

Privacy on the Fly: A Predictive Adversarial Transformation Network for Mobile Sensor Data

Tianle Song¹, Chenhao Lin^{1*}, Yang Cao², Zhengyu Zhao¹, Jiahao Sun¹, Chong Zhang¹, Le Yang¹, and Chao Shen¹

¹Xi'an Jiaotong University, China

²Institute of Science Tokyo, Japan

{tianlesong, sunjiahao}@xjtu.stu.edu.cn, {linchenhao, zhengyu.zhao, zhangchong, yangle15, chaoshen}@xjtu.edu.cn, cao@c.titech.ac.jp

Abstract

Mobile motion sensors such as accelerometers and gyroscopes are now ubiquitously accessible by third-party apps via standard APIs. While enabling rich functionalities like activity recognition and step counting, this openness has also enabled unregulated inference of sensitive user traits, such as gender, age, and even identity, without user consent. Existing privacy-preserving techniques, such as GAN-based obfuscation or differential privacy, typically require access to the full input sequence, introducing latency that is incompatible with real-time scenarios. Worse, they tend to distort temporal and semantic patterns, degrading the utility of the data for benign tasks like activity recognition. To address these limitations, we propose the Predictive Adversarial Transformation Network (PATN), a real-time privacy-preserving framework that leverages historical signals to generate adversarial perturbations proactively. The perturbations are applied immediately upon data acquisition, enabling continuous protection without disrupting application functionality. Experiments on two datasets demonstrate that PATN substantially degrades the performance of privacy inference models, achieving Attack Success Rate (ASR) of 40.11% and 44.65% (reducing inference accuracy to near-random) and increasing the Equal Error Rate (EER) from 8.30% and 7.56% to 41.65% and 46.22%. On ASR, PATN outperforms baseline methods by 16.16% and 31.96%, respectively.

Introduction

Mobile applications can easily access motion sensor data (e.g., accelerometer and gyroscope) through APIs on Android and iOS platforms (Rajguru 2019). This supports various functions like activity recognition (Wang et al. 2019), step counting (Susi, Renaudin, and Lachapelle 2013), and gesture interaction (Brotchie et al. 2022), making motion data central to mobile services. However, this widespread access also raises privacy concerns, as sensor data can reveal personal attributes like identity (Cai et al. 2024; Ferdoooni et al. 2023; Shen et al. 2017), gender (Meena and Sarawadekar 2020), and age (Lin et al. 2023; Miao et al. 2023). Moreover, the availability of open-source privacy inference models (Malekzadeh 2019) makes it easier for third-party applications to exploit such data, potentially infringing

*Corresponding author.

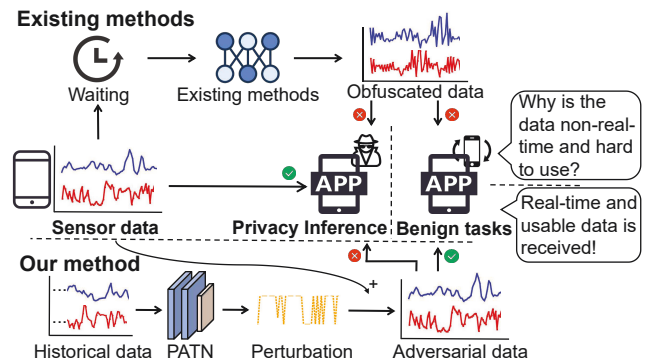


Figure 1: Our method vs. existing methods, addressing two issues: (1) temporal semantic distortion and (2) delayed obfuscation due to segment-wise processing.

on user privacy without their knowledge. These concerns underscore the need for mechanisms that can safeguard user privacy without compromising the utility of motion sensor data.

Existing privacy-preserving approaches, including differential privacy (Kalupahana et al. 2023) and generative model-based obfuscation (Raval, Machanavajjhala, and Pan 2019; Gu et al. 2025), often optimize for specific tasks to preserve utility while preventing private information inference. However, as shown in Figure 1, these methods face two key limitations: (1) Generative methods often compromise temporal semantics by regenerating entire sequences via latent-space sampling, over-smoothing or distorting fine-grained patterns, and degrading utility for tasks requiring precise numerical computations (e.g., smartphone rotation angle estimation). (2) Most cannot satisfy real-time demands, as they buffer complete sensor sequences before transformation, whereas real-world streams must be processed instantly. These challenges highlight the need for methods that protect sensitive information while maintaining both semantic fidelity and real-time usability.

Adversarial perturbations exploit model sensitivity by adding imperceptible noise that induces incorrect outputs. We adopt them as a lightweight way to suppress sensitive patterns while preserving sensor data semantics. Although adversarial attacks on time-series models (Yang et al. 2022;

Pialla et al. 2025) are well studied, most rely on full input segments, making them unsuitable for streaming and unable to address limitation (2). Simply applying perturbations generated from historical data directly to future inputs has limited effectiveness, as it may fail to align with the temporal dynamics of the target sequence. Universal adversarial perturbations (Rathore et al. 2020) often perform poorly on complex, high-dimensional sensor data, further limiting their practicality.

To address these issues, we propose the *Predictive Adversarial Transformation Network* (PATN), which leverages historical sensor data to generate adversarial perturbations for future readings. This enables real-time, zero-latency protection with high consistency to original signals. PATN combines a generative network with a *history-aware top-k optimization* strategy to mitigate temporal misalignment in inference attacks. Experiments on two datasets show PATN achieves 40.11% and 44.65% Attack Success Rate (ASR), reduces Equal Error Rate (EER) to 41.65% and 46.22%, and preserves semantic fidelity for non-sensitive applications. The implementation of PATN is publicly available at <https://github.com/skysky4/PATN>. Our contributions are summarized as follows,

- We propose PATN, the first framework to enable proactive, real-time privacy protection for streaming sensor data by generating future-directed adversarial perturbations, and the first to introduce adversarial perturbations into sensor data privacy protection.
- We introduce a history-aware top- k optimization strategy that effectively counters performance degradation under temporal misalignment attacks.
- Through extensive evaluation, PATN is shown to substantially enhance privacy protection while retaining the usability and semantic fidelity of motion sensor signals.

