

Composition of the surface of Pluto by spectro-imagery. New Horizons mission

Leila Gabasova

► To cite this version:

Leila Gabasova. Composition of the surface of Pluto by spectro-imagery. New Horizons mission. Galactic Astrophysics [astro-ph.GA]. Université Grenoble Alpes [2020-..], 2021. English. NNT: 2021GRALU025. tel-03508177

HAL Id: tel-03508177 https://theses.hal.science/tel-03508177

Submitted on 3 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



THÈSE

Pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ GRENOBLE ALPES

Spécialité : Sciences de la Terre et de l'Univers et de l'Environnement

Arrêtée ministériel : 25 mai 2016

Présentée par

Leila GABASOVA

Thèse dirigée par **Bernard SCHMITT** Directeur de Recherche, Université Grenoble Alpes

préparée au sein de l'Institut de Planetologie et d'Astrophysique de Grenoble dans l'École Doctorale Sciences de la Terre, de l'Environnement et des Planètes

Composition de la surface de Pluton par spectro-imagerie. Mission New Horizons Composition of the surface of Pluto by spectro-imagery. New Horizons mission

Thèse soutenue publiquement le 27 septembre 2021, devant le jury composé de :

Eric QUIRICO IPAG (Grenoble), Président François FORGET LMD (Paris 6), Rapporteur Frédéric SCHMIDT GEOPS (Paris-Saclay), Rapporteur Jocelyn CHANUSSOT GIPSA (Grenoble), Examinateur Silvia PROTOPAPA SWRI (Boulder, USA), Examinatrice Bernard SCHMITT IPAG (Grenoble), Directeur de thèse



Résumé

En 2015, la sonde New Horizons effectua une approche de Pluton et nous permit d'obtenir des informations sans précédent sur sa surface et ses propriétés géologiques. Les mesures effectuées ont déjà permis l'élaboration de cartes globales de réflectance et de topographie et de cartes multispectrales à bande étroite de la composition de la surface. L'analyse des données hyperspectrales du spectromètre infrarouge Ralph/LEISA, qui nous permet d'analyser la composition de la surface, a cependant été confiné jusqu'ici à l'hémisphère de rencontre où nous disposons des données en haute résolution. De plus, ces analyses ont été produites largement de manière qualitative et sans essayer de déterminer la proportion et les modes de mixages des différents composants.

Cette thèse de doctorat utilise une nouvelle technique inspirée de l'imagerie médicale afin de présenter les premières cartes globales de composition de la surface de Pluton. Nous effectuons ensuite une analyse statistique de ces cartes et les comparons avec d'autres bases de données sur la surface de Pluton afin d'en proposer des interprétations géologiques. Ce jeu de données a déjà servi à la modélisation de l'atmosphère de Pluton et semble être un outil prometteur pour l'élaboration de nouvelles descriptions et prédictions du climat plutonien.

Enfin, nous développons une méthode métaheuristique basée sur le recuit simulé pour faire de l'inversion spectrale en utilisant le modèle de transfert radiatif DISORT. Cela nous permet d'établir des résultats quantitatifs sur la composition des sols, et nous présentons les premiers résultats issus de cette méthode. Nous analysons les coûts liés à l'élaboration d'une carte de composition globale à partir de cette méthode, et proposons des améliorations algorithmiques afin d'en améliorer l'efficacité et la précision. Les coûts de calculs étant particulièrement élevés, nous proposons aussi un plan de développement permettant de transformer ces premiers résultats en projet de recherche citoyenne en calcul distribué, qui permettrait d'obtenir à moindre coût une carte globale de précision adéquate d'ici quelques années.

Abstract

In 2015 the New Horizons spacecraft reached the Pluto system and returned unprecedentedly detailed measurements of its surface properties. These measurements have already been integrated into global reflectance, topography and narrow-band multispectral surface composition maps. However, analysis of the hyperspectral data from the Ralph/LEISA infrared spectrometer, which lets us analyse the surface composition, has until now been confined to the high-resolution encounter hemisphere of Pluto, and has mostly been carried out in a qualitative fashion, without a detailed analysis of the proportions and mixture types of the various components.

This thesis uses the technique of intensity-based registration, commonly used in medical imagery and newly applied in planetary science, to present the first global qualitative composition maps for the main materials of Pluto's surface. We then carry out a statistical analysis of these maps and compare them with the other Pluto surface datasets to make geological interpretations of the maps. This dataset has already been used for global atmospheric modeling of Pluto and promises to be a valuable asset for further descriptions and predictions of its climate and behaviour.

We also present a metaheuristic, simulated-annealing-based method for spectral inversion, using the DISORT radiative transfer model. We use this method to develop the qualitative LEISA dataset into quantitative compositional maps, and show the first few results using this method. Finally, we discuss the computational costs and runtime needed to produce the global quantitative compositional map of Pluto, and suggest algorithmic improvements to make the spectral inversion more accurate and more efficient. As the computational costs are very high, we outline a path to developing these first results into a distributed computing project that would make completing the global compositional map feasible in a few years.

Acknowledgements

First and foremost, Bernard, thank you for your tireless guidance in every aspect of my work for the past four years. You've been very patient and understanding in regards to all the unexpected spokes in our wheels on this project, and I hope both we and others can keep developing it into a groundbreaking body of work.

Thank you also to my colleagues at IPAG for providing a welcoming and friendly work environment. I want especially to thank Robin, Jolantha and Ricardo for all the lunches and other friendly shenanigans, and wish two of you courage for the end of the road, and congratulate the third with getting there ahead of us.

To the New Horizons team, thank you for four amazing years of discoveries and for sharing your infectious enthusiasm. I had one of the most exciting times of my life at the Arrokoth flyby, and your work has inspired me not only for the work in this manuscript, but for other hopeful future projects on all manner of planetary topics. I only wish I'd had more opportunities to see you all in person before the end of my thesis.

Massive thanks to Alexandra Elbakyan, someone I've never met or spoken to who has nonetheless been integral to my ability to participate in research — and I am only one of millions of researchers whose life has been transformed by Elbakyan's work.

To my Enka, no amount of thanks can be enough for all your love, support, turtles, encouragement, and math. It's been a very thorny road, and my courage to continue flagged many times but you kept dragging me through the thorns and now we're on the other side where the grass is green and the honorifics are gender-neutral, and I'm incredibly grateful and honestly surprised you put up with all that.

To my N, thank you for cheerleading and love from afar and cross animal dates. S o o n.

To Flo, thank you from the bottom of my heart for all the tea.

To all my other friends and my family, thank you for your love and support. I miss you!

To Wisp, Bramble, Esker, and all the millipedes, thank you for being perfect and beautiful tubes.

Finally (but no less importantly), thank you to the Baltimore Crabs.

Contents

| Acknowledgements 1 | | | | | | | | |
|--------------------|-----------------------|-----------------------------------|--|----|--|--|--|--|
| In | trodu | iction | | 4 | | | | |
| I | I General information | | | | | | | |
| | 1 | Spectro | oscopy in planetary science | 6 | | | | |
| | | 1.1 | Remote sensing | 6 | | | | |
| | | 1.2 | Hyperspectral imaging | 7 | | | | |
| | | 1.3 | Absorption and reflectance spectroscopy | 10 | | | | |
| | | 1.4 | Radiative transfer models | 15 | | | | |
| | 2 | New H | orizons mission to Pluto | 18 | | | | |
| | | 2.1 | Introduction to the Pluto system | 19 | | | | |
| | | 2.2 | The New Horizons probe | 23 | | | | |
| | | 2.3 | 2015 Pluto flyby | 27 | | | | |
| | | 2.4 | Data from the encounter | 29 | | | | |
| II | Gloł | oal map | oregistration | 35 | | | | |
| | 1 | Introd | uction | 35 | | | | |
| | 2 | Data a | nd methods | 35 | | | | |
| | | 2.1 | LEISA approach datasets | 35 | | | | |
| | | 2.2 | Map registration | 38 | | | | |
| | | 2.3 | Registration algorithm details | 40 | | | | |
| | | 2.4 | Validation of registration accuracy | 41 | | | | |
| | 3 | Results | s and analysis | 42 | | | | |
| | | 3.1 | Global maps and their main features | 42 | | | | |
| | | 3.2 | Analysis of correlations between the LEISA maps | 47 | | | | |
| | | 3.3 | Comparison with MVIC maps | 50 | | | | |
| | | 3.4 | Composition relationship with geology | 53 | | | | |
| | | 3.5 | Utility of global composition maps for atmospheric circulation | 56 | | | | |
| 4 9 | | Summ | ary | 57 | | | | |
| | | 4.1 | Global compositional cartography | 57 | | | | |
| | | 4.2 | Evaluation of intensity-based registration | 57 | | | | |
| | | 4.3 | Applications for quantitative global mapping of Pluto | 57 | | | | |
| III | Spec | tral mo | odeling and inverse fit | 58 | | | | |
| | 1 | Background and problem definition | | | | | | |
| | | 1.1 | Summary of previous work | 58 | | | | |
| | | 1.2 | The problem of high-dimensional optimisation | 58 | | | | |
| | 2 | Metho | dology | 59 | | | | |
| | | 2.1 | Radiative transfer models | 59 | | | | |
| | | 2.2 | Metaheuristics | 61 | | | | |
| | | 2.3 | Algorithm description and radiative transfer model setup | 63 | | | | |
| | | 2.4 | The CIMENT grid | 65 | | | | |

| 3 | Results | 66 | | | | |
|---|--|----|--|--|--|--|
| | 3.1 Calibration on synthetic data | 66 | | | | |
| | 3.2 Compositional endmembers | 74 | | | | |
| | 3.3 Profile gradient | 79 | | | | |
| IV Costs and methods for global map inversion | | | | | | |
| 1 | Computational costs and resources | 85 | | | | |
| | 1.1 Upper and lower bounds on computing the global Pluto map | 85 | | | | |
| | 1.2 Can supercomputers help? | 86 | | | | |
| | 1.3 Distributed computing | 87 | | | | |
| 2 | Algorithmic improvements and building a global picture | 88 | | | | |
| | 2.1 Reuse of intermediate spectra | 88 | | | | |
| | 2.2 Reducing surface dimensionality | 89 | | | | |
| | 2.3 Techniques for image segmentation | 89 | | | | |
| | 2.4 Graph colouring | 90 | | | | |
| 3 | Global summary and priorities for moving forward | 91 | | | | |
| List of symbols | | | | | | |
| List of Tables | | | | | | |
| List of Figures | | | | | | |
| Bibliography | | | | | | |
| Résumé de la thèse en français 100 | | | | | | |

Introduction

The most important places on a map are the places we haven't been yet. *The Map of Salt and Stars*, Zeyn Joukhadar

The early days of space exploration were an era of plenty and scarcity. Two global superpowers, the USA and the USSR, were racing to be the first to every exploration milestone, and dedicated more money to the space industry than any country has ever since. Aerospatial engineering flourished, and breakthroughs were duly made, year after year, in terms of safely taking humans to outer space and back again, bringing them further and further out, and keeping them there for longer and longer time periods. Just the fact of being in orbit and on the Moon, and being able to take measurements and bring back samples, gave us unprecedented data and allowed us to confirm or, conversely, demolish, many of our theories about the Solar System. For instance, isotope ratio measurements of the Apollo II lunar regolith samples were key to reconstructing the differentiation of the Moon (Wood et al., 1970), which signified that its early history was significantly more violent and volatile than expected.

These early days, however, were only secondarily about scientific data — the primary goal of these missions was winning geopolitical races and planting flags, and the science done along the way, while groundbreaking, was almost incidental in some ways. This subsidiarity of goals was coupled with the fact that this was also the dawn of the computer era, and our ability to deploy autonomous machines to gather and return information was limited. Data gathered by the Apollo missions relied mostly on the astronauts' firsthand observations and sample return.

These first observations were relatively soon complemented by autonomous missions including the Luna programme of the 1960s and the Mariner and Viking programmes that explored Mars between 1969 and 1975. Even then, their recording and transmission capacities were minuscule compared to modern missions' capabilities. And so, this was an era of want amid plenty — just as every bit of new information was groundbreaking and as likely to overturn existing fields of research as to found new ones, every new dataset shone a light into new informational voids, where we didn't have enough data, or it wasn't sufficiently high-quality, or we didn't send the right kind of instrument to answer the questions that now stood.

As decades passed, the number of space missions grew, as did the abilities of the instruments and the transmissional capacity of the spacecraft. While humans have not returned to the moon in the past fifty years, the probes we have sent to the Moon, to Mars, and to many other more distant bodies can now perform *in situ* spectroscopy (Baer, 2005), record high-resolution video with sound (Maki et al., 2020), and take our most cutting-edge technology with them with every launch. For certain high-interest research targets, such as Mars, we have so much data it becomes functionally impossible to study it all in detail — notably, the Mars Reconnaissance Orbiter had imaged over 99% of the planet's surface with at least a 6 m per pixel resolution by 2017, and its ultra high-resolution HiRISE camera had covered 2% of the surface at 25 cm per pixel (McEwen et al., 2018).

Our data analysis capabilities have grown in tandem with our data gathering, and much of the data analysis is now automated. Still, this analysis needs to be directed by specific theories or models, and as the complexity of the datasets grows, so does that of the tools needed to make sense of them. We see more and more applications of specialised computer science methods such as deep learning, notably in astronomy (Kremer et al., 2017) but also in planetary applications (Silburt et al., 2019). This makes interdisciplinary collaboration with computer science researchers a valuable asset for efficient use of our newly large and complex datasets. This thesis is an example of this sort of collaboration, presenting two different techniques that are little-used in planetary science, and suggesting the use of several others for the eventual continuation of this project.

In our first chapter, we start with an overview of remote sensing techniques and of their applications in planetary science in general, and the use of hyperspectral imaging to study planetary surface compositions in particular. We then introduce the New Horizons NASA mission to Pluto. A geologically fascinating object in its own right, Pluto nevertheless, due to its great distance from us, has until very recently been placed firmly in the data-poor category of planetary science research subjects. One of our chief source of data prior to this mission was imagery from the Hubble Space Telescope, which was not able to resolve Pluto spatially, and so any attempted mapping of the surface had to be done via indirect methods such as observation of its light curve over time. This changed with the New Horizons flyby in 2015, and the marvelously rich and detailed dataset from this encounter has still not been fully analysed. Notably, a global map of the surface composition from New Horizons' hyperspectral imager LEISA has not yet been produced, and the composition and analyses so far have been mostly carried out in a qualitative fashion, with the only quantitative inversion so far (Protopapa et al., 2017) making use of a simplified one-layer surface representation and so not taking into account the possibility of granular mixing or vertical layering.

Chapter II of this thesis is therefore devoted to describing our use of intensity-based registration, a method commonly used in medical imagery but not frequently applied in planetary science, to compose the first global qualitative maps of the main materials present on Pluto's surface. We then carry out a statistical analysis of these maps, correlating them with each other as well as other Pluto datasets, such as the surface reflectance, data from New Horizons' other spectral imager MVIC, and the surface elevation. We use these analyses to produce a set of classified surface maps, and to suggest geological interpretations of the surface material distribution. On a world with a geology as volatile as Pluto's, with significant parts of the surface sublimating and circulating seasonally (Bertrand et al., 2018, 2019), knowing the distribution of the surface composition is critical to studying both its history and its future evolution. Some of the global qualitative maps have already been used to constrain global atmospheric modeling of Pluto (Lewis et al., 2021), and the dataset promises to be a valuable asset for further descriptions and predictions of Pluto's climate and behaviour.

The next step after producing these global qualitative datasets is trying to quantify them in terms of the proportions of the different surface components and the different ways in which they can be mixed, including vertical stratification. In Chapter III we present a metaheuristic method for spectral inversion, which applies a simulated annealing algorithm to attempt to efficiently and accurately fit the surface spectra with quantitative surface representations. Using the DISORT radiative transfer model allows us to model surfaces with vertical stratification, the presence of which is very probable on the surface of an object with seasonal volatile transport. We first calibrate this algorithm on synthetically generated spectra. We then apply it to several regions of interest on the Pluto surface and show the first quantitative results for Pluto's surface composition.

In the final chapter of the thesis, we discuss the computational costs of producing a full global quantitative compositional map of Pluto. We establish these computational costs to be very high, and therefore outline a potential distributed computing project that would make completing the global compositional map feasible in a few years. We also suggest several algorithmic improvements to the inversion algorithm to make it more accurate and efficient, and discuss machine learning and graph-theoretical tools for reducing the high-dimensional problem of spectral inversion to one of more manageable dimensions.

General information

1 Spectroscopy in planetary science

Planetary science is a strongly interdisciplinary field. Many of its fundamental concerns and aims are drawn from earth science, merely broadened in scope beyond the Earth to other planets. The complication, however, of trying to study the properties of other planets, is that *in situ* studies of the sort that taught us much that we know about the Earth are frequently prohibitively difficult and expensive. Our range of techniques is largely restricted to remote sensing, i.e. getting information about the object of study without making physical contact. While a handful of extraterrestrial objects have been the targets of lander missions — notably our own Moon and Mars, but also Venus and several smaller bodies (asteroids and one moon, Titan) — these missions are limited in their in-depth study to a localised area, and serve as a complement to the more global remote-sensing data.

1.1 Remote sensing

Remote sensing, in the sense of spatially resolved observation of a planetary surface, is a relatively modern field. Relying as it does on accurate measurements recorded from a distance, its development was closely tied to that of both photography and flight. While initial forays into the field were made in the mid- to late-nineteenth century, via aerial photography from balloons, it was systematised several decades later with aircraft reconnaissance in WWII and, later, globalised with the development of artificial satellites. This new wealth of global satellite data, notably starting with the launch of Landsat 1 in 1972, incentivised the rapid development of digital analysis of remote sensing data as well as the creation of geographic information systems (GIS) to be able to represent this data in a common framework (Campbell and Wynne, 2011).

Spatially resolved mapping of extraterrestrial surfaces began in the 1960s with imagery of the Moon in preparation for the Apollo missions (Campbell, 2002), and has since been carried out notably in great detail for Mars, first with Mars Express and the Mars Global Surveyor, and subsequently with the Mars Reconnaissance Orbiter. Global mapping of other Solar System bodies has been less comprehensive owing to different complicating factors, such as its thick atmosphere for Venus, its proximity to the Sun for Mercury, and their great distance from the Earth for the outer Solar System planets and their moons.

Remote sensing can be separated into two categories: active, where the instrument sends a signal to its target and studies the signal's reflection, and passive, where the instrument detects reflections of natural signals such as sunlight. The planetary applications of active remote sensing are also limited in scope, as interplanetary distances mean the only way to target other bodies with a signal is sending an expensive space mission (i.e. MESSENGER's and Cassini's radar observation of Venus and Titan respectively). Most remote sensing in planetary science, therefore, is passive observation of how ambient electromagnetic radiation, notably sunlight, interacts with planetary bodies.

Although electromagnetic radiation forms a quantitatively continuous frequency spectrum, as shown in Figure I.1, it interacts with matter in drastically different ways depending on that frequency, which historically has led to referring to different ranges of the spectrum with distinct terminology (e.g. radio, x-rays, gamma rays...) Therefore, when we talk about light, we frequently mean the part of the electromagnetic spectrum that is visible to the human eye, i.e. wavelengths between ~400 and 700 nm. For convenience, however, one can also refer to all electromagnetic radiation as light,



and talk about light interacting with objects even when the radiation in question is in the infrared or ultraviolet range. We will use this terminology from here onwards.

Figure I.1: Schematic of the electromagnetic spectrum showing its different wavelength domains (Philip Ronan for Wikimedia).

1.2 Hyperspectral imaging

Visual remote sensing data is usually represented in the form of an image. This image can have a varying number of dimensions depending on its properties. A monochromatic camera measures light intensity for a single wavelength, or across a wavelength band, and returns a two-dimensional matrix. Each (x,y) location in this matrix (i.e. each pixel) has a corresponding light intensity. A typical RGB colour camera returns three monochromatic images, corresponding to the intensities measured in the red, green, and blue wavelength bands. This makes it a type of **multispectral imager**, which is a wider class of imagers that produce multiple monochromatic images for the same measurement, corresponding to discrete wavelength bands. A **hyperspectral imager**, on the other hand, is an instrument that returns a stack of monochromatic images that each correspond to a point in a continuous wavelength spectrum. This forms what is commonly referred to as a *hyperspectral data cube*, with two of its dimensions being spatial and one spectral.

A spectral imaging sensor can have several different modes of capture (shown in Figure I.2a):

- **spatial scanning** means it registers the entire spectrum at once for one spatial pixel or row, then moves onto the next;
- **spectral scanning** means the full field of view is captured at once for one wavelength, and then the sensor moves onto the next wavelength;
- **spatio-spectral scanning** means the sensor will scan both in the spatial and spectral dimensions;
- and in a non-scanning or staring mode all the data along the spatial and spectral axes will be registered simultaneously.



(a) Different modes of acquisition of a hyperspectral imager, with simultaneous or scanning capture in the spatial and/or spectral dimensions (CC BY-SA 3.0 Wikimedia Commons/Sascha Grusche).

- detector



(b) Comparison of push-broom and whisk-broom spatial scanning modes (modified from Campbell and Wynne, 2011).

Figure I.2: Acquisition and spatial scanning modes of spectral imagers.

Furthermore, a sensor operating in a spatial scanning mode can scan in multiple ways. Two common modes of scanning are **push-broom**, where a one-dimensional sensor array (or, more generally, one row of a two-dimensional sensor array) is situated perpendicular to the direction of movement, and is "pushed" along, capturing the data row by row; and **whisk-broom**, where the sensor array captures one pixel at a time, swinging back and forth across the movement track. These are shown in Figure I.2b. While a push-broom imager will register data much more rapidly, the advantage of a whisk-broom imager is that there is no need for inter-pixel calibration, as all the pixels will have been registered using the same parameters with the same monopixel detector.

The two characteristic elements of a hyperspectral imager are its **spatial** and **spectral resolutions**, which are referred to respectively in units of *pixels* and *spectels*.

Spatial resolution refers to the level of detail per unit of surface an imager is able to resolve. This cannot be given as an absolute value of surface distance, as it also depends on the distance between the imager and its target. Instead, the spatial resolution of an image sensor is expressed in terms of its angular resolution, as its *instantaneous field of view (IFOV)*. The IFOV is a solid angle, typically given in radians per pixel, and can be multiplied by the distance to the target to get the *ground-resolved instantaneous field of view (GIFOV)*, which corresponds to the spatial resolution of the image in meters per pixel¹. Figure 1.3 shows a representation of these values for a simplified case of a single stationary detector element in a nadir orientation.

¹Although a pixel has two dimensions, typically it is square and so has the same resolution along both spatial axes, and so only one distance value is given.



Figure I.3: Schematic of the spatial resolution parameters of a single detector element of an optical sensor.

Spectral resolution refers to an imager's ability to resolve variations in wavelength. A spectrometer registering a signal at a given wavelength is said to have a *spectel* at that wavelength. While a theoretical perfect spectrometer, when receiving a perfectly monochromatic signal of wavelength λ_n , would register it exactly, in reality a spectrometer will have a characteristic response function, typically Gaussian or triangular in shape, which peaks at λ_n . Consequently, the spectral resolution of a spectrometer is described in terms of a spectel's *full width at half maximum (FWHM)*, i.e. the difference between the two points of the response function where its value is half of its maximum intensity. This is accompanied by the spectral sampling interval, which specifies at which intervals the signal is *recorded*, i.e. the spacing between spectels. These different values are illustrated in Figure I.4.



Figure I.4: An example spectral response function for a spectel at wavelength λ_n , with the full width at half maximum (FWHM) and the spectral sampling interval indicated.

Any continuous series of measurements along a range of wavelengths is a spectrum, but different

types of spectra exist and are studied in astronomy, earth, and planetary science. Stellar astronomy commonly uses emission spectroscopy: stars are themselves strong primary sources of electromagnetic radiation, and the relative intensity of its emission as a function of frequency is characteristic of their composition. Planetary bodies, on the other hand, are much dimmer emitters of electromagnetic radiation, and emit at longer wavelengths (thermal infrared). This renders emissivity a poor indicator of planetary surface compositions, as the surface materials will also be strongly attenuating these wavelengths. We are therefore more interested in how planetary surfaces interact with light that's emitted by some other source, e.g. the Sun (or, for extrasolar planets, the star they are orbiting). This is the domain of absorption and reflectance spectroscopy, which we will detail further in the next section.

1.3 Absorption and reflectance spectroscopy

1.3.1 Light and matter interaction

When light interacts with an object, it can be reflected, absorbed, or diffused. Which of these occurs will depend on the wavelength of the light as well as the physical structure and properties of the object. If the object is granular, then the properties in question will be dependent on the size, arrangement, and composition of the grains, and as per Hapke (2012), most if not all of the surfaces we will be observing can be interpreted as a granular material:

[...] most of the surfaces you see consist of particulate materials. Sometimes the particles are loose, as in soils or clouds. Sometimes they are embedded in a transparent matrix, as in paint, which consists of white particles in a colored binder. Or they may be fused together, as in rocks, or tiles which consist of sintered ceramic powder. Even vegetation is a kind of particulate medium in which the "particles" are leaves and stems. These examples show that if we wish to quantitatively interpret the electromagnetic radiation that reaches us, rather than simply form an image from it, it is necessary to consider the scattering and propagation of light within nonuniform media.

To better understand the effect of light absorption and diffusion on the spectrum, we will briefly review the physical mechanisms behind it.

1.3.2 Mechanisms of light absorption

The absorption of light by matter can occur in several distinct ways, depending on the wavelength of the light and the energy of the interaction. The characteristic properties of types of absorption important to planetary science are summarised below.

Rotation. A molecule in a gas or liquid is able to rotate relatively freely, and, if it possesses an electric dipole moment, will change its orientation in an applied electric field. In a solid, a molecule's ability to rotate depends on its shape and its interactions with the environment. Frequently a molecule in a solid will have several stable orientations, and will flip from one to another during an interval called the *relaxation time*. The ratio of the period of the applied field to the relaxation time is described by its dielectric constant. If the ratio is large, and the molecules have time to rotate, the dielectric constant is large. Rotational spectroscopy is typically carried out in the microwave and far infrared wavelength ranges. An application of considerable interest in planetary remote sensing is observation of H₂O: liquid water has a high dielectric constant ($K_{er} = 80$ at 20°C) in the microwave range ($\lambda < 1$ cm), whereas that of water ice is much smaller ($K_{er} = 1.78$ at -20°C).

Molecular vibration. The bonds between the atoms of a molecule can stretch or bend with a characteristic frequency, and so irradiation with light of that frequency will induce the molecule to convert the light's energy into vibration. Typical vibrational frequences are found in the mid-infrared. A molecule can vibrate in multiple ways (vibrational modes): a linear molecule with N atoms will have 3N-5 vibrational degrees of freedom, whereas a nonlinear molecule will have 3N-6. Whether or not vibrational absorption is measurable with spectroscopy will depend on the structure of the molecule and the symmetry of its vibrational modes. For instance, the vibration of a simple symmetrical molecule with one bond such as N₂ will not be observed in the infrared spectrum. CO₂,



on the other hand, while also symmetrical, has a net dipole moment in two of its three vibrational modes, whereas H₂O has one in all three (shown on Figure I.5).

Figure I.5: The three fundamental vibrational modes of CO_2 (left) and H_2O (right). The symmetric stretching mode of CO_2 is IR-inactive, as there is no net dipole moment, but the asymmetric stretching and bending modes are IR-active, and all three vibrational modes of H_2O are IR-active.

Lattice vibration. Ions in a solid lattice are able to vibrate around their equilibrium positions. The two classes of vibration that can occur are those where the positive and negative ions move in sync, and no net dipole moment is produced. This type of vibration is called the *acoustical branch*, and is describes the propagation of shear and pressure waves. The second type of lattice vibration is the *optical branch*: here the positive and negative ions move out of phase, and the net dipole moment means that electromagnetic waves can be absorbed and emitted. The strongest absorption associated with lattice vibrations occurs in the thermal infrared.

Electronic transition. Most absorption in the near-ultraviolet, visible, and near-infrared ranges is facilitated by a transition where an electron is induced to jump from a lower- to a higher-energy state. In terms of the band model of electrons in a solid, the electron moves from the valence band to the conduction band, leaving behind a positively-charged *hole*. Different types of transitional absorption that can occur in solids are crystal-field (in transition elements), charge-transfer (when the electron jumps from one ion to another nearby ion), and excitonic (when a hole attracts a free electron and forms a system called an *exciton*).

1.3.3 The absorption band

An **absorption band** is a dip in intensity for a particular wavelength range which is a result of those wavelengths of light being absorbed by the material. Absorption bands can be of varying shape and depth, depending on the mechanism of absorption, and are typically quantified using the following characteristics:

- the position of the centre of the band along the wavelength axis;
- the depth of the band;
- the integral area of the band;
- the shape of the band.

An ideal absorption band may be described by a Lorentzian (L) or Gaussian (G) shape function, shown in Figure I.6 and mathematically represented as:

$$L = \frac{1}{1+x^2}$$
 (I.1) $G = e^{-(\ln 2)x^2}$ (I.2) with $x = \frac{l-l_0}{w/2}$,

where l_0 is the position at maximum, l is an arbitrary other position, and w is the full width at half maximum, with the unit of p_0 , p and w typically being a wavelength, wavenumber or frequency.

However, this ideal band shape is generally broadened by multiple local effects including Doppler shift and pressure. The final observed shape is a convolution of the profiles of all the implicated mechanisms.



Figure I.6: Comparison of Gaussian (red) and Lorentzian (blue) ideal band shapes.

1.3.4 Single-particle scattering



Figure I.7: Schematic showing the difference between the irradiance (flux per area) and radiance (flux per area per solid angle).

Before developing the more complex case of the diffusion of light by a granular surface, we will describe the geometrical quantities involved for the simpler case of a single particle. First, we must differentiate the two concepts of irradiance and radiance. For a given radiant flux Φ incident on an area dA, the **irradiance** J is the amount of flux received per unit area. The **radiance** or *luminance* I is the flux per unit area as well as per unit solid angle. Figure I.7 illustrates these concepts, and the equations below define J and I respectively.

$$J = \frac{\phi}{dA} \tag{I.3} \qquad I = \frac{\phi}{dA \times d\Omega} \tag{I.4}$$

Now we consider a spherical particle of radius *a* with a certain irradiance *J* incident upon it (shown in Figure I.8). We will refer to the total amount of flux intercepted (or extinguished) by the particle as Φ_E , and define the particle's *extinction cross-section*:

$$\sigma_E = \frac{\Phi_E}{J} \tag{I.5}$$



Figure I.8: Schematic of a spherical particle of radius *a*, with an incident irradiance *J* and scattered radiance *I*.

An amount of Φ_E is absorbed into the particle (Φ_A) and the remainder is scattered (Φ_S). We can therefore define the *absorption* and *scattering cross-sections*:

$$\sigma_A = \frac{\Phi_A}{\Phi_E} \tag{I.6} \qquad \sigma_S = \frac{\Phi_S}{\Phi_E} \tag{I.7}$$

Next, we define the *geometric cross-section* of the particle as $\sigma = \pi a^2$. With this, we can calculate the *extinction*, *absorption*, and *scattering efficiencies*:

$$Q_E = \frac{\sigma_E}{\sigma}$$
 (I.8) $Q_A = \frac{\sigma_A}{\sigma}$ (I.9) $Q_S = \frac{\sigma_S}{\sigma}$ (I.10)

Now we can finally define the useful value of the particle **single-scattering albedo** ω , which we will refer to frequently later, and which corresponds to the ratio of scattered flux to extinguished flux:

$$\omega = \frac{\Phi_S}{\Phi_E} = \frac{\sigma_S}{\sigma_E} = \frac{Q_S}{Q_E} \tag{I.11}$$

Additional useful definitions are the *scattering angle* θ , which is the angle between the direction of the incident flux and the scattering direction, and the *phase angle* $g = \pi - \theta$. Using these, we can define the particle's phase function *P*, which describes the angular pattern into which the incident flux is scattered. *P* can be described equivalently in terms of θ or *g*; with the latter, for a spherical particle *P* is defined as follows:

$$I(P) = J\omega \frac{P(g)}{4\pi} \tag{I.12}$$

If the scattering is isotropic (uniform in every direction), then P(g) = 1. If it is anisotropic, this can be characterised using various asymmetry factors, which will be touched upon further in section 1.4.4.

