Rapid Communications

Random projections in gravitational wave searches of compact binaries

Sumeet Kulkarni,[1] Khun Sang Phukon,[2] Amit Reza,[3] Sukanta Bose,[4,5,*] Anirban Dasgupta,[3] Dilip Krishnaswamy,[6] and Anand S. Sengupta[3]

[1]*Indian Institute of Science Education and Research, Homi Bhabha Road, Pune 411008, India*
[2]*Department of Physics, Indian Institute of Technology, Kanpur 208016, India*
[3]*Indian Institute of Technology Gandhinagar, Gujarat 382355, India*
[4]*Inter-University Centre for Astronomy and Astrophysics, Post Bag 4, Ganeshkhind, Pune 411 007, India*
[5]*Department of Physics & Astronomy, Washington State University, 1245 Webster, Pullman, Washington 99164-2814, USA*
[6]*IBM Research, Bangalore 560045, India*

(Received 19 January 2018; revised manuscript received 26 March 2019; published 31 May 2019)

Random projection (RP) is a powerful dimension reduction technique widely used in analysis of high dimensional data. We demonstrate how this technique can be used to improve the computational efficiency of gravitational wave searches from compact binaries of neutron stars or black holes. Improvements in low-frequency response and bandwidth due to detector hardware upgrades pose a data analysis challenge in the advanced LIGO era as they result in increased redundancy in template databases and longer templates due to a higher number of signal cycles in band. The RP-based methods presented here address both these issues within the same broad framework. We first use RP for an efficient, singular value decomposition-inspired template matrix factorization and develop a geometric intuition for why this approach works. We then use RP to calculate approximate time-domain correlations in a lower dimensional vector space. For searches over parameters corresponding to nonspinning binaries with a neutron star and a black hole, a combination of the two methods can reduce the total on-line computational cost by an order of magnitude over a nominal baseline. This can, in turn, help free up computational resources needed to go beyond current spin-aligned searches to more complex ones involving generically spinning waveforms.

DOI: 10.1103/PhysRevD.99.101503

I. INTRODUCTION

The direct detections of gravitational waves (GWs) from the mergers of black holes and neutron stars [1–6] by Advanced LIGO (aLIGO) [7] and Advanced Virgo (AdV) [8] detectors in the first and second observing runs (O1 and O2, respectively) have launched the era of GW astronomy [9,10]. In the coming years, the global network of ground-based detectors, comprising aLIGO, AdV, KAGRA [11], and LIGO-India [12] will not only increase the detection rate and facilitate the search for their possible electromagnetic counterparts [13–15] but also produce an unprecedentedly large amount of data, which can pose an interesting computational challenge for GW data analysis.

At present, theoretically modeled compact binary coalescence (CBC) waveforms are used as templates to matched-filter [16] the detector data in these searches [17,18]. A brute force computation of this cross correlation with a suitable grid of templates spanning astrophysical ranges of search parameters can be expensive (but see [19–21]). As these detectors are paced through planned upgrades, one expects better sensitivity at low frequencies and an increase in the detector bandwidth. The combined effects of these changes will not only increase the volume of the search parameter space but also result in denser template banks, thereby increasing their redundancy. More cycles of the signal will fall in band and increase their duration. These highlight the need for designing efficient and scalable methods for matched-filtering-based templated CBC searches [22–27].

In a seminal work, Cannon *et al.* [28–30] showed how singular value decomposition (SVD) can mitigate the redundancies in CBC template banks by effectively reducing the number of filters or templates, owing to their strong correlation for similar parameter values, with negligible effect on search performance. We show, however, that the computational cost of SVD factorization does not scale favorably with an increase in bank size. Further, it may not be possible to factorize very large banks *in toto* as it requires prohibitively large random access memory.

Random projection (RP), conceived by the pioneering work of Johnson and Lindenstrauss [31], is a computationally efficient technique for dimension reduction and finds applications in many areas of data science [32]. In this Rapid Communication, we apply this technique to address two key challenges in future CBC searches: handling redundancies in large template databases and efficiently correlating noisy data against long templates.

