Lecture 3: Region Based Vision

Dr Carole Twining

Thursday 18th March

1:00pm – 1:50pm

Segmenting an Image

Assigning labels to pixels (cat, ball, floor)

- Point processing:
  - colour or grayscale values, thresholding
- Neighbourhood Processing:
  - Regions of similar colours or textures
- Edge information (next lecture)
- Prior information: (model-based vision)
  - I know what I *expect* a cat to look like

 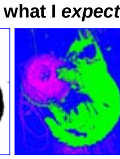 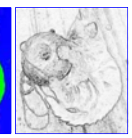

Overview

- Automatic threshold detection
  - Earlier, we did by inspection/guessing
- Multi-Spectral segmentation
  - satellite and medical image data
- Split and Merge
  - Hierarchical, region-based approach
- Relaxation labelling
  - probabilistic, learning approach
- Segmentation as optimisation

Automatic Threshold Selection

Automatic Thresholding: Optimal

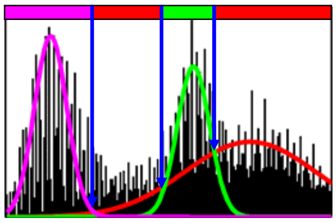

- Assume scene mixture of substances, each with normal/gaussian distribution of possible image values
- Minimum error in probabilistic terms
- But sum of gaussians not easy to find
- Doesn't always fit actual distribution

Automatic Thresholding: Otsu's Method

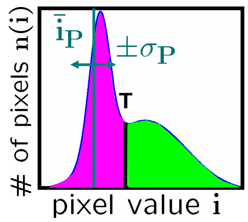

Mean across purples:

$$\bar{i}_P = \frac{1}{N_P}\sum_{i=0}^{T} i \times n(i)$$

Variance for purples:

$$\sigma_P^2 = \frac{1}{N_P}\sum_{i=0}^{T} n(i)\left[i - \bar{i}_P\right]^2$$

Choose T to minimize:

$$N_P\sigma_P^2 + N_G\sigma_G^2$$

- Extend to multiple classes

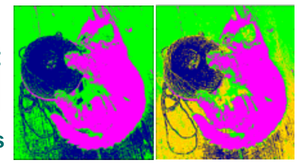

Automatic Thresholding: Max Entropy

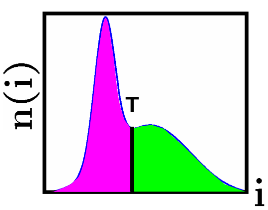

For two sub-populations:

$$p_P(i) = \frac{n(i)}{N_P},\ i < T,$$

$$p_G(i) = \frac{n(i)}{N_G},\ i \geq T.$$

$$\boxed{\text{Entropy: } -\sum p \ln p}$$

Two Entropies:

$$H_P = -\sum_{i<T} p_P(i)\ln p_P(i)\ \&\ H_G = -\sum_{i\geq T} p_G(i)\ln p_G(i)$$

$$\boxed{\text{Minimise: } H_G + H_P \text{ to find } T.}$$

- Makes two sub-populations as peaky as possible

Automatic Thresholding: Example

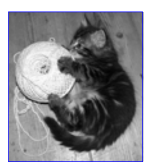

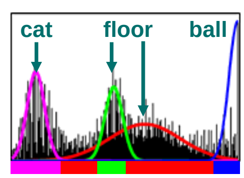

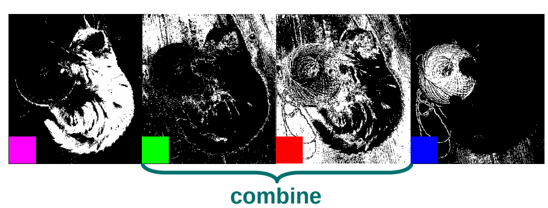

Automatic Thresholding: Summary

- Geometric shape of histogram (bumps, curves etc)
  - Algorithm or just by inspection
- Statistics of sub-populations
  - Otsu & variance
  - Entropy methods
- Model-based methods:
  - Mixture of gaussians
- And so on. > 40 methods surveyed in literature
- Fundamental limit on effectiveness:
  - Never give great result if distributions overlap
- Whatever method, need further processing

Multi-Spectral Segmentation

Multi-spectral Segmentation

- Multiple measurements at each pixel:
  - Satellite remote imaging, various wavebands
  - MR imaging, various imaging sequences
  - Colour (RGB channels, HSB etc)
- Scattergram of pixels in vector space:

