

LVK COLLABORATION · 2026 · HDBANK-2

Gravitational Waves, Match Filtering & Template Banks

An introduction to GW searches — and how we
make template banks faster with normalizing flows.

M. Panzeri

Università di Urbino · INFN Firenze & Urbino · IP2I

AGENDA

Three acts.

01 GWs 101 & History

From Einstein to GW170817 — what they are, how we found them, how lucky we got.

02 Digging Signals out of Noise

The matched filter, the SNR statistic, and the template bank problem.

03 HDBANK-2 & Normalizing Flows

What we built, how we got 10× speedup, and where the ML actually helps.

Spacetime has ripples, we have interferometers

Accelerating masses distort spacetime. Those distortions propagate outward at c . We measure the **strain**:

STRAIN AMPLITUDE

$$h = \frac{\Delta L}{L}$$

QUADRUPOLE RADIATION · TT GAUGE

$$h_{ij} \sim \frac{2G}{c^4 r} \ddot{Q}_{ij}^{\text{TT}}$$


Q_{ij} = traceless quadrupole moment · r = distance to source

At LIGO: $h \sim 10^{-21} \rightarrow \Delta L \sim 10^{-18} \text{ m}$ over 4 km arms. That's 1/1000 the proton diameter. We measure it anyway.

LIGO Livingston



VIRGO



Three kinds of compact mergers.

BBH

Binary Black Hole

Loudest, shortest. Most of the catalog. What GW150914 was.

$$m_1, m_2 \in [3, 500] M_\odot$$

BNS

Binary Neutron Star

Longer inspiral (~100 s in band), quieter. Can glow in EM.
GW170817 was one.

