sensors

MDPI

Article

SCEM: A Structure-Preserving Privacy Framework for Sensor-Generated Time-Series Data

Shehui Jin [1], Qian Pu [1], Shankui Zheng [2] and Haikuo Shen [1,*]

[1] School of Mechanical, Electronic and Control Engineering, Beijing Jiaotong University, No. 3 Shangyuancun, Haidian District, Beijing 100044, China; 23115173@bjtu.edu.cn (S.J.); qpu@bjtu.edu.cn (Q.P.)
[2] Henan Aerospace Rocket Co., Ltd., Jigu Road, Emergency Intelligent Manufacturing Industrial Park, No. 6 Factory Building, Qibin District, Hebi 458030, China; chejunting@fspace.net.cn
* Correspondence: shenhk@bjtu.edu.cn

Abstract

Intelligent sensor systems and Internet of Things (IoT) platforms continuously generate high-value time-series data for downstream analytics and operational monitoring. However, statistical, spectral, and distributional characteristics embedded in these sequences may inadvertently expose sensitive behaviors, operating states, or temporal routines. To address this issue, we propose Structured Constrained Energy Mapping (SCEM), a general and extensible framework for the feature-selective protection of sensor-generated time-series data. Its generality arises from a unified instantiation interface that supports customizable protection mechanisms for heterogeneous sensitive functionals. SCEM follows a four-stage operator workflow, namely feature mapping, structured perturbation, feature recovery, and utility-oriented calibration, with a generalized privacy energy representation supporting the transition from mapping to perturbation. We demonstrate the practical applicability of the framework through five representative instantiations covering spectral amplitudes, mean, variance, quantiles, and high-order moments. Experiments on two sensor-related monitoring benchmarks and one non-sensor temporal benchmark show that SCEM achieves an unweighted mean attack success rate at the 10% reporting threshold (ASR@10%) of 4.69% over the 18 reported SCEM feature–dataset pairs under the evaluated direct feature inference attacks, indicating the reduced recoverability of predefined feature-level information in this protocol. Meanwhile, SCEM maintains competitive forecasting-oriented utility compared with representative protection baselines and the original-data utility reference, achieving a favorable privacy–utility trade-off. Overall, the results suggest that SCEM can serve as a practical model-free preprocessing layer for reducing feature-level leakage while preserving useful temporal structures for intelligent sensor data sharing.

Keywords: sensor privacy; time series; IoT; feature inference; privacy–utility trade-off

Academic Editors: Jinbo Xiong, Shaoen Wu, Periklis Chatzimisios and Mahmoud Daneshmand

Received: 9 June 2026
Revised: 6 July 2026
Accepted: 13 July 2026
Published: 14 July 2026

1. Introduction

Sensor-generated time-series data are continuously produced by intelligent sensor networks, smart meters, industrial sensing systems, wearable and healthcare sensors, environmental monitoring sensors, and Internet of Things (IoT) sensing platforms [1,2]. These data provide a key foundation for edge and cloud sensor analytics, including forecasting, diagnosis, anomaly detection, monitoring, and operational decision-making. For instance, more than 1.2 billion smart meters had been deployed worldwide by 2023 [3], continuously

collecting fine-grained energy consumption data, which may raise privacy concerns [4,5]. At the same time, modern deep learning models increasingly require high-quality temporal data, whose predictive value depends on temporal dependence, statistical regularity, and dynamic evolution [6,7]. This creates a fundamental tension between sensor-generated time-series data sharing and the protection of sensitive information.

The structural nature of sensor-captured time-series data makes this tension particularly difficult to manage [8,9]. Sensitive information is not necessarily stored in individual samples; it may be encoded in frequency-domain characteristics, statistical summaries, temporal correlations, local order statistics, or high-order moments [10,11]. Such feature-level leakage can have direct operational and personal consequences. In smart meter streams, periodic load patterns can reveal occupancy schedules and recurring household activities [12–15]. In industrial monitoring, shifts in statistical or spectral characteristics may expose production cycles, machine utilization patterns, or proprietary operating regimes [16]. In wearable and healthcare sensing, recurring spectral and distributional patterns may disclose health-related states or private behavioral routines [17]. These risks can remain present even when individual samples appear innocuous, because sensitive information may be encoded in aggregate, distributional, and periodic structures. Recent smart meter privacy studies also examine residential energy management and AMI analytics under privacy constraints [18,19]. Therefore, privacy protection for sensor-generated time-series data should not only reduce access to raw values but also control the inferability of sensitive features while preserving structures required for downstream sensor analytics. Feature-level inference therefore represents a practical data security risk in sensor data sharing, rather than merely an abstract statistical concern.

Existing privacy protection methods for sensor-generated time-series data have been developed from several perspectives, including anonymization, differential privacy, synthetic data generation, and frequency-domain perturbation [20–23]. Differential privacy has been widely studied and provides formal sensitivity-based protection mechanisms [24,25]; recent surveys emphasize that the central difficulty is to balance privacy protection with temporal utility [26]. Nevertheless, many existing approaches are either applied directly to raw samples, where raw-value perturbation may damage the temporal continuity and spectral regularity needed for sensor analytics, or are designed for a specific feature type, without an explicit operator-level interface for heterogeneous sensitive functionals [27]. A recent energy-aware time-series analysis further suggests that feature-space and energy-space representations can help to characterize temporal patterns [28], but intelligent sensor systems still require controllable privacy–utility trade-offs before data sharing or cloud analytics.

This paper proposes Structured Constrained Energy Mapping (SCEM), a general and extensible operator interface framework for feature-selective privacy protection in sensor-generated time-series data. The central idea is to represent each sensitive functional through a feature-specific instantiation tuple, perturb the corresponding representation in a structured manner, and recover a feasible protected sequence through feature-specific back-projection and utility-oriented calibration. The five evaluated cases are representative instantiations used to illustrate its extensibility and practical applicability, serving as a foundation for broader feature-specific extensions within the SCEM architecture. This design enables the targeted control of sensitive features while constraining the unnecessary distortion of utility-related statistical, spectral, and temporal structures. The main contributions of this paper are summarized as follows:

(1) We formulate feature inference privacy for sensor-generated time-series data sharing. Instead of focusing only on raw-sample perturbation, this formulation emphasizes

sensitive statistical, spectral, and distributional features that may reveal user behaviors, operating states, or temporal routines in sensing scenarios.

(2) We propose Structured Constrained Energy Mapping (SCEM), a unified operator interface framework for feature-selective privacy protection. SCEM organizes heterogeneous sensitive functionals through a shared four-stage workflow, including feature mapping, structured perturbation, feature recovery, and utility-oriented calibration, supported by a generalized privacy energy representation between mapping and perturbation. This design enables targeted feature protection while constraining the unnecessary distortion of utility-related temporal structures.

(3) We instantiate SCEM for five representative privacy targets, including spectral amplitudes, mean, variance, quantiles, and high-order moments. Experiments on two sensor-related monitoring benchmarks and one non-sensor temporal benchmark show that SCEM reduces feature-level leakage while maintaining competitive forecasting-oriented utility, providing initial evidence of its applicability beyond physical sensor benchmarks under the evaluated settings.

2. Related Work

2.1. Random Perturbation and Differential Privacy

Random perturbation and differential privacy (DP) are fundamental techniques for data privacy protection [29]. Random perturbation protects data by injecting stochastic noise, whereas DP provides a formal mechanism for bounding the influence of individual records through sensitivity-calibrated randomization. For sensor time-series data, however, direct perturbation in the raw sample space may be structure-agnostic [8,21]. When temporal dependence, spectral regularity, and forecasting-oriented utility are important, noise that is not aligned with the semantics of the protected feature may unnecessarily alter non-sensitive structures and reduce downstream performance [9,30,31].

This limitation does not diminish the theoretical value of DP. Rather, it highlights a complementary problem setting in which the privacy target is a predefined statistical, spectral, or distributional functional of a released sequence. In this setting, a useful protection mechanism should first identify the representation in which the target functional is controllable and then modify the corresponding privacy-related component while explicitly constraining utility-related structures. Raw-space noise injection does not, by itself, provide such a functional-specific operator interface. Moreover, a statistic-level DP mechanism and a structure-preserving feature obfuscation mechanism address different privacy goals: the former provides individual-level indistinguishability under a specified neighboring relation, whereas the latter reduces the recoverability of selected functionals under a specified feature inference threat model. These two directions are therefore complementary rather than interchangeable.

2.2. Time-Series Anonymization and Feature Protection

Sensor time-series anonymization and feature protection methods move beyond global random perturbation by exploiting temporal segmentation, clustering, or transformed-domain representations [9,31]. The k-anonymity-based k-nTS method constructs equivalence groups through segmentation and clustering and can preserve selected temporal patterns during anonymization [32,33]. Frequency-domain anonymization and spectral privacy methods are effective when leakage is concentrated in periodic or harmonic components [20,30]. Quantile-oriented and attribute-oriented methods further show that privacy objectives can be expressed through specific statistical targets rather than only through raw-value distortion [34,35].

These approaches are effective in their intended protection spaces, but their internal mechanisms are usually tied to a particular feature family. Segmentation and clustering do not directly specify how a selected mean, variance, quantile, or spectral amplitude should move toward a privacy target. Spectral masking provides explicit control in the frequency domain, but it does not directly define how order statistics or high-order distributional descriptors should be protected. Likewise, a quantile-specific mechanism does not automatically extend to periodic patterns or dispersion-related characteristics. Therefore, the main unresolved issue is not the absence of effective feature-specific methods but the absence of a common construction rule that explains how heterogeneous functionals can be represented, perturbed, recovered, and calibrated under one extensible framework.

2.3. Structure-Aware and Learning-Based Privacy Preservation

Structure-aware perturbation methods seek to reduce utility loss by modifying sensitive information in transformed feature spaces instead of perturbing all raw samples indiscriminately [10,11,36]. In intelligent sensor systems, protection mechanisms should preserve structures required by downstream tasks such as forecasting, monitoring, anomaly detection, and diagnosis [9,21]. Studies on multivariate time-series representation, including transformer-based and convolutional sequence models, also show that downstream usefulness depends strongly on retaining task-relevant temporal structures [37,38].

