FUSE: ADAPTIVE FREQUENCY-TEMPORAL UNIFIED SEMANTIC ATTACKS FOR TIME SERIES FORECASTING

Anonymous authors
Paper under double-blind review

ABSTRACT

Time series forecasting (TSF) models are widely deployed in critical applications, yet their adversarial robustness remains insufficiently explored. Existing attacks mainly rely on time-domain perturbations or untargeted attacks. They overlook the structured vulnerability of forecasting models across spectral and temporal dimensions, and focus only on numerical prediction errors, without analyzing model sensitivity across different representation spaces. In this paper, we propose a **F**requency-temporal **U**nified **SE**mantic attack framework for time series forecasting (**FUSE**). We first reveal that TSF models do not exhibit uniform sensitivity across the frequency spectrum, but instead exhibit varying vulnerability across multiple frequency bands. Based on this observation, we design an adaptive full-spectrum perturbation strategy to identify and exploit critical frequency components. However, frequency-domain perturbations alone are insufficient, as they induce global and stationary disturbances, while TSF models rely on localized and non-uniform temporal dependencies. To address this mismatch, we introduce a frequency-temporal alignment mechanism that couples spectral sensitivity with temporal importance. Finally, we extend adversarial targets from numerical deviation to semantic-level manipulation of trends and periodic patterns. Extensive experiments on four datasets demonstrate that FUSE consistently outperforms existing methods in attack effectiveness, transferability, and semantic controllability across diverse forecasting models and datasets. Our code is available at https://anonymous.4open.science/r/FUSE-3353/.

1 INTRODUCTION

Time series forecasting (TSF) has become a fundamental tool in a wide range of real-world applications, including traffic prediction Li et al. (2017); Yu et al. (2017), energy management Salinas et al. (2020), and financial analysis Fischer & Krauss (2018). As these systems are increasingly deployed in safety-critical environments, understanding their adversarial robustness has emerged as an important research problem Anand & Pappas (2026); Zhang et al. (2025). Recent studies have shown that TSF models are vulnerable to carefully crafted perturbations, which can significantly degrade prediction performance Karim et al. (2020); Xu et al. (2021); Mode & Hoque (2020). However, the underlying vulnerability structure of TSF models remains insufficiently explored, particularly regarding how model sensitivity is distributed across different representation spaces Shen & Li (2025); Govindarajulu et al. (2023); Wu et al. (2022).

Adversarial attacks have been extensively studied in domains such as computer vision Dong et al. (2018); Lin et al. (2019); Guo et al. (2019); Wei et al. (2024), demonstrating that imperceptible perturbations can substantially alter model predictions. Adversarial attacks on TSF remain less explored than those on classification. Existing forecasting attacks mainly optimize temporal perturbations; even frequency-domain approaches typically do not explicitly model heterogeneous frequency sensitivity Feng et al. (2025). For instance, TCA Shen & Li (2025) focuses on temporal correlation to maximize the disruption of sequential dependencies. Additionally, MAPGD-TSF Govindarajulu et al. (2023) adapts multi-step gradient descent for targeted forecasting attacks.

While achieving promising performance, existing methods suffer from two fundamental limitations.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

First, they implicitly assume homogeneous perturbation effectiveness across temporal signals, overlooking the possibility that forecasting models may exhibit structured vulnerability in the spectral space Rahaman et al. (2019); Ortiz-Jimenez et al. (2020). To investigate this, we constrain adversarial perturbations to approximately equal-sized low-, mid-, and high-frequency regions under the same per-instance L_2 budget. As illustrated in Figure 1, attack effectiveness varies substantially across frequency regions, indicating frequency-dependent vulnerability. The full-spectrum attack remains strongest, suggesting that exploitable vulnerabilities span multiple frequency regions rather than a single fixed band.

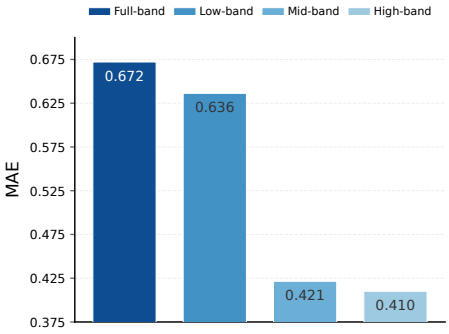

Figure 1: Spectral sensitivity analysis on DLinear under different frequency bands with the same L_2 budget.

Second, existing attacks primarily focus on untargeted attacks, aiming to degrade overall forecasting performance by maximizing prediction error Shen & Li (2025); Wu et al. (2022). While such attacks are effective in reducing accuracy, they neglect an important and realistic threat scenario: controllable manipulation of forecasting behavior. For example, in electricity forecasting, suppressing seasonal peaks may underestimate demand and affect resource allocation. In traffic systems, induced trend reversals may trigger false congestion warnings. In practice, attackers may not merely seek to increase prediction error, but instead intentionally manipulate interpretable forecasting semantics Luo et al. (2022); Mahmood & Elhamifar (2024); Zhang & Liu (2025), such as inducing false trend reversals, amplifying periodic fluctuations, or suppressing meaningful temporal patterns. Since forecasting outputs often serve as downstream decision signals, such semantically manipulated predictions may impose greater risks than simple numerical degradation, which remains largely underexplored.

In this paper, we revisit adversarial robustness in TSF from a unified frequency-temporal perspective and propose **FUSE** (**F**requency-temporal **U**nified **SE**mantic attack), a unified adversarial attack framework for time series forecasting. We first reveal a previously overlooked phenomenon that TSF models exhibit heterogeneous vulnerability across frequency bands. Based on this observation, FUSE introduces an ***Adaptive Frequency Weighting*** strategy to automatically identify and exploit vulnerable spectral components. To bridge the mismatch between global spectral perturbations and localized temporal dependencies, we further design a ***Frequency-temporal Alignment*** mechanism, which couples spectral sensitivity with temporal importance for coordinated perturbation optimization. Furthermore, beyond conventional untargeted attacks, FUSE extends adversarial objectives toward ***Semantic Target Attack***, enabling controllable attacks over trends and periodic patterns. Extensive experiments on four benchmark datasets and multiple forecasting architectures demonstrate that FUSE consistently outperforms existing methods in attack effectiveness, transferability, and semantic controllability. More importantly, our findings reveal previously overlooked structured vulnerabilities of modern forecasting systems, providing new insights into adversarial robustness in time series forecasting.

Our main contributions are summarized as follows:

- We reveal that TSF models exhibit non-uniform vulnerability across frequency bands, uncovering structured spectral sensitivity overlooked by existing adversarial attacks.
- We propose FUSE, a unified frequency-temporal adversarial attack framework that combines adaptive full-spectrum perturbation and frequency-temporal alignment to jointly exploit spectral and temporal vulnerabilities.
- We introduce a semantic targeted attack formulation for TSF, enabling controllable manipulation of trend, seasonality, and spectral periodicity beyond numerical target offsets.
- Extensive experiments across multiple datasets and forecasting models demonstrate superior attack effectiveness, transferability, and semantic controllability compared with state-of-the-art baselines.

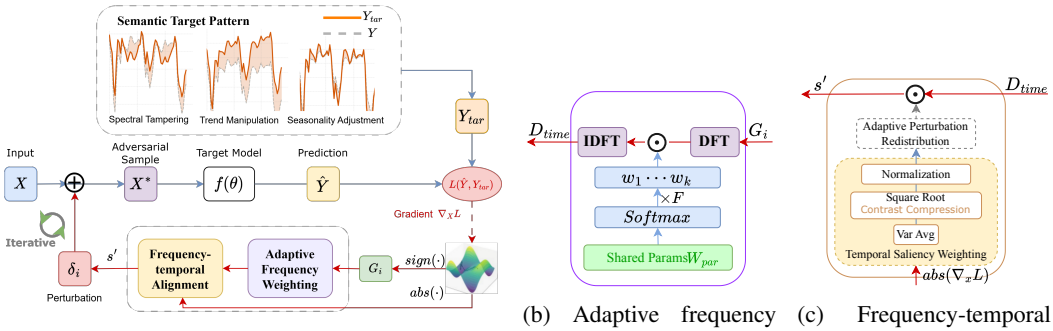

(a) Overall pipeline of FUSE.

(b) Adaptive frequency weighting module.

(c) Frequency-temporal alignment module.

Figure 2: Framework of FUSE for targeted attack. For untargeted attack, the optimization objective is given in Equ. 2.

2 METHODOLOGY

2.1 PROBLEM DEFINITION

Let $f : \mathbb{R}^{L\times M} \to \mathbb{R}^{H\times M}$ denote a pre-trained time series forecasting model. Given an input time series $\mathbf{X} \in \mathbb{R}^{L\times M}$ with a look-back window of L and M variables, *time series forecasting* aims to predict the future states $\widehat{\mathbf{Y}} \in \mathbb{R}^{H\times M}$ with a forecasting horizon of H, *i.e.*, $\widehat{\mathbf{Y}} = f(\mathbf{X})$. *Targeted attack* aims to craft an imperceptible perturbation $\boldsymbol{\delta}$ such that the model's prediction on $\mathbf{X}^* = \mathbf{X} + \boldsymbol{\delta}$ is steered toward a predefined target $\mathbf{Y}_{tar}$, subject to an ℓ_∞ budget ϵ:

$$\min_{\mathbf{X}^*} \mathcal{L}(f(\mathbf{X}^*), \mathbf{Y}_{tar}), \quad \text{s.t. } \|\mathbf{X}^* - \mathbf{X}\|_\infty \le \epsilon. \tag{1}$$

Conversely, *untargeted attack* maximizes the forecasting error relative to the ground truth $\mathbf{Y}$:

$$\max_{\mathbf{X}^*} \mathcal{L}(f(\mathbf{X}^*), \mathbf{Y}), \quad \text{s.t. } \|\mathbf{X}^* - \mathbf{X}\|_\infty \le \epsilon. \tag{2}$$

2.2 FRAMEWORK OVERVIEW

In standard adversarial attacks, the input time series X is fed into the target forecasting model, and perturbations δ are iteratively optimized via gradient updates under a perturbation budget ϵ. As illustrated in Fig. 2a, FUSE extends this conventional pipeline through a unified spectral-temporal optimization strategy consisting of three components. First, **Adaptive Frequency Weighting** identifies vulnerable spectral components and adaptively modulates perturbation updates from a global frequency perspective, enabling the attack to focus on model-sensitive frequency bands. Second, **Frequency-temporal Alignment** bridges spectral vulnerability and temporal importance by redistributing perturbation energy according to localized temporal sensitivity, thereby improving perturbation efficiency under forecasting-specific temporal dependencies. Finally, for targeted attacks, **Semantic Target Attack** guides perturbation optimization toward predefined semantic objectives, such as trend or seasonality manipulation.

Through the coordinated integration of adaptive spectral weighting and temporal alignment, FUSE generates perturbations that are both structurally precise and temporally effective, enabling more reliable untargeted degradation and controllable semantic manipulation.

2.3 PERTURBATION DIRECTION GENERATION

Given a clean input sequence X, adversarial optimization aims to iteratively construct a perturbation δ under a bounded perturbation budget ϵ, such that the perturbed input $X + \delta$ maximally disrupts forecasting behavior or steers predictions toward a predefined target. Following standard iterative attacks, FUSE initializes the optimization with $\delta_0 = \mathbf{0}$ and progressively updates perturbations through gradient-based optimization.

At the i-th iteration, given the current adversarial state $X + \delta_i$, we define the optimization objective as $\mathcal{J} = \mathcal{L}(f(X + \delta_i), Y_{\text{tar}})$ for targeted attacks and $\mathcal{J} = -\mathcal{L}(f(X + \delta_i), Y)$ for untargeted attacks. We normalize the gradient direction using the sign function:

$$\mathbf{G}_i = \text{sign}\left(\nabla_{X+\delta_i} \mathcal{J}\right). \tag{3}$$

FUSE utilizes $\mathbf{G}_i$ as the initial directional guidance and further refines it through frequency-temporal optimization. Specifically, **Adaptive Frequency Weighting** emphasizes vulnerable spectral components, while **Frequency-temporal Alignment** further calibrates perturbations according to temporal importance.

2.4 Adaptive Frequency Weighting

Conventional gradient-based attacks optimize perturbations directly in the temporal domain and lack explicit awareness of spectral vulnerability. However, forecasting models exhibit non-uniform sensitivity across frequency bands due to their spectral characteristics Rahaman et al. (2019); Ortiz-Jimenez et al. (2020). We compare BIM restricted to specific frequency partitions against full-band attacks on the DLinear model. As shown in Figure 1, full-band attacks consistently yield stronger forecasting degradation, indicating that vulnerabilities are distributed across complex spectral combinations rather than isolated bands. This motivates our adaptive mechanism to identify and emphasize sensitive frequencies.

To this end, we introduce a parameter vector $\mathbf{w}_{\text{par}} \in \mathbb{R}^F$ to characterize the spectral vulnerabilities of the target model, where F denotes the number of frequency bins. Shared across the batch, $\mathbf{w}_{\text{par}}$ extracts a global spectral sensitivity by aggregating the gradient information from multiple instances. We convert $\mathbf{w}_{\text{par}}$ into a normalized weighting distribution using Softmax:

$$w_k = F \cdot \frac{\exp(w_{\text{par},k})}{\sum_{j=1}^{F} \exp(w_{\text{par},j})}, \tag{4}$$

where w_k denotes the adaptive importance of the k-th frequency component. The scaling factor F keeps the mean spectral weight at 1.0, enabling relative redistribution across frequency components while avoiding uncontrolled global rescaling. The final perturbation magnitude is enforced by the ℓ_∞ projection during adversarial updates. The weighted spectral representation is then mapped back to the temporal domain via the Inverse Discrete Fourier Transform (IDFT).

To optimize the frequency weights, we formulate the perturbation update as a differentiable process. After generating a trial perturbation, the adversarial instance is re-evaluated by the forecasting model, allowing the attack objective to backpropagate through both the model and perturbation update to $\mathbf{w}_{\text{par}}$. This enables FUSE to progressively identify and emphasize sensitive frequency components during optimization.

2.5 Frequency-temporal Alignment

While **Adaptive Frequency Weighting** identifies vulnerable spectral components from a global perspective, frequency-domain perturbations inherently lack temporal localization due to the global nature of spectral transformations. However, forecasting models exhibit highly non-uniform temporal sensitivity, where different temporal regions contribute unequally to prediction outcomes. To better characterize this phenomenon, we conduct a sensitivity analysis by independently applying localized BIM perturbations to individual timestamps and frequency components. As shown in the sensitivity analysis in Figure 3 (a), forecasting models exhibit highly heterogeneous temporal sensitivity, with vulnerability patterns varying substantially across architectures. In contrast, Figure 3(b)–(d) shows that spectral sensitivity exhibits stronger structural consistency across models. These observations suggest that effective forecasting attacks should combine global frequency-aware perturbation with localized temporal alignment. Motivated by this, we introduce **Frequency-temporal Alignment**, which performs temporal-aware perturbation allocation in two stages: **Temporal Saliency Weighting** and **Adaptive Perturbation Redistribution**.

2.5.1 Temporal Saliency Weighting

We first estimate timestamp-level temporal saliency to identify regions where perturbations are most likely to influence forecasting behavior. Given an input instance $\mathbf{X}_i$, we compute the saliency score g_t at each timestamp t by aggregating gradient magnitudes across all M variables:

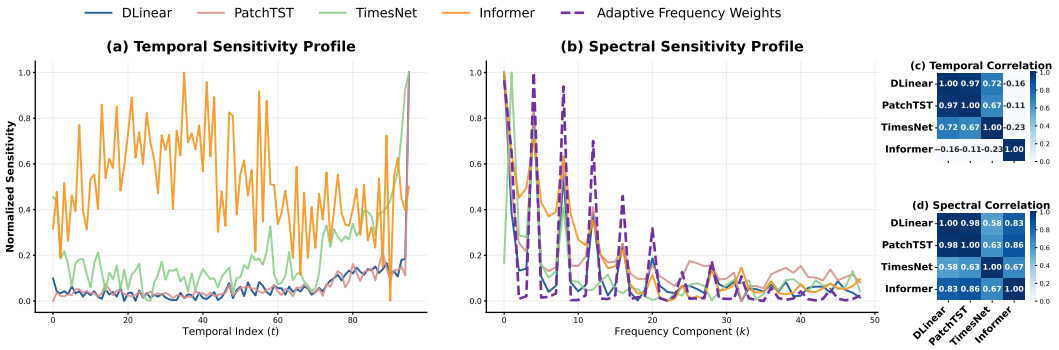

Figure 3: Cross-model temporal and spectral sensitivity. (a)–(b) Sensitivity to BIM perturbations at individual timestamps and frequency components, measured by the normalized increase in MAE. (c)–(d) Cross-model correlations of temporal and spectral sensitivity, respectively.

$$g_t = \frac{1}{M} \sum_{c=1}^{M} |\nabla_{x_{t,c}} \mathcal{L}|, \tag{5}$$

where larger gradients indicate stronger influence on the attack objective.

