

TRANSFORM CODING OF RGB-HISTOGRAMS

Reiner Lenz

*Dept. Science and Engineering, Linköping University, SE-60174 Norrköping, Sweden
reile@itn.liu.se*

Pedro Latorre Carmona

*Departamento de Lenguajes y Sistemas Informáticos, Universidad Jaume I, 12071 Castellón de la Plana, Spain
latorre@lsi.uji.es*

Keywords: RGB-Histograms, Image Databases, Finite Groups, Harmonic Analysis

Abstract: In this paper we introduce the representation theory of the symmetric group $S(3)$ as a tool to investigate the structure of the space of *RGB*-histograms. We show that the theory reveals that typical histogram spaces are highly structured and that these structures originate partly in group theoretically defined symmetries. The algorithms exploit this structure and constructs a PCA like decomposition without the need to construct correlation or covariance matrices and their eigenvectors. We implemented these algorithms and investigate their properties with the help of two real-world databases (one from an image provider and one from a image search engine company) containing over one million images.

1 Introduction

The number and size of image collections is growing steadily and with it the need to organize, search or browse these collections. These collections can also be used to study the statistical properties of large collections of visual data and to derive models of their internal structure.

In this paper we are interested in the understanding of the statistical structure of large image collections and in the design of algorithms for applications where huge numbers of images have to be processed very fast. Therefore our motivation lies in methods that are applicable to huge databases and enable fast response times.

We will investigate the color properties of images using one of the simplest and fastest color descriptors available: the *RGB* histogram (Swain and Ballard, 1991; Hafner et al., 1995). We will analyze the structure of the space of *RGB* histograms and algorithms for their fast processing and *structure reduction*.

The approach we use is based on the observation that compression and fast-processing methods are often tightly related to the underlying structure of the input signal space. This structure can often be described in terms of transformation groups. The best-known

class of algorithms of this type are the FFT-methods based on the group of shift operations. In the signal processing field these methods were generalized to the application of finite groups with applications to filtering, and pattern matching and computer vision. See (Cooley and Tukey, 1965; Holmes, 1979; Lenz, 1994; Lenz, 1995; Lenz, 2007; Rockmore, 2004) for some examples.

3D (like *RGB*) Histograms have a wide variety of applications in image processing, ranging from image indexing and retrieval (Sridhar et al., 2002; Yoo et al., 2002; Geusebroek, 2006; Smeulders et al., 2000) to object tracking (Comaniciu et al., 2003) to cite a few.

In the following we will first argue that a relevant transformation group for the space of *RGB* histograms is the group $S(3)$ of permutations of three objects. We will describe the basic facts from the representation theory of $S(3)$ and investigate the properties of the resulting transforms of histograms. As a result we will see that the generated structures have a PCA like decorrelation property.

We will apply these transforms to very large collections of images. One consisting of 760000 images representing the collection of an image provider and one with 360000 images from the database of an image search engine.

2 Notations and Basic Facts

We first summarize a few facts about permutations and representation theory and then we will describe how to generalize this to representations on spaces of RGB histograms. We only mention the basic facts and the interested reader should consult one of the books in the field, for example (Serre, 1977; Diaconis, 1988; Fulton and Harris, 1991; Fässler and Stiefel, 1992; Chirikjian and Kyatkin, 2000).

The permutations of three objects form the symmetric group $S(3)$. This abstract group comes in several realizations and we will freely change between them. In the most abstract context the permutations π are just elements of $\pi \in S(3)$. We will use it to investigate color images. We describe colors in the RGB coordinate system described by triples (R, G, B) . If we want to denote a triple with some numerical values then we write $(aaa), (aab), (abc)$ in the cases where all three, two or none of the values are equal. If a permutation changes the order within the triple we will simply use the new order of the generic RGB triple as a symbol for the permutation. The permutation (RBG) leaves the first element fixed and interchanges the other two. It should be clear from the context if we mean RGB-triples like (abc) or permutations like (RBG) . We define the special permutations π_c as the cyclic shift $\pi_c = (BRG)$ and π_r as the reflection (RBG) . These two permutations are the generators of $S(3)$ and all others can be written as compositions of these two. The group $S(3)$ has six elements and we usually order them as $\pi_c^0, \pi_c, \pi_c^2, \pi_r, \pi_c \pi_r, \pi_c^2 \pi_r$ or in RGB notation

$$(RGB), (BRG), (GBR), (RBG), (GRB), (BGR)$$

We see that the three even permutations π_c^0, π_c, π_c^2 form a commutative subgroup with the same properties as the group of 0, 120, 240 degrees rotations in the plane. The remaining odd permutations are obtained by preceding the even permutation with π_r .

