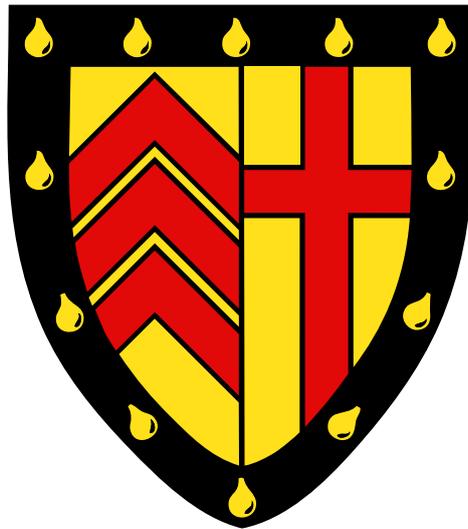


Investigating the Characteristics of Exoplanetary Atmospheres and Interiors



Matthew Conor Nixon

Supervisor: Prof. Nikku Madhusudhan

Institute of Astronomy
University of Cambridge

This thesis is submitted for the degree of
Doctor of Philosophy

Clare College

July 2022

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. I further state that no substantial part of my thesis has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Much of the work presented in this thesis has been published or accepted for publication. The papers from which this thesis draws material are as follows:

- **Chapter 3.** Published in *Monthly Notices of the Royal Astronomical Society* as **Assessment of supervised machine learning for atmospheric retrieval of exoplanets** Nixon & Madhusudhan (2020)
- **Chapter 4.** Accepted for publication in *The Astrophysical Journal* as **AURA-3D: A Three-dimensional Atmospheric Retrieval Framework for Exoplanet Transmission Spectra** Nixon & Madhusudhan (2022)
- **Chapter 5.** Published in *Monthly Notices of the Royal Astronomical Society* as **How deep is the ocean? Exploring the phase structure of water-rich sub-Neptunes** Nixon & Madhusudhan (2021)
- **Chapter 6.** Contains material published in *The Astrophysical Journal Letters* as **The Interior and Atmosphere of the Habitable-zone Exoplanet K2-18b** Madhusudhan, Nixon, Welbanks, Piette, and Booth (2020). and material published in *Astronomy & Astrophysics* as **A planetary system with two transiting mini-Neptunes near the radius valley transition around the bright M dwarf TOI-776** Luque, Serrano, Molaverdikhani, Nixon, Livingston et al. (2021)

The work presented in Chapter 2 uses the AURA retrieval code (Pinhas et al., 2018), which was made available to me in my research group, led by Nikku Madhusudhan. Elements of

this code were adapted for use in Chapters 3 and 4. In Chapter 5, temperature profiles were generated for comparison to my work by Anjali Piette, using her adaptation of the GENESIS code (Gandhi & Madhusudhan, 2017, Piette & Madhusudhan, 2020b). The work presented in Chapter 6 was conducted in collaboration with the co-authors of (Madhusudhan et al., 2020) and (Luque et al., 2021) as listed above. The text and figures which appear in this chapter are adapted from my contributions to these publications. I provided the equation of state for H₂O described in Chapter 5 for the interior models of K2-18b, and later incorporated this equation of state into my own internal structure modelling code which I used throughout Chapter 5 and for the analysis of TOI-776b and c in Chapter 6. Published and submitted works incorporate the comments of coauthors. The use of ‘we’ instead of ‘I’ throughout most of this thesis is a stylistic choice. The spectrum of WASP-96b presented in chapter 7 was prepared by a team led by Néstor Espinoza.

Matthew Conor Nixon

July 2022

Abstract

Investigating the Characteristics of Exoplanetary Atmospheres and Interiors

Matthew Conor Nixon

The characterisation of exoplanets has made rapid progress in recent years, with observations of bulk properties such as mass and radius combining with detailed atmospheric spectroscopy to provide unprecedented insight into the nature of these remote worlds. However, these high-quality observations also require sophisticated modelling and analysis tools in order to maximise the scientific output from the data. In this thesis I present a number of advances in atmospheric modelling and retrieval, as well as internal structure models, which have been used to investigate the properties of a wide range of planets, from hot Jupiters to temperate mini-Neptunes.

I conduct an assessment of the feasibility of supervised machine learning as a tool to carry out atmospheric retrievals of exoplanets. Retrieval methods commonly conduct Bayesian parameter estimation and statistical inference using sampling algorithms such as Markov Chain Monte Carlo or Nested Sampling. Recently several attempts have been made to use machine learning algorithms either to complement or replace fully Bayesian methods in order to improve computational efficiency. However, results from these algorithms sometimes disagree with contemporary Bayesian retrievals. To investigate this, I use the Random Forest supervised machine learning algorithm which has been applied previously for atmospheric retrieval. I extend the machine learning approach to develop a new algorithm, and demonstrate excellent agreement with a Bayesian retrieval of the transmission spectrum of the hot Jupiter HD 209458b. Despite this success, and achieving high computational efficiency, I still find that this machine learning approach is computationally prohibitive for high-dimensional parameter spaces that are routinely explored with Bayesian retrievals with modest computational resources. I discuss the trade offs and potential avenues for the future.

I present AURA-3D, a three-dimensional atmospheric retrieval framework for exoplanet transmission spectra. AURA-3D includes a forward model that enables rapid computation of transmission spectra in 3D geometry for a given atmospheric structure and can, therefore, be used for atmospheric retrievals as well as for computing spectra from General Circulation

Models (GCMs). In order to efficiently explore the space of possible 3D temperature structures in retrievals, I develop a parametric 3D pressure-temperature profile which can accurately represent azimuthally-averaged temperature structures of a range of hot Jupiter GCMs. I apply this retrieval framework to simulated JWST observations of hot Jupiter transmission spectra, obtaining accurate estimates of the day-night temperature variation across the terminator as well as the abundances of chemical species. I demonstrate an example of a model hot Jupiter transmission spectrum for which a traditional 1D retrieval of JWST-quality data returns biased abundance estimates, whereas a retrieval including a day-night temperature gradient can accurately retrieve the true abundances. The forward model also has the capability to include inhomogeneous chemistry as well as variable clouds/hazes. This new retrieval framework opens the field to detailed multidimensional atmospheric characterisation using transmission spectra of exoplanets in the JWST era.

I also present a new internal structure model for super-Earths and mini-Neptunes that enables detailed characterisation of a planet's water component. I use my model to determine how the bulk properties and surface conditions of a water world affect its ocean depth, finding that oceans can be up to hundreds of times deeper than on Earth. I explore the region of mass-radius space in which planets with H-rich envelopes could host liquid H₂O. Such envelopes could contribute significantly to the planet radius while retaining liquid water at the surface, highlighting the exciting potential for habitable conditions to be present on planets much larger than Earth.

I contribute to internal structure models of a number of sub-Neptunes whose atmospheres are set to be observed using JWST. Before such observations take place, it is vitally important to understand the interior structures of these planets, which strongly affects their possible atmospheric compositions. We use the bulk parameters and retrieved atmospheric properties to constrain the internal structure and thermodynamic conditions in the habitable-zone mini-Neptune K2-18b, for which I contribute the H₂O EOS. The constraints on the interior allow multiple scenarios between rocky worlds with massive H/He envelopes and water worlds with thin envelopes. We constrain the mass fraction of the H/He envelope to be $\leq 6\%$; spanning $\lesssim 10^{-5}$ for a predominantly water world to $\sim 6\%$ for a pure iron interior. The thermodynamic conditions at the surface of the H₂O layer range from the supercritical to liquid phases, with a range of solutions allowing for habitable conditions. We also investigate the possible compositions of the pair of planets orbiting the star TOI-776. The bulk densities of TOI-776b and c allow for a wide range of possible interior and atmospheric compositions. However, the models indicate that both planets must have retained a significant atmosphere. Upcoming observations will revolutionise our understanding of these planets, helping to uncover the mysteries of the sub-Neptune population.

Finally, I discuss the latest developments in exoplanet observations, and consider how these advances may further our understanding of worlds beyond our own.

For Teresa Quinn, for convincing me that taking up A Level Physics would be a good use of
my spare time. You were right!

1959–2020

Scientific Acknowledgements

I acknowledge support and funding from the Cambridge Centre for Doctoral Training in Data Intensive Science, which is funded by the Science and Technology Facilities Council (STFC), UK. High-Performance Computing resources required to carry out much of the work presented in this thesis were provided by the Cambridge Service for Data Driven Discovery (CSD3), operated by the University of Cambridge Research Computing Service (www.csd3.cam.ac.uk), provided by Dell EMC and Intel using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/P020259/1) and DiRAC funding from STFC. This work makes use of a variety of software packages and data resources:

- The vast majority of computational work is completed in Python, with extensive use of the packages `NUMPY` (Harris et al., 2020), `SCIPY` (Virtanen et al., 2020) and `MATPLOTLIB` (Hunter, 2007).
- The NASA Astrophysics Data System and the NASA Exoplanet Archive are used to provide access to publications and exoplanet data.
- Chapter 2 uses the `AURA` retrieval code (Pinhas et al., 2018), which implements the MultiNest algorithm (Feroz & Hobson, 2008) through the Python package `PYMULTINEST` (Buchner et al., 2014).
- Chapter 3 uses the Python package `SCIKIT-LEARN` (Pedregosa et al., 2011).
- Chapter 4 uses the Python package `STARRY` (Luger et al., 2019) to generate Figure 4.2.
- Chapter 6 uses the Python package `TERNARY` (Harper et al., 2015) to generate Figure 6.2.

I would like to thank the anonymous referees for their insightful comments on my published papers (Nixon & Madhusudhan, 2020, 2021, 2022) which make up Chapters 3–5 of this thesis. I also thank Stephen Thorp for helpful discussion regarding Figure 4.5 and Anjali Piette for providing model temperature profiles for comparison as shown in Figure 5.6.

Acknowledgements

To begin, I would like to express my sincere thanks to Nikku Madhusudhan for his guidance and mentorship. Madhu's inspiring lectures and encouragement to go ahead with my (extremely last-minute!) PhD application are a huge part of why this all started, and I am delighted to have had the opportunity to work with and learn from him throughout these four years. I would also like to thank Cathie Clarke and Mihkel Kama for their support, and to acknowledge all of my collaborators for their help and insight. I am hugely grateful to my fellow members of the exoplanets group at the IoA, particularly to Luis Welbanks, Anjali Piette, Siddharth Gandhi and Arazi Pinhas for helping me learn the ropes as a new student, and to Adam Langeveld, Savvas Constantinou, Connor Cheverall and Måns Holmberg for providing a friendly and stimulating environment to discuss science.

I am very fortunate to have been able to find plenty of non-academic activities which have made my time at Cambridge infinitely more enjoyable, often revolving around music. I would like to express thanks to the members of Fitz Swing, Dysfunktional, Cambridge University Jazz Orchestra, Moonflower, Syzygie, Quasar, Soft Crunchy Landing, Easy Peeler and Colonel Spanky's Love Ensemble, as well as the many other musicians I have had the pleasure of playing with over the years, several of whom have become some of my best friends. I am also grateful to the committee members of Clare Jazz and Cambridge University Jazz Society for helping organise some incredible gigs. Additionally, I would like to thank Tom Comerford, Zephyr Penoyre, Cat Sinclair and Jeff Jennings for their contribution to games nights and crossword mornings, as well as Frank Paul for his regular provision of fiendish quizzes.

A substantial portion of the research presented in this thesis was completed while living in some degree of lockdown due to the COVID-19 pandemic. I consider myself extremely lucky to have spent much of this time living with some very nice housemates (...nice housemates!) who made the cursed moments far more bearable. I would therefore like to say "Chee" to the members of Hefty House: Jake Bennett, Stephen Thorp, Adam Langeveld, Andy Everall and Regor Marshall, for their steadfast commitment to monthly meals and General Vibing. I hope that in the future we can (not) regroup to Have A Lunch and continue speculating.

I would like to thank my sister, Lucy, and my parents, Claire and Mark, for their support and encouragement. My parents have been there for me throughout my PhD (and everything leading up to it), and there is no way I can express my appreciation for all they have done within a single paragraph, but I would particularly like to say thanks for the (often late-night) proofreading and the (also often late-night) phone calls during some of the more difficult points in my PhD journey.

My final thanks go to Anna Tindall, for being the best. I hope that Chapter 1 lives up to your excitement!

Table of contents

List of figures	xix
List of tables	xxiii
1 Introduction	1
1.1 Discovery of Exoplanets	1
1.1.1 Exoplanet Detection Methods	2
1.2 Planet Formation and Demographics	7
1.2.1 Diversity of Exoplanets	7
1.2.2 Planetary Bulk Composition	10
1.2.3 Planet Formation	13
1.2.4 Habitability and Biosignatures	15
1.3 Characterisation of Exoplanets	16
1.3.1 Internal Structure Models	16
1.3.2 Atmospheric Observations	17
1.3.3 Atmospheric Models	23
1.3.4 Atmospheric Physics and Chemistry	25
1.4 Scope of this Thesis	30
1.4.1 Atmospheric Retrieval of Transmission Spectra	30
1.4.2 Supervised Machine Learning for Atmospheric Retrieval of Exoplanets	30
1.4.3 A Three-Dimensional Retrieval Framework for Exoplanet Transmis- sion Spectra	31
1.4.4 Exploring the Phase Structure of Water-Rich Sub-Neptunes	32
1.4.5 Internal Structures of JWST Targets in the sub-Neptune Regime	32
1.4.6 Discussion and Conclusions	33
2 Atmospheric Retrieval of Transmission Spectra	35
2.1 Model Transmission Spectra	36

2.1.1	Radiative Transfer	36
2.1.2	Transmission Spectra	39
2.1.3	Opacity Sources	41
2.1.4	Atmospheric Structure	46
2.2	Parameter Estimation	47
2.2.1	Grid-based Retrievals	48
2.2.2	Bayesian Inference	49
2.2.3	Machine Learning	52
2.3	Retrieval of a Hot Jupiter Transmission Spectrum	57
2.4	Summary and Discussion	59
3	Supervised Machine Learning for Atmospheric Retrieval of Exoplanets	63
3.1	Introduction	63
3.2	Methods	65
3.2.1	Reproduction of Previous Results	65
3.2.2	Extension of Random Forest Method	71
3.3	Applications	73
3.3.1	Validation	74
3.3.2	Retrieval of WFC3 Transmission Spectrum	76
3.3.3	Addition of Unconstrained Free Parameter	78
3.4	Discussion and Conclusions	82
4	A Three-Dimensional Retrieval Framework for Exoplanet Transmission Spectra	89
4.1	Introduction	89
4.2	Methods	93
4.2.1	Transit Geometry	94
4.2.2	Radiative Transfer	95
4.2.3	A Multidimensional Parametric Temperature Profile	97
4.2.4	Inhomogeneous Chemistry	100
4.2.5	Additional Physical Effects	102
4.2.6	Synthetic Data and Statistical Inference	105
4.2.7	Model Validation	106
4.3	Results	106
4.3.1	Comparison of Parametric Profile with GCM Temperature Profiles	107
4.3.2	Effects of a Day-night Temperature Gradient on Atmospheric Spectra	109
4.3.3	Effect of Chemical Inhomogeneity on Transmission Spectra	110
4.3.4	Retrieval of a Multidimensional Temperature Structure	113

4.3.5	Retrieval of a Synthetic Spectrum from a GCM	118
4.4	Summary and Discussion	120
4.4.1	Applicability of 1D vs 3D Retrievals	122
4.4.2	Complementary Observations using Emission Spectra and Phase Curves	124
4.4.3	Future Developments	124
5	Exploring the Phase Structure of Water-Rich Sub-Neptunes	127
5.1	Introduction	127
5.2	Methods	130
5.2.1	Planetary Structure Equations	131
5.2.2	Equations of State	132
5.2.3	Temperature Profiles	138
5.2.4	Choice of Temperature Profiles in This Study	142
5.3	Model Validation	143
5.3.1	Comparison with Previous Studies	143
5.3.2	Mass–radius Relations	145
5.4	Results	149
5.4.1	Depth of Oceans on Water Worlds	149
5.4.2	Potential for Liquid Water on Mini-Neptunes	152
5.4.3	Diversity of Water World Phase Structures	155
5.4.4	Planets with Mixed Envelopes	160
5.5	Summary and Discussion	162
5.5.1	Potential Habitability of Water Worlds	162
5.5.2	Future Directions and Applications	164
6	Internal Structures of JWST Targets in the Sub-Neptune Regime	167
6.1	Introduction	167
6.2	The Habitable-zone Exoplanet K2-18b	168
6.2.1	Constraints on Internal Structure	170
6.2.2	Atmosphere-ocean Boundary	173
6.2.3	Possible Composition and Origins	175
6.2.4	Potential Habitability	176
6.3	A Pair of Mini-Neptunes Orbiting the Bright M-dwarf TOI-776	177
6.3.1	Planetary Composition and Internal Structure	177
6.4	Future Prospects for Characterisation	179

7	Conclusions and Future Developments	183
7.1	Machine Learning for Atmospheric Retrieval	183
7.2	Three-dimensional Atmospheric Retrieval	184
7.3	Phase Structures of Water-rich Super-Earths	185
7.4	Internal Structures of JWST Targets	186
7.5	Concluding Remarks	186
	References	189

List of figures

1.1	Cumulative exoplanet detections per year, sorted by detection method. . . .	2
1.2	Schematic diagram depicting the radial velocity method.	4
1.3	Transit light curve of the planet HD 209458b, taken from Charbonneau et al. (2000).	6
1.4	Radius-period distribution of known exoplanets.	8
1.5	Mass-period distribution of known exoplanets.	9
1.6	Masses and radii of exoplanets and solar system planets.	11
1.7	Completeness-corrected histogram of planet radii for planets with orbital periods shorter than 100 days, from Fulton et al. (2017).	12
1.8	Transiting planet with different observing modes highlighted.	17
1.9	Schematic representation of a transmission spectrum.	20
1.10	Physical and chemical processes in exoplanetary atmospheres that may be probed by different regions of the electromagnetic spectrum.	26
1.11	Mixing ratios X_i of key molecules present in hydrogen-dominated atmospheres in chemical equilibrium.	28
2.1	Schematic of atmospheric retrieval.	36
2.2	Geometry of an exoplanet atmosphere as observed in transit.	40
2.3	Cross-sections of prominent molecular species in hydrogen-rich atmospheres.	45
2.4	Pressure–temperature profiles commonly employed in retrievals.	48
2.5	Retrieval of the HST STIS+WFC3 transmission spectrum of HD 209458b.	59
2.6	Full posterior probability distribution from retrieval of the transmission spectrum of HD 209458b.	61
3.1	Flowchart describing the extended Random Forest retrieval framework.	65
3.2	Posterior distributions of Nested Sampling retrieval of the WFC3 transmission spectrum of WASP-12b following the methods of Márquez-Neila et al. (2018).	66

3.3	Results of Random Forest retrieval of the WFC3 transmission spectrum of WASP-12b following the methods of Márquez-Neila et al. (2018).	69
3.4	Results of extended Random Forest retrieval of the WFC3 transmission spectrum of WASP-12b, using the same forward model as in Márquez-Neila et al. (2018).	73
3.5	True values versus Random Forest predictions for each parameter in the forward model, for a test set of 2000 synthetic models.	74
3.6	Difference between true and predicted values of P_{ref} and $X_{\text{H}_2\text{O}}$ for synthetic models with $\log X_{\text{H}_2\text{O}} \geq -6$	75
3.7	Results of Nested Sampling retrieval of the WFC3 transmission spectrum of HD 209458b.	77
3.8	Results of extended Random Forest retrieval of the WFC3 transmission spectrum of HD 209458b.	78
3.9	Best-fitting model spectra from Nested Sampling and extended Random Forest retrievals of the HD 209458b WFC3 transmission spectrum.	79
3.10	Results of Nested Sampling retrieval of the WFC3 transmission spectrum of HD 209458b, including CO abundance as a free parameter.	81
3.11	Results of Random Forest retrieval of the WFC3 transmission spectrum of HD 209458b, including CO abundance as a free parameter, following the methods of Márquez-Neila et al. (2018).	82
3.12	Results of extended Random Forest retrieval of the WFC3 transmission spectrum of HD 209458b, including CO abundance as a free parameter.	83
3.13	Marginal posterior distributions for CO abundance retrieved from the transmission spectrum of HD 209458b using different methods.	84
3.14	Comparison of the duration of Nested Sampling retrievals versus extended Random Forest retrievals as a function of spectral resolution.	85
4.1	Schematic of the AURA-3D retrieval framework.	90
4.2	Three-dimensional co-ordinate system adopted for our forward model.	94
4.3	Model transmission spectra of the canonical hot Jupiter HD 209458b with varying temperatures and transition angles.	98
4.4	Comparison between azimuthally-averaged P - T profiles from the terminator regions of a range of GCMs and parametric profiles.	101
4.5	Parametric fits to the 3D P - T profiles presented in Fortney et al. (2010) showing a slice at $\phi = 0$ (through the north-south polar plane).	102
4.6	Differences between spectra with a 1D thermally averaged temperature profile and spectra with different day-night temperature gradients.	103

