Neural Radiance Fields with Physically Based Reflectance for Satellite Images

Lulin Zhang[1], Ewelina Rupnik[2], Tri Dung Nguyen[1], Stéphane Jacquemoud[1], Yann Klinger[1]

[1] Université de Paris, Institut de Physique du Globe de Paris, CNRS, Paris, France - (lzhang, tridnguyen, jacquemoud, klinger)@ipgp.fr
[2] Univ. Gustave Eiffel, IGN-ENSG, LaSTIG – Saint-Mandé, France - ewelina.rupnik@ign.fr

Keywords: neural radiance fields, BRDF, Hapke, albedo, roughness, satellite images

Abstract

Characterising planetary surface optical properties from satellite imagery is key to understanding surface composition and texture. Neural Radiance Fields (NeRF) have recently demonstrated strong potential for 3D scene reconstruction and novel view synthesis, yet most existing formulations rely on empirical or Lambertian reflectance assumptions that limit their physical interpretability. In this work, we extend BRDF-NeRF by replacing its empirical RPV reflectance model with the Hapke model, a physically grounded BRDF widely used in the planetary remote sensing community to characterise surface scattering properties. Our architecture integrates the Hapke rendering equation into a NeRF backbone and jointly estimates four physically interpretable surface parameters - single-scattering albedo w, phase function parameters b and c, and photometric roughness θ - directly from a sparse set of satellite images. We evaluate the capacity of the framework to recover these parameters against two sources of ground truth: Hapke parameters estimated by a competitive inversion method, and reflectance spectra measured in laboratory conditions using a spectroradiometer. Our experiments show a high correlation between the estimated and reference single-scattering albedo, while the roughness parameter θ is systematically underestimated. In terms of 3D reconstruction and novel view synthesis, our Hapke-NeRF performs comparably to BRDF-NeRF with the RPV model, demonstrating that physically grounded reflectance can be incorporated into NeRF without sacrificing geometric accuracy.

1. Introduction

In recent years, several *Neural Radiance Fields* (NeRF) variants have been adapted to remote sensing images Derksen and Izzo (2021). With sufficient satellite viewpoints, NeRF can estimate surface elevations with high fidelity, outperforming traditional stereo-matching methods Marí et al. (2023). To this end, NeRF represents the scene as a set of MLPs and minimises a cost function that compares the observed pixel intensities with those predicted through volumetric rendering within that representation space.

However, some characteristics of remote sensing imagery challenge the *state-of-the-art* NeRF. First, satellite acquisitions typically provide only two or three views, leading to a marked performance decline even when priors are used to constrain the scene volume Zhang and Rupnik (2023). Then, remote sensing images do not follow the simplistic isotropic or specular reflection models, as is the case for natural images Srinivasan et al. (2021). Surface appearance in remote sensing images depends, in part, on the surface material and its reflectance properties. Accounting for this requires adopting some Bidirectional Reflectance Distribution Function (BRDF) within the NeRFs rendering equation Zhang et al. (2025). Last but not least, shadows Derksen and Izzo (2021); Masquil et al. (2025), transient objects Marí et al. (2022) or landscape change are all equally challenging for both NeRF and the traditional stereo matching.

Contributions. In this work, we focus on the BRDF. Unlike BRDF-NeRF Zhang et al. (2025), which employs a semi-empirical reflectance model, we propose a formulation related to physical parameters. Specifically, we adopt the Hapke radiative transfer model Hapke (1981), which is known to efficiently model radiance–surface interactions using a small set of parameters and is widely employed in the remote sensing community to characterise Earth's surface properties Labarre et al. (2017). We further conduct an experimental comparison between the empirical and physical BRDFs and discuss their respective applicability.

Why a physically-based BRDF? Predicting the physical properties of the Earth's surface is essential for analysing its underlying geological processes. For instance, single-scattering albedo can indicate the surface's mineralogical composition, while surface roughness controls the energy balance of bare soils. High-resolution remote sensing imagery, with broad spatial coverage and frequent revisit times, provides a unique data source for deriving such parameters at fine scale. It has been shown that with many images capturing the scene from a wide range of viewing angles, the BRDF estimation and subsequent inversion of surface physical parameters through radiative transfer models are feasible Labarre et al. (2019); Nguyen et al. (2025). In practice, however, acquiring more than a few near-synchronous satellite images of the same area is uncommon. In the following, we propose an approach to directly predict parameters of a physical radiative transfer model while leveraging much fewer images. For clarity, in this work, we use albedo to denote the single-scattering albedo (SSA), and roughness to denote the photometric roughness parameter. Our approach employs NeRF framework illustrated in Figure 4.

2. Related work

2.1 NeRF overview

Neural Radiance Fields (NeRF) (Mildenhall et al. (2020)) represent a scene as a continuous function that maps 3D coordinates and viewing directions to colour and density, enabling high-quality novel view synthesis from multi-view images. However, NeRF typically requires dense input views. When trained with sparse images, it often overfits and produces inaccurate geometry and rendering artefacts. To mitigate

https://doi.org/10.5194/isprs-archives-XLIX-B2-2026-601-2026 |

this issue, recent approaches introduce additional priors, most commonly depth supervision. Both sparse depth (Deng et al. (2022)) and dense depth (Zhang and Rupnik (2023)) have been used to guide geometry learning and improve reconstruction quality under limited viewpoints.
Another limitation of NeRF is that it cannot relight scenes or model material properties. To address this, several works extend NeRF with BRDF modeling, capturing how surfaces reflect light under varying view and lighting conditions. Most BRDF-aware NeRF variants, such as (Bi et al., 2020; Srinivasan et al., 2021; Boss et al., 2021; Verbin et al., 2022), adopt the microfacet BRDF model (Walter et al., 2007), which decomposes reflectance into diffuse and specular components with a surface roughness parameter. While effective for parameterization, microfacet models often fail to capture natural surface effects like backscattering (hotspot). Data-driven BRDFs learned from pre-collected BRDF databases have also been incorporated into NeRF (Zhang et al., 2021), but these databases primarily contain synthetic materials measured under controlled laboratory conditions. Recent work (Zhang et al., 2025) represents the first NeRF approach to incorporate a BRDF model tailored for earth observation scenes by employing the RPV model to capture anisotropic reflectance of natural surface.

