

Unmodelled Detection of Overlapping Gravitational-Wave Signals with the Einstein Telescope

Stijn Heyde, Milan Wils, Francesco Cireddu, Tjonnje Li and Marc Moonen
KU Leuven, Leuven, Belgium

Abstract—The third generation of gravitational-wave detectors, including the Einstein Telescope, will detect hundreds of thousands of compact-binary coalescences per year, so that a rare Galactic core-collapse supernova would arrive in a record likely already convoluted by binary inspirals. Coherent burst pipelines assume a single source per trigger and blend such overlapping signals. This work treats the overlap as a multi-source detection and direction-of-arrival problem on a small, heterogeneous detector network. We first quantify the temporal and time-frequency overlap probability for a realistic one-year inspiral population. For an inspiral cutoff frequency of 3 Hz, overlap is highly likely. We then restructure an existing modular burst pipeline so that every stage carries a list of source hypotheses, and extend the coherent beamforming statistic into a joint maximum-likelihood estimator for two simultaneous sources, with a null-stream criterion that determines how many sources are present per time-frequency bin. For a network consisting of the Einstein Telescope plus Cosmic Explorer plus LIGO-India the method achieves single-digit-degree localisation for single sources and recovers all pairwise two-source combinations, consistently outperforming a strict per-pixel single-source assignment.

Index Terms—burst detection, coherent network analysis, gravitational waves, null stream, overlapping signals, sky localisation

I. INTRODUCTION

The Einstein Telescope (ET) and Cosmic Explorer (CE) will extend gravitational-wave (GW) sensitivity by an order of magnitude and push the low-frequency cutoff to a few hertz [1], [2], [8]. The resulting detection rate reaches $\mathcal{O}(10^5)$ binary neutron star (BNS) and $\mathcal{O}(10^6)$ binary black hole (BBH) coalescences per year [3], with individual BNS inspirals spending hours in band. A Galactic core-collapse supernova (CCSN), occurring about once per 60 years [4], would therefore arrive in a detector record likely already occupied by one or more compact-binary coalescences (CBCs).

Coherent burst pipelines such as coherent WaveBurst [5], X-Pipeline [12], and the Modular Data Analysis Pipeline (MDAP) [6] are built around a single-source assumption: each cluster of excess-power time-frequency (TF) pixels is treated as one event, and signal multiplicity is never propagated downstream. When two unrelated signals share a TF neighbourhood, the cluster potentially blends them and the reconstruction could be biased toward the more phase-coherent contribution. For such a rare event, a missed detection would forfeit a once-in-a-generation observation.

Prior work on overlapping GW signals has focused on the compact-binary–compact-binary case, where both signals are well modelled and matched filtering or joint Bayesian inference techniques are used [7], [11]. The unmodelled case, in which at least one source (here a CCSN) has no waveform template, is far less studied and is the gap addressed here.

From a signal-processing standpoint the problem is multi-source detection and direction-of-arrival (DOA) estimation with a small, heterogeneous detector network. We do not attempt blind source separation, which is ill-suited to the small detector network [7], but instead relax the single-source assumption inside the coherent-likelihood framework. This paper makes three contributions: (i) it quantifies the TF overlap probability for the realistic 3G observing scenario (Section III); (ii) it restructures the MDAP data flow to carry multiple source hypotheses (Section IV); and (iii) it develops a joint maximum-likelihood (ML) localisation statistic with a null-stream source-count criterion (Section V). Section VI validates the method on the ET+CE+A1 network.

II. SIGNAL MODEL

A detector k responds to the two GW polarizations $h_+, h_\times$ through its antenna-pattern functions $F_+^{(k)}, F_\times^{(k)}$ [13]. In the frequency domain its output is

$$\tilde{d}_k(f) = [F_+^{(k)} \tilde{h}_+(f) + F_\times^{(k)} \tilde{h}_\times(f)] e^{-i2\pi f t_k} + \tilde{w}_k(f), \quad (1)$$

