

Training and Deploying a Neural Network for Noise Regression in Gravitational Wave Astronomy

Presenter: Alec Gunny ^{1,3}

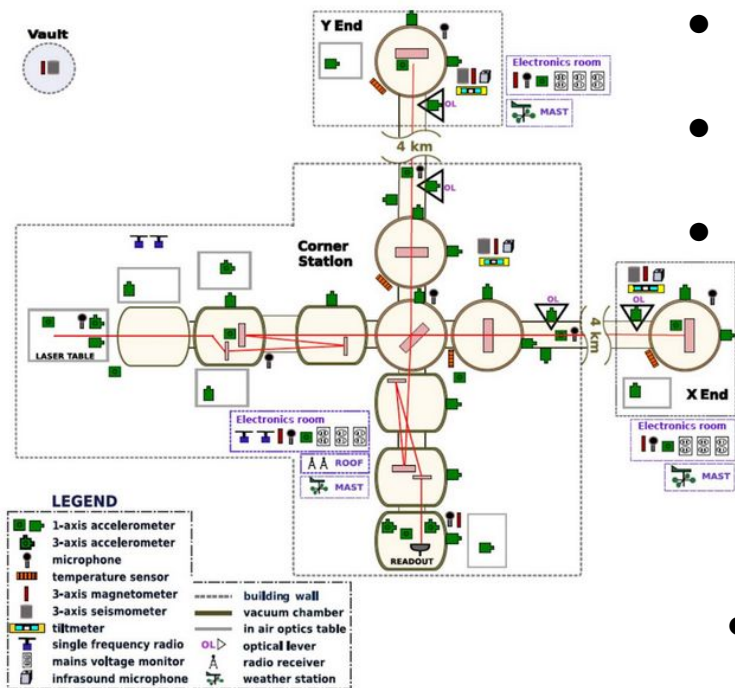
Rich Ormiston ², Michael Coughlin ², Tri Nguyen¹, Erik Katsavounidis ¹, Philip Harris ¹, Dylan Rankin ¹, Jeffrey Krupa ¹, Sang Eon Park ¹, Ethan Marx ¹, Satya Mohapatra ¹, et al.

¹ Massachusetts Institute of Technology

² University of Minnesota


³ Formerly NVIDIA

Gravitational Waves and LIGO



LIGO and Virgo Collaborations, *CQG* **33**, 134001 (2016)

- Large scale astrophysical events cause distortions in spacetime known as gravitational waves
- Tiny amplitude of these distortions makes them difficult to detect
- LIGO - pair of enormous interferometers that use destructively interfering lasers to measure perturbations in spacetime



- Measurement of distortion typically given by unitless quantity “strain”, related to relative change in displacement of objects caught in the wave
- Inferred from intensity of photons detected as GWs distort laser paths and bring them in-phase

Noise, MMA, and FastML

- Environmental noise can degrade the perfect destructive interference of the lasers
- Leads to spurious photon detection, leads to noisy strain measurements
- Makes it difficult to pick out signals with amplitude less than noise, limits detection range
- Auxiliary sensors measure noise for removal


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)