Modern learning-based privacy approaches address this problem from complementary perspectives. Federated learning keeps model training local, which is valuable in medical and IoT applications [39]. Learned representations, privacy-oriented representation learning, and synthetic-data or generative approaches reduce the exposure of original sequences or retain selected distributional properties [37,40,41]. These methods provide strong modeling capacity but generally require training, architecture selection, task-specific optimization, or repeated parameter tuning, and their protection objectives are often expressed at the model, latent representation, or synthetic distribution level rather than as the direct control of an explicitly selected functional.

Beyond temporal data privacy, multimedia security research likewise shows that data protection mechanisms should be evaluated against explicit adversarial models rather than only through visual or statistical distortion. A recent cryptanalysis of a chaos- and S-box-based video cryptosystem identified structural weaknesses under chosen-plaintext, known-plaintext, and chosen-ciphertext attacks [42]. Although multimedia encryption and feature-selective time-series protection address different security goals, this broader evidence reinforces the importance of clearly specifying the attacker's knowledge, protection scope, and structural assumptions when designing data protection mechanisms.

SCEM is positioned as a model-free preprocessing framework that is complementary to federated learning, synthetic data generation, and privacy-preserving representation learning. It provides a direct operator interface for feature-selective modification before time-series data are published or supplied to an external analytical model.

Existing structure-aware methods nevertheless leave three mechanism-level gaps. First, many methods are coupled to a specific model, task, or representation rather than offering a reusable interface for heterogeneous sensitive functionals. Second, the conditions under which perturbation, reconstruction, and calibration produce a valid protected sequence are often implicit. Third, the roles of privacy modification and utility restoration are not always separated, making it difficult to determine whether protection arises from the intended feature perturbation or from subsequent clipping and correction. These gaps motivate an operator-level formulation in which each stage has an explicit responsibility.

2.4. Energy- and Feature-Space Representations

In sensor time-series analysis, energy-related quantities are widely used for feature extraction, anomaly detection, and distributional characterization [43]. Sensor signals are commonly described through frequency-domain amplitudes, squared deviations, statistical moments, and other transformed representations. For power-load and smart-metering data, privacy-relevant behavior may appear as localized spectral components or characteristic distributions that differ from non-sensitive background structures [40]. These observations suggest that privacy protection can be more targeted when it is performed in a representation aligned with the sensitive functional.

Energy-based and feature-space representations have also been explored in privacy-related generative modeling and synthetic smart-grid data analysis [40,41]. However, the term "energy" is used differently across spectral, statistical, and distributional settings. Spectral energy may refer to amplitude or power, variance-related energy may refer to squared centered deviations, and distributional protection may rely on a discrepancy between a feature representation and a target state. Without an explicit operational definition, combining these meanings can make a general framework difficult to interpret and test.

To avoid this ambiguity, SCEM uses a generalized privacy energy concept that is formally defined in Section 3. In SCEM, privacy energy is not assumed to be a universal physical quantity. It is defined as a non-negative, weighted discrepancy between a functional-dependent representation and its privacy reference. Spectral power, centered dispersion energy, and squared feature deviation are treated as different instantiations of this common construction. This definition separates the invariant framework-level role of energy from the functional-specific representation used in each protection mechanism.

2.5. Summary and Motivation

The preceding review reveals four mechanism-level limitations that motivate SCEM. First, raw-space random perturbation and statistic-level DP can provide important privacy protection, but they do not directly specify where a heterogeneous target functional should be represented or how non-target structures should be constrained. This motivates a functional mapping stage that makes the selected sensitive characteristic explicit in an appropriate representation space. Second, anonymization and feature-specific methods are usually effective only in a predefined protection domain, which limits their transfer from spectral features to means, variances, quantiles, or high-order moments. This motivates a generalized privacy energy construction and a structured perturbation stage whose internal realization is allowed to depend on the protected functional.

Third, transformed-domain modification must be paired with a recovery step that converts the perturbed representation into a valid time series. Frequency-domain masking, dispersion redistribution, local quantile adjustment, and moment shaping require different reconstruction rules. This motivates a separate recovery or back-projection stage. Fourth, privacy-oriented modification may introduce value range violations or unintended changes to non-sensitive statistical, spectral, and temporal structures. This motivates an explicit utility-oriented calibration stage governed by application-dependent constraints.

Accordingly, SCEM provides a common but customizable mapping–perturbation–recovery–calibration interface for heterogeneous sensitive functionals. Its scope is feature inference privacy for released time-series data, complementing rather than replacing individual-level DP, federated learning, and model-based synthetic data generation.

3. The Proposed SCEM Framework

3.1. Threat Model and Problem Formulation

Let $X = [x_1, x_2, \ldots, x_T]^\top \in \mathcal{X} \subseteq \mathbb{R}^T$ denote a time-series segment to be released. The segment may be collected by a smart meter, an environmental sensor, an industrial monitoring device, a wearable sensor, or another sensing platform. The objective is to construct a protected sequence $X^+ \in \mathcal{X}$ that reduces the recoverability of predefined sensitive information while retaining the structures required for downstream temporal analytics.

The sensitive information is represented by a measurable functional

$$\phi : \mathcal{X} \to \mathcal{Y}_\phi, \tag{1}$$

where $\mathcal{Y}_\phi$ may be scalar-, vector-, or representation-valued. Examples include the mean or variance, a selected quantile, a vector of high-order moments, a set of spectral amplitudes, an autocorrelation profile, or a domain-specific operational indicator. When several characteristics are considered jointly, they can be written as a vector-valued functional

$$\boldsymbol{\phi}(X) = [\phi_1(X), \phi_2(X), \ldots, \phi_K(X)]^\top. \tag{2}$$

The representative experiments in this work instantiate one functional at a time in order to isolate the behavior of each protection mechanism, whereas the framework itself permits joint functionals through Equation (2).

The adversary is assumed to observe the released sequence X^+, know the type of the protected functional, and construct an estimator $\widehat{\phi} = \mathcal{A}_\phi(X^+)$ using the available domain knowledge. For a functional-dependent discrepancy d_ϕ and tolerance $\tau_\phi > 0$, an inference attempt is regarded as successful when

$$d_\phi\big(\widehat{\phi}, \phi(X)\big) \le \tau_\phi. \tag{3}$$

The tolerance τ_ϕ is functional-dependent because heterogeneous characteristics have different scales, physical meanings, and admissible fluctuation ranges. Consequently, the common normalized threshold used later in the experiments serves as an operational reporting point for method comparison across heterogeneous functionals.

The present study focuses on direct feature inference attacks under a one-shot release setting; broader adaptive and learning-based attacks are discussed in Section 4.8. The SCEM interface defined below supports both deterministic and randomized operators, although the representative implementations evaluated in this study are deterministic for reproducibility.

Utility is described by a collection of non-sensitive or task-related functionals $\Psi = \{\psi_1, \psi_2, \ldots, \psi_M\}$. These functionals may encode value ranges, non-sensitive spectral amplitudes, selected statistical moments, local ordering, or task-related temporal descriptors. Given functional-specific discrepancy measures d_j and tolerances $\boldsymbol{\delta} = [\delta_1, \ldots, \delta_M]^\top$, the utility-feasible set is

$$\Omega_{\mathrm{u}}(X; \Psi, \boldsymbol{\delta}) = \big\{ Z \in \mathcal{X} \;\big|\; d_j\big(\psi_j(Z), \psi_j(X)\big) \le \delta_j,\ j = 1, \ldots, M \big\}. \tag{4}$$

Let $\vartheta_\phi \in \mathcal{Y}_\phi$ denote a privacy target, such as a non-sensitive spectral level, an alternative statistical value, or an admissible target interval. The corresponding privacy-target set is

$$\Omega_{\mathrm{p}}\big(\vartheta_\phi; \varepsilon_\phi\big) = \big\{ Z \in \mathcal{X} \;\big|\; d_\phi\big(\phi(Z), \vartheta_\phi\big) \le \varepsilon_\phi \big\}, \tag{5}$$

where ε_ϕ specifies the desired proximity to the privacy target.

Feature-selective protection is feasible only when the privacy and utility requirements are compatible:

$$\Omega_{\mathrm{p}}\left(\vartheta_{\phi};\varepsilon_{\phi}\right)\cap\Omega_{\mathrm{u}}(X;\Psi,\delta)\neq\varnothing. \tag{6}$$

Equation (6) makes the privacy–utility boundary explicit. If the same functional is simultaneously a sensitive target and an indispensable forecasting feature, arbitrarily strong obfuscation and zero utility loss cannot generally be achieved at the same time. In such cases, the admissible protection strength must be selected according to the application-specific utility tolerance.

The problem considered in this paper is therefore to construct a feature-conditioned operator

$$X^{\dagger}=F_{\phi}\left(X;\Theta_{\phi}\right), \tag{7}$$

where Θ_{ϕ} collects the representation, perturbation, recovery, calibration, and privacy strength parameters. The operator should move the selected sensitive functional toward its privacy target while returning a valid sequence inside the utility-feasible set.

For multivariate sensor data $X\in\mathbb{R}^{T\times D}$, the current implementation applies SCEM channel-wise. Each channel may correspond to a sensor variable, sensing modality, or monitored physical quantity. Joint multivariate protection can be constructed within the same formulation by defining $\boldsymbol{\phi}$ over several channels—for example, through cross-variable correlations, shared periodic patterns, or multichannel operational indicators.

3.2. Generalized Privacy Energy and Overall Design

SCEM provides a common operator interface through which heterogeneous sensitive functionals can be represented, perturbed, recovered, and calibrated. The framework does not require every functional to share an identical closed-form perturbation equation. Instead, each functional-specific instantiation must provide a compatible set of operators with clearly defined responsibilities.

For a selected functional ϕ, let $Z_{\phi}=\mathcal{M}_{\phi}(X)\in\mathcal{Z}_{\phi}$ denote a representation in which the target functional becomes explicit or controllable. Depending on ϕ, Z_{ϕ} may be an amplitude spectrum, a centered dispersion sequence, a normalized signal, a local order statistic neighborhood, or a standardized distributional state.

Let $r_{\phi}\in\mathcal{Z}_{\phi}$ denote a privacy reference in the representation space. SCEM defines the generalized privacy energy as

$$E_{\phi}=\mathcal{E}_{\phi}\left(X;\omega_{\phi},r_{\phi}\right)=\omega_{\phi}\odot\left[\mathcal{D}_{\phi}\left(\mathcal{M}_{\phi}(X),r_{\phi}\right)\right]^{\odot 2}, \tag{8}$$

where $\omega_{\phi}\geq 0$ is a structured weight, $\mathcal{D}_{\phi}$ is a representation-dependent discrepancy operator, $\odot$ denotes the Hadamard product, and $^{\odot 2}$ denotes component-wise squaring. By construction, the generalized privacy energy is component-wise non-negative whenever $\omega_{\phi}\geq 0$, because the discrepancy is squared component-wise before multiplication by the non-negative structured weight. The term "energy" in SCEM therefore denotes a non-negative discrepancy representation relative to a privacy reference. It is not restricted to a single physical unit. Spectral power, centered dispersion energy, and squared functional deviation are specific realizations of Equation (8).

