Feature Descriptors

16-385 Computer Vision (Kris Kitani)

Carnegie Mellon University

Tiny Images

Just downsample it

Simple
Fast
Robust to small affine transformation

What are the problems?

Multi-Scale Oriented Patches (MOPS)

Multi-Image Matching using Multi-Scale Oriented Patches. M. Brown, R. Szeliski and S. Winder. International Conference on Computer Vision and Pattern Recognition (CVPR2005). pages 510-517

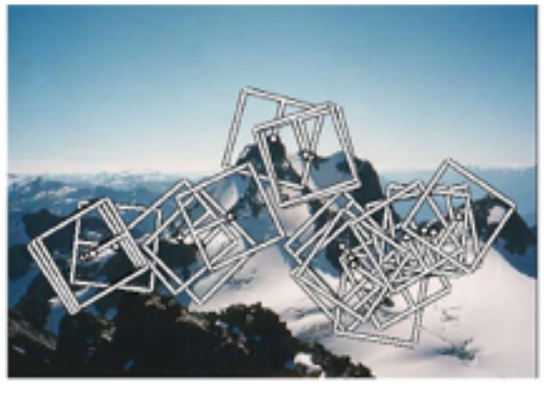

Multi-Scale Oriented Patches (MOPS)

Multi-Image Matching using Multi-Scale Oriented Patches. M. Brown, R. Szeliski and S. Winder. International Conference on Computer Vision and Pattern Recognition (CVPR2005). pages 510-517

Given a feature (x, y, s, θ)

Get 40 x 40 image patch, subsample every 5th pixel
(*what's the purpose of this step?*)

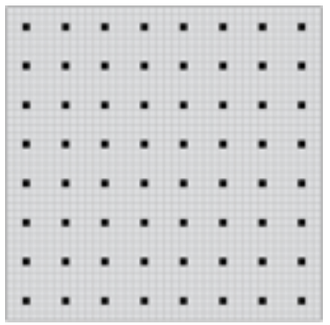

Subtract the mean, divide by standard deviation
(*what's the purpose of this step?*)

Haar Wavelet Transform
(*what's the purpose of this step?*)

Multi-Scale Oriented Patches (MOPS)

Multi-Image Matching using Multi-Scale Oriented Patches. M. Brown, R. Szeliski and S. Winder. International Conference on Computer Vision and Pattern Recognition (CVPR2005). pages 510-517

Given a feature (x, y, s, θ)

Get 40 x 40 image patch, subsample every 5th pixel
(low frequency filtering, absorbs localization errors)

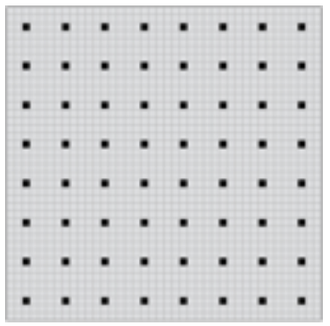

Subtract the mean, divide by standard deviation
(*what's the purpose of this step?*)

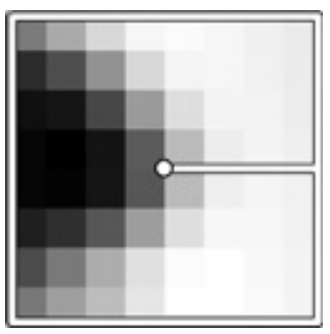

Haar Wavelet Transform
(*what's the purpose of this step?*)

Multi-Scale Oriented Patches (MOPS)

Multi-Image Matching using Multi-Scale Oriented Patches. M. Brown, R. Szeliski and S. Winder. International Conference on Computer Vision and Pattern Recognition (CVPR2005). pages 510-517

Given a feature (x, y, s, θ)

Get 40 x 40 image patch, subsample every 5th pixel
(low frequency filtering, absorbs localization errors)

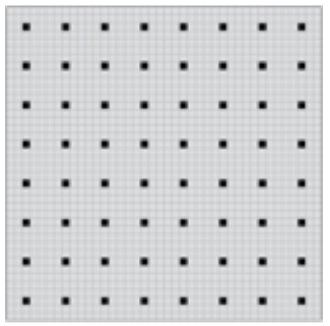

Subtract the mean, divide by standard deviation
(removes bias and gain)

