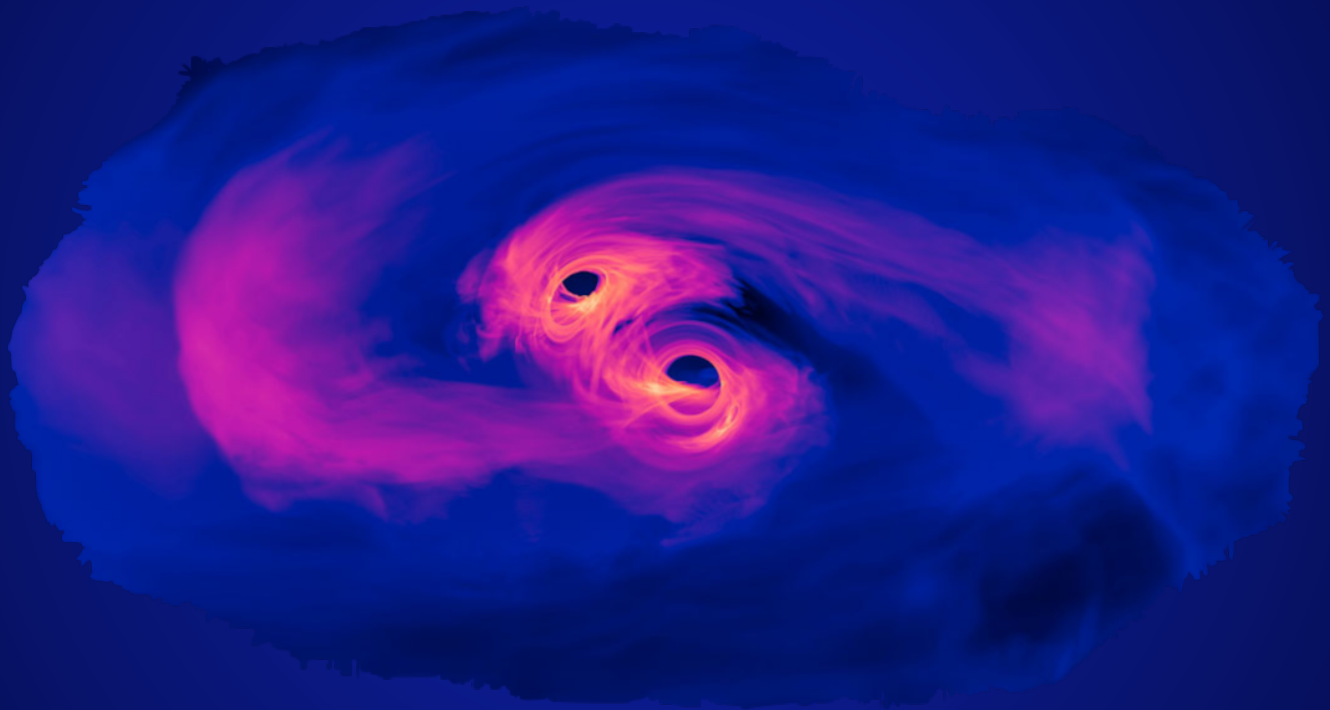

Real-time Gravitational Wave Data Analysis with Machine Learning

Katya Govorkova katyag@mit.edu, Ryan Raikman, Eric A Moreno,
Ethan J Marx, Alec Gunny, William Benoit, Deep Chatterjee, Rafia Omer,
Muhammed Saleem, Dylan S Rankin, Michael W Coughlin, Philip C Harris, Erik Katsavounidis

Gravitational Waves and Their Detection

Accelerating masses produce deformations in space time that we can detect via interferometry

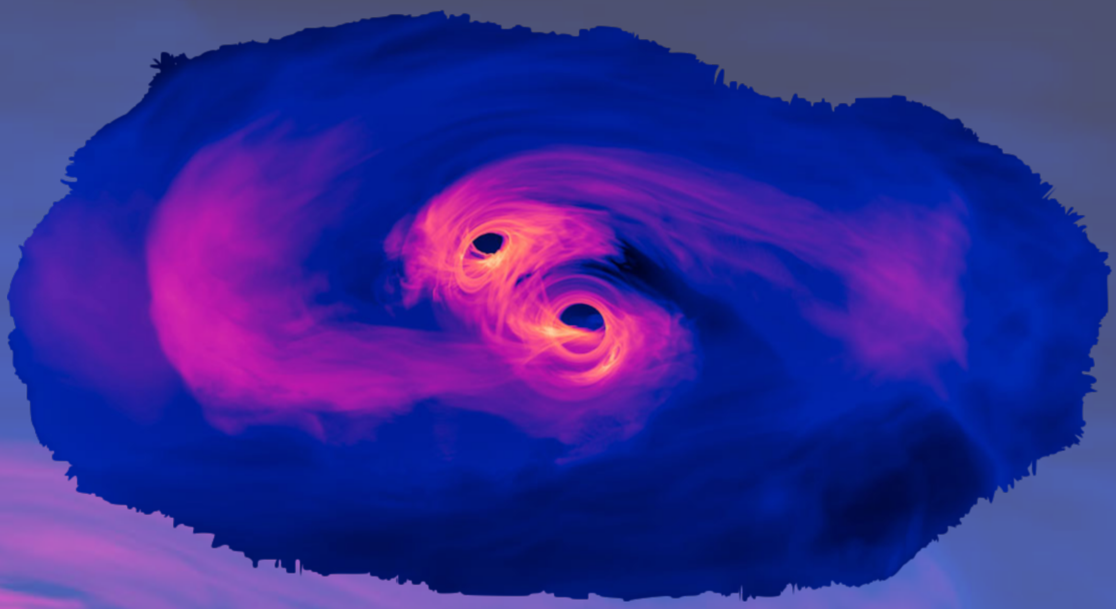

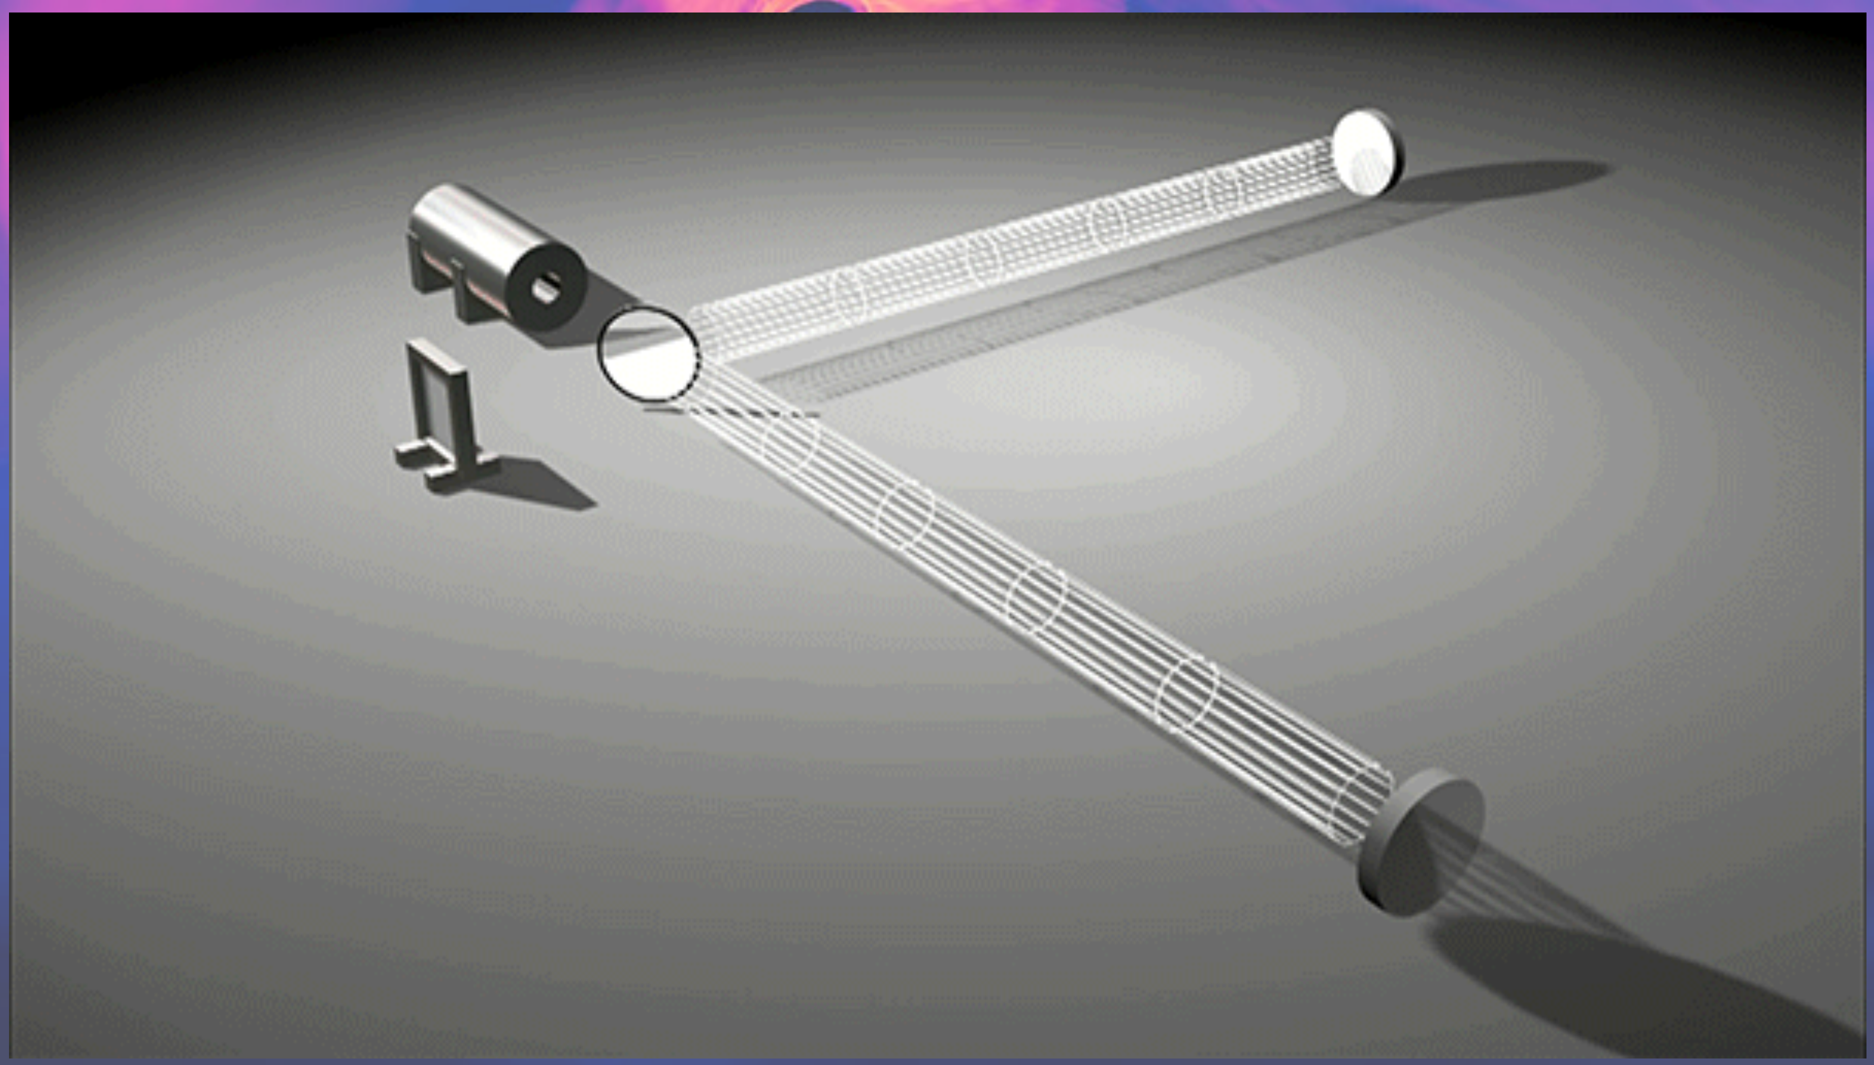

A SIGNAL WILL APPEAR IN AT LEAST TWO INTERFEROMETERS, WITH THE TIME DELAY BECAUSE OF THE DISTANCE BETWEEN THE DETECTORS

Typical GW Data Workflow

~100k aux
channels
16kHz

Detector characterisation

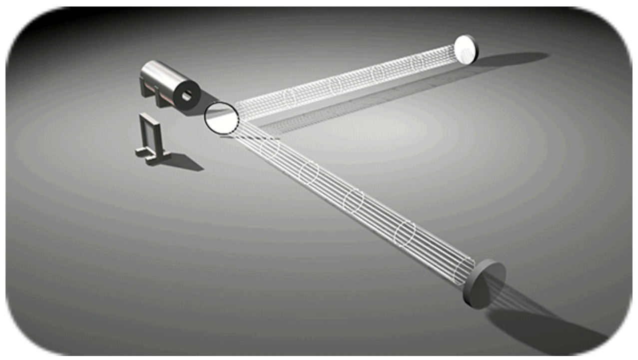

Use info from witness sensors to perform data de-noising

Cleaned Data

Event detection

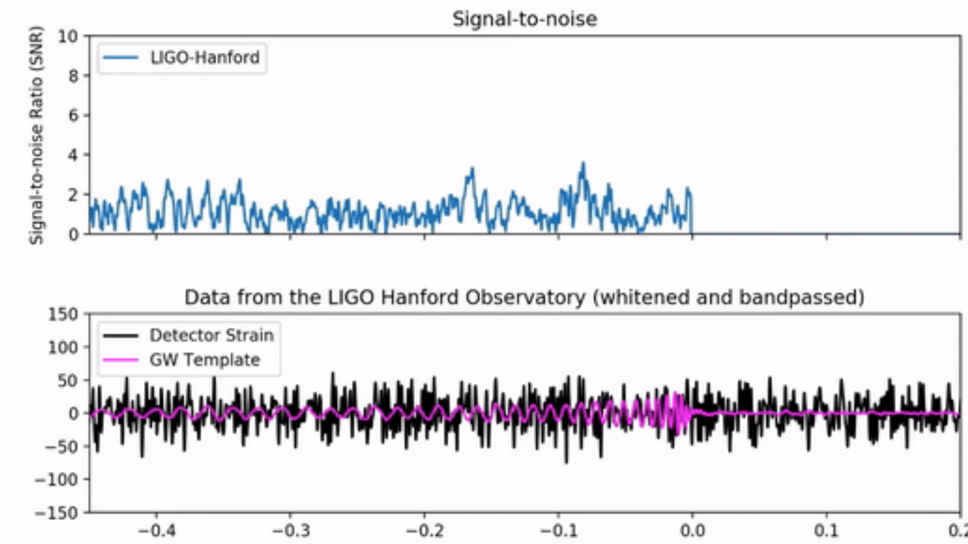

Event

Event characterisation

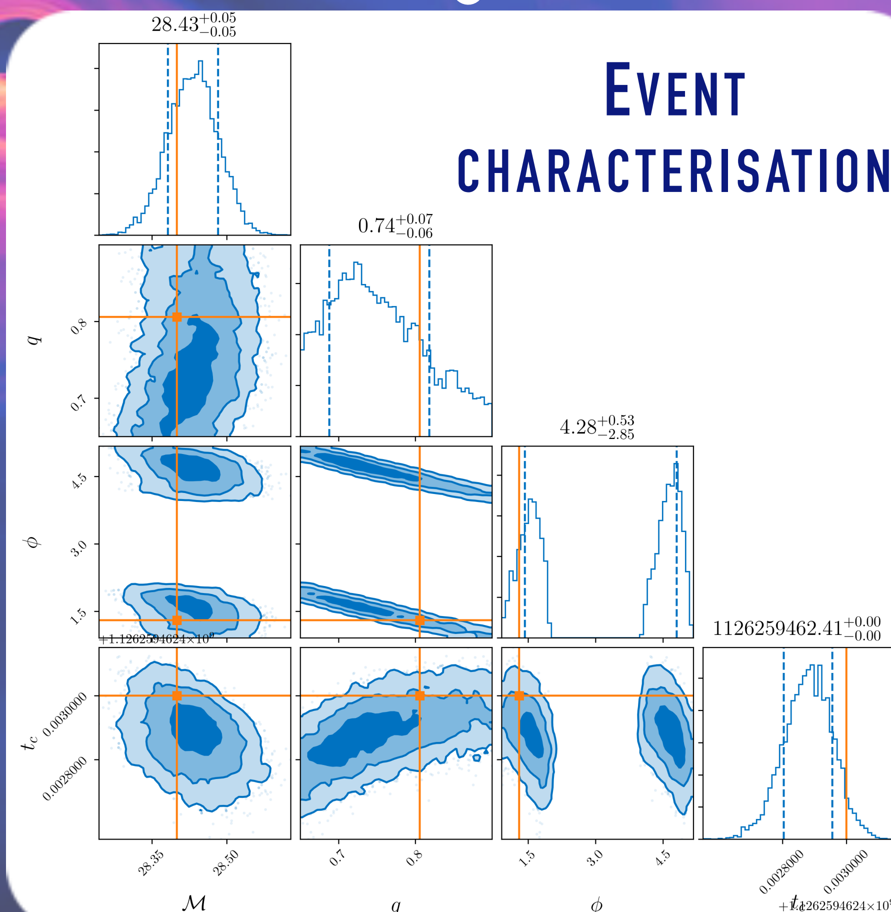

Alert 🚨

Current workflow uses CPU
data grid with rule based algorithms
Challenge is to run this in real-time

Why ML?

- Increasing detector sensitivity → more templates for matched filtering
- Makes ML advantageous in terms of computational cost and latency (and possibly sensitivity) – useful for Multi-messenger astrophysics efforts

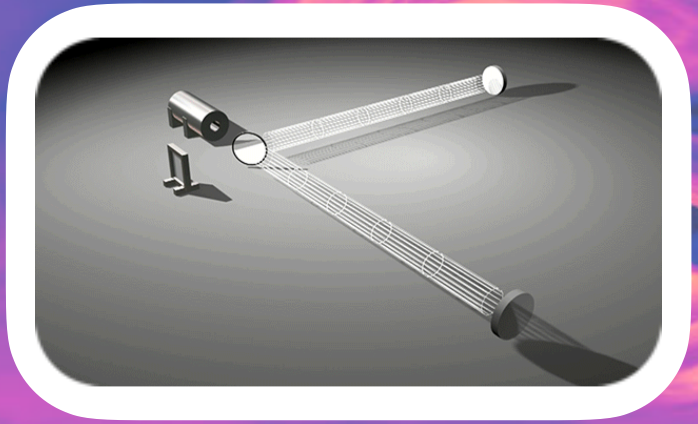

Alert

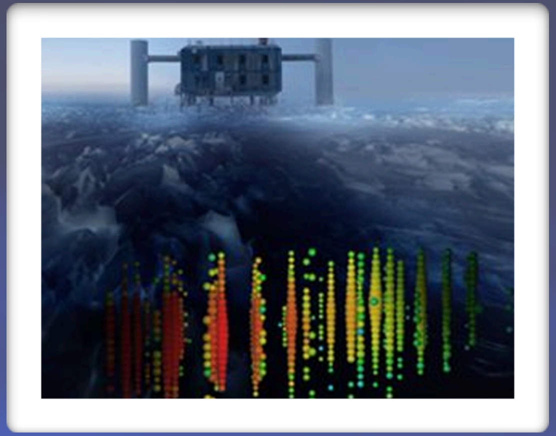

Neutrinos

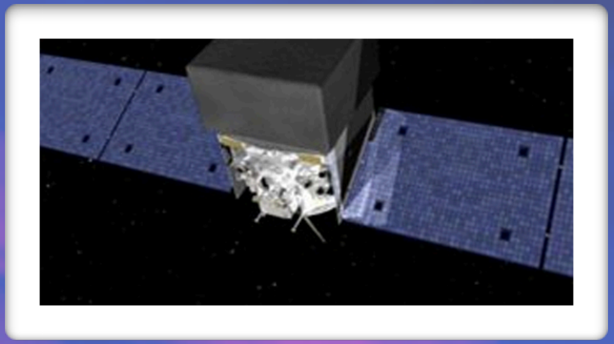

X-rays/gamma-rays

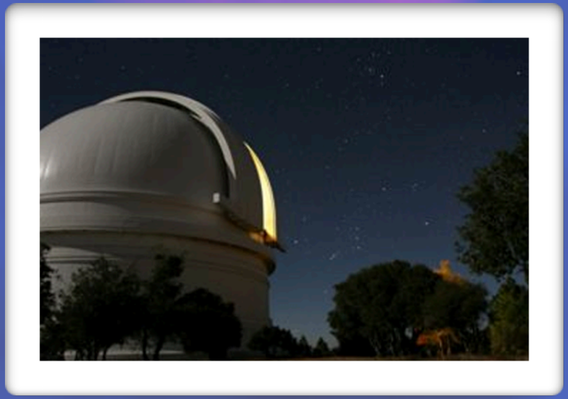

Visible/infrared light

Radio waves

Why ML?

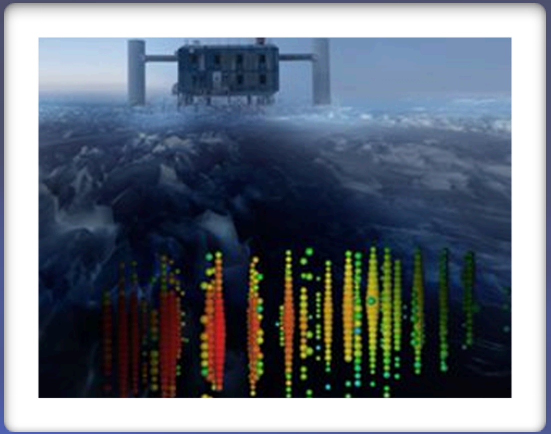

Neutrinos

- Increasing detector sensitivity → more templates for matched filtering
- Makes ML advantageous in terms of computational cost and latency (and possibly sensitivity) — useful for Multi-messenger astrophysics efforts

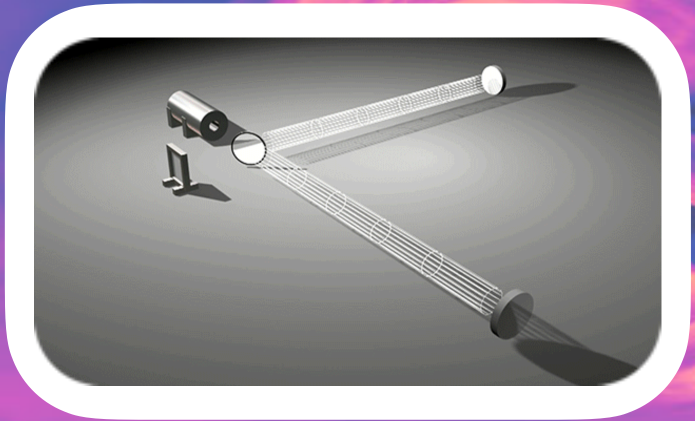

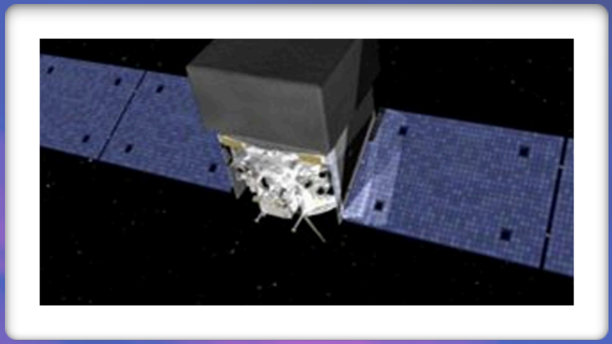

X-rays/gamma-rays

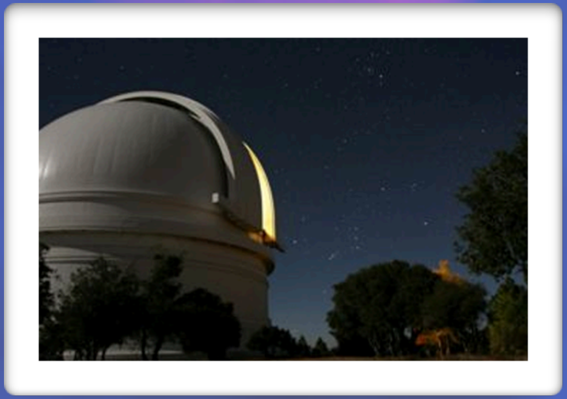

Visible/infrared light

Noise subtraction and downstream algorithms need to work in real-time to capture as much data as possible and satisfy

- High throughput
- Low latency
- Robust to changing data distribution

Radio waves

Future ML-based workflow

Data
16kHz

~100k auxiliary
channels

Detector
characterisation

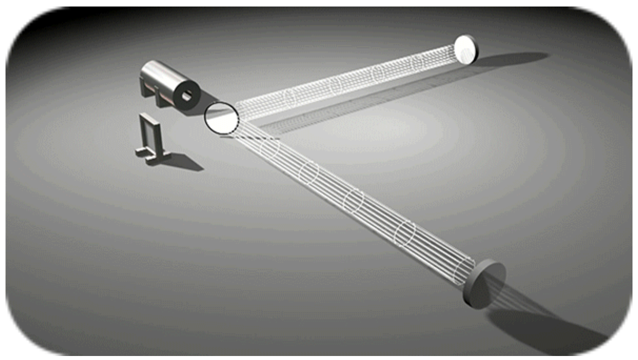

DeepClean
NN based AE
noise subtraction

Cleaned
Data

The output reconstructed from an interferometer contains

$$h(t) = s(t) + n(t)$$

Possible GW signal

Detector noise

The output reconstructed from an interferometer contains

$$h(t) = s(t) + n(t)$$

Possible GW signal

Detector noise

$$n(t) = n_{nw}(t) + n_{w}(t)$$

Non-removable (fundamental noise)
Eg: photon shot noise, thermal noise

Can be reduced only with upgraded design and technology

Source of noise witnessed by dedicated system monitors (witness sensors)

Environmental contamination or technical noise eg: noise arising from the control of suspended optics

DeepClean Denoising

- CNN-based autoencoder to predict the noise using witness channels
- Active-learning: network is fine-tuned at fixed intervals and the new model is hosted alongside stable model on inference service
- DeepClean is capable of denoising the data at ~ 1 s latency - a promising prospect for electromagnetic follow-up of gravitational wave observations