If we consider the triples $(R, G, B)^t$ as vectors x in a three-dimensional vector space then we see that we can describe the effect of the permutations by a linear transformation described by a matrix. In this way the permutations π_c, π_r are associated with the matrices $T_G(\pi)$

$$T_G(\pi_c) = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \quad T_G(\pi_r) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \quad (1)$$

This is the simplest example of a representation of $S(3)$ which is a mapping from the group to matrices so that group operations go over to matrix multiplications. In this case the matrices are of size 3×3

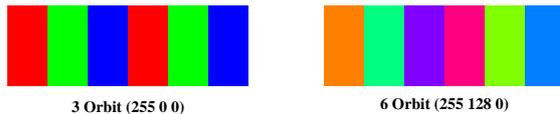


Figure 1: Examples of a three- and a six-orbit

and we say that we have a three-dimensional representation. The elements π_c, π_r generate $S(3)$ and therefore we find that also all six permutation matrices are products of $T_G(\pi_c), T_G(\pi_r)$.

If we apply all six permutations to triples (abc) we obtain the so called orbits. For triples with different values for a, b and c we generate six triples, if we apply them to a triple (abb) we get three triples and the triple (aaa) is invariant under all elements in $S(3)$. The orbits of $S(3)$ have therefore length six, three and one respectively. We denote a general orbit by O and the orbits of length one, three and six by O_1, O_3, O_6 . Two such orbits are illustrated in Fig.1 where each stripe shows one element in the orbit. For the three-orbit the colors are repeated for the odd permutations since the last two values in the RGB triple for the red image are identical.

We can use the concept of an orbit to construct new representations similar to those in Eq. (1). Take the six-orbit O_6 . We describe each element on O_6 by one of the six unit vectors in a six-dimensional vector space. Since permutations map elements in the orbit to other elements in the orbit we see that each permutation π defines a 6×6 permutation matrix $T_6(\pi)$ in the same way as those in Eq. (1). Also here it is sufficient to construct $T_6(\pi_c)$ and $T_6(\pi_r)$. The same construction holds for the three-orbits O_3 . For the one-orbit the matrices are simply the constants $T_1(\pi) = 1$. We denote these vector spaces (defined by the orbits) by V_1, V_3, V_6 .

The row- and column sums of permutation matrices are one and we see that $T(\pi)\mathbf{1} = \mathbf{1}$ where $T(\pi)$ is a permutation matrix and $\mathbf{1} = (1 \dots 1)$ is a vector of suitable length with only elements equal to one. This shows that the subspaces V_k^t of $V_k, (k = 1, 3, 6)$, spanned by $\mathbf{1}$ are invariant under all operations with permutation matrices. These spaces define the *trivial representation* of $S(3)$ (Fulton and Harris, 1991; Fässler and Stiefel, 1992).

Since V_k^t is an invariant subspace of $V_k, (k = 1, 3, 6)$ we see that their orthogonal complements are also invariant and we have thus decomposed the invariant spaces V_k into smaller invariant spaces and each of these subspaces defines a lower-dimensional representation (smaller matrices) of the group. The smallest such invariant spaces define the irreducible representations of the group (for definitions and ex-

amples see (Serre, 1977; Fulton and Harris, 1991; Fässler and Stiefel, 1992)).

The decomposition for the three-dimensional space V_3 is given by the matrix

$$P_3 = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 & 1 & 1 \\ \sqrt{2} & \sqrt{2}\cos(2\pi/3) & \sqrt{2}\cos(4\pi/3) \\ 0 & \sqrt{2}\sin(2\pi/3) & \sqrt{2}\sin(4\pi/3) \end{pmatrix} = \begin{pmatrix} \hat{\mathbf{1}} \\ P_2 \end{pmatrix} \quad (2)$$

where we identify the basis vector for the subspaces V_3^i in the first row. The orthogonal complement is spanned by the remaining two basis vectors and it can be shown that the space V_3^s spanned by these two cannot be split further. This defines another irreducible representation, known as the *standard representation* see (Fulton and Harris, 1991).

