Gravitational-wave inference on time-frequency representations

Assessment of normalizing flows for BBH characterization

Universidad de Oviedo

D. Lanchares, O.G. Freitas,
J. González-Nuevo, J.A. Font

Vniversitat de València

June 23Th 2025

Main goals of our work

- Construct a pipeline that generates synthetic binary black hole data in a time-frequency representation.
- Train a general-purpose, image-based estimator of binary black hole parameters based on generated data, merging neural networks and normalizing flows.
- Test the trained model on suitable events of O3 (catalogues GWTC-2.1 and GWTC-3).

Why use images?

- Image-based learning is at the forefront of research, therefore performant architectures (e.g. ResNet) have already been developed.
- Compact binary coalescences have a distinctive “chirp” signature that models can learn to interpret.
- Frequency-specific noise can be isolated and consequently learnt from more easily.
- Current detector network lends itself nicely to RGB representation.

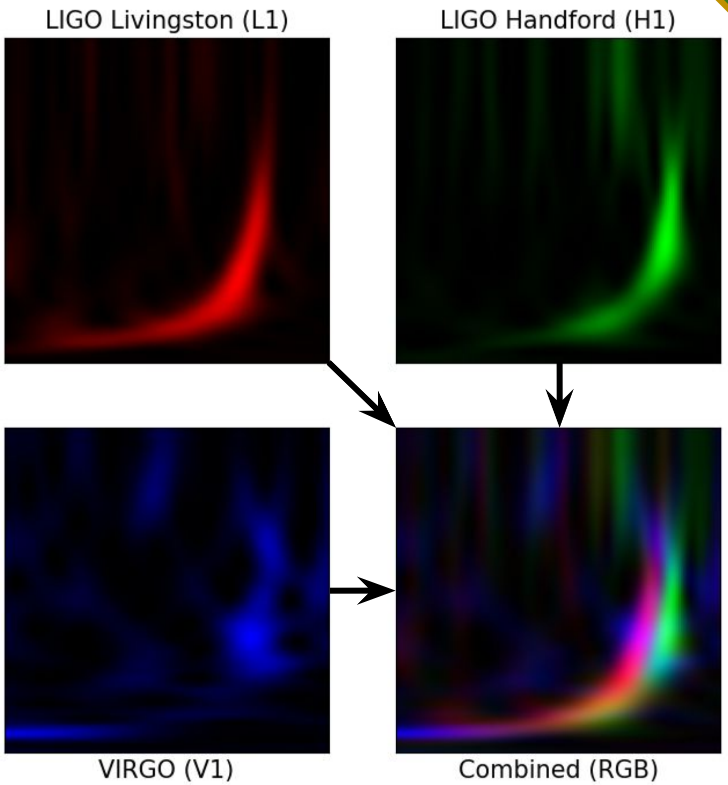

Injection from prototype pipeline

Modeling CBCs: Waveforms and Q-transforms

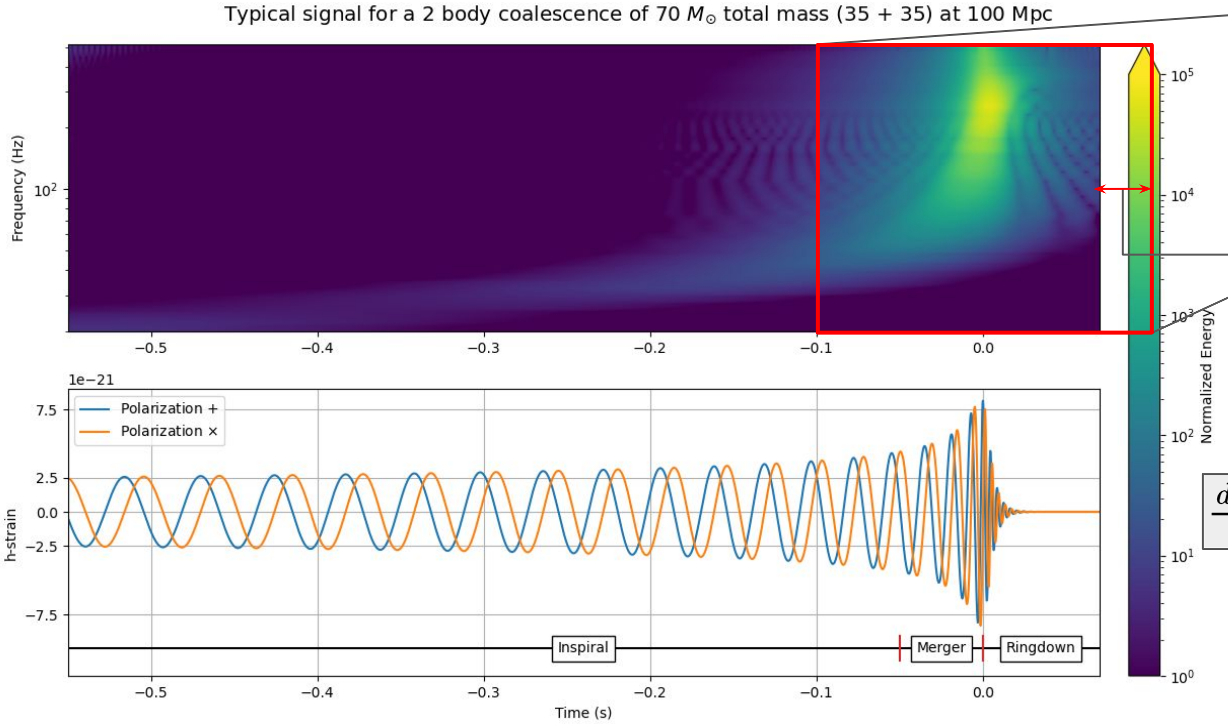

Image window for training

We center the window on the geocent time, so we expect delays. Hence the extra space

$$\frac{df_{GW}}{dt} = \frac{96}{5}\frac{G^{2/3}\boxed{\mathcal{M}^{5/3}}}{c^2}\pi^{8/3} f_{GW}^{11/3} + \mathcal{O}\left(\frac{1}{c^4}\right)$$