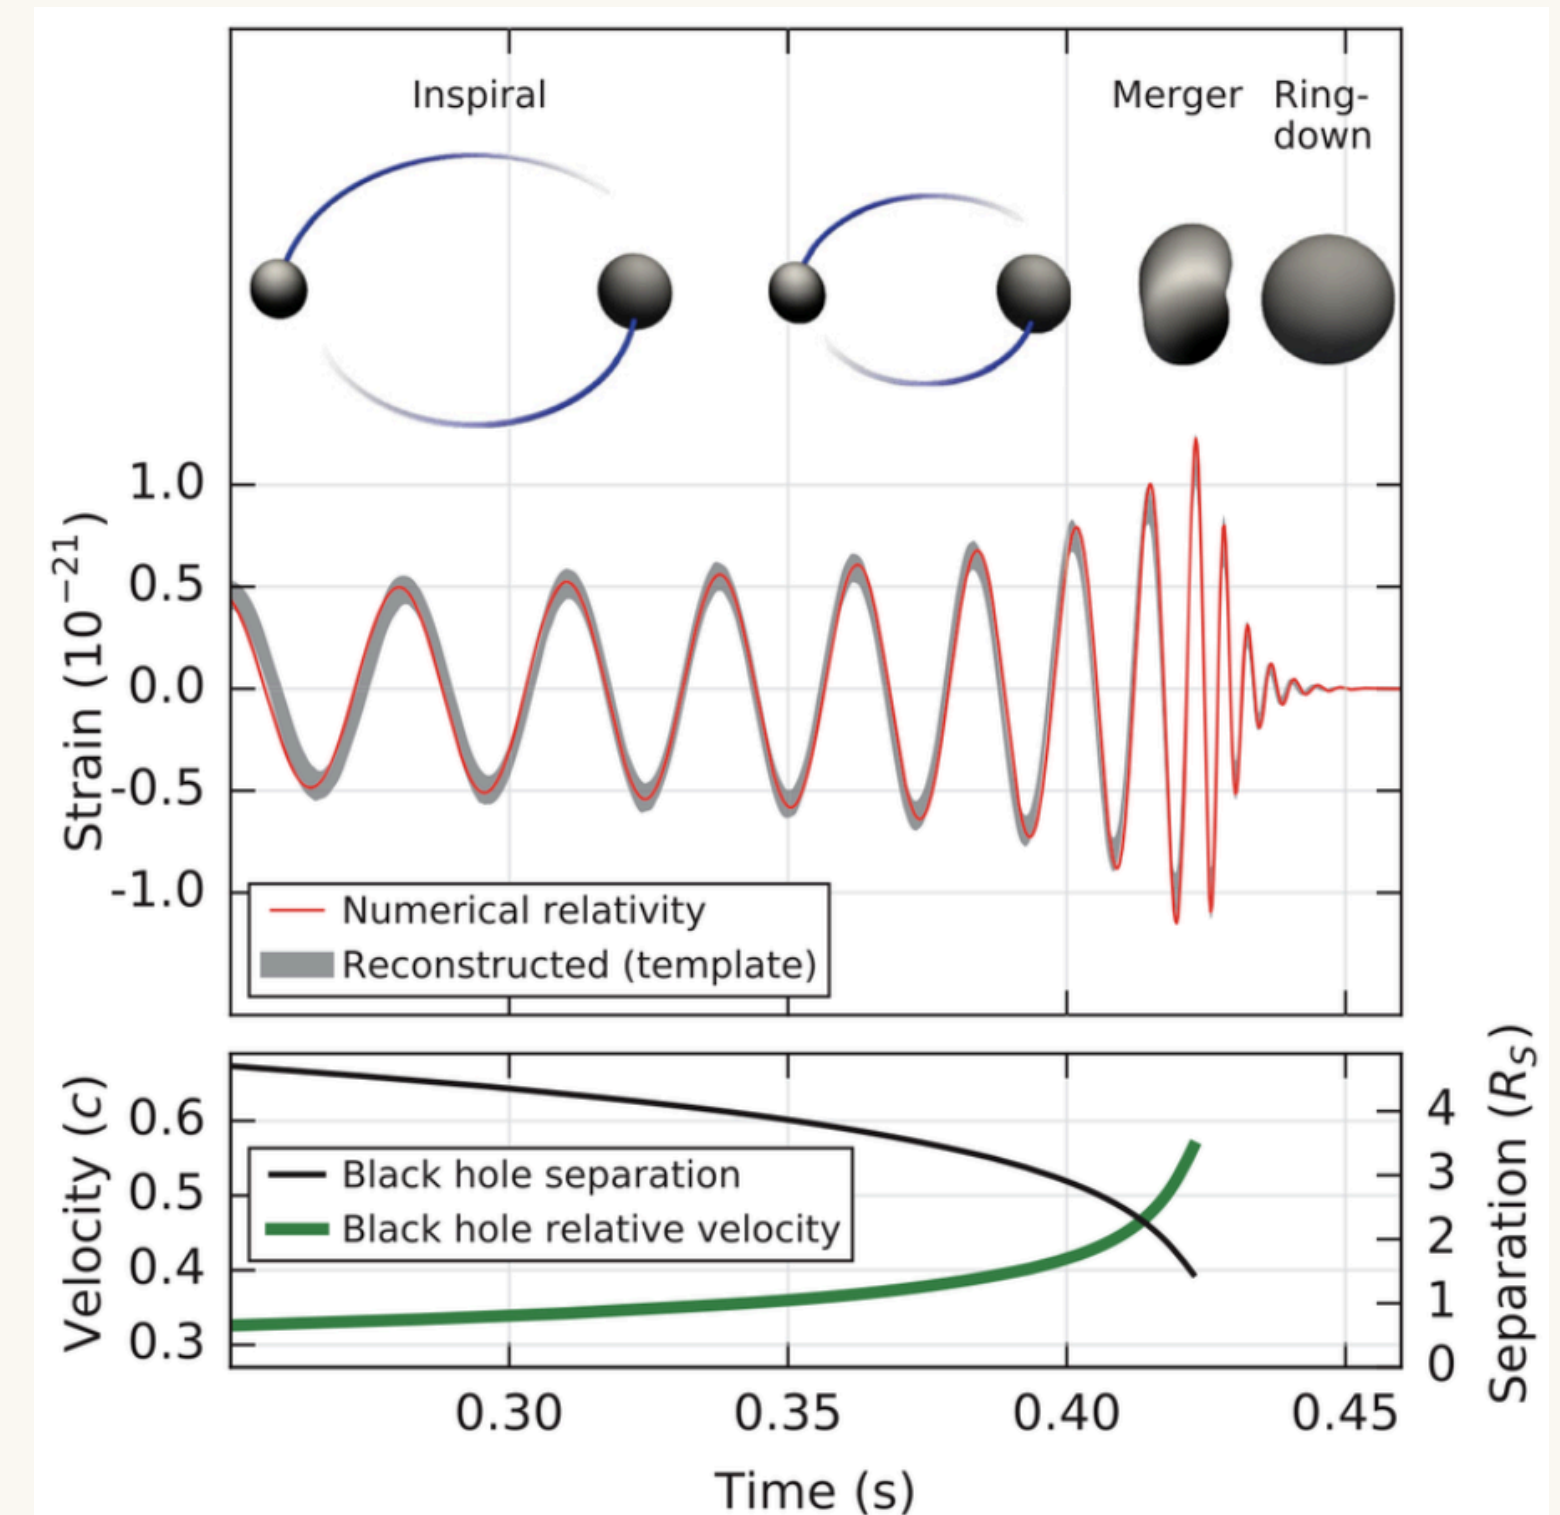
$$m_1, m_2 \in [1, 3] M_\odot$$

NSBH

Neutron Star – Black Hole

Mixed system. Rarer so far. Key target for multi-messenger follow-up.

$$m_1 \in [3, 200], m_2 \in [1, 3] M_\odot$$



Einstein predicted them. Then half-unpredicted them.

1916 — FOUND

Linearising GR, Einstein finds **wave-like solutions** to the field equations: accelerating mass → metric ripples.

1936 — LOST?

Einstein & Rosen submit a paper claiming **GWs don't actually exist**. The reviewer disagrees. Einstein is furious. The reviewer was right.

1957 — CONFIRMED (on paper)

Chapel Hill conference. Feynman's "sticky bead" argument: **GWs carry energy. They're real.** Now we just have to find them.

1960 — THE HUNT BEGINS

Joseph Weber builds the first "**bar detector**" — two aluminium cylinders, 1400 kg each.

1969 — FALSE ALARM

Weber announces detections. The community is thrilled. Then sceptical. Then certain: the signals don't hold up. **GWs remain uncaught.**

AND THEN ...

The first evidence was an orbit getting tighter.

PSR B1913+16 · 1974

Hulse & Taylor discover a **pulsar orbiting a neutron star** —
a perfect cosmic clock in a perfect cosmic lab.

