Matching Image Patches

René Magritte, “Golconda”

Many slides from
Steve Marschner, Alexei Efros

CSC320
Michael Guerzhoy

Template matching

Goal: find  in image

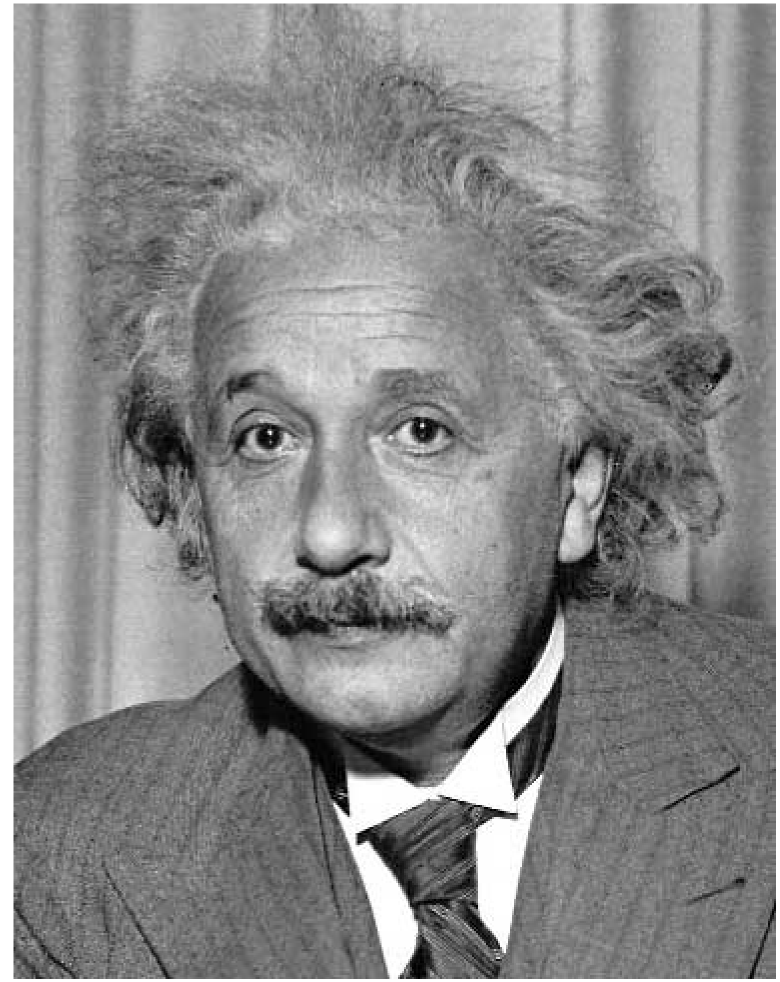

Main challenge: What is a good similarity or distance measure between two patches?

- Dot product
- (Zero-mean) correlation
- Sum Square Difference
- Normalized Cross Correlation