2.2 BRDF from satellite images

The inversion of BRDF models from satellite images presents both significant challenges and opportunities. A wide range of BRDF models have been developed to describe the spectral and directional reflectance of natural and artificial surfaces. These models are generally classified into three categories: physical, empirical, and semi-empirical models. Physical models (Pinty and Verstraete (1991)) rely on rigorously defined physical parameters and provide highly accurate representations of observed scenes, but require a large number of multi-angular observations. Empirical models (Minnaert (1941)) are derived from statistical fits to observed data and do not provide physical interpretability of surface properties. Semi-empirical models (Hapke (1981); Rahman et al. (1993); Lucht et al. (2000)) strike a balance by using a limited number of parameters to approximate the interactions between radiation and surfaces. Among them, the Hapke and RPV models are widely used in remote sensing. The Hapke model offers the advantage of physically meaningful parameters and an ability to predict and describe radiometric measurements. In contrast, the RPV model is mathematically simpler, making it more stable in deep learning frameworks, particularly when only a limited number of training images are available.

Several studies have explored BRDF model inversion using multi-angular satellite imagery. Labarre et al. (2019) inverted the Hapke model using 21 multi-angular Pleiades images with a spatial resolution of 2 m, manually selecting pixels and retrieving BRDFs prior to parameter estimation. Extending this work, Nguyen et al. (2025) performed large-scale Hapke model inversion over the Asal-Ghoubbet rift using Bayesian inversion, enabling broader spatial coverage and a more detailed representation of surface units. More recently, deep learning approaches have been introduced. Zhang et al. (2025) employed NeRF to estimate the RPV model using only three Pleiades images. While this approach is effective for surface estimation and novel view synthesis, the RPV parameters lack direct physical interpretability.

In this work, we aim to bridge this gap by leveraging the NeRF framework to estimate the Hapke model from a limited set of satellite images, combining the advantages of physically-based modeling with the efficiency of learning-based approaches.

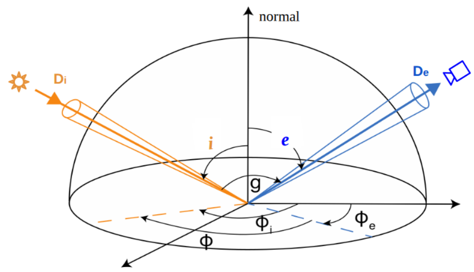

Figure 1. **Directions Definition**. i and e denote the illumination and viewing angles, respectively, while g is the phase angle between them. Φ_i and Φ_e represent the corresponding azimuth angles of i and e defined in the plane orthogonal to the normal. The relative azimuth angle is given by $\Phi = \Phi_i - \Phi_e$.

3. Hapke Reflectance Model

The Hapke reflectance model offers two key features: it is fully characterised by only four parameters, and the parameters carry physical implications, enabling interpretable surface characterisation. Under this model, the colour of a pixel is predicted as:

$$\begin{aligned}\mathbf{c} = \frac{w}{4\mu_0}\frac{\mu_{0e}}{\mu_{0e}+\mu_e} \\ \{p(g,b,c) + H(\mu_{0e})H(\mu_e) - 1\}\, S(i,e,g,\theta)\end{aligned} \tag{1}$$

where w, b, c, θ are the Hapke parameters; illumination i, viewing e and phase g angles are illustrated in Figure 1, $\mu_0 = \cos i, \mu_e = \cos e$ are the so-called cosines of *effective angles* which account for the roughness effect. $H(\mu_{0e})$ and $H(\mu_e)$ are the Ambartsumian–Chandrasekhar functions, which model the cumulative effect of light bouncing multiple times between particles before reaching the sensor. $S(i, e, g, \theta)$ is the shadowing function, $p(g, b, c)$ is the phase function. In what follows, we provide a detailed description of each Hapke parameter, together with the phase and shadowing functions. The full expressions for μ_{0e}, μ_e, $H(\mu_{0e})$, $H(\mu_e)$ and $S(i, e, g, \theta)$ are given in the appendix.

Albedo w The albedo w is the probability that a photon interacting with a particulate medium is scattered rather than absorbed. It ranges from 0 to 1 and depends only on the intrinsic optical properties of the particles, such as their composition and microstructure (e.g., grain size and shape), rather than on measurement conditions. w controls the overall brightness of the surface reflectance and is widely used to characterize material composition and to analyze mixed powders. Moreover, it is the most well-constrained parameter. However, due to its physical meaning, it is inherently correlated with other parameters, so their variations can affect its value.

Parameters b, c and phase function The phase function describes the probability that a photon incident from one direction is scattered into another direction. One of the most widely used phase functions is the two-parameter Henyey-Greenstein (HG) (Equation (2)), and we adopt it in this work. As shown in Figure 2, scattering behaviour is characterised by the parameters b and c, where the former determines scattering asymmetry, and the latter dictates the forward/backward distribution.

