LASER INTERFEROMETER GRAVITATIONAL WAVE OBSERVATORY
- LIGO -
CALIFORNIA INSTITUTE OF TECHNOLOGY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Technical Note **LIGO-P2600326-v1** 2026/08/20

Recovering Gravitational-Wave Sensitivity During LIGO Earthquake Mode Through Witness-Channel Noise Subtraction

Matilda Damon — Mentored by Elenna Capote
California Institute of Technology / LIGO Hanford Observatory

California Institute of Technology
LIGO Project, MC 100-36
Pasadena, CA 91125
Phone (626) 395-2129
Fax (626) 304-9834

Massachusetts Institute of Technology
LIGO Project, Room NW22-295
185 Albany Street
Cambridge, MA 02139
Phone (617) 253-4824
Fax (617) 253-5474

LIGO Hanford Observatory
127124 N Route 10
Richland, WA 99354
Phone (509) 372-8106
Fax (509) 372-8137
E-mail: outreach@ligo-wa.caltech.edu

LIGO Livingston Observatory
19100 LIGO Lane
Livingston, LA 70754
Phone (225) 686-3100
Fax (225) 686-7189
E-mail: outreach@ligo-la.caltech.edu

https://www.ligo.caltech.edu/

Abstract

Earthquakes produce large ground motion that can threaten interferometer stability and reduce gravitational-wave sensitivity. The Advanced Laser Interferometer Gravitational-Wave Observatory (LIGO) uses Earthquake Mode to reduce lock loss during increased seismic activity, but the modified control configuration can introduce or amplify instrumental noise in the calibrated strain signal. This project investigated whether auxiliary control channels could be used to identify and subtract noise coupling into strain during Earthquake Mode at the LIGO Hanford Observatory. I developed a reproducible Python analysis framework and tested three approaches: a custom transfer-function subtraction, GWSubtract linear time-domain subtraction, and the non-stationary noise-subtraction framework (NonSENS),[7] while using coherence, spectral diagnostics, and binary neutron star (BNS) range to evaluate performance. The strongest completed subtraction increased the BNS range from 75.19 Mpc to 86.66 Mpc, an improvement of 11.47 Mpc or approximately 15.3%, while the custom analysis identified the Michelson degree-of-freedom control signal as the dominant linear witness. The best retained NonSENS result reached a BNS range of 83.56 Mpc and demonstrated that model bandwidth and output conditioning strongly affect subtraction performance. These results show that a significant fraction of the Earthquake Mode noise in the analyzed interval was linearly witnessed and removable offline, while also demonstrating the importance of evaluating subtraction using full-band astrophysical sensitivity rather than a single spectral measure.

Contents

Abstract i

1 Introduction 1

2 Background 1
2.1 Seismic isolation, interferometer controls, and Earthquake Mode . . . 1
2.2 Witness-channel noise subtraction . . . 1

3 Objectives 2

4 Methodology 2
4.1 Analysis overview . . . 2
4.2 Earthquake Mode and quiet-reference data . . . 2
4.3 Control-system witness channels . . . 3
4.4 Amplitude spectral density . . . 3
4.5 Coherence analysis . . . 4
4.6 Broad auxiliary-channel search with BRUCO . . . 4
4.7 Transfer-function estimation . . . 4
4.8 Custom GWpy transfer-function subtraction . . . 5
4.9 GWSubtract time-domain subtraction . . . 5
4.10 NonSENS . . . 6
4.11 Reproducible NonSENS summaries . . . 6
4.12 BNS range and subtraction validation . . . 6
4.13 Exploratory control-response and lock-loss analysis . . . 7

5 Results 7
5.1 Earthquake Mode strain differs from quiet operation . . . 7
5.2 Coherence identifies useful control witnesses . . . 7
5.3 Custom sequential subtraction . . . 7
5.4 GWSubtract subtraction produced the largest completed range recovery . . . 9
5.5 NonSENS performance depended on bandwidth and conditioning . . . 9
5.6 Comparison of subtraction methods . . . 9
5.7 Validation results . . . 11

6 Discussion 11

7 Future Work 12

8 Conclusion 12

9 Acknowledgments 13

1 Introduction

The first direct detection of gravitational waves by LIGO was a groundbreaking discovery for physics, and established a new method of observing the Universe.[1] Advanced LIGO uses two laser interferometers with 4-km-long arms and is sensitive to differential arm-length changes far smaller than the diameter of a proton.[2] An incoming gravitational wave changes the relative lengths of the detector's two arms, producing a small change in the interference pattern of the laser light. Extracting this signal, as described previously, requires the detector to suppress environmental and instrumental disturbances by many orders of magnitude.[10,12]

Earthquakes present a significant challenge to detector sensitivity because large ground motion can disturb the interferometer and cause lock loss.[4,5,11] The interferometer is considered locked when its optical cavities and control degrees of freedom are maintained at the operating points required for gravitational-wave observation. A lock loss occurs when the control system can no longer maintain this state, causing the detector to fall out of the observing configuration. During an earthquake, it is therefore important to keep the interferometer locked through high seismic motion and also to preserve the astrophysical sensitivity of the calibrated strain data. This project investigates whether the auxiliary control channels recorded during Earthquake Mode can be used as witnesses of excess instrumental noise and whether that witnessed noise can be subtracted offline.

The primary result is that linearly witnessed control noise could be subtracted to recover more than 11 Mpc of binary-neutron-star (BNS) range using the best completed subtraction pipeline. Some model and filter choices reduced noise over a limited frequency band while worsening the detector's astrophysical sensitivity. For that reason, this project assesses coherence, spectral ratios, out-of-band injection checks, quiet-reference comparisons, and the BNS range together to evaluate the effectiveness of a subtraction.