Slide contents from Derek Hoiem and Alexei Efros

Images as vectors

Matching with filters

Goal: find  in image

Method 0: filter the image with eye patch

$$h[m,n]=\sum_{k,l} g[k,l]\, f[m+k,n+l]$$

f = image
g = filter

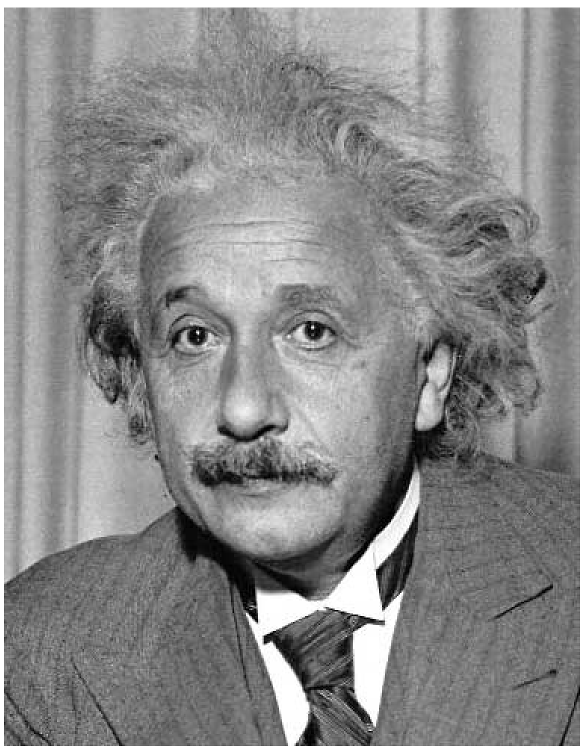

Input

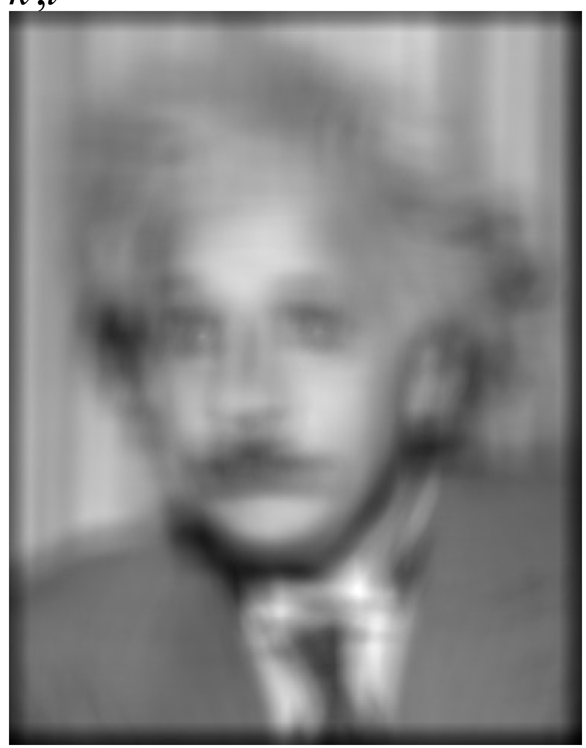

Filtered Image

What went wrong?

Slide by Derek Hoiem

Matching with filters

Goal: find  in image

Method 1: filter the image with zero-mean eye

$$h[m,n]=\sum_{k,l}(f[k,l]-\bar{f})(g[m+k,n+l])$$

mean of f (here, f is the filter)

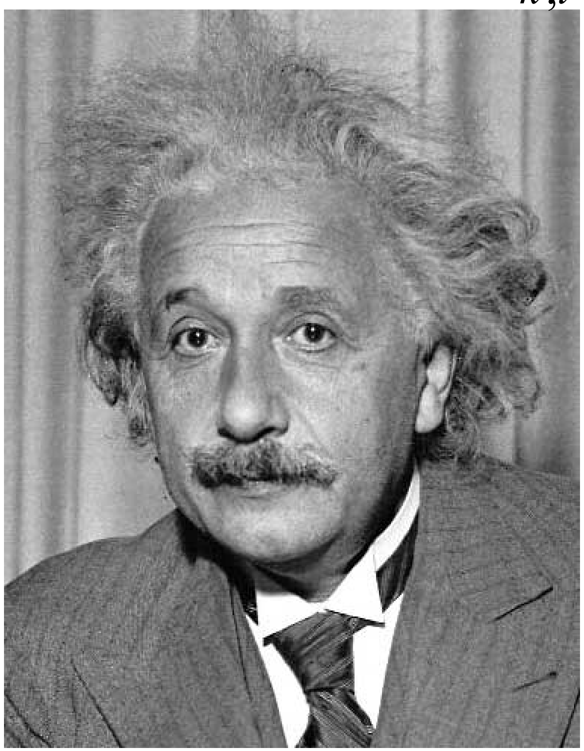

Input

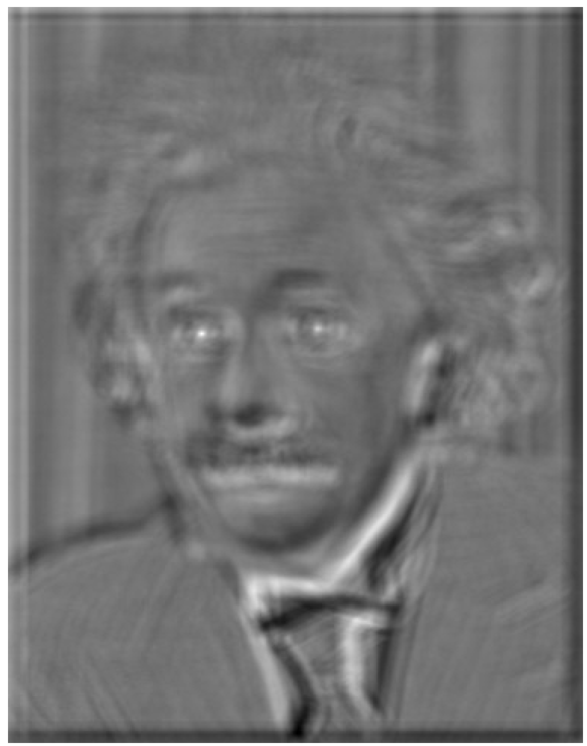

Filtered Image (scaled)

