Inverse Problems with Symmetry: A Functional Perspective on Multi-Reference Alignment

Anna Little

Department of Mathematics, Utah Center for Data Science
University of Utah

ICERM Workshop Bayesian Inverse Problems and UQ - **March 2, 2026**

NSF 2309570, NSF CAREER 2441153

Outline

I. Introduction

1. Biological motivation
2. MRA models

II. Dilation MRA

1. Functional formulation with other group actions
2. Recovery from third moments

III. Inversion Guarantees

1. Functional formulation of simplest model
2. Recovery from second moments

Collaborators


Matthew Hirn


Liping Yin


Mikhail Sweeney

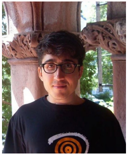
Daniel Sanz-Alonso


Omar Al-Ghattas

Motivation: Cryo-Electron Microscopy

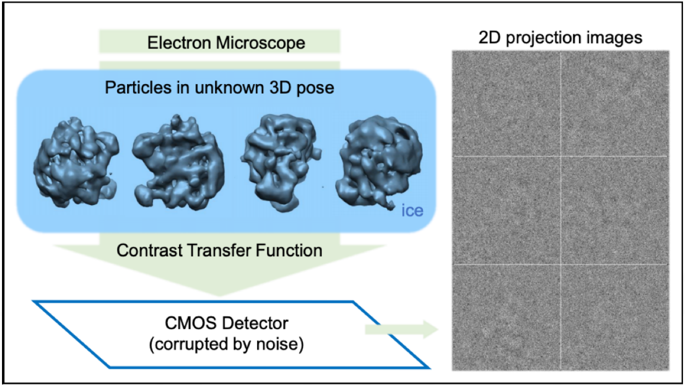

Fan & Zhao 2019

Cryo-EM Imaging Challenges

1. Measurement noise $\implies$ very low signal-to-noise ratio (*SNR*)
2. Random rigid motions
3. Random deformations of the molecule

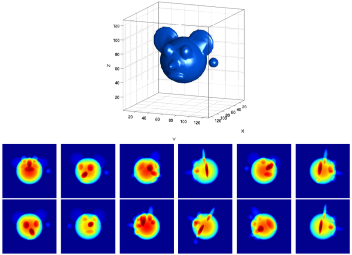

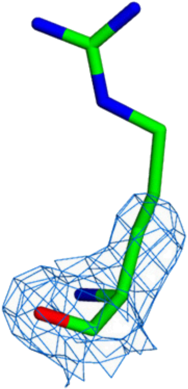

Left: Singer 2018. Right: Palamini, Canciani, & Forneris 2016

Basic set-up of interest:

- **Given:** large number of noisy observations of a target image/signal
- **Goal:** recover a high resolution approximation of the target

"Noise" can mean additive noise or random group action (translation, rotation, rescaling); helpful to utilize **group invariant features**.

Discrete Multi-reference Alignment (MRA)

Set-up: multiple noisy measurements of hidden signal **vector** $f \in \mathbb{R}^L$.

$$\boxed{y_n[i] = f[i - x_n] + \epsilon_n[i], \qquad 1 \leq i \leq L, \qquad 1 \leq n \leq N}$$

Where:

- $L \longleftarrow$ length of hidden signal (dimension)
- $N \longleftarrow$ number of observed signals (sample size)
- $\epsilon \longleftarrow$ Gaussian noise vector ($\eta \sim \text{Normal}(0, \sigma^2 I_L)$)
- $\rho \longleftarrow$ shift distribution of the x_n

Note:

- Translation is cyclic.
- Alternate formulation: f is bandlimited function (finite dimensional vector in Fourier domain)

$$\text{snr} := \frac{\|f\|_2^2}{\sigma^2} \text{ governs difficulty of recovery}$$