2 Background

2.1 Seismic isolation, interferometer controls, and Earthquake Mode

Ground motion is one of the most significant noise sources and disturbances that the detector experiences. Advanced LIGO uses active seismic-isolation platforms combined with multistage pendulum suspensions to reduce ground motion transmission into the test masses.[3] These passive and active isolation systems are effective over most of the detection band, but isolation alone is not sufficient at the lowest frequencies, such as during earthquakes, due to coupling. Feedback control systems sense and actuate the interferometer to keep its optical cavities aligned, resonant, and locked continuously.

Since long-period seismic waves from earthquakes can push the detector out of lock due to resonance, earthquake-specific control strategies were developed at LIGO to reduce lock loss.[4,5] This mode is called Earthquake Mode, and its function is to preserve stable interferometer operation during conditions that would cause lock loss.

The goals of maintaining lock and also astrophysical sensitivity are not in complete alignment. Earthquake Mode uses modified seismic and interferometer control configurations to respond more strongly to the ground motion. When poorly suppressed, this additional motion and control noise can couple into the gravitational-wave readout, increasing the noise in the strain channel and decreasing the detector's overall sensitivity.[8,10,12] This raises the main question of this paper: can the noise introduced during Earthquake Mode's heightened controls be identified and subtracted?

2.2 Witness-channel noise subtraction

The auxiliary detector channels record multiple potential instrumental noise sources, meaning that certain auxiliary channels can act as witnesses. This means the channel measures a disturbance that correlates with

the strain readout channel but does not contain the gravitational-wave signal itself. If the relationship between a witness channel and strain can be modeled, the witness can be filtered to predict its contribution to the detector output, and that predicted contribution can be subtracted.[6,13]

This type of linear witness-channel subtraction has previously improved Advanced LIGO sensitivity. Davis et al. demonstrated that auxiliary control signals can be filtered and subtracted from calibrated strain when their couplings are largely stable and well measured.[6] NonSENS, developed by Vajente et al., extends witness-channel subtraction to non-stationary couplings by modeling both the noise source and variations in its coupling to the target channel.[7] Noise subtraction and control-noise mitigation continue to be important components of Advanced LIGO commissioning and detector performance during observing runs.[8]

3 Objectives

The goal of this project was to determine whether auxiliary control signals could be used to identify and subtract excess strain noise during LIGO Hanford Earthquake Mode, thereby recovering astrophysical sensitivity without compromising detector operation. The specific objectives were the following.

1. Characterize excess strain noise during Earthquake Mode relative to quiet detector operation.
2. Identify auxiliary channels containing predictive information about that noise.
3. Develop and evaluate multiple linear subtraction approaches and quantify their effect on detector sensitivity.
4. Develop a framework that could be extended to other earthquake mode periods and investigate time-dependent noise couplings.

4 Methodology

4.1 Analysis overview

The analysis was performed on archived LIGO Hanford Observatory data using Python, GWpy,[9] NumPy,[15] SciPy,[16] and Matplotlib.[19] The overall workflow was the following.

1. Select representative Earthquake Mode and quiet-reference intervals.
2. Compare calibrated strain and interferometer control behavior.
3. Calculate coherence between candidate witnesses and strain.
4. Search more broadly for coherent auxiliary channels.
5. Estimate frequency-dependent witness-to-strain transfer functions.
6. Construct predicted noise contributions.
7. Subtract those contributions using several methods.
8. Recalculate coherence after subtraction.
9. Test for noise injection inside and outside the modeled band.
10. Re-quantify astrophysical sensitivity using the BNS range.

An important focus of the project was reproducibility. The analysis was progressively reorganized from individual exploratory scripts into reusable functions and configuration-driven runs so that witness lists, time intervals, filter choices, and model parameters could be changed without rewriting the analysis.

4.2 Earthquake Mode and quiet-reference data

The target signal throughout the project was the calibrated LIGO Hanford strain channel

`H1:GDS-CALIB_STRAIN_NOLINES`.

The representative Earthquake Mode interval used for the main final comparison began at GPS 1443544132, and the representative quiet interval used to show the detector noise floor began at GPS 1443540243, one

hour earlier to increase the baseline of similarity between the measurements. Several analyses used longer surrounding stretches for training or filter construction, while the final comparison used common subintervals so the subtraction outputs could be evaluated on the same data.

Because filtering and resampling can introduce edge transients, analysis ranges were trimmed accordingly before sensitivity metrics were calculated. The main comparison plots were displayed beginning near 9 Hz so that the frequency range corresponded to the region in which the calibrated strain and subtraction products were used for interpretation.

4.3 Control-system witness channels

The initial witness set was selected from the Length Sensing and Control (LSC) and Alignment Sensing and Control (ASC) systems because these channels record control signals associated with interferometer degrees of freedom that can couple to the Differential Arm Length or DARM.

The principal LSC witnesses were

- `H1:LSC-MICH_OUT_DQ`,
- `H1:LSC-PRCL_OUT_DQ`, and
- `H1:LSC-SRCL_OUT_DQ`.

MICH describes the Michelson degree of freedom of the central interferometer, while PRCL and SRCL describe the power-recycling-cavity and signal-recycling-cavity longitudinal degrees of freedom.

The selected ASC included hard and soft angular control channels:

- `H1:ASC-DHARD_P_OUT_DQ`,
- `H1:ASC-DHARD_Y_OUT_DQ`,
- `H1:ASC-CHARD_P_OUT_DQ`,
- `H1:ASC-CHARD_Y_OUT_DQ`,
- `H1:ASC-DSOFT_P_OUT_DQ`, and
- `H1:ASC-DSOFT_Y_OUT_DQ`.

Additional arm- and corner-control channels were investigated during the control-response inspection, such as CSOFT, INP1, ASC-MICH, PRC2, SRC1, and SRC2 pitch and yaw signals.