Directly allocating perturbation budgets according to raw gradients may excessively concentrate updates on a few dominant timestamps, reducing optimization stability. To encourage smoother perturbation allocation, we apply square-root normalization:

$$s_t = \frac{\sqrt{g_t}}{\sum_{j=1}^{T} \sqrt{g_j}}. \tag{6}$$

The resulting temporal weights emphasize sensitive temporal regions while maintaining balanced perturbation allocation, enabling globally modulated spectral perturbations to better align with localized forecasting dynamics.

2.5.2 Adaptive Perturbation Redistribution

Introducing non-uniform spectral-temporal weighting changes optimization behavior under the ℓ_∞ constraint. Coordinates receiving larger perturbation weights may quickly reach the perturbation boundary, after which subsequent updates become ineffective. Meanwhile, low-weight coordinates may remain under-optimized, resulting in inefficient utilization of the perturbation budget. To maintain effective perturbation allocation throughout iterative optimization, we further introduce an **Adaptive Perturbation Redistribution** mechanism.

Specifically, we define an active mask $\mathcal{M} \in \{0,1\}^{T \times M}$, where a coordinate (t, c) remains active if $|\delta_{t,c}| < \epsilon - \gamma$ ($\gamma = 10^{-7}$). Based on the number of active coordinates, we compute a redistribution multiplier:

$$\lambda_{\text{red},c} = \frac{T}{\sum_{j=1}^{T} \mathcal{M}_{j,c} + \eta}, \tag{7}$$

which reallocates update magnitude from saturated coordinates to active ones.

To preserve temporal importance during redistribution, we further perform saliency re-normalization over the active set:

$$s'_{t,c} = \left(\frac{\sqrt{g_t}}{\frac{1}{N_{\text{act},c}} \sum_{j \in \mathcal{T}_{\text{act},c}} \sqrt{g_j}} \right) \lambda_{\text{red},c} M_{t,c}, \tag{8}$$

thereby maintaining effective perturbation allocation and preventing optimization stagnation during iterative refinement.

2.6 Perturbation Synthesis and Update

At each iteration i, the directional signal $\mathbf{G}_i$ is transformed into the frequency domain for spectral filtering, yielding the modulated gradient direction:

$$\mathbf{D}_{time} = \text{IDFT}\left(\text{DFT}(\mathbf{G}_i) \odot \mathbf{w}\right), \tag{9}$$

where $\mathbf{w}$ denotes the adaptive frequency weights. The perturbation is then updated via temporal weighting and ℓ_∞ projection:

$$\delta_{i+1} = \text{clip}_\epsilon\{\delta_i - \alpha \cdot \mathbf{D}_{time} \odot \mathbf{s}'\}, \tag{10}$$

yielding the adversarial instance $X^* = X + \delta_I$ after I iterations.

The complete optimization procedure is detailed in Appendix B.

2.7 Semantic Target Attack

Existing targeted attacks on TSF models typically rely on simple numerical offsets, such as uniformly shifting forecasting values by a constant magnitude. Although effective for inducing prediction deviation, such perturbations lack semantic controllability and often fail to reflect realistic manipulation objectives. To pose a more realistic threat, FUSE induces predictions that adhere to predefined semantic patterns with specific physical characteristics, including trend, seasonal, and spectral manipulations.

Given the ground-truth future sequence $\mathbf{Y} \in \mathbb{R}^{T \times M}$, we decompose its structural components into trend and seasonal terms. Based on this decomposition, we define three semantic target patterns:

(1) Trend Manipulation. This objective modifies the long-term forecasting trajectory by weakening or reversing the underlying trend.

(2) Seasonality Adjustment. This objective controls the intensity of periodic fluctuations while preserving the underlying trend.

(3) Spectral Tampering. This objective directly manipulates selected frequency components to suppress or amplify specific periodic structures.

The detailed construction of the three semantic targets is provided in Appendix A.3. Targeted optimization is then performed by minimizing $\mathcal{L}(f(\mathbf{X} + \boldsymbol{\delta}), \mathbf{Y}_{tar})$.

3 Experiment

3.1 Experimental Setup

Datasets. To assess the efficacy of FUSE across diverse scenarios, we conduct experiments on four datasets: ETT datasets (**ETTh1**, **ETTm1**)[1], **Electricity**[2], and **Weather**[3], covering energy, electricity, and meteorology. We use a **6:2:2** train/validation/test split and perform attacks on the test set.

Target Models. We adopt four representative models as target models of adversarial attack, *i.e.,* the MLP-based **DLinear**, the CNN-based **TimesNet**, and two Transformer methods, **PatchTST** (patch-based) and **Informer** (attention-based).

Baselines. We compare FUSE with a comprehensive set of adversarial methods, including: *(1) adversarial attack methods from computer vision*: **FGSM** Goodfellow et al. (2014), **BIM** Kurakin et al. (2018), **MI-FGSM** Dong et al. (2018), and **NI-FGSM** Lin et al. (2019); and *(2) adversarial attack methods for time-series*: **Fre-C&W** Feng et al. (2025), **TCA** Shen & Li (2025), **MAPGD-TSF** Govindarajulu et al. (2023), and **AAIM** Wu et al. (2022). All baselines are implemented according to their original papers to ensure a fair comparison.

Implementation Details. All experiments are implemented with BasicTS[4]. Following prior attacks Kurakin et al. (2018); Dong et al. (2018); Lin et al. (2019); Shen & Li (2025), we use an L_∞ perturbation budget of $\epsilon = 0.2$. Iterative attacks run for $K = 50$ steps with $\alpha = \epsilon/K$, and both the input length and prediction horizon are 96. For targeted semantic attacks, we evaluate three manipulation modes: spectral reduction (P1), trend reduction (P2), and seasonality enhancement (P3), covering spectral, trend, and seasonal structures, respectively.

[1] https://github.com/zhouhaoyi/ETDataset

[2] https://archive.ics.uci.edu/dataset/321/electricityloaddiagrams20112014

[3] https://www.bgc-jena.mpg.de/wetter/

[4] https://github.com/GestaltCogTeam/BasicTS

Table 1: White-box targeted semantic induction on **ETTh1**. **P1/P2/P3**: spectral reduction, trend reduction, and seasonality enhancement. Best/second-best results are **bold**/underlined.

Pat.	Method	MAE (with Adherence %) ↓				MSE (with Adherence %) ↓				ASR % ↑			
		DLinear	PatchTST	TimesNet	Informer	DLinear	PatchTST	TimesNet	Informer	DLinear	PatchTST	TimesNet	Informer
P1	Init.	0.4652(-0.0%)	0.4693(-0.0%)	0.4685(-0.0%)	0.6706(-0.0%)	0.4820(-0.0%)	0.4988(-0.0%)	0.4834(-0.0%)	0.7947(-0.0%)	0.0%	0.0%	0.0%	0.0%
	FGSM	0.3771(18.94%)	0.3943(15.98%)	0.3909(16.56%)	0.5684(15.24%)	0.3698(23.28%)	0.3986(20.09%)	0.3505(27.49%)	0.5889(25.90%)	51.32%	40.80%	59.58%	**100.00%**
	BIM	0.3625(22.08%)	0.3840(18.18%)	0.3622(22.69%)	0.5593(16.60%)	0.3657(24.13%)	0.3950(20.81%)	0.3186(34.09%)	0.5756(27.57%)	62.55%	48.23%	79.96%	**100.00%**
	MI-FGSM	0.3660(21.32%)	0.3870(17.54%)	0.3712(20.77%)	0.5610(16.34%)	0.3660(24.07%)	0.3957(20.67%)	0.3278(32.19%)	0.5775(27.33%)	59.02%	45.56%	75.01%	**100.00%**
	NI-FGSM	0.3636(21.84%)	0.3842(18.13%)	0.3648(22.13%)	0.5593(16.60%)	**0.3651(24.25%)**	**0.3940(21.01%)**	0.3199(33.82%)	0.5761(27.51%)	61.51%	48.05%	77.61%	**100.00%**
	Fre-C&W	0.4239(8.88%)	0.4275(8.91%)	0.4373(6.66%)	0.6110(8.89%)	0.4126(14.40%)	0.4288(14.03%)	0.3970(17.87%)	0.6445(18.90%)	25.18%	24.14%	24.43%	93.38%
	TCA	0.3710(20.26%)	0.3947(15.90%)	0.3668(21.71%)	0.5672(15.41%)	0.3738(22.44%)	0.4047(18.87%)	0.3231(33.16%)	0.5908(25.65%)	58.31%	43.18%	78.21%	98.51%
	MAPGD-TSF	0.3628(22.02%)	0.3843(18.11%)	0.3637(22.36%)	0.5598(16.53%)	0.3657(24.13%)	0.3951(20.78%)	0.3204(33.72%)	0.5763(27.48%)	62.37%	48.08%	79.17%	**100.00%**
	AAIM	0.4510(3.05%)	0.4546(3.13%)	0.4584(2.16%)	0.6646(0.89%)	0.4638(3.78%)	0.4814(3.49%)	0.4685(3.08%)	0.7819(1.61%)	20.08%	19.45%	17.74%	64.60%
	FUSE	**0.3613(22.33%)**	**0.3825(18.50%)**	**0.3563(23.95%)**	**0.5560(17.09%)**	**0.3651(24.25%)**	0.3944(20.93%)	**0.3110(35.66%)**	**0.5695(28.34%)**	**63.63%**	**49.24%**	**81.15%**	**100.00%**
P2	Init.	0.4453(-0.0%)	0.4454(-0.0%)	0.4616(-0.0%)	0.7253(-0.0%)	0.4754(-0.0%)	0.4778(-0.0%)	0.4965(-0.0%)	0.9157(-0.0%)	0.0%	0.0%	0.0%	0.0%
	FGSM	0.3758(15.61%)	0.3934(11.67%)	0.4073(11.76%)	0.6205(14.45%)	0.3819(19.67%)	0.4102(14.19%)	0.3963(20.18%)	0.7181(21.58%)	68.58%	59.20%	55.56%	74.19%
	BIM	0.3590(19.38%)	0.3811(14.44%)	0.3728(19.24%)	0.6104(15.84%)	0.3769(20.72%)	0.4051(15.22%)	0.3563(28.24%)	0.7017(23.37%)	75.46%	66.68%	75.86%	77.09%
	MI-FGSM	0.3626(18.57%)	0.3841(13.76%)	0.3826(17.11%)	0.6122(15.60%)	0.3769(20.72%)	0.4060(15.03%)	0.3670(26.08%)	0.7048(23.03%)	74.34%	65.08%	71.85%	76.76%
	NI-FGSM	0.3599(19.18%)	0.3809(14.48%)	0.3742(18.93%)	0.6103(15.85%)	0.3762(20.87%)	**0.4025(15.76%)**	0.3567(28.16%)	0.6993(23.63%)	75.05%	66.90%	75.01%	77.13%
	Fre-C&W	0.4629(-3.95%)	0.4550(-2.16%)	0.4889(-5.91%)	0.6805(6.18%)	0.4763(-0.19%)	0.4739(0.82%)	0.5309(-6.93%)	0.8352(8.79%)	14.50%	18.11%	17.07%	53.96%
	TCA	0.3619(18.73%)	0.3862(13.30%)	0.3750(18.77%)	0.6165(15.00%)	0.3802(20.02%)	0.4098(14.24%)	0.3586(27.77%)	0.7122(22.22%)	74.30%	64.45%	74.93%	75.31%
	MAPGD-TSF	0.3593(19.32%)	0.3814(14.38%)	0.3744(18.88%)	0.6111(15.75%)	0.3769(20.72%)	0.4053(15.18%)	0.3584(27.82%)	0.7028(23.25%)	75.38%	66.49%	75.05%	76.98%
	AAIM	0.4334(2.67%)	0.4343(2.49%)	0.4494(2.64%)	0.7185(0.94%)	0.4614(2.94%)	0.4656(2.55%)	0.4767(3.99%)	0.9027(1.42%)	30.83%	30.94%	30.57%	39.94%
	FUSE	**0.3575(19.72%)**	**0.3792(14.86%)**	**0.3668(20.54%)**	**0.6070(16.31%)**	**0.3760(20.91%)**	0.4036(15.53%)	**0.3482(29.87%)**	**0.6952(24.08%)**	**76.01%**	**67.65%**	**78.24%**	**77.95%**
P3	Init.	0.4606(-0.0%)	0.4653(-0.0%)	0.4780(-0.0%)	0.7545(-0.0%)	0.5235(-0.0%)	0.5261(-0.0%)	0.5471(-0.0%)	0.9987(-0.0%)	0.0%	0.0%	0.0%	0.0%
	FGSM	0.3922(14.85%)	0.4084(12.23%)	0.4340(9.21%)	0.6553(13.15%)	0.4363(16.66%)	0.4521(14.07%)	0.4720(13.73%)	0.8165(18.24%)	46.00%	27.04%	22.72%	65.94%
	BIM	0.3778(17.98%)	0.3987(14.31%)	0.4008(16.15%)	0.6459(14.39%)	0.4312(17.63%)	0.4479(14.86%)	0.4301(21.39%)	0.7992(19.97%)	63.96%	40.35%	22.40%	72.11%
	MI-FGSM	0.3809(17.30%)	0.4012(13.78%)	0.4095(14.33%)	0.6474(14.19%)	0.4315(17.57%)	0.4486(14.79%)	0.4401(19.56%)	0.8029(19.61%)	60.88%	36.63%	42.88%	71.22%
	NI-FGSM	0.3786(17.80%)	0.3989(14.27%)	0.4014(16.02%)	0.6458(14.41%)	**0.4301(17.84%)**	**0.4463(15.17%)**	0.4287(21.64%)	0.7956(20.34%)	63.22%	38.86%	50.32%	72.89%
	Fre-C&W	0.4842(-5.12%)	0.4836(-3.93%)	0.5238(-9.58%)	0.7112(5.74%)	0.5396(-3.07%)	0.5394(-2.53%)	0.6332(-15.74%)	0.9350(6.38%)	0.22%	0.60%	2.86%	34.10%
	TCA	0.3817(17.14%)	0.4051(12.95%)	0.4046(15.35%)	0.6532(13.43%)	0.4360(16.72%)	0.4551(13.50%)	0.4347(20.55%)	0.8124(18.66%)	61.96%	34.59%	47.94%	66.27%
	MAPGD-TSF	0.3780(17.93%)	0.3990(14.26%)	0.4023(15.85%)	0.6465(14.32%)	0.4313(17.62%)	0.4480(14.84%)	0.4319(21.06%)	0.8004(19.86%)	63.89%	38.75%	48.79%	71.66%
	AAIM	0.4509(2.11%)	0.4561(1.98%)	0.4685(1.99%)	0.7484(0.81%)	0.5128(2.04%)	0.5163(1.86%)	0.5330(2.58%)	0.9866(1.21%)	0.0%	0.0%	0.04%	0.45%
	FUSE	**0.3764(18.28%)**	**0.3972(14.64%)**	**0.3958(17.20%)**	**0.6424(14.86%)**	0.4302(17.82%)	0.4469(15.05%)	**0.4233(22.63%)**	**0.7918(20.72%)**	**65.08%**	**40.94%**	**54.85%**	**74.56%**

Evaluation Metrics. For untargeted attacks, we report MAE, MSE, WAPE, and Degradation Percentage (P). For targeted attacks, we report MAE, MSE, and two **novel metrics**: Attack Success Rate (ASR) and Target Adherence (η). Detailed definitions are provided in Appendix A.4. Higher P, ASR, and η indicate stronger attack performance.

3.2 White-box Attacks

We first evaluate FUSE under the white-box setting, where attackers have full access to model parameters and gradients. We consider two settings: untargeted attacks (Table 4), which evaluate forecasting degradation under a constrained perturbation budget, and targeted attacks (Table 1), which evaluate controllable steering toward predefined semantic objectives.