Initial and Enhanced LIGO
archive (2002-2010)
exceeds 1PB of data



Figure 1

<http://www.ifae.es/eng/magic-gallery.html>



<https://www.nasa.gov/content/fermi/overview>

<https://www.ligo.caltech.edu/image/ligo20150731e>




<https://icecube.wisc.edu/gallery/press/view/1336>

<http://www.ung.si/en/research/cac/projects/auger/>

Meszáros et. al. <https://arxiv.org/pdf/1906.10212.pdf>

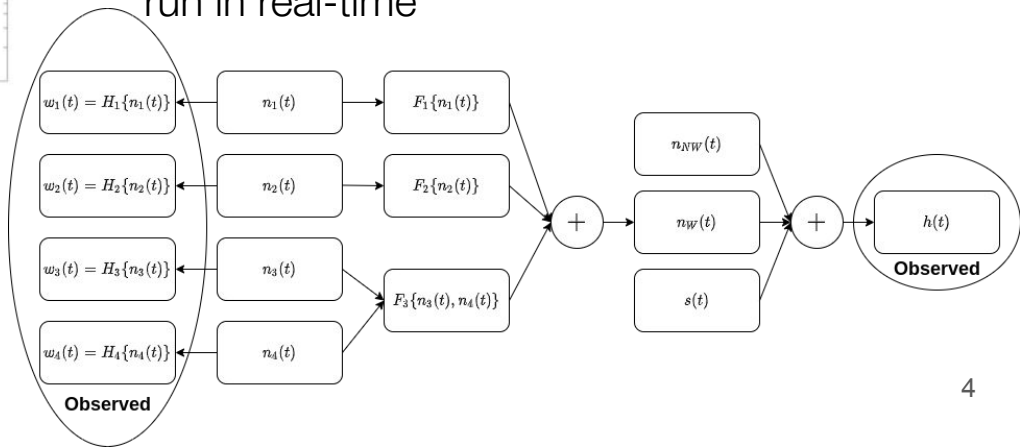
- Multi-messenger astrophysics offers promising insights by comparing different cosmic messengers from same phenomena
- LIGO + VIRGO critical for detecting and locating events to alert other observers
- Noise subtraction and downstream algorithms need to work in real-time to capture as much data as possible

Low Frequency Noise

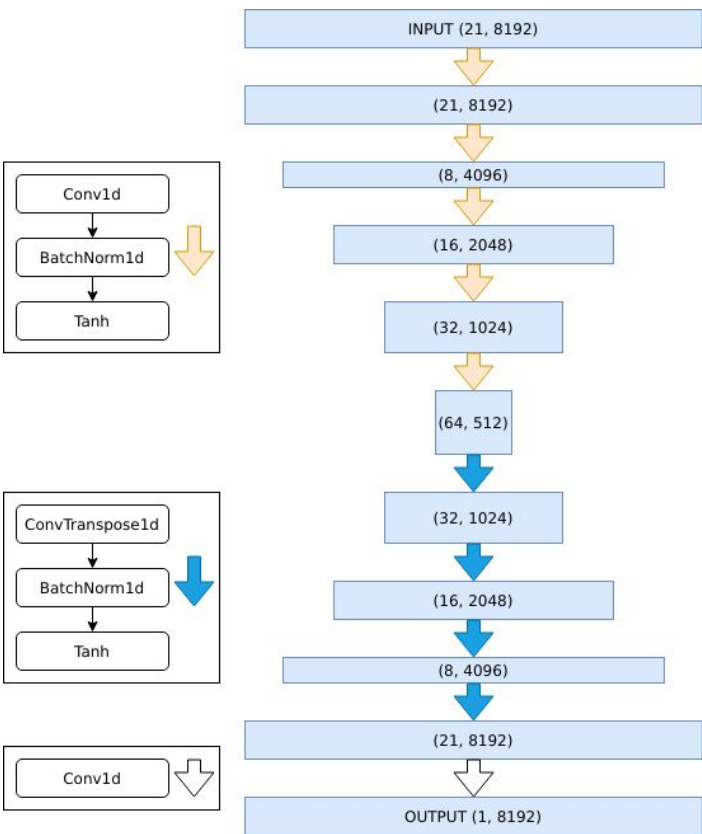


- Unmodelled noise sources below 100 Hz

- Noise is coupled with auxiliary channel measurements, astrophysical signal isn't
- Use auxiliary channels to regress to observed strain noise
- Previously proposed techniques e.g. Wiener filter have limited expressivity and/or can't be run in real-time



