



Conditioning-controlled retrieval of broadband land surface temperature and emissivity from paired ground-based longwave irradiance measurements

Collins Mito

Department of Physics, Faculty of Science and Technology, University of Nairobi, Nairobi, P.O. Box 30197-00100, Kenya

Correspondence: Collins Mito (collins@uonbi.ac.ke)

Received: 12 February 2026 – Discussion started: 2 April 2026

Revised: 6 September 2026 – Accepted: 7 September 2026 – Published: 23 September 2026

Abstract. Accurate retrieval of land surface temperature (T_s) from broadband longwave radiometric measurements is limited by the nonlinear coupling between surface emissivity (ϵ) and temperature, particularly when the available radiative contrast provides insufficient information to distinguish their contributions. A conditioning-controlled framework is presented for jointly retrieving broadband ϵ and T_s directly from high-temporal-resolution paired ground-based upwelling (L) and downwelling (I) longwave irradiance measurements, without requiring externally prescribed emissivity as a retrieval input. The framework constructs a local population of candidate temporal pairs and combines quasi-steady thermal-state screening, numerical and physical admissibility, explicit irradiance-identifiability testing, uncertainty- and information-aware progressive pair selection, physical-state-dependent reciprocal emissivity treatment, fixed eight-update Newton–Raphson refinement, and correlated uncertainty propagation. Field evaluation used 637 datasets from nine SURFRAD and BSRN stations spanning vegetated, mixed, and bare/desert environments. Five datasets lacked the required finite initial irradiances, leaving 632 usable cases, all of which produced final finite retrievals; 628 reached the stable-centre selection criterion and four required the prescribed fallback pathway. Across the 632 field cases, retrieved T_s had a bias of -0.162 K, mean absolute error (MAE) of 0.432 K, root-mean-square error (RMSE) of 0.558 K, and $R^2 = 0.9984$ relative to the ground-based validation-reference temperature. A complementary physics-based known-truth experiment independently prescribed ϵ_{true} and $T_{s,\text{true}}$ for 6000 scenes spanning $0.85 \leq \epsilon_{\text{true}} \leq 0.99$. In the noiseless experiment, 5924 scenes (98.73 %) yielded successful retrievals; among successful cases, emissivity bias,

MAE, and RMSE were 0.0313 , 0.0491 , and 0.0658 , respectively, while the corresponding T_s values were -0.335 , 1.307 , and 3.213 K. Retrieval errors increased toward the lowest-emissivity regime, identifying low-emissivity conditions as the most demanding part of the tested retrieval space. The combined field and known-truth results demonstrate that broadband emissivity and surface temperature can be jointly retrieved from paired longwave irradiance observations when radiative observability and physical applicability are explicitly conditioned before nonlinear inversion.

1 Introduction

Land surface temperature (T_s) is a fundamental variable in studies of surface energy balance, land-atmosphere coupling, hydrology, and climate variability (Dash et al., 2002; Li et al., 2013b, 2023; Wan, 2014). Accurate estimation of T_s is important for quantifying surface fluxes, diagnosing land-atmosphere feedbacks, and evaluating thermal infrared satellite products. Although satellite observations provide broad spatial coverage, ground-based radiometric networks remain essential for algorithm development, physical interpretation, and validation because they provide continuous longwave measurements at high temporal resolution.

At the land-atmosphere interface, measured broadband upwelling longwave irradiance contains both surface thermal emission and reflected atmospheric downwelling irradiance. Conversion of this measurement to T_s is therefore inseparable from broadband surface emissivity. In a single broadband observation, emissivity and temperature are nonlinearly cou-

pled and cannot generally be retrieved independently without additional information. Even when multiple observations are available, the inverse problem can remain weakly constrained if the observations provide insufficient radiative contrast. Small measurement perturbations may then produce disproportionately large changes in retrieved emissivity and, through the emissivity correction, in surface temperature. The problem is therefore one of both radiative physics and parameter observability.

More generally, if N temporal observations are considered while emissivity is treated as common but surface temperature is allowed to vary, the system contains N radiometric equations for N temperature states plus one common emissivity – that is, N equations for $N + 1$ unknowns before additional physical information is introduced. The quasi-steady, physical-admissibility, irradiance-identifiability, and population-information constraints therefore address this structural underdetermination, whereas low irradiance contrast is a separate factor that can further weaken conditioning by reducing measurement sensitivity.

A common approach in ground-based T_s estimation is to prescribe emissivity externally, often from satellite-derived narrowband emissivity products that are subsequently converted to broadband values. Such products have been widely used with SURFRAD measurements (Li et al., 2014; Zhang et al., 2016; Wang et al., 2019). However, externally prescribed emissivities can introduce additional uncertainty through spatial mismatch between a point-scale radiometric observation and a satellite pixel, temporal mismatch between an instantaneous surface state and temporally composited emissivity products, and spectral mismatch between narrowband and broadband emissivity (Gillespie et al., 1998; Hulley and Hook, 2011; Li et al., 2013a). These limitations motivate methods that can infer broadband emissivity directly from the longwave measurements themselves rather than requiring emissivity as an external retrieval input.

High-temporal-resolution ground-based longwave observations provide one possible source of the additional information required for such a retrieval. Over sufficiently short intervals, broadband surface emissivity can reasonably be treated as invariant and the thermodynamic surface state may change only gradually, while atmospheric forcing and the measured longwave irradiances continue to vary. Paired observations can therefore provide two radiative states for the same emissivity and improve the potential observability of the emissivity-temperature system. Ground-based networks such as SURFRAD are particularly suitable for investigating this approach because they provide continuous, well-calibrated upwelling and downwelling longwave measurements at high temporal resolution (Augustine et al., 2000, 2005).

Temporal pairing alone, however, does not guarantee that the added observation contains useful independent information. A pair may satisfy a short temporal-separation requirement yet remain weakly informative if the changes in up-

welling and downwelling irradiance are dominated by common atmospheric forcing. Conversely, increasing the temporal separation can improve radiative contrast but can eventually violate the assumption of a quasi-steady surface state. The relevant problem is therefore not simply how far apart two observations should be, but whether a local population of candidate observations contains enough physically distinguishable information to constrain emissivity without introducing excessive thermodynamic change.

This distinction motivates an explicit separation between thermal-state consistency, irradiance identifiability, and information sufficiency. Thermal-state consistency limits changes in apparent surface temperature over a candidate interval. Irradiance identifiability tests whether the temporal change in upwelling irradiance contains a resolvable surface-sensitive component relative to the contemporaneous change in downwelling irradiance. Information sufficiency then asks whether the local candidate population contains enough effectively independent observations for a stable emissivity state to be identified. These controls address different limitations of the inverse problem and should not be reduced to a single temporal-separation or numerical-convergence criterion.

The numerical solution provides another layer of information but cannot substitute for radiometric observability. Newton–Raphson inversion can solve the nonlinear broadband irradiance equation once a physically meaningful initial state and informative observation pair have been established. The residual and local derivative of the nonlinear equation are useful diagnostics of numerical consistency and sensitivity, but small residuals or completed iterations do not demonstrate that arbitrary observation pairs are well conditioned. Numerical diagnostics are therefore most meaningful when interpreted after physical and information-based screening rather than as the sole basis for retrieval acceptance.

A further issue concerns the physical state used to initialize the nonlinear emissivity inversion. Algebraically related emissivity representations can correspond to different local radiative states and should not necessarily be applied uniformly across all observations. A physically conditioned retrieval should therefore consider not only whether sufficient information is available, but also whether the local longwave state supports the emissivity transformation being applied. This motivates a retrieval architecture in which temporal information selection and physical-state applicability are evaluated before the final nonlinear refinement.

Uncertainty propagation is similarly important but requires careful interpretation. Formal propagation from irradiance uncertainty through the emissivity inversion provides a measure of local measurement sensitivity and enables the contributions of irradiance and emissivity uncertainty to T_s uncertainty to be separated. Such propagated uncertainties are valuable diagnostics of the retrieval geometry, but they are not automatically calibrated confidence intervals. Independent known-truth experiments are required to

determine whether the nominal uncertainty magnitudes adequately represent actual retrieval error across different emissivity regimes.

Validation of broadband emissivity presents an additional challenge. Field surface-temperature references can be constructed from ground-based longwave irradiances using externally specified emissivity, but such references are not fully independent of the radiometric measurements used by the retrieval. They provide realistic field validation of the combined retrieval behaviour, but they do not provide independent broadband emissivity truth. Direct evaluation of emissivity retrieval accuracy therefore requires a complementary experiment in which both emissivity and temperature are prescribed independently before the longwave observations are generated.

