Detection and mitigation of glitches in LISA data: A machine learning approach

Niklas Houba[®],^{*} Luigi Ferraioli[®], and Domenico Giardini[®] Institute of Geophysics, Department of Earth and Planetary Sciences, ETH Zurich, Sonneggstrasse 5, 8092 Zurich, Switzerland

(Received 2 January 2024; accepted 28 March 2024; published 22 April 2024)

The proposed Laser Interferometer Space Antenna (LISA) mission is tasked with the detection and characterization of gravitational waves from various sources in the Universe. This endeavor is challenged by transient displacement and acceleration noise artifacts, commonly called glitches. Uncalibrated glitches impact the interferometric measurements and decrease the signal quality of LISA's time-delay interferometry (TDI) data used for astrophysical data analysis. The paper introduces a novel calibration pipeline that employs a neural network ensemble to detect, characterize, and mitigate transient glitches of diverse morphologies. A convolutional neural network is designed for anomaly detection, accurately identifying and temporally pinpointing anomalies within the TDI time series. Then, a hybrid neural network is developed to differentiate between gravitational wave bursts and glitches, while a long short-term memory (LSTM) network architecture is deployed for glitch estimation. The LSTM network acts as a TDI inverter by processing noisy TDI data to obtain the underlying glitch dynamics. Finally, the inferred noise transient is subtracted from the interferometric measurements, enhancing data integrity and reducing biases in the parameter estimation of astronomical targets. We propose a low-latency solution featuring generalized LSTM networks primed for rapid response data processing and alert service in high-demand scenarios like predicting binary black hole mergers. The research highlights the critical role of machine learning in advancing methodologies for data calibration and astrophysical analysis in LISA.

DOI: 10.1103/PhysRevD.109.083027

I. INTRODUCTION

The Laser Interferometer Space Antenna (LISA) is an ambitious space mission scheduled for launch in the late 2030s, with the objective of detecting gravitational waves in the millihertz band. LISA aims to extend gravitational wave astronomy beyond the capabilities of ground-based detectors like the Laser Interferometer Gravitational Wave Observatory and Virgo [1]. Its measurement strategy involves monitoring the relative motion of free-falling test masses across the spacecraft constellation caused by the time-varying curvature of spacetime [2,3].

LISA is expected to encounter several intrinsic noise sources that challenge its scientific efficacy. One significant disturbance in LISA's measurement system is laser noise, primarily arising from the unequal lengths of the interferometer arms and intrinsic laser frequency instabilities [4]. Ideally, the arms of LISA's interspacecraft interferometers would be equal in length, allowing for the inherent cancellation of laser noise. However, these arms are unequal in practice due to the optical configuration and the orbital dynamics of the constellation [5]. This leads to residual laser noise spoiling the gravitational wave measurements by 7 to 8 orders of magnitude [6,7]. To mitigate the effect, a postprocessing technique called time-delay interferometry (TDI) is applied [8,9]. TDI works by combining timeshifted measurements from multiple spacecraft in a manner that effectively cancels out noise, thereby synthesizing virtual interferometers with equal arm lengths.

Another prevalent disturbance in LISA denotes tiltto-length (TTL) coupling, where spacecraft angular jitter translates into longitudinal displacement measurements [10,11]. The coupling occurs when the angular motion of the spacecraft and optomechanical components induce path length variations due to instrument imperfections, misalignments, and laser wavefront errors, adding another layer of complexity to the noise profile.

Amidst these quasistationary perturbations, LISA contends with transient artifacts, known as glitches. These glitches, first observed in the LISA *Pathfinder* (LPF) mission, are expected to appear in LISA's TDI data, as well [12,13]. Glitches can vary in form, intensity, and duration, and their presence will interfere with accurately characterizing gravitational wave sources. Given a good understanding of suppressing and removing perpetual noise sources like laser noise [14,15] and TTL coupling [16–18], understanding and mitigating glitch artifacts becomes the next challenge in data calibration. The transition to

^{*}niklas.houba@erdw.ethz.ch

handling unpredictable transient glitches during LISA data analysis requires novel methodologies outlined in the paper.

A. Glitches in LISA Pathfinder

The LPF mission, launched in December 2015, marked a milestone in the demonstration of technology critical for LISA [19]. LPF consisted of one spacecraft, and its primary goal was to measure the relative acceleration between two test masses in free fall [20]. The mission surpassed expectations by achieving remarkably high measurement sensitivity, marking a major technological breakthrough and providing validation for the foundational concepts of LISA's design [21,22].

Alongside these successes, LPF faced challenges in the form of transient signals, i.e., glitches, in its interferometric measurements [23]. The glitches, varying significantly in intensity and duration, presented themselves as nonperiodic events uncorrelated with the regular operations of the mission. The glitches' enigmatic nature poses a considerable challenge in data analysis and system understanding, as their origins remain unknown [24]. Extensive research involving simulation studies and on-ground testing has been undertaken, focusing on comprehending the physical processes behind these transient events. LPF's data analysis revealed two predominant types of glitches [24]:

- (1) Fast rise and exponential decay (FRED) events: These glitches were characterized by a sudden, sharp increase in acceleration, followed by a gradual exponential decline back to the baseline. This pattern suggests a rapid energy release followed by a slower dissipation process.
- (2) Sine-Gaussian glitches: These glitches were identified by their distinctive pattern of sequential movements in opposite directions, resembling the shape of a sine wave modulated by a Gaussian envelope. Sine-Gaussian glitches were less frequent than FRED events and appeared to be stochastic in nature, with a notably higher incidence during phases when the system was cooled down. This observation hinted at a potential correlation between environmental conditions and the emergence of glitches, suggesting that external factors might influence their manifestation.

B. Projection to LISA

In the context of the LISA mission, addressing glitches is paramount. Unmitigated glitches in TDI data pose a risk to the mission's primary objective of accurate gravitational wave detection and source parameter estimation [25]. The interference caused by these glitches could potentially lead to estimation biases, misinterpretations of gravitational wave signals, or even mask transient astrophysical signals entirely. The experience gained from handling glitches in the LPF mission plays a crucial role in shaping the strategies for dealing with similar issues in LISA. A significant part of this strategy involves utilizing shapelet models to characterize glitches mathematically [26]. These models are adept at capturing a range of glitch behaviors, including sharp rises, exponential decay, and oscillatory patterns, which align with the characteristics observed in LPF's glitch events.

To effectively identify these glitches, the authors of [24] implemented a matching pursuit algorithm. This iterative method employs matched filtering alongside a parametric model, allowing for a systematic approach to glitch detection. The algorithm's effectiveness was demonstrated through its application to LPF's Δg noise-only measurements recorded from March 2016 to July 2017. The analysis of this data uncovered a variety of glitch events, thereby validating the utility of the approach in deriving statistics about the physical features of the glitch population.

The complexity and unpredictability of glitch morphologies suggest that standard models like first-order shapelets might not be sufficient to capture the entire spectrum of potential glitches [27]. The limitations of first-order models become apparent when confronted with more complex or atypical glitch patterns that may not conform to the shapes seen in LPF measurements. Therefore, while shapelet models and current calibration algorithms offer a solid foundation, they might need to be supplemented with more flexible modeling and mitigation techniques to cover a broader range of glitch morphologies.

C. Overview of the paper

The paper introduces a machine-learning approach to address artifacts in LISA data, proposing a three-step pipeline of connected neural networks for detecting, characterizing, and mitigating transient glitches. Section II deals with the mathematical modeling of glitches and their propagation through TDI. This section forms the theoretical backbone for understanding the nature and impact of glitches in the LISA context. Section III outlines the methodology for detecting anomalies in LISA data. It starts with the fundamentals of convolutional neural networks (CNNs), which form the first step of the proposed pipeline. The section further discusses data preprocessing for robust anomaly detection, simulation procedures, and evaluates the detection performance. Section IV focuses on the characterization of anomalies detected in the previous step. It discusses template matching for glitch characterization and hybrid neural networks to differentiate between glitches and gravitational wave bursts while locating glitch injection points. Section V explores a novel technique for mitigating the effects of glitches. It highlights the application of a long short-term memory (LSTM) network for glitch estimation, which forms the final step of the proposed pipeline before glitch removal. Additionally, the section discusses a low-latency solution. Section VI summarizes the findings of the study and highlights future directions of research.

II. MATHEMATICAL MODELING

The section lays the groundwork for comprehending the glitches LISA's interferometric measurements and TDI data are expected to encounter. We start with a parametric model that encapsulates the transient nature of FRED events, using insights drawn from the LPF mission. The model is further augmented with a stochastic component designed to account for more diverse and unforeseen morphologies of glitches in the context of LISA.

A. Time-domain FRED glitch model

The observations of glitches in LPF measurements have catalyzed the advancement of multiple glitch models. These models have been systematically integrated into the Python package lisaglitch [28] tailored for application within the context of LISA. The paper primarily concentrates on a FRED-like model embodied in lisaglitch, distinguished by a rapid onset and subsequent gradual decay. Mathematically, the model is defined as a one-sided double-exponential function that starts at a specified injection time $t_{inj} \ge 0$ and ramps up to a maximum of $A/(t_{rise} - t_{fall})$ during $t_{rise} > 0$ before flattening off to zero during $t_{fall} > 0$. The glitch signal for $t \ge t_{inj}$ and $t_{rise} \ne t_{fall}$ is

$$g_{\text{FRED}}(t) = A \cdot \frac{e^{\frac{t-t_{\text{inj}}}{t_{\text{rise}}}} - e^{\frac{-t-t_{\text{inj}}}{t_{\text{fall}}}}}{t_{\text{rise}} - t_{\text{fall}}}.$$
 (1)

The model captures the essential features of a FRED-like glitch, including its rise, peak, and decay phases. When $t_{\text{rise}} = t_{\text{fall}}$, the model uses a continuous extension to avoid singularity:

$$g_{\text{FRED}}(t) = A \cdot \frac{t - t_{\text{inj}}}{t_{\text{fall}}^2} \cdot e^{-\frac{t - t_{\text{inj}}}{t_{\text{fall}}}}.$$
 (2)

This extension is valid for $t \ge t_{inj}$ and ensures a smooth transition between the rise and fall phases of the glitch. Figure 1 presents the model evaluated for three sets of parameters. As the glitch maximum is a function of the amplitude *A*, as well as the rising and falling time parameters, the curves are normalized to facilitate visual comparison.

To account for more complex glitch shapes the FRED model will be modified by stochastic fluctuations.