Since detector channels are recorded at different sampling rates, calculating cross-spectral density and coherence required adjusting the signals to compatible time and frequency grids. The initial ASC analysis used an incorrect rate-handling approach in which the DARM target was unnecessarily downsampled for the subtraction itself. That implementation was not used for the final scientific result. In the corrected analysis, the target strain remained at the appropriate rate for the final cleaned output, while resampling was used only when required.

4.4 Amplitude spectral density

Detector noise was visualized using the amplitude spectral density (ASD),

$$\mathrm{ASD}(f) = \sqrt{S_{yy}(f)}, \tag{1}$$

where $S_{yy}(f)$ is the one-sided power spectral density of the strain time series $y(t)$. Welch averaging was used to reduce the variance of the spectral estimate.[14] ASD comparisons were generated for the quiet reference, the original Earthquake Mode strain, and the cleaned output from each subtraction method.

A cleaned ASD below the original earthquake ASD indicates reduced noise at that frequency, but that alone is not sufficient to determine if it will be sufficient for subtraction since a filter can improve one part of the spectrum while injecting noise elsewhere. For this reason, ASD plots were always assessed together with spectral ratios and the BNS range.

4.5 Coherence analysis

The linear relationship between an auxiliary witness $x(t)$ and the strain target $y(t)$ was quantified with magnitude-squared coherence,[6]

$$\gamma_{xy}^2(f) = \frac{|S_{xy}(f)|^2}{S_{xx}(f)S_{yy}(f)}, \tag{2}$$

where $S_{xy}(f)$ is the cross-power spectral density and $S_{xx}(f)$ and $S_{yy}(f)$ are the witness and target power spectral densities. Coherence ranges from 0 to 1. A value near 1 means that a strong stable linear relationship exists at that frequency, while a value near 0 means that the target power is unlikely to be linearly related to that witness.

Coherence was calculated during both the Earthquake Mode period and the quiet reference to identify channels with large amplitudes whose motion was also related to the strain noise.

After each major subtraction stage, coherence was recalculated between the cleaned target and the remaining witnesses. This was important for validation since the coherence should decrease in the frequencies where the subtraction was effective if a witness contribution had actually been removed. Recalculating coherence also helped determine whether another witness still contained independent information that could be removed with a sequential subtraction.

For an ideal single-witness, linear, stationary subtraction, the residual power is related to coherence by approximately

$$S_{\mathrm{res}}(f) \simeq S_{yy}(f)\left[1 - \gamma_{xy}^2(f)\right], \tag{3}$$

so that the expected ASD ratio is

$$\frac{\mathrm{ASD}_{\mathrm{res}}(f)}{\mathrm{ASD}_{\mathrm{orig}}(f)} \simeq \sqrt{1 - \gamma_{xy}^2(f)}. \tag{4}$$

This relation is not exact for correlated witnesses or time-varying couplings, but it provides a useful check on whether an apparent subtraction is reasonable.

4.6 Broad auxiliary-channel search with BRUCO

The initial LSC and ASC witnesses were chosen because of their known roles in interferometer control. To determine if there were other relevant and subtractable coherent channels, I also used BRUCO, a broadband coherence analysis tool used in LIGO detector characterization, to compare the target strain with a large set of auxiliary channels.

The BRUCO run for the selected earthquake interval evaluated thousands of auxiliary channels and ranked coherent channels by frequency. Channels showing around 0.2 coherence or more were inspected as candidate witnesses. The resulting list was then filtered to remove channels that were inappropriate for subtraction or that represented redundant readouts of the same underlying signals.

4.7 Transfer-function estimation

For a linear stationary coupling, the frequency-dependent response from witness to target can be estimated from the cross spectrum,[6,13]

$$H(f) = \frac{S_{xy}(f)}{S_{xx}(f)}. \tag{5}$$

The complex transfer function $H(f)$ contains both magnitude and phase. If $X(f)$ is the Fourier transform

of the witness, the predicted contribution of that witness to strain is

$$\widehat{Y}_n(f) = H(f)X(f), \tag{6}$$

and a cleaned target would be

$$Y_{\text{clean}}(f) = Y(f) - \widehat{Y}_n(f). \tag{7}$$

When the witness and target have low coherence, transfer-function estimates become unreliable. The custom GWPy analysis therefore used coherence and frequency restrictions to avoid interpreting weak correlations as a physical coupling. Both transfer-function magnitude and phase were investigated, and transfer-function smoothness was used as a diagnostic as well.

Control-injection intervals were also identified for MICH, SRCL, PRCL, DHARD pitch and yaw, and CHARD pitch and yaw. These measurements provide a way to compare the response determined with Earthquake Mode data with an intentionally excited control response and to see if they would make a better substitute for subtraction. Due to a large misalignment between the calculated and injected transfer functions, it is likely there are unknown factors during earthquake mode, such as earthquake motion modulating the coupling, suggesting that the injected transfer functions cannot be used. Therefore, the final subtraction results in this paper are based on transfer functions calculated using the earthquake data itself.

4.8 Custom GWpy transfer-function subtraction

The first subtraction pipeline was implemented with GWpy and standard Python tools. Reusable functions were developed for data access, synchronization and resampling, ASD calculation, coherence, transfer-function estimation, witness filtering, subtraction, sequential subtraction, BNS-range calculation, and figure generation.

The custom method was useful for testing witness order. Instead of subtracting all channels at once, a dominant witness could be removed first and the residual strain could then be used as the target for the next stage. After each stage, coherence with the remaining witnesses was recalculated.

The LSC assessments showed that the first subtraction produced most of the available sensitivity recovery, while subsequent channels produced smaller changes.