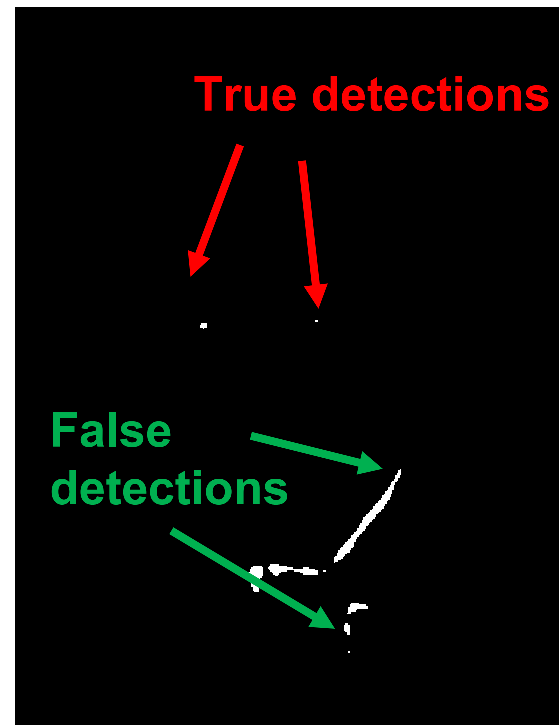

Thresholded Image

Matching with filters

Goal: find  in image

Method 2: SSD

$$h[m,n] = \sum_{k,l} (g[k,l] - f[m+k, n+l])^2$$

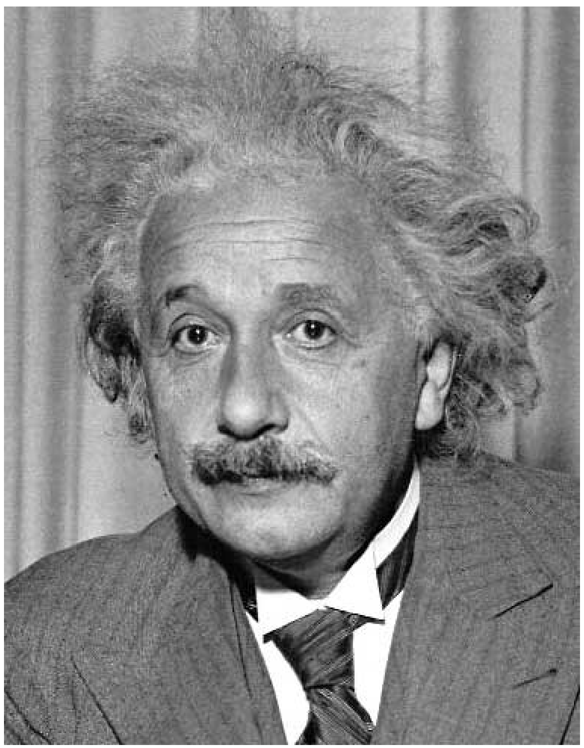

Input

1- sqrt(SSD)

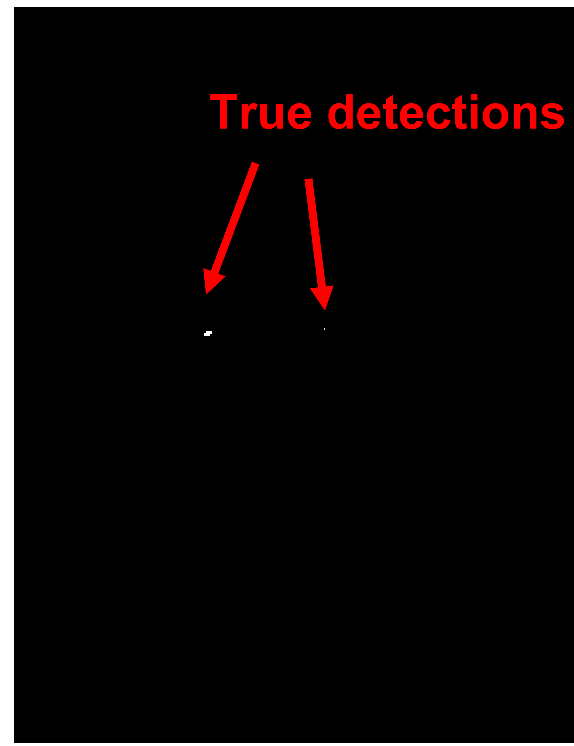

Thresholded Image

Matching with filters

Goal: find  in image

What's the potential downside of SSD?

