CpopT: Fast, Guaranteed Deconvolution and Sensing of Low-Dimensional Families in Noisy Time Series

Davit Shadunts
Department of Electrical Engineering
Columbia University
New York, NY
ds4355@columbia.edu

I. Introduction

Detecting structured signals in noisy time-series data is a fundamental problem across various scientific and engineering domains. The classical way to approach this problem is Matched Filtering (MF) [1], which relies on an exhaustive search over a precomputed template bank. However, as the dimension of the signal manifold grows, this method becomes unsuitable for performing a search. Template Optimization (TpopT) [2] was introduced as an alternative to MF, the idea being to use optimization to efficiently search for best-matching templates instead of testing the whole bank. While TpopT achieves its goal, it makes an assumption that the position of the structured signal within the noisy observation is fixed, making it unsuitable for problems where signals may appear at unknown time shifts.

The purpose of this project is to develop Convolutional TpopT (CpopT), designed to efficiently detect and estimate signals with no previous assumptions about the structured signal location.

II. Technical Approach

A. Problem Setup

We study gravitational wave signals produced by the merging of two black holes. These signals are determined by physical parameters, implicitly define the shape of the gravitational waveform:

- m_1, m_2: masses of the black holes
- s_1, s_2: spins (angular momentum)
- $h(t) \coloneqq$ waveform generated from (m_1, m_2, s_1, s_2)

In our setup, we assume that a single waveform is embedded within a longer observation contaminated by noise. This observed signal can be modeled as:

$$x(t) = h(t - \tau) + \eta \tag{1}$$

where:

- τ: unknown insertion time (shift)
- η: Noise
- $\dim(x(t)) \gg \dim(h(t))$

For the sake of this project we assume that the noise is a random Gaussian noise. Exploring other types of noise, including the ones that will more closely mimic the signals we receive from space, is left for future work. To conduct experiments and explore the behavior of gravitational waves in this setting, we generated a bank containing synthetic gravitational waves. We used a Python package called PyCBC [3], which follows the rules of General Relativity and is a standard library for experimenting with gravitational waves. The bank contains 30000 waves, each has a length of 2048. The length of the noisy observation x(t) is fixed at 10000.

B. Problem Formulation

Let $\mathcal{H} = \{h_i(t)\}_{i=1}^{30000}$ denote the bank of gravitational waveforms, each of length 2048. The observed signal $x(t) \in \mathbb{R}^{10000}$ is assumed to be a noisy mixture containing one such waveform, inserted at an unknown time shift τ.

We aim to recover the gravitational waveform $h_i \in \mathcal{H}$ and the unknown shift $\tau \in \{0, 1, \ldots, D\}$ that best match the noisy observation x. This is naturally posed as a maximization of the cross-correlation between h_i and all 2048-length windows of x:

$$(\hat{h}, \hat{\tau}) = \arg \max_{h_i \in \mathcal{H},\ \tau} \sum_{t=0}^{2047} h_i(t) \cdot x(t + \tau) \tag{2}$$

Equivalently, using inner product notation:

$$(\hat{h}, \hat{\tau}) = \arg \max_{h_i \in \mathcal{H},\ \tau} \langle h_i, x_\tau \rangle \tag{3}$$

where $x_\tau(t) = x(t + \tau)$ is the segment of x starting at position τ. This formulation defines our optimization problem: recover the waveform and its location by maximizing alignment with the observed signal.

Due to the large size of the waveform bank and possible shifts, exhaustively evaluating all combinations is computationally inefficient. We therefore seek efficient methods to evaluate the correlation objective without scanning the entire search space. This includes using FFT-based correlation and dimensionality reduction techniques discussed in the following sections.

C. Computing Correlation

As introduced in Equation (2), the core computation involves evaluating the cross-correlation between a waveform $h \in \mathbb{R}^{2048}$ and a noisy observation $x \in \mathbb{R}^{10000}$, across all valid shifts τ. Performing this directly for all shifts would require thousands of inner product evaluations per waveform, which is computationally intensive.

To mitigate this cost, we leverage the convolution theorem to compute correlations efficiently using the Fast Fourier Transform (FFT). This reduces the computational complexity to $\mathcal{O}(L \log L)$, where L is the length of the padded signals.