B. Stochastic model extension

To comprehensively capture the diverse morphological characteristics of glitches, we have expanded the FRED model by incorporating a stochastic component. This extension reflects the random-walk nature of noise



FIG. 1. Normalized illustration of $g_{FRED}(t)$ for three different sets of parameters. The model will be modified by integrating stochastic fluctuations to accommodate more complex glitch shapes.

transients, characterized by random fluctuations occurring on time scales shorter than $t_{rise} + t_{fall}$. The stochastic element within the model is described as a standard Wiener process W(t), which is a continuous-time stochastic process representing cumulative random fluctuations [29]:

$$n(t) = W(t). \tag{3}$$

To model temporal correlations and prevent unrealistic high-frequency variations, we apply a moving average smoothing to Eq. (3), represented as an integration over a continuous window:

$$n_{\text{filtered}}(t) = \frac{1}{w} \int_{t-w}^{t} n(\tau) d\tau.$$
(4)

Here, w is the window size of the moving average, determining the extent of smoothing.

The FRED model with stochastic shaping denoted as FREDS² combines both the deterministic rise and fall pattern of the FRED model with the smoothed stochastic extension, resulting in a glitch signal that represents both LPF-inspired and unforeseen glitch characteristics:

$$g_{\text{FREDS}^2}(t) = g_{\text{FRED}}(t) \cdot n_{\text{filtered}}(t).$$
(5)

Figure 2 provides a normalized illustration of the extended glitch model for the three sets of parameters from Fig. 1. The figure demonstrates the impact of stochastic variations on the glitch morphology. Equal colors represent glitch pairs with identical FRED model parameters, illustrating that while their general behavior aligns, the introduction of stochastic elements can lead to significant short-term variations.

In this paper, we treat the moving average smoothing window size w as independent of the glitch duration. If it is



FIG. 2. Normalized illustration of $g_{\text{FREDS}^2}(t)$ for the three parameter sets of Fig. 1 and w = 5000 at 4 Hz. The same color denotes glitches with identical FRED model parameters highlighting that, while their broader behavior aligns, short-term variations can differ significantly due to the stochastic extension.

desired in future studies to align the stochastic smoothing process with the inherent glitch timescale, it may be beneficial to adjust the window size proportionately to the sum of t_{rise} and t_{fall} .

The stochastic model extension includes erratic and complex exponential decay distortions and is employed in the paper to simulate glitches in interferometric LISA measurements. It is important to note that introducing the stochastic extension complicates glitch mitigation during TDI data calibration. The stochastic nature of the extended model implies that mere estimation and subtraction of a glitch signal obtained from FRED model matching will be insufficient. The stochastic extension introduces random variations that significantly modify the glitch morphology. Consequently, even when applying FRED model-based glitch correction with the true FRED model parameters, residual glitch noise remains in the TDI data. This residual impact can obscure or distort the gravitational wave signals that LISA aims to detect, thereby reducing the accuracy and reliability of the mission's scientific return.

To address this challenge, a more sophisticated approach is required to predict the stochastic variations in glitch time series. Here, the use of neural networks for glitch time series prediction emerges as a promising solution. Neural networks, with their ability to learn complex patterns and dependencies in data, are well suited for recognizing the intricate and stochastic nature of these glitches. By training a network on a diverse set of glitch time series encompassing a range of stochastic behaviors, a powerful tool can be developed capable of accurately estimating and mitigating the impact of glitches in TDI data.

Before that, we will explore how the derived glitch model can be integrated into LISA's interferometer measurement model. This investigation is essential to understand



FIG. 3. Configuration of the LISA constellation, including the index convention for this paper analog to [30]. MOSA, OB, and TM acronyms stand for moving optical subassembly, optical bench, and test mass. A MOSA mainly accommodates an optical bench and a test mass.

their effect on the various TDI channels that need to be handled by the neural network-driven glitch detection and mitigation pipeline.

C. Acceleration and displacement transients

According to LISA's optical design, there are three distinct interferometers per optical bench: the interspacecraft interferometer, the test mass interferometer, and the reference interferometer. The configuration of the six optical benches of the LISA constellation is depicted in Fig 3. Focusing on optical bench 12 as a representative case, the corresponding measurement models are as follows:

(1) Interspacecraft interferometer:

$$isi_{12}(t) = p_{21:2.1}(t) - p_{12}(t) + H_{12}(t) + N_{12}^{isi}(t).$$
 (6)

(2) Test mass interferometer:

$$\operatorname{tmi}_{12}(t) = p_{13}(t) - p_{12}(t) + \frac{2}{\lambda} \boldsymbol{n}_{1.2} \cdot \boldsymbol{v}_{12}(t) + N_{12}^{\operatorname{tmi}}(t).$$
(7)

(3) Reference interferometer:

$$\mathrm{rfi}_{12}(t) = p_{13}(t) - p_{12}(t) + N_{12}^{\mathrm{rfi}}(t).$$
(8)

These measurements are quantified in units of frequency. In Eq. (6), the term $p_{12}(t)$ represents the laser noise from optical bench 12, which interferes with the laser beam of optical bench 21. The laser noise $p_{21;2.1}(t)$ from optical bench 21 undergoes a time delay, $\tau_{2.1}(t)$, due to the photon travel between optical bench 21 and 12. This delay is incorporated in the model as $p_{21;2.1}(t) \equiv p_{21}(t - \tau_{2.1}(t))$. For the test mass and reference interferometers, outlined in Eqs. (7) and (8), the laser beams from optical bench 12 and the adjacent optical bench 13 on the same spacecraft are combined. As a result, these equations include $p_{13}(t)$ instead of $p_{21;2.1}(t)$. No interspacecraft time delay is required.

In Eq. (6), $H_{12}(t)$ represents the gravitational wave contribution and the terms $N_{12}^{\text{isi}}(t)$, $N_{12}^{\text{tmi}}(t)$, and $N_{12}^{\text{rfi}}(t)$ represent optical path length noise in the respective interferometers.

The test mass interferometer is affected by displacements of the test mass. When considering units of frequency, the test mass velocity $v_{12}(t)$ impacts the measurement, as detailed in Eq. (7). Here, $v_{12}(t)$ is projected onto the line of sight using the laser link unit vector $n_{1,2}$. Analog to [31], it is presumed that the direction of the link remains constant while laser light travels along the arm. The factor 2 accounts for the round-trip journey of the laser beam as it hits and reflects off the test mass, and $\lambda = 1064.5$ nm denotes the nominal wavelength of the laser source.

A complete model of LISA's measurement setup would incorporate additional disturbances, such as clock noise, TTL noise, and backlink fiber noise [32–34]. These components need not be included in the modeling scope of this paper. Nonetheless, they will be regarded as disturbances in the simulation, affecting the performance of the glitch mitigation pipeline. Note that the measurement equations for the remaining five optical benches can be obtained from Eqs. (6) to (8) through a cyclic permutation of the indices.

By establishing the measurement system containing the three interferometers per optical bench in LISA, we can now outline the specific glitch injection points considered in this framework. Optical displacement glitches are taken into account via $N_{12}^{\text{isi}}(t)$, $N_{12}^{\text{tmi}}(t)$, and $N_{12}^{\text{rfi}}(t)$ while test mass acceleration glitches are addressed through the test mass velocity vector $\boldsymbol{v}_{12}(t)$ in Eq. (7). To represent these acceleration glitches, the extended glitch model $g_{\text{FREDS}^2}(t)$ is integrated over time, yielding

$$\boldsymbol{n}_{1.2} \cdot \boldsymbol{v}_{\text{FREDS}^2, 12}(t) \equiv \boldsymbol{v}_{\text{FREDS}^2, 12}(t)$$
$$= \int_0^t g_{\text{FREDS}^2, 12}(\tau) d\tau. \qquad (9)$$

For optical displacement glitches within the various interferometers, we consider differentiating $g_{\text{FREDS}^2}(t)$. The relation between frequency shift and displacement is given by

$$N_{\text{FREDS}^{2},12}^{\text{x}}(t) = \frac{1}{2\pi} \frac{2\pi}{\lambda} \frac{dg_{\text{FREDS}^{2},12}^{\text{x}}(t)}{dt}, \qquad (10)$$

with $x \in \{isi, tmi, rfi\}$. The paper categorizes instrumental glitches into two primary types: test mass acceleration

glitches and optical displacement glitches, also known as phasemeter glitches. This classification aligns with the glitches outlined in [25]. Future studies may consider test mass velocity glitch profiles resulting from closed-loop interactions of the drag-free and attitude control system. The investigation is ongoing and represents an important aspect of research.

D. TDI response to acceleration and displacement transients

Test mass acceleration and displacement artifacts are obscured in the raw interferometer measurements of LISA due to the dominant presence of laser noise. These transient anomalies become visible after TDI is applied. Therefore, understanding the response of the TDI postprocessing algorithm to glitches is crucial for mitigating their impact during astrophysical data analysis.

The initial noise reduction pipeline processes raw measurements of the LISA constellation to obtain noisecalibrated TDI data streams. Before executing TDI, intermediary variables $\xi_{ij}(t)$, $Q_{ij}(t)$, and $\eta_{ij}(t)$ with $i, j \in \{1, 2, 3\}$ and $i \neq j$ are computed in this process. These variables mitigate translational optical displacement noise as well as clock noise and reduce the number of lasers from six to three. The algorithms for calculating $\xi_{ij}(t)$, $Q_{ij}(t)$, and $\eta_{ij}(t)$ are expounded upon in [35]. The $\eta_{ij}(t)$ variables are used as input for TDI.

Note that there are several ways for linearly combining $\eta_{ij}(t)$ variables to suppress laser noise. These linear combinations, known as TDI channels, offer distinct benefits in the context of gravitational wave analysis and instrument diagnosis [36]. The second-generation TDI Michelson variable, which is a four-link combination, is focused on in this paper. TDI second generation (also known as TDI 2.0) suppresses laser noise to the required level for a rotating spacecraft configuration with time-variant arm lengths. For the TDI Michelson channel X(t) associated with spacecraft 1, measurements from the optical benches 12, 13, 21, and 31 are utilized. TDI Michelson X(t) is formulated as follows:

$$X(t) = \eta_{13} - \eta_{12} - \eta_{21;2,1} + \eta_{31;3,1} + \eta_{12;g} - \eta_{13;i} + \eta_{21;Y} - \eta_{31;Z} - \eta_{13;W} + \eta_{12;U} + \eta_{21;M} - \eta_{31;N} - \eta_{12;I} + \eta_{13;K} - \eta_{21;A} + \eta_{31;B}.$$
 (11)

The time argument is omitted in Eq. (11) for readability. The equation incorporates multiple nested delays, essential for laser noise cancellation, with delay arrangements summarized using letter notations for clarity as given in Table I. Similar combinations for Y(t) and Z(t) of space-craft 2 and 3 are obtained by cyclic permutation of measurement and delay indices.