For the six-dimensional space V_6 it can be shown that the decomposition into irreducible representations is given by

$$P_6 = \begin{pmatrix} \hat{\mathbf{1}} & \hat{\mathbf{1}} \\ \hat{\mathbf{1}} & -\hat{\mathbf{1}} \\ P_2 & 0 \\ 0 & P_2 \end{pmatrix} \quad (3)$$

where $\hat{\mathbf{1}}$ represents 3D vectors with entries $\frac{1}{\sqrt{6}}$ and P_2 is the matrix with the two basis vectors defined in Eq.(2).

In the final stage of the construction we describe how the group operates on RGB histograms. We start with an orbit O with elements o . The permutations π are maps $\pi: O \rightarrow O$. Now take a linear function $f: O \rightarrow \mathbb{R}; o \mapsto f(o)$. We then define the new function f^π by $f^\pi(o) = f(\pi^{-1}(o))$. It can be shown that this defines a representation of $S(3)$ on the space of functions on the orbit. These representations can be reduced in the same manner as we did with the representations on the orbit.

For our application the functions of interest are the histograms. We will however modify this idea slightly. We consider a simple example first. Select an orbit O with elements o and assume that we have a probability distribution on O . Since O has finitely many elements this is a histogram h with the properties that $h(o) \geq 0$ and $\sum_{o \in O} h(o) = 1$. Applying a permutation π to the orbit elements defines a new histogram h^π . In the usual framework of representation theory we have orthonormal matrices $T(\pi)$ transforming vectors according to $h \mapsto T(\pi)h$. We thus have two transformations $h \mapsto h^\pi$ and $h \mapsto T(\pi)h$. The first of this preserves the L_1 -norm while the other preserves the L_2 -norm. We avoid this conflict and consider the square-roots of the probabilities instead (see

also (Srivastava et al., 2007)). In the following we use the definition:

$$h(o) = \sqrt{p(o)} \quad (4)$$

where $p(o)$ is the probability of the orbit element o and $h(o)$ is the modified "histogram".

We summarize the construction so far as follows:

- Split the RGB space into subsets X such that the split is compatible with the permutations in $S(3)$. We call the elements of these subsets *bins* and denote them by x .
- For a set of images compute the probabilities $p(x), x \in X$
- Convert and collect them in histogram vectors h with entries $h(x) = \sqrt{p(x)}$.
- Collect bins x that are related by permutations in orbits O_i . This defines a partition $X = \bigcup_i O_i$
- Every orbit O defines a representation of dimension one, three or six
- Split three-dimensional representations into two parts using the matrix P_3 defined in Eq.(2)
- Split six-dimensional representations into four parts using the matrix P_6 defined in Eq.(3)
- Leave the one-dimensional representations as they are
- The final decomposition is now:

$$V = V^t \oplus V^a \oplus V^s \quad (5)$$

where V is the space defined by the bins. The space V^t is the invariant subspace associated with the one-point orbits and the invariant parts (first rows in P_3, P_6) of the three and six point orbits. V^a is the subspace associated with the six-point orbits and depends on the even/odd properties (second row in P_6) of the six point orbits. The V^s part follows the P_2 parts in the three- and six-point orbit transforms, Eqs.(2, 3).

3 Implementation

In the derivation we only required that the split of the RGB space is compatible with the operation of $S(3)$. Since we are interested in fast implementations we will only consider very regular partitions where we split the R, G and B intervals into eight bins each, leading to a 512D RGB histogram. The procedure is easily generalized to all equal partitions of the axis into β bins leading to β^3 dimensional RGB histograms. Due to the exponential growth, values $\beta \leq 8$ are most realistic.

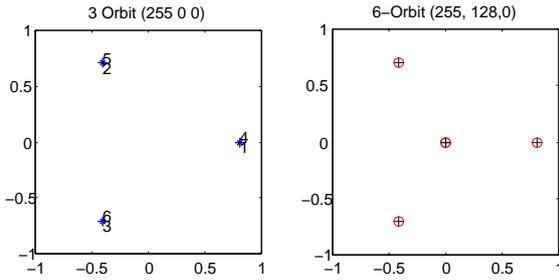


Figure 2: Orbit decompositions for the standard representation block

For eight bins per channel we use octal representations of the bin-number and write (klm) for the number $k + 8l + 64m$. One-point orbits are invariant under all permutations, therefore they represent gray-values (kkk) . Black is characterized by (000) and white by (777) . The three-orbits are given by bin numbers (kll) . Consider as example the images given by the stripes in the left part of Fig. 1. The histogram for the first stripe has a one at position (700) . Applying the six permutations we get the six stripes in this figure and six histograms. Applying the transformation P_3 to the three-orbit section of the histogram space given by $(700), (070), (007)$ we find that the first entry is always one and the positions in the other two dimensions (the P_2 part) transform as in an equal sided triangle as shown in the left part of Figure 2. These two-dimensional vectors transform thus as 120 degrees rotations under permutations.