4.7	Comparison of two theoretical transmission spectra for a HD 209458b-like planet.	107
4.8	Synthetic spectra of the hot Neptune HAT-P-11b with a range of different chemical compositions.	110
4.9	Retrieval of a synthetic JWST spectrum of the hot Jupiter HAT-P-1b.	111
4.10	Input and retrieved multidimensional parametric temperature profiles of a hot Jupiter.	112
4.11	Two-dimensional representations of the input and median retrieved parametric P - T profiles for HAT-P-1b.	113
4.12	Input and retrieved multidimensional parametric temperature profiles of a hot Jupiter with no day-night temperature contrast.	114
4.13	Marginalised posterior probability distributions for the retrieval of a synthetic JWST transmission spectrum of the hot Jupiter HAT-P-1b.	115
4.14	Atmospheric retrieval of a synthetic hot Jupiter transmission spectrum generated using a 3D temperature structure.	116
4.15	Input temperature structure from a GCM and corresponding retrieved parametric temperature profiles.	118
5.1	Architecture of the internal structure model used in this study.	131
5.2	Equations of State for Fe and MgSiO ₃ used in our model.	133
5.3	Data sources for the H ₂ O EOS used in our model and phase diagram of H ₂ O.	134
5.4	Density and specific heat capacity of H ₂ O.	135
5.5	EOS for H/He used in our model for different isotherms.	136
5.6	Comparison of analytic and self-consistent P - T profiles.	139
5.7	Comparison between M - R relations produced by our model and previously published results.	143
5.8	Mass-radius relations for water worlds without H/He envelopes.	146
5.9	M - R relations for planets with H/He envelopes.	148
5.10	Extent of oceans on H ₂ O-rich planets with different surface gravities and temperatures.	149
5.11	Internal structure for one possible interior composition of K2-18b, corresponding to Case 3 from Madhusudhan et al. (2020).	153
5.12	M - R diagram showing the range of planet masses and radii consistent with a H/He envelope and liquid H ₂ O at the HHB.	156
5.13	P - T profiles of the planets shown in Figure 5.14.	157
5.14	Phase structures of H ₂ O-rich planets with different surface conditions and temperature structures.	159

5.15	<i>M–R</i> relations for planets with mixed and unmixed envelopes of fixed mass fraction ($x_{\text{env}} = 0.05$) above an Earth-like nucleus.	161
6.1	Mass–radius diagram showing interior models with three representative compositions that fit the mass and radius of K2-18b.	169
6.2	Ternary diagram showing best-fitting ($\leq 1\sigma$) interior compositions allowed by the mass and radius of K2-18b for two end-member core compositions and interior temperatures.	171
6.3	Envelope vs core mass fraction for models consistent with the mass and radius of K2-18b.	173
6.4	Pressure–density profiles for three possible compositions of K2-18b discussed in Section 6.2.3.	174
6.5	Thermodynamic conditions at the H ₂ O–H/He boundary (HHB) for K2-18b.	175
6.6	Locations of TOI-776b and c on a mass–radius diagram.	178
6.7	H/He vs. H ₂ O mass fractions for the best-fit interior compositions permitted by the masses and radii of TOI-776b and c.	180
7.1	JWST NIRISS transmission spectrum of the hot Jupiter WASP-96b, adapted from its initial presentation by NASA.	187

List of tables

1.1	Stellar classification scheme with corresponding effective temperatures (Morgan et al., 1943, Habets & Heintze, 1981).	10
2.1	Description of priors for retrieval of the transmission spectrum of HD 209458b.	58
3.1	Description of priors for retrievals of the transmission spectrum of WASP-12b.	68
3.2	Description of priors for retrievals of the WFC3 transmission spectrum of HD 209458b.	76
3.3	Comparison of Nested Sampling (NS) and extended Random Forest (RF+) retrieved parameter values from the WFC3 transmission spectrum of HD 209458b.	80
4.1	Chemical abundances used for the day- and nightsides of our HAT-P-11b models.	117
4.2	Description of priors for retrievals of simulated JWST transmission spectra of HAT-P-1b generated using a GCM.	121
4.3	Input and retrieved abundances for the simulated JWST transmission spectrum of HAT-P-1b.	121
5.1	Parameters for the EOS of Fe (ϵ) from Anderson et al. (2001) and MgSiO ₃ (perovskite) from Karki et al. (2000).	132
5.2	Surface conditions and radiative convective boundaries for the model planets considered in Section 5.4.3.	158

Chapter 1

Introduction

“... there is an infinite number of worlds, some like this world, others unlike it.”

– Epicurus, c. 300 BCE

1.1 Discovery of Exoplanets

Astronomers have studied the behaviour of worlds beyond our own since ancient times. Our knowledge of other planets in the solar system can be traced back to at least the 2nd millennium BCE, with the motions of the “naked eye” planets (Mercury, Venus, Mars, Jupiter and Saturn) recorded by ancient Babylonian astronomers (Evans, 1998). The earliest planets with well-defined dates of discovery are the two ice giants of the solar system, Uranus (Herschel & Watson, 1781) and Neptune (Galle, 1846).

The discovery of all currently known solar system planets had therefore concluded in the mid-19th century¹, and even at that time, some astronomers were making (ultimately refuted) claims to have detected planets orbiting other stars (e.g. Jacob, 1855). However, we would have to wait almost 150 years for the first unambiguous identification of extrasolar planets, or exoplanets. There is some debate over which exoplanet discovery can be labelled the “first”, but two key milestones were the identification of a planetary system orbiting the pulsar PSR1257+12 (Wolszczan & Frail, 1992) and the detection of a Jupiter-sized planet closely orbiting the Sun-like star 51 Pegasi (Mayor & Queloz, 1995). These initial discoveries already hinted at a surprising diversity of planetary systems, both in terms of the nature of host stars and the sizes and positions of the planets themselves.

After these initial discoveries, planetary systems continued to be discovered at an ever-increasing rate throughout the following three decades, as shown in Figure 1.1. The variety of

¹We don't talk about Pluto . . .

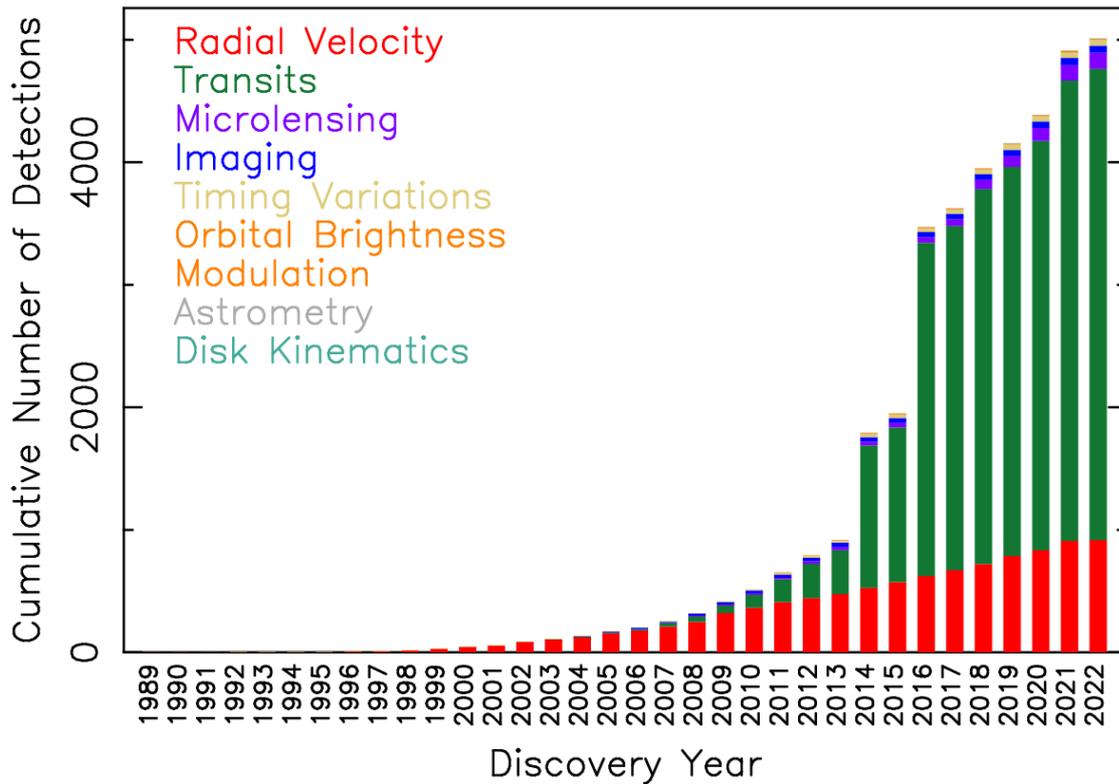


Fig. 1.1 Cumulative exoplanet detections per year, sorted by detection method. Figure taken from the NASA Exoplanet Archive.

methods used to detect new planets also increased (see Section 1.1.1). While many different ground- and space-based telescopes have been employed in the search for distant worlds, one telescope in particular has been responsible for the notable jump in the number of detected planets beginning in 2014. The *Kepler* Space Telescope (Borucki et al., 2010) has been responsible for the detection of over half of all currently-known exoplanets, with astounding yields of 715 detections in 2014 and a further 1 284 detections in 2016. While the *Kepler* mission ended in 2018, other facilities such as the *Transiting Exoplanet Survey Satellite* (TESS, Ricker et al., 2015) have continued to monitor the skies in search of planets, with the total number of confirmed planets in the NASA Exoplanet Archive surpassing 5 000 on 21st March 2022.

1.1.1 Exoplanet Detection Methods

As indicated in Figure 1.1, a wide range of methods have been successfully employed in the search for extrasolar worlds. Here we² summarize five of the most fruitful approaches: radial

²The remainder of this thesis uses the first person plural for stylistic consistency.

velocity, transits, microlensing, direct imaging, and transit timing variations. For a more detailed discussion, see Wright & Gaudi (2013).

1.1.1.1 Radial Velocity

The radial velocity method was used to detect 51 Pegasi b (Mayor & Queloz, 1995), the first exoplanet orbiting a Sun-like star, and was responsible for the majority of exoplanet discoveries until 2009. This method is based on the fact that when a planet orbits a star, the gravitational pull of the planet causes the star to “wobble” around the centre of mass of the star-planet system. We can monitor this wobble by looking for periodic shifts in the absorption lines of the stellar spectrum (see Figure 1.2). The change in wavelength due to redshifted or blueshifted spectral lines is related to the radial velocity of the system through the relativistic Doppler effect:

$$\lambda_{\text{obs}} = \sqrt{\frac{1 + V_r/c}{1 - V_r/c}} \lambda_{\text{em}}, \quad (1.1)$$

where λ_{obs} and λ_{em} are the observed and emitted wavelengths, V_r is the radial velocity of the star, and c is the speed of light. It can be shown that the stellar radial velocity V_r can be written in terms of the following orbital components (see e.g. Murray & Correia, 2010):

$$V_r = K[\cos(\nu + \omega) + e \cos \omega] + \gamma, \quad (1.2)$$

where K is the semi-amplitude of the radial velocity signal, ν is the true anomaly, ω is the argument of pericentre, e is the orbital eccentricity, and γ is the bulk velocity of the centre of mass of the system. The observable quantities K and e can then be related to the masses of the star (M_*) and planet (M_p) using Kepler’s laws:

$$K = \left(\frac{2\pi G}{P\sqrt{1 - e^2}} \right)^{1/3} \frac{M_p \sin i}{(M_p + M_*)^{2/3}}, \quad (1.3)$$

where P is the orbital period and i is the orbital inclination. Radial velocity observations can therefore be used to constrain the mass of a planet. Since the orbital inclination is typically unknown, the observed planet mass is typically reported assuming an edge-on orbit ($\sin i = 1$), which gives the minimum possible mass. The true planet mass will be higher by a factor of $1/\sin i$. For randomly aligned orbits, the median value of $1/\sin i$ is approximately 1.15.

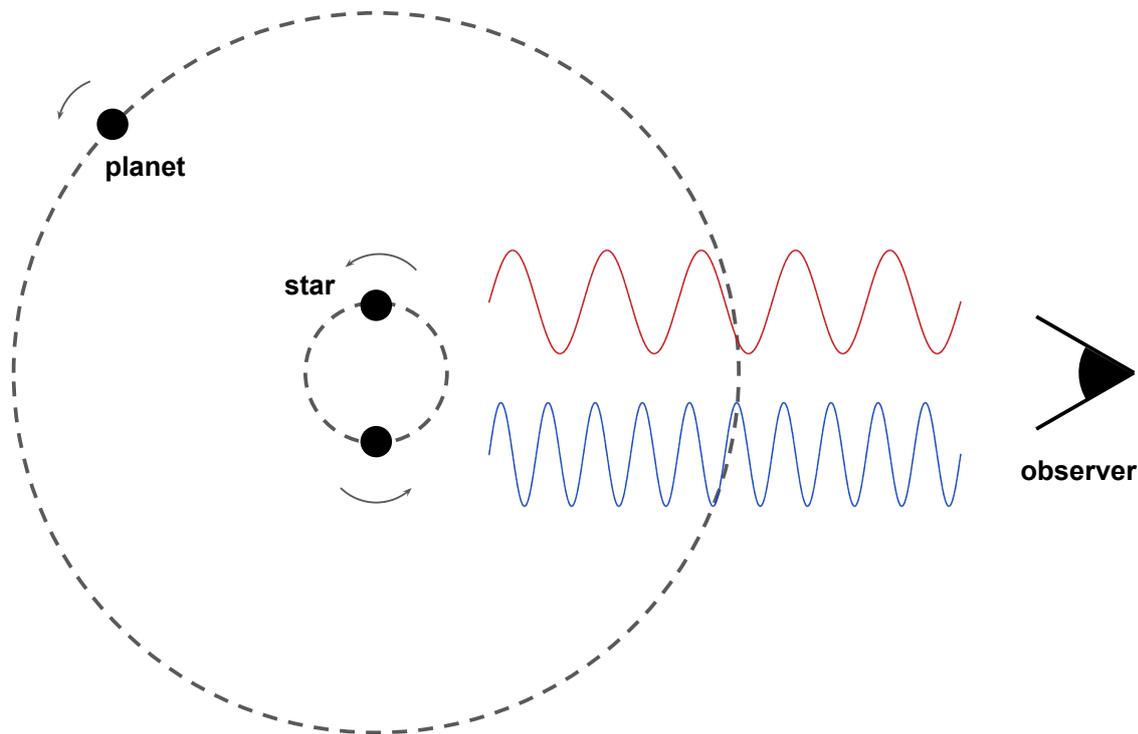


Fig. 1.2 Schematic diagram depicting the radial velocity method. The planet and star orbit a common center of mass. As the star moves away from the observer, stellar absorption lines are redshifted, while as the star moves towards the observer the lines are blueshifted.

It is evident from equation 1.3 that the radial velocity method is best suited to the detection of high-mass planets on short-period orbits. It is therefore unsurprising that the earliest planets detected using this method were hot Jupiters, large planets orbiting in close proximity to their host stars. The observed radial velocity signal of 51 Pegasi b, $K = 59 \text{ m s}^{-1}$, was measured with a precision of 13 m s^{-1} (Mayor & Queloz, 1995). More recent advances have led to measurements as precise as 1 m s^{-1} (Fischer et al., 2016), which has enabled the detection of terrestrial planets in short-period orbits around low-mass stars, such as Proxima Centauri b (Anglada-Escudé et al., 2016). However, current techniques still fall short of detecting Earth-mass planets orbiting nearby stars at habitable-zone distances, which would require a precision of $\sim 0.1 \text{ m s}^{-1}$.

1.1.1.2 Transits

The transit method has now overtaken radial velocity as the most successful method for detecting exoplanets. As a planet orbits its host star, it may pass in front of a portion of the stellar surface, depending on the inclination of the system. This occurrence is called a transit, or primary eclipse. An observer would therefore notice a periodic decrease in the

star's brightness, which can therefore be used to infer the presence of an exoplanet. We can measure the transit depth, δ , by comparing the fluxes observed outside (F_{out}) and inside (F_{in}) the transit (see Chapter 2). Since the flux received is proportional to the observed area of the emitting object, the transit depth is approximately equal to the square of the ratio of the planetary and stellar radii (R_p and R_*):

$$\delta = \frac{F_{\text{out}} - F_{\text{in}}}{F_{\text{out}}} \approx \left(\frac{R_p}{R_*} \right)^2. \quad (1.4)$$

For a more detailed derivation see e.g. Winn (2010). Note that this expression assumes that the planet is entirely opaque, an assumption that will be relaxed when we come to discuss observations of exoplanet atmospheres in Section 1.3.2.

Figure 1.3 shows the transit light curve of HD 209458b, the first exoplanet to be observed transiting its host star (Charbonneau et al., 2000). The transit method has since been used to detect thousands of planets, and was the principal detection method employed by the *Kepler* Space Telescope (Borucki et al., 2010). Like the radial velocity method, the transit method favours giant planets, in this case those with large radii relative to their host star. This method is also more likely to detect planets with short orbital periods, since the probability of a transit event is higher for planets which are closer to their host star (Winn, 2010). A Jupiter-sized planet orbiting a Sun-like star leads to a transit depth $\delta \sim 0.01$. An Earth-sized planet orbiting the same star will yield a much smaller $\delta \sim 10^{-4}$, making these planets much more difficult to detect with this method. However, it is more feasible to detect such planets around smaller host stars, since the fractional change in flux will be larger. This has led to the discovery of terrestrial planets orbiting small stars such as TRAPPIST-1 (Gillon et al., 2017) and LHS 1140 (Dittmann et al., 2017).

1.1.1.3 Microlensing

When a foreground object such as a star or stellar remnant passes very close along our line-of-sight to a more distant star, light from the distant star is split into two different images in a process known as gravitational lensing (Einstein, 1936). These images will be magnified by an amount that varies as the angular separation between the foreground and background objects changes over time. The variation of the magnification over time is known as a microlensing event. In the case where the foreground object is a star with an orbiting planet, it is possible that the gravity of the planet will further distort the light from the background star, allowing the planet to be detected. A detailed description of the mathematics underlying this detection method can be found in Gaudi (2012).

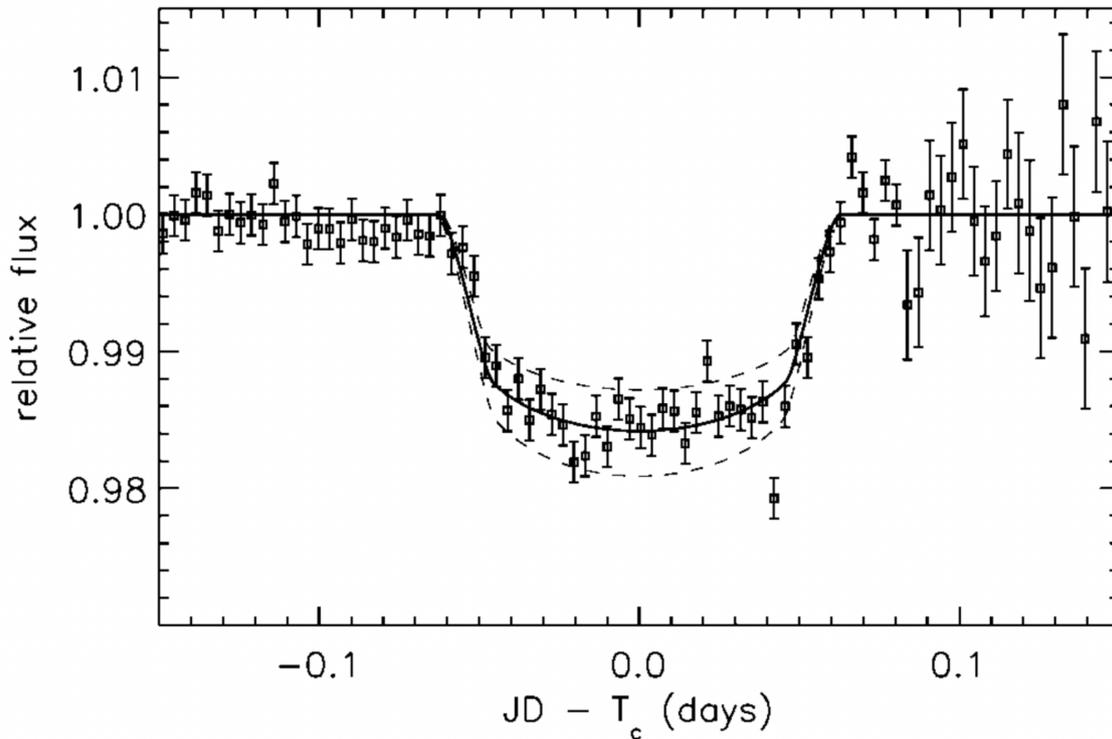


Fig. 1.3 Transit light curve of the planet HD 209458b, taken from Charbonneau et al. (2000). Data are plotted as the relative flux (normalised by F_{out}) as a function of the time from the middle of the transit, T_c .

