“AI” and gravitational waves

[Gemini]

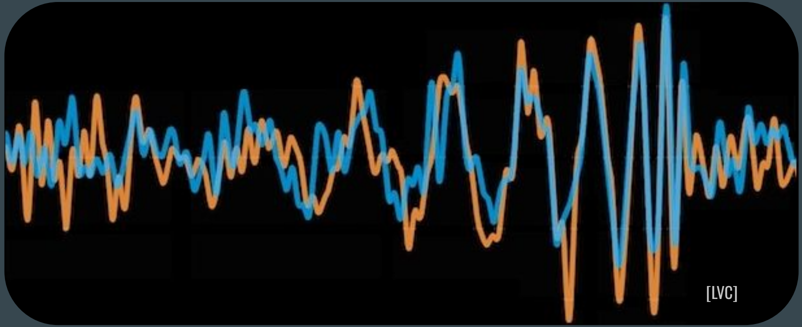

[Gemini]

David Keitel (Universitat de les Illes Balears)

UIB Universitat de les Illes Balears

IAC3 Institut d’Aplicacions Computacionals de Codi Comunitari

EuCAIFCon 2025,
Cagliari, 2025-06-16

LIGO-G2501203-v3

Gravitational Waves

- 100 years from Einstein's prediction to the first detection: 14 September 2015 – 10-year anniversary this September!
- time-varying mass quadrupoles
  → propagating ripples in spacetime (GWs)
- gravity is a weak force, spacetime a "very stiff fabric"
  → need extreme astrophysical sources
  - compact binary coalescences (CBCs)
  - supernovae and other cataclysmic explosions
  - deformed rotating compact objects
  - early universe physics
- A whole new spectrum for observational astronomy!

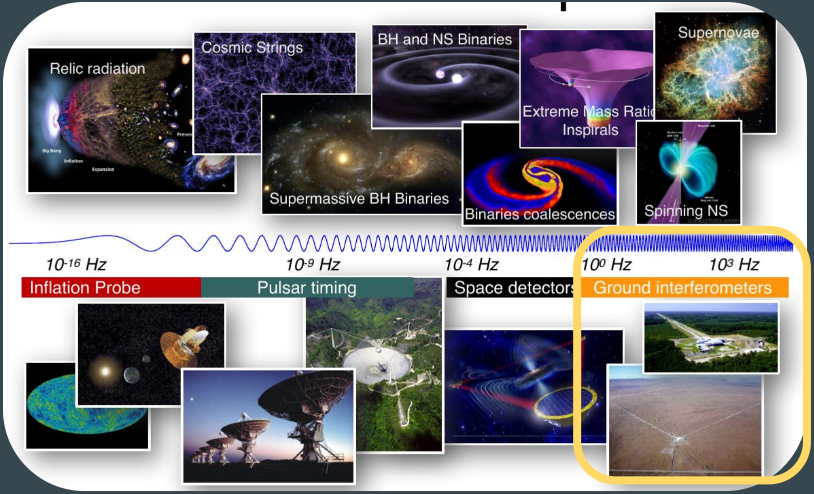

"AI" and gravitational waves (?)

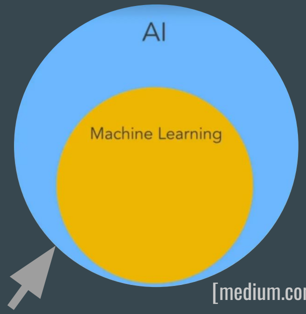

[medium.com]

I don't like this!

- Let's start with an exercise of "How to Lose Friends & Alienate People":

I don't think there have really been any applications of "Artificial INTELLIGENCE" to GWs yet!

- We have seen plenty of great applications of *new data analysis tools* which are *inspired by* the artificial intelligence & machine learning community.
- But none of these really make *decisions* on their own, or replace parts of the human researchers' critical role in designing, training and evaluating their methods.
- In this talk, I'll more often say "ML", for machine learning, rather than "AI".
- Still, let's get to all that cool stuff!

"AI"/ML and gravitational waves

LOADS of great contributed talks at this conference – check them out!

1. State of GW astrophysics
2. "AI"/ML applications:

data analysis

   i. Detector design, operation and characterisation
   ii. Compact binary coalescences
   iii. Bursts
   iv. Continuous waves
   v. Stochastic backgrounds

focus: current LVK detector network

   vi. Beyond LVK: future detectors

for each of these:

- state of the art
- example "AI"/ML highlights *(personal selection, not necessarily representative)*
- open problems and future directions

“AI”/ML and gravitational waves

The field greatly profited from EU COST Action CA17137 “g2net”! g2net.eu

for more details:

Home › Living Reviews in Relativity › Article

Applications of machine learning in gravitational-wave research with current interferometric detectors

Review Article | Open access | Published: 27 February 2025

Volume 28, article number 2, (2025) Cite this article

Elena Cuoco, Marco Cavaglià, Ik Siong Heng, David Keitel & Christopher Messenger, doi:10.1007/s41114-024-00055-8

…but a lot more cool stuff has already been published since we stopped collecting references in 2024!

The LVK detector network and collaboration

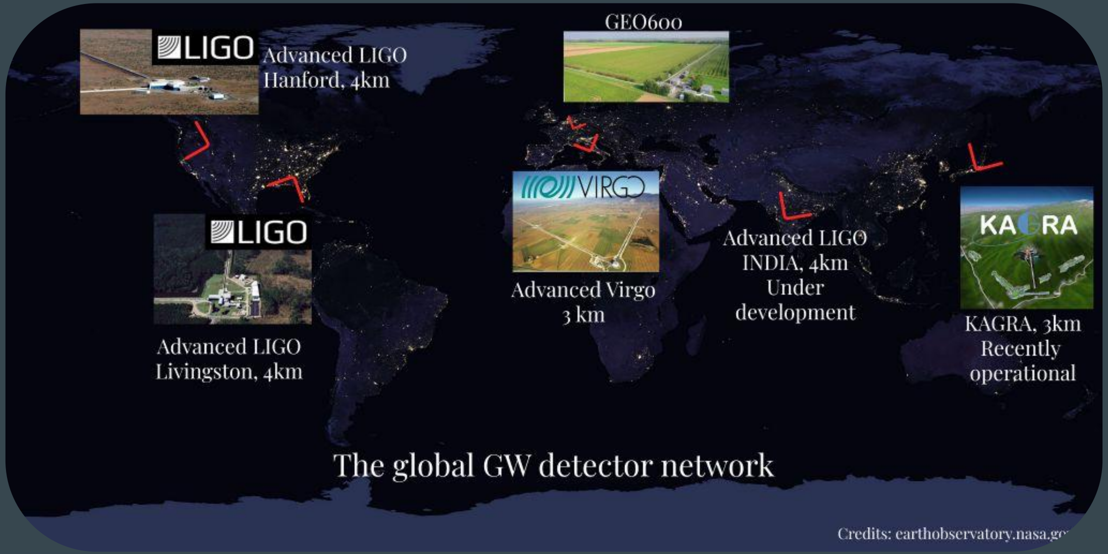

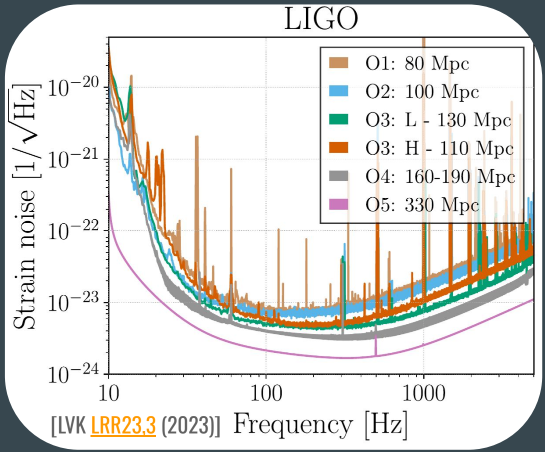

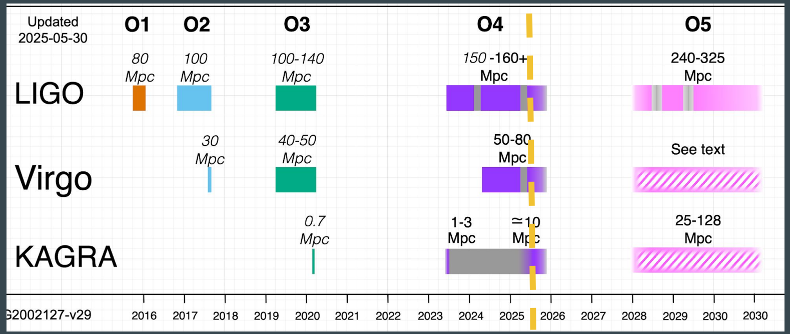

- >2500 scientists from >200 groups on 5 continents
- 53 papers from O3
- 3 from O4 so far