We consider optical bench 12 as an example to assess the impact of glitches on X(t), Y(t), and Z(t). The different

TABLE I. Translation of delay index letters [17].

Index	Delays	Index	Delays
A	2.1, 1.3, 3.1, 1.3, 3.1, 1.2, 2.1	W	1.3, 3.1, 1.2, 2.1
В	3.1, 1.2, 2.1, 1.2, 2.1, 1.3, 3.1	X	2.3, 3.2, 2.1, 1.2
Ι	1.3, 3.1, 1.3, 3.1, 1.2, 2.1	Y	2.1, 1.3, 3.1
Κ	1.2, 2.1, 1.2, 2.1, 1.3, 3.1	Ζ	3.1, 1.2, 2.1
M	2.1, 1.2, 2.1, 1.3, 3.1	f	3.2, 2.3
Ν	3.1, 1.3, 3.1, 1.2, 2.1	g	1.3, 3.1
U	1.2, 2.1, 1.3, 3.1	i	1.2, 2.1

types of glitches introduced in the previous section will reveal distinctive patterns in those TDI channels.

1. Test mass acceleration glitches

For a test mass acceleration glitch in the test mass interferometer of optical bench 12, all terms in Eqs. (6)–(8) are set to zero except for $\frac{2}{\lambda} \boldsymbol{n}_{1.2} \cdot \boldsymbol{v}_{12}(t)$. The scenario is represented as follows:

$$\operatorname{tmi}_{12}(t) = \frac{2}{\lambda} \boldsymbol{n}_{1.2} \cdot \boldsymbol{v}_{\operatorname{FREDS}^2, 12}(t)$$
$$= \frac{2}{\lambda} v_{\operatorname{FREDS}^2, 12}(t), \qquad (12)$$

with $v_{\text{FREDS}^2,12}(t)$ according to Eq. [9]. Following the noise reduction procedure of [35], this leads to the following expressions for $\eta_{ii}(t)$:

$$\eta_{12}(t) = -\frac{1}{\lambda} v_{\text{FREDS}^2, 12}(t), \qquad (13)$$

$$\eta_{21}(t) = -\frac{1}{\lambda} \mathcal{D}_{1.2} v_{\text{FREDS}^2, 12}(t), \qquad (14)$$

with all other terms equal to zero and $\mathcal{D}_{i,j}f(t) \equiv f(t - \tau_{i,j})$. Using Eq. (11), we derive the impact of a test mass acceleration glitch of optical bench 12 on the TDI Michelson channel X(t):

$$X(t) = (\mathcal{D}_g - 1)(\mathcal{D}_i + 1)(\mathcal{D}_W - 1)\frac{v_{\text{FREDS}^2, 12}(t)}{\lambda}.$$
 (15)

The TDI Michelson channels Y(t) and Z(t) are obtained through a cyclic permutation of indices:

$$Y(t) = -2\mathcal{D}_{1,2}(\mathcal{D}_f - 1)(\mathcal{D}_X - 1)\frac{v_{\text{FREDS}^2, 12}(t)}{\lambda}, \quad (16)$$

$$Z(t) = 0. \tag{17}$$

For the analytical investigation presented in this study, we adopt the commonly used assumption that the delay operator is commutative. This assumption is equivalent to treating the arm lengths in the LISA constellation as time invariant.

2. Optical displacement glitches

Next, we will proceed to derive the TDI response to optical displacement glitches. For the interspacecraft interferometer of optical bench 12, all terms in Eqs. (6)–(8) are set to zero except for $N_{12}^{\text{isi}}(t)$. The scenario is represented as follows:

$$isi_{12}(t) = N_{FREDS^2, 12}^{isi}(t),$$
 (18)

with $N_{\text{FREDS}^2,12}^{\text{isi}}(t)$ from Eq. (10). Applying the initial noise reduction pipeline of [35], this translates to

$$\eta_{12}(t) = N_{\text{FREDS}^2 \ 12}^{\text{isi}}(t), \tag{19}$$

with all other terms equal to zero. Using Eq. (11), the impact on X(t) becomes

$$X(t) = -(\mathcal{D}_g - 1)(\mathcal{D}_W - 1)N_{\text{FREDS}^{2}, 12}^{\text{isi}}(t).$$
 (20)

For Y(t) and Z(t), we find

$$Y(t) = \mathcal{D}_{1,2}(\mathcal{D}_f - 1)(\mathcal{D}_X - 1)N_{\text{FREDS}^{2},12}^{\text{isi}}(t), \quad (21)$$

$$Z(t) = 0. \tag{22}$$

In the case of an optical displacement glitch in the test mass interferometer of optical bench 12, the measurement model of Eqs. (6)–(8) reduces to

$$\operatorname{tmi}_{12}(t) = N_{\operatorname{FREDS}^2, 12}^{\operatorname{tmi}}(t).$$
 (23)

This translates to changes in the intermediary variables:

$$\eta_{12}(t) = -\frac{1}{2} N_{\text{FREDS}^2, 12}^{\text{tmi}}(t), \qquad (24)$$

$$\eta_{21}(t) = -\frac{1}{2} \mathcal{D}_{1.2} N_{\text{FREDS}^2, 12}^{\text{tmi}}(t).$$
(25)

With Eq. (11), the impact on X(t) can be obtained:

$$X(t) = \frac{1}{2} (\mathcal{D}_g - 1) (D_i + 1) (\mathcal{D}_W - 1) \cdot N_{\text{FREDS}^2, 12}^{\text{tmi}}(t).$$
(26)

Similarly, for Y(t) and Z(t),

$$Y(t) = -\mathcal{D}_{1.2}(\mathcal{D}_f - 1)(\mathcal{D}_X - 1)N_{\text{FREDS}^2, 12}^{\text{tmi}}(t), \quad (27)$$

$$Z(t) = 0. \tag{28}$$

		Optical glitch			
	Test mass glitch	Interspacecraft interferometer	Test mass interferometer	Reference interferometer	
$\overline{X(t)}$	$(\mathcal{D}_g - 1)(\mathcal{D}_i + 1)(\mathcal{D}_W - 1)$	$-(\mathcal{D}_g-1)(\mathcal{D}_W-1)$	$\frac{1}{2}(\mathcal{D}_g - 1)(D_i + 1)(\mathcal{D}_W - 1)$	$-\mathcal{D}_i(\mathcal{D}_g-1)(\mathcal{D}_W-1)$	
Y(t)	$-2\mathcal{D}_{1,2}(\mathcal{D}_f-1)(\mathcal{D}_X-1)$	$\mathcal{D}_{1,2}(\mathcal{D}_f-1)(\mathcal{D}_X-1)$	$-\mathcal{D}_{1,2}(\mathcal{D}_f-1)(\mathcal{D}_X-1)$	$\mathcal{D}_{1,2}(\mathcal{D}_f-1)(\mathcal{D}_X-1)$	
Z(t)	0	0	0	0	

TABLE II. TDI responses to different types of glitches injected on the optical bench 12.

For the reference interferometer of optical bench 12, it holds:

$$rfi_{12}(t) = N_{FREDS^2, 12}^{rfi}(t).$$
 (29)

Following the approach as before, we obtain

$$\eta_{12}(t) = \frac{1}{2} N_{\text{FREDS}^2, 12}^{\text{rfi}}(t), \qquad (30)$$

$$\eta_{31}(t) = -\frac{1}{2} \mathcal{D}_{1.3} N_{\text{FREDS}^2, 12}^{\text{rfi}}(t), \qquad (31)$$

$$\eta_{13}(t) = \frac{1}{2} N_{\text{FREDS}^2, 12}^{\text{rfi}}(t), \qquad (32)$$

$$\eta_{21}(t) = \frac{1}{2} \mathcal{D}_{1.2} N_{\text{FREDS}^2, 12}^{\text{rfi}}(t).$$
(33)

Substituting into Eq. (11) leads to

$$X(t) = -\mathcal{D}_i(\mathcal{D}_g - 1)(\mathcal{D}_W - 1)N_{\text{FREDS}^2, 12}^{\text{rfi}}(t). \quad (34)$$

Finally, Y(t) and Z(t) are

$$Y(t) = \mathcal{D}_{1.2}(\mathcal{D}_f - 1)(\mathcal{D}_X - 1)N_{\text{FREDS}^2, 12}^{\text{rfi}}(t), \quad (35)$$

$$Z(t) = 0. \tag{36}$$

The TDI glitch responses for the remaining five optical benches can be derived through a cyclic permutation of indices from those detailed for optical bench 12.

From Eqs. (15) to (17) and (26) to (28), it becomes apparent that distinguishing between test-mass velocity glitches and optical displacement glitches within the test mass interferometer of an optical bench is not possible when solely analyzing TDI data. Both glitch types exhibit identical echo patterns in the TDI Michelson set. This agrees with the findings of [31]. Note that here, we assume glitches in the test mass acceleration profile. Under this assumption, a distinction will be attainable.

An interesting observation regarding the occurrence of optical displacement glitches in the reference interferometer is evident when examining Eqs. (30) and (32). The variables $\eta_{12}(t)$ and $\eta_{13}(t)$ are affected identically. Consequently, these two variables, being the only nondelayed

quantities in the definition of X(t), cancel in Eq. (11). Therefore, the effect of an optical displacement glitch in the reference interferometer of optical bench 12 is not immediately measurable in X(t) at the moment of injection. This is also evident from Table II summarizing the TDI responses for optical bench 12. For each glitch category, there is one TDI response, which, when expanded, includes a factor of ± 1 , except for the reference interferometer. Here, in the case of optical bench 12 in X(t), there is the delay $D_i = D_{1,2,2,1}$, and in Y(t), the delay $D_{1,2}$ as a prefactor so that a corresponding glitch is detectable in the TDI set at the earliest by an amount $\tau_{1,2}$ after the actual injection time in the interferometer. This fact will complicate the precise determination of the glitch injection time in the first step of the glitch mitigation pipeline, where the glitch source and, thus, the associated TDI response model are unknown.

From Table II, we see that one of the TDI Michelson combinations remains unaffected by glitches originating from a single optical bench, for example, Z(t) for glitches caused within optical bench 12 (or 21). This characteristic will be valuable for identifying the origin of the glitch in the second step of the glitch mitigation pipeline. Additionally, it aids in differentiating between glitches and transient astrophysical phenomena. Unlike the selective effect observed for glitches, the latter tends to impact all three TDI combinations simultaneously.

In scenarios where noise is present, advanced concepts, such as neural networks, become essential to accurately distinguish between transient astrophysical sources and glitches and determine the glitch source. This is particularly true for noise artifacts masked by quasistationary noise yet still impacting the accuracy of source parameter estimation. These faint glitches will be analyzed in the sections that follow.