The microlensing method for detecting exoplanets was first proposed by Mao & Paczynski (1991) and first demonstrated by Bond et al. (2004), with microlensing surveys such as the Optical Gravitational Lensing Experiment (OGLE, Udalski et al., 1992) typically contributing a handful of planetary detections per year since then. Unlike the radial velocity and transit methods, which favour planets on close-in orbits, microlensing is effective at detecting planets on wide orbits, or even free-floating planets (e.g. Mróz et al., 2018). While it is challenging to derive physical properties of a planet or its host star using this method, useful information can be derived from higher-order effects in microlensing light curves, or via follow-up high resolution imaging (Bennett et al., 2007).

1.1.1.4 Direct Imaging

While the above methods involve indirect inference of the presence of a planet through its effect on a host star or background object, the direct imaging method aims to directly detect reflected light or thermal emission from a planet. Direct imaging of an exoplanet is extremely challenging, due partly to the fact that planets are faint objects, but also to the difficulty

in spatially resolving the planet from its (much brighter) host star. Detections via direct imaging have therefore been mostly limited to young planets on wide orbits, which emit in the near-infrared due to escaped internal heat energy (Bowler & Nielsen, 2018).

The first directly-imaged exoplanet was presented by Chauvin et al. (2004), and direct imaging has also been used to detect multiplanet systems (e.g. Marois et al., 2008). As of January 2022, direct imaging has led to the detection of 59 exoplanets according to the NASA Exoplanet Archive. While the planetary radius can be inferred from the observed flux from the planet, it is not possible to measure a planet's mass via direct imaging.

1.1.1.5 Transit Timing Variations

Transit timing variations (TTVs, Agol et al., 2005) are a means of detecting additional planets in a system in which one transiting planet has already been observed. If the transiting planet is the only planet orbiting its host star, then transit events should occur at evenly-spaced time intervals which can be determined using Kepler's laws. However, if a second planet is also orbiting the host star, the transiting planet will be gravitationally perturbed, leading to variations in the timing of individual transits. By measuring these variations it is possible to infer the mass and orbital properties of the perturbing planet (Nesvorný & Morbidelli, 2008). This method has been successfully employed using *Kepler*, with the earliest planet detected via TTVs presented by Ford et al. (2012). According to the NASA Exoplanet Archive, TTVs have led to the discovery of 22 exoplanets as of January 2022.

1.2 Planet Formation and Demographics

1.2.1 Diversity of Exoplanets

Figures 1.4 and 1.5 show the radius-period and mass-period distributions of known planets. It is evident from this diagram that the exoplanet population is extremely diverse, with discovered planets ranging from Jupiter-sized objects in close orbits around their host stars to planets with sizes between those of Earth and Neptune, which have no solar system analogue. It is important to note that the sample of planets which have been detected is biased by our observational capabilities; it is much easier to detect larger planets with small orbital radii, and so these planets are likely overrepresented in our observed sample relative to their true occurrence rate.

The observed radius distribution of exoplanets does not appear to have any complete gaps, meaning that drawing clear distinctions between different varieties of planet based on size is challenging. However, it can still be informative to group exoplanets into a few broad classes:

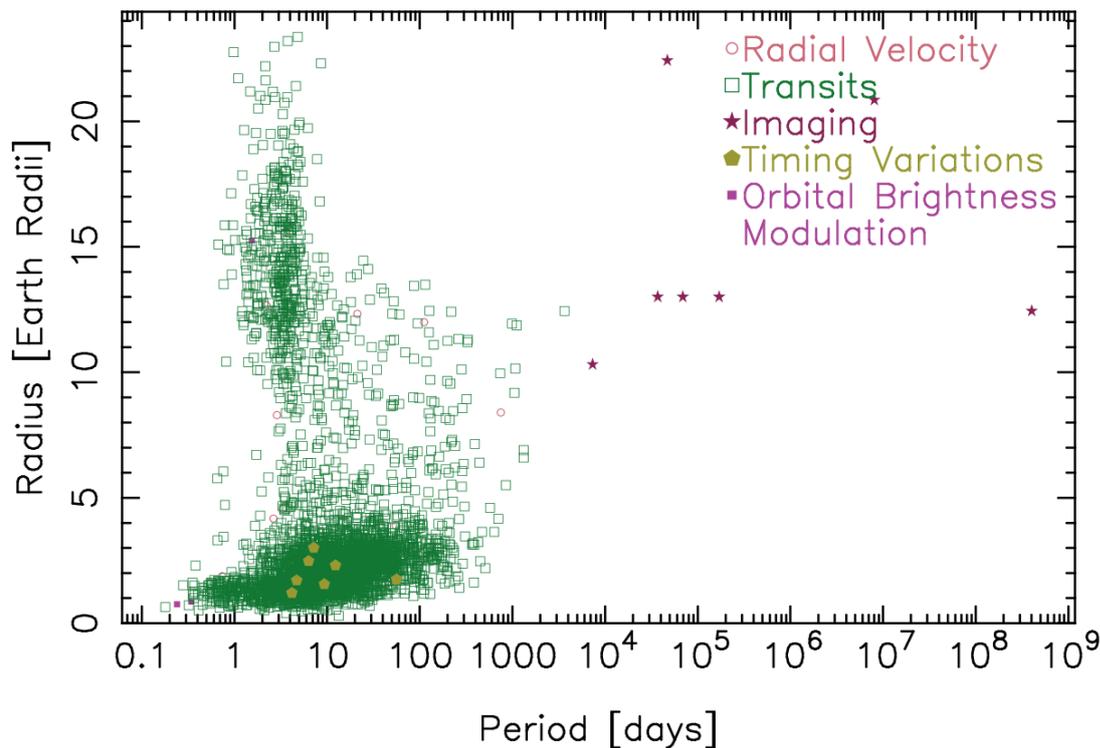


Fig. 1.4 Radius-period distribution of known exoplanets. Figure from the NASA Exoplanet Archive.

- **Gas giants.** These objects have masses and radii comparable to Jupiter and Saturn, with the largest examples having radii up to $\sim 2R_J$ and masses up to $\sim 13M_J$. Higher mass objects are capable of deuterium fusion, and are known as *brown dwarfs*. Gas giants can be further subdivided based on the level of irradiation they receive from their host star. Highly-irradiated giant planets are much closer to their host star than solar system gas giants ($P \lesssim 10$ days) and therefore have much higher temperatures ($T \sim 800\text{--}4000$ K). These planets are often referred to as *hot Jupiters* (or *hot Saturns*). In contrast, less irradiated giant planets have longer orbital periods ($P \gtrsim 100$ days) making them in some sense closer analogues to the giant planets in our solar system. Young, hot giant planets on wider orbits are amenable to detection via direct imaging.
- **Ice giants.** These planets have masses and radii comparable to Uranus ($M = 14.54 M_\oplus$, $R = 4.01 R_\oplus$) and Neptune ($M = 17.15 M_\oplus$, $R = 3.88 R_\oplus$). The majority of detected planets in this size range have short orbital periods ($P \lesssim 100$ days); sub-giants on longer orbital periods are beyond the capabilities of current detection missions, but may be detectable in the future with the advent of facilities such as the *Nancy Grace Roman Space Telescope* (Spergel et al., 2015).

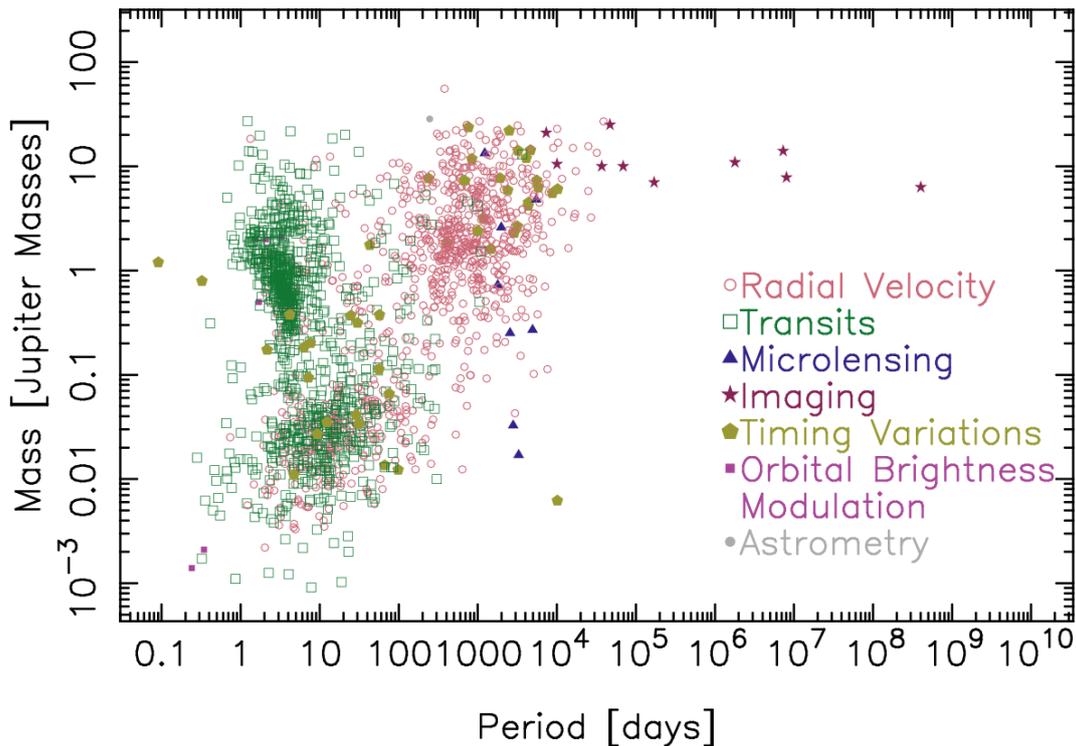


Fig. 1.5 Mass-period distribution of known exoplanets. Figure from the NASA Exoplanet Archive.

- **Sub-Neptunes.** Planets with radii between that of Earth and Neptune ($\sim 1\text{--}4R_{\oplus}$) are typically referred to as *sub-Neptunes*. The sub-Neptune population can be further divided into *super-Earths* and *mini-Neptunes*, with the cutoff between the two classes typically located around $1.5\text{--}2R_{\oplus}$ based on inferred bulk composition (see Section 1.2.2). Like ice giants, sub-Neptunes have mostly been detected at short orbital periods.
- **Rocky planets.** These planets have comparable masses and radii to the terrestrial planets of the solar system. Due to their small size, these planets are generally the most difficult to detect and characterise, and so the extent to which rocky exoplanets have similar properties to small solar system planets remains to be seen.

When detection biases are accounted for, it is believed that the average number of planets orbiting a given star in the galaxy is close to 1, if not higher (Youdin, 2011, Cassan et al., 2012). This does not mean that every star hosts a planet, since many stars are known to host multiple planets. Sub-Neptunes appear to be common, and are estimated to orbit $\sim 50\%$ of Sun-like stars (Mulders, 2018). Cold giant planets have a lower occurrence rate of $\sim 10\%$ (Zhu & Dong, 2021), while hot Jupiters are rarer still, with an occurrence rate of $\sim 1\%$ (Deleuil et al., 2018).

Table 1.1 Stellar classification scheme with corresponding effective temperatures (Morgan et al., 1943, Habets & Heintze, 1981).

Spectral Class	Effective Temperature
O	$\geq 30\,000$ K
B	10 000–30 000 K
A	7 500–10 000 K
F	6 000–7 500 K
G	5 200–6 000 K
K	3 700–5 200 K
M	2 400–3 700 K

Not only are the planets in our galaxy highly diverse, they also orbit a wide range of stars. Planets have been detected orbiting stars of every spectral class, from massive O and A stars to smaller, cooler M dwarfs (see Table 1.1). There are believed to be correlations between stellar properties and the nature of their planets; for example, detection surveys have suggested that massive, metal-rich stars are more likely to host giant planets (Quirrenbach et al., 2011).

1.2.2 Planetary Bulk Composition

Figure 1.6 shows a sample of masses and radii of exoplanets and solar system planets. Once the mass and radius (and therefore the density) of a planet are known, it is possible to compare these quantities to planetary interior structure models in order to determine possible bulk compositions (see Section 1.3.1). Several model mass-radius relationships are overlaid on Figure 1.6 for reference. While the models shown above are very simplistic, they yield some useful information about the possible compositions of various planets. For example, we can see that most terrestrial planets are consistent with a rock+iron composition, and that giant planets must consist largely of hydrogen and helium. Note that a selection of giant planets have larger radii than predicted for even a pure H/He planet; this is known as the *inflated hot Jupiters* problem (Bodenheimer et al., 2001, Guillot & Showman, 2002). A number of possible mechanisms may explain this anomaly, such as the conversion of wind energy to heat or dissipation induced by thermal tides (see e.g. Laughlin, 2018, Guillot et al., 2022).

Although the models shown in Figure 1.6 can provide some insight into the composition of exoplanets, they must be considered with some caution. These models assume uniform

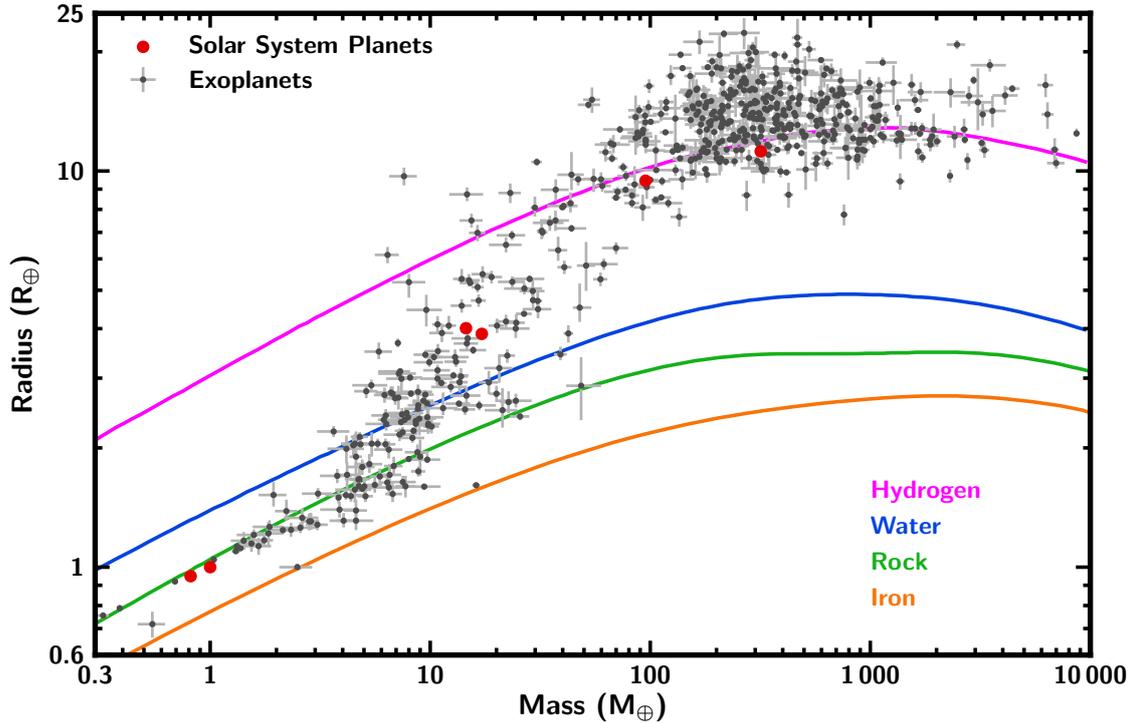


Fig. 1.6 Masses and radii of exoplanets and solar system planets. The magenta, blue, green and orange lines show mass-radius relations from Seager et al. (2007) for model planets made entirely of H, H₂O, MgSiO₃ and Fe respectively. Data for exoplanets with masses and radii measured to $\geq 5\sigma$ taken from the NASA Exoplanet Archive (1st July 2022).

compositions, which are very unlikely to occur in reality: for example, we know that the Earth is made up of a combination of $\sim 30\%$ iron and $\sim 70\%$ rock by mass. When multi-component models are taken into account, a wide range of possible compositions can be used to explain the observed mass and radius of most planets. For example, many sub-Neptunes have masses and radii consistent with a rocky core surrounded by $\sim 1\%$ of H/He by mass, or alternatively with a water-rich planet. The models presented here also fail to take temperature dependence into account, which can strongly affect mass-radius relations, especially for volatile-rich interiors (see Chapter 5).

The challenge posed by the interpretation of sub-Neptune interiors is particularly interesting given that these planets cannot be readily compared to any object in the solar system, and a number of hypotheses exist to explain this population. Figure 1.7 shows the observed radius distribution of sub-Neptunes (Fulton et al., 2017). This distribution is noticeably bimodal, with peaks around $1.3R_{\oplus}$ and $2.4R_{\oplus}$ and a trough near $1.7R_{\oplus}$, known as the *radius valley*. Planets either side of the radius valley are thought to have different interior compositions,

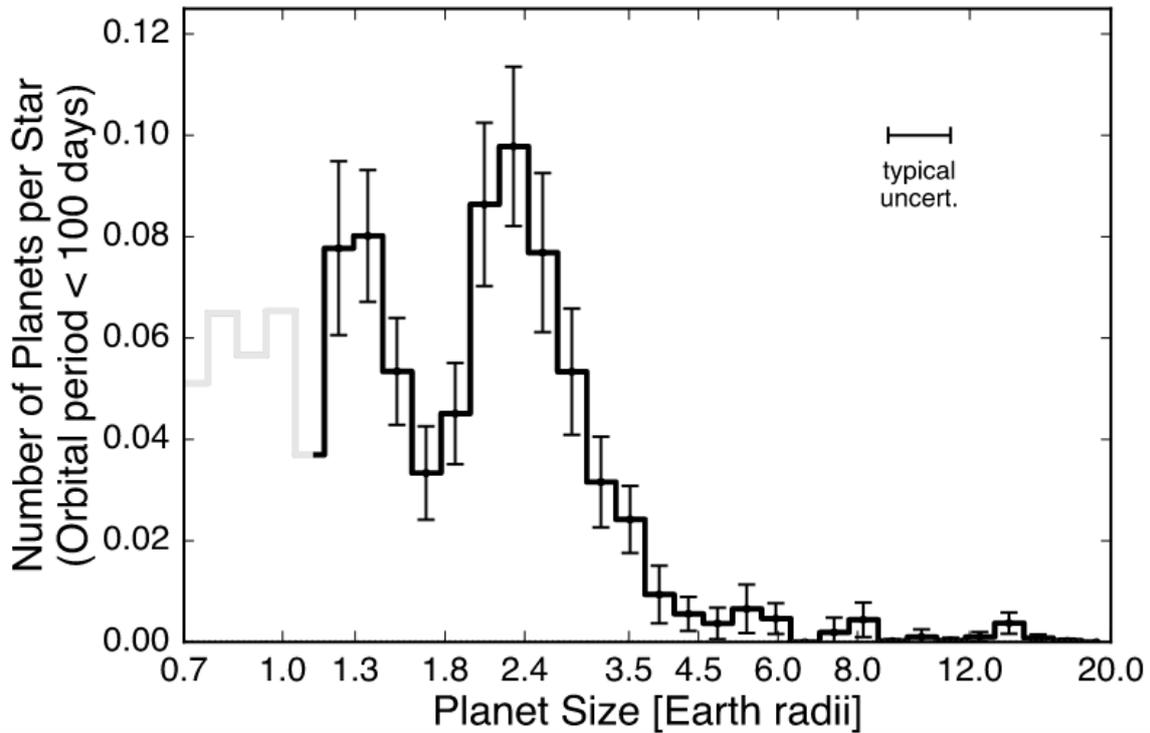


Fig. 1.7 Completeness-corrected histogram of planet radii for planets with orbital periods shorter than 100 days, from Fulton et al. (2017). The light grey region of the histogram for radii smaller than $1.14R_{\oplus}$ suffers from low completeness.

with the smaller planets possessing mostly rocky compositions akin to scaled-up versions of Earth (super-Earths) and the larger planets hosting a thick H/He envelope (mini-Neptunes) (e.g. Rogers, 2015). Planetary evolution models suggest that the radius valley is sculpted by atmospheric loss processes such as photoevaporation (Owen & Wu, 2013) or core-powered mass loss (Gupta & Schlichting, 2019) which strip smaller bodies of their gaseous envelopes while leaving larger atmospheres mostly intact.

