

Exoplanet surface mapping with composite light reflection models

Sweder Blanken

Delft University of Technology

Exoplanet surface mapping with composite light reflection models

by

Sweder Blanken

to obtain the degree of Bachelor of Science
at the Delft University of Technology,
to be defended publicly on August 15th at 10:00 CEST.

Student number: 5082129
Project duration: March 1, 2022 – July 15th, 2022
Thesis committee: Dr. P. Visser, TU Delft, supervisor
Dr. ir. A.J.L. Adam, TU Delft, supervisor
Dr. A. Endo, TU Delft
Dr. J.T. van Essen, TU Delft

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

On the cover: Artist's impression of a rocky exoplanet orbiting a red dwarf star. Credits: NASA, ESA, STScI.

Abstract

In the near future, next-generation telescopes will be able to observe Earth-like exoplanets illuminated by their parent stars for long periods of time. As the distance between exoplanet and observer is enormous, exoplanets will only make up a pixel on our image. However, the observed intensity of this pixel will fluctuate over time as the exoplanet rotates about its axis and orbits around its parent star. These fluctuations in the reflected light contain information about the planet's surface.

Previous researchers have developed a retrieval method known as spin-orbit tomography which uses these intensity fluctuations to construct a surface map of the exoplanet assuming fully diffuse reflection as a model for the reflection of starlight. In this thesis, we aim to build on this method by introducing two new reflection models, namely Fresnel and Lommel-Seeliger reflection, and by using this composite reflection model for the starlight we attempt to reconstruct the surfaces of exoplanets based on their observed light curve. We will derive an analytical expression for the observed intensity of the light reflected off an exoplanet. Next, we will simulate intensity observations along the orbit of an exoplanet with a pseudo-randomly generated surface containing terrain types found on Earth. As we assign a measure of reflectance for each reflection model, we will construct three distinct reflectance maps for one planet, and simulate the observed light curve by linearly transforming these maps. Afterward, we attempt to retrieve the reflectance maps from the observed signal by inverting the transformation.

We show that the retrieval of all three maps is successful if the light curve is observed without noise, even when the number of samples is low. If artificial shot noise is imposed on our signal, we are still able to retrieve glossy reflectance maps for Lambertian and Lommel-Seeliger reflection by truncating the transformation matrix. Lastly, we show that even if the original light signal originates from an exoplanet that only reflects according to the Lambertian model, then the composite reflection model still retrieves the Lambertian map accurately. On the other hand, if the original light curve is simulated by the composite reflection model, then the Lambertian method does not retrieve an accurate surface map.

*Sweder Blanken
Delft, August 2022*

Contents

1	Introduction	1
2	Light reflection models	4
2.1	Coordinate system	4
2.2	Bidirectional reflectance distribution functions	6
2.2.1	Lambertian reflection	7
2.2.2	Fresnel reflection	8
2.2.3	Lommel-Seeliger reflection	8
2.3	Phase functions	10
2.4	Normalisation of the BRDFs	12
3	Surface maps generation	14
3.1	Altitude maps	14
3.1.1	Tetrahedral subdivision	14
3.2	Terrestrial surface types	16
3.3	Albedo assignment and reflectance distribution	17
4	Reflected light simulation	20
4.1	Analytical expression for the reflected light curve	20
4.2	Numerical expression for the reflected light curve	21
4.2.1	Daily rotation of planet	21
4.2.2	Structure of the transformation matrix	21
4.3	Application to generated surface maps	23
4.3.1	Noiseless light curves	23
4.3.2	Light curves with shot noise	24
5	Surface map retrieval	28
5.1	Inversion of the transformation matrix	28
5.1.1	Moore-Penrose pseudo-inverse	28
5.1.2	Singular value decomposition and cutoff ratio	29
5.2	Noiseless retrieval	30
5.3	Retrieval with noise	36
5.4	Comparing the reflection models	40
5.4.1	Retrieval for composite light curves	40
5.4.2	Retrieval for simple Lambertian light curves	40
5.4.3	Relative error comparison	44
6	Conclusion	49
	References	51
A	Appendix	53
A.1	Derivation of Lambertian phase function.	53
A.2	Tetrahedral subdivision algorithm	54
A.3	Code	54

Nomenclature

Abbreviations

Abbreviation	Definition
BRDF	Bidirectional reflectance distribution function
SVD	Singular value decomposition
SVCR	Singular value cutoff ratio
RM	Reflection model
La	Lambertian reflection
Fr	Fresnel reflection
LS	Lommel-Seeliger reflection

Symbols

Symbol	Definition	Unit
$\hat{\mathbf{i}}$	Unit vector from planet to star	
$\hat{\mathbf{o}}$	Unit vector from planet to observer	
$\hat{\mathbf{s}}$	Unit vector from center of planet to point on surface	
$\hat{\mathbf{n}}$	Axial tilt vector of planet	
α	Equinox angle of axial tilt	
β	Obliquity of axial tilt	
γ	Orbital plane tilt	
ρ	Radius of exoplanet	[m]
R	Radius of circular orbit of planet [m]	
ζ	Phase angle of exoplanet	[rad]
ω	Orbital angular velocity	[rad]
Φ	Longitude of planet	[rad]
Θ	Latitude of planet	[rad]
Ω_0	Solid angle of observer	[sr]
\mathcal{D}	Region of the exoplanet's surface that is both illuminated and visible to the observer	
R_z	Rotation about z-axis in orbital system	[rad]
R_y	Rotation about y-axis in orbital system	[rad]
$\hat{\mathbf{s}}'$	rotated unit vector $\hat{\mathbf{s}}$	
θ_i	Zenith angle of incoming light	[rad]
ϕ_i	Azimuth angle of incoming light	[rad]
θ_r	Zenith angle of reflected light	[rad]
ϕ_r	Azimuth angle of reflected light	[rad]
f	Bidirectional reflectance distribution function	
F	Phase function	
P_*	Total stellar luminosity	[W]
a	Model-dependent surface reflectance	
S	Surface type	

Symbol	Definition	Unit
N_{Φ}	Number of longitude facets	
N_{Θ}	Number of latitude facets	
I	Reflected luminous intensity	[cd]
I	Sampled intensity curve	[cd]
M	Reflectance map	
T	Transformation matrix mapping reflectance map to intensity curve	
T⁺	Moore-Penrose pseudoinverse	

Introduction

On October 6 1995, Michel Mayor and Didier Queloz discovered the world's first exoplanet, named Dimidium, using Doppler spectroscopy (Mayor and Queloz, 1995). This achievement earned them the Nobel Prize in Physics in 2019. Since their discovery, the field of exoplanet detection has proven to be very successful, and as of June 2022, over 5,000 exoplanets have been detected in more than 3,700 planetary systems, as shown in Figure 1.1. One of the main motives to study exoplanets is humanity's search for extraterrestrial life. When studying a terrestrial exoplanet, there are a number of signs that indicate the possibility, if not the existence of life on that planet. These include the presence of oxygen or other molecules vital to life in its atmosphere and the existence of liquid oceans or vegetation on its surface. If we are able to construct a map of the planet's surface we might determine its livability, and maybe decide to send probes once interstellar travel becomes possible. This field is known as Exo-cartography, and this thesis will primarily focus on the reflectance or albedo of different terrain types to compose surface maps.




Figure 1.1: Cumulative detection of exoplanets per year (NASA, 2022).

As shown in Figure 1.1, most of the currently known exoplanets are identified through indirect methods of detection like Doppler spectroscopy¹ or transit photometry. These indirect methods analyze the changing light signal from a star to deduce if an orbiting exoplanet exists. However, these detection methods barely tell us anything about the characteristics of the exoplanet itself, apart from its minimal

mass. On the other hand, only a handful of exoplanets have been directly observed by imaging, as this technique is more difficult. Firstly, the reflected light of an exoplanet is very faint when compared to its bright parent star. Secondly, because other planetary systems are many light-years away from our Solar System, if we do observe an exoplanet, they rarely make up more than a pixel on our image, which makes us unable to work out the layout of its surface. Even with the rapid development of next-generation telescopes, both on the Earth (e.g. the Extremely Large Telescope (Gilmozzi and Spyromilio, 2007)), as well as space telescopes (e.g. the recently launched James Webb Space Telescope (Gardner et al., 2006)), this enormous distance between exoplanet and observer renders traditional methods of directly imaging exoplanets almost completely useless. Therefore constructing a surface map of terrestrial exoplanets seems to be beyond our current technological capabilities. Luckily, over the last two decades a technique has been developed to construct surface maps of exoplanets that only uses the observed intensity of the light and hence doesn't require direct imaging.

At first glance, there seems little to discern from one pixel of information. However, the brightness of this pixel changes depending on the part of the planet that is illuminated, which is in turn dependent on the position of the planet with respect to its parent star. Furthermore, like the planets in our own solar system², exoplanets tend to have some intrinsic spin, because they rotate about an axis. This means that during different periods of the exoplanets' day different regions of the exoplanet are illuminated. If one hemisphere of the planet is covered with fresh snow, and the other with liquid water, then the light reflected off the planet is much brighter if the snowy side faces the star, which affects the intensity of the observed light curve. Moreover, if the planet's axial tilt (or obliquity) is non-zero, the section of the planet's surface that is illuminated during a day changes over the exoplanet's year³. More precisely, different longitudes of the surface are illuminated during one planet-day, while different latitudes of the surface are illuminated during the seasons over one planet-year. By repeatedly observing the reflected light from a terrestrial exoplanet during its orbit, one should be able to use the light signal to retrieve a two-dimensional map of its surface. This inversion method is called 'spin-orbit tomography' first introduced by Kawahara and Fujii (2010), and is a fairly new technique in exoplanet research. Some authors (see Cowan and Fujii (2017), Fujii and Kawahara (2012), and Meinke, Stam, and Visser (2022)) started applying the technique by generating random albedo maps of exoplanets (or by using that of the Earth) before simulating its reflected light curve and retrieving the albedo maps through mock observations of the reflected starlight.

Up to now, most research about spin orbit tomography has assumed totally diffusive reflection, otherwise known as Lambertian reflection, to simulate the light reflecting off the exoplanet's surface. In that case, the perceived brightness is invariant to the observer's angle of view, and light scatters equally in all directions off the surface. However, not all types of terrain reflect light the same way. For example, large bodies of water reflect light more specularly or mirror-like at near-grazing incidences, in accordance with the Fresnel equations. Furthermore, the observed light reflected of rocky heavenly bodies such as the moon and Mars appear significantly brighter at low phase angles than predicted by Lambertian reflection. So unlike diffuse reflection, these models depend on the angle of incidence and reflection off the surface.

In this thesis, we attempt to build further on the method of spin-orbit tomography for edge-on observations introduced by Stuger (2021) by introducing two new reflection models, namely Fresnel reflection to account for the mirror-like reflection of water, and Lommel-Seeliger reflection to account for empirical observations of rocky bodies at low phase angles. Effects due to the presence of an atmosphere, such as Rayleigh scattering or overcast, are not incorporated in our model. As a result of this composite reflection, we will make three separate reflectance maps for each planet with each map showing the measure of reflectance for a particular model. We will then investigate if the implementation of these models in our composite simulation makes the retrieval of our planet's surface easier or more difficult by comparing the maps retrieved from both simple and composite light curves.

In Chapter 2 the planetary system used for model purposes along with the necessary parameters are described. Furthermore the different reflection models used for simulation the reflected light off our exoplanets are introduced. Finally, an analytical expression for the observed light curve is formulated.

¹Also known as the Radial Velocity method.

²For example, Mercury completes one full rotation in under 59 Earth-days, while Jupiter does the same in less than 10 hours (Cox, 2015).

³Earth's mean obliquity is about 23 degrees, which causes the seasons (Williams, 1993).

In Chapter 3, we generate pseudo-random maps similar to Earth's surface using a method called tetrahedral subdivision. Each different terrain type on the map is then attributed a measure of reflectance for each of the models. Combining the results of these Chapters, in Chapter 4 the analytical expression for the light curve will be discretized to simulate the reflected light curve we would detect from edge-on observations of irregular planets. When simulating we can also choose to disturb the signal with artificial noise. Finally, in Chapter 5 we attempt to retrieve the surface maps from the light-curves sampled from edge-on observations, both with and without noise. When doing this, we've assumed we knew a priori the exoplanet's axial tilt. Lastly, we evaluate if the assumption of different reflection models improves or impedes our retrieval efforts. Conclusions are drawn in Chapter 6.

2

Light reflection models

In this chapter, the reflection models are introduced that will be used in our research to simulate the observed light curve from an exoplanet, beginning with Lambertian reflection. To account for the mirror-like reflective behaviour of oceans, we include the Fresnel reflection model, and to account for the bright observations of rocky bodies at low phase angles, we include the Lommel-Seeliger reflection model. Firstly, we will describe the orbital system of the planet that we would like to observe, and introduce some useful vector notation. Secondly, we'll introduce the different reflection models using their characteristic reflectance function, known as a BRDF. Next we use these BRDFs to simulate the light curve for a very simplified version of our planetary model. Lastly, we investigate if the introduced reflection models are physically realistic.

2.1. Coordinate system

To model the reflected starlight off an exoplanet, we consider a simple planetary system, in which a single planet with radius ρ and no moons is directly observed as it moves in a circular orbit of radius R around its parent star with orbital angular frequency ω . Furthermore, the exoplanet also rotates about its axis with frequency Ω . We define the orbital plane as the xy -plane such that the origin of our system coincides with the centre of the star, the planet orbits counter-clockwise when viewed from above, and light rays directed to the observer travel parallel to the x -axis. A schematic of this setup can be viewed in Figure 2.1. When talking about incoming rays of light reflecting off the different facets the planet, it's easy to get confused by all the different angles. Therefore we introduce the following unit vectors,

$$\hat{\mathbf{i}} = \begin{bmatrix} \cos \zeta \\ \sin \zeta \\ 0 \end{bmatrix}, \hat{\mathbf{o}} = \begin{bmatrix} \sin \gamma \\ 0 \\ \cos \gamma \end{bmatrix}, \hat{\mathbf{s}} = \begin{bmatrix} \cos \Phi \sin \Theta \\ \sin \Phi \sin \Theta \\ \cos \Theta \end{bmatrix}, \hat{\mathbf{n}} = \begin{bmatrix} \cos \alpha \sin \beta \\ \sin \alpha \sin \beta \\ \cos \beta \end{bmatrix}, \quad (2.1)$$

where $\hat{\mathbf{i}}$ and $\hat{\mathbf{o}}$ point from the center of the exoplanet to the parent star and to the observer, respectively. Here ζ is the phase angle of observation, and γ denotes the tilt of the orbital plane with respect to the viewing plane. If $\gamma = \frac{\pi}{2}$, then the observer is in the orbital plane, and we speak of edge-on observation. Moreover, $\hat{\mathbf{s}}$ is a unit vector pointing to a facet of the planet's surface defined by spherical coordinates Φ and Θ , and $\hat{\mathbf{n}}$ denotes the orientation of the rotational axis of the planet, where α is the equinox angle and β is the planet's obliquity. These vectors and angles are shown in more detail in Figure 2.2.

When studying the incoming and reflecting rays of light from the star, it may be more useful to define $\hat{\mathbf{i}}$, $\hat{\mathbf{o}}$, and $\hat{\mathbf{s}}$ in terms of the incident and reflective angles $(\theta_i, \phi_i$ and $\theta_r, \phi_r)$. If we imagine the surface of the planet to be locally flat, and we orient the coordinate system so that the origin lays on the surface and the x -axis is parallel to the projection of $\hat{\mathbf{i}}$, as is shown in Figure 2.3, then we can rewrite the vectors as:

$$\hat{\mathbf{i}} = \begin{bmatrix} \sin \theta_i \\ 0 \\ \cos \theta_i \end{bmatrix}, \hat{\mathbf{o}} = \begin{bmatrix} \cos(\phi_r - \phi_i) \sin \theta_r \\ \sin(\phi_r - \phi_i) \sin \theta_r \\ \cos \theta_r \end{bmatrix}, \hat{\mathbf{s}} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}. \quad (2.2)$$

Note that these alternate definitions of $\hat{\mathbf{i}}$, $\hat{\mathbf{o}}$, and $\hat{\mathbf{s}}$ are only valid for points on the planet's surface that are in the observable domain \mathcal{D} , i.e. both illuminated by the star and visible to the observer (see Figure 2.1 and Section 2.3). These two different ways to denote $\hat{\mathbf{i}}$ and $\hat{\mathbf{o}}$ allow us to define the phase angle ζ in terms of the angles of the incoming and reflective light rays in one direction and θ_i, θ_r and $\psi = \phi_r - \phi_i$ in terms of the original vectors in the other direction. Namely by combining Equations 2.1 and 2.2, we obtain:

$$\begin{aligned} \cos \theta_i &= \hat{\mathbf{i}} \cdot \hat{\mathbf{s}} \\ \cos \theta_r &= \hat{\mathbf{o}} \cdot \hat{\mathbf{s}} \\ \cos \zeta &= \frac{\hat{\mathbf{i}} \cdot \hat{\mathbf{o}}}{\sin \gamma} = \frac{\cos \theta_i \cos \theta_r + \cos(\phi_r - \phi_i) \sin \theta_i \sin \theta_r}{\sin \gamma} \\ \cos \psi &= \frac{(\hat{\mathbf{i}} \cdot \hat{\mathbf{o}}) - (\hat{\mathbf{i}} \cdot \hat{\mathbf{s}})(\hat{\mathbf{o}} \cdot \hat{\mathbf{s}})}{|\hat{\mathbf{i}} \times \hat{\mathbf{s}}| |\hat{\mathbf{o}} \times \hat{\mathbf{s}}|} \\ &= \frac{\cos \zeta - \cos \theta_i \cos \theta_r}{\sin \theta_i \sin \theta_r}, \end{aligned} \quad (2.3)$$

and for edge-on observations $\sin \gamma$ becomes unity.




Figure 2.1: An overview of the observed planet with radius ρ in a circular orbit of radius R around its parent star, defined in the XY-plane. ζ is the time-dependent phase angle of the planet with the observer, and Ω denotes the angular frequency of the planet. P_1 and P_2 are two points in the planets orbit that are further studied in Figure 2.2.




Figure 2.2: Detailed illustrations of the exoplanet in orbit. **(a)** xy-projection of the planet at point P_1 , where $\hat{\mathbf{n}}$ denotes the rotational axis' orientation, and α its equinox angle. The observable domain, denoted as \mathcal{D} , is also shown. **(b)** xz-projection of the planet at point P_2 , showing the obliquity β and the orbital inclination γ .




Figure 2.3: By reorienting the coordinate system we can express $\hat{\mathbf{i}}$, $\hat{\mathbf{o}}$ and $\hat{\mathbf{s}}$ in terms of the optical angles.

2.2. Bidirectional reflectance distribution functions

A bidirectional reflectance distribution function (or BRDF) is a measure of how much light reflects off different types of opaque surfaces from one angle to another (Nicodemus, 1965). It returns the ratio of radiant flux reflected by a given surface per unit solid angle per unit projected area (or radiance) to the radiant flux per unit surface area (or irradiance). In the most general case, it is an angle-dependent function $f(\theta_i, \phi_i; \theta_r, \phi_r)$, where the first two input angles define the incoming light direction and the latter two define the outgoing light direction from a surface. In reality, the BRDF also depends on the wavelengths of the incoming and reflected light, but this spectral dependence exceeds the scope of this thesis¹. A BRDF f can be formally defined as:

$$f(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{dL(\theta_r, \phi_r)}{dE(\theta_i, \phi_i)}, \quad (2.4)$$

where L denotes the radiance of the light rays (which corresponds to the brightness of the reflected light from surface to detector) and E denotes the irradiance of the incident light. The two terms are

related by the following equation (McCluney, 2014):

$$E = \int_{\Upsilon} L(\theta, \phi) \cos(\theta) d\Omega, \quad (2.5)$$

where Υ denotes a finite solid angle.

All physically realistic BRDFs have some fundamental properties (Duvenhage, Bouatouch, and Kourie, 2013). These are:

1. Positivity: $f(\theta_i, \phi_i; \theta_r, \phi_r) \geq 0$.
2. Helmholtz reciprocity: $f(\theta_i, \phi_i; \theta_r, \phi_r) = f(\theta_r, \phi_r; \theta_i, \phi_i)$.
3. Energy conservation: for all incoming light directions, $\int_{\Omega} f(\theta_i, \phi_i; \theta_r, \phi_r) \cos(\theta_r) d\omega_r \leq 1$.

BRDFs and their reflection models are used to realistically simulate or study the surface appearance of various materials. Domains where these models are used include astronomy, computer graphics, and military and commercial applications. In some situations BRDFs are amplified or empirical models are used when similarity to observed data or visual aesthetics are preferred over a theoretical derivation.

2.2.1. Lambertian reflection

One of the most common reflection models is that of Lambertian reflection, which assumes perfectly diffuse surfaces, i.e. the illuminated surface scatters the light evenly in all directions, and the reflected radiance is independent of the viewing angles θ_r and ϕ_r . This also implies its independence of any incident angle due to Helmholtz reciprocity. Therefore the BRDF for this model (from now on denoted as f_{La}) is simply a constant:

$$f_{La}(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{1}{\pi}. \quad (2.6)$$

Figure 2.4 shows a 3-dimensional plot of the Lambertian reflection BRDF on an opaque surface for different values of θ_r and ϕ_r . For this plot, $\theta_i = \frac{\pi}{3}$ and $\phi_i = \frac{\pi}{4}$.




Figure 2.4: Bidirectional reflectance distribution function for Lambertian reflection, where $\theta_i = \frac{\pi}{3}$ and $\phi_i = \frac{\pi}{4}$. As the BRDF is a constant we get a horizontal plane with value $\frac{1}{\pi}$.

As expected, the Lambertian BRDF remains invariant for different reflecting angles, as the plotted BRDF shows just a horizontal plane with height $\frac{1}{\pi} \approx 0.32$.

Lambertian reflection is one of the most frequently used diffusive reflection models, and the majority of earlier research about exocartography (Fujii and Kawahara (2012), Fan et al., 2019 among others) has used this model as a basis of the light reflecting from terrestrial exoplanets (Fujii and Kawahara (2012), Fan et al., 2019 among others). However, not all surface types reflect light completely diffusely. For instance, if liquid oceans of water are present on the observed planet, then water is reflected more specularly on those parts of the surface, especially at near-grazing incidence. Furthermore, the reflected light from observed heavenly bodies such as the moon and Mars appears much brighter at

¹ See Ferrero et al. (2012) for research on spectral BRDFs.

the limb of the planet than predicted by simple Lambertian reflection (Fairbairn, 2005), therefore we include two other reflection models in our simulation that address each of these problems, namely Fresnel reflection and Lommel-Seeliger reflection.

