

Search strategies for long gravitational-wave transients: Hidden Markov model tracking and seedless clustering

Sharan Banagiri,^{1,*} Ling Sun,^{2,†} Michael W. Coughlin,² and Andrew Melatos³

¹*School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, USA*

²*LIGO Laboratory, California Institute of Technology, Pasadena, California 91125, USA*

³*OzGrav, University of Melbourne, Parkville, Victoria 3010, Australia*



(Received 11 March 2019; published 16 July 2019)

A number of detections have been made in the past few years of gravitational waves from compact binary coalescences. While there exist well-understood waveform models for signals from compact binary coalescences, many sources of gravitational waves are not well modeled, including potential long-transient signals from a binary neutron star postmerger remnant. Searching for these sources requires robust detection algorithms that make minimal assumptions about any potential signals. In this paper, we compare two unmodeled search schemes for long-transient gravitational waves, operating on cross-power spectrograms. One is an efficient algorithm first implemented for continuous wave searches, based on a hidden Markov model. The other is a seedless clustering method, which has been used in transient gravitational wave analysis in the past. We quantify the performance of both algorithms, including sensitivity and computational cost, by simulating synthetic signals with a special focus on sources like binary neutron star postmerger remnants. We demonstrate that the hidden Markov model tracking is a good option in model-agnostic searches for low signal-to-noise ratio signals. We also show that it can outperform the seedless method for certain categories of signals while also being computationally more efficient.

DOI: [10.1103/PhysRevD.100.024034](https://doi.org/10.1103/PhysRevD.100.024034)

I. INTRODUCTION

The discoveries of gravitational waves (GW) by Advanced Laser Interferometer Gravitational Wave Observatory (Advanced LIGO) and Advanced Virgo detectors [1,2] have opened the new window for gravitational wave astrophysics. As of the end of the second observing run (O2), LIGO and Virgo have observed multiple binary black hole (BBH) coalescences [3] and one binary neutron star (BNS) merger GW170817 [4], with the latter marking the advent of multimessenger astronomy [5].

There is also considerable interest in understanding the fate of the BNS merger remnant. In particular, searches have been carried out using O2 interferometric data from LIGO, Virgo, and GEO600 for GW signals from a possible short, intermediate, or long lived remnant of GW170817 with timescales of order of 1 s, 100–1000 s and $\gtrsim 1000$ s respectively [6–9]. Since the nature of the remnant and the exact form of GW emission is unknown, unmodeled searches have played a large role in the analysis. The Stochastic Transient Analysis Multi-detector Pipeline (STAMP), which searches for excess GW power in spectrograms of cross-correlated data, has been employed in both intermediate and long duration searches. These

spectrograms are parsed by pattern recognition algorithms for GW signals. Several such algorithms have been proposed in the past; they can be broadly categorized as seed-based and seedless [10–13]. Seed-based algorithms identify loud seed pixels in the spectrogram (above some threshold) and attempt to grow contiguous clusters from them by adding neighboring pixels. Seedless algorithms pick out clusters by drawing tracks from some predefined template bank. Since they do not depend on initial loud pixels, they are generally more sensitive than seeded algorithms, especially towards toward narrowband waveform models albeit at a higher computational cost [12,14]. Seedless algorithms have to compromise between computational cost and sensitivity to waveform morphology—for example the implementations of seedless algorithms in [6,8] adopt quadratic fitting in the time-frequency spectrograms.

Here, we apply a hidden Markov model (HMM) tracking algorithm—first implemented for continuous gravitational wave searches in Refs. [15–17]—to cross-correlated data. HMM-based tracking in frequency domain provides accurate estimates of the signal frequency at low signal-to-noise ratio (SNR) when a large number of observational samples are available [18,19]. In the GW context, it was applied in the first observing run (O1) of Advanced LIGO to search for continuous waves from the brightest low-mass X-ray binary, Scorpius X-1 [20]. A revised HMM was used to

*banag002@umn.edu

†lssun@caltech.edu

search for signals from a long-lived post-merger remnant of the binary neutron star merger GW170817 [8,21]. We apply this algorithm to the cross-power maps produced by STAMP and make a quantitative comparison between the performance of HMM and the seedless algorithm. We demonstrate that HMM can outperform seedless algorithms for specific waveform models, and in particular for models used in postmerger remnant searches [22].

The rest of the paper is organized as follows. In Sec. II, we briefly describe the cross-power maps, the pixel SNR statistic and the detection statistic. In Sec. III, we describe the two methods being compared—HMM tracking and the seedless clustering algorithms. In Sec. IV, we compare the detection efficiency and computational cost of the two algorithms for a variety of waveform models, and demonstrate that the HMM tracking generally outperforms in both aspects. A summary of the paper is given in Sec. V.

