

The skew spectrum of functions on finite groups and their homogeneous spaces

Risi Kondor

`risi@gatsby.ucl.ac.uk`

Gatsby Unit for Computational Neuroscience
 University College London
 17 Queen Square, WC1N 3AR
 United Kingdom

Abstract

Whenever we have a group acting on a class of functions by translation, the bispectrum offers a principled and lossless way of representing such functions invariant to the action. Unfortunately, computing the bispectrum is often costly and complicated. In this paper we propose a unitarily equivalent, but easier to compute set of invariants, which we call the skew spectrum. For functions on homogeneous spaces the skew spectrum can be efficiently computed using some ideas from Clausen-type fast Fourier transforms.

1 Introduction

Given a function $f: \mathbb{R} \rightarrow \mathbb{C}$ with Fourier transform \hat{f} , it is well known that the bispectrum

$$\hat{q}(k_1, k_2) = \hat{f}(k_1) \hat{f}(k_2) \hat{f}(k_1 + k_2) \quad k_1, k_2 \in \mathbb{R} \quad (1)$$

is invariant to translations of f . The bispectrum is used in various signal and image processing applications which involve translation invariance.

Kakarala generalized (1) to signals on a wide class of non-commutative groups, including all compact, and thus, all finite groups. On a group G the (left-) translate of $f: G \rightarrow \mathbb{C}$ is defined $f^t(x) = f(t^{-1}x)$, and the bispectrum takes the form

$$\hat{q}(\rho_1, \rho_2) = (\hat{f}(\rho_1) \otimes \hat{f}(\rho_2))^{\dagger} C_{\rho_1 \rho_2} \left[\bigoplus_{\rho \in \mathcal{R}} \hat{f}(\rho)^{\oplus m_{\rho}} \right] C_{\rho_1 \rho_2}^{\dagger}, \quad (2)$$

where ρ_1, ρ_2 and ρ are irreducible representations of G , and $C_{\rho_1 \rho_2}$ are the Clebsch-Gordan matrices. Recall that the Fourier transform on a non-commutative group is a collection of matrices indexed by the irreducible representations. The remarkable fact shown by Kakarala is that for a wide class of groups the bispectrum is not only invariant to left translation, but, assuming that each Fourier component $\hat{f}(\rho)$ is invertible, the $\hat{q}(\rho_1, \rho_2)$ matrices also uniquely determine f up to translation [Kakarala, 1992, Theorem 3.2.4]. In other words, given $f, f': G \rightarrow \mathbb{C}$ and their respective bispectra $(\hat{q}(\rho_1, \rho_2))_{\rho_1, \rho_2}$ and $(\hat{q}'(\rho_1, \rho_2))_{\rho_1, \rho_2}$, we have $q(\rho_1, \rho_2) = q'(\rho_1, \rho_2)$ if and only if $f'(x) = f(t^{-1}x)$ for some $t \in G$.

Unfortunately, the bispectrum can be difficult to compute, even on finite groups of moderate size. The total number of entries in the $\widehat{q}(\rho_1, \rho_2)$ matrices is $|G|^2$, which can be a problem. Another issue is the difficulty of deriving the form of the Clebsch Gordan matrices for non-trivial groups. Finally, computing (2) is expensive because it involves the multiplication of large matrices. These factors all limit the practical applicability of the bispectrum method.

In this paper we introduce the **skew spectrum**, which furnishes an alternative set of invariants. The skew spectrum can be computed by Fourier transformation alone, without any explicit Clebsch-Gordan transforms. We show that the skew spectrum is unitarily equivalent to the bispectrum. In particular, the skew spectrum is sufficient to uniquely determine the original function f up to translation.

An important generalization of the above is when f is not a function on G itself, but a function on a homogeneous space of G . This case is likely to be more relevant to applications than the original case of functions on G . The new invariants can be restricted to a wide range of homogeneous spaces, and are closely adapted to Clausen-type FFTs, offering computational economy.

2 Fourier transforms and the bispectrum

In this paper G denotes a finite group of cardinality $|G|$ and \mathcal{R} denotes a complete set of inequivalent irreducible complex-valued matrix representations of G . Since G is finite, \mathcal{R} is a finite set. We use the letter ρ to denote the individual irreducibles, and d_ρ to denote their dimensionalities. Thus, each $\rho \in \mathcal{R}$ is a homomorphism $\rho: G \rightarrow \mathbb{C}^{d_\rho \times d_\rho}$. Without loss of generality, we assume that all $\rho \in \mathcal{R}$ are unitary, i.e., $\rho(x^{-1}) = (\rho(x))^{-1} = \rho(x)^\dagger$.