LVK detectors

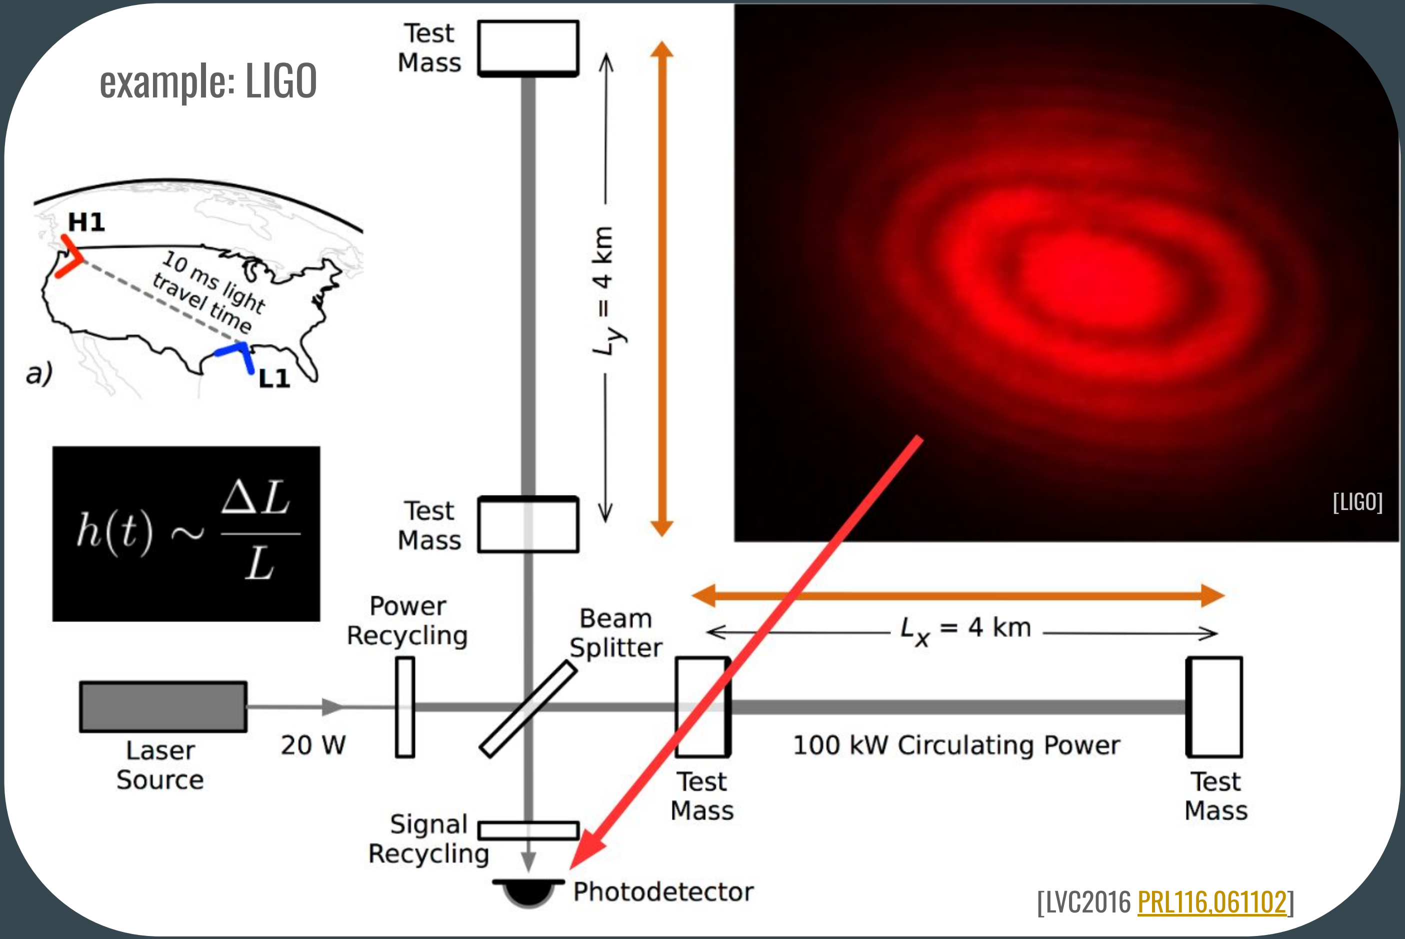

Note: O1 era graphic!

Laser power has since increased, O4 laser offers max input power 125 W.

Many other improvements, including

- new test masses
- frequency-dependent squeezing
- various noise mitigations
- ...

[Capote+2025 PRD111,062002]

Virgo: Acernese+2015 CQG32,024001

KAGRA: Akutsu+2021 PTEP,2021,05A101

LVK search types

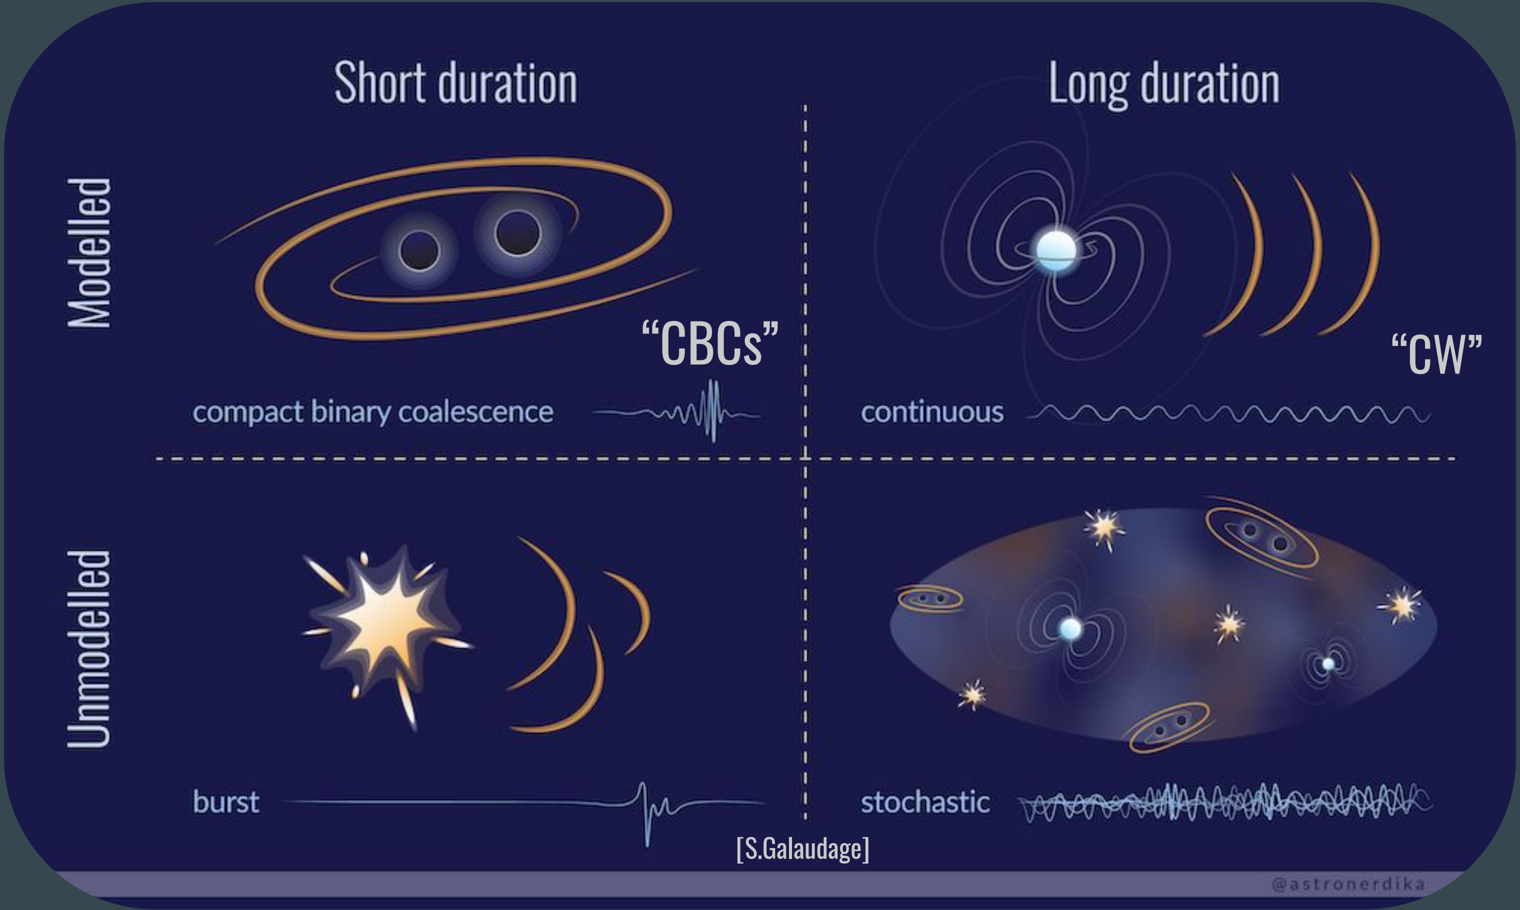

[S.Galaudage]

CBCs

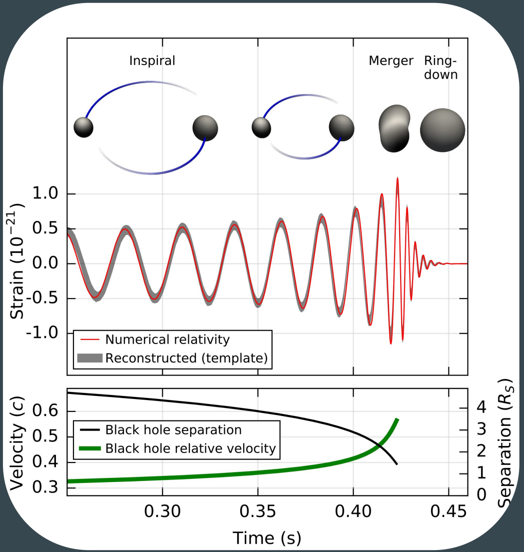

[LVC2016 PRL116,061102 / PRL116,241102]

- evolution of compact objects
- tests of GR in strong-field regime
- “standard siren” cosmography
- nuclear matter at extreme densities