II. CROSS-POWER MAP

Unmodeled transient searches with STAMP are usually done on spectrograms of cross-power. They are constructed by cross-correlating data between two GW detectors in the frequency domain. We follow the definition in Ref. [13] and construct normalized cross-power spectrograms as follows:

$$\rho(t; f, \hat{n}) = \Re \left[\frac{2\tilde{Q}(t; f, \hat{n})\tilde{s}_I^*(t; f)\tilde{s}_J(t; f)}{|\tilde{Q}(t; f, \hat{n})|\sqrt{\frac{1}{2}P_I(t; f)P_J(t; f)}} \right], \quad (1)$$

where $\tilde{s}_{I,J}(t; f)$ is the discrete Fourier transform of data from detector I, J calculated over some segment duration T , and $P_{I,J}(t; f)$ is the noise autopower [23] in detector I, J . Here \Re denotes the real part of a complex number, and $\tilde{Q}(t; f, \hat{n})$ is a complex filter function which helps “point” the search in direction of \hat{n} as seen from Earth, given by,

$$\tilde{Q}(t; f, \hat{n}) = \frac{2 \exp(2\pi i f \hat{n} \cdot \vec{\Delta x}_{IJ}/c)}{\sum_A F_I^A(t; \hat{n}) F_J^A(t; \hat{n})}. \quad (2)$$

Here $F_{I,J}^A(t; \hat{n})$ is the antenna pattern of detector I, J for polarization $A \in \{+, \times\}$, $\vec{\Delta x}_{IJ}$ is the distance between the detectors, and c is the speed of light. We point the reader to Ref. [13] for a derivation of Eqs. (1) and (2). Since $\rho(t; f, \hat{n})$ has been normalized with the noise PSD, it is called the pixel SNR. Cross-power spectrograms such as in Fig. 1 are made by repeating this over many segments of data.

Pattern recognition or clustering algorithms pick out a cluster of pixels in the spectrogram, representing a possible signal, e.g., one monotonically evolving in frequency. The SNRs of these pixels are then summed up to give the detection statistic. The normalized cluster SNR of the track, Γ_c is given by




FIG. 1. An example of a cross-power spectrogram in the frequency-time plane, with a loud simulated signal added to it, which is visible as a narrow track. The color of the pixels indicates the SNR.

$$\Gamma_c = \frac{1}{\mathcal{N}} \sum_c \rho_i, \quad (3)$$

where i indexes over all the pixels in the cluster c , and \mathcal{N} is an empirically chosen normalization factor. Γ_c plays the role of the detection statistic. The cluster with the largest Γ_c will be the trigger for the spectrogram, and the distribution of Γ_c is computed for spectrograms containing pure noise to measure the background for the search. A good clustering algorithm finds an optimal cluster which samples as much of a potential signal as possible, without a corresponding increase in the background.

While the existing STAMP analyses have generally used data from two detectors, the formalism could be extended to more detectors. One way would be running the clustering algorithm separately on cross-power spectrograms computed from each pair of detectors, and demanding the triggers be coherent among the pairs. Alternatively we could run the clustering algorithm on the combined spectrogram from all the pairs of detectors (We refer the readers to Sec. III E of Ref. [13] for more details).

III. CLUSTERING ALGORITHMS

In this section, we describe the two clustering algorithms we use here, HMM tracking (Sec. III A) and seedless clustering (Sec. III B), using the cross-power spectrograms computed as described in Sec. II.

A. HMM tracking

A HMM is a memoryless, probabilistic state automaton based on a Markov process, composed of the hidden state variable $q(t) \in \{q_1, \dots, q_{N_q}\}$ and the measurement variable

$o(t) \in \{o_1, \dots, o_{N_o}\}$ sampled at discrete times $t \in \{t_0, \dots, t_{N_T}\}$. A full description of HMM formulation and the computationally efficient Viterbi algorithm [24] used for solving the HMM can be found in Ref. [15].

We track $q(t) = f_{\text{gw}}(t)$ in a Markov chain, where $f_{\text{gw}}(t)$ is the GW frequency at time t . The discrete hidden states q_i are mapped one-to-one to the frequency bins in the cross-power map, with bin size $\Delta f = m/T$, where m is a coarse-graining integer coefficient (see Sec. IV A). We choose a constant k to satisfy

$$\left| \int_t^{t+T} dt' \dot{f}_{\text{gw}}(t') \right| \leq (k-1)\Delta f, \quad (4)$$

for $0 \leq t \leq T_{\text{obs}}$, where \dot{f}_{gw} is the first time derivative of the GW signal frequency. The HMM emission probability at each discrete time, defined as the likelihood of hidden state q_i being observed in state o_j , is given by [15]

$$L_{o_j q_i} = P[o(t_n) = o_j | q(t_n) = q_i]. \quad (5)$$

Here we leverage the cross-power pixel SNR in Eq. (1), and define the emission probability over each time interval $[t, t+T]$ as

$$L_{o(t) q_i} = P[o(t) | f_i \leq f_{\text{gw}}(t) \leq f_i + \Delta f] \quad (6)$$

$$\propto \exp[\rho(t; f_i)]. \quad (7)$$

We also choose $T \leq 100$ s such that the Earth rotation can be neglected during the interval $[t, t+T]$ in the frequency range of interest. The transition probability of q_i from time t_n to t_{n+1} is defined as [15]