The orbit of (740) , representing the RGB vector $(255, 128, 0)$, are the six stripes in the right part of Figure 1. Using the decomposition defined by P_6 we get two two-dimensional vectors (from the last four rows of the matrix). We see the projection of the six orbit colors to these subspaces in the right part of Figure 2. The points marked with an "o" belong to one subspace, the "+" points to the other.

The coordinates of the projections into the alternating and the standard parts are collected in Table 1

We have now described how to reorganize the histograms so that the different components show simple transformation properties under channel permutations. This is one of the advantages of this approach. The other is the relation to principal component analysis (PCA) that we will explain now.

We start with a simple example. Consider a vector h defined on a three-point orbit. Generate all different versions h^π under permutations and compute the matrix $C = \sum_\pi h^\pi h^{\pi'}$. It is invariant under a re-ordering of the orbit since this will simply rearrange the sum. This is the simplest example of

Alternating Representation

RGB	GBR	BRG	RBG	GRB	BGR
0.408	0.408	0.408	-0.408	-0.408	-0.408

Standard Representation

RGB	GBR	BRG	RBG	GRB	BGR
0.817	-0.408	-0.408	0	0	0
0	-0.707	0.707	0	0	0
0	0	0	-0.408	0.817	-0.408
0	0	0	0.707	0	-0.707

Table 1: Coordinates of projections for six-point orbits

an $S(3)$ -symmetric matrix. We generalize this to the definition of a wide-sense-stationary process as follows: Assume that we have vectors h in a vector space V and the permutations $\pi \in S(3)$ operate on these vectors by $h \mapsto h^\pi$. Assume further that we have a stochastic process with stochastic variable ω and values in $h_\omega \in V$. We define the correlation matrix Σ of this process as $\Sigma = E(h_\omega h_\omega')$ where $E(\cdot)$ denotes the expectation with respect to the stochastic variable ω . Assume further that we have a representation $T(\pi)$ on V .

Definition The stochastic process with correlation matrix Σ is T -wide-sense stationary if $T(\pi)\Sigma = \Sigma T(\pi)$ for all $\pi \in S(3)$.

We will only consider representations for which the matrices T are orthonormal and in this case we have $\Sigma = T(\pi)\Sigma T(\pi)'$ for all $\pi \in S(3)$. But $T(\pi)\Sigma T(\pi)'$ is the correlation matrix of the stochastic process h in the new coordinate system $T(\pi)h$ and we see that wide-sense-stationarity means that the correlation matrix is independent of a certain class of coordinate transforms.

The general theory (Schur's Lemma, (Fässler and Stiefel, 1992)) shows that we can find a matrix U (defining a new basis in the vector space) such that the correlation matrix in the new space is block diagonal. This matrix U depends only on the group $S(3)$:

Theorem 3.1 For an $S(3)$ -symmetric process with correlation matrix Σ we can find a matrix U such that:

$$U\Sigma U' = \begin{pmatrix} \Sigma^t & 0 & 0 \\ 0 & \Sigma^a & 0 \\ 0 & 0 & \Sigma^s \end{pmatrix} \quad (6)$$

This transformation $h \mapsto Uh$ defines a partial principal component analysis of the histogram space by block-diagonalizing the correlation matrix and it is given by the construction described in the previous Section 2.

Database	$E(\ v^f\ ^2)$	$E(\ v^a\ ^2)$	$E(\ v^s\ ^2)$	$\text{Max}(\ v^a\ ^2)$	zero vs. non-zero
PDB	0.606	0.042	0.350	0.312	0.057
SDB	0.678	0.031	0.291	0.278	0.144

Table 2: Contributions of the coefficients

4 Experiments

We implemented the transform described above and used it to investigate the internal structure of two large image databases. One of the image databases (we denote it by **PDB**) contains images from an image provider. It consists of 754034 (watermarked) images and is used on their website. The second database (referred to as **SDB**) contains 322903 images collected from the internet by a commercial image search engine. They are indexed by 31 different keyword categories ranging very general concepts like “beach” to very special like “Claude Monet”. In all experiments we computed first the RGB histograms p for all these image and then the square-root of the histogram entries. In all experiments we used eight bins/channel. We get 8 one-point orbits, 56 three-point orbits and 56 six-point orbits. The dimensions of the blocks described in Eq.(6) are therefore (120,56,336). We apply the transform resulting in the new vector $v = (v^f, v^a, v^s)$ corresponding to the vector spaces in Eq.(5).