Related Works

Data Obfuscation. Privacy-preserving methods for sensor data fall into three categories. Differential privacy (DP) (Kalupahana et al. 2023) adds noise to protect privacy but distorts temporal signals, reducing utility in continuous tasks. Privacy-aware feature extraction (Liu et al. 2019; Li et al. 2021) learns representations that retain task-relevant information while masking sensitive attributes, often via adversarial training. Recent generative approaches, such as GANs (Boutet et al. 2021), VAEs (Hajihassnai, Ardakanian, and Khzaei 2021), and diffusion models (Yang and Ardakanian 2023), produce realistic sensor data but often face semantic drift, affecting fine-grained temporal patterns. For example, DoppelGANger (Lin et al. 2020) captures long-term trends but struggles with accelerometer dynamics, while PrivDiffuser (Yang and Ardakanian 2025) improves privacy-utility trade-offs but may distort temporally sensitive signals. Though generative methods obfuscate private information well, they risk degrading signal semantics, limiting use in sensor-driven applications. Additionally, existing methods cannot meet real-time demands, as they buffer entire sequences before transformation.

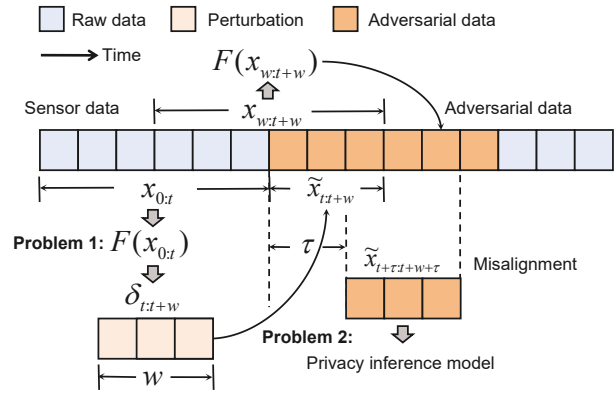


Figure 2: Illustrating the two problems: generating real-time perturbations for continuously arriving sensor data, and addressing temporal misalignment where attacks may occur at arbitrary, unpredictable time points.

Adversarial Attacks. Adversarial attacks on time series have been extensively studied. Early work examined gradient-based attacks, including FGSM (Fawaz et al. 2019; Oregi et al. 2018) and PGD on deep time-series classifiers. Comprehensive surveys such as Adversarial Attacks on Time Series provide taxonomies of these methods and their impacts across various domains (Karim, Majumdar, and Darabi 2020). Özcan et al. introduced AdaptEdge (Khan, Mynuddin, and Nabil 2024), a targeted universal attack framework that demonstrated cross-dataset transferability in smart grid time series. Recent methods, such as SFAttack (Gu et al. 2025) and BlackTreeS (Ding et al. 2023), targeting imperceptible adversarial attacks for time series classification, improve stealth by utilizing localized and frequency-domain perturbations. While these studies highlight the vulnerability of time-series models, they uniformly assume access to complete time-series segments before perturbation, which limits their applicability to streaming sensor data that must be protected in real-time.

Problem Setting

Our objective is to protect user privacy by adding adversarial perturbations to raw sensor data while preserving the semantic integrity of the original signals. Based on this objective, we identify two key problems.

Problem 1: Real-Time Perturbation Generation. The problem we need to solve is to generate real-time, future-directed adversarial perturbations to protect user privacy from inference attacks on motion sensor data. Existing methods typically require access to complete sensor sequences before adding perturbations or reconstructing data, making them incompatible with real-time applications where sensor data arrive continuously and must be processed on the fly.

As shown in Figure 2, we learn a temporal mapping function \mathcal{F} that takes historical motion data $x_{0:t}$ as input and outputs a sequence of perturbations for the upcoming steps:

$$\delta_{t:t+w} = \mathcal{F}(x_{0:t}) \quad (1)$$

The predicted perturbations $\delta_{t:t+w}$ are then applied in

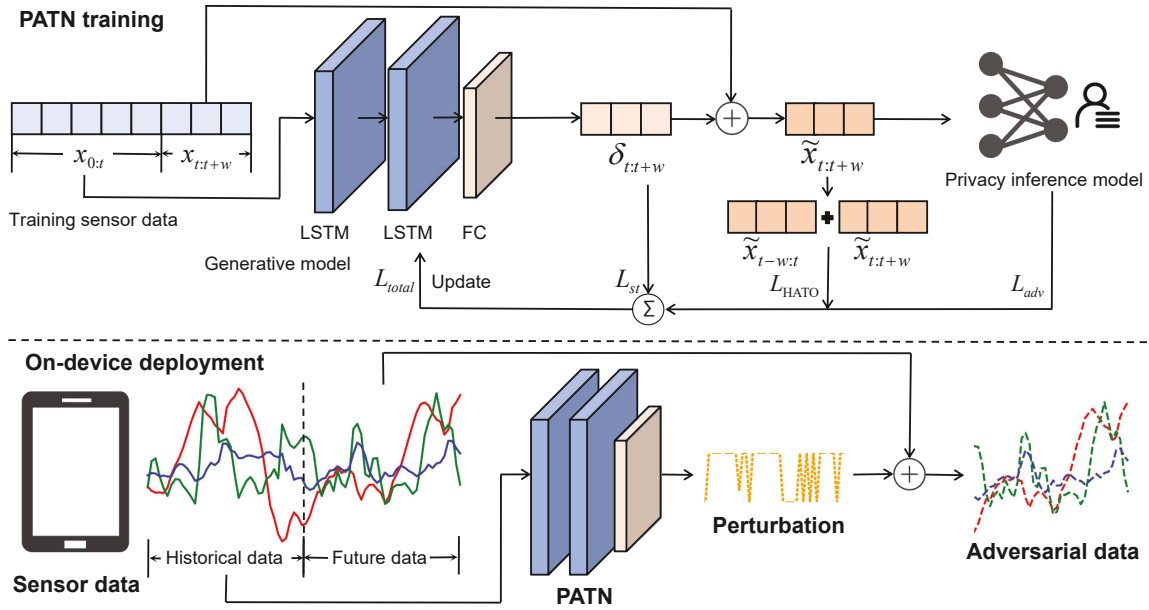


Figure 3: Overview of the PATN framework. The system includes PATN training, where the network learns to generate privacy-preserving perturbations, and on-device deployment, where the trained model runs securely in real time.

real time to the incoming data x_t, x_{t+1}, \dots , ensuring that privacy-sensitive patterns are obfuscated before being accessed by untrusted applications.

Problem 2: Temporal Misalignment Between Defense and Attack. The second problem is ensuring defense effectiveness under temporal misalignment (Figure 2). In practice, adversaries may launch inference attacks at arbitrary time points, while existing defenses generate perturbations aligned to fixed timelines, reducing their effectiveness. Let $\delta_{t:t+w}$ denote the perturbation over window $[t, t+w]$, and suppose the attacker targets a shifted window $\tilde{x}_{t+\tau:t+\tau+w}$ with unknown offset τ , such that $\tilde{x}_{t:t+w} \neq \tilde{x}_{t+\tau:t+\tau+w}$. This misalignment causes partial overlap between the perturbation and attack windows, weakening privacy protection.

Methodology

Overview

To protect user privacy while preserving sensor data semantics, we propose the PATN framework, as shown in Figure 3. Assuming access to open-source privacy inference models and their gradients (white-box), PATN leverages historical sensor data to predict future-directed perturbations, optimized to mislead privacy inference models with minimal semantic distortion. The framework has two components: (1) *PATN training*, where PATN is optimized using three objectives—adversarial effectiveness, temporal robustness, and smoothness regularization—to balance privacy protection and data fidelity; and (2) *On-device deployment*, where the trained network is executed locally on mobile devices to enable zero-latency perturbation generation for real-time sensor streams.

Predictive Adversarial Transformation Network

Perturbation Range Setting. To prevent adversarial perturbations from degrading the semantic integrity of sensor data or disrupting normal system functionality, we constrain their magnitude with carefully designed bounds based on ℓ_∞ normalization. Excessively small perturbations risk ineffective attacks, while overly large ones may distort essential patterns.

We first derive a dimension-wise upper bound based on dataset statistics by computing the mean and standard deviation for each sensor dimension and defining the perturbation limit as 0.05 of these values:

$$\epsilon_d^{\text{data}} = \min(0.05 \times \mu_d, 0.05 \times \sigma_d) \quad (2)$$

where μ_d and σ_d denote the mean and standard deviation of the d -th dimension.

To further ensure imperceptibility, we measure natural sensor variation under static conditions, where 10 users interacted with smartphones fixed on a rigid table. The standard deviation of each feature dimension in this scenario provides $\epsilon_d^{\text{natural}}$, which we combine with the statistical bound to define the final ℓ_∞ constraint:

$$\epsilon_d = \min(\epsilon_d^{\text{data}}, \epsilon_d^{\text{natural}}) \quad (3)$$

This guarantees that perturbations remain within naturally occurring fluctuations, preserving utility while ensuring privacy.

History-Driven Generative Model for Perturbation. To address the challenge of real-time perturbation generation, we propose PATN, a generative framework that produces future adversarial perturbations based solely on historical multivariate sensor data. Unlike conventional adversarial methods that require future observations, PATN forecasts perturbations from past system sensor data, making it suitable for

Algorithm 1: History-Aware Top-k Optimization

```

1: Input: Previous perturbation  $\delta_{t-w:t}$ , current perturbation  $\delta_{t:t+w}$ , input  $x$ , model  $f$ , label  $y_{\text{target}}$ , window size  $w$ , step size  $s$ , top-k  $k$ 
2: Output: Top-k averaged loss
3: Merge perturbations:  $\delta \leftarrow \text{Merge}(\delta_{t-w:t}, \delta_{t:t+w})$ 
4: Generate adversarial input:  $\tilde{x} \leftarrow x + \delta$ 
5: Initialize loss list:  $\mathcal{L} \leftarrow []$ 
6: for each  $t \in \{0, s, 2s, \dots, T - w\}$  do
7:   Extract window:  $\tilde{x}_t \leftarrow \tilde{x}[:, t : t + w, :]$ 
8:   Compute prediction:  $\hat{z}_t \leftarrow f(\tilde{x}_t)$ 
9:   Compute loss:  $\ell_t \leftarrow \text{CrossEntropy}(\hat{z}_t, y_{\text{target}})$ 
10:  Append loss:  $\mathcal{L}.\text{append}(\ell_t)$ 
11: end for
12: Stack all losses:  $\ell_{\text{stack}} \leftarrow \text{stack}(\mathcal{L})$ 
13: Select top- $k$  losses:  $\ell_{\text{topk}} \leftarrow \text{TopK}(\ell_{\text{stack}}, k)$ 
14: Compute final loss:  $\mathcal{L}_{\text{HATO}} \leftarrow \text{mean}(\ell_{\text{topk}})$ 
15: return  $\mathcal{L}_{\text{HATO}}$ 

```

real-world time-series scenarios, where future data is unavailable at inference time.