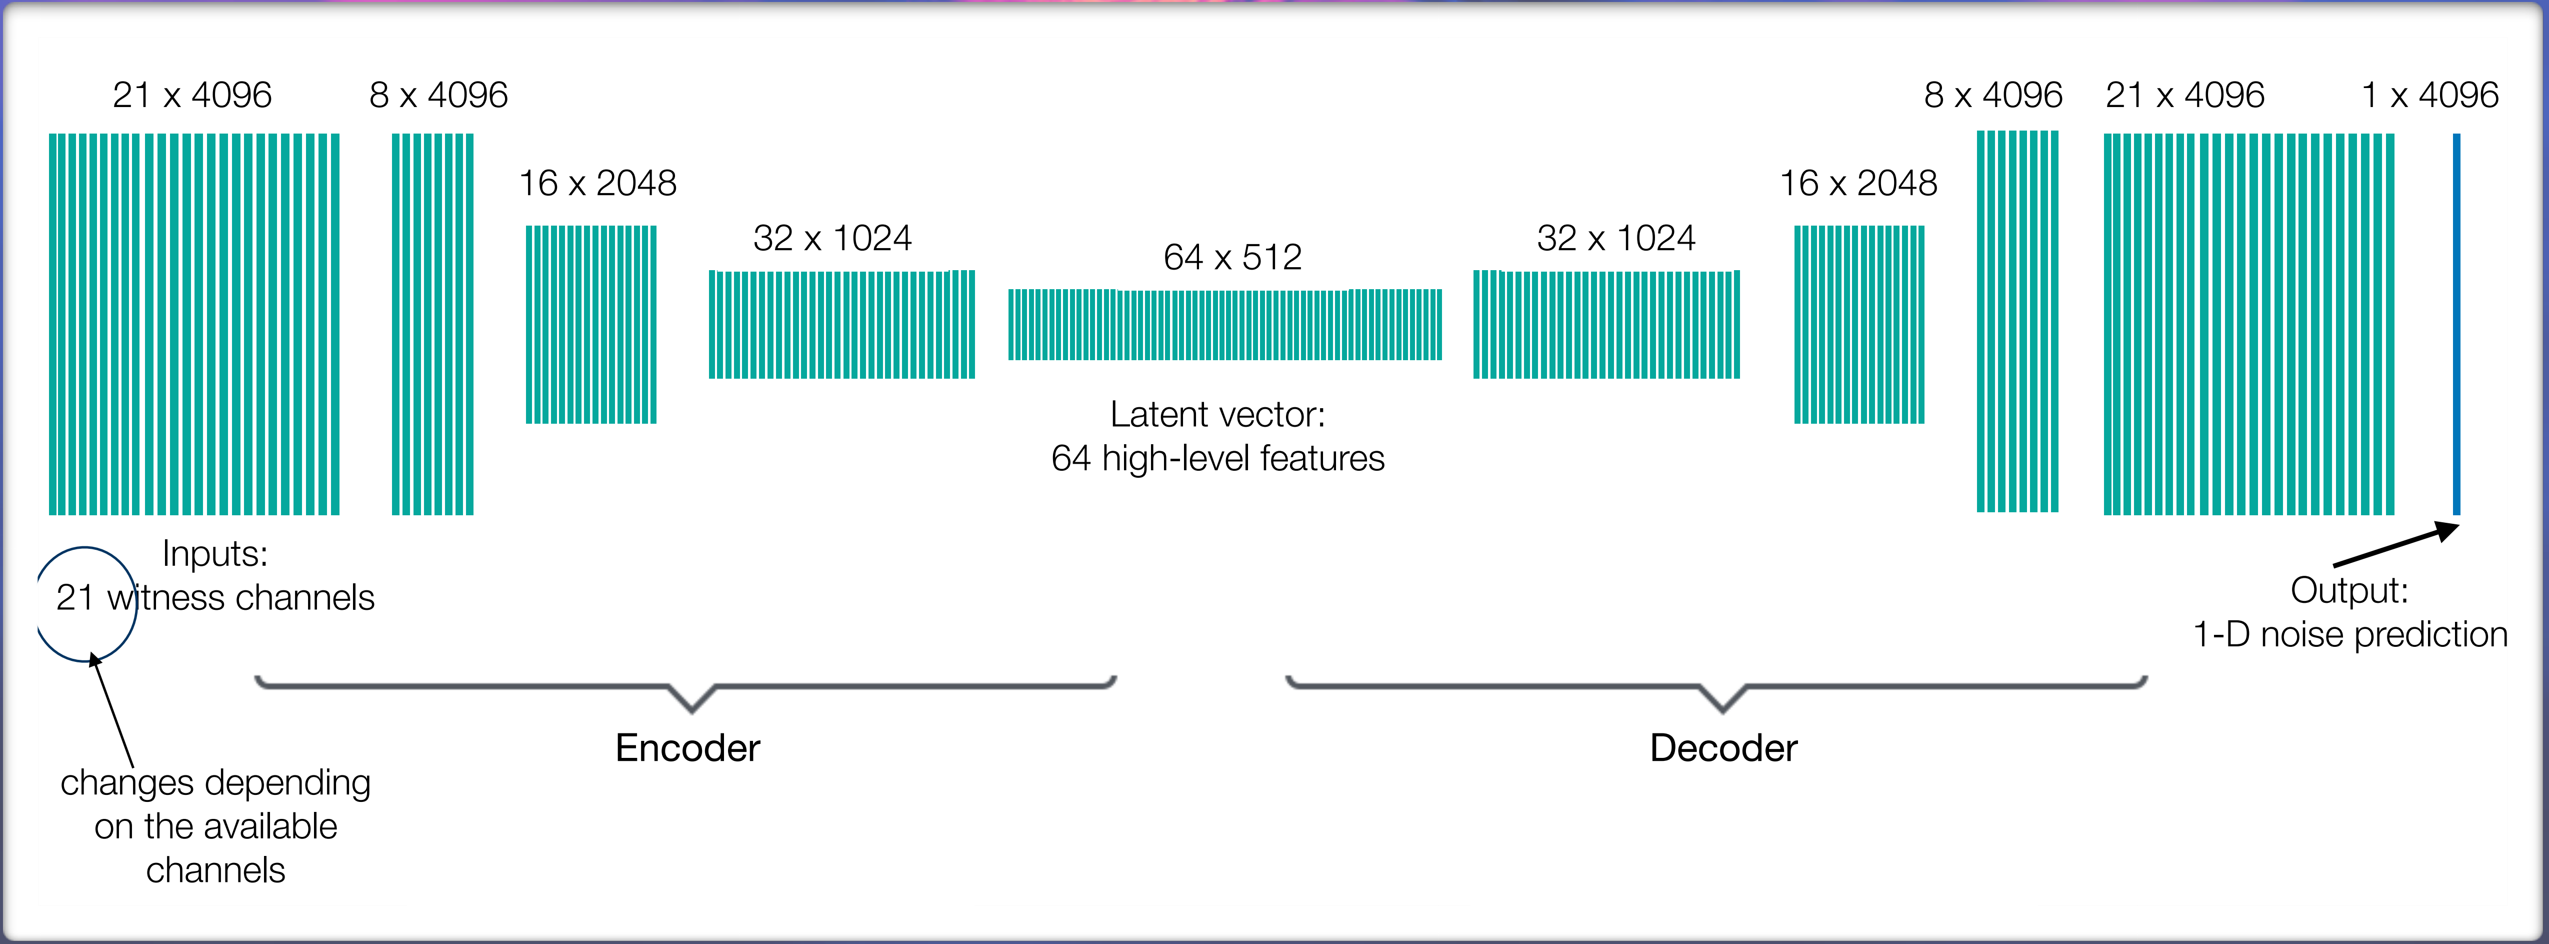

Future ML-based workflow

Data 16kHz

~100k auxiliary channels

Detector characterisation

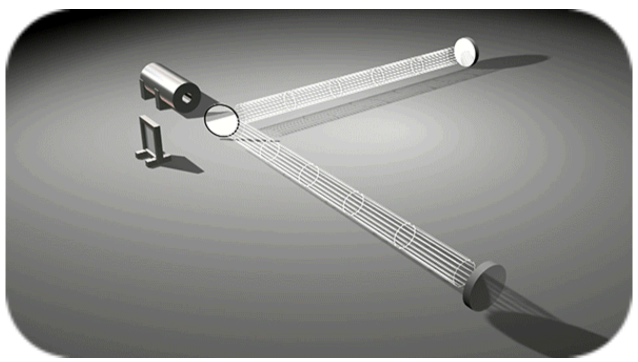

DeepClean
NN based AE
noise subtraction

Cleaned Data

NN-based algos for Event detection

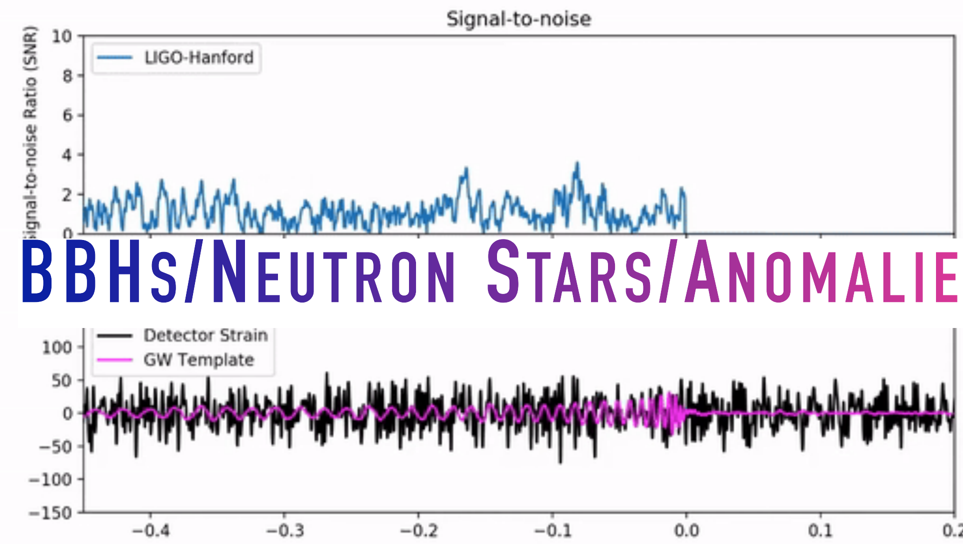

BBHs/Neutron Stars/Anomalies

A-Frame

arXiv2403.18661

- Detecting compact binary coalescences in gravitational wave strain timeseries data using neural networks
- ResNet architecture, maps from detector strain from two interferometers to a scalar neural-network output
- 2-10 times faster than matched filtering CBC pipeline

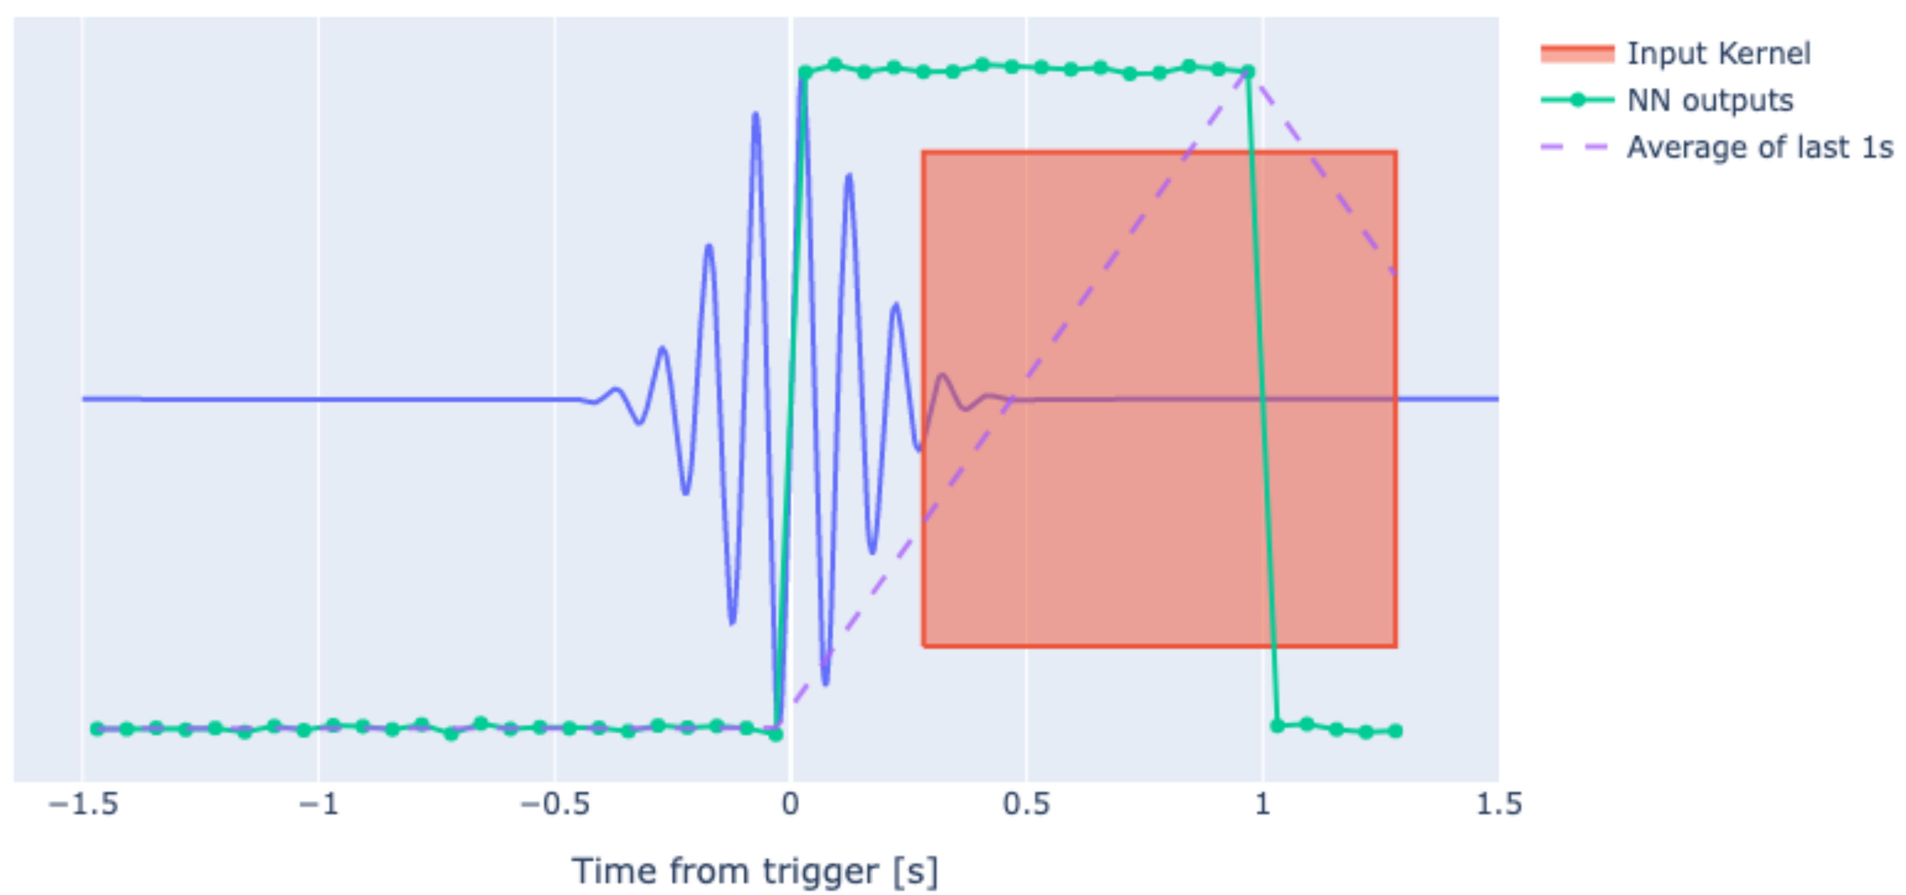

A-Frame performance comparison

arXiv2403.18661

Competitive performance on higher-mass catalog distributions

Work remains to be done for lower masses — alternative architectures or smarter training techniques

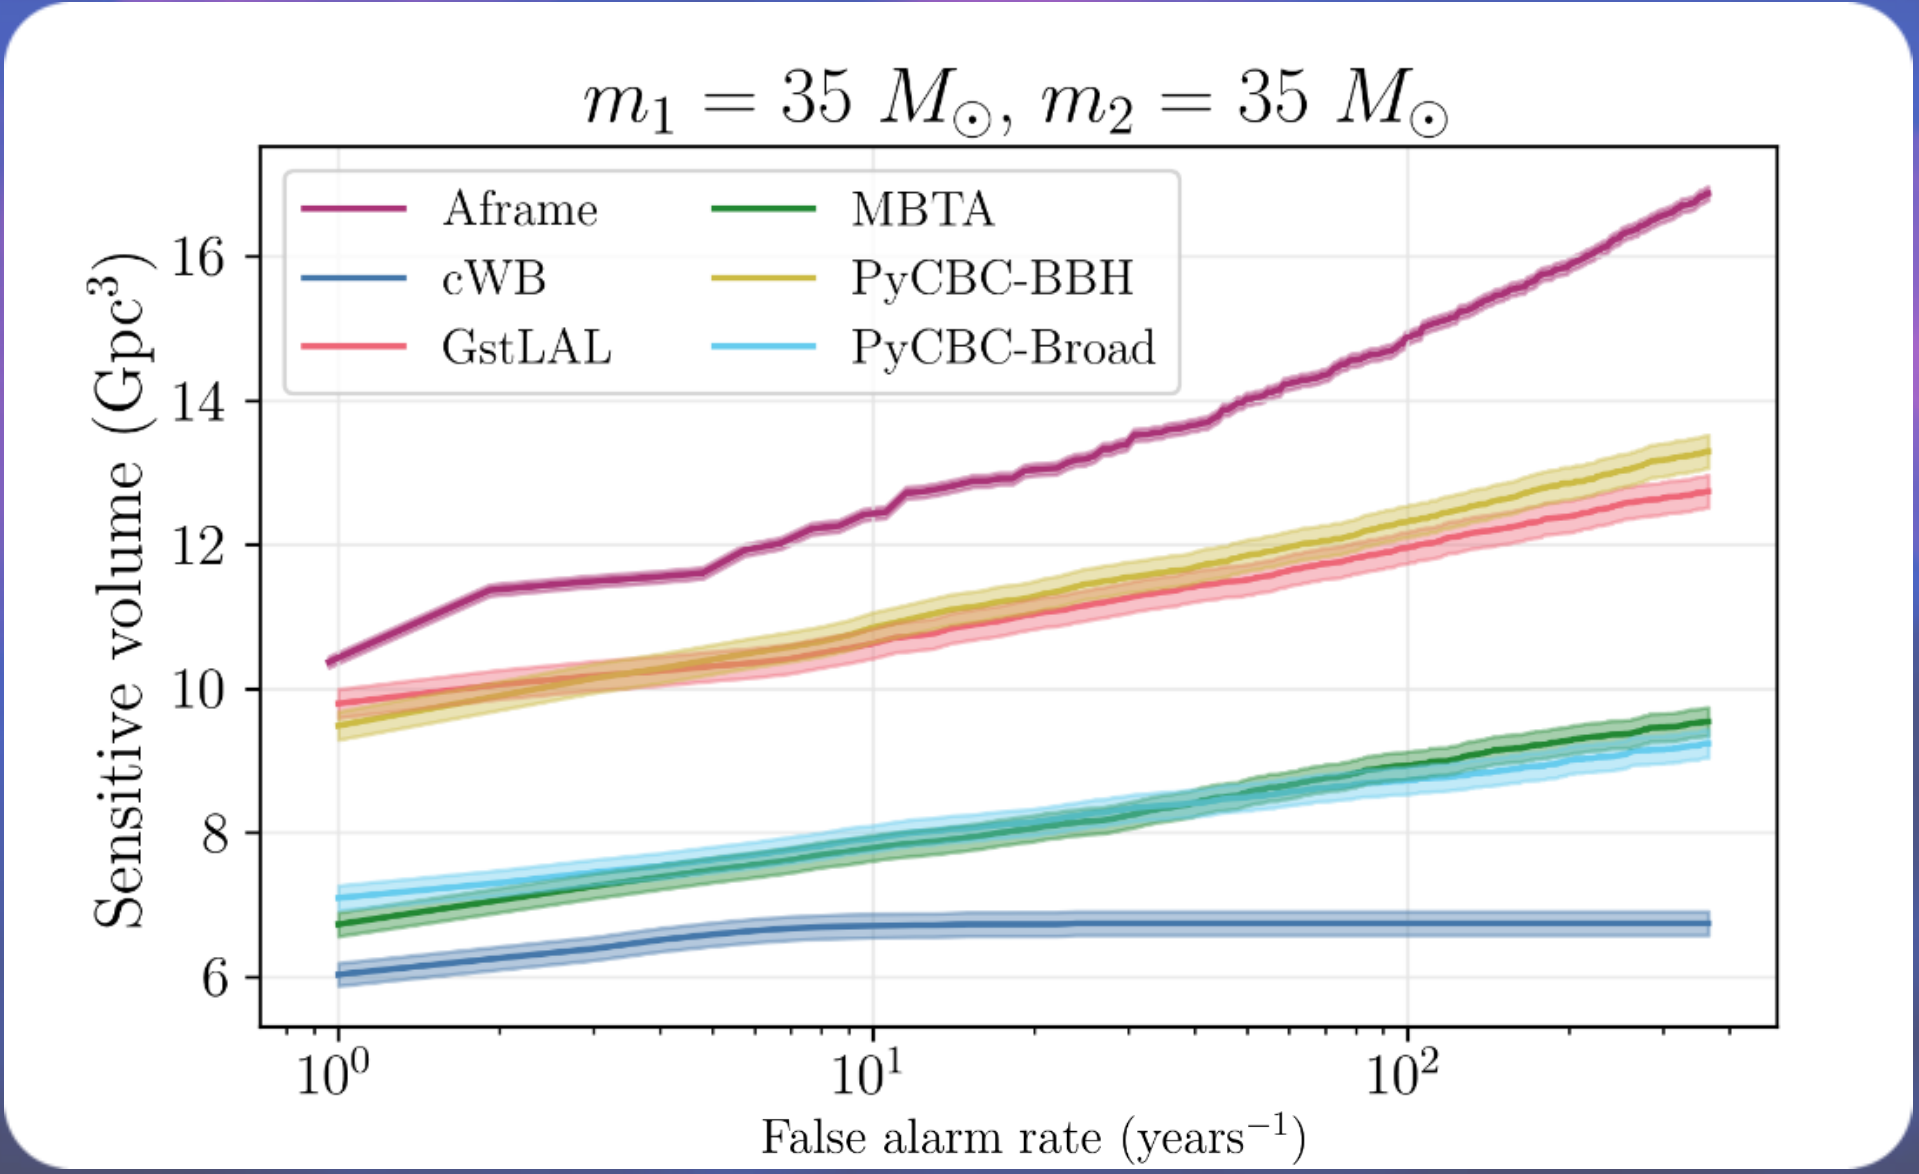

Future ML-based workflow

~100k auxiliary channels

Data 16kHz

Detector characterisation

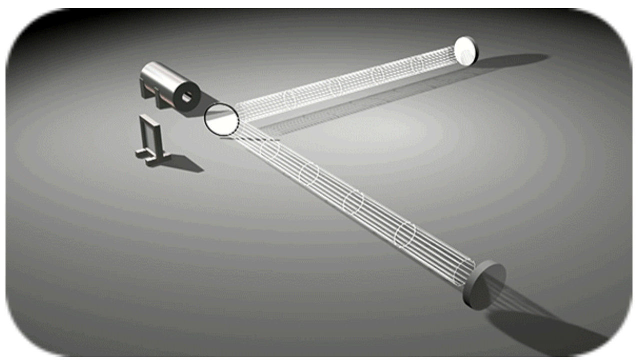

DeepClean
NN based AE
noise subtraction

Cleaned Data

NN-based algos for Event detection

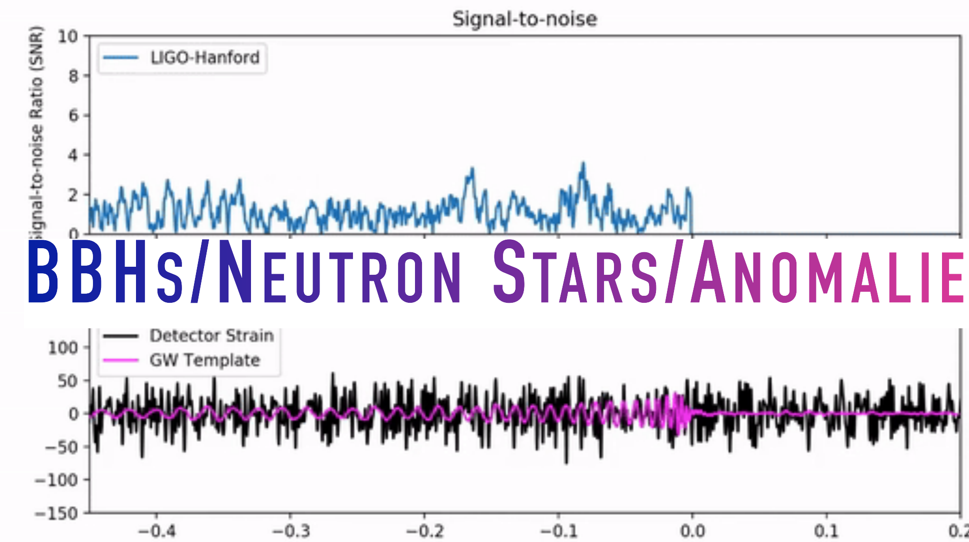

BBHs/Neutron Stars/Anomalies

Known "Unknowns" possible signal sources that are poorly modelled and therefore cannot be easily detected using the match filtering pipeline

Core-collapse supernova (CCSN)

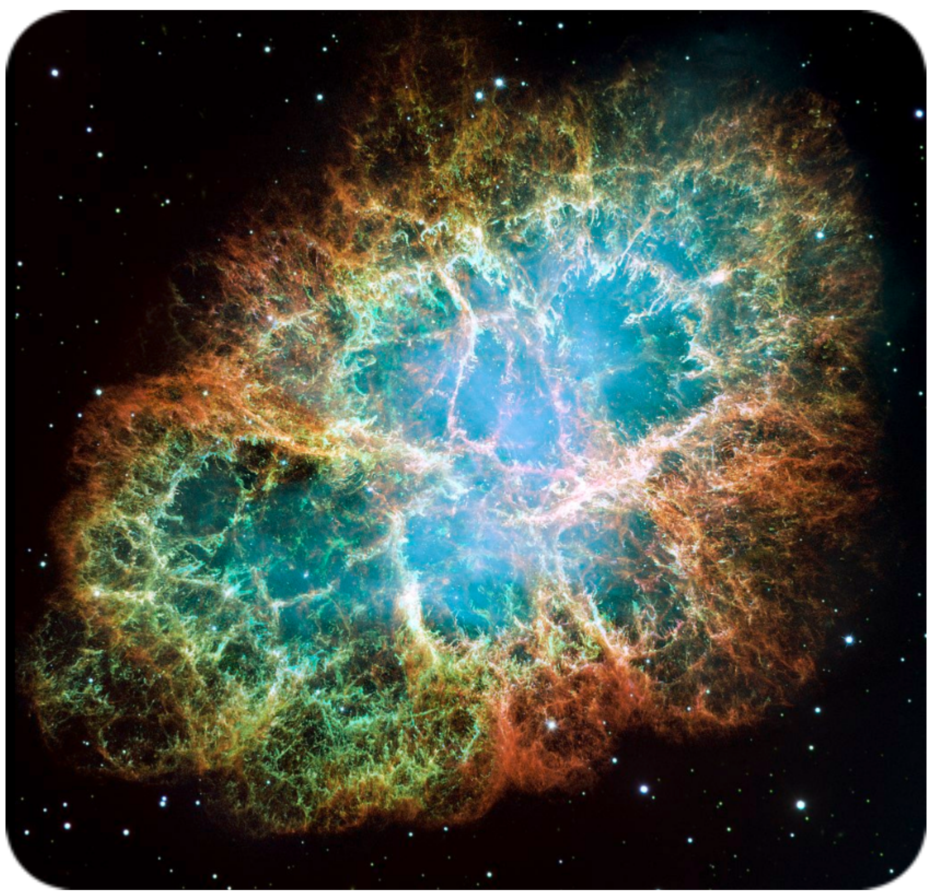

Neutron Star Glitches

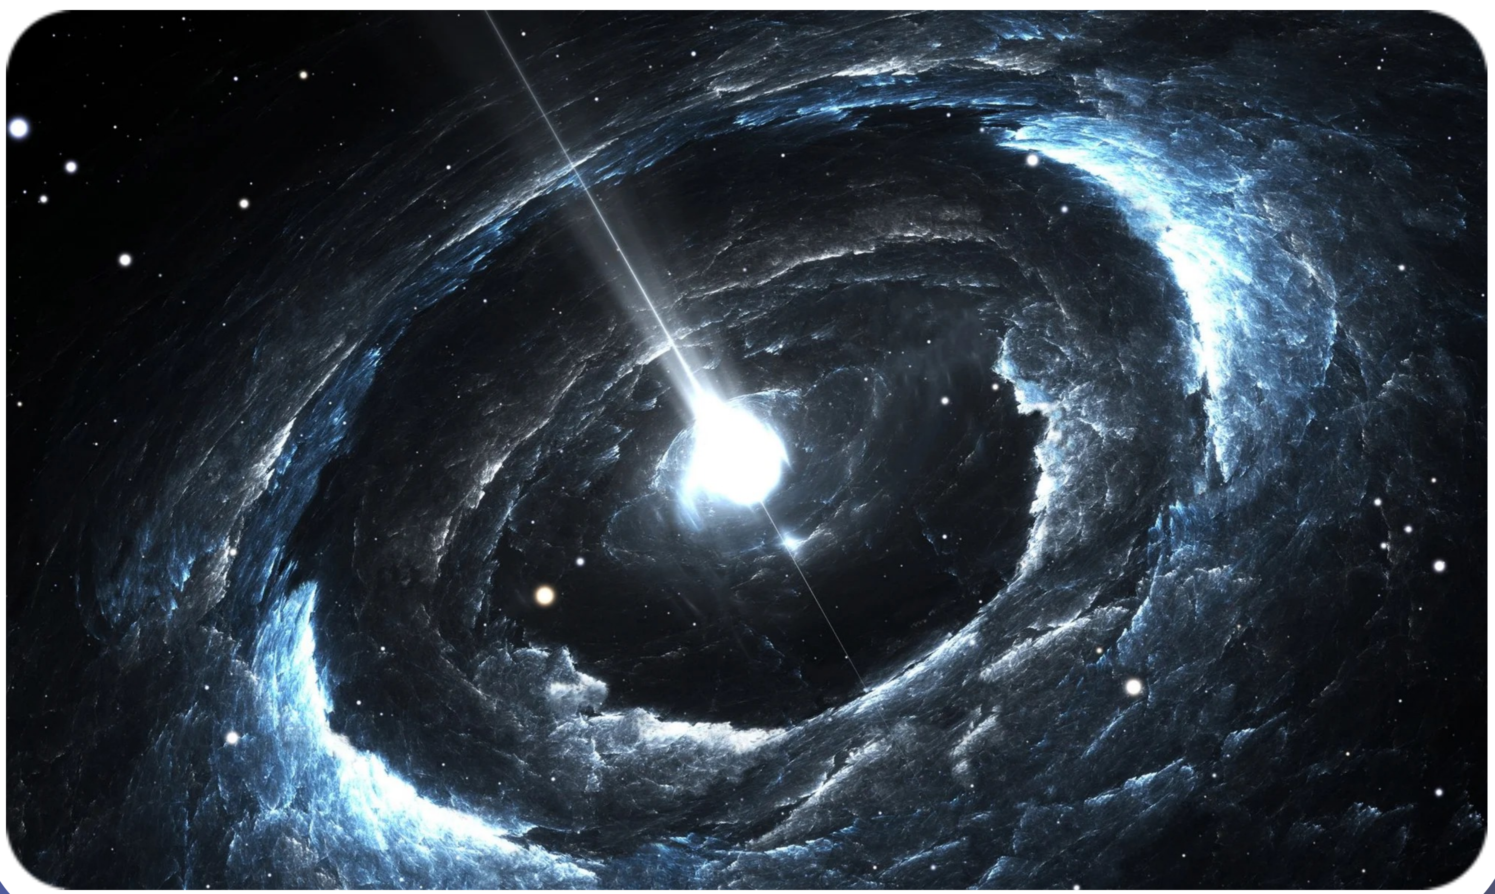

Unknown "Unknowns" new, unexpected GW sources

We refer to them as anomalous and aim to develop a semi-supervised approach which would let us to discover anomalous signals without explicit modelling

?