*Corresponding author.
sukanta@wsu.edu

The primary impact of this work is multifold: (1) Efficient template matrix factorization can be used to address the redundancy problem. This is similar in spirit to the SVD factorization that is at the heart of the "GstLAL"-based inspiral pipeline [27,30,33], but our RP method scales well for very large number of templates embedded in high-dimensional Euclidean space. Such factorizations can be done off-line, in advance of a CBC search. Nonetheless there can be situations when the factors need to be updated on-line, e.g., owing to the nonstationarity of data. Our adaptations will benefit both scenarios. (2) We show the explicit connection between the new factorization scheme and the extant SVD method. This bridges the two approaches and makes it readily usable. (3) The computational challenges arising from correlating noisy data against long templates (also known as the *curse of dimensionality*) is addressed by casting the matched-filtering operation in a lower-dimensional space. For certain types of template banks (e.g., for CBCs with precessing spins), the template matrix may be less amenable to a SVD-like factorization. There the total computational cost can be significantly reduced by using the RP-based correlation alone. (4) Finally, we show that RP-based template matrix factorization and matched-filtering computation in reduced dimension can be combined effectively for efficient CBC searches.

Currently the GstLAL-based inspiral pipeline utilizes time-slicing of templates to improve computational efficiency and also involves spin-aligned templates. Since it is for the first time that the RP is being introduced in GW searches, our primary objective here is to elucidate how its core ideas can help them. This is why we demonstrate application of RP in the simple case of a single slice of data and nonspinning inspiral templates. This simplification notwithstanding, the RP-based methods introduced here can be readily applied to time-sliced data and spin-aligned templates. (A detailed study of that application and the computational advantage so gained will be presented in a future work.)

II. COMPACT BINARY SEARCHES

Consider a CBC search involving a bank of N_T templates over a given parameter space. Following the convention in Ref. [28], let $\mathbf{H}$ denote the $2N_T \times N_s$ template matrix with $2N_T$ rows of real-valued unit-norm whitened filters, each sampled over N_s time points. The template matrix may be viewed as $2N_T$ row vectors embedded in N_s-dimensional Euclidean space $\mathbb{R}^{N_s}$. The complex matched-filter output of the α^{th} template at a specific point in time, against the whitened data $\vec{S}$ is the inner product:

$$\rho_\alpha = (H_{(2\alpha-1)} - iH_{(2\alpha)})\vec{S}^T, \tag{1}$$

where H_α denotes the α^{th} row of $\mathbf{H}$ and $\vec{S}^T$ is the transpose of $\vec{S}$. The signal-to-noise ratio maximized over the initial phase ϕ_0, is given by $|\rho_\alpha|$. In our notation, the H_α's and signal $\vec{S}$ are assumed to be row vectors. The overlap between two templates, when maximized over extrinsic parameters (e.g., the time t_0 and phase ϕ_0 at arrival or coalescence of the signal in band), produces the match. The match between templates with similar intrinsic parameters (such as the compact object masses and spins), can be very high—signifying the rank deficiency of the template matrix. A typical off-line CBC search involves calculating the cross correlation between $\vec{S}$ and every row of $\mathbf{H}$ for a series of relative time shifts, or values of t_0, thereby generating a time series of ρ_α values, for every α. The use of a large number of templates (N_T), each sampled over a large number of points (N_s) amplifies the search's computational cost.

The rank deficiency of $\mathbf{H}$ is exploited in the truncated SVD approach, where every row is approximated as a linear combination of only ℓ of the $2N_T$ right singular vectors with the most dominant singular values. Further, these "basis" vectors are used as eigen-templates against which the data are cross correlated. The left singular vectors of $\mathbf{H}$ and the singular values are combined into a coefficient matrix that is used to reconstruct the approximate signal-to-noise ratio (SNR). The truncation of the basis leads to errors in the approximation of the template waveforms, which further translates to imperfect reconstructions of the SNR. The fractional SNR loss can be measured as a function of the discarded $(2N_T - \ell)$ singular values.

