

Modeling Information Blackouts in Missing Not-At-Random Time Series Data

Aman Sunesh
New York University
as18181@nyu.edu

Allan Ma
New York University
aym9849@nyu.edu

Siddarth Nilol
New York University
sk12590@nyu.edu

Abstract—Large-scale traffic forecasting relies on fixed sensor networks that often exhibit *blackouts*: contiguous intervals of missing measurements caused by detector or communication failures. These outages are typically handled under a Missing At Random (MAR) assumption, even though blackout events may correlate with unobserved traffic conditions (e.g., congestion or anomalous flow), motivating a Missing Not At Random (MNAR) treatment. We propose a latent state-space framework that jointly models (i) traffic dynamics via a linear dynamical system and (ii) sensor dropout via a Bernoulli observation channel whose probability depends on the latent traffic state. Inference uses an Extended Kalman Filter with Rauch–Tung–Striebel smoothing, and parameters are learned via an approximate EM procedure with a dedicated update for detector-specific missingness parameters. On the Seattle inductive loop detector data, introducing latent dynamics yields large gains over naive baselines, reducing blackout imputation RMSE from 7.02 (LOCF) and 5.02 (linear interpolation + seasonal naive) to 4.23 (MAR LDS), corresponding to a $\sim 64\%$ reduction in MSE relative to LOCF. Explicit MNAR modeling provides a consistent but smaller additional improvement on real data (imputation RMSE 4.20; 0.8% RMSE reduction relative to MAR), with similar modest gains for short-horizon post-blackout forecasts (evaluated at 1, 3, and 6 steps). In controlled synthetic experiments, the MNAR advantage increases as the true missingness dependence on latent state strengthens. Overall, temporal dynamics dominate performance, while MNAR modeling offers a principled refinement that becomes most valuable when missingness is genuinely informative.

Index Terms—missing not at random, traffic sensors, imputation, state-space models, Kalman filtering, expectation-maximization

I. INTRODUCTION

Traffic forecasting systems rely on dense networks of fixed sensors (e.g., inductive loop detectors) to provide continuous measurements of traffic speed and flow. In deployed systems, however, sensors frequently experience *blackouts*: contiguous intervals of missing readings lasting from minutes to hours. These outages arise from hardware faults, maintenance, and communication drops, and they create systematic gaps that degrade both reconstruction during the blackout and forecasting immediately after recovery.

A. Blackouts and the Missingness Mechanism

Most pipelines handle missing values implicitly—e.g., by masking missing measurements, skipping updates when readings are absent, or applying simple interpolation. These practices effectively treat missingness as *ignorable* (a MAR-style

assumption), where the probability of missingness depends only on observed information:

$$P(m_t | x_t^{(\text{obs})}, x_t^{(\text{mis})}) = P(m_t | x_t^{(\text{obs})}), \quad (1)$$

where m_t is a missingness indicator and $x_t = (x_t^{(\text{obs})}, x_t^{(\text{mis})})$ denotes measurements partitioned into observed and missing components. Under MAR, the missingness process can be ignored without biasing inference [6], [7].

Blackouts in traffic systems can violate MAR. For example, high congestion, abnormal flow regimes, or network stress may increase the likelihood of communication loss or sensor malfunction. This corresponds to Missing Not At Random (MNAR), where missingness depends on latent traffic conditions:

$$P(m_t | z_t, x_t^{(\text{obs})}) \neq P(m_t | x_t^{(\text{obs})}). \quad (2)$$

If blackouts are MNAR, then the *pattern* of missingness may contain information about the latent traffic state and should be modeled rather than treated as a preprocessing artifact [6]–[8].

B. Research Question

This motivates the central question of our work:

When detectors exhibit contiguous blackouts, does explicitly modeling blackout events as state-dependent (MNAR) improve reconstruction during the outage and short-horizon forecasting after the outage, beyond what is achieved by modeling temporal dynamics alone?

C. Proposed Model and Inference

We model traffic with a low-dimensional latent state $z_t \in \mathbb{R}^K$ following linear dynamics,

$$z_t | z_{t-1} \sim \mathcal{N}(Az_{t-1}, Q), \quad (3)$$

and (when observed) measurements $x_t \in \mathbb{R}^D$,

$$x_t | z_t \sim \mathcal{N}(Cz_t, R). \quad (4)$$

To capture MNAR blackouts, each detector d has a missingness indicator $m_{t,d} \in \{0, 1\}$ with a latent-state-dependent probability,

$$m_{t,d} | z_t \sim \text{Bernoulli}(\pi_{t,d}), \quad (5)$$

$$\text{logit}(\pi_{t,d}) = b_d + \phi_d^\top z_t + \psi_d^\top f_t + \eta_d^\top g_d, \quad (6)$$

$$\pi_{t,d} = \sigma(b_d + \phi_d^\top z_t + \psi_d^\top f_t + \eta_d^\top g_d), \quad (7)$$

where $\sigma(\cdot)$ is the logistic sigmoid, ϕ_d weights the latent state, f_t are time features with weights ψ_d , and g_d are static detector features with weights η_d .