GWAK: GW Anomalous Knowledge

Vanilla Anomaly Detection

K.Govorkova et al MLST
10.1088/2632-2153/ad3a31

The algorithm is inspired by QUAK arXiv2011.03550 from LHC HEP

Use the distance between the input and output as a metric for anomaly detection

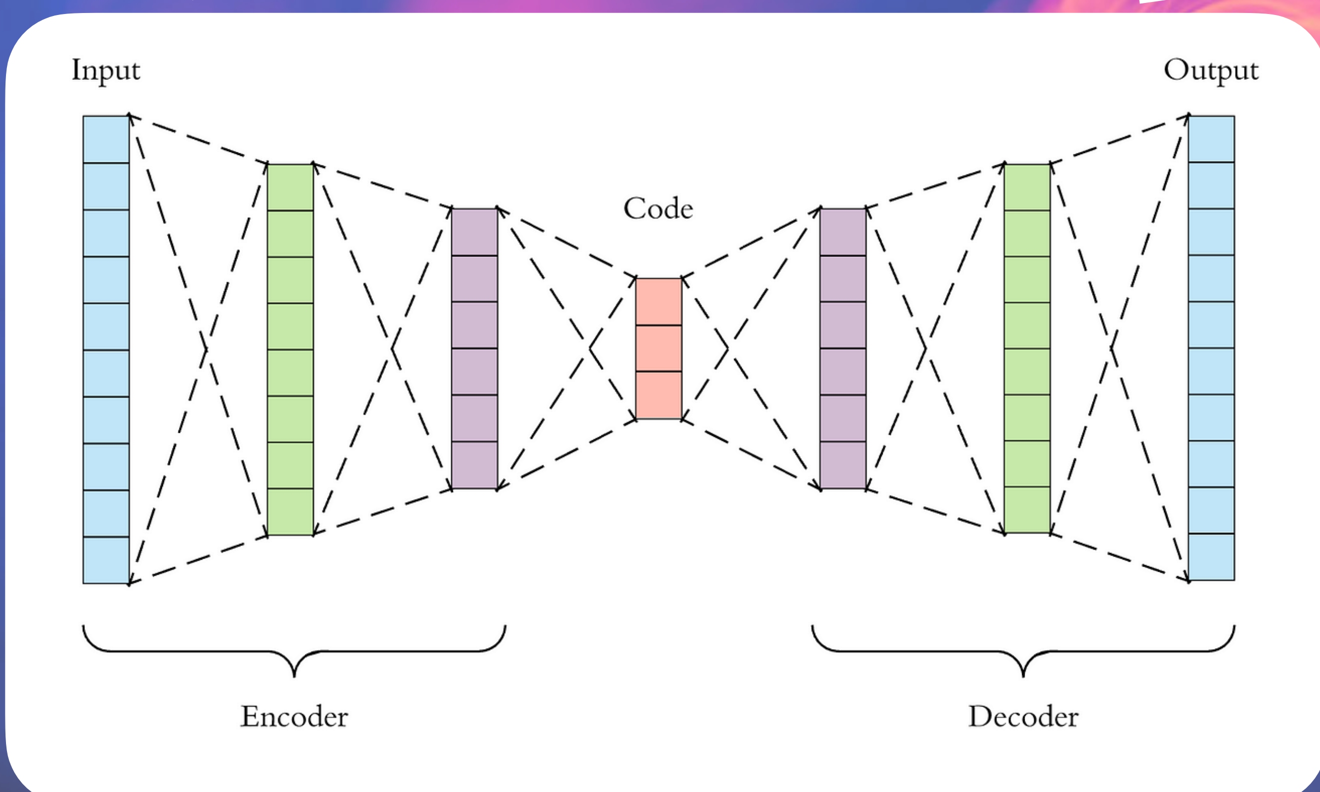

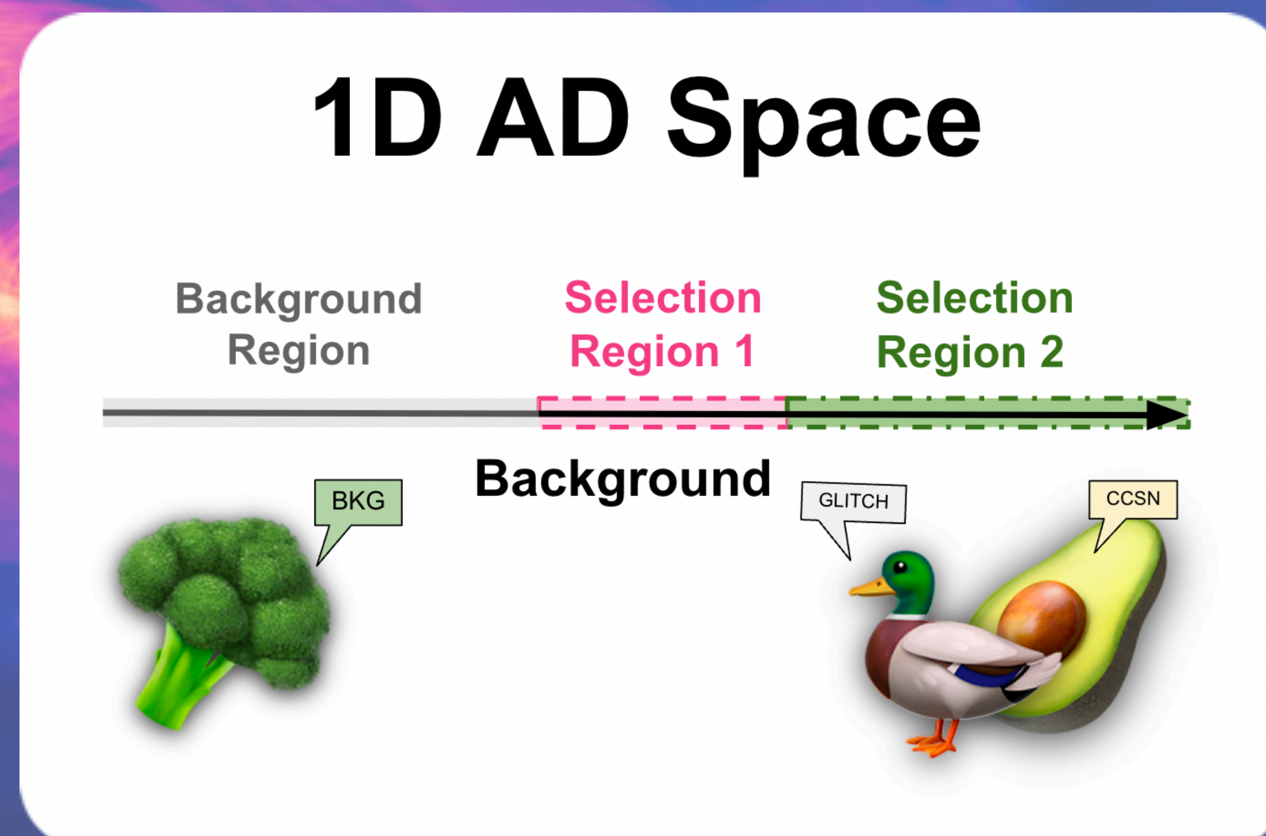

K.Govorkova et al MLST
10.1088/2632-2153/ad3a31

Including more axes, both signal and background, allows to more efficiently select a signal-like anomalies

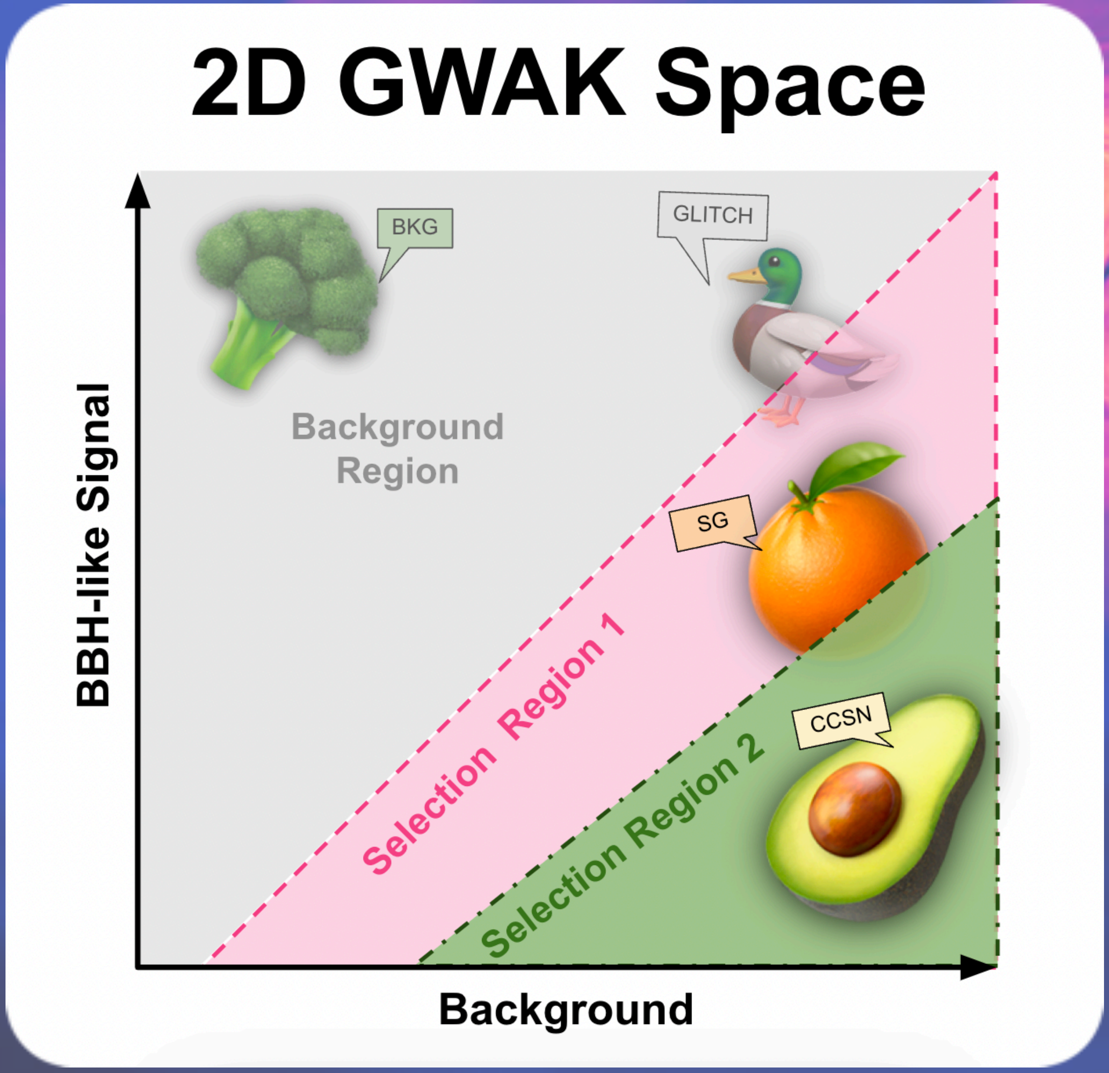

K.Govorkova et al MLST
10.1088/2632-2153/ad3a31

Including more axes, both signal and background, allows to more efficiently select a signal-like anomalies

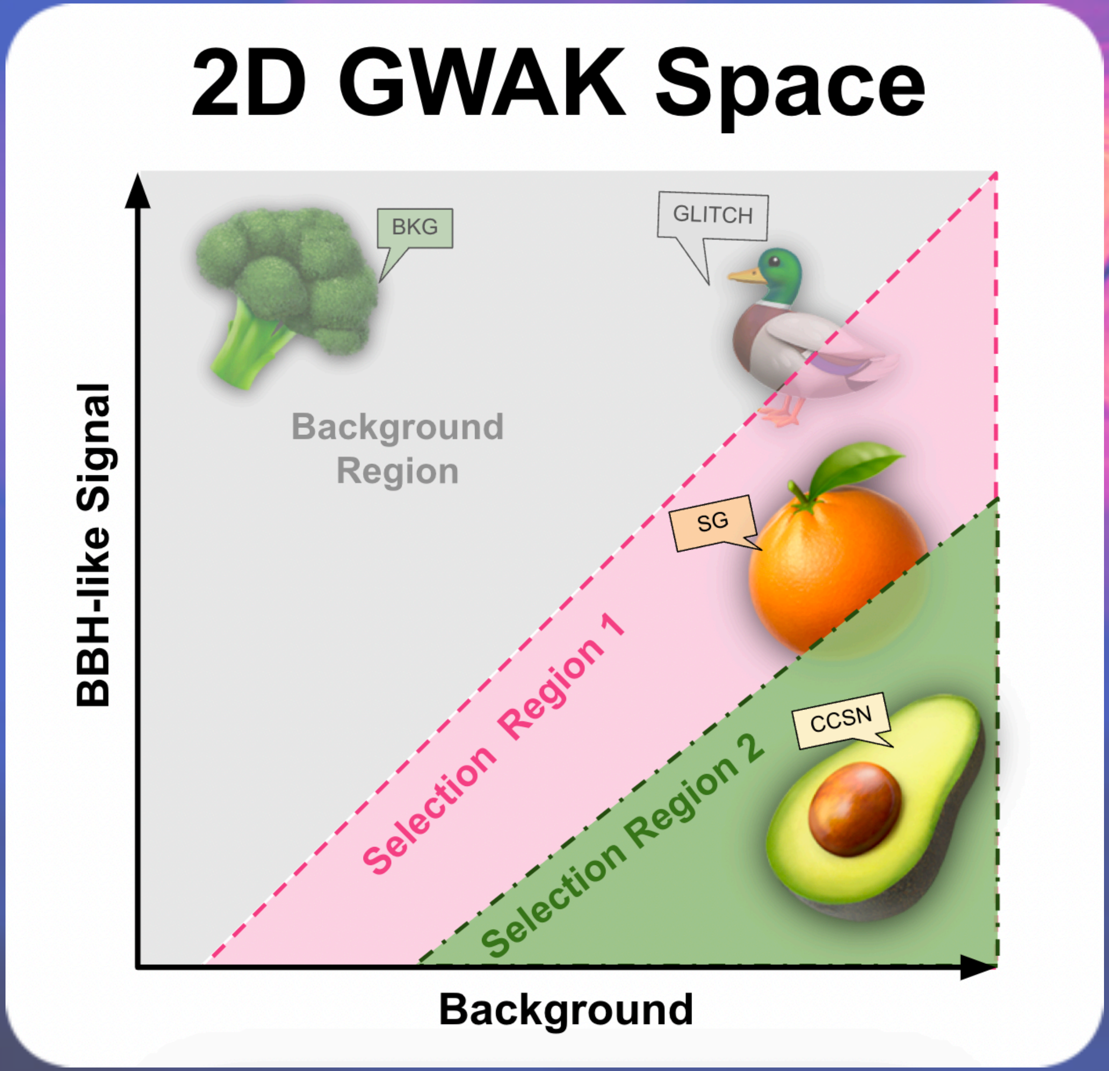

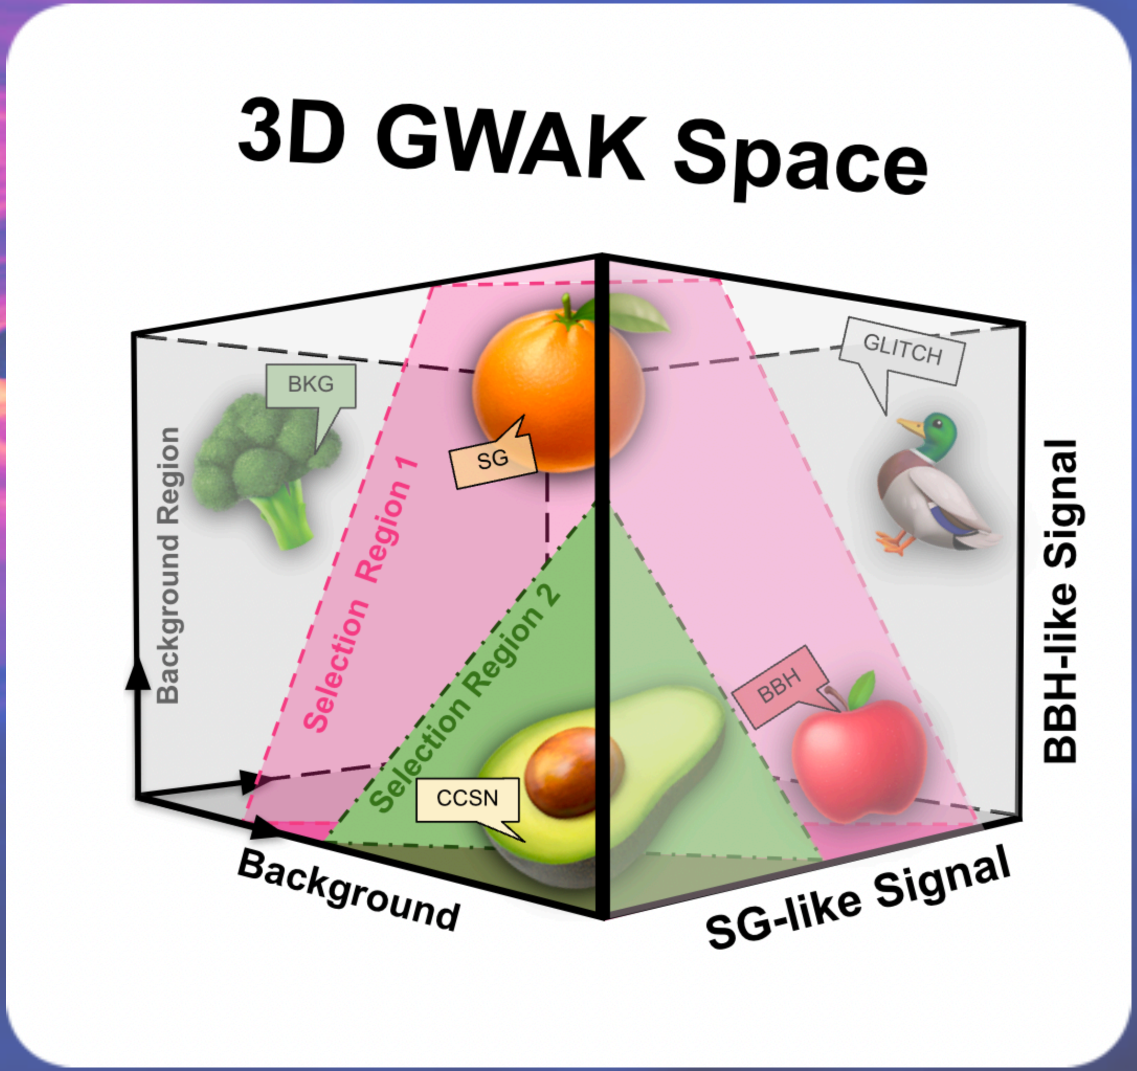

K.Govorkova et al MLST
10.1088/2632-2153/ad3a31

Strain, GWAK metric response and final metric response for Supernova simulated signal

The evaluation of GWAK axes and Pearson correlation with time and on the top right total metric value and FAR are shown as an example of the algorithm's ''reaction'' to an unseen signal

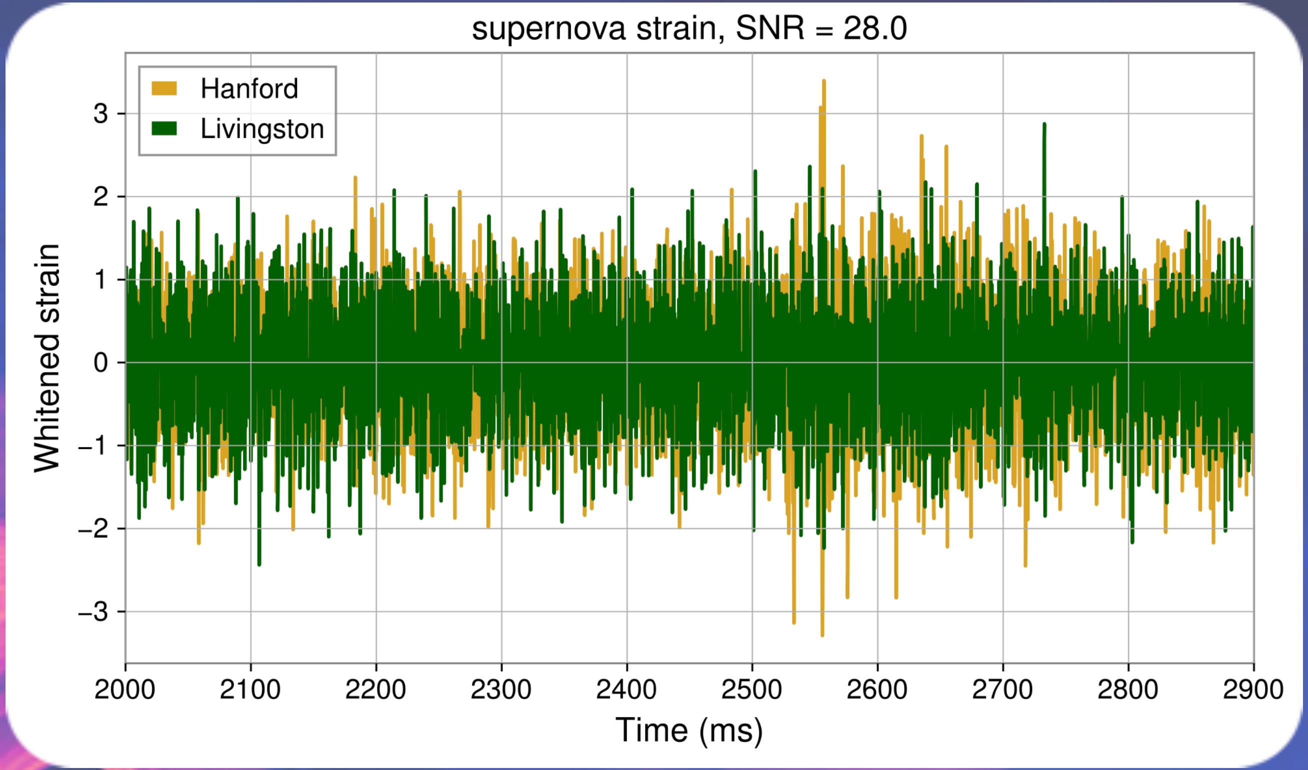

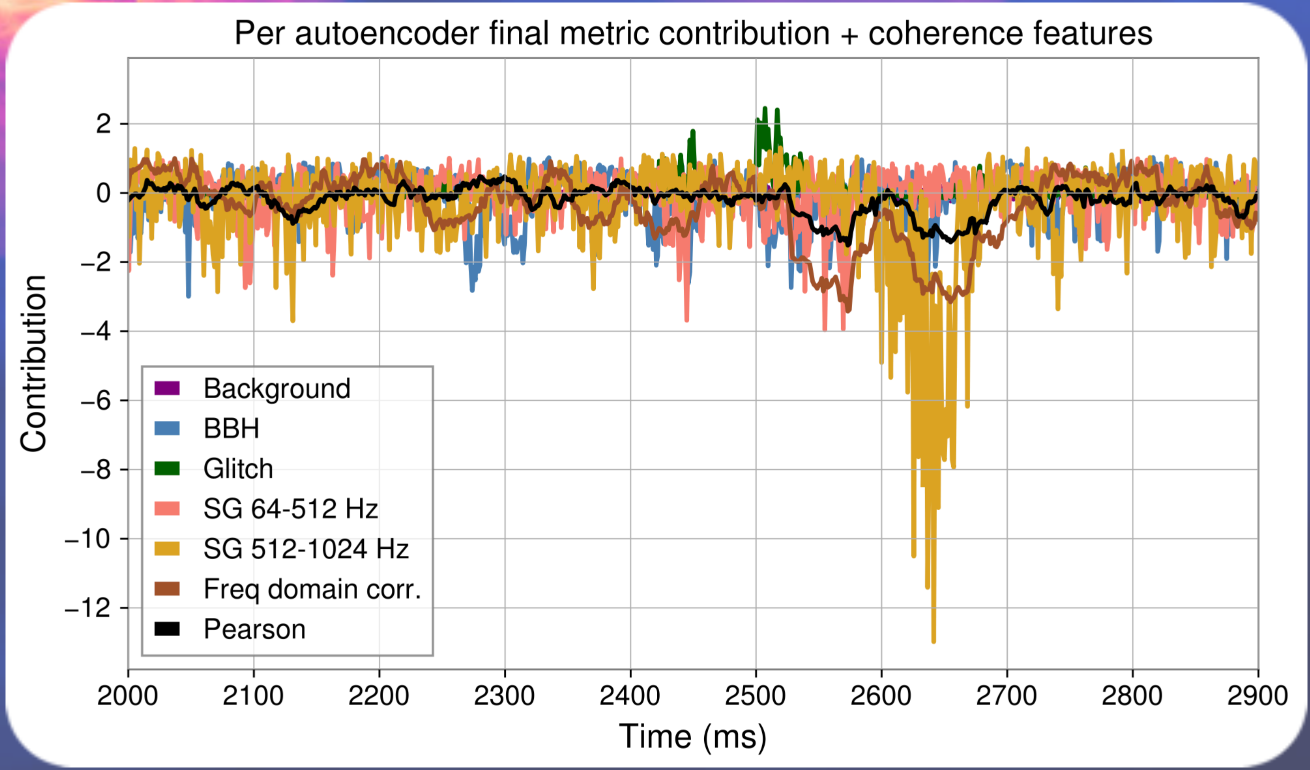

K.Govorkova et al MLST 10.1088/2632-2153/ad3a31

Strain, GWAK metric response and final metric response for Supernova simulated signal