The SVD factorization of the template matrix $\mathbf{H}$ has a time-complexity proportional to $\mathcal{O}(N_T^2 N_s)$, assuming $N_T \leq N_s$. Thus, such factorizations fast become computationally unviable with increasing size of a template bank. Since the entire template matrix can become too large to be saved in single machine memory, a suitable parallel scheme is required to apply SVD to larger banks. SVD-based on-line CBC searches [33,34] work around this problem by splitting the bank into smaller sub-banks that are more amenable to such factorization separately. While the optimal way of partitioning the bank is an open problem, the act of splitting the bank prevents exploitation of the linear dependency of templates across the sub-banks. This is seen in Fig. 1, where we plot β, which is defined as the ratio of the number of basis vectors summed across all the sub-banks to the number of basis vectors from the SVD factorization of the full bank, at a given average fractional loss in accuracy of the reconstructed SNR. By splitting the bank, one effectively ends up requiring many more eigen-templates against which the data are filtered. When extrapolated to realistic template bank sizes of $N_T \approx 10^6$, β can be as large as $\sim 10^2$ at $\langle \frac{\delta\rho}{\rho} \rangle = 10^{-3}$.

The SVD-inspired RP-based factorization presented below addresses this issue and is scalable for large template banks. We also apply RP to calculate the correlations in a lower-dimensional space $\mathbb{R}^k$ where $k \leq N_s$. These correlations could be either between a template and the data as shown in Eq. (1) or between the basis vectors and the data

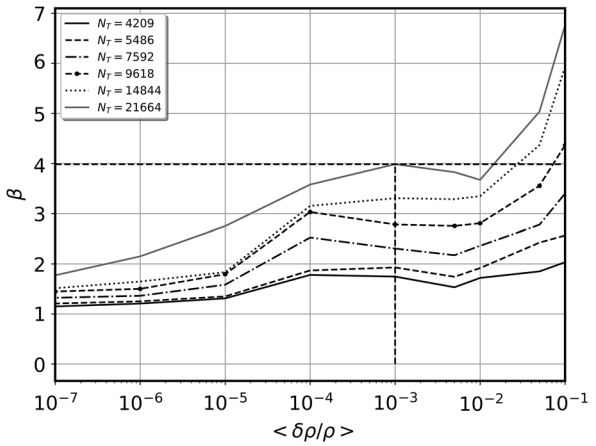

FIG. 1. The factor by which the number of SVD basis vectors increases due to partitioning of the template bank of size N_T into sub-banks of 500 templates each is shown as β on the vertical axis. Results from six different template-bank sizes are shown. For example, the bank with 4209 templates is divided into eight sub-banks of 500 templates each and a ninth one of 209 templates. There β reaches a high of ~ 2 when one tolerates an average fractional loss in SNR of $\langle \delta\rho/\rho \rangle \sim 0.1$. On the other hand, for any of these template banks, as one approaches machine-precision accuracy in SNR reconstruction, $\beta \to 1$ as expected. A practical operating point would be $\langle \delta\rho/\rho \rangle \sim 10^{-3}$. The trend from the six examples shown here indicates that β can be quite large for searches in aLIGO data where $N_T \sim 10^5$.

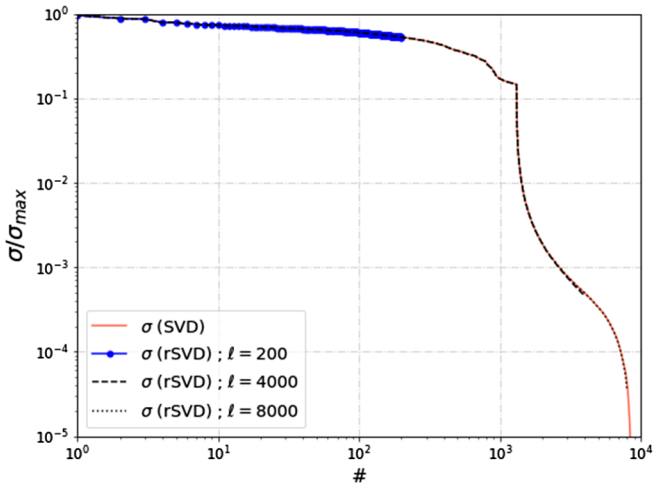

FIG. 2. Comparison of singular values σ for a template matrix $\mathbf{H}$ of size $(2N_T \times N_s) \equiv 9130 \times 65536$, normalized by the maximum singular value $\sigma_{\max}$, as obtained from SVD and RSVD factorization. RSVD is performed in target dimensions $\mathbb{R}^\ell$ where $\ell = 200$, 4000 or 8000. The spectrum of eigenvalues is seen to fall steeply. This example template bank was constructed using a nonspinning signal model for component mass parameters $(m_{1,2})$ in the range $2.5\ M_\odot \leq m_1, m_2 \leq 17.5\ M_\odot$. As seen here, the top-$\ell$ eigenvalues obtained by RSVD agree very well with the spectrum obtained by traditional SVD factorization.

within the SVD paradigm. The full potential of the RP-based methods introduced here can be realized by combining them together. We demonstrate its feasibility with an example.