1.3.5 Bidirectional reflectance

When observing the interaction of light with a surface, we have interest in defining several simple geometric terms to accurately describe what we see. Figure I.9 illustrates the following:

incidence angle *i*: the angle between the normal to the surface and the direction of illumination; **emergence angle** *e*: the angle between the normal and the direction we are observing from; **phase angle** *g*: the angle between the directions of illumination and observation;

azimuthal angle φ : the angle between the surface projections of the directions of illumination and observation.

Also commonly used are the notations $\mu_0 = \cos i$ and $\mu = \cos e$.

As defined previously, the quantity of flux incident per unit of surface is the *irradiance J*, and the flux emitted (and observed) is the *radiance I*. With these two values, we can compute the **reflectance** R(i, e, g) of the surface. We can also represent it in terms of the radiance factor I/F where $F = \pi J$, which is a notation commonly used in astrophysics and planetary science. Finally, we can define the **reflectance factor** REFF by dividing the radiance factor by the cosine of the incidence angle μ_0 .

$$R(i, e, g) = \frac{I}{J} = \frac{I}{\pi F} \quad (1.13) \qquad \frac{I}{F} = \pi \times R(i, e, g) \quad (1.14) \quad REFF = \frac{\pi \times R(i, e, g)}{\mu_0} \quad (1.15)$$

Figure I.9: Schematic showing the photometric terms used to describe the reflectance of a surface.

1.3.6 The reflectance spectrum

The behaviour of a particle or a surface with respect to incident light will vary as a function of that light's wavelength, as per the mechanisms described in 1.3.2. We will thus consider the **reflectance spectrum**, which is a series of reflectances corresponding to the series of wavelengths at which the reflectances was measured. This spectrum consists of an alternating sequence of two types of structure: the *absorption bands* and the *continuum*.

The **continuum** is constituted of the parts of the spectrum between its absorption bands, i.e. the overall curve that the bands are "subtracted" from, and is characterised by the surface's albedo.



Figure I.10: Schematic representation of the changes applied to a radiance *I* traversing a volume *dAdz* due to extinction, scattering and emission.

The relationship between the reflectance spectrum of a surface and its physical properties is given by the **radiative transfer equation (RTE)**. To summarise it in words, the RTE states that as a beam of radiation travels, it loses energy to absorption, gains energy by emission processes, and redistributes energy by scattering. Equation I.16 expresses the general form of the RTE for a radiance *I* traversing a cylindrical volume dAdz in the direction **r** (see schema in Figure I.10). Its principal components are (1) extinction (by absorption and scattering), (2) the sum of radiances incident in other directions **r**' which are scattered into **r**, and (3) thermal emission (multiple scattering and fluorescent emission are not considered).

$$\frac{\partial I(s,\mathbf{r})}{\partial z} = \underbrace{-E(s,\mathbf{r})I(z,\mathbf{r})}_{\Delta I_E \text{ (extinction)}} + \underbrace{\frac{S(z,\mathbf{r})}{4\pi} \int_{4\pi} I(z,\mathbf{r}')P(z,\mathbf{r}',\mathbf{r})d\mathbf{r}'}_{\Delta I_S \text{ (single scattering)}} + \underbrace{F(z,\mathbf{r})}_{\Delta I_F \text{ (thermal emission)}}$$
(I.16)

Frequently, rather than using a physical depth of propagation z, we will define an **optical depth** τ (Equation I.17), a dimensionless distance which corresponds to, for a radiance propagating over a distance z through the medium, a reduction in intensity of $e^{-\tau}$. Substituting z with τ , as well as E and S with the volume single scattering albedo w = S(z)/E(z), we can express the RTE for a horizontally stratified medium in terms of τ , w, and the phase function $P(\tau, g)$ (where g is the phase angle between \mathbf{r} and \mathbf{r}'), as shown in Equation I.18.

$$d\tau = -E(z)dz = -\frac{E(s)ds}{\cos\theta}$$
 where $E = N\sigma Q_E$ (1.17)

$$-\cos\theta \frac{\partial I(\tau,\mathbf{r})}{\partial \tau} = -I(\tau,\mathbf{r}) + \frac{w(\tau)}{4\pi} \int_{4\pi} I(\tau,\mathbf{r}')P(\tau,g')d\mathbf{r}' + J\frac{w(\tau)}{4\pi}P(\tau,g)e^{-\tau/\cos i}$$
(1.18)

Finding an exact analytical solution to the radiative transfer equation has proven intractable, and approximations or numerical methods must be used to solve it. These will be described in the next section.

1.4 Radiative transfer models

Radiative transfer models (RTMs) are theoretical models for solving the RTE, consisting of a set of simplificatory assumptions and/or empirical measurements. There are numerous RTMs, suitable for different cases and applications. We will describe a few of the notable ones here.

1.4.1 Lambert

One of the simplest RTMs is Lambertian reflectance, which simply defines the reflectance as a function of the angle of incidence (Equation I.19). This model is more accurate for bright reflective surfaces, such as snow, and less reliable for dark materials such as soil or vegetation.

$$R = \frac{1}{\pi} A_L \mu_0 \text{ (where } A_L \text{ is the surface albedo)}$$
(I.19)

1.4.2 Hapke

The Hapke model (Equation I.20) expresses bidirectional reflectance via a semi-empirical expression derived from Chandrasekhar's equation of radiative transfer. It takes into account both the opposition effect² in the function B(g) and multiple scattering in the Chandrasekhar function $H(\mu, \omega)$. Its range of application, however, is limited due to the fact that it uses geometrical optics, a set of simplifications that assumes the wavelength is much shorter than the characteristic length of the objects in question. For infrared spectroscopy, the relevant wavelengths are in the 1–5 μ m range,

²The opposition effect is a phenomenon where the reflectance of a surface sharply increases for a phase angle close to 0. It happens due to two physical causes: shadow hiding and coherent backscattter.

which means the size range for accurately modelisable particles is on the order of one to a few tens of microns.

$$R(i, e, g) = \frac{\omega}{4} \frac{1}{\mu_0 + \mu} ((1 + B(g))P(g) + H(\mu_0, \omega)H(\mu, \omega) - 1)$$
(I.20)

1.4.3 Numerical solvers and approximations

A very simple reflectance model such as the Lambertian one can be solved analytically. More complex RTMs, especially those which feature anisotropic scattering, use numerical methods for resolving the RTE. Notable among these are the following:

Monte Carlo: a heuristic method which involves simulating a large number of photons and deriving solutions based on a statistical sampling of their behaviour. This is a powerful algorithm but its accuracy is a direct consequence of its computational intensity.

Two-stream approximation: only radiation propagating along two discrete directions is considered. A notable variant is the Eddington approximation. It is frequently used to calculate radiance in a *plane-parallel* environment, where properties are assumed to only vary in the perpendicular direction.

Discrete ordinates: a method where both the spatial dimensions and the angular variables are discretised. This transforms the integrodifferential RTE into a linear system of ordinary differential equations, and can be thought of as a generalisation of the two-stream approximation to the N-stream case. The system can subsequently be solved using any of the typical iterative methods for solving differential equation systems, e.g. least squares or eigenvalue solvers. One RTM using the discrete ordinate approximation is DISORT, which will be discussed in more detail in Chapter III.

1.4.4 The phase function

While Lambertian reflectance is isotropic, Hapke and other more complex RTMs make use of a **phase function** P(g). As we briefly touched on previously, the phase function describes the anisotropy of a particle or surface, i.e. the angular spread of how it scatters light. A widely used and versatile phase function is the Henyey-Greenstein (HG) function, shown in Equation I.21 in its one-lobed form, with the cosine anisotropy parameter ξ as its main parameter. The two-lobed HG function models backscattering and forward-scattering behaviour separately and gives more accurate results for irregular particles, but is more computationally intensive to use.

$$P_{HG1}(g) = \frac{1 - \xi^2}{(1 + 2\xi \cos g + \xi^2)^{3/2}}$$
(I.21)

1.4.5 Surface representation

Planetary surfaces are complex, with different chemical components mixed and layered in a number of ways. The different kinds of component mixture return different reflectance spectra, which needs to be taken into account in the RTM in order to accurately simulate a surface's reflectance. The RTM parameters affected by this are the single-scattering albedo ω and the optical depth τ_{τ} . There are four basic types of surface component mixture, shown in Figure I.II, and described below.

Spatial (subpixel) mixture: the components are spatially segregated on a subpixel scale and so cannot be resolved into individual pixels. The reflectance spectrum of this pixel is a linear sum of the spectra of the separate components, weighted by their spatial proportion (see Figure I.12a). Note that these spatial components can themselves be mixtures of another kind, granular, molecular, or vertical.

Granular mixture: the grains of the components are mixed together within a layer. The resulting reflectance spectrum is a nonlinear mix of the component spectra, which can be seen in Figure I.12b.

Molecular mixture: the grains of the surface are themselves composed of different molecular compounds, which can take various forms. Some examples are hydrated clays, clathrate hydrates, or solid solutions. The reflectance spectrum of these mixtures is complex and not always easily relateable to the spectra of the pure components, as the intermolecular bonds have their own vibrational frequency and intensity, and spectral bands can be attenuated or displaced, or entirely new spectral bands can appear.

Stratification: the components overlie each other in horizontal layers, and can themselves be mixtures of different types. As Figure I.12c shows, the evolution of the resulting spectrum with the layer thickness is also nonlinear. When the top layer is *optically thin*, sufficient photons will penetrate it and be diffused by the underlying layers that the spectrum will also show their characteristic reflectance bands. Once the layer is sufficiently thick (i.e. *optically thick*), the underlying layer or layers will no longer have an identifiable effect on the spectrum, and so the uppermost layer can be represented as semi-infinite, and the substrata neglected in the model. Stratification is an important mixture type to consider, as it can be used to represent common real-world scenarios which we may expect to see in a chemically complex environment; notably, differential sublimation of a pure component forming a microns-thin layer atop a granular or molecular mixture.



Figure I.11: Schematic diagrams of (A) granular, (B) vertically stratified, and (C) spatially segregated component mixtures (modified from Philippe, 2016). The intramolecular mixture type is not shown, but can occur within any grains of these three types.

In addition to the varying nonlinear effects on the spectrum of these different mixtures, Figure I.12d shows how varying the grain size of pure CO_2 ice affects the shape of its spectrum. Incorporating the effects of grain size, anisotropy, and different mixture types into the radiative transfer model results in a complex surface representation with a large number of free parameters, resolving which numerically is not trivial. Resolution methods will be discussed further in Chapter III.



(a) Evolution of the spectrum of a spatial mixture of CH_4 and H_2O ice, varying from 0 to 100% of H_2O by steps of 5%.



(c) Evolution of the spectrum of an optically thin layer of CH₄ ice (grain size $d = 100 \,\mu$ m) overlying a substrate of H₂O ice ($d = 100 \,\mu$ m), with the thickness of the CH₄ ice layer varying from 0 to 5 mm.



(b) Evolution of the spectrum of a granular mixture of CH_4 and H_2O ice, varying from 0 to 100% of H_2O by steps of 5%.



(d) Evolution of the spectrum of CH_4 ice with increasing grain size.

Figure I.12: Comparison plots showing the effect of different surface mixture types on the evolution of spectra.

1.4.6 Optical constants

In order to represent a specific material in a radiative transfer model, we need its optical properties. There are multiple ways to represent these, but by convention the commonly-used values are the real and imaginary parts of the **complex refractive index**: $\bar{n} = n + ik$. The real part *n* is the *refractive index*, which represents the ratio between the speed of light in the medium and in vacuum, and the imaginary part *k* is the *extinction coefficient*, which represents the attenuation of light as it travels through the medium. *n* and *k* together are termed the **optical constants** of a material, and many databases of optical constants are maintained, such as SSHADE (Solid Spectroscopy Hosting Architecture of Databases and Expertise, see Schmitt et al. (2017b) for documentation). These are not measured directly, but can be derived from experimental values such as transmission spectra, specular reflection, reflectance, and diffusion by aerosol particles. Care must be taken to select the right dataset for the model conditions: optical constants vary with wavelength, but also with solid state (crystalline or amorphous), mixing state, temperature, thermal history, isotopic abundance, crystal system orientation, etc. (Schmitt et al., 1998).

2 New Horizons mission to Pluto

The New Horizons mission is a NASA mission to the outer Solar System, the first flyby target of which was the Pluto system. Its July 2015 flyby returned a wealth of data, in particular greatly advancing our knowledge of Pluto's surface topography, geology, and composition. Much of this data has been processed and analysed, and different maps have been created that show the variation of Pluto's surface properties. However, plenty of work remains to be done — notably, there is no global map of the surface composition, or any quantitative analysis of such.

2.1 Introduction to the Pluto system

2.1.1 General information

Pluto is the ninth largest known object directly orbiting the Sun: its mean radius is 1188.3 km (Nimmo et al., 2017) and its mass is 1.303×10^{22} kg (Stern et al., 2015), or about 0.00218 Earth masses. Its eccentric orbit varies between 30 and 49 AU in distance from the Sun, which means it periodically comes closer to the Sun than Neptune. The large inclination of its orbit (> 17°) in respect to the ecliptic plane of the Solar System sets it apart from the planets. Pluto consists largely of water ice and silicates, and is expected to be partially differentiated into a silicate core and ice crust (see Figure 1.13).



Figure I.13: Pluto and Charon bulk properties and proposed internal structure (modified from McKinnon et al., 2017). PF2 is a hydrated carbonaceous rock mineralogy model based on Solar System abundances revised from the original PF model of Mueller and McKinnon (1988).

Pluto has five known moons. The largest, Charon, has a mean radius of 606 km and a mass of $1.586 \times 10^{\circ}$ kg (12.2% of Pluto's mass). This is sufficient for Charon to have hydrostatic equilibrium, and also causes the Pluto-Charon barycenter to be outside Pluto, making them a binary system. Pluto and Charon are mutually tidally locked. The smaller moons — Nix, Hydra, Kerberos, and Styx — are much smaller than Charon, and are not large enough to be spherical (see Table I.1). All of them were first observed via the Hubble Space Telescope, but while Nix and Hydra were first seen in 2005, shortly before the launch of New Horizons, Kerberos and Styx were only discovered in 2011 and 2012 respectively (Weaver et al., 2016).

| Name | Size (km) | Orbital semimajor axis (10 ³ km) | Orbital period (days) |
|----------|-------------------------|--|------------------------------|
| Styx | $16 \times 9 \times 8$ | 42.6 | 20.2 |
| Nix | 50 × 35 × 33 | 48.7 | 24.9 |
| Kerberos | $19 \times 10 \times 9$ | 57.8 | 32.2 |
| Hydra | 65 × 45 × 25 | 64.7 | 38.2 |

Table I.1: Size and orbital properties of Pluto's small satellites as per Weaver et al. (2016). Note that the orbital periods are close to integer multiples of Charon's 6.4-day orbital period, the ratios being 3:4:5:6 (in the order presented in the table).

While Pluto's crust is composed of water ice, it is largely covered by a relatively thin layer of more volatile ices. Specifically, Pluto's surface is dominated by nitrogen ice, with methane ice, carbon monoxide ice, and dark red organic material as minority components. The composition varies strongly with geography, and is described in more detail in 2.4.3.



DISCOVERY OF THE PLANET PLUTO

Figure I.14: Photographic plates from Clyde Tombaugh's discovery of Pluto (image from Lowell Observatory Archives).

2.1.2 What we knew before New Horizons

The planet Neptune, too far away from Earth to be observed by the unaided eye, was discovered by 1846 following observation of perturbations in the orbit of Uranus by Alexis Bouvard. Bouvard's calculations allowed Neptune's position to be predicted independently by two different astronomers, and it was shortly thereafter observed with a telescope by Johann Galle. Following this, speculation arose that apparent additional perturbation in Uranus's orbit was due to another yet-undiscovered Solar System planet, "Planet X", located beyond Neptune. Percival Lowell financed the initial search for Planet X, which began in 1906 but was not successful until 14 years after his death. In January 1930, Pluto was first observed on images from the Lowell Observatory by Clyde Tombaugh. Tombaugh used a blink comparator to rapidly shift between two photography plates in order to determine if any objects had moved in between the two captures, and discovered a possible moving object with an apparent magnitude of +15.1 in plates taken on January 23 and 29 (shown in Figure I.14).

While Pluto's discovery seemingly corresponded to Lowell's prediction of a trans-Neptunian Planet X, it was quickly determined to be too small to explain the apparent orbital irregularities. After the initial estimation of Pluto's mass as one Earth mass in 1931, it trended downward in subsequent calculations, first via its albedo and then via its relationship with Charon, finally arriving at approximately 0.2% of Earth's mass (Christy and Harrington, 1978). Further searches for Planet X were unsuccessful, and new information from Voyager 2's flyby of Neptune allowed for a revision of Neptune's mass that eliminated the orbital discrepancies and the need for Planet X's existence altogether.

2.1.3 Is it a planet?

Pluto was considered the ninth planet of the Solar System before its "demotion" in 2006 by the International Astronomical Union to the status of dwarf planet. The IAU definition of a celestial body as a "planet" has the following three criteria:

- It is in orbit around the Sun.
- It has sufficient mass to assume a hydrostatic equilibrium (nearly round) shape.
- It has cleared the neighbourhood of its orbit.

Pluto meets the first two of these criteria, but not the third, as it shares its orbital neighbourhood with many other Kuiper belt objects (KBOs), referred to as plutinos. The discovery of several large Solar System bodies such as the trans-Neptunian objects (TNOs) Eris³, Sedna, and Quaoar in the early 2000s, all with sizes and orbits comparable to Plutos, spurred a debate (Basri and Brown, 2006)

³Eris was referred to as the "tenth planet" by NASA at the time of its discovery in 2005 (NASA/JPL, 2005).

about the need for formalising the definition of "planet" in order to determine whether all of these new objects were to be called planets, or whether Pluto needed to be reclassified.

The IAU formed a committee in 2005 to draft this formal definition. The original definition proposal contained only the first two points of the later-adopted definition (IAU, 2006b), which would mean at least three Solar System bodies in addition to the classical nine would be considered planets: the asteroid Ceres, the TNO Eris, and Pluto's moon Charon (this would make Pluto and Charon a double planet). Many further KBOs would also potentially fall under this definition. This draft proposal was contentious within the IAU (Britt, 2006) as well as criticised by popular astronomy writers (Plait, 2006; Matzke, 2006). After several revisions and debates, the final definition of "planet", as stated above, was passed in August 2006 at the 26th General Assembly of the IAU (IAU, 2006a). A dwarf planet was defined as "a celestial body that (a) is in orbit around the Sun, (b) has sufficient mass for its self-gravity to overcome rigid body forces so that it assumes a hydrostatic equilibrium (nearly round) shape [2], (c) has not cleared the neighbourhood around its orbit, and (d) is not a satellite." This category therefore contains Pluto, Ceres, and the large TNOs mentioned above.

2.1.4 Surface geography and chemistry





(a) Evolution of Pluto's visible lightcurve over one decade, based on Hubble Space Telescope observations obtained during 1992–1993 and during 2002–2003 (reproduced from Grundy et al., 2013).

(b) Variation of Pluto's visible lightcurve and integrated band depths and spectral indices with longitude (reproduced from Grundy et al., 2013).

Figure I.15: Comparison of short- and long-term variations visible in Earth-based observations of Pluto.

Pluto's surface cannot be optically resolved with Earth-based hyperspectral instruments, but studying the variation of its lightcurve with time lets us establish the heterogeneity of the surface's composition (Douté et al., 1999; Grundy and Buie, 2001; Grundy et al., 2013). Short-term observations on the scale of a Pluto day (6.4 Earth days) show the longitudinal distribution of terrains, as the sub-Earth point, i.e. the point on Pluto's surface closest to the Earth, moves around Pluto. Longer-term seasonal-scale observations allow us to observe latitudinal variation, as the sub-Earth point migrates from north to south due to Pluto's axial tilt, as well as season-driven evolution of the surface itself. Figure I.15 shows the variation of the lightcurve and several spectral features with longitude as well as the change in the lightcurve in the span of a decade. Figure I.16 shows the effect of Pluto's strong axial tilt of 120° on its seasons.



Figure I.16: Illustration of Pluto's seasons, showing how its axial tilt leads to the formation of distinct northern and southern seasons (CC BY-NC-ND 3.0 Emily Lakdawalla).

Separating latitudinal variation from seasonal surface evolution due to transport of volatiles in this spatially unresolved data is nontrivial, but combining the light curve data with stellar occultation data, which gives us information on the atmosphere, allowed to establish climate evolution models which help with this differentiation (Young, 2013). Coupling these with low spatial-resolution images from the Hubble Space Telescope (Stern et al., 1997; Buie et al., 2010) allowed for the creation of a global compositional map consisting of ~40–55 resolution elements, as shown in Figure I.17. The main terrain types established, using geometric albedo as a threshold, are bright N₂-rich terrains (mostly equatorial), medium-brightness CH_4 -rich terrains (mostly polar), and dark organic material-rich terrains (also mostly equatorial).



Figure I.17: Compositional models of Pluto's surface, based on Hubble Space Telescope data segmented into three albedo-based levels (reproduced from Grundy and Buie, 2001). White areas are N_2 -rich terrains, gray areas are CH_4 -rich terrains, and black areas are organic material-rich terrains, except in the case of the Buie et al. HST map, where shade corresponds to albedo.

2.2 The New Horizons probe

New Horizons is a NASA mission and spacecraft that was launched in January 2006 with the primary goal of studying the Pluto system. It achieved this in a flyby on July 14th, 2015, and then went on to perform a flyby study of 486958 Arrokoth, a much smaller KBO, on January 1st of 2019. New Horizons is the fifth space probe to reach Solar System escape velocity, after the Pioneer and Voyager probes. As it has sufficient power to maintain its instruments operational until the 2030s, the search for further potential study candidates in the Kuiper belt is ongoing.

2.2.1 Mission goals

The scientific goals of the Pluto system flyby were separated into three groups, summarised in Table I.2 as per Stern (2008). The goals termed "required" were the broad global goals of characterising Pluto and Charon's geology and surface composition as well as Pluto's atmosphere. The "highly desired" group mostly features higher-resolution imagery, the characterisation of Pluto's evolution over time, and additional datasets such as surface temperatures and atmospheric chemistry. In the least-priority "desired" group are the refining of data on bulk parameters and orbits, which would be easily obtained with flyby data assuming all went according to plan, and the search for any additional satellites, which of course depended on whether there would be any to find.

The flyby was completely successful, meeting or exceeding all the objectives set out for it (Stern et al., 2018). It was the first mission to be able to map Pluto's surface in detail, and the discovery of active geological features, consisting of volatile ices such as CH_4 and N_2 , allowed us to establish the presence of a seasonal volatile cycle (Grundy et al., 2016; Schmitt et al., 2017a). The details of these discoveries will be further discussed in section 2.4.

Group 1: Required

| Characterise the global geology and morphology of Pluto and Charon |
|---|
| Map the surface composition of Pluto and Charon |
| Characterise the neutral atmosphere of Pluto and its escape rate |
| Group 2: Highly desired |
| Characterize the time variability of Pluto's surface and atmosphere |
| Image Pluto and Charon in stereo |
| Map the terminators of Pluto and Charon with high resolution |
| Map the surface composition of selected areas of Pluto and Charon at high resolution |
| Characterize Pluto's ionosphere and solar wind interaction |
| Search for neutral atmospheric species including H, H ₂ , HCN, and C_xH_y , and other hydrocarbons |
| and nitriles in upper atmosphere |

Search for an atmosphere around Charon

Determine bolometric Bond albedos for Pluto and Charon

Map the surface temperatures of Pluto and Charon

Group 3: Desirable

Characterize the energetic particle environment of Pluto and Charon Refine bulk parameters (radii, masses, densities) and orbits of Pluto and Charon Search for additional satellites and rings

Table I.2: New Horizons science objectives as stated in the January 2001 Announcement of Opportunity for the Pluto-Kuiper Belt Mission (PKB-AO: AO:01-OSS-01).

2.2.2 Instrument payload



Figure I.18: Schematic of the New Horizons spacecraft showing science instrument locations and orientations (reproduced from Flanigan et al., 2016).

New Horizons carries seven scientific instruments: three optical sensors, two plasma sensors, a radiometer, and a dust sensor. Figure I.18 shows their locations on the spacecraft, and a brief description of each instrument as per Weaver et al. (2008) follows.

LORRI (Long-Range Reconnaissance Imager) is a high-resolution monochromatic visible-light imager. It operates between 350 and 950 nm with an angular resolution of 5 μ rad (~1 arcsec). LORRI is the highest-resolution instrument and gives us the most surface detail (best resolution ~ 80 m/px), but cannot register wavelength variations (e.g. colour) that might suggest chemical composition.

Alice is an ultraviolet imaging spectrometer with a spectral range of 46.5–180 nm and a spectral resolution of 1024 bands within that range. The goal of ultraviolet measurements is to study the composition of Pluto's upper atmosphere, establishing vertical profiles of temperature, pressure, and molecular abundances via solar occultation.

Ralph is a two-component visible and infrared spectral imager which provides us with the bulk of the data on Pluto's surface composition, and is described in more detail in 2.2.3.

SWAP (Solar Wind Around Pluto) is a solar wind and plasma spectrometer, registering particles in the energy range of 0.25–0.75 keV. SWAP's purpose is to estimate Pluto's atmospheric escape rate and to study the solar wind interactions of Pluto and Charon.

PEPSSI (Pluto Energetic Particle Spectrometer Science Investigation) overlaps in functionality with SWAP, its measurements of pickup ions⁴ from Pluto's atmosphere also helping estimate the atmospheric escape rate, but it operates in a higher energy range (1 keV–1 MeV) and its primary purpose is to characterise the energetic particle environment at Pluto.

REX (Radio EXperiment) is a radiometer. Its prime purpose is to characterise the pressure and temperature of Pluto's lower atmosphere, which it accomplished via determining the delay of radio signals sent from Earth during occultations.

The **SDC** (Student Dust Counter) is an impact dust detector designed to trace the density of dust in the Solar System along the mission trajectory. It was not a part of the initial New Horizons proposal, but was added by NASA as an education and public outreach experiment. It is the first interplanetary mission instrument designed, built, and operated by students.

None of the instruments have independent scanning platforms, so the entire spacecraft must be reoriented in order to record data. The necessary orientation and mode of motion varies per instrument, and is shown in Figure I.19. The particle instruments, namely SWAP, PEPPSI and SDC, can be performed in any orientation and movement mode, although SDC operates largely during the cruise phase, as thruster firing during active trajectory correction adds a large amount of noise to its data. REX and Alice occultation observations are performed by orienting the radio antenna towards the Sun and Earth across Pluto. LORRI, MVIC framing array, and Alice airglow observations are

⁴Pickup ions are formed when neutral particles within the heliosphere are ionised by solar ultraviolet radiation or charge exchange with solar wind protons (Kallenbach, 2000).

Scanning Images ~ 1 minute **Occultation ~ 40 minutes** (Ralph) (Alice, REX) Scan Sun Direction Earth Pluto Spacecraft Rotation about Z-Axis Staring Images ~ 100 milliseconds (LORRI, Ralph, Alice) **Particle Measurements Anytime** (SWAP, PEPSSI) **Dust Measurements During Cruise** (SDC)

made in staring mode, while MVIC TDI and LEISA measurements are made in pushbroom scanning mode. The FOVs of the visual instruments are compared in Figure I.20.

Figure I.19: The different observation modes of New Horizons instruments (reproduced from Weaver et al., 2008).



Figure I.20: The fields of view (FOVs) of the MVIC, Ralph, Alice airglow, and LORRI instruments, projected onto the sky plane (reproduced from Weaver et al., 2008).

2.2.3 Ralph spectrometer

Ralph is a single telescope feeding data to two instruments: the Multispectral Visible Imaging Camera (MVIC) and the Linear Etalon Imaging Spectral Array (LEISA). A beamsplitter allows IR wavelengths longer than 1.1 μ m to continue to LEISA, while shorter ones are reflected to MVIC. MVIC provides overlapping coverage of the 400–975 nm wavelength range and obtains a maximum spatial resolution of 0.66 km/pixel for the encounter hemisphere as imaged at closest approach. LEISA, although its pixel field-of-view is three times wider than that of MVIC, provides a complete spectrum in the compositionally important 1.25–2.5 μ m range.



Figure I.21: Illustration of time delay integration (TDI) using a four-row CCD (reproduced from McEwen et al., 2007).

MVIC

MVIC is composed of seven independent CCD arrays. One of them is a 5024×128 panchromatic frame transfer array which is used primarily for optical navigation and operates in staring mode. The other six — two panchromatic and four colour channels — have a resolution of 5024 by 32 pixels (with each pixel's IFOV being $20 \times 20 \,\mu$ rad) and operate in time delay integration (TDI) pushbroom scanning mode. Time delay integration is a technique for imaging rapidly moving objects with a high signal-to-noise ratio. It works by transferring the recorded data across the rows of the CCD at the same speed as the imaged object moves across the FOV of the detector. This results in the signal strength being multiplied by the number of rows (see Figure I.21 for an example using four rows and a point source).

The four MVIC colour bands can be distinguished into three broad-band channels — blue (400– 550 nm), red (540–700 nm) and near-IR (780–975 nm) — and one narrow-band (860–910 nm) channel which serves as a CH₄ filter, allowing us to map the abundance of CH₄ on the surface via its absorption band at 0.89 microns. The spectral ranges of these different colour filters are shown in Figure I.22.



Figure I.22: Spectrum of Pluto with the MVIC colour filters for comparison (reproduced from Young et al., 2008).



Figure I.23: Measured resolving power ($\lambda/\Delta\lambda$) of the LEISA filter assembly for the lower-resolution (1.25–2.5 μ m) and higher-resolution (2.1–2.25 μ m) segments (reproduced from Reuter et al., 2008).

LEISA

LEISA is a hyperspectral imager that receives data through a near-IR wedged filter onto a 256 × 256 detector array with a single-pixel IFOV of $62 \times 62 \mu$ rad. A hyperspectral cube is obtained by pushbroom scanning similarly to MVIC TDI imaging, with the frame readout rate synchronised to the rate of the scan. However, where for MVIC TDI scanning the signal is amplified in the same spectral band, for LEISA scanning the motion of the image builds up a spectrum: for every time the array is moved across the target by the width of a pixel, a frame corresponding to a wavelength is read out.

The near-IR filter has two segments: the first allows us to observe the $1.25-2.5 \mu m$ spectral range with an average resolving power of 240, while the second one covers the narrower $2.10-2.25 \mu m$ range with an average resolving power of 560, due to the compositional importance of the N₂ spectral feature at $2.15 \mu m$. Figure I.23 shows the slight but present variation of the actual measured resolving power with wavelength for the two filter segments.

2.3 2015 Pluto flyby

2.3.1 Approach phase

The Pluto system flyby event consisted of a long approach phase (split into three sub-phases) followed by a short closest-encounter phase, where the highest-resolution datasets were taken. This, in turn, was followed by a long tripartite departure phase. Figure I.24 shows an overview of these phases and which science tasks were to be accomplished during which phase. In January of 2015 the spacecraft switched from its low-fuel-consumption spin-stabilised flight mode to the more precise 3-axis pointing mode, as it needed to accurately target Pluto with the LORRI camera to take optical navigation (OpNav) images and compute trajectory corrections. This was the official beginning of the approach phase, but only in February did LORRI's optical resolution begin to exceed that of the Hubble Space Telescope, allowing for higher-resolution data of Pluto to be obtained than we already had with Earth-based observations.

As New Horizons got closer to the Pluto system, increasing image resolution allowed for calculating the system's ephemerides with correspondingly higher accuracy. With this new data, trajection correction maneuvers (TCM) were carried out 126, 30, and 14 days before the closest encounter. It was important to minimise the number of TCMs, as every trajectory modification carries an inherent risk of mistake, and so only three of the identified seven TCM opportunities were used. Figure I.25 shows this, as well as the alternating OpNav and spin maneuver phases of the approach timeline. The goal of the final OpNav phase, after the final TCM, was to improve the knowledge of the time of the closest approach, so as to better time the science measurements. The Core Command sequence, which began 7 days before the closest approach and ran for 9 days, was designed to let the spacecraft operate autonomously and record the encounter data without intervention from Earth, as the light-speed delay was too great for human guidance to be useful (Flanigan et al., 2016).



Figure I.24: An overview of the New Horizons mission science timeline for 2015, prior to and following its close encounter with Pluto in July (modified from NASA/JHU/APL).



