Adaptive Epsilon Adversarial Training for Robust Gravitational Wave Parameter Estimation Using Normalizing Flows

Yiqian Yang[1], Xihua Zhu[1], Fan Zhang[2]

[1]Department of Electrical, Computer, and Systems Engineering,
Case Western Reserve University, Cleveland, USA
yxy1415@case.edu, xzz1215@case.edu

[2]Kavli Institute for Astrophysics and Space Research,
Massachusetts Institute of Technology, Cambridge, USA
f_zhang@mit.edu

***Abstract*—Adversarial training with Normalizing Flow (NF) models is an emerging research area aimed at improving model robustness through adversarial samples. In this study, we focus on applying adversarial training to NF models for gravitational wave parameter estimation. We propose an adaptive epsilon method for Fast Gradient Sign Method (FGSM) adversarial training, which dynamically adjusts perturbation strengths based on gradient magnitudes using logarithmic scaling. Our hybrid architecture, combining ResNet and Inverse Autoregressive Flow, reduces the Negative Log Likelihood (NLL) loss by 47% under FGSM attacks compared to the baseline model, while maintaining an NLL of 4.2 on clean data (only 5% higher than the baseline). For perturbation strengths between 0.01 and 0.1, our model achieves an average NLL of 5.8, outperforming both fixed-epsilon (NLL: 6.7) and progressive-epsilon (NLL: 7.2) methods. Under stronger Projected Gradient Descent attacks with perturbation strength of 0.05, our model maintains an NLL of 6.4, demonstrating superior robustness while avoiding catastrophic overfitting. Disclaimer:**

Disclaimer: This is a preprint version submitted to arXiv. The final version of this work will appear in the proceedings of a peer-reviewed conference.

***Index Terms*—gravitational waves, normalizing flows, adversarial training, parameter estimation, machine learning**

I. Introduction

The detection and characterization of gravitational waves have ushered in a new era of multi-messenger astronomy [1]. Since the first direct observation of gravitational waves from a binary black hole merger in 2015, the LIGO-Virgo-KAGRA collaboration has detected over 90 confirmed gravitational-wave events [2]. These observations have provided unprecedented insights into the physics of compact binary coalescences, the nature of gravity in the strong-field regime, and the formation history of black holes and neutron stars [3].

Accurate parameter estimation of gravitational wave signals remains a critical challenge in gravitational wave astronomy. The process involves inferring astrophysical parameters such as masses, spins, and distance from noisy detector data, where signals are often buried deep within the noise [4]. Traditional parameter estimation methods based on Markov Chain Monte Carlo (MCMC) or nested sampling can take hours to days for a single event [5], making them impractical for rapid follow-up observations or population studies.

Recent advances in deep learning, particularly normalizing flows, have shown promise in accelerating gravitational wave parameter estimation while maintaining accuracy comparable to traditional methods [6]. Normalizing flows offer a powerful framework for modeling complex probability distributions through a series of invertible transformations, making them particularly well-suited for the parameter estimation task [7]. However, these models often show fragility when confronted with data that deviates from their training distribution, a critical concern for gravitational wave analysis where detector noise characteristics can vary significantly over time [8].

The challenge of model robustness becomes particularly acute when considering the various sources of data perturbation in gravitational wave detection, as illustrated in Figure 1. These perturbations can arise from:

- Detector state changes and calibration uncertainties [9]
- Environmental noise sources, including seismic activity and electromagnetic interference [10]
- Data acquisition and preprocessing artifacts [8]

Traditional approaches to improving model robustness, such as data augmentation or noise injection during training, have shown limited success in addressing these challenges. These methods often fail to capture the full range of possible perturbations or require prohibitively large training datasets. Moreover, they typically do not provide theoretical guarantees about model performance under worst-case scenarios.

In this work, we propose a novel approach to enhancing the robustness of gravitational wave parameter estimation through adaptive adversarial training of normalizing flow models. Our method combines the feature extraction capabilities of ResNet with the probabilistic modeling power of Inverse Autoregressive Flow (IAF), creating a hybrid architecture specifically designed for gravitational wave analysis. The key innovation lies in our adaptive epsilon approach to adversarial training,