Because the missingness channel is nonlinear in z_t , exact filtering is intractable. We therefore use approximate inference: an Extended Kalman Filter (EKF) that incorporates both measurement likelihoods and missingness likelihoods, followed by Rauch–Tung–Striebel (RTS) smoothing. Parameters are learned via an approximate EM algorithm, with a dedicated optimization step for $\{\phi_d\}$.

D. What We Find

Empirically, latent dynamics are the dominant source of improvement: on the Seattle loop detector data, a MAR latent dynamical system substantially outperforms naive baselines for blackout imputation. Explicit MNAR modeling yields smaller but consistent gains on real data, while controlled synthetic experiments show that MNAR benefits grow as the true missingness dependence on latent state increases. This supports a nuanced conclusion: *modeling dynamics is essential, and modeling MNAR missingness is most valuable when dropout is genuinely informative.*

E. Contributions

This paper makes four contributions:

- We formulate traffic sensor blackouts as a potentially MNAR process by modeling dropout probabilities as a latent-state-dependent observation channel.
- We develop an approximate inference pipeline combining EKF filtering and RTS smoothing for joint latent-state estimation with state-dependent missingness.
- We train the model using an approximate EM procedure, including detector-specific missingness parameter updates.
- We evaluate blackout imputation and short-horizon post-blackout forecasting on real loop-detector data and controlled synthetic MNAR settings, quantifying when MNAR provides gains beyond MAR dynamics.

II. RELATED WORK

A. Traffic Forecasting and State-Space Inference

Traffic state estimation has long been studied through latent dynamical systems and Kalman-style inference, where missing measurements are typically handled by skipping the measurement update when observations are absent [1]. This implicitly treats missingness as uninformative, aligning with MAR-style assumptions. Modern spatiotemporal deep models (e.g., graph convolutional and recurrent architectures) achieve strong forecasting performance but commonly rely on preprocessing (masking or interpolation) to address missing data [2], [3], again usually without an explicit generative model for the missingness mechanism.

B. Learning with Missingness Under MAR

A large body of work improves robustness to missing observations by incorporating masks and time-since-observation features. GRU-D [4] augments recurrent models with decay dynamics conditioned on missingness patterns, and BRITS [5] performs bidirectional recurrent imputation. While effective in practice, these methods typically treat missingness as auxiliary input rather than explicitly modeling how missingness is generated; this corresponds more closely to MAR-compatible handling than MNAR identification.

C. MNAR Modeling and Informative Missingness

In classical missing-data theory, MNAR mechanisms arise when missingness depends on unobserved variables, requiring explicit modeling to avoid biased inference (e.g., selection models and shared-parameter models) [6]–[8]. Despite their relevance, MNAR ideas are rarely integrated into traffic forecasting and traffic state-space models, even though structured blackouts are ubiquitous in sensor networks.

D. Our Position

We bridge this gap by augmenting a latent linear dynamical system with an explicit Bernoulli missingness channel whose probability depends on the latent traffic state. This yields a principled MNAR formulation while retaining interpretable state-space structure and efficient approximate inference. Our experiments quantify when this additional modeling complexity pays off: gains are modest on real-world loop data but increase in controlled settings where missingness is truly state-dependent.

III. METHODS

We propose an MNAR-aware latent state-space model for traffic sensor blackouts that jointly captures (i) traffic dynamics and (ii) state-dependent sensor dropout. We perform approximate inference with an Extended Kalman Filter (EKF) augmented by a missingness likelihood, followed by Rauch–Tung–Striebel (RTS) smoothing. Parameters are learned via an approximate EM procedure with detector-wise updates for the missingness model.

A. Latent State-Space Model for Traffic Speeds

Let $z_t \in \mathbb{R}^K$ denote a low-dimensional latent traffic state at time t (5-minute intervals), and let $x_t \in \mathbb{R}^D$ denote detector speed measurements. We assume linear-Gaussian dynamics:

$$z_t \mid z_{t-1} \sim \mathcal{N}(Az_{t-1}, Q), \quad (8)$$

$$x_t \mid z_t \sim \mathcal{N}(Cz_t, R), \quad (9)$$

where $A \in \mathbb{R}^{K \times K}$, $Q \in \mathbb{R}^{K \times K}$, $C \in \mathbb{R}^{D \times K}$, and $R \in \mathbb{R}^{D \times D}$. In our implementation, R is constrained to be diagonal to enable efficient updates at large D .

B. MNAR Missingness Model with Intercepts and Covariates

Let $m_{t,d} \in \{0,1\}$ indicate whether detector d is missing at time t ($m_{t,d} = 1$ means missing). To capture informative (MNAR) blackouts, we model dropout probabilities as a logistic function of the latent state, with optional time-varying and detector-specific covariates:

$$\Pr(m_{t,d} = 1 \mid z_t, f_t, g_d) = \sigma(b_d + \phi_d^\top z_t + \psi_d^\top f_t + \eta_d^\top g_d), \quad (10)$$

where $\sigma(\cdot)$ is the sigmoid, b_d is a detector-specific intercept, $\phi_d \in \mathbb{R}^K$ weights the latent state, f_t are time features (e.g., time-of-day/day-of-week encodings) with weights ψ_d , and g_d are static detector features with weights η_d . This formulation reduces to an LDS baseline that treats missingness as ignorable during inference (i.e., does not include the mask likelihood as an observation channel).

