

WREATH PRODUCTS FOR IMAGE PROCESSING

Dennis M. Healy Jr., Gagan Mirchandani, Timothy E. Olson, and Daniel N. Rockmore

Dartmouth College
Department of Mathematics
Hanover, NH 03755
{healy,olson,rockmore}@cs.dartmouth.edu
mirchand@uvm-gen.emba.uvm.edu

ABSTRACT

We present a wreath product approach for matched filtering to detect rotated copies of a template in an image. We view the image as a homogeneous space for a wreath product, a noncommutative symmetry group. The corresponding Fourier analysis has a natural multiresolution structure and accompanying efficient algorithm which we explain and illustrate with an example. The associated matched filter is a new example of the use of a noncommutative convolution for image processing. Numerical experiments are described in which this noncommutative approach outperforms standard Fourier-based methods.

1. INTRODUCTION

The usefulness of matched filtering for signal detection and analysis is well-established. To this end the use of the classical FFT has been pivotal for the calculation of the necessary correlations. One of the fundamental properties of matched filtering via the traditional FFT is the time-independent response of the filter.

In multiple dimensions time-independence becomes spatial-independence. This independence implies that the maximal response of the matched filter will yield the location of the object function. This independence does not extend to rotations of the object function or the filter, however and rotational instabilities of matched filters in multiple dimensions cause many problems.

Here we introduce a Fourier analysis based on viewing the signal as a tree-like homogeneous space for a wreath product, a particular class of finite groups. We outline the corresponding Fourier analysis which has a natural associated multiresolution analysis. We propose to use this Fourier analysis to design matched filtering algorithms which are more robust to rotational invariance than traditional Fourier analysis. This sort of work is in the spirit of [2, 3, 4, 5, 9, 10] all of which espouse the utility and applicability of noncommutative Fourier transforms for pattern recognition and filtering.

D.H., T.O. and D.R. partially supported by ARPA as administered by the AFOSR under contract DOD F4960-93-1-0567. D.R. also supported by NSF DMS Awards 9404275 and 9553134. G.M.'s Permanent address: Department of Mathematics, University of Vermont, Burlington VT

The advantage of this type of analysis is its flexibility. Both the filter and tree structure of the Fourier analysis can be varied in order to increase the ability of the matched filter to detect an anisotropic object. As a consequence, an optimized matched filter of this type will naturally outperform traditional matched filter which comes from a subclass of algorithms for detection. We will present a simple numerical example where a matched filter was designed in this manner (but was not necessarily optimized). We show that this matched filter outperforms a traditional matched filter on rotated copies of the object function. As in the commutative case these algorithms rely on convolution which even in the noncommutative case can still be computed efficiently using various generalizations of the classical FFT (cf. [6] for a survey and extensive bibliography).

While it is true that the increased rotational invariance of our matched filters comes at the price of slightly decreased translational invariance, the translational invariance can be recovered by accepting a slightly higher computational burden.

2. WREATH PRODUCTS AND TREES

Wreath products are finite groups which arise naturally as the automorphism groups of various combinatorial structures [11]. For the purposes of this note we are interested in the specific case of certain wreath products which arise as automorphism groups of **spherically homogeneous trees**. These are rooted trees with constant branching at a fixed distance from the root. If the tree has height n and is such that all nodes at distance k from the root have m_k descendants, then the tree is said to **have type** $\mathbf{m} = (m_0, \dots, m_{n-1})$ and is denoted as $X_{\mathbf{m}}$.

The automorphism group of $X_{\mathbf{m}}$ is an example of a **wreath product**. These are precisely the symmetries of $X_{\mathbf{m}}$ that leave the root fixed and map adjacent vertices to adjacent vertices. To illustrate, consider Figure 1 which displays $X_{(3,4)}$. An automorphism of this graph is given by first choosing any permutation of the nodes of the first level and then independently permuting the nodes of the subtrees. Consequently, in this case, it is clear that any such automorphism can be described as a pair $[\pi_1, \pi_2, \pi_3]; \sigma$ where π_i is in the symmetric group S_4 and σ is in the symmetric group S_3 . The group acts on $X_{(3,4)}$ by first permuting the level one nodes according to σ , leaving the subtrees

fixed, and then independently acting on the resulting subtrees by π_1, π_2 and π_3 respectively. The extension to more general trees X_m is straightforward.

