Constellation Models for Sketch Recognition

D. Sharon and M. van de Panne
University of British Columbia

Published in SIGGRAPH 2006

Presented by Subha Kannan
Nov 19, 2007

Contents

- Introduction
- Related Work
- Constellation Model
- Labeling Likelihood
- Maximum Likelihood Search
- Training Example Sketches
- Results & Discussion
- Review

Introduction

- Goal: Recognition of user drawings – objects, diagrams or maps
- Similar to goals of image interpretation in computer vision
- Apply constellation or “pictorial” model to represent sketch objects

Introduction (contd.)

- Maximum-likelihood labeling for an unlabelled sketch
- Searches through possible label assignments using a multi-pass branch and bound algorithm

Introduction (contd.)

- Output of the algorithm: Set of labels assigned to strokes!
- Useful for a variety of applications
  - To construct parameterized 3D Models
  - Instance models in a 2D or 3D scene
  - Partially interpret a large sketched diagram

Introduction (contd.)

- Assumptions:
  - Similar parts drawn with similar strokes
  - “Mandatory” parts have only instance in the sketch
  - “Optional” parts have multiple instances

Related Work

- Matching treated as Graph Isomorphism problem
  - Using a known model of connectivity
  - Weak at recognizing drawings with disjoint parts
- Probabilistic approach
  - Uses domain-specific libraries of 'Bayesian network fragments'
- Image based techniques

Approach

- Support for model definitions from set of drawn training examples
- Probabilistic framework
- Support for optional parts
- Constellation model – features for sketch recognition
- Efficient multi-stage search strategy

Constellation Model

- A visual model to capture individual and pair-wise features between strokes
- Designed to capture the structure of a particular class of object
- Based on :
  - local features : shape or size of a stroke
  - pairwise features: distances to known parts

Constellation Model (contd.)

Feature vector for local features (F):

1. The x-coordinate of the stroke's bounding box center
2. The y-coordinate of the stroke's bounding box center
3. The length of the bounding box diagonal
4. The angle of the bounding box diagonal

Constellation Model (contd.)

Feature vector for pairwise features (G):

1. Delta x between the strokes
2. Delta y between the strokes
3. The minimum distance between the endpoints of stroke a and any point on stroke b
4. The minimum distance between the endpoints of stroke b and any point on stroke a

Constellation Model (contd.)

- Interaction features are pairwise
  - model scales with number of strokes by $O(n^2)$
- This problem alleviated by allowing to choose the label as mandatory or optional

Constellation Model (contd.)

- Each element of the feature vector F and G have their mean and covariances computed for training data set
- Individual probabilistic model computed for label's features

Example Constellation Model

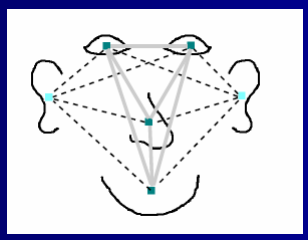

Sketch face
Mandatory – left eye, right-eye, mouth and nose
Optional – left-ear, right-ear

Labeling Likelihood

- For the entire sketch, a likelihood function estimates the probability of an entire labeling L.
- Function maximizes
  - probability of individual label
  - probability of interactions between each label

Labeling Likelihood (contd.)

Probability of a given labeling L = product of individual stroke labeling likelihoods * product of all labelled stroke pair likelihoods

$$P(L|\theta) = \prod_{i=1}^{N}\prod_{\ell=1}^{M} P(\mathcal{F}_i|\theta_\ell)^{\delta_{i\ell}} \prod_{j=1}^{m}\prod_{k=1}^{N} P(\mathcal{G}_{ik}|\theta_{\ell j})^{\delta_{kj}}$$

Probability of stroke i having label l

Probability of stroke i in relation to all the mandatory parts, as measured by pairwise feature vectors

Maximum Likelihood Search

- Tries to maximize the likelihood function defined
- Multi-pass search
  - First pass: labels only mandatory stroke objects
  - Second pass: linear search through optional labels for recognizing remaining unlabelled strokes
- All possible label assignments searched for using “branch-and-bound search tree”

Maximum Likelihood Search (contd.)

- Each node in the search tree represents a partial labeling of the sketch
- Depth of tree corresponds to number of labels applied
- Search advances by choosing the best assignment of mandatory labels

Maximum Likelihood Search (contd.)

- Approach for images with high number of strokes:
  - Multipass thresholding, Hard constraints
- Multipass Thresholding:
  - Bounds the branches of the tree before a full labeling is found.
  - If a new likelihood is worse than the bound, then branches associated with the likelihood's labeling are pruned.
  - If no complete solution is found with the bounding, the bound is loosened until a labeling is found

Training Examples Face Sketches

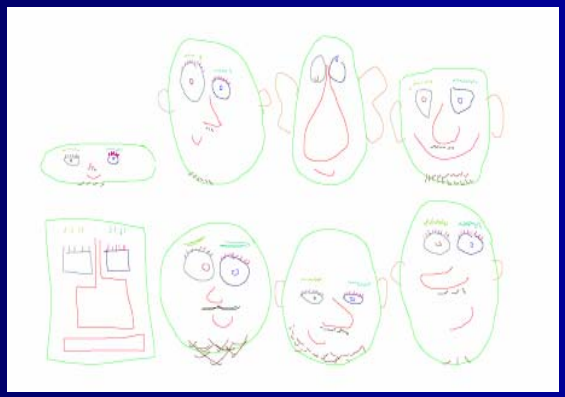

Mandatory : head, left-eye, right-eye, mouth, nose
Optional : left & right pupil, left & right ear, left & right eyebrow, left & right eye lash, moustache, beard

