Scientific Machine Learning for Gravitational Wave Astronomy
Poster Session Abstracts
Poster Session A
June 2 – 3, 2025

Self-Evolving Monte Carlo Tree Search with LLM Heuristics for Gravitational-Wave Signal Detection

He Wang, University of Chinese Academy of Sciences

From applied mathematics and fundamental physics to gravitational wave astronomy and complex systems analysis, scientific discovery relies on decisions driven by prior experience and data assumptions - yet faces persistent challenges in gravitational-wave signal identification. While existing approaches like matched filtering (MF) and deep neural networks (DNNs) have achieved partial success, their limitations emerge through MF's excessive computational demands and dependence on theoretical assumptions requiring predefined waveform templates, and DNNs' black-box architectures that obscure decision logic while introducing hidden biases and overfitting tendencies. We propose a synergistic framework where Monte Carlo Tree Search (MCTS) enables stepwise decision-making through strategic exploration, self-evolving (SelfEvo) optimization dynamically adjusts search policies via continuous performance feedback, and large language model (LLM)-generated heuristics inject domain-aware constraints - creating a closed-loop system that combines structured reasoning with adaptive learning. By encoding LLM-derived physical constraints into MCTS for dynamic search space navigation and enabling SelfEvo-driven adaptive parameterization, our method achieves unmodeled signal reconstruction with faster and better sensitivity compared to MF baselines, while maintaining human-interpretable decision traces unlike DNN approaches. Crucially, we demonstrate the algorithm's capacity to identify microsecond-duration waveform anomalies in simulated GW150914-type events through its hierarchical search decomposition capability, suggesting its potential to transcend conventional template-based searches. Beyond gravitational-wave astronomy, this work establishes a paradigm for AI-generated discovery in scientific computing, where the tripartite synergy of strategic exploration guided by domain knowledge accelerates hypothesis testing while reducing computational resource burdens.

Scalable data-analysis framework for long-duration gravitational waves from compact binaries using Short Fourier Transforms

Rodrigo Tenorio, Università degli Studi di Milano Bicocca

We introduce a framework based on Short Fourier Transforms (SFTs), automatic differentiation, and GPU computing to analyze long-duration gravitational wave signals from compact binaries. Targeted systems include binary neutron stars observed by third-generation ground-based detectors and massive black-hole binaries observed by the LISA space mission. In short, ours is an extremely fast, scalable, and parallelizable implementation of the gravitational-wave inner product, a core operation of gravitational-wave matched filtering. By operating on disjoint data segments, SFTs allow for efficient handling of noise non-stationarities, data gaps, and detector-induced signal modulations. We present a pilot application to early warning problems in both ground- and space-based next-generation detectors. Overall, SFTs reduce the computing cost of evaluating an inner product by three to five orders of magnitude, depending on the specific application, with respect to a non-optimized approach. The inner product is the key building block of all gravitational-wave data treatments; by speeding up this low-level element so massively, SFTs provide an extremely promising solution for current and future gravitational-wave data-analysis problems.

Reconstructing parametric gravitational-wave population fits from non-parametric results without refitting the data
Cecilia Maria Fabbri, University of Nottingham

Combining multiple events into population analyses is a cornerstone of gravitational-wave astronomy. A critical component of such studies is the assumed population model, which can range from astrophysically motivated functional forms to non-parametric treatments that are flexible but difficult to interpret. In practice, the current approach is to fit the data multiple times with different population models to identify robust features. We propose an alternative strategy: assuming the data have already been fit with a flexible model, we present a practical recipe to reconstruct the population distribution of a different model. As our procedure postprocesses existing results, it avoids the need to access the underlying gravitational-wave data again and handle selection effects. Additionally, our reconstruction metric provides a goodness-of-fit measure to compare multiple models. We apply this method to the mass distribution of black-hole binaries detected by LIGO/Virgo/KAGRA. Our work paves the way for streamlined gravitational-wave population analyses by first fitting the data with advanced non-parametric methods and careful handling of selection effects, while the astrophysical interpretation is then made accessible using our reconstruction procedure on targeted models. The key principle is that of conceptually separating data description from data interpretation.

Best of both worlds: integrating principled matched-filtering searches with AI/ML tools
Digvijay Wadekar, Johns Hopkins University

In the infinite data and compute limit, machine learning (ML) methods can be optimal, however this idealistic situation is not often realized in practice. On the other hand, principled data-analysis methods are robust, but they make simplistic assumptions (e.g., the noise is roughly Gaussian). I will show how ML algorithms can enhance matched-filtering pipelines by: (i) generating optimal template banks (ii) weighting templates to downplay unphysical binary configurations (iii) mitigating non-Gaussian noise. Incorporating these advancements in the IAS search pipeline, I will show new detections of black holes in the astrophysically significant pair-instability mass gap and IMBH mass ranges.

