

Pattern Recognition and Machine Learning

Nachwa Aboubakr and James L. Crowley

ENSIMAG 3 - MMIS
Lesson 2

Fall Semester 2018
17 October 2018

Three Challenges for Face Detection with Skin Color

Outline

Project Groups	3
1. Performance Evaluation for Pattern Detectors	4
1.1 True and False Positives and Negatives,	4
1.2 ROC Curves	5
2.3 Precision and Recall	7
3.4 F-Measure	8
4.6 Accuracy	8
2. Face Detection using skin color	9
3. Challenge 1: Detecting skin pixels with color	10
3.1 Base Line: Ratio of RGB histograms	10
3.2 Variation: Detection with Chrominance (normalized color).14	
3.3 Free Parameters	15
4. Challenge 2: Detecting Faces with Skin Color	16
4.1 Baseline: Sliding Window Detector	17
4.2 Variation: Robust Detection using a Gaussian mask.....	19
4.3 Free parameters to test	20
5. Challenge 3: Face Localization	21
5.1 Detection vs Localization (Parameter Estimation)	21
5.2 Baseline: Localization by non-maximum suppression	22
5.3 Variation: Localization by Robust Estimation	22
5.4 Free parameters to test	23

Notation

x_d	A feature. An observed or measured value.
\vec{X}	A vector of features. An observation.
D	The number of dimensions for the vector \vec{X}
$\{C_k\}$	A set of K classes (or class labels).
K	Number of classes
$\vec{X} \in C_k$	Statement that the observation \vec{X} is a member of class C_k
\hat{C}_k	An estimated class label
	For a 2 class detection problem ($K=2$), $C_k \in \{P, N\}$
$R(\vec{X})$	A recognition function
$\hat{C}_k \leftarrow R(\vec{X})$	A recognition function that predicts \hat{C}_k from \vec{X}
$\{\vec{X}_m\}$	Training data for learning.
M	The number of training samples.
$y(\vec{X}_m)$	An annotation (or ground truth) function for $\{\vec{X}_m\} : y(\vec{X}_m) \in \{P, N\}$
$g(\vec{X}_m)$	Discriminant function. $0 \leq g(\vec{X}_m) \leq 1$
$X(i,j)$	An RGB image of size $R \times C$ pixels, 8 bits per color
$P(i,j)$	A probability image of size $R \times C$ pixels. Each pixel contains the probability (or likelihood) that the pixel (i,j) is a part of a face.
$\{X_n(i,j)\}$	A set of N images for training. $\vec{X}_m = X_n(i,j)$ where $m=n \cdot i \cdot j$
N	The number of training images.
$y(X_n(i,j))$	A ground-truth function that tells if pixel (i,j) of image n is part of a face.
M_T	The number of training pixels in the target class
$h(\vec{X})$	A multidimensional histogram of integer features \vec{X}
Q	The number of discrete values for each dimension (quantization) of $h(\vec{X})$
V	The number of cells in the histogram $h(\vec{X})$
$W_n(u,v)$	A rectangular window at (c_i, c_j) and size (width, height) spanning from (l, t) to (b, r)

Project Groups

Nom_patronymique	Prenom	Group Number	Oral Presentation.
BOUZBIBA	LAMYAA		
BRESSOLETTE	PIERRE		
GAENG	JEAN-BAPTISTE	1	
GESLIN	FLORENT		
GRAILLE	FANNY		
GUIBAL	GUSTAVE		
GUILLERMOU	LUCAS		
JOABAR	RUGDY	1	
JOLIVEL	MAXIME		
KLOPP	ANTONIN		
LABOURÉ	JULIEN		
LANGLET	VIRGILE		
LHULLIER	AYMERIC		
MOUHAMADOU	AHAMED	1	
PAVIE	MANON		
ROUVROY	MORGANE		
TILLET	MATHIEU		
TORCHET	AURÉLIEN		

1. Performance Evaluation for Pattern Detectors

1.1 True and False Positives and Negatives,

A pattern detector is a classifier with $K=2$.

Class $k=1$: The target pattern, also known as P or positive

Class $k=2$: Everything else, also known as N or negative.

Pattern detectors are used in computer vision, for example to detect faces, road signs, publicity logos, or other patterns of interest. They are also used in signal communications, data mining and many other domains.

