Capturing Spatiotemporal Dynamics and Predicting LEO ISP Performance Variations

Forecasting Starlink Connectivity: A Data-Driven, Spatiotemporal Analysis Integrating Weather and Satellite Density

Author: Cristian Benghe, Supervisor: Tanya Shreedhar, Professor: Nitinder Mohan

TU Delft

1. Research Questions & Contributions

- Can Starlink **network quality** be accurately predicted using **weather** and **satellite data** through machine learning? What is the achievable resolution for **space** and **time** in these forecasts?
- Does adding meteorological, satellite features and preprocessing improve forecasts and generalization beyond simple baselines?
- Built a **reproducible** ML pipeline with data cleaning, a custom **Weather Index**, and satellite density features, all in an **interactive** global tool that updates predictions daily producing hourly forecasts for internet quality.
- Achieved improved **latency** and **throughput** prediction over baselines; model generalizes well, though **jitter** and **packet loss** remain challenging.

2. Methodology - from Training to Predicting

Data acquisition

MLab (NDT7) > Starlink Data (BigQuery, SQL) - taking median from 10 MLab measurements → Add Latitude & Longitude (from worldcities.csv or country fallback) → Weather Index (WI) Enrichment, Range: 0 - Clear to 1 - Severe, enriched with Meteostat matched by time and location → Compute local satellite density at each point and time using CelesTrak TLEs and SGP4/ Skyfield, counting satellites within 500 km.

Data processing

Timestamp Feature Engineering (hour, day of week, month) → Group by location and 3h windows and remove 38% worst outliers (combined badness score) → Split in 80% train set data and 20% test set data → Machine Learning Models (Random Forest, Gradient Boosting, etc.) - take the best one for each model

MLAB NDT7 (Network Diagnostic Tool) → Meteostat

$$WI = 0.7 \cdot \text{cloudiness} + 0.15 \cdot \text{wind\_speed} + 0.1 \cdot \text{snow} + 0.05 \cdot \text{temp}$$

This index captures how local meteorological conditions impact Starlink signal quality, allowing the model to anticipate weather-driven performance changes. →