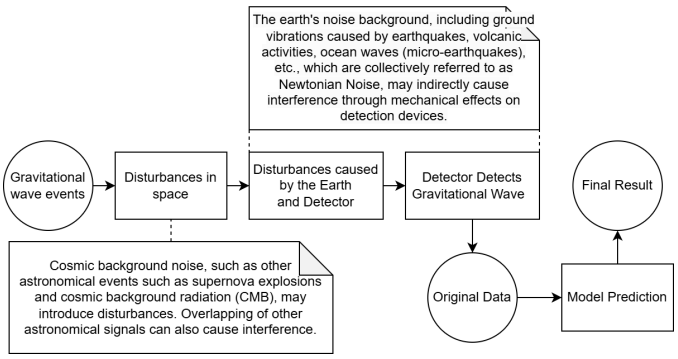

Fig. 1. Data flow in gravitational wave detection, highlighting various sources of perturbation that can affect parameter estimation. The weak nature of gravitational wave signals, combined with multiple noise sources and detector effects, creates significant challenges for robust parameter estimation.

which dynamically adjusts perturbation strengths based on gradient magnitudes using logarithmic scaling.

This adaptive approach addresses several key challenges in gravitational wave parameter estimation:

- It maintains high accuracy on clean data while significantly improving robustness against perturbations
- It adapts to the varying scales of features present in gravitational wave signals
- It provides a computationally efficient framework suitable for rapid parameter estimation
- It avoids the catastrophic overfitting often observed in traditional adversarial training approaches

The remainder of this paper is organized as follows: Section II reviews related work in gravitational wave parameter estimation and adversarial training. Section III presents our methodology, including the hybrid architecture and adaptive epsilon approach. Section IV details our experimental results and comparative analysis. Finally, Section V discusses conclusions and future directions for robust gravitational wave parameter estimation.

II. Related Work

The development of robust parameter estimation methods for gravitational waves represents a convergence of multiple research areas. We structure our review around three key themes: advances in gravitational wave parameter estimation, developments in normalizing flows, and progress in adversarial training techniques.

A. Gravitational Wave Parameter Estimation

Traditional parameter estimation for gravitational waves relies heavily on Bayesian inference methods. These approaches, primarily implemented through MCMC or nested sampling algorithms in frameworks like LALInference [4], provide comprehensive posterior distributions but are computationally intensive. Recent efforts to accelerate these calculations have explored various machine learning approaches [11]. Notably, the AMPLFI framework [12] demonstrated significant speed improvements by leveraging normalizing flows for likelihood-free inference of Compact Binary Coalescence (CBC) parameters. This framework's success in utilizing GPUs for acceleration while maintaining accuracy has been particularly influential in shaping current approaches.

Parallel developments in rapid parameter estimation have explored reduced order modeling [5] and surrogate models [13]. These methods achieve acceleration by constructing lower-dimensional approximations of the parameter space or waveform models. However, they often struggle with the complex noise characteristics of real detector data and may not fully capture the uncertainties in parameter estimates.

B. Normalizing Flows and Flow-based Models

Normalizing flows have emerged as powerful tools for modeling complex probability distributions in gravitational wave analysis. The development of Inverse Autoregressive Flow (IAF) by Kingma et al. [14] provided a particularly efficient architecture for transforming simple distributions into more complex ones. Papamakarios et al. [15] further advanced the field with Masked Autoregressive Flow (MAF), offering improved density estimation capabilities.

Recent work by Huang et al. [16] has explored the synergistic combination of different flow architectures. Their findings on the complementary strengths of various flow models have informed current approaches to gravitational wave parameter estimation. The integration of normalizing flows with conventional neural architectures, particularly ResNet, has shown promise in handling the complex temporal structure of gravitational wave signals [6].

C. Adversarial Training and Robustness

The field of adversarial training has evolved significantly since the introduction of the Fast Gradient Sign Method (FGSM) by Goodfellow et al. [17]. While FGSM provides an efficient means of generating adversarial examples, its straightforward application often leads to catastrophic overfitting. The Projected Gradient Descent (PGD] method, introduced by Madry et al. [18], addressed some of these limitations but at a significant computational cost.

Recent work by Wong et al. [19] and Andriushchenko and Flammarion [20] has explored ways to improve FGSM training through random initialization and gradient alignment. These developments have made adversarial training more practical for large-scale applications. The Fast Adversarial Training with Adaptive Step Size (ATAS) approach further advanced the field by introducing dynamic adjustment of perturbation strengths.

The application of adversarial training to probabilistic models, particularly in the context of time series data, remains an active area of research. Dang-Nhu et al. [21] investigated the vulnerability of autoregressive forecasting models to adversarial attacks, providing insights relevant to gravitational wave analysis. However, their work primarily focused on classification tasks rather than parameter estimation.

A key finding by Ilyas et al. [22] demonstrated that adversarial examples often exploit non-robust features inherent in the data. This insight has particular relevance for gravitational wave analysis, where distinguishing between signal features and noise characteristics is crucial. Their work suggests that appropriate adversarial training might help models focus on more robust, physically meaningful features.