Assume that we have M training samples, $\{\vec{X}_m\}$ along with a function $y(\vec{X}_m)$ that tells if a pixel is in the target class P or N. For face detection, this will be N images $X_n(i,j)$ where $\vec{X}_m = X_n(i,j)$ along with a function, $y(X_n(i,j))$ that tells is a pixel belongs to a face.

The pattern detector is learned as a detection function $g(\vec{X})$ followed by a decision rule, $d()$.

For example, the decision rule can be : if $g(\vec{X})+B \geq 0.5$ then P else N

Observations for which $g(\vec{X})+B > 0.5$ are estimated to be members of the target class. These are POSITIVE or P.

Observations for which $g(\vec{X})+B \leq 0.5$ are estimated to be members of the background. These are NEGATIVE or N.

We can encode the decision function to define our detection function $R(\vec{X}_m)$ as

$$R(\vec{X}) = d(g(\vec{X})+B) = \begin{cases} P & \text{if } g(\vec{X})+B \geq 0.5 \\ N & \text{if } g(\vec{X})+B < 0.5 \end{cases}$$

For training we need ground truth (annotation). For each training sample the annotation or ground truth tells us the real class y_m

$$y_m = \begin{cases} P & \vec{X}_m \in \text{Target - Class} \\ N & \text{otherwise} \end{cases}$$

The Classification can be TRUE or FALSE.

$$\text{if } R(\vec{X}_m) = y_m \text{ then T else F}$$

This gives

$R(\vec{X}_m) = y_m$ AND $R(\vec{X}_m) = P$ is a TRUE POSITIVE or TP

$R(\vec{X}_m) \neq y_m$ AND $R(\vec{X}_m) = P$ is a FALSE POSITIVE or FP

$R(\vec{X}_m) \neq y_m$ AND $R(\vec{X}_m) = N$ is a FALSE NEGATIVE or FN

$R(\vec{X}_m) = y_m$ AND $R(\vec{X}_m) = N$ is a TRUE NEGATIVE or TN

To better understand the detector we need a tool to explore the trade-off between making false detections (false positives) and missed detections (false negatives). The Receiver Operating Characteristic (ROC) provides such a tool.

1.2 ROC Curves

Two-class classifiers have long been used for signal detection problems in communications and have been used to demonstrate optimality for signal detection methods. The quality metric that is used is the Receiver Operating Characteristic (ROC) curve. This curve can be used to describe or compare any method for signal or pattern detection.

The ROC curve is generated by adding a variable Bias term to a discriminant function.

$$R(\vec{X}) = d(g(\vec{X}) + B)$$

and plotting the rate of true positive detection vs false positive detection where $R(\vec{X}_m)$ is the classifier as in lesson 1. As the bias term, B, is swept through a range of values, it changes the ratio of true positive detection to false positives.

For a ratio of histograms, $g(\bar{X}_m)$ is a probability ranging from 0 to 1.

The bias term, B , can act as an adjustable gain that sets the sensitivity of the detector. The bias term allows us to trade False Positives for False Negatives.

B can range from less than -0.5 to more than $+0.5$.

When $B \leq -0.5$ all detections will be Negative.

When $B > +0.5$ all detections will be Positive.

Between -0.5 and $+0.5$ $R(\bar{X})$ will give a mix of TP, TN, FP and FN.

The resulting curve is called a Receiver Operating Characteristics (ROC) curve.

The ROC plots True Positive Rate (TPR) against False Positive Rate (FNR) as a function of B for the training data $\{\bar{X}_m\}$, $\{y_m\}$.

For each training sample, the detection as either Positive (P) or Negative (N)

IF $g(\bar{X}_m) + B > 0.5$ THEN P else N

The detection can be TRUE (T) or FALSE (F) depending on the indicator variable y_m

IF $y_m = R(\bar{X}_m)$ THEN T else F

Combining these two values, any detection can be a True Positive (TP), False Positive (FP), True Negative (TN) or False Negative (FN).

For the M samples of the training data $\{\bar{X}_m\}$, $\{y_m\}$ we can define:

#P as the number of Positives,

#N as the number of Negatives,

#T as the number of True and

#F as the number of False,

From this we can define:

#TP as the number of True Positives,

#FP as the number of False Positives,

#TN as the number of True Negative,

#FN as the number of False Negatives.