4.9 GWSubtract time-domain subtraction

The second approach was based on the established GWSubtract method described by Davis et al.[6] An important change from the first GWPy subtraction code was performing the subtraction on the time-domain strain product rather than subtracting ASD curves.

The approach is the same as the transfer-function method: a witness-to-target coupling is estimated, converted into a filter, applied to the witness, and used to generate a predicted noise time series. The cleaned time series is then

$$y_{\text{clean}}(t) = y(t) - \widehat{n}(t). \tag{8}$$

By saving the cleaned time series the phase information is retained and the output can be treated like other detector data. The final cleaned product could therefore be saved to a gravitational-wave frame file and passed through the same ASD and BNS-range calculations used for the original calibrated strain.

The final all-controls GWSubtract run included the selected LSC and ASC witnesses and provided the strongest completed BNS-range recovery in this project.

4.10 NonSENS

The third subtraction approach used NonSENS, based on the framework developed by Vajente et al. for learning stable digital filters from witness channels and extending subtraction to non-stationary couplings.[7] First, it provides a way to fit multiple witnesses simultaneously rather than requiring a manually chosen sequential order. Second, the framework can be extended to nonlinear or time-varying couplings, which are plausible during high seismic disturbances. Third, unlike the other two methods, NonSENS provides a framework in which subtraction can be constrained and validated to reduce the risk of over-subtraction.[17]

The NonSENS work in this project concentrated on reproducible linear subtraction configurations. Nonlinear NonSENS modeling was explored as a future direction but is not presented as a completed result.

Multiple configurations were tested during development. These included different frequency bands, weighting strategies, witness combinations, and output filters. Broad-band fits, split-band attempts such as 10–100 Hz and 100–300 Hz, a 10–1000 Hz configuration, LSC-only models, a 10–120 Hz LSC configuration, and a 1–150 Hz LSC configuration were also tested. PSD weighting and unweighted optimization were tried, and BNS-sensitive weighting was considered as a possible future optimization method.

A repeated issue was that a model could subtract useful noise inside the intended training band while noise "blew up" outside that region. To address this, the predicted NonSENS noise estimate was conditioned before subtraction. Several low-pass and band-pass choices were tested to balance in-band cleaning and out-of-band injection. A 10–150 Hz conditioned LSC result was selected for the final method comparison.

4.11 Reproducible NonSENS summaries

Because many NonSENS configurations were trained, the final analysis was reorganized around a common code base with a configuration file so that each parameter set could be set in a YAML file rather than through creating new versions of the analysis script.

For each training set, the summary workflow was designed to preserve or report the exact configuration and witness list, the training and evaluation intervals, preprocessing and normalization parameters, model frequency limits and optimization settings, original and cleaned ASDs, the cleaned over original ASD ratio, witness coherence before and after subtraction, BNS range, comparison with other trained models, the predicted noise estimate, and trained filter and model outputs needed to reproduce the subtraction.

A remaining reproducibility issue is that a later attempt to reconstruct the NonSENS noise estimate from saved per-witness filter information did not yet reproduce the original model output accurately. Consequently, the cleaned NonSENS time series and its measured performance are used in this report, but the final conversion of every trained model into independently validated deployment filters remains unfinished. This is discussed further as a limitation and future task.

4.12 BNS range and subtraction validation

The primary astrophysical performance metric was BNS range.[10,12] For a compact-binary signal with frequency-domain waveform $\widetilde{h}(f)$ and detector noise power spectral density $S_n(f)$, the matched-filter signal-to-noise ratio is

$$\rho^2 = 4 \int \frac{|\widetilde{h}(f)|^2}{S_n(f)}\, df. \tag{9}$$

The BNS range turns the detector noise spectrum into the sky- and orientation-averaged distance at which a standard binary-neutron-star inspiral wave would reach the detection threshold.[10,12,18] It is therefore sensitive to the detector noise in the frequencies that contribute most strongly to a neutron-star inspiral rather than to a single frequency bin.

Every candidate subtraction was evaluated using several tests: ASD before and after subtraction, cleaned over original ASD ratio, coherence before and after subtraction, comparison with the quiet reference, checks for noise injection outside the modeled band, and BNS range.

4.13 Exploratory control-response and lock-loss analysis

Another portion of this project investigated the control response to heightened seismic motion. Scripts were developed to compare control-channel amplitudes between quiet and earthquake periods, inspect time-domain peak-to-peak motion, generate spectrograms, and calculate band-limited root-mean-square (BLRMS) motion over relevant seismic and microseismic frequency bands.

This was to determine which controls might require greater authority as ground motion increases and whether particular control degrees of freedom become more active before earthquake-triggered lock loss. Earthquake-tagged lock losses and intervals in which the detector had previously been observing were identified as useful cases for further investigation.

The side-exploration was not developed to the same level as the subtraction comparison. Therefore, no final claim about a universal lock-loss predictor or a specific control threshold is made in this paper.

5 Results

5.1 Earthquake Mode strain differs from quiet operation

The first part of the project compared the strain ASD during the selected Earthquake Mode interval with a quiet reference from one hour prior. The earthquake data showed excess noise relative to the quiet detector. Control-channel time series and spectra also changed during the same interval, which supported further tests on whether the excess strain noise was correlated with interferometer controls.

5.2 Coherence identifies useful control witnesses

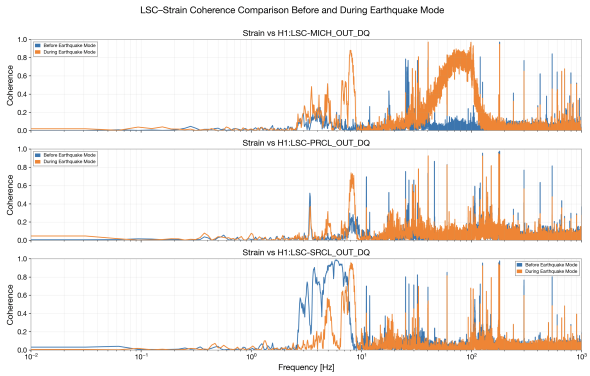
Magnitude-squared coherence was calculated between strain and the selected LSC and ASC controls. The LSC controls presented the most useful broadband subtraction information for the analyzed interval, especially the MICH, PRCL, and SRCL witnesses. Although individual frequency features were still present in the ASC channels, they contained much weaker broadband coherence.