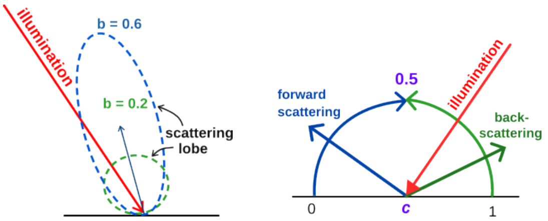

Figure 2. **Phase Function**. The parameter b controls scattering asymmetry: $b \approx 0$ gives isotropic scattering, while $b \to 1$ favors a specific direction. The parameter c governs forward/backward distribution: $c > 0.5$ favors backscattering, $c < 0.5$ favors forward scattering. Modified after Souchon (2012).

$$p(g, b, c) = (1 - c)\,\frac{1 - b^2}{(1 - 2b\cos g + b^2)^{3/2}} + c\,\frac{1 - b^2}{(1 + 2b\cos g + b^2)^{3/2}} \tag{2}$$

Roughness θ and shadowing function S Shadowing function describe how surface irregularities affect light reflection. It accounts for the reduction in reflectance caused by shadows on a rough surface. These shadows arise from two mechanisms: (1) facets that are not illuminated because their orientation exceeds 90 ° relative to the direction of the incident light, and (2) facets that block incoming light from reaching other facets. The reflectance is corrected for this phenomenon of shading by involving the roughness parameter θ, which represents the mean slope angle integrated over all scales from the grain size to the local topography. It controls how strongly shadowing affects reflectance and is mainly related to the surface microstructure and morphology. Ideally, roughness requires measuring slope angles continuously across that full range of scales; however, such comprehensive measurement is not feasible. As a practical compromise, in this work, we consider slope angle calculated from very high resolution DSM as GT.

4. Integration within NeRF Architecture

NeRF backbone NeRF models a continuous volumetric field of a static scene using a fully connected network. Given a 3D point $\mathbf{x} = (x, y, z)$ and viewing direction $\mathbf{D}_e = (d_x, d_y, d_z)$, it predicts volume density σ and colour $\mathbf{c} = (r, g, b)$. For each camera ray $\mathbf{r}(t) = \mathbf{o} + t \cdot \mathbf{D}_e$, N points $\mathbf{x}_p = \mathbf{o} + t_p \cdot \mathbf{D}_e$ are sampled between near and far bounds. The rendered colour is:

$$\mathbf{C}(\mathbf{r}) = \sum_{p=1}^{N} T_p \alpha_p \mathbf{c}_p\ , \tag{3}$$

where $\alpha_p = 1 - e^{-\sigma_p \delta_p}$, $T_p = \prod_{q=1}^{p-1}(1 - \alpha_q)$ and $\delta_p = t_{p+1} - t_p$. α_p denotes opacity and T_p denotes transmittance, jointly determining each sample's contribution. During training, all MLPs' parameters are iteratively optimized through minimisation of the discrepancy between the rendered colour and the observed pixel colour. In this formulation a lambertian surface reflectance assumption is made.

Hapke-NeRF Our NeRF architecture relaxes the Lambertian assumption and consists of three components (see Figure 4): a geometric module, a radiometric module, and a renderer. Surface normals' define the viewing and illumination direction thus their accuracy is critical. Therefore, to provide stable normal estimates before Hapke training begins, the geometric module is first pre-trained under a Lambertian assumption and supervised with low-resolution depth maps from Semi-Global Matching (Pierrot-Deseilligny and Paparoditis, 2006). The Hapke MLPs are then introduced and the renderer $\mathbf{c}$ in Equation (3) is replaced with Equation (1) for joint fine-tuning. Like SpS-NeRF (Zhang and Rupnik (2023)), we (1) employ ray sampling strategy with 64 stratified points along with 64 guided points drawn from a Gaussian distribution, and (2) apply supervision with the ground truth pixel colour and the dense depth information generated from SGM.

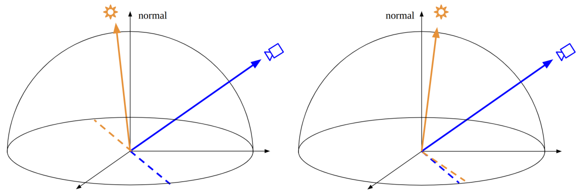

Figure 3. **Differentiability Issue**. Φ changes abruptly between 180° and 0° when the normal undergoes a slight shift between scenarios (a) and (b).

Surface Normal and Differentiability The Hapke rendering equation is not differentiable everywhere due to two distinct issues: (i) the discontinuity of the relative azimuth angle Φ, and (ii) the undefined values of $\tan\theta$, $\cot\theta$, $\cot i$, and $\cot e$ at certain angles, which appear in the modified cosine terms and the shadowing function. The discontinuity of Φ arises from a sign ambiguity in its definition: Φ jumps from 180° to 0° when the surface normal crosses the plane containing the illumination and viewing directions, as illustrated in Figure 3. Since the illumination and viewing directions are fixed, this discontinuity manifests solely through changes in the surface normal during training. To ensure training stability, the surface normal is initialized and supervised using low-resolution depth maps produced by SGM, rather than being optimized freely from scratch. More specifically, the normal from depth is estimated by taking cross products of vectors to its neighbours and averaging them for a more robust result. To address the singularities of $\tan\theta$ and $\cot\theta$, the roughness parameter θ is constrained to $(0°, 60°]$: the strict exclusion of 0° ensures that $\tan\theta$ and $\cot\theta$ remain finite throughout training, while the upper bound reflects the maximum expected mean slope within the scene (Hapke, 1981).

5. Experiments

Dataset We conduct experiments using the Djibouti dataset, located in the Asal-Ghoubbet rift in the Republic of Djibouti. This dataset comprises 21 multi-angular Pléiades 1AB images acquired during a single flyover on January 26, 2013, the corresponding camera and sun positions are illustrated in Figure 5(a). Three areas of interest (AOIs), labelled A, B, and C, are selected for the experiments (Figure 6). Together, these AOIs encompass six samples of varying soil type and roughness. The AOIs were selected so as to maximise the number of available samples unaffected by cloud cover. Soil samples (Figure 7) were collected in the Asal-Ghoubbet rift in February 2016, and their reflectance spectra were measured in the laboratory using a FieldSpec© 4 ASD spectroradiometer (Nguyen et al., 2025). Those spectra represent our ground truth.