This study develops a conditioning-controlled framework for jointly retrieving broadband surface emissivity and T_s from paired ground-based upwelling (L) and downwelling (l) longwave irradiance observations. The framework constructs a local population of temporal pairs and combines quasi-steady thermal-state screening, numerical and physical admissibility, explicit irradiance-identifiability testing, uncertainty- and information-aware progressive pair selection, physical-state-dependent reciprocal emissivity treatment, Newton–Raphson refinement, and correlated uncertainty propagation. Information content is assessed before nonlinear inversion, and the selected production pair is not ranked according to its emissivity value.

Here, “conditioning-controlled” denotes the use of physical-state, irradiance-identifiability, and information-sufficiency controls to determine which observation geometry is passed to the nonlinear solver; it does not refer solely to a numerical condition-number threshold.

The framework is evaluated using two complementary validation strategies. The field experiment contains 637 source datasets from nine SURFRAD and BSRN stations spanning vegetated, mixed, and bare/desert environments; 632 datasets contain the required initial irradiances and form the final field-validation population. A separate physics-based experiment contains 6000 independently prescribed truth scenes spanning $0.85 \leq \epsilon_{\text{true}} \leq 0.99$, with exactly 2000 scenes in each of three predefined emissivity regimes. The synthetic observations are generated using the exact forward broadband longwave relation, and the same truth scenes are evaluated under noiseless and production-noise conditions without using the validation results to retune the retrieval.

The study addresses three related questions: (i) can a local population of paired longwave observations be screened and selected in a way that provides stable emissivity information without prescribing a fixed temporal separation; (ii) how do observable radiative contrast, local nonlinear sensitivity, and propagated retrieval uncertainty relate within the accepted field population; and (iii) can the resulting retrieval recover independently prescribed broadband emissivity and surface temperature across a broadened emissivity range? Section 2

summarizes the physical basis of broadband longwave coupling and apparent surface temperature. Section 3 presents the conditioning-controlled retrieval methodology. Section 4 describes the field and physics-based known-truth validation experiments. Section 5 presents the results and discussion, and Sect. 6 summarizes the principal findings and limitations.

2 Physical and conceptual background

2.1 Broadband surface-atmosphere longwave coupling

Broadband longwave irradiance exchange at the land-atmosphere interface is governed principally by surface temperature, surface emissivity, and atmospheric downwelling emission (Rodgers, 2000). For an opaque surface and for measurements sufficiently close to the surface that additional atmospheric path terms between the surface and radiometer can be neglected, the measured upwelling irradiance L can be written as

$$L = \epsilon \sigma T_s^4 + (1 - \epsilon)l. \quad (1)$$

Here, l is the measured downwelling longwave irradiance, ϵ is broadband surface emissivity, T_s is physical surface temperature, and $\sigma = 5.670374419 \times 10^{-8} \text{ W m}^{-2} \text{ K}^{-4}$ is the Stefan–Boltzmann constant. The first term in Eq. (1) represents thermal emission by the surface, whereas the second represents the reflected fraction of atmospheric downwelling irradiance. If ϵ is externally specified, Eq. (1) can be solved directly for T_s . If both T_s and ϵ are unknown, however, a single observation provides one radiometric equation for two coupled surface variables.

2.2 Apparent surface temperature as an observable state diagnostic

Taking the surface as a unit-emissivity emitter provides an observable radiometric temperature directly from the measured upwelling irradiance:

$$T'_s = \left(\frac{L}{\sigma} \right)^{1/4}. \quad (2)$$

Defining the equivalent radiative temperature of the downwelling irradiance as $T_l = (l/\sigma)^{1/4}$, Eqs. (1) and (2) give

$$T_s'^4 = \epsilon T_s^4 + (1 - \epsilon)T_l^4. \quad (3)$$

Thus, T'_s represents an emissivity-weighted radiometric state formed from the physical surface and atmospheric longwave contributions. It should not in general be interpreted as the physical T_s , nor as a universal upper or lower bound on T_s . Its principal value here is diagnostic: because T'_s is directly observable, changes in T'_s between two nearby observations indicate whether the combined surface radiative state

has changed appreciably. The retrieval therefore uses the difference $T'_s(t_2) - T'_s(t_1)$ as a quasi-steady screening variable rather than assuming that temporal proximity alone guarantees thermal stability.

2.3 Temporal pairing, identifiability, and conditioning

For two observations, the same emissivity can be written in the two radiative balances $L_1 = \epsilon \sigma T_{s,1}^4 + (1 - \epsilon)l_1$ and $L_2 = \epsilon \sigma T_{s,2}^4 + (1 - \epsilon)l_2$. Over a sufficiently short interval, broadband emissivity can be treated as invariant. If the physical surface-temperature change is also sufficiently limited, the two observations provide additional information through their differing radiative states.

The existence of two observations does not by itself make emissivity identifiable. Define $\Delta L = L_2 - L_1$ and $\Delta l = l_2 - l_1$. If the change in L closely follows the change in l , the pair may contain little additional surface-sensitive information even though both irradiances are measured accurately. Conversely, a pair becomes more informative when the upwelling change contains a resolvable component beyond the contemporaneous downwelling change. This motivates the distinction between temporal admissibility and irradiance identifiability.

Observability refers to whether the measurements contain information about the unknown surface state; irradiance identifiability refers more specifically to whether a candidate pair contains a resolvable surface-sensitive contrast relative to changes in atmospheric downwelling irradiance; and numerical conditioning describes the sensitivity of the nonlinear solution to perturbations in the measurements or state variables. These concepts are treated separately in the retrieval. Physical and radiometric information are assessed before the final nonlinear inversion, and Newton residuals and local sensitivities are interpreted as diagnostics of an already selected production state rather than as substitutes for observability.

3 Emissivity and surface temperature retrieval

The retrieval combines radiometric preprocessing, population-based candidate construction and screening, irradiance-identifiability testing, information-aware progressive pair selection, physically conditioned emissivity initialization, nonlinear Newton–Raphson refinement, and uncertainty propagation. The complete production sequence is summarized in Fig. 1.

3.1 Broadband longwave formulation

When emissivity is known, Eq. (1) can be rearranged to give

$$T_s = \left(\frac{L - (1 - \epsilon)l}{\epsilon \sigma} \right)^{1/4}. \quad (4)$$

When ϵ is unknown, Eq. (4) emphasizes the nonlinear coupling between emissivity and temperature. Apparent temperature from Eq. (2) is therefore used only as a short-term state diagnostic, not as the retrieved physical temperature.

3.2 Radiometric preprocessing and candidate temporal-pair construction

The measured $L(t)$ and $l(t)$ series are processed with a third-order Savitzky–Golay filter using an 11-sample frame. Missing-value flags are converted to missing values before processing. Where smoothing requires a complete local sequence, missing samples are temporarily filled by linear interpolation with nearest-value end treatment and the original missing locations are restored immediately after filtering; interpolation therefore enables smoothing but does not create valid observations at originally missing times.

For each field dataset, an initial network-specific time pair defines a validation midpoint t_m . The midpoint is used only as the temporal reference for the subsequent population-based search. All distinct candidate pairs (t_i, t_j) , with $j > i$, are constructed subject to both endpoints lying within ± 10 min of t_m and to a maximum pair separation of 10 min,

$$|t_i - t_m| \leq 10 \text{ min}, \quad |t_j - t_m| \leq 10 \text{ min}, \quad (5)$$

$$0 < |t_j - t_i| \leq 10 \text{ min}. \quad (6)$$

The 10 min interval defines the maximum search geometry rather than a prescribed production separation. The actual search proceeds progressively through radii of 1, 2, ..., 10 min.

3.3 Physical, numerical, and irradiance-identifiability screening

All four irradiances L_1, l_1, L_2, l_2 must be finite. A candidate must also satisfy the quasi-steady apparent-temperature condition

$$|T'_s(t_2) - T'_s(t_1)| \leq 1.0 \text{ K}. \quad (7)$$

Equation (7) limits thermodynamic-state change; satisfying this criterion alone, however, does not establish that the inverse problem is sufficiently observable.