$$A_{q_j q_i} = P[q(t_{n+1}) = q_j | q(t_n) = q_i], \quad (8)$$

which depends on the signal evolution characteristics. Here we consider a model-agnostic, long-transient signal whose frequency rapidly decreases, e.g., a signal from a binary neutron star postmerger remnant. Assuming that the signal frequency can be approximated by a negatively biased random walk, with frequency change over each segment T uniformly distributed in range $[0, (k-1)m/T]$, i.e., $0 \leq f_{\text{gw}}(t_n) - f_{\text{gw}}(t_{n+1}) \leq (k-1)\Delta f$ [see Eq. (4)], we adopt the transition probabilities

$$A_{q_{i-j} q_i} = \frac{1}{k}, \quad (9)$$

with all other entries being zero. In Eq. (9), j takes integer values $0 \leq j \leq k-1$. We can always adjust Eq. (8) in searches for other types of signals. Since we have no independent knowledge of f_{gw} , we choose a uniform prior, viz.

$$\Pi_{q_i} = N_Q^{-1}. \quad (10)$$

The probability that a hidden state path $Q = \{q(t_0), \dots, q(t_{N_T})\}$ gives rise to an observed sequence $O = \{o(t_0), \dots, o(t_{N_T})\}$ via a Markov chain equals

$$P(Q|O) = L_{o(t_{N_T})q(t_{N_T})} A_{q(t_{N_T})q(t_{N_T-1})} \cdots L_{o(t_1)q(t_1)} \times A_{q(t_1)q(t_0)} \Pi_{q(t_0)}. \quad (11)$$

The most probable Viterbi path is the maximum *a posteriori* track, which maximizes $P(Q|O)$. The detection statistic is the cluster SNR as defined by Eq. (3) of the optimal Viterbi path.

B. Seedless clustering

We now briefly describe the seedless clustering algorithm following Refs. [12,14]. The algorithm attempts to pick out the morphology of a potential signal in the cross-power spectrogram by drawing tracks from a template base. In principle the template could have any possible form. In the case of compact binary coalescence sources for example, one could employ very specific templates drawing upon precise models of GW radiation from them [25,26]. Yet, in the presence of uncertainty about the sources and morphology of astrophysical signals, using quadratic Bézier curves is a good tradeoff between sensitivity and computational cost. Quadratic Bézier curves provide good sensitivity to many monotonically evolving waveform models, and have been applied in several searches conducted in the past [6,8,27].

In practice we pick three points (i.e., three pixels) randomly within the spectrogram with the only condition being that the frequency evolution between them be monotonic. The three pixels $N_i = (f_i, t_i)$ are then fit with quadratic curves parametrized by ξ :

$$\begin{bmatrix} f(\xi) \\ t(\xi) \end{bmatrix} = (1-\xi)^2 N_0 + 2(1-\xi)\xi N_1 + \xi^2 N_2. \quad (12)$$

Each Bézier template is defined as one choice of (N_0, N_1, N_2) , which completely describe a quadratic curve in Eq (12). Usually a total number of templates $N_{\text{temp}} \sim 10^6$ are used for a single spectrogram [28]. The SNR for the cluster obtained from one template, Γ_c , is defined as the sum of the SNRs of all pixels along the quadratic curve, and is again calculated using Eq. (3). These are the triggers for the search and the loudest trigger for a spectrogram is picked as the prospective signal candidate for further scrutiny.

IV. SENSITIVITY AND COST

A. Sensitivity analysis

In this section, we compare the sensitivities obtained from HMM tracking and seedless clustering. We make comparisons using two different sizes of spectrograms. The “long duration” spectrograms are 15 000 seconds long, made of short Fourier transforms (SFTs) of 100 seconds of data coarse-grained to 1-Hz frequency resolution (i.e., $m = 100$). The same configuration of spectrograms has been used in Ref. [8] to search for long-duration postmerger remnant signals. The “intermediate duration” spectrograms are 500 seconds long with 1 s SFTs and 1 Hz bin sizes (i.e., $m = 1$). Spectrograms of this intermediate size have been used extensively in the past [29], most recently in Refs. [6,27]. For each configuration, we make a comparison between the two algorithms using (1) Gaussian data recolored to the PSD of the first observing run (O1) of Advanced LIGO, and (2) O2 data from Advanced LIGO Hanford and Livingston detectors, with an unphysical time shift between them [30]. The former sets an ideal scenario for comparison, while the latter aims to accurately capture the impact of non-Gaussian, non-stationary artifacts in real interferometer data [31].

The long and intermediate duration analyses are described in Secs. IV C and IV D, respectively, while the waveform models used for simulated signals are described in IV B. For the sake of simplicity and for reducing the computational cost, we fix the sky position of all simulated signals to be the same as GW170817 [5], and it is also assumed to be known in the search. All comparisons of sensitivities are made at a false alarm probability (FAP) less than 10^{-4} . Finally, the search configurations used for HMM tracking and seedless clustering are listed in Table I, with column 2 and 3 for long and intermediate-duration spectrograms, respectively.