Targeted Attacks. Table 1 reports white-box targeted attacks on ETTh1 under three semantic objectives: spectral reduction (P1), trend reduction (P2), and seasonality enhancement (P3).

(1) FUSE achieves the best or competitive ASR and Target Adherence across all three semantic objectives. For example, on TimesNet, its ASR reaches 81.15%, 78.24%, and 54.85% for P1–P3, respectively, consistently exceeding the strongest baselines. This suggests that frequency–temporal coordination improves directional control toward the intended semantic targets.***(2)*** Performance varies across semantic objectives, with P3 generally being more challenging than P1 and P2. On TimesNet, the ASR of FUSE decreases from 81.15%/78.24% on P1/P2 to 54.85% on P3. This difference can be attributed to the fact that seasonality enhancement requires actively constructing and amplifying periodic structures, whereas reduction-oriented objectives mainly suppress existing components under the same perturbation budget.

Overall, FUSE consistently achieves strong targeted attack performance across all three semantic objectives and forecasting architectures. These results demonstrate that frequency–temporal coordination enables controllable manipulation of forecasting behaviors.

3.3 Transferability Analysis (Black-box Attacks)

To evaluate cross-model transferability, we generate adversarial examples on each surrogate model and evaluate them on the remaining unseen architectures.

Quantitative Results (Table 2). We make two observations. ***(1)*** FUSE shows consistently strong average transferability, ranking first in 11 of the 12 dataset–metric averages. For example, on Weather, FUSE achieves average degradation of 40.78%, 33.47%, and 103.35% in MAE, MSE, and WAPE, respectively. This indicates that the learned frequency–temporal perturbations retain substantial effectiveness across different forecasting architectures. ***(2)*** The advantage is not uniform across individual surrogate models. For instance, on ETTm1, FGSM slightly exceeds FUSE in av-

Table 2: Averaged black-box untargeted transfer attack results. Each column represents the surrogate model used to generate attacks, and the results are the average degradation percentage (%) across the other target models. Best results are in **bold** and second-best are underlined.

Dataset	Method	Degradation Percentage on MAE ↑					Degradation Percentage on MSE ↑					Degradation Percentage on WAPE ↑				
		DLinear	PatchTST	TimesNet	Informer	Average	DLinear	PTST	TimesNet	Informer	Average	DLinear	PatchTST	TimesNet	Informer	Average
ETTh1	FGSM	16.71%	12.94%	**6.47%**	7.80%	10.98%	20.03%	15.60%	**8.47%**	**10.06%**	13.54%	18.28%	14.22%	**6.49%**	**8.45%**	11.86%
	BIM	17.21%	13.22%	5.73%	7.35%	10.88%	20.41%	15.84%	7.65%	9.50%	13.35%	18.81%	14.53%	5.83%	7.95%	11.78%
	MI-FGSM	17.07%	13.16%	6.21%	7.75%	11.05%	20.31%	15.80%	8.17%	10.02%	13.58%	18.67%	14.47%	6.24%	8.42%	11.95%
	NI-FGSM	17.68%	12.66%	5.80%	7.38%	10.88%	20.81%	15.20%	7.82%	9.43%	13.32%	19.35%	13.97%	5.88%	7.89%	11.77%
	Fre-C&W	6.39%	4.96%	3.97%	4.31%	4.91%	7.76%	6.38%	6.17%	5.65%	6.49%	7.40%	6.52%	4.71%	4.57%	5.80%
	TCA	13.36%	11.70%	5.79%	7.50%	9.58%	15.89%	14.07%	7.73%	9.56%	11.81%	14.99%	13.21%	5.90%	8.39%	10.62%
	MAPGD-TSF	17.14%	13.18%	5.82%	7.37%	10.88%	20.33%	15.79%	7.77%	9.55%	13.36%	18.74%	14.48%	5.90%	7.97%	11.77%
	AAIM	1.84%	1.81%	1.06%	0.23%	1.24%	1.84%	1.84%	1.18%	0.30%	1.29%	2.01%	2.08%	1.07%	0.23%	1.35%
	FUSE	**18.11%**	**13.66%**	6.03%	**7.81%**	**11.40%**	**21.28%**	**16.27%**	7.99%	9.86%	**13.85%**	**19.83%**	**15.03%**	6.25%	8.36%	**12.37%**
ETTm1	FGSM	27.29%	**13.46%**	**10.73%**	17.89%	17.34%	33.31%	**15.21%**	**14.34%**	21.46%	**21.08%**	29.18%	**14.80%**	**11.26%**	19.21%	18.61%
	BIM	27.67%	12.46%	9.61%	17.87%	16.90%	33.55%	13.95%	13.29%	21.21%	20.50%	29.60%	13.61%	10.15%	19.19%	18.14%
	MI-FGSM	27.58%	13.20%	10.41%	18.02%	17.30%	33.50%	14.74%	14.09%	21.50%	20.96%	29.50%	14.47%	10.90%	19.37%	18.56%
	NI-FGSM	27.87%	11.93%	9.80%	17.84%	16.86%	33.71%	13.53%	13.43%	21.15%	20.46%	29.82%	13.00%	10.29%	19.15%	18.07%
	Fre-C&W	12.25%	5.81%	5.70%	11.21%	8.74%	14.52%	5.81%	7.39%	11.57%	9.82%	14.50%	7.26%	6.85%	12.89%	10.38%
	TCA	17.54%	13.38%	11.35%	16.76%	14.76%	19.51%	15.10%	14.87%	19.89%	17.34%	19.78%	14.85%	12.21%	18.37%	16.30%
	MAPGD-TSF	26.98%	13.19%	11.43%	16.52%	17.03%	32.40%	14.75%	15.04%	20.45%	20.66%	28.72%	14.53%	12.26%	17.29%	18.20%
	AAIM	4.19%	2.66%	2.07%	2.44%	2.84%	4.31%	2.76%	2.27%	2.64%	3.00%	4.46%	2.97%	2.20%	2.60%	3.06%
	FUSE	**28.16%**	12.49%	10.13%	**18.65%**	**17.36%**	**33.94%**	14.01%	13.89%	**21.87%**	20.93%	**30.15%**	13.60%	10.83%	**20.17%**	**18.69%**
ECL	FGSM	24.63%	16.07%	**11.66%**	**21.42%**	18.44%	29.26%	16.98%	**13.68%**	**24.70%**	**21.15%**	25.79%	16.95%	**12.38%**	**22.65%**	19.44%
	BIM	25.33%	16.63%	8.34%	18.45%	17.19%	29.83%	17.37%	9.49%	20.82%	19.38%	26.56%	17.54%	8.96%	19.40%	18.11%
	MI-FGSM	25.12%	16.44%	10.26%	20.45%	18.07%	29.67%	17.26%	11.94%	23.49%	20.59%	26.33%	17.34%	11.02%	21.59%	19.07%
	NI-FGSM	25.80%	16.35%	8.90%	18.19%	17.31%	30.39%	17.09%	9.66%	20.35%	19.37%	27.07%	17.21%	9.71%	19.07%	18.27%
	Fre-C&W	13.57%	8.46%	8.62%	15.86%	11.62%	15.56%	9.07%	8.35%	15.67%	12.16%	14.88%	9.48%	9.31%	16.92%	12.65%
	TCA	19.12%	16.13%	8.33%	18.46%	15.51%	20.63%	16.79%	9.48%	20.83%	16.93%	20.50%	17.31%	8.95%	19.40%	16.54%
	MAPGD-TSF	25.28%	16.55%	8.23%	18.51%	17.14%	29.82%	17.30%	9.55%	21.02%	19.42%	26.49%	17.45%	8.80%	19.47%	18.05%
	AAIM	2.15%	2.05%	0.54%	0.70%	1.36%	2.46%	2.17%	0.51%	0.87%	1.50%	2.29%	2.14%	0.59%	0.72%	1.43%
	FUSE	**26.73%**	**17.71%**	10.34%	20.33%	**18.77%**	**31.23%**	**18.38%**	11.10%	22.21%	20.73%	**28.06%**	**18.65%**	11.14%	21.36%	**19.80%**
Weather	FGSM	54.57%	29.67%	30.73%	39.35%	38.58%	50.96%	24.12%	21.70%	30.05%	31.71%	113.78%	65.96%	82.48%	104.38%	91.65%
	BIM	55.57%	28.79%	28.58%	37.91%	37.71%	51.74%	22.62%	21.26%	28.42%	31.01%	119.97%	70.31%	78.02%	100.90%	92.30%
	MI-FGSM	55.27%	30.82%	30.87%	39.25%	39.05%	51.54%	24.13%	23.13%	29.79%	32.15%	118.13%	75.75%	80.21%	103.12%	94.30%
	NI-FGSM	56.22%	**33.99%**	30.44%	37.09%	39.44%	**52.66%**	**28.61%**	22.15%	27.34%	32.69%	117.84%	73.65%	89.64%	99.62%	95.19%
	Fre-C&W	42.20%	23.91%	20.30%	30.00%	29.10%	36.11%	18.91%	15.22%	21.32%	22.89%	108.96%	57.45%	56.51%	55.76%	69.67%
	TCA	39.02%	28.82%	29.31%	41.20%	34.59%	27.61%	21.38%	20.23%	27.39%	24.15%	119.98%	70.33%	78.02%	100.91%	92.31%
	MAPGD-TSF	55.46%	28.93%	28.38%	37.74%	37.63%	51.67%	22.88%	21.12%	28.35%	31.00%	119.24%	70.36%	73.86%	99.74%	90.80%
	AAIM	24.55%	11.97%	12.16%	14.57%	15.81%	21.18%	8.70%	8.83%	9.96%	12.17%	36.01%	26.92%	24.37%	28.39%	28.92%
	FUSE	**56.23%**	31.94%	**32.79%**	**42.18%**	**40.78%**	52.35%	25.92%	**24.05%**	**31.56%**	**33.47%**	**120.30%**	**76.68%**	**95.67%**	**120.75%**	**103.35%**

erage MSE degradation (21.08% vs. 20.93%). Such variation reflects the architecture dependence of adversarial transfer, while the stronger overall averages of FUSE suggest that exploiting spectral structures provides more stable cross-model transfer than relying solely on point-wise gradients.

Mechanism of Transferability: Spectral Convergence (Figure 3). The robust transferability of FUSE is rooted in the underlying vulnerability structure of TSF models. As illustrated in Figure 3, while the sensitivity patterns in the temporal domain are often discrete and architecture-specific (Figure 3 (a)), their distributions in the frequency domain exhibit **significant commonality**, particularly within low-frequency bands (Figure 3 (b)). Quantitative analysis in Figures 3 (c) and 3 (d) corroborates this; spectral sensitivity demonstrates a substantially higher Pearson correlation compared to the temporal domain.

Cross-model Spectral Commonality. Figure 3 provides further insight into transferability. Temporal sensitivity varies considerably across architectures, whereas spectral sensitivity exhibits more consistent patterns, especially at low frequencies. The higher cross-model Pearson correlations in the spectral domain further support this observation. Such commonality allows FUSE to emphasize frequency components that remain influential across architectures, improving cross-model transfer.

Weight Visualization. The learned w_{par} assigns larger weights to frequency with stronger model sensitivity (Figure 3(b)). It shows that the attack-adaptive weighting mechanism concentrates optimization on vulnerable spectral components rather than treating all frequencies uniformly.

3.4 Ablation Study

To investigate the individual contribution of each core component in FUSE, we evaluate three ablated variants: (1) **w/o AFW**, which removes adaptive frequency weighting; (2) **w/o Redist**, which disables perturbation redistribution; and (3) **w/o Saliency**, which replaces saliency-guided allocation with uniform redistribution. We additionally include **BIM** as a reference iterative attack. The results in Table 3 are averaged over P1–P3, ETTh1/ETTm1, and DLinear/TimesNet under both white-box and black-box transfer settings. The complete FUSE consistently achieves the highest ASR, reaching **76.46%** and **39.06%** in the white-box and black-box settings, respectively.

Specifically, removing AFW consistently reduces both Target Adherence and ASR, indicating that adaptive spectral weighting improves frequency-domain adaptation and cross-model transferability. The **w/o Redist** variant shows the largest overall degradation, with ASR decreasing to **74.56%** and **37.98%**, suggesting that perturbation redistribution is important for effective budget utilization. Meanwhile, **w/o Saliency** also degrades performance, with ASR dropping to **75.06%** and **38.32%**, showing that saliency-guided allocation further improves where the redistributed perturbation is applied. Overall, AFW, redistribution, and temporal saliency provide complementary benefits in spectral adaptation, perturbation utilization, and temporal allocation.

Table 3: Ablation study of FUSE averaged over **P1–P3**, **ETTh1/ETTm1**, and **DLinear/TimesNet** under both **White-box** and **Black-box** (Transfer) settings. η denotes the Target Adherence (%). Best results are highlighted in **bold**.

Variant	White-box			Black-box (Avg.)		
	MAE (η %) ↓	**MSE (η %) ↓**	**ASR % ↑**	**MAE (η %) ↓**	**MSE (η %) ↓**	**ASR % ↑**
FUSE (Full)	**0.3266(20.61%)**	**0.3082(24.98%)**	**76.46%**	**0.4466(9.04%)**	**0.4621(10.91%)**	**39.06%**
w/o AFW	0.3279(20.29%)	0.3098(24.50%)	75.75%	0.4471(8.94%)	0.4628(10.78%)	38.66%
w/o Redist	0.3299(19.79%)	0.3122(23.88%)	74.56%	0.4479(8.76%)	0.4638(10.56%)	37.98%
w/o Saliency	0.3290(20.01%)	0.3111(24.16%)	75.06%	0.4475(8.85%)	0.4633(10.67%)	38.32%
BIM	0.3307(19.56%)	0.3129(23.62%)	74.83%	0.4477(8.80%)	0.4635(10.62%)	38.16%

4 Related Works

4.1 Adversarial Attacks for Time Series Forecasting

Recent studies have begun to explore adversarial robustness in time series, with a primary focus on classification tasks (TSC) Belkhouja & Doppa (2022); Kim et al. (2025); Wang et al. (2022); Ding et al. (2023); Belkhouja et al. (2022). In contrast, adversarial attacks for time series forecasting (TSF) remain relatively underexplored. Existing methods have introduced task-specific constraints from either temporal or spectral perspectives. For example, Fre-C&W Feng et al. (2025) incorporates frequency-domain regularization, while TCA Shen & Li (2025) imposes temporal consistency constraints. AAIM Wu et al. (2022) utilizes importance measuring to select sparse temporal points for perturbation, and MAPGD-TSF Govindarajulu et al. (2023) employs adaptive step sizes for targeted attacks. Despite these advances, existing TSF-specific attacks rely on fixed frequency constraints Feng et al. (2025), local temporal sparsity Wu et al. (2022), or point-wise updates Govindarajulu et al. (2023), failing to adaptively adjust based on dynamic feedback or exploit the synergy between global spectral vulnerability and local temporal importance. Our FUSE addresses this via coordinated spectral-temporal optimization, automatically learning which frequencies to target and when to intervene through Frequency-temporal Alignment.

4.2 Targeted and Semantic Adversarial Attacks

Targeted attacks aim to steer model outputs toward a specific predefined target trajectory. While CV research has evolved toward semantic-level manipulationLuo et al. (2022); Mahmood & Elhamifar (2024); Zhang & Liu (2025), targeted attacks in TSF are still largely restricted to simple numerical offsets (*e.g.,* $Y + \epsilon$)Govindarajulu et al. (2023). Numerical targeted attacks lack structural plausibility, failing to induce interpretable errors. In practical forecasting scenarios, adversarial behaviors may instead involve altering meaningful physical patterns, such as trend shifts or abnormal periodic fluctuations. FUSE instead performs Semantic Target Attack via three patterns—*Trend Manipulation*, *Seasonality Adjustment*, and *Spectral Tampering*, inducing structurally plausible but erroneous patterns.

5 Conclusion

In this paper, we investigate adversarial robustness in time series forecasting from a unified frequency-temporal perspective and propose FUSE, a novel semantic attack framework for forecasting models. We first reveal that TSF models exhibit non-uniform vulnerability across the frequency spectrum, where different frequency components contribute unequally to adversarial sensitivity. Motivated by this observation, FUSE introduces Adaptive Frequency Weighting to identify vulnerable spectral components and Frequency-temporal Alignment to further coordinate perturbation allocation according to temporal importance. Beyond conventional untargeted attacks, we also design Semantic Target Attack, enabling controllable manipulation of forecasting behaviors such as trend and seasonality. Extensive experiments across multiple datasets and forecasting architectures demonstrate that FUSE consistently outperforms existing baselines in both untargeted degradation and targeted semantic steering.