Note that $\#P = \#TP + \#FN$

And $\#N = \#FP + \#TN$

The True Positive Rate (TPR) is $TPR = \frac{\#TP}{\#P} = \frac{\#TP}{\#TP + \#FN}$

The False Positive Rate (FPR) is $FPR = \frac{\#FP}{\#N} = \frac{\#FP}{\#FP + \#TN}$

The ROC plots the TPR against the FPR as a bias B is swept through a range of values.



When B is less than -0.5, all the samples are detected as N, and both the TPR and FPR are 0. As B increases both the TPR and FPR increase. Normally TPR should rise monotonically with FPR. If TPR and FPR are equal, then the detector is no better than chance.

The closer the curve approaches the upper left corner, the better the detector.

		$y_m = R(\vec{X}_m)$	
		T	F
$d(g(\vec{X}_m) + B > 0.5)$	P	True Positive (TP)	False Positive (FP)
	N	True Negative (TN)	False Negative (FN)

2.3 Precision and Recall

Precision, also called Positive Predictive Value (PPV), is the fraction of retrieved instances that are relevant to the problem.

$$PP = \frac{TP}{TP + FP}$$

A perfect precision score (PPV=1.0) means that every result retrieved by a search was relevant, but says nothing about whether all relevant documents were retrieved.

Recall, also known as sensitivity (S), hit rate, and True Positive Rate (TPR) is the fraction of relevant instances that are retrieved.

$$S = TPR = \frac{TP}{T} = \frac{TP}{TP + FN}$$

A perfect recall score (TPR=1.0) means that all relevant documents were retrieved by the search, but says nothing about how many irrelevant documents were also retrieved.

Both precision and recall are therefore based on an understanding and measure of relevance. In our case, “relevance” corresponds to “True”.

Precision answers the question “How many of the Positive Elements are True?”

Recall answers the question “How many of the True elements are Positive?”

In many domains, there is an inverse relationship between precision and recall. It is possible to increase one at the cost of reducing the other.

3.4 F-Measure

The F-measures combine precision and recall into a single value. The F measures measure the effectiveness of retrieval with respect to a user who attaches 2 times as much importance to recall as precision.

The F_1 score weights recall higher than precision.

4.5 F_1 Score:

$$F_1 = \frac{2TP}{2TP + FP + FN}$$

The F_1 score is the harmonic mean of precision and sensitivity.

This is the geometric mean divided by the arithmetic mean.

4.6 Accuracy

Accuracy is the fraction of test cases that are correctly classified (T).

$$ACC = \frac{T}{M} = \frac{TP + TN}{M}$$

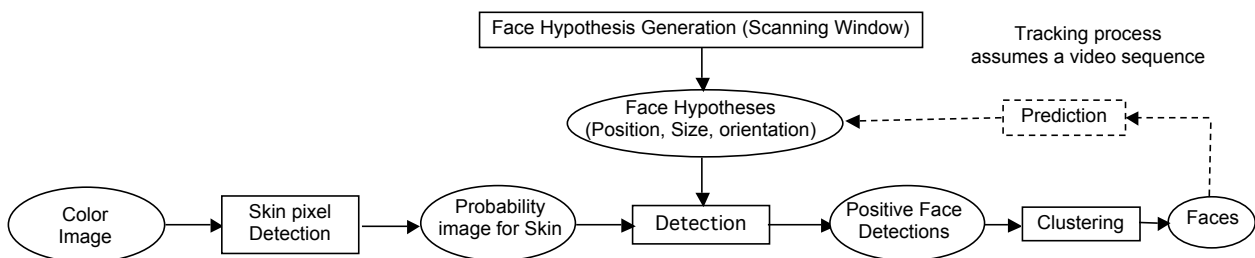
where M is the quantity of test data.

Note that the terms Accuracy and Precision have a very different meaning in Measurement theory. In measurement theory, accuracy is the average distance from a true value, while precision is a measure of the reproducibility for the measurement.

2. Face Detection using skin color

Skin color can be used to construct a simple detector for skin pixels in images. Color skin pixels can then be used to detect and track “blobs” that represent faces, hands and other skin colored regions in images.

The detector works by first computing the probability that each pixel contains skin. A sliding window (Region of Interest or ROI) is then scanned over the image. At each position, a weighted sum of probabilities is determined. Regions for which the weighted sum is above threshold are detected as faces. Adjacent face detections are grouped to form a single face detection.



