

The University of Auckland

CBIR: Colour Features

COMPSCI.708.S1.C
A/P Georgy Gimel'farb

The University of Auckland

Other Opponent Colour Spaces

- **YUV** – PAL Phase Alternating Line TV (most of the European countries, some Asian countries, Australia, and New Zealand)
- Luminance (**Y**) and chrominance (**U, V**):

$$Y = 0.299R + 0.587G + 0.114B$$

$$U = 0.492(B - Y) = -0.147R - 0.289G + 0.436B$$

$$V = 0.877(R - Y) = 0.615R - 0.515G - 0.100B$$

The University of Auckland

Other Opponent Colour Spaces

- **YIQ** – NTSC National Television Systems Committee TV (USA, Canada, and Japan) with the luminance **Y** and chrominance **IQ** (equal to **UV** rotated by 33°):

$$Y = 0.299R + 0.587G + 0.114B$$

$$I = -0.545U + 0.839V = 0.596R - 0.275G - 0.321B$$

$$Q = 0.839U + 0.545V = 0.212R - 0.523G + 0.311B$$
- **YD_bD_r** – SECAM (Sequentiel Couleur a Memoire) TV (France, Russia, Eastern Europe) is the scaled **YUV**:

$$D_b = 3.059U \text{ and } D_r = -2.169V$$

The University of Auckland

Other Opponent Colour Spaces

- **YCbCr** – JPEG and MPEG coding standards: non-negative chrominance components by scaling and shifting the **YUV** co-ordinates:

$$Y = 0.257R + 0.504G + 0.098B + 16$$


$$C_b = -0.148R - 0.291G + 0.439B + 128$$

$$C_r = 0.439R - 0.368G - 0.071B + 128$$
- CIE uniform **Lab** (colour fax) and **Luv** luminance-chrominance colour spaces


The University of Auckland

Other Opponent Colour Spaces


- CIE **Lab** and **Luv** colour spaces compress well in the case of pictorial images:




RGB



L (luminance)



a



b

The University of Auckland

HMMD colour space

- **Hue-Min-Max-Difference**: a new space in MPEG-7
 - Used in the colour structure descriptor (CSD)
 - MPEG-7 supports also the greyscale (intensity only) space, RGB, HSV, and YC_b colour spaces
 - Hue: as in HSV space for MPEG-7
 - Min / Max: $\max\{R, G, B\}$ and $\min\{R, G, B\}$
 - Difference: $\max\{R, G, B\} - \min\{R, G, B\}$
 - Only 3 of the 4 components define the HMMD colour space
 - Intensity: $(\min\{R, G, B\} + \max\{R, G, B\})/2$
 - Chromaticity relates to the difference component

The University of Auckland

Vector Quantisation

- The whole 3D colour space is partitioned into K disjoint subsets, one per code word $c_k; k=1, \dots, K$
- All the colours of one subset are represented by, or quantised to, the same code word c_k
- In a perceptually good palette $C = \{c_1, c_2, \dots, c_K\}$, each subset contains visually similar colours
- Multidimensional clustering* of colours: K -means, fuzzy K -means, EM (Expectation - Maximisation)

Semester 1, 2006 Lecture G6 7

The University of Auckland

Vector Quantisation Algorithm

- Doubling the number of cluster centres until a prescribed limit
- Each iteration t : $K_t = 2^t$ centres $C_t = \{c_{k,t}; k = 1, \dots, K_t\}$
 - $T = 0$ - a single centre $c_{1,0}$ averaging colour vector over an image
 - At each next iteration t each previous centre $c_{k,t-1}; k = 1, \dots, K_{t-1}$, splits into the two new centres:
 - Provisional splitting into $c_{pr,k,t}$ and $c_{tr,K_{t-1}+k,t}$ by multiplying $c_{k,t-1}$ by the fixed factors $(1+w)$ and $(1-w)$, respectively (w - a fixed value), or shifting it to and from the most distant vector in the cluster, etc
 - Assigning each colour vector in the image to the closest cluster centre
 - Forming the new centres $c_{k,t}$ and $c_{K_{t-1}+k,t}$ by averaging the colour vectors assigned to each such cluster

Semester 1, 2006 Lecture G6 8

The University of Auckland

Colour Descriptors

- Colour histogram**: the distribution of colours within a whole image or within a specified region
 - invariance to rotation, translation, and scaling
 - but: it captures no semantic information
- Colour set**: by thresholding the colour histogram
- Colour moments**: the mean colour, variance, etc.
- Colour correlogram**: changes of spatial correlation of pairs of colours with distance