For the final ASC coherence calculation over 10–100 Hz, the mean coherence values were approximately 0.0856 for CHARD pitch, 0.0833 for CHARD yaw, 0.0403 for DHARD yaw, 0.0191 for DHARD pitch, 0.0065 for DSOFT yaw, and 0.0059 for DSOFT pitch. These averages are not intended as significance statistics and can hide narrow coherent structures, but they show that the tested ASC channels were weaker broadband linear predictors of strain compared to the dominant LSC witnesses for the selected event.

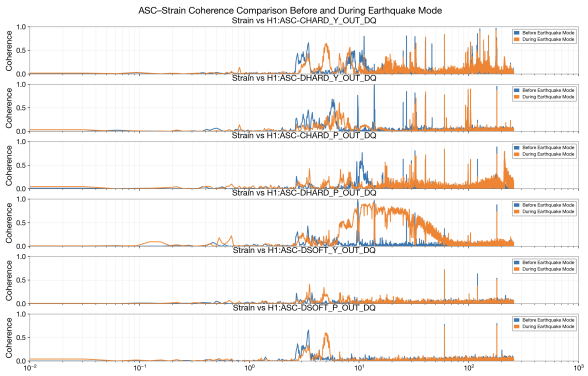
BRUCO provided a complementary broad search across auxiliary channels and confirmed that coherent information was not solely from the manually chosen witnesses. Nonetheless, the final subtraction set was utilizing the LSC and ASC control channels rather than automatically including every channel with a coherent feature, as described previously.

5.3 Custom sequential subtraction

The custom GWpy analysis was first used to sequentially test LSC witnesses. In a representative LSC-only test, the original BNS range was approximately 75.15 Mpc. Subtracting the MICH witness increased the range to approximately 84.73 Mpc. After the additional PRCL and SRCL stages, the final range was approximately 84.36 Mpc, as visible in Figure 2.



(a) LSC witness coherence.



(b) ASC witness coherence.

Figure 1: Magnitude-squared coherence between calibrated strain and selected control witnesses during Earthquake Mode. Coherence identifies frequencies at which a witness contains linearly predictive information about the strain noise.

The most important feature of this result is that the first dominant witness produced most of the recovered range. Additional witnesses did not improve the subtraction much further. As described previously, recalculating coherence after each stage showed whether the remaining witness still contained independent correlated noise and provided a physical check on the sequential order.

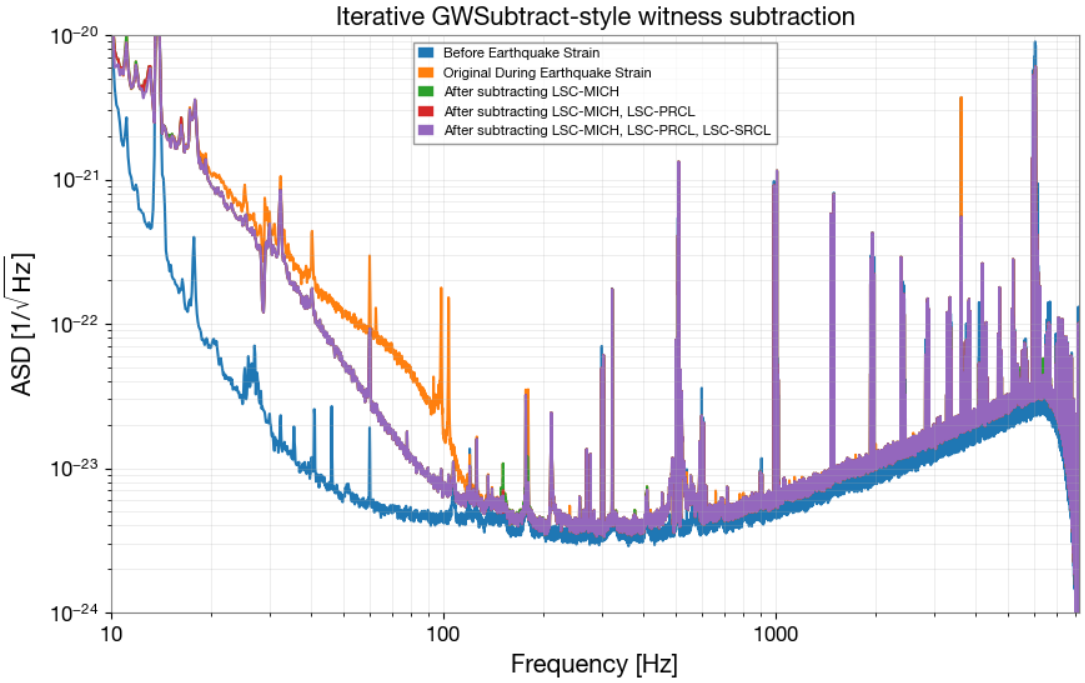


Figure 2: Representative sequential custom subtraction. The dominant witness produces most of the improvement, demonstrating that adding more witnesses does not necessarily continue to improve astrophysical sensitivity.