C. MNAR-Aware EKF Update via Local Gaussianization

Exact filtering is intractable because the missingness likelihood in (10) is Bernoulli-logistic. We therefore incorporate the missingness channel as a *pseudo-observation* by locally Gaussianizing the Bernoulli likelihood around the predicted mean $\mu_{t|t-1}$.

Define logits $\ell_{t,d} = b_d + \phi_d^\top \mu_{t|t-1} + \psi_d^\top f_t + \eta_d^\top g_d$ and probabilities $\pi_{t,d} = \sigma(\ell_{t,d})$. We approximate

$$m_{t,d} \approx \pi_{t,d} + \epsilon_{t,d}, \quad \epsilon_{t,d} \sim \mathcal{N}(0, s_{t,d}), \quad (11)$$

with diagonal variance $s_{t,d}$ chosen either by a moment match $s_{t,d} \approx \pi_{t,d}(1 - \pi_{t,d})$ or as a tuned constant. The Jacobian with respect to z_t is

$$J_{t,d} = \frac{\partial \pi_{t,d}}{\partial z_t} = \pi_{t,d}(1 - \pi_{t,d}) \phi_d. \quad (12)$$

We then perform an EKF-style update that combines: (i) the linear-Gaussian speed likelihood on observed detectors, and (ii) the Gaussianized missingness likelihood on m_t . We additionally include a scalar weight w_{miss} to control the influence of the missingness block.

a) Efficient K -dimensional update: With diagonal R , the speed contribution reduces to sums of $K \times K$ outer products for observed detectors, enabling updates that only solve $K \times K$ linear systems. This makes inference practical even when D is large.

D. RTS Smoothing

After filtering, we apply RTS smoothing using standard LDS recursions:

$$G_t = \Sigma_{t|t} A^\top \Sigma_{t+1|t}^{-1}, \quad (13)$$

$$\mu_{t|T} = \mu_{t|t} + G_t(\mu_{t+1|T} - \mu_{t+1|t}), \quad (14)$$

$$\Sigma_{t|T} = \Sigma_{t|t} + G_t(\Sigma_{t+1|T} - \Sigma_{t+1|t})G_t^\top. \quad (15)$$

We use $(\mu_{t|T}, \Sigma_{t|T})$ for imputation and for the E-step statistics in EM.

E. Approximate EM Learning with Stabilization

We learn parameters via an approximate EM procedure:

a) E-step: Run MNAR-aware EKF + RTS to obtain posterior moments

$$\mathbb{E}[z_t] = \mu_{t|T}, \quad \mathbb{E}[z_t z_t^\top] = \Sigma_{t|T} + \mu_{t|T} \mu_{t|T}^\top. \quad (16)$$

b) M-step (LDS parameters): We update $(\mu_0, \Sigma_0, A, Q, C, R)$ using moment-based updates while respecting missingness: for each detector d , C_d and R_{dd} are computed using only timesteps where $x_{t,d}$ is observed. Cross-covariances $\mathbb{E}[z_t z_{t-1}^\top]$ are approximated using outer products of smoothed means, yielding a practical (though not exact) RTS-EM update. To prevent unstable dynamics during EM, we regularize by shrinking A toward identity and shrinking Q toward an isotropic prior, with an additional cap on $\text{tr}(Q)$.

c) M-step (missingness parameters): We update $\{b_d, \phi_d, \psi_d, \eta_d\}$ by detector-wise gradient ascent on the Bernoulli log-likelihood:

$$\mathcal{L}_d = \sum_{t \in \mathcal{T}_d} [m_{t,d} \log \pi_{t,d} + (1 - m_{t,d}) \log(1 - \pi_{t,d})], \quad (17)$$

where $\pi_{t,d}$ is given by (10). Importantly, artificially masked evaluation-window entries are excluded from \mathcal{T}_d so the missingness model does not learn from our injected masks.

F. Reconstruction and Post-Blackout Forecasting

a) Imputation: We reconstruct speeds using the smoothed latent mean:

$$\hat{x}_t = C \mu_{t|T}. \quad (18)$$

b) k -step forecasting: To forecast after a blackout ends at time index b , we propagate the filtered posterior forward k steps:

$$\mu_{b+k|b} = A^k \mu_{b|b}, \quad (19)$$

$$\Sigma_{b+k|b} = A^k \Sigma_{b|b} (A^\top)^k + \sum_{i=0}^{k-1} A^i Q (A^\top)^i, \quad (20)$$

and map to observation space:

$$x_{b+k} \mid x_{1:b} \approx \mathcal{N}(C \mu_{b+k|b}, C \Sigma_{b+k|b} C^\top + R). \quad (21)$$

G. Evaluation Protocol

We construct standardized blackout evaluation windows by selecting contiguous fully observed segments and masking them to simulate blackouts of varying lengths, enabling direct comparison to ground truth. We use stratified month sampling (equal windows per month) and align forecast targets so that horizons $\{1, 3, 6\}$ are evaluated on the same set of window identifiers. We report RMSE (mph) for (i) imputation inside the blackout and (ii) post-blackout forecasting, and we compute uncertainty intervals via bootstrap resampling over evaluation windows.