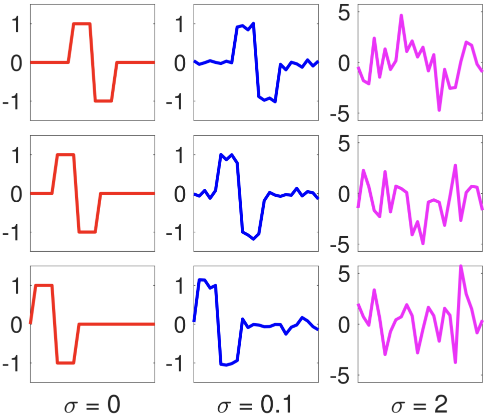

Abbe, Bendory, Leeb, Pereira, Sharon, & Singer 2019

Approaches for Solving Discrete MRA

- **Synchronization**
  First align the signals then average. Fails in low *SNR* regime.
- **Method of Moments**
  Recover f from moments of y ($\mathbb{E}[y]$, $\mathbb{E}[y \otimes y]$, $\mathbb{E}[y \otimes y \otimes y]$).
- **Method of Invariants**
  Recover f from translation-invariant features, i.e. *Fourier invariants*.
- **Expectation Maximization Algorithms**
  Good results but expensive.

Fourier Invariants

Observation: if $y(t) = f(t-x)$, then $y^{\mathrm{ft}}(\omega) = e^{i\omega x} f^{\mathrm{ft}}(\omega)$

Idea: construct Fourier features which "cancel" out the phase factors

For example:

$$(Py)(\omega) = \text{power spectrum} = |y^{\mathrm{ft}}(\omega)|^2$$

$$(By)(\omega_1, \omega_2) = \text{bispectrum} = y^{\mathrm{ft}}(\omega_1) y^{\mathrm{ft}}(-\omega_2) y^{\mathrm{ft}}(\omega_2 - \omega_1)$$

$$\implies Py = Pf, By = Bf$$

Note: the bispectrum is invertible under general conditions (unlike Py).

Bendory, Boumal, Ma, Zhao, Singer (2017); Perry, Weed, Bandeira, Rigollet, Singer (2017)

Fourier Invariants $\subseteq$ Method of Moments

$$\mathbb{E}\left[y^{\mathrm{ft}}(\omega_1)\overline{y^{\mathrm{ft}}(\omega_2)}\right] = f^{\mathrm{ft}}(\omega_1)f^{\mathrm{ft}}(-\omega_2)\rho^{\mathrm{ft}}(\omega_1-\omega_2)$$

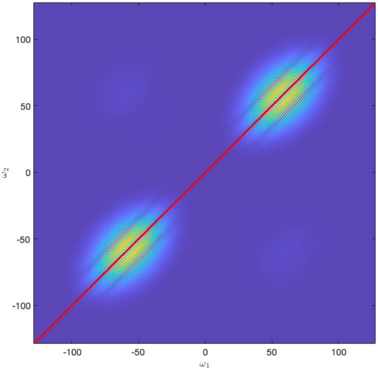

$\rightarrow$ *Plug in* $\omega_2=\omega_1$

$= \text{Power Spectrum}$

$= |f^{\mathrm{ft}}(\omega_1)|^2$

Fourier Invariants $\subseteq$ Method of Moments

$$\mathbb{E}\left[y^{\mathrm{ft}}(\omega_1)\,\overline{y^{\mathrm{ft}}(\omega_2)}\,y^{\mathrm{ft}}(\omega_3)\right] = f^{\mathrm{ft}}(\omega_1)\,f^{\mathrm{ft}}(-\omega_2)\,f^{\mathrm{ft}}(\omega_3)\,\rho^{\mathrm{ft}}(\omega_1-\omega_2+\omega_3)$$

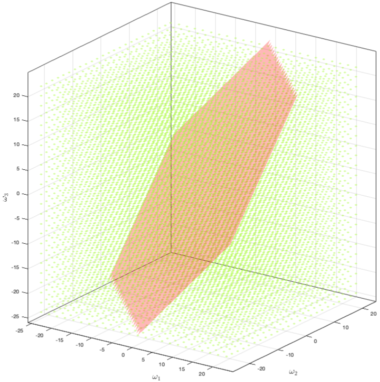

$\rightarrow$ *Plug in* $\omega_3 = \omega_2 - \omega_1$

$= \text{Bispectrum}$

$= f^{\mathrm{ft}}(\omega_1) f^{\mathrm{ft}}(-\omega_2) f^{\mathrm{ft}}(\omega_2 - \omega_1)$

Sample Complexity

Case 1: translation distribution is *aperiodic*

- Sufficient to estimate first and second moments ($\mathbb{E}[y]$, $\mathbb{E}[y \otimes y]$)
- $N \gtrsim \mathrm{snr}^{-2}$ samples needed

Case 2: translation distribution is *periodic* (for example uniform)

- Requires estimation of third moment tensor ($\mathbb{E}[y \otimes y \otimes y]$)
- $N \gtrsim \mathrm{snr}^{-3}$ samples needed

Note: above sample complexity is minimax optimal unless additional structure imposed (e.g. sparsity).