Fast and accurate parameter estimation of high-redshift sources with the Einstein Telescope
Filippo Santoliquido, Gran Sasso Science Institute

The Einstein Telescope (ET), along with other third-generation gravitational wave (GW) detectors, will be a key instrument for detecting GWs in the coming decades. However, analyzing the data and estimating source parameters will be challenging, especially given the large number of expected detections - of order 10^5 per year - which makes current methods based on stochastic sampling impractical. In this work, we use Dingo-IS to perform Neural Posterior Estimation (NPE) of high-redshift events detectable with ET in its triangular configuration. NPE is a likelihood-free inference technique that leverages normalizing flows to approximate posterior distributions. After training, inference is fast, requiring only a few minutes per source, and accurate, as corrected through importance sampling and validated against standard Bayesian inference methods. To confirm previous findings on the ability to estimate parameters for high-redshift sources with ET, we compare NPE results with predictions from the Fisher information matrix (FIM) approximation. We find that FIM underestimates sky localization errors substantially for most sources, as it does not capture the multimodalities in sky localization introduced by the geometry of the triangular detector. FIM also overestimates the uncertainty in luminosity distance by a factor of ∼ 3 on average when the injected luminosity distance is d_L > 10^5 Mpc, further confirming that ET will be particularly well suited for studying the early Universe.

Ringdown mode amplitudes of precessing binary black holes
Francesco Nobili, University of Insubria

The ringdown phase of a binary black-hole merger encodes key information about the remnant properties and provides a direct probe of the strong-field regime of General Relativity. While quasi-normal mode frequencies and damping times are well understood within black-hole perturbation theory, their excitation amplitudes remain challenging to model, as they depend on the merger phase. The complexity increases for precessing black-hole binaries, where multiple emission modes can contribute comparably to the ringdown. In this paper, we investigate the phenomenology of precessing binary black hole ringdowns using the SXS numerical relativity simulations catalog. Precession significantly impacts the ringdown excitation amplitudes and the related mode hierarchy. Using Gaussian process regression, we construct the first fits for the ringdown amplitudes of the most relevant modes in precessing systems.
Preprint available on arxiv: 2504.17021

Measuring correlations in realistic astrophysical populations of gravitational-wave sources with PixelPop
Sofia Alvarez-Lopez, Massachusetts Institute of Technology

The origins of merging compact binaries observed by the LIGO-Virgo-KAGRA gravitational-wave (GW) detectors remain uncertain, with multiple astrophysical processes possibly contributing to the overall merger rate. Such formation channels result in nontrivial correlations in the underlying distribution of source parameters. Current understanding of these sources relies heavily on simplified parametric models that make strong assumptions about the population. PixelPop – a high-resolution Bayesian nonparametric population algorithm – offers a more flexible approach by inferring joint distributions and parameter correlations with minimal assumptions. In this work, we use PixelPop to analyze two-, three-, and four-dimensional correlations in the masses, spins, and redshift for simulated black-hole mergers from population synthesis simulations observable by LIGO-Virgo with current sensitivity. We show that neglecting correlations gives smaller uncertainties but limits and biases astrophysical insights. In contrast, modeling all correlations with PixelPop allows us to reduce biases and correctly measure the astrophysical merger rate in all source parameters without modeling assumptions, at the expense of larger uncertainties. Using the inference results from PixelPop, we also propose a method to distinguish between different formation channels. At current sensitivity, we find that formation channels can only be distinguished if the underlying astrophysical processes differ significantly. However, channels that share several evolutionary stages (like common envelope and stable mass transfer) remain indistinguishable. Given our uncertainty in the formation of GW sources, robust extraction of astrophysics from GW catalogs requires that we allow for potential nontrivial high-dimensional correlations with algorithms like PixelPop

Identifying spacetimes using neural networks
Estuti Shukla, Penn State University

In general relativity, comparing two metric solutions expressed in different coordinate gauges is a non-trivial problem. This poster presents a novel deep learning approach to address this issue in the context of numerical relativity. We demonstrate the application of neural networks to learn coordinate transformations between two metric representations defined on the same manifold and a potential application of how this approach can be used to cross-validate results from different numerical relativity codes.