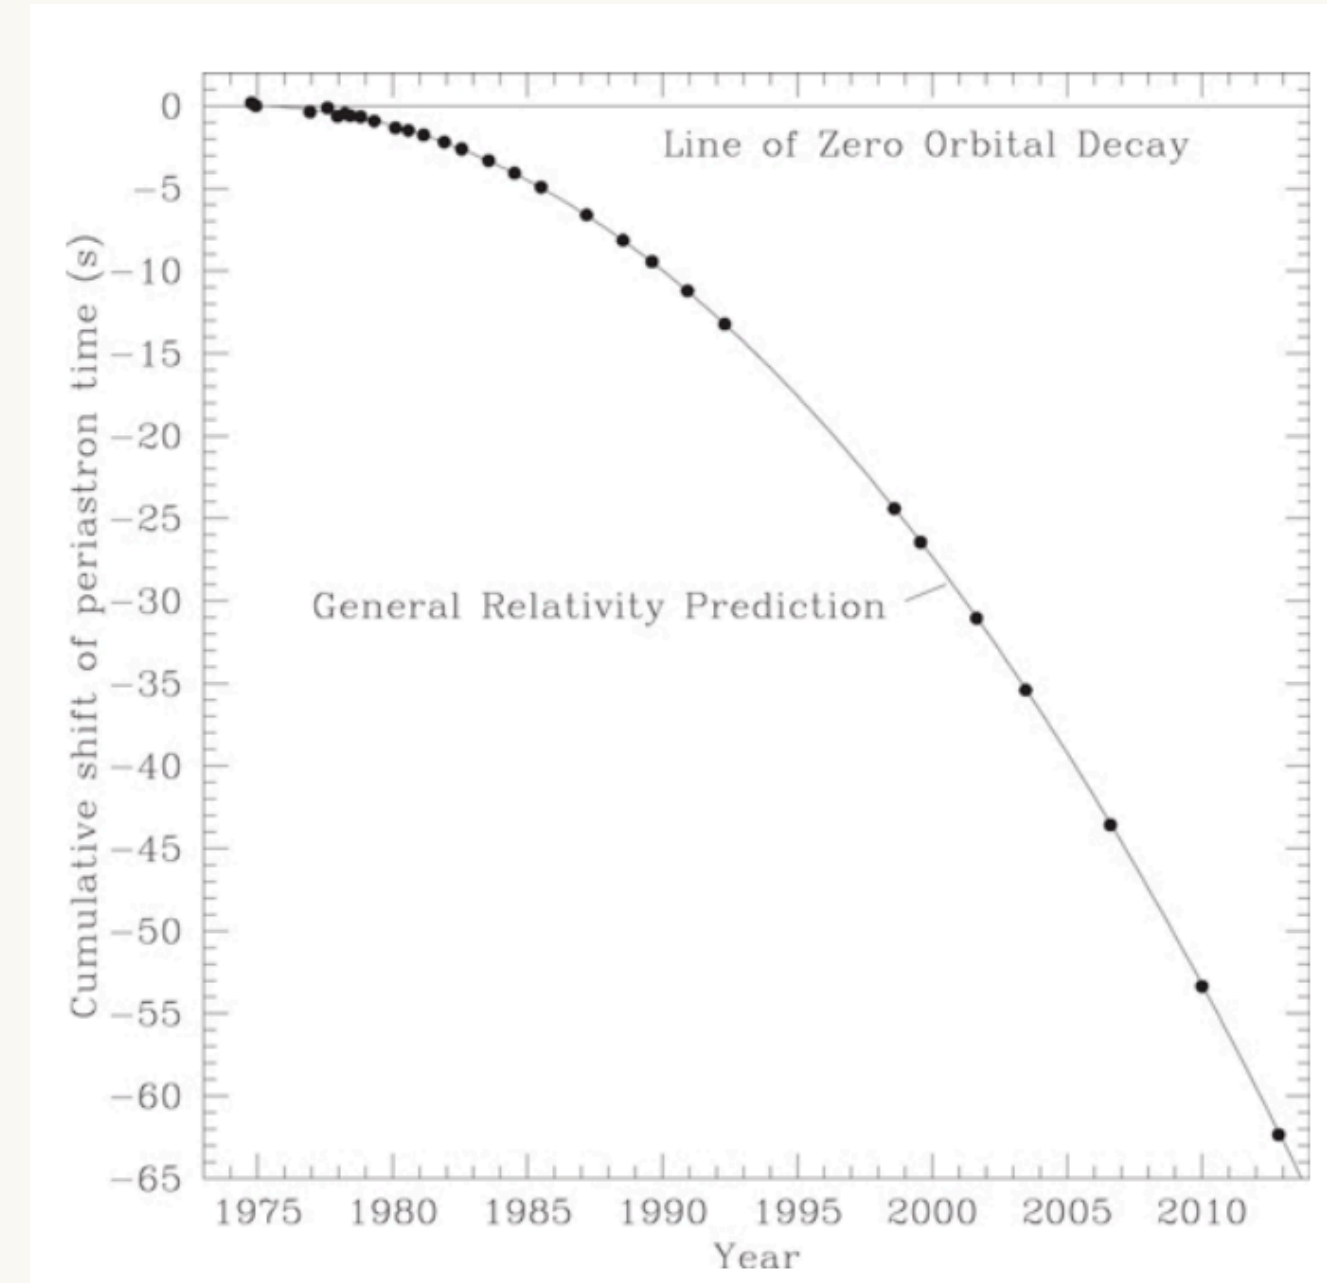
Over 30+ years, **the orbital period shrinks exactly** at the rate
GR predicts, if the system is radiating gravitational waves.

NOBEL PRIZE · 1993

Indirect evidence, rock-solid: the energy has to go somewhere — and
GR tells you exactly where.

Hulse & Taylor 1975

Weisberg, Taylor, Flower 1981



14 Sept 2015 · **GW150914**. We got lucky

Two black holes — $36 M_{\odot}$ and $29 M_{\odot}$ — merged 1.3 billion years ago. $\sim 3 M_{\odot}$ radiated as GWs. LIGO Hanford saw the chirp; 6.9 ms later, Livingston confirmed.