III. RANDOM PROJECTION

The core theoretical idea behind the RP technique is the Johnson-Lindenstrauss (JL) lemma [31], which states that a set of $2N_T$ vectors in $\mathbb{R}^{N_s}$ can be mapped into a randomly generated subspace $\mathbb{R}^k$ of dimension $k \sim \mathcal{O}(\log(2N_T)/\epsilon^2)$ or greater, while preserving all pairwise L_2 norms to within a factor of $(1 \pm \epsilon)$, where $0 < \epsilon < 1$, with a very high probability. Here, ϵ is the mismatch or distortion tolerated in the pairwise L_2 norms between any two filters after projection. Thus, RP also approximately preserves any statistic of the dataset that is characterized by such pairwise distances. The RP of $\mathbf{H}$ onto $\mathbb{R}^k$ produces $\mathbf{H\Omega}$; the accuracy of this data-oblivious transformation depends on the target dimensions and sampling distribution of the $N_s \times k$ projection matrix $\mathbf{\Omega}$. While it is enough to sample the entries independently and identically distributed from a sub-Gaussian distribution, here we choose them independently from a Gaussian distribution with mean zero and variance $1/k$, i.e., $\mathcal{N}(0, 1/k)$, thus producing a Gaussian quasiorthonormal random matrix [35,36], such that $\langle \mathbf{\Omega\Omega}^T \rangle = I$. Results obtained from RP-based processing can vary depending on the actual choice of the distribution (from which elements of $\mathbf{\Omega}$ are drawn), and in a statistical sense, these results arising from different choices of $\mathbf{\Omega}$ are expected to be equivalent due to the quasiorthonormality of the projection. (See Supplemental Material [37] for a geometric explanation.)

IV. RP-BASED TEMPLATE MATRIX FACTORIZATION

The key idea behind an ℓ-truncated SVD approximation of $\mathbf{H}$ is to reconstruct the rows of the template matrix using the top-ℓ right-singular vectors. This approximation works well because $\mathbf{H}$ has a fast-decaying spectrum, as shown in Fig. 2. In making the truncation, one effectively reduces $\mathbf{H}$ to its ℓ-rank approximation $\mathbf{H}^{(\ell)}$ [38,39].[1] Further, for a bank of normalized templates, it is easy to show that the average fractional loss in SNR due to the truncation is given as $\langle \delta\rho/\rho \rangle = \|\mathbf{H} - \mathbf{H}^{(\ell)}\|_F^2 / \|\mathbf{H}\|_F^2$, where $\langle\rangle$ denotes average over the bank of templates and $\|\cdot\|_F$ is the Frobenius norm [38]. However, the existing SVD algorithms do not scale well with increasing dimensions and redundancies of the template database.

Randomized SVD (RSVD) [39] is a RP-based matrix-factorization technique to obtain an ℓ-rank matrix factorization $\mathbf{H}^{(\ell)}$ such that, for some specified $\eta > 0$, $\|\mathbf{H} - \mathbf{H}^{(\ell)}\|_F \leq \min_{\{\mathbf{X}:\ \mathrm{rank}(\mathbf{X}) \leq \ell\}} \|\mathbf{H} - \mathbf{X}\|_F (1 + \eta)$ with

[1]Note that $\mathbf{H}^{(\ell)}$ has the same dimensions as $\mathbf{H}$.

high probability.[2] In one implementation, the RSVD algorithm proceeds by first projecting the individual row vectors in the template matrix $\mathbf{H}$ to $\mathbb{R}^{\ell}$ by using $\bar{\mathbf{\Omega}}_{N_s\times\ell} \in \mathcal{N}(0, 1/\ell)$, thereby yielding $\bar{\mathbf{H}}_{2N_T\times\ell} = \mathbf{H}\bar{\mathbf{\Omega}}$. The latter can be used to perform an SVD-like factorization directly in $\mathbb{R}^{\ell}$ through a series of operations like the ones described below. In passing, we note that while $\bar{\mathbf{H}}$ is an object in a lower-dimensional Euclidean space relative to $\mathbf{H}$, it is not constituted of time-decimated templates.