5.4 GWSubtract subtraction produced the largest completed range recovery

The strongest completed subtraction used the GWSubtract time-domain pipeline with the selected control witnesses. Over the common analysis interval, the original Earthquake Mode strain had a BNS range of

$$R_{\mathrm{BNS,orig}} = 75.1908 \text{ Mpc}, \tag{10}$$

while the cleaned output reached

$$R_{\mathrm{BNS,clean}} = 86.6602 \text{ Mpc}. \tag{11}$$

The improvement was therefore

$$\Delta R_{\mathrm{BNS}} = 11.4694 \text{ Mpc}, \tag{12}$$

corresponding to an increase of approximately 15.3%.

The cleaned strain was preserved as a time series and put in a gravitational-wave frame file. This allowed the subtraction product to be evaluated using the same ASD and BNS-range tools as the original strain rather than only as a modified spectrum.

5.5 NonSENS performance depended on bandwidth and conditioning

The NonSENS work showed that model configuration strongly affects the quality of the subtraction. Broad-band and split-band trials did not all produce useful full-band performance. LSC-only models were more promising than combinations containing less useful witnesses, particularly at lower frequencies.

The best NonSENS result was an earlier LSC subtraction in which the predicted noise contribution was conditioned to around 10–150 Hz. The saved BNS range for this cleaned product was

$$R_{\mathrm{BNS,NonSENS}} = 83.56 \text{ Mpc}. \tag{13}$$

A later series of low-pass trials selected a causal low-pass configuration with a BNS range of

$$R_{\mathrm{BNS,LP}} = 70.962 \text{ Mpc}. \tag{14}$$

The lower range of the selected low-pass result is an important negative result. A subtraction can appear to improve a target region in the ASD while degrading frequencies that contribute enough to the inspiral signal-to-noise ratio to reduce the overall BNS range. Output conditioning was judged by its effect on the full cleaned spectrum and the astrophysical range.

5.6 Comparison of subtraction methods

The final comparison overlaid the original Earthquake Mode ASD, the quiet reference, the best custom subtraction, the GWSubtract subtraction, and the best retained NonSENS output. Figure 3 allows for the assessment of the quiet strain to earthquake mode differences, which method removed the largest portion of the excess noise, and wether any method introduced new noise outside the region it was intended to clean.

The GWSubtract result provided the largest completed BNS-range recovery. The completed NonSENS configuration produced a smaller range recovery and was more sensitive to training-band and output-filter choices, but it established a reproducible framework that can later be extended to time-dependent or nonlinear couplings.

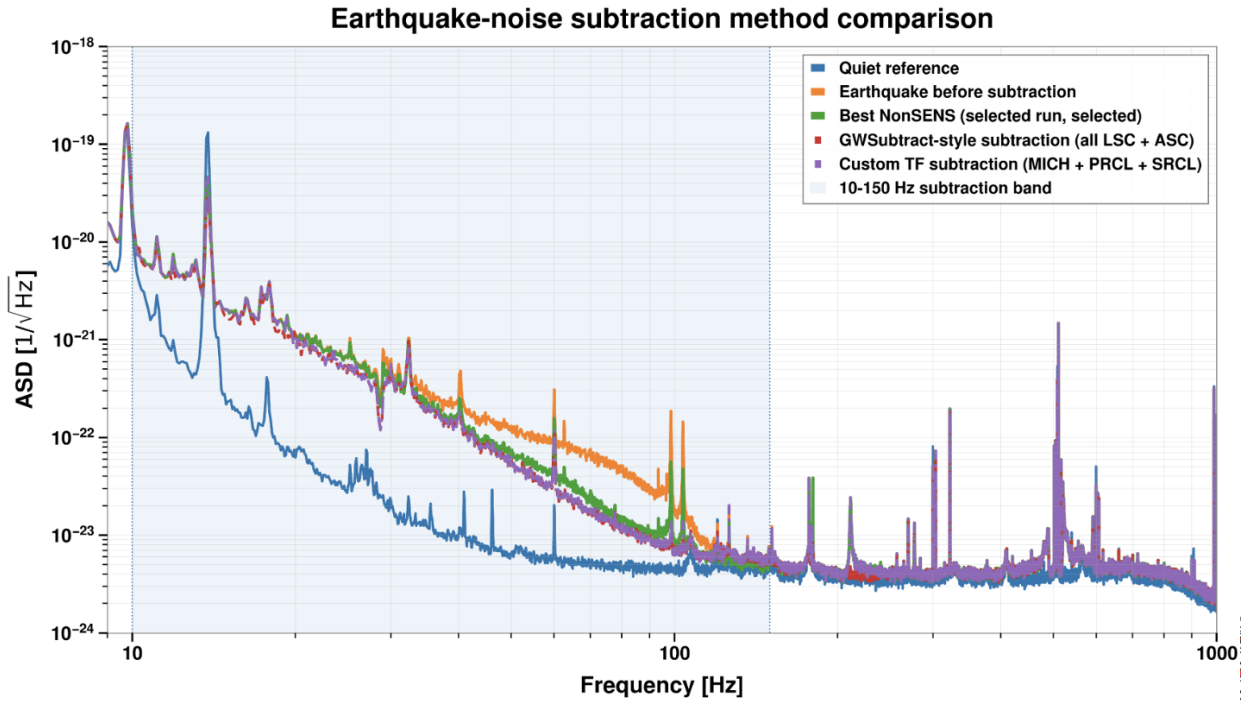


Figure 3: Final comparison of the principal subtraction approaches. The quiet reference shows the approximate non-earthquake sensitivity, while the cleaned spectra show how much of the Earthquake Mode excess could be recovered by each method.

Table 1: Representative completed subtraction results. Some custom-development tests used different internal windows or preprocessing, so the table records achieved performance rather than claiming a perfectly controlled algorithm benchmark.