Chirp mass is the preferred mass representation

Our choice of parameters and priors

- Object's masses: M & q
- Object's spin angular momentum: a_i, θ_i, ϕ_{jl} & ϕ_{12}
- Luminosity distance and orbital inclination: θ_{JN} & d_L
- Phase of coalescence and polarization angle: ϕ_c & Ψ

Parameters: 8 + 4 (precession) = 12

CBC parameter	Prior
$\mathcal{M}$	$\mathcal{U}_{m_1,m_2}\{25, 100\}\mathrm{M}_\odot$
q	$\mathcal{U}_{m_1,m_2}\{\frac{1}{10}, 1\}$
a_i	$\mathcal{U}\{0, 0.99\}$
θ_i	$\mathcal{U}_{\sin}\{0, \pi\}$
ϕ_{jl}	$\mathcal{U}\{0, 2\pi\}$
ϕ_{12}	$\mathcal{U}\{0, 2\pi\}$
θ_{JN}	$\mathcal{U}_{\sin}\{0, \pi\}$
ψ	$\mathcal{U}\{0, \pi\}$
φ	$\mathcal{U}\{0, 2\pi\}$
SNR	$\mathcal{P}_{\alpha=-2}\{7, 30\}$

The prior for the chirp mass $\mathcal{M}$ and mass ratio q are uniform in the mass components, while the priors for θ_i and θ_{JN} are uniform in the sine. Luminosity distance is adjusted to achieve sampled SNR value.

Normalizing Flows for posterior estimation

Regression task:

Bayesian approach:
-Explicit prior and likelihood.
-Per-event analysis cost ~ hours.

Bayes Theorem

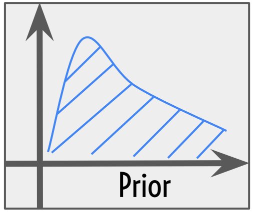

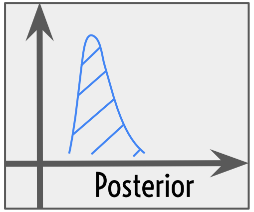

Deep learning

Our approach:
-Likelihood free.
-Prior implicit in training data.
-Per-event analysis cost ~ seconds.

- Normalizing Flows:
  - Fast to sample from.
  - Train like regular networks.
  - Inherently statistical.
  - Equipped for multimodality.
  - Modular in nature (easy to upgrade).