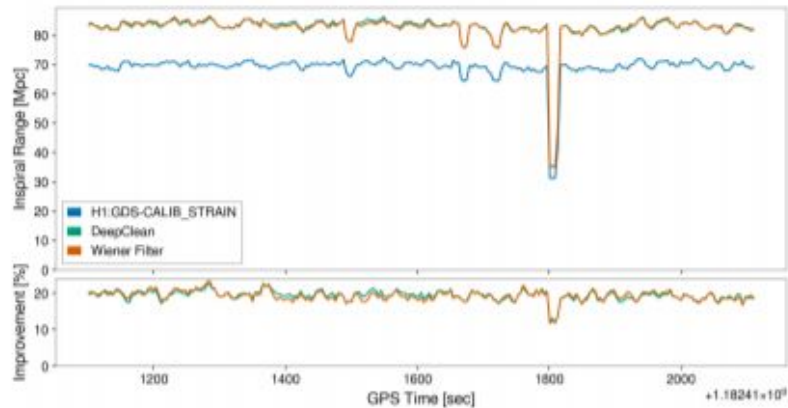


Deep Learning for Noise Regression



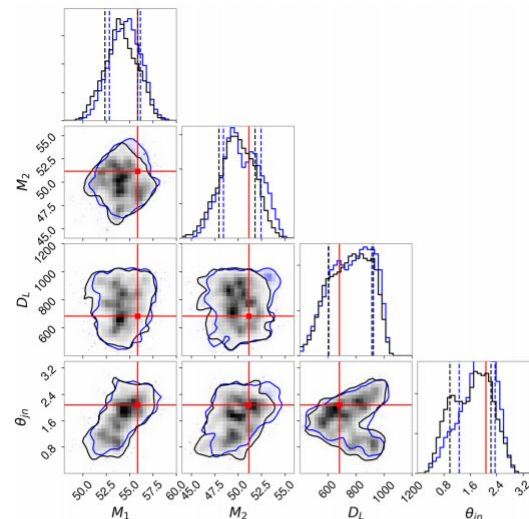
- Neural network can learn complicated nonlinear interactions between auxiliary or “witness” channels
- Fully convolutional network maps from witness measurements to noise estimate
- Regress noise estimates to strain measurements $h(t)$ since signal is independent
- In practice, wide disparities in contributions of various noise sources. Normalize MSE in frequency space by ASD of $h(t)$ to compensate

$$J(r) = w \frac{1}{M} \sum_{i=1}^M \sqrt{\frac{S_{[r,r]}[i]}{S_{[h,h]}[i]}} + (1-w) \sum_{i=1}^N r[i]^2 ; r = h - \mathcal{F}(\vec{w}; \theta)$$