The computation proceeds as follows:

- Zero-pad both h and x to length $L = \text{len}(x) + \text{len}(h) - 1$, ensuring valid linear correlation and avoiding wrap-around effects.
- Compute the correlation over all shifts using:

$$\text{corr}(h, x)[\tau] = \text{DFT}^{-1}\left(\text{DFT}(x) \cdot \text{DFT}(h)^*\right)[\tau] \quad (4)$$

Here, $\text{DFT}(h)^*$ denotes the complex conjugate of the Fourier transform of h. The derivation of Equation (4) is provided in Appendix A.

D. Embedding Waveforms into Low-Dimensional Space

As shown later in the "Experiments" section, the parameters forming the waveform $(m1, m2, s1, s2)$ fail to capture the relationship between different waves, i.e. a small change in these parameters might result in a completely different waveform. Thus, there is a need to embed waveforms into a low-dimensional space where similar waveforms will lie next to each other.

1) Standard MDS Embedding: Let $d(i, j)$ denote a distance metric between waveforms h_i and h_j. In the standard setting, we define this as the Euclidean distance between the raw waveforms:

$$d(i, j) = \|h_i - h_j\| \quad (5)$$

We apply classical Multidimensional Scaling (MDS) to embed each waveform into $\mathbb{R}^2$. MDS maps:

$$h_i \mapsto \xi_i = \begin{pmatrix} \beta_1^{(i)} \\ \beta_2^{(i)} \end{pmatrix} \in \mathbb{R}^2 \quad (6)$$

such that the distances in the embedded space approximate the original distances:

$$\|\xi_i - \xi_j\| \approx d(i, j) \quad (7)$$

While this embedding captures global waveform geometry under Euclidean distance, it fails to account for temporal misalignments, i.e., waves that are similar to a phase shift are considered to be far from each other.

2) Shift-Invariant MDS Embedding: To address this limitation, we redefine $d(i, j)$ using the best-alignment distance:

$$d_{\text{shift}}(i, j) = \min_{\tau} \|h_i - h_j^{(\tau)}\| \quad (8)$$

where $s_j^{(\tau)}$ is the waveform s_j shifted by τ and aligned with s_i to maximize their cross-correlation.

We then apply MDS to the resulting distance matrix:

$$h_i \mapsto \xi_i^{(\text{inv})} = \begin{pmatrix} \beta_1^{(i)} \\ \beta_2^{(i)} \end{pmatrix} \in \mathbb{R}^2 \quad (9)$$

such that:

$$\|\xi_i^{(\text{inv})} - \xi_j^{(\text{inv})}\| \approx d_{\text{shift}}(i, j) \quad (10)$$

This shift-invariant embedding shows improvement over classical MDS, as shown in the Experiments section.

III. Experiments

We visualize and evaluate our method on synthetic gravitational wave signals generated using PyCBC. We begin by showing a sample waveform generated from physical parameters (m_1, s_1):

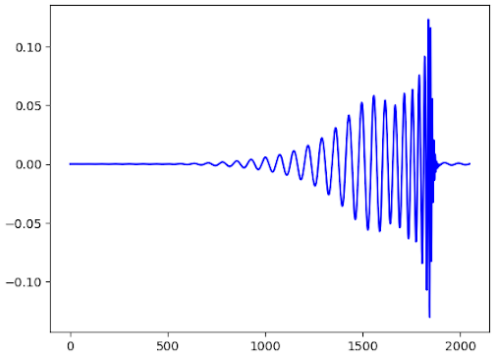

Fig. 1. Generated waveform $h(t)$ from (m_1, m_2, s_1, s_2).

Figure 2 shows the noisy observation $x(t)$, where the waveform is inserted at an unknown shift τ:

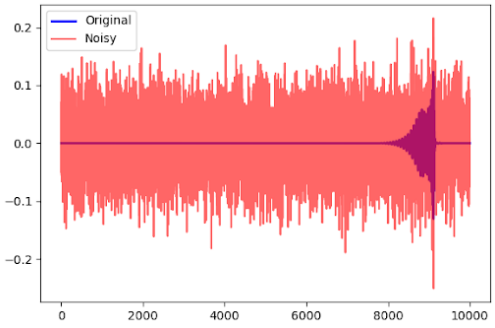

Fig. 2. Observed noisy signal $x(t) = h(t - \tau) + \eta$.

