Edge Detection and Simple Semantic Segmentation

A.L. Yuille (JHU)

Edge Detection and (simple) Semantic Segmentation as examples of low-level vision.

Why do we care about edges?

- A Line Drawing is simple representation of the image (far fewer bits than a normal image).
- They are often sufficient to interpret the entire image (caveats).
- The Line Drawing is composed of edges – an “edge map”.

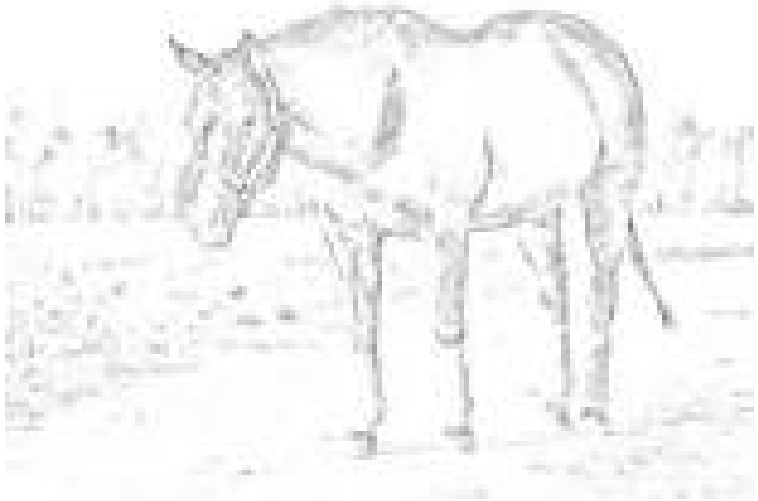

What Does the Line Drawing Represent?

- They represent the boundaries of objects.
- They represent interior edges of objects.
- They represent texture edges.

- Mooney images – a caveat: edges alone are not sufficient. Black and white helps.

But how can we find edges?

- Edge detection applies local operations (filtering) to images in order to detect edges. Local – low-level vision.
- Typically look at a local image patch – 3x3 pixels, or 8 x8 pixels – and decide if there is an edge or not.
- This is binary classification task – machine learning/statistics.

$$\begin{bmatrix} 120 & 118 & 110 & 115 & 116 & 120 \\ 115 & 21 & 20 & 16 & 19 & 121 \\ 112 & 19 & 17 & 18 & 20 & 117 \\ 119 & 118 & 121 & 117 & 116 & 112 \end{bmatrix}$$

(a) Intensity matrix of an image

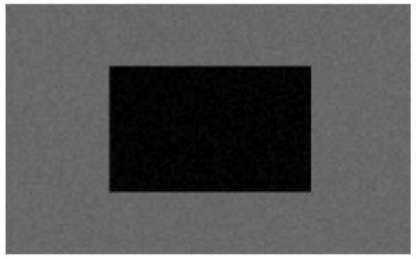

(b) Image of the matrix (150×100)

Typically edges occur at places where the intensity gradient changes.

- Idealized edges: Images I(x). Derivative of Image dI(x)/dx

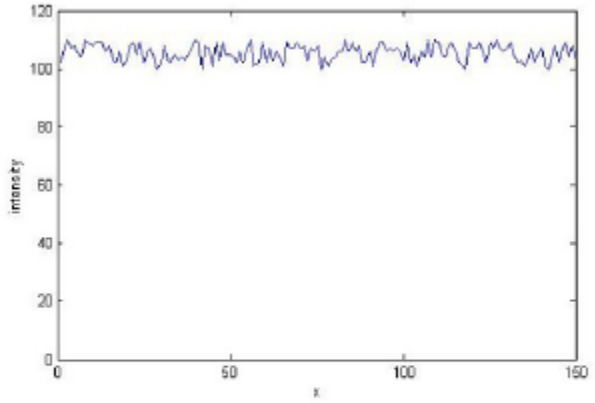

(a) Intensity values (y=5)

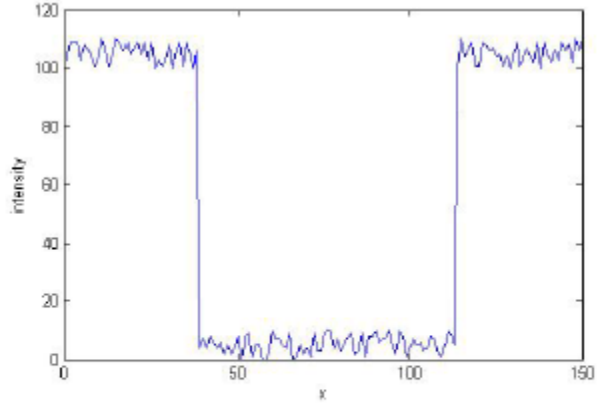

(b) Intensity values (y=50)