TABLE I. Search configurations for long and intermediate-duration spectrograms. The last two rows are parameters for HMM or seedless only. From top down, the parameters stand for the frequency band searched, segment duration, frequency resolution, spectrogram duration, HMM configuration constant [see Eq. (4)], and total number of seedless templates used.

Parameters	Long duration	Intermediate duration
f	30–1800 Hz	30–1800 Hz
T	100 s	1 s
Δf	1 Hz	1 Hz
T_{obs}	15 000 s	500 s
N_T	150	500
k (HMM)	10	10
N_{temp} (Seedless)	10^6	10^6

B. Signal models

1. Magnetar model

For both the long and intermediate-duration spectrograms, we simulate synthetic signals using a neutron star spin-down model, although the search itself is model-agnostic. This model characterizes the gravitational wave radiation from an nonaxisymmetric long-lived post-merger remnant. The remnant might be spinning down due to GW radiation or electromagnetic radiation or some combination thereof. This model has been used in setting limits for post-merger GW emission from GW170817 in both the intermediate-duration and long-duration searches [6,8]. The frequency evolution of the rapidly spinning down signal is given by [32]:

$$f_{\text{gw}}(t) = f_{\text{gw}}(0) \left(1 + \frac{t}{\tau}\right)^{\frac{1}{1-n}}, \quad (13)$$

where n is the braking index defined via $\dot{f}_{\text{gw}} \propto f_{\text{gw}}^n$, $\tau \propto f_{\text{gw}}^{1-n}/(1-n)$ is the spin-down timescale [8], and $f_{\text{gw}}(0)$ is the starting frequency at reference time $t = 0$. The gravitational-wave strain amplitude is given by [32]:

$$h_0(t) = \frac{4\pi^2 G I_{zz} \epsilon f_{\text{gw}}^2(0)}{c^4 D} \left(1 + \frac{t}{\tau}\right)^{\frac{2}{1-n}}, \quad (14)$$

where G is Newton’s gravitational constant, I_{zz} is the principal moment of inertia of the neutron star, ϵ is its equatorial ellipticity, and D is the distance to the source.

In Table II, rows 1–2 and 3–4 list the parameters of the synthetic magnetar signals in the long and intermediate-duration analyses, respectively. For all waveform models, detection efficiency at a particular root-sum-squared strain amplitude h_{rss} is defined as the fraction of simulated signals recovered given a FAP of less than 10^{-4} . In the frequency domain, h_{rss} is defined as

$$h_{\text{rss}} = \sqrt{2 \int_{f_{\text{min}}}^{f_{\text{max}}} df (|\tilde{h}_+(f)|^2 + |\tilde{h}_\times(f)|^2)}, \quad (15)$$

where \tilde{h}_+ and \tilde{h}_\times are strain amplitudes of the waveform in frequency domain for the + and \times polarizations, respectively, and f_{min} and f_{max} are the minimum and maximum frequencies of the waveform in the frequency band being analyzed, respectively.

2. Accretion disk instability model

For intermediate-duration spectrograms, we also test with a different model based on instabilities of accretion disks (ADI) around black holes. These waveforms are parametrized by the mass of the blackhole M_{BH} , dimensionless Kerr spin parameter a^* , and the fraction of mass forming inhomogeneities in the disk η . The inhomogeneities are modeled to behave as a binary system and act as a source of

TABLE II. Parameters of the magnetar and ADI models used to generate synthetic signals in Secs. IV C and IV D. Magnetar models E and M are used for long-duration signal simulations. Magnetar model A and B are used for intermediate-duration signal simulations. The same parameters are used for both Gaussian and time-shifted real interferometer data. The two ADI models are used for intermediate-duration signal simulations, both of which assumes a disk mass of $1.5 M_{\odot}$.

Model	$f_{\text{gw}}(0)$ (Hz)	τ (s)	n	Duration (s)	$\cos \iota$
Magnetar E	1k	10^4	2.5	10^4	1
Magnetar M	2k	10^4	2.5	10^4	1
Magnetar A	1k	10^2	3	10^3	1
Magnetar B	2k	10^2	3	10^3	1

Model	M_{BH}	a^*	η	Duration (s)	f (Hz)
ADI B	$10 M_{\odot}$	0.95	0.2	9.4	110–209
ADI C	$10 M_{\odot}$	0.95	0.04	236	130–251

gravitational waves. We refer to Refs. [29,33,34] and the references therein for more details about these waveforms. The last two rows in Tables II list the parameters of the ADI signals simulated.

Model spectrograms of the waveforms models used in this paper are shown in the Appendix.

C. Long duration

Fig. 2 presents the detection efficiency for HMM tracking and seedless clustering, generated by injecting synthetic signals in Gaussian data recolored with O1 noise PSD (top panels) and real time-shifted O2 data (bottom panels). The dashed and solid curves indicate results from HMM tracking and seedless clustering, respectively. We inject simulated signals at 18 amplitude levels chosen to be uniform in log amplitude. At each amplitude level, 150 simulated signals are used for the Gaussian case and 80 for the real case. These discrete results are then fit to a sigmoid to generate the efficiency curves in Fig. 2.

