

Invariants

Risi Kondor

A classical problem

Let $g \in \mathbb{R}$, and $f: \mathbb{R} \rightarrow \mathbb{C}$. The real numbers, as a group acts on the space of functions on \mathbb{R} by translation:

$$g: f \mapsto f' \quad \text{where} \quad f'(x) = f(x - g).$$

Question: How do we construct functionals $\Upsilon[f]$ that are invariant to this action, i.e., for which $\Upsilon[f] = \Upsilon[f']$ for any f and any g ?

Many applications in signal processing, image analysis, etc..

Solution 1: Autocorrelation

The **autocorrelation** of f is

$$a(x) = \int f(x + y)f(y)dy.$$

Tells us how much f changes when we translate it by an amount y . Clearly invariant to translation:

Solution 2: Power spectrum

The **power spectrum** of f is

$$\hat{a}(\omega) = \int \hat{f}(\omega)^* \hat{f}(\omega) dx = \int |\hat{f}(\omega)|^2 dx.$$

Literally measures the amount of energy in each Fourier mode. Clearly invariant to translation:

$$\begin{aligned} \hat{a}_{f'}(\omega) &= \int (e^{2\pi i \omega g} \hat{f}(\omega))^* (e^{2\pi i \omega g} \hat{f}(\omega)) dx = \\ &= \int \hat{f}(\omega)^* \hat{f}(\omega) dx = \hat{a}_f(\omega). \end{aligned}$$

- In fact, the power spectrum is just the Fourier transform of the autocorrelation.
- Their common limitation: lose all the information in the phase.

3: Triple correlation & bispectrum

The **triple correlation** of f is

$$b(x_1, x_2) = \int f(y - x_1) f(y - x_2) f(y) dy$$

The **bispectrum** of f is

$$\widehat{b}(k_1, k_2) = \widehat{f}(k_1)^* \widehat{f}(k_2)^* \widehat{f}(k_1 + k_2).$$

Reconstructing f from b

$$\widehat{b}(k_1, k_2) = \widehat{f}(k_1)^* \widehat{f}(k_2)^* \widehat{f}(k_1 + k_2).$$

Use the following algorithm to recover f from \widehat{b} :

1. $\widehat{f}(0) = \widehat{b}(0, 0)^{1/3}$
2. $\widehat{f}(1) = e^{i\phi} \sqrt{\widehat{b}(0, 1)/\widehat{f}(0)}$ \rightarrow indeterminacy in ϕ inevitable
- 3.

$$\widehat{f}(k+1) = \frac{\widehat{b}(1, k)}{\widehat{f}(1)^* \widehat{f}(k)^*} \quad k = 2, 3, \dots$$

- If $\widehat{f}(k) \neq 0$ for any k , then \widehat{b} uniquely determines \widehat{f} up to translation.

Invariants on groups

General setup

- G is a group acting on a set X transitively. This means that each $g \in G$ is a map $g: X \rightarrow X$, sending

$$g: x \mapsto gx.$$

- $L(X)$ is a space of functions on X . The action of G on X extends to functions in $L(X)$ by

$$g: f \mapsto f^g \quad \text{where} \quad (f^g)(x) = f(g^{-1}x).$$

(Assume that $gf \in L(X)$ for all $f \in L(X)$ and $g \in G$.)

- A functional $\Upsilon[f]$ is an **invariant** to this action if

$$\Upsilon[f] = \Upsilon[f^g] \quad \forall f \in L(X), \quad \forall g \in G.$$

Examples

- The rotation group $SO(3)$ acts on the sphere S^2 by $x \mapsto Rx$. Consider $L_2(S^2)$...
- The symmetric group acts on the adjacency matrix of a graph by $(i, j) \mapsto (\sigma(i), \sigma(j))$.

The Fourier Transform

Restrict ourselves for now to finite X and finite G . Recall that f induces a function $f \uparrow^G (g) = f(gx_0)$, and the Fourier transform of f is

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

Moreover, if $f^t(g) = f(t^{-1}g)$, then

$$\widehat{f^t}(\rho) = \rho(t) \cdot \widehat{f}(\rho).$$

Noncommutative power spectrum

The **power spectrum** of a function $f: X \rightarrow \mathbb{C}$ is

$$\hat{a}(\rho) = \hat{f}(\rho)^\dagger \cdot \hat{f}(\rho).$$

Clearly invariant because

$$\hat{f}^\tau(\rho)^\dagger \cdot \hat{f}^\tau(\rho) = (\rho_\rho(t) \cdot \hat{f}(\rho))^\dagger (\rho_\rho(t) \cdot \hat{f}(\rho)) = \hat{f}(\rho)^\dagger \cdot \hat{f}(\rho).$$