The evaluation of GWAK axes and Pearson correlation with time and on the top right total metric value and FAR are shown as an example of the algorithm's ''reaction'' to an unseen signal

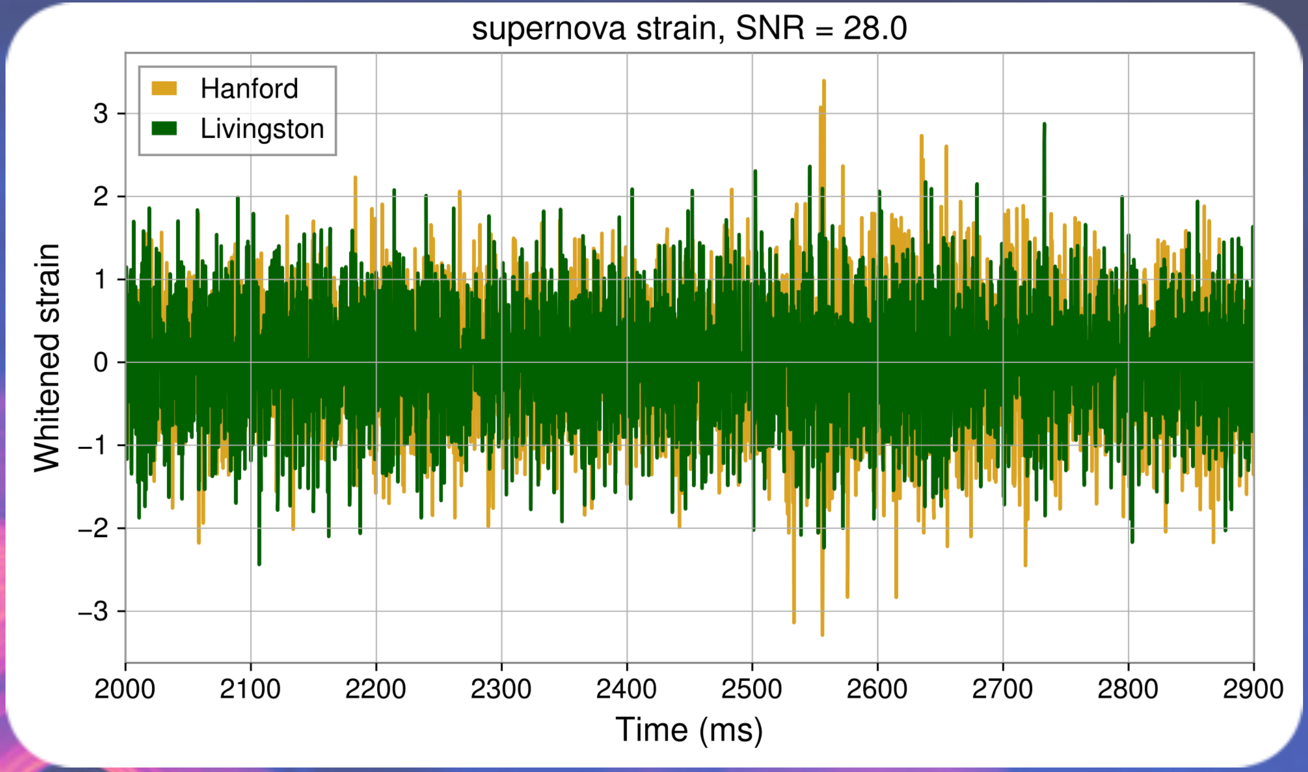

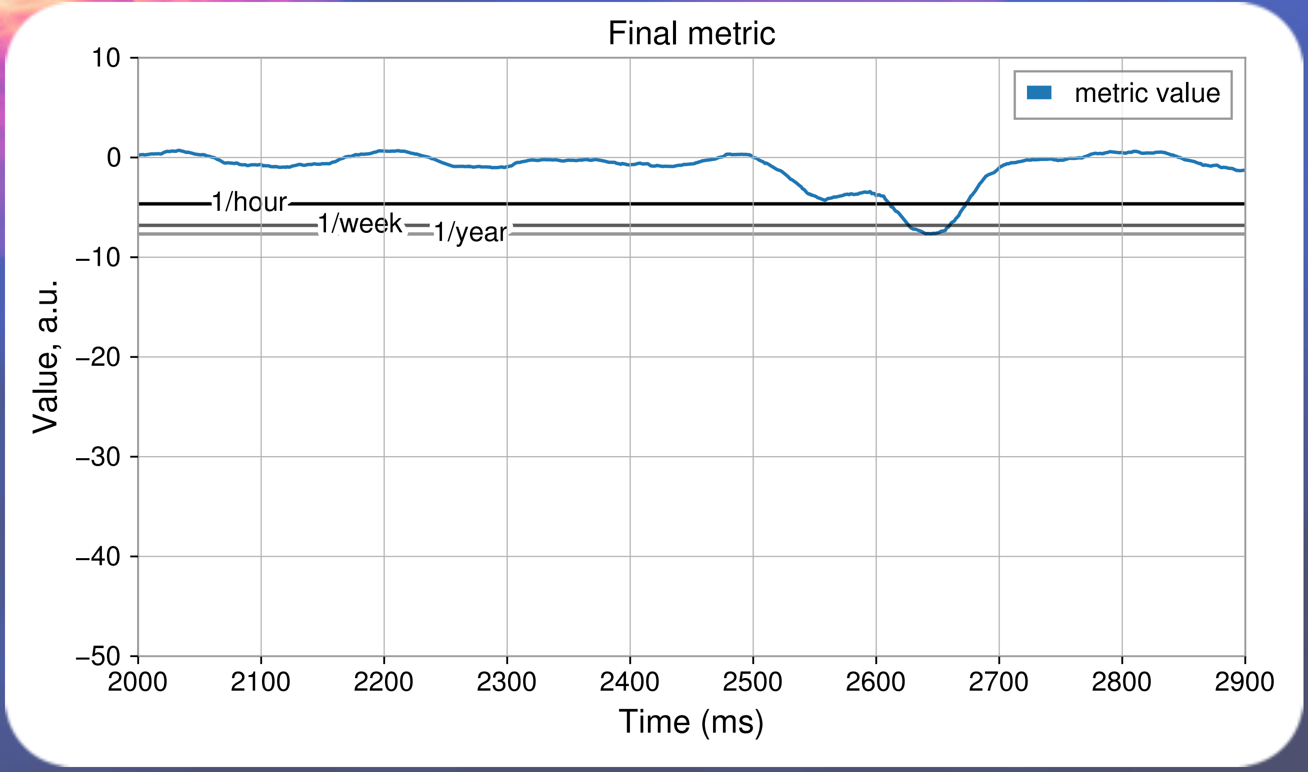

GWAK Detection

GWAK Detection

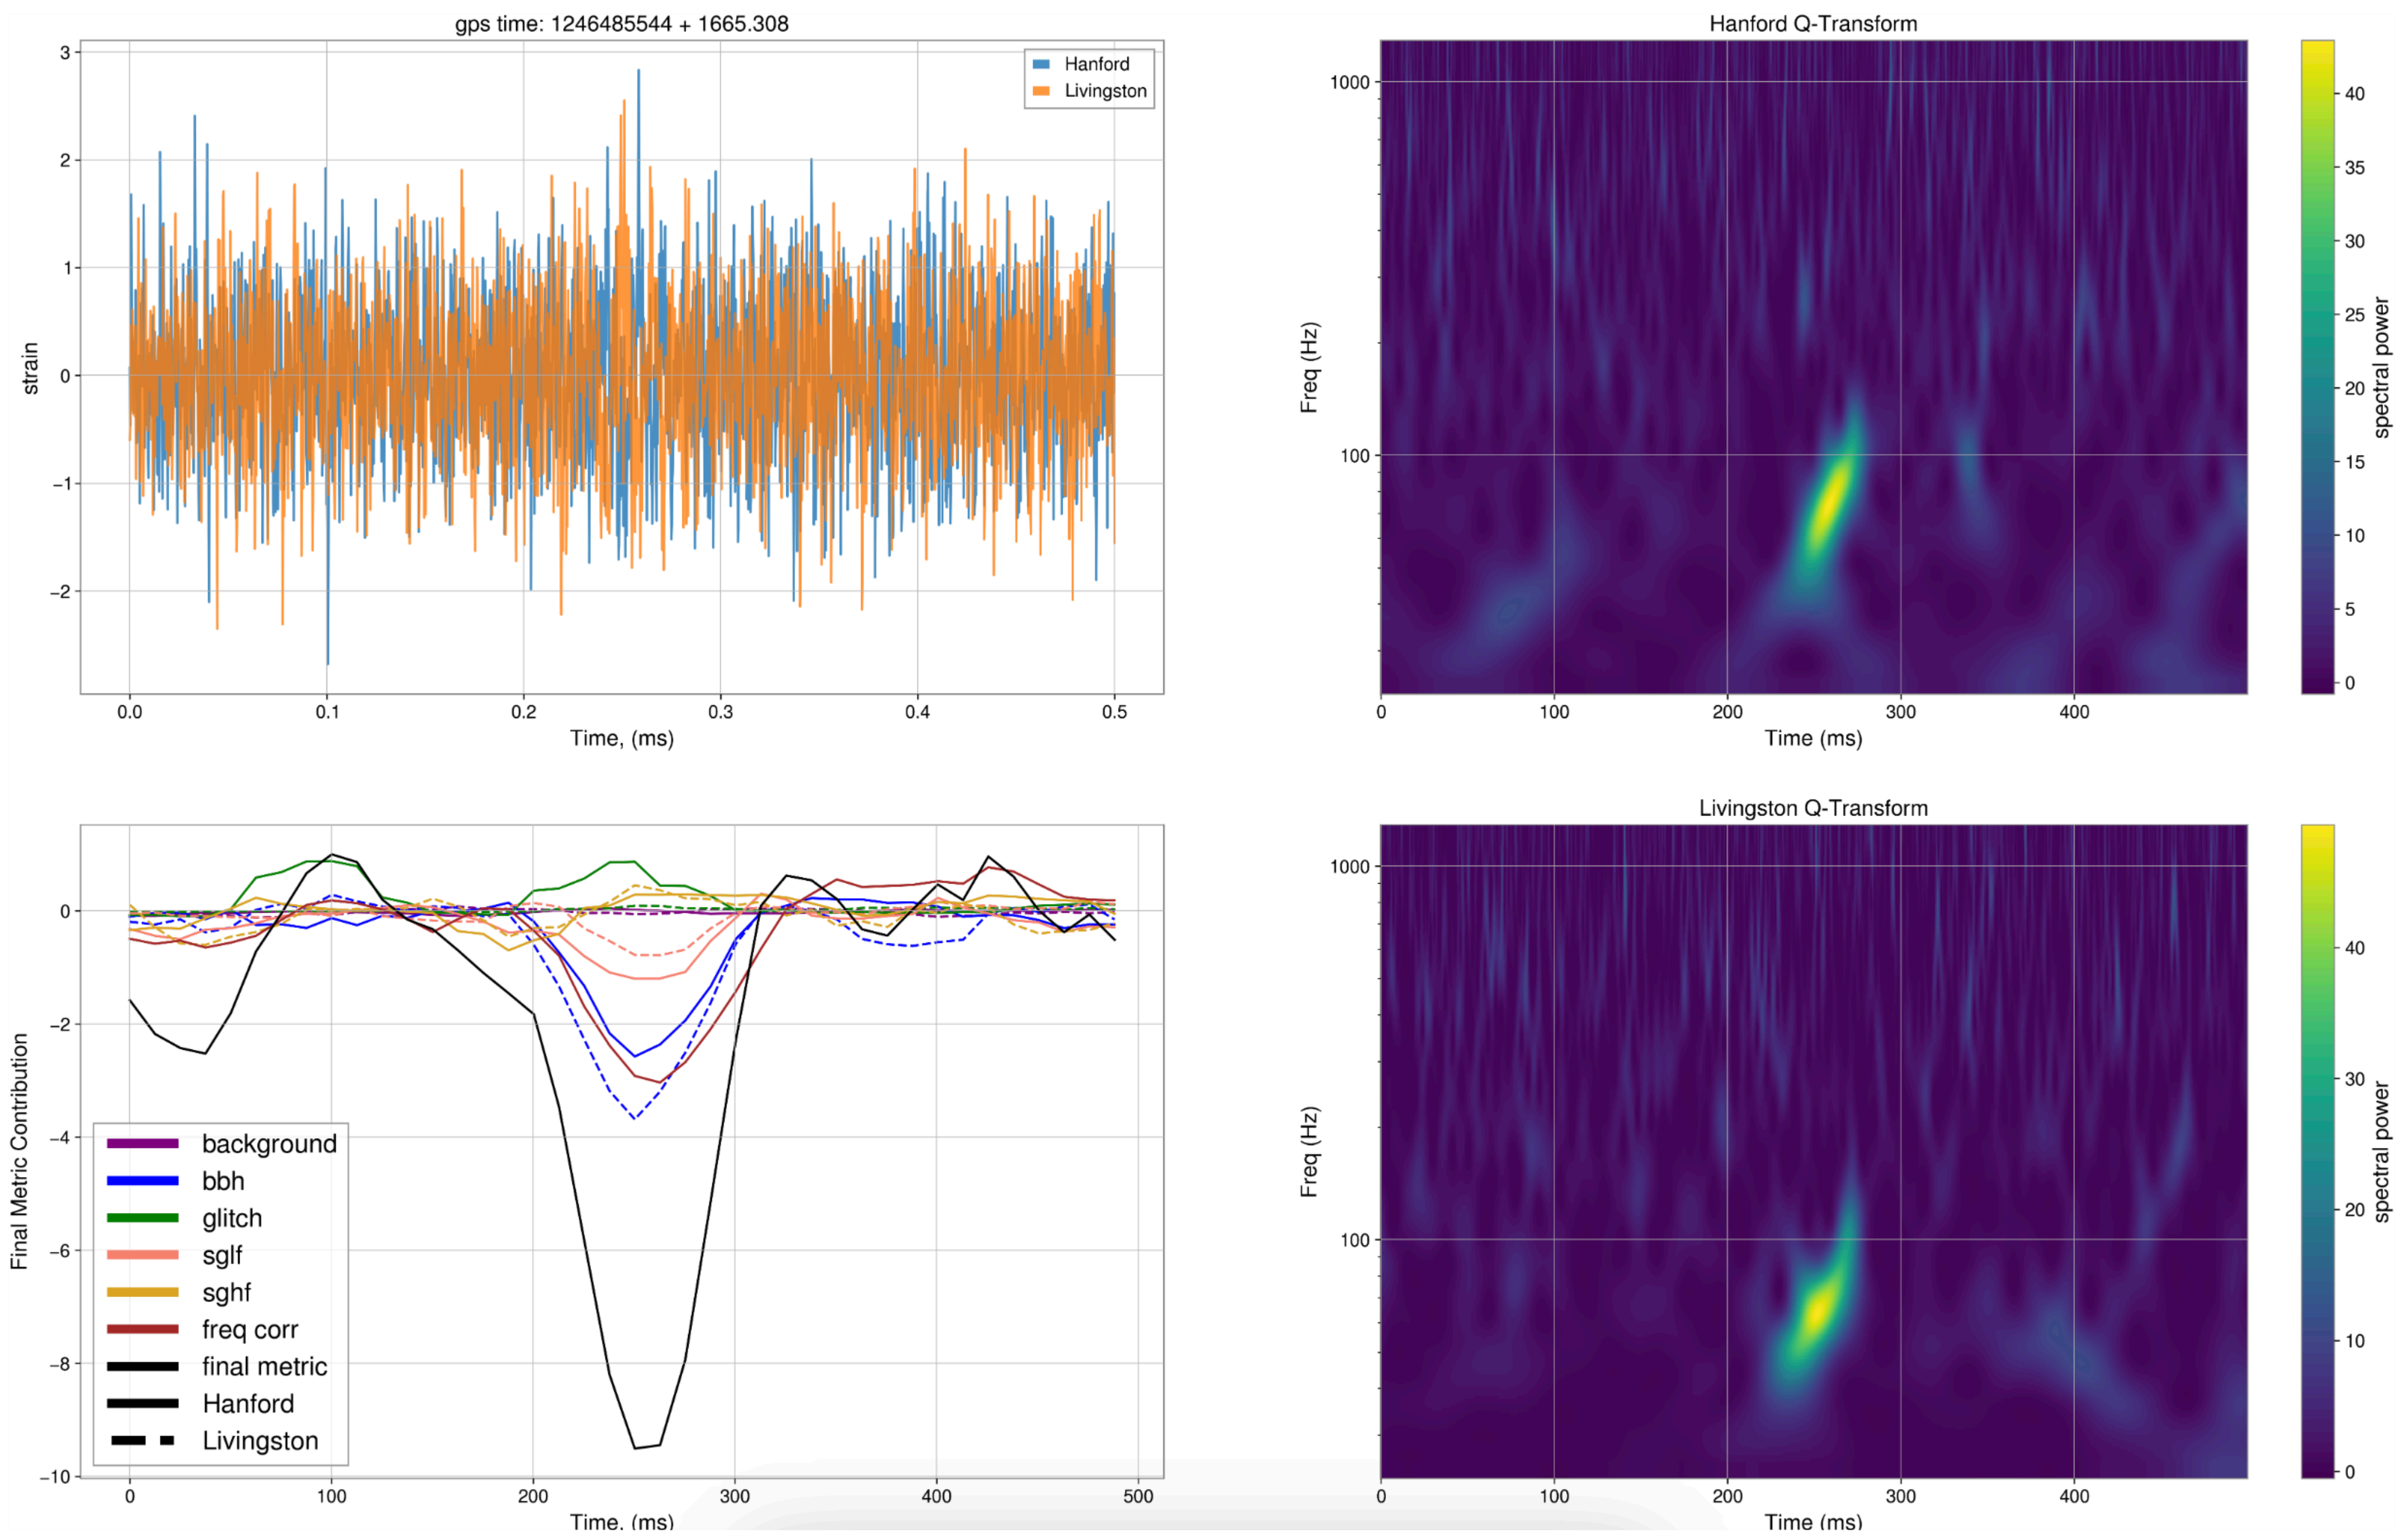

Future ML-based workflow

Data 16kHz

~100k auxiliary channels

Detector characterisation

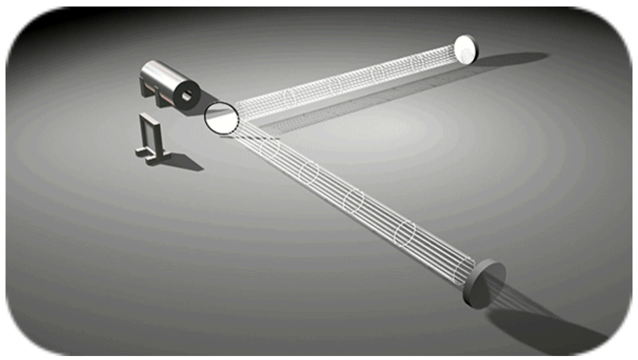

DeepClean
NN based AE
noise subtraction

Cleaned Data

NN-based algos for Event detection

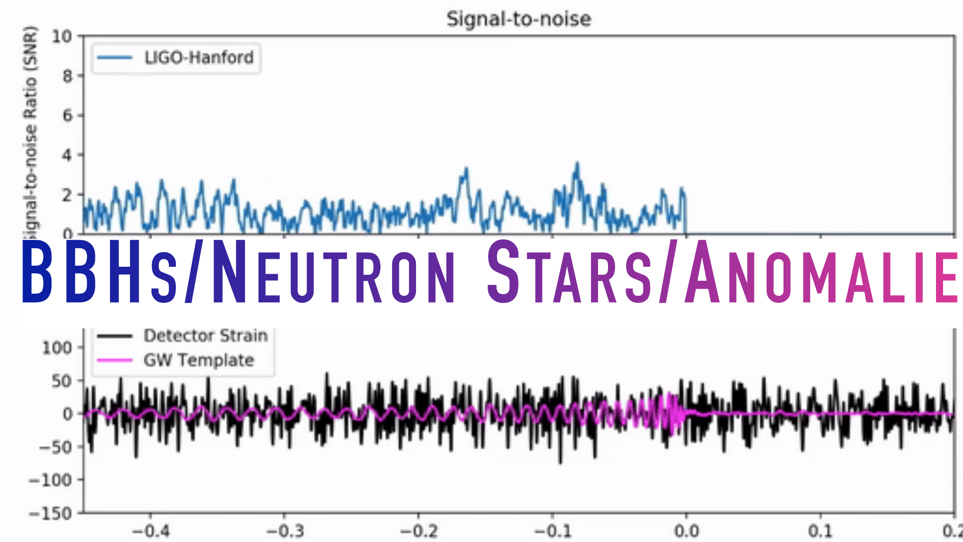

BBHs/Neutron Stars/Anomalies

Event

NF for Event characterisation

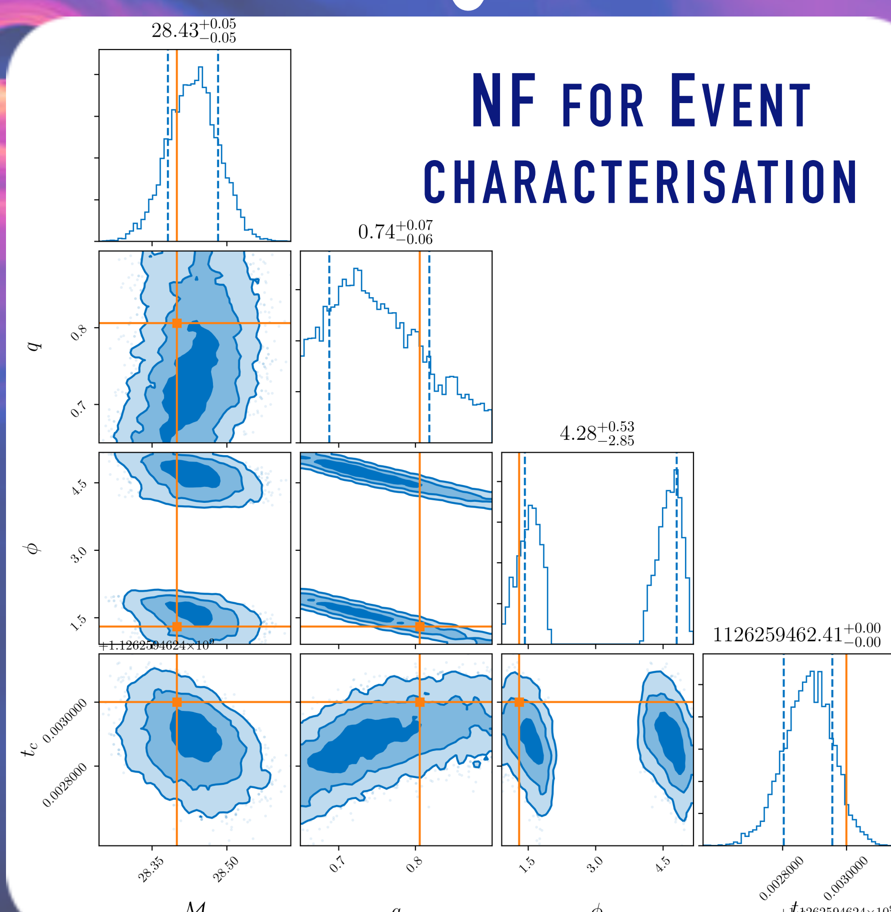

Alert 🚨

AMPLFI: Accelerated Multi-messenger Parameter estimation using Likelihood Free Inference

NeurIPS ML4PS 2023 69 PDF

Perform fast parameter estimation using simulation-based inference

- Simulate data from the likelihood, train neural network to approximate posterior
- Use self-supervision to marginalize symmetries