We denote the set of complex valued functions on G by $L(G)$. Given $f \in L(G)$ and a group element t , the **left-translate** of f by t is the function $f^t \in L(G)$ defined $f^t(x) = f(t^{-1}x)$. A functional $p: f \mapsto p(f)$ is said to be **invariant** to left translation (or just **left-invariant** for short) if $p(f) = p(f^t)$ for any $f \in L(G)$ and any $t \in G$. A set $P = \{p_1, p_2, \dots, p_k\}$ of such invariants is said to be **complete** if $p_1(f) = p_1(f'), \dots, p_k(f) = p_k(f')$ implies $f' = f^t$ for some $t \in G$. The objective of this paper is to find an efficiently computable yet complete set of left-invariant functionals.

2.1 The Fourier transform

The **Fourier transform** of a function $f \in L(G)$ is the collection of matrices

$$\widehat{f}(\rho) = \sum_{x \in G} \rho(x) f(x) \quad \rho \in \mathcal{R}.$$

With respect to the norms

$$\|f\|^2 = \frac{1}{|G|} \sum_{x \in G} |f(x)|^2 \quad \text{and} \quad \|\widehat{f}\|^2 = \frac{1}{|G|^2} \sum_{\rho \in \mathcal{R}} d_\rho \|\widehat{f}(\rho)\|_2^2,$$

(where $\|M\|_2$ denotes the Frobenius norm $(\sum_{i,j=1}^n [M]_{i,j})^{1/2}$) the transform $\mathfrak{F}: f \mapsto \widehat{f}$ is unitary, and its inverse is given by

$$f(x) = \frac{1}{|G|} \sum_{\rho \in \mathcal{R}} d_\rho \operatorname{tr}(\rho(x) \cdot \widehat{f}(\rho)) \quad x \in G.$$

Defining the **convolution** of two functions $f, g \in L(G)$ as

$$(f * g)(x) = \sum_{y \in G} f(xy^{-1}) g(y),$$

it is easy to verify the convolution theorem $\widehat{f * g}(\rho) = \widehat{f}(\rho) \cdot \widehat{g}(\rho)$. Another property of interest is the behavior of the Fourier transform under translation:

$$\widehat{f^t}(\rho) = \sum_{x \in G} \rho(x) f(t^{-1}x) = \sum_{x \in G} \rho(tx) f(x) = \rho(t) \sum_{x \in G} \rho(x) f(x) = \rho(t) \cdot \widehat{f}(\rho). \quad (3)$$

Finally, defining $f^- \in L(G)$ by $f^-(x) = f(x^{-1})$, and letting $*$ denote complex conjugation, we have $\widehat{f^-}(\rho)^* = \widehat{f}(\rho)^\dagger$.

In general, we say that a matrix-valued functional $F: L(G) \rightarrow \mathbb{C}^{D \times d_\rho}$ is ρ -**covariant** if $F(f^t) = \rho(t) \cdot F(f)$ for any $f \in L(X)$ and any $t \in G$. Similarly, $F': L(X) \rightarrow \mathbb{C}^{d_\rho \times d_\rho}$ is ρ -**contravariant** if $F'(f^t) = F'(f) \cdot \rho(t)^\dagger$. Clearly, the adjoint (conjugate transpose) of a ρ -covariant functional is ρ -contravariant.

2.2 The power spectrum

If F is ρ -covariant and F' is ρ -contravariant, then $Q(f) = F'(f) F(f)$ is invariant, i.e., $Q(f^t) = Q(f)$. Thus, the matrices $\widehat{a}(\rho) = \widehat{f}(\rho)^\dagger \widehat{f}(\rho)$ are a natural starting point for finding left-invariant functionals. On closer inspection, the invariance of \widehat{a} should not come as a surprise, since by the convolution theorem \widehat{a} is the Fourier transform of the **autocorrelation**

$$a(x) = \sum_{y \in G} f^-(xy^{-1})^* f(y) = \sum_{y \in G} f(yx^{-1})^* f(y),$$

which is manifestly invariant to left-translation. We call $(\widehat{a}(\rho))_{\rho \in \mathcal{R}}$ the **power spectrum** of f .