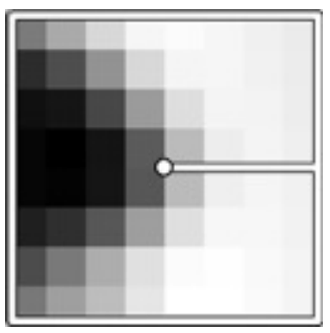

Haar Wavelet Transform
(*what's the purpose of this step?*)

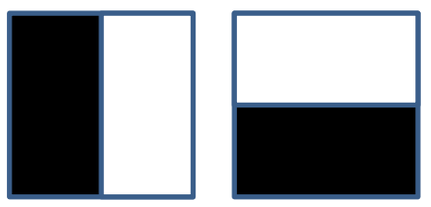

Multi-Scale Oriented Patches (MOPS)

Multi-Image Matching using Multi-Scale Oriented Patches. M. Brown, R. Szeliski and S. Winder. International Conference on Computer Vision and Pattern Recognition (CVPR2005). pages 510-517

Given a feature (x, y, s, θ)

Get 40 x 40 image patch, subsample every 5th pixel
(low frequency filtering, absorbs localization errors)

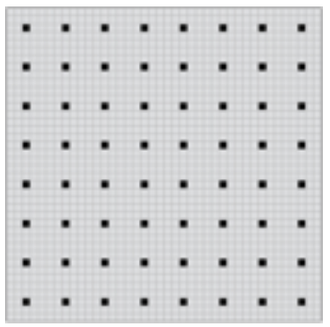

Subtract the mean, divide by standard deviation
(removes bias and gain)

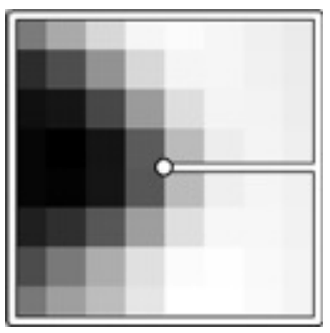

Haar Wavelet Transform
(low frequency projection)

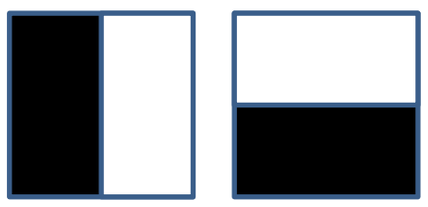

Haar Wavelets

(actually, Haar-like features)