The power spectrum is the FT of the (flipped) autocorrelation function

$$a(h) = \sum_{g \in G} f(gh^{-1})f(g).$$

The noncommutative bispectrum

Recall the Clebsch-Gordan decomposition

$$\rho_1(\sigma) \otimes \rho_2(\sigma) = C_{\rho_1, \rho_2} \left[\bigoplus_{\rho \in R_{\rho_1, \rho_2}} \bigoplus_{i=1}^{c(\rho_1, \rho_2, \rho)} \rho(\sigma) \right] C_{\rho_1, \rho_2}^\dagger.$$

The bispectrum:

$$\widehat{b}_f(\rho_1, \rho_2) = C_{\rho_1, \rho_2}^\dagger \left[\widehat{f}(\rho_1) \otimes \widehat{f}(\rho_2) \right]^\dagger C_{\rho_1, \rho_2} \bigoplus_{\rho \in \Lambda_{\rho_1, \rho_2}} \bigoplus_{i=1}^{c(\rho_1, \rho_2, \rho)} \widehat{f}(\rho)$$

The bispectrum is the FT of the **triple correlation**

$$b(h_1, h_2) = \sum_{g \in G} f(gh_1^{-1}) f(gh_2^{-1}) f(g).$$

Completeness result

Theorem [Kakarala, 1992]. Let f and f' be a pair of complex valued integrable functions on a compact group G . Assume that $\widehat{f}(\rho)$ is invertible for each $\rho \in \mathcal{R}$. Then $f' = f^z$ for some $z \in G$ if and only if $b_f(\rho_1, \rho_2) = b_{f'}(\rho_1, \rho_2)$ for all $\rho_1, \rho_2 \in \mathcal{R}$.

- Generalizes to any Tatsuuma duality group ($ISO(n)$, inhomog. Lorentz Group, etc.)

The skew spectrum

The **skew spectrum** of $f: \mathbb{S}_n \rightarrow \mathbb{C}$ is the collection of matrices

$$\hat{q}_h(\rho) = \hat{r}_h^\dagger(\rho) \cdot \hat{f}(\rho), \quad \rho \in \mathcal{R}_G, \quad \hat{\in} G,$$

with $r_h(g) = f(gh)f(g)$.

Unitarily equivalent to the bispectrum, but sometimes easier to compute [K., 2007]

Optimization problems

Risi Kondor

The Quadratic Assignment Problem

Given $A, A' \in \mathbb{R}^{n \times n}$, the **Quadratic Assignment Problem** is to solve


$$\underset{\sigma \in \mathbb{S}_n}{\text{maximize}} f(\sigma) \quad f(\sigma) = \sum_{i,j=1}^n A_{\sigma(i),\sigma(j)} A'_{i,j}.$$

Equivalently,

$$\underset{\sigma \in \mathbb{S}_n}{\text{maximize}} \text{tr}(P_\sigma A P_\sigma^\top A') \quad [P_\sigma]_{i,j} = \begin{cases} 1 & \text{if } \sigma(j) = i \\ 0 & \text{else.} \end{cases}$$

- Can interpret A and A' as the adjacency matrices of two (weighted) graphs.
- The QAP aims to find the “best way to match” A to A' .
- Most graph-to-graph matching problems reduce to special cases of the QAP: traveling salesman, (sub)-graph isomorphism, graph partitioning, etc.
- \rightarrow NP-hard. Also hard in practice. No PTAS until recently.
- Sometimes written in minimization form.

Classical approach: branch&bound



- Subdivide in the form of a tree based on where vertex 1 is mapped, where vertex 2 is mapped, etc.. Each node corresponds to a coset $\tau_{i_1, \dots, i_k} \mathbb{S}_{n-k}$.
- Define bounds

$$B_{i_1, \dots, i_k} \geq \max_{\sigma \in \tau_{i_1, \dots, i_k} \mathbb{S}_{n-k}} f(\sigma).$$

- Do a directed depth-first search: follow most promising branch, go down to leaf, then backtrack, but never follow branches which are guaranteed to be worse than optimum so far.
- Hard to come up with theoretical performance guarantees. Empirical performance depends critically on the tightness of the bounds.

