SAWS: FORECASTING INACCURATE TIME SERIES WITH SHARPNESS-AWARE WEAK SUPERVISION

Anonymous authors
Paper under double-blind review

ABSTRACT

Inaccurate observations degrade time-series forecasting twice: inaccurate historical observations obscure the dynamics to predict, and inaccurate future observations used as labels degrade supervision. While applying means of observation accuracy improvement can recover useful information, treating resulting improved observations as ground-truth observations can amplify any errors that may have been introduced. We introduce SAWS, a forecasting framework that combines vector-quantized historical retrieval with sharpness-aware weak supervision. A past-only codebook compresses similar historical contexts and their observed dynamics into coordinate-supported prototypes. Retrieved prototypes correct inputs and supply weak targets without accessing the current forecast horizon. A detached, sample-level sharpness score measures the sensitivity of an exponential-moving-average forecaster on observed targets; together with prototype support and dispersion, it controls the influence of weak supervision. Multiscale temporal features preserve the distinction between short-term dynamics and longer-term context without requiring dense attribute-graph propagation. We characterize the bias–variance tradeoff of prototype completion and establish when sharpness weighting reduces weak-label error, explicitly contending with selection bias and confidence miscalibration. A leakage-controlled experimental evaluation offers evidence that SAWS is able to offer better accuracy and reliability on inaccurate data than are state-of-the-art baselines.

1 INTRODUCTION

Time-series forecasting supports decisions in energy (Qiu et al., 2026b), transportation (Zhao et al., 2023), environmental monitoring (Qiu et al., 2026a), and healthcare (Jing et al., 2025). In these settings, sensor outages, asynchronous measurements, and delayed reporting create inaccurate training sequences. Inaccurate data affects both the historical context used to predict the future and the future observations used as labels. A forecaster trained only on visible labels may discard substantial supervision, while a forecaster trained on completed inaccurate labels may learn the completion model's mistakes. The central problem is therefore to recover useful historical information while controlling how strongly uncertain reconstructions influence learning.

Recent imputation methods model temporal and cross-variable dependencies with attention, diffusion, or graph propagation (Du et al., 2023; Tashiro et al., 2021; Cini et al., 2022). However, accurate reconstruction alone does not establish that completed values improve a downstream forecasting objective (Yang et al., 2025). MH-GIN (Liu et al., n.d.) motivates a complementary structural insight: heterogeneous variables change at different rates, so useful dependencies should be represented at multiple temporal scales. We retain this multiscale perspective while targeting the task from within-sequence imputation to directly forecasting a strictly unobserved horizon. This change requires explicit treatment of weak future labels, temporal information boundaries, and the reliability of the resulting training signal.

Challenge 1: How can inaccurate histories and labels be completed without using the future being predicted? Repeated operating regimes provide similar historical contexts with informative continuations. Directly matching every context against every historical window is expensive, while an unrestricted retrieval bank can leak the current target through overlapping windows. Furthermore, different prototype coordinates may have very different observation support. A useful completion

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

mechanism must therefore compress historical similarity, preserve coordinate-wise support, and keep target information outside the retrieval key.

Challenge 2: How should a forecaster learn from uncertain completions? Prototype similarity does not guarantee predictive correctness. A common regime may have several possible futures, and a well-supported prototype can still average incompatible continuations. Uniformly weighting these weak labels can bias the forecaster. Conversely, indiscriminately discarding difficult examples can remove rare but important regimes. The weighting mechanism must expose its assumptions and preserve a direct learning signal from observed labels.

To address Challenge 1, we introduce a vector-quantization-based (VQ) completion module. A codebook learned from historical contexts stores observed-coordinate statistics for both inputs and their continuations. A query retrieves a small set of similar prototypes and aggregates only coordinates with sufficient evidence. Chronological, purged reference prefixes prevent self-retrieval and target leakage. To address Challenge 2, we use a one-sided local sharpness probe at an exponential-moving-average teacher to form detached training weights. The probe is computed on observed labels, never on the weak targets it is intended to assess. Prototype support and dispersion provide a separate reliability signal, and an unweighted observed-label anchor limits the loss of coverage. Our contributions are threefold.

First, we formulate a joint input-and-label completion framework for forecasting under inaccurate observations, with explicit temporal safeguards and coordinate-supported VQ retrieval. Second, we develop sharpness-aware weak supervision that distinguishes observed evidence, prototype reliability, and model sensitivity. Third, we provide conditional guarantees and a controlled evaluation design that tests accuracy, calibration, and end-to-end efficiency. We do not assume that sharpness is a calibrated probability or that quantization is always cheaper than a sparse graph; both are addressed through explicit analysis and matched comparisons.

2 Problem formulation

Let $\mathbf{X}_i \in \mathbb{R}^{L\times C}$ and $\mathbf{Y}_i \in \mathbb{R}^{H\times C}$ be the latent history and future at forecast origin t_i. The observed arrays are

$$\mathbf{X}_i^o = \mathbf{M}_i \odot (\mathbf{X}_i + \boldsymbol{\xi}_i), \qquad \mathbf{Y}_i^o = \mathbf{O}_i \odot (\mathbf{Y}_i + \boldsymbol{\zeta}_i), \tag{1}$$

where $\mathbf{M}_i, \mathbf{O}_i$ are binary masks and $\boldsymbol{\xi}_i, \boldsymbol{\zeta}_i$ allow observed-value corruption. Missing entries are numerically zero-filled but always accompanied by masks. All preprocessing statistics are fitted on a designated initial training prefix and then frozen. The main task forecasts continuous variables; cyclic and categorical covariates may use type-specific encodings when available.

The forecasting risk is the error against the latent future, $R(\theta) = \mathbb{E}[\|f_\theta(\mathbf{X}_i^o, \mathbf{M}_i) - \mathbf{Y}_i\|_F^2/(HC)]$. In controlled experiments, clean evaluator-retained targets are stored before corruption; native data permit evaluation only against available recorded outcomes, which need not equal latent clean values. Naturally unobserved targets have no ground truth and are excluded from numerical evaluation. At inference, neither $\mathbf{Y}_i^o$ nor $\mathbf{O}_i$ is available to the predictor or retrieval key. Training uses observed future labels and weak labels generated from eligible historical reference windows. No guarantee of missing-not-at-random identification is assumed.

3 The proposed framework

3.1 Information boundaries and data preparation

For a raw multivariate series, the context–target support of origin t is $[t - L + 1, t + H]$. Training, validation, and test target intervals are disjoint and chronological. Contexts may use earlier released observations, including earlier observations within a test period, when evaluating rolling forecasting. Such observations update the query context only; they never update the frozen prototype archive. Known calendar covariates are permitted, while realized future sensor values, future observation masks, and future sharpness-score availability are not.

The initial training region is divided into a dictionary-fitting prefix $\mathcal{D}_0$ and a later reference-seed prefix $\mathcal{D}_1$. Prefix $\mathcal{D}_0$ fits the context encoder, codewords, and preprocessing statistics; its windows

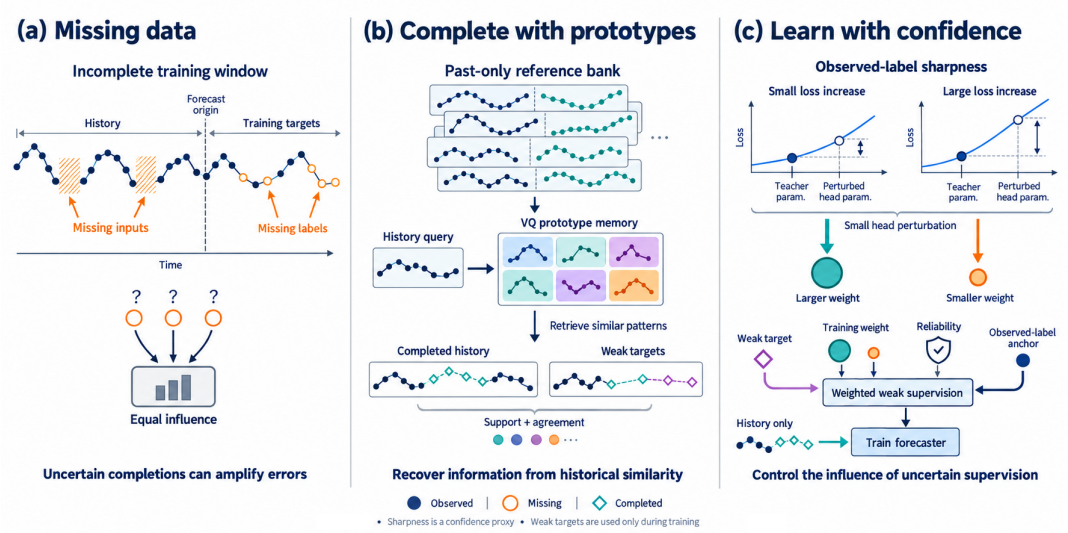


Figure 1: Overview of SAWS. Historical prototypes provide complementary information; observed-label sensitivity controls its training influence. The forecast predictor receives only history-derived quantities. Future labels and their availability masks are used exclusively in training losses and confidence probes.

are excluded from the observed-value prototype payload statistics. Subsequent reference prefixes contain eligible windows from $\mathcal{D}_1$ and later training blocks. The forecaster may use observed labels in the initial regions for warm-up, but no weak targets are generated there without an eligible earlier reference. All baseline training-data budgets include the same warm-up observations.