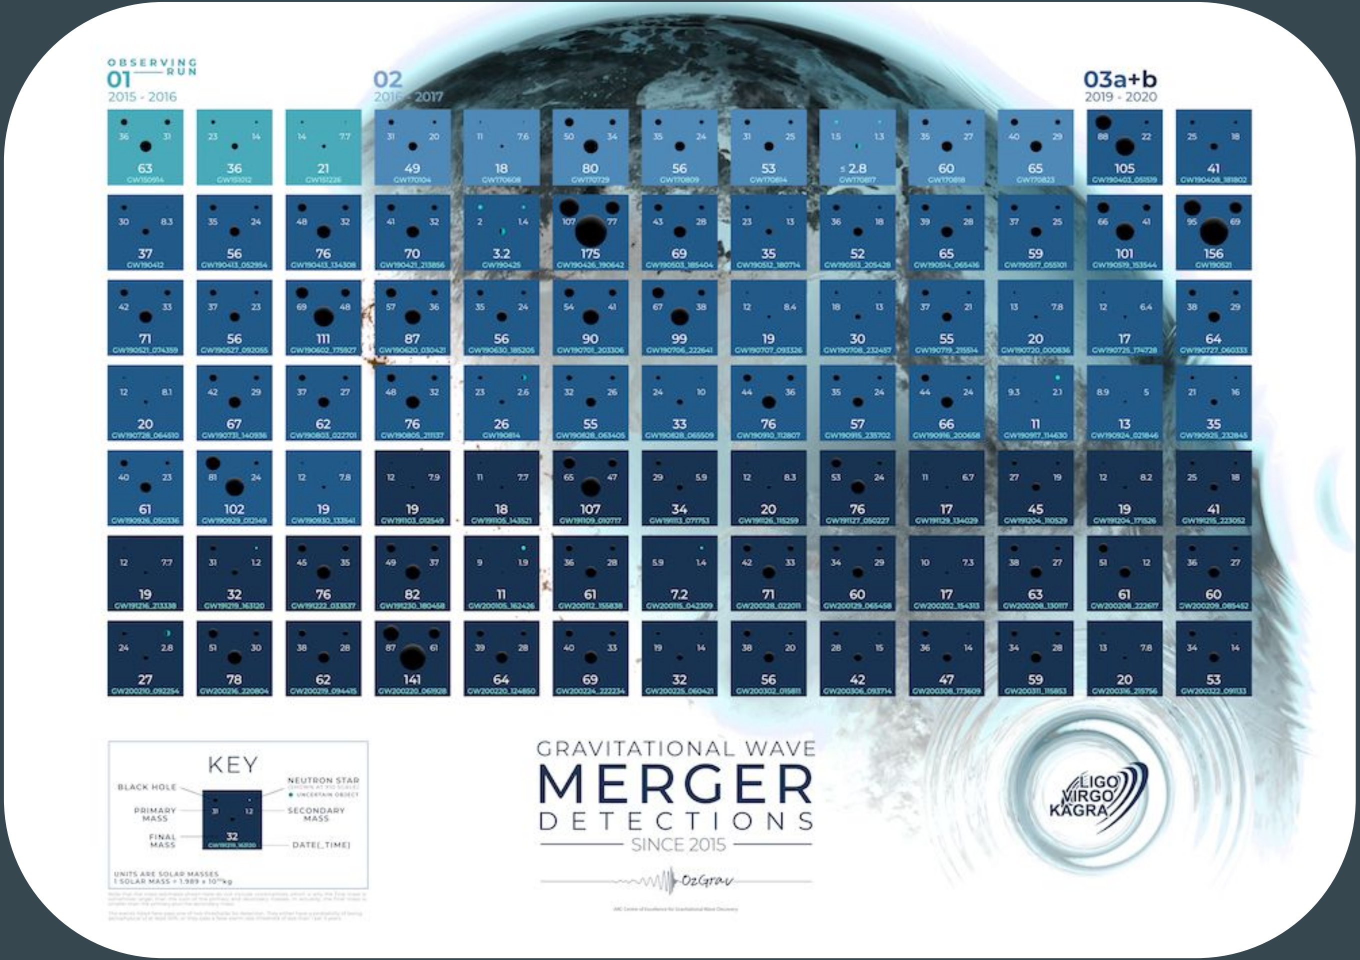

[OzGrav / LVK2024 PRX13,041039]

Beyond CBCs: GW bursts

- less well-modeled GW transients: eccentric BBHs, supernovae, magnetars, cosmic strings, …

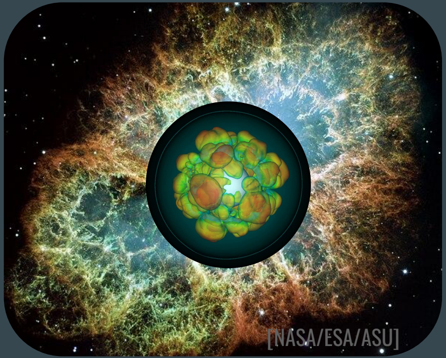

- search with more generic methods: excess power, pattern recognition, …
- No detections so far. (Besides BBHs!)
- Non-detections can still yield physical constraints: nearby supernovae, glitching pulsars, …

Beyond CBCs: stochastic GWs

- Astrophysical backgrounds: overlap of faint, unresolved CBCs
- Cosmological backgrounds = early-universe physics: inflationary tensor modes, phase transitions, …

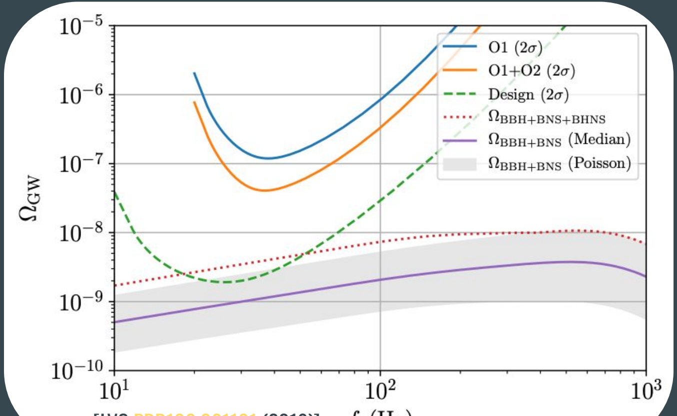

[LVC PRD100,061101 (2019)]

- Tightening upper limits, detection realistic with improving LVK network sensitivity.
- Shout-out to Pulsar Timing Arrays! [IPTA2024 ApJ966:105]

Beyond CBCs: CWs

- GWs from individual spinning neutron stars with non-axisymmetric deformations

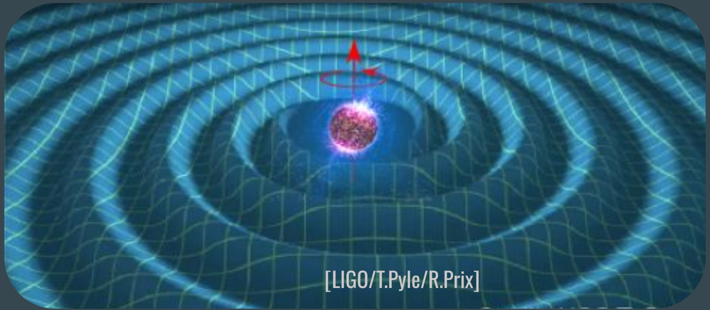
[LIGO/T.Pyle/R.Prix]

- extremely weak, but observable for years
- computationally extremely challenging
- no detection yet → upper limits science so far
- great promise in multi-messenger astronomy
- tests of GR, nuclear matter at extreme densities

Beyond CBCs: new physics

- modified gravity effects
- addressing the Hubble tension
- exotic compact objects (e.g. boson stars)
- primordial black holes
- indirect detection of particle dark matter via boson cloud annihilation
- dark matter direct detection with interferometric detectors

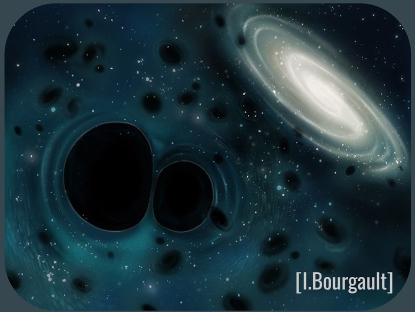
[I.Bourgault]

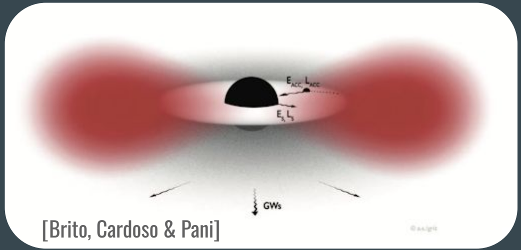

[Brito, Cardoso & Pani]

“AI”/ML and gravitational waves

So, where can it help?

- For some signal types (e.g. CBCs, CWs) we know exactly what we’re looking for, but might not be able to efficiently cover the full generic parameter space with “traditional algorithms”.
- We also search for “unknown knowns” (waveforms that can’t be fully predicted) and “unknown unknowns”.

And why is it difficult?