Method	Starting BNS range	Cleaned BNS range	Interpretation
Custom LSC after MICH	~ 75.15 Mpc	~ 84.73 Mpc	Dominant sequential improvement from MICH
Custom LSC after MICH, PRCL, SRCL	~ 75.15 Mpc	~ 84.36 Mpc	Additional witnesses did not automatically improve the range
GWSubtract all-control run	75.19 Mpc	86.66 Mpc	Largest completed BNS-range recovery
NonSENS, conditioned LSC result	~ 75.15 Mpc	83.56 Mpc	Best retained NonSENS output for final comparison
NonSENS selected causal low-pass trial	~ 75.15 Mpc	70.962 Mpc	Demonstrated that local cleaning can still reduce integrated sensitivity

5.7 Validation results

The project repeatedly showed why multiple validation metrics are necessary. First, coherence before subtraction showed where a witness could remove linearly coupled power. Second, recalculating coherence after each subtraction tested whether the predicted correlated component had actually reduced. Third, the ASD ratio showed whether cleaning was localized to the intended band or whether noise had been introduced elsewhere. Finally, the BNS range determined whether the net spectral change improved observing sensitivity.

Several model and filtering choices appeared promising over certain frequency intervals but were not selected because their full-band behavior or BNS range was worse.

6 Discussion

This work shows that auxiliary-channel subtraction can recover a substantial fraction of the additional strain noise that is recorded by control witnesses through offline subtraction.

The strongest completed result increased the BNS range from 75.19 Mpc to 86.66 Mpc, recovering 11.47 Mpc of astrophysical sensitivity. Since the BNS range weights the spectrum by its contribution to a compact-binary signal it is more accurate in demonstrating improvements than an ASD that is just visibly lower, improving the detector's ability to make observations.

Through the custom sequential analysis, it was visible that the MICH witness produced most of the gain in the representative LSC test. PRCL and SRCL contained additional information, but later subtraction stages did not continue to increase the range significantly. Auxiliary witnesses can be correlated with one another, a second witness may measure some of the same disturbance already removed by the first, and transfer functions estimated from finite data are imperfect.

The coherence results also help explain why the final analysis was focused on LSC witnesses rather than ASC ones. The tested ASC channels contained comparatively weak coherence over 10–100 Hz in the selected event. Although narrow angular-control features may still couple strongly at particular frequencies, the broad LSC couplings were more useful for the completed linear subtraction.

The paper does not use the first ASC transfer-function subtraction results for final conclusions. During development, an incorrect sample-rate treatment was identified in the ASC implementation, and the analysis was corrected so that the final target strain was not unnecessarily downsampled for the subtraction. The development process therefore reinforced the need to validate the actual time series, transfer-function behavior, coherence, and astrophysical range rather than relying on a single plot.

The NonSENS study demonstrated how the modeled frequency band and the conditioning of the predicted noise estimate can be as important as the model itself. A trained filter can represent a useful coupling inside the band where it was constrained while also performing poorly elsewhere. Filtering the predicted noise contribution before subtraction reduced this risk, but the choice of filter introduced a tradeoff between preserving useful subtraction and suppressing out-of-band injection.

The comparison between the 83.56 Mpc conditioned NonSENS output and the 70.962 Mpc selected low-pass trial demonstrates how a filter can produce an apparent improvement in a section of the ASD while worsening the full astrophysical sensitivity. For future optimization, a cost function related more directly to BNS sensitivity may be preferable to optimizing an unweighted or purely PSD-weighted spectral residual.

The NonSENS results should not be interpreted as showing that the framework is worse than the linear GWSubtract method, as the selection of the optimal filter is still in progress. Additionally, NonSENS is especially valuable since it can be extended to time-varying and nonlinear couplings.[7] Which is relevant since during Earthquake Mode both seismic motion and control response change throughout an event.

Several additional limitations should be noted. The main results come from one earthquake rather than a statistical sample. The transfer-function and GWSubtract methods assume approximately linear and sufficiently stationary couplings over their training intervals. Different subtraction approaches were developed

with somewhat different internal preprocessing and training choices, so the current comparison records the best performance of each method rather than a controlled benchmark. Finally, auxiliary-channel subtraction must always be checked for the possibility of removing or distorting a real gravitational-wave signal, especially in the non-NonSENS subtractions.

Earthquake Mode can preserve lock under increased seismic activity, and at least part of the sensitivity loss during this period can be recovered if the relevant instrumental disturbance is witnessed in controls. The work also establishes a reusable analysis framework for whether a channel's measured coupling can be modeled, subtracted, validated, and translated into recovered astrophysical sensitivity.

7 Future Work

During the summer, control-channel amplitudes, spectrograms, BLRMS trends, and earthquake-associated lock losses were investigated to understand which controls respond most strongly as the ground motion increases. This work did not yet produce a sufficiently large earthquake population to claim a universal predictor for lock loss, but it does support an extension of the subtraction study: determine whether the witness-to-DARM transfer function remains stationary throughout an earthquake or changes as the detector approaches the limits of its control authority.

For each earthquake, the analysis could record seismic-band ground motion, interferometer control amplitudes, guardian state, whether lock was preserved, witness coherence, and subtraction performance. Events that enter Earthquake Mode and remain locked could then be compared with events that eventually lose lock. This may address the control-engineering question of why the detector can maintain lock during some earthquakes while others exceed the detector's available control authority. Auxiliary-channel subtraction must also be demonstrated to be safe for astrophysical signals.

Such a study would test whether the sensitivity recovery demonstrated in this project's target event can be generalized into a reliable detector-characterization strategy.

8 Conclusion

This project investigated noise coupling and offline subtraction during LIGO Hanford Earthquake Mode. A modular Python analysis framework was developed to compare earthquake and quiet detector behavior, calculate coherence between strain and auxiliary controls, estimate transfer functions, search for additional coherent witnesses, perform time-domain noise subtraction, and evaluate the cleaned data using both detector and astrophysical metrics.