But... some things remain unsatisfying

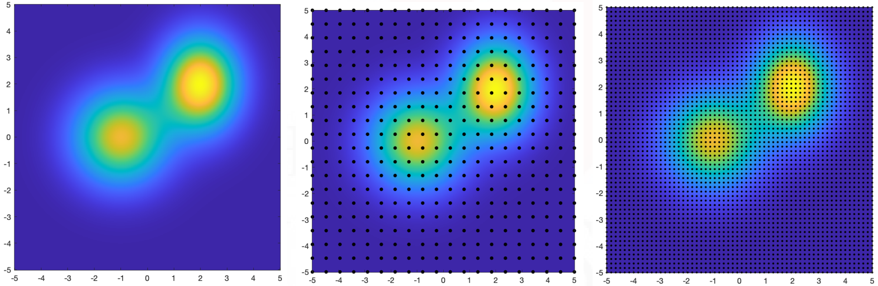

- In motivating applications, f is **function** (not finite dimensional)
- As $L \to \infty$, things should get better (increased resolution)

Taking $L \to \infty$ for Discrete MRA is Challenging

Issues:

- Noise $\epsilon \to$ white noise process (sample paths not defined, $\epsilon \notin L^2$)
- When $L \to \infty$, recovery guarantees can only hold if $\sigma \to 0$; see [1, 2]
- Assumption that consecutive noise values are uncorrelated becomes unrealistic.

Our approach: reformulate MRA in the *functional* regime (remove L).

[1] Romanov, Bendory, & Ordentlich 2021; [2] Dou, Fan, & Zhou 2024

II. Dilation MRA

1. Functional formulation with other group actions
2. Recovery from third moments

Why Consider Dilations?

Previous work:

Mainly focused on random rigid motions

Our focus:

General diffeomorphisms $f(\tau(t))$ (nonrigid motions)

Motivation:

Flexible regions of molecular structures

Simplest case:

Linear diffeomorphisms $\rightarrow$ translations $+$ dilations

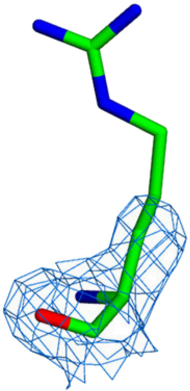

Functional Dilation MRA (1d)

Goal: Recover $f : [-1,1] \to \mathbb{R}$ from N noisy, shifted, dilated observations:

$$\boxed{y_n(t) = \frac{1}{1-\tau_n} f\left(\frac{t-x_n}{1-\tau_n}\right) + \epsilon_n(t)\,, \quad t \in [-1,1]\,, \quad n = 1,\dots,N}$$

Assumptions:

- Observed signals y_n are supported on $[-1,1]$.
- Dilation factors τ_n have uniform distribution with variance η^2.
- Each ϵ_n is a white noise process with variance σ^2.

Functional Dilation MRA (1d)

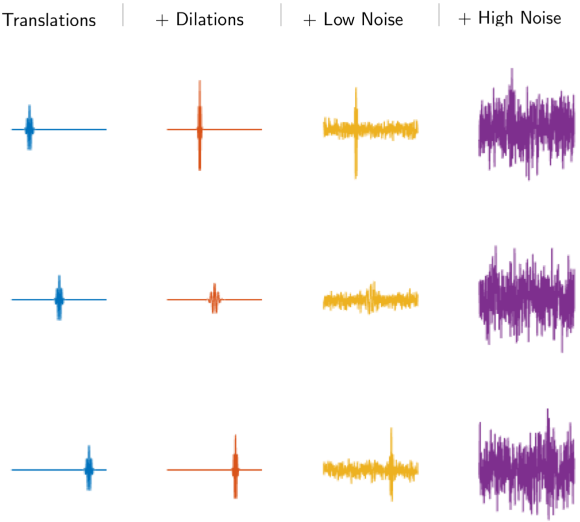

Functional Dilation MRA (2d)

How do dilations affect the Fourier Invariants?