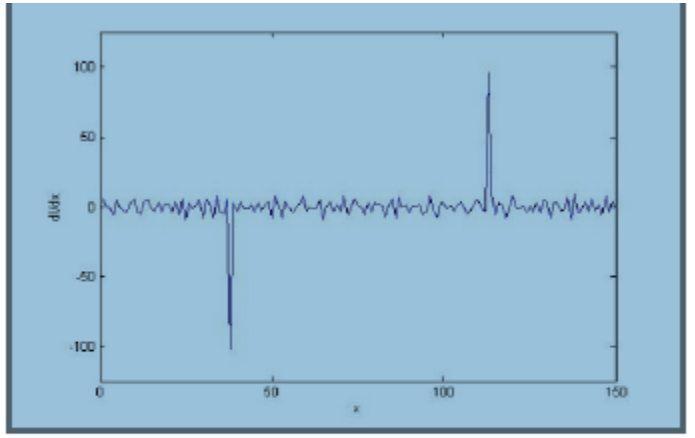

(c) Derivative Filter response (y=50)

But Images are much more complex than this simple picture

- The intensity gradient can be very small at the boundaries of an object. This is surprising to humans since we use context to interpret images. This context can be non-local or high-level (e.g., recognition), later in this course.
- Fox. Local edges – threshold gradient. Steeple Image (ambiguous)


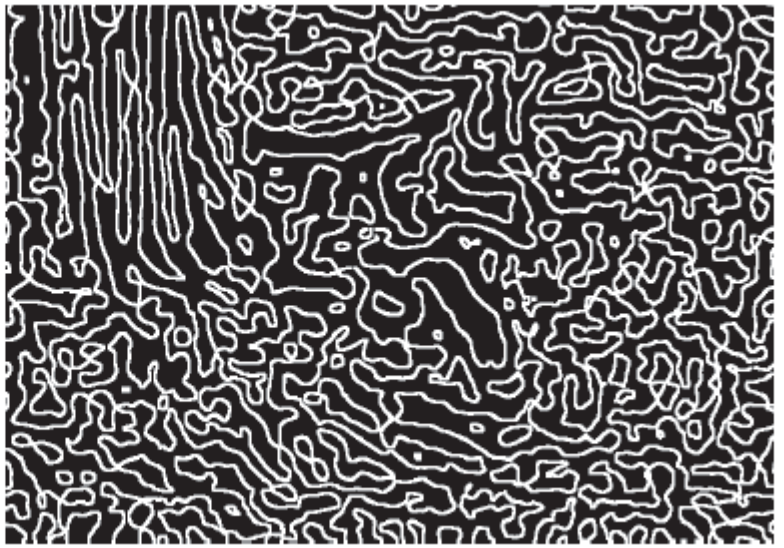


How to Perform Edge Detection.

- Classic approach – define an ideal model of an edge and obtain an optimal edge detector (Canny 1986).

- Treat Edge Detection as a statistical machine learning problem. Requires a dataset with ground truth (positions of edge specified). (Konishi, Yuille, Coughlan, Zhu 1999).

Sowerby Dataset: Groundtruth.