IV. RESULTS

A. Experimental Setup and Evaluation Protocol

Datasets. Our primary evaluation uses the **Seattle Inductive Loop Detector dataset (2015)**, consisting of traffic speed (mph) measured at 5-minute intervals across 147 detectors over one year. The dataset exhibits frequent, structured missingness, including multi-hour to multi-day sensor blackouts, with roughly 5% missing readings overall.

To assess generality under controlled missingness mechanisms, we additionally report experiments on **METR-LA** with synthetically injected blackout events whose occurrence depends on the underlying traffic state (Section IV-C).

Tasks. We evaluate two prediction tasks: (i) *blackout imputation*, measuring reconstruction accuracy within blackout windows; and (ii) *post-blackout forecasting* immediately after blackout termination at horizons of 1, 3, and 6 steps (i.e., 5, 15, and 30 minutes ahead). Forecasting is evaluated as a point prediction at the horizon endpoint (the h -step-ahead value).

Validation windows (aligned across horizons). To ensure strict comparability across horizons, we form an evaluation pool of blackout events that have *matching* imputation and forecasting windows for *all* horizons $\{1, 3, 6\}$ under the same `window_id`. From this aligned pool, we sample **25 blackout windows per calendar month** (stratified by the blackout start timestamp), yielding **300 imputation windows** total. The corresponding 1/3/6-step forecasting sets are constructed by retrieving the matching windows for the same ordered `window_ids`, yielding **300 windows per horizon**.

Leakage-free masking protocol. All methods are evaluated on the same held-out blackout windows using an identical masking protocol: for each selected evaluation blackout interval (detector d , start s , end e), we *artificially mask* $x_{s:e,d}$ in a training panel x^{train} and record an indicator a_t for artificially masked entries. To obtain ground truth, we only mask segments that are fully observed in the original panel; naturally missing readings remain in x^{train} but are never used as evaluation targets. Models are trained/inferred only on x^{train} and evaluated against the original unmasked data x .

Models. We compare:

- **LOCF**: last-observation-carried-forward for both imputation and forecasting.
- **LinearInterp + SeasonalNaive**: linear interpolation within blackout windows (using the last observed pre-blackout and first observed post-blackout values), and a seasonal naive forecast using daily/weekly lags with fallback to the last observed value.
- **MAR LDS**: a linear dynamical system trained via EM with missing observations handled through Kalman filtering and RTS smoothing (missingness is *not* modeled as an observation channel).
- **MNAR LDS**: our proposed model, augmenting the LDS with a latent-state-dependent missingness observation model.

Hyperparameters and training. We found latent dimension $K = 20$ to consistently minimize validation error across

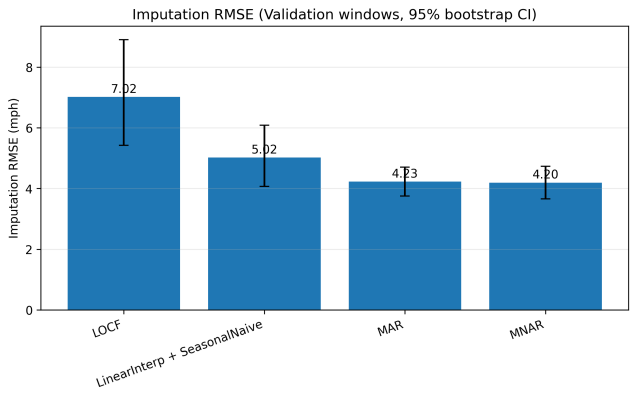


Fig. 1. Seattle Loop (2015) imputation RMSE on aligned validation windows with 95% bootstrap confidence intervals.

TABLE I
SEATTLE LOOP (2015): RMSE (MPH) ON ALIGNED VALIDATION WINDOWS (25 PER MONTH; 300 TOTAL). LOWER IS BETTER.

Method	Impute	1-step	3-step	6-step
LOCF	7.021	7.860	8.465	8.942
LinearInterp + SeasonalNaive	5.024	8.734	8.442	8.881
MAR (LDS)	4.229	4.391	4.349	5.160
MNAR (LDS)	4.195	4.313	4.228	5.104

imputation and short-horizon forecasting. Both LDS variants are trained with **10 EM iterations**; MNAR is **warm-started from the trained MAR parameters** and then run for an additional **10 EM iterations** with missingness parameters updated (two gradient steps per EM iteration, learning rate 10^{-4}). Unless noted otherwise, we report **RMSE (mph)**.

B. Blackout Imputation and Forecasting on Seattle Loop

1) *Main Quantitative Results (Aligned Validation Windows)*: Table I reports the main results for the aligned, stratified validation protocol described above (300 windows per horizon). State-space modeling provides the dominant improvement: moving from LOCF to an LDS reduces error substantially for both imputation and forecasting. On top of this, MNAR provides a consistent refinement over MAR across all tasks, with the largest gains at the intermediate forecast horizon.

Figure 1 visualizes the imputation results with 95% bootstrap confidence intervals, highlighting that the MNAR–MAR gap is small but consistently favors MNAR.

2) *Qualitative Examples: What the Models Do During a Blackout*: Figure 2 provides representative blackout reconstructions under the leakage-free masking protocol. LOCF produces a piecewise-constant fill that cannot track within-blackout evolution, while the LDS-based models infer a smooth latent trajectory that continues through the masked interval. MNAR and MAR are typically close visually, but MNAR tends to make slightly sharper adjustments when the missingness pattern is informative about the latent traffic

regime, consistent with the small but reliable average RMSE gains in Table I.