- We are looking for extremely faint signals in our detector noise: only the loudest CBCs (peak strain ~10^{-21}) can be directly “seen” in the output timeseries.
- The opposite of the original application regime of many ML algorithms (feature-rich and high-contrast).

i. Detector design, operation and characterisation

state of the art

- GW detectors are extremely complex and intricate machines
- near-Gaussian noise floor = superposition of instrumental and environmental noise sources
- plus non-stationary and non-Gaussian components

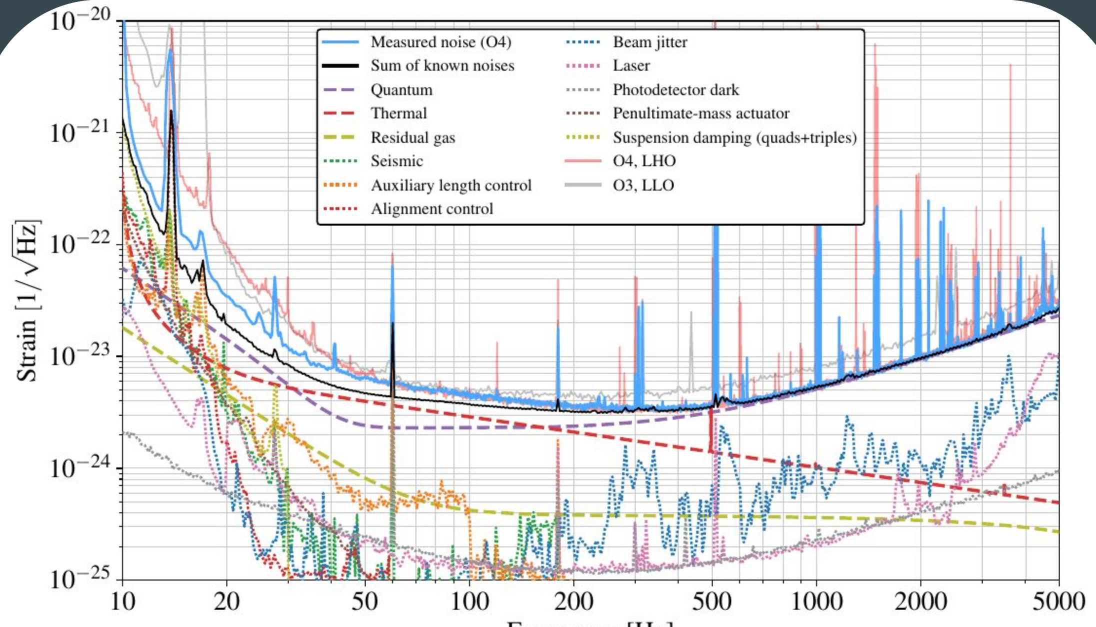

[Capote+2025 PRD111,062002]

(b) Noise budget for the LIGO Livingston Observatory, as of October 2023.

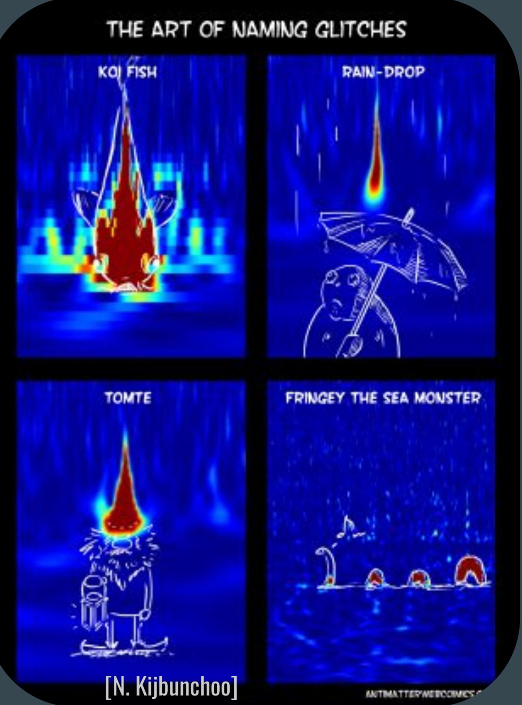

[N. Kijbunchoo]

i. Detector design, operation and characterisation

- ML could offer possibilities for:
    - optimising detector design across an immensely-dimensional parameter space
    - real-time optimization of detector parameters (augmenting the control loops)
    - real-time noise prediction and mitigation: correlations between auxiliary sensors (environmental/instrumental monitors) and the main GW strain channel
    - non-linear noise regression and subtraction after data-taking
    - glitch identification and removal (non-Gaussian transients)

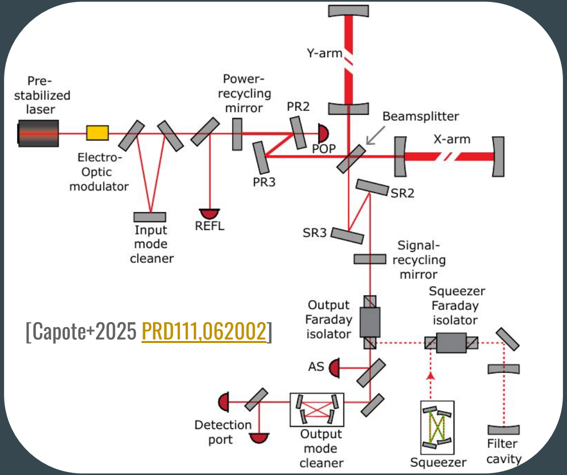

[Capote+2025 PRD111,062002]

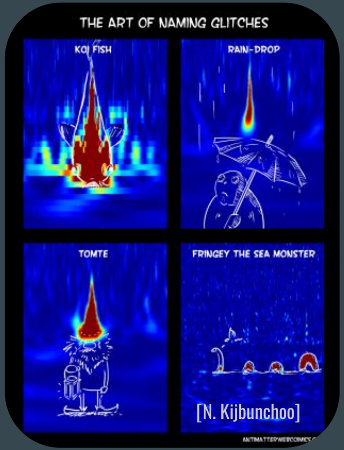

[N. Kijbunchoo]

i. Detector design, operation and characterisation

highlight example

- Some noise components have secure "witness channels": auxiliary sensors that allow monitoring their time-varying strength and subtracting the effect from the GW strain channel
- Vajente+2020 PRD101,042003: "*Machine-learning nonstationary noise out of gravitational-wave detectors*"
  → NonSENS: "Non-linear Noise Subtraction"
  - parameterised model for non-linear relations between channels
  - optimised with gradient descent model (ADAM)
  - O3: non-linear subtraction of narrowband instrumental lines, in particular 60 Hz power line
  - O4: mainly to remove beam jitter noise

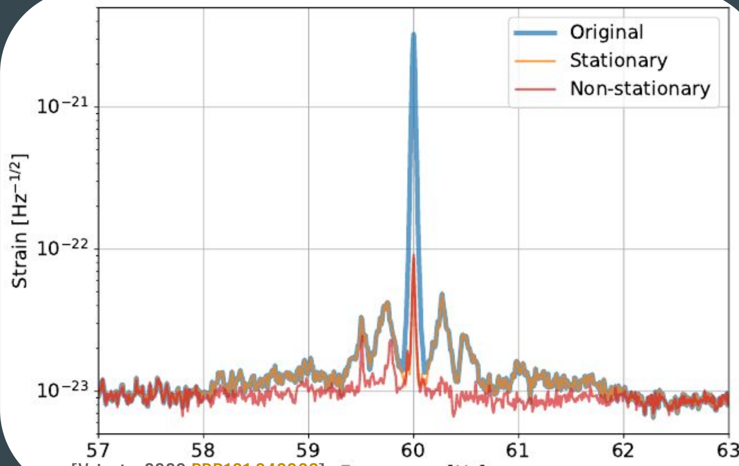

[Vajente+2020 PRD101,042003]

i. Detector design, operation and characterisation

highlight example

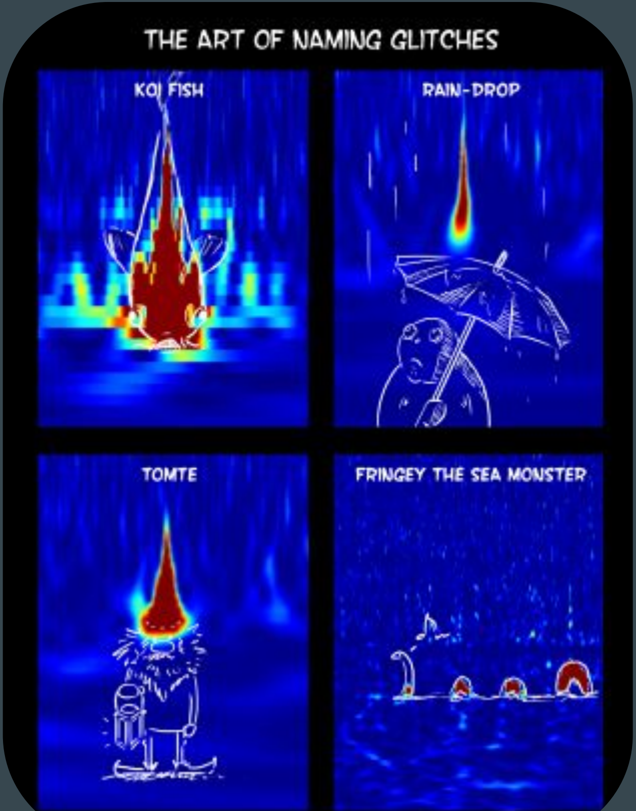

[N. Kijbunchoo]

- Loud, short, broadband, complex-morphology *glitches* are among the most problematic noise artifacts.
- Gravity Spy: synergy of citizen science and machine learning
  - Zevin+ 2017 CQG34,064003, 2023 EPJP139,100
  - triggers flagged by excess power algorithm ("omicron")
  - basic data unit: time-frequency spectrograms
  - initial pre-labeled data to train a CNN for pre-classification
  - volunteers on Zooniverse* confirm/refine classification
  - feedback loop to retrain the network
- results used e.g. in rapid response to online alerts
- actual glitch removal mainly with BayesWave algorithm [Hourihane+2022 PRD106,042006]

[*] zooniverse.org/projects/zooniverse/gravity-spy

i. Detector design, operation and characterisation

open problems & future directions

- How can ML/AI help design the next detector generation
  (Einstein Telescope, Cosmic Explorer, future space missions)...?