Although the hypothesis that the sub-Neptune population consists of rocky planets with and without thick H/He envelopes is in broad agreement with extant observations, it is not the only feasible theory that can explain the existence of these objects. Numerous studies have considered the possibility of water-rich planets (e.g. Léger et al., 2004, Sotin et al., 2007, Grasset et al., 2009), which would have densities comparable to rocky planets with small amounts of H/He. This class of planet is explored in more detail in Chapter 5. While sub-Neptunes with H/He-rich atmospheres have been observed with the *Hubble* Space Telescope (HST, e.g. Benneke et al., 2019), a planet with a water-rich envelope has not yet been identified, although a number of candidate water-worlds do exist (e.g. Luque et al., 2021). It is possible that the sub-Neptune population consists of a mix of water-rich planets

and rocky planets with H/He envelopes, as well as planets with both a H/He envelope and a substantial water component (e.g. Pierrehumbert & Gaidos, 2011, Madhusudhan et al., 2021). The exact distribution of water and rock in a planet's interior depends on a variety of formation conditions, including the ice-to-rock ratio of the formation environment, a quantity which in general is not well constrained from current theory or observations.

1.2.3 Planet Formation

One of the most important goals of planetary and exoplanetary science is to better understand the underlying mechanisms behind planet formation, enabling us to explain how our own world and the many exotic worlds surrounding us came into being. This is a highly ambitious task, and while our knowledge of planet formation processes has progressed rapidly in recent decades, many unsolved questions remain. Insight into planet formation can be obtained by directly observing sites of planet formation, as well as by linking the properties of fully-formed planets to their formation and evolution.

The process of star (and planet) formation begins when material in a giant molecular cloud is pushed out of equilibrium and begins to collapse under its own gravity. As the collapse proceeds, a disk structure is formed due to conservation of angular momentum, through which matter accretes on to the forming star. This structure is known as a *protoplanetary disk*, and typically exists for $\lesssim 10$ million years with an orbital extent of $\gtrsim 30$ au. Some of the material in the disk goes on to form planets, which remain in orbit of the star after the rest of the disk material has either accreted on to the star or been driven away by some other process (see e.g. Williams & Cieza, 2011, Chabrier et al., 2014, Helled et al., 2014, for reviews).

Planets are typically thought to form in disks via one of two mechanisms: *core accretion* (Safronov, 1969, Pollack et al., 1996) or *gravitational instability* (Cameron, 1978, Boss, 1997). In the core accretion process, a solid body is initially built up through an accumulation of dust grains and pebbles, which are known as *planetesimals* once they reach a few km in diameter (Lambrechts & Johansen, 2012, Alibert et al., 2018). This leads to the formation of a proto-planet, which is usually referred to as a *planetary embryo* if it is the seed of a terrestrial or smaller planet, or a *planetary core* if it will go on to become a giant planet. If a core develops before the gaseous disk disperses, then gas may be steadily accreted from the disk, forming an envelope which surrounds the planet. If the mass of the envelope exceeds the core mass, then the envelope will collapse and exponential gas accretion will occur. A more detailed description of this process may be found in Helled & Morbidelli (2021).

The resulting types of planet that are produced according to this mechanism depend on which stages of the formation process are able to take place. Small, rocky planets are formed if the planetary embryo does not become large enough to initiate gas accretion (Wetherill,

1980). Neptunes and sub-Neptunes can be created if a large core develops in time to accrete some gaseous material from the disk, but does not acquire enough gas to trigger runaway accretion (Venturini & Helled, 2017). Finally, giant planets are the outcome of cores for which runaway accretion takes place (D'Angelo et al., 2010). Smaller planets are predicted to form in the inner disk at higher temperatures, with giant planets forming at greater orbital distances. The environment of the outer disk is more conducive to the formation of massive cores due to the presence of ices, which provide additional solid material for the proto-planet (Lecar et al., 2006).

The core accretion theory is consistent with the layout of the solar system, but is challenged by the existence of exoplanets such as hot Jupiters, which should not be able to form at the short orbital distances at which they have been observed. However, the idea that planets could form at large orbital distances and subsequently migrate inwards had already been proposed even before the discovery of hot Jupiters, having been suggested to explain the architecture of the solar system (Goldreich & Tremaine, 1980, Tsiganis et al., 2005). While this theory appears to explain the existence of giant planets with short orbital periods (Lin et al., 1996), other mechanisms including *in situ* formation (Batygin et al., 2016) and planet-planet scattering (Beaugé & Nesvorný, 2012) have also been proposed.

The gravitational instability model provides an alternative pathway for planet formation. According to this theory, planet formation begins when an overdense region in a protoplanetary disk collapses under its own gravity into clumps of gas and dust, which accrue material and contract to form giant planets (Boss, 1997). This process allows for planet formation on a faster timescale than core accretion (e.g. Lagrange et al., 2010) and may explain the existence of massive planets at wide orbital separations (e.g. Marois et al., 2010). However, this theory is unlikely to account for a large number of currently known planets (Kratler & Lodato, 2016).

According to these formation paradigms, the birth and development of a planet should be reflected in its present-day composition. For example, planets which form in the outer disk are expected to have a larger internal H₂O component (Alessi et al., 2017). Furthermore, the evolutionary history of a planet will impact its atmospheric properties. Evolutionary models suggest that planets which form by core accretion should exhibit increasing atmospheric *metallicity* (abundance of elements other than H/He) with decreasing mass (Fortney et al., 2013), a theory which appears to hold for solar system planets (Atreya et al., 2018). As well as overall metallicity, the relative abundances of different elements should also be affected by formation and evolution. A particularly useful diagnostic of a planet's history is its atmospheric carbon-to-oxygen (C/O) ratio. The gas C/O ratio varies strongly throughout the disk as various carbon- and oxygen-bearing species such as H₂O, CO₂ and HCN condense into ices (Wong et al., 2004), and so the atmospheric C/O ratio of a planet should be indicative

of where it formed in the disk (Öberg et al., 2011, Madhusudhan, 2012), and may also indicate whether migration occurred before or after the dispersal of the gaseous disk (Madhusudhan et al., 2014c). Constraining the chemical inventory of a planet's atmosphere through spectroscopic observations can therefore provide extremely valuable insight into the origin and growth of planetary systems.

1.2.4 Habitability and Biosignatures

One of the most tantalising prospects in planetary science is the detection of alien life on another planet. Although no such detection has yet been made, the growing field of astrobiology has improved our understanding of how life on other worlds might behave, and how it could be identified. Much of the study of astrobiology is concerned with determining planetary habitability: the measure of a planet's potential to support life. Many factors can affect habitability, including the nature of the host star, orbital distance from the host star, interactions with moons and other bodies, and galactic phenomena (Meadows & Barnes, 2018). Since all Earth-based life requires liquid water to survive, astrobiologists typically assume that the presence of liquid water is also a necessary condition for alien life. This led to the definition of the *habitable zone* as the range of orbital distances from a star at which a planet would be at the right temperature to host liquid water (Hart, 1979, Kasting et al., 1993). It is important to note that a planet being located in the habitable zone does not mean that it must be habitable, due to the other factors listed above.

Once a potentially habitable exoplanet has been identified, the next step is to search for substances or patterns that could only be of biological origin, which are collectively referred to as *biosignatures* (Des Marais et al., 2008). In the context of atmospheres, biosignatures are typically combinations of gases that would be impossible to produce through abiotic processes. Based on studies of the Earth's atmosphere, a wide range of potential biosignature gases have been proposed, including O₂, O₃, CH₄ and N₂O in various combinations (e.g. Lederberg, 1965, Léger et al., 1996, Des Marais et al., 2002). However, it is very challenging to rule out all possible abiotic sources for a particular atmospheric composition, and several studies have highlighted the possibility of false positives (e.g. Selsis et al., 2002, Wordsworth & Pierrehumbert, 2014). Therefore, it will be extremely important to assess the reliability of any biosignature gas that is detected in future.

Despite the large number of exoplanets detected to date, a habitable-zone Earth-like planet orbiting a Sun-like star has not yet been discovered. This is likely a result of the limitations of current telescopes, rather than a dearth of such planets (Petigura et al., 2013), and future observatories are expected to find the first Earth analogues (Rauer et al., 2014). Habitable-zone terrestrial planets have however been identified orbiting smaller stars such

as M dwarfs (e.g. Anglada-Escudé et al., 2016, Dittmann et al., 2017, Gillon et al., 2017). Current facilities have not been able to unambiguously detect an atmosphere on one of these planets, but it is hoped that upcoming observatories such as JWST and the Extremely Large Telescopes will be able to probe the atmospheres of habitable-zone terrestrial planets orbiting M dwarfs (Snellen et al., 2015, Lustig-Yaeger et al., 2019), meaning there is a real possibility of the first biosignature detection within the next decade.

1.3 Characterisation of Exoplanets

Relating the observations described above to meaningful information about the properties of an exoplanet requires connecting the observed data to theoretical models. Using measurements of bulk properties such as the mass and radius of a planet, we can begin to characterise a planet’s interior using internal structure models. Similarly, spectroscopic and phase curve observations can be combined with atmospheric modelling techniques to learn about properties of the atmosphere.

1.3.1 Internal Structure Models

Internal structure modelling has long been used to link a planet’s composition to its observable bulk properties (mass, radius and equilibrium temperature). These models typically compute the radius of a planet given its mass, composition and temperature profile by solving the planetary structure equations of mass continuity,

$$\frac{dR}{dM} = \frac{1}{4\pi R^2 \rho}, \quad (1.5)$$

where M is the mass of a spherical shell of material internal to a radius R and density ρ , and hydrostatic equilibrium (presented here using mass as the dependent variable),

$$\frac{dP}{dM} = -\frac{GM}{4\pi R^4}, \quad (1.6)$$

where P is the pressure at the shell. Linking these equations requires an equation of state (EOS) for each component of the interior, $\rho = \rho(P, T)$ as well as a pressure-temperature (P - T) profile $T = T(P)$, or simply $\rho = \rho(P)$ for a temperature-independent EOS.

Some of the earliest examples of such models were developed by Zepolsky & Salpeter (1969), who found mass–radius (M – R) relations for zero-temperature spheres made from a range of chemical species. Their methods have subsequently been developed further for exoplanets. Valencia et al. (2006) modelled Mercury to super-Earth sized planets with

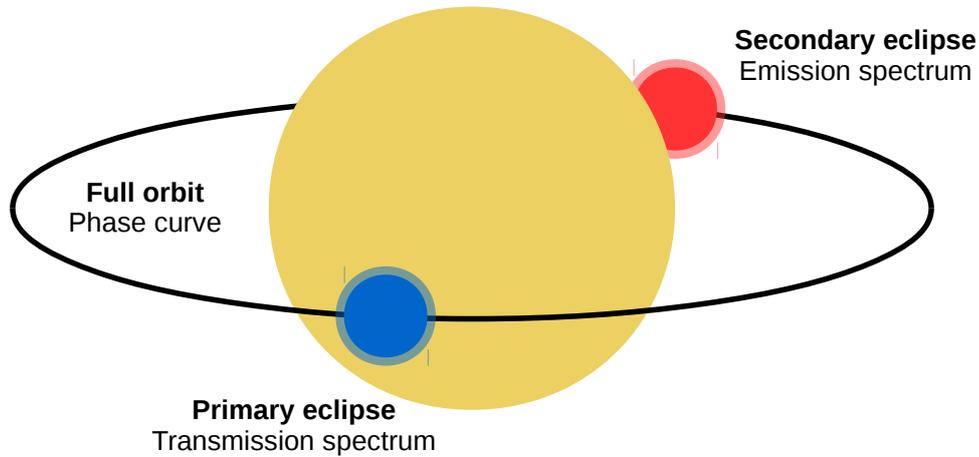


Fig. 1.8 Transiting planet with different observing modes highlighted. At primary eclipse, the wavelength-dependent decrease in the amount of stellar light reaching the observer can be measured (transmission spectroscopy). Just before secondary eclipse, thermal radiation from the dayside of the planet is observed (emission spectroscopy). By observing the system for a whole orbit, a phase curve can be acquired.

different core and mantle compositions incorporating thermal effects, finding that the $M-R$ relation differed depending on whether these planets are primarily rocky or icy. Seager et al. (2007) explored models of isothermal planets consisting of iron, silicates, water, and carbon compounds, as well as H/He, and noted a clear distinction between the radii of planets with gaseous envelopes and those without. $M-R$ relations for planets with sizes varying across several orders of magnitude were also computed by Fortney et al. (2007), who combined an ice+rock interior with a H/He envelope and calculated how planetary evolution affects the interior in order to link the age of a planetary system to the internal structure of its planets. For more detailed reviews, see e.g. Baraffe et al. (2014), Guillot et al. (2022). A technical overview of internal structure modelling is presented in Chapter 5.

1.3.2 Atmospheric Observations

The observational techniques discussed in Section 1.1.1 have been fundamental in the path to exoplanet discovery, with many of these methods providing information about bulk properties of the observed planets, such as mass and radius. While this information alone is sufficient to provide some insight into the possible composition of a planet (see Section 1.3.1), in order to perform detailed characterisation of an exoplanet we must observe its atmosphere. The observational methods described below and shown in Figure 1.8 enable the study of the chemical composition of an atmosphere as well as its temperature structure, and can also

provide information regarding dynamical processes such as winds or the presence of clouds and hazes. Atmospheric observations are a crucial means of determining a planet's formation and evolution history, as well as being one of the most promising avenues in the search for extraterrestrial life.

1.3.2.1 Transmission Spectroscopy

Transmission spectroscopy is closely related to the transit method for exoplanet detection (Section 1.1.1.2), and also involves observing a star-planet system during primary eclipse (while the planet passes in front of the star). Although in the previous discussion, the planet was assumed to be an opaque disk, in reality the planet's atmosphere will also block some of the stellar light passing through it. Consider equation 1.4 for a planet with an atmosphere of height H_a :

$$\delta \approx \left(\frac{R_p + H_a}{R_*} \right)^2 \approx \left(\frac{R_p}{R_*} \right)^2 + \frac{2H_a R_p}{R_*^2}, \quad (1.7)$$

where we have neglected the final term $(H_a/R_*)^2$ since the atmospheric height is assumed to be small compared to the stellar and planetary radii.

We can estimate the total height of the atmosphere by considering the equations of hydrostatic equilibrium, $dP/dz = -\rho g$, which relates pressure P , density ρ and gravity g to the height z in the atmosphere. This equation can be combined with the ideal gas law, $\mu P = k_B \rho T$, where T is the temperature of the atmosphere, μ is the mean molecular weight in kg and k_B is Boltzmann's constant. Assuming that T , ρ and g remain constant throughout the atmosphere, we can solve the combined equations to find that $P \propto e^{-z/H_{sc}}$, where $H_{sc} = k_B T / \mu g$ is the atmospheric scale height, i.e. the increase in height over which the pressure decreases by a factor of e .

The total height of the atmosphere can be approximated as $H_a = N H_{sc}$, where $N \approx 5$ -10 (Madhusudhan et al., 2014a). The atmospheric contribution is small compared to the total transit depth. For example, consider HD 209458b the first transiting exoplanet to be observed (Charbonneau et al., 2000). The planetary and stellar radii are $1.41 R_J$ and $1.2 R_\odot$ respectively (del Burgo & Allende Prieto, 2016), leading to a transit depth $\delta = 0.0146$. To estimate the atmospheric contribution we assume a temperature of 1450 K (Sing et al., 2016), mass of $0.74 M_J$ (del Burgo & Allende Prieto, 2016), implying a surface gravity of 9.2 m s^{-2} , and mean molecular weight of 2.3 amu (appropriate for a Hydrogen-rich atmosphere with solar elemental abundances, Asplund et al., 2009). This yields a scale height $H_{sc} = 564 \text{ km}$, meaning the additional contribution to the transit depth is approximately 8×10^{-4} , assuming

an atmosphere of 5 scale heights. This highlights the high levels of observational precision required to detect the atmosphere of an exoplanet. Additionally, the value of H_{sc} decreases for planets with lower temperatures and higher mean molecular weights, both of which would be expected for an Earth-like atmosphere, thus making atmospheric observations of terrestrial planets even more challenging.

The above calculations assume a uniformly opaque atmosphere. However, in reality the contribution of an atmosphere to the transit depth is a wavelength-dependent effect, since various chemical species present in the atmosphere will absorb more light at certain wavelengths than at others. We can introduce wavelength-dependence into our formulation of the transit depth by considering an effective atmospheric height $H_a = N_\lambda H_{sc}$, where N_λ is the number of scale heights at which the planet is effectively opaque at wavelength λ , which will depend on the composition of the atmosphere. This allows us to construct a wavelength-dependent transmission spectrum, δ_λ .

Note that, due to the observations taking place while the planet passes in front of the host star, the portion of the atmosphere probed in transmission spectroscopy corresponds to the terminator region between the day- and nightside of the planet. Since planets amenable to atmospheric characterisation in this way are typically on short orbital periods, they are also tidally locked with permanent day- and nightsides (Guillot et al., 1996). Therefore the location of the day-night terminator does not change throughout the planet's orbit.

The first detection of an exoplanetary atmosphere was achieved using transmission spectroscopy (Charbonneau et al., 2002), when absorption of atmospheric sodium was detected on the hot Jupiter HD 209458b using the Space Telescope Imaging Spectrograph (STIS) on HST, observing at optical wavelengths. Since then, exoplanet transmission spectra have been obtained for more than 50 exoplanets (Zhang et al., 2020). Observations continue to be acquired from space using HST, with the installation of a new instrument, Wide Field Camera 3 (WFC3), enabling observations in the near IR (Deming et al., 2013). IR transit depths have also been measured using the *Spitzer* Space Telescope (e.g. Tinetti et al., 2007, Désert et al., 2011). Furthermore, transmission spectra are regularly observed using ground-based facilities including the Very Large Telescope (VLT, e.g. Sedaghati et al., 2015, Wilson et al., 2020, Spyros et al., 2021), the Gran Telescopio Canarias (GTC, e.g. Sing et al., 2012, Chen et al., 2017, Murgas et al., 2020) and the *Magellan* Telescope (e.g. Rackham et al., 2017, Weaver et al., 2021).

More than a dozen chemical species have been identified in exoplanet atmospheres using transmission spectroscopy (Guillot et al., 2022). As well as sodium, the alkali metals potassium and lithium have been observed at optical wavelengths (e.g. Chen et al., 2018). A host of other metals including iron, magnesium and calcium have been detected with

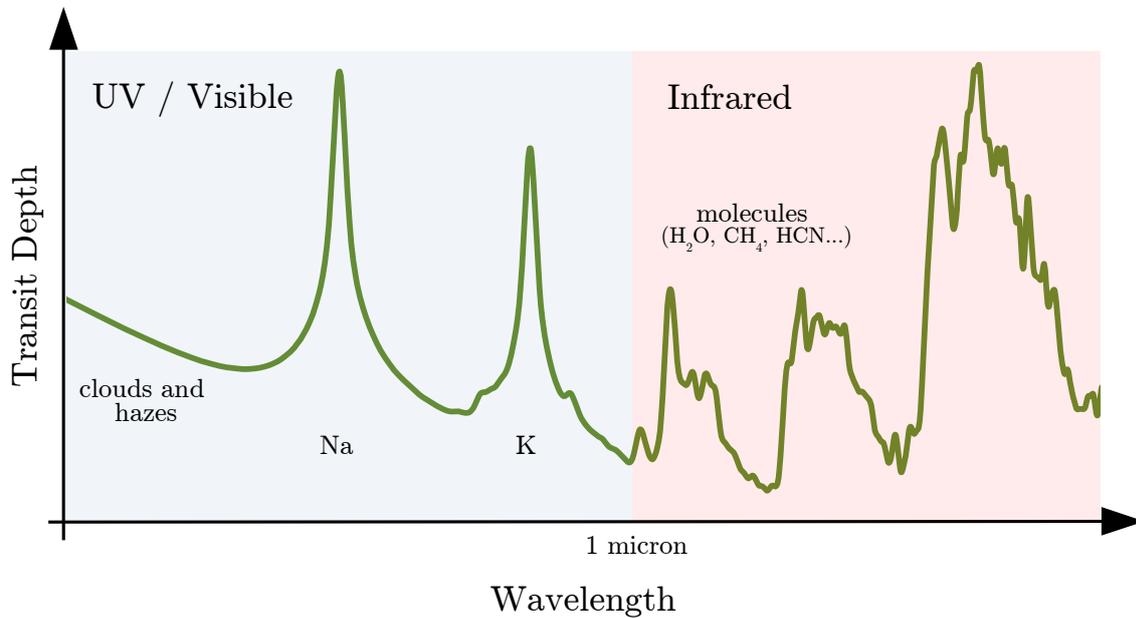


Fig. 1.9 Schematic representation of a transmission spectrum for a typical primary atmosphere, highlighting the wavelengths at which certain features may be observed. Optical wavelengths generally probe atomic species such as Na and K as well as clouds and hazes, whereas IR wavelengths are typically sensitive to molecules. Note however that some molecules such as TiO and VO have features at optical wavelengths.

high-resolution observations from ground-based facilities (e.g. Cabot et al., 2020, Deibert et al., 2021, Gibson et al., 2022). Infrared observations of transmission spectra have led to detections of molecules including H_2O , CO, and HCN (e.g. Snellen et al., 2010, Deming et al., 2013, MacDonald & Madhusudhan, 2017). A schematic representation of a model transmission spectrum is shown in Figure 1.9.