We describe first some of our experiments regarding the statistical properties of these databases and then we illustrate the compression properties of the group theoretical transform in Section 3.

We first computed the norms of the vectors v^f, v^a and v^s , their mean and max value of $\|v^a\|^2$ for **PDB** and **SDB**. The coefficients in v^a (see also Table 1) are given by differences between contributions from even permutations and odd permutations in a six-orbit. If we assume that even and odd permutations are statistically equally likely then we expect the value of these coefficients to be small on average. We also computed the histogram (with 1000 bins) of these norms $\|v^a\|^2$ and computed the ratio between the probability of the first bin (with the small values of the norm) and the sum over the remaining bins (representing the non-zero norm values). The results are collected in Table 2.

This shows that the main contribution comes from v^f and the contributions from v^a are indeed low. We also see that there is a difference between the two databases where the contribution of the v^a is higher for **PDB**. One reason for this could be the higher proportion of cartoon-like images with distinct color distributions in **PDB** as compared with **SDB**.

From the construction we know that the modified histograms and their transforms are unit vectors $\|v\|^2 = \|v^f\|^2 + \|v^a\|^2 + \|v^s\|^2 = 1$. Since $\|v^a\|^2$ is small we conclude that the 3D $v = (v^f, v^a, v^s)$ are concentrated in the neighborhood of one quarter of a great circle of the unit sphere. The length of v^f is also larger than the length of v^s . Based on these heuristic considerations we introduce the following polar coordinate system on the unit vectors given by the norms of the projection vectors:

$$(v^f, v^a, v^s) = (\cos \phi \cos \theta, \cos \phi \sin \theta, \sin \phi) \quad (7)$$

The angle ϕ corresponds to the latitude and we think of it as an indication of the unbalance between the three channels (for a value of zero all the contribution is in the v^f part). The (longitudinal) angle θ is a measure of the contribution of v^a . The (ϕ, θ) -distribution of the images (every dot corresponds to one image) in the **PDB** is shown in the upper plot of Figure 3. The corresponding probability density distribution is shown in the lower plot. This figure shows that the distribution of the images is concentrated around the origin and that the distribution has a banana-like shape in the (ϕ, θ) -space.

The positions of the eight extreme points of the convex hull are marked with filled circles in the upper plot of Figure 3. The images belonging to the eight extreme points of the convex hull are shown in Figure 4.

Theorem 3.1 shows that wide-sense-stationary processes are partially decorrelated by the transform. In the remaining part of this section we will now investigate if the two databases define wide-sense-stationary processes.

We illustrate the effect of the transform on the correlation matrix in Figure 5 showing the contour plots of the correlation matrices computed from the square-root transformed histograms before and after the transformation. It can be clearly seen that the effect of the transformation is a concentration in the first 120 components given by the vectors v^f .

In the following experiment we evaluated the approximation error introduced by reducing the correlation matrix computed from the transformed histograms (Figure 5) to the block-diagonal matrix with block-sizes (120,56, 336). We computed the first 20 eigenvectors of the full correlation matrices and found that they explain about 85% of the summed eigenvalues for both databases. We also computed the 20 eigenvectors for the block-structured correlation matrix. From the construction of the blocks we expect that these eigenvectors of the block-diagonal matrix are elements of the 120, 56 or 336-dimensional