where $t_k = t_{\text{geo}} + \mathbf{x}_k \cdot \hat{n}/c$ is the arrival time at detector position $\mathbf{x}_k$ for a source from sky direction $\hat{n}$, and $\tilde{w}_k$ is detector noise. Whitening each channel by its amplitude spectral density and phase-aligning to the geocentre yields, for a trial direction $\hat{n}$, the K -dimensional data vector $\tilde{\mathbf{d}}(f, t; \hat{n})$. Here K counts the independent GW-sensitive streams: $K_{\text{eff}} = 4$ for the ET+CE+A1 network (two co-located ET streams, one each from CE and LIGO-India A1). The three co-located ET interferometers alone provide only two streams for two polarization unknowns, leaving the likelihood flat over sky position; the geographically separated CE and A1 break this degeneracy through the time-of-flight delays $\mathbf{x}_k \cdot \hat{n}/c$, which is what makes sky localisation, and source separation by sky position, possible. Analysis proceeds per pixel of a short-time Fourier transform (STFT) of the data.




Fig. 1. CCSN-BNS TF overlap probability versus inspiral cutoff frequency $f_{\min}$ (hybrid STFT, five CCSN source curves and their mean, averaged over ET’s detectors and viewing orientations). At full bandwidth $P_{\text{ov}} \approx 87\text{--}91\%$ for the heavier progenitors, dropping sharply once the sub-16 Hz bands are excluded.

III. TIME-FREQUENCY OVERLAP QUANTIFICATION

As a first check we test the temporal overlap, sliding a window with the duration of a CCSN across a day-long observation and counting the placements that land on an in-band inspiral. This overlap is essentially 100% at the few-hertz frequencies ET reaches, since any CCSN in a day-long observation lands on a BNS inspiral. The more relevant question is how much of that coincidence falls in the same TF region. To establish that overlap is a realistic concern we compute the probability that a Galactic CCSN at 10 kpc shares the TF plane with a member of a realistic one-year CBC population [9], [10], for the ET sensitivity [14], in both STFT and discrete wavelet representations. Five CCSN waveforms are used: three neutrino-driven models [23] with progenitor masses 9.6, 15 and 25 $M_{\odot}$ (capturing core bounce, convection, and the standing accretion shock instability) and two magnetorotational models [24] from a 39 $M_{\odot}$ progenitor at initial magnetic fields 10^{10} and 10^{12} G, each taken at equatorial and polar viewing directions. Each signal is reduced to a binary TF mask by a pixel-selection threshold set at the noise floor. Each spectrogram is normalized so that its expected power is unity in pure noise, and every pixel at or above that level (threshold $\sigma = 1$) is marked active, so the mask retains exactly those pixels capable of overlapping. Two representations are compared. In the hybrid STFT the CBC mask uses $\Delta f = 4$ Hz, while the CCSN mask uses $\Delta f = 4$ Hz in its four lowest-frequency bins and $\Delta f = 16$ Hz above; a discrete wavelet transform (sym14, nine levels) provides the second representation. The two masks are mapped onto a common grid, on which the CCSN mask is tested for pixel coincidence against each inspiral in the population, and the fraction of inspiral placements that coincide gives the overlap probability P_{ov} .

In the TF plane, Fig. 1 shows P_{ov} as a function of the inspiral lower cutoff frequency $f_{\min}$. At full bandwidth $P_{\text{ov}} \approx 87\text{--}91\%$ for the heavier progenitors. The lightest progenitor reaches only $\approx 55\%$, but this is a detector-sensitivity




Fig. 2. Modified MDAP architecture (adapted from [6]). Products left of sky localisation are single objects; from sky localisation onward every product is an N -indexed list. The likelihood stage emits per-source coherent maps and a per-pixel single/multi-source classification, clustering produces N pixel masks, and event detection and reconstruction run once per source.

artefact rather than a source property, as one ET interferometer registers too few low-frequency pixels at threshold and the resulting sparse mask pulls the per-interferometer average down. Although the probability of a coincidence is high, each coincidence touches only a small fraction (a few percent) of the CCSN footprint, concentrated below 16 Hz in the energy-dense memory band where the CCSN deposits most of its energy [25], so the overlap is confined but energetically important. BBH overlap is expected to be at least as frequent given the higher BBH rate, so the BNS result acts as a conservative lower bound for both populations. Overlapping signals are therefore highly likely whenever a CCSN occurs, motivating a pipeline that can detect more than one source.

IV. MULTI-SOURCE PIPELINE

The MDAP [6] is a modular chain whose stages exchange fixed data products: TF transform, likelihood generation, clustering, event detection and reconstruction. The single-source assumption is enforced not in any single module but in the data products themselves, each describing one event.