Unfortunately, the matrix elements of $(\widehat{a}(\rho))_{\rho \in \mathcal{R}}$ fall short of forming a complete set of invariants. This is because the spectrum is insensitive to the relative “phase” of the various Fourier components. The natural way to couple the components is to take tensor products, forming higher order spectra.

2.3 Triple correlation and the bispectrum

Recall that for any $\rho_1, \rho_2 \in \mathcal{R}$, the tensor product representation $\rho_1 \otimes \rho_2$ decomposes into irreducible components in the form

$$(\rho_1 \otimes \rho_2)(x) = \rho_1(x) \otimes \rho_2(x) = C_{\rho_1 \rho_2} \left[\bigoplus_{\rho \in \mathcal{R}} \rho(x)^{\oplus m_\rho} \right] C_{\rho_1 \rho_2}^\dagger, \quad (4)$$

where $C_{\rho_1 \rho_2}$ is a unitary matrix called the **Clebsch-Gordan** matrix, $m_\rho = m(\rho_1, \rho_2, \rho) \in \mathbb{Z}^+$ is the multiplicity of ρ in the decomposition, and $M^{\oplus k}$ is a shorthand for $\bigoplus_{i=1}^k M$.

Now consider the tensor product of the Fourier matrices $\widehat{f}(\rho_1)$ and $\widehat{f}(\rho_2)$. Under translation

$$\widehat{f}(\rho_1) \otimes \widehat{f}(\rho_2) \mapsto \widehat{f^t}(\rho_1) \otimes \widehat{f^t}(\rho_2) = (\rho_1(t) \otimes \rho_2(t)) \cdot (\widehat{f}(\rho_1) \otimes \widehat{f}(\rho_2)),$$

hence

$$C_{\rho_1\rho_2}^\dagger (\widehat{f}^t(\rho_1) \otimes \widehat{f}^t(\rho_2)) C_{\rho_1\rho_2} = \left[\bigoplus_{\rho \in \mathcal{R}} \rho(t)^{\oplus m_\rho} \right] C_{\rho_1\rho_2}^\dagger (\widehat{f}(\rho_1) \otimes \widehat{f}(\rho_2)) C_{\rho_1\rho_2}.$$

We say that $C_{\rho_1\rho_2}^\dagger (\widehat{f}(\rho_1) \otimes \widehat{f}(\rho_2)) C_{\rho_1\rho_2}$ is $\bigoplus_\rho \rho^{\oplus m_\rho}$ -covariant. In particular, the **bispectrum**

$$\widehat{b}(\rho_1, \rho_2) = (\widehat{f}(\rho_1) \otimes \widehat{f}(\rho_2))^\dagger C_{\rho_1\rho_2} \left[\bigoplus_{\rho \in \mathcal{R}} \widehat{f}(\rho)^{\oplus m_\rho} \right] C_{\rho_1\rho_2}^\dagger \quad \rho_1, \rho_2 \in \mathcal{R} \quad (5)$$

is invariant to left-translations of f . The remarkable fact proved by Kakarala for compact groups in general is that provided that all $\widehat{f}(\rho)$ are invertible, the matrix elements of $(\widehat{b}(\rho_1, \rho_2))_{\rho_1, \rho_2 \in \mathcal{R}}$ are not only invariant, but also complete [Kakarala, 1992, Theorem 3.2.4][Kakarala, 1993].¹ It is possible to construct even higher order spectra, but there is no need: the bispectrum is already complete.

It is natural to ask what the analog of the autocorrelation is, i.e., what invariant function the bispectrum is the Fourier transform of. To answer this question we consider functions on the direct product group $G \times G$. Recall from general representation theory that the irreducible representations of $G \times G$ are the tensor products of the irreducible representations of G . Thus,

$$\widehat{f}(\rho_1) \otimes \widehat{f}(\rho_2) = \sum_{x_1 \in G} \sum_{x_2 \in G} (\rho_1(x_1) \otimes \rho_2(x_2)) f(x_1) f(x_2)$$

in (5) is the Fourier transform of $u(x_1, x_2) = f(x_1) f(x_2)$, while

$$C_{\rho_1\rho_2} \left[\bigoplus_{\rho \in \mathcal{R}} \widehat{f}(\rho)^{\oplus m_\rho} \right] C_{\rho_1\rho_2}^\dagger = \sum_{x \in G} C_{\rho_1\rho_2} \left[\bigoplus_{\rho \in \mathcal{R}} \widehat{\rho}(x)^{\oplus m_\rho} \right] C_{\rho_1\rho_2}^\dagger f(x) = \sum_{x \in G} (\rho_1(x) \otimes \rho_2(x)) f(x)$$

is the Fourier transform of

$$v(x_1, x_2) = \begin{cases} f(x_1) & \text{if } x_1 = x_2 \\ 0 & \text{otherwise.} \end{cases}$$