For each surviving pair, introduce the intermediate quantities

$$N = L_1 - l_1, \quad (8)$$

$$C_1 = \frac{1}{4}l_1^{-3/4}, \quad C_2 = \frac{1}{4}l_2^{-3/4}, \quad (9)$$

$$M = \frac{1}{2}((L_1 - l_1)C_1 + (L_2 - l_2)C_2), \quad (10)$$

$$J = L_1^{1/4} - \frac{1}{2}(L_1^{1/4} + L_2^{1/4} - l_1^{1/4} - l_2^{1/4}), \quad (11)$$

and $D = J^4 - l_1$, the candidate linear emissivity is

$$\epsilon_{\text{lin}} = \frac{N - 4MJ^3}{D}. \quad (12)$$

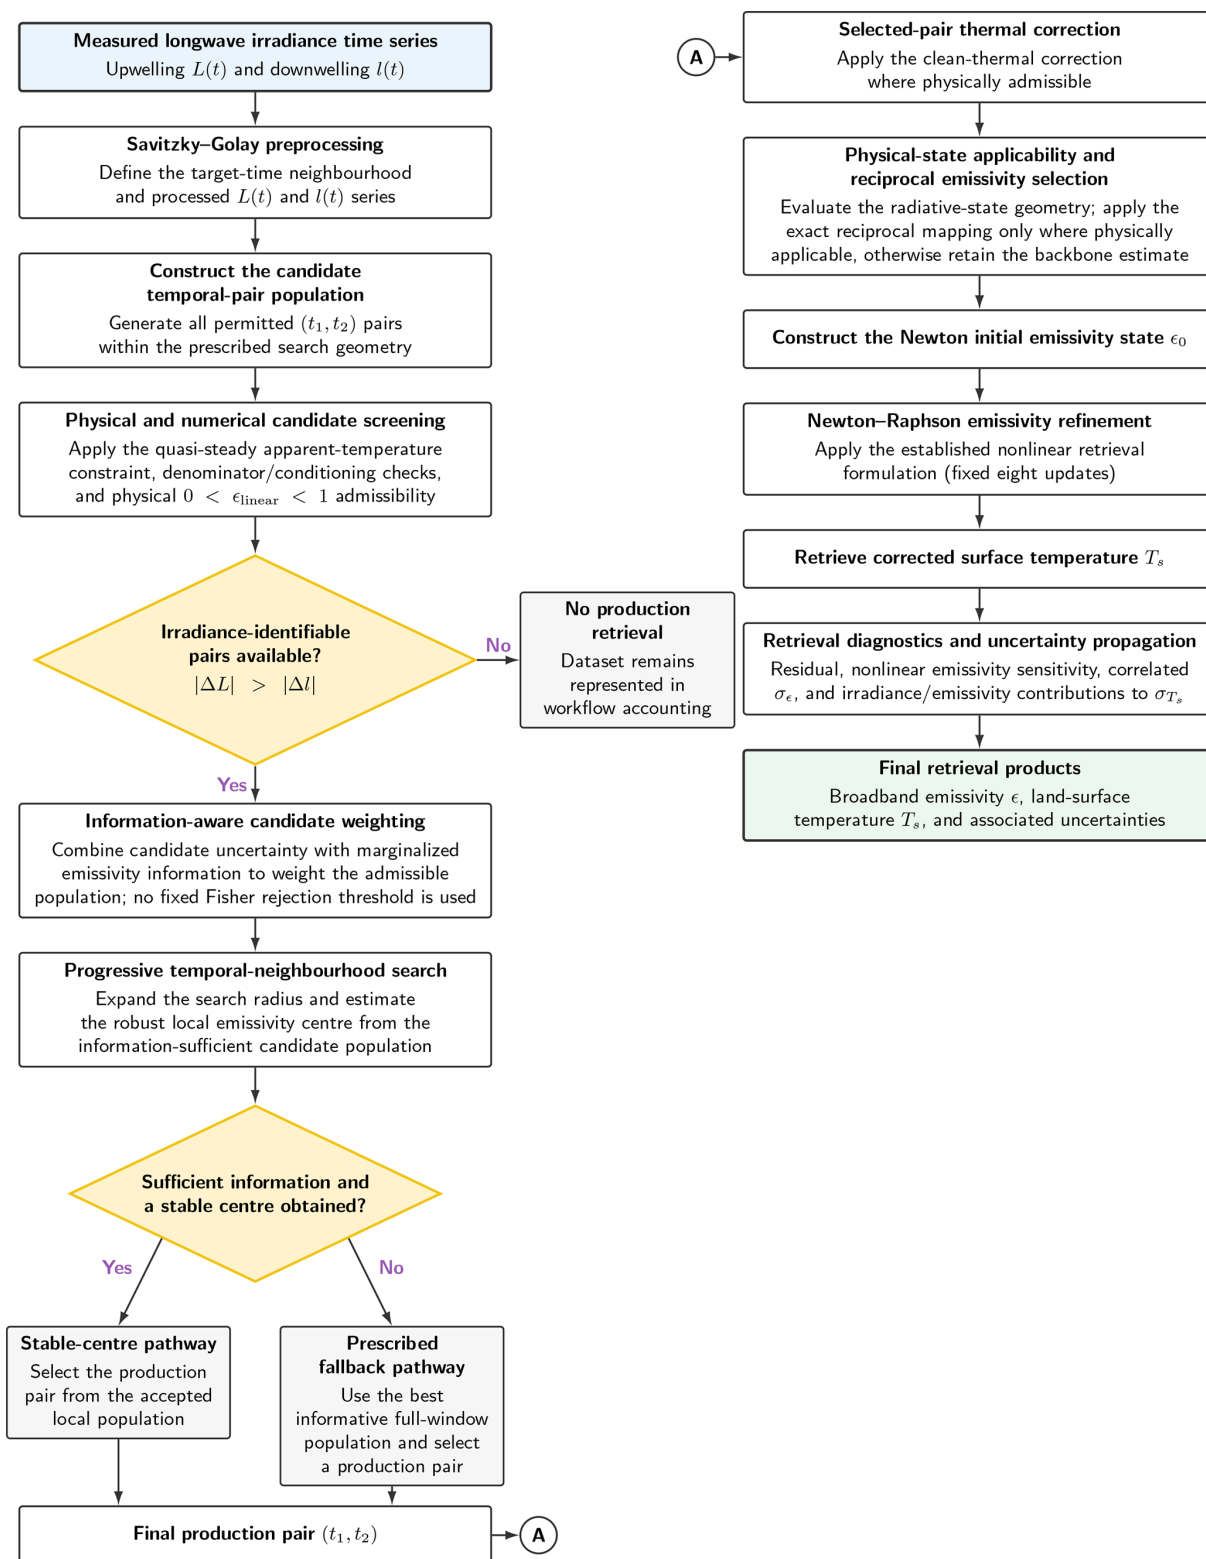


Figure 1. Conditioning-controlled production retrieval architecture, from measured $L(t)$ and $I(t)$ through candidate screening and information-aware pair selection to physical-state conditioning, Newton–Raphson refinement, surface-temperature retrieval, and uncertainty propagation. Matching A connectors indicate continuation from the final production pair to the selected-pair thermal-correction stage.

Numerical screening then uses the relative denominator-conditioning measure

$$\kappa_D = \frac{\max(|J^4|, |l_1|)}{|D|}, \quad (13)$$

and candidates with $\kappa_D > 10^6$ are excluded. The resulting linear estimate must additionally satisfy $0 < \epsilon_{\text{lin}} < 1$; no clipping and no empirical lower emissivity bound are imposed.

Irradiance identifiability is evaluated independently of the emissivity value. With $\Delta L = L_2 - L_1$ and $\Delta l = l_2 - l_1$, the production condition is

$$|\Delta L| > |\Delta l|. \quad (14)$$

Under Eq. (14), the temporal change in measured upwelling irradiance must exceed the contemporaneous change in atmospheric downwelling irradiance. This criterion therefore tests for a resolvable surface-sensitive radiative contrast rather than merely reflecting atmospheric forcing.

3.4 Information-aware progressive temporal-pair selection

Candidate pairs satisfying Sect. 3.3 are evaluated progressively over temporal radii $r = 1, 2, \dots, 10$ min. At radius r , only identifiable candidates whose midpoint lies within r min of t_m and whose temporal separation does not exceed r are included.