A functional-specific SCEM instantiation is represented by the tuple

$$\mathfrak{I}_{\phi}=\left(\phi,\mathcal{M}_{\phi},r_{\phi},\mathcal{E}_{\phi},\mathcal{T}_{\phi},\mathcal{B}_{\phi},\mathcal{C}_{\phi},\Theta_{\phi}\right), \tag{9}$$

where $\mathcal{T}_{\phi}$ is the structured perturbation operator, $\mathcal{B}_{\phi}$ is the recovery or back-projection operator, and $\mathcal{C}_{\phi}$ is the utility-oriented calibration operator.

The complete workflow is written as

$$\begin{aligned} Z_\phi &= \mathcal{M}_\phi(X), \\ E_\phi &= \mathcal{E}_\phi\big(X;\omega_\phi,r_\phi\big), \\ \widetilde{Z}_\phi &= \mathcal{T}_\phi\big(Z_\phi,E_\phi;\theta_\phi,\xi_\phi\big), \\ \widetilde{X} &= \mathcal{B}_\phi\Big(\widetilde{Z}_\phi;X\Big), \\ X^\dagger &= \mathcal{C}_\phi\Big(\widetilde{X};X,\Omega_{\mathrm{u}}\Big). \end{aligned} \tag{10}$$

Here, θ_ϕ controls the perturbation strength or target behavior, and ξ_ϕ is an optional random variable. A fixed ξ_ϕ gives a deterministic instantiation, whereas a sampled ξ_ϕ permits randomized references, weights, perturbation directions, or calibration choices without changing the common operator interface.

As shown in Figure 1, the four operator stages have different responsibilities. The mapping stage determines where the sensitive functional can be represented. The perturbation stage determines how the privacy-related representation should be modified, using the generalized privacy energy quantity where appropriate as an intermediate discrepancy representation rather than as an additional workflow stage. The recovery stage reconstructs a valid time-series sequence from the perturbed representation. The calibration stage enforces range, statistical, spectral, structural, or task-related utility constraints. This separation is the basis of SCEM's generality: the operator interface remains invariant, whereas the internal realization of each operator is allowed to depend on the protected functional.

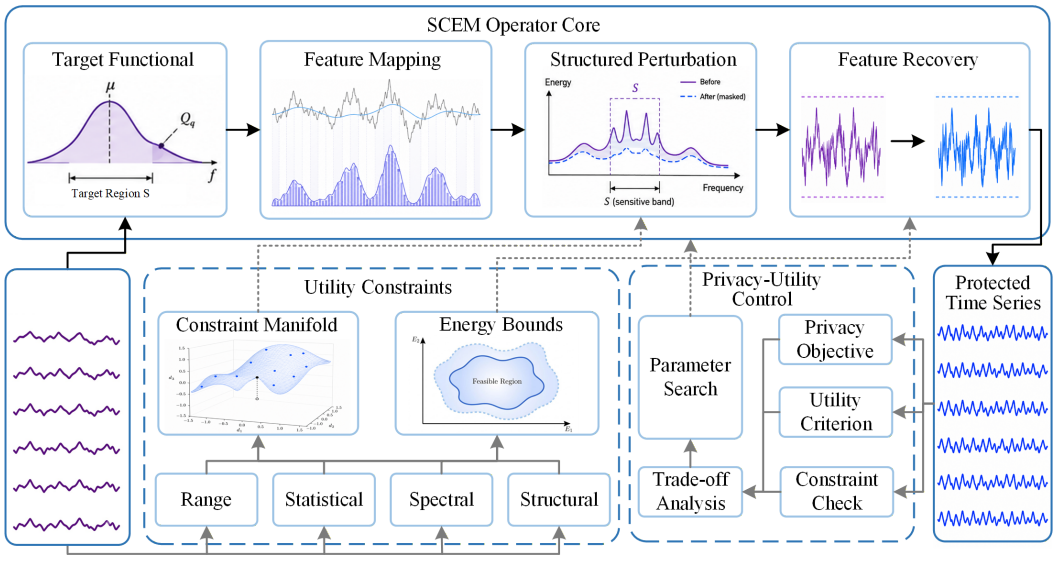

Figure 1. Overall architecture of the SCEM framework. The core workflow consists of target functional specification, feature mapping, structured perturbation, feature recovery, and utility-oriented calibration. The generalized privacy energy quantity is an intermediate discrepancy representation supporting structured perturbation, and utility constraints correspond to the utility-feasible set and calibration stage.

Definition 1 (Admissible SCEM Instantiation)**.** *A tuple $\mathfrak{I}_\phi$ in Equation (9) is an admissible SCEM instantiation if it satisfies the following conditions:*

1. *Functional representability. There exists a map g_ϕ such that $\phi(X) = g_\phi(\mathcal{M}_\phi(X))$.*
2. *Non-negativity and reference consistency. The generalized privacy energy is component-wise non-negative, $\mathcal{E}_\phi(X;\omega_\phi,r_\phi) \geq 0$, and zero discrepancy corresponds to the selected reference state on the weighted support.*

3. *Bounded and controllable perturbation. For admissible parameters, $\mathcal{T}_\phi$ maps bounded representations to bounded representations and varies continuously with the admissible control parameter. A quantitative parameter stability condition is stated in the subsequent theorem.*
4. *Recovery of a valid sequence. The recovery operator satisfies $\mathcal{B}_\phi(\widetilde{Z}_\phi; X) \in \mathcal{X}$. It need not be the exact mathematical inverse of $\mathcal{M}_\phi$, but it must return a valid sequence consistent with the selected instantiation.*
5. *Utility-feasible calibration. The calibrated output satisfies*

$$\mathcal{C}_\phi\left(\widetilde{X}; X, \Omega_{\mathrm{u}}\right) \in \Omega_{\mathrm{u}}(X; \Psi, \delta). \tag{11}$$

6. *Privacy–utility compatibility. The privacy target and utility constraints satisfy Equation (6). Otherwise, the privacy target or utility tolerances must be relaxed.*

These conditions clarify the meaning of framework generality: SCEM obtains its generality from a shared admissible operator interface, while the representation, privacy reference, perturbation rule, recovery operator, and calibration constraints remain customizable for each functional. The calibration condition directly ensures that the final protected sequence belongs to the prescribed utility-feasible set.

For a vector-valued sensitive functional $\boldsymbol{\phi} = [\phi_1, \ldots, \phi_K]^\top$, a joint instantiation can be constructed by block-wise concatenation:

$$\begin{aligned} \mathcal{M}_{\boldsymbol{\phi}}(X) &= \mathrm{concat}\big(\mathcal{M}_{\phi_1}(X), \ldots, \mathcal{M}_{\phi_K}(X)\big), \\ E_{\boldsymbol{\phi}} &= \mathrm{concat}\big(\lambda_1 E_{\phi_1}, \ldots, \lambda_K E_{\phi_K}\big), \qquad \lambda_k \geq 0. \end{aligned} \tag{12}$$

This construction permits combined protection objectives without forcing heterogeneous functionals to have the same representation dimension. The perturbation and calibration operators must then satisfy the corresponding joint privacy and utility constraints.

3.3. Representative Instantiations and Extensibility of SCEM

This subsection presents five representative instantiations spanning spectral, first-order statistical, second-order statistical, order-statistical, and high-order distributional functionals. Although their internal operators are functional-specific, all five cases follow the same mapping–perturbation–recovery–calibration interface supported by a generalized privacy energy representation.

For sensitive frequency amplitudes associated with periodic behavior or recurrent operating patterns, SCEM operates in the spectral domain. The original sequence is transformed as

$$\widehat{X} = \mathcal{F}(X), \qquad A = \left|\widehat{X}\right|, \qquad P = \angle \widehat{X}, \tag{13}$$

where A and P denote the amplitude and phase spectra, respectively. Let S denote the sensitive frequency set, and let $A_{\mathrm{bg},k}$ denote the non-sensitive background amplitude estimated from neighboring frequencies outside S. The frequency privacy energy is

$$E_{\mathrm{f},k} = w_{\mathrm{f},k}\Big(A_k - A_{\mathrm{bg},k}\Big)^2, \qquad k \in S, \tag{14}$$

where $w_{\mathrm{f},k}$ emphasizes the sensitive region.

A reference-oriented contraction is applied to the sensitive amplitudes:

$$\widetilde{A}_k = A_{\mathrm{bg},k} + \alpha_{\mathrm{f}}\Big(A_k - A_{\mathrm{bg},k}\Big), \qquad k \in S, \tag{15}$$

where $\alpha_f \in [0,1]$. Smaller values of α_f move the sensitive amplitudes closer to the background level. Frequencies outside S are left unchanged except for optional boundary smoothing.

The preliminary protected sequence is reconstructed by

$$\widetilde{X} = \mathcal{F}^{-1}\left(\widetilde{A} \odot e^{iP}\right). \tag{16}$$

Range correction and non-sensitive spectral calibration are then applied to limit unintended distortion outside the protected frequency region.

For mean-related functionals, the protected mean level may reveal operating states, load levels, or long-term environmental conditions. Let $u_t \in [0,1]$ denote the normalized sample value, let μ_u denote its mean, and let μ_{ref} denote the target mean. Define $\Delta_\mu = \mu_{\text{ref}} - \mu_u$. A bounded adjustment weight $w_\mu(t) \geq 0$ is selected such that

$$\frac{1}{T}\sum_{t=1}^{T} w_\mu(t) = 1. \tag{17}$$

The preliminary protected value is

$$\widetilde{u}_t = u_t + \Delta_\mu w_\mu(t). \tag{18}$$

Before range projection, Equation (17) yields $\mu(\widetilde{u}) = \mu_u + \Delta_\mu$. The weight design assigns smaller changes to samples near the admissible boundaries and larger changes to samples with greater adjustment capacity. The sequence is subsequently mapped back to the original scale and calibrated with respect to value range and selected non-sensitive moments. Rank consistency can be imposed as an optional constraint when ordering is utility-relevant.