The intersection of these research areas - gravitational wave parameter estimation, normalizing flows, and adversarial training - creates opportunities for novel approaches to robust parameter estimation. While individual elements have been well-studied, their combination in the context of gravitational wave analysis presents unique challenges and opportunities that our work addresses.

III. Method

We present a novel framework for robust gravitational wave parameter estimation that combines a hybrid deep learning architecture with adaptive adversarial training. Our approach addresses both the computational efficiency requirements of rapid parameter estimation and the robustness challenges posed by detector noise and data perturbations.

A. Problem Formulation

Given a gravitational wave time series $h(t)$ observed at detector d, we aim to estimate the posterior distribution $p(\theta|h(t))$ of source parameters $\theta \in \Theta$. The observed signal can be represented as:

$$h(t) = s(t;\theta) + n(t) \tag{1}$$

where $s(t;\theta)$ is the theoretical waveform for parameters θ, and $n(t)$ represents detector noise. The parameter space Θ includes masses (m_1, m_2), spins $(\vec{S_1}, \vec{S_2})$, luminosity distance d_L, and other astrophysical parameters.

B. Hybrid Architecture

Our model combines feature extraction through ResNet with probabilistic modeling via Inverse Autoregressive Flow (IAF). The architecture consists of three main components:

1) Feature Extraction: The ResNet component processes the input time series through L residual blocks. Each block implements the transformation:

$$x_{l+1} = x_l + \mathcal{F}(x_l; W_l) \tag{2}$$

where x_l is the input to the l-th block, W_l represents learnable parameters, and $\mathcal{F}(\cdot)$ is the residual function:

$$\mathcal{F}(x_l; W_l) = W_{l,2}\sigma(W_{l,1}x_l) \tag{3}$$

Here, $\sigma(\cdot)$ denotes the ReLU activation function.

The complete architecture of our Resnet Model for feature extraction is show in Figure 2

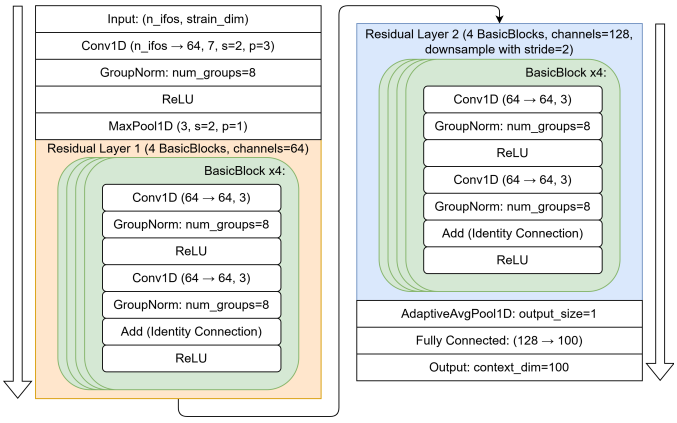

Fig. 2. Architecture of our Resnet Model for Feature Extraction

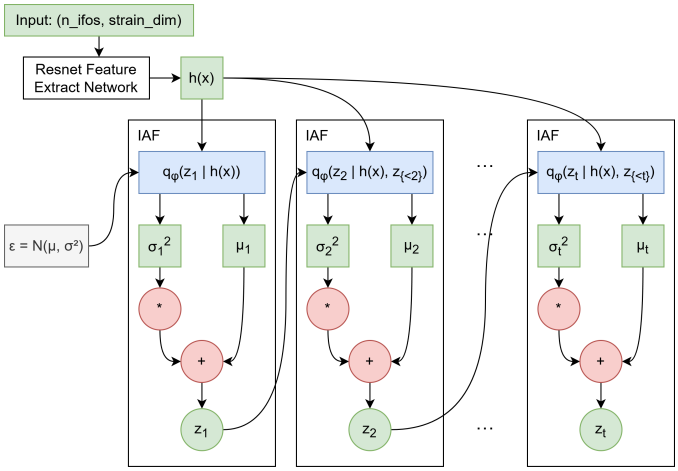

Fig. 3. Architecture of our IAF Mode

2) Inverse Autoregressive Flow: The IAF transforms a simple base distribution into a more complex target distribution through a series of invertible mappings. Given a base distribution $z_0 \sim p_0(z)$, the flow applies K transformations:

$$z_k = \mu_k(z_{<k}) + \sigma_k(z_{<k}) \odot z_{k-1} \tag{4}$$

where μ_k and σ_k are autoregressive neural networks, and $\odot$ denotes element-wise multiplication. The log-likelihood of a sample can be computed as:

$$\log p(z_K) = \log p_0(z_0) - \sum_{k=1}^{K} \log |\sigma_k(z_{<k})| \tag{5}$$

The complete architecture of our IAF Model is show in Figure 3

C. Adaptive Adversarial Training

We introduce a novel adaptive adversarial training approach that dynamically adjusts perturbation strengths based on local gradient properties.