1.3.2.2 Emission Spectroscopy

Unlike transmission spectra, which are obtained as the planet passes in front of the host star, emission spectra are observed just before secondary eclipse, when the planet passes behind the star. At this moment, flux from the dayside of the planet will contribute to the total observed flux from the system. By comparing this to the flux received during secondary eclipse, when only stellar light contributes, we can infer the contribution from the planet. This is typically presented as the ratio of the planetary flux F_p to the stellar flux F_* , known as the eclipse depth. If we assume that the star and planet both emit as blackbodies at temperatures T_* and T_p respectively, then

$$\frac{F_p}{F_*} = \frac{R_p^2 B_\lambda(T_p)}{R_*^2 B_\lambda(T_*)}, \quad (1.8)$$

where $B_\lambda(T)$ is the Planck function:

$$B_\lambda(T) = \frac{2hc^2}{\lambda^5(\exp[hc/\lambda k_B T] - 1)}, \quad (1.9)$$

where h is Planck's constant and c is the speed of light. The signal described in equation 1.8 is maximised for large, hot planets such as hot Jupiters. Emission spectroscopy is typically conducted at infrared wavelengths, since planetary emission peaks at these wavelengths, whereas stellar emission peaks at optical wavelengths. At longer wavelengths (known as the Rayleigh-Jeans limit), the Planck function shown in equation 1.9 can be approximated as $B_\lambda(T) \approx 2ck_B T/\lambda^4$, meaning equation 1.8 becomes

$$\frac{F_p}{F_*} \approx \frac{R_p^2 T_p}{R_*^2 T_*}. \quad (1.10)$$

Considering once again the example of HD 209458b ($R_p = 1.41 R_J$, $T_p = 1450$ K, $R_* = 1.2 R_\odot$, $T_* = 6071$ K; del Burgo & Allende Prieto, 2016), we find an approximate eclipse depth of ~ 0.0035 , which is somewhat smaller than the transit depth for the same system.

Despite the smaller signal, emission spectra are capable of providing useful information about both the composition and temperature structure of a planet's dayside. Emission spectroscopy probes different layers of the atmosphere at different wavelengths: at wavelengths where absorption from chemical species in the upper atmosphere is more prominent, the emergent flux comes from these higher layers of the atmosphere, meaning observations are sensitive to the temperature near the top of the atmosphere. Conversely, at wavelengths where the atmosphere is relatively transparent, the observed flux comes from the deep atmosphere and is sensitive to the temperature in those regions. Emission spectroscopy can therefore be used to determine how temperature varies with height in the atmosphere, as well as providing insight into the atmospheric composition. Observed spectra have led to the identification of thermal inversions in several hot Jupiters (Haynes et al., 2015, Evans et al., 2017, Sheppard et al., 2017) as well as detections of chemical species including H_2O (e.g. Kreidberg et al., 2014b, Line et al., 2016, Evans et al., 2017, Arcangeli et al., 2018, Mansfield et al., 2022).

1.3.2.3 Phase Curves

The two methods discussed above are used to observe specific regions of a planet's atmosphere: the day-night terminator in the case of transmission spectroscopy, and the dayside in the case of emission spectroscopy. In contrast, phase curve observations can be used to obtain longitudinally-resolved properties of the entire surface of an exoplanet. In order to achieve this, the planet is observed throughout its whole orbit, with the observed flux at different times corresponding to different longitudes for tidally-locked planets. Phase curves can then be used to construct temperature maps of a planet (Cowan & Agol, 2008). Furthermore, if phase curves can be observed at multiple wavelengths, then emission spectra can be constructed as described above, but for multiple hemispheres of the planet (Knutson et al., 2009).

While phase curve observations (particularly those taken at multiple wavelengths) would appear to be the most informative class of atmospheric observations, they are typically much more challenging to obtain due to the considerable increase in observing time needed to acquire a full phase curve rather than a transmission or emission spectrum. This can be further hindered by a telescope's inability to observe the same target for long periods of time (e.g. coverage gaps when observing with HST due to the target passing behind the Earth), as well as other effects such as stellar ellipsoidal variations (see e.g. Parmentier & Crossfield, 2018).

Despite these challenges, successful phase curve observations have uncovered a wealth of information about a number of target planets, including the measurement of effects such as the shift of a planet's hotspot away from the substellar point due to winds (e.g. Knutson et al., 2007, Zhang et al., 2018), the temperature contrast between the day- and nightsides of a planet (e.g. Stevenson et al., 2017, Kreidberg et al., 2018, Bell et al., 2021), and inhomogeneities in the distribution of clouds in the atmosphere (e.g. Demory et al., 2013, Shporer & Hu, 2015).

1.3.2.4 Directly Imaged Spectra

Exoplanets which can be observed by direct imaging (Section 1.1.1.4) may also be amenable to direct observations of their emission spectra (Barman et al., 2011). In this case, rather than observing the planet/star flux ratio F_p/F_* , only the planet flux F_p is measured. Similarly to emission spectra obtained through secondary eclipse observations, these spectra are highly sensitive to both the temperature profile and chemical composition of a planet (Biller & Bonnefoy, 2018). However, directly imaged spectra can currently be obtained for just a handful of planets due to the challenges described in Section 1.1.1.4. Furthermore, since direct imaging of an exoplanet does not typically yield a measurement of the planet's mass, it can be difficult to explain the observed spectra using atmospheric models, which require

knowledge of the planet's surface gravity. Therefore it is challenging to infer properties of the atmosphere from directly imaged spectra, although a few chemical species have been detected (e.g. Konopacky et al., 2013, Bonnefoy et al., 2014, Chilcote et al., 2017).

1.3.3 Atmospheric Models

Theoretical models of exoplanet atmospheres used to explain observations typically follow one of two major paradigms: self-consistent forward modelling or inverse modelling, which is also referred to as atmospheric retrieval. A technical explanation of the process for constructing forward models will be presented in subsequent chapters. Here, we briefly summarise the key details of the two approaches, explaining their relative advantages and disadvantages.

1.3.3.1 Self-consistent Forward Models

Self-consistent forward models begin with a set of assumptions about the properties of a planet (such as its surface gravity and orbital semi-major axis) and work forward to create model observables such as spectra using our understanding of the physics and chemistry of atmospheres. The complexity of these models can vary widely depending on the desired application. Some of the earliest self-consistent models of exoplanet atmospheres were adapted from models for solar system planets (e.g. Marley & McKay, 1999) or stellar atmospheres (Seager & Sasselov, 1998). These models typically assumed a 1D atmosphere and would solve for the atmospheric temperature profile using the constraints of radiative-convective and hydrostatic equilibrium, given the level of stellar flux and heating from the planetary interior. This temperature profile and the assumption of chemical equilibrium is then used to determine atomic and molecular abundances, which in turn can alter the temperature profile by absorbing, emitting, or scattering photons. The model iterates over the coupled thermal and chemical properties of the atmosphere until converging to a self-consistent solution. Additional complexity may be added by considering processes such as chemical disequilibrium (e.g. Moses et al., 2011, Venot et al., 2012) or clouds and hazes (e.g. Piette & Madhusudhan, 2020b).

While 1D model parameters typically depend on height only, a number of key atmospheric processes such as dynamics (winds) require models whose properties can vary with latitude and longitude. General Circulation Models (GCMs) are detailed 3D models that simulate atmospheric dynamics, some also incorporating chemical processes along with radiative transport (Showman et al., 2020). Originally developed to simulate the Earth's atmosphere, a wide range of GCMs have been adapted for application to exoplanet atmospheres (Showman

& Guillot, 2002, Cooper & Showman, 2005, Lewis et al., 2010, Rauscher & Menou, 2010, Thrastarson & Cho, 2010, Wordsworth et al., 2011, Polichtchouk & Cho, 2012, Dobbs-Dixon & Agol, 2013, Kataria et al., 2013, Amundsen et al., 2016, Mendonça et al., 2016, Way et al., 2017, Deitrick et al., 2020). These models take various different approaches to computing atmospheric dynamics. In general, GCMs aim to solve the primitive equations of atmospheric dynamics in 3D across the entire planet, which are simplifications of the Navier-Stokes equations that assume hydrostatic equilibrium, a shallow atmosphere compared to the planetary radius, and constant gravity with height. However, some models instead opt to solve the full Navier-Stokes equations (e.g. Dobbs-Dixon & Agol, 2013, Mayne et al., 2014, Amundsen et al., 2016). GCMs have been used to explain phase curve observations, leading to inferences of day-night temperature contrasts and hotspot offsets (see Section 1.3.2.3).

1.3.3.2 Atmospheric Retrieval

In contrast to self-consistent forward modelling, where assumed atmospheric properties are used to compute an observed spectrum, atmospheric retrieval begins by considering a set of observations, and explores the range of model atmospheres that can explain the observed data. In this approach, atmospheric properties such as temperature structure and chemical abundances are described by a set of free parameters which are fed into a heavily simplified forward model. The model is combined with a statistical sampling algorithm which explores the full range of available combinations of parameter values in order to find those which yield the best-fitting models to the data. This procedure results in probability distributions for each model parameter, providing information about the likely properties of the observed atmosphere.

Atmospheric retrieval has become a ubiquitous tool for many kinds of exoplanetary atmospheric observations, including transmission spectra (e.g. Madhusudhan & Seager, 2009, Benneke & Seager, 2012, de Wit & Seager, 2013, Madhusudhan et al., 2014b, Kreidberg et al., 2014b, Waldmann et al., 2015a, Wakeford et al., 2017, Barstow et al., 2017, MacDonald & Madhusudhan, 2017, Pinhas et al., 2018, Zhang et al., 2019, Lacy & Burrows, 2020, Welbanks & Madhusudhan, 2021, Cubillos et al., 2022), emission spectra (e.g. Madhusudhan & Seager, 2009, Madhusudhan et al., 2011, Lee et al., 2012, Line et al., 2013, Waldmann et al., 2015b, Evans et al., 2017, Gandhi & Madhusudhan, 2018, Piette et al., 2022, Harrington et al., 2022), phase curves (e.g. Changeat & Al-Refaie, 2020, Feng et al., 2020, Irwin et al., 2020, Chubb & Min, 2022) and directly imaged spectra (e.g. Lee et al., 2013, Lupu et al., 2016, Lavie et al., 2017, Nayak et al., 2017). Retrieval frameworks also exist to analyse the atmospheres of solar system planets (e.g. Rodgers, 2000, Irwin et al., 2008) and brown dwarfs (e.g. Line et al., 2014, Burningham et al., 2017, Zalesky et al., 2019, Kitzmann et al., 2020,

Piette & Madhusudhan, 2020a). A more detailed description of the atmospheric retrieval algorithm and its application to transmission spectra is presented in Chapter 2.

Atmospheric retrieval requires the computation of a very large number of forward models in order to fully sample the space of possible atmospheres that could explain a given observation. One of the key trade-offs required to make atmospheric retrievals computationally feasible is that the requirement of thermo-chemical equilibrium is relaxed, meaning that the resulting atmospheric models may not be physically plausible. Atmospheric properties inferred from retrieval can be compared to self-consistent forward models in order to assess whether they are reasonable given our knowledge of atmospheric physics and chemistry. Conversely, retrieval results which disagree with self-consistent models could be an indication that further refinement of self-consistent models is required. Therefore, the two approaches can be seen as complementary methods for understanding the atmospheres of distant worlds.

1.3.4 Atmospheric Physics and Chemistry

The observational and modelling techniques described here can be used to infer a wide range of physical and chemical phenomena which occur in planetary atmospheres. The specific atmospheric regions and processes that are probed depend on the type of observation. Here we outline some key components of atmospheres which may be characterised through observations.

1.3.4.1 Chemical Composition

Planetary atmospheres can be grouped into two broad categories based on their chemical composition. Primary atmospheres, which are found around giant planets, are mostly inherited from the planetary nebula in which the planet formed. They consist primarily of hydrogen and helium, with a typical mean molecular weight of ~ 2.3 amu. In contrast, secondary atmospheres found around terrestrial planets are mostly generated through geological and surface processes. They consist of heavier elements such as nitrogen and oxygen, and therefore have a much larger mean molecular weight (e.g. ~ 30 amu for the Earth's atmosphere). As discussed previously, this makes terrestrial atmospheres more difficult to observe than giant planet atmospheres due to their smaller scale height.

Beyond their bulk composition, atmospheres also possess trace amounts of a range of other chemical species. The exact composition is determined by different processes at different locations in the atmosphere. We typically expect the deeper layers of a planet's atmosphere ($P \gtrsim 1$ bar) to be in or near chemical equilibrium. In this case it is possible to calculate the abundances of chemical species in the atmosphere by minimising the Gibbs free energy of

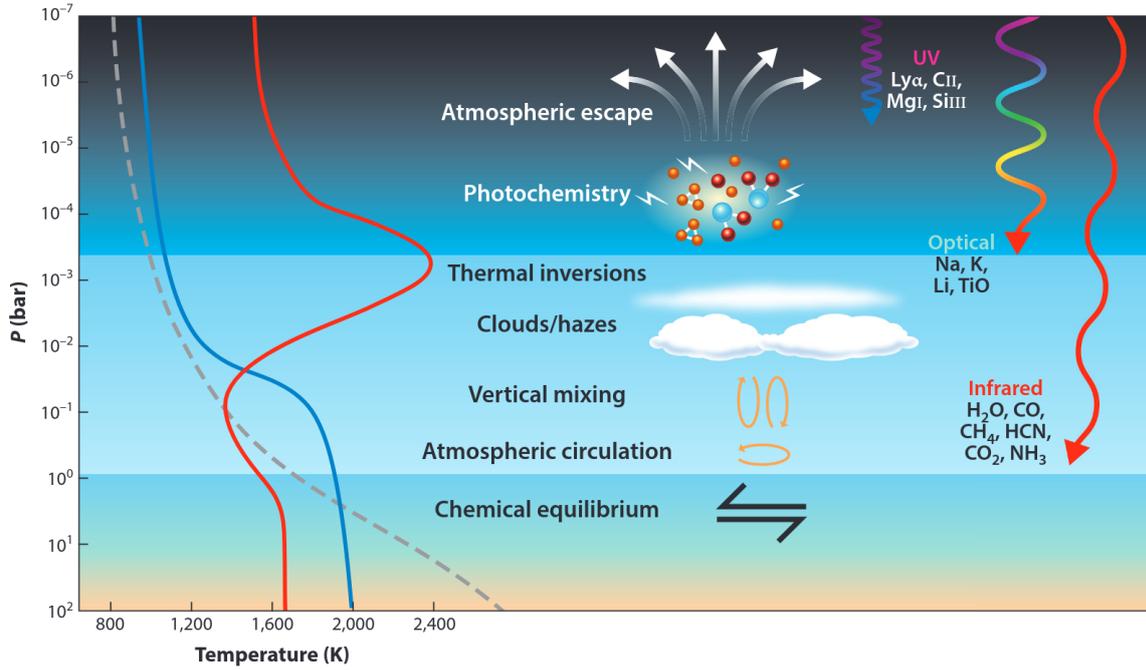


Fig. 1.10 Physical and chemical processes in exoplanetary atmospheres that may be probed by different regions of the electromagnetic spectrum. *Left*: Example temperature profiles for an irradiated planet with a thermal inversion (red), an irradiated planet without a thermal inversion (blue), and a poorly irradiated planet (grey). *Centre*: Processes active at various depths in the atmosphere. *Right*: Penetration depths of UV, optical and IR light, along with chemical species whose signatures are present in those wavelength ranges. Figure from Madhusudhan (2019).

the system given its temperature, pressure, and elemental abundances. This can be shown by considering the first law of thermodynamics for a generalised system:

$$dU = dQ - dW + \sum_i \mu_i dN_i, \quad (1.11)$$

where dU is the change in energy of the system, dQ is the heat supplied to the system, dW is the work done on the system, and μ_i and dN_i are the chemical potential and change in number of particles, respectively, of species i . This can be combined with the second law in the form $TdS \geq dQ$, where S is the entropy of the system, to give

$$dU \leq TdS - PdV + \sum_i \mu_i dN_i, \quad (1.12)$$

noting that $dW = PdV$ where P and V are the pressure and volume of the system. We now define the Gibbs free energy

$$\begin{aligned}
G &= U + PV - TS \\
\Rightarrow dG &= dU + PdV + VdP - TdS - SdT \\
&\leq VdP - SdT + \sum_i \mu_i dN_i,
\end{aligned} \tag{1.13}$$

meaning that at fixed P and T we have $dG \leq \sum_i \mu_i dN_i$, and in a steady state where N_i is constant, $dG \leq 0$, meaning G is minimised. For a system consisting of multiple chemical species, G can be expressed as

$$\frac{G}{RT} = \sum_i N_i \left(\frac{G_{i,0}}{RT} + \ln P + \ln X_i \right), \tag{1.14}$$

where R is the ideal gas constant, $G_{i,0}$ is the Gibbs free energy per mole at some reference pressure P_0 , and $X_i = N_i/N$ is the mixing ratio of species i . The mixing ratios can therefore be calculated by minimising G along with the assumption that the total number of atoms of each element in the system remains constant. A more detailed description of this process can be found in e.g. (Heng et al., 2016).

In primary atmospheres, the most prominent molecules that contribute to observable effects in spectra are those containing oxygen, carbon and nitrogen. Figure 1.11 shows the equilibrium mixing ratios at $P = 1$ bar of these molecules assuming abundances of O, C and N equal to those found in the Sun (Asplund et al., 2009).

In the upper atmosphere, a range of processes can cause chemical abundances to deviate from their equilibrium values. For example, at pressures lower than 1 bar it is possible that the timescale for chemical reactions may exceed the timescale for vertical diffusion in the atmosphere (Allen et al., 1981). The location at which this occurs is called the *quench level* (Visscher & Moses, 2011). Above the quench level, chemical abundances remain mostly constant due to vertical mixing, and so may not follow their equilibrium values. Higher again in the atmosphere, molecules may be subject to photodissociation due to incoming stellar photons, leading to new chemical processes which are collectively known as *photochemistry*. An example of a photochemical process is the formation of ozone in the Earth's atmosphere (Chapman, 1930). Above the photochemistry layer, the kinetic energies of gas particles in the atmosphere may exceed the potential energy of the planet, allowing for atmospheric escape to occur (Owen, 2019).