For variance-related functionals, protection is formulated in the centered dispersion energy space. Let $z_t = x_t - \mu$ and $e_t = z_t^2$, where μ is the sample mean. Let σ^2_{ref} denote the target variance, and let $q_v(t) \geq 0$ satisfy $\sum_{t=1}^{T} q_v(t) = 1$. The required total dispersion change is $\Delta E_v = T(\sigma^2_{\text{ref}} - \sigma^2)$. The redistributed dispersion energy is

$$\widetilde{e}_t = \max(e_t + \Delta E_v q_v(t), \epsilon). \tag{19}$$

The preliminary protected sequence is reconstructed by preserving the original centered sign:

$$\widetilde{x}_t = \mu + \text{sign}(z_t)\sqrt{\widetilde{e}_t}. \tag{20}$$

Mean correction and variance rescaling are applied after reconstruction. When the lower truncation in Equation (19) is inactive, the redistributed total energy matches the specified target exactly; otherwise, the calibration stage restores the intended variance within the prescribed tolerance.

For quantile and median functionals, the protected values are determined by local order statistics and therefore require a representation concentrated around the selected order-statistic region. Let $Q_q(u)$ denote the q-quantile of a normalized sequence u, and let $Q_{q,\text{ref}}$ denote the target quantile. A bounded localization kernel $\kappa_q(\cdot; s_q)$ is used to define

$$\overline{w}_q(t) = \kappa_q\left(\left|u_t - Q_q(u)\right|; s_q\right), \qquad w_q(t) = \frac{\overline{w}_q(t)}{\max_j \overline{w}_q(j) + \epsilon}, \tag{21}$$

where $s_q > 0$ controls the width of the affected neighborhood. The adjustment direction is $d_q = \mathrm{sign}\Big(Q_{q,\mathrm{ref}} - Q_q(u)\Big)$, and the preliminary local adjustment is

$$\widetilde{u}_t = u_t + \eta_q d_q w_q(t), \tag{22}$$

where $\eta_q \geq 0$ controls the adjustment magnitude. Because the weight is concentrated near $Q_q(u)$, the transformation modifies the selected order-statistic region while limiting changes to distant samples. The recovered sequence is projected onto the admissible range and, when required, reassigned through an order-consistent correction. The median is obtained as the special case $q = 0.5$.

For high-order moment functionals, skewness and kurtosis characterize distributional asymmetry and tail behavior. SCEM first standardizes the sequence as $z = (X - \mu\mathbf{1})/(\sigma + \epsilon)$, where μ and σ are the sample mean and standard deviation. A sign-preserving moment-shaping transformation is then applied:

$$z^{\dagger} = \mathrm{sign}(z) \odot (|z| + \epsilon)^{\gamma_m}, \tag{23}$$

where $\gamma_m > 0$ controls the deformation of central and tail samples, and $\gamma_m = 1$ corresponds approximately to the identity transformation. In the evaluated implementation, γ_m is generated from the user-level perturbation strength coefficient p_m through a fixed moment-shaping rule; Section 4.1 therefore reports the user-controlled coefficient p_m. The preliminary reconstruction is

$$\widetilde{X} = \mu\mathbf{1} + \sigma z^{\dagger}. \tag{24}$$

The calibration stage corrects the mean, variance, and admissible range so that the high-order distributional shape is modified without introducing unnecessary first- or second-order drift.

These five cases span spectral, first-order, second-order, order-statistical, and high-order distributional functionals under the same admissible operator interface. Additional functionals can be incorporated by specifying the corresponding representation, reference, perturbation, recovery, and calibration operators.

3.4. Framework Properties and Theoretical Boundaries

This subsection states the properties that follow from the common SCEM interface and distinguishes them from properties that require additional instantiation-specific assumptions.

Let $E_{\phi,\mathrm{ref}}$ denote the selected reference state in the generalized privacy energy space; when Equation (8) measures discrepancy directly from r_ϕ, the natural reference is $E_{\phi,\mathrm{ref}} = 0$.

Proposition 1 (Reference-Oriented Contraction)**.** *For reference-oriented affine energy perturbation, consider the operator*

$$\mathcal{T}_\phi^E(E;\beta_\phi) = \beta_\phi E + (1-\beta_\phi)E_{\phi,\mathrm{ref}}, \qquad \beta_\phi \in [0,1]. \tag{25}$$

Then,

$$\Big\|\mathcal{T}_\phi^E(E;\beta_\phi) - E_{\phi,\mathrm{ref}}\Big\| = \beta_\phi \Big\|E - E_{\phi,\mathrm{ref}}\Big\|. \tag{26}$$

Proof. Subtracting $E_{\phi,\mathrm{ref}}$ from both sides of Equation (25) gives $\mathcal{T}_\phi^E(E;\beta_\phi) - E_{\phi,\mathrm{ref}} = \beta_\phi(E - E_{\phi,\mathrm{ref}})$. Taking the norm yields Equation (26). □

Corollary 1 (Spectral Energy Contraction). *Using Equations (14) and (15), the perturbed spectral privacy energy satisfies*

$$\widetilde{E}_{\mathrm{f},k} = w_{\mathrm{f},k}\left(\widetilde{A}_k - A_{\mathrm{bg},k}\right)^2 = \alpha_{\mathrm{f}}^2 E_{\mathrm{f},k}. \tag{27}$$

Therefore, α_{f} is the amplitude-space contraction factor, α_{f}^2 is the corresponding energy-space contraction factor, and $\beta_{\mathrm{f}} = \alpha_{\mathrm{f}}^2$ in the notation of the preceding proposition.

Theorem 1 (Boundedness and Parameter Stability). *Assume that the input domain and admissible parameter set are bounded and that $\mathcal{M}_\phi$, $\mathcal{E}_\phi$, $\mathcal{T}_\phi$, $\mathcal{B}_\phi$, and $\mathcal{C}_\phi$ are bounded on their respective admissible domains. Then, F_ϕ maps bounded inputs to bounded outputs. Assume further that the operators $\mathcal{M}_\phi$, $\mathcal{E}_\phi$, $\mathcal{T}_\phi$, $\mathcal{B}_\phi$, and $\mathcal{C}_\phi$ are jointly Lipschitz-continuous with respect to their inputs and the admissible parameter θ_ϕ, uniformly over the bounded admissible domains, with all other admissible parameters held fixed. Then, there exists a finite constant L_ϕ such that*

$$\left\|F_\phi(X;\theta_1) - F_\phi(X;\theta_2)\right\| \le L_\phi|\theta_1 - \theta_2|. \tag{28}$$

Proof. Boundedness follows from the composition of bounded maps. The Lipschitz result follows by repeatedly applying the Lipschitz inequality through the complete operator workflow, including the privacy energy construction. The constant L_ϕ is obtained by recursively combining the input-Lipschitz and parameter-Lipschitz constants of the component operators over the bounded admissible domains. □

Proposition 2 (Perturbation–Calibration Displacement Bound). *Assume that d_ϕ is a metric or a norm-induced discrepancy. Let $\widetilde{X}$ denote the recovered sequence before calibration, and define*

$$\begin{aligned} m_\phi &= d_\phi\left(\phi(\widetilde{X}), \phi(X)\right), \\ \kappa_\phi &= d_\phi\left(\phi(X^\dagger), \phi(\widetilde{X})\right). \end{aligned} \tag{29}$$

where m_ϕ is the pre-calibration functional displacement and κ_ϕ is the additional displacement introduced by calibration. Then,

$$d_\phi\left(\phi(X^\dagger), \phi(X)\right) \ge m_\phi - \kappa_\phi. \tag{30}$$

Proof. The result follows from the reverse triangle inequality applied to d_ϕ. □

Equation (30) clarifies the roles of the SCEM stages: mapping identifies the controllable representation, structured perturbation and recovery jointly create the pre-calibration privacy-related displacement, and calibration limits utility distortion.

Remark 1 (Instantiation-Specific Structural Constraints). *Additional structural properties can be incorporated as instantiation-specific utility constraints. For example, when ordering is utility-relevant, the utility-feasible set may additionally require* $\mathrm{argsort}(Z) = \mathrm{argsort}(X)$.

Remark 2 (Deterministic and Randomized Instantiations). *The optional variable ξ_ϕ in Equation (10) allows the randomization of the privacy reference, structured weights, perturbation direction, or calibration choices. The deterministic instantiations used in the present experiments facilitate reproducibility and controllability, while randomized variants can be designed for repeated-observation or mapping dictionary threat models.*

Together, the preceding results establish non-negative generalized privacy energy, reference-oriented contraction, output feasibility, bounded parameter behavior, and a stage-wise displacement bound for admissible SCEM instantiations.

3.5. Privacy–Utility Behavior and Complexity Analysis

SCEM treats privacy protection and utility preservation as coupled objectives. Stronger perturbation generally moves the selected functional farther from its original value or closer to its privacy reference but may increase the distortion of utility-related structures. The admissible protection strength is therefore determined by the compatibility condition in Equation (6), rather than by a universal parameter setting shared across all instantiations.

For mean shifting, variance redistribution, and moment shaping, feature construction, perturbation, recovery, and one calibration pass are linear in the sequence length. With C calibration iterations, the complexity is $\mathcal{O}(T) + \mathcal{O}(CT)$.

For spectral instantiations, the dominant operations are the Fourier transform and inverse reconstruction, giving a complexity of $\mathcal{O}(T \log T) + \mathcal{O}(CT)$.

For quantile-related instantiations, full sorting has complexity $\mathcal{O}(T \log T)$, whereas selection-based or pre-sorted local implementations can reduce the order-statistic search cost. Optional projection onto a small number of task-related bases adds a cost proportional to the number of retained bases.

For a D-channel dataset processed independently, the corresponding channel-wise cost scales linearly with D. SCEM does not require the training of a generative or predictive model as part of the protection mechanism. Its operator-based complexity suggests potential suitability for large-scale preprocessing and resource-constrained deployment, although device-level latency and energy consumption require platform-specific validation.

In summary, SCEM separates privacy targets, utility constraints, and functional-specific operators within a shared admissible interface, providing a basis for extensible structure-preserving time-series protection.