This contribution has been peer-reviewed.
https://doi.org/10.5194/isprs-archives-XLIX-B2-2026-601-2026

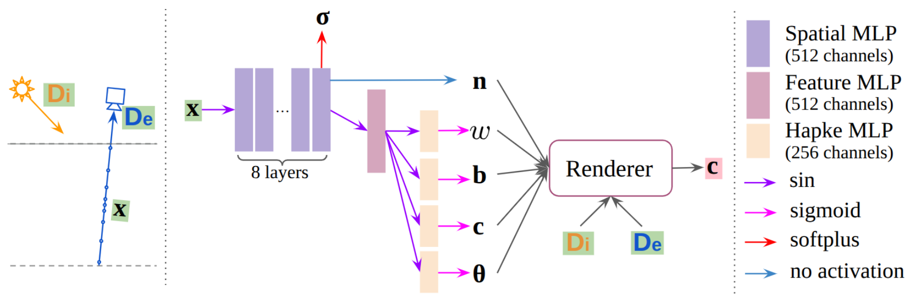

Figure 4. **The Hapke-NeRF Architecture**. The geometric module takes a 3D point location $\mathbf{x}$ as input and passes it through eight spatial MLPs to predict volume density σ, surface normal $\mathbf{n}$, and a shared feature vector. The feature vector is then processed by four independent Hapke MLPs, each predicting one of the Hapke parameters w, b, c, and θ. The renderer combines σ, $\mathbf{n}$, w, b, c, and θ with the illumination and viewing directions $\mathbf{D}_i$ and $\mathbf{D}_e$ to compute the rendered colour $\mathbf{c}$ via Equation (1).

Evaluation Seven images with viewing angles closest to nadir (image IDs 8–14) are used to train the NeRF model. The number of input images is limited to maintain consistent pixel size across views (see Figure 5(b)). For novel view synthesis evaluation, we select an image near the acquisition extremity (image ID 3), which constitutes a challenging extrapolation scenario relative to the training viewpoints. The evaluation covers two aspects:

- Hapke parameter estimation: the retrieved Hapke parameters are compared against the ground truth of Nguyen et al. (2025).
- Novel view synthesis and DSM quality: the fidelity of the synthesised image and the accuracy of the estimated digital surface model (DSM) are assessed against ground truth and the competitive RPV-NeRF (Zhang et al., 2025).

To assess the influence of spatial resolution, two training configurations are considered:

1. Images downsampled by a factor of 5 (GSD = 10 m) for AOIs A, B, and C, matching the spatial resolution of the Bayesian inversion ground truth. Both evaluation aspects above are reported.
2. Original resolution images (GSD = 2 m) for AOI A, evaluated for novel view synthesis and DSM quality only. Results for this configuration are denoted as A_{Z1}.

We also perform an ablation study and compare different BRDF models for volume rendering:

- Lam: Lambertian BRDF, used as a baseline;
- RPV: model implementing the empirical RPV BRDF (Zhang et al., 2025), with the first 20% of training dedicated to Lambertian pretraining to aid convergence;
- Hpk: Hapke BRDF model, similarly pretrained on the Lambertian model for the first 20% of training;
- Hpk_0: Hapke BRDF model, trained from scratch without Lambertian pretraining, serving as an ablation of the pretraining strategy.

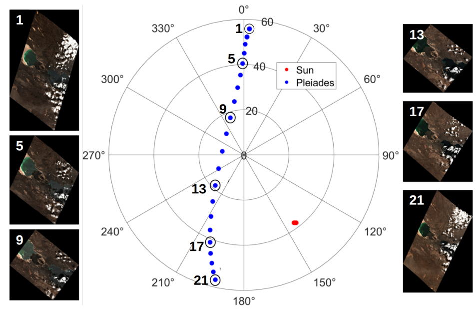

(a) A polar plot illustrating the geometry of the Pleiades acquisition: red dots show the Sun’s directions, blue dots are the viewing directions. The distance from the center represents the zenith angle, the angle around the plot (azimuth) is measured from geographic north. The 7 most nadir images (i.e., images 8–14) are used to train our model.

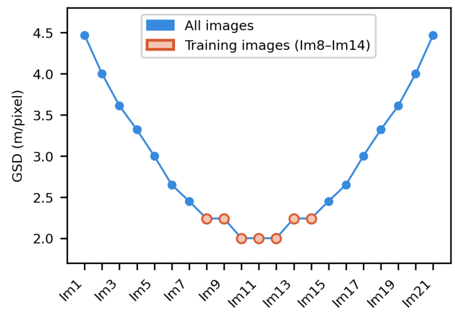

(b) Ground Sampling Distances (GSD) along the acquisition.

Figure 5. **Djibouti Dataset**.

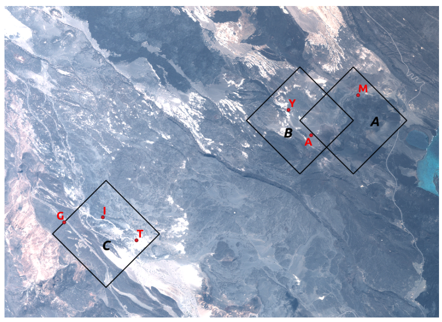

Figure 6. **Evaluation Datasets**. Localization of the AOIs: A, B, C, and soil samples A, Y, M, G, I, T.

Implementation details The albedo is constrained to $[0, 0.8]$, as values exceeding 0.8 are known to degrade the accuracy of roughness estimation (Nguyen, 2024). For configurations involving the Hapke and RPV BRDF models, surface normals are supervised using a low-resolution depth map generated via SGM, rather than analytically derived normals. This choice is motivated by two reasons: first, the SGM-derived depth map is already precomputed to guide NeRF sampling, so extracting surface normals from it adds negligible computational overhead; second, analytical normals are incompatible with the Hapke model, which is not differentiable everywhere.