Use responses of a bank of filters as a descriptor

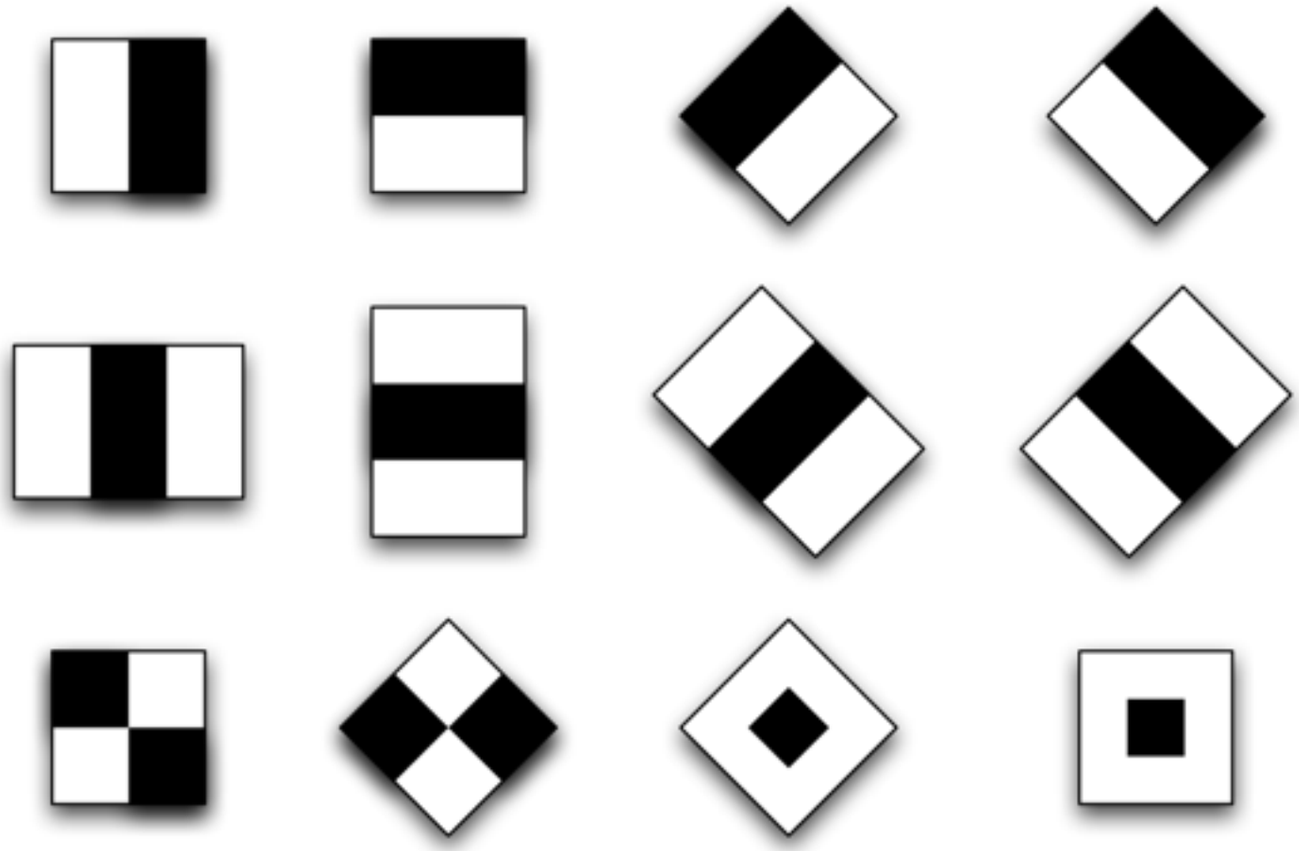

Haar wavelet responses can be computed with filtering

image patch

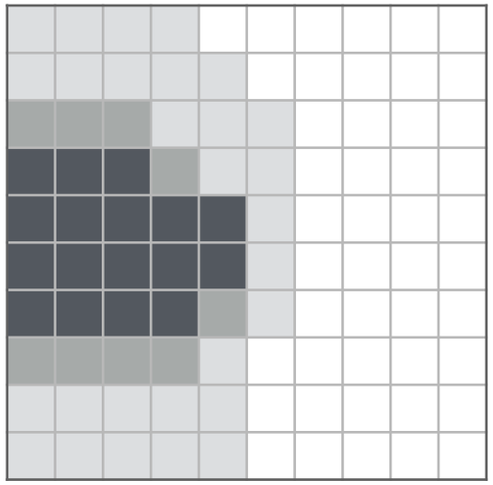

Haar wavelets filters

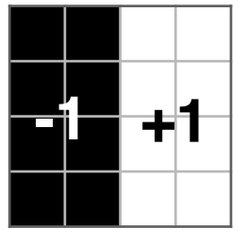

d_x

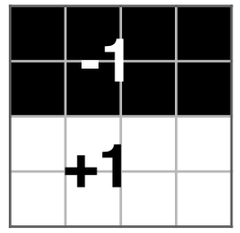

d_y

Haar wavelet responses can be computed
efficiently (in constant time) with integral images

Integral Image

original image

$I(x, y)$

1	5	2
2	4	1
2	1	1

$A(x, y)$

1	6	8
3	12	15
5	15	19

integral image

$$A(x, y) = \sum_{x' \leq x, y' \leq y} I(x', y')$$

Integral Image

original image

$I(x, y)$

1	5	2
2	4	1
2	1	1

$A(x, y)$

1	6	8
3	12	15
5	15	19

integral image

$$A(x, y) = \sum_{x' \leq x, y' \leq y} I(x', y')$$

Can find the **sum** of any block using **3** operations

$$A(x_1, y_1, x_2, y_2) = A(x_2, y_2) - A(x_1, y_2) - A(x_2, y_1) + A(x_1, y_1)$$