PATN utilizes a sequence-to-sequence architecture with an LSTM-based encoder-decoder framework. The encoder processes a multivariate input sequence of length T_{in} , consisting of D -dimensional sensor readings, denoted as $x \in \mathbb{R}^{B \times T_{\text{in}} \times D}$. The LSTM encoder extracts temporal dependencies and condenses the sequence into a fixed-length latent representation, typically the final hidden and cell states. The decoder, another LSTM, autoregressively generates a sequence of adversarial perturbations of length T_{out} , with each step producing a perturbation vector $\delta_i \in \mathbb{R}^D$.

$$\delta_i = W_o h_i + b_o \quad (4)$$

Here, $W_o \in \mathbb{R}^{D \times H}$ and $b_o \in \mathbb{R}^D$ are learnable parameters, and H denotes the hidden dimension of the decoder LSTM. The decoder operates in an autoregressive manner, meaning the output from each previous step can be fed back as input into the next, allowing the network to capture temporal consistency in the generated perturbation sequence. The complete adversarial sequence is defined as:

$$\delta = [\delta_1, \delta_2, \dots, \delta_{T_{\text{out}}}] \in \mathbb{R}^{B \times T_{\text{out}} \times D}, \|\delta\|_{\infty} \leq \epsilon_d \quad (5)$$

This predicted perturbation tensor is subsequently added to the raw future sensor data to construct adversarial examples:

$$\tilde{x}_{\text{adv}} = x_{\text{raw}} + \delta \quad (6)$$

where $x_{\text{raw}} \in \mathbb{R}^{B \times T_{\text{out}} \times D}$ denotes the future clean data sequence and \tilde{x}_{adv} is the resulting adversarial input.

History-Aware Top-k Optimization. To enhance the robustness of adversarial perturbations against Problem 2, we propose history-aware top-k optimization (HATO). This method leverages both the previously generated perturbation (the last perturbation segment) and the current perturbation, aiming to construct a more temporally consistent and generalizable perturbation sequence.

As shown in Algorithm 1, we first concatenate the previous perturbation δ_{prev} with the current perturbation δ_{cur} to

form a longer adversarial perturbation. This merged perturbation is then applied to the clean input sequence. A sliding window with a fixed size and stride is used to extract overlapping temporal segments. For each segment, the privacy inference model f computes a cross-entropy loss with a misleading target label to simulate a privacy-preserving adversarial objective. To address the instability of loss values across different window positions, we adopt a top- k selection strategy. Among the computed losses, we select the k highest values and use their average as the final optimization target.

$$\mathcal{L}_{\text{HATO}} = \frac{1}{k} \sum_{i=1}^k \text{TopK}_i(\mathcal{L}, k) \quad (7)$$

This encourages the perturbation to consistently degrade model performance across multiple subwindows, rather than overfitting to a specific temporal slice. This history-aware top-k loss not only improves the temporal robustness of the perturbation but also enhances its defensive effect under arbitrary window-based evaluation, which is common in real-world privacy-preserving scenarios.

Objective Function. To optimize the PATN model, we define a composite loss function that integrates three complementary objectives.

1. **Misclassification loss \mathcal{L}_{adv} :** This term encourages the adversarial input to mislead the target model (Privacy inference model). It is implemented as a standard cross-entropy loss between the model’s prediction and a misleading target:

$$\mathcal{L}_{\text{adv}} = \mathcal{L}_{\text{CE}}(\text{logits}_{\text{adv}}, y_{\text{target}}) \quad (8)$$

where y_{target} denotes a misleading label, which in a binary classification task corresponds to the opposite of the label output by the privacy inference model.

2. **History-aware top- k loss $\mathcal{L}_{\text{HATO}}$:** As detailed in Equation 7, this loss promotes temporal robustness by encouraging consistent model outputs across a sliding inference window. This helps the perturbation remain effective even under windowed decision-making.

3. **Smoothness loss \mathcal{L}_{st} :** This regularization term penalizes abrupt changes in the perturbation signal over time, thereby improving the visual and temporal coherence of the generated adversarial sequence. It is implemented as the mean squared error (MSE) of the perturbation:

$$\mathcal{L}_{\text{st}} = \text{MSE}(\delta) \quad (9)$$

The total objective is a weighted sum:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{adv}} + \lambda_1 \mathcal{L}_{\text{HATO}} + \lambda_2 \mathcal{L}_{\text{st}} \quad (10)$$

with $\lambda_1 = 0.3$, $\lambda_2 = 0.3$. These weights are chosen to balance attack effectiveness, temporal robustness, and perceptual imperceptibility. When confronted with multiple privacy inference models, we aggregate the $\mathcal{L}_{\text{total}}$ from each model into a unified optimization objective, enabling the method to effectively defend against a diverse set of attacks simultaneously.

Method	MotionSense				ChildShield			
	ASR(%) \uparrow	EER(%) \uparrow	AUC \downarrow	F1_Score \downarrow	ASR(%) \uparrow	EER(%) \uparrow	AUC \downarrow	F1_Score \downarrow
Raw data		8.30	0.979	0.928		7.56	0.972	0.922
DP	14.37	17.46	0.931	0.836	4.12	12.00	0.952	0.887
UAP	9.61	13.53	0.951	0.863	3.17	10.92	0.957	0.894
FGSM	23.95	25.92	0.772	0.802	12.99	19.11	0.889	0.821
PGD	23.95	25.92	0.772	0.802	12.99	19.11	0.889	0.821
PATN(ours)	40.11	41.65	0.662	0.611	44.95	46.22	0.549	0.537

Table 1: Performance comparison with baseline methods.

Experiments

In this section, we describe the experimental setup and evaluate our method via baseline comparisons, utility analysis, ablation studies, and black-box migration tests to assess performance, component contributions, and robustness.

Implementation Details

Datasets. We evaluate our method on two real-world datasets: MotionSense (Malekzadeh et al. 2018) and ChildShield (Lin et al. 2023). Both provide accelerometer and gyroscope data, split 7:3 for training and testing.