Let b be the first origin of a training query block. Its cache admits a reference origin t_j only when $t_j + H < b - L + 1$. This block-level criterion is conservative for later queries in the same block and can be implemented with a single prefix cache. It removes overlapping values before sums, counts, and sums of squares are formed. The dictionary itself is not refitted on a held-out query block. The final evaluation cache uses only reference windows ending before the earliest validation input; it is then frozen for both validation and test. An ablation may use more recent legally available training references, but must disclose any relaxation of the stronger nonoverlap rule.

Normalization and unsupported channels. For each continuous channel, estimate the mean μ_c^0 and standard deviation σ_c^0 from observed values in $\mathcal{D}_0$. Normalize with $\max(\sigma_c^0, \epsilon_\sigma)$. For a channel with no observed prefix values, use the declared identity convention $(\mu_c^0, \sigma_c^0) = (0, 1)$ and an unsupported-statistics flag; report these channels separately. A constant channel uses the variance floor. A missing coordinate with no supported prototype uses the standardized training mean (zero, or the identity fallback), a zero prototype-reliability value, and its missingness indicator. No reconstructed value is silently relabeled as an observation.

3.2 Mask-aware multiscale encoder and forecasting backbone

For each history position, concatenate normalized observed values, observation masks, and clipped log elapsed times. Known static or calendar covariates may be appended. Cyclic quantities are represented by sine/cosine pairs; categorical covariates use embeddings with an explicit missing category. The main forecasting targets are continuous, and heterogeneous-output classification would require an additional output-specific loss beyond the present formulation.

The context encoder first maps per-step features to width d_e. Parallel depthwise temporal convolution branches use declared kernel widths and dilation factors, each followed by a pointwise channel mixer, normalization, and a nonlinear activation. Branches are pooled over observed-history positions and projected to a key of dimension d. The key includes observation-coverage summaries to distinguish data absence from a true zero signal. Pooling returns a defined zero representation plus a

coverage flag when a channel has no observations. All convolution support lies within the available input window; a branch may not read a future target.

The forecast backbone applies the same multiscale construction to the completed values and their auxiliary history channels, but has separately learned parameters. It preserves the temporal representations needed by a linear multihorizon output head, whose outputs reshape to $H \times C$. The output head has a distinct row per horizon–channel coordinate. Backbone-transfer experiments replace this module with DLinear, PatchTST, iTransformer, TimeMixer, or TimesNet (Zeng et al., 2023; Nie et al., 2023; Liu et al., 2024; Wang et al., 2024; Wu et al., 2022), while retaining the same masks, historical completion, and training objective. RevIN (Kim et al., 2022), if used, must compute instance statistics from observed context only and be applied consistently to all matched variants.

3.3 Past-only vector-quantized completion

Context keys and temporal eligibility. A mask-aware multiscale encoder produces $\mathbf{z}_i = E_\phi(\mathbf{X}_i^o, \mathbf{M}_i, \mathbf{\Delta}_i) \in \mathbb{R}^d$, where $\mathbf{\Delta}_i$ contains elapsed times since observed history values. The encoder and K codewords $\{\mathbf{c}_k\}$ are learned on an initial training prefix by artificially hiding observed *context* entries, using a reconstruction objective, straight-through quantization, and exponential-moving-average codeword updates (van den Oord et al., 2017). They are then frozen. This dictionary-fitting prefix is excluded from prototype payload estimation; a subsequent reference prefix supplies the first payload statistics.

For each subsequent training block, a reference cache contains only earlier windows whose *complete* input–target intervals end before the block's earliest input time. On a single series, a reference origin t_j is eligible for query i only if $t_j + H < t_i - L + 1$. Prefix-level purging occurs before aggregation; filtering raw indices after compression would not remove their contribution to a prototype. Validation and test use a frozen, purged training-only cache. Appendix 3.1 specifies the boundary rules.

Observed-coordinate prototype statistics. Assign a reference window by $k(j) = \arg\min_k \|\mathbf{z}_j - \mathbf{c}_k\|_2^2$. For a flattened input or future coordinate q, let v_{jq} be its observed value and b_{jq} its observation indicator. Each code stores

$$n_{kq} = \sum_{j:k(j)=k} b_{jq}, \qquad \mu_{kq} = \frac{\sum_{j:k(j)=k} b_{jq} v_{jq}}{n_{kq}}, \qquad \sigma_{kq}^2 = \frac{\sum_{j:k(j)=k} b_{jq}(v_{jq} - \mu_{kq})^2}{n_{kq}}. \quad (2)$$

Undefined coordinates with $n_{kq} = 0$ remain unsupported. Neither imputed inputs nor pseudo-labels contribute to these statistics. Future values populate prototype *payloads*, but never context keys.

For query i, retrieve the m nearest codewords $\mathcal{K}_i$ and compute

$$a_{ik} = \frac{\exp(-\|\mathbf{z}_i - \mathbf{c}_k\|_2^2 / T_v)}{\sum_{\ell \in \mathcal{K}_i} \exp(-\|\mathbf{z}_i - \mathbf{c}_\ell\|_2^2 / T_v)}, \qquad \bar{a}_{ikq} = \frac{a_{ik} \mathbf{1}[n_{kq} \geq n_{\min}]}{\sum_{\ell \in \mathcal{K}_i} a_{i\ell} \mathbf{1}[n_{\ell q} \geq n_{\min}]}. \quad (3)$$

All coordinate sums below are restricted to $\mathcal{K}_{iq} = \{k \in \mathcal{K}_i : n_{kq} \geq n_{\min} \geq 1\}$. If this set is empty, we set $r_{iq} = 0$ and skip its mean, variance, and effective-support calculations. Otherwise, $\widetilde{v}_{iq} = \sum_k \bar{a}_{ikq} \mu_{kq}$ produces completed input values and weak future targets. A dispersion proxy and effective support are

$$u_{iq} = \sum_k \bar{a}_{ikq} (\mu_{kq} - \widetilde{v}_{iq})^2 + \sum_k \frac{\bar{a}_{ikq}^2 \sigma_{kq}^2}{n_{kq}}, \qquad n_{iq}^{\text{eff}} = \left(\sum_k \frac{\bar{a}_{ikq}^2}{n_{kq}} \right)^{-1}. \quad (4)$$

The first term captures disagreement between continuations; the second estimates uncertainty in their means. We define a support-based reliability

$$r_{iq} = \mathbf{1}[\text{supported}] \, \frac{n_{iq}^{\text{eff}}}{n_{iq}^{\text{eff}} + n_0} \exp(-u_{iq}/T_u - d_i/T_d), \quad d_i = d^{-1} \sum_{k \in \mathcal{K}_i} a_{ik} \|\mathbf{z}_i - \mathbf{c}_k\|_2^2. \quad (5)$$

This is a reliability proxy, not a posterior probability. Unsupported input coordinates use the frozen training mean with reliability zero; unsupported future coordinates receive no weak loss.

