Response to Reviewer #1

Review for "The JUICE 2024 close flyby of the Moon: Thermal assessment from MAJIS":

Overall, the manuscript is well written considering all aspects of the study. While comparing the results, the authors have also highlighted the limitations of each algorithm. However, I have a few concerns that need to be addressed before its acceptance for publication in ANGEO.

We thank the Reviewer for the careful and constructive review of our manuscript and for the detailed technical suggestions. We appreciate the Reviewer's emphasis on improving methodological clarity and on strengthening the quantitative assessment of the retrieval robustness. Following these recommendations, we have revised the manuscript to (i) better document key elements of the Bayesian inversion framework (including the role of the a priori and the definition of covariance matrices), (ii) clarify the treatment of topography and the limits of the ellipsoid-based geometry, (iii) justify the comparison between MAJIS-derived physical temperatures and Diviner brightness-temperature products, and (iv) improve the presentation of several figures for enhanced readability. Overall, we believe that these revisions have significantly strengthened the manuscript and improved its transparency and reproducibility.

Technical comments:

1) The retrievals from the Bayesian approach are reported to be dependent on the a priori. It would be more informative if the authors also reported the degree of dependence. Sensitivity analysis of the retrievals on a priori will add more value to the retrievals. Such a kind of diagnostics is available from the algorithm itself and straightaway tell how much observations (and a priori) are contributing to the retrievals.

Response: We thank the Reviewer for this valuable suggestion. We agree that quantifying the dependence of the Bayesian retrieval on the a priori improves the interpretation of the results and provides an objective measure of retrieval robustness. To address this point, we performed a sensitivity analysis by repeating the inversion using different emissivity priors (ε_0 = 0.70, 0.80, 0.90, and 0.95; Table A1). We quantify the degree of dependence by evaluating the spread of the retrieved mean temperatures across these emissivity priors for each MAJIS observation. For the 4.5–5.56 µm retrieval window, the mean temperature varies by $\Delta T\_mean \approx$ 8.6 K (C1), 2.2 K (C2), 9.2 K (C3), and 12.9 K (C4) across the tested ε_0 values. These results confirm that the sensitivity to the emissivity prior is generally modest for intermediate illumination conditions (C2), while it becomes more significant for extreme cases, particularly for high-temperature conditions where the radiance–temperature relationship is highly nonlinear. We have added a dedicated paragraph in Sect. 2.1, immediately after the paragraph introducing the dependence on the emissivity prior, to explicitly report this quantitative sensitivity.

2) Bayesian inversion is all about giving proper weights to the observations and *a priori*. There is no mention of the covariance matrices used in the retrieval. Although the authors have mentioned the impact of the covariance matrix at many places in the text. Please mention how the covariance matrices (both observation and *a priori*) are prepared. Whether full covariance matrices are used for both observations and *a priori*? If only diagonal matrices are used, then what is done to compensate for the missing covariances? Hyperspectral observations are usually correlated. In practice, only the diagonal observation error

covariance matrix is used, and errors are inflated. For *a priori* covariance matrix, the emissivity may be assumed to be uncorrelated, but temperature and emissivity are highly correlated, so their correlations must be included. Ignoring covariances will result in suboptimal retrievals. Physically inconsistent results mentioned on page no. 13, line no. 272 can be better explained with the help of this information.

Response: We thank the Reviewer for this important and technically well-founded comment. We agree that the construction of the observation error covariance matrix (S_e) and the a priori covariance matrix (S_a) is central to the Bayesian (optimal estimation) inversion and should be described explicitly. We have therefore expanded Sect. 2.1 to detail how both covariance matrices are defined in our implementation.

For the observation error covariance matrix (S_e), we assume uncorrelated errors between spectral channels and adopt a diagonal matrix whose elements are given by the in-flight Noise Equivalent Spectral Radiance (NESR) of MAJIS, evaluated at each wavelength. This choice is motivated by the lack of a validated spectral correlation model for MAJIS radiance noise at the time of writing, and it follows common practice in hyperspectral retrievals when channel-to-channel correlations are not available. We note that the diagonal S_e provides conservative uncertainties because the NESR already includes the dominant instrumental noise sources and increases toward longer wavelengths due to thermal background.

For the a priori covariance matrix (S_a), we apply (i) a temperature prior variance corresponding to ±30 K (soft constraint), and (ii) a wavelength-dependent emissivity prior with physically motivated bounds ($\varepsilon \leq 1$) and smoothness enforced through spectral correlations between neighbouring channels. In practice, S_a includes off-diagonal terms for emissivity, modelled with a Gaussian correlation function whose correlation length is comparable to the spectral sampling/resolution. This regularizes the emissivity solution and mitigates unrealistic channel-to-channel oscillations. The coupling between temperature and emissivity is naturally represented through the Jacobian and the posterior covariance, which captures their correlation in the retrieved state.