MotionSense contains 50Hz motion data from 24 participants using an iPhone 6s during six activities (walking, jogging, sitting, standing, upstairs, and downstairs), enabling gender (male/female) inference. ChildShield includes 60Hz motion data from 1,875 participants across five mobile games, used for age (child/adult) inference.

Evaluation Metric. To evaluate privacy protection, we define the attack success rate (ASR) as the proportion of originally correct predictions by the privacy inference model that are misclassified after applying our perturbation. A higher ASR indicates stronger privacy protection. We also report Equal Error Rate (EER), Area Under the Curve (AUC), and F1_Score to assess the overall impact on adversarial inference performance. For utility evaluation, we adopt two perspectives. First, we measure semantic consistency using ℓ_2 distance, Dynamic Time Warping (DTW), low-frequency (LF) component, and Root Mean Square Error (RMSE) (Wu et al. 2022), where lower values indicate better preservation of signal characteristics. Second, we assess the impact on downstream tasks, specifically step detection (Susi, Renaudin, and Lachapelle 2013) and human activity recognition (HAR) (Mekruksavanich and Jitpattanakul 2021), to quantify the effect of perturbations on typical usage scenarios.

Privacy Inference Models and Baseline Methods. Most privacy inference models for sensor data adopt CNN architectures. In this work, we examine two representative CNN-based models for extracting sensitive user information, following the architectures of Lin et al. (Lin et al. 2023) (age) and Sharshar et al. (Sharshar et al. 2021) (gender). We also evaluate PATN on sequential models (RNN, LSTM), with results reported in the Appendix (Song 2025).

We compare our approach with representative baselines under the same ℓ_∞ -norm constraint. Data-independent methods include differential privacy (DP) (Kalupahana

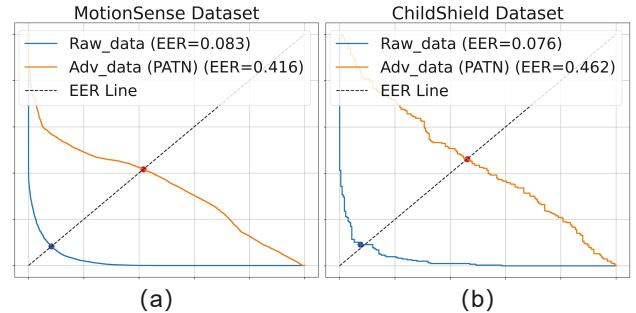


Figure 4: The AUC curve of adversarial data generated by PATN compared to the raw data when applied to the privacy inference model.

et al. 2023) and universal adversarial perturbations (UAP) (Rathore et al. 2020), while data-dependent methods include FGSM (Fawaz et al. 2019) and PGD (Oregi et al. 2018), both leveraging historical sensor data. Generative model-based approaches (Yang and Ardakanian 2025) are excluded from direct comparison, as they severely violate the ℓ_∞ constraint and require full sensor sequences at the current time step, which limits real-time applicability. Nonetheless, we qualitatively compare with PrivDiffuser (Yang and Ardakanian 2025), showing that our method better preserves semantic consistency in perturbed sensor signals (in section : Utility Analysis).

Parameters. The Predictive Adversarial Transformation Network (PATN) takes an input sequence of length $T_{in} = 30$ (each representing 0.5-second intervals of sensor data) and generates an output sequence of length $T_{out} = 10$, corresponding to the input length for the privacy inference model. The input dimension is 6, representing six sensor features, and the LSTM hidden dimension is 64. HATO is applied with a top-k parameter of $k=2$. Training and testing are performed on an online server with an NVIDIA 3090Ti GPU using PyTorch. The Adam optimizer is used, with an initial learning rate of $1e-3$ and a fixed-step decay schedule (halving every 200 epochs), over a total of 600 epochs.

Main Results

Comparison with Baseline Methods. As shown in Table 1 and Figure 4, under the same ℓ_∞ -norm constraint, our proposed PATN achieves superior privacy protection by generating adversarial perturbations that effectively mislead pri-

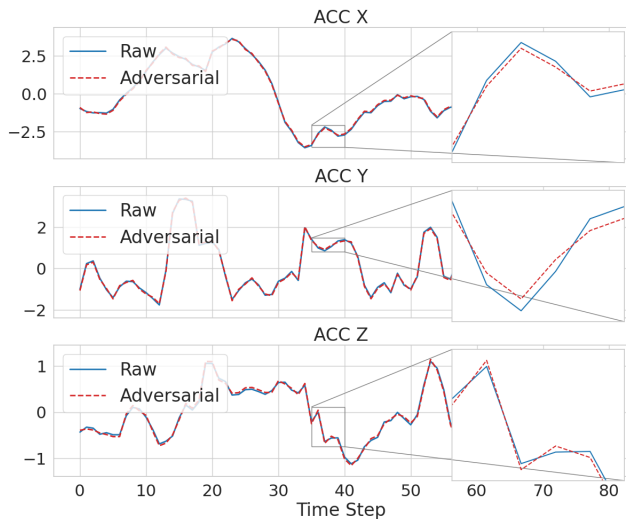


Figure 5: Comparison of Raw and Adversarial sensor data (accelerometer) on IMU Axes.

privacy inference models, targeting gender inference on MotionSense and age inference on ChildShield.

In contrast to traditional approaches like DP and UAP, PATN is dynamic and adaptive, optimizing perturbations specifically to protect privacy. DP and UAP rely on fixed, data-agnostic perturbation strategies that overlook the vulnerabilities of historical data distributions and fail to target real-time or future private attribute inference. PATN, by introducing a predictive adversarial transformation, tailors perturbations to the model’s current predictions and adapts them for future data, making them more effective in confusing privacy inference models. Consequently, PATN achieves substantially higher ASR 40.11% on MotionSense and 44.95% on ChildShield compared to DP (14.37% and 4.12%) and UAP (9.61% and 3.17%).