It is possible to infer the presence of various chemical species in an exoplanetary atmosphere using spectroscopic observations since a given species will preferentially absorb more light at certain wavelengths than at others. This is related to the intrinsic absorption

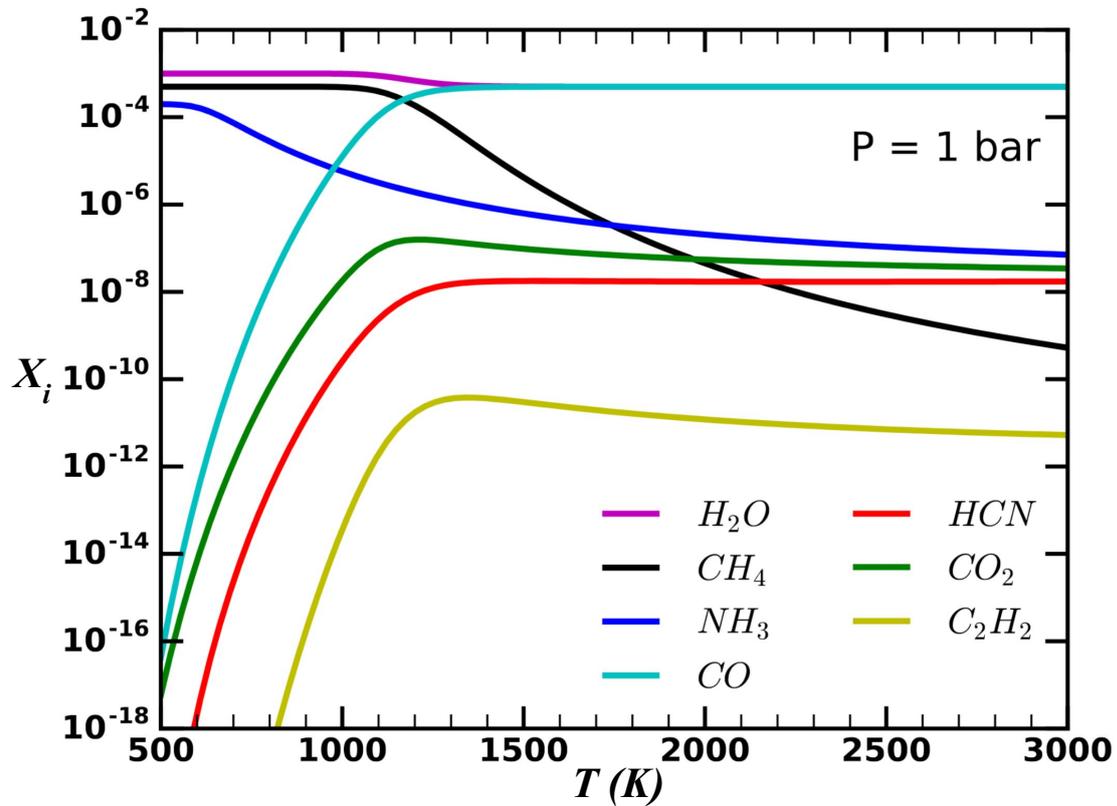


Fig. 1.11 Mixing ratios of key molecules present in hydrogen-dominated atmospheres in chemical equilibrium, assuming solar elemental abundances, at a nominal pressure of 1 bar. Figure adapted from Gandhi & Madhusudhan (2017).

cross-section of the species (see Chapter 2 for further detail). Key absorbers at UV, optical and IR wavelengths are listed in Figure 1.10.

1.3.4.2 Temperature Structure

The temperature structure of a planet's atmosphere is strongly affected by processes such as stellar radiation and the planet's internal heat flux, and is closely linked to chemical and dynamical properties (Guillot, 2010). For hot Jupiters, it is possible to probe how the temperature of a planet's dayside varies with height using emission spectroscopy (see Section 1.3.2.2). Of particular interest is the presence or absence of a thermal inversion, in which temperature increases with increasing height. A number of solar system planets exhibit thermal inversions, caused by the presence of hydrocarbon hazes in giant planets (Moses et al., 2005) or ozone in the Earth's atmosphere (Chapman, 1930). Although these species would not be expected to exist in the atmospheres of hot giant planets, a small number of thermal inversions have been inferred from emission spectra of hot Jupiters (e.g. Haynes et al., 2015,

Evans et al., 2017, Sheppard et al., 2017). Prior to these observations, it had been proposed that UV/visible absorption from species such as TiO and/or VO cause such inversions in highly irradiated planets (Hubeny et al., 2003), although it has also been suggested that other absorbers such as the H⁻ ion could be the cause (Arcangeli et al., 2018).

Temperature variations with latitude and longitude also play an important role in determining the structure of an atmosphere. Early GCMs of hot Jupiters predicted large temperature gradients between the day- and nightsides of the planet as well as an eastward equatorial jet, causing the hottest region of the planet to be shifted away from the substellar point (Showman & Guillot, 2002, Cooper & Showman, 2005). These predictions were later confirmed using phase curve observations (e.g. Knutson et al., 2007). With upcoming facilities, it is possible that effects such as day-night temperature contrasts may be detectable through other kinds of observations such as transmission spectroscopy (see Chapter 4).

1.3.4.3 Clouds and Hazes

The presence of clouds and hazes in an atmosphere, collectively referred to as aerosols, can strongly impact and often inhibit spectroscopic observations. Clouds are formed when vapour condenses onto some nucleus under certain thermodynamical conditions, whereas hazes are formed by non-condensation processes such as photochemistry (Marley et al., 2013). The exact observational signatures of aerosols depend on the wavelengths being observed as well as the nature of the observation. At visible wavelengths, clouds and hazes may lead to a steep downward slope in transmission spectra (Pont et al., 2013) or reflected stellar spectra in emission spectra (Marley et al., 2013). At IR wavelengths clouds and hazes typically lead to muted spectral features in both transmission and emission spectra, with the effect being particularly prominent for transmission spectra (Fortney, 2005).

Clouds and hazes have been proposed to explain atmospheric observations of numerous exoplanets across the mass range. A study of ten hot Jupiters conducted by Sing et al. (2016) invoked the presence of clouds to explain muted features of H₂O compared to what would be expected from chemical equilibrium with solar elemental abundances. However, later studies have suggested that these features may be a result of depleted H₂O abundances, with varied levels of clouds and hazes in the sample (Barstow et al., 2017, Pinhas et al., 2019). Additionally, the presence of high-altitude clouds and hazes has been used to explain flat transmission spectra of sub-Neptunes (e.g. Kreidberg et al., 2014a). These muted spectra have generally been observed for planets with equilibrium temperatures of ~400–600 K, with more prominent features being observed in planets both below (e.g. Benneke et al., 2019) and above (e.g. Guo et al., 2020) this temperature range. This tentative relationship between

haziness and temperature for sub-Neptunes is in agreement with laboratory experiments of haze formation (Yu et al., 2021).

1.4 Scope of this Thesis

In this thesis, we present new developments in atmospheric modelling and retrieval of exoplanet transmission spectra, as well as new internal structure models for water-rich sub-Neptunes. We present a new atmospheric retrieval framework based on machine learning techniques, and demonstrate that it can robustly reproduce the results of existing methods, while improving computational efficiency in some cases. We develop a new atmospheric forward model that can compute the three-dimensional temperature structure of a planet while remaining computationally efficient for retrievals. We also present a new internal structure model for exoplanets with a detailed EOS for H₂O, enabling characterisation of water-rich sub-Neptunes.

1.4.1 Atmospheric Retrieval of Transmission Spectra

In Chapter 2, we introduce a typical atmospheric retrieval framework in detail, focusing on transmission spectroscopy. We begin by presenting the radiative transfer calculations for the forward model, describing key components such as opacity from chemical species, clouds and hazes, and stellar heterogeneity. We subsequently discuss the underlying principles behind Bayesian sampling algorithms, particularly Nested Sampling, which are used to link the forward model to an observed data set. We review key results from atmospheric retrieval analyses and highlight possible limitations.

1.4.2 Supervised Machine Learning for Atmospheric Retrieval of Exoplanets

Atmospheric retrieval of exoplanets from spectroscopic observations requires an extensive exploration of a highly degenerate and high-dimensional parameter space to accurately constrain atmospheric parameters. Retrieval methods commonly conduct Bayesian parameter estimation and statistical inference using sampling algorithms such as Markov Chain Monte Carlo (MCMC) or Nested Sampling. Recently several attempts have been made to use machine learning algorithms either to complement or replace fully Bayesian methods. While these approaches have found some success, they are still at times unable to accurately reproduce results from contemporary Bayesian retrievals. The goal of the work presented

in Chapter 3 is to investigate the efficacy of machine learning for atmospheric retrieval. As a case study, we use the Random Forest supervised machine learning algorithm which has been applied previously for atmospheric retrieval of the hot Jupiter WASP-12b using its near-infrared transmission spectrum. We reproduce previous results using the same approach and the same semi-analytic models, and highlight its limitations. We subsequently extend the machine learning approach to develop a new algorithm that results in a better match to a fully Bayesian retrieval. We combine this new method with a fully numerical model and demonstrate excellent agreement with a Bayesian retrieval of the transmission spectrum of another hot Jupiter, HD 209458b. Despite this success, and achieving high computational efficiency, we still find that the machine learning approach is computationally prohibitive for high-dimensional parameter spaces that are routinely explored with Bayesian retrievals with modest computational resources. We discuss the trade offs and potential avenues for the future.

1.4.3 A Three-Dimensional Retrieval Framework for Exoplanet Transmission Spectra

Atmospheric retrievals of exoplanet transmission spectra allow constraints on the composition and structure of the day-night terminator region. Such retrievals in the past have typically assumed one-dimensional temperature structures which were adequate to explain extant observations. However, the increasing data quality expected from exoplanet spectroscopy with JWST motivates considerations of multidimensional atmospheric retrievals. In Chapter 4 we present AURA-3D, a three-dimensional atmospheric retrieval framework for exoplanet transmission spectra. AURA-3D includes a forward model that enables rapid computation of transmission spectra in 3D geometry for a given atmospheric structure and can, therefore, be used for atmospheric retrievals as well as for computing spectra from GCMs. In order to efficiently explore the space of possible 3D temperature structures in retrievals, we develop a parametric P - T profile which can accurately represent azimuthally-averaged temperature structures of a range of hot Jupiter GCMs. We apply our retrieval framework to simulated JWST observations of hot Jupiter transmission spectra, obtaining accurate estimates of the day-night temperature variation across the terminator as well as the abundances of chemical species. We demonstrate an example of a model hot Jupiter transmission spectrum for which a traditional 1D retrieval of JWST-quality data returns biased abundance estimates, whereas a retrieval including a day-night temperature gradient can accurately retrieve the true abundances. Our forward model also has the capability to include inhomogeneous chemistry as well as variable clouds/hazes. This new retrieval framework opens the field to detailed

multidimensional atmospheric characterisation using transmission spectra of exoplanets in the JWST era.

1.4.4 Exploring the Phase Structure of Water-Rich Sub-Neptunes

Understanding the internal structures of planets with a large H₂O component is important for the characterisation of sub-Neptune planets. The finding that the mini-Neptune K2-18b could host a liquid water ocean beneath a mostly hydrogen envelope motivates a detailed examination of the phase structures of water-rich planets. To this end, in Chapter 5 we present new internal structure models for super-Earths and mini-Neptunes that enable detailed characterisation of a planet's water component. We use our models to explore the possible phase structures of water worlds and find that a diverse range of interiors are possible, from oceans sandwiched between two layers of ice to supercritical interiors beneath steam atmospheres. We determine how the bulk properties and surface conditions of a water world affect its ocean depth, finding that oceans can be up to hundreds of times deeper than on Earth. For example, a planet with a 300 K surface can possess H₂O oceans with depths from 30–500 km, depending on its mass and composition. We also constrain the region of mass–radius space in which planets with H/He envelopes could host liquid H₂O, noting that the liquid phase can persist at temperatures up to 647 K at high pressures of 218– 7×10^4 bar. Such H/He envelopes could contribute significantly to the planet radius while retaining liquid water at the surface, depending on the planet mass and temperature profile. Our findings highlight the exciting possibility that habitable conditions may be present on planets much larger than Earth.

1.4.5 Internal Structures of JWST Targets in the sub-Neptune Regime

JWST is set to revolutionise our understanding of exoplanets smaller than Neptune, providing much higher quality spectra than have so far been accessible. This has been recognised through the allocation of numerous early JWST observing programs dedicated to observations of sub-Neptune atmospheres. Before such observations take place, it is vitally important to understand the interior structures of these planets, which strongly affects their possible atmospheric compositions. In Chapter 6, we present applications of internal structure models to a number of upcoming JWST targets. We use the bulk parameters and retrieved atmospheric properties to constrain the internal structure and thermodynamic conditions in the habitable-zone mini-Neptune K2-18b. The constraints on the interior allow multiple scenarios between rocky worlds with massive H/He envelopes and water worlds with thin envelopes. We constrain the mass fraction of the H/He envelope to be $\lesssim 6\%$; spanning $\lesssim 10^{-5}$ for a

predominantly water world to $\sim 6\%$ for a pure iron interior. The thermodynamic conditions at the surface of the H_2O layer range from the supercritical to liquid phases, with a range of solutions allowing for habitable conditions. We also investigate the possible compositions of the pair of planets orbiting the star TOI-776. The bulk densities of TOI-776b and c allow for a wide range of possible interior and atmospheric compositions. However, we find that both planets must have retained a significant atmosphere, with slightly different envelope mass fractions. We discuss how upcoming observations will revolutionise our understanding of these planets, helping to uncover the mysteries of the sub-Neptune population.

1.4.6 Discussion and Conclusions

Finally, in Chapter 7 we summarise the new developments presented in this thesis and offer some concluding remarks. We consider some of the key questions that remain open in the field of exoplanetary characterisation, and consider potential future avenues of research that may be opened up by upcoming technological developments.

Chapter 2

Atmospheric Retrieval of Transmission Spectra

“Astronomers, like burglars and jazz musicians, operate best at night.”

– Miles Kington

Atmospheric retrieval has been a relatively recent development in exoplanetary science. Initial spectroscopic observations of exoplanets were challenging to interpret due to low resolution and low signal-to-noise (e.g. Tinetti et al., 2007, Grillmair et al., 2008, Swain et al., 2008). Inferences of molecular features or temperature structures such as thermal inversions were made by comparing a small number of forward models to the observations that appeared to provide a good fit. This approach failed to explore broad regions of the parameter space, and did not account for degeneracies between various atmospheric model parameters. The need to introduce a more robust means of inferring atmospheric properties from spectra led to the development of atmospheric retrieval for exoplanets (Madhusudhan & Seager, 2009), inspired by similar techniques which had been employed to analyse Earth-based (Rodgers, 2000) and solar system (Irwin et al., 2008) data. As the quality of available data has improved, retrieval methodologies have evolved in tandem to efficiently find the span of models that can explain new observations (Madhusudhan, 2018).

In this chapter we provide a technical introduction to the theory of atmospheric retrieval, as introduced in Section 1.3.3.2. The algorithm combines a forward modelling code with a parameter estimation method. We largely follow the methodology of the AURA retrieval code (Pinhas et al., 2018). The forward model from AURA is combined with a new parameter estimation scheme using machine learning in Chapter 3, while the forward model is extended to allow for the calculation of three-dimensional model atmospheres in Chapter 4.

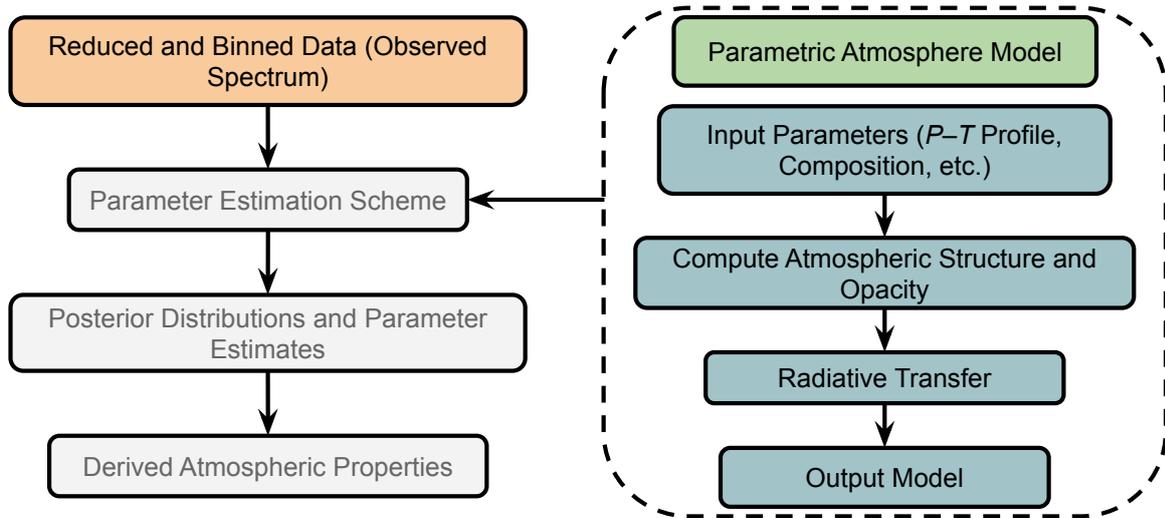


Fig. 2.1 Schematic of atmospheric retrieval. The algorithm combines a parametric model of an atmosphere with a parameter estimation scheme to derive model parameters given an observed spectrum. Typical model parameters describe the temperature structure and chemical composition of the atmosphere as well as other properties such as clouds/hazes. The parameter estimation scheme usually allows for computation of posterior probability distributions for each parameter, which can be used to derive quantities such as elemental abundance ratios from those of molecular abundances.

The process of atmospheric retrieval involves estimating properties of an atmosphere given some observed spectral data. Typically a large number of models are produced and compared to observations in order to find those which can most accurately describe the data (Fortney et al., 2021). The structure of a typical retrieval algorithm is shown in Figure 2.1. In the remainder of this chapter we describe each of the components of a retrieval in detail. We also demonstrate a retrieval of the transmission spectrum of the canonical hot Jupiter HD 209458b, observed using HST (Deming et al., 2013, Sing et al., 2016).

2.1 Model Transmission Spectra

2.1.1 Radiative Transfer

In order to model the spectrum of an astronomical object we must understand how radiation from that object travels to us. For an object emitting energy, the *specific intensity* is defined as

$$I_\nu \equiv \frac{dE}{dA_n dt d\nu d\Omega} \text{erg cm}^{-2} \text{s}^{-1} \text{Hz}^{-1} \text{sr}^{-1}. \quad (2.1)$$

In other words, it is the energy emitted from the source per unit area, per unit time, per unit frequency, per unit solid angle. Here dA_n is the area normal to the beam, which can be written as $dA \cos \vartheta$, where ϑ is the angle between the beam and the normal to the emitting surface. Likewise, the specific intensity received by an observer will be

$$I'_\nu = \frac{dE'}{dA'_n dt dv d\Omega'}. \quad (2.2)$$

For the observer, the solid angle subtended by the source is

$$d\Omega' = \frac{dA \cos \vartheta}{d^2}, \quad (2.3)$$

where d is the distance to the source. Since we assume all of the energy goes to the observer, we can also get an expression for the solid angle subtended by the source

$$d\Omega = \frac{dA' \cos \vartheta'}{d^2}. \quad (2.4)$$

By conservation of energy, we must have $dE = dE'$, so we find that

$$\begin{aligned} \frac{I_\nu dA \cos \vartheta dt dv dA' \cos \vartheta'}{d^2} &= \frac{I'_\nu dA' \cos \vartheta' dt dv dA \cos \vartheta}{d^2} \\ \implies I_\nu &= I'_\nu, \end{aligned} \quad (2.5)$$

so specific intensity does not vary with distance.