3) *State-Space Dynamics Are the Primary Driver*: Compared to LOCF, MAR reduces MSE by *roughly two-thirds* across tasks (e.g., from 49.30 to 17.88 MSE for imputation, and from 61.78 to 19.28 MSE for 1-step forecasting). This large gap shows that explicitly tracking a latent trajectory during outages is substantially more effective than forward filling, even for short horizons.

Moreover, a stronger heuristic baseline (LinearInterp + SeasonalNaive) improves imputation over LOCF (7.02 \rightarrow 5.02 RMSE), but does not improve forecasting, highlighting that reconstruction inside the blackout does not directly translate into accurate post-blackout dynamics.

Figure 3 stratifies imputation error by blackout length within the aligned-window evaluation pool. LOCF degrades sharply as blackout length increases, while LDS-based methods degrade much more gradually, consistent with the benefit of maintaining a latent trajectory through outages.

4) *MNAR Provides a Consistent Refinement Over MAR*: While most gains come from modeling dynamics, MNAR improves over MAR across imputation and all forecast horizons:

$$\Delta\text{RMSE}(\text{MNAR} - \text{MAR}) = (-0.033, -0.078, \\ -0.121, -0.056)$$

for (impute, 1-, 3-, 6-step), respectively. These improvements indicate that the missingness pattern carries additional predictive signal beyond what is captured by the observation model alone, and that treating missingness as a probabilistic observation channel modestly but consistently sharpens latent state estimates during blackouts.

5) *Forecast Horizon Effects*: We observe a mild non-monotonic pattern where 3-step forecasting can outperform 1-step forecasting for both MAR and MNAR (Table I). A plausible explanation is a bias-variance tradeoff at blackout termination: the filtered state at the blackout endpoint can be biased due to prolonged missingness, and short rollouts may partially correct this bias through the learned dynamics, while longer horizons accumulate variance and degrade performance.

Figure 4 visualizes this trend and the associated uncertainty via bootstrap confidence intervals; MNAR remains modestly better than MAR across horizons, though the gap is small relative to window-to-window variability.

6) *EM Training Behavior*: Figure 5 shows the training objective over iterations for MAR and MNAR under our approximate EM pipeline. In our runs, the objective increases steadily and MNAR attains higher values than MAR, suggesting that modeling missingness improves overall fit without introducing instability.

7) *Missingness Diagnostics: Evidence for State-Dependence*: To probe whether missingness is predictable from the latent traffic state, we ran lightweight diagnostics:

Test 1 (onset vs matched control, observed-only features). A logistic classifier using only observed proxies

(last observed speed, short-term variance, and time features) achieves AUC ≈ 0.533 when distinguishing blackout onsets from matched controls, suggesting limited predictive power from raw observed edges alone.

Test 2 (next-step missingness prediction, time-aware split). When augmenting the observed features with smoothed latent state features from the MAR smoother, ROC-AUC increases from ≈ 0.518 (observed-only) to ≈ 0.661 (latent-augmented), indicating that the inferred latent trajectory contains information relevant to predicting missingness. While these diagnostics are not used for model selection, they support the interpretation that missingness is not purely random and can correlate with the underlying system state.

Test 3 (event structure). Blackout lengths are heavy-tailed: the median detector blackout is 37 steps (≈ 3 hours) with a 75th percentile of 84 steps (≈ 7 hours), but rare events extend to extremely long durations. Network-wide outages are much rarer (25 events) and substantially shorter on average, consistent with detector-specific failure modes dominating the missingness structure.

8) *Inference Ablation: Missingness Variance Modeling*: We tested two inference variants for the MNAR missingness observation variance (moment-matched vs constant variance) and found nearly identical performance (differences < 0.003 RMSE across horizons). This indicates that the MNAR gains are not driven by fragile variance tuning, but by the missingness signal itself.

9) *Robustness Across Random Initializations*: Across five independent training runs (different random seeds), MNAR consistently improves imputation RMSE on average, while forecasting differences are near zero relative to run-to-run variability. Concretely, the mean MNAR-MAR delta is approximately -0.026 RMSE for imputation, and within ± 0.03 RMSE for forecasting horizons $\{1, 3, 6\}$, suggesting the MNAR refinement is reliable for reconstruction and does not degrade short-horizon forecasts.

Figure 6 further shows that error is heterogeneous across both blackout length and time-of-day. MNAR typically matches or slightly improves upon MAR across most buckets, indicating that the refinement is not confined to a single regime.

C. Generalization to METR-LA with Synthetic Blackouts

On METR-LA, we inject synthetic blackout events whose occurrence depends on the underlying traffic state, creating a controlled MNAR setting in which missingness carries information about the latent regime. Table II reports the same two tasks as Seattle Loop (blackout imputation and post-blackout forecasting at horizons $\{1, 3, 6\}$).

1) *Main Quantitative Results on METR-LA*: State-space modeling again provides the dominant gain over naive imputation: MAR reduces RMSE substantially relative to LOCF for both reconstruction and forecasting. On top of this, MNAR provides a consistent refinement over MAR for imputation and longer-horizon forecasting, with the largest improvement at the 6-step horizon.

