https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Feasibility and limits of commercial microwave links for opportunistic air-temperature monitoring

Petr Musil[1], Matěj Ištvánek[1], Štěpán Miklánek[1], Jan Hejna[1], and Petr Mlýnek[1]

[1]Brno University of Technology, Faculty of Electrical Engineering and Communications, Department of Telecommunications, Technická 12, 616 00 Brno, Czech Republic

Correspondence: musilp@vutbr.cz

Abstract. Commercial microwave links (CMLs) are widely used as opportunistic rainfall sensors, but their network management telemetry often also includes built-in temperature readings intended for hardware supervision. This study evaluates whether such telemetry can provide useful information on ambient-temperature variability. We combine a controlled laboratory experiment, infrared imaging, a rooftop field comparison with a nearby reference weather station, and blocked-validation regression experiments across multiple roofs and technologies. The results show that CML temperature readings strongly track the temporal variability of nearby air temperature, but they are affected by large offsets caused by device self-heating, solar radiation, convection and radiative exchange with nearby structures, and post-reset warm-up. For the investigated hardware, monthly mean offsets are approximately 24 to 26 °C, confirming that CML telemetry cannot be interpreted as a direct air-temperature measurement. However, the offsets contain systematic components and the temporal co-variability with reference temperature remains strong. We therefore interpret CML temperatures as opportunistic thermal proxies rather than replacements for standard meteorological observations. In the larger multi-roof dataset, blocked validation shows that much of the apparent improvement over a global constant offset is explained by per-endpoint bias removal: the leave-one-month MAE decreases from 7.46 °C for a global offset to 1.97 °C for a per-endpoint offset. The best supervised temporal model further reduces the error against the nearby weather-station reference to 1.70 °C, whereas the best leave-one-roof result remains 2.19 °C, indicating unresolved site-transfer limitations. The errors quantify agreement with nearby reference stations rather than colocated rooftop air-temperature accuracy. Together, the results suggest that CML temperature telemetry has practical potential as a dense environmental thermal proxy, including future urban applications that exploit the range of installation heights in dense networks, provided that device-specific biases and local installation effects are explicitly accounted for. We discuss these limitations and the requirements for scaling the approach to larger heterogeneous networks.

1 Introduction

Commercial microwave links (CMLs) are an integral part of modern telecommunication infrastructure. They provide point-to-point radio connections for data transport and backhaul, and in many countries they form dense terrestrial networks in urban, suburban, and rural areas. In atmospheric science, CMLs are already established as a source of opportunistic rainfall observations because the microwave signal is attenuated by raindrops and atmospheric humidity (Chwala and Kunstmann,

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

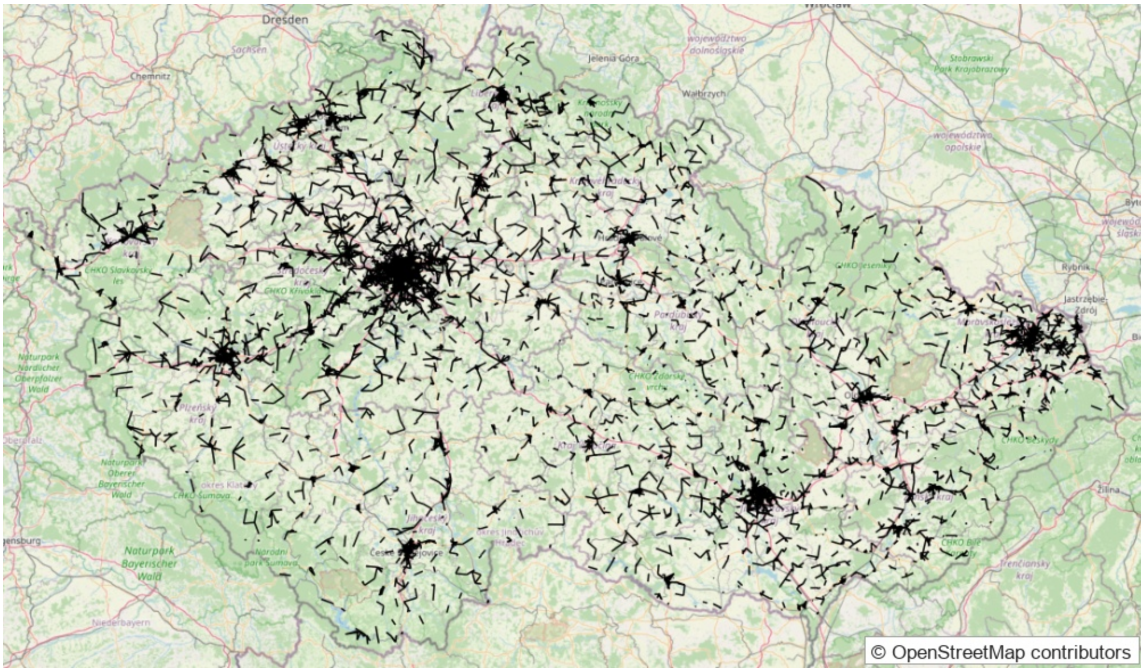

Figure 1. Public E-band point-to-point links (71–76 GHz and 81–86 GHz) in the Czech Republic, generated from the Czech Telecommunication Office database of technical data for fixed point-to-point radio systems (Czech Telecommunication Office, 2026). The background map is © OpenStreetMap contributors, as indicated in the map.

2019; Uijlenhoet et al., 2018; Zheng et al., 2022; Messer et al., 2006). Most of this work uses signal levels or related metadata from commercial backhaul networks (Bianchi et al., 2018; Zinevich et al., 2008; Fencl et al., 2013), while related studies have also explored E-band links and other opportunistic microwave systems for rainfall and water-vapour observations (Fencl et al., 2020; Bubniak et al., 2022; Gianoglio et al., 2023). Figure 1 provides an overview of the spatial distribution of public E-band microwave-link registrations in the Czech Republic.

Most microwave radios also report internal temperature for maintenance and protection of critical components. These temperature values are typically accessible through simple network-management interfaces, such as SNMP, and can therefore be collected continuously without additional field instrumentation. This raises a natural question: can built-in CML temperature telemetry serve as a useful proxy for environmental temperature? The answer is not obvious, because the sensor is embedded inside operating hardware and therefore measures the thermal state of the device rather than free-air temperature.

The broader idea of opportunistic temperature sensing from non-meteorological devices is well established. Previous studies explored vehicle-mounted sensors (Bellavista et al., 2007), optical-fiber infrastructure (Ukil et al., 2012), and especially mobile phones and smartphone batteries (Overeem et al., 2013; Droste et al., 2017; Song et al., 2022; Trivedi et al., 2021). More broadly, opportunistic environmental sensing in dense urban settings has also been demonstrated using heterogeneous infrastructure and crowd-sourced observations (de Vos et al., 2020; Jacoby et al., 2025). In that context, CMLs are not interesting because they are perfect thermometers; rather, they offer a different compromise between density, maintenance, stability, and accessibility. Compared with smartphones or cars, CMLs are stationary, permanently powered, and operated within a relatively controlled technical environment. Compared with many generic IoT devices, their telemetry is already centrally collected and their physical role in atmospheric monitoring is established through rainfall sensing. Community work on opportunistic-sensing data standards has also recognised that CML data acquisition can include additional network-management variables such as outdoor-unit temperature, while noting that the corresponding SNMP object identifiers are not standardised across hardware platforms (Fencl et al., 2024). In that sense, the main motivation for CML temperature-proxy monitoring is not to replace standard weather stations, but to extract an additional atmospheric variable from infrastructure that is already available and already used in environmental applications.

The value of CMLs lies in their combination of (i) high network density, especially in built environments, (ii) fixed position and orientation, which reduces some forms of contextual uncertainty compared with mobile sensors, (iii) operator-managed telemetry streams that can be accessed continuously, and (iv) potential synergy with rainfall monitoring and multi-parameter environmental sensing on the same infrastructure (Nielsen et al., 2024; Angeloni et al., 2024; Blettner et al., 2022; Adirosi et al., 2021). At the same time, CMLs have major drawbacks for temperature monitoring: installation heights differ from standard 2 m meteorological observations, local rooftop conditions can dominate the reading, and offsets can be very large. These limitations are particularly important because CML networks are dense enough to be attractive for urban meteorology. A high-density temperature field can support the study of heat extremes and urban microclimate variability (Bahi et al., 2019; de Vos et al., 2020). Because individual links are mounted at different rooftop and tower heights, a dense CML network does not sample a single 2 m surface but a distribution of heights across the urban canopy. In principle this turns the non-standard installation height from a pure drawback into an opportunity: a sufficiently dense network could support stationary three-dimensional monitoring of the urban thermal field and of heat islands above cities, which fixed 2 m station networks cannot resolve. In addition, if temperature data could be retrieved robustly from operational CMLs, they might complement sparse weather-station observations in regions where classical monitoring is limited, similarly to the way CML rainfall sensing has been investigated in data-sparse environments (Chwala and Kunstmann, 2019; Doumounia et al., 2014). In the Czech Republic alone, the potential scale is large. According to the Czech Telecommunication Office, approximately 24 000 duplex point-to-point links were registered in the unlicensed 60 GHz band as of September 2022 (Czech Telecommunication Office, 2022). In the E-band (71–76 and 81–86 GHz), more than 13 600 point-to-point links have been reported (Ministry of Industry and Trade of the Czech Republic, 2024).

The objective of this paper is therefore to evaluate feasibility, not to claim that CML telemetry provides a standard atmospheric temperature measurement. We ask four questions. First, what physical processes control the difference between internal CML temperature and nearby ambient temperature? Second, how stable is this difference over time? Third, how strongly do unit-specific installation conditions affect the measurement? Fourth, can the available telemetry be used to estimate nearby reference temperature in a useful, though necessarily approximate, way? To answer these questions, we combine laboratory and field observations, compare CML records with a nearby weather station, analyze warm-up and reset behavior, and treat machine learning as an exploratory compensation tool.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

2 Materials and methods

2.1 CML telemetry and hardware background

A CML provides a radio connection between two physical locations within line of sight. Terrestrial microwave units usually operate between about 5 and 40 GHz, with modern deployments extending to even higher bands (Chwala and Kunstmann, 2019). In practice, a radio unit contains several modules that contribute to its thermal behavior: a radio modem, communication interfaces, processing electronics, and power circuitry. Components with higher power consumption, such as amplifiers and processing subsystems, produce heat during operation and are therefore monitored by internal temperature sensors.

In this paper, a CML endpoint denotes one physical radio unit at one end of a point-to-point microwave link. A link consists of two endpoints. The term microwave unit is used when discussing the thermal behavior of the hardware, whereas CML temperature denotes the temperature value reported by the endpoint telemetry. Roof cluster refers to a group of endpoints installed on the same or closely related rooftop site and evaluated together in the blocked spatial-validation experiment.

The exact placement of these sensors depends on vendor and device family. Based on our inspection of several radio platforms, the reported temperature may originate from a sensor located on the printed circuit board (PCB), from an integrated CPU sensor, or from a dedicated sensor near high-power radio circuitry. Consequently, the same ambient conditions can produce different internal temperatures for different devices. Some units show only moderate offsets from ambient air temperature, whereas others exhibit offsets of several tens of degrees Celsius. In the present study, we focus on microwave units whose reported temperature corresponds to a sensor located near the radio modem and therefore reflects internal hardware temperature rather than exposed free-air temperature.

This characteristic already suggests the main challenge of the study. The CML temperature signal cannot be understood as a noisy version of ambient air temperature. Instead, it is a composite thermal response that contains environmental information together with several device-specific offsets. The benefit of CMLs lies not in their sensor placement, but in the availability, continuity, and density of their telemetry.

2.2 Laboratory experiment and thermal imaging

To characterize the role of device self-heating, we performed a controlled indoor experiment with two CML units without antennas, denoted LA and HA. The two units operated as a microwave link pair and their built-in temperature values were read once per minute for approximately 30 min after start-up. Before the experiment, air conditioning was used to stabilize the indoor temperature. The units were linked without a standard field antenna load, but with receive signal levels comparable to a short practical setup. Their electrical power consumption remained stable before and after the link was established (20.5 W and 20.8 W with $\pm 0.5\%$, respectively). In parallel, the units were monitored with an infrared Fluke Ti25 camera. The infrared images are used as qualitative evidence of non-uniform internal and surface heating. Because the camera scale and surface emissivity introduce additional uncertainty, the images are not interpreted as calibrated surface-temperature measurements. The laboratory setup isolates the self-heating contribution, but it does not reproduce outdoor radiative and convective forcing.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Table 1. First-order warm-up fits for the laboratory units and a representative field cold start.