Observation: dilations smear/diffuse the invariants

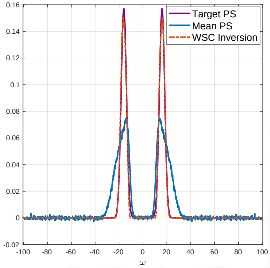

PS (truth + observed)

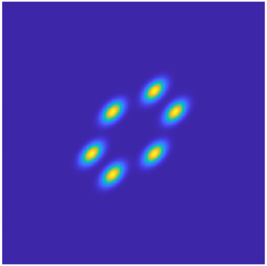

BS (truth)

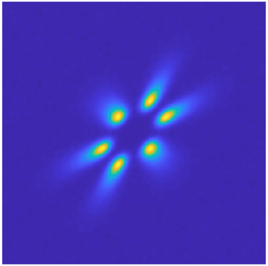

BS (observed)

Strategy: invariants can be recovered from $\widetilde{Bf} = \mathcal{D}\left(\overline{By}\right)$, where $\mathcal{D}$ is known operator and $\overline{By}$ is the empirical mean of the bispectra.

Theoretical Guarantees

Theorem (Hirn, Little, Yin, 2024)

Consider the functional dilation MRA model. Under mild conditions on f, on a finite domain Ω:

$$\mathbb{E}\left[\|Pf - \widetilde{Pf}\|^2_{L^2(\Omega)}\right] \leq C_{f,\Omega}\left(\frac{\eta^2}{N} + \left(\frac{\sigma^4}{N}\right)^{\frac{2}{3}}\right)$$

$$\mathbb{E}\left[\|Bf - \widetilde{Bf}\|^2_{L^2(\Omega)}\right] \leq C_{f,\Omega}\left(\frac{\eta^2}{N} + \left(\frac{\sigma^6}{N}\right)^{\frac{2}{3}}\right)$$

- We obtain *unbiased* estimators of Pf, Bf.
- Estimating Bf with accuracy δ requires $N \gtrsim \eta^2\delta^{-2} + \sigma^6\delta^{-3}$ samples.

Example: $f(t) = e^{-5t^2}\cos(12t)$, $N = 10,000$ $(\eta = 12^{-1/2},\ \sigma = 1.0)$

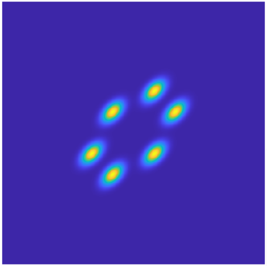

Real(Bf)

(ground truth)

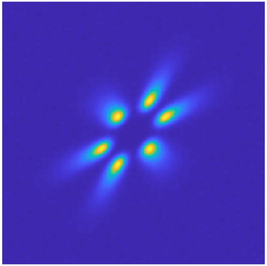

Real$(\overline{By})$

(observed data)

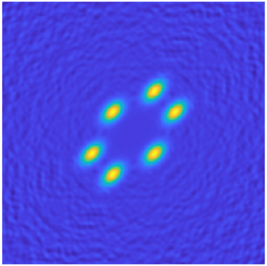

Real$(\widetilde{Bf})$

(recovered approx)

Inverting the Bispectrum

Since we can accurately learn Bf, to recover f we simply need to *invert the bispectrum*.

Algorithms for inverting the bispectrum include:

- non-convex optimization over the manifold of phases
- iterative phase synchronization
- phase unwrapping
- frequency marching
- spectral methods
- tensor decomposition algorithms

See Bendory, Boumal, Ma, Zhao, & Singer, *IEEE TSP*, 2017 for overview.

Challenges

Challenge 1: Unlike PS, **BS does not preserve the norm** of the signal.