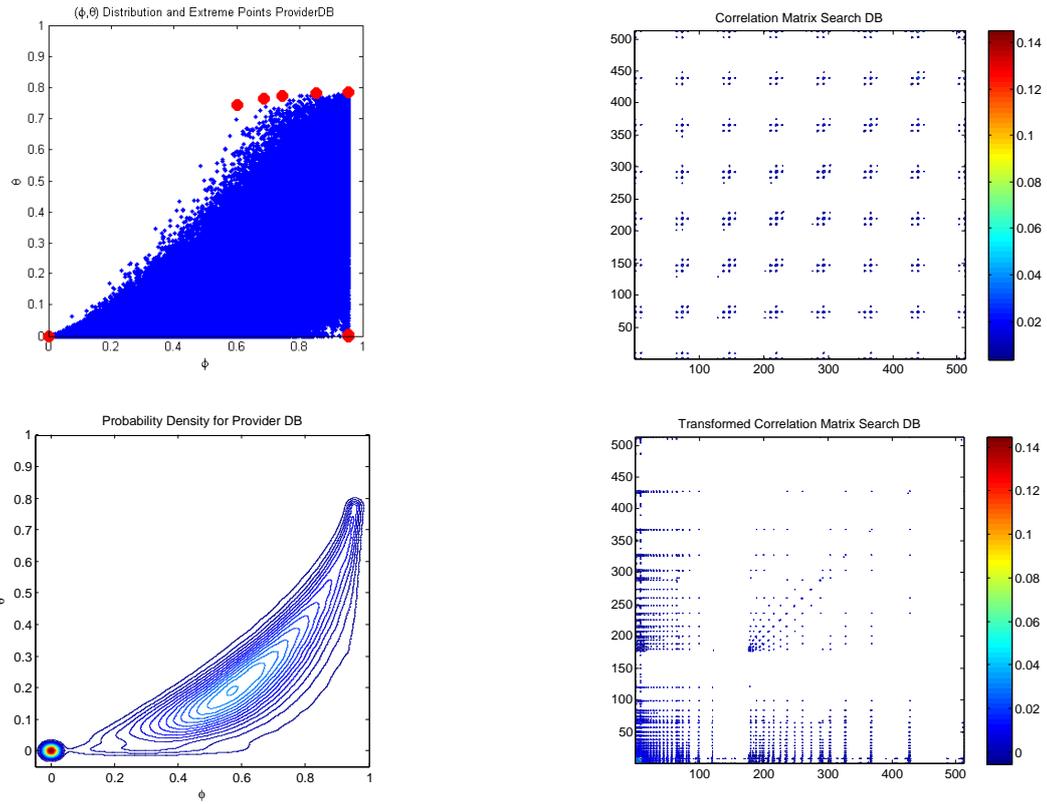


Figure 3: Location and probability distribution from **PDB**

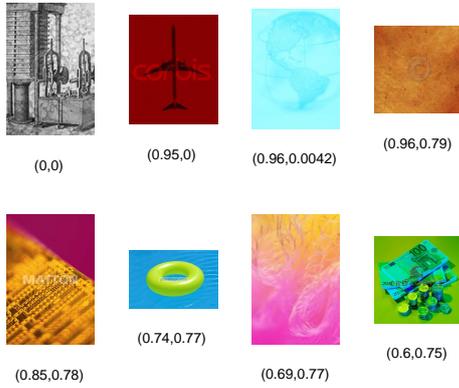


Figure 4: Extreme images in **PDB**

subspaces defined by the blocks. In Table 3 we illustrate the accumulated, normed eigenvalues $A_K = (\sum \gamma_{k=1}^K) / (\sum \gamma_{k=1}^{512})$ where γ_k is the k-th eigenvalue and to which block the different eigenvectors belong. We see the minor role of the coefficients in v^d : the only

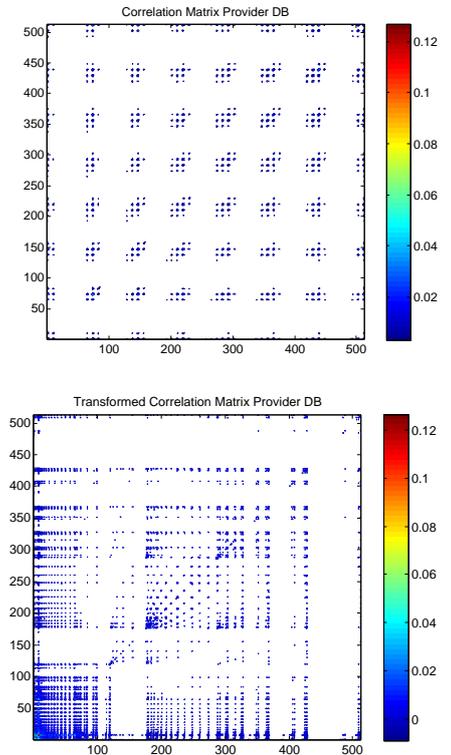


Figure 5: Original and transformed correlation matrices

	1	2	3	4	5	6	7	8	9	10
SDB	0.40	0.51	0.57	0.62	0.65	0.67	0.70	0.72	0.74	0.75
Block	1	1	3	1	1	3	1	3	3	3
PDB	0.41	0.49	0.55	0.59	0.63	0.66	0.68	0.70	0.71	0.72
Block	1	3	1	1	3	1	3	3	2	1

Table 3: Accumulated Eigenvalues and Block-Number for the Block-Diagonal Eigenvectors

eigenvectors from the second block are in positions nine and eleven (last one not shown in Table 3). Also here we see that **PDB** has a higher contribution from v^a .