Figure I.25: Pluto approach schedule with spacecraft mode, OpNav campaigns, TCM opportunities labeled, and TCM opportunities used outlined in green. All dates are in the year 2015 and given in month/day format. Modified from Flanigan et al. (2016).

While the images captured with LORRI were of highest utility to the navigation corrections, other instruments were also operating and recording data during the approach, as is evident from the list of science tasks in Figure I.24. 150 days before the encounter, LORRI was able to resolve the disk of Pluto, while MVIC and LEISA had sufficient resolution to separate Pluto and Charon. This point marked the beginning of a 6.5-day observation sequence, during Pluto and Charon were observed every 12 hours. This covered a complete Pluto rotation period at 30-degree latitudinal spacing. In total there were eight such observational sequences during the approach, starting at 150, 100, 66, 44, 28, 19, 12, and 6 days before closest encounter (Young et al., 2008).

2.3.2 Encounter phase

At the time of the near-encounter flyby, New Horizons was moving at a speed of 13.78 km/s relative to Pluto. The relatively short duration of the near encounter combined with the need to rotate the spacecraft in order to point the different instruments at Pluto (see Figure I.18 for the relative orientations of the instruments) meant that observation opportunities per instrument were limited and scheduled to alternate during the approximately 16 hours of the near-encounter sequence. The higher-resolution MVIC and LORRI instruments were privileged for much of this duration, but LEISA datasets were obtained during two intervals of the near-approach, which formed the "global" encounter hemisphere map, with spatial resolutions of 10 km/px and 5-7 km/px for the first and second datasets respectively. A narrow swath of even higher-resolution LEISA data was obtained during the height of the near encounter, with coverage of about 11% of the surface at 2.7 km/px.

2.4 Data from the encounter

2.4.1 Initial results

While the full volume of data from the Pluto system encounter would take 15 months to download, first results were compiled and published in October of 2015 (Stern et al., 2015). This included a global panchromatic visible light mosaic, with a number of prominent geological features of interest highlighted and given informal names (many of these have by now been formalised). Notable among these is the bright heart-shaped region of Tombaugh Regio, which consists of Sputnik Planitia, a smooth plain of volatile N_2 ice, in its western half, and more rugged uplands in its east. Also of great interest is the dark red equatorial belt of Cthulhu Macula, which encircles Pluto almost entirely, interrupted only by Tombaugh Regio, and is thought to consist of complex hydrocarbons similar to those found elsewhere in the solar system (e.g. Titan). Published images also included enhanced-colour imaging, with MVIC colour filters overlaid on the panchromatic data. Figure I.26 shows an enhanced-colour map of Pluto from Olkin et al. (2017) with these and other globally significant features labelled.



Figure I.26: Global map of Pluto as imaged by New Horizons in enhanced colour, using the MVIC colour filters (two visible at 400–550 and 540–700 nm and one near-IR at 780–975 nm) as blue, green, and red respectively. Reproduced from Olkin et al. (2017).

2.4.2 Albedo and geology

Figure I.27 compares the panchromatic albedo map with the global digital elevation model, both produced by Schenk et al. (2018), allowing us to establish the relative elevations of the abovementioned significant geographic features. Notably, we can see that the bright nitrogen ice expanse of Sputnik Planitia corresponds to a deep basin; closer inspection reveals polygonal structures that imply this is convecting terrain (Trowbridge et al., 2016).

Other interesting geological features include a potential region of windblown CH_4 dunes in northwest Sputnik Planitia (Telfer et al., 2018), H_2O ice blocks probably transported by glacial flow (McKinnon et al., 2016), and dendritic dissected highland terrain in the upper latitudes (Moore et al., 2016).



Figure I.27: Global panchromatic albedo map and digital elevation model of Pluto, produced by Schenk et al. (2018).

2.4.3 Surface composition

MVIC

The colour imagery from MVIC shows that the darker-albedo material on Pluto's surface has a marked reddish tint — we can see this in Figure I.26 as well as in the spectral slope map shown in Figure I.28. In addition to this coarse analysis, we can study the surface methane distribution using the weak CH_4 ice absorption band map produced by Earle et al. (2018) (Figure I.29). Comparisons with LEISA maps are particularly useful due to MVIC's higher spatial resolution complementing LEISA's high spectral resolution, and we carry out these analyses in detail in Chapter II.



Figure I.28: Global mosaic of the MVIC spectral slope map in reflectance percent/100 nm, produced by Earle et al. (2018).



Figure I.29: Global mosaic of the MVIC equivalent width of the 980 nm methane absorption band, produced by Earle et al. (2018).

LEISA

Using the high-resolution LEISA observations, Schmitt et al. (2017a) perform a comprehensive qualitative analysis of the spatial distribution of the various materials present on the surface of the encounter hemisphere. The raw hyperspectral datasets are reduced and denoised using principal component analysis⁵ (results shown in Figure I.30). Schmitt et al. (2017a) then compute values including integrated band depths for CH_4 and N_2 ice as well as spectral indices for presence of H_2O ice and red material. We describe these in more detail in Chapter II.

⁵Principal component analysis is a statistical method for summarising large data matrices using a smaller set of "summary indices", the principal components, calculated such that they represent as much of the dataset's variation as possible in as few components as possible.



Figure I.30: Comparison between original and post-PCA noise-reduced LEISA data (reproduced from Schmitt et al., 2017a).

Figure I.32 shows the high-resolution encounter hemisphere qualitative abundance maps produced by Schmitt et al. (2017a), and Figure I.33 shows a suggested surface representation for interpreting these maps. According to the interpretations in Schmitt et al. (2017a), N₂-rich and CH₄-rich ices are present both separately and mixed together, possibly either as a two-phase system or as stratified phases. N₂-rich ice consists of CO and CH₄ molecules present at low concentration in a ternary molecular mixture with the dominant N₂, according to the phase diagrams. Also present are H₂O ice and a dark red organic material. This analysis is complemented by a first quantitative study using pixel-by-pixel Hapke modelling of the spectra (Protopapa et al., 2017), using a simplified surface representation with one stratum and only granular mixture.

2.4.4 Volatile surface distribution and atmospheric circulation analysis

One of Pluto's more unique features is the intercoupling of its atmospheric circulation and distribution of its surface materials. The low temperatures lead to surface deposition of volatiles such as N_2 and CH_4 , and seasonal temperature and insolation changes induce these volatile deposits to sublimate and circulate to be deposited again elsewhere. These global circulation patterns have been studied and modelised at length by authors including Bertrand and Forget (Bertrand and Forget, 2016; Bertrand et al., 2018, 2019).

One of the main features of Pluto examined in these works is the glacier of Sputnik Planitia, which represents a vast reservoir of N_2 ice — the Sputnik Planitia basin is 1000 km wide and 6–7 km deep. This majority- N_2 rich ice also contains some percentage of CH_4 and CO. Bertrand et al. simulated the evolution of the distribution of nitrogen and methane ice over the history of Pluto and made the following observations:

- The stable accumulation of N₂ in the Sputnik Planitia depression is due to its topography but also its equatorial location (Bertrand and Forget, 2016);
- The N₂ transport cycle is driven by obliquity cycles: high-obliquity periods show intense sublimation in the north of SP, while low-obliquity periods show condensation at the northern

and southern edges of SP combined with sublimation in its center (Bertrand et al., 2018, see Figure I.31);

• CH₄ deposits in equatorial latitudes are perennial and form cold traps for the condensation of N₂, with a similar mechanic also visible seasonally in higher latitudes (Bertrand et al., 2019).

These volatile transport analyses have built models coherent with observations using the highresolution encounter hemisphere data, but further work on global circulation models of Pluto will strongly benefit from global maps of N_2 and CH_4 distribution. We discuss this further in section 3.5 of Chapter 2.



Figure I.31: Sublimation-condensation rates of nitrogen in Sputnik Planitia at different timescales, with figure D showing a full obliquity cycle (2.8 Myrs) (reproduced from Bertrand et al., 2018).

2.4.5 Data processing and analysis to be done

While the datasets produced from the panchromatic LORRI and multispectral MVIC data are reasonably complete, and global mosaics have been made with them, the same had not been done for the hyperspectral LEISA datasets prior to this thesis. The pre-encounter data provides global coverage of the illuminated surface of Pluto, but misregistration between the different pre-encounter images as well as with the high-resolution map meant that work needed to be done to realign them before producing a global mosaic. In addition, once the global datasets were created, we did not have a robust surface representation for accurately quantifying the complex chemical mixtures present in Pluto's different terrains. The project of registering the different datasets to create global hyperspectral mosaics of Pluto is detailed in Chapter II, and the work to develop a spectral inversion model to quantitatively describe the variations of Pluto's surface composition is expounded upon in Chapter III.


Figure I.32: Pluto encounter hemisphere maps from Schmitt et al. (2017a) showing qualitative abundances for CH₄, N₂, CO, and H₂O ices, and organic red material.



Figure I.33: Schematic representation of the evolution and mixing of materials on Pluto's surface (modified from Schmitt et al., 2017a).

CHAPTER II

Global map registration

1 Introduction

The New Horizons flyby of Pluto returned a punctual set of high-resolution imagery of what is referred to as the encounter hemisphere, i.e. the side of Pluto that was oriented towards the probe as it made its closest-approach measurements. This dataset has formed the basis for most of the research and analysis so far of Pluto's surface (Schmitt et al., 2017a; Schenk et al., 2018; Moore et al., 2016, and others).

In addition to the high-resolution data, there exist lower-resolution datasets taken during the approach phase of the mission, which allow us to cover the entirety of the illuminated surface of Pluto (with the southernmost 30° being in polar winter): going backwards from the high-resolution encounter image, each preceding approach image shows Pluto at a slightly greater central longitude thanks to its rotation period of about 6.4 days, and combining them all together gives us global coverage. However, due to imprecision on the pointing information provided by the instrument there are up to several pixels of misalignment between the different datasets, i.e. up to several degrees of latitude and/or longitude, and this must be corrected in order to produce spatially correct global maps.

Creating a map from several different datasets is commonly referred to as producing a *mosaic*, or *mosaicing* the datasets together. Typically one will select one already well-registered dataset to use as the basis for the mosaic, known as the **target**, and will then align the other datasets (the **sources**) with the target, usually based on the position of certain surface features. This is known as *image registration*. Depending on how exactly the datasets are misaligned, they may only need to be shifted or rotated, or they may need to be rescaled or skewed, or even locally warped. Different methods exist for determining how to best do this, which are described in more detail in 2.2.

Due to the relatively low resolution of the LEISA approach datasets, the feature-based methods commonly used for image registration in planetary science are ineffective, and other methods must be explored. Notably, we adopt techniques from medical imagery to produce these global compositional maps.

2 Data and methods

2.1 LEISA approach datasets

During the approach phase each LEISA hyperspectral image fully contains Pluto and is recorded over a short period of time. This means that all channels within that image will have the same misalignment. Therefore, the transformation can be calculated using an arbitrary channel and then applied to all the other channels within the same image, and so we can select a channel or set of channels that features large contrasts and clear patterns for ease of calculation. The CH₄ ice map produced by Schmitt et al. (2017a) using the integrated band depth of the 1.7 μ m band group has precisely these properties, and served as the basis for the registration.

The raw hyperspectral images were manually pre-processed to remove bad pixels, denoise, and subtract the sky value¹, and projected orthographically according to observation geometry data from

¹The "sky" value corresponds to the "background" value of the dataset, i.e. the baseline when there is no signal.

the New Horizons SPICE kernel files². Figure II.1 shows an example image extracted from an approach datacube before and after the pre-processing. An initial low-precision registration was then carried out largely by eye, based on the high-resolution LORRI map that was reprojected to the same viewing geometry as the LEISA cube, in order to provide a reasonable starting point for the higher-precision registration to follow.



Figure II.1: A frame from LEISA approach image #12 (see Table II.1 for further information on the dataset), (a) before and (b) after preprocessing to remove noise, bad pixels, and sky value, and preliminary low-precision registration to (c) the LORRI panchromatic map projected to the same geometry.

In total we selected and processed a dataset of 12 approach images spanning about 5 days prior to the high-resolution images to create the global map (see Table II.1 for detailed information), which was the minimal number of images required for full coverage. Three of the images within that chronological span (MET 298853429, 298940959 and 299064869) were skipped due to either their redundancy with another observation very close to it in time, or spectro-photometric calibration issues. The data's spatial resolution ranges from 30.2 to 394.4 km/px. We used the highest-resolution images that were available for all regions but the resolution degradation between the lowest- and highest-resolution data is notable.

Given Pluto's rotation period of about 6.4 days, the sub-spacecraft longitude (i.e. the longitude of Pluto's surface directly below the spacecraft) changed by about 288.5°, i.e. about 0.80 rotations. On the other hand, the sub-spacecraft latitude only changed by about 5.2°. Figure II.2 shows that this dataset provides sufficient global coverage. This was evaluated using the initial misregistered data, but a comparison to the registered maps (Figures II.7 through II.10) shows that the misregistration is low enough to not substantially change this.

The last, most high-resolution full-disk³ approach image (#12 in Table II.1) was registered using the encounter-phase data as target image (HRI & HR2 in Table II.1, as well as UHR, which corresponds to the limited-coverage highest-resolution strip). This newly registered image was then used to register the previous image at slightly lower spatial resolution, and so on until the first selected approach phase image (e.g.#II to #12, #10 to #11, etc. until #1 to #2). This was done in lieu of registering all the images directly to the high-resolution image, as the overlap between the highest- and lowest-resolution data used was too small to allow good registration.

²SPICE is the NASA-developed information system used by many planetary observation missions to store navigation and observational geometry data. SPICE stands for: **S**pacecraft ephemeris; **P**lanet, satellite, comet, or asteroid ephemerides; Instrument information; orientation information, traditionally called the **C**-matrix; and Events information (NASA/JPL, 2021).

³A full-disk image is one that shows the entire observed object.

| # | Observation time | | Resolution | px/Pluto diameter | Sub-s/c lat | Sub-s/c long |
|-----|------------------|---------------------|------------|-------------------|-------------|--------------|
| | (MET) | (UTC) | (km/px) | | (*N) | (°E) |
| 1 | 298719334 | 2015-07-09 03:46:46 | 394.37 | 6.0 | 43.09 | 87.18 |
| 2 | 298767059 | 2015-07-0917:02:14 | 353.51 | 6.7 | 43.07 | 56.06 |
| 3 | 298824624 | 2015-07-10 09:01:31 | 304.31 | 7.8 | 43.05 | 18.56 |
| 4 | 298854539 | 2015-07-10 17:22:47 | 278.59 | 8.5 | 43.03 | 358.96 |
| 5 | 298891829 | 2015-07-11 03:42:21 | 246.78 | 9.6 | 43.01 | 334.75 |
| 6 | 298939609 | 2015-07-11 16:58:32 | 205.88 | 11.5 | 42.99 | 303.66 |
| 7 | 298995539 | 2015-07-12 08:30:09 | 158.00 | 15 | 42.96 | 267.30 |
| 8 | 299026199 | 2015-07-12 17:01:49 | 131.85 | 18 | 42.92 | 247.35 |
| 9 | 299079314 | 2015-07-13 07:46:25 | 86.37 | 28 | 42.82 | 212.94 |
| 10 | 299105209 | 2015-07-13 13:59:33 | 64.10 | 37 | 42.70 | 196.19 |
| 11 | 299127869 | 2015-07-13 21:18:57 | 44.58 | 53 | 42.49 | 181.68 |
| 12 | 299144829 | 2015-07-14 02:01:50 | 30.23 | 79 | 42.15 | 171.14 |
| HR1 | 299172014 | 2015-07-14 09:33:05 | 6.96 | 342 | 38.52 | 158.62 |
| HR2 | 299172889 | 2015-07-14 09:48:16 | 6.20 | 383 | 37.91 | 158.81 |
| UHR | 299176809 | 2015-07-14 10:56:19 | 2.75 | 865 | 30.73 | 164.17 |

Table II.1: List of Pluto approach-phase LEISA observations used, in chronological order, followed by the two high-resolution encounter-phase observations and the partial highest-resolution observation. The spatial resolution is given at the sub-spacecraft point (i.e.the point on the surface closest to the spacecraft).



Figure II.2: Global coverage obtained with the high-resolution (6-7 km/px) encounter dataset combined with the 12 approach images used (shown prior to registration). The highest-resolution (2.7 km/px) swath is also shown in dark red. The south pole is not shown as it is in polar night in the current season. The sawtoothed aspect of the approach data is due to reprojection of square pixels from orthographic to cylindrical.

As mentioned before, once the transformation has been computed, it can be applied to any or all wavelengths in the hyperspectral image. This allows us to produce correctly registered global maps based on any subset of the LEISA dataset. We generated four main maps for analysis, corresponding to the spectral indices defined by Schmitt et al. (2017a) as follows:

CH₄: integrated CH₄ band depth between 1.589 and 1.833 μ m over a group of 3 CH₄ bands at 1.67, 1.72, and 1.79 μ m.

$$BD(CH_4) = 1 - \frac{\int_{1.589\mu m}^{1.833\mu m} RF(\lambda) d\lambda}{\int_{1.589\mu m}^{1.833\mu m} Cont(\lambda) d\lambda}$$
(II.1)

N₂: integrated N₂ band depth between 2.121 and 2.160 μ m over the N₂ band at 2.15 μ m. The integration has been slightly extended to lower wavelengths compared to Schmitt et al. (2017a), as we realised that the N₂ band wing contributes up to 2.12 μ m in the wing of the strong 2.20 μ m CH₄ band. This resulted in a slightly increased signal-to-noise ratio.

$$BD(N_2) = 1 - \frac{\int_{2.160\mu m}^{2.160\mu m} RF(\lambda) d\lambda}{1.5 * (RF_{2.105\mu m} + RF_{2.13\mu m} + RF_{2.1675\mu m} + RF_{2.1755\mu m})}$$
(II.2)

.

 H_2O : spectral indicator using 10 wavelength bands centred around 1.39 and 2.06 μ m.

$$SI(H_2O) = 1 - \frac{\int_{2.022\mu m}^{2.090\mu m} RF(\lambda)d\lambda}{\int_{1.365\mu m}^{1.410\mu m} RF(\lambda)d\lambda}$$
(II.3)

Red material: spectral indicator using the wavelength ranges around 1.430 and 1.658 μ m.

$$SI(RedMat) = 1 - \frac{\int_{1.447\mu m}^{1.421\mu m} RF(\lambda) d\lambda}{\int_{1.670\mu m}^{1.641\mu m} RF(\lambda) d\lambda}$$
(II.4)

The BD(CH₄), SI(H₂O) and SI(RedMat) indices were calculated after processing the data cubes with a global PCA (over almost all spectral channels), followed by an inverse PC rotation on a selected number of axes to reduce noise and instrument artifacts. The BD(N₂) was calculated after similar processing, but over a restricted spectral range (40 wavelength bands) centred on the band peak (see Schmitt et al., 2017a, for details). For the few lowest-resolution images, the PCA did not reduce the instrument artifacts because the number of pixels was insufficient for statistical evaluation, and so the PCA-processed data was not used.

2.2 Map registration

Image registration involves transforming one or more source images to match a target image. Image registration methods can be classified into two broad categories: feature-based and intensitybased. Feature-based registration is based on identifying and aligning distinctive features (such as points, lines or contours) in the images, and is the most common type of method used in remote sensing (Zitova and Flusser, 2003). However, it is labour-intensive, as feature mapping frequently needs to be done manually by selecting the anchor points, and its application is limited to datasets in which sharp, easy-to-map features (crater rims, cliffs, faults, etc.) are present and still recognisable in the lowest-resolution image used. Moreover, the alignment precision is thus limited to that of the lowest-resolution image, as we cannot map the features with a higher accuracy than that available in the data.

The second type of method, intensity-based registration, involves comparing intensity patterns in the images to be registered. An image is characterised by having a spatial arrangement of brighter and darker areas, and comparing their spatial distribution lets us identify a potential transformation to be made in order to make the images more similar. Different measures of correlation, called metrics, can be used to evaluate their similarity, and the choice of metric will depend on the properties of the dataset.

A **correlation** is a statement of statistical relationship between two variables or datasets. The most commonly-used correlation coefficient is the **Pearson correlation coefficient** *r*, which reflects the strength and direction of a linear relationship between datasets or data samples. It is calculated as follows:

$$r_{xy} = \frac{n\sum x_{i}y_{i} - \sum x_{i}\sum y_{i}}{\sqrt{n\sum x_{i}^{2} - (\sum x_{i})^{2}}\sqrt{n\sum y_{i}^{2} - (\sum y_{i})^{2}}}$$
(II.5)

where *n* is the sample size, and x_i and y_i are individual data points from samples *x* and *y* respectively.

Pearson's correlation coefficient is a convenient tool for a first-order evaluation of whether a relationship may exist between two datasets, but it is not statistically robust, i.e. is sensitive to outliers. It also fails to represent many kinds of nonlinear relationships, as can be seen in Figure II.3 where it is compared with a *nonlinear* correlation coefficient computed from the **mutual information** of two datasets.



Figure II.3: Figure reproduced from Vu et al. (2018). A comparison of the **Pearson's correlation coefficient** (PCC) and a mutual information-derived *nonlinear correlation coefficient* (MICC) computed for several different distributions of two random variables. The MICC is a transformation of MI that normalises it into the [0, 1] range. For each panel, top left shows the PCC in **bold** and top right shows the MICC in *italic*.

Mutual information is a nonlinear measure of correlation that evaluates the mutual dependence between two datasets. More specifically, it quantifies how well one can use one dataset as a predictor for the other. For a 2D image, this means how accurately we can predict the value of a pixel in a specific location based on the value of the pixel in the same location of a different image.

The first step to calculating the mutual information of two datasets is computing their **joint distribution**, which can be represented as a joint histogram. A standard one-dimensional histogram represents the distribution of values in a dataset by assigning to bin number x the frequency of the value x in the dataset. A joint histogram extends this to pairs of values (x,y) in datasets X and Y respectively. The intensity of a pixel located at (a,b) represents how often datapoints of value a in dataset X spatially correspond to datapoints of value b in dataset Y. Figure II.4 shows how a joint histogram may be used to evaluate the misalignment of two images: the relationship is perfectly linear for an image compared with itself, but translating one of the copies makes the relationship more diffuse.



Figure II.4: Figure reproduced from Liang et al. (2014). Illustration of the effect of misregistration, showing the joint histogram of a remote-sensed image with itself. In (a) no transformation is performed. In (b), (c), and (d), one image is translated by 0.5, 1, and 2 pixels, respectively.

Intensity-based registration is underused in planetary science, but is a very common tool in medical imagery processing, where it is used for MRI (magnetic resonance imaging) or CT (X-ray computed tomography) scan images (Pluim et al., 2003). This means most toolkits developed for intensity-based registration are designed around medical uses — but they can be very easily adapted to planetary data (see Liang et al., 2014, for a demonstration of a combined feature- and intensity-based registration technique on Landsat data).

Registration methods may also be categorised according to what kind of transformations they use for aligning the images. The transformation may be global, i.e. linear, with operations such as translation, rotation, scaling, and shear affecting the whole image uniformly, or local/nonlinear, i.e. defined by a warp field whose effect varies across the image. Figure II.5 shows these common transformation types, as per Uchida (2013). The errors in registration of the LEISA data are caused by imprecision in the instrument pointing, resulting in an overall shift of the image and possibly a slight rotation. We can therefore restrict ourselves to global transformations applied to the original orthographic projection.



Figure II.5: Figure reproduced from Uchida (2013). Transformation functions for image registration, with affine transformation being a combination of translation, rotation, scaling, and shear.

2.3 Registration algorithm details

We used intensity-based registration algorithms from SimpleITK (Lowekamp et al., 2013), a userfriendly Python interface to ITK (Insight Segmentation and Registration Toolkit, Schroeder et al., 2003). ITK is a cross-platform open-source framework that provides image segmentation and registration algorithms. It is a powerful and flexible toolkit but using it requires significant expertise in C++ templates, and so an interface was developed that can be accessed in multiple higher-level languages⁴ and that calculates many of the input parameters, such as image dimensions and pixel types, for the user.

As we expect the misalignment between the datasets to be due entirely to imprecisions in the spacecraft pointing information, we restricted ourselves to global similarity transformations (i.e. translation, rotation, and scale). The higher-resolution image was bilinearly resampled⁵ to the lower resolution, then both the source and target images were resampled to 2 km/px to allow for subnative-pixel registration precision.

We used Mattes mutual information metric (Mattes et al., 2001), a popular implementation of mutual information that estimates the MI using a probability density function, and an evolutionary algorithm-based optimiser. Specifically, we used (1+1)-evolutionary optimisation, the simplest kind of evolutionary algorithm that operates on one parent and one descendant. The current transformation (the parent) is slightly modified (mutated), creating a new transformation (the descendant). If the descendant yields a better metric value, then it is selected as the new descendant. The algorithm is probabilistic in its use of a gaussian probability function in defining the possible search-space when generating modifications, and is sometimes referred to as a "probabilistic hill-climber" (Droste et al., 2002).

⁴High-level programming languages, such as Python or Java, are more user-friendly as they are closer to natural languages and implicitly handle low-level aspects of computation such as memory management.

⁵Linear interpolation draws a linear slope between two points to determine the value of a point in between them. Bilinear interpolation is an extension of linear interpolation to two dimensions, where the data is linearly interpolated first along one direction, then along the other direction.

| Metric | Mattes mutual information | | |
|---|--|--|--|
| Number of bins Percentage of pixels sampled | 100 100% | | |
| Optimiser | l+l evolutionary | | |
| Maximum iterations Convergence tolerance Initial search radius Growth factor | 20-100 1.5×10 ⁻⁶ 10 ⁻³ 1.01 | | |

A mask was applied to restrict the metric to the overlapping areas between both images. Table II.2 shows the parameters used for the metric and optimiser.

Table II.2: Registration parameters used for global map creation.

-150° W -140° W -130° W a) b) c) 50° N 50° N 50° N 50° N CH, band depth High : 0.48 Low:0 40° N 40° N 40° N Control points Control points 30° N 30° N 30° N 30° N 20° N 20° N 20° N 20° N -140° W -130° W -150° W -150° W -140° W -130° W

2.4 Validation of registration accuracy

Figure II.6: LEISA registration validation with control points: a) high-resolution data; b) misregistered low-resolution data (observation #10); c) registered low-resolution data.

The registration results were verified in two steps. First, those low-resolution images that had sufficient overlap with the high-resolution data were checked with a set of control points based on features identifiable in the high-resolution map. Figure II.6 shows this for observation #11 as defined in Table II.1. For observations #7–#12, registration was accurate within 0.5 native pixels.

Second, those observations which did not have enough overlap with the encounter hemisphere, or a sufficiently high resolution to do feature-based validation, were visually compared with the global panchromatic reflectance map produced by Schenk et al. (2018), as the 1.7 μ m CH₄ band is strongly correlated with the visible reflectance ($r = 0.655^6$ as calculated using the high-resolution CH₄ map).

 $^{^{6}}$ Here and elsewhere in this chapter *r* is the Pearson correlation coefficient weighted by the cosine of the latitude to account for the area distortion due to the cylindrical projection. The cutoffs for describing the correlation have been chosen as follows, to give a simple linear scale:

| | | $0.25 \le r \le 0.5$ | correlated |
|-----------------------|-------------------|------------------------|--------------------------|
| r < 0.05 | not correlated | $0.5 \leq r < 0.75$ | strongly correlated |
| $0.05 \le r < 0.25$ | weakly correlated | 0.75 < r | very strongly correlated |

This second validation step showed that some latitudinal drift occurred for observations #1–#6. As their resolution was quite low (> 200 km/px), the maximum accuracy of any fit is limited by feature distortion. The transformation tended to overfit in trying to compensate, which was corrected by reducing the number of iterations for these six observations.

3 Results and analysis

The results of the described registration method, while applicable to the full spectral extent of the 12 LEISA images, are shown here through analyses of a set of four principal maps (CH₄ 1.7 μ m integrated band depth, N₂ integrated band depth, H₂O spectral index, and red material spectral index) and two auxiliary maps (CH₄ 1.95 μ m integrated band depth and CH₄ band shift spectral index). First, we look at the global value distribution of these maps and describe the implications of this for the global variations of Pluto's composition. Subsequently, we analyse the correlations between these maps as well as between the LEISA maps and other Pluto New Horizons datasets, such as the visible panchromatic map, the near-IR MVIC maps, and the elevation map, and suggest interpretations of these correlations in terms of Pluto's surface composition and geology.

3.1 Global maps and their main features

Figures II.7 through II.10 show the four registered global maps in simple cylindrical projection, at a sampling resolution of 2 km/px (which corresponds to a significant upsampling of the native resolutions of all approach-phase LEISA observations). Figure II.11 shows the same maps at the same resolution but in North polar orthographic projection. Figures II.12 and II.13 show the latitudinal and longitudinal distributions respectively for these four spectral indices, with binning resolutions of 2 px/° and 1 px/°. The latitudinal distributions have been weighted by the cosine of the latitude to recover the true surface area of the pixels and thus compensate the cylindrical projection. There is data at all longitudes down to 20°S latitude, and more limited longitude coverage down to about 40°S, close to the polar night limit at the time of the encounter. The detection limits and colormap ranges were taken from Schmitt et al. (2017a).

Two additional maps were generated for further analysis — the integrated band depth of the weak CH_4 band at 1.95 μ m (Figure II.14, exact definition of band given in Schmitt et al., 2017a), and a CH_4 band position index (Figure II.15), which is a metric representing the state of CH_4 at the surface, from CH_4 diluted in N₂-rich ice for the smaller index values to CH_4 -rich ice for the highest. Intermediate values mean that both phases are present at the pixel scale, either spatially distributed, or intimately mixed at grain level, or stratified vertically.

3.1.1 CH_4 integrated 1.7 μ m band depth

The CH₄ integrated 1.7 μ m band depth map (Figure II.7) shows that the global longitudinal distribution of CH₄ (shown in Figure II.13) is largely uniform across most of Pluto, albeit with a clear depletion over the longitudinal range of Cthulhu Macula (70°E to 160°E, centered at (0°N, 90°E)). The latitudinal distribution of CH₄ (Figure II.12) is more contrasted, with a global presence in the north polar areas, a more dispersed presence in the mid-northern latitudes and a strong CH₄ signal in a belt between 0 and 30°N from 160°E to 70°E. Both the CH₄-rich north polar cap and the equatorial CH₄ partial belt are clearly visible in the latitudinal distribution of the data in the isolated BT and SP subsets in Figure II.12. The abundance of CH₄ seems also to increase below about 15°S, close to the shadowed south polar region. This last area is separated from the north tropical CH₄ belt by a south tropical belt of red material (see below).



Figure II.7: Global registered CH₄ 1.7- μ m integrated band depth map of Pluto, with the three largest geologic features labelled.

3.1.2 N_2 band depth

The N₂ band depth map (Figure II.8) clearly shows the nitrogen-rich basin of Sputnik Planitia (centered at 20°N, 179°E), easily identifiable as a bump in the centre of the longitudinal distribution (Figure II.13). We also see an asymmetry around 180°E which reflects the east-west contrast in ice composition within SP (the bright and dark plains). N₂ ice is also strongly present in a northern belt between 30°N and 75°N, with a strong peak around 43°N (Figure II.12), as well as in the north tropical CH₄ belt (east of 180°) with a peak around 20°N in phase with the one of CH₄. Its correlation with the 1.7 μ m CH₄ band is not very strong, but positive (r = 0.189). Due to the weakness of the N₂ band and the low signal-to-noise ratio close to the terminator⁷ it is difficult to confirm that an increase of N₂ seen southwest of Sputnik Planitia in the high-resolution part of the map also occurs below about 15–20°S at all longitudes, and that may be correlated with the mid-southern latitude increase of CH₄ ice seen at least between 100°E and 280°E (Figure II.7).



Figure II.8: Global registered N₂ 2.15- μ m band depth map of Pluto, with Sputnik Planitia labelled. The detection limit is at a value of 0.005.

3.1.3 H₂O and red material spectral indices

The H_2O spectral index map (Figure II.9) shows water ice presence between 50°N and 30°S. This overlaps with the global red material belt as seen on Figure II.10, which fills the space between

⁷The terminator is the division between the daylit side and the dark night side of a planetary body.

