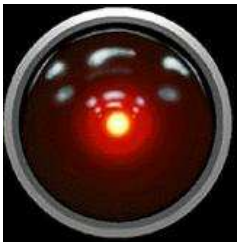

HCI/ComS 575X:
Computational Perception

Instructor: Alexander Stoytchev
http://www.cs.iastate.edu/~alex/classes/2007_Spring_575X/

Face Detection

February 26, 2007

Lecture Plan

- HW3: Due tonight
- Project Updates
- Face Detection: Neural Networks Approach
- Face Detection: Cascades
- OpenCV Demo

Face Detection v.s. Face Recognition

Henry A. Rowley, Shumeet Baluja and Takeo Kanade (1997).

``Rotation Invariant Neural Network-Based Face Detection,''
Carnegie Mellon Technical Report, CMU-CS-97-201.

Rotation Invariant Face Detection

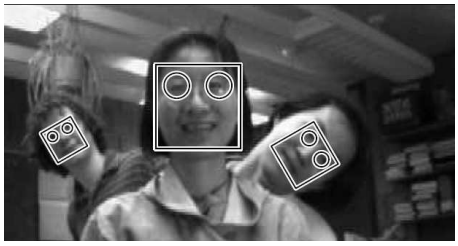

The Algorithm

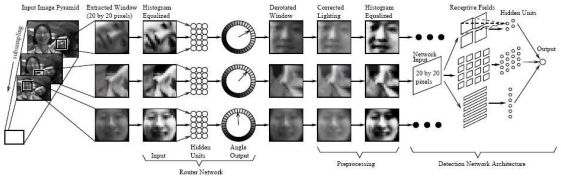

Router Network

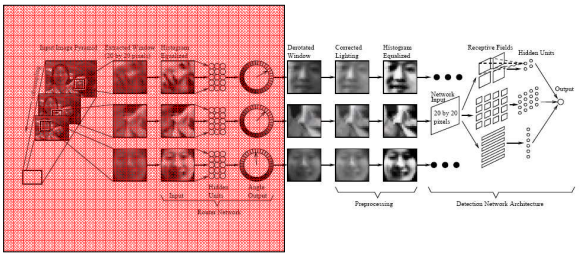

Router Network

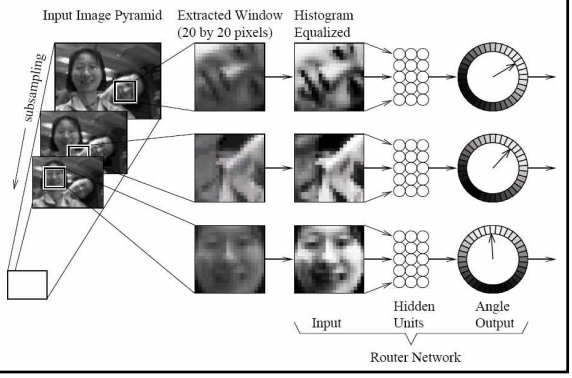

Training Examples

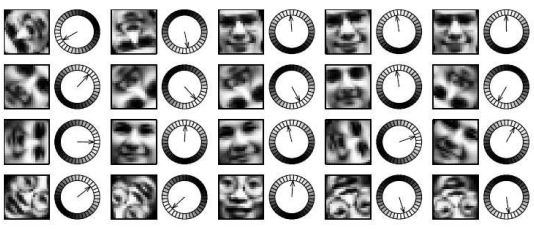

Detection Network

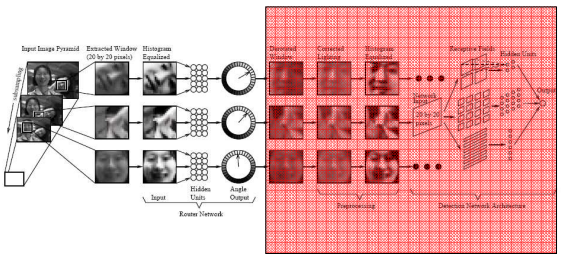

Detection Network

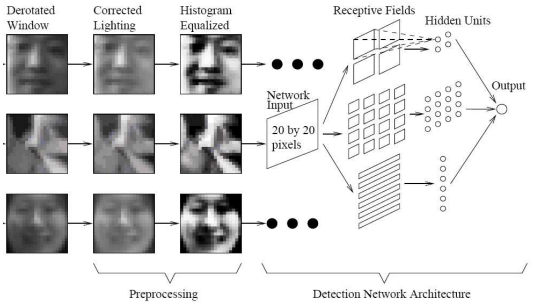

Testing on Images with no Faces

- All detections are automatically false positives
- They are added as negative examples in the training database