Automorphisms of X_m induce permutations on the leaves of X_m . Notice that the leaves are indexed by paths from the root which naturally correspond to sequences of integers $(a_0, a_1, \dots, a_{m-1})$ with $0 \leq a_i < m_i$ indicating that at level i , the j_i^{th} branch is to be traversed. For example in $X_{(3,4)}$, $(2, 0)$ indexes the leftmost leaf of the rightmost subtree. For the purposes of this paper, leaves of the tree are meant to correspond to particular pixels in an image, each level of branching in the tree indicating another level of segmenting of the image. Notice that we are free to specify the particular method of indexing. For example, given an image - a square array of pixels - two possible choices are (1) A nested-grid structure; (2) A raster-scan structure. To illustrate, consider the particular example of a 256×256 image. In the former, the image is successively divided into 4×4 grids, with the subgrids in turn divided into 4×4 grids, and so on. Within 4×4 grid the points are ordered spirally into the center. In this way each point in the image corresponds to a leaf in the tree $X_{(16,16,16,16)}$. (There is of course nothing special about 4). For the raster-scan structure pixels are stored line by line, with the tree-indexing coming from taking successive subsets of lines, and then successive subsets of segments within lines. Our particular example in Section 4 uses this for a 512×512 image. It first divides the lines into four subsets, then each of these into 4 subsets and so on, continuing this theme on the level of lines, where they are divided into fourths, and fourths again, etc. This encodes each pixel as a leaf on the tree $X_{(4,4,4,4,4,4,4,4)}$.

Similarly, even after fixing the indexing structure there is still some freedom of choice in the automorphism group. Given X_m , we may choose any subgroup of the symmetric group S_{m_i} at level i . In particular, we are interested in the case in which the permutations at each level are **cyclic shifts**, although various situations may make other choices appropriate. We will denote by \mathbb{Z}_m the group obtained in this manner.




Figure 1. The spherically homogeneous tree of type $(3,4)$, denoted $X_{(3,4)}$.

3. FOURIER ANALYSIS FOR \mathbb{Z}_m ON X_m

Let L_m denote the vector space of complex-valued functions defined on the leaves of X_m . Notice that this is $m_0 \cdots m_{n-1}$ -dimensional. Here we outline the Fourier analysis for L_m , viewed as a homogeneous space for the group \mathbb{Z}_m . As such we will be working in the context of representation theory for the finite groups \mathbb{Z}_m (cf. [8] for a succinct introduction to the subject). The corresponding approach to data analysis follows the spectral analysis approach introduced by Diaconis (cf. [1]). Put briefly, in analogy with the commutative case, spectral analysis proceeds by consid-

ering the projections of data vectors onto group-invariant subspaces. We will see that in this setting L_m has a natural multiresolution structure.

Example 0: $m = m$. In this case this is the usual action of the cyclic group of order m on m points. Fourier analysis is equivalent to computation of the Fourier transforms

$$\hat{f}(j) = \sum_{k=0}^{m-1} f(k) e^{2\pi i j k / m}.$$

Example 1: $m = (m, n)$. This is our first nonabelian example. Elements of $\mathbb{Z}_{m,n}$ are denoted as pairs $[j; i]$ where $\mathbf{j} = (j_0, \dots, j_{m-1})$ with $0 \leq j_i < n$ and $0 \leq i < m$. Multiplication of group elements is defined by $[k; \ell] \cdot [j; i] = [k \cdot \ell(\mathbf{j}); i + j]$ where $\ell(\mathbf{j})_a = j_{a-\ell}$ and $(k \cdot \ell(\mathbf{j}))_a = k_a + \ell(\mathbf{j})_a$. Additions are performed modulo m and n as is appropriate. It is easy to see that this defines a noncommutative group.