To analyze the geometry of the correlation objective, we visualize the landscape in three different representations of the waveform space:

- The original physical parameter space $(m_1 = m_2, s_1 = s_2)$,
- A 2D MDS embedding using Euclidean waveform distances,
- A 2D shift-invariant MDS embedding using aligned waveform distances.

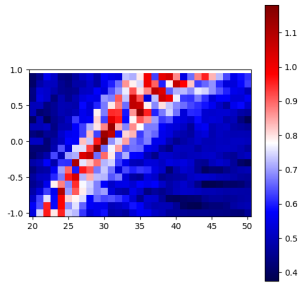

Fig. 3. Objective landscape across original parameter space

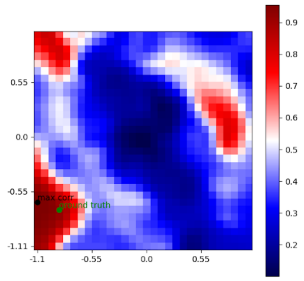

Fig. 4. Objective landscape across Euclidean embedding space

Each point in the landscape (Figure 3) corresponds to a waveform template, and color indicates the maximum cross-correlation score with the observation. The landscape shows that the surface is highly irregular and sensitive to parameter changes, making it unsuitable for gradient-based search.

For the landscapes in the embedding space, the grid dimensions are 30×30, where axes correspond to the (β_1, β_2) coordinates in Equation (6). For each point on the grid, we selected the waveform whose embedding is nearest to the corresponding location in the grid. Figure 4 is constructed using standard MDS. We embed the waveforms into $\mathbb{R}^2$ based on Euclidean distances. This improves the landscape compared to the original parameter space, but there is a clear issue with local maxima due to phase misalignment.

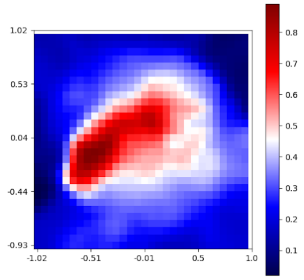

Fig. 5. Objective landscape (smoothed) across shift-invariant embedding space

Figure 5 takes into account phase shifts, which solves the issue with local maxima. However, as shown in Figure 6, the raw landscape remains bumpy and requires smoothing to be usable for optimization.

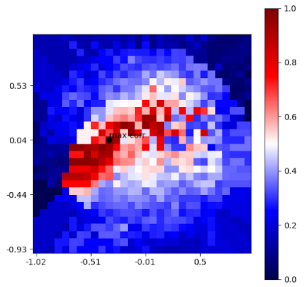

Fig. 6. Unsmoothed shift-invariant landscape.

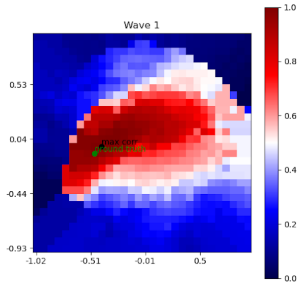

Fig. 7. Shift-invariant landscape + k-NN

To address the irregularities observed in the raw shift-invariant landscape (Figure 6), we propose a strategy based on k-nearest neighbors in the embedding space. Instead of selecting the closest wave with (β_1, β_2) coordinates in the embedding space to represent a square in the grid, we look into the k nearest waveforms for each grid coordinate and select the one that has the highest cross-correlation.

This approach produces a significantly smoother surface, as shown in Figure 7, without the need for any post-processing techniques. However, it comes with the cost of being computationally inefficient. In order to evaluate the best match among the k-nearest neighbors for each grid point, to cover the whole grid, the method requires scanning through the entirety of the waveform bank. As a result, this method has no advantage over the brute-force template matching approach.