6 AI USE STATEMENT

During the preparation of this work, generative AI tools (including ChatGPT and Claude) were used for limited editorial assistance, language refinement, and code debugging. No generative AI tools were used to generate scientific claims, experimental results, or conclusions. All analyses, implementations, interpretations, and final manuscript content were independently verified and approved by the authors, who take full responsibility for the work.

REFERENCES

Sribalaji C Anand and George J Pappas. Adversarial robustness of deep state space models for forecasting. *arXiv preprint arXiv:2604.03427*, 2026.

Taha Belkhouja and Janardhan Rao Doppa. Adversarial framework with certified robustness for time-series domain via statistical features. *Journal of Artificial Intelligence Research*, 73:1435–1471, 2022.

Taha Belkhouja, Yan Yan, and Janardhan Rao Doppa. Dynamic time warping based adversarial framework for time-series domain. *IEEE transactions on pattern analysis and machine intelligence*, 45(6):7353–7366, 2022.

Daizong Ding, Mi Zhang, Fuli Feng, Yuanmin Huang, Erling Jiang, and Min Yang. Black-box adversarial attack on time series classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 7358–7368, 2023.

Yinpeng Dong, Fangzhou Liao, Tianyu Pang, Hang Su, Jun Zhu, Xiaolin Hu, and Jianguo Li. Boosting adversarial attacks with momentum. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 9185–9193, 2018.

Naifu Feng, Lixing Chen, Junhua Tang, Hua Ding, Jianhua Li, and Yang Bai. Fre-cw: Targeted attack on time series forecasting using frequency domain loss. *arXiv preprint arXiv:2508.08955*, 2025.

Thomas Fischer and Christopher Krauss. Deep learning with long short-term memory networks for financial market predictions. *European journal of operational research*, 270(2):654–669, 2018.

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.

Yuvaraj Govindarajulu, Avinash Amballa, Pavan Kulkarni, and Manojkumar Parmar. Targeted attacks on timeseries forecasting. *arXiv preprint arXiv:2301.11544*, 2023.

Chuan Guo, Jacob Gardner, Yurong You, Andrew Gordon Wilson, and Kilian Weinberger. Simple black-box adversarial attacks. In *International conference on machine learning*, pp. 2484–2493. PMLR, 2019.

Fazle Karim, Somshubra Majumdar, and Houshang Darabi. Adversarial attacks on time series. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3309–3320, 2020.

Hoki Kim, Yunyoung Lee, Woojin Lee, and Jaewook Lee. Towards undetectable adversarial attack on time series classification. *Information Sciences*, 715:122216, 2025.

Alexey Kurakin, Ian J Goodfellow, and Samy Bengio. Adversarial examples in the physical world. In *Artificial intelligence safety and security*, pp. 99–112. Chapman and Hall/CRC, 2018.

Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *arXiv preprint arXiv:1707.01926*, 2017.

Jiadong Lin, Chuanbiao Song, Kun He, Liwei Wang, and John E Hopcroft. Nesterov accelerated gradient and scale invariance for adversarial attacks. *arXiv preprint arXiv:1908.06281*, 2019.

Cheng Luo, Qinliang Lin, Weicheng Xie, Bizhu Wu, Jinheng Xie, and Linlin Shen. Frequency-driven imperceptible adversarial attack on semantic similarity. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 15315–15324, 2022.

Hassan Mahmood and Ehsan Elhamifar. Semantic-aware multi-label adversarial attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 24251–24262, 2024.

Gautam Raj Mode and Khaza Anuarul Hoque. Adversarial examples in deep learning for multivariate time series regression. In *2020 ieee applied imagery pattern recognition workshop (aipr)*, pp. 1–10. IEEE, 2020.

Guillermo Ortiz-Jimenez, Apostolos Modas, Seyed-Mohsen Moosavi, and Pascal Frossard. Hold me tight! influence of discriminative features on deep network boundaries. *Advances in Neural Information Processing Systems*, 33:2935–2946, 2020.

Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International conference on machine learning*, pp. 5301–5310. PMLR, 2019.

David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. Deepar: Probabilistic forecasting with autoregressive recurrent networks. *International journal of forecasting*, 36(3):1181–1191, 2020.

Ziyu Shen and Yun Li. Temporal characteristics-based adversarial attacks on time series forecasting. *Expert Systems with Applications*, 264:125950, 2025.

Yanyun Wang, Dehui Du, Haibo Hu, Zi Liang, and Yuanhao Liu. Tsfool: Crafting highly-imperceptible adversarial time series through multi-objective attack. *arXiv preprint arXiv:2209.06388*, 2022.

Hui Wei, Hao Tang, Xuemei Jia, Zhixiang Wang, Hanxun Yu, Zhubo Li, Shin'ichi Satoh, Luc Van Gool, and Zheng Wang. Physical adversarial attack meets computer vision: A decade survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):9797–9817, 2024.

Tao Wu, Xuechun Wang, Shaojie Qiao, Xingping Xian, Yanbing Liu, and Liang Zhang. Small perturbations are enough: Adversarial attacks on time series prediction. *Information Sciences*, 587:794–812, 2022.

Aidong Xu, Xuechun Wang, Yunan Zhang, Tao Wu, and Xingping Xian. Adversarial attacks on deep neural networks for time series prediction. In *2021 10th International Conference on Internet Computing for Science and Engineering*, pp. 8–14, 2021.

Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *arXiv preprint arXiv:1709.04875*, 2017.

Jiawen Zhang, Zhenwei Zhang, Shun Zheng, Xumeng Wen, Jia Li, and Jiang Bian. Are time-series foundation models deployment-ready? a systematic study of adversarial robustness across domains. *arXiv preprint arXiv:2505.19397*, 2025.

Yongliang Zhang and Jing Liu. Smg-diff: Adversarial attack method based on semantic mask-guided diffusion. In *International Conference on Multimedia Modeling*, pp. 44–57. Springer, 2025.

A Experimental Details

A.1 Implementation Details

All experiments are implemented in PyTorch based on the BasicTS[5] framework and conducted on a single NVIDIA GeForce RTX 3090 GPU. Both the input sequence length and prediction horizon are set to 96.

Following prior adversarial attacks, we adopt an L_∞ perturbation budget of $\epsilon = 0.2$. For all iterative attacks, we use $K = 50$ optimization steps with step size

$$\alpha = \frac{\epsilon}{K} = 0.004. \tag{11}$$

For the Adaptive Frequency Weighting (AFW) module, the learnable spectral parameters are optimized using Adam with a learning rate of 0.01.

Unless otherwise specified, the same perturbation constraint and iteration budget are used across attack methods to ensure a fair comparison.

All attacks are performed within the same preprocessing pipeline as the forecasting models. The input sequences and forecasting targets are normalized using the dataset scaler before model inference. Adversarial perturbations are applied only to the primary numerical feature channel, while auxiliary temporal features remain unchanged. After each iterative update, the perturbation is projected back into the feasible L_∞ region.

For frequency-domain operations, we apply the real-valued Fast Fourier Transform along the temporal dimension using orthonormal normalization, i.e., `rFFT`/`irFFT` with `norm='ortho'`.

For targeted semantic attacks, we evaluate three manipulation modes. For **spectral reduction (P1)**, we identify the strongest non-DC frequency component and attenuate it with a reduction strength of 0.85. For **trend reduction (P2)**, the trend component is extracted using Exponential Moving Average (EMA) with smoothing factor 0.2, and the trend manipulation intensity is set to 0.2. For **seasonality enhancement (P3)**, the sequence is decomposed using EMA with smoothing factor 0.2, and the seasonal component is scaled by a factor of 1.35.

For black-box transfer evaluation, adversarial examples are generated on one surrogate forecasting model and evaluated on the remaining unseen architectures. The reported transfer results are averaged over the corresponding target models.

A.2 Baseline Methods

We compare FUSE with eight representative adversarial attack methods, including four general gradient-based attacks and four time-series-oriented attacks.

- **FGSM** Goodfellow et al. (2014). Fast Gradient Sign Method generates adversarial examples using a single gradient-sign update. It is a simple and efficient first-order attack and is widely used as a basic adversarial baseline.
- **BIM** Kurakin et al. (2018). Basic Iterative Method extends FGSM by applying multiple small gradient-sign updates while projecting the perturbation back into the allowed constraint set. It provides a stronger iterative attack baseline than one-step FGSM.
- **MI-FGSM** Dong et al. (2018). Momentum Iterative FGSM incorporates a momentum term into iterative gradient updates to stabilize the optimization direction and reduce oscillations during adversarial example generation.
- **NI-FGSM** Lin et al. (2019). Nesterov Iterative FGSM introduces Nesterov accelerated gradients into the iterative attack process, allowing gradients to be evaluated at a look-ahead point and improving the transferability of adversarial examples.
- **Fre-C&W** Feng et al. (2025). Fre-C&W is an optimization-based adversarial attack for time-series forecasting that incorporates a frequency-domain loss into the C&W attack objective. By constraining the discrepancy between clean and adversarial samples in the frequency domain, it generates adversarial perturbations while preserving spectral similarity.

[5] https://github.com/GestaltCogTeam/BasicTS

- **TCA** Shen & Li (2025). TCA is a temporal-constrained adversarial attack for time-series forecasting. It introduces a time-domain cosine-similarity constraint to preserve the temporal structure between clean and adversarial inputs while optimizing the attack objective.
- **MAPGD-TSF** Govindarajulu et al. (2023). MAPGD-TSF is an optimization-based targeted attack for time-series forecasting. It iteratively updates adversarial perturbations toward predefined target forecasts under bounded perturbation constraints.
- **AAIM** Wu et al. (2022). AAIM is a sparse adversarial attack for time-series models that focuses perturbations on a limited number of influential timestamps, aiming to achieve effective attacks with localized modifications.

A.3 Semantic Target Construction

Given the predicted future sequence $\mathbf{Y} \in \mathbb{R}^{T\times M}$, we decompose it as

$$\mathbf{Y} = \mathbf{T} + \mathbf{S}, \tag{12}$$

where $\mathbf{T}$ and $\mathbf{S}$ denote the trend and seasonal components, respectively.

Trend Manipulation. By extracting the trend component $\mathbf{T}$ via Exponential Moving Average (EMA) and computing the trend increment $\Delta g = t_\tau - t_1$, we construct the target pattern as

$$\mathbf{Y}_{tar} = \mathbf{Y} - (1-\theta)\cdot \mathbf{l} \odot \Delta g, \tag{13}$$

where $\mathbf{l} = [0, \ldots, 1]^\top$ represents a linear temporal ramp and θ controls the trend intensity. Setting $0 < \theta < 1$ weakens the original trend, $\theta = 0$ approximately removes it, $\theta < 0$ reverses its direction, and $\theta > 1$ enhances the trend magnitude.

Seasonality Adjustment. To modify periodic fluctuation intensity while preserving the underlying trend, we define

$$\mathbf{Y}_{tar} = \mathbf{T} + \beta \cdot \mathbf{S}, \tag{14}$$

where β controls seasonal strength. Larger values induce stronger oscillations, while smaller values suppress periodic variation.

Spectral Tampering. Let

$$\mathbf{Z} = \mathrm{DFT}(\mathbf{Y}) \tag{15}$$

denote the spectral representation of the prediction. We identify the top-k dominant frequency components and apply a spectral filtering mask $\mathbf{\Omega}$:

$$\mathbf{Y}_{tar} = \mathrm{IDFT}\left(\mathbf{Z} \odot \mathbf{\Omega}\right), \tag{16}$$

where $\mathbf{\Omega}$ is configured to either suppress ($\omega_k \to 0$) or amplify ($\omega_k > 1$) selected spectral components, resulting in a target pattern with manipulated periodicity.

A.4 Evaluation Metrics

For untargeted attacks, we use MAE, MSE, and WAPE to quantify forecasting error. To measure the relative attack impact, we define the Degradation Percentage (P) as

$$P = \frac{\mathcal{L}(f(\mathbf{X}'), \mathbf{Y}) - \mathcal{L}(f(\mathbf{X}), \mathbf{Y})}{\mathcal{L}(f(\mathbf{X}), \mathbf{Y})} \times 100\%. \tag{17}$$

For targeted attacks, we additionally use two **novel metrics**: Attack Success Rate (ASR) and Target Adherence (η).

ASR considers an attack successful if the adversarial prediction error relative to the target $\mathbf{Y}_{tar}$ is bounded by the model's inherent forecasting error:

$$\mathrm{ASR} = \frac{1}{N}\sum_{i=1}^{N} \mathbb{I}\left(\mathcal{L}(f(\mathbf{X'}_i), \mathbf{Y}_{i,tar}) \le \mathcal{L}(f(\mathbf{X}_i), \mathbf{Y}_i)\right). \tag{18}$$

Algorithm 1 The FUSE Optimization Procedure

Require: Input $\mathbf{X}$; Ground Truth $\mathbf{Y}$; Target Pattern $\mathbf{Y}_{tar}$; Budget ϵ; Step size α; Iterations I.

Initialize $\boldsymbol{\delta}_0 = \mathbf{0}$ and learnable spectral parameters $\mathbf{w}_{par} \in \mathbb{R}^F$;
for $i = 0$ to $I - 1$ **do**
  Compute frequency weights $\mathbf{w}$ as in Eq. (4);
  Obtain data gradient $\mathbf{G}_i$ as in Eq. (3);
  Identify active mask $\mathcal{M}$ and compute redistributed temporal weights $\mathbf{s}'$ via Eq. (7)–(8);
  Calculate modulation signal $\mathbf{D}_{\text{time}} = \text{IDFT}\left(\text{DFT}(\mathbf{G}_i) \odot \mathbf{w}\right)$;
  Construct differentiable update $\boldsymbol{\delta}'_{i+1} = \text{clip}_\epsilon\left\{\boldsymbol{\delta}_i - \alpha\, \mathbf{D}_{time} \odot \mathbf{s}'\right\}$ while preserving the computational graph;
  Evaluate outer loss L_{outer} on the unrolled state $\mathbf{X} + \boldsymbol{\delta}'_{i+1}$;
  Optimize $\mathbf{w}_{par}$ by minimizing L_{outer} using Adam;
  Update $\boldsymbol{\delta}_{i+1}$ by detaching $\boldsymbol{\delta}'_{i+1}$;
end for
return $\mathbf{X}^* = \mathbf{X} + \boldsymbol{\delta}_I$.

Table 4: White-box untargeted attack results. Best results are highlighted in **bold** and second-best are underlined.