Algorithm:

- 1) Compute probability of skin at each pixel.
- 2) Test for faces at possible positions and sizes (scanning Window).
- 3) Cluster adjacent detections.
- 4) Estimate precise face position, size and orientation from clusters of detections.

The function typically returns a "likelihood" estimate for a face at many adjacent positions, orientations and sizes.

We can break the problem down into 3 challenges.

Each challenge requires experimentally determining a set of parameters.

You are asked to determine these parameters experimentally.

3. Challenge 1: Detecting skin pixels with color

3.1 Base Line: Ratio of RGB histograms

The first challenge is to transform a color (RGB) image, $X(i,j)$ into an image where each pixel provides an estimate of the probability of skin, $P(i,j)$.

$$\text{Assume a color image : } X(i,j) = \begin{pmatrix} R \\ G \\ B \end{pmatrix} (i,j)$$

The algorithm will use a lookup table to convert color to probability.

$$P(i,j) \leftarrow L(\vec{X}(i,j))$$

The lookup table is constructed as a ratio of histograms, as explained below.

We can improve the results can be obtained by using a normalized color space as well as by tuning the quantization of the histogram to the data.

In the following, assume that we have a color image, where each pixel (i,j) is a color vector, $X(i,j)$, composed of 3 integers between 0 and 255 representing Red, Green and Blue.

$$X(i,j) = \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

Suppose that we have N color images of size $C \times R$ pixels, $X_n(i,j)$.

This gives a total of $M = C \times R \times N$ pixels.

Call this set of pixels $\vec{X}_m = X_n(i,j)$ where $m = i \cdot j \cdot n$

Suppose that we have a ground truth function $y(X_m)$ that tells us whether each pixel is target (P or Positive) or not target (N or Negative). A subset of M_T pixels that belong to a target class, T .

We allocate two tables $h(\vec{X})$ and $h_T(\vec{X})$ and use these to construct two histograms.

we can also write this as $\forall_m h(\vec{X}_m) = h(\vec{X}_m) + 1$

and $\forall_m y(\vec{X}_m) = P : h_T(\vec{X}_m) = h(\vec{X}_m) + 1 ; M_T \leftarrow M_T + 1$

For a color vector, $\vec{X} = \begin{pmatrix} R \\ G \\ B \end{pmatrix}$ we have two probabilities:

$$P(\vec{X}) = \frac{1}{M} h(\vec{X}) \quad \text{and} \quad P(\vec{X}|T) = \frac{1}{M_T} h_T(\vec{X})$$

Bayes rule tells us that we can estimate the probability that a pixel belongs to target class (Skin) given its color, \vec{X} as:

$$P(\text{Skin} | \vec{X}) = \frac{P(\vec{X} | \text{Skin}) P(\text{Skin})}{P(\vec{X})}$$

$P(\text{Skin})$ is the probability that a pixel belongs to the target class. This can be estimated from the training data by:

$$P(\text{Skin}) = \frac{M_T}{M}$$

From this we can show that the probability that a pixel is skin, is simply the ratio of the two tables.

$$P(\text{Skin} | \vec{X}) = \frac{P(\vec{X} | \text{Skin}) P(\text{Skin})}{P(\vec{X})} = \frac{\frac{1}{M_T} h_T(\vec{X}) \cdot \frac{M_T}{M}}{\frac{1}{M} h(\vec{X})} = \frac{h_T(\vec{X})}{h(\vec{X})}$$

We can use this to compute a lookup table $L_{\text{Skin}}(\vec{X}) = \frac{h_T(\vec{X})}{h(\vec{X})}$

if $h(\vec{X}) = 0$ then $h_{\text{Skin}}(\vec{X}) = 0 = 0$ because $h_{\text{Skin}}(\vec{X})$ is a subset of $h(\vec{X})$.

we will need to test this case to avoid divide by 0.

If we ASSUME that a new image, $X(i,j)$, has similar illumination and color composition then we can use this technique to assign a probability to each pixel by table lookup. The result is an image in which each pixel is a probability $P(i,j)$ that the pixel (i,j) belongs to class skin.

$$P(i,j) = L_{\text{Skin}}(\vec{X}(i,j))$$

Details:

- 1) alternative color codings may provide better recognition.
- 2) Different color quantizations can provide better recognition.

In this example, $h(\vec{X})$ and $h_r(\vec{X})$ are composed of $2^8 \cdot 2^8 \cdot 2^8 = 2^{24}$ cells.