For candidate j , the uncertainty of ϵ_{lin} is propagated from the four irradiances using

$$\sigma_{\epsilon,j}^2 = \mathbf{g}_j^T \mathbf{C}_y \mathbf{g}_j, \quad (15)$$

where $\mathbf{g}_j = \nabla_{\mathbf{y}} \epsilon_{\text{lin}}$ for $\mathbf{y}_j = (L_1, l_1, L_2, l_2)^T$. The covariance matrix is

$$\mathbf{C}_y = \sigma_L^2 \begin{pmatrix} 1 & \rho_{\text{cross}} & \rho_L & \rho_{\text{cross}} \\ \rho_{\text{cross}} & 1 & \rho_{\text{cross}} & \rho_l \\ \rho_L & \rho_{\text{cross}} & 1 & \rho_{\text{cross}} \\ \rho_{\text{cross}} & \rho_l & \rho_{\text{cross}} & 1 \end{pmatrix}, \quad (16)$$

Here, σ_L denotes the standard uncertainty assigned to each longwave irradiance measurement; ρ_L and ρ_l denote the temporal correlations in the upwelling and downwelling measurements, respectively; and ρ_{cross} denotes the assumed cross-channel correlation. Candidate uncertainty contributes the baseline weight

$$w_{D,j} = \frac{1}{\sigma_{\epsilon,j}^2 + \sigma_{\text{floor}}^2}, \quad \sigma_{\text{floor}} = 10^{-4}. \quad (17)$$

To account for the information that remains available for emissivity after allowing T_s to vary, a local two-parameter Fisher matrix is calculated. With the candidate sensitivity matrix $\mathbf{K}_j$ and the temporal upwelling covariance $\mathbf{C}_F$, the Fisher matrix is

$$\mathbf{F}_j = \mathbf{K}_j^T \mathbf{C}_F^{-1} \mathbf{K}_j. \quad (18)$$

Marginalizing over temperature leaves the emissivity information

$$I_{\epsilon|T,j} = F_{\epsilon\epsilon,j} - \frac{F_{\epsilon T,j}^2}{F_{TT,j}}. \quad (19)$$

The scene-wise median then normalizes the valid information values:

$$R_{F,j} = \frac{I_{\epsilon|T,j}}{\text{median}(I_{\epsilon|T})}, \quad w_j = w_{D,j} R_{F,j}. \quad (20)$$

If marginalized Fisher information is unavailable or non-positive, $R_{F,j} = 1$ and the uncertainty-based weight is retained. No candidate is rejected solely because of its Fisher information and no fixed Fisher-information cutoff is introduced.

At each radius, the representative local emissivity state is the weighted median of the candidate ϵ_{lin} population. Because candidate pairs overlap in time, information sufficiency is evaluated from both the number of distinct source observations and an effective pair count. A stage is information sufficient when

$$N_{\text{distinct}} \geq 4, \quad N_{\text{eff}} \geq 2.5. \quad (21)$$

Stability additionally requires the robust emissivity centre to satisfy

$$|\tilde{\epsilon}_r - \tilde{\epsilon}_{r-1}| \leq 0.005, \quad (22)$$

for two consecutive stage transitions. Once an information-sufficient stable stage is reached, the production pair is selected independently of its emissivity value, ranked first by midpoint proximity to the validation midpoint, then by preference for a separation of at least 2 min when otherwise tied, then by longer informative span, and finally by candidate uncertainty. If stability is not demonstrated, a prescribed fallback uses the best information-sufficient stage or, where necessary, the complete identifiable population within the 10 min search geometry.

3.5 Selected-pair thermal correction and physical-state reciprocal emissivity treatment

Once the production pair is fixed, the robust population estimate $\tilde{\epsilon}$ is retained as the production anchor. Two state-ordered linear emissivities are evaluated from the same selected pair, and the selected-pair thermal coefficient is

$$K_T = \frac{\Delta L}{\Delta L - \Delta l}. \quad (23)$$

Applying this coefficient gives the thermal correction and corrected anchor:

$$\Delta \epsilon_T = K_T (\epsilon_{L2} - \epsilon_{L1}), \quad \epsilon_T = \tilde{\epsilon} + \Delta \epsilon_T. \quad (24)$$

The correction is applied only when the ordered emissivities, denominator, and corrected result are finite and physically admissible; otherwise $\epsilon_T = \bar{\epsilon}$.

The local radiative environment of the selected pair is then characterized using the pair-mean irradiances $\bar{L}$ and $\bar{l}$, the surface-sensitive contrast $\bar{A} = \bar{L} - \bar{l}$, and

$$\bar{S}_T = \bar{l} + \frac{\bar{A}}{\epsilon_T}, \quad R_M = \frac{(1 - \epsilon_T)\bar{A}}{\epsilon_T}. \quad (25)$$

For valid surrounding observations in the same production neighbourhood, define the normalized state by

$$x = \frac{S - \bar{S}_T}{R_M}, \quad S = l + \frac{L - l}{\epsilon_T}. \quad (26)$$

A second characterization follows from Stefan–Boltzmann scaling. With $T_T = (\bar{S}_T/\sigma)^{1/4}$, $T_{\text{sky}} = (\bar{l}/\sigma)^{1/4}$, $t = T_{\text{sky}}/T_T$, and $G = \bar{l}/\bar{A}$, define

$$\Psi_{\text{SB}} = 2G - (1 + t + t^2 + t^3), \quad Z = \epsilon_T(1 + t + t^2 + t^3). \quad (27)$$

The framework then considers two exact reciprocal transformations:

$$\mathcal{R}_1(\epsilon_T) = \frac{2\epsilon_T}{1 + \epsilon_T}, \quad \mathcal{R}_2(\epsilon_T) = \frac{\epsilon_T}{2 - \epsilon_T}. \quad (28)$$

A conservative backbone selector determines the default mapping: $\mathcal{R}_1$ is selected for $\epsilon_T \leq 0.963$ and $\mathcal{R}_2$ for $\epsilon_T \geq 0.976$. Within $0.963 < \epsilon_T < 0.976$, $\mathcal{R}_1$ is selected when the outward normalized upwelling contribution is negative or the preceding normalized state satisfies $x_{\text{before}} < -2/5$; otherwise $\mathcal{R}_2$ is retained.

The backbone decision is overridden only for two specific physical-state configurations. In the Stefan–Boltzmann-negative configuration, the mapping is governed by $r_T = (1 - \epsilon_T)/\epsilon_T$ relative to $r_{\text{SB}} = (1 - 0.96246)/0.96246$. In the reciprocal-conductance configuration, the physical substitution is applicable only when

$$x_{\text{before}} < -\frac{1}{6}, \quad (29)$$

The corresponding branch criterion is

$$Z < 2 + \sqrt{2}. \quad (30)$$

Outside these physically applicable configurations, the backbone mapping is retained. The selected reciprocal result becomes the initial emissivity ϵ_0 for nonlinear refinement; if a transformed result is non-finite or falls outside $0 < \epsilon_0 < 1$, the conservative fallback $\epsilon_0 = \epsilon_T$ is used.

3.6 Newton–Raphson emissivity refinement and surface-temperature retrieval

Starting from $\epsilon^{(0)} = \epsilon_0$, define $P_1 = L_1^{1/4}$ and $P_2 = L_2^{1/4}$. At Newton iteration k ,

$$M(\epsilon^{(k)}) = \frac{1}{2} \left(P_1 + P_2 - \left(\frac{L_1 - l_1}{\epsilon^{(k)}} + l_1 \right)^{1/4} - \left(\frac{L_2 - l_2}{\epsilon^{(k)}} + l_2 \right)^{1/4} \right). \quad (31)$$

The nonlinear residual takes the form

$$f(\epsilon) = \epsilon((P_1 - M(\epsilon))^4 - l_1) - (L_1 - l_1). \quad (32)$$

Equation (32) retains the established production sign convention; the correction enters through $P_1 - M(\epsilon)$ and is not replaced by the alternative positive-sign form. The derivative is

$$f'(\epsilon) = (P_1 - M)^4 - l_1 - 4\epsilon(P_1 - M)^3 M'(\epsilon), \quad (33)$$

where

$$M'(\epsilon) = -\frac{1}{8} \left(\left(\frac{L_1 - l_1}{\epsilon} + l_1 \right)^{-3/4} \frac{L_1 - l_1}{\epsilon^2} + \left(\frac{L_2 - l_2}{\epsilon} + l_2 \right)^{-3/4} \frac{L_2 - l_2}{\epsilon^2} \right). \quad (34)$$

Newton–Raphson then advances the emissivity according to

$$\epsilon^{(k+1)} = \epsilon^{(k)} - \frac{f(\epsilon^{(k)})}{f'(\epsilon^{(k)})}. \quad (35)$$

A fixed sequence of up to eight updates is applied. The sequence is interrupted if the derivative becomes non-finite or zero, or if the next emissivity is non-finite or lies outside $0 < \epsilon^{(k+1)} < 1$. The eight-update structure provides a common numerical trajectory and is not described as an empirical stopping tolerance. The Newton–Raphson formulation follows standard nonlinear root-finding practice (Ortega and Rheinboldt, 2000; Nosedal and Wright, 2006).