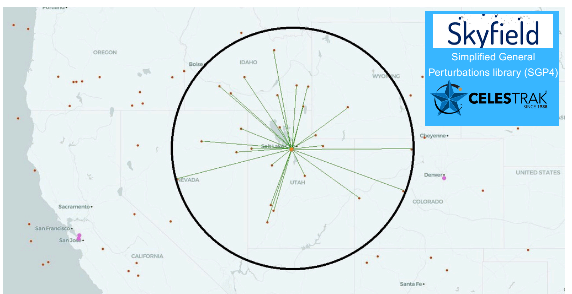

Feature Engineering - fixed positions

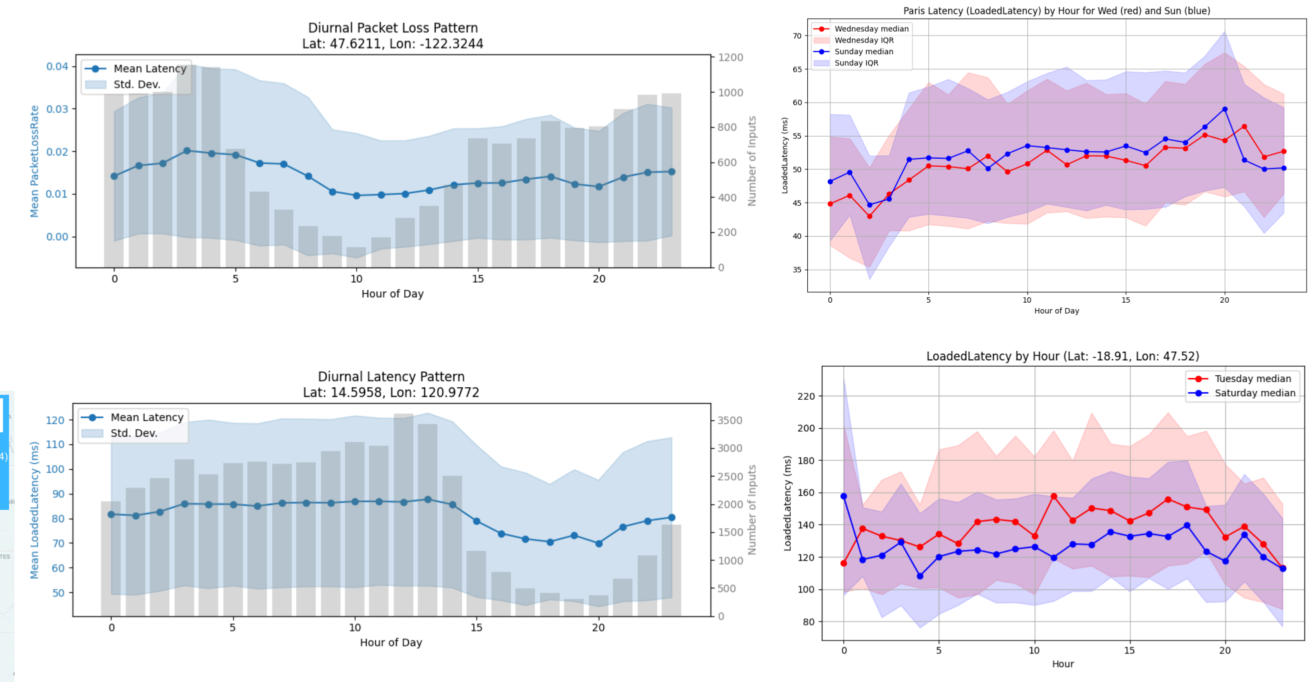

Outlier removal grouped by location and time

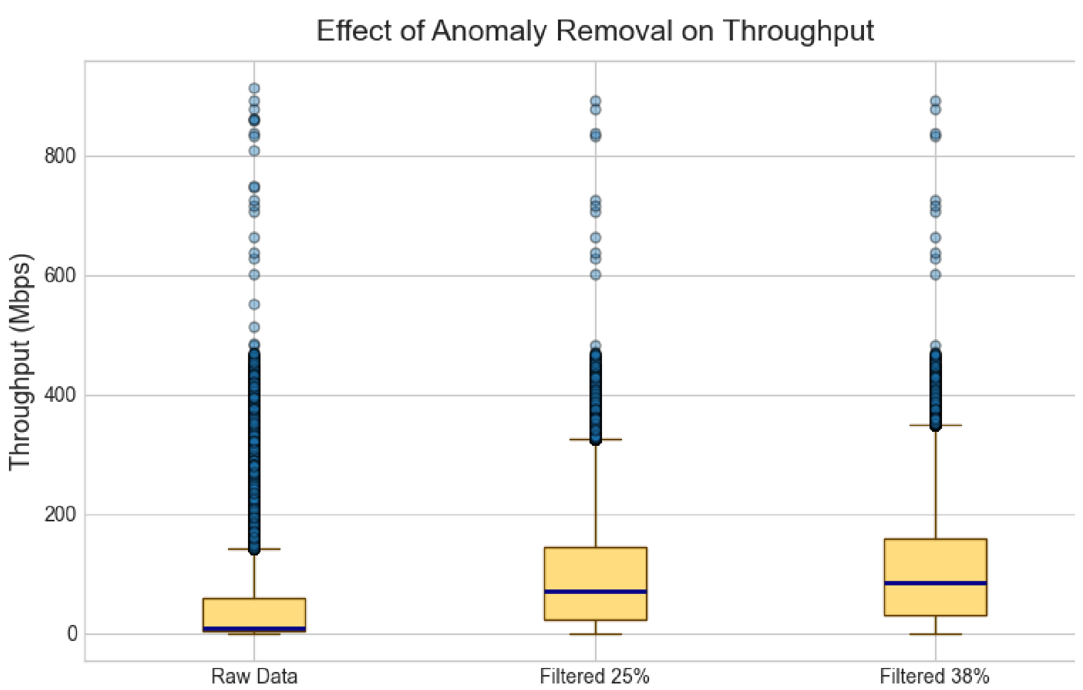

Worst 25% / 38% outliers by composite "badness" score (combining packet loss, latency, jitter, and throughput) score are removed within each 3-hour window, reducing noise and improving model accuracy