Thus, by the convolution theorem (on $G \times G$), (5) is the Fourier transform of $u^- * v$, where $u^-(x_1, x_2) = u(x_1^{-1}, x_2^{-1})$. In long hand, the bispectrum is the Fourier transform of

$$\begin{aligned} \sum_{y_1 \in G} \sum_{y_2 \in G} u^-(x_1 y_1^{-1}, x_2 y_2^{-1})^* v(y_1, y_2) &= \sum_{y_1 \in G} \sum_{y_2 \in G} u(y_1 x_1^{-1}, y_2 x_2^{-1})^* v(y_1, y_2) \\ &= \sum_{y \in G} f(y x_1^{-1})^* f(y x_2^{-1})^* f(y) = b(x_1, x_2), \end{aligned} \quad (6)$$

which is known as the **triple correlation** of f . Just as the bispectrum, provided that all Fourier matrices are invertible, the triple correlation also furnishes a complete set of left-invariant functionals of f . In fact, Kakarala derives the concept of bispectrum from the triple correlation, and not the other way round.

¹Kakarala defines the bispectrum in a slightly different form, equivalent to (5) up to unitary transforms and complex conjugation.

2.4 Computational considerations

The price to pay for the completeness of the bispectrum and the triple correlation is their inflated size. Letting $|G|$ denote the cardinality of G , the triple correlation consists of $|G|^2$ scalars. By the unitarity of the Fourier transform the total number of entries in the bispectrum matrices is the same. The symmetry $b(x_1, x_2) = b(x_2, x_1)$ and its counterpart $\widehat{b}(\rho_1, \rho_2) = \widehat{b}(\rho_2, \rho_1)$ (up to reordering of rows and columns) reduces the number of relevant components to $|G|(|G| + 1)/2$, but to keep the analysis as simple as possible we ignore this constant factor. Another technical detail that we ignore for now is that if all we need is a complete set of invariants, then we can use the matrices $\widehat{b}(\rho_1, \rho_2) C_{\rho_1 \rho_2}$ instead of $\widehat{b}(\rho_1, \rho_2)$, sparing us the cost of multiplying by $C_{\rho_1 \rho_2}^\dagger$ in (5).

The cost of computing the bispectrum involves two factors: the cost of the Fourier transform, and the cost of multiplying the various matrices in (5). Recent years have seen the emergence of fast Fourier transforms for a series of non-commutative groups, including the symmetric group [Clausen, 1989], wreath product groups [Rockmore, 1995], and others [Maslen and Rockmore, 1997]. Typically these algorithms reduce the complexity of Fourier transformation (and inverse Fourier transformation) from $O(|G|^2)$ to $O(|G| \log^k |G|)$ scalar operations for some small integer k . In the following we assume that for whatever group we are working on, an $O(|G| \log^k |G|)$ transform is available. This makes matrix multiplication the dominant factor in the cost of computing the bispectrum.

The complexity of the naive approach to multiplying two D -dimensional matrices is D^3 . Using the well known fact that $\sum_{\rho \in \mathcal{R}} d_\rho^2 = |G|$ (which can be regarded as a corollary to the Fourier transform being an invertible linear map), the cost of computing (5) for given (ρ_1, ρ_2) is $O(d_{\rho_1}^3 d_{\rho_2}^3)$. Summing over all $\rho_1, \rho_2 \in \mathcal{R}$ and assuming that $\sum_{\rho \in \mathcal{R}} d_\rho^3$ grows with some power $1 < \theta < 2$ of $|G|$ (note that even $\sum_{\rho \in \mathcal{R}} d_\rho^4$ grows with at most $|G|^2$) gives an overall complexity bound of $O(|G|^{2\theta})$. Thus, the critical factor in computing the bispectrum is matrix multiplication and not the fast Fourier transform.

As for the triple correlation, (6) involves an explicit sum over G , which is to be computed for all $|G|^2$ possible values of (x_1, x_2) , giving a total time complexity is $O(|G|^3)$.