While FGSM and PGD craft perturbations based on historical data distributions, they lack optimization for future data (i.e., unseen data), limiting their effectiveness in dynamic environments. These methods assume minimal temporal variation in sensor data, which constrains their robustness. As shown in Table 1, FGSM and PGD achieve ASRs of 23.95% on MotionSense and 12.99% on ChildShield, indicating partial success in misleading privacy inference models but without future adaptability. In contrast, PATN explicitly optimizes perturbations for future data, enhancing robustness and privacy protection. It achieves significantly higher ASRs of 40.11% and 44.95% on MotionSense and ChildShield, respectively, demonstrating superior performance in scenarios with evolving data distributions.

Utility Analysis. Our PATN model has a compact size of only 0.365 MB, well within the 2 MB memory constraint for deployment in the TEE. The perturbation generation time is just 0.00036 seconds, much faster than the 1/60-second sensor sampling interval, ensuring real-time applicability.

To assess semantic fidelity, we visualize the raw and perturbed accelerometer signals along individual IMU axes

	DTW↓	ℓ_2 ↓	LF↓	RMSE↓
PATN	0.744	0.162	0.300	0.037
PrivDiffuser	7.058	2.251	3.422	0.503

Table 2: Comparison of Semantic Consistency Between PATN and PrivDiffuser

	Raw data	PrivDiffuser	PATN
Step detection	7916	8683 (+767)	7937 (+21)
HAR-EER(%) ¹	5.08	6.92 (+1.84)	6.57 (+1.49)
HAR-AUC ²	0.987	0.976 (-0.011)	0.981 (-0.006)

¹ HAR-EER(%): human activity recognition (Macro-EER)

² HAR-AUC: human activity recognition (Macro-AUC)

Table 3: Comparison of Raw Data, PrivDiffuser and PATN-perturbed Data in Step Detection and HAR

Model	CNN	ResNet	DenseNet
ASR(%)	36.98	38.09	38.64
Raw data EER(%)	8.30	7.36	7.85
Adv data EER(%)	39.64	38.66	40.20

Table 4: Evaluation of PATN’s effectiveness across multiple privacy inference models.

(Figure 5, gyroscope shown in Appendix (Song 2025)). The adversarial signals show only subtle deviations from the original, indicating minimal impact on the sensor’s temporal semantics, which is crucial for downstream tasks like step detection. We also compute the per-segment difference between raw and adversarial signals in 10-second segments. Table 2 compares PATN with PrivDiffuser across four metrics. PATN achieves a DTW of 0.744, ℓ_2 distance of 0.162, LF difference of 0.300, and RMSE of 0.037, all significantly lower than PrivDiffuser, confirming that PATN better preserves signal fidelity while obfuscating sensitive information.

We conduct experiments on the MotionSense dataset, focusing on walking data for step detection and multi-class human activity recognition (HAR). As shown in Table 3, our method applies subtle ℓ_∞ -bounded perturbations that preserve sensor utility. In step detection, PATN leads to only a negligible increase of 21 steps, showing minimal impact on usage. For HAR, it introduces a small change in EER (+1.49%) and AUC (−0.006), indicating real-time activity inference remains viable.

In contrast, PrivDiffuser, although optimized for HAR task (EER: 6.92%, AUC: 0.976), shows poor generalization to other benign tasks, causing a large step count error (+767 steps). This highlights PATN’s task-agnostic design, achieving a better balance between privacy and fidelity without task-specific tuning.

Multiple Privacy Inference Models. In real-world scenarios, multiple privacy inference models may simultaneously threaten user privacy, so a method like PATN must remain robust against diverse attack strategies. We evaluate PATN on the MotionSense dataset (used for all subsequent

Length T_{in}	10	20	30	40	50
ASR(%)	34.59	37.54	40.11	38.81	30.88
EER(%)	33.34	37.84	41.65	38.87	29.11

Table 5: ASR and EER of different input lengths T_{in} of PATN. We choose $T_{in} = 30$ as our default setting.

Attack Method	Alignment PATN	Misalignment	
		wHATO ¹	w/oHATO ²
ASR(%)	40.11	39.43	30.56
EER(%)	41.65	40.98	33.24

¹ wHATO: with history-aware top-k optimization

² w/oHATO: without history-aware top-k optimization

Table 6: Comparison of ASR and EER with and without History-Aware Top-k Optimization under Misalignment Attacks.

experiments) against three CNN architectures: a standard CNN, a ResNet, and a DenseNet. PATN is jointly optimized against all three models during training to ensure comprehensive defense.

The results, summarized in Table 4, show that PATN effectively mitigates attacks across these architectures. Specifically, PATN achieves ASR from 36.98% to 38.64% and EER between 38.64% and 40.20%, demonstrating consistent defense capability across different privacy inference models. These findings highlight PATN’s versatility and robustness in protecting sensor data privacy.

Ablation Studies

The Impact of Different Input Lengths T_{in} of PATN. As shown in Table 5, PATN’s performance is significantly influenced by the input length T_{in} . As the input length increases, PATN captures a longer history of sensor data, improving perturbation accuracy. The optimal performance is observed at an input length of 30, achieving 40.11% ASR and 41.65% EER. However, while longer input lengths allow PATN to leverage more historical data (each unit representing 0.5 seconds of data), excessively long inputs lead to diminishing returns. For input lengths of 40 and 50, performance degrades, suggesting that processing excessive or redundant historical data may dilute the quality of generated perturbations.

The Impact of History-Aware Top-k Optimization. As shown in Table 6, when subjected to a Misalignment attack with a sliding step length of one second, PATN with HATO achieves 39.43% ASR and 40.98% EER. These results demonstrate HATO’s effectiveness in addressing Problem 2. HATO mitigates this challenge by dynamically adjusting the perturbation generation process with a history-aware mechanism, optimizing perturbations across multiple sliding windows. This ensures that perturbations align with attack timing, maintaining their effectiveness even when attacks occur at different time steps.