Interpretable ML Model Training

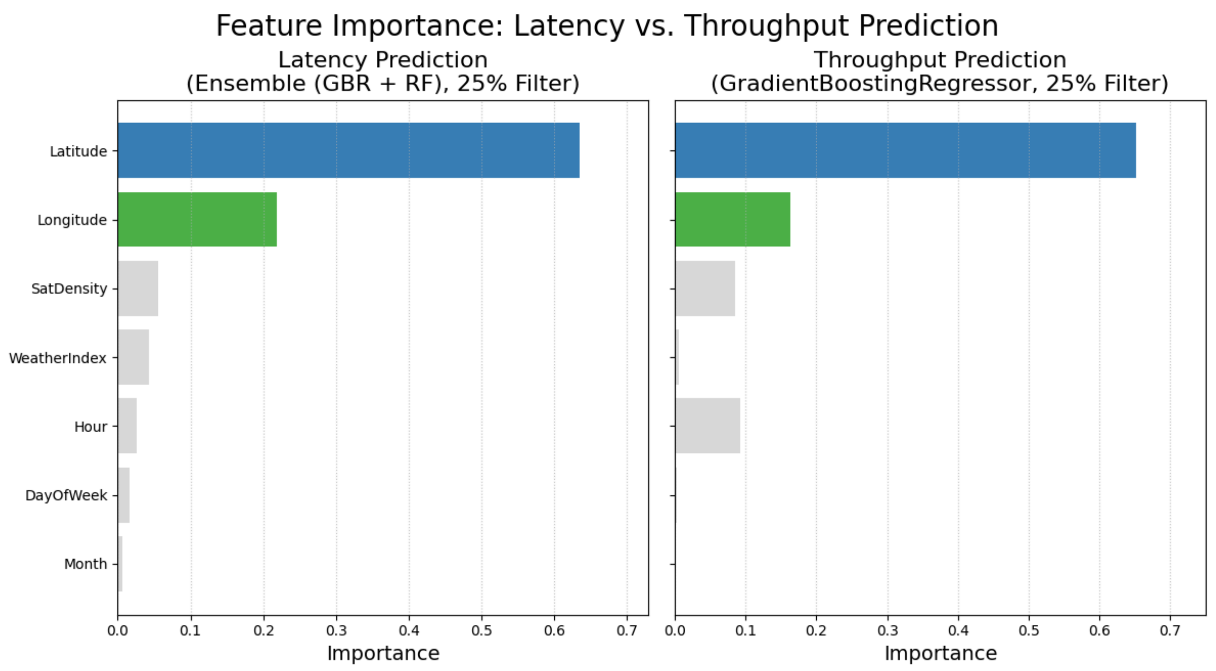

We train on 80% of the data and evaluate on the remaining 20%. For each target, we select and save the best-performing model—Gradient Boosting, Random Forest, or their weighted ensemble.

Use the Median MLab (NDT7) to get all unique cities for 48h predictions starting tomorrow → Get point Lat + Lon using worldcities.csv for all cities → For each location, compute 48 hourly records (next 48 hours), as hour, weekday, and month. → Weather Index (WI) Enrichment, Range: 0 - Clear to 1 - Severe, enriched via OpenWeather matched by time and location → Compute local satellite density per point & time using CelesTrak orbital data (TLEs) and SGP4 propagation via Skyfield, count...