To get a more quantitative measure on how good the block-diagonal eigenvalues approximate the full matrix eigenvectors we computed the approximation values:

$$\gamma_{m,n} = \frac{\sum_{k=1}^n \left\| \tilde{B}_m b_k \right\|}{n} \quad (8)$$

where \tilde{B}_m is the matrix of the m first eigenvectors of the block-diagonal approximation and b_k is the k -th eigenvector of the full correlation matrix. The value of $\gamma_{m,n}$ is a normalized measure on how large part of the n first eigenvectors of the full correlation matrix are projected into the space of the first m eigenvectors of the block-diagonal matrix. Some of the results are shown in Figure 6 where the solid line shows the values of γ_{mm} , $m = 1..20$ (equal number of full and block-diagonal eigenvectors) and the dashed line shows $\gamma_{20,n}$, the result for the full set of 20 eigenvectors of the block-diagonal matrix.

We see that the first 20 eigenvectors explain about 85% of the contributions from the first 20 eigenvectors of the full correlation matrix. Since the contribution from the last eigenvectors to the total variation (see also Table 3) is small we find that the approximation computed from the block-diagonal matrix is probably sufficient for most applications.

5 Extensions and Conclusions

The experiments described so far show that the transform, derived from the assumption that permutations of the R, G and B channels are likely to occur with the same probabilities, lead to a separation of the histogram space into three clearly separated blocks. From an $S(3)$ point of view nothing can be said about the internal structure of these blocks and if we want to transform them we have to use different sources of information about them. In the previous section we

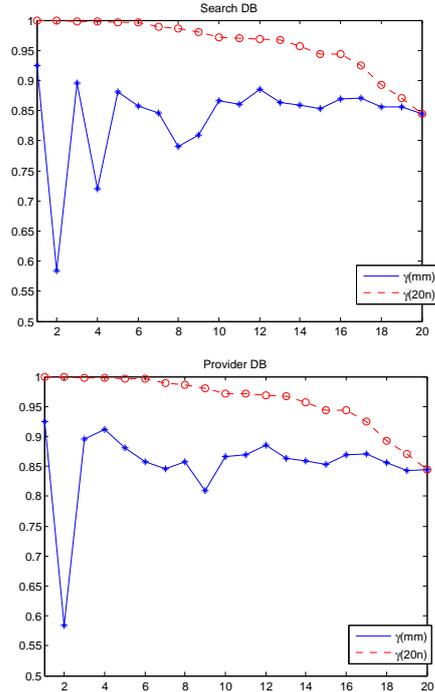


Figure 6: Projection properties of the approximation

used standard PCA to identify important sections of these parts of the space. We also implemented a second transform that takes into account the multilevel structure given by the different bin sizes. All experiments described so far used an eight-bin quantization of the RGB axis resulting in a 512D RGB histogram. If we reduce the number of bins by a factor of 2 we get $4^3 = 64$ bin RGB histograms and at the next level we have only $2 \times 2 \times 2 = 8$ bins left. Bin reduction and channel permutations are two independent processes and care has to be taken combining them into a multilevel-permutation transform. We implemented such a transform which will be described elsewhere. An illustration of the results we achieved is shown in Figure 7 where we reduced the first 120×120 block further using bin-reduction to 4 bins per channel. It can be shown that this defines a split of the 120D space into one subspace of dimension 20 and the rest of dimension 100. The increased concentration of information in the first few components is clearly visible. Summarizing, we conclude that the intuitive assumption that R, G and B channels can be interchanged on average motivates the application of tools from the representation theory of the symmetric group $S(3)$. We implemented a fast transform using tools from representation theory and used them to investigate the structure of two large image