Series	t_0 (°C)	t_∞ (°C)	τ (min)	R^2	G_{th} ($\mathrm{W\,K^{-1}}$)	C_{th} ($\mathrm{J\,K^{-1}}$)
LA	24.8	47.4	6.9 ± 0.2	0.994	0.91	378
HA	24.3	49.5	7.5 ± 0.2	0.996	0.83	373
field CML	24.1	48.0	5.8 ± 0.2	0.996	–	–

To describe start-up transients more physically, we model the microwave unit as a lumped thermal system with effective thermal capacitance C_{th} and effective heat-transfer coefficient G_{th}. The unit temperature t_{MU} then follows

$$C_{\mathrm{th}} \frac{\mathrm{d}t_{\mathrm{MU}}}{\mathrm{d}t} = P_{\mathrm{HW}} - G_{\mathrm{th}} \left(t_{\mathrm{MU}} - t_{\mathrm{A}} \right), \tag{1}$$

where P_{HW} is the internally generated heat and t_{A} is the surrounding air temperature during the experiment. If P_{HW} and t_{A} are approximately constant over the start-up interval, the solution is the standard first-order response

$$t_{\mathrm{MU}}(t) = t_\infty - (t_\infty - t_0) \exp\left(-\frac{t}{\tau} \right), \tag{2}$$

where t_0 is the initial temperature, t_∞ is the asymptotic operating temperature, and

$$\tau = \frac{C_{\mathrm{th}}}{G_{\mathrm{th}}} \tag{3}$$

is the effective thermal time constant. Figure 2 shows that both laboratory units exhibit a similar first-order warm-up behavior after start-up. Rather than treating the observed rise as a purely empirical curve, Eqs. (1)–(3) interpret it as the transient response of a thermally forced system approaching a new equilibrium between internal heat generation and heat loss to the surroundings. Although the laboratory setup does not reproduce outdoor forcing, it isolates the self-heating contribution and provides a physically motivated reference for interpreting field resets and recovery periods. Fitting Eq. (2) to the two laboratory cold starts and to one field cold start gives effective time constants of about 6–8 min (Table 1). The fitted initial temperature t_0 is used as the stabilized room-temperature proxy for estimating G_{th} and C_{th} in the laboratory cases, using the measured electrical power consumption values of $P_{\mathrm{HW}} = 20.5$ and 20.8 W for the LA and HA units, respectively, as proxies for hardware heat dissipation. The fitted time constant τ is therefore the most robust parameter from these transients; the absolute partition into G_{th} and C_{th} is more sensitive to the room-temperature proxy.

Figure 3 confirms that the internal thermal structure of the unit is highly non-uniform. The casing dissipates heat through the heat sink and is therefore cooler than the internal electronics. The hottest areas occur near the modem and the internal thermal path that connects the modem to the housing. The built-in temperature sensor reports a value much closer to these hot spots than to the external air. This observation explains why the temperature signal can still track environmental variability while maintaining a large positive offset.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

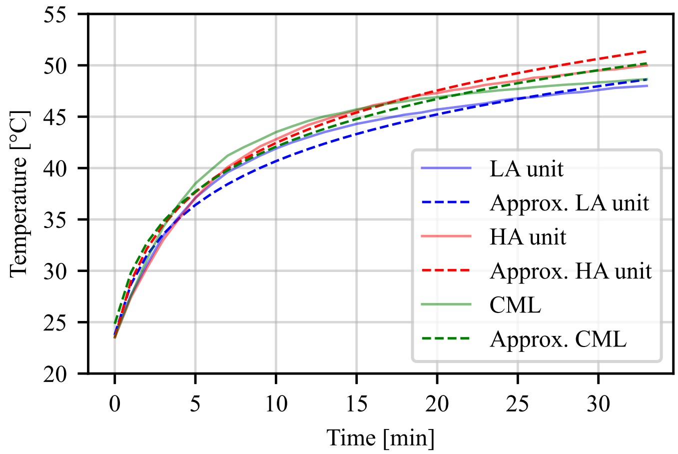

Figure 2. Temperature readings from the laboratory experiment (HA and LA units) and a representative operational field cold start of a microwave device. Dashed curves show first-order fits of Eq. (2).

2.3 Field site and datasets

The field analysis is based on data from a private CML distributor and a nearby reference weather station operated by the Czech Hydrometeorological Institute (CHMI). The main dataset contains five CML units located on a single rooftop approximately 650 m from the weather station. Both the rooftop site and the station are approximately 300 m above sea level. The CML temperature telemetry was recorded at 30 s resolution and aggregated to 10 min intervals to match the temporal resolution of the reference station.

The rooftop environment is important for interpreting the data. CMLs must preserve line of sight and are therefore often installed on roofs or elevated structures, where they are exposed to direct solar radiation, reflections, and local heat sources. Figure 4 shows an example of a real installation with non-uniform shading by nearby structures, while Fig. 5 schematically illustrates how the thermal state of the unit can be influenced by the roof, surrounding walls, and nearby technical equipment.

In addition to the five-unit rooftop dataset, we use a supplementary comparison between opposite ends of the same links. In this comparison, the A units are installed on a roof near heat exchangers, whereas the B units are the opposite units of the same links and are not exposed to these exchangers. This provides an additional perspective on inter-device consistency and local environmental bias. The manuscript therefore combines two complementary data sources: a physically detailed single-site case study used for the thermal interpretation in Sects. 3.1–3.6, and a separate broader 2025 multi-roof dataset used only for the blocked regression analysis in Sect. 3.7. Table 2 summarizes both. The multi-roof dataset is described in full in Sect. 3.7; we introduce it here so that all data sources are defined before the results. To avoid disclosing physical CML locations, Table 3 reports only anonymized roof labels, endpoint counts in the selection metadata, anonymized reference-station labels, and roof-to-reference distances; no coordinates, station names, IP addresses, or real technology names are released. The five selected

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

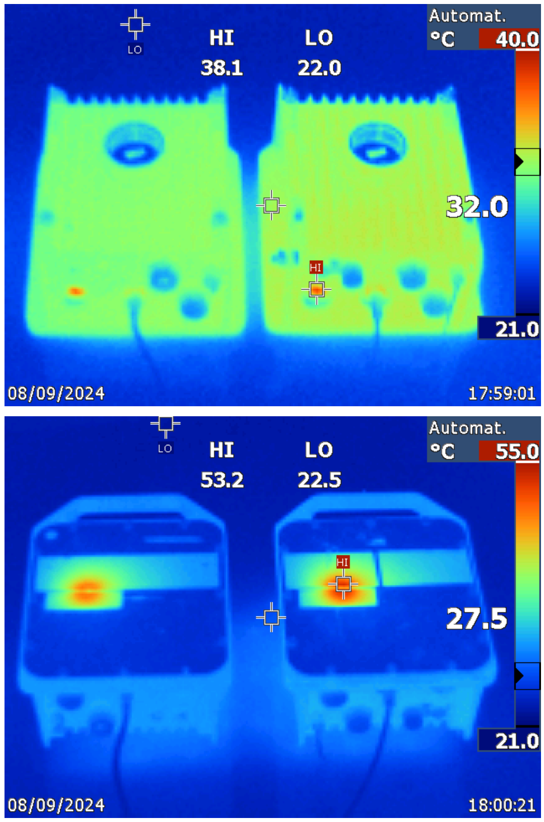

Figure 3. Qualitative infrared images of the experimental microwave units after several hours of operation. The images illustrate strong spatial non-uniformity: the outer casing is cooler than the internal hot spots, and the highest temperatures occur in the region of the modem and internal electronics. The thermograms are not used as calibrated surface-temperature measurements.

roof clusters are associated with four anonymized reference stations, with two clusters sharing the same station, their centroids span approximately 327 km, and their nearest-reference distances range from 156 to 964 m in the available metadata.

2.4 Reduced-order thermal model of microwave-unit temperature

The built-in CML temperature does not represent ambient air temperature directly, because the sensor is embedded in operating hardware and is influenced by both internal heat generation and environmental forcing. Instead of a purely static offset decomposition, we describe the microwave unit as a reduced-order thermal system.

Let $t_{\mathrm{MU}}(t)$ denote the temperature reported by the microwave unit and let $t_{\mathrm{E}}(t)$ denote the local environmental air temperature in the immediate surroundings of the unit. A lumped thermal balance can then be written as

$$C_{\mathrm{th}}\frac{\mathrm{d}t_{\mathrm{MU}}}{\mathrm{d}t} = P_{\mathrm{HW}}(t) + q_{\mathrm{SW}}(t) + q_{\mathrm{LW}}(t) + q_{\mathrm{loc}}(t) - G_{\mathrm{th}}(t)\left[t_{\mathrm{MU}}(t) - t_{\mathrm{E}}(t)\right] + \varepsilon(t), \tag{4}$$

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Figure 4. Example of a real rooftop CML installation. The units experience different solar exposure and shading conditions despite being located at the same site.

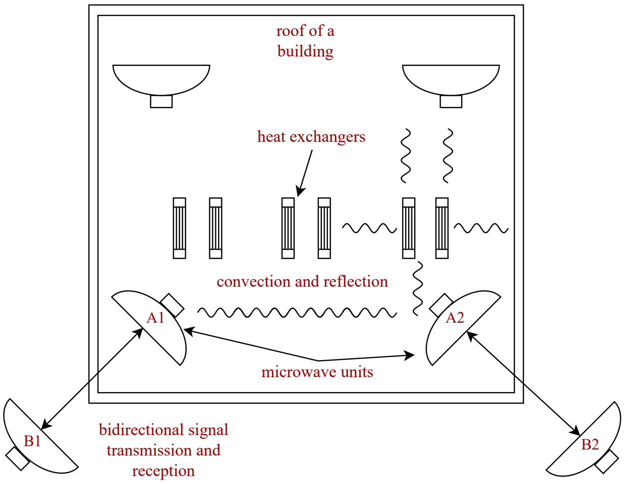

Figure 5. Schematic view of the main local factors affecting the rooftop CML temperature: self-heating of the device, solar loading, reflections, and convective influence from nearby structures and equipment.

where C_{th} is the effective thermal capacitance of the unit, P_{HW} is the heat generated by the hardware, q_{SW} is the shortwave radiative forcing associated with direct and reflected solar radiation, q_{LW} is the external longwave radiative loading from the
160 sky and nearby surfaces, q_{loc} summarizes local site-specific effects such as conductive and convective influence from nearby structures or technical equipment, G_{th} is an *effective* heat-transfer coefficient that lumps convective and linearized longwave exchange with the surrounding air, and $\varepsilon(t)$ is a residual term absorbing unmodeled and higher-order effects. Defining the

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

temperature offset

$$\Delta(t) = t_{\mathrm{MU}}(t) - t_{\mathrm{E}}(t), \tag{5}$$

Eq. (4) can be rewritten as

$$C_{\mathrm{th}} \frac{\mathrm{d}\Delta}{\mathrm{d}t} + G_{\mathrm{th}}(t)\Delta(t) = P_{\mathrm{HW}}(t) + q_{\mathrm{SW}}(t) + q_{\mathrm{LW}}(t) + q_{\mathrm{loc}}(t) - C_{\mathrm{th}} \frac{\mathrm{d}t_{\mathrm{E}}}{\mathrm{d}t} + \varepsilon(t). \tag{6}$$

Equation (6) clarifies that the offset is not constant. It depends not only on internally generated heat and radiative loading, but also on the rate of change of the environmental temperature and on the efficiency of heat exchange with the surroundings. During periods that are approximately quasi-stationary on the thermal time scale of the unit, i.e. when both $\mathrm{d}\Delta/\mathrm{d}t$ and the environmental-temperature tendency are small, the offset reduces to

$$\Delta(t) \approx \frac{P_{\mathrm{HW}}(t) + q_{\mathrm{SW}}(t) + q_{\mathrm{LW}}(t) + q_{\mathrm{loc}}(t)}{G_{\mathrm{th}}(t)}. \tag{7}$$

This quasi-steady form is consistent with the empirical observation that the CML temperature is shifted upward by a large positive offset, while the magnitude of that offset still varies with solar exposure, rooftop environment, and meteorological conditions. In practice, $t_{\mathrm{E}}(t)$ is not observed directly. Instead, the analysis uses the nearby weather-station temperature $t_{\mathrm{WS}}(t)$ as a reference. We therefore write

$$t_{\mathrm{WS}}(t) = t_{\mathrm{E}}(t) + \delta_{\mathrm{site}}(t), \tag{8}$$

where $\delta_{\mathrm{site}}(t)$ represents the mismatch between the local air temperature at the microwave unit and the temperature measured at the reference station. This mismatch includes differences in height, rooftop microclimate, and local exposure. Combining Eqs. (5) and (8) gives

$$t_{\mathrm{MU}}(t) - t_{\mathrm{WS}}(t) = \Delta(t) - \delta_{\mathrm{site}}(t), \tag{9}$$

which shows that the observed difference between the microwave-unit temperature and the weather-station temperature contains both the intrinsic device offset and a site-representativeness term.