After the final valid update, compute the pair-mean apparent temperature as

$$T'_s = \frac{P_1 + P_2}{2\sigma^{1/4}}, \quad (36)$$

The associated emissivity-induced temperature correction becomes

$$\Delta T = \frac{P_1 + P_2 - \left(\frac{L_1 - l_1}{\epsilon} + l_1 \right)^{1/4} - \left(\frac{L_2 - l_2}{\epsilon} + l_2 \right)^{1/4}}{2\sigma^{1/4}}. \quad (37)$$

Finally, the retrieved surface temperature follows the established convention

$$T_s = T'_s - \Delta T. \quad (38)$$

3.7 Conditioning diagnostics and uncertainty propagation

The final nonlinear residual is reported as $R_f = |f(\epsilon)|$, and the minimum local sensitivity encountered over the Newton trajectory is

$$J_{\min} = \min_k |f'(\epsilon^{(k)})|. \quad (39)$$

Using the full covariance structure gives the production emissivity uncertainty

$$\sigma_{\epsilon, \text{corr}} = \frac{\sqrt{\mathbf{a}^T \mathbf{C}_y \mathbf{a}}}{J_{\min}}, \quad \mathbf{a} = (1, 0, 1, 0)^T. \quad (40)$$

An independent-error limiting form is retained only as a diagnostic. Differentiating Eq. (38) with respect to emissivity gives the local temperature sensitivity $\partial T_s / \partial \epsilon$, from which the emissivity-driven contribution is

$$\sigma_{T_s, \epsilon} = \left| \frac{\partial T_s}{\partial \epsilon} \right| \sigma_{\epsilon, \text{corr}}. \quad (41)$$

The correlated irradiance-driven contribution follows from

$$\sigma_{T_s, L} = \left| \frac{L_1^{-3/4} + L_2^{-3/4}}{8\sigma^{1/4}} \right| \sigma_L \sqrt{(1 \quad 1) \begin{pmatrix} 1 & \rho_L \\ \rho_L & 1 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix}}, \quad (42)$$

Combining the irradiance- and emissivity-driven terms gives the total propagated temperature uncertainty

$$\sigma_{T_s, \text{corr}} = \sqrt{\sigma_{T_s, L}^2 + \sigma_{T_s, \epsilon}^2}. \quad (43)$$

These quantities are propagated retrieval uncertainties and are not assumed to be calibrated confidence intervals; empirical coverage is evaluated separately using the known-truth experiment. The interpretation of local sensitivity and uncertainty follows standard inverse-problem reasoning (Taranola, 2005; Menke, 2012; Aster et al., 2019).

4 Experimental setup and validation framework

Two complementary validation strategies are used (Fig. 2). The field experiment uses ground-based longwave observations from nine SURFRAD and BSRN stations spanning vegetated, mixed, and bare/desert environments. The second is a physics-based known-truth experiment in which emissivity and surface temperature are prescribed independently of the retrieval and broadband longwave observations are generated from exact forward physics. Both validation arms are processed with the same pre-specified conditioning-controlled retrieval formulation.

4.1 Ground-based field-validation dataset

The field population combines observations from seven SURFRAD stations and two BSRN stations: Bondville (BON), Desert Rock (DRA), Fort Peck (FPK), Goodwin Creek (GWN), Penn State (PSU), Sioux Falls (SXF), Table Mountain (TBL), Cabauw (CAB), and Gobabeb (GOB). Their composition is summarized in Table 1.

A total of 637 source datasets entered the workflow audit. Five datasets (one BON case, one CAB case, and three GOB cases) did not contain the required finite initial irradiances and were excluded, leaving 632 usable field datasets. All 632 usable cases produced final finite production retrievals.

4.2 Ground-based surface-temperature validation reference

The field comparison uses a ground-based validation-reference temperature $T_{s, \text{ref}}$. It is derived from station-calibrated ground-based longwave thermal-radiometer observations combined with externally specified broadband emissivity,

$$T_{s, \text{ref}} = \left(\frac{L_{\text{ref}} - (1 - \epsilon_{\text{ref}}) L_{\text{ref}}}{\epsilon_{\text{ref}} \sigma} \right)^{1/4}. \quad (44)$$

The field reference and retrieval are therefore not statistically independent in their irradiance measurements: both ultimately involve ground-observed longwave irradiance. The distinction is that ϵ_{ref} is externally specified and is not supplied to the proposed emissivity retrieval. The field arm consequently validates combined emissivity-temperature retrieval behaviour under real observations rather than providing independent broadband-emissivity truth.

For the SURFRAD reference records, externally supplied broadband emissivity information was used to derive the validation-reference temperatures. For Cabauw and Gobabeb, the validation-reference temperature was constructed from the nearest 1 min upwelling and downwelling longwave observations at the Landsat acquisition time together with the corresponding Landsat-8 3×3 mean emissivity.

The externally specified emissivities used only to construct the field temperature reference are themselves subject to spectral, spatial, temporal, and representativeness uncertainty. The field comparison is therefore interpreted as a realistic benchmark of the combined emissivity-temperature retrieval rather than as independent emissivity truth; direct broadband-emissivity accuracy is evaluated only in the prescribed-truth experiment.

4.3 Field performance metrics and workflow accounting

For each field retrieval, the signed temperature error is $e_{T,i} = T_{s,i}^{\text{ret}} - T_{s,i}^{\text{ref}}$. Performance is summarized by bias, MAE, RMSE, standard deviation of signed error, and squared Pear-




Figure 2. Field and physics-based known-truth validation architecture. The field arm retains 632 usable datasets from 637 source datasets; the synthetic arm contains 6000 prescribed truth scenes evaluated as noiseless (G1) and production-noise (G2) observations. Both arms use the same conditioning-controlled retrieval framework.

son correlation coefficient R^2 . A separate read-only audit tracks candidate geometry, finite irradiance, quasi-steady screening, numerical/physical admissibility, irradiance identifiability, information sufficiency, production-pair selection, Newton progression, and final retrieval. The audit does not impose additional rejection criteria and does not modify the

production retrieval. Detailed scene- and pair-level accounting is given in Appendix A.

Across the 632 usable field scenes, the audit contains 88 154 geometry-valid candidate pairs, of which 88 135 have finite irradiances, 81 682 satisfy the quasi-steady and denominator-validity stages, 68 525 retain a physically ad-

Table 1. Validation datasets and station composition. The initial-irradiance exclusion refers only to the prescribed validation-pair availability check.

Network	Station	Code	Source datasets	Excluded at initial irradiance check	Final retrievals	Retention (%)
SURFRAD	Bondville	BON	70	1	69	98.57
BSRN	Cabauw	CAB	48	1	47	97.92
SURFRAD	Desert Rock	DRA	50	0	50	100.00
SURFRAD	Fort Peck	FPK	75	0	75	100.00
BSRN	Gobabeb	GOB	194	3	191	98.45
SURFRAD	Goodwin Creek	GWN	50	0	50	100.00
SURFRAD	Penn State	PSU	26	0	26	100.00
SURFRAD	Sioux Falls	SXF	40	0	40	100.00
SURFRAD	Table Mountain	TBL	84	0	84	100.00
TOTAL	All stations	ALL	637	5	632	99.22

missible linear-emissivity state, and 63 190 satisfy the irradiance-identifiability condition. At scene level, all 632 usable datasets contain at least one admissible and identifiable pair, all 632 enter Newton–Raphson refinement, and all 632 complete eight valid Newton updates. Appendix A provides the corresponding station-wise and stage-by-stage accounting.

4.4 Physics-based known-truth synthetic validation

The synthetic validation contains $N_{\text{syn}} = 6000$ independent truth scenes. A fixed random seed, `rng(858)`, is used for reproducibility. Exactly 2000 scenes are generated in each of three emissivity regimes: $0.85 \leq \epsilon_{\text{true}} < 0.92$, $0.92 \leq \epsilon_{\text{true}} < 0.97$, and $0.97 \leq \epsilon_{\text{true}} \leq 0.99$. The initial true surface temperature is prescribed independently over 270–330 K. Synthetic temporal forcing is empirically grounded by bootstrap sampling observable temporal-driver states derived from the 632 field cases.

For each time sample, the true surface emission is $S(t) = \sigma T_s^{\text{true}}(t)^4$, and exact upwelling irradiance is generated as

$$L(t) = \epsilon_{\text{true}} S(t) + (1 - \epsilon_{\text{true}}) l(t).$$
 (45)

Thus, ϵ_{true} enters only through the forward model and is never derived from, fitted to, or conditioned on the retrieval result.

4.5 Noiseless and production-noise synthetic experiments

Each truth scene is used for two paired experiments. G1 contains the exact physically generated longwave observations with no additional measurement perturbation and therefore isolates retrieval/inversion behaviour. G2 uses the same 6000 truth scenes but adds a correlated production-noise realization drawn from the covariance structure defined in Eq. (16) before unchanged Savitzky–Golay processing. Apart from

this perturbation, G1 and G2 use identical truth states and the same retrieval implementation.

All 6000 generated scenes remain in the denominator used to compute retrieval success:

$$f_{\text{success}} = \frac{N_{\text{successful}}}{6000}.$$
 (46)

Bias, MAE, and RMSE require a finite retrieved value and are therefore calculated over successful retrievals only. For T_s , error is evaluated against the true temperature corresponding to the selected production pair.

5 Results and discussion

The results follow the physical and numerical sequence through which retrieval reliability is established. Temporal-pair selection and stability are examined first, followed by reciprocal-state applicability, Newton behaviour, the contrast-sensitivity-uncertainty relationship, field emissivity behaviour, field surface-temperature validation, and physics-based known-truth validation.

