TelcoTemp: Opportunistic Air Temperature Monitoring from Operational Microwave Link Networks

Matej Istvanek, Petr Musil

Department of Telecommunications, Brno University of Technology, Brno, Czechia

Key idea

TelcoTemp estimates near-surface air temperature from operational commercial microwave links (CMLs). The central idea is to treat telecom equipment as an opportunistic sensor network and use a sequence model to translate link-side measurements into temperature estimates that can be compared with meteorological stations and then mapped spatially.

The poster focuses on the neural-network part of the system: how CML time series are transformed into model-ready sequences, how the LSTM predicts endpoint temperature, and how the predictions behave when evaluated against nearby meteo-station observations.

Why CML data?

- CML endpoints already contain temperature sensors inside operational telecommunication infrastructure.
- Dense urban link networks can complement sparse meteorological stations and improve spatial detail.
- Opportunistic sensing reuses existing infrastructure instead of requiring a new dedicated sensor network.
- The application provides a repeatable pipeline from raw measurements to maps that can be run continuously or for historical time ranges.

Estimation task

Input	Role in the model
CML measurements	Device temperature and received signal level form the main temporal sequence.
Time context	Hour, day, and daylight features describe diurnal heating and cooling.
Static context	Azimuth and altitude describe endpoint installation geometry without exposing endpoint identity.
Technology label	A learned embedding accounts for different microwave-link hardware families.

CML inference

- Measurements are read from InfluxDB and enriched with link metadata from MariaDB/MySQL.
- Runtime features include time, daylight, coordinates, altitude, technology, and device temperature.
- Rolling windows are grouped by CML endpoint, where each endpoint is represented by `Link_ID` + `Side`.
- Each valid sequence is evaluated by a PyTorch LSTM checkpoint and converted back to degrees Celsius using the stored scaler bundle.

LSTM temperature model

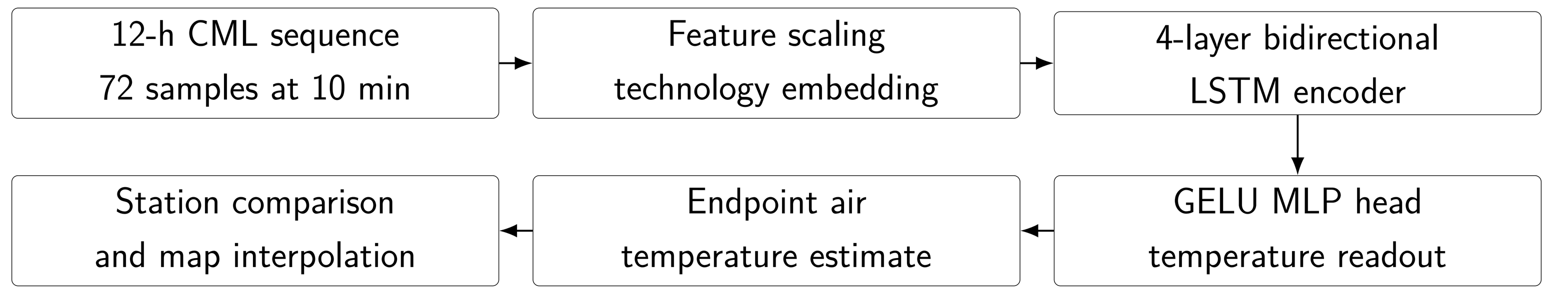

Figure 1. Inference path of the deployed `lstm_v6` model. The model uses temporal dynamics and link context to estimate air temperature at each CML endpoint.

Network configuration

Parameter	`lstm_v6`
Sequence window	72 samples, 10 min sampling, 12 h history
LSTM core	4 layers, hidden size 96, bidirectional
Regularization	Dropout 0.3
Categorical context	Technology embedding, dimension 8
Prediction head	MLP hidden size 96 with GELU activation
Supported technologies	4

Why temporal context matters

L	History	Best LSTM MAE	BL MAE
12	2.0 h	1.411	2.415
20	3.3 h	1.353	2.410
36	6.0 h	1.225	2.395
72	12.0 h	1.184	2.383

- Longer windows consistently reduce error, supporting the thermal-memory interpretation of the unit.
- The best completed window uses 72 samples, i.e. 12 hours of recent 10-min telemetry.
- The linear baseline sees the same flattened history, so the gain comes from nonlinear recurrent sequence modeling.

Benchmark protocol and interpolation

- The principal non-TRSL corpus contains 13,082,447 aligned endpoint-time rows from four technology classes – endpoints are paired with nearby reference meteorological stations to define the supervised target.
- Train, validation, and test sets use leakage-aware temporal blocks; valid windows cannot cross splits.
- Point temperature estimates are interpolated and masked onto a regular grid covering the Czechia
- TelcoTemp uses regression kriging: a regression model captures the elevation-related temperature trend, while kriging models the remaining spatial residuals.
- Duplicate or near-duplicate point locations (usually roofs) are merged before interpolation using the median temperature to improve numerical stability.

Main testing results

Dataset	Test windows	LSTM MAE	BL MAE	Improvement
non-TRSL	2.21 M	1.184 °C	2.383 °C	50.3%
TRSL-enabled	1.62 M	1.170 °C	2.121 °C	44.9%

- On the principal benchmark, the selected BiLSTM-Deep model reaches RMSE 1.740 °C versus 3.069 °C for the baseline.
- Limitation: no "real" groundtruth data, different height and env conditions to meteo stations, various CML technologies may report different behaviour

Application

- TelcoTemp CLI application is currently running on BUT internal server.
- The source code is available at: `https://github.com/telcoSense/telcotemp`.
- Outputs are available at our TelcoSense platform: `https://telcosense.cz/temp`.

