Benchmarking Multivariate Time-Series Imputation in 6G Networks

Author: Alessandro Neri | **Supervisors:** Rihan Hai, Yuandou Wang

1. The Problem: Missing time-series data from 6G datasets

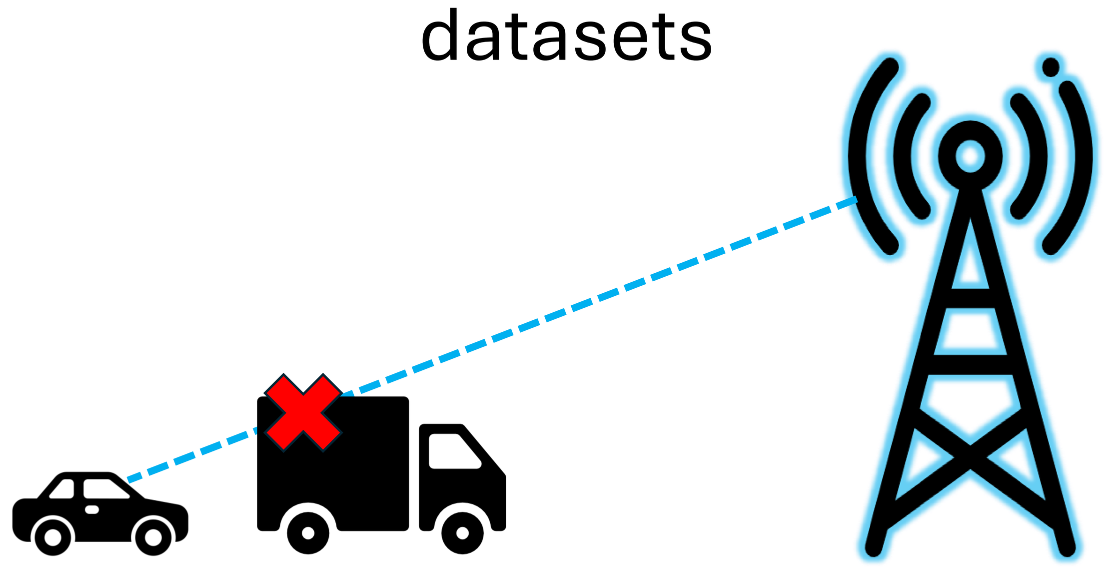

6G mmWave provides massive bandwidth but suffers from physical blockages. These create bursty telemetry outages that blind live network orchestrators.

2. Research Questions

- **RQ1:** Which architecture reconstructs data best?
- **RQ2:** How do gap size and frequency impact performance?
- **RQ3:** What is the Pareto-optimal model for live edge routing?

3. Evaluated Architectures

Nearest Neighbor	Baseline, closest temporally available value
Kalman Filter	Classical, recursive state-space updates
BRITS	Recurrent (RNN), merges past and future context
CSDI	Generative AI, conditional reverse-denoising
TimesNet	Convolutional (CNN), multi-periodic 2D FFTs

4. Experimental Setup

- **Datasets:** EURECOM 5G Core AMF Telemetry & Python Web Server
- **Gap Injection:** Parametrized to test different scenarios
  - **Missing Ratios:** 10%, 25%, 40%, and 50%
  - **Gap Sizes:** 1, 5, 10, 20, 30, and 60 seconds
- **Evaluation Metrics:**
  - **RMSE** for point-wise reconstruction accuracy
  - **Latency** measured in milliseconds (only inference, no training)