The permutation action on the leaves is defined similarly. As described above, any leaf is denoted by a pair (a, b) and the action is defined by setting $[j; i](a, b) = (a + i, b_{j_{a+i}})$. Fourier analysis of the vector space $L_{(m,n)}$ proceeds by considering $L_{(m,n)}$ as a **representation space** for $\mathbb{Z}_{m,n}$. Elements of $\mathbb{Z}_{m,n}$ act on functions in $L_{(m,n)}$ by $(\alpha f)(x) = f(\alpha^{-1}x)$ for $\alpha \in \mathbb{Z}_{m,n}$. In this context, $L(m, n)$ has a direct sum decomposition into subspaces $L(m, n) = V_0 \oplus \cdots \oplus V_{m-1} \oplus W_1 \oplus \cdots \oplus W_{n-1}$ such that each V_a is one-dimensional and each W_b is m -dimensional and each V_a and W_b is $\mathbb{Z}_{m,n}$ -invariant (i.e. $[j; i]V_a = V_a$ and $[j; i]W_b = W_b$). Fourier analysis in this setting amounts to computing the projections of f onto these subspaces [1]. We now describe the equivalent computation (for a natural choice of basis for the subspaces which comes from representation theory).

Given a function $f \in L_{(m,n)}$ define the **Radon transform** $R : L_{(m,n)} \rightarrow L_{(m)}$ by

$$Rf(a) = \sum_{j=0}^{n-1} f(a, j).$$

Also, define **restriction maps** $S_a : L_{(m,n)} \rightarrow L_{(n)}$ by

$$S_a f(j) = f(a, j).$$

Then we have the following theorem.

Theorem 1 (1) For $0 \leq a < m$, projection of $f \in L_{(m,n)}$ onto V_a is equivalent to the computation of $\widehat{Rf}(a)$.

(2) For $1 \leq j < n$, projection of $f \in L_{(m,n)}$ onto W_j is equivalent to the computation of $\widehat{S_a f}(j)$, $0 \leq a < m$.

Consequently, the appropriate form of Fourier inversion rewrites f as the weighted sum of coarse-scale information (average over blocks) plus detail, given by restriction to the subtrees. This approach generalizes in a straightforward way to trees of arbitrary branching. Appropriately defined Radon transforms map functions from one scale to the next coarser scale, with restriction maps moving in the opposite direction. This yields a natural type of multiresolution analysis and in fact applied appropriately, gives another formulation of the usual Haar wavelet construction. It is not difficult to see that using appropriately scheduled FFT's,

the entire analysis of a function on $L_{(m)}$ can be accomplished in $O(|L_m| \log |L_m|)$ operations (cf. [7], esp. Section 5.1).

4. MULTIREOLUTION DECOMPOSITION

As a quick illustration, we consider here the Fourier decomposition of $f \in L_m$, where we have used the pixel indexing for a 512×512 image given by the raster-scan indexing described above. Thus, f is defined on the leaves of the tree X_m with height $n = 9$ and $m_i = 4, i = 1, \dots, 9$. Note that the sequence f viewed as a vector \mathbf{f} has dimension $512^2 \times 1$, L_m decomposes into the direct sum of 28 Z_m -invariant subspaces and all groups that comprise Z_m are cyclic and of order 4. The irreducible subspaces are of dimension 4^k , $k = 0, \dots, 8$ with 3 irreducible subspaces of each dimension, except for $k = 0$ in which case there are 4. For every such set of 3 subspaces, transform symmetry results in real and complex conjugate spectra. Using the natural choice of basis, the multiresolution spectrum \mathbf{v} can be seen to be generated from a one-dimensional transform $\mathbf{v} = \mathbf{U}\mathbf{f}$ where \mathbf{U} is unitary. Determination of \mathbf{v} follows Theorem 1.

For a vertically scanned raster image of ‘Lenna’, Figures 2 and 3 show respectively the magnitude and phase of the spectrum \mathbf{v} displayed with a similar scan.




Figure 2: Magnitude of raster-scan wreath product Fourier transform

Figure 4 shows an example of a reconstructed image, when all frequency bands except those corresponding to the 3 subspaces at the finest detail are set to zero.

5. CONVOLUTION

In this representation theoretic setting convolution of functions on the leaves will reflect the choice of symmetry group. For a tree of one level, which is the same as a sequence of




Figure 3: Phase of raster-scan wreath product Fourier transform

points, the usual cyclic convolution is obtained and is performed efficiently via the FFT. For higher branching trees the nonabelian nature of the symmetry group comes into play. Convolution for an arbitrary group is equivalent to multiplication in the associated group algebra. It is from this definition that the appropriate analogue of convolution of functions $f, g \in L_m$ is given by

$$f \star g(x) = \sum_{\alpha \in \mathbb{Z}_m} f(\alpha x_0) g(\alpha^{-1} x)$$

where x_0 denotes a fixed basepoint, say the leftmost leaf.