Semester 1, 2006 Lecture G6 9

The University of Auckland

Colour Histogram

- Distribution of colours within a whole image or a specified region
- Invariant to rotation, translation, and scaling of an object
- Does not capture semantic information: two images with similar colour histograms can possess totally different contents
- Quantised HSI (HSV) colour space is typically used to make the search partially invariant to illumination and object viewpoints
 - In the HSI (HSV) space, an Euclidean or similar component-wise distance between the colours specifies colour similarity quite well
 - The YUV colour space is also often used since it is standard for the MPEG video compression

Semester 1, 2006 Lecture G6 10

The University of Auckland

Colour Histogram

- Colour histogram $h(\text{image}) = (h_k(\text{image}); k=1, \dots, K)$ is a K -dimensional vector
 - each component $h_k(\text{image})$: the relative number of pixels of colour c_k in the image (i.e. the fraction of pixels being most similar to the corresponding representative colour)
 - To build the colour histogram, the image colours should be transformed to an appropriate colour space and quantised according to a particular codebook of the size K
 - Histograms for image regions roughly account for spatial colour distribution

Semester 1, 2006 Lecture G6 11

The University of Auckland

Colour Bins for an RGB Histogram

11x11x11 RGB colour space (from <http://www.owl.net.nco.edu/~elec301/Projects02/art5pycolor.html>)

Semester 1, 2006 Lecture G6 12

The University of Auckland

Colour Histogram



6x6 histogram (sphere-size-coded bin values)
From i-plugins.sourceforge.net/q-vtk/color-space/

Semester 1, 2006 Lecture G6 13

The University of Auckland

Matching Colour Histograms

- **QBE**: the database image compared to the query has to be re-quantised by finding for every pixel the closest colour in the query codebook
- Distance between two normalised histograms:

$$D(\mathbf{h}, \mathbf{h}') = \begin{cases} \sum_{k=1}^K |h_k - h'_k| \\ \sum_{k=1}^K (h_k - h'_k)^2 \\ 1 - \sum_{k=1}^K \min\{h_k, h'_k\} \equiv 0.5 \sum_{k=1}^K |h_k - h'_k| \end{cases}$$

Semester 1, 2006 Lecture G6 14

The University of Auckland

Matching Colour Histograms

- **Quadratic-form metrics**
 - comparing all colour pairs to account for cross-relations of different colour clusters:

$$D(\mathbf{h}, \mathbf{h}') = (\mathbf{h} - \mathbf{h}')^T \mathbf{A} (\mathbf{h} - \mathbf{h}')$$
 - $\mathbf{A} = [a_{ij}]$ is the positive definite symmetric $K \times K$ matrix
 - Its components $a_{ij} = a_{ji}$ specify the dissimilarity between the codewords \mathbf{c}_i and \mathbf{c}_j for the histogram components with indices i and j
- **Special case**: Mahalanobis distance: $\mathbf{A} = \Sigma^{-1}$ Covariance matrix

Semester 1, 2006 Lecture G6 15

The University of Auckland

MPEG-7 Colour Descriptors

Dominant Colour Descriptor

Scalable Colour Descriptor (SCD)

Colour Structure Descriptor (CSD)

Colour Layout Descriptor

Y, C_b, C_r
Y only (monochrome)
RGB
HSV
HMMD

Group of frames / pictures histogram (GoP)

- Generic colour histogram accurately captures the colour distribution for image search and retrieval but has too many independent characteristics to choose, e.g. colour space, quantisation in that space, and quantisation of the histogram values
- For the interoperability between different MPEG-7 descriptors: the histogram-based descriptors are limited to the SCD and the CSD

Semester 1, 2006 Lecture G6 16

The University of Auckland

Scalable Colour Descriptor

- HSV colour space uniformly quantised to 16 ... 256 bins
 - Histogram values: truncated to 11 bits; non-linearly mapped to 4 bits
- Haar transform coding: $a, b \Rightarrow a+b, a-b$
 - Repetitive HT: 256 bins - 128, 64, ... bins; sign-alternate high-pass coefficients: a truncated integer representation
- Similarity matching of the SCD histograms: the absolute distance
 - The same matching for the Haar coefficients
 - Fast coarse-to-fine matching
- **GoP descriptor**: extending the SCD to a collection of images