5. Methodology

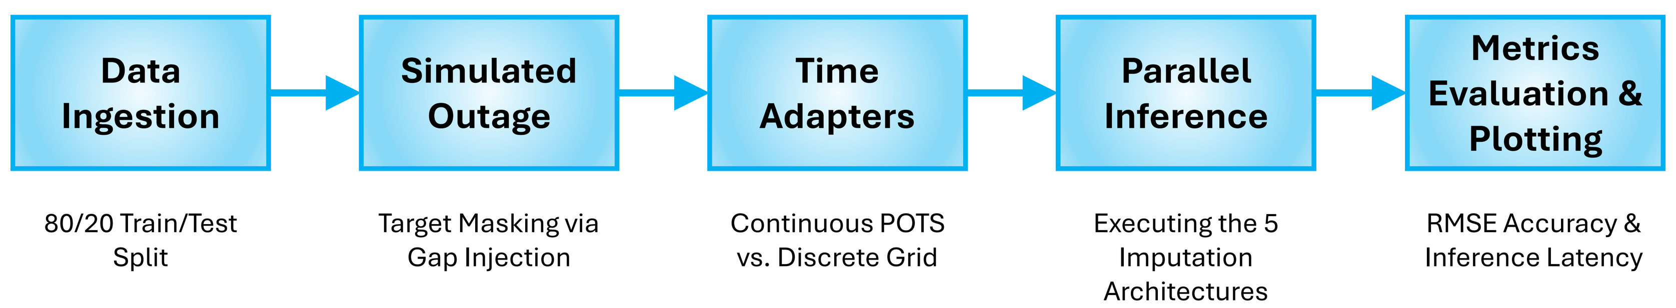

The end-to-end Airflow pipeline, datasets, and a custom interactive **Streamlit dashboard** are fully open-sourced.