2.2.2. Fresnel reflection

The Fresnel reflection model is used for mirror-like reflections. A purely specular reflector has BRDF zero unless $\theta_r = \theta_i$ and $\phi_r = \phi_i + \pi$, leading to the BRDF of:

$$f_{Me}(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{\delta(\theta_r - \theta_i) \delta(\phi_r - \phi_i - \pi)}{\cos \theta_r \sin \theta_r}, \quad (2.7)$$

here the subscript *Me* stands for 'metallic' reflection, and δ denotes the (finite) Dirac delta function:

$$\delta(x) = \begin{cases} 1, & x = 0 \\ 0, & x \neq 0. \end{cases} \quad (2.8)$$

This BRDF describes the reflection off perfectly reflecting opaque surfaces. In practice, semi-transparent surfaces like water transmit some of the incident light. In that case the BRDF is adjusted to:

$$f_{Fr}(\theta_i, \phi_i; \theta_r, \phi_r) = R(\theta_i) \frac{\delta(\theta_r - \theta_i) \delta(\phi_r - \phi_i - \pi)}{\cos \theta_r \sin \theta_r}, \quad (2.9)$$

here $R(\theta_i) = \frac{1}{2}(R_s + R_p)$ denotes the Fresnel power reflection coefficient for natural light, and

$$R_s(\theta_i) = \left| \frac{\cos \theta_i - \sqrt{n^2 - \sin^2 \theta_i}}{\cos \theta_i + \sqrt{n^2 - \sin^2 \theta_i}} \right|^2 \quad (2.10)$$

and

$$R_p(\theta_i) = \left| \frac{n^2 \cos \theta_i - \sqrt{n^2 - \sin^2 \theta_i}}{n^2 \cos \theta_i + \sqrt{n^2 - \sin^2 \theta_i}} \right|^2 \quad (2.11)$$

denote the reflectance for s-polarized and p-polarized light, respectively. These Fresnel coefficients are also where the subscript *Fr* comes from. Note that R_s and R_p are at their maximum² when $\theta_i = \frac{\pi}{2}$, so that the rays of light barely graze the surface. Furthermore, n denotes the refractive index of the reflecting surface. In our case this is water, hence $n \approx 1.33$. In our simulation, the Dirac delta in equation 2.9 is replaced by a narrow Gaussian with variance 0.01 (unless specified otherwise) to prevent the BRDF from always returning zero. Since we only simulate a finite number of different surface facets and reflective angles, it could happen that for instance θ_i and θ_r are never exactly the same. Introducing this Gaussian means our modeled reflected light becomes more glossy (and this glossiness depends on the variance σ^2 of the Gaussian), which might resemble the empirical light-reflective properties of large bodies of water.

Figure 2.5 shows a 3-dimensional plot of the Fresnel reflection BRDF for a purely specular reflector for different values of θ_r and ϕ_r , while Figure 2.6 does the same for reflection off water. For these plots, $\theta_i = \frac{\pi}{3}$ and $\phi_i = \frac{\pi}{4}$. As we can see from the figure, the BRDFs for Fresnel reflection show a high peak when $\theta_r = \theta_i$ and $\phi_r = \phi_i + \pi$ and zero reflectance everywhere else.

2.2.3. Lommel-Seeliger reflection

As discussed earlier, the reflected light we observe from terrestrial bodies in our Solar system with regolith on the surface appear significantly brighter than what the simple Lambertian model would predict when the phase angle ζ of observation is small. This phenomenon is called opposition surge or the Seeliger effect, and the corresponding empirical reflection model is known as Lommel-Seeliger reflection as described by Seeliger (1887) and Fairbairn (2005). It's BRDF is

$$f_{LS}(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{2}{3\pi^2} \frac{\sin \zeta + (\pi - \zeta) \cos \zeta}{\cos \theta_i + \cos \theta_r}, \quad (2.12)$$

²The analytical global maxima of the Fresnel coefficient $R_s(\theta_i)$ are actually at $\theta_i = (2k + 1)\pi$, $k \in \mathbb{Z}$, but the incident zenith angle θ_i cannot exceed $\frac{\pi}{2}$.




Figure 2.5: Bidirectional reflectance distribution function for perfect specular reflection, where $\theta_i = \frac{\pi}{3}$ and $\phi_i = \frac{\pi}{4}$. The plot shows a narrow peak when $\theta_r = \theta_i$ and $\phi_r - \phi_i = \pi$.




Figure 2.6: Bidirectional reflectance distribution function for specular reflection off water described by the Fresnel coefficients, where $\theta_i = \frac{\pi}{3}$ and $\phi_i = \frac{\pi}{4}$. The plot also shows a narrow peak when $\theta_r = \theta_i$ and $\phi_r - \phi_i = \pi$. Although similar to Figure 2.5, notice the large difference in the height of the peak, also visible in the colorbar.

where ζ depends on the incident angles as shown in Equation 2.3. Figure 2.7 shows a 3-dimensional plot of the Lommel-Seeliger reflection BRDF for different values of θ_r and ϕ_r . As the Lommel-Seeliger behaves especially interestingly when the planet is in opposition to the Sun. We assume $\theta_i = 0$ and $\phi_i = \frac{\pi}{4}$ in this plot. We find that the Lommel-Seeliger reflectance is highest for small angles of θ_r . This was expected, as empirical data shows the Seeliger effects occurs near normal incidence. However it is not known why Figure 2.12 shows a maximum for θ_i . Lastly, the figure shows that the Lommel-Seeliger BRDF is invariant for different values of ϕ_r , much like the Lambertian BRDF.




Figure 2.7: Bidirectional reflectance distribution function for Lommel-Seeliger reflection, where $\theta_i = 0$ and $\phi_i = \frac{\pi}{4}$. The reflectance drops heavily as θ_r , but remains unchanged for varying ϕ_r .

2.3. Phase functions

In order to further study these reflection models in the context of reflected light from exoplanets in orbit, we first consider a simplified version of our simulation. Imagine a homogeneous spherical planet with radius ρ and orbital radius R that scatters all incident light according to some reflection model RM , in a circular orbit of radius R around its parent star, which emits light in all directions. Furthermore, this planet is observed edge-on, so $\gamma = \frac{\pi}{2}$. If the total power or luminosity of the starlight is P_* , then the power of the light that reached the exoplanet is $\frac{P_*}{4\pi R^2}$. Sadly, not all of this intercepted light is observable. Firstly, only one hemisphere of the planet is illuminated by the star at all times. Secondly, only one hemisphere of the planet is visible to us at all times. The intersection of these regions, namely the part of the planet's surface that is both illuminated and visible to us, is the surface region that actually contributes to the light-curve we observe, and is therefore called the observable domain, henceforth denoted as \mathbb{D} :

$$\mathbb{D} = \left\{ \hat{\mathbf{s}} : \hat{\mathbf{s}} \cdot \hat{\mathbf{i}} > 0 \wedge \hat{\mathbf{s}} \cdot \hat{\mathbf{o}} > 0 \right\}. \quad (2.13)$$

The fraction of received stellar power that's 'intercepted' by the planet is proportional to the frontal surface, given by $(\hat{\mathbf{s}} \cdot \hat{\mathbf{i}})d^2\hat{\mathbf{s}}$, where $d^2\hat{\mathbf{s}} = \rho^2 \sin\Theta d\Phi d\Theta$ is a small part of the planet's surface. The fraction of reflected stellar power that reaches the observer is proportional to $(\hat{\mathbf{s}} \cdot \hat{\mathbf{o}})f_{RM}d^2\Omega_o$, where f_{RM} is the reflectance distribution function by which the light is modelled, and $d^2\Omega_o$ is the solid angle of the observer. Multiplying all these terms gives the contribution of the light reflected a small part of the surface area:

$$\frac{P_*}{4\pi R^2} (\hat{\mathbf{s}} \cdot \hat{\mathbf{i}})(\hat{\mathbf{s}} \cdot \hat{\mathbf{o}})f_{RM}d^2\hat{\mathbf{s}}d^2\Omega_o. \quad (2.14)$$

Since we are only interested in the intensity of the observed light relative to the total intensity that reached the planet (denoted by $\frac{P_*d^2\Omega_o}{4\pi}$), it is customary to divide by this factor. To obtain the observed light-curve for a given time, we simply take this result and integrate it over the observable domain \mathbb{D} . This shows that the light-curve is simply the prefactor $\frac{1}{R^2}$ multiplied by some function F that is dependent on both the position of the planet in orbit³ (parameterized by the phase angle ζ) as well as the choice of reflection model RM . This new function is called the phase function, and is given by:

$$\begin{aligned} F_{RM}(\zeta) &= 4 \iint_{\mathbb{D}} f_{RM}(\theta_i, \phi_i; \theta_r, \phi_r) (\hat{\mathbf{s}} \cdot \hat{\mathbf{i}})(\hat{\mathbf{s}} \cdot \hat{\mathbf{o}}) d^2\hat{\mathbf{s}} \\ &= 4 \iint_{\mathbb{D}} f_{RM}(\theta_i, \phi_i; \theta_r, \phi_r) \cos(\Phi - \zeta) \cos\Phi \sin^2\Theta d^2\hat{\mathbf{s}} \\ &= 4 \iint_{\mathbb{D}} f_{RM}(\theta_i, \phi_i; \theta_r, \phi_r) \cos\theta_i \cos\theta_r d^2\hat{\mathbf{s}}. \end{aligned} \quad (2.15)$$

The prefactor 4 comes from the fact that, for any phase angle ζ , the integral of the phase function F over a unit sphere should return the surface area of 4π :

$$\oiint_{\circlearrowleft} F(\zeta) d^2\hat{\mathbf{i}} = 4\pi. \quad (2.16)$$

Here $\hat{\mathbf{i}}$ is a unit vector and not the vector pointing to the star. For Lambertian reflection the phase function can be derived analytically, see section A.1 in the Appendix for it's derivation:

$$F_{La}(\zeta) = \frac{8}{3\pi} [\sin\zeta + (\pi - \zeta) \cos\zeta] \quad (2.17)$$

Figure 2.8 shows both the analytical derivation of the Lambertian phase function for different values of ζ as well as a numerical approximation for a planet with unit radius, and 50 discrete surface facets, that reflects all light diffusely. As the plot shows, even with a small number of distinct surface facets, the numerical approximation fits the analytical curve very well. The phase functions of all our reflection models for different values of ζ for a totally reflecting planet with unit radius and 20.000 discrete surface facets are shown in Figure 2.9.




Figure 2.8: Analytical and numerical phase functions for the Lambertian reflection model. The numerical phase function is computed by modelling the reflected light off a homogeneous planet with only 50 equi-rectangular surface facets, while the analytical function is sampled 100 times over the shown range.




Figure 2.9: Phase functions F_{RM} for the different reflection models for a high-resolution, totally reflecting planet in a circular orbit. Compared to Figure 2.8, the Lambertian phase function is a lot smoother due to the high number of facets.

When studying Figure 2.9, we see that the reflected intensity off the planet according to the Lambertian and Lommel-Seeliger reflection models is highest for small angles of the phase angle ζ as the intersection of the visible and illuminated regions of the earth is biggest, or in other words, when the size of the observable domain \mathfrak{D} is at its maximum, and so we observe the most diffusely reflected light in that instance of the planet's orbit. As we increase ζ , \mathfrak{D} becomes smaller, and so we the perceived intensity of the light curve also decreases. When $\zeta \approx \pi$, the exoplanet passes between the observer and the star, which is known as an astronomical transit, and so we don't observe any reflected light as expected.

On the other hand, the Fresnel reflection light curve remains quite small for small ζ , but increases significantly when ζ approaches π . This can be explained by the fact that for large values of ζ the incident light hits the light at a greater incident zenith angle θ_i , and for phase angles very close to π $\theta_i \approx \frac{\pi}{2}$, and so the Fresnel coefficients R_s and R_p are maximized. However, if ζ is too high, then the region of the planet's surface able to reflect light mirror-like becomes too small, and the observed intensity drops steadily again. This drop-off is believed to be caused by the approximation of our Gaussian,

as a perfect delta function would produce a maximum phase function at $\zeta = \pi$. There appears to be an optimal phase angle around $\zeta = \frac{11\pi}{12}$ where the balance between high Fresnel reflectance and sufficiently large \gg is just right for a maximal phase function F_{Fr} .

Figure 2.9 also shows that the chosen BRDF for Lommel-Seeliger reflection produces a phase function F_{LS} that is consistently smaller than the Lambertian phase function, even for low ζ . This is not congruent with our expectations, as the Lommel-Seeliger reflection model was implemented to account for the brighter observations due to the opposition surge. This could be an indication that the current reflection models are not yet perfectly implemented in our system. Indeed, we still need to verify if the chosen BRDFs are normalised.

2.4. Normalisation of the BRDFs

Now that we've introduced the different reflection models for our light, it's time to investigate if they are physically realistic. Recall that one of the fundamental properties of physical BRDFs is energy conservation, i.e. for all incoming light directions (θ_i, ϕ_i) the surface never reflects more light than is incident, so

$$\int_{\Omega} f_{RM}(\theta_i, \phi_i; \theta_r, \phi_r) \cos(\theta_r) d\omega_r \leq 1. \quad (2.18)$$

We can rewrite this integral A for our simplified simulation with the unit sphere: for all θ_i ,

$$A = \int_0^{\pi/2} \int_0^{2\pi} \sin \theta_r \cos \theta_r f_{RM}(\theta_i, \theta_r, \psi) d\psi d\theta_r \leq 1, \quad (2.19)$$

where $\psi = \phi_r - \phi_i$. Furthermore, if we assume the sphere reflects all light, then the inequality becomes an equality. Because of this, the used BRDFs should first be normalised so that they satisfy the energy conservation property. If we find that the BRDF for a given reflection model cannot satisfy Equation 2.19 for all θ_i , for instance because the model always assumes some light is absorbed, then we choose the normalisation factor to be as large as possible while still satisfying the inequality, i.e. such that the integral ≤ 1 for all θ_i .

Figure 2.10 shows the numerically computed normalisation integral 2.19 of the different BRDFs as a function of θ_i . As we can see, the BRDF of the Lambertian and the purely specular reflection model are normalised since they are constantly 1, save for some steep drop-offs at very low and very high angles of θ_i for the specular model again probably caused by our Gaussian approximation. It follows that their BRDFs do not need an additional normalisation factor.

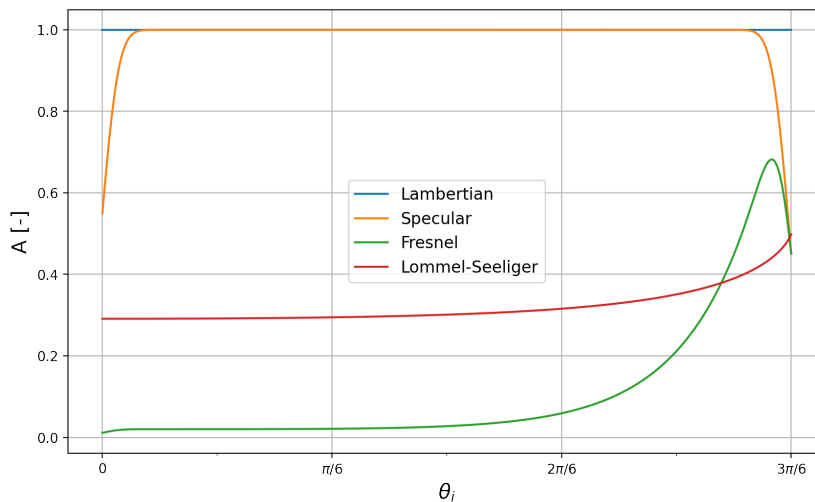


Figure 2.10: Normalisation integral A of the different BRDFs as a function of θ_i .

³Also on its rotation, but we'll discuss that in Chapter 4 for planets with irregular surface maps.

Moreover, from Figure 2.10 we also find that the normalisation integral for Fresnel and Lommel-Seeliger are not constant, but depend on incident angle θ_i . For the Fresnel model, this is caused by the Fresnel coefficients, as those are the only difference between the Fresnel and the purely specular reflection model. For the Lommel-Seeliger model, it is believed that the odd normalisation curve is caused by BDRF's dependence of the planet's phase angle ζ , which is not the case for the other reflection models. We can conclude that empirical model of Lommel-Seeliger reflection is quite hard to simulate, as an analytical derivation for its BRDF has yet to be found. The Fresnel normalisation curve was found to have a maximum of about 0.68, while the Lommel-Seeliger normalisation curve had a maximum close to 0.5. Therefore the proper normalisation factors for those BRDFs should be equal to 1.47 and 2 respectively. It follows that the normalised BRDFs of our considered reflection models are:

$$\begin{aligned}
 f_{La}(\theta_i, \phi_i; \theta_r, \phi_r) &= \frac{1}{\pi} \\
 f_{Fr}(\theta_i, \phi_i; \theta_r, \phi_r) &= 1.47R(\theta_i) \frac{\delta(\theta_r - \theta_i) \delta(\phi_r - \phi_i - \pi)}{\cos \theta_r \sin \theta_r} \\
 f_{LS}(\theta_i, \phi_i; \theta_r, \phi_r) &= \frac{4}{3\pi^2} \frac{\sin \zeta + (\pi - \zeta) \cos \zeta}{\cos \theta_i + \cos \theta_r}.
 \end{aligned} \tag{2.20}$$

Figure 2.11 shows the normalised phase functions for our homogeneous spherical planet. As expected, the Lommel-Seeliger phase function now more closely resembles the Lambertian one. Note that the planet already fully scatters all light, and so the Seeliger effect is minimal in this case. Going forward, when discussing or implementing our BRDFs, we will always use their normalised versions.

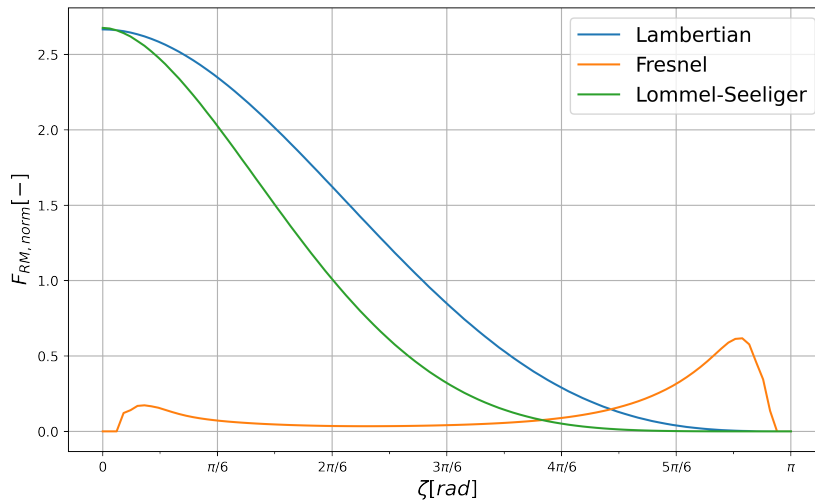


Figure 2.11: Normalised phase function for the totally reflecting unit sphere.

3

Surface maps generation

To find out if we can retrieve surface maps of terrestrial exoplanets from their observed light curve, we first need a sample of planets and their maps to try out our method. As the number of terrestrial planets in our Solar System is small and the current data of light curves of directly observed exoplanets is also not sufficient, we propose to create our own fictional exoplanets, simulate their light curves from their fabricated surface maps, and investigate if we are able to retrieve their surface maps using our method. To generate these surfaces, we first create a pseudo-randomly generated 3D altitude map with finite resolution using a method known as tetrahedral subdivision, as proposed in Mogensen (2009). We then use this altitude map to divide the planet's surface into different types of terrain likely to be found on habitable planets, like vegetation and liquid oceans. For each terrain type, we give a standard albedo as a measure of light reflectance. To incorporate our composite reflection model, we also assign a distribution to each surface type to indicate how much of their reflected light is scattered according to Lambertian reflection, Lommel-Seeliger reflection, or the Fresnel equations. Finally we are able to produce three distinct surface maps per planet, with each map indicating the measure of Lambertian, specular, and Lommel-Seeliger reflection per unit of area.

3.1. Altitude maps

There are some properties about the nature of landscapes one should respect when generating height maps. Firstly, real landscapes are continuous, i.e. small changes in position should give small changes in altitude. This means that picking random heights for each unit surface would bring about large discontinuities and consequently superficial maps. Secondly, landscapes have roughly the same degree of detail at all levels of magnification (Mandelbrot and Mandelbrot, 1982), so the map should be self-similar to a degree. Furthermore, to create a 3-dimensional map for the surface of an exoplanet with a discrete number of facets, we can't simply generate a 2D height map and 'wrap' it around a sphere, as that wrapping would also bring about discontinuities or artefacts in at least one or two points on the sphere, e.g. near the poles. Instead we will create 3-dimensional spherical height maps, that can be projected onto a flat surface for display purposes.

3.1.1. Tetrahedral subdivision

To generate multiple convincing natural maps we need a method to consistently produce random 3D height-maps that meet the aforementioned properties. The method (as described in Mogensen (2009)) used in this thesis is known as tetrahedral subdivision. It returns a pseudo-randomly generated altitude for each point p_j of a spherically distributed set. The algorithm starts by constructing a tetrahedron around the set of points, and assigning an altitude value h_j to each vertex v_j of the tetrahedron, pseudo-randomly generated by a predetermined seed s_j . Next, the longest edge of the tetrahedron is halved by creating a new vertex point, with its own altitude and seed being determined by the altitudes and seeds of its adjacent vertices. Thereafter, the tetrahedron is halved by joining the new vertex with its nonadjacent vertices, thereby creating two new tetrahedra. The process is then repeated for the tetrahedron that contains the point p_j , until the length of the longest edge passes below a predetermined value. The altitude h_j is then calculated by averaging the altitudes of the vertices of the final

tetrahedron. After all points p_j are given an altitude, these values are scaled between 0 and 1. A more detailed description of the subdivision algorithm can be found in section A.2 of the Appendix. A visual representation of algorithm's first step is shown in Figure 3.1. Here the longest edge lies between v_1 and v_2 , so that edge is halved by creating a new vertex v_{new} , thereby forming two new tetrahedra. Furthermore, p_j lies in the tetrahedron constructed by v_1, v_3, v_4 and v_{new} , so for the next step, $v_2 = v_{\text{new}}$.




Figure 3.1: First step of the tetrahedral subdivision to generate random altitude maps. Note that p_j lies on the surface of the unit sphere. In the next step, the tetrahedron constructed by v_1, v_3, v_4 and v_{new} will be further subdivided.