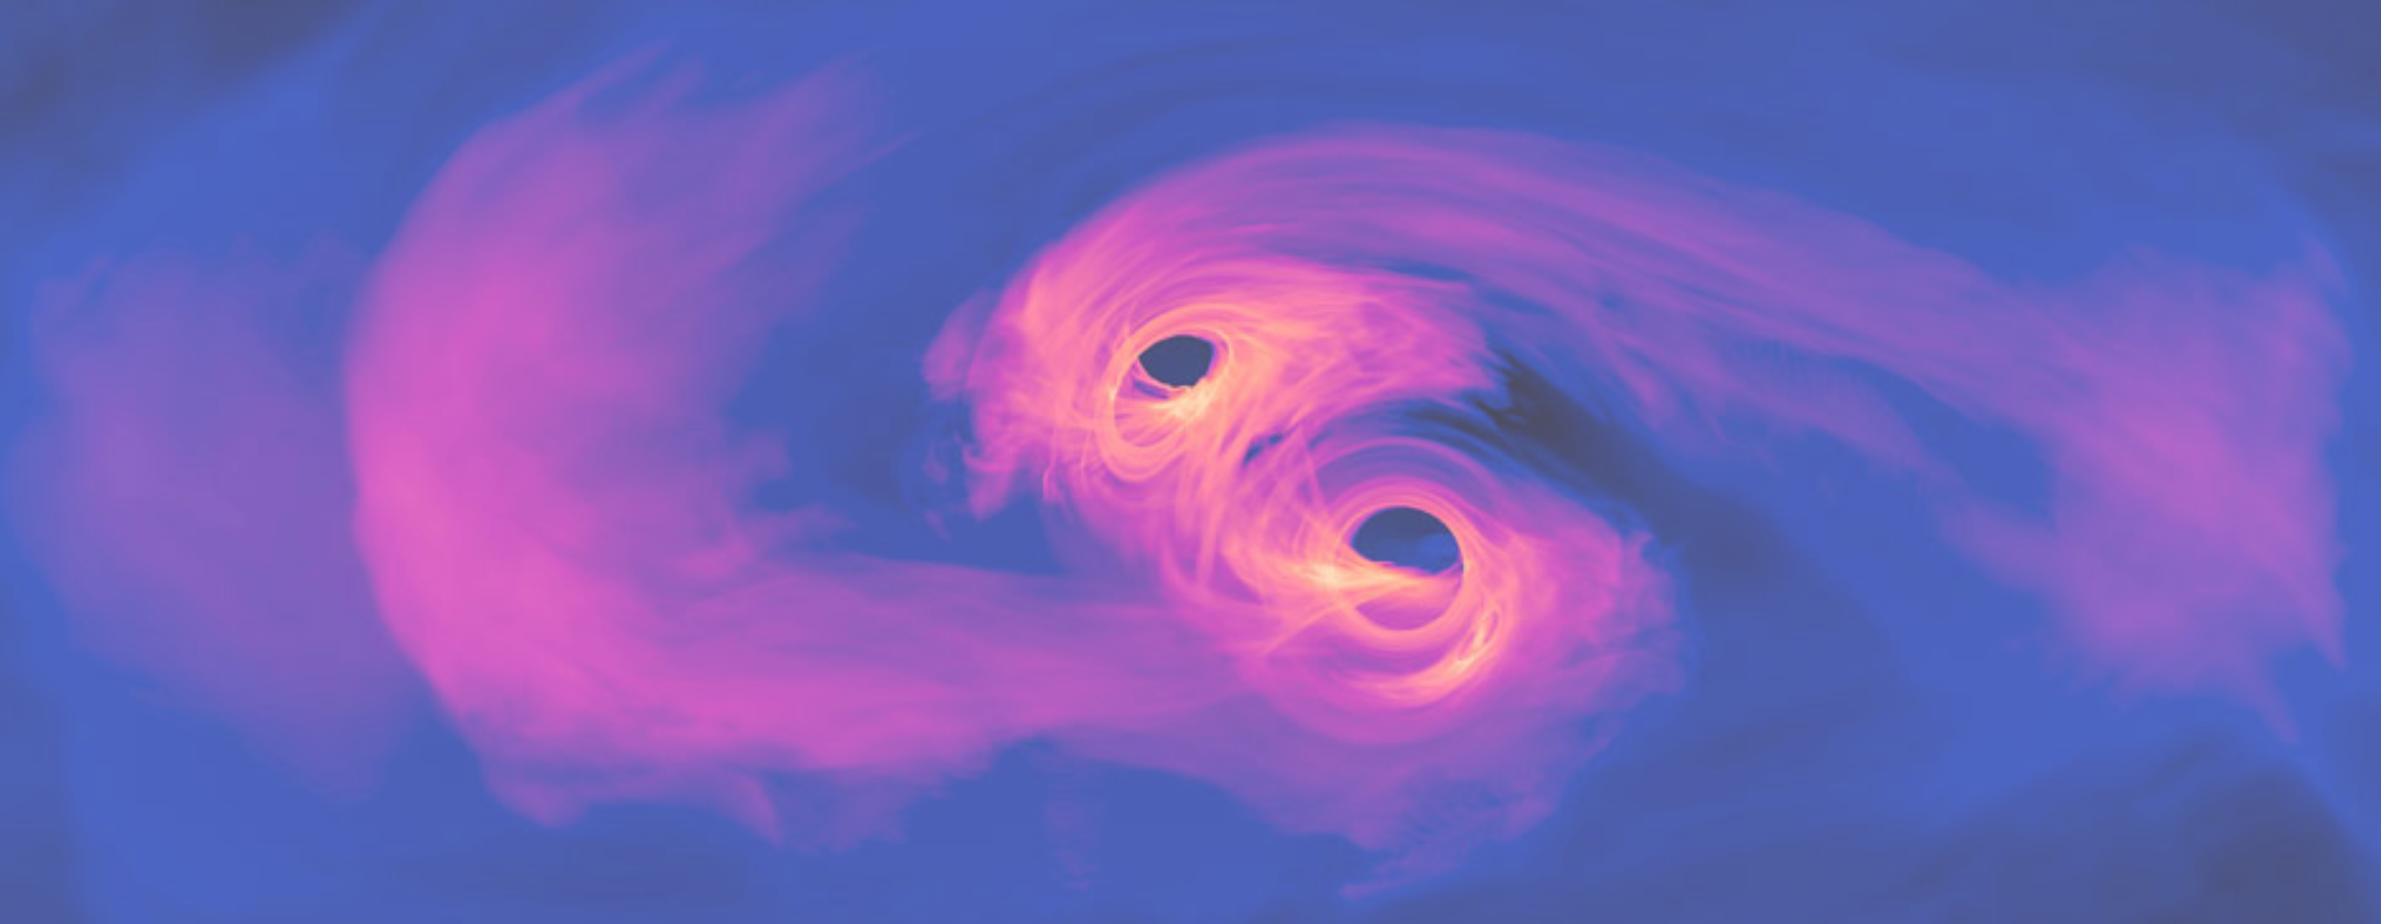

AMPLFI: Accelerated Multi-messenger Parameter estimation using Likelihood Free Inference

Perform fast parameter estimation using simulation-based inference

- Simulate data from the likelihood, train neural network to approximate posterior
- Use self-supervision to marginalize over coalescence time
- Normalizing flows (invertible transforms map simple distribution to complex distribution) embed broad knowledge of waveforms

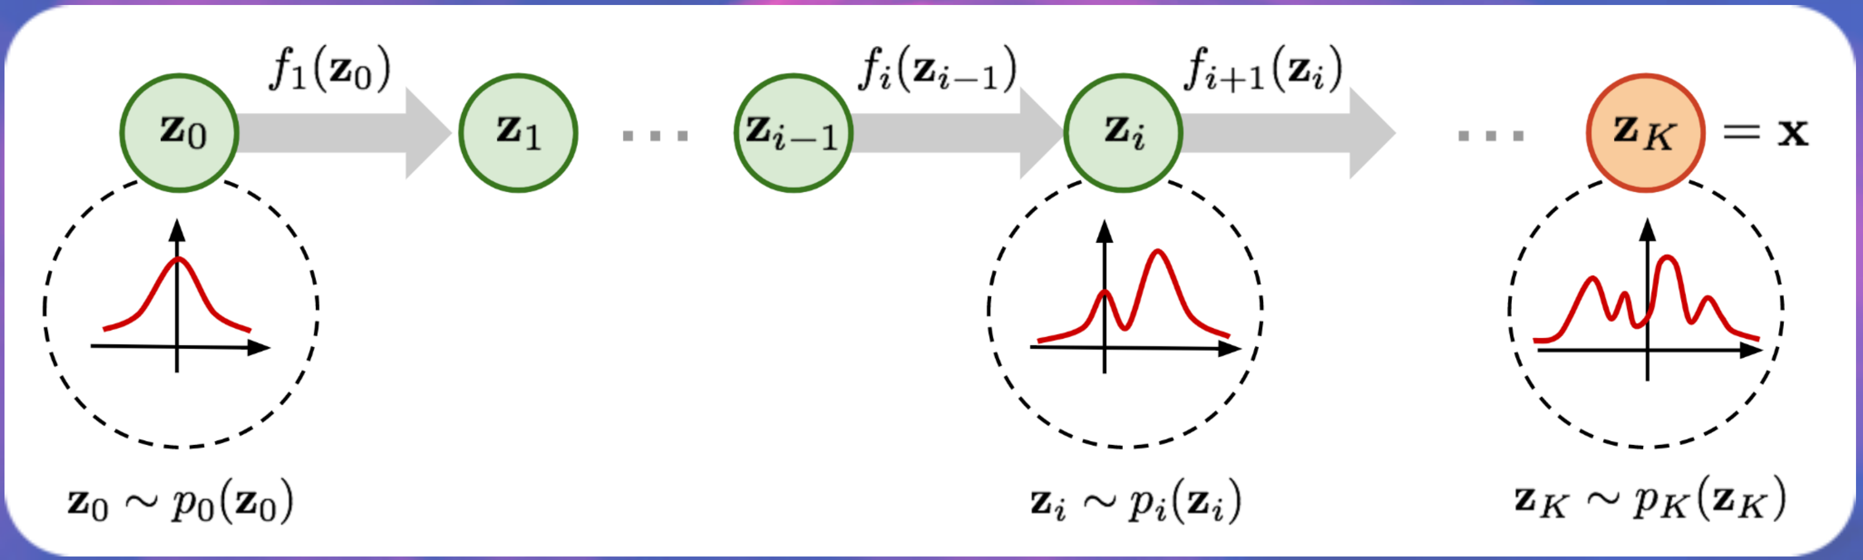

AMPLFI: Accelerated Multi-messenger Parameter estimation using Likelihood Free Inference

NeurIPS ML4PS 2023 69 PDF

Perform fast parameter estimation using simulation-based inference

- Simulate data from the likelihood, train neural network to approximate posterior
- Use self-supervision to marginalize over coalescence time
- Normalizing flows (invertible transforms map simple distribution to complex distribution) embed broad knowledge of waveforms

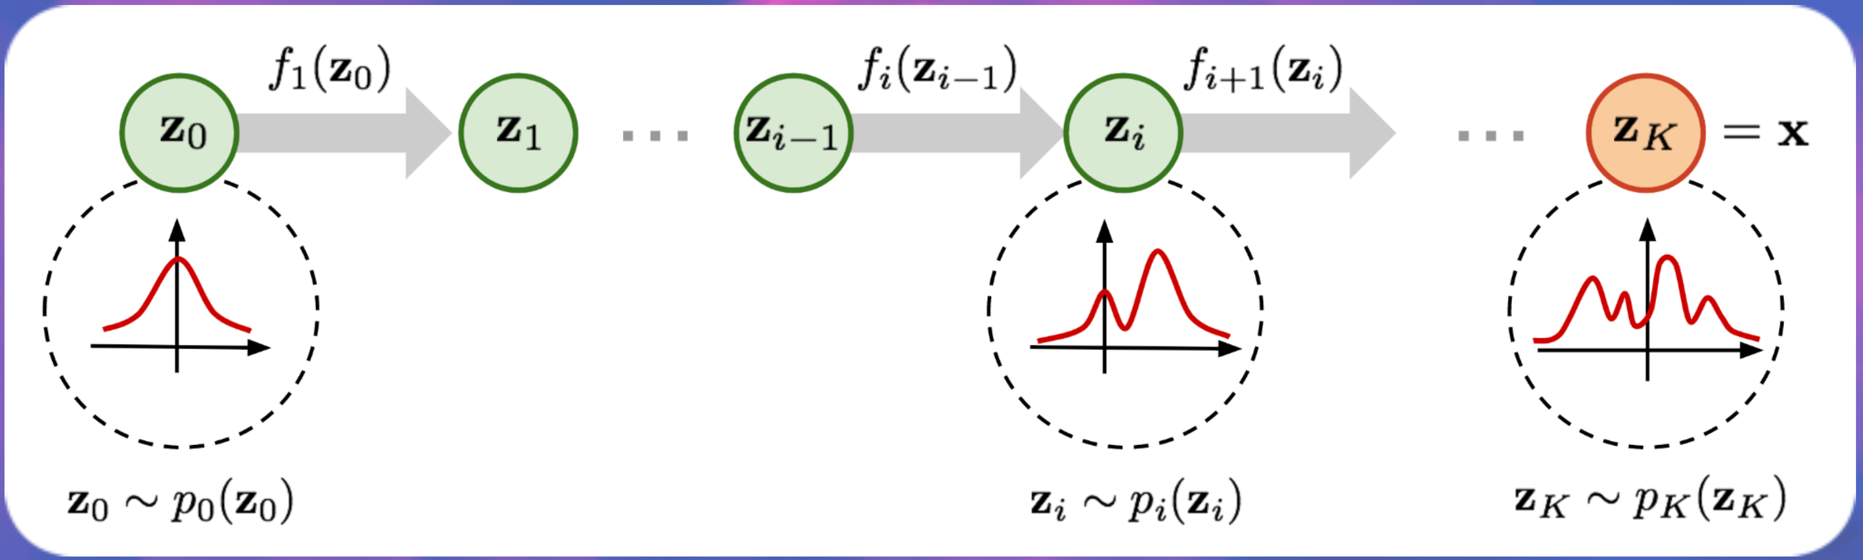

- PE done in seconds!

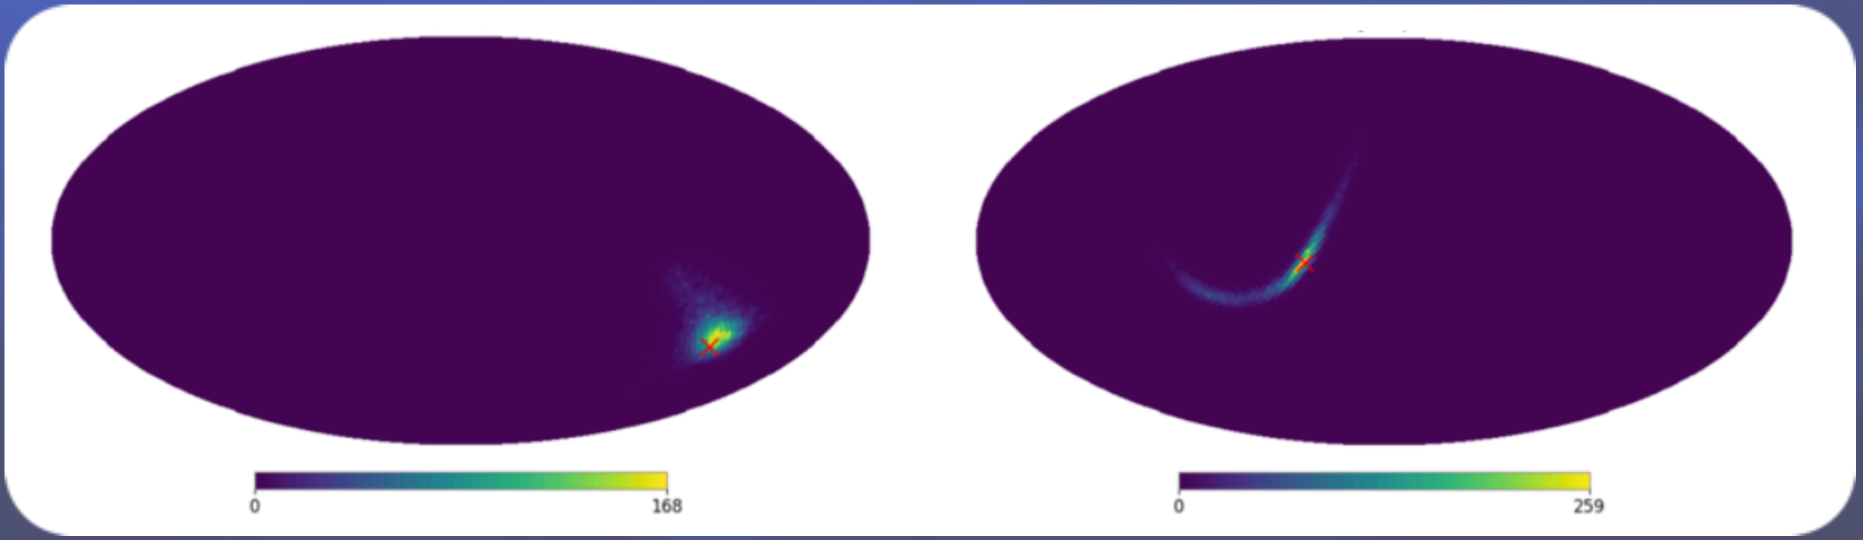

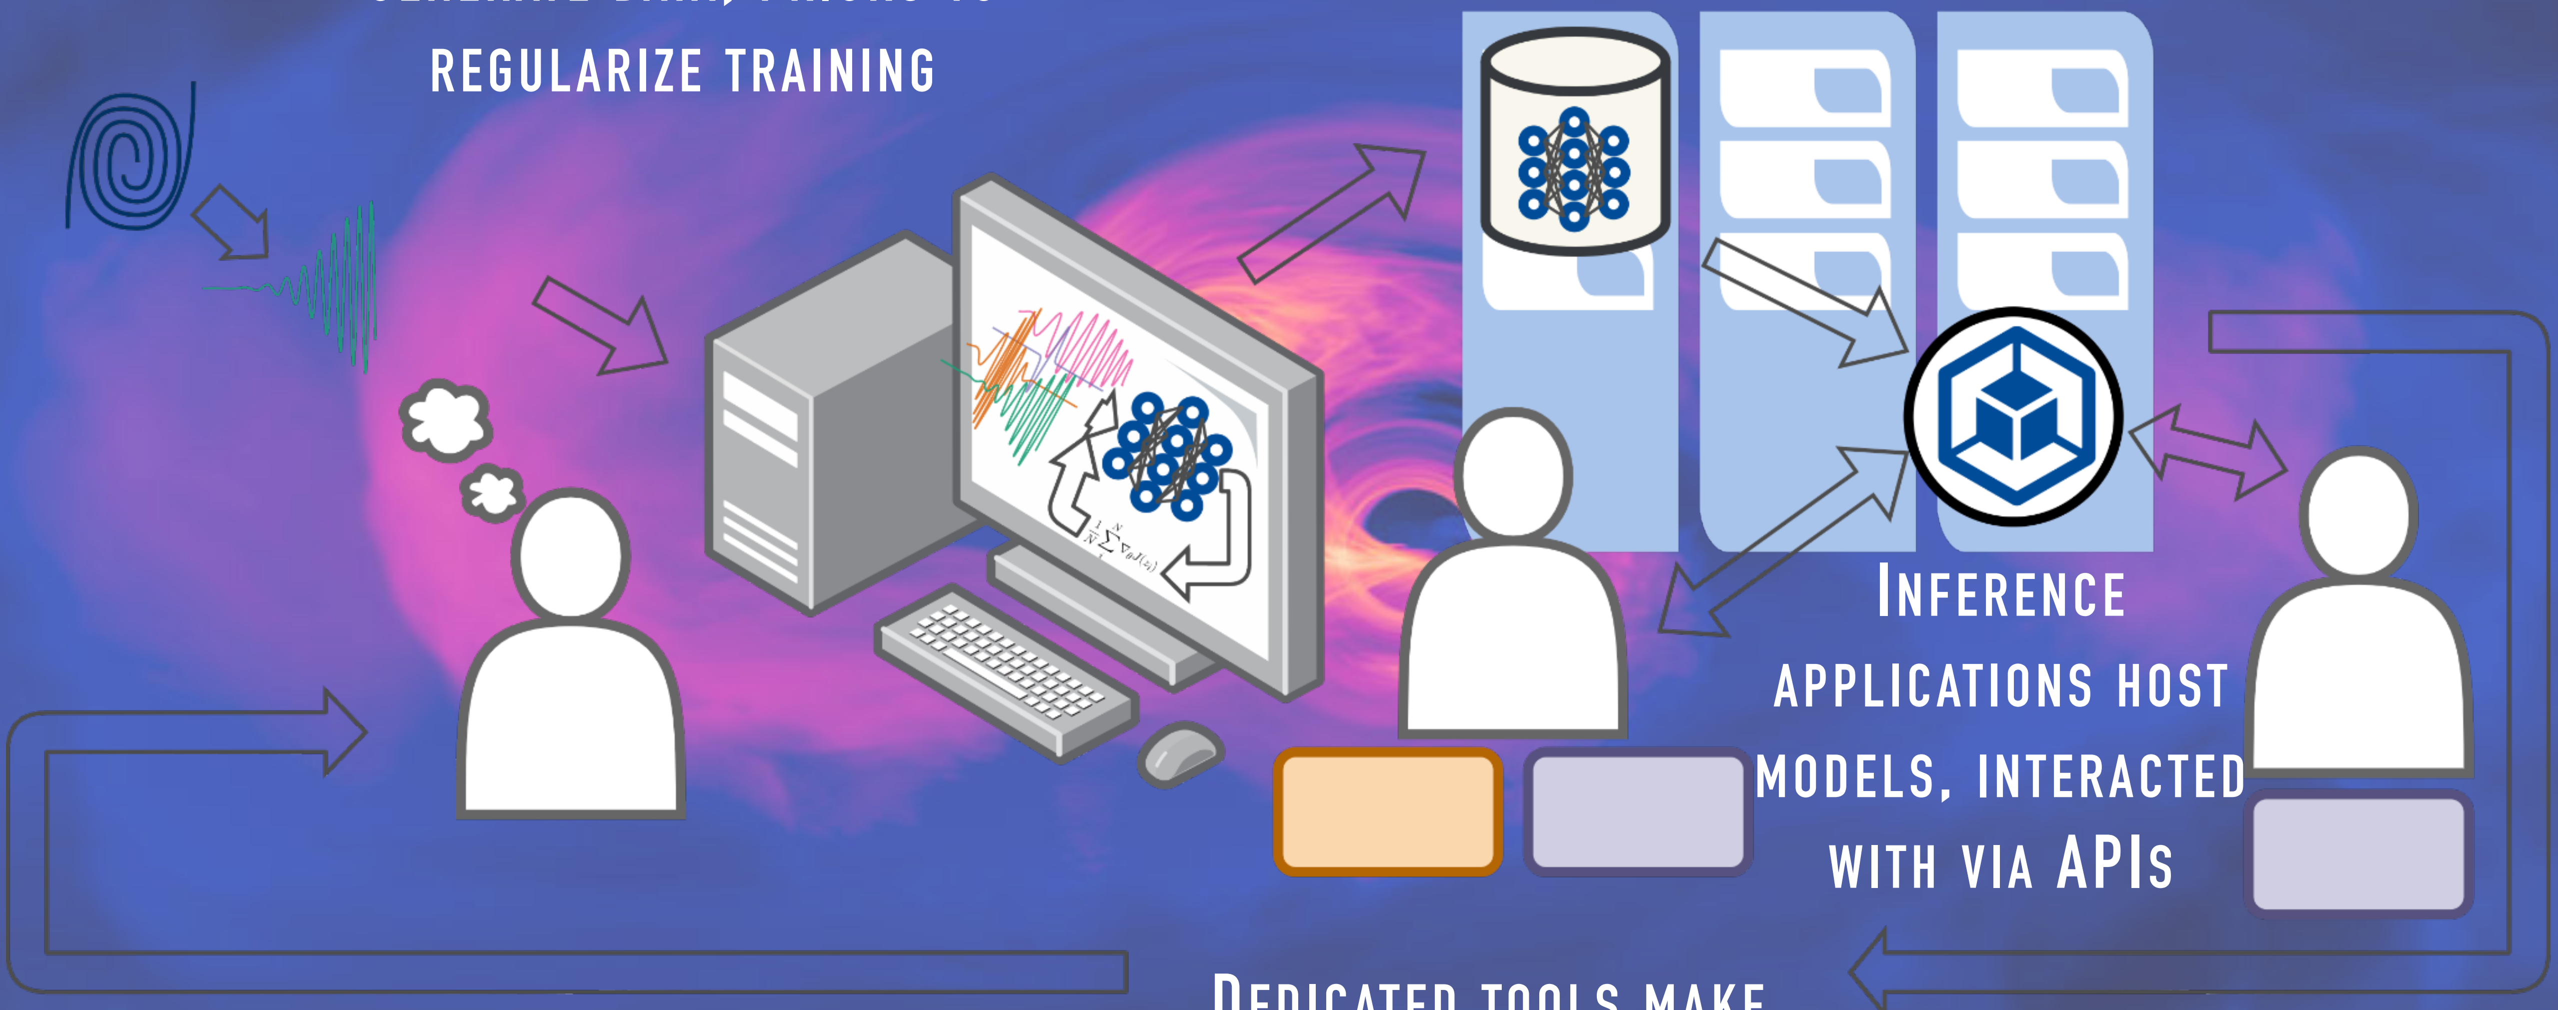

Scientist uses simulations to generate data, priors to regularize training

Models are distributed and versioned in centralized repositories

Inference applications host models, interacted with via APIs

Dedicated tools make iteration/exploration frictionless

Heterogeneous computing

Scalability

ml4gw & Hermes

ml4gw — torch utilities for training neural networks in gravitational wave physics applications

Fast data loading	GPU-friendly implementations of common analysis operations	Allowing for more robust use of simulations

hermes — a set of APIs for assisting in the acceleration, export, serving, and requesting of models using Triton Inference Server

Distribute models using centralized repos	Perform inference with an off-the-shelf application - NVIDIA Triton	Users interact via lightweight client APIs, abstracting implementation details

— Using those tools, we were able to run the algorithms on 64 GPU server seamlessly!

Smooth integration into online!

GraceDB (USER TESTING) Public Alerts Latest Search Notifications Pipelines Documentation Logout

Authenticated as: Katya Govorkova

G1783271
Neighbors
Log Messages
Full Event Log

G1783271

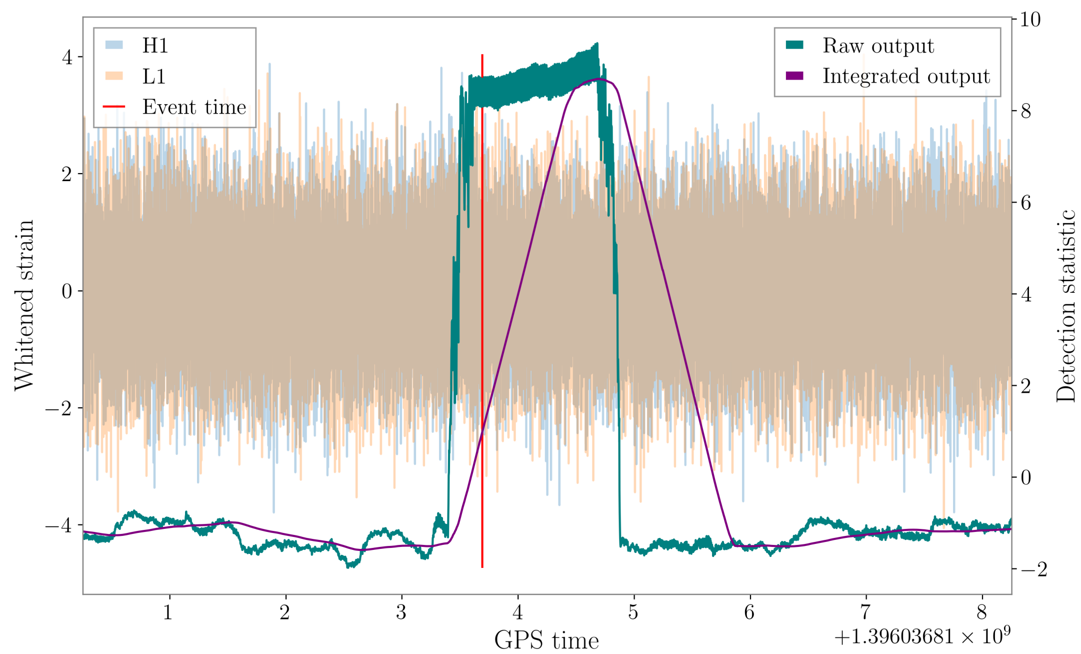

Basic Event Information

UID	G1783271
Labels	
Group	CBC
Pipeline	aframe
Search	AllSky
Instruments	['H1', 'L1']
Event Time	139[redacted]
FAR (Hz)	3.087e-08
FAR (yr^{-1})	1 per 1.0264 years
Latency (s)	3.524
Links	Data
Submitted	2024[redacted] UTC

SMOOTH INTEGRATION INTO ONLINE!