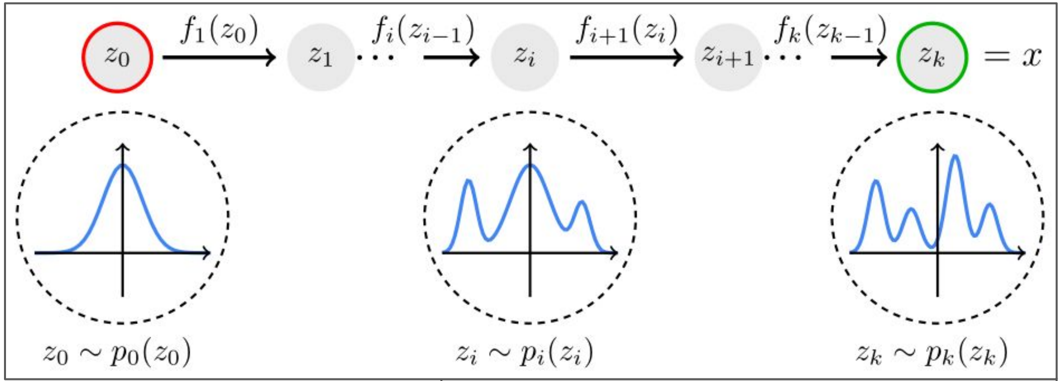

Adapted from https://tikz.net/normalizing-flow/

Our GP12 net & flow model

- Embedding ResNet:
    - PyTorch ResNet-18 implementation.
    - Trained from scratch.
- Normalizing Flow:
    - Similar structure to DINGO.
    - Using glasflow's nflows fork.
    - Base: RQ-coupling + MAF, with (2+2) hidden layers per block.
    - Middle, final: Perm. + LULinear.
    - 16 flow steps (repetitions) and 5 blocks: 16·5·4 = 320 hidden layers.
- Scaling parameters to ~ [0,1]:
    - Downscaled at training time, upscaled at sampling time.