One advantage of this method is that points that lay close to each other on the sphere share the same sequence of tetrahedra up to a certain iteration of the algorithm. Therefore their altitudes will not differ that much when compared to the altitude of a point on the other side of the sphere, which has a completely different sequence of tetrahedra. This result solves the problem of discontinuities at the poles and edges that 2-dimensional height maps would have brought about, and so we can produce continuous landscapes as discussed in section 3.1. Furthermore, although the altitudes of the point on a sphere are all randomly generated, they all fully depend on the original set of four seeds s_j of the vertices of the first tetrahedron. This means that repeating the algorithm always produces the same height map, provided the seeds stay the same. The ability to reproduce the same altitude maps during the span of our research is necessary for making observations and attempting to retrieve the surface maps from their generated light curves. In total we will generate 4 planets with different altitude maps for our research.

Figure 3.2 shows the projected 2-dimensional altitude maps of two generated planets. The used scheme to map the spherical facets to the flat map is by projecting the facets equirectangularly, which preserves the verticality of meridians on the map at the expense of a distortion of the equal areas of the facets, especially towards the poles. Unless otherwise stated, we let the number of different longitudes $N_{\Phi} = 80$ and the number of latitudes $N_{\Theta} = 40$, as a good midway point between a high resolution map and a low computing time. The corresponding Mollweide projections of the maps, which we will use henceforth to visualize our surface maps, are shown on the right. Note that we manually set the altitudes of facets near the poles of the planet to 1 to later emulate snowy polar caps.

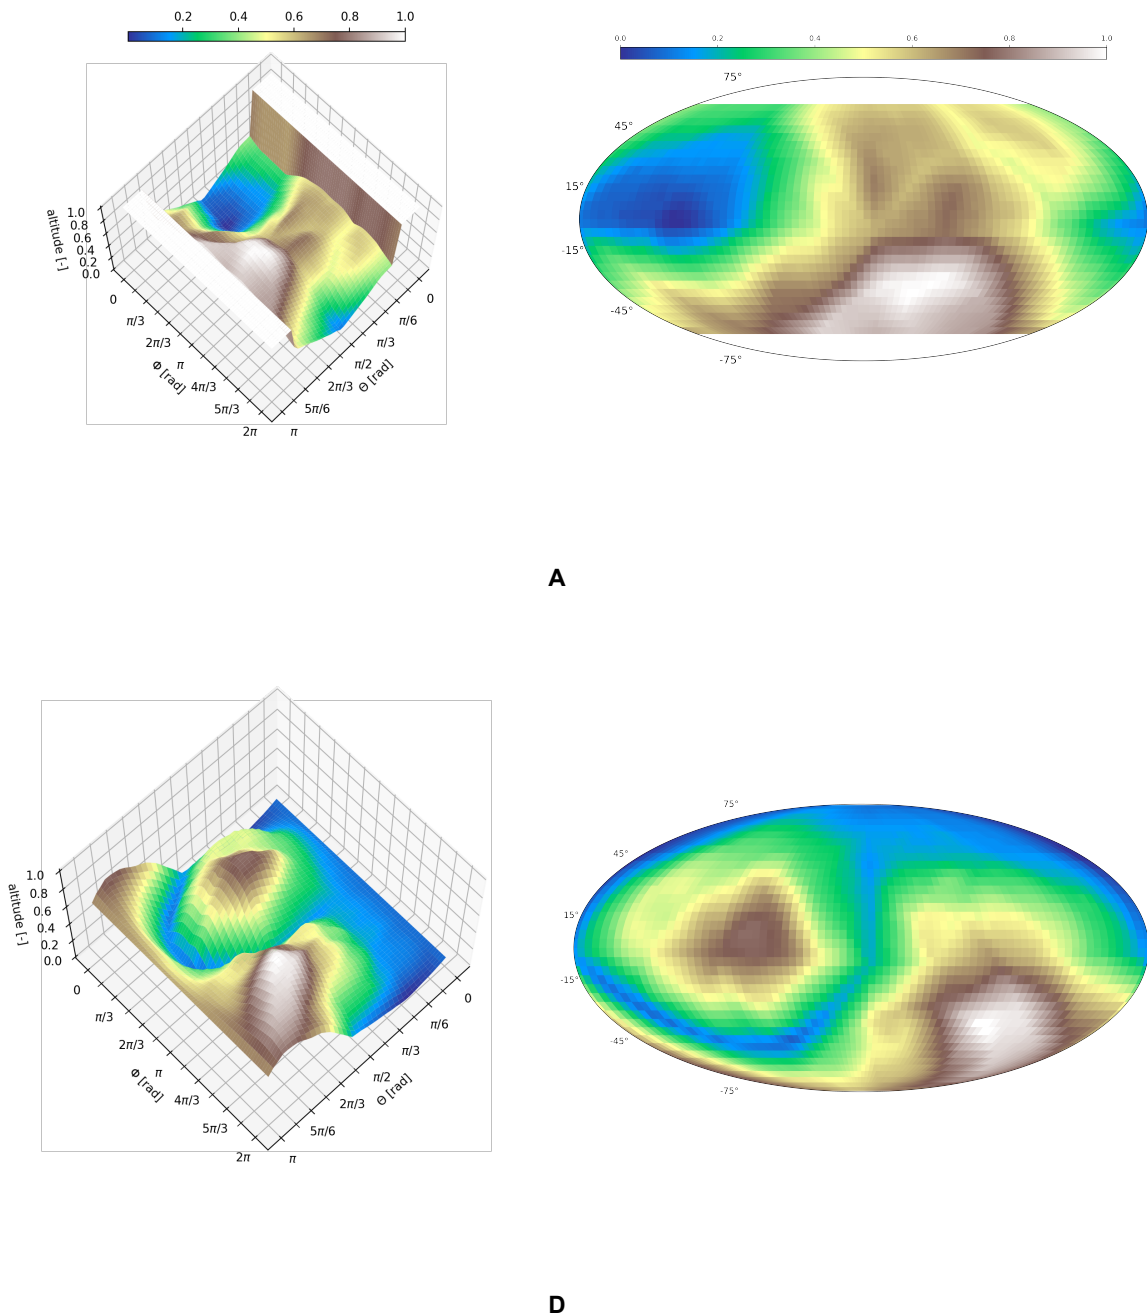


Figure 3.2: Left: the generated altitude maps for planets A and D using tetrahedral subdivision. Right: Their corresponding Mollweide projection.

3.2. Terrestrial surface types

After generating the 3D altitude map, we use these varying altitudes to divide the exoplanet's surface into different types of terrains. The ultimate goal of our research is to see if we can retrieve these terrains accurately from the light reflected off our planet. As the most interesting exoplanets are those where life exists (or could exist), we primarily want to discern surface types akin to habitable terrestrial planets such as our own Earth, especially bodies of liquid water¹. The chosen surface types are, in

increasing order of elevation: ocean, grasslands, forest, regolith, and snow. The order of these terrain types is chosen to imitate the landscape of the Earth: liquid water is attracted by Earth's gravity and flows to the lowest parts of the spheroid's surface, while snowy peaks only occur on the tips of mountain ranges, where temperature drops below the freezing point.

Table 3.1 shows the different terrestrial types S studied in our simulation, combined with the distribution of (normalised) altitude values h_j that are mapped to this surface type. Figure 3.3 shows the Mollweide projections of our 4 planets with added pseudocolors to distinguish the 5 surface types

S	L	U
Ocean	0	0.5
Grasslands	0.5	0.625
Forest	0.625	0.75
Regolith	0.75	0.875
Snow	0.875	1

Table 3.1: List of surface types and the part of the altitude values that are mapped to it. If $L \leq h_j < U$, then p_j is mapped to surface type S .

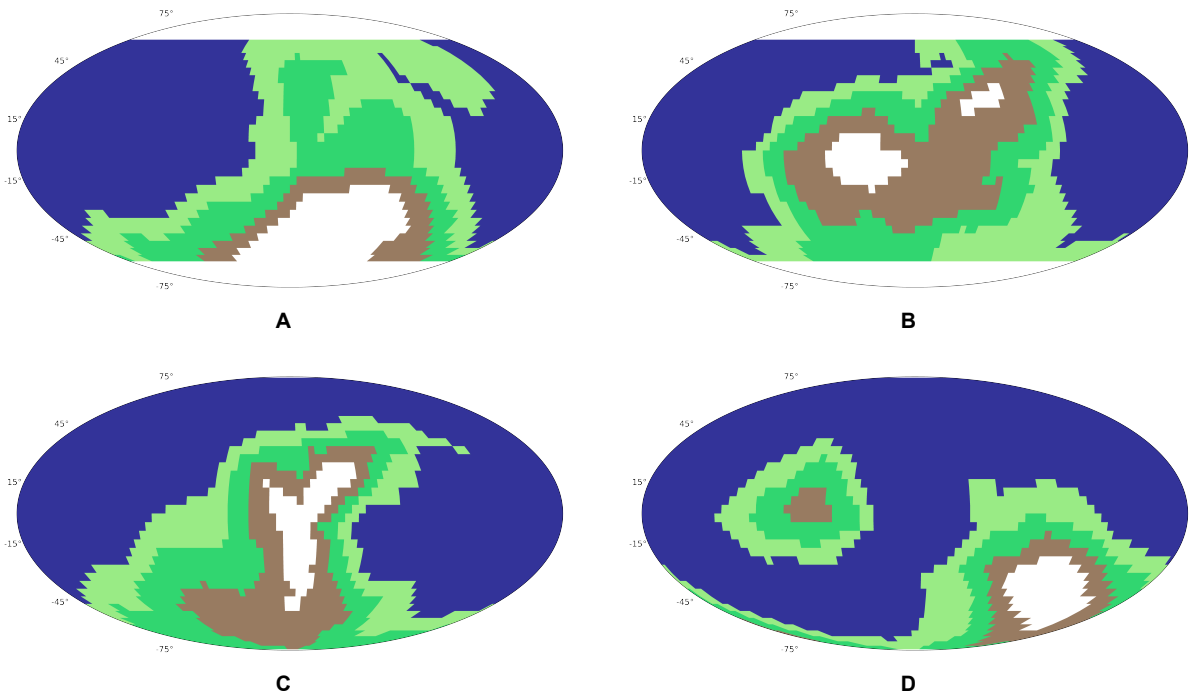


Figure 3.3: Pseudo-color map of planets A,B,C, and D. For this map, blue = ocean, light green = grasslands, dark green = forest, brown = regolith, and white = snow.

3.3. Albedo assignment and reflectance distribution

The penultimate step in our map generation is to assign a quantity of reflectance to each surface type S . This will allow us to later discern the different surface types by sampling the observed reflected light curve. When simulating the phase function in Chapter 2, we have only considered the case where our planet is homogeneous and reflects all light exactly the same. However, our generated planets have irregular surfaces, and not all surface types reflect light equally. For example, white snow reflects more incident light than dark blue water of oceans. One quantity of reflectance is the ratio of the stellar intensity of the light reflected diffusely off a particular surface to the total intensity of the light from the planet that reaches the planet's surface. This ratio is also known as the average stellar reflectance or (bond) albedo (McCluney, 2014), and is scaled from 0, for a black body that absorbs all incident light, to 1 for a body that reflects all light. Since more reflective material (e.g. snow) has a higher albedo than

absorbing material (e.g. water), the intensity of the light that different terrain types reflect will also be different, and so a planet with an irregular surface will also have a varying light curve. This will form the basis of our retrieval method. The surface reflectance for different terrain types is not only governed by the properties of the surface itself, but also by the wavelength of the spectral distribution of the incident radiation², and so for each wavelength a spectral albedo can be analyzed. However this exceeds the scope of our thesis, and so when discussing albedo we consider the average reflectance of the entire spectrum of incident radiation. Table 3.2 shows the albedo for each of the generated terrain types S .

S	a_0
Ocean	0.06
Forest	0.15
Grasslands	0.25
Regolith	0.5
Snow	0.8

Table 3.2: List of the generated surface types, along with their average reflectance or typical albedo a_0 (NSIDC (2022) and Markwart and Castaner (2003)).

The typical bond albedo allows us to distinguish the different surface types based on their ability to reflect light diffusely. However we also want to incorporate other reflection models, and different surface types reflect light in different ways. For instance, liquid oceans reflect light more according to Fresnel reflection, while regolith reflects light more according to the Lommel-Seeliger model. This means that the simple albedo a_0 alone will not suffice, and we also need to incorporate model-dependent reflectance. To solve this issue, we introduce a distribution to each surface type to indicate how much of the reflected light is scattered in accordance with each of the three reflection models. To calculate how much reflected light scatters off a certain surface facet according to a certain reflection model, we simply multiply the bond albedo a_0 of the type of the facet with the appropriate reflectance fraction P_{RM} . Therefore, the quantity of model-dependent reflectance a_{RM} is given by:

$$a_{RM}(p_j) = a_0(p_j)P_{RM}. \quad (3.1)$$

The reflectance distribution, along with the different values of the bond albedo, are shown in Table 3.3. Note that the bond albedo for oceans of 0.06 denoted the incorrect diffusive reflectance of water, and so for specular reflectance we slightly increased it to 0.1.

S	a_0	P_{LaL}	P_{Fr}	P_{LS}
Ocean	0.1	0	1	0
Forest	0.15	1	0	0
Grasslands	0.25	0.5	0	0.5
Regolith	0.5	0.3	0	0.7
Snow	0.8	0.9	0	0.1

Table 3.3: List of the generated surface types, along with their standard albedo a_0 and the fraction of reflected light that is scattered in accordance with Lambertian diffusion P_{La} , the Fresnel equations P_{Fr} , or Lommel-Seeliger diffusion model P_{LS} .

Finally we are able to generate three different surface maps per planet, with each map showing the fraction of light each unit surface reflects according to each reflection model. These surface maps will allow us to generate a time-dependent light curve of the planet, and the goal of our research will be to see if we can retrieve all three maps from the observed light curve. The three reflectance maps for each of the 4 planets are shown in Figure 3.4.

¹The range of orbits around a star within which liquid water can exist on its surface is called the habitable or Goldilocks zone (Hegner, 2020).

²For example, The spectral albedo of grass for wavelengths of 500 nm is greater than that for 700 nm, and so it reflects green light better than red light. Therefore we perceive grasslands as green.





Figure 3.4: Reflectance maps of each planet for each of the three reflection models. Left: Lambertian reflectance. Center: Fresnel reflectance. Right: Lommel-Seeliger reflectance.



Reflected light simulation

Now that we can generate planet surfaces with their different reflectance maps, it's time to simulate what the observed light of such a planet would look like. As discussed in Chapter 1, the light reflected off an exoplanet is very dim compared to the emitted light of its parent star, and because the distance between exoplanet and observer is so great, it is impossible to resolve features of the surface by direct imaging. Instead, the only quantity we would be able to measure is the total intensity of the reflected light that reaches the observer. This intensity, henceforth denoted as $I(t)$, depends on the observable domain \mathcal{D} as defined in Chapter 2, which itself depends on the exoplanet's location, given by $\zeta = \omega t$, and its orientation, given by Ωt . Furthermore, since each studied reflection model deals with its own reflectance map, we prefer to calculate the observed intensity for these models separately. The total intensity for the composite model is then simply the sum of the three calculated intensities. The energy conservation property of Section 2.2 is then satisfied as a consequence of the distribution of the reflectances in Section 3.3.

We will first formulate a complete analytical expression for the time dependent intensity for irregular planets by combining our results about the different phase functions obtained in Section 2.3 and our analysis of the surface reflectances obtained in Chapter 3. Next, we discretize this to form a numerical expression to be able to work with a finite number of surface facets during a finite number of time steps. This step requires us to describe the planet's rotation in terms of our fixed coordinate system. Thereafter we apply this expression to our own generated surface maps. Finally, we investigate how artificial noise can disturb the light curve.

4.1. Analytical expression for the reflected light curve

In Section 2.3 we showed that the light curve for an exoplanet with unit albedo and radius 1 in a circular orbit with radius R for some BRDF f and for some time t is simply the (normalised) BRDF multiplied by the prefactor $\frac{1}{R^2}$ and integrated over the visible domain:

$$I_{\text{RM}}(t) = \frac{1}{R^2} \iint_{\mathcal{D}(t)} f_{\text{RM}}(\theta_i, \phi_i; \theta_r, \phi_r) \cos \theta_i \cos \theta_r \sin \Theta d\Theta d\Phi. \quad (4.1)$$

In order to derive the analytical expression for the reflected light curves of more complex planets with radius ρ and possessing three surface maps M_{RM} , we need to multiply each unit surface facet $d^2\hat{\mathbf{s}}$ by the reflectance a_{RM} at that particular surface. Note that $d^2\hat{\mathbf{s}}$ now also gives rise to a prefactor ρ^2 , as the planet no longer has unit radius:

$$\begin{aligned} I_{\text{RM}}(t) &= \frac{1}{R^2} \iint_{\mathcal{D}(t)} f_{\text{RM}}(\theta_i, \phi_i; \theta_r, \phi_r) (\hat{\mathbf{s}} \cdot \hat{\mathbf{i}}) (\hat{\mathbf{s}} \cdot \hat{\mathbf{o}}) a_{\text{RM}}(\hat{\mathbf{s}}) d^2\hat{\mathbf{s}} \\ &= \frac{\rho^2}{R^2} \iint_{\mathcal{D}(t)} f_{\text{RM}}(\theta_i, \phi_i; \theta_r, \phi_r) \cos \theta_i \cos \theta_r a_{\text{RM}}(\Theta, \Phi) \sin \Theta d\Theta d\Phi \end{aligned} \quad (4.2)$$

The total light intensity $I_{\text{TOT}}(t)$ is then found by adding I_{La} , I_{Fr} , and I_{LS} .

4.2. Numerical expression for the reflected light curve

For simulating light curves reflected off discrete surface maps \mathbf{M}_{RM} generated in Chapter 3, we need to adjust the above integral in order to carry out numerical integration over a finite amount of surface facets $d^2\hat{\mathbf{s}}$ for a finite number of time steps t_i . If we let $N_\Phi \times N_\Theta$ be the number of surface facets, i.e. the longitudinal resolution times the latitudinal resolution, and we let N_t denote the number of performed observations, then we can represent $I_{\text{RM}}(t)$ as a vector $\mathbf{I}_{\text{RM}} \in \mathbb{R}^{N_t}$ given by:

$$\mathbf{I}_{\text{RM}}[t] = \mathbf{T}_{\text{RM}}\mathbf{M}_{\text{RM}}, \quad (4.3)$$

where $\mathbf{M} \in \mathbb{R}^{N_\Phi \times N_\Theta}$ is the vector containing the reflectance of all the facets on the planet:

$$\mathbf{M} = [a(1, 1), \dots, a(1, N_\Theta), a(2, 1), \dots, a(N_\Phi, N_\Theta)]^T, \quad (4.4)$$

and $\mathbf{T} \in \mathbb{R}^{N_t} \times (\mathbb{R}^{N_\Phi \times N_\Theta})$ a transformation matrix. The next step is to find the entries of \mathbf{T} .

4.2.1. Daily rotation of planet

Up to this point, we have assumed that the observable domain \mathcal{D} of the exoplanet only changes due to the planet's revolution around its parent star, determined by its phase angle $\zeta = \omega t$. However, as discussed earlier, we also expect the planet to rotate about some axis $\hat{\mathbf{n}}$, defined by equinox angle α and obliquity β , with some angular frequency Ω . This means that for planets with irregular surface maps the observable domain \mathcal{D} also depends on the planet's time-dependent rotation, and therefore we need to implement the axial tilt in our transformation matrix \mathbf{T} .

For each time step t_i , the planet rotates about its rotational axis, and all facets rotate with it. This means that, in order for the same spherical coordinates Φ and Θ to always describe the same facet, we need to rotate the vector $\hat{\mathbf{s}}$ pointing to that facet on each time step so that it always points to the same facet with respect to the fixed orbital coordinate system. We accomplish this by first rotating the planet Ωt_i radians about the z-axis of our fixed planetary system (i.e. that doesn't rotate along), then rotating it β radians about its fixed y-axis, before rotating it α radians about its fixed z-axis again. Using this sequence of extrinsic rotations, the formula for the new facet vector $\hat{\mathbf{s}}'$ is given by:

$$\begin{aligned} \hat{\mathbf{s}}' &= R_z(\alpha)R_y(\beta)R_z(\Omega t_i)\hat{\mathbf{s}} \\ &= \begin{bmatrix} \cos \alpha & -\sin \alpha & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos \beta & 0 & \sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & \cos \beta \end{bmatrix} \begin{bmatrix} \cos \Omega t_i & -\sin \Omega t_i & 0 \\ \sin \Omega t_i & \cos \Omega t_i & 0 \\ 0 & 0 & 1 \end{bmatrix} \hat{\mathbf{s}} \end{aligned} \quad (4.5)$$

Note that the same resulting tilt can be also achieved by first rotating the planet α radians about the z-axis of the mobile XYZ-frame attached to the planet, then rotating β radians about the rotated y-axis before rotating Ωt_i radians about the rotated z-axis again, the latter now being equivalent to the rotational axis $\hat{\mathbf{n}}$. This intrinsic rotation sequence is done with so-called Euler angles (See Taylor (2005)), and is illustrated in Figure 4.1.