Length T_{priv}	10	20	30	40	50
ASR(%)	40.11	37.31	39.11	37.61	35.54
Raw data EER(%)	8.30	7.98	6.44	5.86	5.83
Adv data EER(%)	41.65	41.46	43.02	41.15	38.24

Table 7: ASR and EER of different input lengths T_{priv} on privacy inference model.

Black-box model	MobileNet	Xception	FCN
ASR(%)	29.43	36.57	33.48
Raw data EER(%)	11.57	4.97	11.78
Adv data EER(%)	36.76	37.89	38.79

Table 8: Evaluation of PATN’s effectiveness across multiple black-box architectural models.

Black-box Migration Studies

Cross-Input-Length Migration on Privacy Inference Models. To assess the transferability and robustness of our adversarial perturbations, we train a single PATN with a fixed output length of $T_{out} = 10$, and use it to generate adversarial perturbations. These perturbations are then used to attack multiple black-box privacy inference models, each independently trained with a different input length ($T_{priv} \in \{20, 30, 40, 50\}$). As shown in Table 7, the perturbations remain highly effective across all models, with EER ranging from 38.24% to 43.02%. These results demonstrate that our PATN-generated perturbations exhibit strong generalization ability and temporal adaptability, effectively compromising models with varying receptive fields without requiring re-optimization.

Diverse Black-Box Privacy Inference Model Architectures. To further evaluate PATN’s transferability, we test its robustness against three unseen black-box privacy inference models: MobileNet, Xception, and a fully convolutional network (FCN). Although trained in a white-box setting with three CNNs (Section), these models have entirely different architectures. As shown in Table 8, PATN maintains consistent performance, achieving ASR of 29.43%, 36.57%, and 33.48%, and EER of 36.76%, 37.89%, and 38.79% on MobileNet, Xception, and FCN, respectively, demonstrating satisfactory privacy protection even against unseen or more complex models.

Conclusion

In this paper, we propose PATN, a history-driven perturbation generation framework that leverages past sensor data to synthesize future adversarial traces. By conditioning a generative network on historical signals, PATN enables real-time, zero-latency obfuscation while preserving fidelity to the original data. Perturbations can be applied instantly to incoming data, maintaining alignment with natural sensor patterns and mitigating inference-based privacy risks. For future work, we plan to evaluate PATN in black-box settings across more model architectures and extend its applicability to additional privacy-sensitive attributes, including but not limited to gender and age.

Acknowledgments

The authors sincerely thank all participants of the peer review. Special appreciation is extended to Xuanqi Gao, Yuhan Zhi, Weipeng Jiang, Yingmao Miao, Chen Ma, Yulong Yang, Yuchen Ren, and Chaoyi Wang for their valuable support and insightful suggestions that contributed to the completion of this work. This work is supported by the National Key Research and Development Program of China (2023YFB3107401), the National Natural Science Foundation of China (62402377, 62536002, T2341003, 62521002, 62376210, 62161160337, 62132011, U24B20185, U21B2018, 62206217), the Shaanxi Province Key Industry Innovation Program (2023-ZDLGY-38), JSPS KAKENHI JP23K24851, JST PRESTO JPMJPR23P5, JST CREST JPMJCR21M2, JST NEXUS JPMJNX25C4. Thanks to the New Cornerstone Science Foundation and the Xplorer Prize.