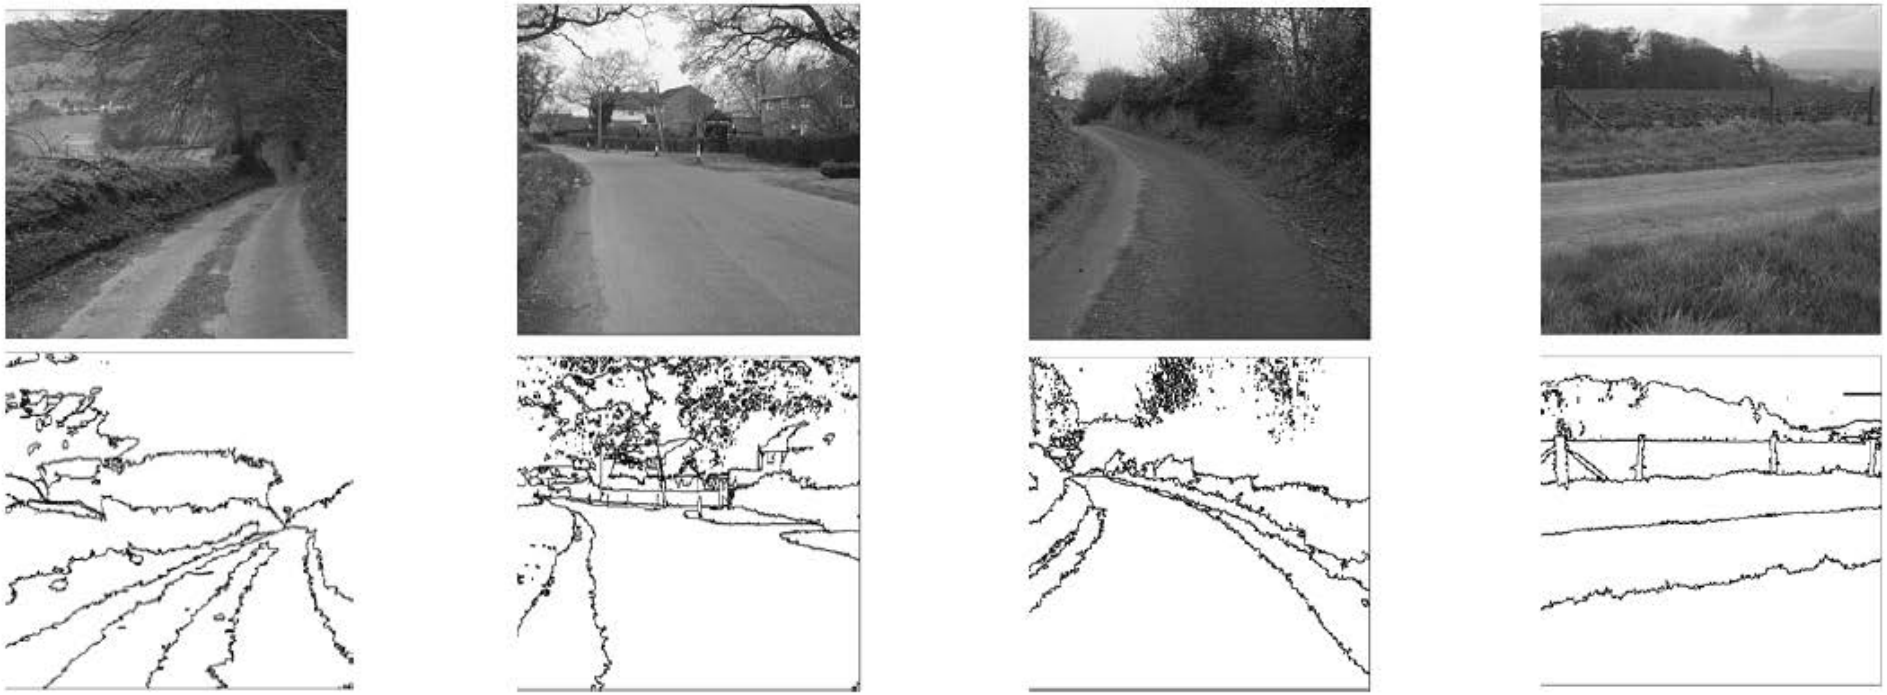

Figure 5. Upper Panels: the data images. Lower Panels: the groundtruth edge maps

Statistical Edge Detection

- Konishi, Yuille, Coughlan, Zhu. CVPR. 1999.
- Learn conditional probability distributions of image features conditioned on whether there is an edge or not (on-edge, off-edge).
- Use log-likelihood ratio test to detect edges.
- 1. Probability distributions. 2. Image. 3. Ground Truth. 4. Output.

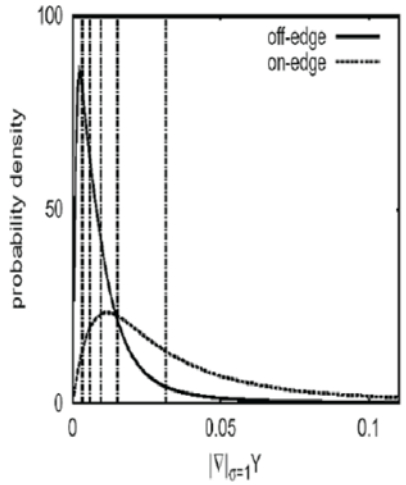

Different Datasets.