4.2.2. Structure of the transformation matrix

After having rotated the surface vectors $\hat{\mathbf{s}}$ to $\hat{\mathbf{s}}'$, we can construct the entries of the different transformation matrices \mathbf{T}_{RM} . If we take a single facet of the planet, defined by $\hat{\mathbf{s}}'_j$, at certain time t_i , and we assume it lies in \mathcal{D} , then the reflected intensity for that point for a certain reflection model RM is given by:

$$\mathbf{T}_{\text{RM},i,j} a_{\text{RM}} = \frac{\rho^2}{R^2} (\hat{\mathbf{s}}'_j \cdot \hat{\mathbf{i}}) (\hat{\mathbf{s}}'_j \cdot \hat{\mathbf{o}}) f_{\text{RM}} \sin \Theta \Delta \Phi \Delta \Theta a_{\text{RM}} \quad (4.6)$$

The last step is to check if our discrete facet actually lies in \mathcal{D} . Using equation 2.13, we define the following indicator function $\mathbb{1}_{\mathcal{D}}$ to determine if a facet lies in the observable domain:

$$\mathbb{1}_{\mathcal{D}} = \begin{cases} 1 & \text{if } \hat{\mathbf{s}}' \cdot \hat{\mathbf{i}} > 0 \text{ and } \hat{\mathbf{s}}' \cdot \hat{\mathbf{o}} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (4.7)$$

Combining 4.6 and 4.7, we can compute the entry of $\mathbf{T}_{\text{RM},i,j}$ for a single point p_j at a certain time t_i :




Figure 4.1: Schematic illustrating the ZY'Z'' intrinsic rotation sequence implemented to rotate the facet vectors $\hat{\mathbf{s}}$ with respect to the original orbital coordinate system, here depicted in black. The resulting coordinate system for $\hat{\mathbf{s}}$ is depicted in violet. Figure source code: Dorian (2022).

$$\begin{aligned}
 \mathbf{T}_{\text{RM},i,j} &= \frac{\rho^2}{R^2} (\hat{\mathbf{s}}'_j \cdot \hat{\mathbf{i}}) (\hat{\mathbf{s}}'_j \cdot \hat{\mathbf{o}}) \mathbb{1}_{\text{D}} f_{\text{RM}} \sin \Theta \Delta \Phi \Delta \Theta \\
 &= \frac{\rho^2}{R^2} \left[R_z(\alpha) R_y(\beta) R_z(\Omega t_i) \begin{bmatrix} \cos \Phi_j \sin \Theta_j \\ \sin \Phi_j \sin \Theta_j \\ \cos \Theta_j \end{bmatrix} \cdot \begin{bmatrix} \cos \omega t_i \\ \sin \omega t_i \\ 0 \end{bmatrix} \right] \\
 &\quad \left[R_z(\alpha) R_y(\beta) R_z(\Omega t_i) \begin{bmatrix} \cos \Phi_j \sin \Theta_j \\ \sin \Phi_j \sin \Theta_j \\ \cos \Theta_j \end{bmatrix} \cdot \begin{bmatrix} \sin \gamma \\ 0 \\ \cos \gamma \end{bmatrix} \right] \mathbb{1}_{\text{D}} f_{\text{RM}} \sin \Theta_j \Delta \phi \Delta \Theta, \quad (4.8) \\
 t_i &= i \Delta t, \quad \Phi_j = \frac{j}{N_\Phi} \Delta \Phi, \quad \Theta_j = (j \bmod N_\Theta) \Delta \Theta, \\
 \Delta t &= \frac{1 \text{ year}}{N_t}, \quad \Delta \Phi = \frac{2\pi}{N_\Phi}, \quad \Delta \Theta = \frac{\pi}{N_\Theta}
 \end{aligned}$$

As every column of \mathbf{T}_{RM} corresponds to a different point p_j on the exoplanet's surface, and every row corresponds to a different time t_i , it follows that the columns and rows of \mathbf{T}_{RM} are all linearly independent, and therefore the rank of \mathbf{T}_{RM} is simply the lesser of the number of rows and columns:

$$\text{rank}(\mathbf{T}_{\text{RM}}) = \min \{N_t, N_\Phi \times N_\Theta\}. \quad (4.9)$$

By multiplying \mathbf{T}_{RM} with the surface map \mathbf{M}_{RM} , we obtain the intensity for one reflection model, as stated in 4.3. To calculate the total intensity \mathbf{I}_{TOT} of the composite light curve, we simply add \mathbf{I}_{La} , \mathbf{I}_{Fr} and \mathbf{I}_{LS} together. Furthermore, we can also rewrite this as a bigger matrix multiplication by concatenating the transformation matrices and the surface maps:

$$\mathbf{I}_{\text{TOT}} = \left[\mathbf{T}_{\text{La}} \quad \vdots \quad \mathbf{T}_{\text{Fr}} \quad \vdots \quad \mathbf{T}_{\text{LS}} \right] \cdot \begin{bmatrix} \mathbf{M}_{\text{La}} \\ \dots \\ \mathbf{M}_{\text{Fr}} \\ \dots \\ \mathbf{M}_{\text{LS}} \end{bmatrix}. \quad (4.10)$$

In that case, the rank of the concatenated transformation matrix \mathbf{T}_{TOT} becomes:

$$\text{rank}(\mathbf{T}_{TOT}) = \min \{N_t, 3(N_\Phi \times N_\Theta)\}. \quad (4.11)$$

4.3. Application to generated surface maps

Now that we have a numerical expression for the observed intensity, we can apply it to our generated maps. To do this, we have to make some assumptions regarding the properties of the exoplanet and its orbit around the parent star. Since we'd like to study planets similar to Earth, we set the planetary radius ρ and the radius R of its circular orbit to that of the Earth, namely $6.378 \cdot 10^3$ m and $1.496 \cdot 10^6$ m respectively. One might object to this by stating that planetary orbits are actually elliptical. However, Earth's eccentricity is only about 0.06 (Berger and Loutre, 1991), which makes its orbit essentially circular, and so a low-eccentricity orbit for our exoplanet is justified. For the rotational and orbital period ($\frac{2\pi}{\Omega}$ and $\frac{2\pi}{\omega}$, we also emulate the Earth, i.e. 365 sidereal days per year and 24 hours per sidereal day. As discussed in Chapter 3, we let $N_\Phi = 80$ and $N_\Theta = 40$, as a reasonable trade-off between resolution and computing time. Furthermore, in the most ideal case, we effortlessly sample the intensity every hour throughout the exoplanet's year, so $\Delta t = 1$ hour, and therefore $N_t = 8760$. We compare this with the less ideal case with less frequent sampling, namely only sampling in 200 periods spread evenly across the planet-year for 17 hours per period, so $N_t = 3400$. In Chapter 5 we'll also see how our model fares retrieving the surface maps if the light curve $\mathbf{I}[t]$ is sampled less frequently.

4.3.1. Noiseless light curves

Frequent observations

We first simulate our observed light curve for edge-on observations in the ideal situation, i.e. no noise and hourly observations. Figure 4.2 shows the reflected intensity I off planet A for the different reflection models during a complete orbit around its parent star. Note that just as in Chapter 2, the observations start when the planet's phase angle ζ is 0, which means the planet is opposite its star, and so its illuminated light is entirely visible from Earth. When $\zeta \approx \pi$, the exoplanet passes in an astronomical transit between the observer and the star.

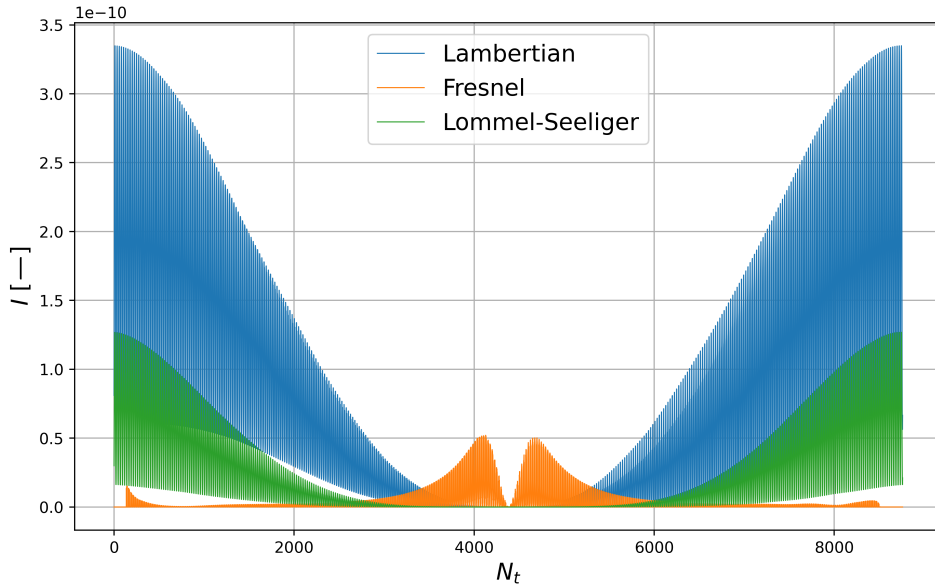


Figure 4.2: Relative intensities of the different light curves \mathbf{I}_{RM} of planet A for one full planet-year as a function of time, sampled $N_t = 8760$ times evenly from noiseless edge-on observations. Sampling starts at $\zeta = 0$.

As expected, the intensity for the Lambertian and Lommel-Seeliger reflection model is highest for ζ close to 0 and close to 2π , as that is when the observable domain \mathcal{D} is at its maximum. During the astronomical transit ($\zeta \approx \pi$), \mathcal{D} is nearly vanished, and so there is practically no diffuse reflection. On the other hand, the intensity for the Fresnel reflection model I_{Fr} increases significantly when ζ

approaches π , since in this case the incident zenith angle $\theta_i \approx \pi$ and so the Fresnel coefficients R_s and R_p are at their maximum.

Figure 4.3 shows the total reflected intensity I_{TOT} for the ideal situation, sampled every hour in the exoplanet's year. We can clearly see that the noiseless intensity can be described by 2 significant oscillations. The larger oscillation with a period of 1 year is caused by the exoplanet's orbit around the star, and the smaller oscillations are caused by the planet's daily rotation. To demonstrate this more clearly, Figure 4.4 shows an enhanced version of the composite light curve for planet D during the 7th and 8th day of its orbit around its parent star. In this plot we can clearly distinguish two peaks of different intensity that recur each exoplanet-day. When comparing this to the surface map of planet D, we can see that these peaks are a result of the two continents on the planet. The larger intensity peak is caused by the starlight reflected off the bigger continent, while the smaller peak is produced by the reflected light off the smaller island.




Figure 4.3: Relative intensity of the composite light curve I_{TOT} of planet A for one full planet-year as a function of time, sampled $N_t = 8760$ times evenly from noiseless edge-on observations.

Less frequent observations

Figure 4.5 shows the total reflected intensity I_{TOT} when sampled only 3400 times. Compared to 4.3 we clearly see the change in sampling rate, and so we obtain fewer datapoints and therefore less information about the exoplanet's surface.

4.3.2. Light curves with shot noise

In order to make the observed intensities more realistic, we consider the faint light rays reflected off our exoplanet. These rays are composed of a finite number of discrete photons, which are emitted from the star at random intervals, and therefore also observed by us at random intervals. This random occurrence can be described by a Poisson distribution, with its parameter λ denoting the average number of photons that reach the observer, henceforth denoted as N_{ph} . This average is dependent of the type of telescope, and in the case of the recently launched James Webb Telescope it is calculated that $N_{ph} = 187$ (Stuger, 2021). This means that the possibility of measuring a certain number of photons per hour is given by:

$$P(n \text{ observed photons in one hour}) = \frac{N_{ph}^n e^{-N_{ph}}}{n!} \quad (4.12)$$

As the planet is several light-years away, only a handful of photons eventually reach our telescope every hour, and so the observed intensity will vary significantly due to this randomness. This effect is

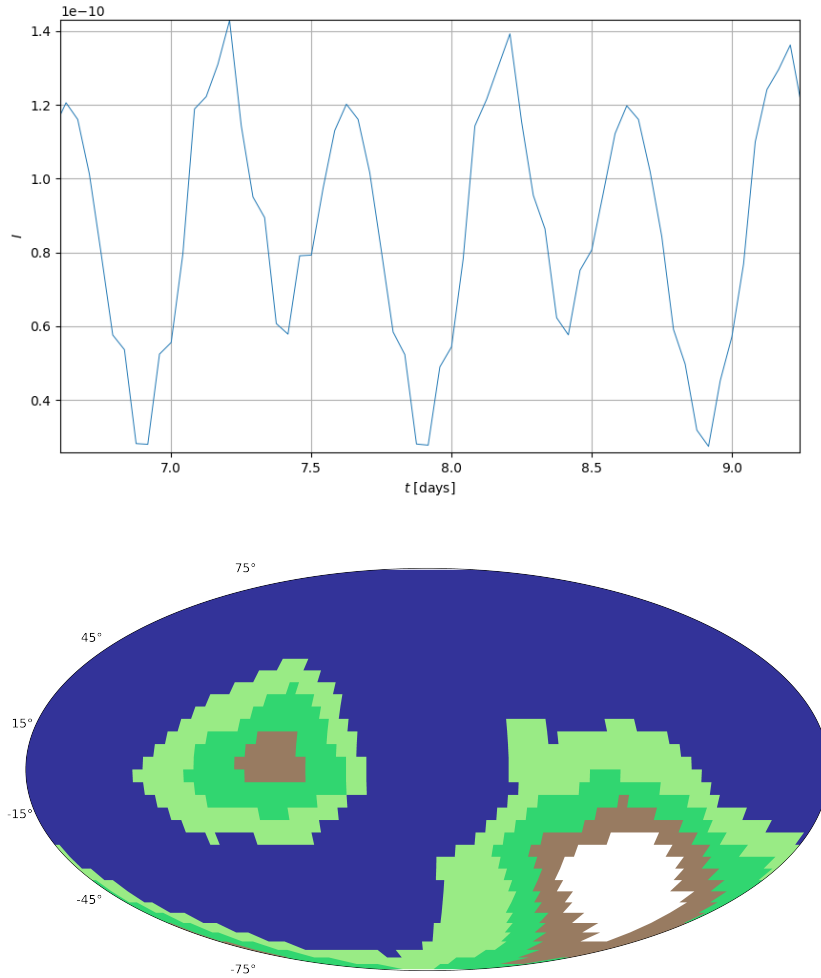


Figure 4.4: Above: Enhanced view of the composite intensity curve for planet D during the 7th and 8th day of its orbit. Below: the generated pseudo-colormap of planet D. The two periodic peaks in the light curve correspond to the two continents on the surface.

known as shot noise, and it will be implemented our simulation to study how our model performs with actual observations of the reflected light curves.

If we let $\mathbf{Ps}_{N_{ph}} \in \mathbb{R}^{N_t}$ be a vector that draws the number of observed photons for each time t from a Poisson distribution with mean N_{ph} , it follows that the total intensity for a particular reflection model disturbed by shot noise can be described by:

$$\mathbf{I}_{RM,N} = \frac{\mathbf{I}_{RM}^T \mathbf{Ps}_{N_{ph}}}{N_{ph}} \quad (4.13)$$

Figure 4.6 and Figure 4.7 shows the total reflected intensity with added shot noise for $N_t = 8760$ and $N_t = 3400$ respectively. In each case, $N_{ph} = 187$.




Figure 4.5: Relative intensity of the composite light curve I_{TOT} of planet A for one full planet-year as a function of time, sampled $N_t = 3400$ times evenly from noiseless edge-on observations.




Figure 4.6: Relative intensity of the composite light curve I_{TOT} of planet A for one full planet-year as a function of time, sampled $N_t = 8760$ times evenly from edge-on observations with artificial shot noise with $N_{ph} = 187$.

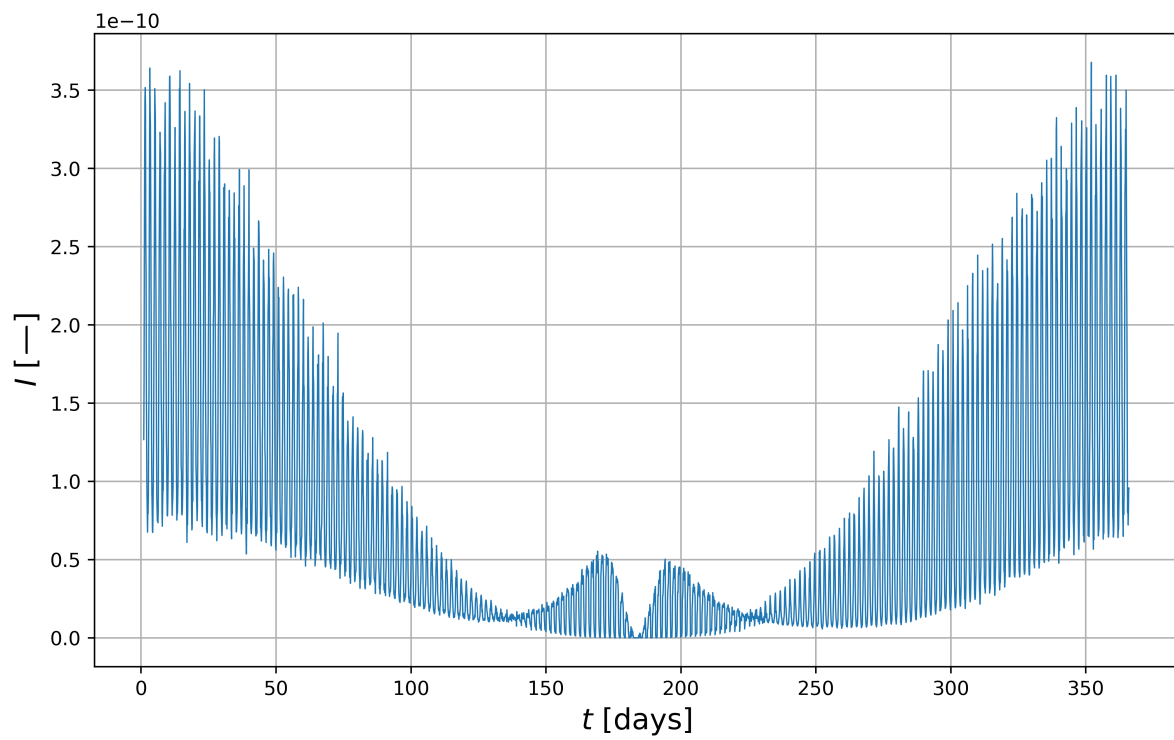


Figure 4.7: Relative intensity of the composite light curve I_{TOT} of planet A for one full planet-year as a function of time, sampled $N_t = 3400$ times evenly from edge-on observations with artificial shot noise with $N_{ph} = 187$.

5

Surface map retrieval

In this final chapter, we attempt to replicate the results of several researchers by retrieving the reflectance maps based on the observed intensity signal $\mathbf{I}[t]$. However, in contrast to other authors who have only assumed Lambertian reflection, we attempt to retrieve the 3 distinct reflectance maps \mathbf{M}_{RM} (as introduced in Chapter 3 for each reflection model from the single light curve. We will achieve this by inverting the transformation of these maps introduced in Chapter 4. As the matrix \mathbf{T}_{TOT} is not necessarily square and therefore not invertible, we'll have to resort to a more general form of the inverse matrix, namely the Moore-Penrose inverse. We will carry out this transformation inversion for different values of the obliquity β and equinox angle α , as well as for varying measurement rates of the light curve. By doing this, we assume we know a priori for each observed planet what these axial tilt angles are¹. Next, we'll study how our model fares retrieving the reflectance maps when artificial shot noise is added, both for frequent and less frequent sampling. Lastly, we'll compare our new composite model with the traditional model of purely Lambertian reflection. To do this, we'll first have both models retrieve the reflectance map of a planet that only reflects Lambertian-like, and compare their retrieved maps with the original. Next, we'll do the same, but now for a planet that reflects according to our composite model as assumed in Chapter 3.

5.1. Inversion of the transformation matrix

5.1.1. Moore-Penrose pseudo-inverse

As concluded in Chapter 4, the observed intensity \mathbf{I}_{TOT} can be simply expressed as the matrix multiplication of the (concatenated) transformation matrix \mathbf{T}_{TOT} and the concatenated reflectance maps \mathbf{M}_{RM} . In other words, $\mathbf{I}_{TOT} = \mathbf{T}_{TOT}\mathbf{M}_{TOT}$. Therefore, in order to invert the transformation, one could argue to just multiply the intensity by the inverse of \mathbf{T} to obtain \mathbf{M} . However, as the shape of \mathbf{T}_{TOT} is defined by N_t, N_Φ and N_Θ , it does not necessarily have to be square, and so a perfect inverse matrix \mathbf{T}^{-1} cannot always be defined. Fortunately, a particular generalization of the inverse for non-square matrices does always exist, known as the Moore-Penrose inverse \mathbf{T}^+ . This pseudoinverse is unique for every matrix and can be used to find the best fitting solution for the combined reflectance maps \mathbf{M}_{TOT} .

Let $\mathbf{T} : \mathbb{R}^{N_\Phi \times N_\Theta} \rightarrow \mathbb{R}^{N_t}$ be the linear transformation matrix for a particular reflection model. Its corresponding reflectance map map \mathbf{M} can be decomposed as the sum of a vector in $\text{Row}(\mathbf{T})$, named \mathbf{M}_r and a vector in $\text{Null}(\mathbf{T})$, named \mathbf{M}_n . Similarly, any observed light curve \mathbf{I} can be decomposed as the sum of a vector in $\text{Row}(\mathbf{T}^T) = \text{Col}(\mathbf{T})$, named \mathbf{I}_c , and a vector in $\text{Null}(\mathbf{T}^T)$, named \mathbf{I}_n . However, since all our intensities are computed by Equation 4.10, $\mathbf{I} \in \text{Col}(\mathbf{T})$ unless artificial noise is added, so $\mathbf{I}_c = \mathbf{I}$. In this case, \mathbf{T} maps \mathbf{M}_r to \mathbf{I} , while \mathbf{M}_n is mapped to $\mathbf{0}$. If \mathbf{T} is not full rank, i.e. if $\text{Null}(\mathbf{T}) \neq \{\mathbf{0}\}$, any multiple of $\mathbf{M}_n \in \text{Null}(\mathbf{T})$, and so we cannot find a unique \mathbf{M} such that $\mathbf{I} = \mathbf{T}\mathbf{M}$. However, as any $\mathbf{I}_c \in \text{Col}(\mathbf{T})$ is the mapping of one and only one $\mathbf{M}_r \in \text{Row}(\mathbf{T})$, the transformation from $\text{Row}(\mathbf{T})$ to $\text{Col}(\mathbf{T})$ is actually bijective and therefore invertible. This inverting operator is called the Moore-Penrose pseudoinverse \mathbf{T}^+ , and is defined as follows:

¹See Stuger (2021) for the recovery of the axial tilt in spin-orbit tomography using Euclidean distance.

$$\mathbf{T}^+\mathbf{I} = \begin{cases} \mathbf{M}_r & \text{if } \mathbf{I} \in \text{Col}(\mathbf{T}) \\ \mathbf{0} & \text{if } \mathbf{I} \in \text{Null}(\mathbf{T}^T) \end{cases} \quad (5.1)$$

Therefore, for any reflectance map \mathbf{M} , we have that

$$\mathbf{T}^+\mathbf{I} = \mathbf{T}^+(\mathbf{I}_c + \mathbf{I}_n) = \mathbf{M}_r + \mathbf{0} = \mathbf{M}_r, \quad (5.2)$$

and since $\text{Row}(\mathbf{T})$ and $\text{Null}(\mathbf{T})$ are orthogonal vector spaces, it follows that \mathbf{M}_r is the best approximation for \mathbf{M} , as it is the solution with minimal Euclidean distance to the true map. Furthermore, surface facets of the planet that are never observed² and therefore contribute no information to \mathbf{I} will be given zero reflectance when retrieving the map.