5.1 Results and discussion

5.1.1 Hapke parameter estimation For Hapke parameter evaluation, we focus on albedo w and roughness θ, as these are the only two parameters for which independent ground truth is available: w is validated against reflectance spectra measured by spectroradiometer, and θ against mean slope angles derived from a very high resolution DSM. The phase function parameters b and c, while physically interpretable in terms of scattering behaviour, have no direct field or laboratory equivalent and are therefore excluded from quantitative evaluation.

Ground truth The reference data was published in Nguyen et al. (2025) and includes:

- A Bayesian inversion of the Hapke model over the Asal-Ghoubbet rift, performed on green band of 21 Pleiades images, producing four maps of the Hapke parameters w, b, c, and θ.
- Soil and rock samples collected in the field (Figure 7), with corresponding reflectance spectra measured in the laboratory over the wavelength range 400–2500 nm using a FieldSpec© 4 ASD spectroradiometer. The spectral resolution is 1 nm, and measurements were conducted under two configurations: nadir viewing geometry and illumination at 15° from the vertical.
- High-resolution digital Surface models (DSMs) at millimeter-scale resolution generated from drone-based imagery over our six samples; the DSMs serve to calculate the mean slope angle as roughness GT.

Results The evaluation is conducted at two scales. At the global scale, the estimated albedo w and roughness θ maps are compared against those obtained by Bayesian inversion over the full extent of AOIs A, B, and C. At the local scale, the estimated w values at six sampling locations are validated against spectroradiometer measurements, and θ is compared against reference values derived from drone-based DSMs.

Global scale evaluation Figure 9 and Figure 10 present the albedo and roughness maps over AOIs A, B, and C obtained using RPV, Hpk, Hpk_0, along with the GT. The corresponding quantitative metrics are reported in Table 1. Our results show a strong correlation between the albedo predicted by Hpk and Hpk_0 and the GT. This is expected, as albedo is closely related to image intensity, which is directly observed within our NeRF. The albedo estimated using the RPV model is less consistent with the GT, which can be explained by the fact that the definition of albedo in RPV differs from that in the Hapke model. In contrast, the agreement between predicted roughness and GT is less pronounced. Comparing Hpk with Hpk_0, the former yields a more accurate estimate of albedo, whereas the latter better captures surface roughness. This indicates an inherent competition during training, where the estimations of albedo and roughness tend to conflict, often leading to improved accuracy in one at the expense of the other relative to the ground truth. Note that the RPV model does not provide an explicit estimate of surface roughness.

		A	B	C
Albedo w	RPV	0.877	0.865	0.840
	Hpk_0	0.366	0.281	0.229
	Hpk	0.374	0.261	0.214
Theta θ	Hpk_0	0.987	0.601	0.667
	Hpk	1.072	0.846	0.829

Table 1. **Global scale evaluation**. Normalised RMSE of albedo and roughness computed as the RMSE divided by the true value of the parameter.

Local scale evaluation of albedo w Since albedo cannot be measured directly, it is evaluated through reflectance spectra at the red, green, and blue bands and compared against spectroradiometer measurements, as reported in Table 2. Overall, the estimates are in close agreement with the ground truth, with the exception of samples T and I. We attribute these discrepancies to the heterogeneous composition of these samples, which consist of multiple components (see Figure 7): while the ground truth averages spectroradiometer measurements taken separately for each component, our method treats the mixture as a single homogeneous material.

Local scale evaluation of roughness θ For roughness, our estimated values are directly compared with GT derived from drone-based DSMs, as shown in Table 3. The results exhibit a noticeable discrepancy with respect to the GT, which is consistent with the findings from the global evaluation. Overall, our method tends to overestimate the roughness parameter θ.

5.1.2 Novel view synthesis and DSM quality

Novel view synthesis Synthesised image quality is assessed using PSNR, SSIM, and LPIPS. We additionally report the aggregate metric AVG (Barron et al., 2021), which compresses these three scores into a single value via the geometric mean of their corresponding error terms, enabling compact comparison across methods:

$$\mathrm{AVG} = \frac{1}{3} \cdot (10^{-\frac{\mathrm{PSNR}}{10}} + \sqrt{1 - \mathrm{SSIM}} + \mathrm{LPIPS})$$

The qualitative and quantitative results are presented in Figure 8 and Table 4. Across AOIs A, B, and C, RPV achieves the best performance among the four method, followed by Hpk_0, Hpk, and Lam. We attribute the slightly lower performance of the Hapke-based methods to their physically grounded parameterisation, which, unlike RPV, constrains the model's capacity to fit the data by encoding realistic physical interactions between parameters. Comparing A and A_{Z1}, lower spatial resolution expectedly degrades performance across all methods, with the most pronounced effect observed for the Lambertian baseline.

DSM quality Table 4 reports the mean absolute error (MAE) of the NeRF-derived DSMs with respect to a ground truth DSM generated from high-resolution Pleiades 1AB panchromatic images using SGM at a 0.5m GSD. Overall, Hpk_0 achieves the best performance, which we attribute to two factors: (1) the suitability of the Hapke model for representing the BRDF of natural surfaces, and (2) the absence of Lambertian pretraining, which gives Hpk_0 greater freedom to optimise for colour and

Figure 7. **Soil Samples**. For each sample, 2-4 images are shown. The first is an in situ photograph, with the sample ID indicated in the top-left corner. The remaining images show the individual soil, rock, or powder components collected in the field and subsequently measured in the laboratory using a spectroradiometer. For samples comprising multiple components, the reflectance spectra are averaged prior to comparison with our estimates.