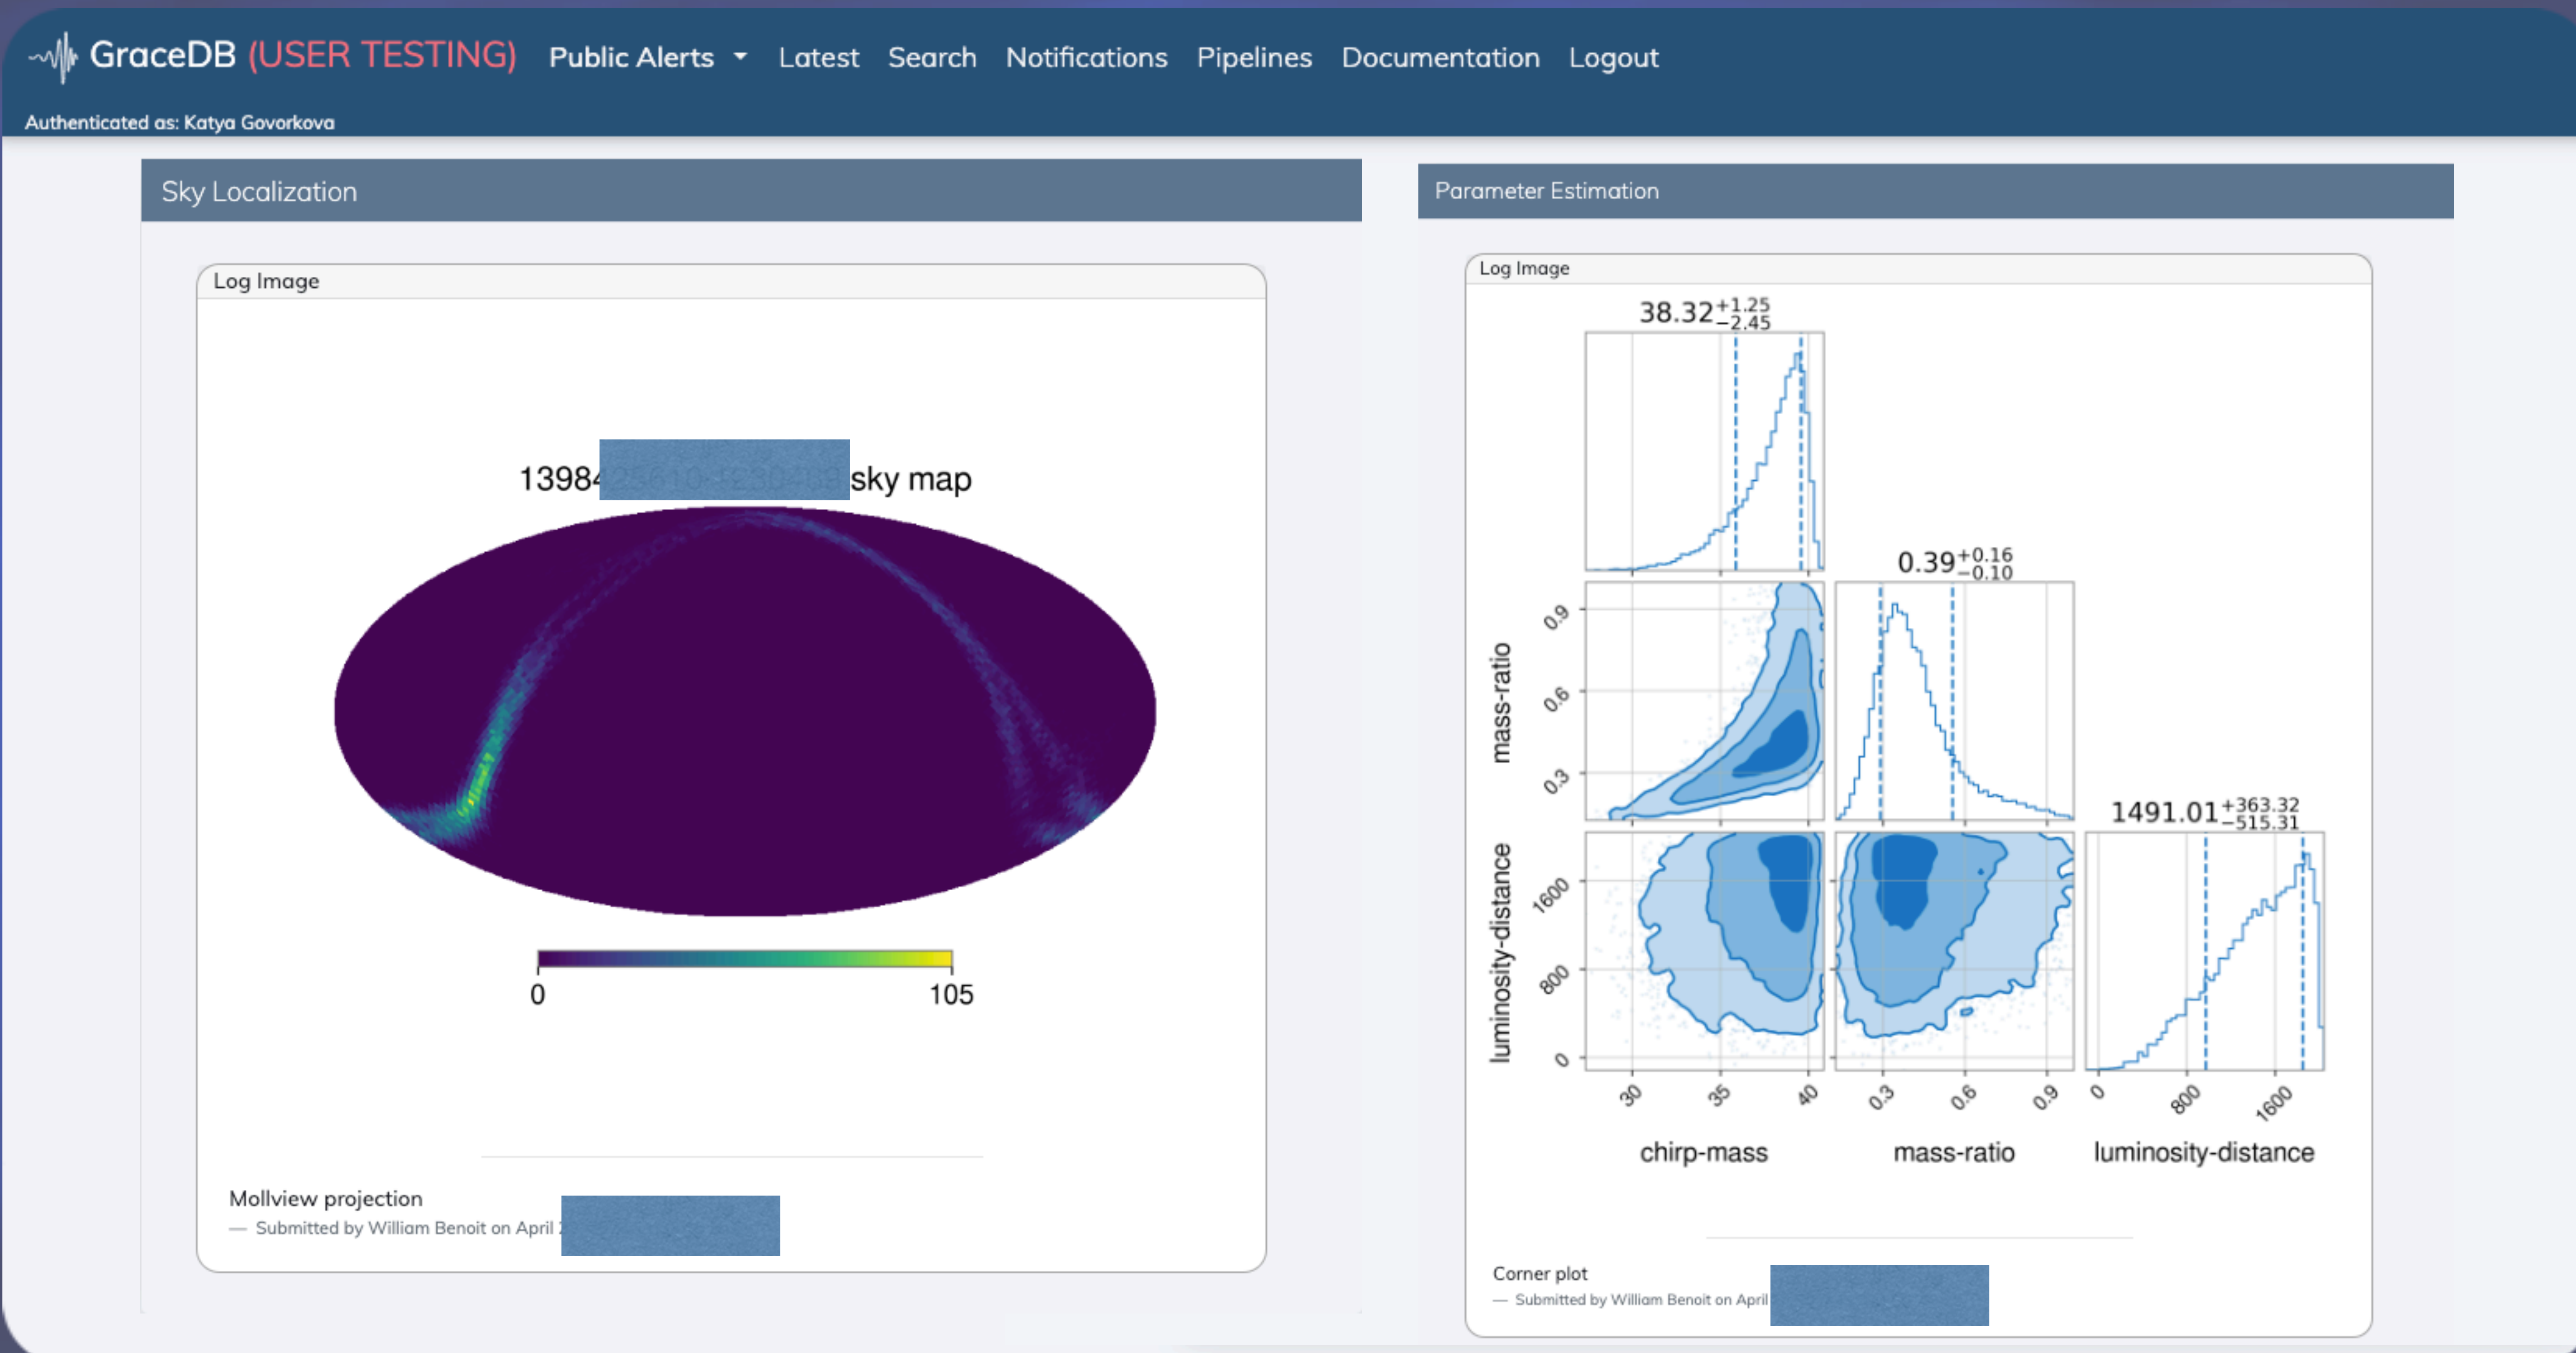

Conclusions

Nature Astronomy doi.org/10.1038/s41550-022-01651-w

To enable a complete AI pipeline, we have developed github.com/ML4GW — a set of comprehensive tools for ML pipeline in GW physics which allows to perform

- Modelled and unmodelled searches
- Run efficiently Offline
- Run Online with low latency
- Seamless development and fast deployment of NN-based algorithms
- Small computation footprint and optimised heterogeneity

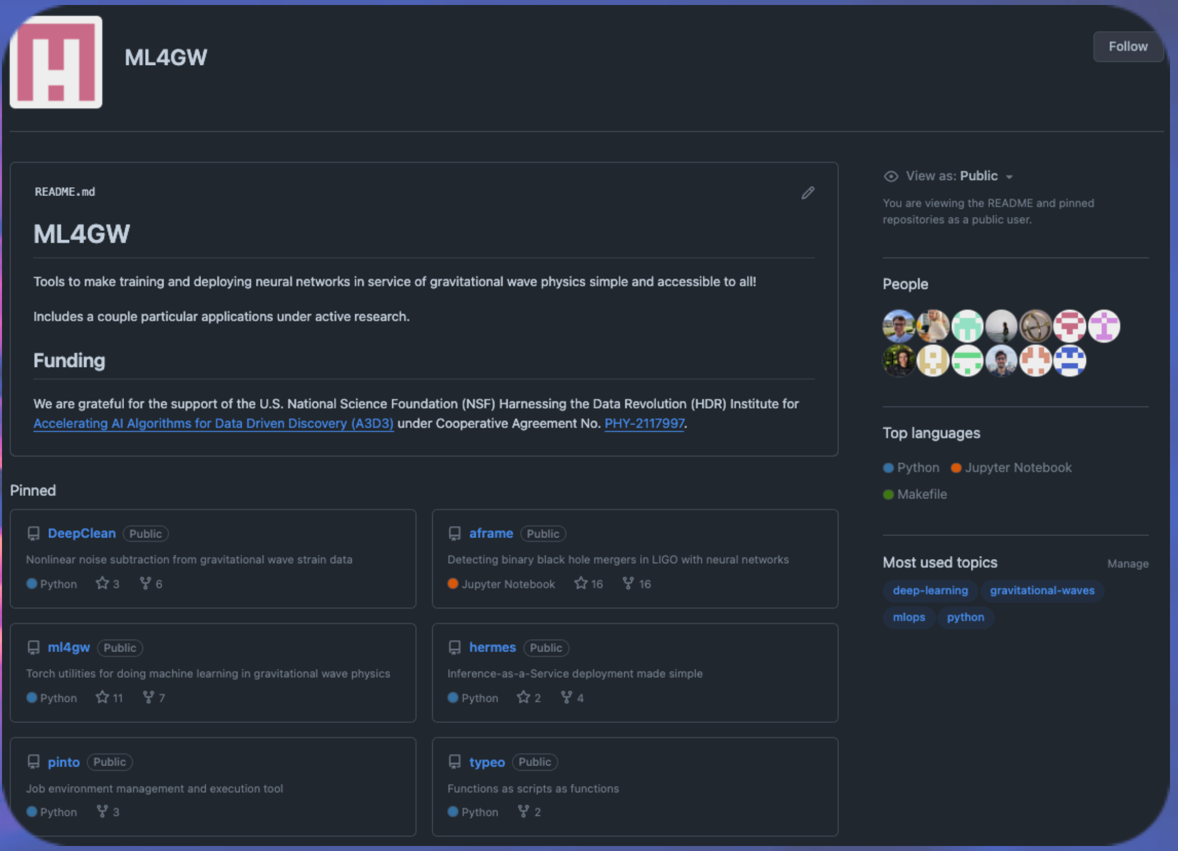

— Looking to invite many others to build on our work!

We run open weekly meetings and everyone is welcome to join

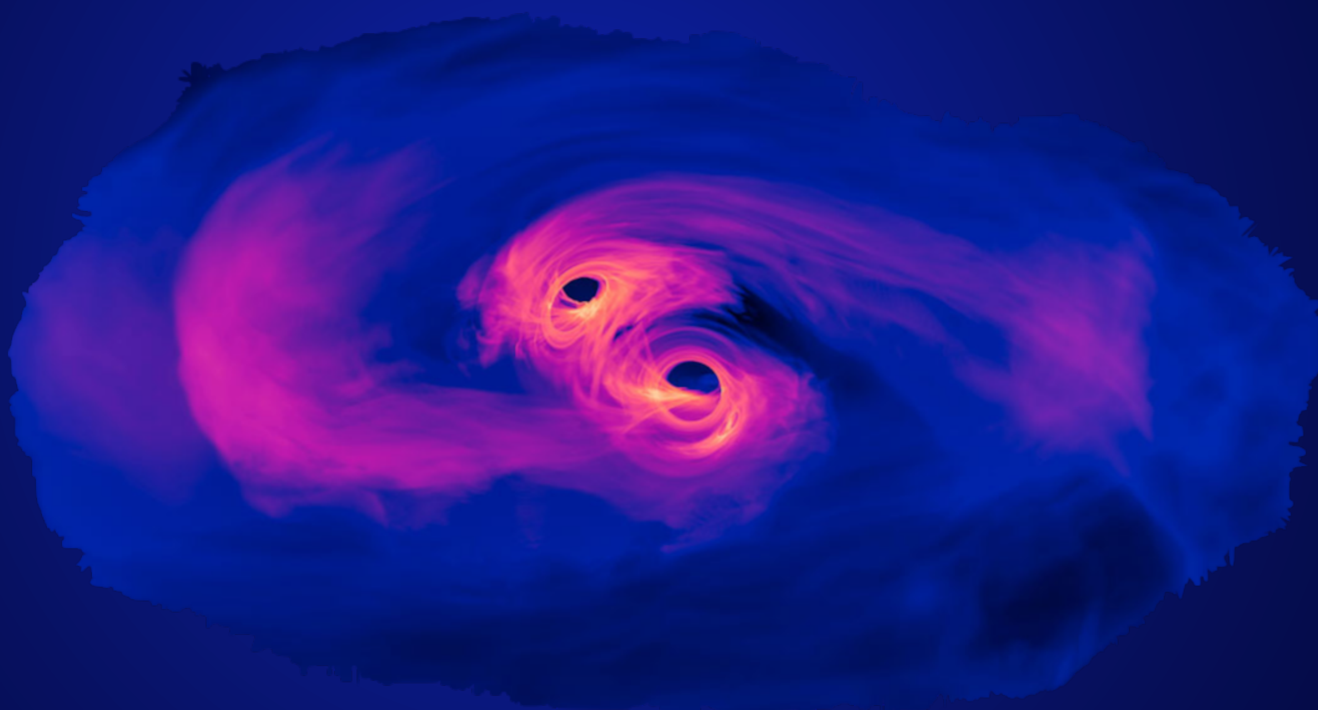

BACKUP

Gravitational-wave detector data

Continuous time series (1Hz, 128Hz … 16kHz)

Gravitational Wave channel
~20GB/day (per instrument)

Physical Environment Monitors
(seismometers, accelerometers, magnetometers, microphones etc)

Internal Engineering Monitors
(sensing, housekeeping, status etc)

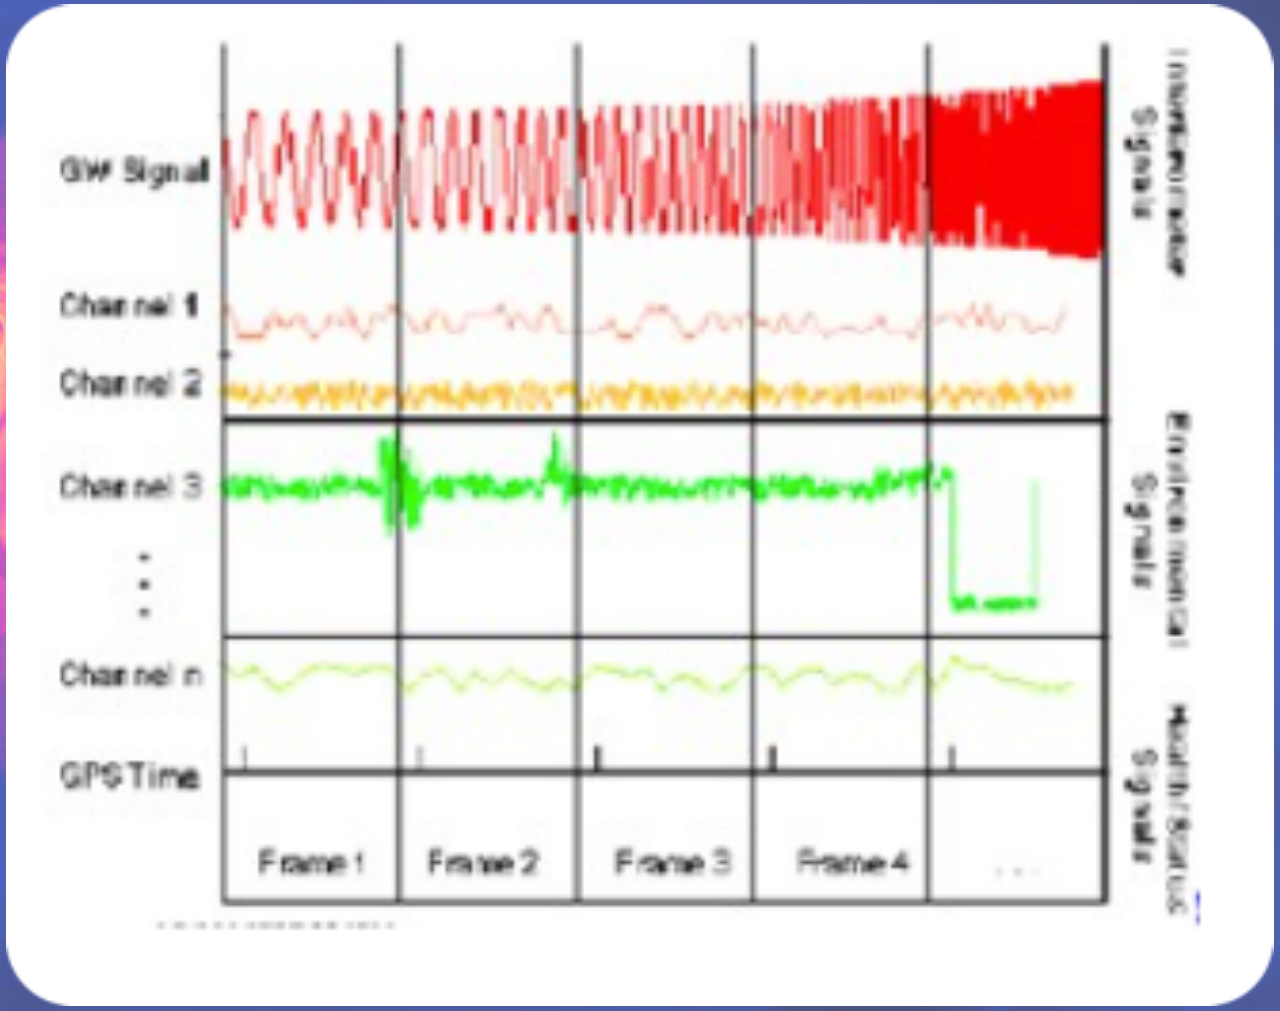

Together with various intermediate data products >2TB/day (per instrument)

A-Frame performance comparison

[Zenodo]

- Competitive performance on higher-mass catalog distributions
- Work remains to be done for lower masses – alternative architectures or smarter training techniques

$$V(\mathcal{F}) = \int \mathrm{d}\mathbf{x}\, \mathrm{d}\theta\ \epsilon(\mathcal{F}; \mathbf{x}, \theta)\phi(\mathbf{x}, \theta)$$

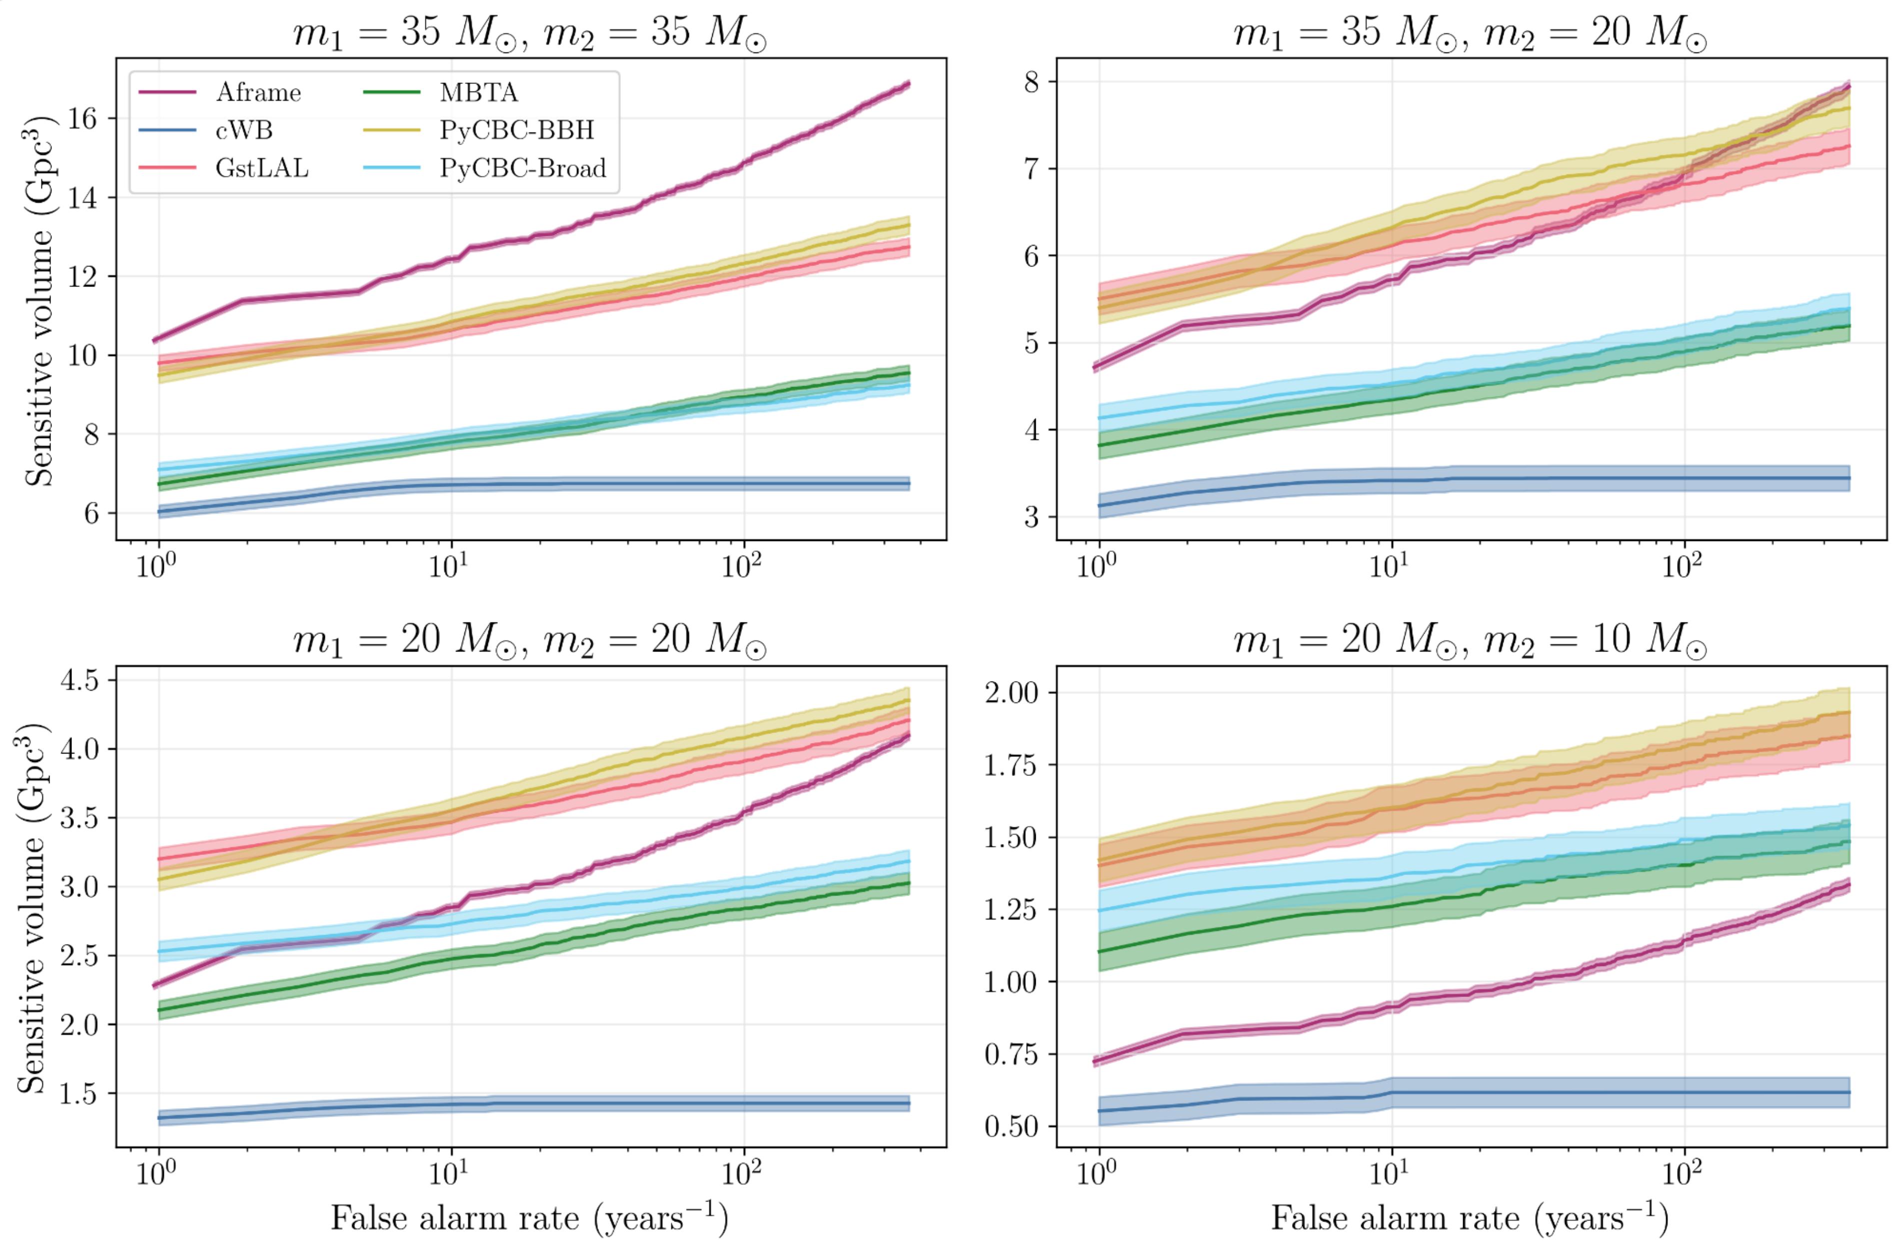

ARXIV2309.11537

We choose LSTM architecture to properly handle sequential data with temporal dependencies