Three subtraction approaches were studied: a custom GWpy transfer-function method, a GWSubtract linear time-domain subtraction, and NonSENS. The custom method showed that the LSC MICH witness accounted for the majority of the recoverable linearly coupled noise in the selected event. The GWSubtract analysis produced the strongest completed sensitivity recovery, increasing the BNS range from 75.19 Mpc to 86.66 Mpc, a gain of 11.47 Mpc or approximately 15.3%. NonSENS testing demonstrated the importance of training-band choice and output conditioning. The best result reached 83.56 Mpc, while an alternate low-pass result reached only 70.962 Mpc.

The project also established that subtraction cannot be validated from the ASD alone. Coherence before and after subtraction, ASD ratios, quiet-reference comparisons, out-of-band noise-injection checks, and BNS range provide helpful tests of whether a cleaned result is physically and astrophysically meaningful.

The next steps are to extend the analysis to a larger population of earthquakes, quantify how control couplings change over time and near lock loss, test nonlinear or explicitly non-stationary NonSENS models, complete validation of filters, and potentially use software injections to verify signal safety. Together, these extensions would determine whether offline Earthquake Mode noise subtraction can be generalized into a

reliable detector-characterization or low-latency sensitivity-recovery strategy.

Earthquake Mode allows LIGO to continue operating through disturbances that would likely cause lock loss otherwise. Carefully validated witness-channel subtraction can then recover part of the astrophysical sensitivity that is lost while the detector is operating in that configuration.

9 Acknowledgments

This work would not be possible without Elenna Capote due to her mentorship, technical guidance, and support throughout this project, or without the scientists, engineers, commissioners, operators, and staff at the LIGO Hanford Observatory for instruction, discussions, and computing help, as well as members of the LIGO Scientific Collaboration whose software and detector-characterization infrastructure made this work possible.

This work was carried out through the California Institute of Technology Student-Faculty Programs Summer Undergraduate Research Fellowship program.

References

[1] LIGO Scientific Collaboration and Virgo Collaboration, "Observation of Gravitational Waves from a Binary Black Hole Merger," *Physical Review Letters* **116**, 061102 (2016). doi:10.1103/PhysRevLett.116.061102.

[2] J. Aasi et al., "Advanced LIGO," *Classical and Quantum Gravity* **32**, 074001 (2015). doi:10.1088/0264-9381/32/7/074001.

[3] F. Matichard et al., "Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance," *Classical and Quantum Gravity* **32**, 185003 (2015).

[4] S. Biscans et al., "Control strategy to limit duty cycle impact of earthquakes on the LIGO gravitational-wave detectors," *Classical and Quantum Gravity* **35**, 055004 (2018). doi:10.1088/1361-6382/aaa4aa.

[5] E. Schwartz et al., "Improving the robustness of the Advanced LIGO detectors to earthquakes," *Classical and Quantum Gravity* **37**, 235007 (2020).

[6] D. Davis, T. J. Massinger, A. P. Lundgren, J. C. Driggers, A. L. Urban, and L. K. Nuttall, "Improving the Sensitivity of Advanced LIGO Using Noise Subtraction," *Classical and Quantum Gravity* **36**, 055011 (2019). doi:10.1088/1361-6382/ab01c5.

[7] G. Vajente et al., "Machine-learning non-stationary noise out of gravitational wave detectors," *Physical Review D* **101**, 042003 (2020). doi:10.1103/PhysRevD.101.042003.

[8] E. Capote et al., "Advanced LIGO detector performance in the fourth observing run," *Physical Review D* **111**, 062002 (2025). doi:10.1103/PhysRevD.111.062002.

[9] D. M. Macleod, J. S. Areeda, S. B. Coughlin, T. J. Massinger, and A. L. Urban, "GWpy: A Python package for gravitational-wave astrophysics," *SoftwareX* **13**, 100657 (2021). doi:10.1016/j.softx.2021.100657.

[10] D. V. Martynov et al., "The Sensitivity of the Advanced LIGO Detectors at the Beginning of Gravitational Wave Astronomy," *Physical Review D* **93**, 112004 (2016). doi:10.1103/PhysRevD.93.112004.

[11] M. Coughlin et al., "Limiting the Effects of Earthquakes on Gravitational-Wave Interferometers," *Classical and Quantum Gravity* **34**, 044004 (2017). doi:10.1088/1361-6382/aa5a60.

[12] A. Buikema et al., "Sensitivity and Performance of the Advanced LIGO Detectors in the Third Observing Run," *Physical Review D* **102**, 062003 (2020). doi:10.1103/PhysRevD.102.062003.

[13] N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications* (MIT Press, 1949). doi:10.7551/mitpress/2946.001.0001.

[14] P. D. Welch, "The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms," *IEEE Transactions on Audio and Electroacoustics* **15**, 70–73 (1967). doi:10.1109/TAU.1967.1161901.

[15] C. R. Harris et al., "Array Programming with NumPy," *Nature* **585**, 357–362 (2020). doi:10.1038/s41586-020-2649-2.

[16] P. Virtanen et al., "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods* **17**, 261–272 (2020). doi:10.1038/s41592-019-0686-2.

[17] M. Saleem et al., "Demonstration of Machine Learning-Assisted Low-Latency Noise Regression in Gravitational Wave Detectors," *Classical and Quantum Gravity* **41**, 195024 (2024). doi:10.1088/1361-6382/ad708a.

[18] LIGO Scientific Collaboration and Virgo Collaboration, "GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral," *Physical Review Letters* **119**, 161101 (2017). doi:10.1103/PhysRevLett.119.161101.

[19] J. D. Hunter, "Matplotlib: A 2D Graphics Environment," *Computing in Science & Engineering* **9**, 90–95 (2007). doi:10.1109/MCSE.2007.55.