Table 2. **Local scale evaluation of albedo**. Soil reflectance with nadir view and illumination at 15 ° from the vertical at RGB channels. Ours are rendered using Hapke parameters (Hpk). GTs are mesured with FieldSpec ©4 ASD spectroradiometer.

AOI	Sample	Ours (Hpk)			GT			Ours – GT		
		B	G	R	B	G	R	B	G	R
A	M	0.10	0.12	0.13	0.08	0.10	0.11	0.02	0.02	0.02
B	Y	0.11	0.12	0.14	0.10	0.12	0.14	0.01	0.00	0.00
	A	0.12	0.14	0.16	0.11	0.14	0.16	0.01	0.00	0.00
C	G	0.09	0.10	0.12	0.08	0.11	0.14	0.01	-0.01	-0.02
	T	0.22	0.26	0.32	0.21	0.34	0.44	0.01	-0.07	-0.12
	I	0.06	0.07	0.07	0.23	0.29	0.33	-0.17	-0.22	-0.26

Table 3. **Local scale evaluation of roughness**. Our roughness compared against GT extracted from drone-derived DSMs.

AOI	Sample	Ours (Hpk)	GT	Ours – GT
A	M	37.38	16.74	20.64
B	A	43.98	22.91	21.07
	Y	43.17	12.21	30.96
C	G	35.56	25.11	10.45
	T	12.37	18.64	-6.27
	I	42.21	25.02	17.19

Table 4. **Novel view synthesis and DSM quality – quantitative evaluation**. Experiments are ran with Pleiades 1AB images downsampled by a factor of 5 (i.e., GSD = 10m), except for A_{Z1}, which used the original images (GSD = 2m). Seven images are used for training across all the configurations.

AOI	Method	LPIPS ↓	PSNR ↑	SSIM ↑	AVG ↓	MAE ↓
A	Lam	0.11	22.77	0.76	0.20	1.76
	Hpk	0.06	29.02	0.86	0.14	1.71
	Hpk_0	0.03	29.68	0.86	0.14	1.68
	RPV	0.03	30.27	0.86	0.14	1.73
A_{Z1}	Lam	0.18	22.09	0.74	0.23	1.44
	Hpk	0.06	31.90	0.92	0.12	1.38
	Hpk_0	0.06	33.27	0.93	0.11	1.36
	RPV	0.06	32.69	0.93	0.11	1.37
B	Lam	0.04	22.23	0.79	0.17	2.29
	Hpk	0.05	26.64	0.85	0.15	2.07
	Hpk_0	0.03	27.30	0.86	0.14	1.96
	RPV	0.03	28.23	0.88	0.13	2.08
C	Lam	0.04	21.41	0.80	0.16	1.96
	Hpk	0.03	26.49	0.86	0.14	1.72
	Hpk_0	0.06	26.68	0.83	0.16	1.68
	RPV	0.02	27.49	0.88	0.12	1.77

depth consistency jointly. Hpk and RPV perform comparably, both clearly outperforming Lam. Comparing configurations A and A_{Z1}, the latter yields significantly better DSM accuracy, albeit at a substantially higher computational cost (approximately $25\times$).

6. Conclusion

This work investigates the integration of the Hapke reflectance model into a NeRF framework to estimate Hapke parameters from Pleiades 1AB images. The results are compared to GT obtained from Bayesian inversion, spectroradiometer measurements, and drone-derived DSMs. In addition, we compare our approach with the RPV reflectance model and analyze the performance of both methods in terms of synthesized images and surface estimation. We show that our Hapke-based NeRF can reliably estimate single scattering albedo from just seven satellite views, albeit with a tendency to overestimate surface roughness. RPV slightly outperforms our approach in novel view synthesis due to its greater parametric flexibility, whereas the physical constraints embedded in the Hapke model limit adaptability but yield superior DSM predictions. We note, however, that the competing colour and depth supervision objectives create a training tension where favoring one degrades the other.

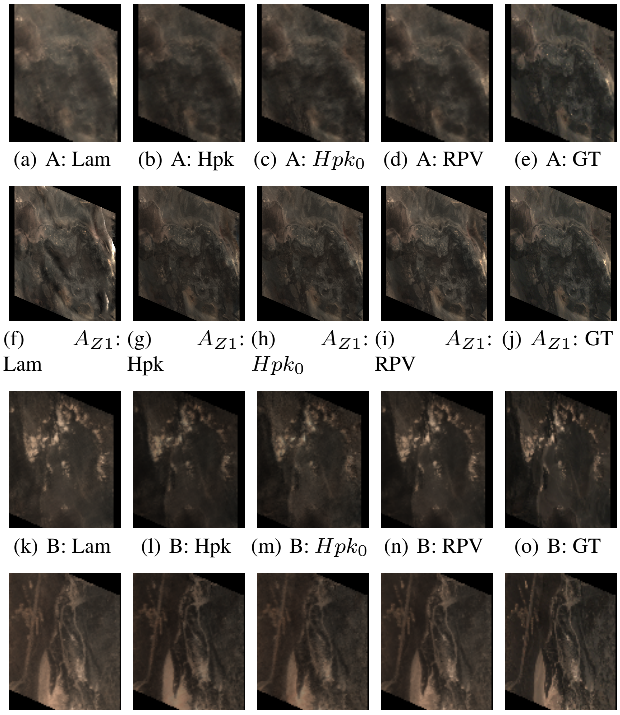

(a) A: Lam (b) A: Hpk (c) A: Hpk_0 (d) A: RPV (e) A: GT
(f) A_{Z1}: Lam (g) A_{Z1}: Hpk (h) A_{Z1}: Hpk_0 (i) A_{Z1}: RPV (j) A_{Z1}: GT
(k) B: Lam (l) B: Hpk (m) B: Hpk_0 (n) B: RPV (o) B: GT
(p) C: Lam (q) C: Hpk (r) C: Hpk_0 (s) C: RPV (t) C: GT

Figure 8. **Novel view synthesis – qualitative evaluation.** RPV performs best, followed by Hpk_0, Hpk, and Lam. Comparing A with A_{Z1} shows that reduced spatial resolution leads to worse performance for all methods.
.