We define this as the Capacity of the histogram, V

$$V = 2^8 \cdot 2^8 \cdot 2^8 = 2^{24} \text{ cells.}$$

In general, $V = Q^D$ where Q is the number of values per feature and D is the number of features.

This can result in a very sparse histogram. The reliability can be improved by using more training images from more cases.

A naive statistics view says to have at least 10 training samples for histogram cell.

That is $M \geq 10 V$. However, it is often easier to work with powers of 2.

For example, $2^3 = 8 \approx 10$ which is approximately 10.

This suggest that we required $M \geq 8 V = 2^3 V$.

Thus we would need $2^3 \cdot 2^{24} = 2^{27}$ training pixels. 2^7 Meg.

(Note that a 1024 x 1024 image contains 2^{20} pixels. This is the definition of 1 Meg)

Obtaining 2^7 Meg pixels is not a problem for $P(\vec{X}) = \frac{1}{M} h(\vec{X})$ but may be a problem for training the histogram of target pixels $P(\vec{X} | Skin) = \frac{1}{M_r} h_r(\vec{X})$.

A more realistic view is that the training data must contain a variety of training samples that reflect that variations in the real world.

What can we do? Two approaches:

We can reduce the number of values, Q , for each feature, or we can reduce the number of features.

For example, for many color images, $Q=32$ color values are sufficient to detect objects. We simply divide each color R, G, B by 8.

$$R' = \text{Trunc}(R/8), \quad G' = \text{Trunc}(G/8), \quad B' = \text{Trunc}(B/8).$$

We can also use our knowledge of physics to look for features that are "invariant".

3.2 Variation: Detection with Chrominance (normalized color).

Luminance captures local surface orientation (3D shape) while Chrominance is a signature for object pigment (identity). The color of pigment for any individual is generally constant. Luminance can change with pigment density (eg. Lips), and skin surface orientation, but chrominance will remain invariant.

Several methods exist to transform the (RGB) color pixels into a color space that separates Luminance from Chrominance.

$$\begin{pmatrix} L \\ c_1 \\ c_2 \end{pmatrix} \leftarrow \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

A popular space for skin detection is computed by dividing R and G by luminance. These are often called "r" and "g" in the literature.

Luminance: $L = R + G + B$

$$\text{Chrominance : } \quad r = c_1 = \frac{R}{R + G + B} \quad g = c_2 = \frac{G}{R + G + B}$$

The terms r and g have values between 0 and 1. To count statistics we need integer values. For this, it is common to represent r and g as integer numbers coded with Q values between 0 and Q-1 by :

$$r = \text{trunc} \left((Q-1) \cdot \frac{R}{R + G + B} \right) \quad g = \text{trunc} \left((Q-1) \cdot \frac{G}{R + G + B} \right)$$

From experience, $Q = 32$ color values seems to work well for skin with most web cameras, but this may change for certain data sets and should be experimentally verified.

Thus we can use a normalized vector $\vec{X} = \begin{pmatrix} r \\ g \end{pmatrix}$ as an invariant color signature for detecting skin in images.

Suppose we have a set of N training images $\{X_n(i,j)\}$ of size $R \times C$ where each pixel is an RGB color vector. This gives a total of $M = N \times I \times J$ color pixels.

Suppose that M_T of these are labeled as skin pixels i.e. $y(X_n(i,j)) = P$

We allocate two tables: $h(r,g)$ and $h_T(r,g)$ of size $Q \times Q$.

As before:

For all i,j,n in the training set $\{X_n(i,j)\}$:

BEGIN

$$r = \text{trunc}\left((Q-1) \cdot \frac{R}{R+G+B}\right) \quad g = \text{trunc}\left((Q-1) \cdot \frac{G}{R+G+B}\right)$$

$$h(r,g) = h(r,g) + 1$$

$$\text{IF } (y(X_n(i,j)) = P) \text{ THEN } h_T(r,g) = h_T(r,g) + 1 ; M_T \leftarrow M_T + 1$$

END

As before, we can obtain a lookup table $L_{\text{skin}}(r,g)$ that gives the probability that a pixel is skin.

$$L_{\text{skin}}(r,g) = \frac{h_T(r,g)}{h(r,g)}$$

Given a new RGB image $C(i,j)$: For all i,j :

$$r = \text{trunc}\left((Q-1) \cdot \frac{R}{R+G+B}\right) \quad g = \text{trunc}\left((Q-1) \cdot \frac{G}{R+G+B}\right)$$

$$P(i,j) = L_{\text{skin}}(r,g)$$

Where $P(i,j)$ is an image in which each pixel is a probability value from 0 to 1.