3.4 DICTIONARY TRAINING AND PAYLOAD CONSTRUCTION

For a dictionary-training context, sample an artificial holdout mask $\mathbf{A}_i \leq \mathbf{M}_i$ with nonempty support and form $\mathbf{M}_i^- = \mathbf{M}_i \odot (1 - \mathbf{A}_i)$. The key $\mathbf{z}_i$ is computed from this reduced observation set. Quantization and a straight-through representation are

$$k_i = \arg\min_k \|\mathbf{z}_i - \mathbf{c}_k\|_2^2, \qquad \mathbf{z}_i^{\mathrm{st}} = \mathbf{z}_i + \mathrm{sg}(\mathbf{c}_{k_i} - \mathbf{z}_i). \tag{6}$$

A decoder D_ξ reconstructs artificially held-out observed context entries:

$$\mathcal{L}_{\mathrm{dict}} = \frac{\|\mathbf{A}_i \odot (D_\xi(\mathbf{z}_i^{\mathrm{st}}) - \mathbf{X}_i^o)\|_F^2}{|\mathbf{A}_i|} + \lambda_{\mathrm{commit}} \|\mathbf{z}_i - \mathrm{sg}(\mathbf{c}_{k_i})\|_2^2. \tag{7}$$

This loss uses neither unobserved context values nor future labels. The codewords follow EMA count and sum updates,

$$\begin{aligned} N_k &\leftarrow \nu N_k + (1-\nu) \sum_i \mathbf{1}[k_i = k], \\ S_k &\leftarrow \nu S_k + (1-\nu) \sum_{i:k_i=k} \mathbf{z}_i, \qquad \mathbf{c}_k \leftarrow S_k / \max(N_k, \epsilon_N). \end{aligned} \tag{8}$$

Dead codewords may be reinitialized from supported $\mathcal{D}_0$ keys during dictionary fitting only. After fitting, freeze E_ϕ and all $\mathbf{c}_k$. Reference contexts are then encoded with their actual masks, assigned to the frozen cells, and summarized using only their observed values. For each coordinate, store a count, sum, and sum of squares. Prefix snapshots can be maintained incrementally; training-only weak completions are cached by query block. No pseudo-target is added to the archive, preventing recursive reinforcement of generated values.

Equation 4 uses an empirical dispersion proxy. Its within-cell variance estimate is not claimed to be an unbiased estimator under arbitrary observation mechanisms. Coordinate sums range only over cells with $n_{kq} \geq n_{\min}$. Before computing any mean or variance, an empty support set returns the explicit unsupported branch. This avoids interpreting zero times an undefined statistic as a valid contribution.

3.5 FORECASTING WITH COMPLETED HISTORIES

Observed inputs are retained:

$$\mathbf{X}_i^c = \mathbf{M}_i \odot \mathbf{X}_i^o + (1 - \mathbf{M}_i) \odot \widetilde{\mathbf{X}}_i, \qquad \widehat{\mathbf{Y}}_i = f_\theta(\mathbf{X}_i^c, \mathbf{M}_i, \mathbf{R}_i^x, \mathbf{\Delta}_i). \tag{9}$$

Here $\mathbf{R}_i^x$ contains prototype reliability only for missing context coordinates. The default forecaster uses parallel dilated depthwise temporal convolutions, learned channel mixing, and a linear multi-horizon head. The scales describe multiple history receptive fields rather than fixed AIS-specific attribute groups. Completed values, masks, and reliability remain distinct input channels. Appendix 3.2 defines the architecture; stronger forecasting backbones can use the same completion and weighting interface. Weak future targets are never fed into the predictor.

3.6 SHARPNESS-AWARE TRAINING WEIGHTS

Let $\bar{\theta}$ be an exponential-moving-average (EMA) teacher, and ϑ its forecasting-head parameters. For a window with at least one observed future coordinate, define

$$\ell_i^o(\bar{\theta}) = \frac{1}{2|\mathbf{O}_i|} \sum_{q:O_{iq}=1} \left(f_{\bar{\theta}}(\mathbf{X}_i^c, \ldots)_q - Y_{iq}^o\right)^2, \qquad \mathbf{g}_i = \nabla_\vartheta \ell_i^o(\bar{\theta}). \tag{10}$$

For each head parameter block b, set $A_b = (\|\vartheta_b\|_2 / \sqrt{p_b} + \epsilon_a) I$, where p_b is its size. With this scale matrix held fixed during a probe,

$$\boldsymbol{\delta}_i = \rho \frac{\mathbf{A}^2 \mathbf{g}_i}{\|\mathbf{A}\mathbf{g}_i\|_2 + \epsilon_g}, \qquad s_i = \frac{[\ell_i^o(\bar{\theta} + \boldsymbol{\delta}_i) - \ell_i^o(\bar{\theta})]_+}{c_{\mathrm{loss}} + \epsilon_\ell}, \qquad w_i = \exp(-s_i / T_s). \tag{11}$$

Only the head is perturbed; backbone features are cached and detached. The shared scale c_{loss} is an EMA of the batch-median observed probe loss. We deliberately avoid division by the window's

Algorithm 1: Training and inference for SAWS

Input Chronological data, masks, context length L, horizon H, dictionary/reference prefixes, and hyperparameters.
1 Fit preprocessing, the context encoder, and VQ codewords on $\mathcal{D}_0$ with Equation 7; freeze them.
2 Build purged reference-prefix payload caches from eligible observed windows outside $\mathcal{D}_0$. Store counts, sums, and sums of squares only.
3 Retrieve supported prototypes for each later training block. Cache completed histories, weak targets, their supports, and r_{iq}.
4 Warm up the forecasting student on observed targets; initialize the EMA teacher.
5 **For each training epoch:**
6 If a score refresh is due, freeze a teacher/scale snapshot; compute one-sided head probes on observed targets; cache detached weights.
7 For each batch, compute predictions from completed histories and history-only auxiliary features.
8 Form the observed anchor and confidence-mass-preserving weak loss; omit empty terms; update the student.
9 Update the EMA teacher and the loss-scale accumulator; validate only on the allowed validation target set.
10 Freeze the selected forecaster and the final purged training-only reference cache.
Test Encode the current incomplete history, retrieve input completions, and forecast the entire horizon. No future labels, future masks, sharpness probe, or archive update is used.

own loss, which can give very large residuals spuriously small relative sharpness (Appendix B.2). Equation 11 is a one-step lower estimate of local worst-case head loss increase, not an exact full-network sharpness measure.

Scores are clipped to a declared maximum $s_{\max}$, refreshed periodically after an observed-label warm-up, and treated as constants during student optimization. Each refresh freezes one teacher snapshot, parameter scale, and shared loss scale for all windows; probes use deterministic evaluation mode. If no future label is observed, the sharpness probe is unavailable: we set the multiplicative factor to its neutral value $w_i = 1$ and rely on r_{iq} alone. This convention does not certify confidence. Score availability is a training-only variable and is never an input feature.

3.7 Observed-label anchoring and weak supervision

Let $\mathcal{B}_o$ be windows with observed labels, and let $U_{iq} = (1-O_{iq})\mathbf{1}[r_{iq} > 0]$ select supported missing targets. An observed-label anchor and a reweighted observed loss are combined as

$$\mathcal{L}_o = (1-\alpha)\frac{\sum_{i\in\mathcal{B}_o} \ell_i^o(\theta)}{|\mathcal{B}_o|} + \alpha\frac{\sum_{i\in\mathcal{B}_o} w_i \ell_i^o(\theta)}{\sum_{i\in\mathcal{B}_o} w_i}, \quad 0 \le \alpha < 1. \tag{12}$$

The weak loss is coordinate-weighted:

$$\mathcal{L}_u = \frac{\sum_{i,q} U_{iq} w_i r_{iq} \frac{1}{2}(\widehat{Y}_{iq} - \widetilde{Y}_{iq})^2}{\sum_{i,q} U_{iq}}, \qquad \mathcal{L} = \mathcal{L}_o + \lambda_u \mathcal{L}_u. \tag{13}$$

Empty terms are omitted. Dividing the weak loss by its eligible-coordinate count preserves absolute confidence mass: uniformly unreliable weak targets reduce its total influence. The weighted observed term uses stable log-weight normalization. All weights, retrieval statistics, and pseudo-targets are stop-gradient quantities. The unweighted component of Equation 12 preserves supervised influence from difficult observed windows, while λ_u is ramped up after warm-up. The whole process of SAWS is described in Algorithm 1.

3.8 Complexity

For a d-dimensional key and payload size $Q = (L+H)C$, exact codeword search and top-m payload aggregation cost $O(Kd + mQ)$ per window, after encoding. The compressed memory uses

$O(Kd + KQ)$ storage for codewords and sufficient statistics, compared with $O(N_r d + N_r Q)$ for N_r raw reference windows. It is advantageous when $K \ll N_r$. An ℓ-layer dense graph with V nodes typically incurs $O(\ell V^2 d + \ell V d^2)$, whereas sparse graphs replace V^2 by their edge count. For fixed K, m, L, H, d and scale count S, with $V = SC$, the retrieval/aggregation work is linear in C while dense graph-edge propagation is quadratic. Sparse graphs can also be linear, so no universal VQ–GNN ordering follows. Codebook fitting, prefix-cache construction, and amortized sharpness refreshes are included in the measured cost. Algorithm 1 specifies the complete procedure.

4 Analysis: when does weak supervision help?

The analysis conditions on a frozen codebook and teacher. Its conclusions concern prototype estimation and population objectives; finite-sample generalization and convergence of adaptive training require additional assumptions. Proofs and counterexamples appear in Appendix B.