Reflecting the noncommutative nature of the situation, the Fourier transform of the convolution is not simply the pointwise product of the individual coefficients as in the commutative case. The following Theorem gives the result for a tree of height 2. The generalization to arbitrary spherically homogeneous trees is straightforward.

Theorem 2 (1) For $0 \leq a < m$, projection of $f \star g \in L_{(m,n)}$ onto V_a is equivalent to the computation of $\widehat{Rf}(a)\widehat{Rg}(a)$.

(2) For $1 \leq j < n$, projection of $f \star g \in L_{(m,n)}$ onto W_j is equivalent to the computation of $\widehat{S_a f}(j)\widehat{S_a g}(0)$, $0 \leq j < m$.

6. NUMERICAL EXPERIMENTS

Our preliminary numerical experiments confirmed our belief that properly designed Fourier analysis as described above may be more robust to the rotation than traditional matched filtering techniques. Our simple experiment compared the maximum output of a traditional matched filter, with the filter in the upper righthand portion of Figure 5.

We embedded the 8×4 templates shown in Figure 1 in a 256×256 array. The bottom 4 templates are rotations




Figure 4: Fine-scale reconstruction using raster-scan wreath product Fourier transform

of the original “T” in the upper left hand corner by odd multiples of 90 degrees and show the effect of discretization. We use the nested grid indexing explained above in which each point is naturally encoded by a leaf in $X_{(16,16,16,16)}$. The levels reflect successive subdivisions of the image into successions of 4×4 grids, with a clockwise spiral ordering on the blocks. Notice that with this ordering the effect of cyclic shift is to rotate the blocks and spiral them towards the center. Again, other orderings are possible.

We compared the outputs of the correlations of the template in the upper right against the other 7 rotations. With the templates taking on 0 and 1 values, the output of the traditional and non-traditional matched filter was 6 when the object function and the filter corresponded to the top right “T”. To test the rotational sensitivity of the traditional matched filter against the wreath product matched filter, we then “correlated” the top right filter against all 8 digitized rotations of the template which are shown in Figure 5.

These preliminary results lead us to believe that these structures exhibit more rotational sensitivity than a traditional Fourier based matched filter. In 5 of the 7 experiments, the new matched filter obtained a higher peak value than the traditional matched filter, with each filter obtaining the same value in the remaining 2 experiments. In all experiments where there was a difference, the new matched filter obtained a 5 and the old a 4, or 3 and 2, etc (the output is always integer). This experiment suggests, that optimally designed filtering processes which take advantage of these more general Fourier structures can be made to outperform traditional matched filters.

7. REFERENCES

[1] P. Diaconis. *Group Representations in Probabil-*




Figure 5: Above we illustrate the “T” template and its 8 rotations.

- ity and Statistics*, IMS, Hayward, CA (1988).
- [2] D. Eberly and D. Wenzel, Adaptation of group algebras to signal and image processing, *CVGIP: Graphical Models and Image Processing* **53**(4) (1991), 340-348.
- [3] R. Holmes, Signal processing on finite groups, *Technical Report 873, MIT, Lincoln Laboratory*, (1990).
- [4] R. Lenz, Using representations of the dihedral groups in the design of early vision filters, *Proc. ICASSP '93, Vol. 5* pp. 165-168.
- [5] R. Lenz, Group theoretical transforms and their applications in image coding, *Proc. ICASSP '95, Vol. 4* pp. 2355-2358.
- [6] D. Maslen and D. Rockmore, Generalized FFTs – a survey of recent results, submitted for publication, (1995).
- [7] D. Rockmore, Fast Fourier transforms for wreath products. *Applied and Computational Harmonic Analysis* **2** (1995), 279-292.
- [8] J. P. Serre, *Linear Representations of Finite Groups*, Springer-Verlag, New York, 1977.
- [9] L. Stiller, *Exploiting Symmetry on Parallel Architectures* Doctoral Dissertation, Dept, of Computer Science, Johns Hopkins University, (1995).
- [10] A. S. Willsky, On the algebraic structure of certain partially observable finite-state Markov processes. *Inform. Contr.* **38**, 179-212 (1978).
- [11] C. Wells. Some applications of the wreath product construction. *Amer. Math. Monthly* **83**, 317-338 (1976).