10°N and 20°S except at Sputnik Planitia — although small quantities of red material are indeed visible in the dark plains of SP. There are only small differences between the red material and H_2O distributions, as they seem well-correlated when the red material is moderately abundant but start to be anti-correlated when it is very abundant, as in the centre-east of Cthulhu Macula (115–160°E), in the east of Krun Macula (220°E), and in Balrog Macula (280°E). This correlation was already evident from the high-resolution map, but is now confirmed globally (r = 0.490).

A small bump of red material abundance also occurs just above 30°N (Figure II.12). The overabundance of H₂O both between 70–160°E and 195–220°E is due to a combination of the presence of large maculae (Cthulhu, Krun) at south tropical latitudes and concentrated water ice spots at midnorthern latitudes (30°N) at these longitudes. As expected, both the H₂O ice and the red material are very strongly anti-correlated with the 1.7 μ m CH₄ band (r = -0.751 and r = -0.757 respectively).



Figure II.9: Global registered H_2O spectral index map of Pluto, with Cthulhu Macula and Sputnik Planitia labelled. The detection limit is at an index value of -0.23.



Figure II.10: Global registered red material spectral index map of Pluto, with Cthulhu Macula and Sputnik Planitia labelled. The detection limit is at an index value of -1.3.



(a) Global registered CH_4 1.7- μ m integrated band depth map of Pluto. (b) Global registered N_2 2.15- μ m band depth map of Pluto. The detection limit is at 0.005.



(c) Global registered H₂O spectral index map of Pluto. The detection (d) Red material spectral index map of Pluto. The detection limit is at -0.23. at -1.3.

Figure II.11: Registered global LEISA maps in north polar orthographic projection.



Figure II.12: Distribution of LEISA dataset band depth and index values as a function of latitude, binned at 2 pixels per degree. The distribution is computed for the global dataset, for the Cthulhu Macula longitude range (70–160°E, labelled CM), for the Sputnik Planitia longitude range (160–240°E, labelled SP), and for the bladed terrain range (0–70 and 240–360°E, labelled BT). The points represent the mean value over longitude for that latitude range and the error bars are the standard deviation.



Figure II.13: Distribution of LEISA dataset band depth and index values as a function of longitude, binned at 1 pixel per degree. The distribution is computed for the global dataset, for the mid-to-equatorial latitudes (0–45°, labelled EQ), and for the mid-to-polar latitudes (45–90°, labelled POL). The points represent the mean value over latitude for that longitude range and the error bars are the standard deviation. The distributions are weighted by cos(lat) to compensate for the cylindrical projection.



Figure II.14: Global registered CH₄ map of Pluto (1.95 μ m integrated band depth).



Figure II.15: Global registered CH₄ band position index (CH₄ state index) map of Pluto. The lower index limit is at a value of 36.5.

3.2 Analysis of correlations between the LEISA maps

In addition to computing the simple correlation coefficients between the different LEISA maps, we also computed the two-dimensional joint correlations in order to identify spatially localised populations. The more significant of these correlations are shown as hexagonal bin plots in Figure II.16, and may be compared with some of the same calculations carried out for the high-resolution hemisphere by Schmitt et al. (2017a).

Figure II.16a shows multiple trends within the global positive correlation between the 2.15- μ m N₂ and 1.7- μ m CH₄ integrated band depths:

- a vertical trend, with variable CH₄ band depth and negative N₂ band depth (marked in red) that corresponds to fine- to coarse-grained CH₄-rich ice.
- a negative correlation trend, going from strong CH₄ values and absent N₂ to strong N₂ and weak CH₄ values (marked in orange), corresponding to coarse-grained N₂-rich ice containing low to medium amounts of dissolved CH₄ (< 1%).
- a horizontal trend with high CH₄ band depth and variable N₂ values (marked in yellow), that corresponds to very coarse-grained N₂-rich ice (> 20 cm) with a high CH₄ content (> 1%).

These trends correspond largely to those seen in Figure 18A of Schmitt et al. (2017a), and we used the terrain classes identified by the latter as well as a Gaussian mixture clustering model (Li



et al., 2013) as implemented in the python module scikit-learn (Pedregosa et al., 2011) to generate a classification of the global map, shown in Figure II.18.

Figure II.16: Hexagonal bin plots showing correlations between different LEISA maps. The detection thresholds for each map are marked with a red dotted line, and correlation trends are marked with solid lines.

Figure II.16b shows the correlation between the N_2 integrated band depth map and the CH_4 band position index map (i.e. the CH_4 state map). We observe a global anti-correlation (marked in orange) that can be classified into a spectrum of terrains going from N_2 -rich to CH_4 -rich, as shown by Schmitt et al. (2017a) in their figures 22 and 39 for the encounter hemisphere. We compute a similar classification for the global map and show it in Figure II.19.

Figure II.16c shows a simple strong linear correlation between the H_2O and red material spectral indices. A global segmentation of the four types of terrains shown visible in this correlation plot is presented in Figure II.17. The pixels which have a positive value for both datasets correspond to the top right (yellow) quadrant. There is, however, a number of pixels in Figure II.16c displaying an anti-correlation, i.e. a positive presence of red material but without detection of H_2O ice (top left quadrangle, blue). They correspond mostly to the centre of the maculae, where the red material either fully covers the water ice crust or is mixed with sufficient quantities of N_2 and CH_4 ice to attenuate the water ice signal. On the other hand, there are clearly only very few pixels with pure water ice without red material (bottom right quadrangle, green). The bottom left quadrangle (pur-

ple) corresponds to the terrains fully covered by N₂-rich and/or CH₄-rich volatile ices, thus hiding the non-volatile materials. We can clearly see this anti-correlation between volatile ices and red material in Figure II.16d, trend line marked in orange.



Figure II.17: Global classified map based on the correlation plot of the H₂O and RedMat spectral indices, defined by the dataset detection thresholds.



Figure II.18: Global classified map based on the correlation plot of the N₂ and 1.7 μ m CH₄ integrated band depths, calculated with a Gaussian mixture clustering model.



Figure II.19: Global classified map based on the correlation plot of the CH₄ band position index (CH₄ state index) and N₂ integrated band depth, calculated with a Gaussian mixture model.

3.3 Comparison with MVIC maps

We refer to Figures I.29 and I.28 of Chapter I for the methane narrow-band (980 nm) equivalent width map and the global spectral slope map obtained by Earle et al. (2018) from the analysis of the MVIC dataset. Figure II.20 shows the 2D correlations between these two maps and several of the LEISA maps.



Figure II.20: Hexagonal bin plots showing correlations between different LEISA maps and the MVIC maps. The detection thresholds for each map are marked with a red dotted line.

The LEISA strong-band 1.7 μ m CH₄ band depth map is somewhat correlated with the MVIC 980 nm-band CH₄ map (r = 0.445), but Figure II.20a shows that there is no single strong trend to be identified in the correlation plot. We expected the weak-band 1.95 μ m CH₄ band depth map to be more strongly correlated with the MVIC 980 nm band, and while the overall correlation is slightly stronger (r = 0.531), the correlation plot (Figure II.20b) again shows no particularly strong trends. In particular, both LEISA band depth maps display areas with abundant methane where MVIC did not see any, or at very weak levels (bottom right quadrangles and in particular the yellow spot along the right part of the red horizontal line in both figures). Conversely, over a number of pixels MVIC has positive detection of CH₄ where LEISA detects none of the CH₄ bands (upper left quadrangle). Figure II.21 shows which spatial regions these quadrangles correspond to. We can see that the LEISA-only values (in green) are present almost exclusively in the high-resolution data, while conversely the MVIC-only values (in blue) are seen largely in the low-resolution part of the map. This seems to suggest that while the encounter-phase LEISA measurements are more sensitive to CH₄ than MVIC, the signal averaging that occurs at lower resolutions means this sensitivity is lost. This may be due to the spatial scale of this type of CH₄-rich area falling below the spatial resolution of the LEISA data

faster than for MVIC, as the spatial resolution of the MVIC map is typically 3 times higher than for our global LEISA map.

In Figure II.20c and d we see that the MVIC slope map is correlated relatively strongly with the red material spectral index (r = 0.675) as well as the H₂O spectral index (r = 0.529), with overall similar distributions. The latter figure is consistent with the correlation and terrain classification calculated by Schmitt et al. (2017a) for the H₂O spectral index, and we adapt their terrain definitions to produce a similar classification (Figure II.22).



Figure II.21: Global classified map based on the correlation plot of the LEISA CH₄ 1.95 μ m band depth map and the MVIC CH₄ 980 nm band map, defined by the dataset detection thresholds.



Figure II.22: Global classified map based on the correlation plot of the H₂O spectral index and MVIC slope, defined from a Gaussian mixture model and the segmentation in (Schmitt et al., 2017a).

Earle et al. (2018) also developed a terrain classification based on the MVIC CH₄ and slope maps (see their Fig. 15). The six terrain types established are: **pure** CH₄-**rich**; CH₄-**rich diluted with** N₂; N₂- **or** CO-**rich areas with some** CH₄; **mix of** CH₄ **and tholins**; **moderate tholin deposits with some** CH₄; and **substantial tholin deposits with little/no** CH₄. Figure II.23 shows histograms for the four LEISA band depths (N₂, 1.7 μ m CH₄) and spectral indices (H₂O and red material) within these six terrain types.



Figure II.23: Histograms of the value distribution of the four LEISA maps (N₂ and 1.7 μ m CH₄ band depths, H₂O and red material spectral index) for the six MVIC terrain types, weighted to compensate for distortion due to the cylindrical projection and normalised to the terrain area.

The global methane-rich belt between 0 and 30°N largely corresponds to the "pure CH₄-rich" terrain type on the MVIC terrain map (band depth mean $\mu = 0.359$, $\sigma = 0.056$), and the diluted methane presence at mid-northern latitudes relatively well corresponds to the "CH₄-rich diluted with N₂" terrain ($\mu = 0.352$, $\sigma = 0.091$), as shown in Figure II.23a. The histograms extracted from the N₂ band depth map (Figure II.23b) show N₂ presence within both the "N₂-rich with some CH₄" and "CH₄-rich with N₂" MVIC terrain types, but also in all other MVIC terrain types. As expected, the MVIC channels are only poorly sensitive to the presence of N₂ ice through the general anticorrelation between the presence of volatile ices and the red slope produced by the red material, due to the segregation of both components through sublimation-condensation cycles. Even if the LEISA N₂ band depth is not a direct measurement of the abundance of nitrogen ice at the surface, and is sensitive to other parameters such as grain size and CH₄ abundance, it is far more sensitive to its presence than the MVIC data. Looking at the N₂ and CH₄ band depth-based classification in Figure II.18, we can see several distinct regions corresponding to various coarse-grained N₂-rich ices with high to medium CH₄ content that are not mapped as containing N₂ ice using the MVIC CH₄-band/red slope classification alone.

The LEISA red material spectral index map (Figure II.10, histograms in Figure II.23c) corresponds primarily to the "substantial tholin deposits" terrain type (spectral index mean $\mu = -0.205$, $\sigma = 0.274$), while the spectral index for the two tholin-CH₄ mixed terrain types is significantly lower ("moderate tholin deposits, some CH₄": $\mu = -0.576$, $\sigma = 0.348$; "mix of CH₄ and tholins": $\mu = -0.706$, $\sigma = 0.463$). Our red material indices are thus well-correlated in the terrains dominated by this material. The H₂O ice spectral index map exhibits a similar distribution to the red material map (Figure II.23d), with the strongest presence in the "substantial tholin deposits" terrain ($\mu = 0.090$, $\sigma = 0.212$) and in the "moderate tholin deposits, some CH₄" terrain ($\mu = 0.064$, $\sigma = 0.208$). This is consistent with the LEISA-based classification in Figure II.22 and confirms that H₂O ice is mostly spatially coincident with the red material.

3.4 Composition relationship with geology

We are able to evaluate the relationship between the LEISA composition maps and the geology of Pluto's surface using the LORRI panchromatic reflectance map and the digital elevation model (both shown in Figure I.27) produced by Schenk et al. (2018).

While both the LEISA CH₄ and N₂ maps (Figures II.7 and II.8) are globally anti-correlated with altitude, this is largely due to the massive presence of the convective Sputnik Planitia basin, which is N₂-rich with some diluted CH₄. The two maps are strongly anti-correlated with altitude for the area of the DEM containing Sputnik Planitia (between 140 and 200°E and south of 40°N, $r(CH_4) = -0.501$ and $r(N_2) = -0.569$), but outside this area CH₄ is positively correlated with the DEM (r = 0.211), while the N₂ anti-correlation is much weaker than it is within SP (r = -0.161). The two separate populations of CH₄ and N₂ terrains can also be clearly seen in the 2D hexagonal bin correlation plots (see II.24).



Figure II.24: Hexagonal bin plots showing correlations between different LEISA maps and the digital elevation map. The detection thresholds for each map are marked with a red dotted line.

These CH₄-rich elevations appear to mostly correspond to the bladed terrains as described in Moore et al. (2018), and can be visually identified in the LEISA 1.7 μ m CH₄ map as the methane-rich belt between 0 and 30°N. The N₂-rich medium-altitude areas, accordingly, correspond to the smooth plains present in the bright pitted uplands as described in Moore et al. (2016), and are identifiable between 30°N and the equator as well as intermittently between 0 and 30°S. Figure II.25 and II.26



show overlays of the CH₄ and N₂ maps respectively over the DEM, where these terrains can be seen.

Figure II.25: Global LEISA CH₄ 1.7 μ m band depth map overlaid over the digital elevation model of Pluto.



Figure II.26: Global LEISA N₂ 2.15 μ m band depth map overlaid over the digital elevation model of Pluto.

The CH₄ and N₂ presence in this belt extends fully across the low-resolution hemisphere, and so we may posit that the "bright pitted upland and bladed terrain"-type landscape also continues eastward across Pluto until approximately 80° E, although it is mostly coated with dark red material in the southern part of the belt between 0 and 30° S.

Calculating the correlations between the datasets as a function of latitude (see II.27) shows that in addition to the equatorial methane belt, the CH₄ 1.7 μ m band depth is also correlated with altitude at higher latitudes (between 40 and 80°N, peaking at 60°N). These CH₄-rich uplands match the dissected and eroded terrains described in Moore et al. (2016), while the lower-altitude N₂-rich areas within them correspond to the smooth infill of the pits and depressions contained in these terrains.

Conversely, while H_2O ice is globally weakly correlated with altitude (r = 0.118), the latitude- and longitude-dependent correlation plots show that elevated water-ice terrains are found primarily in two latitudinal belts, at about 30°N and 10°S, and that within the encounter hemisphere they are restricted to the 150-180°E range. This suggests that the global H_2O -altitude correlation is due mostly to the water-ice mountain ranges bordering the west side of Sputnik Planitia (Howard et al., 2017; White et al., 2017), and water ice found elsewhere corresponds mostly to partially exposed substrate terrain, underlying either a layer of red material at sub-equatorial southern latitudes or



CH₄-rich uplands at northern latitudes. This is supported by computing the H₂O–DEM correlation for the SP area (positive, r = 0.448) separately from the rest of the DEM (weakly negative, r = -0.221).

Figure II.27: Correlations between the datasets used as a function of latitude, grouped by LEISA datasets. For each pair of datasets, the Pearson correlation coefficient is calculated for each row of data points, corresponding to a range of 1.7 degrees of latitude.



Figure II.28: Correlations between the datasets used as a function of longitude, grouped by LEISA datasets. For each pair of datasets, the Pearson correlation coefficient is calculated for each row of data points, corresponding to a range of 0.65 degrees of longitude. The calculation is weighted to compensate for latitude-based distortion due to the cylindrical projection.

3.5 Utility of global composition maps for atmospheric circulation

As summarised in section 2.4.4 of Chapter I, Bertrand et al. have dedicated multiple works (Bertrand and Forget, 2016; Bertrand et al., 2018, 2019) to simulating the long-term volatile transport history that would result in a present-day ice reservoir distribution coherent with New Horizons observations. These studies have relied on composition information from both the high-resolution hemisphere LEISA data and the global MVIC maps, but have not been able to take into account global composition maps from hyperspectral data, which will provide more accurate information on both the latitudinal and longitudinal asymmetries of the volatile ice reservoirs as well as on their total volume. Notably, Bertrand et al. (2018) mention that the location of perennial N₂ ice deposits should depend on the brightness of mid-to-polar CH₄ deposits, which can be more easily quantified with a global map. Also, Bertrand et al. (2019) predict the deposition pattern of N₂ outside of Sputnik Planitia to consist of very localized patchy deposits in the equatorial deep depressions of the BT and mid-latitude deposits (\pm 30–60°). This was consistent with the encounter hemisphere maps released by Schmitt et al. (2017a) and Protopapa et al. (2017) but can now be verified globally and can be easily observed to be true from the latitudinal map distributions (Figure II.12). All types of global or local energy balance models or climatic models will benefit from these new global maps of the different materials present at the surface of Pluto, either as an input to calculate the current thermal balance of the planet, or as a constraint for the current volatile ice distribution to be obtained (e.g. Lewis et al., 2021; Johnson et al., 2021; Bertrand et al., 2020). For a specific example, as on Pluto the general atmospheric circulation is driven by the N₂ condensation–sublimation flows, the global maps of distribution of surface N₂ ice are key for simulating realistic atmospheric circulation.

4 Summary

4.1 Global compositional cartography

We have co-registered a dataset of low-resolution LEISA hyperspectral images, recorded during the approach phase, with the high-resolution closest-approach data from the New Horizons encounter with Pluto, and produced global maps for the following spectral indicators:

- 1. CH_4 1.7 μ m integrated band depth 3. H_2O spectral index
- 2. N₂ 2.15 μ m integrated band depth

We compare these maps with other global Pluto datasets based on LORRI (topography and panchromatic reflectance, Schenk et al., 2018) and MVIC measurements (980 nm CH₄ band and spectral slope, Earle et al., 2018), as well as previous works on the high-resolution LEISA dataset (Schmitt et al., 2017a; Protopapa et al., 2017). We are able to globally confirm the latitudinal composition

4. Red material spectral index

et al., 2017a; Protopapa et al., 2017). We are able to globally confirm the latitudinal composition variation proposed by Protopapa et al. (2017) — a red carbonaceous material belt between 0 and 30°S, a CH_4 -rich ice belt between 0 and 30°N, a mixture of N₂-rich terrains with some CH_4 -rich areas between 30 and 60°N, and a CH_4 -rich North polar area — which corresponds well to the categorisation of CH_4 and red material terrains based on MVIC data presented in Earle et al. (2018). In addition, the LEISA data allows us to correctly globally map the presence of N₂ and H₂O ices.

Based on the composition maps' correlation with geology and topography we also propose geological interpretations for compositional features in the anti-encounter hemisphere, such as:

- presence of CH₄-rich dissected/eroded terrain in high latitudes;
- further evidence of a largely continuous equatorial belt of CH₄-rich bladed terrain punctuated with N₂-rich lowlands;
- presence of H₂O ice at longitudes outside Sputnik Planitia corresponding to exposed substrate.

4.2 Evaluation of intensity-based registration

The work presented in this chapter is one of the first applications of intensity-based registration in planetary cartography, and shows the promising potential of this tool. Figure II.6 shows the subpixel accuracy of the technique in the regions that are verifiable with control points based on the high-resolution data, and this shows that the accuracy remains constant when co-registering maps with resolutions as low as 158 km/px (the lowest tested resolution) and no easily-delineable features. While intensity-based registration alone becomes insufficient when working with significantly lower resolutions, it can be easily combined with a large-scale feature-based algorithm that helps maintain common-sense constraints (e.g. making sure the object being mapped remains within its possible range of boundaries) to create a fully automated all-purpose registration algorithm.

4.3 Applications for quantitative global mapping of Pluto

The maps presented in this chapter provide a very valuable dataset in themselves, allowing for the first time for a global look at Pluto's compositional variability. We have carried out a set of statistical and comparative analyses that is only the first step of making good use of this data. The most important next step is finding a way to quantitatively translate these spectra into an accurate surface model that shows the proportions and mixture types of all the different chemical components present. We make a foray into this quantification in the following chapter, which demonstrates an application of simulated annealing to spectral inversion, and presents results for several data points from the LEISA maps.

Spectral modeling and inverse fit

1 Background and problem definition

1.1 Summary of previous work

Our work on map registration as described in Chapter II resulted in a complete global hyperspectral dataset, with a spectrum for each pixel of Pluto observed by New Horizons (excepting those in polar night at the time of the flyby). We can then study these spectra directly, using band depths and spectral indices such as band ratios to qualitatively identify which chemical species are present. This was done by Schmitt et al. (2017a) for the high-resolution encounter hemisphere data, and these analyses are extended to the global data in section 3 of Chapter II. Due to the complexity of the surface, we must turn to spectral inversion to get quantitative information on the proportions between different surface components, on their texture, and how they're spatially arranged. Section 1.4 of Chapter I gives a brief introduction to radiative transfer models, and what assumptions and approximations they tend to make.

Depending on the model used and the assumptions made, the surface representation can be simple or complex, and the size of the problem that needs to be solved can range from a handful to several dozen dimensions. Protopapa et al. (2017) quantified the encounter hemisphere spectra using a Hapke model with four possible surface materials — H_2O ice, N_2 -rich ice in a molecular mixture with CH_4 ice (0–5% of CH_4), CH_4 ice saturated with N_2 (at about 3%), and Titan tholins — spatially mixed in a single layer. This is an eight-dimensional problem, with a grain size parameter for each of the four components, three surface proportion parameters (the fourth one not being free as they all sum to 1), and the molecular concentration of CH_4 . Solving this is computationally expensive but still doable with a curve-fitting algorithm; Protopapa et al. (2017) use the Levenberg-Marquardt algorithm, which is reasonably robust but converges to only a local minimum.

The analysis of Protopapa et al. (2017) allows them to establish global zoning of Pluto. However, the simplified one-layer model fails to obtain accurate results notably in Sputnik Planitia, where a single-layer model with spatial mixing of N_2 - and CH_4 -rich components at thermal equilibrium cannot reproduce the N_2 band at appropriate temperatures. As Protopapa et al. suggest in their discussion, vertical layering may allow for all the conditions to be met and for the spectrum to be accurately reproduced. The goal of this thesis chapter is to set out the problem of spectral inversion with a complex multilayer model with multiple mixing modes, to study the ruggedness of the parameter space and multiplicity of solutions, and to determine the computational cost and feasibility of running this process globally for Pluto.

1.2 The problem of high-dimensional optimisation

In order to accurately represent the variability of Pluto's surface and construct a realistic picture of what surfaces correspond to the spectra we observe, a complex surface representation with a very high number of independent parameters is required. Let us consider the basic materials that have been identified on the surface: N₂, CH₄, CO, and H₂O ice, and the dark red organic matter. The first three of these may be present either in a pure state, or in a solid solution with each other. We thus have the following components: H₂O ice; pure CH₄ ice; N₂:CH₄:CO solid solution, which may be saturated or unsaturated with CH₄; and red matter. These each have a set of independent parameters to serve as RTM input: grain size, scattering anisotropy, molecular proportion. They also all may further be mixed granularly or spatially, or stratified. This introduces the parameters of granular proportion, spatial proportion, and layer thickness. If we allow all six components to be present simultaneously in any combination of spatial, stratified, and granular mixtures, we arrive at a dimensionality of more than five hundred, which is a completely intractable problem. If we limit the number of elements in a given mixture to two (i.e. there can be two spatial components, which can each have two vertical layers, which in turn can have two types of granular material), we still have a 35-dimensional parameter space to search, with some of the parameters varying between 0 and 1 and others by many orders of magnitude (e.g. the grain size, which can be anywhere between several nanometers and several meters).

If the parameter space were perfectly smooth and the error always sloped down towards one particular parameter set for a given spectrum, the high dimensionality of the search space would not be a concern. Unfortunately, in our case the parameter space is quite rugged — as briefly shown in Chapter I, the effects of different mixing modes and granular properties of these different compounds are complex and interact in non-obvious ways, and changing different parameters can modulate the shape of the spectrum in similar ways. This may lead to a multiplicity of equally valid solutions for a given spectrum. This is a problem that can potentially be mitigated in different ways, but may also be intractable, and may mean that there is no way to quantitatively determine the structure of the surface with a high enough degree of accuracy and precision.

In this chapter, we will first establish the limits of the inversion algorithm in terms of fit accuracy and efficiency, using synthetic high-resolution spectra generated by the same radiative transfer model that we are using for the inversion. We will then introduce complications such as noise and downsampling in order to more closely approximate the quality of the New Horizons data. Having set bounds on the precision and computing requirements of the fitting algorithm, we will then apply it to relatively simple New Horizons datapoints ("endmembers").

2 Methodology

Continuing from Chapter I's introduction to radiative transfer models, we will describe the structure and capabilities of two complex RTMs capable of supporting multiple layers and mixing modes, and then describe the problem-solving algorithms required to efficiently use these RTMs for spectral inversion.

2.1 Radiative transfer models

During this thesis we worked principally with two radiative transfer models, Spectrimag and DISORT. Spectrimag was the initial choice, as it is less resource-intensive than DISORT while also allowing for the different mixing modes we wanted to study. Unfortunately, numerical instabilities and undocumented modifications to the code made the model unusable. For posterity, we still provide a brief description of Spectrimag in the following section, but all of the quantitative results in this chapter were computed with DISORT.

2.1.1 Spectrimag

Spectrimag is a bidirectional reflectance RTM developed by S. Douté and described in Douté and Schmitt (1998). It is an improvement on the Hapke model, using a similar semi-empirical approach but allowing also for multiple scattering within a single layer. The model also allows for a two-layered surface, which originally used its own adding method but makes a call to DISORT in the version available to us. In theory the Spectrimag model allows for accurate spectral modeling while being much faster than DISORT for the one-layer scenario. In practice, however, some of the mixture modes were nonfunctional and the model in its current state did not appear to handle anisotropy in a manner consistent with its original published benchmarks. Due to poor archiving and versioning we were unable to procure a version of Spectrimag that would produce results corresponding to the benchmarks, or that would handle the mixture types we needed to simulate, and so we had to abandon it in favour of using DISORT for all calculations in late 2019.

2.1.2 DISORT/CDISORT

DISORT is a discrete-ordinate radiative transfer model originally written in Fortran 77 and rewritten in C by T. Dowling as detailed in Buras et al. (2011). It is designed for work with vertically inhomogeneous layered media, and aims to provide robust and accurate computation at the expense of computational speed. It is frequently used as a benchmark to evaluate the accuracy of other radiative transfer models, and so we opted to use it directly rather than search for other RTMs which would be faster but may be less numerically stable.

As described in Chapter I, the discrete ordinate approximation assumes a set of discrete values for the spatial and spectral dimensions in use. We are also using the plane parallel approximation. In reality the planetary surface we are observing is the surface of a spheroid, and so has a slight curvature, but we assume that the surface-atmosphere boundary and all other potential layer boundaries are parallel planes, and only consider the vertical dimension.

DISORT's implementation of the discrete ordinate approximation, and its solution of the differential equation system, is well-described in the DISORT report (Stamnes et al., 2000), and will not be reiterated here. We reproduce, as Figure III.1, the schematic from the report that illustrates the layered structure of the model and the input parameters that it uses to calculate the reflectance. The figure shows that each layer has as parameters its single-scattering albedo ω and phase function *P*, while the optical depth τ and temperature *T* are defined at the layer boundaries. We neglect thermal emission calculations in our usage of the model, as we are only working in the spectral range of the reflected solar component, so only ω , *P*, and τ are applicable.



Figure III.1: Schematic illustration of a multilayered optical medium as modeled by DISORT, with an incident beam and diffused intensity at the top boundary, and reflection and thermal emission at the bottom boundary. Cumulative optical depth τ , temperature *T* (neglected in this thesis), intensity, and flux are defined at layer interfaces while single-scattering albedo ω and phase function *P* are defined as layer averages (reproduced from Stamnes et al., 2000).

Computing the input parameters

The only parameters that we input to DISORT are the number of layers, ω and τ for each layer, and the viewing geometry (i.e. angles of incidence and emergence and the azimuthal angle). To do this, we first need to derive ω and τ from the parameters of our surface representation: layer

thickness, grain size, proportion of different components, their mixture modes, etc. As we established in Equation I.11 of Chapter I, the single scattering albedo corresponds to the ratio of the scattering and extinction effective cross-sections:

$$\omega = \frac{\sigma_S}{\sigma_E}$$

Since the particles comprising the surface are closely packed together, we are assuming that any light that gets scattered by a particle will be fully extinguished before it can re-emerge, i.e. we are neglecting non-extinctive diffraction.

In our study case, the particles comprising the surface are closely packed together and relatively large compared to the light wavelength. As per Hapke (2012, p. 417), this means we can assume extinction efficiency $Q_E = 1$ and assimilate σ_E to the geometric cross-section of the particle, with d being the grain diameter, as shown in Equation III.1. The scattering cross-section is defined similarly (Equation III.2), with a multiplicative factor of Q_S , which is the scattering efficiency as defined in Equation I.10.

$$\sigma_E = \frac{\pi d^2}{4} \tag{III.1} \qquad \sigma_S = \frac{Q_S \pi d^2}{4} \tag{III.2}$$

 Q_S can be computed as shown in Equation III.3, with S_e and S_i referring respectively to the fraction of light reflected from the particle's surface externally, and the fraction reflected internally. They, in turn, can be computed by integrating the Fresnel equations (Eq. III.4). R_{\perp} and $R_{//}$ are the Fresnel coefficients, for which we follow Douté (1998) in using analytical expressions given in Mahan (1956).

$$Q_{S} = S_{e} + (1 - S_{e}) \frac{(1 - S_{i})\theta}{1 - S_{i}\theta}$$
(III.3)
$$S_{e} = \int_{0}^{\frac{\pi}{2}} (R_{\perp}(i) + R_{//}(i)) \cos i \sin i \, di$$
(III.4)
$$S_{i} = \int_{0}^{\frac{\pi}{2}} (R'_{\perp}(i) + R'_{//}(i)) \cos i \sin i \, di$$

This lets us calculate $\omega = \sigma_s/\sigma_E$. We also compute τ by integrating Equation I.17 and obtain the expression given in Equation III.5, with *N* being the granular medium's number density, γ its compacity, and *h* the layer thickness.

$$\tau = -\frac{N\sigma_E \ln(1-\gamma)h}{\gamma} \qquad (III.5) \qquad \qquad N = \frac{6\gamma}{\pi d^3} \qquad (III.6)$$

2.2 Metaheuristics

A heuristic is a problem-solving algorithm designed to work with limited time and knowledge, in situations where finding the optimal solution would take a prohibitively long amount of time. It arrives at a approximate but sufficiently good solution in a reasonable amount of time. The simplest, most fundamental heuristic is trial and error. For a more complex example, consider the well-known travelling salesman problem:

Given a list of cities and the distances between each pair of cities, what is the shortest possible route that visits each city and returns to the origin city?

This is an NP-hard problem¹. The time required to find the optimal solution to travelling salesman via an exhaustive search is within a polynomial factor of O(n!), the factorial of the number of

 $T(n)\in O(n^k)$

¹A definition from computational complexity theory: an NP-hard problem is a problem that is at least as hard as every problem in the NP (non-deterministic polynomial) class of problems, and is likely not resolvable in polynomial time. A problem resolvable in polynomial time will have an upper bound on its running time of a polynomial expression relative to the size of the input:

cities, so it quickly becomes impractical. A simple heuristic that can find a good solution for randomly distributed cities on a Euclidean plane is the nearest-neighbour algorithm, which on average returns a path that is 25% longer than the shortest possible path but reduces the complexity to $O(n^2)$.

The disadvantage of using a fixed heuristic is that it will generally be designed to run optimally on a particular dataset, and the assumptions it makes in order to do that mean that it may not be applicable to other datasets and situations. For complex problems where we do not have sufficient information to manually design a good heuristic, we can turn to **metaheuristics**. A metaheuristic is a higher-level heuristic that is designed to *search for* a heuristic that will efficiently optimise the problem and find a good enough global solution in a reasonable runtime. It makes relatively few assumptions about the dataset and the shape of the search space, and frequently uses *stochastic optimisation* (generating random values to test against).

2.2.1 Simulated annealing: definition, algorithm, approximations

Simulated annealing (SA) is a commonly-used optimisation metaheuristic, where a local search for a better solution is combined with a probability function that allows moving to a worse solution. This prevents the algorithm from converging to a local minimum that isn't the best global solution. A simplified schematic of how simulated annealing works is shown in Figure III.2.



