Lecture 3
Day 1.

Statistical Edge Detection.

Edge Detection. Canny Model.

Step Edge Model:

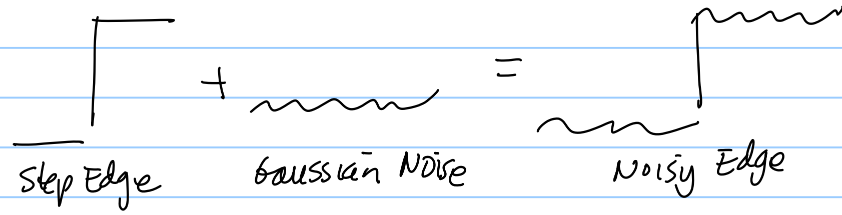

Solution:
- Apply a smoothed derivative filter.
- Threshold response
- Non-Maximal Suppression

E.G.
One-Dimension $\left| \frac{d}{dx} G_{\sigma} * I(x) \right| > T$, Threshold T.
Gaussian σ

If multiple responses above threshold
Select the maximum response and remove the rest.

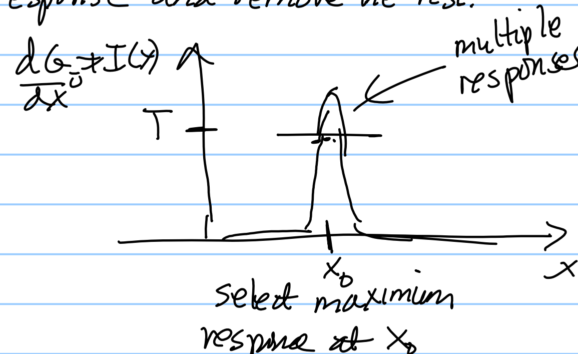

In Two-Dimensions
Take Derivatives in many directions — eg.
Select maximum response.

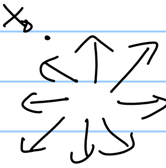

Note: this method assumes a model of an edge.
It does not try to learn an edge model from data.

Statistical Edge Detection:

Get dataset of images $\{ I^{\mu}(x) : \mu \in \Lambda \}$,
Pixels are labelled $\omega^{\mu}(x) = 1$, pixel x in μ^{th} image is an edge
" " " " " is not an edge.

Learn conditional distributions:
$$P(\underline{F} * I(x) \mid \omega(x) = 1)$$
$$P(\underline{F} * I(x) \mid \omega(x) = 0)$$
$\underline{F} = (F_1, \ldots, F_m)$ derivative filters

Example: Filter $|\nabla I| = \sqrt{\left(\frac{\partial I}{\partial x}\right)^2 + \left(\frac{\partial I}{\partial y}\right)^2}$

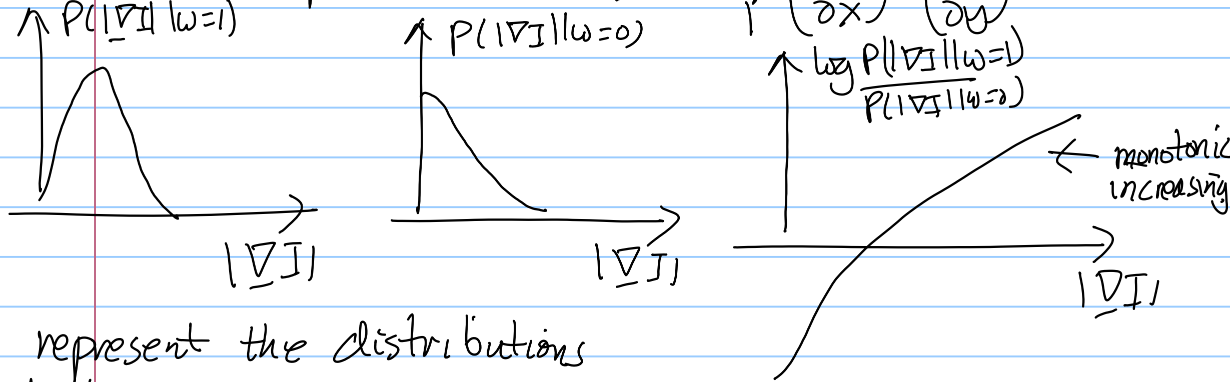

represent the distributions by histograms.

To detect edges: Bayes decision

$$P(\omega \mid |\nabla I|) = \frac{P(|\nabla I| \mid \omega) P(\omega)}{P(|\nabla I|)}$$

Label a pixel as edge $\omega = 1$, if $P(\omega=1 \mid |\nabla I|) > P(\omega=0 \mid |\nabla I|)$

$\omega = 0$, otherwise.

This requires knowing $P(\omega)$.

Estimate from data $\quad P(\omega=1) \approx 0.08$

$P(\omega=0) \approx 0.92$

Because $\log \frac{P(|\nabla I| \mid \omega=1)}{P(|\nabla I| \mid \omega=0)}$ is monotonic in $|\nabla I|$

the Bayes decision is equivalent to thresholding $|\nabla I|$ which is similar to Canny.

But, suppose we apply several different filters.

$$P(|\nabla I|, |\nabla G_{\sigma_1} * I|, |\nabla G_{\sigma_2} * I| \mid \omega).$$

Learn these joint distributions from training examples. Then apply Bayes.

This does much better than Canny when evaluated on difficult datasets.

Statistical edge detection can combine edge cues from multiple filters in an optimal manner.

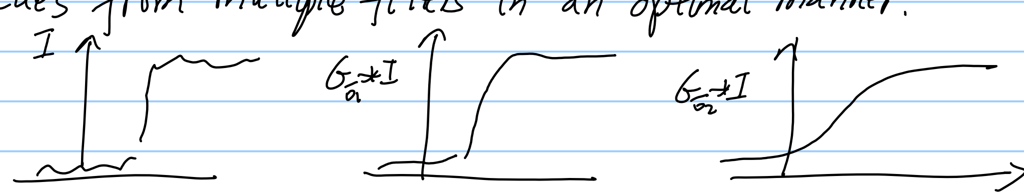

Statistical edge detection can combine information from multiple scales, or from multiple orientations.

<u>Problem</u> → if we represent $P(\underline{F}|\omega)$ by histograms, then we can need $O(k^m)$ data where k is no. of histogram bins and m is the number of filters. Requires too much data, and becomes impractical for large m. $(m \gg 10)$.

<u>Note</u>: less data is needed if we know a parameterized form for the distributions (typically polynomial no. of parameters). But we do not know the parameterized form of the distribution.

<u>Learning</u>: Learn on Training Dataset
evaluate on Test Dataset
Cross-Validation.

Bayes is a special case of Bayes Decision Theory. This includes a loss function, so that we can penalize the expected loss.

<u>Note</u>: that the Statistical approach described above is a generative model on filters response, but is not a generative model on images.

Alternative related approach (Martin, Fowlkes, Malik)
Extract features $\phi_1(I(x)), \phi_2(I(x)), .., \phi_m(I(x))$
Learn distribution $$P(\omega | \phi_1 .. \phi_m) = \frac{e^{\omega \sum_{i=1}^{m} \lambda_i \phi_i(I(x))}}{e^{\sum_{i=1}^{m} \lambda_i \phi_i(I(x))} + e^{-\sum_{i=1}^{m} \lambda_i \phi_i(I(x))}}$$
with $\omega \in \{\pm 1\}$,
Regresim model - learn parameters $\{\lambda_i\}$.