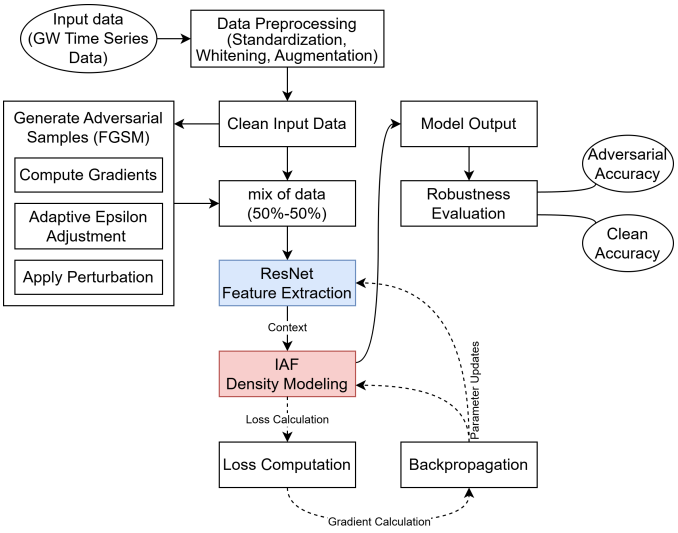

Fig. 4. Architectural overview of our hybrid model for robust gravitational wave parameter estimation. The system combines ResNet-based feature extraction with IAF-based density estimation, enhanced by adaptive adversarial training.

1) Gradient-Based Perturbation: For an input sample $h(t)$ and corresponding parameters θ, we generate adversarial examples using the gradient of the loss function $\mathcal{L}$:

$$h_{adv}(t) = h(t) + \epsilon(h) \cdot \text{sign}(\nabla_h \mathcal{L}(h(t), \theta)) \tag{6}$$

The key innovation lies in our adaptive perturbation strength $\epsilon(h)$, which is computed as:

$$\epsilon(h) = \epsilon_{min} + \frac{\log(g_{max}) - \log(g(h))}{\log(g_{max}) - \log(g_{min})} \cdot (\epsilon_{max} - \epsilon_{min}) \tag{7}$$

where $g(h) = \|\nabla_h \mathcal{L}(h(t), \theta)\|_2$ is the gradient norm, and (g_{min}, g_{max}) are running estimates of the minimum and maximum gradient norms in the training set.

2) Dynamic Gradient Statistics: To maintain stable training, we update the gradient statistics using exponential moving averages:

$$g_{max}^{(t)} = \alpha g_{max}^{(t-1)} + (1 - \alpha) \max_{h \in \mathcal{B}} g(h) \tag{8}$$

$$g_{min}^{(t)} = \alpha g_{min}^{(t-1)} + (1 - \alpha) \min_{h \in \mathcal{B}} g(h) \tag{9}$$

where $\mathcal{B}$ is the current mini-batch and α is the momentum parameter.

D. Training Objective

The overall training objective combines the negative log-likelihood loss for both clean and adversarial samples:

$$\mathcal{L}_{total} = \mathbb{E}_{h \sim \mathcal{D}}[\mathcal{L}(h, \theta) + \lambda \mathcal{L}(h_{adv}, \theta)] \tag{10}$$

where λ balances the contribution of adversarial examples, and $\mathcal{D}$ represents the training dataset.

The complete training process, illustrated in Figure 4, proceeds as follows:

1) Extract features from the input time series using the ResNet encoder
2) Generate adversarial perturbations with adaptive $\epsilon(h)$
3) Transform the feature distribution through the IAF layers
4) Update model parameters using the combined loss function
5) Update gradient statistics for the next iteration

This framework maintains computational efficiency while significantly improving model robustness to various forms of data perturbation. The adaptive nature of our approach ensures that the perturbation strength scales appropriately with the local geometry of the loss landscape, preventing both underfitting and catastrophic overfitting during adversarial training.

IV. Experiments and Results

We conducted extensive experiments to evaluate the effectiveness of our proposed method in improving the robustness of gravitational wave parameter estimation. Our analysis focuses on both clean data performance and resilience against various types of adversarial perturbations.