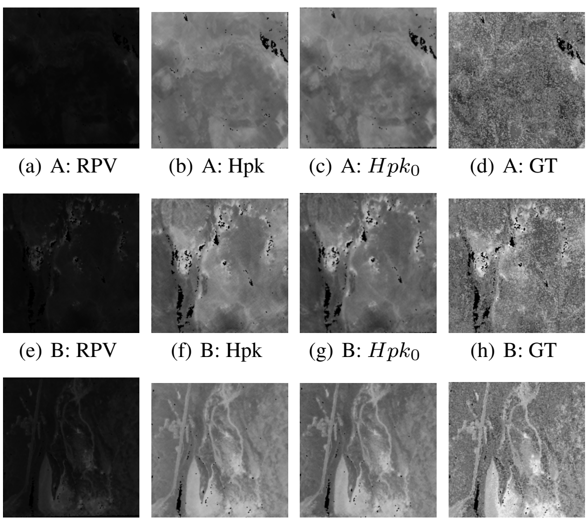

(a) A: RPV (b) A: Hpk (c) A: Hpk_0 (d) A: GT
(e) B: RPV (f) B: Hpk (g) B: Hpk_0 (h) B: GT
(i) C: RPV (j) C: Hpk (k) C: Hpk_0 (l) C: GT

Figure 9. **Albedo w maps**. Hpk and Hpk_0 produce albedo estimates close to the GT, while RPV is less consistent due to its different albedo definition.

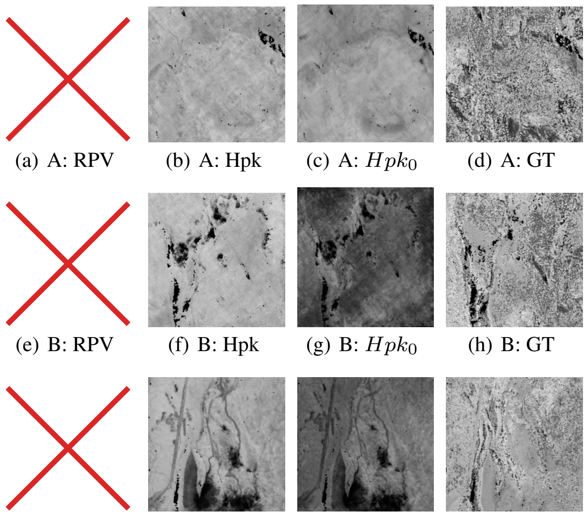

(a) A: RPV (b) A: Hpk (c) A: Hpk_0 (d) A: GT
(e) B: RPV (f) B: Hpk (g) B: Hpk_0 (h) B: GT
(i) C: RPV (j) C: Hpk (k) C: Hpk_0 (l) C: GT

Figure 10. **Roughness θ maps.** Our roughness estimates differ from the GT, likely due to the limited seven training images.

7. Appendix: details of Hapke equation

Ambartsumian–Chandrasekhar H Function

$$r_0 = \frac{1-\sqrt{1-w}}{1+\sqrt{1-w}} \tag{4}$$

$$H(x) = \left[1 - wx\left(r_0 + \frac{1-2r_0x}{2}\ln\frac{1+x}{x}\right)\right]^{-1} \tag{5}$$

Modified Cosine Terms

For $i < e$:

$$\mu_{0e} = \chi(\theta)\left[\cos i + \sin i \tan\theta \frac{\cos\Phi\, E_2(e) + \sin^2(\Phi/2)E_2(i)}{2 - E_1(e) - (\psi/\pi)E_1(i)}\right] \tag{6}$$

$$\mu_e = \chi(\theta)\left[\cos e + \sin e \tan\theta \frac{E_2(e) - \sin^2(\Phi/2)E_2(i)}{2 - E_1(e) - (\Phi/\pi)E_1(i)}\right] \tag{7}$$

For $i \geq e$:

$$\mu_{0e} = \chi(\theta)\left[\cos i + \sin i \tan\theta \frac{E_2(i) - \sin^2(\Phi/\pi)E_2(e)}{2 - E_1(i) - (\Phi/\pi)E_1(e)}\right] \tag{8}$$

$$\mu_e = \chi(\theta)\left[\cos e + \sin e \tan\theta \frac{\cos\Phi\, E_2(i) - \sin^2(\Phi/\pi)E_2(e)}{2 - E_1(i) - (\Phi/\pi)E_1(e)}\right] \tag{9}$$

$$\chi(\theta) = \frac{1}{\sqrt{1+\pi\tan^2\theta}} \tag{10}$$

$$E_1(x) = \exp\left(-\frac{2}{\pi}\cot\theta\cot x\right) \tag{11}$$

$$E_2(x) = \exp\left(-\frac{1}{\pi}\cot^2\theta\cot^2 x\right) \tag{12}$$

Shadowing Function

$$\eta(x) = \chi(\theta)\left[\cos x + \sin x \tan\theta \frac{E_2(x)}{2 - E_1(x)}\right] \tag{13}$$

For $i < e$:

$$S = \frac{\mu_e}{\eta(e)}\frac{\mu_0}{\eta(i)}\chi(\theta)\left[1 - f(\Phi) + f(\Phi)\chi(\theta)\frac{\mu_0}{\eta(i)}\right] \tag{14}$$

For $i \geq e$:

$$S = \frac{\mu_e}{\eta(e)}\frac{\mu_0}{\eta(i)}\chi(\theta)\left[1 - f(\Phi) + f(\Phi)\chi(\theta)\frac{\mu_0}{\eta(e)}\right] \tag{15}$$

$$f(\Phi) = \exp\left(-2\tan\frac{\Phi}{2}\right) \tag{16}$$

References

Barron, J. T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P. P., 2021. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. Proceedings of the IEEE/CVF international conference on computer vision, 5855–5864.