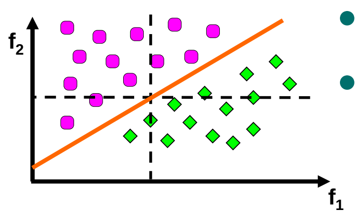

- Can't separate using single measurement
- Can using multiple

Multi-Spectral Segmentation:Example

Spectral Bands

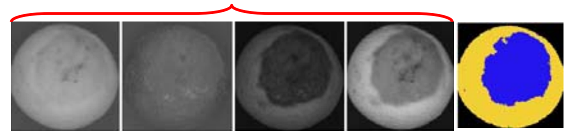

Over-ripe Orange

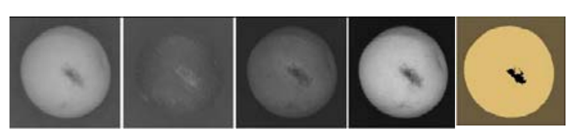

Scratched Orange

Multispectral Image Segmentation by Energy Minimization for Fruit Quality Estimation:
Martínez-Usó, Pla, and García-Sevilla, **Pattern Recognition and Image Analysis , 2005**

Split and Merge

Split and Merge

- Obvious approaches to segmentation:
  - Start from small regions and stitch them together
  - Start from large regions and split them

  Combine
- Start with large regions , split non-uniform regions
  - e.g. variance σ^2 > threshold
- Merge similar adjacent regions
  - e.g. combined variance σ^2 < threshold

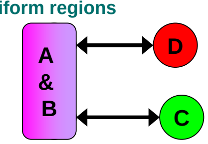

- Region adjacency graph
  - housekeeping for adjacency as regions become irregular
  - regions are nodes, adjacency relations arcs
  - simple update rules during splitting and merging

Split and Merge

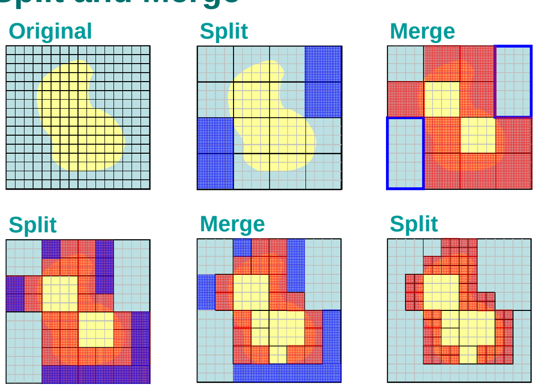

Split and Merge: Example

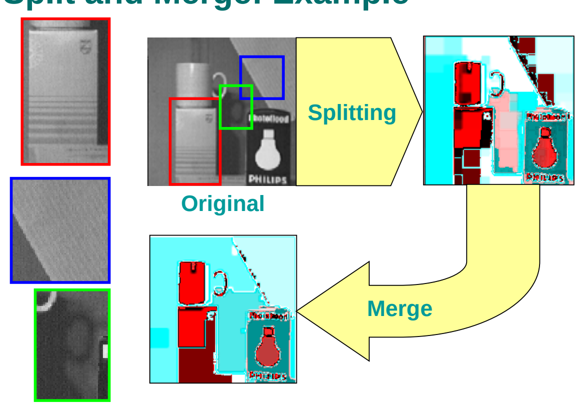

Relaxation Labelling

Aside: Conditional Probability

probability of A given that B is the case

$P(A \mid B)$

- P(pet) = $\frac{(\square+\square)}{\boxed{ALL}}$ etc
- P(pet | mammal) = $\frac{\square}{(\square+\square)}$
- P(mammal | pet) = $\frac{\square}{(\square+\square)}$
- Bayes Theorem:

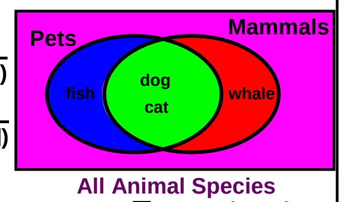

$$\frac{\square}{(\square+\square)} \times \frac{(\square+\square)}{\boxed{ALL}} = \frac{\square}{(\square+\square)} \times \frac{(\square+\square)}{\boxed{ALL}}$$

- P(pet | mammal)P(mammal) = P(mammal | pet)P(pet)

Relaxation Labelling:

- Image histogram, object/background

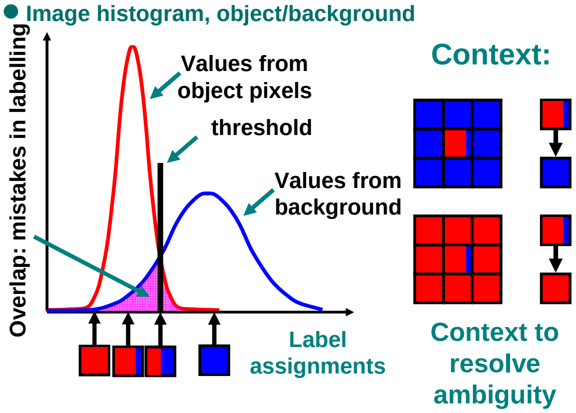

Relaxation Labelling

- Evidence for a label at a pixel:
  - Measurements at that pixel (e.g., pixel value)
  - Context for that pixel (i.e., what neighbours are doing)
- Iterative approach, labelling evolves
- Soft-assignment of labels:

  Possible labels: $\{\mathbf{l}_\mu : \mu = \mathbf{1}, \ldots \mathbf{n}\}$

  $\mathrm{P_i}(\mu)$: Probability that pixel i has label l_μ.