We see that HMM tracking outperforms seedless clustering ($N_{\text{temp}} = 10^6$). For example, the strain h_{rss} needed by HMM for an efficiency of 0.9 is lower than seedless by about a factor of two in all cases in Fig. 2. A part of the gain comes from the fact that magnetar signal curves are nonquadratic. The quadratic Bézier curves used in seedless clustering do not fit the signal well, while the HMM tracking does not assume a particular shape of the signal curve. One can expect that a better template (for instance a template made of cubic Bézier curves or from the waveform model itself) will give better sensitivity, albeit at a substantial increase in computational cost or loss in sensitivity to other waveforms.

We also compare the sensitivity obtained here using HMM tracking with cross-power spectrograms to the existing HMM method used in Ref. [8], which operates on normalized power in SFTs summed over multiple detectors, i.e., $\sum_X \tilde{x}_i^X \tilde{x}_i^{X*}$, where i indexes the frequency bins of the normalized SFT \tilde{x} , and X indexes the detector [21]. With simulations done on time-shifted O2 data, we

compute the 90% sensitive distance, $d^{90\%}$, i.e., the largest distance at which 90% of the injected signals can be recovered for HMM run on cross-power spectrograms. Using the same moment of inertia of $4.38 \times 10^{38} \text{ kg m}^2$ and the maximum possible ϵ as described in Ref. [8], we obtain $d^{90\%} \approx 0.5$ Mpc and 0.2 Mpc for magnetar E and M models, respectively. This is done by calculating h_0 corresponding to the h_{rss} value required for 90% efficiency, and using Eq. (14) to convert h_0 to the limit on distance. While these are still astrophysically unrealistic distances for a source like GW170817, they are significantly better than the $d^{90\%}$ values ($d^{90\%} \approx 0.064$ Mpc and 0.035 Mpc) quoted for similar waveforms in [8] using the existing HMM method.

Note that there are some differences between the simulations in this paper and in Ref. [8]: (1) We use braking index of $n = 2.5$ here as opposed to $n = 5$ in [8]; (2) We use $\cos \iota = 1$ (the inclination of the source) here as opposed to randomized $\cos \iota$ in [8]; (3) The FAP in this paper and in [8] are $\leq 10^{-4}$ and 10^{-2} , respectively. Although setting $\cos \iota = 1$ can improve $d^{90\%}$ by a factor of 2–3 compared to randomized $\cos \iota$, $d^{90\%}$ for signals with $n = 5$ are generally better than $n = 2.5$ by a factor of < 2 [21]. Combined with the much more stringent FAP adopted in this paper, these results demonstrate that using HMM operated on cross-correlated spectrograms outperforms its usage on incoherent SFT powers from two detectors. The improvement is not unexpected given that the cross-power SNR statistic demands that the phase difference of the signals between two detectors be consistent with the sky position [see Eq. (1)], whereas the SFT power spectrograms effectively marginalize over the phase and the sky position [21]. A more detailed study of the difference between cross-power and SFT spectrograms is out of the scope of this paper.

D. Intermediate duration

In the intermediate duration search, we inject simulated signals at 18 amplitude levels for the magnetar model and at




FIG. 2. Detection efficiencies (i.e., the rates of correctly recovering injections) for long-duration simulations based on the magnetar models E and M. The curves shown are generated from sigmoid fits of the discrete injection results. HMM tracking generally performs better than seedless clustering for these waveforms. The top panels are from simulations in Gaussian noise. The bottom panels are with simulations of the same waveform injected into time-shifted real O2 data. The results demonstrate that the gain in sensitivity from HMM tracking is not affected by non-Gaussian, non-stationary noise in real interferometric data. We note that the Gaussian data is recolored with O1 noise PSD, and the real data is from the O2 run. The colored regions represent 1σ binomial uncertainty in detection efficiency.

22 amplitude levels for the ADI model, uniform in log amplitude. At each amplitude level, 100 simulated signals are injected. Discrete results are fit to a sigmoid to generate the efficiency curves in Figs. 3 and 4.

Figure 3 shows the relative performance of HMM and seedless clustering using magnetar models A and B, in both recolored Gaussian noise (O1 PSD) and O2 real interferometric data. HMM tracking still outperforms seedless clustering, although in these intermediate-duration simulations in real data, the sensitivity gain from HMM is not as significant as the long-duration search.

The same intermediate-duration search configuration (column 3 in Table I) is used for recovering accretion-disk instability (ADI) simulations. For illustration purpose, we perform the ADI simulations in Gaussian noise only. The detection efficiencies for these waveforms are shown in Fig. 4. For the ADI models, the recovery efficiencies from HMM tracking and seedless clustering are generally comparable. Unlike the magnetar models described above, the ADI signal waveforms are better sampled by quadratic Bézier curves, and hence seedless is expected to produce similar sensitivity as to HMM.




