

Principles of Safe Autonomy

Lecture 4: Perception, edge detection




Sayan Mitra

slides from Svetlana Lazebnik



GEM platform

Autonomy pipeline



Sensing

Physics-based models of camera, LIDAR, RADAR, GPS, etc.

Perception

Programs for object detection, lane tracking, scene understanding, etc.

Decisions and planning

Programs and multi-agent models of pedestrians, cars, etc.

Control

Dynamical models of engine, powertrain, steering, tires, etc.





Perception

Programs for object detection, lane tracking, scene understanding, etc.



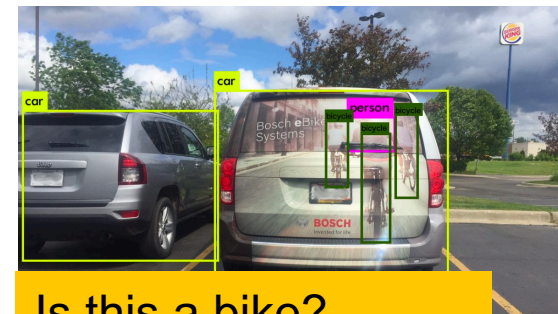


Outline

- Linear filtering
- **Edge detection**
- Assumptions in simple safety model (read)



Perception: EM to objects

Problem: Process electromagnetic radiation from the environment to construct a *model* of the world, so that the constructed model is close to the real world