E. TDI response to gravitational wave bursts

Gravitational wave bursts represent key astrophysical signals in LISA [37]. It is essential to differentiate and preserve these signals in the data during calibration rather than erroneously removing them as noise artifacts. Therefore, the glitch detection and mitigation pipeline aims to identify and discern gravitational wave bursts accurately. Achieving this requires modeling the influence of gravitational wave bursts on LISA's TDI data. The impact of gravitational wave bursts on the change in interferometric path length was presented in [13] and is briefly summarized below. The observed path length change, denoted as $\delta l_{i,j}(t)$, in the link between spacecraft *i* and *j* is as follows:

$$\delta l_{i,j}(t) = D_{i,j} \colon \int_{\xi_i}^{\xi_j} h(t) dt, \qquad (37)$$

where the colon indicates full contraction between tensors, with *t* being the Solar System barycenter time. The wave variable $\xi_i = t_i - \mathbf{k} \cdot \mathbf{x}_i$ defines surfaces of constant phase for the gravitational wave. The detector tensor $D_{i,j}$ is given by

$$D_{i,j} = \frac{1}{2} (\mathbf{r}_{i,j} \otimes \mathbf{r}_{i,j}) / (1 - \mathbf{k} \cdot \mathbf{r}_{i,j}), \qquad (38)$$

where k is the gravitational wave propagation direction, and $r_{i,j}$ is the unit-separation vector between the spacecraft pointing from i to j. The gravitational wave tensor h(t) comprises two polarization states, $h_+(t)$ and $h_{\times}(t)$:

$$h(t) = h_+(t)e'_+(\psi,\theta,\phi) + h_\times(t)e'_\times(\psi,\theta,\phi), \quad (39)$$

where e'_+ and e'_{\times} are polarization tensors rotated by the polarization angle $\psi \in [0, \pi]$, while θ and ϕ denote the source position in spherical polar coordinates stated in the Solar System barycentric frame [13].

The FREDS² model is employed for modeling the gravitational wave polarizations. To simulate the instrument measurement, we make use of the Python packages lisagwresponse and lisainstrument [38,39]. The packages enable simulating the response of the six interspacecraft interferometers to the gravitational wave signal of interest. The resulting data is then processed by the PyTDI software [40] to obtain the manifestation of the gravitational wave burst in the TDI Michelson channels.

III. ANOMALY DETECTION

This section is dedicated to robust anomaly detection, representing the initial step of the glitch mitigation pipeline developed for LISA. Classical anomaly detection employs diverse methods, including statistical techniques, clustering-based approaches, and dimensionality reduction strategies, each with its own strengths tailored for specific scenarios [41,42].

Within the scope of this paper, we define an anomaly as a transient signal that can signify either a gravitational wave source or an instrument glitch. Such an anomaly may exhibit complicated features and be obscured by quasistationary noise. In this context, the capabilities of CNNs appear highly promising. The adeptness in managing anomalies with unforeseen morphologies and noisy environments positions CNNs as a compelling choice for the anomaly detection step of the glitch mitigation pipeline. The applicability for LISA is studied in the following.

A. Convolutional neural networks

CNNs are specialized feed-forward neural networks that significantly impacted the evolution of machine learning with applications ranging from image and video recognition to natural language processing and financial time series analysis [43]. A notable advantage of CNNs is their ability to counteract issues such as vanishing and exploding gradients, commonly encountered in traditional neural networks, by applying weight regularization across reduced connections. One of the hallmarks of CNNs is their reduced need for data preprocessing compared to other image analysis techniques. CNNs learn to optimize their filters autonomously, a significant advantage in fields where feature extraction is complex [44].

In traditional neural networks, particularly those fully connected, each neuron in a layer is linked to all neurons in the subsequent layer. This can lead to a tremendous number of weights, especially for larger inputs like high-resolution images. CNNs drastically reduce this number by employing convolutional layers. These layers use cascaded convolution kernels to process the input in small segments, significantly scaling down the number of required neurons and, consequently, the number of free parameters [45].

Figure 4 illustrates the CNN architecture with TensorFlow attributes [46] applied in the research presented. The architecture is tailored to enable the precise detection of anomalies within LISA's TDI data. The network processes TDI data streams as images, thus capitalizing on the pattern recognition capabilities of CNNs.

The CNN structure presented in Fig. 4 initiates with two convolutional layers, each containing 32 filters of size 3×3 . This is followed by four convolutional layers, each with 64 filters. All convolutional layers employ the rectified linear unit (ReLU) activation function defined as $f(x) = \max(0, x)$. Utilizing a nonlinear activation function allows the network to learn and represent nonlinear relations



FIG. 4. CNN architecture and attributes employed for anomaly detection and anomaly injection time estimation. The architecture and subsequent network architectures are visualized using the open-source tool Netron [47].

between the network's inputs and outputs. The nonlinearity is important as the TDI data streams will contain complex patterns with nonlinear characteristics.

Each pair of convolutional layers is followed by a 2×2 maximum pooling layer. In CNNs, maximum pooling layers are used for reducing the spatial dimensions of the input data, which in turn diminishes the computational load and the number of parameters. Maximum pooling layers operate on the output of the convolutional layers, known as feature maps, with the primary goal of downsampling these maps while retaining relevant information. The process involves a specific-sized window that slides over the feature map. As it moves, typically in strides of two pixels, the pooling layer performs a maximum operation. This operation selects the maximum value from the pixels under the window, representing the entire region on the output feature map. As a result, the spatial dimensions of the feature map, both height and width, can be reduced. The reduction process is vital for preserving the most important features in the data. Features with higher numerical values in the feature maps, indicative of critical information like edges or specific textures, are prioritized. By focusing on these maximum values, the pooling layer ensures the retention of salient parts of the information. Note that pooling layers contribute to the model's generalizability and robustness. By abstracting the features learned and lowering the resolution of the feature maps, maximum pooling layers help in reducing the risk of overfitting. Consequently, the network is expected to perform better on unseen data [48,49].

The network of Fig. 4 then proceeds to a flattening step, where the two-dimensional feature maps are converted into a one-dimensional vector compatible with fully connected layers. The flattened data is fed into three successive fully connected dense layers with 64, 64, and 32 nodes, respectively. The function of these layers is to refine the feature representation further, focusing on subtle aspects that might indicate anomalies.

B. Time series to image conversion

The preprocessing of TDI time series data into an image format is critical in enabling its analysis with CNNs. The inherent structure and patterns of instrumental artifacts expected in TDI lend themselves well to image-based analysis.

The conversion process begins with segmenting TDI data streams into intervals of 7.5 min each. The duration strikes a pragmatic balance, being sufficiently brief to ensure that, ideally, each segment encapsulates no more than one anomaly while also being long enough to include a mix of anomalous and nonanomalous data. This approach enables the network to be trained on transient artifacts embedded in quasistationary noise.

Considering a sampling rate of 4 Hz, the 7.5-min segment translates to an initial image width of 1800 pixels,

calculated as $7.5 \text{ min} \times 60 \text{ sec} \times 4 \text{ samples/sec}$. To accommodate a subsequent resizing of the image to a targeted value of 256 pixels, we opt for a slightly reduced initial image width of 1792 pixels. This adjustment ensures that the resizing process can be executed with an integer factor, thereby reducing interpolation errors.

Following the segmentation process, the data stream is subjected to a moving average filter with a window size of 20. The window size is chosen to smooth short-term fluctuations while preserving the data's longer-term trends, which are important for anomaly detection. Subsequently, the values are normalized, ensuring their range is confined within the interval [0, 1]. This is achieved by subtracting the minimum value from each element and dividing by the range of values, i.e., the difference between the maximum and minimum values.

For graphical representation, the normalized and filtered data is mapped onto an image matrix. We select an image height of 16 pixels and perform the time series to image (TTI) conversion for the quasi-orthogonal TDI combination:

$$E(t) = \frac{X(t) - 2Y(t) + Z(t)}{\sqrt{6}} \equiv E_Y(t), \qquad (40)$$

as derived in [50] and its variant:

$$E_Z(t) = \frac{Y(t) - 2Z(t) + X(t)}{\sqrt{6}}.$$
(41)

Each of these combinations is represented as an image of size 16×1792 pixels, is obtained by stacking the two images.

Note that $E_Y(t)$ and $E_Z(t)$ represent six link combinations containing all three TDI Michelson variables in parallel. This ensures that glitches are captured in both $E_Y(t)$ and $E_Z(t)$ regardless of their injection point throughout the LISA constellation. This applies to gravitational wave bursts as well and proves particularly valuable in the detection step, where no differentiation is made between bursts and glitches.

Note that the quasiorthogonal TDI variable A(t) = $(X(t) - Z(t))/\sqrt{2}$ [50] also involves six laser links but only two TDI Michelson variables. It is, thus, less preferable in the process of obtaining consistent features for image generation. To illustrate, in the event of a test mass acceleration glitch of optical bench 21, A(t) does not immediately reflect the impact since it lacks the inclusion of the associated TDI Michelson variable Y(t). This leads to a delayed response in A(t) due to the time it takes for photons to travel between spacecraft 2 and 3 or spacecraft 2 and 1. In contrast, $E_Y(t)$ and $E_Z(t)$ both immediately capture the effect of the glitch as they include Y(t) in their formulation. Using A(t) and $E_Y(t)$ or A(t) and $E_Z(t)$ together can lead to glitches not consistently mapping to the same pixel position in the TTI conversion, depending on their source. This justifies the choice of $E_Y(t)$ and $E_Z(t)$. The quasiorthogonal TDI variable T(t), expressed as $(X(t) + Y(t) + Z(t))/\sqrt{3}$ [50], was initially considered as an alternative to either $E_Y(t)$ or $E_Z(t)$. However, it is observed that test mass acceleration glitches are less pronounced in T(t). A more detailed examination of this observation based on analytical TDI response models is omitted here.

Following the TTI conversion, the image width undergoes resizing for reasons primarily related to computational efficiency. The width is reduced to 256 pixels. The reduction to 256 pixels facilitates faster training times for the CNN but also tends to boost the detection performance by making artifacts more evident in the resized images with unchanged height. The resizing step makes use of the Python imaging library.

