Image Segmentation

CMSC 828J – David Jacobs

What the course is about

- Image Segmentation
  - Dividing images up into meaningful chunks.
  - This is central in vision.

  Other processes seem to require segmentation, such as recognition, stereo, motion.

  A lot of interesting recent progress.

  Still, very difficult.

Who is course intended for?

- Meant to provide an overview of lower-level vision.
  - Maybe 1/3 of vision (no reconstruction or recognition).
  - But much of the course is basic material every vision researcher should know.
- Some topics relevant to AI or Numerical Analysis.
- It’s pretty mathematical.

What we’ll talk about today.

- What I think are the key issues.
- How this course is organized pedagogically.
- Course details (eg., grades).

Segmentation is hard to define

- Divide world up into objects?
  - What is an object?
  - Objects involve semantics. Often no obvious connection between objects and images.
- But we want segmentation to be low-level. Divide world into uniform surfaces?
  - Better defined but still tough
- Divide world into regions with uniform image properties.
  - This is what we do
  - But it only approximates what we want.

My take on segmentation work:

Segmentation involves combining local, overlapping cues and global information into a global solution.

This is why image segmentation is different from unsupervised learning. There is a neighborhood structure, and local constraints.

Solving these kinds of problems involves some kind of *propagation* or *diffusion* of information, across the image. Most of the class will discuss methods that do this.

Example 1: Anisotropic diffusion. To explain this, first I have to explain diffusion. Take a 1D slice through an image, and diffuse pixel intensities. Now do it anisotropically. You wind up with intensities that provide the segmentation.

Example 2: Diffusion of probabilities. Consider a contour. We can guess the position and orientation of a missing part of it. We are confident near the contour, but we get less confident as we move further away. We can think of this as our having an implicit pdf about the contour position. This is diffusing in space/orientation. Solving this distribution is equivalent to solving a diffusion equation.

Example 3: Linear relaxation labeling with a bunch of contour fragments. Each contour passes confidence about its being figure to neighbors. We solve this by taking the leading eigenvalue of a matrix. This is also the steady state solution to a random walk/diffusion.

Example 4: Belief propagation. We can do the same thing, but separate information based on direction. If they don't form a closed contour, this computes the probability each contour is figure.

Example 5: Markov Random Fields. This is the explicit

formulation of overlapping local constraints. It is exactly what we want, but unfortunately it is also intractable to solve. But approximations exist.

Example 6: Graph cuts. Here we formulate a global optimization that respects local constraints, and solve it approximately with some algorithms. It turns out this is also a solution to a kind of random walk problem.

Example7: Graph algorithms, such as shortest path. We can use these to find the best boundary of an object. This gives us global optimum, but depends on having local constraints, so we can use dynamic programming. In fact, finding the shortest path is a lot like solving a diffusion problem, except we use max instead of sum.

Example 8: Algebraic multigrid. We create a multiscale representation by combining together pieces of image that best follow local constraints. This is a way of building global solution with local constraints. It's also in some ways equivalent to graph cuts and to anisotropic diffusion.

Example 9: Texture. This also raises the question of scale. Two neighboring regions are likely from the same object if they have the same texture. But texture must be computed at some, unknown scale. This is really more of just a paradox we must cope with, not an example.

Is this really the right way to do segmentation? Is it really all about overlapping local constraints? I don't think so, but this is all we know right now.

What is missing:

1. Non-local constraints. Symmetry, for example. Or similarity to familiar objects. Unfortunately, we won't talk much about this because we don't understand it very well.
2. 3D. Objects and surfaces are 3D things projected into 2D. For example, real surfaces have smooth shading, but not constant shading. We'd really like to divide images into regions of uniform albedo, not intensity. But we don't know much about this either.
3. Non-local completion
4. What is the right shape prior?

One side of this course is to teach all these approaches to segmentation.

The other is to teach fundamental tools of Low-level vision.

1. Image Processing – Fourier transforms, convolutions and wavelets
2. Diffusion processes
3. Learning techniques – E-M, belief propagation, Markov processes and MRFs
4. Texture
5. Optical Flow

Course Pedagogy

- Segmentation is not solved, and most current work probably won't last.
  - So course is divided into two parts. First is stuff of fundamental importance to vision and segmentation. This is mostly math, and I'll mostly lecture on it. Second are current research papers. We must understand these to see how to make progress. But we must look critically at these. So we must discuss them.

A tour of the syllabus

How this might change

- Discussion topics are up for grabs.
  - Happy to consider suggestions for other topics.

Requirements (1)

- Prepare for discussions.
- Paper reviews.
  - On classes where papers are discussed, you must turn in a 1 page review of one paper before class.
  - One paragraph summarizing main points.
    - Doing this well is enough for a B.
  - One paragraph critiquing ideas, suggesting new directions.
    - Do this well for an A.
  - Discussion. Join in a discussion of the topic.
  - 10% of grade.

Requirements (2)

- Paper expert
  - Students will be assigned to become experts on two papers.
    - **Short** summary.
    - Answer our questions.
    - Any Insights culled from extra reading.
    - Analysis and topics for discussion.
  - 15% of grade. You may be graded as a group.
  - Each student will do this once.

Requirements (3)

- Quizzes, Midterm and Final
  - Oral quiz, early in the semester.
  - Midterm (may be one week take home).
  - Final during exam period.
  - Will cover materials in lectures.
  - 50% of grade.
  - This is mostly math. You cannot possibly get an A without understanding the math in the class.

Requirement (4)

- Implementations 25% Choose 1
  - Problem sets
    - Implementations of algorithms we study.
    - Not certain, but might be: Perona-Malik diffusion, normalized cut and E-M.
  - Implementation project and write-up. Can be groups.
    - Discuss with me first.
    - Implement existing and/or novel algorithms
    - Experiment on real-world data.

Your Background

- Calculus, linear algebra, probability is essential.
- Math that makes you really learn these topics is important.
- Other math very helpful: functional analysis (Fourier transforms), wavelets, geometry, stochastic processes, optimization.
- Knowledge of vision may help a little.