What is the sum of the bottom right 2x2 square?

$$A(x_1, y_1, x_2, y_2) = A(x_2, y_2) - A(x_1, y_2) - A(x_2, y_1) + A(x_1, y_1)$$

$I(x, y)$

1	5	2
2	4	1
2	1	1

image

$A(x, y)$

1	6	**8**
3	12	15
5	15	**19**

integral image

$$\begin{aligned} A(1,1,3,3) &= A(3,3) - A(1,3) - A(3,1) + A(1,1) \\ &= 19 - 8 - 5 + 1 \\ &= 7 \end{aligned}$$

Given an image patch, compute filter responses

filter bank (20 Haar wavelet filters)

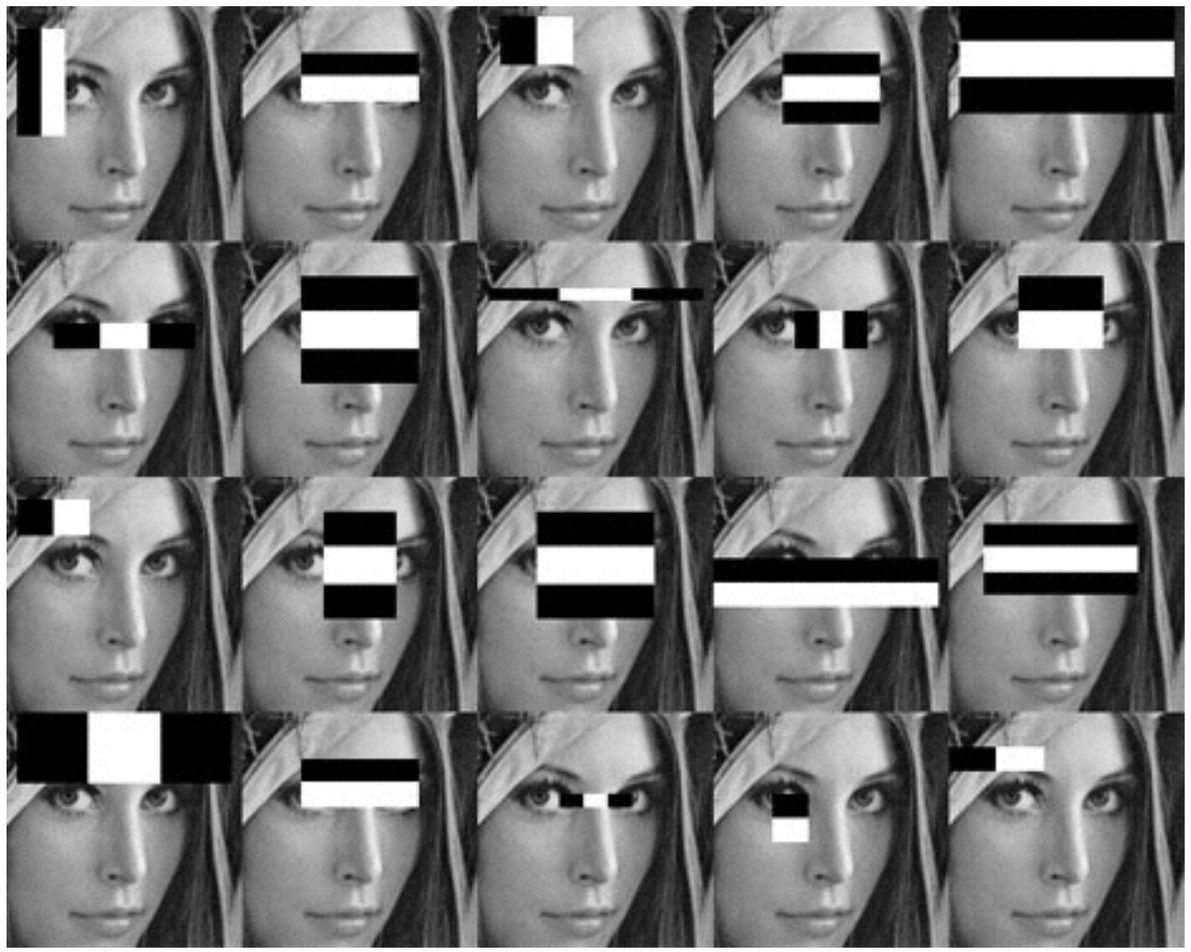

Responses are usually computed at specified location as a face patch descriptor

vector of filter responses

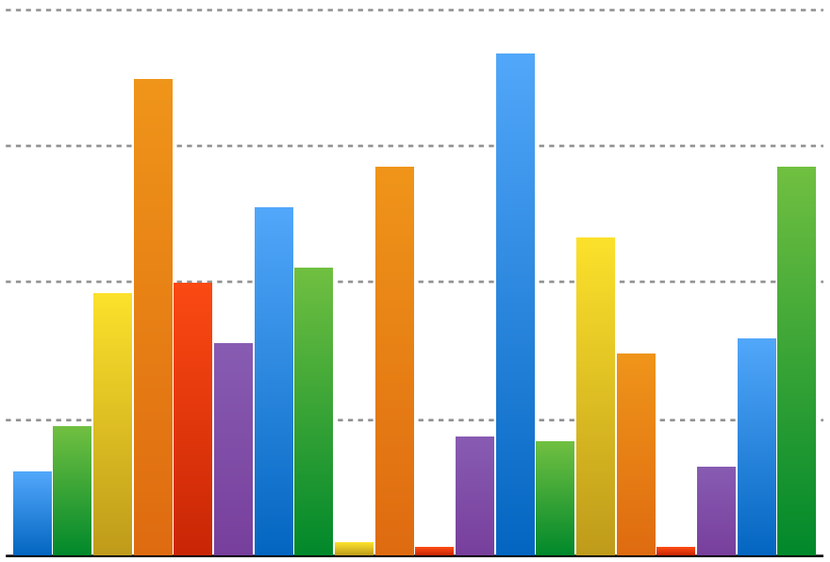