Examples of the stacked and resized images of TDI $E_Y(t)$ and $E_Z(t)$ time segments are given in Fig 5, where the ground truth label highlights the anomaly injection time across varying anomaly dynamics. Each subfigure displays dual-channel information from $E_Y(t)$ and $E_Z(t)$, as explained earlier, with the upper and lower half corresponding to $E_Y(t)$ and $E_Z(t)$, respectively. Figure 5(a) shows an anomaly that eludes visual detection in the presence of noise despite its several-minute span. In Fig. 5(b), $E_Y(t)$ reveals a slow-dynamic transient that stands out against the background. The same holds for $E_Z(t)$ in Fig. 5(c). Figure 5(d) marks an example of an anomaly



FIG. 5. Example outputs of the dual-channel TTI conversion process, including ground truth labels that highlight the anomaly injection time. The subfigures illustrate the diverse ways in which an anomaly may manifest. Note that in (a), the anomaly is scarcely discernible to the human observer. This prompts the exploration of CNN capabilities in detecting such subtle artifacts in LISA's TDI data.

being readily apparent against the background noise in both the $E_Y(t)$ and $E_Z(t)$ combinations. The CNN will be trained with these images across a spectrum of different anomaly dynamics and anomaly-to-noise ratios, a process whose efficacy will be evaluated in the following section.

C. Detection performance

The training and testing TDI data sets of the CNN model are generated using a suite of Python packages developed by the LISA consortium, including lisaglitch, lisainstrument, lisagwresponse, and PyTDI [28,38-40]. Here, the FREDS² model is locally integrated into lisaglitch to simulate glitches with stochastic fluctuations, as described in Sec. IIB. The instrument response is acquired with lisainstrument, which is then fed into PyTDI to obtain the TDI variables at 4 Hz in units of frequency. The simulation assumes static yet unequal distances between the spacecraft to streamline computation, an important aspect to be considered when generating large data sets. The noise profile in the TDI variables adheres to the SciRDv1 model outlined in [37]. Glitch characteristics are varied, with durations extending from a few seconds to 15 min and strengths ranging from 5% to 200% of the maximum quasistationary noise amplitude. The rising time of these glitches is set to fluctuate between 1 sec and 200 sec. Note that the window size of the moving average filter equals the one used in Fig. 2.

The CNN of Fig. 4 is applied in two stages. The initial stage classifies time segments as anomalous or nonanomalous, while the second stage focuses on predicting the injection times of anomalies. This dual-stage process improves the accuracy in predicting glitch injection times but requires training the network on two distinct data sets: one with a mixture of anomalous and nonanomalous segments and another exclusively containing anomalous data. Nonanomalous data are uniquely labeled with negative injection times. Consequently, a predicted negative injection time is evaluated as a nonanomalous data segment.

For the classification task, the training set comprises 250,000 samples, containing a randomly generated mix of quasistationary and transient noise time series. This data set includes both anomalous data (such as gravitational wave bursts and instrumental glitches from all test masses and interferometers) and nonanomalous data. The testing set, containing 20,000 samples, follows a similar composition. In the second stage, focusing on anomaly injection time prediction, the training and testing sets contain 250,000 and 20,000 samples, as well, consisting solely of anomalous events, as the first stage is tasked to filter out nonanomalous TDI segments.

The training sets for our pipeline were constructed without relying on prior knowledge or assumptions about the distributions and characteristics of astrophysical signals and instrumental glitches, although theoretical models and simulations could inform our understanding of their distribution and variability. In this initial stage of developing a machine-learning-based glitch detection and mitigation pipeline, we opted not to incorporate such priors into the training set generation process. Instead, a balanced training set is applied, ensuring equal representation across the various identified classes. This balanced representation minimizes the risk of bias towards more frequently occurring signals or glitches, thereby supporting the model's generalizability. Although leveraging realistic knowledge about the glitch and gravitational wave burst distributions in the training dataset could amplify the model's performance in practice, this approach requires a thorough examination to understand its performance impact, demanding rigorous scrutiny. It will be considered at a later stage.

The determination of training set sizes in this paper is based on comprehensive evaluations. We assessed the impact of varying training sizes on testing performance, focusing on identifying a threshold beyond which no significant improvements were observed. The selected training set sizes represent these thresholds, that is, the minimum necessary training set size to achieve robust model performance without yielding increasing returns from further data augmentation. This approach reflects a pragmatic strategy, aiming to balance computational demands and the performance of the machine learning pipeline.

During both training phases, we employ the Adam optimizer, setting the learning rate at a constant 0.0005. The loss function is the mean squared error (MSE), and the network undergoes training over 50 epochs with a batch size of 16. The MSE serves as a commonly used metric in assessing the performance of models, particularly within regression analysis. The MSE can be derived from the square of the Euclidean distance between two vectors in a multidimensional space. These vectors represent the true anomaly injection times converted to horizontal pixel positions along the width axis of the respective input image, x_i , and the predicted horizontal pixel positions, \hat{x}_i , across the instances of the dataset. The MSE, denoted by L, for a dataset with n samples is defined as follows:

$$L = \frac{1}{n} \sum_{i=1}^{n} (\hat{x}_i - x_i)^2.$$
(42)

Here, $(\hat{x}_i - x_i)^2$ represents the squared Euclidean distance for each prediction. The overarching objective of minimizing the MSE during model training translates into reducing the Euclidean distance between the model's predictions and the actual anomaly injections.

1. True negative and true positive detections

Considering the testing data set of 20,000 TTIconverted TDI samples, the CNN exhibits a 100% rate for correctly detecting nonanomalous data as nonanomalous; 99.99915% of anomalous events are detected as such, while 0.00085% of anomalous data are misclassified as nonanomalous.

2. Estimation of anomaly injection times

The proficiency of the CNN in estimating anomaly injection times is presented in Fig. 6 for the examples of Fig. 5. The estimation accuracy for Fig. 6(a) is remarkable. To validate this performance, comparable scenarios were tested. Three of them are depicted in the Appendix. Overall, the estimation error maintains a zero mean with a standard deviation of 3.995 sec, considering the testing data set of the second stage, which comprises 20,000 samples.

Figure 7 presents two scenarios for cases with larger estimation errors. The primary challenge in injection time estimation arises with slow-rising anomalies, which are less distinct in their initial phases due to the differentiating behavior of TDI 2.0. It turns out that pinpointing the onset of such anomalies becomes difficult if the glitch peak is reached more than 20 sec postinjection. Conversely, anomalies with less than 20 sec rising time show estimation errors of less than 1 sec owing to more substantial changes in the statistical properties of $E_Y(t)$ and $E_Z(t)$ around the anomaly injection. This also holds for artifacts embedded in quasistationary noise that remain undetectable through visual inspection.

Besides slow-rising artifacts, the network's estimation accuracy is impacted by optical displacement glitches from LISA's reference interferometers. These glitches are labeled



FIG. 6. Estimated anomaly injection times for the four scenarios presented in Fig. 5, contrasted with the corresponding ground truth data.



FIG. 7. Estimated anomaly injection times compared to ground truth data in scenarios with larger estimation errors.

in the training data according to their injection time in the interferometer, but they manifest in the TDI data with an 8.3-sec delay, increasing network confusion when handling various types of artifacts simultaneously.

In conclusion, while the CNN effectively detects anomalies and estimates injection times, determining the onset of faint and slow-rising artifacts with subsecond accuracy is challenging for the current CNN design. Note that the subsequent steps of the glitch mitigation pipeline will operate based on the estimation performance presented here.

IV. ANOMALY CHARACTERIZATION

In the previous section, we detected anomalies in LISA's TDI data. Building upon this, the focus shifts to the task of anomaly characterization. This involves differentiating between astrophysical sources and instrumental glitches and accurately identifying the source of these glitches. Such a source could be either a test mass glitch (with the need to specify the particular test mass involved) or an optical displacement glitch (pinpointing the exact interferometer and optical bench). We integrate a neural network with a glitch template matching methodology to address this challenge, leveraging the glitch TDI response models derived in Sec. II D.

A. Glitch template matching

In gravitational wave astrophysics, template matching is a systematic approach for contrasting observed data with a predefined set of waveforms, either theoretical or empirical [51,52]. This process is instrumental in recognizing and characterizing potential gravitational wave events. We adapt this method to the context of glitch localization within LISA's TDI data. Therefore, we introduce the glitchto-noise ratio (GNR) ρ in Eq. (43) analogously to the LISA signal-to-noise ratio provided by [53]:

$$\rho = \frac{\langle d(t)|g(t,\theta)\rangle}{\sqrt{\langle g(t,\theta)|g(t,\theta)\rangle}}.$$
(43)

Here, the scalar product $\langle x(t)|y(t)\rangle$ between two timedomain signals x(t) and y(t) is defined by

$$\langle x(t)|y(t)\rangle = 4\mathcal{R}\left(\int_0^\infty \frac{\tilde{x}(f)\tilde{y}^*(f)}{S(f)}df\right),\qquad(44)$$

with $\tilde{x}(f)$ representing the Fourier transform of x(t), $\tilde{x}^*(f)$ as the complex-conjugate of $\tilde{x}(f)$, and S(f) signifying the one-sided noise power spectral density. In this context, the data d(t) refers to the TDI variable of interest, while $g(t, \theta)$ denotes the glitch response in d(t) characterized by a set of parameters θ .

In the case of a deterministic FRED glitch, it would be possible to infer the intrinsic FRED model parameters $\theta = \{t_{inj}, t_{rise}, t_{fall}\}$ of Eq. (1) from d(t) when knowing the glitch injection point. For example, the parameter inference could be performed through maximum likelihood estimation utilizing Eq. (43) as proposed in [54] for astrophysical targets.

When dealing with the FREDS² model and an unknown glitch source, we confront a more intricate challenge. To determine the glitch source, we first approximate the $FREDS^2$ glitch of Eq. (5) by the FRED model of Eq. (1) adopting the average of the expected minimum and maximum values for the rising and falling time parameters, along with the injection time estimate derived from the preceding section. This results in a very rough approximation of the actual glitch transient. Then, the approach involves evaluating Eq. (43) for all 24 feasible glitch sources across the TDI variables X(t), Y(t), and Z(t)knowing that the Fourier transform of $q(t, \theta)$ is given by the Fourier transform of the glitch approximation and the TDI response models derived in Sec. II D. This computation yields a GNR triplet $\{\rho_X, \rho_Y, \rho_Z\}$ for each glitch injection point.

Selecting the triplet with the highest aggregate GNR could serve as a starting point for predicting the glitch origin. This method will prove insufficient given the stochastic nature of glitch shapes, substantial uncertainties in the model parameters and the model itself, as well as quasistationary noise. Therefore, we feed the 24 GNR triplets into a hybrid classification neural network, which maps the inputs to specific interferometers or test masses or identifies them as gravitational wave bursts.