Two changes therefore reach deeper than swapping a module. The original chain has no notion of source direction, so a new sky-localisation module (Section V) is added ahead of the likelihood stage, and the data flow is refactored (Fig. 2) so that from sky localisation onward every product is an N -indexed list rather than a single object. Each source i then carries a descriptor $(\hat{n}_i, \psi_i, \chi_i, \text{model_type}_i)$ where ψ is the polarization angle and χ the cross-polarization fraction, that fixes its signal model. The model_type selects either a one-dimensional signal subspace for a CBC, whose elliptical polarization collapses its response onto a single effective steering vector with (ψ, χ) assumed circular or estimated, or the full two-dimensional subspace for an unmodelled CCSN. These changes touch the pipeline backbone while leaving its modular structure intact, so individual stages can still be replaced independently.

Simultaneous BBH and CCSN injection and reconstruction




Fig. 3. End-to-end simultaneous BBH and CCSN injection and reconstruction for one of the K detectors (ET1) under realistic ET-D noise. *Left*: injected components: BBH (a), CCSN (b), noise (c), and combined strain (d). *Right*: per-source reconstruction overlaid on the injected signal for BBH (e) and CCSN (f). Both sources are recovered from a single blended data stream.

Within that structure the downstream stages receive new multi-source implementations. Multi-source data are injected by projecting each source’s polarizations onto every detector through its antenna patterns and per-detector delays, summing the contributions and adding PSD-coloured noise, with externally simulated CCSN waveforms read through the same interface as the internal CBC generator. The model-aware likelihood stage applies the per-source coherent statistic and null-stream classification of Section V uniformly to every TF pixel, evaluating all N hypotheses in one pass and emitting N coherent-energy maps together with a per-pixel single-versus multi-source label, with no energy threshold applied yet. Clustering then applies the standard cWB core-plus-halo construction [5], [6] independently to each of the N per-source maps, thresholding at a per-map percentile and dilating into a halo. Pixels in two or more of the resulting clusters form the multi-source set. Event detection and reconstruction run once per source. A $K \times 2$ solve recovers the two polarizations at single-source pixels, while overlap pixels use a joint $K \times 2N$ least-squares solve $\hat{\mathbf{h}} = (A^H A)^{-1} A^H \mathbf{d}$, exactly determined for $K_{\text{eff}} = 4$ and $N = 2$. An end-to-end test with a simultaneous BBH and CCSN injection (Fig. 3) identifies, separates and reconstructs both sources in a single pipeline pass. This is a proof of concept rather than a tuned pipeline. Each stage works end-to-end on multi-source data, but the clustering rules, threshold calibrations and model-type assignment leave room for future refinement that the modular structure is meant for.

V. JOINT MAXIMUM-LIKELIHOOD LOCALISATION

A. Single-source coherent statistic

For a single source, a signal from $\hat{n}$ enters the whitened data only through the two antenna patterns, so it lies on the two-dimensional plane spanned by the noise-weighted response vectors $\tilde{\mathbf{F}}_{+, \times} = (F_{+, \times}^{(1)}/\text{ASD}_1, \dots, F_{+, \times}^{(K)}/\text{ASD}_K)^\top$, which weight each detector inversely by its noise floor. For whitened Gaussian noise the log-likelihood reduces to the least-squares cost $-\|\tilde{\mathbf{d}} - \tilde{\mathbf{F}}_+ \tilde{h}_+ - \tilde{\mathbf{F}}_\times \tilde{h}_\times\|^2$, whose maximizer is the orthogonal projection of $\tilde{\mathbf{d}}$ onto that plane. At a wrong trial direction the phase-alignment is incomplete, so the projection interferes destructively and residual power leaks into the orthogonal complement. With an orthonormal basis $(\hat{\mathbf{e}}_+, \hat{\mathbf{e}}_\times)$ obtained by Gram–Schmidt, the coherent energy is

$$E_{\text{coh}}(f, t; \hat{n}) = |\hat{\mathbf{e}}_+^H \tilde{\mathbf{d}}|^2 + |\hat{\mathbf{e}}_\times^H \tilde{\mathbf{d}}|^2, \quad (2)$$

and the per-pixel sky statistic, following the coherent-network framework [5], [15], is $\Lambda_{\text{single}} = E_{\text{coh}}^2/E_{\text{total}}$ with $E_{\text{total}} = \|\tilde{\mathbf{d}}\|^2$. The factor $1/E_{\text{total}}$ weights the coherent energy by the coherent fraction, suppressing pixels dominated by incoherent power. This statistic belongs to the cWB family [5]. The form of [16] instead factors out a separate network-correlation quality factor, tuned for the two-detector LIGO network and not adopted here for the heterogeneous 3G network, and the present statistic is not algebraically identical to it. Accumulated over the pixels of a segment, $\Lambda_{\text{single}}(\hat{n})$ peaks at the source direction. Sky positions with higher network sensitivity