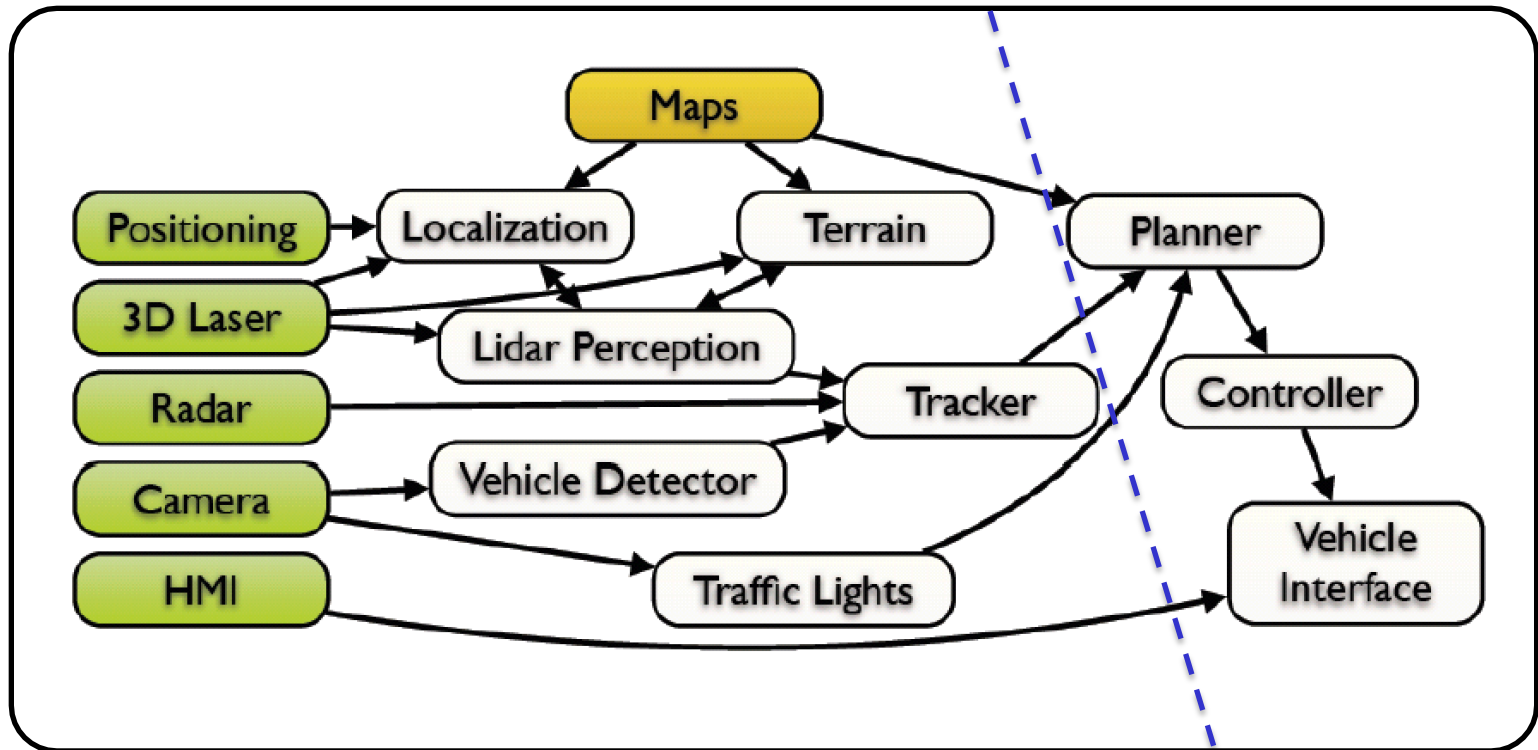


Challenging for computers: millions of years of evolution

Ill-defined problem: impossibility