Figure 2 compares the singular values obtained by the RP-based factorization against those from a direct SVD factorization. As seen there, it is typically sufficient to take $\ell \ll N_T$. (Since $N_T \leq N_s$, as mentioned above, it follows that $\ell \ll N_s$ as well.) In fact, the numerical value of ℓ chosen in RSVD may be smaller than the theoretical JL bound prescribed for preserving pairwise L_2 distances between the $2N_T$ rows to a ϵ-distortion factor. Working with the reduced sized matrix $\bar{\mathbf{H}}$ leads to significant computational savings, while producing a decomposition that closely approximates the optimal ℓ-rank factorization of $\mathbf{H}$. The optimum choice of ℓ depends on the shape of the eigenvalue spectrum. In the Monte Carlo simulations presented in the Supplemental Material [37], we choose $\ell = 200$. The corresponding average SNR loss for a set of 500 CBC signals added to simulated aLIGO noise is $\langle\delta\rho/\rho\rangle = 2\times10^{-4}$ in that study (see Fig. 3 in Supplemental Material [37]).

RSVD thus proceeds by obtaining a set of orthogonal bases for the column space of $\bar{\mathbf{H}}$ by using a thin-QR decomposition [38]: $\bar{\mathbf{H}} = \mathbf{QR}$, where $\mathbf{Q}$ is an orthonormal matrix with dimensions $2N_T \times \ell$. The approximate rank-ℓ decomposition is then obtained as $\mathbf{H}^{(\ell)} = \mathbf{Q}(\mathbf{Q}^T\mathbf{H}) = \mathbf{QB}$, where $\mathbf{B}_{\ell\times Ns} \equiv \mathbf{Q}^T\mathbf{H}$ is a matrix that defines the orthonormal projection of the template waveforms into the compressed subspace. It is clear that one can use the ℓ rows of $\mathbf{B}$ as the surrogate templates, which in turn can be used to correlate against the detector data $\vec{S}$. These can be further combined with $\mathbf{Q}$ to reconstruct ρ in $\mathbb{R}^{N_s}$. We can thus use the QB decomposition itself to improve the efficiency of both the time and frequency domain searches by constructing $\mathbf{H}$ appropriately, with templates from the corresponding domains.

Instead of randomly projecting the column space of $\mathbf{H}$, the method can be generalized by applying RP on both the row and column spaces [39]. This bilateral RSVD method is particularly useful when both N_s and N_T are very large.

V. RECONSTRUCTION OF SNR

The rank-ℓ matrix factorization of $\mathbf{H}$ using RSVD is given by $\mathbf{H}^{(\ell)} = \mathbf{QB}$. Thus, the SNR ρ'_α, for any given t_0, can be reconstructed in $\mathbb{R}^{N_s}$ as

$$\begin{aligned}\rho'_\alpha &= (H^{(\ell)}_{(2\alpha-1)} - iH^{(\ell)}_{(2\alpha)})\vec{S}^T \\ &= \sum_{\nu=1}^{\ell}(Q_{(2\alpha-1)\nu} - iQ_{(2\alpha)\nu})(B_\nu\vec{S}^T). \end{aligned} \tag{2}$$

Using the Pythagoras theorem, and the fact that $\|\mathbf{H}\|_F^2 = 2N_T$, it is easy to show that the average fractional loss of SNR is given by

$$\left\langle\frac{\delta\rho}{\rho}\right\rangle \leq \frac{\|\mathbf{H}\|_F^2 - \|\mathbf{H}^{(\ell)}\|_F^2}{\|\mathbf{H}\|_F^2} = 1 - \frac{\sum_{\mu=1}^{\ell}\sigma_\mu^2}{2N_T}, \tag{3}$$

where σ_μ are the eigenvalues of $\mathbf{H}^{(\ell)}$. For the example discussed in Fig. 2, $\sum_{\mu=1}^{\ell}\sigma_\mu^2/(2N_T) < 1$ but approaches unity monotonically with increasing ℓ. The right-hand side of Eq. (2) can be calculated efficiently by evaluating the Frobenius norm of $\mathbf{B}$ directly (i.e., without explicitly finding the eigenvalues of $\mathbf{H}^{(\ell)}$ first). Thus, the QB decomposition can indeed serve as a stand-in replacement for the SVD factorization. (For an efficient method of explicitly calculating the SVD factors from the RP-based factorization see Supplemental Material [37].)