Dataset	Method	MAE (with Degradation Percentage %) ↑				MSE (with Degradation Percentage %) ↑				WAPE (with Degradation Percentage %) ↑			
		DLinear	PatchTST	TimesNet	Informer	DLinear	PatchTST	TimesNet	Informer	DLinear	PatchTST	TimesNet	Informer
ETTh1	Clean	0.3884(-0.0%)	0.3911(-0.0%)	0.4054(-0.0%)	0.6817(-0.0%)	0.3845(-0.0%)	0.3876(-0.0%)	0.3983(-0.0%)	0.8030(-0.0%)	0.5625(-0.0%)	0.5702(-0.0%)	0.5882(-0.0%)	0.8778(-0.0%)
	FGSM	0.5155(↑32.72%)	0.4863(↑24.34%)	0.5321(↑31.25%)	0.8241(↑20.89%)	0.5205(↑35.37%)	0.4928(↑27.14%)	0.5946(↑49.28%)	1.0848(↑35.09%)	0.7623(↑35.52%)	0.7181(↑25.94%)	0.7740(↑31.59%)	1.0877(↑23.91%)
	BIM	0.5194(↑33.73%)	0.4930(↑26.05%)	0.5685(↑40.23%)	0.8337(↑22.30%)	0.5224(↑35.86%)	0.4970(↑28.22%)	0.6620(↑66.21%)	1.0999(↑36.97%)	0.7682(↑36.57%)	0.7279(↑27.66%)	0.8251(↑40.28%)	1.1064(↑26.04%)
	MI-FGSM	0.5183(↑33.44%)	0.4905(↑25.42%)	0.5530(↑36.41%)	0.8323(↑22.09%)	0.5220(↑35.76%)	0.4956(↑27.86%)	0.6316(↑58.57%)	1.0997(↑36.95%)	0.7665(↑36.27%)	0.7243(↑27.01%)	0.8041(↑36.71%)	1.1026(↑25.61%)
	NI-FGSM	0.5207(↑34.06%)	0.4907(↑25.47%)	0.5722(↑41.14%)	0.8348(↑22.45%)	0.5236(↑36.18%)	0.4945(↑27.58%)	0.6693(↑68.04%)	1.1013(↑37.15%)	0.7702(↑36.92%)	0.7245(↑27.06%)	0.8293(↑40.99%)	1.1083(↑26.26%)
	Fre-C&W	0.4304(↑10.81%)	0.4292(↑9.74%)	0.4715(↑16.30%)	0.7789(↑14.26%)	0.4354(↑13.24%)	0.4345(↑12.10%)	0.5029(↑26.26%)	0.9956(↑23.99%)	0.6401(↑13.80%)	0.6345(↑11.28%)	0.6925(↑17.70%)	1.0365(↑18.08%)
	TCA	0.4830(↑24.37%)	0.4801(↑22.75%)	0.5680(↑40.10%)	0.8220(↑20.57%)	0.4853(↑26.21%)	0.4848(↑25.08%)	0.6609(↑65.94%)	1.0774(↑34.17%)	0.7131(↑26.77%)	0.7111(↑24.72%)	0.8241(↑40.11%)	1.0898(↑24.16%)
	MAPGD-TSF	0.5190(↑33.62%)	0.4924(↑25.89%)	0.5641(↑39.14%)	0.8328(↑22.16%)	0.5222(↑35.82%)	0.4966(↑28.12%)	0.6533(↑64.02%)	1.0985(↑36.80%)	0.7675(↑36.44%)	0.7269(↑27.49%)	0.8191(↑39.25%)	1.1045(↑25.83%)
	AAIM	0.4047(↑4.20%)	0.4064(↑3.91%)	0.4183(↑3.18%)	0.6886(↑1.01%)	0.3998(↑3.98%)	0.4015(↑3.59%)	0.4143(↑4.02%)	0.8151(↑1.51%)	0.5881(↑4.55%)	0.5935(↑4.09%)	0.6081(↑3.38%)	0.8880(↑1.16%)
	FUSE	**0.5240**(↑**34.91%**)	**0.4979**(↑**27.31%**)	**0.5901**(↑**45.56%**)	**0.8404**(↑**23.28%**)	**0.5253**(↑**36.62%**)	**0.5001**(↑**29.02%**)	**0.7039**(↑**76.73%**)	**1.1081**(↑**38.00%**)	**0.7755**(↑**37.78%**)	**0.7358**(↑**29.04%**)	**0.8539**(↑**45.17%**)	**1.1212**(↑**27.73%**)
ETTm1	Clean	0.3494(-0.0%)	0.3389(-0.0%)	0.3591(-0.0%)	0.4032(-0.0%)	0.3307(-0.0%)	0.3134(-0.0%)	0.3304(-0.0%)	0.3703(-0.0%)	0.5298(-0.0%)	0.5201(-0.0%)	0.5496(-0.0%)	0.6006(-0.0%)
	FGSM	0.5135(↑46.97%)	0.5152(↑52.02%)	0.5304(↑47.70%)	0.5491(↑36.19%)	0.4957(↑49.89%)	0.5330(↑70.07%)	0.6288(↑90.31%)	0.5923(↑59.95%)	0.8004(↑51.08%)	0.8117(↑56.07%)	0.8052(↑46.51%)	0.8258(↑37.50%)
	BIM	0.5174(↑48.08%)	0.6015(↑77.48%)	0.5973(↑66.33%)	0.5580(↑38.39%)	0.4976(↑50.47%)	0.6356(↑102.81%)	0.7830(↑136.99%)	0.6013(↑62.38%)	0.8072(↑52.36%)	0.9555(↑83.71%)	0.8987(↑63.52%)	0.8418(↑40.16%)
	MI-FGSM	0.5163(↑47.77%)	0.5816(↑71.61%)	0.5706(↑58.90%)	0.5554(↑37.73%)	0.4971(↑50.32%)	0.6097(↑94.54%)	0.7222(↑118.58%)	0.5997(↑61.95%)	0.8053(↑52.00%)	0.9210(↑77.08%)	0.8613(↑56.71%)	0.8368(↑39.33%)
	NI-FGSM	0.5190(↑48.54%)	0.5712(↑68.55%)	0.5994(↑66.92%)	0.5573(↑38.21%)	0.4988(↑50.83%)	0.5994(↑91.26%)	0.7809(↑136.35%)	0.5984(↑61.60%)	0.8109(↑52.88%)	0.9089(↑74.75%)	0.9012(↑63.99%)	0.8411(↑40.04%)
	Fre-C&W	0.4084(↑16.89%)	0.4876(↑43.88%)	0.4855(↑35.20%)	0.5445(↑35.04%)	0.3943(↑19.23%)	0.4874(↑55.52%)	0.5458(↑65.19%)	0.5122(↑38.32%)	0.6394(↑20.69%)	0.7800(↑49.97%)	0.7575(↑37.83%)	0.7845(↑30.62%)
	TCA	0.4471(↑27.96%)	0.5906(↑74.27%)	0.5938(↑65.35%)	0.5240(↑29.95%)	0.4246(↑28.38%)	0.6205(↑97.98%)	0.7760(↑134.86%)	0.5453(↑47.25%)	0.6980(↑31.75%)	0.9360(↑79.96%)	0.8936(↑62.60%)	0.7886(↑31.30%)
	MAPGD-TSF	0.5169(↑47.94%)	0.5966(↑76.05%)	0.5897(↑64.21%)	0.5573(↑38.22%)	0.4973(↑50.39%)	0.6293(↑100.79%)	0.7657(↑131.74%)	0.6004(↑62.15%)	0.8064(↑52.20%)	0.9472(↑82.12%)	0.8874(↑61.46%)	0.8406(↑39.96%)
	AAIM	0.3835(↑9.76%)	0.3715(↑9.62%)	0.3816(↑6.27%)	0.4181(↑3.70%)	0.3599(↑8.82%)	0.3432(↑9.51%)	0.3621(↑9.59%)	0.3890(↑5.05%)	0.5860(↑10.61%)	0.5729(↑10.15%)	0.5826(↑6.00%)	0.6243(↑3.95%)
	FUSE	**0.5217**(↑**49.31%**)	**0.6195**(↑**82.80%**)	**0.6334**(↑**76.39%**)	**0.5632**(↑**39.68%**)	**0.4999**(↑**51.16%**)	**0.6584**(↑**110.08%**)	**0.8653**(↑**161.89%**)	**0.6044**(↑**63.22%**)	**0.8148**(↑**53.79%**)	**0.9876**(↑**89.89%**)	**0.9508**(↑**73.00%**)	**0.8506**(↑**41.63%**)
ECL	Clean	0.2705(-0.0%)	0.2587(-0.0%)	0.2779(-0.0%)	0.3672(-0.0%)	0.1996(-0.0%)	0.1850(-0.0%)	0.1729(-0.0%)	0.2815(-0.0%)	0.3432(-0.0%)	0.3302(-0.0%)	0.3525(-0.0%)	0.4697(-0.0%)
	FGSM	0.4182(↑54.60%)	0.3764(↑45.50%)	0.7069(↑154.37%)	0.8812(↑139.98%)	0.3169(↑58.77%)	0.2789(↑50.76%)	0.8320(↑381.20%)	1.2138(↑331.19%)	0.5403(↑57.43%)	0.4878(↑47.73%)	0.8783(↑149.16%)	1.1410(↑142.92%)
	BIM	0.4218(↑55.93%)	0.3834(↑48.20%)	1.0833(↑289.82%)	1.1197(↑204.93%)	0.3184(↑59.52%)	0.2834(↑53.19%)	1.8503(↑970.16%)	1.8177(↑545.72%)	0.5455(↑58.95%)	0.4980(↑50.82%)	1.3424(↑280.82%)	1.4468(↑208.03%)
	MI-FGSM	0.4207(↑55.53%)	0.3811(↑47.31%)	0.9919(↑256.93%)	1.0348(↑181.81%)	0.3180(↑59.32%)	0.2819(↑52.38%)	1.5925(↑821.05%)	1.5901(↑464.87%)	0.5439(↑58.48%)	0.4947(↑49.82%)	1.2343(↑250.16%)	1.3379(↑184.84%)
	NI-FGSM	0.4231(↑56.41%)	0.3813(↑47.39%)	1.0902(↑292.30%)	1.1280(↑207.19%)	0.3192(↑59.92%)	0.2815(↑52.16%)	1.8670(↑979.81%)	1.8153(↑544.87%)	0.5472(↑59.44%)	0.4948(↑49.85%)	1.3515(↑283.40%)	1.4609(↑211.03%)
	Fre-C&W	0.3458(↑27.84%)	0.3217(↑24.35%)	1.0489(↑277.44%)	0.9303(↑153.35%)	0.2591(↑29.81%)	0.2371(↑28.16%)	1.7058(↑886.58%)	1.4079(↑400.14%)	0.4480(↑30.54%)	0.4178(↑26.53%)	1.3111(↑271.94%)	1.1828(↑151.82%)
	TCA	0.3662(↑35.38%)	0.3637(↑40.66%)	1.0833(↑289.80%)	1.1198(↑204.96%)	0.2690(↑34.77%)	0.2623(↑41.80%)	1.8503(↑970.16%)	1.8181(↑545.87%)	0.4741(↑38.14%)	0.4730(↑43.24%)	1.3434(↑280.83%)	1.4471(↑208.08%)
	MAPGD-TSF	0.4213(↑55.76%)	0.3826(↑47.90%)	1.0512(↑278.27%)	1.0886(↑196.46%)	0.3182(↑59.41%)	0.2829(↑52.89%)	1.7604(↑918.19%)	1.7318(↑515.19%)	0.5448(↑58.73%)	0.4969(↑50.49%)	1.3037(↑269.84%)	1.4069(↑199.54%)
	AAIM	0.2857(↑5.62%)	0.2737(↑5.80%)	0.3312(↑19.18%)	0.3861(↑5.15%)	0.2108(↑5.61%)	0.1963(↑6.11%)	0.2306(↑33.37%)	0.3034(↑7.78%)	0.3636(↑5.94%)	0.3503(↑6.09%)	0.4178(↑18.52%)	0.4936(↑5.09%)
	FUSE	**0.4271**(↑**57.89%**)	**0.3997**(↑**51.62%**)	**1.1639**(↑**318.82%**)	**1.2562**(↑**242.10%**)	**0.3211**(↑**60.87%**)	**0.2880**(↑**55.68%**)	**2.0931**(↑**1110.58%**)	**2.2334**(↑**693.39%**)	**0.5528**(↑**61.07%**)	**0.5081**(↑**53.88%**)	**1.4384**(↑**308.06%**)	**1.6235**(↑**245.65%**)
Weather	Clean	0.2317(-0.0%)	0.2061(-0.0%)	0.2087(-0.0%)	0.2454(-0.0%)	0.2080(-0.0%)	0.1733(-0.0%)	0.1669(-0.0%)	0.1982(-0.0%)	0.5647(-0.0%)	0.4930(-0.0%)	0.4996(-0.0%)	0.5310(-0.0%)
	FGSM	0.4367(↑88.48%)	0.3140(↑52.35%)	0.5346(↑156.16%)	0.4583(↑86.76%)	0.3497(↑68.13%)	0.2568(↑48.18%)	0.5806(↑247.87%)	0.4525(↑128.30%)	1.8234(↑220.90%)	1.0109(↑105.05%)	1.4532(↑190.87%)	1.0074(↑89.50%)
	BIM	0.4403(↑90.03%)	0.4879(↑136.73%)	0.6763(↑224.05%)	0.5285(↑115.36%)	0.3513(↑68.89%)	0.4240(↑144.66%)	0.8718(↑422.35%)	0.5760(↑190.62%)	1.8707(↑231.27%)	1.9438(↑294.28%)	1.8705(↑274.40%)	1.1784(↑121.67%)
	MI-FGSM	0.4393(↑89.60%)	0.4654(↑125.81%)	0.6485(↑210.73%)	0.5056(↑106.03%)	0.3509(↑68.70%)	0.3954(↑128.16%)	0.8079(↑384.06%)	0.5328(↑168.82%)	1.8595(↑229.29%)	1.8600(↑277.28%)	1.7614(↑252.56%)	1.1237(↑111.38%)
	NI-FGSM	0.4409(↑90.29%)	0.4104(↑99.13%)	0.6694(↑220.75%)	0.5112(↑108.31%)	0.3517(↑69.09%)	0.3486(↑101.15%)	0.8415(↑404.19%)	0.5445(↑174.72%)	1.8775(↑232.48%)	1.3696(↑177.81%)	1.5924(↑286.79%)	1.1398(↑114.41%)
	Fre-C&W	0.3497(↑50.93%)	0.4081(↑98.81%)	0.5083(↑143.56%)	0.4274(↑74.16%)	0.2861(↑37.55%)	0.3445(↑98.79%)	0.5747(↑244.34%)	0.4651(↑134.66%)	1.5275(↑170.50%)	1.6221(↑229.03%)	1.4170(↑183.63%)	0.9964(↑87.43%)
	TCA	0.3354(↑44.77%)	0.4667(↑126.44%)	0.5986(↑186.84%)	0.3988(↑62.51%)	0.2650(↑27.40%)	0.4005(↑131.13%)	0.7425(↑344.88%)	0.4184(↑111.11%)	1.5392(↑172.57%)	1.8844(↑282.23%)	1.8310(↑266.49%)	0.9653(↑81.58%)
	MAPGD-TSF	0.4399(↑89.85%)	0.4868(↑136.22%)	0.6635(↑217.92%)	0.5219(↑112.69%)	0.3511(↑68.80%)	0.4239(↑144.63%)	0.8423(↑404.70%)	0.5635(↑184.30%)	1.8643(↑230.14%)	1.9253(↑290.53%)	1.8126(↑262.80%)	1.1625(↑118.68%)
	AAIM	0.3197(↑37.98%)	0.2981(↑44.64%)	0.3044(↑45.86%)	0.3566(↑16.79%)	0.2574(↑23.75%)	0.2384(↑37.56%)	0.2630(↑57.58%)	0.2329(↑17.51%)	1.1169(↑97.79%)	0.8016(↑62.60%)	0.7296(↑46.04%)	0.6328(↑19.04%)
	FUSE	**0.4425**(↑**90.98%**)	**0.5093**(↑**147.11%**)	**0.6968**(↑**233.88%**)	**0.5679**(↑**131.42%**)	**0.3544**(↑**70.38%**)	**0.4499**(↑**159.61%**)	**0.9098**(↑**445.12%**)	**0.6573**(↑**231.63%**)	**1.8897**(↑**234.64%**)	**2.0277**(↑**311.30%**)	**1.9799**(↑**296.30%**)	**1.2641**(↑**137.79%**)

Target Adherence (η) measures the relative improvement toward the target pattern:

$$\eta = \frac{\mathcal{L}(f(\mathbf{X}), \mathbf{Y}_{tar}) - \mathcal{L}(f(\mathbf{X}'), \mathbf{Y}_{tar})}{\mathcal{L}(f(\mathbf{X}), \mathbf{Y}_{tar})} \times 100\%. \tag{19}$$

Here, $\mathcal{L}$ denotes the MAE loss. Higher P, ASR, and η indicate stronger attack performance.

B Algorithm Details

The complete optimization procedure of FUSE is summarized in Algorithm 1. At each iteration, FUSE first computes the adaptive frequency weights and the input gradients, then redistributes the temporal perturbation budget according to the active mask and temporal saliency. The spectral parameters are optimized through the outer objective, while the adversarial perturbation is iteratively updated under the L_∞ constraint.

C Additional White-box Results

C.1 White-box Untargeted Attacks

Analysis. Table 4 reports the complete white-box untargeted attack results across four datasets and four forecasting models.

(1) Classic gradient-based attacks, particularly BIM and NI-FGSM, remain highly competitive and often outperform several TSF-specific baselines, indicating that dense iterative gradient updates are already strong baselines for forecasting attacks. Nevertheless, FUSE consistently achieves stronger degradation. For example, on ETTm1 with TimesNet, FUSE increases MAE and MSE by 76.39% and 161.89%, respectively, compared with 66.92% and 136.35% for NI-FGSM.

(2) FUSE also consistently outperforms the TSF-oriented attacks. On ECL with TimesNet, FUSE increases MSE by 1110.58%, compared with 886.58%, 970.16%, and 918.19% for Fre-C&W, TCA, and MAPGD-TSF, respectively. These results indicate that coordinating adaptive spectral weighting with saliency-guided temporal redistribution provides additional attack effectiveness beyond purely temporal or fixed frequency-domain objectives.