Acknowledgements

This output was developed within the project Spatial Air Temperature Monitoring Using Microwave Links Data (SS07020434), financed with the state support of the Technology Agency of the Czech Republic and the Ministry of the Environment of the Czech Republic. The platform is further developed and applied within the project Multidimensional Data Analysis for the Management of River Ecosystems through Multiscale Automatic Tools (TQ14000002), financed with the state support of the Technology Agency of the Czech Republic within the SIGMA Programme.

Example hourly products

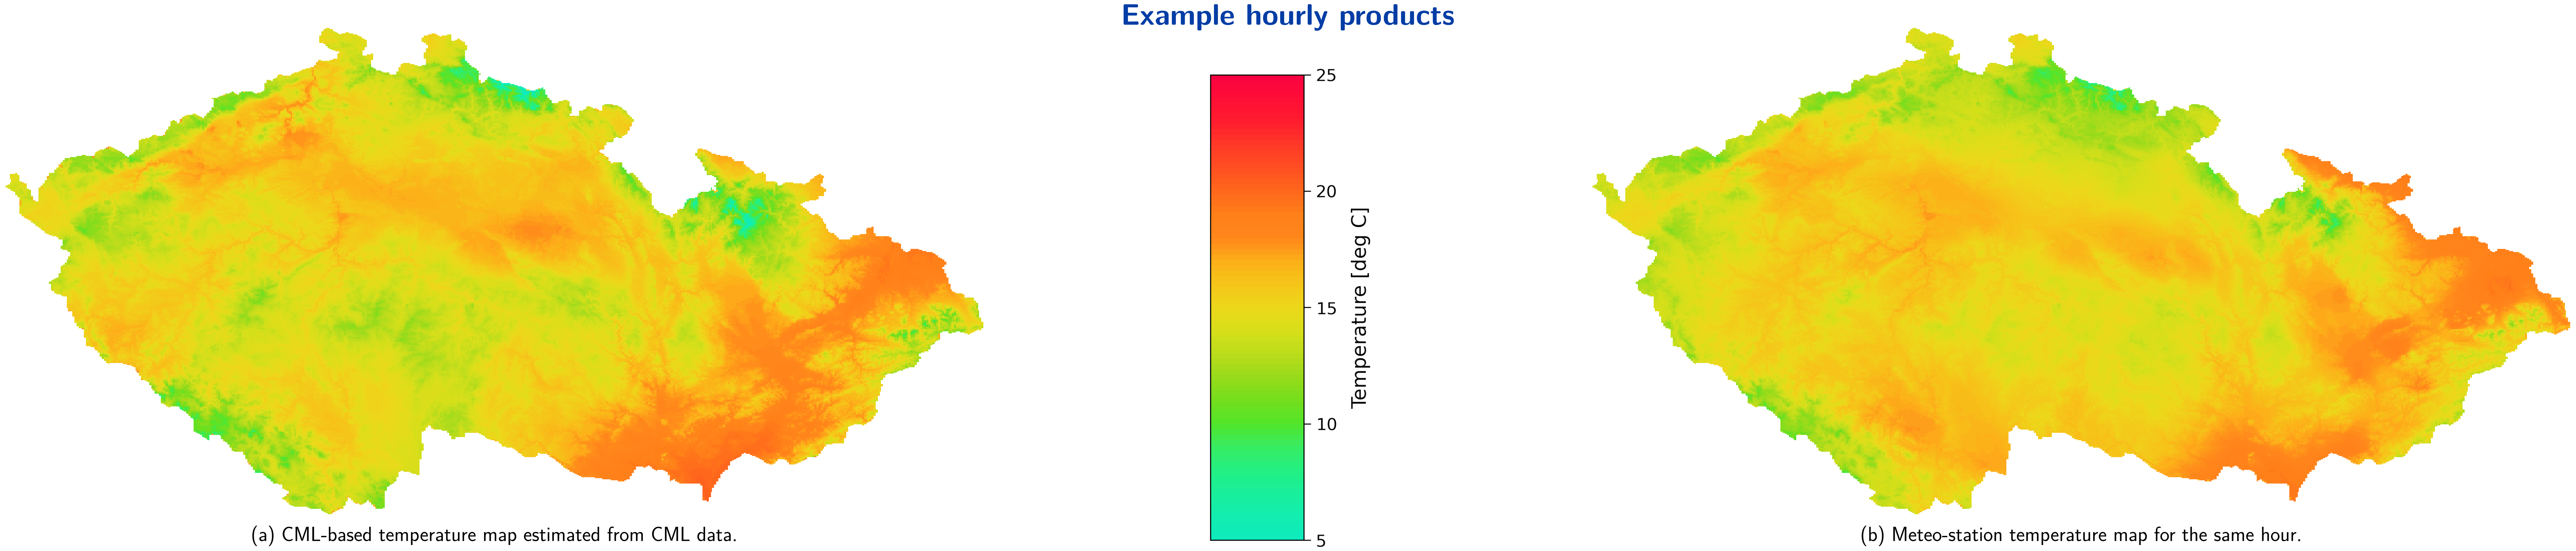

(a) CML-based temperature map estimated from CML data.

(b) Meteo-station temperature map for the same hour.

Figure 2. Representative TelcoTemp outputs for 2025-08-01 06:00. The two products support direct visual comparison between opportunistic CML sensing and station-based interpolation.