SNR 24

vs threshold ~ 8

440 Mpc

~ 1.6 Gly

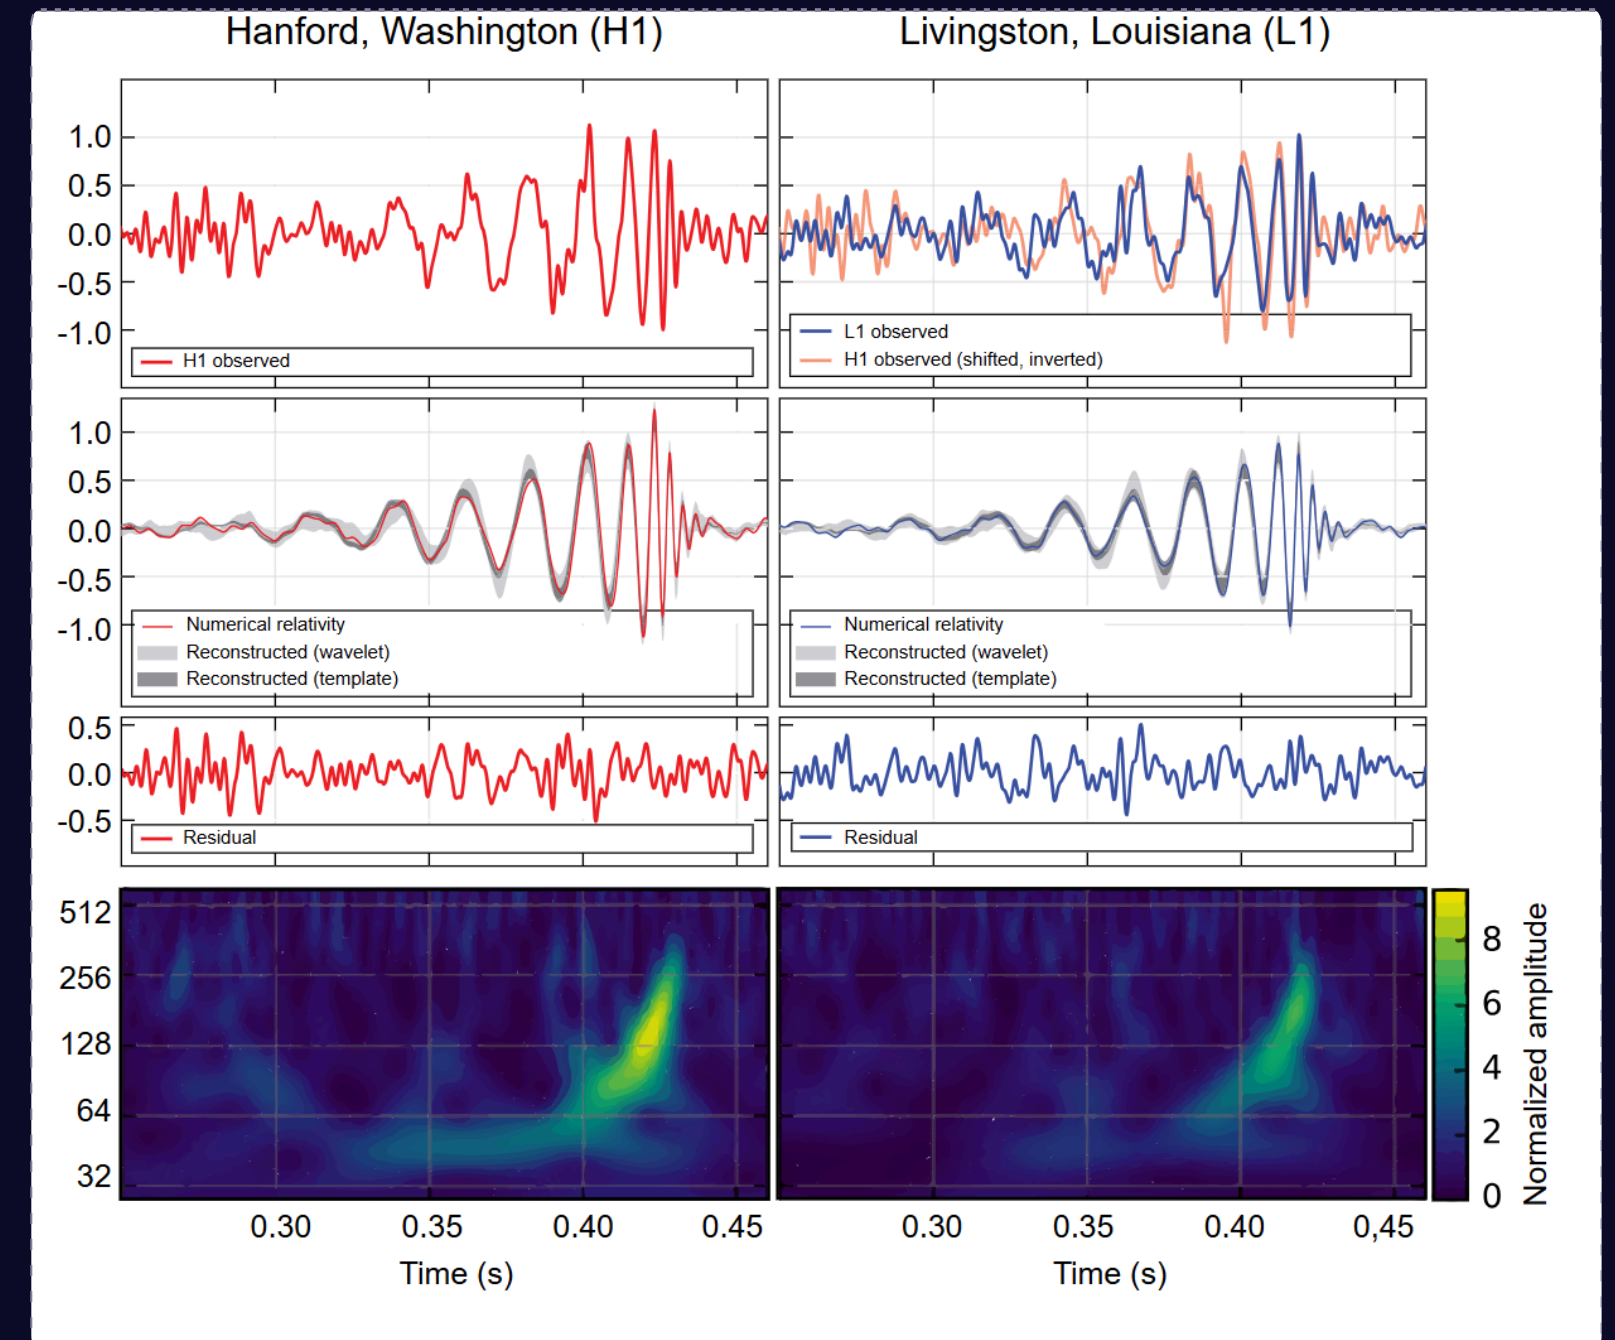
6.9 ms

Δt H1 \rightarrow L1

WE WERE VERY LUCKY

O1 had been running for just **2 days**. The binary was nearly **face-on** ($\sim 150^\circ$ inclination), maximising the amplitude. SNR 24 is $\sim 3\times$ the threshold - $3\times$ farther and we'd have barely seen it. **100 years after the prediction, almost exactly.**

By B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration)



17 Aug 2017 · **GW170817**. We got lucky, again

Two neutron stars — $\sim 1.17 + 1.36 M_{\odot}$ — in NGC 4993 (Hydra), **40 Mpc** away. Signal in band for ~ 100 s before merger.

- **t = 0** GW170817 detected · H1 + L1 + Virgo
SNR: H1 \approx 26, L1 \approx 26, V1 \approx 2.5
- **+1.74 s** GRB 170817A · Fermi-GBM + INTEGRAL
First confirmed GW + short GRB association
- **+11 h** Kilonova AT2017gfo · 70+ telescopes
r-process nucleosynthesis confirmed — gold, platinum forged here

VERY LUCKY, AGAIN

O2 ended **8 days later** (Aug 25). All 3 detectors simultaneously online. Virgo's V1 SNR \approx 2.5 shrank the sky from thousands of deg² to **28 deg²**, crucial for EM follow-up. The jet was aimed roughly at Earth.