Recognized Face Sketches

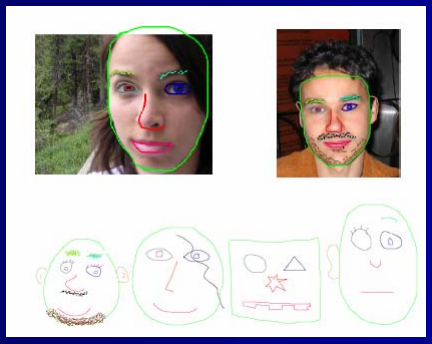

Training Examples Flower Sketches

Recognized Flower Sketches

Training Examples Sailboat Sketches

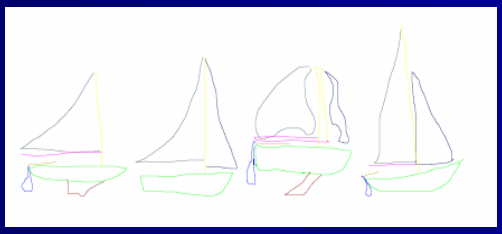

Recognized Sailboat Sketches

Results & Discussion

- Tested the method on 5 classes of objects

class	mandatory labels	optional labels
faces	5	10
flowers	2	5
sailboats	3	5
airplanes	3	4
characters	7	8

Results & Discussion (contd.)

- Recognition Time – 0.01 – 2.5 s
- Multi-pass thresholding/bounding greatly reduces time required to find a complete labeling

class	num strokes	with multipass (s)	without multipass (s)
face	103	1.242	> 9 hours
flower	54	0.46	0.98
sailboat	8	0.02	0.03
airplane	21	0.08	0.1
character	18	0.12	126.69

Results & Discussion (contd.)

- Failure Modes:

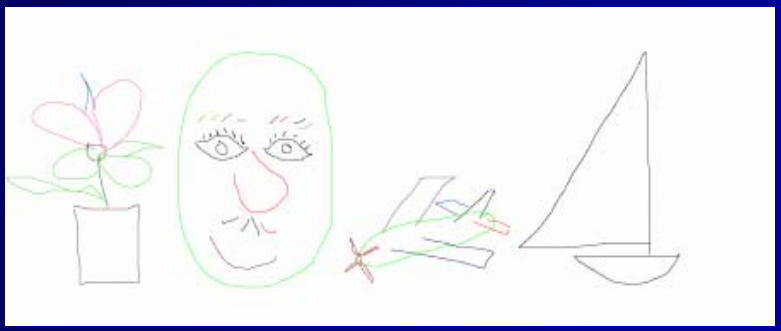

Results & Discussion (contd.)

- Reasons for failure:
  - Inability to find mandatory strokes due to hard constraints – Rare, lack of training data
  - Mislabeling mandatory strokes – leads to havoc with rest of the strokes – Unusual strokes like long, bushy hair on faces
  - Mislabeling optional strokes – If there are few mandatory strokes, such as the sailboats or flowers

Review

- What about accuracy of the results?
- Assumption that mandatory labels drawn with one stroke – might not always hold true!
- Only recognizes from pre-identified class of objects.
- Too many constraints – too few or too many mandatory/optional strokes!

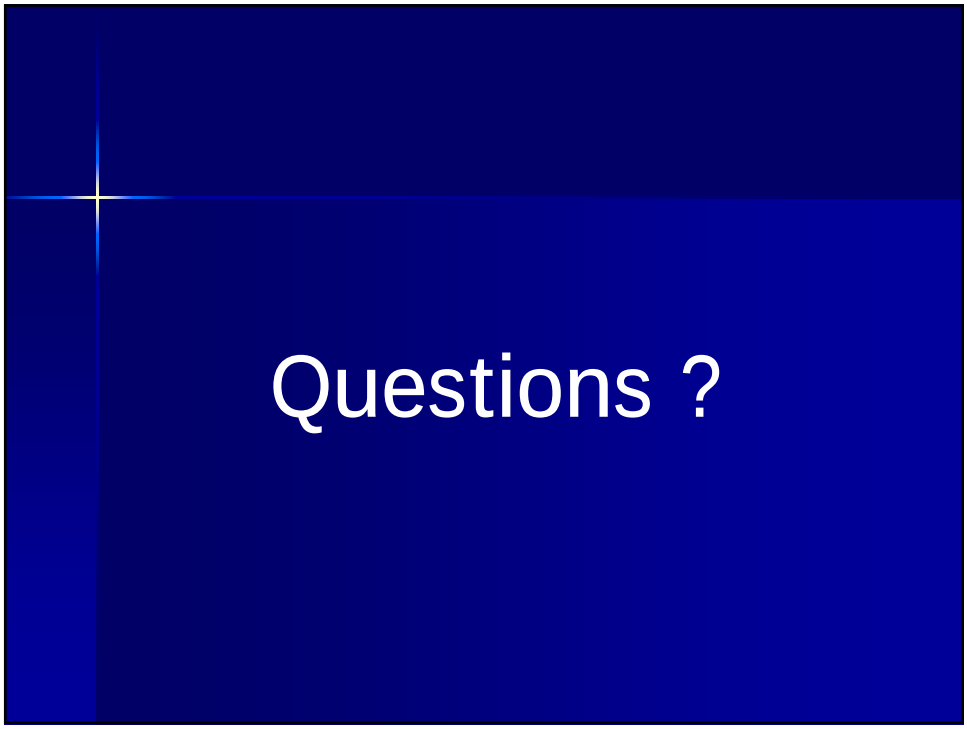
Questions ?