We clarify in the revised manuscript that the remaining physically inconsistent emissivity behaviour near the crossover region is primarily linked to the intrinsic temperature–emissivity degeneracy and to local observational conditions (e.g. low thermal signal and/or high incidence angle), rather than to a failure of the Bayesian formulation. We have added this information in Sect. 2.1 and referenced Rodgers (2000) for the formalism.

3) It is not clear whether topographic corrections are implemented in the Bayesian approach for retrieval. The sentence on page no. 6, line 184, “Topographic effects may be mitigated…” is creating doubt. It’s not certain. Further, the sentence on page no. 7., line no. 185, “However, while such corrections are critical… a few hundred meters.” is creating more suspicion. On page no. 13, line no. 298, “The Moon’s global shape can be approximated by a smooth ellipsoid… rather than geometric scattering.” emphasize that topographic effects are not important at MAJIS observations’ spatial resolution, but on page no. 16, line no. 370, “This discrepancy likely arises from unresolved topographic effects… can dominate the measured radiance and produce local deviations from the expected thermal behavior.” and Page no. 23, line no. 491, “At higher… in temperature retrievals”, the authors attribute the retrieval errors to local topography.

Response: We thank the Reviewer for highlighting this ambiguity. We agree that the manuscript could be interpreted as contradictory regarding the role of topography. We clarify here that no explicit topographic correction is implemented in the Bayesian retrieval (Tosi et al., 2014, Sect. 2.1

of the paper) and in the empirical thermal correction (Li and Milliken, 2016, Sect. 2.2 of the paper), i.e. we do not use a DEM to compute local slopes, local incidence/emission angles, or sub-pixel shadowing corrections. The MAJIS data are georeferenced using SPICE-based geometry and an ellipsoidal lunar reference surface. However, a LOLA-derived digital roughness map is used in Sect. 2.3 to properly inform the thermal model, following the framework of Wohlfarth et al. (2023).

The statement that an ellipsoidal approximation is acceptable refers to the global mapping scale and to the fact that the MAJIS pixel size during the flyby is of the order of a few hundred metres, so that large-scale trends in temperature and emissivity can be robustly interpreted without DEM-based corrections. However, local discrepancies between MAJIS retrievals and Diviner-derived products can still arise from unresolved sub-pixel topographic effects (e.g. mixed slopes and shadowing within a pixel), which are not represented in the ellipsoid-based geometry and can locally bias the measured radiance. We have revised the relevant sentences in Sect. 2.1 to explicitly state that topographic corrections are not applied, and to distinguish between large-scale mapping robustness and local residual effects.

4) The retrievals from the three algorithms mentioned in the study provide the physical lunar surface temperature, but LRO Diviner provides bolometric temperature. Authors have compared the surface temperature retrievals against the Diviner peak brightness temperature (determined by parabolic fitting of channels 3, 4, and 5). It is still brightness temperature not physical surface temperature. Authors should add a justification for making such a comparison and should also mention this while discussing discrepencies in the results.

Response: We thank the Reviewer for this important clarification. We agree that MAJIS retrievals aim at estimating a physical surface temperature (together with wavelength-dependent emissivity), whereas Diviner products provide brightness temperatures (per channel) and a peak brightness temperature derived from a parabolic fit of channels 3–5 around the Christiansen Feature (CF). These quantities are not strictly equivalent and should not be interpreted as identical physical temperatures.

Our comparison is therefore intended as a consistency check using Diviner peak brightness temperature as a practical proxy for the surface thermal state at the time of observation, rather than a one-to-one validation of absolute thermodynamic temperature. The Diviner peak brightness temperature is particularly suitable for this purpose because it is derived near the CF, where the effective emissivity is close to unity and the dependence on emissivity is reduced compared to single-channel brightness temperatures. We have revised the manuscript to explicitly state this limitation and to clarify that residual discrepancies between MAJIS and Diviner can partly arise from the intrinsic difference between physical temperature estimates and brightness-temperature products, in addition to spatial resolution and unresolved topographic effects.

5) In the "Roughness-informed thermal model" used in this study, it is not clear from the text whether whole model (Reflectance + Emission) given by Wohlfarth et al. (2023) is implemented or only the emission model is used. Wohlfarth et al. (2023) used Hapke model for reflectance which involves various parameters. Moreover, single scattering albedo is used as a free parameter in reflectance model. If whole model is used then the results are self-explanatory.

Response: We thank the Reviewer for this useful remark. In the original version of the manuscript, we implemented only the thermal formulation presented in Wohlfarth et al. (2023). To improve the robustness of the analysis, we have revised this implementation by developing and applying a combined reflectance and thermal emission model following the full framework described in

Wohlfarth et al. (2023). For the reflected component of the signal, we adopt the definitions provided in Hapke (2005). In addition, for consistency with previous studies (e.g., Bandfield et al., 2015; Müller et al., 2021; Wohlfarth et al., 2023), we derived an average estimate of the cm–mm scale roughness for selected areas within C4, which is characterized by high signal-to-noise ratio (SNR) and stronger thermal emission compared to C1–C3. This approach provides a more representative description of the actual sub-pixel roughness than the meter-scale roughness derived from LRO/LOLA data. In light of these changes, Section 2.3 of the manuscript has been updated accordingly, including revised figures and tables (also in the Appendix), and the analysis has been repeated using this new approach. The Hapke parameters adopted for the implementation of the full model are taken from Li and Li (2011).