6. RQ1 - Imputation Accuracy

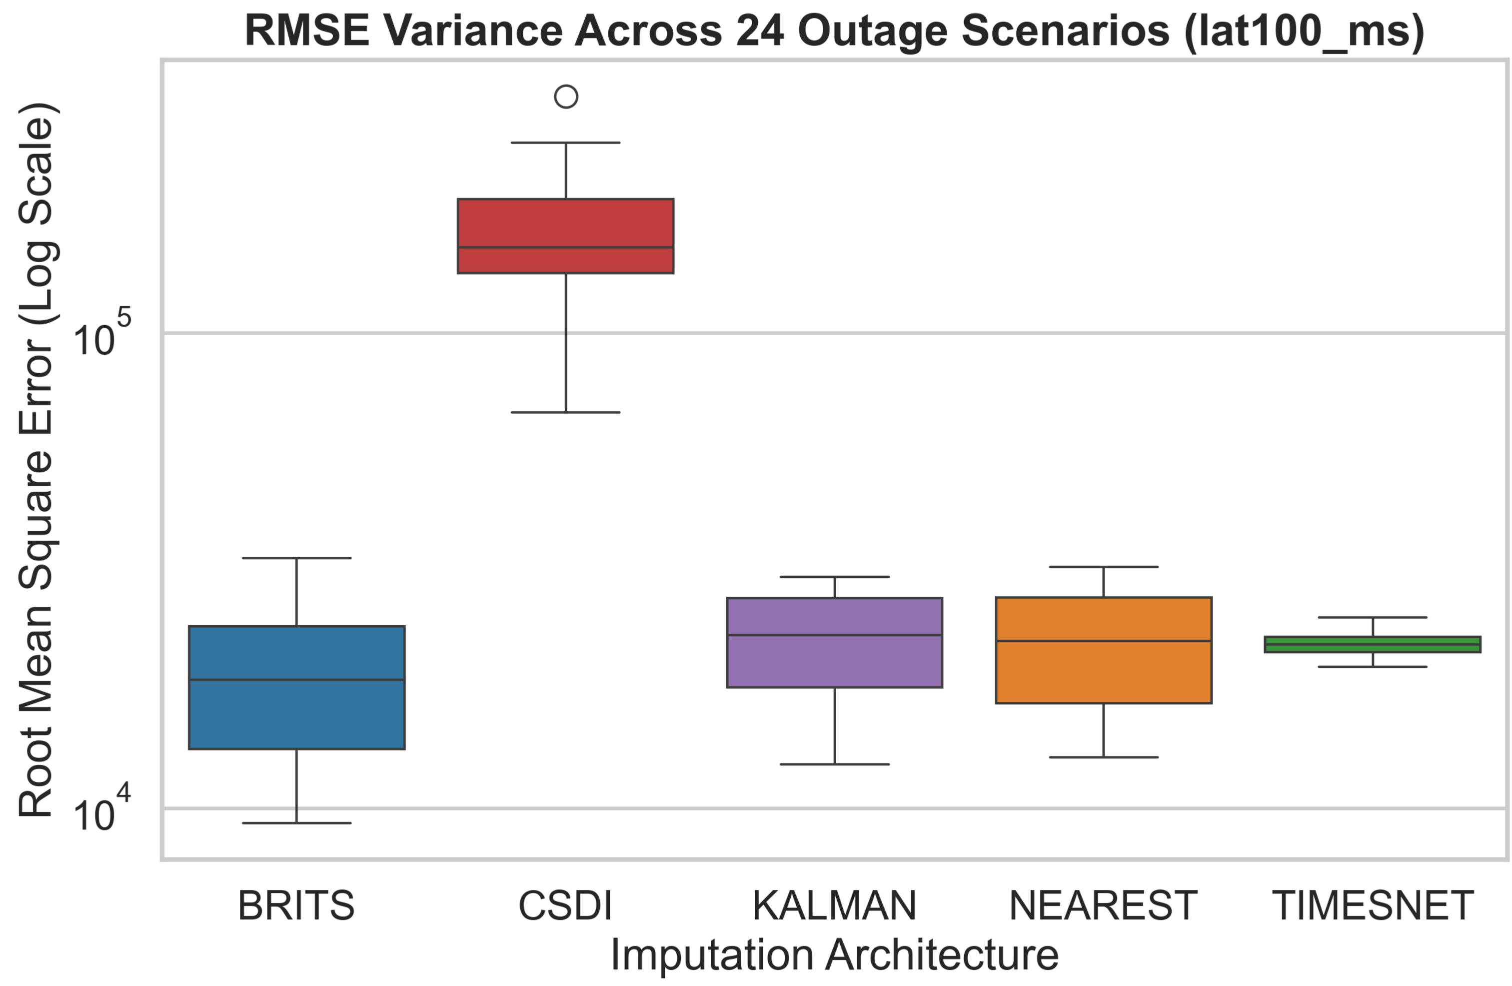

1. **BRITS** achieves the highest average reconstruction fidelity
2. The generative **CSDI** model exhibits massive instability