Fig. 4. Per-injection angular error (heatmap) for the BBH sweep at 20 Gpc, overlaid with iso-contours of the band-integrated ET+CE+A1 network sensitivity (normalized to its peak). High-error injections fall inside the low-sensitivity regions, confirming that localisation is governed by the network antenna response rather than the signal itself.

yield more coherent energy and a sharper peak, so localisation quality follows the antenna response across the sky (Fig. 4).

B. Null streams and source counting

The $K - 2$ directions orthogonal to the polarization plane form the null streams [17], carrying noise but no single-source signal. Their energy $E_{\text{null}} = E_{\text{total}} - E_{\text{coh}}$ vanishes (up to noise) at the true direction of a single source but retains residual power when two signals overlap, because no single 2D subspace absorbs both. The minimum null fraction over the sky, $\xi = \min_{\hat{n}} E_{\text{null}}/E_{\text{total}}$, is a natural per-pixel source-count classifier. Pixels with $\xi < \varepsilon_{\text{null}}$ are treated as single-source, the rest as overlap, with $\varepsilon_{\text{null}}$ calibrated on a noise-free coincident injections of two sources.

C. Joint two-source estimator

At overlap pixels the data follow a two-source model with true directions $\hat{n}_1$ and $\hat{n}_2$. The two directions must be estimated jointly: a cheaper CLEAN-style scheme [19] that locates one source, subtracts its beam, and iterates fails. For two-source data the single-source coherent energy at a trial direction $\hat{n} \neq \hat{n}_1, \hat{n}_2$ obeys Cauchy–Schwarz, with equality when the steering vector aligns with the mixture, placing the maximum at a *ghost* direction that is neither true source. On a small co-located network that equality is attainable exactly, whereas on the heterogeneous ET+CE+A1 network the inter-detector phases and magnitudes cannot both match the mixture at any single direction, so the bound is never reached. The scheme still fails, because the ghost peak nonetheless exceeds the coherent energy at the true directions, where both phase and magnitude are misaligned with the mixture, so a greedy iteration locks onto a ghost and never converges.

We therefore use a joint maximum-likelihood formulation. With combined $K \times 4$ steering matrix $A = [A_1(\hat{n}_1) \ A_2(\hat{n}_2)]$, where each $A_i = [\mathbf{a}_+^{(i)} \ \mathbf{a}_\times^{(i)}]$ contains the per-source analogues of $\tilde{\mathbf{F}}_{+, \times}$ with the time-of-flight phase factor $e^{-i2\pi f \tau_k(\hat{n}_i)}$ folded in (since the data cannot be phase-aligned to two sources at once), the joint ML coherent energy is the projection energy of $\tilde{\mathbf{d}}$ onto the column space of A ,

$$E_{\text{coh}}(\hat{n}_1, \hat{n}_2) = \tilde{\mathbf{d}}^H A (A^H A)^{-1} A^H \tilde{\mathbf{d}}, \quad (3)$$

giving $\Lambda_{\text{joint}} = E_{\text{coh}}^2(\hat{n}_1, \hat{n}_2)/E_{\text{total}}$. For $K_{\text{eff}} = 4$ two free-polarization sources exhaust the data space ($2 + 2 = 4$) and leave no noise dimension. Constraining each CBC to its elliptical polarization, $\tilde{h}_\times = \pm i\chi \tilde{h}_+$, collapses its 2D subspace onto the 1D span of an effective steering vector $\mathbf{a}_{\text{eff}} = \mathbf{a}_+ \pm i\chi \mathbf{a}_\times$, restoring a noise dimension and making the joint problem well-posed. A Tikhonov term [18] regularizes $A^H A$ where the steering vectors lose independence.