MULTI-MESSENGER ASTRONOMY · FIRST DETECTION

- **Gravitational waves** · LIGO + Virgo
- **Gamma-ray burst** · Fermi + INTEGRAL
- **Kilonova + X-ray + radio** · 70+ obs.

The gold in your jewellery was forged in events like this.

40 Mpc
NGC 4993 · Hydra

28 deg²
Sky localisation

Abbott et al. (LIGO/Virgo) 2017, PRL 119, 161101

Now we have a **catalog**.

~300

confirmed events

End of O4a. Mostly BBH.

O1 → O4

observing runs

Since Sept 2015. O5 coming soon.

5

detectors

LHO, LLO, Virgo, KAGRA, GEO600 (the OG).

1/2 day


candidate rate O4

Detection is routine. Science is now populations + precision GR tests.

PART 02 · THE PROBLEM

Signal : noise $\approx 1 : 1000$.

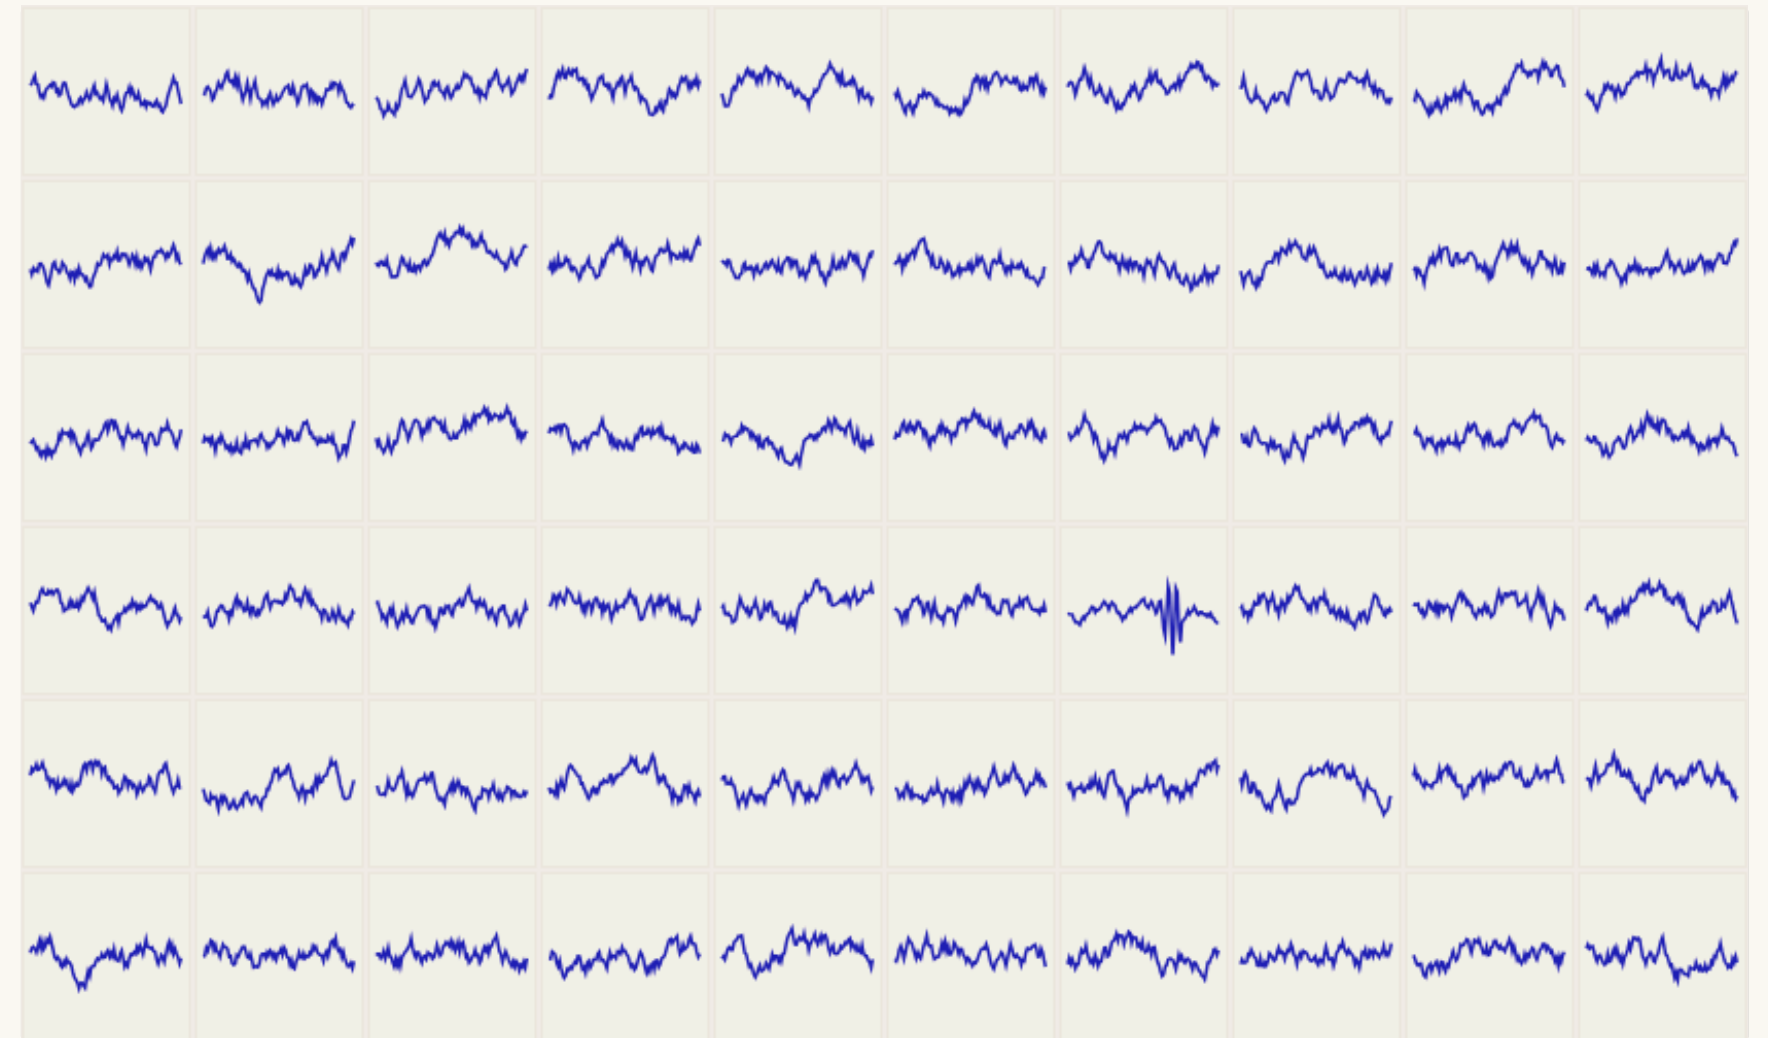
Raw strain data looks like nothing. Every seismic tremor, every laser glitch (every sneeze) – all louder than most real astrophysical events.