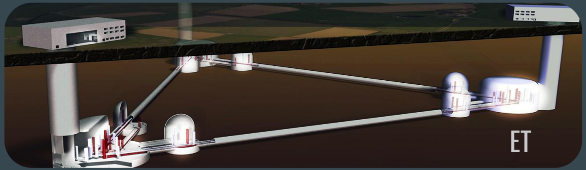

[Abac+ 2503.12263]

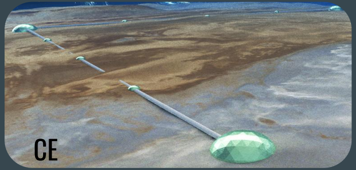

[Evans+ 2109.09882]

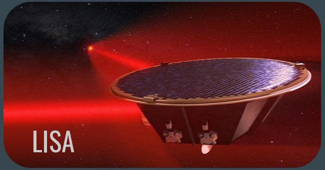

[Amaro-Seoane+2020]

- Can we embed more ML/AI into the day-to-day detector operation?
  (control loops, lock acquisition and loss prevention, ...) (→ e.g. YOLO point absorber detection, Goode+ 2411.16104)
- More production uses for improved noise subtraction and glitch mitigation?
  (e.g. DeepClean [Saleem+2024 CQG41,195024], DeepExtractor [Dooney+ 2501.18423])
- ML/AI for hunting narrow spectral lines, which especially affect long-duration signal searches?
- Realistic noise simulation (e.g. Gengli glitch generator: Lopez+ 2205.09204)
- Improved automation of calibration and detector characterisation tasks:
  currently severely person-power limited and a key dependency for all LVK observational results

ii. Compact binary coalescences

state of the art

- Signal waveforms can be predicted from General Relativity (models based on analytical+numerical results)
- "Searches": find candidates and estimate their significance:
  - multiple matched-filter pipelines (fixed template banks)
  - weakly-modelled pipelines too
- "Parameter estimation": Bayesian inference
- Challenges:
  - full generic parameter space coverage (e.g. orbital precession and eccentricity)
  - search efficiency in periods affected by non-stationary noise
  - computational cost of full Bayesian inference
  - robustness of Bayesian inference in the presence of noise glitches
  - latency for public alerts (enabling telescope follow-up)

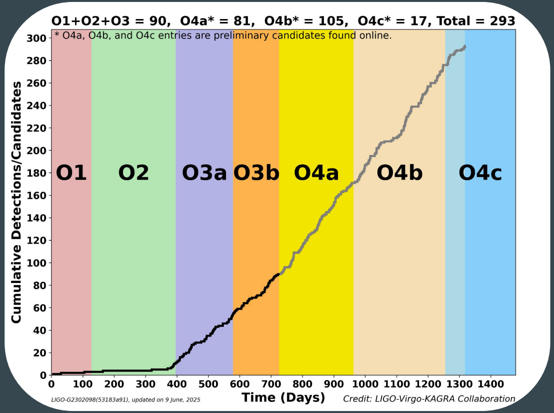

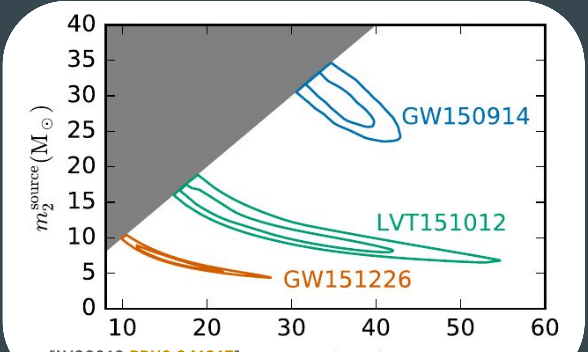

[LVC2016 PRX6,041015]

a review:
Chatziioannou+
2409.02037

ii. Compact binary coalescences – searches

highlight example

- Main promise of ML: front-load computational cost to training phase, find candidates even faster
- GW g2net-Kaggle challenge* and MLGWSC-1 [Schäfer+2023 PRD107,023021]: standardised data sets to compare ML solutions to each other, and standard matched filter
- AresGW** [Nousi+2023 PRD108,024022, Kolonari+2025 MLST6,015054], based on ResNet: strong performance on MLGWSC-1, 8 new GW candidates reported from O3 data
- SAGE*** [Nagarajan&Messenger 2501.13846], OSNet feature extractor + ResNet/CBAM classifier: further improvements on MLGWSC-1 over AresGW and matched filter
  - paper also highlights 11 types of *biases that challenge CBC detection with ML*: training set construction, spectral bias, etc
- Gains can come from certain parameter space regions and/or better performance in noisy data.
- Caveat: ML submissions often optimised to the specific parameter space of the challenge, which could also be done to improve performance of standard methods! (e.g. Kumar&Dent 2024 PRD110,043036)

More examples:
Trovato+2024 CQG41,125003
Marx+ 2403.18661

[*] kaggle.com/c/g2net-gravitational-wave-detection (2021) | [**] github.com/vivinousi/gw-detection-deep-learning | [***] github.com/nnarenraju/sage

ii. Compact binary coalescences – inference

highlight example

- DINGO [Dax+2021 PRL127,241103, 2023 PRL130,171403]: neural posterior estimation (with normalising flows) in seconds–minutes instead of hours–days per event

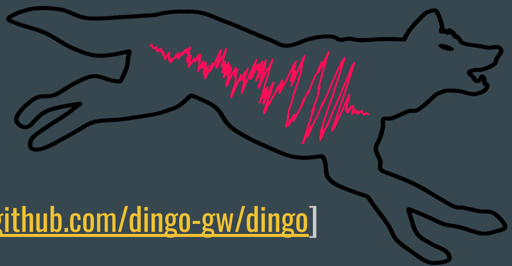

[github.com/dingo-gw/dingo]

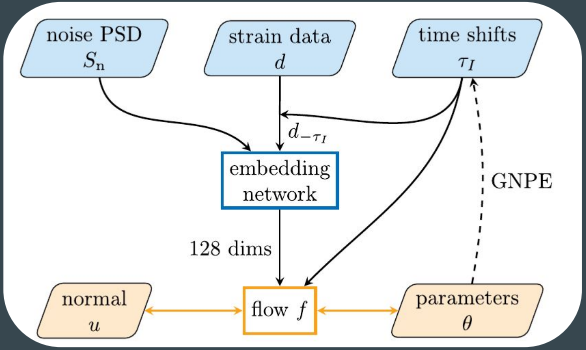

[Dax+2021 PRL127,241103]

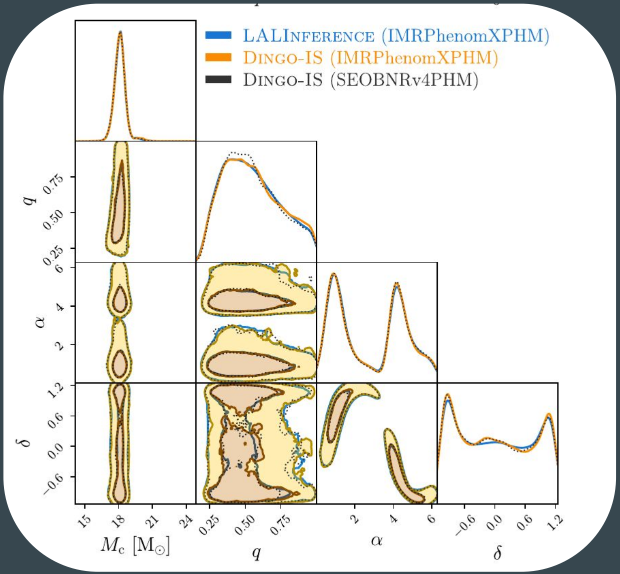

[Dax+2023 PRL130,171403]

- initially working best for high-mass, short binary-black-hole signals, now also extended to binary neutron stars [Dax+2025 Nature 639,49-53]

- special promise for otherwise extremely expensive waveforms, e.g. including orbital eccentricity [Gupte+ 2404.14286]

Other examples:
Nessai: Williams+2021 PRD103,103006
Peregrine: Bhardwaj+2023 PRD108,042004
AMPLFI: Chatterjee+ 2407.19048

ii. Compact binary coalescences

- optimal network architectures and training methods to deal with the typical kinds of biases identified by 2501.13846 and with the full complexities of real detector data
- fair comparisons between ML and "traditional" search algorithms, avoiding fine-tuning
- finding the right mix for fruitful coexistence of fast neural and "full" Bayesian inference
- passing detailed LVK scientific&code review and operational stability criteria for production runs, including low-latency alerts (gracedb.ligo.org | emfollow.docs.ligo.org/userguide)
- ML in waveform modeling itself
- future detectors:
  - longer signal durations
    (e.g. Hu+2412.03454, Dax+2025 Nature 639,49-53)
  - huge detection rates (→ overlapping signals!)
    (e.g. Langendorff+2023 PRL.130.171402, Alvey+ 2308.06318, Santoliquido+ 2504.21087)

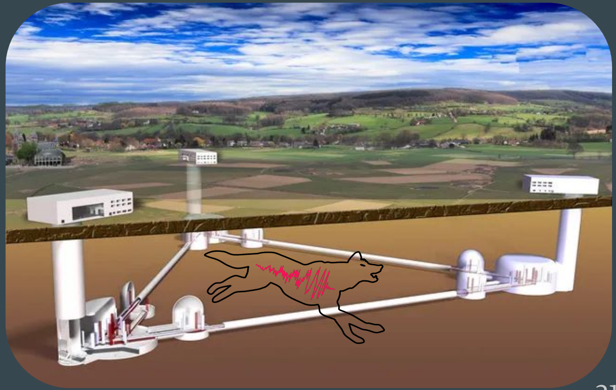

iii. GW bursts

state of the art

- less well-modeled GW transients: eccentric BBHs, supernovae, magnetars, cosmic strings, …
- … and unknown unknowns!
- most LVK algorithms based on some form of *excess power*
  and searches for correlated structures in time-frequency spectrograms
- also possible *coherently* across multiple detectors
- basically: anomaly detection and pattern recognition
- weakly-modeled techniques,
  such as wavelet decomposition,
  also allow *signal reconstruction*
- everything already on a continuum
  towards ML *(e.g. are KDEs ML…?)*