One should note that when unique identifiability of f is not an absolute necessity, it is possible to truncate \widehat{f} according to some “low pass filtering” scheme in the interest of computational efficiency. This is easy to implement in the bispectrum, but not in the triple correlation.

3 The skew spectrum

The computational demands of the triple correlation and the bispectrum are of serious concern in applications. The former involves an explicit summation over G for each (x_1, x_2) pair, while the latter involves multiplying together large matrices and the non-trivial issue of computing Clebsch Gordan coefficients. In this section we derive a third, unitarily equivalent, set of invariants, which, in some sense are a combination of the two.

Our starting point is the collection of functions $r_z: G \rightarrow \mathbb{C}$ indexed by $z \in G$ and defined

$$r_z(x) = f(x) f(xz).$$

Introducing the concept of **left z -diagonal slice** $g \downarrow_z(x) = g(x, zx)$ and **right z -diagonal slice** $g \uparrow_z(x) = g(x, xz)$ of a general function $g: G \times G \rightarrow \mathbb{C}$, and recalling our previous definition

$u(x_1, x_2) = f(x_1) f(x_2)$, we may write $r_z = u|_z$. The $\widehat{r}_z(\rho)$ Fourier components are ρ -covariant (with respect to the action of G on f), since

$$\widehat{r}_z^t(\rho) = \sum_{x \in G} \rho(x) f(t^{-1}x) f(t^{-1}xz) = \sum_{x \in G} \rho(tx) f(x) f(xz) = \rho(t) \cdot \widehat{r}_z(\rho).$$

This immediately gives rise to the left-translation invariant matrices

$$\widehat{q}_z(\rho) = \widehat{r}_z(\rho)^\dagger \cdot \widehat{f}(\rho) \quad \rho \in \mathcal{R}. \quad (7)$$

The collection of matrices $\{\widehat{q}_z(\rho)\}_{\rho \in \mathcal{R}, z \in G}$ we call the **skew spectrum** of f .

To see that the skew spectrum and the bispectrum are unitarily equivalent, it is sufficient to observe that by the convolution theorem

$$\begin{aligned} q_z(x) &= \sum_{y \in G} r_z^-(xy^{-1})^* f(y) = \sum_{y \in G} r_z(yx^{-1})^* f(y) \\ &= \sum_{y \in G} f(yx^{-1})^* f(yx^{-1}z)^* f(y) = b(x, z^{-1}x) = b|_{z^{-1}}(x), \end{aligned}$$

and that $\{b|_z\}_{z \in G}$ together make up the triple correlation b . We thus have the following theorem.

Theorem 1 *Let f and f' be complex valued functions on a finite group G and let \mathcal{R} be a complete set of inequivalent irreducible representations of G . Assume that $\widehat{f}(\rho)$ is invertible for all $\rho \in \mathcal{R}$. Let the skew spectrum \widehat{q}_z be defined as in (7). Then $f' = f^t$ for some $t \in G$ if and only if $\widehat{q}'_z(\rho) = \widehat{q}_z(\rho)$ for all $z \in G$ and all $\rho \in \mathcal{R}$.*

From a computational perspective, the skew spectrum involves $|G| + 1$ separate Fourier transforms followed by $|G|$ sequences of multiplying $\{\mathbb{C}^{d_\rho \times d_\rho}\}_{\rho \in \mathcal{R}}$ matrices. Using the notations of the previous section, the cost of the former is $O(|G|^2 \log^k |G|)$, while the cost of the latter is $O(|G|^{\theta+1})$, improving on both the triple correlation and the bispectrum.

4 Homogeneous spaces

Real world problems often involve functions on homogeneous spaces of groups as opposed to functions on the groups themselves. Recall that a **homogeneous space** of a finite group G is a set S on which G acts by $s \mapsto x(s)$ in such a manner that for any given $s_0 \in S$, $\{x(s_0) \mid x \in G\}$ sweeps out the entire set. The group elements stabilizing s_0 form a subgroup H , making it possible to identify any $s \in S$ with some coset xH . Thus, S itself is identified with the **quotient space** G/H of left cosets of H in G . A **transversal** for G/H is a set containing exactly one group element from each xH coset. By abuse of notation we use the symbol G/H for the transversals as well as the quotient space. The right cosets $\{Hx\}_{x \in G}$ also form a quotient space (with respect to the right action of G) and this we denote $H \backslash G$. We examine functions on such spaces in the next section.