5.1 Temporal-pair selection and stability

Figure 3 characterizes the progressive selector over the complete 632-case field population. Of the 632 retrievals, 628 were selected through the stable-centre pathway and only four required fallback selection. The median selected production-pair separation was 3.0 min. These results show that, for the overwhelming majority of cases, the progressive search identified a locally stable emissivity population without exhausting the prescribed fallback pathway.

The relationship in Fig. 3d is used as an empirical check for systematic displacement of the emissivity solution with increasing pair separation. Within the retained population, no strong monotonic shift is evident. Quasi-steady screening is therefore interpreted as one component of the selection chain




Figure 3. Temporal-pair selection and progressive information sufficiency for the 632 field retrievals. **(a)** Selected production-pair separation. **(b)** Accepted search radius, distinguishing stable-centre and fallback pathways. **(c)** Progressive growth of information-sufficient and stable-stage populations with search radius. **(d)** Selected pair separation versus retrieved emissivity.

rather than a complete observability test; stable production selection also requires irradiance identifiability and sufficient independent information.

5.2 Physical-state applicability of the reciprocal emissivity formulations

Following selected-pair thermal correction, the retrieval evaluates whether the surrounding radiative state supports one of the exact reciprocal emissivity transformations. Figure 4 presents the two physical configurations for which the applicability boundaries have a direct radiative interpretation.

The Stefan–Boltzmann-controlled configuration occurs in 102 of the 632 field cases, while 58 cases occupy the reciprocal-conductance configuration. The remaining 472 cases exhibit other production-state signatures and are not represented in these two diagnostic panels. The reciprocal transformation is therefore state dependent rather than universally applied, and the branch choice is made before comparison with the field validation reference.

5.3 Newton numerical behaviour within the retained population

Figure 5 evaluates Newton–Raphson behaviour after candidate screening, production-pair selection, and physical-state conditioning. Figure 5a relates the final absolute residual to the minimum local emissivity sensitivity, while Fig. 5b shows the distribution of that sensitivity across the retained field population.

The figure characterizes the numerical state of accepted production retrievals and should not be interpreted as evidence that arbitrary temporal pairs are well conditioned. Weakly identifiable and numerically inadmissible candidates are addressed earlier in the workflow. Small residuals establish consistency with the nonlinear equation at the retrieved state, while the derivative magnitude quantifies local sensitivity; neither quantity alone is treated as a complete observability criterion.

5.4 Radiative contrast, inversion sensitivity, and propagated uncertainty

Figure 6 links the selected-pair upwelling contrast $|\Delta L|$ to local nonlinear sensitivity and propagated uncertainty. Across the 632 retrievals, the Spearman correlation between




Figure 4. Physical-state applicability of the reciprocal emissivity formulations. (a) Stefan–Boltzmann-controlled configuration expressed in the $r_T - \Psi_{SB}$ state space. (b) Reciprocal-conductance configuration expressed in the $x_{\text{before}} - Z$ state space. The plotted boundaries are the pre-specified physical applicability conditions used by the production selector.




Figure 5. Newton numerical behaviour within the retained 632-case field population. (a) Final absolute nonlinear residual versus minimum absolute emissivity sensitivity encountered during the fixed eight-update trajectory. (b) Distribution of the minimum absolute sensitivity.

$|\Delta L|$ and correlated emissivity uncertainty is -0.4026 , while the correlation between $|\Delta L|$ and minimum sensitivity is $+0.4026$. Minimum sensitivity and correlated emissivity uncertainty have a rank correlation of -1.0000 .

The first two relationships show that larger selected-pair upwelling contrast is generally associated with greater emissivity sensitivity and lower propagated emissivity uncertainty. The perfect rank anticorrelation between sensitivity and σ_ϵ follows directly from the mathematical structure of the local uncertainty formulation and is not treated as an independent validation statistic. Figure 6d extends the conditioning pathway to T_s through its irradiance- and emissivity-driven uncertainty components. These are propagated uncertainties; their empirical calibration is evaluated separately in Appendix C.

5.5 Field behaviour of retrieved broadband emissivity

Figure 7 shows retrieved broadband-emissivity distributions for the nine stations. All 632 production retrievals are included, with station medians and interquartile ranges. The distributions are interpreted descriptively rather than as field emissivity-validation errors because a common independent broadband-emissivity truth is not available for the complete field population.

5.6 Nine-station field surface-temperature validation

Figure 8 compares retrieved T_s with the ground-based validation-reference T_s for all 632 cases. Across the complete field population, the production retrieval has a bias of -0.162 K, MAE of 0.432 K, RMSE of 0.558 K, and $R^2 = 0.9984$.




Figure 6. Conditioning and propagated uncertainty. (a) Selected-pair upwelling contrast $|\Delta L|$ versus correlated emissivity uncertainty. (b) $|\Delta L|$ versus minimum nonlinear emissivity sensitivity. (c) Minimum sensitivity versus correlated emissivity uncertainty. (d) Irradiance-driven, emissivity-driven, and total correlated T_s uncertainty components.




Figure 7. Station-wise distributions of retrieved broadband surface emissivity for the 632 field retrievals. Individual retrievals are shown together with the station median and interquartile range. The distributions characterize field retrieval behaviour and are not interpreted as direct emissivity-validation errors.




Figure 8. Field surface-temperature validation for all 632 retrievals. **(a)** Retrieved T_s versus ground-based validation-reference T_s for the nine stations, with the 1 : 1 relation shown for reference. **(b)** Station-wise distributions of signed error $T_{s,\text{ret}} - T_{s,\text{ref}}$.

The low overall bias indicates little systematic displacement relative to the field reference, while the MAE and RMSE show that the typical retrieval-reference difference remains well below 1 K. Station-wise values demonstrate that the aggregate result is assembled from distinct site populations rather than repeated observations from one homogeneous environment. Station-wise validation statistics are summarized in Table 2. The reference temperature nevertheless shares ground-based longwave measurements with the retrieval and is therefore not treated as fully independent radiometric truth.

5.7 Physics-based known-truth validation

The principal known-truth test is the G1 noiseless experiment shown in Fig. 9 and summarized in Table 3. Of the 6000 independently prescribed truth scenes, 5924 yield successful finite retrievals, corresponding to 98.73 %. All 6000 generated scenes remain in the success denominator.

Across successful retrievals, overall emissivity bias, MAE, and RMSE are 0.03129, 0.04908, and 0.06578, respectively. The corresponding T_s bias, MAE, and RMSE are -0.335 , 1.307, and 3.213 K. Performance varies across emissivity regimes: the largest positive emissivity bias occurs in the lowest-emissivity interval, while the highest-emissivity interval has a smaller absolute bias and lower emissivity RMSE. The method therefore maintains high retrieval availability across $0.85 \leq \epsilon_{\text{true}} \leq 0.99$, but low-emissivity conditions remain the most demanding part of the tested retrieval space.

The paired G2 experiment applies the production-noise model to the same truth scenes. Detailed G1–G2 comparisons are retained in Appendix B because they quantify robustness to realistic measurement perturbations rather than constituting an additional independent validation population.

5.8 Integrated interpretation and limitations

The results support a conditioning pathway in which retrieval reliability arises from complementary controls rather than Newton behaviour alone. Temporal screening limits strong thermodynamic-state changes; irradiance identifiability requires surface-sensitive contrast; progressive population analysis requires sufficient independent information and centre stability; physical-state diagnostics determine reciprocal applicability; and Newton refinement solves the resulting nonlinear emissivity equation.

The field and synthetic experiments provide complementary evidence. The field observations retain real instrument behaviour, atmospheric variability, temporal variability, and surface heterogeneity, but the validation-reference T_s shares longwave radiometric information with the retrieval. The synthetic experiment provides independently prescribed emissivity and temperature truth but necessarily represents a controlled physical model of the observed environment. Confidence in the retrieval therefore derives from the consistency of the two lines of evidence rather than from either experiment alone.

Two further limitations are important. First, the conditioning and Newton diagnostics describe the screened production population and should not be extrapolated to arbitrary longwave temporal pairs. Second, propagated uncertainty quantifies local measurement and emissivity sensitivity but is not automatically a calibrated confidence interval. Known-truth uncertainty coverage is therefore evaluated separately in Appendix C. The synthetic validation was evaluated without subsequent retuning of production thresholds, pair-selection rules, reciprocal-state logic, or uncertainty parameters using the synthetic outcomes.

Table 2. Field surface-temperature validation performance by station and overall.