Proposition 4.1 (Prototype error decomposition)**.** *For one future coordinate, suppose reference observations are ignorable within fixed VQ cells, have conditionally mean-zero corruption, and form independent reference samples independent of the query future. Let cell means and variances be μ_k, σ_k^2, observed counts be $n_k > 0$, and response-independent retrieval weights satisfy $a_k \geq 0$, $\sum_k a_k = 1$. Then*

$$\mathbb{E}[(\widetilde{Y} - Y)^2 \mid x, \{n_k\}] = \mathrm{Var}(Y \mid x) + \left(\sum_k a_k \mu_k - \mathbb{E}[Y \mid x] \right)^2 + \sum_k \frac{a_k^2 \sigma_k^2}{n_k}. \tag{14}$$

This identifies irreducible future variation, prototype approximation bias, and finite-support variance. Increasing support can reduce variance but cannot remove a biased or multimodal continuation. Dependence among reference estimates replaces the final term by a covariance quadratic form; reference–query dependence introduces additional covariance terms. Purging alone does not establish independence.

Proposition 4.2 (Calibrated sharpness weighting)**.** *Within a score-available stratum, let $e = (\widetilde{Y} - Y)^2$, $r \in [0, 1]$, and a bounded score $s \geq 0$ define $dQ_\beta / dP = re^{-\beta s} / \mathbb{E}_P[re^{-\beta s}]$, with positive denominator and $\mathbb{E}_P[e] < \infty$. If $\mathbb{E}_{Q_0}[e \mid s]$ is nondecreasing in s, then*

$$\frac{d}{d\beta} \mathbb{E}_{Q_\beta}[e] = -\,\mathrm{Cov}_{Q_\beta}(s, e) \leq 0. \tag{15}$$

The condition asserts that larger sharpness predicts larger weak-label error after prototype reliability is accounted for. It is a testable assumption, not implied by Equation 11. Entirely unlabeled windows are evaluated separately, since their sharpness score is unavailable.

Theorem 4.3 (Risk distortion and selection tradeoff)**.** *For fixed weights defining Q, assume $|f(X)|, |\widetilde{Y}|, |Y| \leq B$, and let $\varepsilon_Q = \mathbb{E}_Q[(\widetilde{Y} - Y)^2]$. If $\widehat{f}$ is an ϵ_{opt}-optimal minimizer of the weak half-squared population risk, then*

$$R_P(\widehat{f}) - \inf_{f \in \mathcal{F}} R_P(f) \leq 4B\sqrt{\varepsilon_Q} + 4B^2\, \mathrm{TV}(P, Q) + \epsilon_{\mathrm{opt}}, \tag{16}$$

where $R_P(f) = \frac{1}{2}\mathbb{E}_P[(f(X) - Y)^2]$.

Thus, reducing weak-label error can be offset by selecting an unrepresentative training population. This explains the need for an observed-label anchor and coverage diagnostics. For the actual anchored objective, let P denote eligible missing coordinates and $p = \mathbb{E}_P[wr]$. Replacing its weak targets by clean targets changes the population objective uniformly by at most $2\lambda_u p B \sqrt{\varepsilon_Q}$ (Appendix B.4). No unconditional forecast improvement or universal noise-detection guarantee is claimed.

5 Experimental protocol

Evaluation status. Numerical experiments have not yet been executed. All result cells are intentionally empty; the following protocol specifies falsifiable comparisons without claiming empirical gains.

Table 1: Comparison of forecasting performance on four datasets with various window length when missing ratio $r = 0.05$. Entries with '–' indicate the experiment can not be done due to out-of-memory issue.

Data	ℓ_L	Metric ↓	BRITS	GRU-D	Trans.	Auto.	BiTGraph	S4 (Mean)	S4 (Ffill)	S4 (Decay)	S4 (SAITS)	S4M	SAWS
Electricity	96	MAE	0.431	0.422	0.397	0.366	0.377	0.379	0.380	0.391	0.427	0.367	0.355
		MSE	0.365	0.349	0.304	0.261	0.208	0.302	0.301	0.308	0.350	0.277	0.267
	192	MAE	0.451	0.470	0.394	0.364	0.373	0.360	0.360	0.371	0.403	0.351	0.341
		MSE	0.382	0.404	0.301	0.252	0.271	0.264	0.265	0.272	0.300	0.252	0.240
	384	MAE	0.618	0.466	0.405	0.353	0.383	0.361	0.361	0.372	0.406	0.358	0.339
		MSE	0.614	0.402	0.313	0.242	0.283	0.265	0.263	0.277	0.315	0.259	0.242
	768	MAE	0.631	0.480	0.403	0.356	0.383	0.366	0.374	0.371	0.399	0.357	0.344
		MSE	0.636	0.431	0.323	0.260	0.282	0.268	0.283	0.276	0.308	0.269	0.253
ETTh1	96	MAE	0.694	0.612	0.583	0.582	0.573	0.637	0.638	0.619	0.673	0.572	0.566
		MSE	0.902	0.724	0.650	0.649	0.652	0.757	0.761	0.716	0.834	0.624	0.611
	192	MAE	0.812	0.614	0.638	0.584	0.593	0.616	0.613	0.597	0.651	0.568	0.552
		MSE	1.161	0.725	0.812	0.641	0.712	0.684	1.616	0.659	0.783	0.560	0.552
	384	MAE	0.818	0.604	0.612	0.586	0.599	0.603	0.601	0.604	0.628	0.585	0.564
		MSE	1.164	0.703	0.685	0.636	0.718	0.662	0.675	0.681	0.718	0.614	0.607
	768	MAE	0.817	0.623	0.745	0.639	0.632	0.614	0.622	0.622	0.642	0.601	0.589
		MSE	1.162	0.735	1.025	0.732	0.812	0.678	0.708	0.717	0.730	0.650	0.641
Weather	96	MAE	0.404	0.405	0.423	0.495	0.538	0.409	0.392	0.389	0.433	0.340	0.331
		MSE	0.334	0.343	0.355	0.437	0.545	0.365	0.332	0.333	0.392	0.270	0.257
	192	MAE	0.414	0.382	0.421	0.506	0.445	0.367	0.355	0.350	0.401	0.327	0.314
		MSE	0.354	0.312	0.352	0.453	0.416	0.299	0.278	0.279	0.332	0.265	0.253
	384	MAE	0.450	0.382	0.403	0.515	0.657	0.361	0.343	0.345	0.407	0.340	0.331
		MSE	0.400	0.315	0.322	0.471	0.695	0.293	0.268	0.269	0.332	0.257	0.235
	768	MAE	0.471	0.390	0.403	0.527	0.626	0.347	0.347	0.341	0.395	0.340	0.333
		MSE	0.424	0.322	0.335	0.504	0.661	0.276	0.277	0.269	0.321	0.270	0.254
Traffic	96	MAE	0.673	0.481	0.423	0.436	0.518	0.449	0.440	0.441	0.501	0.443	0.432
		MSE	1.231	0.904	0.764	0.702	0.906	0.822	0.816	0.812	0.871	0.780	0.771
	192	MAE	0.673	0.501	0.414	0.392	0.520	0.399	0.394	0.394	0.447	0.385	0.372
		MSE	1.190	0.923	0.704	0.636	0.881	0.703	0.699	0.724	0.771	0.687	0.667
	384	MAE	0.671	0.502	0.393	0.395	0.481	0.385	0.421	0.386	0.442	0.390	0.381
		MSE	1.188	0.950	0.693	0.649	0.791	0.711	0.754	0.716	0.777	0.703	0.699
	768	MAE	0.673	0.515	0.443	0.402	–	0.392	0.384	0.382	0.441	0.386	0.376
		MSE	1.202	0.969	0.753	0.669	–	0.701	0.697	0.703	0.771	0.699	0.686

Datasets and splits. The primary suite uses ETTm1 (Qiu et al., 2026a), ETTm2 (Jing et al., 2025), Electricity (Wu et al., 2026), and Weather (Qiu et al., 2026b). Dataset versions, origins, and native observation rates are recorded in Appendix C. Chronological splits are fixed before windows and corruption masks are constructed. Each target interval lies within one split; retrieval uses the purged prefix protocol of Section 3.3. We evaluate multiple context lengths and horizons, with all methods sharing the same masks, forecast origins, and validation-label budget.

Settings and metrics. Input missingness and training-label availability vary independently. The suite includes random entries, contiguous gaps, channel outages, covariate-dependent masking, value-dependent masking, and separately controlled input/label corruption. Artificially hidden values are retained only by the evaluator. MSE and MAE are computed on common known targets, normalized using frozen training-prefix statistics, and macro-averaged over channels. Native-missingness scores estimate observed-outcome risk. Paired training/corruption seeds and temporal block-bootstrap confidence intervals are required; overlapped forecast windows are not treated as independent.