We denote the space of complex valued functions on $S = G/H$ by $L(G/H)$. Any $f \in L(G/H)$ extends naturally to a **right H -invariant** function $f \uparrow^G$ on G by $f \uparrow^G(x) = |H|^{-1} f(xH)$. Conversely, any $g \in L(G)$ may be restricted to $L(G/H)$ by $g \downarrow_{G/H}(xH) = \sum_{h \in H} g(xh)$. In accordance with

this correspondence, for $z \in G$ we define the **left-translate** of $f \in L(G/H)$ as the function $f^z \in L(G/H)$ given by $f^z(xH) = f(z^{-1}xH)$, and the Fourier transform of f as

$$\widehat{f}(\rho) = \sum_{x \in G} \rho(x) f(xH) = \left[\sum_{x \in G/H} \rho(x) f(xH) \right] \cdot \left[\sum_{h \in H} \rho(h) \right]. \quad (8)$$

Clearly, $\widehat{f}(\rho) = \widehat{f \uparrow G}(\rho)$, and $\widehat{f^z}(\rho) = \rho(z) \cdot \widehat{f}(\rho)$ remains valid for functions on G/H .

We denote by $\rho \downarrow_H$ the restricted representations $\rho: H \rightarrow \mathbb{C}^{d_\rho \times d_\rho}$ given by $\rho \downarrow_H(h) = \rho(h)$. While ρ is always an irreducible representation of G , in general $\rho \downarrow_H$ is reducible, i.e.,

$$\rho \downarrow_H(h) = U_\rho \left[\bigoplus_{\eta} \eta(h) \right] U_\rho^\dagger \quad h \in H, \quad (9)$$

where U_ρ is a unitary matrix and η runs over some well-defined subset of irreducible representations of H , possibly with repeats. If $U_\rho = I$ we say that ρ is **adapted** to H or, equivalently, that it is expressed in a **Gelfand-Tsetlin** basis with respect to G/H . In the following we assume that each $\rho \in \mathcal{R}$ is not only unitary but also H -adapted. It is possible to show that one may always choose \mathcal{R} so as to satisfy these conditions.

For the trivial representation $\eta_{\text{triv}}(h) = 1$ we have $\sum_h \eta_{\text{triv}}(h) = |H|$, hence by the unitarity of the Fourier transform over H for any other irreducible we must have $\sum_h \eta(h) = 0$. Plugging (9) (with $U_\rho = I$) back in (8) then shows that the Fourier transform of functions in $L(G/H)$ have a very special form: only those columns of the $\widehat{f}(\rho)$ matrices are non-zero which correspond to the trivial representation of H in (9).

The consequences of this type of sparsity are two-fold. On the one hand, except for the trivial case $H = e$, the sparsity implies that the non-singularity condition required for Kakarala's completeness result and derived results such as Theorem 1 are *always* violated when working over homogeneous spaces. The bispectrum, triple correlation and skew spectrum remain useful invariants, but we can no longer trust that they uniquely determine the original function up to translation. On the other hand, the sparsity suggests that the invariants can be computed much faster than when we were working over the full group G .

Letting n_ρ be the multiplicity of the trivial representation in the summation on the right hand side of (9), each bispectrum component (with the final $C_{\rho_1 \rho_2}^\dagger$ omitted)

$$\widehat{b}(\rho_1, \rho_2) = (\widehat{f}(\rho_1) \otimes \widehat{f}(\rho_2))^\dagger C_{\rho_1 \rho_2} \bigoplus_{\rho \in \mathcal{R}} \widehat{f}(\rho)^{\oplus m_\rho} \quad (10)$$

will only have $n_{\rho_1} n_{\rho_2}$ non-zero rows and a similarly small number of non-zero columns.

Since $y(hx)^{-1}H = yx^{-1}hH = yx^{-1}H$ for any $h \in H$, the triple correlation

$$b(x_1, x_2) = \sum_{y \in G} f(yx_1^{-1}H)^* f(yx_2^{-1}H)^* f(yH)$$

will be a left $H \times H$ -invariant function, i.e., effectively $b \in L(H \backslash G \times H \backslash G)$. The summation still extends over all of $y \in G$, however, giving a total complexity of $|G|^3 / |H|^2$.