B. Hybrid neural network model

The neural network employed for anomaly characterization is hybrid in two key respects. Firstly, it incorporates model information by utilizing frequency-domain TDI glitch responses from Sec. II D while avoiding explicit waveform models for classifying gravitational wave bursts. This design enables the network to accommodate unexpected burst morphologies without relying on assumptions for waveform parameters. Secondly, the network combines wide and deep learning components, forming a so-called wide and deep neural network (WDNN). The WDNN architecture merges the strengths of both wide and deep



FIG. 8. WDNN architecture and attributes employed for anomaly characterization.

components: the wide part consists of a linear model that directly connects input features to the output, excelling in memorizing feature interactions, while the deep component comprises multiple layers of nonlinear transformations, adept at capturing complex feature interactions [55]. The architecture of the WDNN is illustrated in Fig. 8.

The neural network model comprises a wide and deep component. The wide component has two layers with 64 and 32 neurons: ReLU activation, L2 regularization, and batch normalization. The deep component consists of three layers with 64, 64, and 32 neurons, using ReLU activation, L2 regularization, and batch normalization, as well. L2 regularization mitigates overfitting by adding a penalty term to the loss function, encouraging smaller weight magnitudes. Batch normalization improves training stability and accelerates convergence by normalizing the activations within each sub-batch, reducing internal covariate shift and enabling higher learning rates during training. The final layer employs the softmax activation function. The softmax activation function is given by

$$f(x_i) = \frac{\exp(x_i)}{\sum_{j=1}^{n} \exp(x_j)},\tag{45}$$

for $x_i \in X$, where $X \in \mathbb{R}^n$ denotes an input vector with $n \ge 1$. Equation (45) is commonly used in the context of classification tasks [56,57]. It takes the raw output scores from the previous network layers and converts them into a probability distribution over multiple classes, justifying its application for the glitch characterization network.

C. Characterization performance

To train and test the WDNN, TDI data sets are generated, simulating both glitch and gravitational wave events with properties described in the previous section. The GNR triplets calculated for each TDI glitch model of each simulation run are normalized and labeled using one-hot encoding according to Table III in the Appendix. One-hot encoding categorizes each sample into one of 25 classes, encompassing 24 possible glitch injection points and a class for gravitational wave bursts. We apply a data set comprising 150,000 training samples and 20,000 testing samples for the network's training and evaluation.

The performance of the network is gauged using a confusion matrix, which offers a detailed view of the model's classification accuracy across the existing classes. The matrix is presented in Fig. 9. We find gravitational wave bursts correctly classified in 98.187% of all cases, while 0.065% of FREDS² glitches are misclassified as gravitational wave bursts. Overall, the glitch injection points are correctly characterized with a 92.085% rate. The hybrid network model demonstrates a remarkable increase in classification correctness for stochastic glitch morphologies compared to when selecting the injection point without machine learning solely based on the highest GNR triplet sum. In this case, only 68.613% of the glitch origins would be correctly identified.

The distinctive pattern observed in the three off diagonals in Fig. 9 reveals misclassifications of optical displacement glitches between the interspacecraft interferometer ij and the reference interferometer *ji* of two opposite spacecraft. Optical benches of opposite spacecraft inject their glitches into the same bidirectional laser link. Consequently, they share the same glitch-unaffected TDI variable, such as TDI 2.0 Michelson Z(t) for glitches originating from optical benches 12 and 21, see Table II. Assuming equal arm lengths, the response to a glitch of the interspacecraft interferometer 12 in X(t) is given by $-(\mathcal{D}^2 - 1) \times$ $(\mathcal{D}^4 - 1)$, while the response to a glitch of the reference interferometer 21 in X(t) would be $\mathcal{D}(\mathcal{D}^2 - 1)(\mathcal{D}^4 - 1)$. Here, the primary distinction lies in the additional delay \mathcal{D} in the response related to the reference interferometer 21. Considering that the standard deviation of the estimated injection times is 3.995 sec, and by that approximately half the order of the light travel times between two spacecraft, there are instances in which an optical glitch of the interspacecraft interferometer *ij* may resemble an optical glitch of the reference interferometer *ji* for the network. This occurs when the injection time has been overestimated. Similarly, an optical glitch of the reference interferometer *ji* can be mistaken for an optical glitch of the interspacecraft interferometer ij when the injection time has been underestimated.



FIG. 9. Confusion matrix demonstrating the classification accuracy of the hybrid network for the classes outlined in Table III. The matrix differentiates between the true classes, represented by rows, and the predicted classes, indicated by columns, offering a detailed assessment of the model's ability to classify anomalous TDI data segments.

Note that confusion between, for instance, optical glitches of the interspacecraft interferometer 12 and reference interferometer 12 are generally avoided due to their TDI response difference characterized by two delays $\mathcal{D}_i =$ $\mathcal{D}_{1,2,2,1} = \mathcal{D}^2$ (where the latter notation applies for the equal arm length assumption), as outlined in Tables I and II. Such confusion would only arise if the injection time estimation error were to increase to the order of two times the light travel time, which, fortunately, is not the case. On the contrary, the misclassification of optical glitches between the interspacecraft and reference interferometer of opposite optical benches would diminish with further enhancements of the CNN's performance. For illustration, another confusion matrix is provided in Fig. 10, with the difference to Fig. 9 being a reduction in the standard deviation of the injection time estimation error from 3.995 sec (the performance of the previous section) to 3.0 sec.

Finally, in Fig. 11, we utilize a semilogarithmic representation of the receiver operating characteristic (ROC) curve, adapted for multiclass classification using the onevs-rest approach. Similar to confusion matrices, ROC curves are a commonly employed metric for evaluating the performance of classification networks and are consequently presented here for reference. They visually represent the model's proficiency for binary classification in distinguishing between positive and negative classes across



FIG. 10. Confusion matrix for a standard deviation error in the injection times reduced from 3.995 sec of Fig. 9 to 3.0 sec. This reduction results in a decrease of misclassifications within the off diagonals.

a spectrum of thresholds, plotting the true positive rate against the false positive rate. The area under the curve (AUC) measures the classifier's overall efficacy. An AUC value of 0.5 represents random guessing, while an AUC of 1.0 indicates perfect classification. In the one-vs-rest method, described, for example, in [58], each class is treated as positive, with the rest as negative, facilitating an assessment of the model's ability to differentiate between various anomaly types.

The ROC curves of Fig. 11 are categorized into three groups. Group I specifically relates to the misclassifications represented by the off-diagonal elements in Fig. 9, occurring at a false positive rate of less than 0.04.

V. GLITCH MITIGATION

After determining the glitch characteristics, including the identification of the injection point and timing through machine learning, we move to the final step, that is, glitch mitigation. The section focuses on estimating and removing the glitch from the TDI data. A LSTM network plays the key role in this process.

A. Long short-term memory neural networks

LSTM networks are a subset of recurrent neural networks (RNNs), distinguishing themselves from traditional feed-forward neural networks. In contrast to the onedirectional flow of information from input to output in



FIG. 11. Semilogarithmic ROC curves representing the model's proficiency for binary classification using the one-vs-rest approach. Translations of class labels are provided in Table III. Group I corresponds to the off-diagonal elements in Fig. 9.

feed-forward networks, RNNs, including LSTMs, incorporate looping connections that preserve a hidden state or memory of prior inputs [59,60]. In contrast to CNNs, which are proficient in spatial recognition within images, LSTMs can make continuous predictions of time series. LSTMs are distinct from standard RNNs in their ability to circumvent long-term dependency. The core of their architecture lies in a system of gates-specifically, the input, forget, and output gates. These gates function as regulators of data flow within the network, enabling LSTMs to make informed decisions about what information to store, update, or discard at each step in the sequence [61]. This selective memory process could make LSTMs adept at handling time-dependent and stochastic glitch patterns in TDI data. The input gate controls the extent to which new information flows into the cell state, the internal memory of the network. The forget gate decides which information is no longer relevant and should be removed from the cell state. Finally, the output gate determines what part of the cell state should be used to compute the output at each time step. This layered approach to information processing allows LSTMs to preserve relevant historical data.

The LSTM network applied in this paper for glitch time series prediction is detailed in Fig. 12. It consists of multiple stacked LSTM layers, each with 50 neurons and tanh activation functions. The network culminates in a dense layer with a single neuron for continuous value prediction. It takes as input the two TDI variables affected by the glitch and produces the glitch dynamics responsible for the observed TDI data, essentially functioning as a TDI inverter from the perspective of glitches.

B. Differential evolution of FRED model parameters

The FREDS² model of Eq. (5), blending deterministic FRED characteristics with stochastic fluctuations, presents considerable complexity. Therefore, the LSTM network shall primarily focus on the stochastic elements while incorporating as much *a priori* knowledge of the deterministic glitch properties as possible. Since the glitch origin is determined by the probability presented in the previous section, we can now refine our understanding of the deterministic glitch components. Therefore, differential evolution is employed for optimizing the intrinsic FRED



FIG. 12. LSTM architecture and attributes employed for glitch time series prediction. The parameter *c* represents the length of the TDI inputs, which depends on the FRED model parameters acquired pretraining via differential evolution optimization.

model parameters t_{rise} and t_{fall} of Eq. (1). This step is vital for providing a more accurate description of the glitch's deterministic nature, which the LSTM network will incorporate into its training process.

The optimization is performed by maximizing the GNR, as defined in Eq. (43). Remember, $g(t, \theta)$ denotes the TDI response to the glitch. Therefore, the glitch injection point is essential for this computation. By integrating all the information from the previous steps of the pipeline, we gain a comprehensive understanding of the glitch characteristics, which are to be made available to the LSTM. This includes selecting TDI variables based on the injection point and guiding the training with tailored data sets. The intrinsic FRED model parameters t_{rise} and t_{fall} are estimated similarly to the approach presented in [53] for estimating parameters of galactic binaries via differential evolution. For that, we make use of the implementation provided by SciPy [62].

C. Mitigation performance

Due to the complexity of the glitch shapes focused on in this paper, the LSTM network of Fig. 12 is retrained whenever there is a change in the glitch's injection point or FRED parameters. This process is computationally intensive but leads to accurate predictions. We suggest a secondary strategy to enable a quicker response for high-demand scenarios in a low-latency sense: training the LSTM with a data set encompassing a range of FRED model parameters. This approach offers ready-to-deploy generalized networks with a compromise in accuracy, as presented later.

The network is trained using the Adam optimizer with a learning rate of 0.001 and the MSE as the loss function.

Model checkpoints save the weights with the lowest loss on a validation set, ensuring peak performance. This approach is consistent with those followed for the first two neural networks in our pipeline.

A first example showcasing the glitch prediction performance is presented in Fig. 13. The figure depicts the glitch time series prediction and the TDI inputs applying the high-accuracy LSTM model. The model was trained using a data set comprising 100,000 samples. The TDI training data set exclusively incorporates quasistationary noise and FREDS² glitches, with t_{rise} and t_{fall} estimates derived from the previously explained differential evolution process as well as the t_{inj} estimate obtained from the CNN. The t_{rise} and t_{fall} model parameters correspond to the glitch depicted in Fig. 6(a). Note that the preceding step in the pipeline successfully identified this anomaly as a glitch within the interspacecraft interferometer of optical bench 12.