Ideally one would like to use $\langle\delta\rho/\rho\rangle$ as the control parameter and solve Eq. (3) for the optimum value of ℓ. However, this is a hard problem and in practice the value is set by a process of trial and error, which thankfully can be done off-line even when the computation in Eq. (2) is conducted on-line.

A naive implementation of matched-filter in time-domain can be very expensive, with a complexity of $\mathcal{O}(N_s^2)$ per template for N_s time-shifts. Of course, the Fast Fourier transform (FFT) can reduce this to $\mathcal{O}(N_s \log N_s)$. It is however more efficient instead to first project the two aforementioned whitened time-series vectors in $\mathbb{R}^{N_s}$ to a random k-dimensional ($k \ll N_s$) subspace and then calculate the match (circular cross correlations), as seen in Fig. 2 of the Supplemental Material [37]. In fact, for the template part, one can directly project the rows of the $\mathbf{B}$ matrix (which serve as surrogate templates) to R^k ($k \leq N_s$). In this context, RP reduces the complexity of calculating the matches by a factor N_s/k. (See Supplemental Material [37] for how FFT-like algorithms enable its fast computation [38].)

VI. COMPUTATIONAL COMPLEXITY ANALYSIS

The straightforward SVD factorization of $\mathbf{H}$ requires $\mathcal{O}(N_T^2N_s)$ floating-point operations, assuming $N_T \leq N_s$. In comparison, the cost of the RP matrix factorization is $\mathcal{O}(\ell N_T N_s + (\ell^2 N_T - \frac{2}{3}\ell^3) + \ell N_T N_s + \ell N_s)$. In this last expression, we have included partial contributions from first projecting the template matrix to $\mathbb{R}^{\ell}$, then taking the thin-QR decomposition of $\bar{\mathbf{H}}$ using Householder's method [40], followed by the cost of constructing $\mathbf{B}$ and calculating its Frobenius norm, respectively. For practical cases, one

[2]Note that the values of η and ϵ can be different.

expects $\ell \ll N_s$, due to which the cost of factorizing $\mathbf{H}$ can be orders of magnitude less than a full SVD factorization. This advantage is not just realized off-line, but can also directly impact the total on-line cost of the searches owing to a lower value of ℓ alone: Figure 1 shows that for moderate sized banks one effectively ends up using ∼3–4 times fewer surrogate templates in the on-line portion of the search from the new RP-based factorization. This improvement is expected to be higher for larger banks.

For on-line searches, the number of floating point operations per second (flops) in our method is $N_{\text{flops}} = (2\ell k f_s + 2\ell N_T f_s + k f_s)$. The first term is the number of floating-point operations required for computing the cross correlation between the surrogate templates (rows of $\mathbf{B}$) and the data vector; the second term is the cost of reconstruction of the SNR for every template; and the third term is the cost of projecting the data vector into the lower-dimensional space. In the SVD-only method, the expression for N_{flops} is analogous, except that instead of the last term above, it has a down-sampling cost that is similarly insignificant as our projection cost. The primary difference between the two methods is that owing to our use of RSVD and RP, ℓ and k are less than the number of basis templates and the number of time samples of data used, respectively, in the SVD-only method. For the crucial last couple of seconds of the cross-correlation analysis for CBC signals we have evaluated that our method is an order of magnitude faster than the SVD-only method.

VII. CONCLUSION

In summary, here we introduced random projection-based techniques that hold promise for factorization of large template matrices and cross correlation of templates with data in a scalable and computationally efficient way, which can aid more complex searches, such as of CBCs with generic spins, and, hence, improve the chances for new discoveries.

ACKNOWLEDGMENTS

We would like to thank Surabhi Sachdev for carefully reading the manuscript and making useful comments. This work is supported in part by DST's SERB Grants No. EMR/2016/007593 and No. DST/ICPS/CLUSTER/Data Science/General/T-150, NSF Grant No. PHY-1506497, and the Navajbai Ratan Tata Trust. A large set of data analysis studies were performed on the Sarathi computing cluster at IUCAA.