As for the skew spectrum, the first fact to note is that the z index can be restricted to one element from each $HxH = \{h_1 x h_2 \mid h_1, h_2 \in H\}$ **double coset**. A transversal for such double cosets we denote $H \backslash G / H$.

Theorem 2 Let H be subgroup of a finite group G and let $f \in L(G/H)$. Then the skew spectrum \widehat{q}_f is uniquely determined by its subset of components $\{ \widehat{q}_z(\rho) \mid \rho \in \mathcal{R}, z \in H \setminus G/H \}$.

Proof. For any $h \in H$,

$$r_{zh}(x) = f(xzhH)f(xhH) = f(xzH)f(xH) = r_z(x),$$

so $\widehat{q}_{zh}(\rho) = \widehat{r}_z(\rho)^\dagger \widehat{f}(\rho) = \widehat{r}(\rho)^\dagger \widehat{f}(\rho) = q_z(\rho)$. Now for $h' \in H$

$$r_{h'z}(x) = f(xh'zH)f(xH) = f(xh'zH)f(xh'H) = r_z^{(h'-1)}(x),$$

where $r_z^{(t)}$ denotes the right-translate $r_z^{(t)}(x) = r_z(xt^{-1})$. Thus, by the right-translation property

$$\widehat{f^{(t)}}(\rho) = \sum_{x \in G} \rho(x) f(xt^{-1}) = \sum_{x \in G} \rho(xt) f(x) = \left[\sum_{x \in G} \rho(x) f(x) \right] \rho(t) = \widehat{f}(\rho) \cdot \rho(t)$$

of the Fourier transform,

$$\widehat{q}_{h'z}(\rho) = (\widehat{r}_z(\rho) \rho(h'^{-1}))^\dagger \widehat{f}(\rho) = \rho(h') \widehat{r}_z(\rho)^\dagger \widehat{f}(\rho) = \rho(h') \widehat{q}_z(\rho),$$

so $\widehat{q}_{h'z}(\rho)$ and $\widehat{q}_z(\rho)$ albeit not equal, are related by an invertible linear mapping. \blacksquare

What the best way to compute $\widehat{q}_z(\rho)$ is strongly depends on G and H . One possible algorithm is based on the decomposition

$$\widehat{f}(\rho) = \sum_{y \in G/H} \rho(y) \bigoplus_{\eta} \sum_{h \in H} \eta(h) f(yh) = \sum_{y \in G/H} \rho(y) \bigoplus_{\eta} \widehat{f}_y(\eta)$$

where the f_y are functions on H defined $f_y(h) = f(yh)$ and \widehat{f}_y are their Fourier transforms

$$\widehat{f}_y(\eta) = \sum_{h \in H} \eta(h) f_y(h).$$

This is akin to a Clausen-type ‘‘matrix separation of variables’’ fast Fourier transform (FFT) tailored to the $G > H > 1$ chain. Because $f \in L(G/H)$, all but the $\widehat{f}_y(\eta_{\text{triv}})$ component of \widehat{f}_y vanish, and \widehat{f} can be computed efficiently by restricting the FFT to the descendants of the $\widehat{f}_y(\eta_{\text{triv}})$ components (see the section on ‘‘partial Fourier transforms’’ in the documentation to **SnOB** [Kondor, 2006] for details). Letting g be the right-translated function $g(x) = f^{(z^{-1})}(x) = f(xz)$ we can then compute $\widehat{g}(\rho) = \widehat{f}(\rho) \rho(z^{-1})$ for each ρ and run the FFT backwards to get $\widehat{g}_y(\eta)$ for each $y \in G/H$. Multiplying each of these by the corresponding $f(y)$ yields $\widehat{g}_y(\eta) f(y)$, which is exactly the Fourier transform of r_z restricted to yH , so one more forward transform yields $\widehat{r}_z(\rho)$, which we can plug directly into (7). The exact complexity of this procedure depends on the groups G and H .