Bi, S., Xu, Z., Srinivasan, P., Mildenhall, B., Sunkavalli, K., Hasan, M., Hold-Geoffroy, Y., Kriegman, D., Ramamoorthi, R., 2020. Neural Reflectance Fields for Appearance Acquisition. arXiv preprint arXiv:2008.03824.

Boss, M., Braun, R., Jampani, V., Barron, J. T., Liu, C., Lensch, H. P. A., 2021. NeRD: Neural Reflectance Decomposition from Image Collections. ICCV, 12684–12694.

Deng, K., Liu, A., Zhu, J.-Y., Ramanan, D., 2022. Depth-supervised nerf: Fewer views and faster training for free. CVPR, 12882–12891.

Derksen, D., Izzo, D., 2021. Shadow neural radiance fields for multi-view satellite photogrammetry. CVPRW, 1152–1161.

Hapke, B., 1981. Bidirectional reflectance spectroscopy: 1. Theory. J. Geophys. Res. Solid Earth, 86(B4), 3039–3054.

Labarre, S., Ferrari, C., Jacquemoud, S., 2017. Surface roughness retrieval by inversion of the Hapke model: A multiscale approach. Icarus, 290, 63–80.

Labarre, S., Jacquemoud, S., Ferrari, C., Delorme, A., Derrien, A., Grandin, R., Jalludin, M., Lemaitre, F., Métois, M., Pierrot-Deseilligny, M. et al., 2019. Retrieving soil surface roughness with the Hapke photometric model: Confrontation with the ground truth. Remote Sens. Environment, 225, 1–15.

Lucht, W., Schaaf, C. B., Strahler, A. H., 2000. An algorithm for the retrieval of albedo from space using semiempirical BRDF models. IEEE Trans. on Geosci. Remote sensing, 38(2), 977–998.

Marí, R., Facciolo, G., Ehret, T., 2022. Sat-NeRF: Learning multi-view satellite photogrammetry with transient objects and shadow modeling using RPC cameras. CVPRW, 1311–1321.

Marí, R., Facciolo, G., Ehret, T., 2023. Multi-date earth observation nerf: The detail is in the shadows. CVPRW, 2034–2044.

Masquil, E., Marí, R., Ehret, T., Meinhardt-Llopis, E., Musé, P., Facciolo, G., 2025. S-eo: A large-scale dataset for geometry-aware shadow detection in remote sensing applications. Proceedings of the Computer Vision and Pattern Recognition Conference, 2383–2393.

Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., Ng, R., 2020. NeRF: Representing scenes as neural radiance fields for view synthesis. ECCV, 99–106.

Minnaert, M., 1941. The reciprocity principle in lunar photometry. Astrophys. J., 93, 403–410.

Nguyen, D., Jacquemoud, S., Lucas, A., Douté, S., Ferrari, C., Coustance, S., Marcq, S., Meygret, A., 2025. Mapping the surface properties of the Asal-Ghoubbet rift by massive inversion of the Hapke model on Pleiades multiangular images. Remote Sensing of Environment, 322, 114691.

Nguyen, T. D., 2024. Photometric modelling to characterize the planetary surfaces in the solar domain. PhD thesis, Université Paris Cité.

Pierrot-Deseilligny, M., Paparoditis, N., 2006. A multiresolution and optimization-based image matching approach: An application to surface reconstruction from spot5-hrs stereo imagery. ISPRS Workshop On Topographic Mapping From Space, 36(1).

Pinty, B., Verstraete, M. M., 1991. Extracting information on surface properties from bidirectional reflectance measurements. J. Geophys. Res. Atmospheres, 96(D2), 2865–2874.

Rahman, H., Pinty, B., Verstraete, M. M., 1993. Coupled surface-atmosphere reflectance (CSAR) model: 2. Semiempirical surface model usable with NOAA advanced very high resolution radiometer data. J. Geophys. Res. Atmospheres, 98(D11), 20791–20801.

Souchon, A., 2012. Influence des phases amorphes dans la réponse optique des régolites planétaires: caractérisation des propriétés physiques et application à l’étude géologique de la Lune. PhD thesis, Université Paul Sabatier-Toulouse III.

Srinivasan, P. P., Deng, B., Zhang, X., Tancik, M., Mildenhall, B., Barron, J. T., 2021. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. CVPR, 7495–7504.

Verbin, D., Hedman, P., Mildenhall, B., Zickler, T., Barron, J. T., Srinivasan, P. P., 2022. Ref-NeRF: Structured View-Dependent Appearance for Neural Radiance Fields. CVPR.

Walter, B., Marschner, S. R., Li, H., Torrance, K. E., 2007. Microfacet models for refraction through rough surfaces. Proceedings of the 18th Eurographics conference on Rendering Techniques, 195–206.

Zhang, L., Rupnik, E., 2023. SparseSat-NeRF: Dense Depth Supervised Neural Radiance Fields for Sparse Satellite Images. ISPRS Ann.

Zhang, L., Rupnik, E., Nguyen, T. D., Jacquemoud, S., Klinger, Y., 2025. BRDF-NeRF: Neural radiance fields with optical satellite images and BRDF modelling. International Journal of Applied Earth Observation and Geoinformation, 143, 104747.

Zhang, X., Srinivasan, P. P., Deng, B., Debevec, P., Freeman, W. T., Barron, J. T., 2021. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. ACM Trans. on Graph. (ToG), 40(6), 1–18.