Probabilities can be expressed, of course be expressed as integers by multiplying by a quantization value (Q).

3.3 Free Parameters

A number of parameters remain unspecified. Specification of these parameters generally depends on the application domain and the training data. Often these parameters make it possible to trade off error rates for computing time. To properly evaluate the algorithm, you should measure performance and compare performance over a range of parameters.


Free parameters for color based skin detection include:

- The use of color coding (eg, RGB, rg, other codings for chrominance)
- Value of Q – the quantification for color color values.
- Training Data – different training and test regimes with variations such as the number of folds, and whether to train with a subset of the folds or all folds.

4. Challenge 2: Detecting Faces with Skin Color

We can detect faces from the mass of skin detections within a “Region of Interests” (ROI). The ROI is determined either by a-prior knowledge or by some form of search procedure. In many cases this is based on tracking of targets in the scene. This is not possible for our example, because our training and test data are restricted to static images.

In many vision detectors the ROI is simply a sliding window that scans the entire image. This is the technique used, for example, with the Viola Jones Face Detector.



Generally both the size and the position of the ROI are scanned. Size can be on a linear scale or a logarithmic scale. There are good reasons to use a logarithmic scale for size (e.g. Fractional powers of 2), for covering large ranges of scales, but this is beyond the scope of this lecture.

A common representation for the ROI is a rectangle represented by four coordinates: (top, left, bottom, right) or (t, l, b, r)

- t - "top" - first row of the ROI.
- l - "left" - first column of the ROI.
- b - "bottom" - last row of the ROI
- r - "right" - last column of the ROI.

(t,l,b,r) can be seen as a bounding box, expressed by opposite corners $(l,t), (r,b)$.

The center position, width and height are simply

$$c_i = \frac{l+r-1}{2}, \quad c_j = \frac{b+t-1}{2}, \quad w = r-l+1, \quad h = b-t+1$$

In many cases it is more convenient to fix c_i, c_j, w, h and calculate (l, t, r, b)

In this case the bounding box ROI is:

$$t = c_j - \frac{h-1}{2}, \quad b = c_j + \frac{h-1}{2}, \quad l = c_i - \frac{w-1}{2}, \quad r = c_i + \frac{w-1}{2}$$

The ratio of w and h can be fixed or variable. It is convenient to assure that w and h are odd numbers so that the boundaries fall on an integer pixel.

For detection in static images, we will need to test a range of ROIs with different positions and sizes. Note that with sufficient computing power, all windows can be processed in parallel.

In a scanning window system, the scanning step size, s , may also be varied. This may be the same size for row and column direction, or different step sizes may be used (say s_i and s_j)

Many engineers assume a step size of 1 pixel. However when the discriminant function is highly correlated at adjacent pixels (as is the case with color skin detection), then the resulting detection is very smooth. In this case a larger step size may be sufficient.

The largest reasonable value for s is half the width (and half the height if different step sizes in row and column are used). Beyond this the detection degrades rapidly because of aliasing effects, as can be demonstrated using signal processing techniques that are beyond the scope of this class.

When computing time is an issue, the detection algorithm can be optimized using a hierarchical search, starting with a large step size (Typically a power of 2) and recursively searching at higher step sizes near detected points.

4.1 Baseline: Sliding Window Detector

We can compute the ROI parameters for a face as a bounding box for an ellipse.

Let us define a face hypothesis as $\bar{X} = \begin{pmatrix} c_i \\ c_j \\ w \\ h \end{pmatrix}$

Since faces are generally vertical, we may assume that the ellipse is oriented with the major axis aligned with the row direction (j). The bounding box ROI is:

$$t = c_j - \frac{h-1}{2}, \quad b = c_j + \frac{h-1}{2}, \quad l = c_i - \frac{w-1}{2}, \quad r = c_i + \frac{w-1}{2}$$

The likelihood of a face at a position (c_i, c_j) of size (w, h) is:

$$g(\bar{X}) = \frac{1}{w \cdot h} \sum_{i=l}^r \sum_{j=t}^b P(i, j)$$

We can bias this likelihood to make a decision:

$$\text{IF } g(\vec{X}_m) + B > 0.5 \text{ THEN } R(\vec{X}_m) = P \text{ else } R(\vec{X}_m) = N$$

And of course

$$\text{IF } R(\vec{X}_m) = y(\vec{X}_m) \text{ THEN T else F.}$$

This technique will detect “faces” for a range of positions and sizes. As long as the detections overlap with the face ellipse in the data-base they are TRUE detections.