- For high frequency signals, there are lots of cancellations.
- Figure: Stability of PS vs BS recovery for $f^{\mathrm{ft}}(\omega) = e^{-(\omega-\mu)^2}$ $(N = 10000,\ \sigma = .5,\ \eta = 12^{-\frac{1}{2}})$

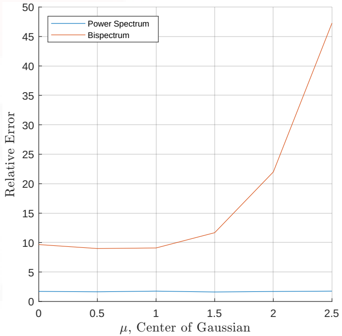

Challenges

Challenge 2: Theoretical guarantees for inversion of **continuous** bispectra are lacking in the literature.

Exception: when f is **bandlimited**, can be recovered from discrete vector of Fourier series frequencies; see:

1. Dou, Fan, and Zhou, *Annals of Statistics*, 2024
2. Perry, Weed, Bandeira, Rigollet, and Singer, *SIMODS*, 2019

Bispectrum and Hidden Signal Recovery

Example: $f(t) = e^{-5t^2}\cos(12t)$

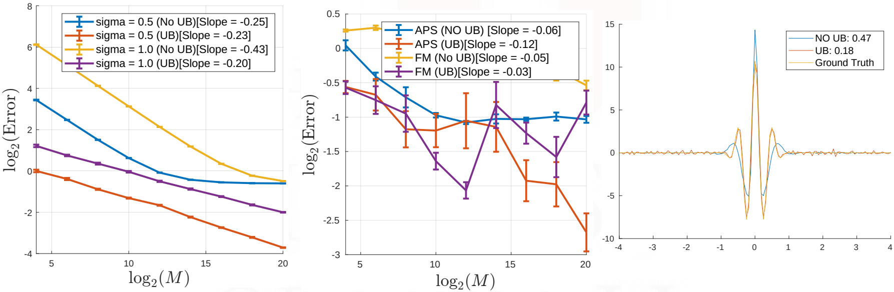

(a) Bispectrum error (b) Reconstruction error (c) Recovered signal

Key Takeaways for Dilation MRA

- We consider a generalized MRA problem which includes random translation, random dilation, and additive noise.
- One of the first works considering action of a non-compact group.

Key Question: Can we obtain theoretical guarantees for inversion of moments/Fourier invariants in the *functional* regime?

III. Inversion Guarantees

1. Functional formulation of simplest model
2. Recovery from second moments

Functional Multi-Reference Alignment

Goal: Recover $f : [-1,1]^d \to \mathbb{R}$ from N noisy, shifted observations:

$$\boxed{y_n(t) = f(t - x_n) + \epsilon_n(t)\,, \quad t \in [-1,1]^d, \quad n = 1, \ldots, N}$$

Assumptions:

- Observed signals y_n are supported on $[-1,1]^d$.
- The shifts x_n are iid with distribution ρ.
- Each ϵ_n is a centered Gaussian process with known covariance k_ϵ.

Connecting MRA & Deconvolution

MRA Model:

$$\begin{cases} \text{Space:} & y(t) = f(t-x) + \text{ noise} \\ \text{Frequency:} & y^{\text{ft}}(\omega) = e^{-i\omega x} f^{\text{ft}}(\omega) + \text{ noise} \end{cases}$$

$x \sim \rho$

(shift distribution)

Second Moments (noise = 0):

$$\begin{aligned} \mathbb{E}[y^{\text{ft}}(\omega_1)\overline{y^{\text{ft}}(\omega_2)}] &= \int e^{-i\omega_1 x} f^{\text{ft}}(\omega_1) e^{i\omega_2 x} f^{\text{ft}}(-\omega_2)\rho(x)\; dx \\ &= f^{\text{ft}}(\omega_1) f^{\text{ft}}(-\omega_2)\rho^{\text{ft}}(\omega_1 - \omega_2) \\ &:= \Psi(\omega_1, \omega_2) \end{aligned}$$

Connecting MRA & Deconvolution

Repeated Measurements Model:

$$\begin{cases} Y_1 = x + \epsilon_1 \\ Y_2 = x + \epsilon_2 \end{cases} \qquad x \sim \rho, \quad \epsilon_1, \epsilon_2 \overset{\text{iid}}{\sim} f$$

