Retrieval Approach to Object Matching in Videos, 2003: http://www.robots.ox.ac.uk/~vgg/publications/papers/sivic03.pdf
- TF-IDF information: https://en.wikipedia.org/wiki/Tf%E2%80%93idf

Summary

- BoVW is an approach to compactly describe images and compute similarities between images
- Based in a set of visual words
- Images become histograms of visual word occurrences
- TF-IDF weighting for increasing the influence of expressive words
- Similarity = histogram similarity
- Cosine distance

Small Project

Task Description

- **Task:** Realize a visual place recognition system using BoVW
- **Input:**
  - Database of images
  - Query image(s)
- **Output:**
  - The most similar 10 images to the query image
- Implementation in C++

Hints

- Read/write features in **binary** files for loading/saving the descriptor values
- Test k-means with **tiny 2D examples**
- **k-means** without FLANN **will be slow**
- FLANN = Fast approximate NN search
- FLANN is an approximation and it is **non-deterministic** (output varies)
- **Dictionary size** to start with: 1000
- Visualize results by writing **simple html** files and display them with your browser

Data

- Download:
  https://uni-bonn.sciebo.de/s/c2d0a1ebbe575fdba2a35a8033f1e2ab

Freiburg dataset

- gps_info.txt (GPS w/ timestamps)
- image-timestamps.txt (image timestamps)
- imageCompressedCam0_00000000.png
- …
- imageCompressedCam0_000NNNN.png

Data Example

```
1337850707.108278
1337850707.558294
1337850708.058266
1337850708.508274
1337850709.008243
1337850709.458249
1337850709.958227
1337850710.408230
1337850710.908209
1337850711.358219
1337850711.858189
1337850712.358196
```

```
-0.103          GPS_RAW                 gps             time=1337850706.695046,utctime=[3x1]
{9,11,23},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
0.798           GPS_RAW                 gps             time=1337850707.596694,utctime=[3x1]
{9,11,24},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
1.700           GPS_RAW                 gps             time=1337850708.498274,utctime=[3x1]
{9,11,25},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
2.902           GPS_RAW                 gps             time=1337850709.700428,utctime=[3x1]
{9,11,26},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
3.704           GPS_RAW                 gps             time=1337850710.502017,utctime=[3x1]
{9,11,27},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
4.506           GPS_RAW                 gps             time=1337850711.303786,utctime=[3x1]
{9,11,28},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
5.908           GPS_RAW                 gps             time=1337850712.706173,utctime=[3x1]
{9,11,29},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
6.710           GPS_RAW                 gps             time=1337850713.507768,utctime=[3x1]
{9,11,30},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
7.514           GPS_RAW                 gps             time=1337850714.312165,utctime=[3x1]
{9,11,31},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
8.920           GPS_RAW                 gps             time=1337850715.718731,utctime=[3x1]
{9,11,32},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
9.723           GPS_RAW                 gps             time=1337850716.520946,utctime=[3x1]
{9,11,33},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
10.524          GPS_RAW                 gps             time=1337850717.322499,utctime=[3x1]
{9,11,34},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
11.927          GPS_RAW                 gps             time=1337850718.724799,utctime=[3x1]
{9,11,35},lat=48.014315,lon=7.832350,qual=1,sats=8,hdop=1.3
```

Next Steps

1. Read the Jupyter notebook by Vysotska
2. Read "*Video Google: A Text Retrieval Approach to Object Matching in Videos*" by Sivic and Zisserman
3. Identify the key components to implemennt
4. Identify dependencies as well as inputs and outputs between components
5. Create a schedule and assign tasks
6. Go!

Rules

- Team work in teams of two students
- Code all components yourself

Two exceptions:

1. Use OpenCV only for loading/displaying images and for extracting SIFT features

2. If your nearest neighbor queries are too slow, use approximate NN techniques (FLANN - Fast Approximate Nearest Neighbor Search in OpenCV 2.4+)