4. Experimental Design and Result Analysis

4.1. Experimental Setup

To evaluate SCEM, we use two sensor-related monitoring benchmarks, ETTm1 and Weather, and one non-sensor temporal benchmark, Exchange Rate, for cross-domain evaluation [44,45]. ETTm1 contains 7 variables collected at 15 min intervals with 69,680 time steps; Weather contains 21 meteorological variables recorded at 10 min intervals with 52,696 time steps; and Exchange Rate contains 8 daily exchange rate variables with 7588 time steps. ETTm1 and Weather represent energy or industrial monitoring and environmental sensing scenarios, respectively, whereas Exchange Rate is included to examine whether the operator interface design remains applicable outside physical sensing scenarios under the same evaluated protocol.

SCEM is evaluated through five representative instantiations covering sensitive frequency amplitudes, mean, variance, median, and high-order moments. The reported experiments instantiate one sensitive functional at a time so that the behavior of each protection mechanism can be examined separately. For multivariate datasets, SCEM is applied channel-wise to a fixed set of target variables selected before evaluation. The same variables, input sequences, attack windows, and downstream forecasting protocol are used for all protection methods. Statistical feature inference attacks were evaluated over normalized windows. For ETTm1 and Weather, the analysis used the first 20 non-overlapping windows, each containing 2000 samples; for Exchange Rate, the analysis used 20 sliding windows of length 2000 with stride of 200.

Three representative baselines covering random perturbation, statistic-level differential privacy, and time-series anonymization are considered. Additive Gaussian Noise (AGN) is implemented as

$$X_{\mathrm{AGN}} = X + \eta, \qquad \eta_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(b, \sigma^2), \qquad b = x_{\max} - x_{\min}, \quad \sigma = 0.1b, \tag{31}$$

with random seed 42. Thus, b is the Gaussian mean bias set to the data range, rather than a clipping bound. The statistic-level sliding-window Laplace baseline perturbs window means with nominal privacy budget $\epsilon_{\mathrm{DP}} = 0.5$, window length 48, and scale $\Delta f/\epsilon_{\mathrm{DP}}$, where $\Delta f = (x_{\max} - x_{\min})/48$. Because overlapping windows repeatedly query shared samples, this method is treated as a heuristic statistic-level DP reference rather than a strict sequence-level DP mechanism. The k-nTS baseline uses segmentation and clustering-based anonymization with $k = 3$, window length $m = 20$, and step size 10 [32].

Min–max normalization is applied before downstream model training. For local statistical feature inference evaluation, normalization is performed independently within each attack window, as defined in Equation (33). The sensitivity experiment varies the reference-oriented control parameter over $\alpha \in \{0.1, 0.2, \ldots, 0.9\}$. All reported parameters are fixed before evaluation and are not tuned separately for individual evaluation windows.

For the representative ETTm1 spectral instantiation, the sensitive frequency is set to 2 cycles/day with a sensitive-band tolerance of ± 0.03 cycles/day. The background spectrum is estimated from the neighboring region within ± 0.1 cycles/day, excluding the sensitive band. Sensitive amplitudes are adjusted using the background mean and standard deviation with a structural scaling coefficient of 0.5, followed by Gaussian smoothing with $\sigma = 0.3$. The protected sequence is reconstructed by inverse FFT while retaining the original phase and is subsequently calibrated to preserve the original mean.

For the representative statistical SCEM instantiations, the mean-related configuration used the proportional adjustment mode. The shift coefficients were 0.75 for ETTm1, 1.0 for Weather, and 0.5 for Exchange Rate, corresponding to target means $\mu_{\mathrm{ref}} = 1.75\mu$, $\mu_{\mathrm{ref}} = 2.0\mu$, and $\mu_{\mathrm{ref}} = 1.5\mu$, respectively. The variance-related configuration used $\gamma_{\mathrm{v}} = 0.5$ for ETTm1 and $\gamma_{\mathrm{v}} = 1.0$ for Weather and Exchange Rate, yielding target variances $\sigma^2_{\mathrm{ref}} = 1.5\sigma^2$ and $\sigma^2_{\mathrm{ref}} = 2.0\sigma^2$, respectively. For median protection, the target quantile was fixed at $q = 0.5$, with a localization width parameter of $s_q = 0.5$, a privacy energy contraction coefficient of 0.3, and a back-projection step ratio of 0.5 for all three datasets; the orthogonal perturbation amplitude was 200 for ETTm1 and Weather and 20 for Exchange Rate. The high-order moment instantiation used a unified perturbation strength coefficient of $p_m = 0.1$ for the skewness- and kurtosis-related transformations. These settings were fixed before evaluation and were not tuned separately for individual evaluation windows.

The downstream utility evaluator is a one-layer LSTM implemented in PyTorch. For each prediction sample, the input length is set to 288 time steps, and the output is a 24-step-ahead prediction vector. The input dimension equals the number of variables used by the corresponding dataset, and the hidden dimension of the LSTM is set to 50. The last hidden output of the LSTM is passed to a fully connected linear layer to generate the 24-dimensional prediction output. The batch size is 32, and the model is trained for 50 epochs. Adam is used as the optimizer, with a learning rate of 0.001, and the mean squared error is used as the training loss. The samples are divided chronologically into 80% training data and 20% validation data after min–max normalization to $[0, 1]$. The training mini-batches are shuffled, whereas the validation mini-batches are not shuffled. No dropout, early stopping, learning rate scheduler, or dataset-specific architecture adjustment is used in the reported experiments. The same architecture, data split, and training protocol are used for the original and all protected datasets.

Feature-level privacy is evaluated using the attack success rate (ASR), relative difference rate (RDR), and peak ratio threshold (PRT). Forecasting-oriented utility is assessed using the R^2, MAE, RMSE, and MAPE (%). All experiments are conducted on a workstation equipped with an Intel Core i5-14600KF CPU and an NVIDIA RTX 5070 GPU with 12 GB memory. The implementation uses Python 3.9, NumPy 1.24, SciPy 1.10, Scikit-learn 1.2, and Matplotlib 3.7.

4.2. Evaluation Metrics

We evaluate protection from two complementary perspectives: the recoverability of predefined sensitive functionals and the preservation of forecasting-oriented utility.

The primary privacy metric is the attack success rate:

$$\mathrm{ASR} = \frac{N_{\mathrm{success}}}{N_{\mathrm{total}}} \times 100\%, \tag{32}$$

where N_{success} and N_{total} denote the numbers of successful inference trials and total evaluated trials, respectively. A lower ASR indicates lower recoverability of the selected functional under the evaluated attack protocol.

For statistical feature inference, all features are computed after window-wise min–max normalization:

$$x_{\mathrm{norm}} = \frac{x - \min(X_w)}{\max(X_w) - \min(X_w) + \epsilon}, \tag{33}$$

where X_w is the local evaluation window and ϵ is a small numerical constant. This normalization reduces the influence of trivial scale differences and yields a more demanding feature inference setting focused on a normalized statistical structure.

For a statistical functional ϕ, the normalized relative discrepancy is

$$\Delta_\phi(X_p, X_o) = \frac{\left|\phi(X_p^{\mathrm{norm}}) - \phi(X_o^{\mathrm{norm}})\right|}{|\phi(X_o^{\mathrm{norm}})| + \epsilon} \times 100\%, \quad \phi \in \{\mu, \sigma^2, Q_{0.5}, \mathrm{Skew}, \mathrm{Kurt}\}. \tag{34}$$

A statistical inference trial is counted as successful when

$$\Delta_\phi(X_p, X_o) \leq \tau_\phi. \tag{35}$$

Because Δ_ϕ is expressed as a relative percentage discrepancy, the choice $\tau_\phi = 10\%$ provides a dimensionless common numerical operating point for comparative reporting across the evaluated statistical functionals. It is not interpreted as a universal physical tolerance. In an application-specific deployment, a functional-dependent attack tolerance vector $\boldsymbol{\tau} = [\tau_{\phi_1}, \ldots, \tau_{\phi_K}]^\top$ should be selected according to sensor precision, natural variability, domain risk, and the operational meaning of each protected functional. The resulting metric is denoted as ASR@10%.

For frequency-domain inference, let A_{target} be the amplitude of the target sensitive frequency and $A_{\max}$ the maximum spectral amplitude. The peak ratio is

$$\mathrm{PRT} = \frac{A_{\mathrm{target}}}{A_{\max}} \times 100\%. \tag{36}$$

A frequency-domain inference trial is counted as successful when PRT $\geq \tau_f$. The same numerical value, $\tau_f = 10\%$, is used as the reporting point for frequency-domain inference. However, τ_f measures spectral prominence through PRT rather than the relative statistical reconstruction error. Therefore, the statistical and frequency-domain criteria share a common numerical operating point but not a common physical interpretation.

The relative difference rate is $\mathrm{RDR}_\phi = \Delta_\phi(X_p, X_o)$. For a privacy-sensitive functional, a larger RDR indicates stronger modification of the original functional. For a utility-related non-sensitive functional, a smaller RDR indicates better preservation. RDR is therefore interpreted together with the ASR and downstream forecasting results rather than as a standalone measure.

Forecasting-oriented utility is measured using a fixed LSTM evaluator [46]; a higher R^2 and lower MAE, RMSE, and MAPE (%) indicate better predictive utility under the adopted protocol.

4.3. Forecasting-Oriented Utility Analysis

The utility analysis examines temporal trajectories, feature distributions, low-dimensional visualizations, feature-level shifts, and downstream forecasting results.

Figure 2 shows that SCEM modifies the designated sensitive characteristics while retaining the overall temporal evolution of the evaluated sequences.

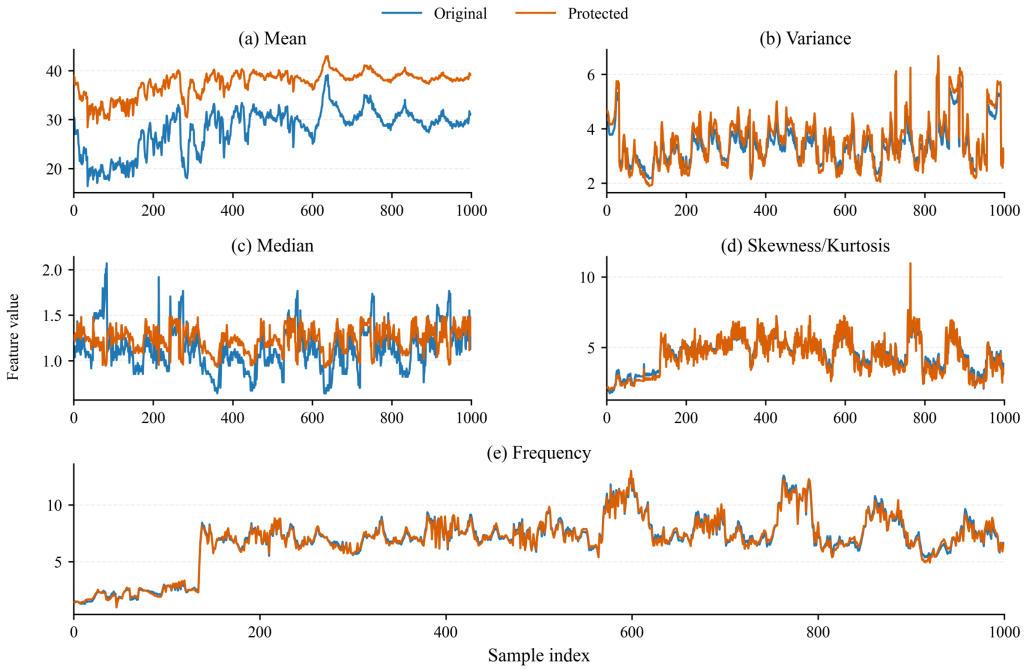

Figure 2. Temporal comparison of original and protected features over the evaluated segment.