Masks and repeated windows. Observation masks are generated on the raw timeline before extracting overlapping windows and reused wherever the same timestamp, channel, and observation role recur. A missing future label is never counted as observed in another window's future prototype payload. The main controlled protocol allows separate context and label-release masks, representing incomplete input acquisition and limited supervised labels; their realized observation budgets are reported separately. A stricter fully-concealed-value experiment removes a withheld measurement from every training view, inducing coupled input/label missingness. These settings must not be conflated. Evaluator-only test truth is excluded from every training, codebook, and tuning operation.

Baselines and main comparison. We compare our proposed method, SAWS, with two main groups of baseline methods: state space model (SSM) based baselines and other state-of-the-art and classical methods for handling missing data. The SSM-based baseline group includes S4 (Mean), S4 (Ffill) (Du et al., 2023), S4 (Decay), S4 (SAITS), and S4M (Jing et al., 2025). These methods impute

missing data using strategies such as global mean, last observation, a decay mechanism based on these statistics, the superior imputation method SAITS (Du et al., 2023), or end-to-end method (Jing et al., 2025). The other methods include classic RNN-based methods like GRUD (Che et al., 2016), LSTM-based methods such as BRITS (Cao et al., 2018), the top-performing Transformer-based methods Transformer (Vaswani et al., 2017) and Autoformer (Wu et al., 2021), and the end-to-end method BiTGraph (Chen et al., 2023) and Timeseed (Wu et al., 2026), which is specifically designed for missing data prediction.

Table 1 includes direct comparison. A common-backbone panel compares VQ against raw exemplar retrieval and dense/sparse graph aggregation using identical information. Missing labels are excluded from baseline losses unless a declared weak-supervision adaptation is used. Key observations are as follows. First, our method can achieve state-of-the-art results on almost all datasets. Real-world time series data usually contain different missing types. This shows that our method is able to better tackle the inaccurate data problems with VQ and sharpness-aware learning.

6 Related work

Forecasting with inaccurate observations. GRU-D incorporates missingness and elapsed-time information (Che et al., 2016); GRIN, SAITS, and CSDI address multivariate imputation using graphs, attention, and diffusion (Cini et al., 2022; Du et al., 2023; Tashiro et al., 2021). CoIFNet integrates imputation and forecasting, whereas CRIB questions whether a reconstruction objective necessarily preserves forecasting information (Tang et al., 2025; Yang et al., 2025). S4 (Du et al., 2023) and S4M (Jing et al., 2025) impute missing data using strategies such as global mean, last observation, a decay mechanism based on these statistics, the superior imputation method SAITS or end-to-end. We therefore compare joint learning against both impute-then-forecast pipelines and a completion-free forecaster. Our distinction is the combination of past-only prototype completion for inaccurate *labels* and an explicit mechanism controlling their training influence.

Multiscale models and quantized representations. PatchTST, iTransformer, and TimeMixer provide complementary temporal and cross-variable forecasting architectures (Nie et al., 2023; Liu et al., 2024; Wang et al., 2024). MH-GIN uses multiscale heterogeneous graph propagation for AIS imputation (Liu et al., n.d.). VQ-VAE learns discrete representations through vector quantization (van den Oord et al., 2017), and Sparse-VQ already applies quantization to time-series forecasting (Zhao et al., 2024). We do not claim quantization or multiscale forecasting as new in isolation. Our codebook is a compressed, observation-supported memory of context–continuation pairs, with a past-context-only retrieval key and separate input-completion and weak-target roles.

Sharpness and unreliable supervision. SAM and ASAM optimize neighborhood-sensitive objectives (Foret et al., 2020; Kwon et al., 2021); SAMformer applies sharpness-aware optimization to forecasting (Ilbert et al., 2024). Co-teaching and learned example reweighting mitigate unreliable labels (Han et al., 2018; Ren et al., 2018), and SANER studies noisy-label learning through SAM-related gradient downweighting (Luong et al., 2024). SAWS instead uses a detached, periodically refreshed sharpness score to weight individual forecasting windows and their retrieved weak targets. SAM remains a separate optimizer baseline. Because sharpness depends on parameterization and can misidentify difficult clean examples (Dinh et al., 2017), we use scale-normalized probes, observed-label anchoring, and direct reliability diagnostics.

7 Conclusion and limitations

SAWS combines past-only VQ completion with sharpness-aware weak supervision for forecasting inaccurate time series. Its analysis isolates prototype bias, support variance, and the selection cost of confidence weighting. Prototype averaging can erase rare futures; head-local sensitivity can suppress difficult clean windows; and wholly unobserved channels remain unidentified without additional information. Prefix freezing also limits adaptation to new regimes. The specified experiments must establish whether these tradeoffs yield improved forecasting and lower total cost; the present draft reports no empirical outcomes.

AI USE STATEMENT

We only used generative AI tools to assist with the writing of this paper. All AI-assisted text were carefully reviewed and verified by the authors. The research ideas, methodology, experimental design, analysis, and scientific claims were conducted independently by the authors. We take full responsibility for the final content of this work.

REPRODUCIBILITY STATEMENT

In this work, we have made every effort to ensure reproducibility. In Sections 3.3–3.7, from input to output, we provide a detailed description of our method, including model structure, model training, and anomaly scores. In Appendix C, we describe all the datasets in detail. In Appendix A, we summarize all the implementation details and default hyper-parameters, including batch size, window size, learning rate, etc. Our codes will be available upon acceptance.

ETHICS STATEMENT

Our work use existing benchmark datasets and codes, without involving human subjects or the intentional collection of personally identifiable information. The proposed method is intended for general-purpose time series forecasting research.

REFERENCES

Wei Cao, Dong Wang, Jian Li, Hao Zhou, Lei Li, and Yitan Li. BRITS: Bidirectional recurrent imputation for time series. In *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://arxiv.org/abs/1805.10572.

Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural networks for multivariate time series with missing values. *arXiv preprint arXiv:1606.01865*, 2016. URL https://arxiv.org/abs/1606.01865.

Xiaodan Chen, Xiucheng Li, Bo Liu, and Zhijun Li. Biased temporal convolution graph network for time series forecasting with missing values. In *The Twelfth International Conference on Learning Representations*, 2023.

Andrea Cini, Ivan Marisca, and Cesare Alippi. Filling the G_ap_s: Multivariate time series imputation by graph neural networks. In *International Conference on Learning Representations*, 2022. URL https://arxiv.org/abs/2108.00298.

Laurent Dinh, Razvan Pascanu, Samy Bengio, and Yoshua Bengio. Sharp minima can generalize for deep nets. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1019–1028, 2017. URL https://proceedings.mlr.press/v70/dinh17b.html.

Wenjie Du, David Cote, and Yan Liu. SAITS: Self-attention-based imputation for time series. *Expert Systems with Applications*, 219:119619, 2023. doi: 10.1016/j.eswa.2023.119619. URL https://arxiv.org/abs/2202.08516.

Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *International Conference on Learning Representations*, 2020. URL https://arxiv.org/abs/2010.01412.

Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://arxiv.org/abs/1804.06872.

Romain Ilbert, Ambroise Odonnat, Vasilii Feofanov, Aladin Virmaux, Giuseppe Paolo, Themis Palpanas, and Ievgen Redko. SAMformer: Unlocking the potential of transformers in time series forecasting with sharpness-aware minimization and channel-wise attention. In *Proceedings of*

the 41st International Conference on Machine Learning, volume 235 of *Proceedings of Machine Learning Research*, pp. 20924–20954, 2024. URL https://proceedings.mlr.press/v235/ilbert24a.html.

Peng Jing, Meiqi Yang, Qiong Zhang, and Xiaoxiao Li. S4M: S4 for multivariate time series forecasting with missing values. In *The Thirteenth International Conference on Learning Representations*. OpenReview.net, 2025. URL https://openreview.net/forum?id=BkftcwIVmR.

Taesung Kim, Jinhee Kim, Yunwon Tae, Cheonbok Park, Jang-Ho Choi, and Jaegul Choo. Reversible instance normalization for accurate time-series forecasting against distribution shift. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=cGDAkQo1C0p.

Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. ASAM: Adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 5905–5914, 2021. URL https://proceedings.mlr.press/v139/kwon21b.html.

Hengyu Liu, Tianyi Li, Yuqiang He, Kristian Torp, Yushuai Li, and Christian S. Jensen. MH-GIN: Multi-scale heterogeneous graph-based imputation network for AIS data (extended version), n.d. Unpublished manuscript.

Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. itransformer: Inverted transformers are effective for time series forecasting. 2024. URL https://arxiv.org/abs/2310.06625.