To validate that this smoother region captures the correct waveform, we compare the ground truth waveform and the best match obtained via the k-NN approach after alignment:

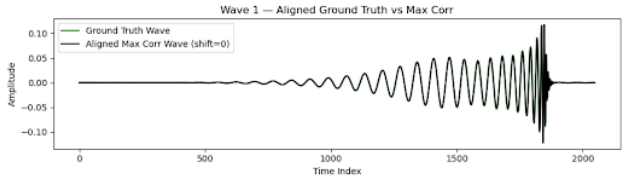

Fig. 8. Alignment of the ground truth waveform and the best match from the k-NN refined landscape.

Figure 8 shows that the wave corresponding to the maximum correlation point in the grid very closely approximates the ground truth.

IV. Discussion and Future Work

While the shift-invariant embedding shows the highest potential by resolving phase misalignment issues, it remains surprisingly bumpy in the region around the maximum correlation point. This raises a question: why do some waveforms that are embedded close together in the shift-invariant space still show low correlation?

To investigate this, we examine the similarity between waveforms near the maximum correlation point. Figure 9 compares the aligned ground truth waveform with the worst neighbor (in terms of correlation) among the $k = 33$ nearest waveforms in the embedding space.

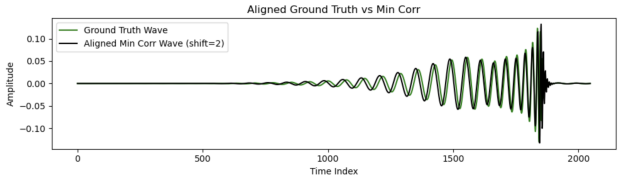

Fig. 9. Aligned ground truth vs. worst correlated waveform among $k = 33$ nearest neighbors in embedding space.

Even though the waves were selected from nearby coordinates in the embedding space, the correlation statistics in Figure 10 reveal that some are poorly correlated with both the noisy observation and the ground truth.

```
Correlation between X and Aligned Worst Corr Wave:                    0.4855
Correlation between X and Aligned Aligned Max Corr Wave:              1.0062
Correlation between X and Ground Truth Wave:                          1.0178
------------------------------------------------------------------------------
Correlation between Aligned Worst Corr Wave and Aligned Max Corr Wave:   0.4464
Correlation between Ground Truth Wave and Aligned Max Corr Wave:         0.9891
Correlation between Ground Truth Wave and Aligned Worst Corr Wave:       0.5093
------------------------------------------------------------------------------
Euclidean distance between Aligned Max Corr and Aligned Worst Corr Wave: 1.0522
Euclidean distance between Ground Truth and Aligned Max Corr Wave:       0.1479
Euclidean distance between Ground Truth and Aligned Worst Corr Wave:     0.9906
```

Fig. 10. Correlation and Euclidean distance statistics between X (noisy observation), ground truth, and best and worst waves (in terms of correlation with X). Poorly correlated waves can still appear close in the embedding space.

Experiments suggest that nearby waves in the shift-invariant embedding space are not necessarily similar. Future work should investigate the root causes of this phenomenon. A better embedding approach may be necessary, one that will be explicitly targeting the objective of embedding similar waves together. Additionally, future efforts could develop an optimization-based approach for waveform selection that avoids exhaustive search.

V. Acknowledgements

This is work done with Professor John Wright, Shiyu Wang, Professor Zsuzsa Márka and Professor Szabolcs Márka

Appendix A
Derivation of Correlation via FFT

The purpose of this section is to provide a derivation of Equation (4). Here I re-derived the formula by referring to Professor John Wright's notes on CpopT, and tried to provide some intuitive understanding of the relationship between correlation and convolution.

Let's start with the definition of convolution between two signals x and h:

$$(x * h)[\tau] = \sum_{n=0}^{L-1} x[n] \cdot h[\tau - n] \tag{11}$$

For correlation, we're essentially computing the convolution of signal x with the time-reversed version of signal h. Let's denote the time-reversed version of h as $\tilde{h}[n] = h[-n]$. Then the correlation can be expressed as:

$$\begin{aligned} \text{corr}(h, x)[\tau] &= (x * \tilde{h})[\tau] \\ &= \sum_{n=0}^{N-1} x[n] \cdot \tilde{h}[\tau - n] \\ &= \sum_{n=0}^{N-1} x[n] \cdot h[n - \tau] \end{aligned} \tag{12}$$

This formulation shows the connection between correlation and convolution, allowing us to leverage the efficient computational properties of convolution.

