Unbiasing Procedures for Scale-invariant Multi-reference Alignment

Matthew Hirn
Michigan State University
mhirn@msu.edu

Anna Little
University of Utah
little@math.utah.edu

***Abstract*—This article discusses a generalization of the 1-dimensional multi-reference alignment problem. The goal is to recover a hidden signal from many noisy observations, where each noisy observation includes a random translation and random dilation of the hidden signal, as well as high additive noise. We propose a method that recovers the power spectrum of the hidden signal by applying a data-driven, nonlinear unbiasing procedure, and thus the hidden signal is obtained up to an unknown phase. An unbiased estimator of the power spectrum is defined, whose error depends on the sample size and noise levels, and we precisely quantify the convergence rate of the proposed estimator. The unbiasing procedure relies on knowledge of the dilation distribution, and we implement an optimization procedure to learn the dilation variance when this parameter is unknown. Our theoretical work is supported by extensive numerical experiments on a wide range of signals.**

I. Introduction

In this article we analyze the following generalization of classic MRA, where signals are also corrupted by a random scale change (i.e. dilation) in addition to random translation and additive noise.

Model 1 (Noisy dilation MRA data model)**.** *The noisy dilation MRA data model consists of M independent observations of a compactly supported, real-valued signal $f \in \mathbf{L}^2(\mathbb{R})$:*

$$y_j(x) = f\left((1-\tau_j)^{-1}(x-t_j)\right) + \varepsilon_j(x)\,, \quad 1 \le j \le M\,. \tag{1}$$

In addition, we assume:

(i) $Pf \in \mathbf{C}^0$.
(i) *supp*$(y_j) \subseteq [-\frac{1}{2}, \frac{1}{2}]$ *for* $1 \le j \le M$.
(ii) $\{t_j\}_{j=1}^M$ *are independent samples of a random variable* $t \in \mathbb{R}$.
(iii) $\{\tau_j\}_{j=1}^M$ *are independent samples from a uniformly distributed random variable* τ *satisfying:*

$$\tau \in \mathbb{R} \quad , \quad \mathbb{E}(\tau) = 0 \quad , \quad Var(\tau) = \eta^2 \le 1/12.$$

(iv) $\{\varepsilon_j(x)\}_{j=1}^M$ *are independent white noise processes on* $[-\frac{1}{2}, \frac{1}{2}]$ *with variance* σ^2.

Model 1 is a first step towards studying more general diffeomorphisms $f(\tau(x))$, since it considers the case when $\tau(x)$ is an affine function. This is relevant to molecular imaging applications since the flexible regions of macro-molecular structures create diffeomorphisms of the underlying shape. Dilations are also highly relevant in imaging applications.

However solving Model 1 is highly challenging. Dilations cause instabilities in the high frequencies of a signal, where even a small dilation can lead to a large perturbation of the frequency values. Ideally, one would like to compute a representation which (1) is both translation and dilation invariant, (2) allows for the additive noise to be removed by averaging, and (3) is invertible with a numerically stable algorithm. However there is a tension between achieving (1) and (3), since the more invariants which are built into the representation, the harder it will be to invert the representation and obtain the underlying signal. In this article we propose the following compromise: we do not define a dilation invariant representation, but propose a method for dilation unbiasing which can be achieved with a numerically stable algorithm; we learn the power spectrum of the hidden signal instead of the hidden signal itself, thus reducing to a phase retrieval problem.

II. Results

From noisy data obtained from Model 1, we construct the following estimator of the power spectrum of the hidden signal:

$$\begin{aligned}(\widetilde{Pf})(\omega) := (I - L_{C_0})^{-1} C_1 L_{C_2} \\ \left[3(\overline{Py} * \phi_L)(\omega) + \omega\left(\overline{Py} * \phi_L\right)'(\omega)\right]\,.\end{aligned} \tag{2}$$

where $\overline{Py} = \frac{1}{M}\sum_{j=1}^M Py_j - \sigma^2$ is the empirical mean of the noisy power spectra, $\phi_L(\omega) = (2\pi L^2)^{-\frac{1}{2}} e^{-\frac{\omega^2}{2L^2}}$ is a Gaussian filter with width L, $(L_C g)(\omega) = C^3 g(C\omega)$ is a dilation operator, and C_0, C_1, C_2 are constants depending on η. Theorem 1 bounds the error of this estimator.

Theorem 1. *Assume Model 1, the estimator* $(\widetilde{Pf})(\omega)$ *defined in* (2)*,* $Pf \in \mathbf{C}^3(\mathbb{R})$*, and that* $\omega^k (Pf)^{(k)}_{\max}(\omega) \in \mathbf{L}^2(\mathbb{R})$ *for* $k = 2, 3$ *where* $(Pf)^k_{\max}(\omega) = \max_{\xi \in [\omega/2, 2\omega]} |(Pf)^k(\xi)|$*. Then*

$$\mathbb{E}\left[\|Pf - \widetilde{Pf}\|^2_{\mathbf{L}^2(\Omega)}\right] \lesssim C_{f,\Omega}\left(\frac{\eta^2}{M} + L^4 + \frac{\sigma^2 \vee \sigma^4}{L^2 M}\right)\,.$$

The following corollary illustrates how a proper choice of L yields an unbiased estimator of Pf; its MSE converges at rate $M^{-\frac{2}{3}}$.

Corollary 1. *Let the assumptions of Theorem 1 hold and in addition let* $\sigma \ge 1$ *and* $L = \left(\frac{\sigma^4}{M}\right)^{\frac{1}{6}}$*. Then:*

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

III. Conclusion

This article considers a generalization of MRA which includes random dilations in addition to random translations and additive noise. The proposed method has several desirable properties compared with previous work. The bias due to dilations is eliminated as the sample size increases. In addition, the method is numerically stable, as the unbiasing procedure operates directly on the power spectrum, rather than features derived from the power spectrum. There are many compelling directions for future research. By extending the inversion unbiasing procedure to operate on the bispectrum, full signal recovery should be possible with an additional computational cost. In addition, preliminary work suggests that inversion unbiasing can be extended to a broad class of dilation distributions as long as their underlying density functions are known. Thus innovative methods for robustly learning the dilation distribution are critical for these methods to become competitive for real world applications.