of defining meaning “car”, “bicycle”, etc.



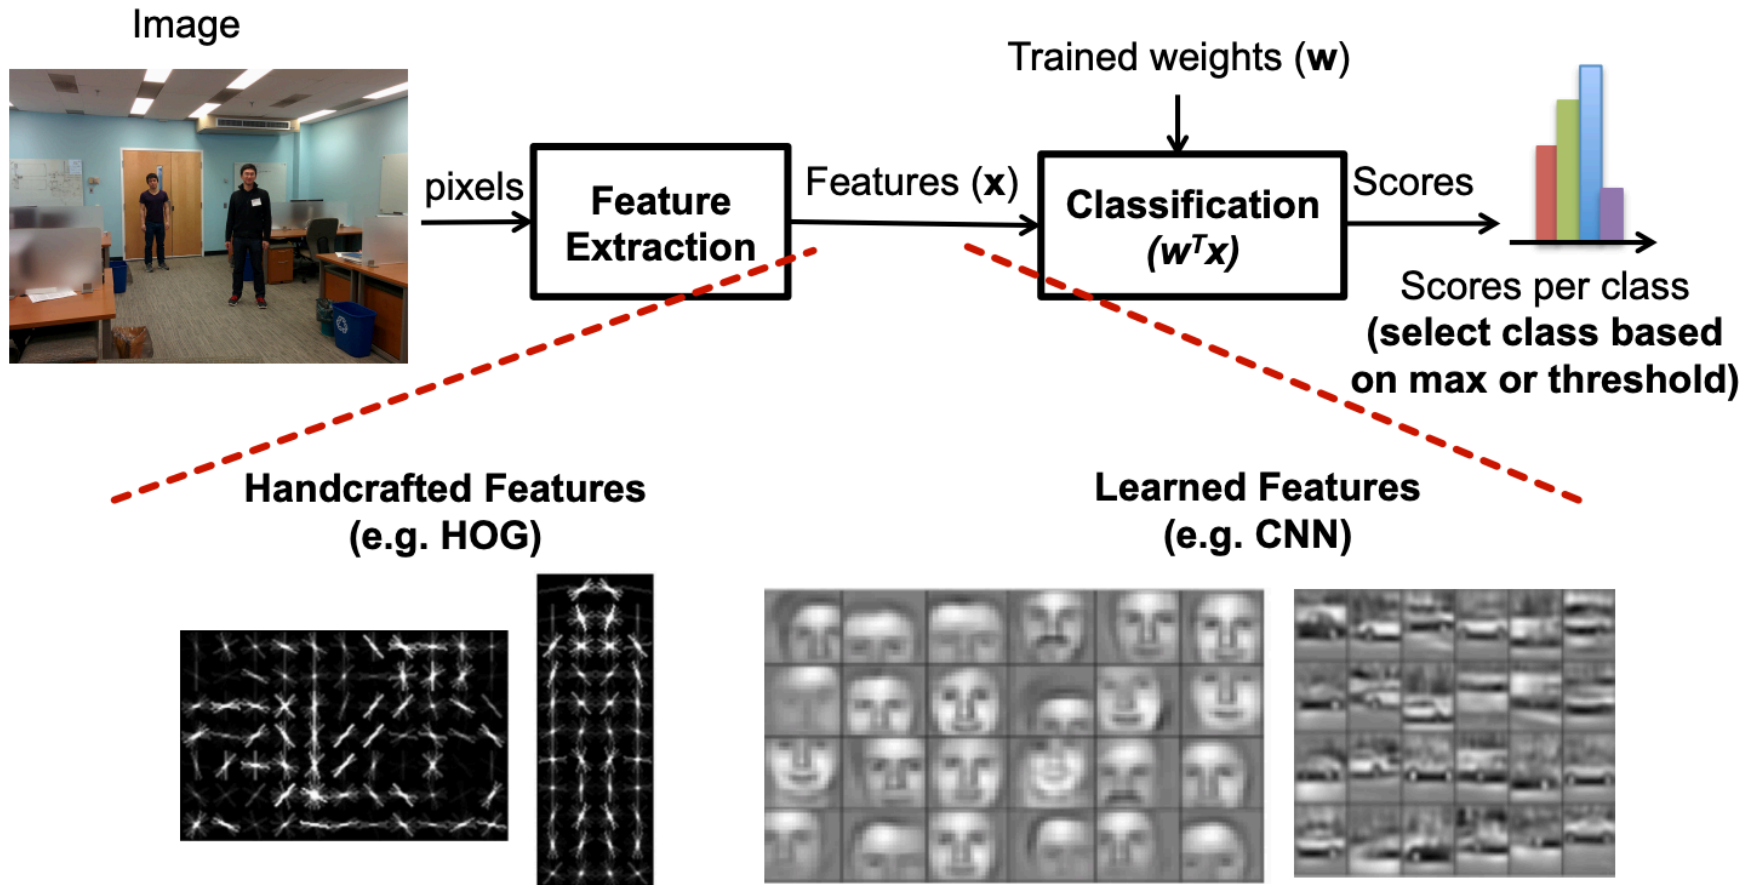
A practical perception pipeline in an AV