5 Right invariance

The **right-translate** of $f: G \rightarrow \mathbb{C}$ by $z \in G$ is defined $f^{(z)}(x) = f(xz^{-1})$, and a functional $p: f \mapsto p(f)$ is said to be **right-invariant** if $p(f^{(t)}) = p(f)$ for any f and any $t \in G$. Everything that we described in sections 2 and 3 have natural right-invariant analogs. The Fourier transform obeys $\widehat{f}^{(t)}(\rho) = \widehat{f}(\rho) \rho(t)$ and the corresponding invariant power spectrum is $\widehat{q}^R(\rho) = \widehat{f}(\rho) \widehat{f}(\rho)^\dagger$, which is the Fourier transform of the “right-autocorrelation” $q^R(x) = \sum_{y \in G} f(y) f(xy)$. Similarly, the “right-bispectrum” is

$$\widehat{b}^R(\rho_1, \rho_2) = (\widehat{f}(\rho_1) \otimes \widehat{f}(\rho_2)) C_{\rho_1 \rho_2} \left[\bigoplus_{\rho \in \mathcal{R}} \widehat{f}(\rho)^\dagger \right] C_{\rho_1 \rho_2}^\dagger,$$

which is the Fourier transform of $b^R(x_1, x_2) = \sum_{y \in G} f(x_1 y) f(x_2 y) f(y)$. For the skew spectrum we must change r_z to $r_z^R(x) = u|_z(x) = f(x) f(zx)$, and let $\widehat{q}_z^R(\rho) = \widehat{r}_z^R(\rho) \widehat{f}(\rho)^\dagger$, which is the Fourier transform of $q_z^R = \sum_{y \in G} f(xy) f(zxy) f(y)$. The rest of the analysis goes through exactly as in the left-invariant case.

Generalizing section 4 is more interesting, because our choice of G/H over $H \backslash G$ breaks the left-right symmetry. For notational convenience, instead of constructing right-invariant functionals of $f \in L(G/H)$, we discuss the analogous problem of left-invariant functionals of $f \in L(H \backslash G)$.

Left-translating $f \in L(H \backslash G)$ by $t \in H$ leaves it invariant, effectively reducing the set of transformations of interest from G to G/H . In this sense we are in a easier situation than in section 4. Indeed, despite that the Fourier matrices of f are just as rank deficient as in the previous section, in this case Kakarala’s completeness result is salvagable.

Theorem 3 [Kakarala, 1992, Theorem 3.3.6] *Let H be any closed subgroup of a compact group G , and let $f \in L_1(H \backslash G)$ be such that for all $\rho \in \mathcal{R}$, the matrix rank of $\widehat{f}(\rho)$ is equal to the multiplicity of η_{triv} in the decomposition of ρ into irreducible representations of H . Then $\widehat{b}_f = \widehat{b}_{f'}$ for some $f' \in L_1(G)$ if and only if there exists some $t \in G$ such that $f' = (f \uparrow^G)^t$.*

As a corollary, the skew spectrum will all so be complete. However, the skew spectrum is a much larger object than before because (a) there is no obvious way to restrict z like in Theorem 2 (b) instead of being columns-sparse, the Fourier matrices are row-sparse, hence as long as the decomposition of $\rho \downarrow_H$ contains η_{triv} with multiplicity at least one, $\widehat{q}_z(\rho)$ will be a full $d_\rho \times d_\rho$ matrix.

6 Conclusions

The bispectrum is an elegant way of constructing invariants of functions on finite groups with respect to translation, but from a computational point of view it is not necessarily attractive. We discussed the theory behind the bispectrum from a new angle, and arrived at a unitarily equivalent set of invariants, which we named the skew spectrum. Not only is the skew spectrum easier to compute, it also meshes naturally with the structure of Clausen-type fast Fourier transforms, promising efficient specialized algorithms for computing it on specific groups or homogeneous spaces of groups.

Acknowledgements

I would like to thank Ramakrishna Kakarala for providing me with a hard copy of his thesis.

References

M. Clausen. Fast generalized Fourier transforms. *Theor. Comput. Sci.*, pages 55–63, 1989.

R Kakarala. *Triple corelation on groups*. PhD thesis, Department of Mathematics, UC Irvine, 1992.

R. Kakarala. A group theoretic approach to the triple correlation. In *IEEE Workshop on higher order statistics*, pages 28–32, 1993.

Risi Kondor. \mathbb{S}_n ob: a C++ library for fast Fourier transforms on the symmetric group, 2006. Available at <http://www.cs.columbia.edu/~risi/Snob/>.

D. Maslen and D. Rockmore. Generalized FFTs — a survey of some recent results. In *Groups and Computation II*, volume 28 of *DIMACS Ser. Discrete Math. Theor. Comput. Sci.*, pages 183–287. AMS, Providence, RI, 1997.

D. Rockmore. Fast fourier transforms for wreath products. *J. Applied and Computational Harmonic Analysis*, 2:279–292, 1995.