Boundary Segmentation

CSC 249/449
Sensory Motor Systems
Randal C. Nelson

Boundary Segmentation (Edge Linking)

Topics:

- Introduction
- Binary contour following
- Local boundary adjustment
- Hough transform techniques
- Graph search methods
- Boundary growing methods
- Higher level and hierarchical grouping

Introduction

A basic issue is the degree and form of prior information available. There are a number of possibilities.

- Prior segmentation into regions
- Approximate location of boundary
- Location of points on boundary
- Functional form of boundary
- Expected density of boundary points
- No specific information available

Boundaries can be represented in several forms. Different algorithms operate on different representations.

- A chain of adjacent pixels (under some connectivity)
- A chain of "crack" edges between pixels
- A chain of line segments with discrete pixel endpoints
- Parametric forms. e.g. lines, quadrics, or splines

Binary contour following

Prior knowledge consists of a segmentation of the scene, either into a binary image or labeled regions. The task is a straightforward, algorithmic conversion between related representations.

- Basic binary algorithm to produce pixel chain is to start with a pair of points ordered clockwise on the boundary (current and previous point), run a clockwise search around neighbors of the current point starting at the previous point to find the next point, and iterate.
- In the above, a decision must me made whether to generate an inner or an outer boundary. This could get confusing if multiple labels are present. Similar algorithms can generate crack edges, which don't have the inner/outer distinction.
- Main complications involve discrete topology and how ambiguous connections (e.g. diagonally adjacent) are handled. This can get tricky if multiple labels are involved. However, if you don't particularly care what the results are in borderline cases, and don't need mathematically consistent topology in the output, you can usually sweep these issues under the rug.

Local Boundary Adjustment

Prior knowledge consists of an initial idea of where the boundary lies. We want to adjust it to match the image. Situations where this comes up include tracking and interactive image processing.

- Divide and conquer: Search perpendicular to the midpoint of an initial segment for closest strong gradient, split segment into two, and iterate.
- Local adjustment: Make local search perpendicular to initial boundary approximation to find high-gradient image points that provide an updated position. Can include gradient direction in match as well for sharper result.
- Snakes: The Cadillac of boundary adjustment methods. Minmize an energy consisting of an internal "bending" term and an external "image match" term. Forms and tradeoff factors are up to the implementor. A classic internal energy measure is to integrate the square of the local curvature. A classic image energy measure is to integrate the gradient magnitude under the curve, possibly modulated by the angle between the curve and the gradient direction.

Hough Transform Methods

Prior knowledge consists of information to the effect that the boundaries of interest are described by a functional form with only a few degrees of freedom (e.g. we know that the boundaries form straight lines or circles).

- The Hough tranform approach consists of accumulating votes in a parameter space that spans the relevant functional form, and then detecting concentrations of votes (peaks) that indicate the presence of particular instances of the functional form.
- Classically, each edge point (and more generally, each feature) in an image generates one or more votes in the parameter space.
- Results are often equivalent to what could be obtained by correlation with a set of masks, but with added efficiency arising from from sparsity of feature points, and effective use of additional constraints to reduce the number of votes that each feature generates.
- Technique generalizes to detection of arbitrary template shapes, where the parameters are used to represent the position (and possibly scaling and rotation) of the template in the image.
- Issues: Scaling of cells that accumulate votes, peak detection algorithms, reducing the number of hypotheses each feature point votes for. Parameter spaces of more than two or three dimensions are typically too large to represent explicitly, and require additional strategies such as projection and hypothesis lists.

Graph Search Methods

Prior knowledge frequently consists of information about endpoint location. Also sometimes used without specific prior information.

- The graph search approach attempts to cast the contour segmentation problem as a minimum cost path problem that can be addressed with classic search algorithms.
- Classical formulation assigns a linkage cost between all pairs of neighboring pixels; generally these are a function of the intensity and direction of the local gradient. The algorithm is then set up to find the lowest cost path between specified points (often supplied by an interactive user). This formulation can be solved efficiently with a dynamic programming approach
- Costs can also be formulated on paths so as to penalize abrupt changes in angle, and to allow gaps to be bridged; however such introduction of less local terms may make efficient implementation harder to obtain. In some cases, heuristic search procedures may be necessary.
- Graph search can be adapted to segment contours in the absence of pre-defined start and end points by combining a method for identifying starting locations with some sort of termination condition. The main difficulty with these approaches is overcoming the "short sightedness" inherent in the local linking criteria. This leads to contours terminated by short branches and minor gaps, and can connect pieces that do not appear plausibly grouped at a high level because some local glitch bridged them.

Boundary Growing Methods

Used primarily for segmentation without specific prior information, though these can also be used in supervised mode.

- Basic idea is to find locations that a boundary is likely to pass through (e.g. high gradient points) and attempt to extend a boundary in both directions from such a point until it closes on itself, fades out, terminates at a previously mapped segment, encounters a corner, etc... The process can then be iterated, until no good, unincorporated starting points remain.
- Simplest method is to adapt the binary boundary following algorithm: Start at a high-gradient point, and link on the highest gradient neighbor subject to some constraints to prevent worming about in a high gradient neighborhood (e.g. look only within 45 degrees of a current direction of growth). This does not actually work very well - it is too short sighted.
- Performance can be improved by incorporating longer range information, for example from correlation with sizable masks representing locally smooth lines or curves, or responding to termination situations. Measures of contour smoothness can also be incorporated. Directional as well as intensity information is generally valuable here.
- Ultimately, a growth energy function can be defined with internal and external energy components, that (hopefully) correctly balances the issues of smoothness versus attraction to strong local gradient, and termination evidence against the goal of producing extended structures. Decisions on whether and how to extend a contour can be made on the basis of this energy function, which thus incorporates information from an extended neighborhood.

Higher-level and Hierarchical Grouping

Sometimes it is desired to associate boundary fragments that are interrupted by occluding obstacles, confusing local structure, or broken at corners, in order to produce longer boundaries or closed contours.

- Processes that attempt to accomplish this sort of higher-level association go under the general term of "perceptual grouping". This is a subject in itself, and covers other sorts of relationships than smooth extension of curves and contour closure. These include parallel structure, symmetry, similarity, familiarity, and other so-called "Gestalt laws".
- Basic techniques include direct search for consistently linked chains of linear features, relaxation methods for finding globally consistent arrangements, and various statistical clustering methods.
- Some of these methods can be used hierarchically to produce extended contours from relatively small fragments, starting with few-pixel groups or even individual pixels. This constitutes another group of methods for base-level boundary segmentation.