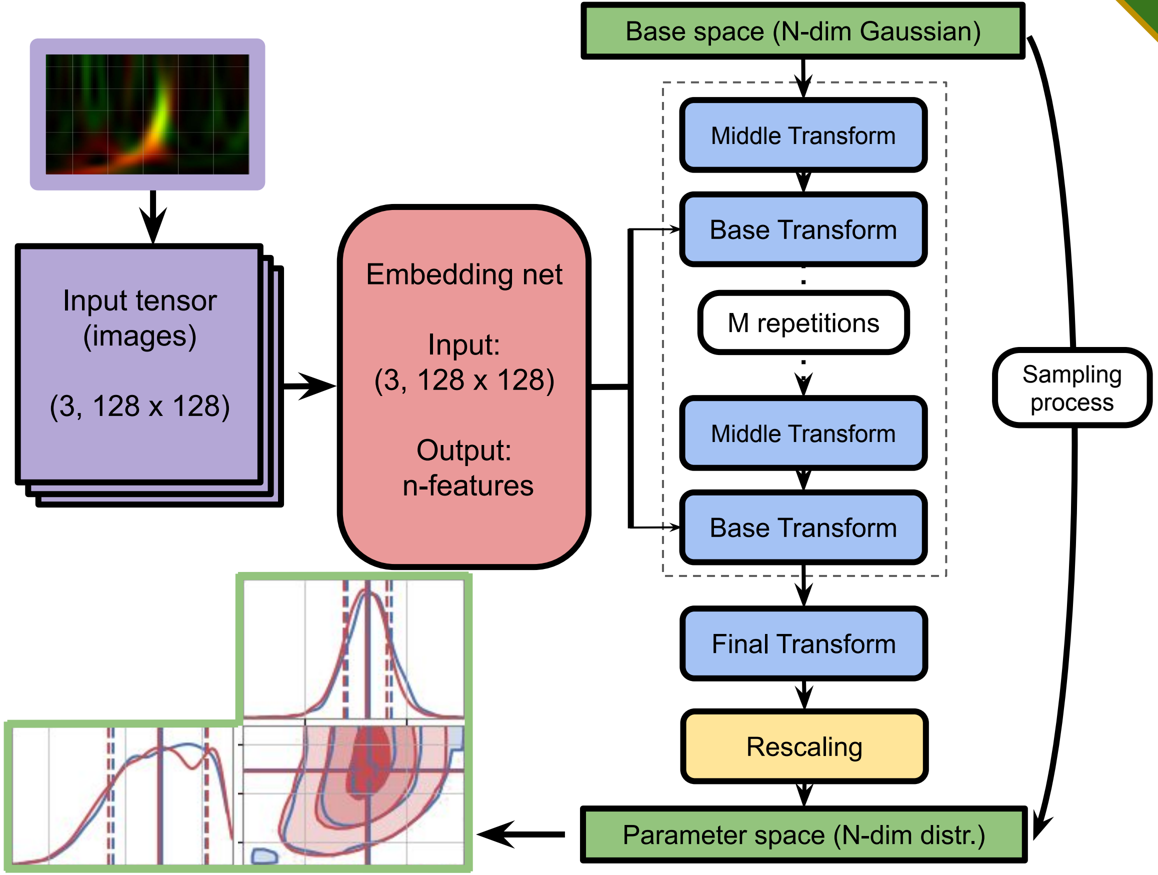

Catalogue comparison I: GWTC-2.1

- Filtered events: 24
  - GWTC-2.1 & GWTC-3
  - Within our priors.
  - 3-detector configuration.
- Publicly released samples:
  - IMRPhenomXPHM approximant.
- Metric:
  - JS Divergence.
  - In nat: [0, ln(2)]

	$\mathcal{M}$	q	a_1	a_2	θ_1	θ_2	ϕ_{JL}	ϕ_{12}	d_L	θ_{JN}	ϕ	Ψ
GW190413_134308	9.0	0.9	1.4	0.4	6.4	0.1	0.6	0.2	14.5	8.2	1.1	0.4
GW190503_185404	25.8	12.5	3.7	0.6	0.9	0.3	0.9	0.3	9.8	36.7	14.8	0.2
GW190513_205428	50.6	7.0	0.8	0.8	8.1	0.8	0.3	0.5	18.0	16.0	4.1	0.2
GW190517_055101	12.5	2.8	22.5	0.3	19.6	0.1	2.9	0.4	3.2	9.2	17.6	3.5
GW190519_153544	16.0	5.5	1.4	0.3	49.4	16.4	1.2	0.2	32.8	6.1	0.5	6.3
GW190602_175927	2.8	6.3	4.5	0.8	18.7	6.1	0.6	2.4	8.0	32.7	2.8	0.2
GW190701_203306	30.0	14.1	6.9	0.4	0.3	0.8	0.7	0.2	21.7	45.3	2.2	0.2
GW190706_222641	4.6	8.2	5.1	0.1	13.9	0.1	0.5	0.1	13.6	11.8	28.5	0.4
GW190727_060333	47.4	5.0	0.9	1.8	23.9	4.7	1.1	1.5	7.5	16.3	0.8	0.3
GW190803_022701	21.0	23.9	1.5	0.7	5.8	1.4	0.8	0.2	27.2	21.7	2.4	0.1
GW190828_063405	50.4	22.0	0.2	1.5	13.3	3.6	0.9	1.2	14.6	24.6	10.7	3.5
GW190915_235702	60.7	6.0	1.1	1.0	14.6	1.6	2.3	2.7	14.8	20.2	4.6	0.4
GW190916_200658	0.8	1.0	0.2	0.1	8.4	2.6	0.1	0.0	49.5	10.1	2.6	0.1
GW190926_050336	5.7	1.9	7.1	0.5	0.5	0.7	0.2	1.4	31.1	1.6	1.6	1.7
GW190929_012149	10.5	8.2	18.7	2.4	9.0	0.2	0.9	3.0	32.9	12.6	2.4	2.1

JS Divergence$[\times 10^{-2}$ nat$]$ (colour scale 0–70)