Method 2: SSD

$$h[m,n]=\sum_{k,l}(g[k,l]-f[m+k,n+l])^2$$

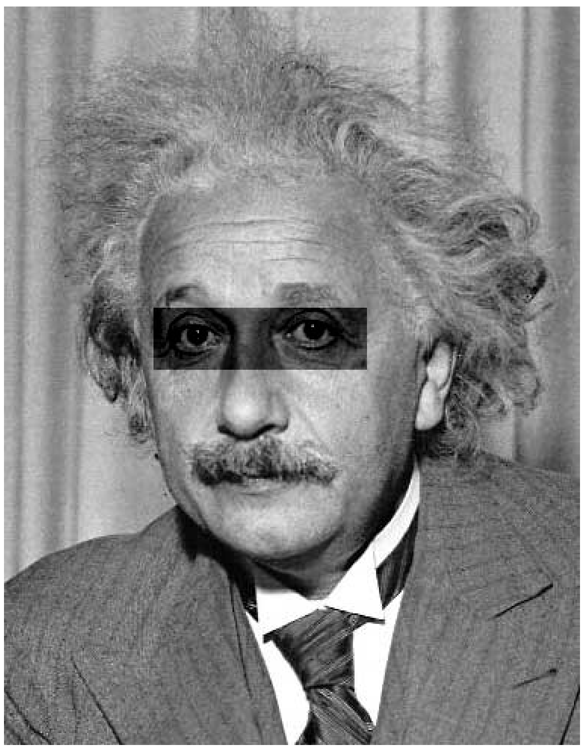

Input

1- sqrt(SSD)

Slide by Derek Hoiem

Matching with filters

Goal: find  in image

Method 3: Normalized cross-correlation

mean template

mean image patch

$$h[m,n]=\frac{\sum_{k,l}(g[k,l]-\bar{g})(f[m+k,n+l]-\bar{f}_{m,n})}{\left(\sum_{k,l}(g[k,l]-\bar{g})^2\sum_{k,l}(f[m+k,n+l]-\bar{f}_{m,n})^2\right)^{0.5}}$$

