Scientific Machine Learning for Gravitational Wave Astronomy
Poster Session Abstracts

Poster Session B
June 3 – 4, 2025

Super-Resolution for dark matter simulations

John Brennan, Maynooth University

Numerical simulations are a cornerstone in theoretical astrophysics and cosmology, enabling researchers to model complex processes that are impractical for analytical or observational study. However, achieving high spatial resolution in numerical simulations requires significant computational resources. As cosmological and astrophysical processes can span extreme dynamical scales, resolution requirements often render direct modeling impossible, forcing researchers to make a trade-off between accuracy and computational cost. To mitigate this, researchers use subgrid models to approximate unresolved processes, such as star and black hole formation, which would otherwise be excluded from their simulations due to resolution limitations.
In recent decades, advancements in artificial intelligence, particularly neural networks, have found remarkable applications in many areas of science. In the field of computational astrophysics, neural networks offer new approaches to model subgrid physics with substantial boosts in performance. For instance, super-resolution models trained to artificially enhance the resolution of data, such as images, can be used to similarly enhance the spatial resolution of numerical simulations, providing an alternative approach to generating high resolution datasets. In this poster, we present results from recent work on applying super-resolution techniques to model physical quantities that are not resolved in simulations.

Towards a Foundational AI Model for Gravitational Wave Astronomy

Chayan Chatterjee, Vanderbilt University

As gravitational wave detectors become more advanced and sensitive, the number of signals recorded by Advanced LIGO and Virgo from merging compact objects is expected to rise dramatically. This surge in detection rates necessitates the development of adaptable, scalable, and efficient tools capable of addressing a wide range of tasks in gravitational wave astronomy. Foundational AI models present a transformative opportunity in this context by providing a unified framework that can be fine-tuned for diverse applications while leveraging the power of large-scale pre-training. In this work, we explore how advanced transformer models, specifically OpenAI's Whisper, can be adapted as a foundational model for gravitational wave data analysis. By fine-tuning Whisper's encoder model—originally trained on extensive audio data—and combining it with neural networks for specialized tasks, we achieve reliable results in detecting astrophysical signals and classifying transient noise artifacts or 'glitches'. This represents the first application of open-source transformer models, pre-trained on unrelated tasks, for gravitational wave research, demonstrating their potential to enable versatile and efficient data analysis in the era of rapidly increasing detection rates.

Exact Outgoing Boundary conditions for Teukolsky equation in Boyer-Lindquist coordinates

Som Bishoyi, University of Massachusetts Dartmouth

Next generation gravitational wave(GW) detectors such as LISA would hunt for GW signals from extreme mass ratio inspirals(EMRIS). These systems comprise of a smaller black hole orbiting a much larger black hole; the mass ratio being 10^{-3} to 10^{-6}. Modeling such systems requires accurately solving the Teukolsky equation. Numerically solving the Teukolsky equation in Boyer-Lindquist coordinates is challenging due to the long range potential caused due to the peeling properties of the Weyl scalars. In this work, we develop exact outgoing boundary conditions for the Teukolsky equation using Boundary Kernels. Using this technique, one can arbitrarily choose any finite value of radial coordinate and impose purely outgoing solutions at that boundary. This prevents any spurious reflections and unphysical growth of the solution at late times and the need to causally disconnect the boundary by placing it at large values of the radial coordinate.

Hyperspectral Time-Series Classification via Gramian Angular Difference Fields and Deep CNNs

Sanaz Hami Hassan Kiyadeh, The University of Alabama

We present a unified framework for classifying hyperspectral time-series data by transforming each 1D spectral signature into a 2D Gramian Angular Difference Field (GADF). This encoding preserves temporal dependencies while enabling convolutional neural networks to extract spatial patterns from spectral signals. We fine-tune a ResNet101 model on GADF images, achieving up to 99% accuracy with under 1% misclassification across several hyperspectral benchmarks. Comparative evaluations with classical models such as KNN and linear SVM show superior accuracy and faster convergence in the transformed space. To address data scarcity, we integrate unsupervised feature learning through nonnegative matrix factorization and autoencoders, which enhance performance under limited label availability. Our results demonstrate that angular-difference encoding yields nearly linearly separable representations, making even simple classifiers effective, while deep networks capture higher-order dependencies for robust segmentation. This approach offers a lightweight, interpretable pipeline for hyperspectral classification and opens the door to hybrid CNN–Transformer architectures and advanced oversampling strategies for imbalanced classes.

