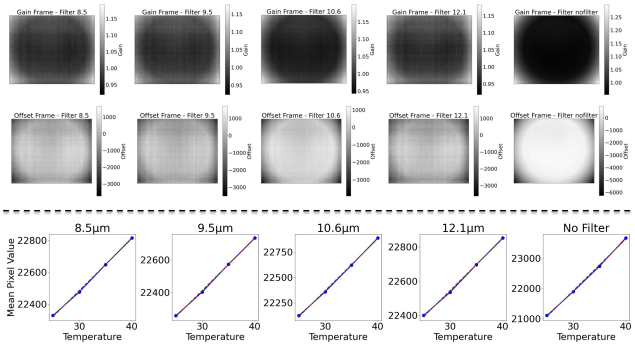

Figure 7. Top: Shows the gain and offset correction images for the various spectral bands. Unlike [22] which only performs offset correction, we perform both gain and offset correction for accurate recovery of incoming radiation. Bottom: Shows blackbodies at different temperatures and their corresponding pixel values in blue, with the exponential Sakuma-Hattori fit [29] and linear fit for the pixel values is shown in dotted red and green respectively. As seen, both temperature to camera counts curve looks linear and the difference between the Sakuma-Hattori and linear curve is less than a couple of counts, which is much below the noise floor of the thermal camera.

A. Calibration and Pre-Processing

In this section, we provide more details on our camera calibration procedure.

Gain-Offset Correction: We begin by applying gain and offset correction using the gain and offset frames shown in Figure 7. To calibrate, we position a blackbody at various temperatures in front of the camera, ensuring it fills the entire field of view. The true pixel value is determined as the mean pixel count across the entire image. For each pixel, we compute gain and offset values so that the corrected output matches this mean. This calibration is performed in a least-squares sense across multiple blackbody temperatures.

Thermal Camera Calibration: Similar to gain-offset correction, we use blackbody videos at different temperatures to perform radiometric calibration of thermal cameras. Figure 7 presents the calibration results across various spectral bands, comparing a Planckian Sakuma-Hattori fit [29] with a linear fit. As observed, the camera response is nearly linear for objects at ambient temperatures.

B. More Experimental Results

Ablation Studies: Figure 8 presents the ablation results of our optimization, where each loss term is individually removed from the pipeline. As shown, the reconstruction loss has the most significant impact, while the smoothing and Huber losses have a comparatively smaller effect on overall accuracy. Initially, optimizing noise parameters does not provide a noticeable improvement, but at lower SNR values, failing to account for noise leads to higher errors. Fig. 8 represents the SNR level breakdown of Tab. 1 shown in the main paper.

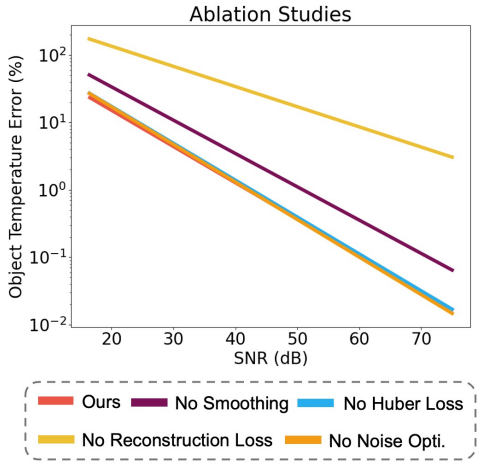

Figure 8. Ablating the loss terms in simulation: We individually disable each loss term from the full optimization pipeline. Note the log scale on the plots. Reconstruction loss is obviously the most important term followed by smoothing, huber and noise optimization. The noise term is useful at high noise levels.

Breakdown of Simulation Evaluation: Fig. 10 provides a detailed breakdown of the simulation results summarized in Fig. 2 in the main paper.

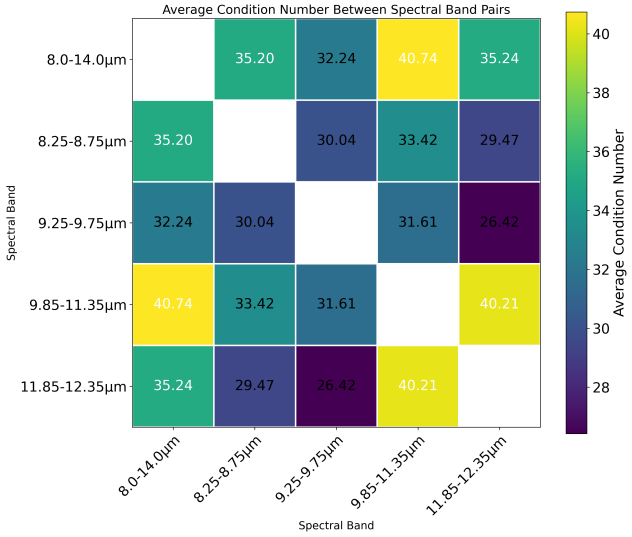

Figure 9. Condition numbers of emissivity matrix $\mathbf{E}$ in Eq. 15 for different spectral band pairs, averaged over materials in the spectral library [3, 21].

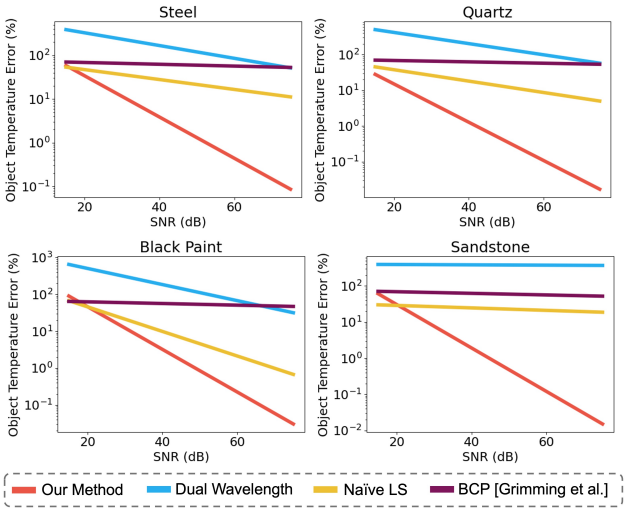

Figure 10. Comparison of our method with a recent BCP [9] technique to remove reflections, a naive multi-wavelength approach and a traditional dual-wavelength pyrometry technique [2, 39] using simulated thermal videos of different materials sourced from spectral library [3, 21]. For naive least squares, we run the optimization with five initializations and select one that achieved the least objective compared to ground truth (which we will not have access to at test time). At high noise levels, all methods have a large error as the problem is too under constrained. As noise decreases to more reasonable levels, our method performs significantly better. Note the log scale on the plots.

Choice of Spectral Filters: The selection of spectral filters for thermography plays a crucial role in distinguishing emitted and reflected light. Figure 9 illustrates the condition number of the matrix $\mathbf{E}$ from Eq. 15 for different spectral filter choices in our experiments, using emissivities sourced from the spectral library [21]. The choice of spectral bands involves a trade-off between the condition number of the resulting emissivity matrix and the noise introduced by each filter. Generally, wider spectral bands reduce noise levels; hence, we select 8–14μm as one of the spectral bands. Additionally, Figure 9 shows that the 9.5μm central wavelength filter achieves the lowest condition number, indicating a better-conditioned linear system for our application.