Hoang-Chau Luong, Quang-Thuc Nguyen, Dat Ba Tran, and Minh-Triet Tran. Understanding SAM's robustness to noisy labels through gradient down-weighting. *arXiv preprint arXiv:2411.17132*, 2024. URL https://arxiv.org/abs/2411.17132. Revised March 2026.

Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. In *International Conference on Learning Representations*, 2023. URL https://arxiv.org/abs/2211.14730.

Xiangfei Qiu, Xvyuan Liu, Tianen Shen, Xingjian Wu, Hanyin Cheng, Bin Yang, and Jilin Hu. SEER: transformer-based robust time series forecasting via automated patch enhancement and replacement. *ICML 2026*, 2026a.

Xiangfei Qiu, Yuhan Zhu, Zhengyu Li, Hanyin Cheng, Xingjian Wu, Chenjuan Guo, Bin Yang, and Jilin Hu. DAG: A dual causal network for time series forecasting with exogenous variables. *ICML 2026*, 2026b.

Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4334–4343, 2018. URL https://proceedings.mlr.press/v80/ren18a.html.

Kai Tang, Ji Zhang, Hua Meng, Minbo Ma, Qi Xiong, Fengmao Lv, Jie Xu, and Tianrui Li. CoIFNet: A unified framework for multivariate time series forecasting with missing values. *arXiv preprint arXiv:2506.13064*, 2025. URL https://arxiv.org/abs/2506.13064.

Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. CSDI: Conditional score-based diffusion models for probabilistic time series imputation. In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://arxiv.org/abs/2107.03502.

Adedolapo Aishat Toye, Asuman Celik, and Samantha Kleinberg. Benchmarking missing data imputation methods for time series using real-world test cases. In *Proceedings of the Sixth Conference on Health, Inference, and Learning*, volume 287 of *Proceedings of Machine Learning Research*, pp. 480–501, 2025. URL https://proceedings.mlr.press/v287/toye25a.html.

Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://arxiv.org/abs/1711.00937.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.

Saverio Vito. Air quality. UCI Machine Learning Repository, 2008. URL https://archive.ics.uci.edu/dataset/360/air+quality.

Shiyu Wang, Haixu Wu, Xiaoming Shi, Tengge Hu, Huakun Luo, Lintao Ma, James Y. Zhang, and Jun Zhou. TimeMixer: Decomposable multiscale mixing for time series forecasting. In *International Conference on Learning Representations*, 2024. URL https://arxiv.org/abs/2405.14616.

Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. In *Advances in Neural Information Processing Systems*, 2021.

Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. TimesNet: Temporal 2D-variation modeling for general time series analysis. In *International Conference on Learning Representations*, 2022. URL https://arxiv.org/abs/2210.02186.

Zhaowang Wu, Kaixin Deng, and Hua Yan. Timeseed: Effective time series forecasting with sparse endogenous variables. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=p6VlvdOky9.

Jie Yang, Yifan Hu, Kexin Zhang, Luyang Niu, Philip S. Yu, and Kaize Ding. Revisiting multivariate time series forecasting with missing values. *arXiv preprint arXiv:2509.23494*, 2025. URL https://arxiv.org/abs/2509.23494. Revised January 2026.

Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. Are transformers effective for time series forecasting? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 11121–11128, 2023. doi: 10.1609/aaai.v37i9.26317. URL https://ojs.aaai.org/index.php/AAAI/article/view/26317.

Kai Zhao, Chenjuan Guo, Yunyao Cheng, Peng Han, Miao Zhang, and Bin Yang. Multiple time series forecasting with dynamic graph modeling. *Proc. VLDB Endow.*, 17(4):753–765, 2023. doi: 10.14778/3636218.3636230. URL https://www.vldb.org/pvldb/vol17/p753-zhao.pdf.

Yanjun Zhao, Tian Zhou, Chao Chen, Liang Sun, Yi Qian, and Rong Jin. Sparse-VQ Transformer: An FFN-free framework with vector quantization for enhanced time series forecasting. *arXiv preprint arXiv:2402.05830*, 2024. URL https://arxiv.org/abs/2402.05830.

Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 11106–11115, 2021. URL https://arxiv.org/abs/2012.07436.

A Implementation details

A.1 Observed teacher, sharpness refreshes, and weight mass

Warm up the forecasting student using observed-target loss only, and initialize the EMA teacher from the student. After each student update,

$$\bar{\theta} \leftarrow \kappa\bar{\theta} + (1-\kappa)\theta. \tag{17}$$

The weak-loss coefficient is increased from zero over a declared ramp period. Every R training epochs, freeze one teacher checkpoint, its block scale matrix $\mathbf{A}$, and the shared scalar c_{loss} for a complete score-refresh pass. Use evaluation mode, identical detached features for the base and perturbed heads, and no batch-normalization-state updates. EMA loss-scale updates accumulate for the *next* refresh, so cached scores are comparable within the current pass.

For a batch with no observed target coordinates, omit the observed loss. For a batch with no supported missing coordinates, omit the weak loss. If both are empty, skip the optimization step. A zero head gradient gives $\boldsymbol{\delta}_i = 0$. Clip the resulting sharpness score to $[0, s_{\max}]$, cache w_i, and apply stop-gradient to weights and weak targets. Fully unlabeled windows receive the neutral sharpness multiplier and are tracked as a separate stratum.

The distinction between relative weighting and total confidence is important. Let $\mathcal{U} = \{(i,q) : U_{iq} = 1\}$, $v_{iq} = w_i r_{iq}$, and

$$p_{\mathcal{B}} = \frac{1}{|\mathcal{U}|}\sum_{(i,q)\in\mathcal{U}} v_{iq}. \qquad \mathcal{L}_u = p_{\mathcal{B}}\frac{\sum_{(i,q)\in\mathcal{U}} v_{iq}\frac{1}{2}(\widehat{Y}_{iq} - \widetilde{Y}_{iq})^2}{\sum_{(i,q)\in\mathcal{U}} v_{iq}}. \tag{18}$$

This is Equation 13 written as a retained-mass gate times a normalized tilted loss. It reduces the total weak contribution when all eligible weak targets have low confidence; normalization by total confidence alone would cancel that effect. For stability, compute the equivalent unnormalized numerator divided by $|\mathcal{U}|$ directly. The observed anchor remains present regardless of weak-label confidence.

A.2 Training-cost accounting

Let N_r be the number of reference windows and N_t the number of scored training windows. After dictionary fitting, assigning references costs $O(N_r K d)$ and accumulating payloads costs $O(N_r Q)$. For a batch of B_b queries, a materialized exact-search distance matrix requires $O(B_b K)$ temporary storage; chunked search trades additional launch overhead for lower peak memory. An exact top-m selection adds its selection cost to $O(Kd)$ distance computation. Report the actual implementation, including whether selection uses a heap, partial sort, or fused kernel.

For the default linear output head, per-window gradients and perturbed predictions can be computed in batches from detached features without a backbone backward pass. A refresh still requires teacher feature computation and additional head evaluation for the scored windows; its cost is not zero. The amortized total includes all refreshes, all dictionary fitting and prefix-cache work, and any hyperparameter-search budget. Static codeword storage alone is not an adequate memory metric. Increasing K or m can improve approximation while increasing search, payload access, and cache-construction costs.

B Theoretical Analysis and Proofs

We analyze a fixed codebook, a fixed reference-memory construction, and a frozen teacher used to assign supervision weights. The results separate prototype estimation error, reliability-based selection, and the distortion caused by replacing unavailable targets with weak labels. They do not assert convergence of the joint adaptive training procedure. For clarity, we consider one flattened future coordinate q; the same arguments apply coordinatewise and can be averaged over coordinates.

B.1 Prototype completion with partially observed references

Let x denote the information available to a query's past-only retrieval key. Let Y_q be its desired future target, $\mu_q(x) = \mathbb{E}[Y_q \mid x]$, and $v_q(x) = \mathrm{Var}(Y_q \mid x)$. For a fixed VQ cell k, let $\mu_{kq} =$

$\mathbb{E}[Y_q \mid K = k]$. Among the reference windows assigned to this cell, let $n_{kq} > 0$ be the number with coordinate q observed, and let $\widehat{\mu}_{kq}$ be their observed-coordinate sample mean. The completion is

$$g_q(x) = \sum_{k \in \mathcal{K}_q(x)} a_k(x)\widehat{\mu}_{kq}, \qquad a_k(x) \geq 0, \qquad \sum_{k \in \mathcal{K}_q(x)} a_k(x) = 1. \tag{19}$$

Unsupported cells are excluded before normalizing the weights. If no retrieved cell has observed support, the method abstains from using that coordinate as a pseudo-target.