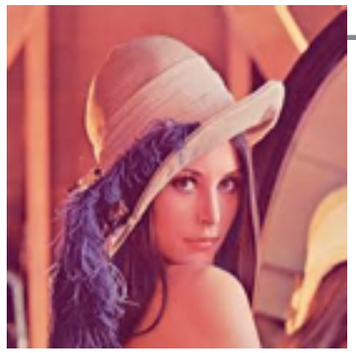

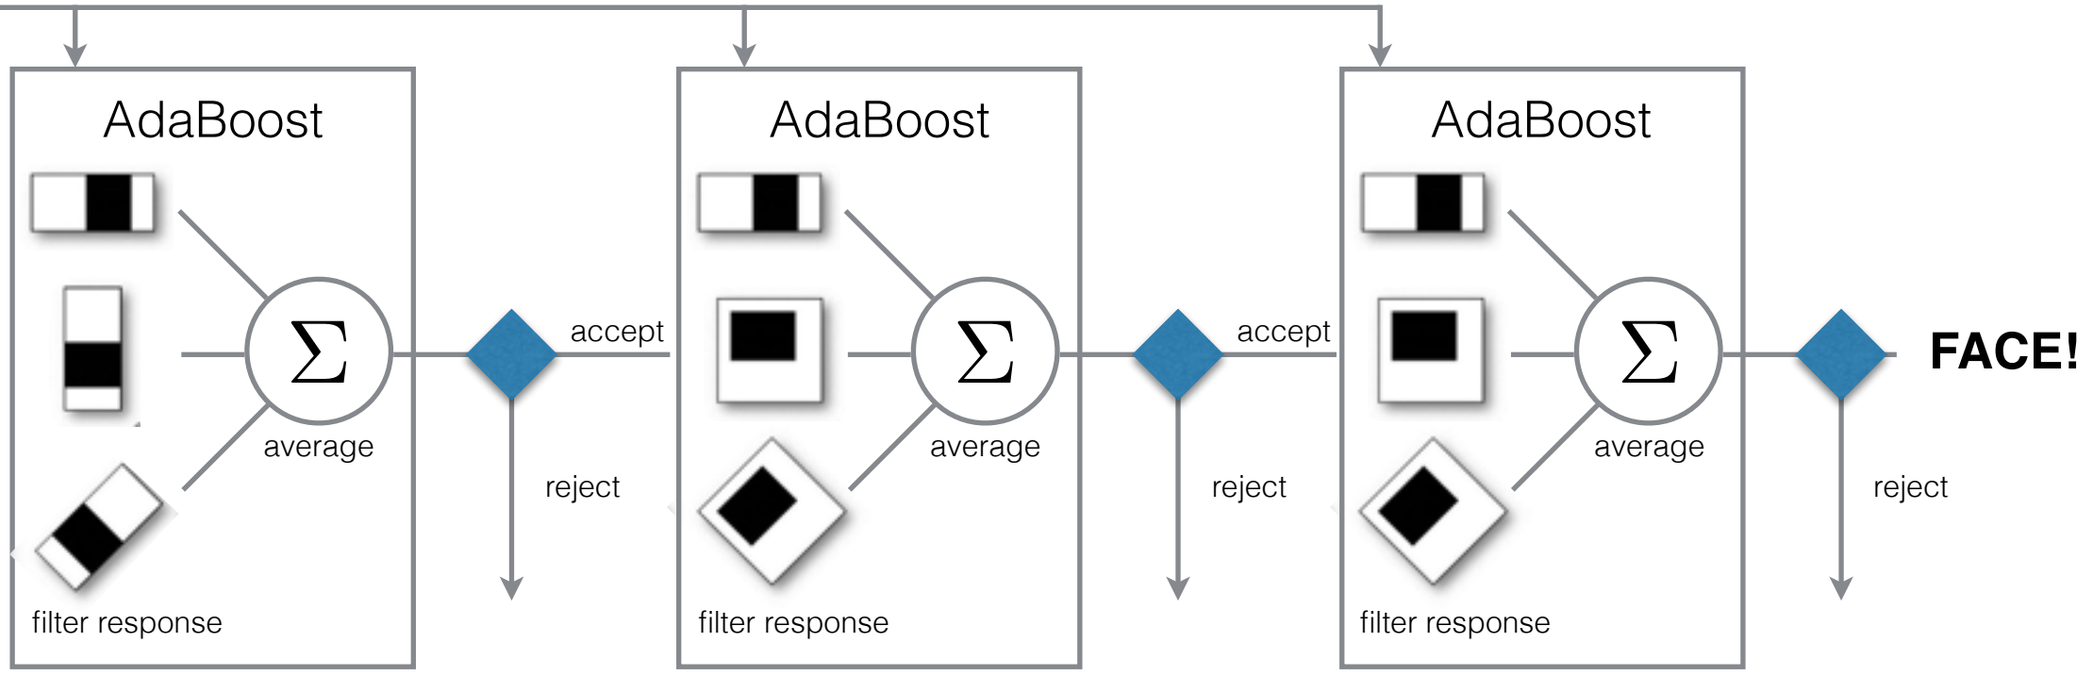
AdaBoost
AdaBoost
AdaBoost
accept
accept
FACE!
reject
reject
reject
average
average
average
filter response
filter response
filter response


Stage 1/2 (9%)
Feature 3/3

Given an image patch, compute filter responses

filter bank (20 Haar wavelet filters)