- Hard (Sowerby) and Easy (S. Florida) Datasets.

Soweby consists of outdoor images.
Much texture and vegetation.

South Florida consist of indoor images.
Very little texture.

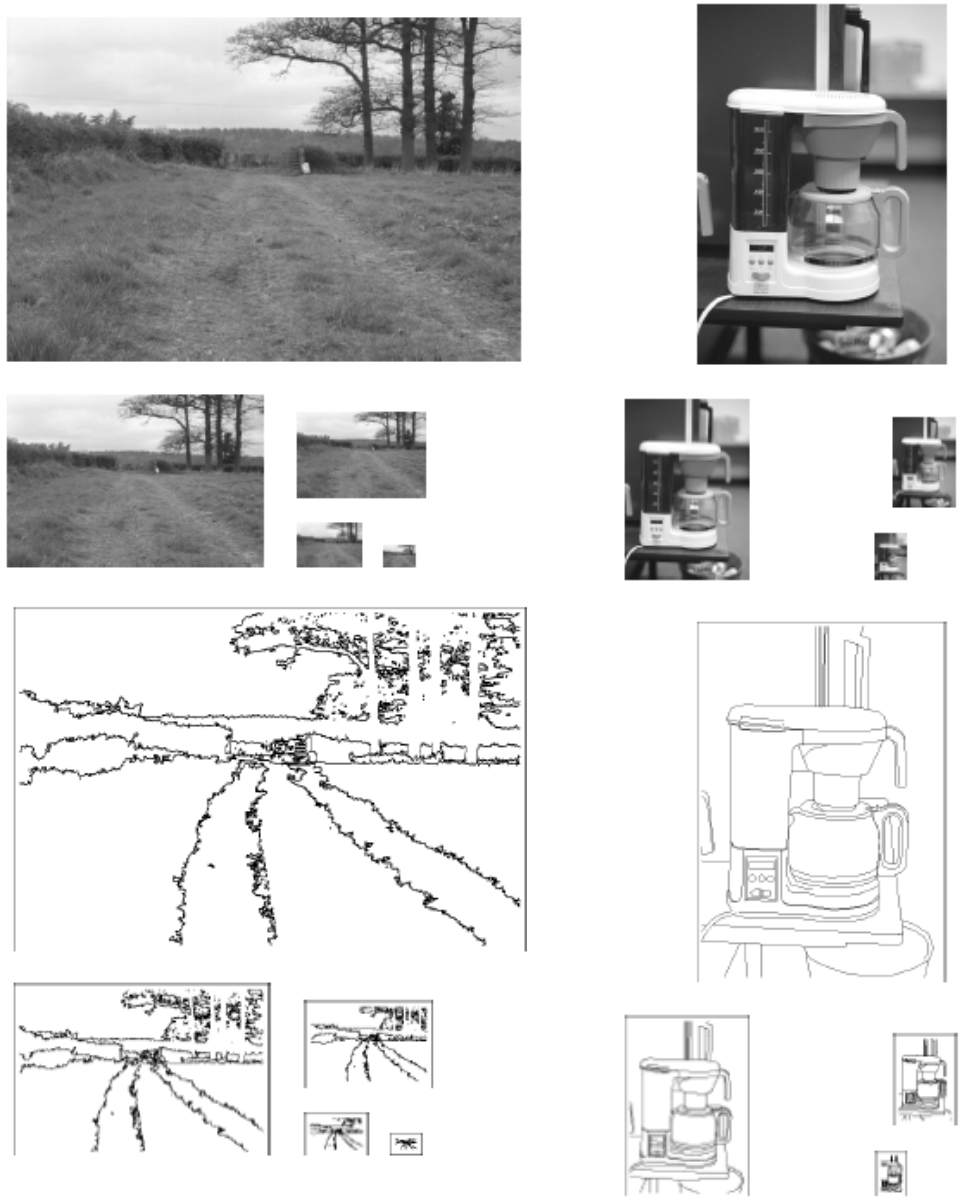

Semantic Segmentation: What else can you classify locally?

- What about “sky”, “vegetation”, “water”?
- These are roughly homogeneous. E.g., all parts of sky are similar.
- Colour and Texture.

Label Regions Sowerby

- Konishi and Yuille. CVPR. 2018.
- Output examples.