New extension for black-hole remnant surrogate models at extreme mass ratios

Matteo Boschini, University of Milano-Bicocca

I present a new extension of black-hole remnant surrogate models for quasi-circular precessing binaries. This development advances previous efforts to predict post-merger properties across a wide region of the binary black hole parameter space. The new model is based on a new set of numerical relativity simulations up to mass ratio $q<8$ and combines them with analytic training data generated at extreme mass ratios. This solution allows to greatly enhance the regime of validity of the surrogate, allowing predictions on the entire mass ratio domain. Additionally, the model addresses the problem of extending the fit for the remnant recoil, associated with the anysotropic emission of gravitational waves. This work represents a further step toward robust, efficient modeling of binary black hole mergers across the full astrophysical parameter space.

Training techniques for learning orbital dynamics of binary black hole systems

Pranav Vinod, University of Massachusetts Dartmouth

One important astrophysical application of general relativity is modeling binary black hole systems. While numerical relativity solves Einstein’s equations for these systems, it remains resource-intensive, motivating faster data-driven alternatives. Previous work has shown that neural ordinary differential equations (NODEs) can learn the underlying dynamics directly from gravitational wave data by solving a constrained optimization over plausible physical models. This approach, however, requires solving potentially expensive ODEs multiple times throughout the training procedure, thereby limiting the network size and accuracy. Our modified approach employs a feed-forward neural network (NN) trained in two stages. First, the NN is trained directly to approximate the ODEs right-hand side without considering any physics of the problem. This pre-trained NN can be refined by solving a physics-informed constrained optimization using waveform data. Preliminary results indicate successful training across various orbits, with errors nearing numerical round-off errors. This enhanced accuracy enables tackling new problem types. For example, we show that the resulting NODE accurately extrapolates to long-time durations, can be used close to the problem's separatrix, and can be applied to more complex dynamics such as zoom-whirl orbits where multiple distinct timescales appear. These capabilities open the door to broader applications—for instance, accelerating eccentricity reduction in numerical relativity simulations. We present preliminary results illustrating these advantages.

Time Domain Neural Posterior Estimation for Black Hole Ringdowns

Ashwin Girish, University of Rhode Island

We present a time-domain implementation of Neural Posterior Estimation (NPE) for gravitational-wave ringdown signals from black hole mergers. Traditional inference methods in the frequency domain may lose information critical for short, noisy signals, while time-domain approaches preserve the full structure of the waveform. Our method leverages simulation-based inference techniques with amortized neural networks to perform likelihood-free Bayesian parameter estimation. Specifically, we target the characterization of quasi-normal modes in the post-merger ringdown phase. We explore both Gaussian-based and normalizing flow architectures trained on whitened, truncated segments of simulated ringdown data. We evaluate the accuracy and robustness of posterior recovery across a range of signal-to-noise ratios, addressing challenges such as overfitting and model calibration.

Data-Driven Approaches to Modeling Eccentric Binary Black-Hole Mergers and Building Template Banks

Tousif Islam, Kavli Institute for Theoretical Physics

The presence of orbital eccentricity in compact binaries introduces additional modulations in the amplitudes, phases, and frequencies of their gravitational-wave signals, complicating both detection (through template-bank construction) and source characterization (via fast yet accurate Bayesian inference). Leveraging data-driven techniques—such as singular value decomposition, Gaussian process regression, and k-means clustering—we (i) develop gwMiner, a framework that decomposes oscillatory eccentric spherical-harmonic modes into several monotonic eccentric modes; (ii) construct gwharmone, a reduced-order model for the four dominant eccentric modes of non-spinning binaries; and (iii) employ gwharmone to generate an efficient template bank that explicitly incorporates the effects of eccentricity.

Emulating compact binary population synthesis simulations with robust uncertainty quantification using Bayesian normalizing flows.

Anarya Ray, Northwestern University

Population synthesis simulations of compact binary mergers are crucial in extracting astrophysical insights from an ensemble of gravitational wave observations. However, realistic simulations can be costly to implement for a dense grid of initial conditions. Previous studies have shown normalizing flows can emulate the simulated distribution function of binary parameters and thereby enable empirical constraints on the astrophysical initial conditions and branching fractions of various formation channels given data from a catalog of GW observations. In this project I develop a method for quantifying model uncertainties in the emulators by introducing Bayesian Normalizing flows, a density estimator constructed from Bayesian neural networks. I demonstrate the accuracy of the estimated uncertainties for simulations of binary black hole populations formed through common envelope evolution and outline the applications of the proposed methodology in the context of simulation based inference now marginalized over model and data uncertainties.