Combining the quasi-steady offset (7) with Eqs. (5) and (8) gives an explicit approximate form for the reference temperature,

$$t_{\mathrm{WS}}(t) \approx t_{\mathrm{MU}}(t) - \underbrace{\left[\frac{P_{\mathrm{HW}}}{G_{\mathrm{th}}} - \delta_{\mathrm{site}}\right]}_{\text{quasi-static, site/device specific}} - \underbrace{\frac{q_{\mathrm{SW}}(t) + q_{\mathrm{LW}}(t) + q_{\mathrm{loc}}(t)}{G_{\mathrm{th}}}}_{\text{time-varying radiative loading}}. \tag{10}$$

Equation (10) motivates the regression design of Sect. 3.7: the first bracket is a per-device or per-roof constant, estimated by the per-endpoint offset baseline, whereas the time-varying radiative term is proxied by the daylight indicator, cyclic diurnal and seasonal features, azimuth, and altitude. The supervised models therefore estimate a physically motivated additive correction rather than an arbitrary statistical fit.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Table 2. Overview of the two datasets used in this study. Dataset 1 supports the physical interpretation; Dataset 2 is used only for the blocked-validation regression.

	Dataset 1 (case study)	Dataset 2 (regression)
Period	2021-10 to 2022-09	2025 (full year)
Used in	Sects. 3.1–3.6	Sect. 3.7
Rooftop sites	1	5 clusters
Endpoints	5 (+ A/B link ends)	35
Hardware families	1 (modem-coupled sensor)	4 (A–D)
Reference station	1 CHMI, ∼650 m	4 anonymized CHMI stations, 156–964 m
Resolution	30 s → 10 min	10 min
Matched samples	∼52 500	1 418 747
Role	physical interpretation	blocked-validation skill

Table 3. Privacy-safe summary of the five 2025 roof clusters used to construct the regression dataset. Shared reference-station labels indicate shared stations.

Roof cluster	Candidate endpoints	Reference station	Distance to reference (m)
R1	17	S4	964
R2	15	S2	775
R3	9	S2	156
R4	7	S3	675
R5	5	S1	885

2.5 Evaluation methodology

The field analysis addresses four aspects of feasibility. First, we compare monthly mean temperatures and one-year histograms of CML and weather-station data to quantify the overall offset and distributional similarity. Second, we compute monthly Pearson, Kendall, and Spearman correlations to assess temporal agreement between the two sources. Third, we examine specific situations that reveal physical controls on the offset, including diurnal cycles, different seasons, precipitation events, and device resets. Fourth, we evaluate several blocked regression models as an exploratory way of compensating the offset.

To synchronize the data, CML measurements at 30 s resolution were averaged to 10 min intervals and aligned with the CHMI weather-station records. The resulting time series were then split into monthly subsets. We additionally inspected representative daytime, nighttime, wet, and dry situations to illustrate how solar exposure and precipitation affect the relationship between CML and reference temperature.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Because the paper focuses primarily on physical interpretation, the machine-learning analysis is intentionally secondary. In the context of Eq. (6), the predictors used in Sect. 3.7 can be interpreted as proxies for the dominant forcing terms: the microwave-unit temperature captures the current thermal state, the daylight indicator together with the cyclic seasonal and diurnal terms proxies radiative forcing, azimuth and altitude modulate solar exposure and installation height, and the received signal level may indirectly reflect wet conditions and associated cooling. Device, roof, and technology identifiers additionally account for persistent site- and hardware-specific offsets. We therefore compare training-only constant-offset and clock-only baselines, linear support-vector regression, histogram gradient-boosted trees, and XGBoost under blocked temporal and roof-based validation.

For N matched samples with estimates $\hat{t}_i$, reference values t_i and reference mean $\bar{t}$, we report

$$\mathrm{MAE} = \frac{1}{N}\sum_i |\hat{t}_i - t_i|, \qquad \mathrm{bias} = \frac{1}{N}\sum_i (\hat{t}_i - t_i), \tag{11}$$

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_i (\hat{t}_i - t_i)^2}, \qquad \mathrm{R}^2 = 1 - \frac{\sum_i (\hat{t}_i - t_i)^2}{\sum_i (t_i - \bar{t})^2}. \tag{12}$$

3 Results

3.1 Warm-up and self-heating in laboratory and field data

The controlled experiment demonstrates that the internal temperature is strongly affected by self-heating. Figure 2 shows that the two laboratory units exhibit almost identical warm-up dynamics after power-on, despite a small hardware revision difference. The field cold-start example follows the same overall pattern, which supports the interpretation that device recovery after restart is governed mainly by internal thermal stabilization rather than by ambient air temperature alone.

The infrared images in Fig. 3 further explain the origin of this behavior. The built-in temperature sensor is thermally coupled to the electronics and therefore reflects the internal state of the device. The outer housing and the external environment affect the measurement indirectly through heat exchange rather than direct exposure of the sensor to free air.

3.2 Environmental forcing and daily behavior

Figure 6 compares the temperature of one microwave unit with the nearby weather station during two spring days. The two curves follow the same large-scale temperature evolution, including cooling and warming periods, but the CML values remain shifted and show short deviations. This illustrates the central result of the paper: ambient variability is visible in the CML temperature record, but it is superimposed on a substantial offset and local perturbations.

The influence of solar exposure is even clearer when comparing two units mounted on the same roof but facing different directions. Figure 7 shows that the daily temperature maxima occur at different times for two rooftop units with different azimuths. One unit reaches its maximum in the morning, the other in the afternoon, whereas the weather-station temperature follows a smoother diurnal course. This confirms that orientation and shading modulate the shortwave radiative forcing term and can alter the apparent relationship between CML and ambient temperature even at the same site.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

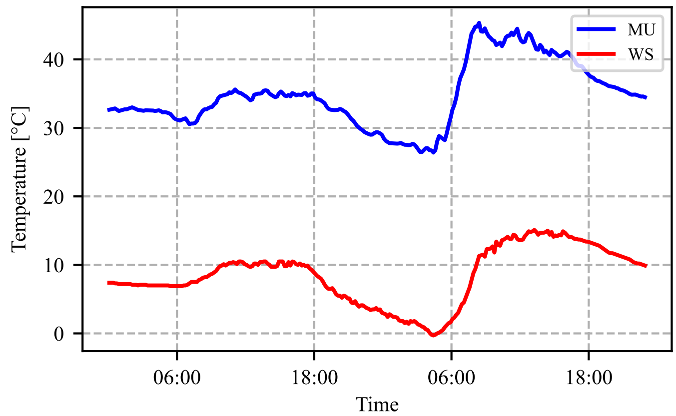

Figure 6. Comparison of microwave-unit temperature and weather-station temperature during two spring days. The overall evolution is similar, but a pronounced positive offset remains.

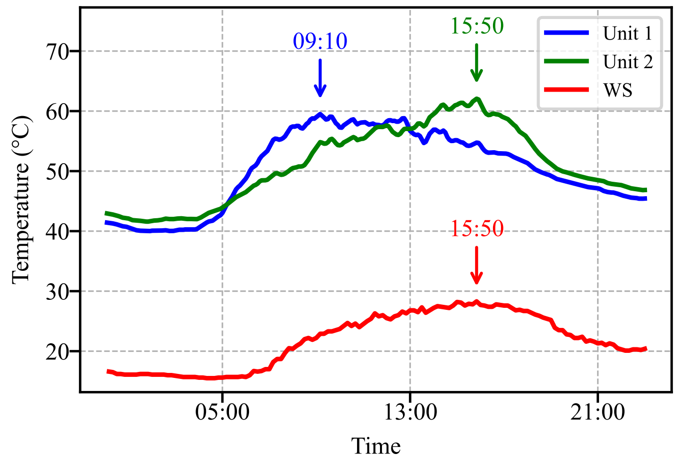

Figure 7. Daily temperatures of two CML units with different azimuths and of the nearby weather station. Different unit orientations produce different times of daily maximum temperature, illustrating orientation-dependent solar forcing at the same rooftop site.

Seasonal comparison leads to the same conclusion from another perspective. Figure 8 compares a CML temperature record with the weather-station temperature during summer and winter days. During winter, solar forcing is weaker and the environment cools the hardware more effectively, so the two curves are more similar. During summer, stronger radiative forcing increases the daytime offset and makes the CML temperature more nonlinear with respect to the reference.

Figure 9 shows another physically meaningful situation: a rainfall event. Both the weather station and the CML temperature record respond to the event, and the microwave unit cools during and after rain. This is consistent with the environmental temperature drop before precipitation and with the enhanced cooling of the wet device surface, in line with the broader literature

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

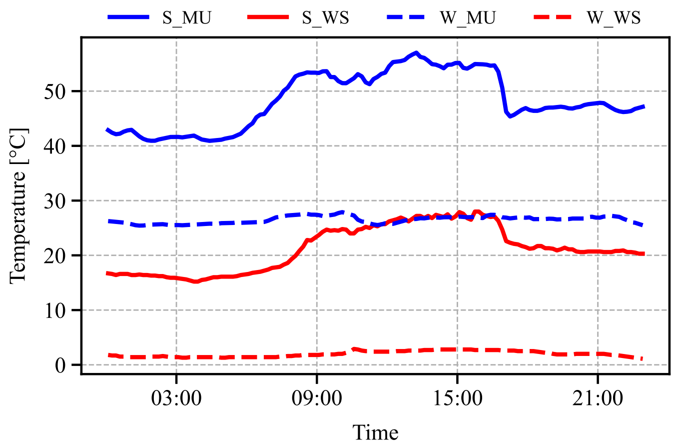

Figure 8. Comparison of a CML temperature record and the nearby weather-station temperature during summer and winter days. The example illustrates the seasonal dependence of the CML–reference relationship, with stronger daytime departures during summer.

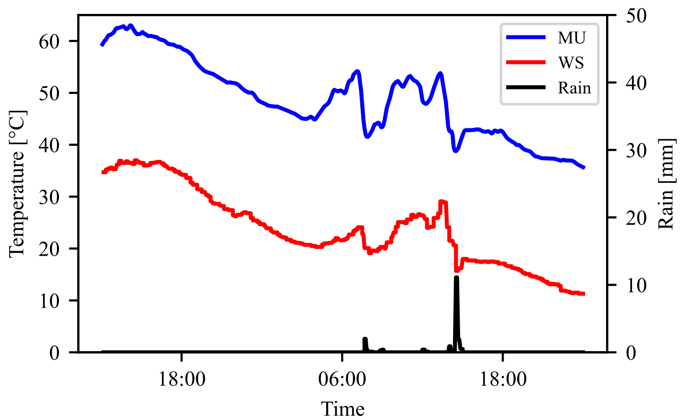

Figure 9. Temperature of a microwave unit and the nearby weather station during a rainfall event. The CML record reflects event-related cooling in addition to its baseline offset, consistent with rainfall-associated cooling and enhanced heat loss from the wet device surface.

on rainfall-related cooling effects (Liu et al., 2022). The persistence of cooling after rainfall also suggests that evaporation and the physical presence of water on the unit contribute to the short-term thermal response.

3.3 Height difference and local representativeness

240 One important limitation of the comparison is the difference in measurement height and local exposure between the CML and the weather station. Weather stations typically measure air temperature at 2 m above the ground, whereas rooftop or tower-mounted CMLs are located substantially higher and within a different micro-environment. Figure 10 illustrates this mismatch conceptually.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

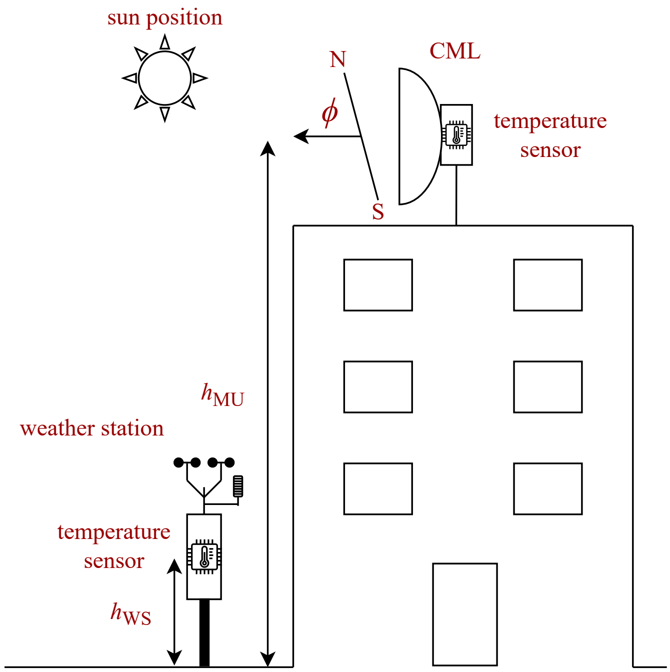

Figure 10. Conceptual difference between the standard weather-station measurement height and the installation height of a rooftop microwave unit.