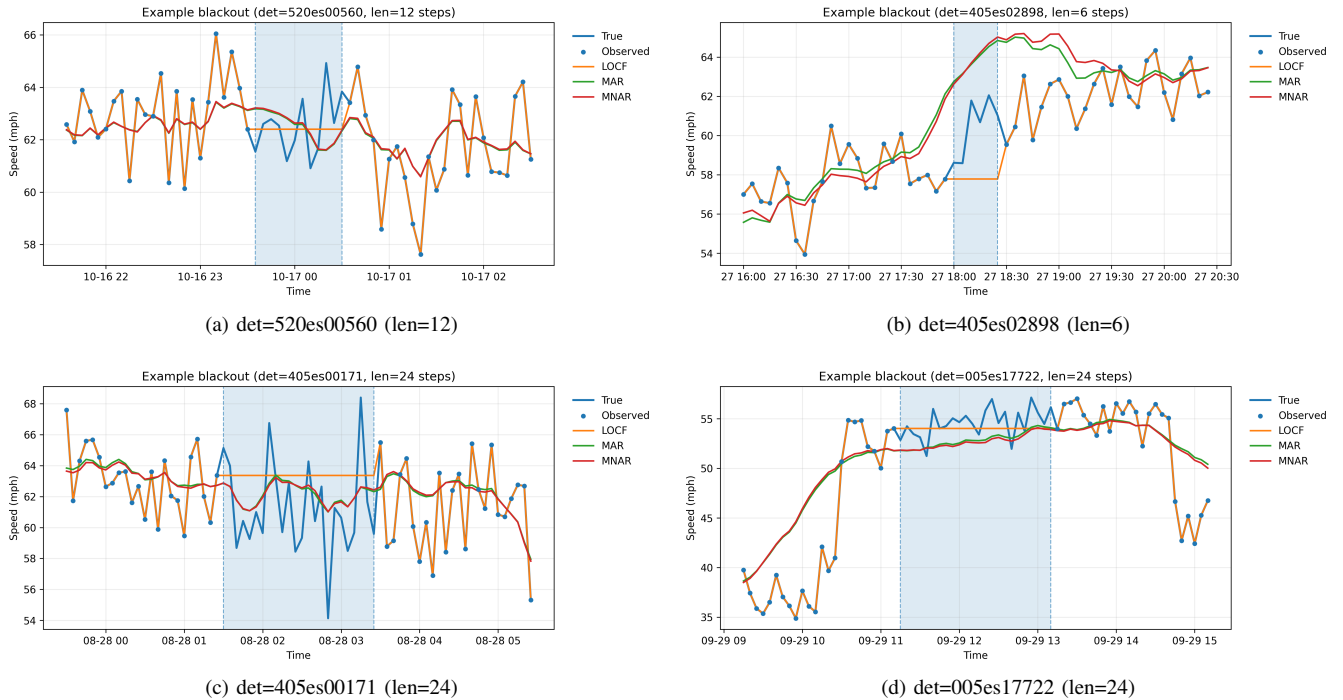


Fig. 2. Qualitative reconstructions on Seattle Loop. Shaded regions denote artificially masked blackout intervals. LOCF freezes at the last observed value, while LDS models infer a latent trajectory through the outage; MNAR and MAR are often close but MNAR can yield slightly sharper latent corrections when missingness correlates with the underlying state.

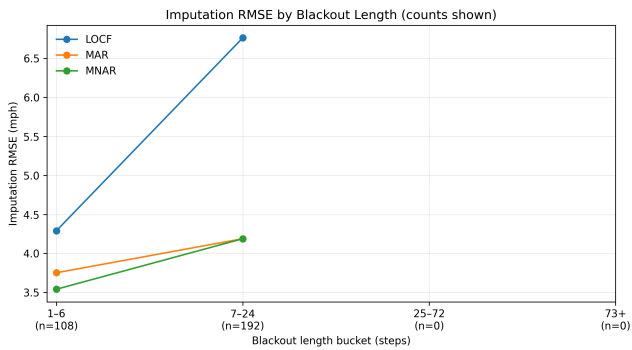


Fig. 3. Imputation RMSE by blackout length bucket (counts shown for the aligned evaluation pool). LOCF error grows rapidly with longer blackouts, while LDS-based imputations are substantially more stable.

TABLE II
METR-LA (SYNTHETIC STATE-DEPENDENT BLACKOUTS): RMSE ON HELD-OUT BLACKOUT WINDOWS (LOWER IS BETTER).

Method	Impute	1-step	3-step	6-step
LOCF	9.433	9.855	9.262	9.147
MAR (LDS)	5.483	4.700	5.320	5.114
MNAR (LDS)	5.355	4.700	5.238	4.790

MNAR improves over MAR by

$$\Delta \text{RMSE}(\text{MNAR} - \text{MAR}) = (-0.129, -0.001, -0.082, -0.324)$$

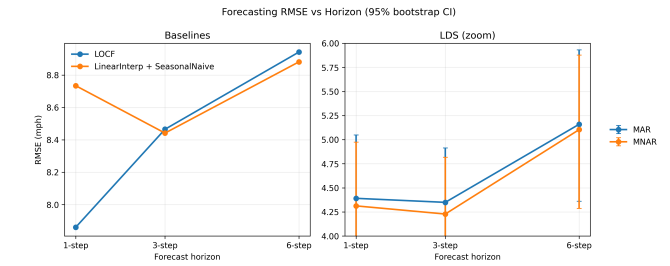


Fig. 4. Seattle Loop forecasting RMSE vs horizon with 95% bootstrap confidence intervals. State-space modeling dominates the gain over heuristic baselines; MNAR provides a small, consistent refinement over MAR across horizons.

for (impute, 1-, 3-, 6-step), respectively. The near-zero 1-step delta suggests that immediate post-blackout prediction is dominated by the shared LDS dynamics, while the larger 6-step improvement indicates that state-dependent missingness can meaningfully sharpen latent state estimates in ways that matter more for longer rollouts.