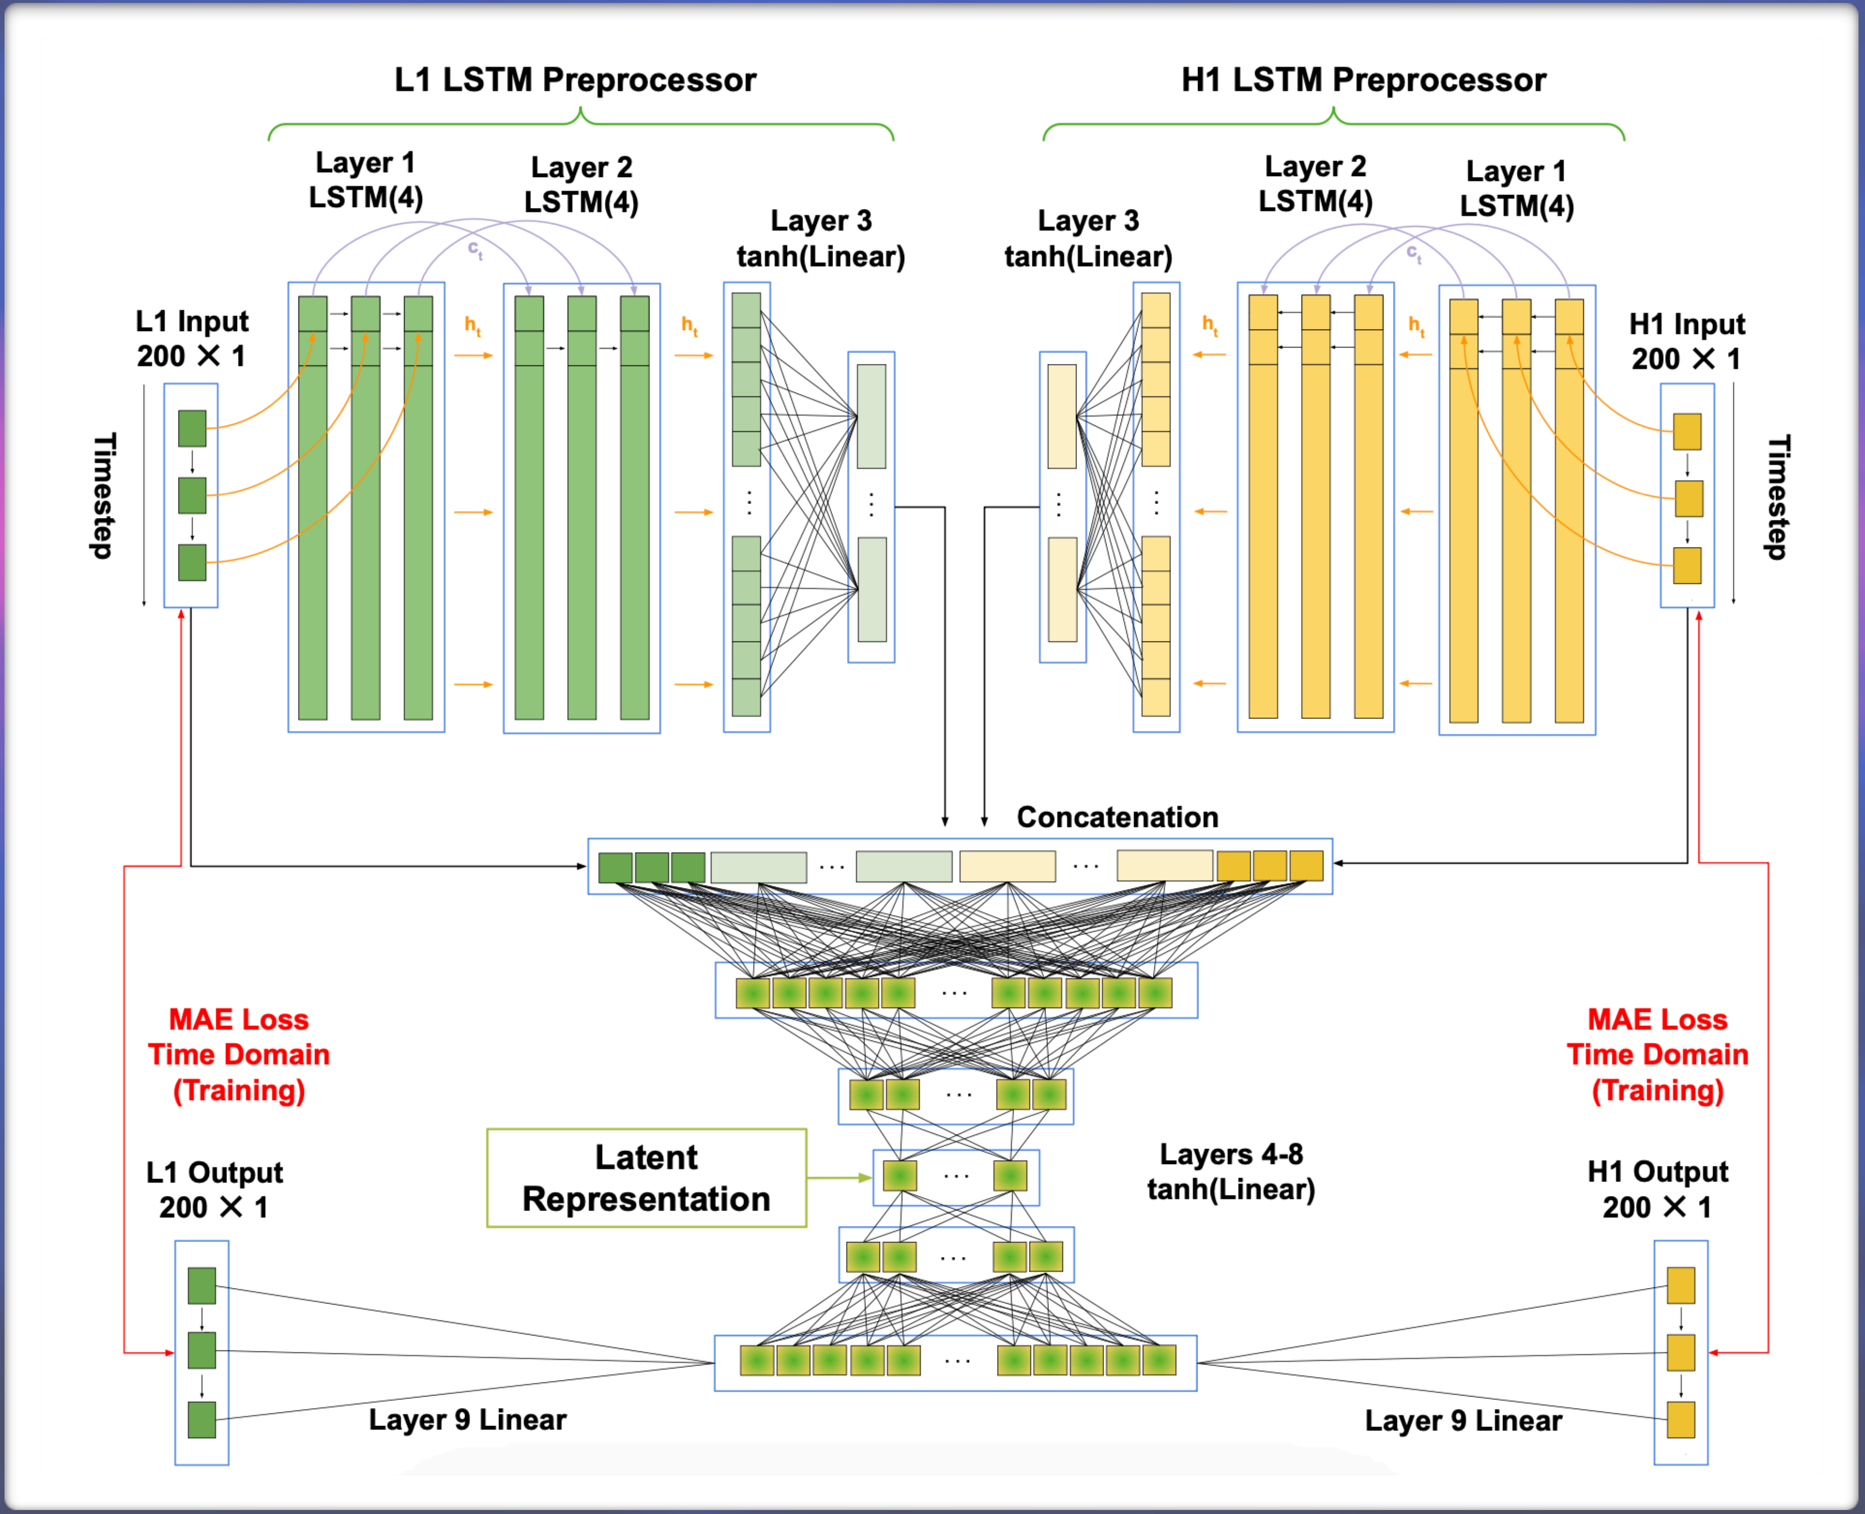

GWAK Dense

ARXIV2309.11537

We choose Dense architecture for backgrounds to properly handle sequential data without temporal dependencies

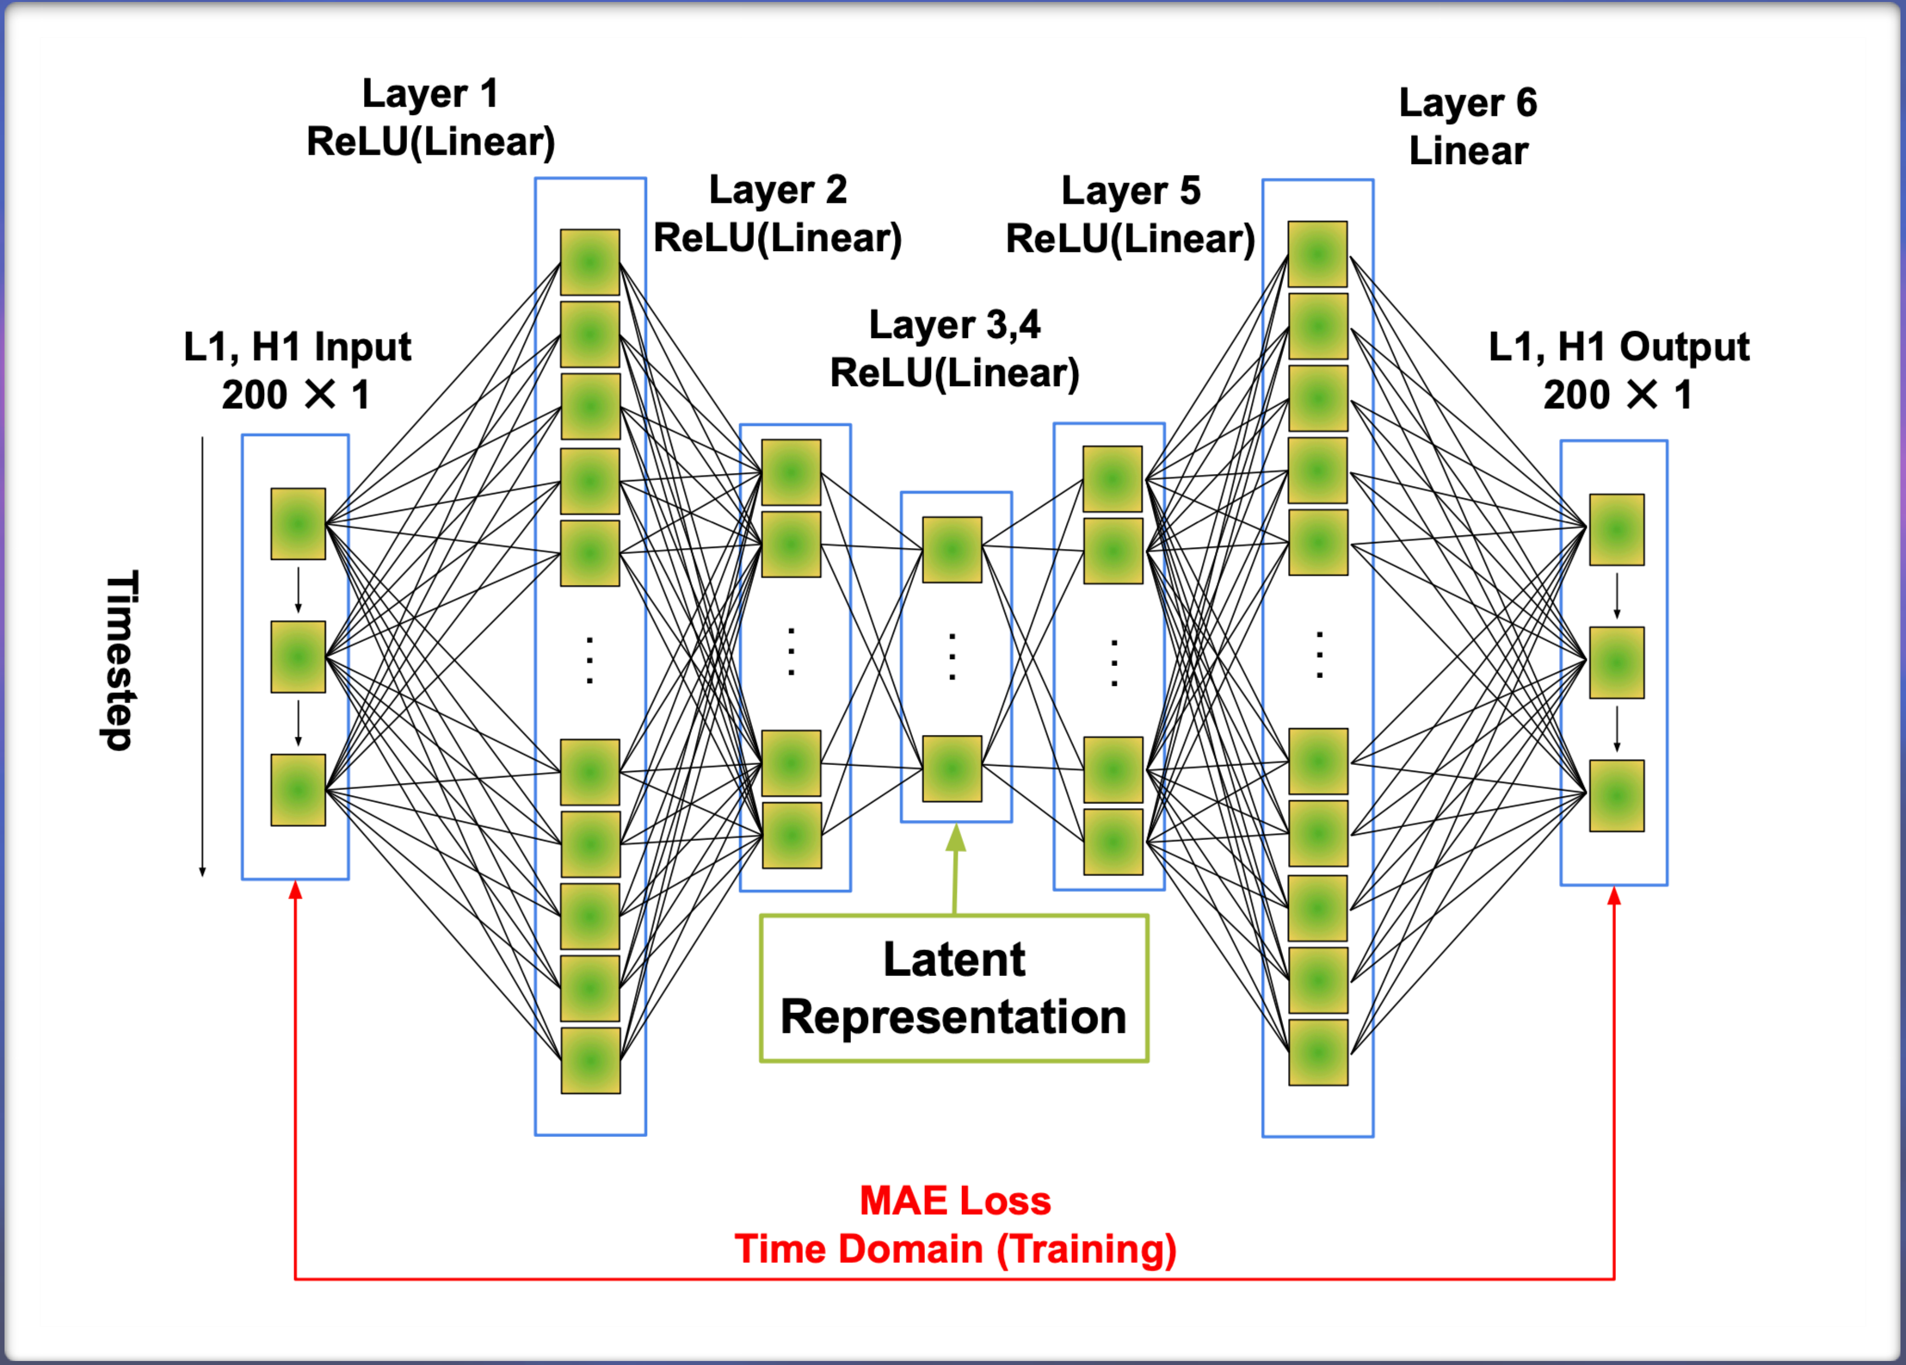

"Signal" dataset priors

arXiv2309.11537

Sampling parameters and priors for BBH (top) and sine-Gaussian (bottom) injections.

	Parameter	Prior	Limits	Units
BBH	m_1	-	$(5, 100)$	$M_\odot$
	m_2	-	$(5, 100)$	$M_\odot$
	Mass ratio q	Uniform	$(0.125, 1)$	-
	Chirp mass M_c	Uniform	$(25, 100)$	$M_\odot$
	Tilts $\theta_{1,2}$	Sine	$(0, \pi)$	rad.
	Phase ϕ	Uniform	$(0, 2\pi)$	rad.
	Right Ascension	Uniform	$(0, 2\pi)$	rad.
	Declination δ	Cosine	$(-\pi/2, \pi/2)$	rad.
sine-Gaussian	Q	Uniform	$(25, 75)$	-
	Frequency	Uniform	$(64, 512)$ and $(512, 1024)$	Hz
	Phase ϕ	Uniform	$(0, 2\pi)$	rad.
	Eccentricity	Uniform	$(0, 0.01)$	-
	Declination δ	Cosine	$(-\pi/2, \pi/2)$	rad.
	Right Ascension	Uniform	$(0, 2\pi)$	rad.
	Ψ	Uniform	$(0, 2\pi)$	rad.

ARXIV2309.11537

Example of GWAK classes: glitch and background strains

The light blue shading highlights an example region that is passed as input to the autoencoders for training

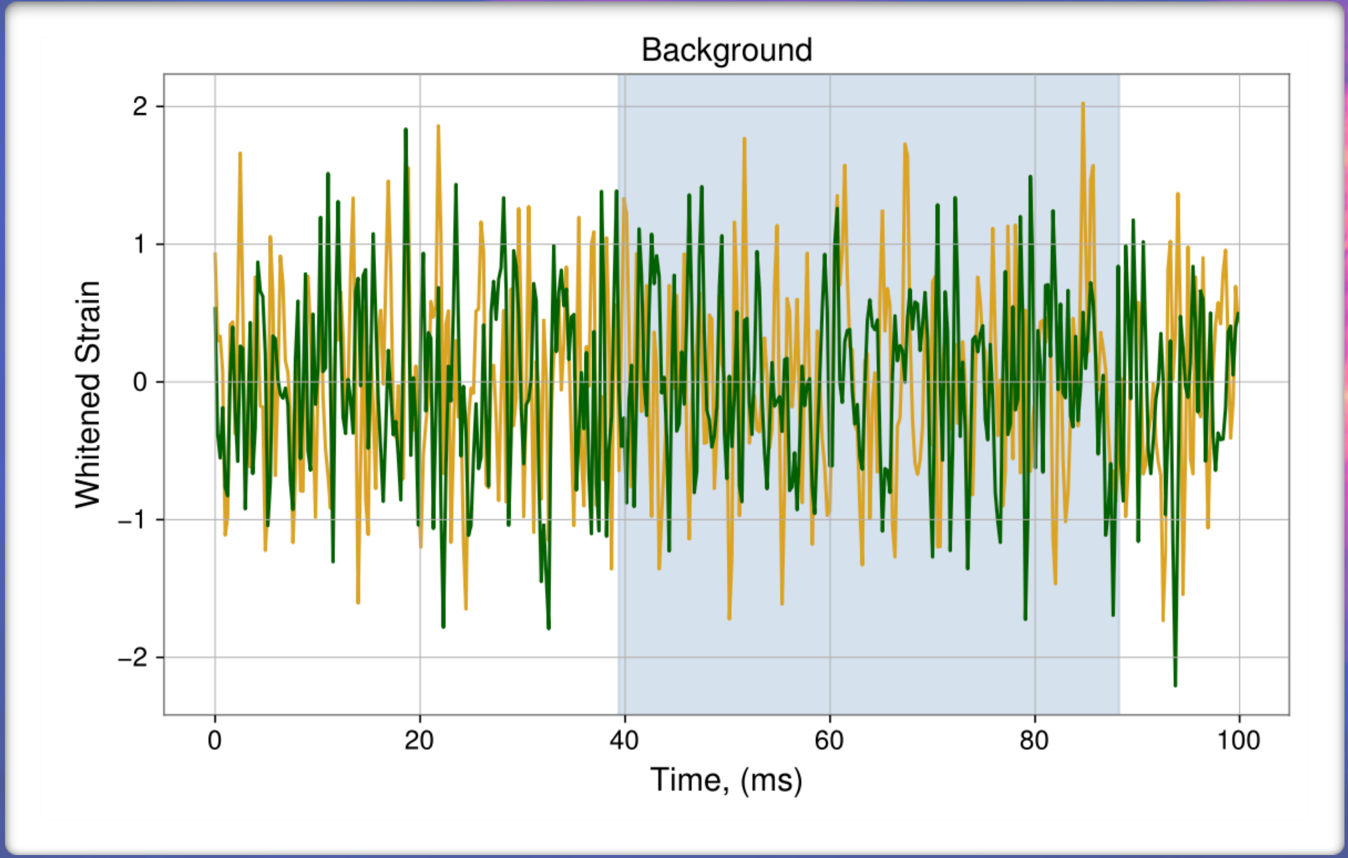

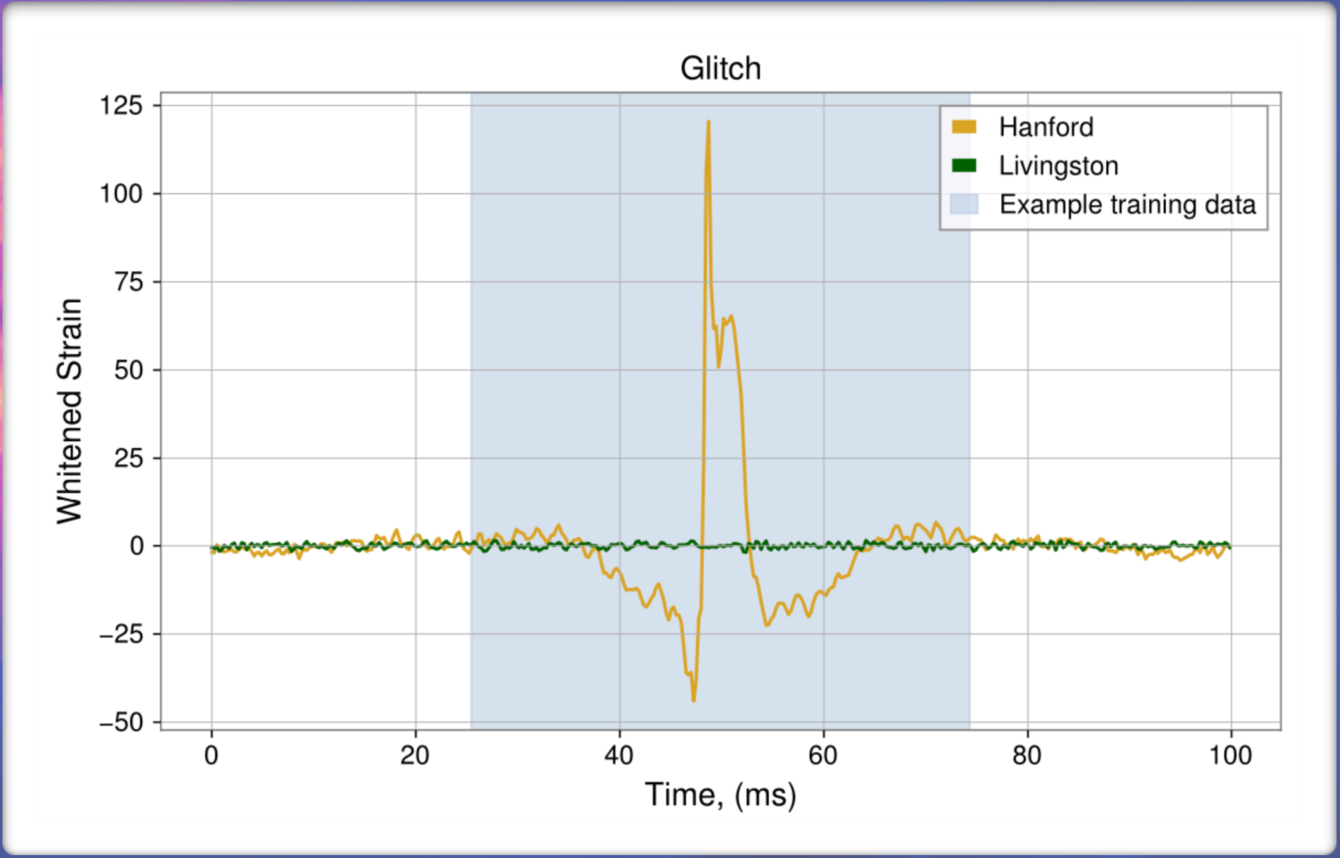

Signal datasets

ARXIV2309.11537

Example of signal-like classes: BBH and sine-Gaussian strains from Livingston and Hanford

The light blue shading highlights an example region that is passed as input to the autoencoders for training