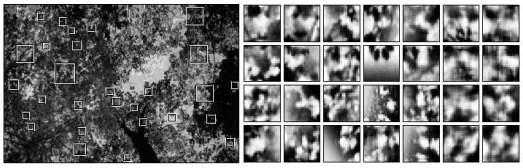

Results

System	Upright Test Set Detect %	Upright Test Set # False	Rotated Test Set Detect %	Rotated Test Set # False
Network 1	89.6%	4835	91.5%	2174
Network 2	87.5%	4111	90.6%	1842
Net 1 → Postproc	85.7%	2024	89.2%	854
Net 2 → Postproc	84.1%	1728	87.0%	745
Postproc → AND	81.6%	293	85.7%	119

Face Detection Movies

Web Demo of Face Detection

http://demo.pittpatt.com/

Face Detection Using Cascades

Paul Viola and Michael Jones (2001).

``Robust Real-time Object Detection'',
Second International Workshop on Statistical
and Computational Theories of Vision
Modeling, Learning, Computing, and
Sampling, Vancouver, Canada, July 13, 2001.

Classical Face Detection

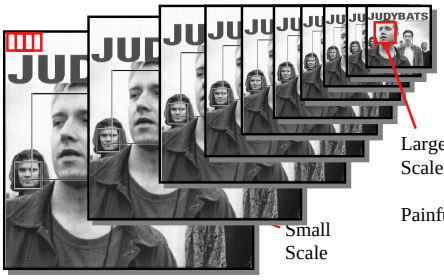

Painful!

Viola/Jones Face Detector

- Technical advantages:
  - Uses lots of very simple box features, enabling an efficient image representation
  - Scales features rather than source image
  - Cascaded classifier is very fast on non-faces
- Practical benefits:
  - Very fast, compact footprint
  - You don't have to implement it! (should be in latest version of OpenCV)

This next set of slides is from:

Robust Real-time Object Detection
by
Paul Viola and Michael Jones
ICCV 2001 Workshop on Statistical and Computation Theories of Vision

Presentation by Gyozo Gidofalvi
Computer Science and Engineering Department
University of California, San Diego
gyozo@cs.ucsd.edu
October 25, 2001

Object detection task

- Object detection framework: Given a set of images find regions in these images which contain instances of a certain kind of object.
- Task: Develop an algorithm to learn an fast and accurate method for object detection.

To capture ad-hoc domain knowledge classifiers for images do not operate on raw grayscale pixel values but rather on values obtained from applying simple filters to the pixels.

Definition of simple features for object detection

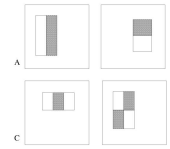

3 rectangular features types:

- *two-rectangle feature* type (horizontal/vertical)
- *three-rectangle feature* type
- *four-rectangle feature* type

Using a 24x24 pixel base detection window, with all the possible combination of horizontal and vertical location and scale of these feature types the full set of features has 49,396 features.

The motivation behind using rectangular features, as opposed to more expressive steerable filters is due to their extreme computational efficiency.

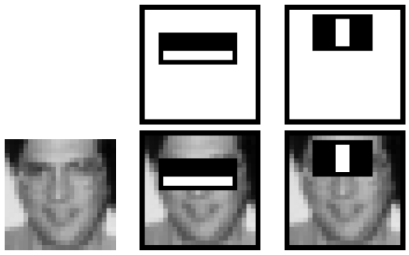

Integral image

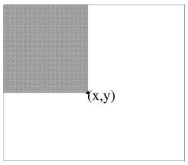

Def: The *integral image* at location (*x,y*), is the sum of the pixel values above and to the left of (*x,y*), inclusive.

Using the following two recurrences, where *i*(*x,y*) is the pixel value of original image at the given location and *s*(*x,y*) is the cumulative column sum, we can calculate the integral image representation of the image in a single pass.

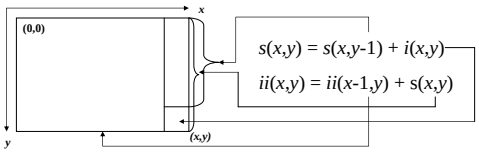

Rapid evaluation of rectangular features

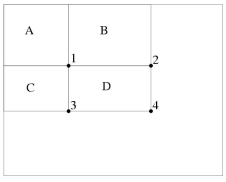

Using the integral image representation one can compute the value of any rectangular sum in constant time.

For example the integral sum inside rectangle D we can compute as:

$$ii(4) + ii(1) - ii(2) - ii(3)$$

As a result two-, three-, and four-rectangular features can be computed with 6, 8 and 9 array references respectively.

