Denoising gravitational wave signals with a variational autoencoder

P. Bacon [1] M. Bejger [2] A. Trovato [1]

[1] Laboratoire Astroparticule et Cosmologie (CNRS), Paris
[2] Nicolaus Copernicus Astronomical Center, Warsaw

March 5, 2020

APC, CNRS, FRANCE

Université de Paris

What this presentation is about...

Remove noisy components from injected GW signals (compact binaries) in real interferometer strain data.

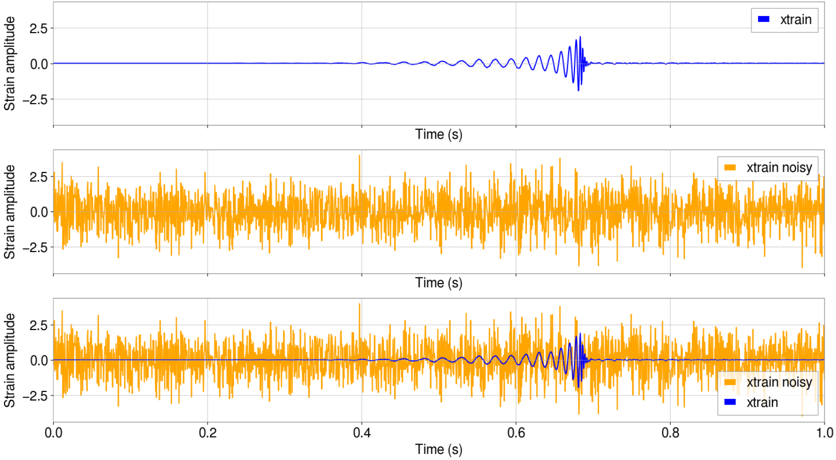

Context

- Actual gravitational waves (GW) searches for compact binary coalescence (CBC) signals mainly rely on the **gaussian noise hypothesis**. How about dealing with the non gaussian part ?
- Low-latency searches are indispensable as the detection rate is expected to increase in next generation instruments (electromagnetic follow-up).
- Model based searches are optimal. However a model is not available for all GW sources (parameter space is partially covered).

$\rightarrow$ **Deep learning** (DL)

Convolutional denoising autoencoders (DAE)

The novel approach we propose:

- Recent applications of DL in GW astronomy involve classifiers $\rightarrow$ **regression problem** (denoising)
- Rule of thumb: use recurrent networks (ex: LSTM) for timeseries and use convolutional networks for images
  $\rightarrow$ use **1D convolutional network** to perform on strain data from GW detectors.
- Point estimate is useless $\rightarrow$ **Bayesian framework** offers a probabilistic interpretation.
  Uncertainty is the key ingredient to inference/decision making.

Data analysis & sparsity

Usually DAE are bottleneck-shaped so as to enforce a **sparse representation**.

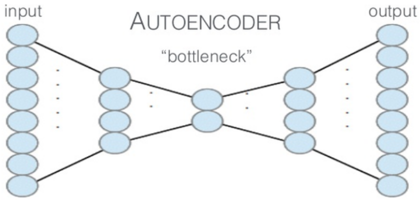

Credits: quora.com

- bottleneck enforces DAE to perform a dimension reduction
- Sparsity is crucial when dealing with noise:
  high coeffs dictionnary elements are **less prone to noise fluctuations**.

Variational autoencoders (VAE)

Add Bayesian framework on top of it:

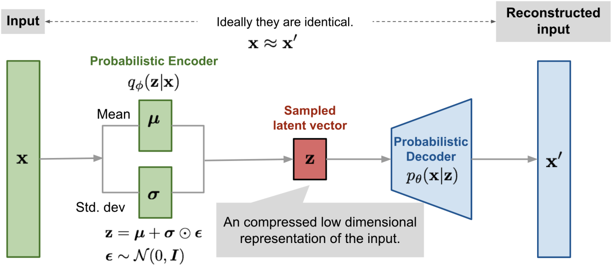

Credits: lilianweng.github.io

And minimize the loss function:

$$L_{\mathrm{BETA}}(\phi, \beta) = -\mathbb{E}_{\mathbf{z}\sim q_\phi(\mathbf{z}|\mathbf{x})} \log p_\theta(\mathbf{x}|\mathbf{z}) + \beta D_{\mathrm{KL}}(q_\phi(\mathbf{z}|\mathbf{x}) \| p_\theta(\mathbf{z}))$$

Decoder (reconstruction term)

Encoder (regularisation term)

Dataset, training & prediction

Dataset:

- Injections: GW signals with $f_{low} = 30Hz$
  indiv. masses in $[10, 30]M_{\odot}$ signal-to-noise ratio (SNR) in $[5, 20]$.
- Input: whitened GW signals + real interferometric O1 noise
- Output: whitened GW signals

Training

- 3x100 epochs with three distinct chunks of 1000 injections each.
- Flat signal probability: make method robust to near gaussian noise.
- Prevent overfitting: low learning rate & monitor train/test losses.
- Traing time: $\sim$ 0.5d on `AMD Ryzen 7 PRO` CPU

Predicting by passing N times the same noisy signal to the VAE then compute μ and σ. $\rightarrow$ Equivalent to predicting distributions