Second Moments:

$$\mathbb{E}[e^{-i\omega_1 Y_1 + i\omega_2 Y_2}] = \int\int\int e^{-i\omega_1(x+y)+i\omega_2(x+z)}\rho(x)f(y)f(z)\,dx\,dy\,dz$$

$$= \left(\int e^{-i\omega_1 y} f(y)\,dy\right)\left(\int e^{i\omega_2 z} f(z)\,dz\right)\left(\int e^{-i(\omega_1-\omega_2)x}\rho(x)\,dx\right)$$

$$= f^{\text{ft}}(\omega_1) f^{\text{ft}}(-\omega_2)\rho^{\text{ft}}(\omega_1 - \omega_2)$$

$$= \Psi(\omega_1, \omega_2) \leftarrow \textit{Exact same object as in MRA!}$$

Can f be Recovered from Second Moments Ψ?

Theorem (Kotlarski's Identity)

Let $f, \rho \in L^2([-1,1]^d)$ and assume ρ is pdf of a centered random variable and that f is bounded, $f^{ft}(\omega) \neq 0$ for all ω, and $f^{ft}(0) = 1$. Then:

$$f^{ft}(\omega) = \exp\left(\int_0^1 \frac{\nabla_1 \Psi(\alpha\omega, \alpha\omega)}{\Psi(\alpha\omega, \alpha\omega)} \cdot \omega \, d\alpha\right).$$

Note:

- Kotlarski's identity for $d = 1$ is well known in the econometrics literature.
- Different versions can target recovery of either f or ρ from Ψ.
- Generalization to general d is straight-forward (though not published).

Kotlarski Visualization ($d = 1$)

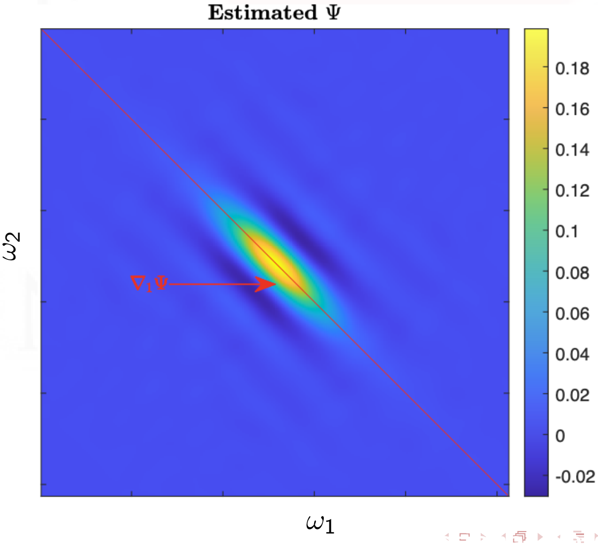

Solving Functional MRA with Kotlarski

We can thus solve functional MRA with the following pipeline:

1. Estimate Ψ from the given data.
2. Estimate f^{ft} using Kotlarski's Identity.
3. Estimate the signal f via deconvolution.

Accuracy will be governed by:

$$\boxed{\text{snr} := \frac{\|f\|_\infty^2}{\|k_\epsilon\|_\infty^{1/2}\mathbb{E}[\sup_{t\in D}\epsilon(t)]}}$$

where $k_\epsilon(t_1, t_2) = \mathbb{E}[\epsilon(t_1)\epsilon(t_2)]$ is the auto-covariance of the noise.

Algorithm 2.1 MRA via Deconvolution

1: **Input:** Data $\{y_n\}_{n=1}^N$; Covariance function k_η; Kernel K; Bandwidth parameter $h > 0$.