7. RQ2 - Degradation under Prolonged Outages

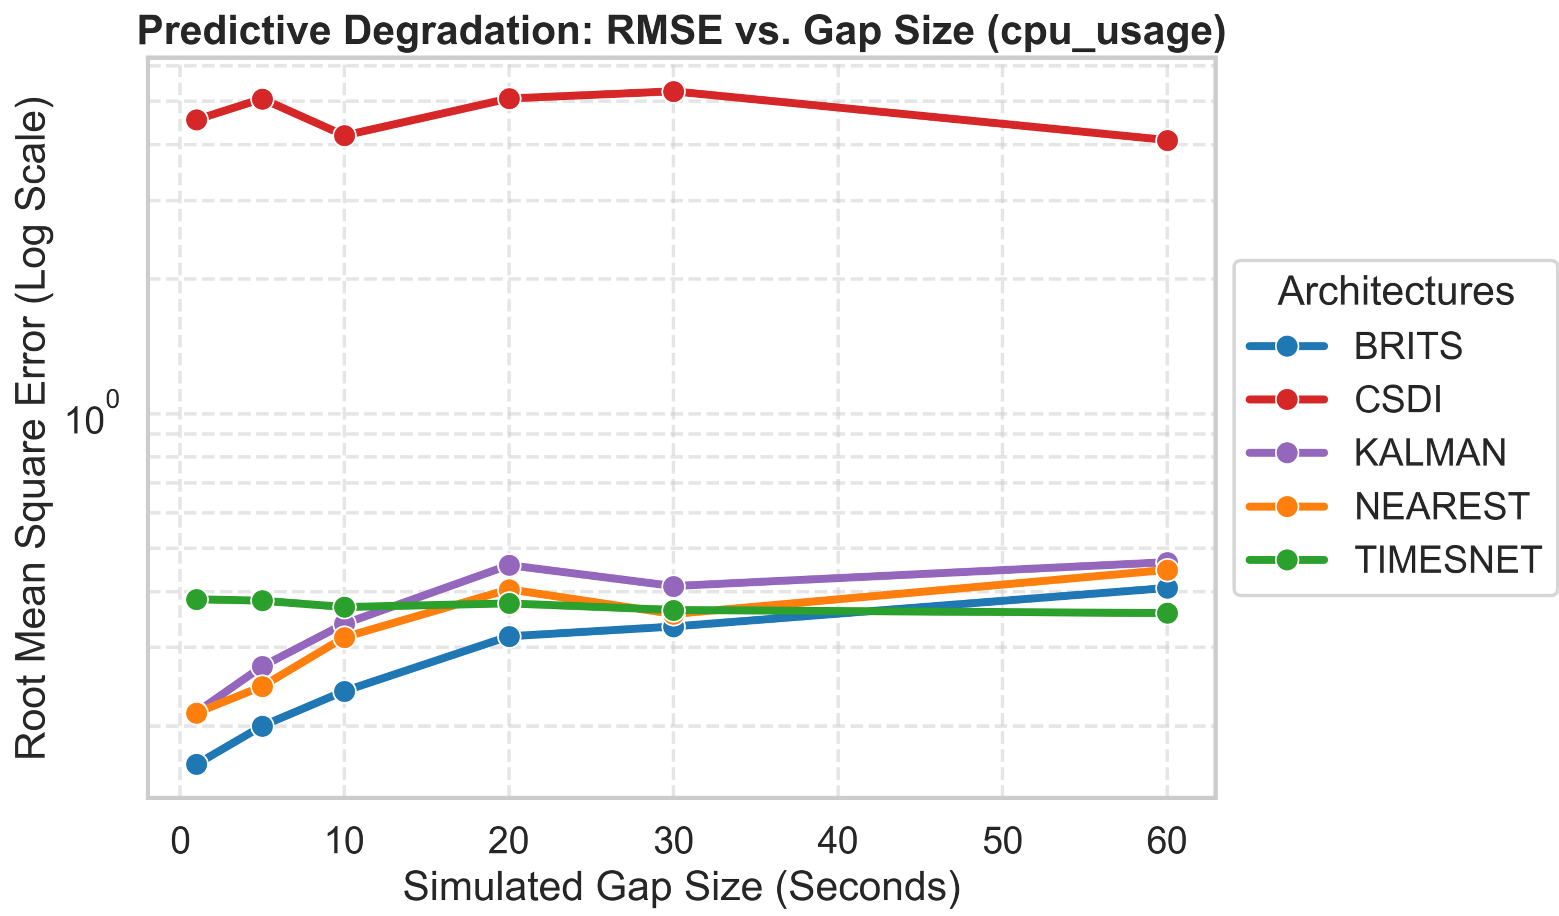

1. Degradation is **not strictly linear**
2. TimesNet & CSDI keep same accuracy across various scenarios
3. BRITS, Nearest Neighbor and Kalman Filter stabilize from 20-second gaps onward