Figure 3 shows visible shifts in selected sensitive distributions together with substantial overlap in several utility-related characteristics.

Figure 4 provides qualitative evidence that SCEM retains a low-dimensional organization that is closer to the original data than the evaluated baselines.

For Figure 5, the displayed feature shift value is defined as the unscaled relative ratio

$$S_\phi(X_p, X_o) = \frac{\left|\phi(X_p) - \phi(X_o)\right|}{|\phi(X_o)| + \epsilon}. \tag{37}$$

Thus, a displayed value of 1.502 corresponds to a relative shift of 150.2%, whereas 0.199 corresponds to 19.9%. The numerical annotations show the original untransformed relative ratios, whereas the heatmap colors are displayed using logarithmic normalization to improve visual discrimination across different magnitudes. The values are comparable within the same functional and dataset, but their absolute magnitudes should not be compared across heterogeneous functionals without considering their different scales and meanings.

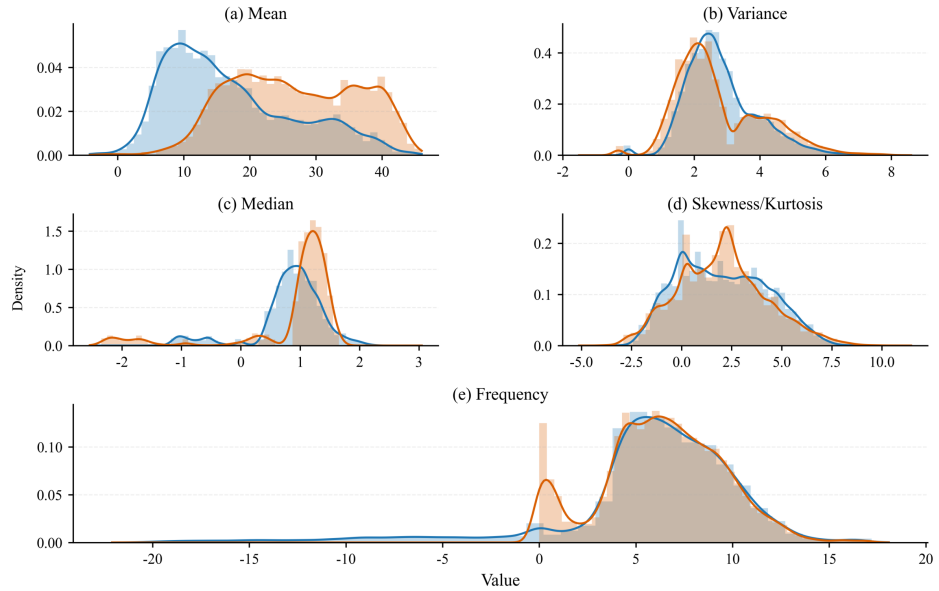

Figure 3. Distributional comparison of original and protected features.

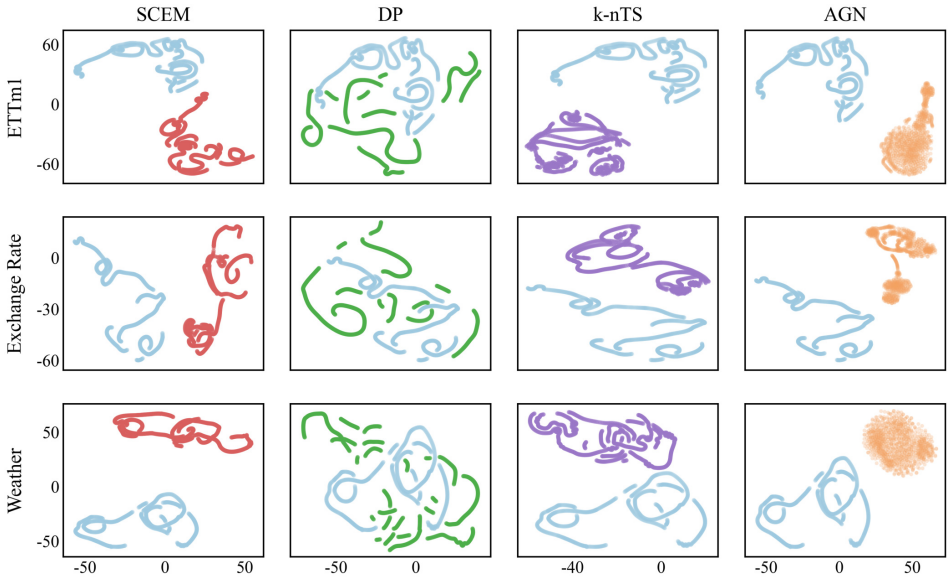

Figure 4. t-SNE visualization of original and protected time-series representations. The visualization provides a qualitative comparison of the low-dimensional organization of original and protected time-series representations.

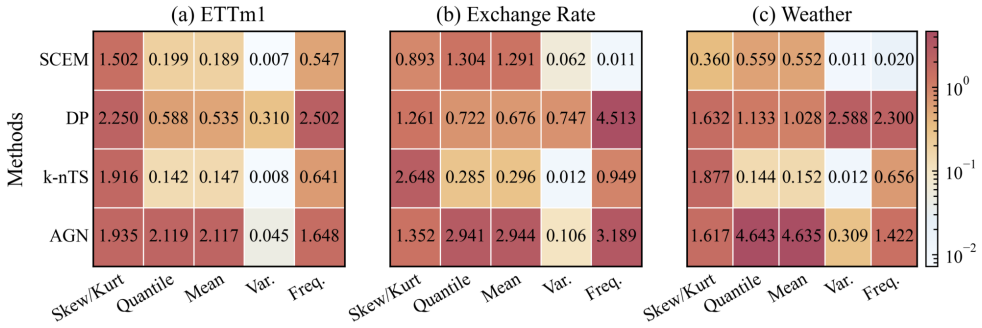

Figure 5. Relative feature shift ratio across datasets and protection methods, calculated using Equation (37). The numerical annotations are original untransformed relative ratios; for example, 1.0 corresponds to a 100% relative shift. Heatmap colors use logarithmic normalization to improve visual discrimination across different magnitudes.

Figure 5 shows that SCEM produces non-uniform changes across features rather than applying a single global distortion pattern. Larger shifts appear for designated privacy-sensitive components, whereas several utility-related characteristics remain closer to their original values. This descriptive result is consistent with the feature-conditioned mapping, perturbation, recovery, and calibration design.

Table 1 shows that SCEM retains competitive forecasting-oriented utility across the three datasets. Its performance remains equal or close to that of the original-data reference on Weather and Exchange Rate and is substantially better than DP and AGN on several ETTm1 metrics. Although k-nTS is competitive on selected metrics, the results, together with Table 2, support a favorable privacy–utility balance rather than uniformly optimal forecasting accuracy.

Table 1. Forecasting-oriented utility results based on LSTM prediction metrics.

Method	Weather				ETTm1				Exchange Rate			
	R^2	MAE	RMSE	MAPE (%)	R^2	MAE	RMSE	MAPE (%)	R^2	MAE	RMSE	MAPE (%)
Original	0.79	4.95	7.56	1.11	0.90	0.74	1.07	14.30	0.85	0.02	0.03	2.75
SCEM	0.79	5.39	7.93	1.24	0.89	1.28	1.29	4.80	0.85	0.02	0.03	2.85
DP	0.69	5.89	9.14	1.32	0.40	11.90	15.60	95.90	0.02	0.13	0.17	19.10
k-nTS	0.66	6.46	9.47	1.45	0.78	0.72	1.13	5.40	0.71	0.01	0.01	1.49
AGN	0.64	6.72	9.73	1.51	0.23	4.28	5.37	7.50	0.54	0.05	0.06	3.68

Table 2. Attack success rate at the common numerical operating point of 10% (ASR@10%) under the evaluated feature inference protocol.

	Weather				ETTm1				Exchange Rate			
	SCEM	DP	k-nTS	AGN	SCEM	DP	k-nTS	AGN	SCEM	DP	k-nTS	AGN
Freq	9.52	57.14	100.00	95.24	7.41	81.48	96.30	100.00	12.50	25.00	95.83	91.67
Mean	5	40	100	65	15	40	20	60	0	65	80	85
Var	0	15	40	5	0	15	20	0	10	0	25	0
Median	15	45	0	45	0	15	20	10	10	20	0	20
Skew	0	0	80	0	0	0	10	5	0	5	40	0
Kurt	0	10	65	0	0	5	5	0	0	0	10	0

4.4. Privacy Attack Experiment

The privacy experiment evaluates whether predefined statistical and spectral functionals remain recoverable from the released sequence under the one-shot feature inference threat model in Section 3.1. Statistical feature inference targets the mean, variance, median, skewness, and kurtosis, while frequency-domain inference tests whether the designated spectral component remains prominent after protection. The primary result is the ASR@10%, with lower values indicating lower recoverability under this protocol.

Table 2 shows that the evaluated SCEM instantiations achieve ASR@10% values no higher than 15% across the 18 reported feature–dataset pairs, with an unweighted mean of 4.69%. Some baselines obtain equal or lower ASR values on individual pairs; therefore, the privacy results should be interpreted jointly with the forecasting utility in Table 1. The reduced recoverability is consistent with the intended roles of spectral contraction, variance redistribution, quantile adjustment, and moment shaping.