The problem with this technique is that faces are not square. The ROI gives equal weight to corners as the center. We can do better by weighting the pixels using a Gaussian function. This is called a “robust estimator” because it tends to reject outliers.

4.2 Variation: Robust Detection using a Gaussian mask.

In machine vision, fixation serves to reduce computational load and reduce errors by focusing processing on parts of the image that are most likely to contain information. A commonly practice is to use a Gaussian Window to suppress signals outside of the fixated region. A similar technique is used to suppress outliers for robust estimation.

A Gaussian window has the form:

$$G(i, j; \vec{\mu}, \Sigma) = e^{-\frac{1}{2}(\vec{P}-\vec{\mu})^T \Sigma^{-1}(\vec{P}-\vec{\mu})}$$

Where $\vec{P} = \begin{pmatrix} i \\ j \end{pmatrix}$ represents image positions, and $\vec{\mu} = \begin{pmatrix} \mu_i \\ \mu_j \end{pmatrix}$ is the center position of the window.

The covariance matrix of the window is: $\Sigma = \begin{pmatrix} \sigma_i^2 & \sigma_{ij} \\ \sigma_{ij} & \sigma_j^2 \end{pmatrix}$

The coefficients have a value of 1 at the center of the mask, and taper to 0.1 at a distance of 2σ and to 0.01 at a distance of 3σ .

We can use the parameters of the window to define a face hypothesis. $\vec{X} = \begin{pmatrix} \mu_i \\ \mu_j \\ \sigma_i^2 \\ \sigma_j^2 \\ \sigma_{ij} \end{pmatrix}$

Normally the Gaussian mask should be computed within a ROI determined by a bounding box. Typically the bounding box should be at least 2σ (standard deviations), but if speed is more important the false negatives, then computation time can be reduced by using a smaller ROI, down to as small as 1 standard deviation.

We can define the ROI as a 2σ bounding box:

$$t = c_j - 2\sigma_j, \quad b = c_j + 2\sigma_j, \quad l = c_i - 2\sigma_i, \quad r = c_i + 2\sigma_i$$

To evaluate a face hypothesis, we compute the face likelihood as a weighed sum of the skin probabilities by the Gaussian mask.

$$g(\vec{X}) = \sum_{i=1}^r \sum_{j=1}^b P(i, j) G(i, j; \vec{X})$$

As before, the discriminant, $g(\vec{X})$, has a value between 0 and 1. This is the likelihood that a face may be found at \vec{X} .

As before, faces are detected as:

$$\text{IF } g(\vec{X}_m) + B > 0.5 \text{ THEN } R(\vec{X}_m) = P \text{ else } R(\vec{X}_m) = N$$

4.3 Free parameters to test

As with color skin detection a number of parameters remain unspecified. To properly evaluate the algorithm, you should measure performance and compare performance over a range of parameters.

Free parameters for color face detection include:

For a simple scanning window, these include:

- Width and height of ROI
- Range of positions
- Step size for scanning windows
- The percentage of overlap with the ground truth that is considered a TRUE detection.

For a Gaussian detection window, parameters also include

- The width and height of the ROI compared to the standard deviations of the principal axis.
- The range and step sizes for orientation of the Gaussian window.

5. Challenge 3: Face Localization

5.1 Detection vs Localization (Parameter Estimation)

Detection: A face is considered to be “detected” if a positive detection is found at a position that overlaps the face in the ground truth. The simplest form of test is to verify that the location of the face is within the ellipse of the ground truth label for a face in the image.

For a face hypothesis at $\vec{X} = \begin{pmatrix} c_i \\ c_j \\ w \\ h \end{pmatrix}$ a face is detected if $R(\vec{X})=P$

The detection is TRUE if $R(\vec{X})=y(\vec{X})$

Using the ground truth $y(\vec{X}) = \begin{cases} P & \text{if } \left(\frac{(x-c_x)^2}{r_a^2} + \frac{(y-c_y)^2}{r_b^2} \leq 1 \right) \\ N & \text{otherwise} \end{cases}$

The Hypothesis is a TRUE POSITIVE detection if $y(\vec{X})=P$ and $R(\vec{X})=y(\vec{X})$.