V. CONCLUSION

In this paper, we studied structured sensor blackouts in traffic time series and proposed an MNAR-aware state-space model that treats dropout indicators as an additional observation channel whose probability depends on the latent traffic state. Using approximate inference (EKF/RTS) and an EM-style learning procedure, we find that modeling temporal dynamics is the primary driver of performance, while explicit

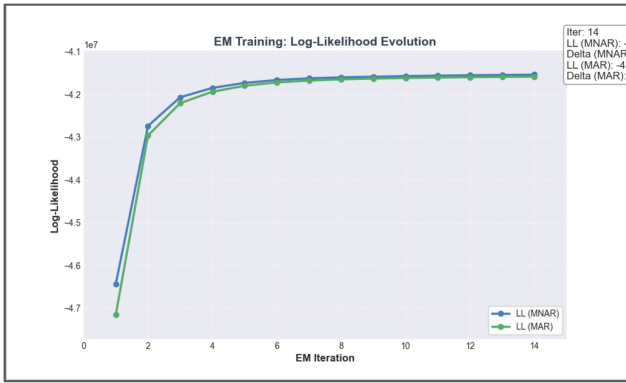


Fig. 5. Approximate-EM training objective (log-likelihood) for MAR and MNAR on Seattle Loop. In our runs it increases steadily and MNAR attains higher values.

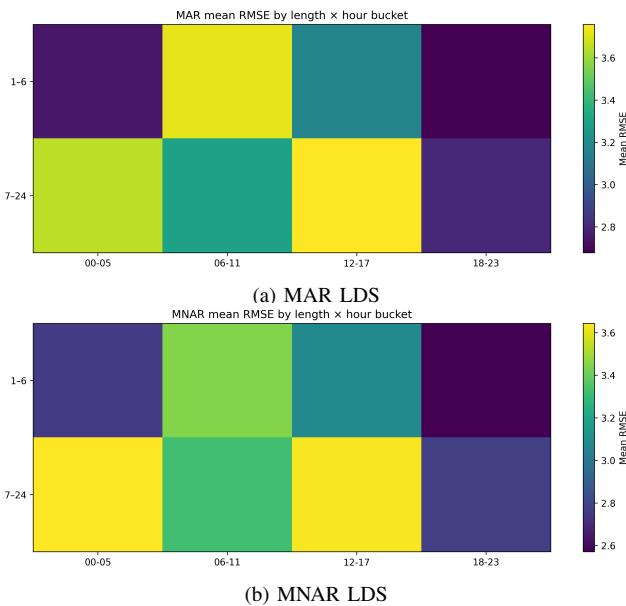


Fig. 6. Mean imputation RMSE stratified by blackout length bucket and hour-of-day bucket (Seattle Loop, aligned evaluation pool). MNAR generally matches or modestly improves upon MAR across buckets, suggesting the refinement is not limited to a single time regime.

MNAR modeling provides a principled refinement when missingness is informative.

On the Seattle Loop dataset, introducing latent dynamics yields large gains over heuristic baselines, reducing imputation RMSE from 7.02 (LOCF) and 5.02 (LinearInterp + Seasonal-Naive) to 4.23 (MAR LDS). Incorporating MNAR missingness further improves imputation to 4.20 and yields consistent (though modest) improvements in post-blackout forecasting, with the clearest gain at the intermediate horizon (3-step RMSE $4.35 \rightarrow 4.23$). On METR-LA with injected state-dependent blackouts, the MNAR advantage becomes more pronounced, improving imputation RMSE from 5.48 to 5.35 and 6-step forecasting RMSE from 5.11 to 4.79, consistent with the hypothesis that MNAR modeling helps most when dropout is causally linked to the underlying system state.

This work has several limitations. Our dynamics are linear-Gaussian, the missingness model uses a simple logistic parameterization, and inference relies on local linearization, which may be insufficient in strongly nonlinear regimes. In addition, we do not incorporate richer covariates (e.g., incidents, weather, or detector health) that could better explain dropout.

Future work includes extending MNAR-aware inference to nonlinear or graph-structured state-space models, using variational inference to better capture posterior uncertainty, and enriching the missingness model with explicit covariates (e.g., time-of-day/week, network-wide outage indicators, and detector health statistics) to improve identification and robustness in real deployments.