Validation

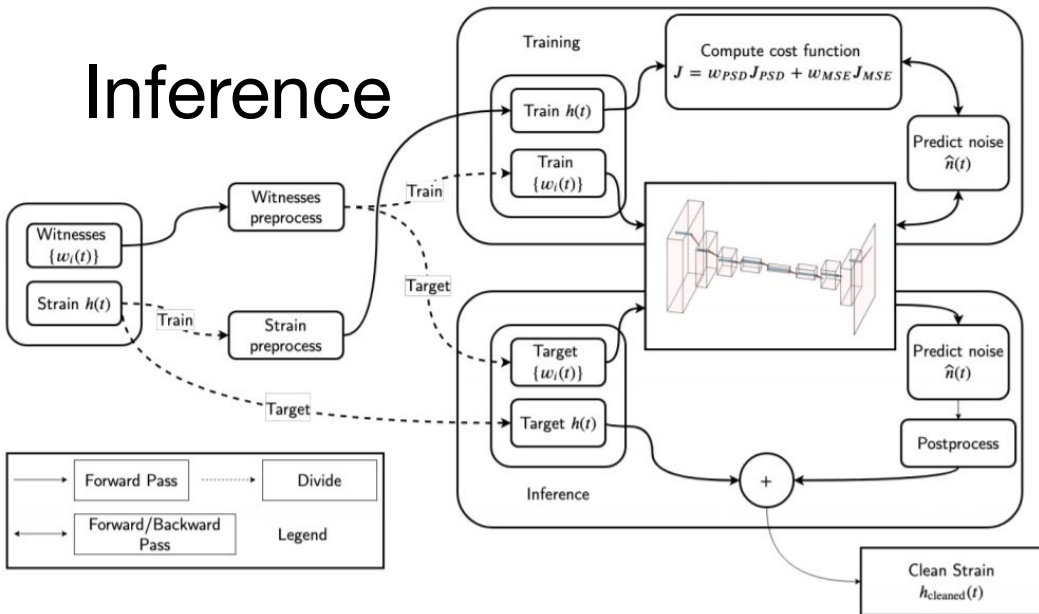


- No corruption of astrophysical signal



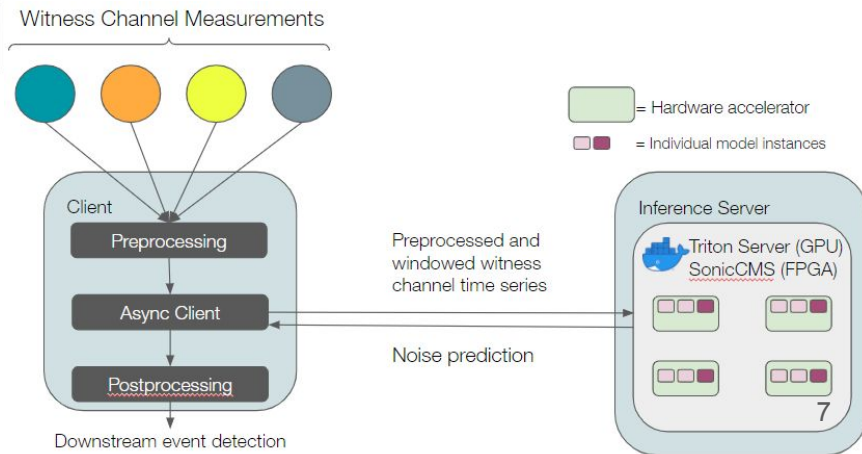
- Consistency with existing explicitly modelled noise removal mechanisms

Inference



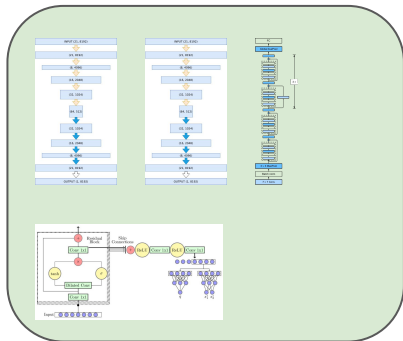

- Resample and center witness data
- Use trained model to estimate noise
- Uncenter and bandpass filter
- Subtract from strain data

- Implement steps as asynchronous processes to maximize throughput
- Implement model inference on dedicated inference server using accelerated hardware/software




Inference-as-a-Service

- Portable




- Framework and architecture agnostic
- Critical for applications like DeepClean that need frequent retraining




- Co-locate downstream models for better resource allocation/autoscaling
- Manage and accelerate end-to-end latency of DeepClean + downstream algorithms to meet requirements

Inference-as-a-Service - DeepClean challenges




- Frame width dictated at train time

Inference-as-a-Service - DeepClean challenges



- Stride between subsequent frames is an inference time parameter that affects estimation quality and arrival rate to inference pipeline
- If pipeline throughput can't meet arrival rate, frames pile up and queue latency explodes
- Current pipeline running with 2 copies of model, frame stride of 2 ms, batch size of 8, achieving throughput of ~450 frames / s
- Working on custom backend for streaming as well as tools to explore cost landscape



- Batching subsequent frames linearly reduces arrival rate
- Introduces unavoidable latency
- Makes streaming picture nontrivial: send duplicate data or build custom backend to batch on server side
- High throughput ML inference critical to mitigating these issues

Summary

- DeepClean low-frequency noise estimation can increase our capacity to detect and analyze astrophysical events
- Inference-as-a-service deployment represents a powerful model for accelerating the pace at which new architectures and applications can be adopted
- Further optimization and tools for exploring the relevant parameters will allow each use case to fit their own latency/throughput/cost constraints

Thank You