Figure III.2: A schematic of how simulated annealing optimisation finds the global solution for a one-dimensional function f(x) (modified from Ghasemalizadeh et al., 2016).

The name comes from the technique of annealing in metallurgy. Annealing involves heating a material above its crystallisation temperature, then slowly cooling it. This allows enough time for the particles to arrange themselves in a minimal-energy state, which produces larger crystals and reduces defects in the material. In simulated annealing, the probabilistic element is controlled by a "temperature" variable, which slowly decreases over time.

Let us consider the simplified case of optimising a spectral inversion for just one wavelength, where we look at the measured reflectance for that wavelength and try to determine the surface that would produce this reflectance. The simulated annealing algorithm has a target reflectance value, a set of parameters that are input into an RTM, and the reflectance computed by the RTM. In a given iteration, the algorithm has an initial set of parameters $\{x_0\}$, to which it introduces a random perturbation to obtain a parameter set $\{x\}$. It then obtains the corresponding reflectance $R\{x\}$ from the RTM, computes the error between $R\{x\}$ and the target reflectance, and checks how it differs from the error corresponding to the previously computed reflectance $R\{x_0\}$. If the new value is closer to the goal than the old value, then it is selected as the value to keep, and the algorithm continues to the next iteration. If this is not the case, then the algorithm may still select the new value with a probability A controlled by the temperature parameter T, and also continue to the next iteration. As T decreases with iterations, the probability of selecting a locally worse solution also decreases.

Figure III.3 shows a flowchart summarising this process.



Figure III.3: A flowchart showing how the simulated annealing algorithm applies to a spectral inversion problem.

2.3 Algorithm description and radiative transfer model setup

2.3.1 Scales of parameter effect

The basic form of the simulated annealing algorithm applies a random perturbation to all the parameters in every step. If they all had a linear effect on the spectrum, this would be reasonable. However, different parameters act on it on different scales: for instance, the anisotropy factor moreor-less uniformly affects the whole spectrum, "squeezing" or "stretching" its intensity, whereas the grain size of a minor component may only affect the shape of a specific narrow band. This means that perturbing all the parameters at once is inefficient: if we change a large-scale parameter in a way that worsens the fit, an improvement on one narrow band will be lost in the overall error increase. There are multiple solutions to this. First, we can separate the parameters into subsets by the size of their effect, only varying parameters of similar effect scale at the same time. Second, we can compute the error on the derivative of the spectrum, which lets us evaluate changes between wavelengths and the overall shape of the spectrum, neglecting large-scale uniform changes like those induced by the anisotropy of the scattering function. In our working algorithm, we implement both of these solutions, with the algorithm phases proceeding as follows:

- 1. **Shape fit:** we perturb the full set of parameters, and compute the root mean square error (RMSE) between the derivative of the goal spectrum and that of the computed spectrum.
- 2. **Amplitude fit:** We only perturb the anisotropy factor ξ , and compute the RMSE directly on the spectrum.
- 3. **Full fit:** We perturb the parameters other than the anisotropy factor, and compute the RMSE on the spectrum.

This sequence allows us to fit the large-scale and small-scale parameters separately, without the former overwhelming the latter. Section 3.1.1 describes how we determined the optimal relative phase lengths.

We also performed a number of test runs with further granularity for specific surface compositions, where the error was calculated for specific bands and parameters were further separated by their effect in these bands, but this was not widely applied due to lack of time to specifically calibrate it to each scenario.

2.3.2 Surface representation

As described previously, Schmitt et al. (2017a) and others have established the presence on Pluto's surface of the following major components:

- H₂O ice;
- pure CH₄ ice;
- a solid solution of N₂, CH₄, and CO ices, which can be N₂-rich or CH₄-rich;
- dark red organic material (hereafter referred to as RM).

Randomly mixing all these components in all possible mixture modes is computationally prohibitive and inefficient, as many of them are not likely to occur together, or not likely to be mixed in particular ways. In order to reduce the dimensionality of the problem, we will define several types of terrain, each of which have certain components present and not others, in mixture modes that qualitatively correspond to the data. These are:

- 1. pure CH₄ ice;
- 2. N₂-rich N₂:CH₄:CO solid solution covered by an optically thin layer of pure CH₄ ice;
- stratified layers of N₂:CH₄:CO solid solution, corresponding to a vertical grain size or molecular proportion gradient;
- 4. H₂O ice in a spatial or granular mixture with RM.

Terrains 1 through 3 may also have a small amount of H_2O ice added in a spatial mixture.

2.3.3 Optical constants

As noted in the previous section, we need the optical constants *n* and *k* to compute the DISORT input parameters. For the different ices present on Pluto's surface, we have laboratory measurements of these optical constants in the appropriate temperature conditions via the SSHADE database (Schmitt et al., 2017b). The measurements were carried out over a wide range of temperatures (including the range of 39–41K which are the optical constants we use for our model) with a Nicolet 800 spectrometer, of which the properties are detailed in Quirico and Schmitt (1997b). Table III.1 lists the materials and temperatures corresponding to the optical constants we use, and a list of references for the datasets is included in the bibliography. The optical constants of pure CH₄ ice were characterised in Grundy et al. (2002), those of N₂ ice in Grundy et al. (1993), and those of H₂O ice in Grundy and Schmitt (1998). CO dissolved in N₂ and CH₄ dissolved in N₂ were characterised respectively in Quirico and Schmitt (1997a) and Quirico and Schmitt (1997b).

| Material | Temperature | SSHADE dataset |
|---|-------------|--------------------------|
| pure crystalline CH ₄ -I | 39K | SPECTRUM_BS_20130114_004 |
| CH_4 in solid solution in β -N ₂ crystal | 41K | SPECTRUM_BS_20130103_005 |
| CO in solid solution in β -N ₂ crystal | 41K | SPECTRUM_BS_20130103_015 |
| β -N ₂ crystal | 41K | SPECTRUM_BS_20120925_005 |
| crystalline H ₂ O-Ih | 40K | SPECTRUM_BS_20120924_014 |



We do not currently have laboratory measurements that give an adequate fit for the RM. Fayolle et al. (2021) attempts to use synthetic Pluto tholins as an analogue for the Pluto RM, but the fit is not sufficiently good to be used in our model, so we have no physical analogue to this component. A stopgap experimental solution we used was inverting the observed RM spectra to obtain synthetic optical constants that could be used as a "component" in the radiative transfer model, but as of the writing of this manuscript this is not yet integrated into the model and fitting algorithm.

2.3.4 Deresolution

The radiative transfer model produces spectra at the spectral resolution of the optical constant files. For the optical constants we're using, the sampling resolution is of 0.482 cm⁻¹ in terms of wavenumber. We restrict the computation to the wavenumber range of 3500–8500 cm⁻¹, which corresponds to a wavelength range of 1.176–2.857 μ m, and a corresponding spectral resolution range of 0.0667–0.393 nm. This gives us a resolving power of between 7269 and 17631. As mentioned in Chapter I, the spectral resolution of the LEISA instrument is much lower, with an average resolving power of 240. This means we need to downsample the synthetic spectra in order to be able to compare them with the LEISA data.

To correctly downsample the spectrum, we need to convolve it with a spectral response function similar to that of LEISA, which is approximately Gaussian. As seen in section 1.2 of Chapter I, the spectral resolution of a hyperspectral imager is defined as the full width half maximum (FWHM) of the Gaussian. The Gaussian density function has the following form as a function of the standard deviation σ and the mean μ :

$$f(\lambda) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(\lambda-\mu)^2}{2\sigma^2}\right]$$
(III.7)

For the FWHM, at the half-maximum points λ_0 , this becomes:

$$\exp\left[-\frac{(\lambda_0 - \mu)^2}{2\sigma^2}\right] = \frac{1}{2} = 2^{-1} \implies \lambda_0 = \pm \sigma \sqrt{2\ln 2} + \mu$$
(III.8)

From this, we derive the relation $FWHM = 2\sqrt{2 \ln 2\sigma} \approx 2.35\sigma$. We can then construct a Gaussian curve that corresponds to the LEISA spectral resolution, and convolve it with the synthetic spectrum to produce a spectrum that is directly comparable with the LEISA data (see Figure III.4).



Figure III.4: An example of a high-resolution randomly generated spectrum, compared with the same spectrum convolved to the New Horizons spectral resolution.

Downsampling of optical constants

As the spectrum calculation is done independently for each wavelength, the computation time is directly proportional to the resolving power of the optical constants. The resolving power of the LEISA data being significantly lower than that of the optical constant data gave us an opportunity to save computational time via a preliminary downsampling step. We reduce the initial optical constant sampling step by a factor of four, going from 0.482 cm⁻¹ to 1.929 cm⁻¹, and also reduce their resolution by the same factor by convolving them with a Gaussian function with a FWHM of 3.86 cm⁻¹. This reduces computation time by a factor of four, has no appreciable effect on fit accuracy in all tests, and is not discernible by eye at the scale of our typical figure output, as the optical constant curve goes from 10371 to 2593 points.

2.4 The CIMENT grid

All the calculations in this section were carried out on a cluster of the CIMENT² grid, itself a part of the GRICAD³ CNRS support unit which centralises the high-power computing resources available to researchers in the Grenoble area. We used the Froggy cluster, a BullX DLC supercomputer designed for running many parallel short-term computing tasks. Froggy has 190 nodes with 2 Intel Sandy Bridge EP E5-2670 8-core processors in each node and 64GB of RAM per node. Further information about the computing environment can be found in the GRICAD documentation at https://gricad-doc.univ-grenoble-alpes.fr.

²Calcul Intensif/ Modélisation/ Expérimentation Numérique et Technologique — High-Performance ComputIng/ Modéling/ Numerical and Technological experimentation.

³Grenoble Alpes Recherche - Infrastructure de Calcul Intensif et de Données — Grenoble Alpes Research Scientific Computing and Data Infrastructure.

3 Results

3.1 Calibration on synthetic data

Our ability to fit synthetic spectra to real data depends on the following: 1) whether or not there exists a set of parameters with which our chosen surface representation will produce a spectrum that matches the data; 2) whether or not our fitting algorithm can efficiently find this set of parameters. We cannot evaluate the first condition in advance, but only try and succeed or fail. We can, however, evaluate the second condition, by generating random synthetic spectra and trying to fit them using the optimisation algorithm.



Figure III.5: Synthetic spectra used for model validation and calibration. Top: synthetic spectra, bottom: corresponding surface representations.

For this section, we use four random spectra with compositions of increasing complexity, shown in Figure III.5 along with schematics of their corresponding surface representations. The exact model parameters are detailed in Table III.2. First, we calibrate the fitting parameters such as phase length and annealing temperature, to ensure that the algorithm runs optimally (section 3.1.1). Then, we

| | component | spatial % | granular % | thickness | grain size | ξ | CH_4 % | CO % |
|---|-------------------------------------|-----------|------------|-----------|------------|--------|----------|------|
| Spectrum 1: pure CH ₄ ice | | | | | | | | |
| | CH_4 | 100% | 100% | ∞ | 0.255 mm | 0.066 | | |
| Spectrum 2: solid solution N ₂ :CH ₄ :CO ice | | | | | | | | |
| | N ₂ :CH ₄ :CO | 100% | 100% | ∞ | 442.6 mm | -0.218 | 0.4% | 3.4% |
| Spectrum 3: CH ₄ stratified over N ₂ :CH ₄ :CO | | | | | | | | |
| L1 | CH_4 | 100% | 100% | 0.296 mm | 0.135 mm | 0.157 | | |
| L2 | N ₂ :CH ₄ :CO | 100% | 100% | ∞ | 732.9 mm | -0.097 | 4.4% | 1.1% |
| Spectrum 4: CH_4 stratified over granular [N_2 : CH_4 : CO , CH_4], with spatial H_2O | | | | | | | | |
| Al | H ₂ O | 34% | 100% | ∞ | 3.51 mm | -0.648 | | |
| A2:L1 | CH_4 | 66% | 100% | 0.1 mm | 0.013 mm | 0.506 | | |
| A2:L2 | N ₂ :CH ₄ :CO | 66% | 19.5% | m | 4.53 mm | 0.536 | 1.5% | 1.4% |
| | CH_4 | 0070 | 80.5% | | 0.8 mm | 0.536 | | |

evaluate the algorithm's efficiency and precision limits, first by computing best fits for clean highresolution spectra, then adding noise and downsampling.

Table III.2: Surface structure and composition for the four synthetic spectra used for model validation and calibration. L1 and L2 refer to stratified layers; A1 and A2 refer to horizontally spatially arranged areas. ξ is the anisotropy factor used for the layer's phase function.

3.1.1 Algorithm phase lengths

As described in section 2.3.1, our algorithm runs in three phases: (1) shape fit, (2) amplitude fit, and (3) full fit. To determine the optimal lengths of these phases, we run a number of fits on the four synthetic spectra with different sets of phase lengths, and compute the error distributions for each.



Figure III.6: Best RMSE for runs of the three-phase algorithm for different relative phase lengths, for a pure CH₄ ice spectrum (in blue) and for a N₂:CH₄:CO solid solution (in green). The values shown are medians over 40 5000-iteration runs, with error bars showing quartiles. The simulated annealing parameters used are T = 400 and p = 0.2.

Figure III.6 shows the median RMSE for different relative lengths of the three phases, with the

split between first and second phase and between the second and third phase placed respectively at 20 and 40%, 30 and 40%, 40 and 50%, 40 and 60%, 60 and 70%, 60 and 90%, and 80 and 90% of the total runtime. These are computed over 40 runs with 5000 iterations per run, for the two simplest calibration spectra: pure CH_4 and $N_2:CH_4:CO$ solid solution.

We can see that the robustness of the fit as a function of the phase lengths is not consistent between the two spectra: while a 60/90% runtime split produces the lowest-error results for the pure CH₄ spectrum, for the solid solution spectrum the most accurate fit is obtained with a 40/50% split. We consider the optimal parameters for the more complex surface representation to be of higher importance, as this configuration is likely to remain more robust for even more complex surface representations and for fitting real data. We therefore perform all subsequent computations with the three phases running between 0-40%, 40-50%, and 50-100% of the runtime respectively.

3.1.2 Calibrating simulated annealing parameters

The next calibration step is to determine the optimal values of the two parameters that control the simulated annealing algorithm: the temperature T and the perturbation p. We run 40 sets of 5000-iteration fits for each of the four calibration spectra. We compute the RMSE distribution for these fits, and we also compute a synthetic value which we call the *error score*.

The error score is a way to evaluate the specificity of the best solution: for the more complex spectra, it is possible that different combinations of parameters will result in fits which are very close to the target spectrum but for which the physical parameters do not correspond to the target surface. To determine the error score, we compare the different surface parameters (anisotropy, area fraction, layer thickness, grain size, granular proportion, and molecular proportion) between the target and best fit composition, and calculate a weighted sum where the weight is the relative size of the part of the surface area considered.

Figure III.7 shows the distributions of the best RMSE for the different *T* and *p* values tested, and Figure III.8 shows the distributions of the corresponding error scores. As with the phase length calibrations, we can see that the relative robustness of the different (T, p) parameter sets varies somewhat with the complexity of the spectrum. We find that the parameter set (T = 100, p = 0.05) is the most stable in terms of returning a low RMSE and a low error score across all the spectra. We therefore perform all subsequent computations with this parameter set.



Figure III.7: Best RMSE for runs of the simulated annealing algorithm for different values of *T* and *p*, with *T* between 0 and 800 and *p* between 0.01 and 0.4. The values shown are medians over 40 5000-iteration runs, with error bars showing quartiles.



Figure III.8: Error score values for runs of the simulated annealing algorithm for different values of *T* and *p*, with *T* between 0 and 800 and *p* between 0.01 and 0.4. The values shown are medians over 40 5000-iteration runs, with error bars showing quartiles. Spectra 1 and 2 are shown on a log scale, while spectra 3 and 4 are shown on a linear scale, to facilitate distinguishing the different cases.

3.1.3 Run duration

During the parameter calibration testing, we note that we do not achieve fits with a better RMSE than a few percent for the two more complex spectra. This could be due to either insufficient runtime or inherent inability of the problem-solving algorithm to efficiently solve a problem of very high dimensionality. We can test this by running the algorithm for varying numbers of iterations.

Figure III.9 shows the median RMSE and error distribution for a range of iterations between 100 and 20000, the interpretation of which is nontrivial. Figure III.10 shows the error score medians and distributions, the behavior of which is quite similar to that of the RMSE and so will not be discussed separately, but is presented for completeness.

Up to 5000 iterations, increasing the number of iterations improves the median RMSE for all spectra. For spectrum 1, the simplest case, we achieve a median RMSE of below 0.1% for 500+ iterations and appear to reliably reach convergence at 1000+ iterations, with median RMSE remaining at around 0.03% thereonwards. For spectrum 2, the median RMSE drops below 1% for 500+ iterations, and continues to slowly drop with increasing iterations, reaching 0.38% for 20000 iterations.

For the more complex spectra 3 and 4, while increasing iterations to several thousand achieves better fits than with lower values, the median RMSE seems to plateau at about 4 and 5% respectively for these spectra regardless of number of iterations. This suggests that the fitting algorithm cannot obtain reliable convergence beyond several percent of precision for spectra of this complexity, regardless of runtime. This does not mean that closer fits cannot be obtained at all — the best RMSE obtained for these spectra across all testing phases are 0.6% for spectrum 3 and 1.2% for spectrum 4. However, this quality of result cannot be guaranteed and may require hundreds of runs for one spectrum.

The general theory of simulated annealing states that with a sufficiently slowly decreasing temperature function convergence is guaranteed (Granville et al., 1994). However, without highly detailed knowledge of the function space⁴ it is not possible to give a bound on the needed temperature. As such, the most promising method we have is to launch multiple parallel runs to explore different

⁴Ideally the rugosity of this function space could be characterised using complex analysis of the DISORT model. However, the mathematics behind the model do not lend themselves to this kind of analysis and this characterisation would require years of professional mathematical work.




Figure III.9: Best RMSE for runs of the simulated annealing algorithm for different numbers of iterations, ranging between 100 and 20000. The values shown are medians over 20 runs, with error bars showing quartiles.



Figure III.10: Error score values for runs of the simulated annealing algorithm for different numbers of iterations, ranging between 100 and 20000. The values shown are medians over 20 runs, with error bars showing quartiles. Spectra 1 and 2 are shown on a log scale, while spectra 3 and 4 are shown on a linear scale, to facilitate distinguishing the different cases.

3.1.4 Testing with noise

Real hyperspectral data, such as that from the New Horizons mission, has a certain unavoidable amount of instrumental and environmental noise. We therefore need to verify whether our fitting algorithm works as well with noisy spectra as it does with the clean synthetic ones. To do this, we apply a Gaussian noise to spectra 1 and 2, and try to fit them with the established model parameters. We test both additive and multiplicative noise, with amplitudes of 1% and 5%. Figure III.11 shows the median RMSE and error scores for 20 runs of 5000 iterations each for each noise case. Figure III.12 shows example best fits for each noise case. We show only spectrum 1 for the additive noise and only spectrum 2 for the multiplicative noise, as the results are very similar for both.

Before interpreting the RMSE results, we add a corrective factor. Due to the noise being Gaussian and independent from wavelength, there is an inherent noise "floor" that we need to account for. We compute this by calculating the RMSE of a noisy spectrum with itself, and subtracting it from the raw RMSE values. This is approximately 1% and 5% for the additive noise scenario for both spectra, and a little less for the multiplicative noise: 0.77% and 3.75% for spectrum 1 and 0.65% and 3.25% for spectrum 2 respectively for the 1% and 5% percent noise cases.

For the 1% noise case, median corrected RMSE remains at or below 2% for both spectra. For the 5% multiplicative case, the median RMSE is below 5%. For additive noise with 5% amplitude, the RMSE rises above 10%. However, we also observe that the error score remains quite low for all cases — < 1 for all but the 5% additive noise case for spectrum 1. Both this and the tight clustering of the median RMSE values suggests that the fit remains relatively accurate and the algorithm works well even for these high-noise cases.



Figure III.11: RMSE and error score values for runs of the simulated annealing algorithm with different amounts of Gaussian noise applied to synthetic spectra 1 and 2. Values shown are medians over 20 runs of 5000 iterations each. Error bars show quartiles for the error score values, and have been omitted for RMSE, as they are too small to see clearly. The RMSE values are shown for both the raw RMSE and RMSE corrected for the noise floor.



Figure III.12: Example best fits generated in testing with noisy spectra, shown on spectrum 1 for the 1% and 5% additive noise cases, and on spectrum 2 for the 1% and 5% multiplicative noise cases. In green are the goal spectra with added noise, and in dark blue are the best fits.

3.1.5 Fitting a simple spectrum with a complex surface representation

In the previous tests, the fitting is run with the same surface representation as the spectrum was generated with, i.e. the same number of components, layers, areas, etc. This is done because computational time increases with complexity, roughly doubling between spectrum 1 (1.9 hours for 5000 iterations) and spectrum 4 (4.3 hours for 5000 iterations). With tens of runs per calibration case, only running as complex a case as necessary saves many days of computing time.

However, a robust model will be able to fit a simple spectrum (i.e. spectrum 1) with a complex surface, reducing the unnecessary components to a negligible fraction. For instance, fitting spectrum 1 (pure CH_4 ice) with the surface model of spectrum 4 will ideally result in a surface where the fraction of the Al area (H_2O ice in a spatial mixture) will be negligible, and the A2:L1 layer (pure CH_4) will be optically thick. We test this exact case and obtain good results, with a median of 0.7% RMSE (QI = 0.5%, Q3 = 1.2%) over 10 runs, and the superfluous surface components reduced to very small proportions as expected.

Figure III.13 shows the best fit and Table III.3 lists its exact surface parameters: we can see that the spatial H_2O component is at 1.5% and the CH_4 layer is approximately 27 grains thick, which makes it optically thick and the layer below not visible in the spectrum. The grain size of the CH_4 layer is exactly that of the target spectrum, and while the anisotropy parameter ξ diverges slightly, the surface remains slightly forward-scattering.



Figure III.13: Best fit of synthetic spectrum 1 using the more complex surface representation of spectrum 4. Table III.3 lists the model parameters.

| | component | spatial % | granular % | thickness | grain size | ξ | CH ₄ % | CO % |
|-------------------|-------------------------------------|------------------|------------|-----------|------------|--------|-------------------|-------|
| Target spectrum 1 | | | | | | | | |
| | CH_4 | 100% | 100% | 8 | 0.255 mm | 0.066 | | |
| Best fit | t | | | | | | | |
| Al | H ₂ O | 1.5% | 100% | 8 | 15 399 mm | -0.516 | | |
| A2:L1 | CH_4 | 98.5% | 100% | 6.979 mm | 0.255 mm | 0.137 | | |
| 12.12 | N ₂ :CH ₄ :CO | 98 5% | 10.5% | m | 1.258 mm | 0.619 | 0.17% | 0.01% |
| A2.L2 | CH_4 | J J .J /0 | 89.5% | ~ | 0.023 mm | 0.619 | | |

Table III.3: Best fit of synthetic spectrum 1 using the more complex surface representation of spectrum 4. Figure III.13 shows the corresponding spectrum.

3.1.6 Component selection

It is possible to run the fitting algorithm with a maximally complex surface and allow the surface components which do not correspond to any spectral characteristics to be reduced to negligible proportions by the optimisation. In practice, however, this not only corresponds to a massive inflation of required computing time, but also makes the fitting problem very hard to solve. As we saw with the parameter calibration on the synthetic spectrum 4, fitting which corresponds to a 20-dimensional problem, the median RMSE plateaus at several percent, and cannot be deterministically improved on. Using a model with all possible components present in all of their possible mixture types would amount to solving at least a 35-dimensional problem, and getting an accurate fit would be extremely hard if not impossible.

With this in mind, we need to determine a way to reduce the dimensionality of the problem by selecting appropriate components for each spectrum we are fitting. There are multiple ways to do this. First, we can do the reduction manually by qualitatively studying the spectrum and seeing what spectral characteristics are present. This is doable for a small number of datapoints, but is prohibitive for the millions of pixels that comprise the global map of Pluto. Secondly, we can provide the algorithm with a number of different possible surface representations, run a relatively short fit with these possibilities, and proceed with the one which has the lowest error. We implemented this at an interim stage of the thesis work, with a simplified Spectrimagderived model which only allowed for a granular mix of two out of the six possible components. Results were satisfactory, with the algorithm selecting the correct pair of components out of all the possible pairs. We have not, however, yet implemented this with the more complex surface modeling which uses DISORT for the computational backend and allows for all the different mixture types including stratification.

An automatic selection algorithm for this complex surface model would need to be iterative, selecting first a subset of components, then trying different mixture types of these components to see which was more appropriate. It would also require multiple runs per each test case; as we saw in the parameter calibration, for complex enough spectra the closeness of the fit cannot be deterministically guaranteed, and running just one fit for each test case has a high chance of not being representative. With all of this in mind, it may be that a naive selection algorithm is just as computationally intensive as trying to fit a fully complex surface model. For further theoretical discussion of how to make the global mapping of Pluto practical and of what automation techniques may be useful, see Chapter IV.

3.2 Compositional endmembers



Figure III.14: Location of the two LEISA endmembers, as well as those of the endpoints of the compositional gradient profile studied in section 3.3. These are shown on an orthographic projection of the LORRI panchromatic surface albedo map, centered at 60°N, 150°E.

Having determined an optimal set of parameters for the simulated annealing algorithm, and having established bounds on its precision using synthetic spectra, we proceed to working with the New Horizons LEISA data. We select several groups of pixels which correspond to compositional "endmembers", suggested by a qualitative analysis to have a relatively simple surface structure with only one or two components present. Due to time constraints, we were only able to finish working with two of these endmembers: "North Pole" and "Sputnik". Figure III.14 shows these areas' ap-

proximate location on Pluto's surface, and Table III.4 lists their exact coordinates and observational geometry. After fitting these endmembers, we also study a profile of 14 points along a progressive transition between 2 endpoints near the North Pole endmember in hopes of quantifying a compositional and textural gradient; this is detailed in the subsequent section. The spatial location and the illumination-observation geometry of the two extreme endpoints of the profile are also listed in Table III.4.

| Name | Latitude (°) | Longitude (°) | i (°) | e (°) | φ (°) | g (°) |
|--------------------------------|--------------|---------------|-------|--------------|--------------|--------------|
| Sputnik | 26.51 | 192.78 | 50.86 | 31.30 | 14.82 | 21.75 |
| North Pole | 69.98 | 63.30 | 36.02 | 56.32 | 13.25 | 22.36 |
| Profile: CH ₄ -rich | 65.50 | 152.21 | 16.96 | 27.44 | 54.87 | 22.15 |
| Profile: N ₂ -rich | 65.26 | 156.20 | 18.03 | 26.99 | 55.80 | 22.14 |

Table III.4: Location and observation geometry of the two LEISA endmembers, as well as those of the endpoints of the compositional gradient profile studied in section 3.3. The observational geometry angles are: the incidence angle *i*, the emergence angle *e*, the azimuth φ and the phase angle *q*.

3.2.1 North Pole

The North Pole endmember is located at 69.98°N and 63.3°E, in the mid-latitude northern highlands, which we had previously classified as very coarse-grained N₂-rich ice with high CH₄ content. In our initial work with Spectrimag, we attempted to fit it with a simple one-component pure CH₄ or N₂:CH₄:CO ice surface. This did not result in an adequate fit on all CH₄ bands with the same composition, and so we fitted the three principal groups of bands separately, as shown in Figure III.15. The three different spectra that fit the bands centered around 1.35, 1.7 and 2.3 μ m, with each spectrum corresponding to a granular mixture of pure CH₄ and molecular N₂:CH₄, form a compositional gradient with the percentage of CH₄ dissolved into the N₂ ice increasing with wavelength (see Table III.5 for details).



Figure III.15: Best fit spectra for each CH_4 band group with the single-layer Spectrimag model for the North Pole endmember. The surface representation is a granular mix of pure CH_4 ice and $N_2:CH_4$ molecular mix. Table III.5 shows the exact surface parameters.

| band group | 1.35 μm | | | 1.7 μm | 2.3 μm | | |
|------------|---------|--|--------|---|--------|--|--|
| component | CH_4 | N ₂ : C H ₄ (0.1%) | CH_4 | N ₂ : CH ₄ (0.8%) | CH_4 | N ₂ :CH ₄ (1.8%) | |
| proportion | 95.8% | 4.2% | 6.2% | 93.8% | 33.4% | 66.6% | |
| d | 5 mm | 153 μm | 5 mm | 5.3 cm | 1.3 mm | 3.6 mm | |
| ξ | -0.313 | -0.313 | -0.354 | -0.354 | 0.011 | 0.011 | |

Table III.5: Compositions for the best fit results for each CH_4 band group with the single-layer Spectrimag model for the North Pole endmember. The surface representation is a granular mix of pure CH_4 ice and N_2 : CH_4 molecular mix. All proportions are mass proportions; the proportion after N_2 : CH_4 is the proportion of CH_4 in N_2 . Figure III.15 shows the corresponding spectra.

The fact of different compositions being needed to match the different band groups suggests vertical stratification, as bands of different intensities correspond to reflectance from different substrate depths. Therefore, once we had access to a RTM that could include stratification (DISORT), we modeled this endmember with a thin layer of pure CH₄ on top of N₂:CH₄:CO substrate, and eventually added a minority spatial component of H₂O to dampen overall reflectance. Figure III.16 shows the lowest-error fit (RMSE=3.27%) overlaid on the North Pole target spectrum and a schematic of the corresponding surface model, and Table III.6 lists its properties. We see a near-anisotropic layer of pure CH₄ ice that is ~4–5 grains thick, with a grain size of 1.2 mm, atop N₂:CH₄:CO ice with 0.13% of CH₄ and 0.3% of CO. The spatial proportion of H₂O ice is ~6%, with very fine grains of 23 μ m.



Figure III.16: The best fit over 20 5000-iteration runs of the simulated annealing algorithm using the DISORT RTM for the North Pole endmember (top) and a schematic of its corresponding surface representation (bottom). The exact parameters are also given in Table III.6.

This relatively thick layer of pure CH_4 corresponds relatively well to the Schmitt et al. (2017a) interpretation of an optically thick crust of CH_4 ice having formed atop a substrate of N_2 -rich ice due to sublimation of N_2 . The presence of H_2O component is likely to correspond to a small exposed area of H_2O "bedrock", H_2O being much less volatile than the other ices. As Pluto's surface is a complex and dynamic environment, these components do not necessarily have to be in thermodynamic equilibrium, and may represent a transient state. The seasonal cycle of CH_4 as studied by Bertrand et al. (2019) situates us in the spring N_2 sublimation phase, which corresponds well to our observations of a pure CH_4 lag layer on the surface produced by compositional differentiation of CH_4 and the much more volatile N_2 , and CO.

| | component | spatial % | thickness | grain size | ξ | CH_4 % | CO % |
|-------|-------------------------------------|-----------|-----------|------------|--------|----------|-------|
| Al | H ₂ O | 5.9% | ∞ | 23 µm | -0.692 | | |
| A2:L1 | CH_4 | 94 1% | 5.76 mm | 1.19 mm | 0.012 | | |
| A2:L2 | N ₂ :CH ₄ :CO | 74.170 | ∞ | 0.12 mm | 0.106 | 0.13% | 0.33% |

Table III.6: The surface representation parameters for the best fit with the DISORT RTM for the North Pole endmember. There are two spatially distinct components (Al and A2), with A2 being composed of two vertical layers: pure CH_4 (A2:L1) and solid solution $N_2:CH_4:CO$ (A2:L2). The corresponding spectrum is shown in Figure III.16.

3.2.2 Sputnik Planitia

The Sputnik endmember, as the name suggests, is located in the vast nitrogen plain of Sputnik Planitia, in the low northern latitudes (26.51°N), and close to the sub-spacecraft point of the New Horizons encounter (192.78°E). In the surface composition of this endmember, we expect to see a majority component of nitrogen ice, with small amounts of CH_4 diluted in it.

In simple model testing, as with the North Pole endmember, we fail to obtain a good fit with a single N_2 :CH₄:CO ice component, and must add a thin layer of CH₄ ice above it. As with the North Pole endmember area, this can correspond to a layer left behind by sublimated N_2 or, conversely, may represent a seasonal condensed CH₄ frost. To represent the possibility of partial coverage by the thin CH₄ stratum, we allow the ice thickness to vary below 1 grain diameter, with i.e. a layer thickness of 0.5 grains having a physical interpretation of a 1-grain layer that covers 50% of the surface. We obtain a reasonably good fit with this surface representation, with the lowest RMSE at 4.57% over 20 5000-iteration runs of the model. Figure III.17 shows this best fit spectrum overlaid on the target Sputnik spectrum as well as a schematic of the corresponding surface representation. Table III.7 lists the properties of the surface.