REFERENCES

- [1] Y. Wang and M. Papageorgiou, “Real-time freeway traffic state estimation based on extended Kalman filter: A general approach,” *Transportation Research Part B: Methodological*, vol. 39, no. 2, pp. 141–167, 2005.
- [2] W. Li, M. Jiang, Y. Chen, and M. C. Lin, “Estimating urban traffic states using iterative refinement and Wardrop equilibria,” *IET Intelligent Transport Systems*, vol. 12, no. 8, pp. 875–883, 2018.
- [3] B. Yu, H. Yin, and Z. Zhu, “Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting,” in *Proc. 27th Int. Joint Conf. Artificial Intelligence (IJCAI)*, 2018, pp. 3634–3640.
- [4] Z. Che, S. Purushotham, K. Cho, D. Sontag, and Y. Liu, “Recurrent neural networks for multivariate time series with missing values,” *Scientific Reports*, vol. 8, no. 1, Art. no. 6085, 2018.
- [5] W. Cao, D. Wang, J. Li, H. Zhou, L. Li, and Y. Li, “BRITS: Bidirectional recurrent imputation for time series,” *arXiv:1805.10572*, 2018.
- [6] D. B. Rubin, “Inference and missing data,” *Biometrika*, vol. 63, no. 3, pp. 581–592, 1976.
- [7] R. J. A. Little and D. B. Rubin, *Statistical Analysis with Missing Data*, 2nd ed. Hoboken, NJ, USA: Wiley, 2002.
- [8] P. Diggle and M. Kenward, “Informative drop-out in longitudinal data analysis,” *J. Roy. Statist. Soc. Ser. C (Appl. Statist.)*, vol. 43, no. 1, pp. 49–93, 1994.

APPENDIX A

REPRODUCIBILITY AND EVALUATION DETAILS

A. Code

Our implementation (model, training, evaluation, and plotting) is available at: <https://github.com/BlackoutBayes/Modeling-Information-Blackouts-in-MNAR-Time-Series>.

B. Aligned validation windows

We evaluate blackout imputation and post-blackout forecasting using leakage-free masking on a fixed set of held-out windows. To ensure strict comparability across horizons, we construct an aligned pool where each `window_id` has (i) an imputation blackout interval and (ii) forecast targets for all horizons $\{1, 3, 6\}$. From this aligned pool, we sample 25 windows per month (stratified by blackout start time), yielding 300 windows total, and reuse the same ordered `window_ids` for each forecasting horizon.

C. Training protocol

Both LDS variants are trained for 10 EM iterations. MNAR is warm-started from the converged MAR parameters and trained for an additional 10 EM iterations with missingness updates enabled (two gradient steps per EM iteration, learning

rate 10^{-4}). Artificially masked evaluation-window entries are excluded from missingness updates to avoid leakage.

APPENDIX B
DATASETS

- **Seattle Inductive Loop Detector Dataset (2015):** <https://github.com/zhiyongc/Seattle-Loop-Data>
- **METR-LA:** distributed as `metr-la.h5` and used for synthetic blackout experiments.

APPENDIX C
SYNTHETIC MNAR-STRENGTH SWEEP

To validate that MNAR modeling helps most when missingness is genuinely state-dependent, we sweep the dependence strength α in a controlled synthetic setting. Results are summarized in Table III, reporting mean \pm std RMSE across seeds and $\Delta\text{RMSE} = \text{RMSE}_{\text{MNAR}} - \text{RMSE}_{\text{MAR}}$ (negative favors MNAR). We observe that MNAR tends to yield larger improvements under stronger state dependence, though variability increases at higher α .

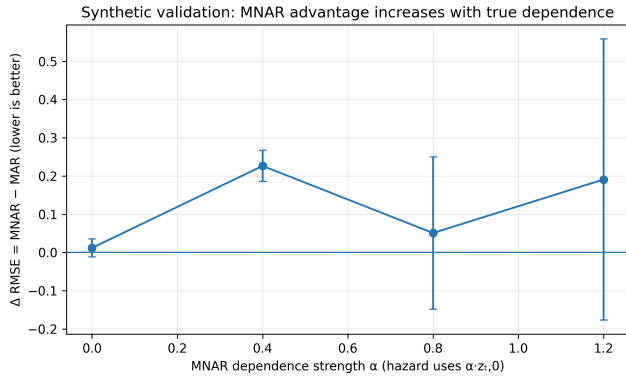


Fig. 7. Synthetic validation: MNAR advantage increases as the true state-dependence strength α increases. Points show mean $\Delta\text{RMSE} = \text{RMSE}_{\text{MNAR}} - \text{RMSE}_{\text{MAR}}$ across seeds; error bars show ± 1 std.

TABLE III
SYNTHETIC MNAR-STRENGTH SWEEP: RMSE (MEAN \pm STD) ACROSS RANDOM SEEDS. ΔRMSE DENOTES $\text{RMSE}_{\text{MNAR}} - \text{RMSE}_{\text{MAR}}$ (NEGATIVE FAVORS MNAR).

α	MAR RMSE	MNAR RMSE	ΔRMSE
0.0	3.741 \pm 0.026	3.730 \pm 0.023	-0.011 \pm 0.007
0.4	3.858 \pm 0.082	3.630 \pm 0.063	-0.227 \pm 0.033
0.8	3.750 \pm 0.445	3.698 \pm 0.538	-0.052 \pm 0.109
1.2	3.849 \pm 0.510	3.659 \pm 0.946	-0.190 \pm 0.372