The method for retrieving the different reflectance maps from a composite light curve \mathbf{I}_{TOT} is analogous, but in that case we take the pseudo-inverse of the concatenated transformation matrix to obtain the (closest) solution for the concatenated reflectance map \mathbf{M}_{TOT} :

$$\begin{bmatrix} \mathbf{T}_{La} & \vdots & \mathbf{T}_{Fr} & \vdots & \mathbf{T}_{LS} \end{bmatrix}^+ \mathbf{I}_{TOT} = \begin{bmatrix} \mathbf{M}_{La} \\ \vdots \\ \mathbf{M}_{Fr} \\ \vdots \\ \mathbf{M}_{LS} \end{bmatrix}_r. \quad (5.3)$$

The desired reflectance maps $\mathbf{M}_{RM,r}$ can then be obtained by slicing $\mathbf{M}_{TOT,r}$ accordingly.

5.1.2. Singular value decomposition and cutoff ratio

The Moore-Penrose pseudo-inverse \mathbf{T}^+ can be numerically computed by decomposing or factorizing the $N_t \times (N_\Phi \times N_\Theta)$ matrix \mathbf{T} into three matrices:

$$\mathbf{T} = \mathbf{U}\Sigma\mathbf{V}^T, \quad (5.4)$$

where \mathbf{U} is a $N_t \times N_t$ unitary matrix, \mathbf{V} is a $(N_\Phi \times N_\Theta) \times (N_\Phi \times N_\Theta)$ unitary matrix, and Σ is a $N_t \times (N_\Phi \times N_\Theta)$ diagonal matrix with non-negative real numbers on its diagonal. These diagonal entries are uniquely determined by \mathbf{T} , and are called the singular values of \mathbf{T} , henceforth denoted as σ_i . This is also why this factorization is called the singular value decomposition or SVD. The number of non-zero σ_i is equal to the rank of \mathbf{T} , which is just $\min\{N_t, 3(N_\Phi \times N_\Theta)\}$, since the rows and columns of \mathbf{T} are linearly independent, as we have seen in Chapter 4. Figure 5.1.2 shows a schematic of the singular value decomposition of a matrix \mathbf{T} , where r denotes the rank of \mathbf{T} .

$$\mathbf{T} = \mathbf{U}\Sigma\mathbf{V}^T$$

$$= \begin{pmatrix} \underbrace{\begin{pmatrix} | & & | & & | & & | \\ u_1 & & u_r & u_{r+1} & & & u_m \\ | & & | & & | & & | \end{pmatrix}}_{\text{Col}(\mathbf{T})} & \dots & \underbrace{\begin{pmatrix} | & & | & & | \\ u_{r+1} & & u_m \\ | & & | \end{pmatrix}}_{\text{Null}(\mathbf{T}^+)} \end{pmatrix} \begin{pmatrix} \sigma_1 & & & & & & \\ & \ddots & & & & & \\ & & \sigma_r & & & & \\ & & & 0 & & & \\ & & & & \ddots & & \\ & & & & & 0 & \end{pmatrix} \begin{pmatrix} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{pmatrix} \begin{matrix} v_1^T \\ v_r^T \\ v_{r+1}^T \\ v_n^T \end{matrix} \left. \begin{matrix} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{matrix} \right\} \begin{matrix} \text{Row}(\mathbf{T}) \\ \text{Null}(\mathbf{T}) \end{matrix}$$

Using the decomposition, the pseudo-inverse of the real matrix \mathbf{T} can be easily computed by the following equation (Golub and Van Loan, 2013):

²For instance, if $\alpha = 0$, $\beta = \frac{\pi}{2}$, we would only ever observe the northern hemisphere of the planet, so the entire southern hemisphere would be mapped to 0 reflectance.

$$\mathbf{T}^+ = \mathbf{V}\Sigma^+\mathbf{U}^T. \quad (5.5)$$

For the diagonal matrix Σ , its pseudo-inverse Σ^+ is computed by taking the reciprocal of every non-zero singular value and then transposing the matrix.

When numerically computing the SVD and the subsequent pseudo-inverse of a transformation matrix \mathbf{T} , it's possible due to rounding errors that some singular values σ_{ii} that are actually equal to 0 are instead very small positive numbers. These small non-zero values will then erroneously produce extremely large diagonal entries when calculating Σ^+ , which causes additional errors in the calculation of \mathbf{T}^+ . As such, \mathbf{I}_n no longer gets mapped to $\mathbf{0}$, but instead to some other vector in $\mathbb{R}^{N_\Phi \times N_\Theta}$. This can lead to enormous errors in the retrieval of the reflectance maps \mathbf{M}_{RM} , especially if the observed light curve has added noise or isn't measured frequently. To counteract this, we propose to truncate Σ before inverting, i.e. to reassign all singular values of Σ smaller than some fraction of $\max\{\sigma_i\}$ to 0, such that the number of non-zero singular values lies close to the actual rank of \mathbf{T} . This threshold is called the Singular Value Cutoff Ratio or SVCR and lies between 0 and 1. One has to be careful when choosing the SVCR. Setting it too low will not get rid of all the 'bad' singular values, and so noise or faulty rounding errors will still disturb reflectance map retrieval. Setting it too high, and some of the 'good' singular values determined by \mathbf{T} will also be set to 0, which will result in a loss of information and a less resolved reflectance map. As we will see, the correct choice of SVCR is very dependent of situational circumstances, such as the spread of Gaussian noise, the amount of measurements made, and the axial tilt. Therefore, one has to resort to manually adjusting the SVCR to obtain the clearest reflectance map possible.

5.2. Noiseless retrieval

Figure 5.1 and Figure 5.2 show the resulting retrieval when using our method for frequent edge-on noiseless observations of our exoplanet with equinox angles $\alpha = \frac{\pi}{6}$ and $\alpha = \frac{\pi}{2}$ respectively. In this case, the number of sampling points $N_t = 8760$ is slightly smaller than three times number of surface facets $N_\Phi \times N_\Theta = 3200$. The retrieved maps for Lambertian, Fresnel, and Lommel-Seeliger reflectance are shown for different values of the obliquity β , and in each case the optimal SVCR that produces the clearest image is used. Figure 5.3 and Figure 5.4 show the same figures for less frequent observations, when the number of sampling points $N_t = 3400$ is much smaller than $3(N_\Phi \times N_\Theta)$.

As can be seen in Figure 5.1 and 5.2, our model is able to retrieve the reflectance maps very accurately from the noiseless light curves when we sample frequently, especially the maps indicating Lambertian and Lommel-Seeliger reflectance. We can easily distinguish the different terrain types present in the generated map. As predicted in section 5.1, the parts of the planet's surface that are never in the observable domain are attributed zero reflectance by the Moore-Penrose inverse transformation. This happens for instance in Figure 5.3 when $\alpha = \frac{\pi}{6}$ and $\beta = \frac{\pi}{4}$, such that the unobservable southern polar ice caps are not retrieved in our reflectance maps. Apart from this hidden area, the retrieved maps are nearly identical to the original, save for some scaling errors. Furthermore, if $\alpha = \frac{\pi}{2}$, then the planet's rotational axis is perpendicular to $\hat{\mathbf{o}}$, and so every point on the planet's surface is visible to us at least once, so we are able to retrieve the entire surface.

When studying the retrieved Fresnel reflectance maps more closely, we can see that there is a lateral band surrounding the planet's equator where the surface is retrieved quite accurately, but outside of this band the retrieval is more noisy. Remarkably enough, this border lies at about β radians above and below the equator, which is especially apparent if $\alpha = \frac{\pi}{2}$. This effect can be attributed to the glint-like behaviour of the Fresnel reflection model, in which the reflected light is only observable if $\theta_r = \theta_i$ and $\phi_r = \phi_i + \pi$ as we discussed in Chapter 2. For edge-on observations, this is only possible on the latitudes of the planet at which the star can be seen directly overhead, i.e. the subsolar points. The analogous region on Earth are the tropics, and the latitude of the enclosing tropical circles is equal to the obliquity β . The fact that regions outside of this tropical band can still be retrieved, albeit with a little less certainty, can be attributed to the fact that the BRDF for Fresnel reflection is approximated by a Gaussian as noted in Chapter 2.

In Figure 5.2 we can also see that for larger angles of obliquity β , a smaller strip around the equator suddenly becomes much more noisy, and we cannot distinguish ocean from continent. This can be explained by the fact that for $\beta \approx \frac{\pi}{2}$ the light reflected from this strip according to Fresnel reflection is

only observed if the phase angle ζ of the planet is very small, i.e. when the planet is directly behind the star in eclipse. Note that it is brevity of this moment that causes this blind strip, and not the fact that the star obscures the planet, as we haven't incorporated that in our model. As expected, the noiseless retrieval for frequently sampled light curves doesn't require a high SVCR to obtain a clear picture.

When comparing these results to those when the intensity is sampled less frequently, the first thing we notice is that the retrieved maps in Figure 5.3 and 5.4 are less accurate than their frequent-sampling counterparts. This difference makes sense intuitively, as more measurements and data points provides more information about the planet's surface, and will subsequently lead to less uncertainty and a clearer picture of all 3 reflectance maps. In the context of the transformation matrix \mathbf{T}_{TOT} , sampling less frequently than three times the number of surface facets causes the rank of \mathbf{T}_{TOT} to be determined by the number of sampling points N_t . It follows that sampling less frequently reduces the rank of \mathbf{T}_{TOT} , which in turn increases its nullity, so that the transformation is not injective, i.e. multiple reflectance maps \mathbf{M}_{TOT} would produce the same composite light curve I_{TOT} . Therefore the retrieved maps are not necessarily the originals, and so perfect retrieval is not possible. This uncertainty in our map retrieval decreases as we increase the number of sampling points N_t , up to three times the number of surface facets $3(N_\Phi \times N_\Theta)$, which then becomes the maximum rank of \mathbf{T}_{TOT} . As $N_t = 8760$ lies much closer to this number, its transformation matrix has a higher rank, and so the Moore-Penrose inversion operates more like the actual matrix inversion for full rank matrices, and therefore we retrieve less ambiguous maps. The other peculiarities we observed in the frequent-sampling results, such as the tropical β -band and the blind strip for $\beta = \frac{\pi}{2}$ are also visible when we sample less frequently, and their explanations remain the same. Furthermore we can see that larger angles of β provide the best retrieval accuracy for Lambertian and Lommel-Seeliger reflection. Lastly, even with significantly less samples, we still don't need a large SVCR to obtain a coherent set of maps, with a cutoff ratio of 10^{-5} being the minimum when $\alpha = \frac{\pi}{6}$ and $\beta = \frac{\pi}{6}$.

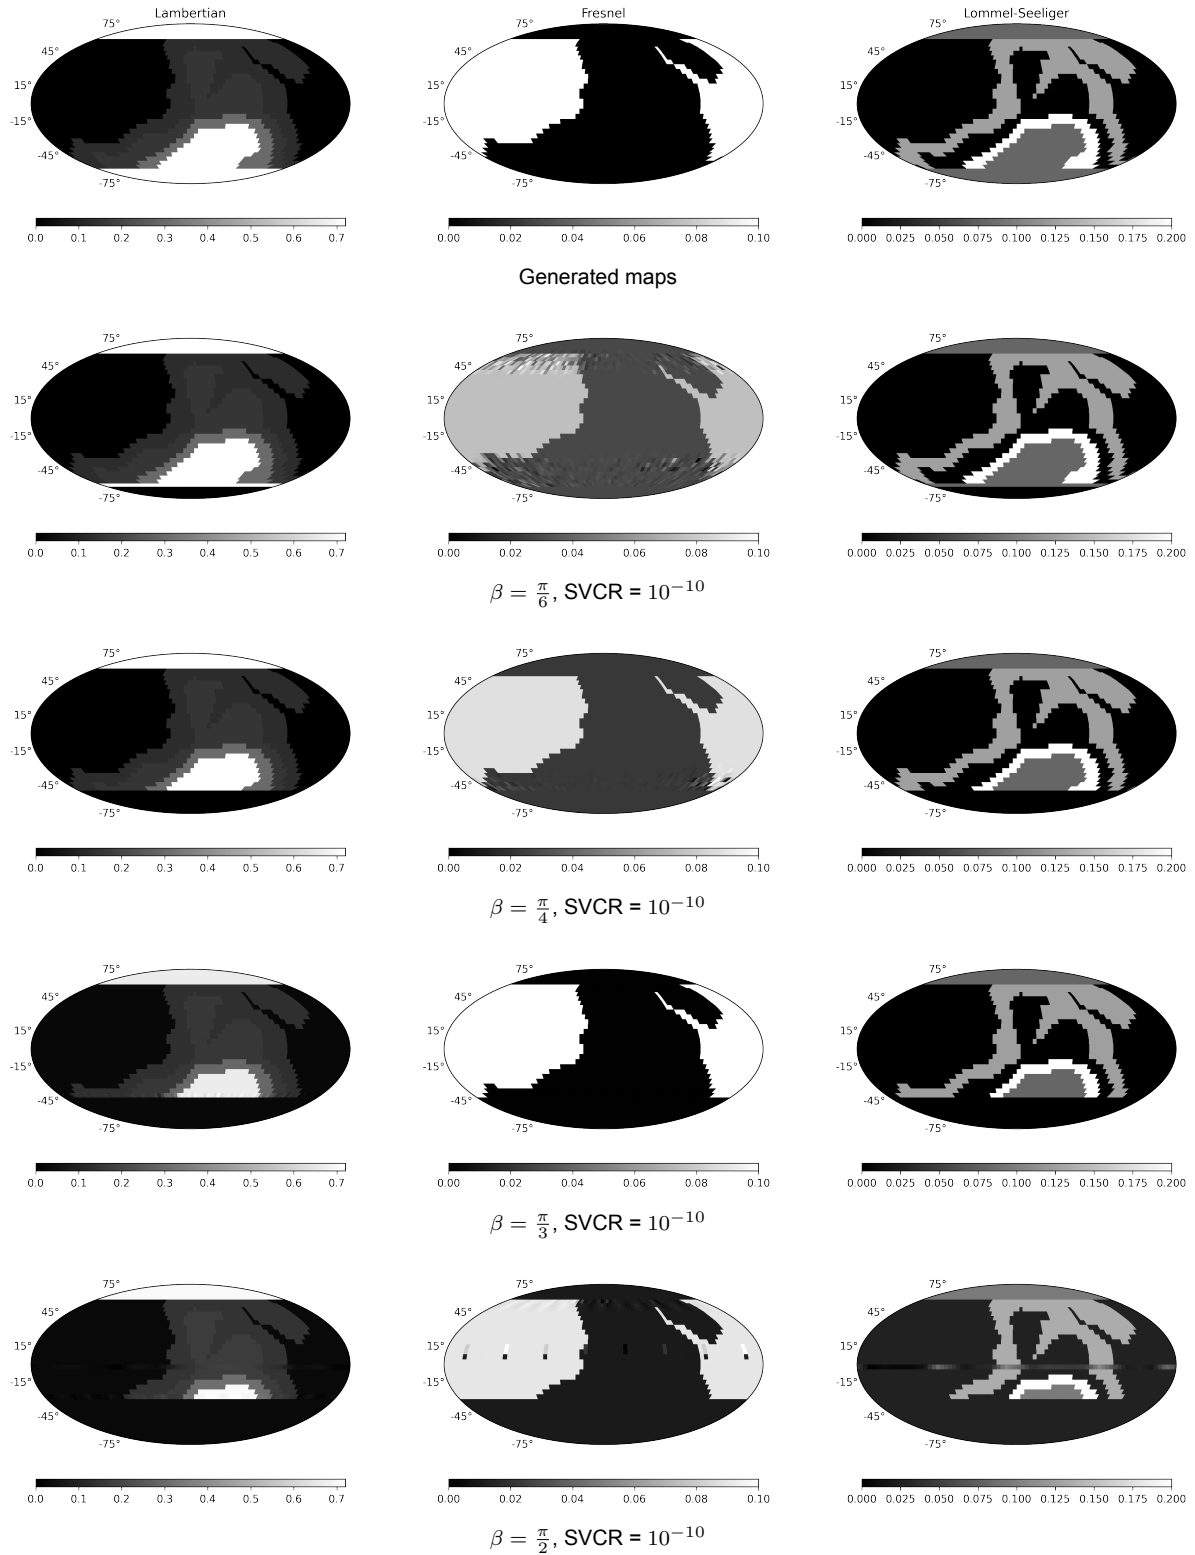


Figure 5.1: The original reflectance maps of Planet A, together with the reflectance maps retrieved from noiseless composite light curves for different obliquities β where $\alpha = \frac{\pi}{6}$ and $N_t = 8760$. The unretrievable part of the surface on the southern hemisphere grows as the obliquity β increases. As there is no noise in our signal, the SVCR can be taken very low while still retrieving near perfect images.

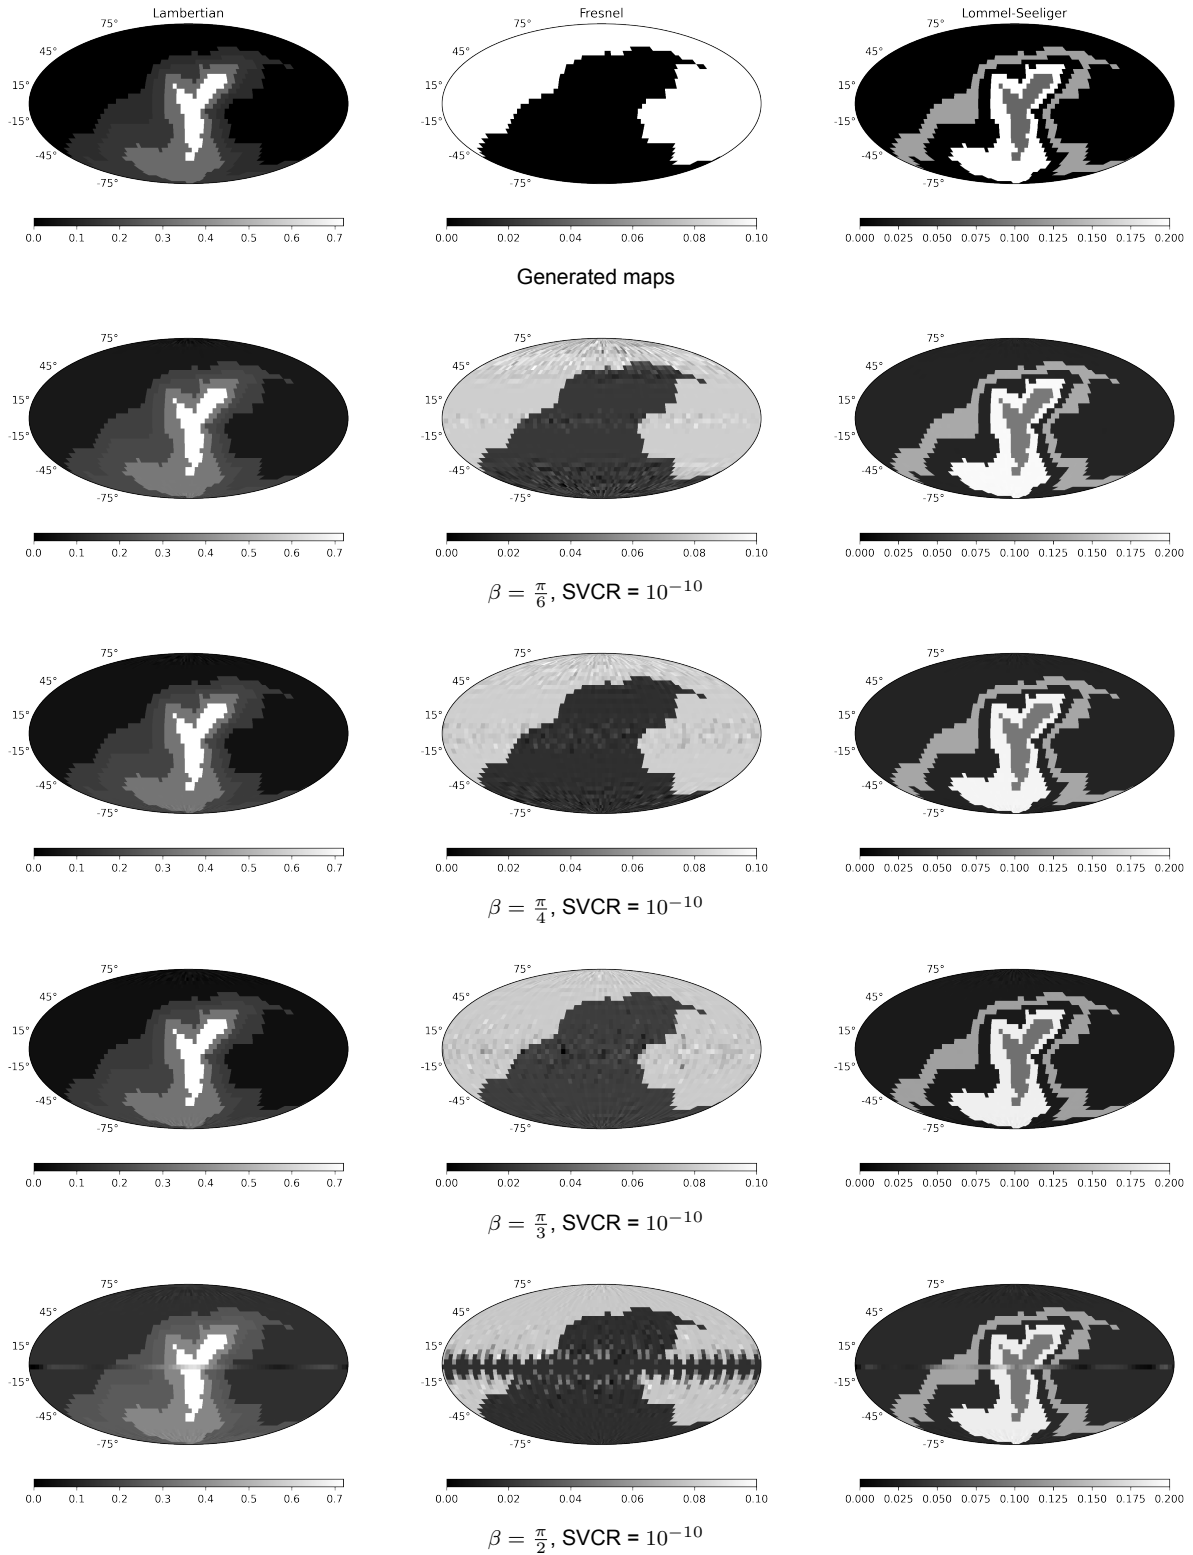


Figure 5.2: The original reflectance maps of Planet C, together with the reflectance maps retrieved from noiseless composite light curves for different obliquities β where $\alpha = \frac{\pi}{2}$ and $N_t = 8760$. As $\alpha = \frac{\pi}{2}$, every surface is visible to us at least once, and hence we are able to retrieve the entire Lambertian and Lommel-Seeliger maps. The band of width 2β around the equator where Fresnel retrieval is possible is also visible. The blind strip near the equator on the Fresnel map might come from the small time window that this region is illuminated specularly.