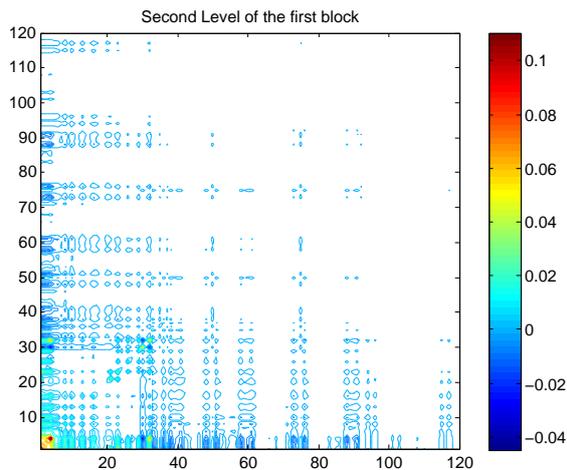


Figure 7: Multilevel transform of the first block (PDB)

databases and to develop fast PCA-like compression methods.

ACKNOWLEDGEMENTS

The support of the Swedish Research Council and the Ministry of Education and Science of the Spanish Government through the *DATASAT* project (*ESP-2005-00724-C05-C05*) and the *MIPRCV* project (*CSD2007-00018*) are gratefully acknowledged. Pedro Latorre Carmona is a *Juan de la Cierva* Programme researcher (Ministry of Education and Science). The databases were provided by Picsearch AB, Stockholm and Matton AB, Stockholm.

REFERENCES

- Chirikjian, G. S. and Kyatkin, A. B. (2000). *Engineering applications of noncommutative harmonic analysis :with emphasis on rotation and motion groups*. CRC Press, Boca Raton, FL.
- Comaniciu, D., Ramesh, V., and Meer, P. (2003). Kernel-based object tracking. *IEEE-TPAMI*, 25:564–577.
- Cooley, J. W. and Tukey, J. W. (1965). An algorithm for machine calculation of complex fourier series. *Mathematical Computations*, 19:297–301.
- Diaconis, P. (1988). *Group representation in probability and statistics* /. Inst. Math. Stat., Hayward, Calif. .:
- Fässler, A. and Stiefel, E. (1992). *Group theoretical methods and their applications*. Birkhuser Boston.
- Fulton, W. and Harris, J. (1991). *Representation Theory*. Springer, New York.
- Geusebroek, J. M. (2006). Compact object descriptors from local colour invariant histograms. In *Proc. British Machine Vision Conf.*
- Hafner, J., Sawhney, H. S., Equitz, W., Flickner, M., and Niblack, W. (1995). Efficient color histogram indexing for quadratic form distance functions. *IEEE-TPAMI*, 17(7):729–736.
- Holmes, R. B. (1979). Mathematical foundations of signal processing. *SIAM Review*, 21(3):361–388.
- Lenz, R. (1994). *Group Theoretical Transforms in Image Processing*. LNCS, Springer <http://www.itn.liu.se/~reile/LNCS413>.
- Lenz, R. (1995). Investigation of receptive fields using representations of dihedral groups. *J. Vis. Comm. Im. Rep.*, 6(3):209–227.
- Lenz, R. (2007). Crystal vision-applications of point groups in computer vision. *Proc. ACCV 2007*, volume 4844, of LNCS, page 744–753, Springer.
- Rockmore, D. (2004). Recent progress and applications in group fts. In Byrnes, J. and Ostheimer, G., editors, *Computational noncommutative algebra and applications*. Kluwer.
- Serre, J. P. (1977). *Linear representations of finite groups*. Springer.
- Smeulders, A. W. M., Worring, M., Santini, S., Gupta, A., and R., J. (2000). Content based image retrieval at the end of the early years. *IEEE TPAMI*, 22:1349–1380.
- Sridhar, V., Nascimento, M. A., and Li, X. (2002). Region-based image retrieval using multiple-features. In *LNCS, 2314*. Springer.
- Srivastava, A., Jermyn, I., and Joshi, S. (2007). Riemannian analysis of probability density functions with applications in vision. In *2007 IEEE CVPR'07*, Minneapolis, MN.
- Swain, M. and Ballard, D. (1991). Color indexing. *Int J Comput Vision*, 7(1):11–32.
- Tran, L. V. and Lenz, R. (2005). Compact colour descriptors for colour-based image retrieval. *Signal Processing*, 85(2):233–246.
- Yoo, H. W., Jang, D. S., Jung, S. H., and Park, J. H. (2002). Visual information retrieval system via content-based approach. *Pattern Recognition*, 35:749–769.