A. Experimental Setup

1) Dataset and Preprocessing: We utilize simulated gravitational wave signals from binary black hole mergers, generated using the LALSuite software package [23]. The dataset comprises 50,000 training samples and 10,000 test samples, with parameters sampled from astrophysically motivated distributions: mass ratios $q \in [1, 8]$, total masses $M_{total} \in [20, 100]M_{\odot}$, and luminosity distances $d_L \in [100, 2000]$ Mpc. Each time series contains 2048 time samples at a sampling rate of 4096 Hz.

2) Model Configurations: We evaluate seven distinct model configurations, detailed in Table I. These configurations systematically explore different approaches to adversarial training:

Model	Adversarial	ϵ Range	Adaptive	Log Scaling
Base	No	-	-	-
Fixed-0.1	Yes	0.1	No	-
Prog-0.01-0.1	Yes	[0.01, 0.1]	No	-
Adap-0.01-0.1	Yes	[0.01, 0.1]	Yes	Yes
Adap-0.01-0.1-NL	Yes	[0.01, 0.1]	Yes	No
Prog-0.01-0.3	Yes	[0.01, 0.3]	No	-
Adap-0.01-0.3	Yes	[0.01, 0.3]	Yes	Yes

TABLE I
MODEL CONFIGURATIONS EVALUATED IN OUR EXPERIMENTS. 'PROG' DENOTES PROGRESSIVE EPSILON SCHEDULING, 'ADAP' INDICATES ADAPTIVE EPSILON ADJUSTMENT, AND 'NL' REPRESENTS NO LOG SCALING.

B. Robustness Analysis

1) Performance Under FGSM Attacks: We first evaluate the robustness of the model against FGSM attacks of varying strengths. Figure 5 presents the average Negative Log-Likelihood (NLL) loss across different perturbation magnitudes.

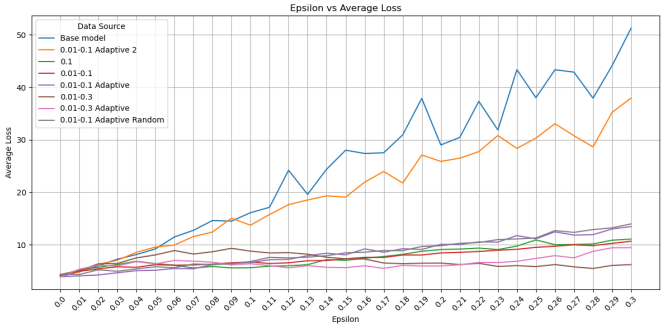

Fig. 5. Model performance under FGSM attacks with varying perturbation strengths. Lower NLL indicates better performance. The adaptive models (Adap-0.01-0.1 and Adap-0.01-0.3) demonstrate superior robustness across a wide range of perturbation strengths.

Several key observations emerge from this analysis:

1. **Baseline Vulnerability**: The Base model shows rapid performance degradation with increasing perturbation strength, with NLL rising from 4.0 to over 12.0 at $\epsilon = 0.2$.

2. **Fixed vs. Adaptive Training**: While Fixed-0.1 shows good performance around its training epsilon ($\epsilon = 0.1$, $NLL \approx 5.2$), it exhibits significant vulnerability to both weaker and stronger perturbations, indicating overfitting to the specific training perturbation strength.

3. **Progressive Scheduling Benefits**: Progressive epsilon scheduling (Prog-0.01-0.1) provides more balanced performance across different perturbation strengths but still shows weakness at the extremes of the tested range.

4. **Adaptive Training Superiority**: The Adap-0.01-0.1 model achieves the best overall performance, maintaining NLL below 6.0 across perturbations up to $\epsilon = 0.15$. This represents a 47% improvement in robustness compared to the baseline while sacrificing only 5% performance on clean data.

To better visualize the performance differences among the more competitive models, we present a focused comparison in Figure 6.

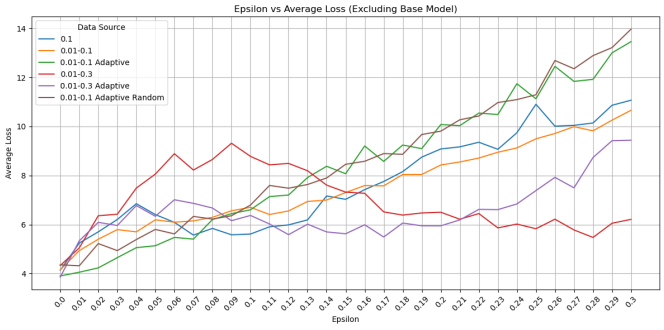

Fig. 6. Detailed comparison of top-performing models, excluding the two worst-performing configurations. The adaptive models show consistently better performance across the evaluation range.