Challenges for learning a classification function

- Given a feature set and labeled training set of images one can apply number of machine learning techniques.
- Recall however, that there is 45,396 features associated with each image sub-window, hence the computation of all features is computationally prohibitive.
- Hypothesis: A combination of only a small number of these features can yield an effective classifier.
- Challenge: Find these discriminant features.

Performance of 200 feature face detector

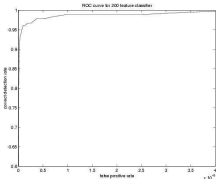

The ROC curve of the constructed classifies indicates that a reasonable detection rate of 0.95 can be achieved while maintaining an extremely low false positive rate of approximately 10^{-4}.

• First features selected by AdaBoost are meaningful and have high discriminative power

• By varying the threshold of the final classifier one can construct a two-feature classifier which has a detection rate of 1 and a false positive rate of 0.4.

Speed-up through the Attentional Cascade

• Simple, boosted classifiers can reject many of negative sub-windows while detecting all positive instances.

• Series of such simple classifiers can achieve good detection performance while eliminating the need for further processing of negative sub-windows.

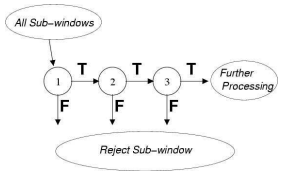

Processing in / training of the Attentional Cascade

Processing: is essentially identical to the processing performed by a degenerate decision tree, namely only a positive result from a previous classifier triggers the evaluation of the subsequent classifier.

Training: is also much like the training of a decision tree, namely subsequent classifiers are trained only on examples which pass through all the previous classifiers. Hence the task faced by classifiers further down the cascade is more difficult.

To achieve efficient cascade for a given false positive rate F and detection rate D we would like to minimize the expected number of features evaluated N:

$$N = n_0 + \sum_{i=1}^{K}\left(n_i \prod_{j<i} p_j\right)$$

Since this optimization is extremely difficult the usual framework is to choose a minimal acceptable false positive and detection rate per layer.

Experiments (dataset for training)

- 4916 positive training example were hand picked aligned, normalized, and scaled to a base resolution of 24x24
- 10,000 negative examples were selected by randomly picking sub-windows from 9500 images which did not contain faces

Experiments cont. (structure of the detector cascade)

• The final detector had 32 layers and 4297 features total

Layer number	1	2	3 to 5	6 and 7	8 to 12	13 to 32
Number of feautures	2	5	20	50	100	200
Detection rate	100%	100%	-	-	-	-
Rejection rate	60%	80%	-	-	-	-

• Speed of the detector ~ total number of features evaluated
• On the MIT-CMU test set the average number of features evaluated is 8 (out of 4297).
• The processing time of a 384 by 288 pixel image on a conventional personal computer about .067 seconds.
• Processing time should linearly scale with image size, hence processing of a 3.1 mega pixel images taken from a digital camera should approximately take 2 seconds.

Results

Testing of the final face detector was performed using the MIT+CMU frontal face test which consists of:

• 130 images

• 505 labeled frontal faces

Results in the table compare the performance of the detector to best face detectors known.

False detections	10	31	50	65	78	95	110	167	422
Viola-Jones	78.3%	85.2%	88.8%	89.8%	90.1%	90.8%	91.1%	91.8%	93.7%
Rowley-Baluja-Kanade	83.2%	86.0%	-	-	-	89.2%	-	90.1%	89.9%
Schneiderman-Kanade	-	-	-	94.4%	-	-	-	-	-
Roth-Yang-Ajuha	-	-	-	-	94.8%	-	-	-	-

Rowley at al.: use a combination of 1wo neural networks (simple network for prescreening larger regions, complex network for detection of faces).

Schneiderman at al.: use a set of models to capture the variation in facial appearance; each model describes the statistical behavior of a group of wavelet coefficients.

Results cont.

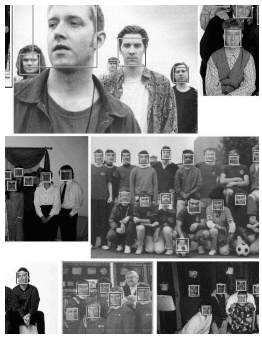

Conclusion

- The paper presents general object detection method which is illustrated on the face detection task.
- Using the integral image representation and simple rectangular features eliminate the need of expensive calculation of multi-scale image pyramid.
- Simple modification to AdaBoost gives a general technique for efficient feature selection.
- A general technique for constructing a cascade of homogeneous classifiers is presented, which can reject most of the negative examples at early stages of processing thereby significantly reducing computation time.
- A face detector using these techniques is presented which is comparable in classification performance to, and orders of magnitude faster than the best detectors know today.

Live OpenCV Demo of Face Detection Using Cascades

THE END