This does not invalidate the comparison, but it changes its interpretation. The weather station is not a perfect representation
245 of the immediate thermal environment of the microwave housing. Instead, it provides a nearby reference for the local air-temperature field. The agreement between the two signals therefore indicates that the CML carries a regional atmospheric signal, even though the absolute difference contains local rooftop effects.

3.4 Inter-device consistency and site-specific effects

Inter-device consistency is a key question for large-scale applicability. Figure 11 compares year-long daily means of the A
250 and B units, which are opposite ends of the same links but experience different local conditions. The A units are installed near heat exchangers, while the B units are not. Both groups follow the same annual temperature evolution, but the A units remain systematically warmer. This behavior is particularly visible in colder periods and indicates a persistent site-specific contribution to the offset.

This result is important for two reasons. First, it confirms that a large part of the offset is systematic rather than random.
255 Second, it shows that transferability across sites cannot be taken for granted even for the same hardware family. Local rooftop conditions can produce stable but different biases that need to be taken into account in any scalable temperature-estimation framework.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

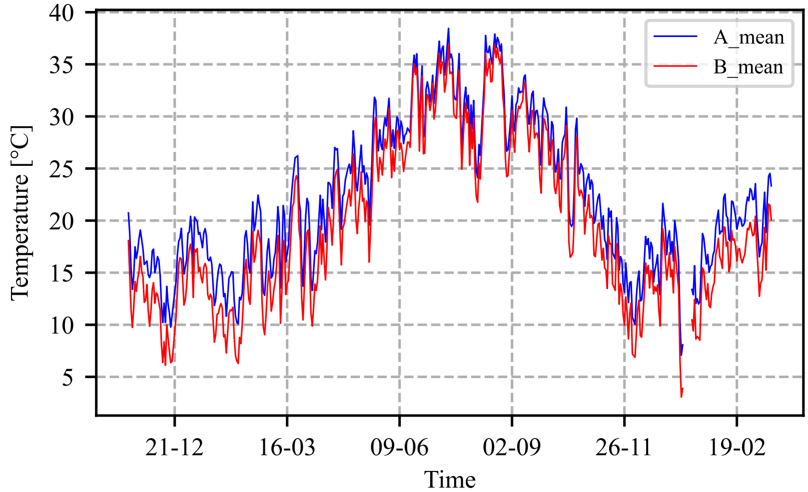

Figure 11. Year-long daily mean temperatures of opposite ends of the same links. Units A were located near heat exchangers, whereas units B were not.

Table 4. Median, mean, and standard deviation (SD) of recovery duration τ_r after reset events. Values are in minutes.

Date	Median τ_r	Mean τ_r	SD τ_r
2023-06	20.50	20.65	5.59
2023-12	19.25	19.29	4.07

3.5 Reset behavior and recovery duration

Operational telemetry also contains irregular resets due to maintenance, outages, or technical problems. After restart, the unit must warm up before its temperature returns to a stable operating regime. Figure 12 shows examples of recovery periods after both short outages and cold starts in June and December. Zero minutes denotes the moment when telemetry resumed.

Table 4 summarizes 40 reset events recorded in June and December 2023. The median and mean recovery times are close to 20 min in both months. The similarity between summer and winter suggests that recovery depends mainly on the pre-reset hardware state and outage duration rather than on season alone. This supports a practical recommendation: CML temperature data should be excluded for at least 20 min after restart.

3.6 Statistical comparison with the nearby weather station

The one-year histograms in Fig. 13 show that the overall distribution of CML temperatures resembles the distribution of the nearby weather-station temperatures, but is shifted toward much higher values and has a somewhat larger spread. This already indicates that a simple constant correction cannot fully describe the relationship between the two time series.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

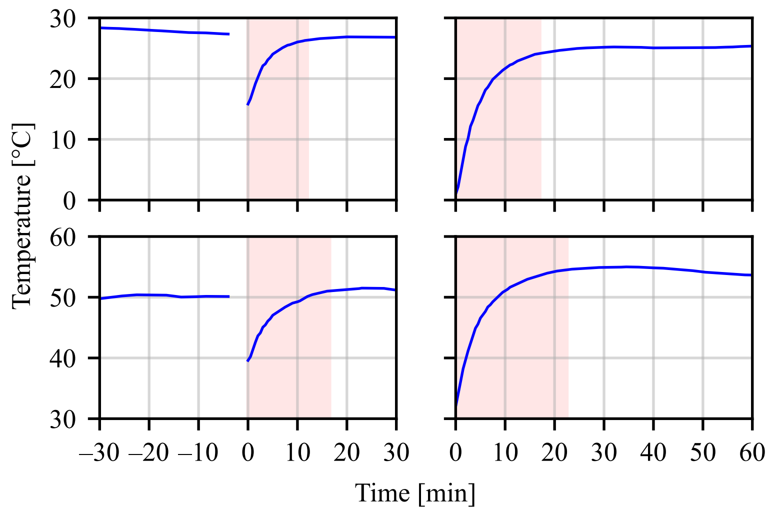

Figure 12. Examples of recovery duration τ_r after short outages and cold starts captured in summer and winter. Shaded intervals indicate the estimated recovery periods used to derive τ_r.

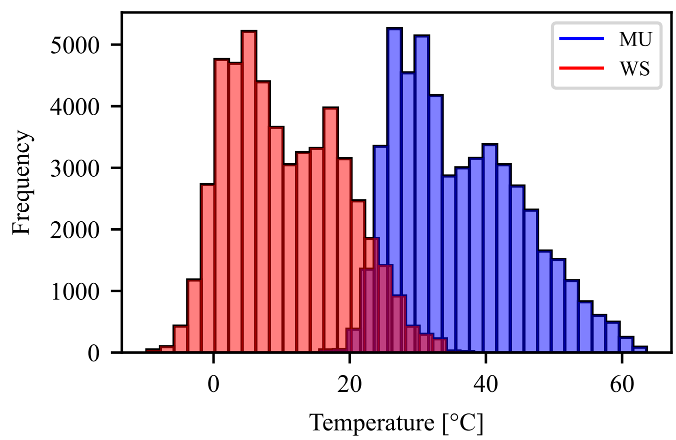

Figure 13. One-year histograms of temperature measured by a weather station and by a microwave unit. The shapes are similar, but the CML temperature is shifted to significantly higher values.

Table 5 quantifies the monthly mean values for one representative microwave unit and the nearby weather station. We report one unit here for readability, while inter-device behaviour is addressed separately in Sect. 3.4 and the broader blocked regression analysis uses a larger multi-endpoint dataset. The monthly offset remains between 24.34 and 26.41 °C, with an average of approximately 25.6 °C. This offset is very large in absolute terms, which confirms that the device measures its internal thermal state rather than ambient air temperature. However, the month-to-month variation of the mean offset is relatively small. The bias therefore contains a stable component in addition to short-term fluctuations.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Table 5. Monthly mean temperatures of a representative CML unit and the nearby weather station, including standard deviation (SD) and mean offset t_{Δ}.

Date	t_{MU}	SD t_{MU}	t_{WS}	SD t_{WS}	t_{Δ}
2021-10	34.90	4.91	9.15	4.98	25.75
2021-11	29.99	2.82	4.75	2.69	25.24
2021-12	27.14	3.55	2.08	4.11	25.06
2022-01	26.85	3.09	2.33	3.33	24.52
2022-02	28.97	2.86	4.63	3.16	24.34
2022-03	30.77	5.86	4.86	5.75	25.91
2022-04	33.52	5.74	7.87	4.92	25.65
2022-05	42.41	5.55	16.06	4.95	26.34
2022-06	46.80	6.11	20.48	5.05	26.33
2022-07	46.29	6.36	20.01	5.18	26.28
2022-08	47.11	6.22	20.71	4.90	26.41
2022-09	39.49	5.60	13.59	4.61	25.90

Using the laboratory estimate of G_{th} from Table 1 and $P_{\mathrm{HW}} \approx 20.5$–$20.8\,\mathrm{W}$, the self-heating floor predicted by Eq. (7) is approximately 23–25 °C. This value should be treated as an order-of-magnitude estimate for the characterized laboratory hardware, because the absolute G_{th} estimate depends on the room-temperature proxy discussed above. Taking the mean laboratory estimate gives about 24 °C, or roughly 90–95 % of the observed mean offset of 25.6 °C for this representative unit. The remaining few degrees, and most of the diurnal and seasonal variation of the offset, are therefore attributable to radiative loading, local site effects, and representativeness differences between the rooftop unit and the reference station.

Table 6 shows monthly correlations between CML and reference temperature. The Pearson coefficient ranges from 0.78 to 0.92, which demonstrates strong temporal agreement despite the large absolute offset. These values support the central interpretation of the paper: the CML record contains a robust atmospheric signal, but that signal is superimposed on a large and installation-dependent thermal bias.

3.7 Temperature estimation

Whereas the preceding subsections focus on physical interpretation using the detailed single-site case study, the temperature-estimation analysis uses a broader operational dataset designed specifically for blocked regression experiments. The regression dataset is not intended to extend the single-site physical case study directly; rather, it provides an independent operational test of whether the telemetry–reference relationship remains useful under blocked temporal and spatial validation.

For the temperature-estimation analysis, we constructed a supervised regression dataset from operational telemetry for five roof clusters and one complete calendar year, 2025. Each row represents a matched 10-min record from one CML endpoint and the nearby weather-station reference. The monthly export used the device uptime telemetry to exclude samples during the first 20 min after a detected restart before the regression matrix was written; the uptime values themselves were then dropped from the exported feature matrix. After filtering and matching, the dataset contained 1 418 747 samples from 35 CML endpoints.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Table 6. Monthly correlation coefficients between CML temperature and the nearby weather-station temperature.

Date	Pearson	Kendall	Spearman	No. of values
2021-10	0.84	0.70	0.86	4463
2021-11	0.78	0.61	0.77	4320
2021-12	0.91	0.74	0.89	4464
2022-01	0.88	0.70	0.85	4464
2022-02	0.81	0.70	0.84	4032
2022-03	0.90	0.76	0.91	4457
2022-04	0.91	0.77	0.92	4320
2022-05	0.88	0.72	0.89	4464
2022-06	0.89	0.74	0.91	4318
2022-07	0.92	0.78	0.93	4464
2022-08	0.90	0.74	0.90	4464
2022-09	0.88	0.74	0.89	4320

The endpoints belonged to four anonymized technology groups: Technology A (428 825 samples), Technology B (369 312), Technology C (320 324), and Technology D (300 286). The target variable is the nearby weather-station temperature t_{WS}; it is therefore a remote reference, not a colocated measurement of the true rooftop air temperature. In the selection metadata, the five roof clusters are 156–964 m from their nearest reference station and use four distinct anonymized reference stations, with two clusters sharing one station (Table 3); station identifiers are not retained in the releasable feature matrix. Although additional data were available, we selected sites that provided sufficient variability in communication technologies, proximity to the weather station, and a sufficient number of endpoints.

The predictor matrix contained the microwave-unit temperature t_{MU}, a binary daylight indicator, unit azimuth and altitude, anonymized technology identifier, CML endpoint identifier, and roof identifier. Seasonal and diurnal cycles were encoded as

$$\left(\sin\frac{2\pi D}{366}, \cos\frac{2\pi D}{366}, \sin\frac{2\pi H}{24}, \cos\frac{2\pi H}{24}\right), \tag{13}$$

for day-of-year D and hour-of-day H. Transmitted- and received-power variables were included when available. Numerical predictors were median-imputed and standardized using parameters estimated from the training part of each fold only. Categorical predictors were imputed from the training fold and one-hot encoded with unseen held-out categories ignored during prediction. This fold-wise preprocessing is important because both roof identity and device identity can otherwise leak information from the test block into the fitted model; in the leave-one-roof protocol, the held-out roof is therefore never available as a usable categorical level at prediction time.

We used three blocked validation protocols. In the last-month holdout, all samples from January–November 2025 were used for training and all samples from December 2025 were used for testing. In leave-one-month-out validation, each calendar month was withheld once, yielding 12 temporal folds. In leave-one-roof-out validation, each roof cluster was withheld once, yielding five spatial-transfer folds. We did not use a randomly shuffled train–test split, because adjacent 10-min samples are strongly autocorrelated and would not provide an independent test of generalization. The temporal protocols should therefore

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

be interpreted as extrapolation in time for previously observed endpoints, whereas leave-one-roof-out is the stricter test of transfer to unseen roof clusters. Because hardware families may differ in thermal offset and transferability, the main experiment pooled all four technologies and used technology as a categorical predictor. A second leave-one-month experiment was fitted only on the largest technology group, Technology A, to quantify how much of the prediction error is associated with pooling heterogeneous hardware.