*Finding a chirp in raw detector data is like playing **Where's Waldo (Où est Charlie)** — except there's a million people per page, Waldo changes outfit every time, and you have a millisecond per page.*

We need a method to find signals of a *known shape* in noise. That method is: **matched filtering**.

60 DETECTOR SEGMENTS — ONE CONTAINS A REAL CHIRP. CAN YOU FIND IT?



Matched filtering finds it automatically .

Slide a known shape across the data.

STEP 01

Pick a template — a GR-predicted waveform

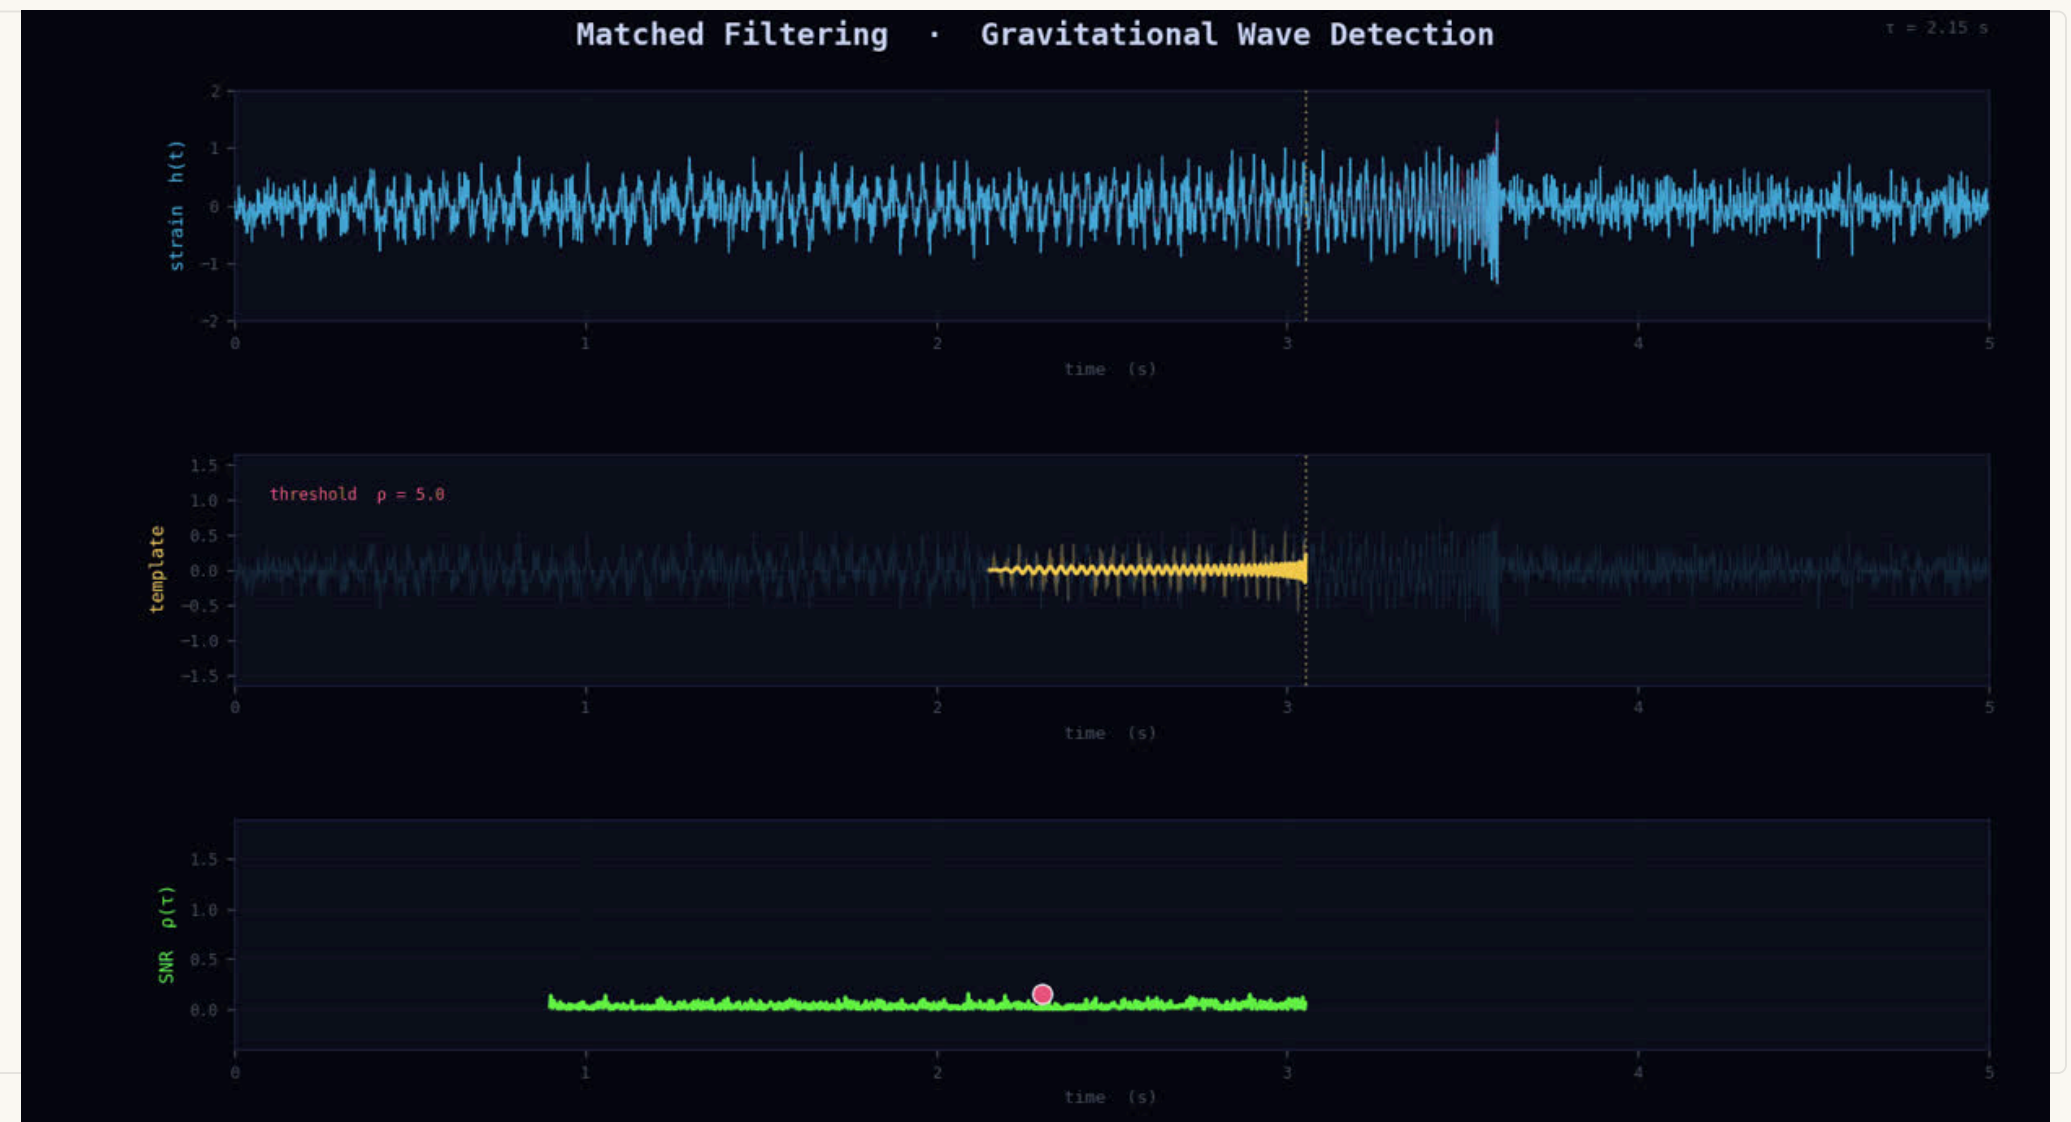
STEP 02

Correlate it with the data at every time shift

STEP 03

A sharp spike in the SNR output → chirp found

Under Gaussian noise, this maximises SNR for a known signal — the **Neyman–Pearson lemma** guarantees it.



The math, in two equations.

NOISE-WEIGHTED INNER PRODUCT

$$\langle a | b \rangle = 4 \operatorname{Re} \int_0^\infty \frac{\tilde{a}^*(f) \tilde{b}(f)}{S_n(f)} df$$

Weighted by the one-sided noise PSD $S_n(f)$. Frequencies where the detector is noisy count less.

SIGNAL-TO-NOISE RATIO

$$\rho(t_0) = \frac{\langle d | h_{t_0} \rangle}{\langle h | h \rangle}$$

Template h time-shifted by t_0 , correlated against data d , normalised. ρ is the detection statistic.

The catch

"Known signal shape." But the shape depends on parameters we **don't know**: masses, spins, orbital eccentricity...

So we don't filter with one template. We filter with **hundreds of thousands.**

NEXT PROBLEM

How many templates? Where to place them? How to do this fast enough to keep up with the data?

PART 02 · TEMPLATE BANKS

A template bank is a **grid** over parameter space (m_1, m_2).