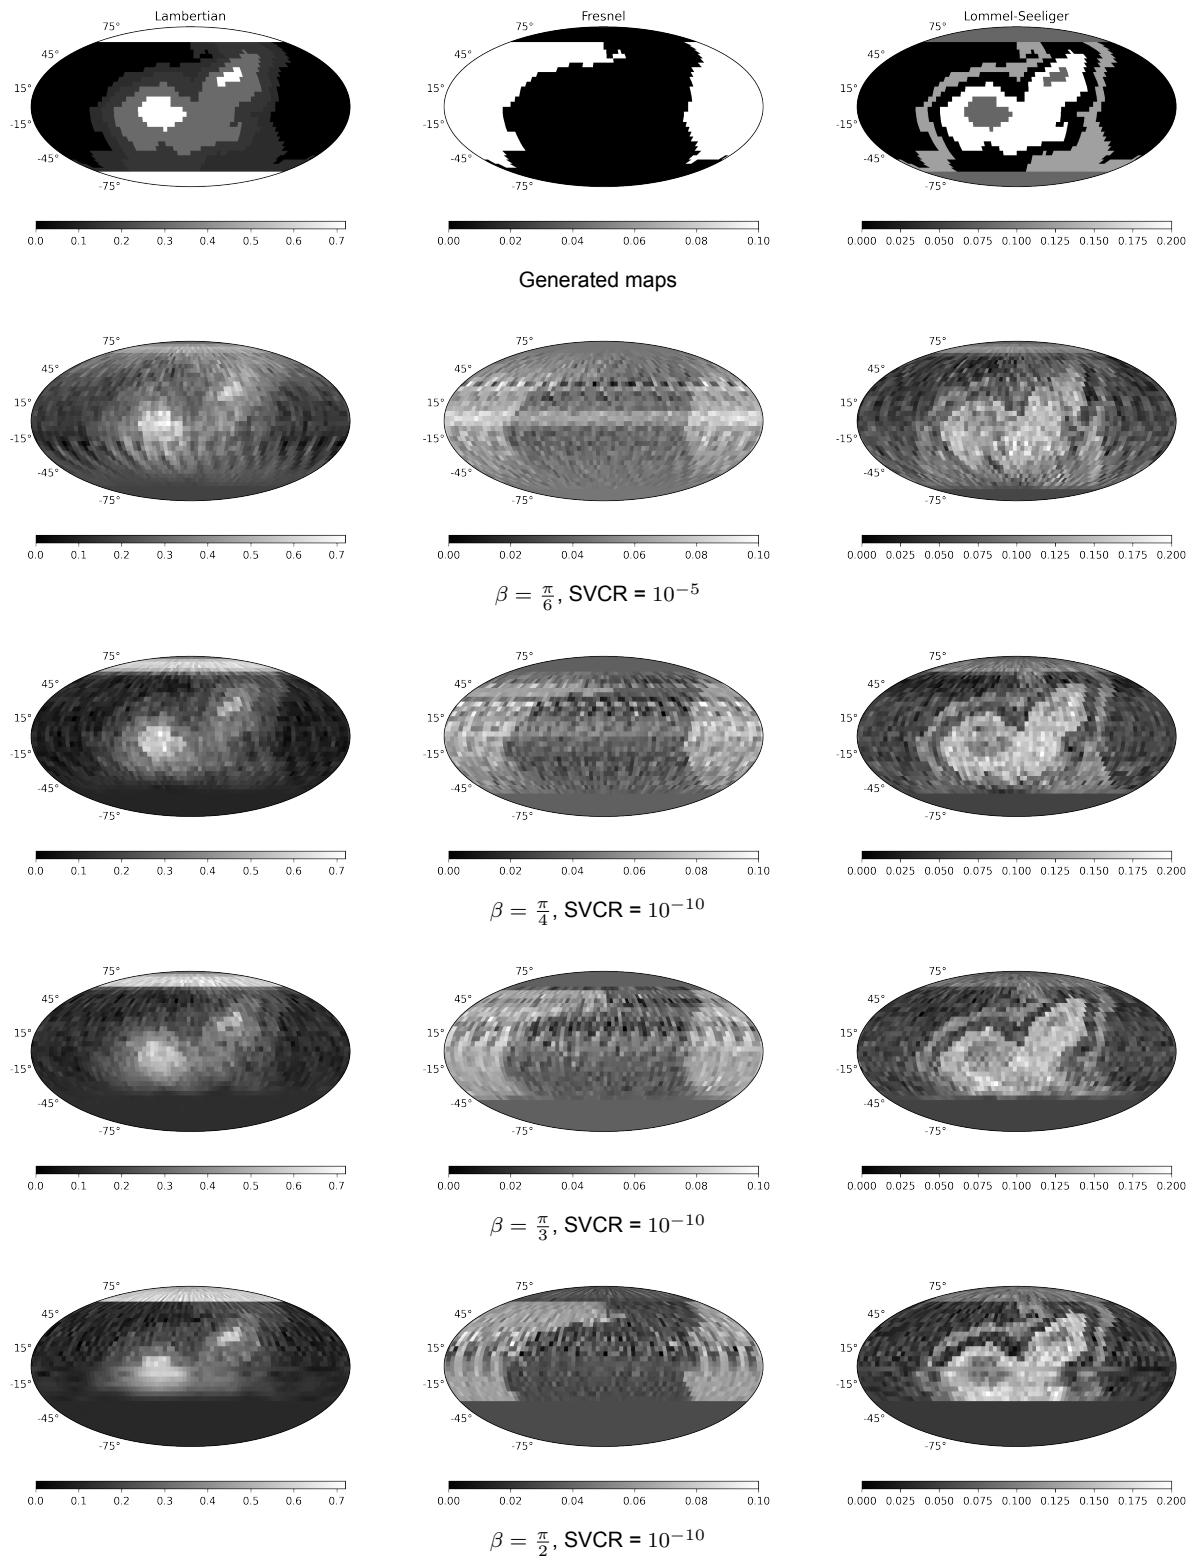


Figure 5.3: The original reflectance maps of Planet B, together with the reflectance maps retrieved from noiseless composite light curves for different obliquities β where $\alpha = \frac{\pi}{6}$ and $N_t = 3400$. As the amount of samples is lower, the quality of the retrieved maps is quite worse, however we are still able to discern the different terrain types, especially from the Lambertian and Lommel-Seeliger maps if β is large, and so the SVCR can remain quite low for accurate retrieval.

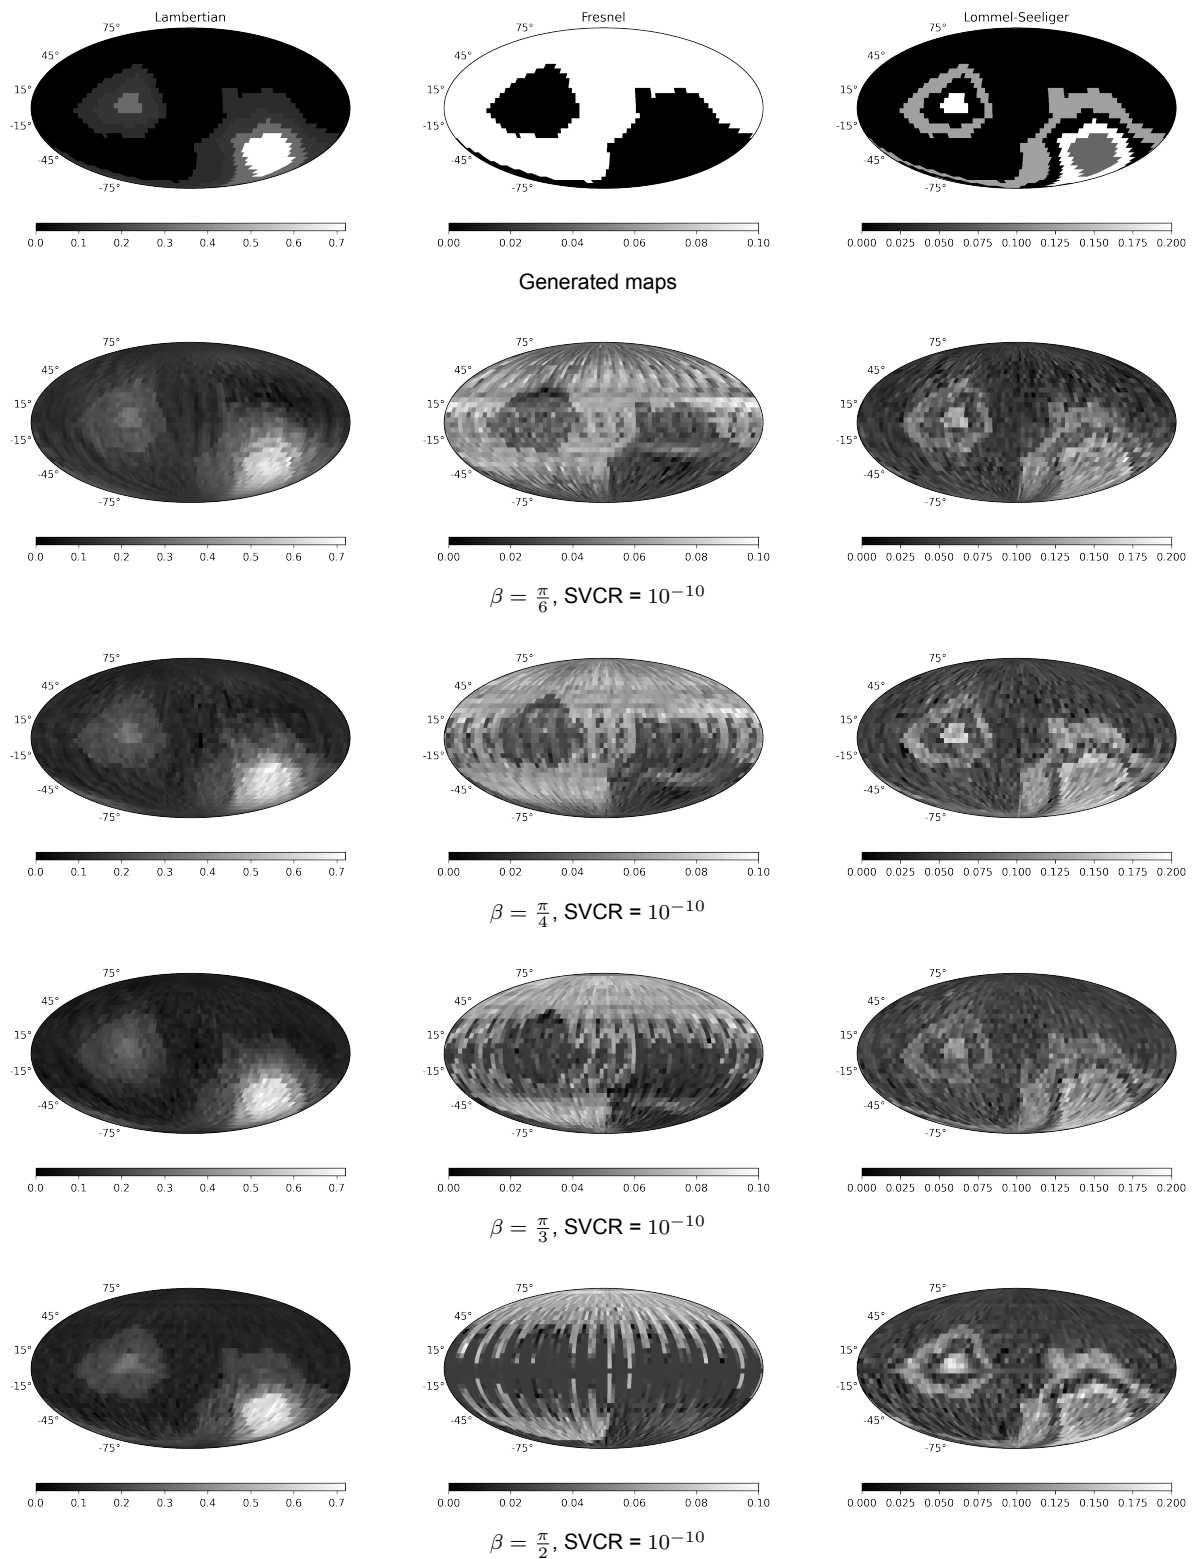


Figure 5.4: The original reflectance maps of Planet D, together with the reflectance maps retrieved from noiseless composite light curves for different obliquities β where $\alpha = \frac{\pi}{2}$ and $N_t = 3400$. The 2β band on the Fresnel map occurs again, as well as the blind strip for large β .

5.3. Retrieval with noise

Figure 5.5 shows the resulting reflectance maps retrieval when using our method for edge-on observations of our exoplanets with added photon noise as explained in Chapter 4 for $N_t = 8760$, while Figure 5.6 shows the same with less information ($N_t = 3400$). In each case, $N_{ph} = 187$, $\alpha = \frac{\pi}{2}$, $\beta = \frac{\pi}{6}$, and the retrieved reflectance maps are shown for varying values of the SVCR.

When studying both figures, it becomes apparent that even a little bit of noise heavily interferes with our model. Even when sampling frequently, retrieving the reflectance maps proves to be very difficult. In the best case, we are only able to resolve the rough outlines of the continent in the Lambertian and Lommel-Seeliger reflectance maps. Retrieving the Fresnel map that would divide the oceans from continents is also not possible. It has to be said that the retrieval of the maps with similar accuracy to that of Figures 5.1 - 5.4 only became possible when we modeled Gaussian noise instead with a small variance, which is an approximation for the photon noise for high values of N_{ph} . We believe this inaccuracy is due to the relatively small rank of the transformation matrix \mathbf{T}_{TOT} in each scenario, which leads to a high uncertainty of the retrieved maps. Furthermore, for the Fresnel map specifically, it's true that its contribution to the observed light curve is already quite small³, and so when artificial noise is added this information is most susceptible to interference.

Another clear result is that truncating the diagonal matrix Σ with the Singular value cutoff ratio now becomes essential to retrieve a somewhat accurate map, as predicted in 5.1.2. For both sampling rates, the optimal SVCR for the clearest picture seems to lie around 0.05. It was found that this correlation also holds for different values of N_{ph} . Any lower than that, and not all the 'bad' singular values are removed, such that the retrieval is still not successful. Much higher than, and the information of the original maps is partly lost as well, which is especially evident in the retrieved reflectance maps of the low-amplitude Fresnel reflection for SVCR = 0.5 and SVCR = 0.1.

Another interesting question is how the average number of photons N_{ph} affects the retrieval of our maps. Figure 5.7 shows the retrieved reflectance maps for different values of N_{ph} when $\alpha = \frac{\pi}{2}$, $\beta = \frac{\pi}{6}$, and $N_t = 3400$. As expected, the accuracy of the retrieved maps increases if the average number of photons increases, as more photons corresponds with more information about the exoplanet's surface. However, this increase in accuracy is not very great, as the shot noise is modelled by the Poisson distribution, and so a higher average also means a higher variance when sampling, which causes the intensity to fluctuate more, impeding our retrieval efforts.

³Recall that $\max\{a_{Fr}\} = 0.1$

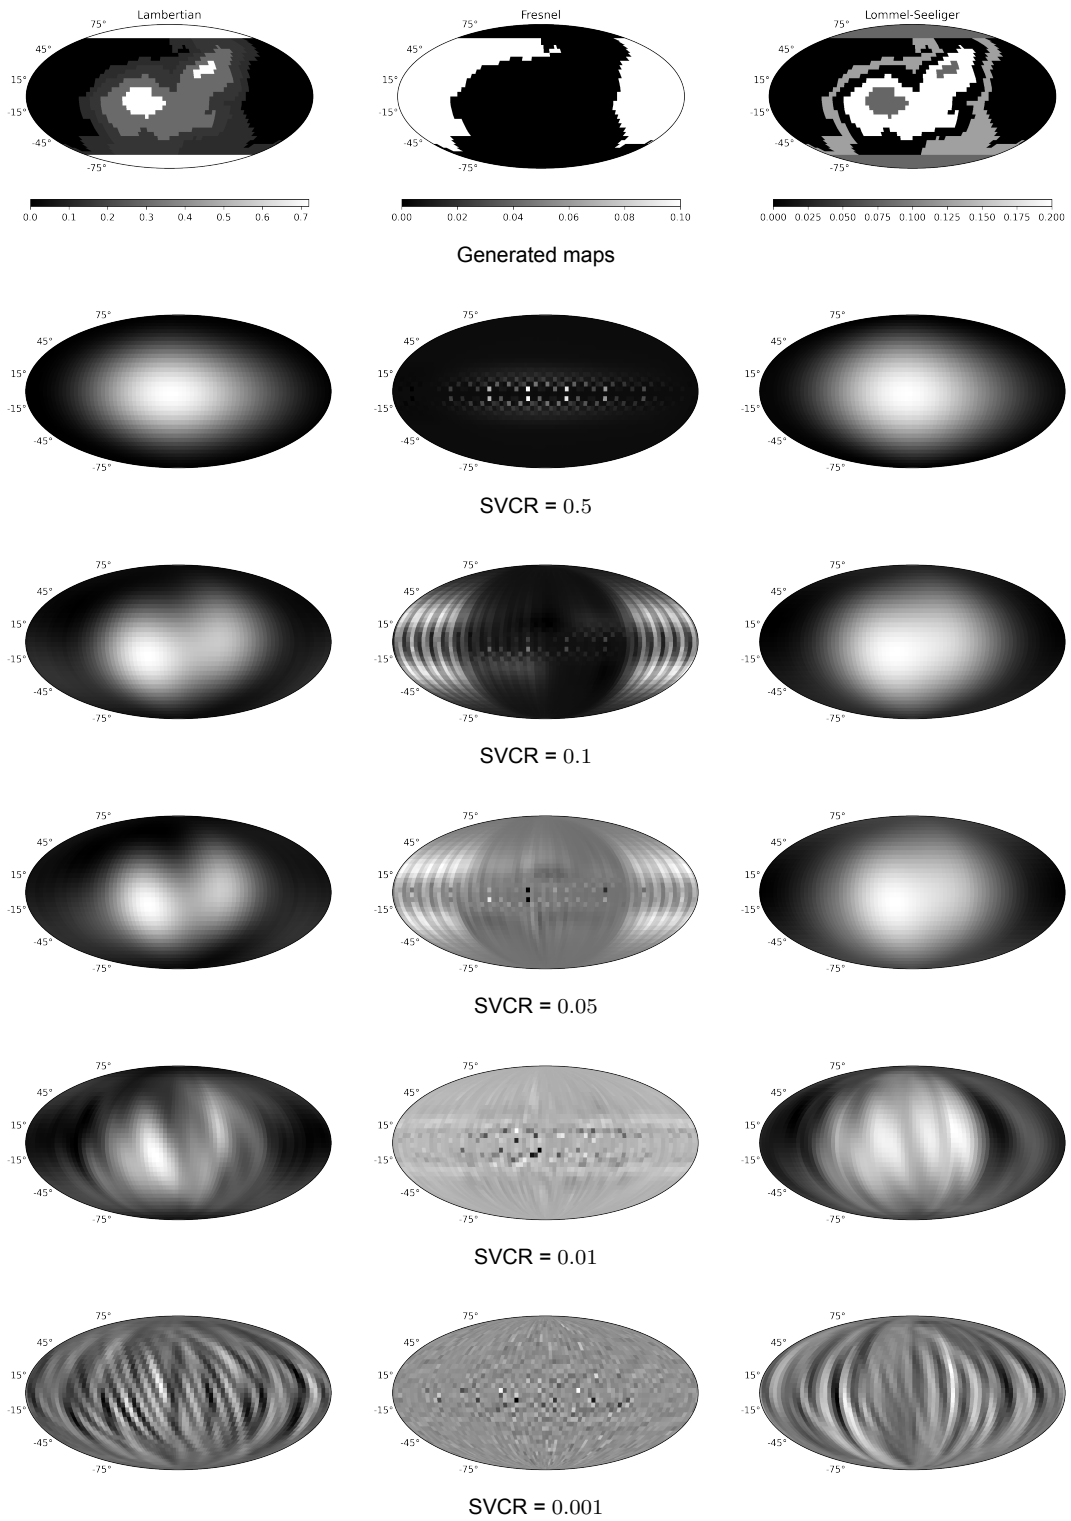


Figure 5.5: Reflectance maps of Planet B retrieved from composite light curves with added shot noise with $N_{ph} = 187$. Furthermore, $\alpha = \frac{\pi}{2}$, $\beta = \frac{\pi}{6}$ and $N_t = 8760$.




Figure 5.6: Reflectance maps of Planet B retrieved from composite light curves with added shot noise with $N_{ph} = 187$. Furthermore, $\alpha = \frac{\pi}{2}$, $\beta = \frac{\pi}{6}$ and $N_t = 3400$.