  $\sum_\mu \mathrm{P_i}(\mu) \equiv \mathbf{1}$. normalised probability.
- Soft-assignment allows you to consider all possibilities
- Let context act to find stable solution

Relaxation Labelling

- **Compatibility:**

  Pixels i and j, labels μ and ν:

no effect	$c_{i,j}(\mu,\nu) = 0$	If not neighbours
support	$c_{i,j}(\mu,\mu) = \alpha$	Neighbours & same label
oppose	$c_{i,j}(\mu,\nu) = -\alpha$ if $\mu \neq \nu$	Neighbours & different label

- **Contextual support for label μ at pixel i :**

$$s_i(\mu) \propto \boxed{\sum_{j \neq i}} \boxed{\sum_{\nu}} \boxed{c_{i,j}(\mu,\nu)} \boxed{P_j(\nu)}$$

look at all other pixels

all possible labels & how strong

degree of compatibility

Relaxation Labelling:

- **Update soft labelling given context:**

$$P_i(\mu) \Leftarrow A_i P_i(\mu)(1 + s_i(\mu))$$

A_i chosen so sums to 1 at i.

- **The more support, more likely the label**
- **Iterate**

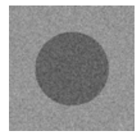

Noisy Image

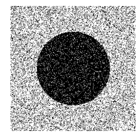

Threshold labelling

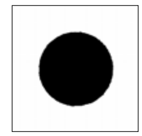

After iterating

Relaxation Labeling:

- **Value of α alters final result**

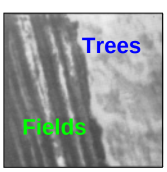

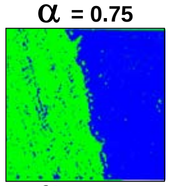

α = 0.75

Initialisation

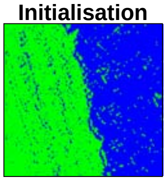

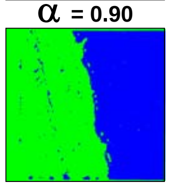

α = 0.90

Segmentation as Optimisation

Slide 25

Segmentation as Optimisation

Image: $\mathcal{I}$, value at pixel i: $\mathcal{I}(i)$

Label Image: L, label at pixel i: $L(i)$

Label configuration in neighbourhood of i: $l(i)$

- Maximise probability of labelling given image:

$$P(L|\mathcal{I}) = \prod_i \underbrace{P(L(i)|\mathcal{I}(i))}_{\text{label at i given value at i}} \underbrace{P(L(i)|l(i))}_{\text{label at i given labels in neighbourhood of i}}$$

- Re-write by taking logs, minimise cost function:

$$C(L,\mathcal{I}) = \sum_i [\underbrace{-\log P(L(i)|\mathcal{I}(i))}_{\text{label-data match}} \underbrace{-\log P(L(i)|l(i))}_{\text{label consistency}}]$$

- How to find the appropriate form for the two terms.
- How to find the optimum.

Slide 26

Segmentation as Optimisation

$P(L(i)|l(i))$ label consistency

$P(L(i)|\mathcal{I}(i))$ label-data match

- Exact form depends on type of data
- Histogram gives: $p(\mathcal{I}(i))$
- Model of histogram $P(L(i)|\mathcal{I}(i))$

(e.g., sum of gaussians, relaxation case)

Learning approach:

- Explicit training data (i.e., similar labelled images)
- Unsupervised, from image itself (e.g., histogram model):

Expectation/Maximization

- •Given labels, construct model
- •Given model, update labels
- •Repeat

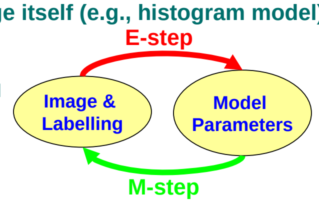

Slide 27

Segmentation as Optimisation

- General case:

Cost function: $C(L,\mathcal{I}) = \boxed{\text{label-data match term}} + \boxed{\text{label consistency}}$

- High-dimensional search space, local minima
- Analogy to statistical mechanics
  - crystalline solid finding minimum energy state
  - stochastic optimisation
  - simulated annealing
- Search:
  - Downhill
  - Allow slight uphill

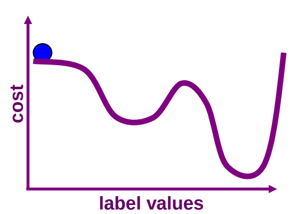

Slide 28

Segmentation as Optimisation