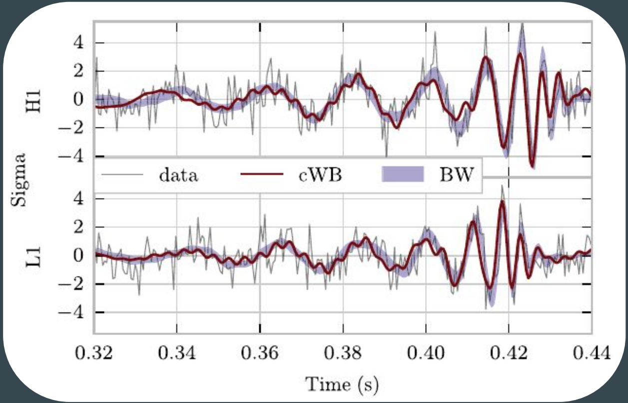

[LVC2016 PRD93,122004]

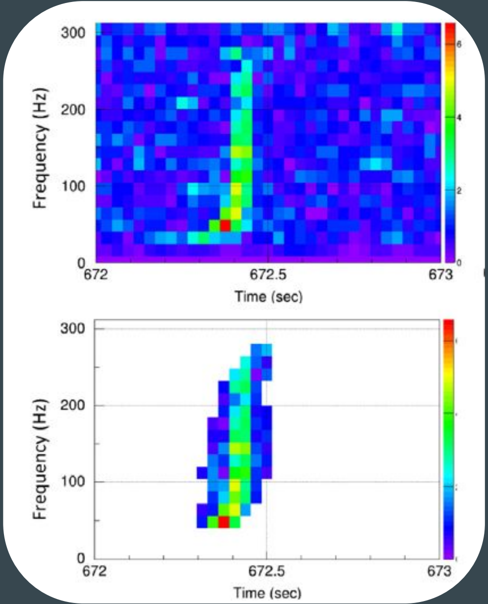

[Drago+2021 JsoftX14,100678]

iii. GW bursts

highlight example

- MLy pipeline
  - Skliris+2024 [PRD110,104034]
  - [git.ligo.org/mly/mly]
- dual architecture for coincidence and coherent modes across detectors
- first tested on LIGO-Virgo O2 data
- now LVK-reviewed and running "in production" on O4 data [[emfollow.docs.ligo.org/userguide/analysis/searches.html#unmodeled-search]]

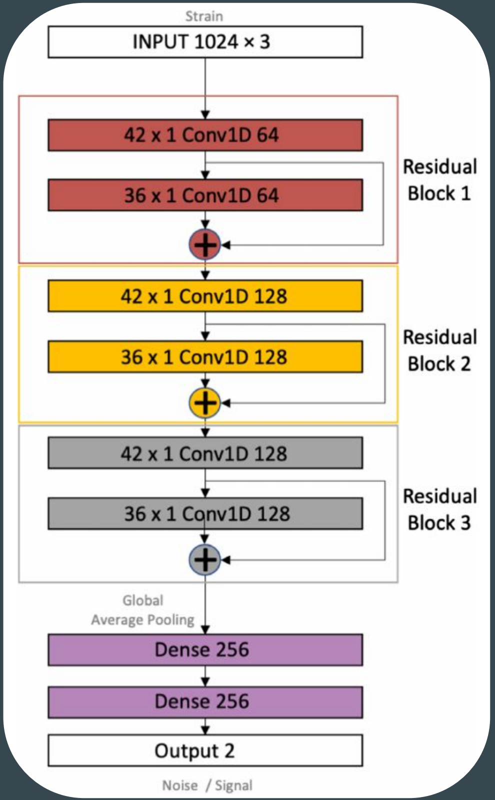

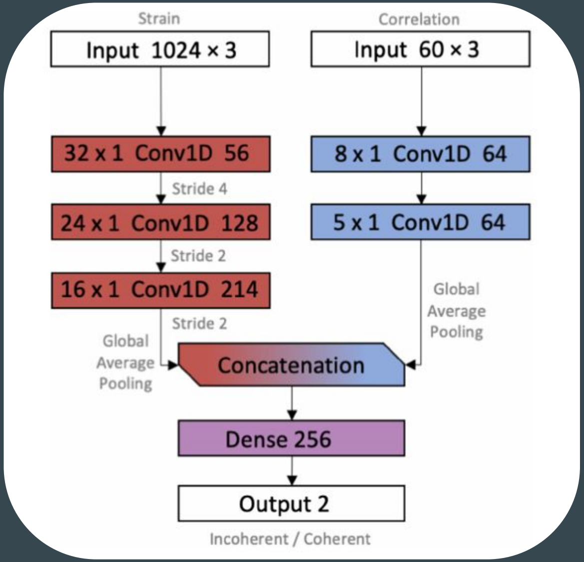

iii. GW bursts

open problems & future directions

- extended real-world testing
- besides pure ML pipelines like MLy, also “traditional” ones getting enhanced with ML ingredients, e.g. XGBoost postprocessing for cWB [gwburst.gitlab.io] – Mishra+2021 PRD104,023014
  → used on O3 data in Szczepańczyk+2023 PRD107,062002, Mishra+2025 PRD111,023054
- bridging the gap between “modelled” and “unmodelled burst” analyses
  for complicated sources like supernovae, with simulation-based inference etc.
- pure anomaly detection frameworks for the known unknowns
  (e.g. GWAK, Raikman+2025 MLST5,025020 and 2412.19883)
- interpretable/explainable AI to understand what is being detected?

iv. Continuous Waves

state of the art

- simple signal model → matched filtering
- optimal fully-coherent analysis possible for known pulsars with full timing model from EM observations
- computationally extremely challenging for *unknown* sources: large parameter space and extremely fine required grid resolution
- semi-coherent methods provide best tradeoff so far between sensitivity and computing cost
- similar issues for long-duration CW-like transients from glitching pulsars, BNS remnants, …
- g2net-Kaggle challenge[*] mostly produced GPU-optimised variants of "traditional" semi-coherent methods

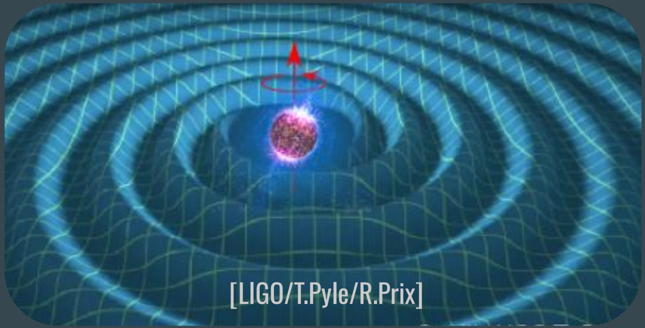
[LIGO/T.Pyle/R.Prix]

a review: Riles2023 LRR26,3

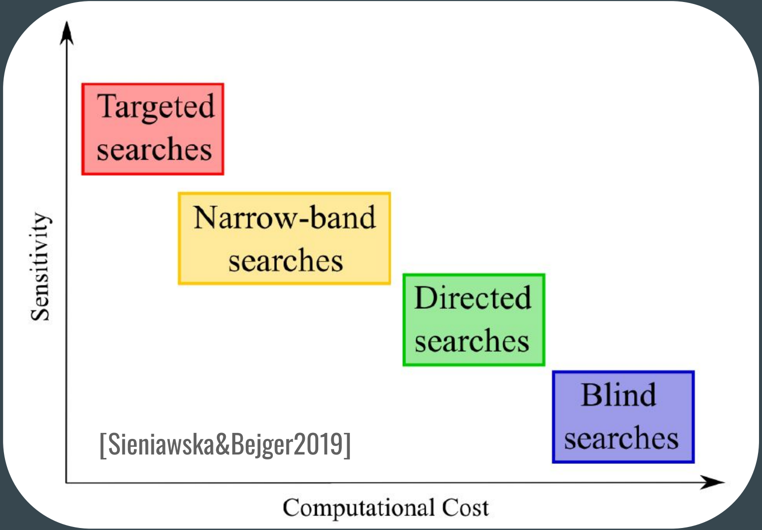

[Sieniawska&Bejger2019]

[*] kaggle.com/competitions/g2net-detecting-continuous-gravitational-waves (2023)

iv. Continuous Waves

highlight example

- Joshi&Prix 2023: "*Novel neural-network architecture for continuous gravitational waves*" [PRD108,063021]
- Identified the key challenges of neural networks applied to CWs:
    - signals not only faint, but spread across long durations, with low local contrast and rich structure
    - morphology changes across parameter space, Doppler shifts become more challenging at high frequencies
- For durations up to 10 days, customised CNNs can almost reach matched-filter performance, but not yet quite.
- Joshi&Prix 2024 [PRD110,124071]: can also generalise to a single network trained across 20–1000 Hz