The *flux* emitted by an object is defined as

$$F_\nu \equiv \int_{\Omega} I_\nu d\Omega, \quad (2.6)$$

meaning

$$dF_\nu = \frac{dE}{dA_n dt dv}. \quad (2.7)$$

Strictly speaking, F_ν is a vector quantity pointing in the direction $\hat{\mathbf{n}}$ of the beam of radiation, but this is not always written. Using the area element in spherical coordinates, $d\Omega = \sin \vartheta d\vartheta d\phi$,

we have

$$\begin{aligned}
 F_\nu &= \int_{\Omega} I_\nu d\Omega \hat{\mathbf{n}} \cdot \hat{\mathbf{z}} \\
 &= \int_0^{2\pi} \int_0^{\frac{\pi}{2}} I_\nu \sin \vartheta \cos \vartheta d\vartheta d\phi \\
 &= 2\pi \int_0^1 \mu I_\nu d\mu,
 \end{aligned} \tag{2.8}$$

where we have used the substitution $\mu = \cos \vartheta$. In the case where I_ν is isotropic we therefore have

$$F_\nu = \pi I_\nu. \tag{2.9}$$

The above equations dealt with the surface flux, meaning the flux emitted from the source. Let us now consider the observed flux. When we observe a distant spherical object such as a star or planet we see a circle. We want to integrate along the radius of this circle, which is given by $x = R \sin \vartheta$. The incremental area is then $dA = 2\pi x dx$, so the solid angle subtended by the object on the sky will be

$$\begin{aligned}
 d\Omega &= \frac{2\pi x dx}{d^2} \\
 &= \frac{2\pi R \sin \vartheta R \cos \vartheta d\vartheta}{d^2}
 \end{aligned} \tag{2.10}$$

where d is the distance to the object. Substituting this into the definition of flux, we obtain

$$\begin{aligned}
 F_\nu &= 2\pi \int_0^{\frac{\pi}{2}} I_\nu \frac{R^2 \sin \vartheta \cos \vartheta}{d^2} d\vartheta \\
 &= \pi I_\nu \frac{R^2}{d^2},
 \end{aligned} \tag{2.11}$$

so we see that flux is dependent on both the distance to the source and the radius of the source. We want to know how I_ν changes as it passes through a medium. Consider a beam of light passing through a medium of thickness dl , cross-sectional area dA and density ρ . This medium has mass defined as $dm = \rho dA ds$. The specific intensity going into the medium is I_ν , and coming out of the medium is $I'_\nu = I_\nu + dI_\nu$. The three main processes to take into account are absorption, scattering and emission. The change in energy due to extinction

(absorption and scattering) is

$$dE_{ex} = k_\nu I_\nu d\Omega dv dt dm, \quad (2.12)$$

where k_ν is the extinction coefficient, defined per unit mass of material in the medium. For emission we have

$$dE_{em} = j_\nu I_\nu d\Omega dv dt dm, \quad (2.13)$$

with emission coefficient j_ν . Note that, unlike absorption and scattering, emission does not depend on I_ν (since the light is coming from the medium, not the source) and it adds to the specific intensity, whereas the other two processes are extinction mechanisms. Putting all this together we get

$$\begin{aligned} dE &= dE_{em} - dE_{ex} \\ &= [j_\nu - k_\nu I_\nu] d\Omega dv dt \rho dA dl \\ \implies dI_\nu &= [j_\nu - k_\nu I_\nu] \rho dl. \end{aligned} \quad (2.14)$$

This is the equation of radiative transfer. In the case where there is no emitting source ($j_\nu = 0$), the equation reduces to

$$\begin{aligned} \frac{dI_\nu}{\rho dl} &= -k_\nu I_\nu \\ \implies I_\nu &= I_0 e^{-\int k_\nu \rho dl} \\ &= I_0 e^{-\tau_\nu} \end{aligned} \quad (2.15)$$

where we have defined the optical depth, $\tau_\nu = \int k_\nu \rho dl$. This can also be expressed as $\tau_\nu = \int n \sigma_\nu dl$ where n is the number density of particles in the medium and σ_ν is the absorption cross-section.

2.1.2 Transmission Spectra

We have already derived an expression for the observed flux from a source, such as a star, of radius R at distance d :

$$F_\lambda = \pi I_\lambda \frac{R^2}{d^2}, \quad (2.16)$$

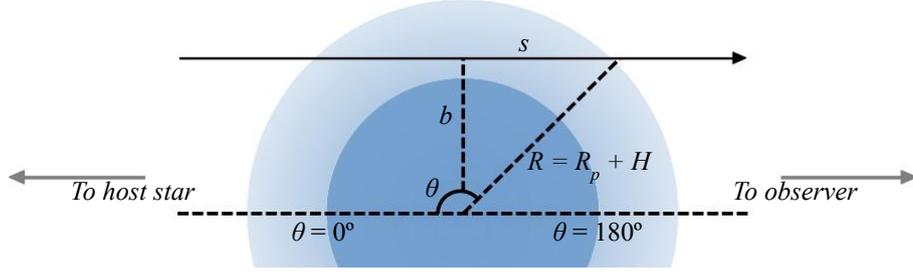


Fig. 2.2 Geometry of an exoplanet atmosphere as observed in transit, from Nixon & Madhusudhan (2022). A ray passes through the atmosphere along the path s , travelling through the day-night terminator region which covers a range of values of the zenith angle θ .

where we have now changed from frequency to wavelength dependence using the fact that $c = \nu\lambda$ for electromagnetic radiation. Now we consider the case where we are observing a star with a planet in front of it, so that the planet blocks some of the stellar radiation (see Figure 2.2). We will assume a 90° inclination. Our expression for the flux is now

$$F_\lambda = \int_0^{R_{pl}} I_{\lambda,pl} \frac{2\pi x}{d^2} dx + \int_{R_{pl}}^{R_*} I_{\lambda,*} \frac{2\pi x}{d^2} dx. \quad (2.17)$$

Since the night side of the planet will be facing us, the first term is negligible, so we can approximate

$$F_\lambda \approx \pi I_{\lambda,*} \frac{R_*^2 - R_{pl}^2}{d^2}. \quad (2.18)$$

We can use this to find the transit depth, i.e. the (normalised) difference between the flux in and out of transit, as previously seen in Chapter 1:

$$\Delta(\lambda) = \frac{F_{\lambda,out} - F_{\lambda,in}}{F_{\lambda,out}} = \frac{R_{pl}^2}{R_*^2}. \quad (2.19)$$

Now suppose the planet has an atmosphere extending to a height H above the surface of the planet. Unlike the derivation in Chapter 1, we will not assume that the atmosphere is entirely opaque. The in-transit flux then becomes

$$F_{\lambda,in} \approx \int_{R_{pl}}^{R_{pl}+H} I_{\lambda,pl}^a \frac{2\pi x}{d^2} dx + \int_{R_{pl}+H}^{R_*} I_{\lambda,*} \frac{2\pi x}{d^2} dx, \quad (2.20)$$

where $I_{\lambda,pl}^a$ is the intensity of light passing through the planetary atmosphere. We have again assumed that the $I_{\lambda,pl}$ term is negligible. Using equation 2.15 we can calculate the difference between the flux in and out of transit:

$$\begin{aligned}
F_{\lambda,out} - F_{\lambda,in} &= \int_0^{R_*} I_{\lambda,*} \frac{2\pi x}{d^2} dx - \int_{R_{pl}}^{R_{pl}+H} I_{\lambda,*} e^{-\tau_\lambda(b)} \frac{2\pi x}{d^2} dx - \int_{R_{pl}+H}^{R_*} I_{\lambda,*} \frac{2\pi x}{d^2} dx \\
&= \int_0^{R_{pl}+H} I_{\lambda,*} \frac{2\pi x}{d^2} dx - \int_{R_{pl}}^{R_{pl}+H} I_{\lambda,*} e^{-\tau_\lambda(b)} \frac{2\pi x}{d^2} dx \\
&= \int_0^{R_{pl}} I_{\lambda,*} \frac{2\pi x}{d^2} dx + \int_{R_{pl}}^{R_{pl}+H} I_{\lambda,*} \left(1 - e^{-\tau_\lambda(b)}\right) \frac{2\pi x}{d^2} dx \\
&= I_{\lambda,*} \frac{2\pi R_{pl}^2}{d^2} + \int_0^H I_{\lambda,*} \left(1 - e^{-\tau_\lambda(b)}\right) \frac{2\pi(R_{pl} + b)}{d^2} db, \tag{2.21}
\end{aligned}$$

where in the last line we have made the change of variable $x = R_{pl} + b$, meaning b is the impact parameter, or perpendicular height above the planetary surface. Dividing by the in-transit flux we obtain the transit depth:

$$\Delta_\lambda = \frac{R_{pl}^2 + A_\lambda}{R_*^2}, \tag{2.22}$$

where A_λ is the atmospheric depth, given by

$$A_\lambda = 2 \int_0^H (R_{pl} + b) \left(1 - e^{-\tau_\lambda(b)}\right) db. \tag{2.23}$$

A similar derivation, along with the derivation of the form of an emission spectrum, can be found in Seager & Deming (2010).

2.1.3 Opacity Sources

In order to compute A_λ , the optical depth $\tau_\lambda(b)$ must be found. This is dependent on the chemistry and the cloud properties in the atmosphere. The total optical depth at each impact parameter and wavelength can be written as an integral along a ray path s :

$$\tau_\lambda(b) = \int \mu_\lambda(s, b) ds, \tag{2.24}$$

where μ_λ is the *attenuation coefficient*. The total attenuation coefficient can be computed by summing contributions from various opacity sources:

$$\mu_\lambda = \sum_{i \in \mathcal{O}} \mu_{\lambda,i}, \quad (2.25)$$

where \mathcal{O} is the set of opacity sources in the model. Here we consider three of the most important sources of opacity present in most atmospheres: extinction from chemical species, collision-induced absorption and cloud/haze opacity. Additional effects such as stellar heterogeneity and refraction can also affect the transit depth, and are discussed in Chapter 4.

2.1.3.1 Trace chemical species

For a given chemical species i , the attenuation coefficient can be written in terms of the mass density ρ_i and opacity κ_i , or equivalently in terms of number density n_i and absorption cross-section σ_i :

$$\mu_{\lambda,i}(P, T) = \rho_i \kappa_{\lambda,i}(P, T) = n_i \sigma_{\lambda,i}(P, T). \quad (2.26)$$

The number density of a species is related to its volume mixing ratio, $X_i = n_i/n_{\text{tot}}$, where n_{tot} is the total number density of the gas.

The wavelength-dependent absorption cross-sections of different chemical species are derived using knowledge of the electronic transitions that can occur when photons are absorbed by that species. A selection of molecular cross-sections is shown in Figure 2.3. Cross-sections also vary with temperature and pressure, which contribute to the broadening of transition lines (Gandhi & Madhusudhan, 2017). In order to find the cross-section for a given species at some wavelength, temperature and pressure, a list of transition lines and strengths for that species is required. The model presented in this chapter includes line lists from the HITEMP database for H₂O, CO and CO₂ (Rothman et al., 2010) and the Exomol database for HCN, CH₄ and NH₃ (Tennyson et al., 2016). Cross-sections for Na and K are adopted from Welbanks et al. (2019) using line data from Allard et al. (2019) for Na and Allard et al. (2016) for K. The procedure for calculating cross-sections from line lists is summarised here; more detailed descriptions can be found in Hedges & Madhusudhan (2016) and Gandhi & Madhusudhan (2017).

Line list data bases typically give Einstein coefficients with energies and degeneracies for each state. These can be converted to line strengths at some reference temperature T_{ref} using the following formula (Rothman et al., 2013):

$$S_{i,j}(T_{\text{ref}}) = \frac{A_{i,j}}{8\pi c \nu_{i,j}^2 Q(T_{\text{ref}})} g_i \exp\left[-\frac{hcE_j}{k_B T_{\text{ref}}}\right] \left(1 - \exp\left[-\frac{hc\nu_{i,j}}{k_B T_{\text{ref}}}\right]\right), \quad (2.27)$$

where $A_{i,j}$ is the Einstein coefficient for spontaneous emission for the transition between states i and j , g_i is the upper state degeneracy, E_j is the lower state energy, $\nu_{i,j}$ is the transition frequency between states i and j , and h is Planck's constant. The partition function $Q(T)$ is given by

$$Q(T) = \sum_j g_j \exp\left[-\frac{E_j}{k_B T}\right], \quad (2.28)$$

where g_j is the lower state degeneracy. The line strength at any temperature T can be converted from the strength at T_{ref} shown in equation 2.27 with the formula

$$S_{i,j}(T) = S_{i,j}(T_{\text{ref}}) \frac{Q(T_{\text{ref}})}{Q(T)} \frac{\exp(-hcE_j/k_B T)}{(-hcE_j/k_B T_{\text{ref}})} \frac{[1 - \exp(-hc\nu_{i,j}/k_B T)]}{[1 - \exp(-hc\nu_{i,j}/k_B T_{\text{ref}})]}. \quad (2.29)$$

To obtain cross-sections from line strengths, the effects of thermal and pressure broadening must be taken into account. Thermal broadening is a result of the different velocities of atoms or molecules which lead to Doppler shifting of transition lines, ultimately blurring the line into a Gaussian shape with a profile f_G , which for a molecule of mass m with a line centred at frequency ν_0 is expressed as

$$f_G(\nu - \nu_0) = \frac{1}{\gamma_G \sqrt{\pi}} \exp\left(-\frac{(\nu - \nu_0)^2}{\gamma_G^2}\right), \quad (2.30)$$

$$\gamma_G \equiv \sqrt{\frac{2k_B T}{m}} \frac{\nu_0}{c}. \quad (2.31)$$

The line is also broadened due to gas pressure, yielding a Lorentzian profile at pressure P :

$$f_L(\nu - \nu_0) = \frac{1}{\pi} \frac{\gamma_L}{(\nu - \nu_0)^2 + \gamma_L^2}, \quad (2.32)$$

$$\gamma_L \equiv \left(\frac{T_{\text{ref}}}{T}\right)^n P \sum_b \gamma_{L,b} p_b, \quad (2.33)$$

where n is a temperature scaling factor, p_b is the partial pressure of a broadening molecule b and $\gamma_{L,b}$ is the Lorentzian half width half maximum for molecule b (see e.g. Rothman et al., 1998, for further detail). The full broadening is obtained through the convolution of

the Gaussian and Lorentzian profiles, yielding the Voigt function

$$f_V(\nu - \nu_0) = \int_{-\infty}^{\infty} df_G(\nu' - \nu_0) f_L(\nu - \nu') \nu', \quad (2.34)$$

and the cross-section at frequency ν for a particular line is given by

$$\sigma_\nu = S(T) f_V(\nu). \quad (2.35)$$

2.1.3.2 Collision-induced absorption

Inelastic collisions between chemical species can induce quantum transitions that also contribute to the opacity of an atmosphere. This phenomenon is known as *collision-induced absorption* (CIA). In the atmospheres of giant planets, the main sources of CIA are H₂-H₂ and H₂-He interactions, which contribute the following attenuation coefficient:

$$\mu_{\text{CIA}} = X_{\text{H}_2} n_{\text{tot}}^2 [X_{\text{H}_2} \sigma_{\text{H}_2-\text{H}_2}(\lambda, T) + X_{\text{He}} \sigma_{\text{H}_2-\text{He}}(\lambda, T)]. \quad (2.36)$$

The relevant cross-sections, $\sigma_{\text{H}_2-\text{H}_2}$ and $\sigma_{\text{H}_2-\text{He}}$, are obtained from the HITRAN database (Richard et al., 2012). The mixing ratios of H₂ and He are not treated as free parameters in the model; they are instead determined by assuming a solar composition of $X_{\text{He}}/X_{\text{H}_2} = 0.17$ (Asplund et al., 2009) and using the fact that the sum of all mixing ratios must equal unity, so

$$X_{\text{H}_2} = \frac{1 - \sum_s X_s}{1 + X_{\text{He}}/X_{\text{H}_2}}, \quad (2.37)$$

where the sum is over all species s other than H₂ and He.

2.1.3.3 Cloud and haze opacity

As discussed in Section 1.3.4.3, clouds and hazes can have a considerable impact on the shape of a transmission spectrum. It is therefore important to incorporate a prescription for clouds and hazes into any retrieval framework. In this chapter, we introduce the parameterisation from MacDonald & Madhusudhan (2017) following Line & Parmentier (2016) that enables consideration of inhomogeneous cloud cover across the terminator region. Later, in Section 4.2.5.1, we present an extension of this parameterisation developed by Welbanks & Madhusudhan (2021).

We assume that clouds consist of an opaque cloud deck at pressures above some P_{cld} and scattering due to hazes pressures below (meaning altitudes above) P_{cld} . The contribution to

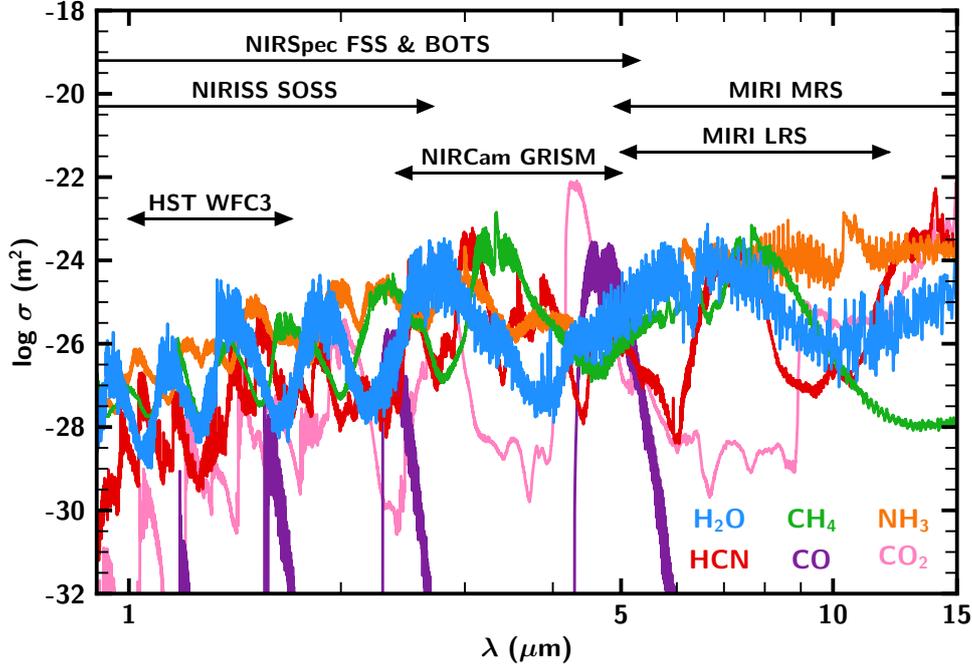


Fig. 2.3 Cross-sections of prominent molecular species in hydrogen-rich atmospheres at $P = 1$ bar and $T = 1000$ K. Spectral ranges of HST WFC3 are shown for comparison, as well as the following JWST instruments: the Near Infrared Spectrograph (NIRSpec), the Near Infrared Imager and Slitless Spectrograph (NIRISS), the Near Infrared Camera (NIRCам) and the Mid Infrared Instrument (MIRI).

the optical depth for $P \geq P_{\text{cld}}$ is infinity. For $P < P_{\text{cld}}$ the attenuation coefficient is given by

$$\mu_{\lambda, \text{haze}}(\lambda, P, T) = X_{\text{H}_2} n_{\text{tot}}(P, T) \sigma_{\lambda, \text{haze}}, \quad (2.38)$$

$$\sigma_{\lambda, \text{haze}} = a \sigma_0 \left(\frac{\lambda}{\lambda_0} \right)^\gamma, \quad (2.39)$$

where a is the Rayleigh enhancement factor, γ is the scattering slope, and $\sigma_0 = 5.31 \times 10^{-31} \text{ m}^2$ is the cross-section due to H_2 Rayleigh scattering at a reference wavelength $\lambda_0 = 3.5 \times 10^{-7} \text{ m}$ (Dalgarno & Williams, 1962).

This prescription allows for partial cloud cover as discussed in Line & Parmentier (2016) by introducing another free parameter $\bar{\phi}$ that determines the fraction of the atmosphere which is covered by clouds. This alters the expression for transit depth so that we obtain a linear

combination of the transit depths computed both with and without clouds:

$$\Delta = \bar{\phi}\Delta_{\text{cloudy}} + (1 - \bar{\phi})\Delta_{\text{clear}}. \quad (2.40)$$

The addition of clouds to the atmospheric forward model therefore introduces four free parameters: P_{cld} , a , γ and $\bar{\phi}$.

2.1.4 Atmospheric Structure

In order to compute the optical depth at a given height in the atmosphere, the number density of particles at that height must be known. Additionally, cross-sections of chemical species are dependent on pressure and temperature, which also varies with height. Therefore, the temperature, density and pressure of the atmosphere must be computed as a function of height in order to model a transmission spectrum.

Solving for the structure of an atmosphere requires the equation of hydrostatic equilibrium and an equation of state, in this case the ideal gas law (see Section 1.3.2.1), which can be combined to yield

$$\frac{dP}{dz} = \frac{\mu m_H g}{k_B T} P = \frac{P}{H_{\text{sc}}}, \quad (2.41)$$

where H_{sc} is the atmospheric scale height, as defined in Chapter 1. In order to solve equation 2.41 a reference pressure P_{ref} must be specified. This is typically given as the pressure at the planetary radius ($z = 0$).