The Fourier approach to the QAP

Fundamental Observation

Consider the Fourier transform of the objective function:

$$\hat{f}(\lambda) = \sum_{\sigma \in \mathbb{S}_n} f(\sigma) \rho_\lambda(\sigma).$$

Theorem [Rockmore et al., 2002]. If f is the QAP objective function, then $\hat{f}(\lambda) = 0$ unless $\lambda \in \left\{ \begin{array}{c} \square\square\square\square\square\square\square\square \\ \square\square\square\square\square\square\square\square \\ \square\square\square\square\square\square\square\square \\ \square\square\square\square\square\square\square\square \end{array} \right\}$.

Proposition [K., 2010]. $\hat{f}(\begin{array}{c} \square\square\square\square\square\square\square\square \\ \square\square\square\square\square\square\square\square \end{array})$ and $\hat{f}(\begin{array}{c} \square\square\square\square\square\square\square\square \\ \square\square\square\square\square\square\square\square \end{array})$ are rank one matrices.

Question: Why is this?

Graph Functions

\mathbb{S}_n acts transitively on the off-diagonal entries of A , so it induces a function

$$g_A(\sigma) = A_{\sigma(n), \sigma(n-1)}$$

with the property $\pi \circ (g_A) = g_{(\pi \circ A)}$ (note $g_A(\sigma\tau) = g_A(\sigma) \forall \tau \in \mathbb{S}_{n-2}$).

Graph correlation

Proposition. The QAP objective function can be written in the form

$$f(\sigma) = \frac{1}{(n-2)!} \sum_{\tau \in \mathbb{S}_n} g_A(\sigma\tau) g_{A'}(\tau).$$

Proof.

$$\begin{aligned} \sum_{i,j=1}^n A_{\sigma(i),\sigma(j)} A'_{i,j} &= \frac{1}{(n-2)!} \sum_{\pi \in \mathbb{S}_n} A_{\sigma\pi(n),\sigma\pi(n-1)} A'_{\pi(n),\pi(n-1)} = \\ &= \frac{1}{(n-2)!} \sum_{\pi \in \mathbb{S}_n} g_A(\sigma\pi) g_{A'}(\pi). \end{aligned}$$

Graph correlation

Proposition. The QAP objective function can be written in the form

$$f(\sigma) = \frac{1}{(n-2)!} \sum_{\tau \in \mathbb{S}_n} g_A(\sigma\tau) g_{A'}(\tau).$$

Corollary. The Fourier transform of the QAP objective is of the form

$$\widehat{f}(\lambda) = \frac{1}{(n-2)!} \widehat{g}_A(\lambda) \cdot (\widehat{g}_{A'}(\lambda))^\top, \quad \lambda \vdash n. \quad (1)$$


In particular, $\widehat{f}(\lambda)$ is identically zero unless $\lambda = (n), (n-1, 1), (n-2, 2)$ or $(n-2, 1, 1)$, and $\widehat{f}((n-2, 2))$ and $\widehat{f}((n-2, 1, 1))$ are dyadic (rank one) matrices.

Question: Is this useful for anything?

Fourier space branch & bound

- First compute \widehat{g}_A and $\widehat{g}_{A'}$ $\rightarrow O(n^2)$ space, $O(n^3)$ time.
- Compute \widehat{f} . $\rightarrow O(n^2)$ space, $O(n^3)$ time.
- Do branch and bound on the same coset tree as before.
- Compute bounds directly in Fourier space.

Restriction to cosets in Fourier space



$$\hat{f}_i(\lambda) = \sum_{\lambda' \in \lambda^n} \frac{d_{\lambda'}}{n\lambda} \left[\rho_{\lambda}(in)^{\top} \cdot \hat{f}(\lambda) \right]_{\lambda'}$$

Fourier space bounds

Proposition. For any $f: \mathbb{S}_n \rightarrow \mathbb{R}$

$$\max_{\sigma \in \mathbb{S}_n} f(\sigma) \leq \frac{1}{n!} \sum_{\lambda \vdash n} d_\lambda \|\widehat{f}(\lambda)\|_{\text{tr}}.$$

Proof.

$$f(\sigma) = \frac{1}{n!} \sum_{\lambda \vdash n} d_\lambda \text{tr}[\widehat{f}(\lambda) \cdot \rho_\lambda(\sigma)^{-1}]$$

For any orthogonal matrix O , $\text{tr}(MO) \leq \|M\|_{\text{tr}}$.

- This bound is computed in time $O(n^3)$ (requires an SVD).