Results - SNR=17

$m_1 = 21 M_\odot, m_2 = 15 M_\odot$

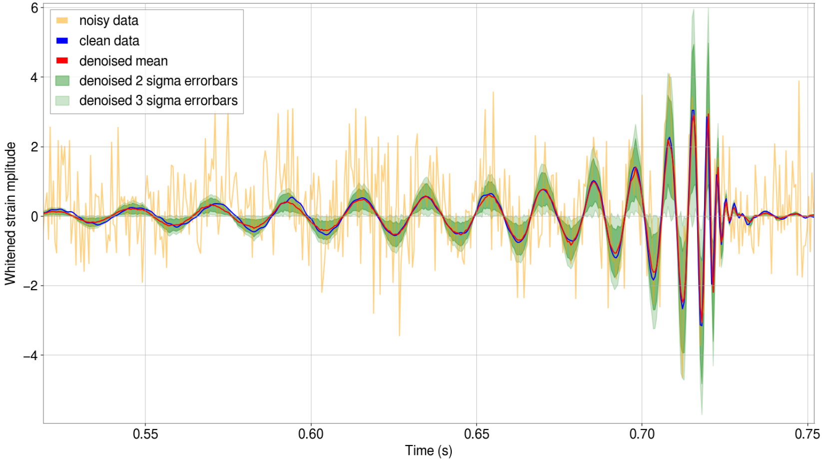

Results - SNR=10

$$m_1 = 11M_\odot, m_2 = 15M_\odot$$

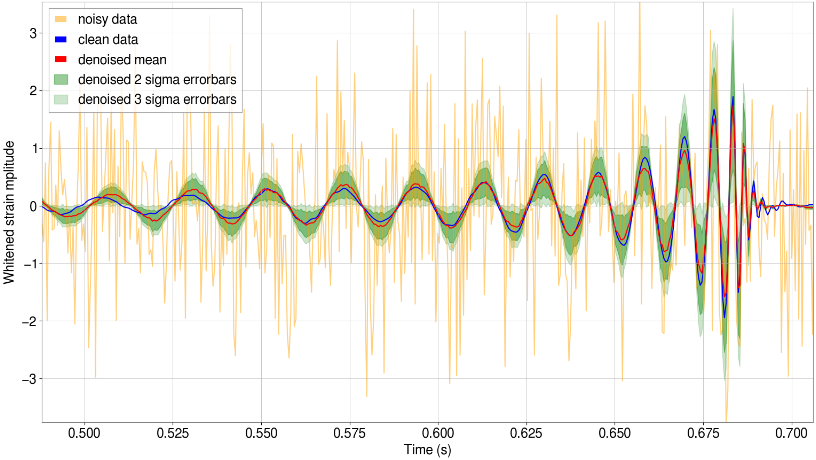

You said uncertainties ?

From theory:

Output probability of a VAE is the probability of the data $\mathbf{x}$ *knowing the learnt latent representation* $\mathbf{z}$

but it leads to the following remarks:

- In the context of denoising: $q_\phi(\mathbf{z}|\mathbf{x}) = \mathbf{q}_\phi(\mathbf{z}_{\textbf{noisy}}|\mathbf{x}_{\textbf{noisy}})$ and subsequently $p_\theta(\mathbf{x}|\mathbf{z}) = \mathbf{p}_\theta(\mathbf{x}|\mathbf{z}_{\textbf{noisy}})$
- VAE loss function design suggests uncertainty is driven by the parameter β: make it trainable ! (ongoing work)

  ex: previous example with $SNR = 17$ has
  $p(signal \in 2\sigma region) = 50\%$ and $p(signal \in 3\sigma region) = 74\%$
  $\rightarrow$ **How to interpret this ?**

Highlights from round table

- Gaussian models for decoder output may not fit well with real posterior distribution
- Flat probability in training may depopulate some regons of the parametr space: increase aleatoric uncertainty
- Looking forward to invetigating tensorflow probability

Conclusions & References

Take away ideas:

- Convolutional layers are successfully applied to regresion problems involving timeseries.
- VAE elegantly combine deep learning efficiency and the Bayesian framework.
- Marge for making β trainable and see whether it helps in interpreting.

References:

- Kingma, D. P. & Welling, M. *Auto-Encoding Variational Bayes*. arXiv preprint arXiv:1312.6114, 2013
- Higgins *et al.*, *beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework*, 2017

Backup slide: LIGO O1 events

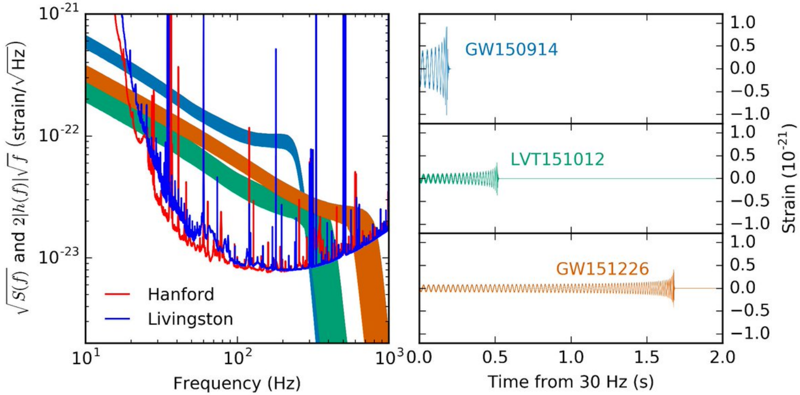

Optimal signal-to-noise ρ : $$\rho^2 = \int_0^\infty \left(\frac{2|\tilde{h}(f)|\sqrt{f}}{\sqrt{S_n(f)}} \right)^2 d\ln(f)$$

(GW150914: $\rho \simeq 24$, GW151226: $\rho \simeq 13$, LVT151012: $\rho \simeq 10$)