For the following result, the codebook is learned on a separate split. Conditional on the codebook, reference assignments, and observation masks, the observed values in cell k have mean μ_{kq} and variance σ^2_{kq}, and the reference values are independent. The reference sample is also conditionally independent of the query's future target. The aggregation weights may depend on x, the fixed codebook, and support counts, but not on the reference response values. These conditions require ignorable observation within cells. If observed values are themselves corrupted, they additionally require conditionally mean-zero corruption. MAR conditional on a finer context does not automatically imply ignorability after that context is coarsened into VQ cells.

Proposition B.1 (Prototype bias–variance decomposition)**.** *Under the preceding conditions,*

$$\begin{aligned}\mathbb{E}[(g_q(x) - Y_q)^2 \mid x, \{n_{kq}\}] = v_q(x) + \left(\sum_k a_k(x)\mu_{kq} - \mu_q(x)\right)^2 \\ + \sum_k \frac{a_k(x)^2 \sigma^2_{kq}}{n_{kq}}.\end{aligned} \tag{20}$$

The conditioning on the fixed codebook, reference assignments, and masks is suppressed in the notation. For MSE against $\mu_q(x)$, the first term is absent.

Proof. Write $\widehat{\mu}_{kq} = \mu_{kq} + \varepsilon_{kq}$ and $Y_q = \mu_q(x) + \xi_q$. Conditional on the stated quantities, $\mathbb{E}[\varepsilon_{kq}] = 0$, $\mathbb{E}[\varepsilon^2_{kq}] = \sigma^2_{kq}/n_{kq}$, $\mathbb{E}[\xi_q] = 0$, and $\mathbb{E}[\xi_q^2] = v_q(x)$. Therefore,

$$g_q(x) - Y_q = \left(\sum_k a_k(x)\mu_{kq} - \mu_q(x)\right) + \sum_k a_k(x)\varepsilon_{kq} - \xi_q. \tag{21}$$

After squaring and taking expectations, all cross terms vanish by centering and conditional independence. The remaining three terms are exactly those in Equation (20). □

The decomposition identifies irreducible future variation, approximation bias, and uncertainty from finite observed support. If $\sigma^2_{kq} \leq \sigma^2_{\max,q}$, its variance term is at most $\sigma^2_{\max,q}/n_{\mathrm{eff},q}(x)$, where

$$n_{\mathrm{eff},q}(x) = \left(\sum_k \frac{a_k(x)^2}{n_{kq}}\right)^{-1}. \tag{22}$$

Retrieving more prototypes can reduce this variance term while increasing approximation bias. The proposition therefore does not claim monotonic improvement with codebook size or retrieval count.

Dependence and observation bias. Independent reference samples, or independent block-level units counted as such, are a sufficient condition for the variance formula; nonoverlapping windows alone do not establish independence of a time series. With correlated reference means, replace the final term of Equation (20) by $a(x)^\top \Sigma_q a(x)$, where $[\Sigma_q]_{kl} = \mathrm{Cov}(\widehat{\mu}_{kq}, \widehat{\mu}_{lq})$. If reference and query targets are dependent, their covariance must also be retained. Under nonignorable observation, the observed-cell mean $\mu^{\mathrm{obs}}_{kq} = \mathbb{E}[Y^{\mathrm{obs}}_q \mid K = k, M_q = 1]$ replaces μ_{kq}, and the approximation term also contains missingness-selection or corruption bias. Increased support and cross-fitting do not by themselves eliminate that bias.

B.2 A SANITY CALCULATION FOR THE SHARPNESS PROBE

Let ℓ_i be the frozen teacher's half squared forecasting loss, averaged only over genuinely observed future coordinates in window i. With head parameters $\bar{\theta}_h$, fixed symmetric parameter scales A, and $g_i = \nabla_{\bar{\theta}_h} \ell_i$, the one-step probe is

$$\delta_i = \rho \frac{A^2 g_i}{\|A g_i\|_2 + \epsilon}, \qquad s_i = \frac{[\ell_i(\bar{\theta}_h + \delta_i) - \ell_i(\bar{\theta}_h)]_+}{c_{\text{loss}} + \epsilon}. \tag{23}$$

The positive scale c_{loss} is shared across the windows being scored and detached from gradient computation. It can be estimated from a running training-loss statistic. It must not be the same window's own residual-dependent loss. For invertible A, the perturbation is feasible for $\|A^{-1}\delta_i\|_2 \le \rho$. This is a one-step head-local sensitivity measure, rather than an exact maximization over a parameter ball or a full-network certificate. When no future coordinate is observed, the sharpness modulation is disabled; prototype reliability still determines whether to use the weak target.

Proposition B.2 (Linear-head residual sensitivity)**.** *Consider $\ell_y(\theta) = \frac{1}{2}(\theta^\top h - y)^2$ for fixed features h and a fixed symmetric scale matrix A. Let $e = \theta^\top h - y$ and $a = \|Ah\|_2$. For $e \neq 0$, $a > 0$, and the unregularized direction $\delta = \rho A^2 \nabla \ell_y / \|A \nabla \ell_y\|_2$,*

$$\ell_y(\theta + \delta) - \ell_y(\theta) = \rho |e| a + \frac{\rho^2 a^2}{2}. \tag{24}$$

For fixed h, A, ρ, this increase is strictly increasing in $|e|$. With the regularized denominator in Equation (23)*, the increase is nondecreasing in $|e|$.*

Proof. Since $\nabla \ell_y = eh$, the unregularized direction satisfies $\delta = \rho \operatorname{sign}(e) A^2 h / a$ and $h^\top \delta = \rho \operatorname{sign}(e) a$. Expanding $\frac{1}{2}(e + h^\top \delta)^2 - \frac{1}{2}e^2$ gives Equation (24). For $\epsilon > 0$, write $c = \rho |e| a^2 / (|e| a + \epsilon)$. Then $h^\top \delta = \operatorname{sign}(e) c$, so the loss increase is $|e| c + c^2/2$. Both c and $|e|$ are nonnegative and nondecreasing functions of $|e|$, which proves the second assertion. □

At zero gradient, the implemented first-order probe uses $\delta = 0$ and does not explore second-order curvature. The proposition motivates the one-sided probe but does not equate a large residual with corruption: difficult valid windows and distribution shifts can also have large residuals. In nonlinear forecasting heads, feature sensitivity and curvature introduce further dependencies.

Why a symmetric probe is insufficient. For the same linear model and any fixed direction u,

$$\frac{\ell_y(\theta + \rho u) + \ell_y(\theta - \rho u)}{2} - \ell_y(\theta) = \frac{\rho^2 (u^\top h)^2}{2}, \tag{25}$$

which is independent of y. Symmetric curvature alone therefore cannot universally identify target corruption. Furthermore, dividing the one-sided increase in Equation (24) by $\frac{1}{2}e^2 + \tau$ makes the score tend to zero as $|e| \to \infty$. A shared scale avoids this particular reversal.

B.3 EXPONENTIAL WEIGHTING UNDER AN EXPLICIT CALIBRATION CONDITION

Let P be a specified target-coordinate population within a score-available stratum, let $E = (g(X) - Y)^2$ denote its weak-label error, and let $r \in [0, 1]$ be prototype support/dispersion reliability. For $s \in [0, s_{\max}]$, inverse temperature $\beta = 1/T \ge 0$, and $\mathbb{E}_P[r] > 0$, define

$$w_\beta = r e^{-\beta s}, \qquad p_\beta = \mathbb{E}_P[w_\beta], \qquad \frac{dQ_\beta}{dP} = \frac{w_\beta}{p_\beta}. \tag{26}$$

The baseline measure Q_0 already incorporates prototype reliability. No independence between r, s, and E is assumed.

Proposition B.3 (Calibrated sharpness weighting)**.** *Assume E is integrable under P and that $m(t) = \mathbb{E}_{Q_0}[E \mid s = t]$ has a nondecreasing version. Then $\eta(\beta) = \mathbb{E}_{Q_\beta}[E]$ is nonincreasing, with*

$$\eta'(\beta) = -\operatorname{Cov}_{Q_\beta}(s, E) \le 0. \tag{27}$$

Proof. Because s is bounded, differentiation under the expectation is valid. The quotient rule applied to $\eta(\beta) = \mathbb{E}_P[re^{-\beta s}E]/\mathbb{E}_P[re^{-\beta s}]$ gives the equality in Equation (27). The density change from Q_0 to Q_β depends only on s, so the conditional distribution of E given s is unchanged. It follows that $\mathrm{Cov}_{Q_\beta}(s, E) = \mathrm{Cov}_{Q_\beta}(s, m(s))$. For independent S, S' with the Q_β-distribution of s,

$$2\,\mathrm{Cov}_{Q_\beta}(s, m(s)) = \mathbb{E}[(S - S')(m(S) - m(S'))] \geq 0, \tag{28}$$

where monotonicity gives the final inequality. □

The calibration condition is an assumption to be assessed on held-out observed coordinates, not a consequence of the sharpness definition. Transporting calibration from observed to truly missing coordinates requires a corresponding observation/transport assumption. The result concerns normalized weak-label error, not final forecasting accuracy. A positive additive floor on the pseudo-target weights would change the family in Equation (26) and need not preserve its monotonicity. An independently retained observed-target anchor is compatible with this statement.