Catalogue comparison II: GWTC-3

	$\mathcal{M}$	q	a_1	a_2	θ_1	θ_2	ϕ_{JL}	ϕ_{12}	d_L	θ_{JN}	ϕ	Ψ
GW191127_050227	10.3	8.2	3.4	0.1	16.6	3.4	0.6	0.3	25.4	21.9	1.0	0.3
GW191230_180458	14.7	0.9	0.4	0.2	0.7	0.7	0.3	0.5	5.2	5.8	2.7	0.6
GW200129_065458	40.2	5.9	16.9	1.3	6.6	2.1	16.6	1.2	17.2	54.8	10.9	2.5
GW200208_130117	45.8	5.7	0.2	0.7	0.7	0.5	0.6	0.4	13.8	34.0	1.8	1.1
GW200209_085452	36.6	6.1	0.8	0.2	0.5	0.3	0.8	0.6	17.5	18.3	4.0	1.0
GW200216_220804	3.1	2.0	0.4	0.2	11.0	1.6	1.9	0.6	29.1	21.7	0.3	0.1
GW200219_094415	43.6	4.2	8.6	1.4	7.3	0.2	0.7	1.1	24.6	15.0	4.6	1.5
GW200224_222234	5.3	0.5	1.4	0.1	3.9	0.6	10.6	0.4	42.9	22.2	4.3	4.4
GW200311_115853	57.5	12.9	7.1	2.4	18.8	0.8	3.7	0.4	23.7	40.9	11.4	2.3

JS Divergence $[\times 10^{-2}$ nat$]$ (colour scale 0–70: 0, 10, 20, 30, 40, 50, 60, 70)

GW191230

- Ideal scenario.
  - Visually clear chirp signal.
  - Median of our mass prior.
- All estimated values are within the 90% regions of the LVK posteriors.
  - Accurate mass estimation
  - Strong agreement on spin parameters.
  - Distance is somewhat underestimated, obfuscating inclination degeneracy.