The model comparison includes two baselines and three simple supervised regressors. The first baseline is a constant-offset correction,

$$\hat{t}_{\mathrm{WS}} = t_{\mathrm{MU}} - \mathrm{median}_{\mathrm{train}}(t_{\mathrm{MU}} - t_{\mathrm{WS}}), \tag{14}$$

which tests whether the relationship can be explained by subtracting a stable hardware offset estimated from the training fold only. In the pooled experiment this correction uses a single fold-specific scalar estimated from all training samples rather than a separate offset for each endpoint; it is therefore a deliberately strict baseline for heterogeneous deployment, not an optimized per-device recalibration scheme. The second baseline is a clock-only ridge regression using only seasonal and diurnal variables; it tests how much skill can be obtained from the annual and daily temperature cycle without using CML temperature itself. The supervised models are a linear support-vector regressor, histogram gradient-boosted trees (HistGBT), and an XGBoost regressor. This selection provides one regularized linear model and two nonlinear boosted-tree alternatives without introducing neural-network complexity. The support-vector and histogram gradient-boosting models follow the standard scikit-learn implementation (Pedregosa et al., 2011); the support-vector framework is described by Chang and Lin (2011) and Cortes and Vapnik (1995), and XGBoost follows Chen and Guestrin (2016). Where model tuning was used, hyperparameters were selected only within the training portion of each blocked fold using month-grouped inner cross-validation. To keep the experiment reproducible at paper scale, the fitted training set for each supervised model was capped by stratified sampling over month–roof groups, whereas all metrics were computed on the complete held-out blocks.

The pipeline also evaluates a per-endpoint offset baseline for the two temporal protocols,

$$\hat{t}_{\mathrm{WS}} = t_{\mathrm{MU}} - \mathrm{median}_{i\in\mathrm{train}(e)}\left(t_{\mathrm{MU},i} - t_{\mathrm{WS},i}\right), \tag{15}$$

where the median offset is estimated separately for each endpoint e using only the training fold. This baseline is intentionally not defined for leave-one-roof validation, because all endpoints on the withheld roof are unseen during training. In the configured run, the reported supervised models were capped at 50000 stratified training samples per fold; all metrics were still evaluated on the full held-out blocks.

The blocked-validation results for the pooled multi-technology dataset are reported in Table 7. The table separates the temporal holdout protocols from the leave-one-roof protocol, because the latter tests transfer to an unseen roof cluster rather than interpolation in time within already observed roof clusters.

The role of hardware heterogeneity is summarized in Table 8, which compares the same leave-one-month protocol for the pooled dataset and for Technology A only. This comparison is not intended to select a preferred deployment subset; it is used to separate temporal predictability from the additional variability introduced by mixing hardware groups.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Table 7. Blocked-validation performance for estimating the weather-station reference temperature from CML telemetry using all four technology groups. MAE, RMSE, and bias are in °C. Baseline parameters were estimated from the training fold only; the per-endpoint offset is defined only for temporal protocols with previously observed endpoints. The reported supervised models used up to 50000 stratified training samples per fold, and all metrics were evaluated on the complete held-out blocks.

Protocol	Model	Folds	MAE	RMSE	Bias	R^2
last-month	Constant offset	1	7.28	8.26	0.10	-3.85
last-month	Per-endpoint offset	1	1.59	2.07	0.05	0.70
last-month	Clock-only ridge	1	3.16	3.97	-0.87	-0.12
last-month	Linear SVR	1	1.57	1.99	-0.31	0.72
last-month	HistGBT	1	1.91	2.65	-0.45	0.50
last-month	XGBoost	1	1.89	2.67	-0.38	0.49
leave-one-month	Constant offset	12	7.46	8.58	1.33	-2.86
leave-one-month	Per-endpoint offset	12	1.97	2.60	0.35	0.64
leave-one-month	Clock-only ridge	12	3.19	4.01	0.01	0.17
leave-one-month	Linear SVR	12	1.70	2.19	0.11	0.74
leave-one-month	HistGBT	12	2.06	2.63	0.21	0.62
leave-one-month	XGBoost	12	2.10	2.65	0.22	0.63
leave-one-roof	Constant offset	5	9.44	10.25	1.80	-0.58
leave-one-roof	Clock-only ridge	5	3.08	3.90	0.00	0.79
leave-one-roof	Linear SVR	5	4.49	5.19	-1.17	0.62
leave-one-roof	HistGBT	5	2.19	2.68	-0.02	0.89
leave-one-roof	XGBoost	5	2.41	2.92	0.24	0.86

Table 8. Leave-one-month performance for the pooled multi-technology dataset and for the largest single technology group (Technology A). MAE and RMSE are in °C.

Scope	Model	Folds	MAE	RMSE	R^2
all	Constant offset	12	7.46	8.58	-2.86
all	Per-endpoint offset	12	1.97	2.60	0.64
all	Clock-only ridge	12	3.19	4.01	0.17
all	Linear SVR	12	1.70	2.19	0.74
all	HistGBT	12	2.06	2.63	0.62
all	XGBoost	12	2.10	2.65	0.63
Technology A	Constant offset	12	4.05	5.47	-0.97
Technology A	Per-endpoint offset	12	1.82	2.42	0.67
Technology A	Clock-only ridge	12	3.08	3.88	0.18
Technology A	Linear SVR	12	1.53	1.97	0.77
Technology A	HistGBT	12	1.67	2.17	0.71
Technology A	XGBoost	12	1.65	2.14	0.72

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Table 9. Distribution of endpoint-specific median offsets $t_{\mathrm{MU}} - t_{\mathrm{WS}}$ in the 2025 regression dataset, grouped by anonymized technology. The range reports the 10th–90th percentile across endpoints within each technology group.

Technology	Endpoints	Median offset (°C)	10–90 % range (°C)
Technology A	13	8.2	4.2–13.2
Technology B	8	26.4	25.2–27.6
Technology C	8	19.6	17.4–23.9
Technology D	6	11.8	10.4–12.9

The physical parameters in Table 1 were obtained from the laboratory units and the Dataset 1 field example, not from all four anonymized hardware groups in Dataset 2. Therefore, Eq. (10) should be read as a physically motivated regression structure rather than as a fully identified thermal model for every Dataset 2 technology. The reduced-order model nonetheless explains the mechanism behind the inter-technology spread: because the quasi-steady floor scales as $P_{\mathrm{HW}}/G_{\mathrm{th}}$ (Eq. 7), hardware families with different power consumption and heat-sinking are expected to have systematically different offsets. This is consistent with the technology-dependent medians in Table 9: the model predicts the existence and ordering of these offsets even though it does not fix their absolute magnitude for each unseen device. As a direct bridge between the two parts of the study, Table 9 summarizes the distribution of endpoint-specific median offsets in the 2025 regression matrix. The large differences between anonymized technology groups show that the persistent hardware term is not a minor nuisance: it is one of the dominant quantities that the per-endpoint baseline and the technology predictor absorb. The within-group spread is equally informative for deployment. For some technologies the endpoint offsets are tightly clustered (Technology B spans only 25.2–27.6 °C across eight endpoints), whereas for others they are broad (Technology A spans 4.2–13.2 °C). Where the within-technology spread is small, a technology- or hardware-family-level offset prior obtained from metadata alone could recover much of the per-device correction without a colocated reference at every link, whereas a broad spread implies that device-specific calibration is unavoidable. The gap between such a metadata-level prior and the per-endpoint offset therefore measures the calibration effort required to scale the approach to uncalibrated devices. We do not fit a dedicated per-technology baseline here, but this reading identifies a realistic intermediate option between the global offset, which is clearly insufficient (Table 7), and the per-endpoint offset, which presupposes local ground truth at each device.

To expose the seasonal stability of the temporal validation, Table 10 reports month-wise MAE for the leave-one-month protocol. The table makes clear that the blocked evaluation does not depend on a single favourable month: the per-endpoint offset remains between 1.59 and 2.26 °C, while the linear SVR remains between 1.39 and 2.16 °C. The clock-only baseline remains near 2.6–4.3 °C, and the global constant-offset baseline remains above 7 °C throughout.

Table 7 shows that the global constant-offset correction is not sufficient for the pooled multi-technology dataset. Its leave-one-month MAE is 7.46 °C, which confirms that a single stable offset explains only part of the relationship between internal unit temperature and the weather-station reference. The per-endpoint offset reduces this error to 1.97 °C, showing that most of the apparent gain over the global offset comes from removing persistent device-specific bias. The linear SVR gives the

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Table 10. Month-wise MAE (°C) for the leave-one-month validation protocol on the pooled multi-technology dataset.

Held-out month	Offset	Per-endpoint	Clock-only	Linear SVR	HistGBT	XGBoost
2025-01	7.08	2.11	3.02	1.79	2.17	1.95
2025-02	7.73	2.20	3.17	2.16	2.55	2.57
2025-03	7.69	2.26	3.27	1.81	2.04	2.08
2025-04	7.55	2.16	3.70	1.87	2.13	2.39
2025-05	7.53	2.02	4.34	1.69	2.37	2.45
2025-06	7.58	2.03	3.02	1.71	1.69	1.69
2025-07	7.61	1.88	2.93	1.39	1.58	1.64
2025-08	7.61	2.13	3.44	1.71	2.03	2.28
2025-09	7.25	1.84	2.61	1.43	2.12	1.90
2025-10	7.25	1.77	2.61	1.64	1.93	1.93
2025-11	7.34	1.68	3.01	1.65	2.16	2.38
2025-12	7.28	1.59	3.16	1.57	1.91	1.89

best temporal generalization, with a leave-one-month MAE of 1.70 °C and a last-month MAE of 1.57 °C. Its margin over the per-endpoint baseline is therefore modest, about 0.27 °C for leave-one-month validation, but it indicates that the telemetry and auxiliary predictors contain information beyond a static per-device correction, especially for within-month variability. HistGBT and XGBoost do not improve on the linear model under blocked temporal validation, with leave-one-month MAE values of 2.06 and 2.10 °C, respectively. This does not imply that the thermal response is physically linear; rather, after clock, device, roof, and technology predictors are included, the extra flexibility of these nonlinear models does not transfer better to unseen time blocks in this dataset. The leave-one-roof protocol has a different interpretation. The per-endpoint offset is undefined there because the held-out endpoints are not observed in training. The clock-only baseline is already strong in this setting (3.08 °C), because it avoids site-specific CML telemetry and relies only on the seasonal and diurnal cycle. HistGBT and XGBoost nevertheless reduce the spatial-transfer MAE to 2.19 and 2.41 °C, respectively, compared with 4.49 °C for the linear SVR and 9.44 °C for the global offset baseline. The improvement over clock-only is therefore meaningful but modest, and it should be read as evidence that nonlinear tabular models can add CML-specific information in roof transfer, not as proof that spatial transfer is solved.

The R^2 values in Table 7 should not be compared directly across validation protocols. Because $\mathrm{R}^2 = 1 - \mathrm{MSE}/\mathrm{Var}(t_{\mathrm{WS}})$ is evaluated on the held-out block, it depends on the variance of the target within that block. A single withheld month covers only part of the annual cycle and is dominated by day-to-day synoptic variability (target standard deviation of roughly 4–5 °C), whereas a withheld roof retains the full year and therefore a much larger seasonal variance (roughly 8–9 °C). The clock-only baseline illustrates this directly: its RMSE is almost identical in the two protocols (4.01 versus 3.90 °C), yet its R^2 rises from 0.17 to 0.79 purely because the denominator changes. We therefore base all cross-protocol comparisons on MAE and RMSE, which are expressed on the physical temperature scale, and interpret R^2 only within a single protocol.

The Technology A-only experiment in Table 8 has lower error for the global offset baseline, the per-endpoint offset baseline, and all supervised models, while the clock-only baseline changes only slightly because it does not use CML telemetry. This

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

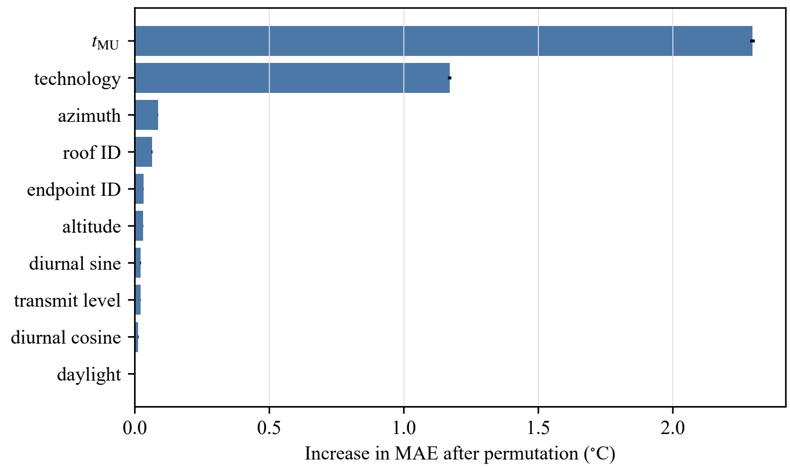

Figure 14. Held-out permutation importance for HistGBT in the December 2025 last-month validation block.

is expected because removing hardware heterogeneity reduces the spread of device offsets. The improvement should not be interpreted as a stronger operational result, however, because the single-technology subset covers fewer CML endpoints and fewer roof clusters. The pooled experiment is therefore the more relevant stress test for heterogeneous network deployment, whereas the single-technology experiment shows the performance attainable when hardware variability is reduced.