The spectrograms of X(t) for the scenario of Fig. 13 are presented in Fig. 14. Spectrograms serve as a visual representation of the temporal evolution of frequency content in a nonstationary signal. The figure illustrates the spectrogram before and after calibration for the achieved estimation performance. It also includes the spectrogram of



FIG. 13. LSTM neural network prediction for an optical displacement glitch injected in the interspacecraft interferometer of optical bench 12, alongside the true glitch and the LSTM inputs.



FIG. 14. Semilogarithmic time-frequency spectrograms for the scenario depicted in Fig. 13, showcasing the glitch-free, glitch-uncalibrated, and glitch-calibrated data using the glitch-affected TDI Michelson X(t) channel as an illustrative example. Brighter colors indicate increased noise levels. The vertical red line denotes the moment of glitch injection.



FIG. 15. Performance comparison of the LSTM neural network in predicting the glitch of Fig. 13 under conditions of amplified stochastic fluctuations. The TDI inputs for the LSTM predictions are not depicted. The spectrograms for (a)–(d) can be found in the Appendix.

the glitch-free case, representing an ideal scenario with perfect glitch mitigation. The glitch appears prominently as a bright yellow vertical stripe before calibration. It is worth noting that for frequencies greater than 1 Hz, the TDI 2.0 zero locations appear as faint horizontal blue lines. These zeros are not discernible for frequencies below 1 Hz due to the semilogarithmic plotting and particular spectrogram calculation. To make these zero locations visible in the lower frequency range, one needs to increase the segment length of the Fourier calculation. Nonetheless, such an adjustment would further blur the contribution of the glitch. The selected segment length was carefully chosen to optimize the appearance of the glitch.

Figure 15 presents additional realizations of this glitch with amplified levels of stochastic fluctuations. The corresponding spectrograms can be found in the Appendix.

Figure 16 provides a comparison with predictions obtained from the low-latency LSTM neural network. Although the quick-response variant may not attain the same level of precision as the high-accuracy alternative, it



FIG. 16. Performance evaluation of the LSTM neural network in the quick response mode considerable for low-latency pipeline applications with reduced emphasis on accuracy. The spectrograms for (e)-(g) can be found in the Appendix.

stands as a practical option for TDI glitch calibration in real-time situations requiring swift action and alert services. Additionally, it may offer benefits in Athena-LISA interactions, as suggested in [63].

In the final step of the calibration pipeline, the prediction for the specific glitch under consideration is subtracted from the affected interferometer measurement. Subsequently, the glitch-calibrated TDI data streams are calculated. To evaluate the effectiveness of the mitigation process, we subject the glitch-calibrated TDI data to a reanalysis by the anomaly detection step. In 93.93% of cases, the network classifies the calibrated data set as anomaly free. One could explore the potential of iterating through the pipeline to increase this value further. In each iteration, a prediction of the (residual) glitch in the underlying TDI channels would be generated. The final step would involve subtracting the sum of these predictions from the original data.

Training the LSTM network requires significantly more time than training the other two network architectures. As a result, we have yet to perform tests for all glitch injection points and the entire parameter space. So far, we have run two additional tests, one involving a test mass acceleration glitch associated with optical bench 32 and another focusing on an optical displacement glitch related to the reference interferometer of optical bench 13. We conducted a thorough analysis of all three scenarios using two different glitch parameter sets for each, one with a 20-sec duration and the other lasting 13 min. All three scenarios exhibited comparable glitch residuals, with the 13-min glitch predictions slightly less accurate. To assess the overall mitigation performance, further simulations must be conducted, encompassing various glitch types and a broader range of deterministic FRED model parameters. Additionally, it would be advantageous to compare the residual glitch noise against a specific requirement to ascertain the pipeline's global performance.

VI. CONCLUSION

The study introduced a data calibration pipeline designed to detect, characterize, and mitigate glitches for the LISA mission. The pipeline is structured around an ensemble of three neural networks. An overview of the glitch detection and mitigation process is given in Fig. 17.

The first component of our approach employs a CNN with the capability to identify anomalies within LISA's TDI time series data. The CNN correctly identified nonanomalous data 100% of the time and achieved a 99.99915% accuracy in detecting anomalous events. It faced challenges in precisely pinpointing the onset of faint and slow-rising anomalies.

In the second phase, we employed a hybrid neural network, which distinguishes between gravitational wave bursts and glitches while accurately identifying the locations where glitches have been injected. The hybrid neural



FIG. 17. Outline of the glitch detection and mitigation process as detailed in this paper, with the steps involving machine learning highlighted in dark blue.

network correctly classified gravitational wave bursts in 98.187% of cases and characterized glitch injection points with a 92.085% success rate, surpassing a conventional reference method by roughly 20%.

Then, the calibration pipeline incorporates a LSTM neural network to estimate glitch time series based on noisy TDI input data and information obtained by the preceding networks. The LSTM network is designed to operate in a high accuracy or quick response mode, depending on the configuration of the training data sets.

The final step involves subtracting the estimated stochastic glitch transient from the affected interferometric measurement to obtain glitch-calibrated TDI variables. A reanalysis of the calibrated TDI data set was performed to verify the overall performance of our approach. The calibrated TDI data streams were classified as anomaly free at 93.93%. Iteratively executing the pipeline could further increase this value and is under investigation.

Future efforts will focus on tackling the challenge of overlapping gravitational wave signals and glitches, as this is another area that demands careful attention. Moreover, the exploration of tensor processing units (TPUs) is underway. TPUs are specialized hardware accelerators designed for machine learning applications. Renowned for their processing power and efficiency in managing computational workloads, they are expected to further enhance the performance of the calibration pipeline, primarily by enabling the handling of larger training data sets and expediting the evaluation of diverse network architectures.

ACKNOWLEDGMENTS

The authors thank Michele Vallisneri, Stefan Strub, Martina Muratore, and Jean-Baptiste Bayle for the valuable discussions. Moreover, Jean-Baptiste Bayle and Martin Staab for their practical support with lisagwresponse and PyTDI. The authors acknowledge support from GW-Learn, a project funded through a Sinergia grant from the Swiss National Science Foundation.

APPENDIX

We conducted additional evaluations to validate the performance of anomaly injection time estimation under

conditions similar to those in Fig. 6(a). Figure 18 displays three such comparable scenarios where visual detection of the exact injection time is challenging but managed by the CNN of Fig. 4. The spectrograms for the TDI 2.0 Michelson X(t) data sets of Figs. 15 and 16 are given by Figs. 19–25. The translation of neural network labels is given by Table III.



FIG. 18. Examples of estimated anomaly injection times for glitch characteristics comparable to Fig. 6(a).



FIG. 19. Spectrograms of X(t) for case (a) of Fig. 15.



FIG. 20. Spectrograms of X(t) for case (b) of Fig. 15.



FIG. 21. Spectrograms of X(t) for case (e) of Fig. 16.



FIG. 22. Spectrograms of X(t) for case (c) of Fig. 15.



FIG. 23. Spectrograms of X(t) for case (f) of Fig. 16.



FIG. 24. Spectrograms of X(t) for case (d) of Fig. 15.



FIG. 25. Spectrograms of X(t) for case (g) of Fig. 16.

Label	Optical bench	Injection point		
0	12	Interspacecraft interferometer		
1	12	Test mass interferometer		
2	12	Reference interferometer		
3	13	Interspacecraft interferometer		
4	13	Test mass interferometer		
5	13	Reference interferometer		
6	21	Interspacecraft interferometer		
7	21	Test mass interferometer		
8	21	Reference interferometer		
9	23	Interspacecraft interferometer		
10	23	Test mass interferometer		
11	23	Reference interferometer		
12	31	Interspacecraft interferometer		
13	31	Test mass interferometer		
14	31	Reference interferometer		
15	32	Interspacecraft interferometer		
16	32	Test mass interferometer		
17	32	Reference interferometer		
18	12	Test mass		
19	23	Test mass		
20	31	Test mass		
21	13	Test mass		
22	21	Test mass		
23	32	Test mass		
24	Grav	Gravitational wave burst		

TABLE III. Translation of neural network labels.

- P. Amaro-Seoane, H. Audley, S. Babak, J. Baker, E. Barausse, P. Bender, E. Berti, P. Binetruy, M. Born, D. Bortoluzzi *et al.*, Laser interferometer space antenna, arXiv:1702.00786.
- [2] J. Thorpe, R. J. Cruz, M. Hartman, and G. Mueller, Armlocking in a LISA-like hardware model: A status report, AIP Conf. Proc. 873, 661 (2006).
- [3] D. Shaddock, Space-based gravitational wave detection with LISA, Classical Quantum Gravity 25, 114012 (2008).
- [4] S. Dhurandhar, K. R. Nayak, and J. Vinet, Time-delay interferometry for LISA with one arm dysfunctional, Classical Quantum Gravity 27, 135013 (2010).
- [5] W. Martens and E. Joffre, Trajectory design for the ESA LISA mission, J. Astronaut. Sci. 68, 402 (2021).
- [6] J.-B. Bayle, M. Lilley, A. Petiteau, and H. Halloin, Analytic model and simulations of residual laser noise after timedelay interferometry in LISA, instrumentation and methods for astrophysics (2018), https://api.semanticscholar.org/ CorpusID:126324850.
- [7] J.-B. Bayle, M. Lilley, A. Petiteau, and H. Halloin, Effect of filters on the time-delay interferometry residual laser noise for LISA, Phys. Rev. D 99, 084023 (2018).

- [8] M. Tinto, F. B. Estabrook, and J. W. Armstrong, Timedelay interferometry for LISA, Phys. Rev. D 65, 082003 (2002).
- [9] M. Tinto and S. V. Dhurandhar, Time-delay interferometry, Living Rev. Relativity **8**, 4 (2005).
- [10] M.-S. Hartig, S. Schuster, G. Heinzel, and G. Wanner, Nongeometric tilt-to-length coupling in precision interferometry: Mechanisms and analytical descriptions, J. Opt. 25, 055601 (2023).
- [11] M.-S. Hartig, S. Schuster, and G. Wanner, Geometric tilt-tolength coupling in precision interferometry: Mechanisms and analytical descriptions, J. Opt. 24, 065601 (2022).
- [12] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, A. M. Cruise *et al.*, Beyond the required LISA free-fall performance: New LISA Pathfinder results down to 20 μHz, Phys. Rev. Lett. **120**, 061101 (2018).
- [13] T. Robson and N. J. Cornish, Detecting gravitational wave bursts with LISA in the presence of instrumental glitches, Phys. Rev. D 99, 024019 (2019).
- [14] G. de Vine, B. Ware, K. McKenzie, R. E. Spero, W. M. Klipstein, and D. A. Shaddock, Experimental demonstration

of time-delay interferometry for the laser interferometer space antenna, Phys. Rev. Lett. **104**, 211103 (2010).