BHPTNRSur2dq1e3: A surrogate model for spinning, intermediate-mass-ratio binaries based on perturbation theory calibrated to numerical relativity

Ritesh Bachhar, University Of Rhode Island

Surrogate models (or reduced-order models) offer the ability to faithfully and rapidly produce gravitational waves specified through computationally expensive, theory-based models. In recent years, there has been a surge of interest in building surrogate models in the waveform modeling community because of the usefulness of these models for multi-query data analysis studies. In this talk, I will provide a brief overview of the point particle black hole perturbation theory (ppBHPT) framework for generating gravitational waveforms and summarize the key algorithms for building a surrogate model using these waveforms. Specifically, I'll present a waveform model, BHPTNRSur2dq3 - a surrogate model of gravitational waves emitted from a binary black hole (BBH) for the comparable-to-intermediate mass ratio regime with aligned spin on the heavier mass. A key element of our work is the the calibration of ppBHPT waveforms to match the waveforms computed by numerical relativity (NR). We show the calibrated ppBHPT-based model is accurate even in the comparable mass regime despite being well outside the formal domain of validity for perturbation theory. Our model will help to bridge the gap between comparable to extreme mass ratio systems, where state-of-the-art models are currently lacking.

Rapid Identification and Classification of Eccentric Gravitational Wave Inspirals with Machine Learning

Adhrit Ravichandran, CSCDR, University of Massachusetts Dartmouth

Current templated searches for gravitational waves (GW) from compact binary coalescences (CBC) assume that the binaries have circularized by the time they enter the sensitivity band of the LIGO-Virgo-KAGRA (LVK) network. However, certain formation channels predict that in future observing runs, a fraction of detectable binaries could enter the sensitivity band with a measurable eccentricity. But constraining eccentricity for each GW event with Bayesian parameter estimation is computationally expensive. We therefore propose a proof-of-concept separable CNN that acts on qscans to segregate candidates as either eccentric or non-eccentric, which can then streamline downstream LIGO analysis.

labrador: an easy-to-train machine-learning tool for gravitational wave inference

Javier Roulet, Caltech

Fast and reliable inference of gravitational wave source parameters is crucial for identifying short-lived electromagnetic counterparts and avoiding systematic biases in large event catalogs. Neural posterior estimation has recently emerged as a powerful inference method, where a model is trained on simulations at substantial computational cost and then enables extremely fast and inexpensive inference at test time. Here, we extend this approach by incorporating domain-specific physical insights and methods in the model architecture. These include compressing the data by heterodyning against a reference waveform selected via likelihood maximization with analytical approximations; removing parameter degeneracies through tailored coordinate systems; and eliminating symmetry-induced multimodalities by folding the parameter space. As a result the network becomes approximately equivariant to changes in the source parameters, leading to improved interpretability and reduced training cost. Our implementation, called `labrador`, can be trained on a 1-day timescale on an A100 GPU.

Mapping Systematic Effects of Gravitational-Wave Parameter Estimation with DINGO

Samuel Clyne, URI

As the sensitivity of ground based gravitational wave detectors continues to improve, discrepancies between waveform models and the true signals are expected to lead to larger biases in posteriors obtained via Bayesian Inference. This study leverages the Deep Inference for Gravitational Wave Observation (DINGO[1]) code to rapidly estimate posterior distributions for synthetic gravitational wave events using normalizing flow neural networks. We analyze mock binary black hole mergers at design sensitivity for LIGO and Virgo, performing inference using models trained on phenomenological, effective-one-body, and NR surrogate waveform families. We analyze NR surrogate mock signals with each model while varying key intrinsic parameters in order to identify the systematic effects of waveform models on binary black hole parameter estimation at various points in parameter space. We assess the similarity of the mock posterior distributions and introduce a novel multivariate definition of bias based on the Mahalanobis distance. We find large disagreements between posteriors at high spin and inclination, where higher order modes have greater effects on the waveform model.

Building Next-Generation Gravitational Wave Searches

Kanchan Soni, Syracuse University

Gravitational-wave detectors such as LIGO and Virgo have transformed our understanding of compact binary mergers by enabling the detection of signals from colliding neutron stars and black holes. Next-generation observatories, such as Cosmic Explorer, will offer orders-of-magnitude improvements in sensitivity, leading to near-continuous detections and access to previously unobservable sources, including primordial black holes. Its enhanced sensitivity at low frequencies will also improve our ability to detect eccentric and precessing binary mergers, shedding light on their astrophysical formation channels. These advancements, however, pose significant computational challenges, as current search algorithms are not optimized for efficiently detecting signals lasting from minutes to hours in the detector's sensitivity band. To address this, we aim to develop a scalable, low-latency search pipeline that is robust to noise transients and capable of performing Bayesian inference in near real-time. By integrating machine learning-assisted hierarchical search strategies with rapid evidence evaluation, our goal is to build a fully coherent Bayesian pipeline tailored for the demands of next-generation gravitational-wave astronomy.