The prediction diagnostics are shown in Figs. 15–17. Figure 15 compares observed and estimated December 2025 temperatures for three representative models: the training-only constant-offset correction, the best temporal linear model (linear SVR), and the strongest boosted-tree model (XGBoost), with all panels showing the same sampled held-out observations from left to right in that order. Figure 16 shows the same December holdout for the linear SVR with points coloured by anonymized technology group, and Fig. 17 shows a representative one-week time series from the holdout for the two baselines and three supervised models. The summary of MAE across validation protocols is shown in Fig. 18, and the pooled-versus-single-technology comparison is shown in Fig. 19.

To connect the predictors back to the reduced-order thermal interpretation, Fig. 14 reports held-out permutation importance for HistGBT in the December 2025 last-month block, where the bars show the increase in MAE after permuting each original input column and categorical variables are permuted as whole columns before one-hot encoding. The dominant predictor is the microwave-unit temperature itself: permuting t_{MU} increases MAE by about 2.30 °C. The anonymized technology group is second (1.17 °C), showing that persistent hardware-family offsets remain important even after pooling. Azimuth, roof ID, endpoint ID, altitude, transmit/receive-level variables, and diurnal terms have smaller but non-zero effects. This ranking supports the interpretation that the regression models primarily correct a device/hardware thermal state and then add smaller exposure- and timing-dependent adjustments, rather than learning temperature from clock variables alone.

The leave-one-roof results should be interpreted separately from the temporal results. When a roof cluster is withheld in the pooled experiment, the global constant-offset and linear SVR errors increase to 9.44 and 4.49 °C, respectively, while the

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

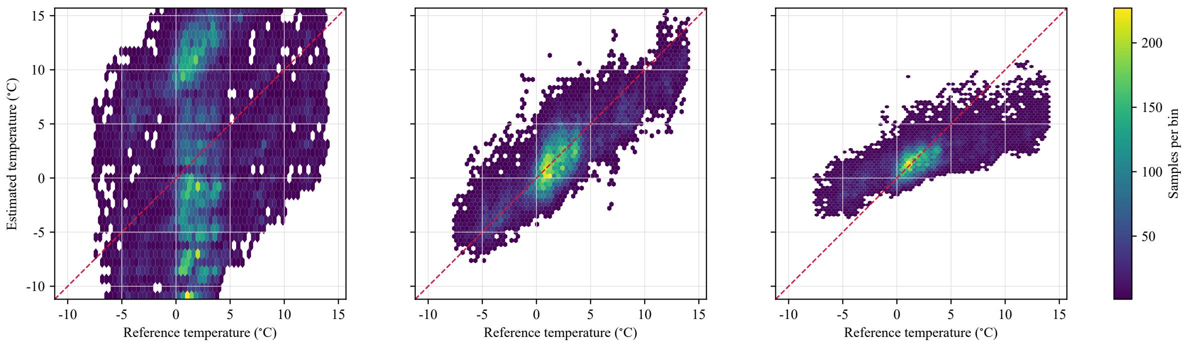

Figure 15. Observed weather-station temperature versus estimated temperature for the last-month blocked holdout (December 2025). The dashed red line indicates the 1:1 relationship corresponding to perfect agreement.

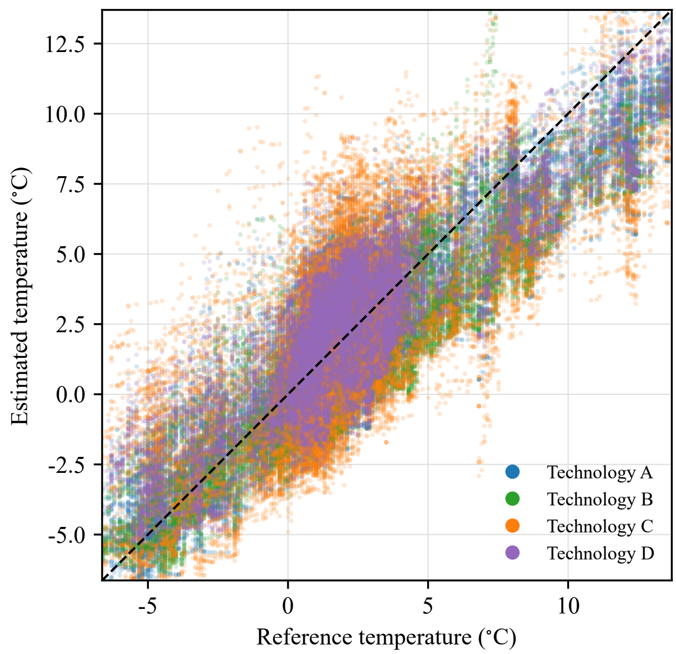

Figure 16. Linear SVR estimates in the December 2025 holdout, coloured by anonymized technology group. The figure illustrates that the pooled validation set contains several hardware families rather than a single homogeneous device type.

clock-only baseline remains at 3.08 °C because it does not rely on site-specific CML telemetry. HistGBT and XGBoost reduce this spatial-transfer error to 2.19 and 2.41 °C, respectively, suggesting that nonlinear interactions can help when the test block is an unseen roof rather than an unseen month. The five folds are heterogeneous (Table 11), and all roofs are within the same broad regional climate, so the clock-only baseline partly tests the ability to reconstruct regional temperature from calendar

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

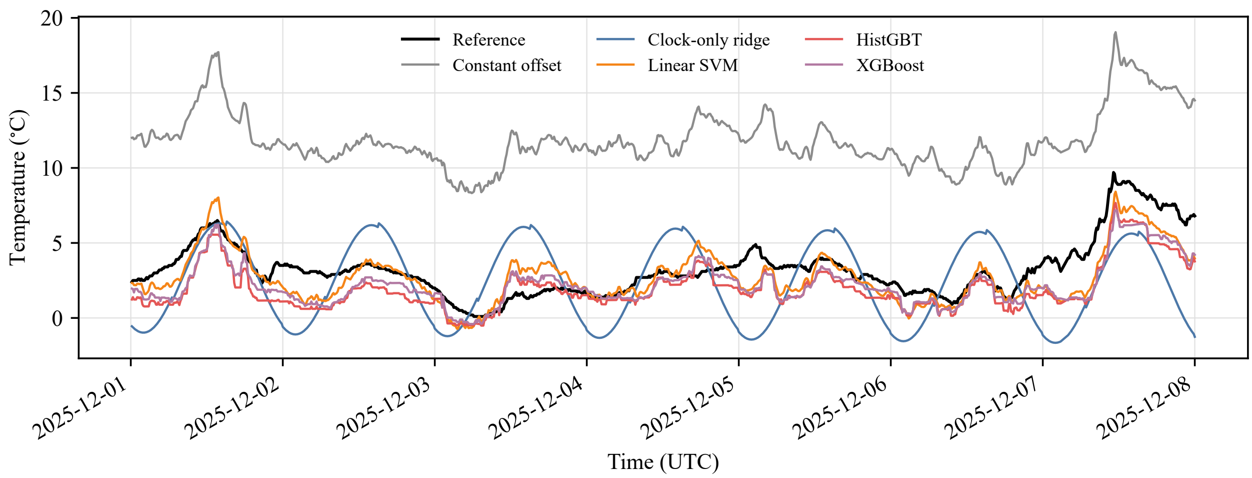

Figure 17. Example one-week time series from the December 2025 holdout for the two baselines, the three supervised regression models, and the nearby reference-station temperature.

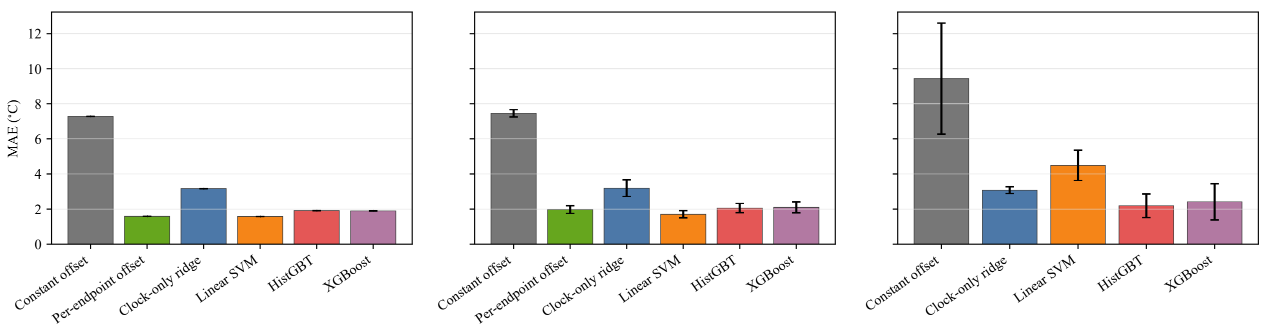

Figure 18. Mean absolute error by validation protocol. Panels from left to right show last-month, leave-one-month, and leave-one-roof validation. Error bars show the standard deviation across folds where multiple folds are available.

variables. Because the per-endpoint offset is unavailable for unseen endpoints, this protocol mainly tests whether transferable relationships can be learned from metadata, clock terms, and telemetry rather than from endpoint-specific calibration. The improvement over clock-only is meaningful but modest, so this result should be presented as partial evidence for transferable CML information, not as a completed transfer solution. The comparison is scientifically useful because it identifies the main limitation of the current dataset: the CML signal contains recoverable temperature information, but robust deployment requires calibration across more roofs, hardware families, and exposure types.

Overall, the blocked validation shows that CML telemetry contains useful information beyond a global constant-offset correction and the seasonal/diurnal clock signal for estimating the nearby reference-station temperature in temporal validation. The per-endpoint baseline in Eq. (15) is the stricter test for already observed devices, because it removes the dominant learned

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

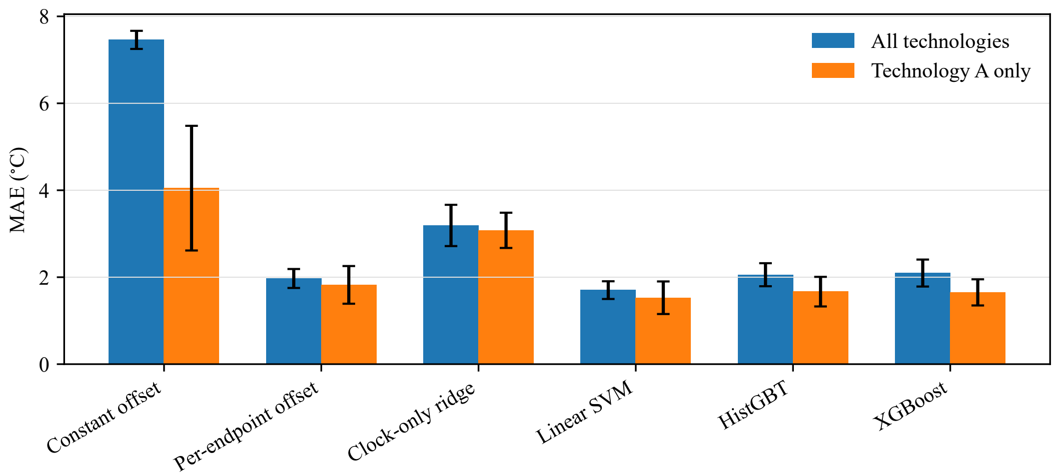

Figure 19. Leave-one-month MAE for the pooled multi-technology experiment and the Technology A-only experiment.

Table 11. Spread of MAE across the five leave-one-roof folds in the pooled multi-technology experiment. Values are in °C.

Model	Folds	Mean MAE	Min–max MAE	SD
Constant offset	5	9.44	6.92–14.39	3.16
Clock-only ridge	5	3.08	2.90–3.39	0.19
Linear SVR	5	4.49	3.07–5.18	0.86
HistGBT	5	2.19	1.56–3.02	0.67
XGBoost	5	2.41	1.61–4.07	1.03

device intercept before any supervised model is credited with additional skill. In leave-one-month validation the linear SVR improves on this stricter baseline by 0.27 °C in MAE, which is a much smaller but more meaningful margin than the comparison with the global offset. In roof-transfer validation, boosted trees also improve on the clock-only baseline, but the smaller margin shows that site transfer remains unresolved. For this reason, we treat the regression analysis as a supporting feasibility result and report all errors as agreement with a nearby weather-station reference rather than as true local rooftop air-temperature accuracy.