FIG. 3. Detection efficiencies for intermediate-duration simulations based on the magnetar models A and B. The curves shown are generated from sigmoid fits of the discrete injection results. HMM tracking generally performs better than seedless clustering. The top panels are from simulations in Gaussian noise. The bottom panels are with simulations of the same waveform injected into time-shifted real O2 data. For intermediate-duration signals, the gain in sensitivity from HMM tracking is slightly affected by non-Gaussian, non-stationary noise in real interferometric data. The colored regions represent 1σ binomial uncertainty in detection efficiency.

E. Computing cost

We have demonstrated that HMM tracking provides detection efficiencies better than or at least comparable to seedless clustering in the parameter space considered here. In this section, we show that HMM tracking significantly outperforms seedless clustering with respect to run time and computational cost, and briefly explain the reason.

We have tested the run time of both HMM and seedless methods (with identical scenarios and configurations) for both the intermediate and long-duration spectrograms. Over 2500 realizations of the intermediate-duration spectrograms, the median run time of HMM tracking was ~ 35 times shorter than seedless clustering with 10^6 templates

(see Table III). A similar test for the long-duration spectrograms produces an improvement of a factor of ~ 23 for HMM compared to seedless. The gain is less significant in long-duration spectrograms probably because the coarse-grained long-duration spectrograms consist of a smaller number of pixels, N_{pixel} , compared to the intermediate-duration spectrogram (see N_{pixel} in Table III). Although the amount of raw data in the 15 000 second long spectrograms is much larger than the intermediate-duration ones, coarse-graining is generally employed [35] to reduce the volume of data being analyzed. Coarse-graining is done by averaging both the cross and auto-power in finer frequency bins to give coarser bins. While averaging reduces the

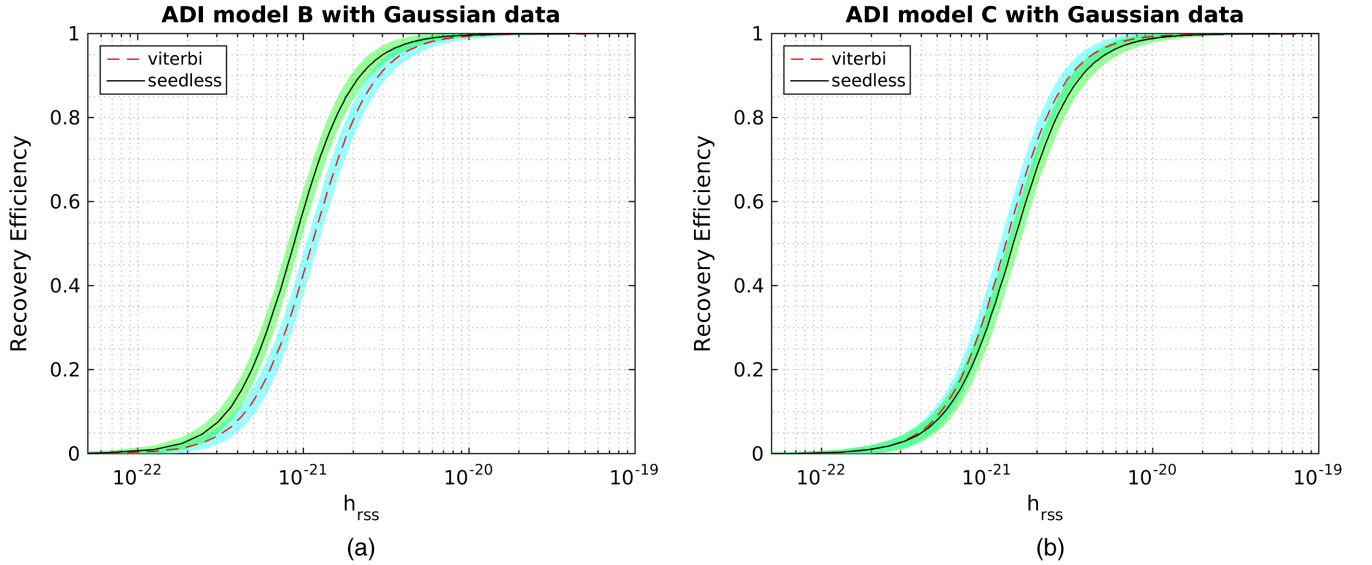


FIG. 4. Detection efficiencies for intermediate-duration simulations based on the ADI models B and C in Gaussian noise. The curves shown are generated from sigmoid fits of the discrete injection results. The performance of two methods is generally comparable. The colored regions represent 1σ binomial uncertainty in detection efficiency.

computational cost of analyzing long duration spectrograms, it also leads to a loss of sensitivity. The natural Fourier transform frequency resolution of $1/T = 10$ mHz ($T = 100$ s) in the long-duration spectrograms would require a prohibitively large number of seedless templates for analysis. The efficiency of HMM tracking makes it a promising tool to run deeper searches over spectrograms with finer frequency resolution.