Over a segment, pixels above the noise floor are selected and each is classified by ξ . The two statistics are accumulated into a single surface over the source pair,

$$\Lambda_{\text{total}}(\hat{n}_1, \hat{n}_2) = \sum_{(f,t) \in \mathcal{S}} \max(\Lambda_{\text{single}}(\hat{n}_1), \Lambda_{\text{single}}(\hat{n}_2)) + \sum_{(f,t) \in \mathcal{M}} \Lambda_{\text{joint}}(\hat{n}_1, \hat{n}_2), \quad (4)$$

where $\mathcal{S}$ and $\mathcal{M}$ are the single- and multi-source pixel sets. The max on single-source pixels assigns each pixel’s evidence to the better-fitting direction without double-counting, which would otherwise inflate the diagonal $\hat{n}_1 = \hat{n}_2$ and bias the solution toward nearby pairs. The inverse $(A^H A)^{-1}$ in (3) becomes singular when the steering vectors lose independence (coincident directions, or a CBC lying in the CCSN subspace). The Tikhonov term keeps the surface finite there and recovers the single-source statistic on that locus. The peak of (4), evaluated on a coarse-to-fine grid over the pair, gives the two sky positions. Per-source marginal maps are slices of the surface with the other source fixed at its found position. Credible regions are calibrated empirically rather than from a posterior distribution [21]. For each of N injections the fractional sky area whose statistic equals or exceeds the value at the true direction is recorded, and the X -th percentile of this distribution is the $X\%$ credible area, which by construction contains the true direction in $X\%$ of injections.

D. Validating the polarization constraint

The circular assumption sets $\chi = 1$ (face-on) for every CBC, the strongest form of the constraint. Its cost is bounded by comparing three treatments on single-source injections: *circular* itself ($\chi = 1$), an unattainable *oracle* given the true (ψ, χ) , and a *scan* that grid-maximizes the statistic over $\psi \in [0, \pi/2)$ and $\chi \in [-1, 1]$. A single pixel admits a closed-form (ψ, χ) maximizer but is too noisy to use, so the scan aggregates over pixels. The scan beats the circular assumption only for inclinations $\theta_{jn} \gtrsim 60^\circ$ and only out to roughly 10 Gpc (BBH) or 1 Gpc (BNS), beyond which the estimates are too noisy and circular is no worse. For two overlapping sources the scan is wrapped in a hard, classification-style expectation-maximization coordinate ascent [20] that alternates a winner-take-all pixel-to-source assignment with a per-source (ψ, χ) scan. Because the gains are marginal and confined to nearby high-inclination sources, the validation below adopts the circular assumption for CBCs, with the scan retained as an optional implementation in the pipeline.

VI. RESULTS

All runs use the ET+CE+A1 network with PSD-coloured Gaussian noise. CBC sources use the circular polarization assumption adopted above. Single-source performance is measured over 288 sky positions, area-weighted to represent a uniform sky draw. The CCSN is fixed at its standard Galactic reference distance of 10 kpc, and the CBC reference distances are chosen to give comparable network SNR, 10 Gpc for the BBH and 1 Gpc for the BNS. In the multi-source runs the CBCs stay at these reference distances while the CCSN sky position is drawn from a Galactic exponential thin-disk distribution following Adams et al. [22]. The two sources are timed to maximize their TF overlap, the worst case for the joint solver.

A. Single-source localisation

At these reference distances localisation reaches single-digit-degree median accuracy for all three source types (Table I). BBH and BNS behave near-identically, and the localisation quality depends strongly on the network’s sky-dependent sensitivity, with the largest errors in its low-sensitivity regions (Fig. 4). The CCSN keeps a small angular error but a large 90% credible region, because its credible area spreads over disjoint sky clusters rather than one contiguous patch.

B. Multi-source localisation

For two sources (Table II) all three pairwise combinations are recovered. The two degradation patterns differ. CBC+CBC is tail-heavy (the median roughly doubles while the 90% credible region grows far more, so a fraction of injections are lost to mutual interference while the rest retain near single-source accuracy), whereas CBC+CCSN degrades the full distribution.