References

- Boutet, A.; Frindel, C.; Gambs, S.; Jourdan, T.; and Ngeueveu, R. C. 2021. DYSAN: Dynamically sanitizing motion sensor data against sensitive inferences through adversarial networks. In *Proceedings of the 2021 ACM asia conference on computer and communications security*, 672–686.
- Brotchie, J.; Shao, W.; Li, W.; and Kealy, A. 2022. Leveraging self-attention mechanism for attitude estimation in smartphones. *Sensors*, 22(22): 9011.
- Cai, Y.; Zhang, Z.; Gui, J.; Liu, B.; Zhao, X.; Li, R.; Li, Z.; and Li, D. 2024. {FAMOS}: Robust {Privacy-Preserving} Authentication on Payment Apps via Federated {Multi-Modal} Contrastive Learning. In *33rd USENIX Security Symposium (USENIX Security 24)*, 289–306.
- Ding, D.; Zhang, M.; Feng, F.; Huang, Y.; Jiang, E.; and Yang, M. 2023. Black-box adversarial attack on time series classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 7358–7368.
- Fawaz, H. I.; Forestier, G.; Weber, J.; Idoumghar, L.; and Muller, P.-A. 2019. Adversarial attacks on deep neural networks for time series classification. In *2019 International joint conference on neural networks (IJCNN)*, 1–8. IEEE.
- Fereidooni, H.; König, J.; Rieger, P.; Chilese, M.; Gökbakan, B.; Finke, M.; Dmitrienko, A.; and Sadeghi, A.-R. 2023. Authentisense: A scalable behavioral biometrics authentication scheme using few-shot learning for mobile platforms. *arXiv preprint arXiv:2302.02740*.
- Gu, W.; Zhong, R.; Zhang, J.; and Lyu, M. R. 2025. Towards Imperceptible Adversarial Attacks for Time Series Classification with Local Perturbations and Frequency Analysis. *arXiv preprint arXiv:2503.19519*.
- Hajihassnani, O.; Ardakanian, O.; and Khazaei, H. 2021. ObscureNet: Learning attribute-invariant latent representation for anonymizing sensor data. In *Proceedings of the international conference on internet-of-things design and implementation*, 40–52.
- Kalupahana, A. I. K.; Balaji, A. N.; Xiao, X.; and Peh, L.-S. 2023. Serandip: Leveraging inherent sensor random noise for differential privacy preservation in wearable community sensing applications. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 7(2): 1–38.
- Karim, F.; Majumdar, S.; and Darabi, H. 2020. Adversarial attacks on time series. *IEEE transactions on pattern analysis and machine intelligence*, 43(10): 3309–3320.
- Khan, S. U.; Mynuddin, M.; and Nabil, M. 2024. Adapt-edge: Targeted universal adversarial attacks on time series data in smart grids. *IEEE Transactions on Smart Grid*, 15(5): 5072–5086.
- Li, A.; Guo, J.; Yang, H.; Salim, F. D.; and Chen, Y. 2021. Deepobfuscator: Obfuscating intermediate representations with privacy-preserving adversarial learning on smartphones. In *Proceedings of the International Conference on Internet-of-Things Design and Implementation*, 28–39.
- Lin, C.; Song, T.; Miao, Y.; He, J.; Yang, M.; Wang, H.; and Shen, C. 2023. ChildShield: An Implicit and Continuous Child Identification System on Smartphones. *IEEE Transactions on Dependable and Secure Computing*, 21(4): 2257–2272.
- Lin, Z.; Jain, A.; Wang, C.; Fanti, G.; and Sekar, V. 2020. Using gans for sharing networked time series data: Challenges, initial promise, and open questions. In *Proceedings of the ACM internet measurement conference*, 464–483.
- Liu, S.; Du, J.; Shrivastava, A.; and Zhong, L. 2019. Privacy adversarial network: representation learning for mobile data privacy. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 3(4): 1–18.
- Malekzadeh, M. 2019. MotionSense Dataset for Human Activity and Attribute Recognition. Accessed: 2025-07-20.
- Malekzadeh, M.; Clegg, R. G.; Cavallaro, A.; and Haddadi, H. 2018. Protecting Sensory Data Against Sensitive Inferences. In *Proceedings of the 1st Workshop on Privacy by Design in Distributed Systems, W-P2DS'18*, 2:1–2:6. New York, NY, USA: ACM. ISBN 978-1-4503-5654-1.
- Meena, T.; and Sarawadekar, K. 2020. Gender recognition using in-built inertial sensors of smartphone. In *2020 IEEE REGION 10 CONFERENCE (TENCON)*, 462–467. IEEE.
- Mekruksavanich, S.; and Jitpattanakul, A. 2021. Lstm networks using smartphone data for sensor-based human activity recognition in smart homes. *Sensors*, 21(5): 1636.
- Miao, Y.; Tian, Q.; Lin, C.; Song, T.; Zhou, Y.; Zhao, J.; Gao, S.; Yang, M.; and Shen, C. 2023. Learning Heuristically-Selected and Neurally-Guided Feature for Age Group Recognition Using Unconstrained Smartphone Interaction. In *IJCAI*, 3029–3037.
- Oregi, I.; Del Ser, J.; Perez, A.; and Lozano, J. A. 2018. Adversarial sample crafting for time series classification with elastic similarity measures. In *International Symposium on Intelligent and Distributed Computing*, 26–39. Springer.
- Pialla, G.; Ismail Fawaz, H.; Devanne, M.; Weber, J.; Idoumghar, L.; Muller, P.-A.; Bergmeir, C.; Schmidt, D. F.;

- Webb, G. I.; and Forestier, G. 2025. Time series adversarial attacks: an investigation of smooth perturbations and defense approaches. *International Journal of Data Science and Analytics*, 19(1): 129–139.
- Rajguru, A. S. 2019. *Sensor monitor for android*. Ph.D. thesis, California State University, Sacramento.
- Rathore, P.; Basak, A.; Nistala, S. H.; and Runkana, V. 2020. Untargeted, targeted and universal adversarial attacks and defenses on time series. In *2020 international joint conference on neural networks (IJCNN)*, 1–8. IEEE.
- Raval, N.; Machanavajjhala, A.; and Pan, J. 2019. Olympus: Sensor privacy through utility aware obfuscation. *Proceedings on Privacy Enhancing Technologies*.
- Sharshar, A.; Fayez, A.; Ashraf, Y.; and Gomaa, W. 2021. Activity with gender recognition using accelerometer and gyroscope. In *2021 15th International Conference on Ubiquitous Information Management and Communication (IMCOM)*, 1–7. IEEE.
- Shen, C.; Li, Y.; Chen, Y.; Guan, X.; and Maxion, R. A. 2017. Performance analysis of multi-motion sensor behavior for active smartphone authentication. *IEEE Transactions on Information Forensics and Security*, 13(1): 48–62.
- Song, T. 2025. PATN. <https://github.com/skysky4/PATN>. GitHub repository.
- Susi, M.; Renaudin, V.; and Lachapelle, G. 2013. Motion mode recognition and step detection algorithms for mobile phone users. *Sensors*, 13(2): 1539–1562.
- Wang, J.; Chen, Y.; Hao, S.; Peng, X.; and Hu, L. 2019. Deep learning for sensor-based activity recognition: A survey. *Pattern recognition letters*, 119: 3–11.
- Wu, T.; Wang, X.; Qiao, S.; Xian, X.; Liu, Y.; and Zhang, L. 2022. Small perturbations are enough: Adversarial attacks on time series prediction. *Information Sciences*, 587: 794–812.
- Yang, W.; Yuan, J.; Wang, X.; and Zhao, P. 2022. TSadv: Black-box adversarial attack on time series with local perturbations. *Engineering Applications of Artificial Intelligence*, 114: 105218.
- Yang, X.; and Ardakanian, O. 2023. Privacy through diffusion: A white-listing approach to sensor data anonymization. In *Proceedings of the 5th Workshop on CPS&IoT Security and Privacy*, 101–107.
- Yang, X.; and Ardakanian, O. 2025. PrivDiffuser: Privacy-Guided Diffusion Model for Data Obfuscation in Sensor Networks. *Proceedings on Privacy Enhancing Technologies*.