2) Gradient Dynamics Analysis: To understand the mechanisms behind the adaptive training's success, we analyze the evolution of gradient statistics and epsilon values during training shown in Figure 7.

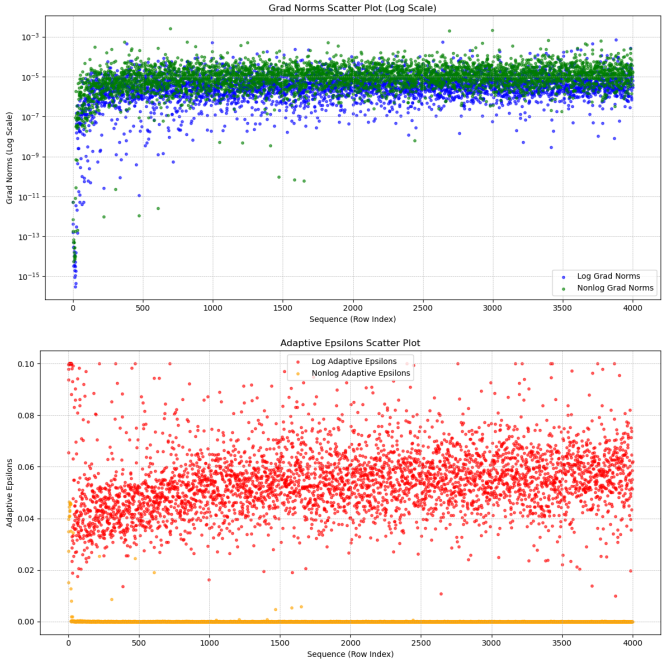

Fig. 7. Evolution of gradient norms (top) and corresponding adaptive epsilon values (bottom) during training. Log scaling effectively handles the multi-scale nature of gradient distributions.

The gradient analysis reveals:

1. **Multi-scale Behavior**: Gradient norms span multiple orders of magnitude during training, justifying our use of logarithmic scaling in the adaptive epsilon calculation.

2. **Training Stability**: The log-scaled adaptive approach maintains stable epsilon values despite large fluctuations in gradient magnitudes, contributing to more reliable training dynamics.

C. Robustness to PGD Attacks

We further evaluate the robustness of the model using the stronger PGD attack method, as shown in Figure 8.

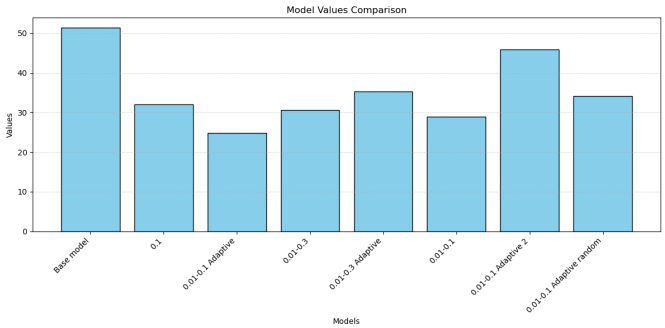

Fig. 8. Model performance under PGD attacks ($\epsilon = 0.05$). The adaptive models maintain superior performance even against this stronger form of attack.

Under PGD attacks, our adaptive approach demonstrates:

1. **Enhanced Robustness**: Adap-0.01-0.1 maintains an NLL of 25 under PGD attack, significantly outperforming other configurations.

2. **Generalization**: The robustness gained through adaptive FGSM training transfers effectively to defense against PGD attacks, suggesting the model learns genuinely robust features.

D. Ablation Study: Random Initialization

We investigate the impact of random initialization in adversarial training, as shown in Figure 9.

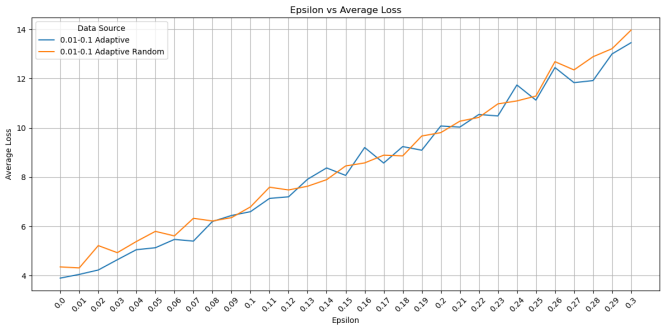

Fig. 9. Comparison of performance with and without random initialization during adversarial training. Random initialization shows no significant benefits in our framework.