2: **Estimate Ψ:** Set

$$\widehat{\Psi}(\omega,\omega) := \frac{1}{N}\sum_{n=1}^{N} y_n^{\mathrm{ft}}(\omega)\overline{y_n^{\mathrm{ft}}(\omega)} - k_\eta^{\mathrm{ft}}(\omega,\omega), \qquad \widetilde{\Psi}(\omega,\omega) := \frac{\widehat{\Psi}(\omega,\omega)}{1 \wedge \sqrt{N}|\widehat{\Psi}(\omega,\omega)|}.$$

3: **Estimate f^{ft}:** Set

$$\widehat{f}^{\mathrm{ft}}(\omega) := \exp\left(\int_0^1 \frac{\nabla_1 \widehat{\Psi}(\alpha\omega,\alpha\omega)}{\widetilde{\Psi}(\alpha\omega,\alpha\omega)}\cdot\omega \, d\alpha\right).$$

4: **Deconvolve:** Set

$$\widehat{f}(t) := \frac{1}{(2\pi)^d}\int_{\mathbb{R}^d} e^{i\omega\cdot t}\widehat{f}^{\mathrm{ft}}(\omega)K^{\mathrm{ft}}(h\omega)\,d\omega.$$

5: **Output**: Approximation $\widehat{f}$ to the hidden signal f.

Theoretical Guarantees

Theorem (High Noise Regime)

Assume $f^{\mathrm{ft}}(\omega)$ *decays like* $\|\omega\|^{-\beta}$ *and that* $\|k_\epsilon\|_\infty^{1/2} \geq \|f\|_\infty$. *Then for* N *large enough and optimal bandwidth* h,

$$\frac{\|f^{\mathrm{ft}} - f\|_\infty}{\|f\|_\infty} \lesssim \left(\frac{\sqrt{\tau}\|k_\epsilon\|_\infty \vee \mathrm{snr}^{-1}}{\sqrt{N}} \right)^{\frac{\beta-d}{4\beta+1}} .$$

with probability at least $1 - c_1 e^{-c_2\tau}$.

Comments:

- $N \gtrsim \mathrm{snr}^{-2}$ samples needed (since second moments sufficient)
- Theoretical convergence rate is $O(N^{-(\frac{\beta-d}{8\beta+2})})$, i.e. $O(N^{-\frac{1}{8}})$ as $\beta \to \infty$

Impact of Noise on Sample Complexity

Example: squared exponential kernel

$$k_\epsilon^{\mathrm{SE}}(t_1, t_2) = \sigma^2 \exp\left(-\frac{\|t_1 - t_2\|_2^2}{2\lambda^2}\right) \quad , \quad \begin{cases} \sigma & = \text{size of noise} \\ \lambda & = \text{correlation length-scale} \end{cases}$$

In terms of σ, λ, the sample size requirement $N \gtrsim \mathrm{snr}^{-2}$ becomes

$$N \gtrsim \sigma^4 \quad , \quad N \gtrsim \log \lambda^{-d}$$

Note: discrete MRA requires $N \gtrsim \sigma^6$ or $N \gtrsim \sigma^4$ depending on ρ

Experiments on Synthetic Signals

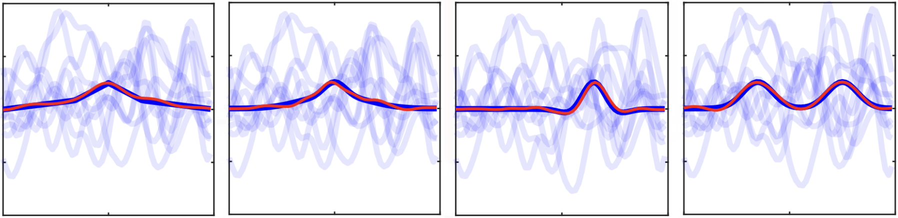

(a) f_1, $\beta = 2$ **(b)** f_2, $\beta = 4$ **(c)** f_3, SS **(d)** f_4, SS

Figure 2: True signal (thick blue), estimated signal (red), & noisy observations (shadowed blue) with $\sigma = 1$, $\lambda = 0.1$