4.5. Feature-Selective Privacy Sensitivity Analysis

The sensitivity analysis varies the reference-oriented coefficient over $\alpha \in [0.1, 0.9]$ to examine gradual feature-level responses. Here, α is a common sensitivity analysis index

rather than an identical physical parameter shared by all SCEM instantiations, and smaller values represent stronger movement toward the reference state.

Figure 6 presents an α-sensitivity scan rather than a pass–fail test against fixed privacy target values. The purpose of this analysis is to examine how the protected statistical and spectral characteristics respond when the privacy control parameter varies from 0.1 to 0.9. The results show that different sensitive characteristics exhibit feature-specific response trends: the statistical descriptors change gradually without abrupt oscillations, while the sensitive frequency amplitude varies systematically with respect to the selected background reference.

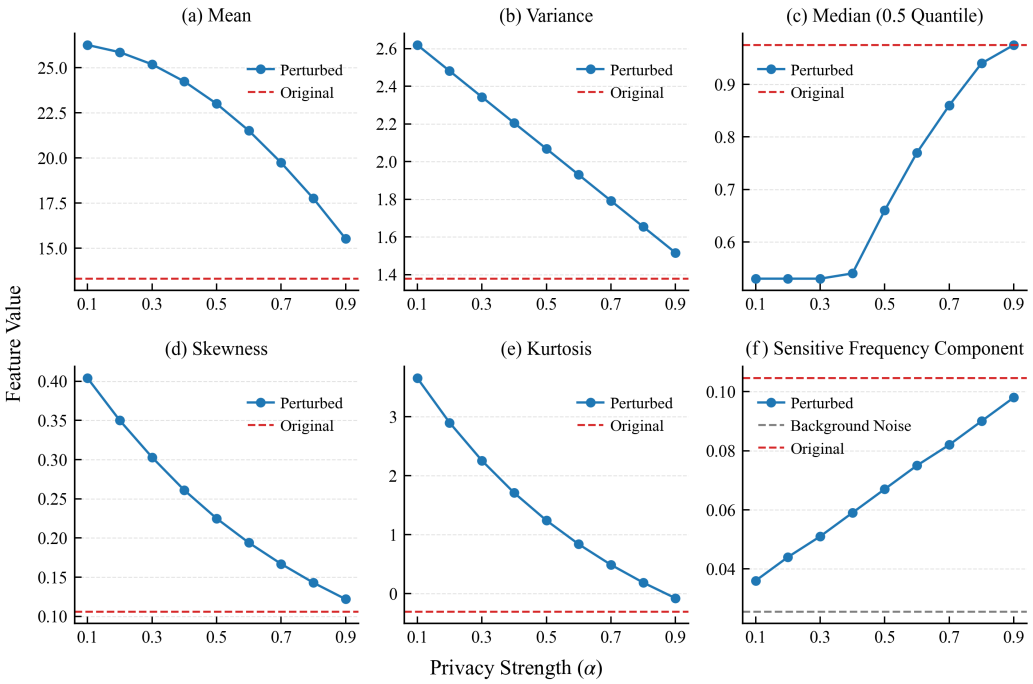

Figure 6. Sensitivity analysis of statistical and frequency characteristics under different values of the reference-oriented privacy control parameter α. The curves show how each protected feature responds to gradually changing privacy strength. The horizontal reference lines indicate the original statistical values and the background noise level in the frequency subplot, serving as comparative baselines rather than hard pass–fail privacy thresholds.

The sensitivity results therefore support the controllability of SCEM through feature-specific responses to the shared privacy control principle. In this analysis, the reference lines are used only as comparative baselines for interpreting the response trends under different privacy strengths, not as universal thresholds for judging whether protection is successful.

4.6. Overall Privacy–Utility Trade-Off

To summarize privacy and utility jointly, the privacy coordinate is calculated as the unweighted mean ASR@10% over the six attacks in Table 2. The utility coordinate is the non-negative relative RMSE increase over the original-data reference:

$$\Delta_{\mathrm{RMSE}} = \max\left(0, \frac{\mathrm{RMSE}_{\mathrm{method}} - \mathrm{RMSE}_{\mathrm{original}}}{\mathrm{RMSE}_{\mathrm{original}} + \epsilon}\right) \times 100\%. \tag{38}$$

The value is clipped at zero because the ordinate measures utility loss, while the equal-weight average provides only an aggregate view alongside the feature-specific results.

Figure 7 places SCEM in a favorable low-leakage and low-to-moderate utility loss region under the adopted aggregate metrics.

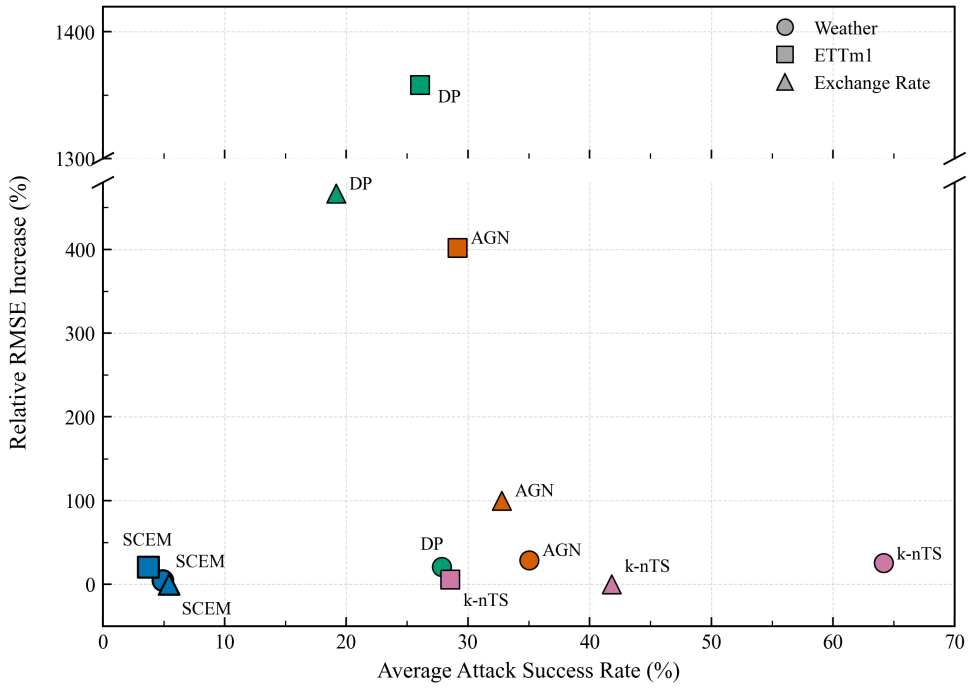

Figure 7. Aggregate privacy–utility trade-off. The x-axis is the unweighted mean ASR@10% over six evaluated attacks, and the y-axis is the non-negative relative RMSE increase over the original-data reference. Lower-left positions indicate a more favorable aggregate trade-off.

4.7. Implications for Intelligent Sensor Data Sharing

SCEM can serve as a feature-selective preprocessing layer before time-series data are released to external analysts or cloud-based models. In smart metering, industrial monitoring, environmental sensing, and related IoT settings, its operator interface allows data publishers to modify selected sensitive functionals while constraining utility-related characteristics within prescribed tolerances.

4.8. Scope and Limitations

The current evaluation focuses on deterministic, channel-wise SCEM instantiations applied to one predefined functional at a time under direct feature inference attacks. Joint cross-functional and cross-channel protection, randomized instantiations against repeated-observation or surrogate-model attacks, formal individual-level privacy guarantees, and learned attack models require additional mechanism-specific evaluation. SCEM does not provide an individual-level DP guarantee in the reported instantiations, and the statistic-level Laplace baseline is included as a heuristic reference. Device-level latency, energy consumption, and high-dimensional stress testing also remain topics for future deployment studies. Accordingly, the empirical conclusions are limited to the reported datasets, feature-specific instantiations, forecasting evaluator, and attack protocol.

5. Conclusions and Future Work

This paper presents SCEM, a unified operator interface framework for feature-selective time-series privacy protection. Its shared mapping–perturbation–recovery–calibration interface supports spectral amplitudes, mean, variance, quantiles, and high-order moments while constraining utility-related structures.

The framework analysis establishes conditional properties including non-negative privacy energy representation, output feasibility, bounded behavior, and parameter stability. Across ETTm1, Weather, and Exchange Rate, SCEM achieved ASR@10% values no higher than 15%, with an unweighted mean of 4.69% over the 18 reported pairs, while maintaining competitive forecasting-oriented utility under the evaluated protocol. Future work will examine joint functionals, randomized instantiations, formal privacy integration, learned attacks, and edge-oriented deployment.

Author Contributions: Conceptualization, S.J. and Q.P.; methodology, S.J.; investigation, S.Z.; validation, S.Z.; visualization, S.Z.; writing—original draft preparation, S.J.; writing—review and editing, Q.P. and H.S.; supervision, H.S.; project administration, H.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The benchmark datasets used in this study are publicly available. The implementation code for reproducing the main results is publicly available at https://github.com/jshe6427-cpu/SCEM-time-series-privacy (accessed on 12 July 2026).