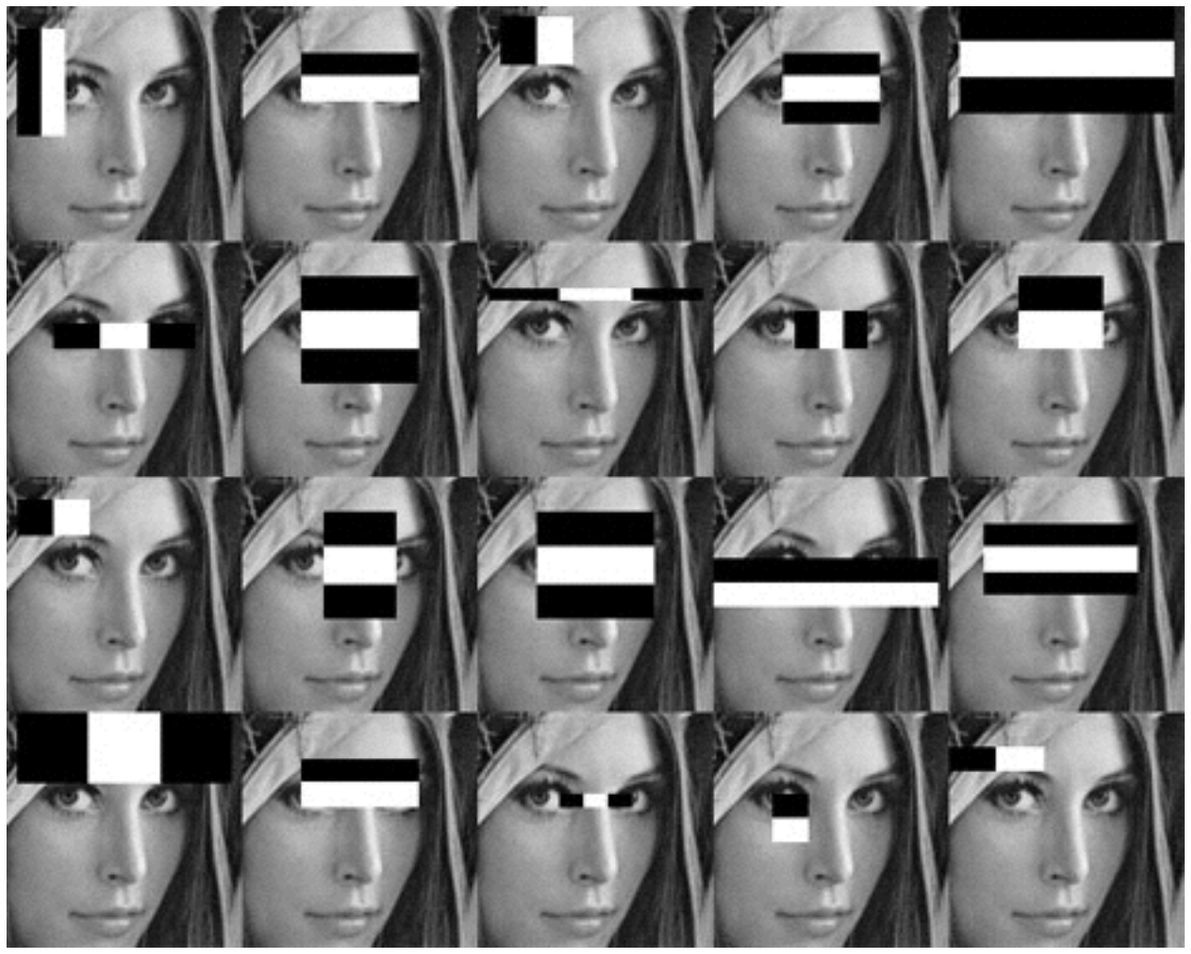

Responses are usually computed at specified location as a face patch descriptor

vector of filter responses

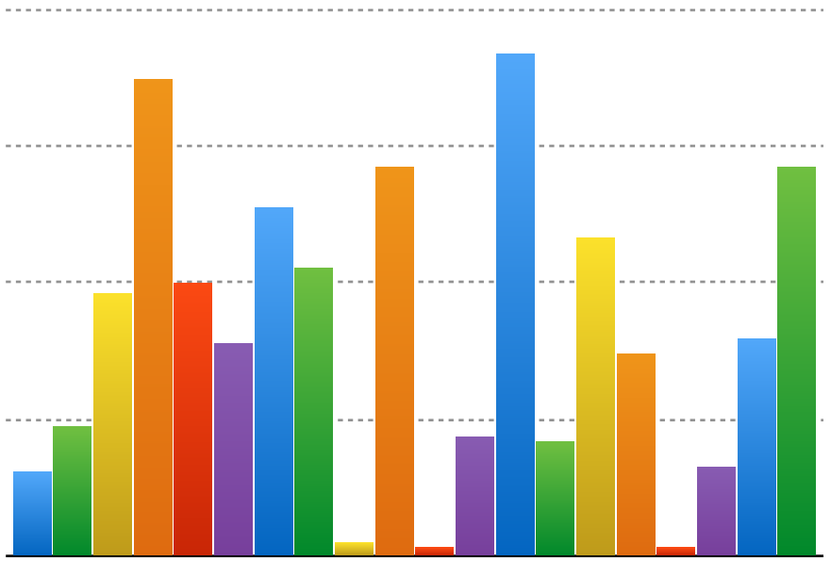

When will this feature descriptor fail?