To probe the transition between regimes we fix the CCSN at 10 kpc and sweep the CBC companion in distance, comparing three configurations: the full joint ML with two-source pixels, a strict per-pixel single-source assignment, and a single-source mode that returns only the dominant source. At close CBC distance the single-source mode locks onto the loud CBC and misses the CCSN entirely. At large CBC distance the CBC drops below detectability, where strict assignment corrupts the CCSN by forcing a two-source fit onto single-source data, while joint ML keeps recovering the CCSN accurately. The BNS companion transitions at smaller distances than the BBH one, reflecting its lower network SNR. Across the sweep the joint ML estimator consistently outperforms strict assignment, confirming the value of two-source pixels.

The two-source mode always returns two positions and never decides the source count. Running a single source through it shows that the spurious second position can still be flagged. The higher-likelihood position concentrates at the true sky location in every trial while the lower-likelihood position is essentially random, so the per-position likelihood reliably identifies the spurious second source. The real source is still recovered with a median error of only 2.5° , close to the 1.9° single-source baseline.

TABLE I
SINGLE-SOURCE LOCALISATION

	BBH			BNS		CCSN
	10 Gpc	20 Gpc	30 Gpc	1 Gpc	2 Gpc	10 kpc
Median err. ($^\circ$)	1.9	5.8	15.9	1.3	5.5	1.3
90th pct. ($^\circ$)	6.1	21.3	115	5.5	16.8	12.4
50% CR area	12	64	402	12	68	12
90% CR area	35	375	5359	35	631	1463

ET+CE+A1 with PSD-coloured Gaussian noise, $N = 288$ injections per column area-weighted for a uniform sky draw, areas in deg^2 . Columns are a $20+20 M_\odot$ BBH at 10, 20 and 30 Gpc, a $1.4+1.4 M_\odot$ BNS at 1 and 2 Gpc, and a Galactic CCSN at 10 kpc. All metrics degrade with distance, and BBH and BNS behave near-identically. The CCSN keeps a small angular error but an exceptionally large 90% credible region, because its credible area spreads over disjoint sky clusters rather than one contiguous patch.

TABLE II
MULTI-SOURCE LOCALISATION

	BNS+BBH		BBH+CCSN		BNS+CCSN	
	BNS	BBH	BBH	CCSN	BNS	CCSN
N	150	150	50	50	50	50
Median err. ($^\circ$)	2.2	4.2	7.7	7.6	8.2	7.9
90th pct. ($^\circ$)	7.2	22.8	73	153	84	82
50% CR area	17	23	41	221	88	1086
90% CR area	82	280	20936	32933	16845	38041

Three pairwise combinations on ET+CE+A1, areas in deg^2 , with the number of injections N per combination as listed. The BBH and BNS sit at their single-source reference distances of 10 and 1 Gpc and the CCSN sky position is drawn from a Galactic thin-disk distribution [22], with sources timed to maximize TF overlap, the worst case for the joint solver. Relative to the single-source baseline (Table I), CBC+CBC (BNS+BBH) is tail-heavy, the median roughly doubling while the 90% credible region grows far more, whereas CBC+CCSN degrades the full distribution.

VII. CONCLUSION

We have shown that signal overlap is highly likely for the third-generation observing scenario, restructured a modular burst pipeline to carry multiple source hypotheses end to end, and developed a joint ML localisation statistic with a null-stream source-count criterion that handles one or two simultaneous sources on a small detector network. The method localises single sources to a few degrees and recovers all pairwise overlap combinations while outperforming a strict single-source pixel assignment.

Localisation quality depends strongly on the network’s sky-dependent sensitivity. BBH and BNS behave near-identically, and the high-error sky directions coincide with the network’s low-sensitivity regions (Fig. 4). The pipeline is a proof of concept rather than a tuned system. Its clustering stage is not yet tuned for diffuse bursts, which limits CCSN reconstruction, its detection significance under realistic non-stationary backgrounds remains to be calibrated, and parameters such as the STFT window length, the likelihood-statistic variant, and the sky-grid resolution were left at their initial values for compute-budget reasons. The modular design is built so that each of these refinements can be taken on module by module as research progresses.