B.4 Weak-label risk distortion and selection shift

Fix a measurable weight $w \in [0, 1]$ with $p = \mathbb{E}_P[w] > 0$ and define Q by $dQ/dP = w/p$. Assume $|Y|, |g(X)|, |f(X)| \leq B$ for all $f \in \mathcal{F}$. This boundedness is an explicit theoretical condition, for example when standardized target, weak-target, and prediction values are all bounded or explicitly clipped. Clipping only the loss would not establish these value bounds. Use half squared loss and define

$$R_Q(f) = \frac{1}{2}\mathbb{E}_Q[(f(X)-Y)^2], \qquad \widetilde{R}_Q(f) = \frac{1}{2}\mathbb{E}_Q[(f(X)-g(X))^2], \qquad \eta = \mathbb{E}_Q[(g(X)-Y)^2]. \tag{29}$$

Theorem B.4 (Weighted weak-supervision risk distortion)**.** *Under the preceding conditions,*

$$\sup_{f\in\mathcal{F}} |\widetilde{R}_Q(f) - R_Q(f)| \leq 2B\sqrt{\eta}. \tag{30}$$

Suppose $\widehat{f}$ satisfies $\widetilde{R}_Q(\widehat{f}) \leq \inf_{f\in\mathcal{F}} \widetilde{R}_Q(f) + \varepsilon_{\mathrm{opt}}$. Then

$$R_Q(\widehat{f}) - \inf_{f\in\mathcal{F}} R_Q(f) \leq 4B\sqrt{\eta} + \varepsilon_{\mathrm{opt}}, \tag{31}$$

and the original-population excess risk satisfies

$$R_P(\widehat{f}) - \inf_{f\in\mathcal{F}} R_P(f) \leq 4B\sqrt{\eta} + 4B^2\,\mathrm{TV}(P, Q) + \varepsilon_{\mathrm{opt}}, \tag{32}$$

where

$$\mathrm{TV}(P, Q) = \frac{1}{2}\mathbb{E}_P\left[\left|\frac{w}{\mathbb{E}_P[w]} - 1\right|\right]. \tag{33}$$

Proof. The pointwise half-loss difference satisfies

$$\frac{1}{2}|(f-g)^2 - (f-Y)^2| = \frac{1}{2}|(Y-g)(2f-g-Y)| \leq 2B|g-Y|. \tag{34}$$

Taking Q-expectations and using Cauchy–Schwarz yields Equation (30). For any $f \in \mathcal{F}$,

$$\begin{aligned} R_Q(\widehat{f}) &\leq \widetilde{R}_Q(\widehat{f}) + 2B\sqrt{\eta} \\ &\leq \widetilde{R}_Q(f) + \varepsilon_{\mathrm{opt}} + 2B\sqrt{\eta} \\ &\leq R_Q(f) + \varepsilon_{\mathrm{opt}} + 4B\sqrt{\eta}. \end{aligned} \tag{35}$$

Taking the infimum over f proves Equation (31). The half squared loss takes values in $[0, 2B^2]$, hence $|R_P(f) - R_Q(f)| \leq 2B^2\,\mathrm{TV}(P, Q)$ uniformly in f. Applying this inequality once to $\widehat{f}$ and once to the comparison infimum gives Equation (32). □

The theorem applies to any fixed measurable weighting rule, including a rule depending on observed training targets. It requires no claim that sharpness is itself calibrated. Combining it with Proposition B.3 shows a tradeoff: calibrated weighting can decrease η, while concentrating weights can increase $\mathrm{TV}(P, Q)$. The theorem consequently does not justify an arbitrarily small temperature. It is a population-objective result. A finite-sample generalization claim would additionally require control of the forecast class, the learned weighting rule, and temporal dependence.

Observed-target anchor. For the actual separated objective, let P be the population of eligible missing coordinates ($U = 1$), and let R_{obs} be the observed-target term and define

$$J_{\mathrm{weak}}(f) = R_{\mathrm{obs}}(f) + \frac{\lambda}{2}\mathbb{E}_P[w(f-g)^2],$$
$$J_{\mathrm{oracle}}(f) = R_{\mathrm{obs}}(f) + \frac{\lambda}{2}\mathbb{E}_P[w(f-Y)^2]. \tag{36}$$

Here the generic w is the combined method weight $w_i r_{iq}$, conditional on $U_{iq} = 1$. Set $e_w = \mathbb{E}_P[w(g-Y)^2]$. Because the observed-target term cancels,

$$\sup_{f\in\mathcal{F}} |J_{\mathrm{weak}}(f) - J_{\mathrm{oracle}}(f)| \le 2\lambda B\sqrt{pe_w}. \tag{37}$$

An $\varepsilon_{\mathrm{opt}}$-optimal population minimizer of J_{weak} therefore has J_{oracle}-excess at most $4\lambda B\sqrt{pe_w} + \varepsilon_{\mathrm{opt}}$. If the pseudo-target term is normalized by p, replace $\sqrt{pe_w}$ by $\sqrt{\eta}$. This comparison addresses weak-target substitution and makes no additional assumption about the sampling distribution of R_{obs}; it does not establish that noisy observed targets have been corrected or that their distribution equals P.

Weight concentration. We recommend reporting retained weight mass and the diagnostic $n_{\mathrm{eff}} = (\sum_i w_i)^2 / \sum_i w_i^2$ together with forecast accuracy. This quantity measures weight concentration; by itself it does not establish independence or a generalization bound for temporal data and learned weights.

C Detailed evaluation and reporting specification

C.1 Datasets, preprocessing, and measurement targets

The controlled benchmark panel consists of ETTm1, ETTm2, Electricity, and Weather, using the data distributions adopted by the cited forecasting benchmarks (Zhou et al., 2021; Nie et al., 2023). The ETT source is the authors' official ETDataset repository.[1] Each released archive, timestamp parser, channel list, and preprocessing transformation must be identified before running the comparison. Do not silently merge results from different processed versions.

The native-missingness panel uses UCI Air Quality (Vito, 2008), whose source documents missing measurements with a sentinel value. Convert documented missing sentinels into explicit masks before scaling; do not interpret zero concentrations as missing without a documented rule. Date and time fields become time indices or known covariates, not numerical forecast channels. The panel is used for forecasting available future sensor/reference measurements, with channels and observation counts disclosed. No claim is made about errors at naturally unobserved targets.

Controlled corruptions are applied to an evaluator-retained version of originally available values. A clean evaluation target means the recorded value before the *additional* synthetic corruption; it is not a claim that the original sensor measurement is physically noise-free. Native noisy data can support only risk against the recorded outcome unless an independent reference measurement is available.

Table 2: Dataset statistics.

Data	Variables	Time steps	Granularity
Electricity	321	26,304	1 hour
ETTh1	7	17,420	15 min
Weather	21	52,696	10 min
Traffic	862	17,544	1 hour

C.2 Missingness and corruption mechanisms

Let p_x denote additional input missingness and p_y training-label observation probability. These quantities are varied independently in the label-release protocol; the fully-concealed protocol re-

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

ports their induced dependence. All raw masks, seeds, and timestamps are shared across compared methods.

- **Independent entries:** sample a Bernoulli observation mask for each timestamp–channel measurement and observation role.
- **Contiguous blocks:** sample start positions and block lengths per channel until the desired additional missing fraction is reached; report realized fractions and gap-length quantiles.
- **Channel outages:** remove a subset of channels for sustained intervals. Distinguish temporary outages from channels absent from the entire reference archive.
- **Covariate-dependent observation:** choose observation probability from a fixed logistic function of known calendar or previously observed context features. Parameters are selected without test outcomes.
- **Value-dependent observation:** the evaluator masks values using a prespecified function of their hidden magnitude or tail membership. The learner receives only the resulting mask. Such data are not claimed to satisfy the theoretical ignorability condition.
- **Observed-value noise:** add zero-mean Gaussian perturbations scaled by frozen channel statistics, and separately test sparse impulses or biased sensor drift. Input corruption and label corruption are separate axes.

Missingness-mechanism sensitivity is necessary because rankings can change across observation patterns (Toye et al., 2025). The suite includes matched train/test corruption, missingness-shift, and clean-test-after-corrupted-training conditions. A clean complete-data reference is an additional-information comparison and is labeled separately from methods constrained to the corrupted observations.