Figure III.17: The best fit over 20 5000-iteration runs of the simulated annealing algorithm using the DISORT RTM for the Sputnik endmember (top) and a schematic of its corresponding surface representation (bottom). The sub-grain layer height of 1.3 mm is interpreted as partial (68%) coverage by a single-grain 1.9-mm thick layer. The exact parameters are also given in Table III.7.

| | component | thickness | grain size | ξ | CH ₄ % | CO % |
|----|-------------------------------------|-----------|------------|--------|-------------------|-------|
| L1 | CH_4 | 1.29 mm | 1.94 mm | -0.121 | | |
| L2 | N ₂ :CH ₄ :CO | ∞ | 47.47 mm | 0.8 | 0.036% | 0.02% |

Table III.7: The surface representation parameters for the best fit with the DISORT RTM for the Sputnik endmember. The surface components are vertically stratified, consisting of a thin pure CH_4 layer (L1) and a solid solution of N₂: CH_4 :CO (L2). The sub-grain layer height of 1.3 mm is interpreted as partial (68%) coverage by a single-grain 1.9-mm thick layer. The corresponding spectrum is shown in Figure III.17.

In summary, the surface consists of a very thin partial pure CH_4 layer with a grain size of about 2 mm, which is slightly retrodiffusive and has a coverage of about 68%. The thick underlying N₂-rich ice layer is very coarsely grained (d = 47.5 mm), and has a very low proportion of CH_4 (0.04%). This is consistent with the observations of Sputnik Planitia in general, and the identification of White et al. (2017) of this region as belonging to the lightly pitted plains geological unit is consistent with the presence of a sublimation-originated thin CH_4 layer.

A complication of the Sputnik endmember fitting is that the wavelength data in the hyperspectral cube appear to be misaligned by approximately one spectel (one unit of spectral sampling), which

initially resulted in the model failing to fit the spectra correctly. We are not certain of the source of this miscalibration, but once the data was shifted by one spectel to the higher wavelength, the fit worked correctly.

3.3 Profile gradient

After successfully fitting the compositional endmembers, we increase the complexity of our task by trying to fit a series of 14 points that correspond to a progressive transition of the spectra between two local endmembers, the composition and possibly also the texture smoothly evolving from one end of the series to the other. Table III.4 shows the coordinates and observation geometry of the two endpoints of this series, and Figure III.14 shows them on a global map. Figure III.18 shows all the spectra for this series, as well as their placement on the CH_4 and N_2 band depth maps.



Figure III.18: The spectra corresponding to the 14 points of the profile dataset (top) and their locations overlaid on the 1.7- μ m CH₄ and 2.15- μ m N₂ band depth maps (bottom). The maps are in their native orthographic projection and centered at around (65°N,154°E), with a spatial resolution of ~ 3.3°/px.

The profile points are located in Voyager Terra at approximately 65° N and between 152 and 156°E. This places them in the same mid-latitude belt as the North Pole endmember, but significantly more to the east, directly to the north of the western edge of Sputnik Planitia. This is an eroded patchwork terrain with variable geology, classified by Moore et al. (2016) as fretted or washboard terrain. Compositionally it is dominated by CH₄ ice, with N₂ ice present at the bottom of craters and other depressions, and potentially admixed with H₂O ice spatially (Schmitt et al., 2017a). The

profile approximately follows the slope from a point on top of a ridge (1200 m) to the bottom of a crater (-3000 m).

Since the terrain of this profile belongs to the same compositional group as the North Pole endmember, we attempt to fit it with the same surface representation. We have some difficulty attaining convergence with this model verfor some of the profile points, with it working well for the first two thirds of the spectrum, but failing for the latter third, as the N₂ band depth increases. We therefore also test a different surface representation consisting of two layers of N₂:CH₄:CO solid solution, providing a grain size and molecular proportion gradient (with a small spatial component of water ice).

In fitting this dataset with these quite complex surface models, we encountered the problem of multiplicity of solutions. Figure III.19 shows two spectra which fit the target spectrum 14 similarly well, but have different compositions.



Figure III.19: The spectra and surface representations for two close fits for the profile point 14. There are two spatially distinct components (Al and A2), with A2 being composed of two stratified layers, both of which are solid solution N_2 :CH₄:CO (A2:L2) ice. The precise parameters are listed in Table III.8.

Both spectra A and B have the same surface representation structure, consisting of an optically thin single-, but very large-grain N_2 :CH₄:CO ice layer atop a different layer of N_2 :CH₄:CO ice with a different grain size and molecular proportions, as well as a minority spatial component of H₂O ice. However, spectrum A has much smaller grains in the top layer than spectrum B (and inversely for the bottom layer), and the vertical CH₄ concentration gradient is inverted between the two spectra, with spectrum A having a more CH₄-rich top layer with a more diluted substrate, whereas spectrum 14B has the less CH₄-rich layer on top.

We note that in the continuum, the two spectra diverge from the goal in opposing directions, and that potentially combining the two different surface representations spatially, with three spatial components (H_2O , pure CH_4 atop N_2 : CH_4 :CO, and two layers of N_2 : CH_4 :CO), could provide a more accurate fit. We do not put this theory into practice due to lack of time and complexity of the resulting surface representation.

| | component | spatial % | thickness | grain size | ξ | CH ₄ % | CO % |
|---------|-------------------------------------|-----------|-----------|------------|--------|-------------------|-------|
| Fit 14A | | | | | | | |
| Al | H ₂ O | 4% | ∞ | 140 µm | 0.0 | | |
| A2:L1 | N ₂ :CH ₄ :CO | 96% | 60 mm | 60 mm | -0.750 | 0.9% | 0.0% |
| A2:L2 | N ₂ :CH ₄ :CO | 2070 | ∞ | 30 mm | -0.750 | 0.001% | 0.01% |
| Fit 14B | | | | | | | |
| Al | H ₂ O | 10.6% | ∞ | 23 µm | -0.367 | | |
| A2:L1 | N ₂ :CH ₄ :CO | 89 4% | 324.1 mm | 160.5 mm | -0.730 | 0.12% | 0.02% |
| A2:L2 | N ₂ :CH ₄ :CO | 07.470 | 8 | 0.95 mm | -0.060 | 0.22% | 0.41% |

Table III.8: The surface representation parameters for two close fits for the profile point 14. There are two spatially distinct components (Al and A2), with A2 being composed of two stratified layers, both of which are N₂:CH₄:CO ice. The corresponding spectrum is shown in Figure III.19.

Both of the shown surface models are capable of eventually accurately fitting the profile spectra. However, they did not guarantee an accurate fit even with tens of runs per point. We eventually had to use both types of surface representation to closely fit the full profile. We switch from the pure CH_4 upper layer to the diluted CH_4 layer for point 10 and subsequent spectra, where the sharpness of the N₂ band does not allow for a fit with pure CH_4 . This point corresponds to the edge of the very sharp increase of the N₂ band depth (see Figure III.18).

We attempted to provide a series of fits that shows a smooth compositional evolution from point to point, while having a relatively low (<10%) RMSE, and adjusted the final parameter sets by hand to maximise fit accuracy. However, considering the multiplicity of solutions, we cannot say that this is the one correct interpretation of the data, and probably a mix of both types of surface representations occurs around the transition edge. Integrating these points into a global compositional map will require comparisons with neighbouring terrains in order to provide a globally consistent and geologically realistic picture. We will further discuss methods and considerations for this in Chapter IV.

The best fits for the profile points are shown with their corresponding target spectra in Figures III.20a–b. Table III.9 lists the surface properties of these best fit spectra, and Figure III.21 shows the evolution of the granulometry and CH₄ content as we progress along the profile. While adjusting the best fit parameters, we noticed that some of the parameters — specifically those pertaining to the water ice component — varied very little between points of the profile and tended to converge to the same value, so we regressed these to the median value and kept them constant for the full set of points in order to reduce the dimensionality of the problem. These values are not included in the parameter table for readability, but they are as follows: grain size $d = 14 \ \mu m$, spatial proportion f = 0.05, anisotropy parameter $\xi = -0.795$.

The observed fits are consistent with the description of the eroded terrain in Voyager Terra. The first spectra, while appearing to be fully dominated by CH_4 ice (Schmitt et al., 2017a; Protopapa et al., 2017), are in fact consistent with a thin layer of CH_4 ice atop a very coarse-grained layer of N_2 ice ($d \sim 1-3$ cm) with only a fraction of a percent of CH_4 diluted in it (varying from 0.5% to 0.01% between points 1 and 14 of the profile).

Specifically, for the relatively CH₄-rich end of the profile, we see an optically thin layer of CH₄ ice (h = 2 mm) which does not reach more than 2 grains of thickness, and which diminishes in thickness (and also grain size) as we move east towards the N₂-rich end of the profile. This layer eventually disappears and is replaced by a vertical CH₄ concentration gradient in the N₂ solid solution, which can in practice be explained by an in-progress sublimation of N₂ ice, which has raised the CH₄ concentration in the remaining ice, but has not yet formed a pure CH₄ crust.



Figure III.20a: Best fits for spectra 1–8 of the Voyager Terra profile. Table III.9 shows the corresponding surface parameters.



Figure III.20b: Best fits for spectra 9–14 of the Voyager Terra profile. Table III.9 shows the corresponding surface parameters.

| | top layer | | | | | | bot | tom layer | | |
|----|---------------|---------------|-------|----------|------|---------------|-------|-----------|------|----------|
| # | <i>h</i> (mm) | <i>d</i> (mm) | ξ | CH_4 % | CO % | <i>d</i> (mm) | ξ | CH_4 % | CO % | RMSE (%) |
| 1 | 2.03 | 1.773 | -0.62 | 100 | 0 | 20 | -0.73 | 0.5 | 0.01 | 3.68 |
| 2 | 2.23 | 1.791 | -0.63 | 100 | 0 | 30 | -0.79 | 0.15 | 0.01 | 3.73 |
| 3 | 2.23 | 1.791 | -0.63 | 100 | 0 | 30 | -0.79 | 0.15 | 0.01 | 4.37 |
| 4 | 1.469 | 1.168 | -0.65 | 100 | 0 | 35 | -0.79 | 0.12 | 0.01 | 2.97 |
| 5 | 1.337 | 0.95 | -0.65 | 100 | 0 | 30 | -0.79 | 0.08 | 0.01 | 4.66 |
| 6 | 1.264 | 0.892 | -0.73 | 100 | 0 | 15 | -0.79 | 0.056 | 0.01 | 3.25 |
| 7 | 1.199 | 0.828 | -0.73 | 100 | 0 | 30 | -0.79 | 0.028 | 0.01 | 3.70 |
| 8 | 0.726 | 0.536 | -0.79 | 100 | 0 | 10 | -0.79 | 0.013 | 0.01 | 3.28 |
| 9 | 0.707 | 0.538 | -0.79 | 100 | 0 | 10 | -0.79 | 0.02 | 0.01 | 3.84 |
| 10 | 30 | 22 | -0.75 | 1 | 0.1 | 35 | -0.75 | 0.02 | 0.1 | 4.40 |
| 11 | 30 | 25 | -0.75 | 0.8 | 0.1 | 35 | -0.75 | 0.018 | 0.1 | 5.01 |
| 12 | 30 | 25 | -0.75 | 0.8 | 0.1 | 35 | -0.75 | 0.015 | 0.1 | 4.04 |
| 13 | 35 | 30 | -0.75 | 0.8 | 0.1 | 35 | -0.75 | 0.01 | 0.1 | 4.21 |
| 14 | 40 | 35 | -0.75 | 0.8 | 0.1 | 40 | -0.75 | 0.01 | 0.1 | 4.02 |

Table III.9: Best fit parameters for the 14 points of the Voyager Terra profile. Figures III.20a–b show the spectra. The dotted line marks the transition between surface representations, with spectra 1–9 using the pure CH_4 top layer model and spectra 10–14 using the N_2 : CH_4 :CO top layer model.

The sharp transition between two types of terrain in our model is a necessity driven by the need to reduce the complexity of the problem and the computation time required. In reality, it is probably more gradual, with cohabitation of both terrain types in one pixel with a progressive change in their spatial proportions. Ideally, in the future the transition region needs to be fitted with a more complex representation. In addition to the spatial mix, another possibility is the presence of three layers, with a very thin layer of fine-grained CH_4 overlying two layers of $N_2:CH_4:CO$ ice, which a vertical composition and texture gradient.



Figure III.21: The evolution of the grain sizes, layer thickness, and CH_4 content of the surface representation along the Voyager Terra profile. The dotted line marks the transition between the surface representations used, with spectra 1–9 using the pure CH_4 top layer model and spectra 10–14 using the N_2 : CH_4 :CO top layer model.

3.3.1 Next steps towards constructing a global map

The work done in this chapter on quantitatively fitting real LEISA data for the surface of Pluto with a stratified model obtained satisfactory results, but required a large amount of manual adjustment of model parameters and a preliminary assessment of each spectrum to determine the best surface representation to use for it before running the inversion process. This would represent a prohibitive amount of work for the global map of Pluto, and in order for this project to continue we need to find ways to automate more of the fitting process and to reduce the dimensionality of the problem. We discuss this in more detail in the subsequent chapter.

Costs and methods for global map inversion

1 Computational costs and resources

1.1 Upper and lower bounds on computing the global Pluto map

There are many variables influencing the computation time needed to produce a quantitative global map of Pluto with DISORT, such as the efficiency of the algorithm (the number of iterations and of runs needed for an accurate fit for each pixel) and the complexity of the surface. We can compute a lower and upper bound for this based on the results of calibration on synthetic spectra as well as the endmember fitting, detailed respectively in sections 3.1 and 3.2 of Chapter III.

The global map of Pluto produced from LEISA data as detailed in Chapter II, at a resolution of 2 km/px, contains slightly over 36 million¹ non-empty pixels. This number does represent a significant amount of oversampling for most of the data. Only the very high-resolution encounter swath that represents about 11% of the surface has a resolution of 2.7 km/px, while the rest of the encounter hemisphere is imaged at ≈ 6 km/px. The anti-encounter hemisphere map is constructed from approach data as described in Chapter II, with resolutions ranging from 30 to 400 km/px. Based on this, we estimate the number of real-data pixels that require individual inversion as 3.75 million, an order of magnitude lower than the oversampled high-resolution map.

The computation time of this inversion varies drastically depending on the complexity of the spectrum. To establish a lower bound on calculation time, we assume the following:

- all the spectra can be modeled with a simple one-component surface;
- one in two runs results in an accurate fit;
- 1000 iterations are sufficient for an accurate fit.

This corresponds to a runtime of 48 minutes per spectrum on a node of the Froggy cluster, so about 342.5 years of computing time for the full map. Froggy has 190 computing nodes available to users, so if we were permitted to use them all 24/7 (which is not possible as there are many other users of the CIMENT grid), the full map computation would still take almost two years. In reality, this needs to be multiplied by at least 4 to account for the projects of other users, as even the largest CIMENT projects do not use more than 25% of the available CPU time. We can see that even in these extremely simplified lower-bound conditions this is not a realistically accomplishable PhD thesis project. To determine whether it is a project that is potentially attainable with more resources, we also compute the upper bound of the needed calculation time, assuming the following:

- a complex surface with both spatial mixing and stratification is required to accurately fit each pixel;
- we only obtain an accurate result for 1 in 10 cases;
- at least 5000 iterations are necessary for convergence.

This means that we need approximately 43 hours per pixel to obtain an accurate fit, which corresponds to 96.9 years of full 24/7 occupation of the CIMENT Froggy cluster to compute the full Pluto map. To see whether we can reduce this to a duration significantly smaller than the human lifespan, we turn to European supercomputing clusters, which we describe in the next section.

¹36,107,188 is the exact number of pixels.

1.2 Can supercomputers help?

The CIMENT cluster is what the CNRS refers to as a "computing mesocenter", a medium-sized regional supercomputing center that provides high-performance computing resources to one or several research institutions. A mesocenter typically provides tens or hundreds of teraFLOPs², and each of Froggy's 190 nodes has on average close to 0.5 TFLOPs.

To go beyond this, two options are available. The first option is to use a different kind of hardware, typically graphics processing units (GPUs), although FPGAs and ASICs³ would potentially be an option. The second, more promising one, is to use supercomputers or large computing clusters made of many nodes, similar to Froggy on a larger scale. The advantage of GPUs is that they can provide massive speed improvements for parallel computing tasks. For example, Nvidia's latest generation high-end datacenter GPU, the Al00 (costing around $20k\in$), has a nominal performance peak of 19.5 teraFLOPs. Two problems prevent efficiently using this kind of hardware. The first is that this type of graphics processor is optimised for certain simple operations (and some specialised AI tensor operations for the Al00), which means that they might not efficiently support an algorithm such as DISORT. The second issue, related to the first, is linked to memory usage. DISORT requires around 25MB of maximum RAM per run, which is negligible for a CPU. However, this is at least a few times higher than what most GPU nodes have access to, preventing us from using this solution (as well as ASICS and FPGAs which suffer from the same issues).

The second option is to use a more powerful supercomputer. For large projects that require more resources than can be provided by a mesocenter, researchers can submit proposals to one of the national or European supercomputing centers. Table IV.1 compares the machines available to us for this thesis — a Dell laptop computer and the Froggy cluster — to the fastest supercomputer in France available for research projects (Joliot-Curie, housed at CEA's Very Large Computing Center) as well as, for informational purposes, to the fastest supercomputer in the United States (Summit, housed at the Oak Ridge National Laboratory in Tennessee).

We can see that in the worst-case scenario, the global map computation requires nearly half a year of full-capacity supercomputer time. In the best case, the computation is within the scope of a GENCI⁴ project. However, we expect most of the pixels to be closer to the more complex upperbound case, and so we cannot submit a proposal that will guarantee results within the allocated computational time. While GENCI allocates 22 billion CPU hours annually, on average a project only uses 5 million CPU hours, i.e. 0.02% of the full capacity, and we cannot reasonably request to use 40% of the full capacity, as the worst-case computation would require.

| Computer | Country | Institute | Speed | Runtim | e (years) |
|--------------------|---------------|-------------|----------|--------|-----------|
| | | | (TFLOPs) | min | max |
| Dell Latitude 7480 | France | IPAG | 0.133 | 642.3 | 38777.4 |
| Froggy (190 nodes) | France | CIMENT | 88.25 | 1.80 | 97.66 |
| Joliot-Curie | France | CEA | 22000 | 0.007 | 0.39 |
| Summit | United States | DOE/SC/ORNL | 148 600 | 0.001 | 0.06 |

Table IV.1: A comparison of the computing power of the machines available for this thesis to the fastest supercomputers in France and the United States, and the estimated runtime in the best and worst cases for the completion of the global map computation on each. The full runtime for the two lower-power machines is extrapolated directly from the algorithm runtime for one pixel, while that for HPC clusters is estimated from the computation speed ratios compared to Froggy. The speed for the Dell computer corresponds to the Geekbench 4 floating-point SGEMM benchmark (this value is not used for the full runtime computation and is only given for informational purposes). The speed for the Froggy cluster is that given in the user documentation, and the speed for the two HPC clusters corresponds to the LINPACK benchmark.

²Floating-point operations per second (FLOPs) are a common way to measure the performance of supercomputers. The CPU of a modern laptop computer may have a speed of several gigaFLOPs, while the most powerful supercomputers have recently achieved speeds of hundreds of petaFLOPs.

³Field Programmable Gate Arrays and Application-Specific Integrated Circuits are more specialised hardware that are optimised for a single task and thus less generic than GPUs (themselves less generic than CPUs).

⁴GENCI (Grand équipement national de calcul intensif) is a company co-owned by the French state and several large research organisations, whose purpose is to coordinate and allocate usage of the French high-performance computing centres.

1.3 Distributed computing

Centralised supercomputing clusters are not the only way to solve computationally expensive research problems. If a problem can be split into many independent parts, it can be distributed to a grid of individual computers which can all work on the problem in parallel. This can take the form of a citizen science project where volunteers allow the computation to use the idle CPU time of their machines. Distributed computing projects have been used for a varied range of problems such as climate modeling, protein folding, and searching for extraterrestrial radio signals. Notable projects are the Great Internet Mersenne Prime Search, one of the very first internet-based volunteer computer projects, launched in 1996; SETI@home, which pioneered using distributed computing to analyse large volumes of astronomical data, and ran for over 20 years before going on indefinite hiatus in March 2020; and Folding@home, a protein-folding project that is the first computing system to reach a speed measured in exaFLOPs.

At its peak computing speed of 2.43 exaFLOPs, the Folding@home infrastructure would have been able to compute the global map in one or two days. This computational peak was reached in April 2020 and driven by interest in finding solutions for the COVID-19 pandemic; the typical computing power provided by the Folding@home network is in 100–200 petaFLOPs range⁵. While most distributed computing networks are less powerful, they are still comparable to a national-level supercomputer. A distributed computing project would be significantly easier and cheaper to launch than submitting a successful project proposal to a national supercomputer, and it would also be much more flexible in terms of extending its end-date and requesting more computational resources.

1.3.1 Setting up a project with the BOINC infrastructure

Ready infrastructure exists for researchers interested in launching a distributed computing project. The BOINC (Berkeley Open Infrastructure for Network Computing) system was initially designed to support SETI@home, but has been generalised for use as a platform for many other applications. Its total 24-hour processing average as of July 6th 2021 was 25.1 petaFLOPs⁶, but it has recently been as high as 41.6 petaFLOPs (as of March 2020). The initial setup cost of a BOINC project is low. It requires a server computer to host the data and provide a hub for clients to connect to, and setting up a public-facing website that describes the project to interested volunteers.

The BOINC documentation estimates the cost of running a BOINC project at \$125,000⁷. However, this is largely comprised of system administrator hours, with the barebone hardware setup cost being at most several thousand euros. The software setup and maintenance costs will then be dictated by the availability of a research technician for the project.

Once the project is set up, the most important factor for success is publicising it. Many of the largest citizen science projects active today are space science projects, which evidences the public's interest in contributing to space research. This is a promising sign for the potential buy-in for a distributed project to compute the map of Pluto. More concretely, this PhD thesis was carried out as part of the New Horizons mission, and we can therefore make use of the NASA publicity apparatus to gain a wider reach and to make as many people aware of the mapping project as possible.

1.3.2 Using an independent infrastructure

Using a readymade infrastructure such as BOINC greatly streamlines the initial setup. However, it means we are constrained to the BOINC project architecture, which consists of downloading and processing independent work units. In our case this would mean independently processing each pixel and sending the best fit results back to the central hub as they are computed. This becomes inconvenient if we want to put into practice algorithm optimisation strategies that involve comparison between different pixels (see section 2 for an overview of some of these strategies), or that involve a certain degree of active user input, e.g. having the volunteers visually assess the accuracy of a fit to see if the error is due to noise in the data or to bad fit on a compositional band. This may make it worthwhile to develop an independent service for this project that is not bounded by existing infrastructure in terms of what a user can do or have access to.

⁵This is estimated from snapshots of the Folding@home statistics webpage at https://stats.foldingathome.org/os throughout 2021.

⁶The BOINC 24-hour computing statistics can be found at https://boinc.berkeley.edu/computing.php.

 $^{^{7} \}text{Estimated for a CPU load of 100 TFLOPs for one year, from \texttt{https://boinc.berkeley.edu/trac/wiki/BoincOverview.}}$

A promising initial concept for this independent service would be to put the global hyperspectral map online and to allow users to select a pixel to work on and submit potential fits for. This sort of setup lends itself well to gamification. For instance, users could select a compositional unit on the map, consisting of adjacent pixels with good fits with the same surface model, and try to expand its borders by trying to find good fits with the same surface model for pixels around its edges. Giving individual users credit for accurately fitting a pixel, potentially coupled with a leaderboard, could also be motivational.

2 Algorithmic improvements and building a global picture

The beginning of this chapter lays out the computational requirements for the quantitative inversion of a global map of Pluto, based on the calculation cost for the test synthetic spectra and compositional endmembers shown in Chapter III. Chapter III also touches upon the problem of selecting the correct surface representation for each pixel, which we implemented in simplified form for single-layer granular mixture but did not have the time to reimplement for the more complex surface representation that DISORT allows us to model. Here we discuss this problem in more detail and describe some algorithms that can be used to solve it.

2.1 Reuse of intermediate spectra

As the fitting algorithm runs on a pixel, it generates thousands of intermediate spectra, which are largely discarded due to not being an accurate fit for the pixel in question. The generation of the spectrum is the most computationally expensive step of an iteration, taking 1.5–3 seconds on a Froggy node, whereas the RMSE computation takes around a tenth of a millisecond for a spectrum at LEISA spectral resolution. As it is quite possible that the spectrum would be a good fit for a different pixel of the map, discarding it after one comparison is inefficient, and ideally it should be compared to many other pixels.

To avoid greatly extending the computational needs of the project, the algorithm should not spend more than 10% of its runtime on these comparisons. This means that for each generated spectrum, we can run 1000–2000 comparisons to other pixels. In the simplest implementation of this process, we can select these pixels randomly from the map. However, this only constitutes about 0.3% of the total number of pixels, and so even if there is a pixel in the map for which the generated spectrum is a good match, our chance of finding it is using this procedure is quite low. We can improve this by precomputing a tree-like structure of histograms for the LEISA spectra, which would be generated and categorised as follows:

- 1. the spectrum is split into a small set of bands, and the spectels within each band are averaged (shown in Figure IV.1 for the Sputnik and North Pole endmembers);
- 2. the range of the values in each of these bands is split into equal-sized bins (e.g. 20, with each containing 5% of the range);
- for each band, the spectrum is assigned to the two bins closest to its average in that band (two bins are selected to avoid discontinuities);
- 4. this assignment is stored in a tree-like structure, such that one can query it with a list of band averages (calculated from a generated spectrum) and get exactly the spectra that are in the same bins as it on all bands.

This allows us to quickly obtain the most likely candidates that a generated spectrum could fit, and can be adjusted for speed by modulating the size of the bins and the number of bands. Assuming six bands, the returned set of candidate spectra should hopefully always be in the tens to hundreds range. This could be further improved by having bands of non-uniform width, with some centered on specific spectral characteristics such as absorption bands. Having non-uniform bin size could also reduce the number of candidates by splitting them more equally into different bins.

An additional use for this global check is potentially having a better-than-random start point for fitting those pixels later on: even if there are no previously-computed intermediate spectra that fit a given pixel very closely, one that is close enough can be used as a starting point for the fitting algorithm for that pixel, which will improve convergence likelihood.



Figure IV.1: Example of band averages for testing spectrum similarity using the Sputnik and North Pole endmembers, with six equal-width bands.

2.2 Reducing surface dimensionality

As we have established, naively fitting a spectrum with a maximally complex surface representation, with all of the different possible components present in every kind of possible mixture, is an at least 35-dimensional problem which is extremely computationally intensive to try to solve and the complexity of which may not allow for solving it in any reasonable amount of time. We are therefore interested in identifying what components may or may not be in a pixel prior to running the fitting algorithm. We can approach this with a number of complementary tools, starting with image segmentation.

2.3 Techniques for image segmentation

The raw data returned by a remote sensing instrument is generally a spatially distributed set of values in a continuous range. In addition to studying these values directly, there is value to splitting them into categories by similarity of by some set of threshold criteria to produce a *classification* of the image, in order to more easily see patterns in the spatial distribution of the data. The most important problem in map classification is the question of *classification schema*, i.e. what rules or thresholds to use to produce the most informative set of classes from a dataset. To be useful, a classification schema must ideally be *mutually exclusive* (each datum belongs to only one class) and *collectively exhaustive* (all data belong to at least one class). We will briefly introduce several different classification schema and their relative utilities.

The simplest type of schema is one that can be applied to a two-dimensional dataset, with a single value assigned to each spatial location. This is **quantitative thresholding**, where we divide the full range of values into several classes. These can be equal intervals, quantiles, or be based on one of several thresholding rulesets commonly used in cartography.

For higher-dimensional datasets, manually selecting all the applicable thresholds for each class becomes unwieldy, and automatic classification methods are used. These are generally split into two categories: **supervised**, where the algorithm is given samples that correspond to each desired class, and from which it builds definitions that it then tries to match the remaining datapoints to; and **unsupervised**, where the algorithm is given a desired number of classes and tries to find a way to split the data into that number of classes.

2.3.1 Automatic classification to generate seed pixels

We briefly touched upon map classification in Chapter II, where we produced classified maps of Pluto using a Gaussian mixture model. This was an example of an unsupervised classification method; given a number of classes, it found clusters in the data which we then characterised based on their properties. We can try to use these classified maps to reduce the dimensionality of the model inversion. If a pixel is adjacent to and in the same class as another pixel for which we have already determined a surface representation, such as one of the endmembers described in Chapter III, then we can try to fit it with the same surface representation (varying the parameters) and see whether we can also obtain an accurate fit for this pixel.

This similarity-based approach means that many "seed" pixels need to be identified in order to provide sufficient points of comparison for fitting all the other pixels. The surface representations we use are sufficiently complex and the RTM's response to small changes is sufficiently sensitive that there will most likely be pixels for which we fail to obtain an accurate fit via comparison with seed pixels.

Conversely, we have several classification schema for these maps and the relationships between pixels may be different when examined through different schema, and so we may have multiple possible seeds for a given pixel. The accuracy of the fits will be unavoidably limited by the low resolution and noisiness of the data, and so we may run into cases where we attempt to fit a pixel with several different seeds and obtain similarly accurate results with more than one of them. This is a situation that can be resolved using graph-theoretical tools, which we will describe in the next section.

2.4 Graph colouring

2.4.1 What is a graph?

Any structure with elements that are linked can be represented as a graph. A graph consists of vertices, or nodes, which are connected with edges. The field of graph theory studies properties of graphs. Graphs can be used to represent communication networks, the organisation of data, and many other kinds of structures. Relevantly in our case, an image, which is a two-dimensional grid⁸, can be represented as a graph where each pixel is a node, connected by edges to the nodes corresponding to adjacent pixels. This means we can use graph-theoretical tools to help determine relationships between pixels.

2.4.2 Hyperspectral images as graphs

Graph-theoretical tools form an essential part of many kinds of image processing, such as object recognition. Some researchers have applied graph theory to hyperspectral data in particular, notably Bai et al. (2013), who use graph cut theory to classify remote sensing hyperspectral images. They note that while individually classifying the pixels in a hyperspectral dataset is simple and robust, it results in a disarranged list of individual spectral signals with no spatial correlation, and leaves many tiny regions that need human labor to identify. The authors therefore propose an application of graph cut theory, reducing the problem of finding the best labelling for the pixels to one of finding a minimum-weight set of edges that cuts a graph into two connected components.

2.4.3 Improper list colouring to identify classes

The approach mentioned earlier with the seed pixels can serve as inspiration for the elaboration of a global map. Let's suppose that, using the seed pixels and progressing in a greedy fashion, we have obtained a few candidate compositions for each pixel. The next step is to find a most-likely composition for large swathes of the map. We are in a problem related to list-colouring, where to each vertex of a graph must be assigned a colour from a given set (the set depending on the vertex).

Unlike most colouring problems, we are seeking an improper colouring, and more precisely one that minimises the number of discontinuities. If we create a discontinuity-cost function D which assigns a small cost to an edge in the map depending on the difference between the compositions

⁸For hyperspectral data each element of the grid contains a spectrum.

of the two adjacent vertices, we can formalise the problem in the following way, with *v* referring to a vertex on the map and *e* referring to an edge between adjacent vertices:

Select a colour assignment
$$F(v \rightarrow \text{colours}(v))$$
 minimising $\sum_{e} D(e) + \sum_{v} RMSE(v)$

Although empirical validation would be needed to confirm it, this is typically the kind of problem that can be efficiently solved with either linear programming or even hopefully greedy algorithms.

3 Global summary and priorities for moving forward

As established in the first section of this chapter, an accurate quantitative mapping of Pluto's surface composition is an arduous and computationally costly task. Over the course of a four-year PhD program, we have established the bases for this mapping project to build upon. As detailed in Chapter II, we expanded the qualitative mapping by Schmitt et al. (2017a) of the Pluto encounter hemisphere to a global hyperspectral dataset. We applied the qualitative analysis from Schmitt et al. (2017a) to this global dataset, producing band depth and spectral index maps, and carrying out statistical analyses on these maps that allowed us to infer potential geological interpretations of the data.