- [15] S. Mitryk, V. Wand, and G. Mueller, Verification of timedelay interferometry techniques using the University of Florida LISA interferometry simulator, Classical Quantum Gravity 27, 084012 (2010).
- [16] S. Paczkowski, R. Giusteri, M. Hewitson, N. Karnesis, E. D. Fitzsimons, G. Wanner, and G. Heinzel, Postprocessing subtraction of tilt-to-length noise in LISA, Phys. Rev. D 106, 042005 (2022).
- [17] N. Houba, S. Delchambre, T. Ziegler, G. Hechenblaikner, and W. Fichter, LISA spacecraft maneuver design to estimate tilt-to-length noise during gravitational wave events, Phys. Rev. D 106 (2022).
- [18] N. Houba, S. Delchambre, G. Hechenblaikner, T. Ziegler, and W. Fichter, Time-delay interferometry infinity for tiltto-length noise estimation in LISA, Classical Quantum Gravity 40, 107001 (2023).
- [19] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, M. Cruise *et al.*, LISA Pathfinder, in *Lepton Photon Interactions at High Energies* (World Scientific, Singapore, 2020).
- [20] D. Bortoluzzi, L. Baglivo, M. Benedetti, F. Biral, P. Bosetti, A. Cavalleri, M. Da Lio, M. De Cecco, R. Dolesi, M. Lapolla, W. Weber, and S. Vitale, LISA Pathfinder test mass injection in geodesic motion: Status of the on-ground testing, Classical Quantum Gravity 26, 094011 (2009).
- [21] M. Armano, H. Audley, G. Auger, J. T. Baird, M. Bassan, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, M. Caleno *et al.*, Sub-femto-*g* free fall for space-based gravitational wave observatories: LISA Pathfinder results, Phys. Rev. Lett. **116**, 231101 (2016).
- [22] G. Wanner, Space-based gravitational wave detection and how LISA Pathfinder successfully paved the way, Nat. Phys. 15, 200 (2019).
- [23] M. Armano, H. Audley, J. Baird, P. Binetruy, M. Born, D. Bortoluzzi, E. Castelli, A. Cavalleri, A. Cesarini, V. Chiavegato *et al.* (LISA Pathfinder Collaboration), Transient acceleration events in LISA Pathfinder data: Properties and possible physical origin, Phys. Rev. D **106**, 062001 (2022).
- [24] Q. Baghi, N. Korsakova, J. Slutsky, E. Castelli, N. Karnesis, and J.-B. Bayle, Detection and characterization of instrumental transients in LISA Pathfinder and their projection to LISA, Phys. Rev. D 105, 042002 (2022).
- [25] A. Spadaro, R. Buscicchio, D. Vetrugno, A. Klein, D. Gerosa, S. Vitale, R. Dolesi, W. J. Weber, and M. Colpi, Glitch systematics on the observation of massive black-hole binaries with LISA, Phys. Rev. D 108, 123029 (2023).
- [26] LDC Working Group, LISA data challenge: Spritz, LISA-LCST-XXX-TN-001 (2022), https://lisa-ldc.lal.in2p3.fr/ static/data/pdf/LDC-manual-Spritz.pdf.
- [27] S. Hourihane, K. Chatziioannou, M. Wijngaarden, D. Davis, T. Littenberg, and N. Cornish, Accurate modeling and mitigation of overlapping signals and glitches in gravitational-wave data, Phys. Rev. D 106, 042006 (2022).
- [28] J.-B. Bayle, E. Castelli, and N. Korsakova, LISA glitch (Zenodo, 2021), 10.5281/zenodo.6452903.
- [29] B. C. Levy, Wiener process and white Gaussian noise, in Random Processes with Applications to Circuits and

Communications (Springer International Publishing, New York, 2019), p. 207–234.

- [30] N. Houba, Tilt-to-Length Noise Estimation and Reduction Algorithms for Spaceborne Gravitational-Wave Observatories, Fortschrittsberichte des Instituts für Flugmechanik und Flugregelung, Vol. 15 (Shaker Verlag, Düren, 2023), ISBN 9783844092813.
- [31] J.-B. Bayle, Simulation and data analysis for LISA instrumental modeling, time-delay interferometry, noisereduction performance study, and discrimination of transient gravitational signals, Ph.D. thesis, Universite Paris Diderot, Laboratoire Astroparticules et Cosmologie, 2019.
- [32] O. Hartwig, Instrumental modelling and noise reduction algorithms for the Laser Interferometer Space Antenna, 10.15488/11372 (2021).
- [33] N. Houba, S. Delchambre, T. Ziegler, and W. Fichter, Optimal estimation of tilt-to-length noise for spaceborne gravitationalwave observatories, J. Guid. Control Dyn. 45, 1 (2022).
- [34] R. Fleddermann, C. Diekmann, F. Steier, M. Tröbs, G. Heinzel, and K. Danzmann, Sub-pm $\sqrt{\text{Hz}^{-1}}$ non-reciprocal noise in the LISA backlink fiber, Classical Quantum Gravity **35**, 075007 (2018).
- [35] M. Otto, Time-delay interferometry simulations for the laser interferometer space antenna, 10.15488/8545 (2015).
- [36] M. Muratore, D. Vetrugno, and S. Vitale, Revisitation of time delay interferometry combinations that suppress laser noise in LISA, Classical Quantum Gravity 37, 185019 (2020).
- [37] LISA Science Study Team, LISA science requirements document, Technical Report ESA-L3-EST-SCI-RS-001, European Space Agency, 2018.
- [38] J.-B. Bayle, LISA GW response (Zenodo, 2022), 10.5281/ zenodo.6423435.
- [39] J.-B. Bayle, O. Hartwig, and M. Staab, LISA instrument (Zenodo, 2022), 10.5281/zenodo.6447389.
- [40] J.-B. Bayle and M. Staab, PyTDI (Zenodo, 2023), 10.5281/ zenodo.6351736.
- [41] C. León, Detecting anomalous payments networks: A dimensionality-reduction approach, Latin Am. J. Cent. Banking 1, 100001 (2020).
- [42] M. Imani, Anomaly detection from hyperspectral images using clustering based feature reduction, J. Indian Soc. Remote Sens. 46, 1389 (2018).
- [43] W. Rawat and Z. Wang, Deep convolutional neural networks for image classification: A comprehensive review, Neural Comput. 29, 2352 (2017).
- [44] L. Alzubaidi *et al.*, Review of deep learning: Concepts, CNN architectures, challenges, applications, future directions, J. Big Data 8, 53 (2021).
- [45] E. Golinko, T. Sonderman, and X. Zhu, Learning convolutional neural networks from ordered features of generic data, in 2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA) (IEEE, Orlando, Florida, 2018), pp. 897–900.
- [46] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin *et al.*, TensorFlow: Large-scale machine learning on heterogeneous systems (2015), software available from tensorflow.org.
- [47] L. Roeder, Netron, https://github.com/lutzroeder/netron (2023), Netron is a viewer for neural network, deep learning and machine learning models.

- [48] N. Galanis, P. Vafiadis, K.-G. Mirzaev, and G. Papakostas, Convolutional neural networks: A roundup and benchmark of their pooling layer variants, Algorithms 15, 391 (2022).
- [49] M. A. Mahmoudi, A. Chetouani, F. Boufera, and H. Tabia, Learnable pooling weights for facial expression recognition, Pattern Recogn. Lett. 138, 644 (2020).
- [50] T. Prince, M. Tinto, and S. Larson, LISA optimal sensitivity, Phys. Rev. D 66, 122002 (2002).
- [51] L. Sampson, N. Cornish, and N. Yunes, Mismodeling in gravitational-wave astronomy: The trouble with templates, Phys. Rev. D 89, 064037 (2013).
- [52] P. Ajith, S. Babak, Y. Chen, M. Hewitson, B. Krishnan, A. M. Sintes, J. T. Whelan, B. Brügmann, P. Diener, N. Dorband, J. Gonzalez, M. Hannam, S. Husa, D. Pollney, L. Rezzolla, L. Santamaría, U. Sperhake, and J. Thornburg, Template bank for gravitational waveforms from coalescing binary black holes: Nonspinning binaries, Phys. Rev. D 77, 104017 (2008).
- [53] S. H. Strub, L. Ferraioli, C. Schmelzbach, S. C. Stähler, and D. Giardini, Accelerating global parameter estimation of gravitational waves from Galactic binaries using a genetic algorithm and GPUs, Phys. Rev. D 108, 103018 (2023).
- [54] S. H. Strub, L. Ferraioli, C. Schmelzbach, S. C. Stähler, and D. Giardini, Bayesian parameter estimation of Galactic binaries in LISA data with Gaussian process regression, Phys. Rev. D 106, 062003 (2022).
- [55] H.-T. Cheng, L. Koc, J. Harmsen, T. Shaked, T. Chandra, H. Aradhye, G. Anderson, G. Corrado, W. Chai,

M. Ispir *et al.*, *Wide & Deep Learning for Recommender Systems* (Association for Computing Machinery, New York, NY, USA, 2016), pp. 7–10.

- [56] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, England, 2016).
- [57] S. R. Dubey, S. K. Singh, and B. B. Chaudhuri, Activation functions in deep learning: A comprehensive survey and benchmark, Neurocomputing 503, 92 (2022).
- [58] P. Gimeno, V. Mingote, A. Ortega, A. Miguel, and E. Lleida, Generalising AUC optimisation to multiclass classification for audio segmentation with limited training data, IEEE Signal Process. Lett. 28, 1135 (2021).
- [59] S. Hochreiter and J. Schmidhuber, Long short-term memory, Neural Comput. 9, 1735 (1997).
- [60] Y. Yu, X. Si, C. Hu, and J. xun Zhang, A review of recurrent neural networks: LSTM cells and network architectures, Neural Comput. 31, 1235 (2019).
- [61] F. A. Gers, J. Schmidhuber, and F. Cummins, Learning to forget: Continual prediction with LSTM, Neural Comput. 12, 2451 (2000).
- [62] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright *et al.*, sciPy 1.0: Fundamental algorithms for scientific computing in Python, Nat. Methods 17, 261 (2020).
- [63] A. C. Fabian, M. Guainazzi, P. Mcnamara, L. Piro, and N. R. Tanvir, Athena-LISA synergies (2019), https://api .semanticscholar.org/CorpusID:218551158.