REFERENCES

- [1] M. Maggiore *et al.*, “Science case for the Einstein Telescope,” *JCAP*, vol. 2020, no. 03, p. 050, 2020.
- [2] M. Branchesi *et al.*, “Science with the Einstein Telescope: a comparison of different designs,” *JCAP*, vol. 2023, no. 07, p. 068, 2023.
- [3] N. Singh *et al.*, “Compact-binary coalescence rates for the third-generation detector network,” 2022, arXiv:2112.04058.
- [4] K. Rozwadowska, F. Vissani, and E. Cappellaro, “On the rate of core collapse supernovae in the Milky Way,” *New Astron.*, vol. 83, p. 101498, 2021.
- [5] S. Klimenko *et al.*, “Coherent method for detection of gravitational wave bursts,” *Class. Quantum Grav.*, vol. 25, p. 114029, 2008.
- [6] A. Mannaert, “A modular data analysis pipeline for unmodelled gravitational-wave bursts,” M.S. thesis, KU Leuven, 2025.
- [7] F. Badaracco *et al.*, “Overlapping signals in next-generation gravitational-wave detectors,” *New Astron. Rev.*, vol. 99, p. 101707, 2024.
- [8] M. Evans *et al.*, “A horizon study for Cosmic Explorer: science, observatories, and community,” 2021, arXiv:2109.09882.
- [9] I. C. F. Wong and F. Cireddu, “gwmock: Python package for mock gravitational-wave samples,” [Online]. Available: <https://pypi.org/project/gwmock/>
- [10] F. Cireddu, “GW mock data challenge v1: population files,” Zenodo, 2025, doi:10.5072/zenodo.413548.
- [11] P. Reilton *et al.*, “Addressing the challenges of detecting time-overlapping compact binary coalescences,” *Phys. Rev. D*, vol. 106, p. 104045, 2022.
- [12] P. J. Sutton *et al.*, “X-Pipeline: an analysis package for autonomous gravitational-wave burst searches,” *New J. Phys.*, vol. 12, p. 053034, 2010.
- [13] M. Drago, “Search for transient gravitational wave signals with unknown waveform in the LIGO-Virgo network using a fully coherent algorithm,” Ph.D. dissertation, Univ. Padova, 2010.
- [14] S. Hild *et al.*, “Sensitivity studies for third-generation gravitational wave observatories,” *Class. Quantum Grav.*, vol. 28, p. 094013, 2011.
- [15] S. Klimenko, S. Mohanty, M. Rakhmanov, and G. Mitselmakher, “Constraint likelihood analysis for a network of gravitational wave detectors,” *Phys. Rev. D*, vol. 72, p. 122002, 2005.
- [16] S. Klimenko *et al.*, “Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors,” *Phys. Rev. D*, vol. 93, p. 042004, 2016.
- [17] S. Chatterji, A. Lazzarini, L. Stein, P. J. Sutton, A. Stroeer, and M. Tinto, “Coherent network analysis technique for discriminating gravitational-wave bursts from instrumental noise,” *Phys. Rev. D*, vol. 74, p. 082005, 2006.
- [18] A. N. Tikhonov, “Solution of incorrectly formulated problems and the regularization method,” *Soviet Math. Dokl.*, vol. 4, pp. 1035–1038, 1963.
- [19] J. A. Högbom, “Aperture synthesis with a non-regular distribution of interferometer baselines,” *Astron. Astrophys. Suppl. Ser.*, vol. 15, pp. 417–426, 1974.
- [20] G. Celeux and G. Govaert, “A classification EM algorithm for clustering and two stochastic versions,” *Comput. Statist. Data Anal.*, vol. 14, no. 3, pp. 315–332, 1992.
- [21] T. Sidery *et al.*, “Reconstructing the sky location of gravitational-wave detected compact binary systems,” *Phys. Rev. D*, vol. 89, p. 084060, 2014.
- [22] S. M. Adams, C. S. Kochanek, J. F. Beacom, M. R. Vagins, and K. Z. Stanek, “Observing the next Galactic supernova,” *Astrophys. J.*, vol. 778, no. 2, p. 164, 2013.
- [23] A. Mezzacappa *et al.*, “Core collapse supernova gravitational wave emission for progenitors of 9.6, 15, and 25 $M_{\odot}$,” *Phys. Rev. D*, vol. 107, p. 043008, 2023.
- [24] J. Powell, B. Müller, D. R. Aguilera-Dena, and N. Langer, “Three dimensional magnetorotational core-collapse supernova explosions,” *MNRAS*, 2023, arXiv:2212.00200.
- [25] A. Mezzacappa and M. Zanolin, “Gravitational waves from neutrino-driven core-collapse supernovae: predictions, detection, and parameter estimation,” 2024, arXiv:2401.11635.