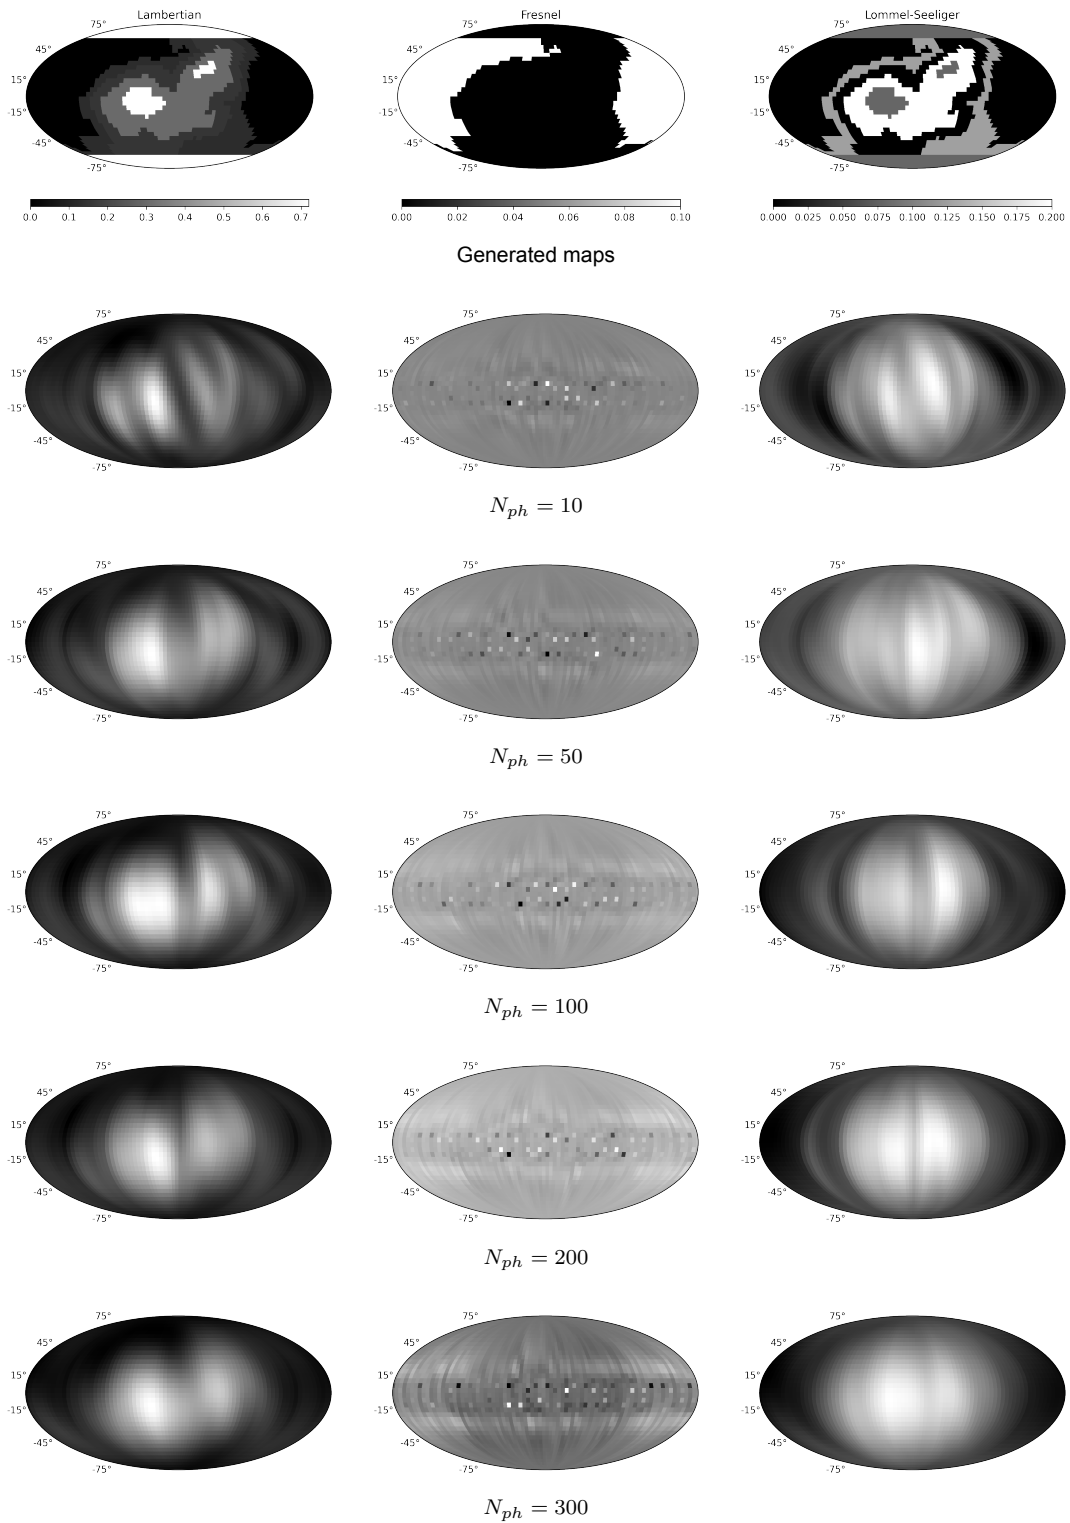


Figure 5.7: Reflectance maps of Planet B retrieved from composite light curves with added shot noise for different number of average photons N_{ph} . Furthermore, $\alpha = \frac{\pi}{2}$, $\beta = \frac{\pi}{6}$ and $N_t = 3400$.

5.4. Comparing the reflection models

As mentioned throughout this thesis, previous research around Exocartography has assumed completely diffusive reflection following the Lambertian model, and so the transformation matrices were constructed according to this simple model. In this thesis, we have introduced a new model of composite reflection that accounts for the mirror-like reflection of water and the bright reflection of rocky terrains. In this final part, we investigate and compare how well these two reflection models manage to retrieve the surface maps in different scenarios, to conclude if one model performs better than the other in general. To achieve this we first let both the simple and the composite model retrieve the maps from the same sampled noiseless signal of light reflected according to the composite model. We then compare the relative error in the light curve produced by the retrieved maps against the original light curve for different values of SVCR and axial tilt angles. Next we do the same for a sampled signal of light reflected according to the simple model, which earlier research suggests is the correct one.

5.4.1. Retrieval for composite light curves

First, let's assume that the light we observe is indeed reflected off the exoplanet in accordance with our composite reflection model, i.e. a (weighted) combination of Lambertian, Fresnel, and Lommel-Seeliger reflection. In that case, the real intensity is equal to the total transformation matrix multiplied with the concatenated reflection maps:

$$\mathbf{I}_{TOT} = \mathbf{T}_{TOT} \mathbf{M}_{TOT}. \quad (5.6)$$

Therefore, if we assume composite reflection when retrieving, the retrieved reflectance maps $\mathbf{M}_{r,TOT}$ can be computed by:

$$\mathbf{M}_{r,TOT} = \mathbf{T}_{TOT}^+ \mathbf{I}_{TOT} = \mathbf{T}_{TOT}^+ \mathbf{T}_{TOT} \mathbf{M}_{TOT} \quad (5.7)$$

However, if we assume simple Lambertian reflection when retrieving, the retrieved reflectance maps are now computed by:

$$\mathbf{M}_{r,La} = \mathbf{T}_{La}^+ \mathbf{I}_{TOT} = \mathbf{T}_{La}^+ \mathbf{T}_{TOT} \mathbf{M}_{TOT} \quad (5.8)$$

Figure 5.8 shows the retrieved reflectance maps from a noiseless composite light curve, assuming either simple or composite reflection. In each case, $N_t = 3400$, $\alpha = \frac{\pi}{2}$, and $\beta = \frac{\pi}{3}$, and the retrieved reflectance maps are shown for varying values of the SVCR. In the case of simple retrieval, which only returns the Lambertian reflected map, the retrieved Fresnel and Lommel-Seeliger maps are manually assigned to zero to complete the figure.

When studying the retrieved reflectance maps in Figure 5.8, we find that the Lambertian model does not accurately retrieve the original Lambertian map if the sampled light curve is composite. For large SVCRs, both the simple and composite retrieval produce unclear Lambertian maps, which appear nearly identical. However, the light curve is sampled without noise, and so a smaller SVCR should be warranted. For $SVCR = 10^{-5}$, the composite retrieval is able to retrieve all three reflectance maps quite accurately, while the simple model is unable to retrieve the original Lambertian map. As the simple model assumes all information in the light curve \mathbf{I}_{TOT} to only come from a single Lambertian map, it erroneously interprets the reflected light contributed by the other 2 models as Lambertian, and so it produces a seemingly random Lambertian map, not unlike those when we attempted to retrieve the reflectance maps with added noise. It follows that the Lambertian retrieval model does not achieve accurate retrieval if the sampled light curve is composite.

5.4.2. Retrieval for simple Lambertian light curves

Next, let's assume that the observed light is reflected off the exoplanet solely according to the simple Lambertian reflection, as most studies have. In that case, the real intensity is computed by:

$$\mathbf{I}_{La} = \mathbf{T}_{La} \mathbf{M}_{La}. \quad (5.9)$$

Therefore, if we assume composite reflection when retrieving, the retrieved reflectance maps $\mathbf{M}_{r,TOT}$ can be computed by:

$$\mathbf{M}_{r,La} = \mathbf{T}_{La}^+ \mathbf{I}_{La} = \mathbf{T}_{La}^+ \mathbf{T}_{La} \mathbf{M}_{La} \quad (5.10)$$

However, if we assume composite reflection when retrieving, the retrieved reflectance maps are now computed by:

$$\mathbf{M}_{r,TOT} = \mathbf{T}_{TOT}^+ \mathbf{I}_{La} = \mathbf{T}_{TOT}^+ \mathbf{T}_{La} \mathbf{M}_{La} \quad (5.11)$$

Figure 5.9 shows the retrieved reflectance maps from a noiseless simple light curve, assuming either simple or composite reflection. In each case, $N_t = 3400$, $\alpha = \frac{\pi}{2}$, and $\beta = \frac{\pi}{3}$, and the retrieved reflectance maps are shown for varying values of the SVCR. As before, in the case of simple retrieval, the retrieved Fresnel and Lommel-Seeliger maps are manually assigned to zero to complete the figure.

When analyzing the retrieved reflectance maps in Figure 5.9, we find that the composite model accurately retrieves the Lambertian surface maps, even if the original light curve was modelled by simple Lambertian reflection. For large SVCRs, the simple and composite retrieval again produce nearly identical glossy Lambertian maps, while for a SVCR of 10^{-5} , the simple retrieval retrieves an almost perfect Lambertian map. Interestingly enough, the composite model also retrieves a clear Lambertian map, albeit a little less accurate. The retrieved Lommel-Seeliger reflectance map, which should be zero for perfect retrieval, bears some resemblance to the retrieved Lambertian map. This is not surprising, as the Lommel-Seeliger reflection model is only a variation of diffusive reflection, and so a part of the Lambertian contribution to the light curve is wrongly interpreted as Lommel-Seeliger light. The retrieved Fresnel map is even more noisy and bears no resemblance to any map, not even to the retrieved Fresnel map if the light was reflected compositely. However the retrieved Lambertian map is accurate enough to distinguish continents from oceans and snowy peaks.

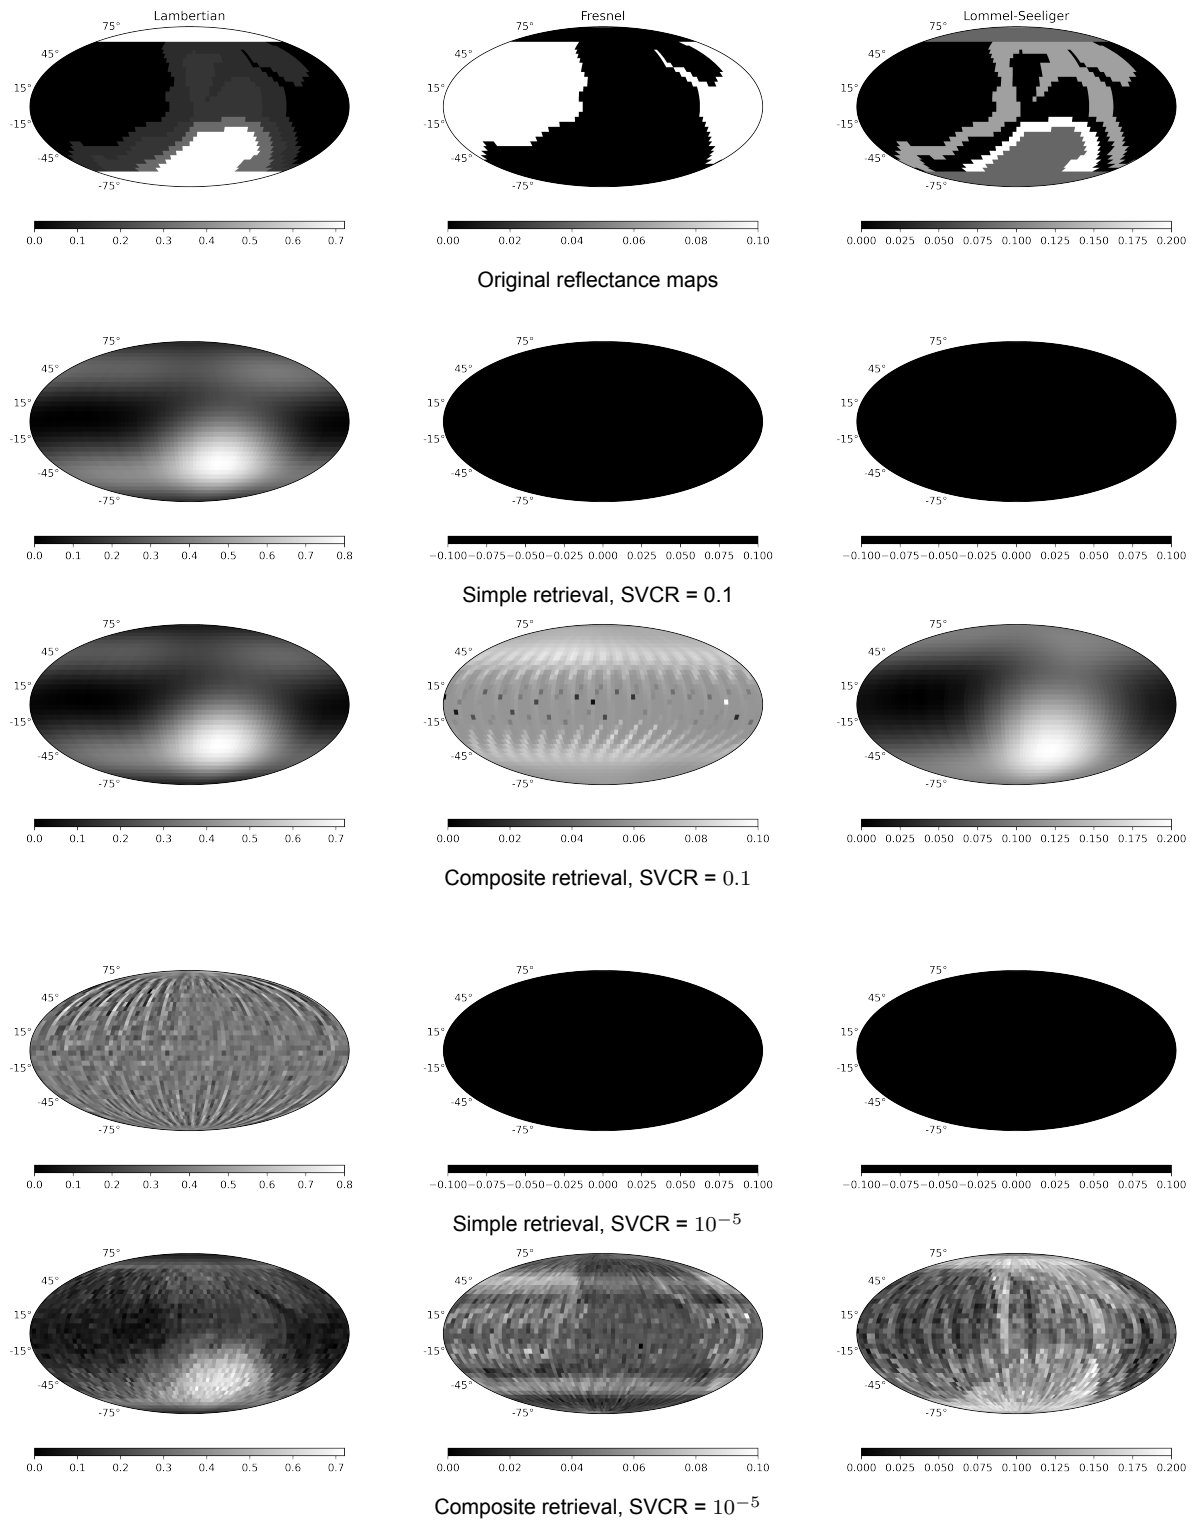


Figure 5.8: Reflectance maps of Planet A retrieved using the simple or composite reflection model from a noiseless composite light curve. Furthermore, $\alpha = \frac{\pi}{2}$, $\beta = \frac{\pi}{3}$ and $N_t = 3400$.

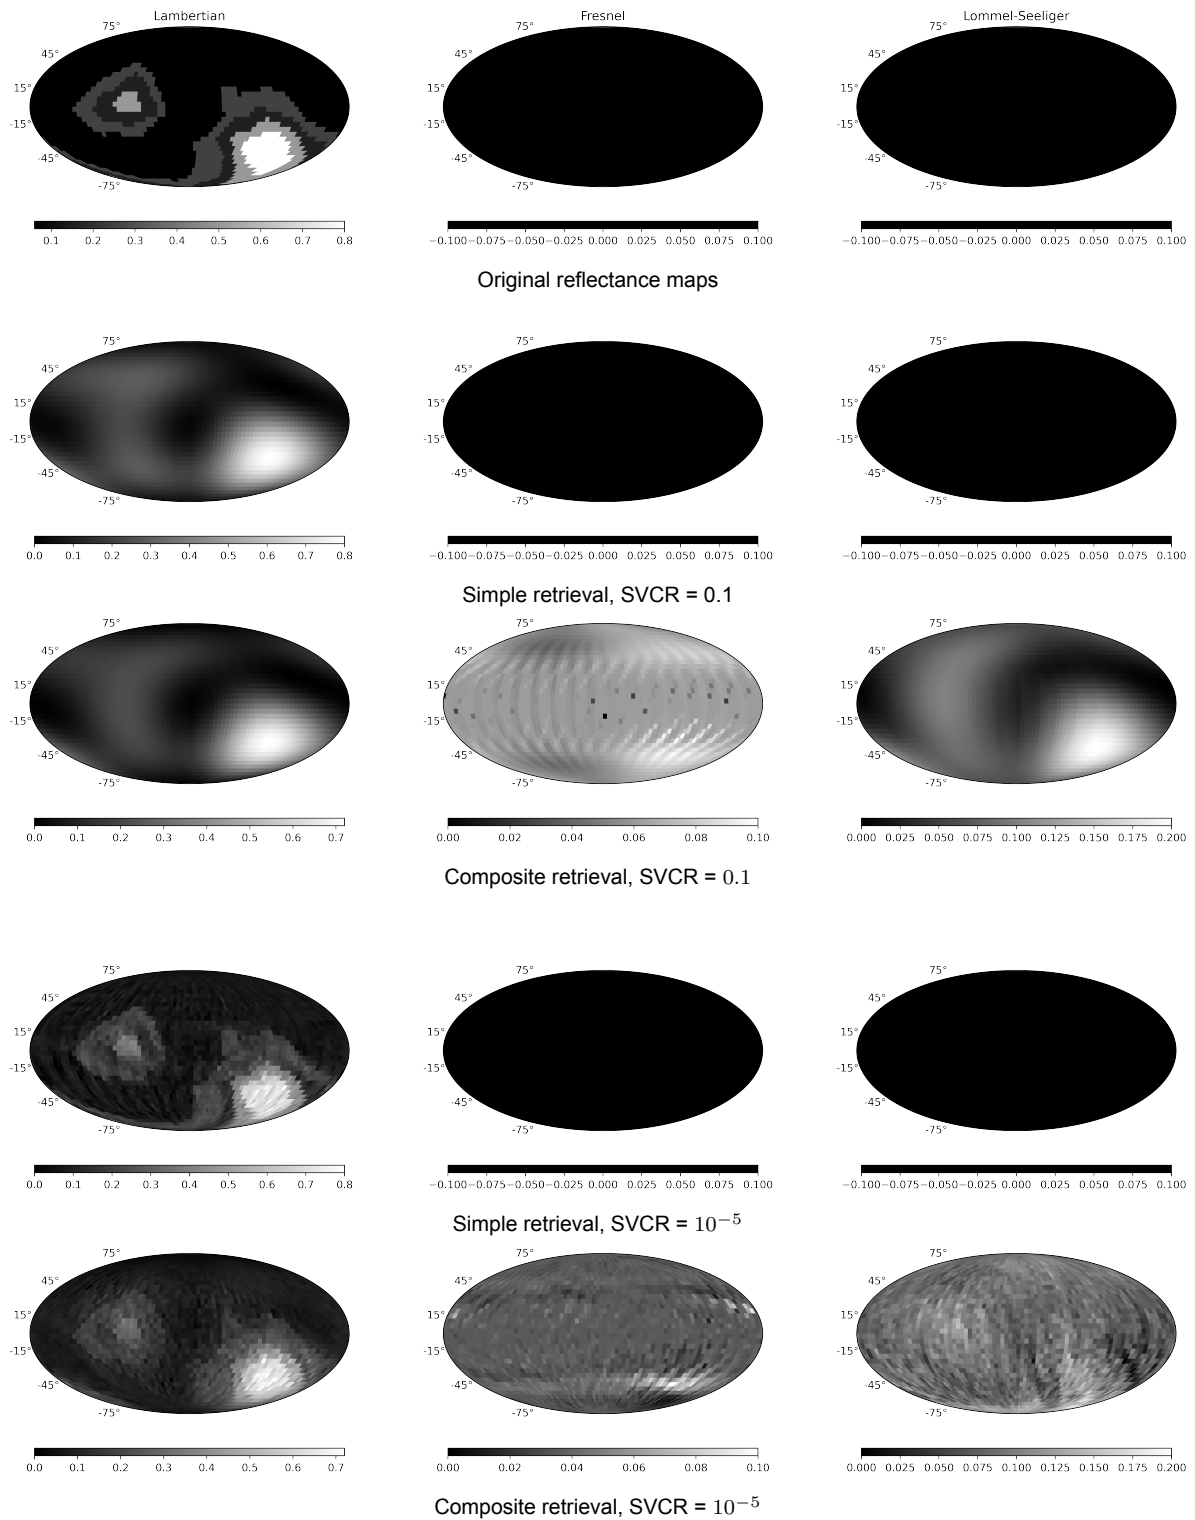


Figure 5.9: Reflectance maps of Planet D retrieved using the simple or composite reflection model from a noiseless simple light curve. Furthermore, $\alpha = \frac{\pi}{2}$, $\beta = \frac{\pi}{3}$ and $N_t = 3400$.