Surrogate models for long Eccentric Binary Black hole waveforms
Akash Maurya, International Centre for Theoretical Sciences, Bengaluru, India

Orbital eccentricity leaves a strong imprint in the gravitational waves (GWs) emitted by coalescing compact binaries. Its presence indicates at dynamically assembled binaries in dense stellar environments like globular clusters, nuclear star clusters, etc., thus making it an excellent tracer of their formation channels. As a result, the GW community has lately begun to focus on including eccentricity in GW data analysis. One such effort in this direction is our recently developed inspiral-merger-ringdown (IMR) waveform model with higher-order GW modes, ESIGMAHM, for black holes on moderately eccentric orbits with non-precessing spins.
While ESIGMAHM matches other eccentric IMR models in evaluation speed, it still requires approximately a week to complete Bayesian parameter estimation study for a vanilla GW150914-like event on considerable computing resources. This problem will be exacerbated even more in future as our detectors become more sensitive at lower frequencies, necessitating even longer waveforms to be computed for the analyses. To address this, we present a novel parameterization to build a scalable, long-duration, yet efficient surrogate model of ESIGMAHM capable of analyzing BBH events from frequencies as low as 10Hz, thus making ESIGMAHM a valuable eccentric IMR waveform model for the current and future GW data analysis needs.

Unraveling the eccentric binary black hole mergers with ESIGMAHM
Kaushik Paul, Indian Institute of Technology Madras

The sensitivities of the LIGO detectors are continuously upgraded, and as a consequence, we hope to see a gradual increase in the number of observed events. A small fraction of these are expected to be detected with non-negligible orbital eccentricities. Further, the BHs in binaries are also expected to have significant spins. Neglecting these effects in templates may not only reduce the detection efficiency, it may also lead to a biased estimation of binary parameters. Here, we present a spinning, eccentric waveform model which, apart from the dominant quadrupole modes, also includes a number of higher-order modes. Built on an earlier framework, ENIGMA, which could be used to generate waveforms suitable for analysing BBHs with non-spinning components, we upgrade the model by adding spins and higher modes and name it ESIGMAHM. Our model is capable of describing positively aligned-spin BBH systems in moderately eccentric orbits with an overlap of > 97% with NR simulations from the SXS catalog. We show that current LIGO searches will lose more than 10% of optimal SNR for about 20% of all eccentric sources by using only quasi-circular waveform templates. Our model is presented in Phys. Rev. D 111, 084074.

Interpretable Machine Learning for Dynamical Dark Energy Models
Giulia Borghetto, University of Swansea

Recent cosmological observations suggest possible deviations from a cosmological constant, pointing toward a dynamical nature of dark energy. Quintessence models, which assume a slowly rolling scalar field, provide a compelling theoretical framework to explain this late-time evolution in the dark energy equation of state. However, identifying the correct form of the quintessence potential remains a major challenge, due to both theoretical constraints and the vast landscape of functional possibilities. In this work, we explore the use of symbolic regression—an interpretable machine learning technique—to discover viable quintessence potentials directly from observational data. By searching over analytical expressions rather than fitting predefined forms, symbolic regression offers a data-driven approach to model selection that retains physical

interpretability. Furthermore, the framework developed here may also be applied to scenarios where scalar fields influence the propagation of gravitational waves, potentially offering new insights into forthcoming observational data.

On the Generalization Ability of Deep Learning Networks in Gravitational Wave Detection
Paola C. M. Delgado, CEICO - FZU - Czech Academy of Sciences

We explore the generalization capacity of deep neural networks trained exclusively on General Relativity (GR) waveforms to detect beyond GR (bGR) signals. Using simulated data generated with parameterized post-Einsteinian (ppE) modifications in both the amplitude and phase of the waveforms, we assess the performance of a convolutional neural network in distinguishing injected signals from noise. Our main goal is to understand whether and what kind of bGR physics can be learned by the NN from GR training.

Striking a Chord with Spectral Sirens: multiple features in the compact binary population correlate with H0
Utkarsh Mali, CITA

Spectral siren measurements of the Hubble constant (H0) rely on correlations between observed detector-frame masses and luminosity distances. Features in the source-frame mass distribution can induce these correlations. It is crucial, then, to understand (i) which features in the source-frame mass distribution are robust against model (re)parametrization, (ii) which features carry the most information about H0, and (iii) whether distinct features independently correlate with cosmological parameters. We study these questions using real gravitational-wave observations from the LIGO-Virgo-KAGRA Collaborations' third observing run. Although constraints on H0 are weak, we find that current data reveals several prominent features in the mass distribution, including peaks in the binary black hole source-frame mass distribution near $\sim 9\ M_\odot$ and $\sim 32 M_\odot$ and a roll-off at masses above $\sim 46 M_\odot$. For the first time using real data, we show that all of these features carry cosmological information and that the peak near $\sim 32 M_\odot$ consistently correlates with H0 most strongly. Introducing model-independent summary statistics, we show that these statistics independently correlate with H0, exactly what is required to limit systematics within future spectral siren measurements from the (expected) astrophysical evolution of the mass distribution