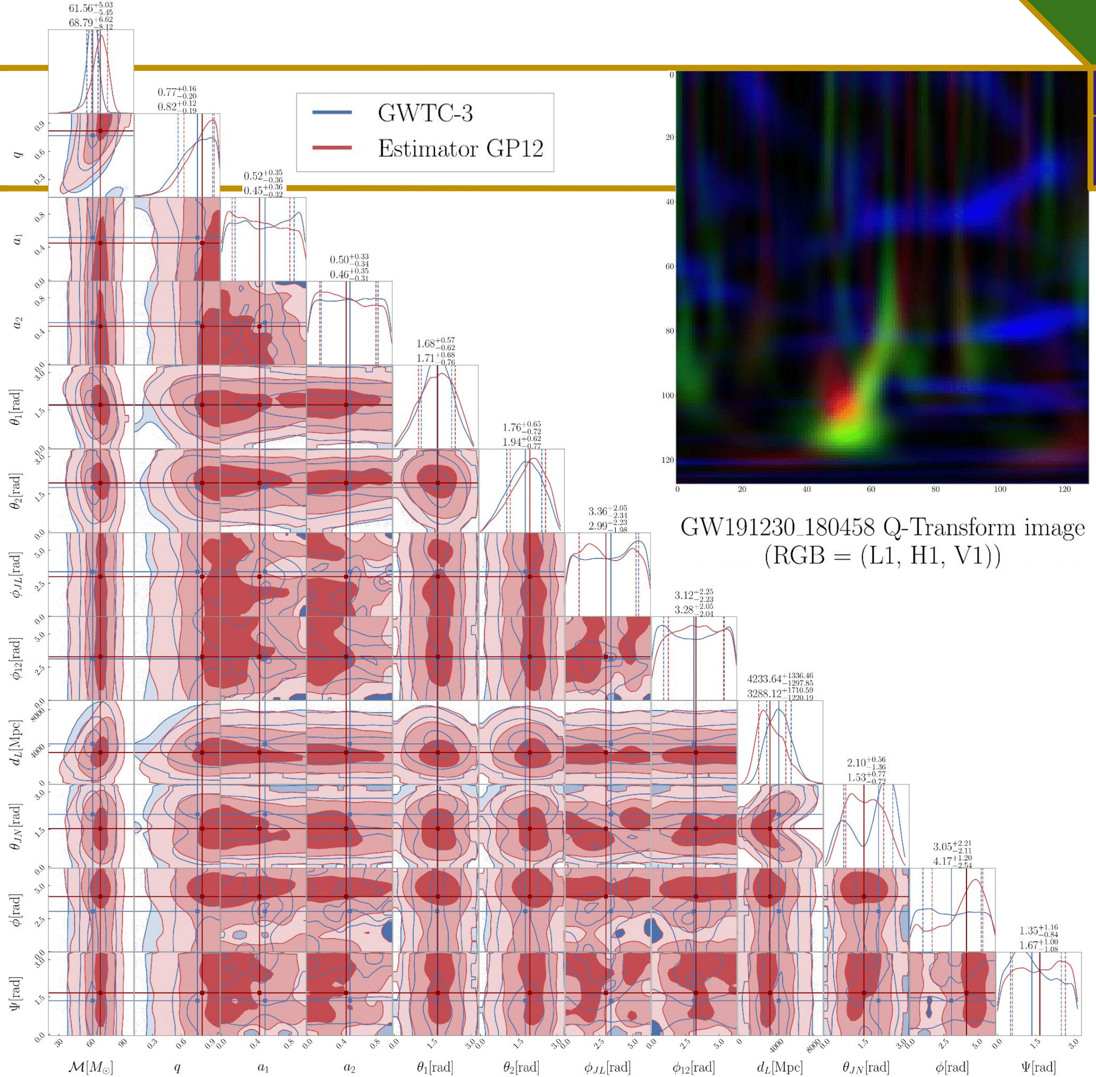

GW191230_180458 Q-Transform image (RGB = (L1, H1, V1))

To conclude

- These technologies can accurately model the high dimensional parameter space of a BBH.
- Estimations are largely compatible with reference data, especially where the latter is hesitant.
- Poor performances were found mainly in the estimation of the chirp mass, albeit only for low-mass events with chirp mass ~ $30M_\odot$, and of the luminosity distance, which our model tends to underestimate.
- A fast-training, simple model as GP12 can produce large amounts of posterior samples of an unknown target distribution in the order of tenths of seconds (0.5 s/10^4 samples on V100 GPU).
- These models will become necessary as detection rates increase.

Future work

- Various open fronts:
  - Generalizing architecture to include auxiliary PSD data.
  - Incorporating time:
    - Training around t_c for event-specific model.
    - Developing mod 24h time for generic model.
  - Using SNR as a parameter to gauge network sensitivity.
  - Multi-band analysis of sources with particularly complex or less well-modelled frequency evolutions.

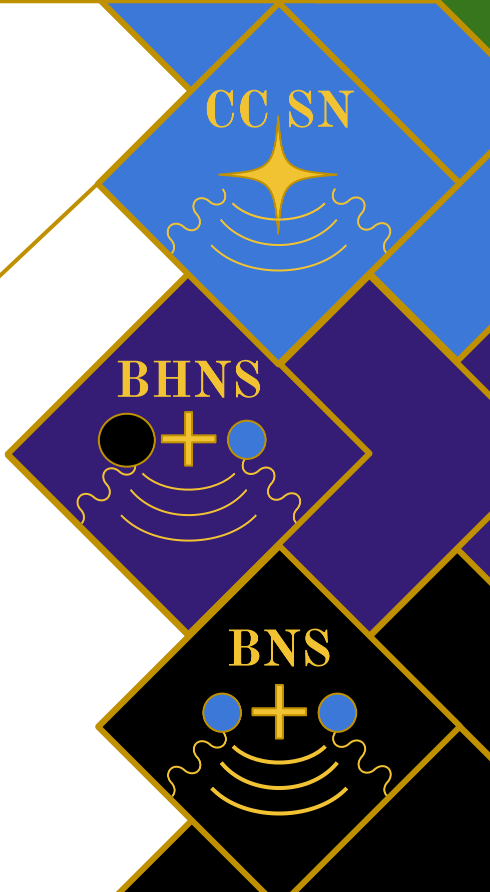

Notes and disclaimers

- This presentation stems from research carried out for the presenter's Master Thesis, detailed in "*Assessment of normalizing flows for parameter estimation on time-frequency representations of gravitational-wave data*" (https://arxiv.org/abs/2505.08089).
- The code to reproduce these results is available at https://github.com/Daniel-Lanchares/dtempest.

Thank you for your attention