Overall, FUSE achieves the strongest degradation across all 48 dataset–model–metric settings in Table 4, showing consistent effectiveness across datasets, forecasting architectures, and evaluation metrics.

C.2 White-box Targeted Attacks on ETTm1

Table 5: White-box targeted semantic induction results on **ETTm1**. **P1**, **P2**, and **P3** denote *Spectral Reduction*, *Trend Reduction*, and *Seasonality Enhancement*. **Init.** represents the initial distance between the clean prediction and the target pattern. Best results are in **bold** and second-best are underlined.

Pat.	Method	MAE (with Adherence %) ↓				MSE (with Adherence %) ↓				ASR % ↑			
		DLinear	PatchTST	TimesNet	Informer	DLinear	PatchTST	TimesNet	Informer	DLinear	PatchTST	TimesNet	Informer
P1	Init.	0.3758(-0.0%)	0.3887(-0.0%)	0.3844(-0.0%)	0.4084(-0.0%)	0.3258(-0.0%)	0.3526(-0.0%)	0.3373(-0.0%)	0.3401(-0.0%)	0.0%	0.0%	0.0%	0.0%
	FGSM	0.3289(12.48%)	0.3486(10.32%)	0.3423(10.95%)	0.3476(14.89%)	0.2821(13.41%)	0.3017(14.44%)	0.2570(23.81%)	0.2988(12.14%)	64.67%	49.22%	61.40%	97.55%
	BIM	0.3054(18.73%)	0.2893(25.57%)	0.2823(26.56%)	0.3321(18.68%)	0.2758(15.35%)	0.2639(25.16%)	0.1954(42.07%)	0.2840(16.49%)	75.37%	72.24%	86.31%	99.85%
	MI-FGSM	0.3111(17.22%)	0.3020(22.31%)	0.2970(22.74%)	0.3362(17.68%)	0.2766(15.10%)	0.2711(23.11%)	0.2081(38.30%)	0.2880(15.32%)	73.00%	68.60%	82.63%	99.66%
	NI-FGSM	0.3097(17.59%)	0.2982(23.28%)	0.2894(24.71%)	0.3345(18.10%)	0.2771(14.95%)	0.2693(23.62%)	0.1995(40.85%)	0.2841(16.46%)	73.72%	69.86%	83.64%	99.81%
	Fre-C&W	0.3703(1.46%)	0.3907(-0.51%)	0.3606(6.19%)	0.4367(-6.93%)	0.3185(2.24%)	0.3402(3.52%)	0.2896(14.14%)	0.4367(-28.40%)	40.69%	25.55%	50.23%	22.78%
	TCA	0.3054(18.73%)	0.2893(25.57%)	0.2823(26.56%)	0.3322(18.66%)	0.2758(15.35%)	0.2639(25.16%)	0.1954(42.07%)	0.2841(16.46%)	75.37%	72.24%	86.31%	99.84%
	MAPGD-TSF	0.3053(18.76%)	0.2890(25.65%)	0.2818(26.69%)	0.3323(18.63%)	0.2758(15.35%)	0.2639(25.16%)	0.1951(42.16%)	0.2841(16.46%)	75.42%	72.37%	86.35%	99.82%
	AAIM	0.4091(-8.86%)	0.4220(-8.57%)	0.4148(-7.91%)	0.4216(-3.23%)	0.3769(-15.68%)	0.4154(-17.81%)	0.3858(-14.38%)	0.4089(-20.23%)	23.19%	19.88%	26.90%	10.39%
	FUSE	**0.3047(18.92%)**	**0.2834(27.09%)**	**0.2735(28.85%)**	**0.3295(19.32%)**	**0.2753(15.50%)**	**0.2592(26.49%)**	**0.1856(44.97%)**	**0.2804(17.55%)**	**75.69%**	**74.31%**	**87.97%**	**99.91%**
P2	Init.	0.3574(-0.0%)	0.3509(-0.0%)	0.3728(-0.0%)	0.4206(-0.0%)	0.3261(-0.0%)	0.3199(-0.0%)	0.3404(-0.0%)	0.3851(-0.0%)	0.0%	0.0%	0.0%	0.0%
	FGSM	0.2785(22.08%)	0.2798(20.26%)	0.3228(13.41%)	0.3276(22.11%)	0.2356(27.75%)	0.2201(31.20%)	0.2507(26.35%)	0.2557(26.59%)	91.12%	86.39%	72.62%	89.55%
	BIM	0.2481(30.58%)	0.2133(39.21%)	0.2554(31.49%)	0.3112(26.01%)	0.2276(30.20%)	0.1782(44.30%)	0.1844(45.83%)	0.2421(37.13%)	95.36%	97.16%	94.65%	91.37%
	MI-FGSM	0.2546(28.76%)	0.2266(35.42%)	0.2686(27.95%)	0.3156(24.96%)	0.2283(30.00%)	0.1858(41.92%)	0.1943(42.92%)	0.2455(36.25%)	94.72%	96.35%	93.04%	90.97%
	NI-FGSM	0.2513(29.69%)	0.2210(37.02%)	0.2592(30.47%)	0.3138(25.39%)	0.2283(30.00%)	0.1825(42.95%)	0.1852(45.59%)	0.2429(36.92%)	94.99%	96.42%	94.03%	91.22%
	Fre-C&W	0.3755(-5.06%)	0.3966(-13.02%)	0.3712(0.43%)	0.4084(2.90%)	0.3238(0.71%)	0.3434(-7.34%)	0.3283(3.55%)	0.3737(2.96)	33.03%	23.67%	42.75%	48.65%
	TCA	0.2481(30.58%)	0.2133(39.21%)	0.2554(31.49%)	0.3113(25.99%)	0.2276(30.20%)	0.1782(44.30%)	0.1844(45.83%)	0.2421(37.13%)	95.36%	97.16%	94.65%	91.41%
	MAPGD-TSF	0.2480(30.61%)	0.2131(39.27%)	0.2550(31.60%)	0.3114(25.96%)	0.2276(30.20%)	0.1783(44.26%)	0.1841(45.92%)	0.2422(37.11%)	95.36%	97.19%	94.66%	91.36%
	AAIM	0.3451(3.44%)	0.3339(4.84%)	0.3562(4.45%)	0.4085(2.88%)	0.3112(4.57%)	0.2997(6.31%)	0.3162(7.11%)	0.3697(4.00%)	57.34%	61.29%	57.83%	54.03%
	FUSE	**0.2476(30.72%)**	**0.2071(40.98%)**	**0.2479(33.50%)**	**0.3082(26.72%)**	**0.2271(30.36%)**	**0.1722(46.17%)**	**0.1751(48.56%)**	**0.2384(38.09%)**	**95.41%**	**97.46%**	**95.82%**	**91.68%**
P3	Init.	0.3828(-0.0%)	0.3709(-0.0%)	0.3890(-0.0%)	0.4328(-0.0%)	0.3829(-0.0%)	0.3632(-0.0%)	0.3795(-0.0%)	0.4240(-0.0%)	0.0%	0.0%	0.0%	0.0%
	FGSM	0.3052(20.27%)	0.3076(17.07%)	0.3455(11.18%)	0.3476(19.69%)	0.2886(24.63%)	0.2716(25.22%)	0.2931(22.77%)	0.2988(2.1%)	95.32%	80.73%	61.55%	97.55%
	BIM	0.2784(27.27%)	0.2448(34.00%)	0.2801(27.99%)	0.3321(23.27%)	0.2807(26.70%)	0.2272(37.44%)	0.2244(40.87%)	0.2840(2.1%)	99.96%	100.00%	99.65%	99.85%
	MI-FGSM	0.2840(25.81%)	0.2566(30.82%)	0.2922(24.88%)	0.3362(22.32%)	0.2815(26.48%)	0.2348(35.35%)	0.2334(38.50%)	0.2880(2.1%)	99.90%	100.00%	99.13%	99.66%
	NI-FGSM	0.2819(26.36%)	0.2521(32.03%)	0.2833(27.17%)	0.3345(22.72%)	0.2813(26.53%)	0.2310(36.40%)	0.2254(40.61%)	0.2844(2.1%)	99.95%	100.00%	99.57%	99.81%
	Fre-C&W	0.4054(-5.90%)	0.4299(-15.91%)	0.4020(-3.34%)	0.4367(-0.90%)	0.3861(-0.84%)	0.4108(-13.11%)	0.3851(-1.48%)	0.4367(2.1%)	7.39%	10.81%	22.09%	22.78%
	TCA	0.2784(27.27%)	0.2448(34.00%)	0.2801(27.99%)	0.3322(23.24%)	0.2807(26.70%)	0.2272(37.44%)	0.2244(40.87%)	0.2841(2.1%)	99.96%	100.00%	99.65%	99.84%
	MAPGD-TSF	0.2783(27.30%)	0.2446(34.05%)	0.2798(28.07%)	**0.2841(8.3%)**	0.2807(26.69%)	0.2273(37.42%)	0.2241(40.95%)	0.4751(2.1%)	99.96%	100.00%	99.71%	99.82%
	AAIM	0.3706(3.19%)	0.3560(4.02%)	0.3736(3.96%)	0.4216(6.4%)	0.3670(4.15%)	0.3442(9.30%)	0.3557(6.27%)	0.4089(4.9%)	9.11%	16.07%	18.91%	10.39%
	FUSE	**0.2778(27.43%)**	**0.2392(35.51%)**	**0.2740(29.56%)**	0.3292(25.1%)	**0.2801(26.85%)**	**0.2212(39.10%)**	**0.2164(42.98%)**	**0.2798(24.6%)**	**99.99%**	**100.00%**	**99.73%**	**99.90%**

D Black-box Transfer Results

D.1 Transfer Results on ETTh1

Analysis. Table 6 reports the black-box targeted transfer results on ETTh1 under the three semantic objectives.

(1) FUSE demonstrates competitive transferability across different surrogate architectures. Under P1, it achieves the best average MSE and ASR, reaching 0.4999 and 44.85%, respectively, while maintaining an average MAE of 0.4797, which is close to the best result. Under P3, FUSE obtains the best average MAE (0.5177) and the highest average ASR (10.55%), indicating that its transfer advantage remains observable even for the more challenging seasonality-enhancement objective.

(2) The transferability varies substantially across semantic objectives. P1 achieves the highest overall transfer success, whereas P3 is considerably more difficult, with average ASR values close to 10% even for the strongest methods. For P2, FUSE remains highly competitive, achieving an average ASR of 40.64%, close to the best result of 40.74% obtained by NI-FGSM. These results indicate that black-box targeted transfer depends not only on the surrogate model but also strongly on the type of semantic manipulation.

Table 6: Black-box targeted semantic induction results on **ETTh1**. **P1**: Spectral Reduction, **P2**: Trend Reduction, **P3**: Seasonal Enhancement. Each column represents the **surrogate model** used to craft attacks; results are the averaged performance across all other target models. Best results are in **bold** and second-best are underlined.

Pat.	Method	MAE (with Adherence %) ↓					MSE (with Adherence %) ↓					ASR % ↑				
		DLin.	PTST	TNet	Inf.	Avg.	DLin.	PTST	TNet	Inf.	Avg.	DLin.	PTST	TNet	Inf.	Avg.
P1	FGSM	0.4834(10.43%)	0.4869(9.62%)	0.5101(4.75%)	0.4577(2.15%)	0.4845(6.74%)	0.5120(13.59%)	0.5120(14.69%)	0.5403(9.10%)	0.4527(7.29%)	0.5042(11.17%)	54.51%	53.36%	42.19%	21.06%	42.78%
	BIM	0.4800(11.17%)	0.4844(10.17%)	0.5055(5.71%)	0.4519(3.39%)	0.4804(7.61%)	0.5109(13.87%)	0.5114(14.84%)	0.5365(9.78%)	0.4461(6.64%)	0.5012(11.78%)	55.96%	55.19%	44.66%	22.18%	44.50%
	MI-FGSM	0.4802(11.11%)	0.4846(10.11%)	0.5064(5.50%)	0.4545(2.84%)	0.4814(7.39%)	0.5102(13.97%)	0.5108(14.94%)	0.5367(9.74%)	0.4487(8.11%)	0.5016(11.69%)	55.94%	54.82%	44.13%	21.74%	44.16%
	NI-FGSM	0.4803(11.11%)	0.4839(10.25%)	**0.5048(5.86%)**	**0.4496(3.88%)**	**0.4796(7.78%)**	0.5109(13.86%)	**0.5106(14.85%)**	0.5356(9.99%)	**0.4431(9.25%)**	0.5000(12.01%)	55.91%	**55.68%**	**45.09%**	**22.64%**	44.83%
	Fre-C&W	0.5100(5.29%)	0.5054(5.82%)	0.5199(2.82%)	0.4612(1.40%)	0.4991(3.83%)	0.5442(8.03%)	0.5321(11.03%)	0.5573(5.99%)	0.4615(5.51%)	0.5238(7.64%)	40.65%	43.58%	37.58%	17.71%	34.88%
	TCA	0.4800(11.18%)	0.4843(10.17%)	0.5055(5.71%)	0.4519(3.40%)	0.4804(7.61%)	0.5108(13.88%)	0.5114(14.85%)	0.5364(9.80%)	0.4461(6.64%)	0.5012(11.79%)	56.20%	55.20%	44.60%	22.19%	44.55%
	MAPGD-TSF	0.4800(11.17%)	0.4844(10.17%)	0.5054(5.72%)	0.4519(3.39%)	0.4804(7.61%)	0.5110(13.86%)	0.5115(14.84%)	0.5364(9.80%)	0.4461(8.63%)	0.5012(11.78%)	56.04%	55.19%	44.64%	22.18%	44.51%
	AAIM	0.5292(1.42%)	0.5272(1.59%)	0.5323(0.58%)	0.4667(0.22%)	0.5138(0.95%)	0.5836(0.43%)	0.5770(2.99%)	0.5877(0.35%)	0.4865(0.48%)	0.5586(1.24%)	31.85%	32.03%	30.10%	15.25%	27.31%
	FUSE	**0.4792(11.33%)**	**0.4838(10.27%)**	0.5053(5.74%)	0.4503(3.73%)	0.4797(7.77%)	**0.5097(14.09%)**	0.5108(14.95%)	**0.5355(9.95%)**	0.4436(9.15%)	**0.4999(12.04%)**	**56.49%**	55.62%	45.06%	22.25%	**44.85%**
P2	FGSM	0.4986(9.34%)	0.5183(5.66%)	0.5360(0.52%)	0.4471(0.80%)	0.5000(4.03%)	0.5554(13.44%)	0.5906(7.29%)	0.6119(2.08%)	0.4625(4.28%)	0.5551(6.77%)	55.31%	44.42%	29.73%	22.45%	37.98%
	BIM	0.4954(10.09%)	0.5150(6.22%)	**0.5325(1.31%)**	0.4441(1.47%)	0.4967(4.77%)	0.5550(13.61%)	0.5887(7.75%)	0.6093(2.66%)	0.4599(4.83%)	0.5532(7.21%)	58.03%	46.60%	32.75%	24.24%	40.40%
	MI-FGSM	0.4953(10.09%)	0.5156(6.09%)	0.5331(1.14%)	0.4457(1.11%)	0.4974(4.61%)	0.5539(13.77%)	0.5890(7.68%)	**0.6090(2.65%)**	0.4609(4.61%)	0.5532(7.18%)	57.67%	46.25%	32.07%	23.11%	39.77%
	NI-FGSM	0.4959(9.99%)	**0.5136(6.51%)**	0.5325(1.33%)	**0.4431(1.68%)**	**0.4963(4.88%)**	0.5556(13.52%)	**0.5860(8.34%)**	0.6094(2.70%)	**0.4590(5.02%)**	**0.5525(7.37%)**	57.80%	**47.43%**	**32.92%**	**24.79%**	**40.74%**
	Fre-C&W	0.5527(-1.76%)	0.5488(-1.26%)	0.5413(-1.01%)	0.4609(-2.24%)	0.5259(-1.57%)	0.6290(0.22%)	0.6263(0.03%)	0.6175(0.22%)	0.4881(-0.99%)	0.5902(-0.13%)	23.61%	25.17%	25.13%	17.89%	22.95%
	TCA	0.4955(10.08%)	0.5149(6.24%)	0.5326(1.30%)	0.4441(1.47%)	0.4968(4.77%)	0.5551(13.59%)	0.5885(7.78%)	0.6094(2.64%)	0.4599(4.83%)	0.5532(7.21%)	58.08%	46.68%	32.59%	24.33%	40.42%
	MAPGD-TSF	0.4954(10.08%)	0.5149(6.24%)	0.5326(1.29%)	0.4441(1.47%)	0.4968(4.77%)	0.5551(13.60%)	0.5885(7.77%)	0.6095(2.63%)	0.4599(4.83%)	0.5532(7.21%)	58.06%	46.71%	32.70%	24.30%	40.44%
	AAIM	0.5384(1.24%)	0.5384(1.23%)	0.5368(0.39%)	0.4500(0.18%)	0.5159(0.76%)	0.6227(1.43%)	0.6221(1.40%)	0.6194(0.68%)	0.4816(0.35%)	0.5864(0.96%)	31.10%	30.98%	28.11%	23.33%	28.38%
	FUSE	**0.4946(10.25%)**	0.5142(6.41%)	0.5326(1.29%)	0.4441(1.48%)	0.4963(4.86%)	**0.5538(13.83%)**	0.5876(7.97%)	0.6095(2.63%)	0.4599(4.83%)	0.5527(7.32%)	**58.37%**	47.22%	32.70%	24.27%	40.64%
P3	FGSM	0.5283(7.56%)	0.5353(6.02%)	0.5521(1.64%)	0.4668(0.25%)	0.5206(3.87%)	0.6314(0.84%)	0.6444(7.77%)	0.6613(3.70%)	0.5259(1.19%)	0.6157(5.63%)	20.66%	11.93%	2.03%	1.03%	8.91%
	BIM	0.5258(8.11%)	0.5327(6.59%)	0.5487(2.37%)	0.4648(0.67%)	0.5180(4.44%)	0.6305(10.05%)	0.6431(8.08%)	0.6581(4.34%)	0.5239(1.56%)	0.6139(6.01%)	22.73%	15.00%	2.00%	1.30%	10.26%
	MI-FGSM	0.5258(8.11%)	0.5333(6.48%)	0.5489(2.30%)	0.4660(0.43%)	0.5185(4.33%)	0.6303(10.09%)	0.6435(8.00%)	**0.6574(4.40%)**	0.5250(1.36%)	0.6141(5.96%)	23.03%	14.53%	**2.42%**	1.12%	10.28%
	NI-FGSM	0.5262(8.04%)	**0.5320(6.74%)**	0.5488(2.36%)	**0.4644(0.77%)**	0.5178(4.48%)	0.6306(10.05%)	**0.6415(8.37%)**	0.6584(4.35%)	**0.5233(1.68%)**	**0.6134(6.11%)**	22.19%	**15.80%**	2.13%	**1.40%**	10.38%
	Fre-C&W	0.5786(-2.52%)	0.5754(-2.45%)	0.5670(-1.79%)	0.4793(-2.43%)	0.5501(-2.30%)	0.7004(-1.71%)	0.7007(-2.33%)	0.6855(-1.19%)	0.5442(-2.23%)	0.6577(-1.87%)	0.96%	1.31%	1.41%	0.73%	1.10%
	TCA	0.5259(8.11%)	0.5327(6.59%)	**0.5486(2.38%)**	0.4648(0.67%)	0.5180(4.44%)	0.6308(10.03%)	0.6431(8.08%)	0.6579(4.36%)	0.5239(1.56%)	0.6139(6.01%)	22.78%	14.90%	2.16%	1.32%	10.29%
	MAPGD-TSF	0.5258(8.12%)	0.5327(6.60%)	0.5486(2.37%)	0.4648(0.67%)	0.5180(4.44%)	0.6305(10.06%)	0.6430(8.10%)	0.6579(4.36%)	0.5239(1.56%)	0.6138(6.02%)	22.79%	14.96%	2.17%	1.29%	10.30%
	AAIM	0.5614(0.93%)	0.5598(0.97%)	0.5588(0.33%)	0.4673(0.14%)	0.5368(0.59%)	0.6853(0.96%)	0.6842(1.00%)	0.6801(0.47%)	0.5311(0.20%)	0.6452(0.66%)	0.01%	0.07%	0.02%	0.00%	0.03%
	FUSE	**0.5253(8.22%)**	0.5322(6.71%)	0.5486(2.37%)	0.4648(0.67%)	**0.5177(4.49%)**	**0.6297(10.20%)**	0.6423(8.23%)	0.6579(4.36%)	0.5239(1.56%)	0.6135(6.09%)	**23.33%**	15.43%	2.17%	1.29%	**10.55%**