4 Discussion

The most important result of this study is that CML temperature telemetry contains a strong and physically interpretable atmospheric signal, but only after acknowledging that it is not a direct measurement of air temperature. The internal sensor is embedded in active electronics and the offset is large. This offset is not a minor calibration artifact; for the investigated

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

hardware, it is on the order of 25 °C. In that sense, the concern about practicality is justified. Any realistic operational use of CML temperature data must include device-specific correction, exposure-aware interpretation, and quality control after resets.

At the same time, the paper shows why CMLs remain interesting despite this limitation. They are fixed, operator-managed devices with continuous telemetry and established relevance for environmental sensing through rainfall monitoring. The ability to extract thermal information from the same network is valuable even if the signal must be interpreted as a thermal proxy rather than a standard meteorological observation. This is particularly attractive in urban environments, where network density is high and where thermal heterogeneity matters for heat monitoring and microclimate analysis. It is also attractive in multi-sensor applications, where temperature information from the same CML infrastructure could help interpret or quality-control rainfall products.

To judge whether this is practically useful, the reported error must be compared with the magnitude of the phenomena of interest. Urban heat-island and intra-city microclimate signals are typically of the order of several degrees Celsius (Bahi et al., 2019), so a per-endpoint agreement of about 1.7–2.0 °C with a nearby reference places CML telemetry at the threshold of being informative for mapping such gradients, rather than for precise point measurement. This level of agreement is of the same order as the performance reported for other opportunistic air-temperature retrievals, such as smartphone-battery-based estimates (Overeem et al., 2013; Droste et al., 2017; Song et al., 2022), but it is obtained from fixed, permanently powered, centrally managed infrastructure rather than from moving and intermittently available consumer devices. We emphasise that in our case the reported agreement is also limited by the comparison itself: the CML and the reference station are separated by hundreds of metres and differ in height, so a non-negligible part of the reported error reflects genuine differences between the two measurement points rather than device error alone.

The field results also make clear that measurement consistency is the central practical issue. The comparison of two units with different azimuths demonstrates that even colocated devices can behave differently over the day because of solar exposure and shading. The A/B comparison further shows that local infrastructure, such as heat exchangers, can generate persistent site-specific bias. These are not random errors; they are systematic local influences that can in principle be modeled, but only if adequate metadata and calibration data are available.

Another limitation is the representativeness of the weather-station reference. The station is horizontally separated by about 650 m from the rooftop site and measures temperature at standard height, whereas the microwave units are mounted higher and within a different micro-environment. Under stable atmospheric conditions, vertical gradients can influence the comparison; in urban settings, local horizontal gradients can also be non-negligible. The present analysis therefore does not validate rooftop CML telemetry against a perfect colocated 2 m reference. Instead, it demonstrates that the CML signal carries robust information about nearby ambient-temperature variability under realistic operational constraints. This distinction is especially important for the regression metrics: the reported MAE quantifies agreement with a nearby reference station, not error relative to the unknown true air temperature in the immediate surroundings of the radio housing. Although a full quantification of δ_{site} would require colocated sensors, its order of magnitude can be bracketed. Sub-kilometre horizontal air-temperature differences in built environments are commonly of the order of a few tenths to about one degree Celsius, increasing under calm, clear-sky conditions and across strong surface contrasts (de Vos et al., 2020). The dataset itself also offers an internal handle on this

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Table 12. Indicative contributions to the difference between the microwave-unit temperature and the true local or reference air temperature. Ranges are approximate and specific to the investigated hardware and sites.

Contribution	Typical magnitude (°C)
Self-heating floor $P_{\mathrm{HW}}/G_{\mathrm{th}}$	23–25
Solar/radiative loading $(q_{\mathrm{SW}} + q_{\mathrm{LW}})/G_{\mathrm{th}}$	0–several
Local site effects $q_{\mathrm{loc}}/G_{\mathrm{th}}$	0–several
Site representativeness δ_{site} (height and 0.2–1.0 km separation)	0–few
Sensor quantization and telemetry rounding	$\lesssim$1
Post-reset transient after filtering	small residual

term, because two of the roof clusters are referenced to the same station (S2) at distances of 156 and 775 m (Table 3); comparing agreement across such shared-station pairs provides a route to bounding the reference-separation contribution without additional instrumentation. Taken together, this indicates that δ_{site} contributes at most a few degrees and that a meaningful fraction of the reported MAE originates from the spatial separation between the CML and the reference rather than from the telemetry alone. Two further mechanisms act in the same direction. First, the reference itself is not error-free: standard naturally ventilated radiation shields develop daytime errors of order several tenths to more than one degree Celsius under high solar irradiance and weak wind, as shown both by classical error analyses and by recent shield intercomparisons (Nakamura and Mahrt, 2005; García Izquierdo et al., 2024; Jin et al., 2026), precisely the conditions under which the CML offset also varies most. Part of the residual CML–station disagreement is therefore attributable to the reference instrument rather than to the telemetry. Second, CML units are typically mounted on exposed roofs and towers, which are windier than the sheltered 2 m environment of the station; the air around the unit can therefore change temperature faster (lower local thermal inertia) and is not perfectly synchronized in time with the reference, so short-lived gradients and timing mismatches inflate the instantaneous error even when the two sensors track the same regional air mass.

We stress that these caveats do not weaken the present result: even with a remote reference, the CML telemetry already carries useful temperature information at the level of reference agreement demonstrated here. They instead show where the remaining error comes from and how it can be reduced. A natural next step is therefore to install accurate, low-maintenance thermometers in the immediate vicinity of the radio units. This would remove most of the reference-separation and radiation-shield contributions and isolate regional air-temperature variability from strictly local rooftop effects, and we expect it to improve the achievable agreement rather than merely provide a cleaner diagnosis.

Table 12 summarizes the main contributions as an indicative error budget. The entries should be read as order-of-magnitude ranges for the present hardware and sites, not as independently separable corrections. The self-heating term dominates the absolute offset, while radiative loading, local rooftop conditions, and representativeness control much of the remaining variability.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

The role of machine learning should also be interpreted carefully. The blocked-validation results show that the bias structure is partly learnable and that CML temperature provides useful information beyond a global constant-offset correction and beyond the seasonal and diurnal clock signal for estimating the nearby reference-station temperature. However, most of the global-offset improvement comes from device-specific intercepts, so the per-endpoint baseline isolates the remaining information beyond a static per-device correction. In the present data this remaining gain is about 0.27 °C for leave-one-month validation. The ranking of the supervised models depends on the validation target: the linear SVR gives the strongest blocked temporal performance, whereas HistGBT is best for leave-one-roof transfer and XGBoost is similar but slightly worse in the pooled spatial-transfer test. The strong clock-only performance in the leave-one-roof protocol is an important warning that seasonal and diurnal structure alone can explain a non-negligible part of the reference-station variability. The boosted-tree improvement over this baseline suggests that CML telemetry still contributes useful information, but the margin is not large enough to treat roof-to-roof transfer as resolved. A scalable application would require multiple sites, multiple hardware types, a broader range of orientations and local conditions, and explicit tests of inter-site transferability. More generally, this study suggests that any future model should combine shared predictors, such as diurnal and seasonal timing, with unit- or site-specific terms that account for exposure and hardware family. More broadly, the staged baseline ladder used here (a global offset, then technology- and endpoint-level offsets, and finally dynamic supervised models) is itself a transferable way to report opportunistic-sensing skill honestly, because it separates the large but trivial gain from device-specific intercepts from the smaller gain that genuinely reflects recoverable atmospheric information.

The reduced-order model also explains why short-term transients after restart must be treated separately from quasi-steady operation: immediately after a reset, the derivative term in Eq. (6) is non-negligible, so the observed temperature cannot be interpreted using a stationary offset assumption. The first-order fits in Table 1 give $\tau \approx 6$–8 min, and recovery to within about 5 % of the new equilibrium requires roughly 3τ, consistent with the approximately 20 min recovery durations in Table 4. A fully constrained identification of all forcing terms, especially radiative and convective forcing outdoors, would still require additional field measurements. The model should therefore be read as a physically motivated framework that organizes the observations and indicates what needs to be measured next.

Finally, the scope of the evidence should be stated clearly. The physically detailed case study is centred on one rooftop environment and one nearby weather-station reference, while the broader machine-learning experiment extends to 35 endpoints, five roof clusters, and four anonymized hardware groups. That is a useful step beyond a single-device proof of concept, but it is still not a network-scale validation. The main limitations of this study are the use of nearby rather than colocated reference air-temperature measurements, the restricted release of operator telemetry, the limited number of roof clusters in the spatial-transfer experiment, and the anonymization of hardware and site metadata. These limitations do not affect the main feasibility conclusion, but they constrain the interpretation of the reported errors as reference-station agreement rather than absolute rooftop air-temperature accuracy. Therefore, this is a feasibility study that establishes physical plausibility, quantifies the dominant limitations, and shows that careful blocked validation yields defensible agreement levels against a nearby reference.

The underlying concept is not limited to microwave radios. Many outdoor technical devices report an internal, board-level, battery, processor, or casing temperature for supervision rather than for meteorology. If such devices are fixed, continuously

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

powered or frequently sampled, and exposed to the outdoor environment, their thermal state should contain a mixture of self-heating, radiative loading, ventilation, local installation effects, and ambient-air variability similar in structure to Eq. (6). This suggests a broader class of opportunistic thermal proxies from outdoor infrastructure. However, the present results also show the necessary conditions for using such proxies responsibly: the reported temperature must not be interpreted as air temperature directly, device-specific offsets and warm-up periods must be handled explicitly, and transfer to new device families or installation types requires exposure metadata and independent validation.

Overall, the results support a cautious but positive conclusion. CMLs should not be presented as replacement thermometers, yet their telemetry can act as a promising opportunistic thermal proxy when its limitations are treated explicitly. Their main strengths are not sensor quality, but infrastructure availability, density, and the possibility of combining several environmental observables within the same telecommunication network.

5 Conclusions

This study evaluated the feasibility of using built-in temperature telemetry from CMLs as a source of opportunistic temperature information. The results show that CML temperatures strongly follow the temporal evolution of nearby ambient temperature, but they are shifted by a large and physically meaningful offset caused by self-heating, solar loading, local rooftop environment, and post-reset recovery. For the investigated hardware, the monthly mean offset is approximately 24 to 26 °C, which makes clear that the reported temperature is not a direct air-temperature measurement.

Despite this limitation, the temporal co-variability with the weather-station signal is strong, and much of the offset appears systematic rather than random. In the blocked regression analysis, the global constant-offset baseline remained insufficient for the main temporal validation, while a per-endpoint offset reduced the leave-one-month MAE to 1.97 °C. The best blocked temporal model (linear SVR) further reduced this to 1.70 °C against the nearby reference station, showing a modest gain beyond static per-device calibration. The best leave-one-roof result (HistGBT) was 2.19 °C, improving on the clock-only baseline but still indicating unresolved site transfer. These values should be interpreted as agreement with a remote reference, not as direct estimates of rooftop air-temperature accuracy. They also represent performance for the specific validation tasks used here: blocked extrapolation in time for already observed endpoints and, separately, transfer to unseen roof clusters within the same broader network.

The present results therefore support CMLs as a plausible supplementary temperature-sensing layer, especially where dense fixed infrastructure is already available and where the same network is of interest for rainfall monitoring. At the same time, the cross-roof experiment shows that site transfer remains the central challenge. The telemetry already provides useful temperature information at the level of reference agreement demonstrated here; future work should focus on colocated reference sensors at radio sites, which we expect to improve agreement by removing the reference-separation and radiation-shield contributions of the present remote reference, together with more diverse datasets covering multiple vendors and installation types, and explicit evaluation of transferability across different environments. Dedicated outdoor thermal experiments should also be used to constrain radiative and convective forcing terms beyond the first-order warm-up time constant estimated here. Under those

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

conditions, CML temperature data could contribute to high-density urban thermal monitoring, including future attempts at three-dimensional urban thermal-field monitoring that exploit the range of installation heights across a dense network, and to multi-parameter environmental sensing from existing telecommunication networks. The same physical interpretation may also be useful for other fixed outdoor devices that routinely report internal or casing temperature, provided that their non-meteorological thermal biases are calibrated rather than ignored.

Code and data availability. Raw operational telemetry cannot be released because it contains operator-sensitive infrastructure information, including exact link locations and metadata that could reveal network topology, assets, or customers. The source network includes not only conventional mobile-backhaul infrastructure, such as base-station and tower links, but also links serving business-to-business customers. Publishing exact endpoint coordinates could therefore disclose customer premises and commercial connectivity relationships, which is a stronger privacy and confidentiality constraint than in many publicly released telco-oriented CML datasets that mainly describe operator infrastructure. To support transparency, an anonymized reduced preview dataset is available from Zenodo (Mlýnek et al., 2026) at https://doi.org/10.5281/zenodo.20718359. The archive contains a CSV preview file documenting the feature structure used in the regression analysis and allowing inspection of example relationships between CML unit temperature, auxiliary predictors, and nearby weather-station reference temperature. It is not the full operational 2025 regression matrix and cannot be used to reproduce the reported blocked-validation metrics exactly, because the full matrix contained multiple roof clusters, endpoint identifiers, and spatial metadata required for the leave-one-roof protocol and the per-endpoint offset baseline. The blocked-validation and analysis code compatible with these confidentiality constraints is available at https://github.com/TelcoSense/cml-temperature-network. Richer data may be made available on reasonable request under a data-use agreement, subject to data-owner approval.