[1] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), Phys. Rev. Lett. **116**, 061102 (2016).
[2] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), Phys. Rev. Lett. **116**, 241103 (2016).
[3] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), Phys. Rev. X **6**, 041015 (2016).
[4] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), Phys. Rev. Lett. **118**, 221101 (2017).
[5] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), Phys. Rev. Lett. **119**, 141101 (2017).
[6] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), Phys. Rev. Lett. **119**, 161101 (2017).
[7] J. Aasi *et al.* (LIGO Scientific Collaboration), Classical Quantum Gravity **32**, 074001 (2015).
[8] F. Acernese *et al.* (Virgo Collaboration), Classical Quantum Gravity **32**, 024001 (2015).
[9] B. P. Abbott *et al.* (Virgo, Fermi-GBM, INTEGRAL, and LIGO Scientific Collaborations), Astrophys. J. **848**, L13 (2017).
[10] B. P. Abbott *et al.*, Astrophys. J. **848**, L12 (2017).
[11] T. Akutsu *et al.* (KAGRA Collaboration), arXiv:1710.04823.
[12] C. S. Unnikrishnan, Int. J. Mod. Phys. D **22**, 1341010 (2013).
[13] J. Rana, A. Singhal, B. Gadre, V. Bhalerao, and S. Bose, Astrophys. J. **838**, 108 (2017).
[14] V. Srivastava, V. Bhalerao, A. P. Ravi, A. Ghosh, and S. Bose, Astrophys. J. **838**, 46 (2017).
[15] S. Ghosh and S. Bose, arXiv:1308.6081.
[16] C. W. Helstrom, *Elements of Signal Detection and Estimation* (Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1995).
[17] B. J. Owen, Phys. Rev. D **53**, 6749 (1996).
[18] B. S. Sathyaprakash and S. V. Dhurandhar, Phys. Rev. D **44**, 3819 (1991).
[19] R. Prix, Classical Quantum Gravity **24**, S481 (2007).
[20] I. W. Harry, B. Allen, and B. S. Sathyaprakash, Phys. Rev. D **80**, 104014 (2009).
[21] S. Roy, A. S. Sengupta, and N. Thakor, Phys. Rev. D **95**, 104045 (2017).
[22] T. Cokelaer, Phys. Rev. D **76**, 102004 (2007).
[23] B. Abbott *et al.* (LIGO Scientific Collaboration), Phys. Rev. D **78**, 042002 (2008).
[24] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel, Phys. Rev. D **89**, 024010 (2014).
[25] S. Babak, Classical Quantum Gravity **25**, 195011 (2008).
[26] G. M. Manca and M. Vallisneri, Phys. Rev. D **81**, 024004 (2010).
[27] S. Privitera, S. R. P. Mohapatra, P. Ajith, K. Cannon, N. Fotopoulos, M. A. Frei, C. Hanna, A. J. Weinstein, and J. T. Whelan, Phys. Rev. D **89**, 024003 (2014).
[28] K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, and A. J. Weinstein, Phys. Rev. D **82**, 044025 (2010).

[29] K. Cannon, C. Hanna, and D. Keppel, Phys. Rev. D **84**, 084003 (2011).
[30] K. Cannon *et al.*, Astrophys. J. **748**, 136 (2012).
[31] W. B. Johnson and J. Lindenstrauss, in *Conference in Modern Analysis and Probability (1982, New Haven, Conn.)*, Contemporary Mathematics Vol. 26 (American Mathematical Society, Providence, RI, 1984), pp. 189–206.
[32] E. Bingham and H. Mannila, in *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, New York, 2001), pp. 245–250.
[33] C. Messick *et al.*, Phys. Rev. D **95**, 042001 (2017).
[34] K. Cannon *et al.*, Astrophys. J. **748**, 136 (2012).
[35] S. Dasgupta and A. Gupta, Random Struct. Algorithms **22**, 60 (2003).
[36] S. Dasgupta, in *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence* (Morgan Kaufmann Publishers Inc., San Francisco, 2000), pp. 143–151.
[37] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevD.99.101503 for a geometric explanation for why RSVD preserves the top singular subspaces and a demonstration of how RP can be used to compute the SNR time-series directly in the target space.
[38] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. (Johns Hopkins University Press, Baltimore, MD, USA, 1996).
[39] N. Halko, P.-G. Martinsson, and J. A. Tropp, SIAM Rev. **53**, 217 (2011).
[40] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, 3rd ed. (Cambridge University Press, New York, NY, USA, 2007).