6) Figure A1. in Appendix A, (a) where is the noise plot? The plot in figure (b) is skeptical, how can the posterior uncentainty be so low (σ ≈0.01 K)? It is given as (see Rodgers 2000):

$$\boldsymbol{S} = (\boldsymbol{S}_a^{-1} + \boldsymbol{K}^T \boldsymbol{S}_\varepsilon^{-1} \boldsymbol{K})^{-1}$$

Is it verified? It can only happen in the case of almost perfect observations. It is difficult to say anything without having the knowledge about the observation and *a priori* errors used in the retrieval.

Response: We thank the Reviewer for this useful technical remark. In Fig. A1(a), the noise curve (NESR) is plotted in light blue. However, for the selected best-case pixel its amplitude is much smaller than the measured/modelled radiance in the 4.5–5.5 µm interval, and therefore it appears nearly overlapped by the other curves at this plotting scale.

Regarding the very low posterior uncertainty shown in Fig. A1(b) ($\sigma \approx 0.01$ K), we agree that such a narrow posterior can only occur in particularly favourable conditions. In the original version, Fig. A1 was intentionally constructed using a "best-case" pixel from MAJIS observation C3, selected to provide a clear illustration of the Bayesian retrieval behaviour under high signal-to-noise conditions. In this regime, the posterior covariance matrix can yield very small formal uncertainties when the radiance sensitivity to temperature is strong and the observational noise is low.

To address the Reviewer's concern and improve transparency, we have now revised Appendix A by updating Fig. A1 to include both (i) a best-case example and (ii) a representative suboptimal pixel from the same MAJIS cube, characterized by a significantly larger formal uncertainty. The revised figure therefore illustrates how the posterior temperature distribution broadens under less favourable observational conditions, as expected from the Rodgers (2000) formulation.

Minor comments:

1) Page no. 5, line no. 138, What does "Kirchhoff-derived estimate" mean? Kirchoff's law is $r = 1 - \varepsilon$. Both r and ε are unknown. Is estimate of r available from any other source?

Response: We thank the Reviewer for pointing out this ambiguity. In our retrieval, "Kirchhoff-derived estimate" does not refer to an external measurement of reflectance. Rather, it refers to an internal consistency check based on Kirchhoff's law applied in a spectral, directional sense under the standard assumptions of thermal equilibrium and an opaque surface (negligible transmittance). In this framework, spectral directional emissivity equals the corresponding spectral directional absorptivity; for an opaque surface this can be written as

$$\varepsilon_d(\lambda) = 1 - r_{hd}(\lambda),$$

where r_{hd} is the hemispherical–directional spectral reflectance (consistent with the directional emissivity definition). Practically, we compute a Kirchhoff-consistent reflectance proxy $r_{hd}(\lambda) = 1 - \varepsilon_d(\lambda)$ from the retrieved emissivity and use it only as a diagnostic to flag spectral segments near the crossover region where the inversion becomes ill-conditioned and may produce non-physical emissivity behaviour. We have revised the manuscript to clarify the emissivity/reflectance definitions and to avoid implying that reflectance is independently known or that a bolometric Kirchhoff relation is used.

2) Figure 13(b), there are three plots in this figure, all shown in shades of gray. It is difficult to identify which is what. Authors should consider revising this figure for better presentation.

Response: We thank the Reviewer for this suggestion. We agree that the three MAJIS-derived datasets in Fig. 13 were difficult to distinguish in grayscale. We have revised Fig. 13 to improve readability by using colors other than grey. The updated figure now allows a better identification of MAJIS-derived temperature values trends for each of the three approaches.

3) Figure 14(d), I think line plots showing the mean along with error bars would be better. It requires binning in latitude and calculate the values (mean and standard deviation, or median and mean absolute deviation) in each bin.

Response: We thank the Reviewer for this helpful recommendation. We have revised Fig. 14(d) accordingly by binning the data in longitude and plotting the binned central tendency together with uncertainty bars. Specifically, for each latitude bin we compute the mean value and its dispersion (standard deviation), which is shown as error bars. This revised representation improves clarity and makes the latitude trends easier to interpret compared to the original scatter plot.

4) I think it should be "Kelvin" not "kelvin", throughout the text.

Response: We thank the Reviewer for the remark. According to SI conventions, the unit name is written in lowercase ("kelvin"), while the unit symbol is uppercase ("K"). We have ensured consistency throughout the manuscript, using "K" for temperature values and "kelvin" for text/captions.