Slide by Derek Hoiem

Matching with filters

Goal: find  in image

Method 3: Normalized cross-correlation

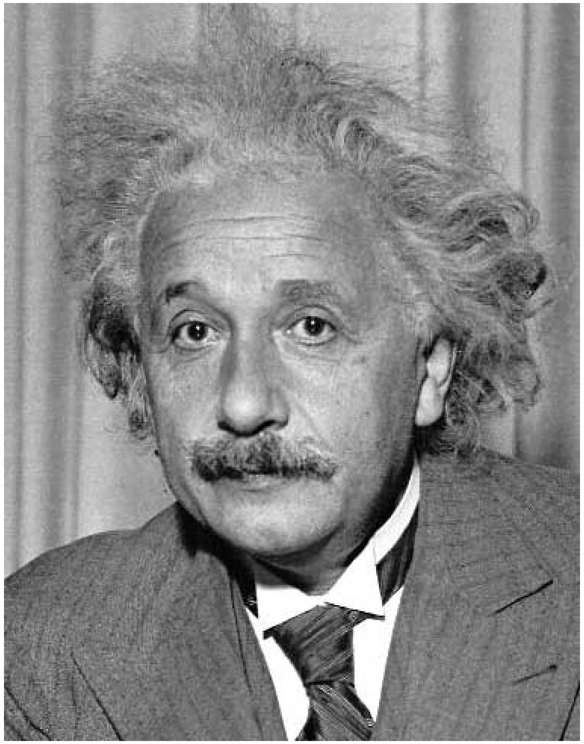

Input

Normalized X-Correlation

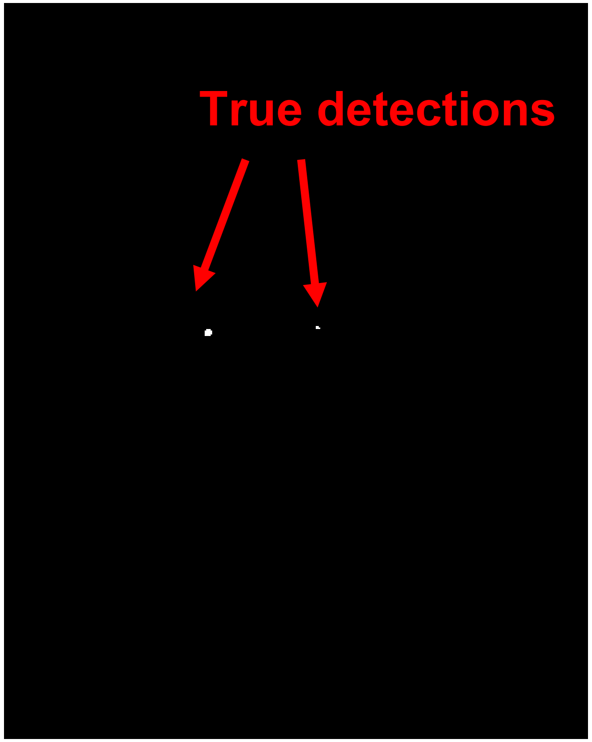

Thresholded Image

Matching with filters

Goal: find  in image

Method 3: Normalized cross-correlation

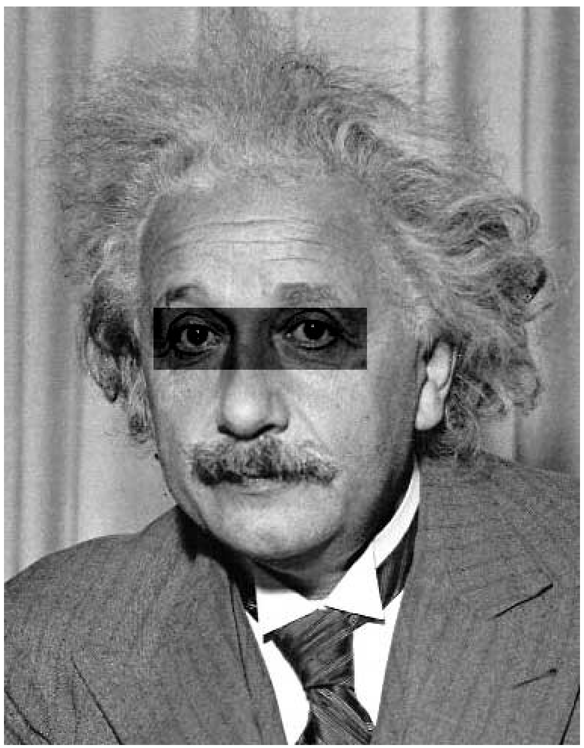

Input

Normalized X-Correlation

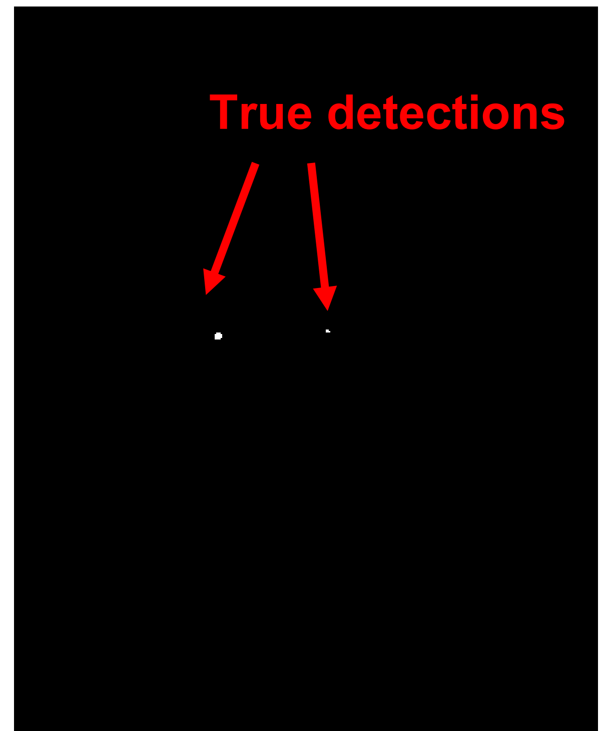

Thresholded Image

Q: What is the best method to use?

A: Depends

Zero-mean filter: fastest but not a great matcher

SSD: next fastest, sensitive to overall intensity

Normalized cross-correlation: slowest, invariant to local average intensity and contrast

Slide by Derek Hoiem

Image half-sizing

This image is too big to fit on the screen. How can we reduce it?

How to generate a half-sized version?