We now briefly discuss the reason why using HMM shows a significant improvement in computational cost. The HMM tracking uses the dynamic programming algorithm, Viterbi, which reduces the total number of comparisons required to find the optimal path from $N_Q^{N_T+1}$ to $(N_T + 1)N_Q^2$ in a spectrogram with $N_{\text{pixel}} = N_Q N_T$ [15,19]. When matrix $A_{q_j q_i}$ only contains ten nonzero terms along the diagonal, the total number of comparisons reduces to $10N_{\text{pixel}}$. Hence in the intermediate and long duration spectrograms, the total numbers of comparisons are 8.9×10^6 and 2.7×10^6 , respectively. As a dynamic

process, at each step, the algorithm only records $10N_Q = 1.8 \times 10^4$ paths for both configurations, but effectively ensures that the optimal one is kept. This is significantly more efficient than fitting 10^6 curves and summing up the N_T SNR pixels for each curve as is done in seedless.

V. CONCLUSION

In this paper, we describe two clustering strategies for long-transient gravitational-wave searches, both operating on precalculated, cross-power spectrograms. We conduct a large number of simulations—about $\sim 16\,000$, and $\sim 21\,000$ for the long-duration (15 000 s) and intermediate-duration (500 s) signals, respectively—in both Gaussian noise and real interferometric data. In the simulations and comparison carried out in this paper, we mainly focus on the magnetar model that is adopted in BNS post-merger remnant and other long-transient searches. We have demonstrated that HMM tracking can produce detection efficiency better than or at least similar to seedless clustering, and reduce the computing cost significantly, based on the same sets of spectrograms. HMM tracking can also be applied to track a variety of signal models in addition to the ones tested above, e.g., nonmonotonic signals, by adjusting transition probabilities (see Ref. [17]). In addition, a small improvement in the sensitive distance will give a relatively large improvement in sensitive volume. With much lower computational requirements, HMM tracking method can be a good option in unmodeled all-sky searches for long-transient signals. In model-agnostic, computationally challenging searches, HMM tracking can prove to be a superior strategy to parse spectrograms.

The HMM tracking algorithm operated on cross-power spectrograms also outperforms the existing HMM tracking

TABLE III. Total number of pixels and median run time of HMM tracking and seedless clustering for intermediate and long duration spectrograms (over 2500 realizations). We note that the tests were run on a computing cluster with machines containing various intel CPU generations, and the run time depends on the CPU architecture. Hence the improvement ratio in the last column is of more interest in the comparison.

	N_{pixel}	HMM (s)	Seedless (s)	Ratio
Intermediate	8.9×10^5	4	145	36.25
Long	2.7×10^5	1.6	37.8	23.63

operated on SFT power spectrograms used in previous searches for long-duration BNS postmerger signals [8]. This new implementation can serve as a more sensitive and efficient alternative in future analyses of the same kind. Finally, this work might also open the window to otherwise prohibitively expensive all-sky long-duration searches (with $\sim 10^4$ -s or longer spectrograms), given the significant reduction in computational cost by using HMM tracking.

ACKNOWLEDGMENTS

We are grateful to Maxime Fays, Rich Ormiston, Stuart Anderson and Vuk Mandic for comments and informative discussions. S. B. acknowledges support in part by the Hoff Lu Fellowship at the university of Minnesota, and by NSF Grant No. PHY-1806630. L. S. is a member of the LIGO Laboratory. M. W. C. is supported by the David and Ellen Lee Postdoctoral Fellowship at the California Institute of Technology. The authors are thankful for the computing resources provided by LIGO Laboratory. LIGO was constructed by the California Institute of Technology and Massachusetts Institute of Technology with funding from the National Science Foundation, and operates under cooperative agreement No. PHY-0757058. Advanced LIGO was built under Grant No. PHY-0823459. The research was also supported by Australian Research Council (ARC) Discovery Project DP170103625 and the ARC Centre of Excellence for Gravitational Wave Discovery CE170100004. This paper carries LIGO Document Number LIGO-P1900047.




FIG. 5. Sample long-duration spectrogram of the magnetar model E. The spectrogram is for time duration 15000 s and frequency band 30–1800 Hz. The signal is visible from about 1000–500 Hz.

APPENDIX: SPECTROGRAMS

We show sample spectrograms of the waveform models used in this study in Fig. 5 for the long-duration spectrograms and Figs. 6–7 for the intermediate-duration ones.




FIG. 6. Sample intermediate-duration spectrogram of the magnetar model A. The spectrogram is for time duration 500 s and frequency band 30–1800 Hz. The signal is visible from about 1000–500 Hz. The horizontal bars and vertical lines are noisy frequencies which have been notched out or segments which have been vetoed by data-quality cuts.




FIG. 7. Sample intermediate-duration spectrogram of the ADI model C. The spectrogram is for time duration 500 s and frequency band 30–1800 Hz. A subband 30–1000 Hz is displayed here in order to show the injected signal more clearly. The horizontal bars are noisy frequencies which have been notched out.