The bank must be dense enough that any real signal falls within a small "mismatch" of some template.

Less mismatch: **highest recovered SNR**.

Too few templates → miss signals

Too many templates → can't run in time, too noisy

TEMPLATE PLACEMENT · (m_1, m_2) SPACE



Geometric

Stochastic

Hybrid

APPROACH	STRENGTH	WEAKNESS
Geometric	Fast (lattice)	Metric only locally valid — flat-space approximation breaks in curved param. regions
Stochastic	Flexible	Overcover — random placement wastes templates in dense regions
Hybrid	Both ✓	Hannah Montana energy — stochastic seed → geometric refinement

HDBank by S.Roy et al, <https://arxiv.org/pdf/1702.06771>

HDBANK-2: faster hybrid bank with by a **normalizing flow**.

01 · ENGINEERING

Parallelization

Jobs split by chirp mass, farmed via HTCondor. Sub-banks merged at the end.

[HTCondor](#) · [N parallel jobs](#)

02 · ALGORITHMIC

Smarter Placement

Dynamic matching with two parameters, extra random points, waveform caching.

[Dynamic match](#) · [Caching](#)

03 · MACHINE LEA

Normalizing Flow Seeding

Learn the target template density, sample from it. Smarter starting points → fewer wasted templates.

[NF](#) · [BNS](#) / [NSBH](#) / [BBH](#)

NF: Invertible differentiable maps that turn a Gaussian into any target density. We know the probability of every sample — ideal for placing templates where they're actually needed.

PART 03 · DOES IT WORK?

For sure  – about **50× faster** (10× using NF).

1 h 15 m

total wall time - 6 h (using NF) - vs >2day with standard HDBANK

Less Templates

10% less using NF, 7% less without it.