Station	Code	<i>N</i>	Bias (K)	SD (K)	MAE (K)	RMSE (K)	<i>R</i> ²
Bondville	BON	69	−0.207	0.326	0.321	0.385	0.9992
Cabauw	CAB	47	0.284	0.445	0.397	0.524	0.9984
Desert Rock	DRA	50	0.020	0.664	0.569	0.658	0.9977
Fort Peck	FPK	75	−0.099	0.465	0.396	0.472	0.9993
Gobabeb	GOB	191	−0.238	0.481	0.443	0.536	0.9944
Goodwin Creek	GWN	50	−0.332	0.680	0.464	0.751	0.9942
Penn State	PSU	26	−0.289	0.410	0.372	0.495	0.9985
Sioux Falls	SXF	40	−0.275	0.505	0.465	0.570	0.9988
Table Mountain	TBL	84	−0.174	0.604	0.449	0.625	0.9979
All stations	ALL	632	−0.162	0.534	0.432	0.558	0.9984




Figure 9. Physics-based known-truth validation for the G1 noiseless experiment. (a) Retrieved versus prescribed broadband emissivity for successful retrievals. (b) Retrieved T_s versus the true surface temperature corresponding to the selected production pair. The success denominator is all 6000 generated scenes; plotted error statistics use successful retrievals only.

Table 3. Physics-based known-truth validation for the G1 noiseless experiment, reported overall and by predefined emissivity regime.

True ϵ range	Generated scenes	Successful retrievals	Success (%)	ϵ bias	ϵ MAE	ϵ RMSE	T_s bias (K)	T_s MAE (K)	T_s (K)
Overall	6000	5924	98.73	0.03129	0.04908	0.06578	−0.335	1.307	3.213
0.85–0.92	2000	1967	98.35	0.08581	0.08778	0.09352	−1.754	2.043	2.522
0.92–0.97	2000	1976	98.80	0.02213	0.03575	0.05043	−0.058	1.001	4.337
0.97–0.99	2000	1981	99.05	−0.01370	0.02395	0.04148	0.797	0.882	2.407

6 Conclusions

This study presents a conditioning-controlled framework for retrieving broadband surface emissivity and land surface temperature directly from paired ground-based upwelling and downwelling longwave irradiance measurements. The inverse problem is treated as an information-limited radiometric retrieval in which thermal-state consistency, numerical admissibility, irradiance identifiability, information sufficiency, physical-state applicability, nonlinear refinement, and uncertainty propagation are incorporated explicitly into the retrieval workflow. The framework therefore addresses observability before nonlinear inversion rather than relying on convergence behaviour alone as evidence of retrieval reliability.

A population-based progressive temporal search identifies production pairs from the available local observations. Candidate pairs are screened for quasi-steady apparent surface temperature, numerical and physical admissibility, and resolvable surface-sensitive irradiance contrast. Their local information content is then evaluated using propagated candidate uncertainty together with marginalized emissivity information, while the final production pair is selected without ranking candidates according to their emissivity value. Among the 632 usable field datasets, 628 satisfied the stable-centre selection criterion and only four required the prescribed fallback pathway; the median selected-pair separation was 3 min.

The framework also incorporates an explicit physical-state treatment of the reciprocal emissivity formulations. Reciprocal transformation is applied conditionally according to the radiative state surrounding the selected observation pair rather than as a universal algebraic correction. The resulting state-conditioned emissivity provides the initial state for the fixed eight-update Newton–Raphson refinement. Numerical residuals and local sensitivities are consequently interpreted within a population that has already passed the preceding physical and information-based conditioning stages.

Field validation comprised 637 source datasets from nine SURFRAD and BSRN stations spanning vegetated, mixed, and bare/desert environments. Five datasets lacked the required finite initial validation-pair irradiances, leaving 632 usable cases, all of which produced final finite retrievals. Across this population, retrieved T_s had an overall bias of -0.162 K, MAE of 0.432 K, RMSE of 0.558 K, and $R^2 = 0.9984$ relative to the ground-based validation-reference temperature.

Direct broadband-emissivity validation was provided by the complementary physics-based synthetic experiment. Six thousand truth scenes were generated, with exactly 2000 scenes in each of the predefined emissivity regimes $0.85 \leq \epsilon < 0.92$, $0.92 \leq \epsilon < 0.97$, and $0.97 \leq \epsilon \leq 0.99$. In the noiseless known-truth experiment, 5924 of 6000 scenes were successfully retrieved (98.73 %). Across successful retrievals, the overall emissivity bias, MAE, and RMSE were 0.0313 , 0.0491 , and 0.0658 , respectively, while the corresponding T_s bias, MAE, and RMSE were -0.335 , 1.307 , and 3.213 K. Retrieval performance varied across emissivity regimes, with the lowest-emissivity interval showing the largest positive emissivity bias and the highest emissivity RMSE.

The field and synthetic experiments provide complementary evidence. The field observations contain real atmospheric variability, instrument behaviour, temporal variability, and surface heterogeneity, but the validation-reference temperature shares ground-based longwave radiometric information with the retrieval. The synthetic experiment provides independently prescribed emissivity and temperature truth but necessarily represents a controlled forward model of the observed environment. Their combined interpretation therefore provides stronger evidence than either experiment alone.

Several limitations remain. The field temperature reference is not independent of all longwave measurements used by the retrieval. The synthetic experiment reproduces empirically grounded short-term radiative forcing but cannot represent every unmodelled process or representativeness error present at real field sites. Retrieval errors also increase in portions of the lower-emissivity truth space, indicating that weak-information conditions remain challenging. Further work should focus on validation against independent broadband emissivity and surface-temperature references, extension to additional low-emissivity surface types and radiative regimes, and continued evaluation of uncertainty calibration under weak-information conditions.

Overall, the main contribution is an integrated retrieval architecture in which observability is assessed before inversion, physical applicability is treated explicitly, numerical refinement is diagnosed rather than assumed, and validation combines real-field performance with independently prescribed known truth. The results demonstrate that broadband emissivity and land surface temperature can be jointly retrieved from high-temporal-resolution paired longwave irradiance measurements without requiring externally prescribed emissivity as an input to the retrieval itself.

Appendix A: Workflow-stage accounting

Appendix A provides the detailed field sample and candidate-pair accounting underlying the compact audit shown in Fig. 2. The accounting is read-only and does not alter the retrieval.




Figure A1. Workflow-stage accounting for the field dataset. (a) Scene-level progression from source datasets to final retrievals. (b) Candidate-pair retention through the physical, numerical, and irradiance-identifiability screening stages.

Table A1. Detailed overall field workflow-stage accounting.

Audit stage	Count
Source files seen	637
Files with validation-time constraint	637
Files with valid data record/layout	637
Files with finite initial validation-pair irradiances	632
Scenes with ≥ 1 admissible linear-epsilon pair	632
Scenes with ≥ 1 irradiance-identifiable pair	632
Scenes selected by stable-centre criterion	628
Scenes selected through fallback	4
Scenes with finite production pair	632
Scenes entering Newton–Raphson	632
Scenes with ≥ 1 valid Newton update	632
Scenes completing 8 valid Newton updates	632
Scenes with invalid Newton break	0
Final finite retrievals	632
PAIR COUNT: geometry-valid	88 154
PAIR COUNT: finite-irradiance	88 135
PAIR COUNT: quasi-steady	81 682
PAIR COUNT: denominator-valid	81 682
PAIR COUNT: conditioned	81 651
PAIR COUNT: finite linear epsilon	81 651
PAIR COUNT: $0 < \text{linear epsilon} < 1$	68 525
PAIR COUNT: irradiance-identifiable	63 190

Table A2. Station-wise field workflow accounting.

Site	Source	Initial finite	Admissible	Identifiable	Stable	Fallback	Prod. pair	8 updates	Final
BON	70	69	69	69	69	0	69	69	69
CAB	48	47	47	47	47	0	47	47	47
DRA	50	50	50	50	50	0	50	50	50
FPK	75	75	75	75	75	0	75	75	75
GOB	194	191	191	191	190	1	191	191	191
GWN	50	50	50	50	49	1	50	50	50
PSU	26	26	26	26	26	0	26	26	26
SXF	40	40	40	40	39	1	40	40	40
TBL	84	84	84	84	83	1	84	84	84
ALL	637	632	632	632	628	4	632	632	632

Appendix B: Synthetic realism and production-noise robustness




Figure B1. Field-versus-synthetic selected-pair geometry. Field quantiles are compared with G1 noiseless and G2 production-noise synthetic quantiles for selected-pair separation and the principal longwave contrast variables.




Figure B2. Paired G1-versus-G2 production-noise robustness for the same prescribed truth scenes. The panels compare the change in emissivity error and the change in error when the pre-specified production-noise model is added before preprocessing.

Table B1. Synthetic performance for G1 noiseless and G2 production-noise experiments, reported overall and by predefined emissivity regime.