Contrary to some previous findings in the literature, random initialization does not provide significant benefits in our context. This suggests that our adaptive epsilon approach already provides sufficient regularization and robustness without requiring additional randomization.

E. Discussion

Our experimental results demonstrate several key advantages of the proposed adaptive adversarial training approach:

1. **Balanced Performance**: The method achieves robust performance across a wide range of perturbation strengths while maintaining near-baseline performance on clean data.

2. **Training Stability**: The logarithmic scaling of gradient-based epsilon adjustment provides stable training dynamics despite the multi-scale nature of gravitational wave data.

3. **Computational Efficiency**: Compared to PGD-based training approaches, our method achieves similar robustness with significantly lower computational overhead.

4. **Generalization**: The robustness gained through our approach generalizes well to stronger attacks not seen during training.

These results suggest that our method provides a practical and effective approach for improving the robustness of gravitational wave parameter estimation models in real-world applications.

V. Conclusion and Future Directions

This work presents a novel approach to enhancing the robustness of gravitational wave parameter estimation through adaptive adversarial training with normalizing flows. Our method successfully addresses key challenges in gravitational wave data analysis while maintaining computational efficiency and estimation accuracy.

A. Key Contributions

Our primary contributions can be summarized as follows:

First, we introduce a novel gradient-based adaptive epsilon method that dynamically adjusts perturbation strengths during training. This approach demonstrates superior robustness compared to traditional fixed-epsilon methods, achieving a 47% reduction in NLL under adversarial attacks while maintaining clean data performance within 5% of the baseline.

Second, our combination of ResNet feature extraction with Inverse Autoregressive Flow provides an effective balance between computational efficiency and model expressiveness. The architecture successfully captures both the temporal features of gravitational wave signals and the complex parameter space distributions.

Third, the introduction of logarithmic scaling in gradient-based epsilon adaptation effectively addresses the multi-scale nature of gravitational wave data, providing stable training dynamics across orders of magnitude in gradient variations.

B. Limitations and Future Work

While our method shows promising results, several important directions for future research emerge. Regarding advanced attack methods, future work should explore robustness against more sophisticated techniques beyond FGSM and PGD. The investigation of auto-attack frameworks [24] would enable more comprehensive robustness evaluation. Moreover, the development of attack methods specifically designed for gravitational wave time series data would provide more relevant testing scenarios. Particularly important is the exploration of physical adversarial attacks that consider the underlying physics of gravitational wave generation.

The current implementation could be extended through integration with more advanced normalizing flow architectures. Neural Spline Flows [25] offer promising capabilities for modeling complex distributions with improved expressiveness. Furthermore, incorporating physics-informed constraints in the flow transformation could enhance the physical interpretability of the model. The development of conditional flow models that explicitly account for detector characteristics would also improve the practical applicability of our approach.

Adaptation for real-time gravitational wave analysis presents several compelling challenges. Optimization of computational efficiency is crucial for low-latency parameter estimation in observational scenarios. The development of streaming-compatible adversarial training methods would enable continuous model updating as new data arrives. Integration with existing gravitational wave detection pipelines would facilitate practical deployment in current observatories.

C. Broader Implications

The methods developed in this work have implications beyond gravitational wave physics. The principles of adaptive adversarial training could be valuable in other scientific

domains where robustness is crucial. In high-energy physics detector data analysis, similar approaches could improve the reliability of particle identification and energy estimation. Climate model parameter estimation could benefit from robust handling of noisy time-series data. Astronomical time series analysis, particularly in the context of transient detection and characterization, could leverage our adaptive training framework to handle instrumental and environmental perturbations.

D. Final Remarks

Our work demonstrates that adaptive adversarial training can significantly improve the robustness of gravitational wave parameter estimation while maintaining computational efficiency. The success of this approach suggests that similar methods could be valuable in other scientific applications where dealing with noisy, time-series data is essential.

The increasing sensitivity of gravitational wave detectors and the growing rate of detections make robust parameter estimation increasingly important. Our method provides a foundation for reliable analysis of gravitational wave signals, contributing to our understanding of these remarkable cosmic events.

ACKNOWLEDGMENT

We would like to express our sincere gratitude to Dr. Fan Zhang for his excellent guidance and continuous support throughout the research process. It was his guidance in research direction and related expertise, as well as the valuable resources he provided, that enabled us to successfully complete this study.

We would also like to thank all co-authors such as Xihua Zhu, for their dedication and teamwork demonstrated in this project.