Minimal match improves

Tested over 9000 injections

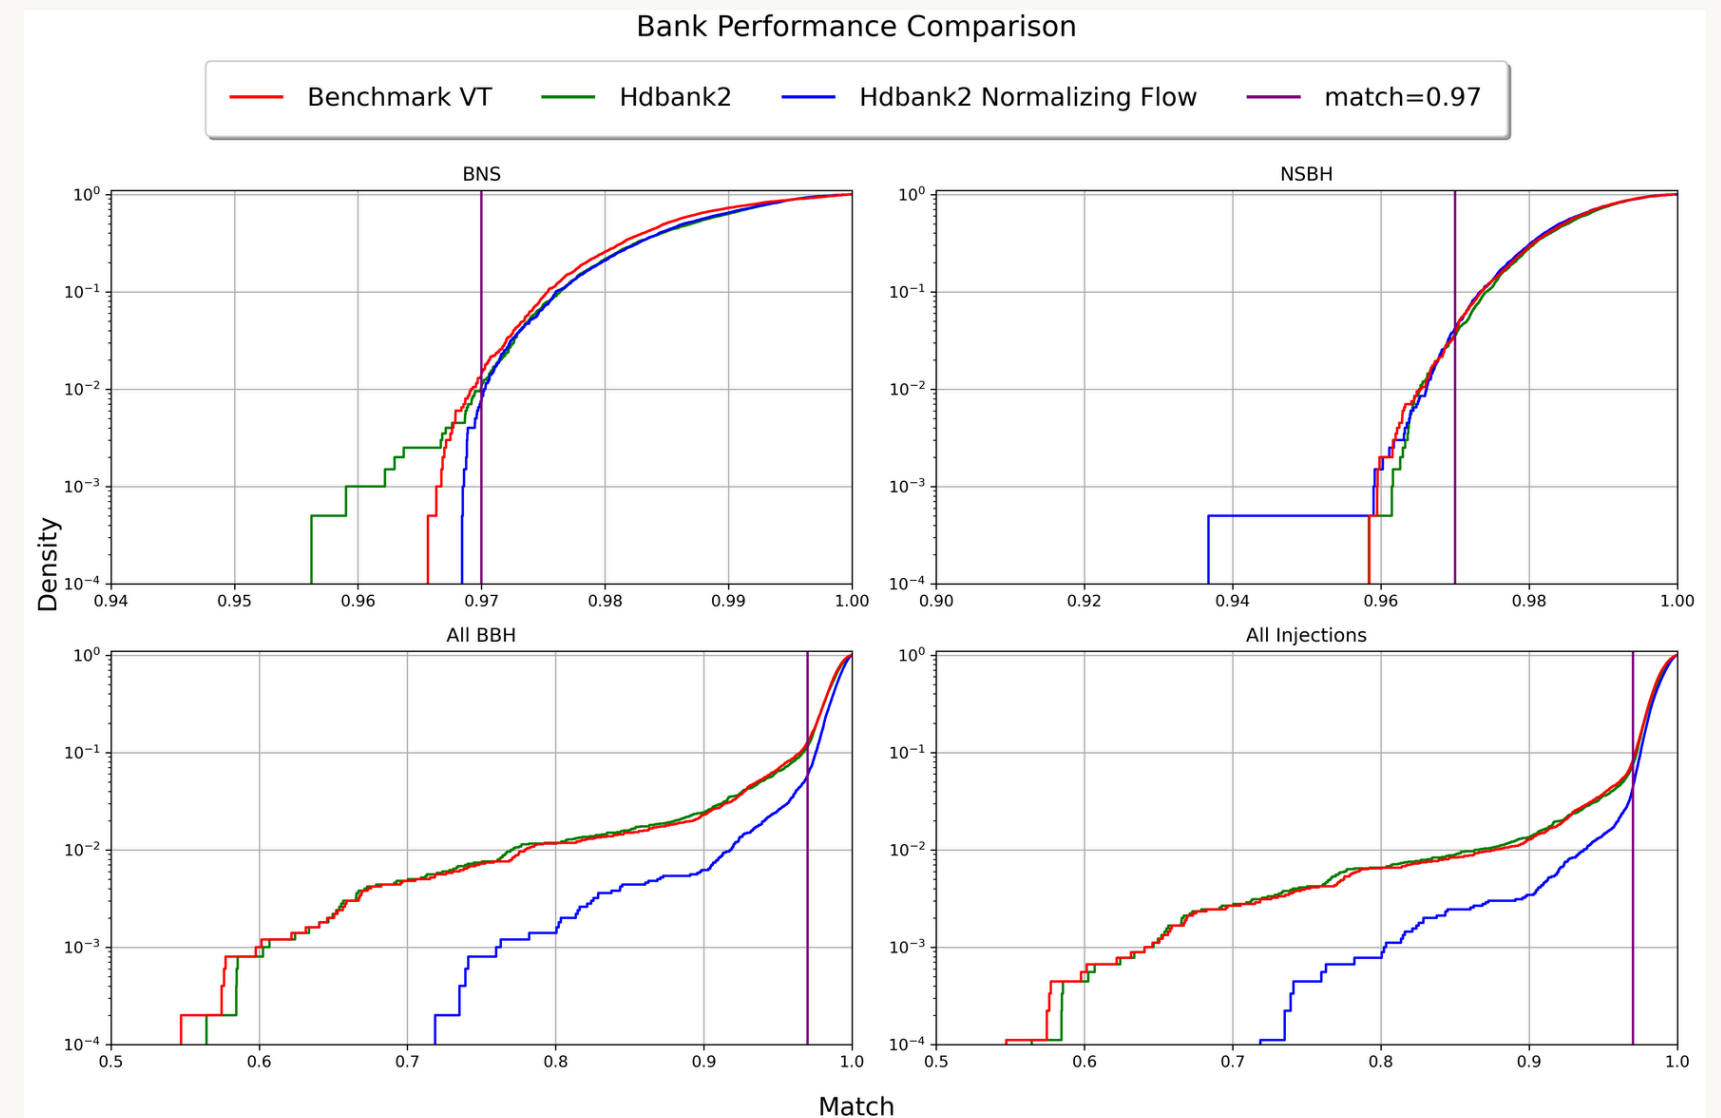
TAKEAWAYS

NF-seeded bank: fewer templates, better coverage.

Submittable to HTCondor out of the box.

ML actually helped — and we checked!

CUMULATIVE MATCH · BANK VERIFICATION · O4a VT NOISE CURVE



Thanks · Questions · Coffee

GWs → interferometers → matched filtering → template banks → HDBANK-2.

m.panzeri@campus.uniurb.it

Università di Urbino · INFN Firenze & Urbino · IP2I

HDBANK-2 · Panzeri, Fabrizi, Guidi