bins	16	32	64	128	256
H	4	8	8	8	16
S	2	2	2	4	4
V	2	2	4	4	4

Semester 1, 2006 Lecture G6 17

The University of Auckland

Colour Structure Descriptor

- HMMD colour space + an 8 x 8 structuring element
 - Colour structure histogram $\mathbf{h}_{cs} = (h_{cs,k} : k=0,1,\dots,K-1)$ has in each bin k the number of structuring elements in the image containing one or more pixels with colour \mathbf{c}_k
 - Bin values: normalised to the range [0.0, 1.0] and quantised to 8 bits / bin
 - 184-bin CSD is approximated with 120, 64, or 32 bins
 - Subsampling of the images larger than 2^{16} pixels: the 8 x 8 structuring element with the spatial extent $E \times E$:

$$E = 8S; S = 2^p; p = \max\{0, \text{int}\{0.5 \log_2(WH)\} - 8\}$$
 where W and H denote the width and height of an image, respectively, and $\text{int}\{z\}$ is the closest to z integer

Semester 1, 2006 Lecture G6 18

The University of Auckland

Dominant Colour Descriptor

- Representative colours from each image instead of being fixed in the colour space:

$$DCD = \{ \{ \mathbf{c}_i, p_i, v_i \} : i=1, \dots, n \}, s \}; p_1 + \dots + p_n = 1$$
 - n : the total number of colour clusters in the image region: $1 \leq n \leq 8$; typically, 3 - 4 colours describe satisfactorily the colour region
 - Attribute triple**: 3D colour vector \mathbf{c}_i , its percentage p_i , and its variance v_i
 - s : the overall colour spatial coherency
 - Colour clustering** minimises the weighted scatter D_i in each cluster C_i :

$$D_i = \sum_{\substack{(x,y) \in C_i \\ \mathbf{c} \in C_i}} w_{x,y} \| \mathbf{c}(x,y) - \mathbf{c}_i \|^2 \Rightarrow \mathbf{c}_i = \frac{\sum_{\substack{(x,y) \in C_i \\ \mathbf{c} \in C_i}} w_{x,y} \mathbf{c}(x,y)}{\sum_{\substack{(x,y) \in C_i \\ \mathbf{c} \in C_i}} w_{x,y}}$$

K-means-like algorithm

perceptual weight depending on local pixel statistics

Semester 1, 2006 Lecture G6 19

The University of Auckland

Dominant Colour Descriptor

- Spatial coherency**: overall spatial homogeneity of the dominant colours in the image
 - Spatial coherency of a dominant colour: the normalised average number of connected pixels of this colour (computed using a 3x3 mask)
 - Overall coherency: a linear combination of the individual spatial ones weighed with the corresponding percentages p_i
- Dissimilarity between DCDs with no optional variance/coherency**:

$$DCD_1 = \{ \{ \mathbf{c}_{1i}, p_{1i} \} : i=1, \dots, n_1 \}; DCD_2 = \{ \{ \mathbf{c}_{2i}, p_{2i} \} : i=1, \dots, n_2 \}$$

$$D^2(DCD_1, DCD_2) = \sum_{i=1}^{n_1} p_{1i}^2 + \sum_{j=1}^{n_2} p_{2j}^2 - 2 \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} a_{1i,2j} p_{1i} p_{2j}$$

Similarity coefficient depending on the Euclidean distance between the two colours

Semester 1, 2006 Lecture G6 20

The University of Auckland

Colour Layout Descriptor (CLD)

- Representative colours on an 8x8 grid \Rightarrow a discrete cosine transform (DCT) \Rightarrow encoding of the resulting coefficients
 - Input image \Rightarrow into 8x8 blocks \Rightarrow average colours in the Y, C_r, C_b space
 - Average colours \Rightarrow a series of 8x8 DCT coefficients (independently for Y, C_r , and C_b components) \Rightarrow a few low-frequency coefficients by zigzag scanning and quantisation:

$$CLD = \{ \{ \Psi_{Y,i}, \Psi_{C_r,i}, \Psi_{C_b,i} \} : i = 1, \dots, v \}$$

- $\Psi_{...i}$ - the i -th DCT coefficient of the corresponding colour component
- v - the number of the coefficients (generally different for each component)

$$D(CLD, CLD') = \sqrt{\sum_{i=1}^v w_{Y,i} (\Psi_{Y,i} - \Psi'_{Y,i})^2} + \sqrt{\sum_{i=1}^v w_{C_r,i} (\Psi_{C_r,i} - \Psi'_{C_r,i})^2} + \sqrt{\sum_{i=1}^v w_{C_b,i} (\Psi_{C_b,i} - \Psi'_{C_b,i})^2}$$

Larger weights - the to lower frequency coefficients

Semester 1, 2006 Lecture G6 21