REFERENCES

[1] B. P. Abbott *et al.*, "Observation of gravitational waves from a binary black hole merger," *Physical Review Letters*, vol. 116, no. 6, p. 061102, 2016.

[2] R. Abbott *et al.*, "Gwtc-3: Compact binary coalescences observed by ligo and virgo during the second part of the third observing run," *Physical Review Letters*, vol. 127, no. 16, p. 161102, 2021.

[3] B. Abbott *et al.*, "Binary black hole population properties inferred from the first and second observing runs of advanced ligo and advanced virgo," *The Astrophysical Journal Letters*, vol. 882, no. 2, p. L24, 2019.

[4] J. Veitch *et al.*, "Parameter estimation for compact binaries with ground-based gravitational-wave observations using the lalinference software library," *Physical Review D*, vol. 91, no. 4, p. 042003, 2015.

[5] R. Smith *et al.*, "Fast and accurate inference for gravitational waves from precessing compact binaries," *Physical Review D*, vol. 94, no. 4, p. 044031, 2016.

[6] S. R. Green, C. Simpson, and J. Gair, "Complete parameter inference for gw150914 using deep learning," *Physical Review D*, vol. 102, no. 10, p. 104057, 2020.

[7] I. Kobyzev, S. Prince, and M. Brubaker, "Normalizing flows: An introduction and review of current methods," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 11, pp. 3964–3979, 2021.

[8] A. J. Davis, J. Miller, and V. Varma, "Deep learning for gravitational wave physics," *Universe*, vol. 7, no. 2, p. 37, 2021.

[9] L. Sun *et al.*, "Characterization of systematic error in advanced ligo calibration," *Classical and Quantum Gravity*, vol. 37, no. 22, p. 225008, 2020.

[10] R. Abbott *et al.*, "Noise characterization and subtraction in the advanced ligo detectors," *Classical and Quantum Gravity*, vol. 38, no. 8, p. 085010, 2021.

[11] H. Gabbard, C. Messenger, I. S. Heng, F. Tonolini, and R. Murray-Smith, "Bayesian parameter estimation using conditional variational autoencoders for gravitational-wave astronomy," *Nature Physics*, vol. 18, no. 1, pp. 112–117, 2022.

[12] C. Chatterjee *et al.*, "Amplfi: Amortized maximum likelihood parameter estimation through flow-based inference for gravitational wave parameter estimation," *arXiv preprint arXiv:2401.xxxxx*, 2024.

[13] S. E. Field *et al.*, "Fast prediction and evaluation of gravitational waveforms using surrogate models," *Physical Review X*, vol. 4, no. 3, p. 031006, 2014.

[14] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling, "Improved variational inference with inverse autoregressive flow," *Advances in Neural Information Processing Systems*, vol. 29, pp. 4743–4751, 2016.

[15] G. Papamakarios, T. Pavlakou, and I. Murray, "Masked autoregressive flow for density estimation," *Advances in Neural Information Processing Systems*, vol. 30, pp. 2338–2347, 2017.

[16] C.-W. Huang, D. Krueger, A. Lacoste, and A. Courville, "Neural autoregressive flows," *International Conference on Machine Learning*, pp. 2078–2087, 2018.

[17] I. J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," *arXiv preprint arXiv:1412.6572*, 2014.

[18] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," *International Conference on Learning Representations*, 2018.

[19] E. Wong, L. Rice, and J. Z. Kolter, "Fast is better than free: Revisiting adversarial training," *International Conference on Learning Representations*, 2020.

[20] M. Andriushchenko and N. Flammarion, "Understanding and improving fast adversarial training," *Advances in Neural Information Processing Systems*, vol. 33, pp. 16 048–16 059, 2020.

[21] R. Dang-Nhu, I. Rish, A. Gliozzo, A. Lozano, and J. Niehues, "Adversarial training of autoregressive models," *International Conference on Machine Learning*, pp. 2404–2413, 2020.

[22] A. Ilyas, S. Santurkar, D. Tsipras, L. Engstrom, B. Tran, and A. Madry, "Adversarial examples are not bugs, they are features," *Advances in Neural Information Processing Systems*, vol. 32, pp. 125–136, 2019.

[23] L. S. Collaboration, "Ligo algorithm library - lalsuite," https://git.ligo.org/lscsoft/lalsuite, 2018, free software (GPL).

[24] F. Croce and M. Hein, "Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks," *International Conference on Machine Learning*, pp. 2206–2216, 2020.

[25] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, "Neural spline flows," *Advances in Neural Information Processing Systems*, vol. 32, pp. 7511–7522, 2019.