D.2 TRANSFER RESULTS ON ETTM1

Analysis. Table 7 reports the black-box targeted transfer results on ETTm1 under the three semantic objectives.

(1) FUSE remains competitive across different surrogate models and semantic objectives, although the transfer gains are smaller than those observed in the white-box setting. For P1, FUSE achieves the best average MAE and MSE (0.3999 and 0.3610), while obtaining an average ASR of 35.19%, which is close to the best result of 35.33% achieved by NI-FGSM. For P3, FUSE obtains the best average MAE (0.3718) and achieves an average ASR of 33.79%, again remaining close to the strongest baseline.

(2) Transfer performance varies across semantic objectives. P2 generally exhibits higher ASR than P1 and P3, with FUSE reaching an average ASR of 61.36%. In contrast, P1 shows negative Target Adherence for several surrogate models, indicating that spectral reduction is more difficult to transfer consistently across heterogeneous forecasting architectures. Overall, these results suggest that the effectiveness of targeted transfer attacks depends strongly on both the surrogate architecture and the semantic objective.

Table 7: Black-box targeted semantic induction results on **ETTm1**. **P1**: Spectral Reduction, **P2**: Trend Reduction, **P3**: Seasonal Enhancement. Each column represents the **surrogate model** used to craft attacks; results are the averaged performance across all other target models. Best results are in **bold** and second-best are underlined.

Pat.	Method	MAE (with Adherence %) ↓					MSE (with Adherence %) ↓					ASR % ↑				
		DLin.	PTST	TNet	Inf.	Avg.	DLin.	PTST	TNet	Inf.	Avg.	DLin.	PTST	TNet	Inf.	Avg.
P1	FGSM	0.3806(3.38%)	0.4099(-5.24%)	0.4157(-6.37%)	0.4156(-8.51%)	0.4055(-4.19%)	0.3380(1.49%)	0.3737(-11.70%)	0.3678(-8.32%)	0.3805(-12.34%)	0.3650(-7.72%)	47.27%	33.09%	26.08%	23.14%	32.39%
	BIM	0.3739(5.12%)	0.4074(-4.55%)	0.4092(-4.69%)	0.4098(-6.99%)	0.4000(-2.78%)	0.3350(2.36%)	0.3705(-10.74%)	0.3655(-7.63%)	0.3746(-10.61%)	0.3614(-6.66%)	50.48%	33.68%	29.78%	26.31%	35.06%
	MI-FGSM	0.3752(4.78%)	**0.4063(-4.29%)**	0.4100(-4.90%)	0.4121(-7.59%)	0.4009(-3.00%)	0.3357(2.16%)	0.3708(-10.82%)	**0.3637(-7.10%)**	0.3768(-11.27%)	0.3617(-6.76%)	49.94%	**34.90%**	29.49%	25.02%	34.84%
	NI-FGSM	0.3759(4.60%)	0.4076(-4.61%)	0.4089(-4.60%)	**0.4077(-6.44%)**	0.4000(-2.76%)	0.3375(1.61%)	0.3714(-11.00%)	0.3641(-7.22%)	**0.3718(-9.79%)**	0.3612(-6.60%)	49.70%	33.91%	**30.32%**	**27.38%**	**35.33%**
	Fre-C&W	0.4023(-2.13%)	0.4248(-9.10%)	0.4212(-7.78%)	0.4363(-13.93%)	0.4212(-8.24%)	0.3625(-5.65%)	0.3872(-15.77%)	0.3804(-12.03%)	0.4142(-22.31%)	0.3861(-13.94%)	35.54%	33.78%	22.21%	15.98%	24.38%
	TCA	0.3739(5.12%)	0.4073(-4.55%)	0.4092(-4.69%)	0.4098(-6.99%)	0.4001(-2.78%)	0.3350(2.35%)	**0.3705(-10.74%)**	0.3655(-7.63%)	0.3746(-10.61%)	0.3614(-6.66%)	50.47%	33.67%	29.82%	26.28%	35.06%
	MAPGD-TSF	0.3739(5.12%)	0.4074(-4.57%)	0.4091(-4.67%)	0.4098(-6.99%)	0.4001(-2.78%)	0.3350(2.34%)	0.3706(-10.77%)	0.3653(-7.59%)	0.3746(-10.62%)	0.3614(-6.66%)	50.49%	33.62%	29.86%	26.30%	35.07%
	AAIM	0.4319(-9.69%)	0.4296(-10.32%)	0.4338(-11.00%)	0.4297(-12.20%)	0.4312(-10.80%)	0.4082(-18.90%)	0.3979(-19.00%)	0.4055(-19.41%)	0.4125(-21.79%)	0.4060(-19.77%)	22.74%	22.63%	20.51%	17.25%	20.78%
	FUSE	**0.3738(5.14%)**	0.4074(-4.57%)	**0.4087(-4.57%)**	0.4096(-6.94%)	**0.3999(-2.73%)**	**0.3349(2.38%)**	0.3706(-10.77%)	0.3648(-7.42%)	0.3736(-10.32%)	**0.3610(-6.53%)**	**50.56%**	33.62%	30.26%	26.34%	35.19%
P2	FGSM	0.3381(11.62%)	0.3661(4.64%)	0.3753(0.18%)	0.3540(1.75%)	0.3584(4.55%)	0.2887(17.54%)	0.3277(6.65%)	0.3295(4.17%)	0.3052(7.19%)	0.3128(8.89%)	71.89%	57.38%	44.47%	45.91%	54.91%
	BIM	0.3287(14.23%)	0.3605(6.21%)	0.3671(2.48%)	0.3486(3.27%)	0.3512(6.55%)	0.2840(18.97%)	0.3236(7.86%)	0.3248(5.57%)	0.3008(8.55%)	0.3083(10.24%)	**77.65%**	61.54%	53.45%	51.93%	61.39%
	MI-FGSM	0.3307(13.67%)	**0.3600(6.34%)**	0.3689(1.95%)	0.3509(2.61%)	0.3526(6.14%)	0.2852(18.60%)	**0.3235(7.88%)**	0.3249(5.55%)	0.3026(7.99%)	0.3090(10.01%)	76.40%	**62.64%**	51.47%	49.08%	59.90%
	NI-FGSM	0.3295(14.04%)	0.3602(6.30%)	0.3676(2.33%)	**0.3472(3.66%)**	**0.3511(6.58%)**	0.2849(18.71%)	0.3237(7.87%)	**0.3239(5.85%)**	**0.2990(9.10%)**	**0.3079(10.38%)**	77.07%	62.56%	52.55%	**53.58%**	**61.44%**
	Fre-C&W	0.3914(-2.67%)	0.3913(-2.03%)	0.3853(-2.50%)	0.3737(-3.71%)	0.3854(-2.73%)	0.3458(0.81%)	0.3562(-1.58%)	0.3483(-1.40%)	0.3339(-1.53%)	0.3460(-0.93%)	35.62%	35.92%	35.49%	33.19%	35.05%
	TCA	0.3287(14.24%)	0.3605(6.21%)	**0.3671(2.48%)**	0.3486(3.27%)	0.3512(6.55%)	0.2839(18.98%)	0.3236(7.87%)	0.3248(5.57%)	0.3008(8.55%)	0.3083(10.24%)	77.61%	62.56%	**53.46%**	51.90%	61.38%
	MAPGD-TSF	0.3287(14.24%)	0.3605(6.21%)	0.3671(2.47%)	0.3486(3.27%)	0.3512(6.55%)	0.2840(18.95%)	0.3236(7.87%)	0.3248(5.59%)	0.3008(8.54%)	0.3083(10.24%)	77.57%	62.53%	53.36%	51.90%	61.34%
	AAIM	0.3750(1.73%)	0.3792(1.16%)	0.3744(0.50%)	0.3586(0.48%)	0.3718(0.97%)	0.3400(2.48%)	0.3448(1.66%)	0.3389(1.41%)	0.3249(1.20%)	0.3372(1.69%)	49.13%	47.90%	45.39%	45.65%	47.02%
	FUSE	**0.3285(14.28%)**	0.3605(6.23%)	0.3671(2.47%)	0.3486(3.27%)	0.3512(6.56%)	**0.2839(19.00%)**	0.3235(7.89%)	0.3248(5.59%)	0.3008(8.54%)	0.3082(10.26%)	77.62%	62.58%	53.36%	51.90%	61.36%
P3	FGSM	0.3599(9.69%)	0.3843(4.37%)	0.3930(0.59%)	0.3771(0.97%)	0.3786(3.91%)	0.3379(13.42%)	0.3725(5.88%)	0.3729(4.45%)	0.3549(8.41%)	0.3595(7.29%)	56.15%	27.76%	10.03%	17.73%	27.91%
	BIM	0.3514(11.95%)	0.3791(5.73%)	0.3846(2.79%)	0.3724(2.22%)	0.3719(5.67%)	0.3321(14.95%)	0.3684(6.94%)	0.3673(5.89%)	0.3510(6.45%)	0.3547(8.56%)	64.63%	35.78%	13.62%	20.60%	33.66%
	MI-FGSM	0.3532(11.46%)	**0.3789(5.76%)**	0.3864(2.31%)	0.3744(1.69%)	0.3733(5.31%)	0.3337(14.52%)	0.3687(6.85%)	0.3675(5.85%)	0.3525(6.04%)	0.3556(8.32%)	62.88%	36.66%	12.91%	19.57%	33.01%
	NI-FGSM	0.3523(11.72%)	0.3790(5.76%)	0.3853(2.62%)	**0.3713(2.53%)**	0.3719(5.66%)	0.3327(14.80%)	0.3685(6.92%)	**0.3669(6.01%)**	**0.3496(6.84%)**	**0.3544(8.64%)**	63.83%	**36.78%**	13.50%	**21.17%**	**33.82%**
	Fre-C&W	0.4122(-3.76%)	0.4111(-2.41%)	0.4085(-3.38%)	0.3972(-4.30%)	0.4073(-3.46%)	0.3933(-1.11%)	0.4040(-2.14%)	0.4010(-2.85%)	0.3851(-2.65%)	0.3959(-2.19%)	9.84%	4.55%	3.59%	6.18%	6.04%
	TCA	0.3514(11.94%)	0.3791(5.73%)	**0.3846(2.80%)**	0.3724(2.22%)	0.3719(5.67%)	0.3321(14.94%)	0.3684(6.94%)	0.3673(5.90%)	0.3510(6.45%)	0.3547(8.56%)	64.64%	35.79%	**13.68%**	20.65%	33.69%
	MAPGD-TSF	0.3513(11.95%)	0.3791(5.72%)	0.3846(2.78%)	0.3724(2.22%)	0.3719(5.67%)	0.3321(14.95%)	0.3685(6.92%)	0.3673(5.90%)	0.3511(6.44%)	0.3547(8.55%)	64.61%	35.76%	13.64%	20.67%	33.67%
	AAIM	0.3919(1.46%)	0.3976(0.99%)	0.3941(0.33%)	0.3794(0.39%)	0.3908(0.79%)	0.3810(2.08%)	0.3899(1.42%)	0.3856(1.14%)	0.3712(0.06%)	0.3813(1.42%)	7.01%	3.74%	2.81%	3.25%	4.20%
	FUSE	**0.3513(11.97%)**	0.3790(5.76%)	0.3846(2.78%)	0.3724(2.22%)	**0.3718(5.68%)**	**0.3319(15.00%)**	**0.3683(6.96%)**	0.3673(5.90%)	0.3511(6.44%)	0.3547(8.57%)	**64.73%**	36.11%	13.64%	20.67%	33.79%

E COMPLETE ABLATION STUDIES

E.1 WHITE-BOX ABLATION

Table 8: Complete white-box ablation results on **ETTh1** and **ETTm1** with **DLinear** and **TimesNet**. P1/P2/P3 denote spectral reduction, trend reduction, and seasonality enhancement. Best/second-best results are **bold**/underlined.