Error Decay

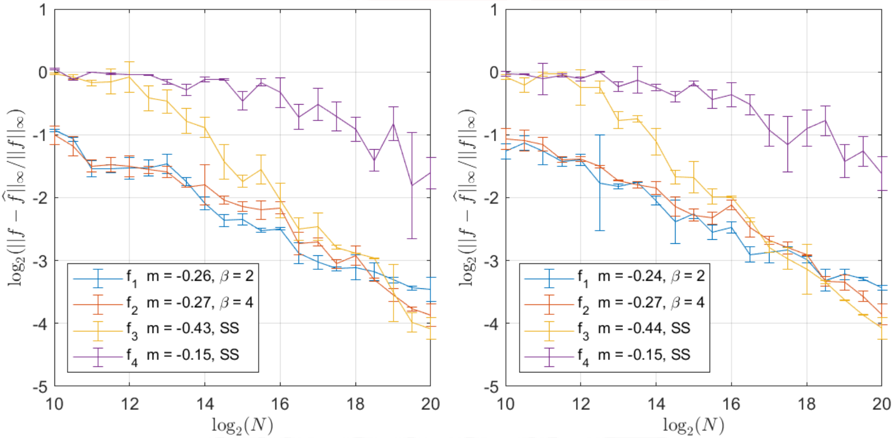

Figure 3: Error decay with varying N; slopes m are for least-squares fit. Left: shifts from ρ_1 (uniform). Right: shifts from ρ_2 (aperiodic).

Observations

Convergence rate

- Nearly identical recovery for different translation distributions.
- Empirical decay rate is better than the theoretical one.
- As expected, rate improves with more smoothness.

But what about f_4?

f_4 **is super smooth but has slower decay. . .**

- f_4 much harder because f_4^{ft} vanishes.
- Kotlarski can be generalized to allow $f^{\mathrm{ft}}(\omega_n) = 0$ for isolated zeros ω_n (we analyze $d = 1$ case; similar theorem with slower rate)
- Previous (discrete) MRA results require $f^{\mathrm{ft}}(\omega) \neq 0$ for **all** frequencies in bandlimit.

Key Takeaways for Inversion

- We present the first full recovery guarantees for functional MRA in the literature for general (i.e. non-bandlimited) functions.
- Functional perspective is useful for these types of problems as it is
  1. more natural for biological applications.
  2. allows one to leverage signal smoothness & tools from deconvolution.
- Functional formulation allows for recovery of functions with vanishing Fourier transform.

Publications on MRA

1. Hirn, Little. "Wavelet Invariants for Statistically Robust Multi-Reference Alignment." *Information and Inference: A Journal of the IMA*, Vol. 10, Issue 4, 2021
2. Hirn, Little. "Power Spectrum Unbiasing for Dilation-Invariant Multi-reference Alignment." *Journal of Fourier Analysis and Applications*, Vol. 29, No. 4, 2023
3. Yin, Little, Hirn. "Bispectrum Unbiasing for Dilation-Invariant Multi-reference Alignment." *IEEE Transactions on Signal Processing*, Vol. 72, pp. 3761-3775, 2024
4. Al-Ghattas, Little, Sanz-Alonso, Sweeney. "Functional Multi-reference Alignment via Deconvolution." Arxiv preprint at link, 2025.

QR code for arxiv pre-print:

What's Next? Multi-target Detection

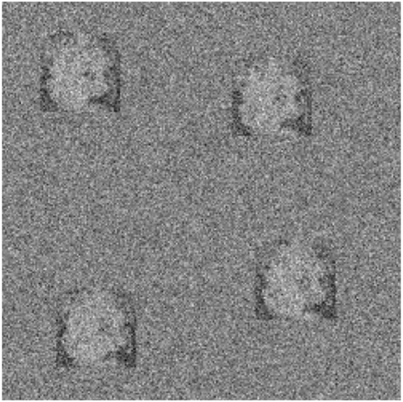

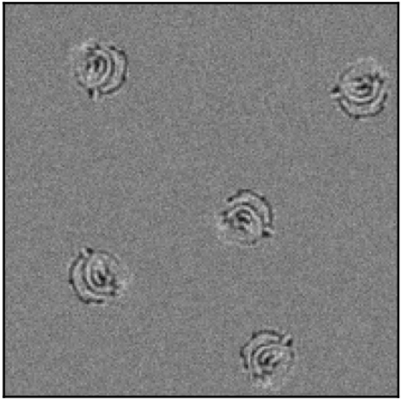

Figure 4: Signal repeats at unknown locations in one large image. Left: MTD (Bendory, Boumal, Leeb, Levin & Singer 2019). Right: MTD with rotations (Marshall, Lan, Bendory & Singer 2020).

Thank you!