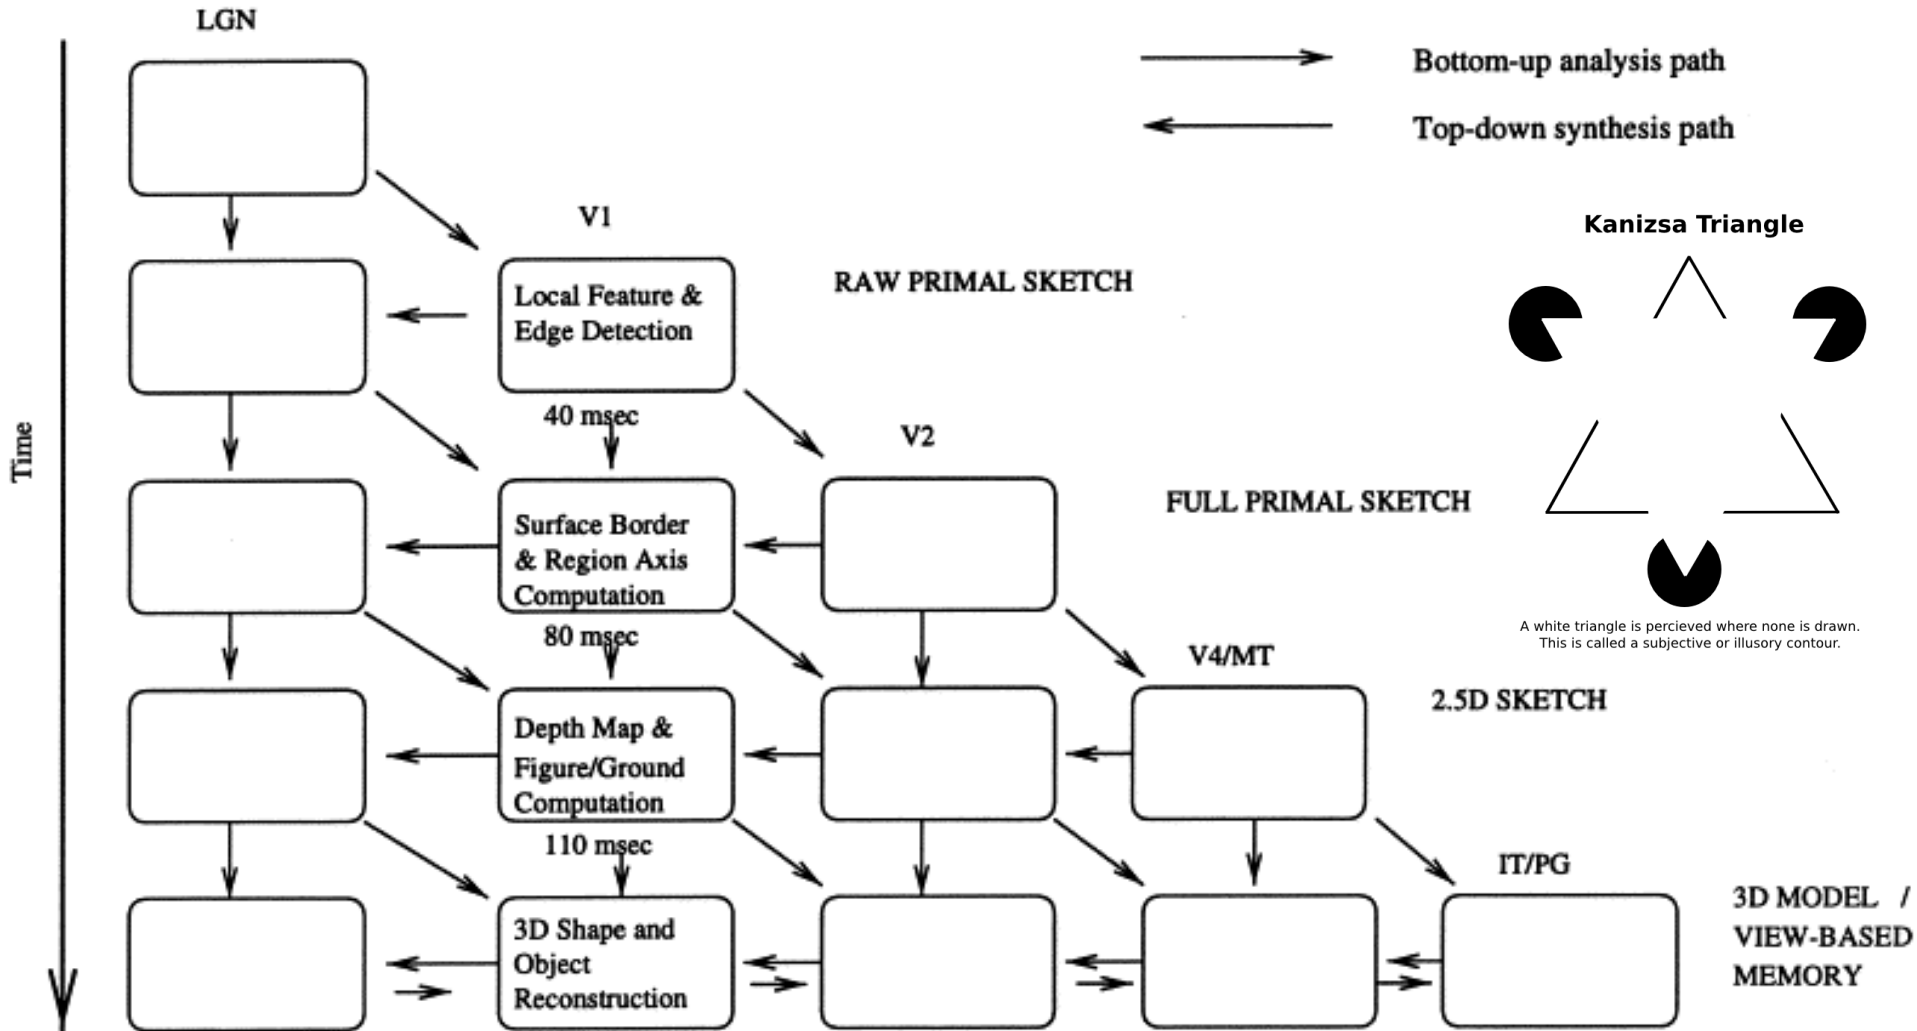
This architecture from a slide from M. James of Toyota Research Institute, North America