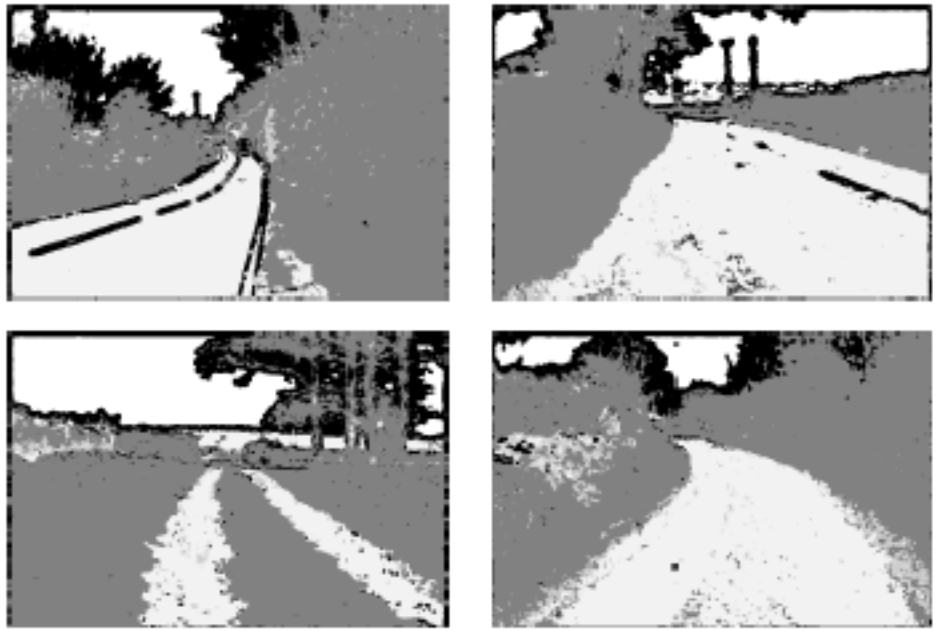

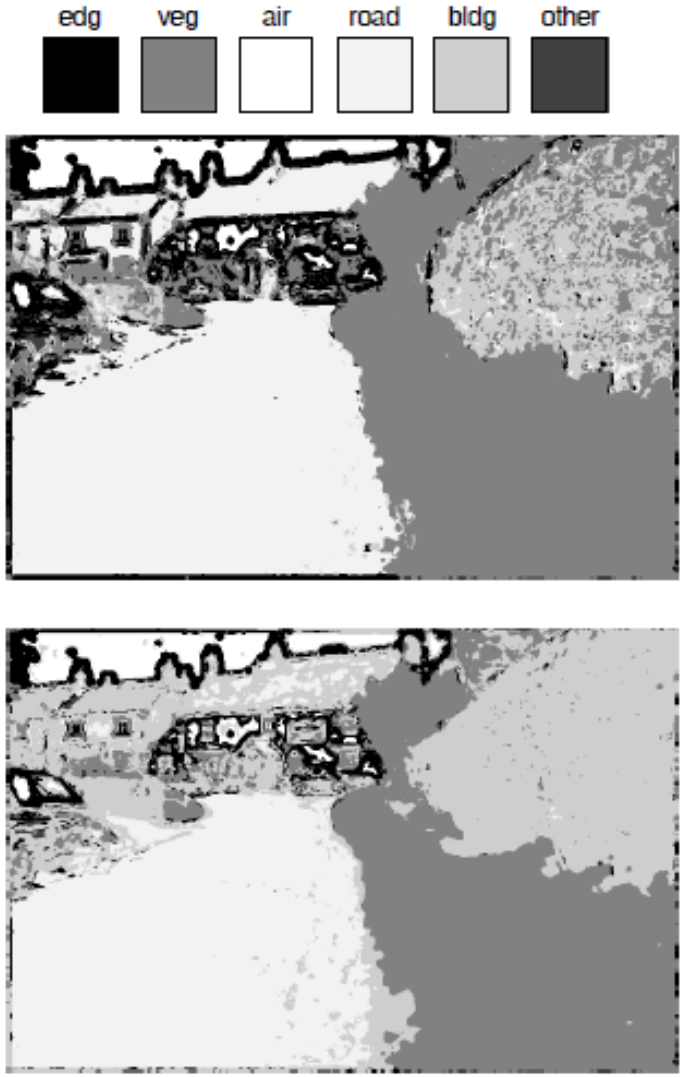

Label Regions: San Francisco