Experiment	True ϵ range	Generated	Successful	Success (%)	ϵ bias	ϵ MAE	ϵ RMSE	T_s bias (K)	T_s MAE (K)	T_s RMSE (K)
G1 noiseless	Overall	6000	5924	98.73	0.03129	0.04908	0.06578	−0.335	1.307	3.213
G1 noiseless	0.85–0.92	2000	1967	98.35	0.08581	0.08778	0.09352	−1.754	2.043	2.522
G1 noiseless	0.92–0.97	2000	1976	98.80	0.02213	0.03575	0.05043	−0.058	1.001	4.337
G1 noiseless	0.97–0.99	2000	1981	99.05	−0.01370	0.02395	0.04148	0.797	0.882	2.407
G2 production noise	Overall	6000	5939	98.98	0.03181	0.04885	0.06568	−0.376	1.311	2.712
G2 production noise	0.85–0.92	2000	1970	98.50	0.08604	0.08810	0.09380	−1.763	2.066	2.621
G2 production noise	0.92–0.97	2000	1981	99.05	0.02253	0.03573	0.05025	−0.143	0.953	2.744
G2 production noise	0.97–0.99	2000	1988	99.40	−0.01269	0.02304	0.04067	0.766	0.920	2.770

Appendix C: Detailed emissivity-range and uncertainty diagnostics




Figure C1. Detailed synthetic response across 0.01 emissivity bins, showing retrieval success, emissivity RMSE, and T_s RMSE for G1 noiseless and G2 production-noise experiments.

Table C1. Synthetic uncertainty calibration and nominal 1-sigma coverage for emissivity and T_s .

Experiment	True ϵ range	$N\epsilon$	ϵ within 1σ (%)	Median $ z\epsilon $	NT_s	T_s within 1σ (%)	Median $ zT_s $
G1 noiseless	Overall	5924	25.9	3.119	5924	40.2	1.472
G1 noiseless	0.85–0.92	1967	4.5	7.138	1967	10.9	3.533
G1 noiseless	0.92–0.97	1976	22.1	2.270	1976	46.9	1.090
G1 noiseless	0.97–0.99	1981	51.0	0.949	1981	62.7	0.500
G2 production noise	Overall	5939	25.9	3.089	5939	38.7	1.517
G2 production noise	0.85–0.92	1970	4.7	7.186	1970	11.7	3.615
G2 production noise	0.92–0.97	1981	22.2	2.323	1981	45.0	1.130
G2 production noise	0.97–0.99	1988	50.5	0.968	1988	59.3	0.760

Code and data availability. The public-source observations used in this study are obtained from SURFRAD and BSRN data services. The derived validation datasets and analysis/retrieval code supporting this study are archived in the frozen Version 1.0 reproducibility record on Zenodo at <https://doi.org/10.5281/zenodo.22684633> (Mito, 2026). Raw third-party SURFRAD and BSRN observations are not redistributed in the archive and remain available from their respective public data services.

Competing interests. The author has declared that there are no competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. The authors bear the ultimate responsibility for providing appropriate place names. Views expressed in the text are those of the authors and do not necessarily reflect the views of the publisher.

Acknowledgements. The author gratefully acknowledges the organizations and data providers that maintain the SURFRAD and BSRN radiation records used in this study, and Mohannad Loho of Damascus University for providing part of the SURFRAD validation data.

Review statement. This paper was edited by Anthony Bucholtz and reviewed by Prasanjit Dash and one anonymous referee.

References

- Aster, R. C., Borchers, B., and Thurber, C. H.: Parameter Estimation and Inverse Problems, 3rd edn., Elsevier Academic Press, Cambridge, MA, USA, <https://doi.org/10.1016/C2015-0-02458-3>, 2019.
- Augustine, J. A., DeLuisi, J. J., and Long, C. N.: SURFRAD – A national surface radiation budget network for atmospheric research, B. Am. Meteorol. Soc., 81, 2341–2357, [https://doi.org/10.1175/1520-0477\(2000\)081<2341:SANSRB>2.3.CO;2](https://doi.org/10.1175/1520-0477(2000)081<2341:SANSRB>2.3.CO;2), 2000.
- Augustine, J. A., Hodges, G. B., Cornwall, C. R., Michalsky, J. J., and Medina, C. I.: An update on SURFRAD – The GCOS surface radiation budget network for the continental United States, J. Atmos. Ocean. Tech., 22, 1460–1472, <https://doi.org/10.1175/JTECH1806.1>, 2005.
- Dash, P., Göttsche, F.-M., Olesen, F.-S., and Fischer, H.: Land surface temperature and emissivity estimation from passive sensor data: Theory and practice; current trends, Int. J. Remote Sens., 23, 2563–2594, <https://doi.org/10.1080/01431160110115041>, 2002.
- Gillespie, A. R., Rokugawa, S., Matsunaga, T., Cothorn, J. S., Hook, S. J., and Kahle, A. B.: A temperature and emissivity separation algorithm for Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) images, IEEE T. Geosci. Remote, 36, 1113–1126, <https://doi.org/10.1109/36.700995>, 1998.
- Hulley, G. C. and Hook, S. J.: Generating consistent land surface temperature and emissivity products between ASTER and MODIS data for Earth science research, IEEE T. Geosci. Remote, 49, 1304–1315, <https://doi.org/10.1109/TGRS.2010.2063034>, 2011.
- Li, S., Yu, Y., Sun, D., Tarpley, D., Zhan, X., and Chiu, L.: Evaluation of 10-year AQUA/MODIS land surface temperature with SURFRAD observations, Int. J. Remote Sens., 35, 830–856, <https://doi.org/10.1080/01431161.2013.873149>, 2014.
- Li, Z.-L., Wu, H., Wang, N., Qiu, S., Sobrino, J. A., Wan, Z., Tang, B.-H., and Yan, G.: Land surface emissivity retrieval from satellite data, Int. J. Remote Sens., 34, 3084–3127, <https://doi.org/10.1080/01431161.2012.716540>, 2013a.
- Li, Z.-L., Tang, B.-H., Wu, H., Ren, H., Yan, G., Wan, Z., Trigo, I. F., and Sobrino, J. A.: Satellite-derived land surface temperature: Current status and perspectives, Remote Sens. Environ., 131, 14–37, <https://doi.org/10.1016/j.rse.2012.12.008>, 2013b.

- Li, Z.-L., Wu, H., Duan, S.-B., Zhao, W., Ren, H., Liu, X., Leng, P., Tang, R., Ye, X., Zhu, J., Sun, Y., Si, M., Liu, M., Li, J., Zhang, X., Shang, G., Tang, B.-H., Yan, G., and Zhou, C.: Satellite remote sensing of global land surface temperature: Definition, methods, products, and applications, *Rev. Geophys.*, 61, e2022RG000777, <https://doi.org/10.1029/2022RG000777>, 2023.
- Menke, W.: *Geophysical Data Analysis: Discrete Inverse Theory*, 3rd edn., Academic Press, Burlington, MA, USA, <https://doi.org/10.1016/C2011-0-69765-0>, 2012.
- Mito, C.: Reproducibility archive for conditioning-controlled retrieval of broadband land surface temperature and emissivity from paired ground-based longwave irradiance measurements, Zenodo [code], <https://doi.org/10.5281/zenodo.22684633>, 2026.
- Nocedal, J. and Wright, S. J.: *Numerical Optimization*, Springer, New York, NY, USA, <https://doi.org/10.1007/978-0-387-40065-5>, 2006.
- Ortega, J. M. and Rheinboldt, W. C.: *Iterative Solution of Nonlinear Equations in Several Variables*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, <https://doi.org/10.1137/1.9780898719468>, 2000.
- Rodgers, C. D.: *Inverse Methods for Atmospheric Sound-ing: Theory and Practice*, World Scientific, Singapore, <https://doi.org/10.1142/3171>, 2000.
- Tarantola, A.: *Inverse Problem Theory and Methods for Model Parameter Estimation*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, <https://doi.org/10.1137/1.9780898717921>, 2005.
- Wan, Z.: New refinements and validation of the collection-6 MODIS land-surface temperature/emissivity product, *Remote Sens. Environ.*, 140, 36–45, <https://doi.org/10.1016/j.rse.2013.08.027>, 2014.
- Wang, M., Zhang, Z., Hu, T., and Liu, X.: A practical single-channel algorithm for land surface temperature retrieval: Application to Landsat series data, *J. Geophys. Res.-Atmos.*, 124, 299–316, <https://doi.org/10.1029/2018JD029330>, 2019.
- Zhang, Z., He, G., Wang, M., Long, T., Wang, G., and Zhang, X.: Validation of the generalized single-channel algorithm using Landsat 8 imagery and SURFRAD ground measurements, *Remote Sens. Lett.*, 7, 810–816, <https://doi.org/10.1080/2150704X.2016.1190475>, 2016.