We know that convolution in time domain is multiplication in frequency domain:

$$(x * h)[\tau] = \text{DFT}^{-1}\{\text{DFT}\{x\} \cdot \text{DFT}\{h\}\}[\tau] \tag{13}$$

To apply this to our correlation calculation, we need to determine how the DFT of the time-reversed signal $\tilde{h}$ relates to the DFT of the original signal h. For a signal of length L, we can derive this relationship:

$$\text{DFT}_L\{\tilde{h}[n]\} = \sum_{n=0}^{L-1} \tilde{h}[n] e^{-j\frac{2\pi}{L}kn} \tag{14}$$

$$= \sum_{n=0}^{L-1} h[-n] e^{-j\frac{2\pi}{L}kn} \tag{15}$$

Let's make a change of variables by setting $m = -n + L$. This gives us:

$$\text{DFT}_L\{\tilde{h}[n]\} = \sum_{n=0}^{L-1} h[-n] e^{-j\frac{2\pi}{L}kn} \tag{16}$$

$$= \sum_{m=L}^{1} h[m - L] e^{-j\frac{2\pi}{L}k(L-m)} \tag{17}$$

$$= \sum_{m=1}^{L} h[m - L] e^{-j\frac{2\pi}{L}kL} e^{j\frac{2\pi}{L}km} \tag{18}$$

$$= \sum_{m=1}^{L} h[m - L] e^{-j2\pi k} e^{j\frac{2\pi}{L}km} \tag{19}$$

$$= \sum_{m=1}^{L} h[m - L] e^{j\frac{2\pi}{L}km} \tag{20}$$

since $e^{-j2\pi k} = 1$ for $k \in \mathbb{Z}$

The sum can be rewritten with adjusted indices to match our usual DFT form:

$$\mathrm{DFT}_L\{\tilde{h}[n]\} = \sum_{m=0}^{L-1} h[m] e^{j\frac{2\pi}{L}km} \tag{21}$$

$$= \left(\sum_{m=0}^{L-1} h[m] e^{-j\frac{2\pi}{L}km}\right)^* \tag{22}$$

$$= \mathrm{DFT}_L^*\{h[n]\} \tag{23}$$

Therefore, the DFT of a time-reversed signal equals the complex conjugate of the original signal's DFT.

Substituting this into the convolution theorem, our correlation becomes:

$$\mathrm{corr}(h, x)[\tau] = (x * \tilde{h})[\tau] \tag{24}$$

$$= \mathrm{DFT}^{-1}\{\mathrm{DFT}\{x\} \cdot \mathrm{DFT}\{\tilde{h}\}\}[\tau] \tag{25}$$

$$= \mathrm{DFT}^{-1}\{\mathrm{DFT}\{x\} \cdot \mathrm{DFT}^*\{h\}\}[\tau] \tag{26}$$

Intuitively, correlation is pretty similar to convolution. Except, rather than following the usual "flip, shift, multiply, add", correlation does not want to flip the signal, so we "double flip" the signal to cancel out this step in convolution. This is mathematically equivalent to convolving signal x with a time-flipped version of h. The time-flipping operation (replacing $h[n]$ with $h[-n]$) effectively creates a "template" that we slide across x to find where it best matches.

In the frequency domain, this time-reversal operation translates to taking the complex conjugate of the Fourier transform. This allows us to compute correlations efficiently using FFT, reducing the computational complexity to $O(L \log L)$.

REFERENCES

[1] B. J. Owen and B. S. Sathyaprakash, "Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement," *Physical Review D*, vol. 60, no. 2, Jun. 1999. [Online]. Available: http://dx.doi.org/10.1103/PhysRevD.60.022002

[2] J. Yan, S. Wang, X. R. Wei, J. Wang, Z. Márka, S. Márka, and J. Wright, "Tpopt: Efficient trainable template optimization on low-dimensional manifolds," 2023. [Online]. Available: https://arxiv.org/abs/2310.10039

[3] A. Nitz, "gwastro/pycbc: v2.3.3 release of PyCBC," https://doi.org/10.5281/zenodo.10473621, Jan. 2024, zenodo, Jan. 09, 2024.