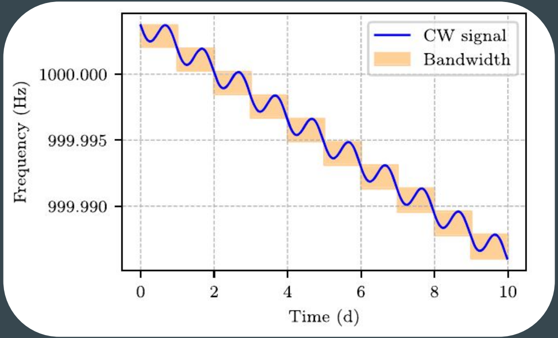

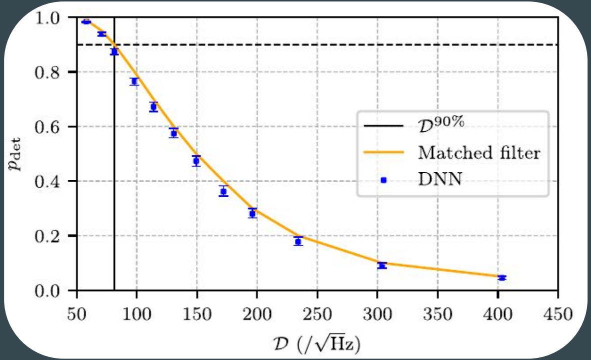

iv. Continuous Waves(-like long transients)

highlight examples

- BNS merger remnants: rapid spindown

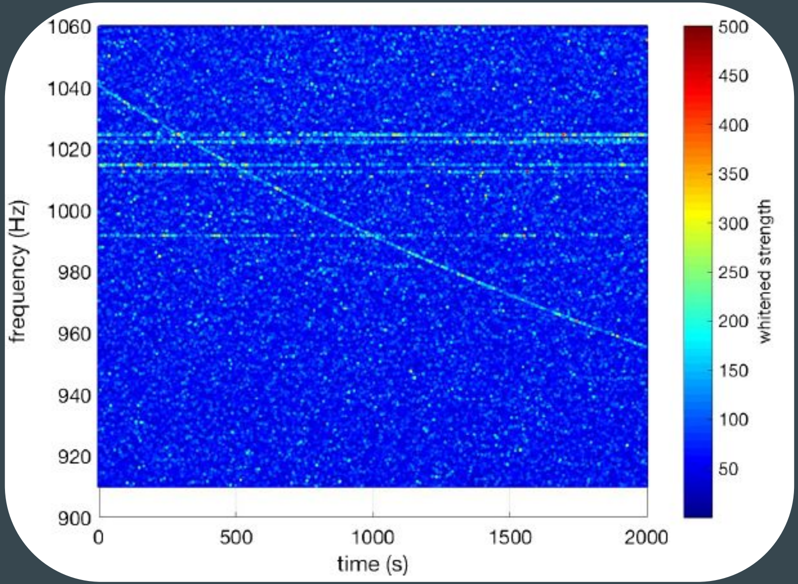

- Miller+2019 PRD100,062005: *How effective is machine learning to detect long transient gravitational waves from neutron stars in a real search?*
- using CNNs on spectrograms

- pulsar glitches can trigger CW-like transients of unknown duration
- Modafferi+2023 PRD108,023005: *Convolutional neural network search for long-duration transient gravitational waves from glitching pulsars*
- hybrid approach: CNN on matched-filter intermediate data products

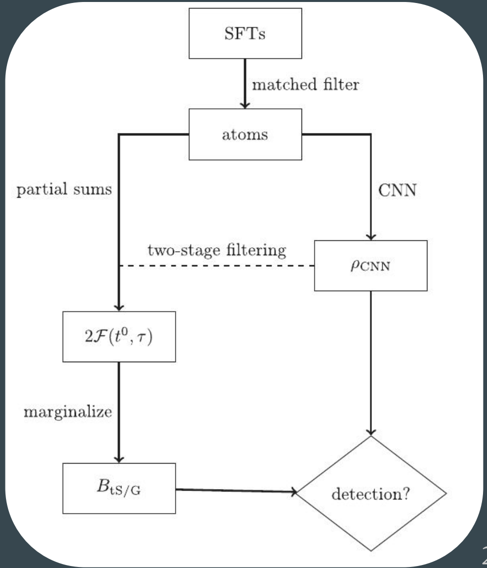

iv. Continuous Waves

- still working towards a "first detection" with *any* method ("traditional" or ML)
- immense sensitivity gap between optimal fully-coherent matched filter and what is computationally feasible over large parameter spaces (factors 5–50 in "depth" below the detector noise floor)

  → in principle, ML methods could reduce this!
- also, neutron stars are known to be "messy" → make methods more robust to signal deviations?
- but need to overcome the challenges identified by PRD108,063021 and others:
  - very faint signals…
  - …with even fainter local contrast
  - …and complex morphologies
  - …that vary strongly across parameter space

v. Stochastic signals and backgrounds

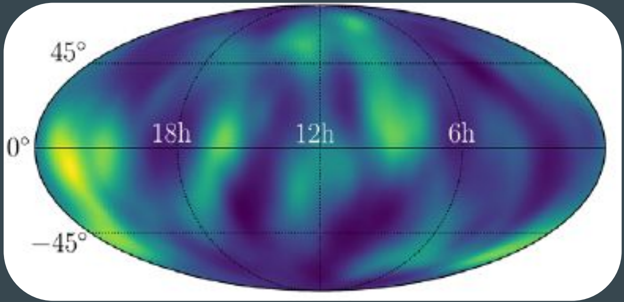

[LVK 2021 PRD104,022005]

- persistent signals without deterministic models
- state of the art: primarily *cross-correlation* between 2+ detectors
- already computationally very efficient
- key challenge: controlling correlated noise sources
- not many example applications of ML to this yet
- open problems & future directions:
  - ML noise mitigation?
  - Early-universe physics through simulation-based inference?
  - intermittent, non-Gaussian backgrounds: enabling optimal Bayesian-style search for stochastic background from faint CBC sources? [Smith&Thrane2018 PRX8,021019]
  - overlap with "burst" and CW-like searches for long-duration transients, with possibly rather complicated waveforms (newborn neutron stars, magnetars, ...)

a review: Remortel+2023 PPNP128,104003

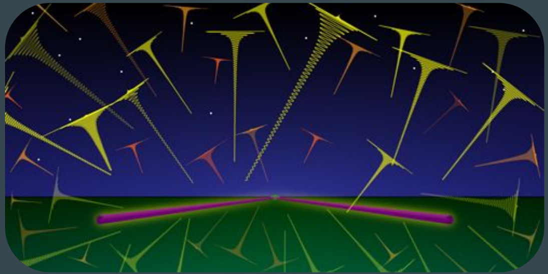

[APS/A.Stonebraker]

vi. Beyond LVK: future detectors

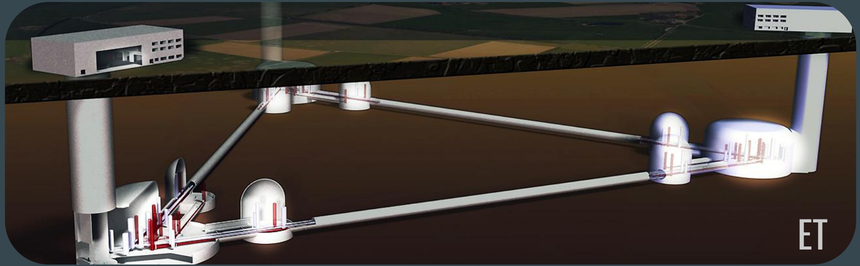

European Einstein Telescope
– in Sardinia...?
[Abac+ 2503.12263]

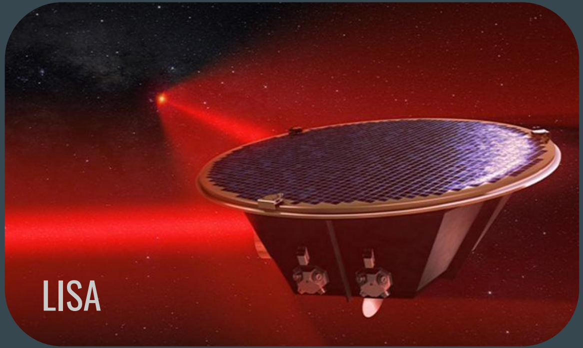

ESA
(+NASA junior partner?)
[Amaro-Seoane+2020]

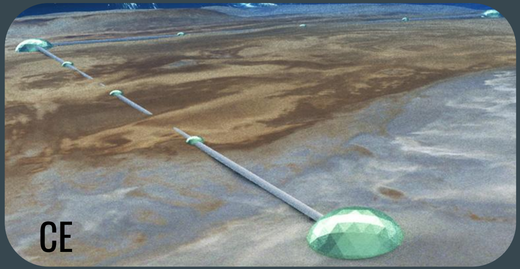

Cosmic Explorer (US?)
[Evans+ 2109.09882]

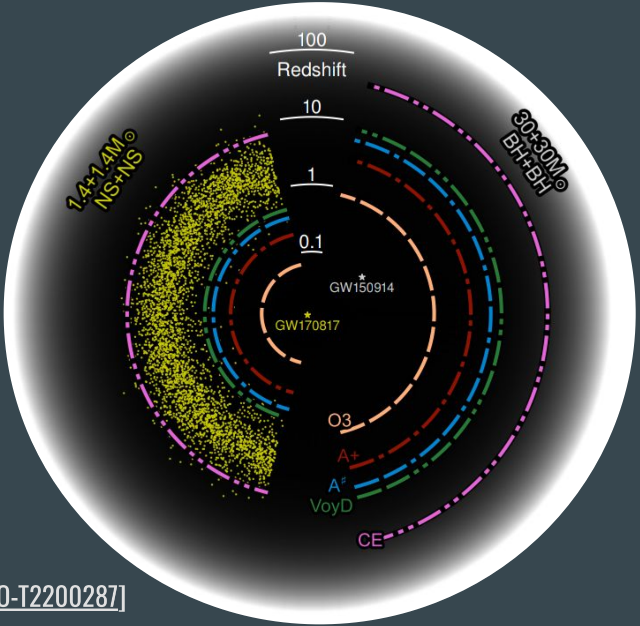

[LIGO-T2200287]

vi. Beyond LVK: Einstein Telescope

- underground cryogenic 3G detector
- primary design: 10 km triangle, single site
- alternative under discussion:
  two sites with 10 km L-shaped detectors
- Abac+ “The Science of the Einstein Telescope” 2503.12263