Recognition pipeline (more details in next lecture)



Natural vision “pipeline”



The role of the primary visual cortex in higher level vision, Tai Sing Lee, David Mumford, Richard Romero, Victor A.F.Lamme



Edge detection




[Winter in Kraków photographed by Marcin Ryczek](#)




Edge detection

- **Goal:** Identify sudden changes (discontinuities) in an image
- Intuitively, edges carry most of the semantic and shape information from the image
 - E.g., Lanes, traffic signs, cars



Edge detection

- An edge is a place of rapid change in the image intensity function



Derivatives with convolution

For 2D function $f(x,y)$, the partial derivative w.r.t x is:

$$\frac{\partial f(x, y)}{\partial x} = \lim_{\varepsilon \rightarrow 0} \frac{f(x + \varepsilon, y) - f(x, y)}{\varepsilon}$$

For discrete data, we can approximate using finite differences:


$$\frac{\partial f(x, y)}{\partial x} \approx \frac{f(x + 1, y) - f(x, y)}{1}$$

To implement the above as convolution, what would be the associated filter?

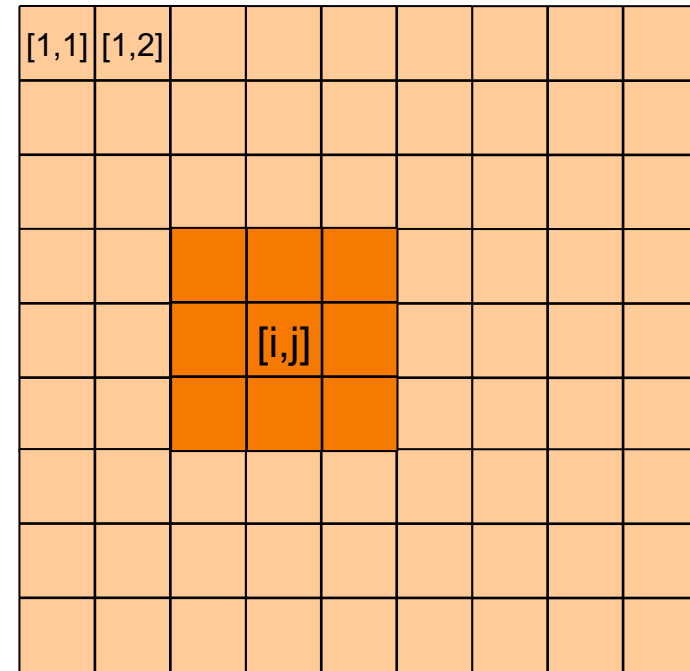


Convolution

convolution
mask $g[,]$



image[i,j]




Output or convolved image

$$f = g * \text{img}$$

$$f[i,j] = \underline{-1 \cdot \text{img}[i,j-1]} + \underline{1 \cdot \text{img}[i,j]}$$



Partial derivatives of an image



$$\frac{\partial f(x, y)}{\partial x}$$

-1	1
----	---

$$\frac{\partial f(x, y)}{\partial y}$$

-1
1

1
-1

Which shows changes with respect to x?



Finite difference filters

Other approximations of derivative filters exist:

Prewitt: $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \\ -1 & -1 & -1 \end{bmatrix}$

Sobel: $M_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$

Roberts: $M_x = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$; $M_y = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$






Kahoot!



Image gradient

The gradient of an image: $\nabla f = \left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right]$



The gradient points in the direction of most rapid increase in intensity

- How does this direction relate to the direction of the edge?

The gradient direction is given by $\theta = \tan^{-1} \left(\frac{\partial f}{\partial y} / \frac{\partial f}{\partial x} \right)$



The edge strength is given by the gradient magnitude (norm)

$$\|\nabla f\| = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$



Effects of noise


Consider a single row or column of the image



Where is the edge?



Solution: smooth first




- To find edges, look for peaks in $\frac{d}{dx}(f * g)$




Derivative theorem of convolution



- Differentiation is convolution, and convolution is associative: $\frac{d}{dx}(f * g) = f * \frac{d}{dx}g$
- This saves us one operation:




Derivative of Gaussian filters



x-direction




y-direction




Which one finds horizontal/vertical edges?