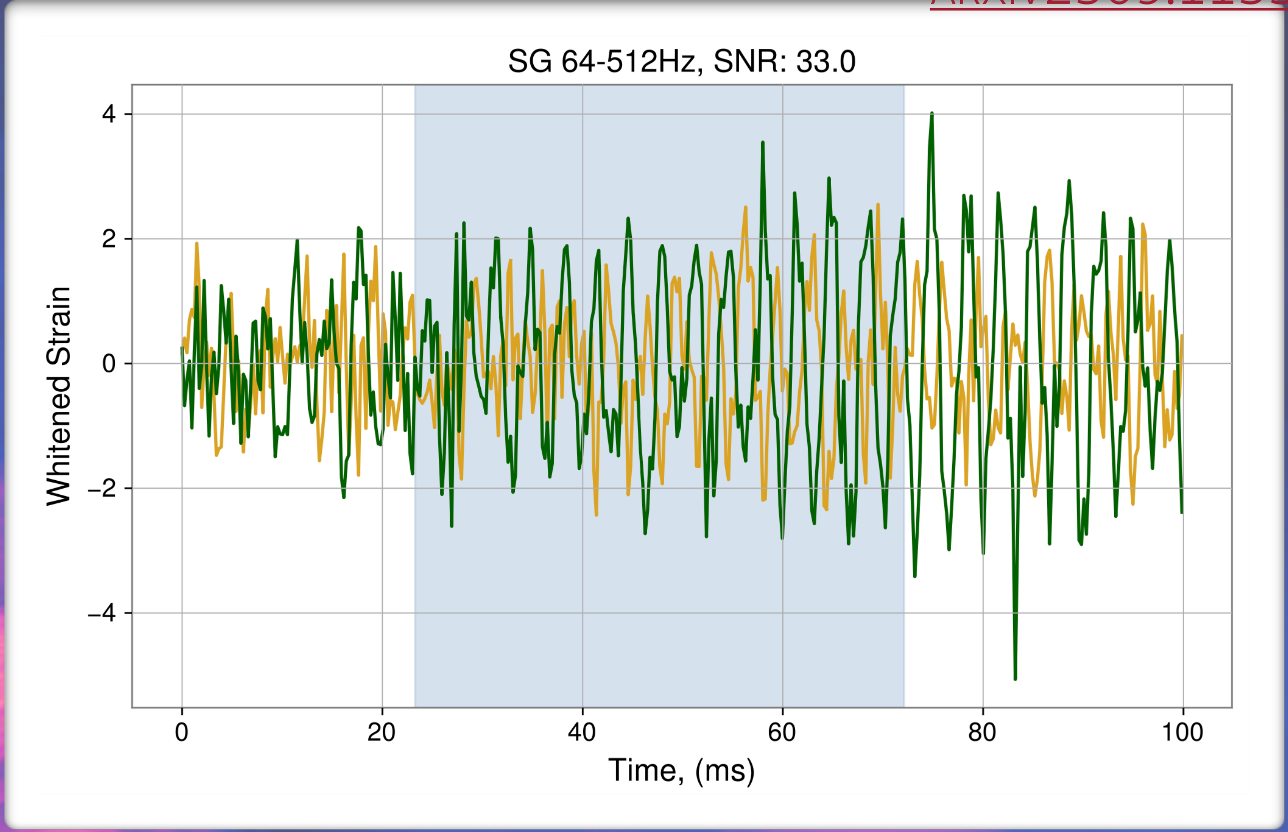

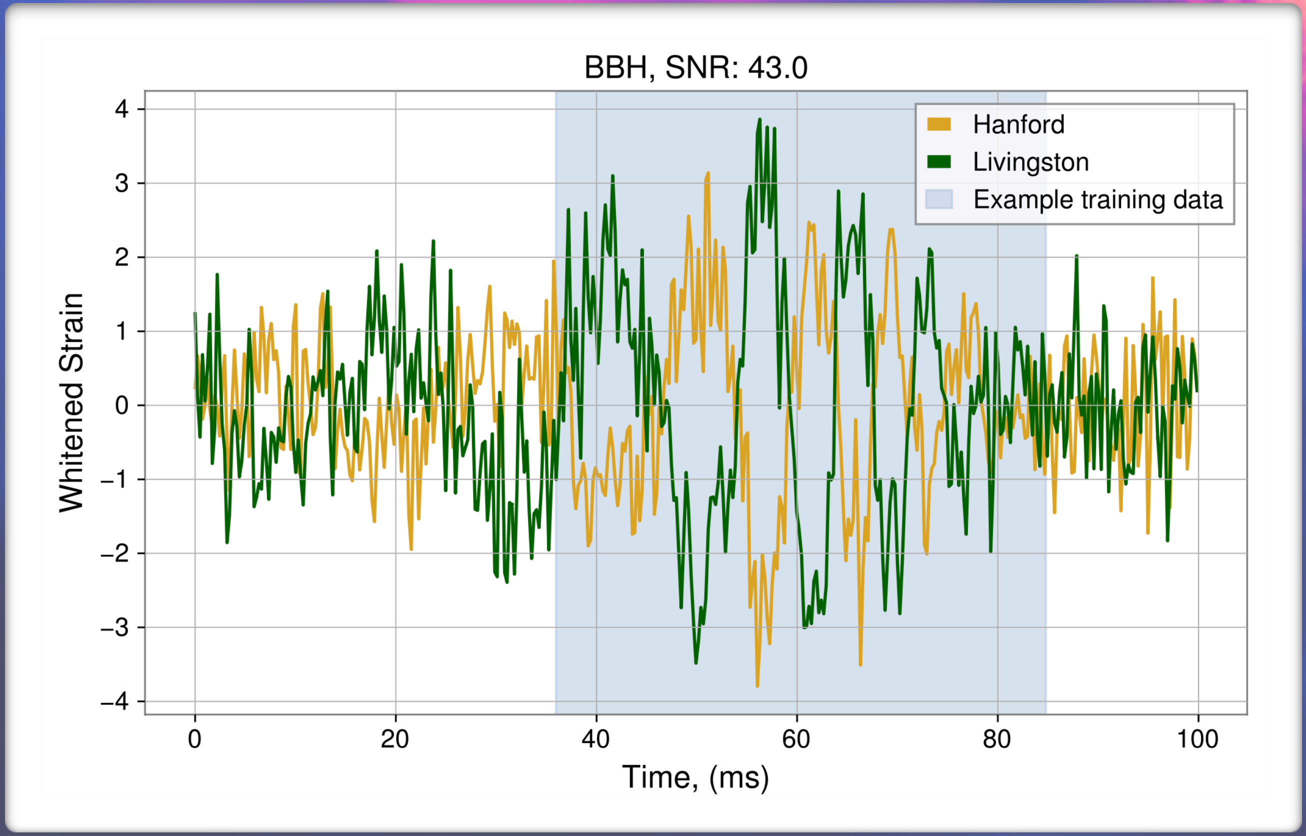

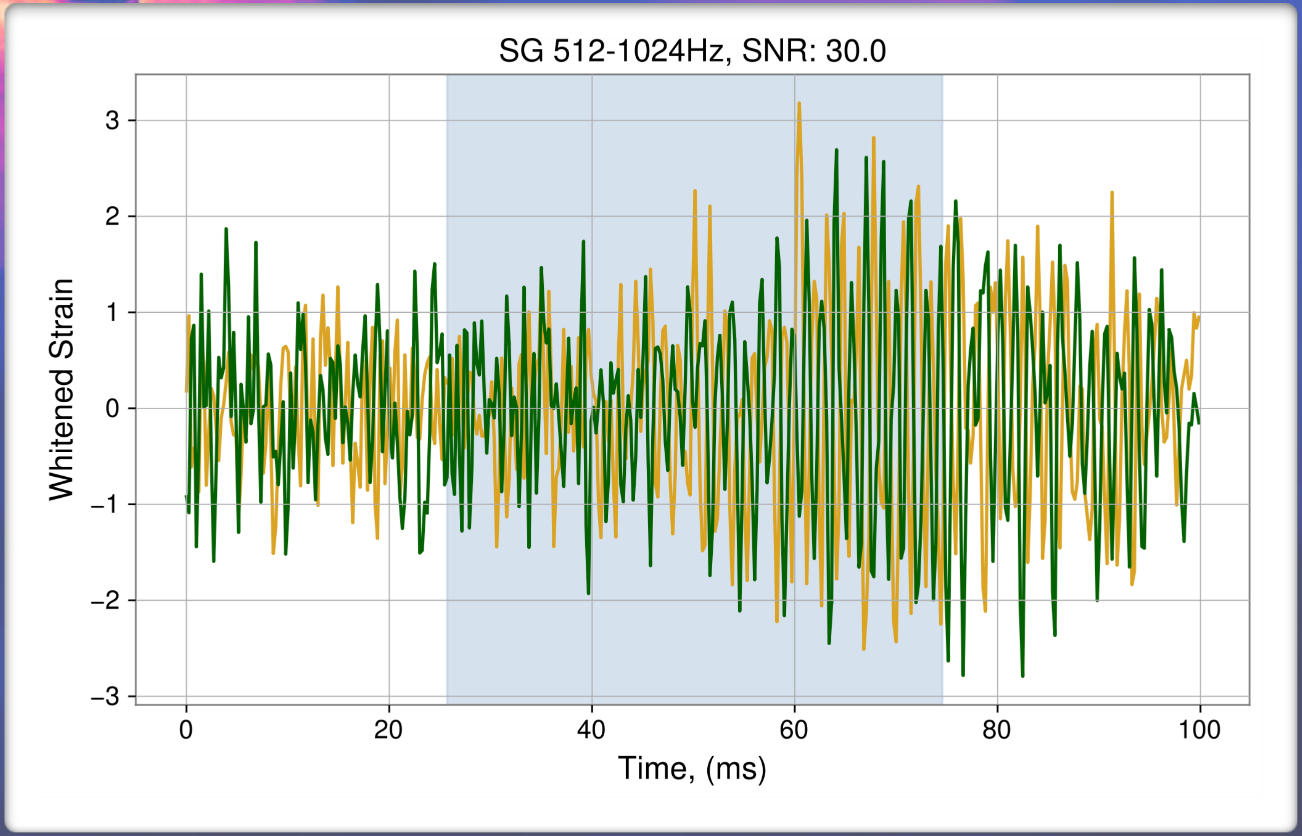

ARXIV2309.11537

EXAMPLE OF SIGNAL-LIKE CLASSES: SUPERNOVA AND WHITE NOISE BURST STRAINS FROM LIVINGSTON AND HANFORD

THOSE ANOMALIES ARE NOT USED TO CREATE THE GWAK

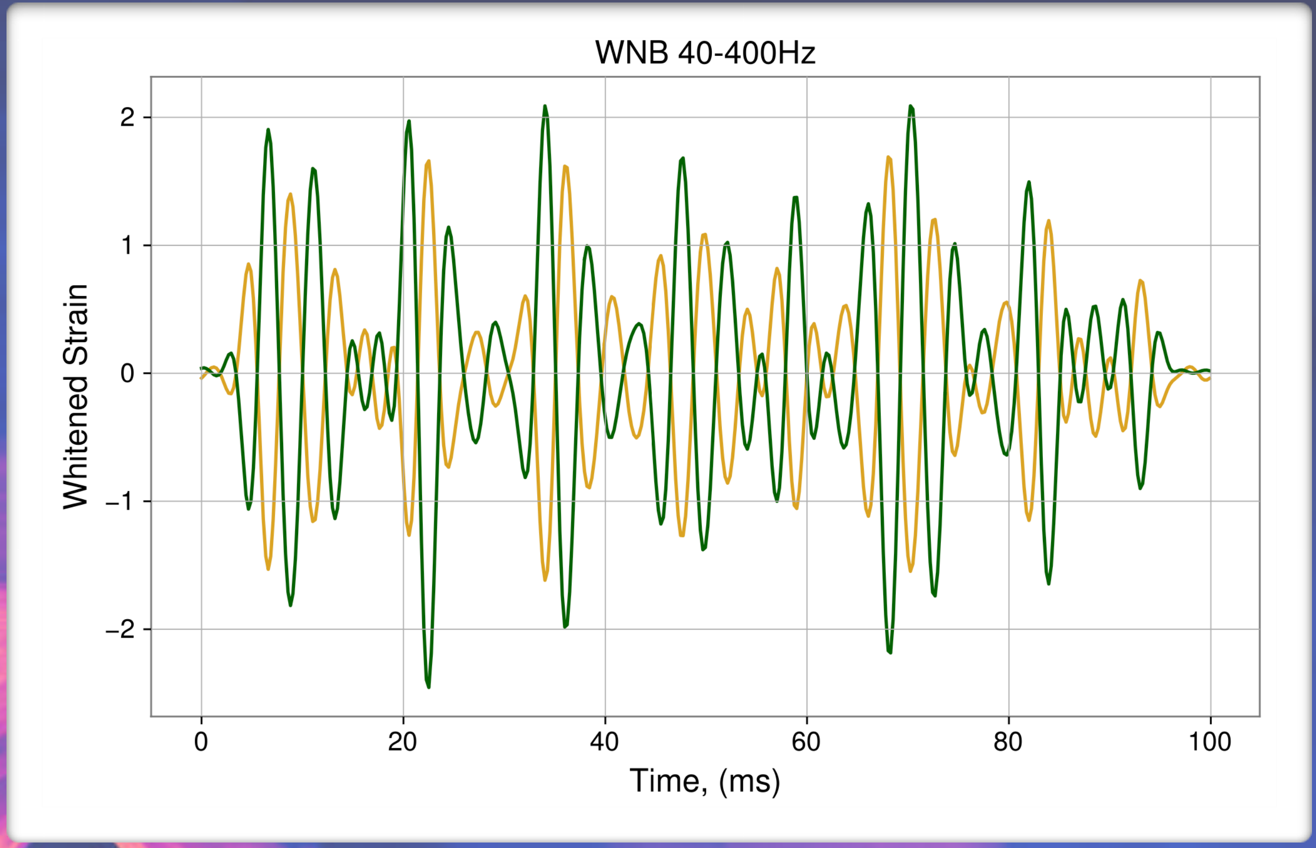

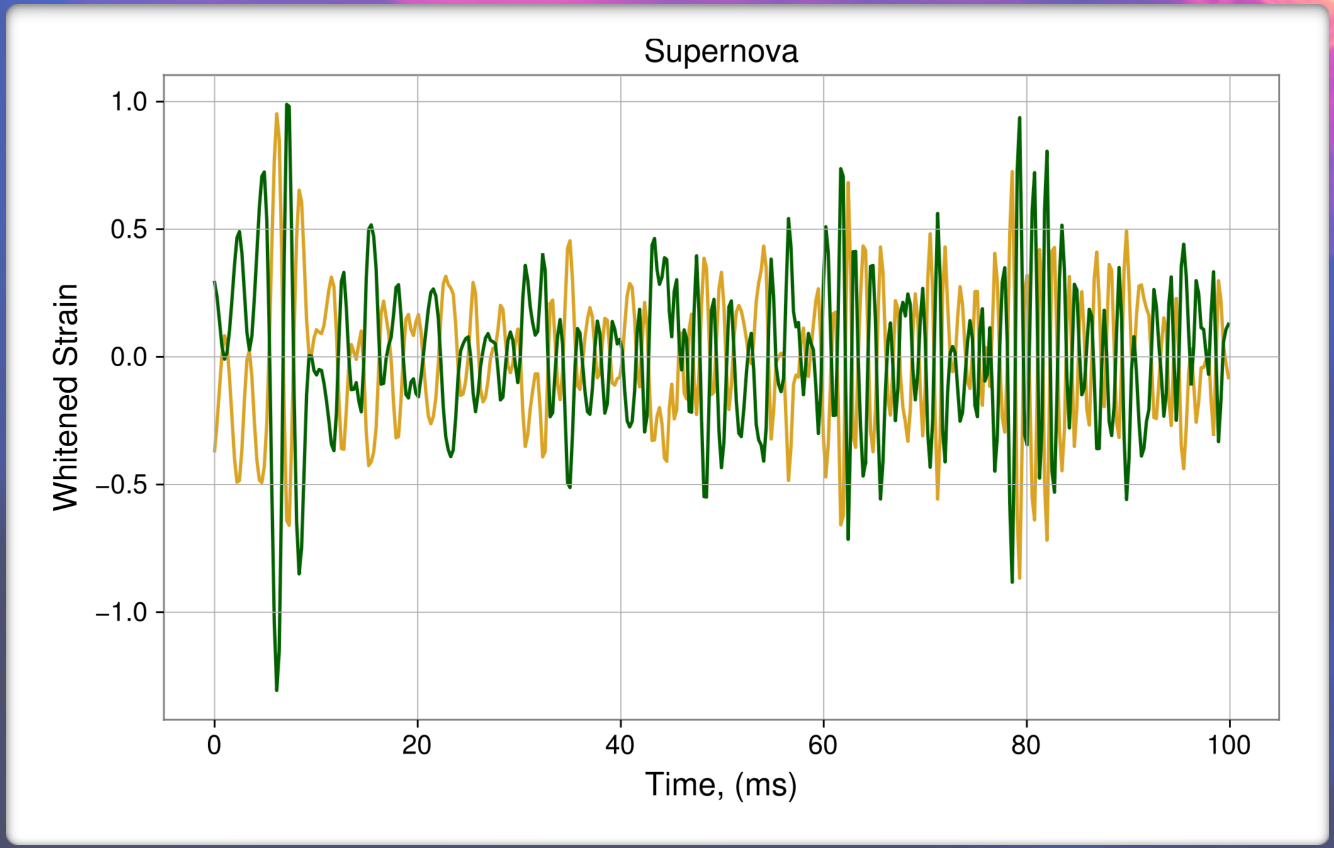

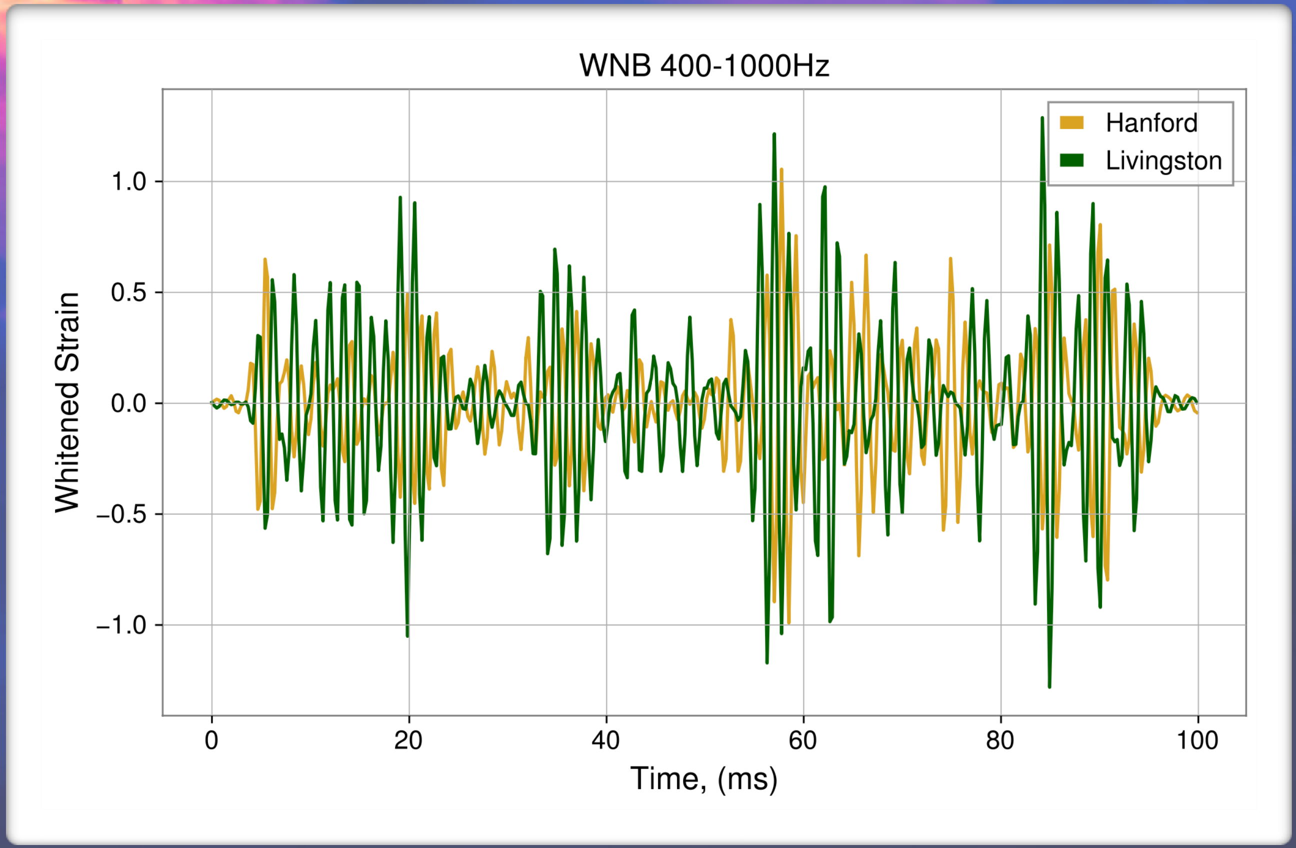

Autoencoder recreations

arxiv2309.11537

Example of recreation on injected BBH signal, with the noise-less template shown as well

The recreation of the BBH autoencoder follows closely the original signal injection

While background, glitches, SG 64-512 Hz and SG 512-1024 Hz fail to reconstruct the injected BBH signal

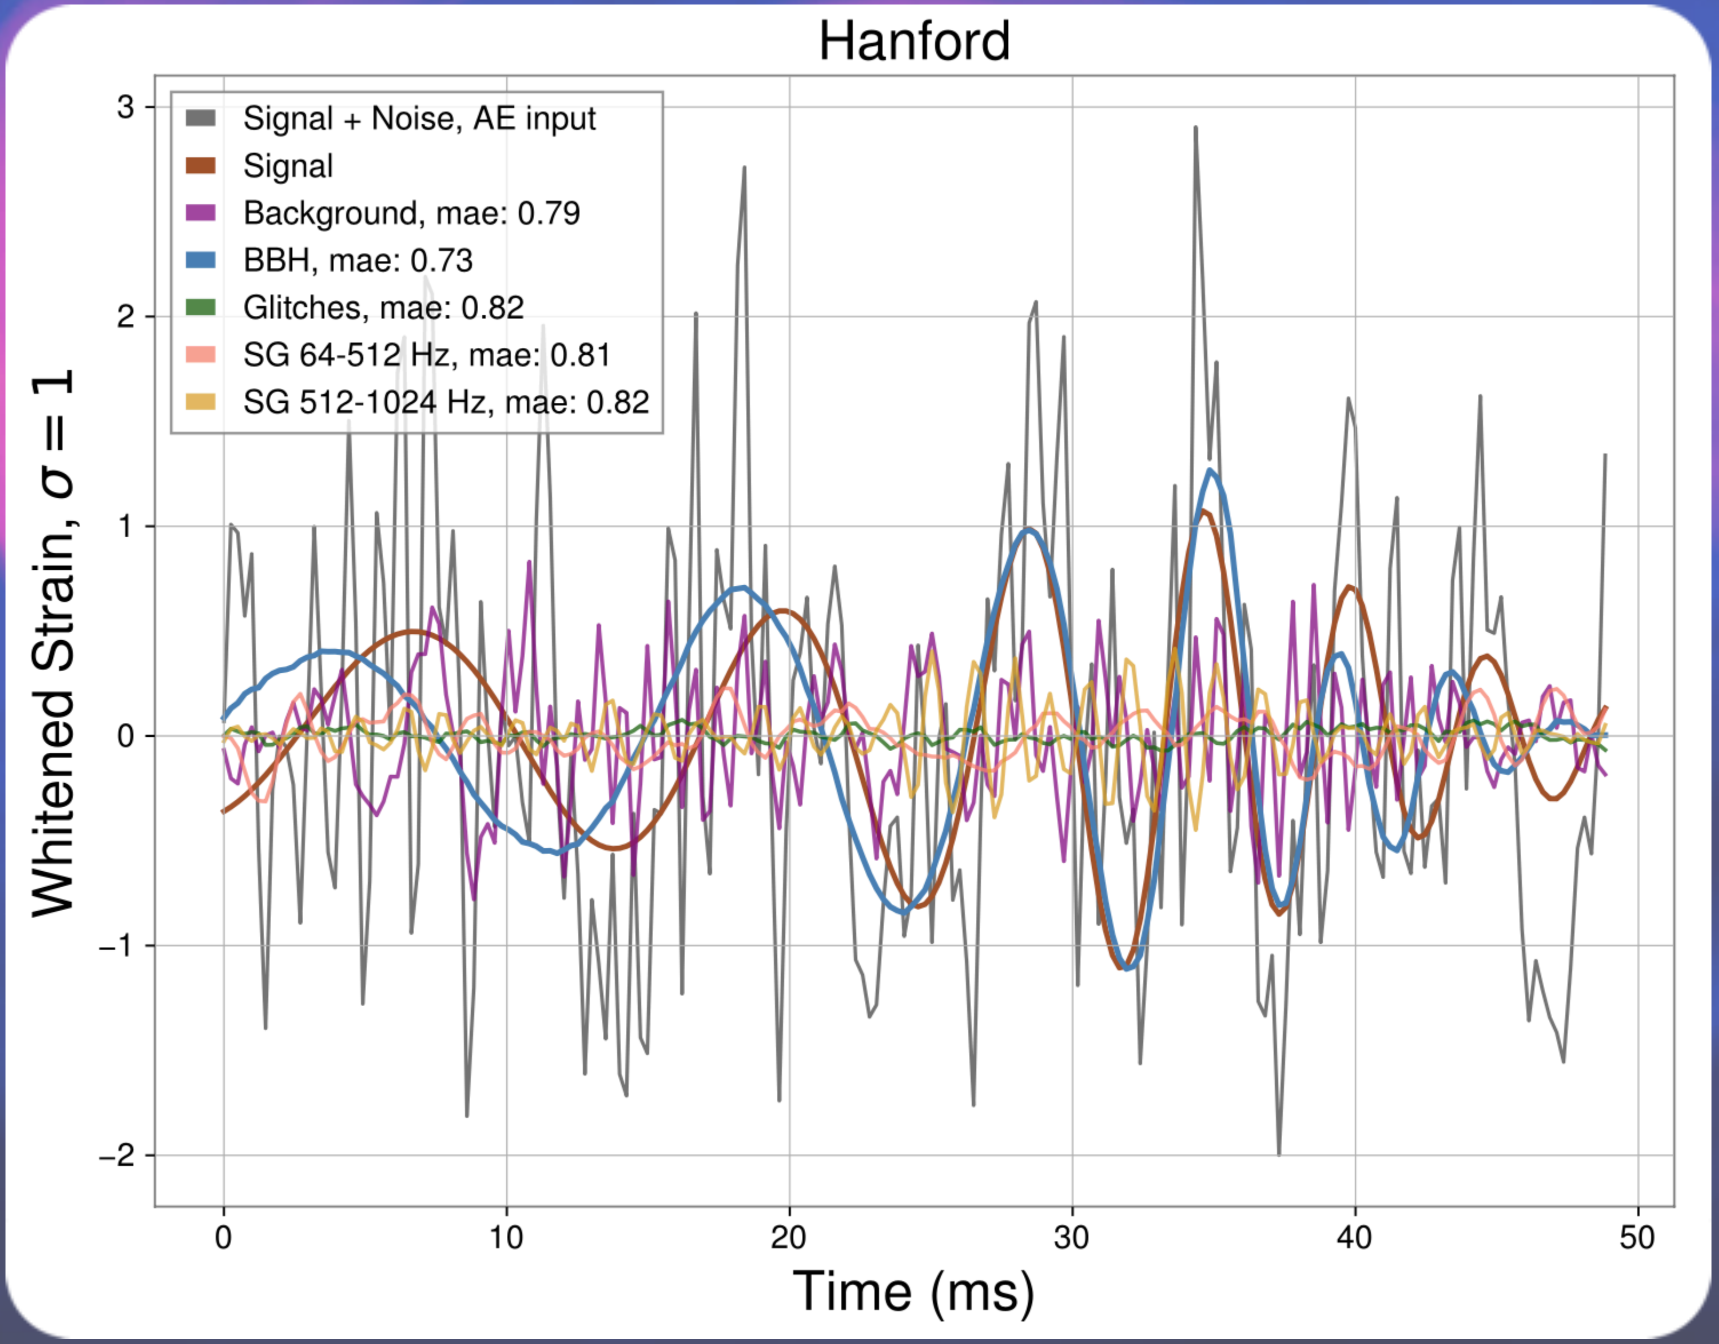

The GWAK Efficiency

arxiv2309.11537

The final metric as a function of SNR for GWAK axes training signals, BBH, SG 64-512 Hz, SG 512-1024 Hz and for potential anomalies, WNB 40-400 Hz, WNB 400-1000 Hz, and Supernova