IMAGES

For image requests, please contact me at U N D I S C L O S E D and be sure to specify which image and the pixel-dimension needed.

SUPPORT

Support this project by shopping at the Privacy Gift Shop

Face

Once computer vision programs detect a face, they can extract data about your emotions, age, and identity.

See how a face is detected

Anti Face

This face is unrecognizable to several state-of-art face detection algorithms.

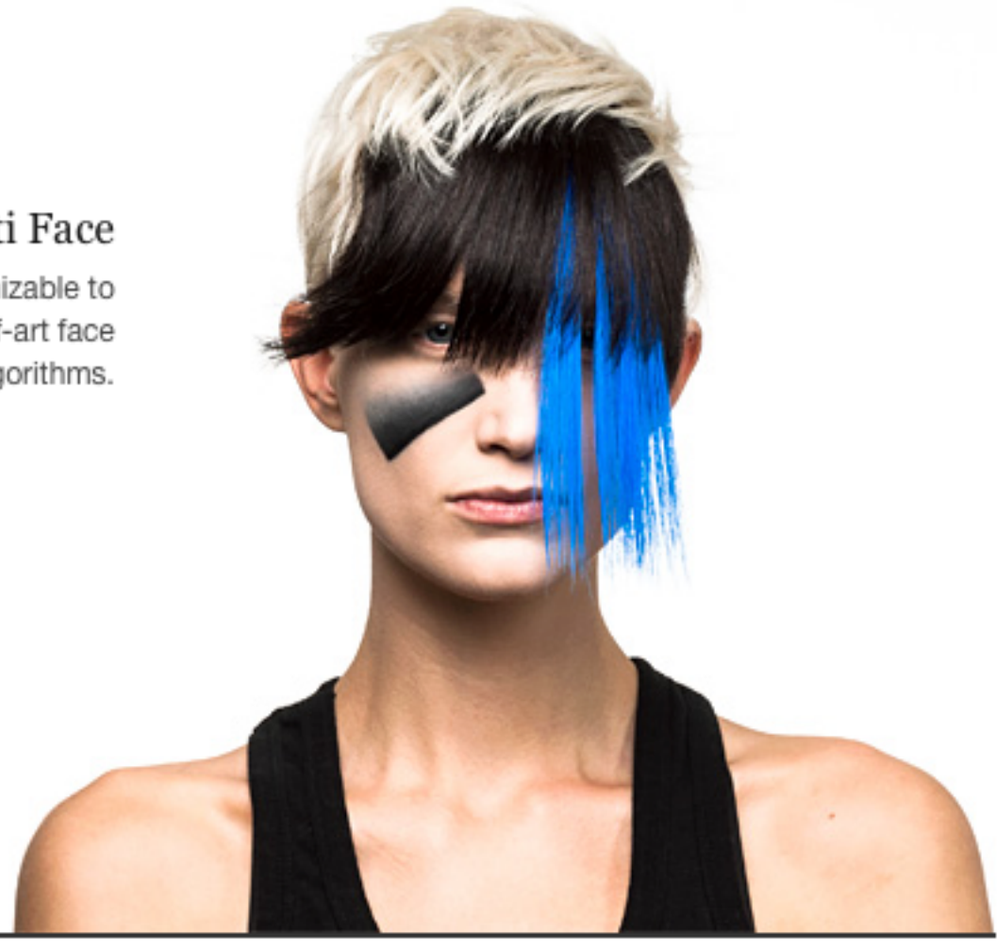

Camouflage from face detection.

CV Dazzle explores how fashion can be used as camouflage from face-detection technology, the first step in automated face recognition.

From all appearances, deception has always been critical to daily survival—for human and non-human creatures alike—and,