We then developed a metaheuristic algorithm to fit a complex surface model, with several different potential components mixed in different ways, including stratification, to the spectra, thus setting the stage for a global quantitative map. We validated the accuracy and efficiency of this algorithm with synthetic data, and computed quantitative surface representations for several regions of interest on Pluto's surface, notably in the N₂-rich plain of Sputnik Planitia and in the mid-latitude dissected CH₄ highland terrains. These real surface fits required multilayered surface representations, with vertical stratification in the composition.

Throughout this thesis, we have implemented efficient techniques for data analysis that are wellknown and used in other fields, such as intensity-based registration for medical imagery, but are largely overlooked in planetary science. The continuous improvements in space mission instrument loadouts and transmission methods inevitably lead to a relentless growth in size and complexity of the datasets they return, and in turn to the need for more and more powerful tools to analyse these datasets. With this in mind, we expect to continue to develop this global mapping project as a distributed computing program, and to further deepen its interdisciplinarity with the application of theoretical computer science tools to reduce the computational cost and algorithmic complexity of the problem.

The number one prority for continuing the project will be documentation and open-sourcing of the simulated annealing algorithm, and either development of a distributed computing infrastructure from scratch or the adoption of an existing one such as BOINC. This is a task projected for at most a few months, and subsequent work shall be less time-intensive, requiring only maintenance of the servers and the project infrastructure. The success of the project from thereonwards shall be predicated largely on its publicity and institutional support from both French institutions and NASA via the New Horizons team.

List of symbols

| Quantity | Unit | Symbol |
|--|---------------------|-----------|
| Radiative transfer | | |
| Wavelength | m | λ |
| Flux | W | Φ |
| Irradiance | W/m^2 | J |
| Radiance | $W/(m^2 \times sr)$ | Ι |
| Cross-section | m ² | σ |
| Reflectance | | R |
| Phase function | | Р |
| Optical geometry | | |
| Incidence angle | rad, ° | i |
| Emergence angle | rad, ° | е |
| Phase angle | rad, ° | g |
| Scattering angle | rad, ° | θ |
| Azimuthal angle | rad, ° | φ |
| Surface properties | | |
| Single-scattering albedo | | ω |
| Optical depth | | τ |
| Henyey-Greenstein anisotropy parameter | | ξ |
| too Grain size | m | d |
| Layer thickness | m | h |
| Optical constants | | |
| Refractive index | | п |
| Extinction coefficient | | k |
| Simulated annealing | 5 | |
| Perturbation factor | | р |
| Temperature | | Т |
| Statistics | | |
| Pearson correlation coefficient | | r |

List of Tables

| I.1 I.2 | Size and orbital properties of Pluto's small satellites | 19 23 |
|--------------|--|----------|
| II.1 II.2 | List of Pluto approach-phase LEISA observations used | 37 41 |
| III.1 | Physical properties of ice optical constants | 64 |
| III.2 | Surface structure and composition for the synthetic calibration spectra | 67 |
| III.3 | Best fit of simple spectrum with complex surface | 73 |
| III.4 | Location and observation geometry of endmembers and profile endpoints | 75 |
| III.5 | Best fits with Spectrimag for North Pole endmember | 76 |
| III.6 | Model parameters for DISORT fit of North Pole endmember | 77 |
| III.7 | Model parameters for DISORT fit of Sputnik endmember | 78 |
| III.8 | Example of solution multiplicity for profile point fit | 81 |
| III.9 | Voyager Terra profile best fit spectra | 83 |
| IV.1 | Comparison of speed and map runtime for different computers and clusters | 86 |

List of Figures

| I.1 | The electromagnetic spectrum | 7 |
|-------|---|----|
| I.2 | Acquisition and spatial scanning modes of spectral imagers. | 8 |
| I.3 | The spatial resolution parameters of an optical sensor | 9 |
| I.4 | An example spectral response function with FWHM | 9 |
| I.5 | The fundamental vibrational modes of CO_2 and H_2O | 11 |
| I.6 | Comparison of Gaussian and Lorentzian band shapes | 12 |
| I.7 | Irradiance and radiance | 12 |
| I.8 | Single-particle scattering schema | 13 |
| I.9 | Photometric angles | 14 |
| I.10 | Radiative transfer equation schematic | 14 |
| I.11 | Surface representation: mixture types | 17 |
| I.12 | The effect on spectra of different surface mixture types | 18 |
| I.13 | Pluto and Charon bulk properties and proposed internal structure | 19 |
| I.14 | Photographic plates from Clvde Tombaugh's discovery of Pluto | 20 |
| L15 | Short- and long-term variations in Earth-based observation of Pluto | 21 |
| L16 | Pluto's seasons | 22 |
| L17 | Compositional models of Pluto's surface based on Hubble Space Telescope data | 22 |
| I.18 | Schematic of New Horizons showing science instruments | 24 |
| L19 | The different observation modes of New Horizons instruments | 25 |
| I.20 | The fields of view of New Horizons optical instruments | 25 |
| I.21 | Illustration of time delay integration (TDI) | 26 |
| L22 | Spectrum of Pluto with the MVIC colour filters for comparison | 26 |
| L23 | Resolving power of the LEISA filter assembly as a function of wavelength | 27 |
| L24 | An overview of the New Horizons mission science timeline for 2015 | 28 |
| L25 | Pluto approach schedule | 28 |
| L26 | Global map of Pluto as imaged by New Horizons in enhanced colour | 29 |
| L27 | Albedo map and digital elevation model of Pluto | 30 |
| L28 | MVIC spectral slope map | 31 |
| I.29 | MVIC 980 nm CH ₄ band map | 31 |
| 1.30 | Effect of PCA on LEISA data | 32 |
| 131 | N_{2} transport in Sputnik Planitia at different timescales | 33 |
| 132 | Pluto encounter hemisphere maps showing qualitative chemical abundances | 34 |
| I 33 | Schema of the evolution and mixing of materials on Pluto's surface | 34 |
| 1.22 | benefiti of the evolution and mixing of materials on Fraco's surface | 21 |
| II.1 | A frame from a LEISA approach image before and after preprocessing | 36 |
| II.2 | Global coverage obtained with the encounter and approach datasets combined | 37 |
| II.3 | A comparison of PCC and MICC similarity metrics for a set of sample distributions | 39 |
| II.4 | Illustration of the effect of misregistration | 39 |
| II.5 | Transformation functions for image registration | 40 |
| II.6 | LEISA registration validation with control points | 41 |
| IL7 | Global registered CH_4 1.7-um band depth map of Pluto | 43 |
| IL8 | Global registered N ₂ 2.15- μ m band depth map of Pluto | 43 |
| II.9 | Global registered H ₂ O spectral index map of Pluto | 44 |
| II.10 | Global registered red material spectral index map of Pluto | 44 |
| II.]] | Registered global LEISA maps in north polar orthographic projection | 45 |
| II.12 | Distribution of LEISA dataset values as a function of latitude | 46 |
| | | |

| II.13 | Distribution of LEISA dataset values as a function of longitude | 46 |
|----------------|--|----|
| II.14 | Global registered CH ₄ map of Pluto (1.95 μ m integrated band depth) | 47 |
| II.15 | Global registered CH ₄ band position index (CH ₄ state index) map of Pluto | 47 |
| II.16 | Hexagonal bin plots correlating different LEISA maps | 48 |
| II.17 | Classified map correlating the H ₂ O and RedMat spectral indices | 49 |
| II.18 | Classified map correlating the N ₂ and 1.7 μ m CH ₄ bands | 49 |
| II.19 | Classified map correlating the CH ₄ band position index and the N ₂ band | 49 |
| II.20 | Hexagonal bin plots correlating LEISA and MVIC maps | 50 |
| II.21 | Classified map correlating LEISA 1.95 μ m and MVIC 0.98 μ m CH ₄ bands | 51 |
| II.22 | Global classified map correlating the H ₂ O spectral index and MVIC slope | 51 |
| II.23 | Value distribution of the four LEISA maps for the six MVIC terrain types | 52 |
| II.24 | Hexagonal bin plots correlating different LEISA maps and the DEM | 53 |
| II.25 | LEISA CH_4 1.7 μ m band depth map overlaid on the DEM | 54 |
| II.26 | LEISA N ₂ 2.15 μ m band depth map overlaid on the DEM | 54 |
| II.27 | Correlations between all datasets as a function of latitude | 55 |
| II.28 | Correlations between all datasets as a function of longitude | 56 |
| III 1 | Schema of multilayered medium modeled by DISORT | 60 |
| III.1 III.2 | Simulated appealing diagram | 62 |
| III.2 III.3 | Simulated annealing flowchart | 63 |
| III.2 III 4 | High-resolution and convolved spectrum example | 65 |
| III.1 | Synthetic spectra used for model validation and calibration | 66 |
| III.5 III.6 | Calibration of different relative fitting phase lengths | 67 |
| III.0 III 7 | Calibration of simulated annealing parameters: RMSE | 68 |
| III.7 | Calibration of simulated annealing parameters: error score | 69 |
| III.0 | Calibration of simulated annealing run duration: RMSE | 70 |
| IIL10 | Calibration of simulated annealing parameters: error score | 70 |
| III.II | Testing fit with noise | 71 |
| III.12 | Example fits to noisy spectra | 72 |
| III.13 | Best fit of simple spectrum with complex surface | 73 |
| III.14 | Location of endmembers and profile points | 74 |
| III.15 | Best fits with Spectrimag for North Pole endmember | 75 |
| III.16 | Best DISORT fit for North Pole endmember | 76 |
| III.17 | Best DISORT fit for Sputnik endmember | 78 |
| III.18 | Upland terrain profile spectra and map | 79 |
| III.19 | Example of two close profile point fits with different surface models | 80 |
| III.20 | Voyager Terra profile best fit spectra | 82 |
| III.21 | Profile parameter evolution | 84 |
| IV 1 | Band averages example for similarity testing | 80 |
| 1 4 . 1 | build averages example for similarity testing | 57 |

Bibliography

- Baer, J. W. (2005). "The Deep Impact mission and instruments". In: *Tribute to Warren Smith: A Legacy in Lens Design and Optical Engineering*. Ed. by R. E. Fischer. SPIE. DOI: 10.1117/12.624559.
- Bai, J., S. Xiang, and C. Pan (2013). "A Graph-Based Classification Method for Hyperspectral Images". In: *IEEE Transactions on Geoscience and Remote Sensing* 51.2, pp. 803–817. DOI: 10.1109/tgrs. 2012.2205002.
- Basri, G. and M. E. Brown (2006). "Planetesimals to Brown Dwarfs: What is a Planet?" In: Annual Review of Earth and Planetary Sciences 34.1, pp. 193–216. DOI: 10.1146/annurev.earth.34. 031405.125058.
- Bertrand, T. et al. (2018). "The nitrogen cycles on Pluto over seasonal and astronomical timescales". In: *Icarus* 309, pp. 277–296. DOI: 10.1016/j.icarus.2018.03.012.
- Bertrand, T. et al. (2019). "The CH₄ cycles on Pluto over seasonal and astronomical timescales". In: *Icarus* 329, pp. 148–165. DOI: 10.1016/j.icarus.2019.02.007.
- Bertrand, T. et al. (2020). "Pluto's beating heart regulates the atmospheric circulation: results from high resolution and multi-year numerical climate simulations". In: *Journal of Geophysical Research: Planets* 125.2. DOI: 10.1029/2019je006120.
- Bertrand, T. and F. Forget (2016). "Observed glacier and volatile distribution on Pluto from atmospheretopography processes". In: *Nature* 540.7631, pp. 86–89. DOI: 10.1038/nature19337.
- Britt, R. R. (2006). *Pluto May Get Demoted After All*. Ed. by Space.com. URL: https://www.space.com/2764-pluto-demoted.html.
- Buie, M. W. et al. (2010). "Pluto and Charon with the Hubble Space Telescope. II. Resolving changes on Pluto's surface and a map for Charon". In: *The Astronomical Journal* 139.3, pp. 1128–1143. DOI: 10.1088/0004-6256/139/3/1128.
- Buras, R., T. Dowling, and C. Emde (2011). "New secondary-scattering correction in DISORT with increased efficiency for forward scattering". In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 112.12, pp. 2028–2034. DOI: 10.1016/j.jqsrt.2011.03.019.
- Campbell, B. A. (2002). Radar remote sensing of planetary surfaces. Cambridge University Press.
- Campbell, J. B. and R. H. Wynne (2011). Introduction to remote sensing. Guilford Press.
- Christy, J. W. and R. S. Harrington (1978). "The satellite of Pluto". In: *The Astronomical Journal* 83, p. 1005. DOI: 10.1086/112284.
- Douté, S. (1998). "Modelisation numerique de la reflectance spectrale des surfaces glacées du systeme solaire. Application a l'analyse de spectres de Triton et Pluton et au traitement d'images hyperspectrales NIMS de Io". PhD thesis. Paris 7.
- Douté, S., B. Schmitt, and E. Quirico (1999). "Evidence for methane segregation at the surface of Pluto". In: *Icarus* 142.2, pp. 421–444. DOI: 10.1006/icar.1999.6226.
- Douté, S. and B. Schmitt (1998). "A multilayer bidirectional reflectance model for the analysis of planetary surface hyperspectral images at visible and near-infrared wavelengths". In: *Journal of Geophysical Research: Planets* 103.E13, pp. 31367–31389. DOI: 10.1029/98je01894.
- Droste, S., T. Jansen, and I. Wegener (2002). "On the analysis of the (1+1) evolutionary algorithm". In: *Theoretical Computer Science* 276.1-2, pp. 51–81. DOI: 10.1016/s0304-3975(01)00182-7.
- Earle, A. M. et al. (2018). "Methane distribution on Pluto as mapped by the New Horizons Ralph/MVIC instrument". In: *Icarus* 314, pp. 195–209.

- Fayolle, M. et al. (2021). "Testing tholins as analogues of the dark reddish material covering Pluto's Cthulhu region". In: *Icarus* 367, p. 114574. DOI: 10.1016/j.icarus.2021.114574.
- Flanigan, S. H. et al. (2016). "Destination Pluto: New Horizons performance during the approach phase". In: *Acta Astronautica* 128, pp. 33–43. DOI: 10.1016/j.actaastro.2016.02.029.
- Ghasemalizadeh, O., S. Khaleghian, and S. Taheri (2016). "A review of optimization techniques in artificial networks". In: *International Journal of Advanced Research* 4.9, pp. 1668–1686. DOI: 10. 21474/ijar01/1627.
- Granville, V., M. Krivanek, and J.-P. Rasson (1994). "Simulated annealing: a proof of convergence". In: IEEE Transactions on Pattern Analysis and Machine Intelligence 16.6, pp. 652–656. DOI: 10.1109/ 34.295910.
- Grundy, W. M., M. W. Buie, and J. R. Spencer (2002). "Spectroscopy of Pluto and Triton at 3-4 Microns: Possible Evidence for Wide Distribution of Nonvolatile Solids". In: *The Astronomical Journal* 124.4, pp. 2273–2278. DOI: 10.1086/342933.
- Grundy, W. M. and B. Schmitt (1998). "The temperature-dependent near-infrared absorption spectrum of hexagonal H₂O ice". In: *Journal of Geophysical Research: Planets* 103.Ell, pp. 25809–25822. DOI: 10.1029/98je00738.
- Grundy, W. M. et al. (2016). "Surface compositions across Pluto and Charon". In: *Science* 351.6279, aad9189-aad9189. DOI: 10.1126/science.aad9189.
- Grundy, W. and M. Buie (2001). "Distribution and evolution of CH₄, N₂, and CO ices on Pluto's surface: 1995 to 1998". In: *Icarus* 153.2, pp. 248–263. DOI: 10.1006/icar.2001.6684.
- Grundy, W., B. Schmitt, and E. Quirico (1993). "The Temperature-Dependent Spectra of α and β Nitrogen Ice with Application to Triton". In: *Icarus* 105.1, pp. 254–258. DOI: 10.1006/icar. 1993.1122.
- Grundy, W. et al. (2013). "Near-infrared spectral monitoring of Pluto's ices: Spatial distribution and secular evolution". In: *Icarus* 223.2, pp. 710–721. DOI: 10.1016/j.icarus.2013.01.019.
- Hapke, B. (2012). Theory of reflectance and emittance spectroscopy. Cambridge University Press.
- Howard, A. D. et al. (2017). "Present and past glaciation on Pluto". In: *Icarus* 287, pp. 287–300. DOI: 10.1016/j.icarus.2016.07.006.
- IAU (2006a). "Resolutions B5 and B6: "Definition of a Planet in the Solar System" AND "Pluto"". In: 26th General Assembly.
- (2006b). The IAU draft definition of "planet" and "plutons". URL: https://www.iau.org/news/ pressreleases/detail/iau0601/.
- Johnson, P. E. et al. (2021). "Modeling Pluto's minimum pressure: implications for haze production". In: *Icarus* 356, p. 114070. DOI: 10.1016/j.icarus.2020.114070.
- Kallenbach, R. (2000). "Pick-up Ion Measurements in the Heliosphere A Review". In: Astrophysics and Space Science 274.1/2, pp. 97–114. DOI: 10.1023/a:1026587620772.
- Kremer, J. et al. (2017). "Big Universe, Big Data: Machine Learning and Image Analysis for Astronomy". In: *IEEE Intelligent Systems* 32.2, pp. 16–22. DOI: 10.1109/mis.2017.40.
- Lewis, B. L. et al. (2021). "Distribution and energy balance of Pluto's nitrogen ice, as seen by New Horizons in 2015". In: *Icarus* 356, p. 113633. DOI: 10.1016/j.icarus.2020.113633.
- Li, W., S. Prasad, and J. E. Fowler (2013). "Hyperspectral image classification using Gaussian mixture models and Markov random fields". In: *IEEE Geoscience and Remote Sensing Letters* 11.1, pp. 153– 157.
- Liang, J. et al. (2014). "Automatic registration of multisensor images using an integrated spatial and mutual information (SMI) metric". In: *IEEE Transactions on Geoscience and Remote Sensing* 52.1, pp. 603–615.
- Lowekamp, B. C. et al. (2013). "The design of SimpleITK". In: Frontiers in Neuroinformatics 7, p. 45.
- Mahan, A. I. (1956). "Reflection and refraction at oblique incidence on a dielectric-metallic interface as a boundary value problem in electromagnetic theory". In: *Journal of the Optical Society of America* 46.11, p. 913. DOI: 10.1364/josa.46.000913.

- Maki, J. N. et al. (2020). "The Mars 2020 Engineering Cameras and Microphone on the Perseverance Rover: A Next-Generation Imaging System for Mars Exploration". In: *Space Science Reviews* 216.8. DOI: 10.1007/s11214-020-00765-9.
- Mattes, D. et al. (2001). "Nonrigid multimodality image registration". In: *Medical Imaging 2001: Image Processing*. Vol. 4322. SPIE, pp. 1609–1621.
- Matzke, N. (2006). Wherein I argue emotionally about the definition of "planet". Ed. by P. T. blog. URL: https://pandasthumb.org/archives/2006/08/wherein-i-argue.html.
- McEwen, A. S. et al. (2018). "The Future of MRO/HiRISE". In: *Lunar and Planetary Science Conference*, p. 1301.
- McEwen, A. S. et al. (2007). "Mars Reconnaissance Orbiter's High Resolution Imaging Science Experiment HiRISE". In: *Journal of Geophysical Research* 112.E5. DOI: 10.1029/2005je002605.
- McKinnon, W. B. et al. (2016). "Convection in a volatile nitrogen-ice-rich layer drives Pluto's geological vigour". In: *Nature* 534.7605, pp. 82–85. DOI: 10.1038/nature18289.
- McKinnon, W. B. et al. (2017). "Origin of the Pluto-Charon system: Constraints from the New Horizons flyby". In: *Icarus* 287, pp. 2-ll. DOI: 10.1016/j.icarus.2016.11.019.
- Moore, J. M. et al. (2016). "The geology of Pluto and Charon through the eyes of New Horizons". In: *Science* 351.6279, pp. 1284–1293. DOI: 10.1126/science.aad7055.
- Moore, J. M. et al. (2018). "Bladed terrain on Pluto: Possible origins and evolution". In: *Icarus* 300, pp. 129–144. DOI: 10.1016/j.icarus.2017.08.031.
- Mueller, S. and W. B. McKinnon (1988). "Three-layered models of Ganymede and Callisto: Compositions, structures, and aspects of evolution". In: *Icarus* 76.3, pp. 437–464. DOI: 10.1016/0019– 1035(88)90014–0.
- NASA/JPL (2005). NASA-Funded Scientists Discover Tenth Planet. URL: https://www.jpl.nasa. gov/news/nasa-funded-scientists-discover-tenth-planet.
- (2021). The SPICE concept. Ed. by Navigation and A. I. Facility. URL: https://naif.jpl.nasa. gov/naif/spiceconcept.html.
- Nimmo, F. et al. (2017). "Mean radius and shape of Pluto and Charon from New Horizons images". In: *Icarus* 287, pp. 12–29. DOI: 10.1016/j.icarus.2016.06.027.
- Olkin, C. B. et al. (2017). "The global color of Pluto from New Horizons". In: *The Astronomical Journal* 154.6, p. 258. DOI: 10.3847/1538-3881/aa965b.
- Pedregosa, F. et al. (2011). "Scikit-learn: Machine learning in Python". In: Journal of Machine Learning Research 12.Oct, pp. 2825–2830.
- Philippe, S. (2016). "Microphysique des processus saisonniers des glaces de Mars et Pluton : suivi par télédétection hyperspectrale et étude expérimentale". PhD thesis. Université Grenoble Alpes.
- Plait, P. (2006). Congratulations! It's a planet! Ed. by Slate.com. URL: https://slate.com/technology/ 2006/08/congratulations-its-a-planet.html.
- Pluim, J. P. W., J. B. A. Maintz, and M. A. Viergever (2003). "Mutual-information-based registration of medical images: a survey". In: *IEEE Transactions on Medical Imaging* 22.8, pp. 986–1004.
- Protopapa, S. et al. (2017). "Pluto's global surface composition through pixel-by-pixel Hapke modeling of New Horizons Ralph/LEISA data". In: *Icarus* 287, pp. 218–228. DOI: 10.1016/j.icarus. 2016.11.028.
- Quirico, E. and B. Schmitt (1997a). "A Spectroscopic Study of CO Diluted in N₂ Ice: Applications for Triton and Pluto". In: *Icarus* 128.1, pp. 181–188. DOI: 10.1006/icar.1997.5710.
- (1997b). "Near-Infrared Spectroscopy of Simple Hydrocarbons and Carbon Oxides Diluted in Solid N₂ and as Pure Ices: Implications for Triton and Pluto". In: *Icarus* 127.2, pp. 354–378. DOI: 10.1006/icar.1996.5663.
- Reuter, D. C. et al. (2008). "Ralph: A Visible/Infrared Imager for the New Horizons Pluto/Kuiper Belt Mission". In: *Space Science Reviews* 140.1-4, pp. 129–154. DOI: 10.1007/s11214-008-9375-7.
- Schenk, P. M. et al. (2018). "Basins, fractures and volcanoes: Global cartography and topography of Pluto from New Horizons". In: *Icarus* 314, pp. 400–433. DOI: 10.1016/j.icarus.2018.06.008.

- Schmitt, B. et al. (1998). "Optical properties of ices from UV to infrared". In: *Astrophysics and Space Science Library*. Springer Netherlands, pp. 199–240. DOI: 10.1007/978-94-011-5252-5_9.
- Schmitt, B. et al. (2017a). "Physical state and distribution of materials at the surface of Pluto from New Horizons LEISA imaging spectrometer". In: *Icarus* 287, pp. 229–260. DOI: 10.1016/j.icarus.2016.12.025.
- Schmitt, B. et al. (2017b). SSHADE: "Solid Spectroscopy Hosting Architecture of Databases and Expertise". en. DOI: 10.26302/SSHADE.
- Schroeder, W., L. Ng, and J. Cates (2003). The ITK software guide.
- Silburt, A. et al. (2019). "Lunar crater identification via deep learning". In: *Icarus* 317, pp. 27–38. DOI: 10.1016/j.icarus.2018.06.022.
- Stamnes, K. et al. (2000). DISORT, a General-Purpose Fortran Program for Discrete-Ordinate-Method Radiative Transfer in Scattering and Emitting Layered Media: Documentation of Methodology.
- Stern, S. A., M. W. Buie, and L. M. Trafton (1997). "HST high-resolution images and maps of Pluto". In: *The Astronomical Journal* 113, p. 827. DOI: 10.1086/118304.
- Stern, S. A. et al. (2015). "The Pluto system: Initial results from its exploration by New Horizons". In: *Science* 350.6258, aad1815–aad1815. DOI: 10.1126/science.aad1815.
- Stern, S. A. (2008). "The New Horizons Pluto Kuiper Belt Mission: An Overview with Historical Context". In: *Space Science Reviews* 140.1-4, pp. 3–21. DOI: 10.1007/s11214-007-9295-y.
- Stern, S. A. et al. (2018). "The Pluto system after New Horizons". In: *Annual Review of Astronomy and Astrophysics* 56.1, pp. 357–392. DOI: 10.1146/annurev-astro-081817-051935.
- Telfer, M. W. et al. (2018). "Dunes on Pluto". In: Science 360.6392, pp. 992–997. DOI: 10.1126/science.aao2975.
- Trowbridge, A. J. et al. (2016). "Vigorous convection as the explanation for Pluto's polygonal terrain". In: *Nature* 534.7605, pp. 79–81. DOI: 10.1038/nature18016.
- Uchida, S. (2013). "Image processing and recognition for biological images". In: *Development, Growth* & *Differentiation* 55.4, pp. 523–549. DOI: 10.1111/dgd.12054.
- Vu, T., A. Mishra, and G. Konapala (2018). "Information entropy suggests stronger nonlinear associations between hydro-meteorological variables and ENSO". In: *Entropy* 20.1, p. 38. DOI: 10. 3390/e20010038.
- Weaver, H. A. et al. (2008). "Overview of the New Horizons Science Payload". In: Space Science Reviews 140.1-4, pp. 75–91. DOI: 10.1007/s11214-008-9376-6.
- Weaver, H. A. et al. (2016). "The small satellites of Pluto as observed by New Horizons". In: *Science* 351.6279, aae0030–aae0030. DOI: 10.1126/science.aae0030.
- White, O. L. et al. (2017). "Geological mapping of Sputnik Planitia on Pluto". In: *Icarus* 287, pp. 261–286. DOI: 10.1016/j.icarus.2017.01.011.
- Wood, J. A. et al. (1970). "Lunar anorthosites and a geophysical model of the moon". In: *Proceedings* of the Apollo 11 Lunar Science Conference. Ed. by A. A. Levinson. Vol. 34. 12. Pergamon Press, pp. 1367–1372. DOI: 10.1016/0016-7037(70)90051-7.
- Young, L. A. (2013). "Pluto's seasons: new predictions for New Horizons". In: *The Astrophysical Journal* 766.2, p. L22. DOI: 10.1088/2041-8205/766/2/122.
- Young, L. A. et al. (2008). "New Horizons: anticipated scientific investigations at the Pluto system". In: *Space Science Reviews* 140.1-4, pp. 93–127. DOI: 10.1007/s11214-008-9462-9.
- Zitova, B. and J. Flusser (2003). "Image registration methods: a survey". In: *Image and Vision Computing* 21.11, pp. 977–1000.

Résumé de la thèse en français

Introduction

Les premiers jours de l'exploration spatiale furent à la fois un temps d'abondance et de pénurie. Un investissement financier massif alimenté par deux superpuissances en compétition assurait des révolutions techniques annuelles afin d'atteindre en premier les étapes marquantes de la *conquête spatiale* : envoyer des humains dans l'espace, de plus en plus loin et de plus en plus longtemps, si possible en les ramenant sans danger. Le simple fait d'atteindre la Lune et d'en rapporter des échantillons nous fournit des données sans précédents pour confirmer ou contredire de nombreuses hypothèses sur l'évolution du système solaire. Par exemple, les mesures des ratio d'isotopes des échantillons de régolithe lunaire nous permit d'avoir une meilleure compréhension de la différenciation lunaire (Wood et al., 1970), plus violente et volatile qu'anticipée.

Cependant, ces jours étaient aussi marqués par l'importance seulement secondaire apportée aux questions scientifiques — les buts principaux étant géopolitiques. Les recherches effectuées étaient ainsi presque secondaires au vu des décideurs, même si elles étaient révolutionnaires. Elles ne reposaient que partiellement sur notre capacité naissante à envoyer des machines pour analyser le reste du système solaire et dépendaient plus des observations individuelles des astronautes envoyés observer l'espace. Ces observations furent progressivement remplacées par des missions autonomes : programme Luna des années 1960, puis Mariner et Viking sur Mars. Cependant, les capacités d'analyse et de transmission de ces missions étaient infimes par rapport aux capacités actuelles et les données étaient donc rares. Chaque information supplémentaire était critique et pouvait transformer des champs entiers de recherche, chaque mission permettant d'éclairer nos immenses lacunes, même si les données partielles n'étaient pas capturées avec les instruments idéaux ou en qualité suffisante.

Au fil des décennies, les missions se multiplièrent et la richesse de l'instrumentation explosa en tandem avec les capacités de communication des sondes. Même si cela fait plus de 50 ans qu'un humain n'a pas posé pied sur la lune, les sondes envoyés sur des corps de plus en plus lointains nous donnent aujourd'hui accès à de l'audio-vidéo en haute qualité (Maki et al., 2020), et de la spectroscopie (Baer, 2005), au-delà des capacités humaines. Pour certaines cibles de recherche comme Mars, nous avons aujourd'hui tellement de données qu'il devient presque impossible de l'étudier de manière exhaustive : le MRO a couvert plus de 99% de la surface avec une résolution de 6m par pixel, et une cartographie à 25 cm par pixel est en cours (McEwen et al., 2018).

Nos capacités de traitement des données et d'analyse ont heureusement augmenté avec cette abondance de données et une partie du traitement est désormais automatisé. Il doit cependant être encore guidé par des modèles, qui doivent être raffinés au fur et à mesure que la complexité des bases de données augmente. On voit ainsi de plus en plus d'applications de méthodes spécialisées en informatique, notamment l'apprentissage profond (Kremer et al., 2017; Silburt et al., 2019). Cela rend la collaboration interdisciplinaire avec l'informatique d'autant plus importante. Cette thèse de doctorat est ainsi un essai de collaboration qui vise à démontrer l'intérêt de deux techniques jusqu'ici rares en planétologie et introduit plusieurs pistes d'autres méthodes ayant des applications potentielles dans le domaine.

Chapitre I

Le premier chapitre commence par une présentation des techniques de télédétection et de leurs applications à la planétologie, en particulier en ce qui concerne les méthodes d'imagerie hyperspectrales visant à étudier la composition de la surface.

Les données visuelles issues de la télédétection ont un nombre de dimensions variables en fonction de la nature des capteurs et sont généralement présentées sous la forme d'images. Par exemple, un capteur monochromatique mesure l'intensité lumineuse pour une seule longueur d'onde — ou sur une bande de longueurs d'ondes — et renvoie une matrice où chaque paire de coordonnées (x,y) correspond à un pixel dont la valeur est l'intensité lumineuse. Un appareil photo RVB standard va renvoyer trois images monochromes, correspondant aux intensités sur les bandes rouge, verte et bleue et est donc un capteur multispectral. L'imagerie *hyperspectrale* est une généralisation où le capteur renvoie une série d'images monochromes, chacune associée à une longueur d'onde dans un large spectre continu. On obtient ainsi un *cube de données hyperspectrales* en trois dimensions, deux spatiales et une spectrale.

Chaque série continue de mesures selon une bande de fréquence correspond donc à un spectre, en sachant que des spectres de types différents sont étudiés en astronomie, géologie et sciences planétaires. L'astronomie stellaire utilise fréquemment la spectroscopie d'émission, l'objet d'étude étant une source intense de radiation électromagnétique dont l'intensité relative est caractéristique de la composition. Au contraire, les corps planétaires émettent beaucoup moins de radiation et sur des longueur d'ondes plus élevées (infrarouge thermique), ce qui limite l'utilité de ces données dans l'étude de la composition de ces corps. Les méthodes qui nous intéressent se focalisent donc sur l'analyse de la lumière émise par d'autres sources (comme le soleil) après son interaction avec la surface planétaire. C'est le domaine de la spectroscopie d'absorption et de réflectance.