Get all resolution level 2 H3 hexagons central points (lat+lon) → Filter out using h3-country-coverage.json all hexes above Starlink banned countries, ocean or >710 km away from points where we have predictions

Input data in the best ML Model for each prediction (PacketLoss, LoadedLatency, Jitter, Throughput) → Create 2 JSONs (for both city and hex view) for frontend visualization

MapLibre

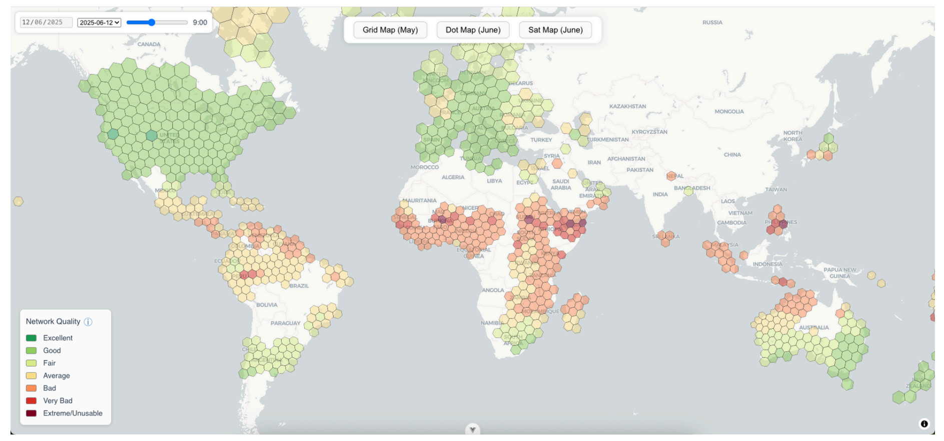

Predictions are visualized on an interactive web frontend using the H3 hexagonal grid/dots system and MapLibre for global mapping. For each day, the system automatically updates performance forecasts at both hexagonal (regional) and city-level resolution, enabling users to explore predicted Starlink quality across space and time. Data and maps refresh daily with the latest weather and satellite data for up-to-date insights.

3. Models evaluation

We evaluate our models using real Starlink measurements, benchmarking against both simple and advanced baselines. Predictions are matched to ground truth using median values within a 3-hour window for each location and hour. Careful anomaly filtering, weather, and satellite enrichment improve accuracy—especially for latency. Ensemble models achieve up to 0.60 explained variance R^2 for latency, outperforming all baselines, but gains for jitter and packet loss are more modest.

For all cities create 24 entries for 20th and 24th of May → Weather Index (WI) Enrichment - Meteostat; Sat. Density Enrichment (CelesTrak + SG4/Skyfield) → Predict for each unique city (lat+lon) per hour for 20th of May and 24th of May

MLab (NDT7) > Starlink Data (BigQuery, SQL, median) - data for ground truth 20th and 24th of May → For each prediction: Find Ground Truth rows within 1.5h (3h window) and with same location and take the median as the ultimate real value