Conflicts of Interest: Author Shankui Zheng was employed by the company Henan Aerospace Rocket Co., Ltd. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Kua, J.; Hossain, M.B.; Natgunanathan, I.; Xiang, Y. Privacy Preservation in Smart Meters: Current Status, Challenges and Future Directions. *Sensors* **2023**, *23*, 3697. [CrossRef] [PubMed]
2. Asghar, M.R.; Dán, G.; Miorandi, D.; Chlamtac, I. Smart Meter Data Privacy: A Survey. *IEEE Commun. Surv. Tutor.* **2017**, *19*, 2820–2835. [CrossRef]
3. Ali, M.; Suchismita, M.; Ali, S.S.; Choi, B.J. Privacy-Preserving Machine Learning for IoT-Integrated Smart Grids: Recent Advances, Opportunities, and Challenges. *Energies* **2025**, *18*, 2515. [CrossRef]
4. Voyez, A.; Allard, T.; Avoine, G.; Cauchois, P.; Fromont, E.; Simonin, M. The Privacy Cost of Fine-Grained Electrical Consumption Data. *Sci. Rep.* **2025**, *15*, 17391. [CrossRef] [PubMed]
5. Shanmugarasa, Y.; Chamikara, M.A.P.; Paik, H.Y.; Kanhere, S.S.; Zhu, L. Local Differential Privacy for Smart Meter Data Sharing. *arXiv* **2023**, arXiv:2311.04544.
6. Xiao, K.; Li, Y.; Dong, Y.; Yang, W.; Yao, B.; Chen, L. Nonlinear time domain and multi-scale frequency domain feature fusion for time series forecasting. *Sci. Rep.* **2025**, *15*, 30024. [CrossRef] [PubMed]
7. Yuan, X.; Chen, J.; An, D.; Wan, J. Time Series Anomaly Detection Based on Entropy-Sparsified Time-Frequency Fusion and MsRwGWO Meta-Optimization. *Int. J. Adv. Comput. Sci. Appl.* **2026**, *17*, 994–1009. [CrossRef]
8. Fan, L.; Xiong, L. Adaptively Sharing Time-Series with Differential Privacy. *arXiv* **2012**, arXiv:1202.3461.
9. McElroy, T.; Roy, A.; Hore, G. FLIP: A Utility Preserving Privacy Mechanism for Time Series. *J. Mach. Learn. Res.* **2023**, *24*, 1–29.
10. Song, S.; Wang, Y.; Chaudhuri, K. Pufferfish Privacy Mechanisms for Correlated Data. In *Proceedings of the 2017 ACM International Conference on Management of Data*; Association for Computing Machinery: New York, NY, USA, 2017; pp. 1291–1306. [CrossRef]
11. Song, S.; Chaudhuri, K. Composition Properties of Inferential Privacy for Time-Series Data. *arXiv* **2017**, arXiv:1707.02702.
12. Nazir, I.; Mushtaq, N.; Amin, W. Smart Grid Systems: Addressing Privacy Threats, Security Vulnerabilities, and Demand–Supply Balance (A Review). *Energies* **2025**, *18*, 5076. [CrossRef]
13. Radovanovic, D.; Fernandez, J.D.; Schirl, M.; Eibl, G.; Unterweger, A.; Menci, S.P. Inferring the Hidden: Privacy Risks of Microaggregation in Smart Meter Data. *SIGENERGY Energy Inform. Rev.* **2025**, *5*, 179–195. [CrossRef]
14. Leukam Lako, F.; Lajoie-Mazenc, P.; Laurent, M. Privacy-Preserving Publication of Time-Series Data in Smart Grid. *Secur. Commun. Netw.* **2021**, *2021*, 6643566. [CrossRef]
15. Wang, H.; Zhang, J.; Lu, C.; Wu, C. Privacy Preserving in Non-Intrusive Load Monitoring: A Differential Privacy Perspective. *arXiv* **2020**, arXiv:2011.06205.
16. Lalos, A.S.; Vlachos, E.; Berberidis, K.; Fournaris, A.P.; Koulamas, C. Privacy Preservation in Industrial IoT via Fast Adaptive Correlation Matrix Completion. *IEEE Trans. Ind. Inform.* **2020**, *16*, 7765–7773. [CrossRef]

17. Tian, M.; Chen, B.; Guo, A.; Jiang, S.; Zhang, A.R. Reliable Generation of Privacy-preserving Synthetic Electronic Health Record Time Series via Diffusion Models. *arXiv* **2024**, arXiv:2310.15290.
18. Tian, Y.; Shi, Y.; Zhang, Y. Differentially Private System for Residential Energy Management via Markov Decision Process. *Tsinghua Sci. Technol.* **2026**, *31*, 2694–2706. [CrossRef]
19. Westrich, B. Privacy-Preserving Analytics for Smart Meter (AMI) Data: A Hybrid Approach to Comply with CPUC Privacy Regulations. *arXiv* **2025**, arXiv:2505.08237.
20. Rougé, P.; Moukadem, A.; Dieterlen, A.; Boutet, A.; Frindel, C. Anonymizing motion sensor data through time-frequency domain. In *Proceedings of the 2021 IEEE 31st International Workshop on Machine Learning for Signal Processing (MLSP)*; IEEE: Piscataway, NJ, USA, 2021; pp. 1–6. [CrossRef]
21. Fioretto, F.; Van Hentenryck, P. OptStream: Releasing Time Series Privately. *J. Artif. Intell. Res.* **2019**, *65*, 423–456. [CrossRef]
22. Zhang, X.; Khalili, M.M.; Liu, M. Differentially Private Real-Time Release of Sequential Data. *ACM Trans. Priv. Secur.* **2022**, *26*, 1–29. [CrossRef]
23. Li, X.; Qin, Z.; Ren, K.; Gong, C.; Feng, S.; Hong, Y.; Wang, T. Delay-Allowed Differentially Private Data Stream Release. In *Proceedings of the Network and Distributed System Security Symposium*; Internet Society: Reston, VA, USA, 2025. [CrossRef]
24. Rastogi, V.; Nath, S. Differentially Private Aggregation of Distributed Time-Series with Transformation and Encryption. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*; Association for Computing Machinery: New York, NY, USA, 2010; pp. 735–746. [CrossRef]
25. Joye, M.; Libert, B. A Scalable Scheme for Privacy-Preserving Aggregation of Time-Series Data. In *Financial Cryptography and Data Security*; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 2013; Volume 7859, pp. 111–125. [CrossRef]
26. Mao, Y.; Ye, Q.; Wang, Q.; Hu, H. Differential Privacy for Time Series: A Survey. *IEEE Data Eng. Bull.* **2024**, *48*, 67–92.
27. Schuchardt, J.; Dalirrooyfard, M.; Guzelkabaagac, J.; Schneider, A.; Nevmyvaka, Y.; Günnemann, S. Privacy Amplification by Structured Subsampling for Deep Differentially Private Time Series Forecasting. *arXiv* **2025**, arXiv:2502.02410.
28. Ma, X.; Hong, X.; Li, W.; Lu, S. Energy-Aware Pattern Disentanglement: A Generalizable Pattern Assisted Architecture for Multi-task Time Series Analysis. *arXiv* **2025**, arXiv:2504.14209.
29. Dwork, C. Differential Privacy. In *Encyclopedia of Cryptography, Security and Privacy*; Jajodia, S., Samarati, P., Yung, M., Eds.; Springer Nature: Cham, Switzerland, 2025; pp. 649–652. [CrossRef]
30. Parker, K.; Hale, M.; Barooah, P. Spectral Differential Privacy: Application to Smart Meter Data. *IEEE Internet Things J.* **2022**, *9*, 4987–4996. [CrossRef]
31. Mao, Y.; Ye, Q.; Hu, H.; Wang, Q.; Huang, K. PrivShape: Extracting Shapes in Time Series under User-Level Local Differential Privacy. In *Proceedings of the 2024 IEEE 40th International Conference on Data Engineering*; IEEE: Piscataway, NJ, USA, 2024; pp. 1739–1751. [CrossRef]
32. Bale, C.D.; Schneider, M.J.; Lee, J. Can we protect time series data while maintaining accurate forecasts? *Int. J. Forecast.* **2026**, *42*, 297–314. [CrossRef]
33. Liu, L.; Li, J.; Lv, J.; Wang, J.; Zhao, S.; Lu, Q. Privacy-Preserving and Secure Industrial Big Data Analytics: A Survey and the Research Framework. *IEEE Internet Things J.* **2024**, *11*, 18976–18999. [CrossRef]
34. Gillenwater, J.; Joseph, M.; Kulesza, A. Differentially Private Quantiles. In *Proceedings of the 38th International Conference on Machine Learning. PMLR*; JMLR: Cambridge MA, USA, 2021; Volume 139, pp. 3713–3722.
35. Erdemir, E.; Dragotti, P.L.; Gündüz, D. Active Privacy-Utility Trade-Off against Inference in Time-Series Data Sharing. *IEEE J. Sel. Areas Inf. Theory* **2023**, *4*, 159–173. [CrossRef]
36. Kifer, D.; Machanavajjhala, A. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data, SIGMOD '11*; ACM: New York, NY, USA, 2011; pp. 193–204. [CrossRef]
37. Zerveas, G.; Jayaraman, S.; Patel, D.; Bhamidipaty, A.; Eickhoff, C. A Transformer-based Framework for Multivariate Time Series Representation Learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD '21*; ACM: New York, NY, USA, 2021; pp. 2114–2124. [CrossRef]
38. Bai, S.; Kolter, J.Z.; Koltun, V. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. *arXiv* **2018**, arXiv:1803.01271.
39. Abbas, S.R.; Abbas, Z.; Zahir, A.; Lee, S.W. Federated Learning in Smart Healthcare: A Comprehensive Review on Privacy, Security, and Predictive Analytics with IoT Integration. *Healthcare* **2024**, *12*, 2587. [CrossRef] [PubMed]
40. Catarino, A.; Melo, R.; Abreu, R.; Cruz, L. Evaluating Privacy-Utility Tradeoffs in Synthetic Smart Grid Data. *arXiv* **2025**, arXiv:2506.11026.
41. Chai, S.; Chadney, G.; Avery, C.; Grunewald, P.; Van Hentenryck, P.; Donti, P.L. Defining "Good": Evaluation Framework for Synthetic Smart Meter Data. *arXiv* **2024**, arXiv:2407.11785.
42. Lin, Y.; Wei, Y.; Chen, D.; Li, Y.; Erkan, U.; Toktas, A.; Gao, S.; Zhang, Y. Cryptanalysis and Improvement of a Video Cryptosystem via Chaos and S-Box. *ACM Trans. Multimed. Comput. Commun. Appl.* **2026**, *22*, 1–28 . [CrossRef]

43. Tillman, R.E.; Balch, T.; Veloso, M. Privacy-Preserving Energy-Based Generative Models for Marginal Distribution Protection. 2023. Available online: https://openreview.net/forum?id=vTsfup5ll6 (accessed on 12 July 2026).
44. Lai, G.; Chang, W.C.; Yang, Y.; Liu, H. Modeling Long- and Short-Term Temporal Patterns with Deep Neural Networks. In *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR '18*; ACM: New York, NY, USA, 2018; pp. 95–104. [CrossRef]
45. Zhou, H.; Zhang, S.; Peng, J.; Zhang, S.; Li, J.; Xiong, H.; Zhang, W. Informer: Beyond efficient transformer for long sequence time-series forecasting. *Proc. AAAI Conf. Artif. Intell.* **2021**, *35*, 11106–11115. [CrossRef]
46. Hochreiter, S.; Schmidhuber, J. Long Short-Term Memory. *Neural Comput.* **1997**, *9*, 1735–1780. [CrossRef] [PubMed]