5.4.3. Relative error comparison

In order to determine how accurate the retrieved maps are, apart from just searching for similarities and differences with the generated maps visually, we simulate the noiseless light curve that these maps would produce, and compute the (relative) difference between the observed light curve and the simulated curve to quantitatively determine its error. Suppose a observed light curve \mathbf{I} is assumed to be generated in accordance with a reflection model RM , then the retrieved reflectance map is given by:

$$\mathbf{M}_{r,RM} = \mathbf{T}_{RM}^+ \mathbf{I}. \quad (5.12)$$

It follows that the light curve simulated by the retrieved reflectance map $\mathbf{M}_{r,RM}$ is given by:

$$\mathbf{I}_{r,RM} = \mathbf{T}_{RM} \mathbf{M}_{r,RM} = \mathbf{T}_{RM} \mathbf{T}_{RM}^+ \mathbf{I}, \quad (5.13)$$

and therefore the relative error between observed and simulated light curve, henceforth denoted as E is given by:

$$E = \frac{|\mathbf{I} - \mathbf{I}_{r,RM}|^2}{|\mathbf{I}|^2} = \frac{|(1 - \mathbf{T}_{RM} \mathbf{T}_{RM}^+) \mathbf{I}|^2}{|\mathbf{I}|^2} \quad (5.14)$$


Figure 5.10 and Figure 5.11 show the relative error E for simulated light curves for noiseless composite light curves, while Figure 5.12 and Figure 5.13 show the same error for noiseless simple light curves. In each Figure, the upper image is a scatter plot where the relative error E is plotted as a function of the obliquity β for different values of α , while the lower image is a bar plot for $\alpha = \frac{\pi}{6}$ and $\frac{\pi}{2}$ specifically. In each plot, $N_t = 3400$, and the relative error is computed for different values of SVCR. Furthermore, the labelname 'La' denotes that the signal was treated as a simple Lambertian light curve when retrieving the reflectance maps and simulating its light curve, while the name 'TOT' denotes that the signal was treated as a composite light curve.

As can be seen in the Figures 5.10 to 5.13, the relative error is smallest if the observed light curve is treated as from the same model as it actually is, e.g. when a Lambertian light curve is retrieved and simulated according to the Lambertian reflection model and likewise for composite signals. This makes sense intuitively. Another finding is that E also decreases for smaller values of SVCR, even when the light curve is retrieved and simulated wrongly. One may believe this implies the retrieved maps are more accurate for smaller SVCRs, which we know isn't true from Figure 5.8. Therefore one should always look to the retrieved maps themselves to conclude if the retrieval was accurate.

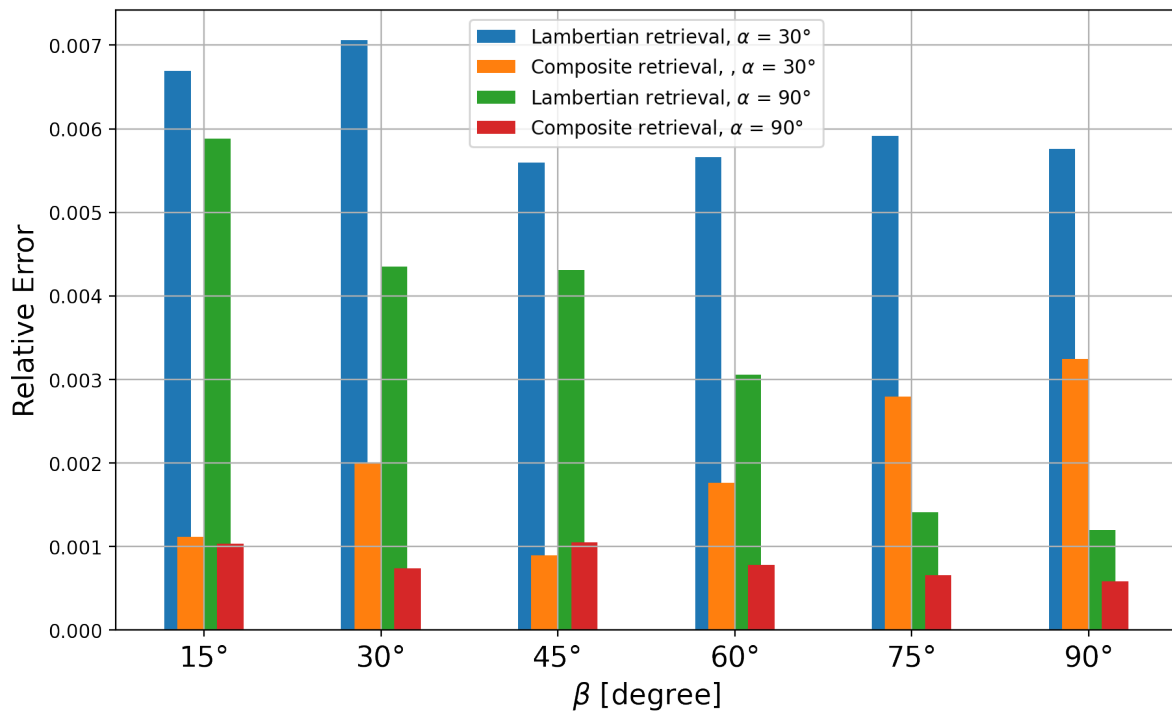
When comparing the different cases for the original light, i.e. if the planet actually reflects light simply or compositely, we can see that the relative errors for simple and composite retrieval are about the same if the original light curve is reflected according to the Lambertian model, with only some outliers for $\beta = 0$ and $\beta = \frac{\pi}{2}$. Indeed, if $SVCR = 0.1$ then the biggest error ratio between composite retrieval and Lambertian retrieval, i.e. $\max \left\{ \frac{E_{TOT}(\alpha, \beta)}{E_{La}(\alpha, \beta)} \right\}$, is about 4.7 for $\alpha = \frac{\pi}{6}$ and $\beta = \frac{\pi}{12}$. On the other hand, if the original light curve is composite, then the relative error E is far greater if the curve is retrieved and simulated according to the Lambertian model than according to the composite model. This is especially evident in Figure 5.11 (a), where the composite relative errors all seem to lie on the x-axis. If we compute the maximum error ratio $\max \left\{ \frac{E_{La}(\alpha, \beta)}{E_{TOT}(\alpha, \beta)} \right\}$ ⁴ again for $SVCR = 0.1$, we get a value of 7.9 for $\alpha = \frac{\pi}{3}$ and $\beta = \frac{\pi}{6}$. We can conclude from these figures that the retrieval using the composite retrieval model generally produces the smallest relative error.

The results of these figures, combined with our findings from Figures 5.8 and 5.9, we conclude that reflectance map retrieval using the composite reflection model outperforms the simple Lambertian reflection model in general.

⁴Notice the difference between the two formulas for the maximum relative error ratio.




(a)




(b)

Figure 5.10: Above: Scatter plots of the relative error of the light curves produced by the retrieved maps for composite light for different axial tilt angles α and β , where $N_t = 3400$ and $SVCR = 0.1$. Here 'La' stands for Lambertian retrieval, while 'TOT' stands for composite retrieval. Below: corresponding bar plot for $\alpha = \frac{\pi}{6}$ and $\frac{\pi}{2}$.




(a)

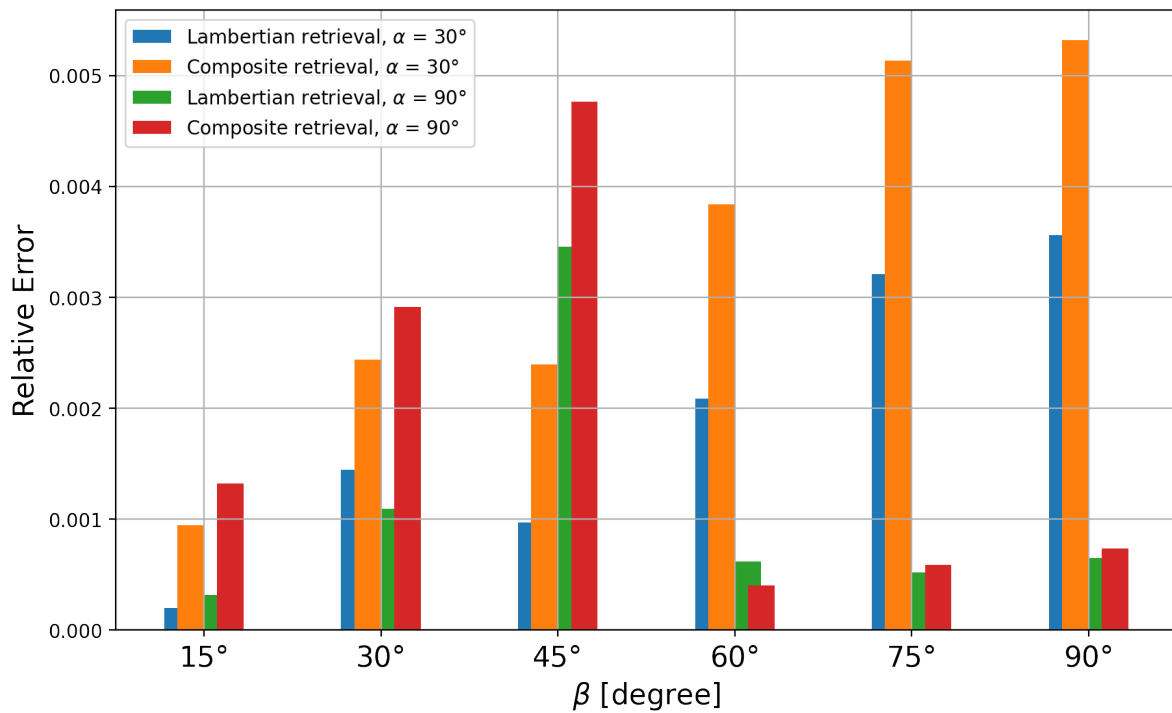


(b)

Figure 5.11: Above: Scatter plots of the relative error of the light curves produced by the retrieved maps for composite light for different axial tilt angles α and β , where $N_t = 3400$ and $\text{SVCR} = 10^{-5}$. Here 'La' stands for Lambertian retrieval, while 'TOT' stands for composite retrieval. Below: corresponding bar plot for $\alpha = \frac{\pi}{6}$ and $\frac{\pi}{2}$.




(a)




(b)

Figure 5.12: Above: Scatter plots of the relative error of the light curves produced by the retrieved maps for Lambertian light for different axial tilt angles α and β , where $N_t = 3400$ and $SVCR = 0.1$. Here 'La' stands for Lambertian retrieval, while 'TOT' stands for composite retrieval. Below: corresponding bar plot for $\alpha = \frac{\pi}{6}$ and $\frac{\pi}{2}$.



(a)



(b)

Figure 5.13: Above: Scatter plots of the relative error of the light curves produced by the retrieved maps for Lambertian light for different axial tilt angles α and β , where $N_t = 3400$ and $\text{SVCR} = 10^{-5}$. Here 'La' stands for Lambertian retrieval, while 'TOT' stands for composite retrieval. Below: corresponding bar plot for $\alpha = \frac{\pi}{6}$ and $\frac{\pi}{2}$.

6

Conclusion

In this thesis we attempted to retrieve the surface maps of Earth-like exoplanets using temporal variations in the intensity of the light of the parent star reflected off the exoplanets surface. We expanded on the known inversion technique of spin-orbit tomography with Lambertian reflection by introducing two new reflection models, namely Fresnel reflection to account for the mirror-like behaviour of large bodies of water, and Lommel-Seeliger reflection as an empirical model congruent with observations of rocky heavenly bodies at low phase angles. These reflection models were studied by analyzing their bidirectional reflectance distribution functions and their phase functions.

Using the pseudo-random algorithm of tetrahedral subdivision, we were able to generate realistic altitude maps for multiple planets. Using these altitudes, we then divided the planet's surface into different terrain types likely present on habitable planets, such as oceans and grasslands. Next we assigned a measure of reflectance to each surface type for all 3 reflection models, and so we could construct three distinct reflectance maps.

Combining the generated reflectance maps with the analyzed phase angles, we were able to successfully construct the transformation matrix for discrete surface maps and subsequently generate mock light curves of the composite intensity of the reflected starlight, in which the contributions of the different reflection models were clearly visible. The light curve could be simulated for varying angles of the rotational axis of the planet, as well as for different number of samples taken. In order to emulate the stochastic nature of photon observation, we also included the option to impose shot noise on the signal.

We applied the Moore-Penrose pseudoinverse combined with a cutoff for the singular values in order to invert the transformation matrix and retrieve all three reflectance maps. This expanded approach proved successful when retrieving the maps from noiseless light curves, even if the the number of samples was significantly smaller than three times the planet surface's resolution. The accuracy of the retrieved maps for Lambertian and Lommel-Seeliger reflection increased for larger values of the planets obliquity β , especially if the equinox angle α was large. The retrieved Fresnel reflectance maps were only accurate on the lateral band around the equator with width 2β , as that is the region from which the specularly reflected light can reach the observer. Furthermore, for β close to $\frac{\pi}{2}$, the Fresnel map showed a strip where our model fails to retrieve the surface, as this region is only illuminated head-on for a very short instance in the planets orbit when the phase angle is nearly zero.

When artificial shot noise was added, the map retrieval was much poorer, and a high cutoff ratio of the singular values was needed to retrieve glossy reflectance maps for Lambertian and Lommel-Seeliger reflection. It was found that the optimal cutoff ratio was about 0.05. Retrieving the Fresnel reflectance map from noisy light curves proved especially difficult, as its contribution to the light curve is smallest and therefore easiest interpreted as noise.

We compared our new composite model with the one used by earlier research that only implements Lambertian reflection. We simulated a noiseless light curve assuming only Lambertian reflection, and then attempted to retrieve the single reflectance map with both the Lambertian and the composite transformation. We found that our composite model correctly retrieved the original Lambertian map, even for small values of the cutoff ratio, while the erroneously retrieved Fresnel and Lommel-Seeliger maps

were essentially meaningless. Furthermore, when calculating the relative error of the light curve simulated by the retrieved maps, if the original light curve was produced according to Lambertian reflection, we found that the relative error of the erroneously simulated composite signal was not bad, with a maximum relative error ratio of about 4.7

On the other hand, if the original light curve was constructed with our composite model, then the Lambertian reflection model was unable to retrieve an accurate Lambertian map for small singular values cutoff ratios, as the Lambertian retrieval is believed to treat the Fresnel and Lommel-Seeliger intensity contributions as noise. Moreover, when comparing the re-simulated light curves, it was found that the Lambertian retrieval produced significantly worse errors than if the light curve was retrieved and re-simulated compositely, with a maximum relative error ratio of 7.9. Therefore we conclude that the composite reflection model appears to be a viable tool for retrieving exoplanet surface maps, even if the original light curve is not identical to our expectations.


We have found the following recommendations for further research:

- Introduce the wavelength of the incident and reflected light rays as a parameter for both the bidirectional reflectance distribution functions as well as the surface albedos and attempt to better retrieve the different surface types through spectral analysis.
- Include model atmospheres on simulated planets and study how it affects the incident and reflected light through diffraction.
- Simulate elliptical orbits and introduce the eccentricity of the planet as a parameter to produce more realistic light curves.
- Include the Relativistic Doppler effect when the exoplanet approaches or recedes from the observer in edge-on observations, and try to retrieve the duration of its orbital period empirically.
- Study how the composite model fares retrieving the reflectance maps for different inclination angles, e.g. when the planet's orbit is observed face-on.
- Let each surface type have a range of reflectance values for each reflection model, instead of only 1 value per surface type.
- Implement an equal-area pixelization scheme for the surface facets, such as HEALPIX or the Fibonacci spiral, instead of the equi-rectangular facet scheme used in this thesis.

References

- [1] André Berger and Marie-France Loutre. “Insolation values for the climate of the last 10 million years”. In: *Quaternary Science Reviews* 10.4 (1991), pp. 297–317.
- [2] Nicolas B Cowan and Yuka Fujii. “Mapping exoplanets”. In: *arXiv preprint arXiv:1704.07832* (2017).
- [3] Arthur N Cox. *Allen’s astrophysical quantities*. Springer, 2015.
- [4] Dorian. *Représentation des angles d’Euler avec TikZ*. 2022. URL: <https://nsidc.org/cryosphere/seaice/processes/albedo.html> (visited on 07/14/2022).
- [5] Bernardt Duvenhage, Kadi Bouatouch, and Derrick G Kourie. “Numerical verification of bidirectional reflectance distribution functions for physical plausibility”. In: *Proceedings of the South African Institute for Computer Scientists and Information Technologists Conference*. 2013, pp. 200–208.
- [6] Maxwell B Fairbairn. “Planetary photometry: The Lommel-Seeliger law”. In: *Journal of the Royal Astronomical Society of Canada* 99 (2005), p. 92.
- [7] Siteng Fan et al. “Earth as an exoplanet: A two-dimensional alien map”. In: *The Astrophysical Journal Letters* 882.1 (2019), p. L1.
- [8] Alejandro Ferrero et al. “Spectral and geometrical variation of the bidirectional reflectance distribution function of diffuse reflectance standards”. In: *Applied Optics* 51.36 (2012), pp. 8535–8540.
- [9] Yuka Fujii and Hajime Kawahara. “Mapping Earth analogs from photometric variability: spin–orbit tomography for planets in inclined orbits”. In: *The Astrophysical Journal* 755.2 (2012), p. 101.
- [10] Jonathan P Gardner et al. “The James Webb Space Telescope”. In: *Space Science Reviews* 123.4 (2006), pp. 485–606.
- [11] Roberto Gilmozzi and Jason Spyromilio. “The European Extremely Large Telescope (E-ELT)”. In: *The Messenger* 127.11 (2007), p. 3.
- [12] Gene H Golub and Charles F Van Loan. *Matrix computations*. JHU press, 2013.
- [13] Ian von Hegner. “A limbus mundi elucidation of habitability: the Goldilocks Edge”. In: *International Journal of Astrobiology* 19.4 (2020), pp. 320–329.
- [14] Hajime Kawahara and Yuka Fujii. “Global mapping of Earth-like exoplanets from scattered light curves”. In: *The Astrophysical Journal* 720.2 (2010), p. 1333.
- [15] Benoit B Mandelbrot and Benoit B Mandelbrot. *The fractal geometry of nature*. Vol. 1. WH Freeman New York, 1982.
- [16] Tom Markvart and Luis Castaner. “Practical handbook of photovoltaics. Fundamentals and applications”. In: (2003).
- [17] Michel Mayor and Didier Queloz. “A Jupiter-mass companion to a solar-type star”. In: *Nature* 378.6555 (1995), pp. 355–359.
- [18] William Ross McCluney. *Introduction to radiometry and photometry*. Artech House, 2014.
- [19] K Meinke, DM Stam, and PM Visser. “Exoplanet Cartography using Convolutional Neural Networks”. In: *arXiv preprint arXiv:2204.11821* (2022).
- [20] Torben Ægidius Mogensen. “Planet map generation by tetrahedral subdivision”. In: *International Andrei Ershov Memorial Conference on Perspectives of System Informatics*. Springer. 2009, pp. 306–318.
- [21] NASA. *Exoplanet Archive*. 2022. URL: <https://exoplanetarchive.ipac.caltech.edu/> (visited on 07/14/2022).

-
- [22] Fred E Nicodemus. “Directional reflectance and emissivity of an opaque surface”. In: *Applied optics* 4.7 (1965), pp. 767–775.
- [23] NSIDC. *Albedo*. 2022. URL: <https://nsidc.org/cryosphere/seaice/processes/albedo.html> (visited on 07/14/2022).
- [24] Hugo Seeliger. *Zur Theorie der Beleuchtung der grossen Planeten, insbesondere des Saturn*. Vol. 16. Verlag der k. Akademie in Comm. bei G. Franz, 1887.
- [25] Samuel Stuger. “Exoplanet surface mapping using scattered light curves”. In: (2021).
- [26] John Robert Taylor. *Classical mechanics*. 531 TAY. 2005.
- [27] George E Williams. “History of the Earth’s obliquity”. In: *Earth-Science Reviews* 34.1 (1993), pp. 1–45.



Appendix

A.1. Derivation of Lambertian phase function.

The derivation for the phase function for Lambertian reflectance, as described in Section 2.3, can be found below.

$$\begin{aligned}
 F_L(\zeta) &= 4 \iint_{\mathcal{D}} f_L(\theta_i, \phi_i; \theta_r, \phi_r) \cos(\Phi - \zeta) \cos \Phi \sin^2 \Theta d^2 \hat{\mathbf{s}} \\
 &= \frac{4}{\pi} \int_{-\pi/2+\zeta}^{\pi/2} \int_0^{\pi} \cos(\Phi - \zeta) \cos \Phi \sin^2 \Theta \sin \Theta d\Theta d\Phi \\
 &= \frac{4}{\pi} \int_0^{\pi} \sin^2 \Theta \sin \Theta d\Theta \int_{-\pi/2+\zeta}^{\pi/2} \cos(\Phi - \zeta) \cos \Phi d\Phi \\
 &= \frac{4}{\pi} \int_0^{\pi} (1 - \cos^2 \Theta) \sin \Theta d\Theta \int_{-\pi/2+\zeta}^{\pi/2} \cos(\Phi - \zeta) \cos \Phi d\Phi \\
 &= \frac{4}{\pi} \int_1^{-1} (u^2 - 1) du \int_{-\pi/2+\zeta}^{\pi/2} \cos(\Phi - \zeta) \cos \Phi d\Phi \\
 &= \frac{4}{\pi} \left[\frac{u^3}{3} - u \right]_1^{-1} \left[\frac{2\Phi \cos \zeta - \sin(\zeta - 2\Phi)}{4} \right]_{-\pi/2+\zeta}^{\pi/2} \\
 &= \frac{16}{3\pi} \left(\frac{\sin \zeta + (\pi - \zeta) \cos \zeta}{2} \right) \\
 &= \frac{8}{3\pi} [\sin \zeta + (\pi - \zeta) \cos \zeta]
 \end{aligned} \tag{A.1}$$

A.2. Tetrahedral subdivision algorithm

The detailed description of the tetrahedral subdivision algorithm, as introduced in Section 3.1.1, can be found below.

Algorithm 1 Tetrahedral subdivision

Input: A target point p on the sphere and four vertices v_1, \dots, v_4 of a tetrahedron, each with the following information:

- Coordinates (x_i, y_i, z_i)
- Seed s_i for the pseudo-random generator
- Altitude value h_i

Repeat:

1. Re-order the vertices so that the longest edge e_{max} of the tetrahedron is between v_1 and v_2
2. Create new vertex v_m by:

- $(x_m, y_m, z_m) = ((x_1 + x_2)/2, (y_1 + y_2)/2, (z_1 + z_2)/2)$
- $s_m = random((s_1 + s_2)/2)$
- $l = |e_{max}|$
- $h_m = (h_1 + h_2)/2 + 0.01 \times random(s_m) \times l^{\frac{3}{2}}$

where $random(s)$ generates a pseudo-random value from a seed s

3. **if** p is contained in the tetrahedron defined by v_m, v_1, v_3 and v_4 **then**
4. **set** $v_2 = v_m$
5. **else**
6. **set** $v_1 = v_m$
7. **end if**

Stop if: l is small enough

Output: $(h_1 + h_2 + h_3 + h_4)/4$, the final altitude for point p .

A.3. Code

All of the code that has been used in this report is available on request via mail to: S.W.S.M.Blanken@student.tudelft.nl