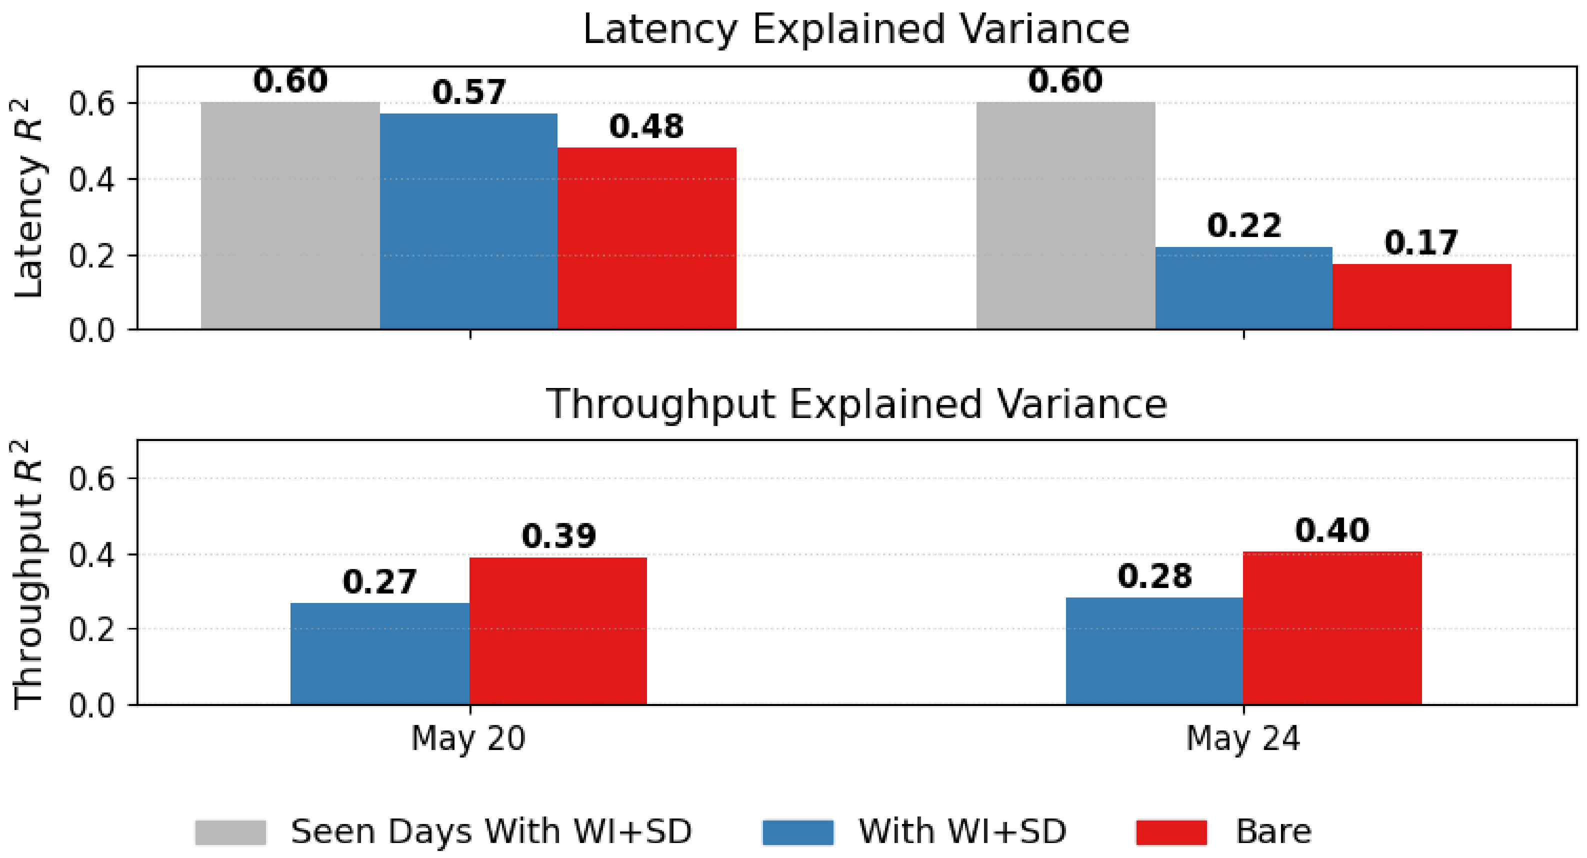

R^2 scores for predictions on two fully unseen days. Weather and satellite features significantly boost latency prediction, but not throughput. All results use robust evaluation with real MLab data and median-matched ground truth.

4. Future Work & Conclusions

- Weather and satellite density features greatly improve latency prediction, but throughput, jitter, and packet loss remain challenging.
- Frequent **retraining** and **aggresive data cleaning** are paramount for reliable forecasts especially for latency and throughput.
- The interactive tool delivers real-time, global Starlink performance predictions at city and grid-cell resolution, validated on unseen days using a strict, median-matching ground truth methodology.

How can we improve the model in the future?

- Add features like space weather and distance to Starlink ground stations (POPs) for better throughput prediction.
- Improve satellite density metric using actual antenna field-of-view.
- Test and adapt the approach for new regions (e.g., India) and longer forecast horizons.
- Quantify uncertainty and support more advanced ML techniques for better generalization.