Quand un rayon interagit avec un objet, il peut être réfléchi, absorbé ou diffusé, selon la longueur d'onde et les caractéristiques physiques de l'objet. Si ce dernier est granulaire — ce que l'on peut supposer être le cas de l'immense majorité des objets selon Hapke (2012) — alors les propriétés seront dépendantes de la taille et de l'organisation structurelle des grains, en plus de leur composition moléculaire.

Si l'on considère la variation de réflectance pour un grain donné en fonction de la longueur d'onde, on se retrouve avec un *spectre de réflectance*, généralement obtenu par une série de mesure. Ce spectre consiste en des séquences alternantes de *bandes d'absorption* et de *continuum*. Ces *bandes d'absorption* correspondent à des chutes d'intensité sur de petite bandes de fréquence, dont la forme et la profondeur dépendent du mécanisme d'absorption par le matériau. Ces bandes sont quantifiées en fonction de la position du centre de la bande, de sa largeur, de sa forme et de son aire.

Le rapport entre le spectre de réflectance d'une surface et ses propriétés physiques est gouverné par l'équation de transfert radiatif. Cette équation établit que lorsqu'un rayon traverse un milieu, il perd de l'énergie par absorption, en gagne par émission du milieu et la redistribue le long du spectre par diffusion. Hélas, trouver une solution analytique aux équations de transfert radiatifs forme un problème à priori insoluble avec les méthodes mathématiques contemporaines. Des outils numériques et des méthodes d'approximation sont donc nécessaires.

La première étape est de formuler des *modèles de transfert radiatif* (MTR) qui consistent en un ensemble d'hypothèses simplificatrices (ainsi que de mesures empiriques). Une source de complexité est que les surfaces planétaires ont de nombreuses composantes moléculaires avec des arrangements et mélanges variés, dont la structure affecte le spectre de réflectance, ainsi que l'albédo de diffusion simple ω et la profondeur optique τ . Les hypothèses faites par les modèles concernent déjà les *mélanges* dans la surface, qui sont simplifiés sous la forme de 4 modes principaux :

- Mélange spatial : les composants sont séparés mais à une échelle sub-pixel et ne peuvent donc pas être isolés comme pixels individuels. Le spectre de réflectance est donc une combinaison linéaire des différents composants, pondérés par leur proportion spatiale. Chaque composant peut de plus être lui-même un mélange d'un autre type ci-dessous.
- Mélange granulaire : les grains de plusieurs composants sont mélangés au sein d'une seule couche, donnant lieu à un spectre final correspondant à un mélange non-linéaire des spectres de chaque composant.
- Mélange moléculaire : les grains sont eux-mêmes composés de mélanges de plusieurs molécules, comme par exemple les clathrates de gaz. Les réflectances de ces mélanges ne sont pas directement calculables à partir des réflectances de composantes pures : les liaisons intermoléculaires ont leurs propres fréquences et intensités de vibration, ce qui peut atténuer ou déplacer les bandes spectrales, voire en créer de nouvelles.
- Stratification : les composants sont superposés en couches horizontales, chaque couche pouvant être elle-même un mélange des types ci-dessus. L'évolution du spectre selon l'épaisseur

de la couche n'est pas linéaire. Quand la couche supérieure est fine, suffisamment de photons la traversent pour que les bandes caractéristiques des spectres des couches inférieures transparaissent. Au contraire, si la couche supérieure est suffisamment épaisse, les couches inférieures ont un impact négligeable sur le modèle. La stratification est utile pour représenter des scénarios communs dans les environnements complexes, comme lors de la sublimation d'un composant qui forme une couche de quelques microns au dessus d'un mélange granulaire ou moléculaire.

En plus des effets non-linéaires de ces mélanges sur les spectres, il faut aussi considérer l'impact de la taille des grains et de l'anisotropie. Cela donne donc des MTR intégrant des mélanges de mélanges, et donc un nombre de paramètres indépendants très élevé (plus de 50). La résolution numérique de ces modèles est complexe et fait l'objet du chapitre III.

La dernière partie de ce chapitre décrit la mission New Horizons de la NASA ayant pour but d'observer Pluton. Cette planète naine avait jusqu'ici reçu très peu d'attention scientifique à cause de la pénurie de données à son sujet. La source principale avant cette mission était le télescope Hubble, qui ne permettait pas d'avoir une résolution spatiale sur Pluton. Tout essai de cartographie reposait donc sur des méthodes indirectes d'analyse de son spectre au cours du temps. Le survol de 2015 permit de changer complètement nos méthodes en fournissant une multitude de données permettant d'augmenter notre connaissance de la surface de Pluton : sa topographie, sa géologie et sa composition. Les données qui nous concernent viennent principalement de deux instruments, l'imageur multi-bandes visible et infra-rouge (MVIC) et l'imageur spectromètre proche infrarouge (LEISA). Les longueurs d'ondes plus grandes que 1.1 μ m sont traitées par LEISA et le reste par MVIC. Ce dernier permet d'obtenir une résolution spatiale maximale de 0.66 km/pixel pour l'hémisphère survolé, alors que LEISA a un cône de vision trois fois plus large et fournit des données spectrales importantes dans la bande 1.25–2.5 μ m.

Une grande partie de ces données a déjà été traitée et analysée et plusieurs cartes ont pu être produites, permettant de montrer les variations de la surface de Pluton. Cependant, il reste encore du travail à accomplir et nous ne disposons pas encore d'une carte globale de la composition superficielle qui utiliserait LEISA. De plus, les analyses compositionnelles ont jusqu'ici principalement été menées de manière qualitative.

Chapitre II

Les travaux de Schmitt et al. (2017a), basés sur les données en haute résolution de LEISA, offrirent une analyse qualitative de la distribution spatiale des différents composants présents à la surface de Pluton. Les données hyperspectrales furent traitées en utilisant l'analyse en composantes principales, afin de calculer la profondeur de bande intégrée du méthane et de l'azote solide, ainsi que les indices spectraux correspondant à la présence de glace d'eau et de matière rouge.

En plus des données en haute résolution, nous possédons aussi des images acquises durant la phase d'approche de la mission, avec des informations sur la totalité de la partie éclairée de la surface de Pluton (avec les 30° du sud en hiver polaire). En procédant à rebours depuis les images en haute résolution, chaque image précédente correspond à une longitude centrale légèrement supérieure, par l'effet de la période de rotation de 6.4 jours. Combiner ces différentes images nous permet d'obtenir une première carte globale, mais les imprécisions dues aux instruments font qu'il y a un décalage de plusieurs pixels entre les différentes séries de données, correspondant parfois à quelques degrés de longitude/latitude. Ce chapitre se consacre au traitement de ces données afin de résoudre ce décalage.

La création d'une carte unique à partir de plusieurs sources s'appelle le mosaïquage et le processus part généralement d'une image bien calée sur laquelle on aligne ensuite les autres sources de données. Selon les disparités entre les sources, il peut être suffisant de décaler linéairement ou de pivoter certaines sources, mais il est parfois nécessaire d'utiliser des homothéties, voire des transformations non-linéaires et des homomorphismes locaux. Deux méthodes principales d'alignement existent. La première se base sur certains éléments distinctifs connus (points, lignes, contours) et est la plus fréquemment utilisée dans le domaine (Zitova and Flusser, 2003). Elle demande cependant un travail intense car elle doit être effectuée à la main en cherchant les meilleurs éléments distinctifs aux frontières nettes (bords de cratères, falaises). À cause de la basse résolution des images d'approche, cette méthode n'est pas efficiente, ce qui nécessite de se tourner vers des techniques provenant de l'imagerie médicale pour produire les cartes globales.

Le deuxième type de méthode se base sur les régularités des motifs d'intensité des images. Chaque image est caractérisée par un arrangement spatial de zones sombres et claires. La comparaison des distributions correspondantes permet d'identifier des transformations potentielles permettant de rendre les images plus similaires. Une première étape est de choisir la mesure de corrélation appropriée (appelée *métrique*), ce choix dépendant de la nature des données.

En raison de l'origine de cette méthode et de son utilisation principale, la plupart des outils logiciels sont optimisés pour les images médicales mais ils peuvent facilement être adaptés à la planétologie et ont déjà été utilisés à cet escient (Liang et al., 2014, see). Pour les décalages nous concernant, il y a des imprécisions dues principalement à une erreur de visée de l'instrument et donc une translation, assortie potentiellement d'une légère rotation. On peut donc se restreindre à ce type de transformations globales simples, en normalisant les données pour être toujours dans la projection orthographique originale.

Au final, nous avons sélectionné et traité 12 images d'approche toutes prises lors des cinq jours d'approche avant la prise des images en haute résolution. Cette période est la plus petite qui permet d'obtenir une couverture globale et nous avons choisi d'ignorer trois images aussi prises lors de cette période car elles étaient redondantes ou souffraient de problèmes de calibration spectrophotométrique. Nous avons utilisé les données en haute résolution quand c'était possible, mais la dégradation entre les zones en haute (30km/px) et basse (394km/px) résolution est visible.

Une fois la transformation calculée, elle peut être appliquée uniformément à toutes les longueur d'onde de l'image hyperspectrales. Cela nous permet de produire les premières mosaïques calées à partir de tout sous-ensemble de données de LEISA. Nous avons produit 4 cartes principales, en se basant sur les indices spectraux définis par Schmitt et al. (2017a) :

CH₄: bande de CH₄ de profondeur intégrée entre 1.589 et 1.833 μ m, pour un groupe de 3 bandes de CH₄ à 1.67, 1.72 et 1.79 μ m.

$$BD(CH_4) = 1 - \frac{\int_{1.589\mu m}^{1.833\mu m} RF(\lambda)d\lambda}{\int_{1.589\mu m}^{1.833\mu m} Cont(\lambda)d\lambda}$$
(IV.1)

N₂: bande de N₂ de profondeur intégrée entre 2.12l et 2.160 μm pour la bande de N₂ à 2.15 μm. Cette intégration a été légèrement étendue vers les bassses longueurs d'ondes par rapport à Schmitt et al. (2017a). Nous nous sommes rendu compte que cette bande contribue jusqu'à 2.12 μm à la forte bande de CH₄ de 2.20 μm, ce qui améliore légèrement le rapport signal/bruit.

$$BD(N_2) = 1 - \frac{\int_{2.105\mu m}^{2.160\mu m} RF(\lambda) d\lambda}{1.5 * (RF_{2.105\mu m} + RF_{2.13\mu m} + RF_{2.1675\mu m} + RF_{2.1755\mu m})}$$
(IV.2)

2100

 H_2O : indicateur spectral utilisant 10 bandes centrées autour de 1.39 et 2.06 μ m.

$$SI(H_2O) = 1 - \frac{\int_{2.022\mu m}^{2.090\mu m} RF(\lambda)d\lambda}{\int_{1.365\mu m}^{1.410\mu m} RF(\lambda)d\lambda}$$
(IV.3)

Matière rouge: indicateur spectral utilisant les longueurs d'ondes autour de 1.430 et 1.658 μ m.

$$SI(RedMat) = 1 - \frac{\int_{1.447\mu m}^{1.421\mu m} RF(\lambda) d\lambda}{\int_{1.670\mu m}^{1.641\mu m} RF(\lambda) d\lambda}$$
(IV.4)

Après avoir compilé ces quatre cartes globales ainsi que deux cartes auxiliaires, nous avons étudié les valeurs globales sur ces cartes et les conséquences sur les variations de la composition superficielle de Pluton. Nous avons aussi analysé les corrélations entre ces cartes ainsi qu'avec les autres cartes de LEISA et les autres cartes produites par la mission New Horizons, telles que la carte panchromatique visible, la carte d'altitude et la carte en IR-proche du MVIC, et en avons interprété ces résultats afin d'améliorer notre compréhension de la composition superficielle. Nous avons notamment pu confirmer la variation de la composition latitudinale proposée par Protopapa et al. (2017) et la catégorisation du méthane et des terrains de matière rouge proposée par Earle et al. (2018) à partir de données MVIC. Ces données nous ont aussi permis de cartographier la présence globale de glaces d'eau et d'azote solide.

En se basant sur ces cartes de composition et les corrélations avec les connaissances topographiques et géologiques, nous proposons aussi des interprétations géologiques d'aspects compositionnels de l'autre hémisphère :

- la présence de terrains érodés avec des glaces de méthanes dans les hautes latitudes,
- une indication d'une bande équatoriale quasi-continue de terrains à lames riches en méthane, ponctuée de plaines riches en azote solide,
- la présence de glaces d'eau à des longitudes en dehors de la plaine Spoutnik, correspondant à un substrat exposé.

Le travail présenté dans ce chapitre est une des premières utilisations de cette méthode en cartographie planétaire et montre le potentiel de l'outil, avec une précision allant au-delà du pixel (précision confirmée par traitement manuel à partir des données en haute résolution). Cette méthode reste efficace même avec des résolutions de 158 km/px en absence d'éléments distinctifs. Elle ne peut pas être utilisée telle quelle sur des résolutions si basses mais en la combinant à des contraintes de bon sens (bornant les décalages possibles), il devient possible de créer un algorithme automatique de recalage à utilisations multiples. Sur un corps dont la géologie est aussi volatile que Pluton, dont une fraction considérable se sublime et circule à un rythme saisonnier (Bertrand et al., 2018, 2019), connaître la distribution de la composition superficielle est un élément critique à la fois pour comprendre son histoire et pour prédire son futur. Certaines des cartes produites ont d'ailleurs déjà été utilisées afin de contraindre les modèles atmosphériques de Pluton (Lewis et al., 2021).

Chapitre III

Après avoir obtenu des données qualitatives, l'étape suivante est d'essayer d'obtenir des analyses quantitatives, en terme de proportion des différents composants à la surface et de leur modes de mélange. Ce troisième chapitre présente une méthode méta-heuristique d'inversion spectrale, qui utilise l'algorithme de recuit simulé afin de trouver efficacement une composition candidate dont le spectre correspond au spectre observé sur un pixel donné. Nous utilisons le modèle de transfert radiatif DISORT, qui nous permet de prendre en compte les modèles stratifiés qui jouent très probablement un rôle à la surface d'un corps exhibant du transport volatil saisonnier. Nous commençons par calibrer la méthode sur des spectres synthétiques avant de l'appliquer à plusieurs régions cibles de la surface de Pluton et fournissons ainsi les premiers résultats quantitatifs sur la composition superficielle.

Selon le modèle utilisé et les hypothèses simplificatrices sur la représentation superficielle, l'inversion de spectre est un problème d'optimisation dont l'espace de solution va de quelques dimensions à plusieurs douzaines de dimensions. Les travaux de Protopapa et al. (2017) avaient simplifié le problème en utilisant un modèle de Hapke à quatre composantes chimiques dans un mélange spatial. Ce problème à huit dimensions était coûteux mais solvable en utilisant des algorithmes de fitting (comme l'algorithme de Levenberg-Marquardt, qui ne converge cependant que vers un minimum local). Cela permit d'avoir des cartes de zones de composition globale, mais ne permet pas d'avoir de résultats précis, notamment sur la plaine Spoutnik où leur modèle n'arrive pas à reproduire la bande N₂, les auteurs suggérant donc d'utiliser une stratification dans un modèle plus complet. L'objectif de ce chapitre est triple : formaliser le problème de l'inversion de spectre dans un modèle complexe à plusieurs couches et modes de mélanges, étudier la régularité de l'espace de paramètres et la multiplicité des solutions et enfin déterminer le coût calculatoire et la faisabilité de cette méthode dans le but d'obtenir une carte quantitative globale de la composition superficielle de Pluton.

Afin de représenter avec précision la variabilité de la surface de Pluton, correspondant à la diversité de spectres observés, le modèle utilisé requiert un grand nombre de paramètres. Tout d'abord, il y a les composants élémentaires sur le plan moléculaire : des glaces d'azote, de méthane, de monoxyde de carbone et d'eau, ainsi que la matière organique rouge. Les trois premiers peuvent être présents purs ou au sein de mélanges solides, avec une proportion variable de méthane. Chacun de ces composants a aussi plusieurs propriétés utilisées par le MTR : taille de grains, anisotropie et proportion moléculaire (à part pour les composants purs). Enfin, ces composants peuvent être mélangés spatialement, granulairement ou en strates, ce qui introduit des paramètres de proportion granulaire et spatiale et d'épaisseur de couches. Au total, l'espace de solutions a approximativement 60 dimensions, avec certains paramètres variant entre 0 et 1, et d'autres couvrant plusieurs ordres de grandeur (comme la taille des grains, qui varie entre quelques nanomètres et quelques mètres).

Le MTR utilisé est DISORT, écrit au départ en Fortran 77 et plus récemment adapté en C par T. Dowling Buras et al. (2011). Il est conçu pour modéliser des couches verticales hétérogènes et pour donner des spectres précis et robustes, en contrepartie d'un coût de calcul élevé. Il sert donc souvent d'étalon pour évaluer la précision d'autres MTR plus efficaces. Après des soucis d'instabilité numérique avec Spectrimag, nous avons décidé d'utiliser DISORT directement plutôt que de trouver un autre compromis plus efficace mais moins robuste.

DISORT est un outil de calcul de spectre à partir de données compositionnelles mais nous voulons faire une inversion et retrouver la composition à partir du spectre. Ce problème souffre, hélas, d'une explosion combinatoire et aucun algorithme existant ne permet de le résoudre précisément. Il existe heureusement un ensemble de méthodes permettant de faire de l'inversion sur des problèmes abstraits, et les méta-heuristiques sont un outil privilégié du domaine quand on cherche une solution suffisamment proche de l'optimum en temps raisonnable. Là où des heuristiques spécifiques sont optimisées sur des problèmes donnés (voire des configurations précises), les méta-heuristiques ont un caractère plus général et ne font pas d'hypothèse sur la structure de l'espace de solutions, qu'elles explorent de manière stochastique.

La méta-heuristique choisie s'appelle le recuit simulé, qui procède en opérant une optimisation locale tout en maintenant une certaine probabilité appelée température — qui diminue avec le temps — d'évoluer vers des solutions plus coûteuses. Cela permet d'éviter de se retrouver bloqué dans un optimum local qui n'est pas l'optimum global. Des théorèmes généraux permettent de prouver que, à condition de bien choisir la fonction d'évolution de la température, l'algorithme converge vers la solution optimale avec probabilité 1 (Granville et al., 1994). Cependant, cela requiert généralement la connaissance de propriétés sur l'espace des solutions. Idéalement, la fonction du spectre dans l'espace de solutions pourrait être étudiée analytiquement pour avoir des bornes sur ses gradients. Cependant, le modèle DISORT n'étant pas analytique lui-même, cela prendrait au minimum plusieurs années de travail mathématique complexe. Afin de trouver une bonne fonction d'évolution de la température, nous avons donc utilisé une fonction standard dans la littérature, en optimisant ses paramètres à partir d'essais en parallèle pour trouver le plus efficace.

En pratique, on commence l'algorithme avec une composition aléatoire et on calcule le spectre correspondant. On modifie alors légèrement les paramètres de la composition et on calcule un deuxième spectre. Si ce deuxième spectre est plus proche de la cible, on prend la deuxième composition comme base et on itère. Sinon, on choisit au hasard entre la première et la deuxième composition, où la probabilité de prendre la deuxième est d'autant plus faible que son spectre est pire que celui du premier. Plus l'algorithme itère, plus cette probabilité de choisir une solution moins bonne devient faible. La question de comment décider quel spectre est meilleur (et à quel point) dépend de la métrique utilisée, et donc aussi de la phase de l'algorithme. Notre méthode standard sera de comparer la racine de l'écart quadratique moyen entre les deux spectres.

La version simple du recuit simulé perturbe tous les paramètres de manière linéaire à chaque itération. Cela serait raisonnable s'ils avaient une influence linéaire sur le spectre mais ce n'est pas le cas. Par exemple, l'anisotropie affecte le spectre dans sa totalité de manière quasi-uniforme en appliquant un facteur multiplicatif sur l'intensité. Au contraire, la taille des grains n'affecte généra-lement le spectre que sur quelques bandes étroites. Cela signifie que perturber tous les paramètres à la fois est à priori inefficace : la perturbation d'un paramètre global rend caduque toutes les optimi-sations sur les paramètres ayant des impacts locaux. Une solution est de séparer les paramètres en sous-ensembles en fonction de la localité de leurs effets. Une autre option est de calculer la dérivé des spectres, ce qui nous permet d'évaluer les changements sur la forme du spectre en ignorant les facteurs uniformes non-locaux comme ceux induits par l'anisotropie. L'algorithme que nous proposons a donc trois phases qui nous permettent d'optimiser à la fois les paramètres locaux et globaux sans que l'une des catégories n'écrase l'autre :

- 1. Fitting de forme : on perturbe l'ensemble des paramètres et calcule la racine de l'écart quadratique moyen (REQM) entre la dérivée du spectre et la dérivée du spectre cible.
- 2. Fitting d'amplitude : On ne perturbe que le facteur d'anisotropie ξ et on calcule la REQM entre le spectre calculé et le spectre cible.
3. Fitting global : On perturbe tous les paramètres sauf l'anisotropie et on calcule la REQM entre le spectre calculé et le spectre cible.

Notre capacité à calculer des spectres synthétiques correspondant aux données réelles dépend de deux facteurs. Premièrement, existe-t-il des paramètres tels que notre modèle de représentation arrive à produire un spectre approprié ? Deuxièmement, l'algorithme est-il capable de trouver efficacement cet ensemble de paramètres ? La première condition n'est pas à priori testable en avance mais on peut évaluer la deuxième en créant des spectres synthétiques et en cherchant à résoudre ces problèmes artificiels d'inversion (en supposant donc que les spectres synthétiques sont suffisamment représentatifs des spectres réels).

Durant la phase de calibration des paramètres, nous avons observé que nous n'arrivions pas à obtenir des REQM inférieures à quelques pourcent sur les spectres complexes. Cela peut être dû à un temps de calcul insuffisant pour la convergence ou à une limite inhérente de l'algorithme sur la résolution des cas en très haute dimension. Nous avons testé cela, mais même en augmentant le nombre d'itérations, la qualité du spectre augmente légèrement alors que la REQM médiane stagne avec un plateau de 4-5%. Cela suggère que l'algorithme ne peut pas converger au-delà de cette limite sur des spectres complexes. La meilleure REQM obtenue sur l'ensemble des phases de calibration était de 0.6% mais une telle performance est impossible à garantir et requiert des centaines d'exécutions (chacune ayant des milliers d'itérations).

Une fois effectué le choix des paramètres, validés par expérience synthétique, nous avons utilisé l'algorithme sur les données LEISA. Nous avons sélectionné plusieurs groupes de pixels correspondant à des "endmembers", dont une analyse qualitative semble indiquer qu'ils ont une composition relativement simple avec un ou deux composants. À cause de contraintes temporelles, nous n'avons pu terminer le traitement que de deux de ces endmembers : "Spoutnik" et "Pôle Nord".

L'endmember du pôle nord est situé à 69.98°N, 63.3°E et avait été classé initialement comme une glace d'azote à gros grains et haute teneur en méthane. Nos calculs montrent l'existence d'une couche de méthane pure de 4-5 grains d'épaisseur (5-6 mm), en accord avec les hypothèses formulées dans Schmitt et al. (2017a). Nous montrons aussi que la présence de glace d'eau correspond probablement à une partie exposée du substrat, l'eau étant beaucoup moins volatile que les autres glaces. Vu la dynamique complexe de la surface de Pluton, on ne peut pas supposer avoir un équilibre thermodynamique et cet état observé peut être transitoire. Le cycle saisonnier du méthane étudié par Bertrand et al. (2019) nous situe en printemps avec la sublimation de l'azote, en accord avec nos observations. Les travaux sur Spoutnik mènent à des résultats similaires.

Après avoir obtenu des compositions pour les endmembers, nous avons étendu la tâche avec une tentative d'inversion d'une série de 14 points correspondant à une ligne de mélange spectrale, évoluant continûment entre les deux extrémités. Celles-ci sont situées à Voyager Terra autour de 65°N et entre 152 et 156°E. Ce terrain avait été décrit par Moore et al. (2016) comme composé de gros grains d'azote solide avec une faible proportion de méthane, potentiellement mélangés à de la glace d'eau (Schmitt et al., 2017a). Nos calculs confirment aussi ces observations initiales : on voit des grains épais ($d \approx 1 - 3$ cm) de glace riche en azote, avec une proportion de méthane variant entre 0.01% et 1% au long de la série de points. Sur la partie riche en méthane, on voit aussi une couche fine de méthane (au plus 2 grains d'épaisseur), qui diminue en taille des grains vers le sudouest. Cette couche est progressivement remplacée par un gradient de concentration en méthane au sein de la solution solide d'azote, ce qu'on peut expliquer par une sublimation en cours de cet azote solide, qui augmente la concentration du méthane dans la glace subsistante sans pour autant former une couche de méthane pur.

Le travail effectué au sein de ce chapitre sur l'inversion de spectre est satisfaisant quant à la qualité des compositions, mais a requis une quantité trop élevée d'interventions humaines afin d'ajuster les paramètres et de trouver le meilleur modèle pour chaque pixel, en plus d'un coût de calcul prohibitif. Cela formerait une quantité de travail trop élevée pour l'étendre tel quel à une carte globale et il faut donc automatiser ce processus et améliorer l'efficacité des algorithmes, ce qui est l'objet du chapitre suivant.

Tous les calculs de ce chapitre ont été effectués sur un cluster appartenant au GRICAD (Grenoble Alpes Recherche - Infrastructure de Calcul Intensif et de Données) du CNRS. Le cluster utilisé avait 190 noeuds, chacun composé de 2 Intel Sandy Bridge EP E5-2670 à 8 coeurs et 64GO de RAM.

Chapitre IV

Le dernier chapitre de cette thèse étudie les coûts calculatoires liés à la production d'une carte globale de composition. Cela inclut des estimations (hautes et basses) et des pistes de réflexion sur les manières de résoudre les obstacles présentés à court et moyen terme.

La carte de Pluton basée sur les données de LEISA contient autour de 36 millions de pixels nonvides, correspondant à un suréchantillonage non-négligeable sur de nombreuses zones. Nous estimons le nombre de pixels uniques à 3.75 millions une fois ce suréchantillonage éliminé. En utilisant la calibration sur les spectres synthétiques et les endmembers pour avoir une idée du temps de calcul maximum et minimum, on obtient dans le meilleur des cas 48 minutes de calcul par spectre (chaque calcul étant fait sur un noeud du cluster). Même en ayant un accès exclusif à la totalité des noeuds du cluster, cela requerrait près de deux ans (342.5 ans avec un seul noeud). Cela n'est pas réaliste au sein d'un projet de thèse (financièrement, écologiquement ou humainement), surtout que cette estimation se base sur le meilleur cas observé, le pire cas requérant un siècle d'utilisation continue du cluster entier. Même en ayant accès aux supercalculateurs les plus puissants existants, le temps de calcul requis se chiffre en mois dans le pire cas.

En l'état, nous proposons donc d'étudier plusieurs pistes qui permettraient la production d'une carte globale en quelques années. Une première piste est d'utiliser les plateformes de calcul citoyen, comme pour les projets de recherche de nombres premiers, de modélisation climatique ou de pliage de protéines. Certains de ces projets ont duré plus de 20 ans, avec des pics de puissance de calcul atteignant 2.43 exaFLOPs (de quoi résoudre le pire cas en moins de deux jours). La plupart des projets de ce type n'atteignent pas ce niveau, mais restent comparables en puissance à la plupart des supercalculateurs nationaux. Lancer un tel projet est surtout plus facile et moins coûteux, et plus flexible en terme d'extension de délai ou de besoins flous en ressources. La priorité dans cette direction serait de documenter le code de l'algorithme d'inversion de spectre et de le distribuer en logiciel libre. Il faudrait ensuite ou bien utiliser une plateforme de science citoyenne pré-existante (comme BOINC), ou bien créer un projet ad-hoc, qui pourrait offrir une plus grande latitude dans l'évolution. Cela permettrait de profiter d'améliorations algorithmiques mais surtout d'une ludification du processus, en profitant de l'intérêt public fort pour tout ce qui touche au spatial (et la publicité qu'offre l'affiliation aux institutions françaises et à la NASA).

Nous voyons enfin deux potentielles améliorations algorithmiques. La première se base sur l'utilisation des spectres intermédiaires produits lors d'une inversion de spectre. La génération du spectre est la partie la plus coûteuse du processus (au moins 10000 fois plus coûteux que les calculs de REQM). Certains de ces spectres pourraient être des solutions (ou des pistes de solutions) pour d'autres spectres à inverser. Si on veut que les comparaisons entre spectres ne prennent pas plus de 10% du coût total, cela nous permet de comparer chaque spectre synthétique à plus de 1000 spectres réels. En proportion et sans amélioration, cela n'est pas directement utile. Nous proposons tout de même une première fonction de hachage (non cryptographique) qui transforme chaque spectre afin de lui assigner une position dans un arbre de décision. À partir de cet arbre, à chaque fois qu'on calcule un nouveau spectre on peut efficacement trouver l'élément approprié dans l'arbre, correspondant à chaque fois à une listes des seuls pixels candidats (dont le spectre est donc proche du spectre initial). Selon nos calculs, il doit être possible d'obtenir en moyenne quelques dizaines de candidats et cette méthode peut encore être grandement raffinée (avant d'être validée empiriquement).

La deuxième piste algorithmique est basée sur un prétraitement. Le modèle le plus complexe que nous autorisons dans l'inversion de spectre a jusqu'à 60 dimensions quand tous les composants sont présents. Cependant, une majorité des zones n'a probablement pas besoin de modèles si complexes et nous pouvons potentiellement augmenter l'efficacité en limitant la complexité autorisée sur ces zones. Nous avions brièvement mentionné au chapitre II la production de cartes classifiées de Pluton en utilisant un modèle de mélanges Gaussiens, qui permettait de créer un ensemble de clusters de données correspondant à des propriétés corrélées. Nous pouvons désormais utiliser ces clusters pour simplifier le calcul d'inversion de spectre. Si un pixel est proche d'un autre pixel appartenant au même cluster et pour lequel nous avons déjà le spectre, on peut choisir d'utiliser le même modèle (tout en modifiant les paramètres).

Cette méthode requiert un certain nombre de pixels initiaux qui doivent être traités en amont afin de servir de points de comparaisons aux autres pixels. Elle n'a pas non plus vocation à être universelle : elle ne marchera efficacement que tant que l'on n'est pas sur une frontière. Vu le bruit inhérent des spectres inversés et la grande sensibilité au bruit du MTR, il est possible qu'il n'y ait aucun pixel candidat parmi les pixels initiaux, tout comme il est possible qu'il y en ait plusieurs. Dans cette situation, on peut utiliser une variante des algorithmes de coloration de graphe pour résoudre les ambiguïtés, cette fois en utilisant des outils de programmation linéaire.

Comme établi au début de ce chapitre, produire une carte globale de composition aurait été ambitieux et coûteux dans le meilleur des cas. Au cours de cette thèse, nous avons cependant posé les bases de ce projet cartographique, en étendant les cartes qualitatives de Schmitt et al. (2017a) afin d'obtenir une base de donnée hyperspectrale globale, tout en effectuant des analyses et des interprétations géologiques sur ces données. Nous avons ensuite développé une méthode métaheuristique afin de faire de l'inversion de spectre sur un modèle complexe de surface, dont nous avons évalué l'efficacité et la fidélité sur des données synthétiques. Nous avons enfin utilisé cette méthode pour analyser certaines régions cibles de Pluton, comme la plaine Spoutnik, où la complexité du terrain empêche une inversion précise par des modèles moins complets. Tout au long de cette thèse, nous avons implémenté des techniques d'analyses de données qui sont connues et fréquemment utilisées dans d'autres domaines (comme le recalage d'images basé sur l'intensité, utilisé en imagerie médicale), alors même qu'elles sont souvent ignorées en planétologie. Le développement continu des instruments de capture de données spatiales nous emmène inéluctablement vers un monde où les données sont riches et requièrent des outils d'autant plus complexes afin de les analyser rigoureusement. Nous espérons que nous aurons l'occasion de poursuivre ce projet en renforçant l'interdisciplinarité en informatique théorique et planétologie, et en développant l'aspect de science citoyenne.