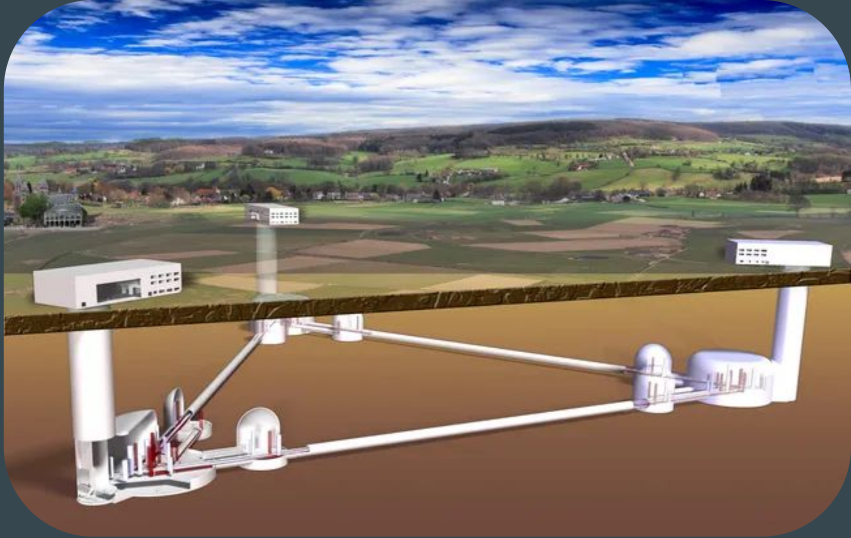

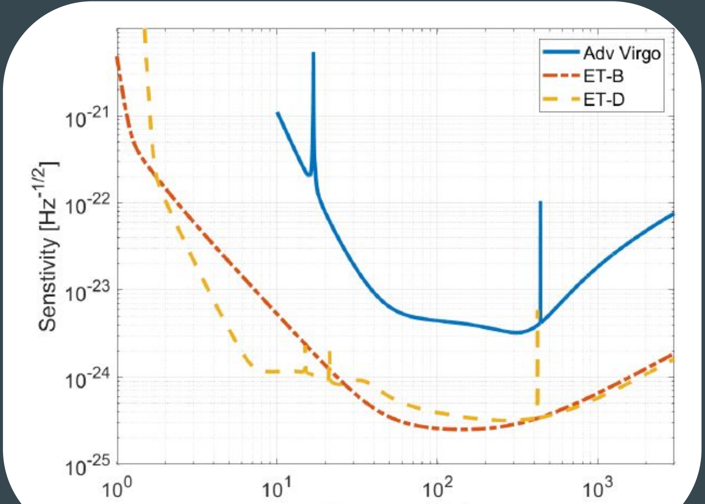

[Maggiore+ JCAP03(2020)050]

- Sardinia?
- Meuse-Rhine?
- ~~Hungary?~~
- East Germany?

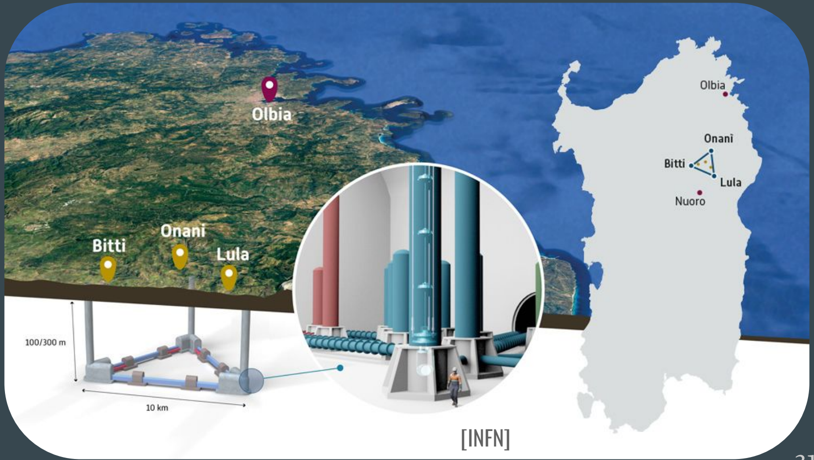

[INFN]

vi. Beyond LVK: future detectors

- Increased detection rates and signal complexities will make ML tools indispensable.
- Low-frequency sensitivity of ET&CE:
  CBCs become long-duration signals
  ⟶ new challenges in waveform accuracy
  and algorithm efficiency

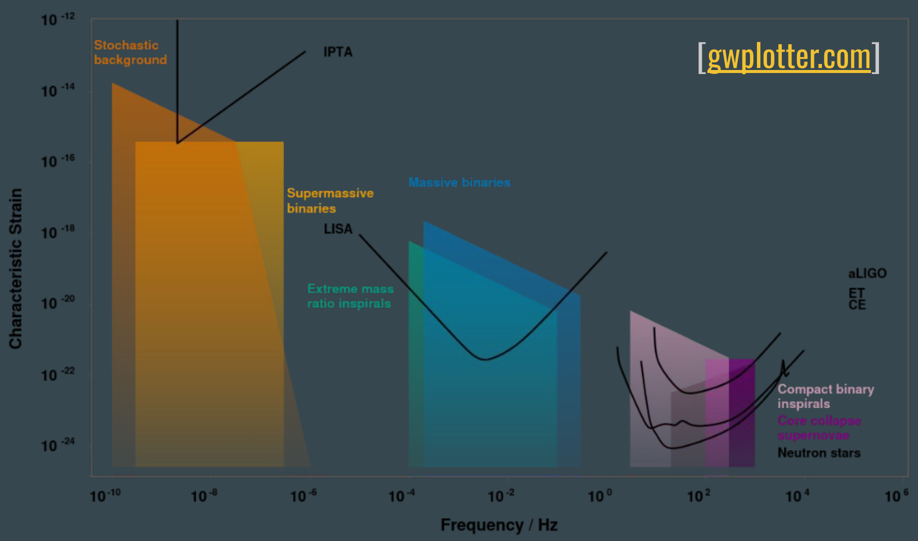

- LISA opens up entirely new signal types
  [Afshordi+ 2311.01300],
  especially extreme mass-ratio inspirals (EMRIs)
  - very long and complex waveforms,
    big challenge to model and to analyse
  - many works already exploring neural networks for these

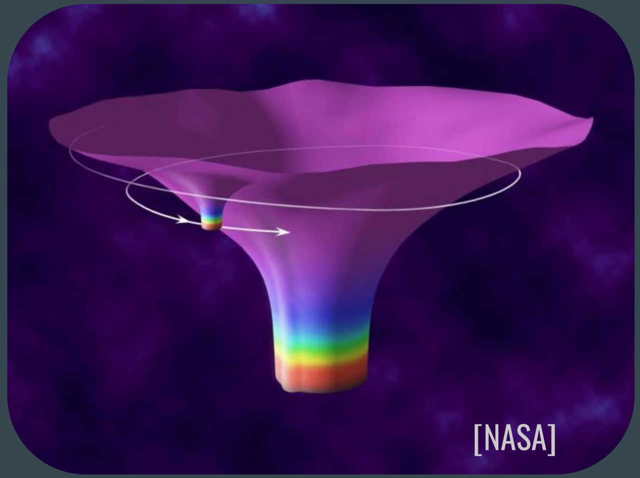

Conclusions: “AI” and gravitational waves (?)

- Though often in early stages, ML methods are already making an impact across GW science.
- Recent phase transition from “proofs of concept” to the first noise mitigation, signal search and parameter estimation ML algorithms productively contributing to the LVK science outputs.
- Bright future ahead applying more advanced methods from the AI/ML community and penetrating further into the realm of “production analyses”.
- No traces of real “AI” yet – if anything, human researchers are more than ever required to use their own intelligence to define well-posed problems for applying ML, and to design fair assessment benchmarks against “traditional methods” and for successful real-world applications.
- Increased detection rates and signal complexities with final stages of LVK network, and especially with future GW detectors, will make ML tools indispensable.
- For more details, see LRR by Cuoco, Cavaglià, Heng, Keitel & Messenger, doi:10.1007/s41114-024-00055-8

  (but lots of great recent work not yet included!)

Home › Living Reviews in Relativity › Article

Applications of machine learning in gravitational-wave research with current interferometric detectors

Review Article | Open access | Published: 27 February 2025

Volume 28, article number 2, (2025) Cite this article

Acknowledgments

This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation. LIGO Laboratory and Advanced LIGO are funded by the United States National Science Foundation (NSF) as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. KAGRA is supported by Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan Society for the Promotion of Science (JSPS) in Japan; National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS) and National Science and Technology Council (NSTC) in Taiwan.

David Keitel is supported by the Universitat de les Illes Balears (UIB); the Spanish Agencia Estatal de Investigación grants CNS2022-135440, PID2022-138626NB-I00, RED2022-134204-E, RED2022-134411-T, funded by MICIU/AEI/10.13039/501100011033, the European Union NextGenerationEU/PRTR, and the ERDF/EU; and the Comunitat Autònoma de les Illes Balears through the Servei de Recerca i Desenvolupament and the Conselleria d'Educació i Universitats with funds from the Tourist Stay Tax Law (PDR2020/11 - ITS2017-006), from the European Union - NextGenerationEU/PRTR-C17.I1 (SINCO2022/6719) and from the European Union - European Regional Development Fund (ERDF) (SINCO2022/18146).