8. RQ3 - Pareto-Optimal Trade-off

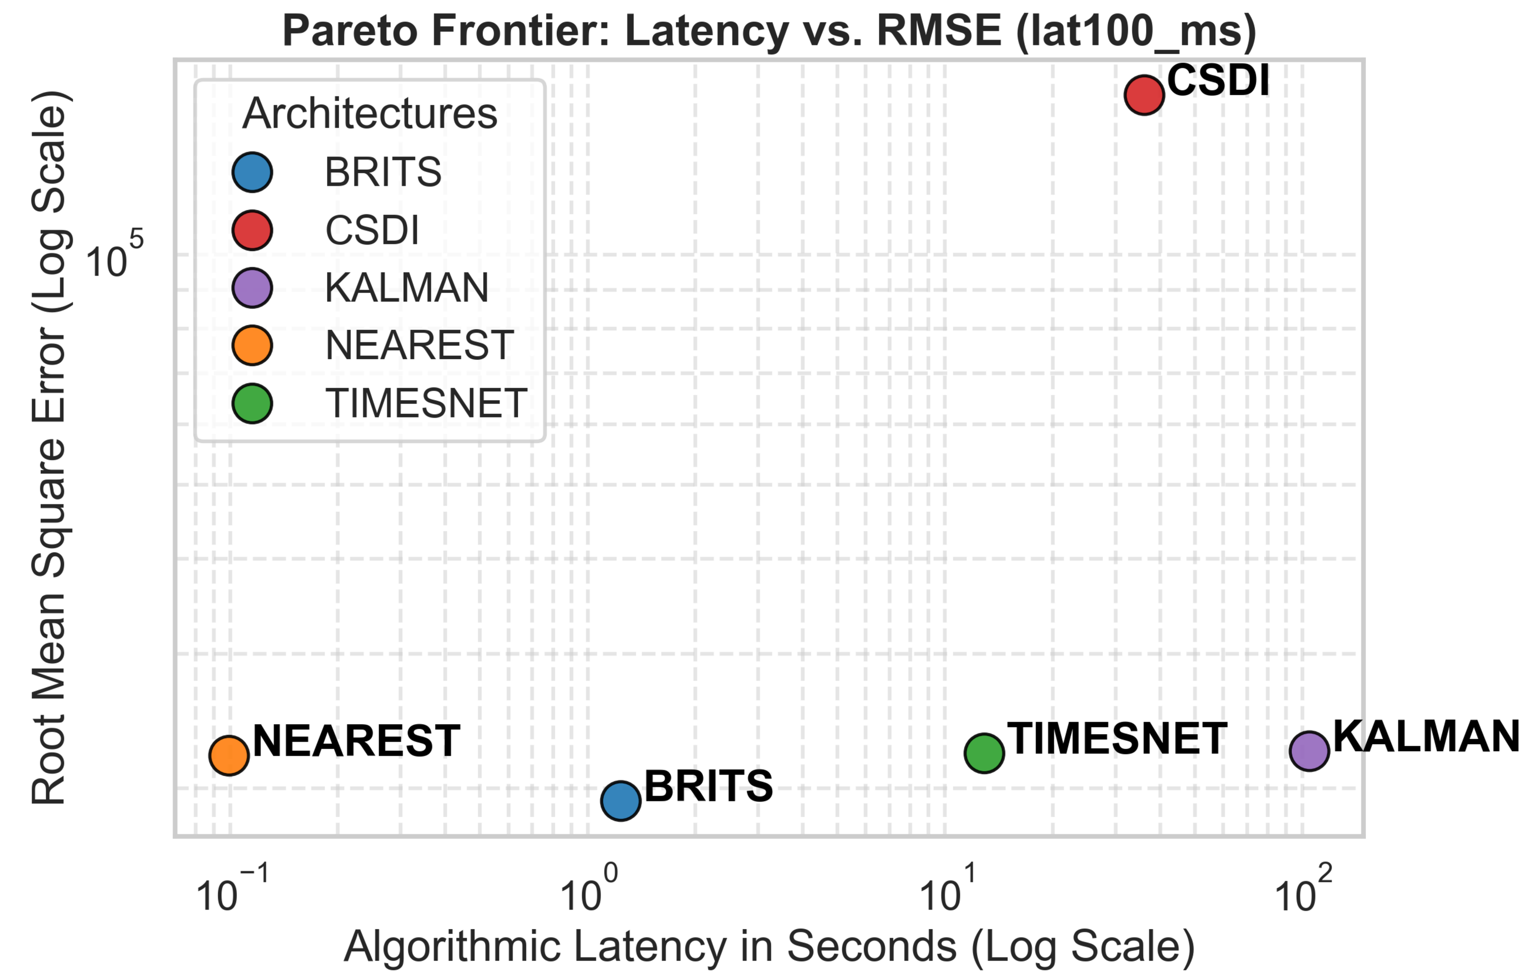

1. **Nearest Neighbor** dominates immediate latency (<250ms)
2. **BRITS** dominates inference accuracy while keeping a competitive latency (1.2s per inference)
3. **Kalman Filter** suffers severe sequential bottlenecks (1.7 mins per inference)

9. Conclusions

- **The Pareto-Optimal Winners: Nearest Neighbor** dominates ultra-low latency requirements, while **BRITS** represents the ideal deep learning compromise for high-fidelity spike reconstruction.
- **The Generative Failure:** Diffusion models (**CSDI**) fundamentally fail in this domain; their lack of awareness regarding physical boundaries results in exponential reconstruction errors.
- **POTS vs. Discretization:** Architectures that process continuous timestamps (POTS) bypass the discretization penalties that cause latency bottlenecks in other models, such as the **Kalman Filter**.

10. Limitations and Future Work

- **Online vs Offline Inference:** The models achieving peak accuracy (**BRITS**, **Nearest Neighbor**) utilized *offline, batched* processing to merge past and future context. However, live 6G edge routing algorithms would have no access to future observations.
- **Edge Hardware Deployment:** To measure true operational latency and bottlenecks, future experiments should be deployed directly onto edge devices, such as **Raspberry Pi clusters**.