Author contributions. PM conceived the study. MI, SM, and JH performed the experiments and analyses. PM and MI prepared the manuscript with input from all co-authors. PM and PML supervised the work and acquired funding.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. AI-assisted editing with OpenAI Codex was used during manuscript preparation for language polishing, consistency checks, and LaTeX-editing support. All scientific content, analyses, interpretations, and conclusions were reviewed and verified by the authors.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Financial support. This research was supported by the Technology Agency of the Czech Republic (TA CR) under the Environment for Life Programme, project SS07020434 "Spatial Air Temperature Monitoring Using Microwave Links Data". The publication was supported by the project "Energy Conversion and Storage" (CZ.02.01.01/00/22_008/0004617), funded through the Johannes Amos Comenius Programme, call Excellent Research.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

References

Adirosi, E., Facheris, L., Giannetti, F., Scarfone, S., Bacci, G., Mazza, A., Ortolani, A., and Baldini, L.: Evaluation of Rainfall Estimation Derived From Commercial Interactive DVB Receivers Using Disdrometer, Rain Gauge, and Weather Radar, IEEE Transactions on Geoscience and Remote Sensing, 59, 8978–8991, https://doi.org/10.1109/TGRS.2020.3041448, 2021.

Angeloni, S., Adirosi, E., Sapienza, F., Giannetti, F., Francini, F., Magherini, L., Valgimigli, A., Vaccaro, A., Melani, S., Antonini, A., and Baldini, L.: Enhanced Estimation of Rainfall From Opportunistic Microwave Satellite Signals, IEEE Transactions on Geoscience and Remote Sensing, 62, 1–12, https://doi.org/10.1109/TGRS.2023.3349100, 2024.

Bahi, H., Radoine, H., and Mastouri, H.: Urban Heat Island: State of the Art, in: 2019 7th International Renewable and Sustainable Energy Conference (IRSEC), pp. 1–7, https://doi.org/10.1109/IRSEC48032.2019.9078329, 2019.

Bellavista, P., Magistretti, E., Lee, U., and Gerla, M.: Standard Integration of Sensing and Opportunistic Diffusion for Urban Monitoring in Vehicular Sensor Networks: The MobEyes Architecture, in: 2007 IEEE International Symposium on Industrial Electronics, pp. 2582–2588, Vigo, Spain, https://doi.org/10.1109/ISIE.2007.4375015, 2007.

Bianchi, B., van Leeuwen, P. J., Hogan, R. J., and Berne, A.: Rainfall nowcasting by combining radars, microwave links and rain gauges, https://arxiv.org/abs/1810.11811, 2018.

Blettner, N., Chwala, C., Haese, B., Hörning, S., and Kunstmann, H.: Combining Commercial Microwave Link and Rain Gauge Observations to Estimate Countrywide Precipitation: A Stochastic Reconstruction and Pattern Analysis Approach, Water Resources Research, 58, e2022WR032 563, https://doi.org/10.1029/2022WR032563, e2022WR032563, 2022.

Bubniak, M., Musil, P., Mlýnek, P., Benesl, L., and Rusz, M.: Design and Evaluation of the Application for a Measurement of Rainfall via Microwave Links, in: 2022 14th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT), pp. 8–13, Valencia, Spain, https://doi.org/10.1109/ICUMT57764.2022.9943528, 2022.

Chang, C.-C. and Lin, C.-J.: LIBSVM: A library for support vector machines, ACM Transactions on Intelligent Systems and Technology, 2, 1–27, https://doi.org/10.1145/1961189.1961199, 2011.

Chen, T. and Guestrin, C.: XGBoost: A Scalable Tree Boosting System, in: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 785–794, https://doi.org/10.1145/2939672.2939785, 2016.

Chwala, C. and Kunstmann, H.: Commercial microwave link networks for rainfall observation: Assessment of the current status and future challenges, WIREs Water, 6, e1337, https://doi.org/10.1002/wat2.1337, e1337 WATER-363.R2, 2019.

Cortes, C. and Vapnik, V.: Support-Vector Networks, Machine Learning, 20, 273–297, https://doi.org/10.1007/BF00994018, 1995.

Czech Telecommunication Office: Evidence of stations according to VO-R, https://ctu.gov.cz/en/evidence-links-according-vo-r, web portal and explanatory page for station/link registration; accessed 2026-03-19, 2022.

Czech Telecommunication Office: Technické údaje pevných rádiových systémů typu bod-bod v pásmech 71–76 GHz a 81–86 GHz, https://ctu.gov.cz/vyhledavaci-databaze/technicke-udaje-pevnych-radiovych-systemu-typu-bod-bod-v-pasmech-71-76-ghz-a-81-86-ghz/vyhledavani, public search database; accessed 2026-06-30, 2026.

de Vos, L. W., Droste, A. M., Zander, M. J., Overeem, A., Leijnse, H., Heusinkveld, B. G., Steeneveld, G. J., and Uijlenhoet, R.: Hydrometeorological Monitoring Using Opportunistic Sensing Networks in the Amsterdam Metropolitan Area, Bulletin of the American Meteorological Society, 101, 167–185, https://doi.org/10.1175/BAMS-D-19-0091.1, 2020.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Doumounia, A., Gosset, M., Cazenave, F., Kacou, M., and Zougmore, F.: Rainfall monitoring based on microwave links from cellular telecommunication networks: First results from a West African test bed, Geophysical Research Letters, 41, 6016–6022, https://doi.org/10.1002/2014GL060724, 2014.

Droste, A. M., Pape, J. J., Overeem, A., Leijnse, H., Steeneveld, G. J., and Uijlenhoet, R.: Crowdsourcing Urban Air Temperatures through Smartphone Battery Temperatures in São Paulo, Brazil, Journal of Atmospheric and Oceanic Technology, 34, 1853–1866, https://doi.org/10.1175/JTECH-D-16-0150.1, 2017.

Fencl, M., Rieckermann, J., Schleiss, M., Stránský, D., and Bareš, V.: Assessing the potential of using telecommunication microwave links in urban drainage modelling, Water Science and Technology, 68, 1810–1818, https://doi.org/10.2166/wst.2013.429, 2013.

Fencl, M., Dohnal, M., Valtr, P., Grabner, M., and Bareš, V.: Atmospheric observations with E-band microwave links – challenges and opportunities, Atmospheric Measurement Techniques, 13, 6559–6578, https://doi.org/10.5194/amt-13-6559-2020, 2020.

Fencl, M., Nebuloni, R., Andersson, J. C. M., Bares, V., Blettner, N., Cazzaniga, G., Chwala, C., Colli, M., de Vos, L., El Hachem, A., Galdies, C., Giannetti, F., Graf, M., Jacoby, D., Habi, H. V., Musil, P., Ostrometzky, J., Roversi, G., Sapienza, F., Seidel, J., Spackova, A., van de Beek, R., Walraven, B., Wilgan, K., and Zheng, X.: Data formats and standards for opportunistic rainfall sensors, Open Research Europe, 3, https://doi.org/10.12688/openreseurope.16068.2, 2024.

García Izquierdo, C., Hernandez, S., Parrondo, M., Casas, A., Viola, A., Mazzola, M., Merlone, A., and Roulet, Y.-A.: COAT Project: Intercomparison of Thermometer Radiation Shields in the Arctic, Atmosphere, 15, 841, https://doi.org/10.3390/atmos15070841, 2024.

Gianoglio, C., Colli, M., Zani, S., and Caviglia, D. D.: An Online Training Procedure for Rain Detection Models Applied to Satellite Microwave Links, IEEE Geoscience and Remote Sensing Letters, 20, 1–5, https://doi.org/10.1109/LGRS.2023.3328718, 2023.

Jacoby, D., Yu, S., Hu, Q., Hine, Z., Johnson, R., Ostrometzky, J., Kadota, I., Zussman, G., and Messer, H.: OpenMesh: Wireless Signal Dataset for Opportunistic Urban Weather Sensing in New York City, Earth System Science Data Discussions, 2025, 1–27, https://doi.org/10.5194/essd-2025-238, preprint, in review, 2025.

Jin, W., Zhou, Y., Tang, J., and Amdadul, H. M.: Design and Experimental Validation of a High-Accuracy Naturally Ventilated Radiation Shield for Near-Surface Air Temperature Observation, Atmosphere, 17, 272, https://doi.org/10.3390/atmos17030272, 2026.

Liu, W., Dong, S., Zheng, J., Liu, C., Wang, C., Shangguan, W., Zhang, Y., and Zhang, Y.: Quantifying the Rainfall Cooling Effect: The Importance of Relative Humidity in Guangdong, South China, Journal of Hydrometeorology, 23, 875–889, https://doi.org/10.1175/JHM-D-21-0155.1, 2022.

Messer, H., Zinevich, A., and Alpert, P.: Environmental Monitoring by Wireless Communication Networks, Science, 312, 713–713, https://doi.org/10.1126/science.1120034, 2006.

Ministry of Industry and Trade of the Czech Republic: Analýza současného a budoucího využití pásma 42 GHz, https://mpo.gov.cz/assets/cz/e-komunikace-a-posta/studie-k-narodnimu-planu-obnovy/2024/6/Studie-3.pdf, study published 2024-06-22, 2024.

Mlýnek, P., Bouzek, K., Paušová, Š., Istvanek, M., and Musil, P.: Dataset of "Feasibility and limits of commercial microwave links for opportunistic air-temperature monitoring", https://doi.org/10.5281/zenodo.20718359, 2026.

Nakamura, R. and Mahrt, L.: Air Temperature Measurement Errors in Naturally Ventilated Radiation Shields, Journal of Atmospheric and Oceanic Technology, 22, 1046–1058, https://doi.org/10.1175/JTECH1762.1, 2005.

Nielsen, J. M., van de Beek, C. Z. R., Thorndahl, S., Olsson, J., Andersen, C. B., Andersson, J. C. M., Rasmussen, M. R., and Nielsen, J. E.: Merging weather radar data and opportunistic rainfall sensor data to enhance rainfall estimates, Atmospheric Research, 300, 107 228, https://doi.org/10.1016/j.atmosres.2024.107228, 2024.

https://doi.org/10.5194/egusphere-2026-3888
Preprint. Discussion started: 14 August 2026

Overeem, A., Robinson, J. C. R., Leijnse, H., Steeneveld, G. J., and Horn, B. K. P.: Crowdsourcing urban air temperatures from smartphone battery temperatures, Geophysical Research Letters, 40, 4081–4085, https://doi.org/10.1002/grl.50786, 2013.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E.: Scikit-learn: Machine Learning in Python, Journal of Machine Learning Research, 12, 2825–2830, 2011.

Song, K., Liu, X., and Gao, T.: Potential Application of Using Smartphone Sensor for Estimating Air Temperature: Experimental Study, IEEE Internet of Things Journal, 9, 14 300–14 306, https://doi.org/10.1109/JIOT.2021.3063488, 2022.

Trivedi, A., Bovornkeeratiroj, P., Breda, J., Shenoy, P., Taneja, J., and Irwin, D.: Phone-based ambient temperature sensing using opportunistic crowdsensing and machine learning, Sustainable Computing: Informatics and Systems, 29, 100 479, https://doi.org/10.1016/j.suscom.2020.100479, 2021.

Uijlenhoet, R., Overeem, A., and Leijnse, H.: Opportunistic remote sensing of rainfall using microwave links from cellular communication networks, WIREs Water, 5, e1289, https://doi.org/10.1002/wat2.1289, 2018.

Ukil, A., Braendle, H., and Krippner, P.: Distributed Temperature Sensing: Review of Technology and Applications, IEEE Sensors Journal, 12, 885–892, https://doi.org/10.1109/JSEN.2011.2162060, 2012.

Zheng, X., Messer, H., Wang, Q., Xu, T., Qin, Y., and Yang, T.: On the potential of commercial microwave link networks for high spatial resolution rainfall monitoring in urban areas, Atmospheric Research, 277, 106 289, https://doi.org/10.1016/j.atmosres.2022.106289, 2022.

Zinevich, A., Alpert, P., and Messer, H.: Estimation of rainfall fields using commercial microwave communication networks of variable density, Advances in Water Resources, 31, 1470–1480, https://doi.org/10.1016/j.advwatres.2008.03.003, hydrologic Remote Sensing, 2008.