Pattern	Method	MAE (η%) ↓	MSE (η%) ↓	ASR (%) ↑
		ETTh1 – DLinear		
P1	**FUSE (Full)**	**0.3614(22.32%)**	**0.3650(24.27%)**	**63.56%**
	w/o AFW	0.3621(22.17%)	0.3659(24.09%)	62.94%
	w/o Redist	0.3633(21.91%)	0.3672(23.82%)	61.98%
	w/o Saliency	0.3627(22.04%)	0.3665(23.96%)	62.47%
	BIM	0.3625(22.08%)	0.3657(24.13%)	62.55%
P2	**FUSE (Full)**	**0.3575(19.72%)**	**0.3760(20.91%)**	**76.01%**
	w/o AFW	0.3584(19.52%)	0.3771(20.68%)	75.28%
	w/o Redist	0.3598(19.20%)	0.3788(20.32%)	74.31%
	w/o Saliency	0.3591(19.36%)	0.3780(20.49%)	74.86%
	BIM	0.3590(19.38%)	0.3769(20.72%)	75.46%
P3	**FUSE (Full)**	**0.3764(18.28%)**	**0.4302(17.83%)**	**65.08%**
	w/o AFW	0.3770(18.15%)	0.4309(17.69%)	64.41%
	w/o Redist	0.3782(17.89%)	0.4322(17.44%)	63.52%
	w/o Saliency	0.3776(18.02%)	0.4314(17.59%)	63.97%
	BIM	0.3778(17.98%)	0.4312(17.63%)	63.96%
		ETTh1 – TimesNet		
P1	**FUSE (Full)**	**0.3563(23.95%)**	**0.3110(35.66%)**	**81.15%**
	w/o AFW	0.3578(23.63%)	0.3130(35.25%)	80.45%
	w/o Redist	0.3605(23.05%)	0.3164(34.55%)	79.70%
	w/o Saliency	0.3593(23.31%)	0.3148(34.88%)	80.05%
	BIM	0.3622(22.69%)	0.3186(34.09%)	79.96%
P2	**FUSE (Full)**	**0.3668(20.54%)**	**0.3482(29.87%)**	**78.24%**
	w/o AFW	0.3681(20.26%)	0.3503(29.45%)	77.45%
	w/o Redist	0.3704(19.76%)	0.3532(28.86%)	76.30%
	w/o Saliency	0.3695(19.95%)	0.3520(29.10%)	76.65%
	BIM	0.3728(19.24%)	0.3563(28.24%)	75.86%
P3	**FUSE (Full)**	**0.3958(17.20%)**	**0.4233(22.63%)**	**54.85%**
	w/o AFW	0.3972(16.90%)	0.4252(22.28%)	53.75%
	w/o Redist	0.3997(16.38%)	0.4283(21.71%)	51.70%
	w/o Saliency	0.3988(16.57%)	0.4271(21.93%)	52.10%
	BIM	0.4008(16.15%)	0.4301(21.39%)	50.35%
		ETTm1 – DLinear		
P1	**FUSE (Full)**	**0.3043(21.76%)**	**0.2760(23.23%)**	**75.76%**
	w/o AFW	0.3052(21.54%)	0.2770(22.95%)	75.12%
	w/o Redist	0.3067(21.16%)	0.2787(22.48%)	74.10%
	w/o Saliency	0.3059(21.36%)	0.2778(22.73%)	74.62%
	BIM	0.3060(21.34%)	0.2776(22.78%)	74.55%
P2	**FUSE (Full)**	**0.2467(19.11%)**	**0.2268(20.69%)**	**95.57%**
	w/o AFW	0.2474(18.89%)	0.2277(20.38%)	94.91%
	w/o Redist	0.2488(18.43%)	0.2292(19.86%)	93.98%
	w/o Saliency	0.2481(18.66%)	0.2284(20.14%)	94.41%
	BIM	0.2480(18.69%)	0.2282(20.21%)	94.36%
P3	**FUSE (Full)**	**0.3430(18.53%)**	**0.3290(18.16%)**	**64.20%**
	w/o AFW	0.3438(18.34%)	0.3300(17.91%)	63.55%
	w/o Redist	0.3454(17.96%)	0.3318(17.46%)	62.50%
	w/o Saliency	0.3446(18.15%)	0.3309(17.69%)	63.02%
	BIM	0.3449(18.08%)	0.3312(17.61%)	62.91%
		ETTm1 – TimesNet		
P1	**FUSE (Full)**	**0.2510(24.18%)**	**0.1565(34.12%)**	**92.10%**
	w/o AFW	0.2530(23.57%)	0.1591(33.04%)	91.88%
	w/o Redist	0.2548(23.03%)	0.1612(32.15%)	90.74%
	w/o Saliency	0.2541(23.24%)	0.1604(32.49%)	91.18%
	BIM	0.2589(21.79%)	0.1650(30.54%)	90.99%
P2	**FUSE (Full)**	**0.2222(22.84%)**	**0.1466(31.06%)**	**98.12%**
	w/o AFW	0.2243(22.12%)	0.1495(29.70%)	97.74%
	w/o Redist	0.2265(21.35%)	0.1518(28.60%)	96.93%
	w/o Saliency	0.2257(21.63%)	0.1508(29.07%)	97.21%

Continued on next page

Table 8 continued

Pattern	Method	MAE (η%) ↓	MSE (η%) ↓	ASR (%) ↑
	BIM	0.2295(20.31%)	0.1546(27.29%)	97.66%
P3	**FUSE (Full)**	**0.3378(18.92%)**	**0.3096(21.34%)**	**72.84%**
	w/o AFW	0.3401(18.37%)	0.3124(20.63%)	71.52%
	w/o Redist	0.3442(17.38%)	0.3178(19.26%)	68.96%
	w/o Saliency	0.3425(17.79%)	0.3155(19.84%)	70.12%
	BIM	0.3457(17.02%)	0.3196(18.80%)	69.34%

Analysis. Table 8 reports the complete white-box ablation results across two datasets, two forecasting architectures, and three semantic objectives.

(1) The complete FUSE consistently achieves the strongest performance across all evaluated settings, showing that the three components provide complementary benefits. Removing AFW leads to a consistent degradation in both Target Adherence and ASR. For example, on ETTh1 with TimesNet under P1, removing AFW decreases ASR from 81.15% to 80.45% and reduces the MSE adherence from 35.66% to 35.25%. This indicates that adaptive frequency weighting helps the attack allocate spectral emphasis more effectively.

(2) Perturbation redistribution has the largest overall impact among the three components. On ETTm1 with TimesNet under P3, removing redistribution decreases ASR from 72.84% to 68.96%, while MAE adherence drops from 18.92% to 17.38%. Removing temporal saliency also consistently weakens the attack, but the degradation is generally smaller than that caused by removing redistribution. This suggests that redistribution primarily improves perturbation-budget utilization, whereas temporal saliency further determines where the redistributed perturbation should be emphasized.

Overall, AFW, perturbation redistribution, and temporal saliency contribute to spectral adaptation, budget utilization, and temporal allocation, respectively. Their combination yields the most consistent targeted attack performance across datasets, models, and semantic objectives.

E.2 Black-box Transfer Ablation

Table 9: Complete black-box transfer ablation results on **ETTh1** and **ETTm1** with **DLinear** and **TimesNet**. P1/P2/P3 denote spectral reduction, trend reduction, and seasonality enhancement. Best/second-best results are **bold**/underlined.

Pattern	Method	MAE (η%) ↓	MSE (η%) ↓	ASR (%) ↑
		ETTh1 – DLinear		
P1	**FUSE (Full)**	**0.4796(10.14%)**	**0.5104(12.31%)**	**56.17%**
	w/o AFW	0.4801(10.04%)	0.5111(12.19%)	55.82%
	w/o Redist	0.4810(9.87%)	0.5123(11.98%)	55.18%
	w/o Saliency	0.4806(9.95%)	0.5117(12.08%)	55.49%
	BIM	0.4808(9.91%)	0.5119(12.05%)	55.36%
P2	**FUSE (Full)**	**0.4949(9.64%)**	**0.5543(11.92%)**	**58.22%**
	w/o AFW	0.4954(9.54%)	0.5550(11.80%)	57.88%
	w/o Redist	0.4960(9.42%)	0.5558(11.65%)	57.35%
	w/o Saliency	0.4957(9.48%)	0.5554(11.72%)	57.61%
	BIM	0.4959(9.44%)	0.5556(11.68%)	57.48%
P3	**FUSE (Full)**	**0.5256(8.92%)**	**0.6301(10.74%)**	**23.13%**
	w/o AFW	0.5260(8.84%)	0.6308(10.63%)	22.78%
	w/o Redist	0.5267(8.72%)	0.6317(10.48%)	22.18%
	w/o Saliency	0.5263(8.79%)	0.6312(10.56%)	22.49%
	BIM	0.5265(8.75%)	0.6314(10.53%)	22.36%
		ETTh1 – TimesNet		
P1	**FUSE (Full)**	**0.5054(9.72%)**	**0.5357(11.88%)**	**44.76%**
	w/o AFW	0.5060(9.61%)	0.5365(11.75%)	44.28%
	w/o Redist	0.5071(9.41%)	0.5379(11.52%)	43.46%
	w/o Saliency	0.5066(9.50%)	0.5372(11.63%)	43.89%
	BIM	0.5069(9.45%)	0.5376(11.57%)	43.71%
P2	**FUSE (Full)**	**0.5339(8.86%)**	**0.6110(10.94%)**	**32.01%**
	w/o AFW	0.5345(8.75%)	0.6118(10.82%)	31.58%
	w/o Redist	0.5354(8.59%)	0.6130(10.64%)	30.82%
	w/o Saliency	0.5349(8.68%)	0.6124(10.73%)	31.21%

Continued on next page

Table 9 continued

Pattern	Method	MAE (η%) ↓	MSE (η%) ↓	ASR (%) ↑
	BIM	0.5352(8.63%)	0.6127(10.68%)	31.05%
P3	**FUSE (Full)**	**0.5496(7.86%)**	**0.6593(9.48%)**	**2.74%**
	w/o AFW	0.5500(7.79%)	0.6599(9.39%)	2.42%
	w/o Redist	0.5508(7.66%)	0.6610(9.24%)	1.86%
	w/o Saliency	0.5504(7.72%)	0.6605(9.31%)	2.12%
	BIM	0.5506(7.69%)	0.6608(9.27%)	1.98%
	ETTm1 – DLinear			
P1	**FUSE (Full)**	**0.3829(9.21%)**	**0.3415(11.08%)**	**48.32%**
	w/o AFW	0.3832(9.14%)	0.3419(10.98%)	47.98%
	w/o Redist	0.3838(9.00%)	0.3426(10.80%)	47.39%
	w/o Saliency	0.3835(9.07%)	0.3422(10.90%)	47.68%
	BIM	0.3836(9.05%)	0.3424(10.85%)	47.55%
P2	**FUSE (Full)**	**0.3437(9.36%)**	**0.3009(11.24%)**	**75.02%**
	w/o AFW	0.3440(9.28%)	0.3013(11.12%)	74.71%
	w/o Redist	0.3446(9.12%)	0.3020(10.91%)	74.18%
	w/o Saliency	0.3443(9.20%)	0.3017(11.00%)	74.43%
	BIM	0.3444(9.17%)	0.3018(10.97%)	74.34%
P3	**FUSE (Full)**	**0.3796(8.74%)**	**0.3497(10.18%)**	**17.02%**
	w/o AFW	0.3800(8.65%)	0.3502(10.05%)	16.68%
	w/o Redist	0.3808(8.45%)	0.3511(9.82%)	16.02%
	w/o Saliency	0.3804(8.55%)	0.3506(9.95%)	16.35%
	BIM	0.3806(8.50%)	0.3509(9.88%)	16.21%
	ETTm1 – TimesNet			
P1	**FUSE (Full)**	**0.4137(8.76%)**	**0.3723(10.64%)**	**36.18%**
	w/o AFW	0.4142(8.65%)	0.3730(10.47%)	35.66%
	w/o Redist	0.4152(8.43%)	0.3742(10.18%)	34.84%
	w/o Saliency	0.4147(8.54%)	0.3736(10.33%)	35.25%
	BIM	0.4150(8.47%)	0.3739(10.25%)	35.03%
P2	**FUSE (Full)**	**0.3742(8.91%)**	**0.3297(10.73%)**	**57.82%**
	w/o AFW	0.3748(8.76%)	0.3305(10.51%)	57.34%
	w/o Redist	0.3758(8.52%)	0.3317(10.18%)	56.49%
	w/o Saliency	0.3753(8.64%)	0.3311(10.35%)	56.91%
	BIM	0.3756(8.57%)	0.3314(10.27%)	56.70%
P3	**FUSE (Full)**	**0.3761(8.31%)**	**0.3504(9.76%)**	**17.28%**
	w/o AFW	0.3766(8.19%)	0.3510(9.61%)	16.81%
	w/o Redist	0.3775(7.97%)	0.3521(9.33%)	15.94%
	w/o Saliency	0.3770(8.09%)	0.3515(9.48%)	16.37%
	BIM	0.3773(8.02%)	0.3518(9.41%)	16.18%

Analysis. Table 9 reports the complete black-box transfer ablation results across two datasets, two forecasting architectures, and three semantic objectives.

(1) The complete FUSE consistently achieves the strongest transfer performance across all evaluated settings, indicating that the three components remain complementary under cross-model transfer. Removing AFW causes a consistent reduction in Target Adherence and ASR. For example, on ETTh1 with DLinear under P1, removing AFW decreases ASR from 56.17% to 55.82%, while MSE adherence drops from 12.31% to 12.19%. This suggests that adaptive spectral weighting also contributes to transferability by emphasizing frequency components that remain effective across different forecasting models.

(2) Perturbation redistribution exhibits the largest overall impact on black-box transfer performance. On ETTm1 with TimesNet under P3, removing redistribution decreases ASR from 17.28% to 15.94%, while MAE adherence decreases from 8.31% to 7.97%. Removing temporal saliency also consistently reduces transfer performance, but its effect is generally smaller than removing redistribution. These observations indicate that effective reuse of the perturbation budget is particularly important for cross-model transfer, while temporal saliency further improves the placement of the redistributed perturbations.

Overall, the black-box ablation results are consistent with the white-box observations: AFW, perturbation redistribution, and temporal saliency contribute to spectral adaptation, budget utilization, and temporal allocation, respectively. Their combination yields the most stable transfer performance across datasets, models, and semantic objectives.

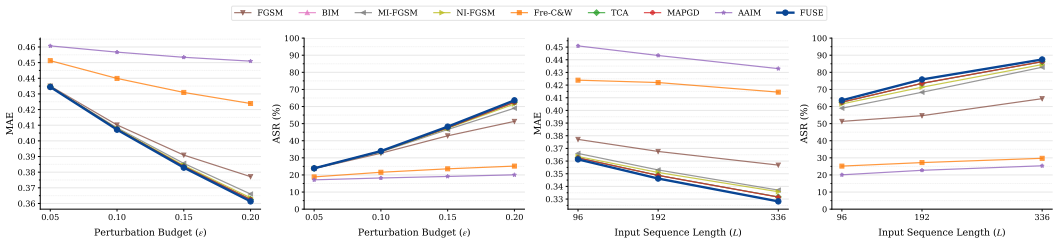

Figure 4: Sensitivity analysis of FUSE and baselines on the ETTh1 dataset. We evaluate the impact of the perturbation budget $\epsilon \in \{0.05, 0.1, 0.15, 0.2\}$ and the input sequence length $L \in \{96, 192, 336\}$ on MAE and ASR. Lower MAE and higher ASR indicate stronger targeted attack performance.

F Hyper-parameter Analysis

We further investigate the sensitivity of FUSE to two key experimental settings: the perturbation budget ϵ and the input sequence length L.

Effect of Perturbation Budget. As shown in Figure 4, increasing the perturbation budget from $\epsilon = 0.05$ to $\epsilon = 0.20$ generally improves the attack effectiveness of all methods, resulting in higher ASR and lower MAE. FUSE consistently maintains strong performance across different perturbation budgets and achieves more than 60% ASR when $\epsilon = 0.20$. A larger perturbation budget provides greater optimization flexibility for adaptive spectral weighting and temporal perturbation allocation, allowing the attack to more effectively steer predictions toward the target pattern.

Effect of Input Sequence Length. We further examine the influence of the look-back window $L \in \{96, 192, 336\}$. As the input sequence becomes longer, ASR generally increases while MAE decreases. Longer input sequences provide richer temporal and spectral information, enabling frequency–temporal optimization to exploit a broader context when constructing adversarial perturbations. These results indicate that FUSE remains effective under different input lengths and can benefit from longer historical contexts.