A large number of TP faces will be detected for each ground truth face. If the search space includes multiple scales and orientations, than the number of detected faces will be even larger. All of these correspond to the same face!

Localization (more precisely parameter estimation) specifies precisely where, and with what parameters, the face may be found. There should be only ONE face located for each true face. The face should be located at a position, size and orientation as close to the true face as possible. When searching at multiple scales and orientations, each detection has the form of an error vector.

A distance metric, relating degrees and scale change to position is required to reduce this to a single number. For discrete samples the distance metric can provided implicitly by the sample step size in scale, orientation and position.

We can obtain a location (or parameter estimation) from multiple positive detections by suppressing detections for which the discriminant, $g(\vec{X})$ is not a local maximum. This requires specifying a measure for locality.

We can also estimate the parameters of the face from the moments of a “cloud” of detections at multiple positions, scales and orientations. This is the same robust estimation technique that we used above, extended to robustly estimate position, size and orientation.

5.2 Baseline: Localization by non-maximum suppression

In order to suppress non-maximal detections, the easiest method is to build a list of detections hypotheses, $\{\vec{X}\}$ over the desired range of positions, scales, orientations and any other parameters, and then filter this list to remove any hypothesis for which the discriminant not a local maximum. (A maximum within a some distance R.)

$\forall \vec{X}$

$\forall \vec{X}_i, \vec{X}_j \in \{\vec{X}\} : \text{IF } \text{Dist}(\vec{X}_i, \vec{X}_j) < R \text{ AND } g(\vec{X}_i) < g(\vec{X}_j) \text{ THEN } \{\vec{X}\} \leftarrow \{\vec{X}\} - \vec{X}_i$

This requires defining some notion of distance that includes position, scale and orientation. This is generally done with regard to an expected range of positions, orientations and scales over which a single face would be detected. For example, using a Mahalanobis distance (Distance normalized by covariance).

5.3 Variation: Localization by Robust Estimation

We can use robust estimation to estimate the most likely parameters from a set of detections. To do this, we will calculate the weighted moments from the set of detections.

Suppose that we have a set of N detections: $\{\vec{X}_n\}$ and that for each detection we have a discriminant $g(\vec{X}_n)$.

The mass of the detections is $M = \sum_{n=1}^N g(\vec{X}_n)$. This is the zeroth moment of the set of detections.

The "expected value" is $E\{\vec{X}\} = \frac{1}{M} \sum_{n=1}^N g(\vec{X}_n) \cdot \vec{X}_n$

This is the first moment (or center of gravity) of the values of $\{\vec{X}_n\}$.

The expected value is a vector of first moments for each parameter:

For D parameters, the center of gravity is a vector

$$\bar{\mu} = E\{\bar{X}\} = \frac{1}{M} \sum_{n=1}^N g(\bar{X}_n) \bar{X}_n = \begin{pmatrix} \mu_1 \\ \mu_2 \\ \dots \\ \mu_D \end{pmatrix} = \begin{pmatrix} \frac{1}{M} \sum_{n=1}^N g(\bar{X}_n) X_{1n} \\ \frac{1}{M} \sum_{n=1}^N g(\bar{X}_n) X_{2n} \\ \dots \\ \frac{1}{M} \sum_{n=1}^N g(\bar{X}_n) X_{Dn} \end{pmatrix}$$

The vector $\bar{\mu}$ the vector of weighted averages for the components of \bar{X}_n .

If there is only one face inside the set $\{\bar{X}_n\}$ of positive detections then $\bar{\mu}$ provides the most likely estimate for set of parameters the detection, (e.g. position, size and orientation).

In the case of multiple faces, the best estimate for each of the faces can be determined using the Expectation Maximization (EM) algorithm. Alternatively, the estimation may be limited to faces detections within a small region of the image.

5.4 Free parameters to test

As with color face detection a number of parameters remain unspecified. To properly evaluate the algorithm, you should measure performance and compare performance over a range of parameters.

Free parameters for color face detection include:

For a simple scanning window, these include:

- Number of Faces in the image
- Range of positions, size and orientations to test
- Distance to use for non-maximal suppression.
- Size of the region used to cluster faces detections.