- [1] J. Aasi *et al.* (LSC Collaboration), *Classical Quantum Gravity* **32**, 074001 (2015).
- [2] F. Acernese *et al.* (Virgo Collaboration), *Classical Quantum Gravity* **32**, 024001 (2015).
- [3] B. P. Abbott *et al.*, arXiv:1811.12907.
- [4] B. P. Abbott *et al.*, *Phys. Rev. Lett.* **119**, 161101 (2017).
- [5] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations *et al.*), *Astrophys. J.* **848**, L12 (2017).
- [6] B. P. Abbott *et al.* (LVC Collaboration), *Astrophys. J. Lett.* **851**, L16 (2017).
- [7] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. X* **9**, 011001 (2019).
- [8] B. P. Abbott *et al.*, *Astrophys. J.* **875**, 160 (2019).
- [9] M. H. P. M. van Putten and M. Della Valle, *Mon. Not. R. Astron. Soc.* **482**, L46 (2019).
- [10] P. Tanner and T. Eric, LIGO-11200204-v1: Burstegard: A hierarchical clustering algorithm.
- [11] R. Khan and S. Chatterji, *Classical Quantum Gravity* **26**, 155009 (2009).
- [12] E. Thrane and M. Coughlin, *Phys. Rev. D* **88**, 083010 (2013).
- [13] E. Thrane *et al.*, *Phys. Rev. D* **83**, 083004 (2011).
- [14] E. Thrane and M. Coughlin, *Phys. Rev. D* **89**, 063012 (2014).
- [15] S. Suvorova, L. Sun, A. Melatos, W. Moran, and R. J. Evans, *Phys. Rev. D* **93**, 123009 (2016).
- [16] S. Suvorova, P. Clearwater, A. Melatos, L. Sun, W. Moran, and R. J. Evans, *Phys. Rev. D* **96**, 102006 (2017).
- [17] L. Sun, A. Melatos, S. Suvorova, W. Moran, and R. J. Evans, *Phys. Rev. D* **97**, 043013 (2018).
- [18] R. Streit and R. Barrett, *IEEE Trans. Acoust. Speech Signal Process.* **38**, 586 (1990).
- [19] B. G. Quinn and E. J. Hannan, *The Estimation and Tracking of Frequency* (Cambridge University Press, Cambridge, England, 2001), p. 266.
- [20] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), *Phys. Rev. D* **95**, 122003 (2017).
- [21] L. Sun and A. Melatos, *Phys. Rev. D* **99**, 123003 (2019).
- [22] We focus on the comparison between HMM tracking and seedless clustering because seedless algorithms generally produce better sensitivity than seeded algorithms. We do not discuss seeded algorithms in this paper.
- [23] A common way to estimate the noise autopower is $P_I(t; f) = 2\overline{|s_I(t; f)|^2}$, where the overline indicates an average over the time segments neighboring t at frequency f .
- [24] A. Viterbi, *IEEE Trans. Inf. Theory* **13**, 260 (1967).
- [25] M. Coughlin, E. Thrane, and N. Christensen, *Phys. Rev. D* **90**, 083005 (2014).
- [26] M. Coughlin, P. Meyers, E. Thrane, J. Luo, and N. Christensen, *Phys. Rev. D* **91**, 063004 (2015).
- [27] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Classical Quantum Gravity* **35**, 065009 (2018).
- [28] While N_{temp} is tunable parameter, $\mathcal{O}(10^6)$ templates are somewhat of an optimum value. For example using ten times more templates does not result in a substantial increase sensitivity for many waveform models as shown in Ref. [12].
- [29] B. P. Abbott *et al.* (Virgo and LIGO Scientific Collaborations), *Phys. Rev. D* **93**, 042005 (2016).
- [30] The time-shift or the time-difference between the data from the detectors is set to be greater than the light travel time and the segment duration. This helps avoid any correlations between true gravitational wave signals in the data.
- [31] Some existing simulation results using Gaussian data recolored to the O1 PSD have been used in this study to understand the impact of non-Gaussian, nonstationary noise. Qualitatively, we expect the detection efficiencies obtained from the colored O2 Gaussian data to be comparable to or slightly better than those from the colored O1 Gaussian data. It is difficult to fully quantify the effect of glitches and other non-Gaussian artifacts in real data, and the loss of efficiency because of them. It depends on the types of glitches, the waveform model being considered, the efficacy of the glitch rejection algorithms used and the sampling algorithms. We defer any such analysis to a future work.
- [32] N. Sarin, P. D. Lasky, L. Sammut, and G. Ashton, *Phys. Rev. D* **98**, 043011 (2018).
- [33] M. H. P. M. van Putten, A. Levinson, H. K. Lee, T. Regimbau, M. Punturo, and G. M. Harry, *Phys. Rev. D* **69**, 044007 (2004).
- [34] M. H. P. M. van Putten, *Phys. Rev. Lett.* **87**, 091101 (2001).
- [35] E. Thrane, V. Mandic, and N. Christensen, *Phys. Rev. D* **91**, 104021 (2015).