Derivative of Gaussian filters



x-direction



y-direction



Are these filters separable?



Recall: Separability of the Gaussian filter


$$\begin{aligned} G_{\sigma}(x, y) &= \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \\ &= \left(\frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{x^2}{2\sigma^2}\right) \right) \left(\frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{y^2}{2\sigma^2}\right) \right) \end{aligned}$$

The 2D Gaussian can be expressed as the product of two functions, one a function of x and the other a function of y

In this case, the two functions are the (identical) 1D Gaussian



Scale of Gaussian derivative filter



1 pixel

3 pixels

7 pixels


Smoothed derivative removes noise, but blurs edge
Also finds edges at different “scales”



Review: Smoothing vs. derivative filters



Smoothing filters

- Gaussian: remove “high-frequency” components; “low-pass” filter
- Can the values of a smoothing filter be negative?
- What should the values sum to?
 - **One**: constant regions are not affected by the filter



Derivative filters

- Derivatives of Gaussian
- Can the values of a derivative filter be negative?
- What should the values sum to?
 - **Zero**: no response in constant regions




Building an edge detector

Original Image



original image

Edge Image



final output

norm of the gradient

$$\|\nabla f\| = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$



Building an edge detector



How to turn these thick regions of the gradient into curves?

Thresholded norm of the gradient



Non-maximum suppression




Another problem:
pixels along this
edge didn't survive
thresholding



Hysteresis thresholding

Use a high threshold to start edge curves, and a low threshold to continue them.



Hysteresis thresholding



original image



**high threshold
(strong edges)**



**low threshold
(weak edges)**



hysteresis threshold

Source: L. Fei-Fei



Recap: Canny edge detector

1. Compute x and y gradient images
2. Find magnitude and orientation of gradient
3. **Non-maximum suppression:**
 - Thin wide “ridges” down to single pixel width
4. **Linking and thresholding (hysteresis):**
 - Define two thresholds: low, high
 - Use the high threshold to start edge curves and the low threshold to continue them


opencv: `canny(image, th1, th2)`

J. Canny, [A Computational Approach To Edge Detection](#), IEEE Trans. Pattern Analysis and Machine Intelligence, 8:679-714, 1986.

A Computational Approach to Edge Detection

JOHN CANNY, MEMBER, IEEE

Abstract—This paper describes a computational approach to edge detection. The success of the approach depends on the definition of a *comprehensive* set of goals for the computation of edge points. These goals must be precise enough to delimit the desired behavior of the detector while making minimal assumptions about the form of the solution. We define detection and localization criteria for a class of edges, and present mathematical forms for these criteria as functionals on the operator impulse response. A third criterion is then added to ensure that the detector has only one response to a single edge. We use the criteria in numerical optimization to derive detectors for several common image features, including step edges. On specializing the analysis to step edges, we find that there is a natural uncertainty principle between detection and localization performance, which are the two main goals. With this principle we derive a single operator shape which is optimal at any scale. The optimal detector has a simple approximate implementation in which edges are marked at maxima in gradient magnitude of a Gaussian-smoothed image. We extend this simple detector using operators of several widths to cope with different signal-to-noise ratios in the image. We present a general method, called feature synthesis, for the fine-to-course integration of information from operators at different scales. Finally we show that step edge detector performance improves considerably as the operator point spread function is extended along the edge. This detection scheme uses several elongated operators at each point, and the directional operator outputs are integrated with the gradient maximum detector.



Summary

- Convolution as translation invariant linear operations on signals and images
- Definition of convolution and its properties (associativity, commutativity, etc.)
- Artifacts of of hard-edge kernels
- Gaussian kernel, its definition and properties (separability)
- Median filter, sharpening
- Derivatives as convolution (Sobel, etc.)



Outline

- Linear filtering
- Edge detection
- Assumptions in simple safety model (read)



Sharpening

What does blurring take away?



-



=



Let's add it back:



+



=




Unsharp mask filter

$$f + \alpha(f - f * g) = (1 + \alpha)f - \alpha f * g = f * ((1 + \alpha)e - g)$$

image


blurred image

unit impulse
(identity)




unit impulse

−



Gaussian

≈



Laplacian of Gaussian