Solving equation 2.41 also requires temperature to be known as a function of pressure. Forward models such as the self-consistent models described in Section 1.3.3.1 typically calculate a pressure–temperature (P – T) profile using the constraints of radiative-convective equilibrium. However, this is too computationally expensive for retrievals, where a very wide range of models need to be explored to find those which best explain an observed data set. Therefore, a number of flexible, analytic approaches to calculating P – T profiles have been developed. The simplest prescription is to assume an isothermal atmosphere, a method which has been adopted in a number of retrieval studies (e.g. Waldmann et al., 2015a, Zhang et al., 2019). Other retrieval frameworks (e.g. Benneke & Seager, 2012, Line et al., 2013) have implemented the analytic P – T introduced by Guillot (2010), which is a solution to the equations of radiative transfer under certain assumptions. This takes the form

$$T^4 = \frac{3T_{\text{int}}^4}{4} \left[\frac{2}{3} + \tau \right] + \frac{3T_{\text{irr}}^4}{4} f \left[\frac{2}{3} + \frac{1}{\gamma_r \sqrt{3}} + \left(\frac{\gamma_r}{\sqrt{3}} - \frac{1}{\gamma_r \sqrt{3}} \right) \exp(-\gamma_r \tau \sqrt{3}) \right] \quad (2.42)$$

where T_{irr} is the irradiation temperature, arising due to the irradiation intensity on the planet from the host star, T_{int} is the intrinsic temperature which characterises the planet's internal heat flux, and γ_r is the ratio of visible to thermal opacities. In this case, the temperature is linked to pressure through the optical depth τ . Another commonly implemented P - T profile (e.g. Pinhas et al., 2018, Blecic et al., 2021) is the parametric prescription developed by Madhusudhan & Seager (2009), which takes the form

$$\begin{aligned} P &= P_0 e^{\alpha_1 \sqrt{T-T_0}}, & P_0 < P < P_1, \\ P &= P_2 e^{\alpha_2 \sqrt{T-T_2}}, & P_1 < P < P_3, \\ T &= T_3, & P > P_3, \end{aligned} \quad (2.43)$$

where P_0 is the pressure at the top of the atmosphere. This profile can also be expressed as:

$$T(P) = \begin{cases} T_0 + \left(\frac{\ln(P/P_0)}{\alpha_1} \right)^2, & P_0 < P < P_1, \\ T_2 + \left(\frac{\ln(P/P_2)}{\alpha_2} \right)^2, & P_1 < P < P_3, \\ T_2 + \left(\frac{\ln(P_3/P_2)}{\alpha_2} \right)^2, & P > P_3. \end{cases} \quad (2.44)$$

As well as P_0 , this P - T profile requires nine further parameters to be specified: α_1 , α_2 , T_0 , T_1 , T_2 , T_3 , P_1 , P_2 and P_3 . However, the requirement that that the resulting profile should be continuous allows three of these to be eliminated, leaving a total of six free parameters which define the shape of the P - T profile. It has been demonstrated that this prescription is capable of matching a wide range of temperature profiles from the solar system and self-consistent models, and can describe atmospheres both with and without thermal inversions.

Further discussion of P - T profiles for atmospheric retrieval, including profiles that vary with longitude and latitude in the planet as well as height, is presented in Chapter 4.

2.2 Parameter Estimation

Alongside an atmospheric forward model, the other key component of a retrieval framework is a parameter estimation scheme, which explores the space of combinations of model parameters in order to find which models (and therefore which parameters) best describe the data at hand. A wide variety of methods have been used in different retrieval algorithms, with the most prominent approaches summarised here.

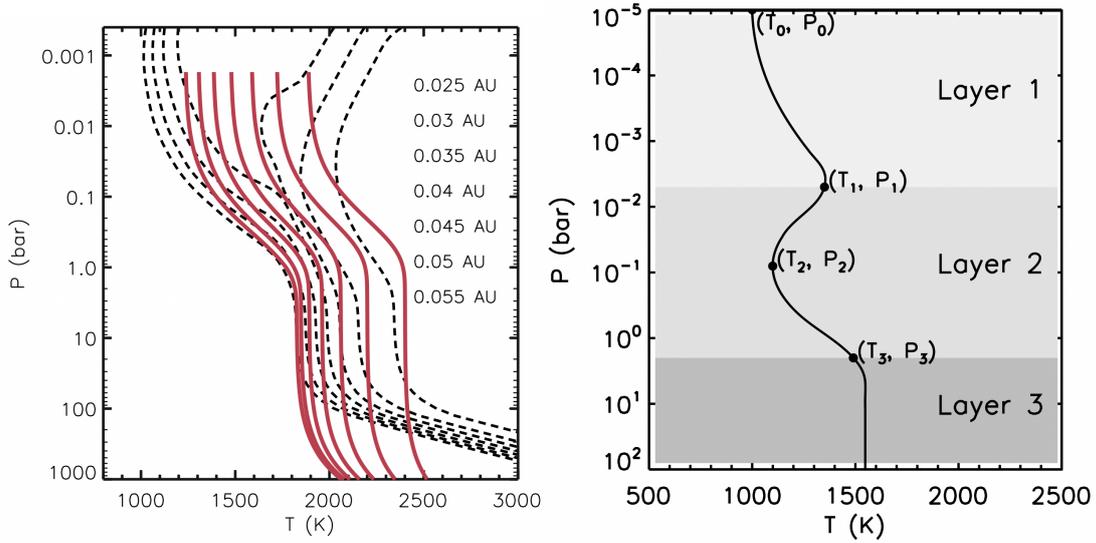


Fig. 2.4 Pressure–temperature (P – T) profiles commonly employed in retrievals. *Left*: Analytic P – T profiles from Guillot (2010), calculated for a theoretical hot Jupiter at a range of orbital distances, shown as red lines. Black dashed lines show self-consistently modelled P – T profiles for the same theoretical planet from Fortney et al. (2008). *Right*: Parametric P – T profile from Madhusudhan & Seager (2009). The profile includes a thermal inversion layer (Layer 2) in cases where $P_2 > P_1$. The profile is isothermal for $P > P_3$ (Layer 3).

2.2.1 Grid-based Retrievals

The earliest atmospheric retrievals explored the parameter space using a large pre-computed grid of forward models. Madhusudhan & Seager (2009) produced a grid of 10^7 forward models for a 10-dimensional parameter space, with six free parameters describing the temperature profile and four parameters for molecular abundances (H_2O , CO , CH_4 and CO_2). The grid did not simply use evenly spaced values for each parameter, since acquiring a reasonable resolution in 10 dimensions would require many more models to be calculated ($\geq 10^{10}$). The full grid of models instead was created by producing a number of coarse grids in order to gain some understanding of the parameter space, before producing a more detailed set of models where a better fit to the data could be expected. For each model in the grid, the goodness of fit to the data was computed through the following statistic:

$$\chi^2 = \frac{1}{N_{\text{obs}}} \sum_{i=1}^{N_{\text{obs}}} \frac{(\hat{y}_i - \bar{y}_i)^2}{\sigma_i^2}, \quad (2.45)$$

where \bar{y}_i and σ_i are the mean and standard deviation of the observed data point i , \hat{y}_i is the value of the model data point i , and N_{obs} is the number of data points in the spectrum. Models with the lowest values of ξ^2 provide a better fit to the data, allowing for statistical constraints to be placed on atmospheric parameters given the data at hand.

While grid-based retrievals were revolutionary when first implemented, they are far from optimised for searching high-dimensional parameter spaces, and would struggle to achieve the computational efficiency required to incorporate more complex models that may need additional free parameters to describe the presence of extra molecules or phenomena such as clouds and hazes. It has therefore been necessary to adopt more formal parameter estimation methods, as described below.

2.2.2 Bayesian Inference

Bayesian inference is an application of Bayes' theorem (Bayes, 1763) to determine the probability distribution of a set of model parameters θ given some data d and a model \mathcal{M} :

$$p(\theta|d, \mathcal{M}) = \frac{p(d|\theta, \mathcal{M})p(\theta|\mathcal{M})}{p(d|\mathcal{M})}, \quad (2.46)$$

where $p(\theta|d, \mathcal{M})$ is the probability of the model parameters given the data, known as the *posterior distribution*, $p(d|\theta, \mathcal{M})$ is the probability of the data given the model parameters, known as the *likelihood function*, $p(\theta|\mathcal{M})$ is the *prior* probability distribution of the model parameters independent of the data, and $p(d|\mathcal{M})$ is the probability of the data based on the model independent of the parameters, known as the Bayesian *evidence*. Note that each of the other quantities are also conditioned on the choice of model. Often the likelihood function is denoted by \mathcal{L} and the evidence is denoted by \mathcal{Z} .

In order to compute the posterior, the prior and likelihood are therefore required. Prior distributions for each parameter are specified in advance, and are typically chosen so as to be as uninformative as possible, using uniform or log-uniform priors covering a wide range of possible values of each parameter. The likelihood function is then evaluated by sampling the prior space and evaluating how well a model generated with a given set of parameters fits the observed data. For data with independently distributed Gaussian errors (which is commonly assumed for atmospheric retrieval), the likelihood function is defined as

$$\mathcal{L} = \prod_i^{N_{\text{obs}}} \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left(-\frac{(\hat{y}_i - \bar{y}_i)^2}{2\sigma_i^2}\right). \quad (2.47)$$

The Bayesian evidence \mathcal{Z} is a normalising factor in equation 2.46 and is not necessary when carrying out parameter estimation, since it is independent of the model parameters. On its own, an evidence is not particularly interpretable, since it is just a single number. However, by comparing the evidences obtained through parameter estimation using different models, the relative validity of the models for analysing the data at hand can be assessed, with a larger evidence indicating that a model is more appropriate for the data set. It is worth noting however that the evidence also depends strongly on the choice of priors, which could vary between different model considerations, and so model comparison using Bayesian evidences should be viewed with some caution (see e.g. Schad et al., 2022, for a more detailed discussion).

A number of algorithms exist that can be used to sample the model parameter space for the purposes of Bayesian inference. The key methods which have been applied to atmospheric retrieval are summarised below.

2.2.2.1 Optimal Estimation

The Optimal Estimation algorithm optimises the likelihood function using a non-linear least squares minimisation routine (see e.g. Rodgers, 2000). This approach allows priors to be specified assuming a Gaussian-distributed prior covariance matrix, which is useful for inverse modelling of Earth and solar system atmospheres, where the prior can be strongly informed by direct measurements (Irwin et al., 2008). However, the low resolution and low signal-to-noise common in exoplanetary data is not typically compatible with Optimal Estimation methods (Line et al., 2013). Furthermore, this approach is less suited to large, multimodal parameter spaces with strong degeneracies, as is often the case for atmospheric retrievals of exoplanets, and also assumes that the posterior takes the form of a multimodal Gaussian distribution, which may not be accurate. Despite these disadvantages, there are a number of studies which have successfully implemented Optimal Estimation techniques for atmospheric retrieval (e.g. Lee et al., 2012, Barstow et al., 2017).

2.2.2.2 Markov chain Monte Carlo

Markov chain Monte Carlo (MCMC) is a widely-used sampling algorithm in many areas of astronomy and other scientific fields (e.g. Everall et al., 2019, Kasim et al., 2019, Valderrama-Bahamóndez & Fröhlich, 2019). Unlike Optimal Estimation, MCMC does not make any assumptions about the form of the posterior distribution, but instead explores and evaluates the posterior by sampling the parameter space. The sampling uses a random walk procedure, starting from a random location in the parameter space. A new state of parameter values is

drawn from some proposal distribution, and the likelihoods of the current and new states are compared to one another. If the new state has a higher likelihood than the current one, then the new state replaces the current state, and if the new state does not have an improved likelihood then it is either rejected or accepted with some user-defined probability. This process constitutes a Markov chain since the decision to change states depends only on the current state and the new state being considered. The specific proposal distribution used to find new states varies depending on user preference as well as the specific MCMC algorithm that is implemented, with common examples being the Metropolis-Hastings algorithm or Gibbs sampling (see e.g. Trotta, 2017, for a more detailed review). Given enough samples, this approach will converge to the posterior distribution.

Since MCMC is capable of efficiently computing non-Gaussian posterior distributions, it is well-suited to atmospheric retrieval and has been implemented in numerous retrieval frameworks (e.g. Madhusudhan et al., 2011, Line et al., 2013, Wakeford et al., 2017). The approach does have some drawbacks; for example, properties of the proposal distribution which must be defined before running the algorithm are often difficult to determine in advance and can lead to convergence issues, and the algorithm can struggle to compute multimodal posteriors (Hogg & Foreman-Mackey, 2018). Furthermore, MCMC is not optimised for calculating the Bayesian evidence, which is not important for parameter estimation but can be useful for model comparison.

2.2.2.3 Nested Sampling

Nested Sampling (Skilling, 2004, 2006) differs from the other algorithms described above in that its main goal is to calculate the Bayesian evidence for a model, with the posterior distributions of model parameters being computed as a by-product. Nested Sampling works by selecting a number of random locations drawn from the prior, known as *live points*. For each live point, the likelihood is calculated, and the points are placed in ascending likelihood order. The point with the smallest likelihood is discarded, and replaced with a new point whose likelihood must be higher than that of the discarded point. This process is repeated, with the set of live points occupying an ever-shrinking region of the parameter space with increasing likelihoods. The evidence is given by a weighted sum of the likelihoods of the sampled points, which will converge after a sufficient number of iterations. Since the likelihood of each point must be calculated in order to compute the evidence, it is trivial to also obtain the posterior probability distribution. A recent review of the algorithm can be found in Ashton et al. (2022).

The method for selecting new live points varies between different implementations of the Nested Sampling algorithm. One of the most commonly used implementations for

atmospheric retrieval is MultiNest (Feroz et al., 2009, Feroz & Skilling, 2013), which samples the parameter space using an ellipsoidal rejection sampling scheme. In this approach, the next live point is found from within a set of ellipsoids defined by the iso-likelihood contour of the minimum likelihood. These ellipsoids shrink as the minimum likelihood increases with each iteration. MultiNest has been shown to be effective for models with moderate numbers of free parameters ($\lesssim 80$, Handley et al., 2015) and for evaluating multimodal posterior distributions (Feroz & Skilling, 2013), making it well-suited to atmospheric retrieval problems and leading to its incorporation into a large number of retrieval frameworks (e.g. Benneke & Seager, 2013, Line et al., 2015, Waldmann et al., 2015a, Lavie et al., 2017, Pinhas et al., 2018). Other variations of Nested Sampling have also been used for retrievals, such as PolyChord (Handley et al., 2015), which has been shown to give comparable retrieval results to MultiNest (Welbanks & Madhusudhan, 2021). PolyChord uses a different sampling method that is more efficient for models with $\gtrsim 80$ free parameters (Handley et al., 2015), which is not required for contemporary atmospheric retrievals but may become necessary in the future as increasingly complex models are required to explain more detailed data sets.

2.2.3 Machine Learning

The past decade has witnessed a huge increase in the application of machine learning techniques in many areas of science (Carleo et al., 2019). The goal of a machine learning algorithm is to automatically perform a task involving inference or prediction using a large amount of data. These algorithms can be broadly grouped into two varieties:

- Supervised learning algorithms, in which the algorithm learns from a data set that has been labelled in some way, and then tries to predict the labels for some new, unlabelled data. Examples of these algorithms include Support Vector Machines, most Neural Networks and Random Forests.
- Unsupervised learning algorithms, in which the algorithm tries to find patterns in unlabelled data, and then searches for similar patterns in new data. Examples include anomaly detection and clustering algorithms.

Since most existing applications of machine learning to the problem of atmospheric retrieval focus on supervised learning, a more formal description of a general supervised learning problem is presented: we define a set X of samples \mathbf{x} and another set Y of labels \mathbf{y} . Each label $\mathbf{y} \in Y$ corresponds to a specific sample $\mathbf{x} \in X$. The samples are typically vectors whose components are called features. The labels can take either categorical or numerical forms. X and Y together make up the training data set for the problem.

The task of the supervised machine learning algorithm is to find the relationship between the input data X and outputs Y . If we assume that in reality, there is some function f satisfying $f(\mathbf{x}) = \mathbf{y}$, then the algorithm should produce a function \hat{f} such that $\hat{\mathbf{y}} = \hat{f}(\mathbf{x})$ is as close to the true value of \mathbf{y} as possible. The goodness-of-fit between $\hat{\mathbf{y}}$ and $\hat{f}(\mathbf{x})$ can be evaluated using loss functions such as the statistic shown in equation 2.45. Once such a function has been learned, the algorithm can be used to predict the label of a new sample not in the training data set for which there is no previously known label.

Machine learning techniques are potentially well-suited to the problem of atmospheric retrieval, which involves determining the relationship between input data in the form of a spectrum and outputs in the form of model parameters. It is possible that new approaches to atmospheric retrieval using machine learning could be more computationally efficient, and that they could alleviate some of the problems associated with the parameter estimation methods presented previously. For example, Bayesian sampling algorithms tend to become exponentially slower as the dimension of the parameter space increases, a phenomenon known as the *curse of dimensionality* (see e.g. Robert & Casella, 2004). This could be less of an issue if a machine learning algorithm is used, allowing for more efficient scaling to higher dimensions. A selection of machine learning algorithms applied to atmospheric modelling and retrieval are introduced below.

2.2.3.1 Deep Belief Networks

Deep Belief Networks (DBNs, Hinton et al., 2006) are a class of Neural Network composed of several unsupervised components known as Restricted Boltzmann Machines (Smolensky, 1986) and a final supervised regression or classification component. Neural Networks are a subset of machine learning algorithms inspired by the biological networks in the brain. A Restricted Boltzmann Machine (RBM) is a simple, two-layer network that is able to learn probability distributions over inputs. It can be visualised as a bipartite graph consisting of a vector of visible inputs \mathbf{x} and a vector of hidden units \mathbf{y} . All of the inputs are connected to all of the hidden units, but there are no connections between units in the same layer. This means that units in the same layer can be considered as independent and so we can write the probability of the input units given the hidden units, and the hidden units given the input units, as

$$P(\mathbf{x}|\mathbf{y}) = \prod_i P(x_i|\mathbf{y}), \quad (2.48)$$

$$P(\mathbf{y}|\mathbf{x}) = \prod_j P(y_j|\mathbf{x}). \quad (2.49)$$

The input and hidden layers are connected by defining a matrix \mathcal{W} of weights, where element \mathcal{W}_{ij} connects input element x_i to hidden element y_j . The input and hidden layers are assigned bias vectors \mathbf{a} and \mathbf{b} respectively. The hidden nodes are always binary variables, and the input nodes can either be binary or continuous. Assuming the input variables are binary, the probability distribution of each hidden node given the input vector is

$$P(y_j = 1|\mathbf{x}) = \varsigma\left(b_j + \sum_i \mathcal{W}_{ij}x_i\right), \quad (2.50)$$

where ς is the sigmoid function:

$$\varsigma(x) = \frac{1}{1 + e^{-x}}. \quad (2.51)$$

The probability distribution of an input node given a hidden node is

$$P(x_i = 1|\mathbf{y}) = \varsigma\left(a_i + \sum_j \mathcal{W}_{ij}y_j\right). \quad (2.52)$$

We now define the *energy* of a given pair of input and hidden vectors to be

$$E(\mathbf{x}, \mathbf{y}) = -\mathbf{a} \cdot \mathbf{x} - \mathbf{b} \cdot \mathbf{y} - \mathbf{y}^T \mathcal{W} \mathbf{x}, \quad (2.53)$$

and we define a *partition function*

$$Z = \sum_{\mathbf{x}, \mathbf{y}} e^{-E(\mathbf{x}, \mathbf{y})}. \quad (2.54)$$

This allows us to define the joint probability distribution over input and hidden vectors:

$$P(\mathbf{x}, \mathbf{y}) = \frac{1}{Z} e^{-E(\mathbf{x}, \mathbf{y})}. \quad (2.55)$$

Finally, the distribution of the inputs is given by summing over all possible vectors of hidden units:

$$P(\mathbf{x}) = \frac{1}{Z} \sum_{\mathbf{y}} e^{-E(\mathbf{x}, \mathbf{y})}. \quad (2.56)$$

If the input variables are instead continuous we instead have

$$P(y_j = 1|\mathbf{x}) = \varsigma\left(b_j + \sum_i \mathcal{W}_{ij} \frac{x_i}{\sigma_i}\right), \quad (2.57)$$

$$P(x_i|\mathbf{y}) = \mathcal{N}\left(x_i; a_i + \sum_j \mathcal{W}_{ij} y_j, \sigma_i^2\right), \quad (2.58)$$

$$E(\mathbf{x}, \mathbf{y}) = \sum_i \frac{(x_i - a_i)^2}{2\sigma_i^2} - \sum_j b_j y_j - \sum_{i,j} \frac{x_i}{\sigma_i} \mathcal{W}_{ij} y_j, \quad (2.59)$$

where \mathcal{N} represents a normal distribution, and σ_i is the (assumed Gaussian) noise on data point x_i . The RBM is trained by adjusting the weights and biases in order to maximise the log-likelihood $P(\mathbf{x}|\mathcal{W}, \mathbf{a}, \mathbf{b})$ using gradient descent (see e.g. Roux & Bengio, 2008, for further detail).

DBNs were one of the first machine learning algorithms to be applied in the field of exoplanet atmospheres (Waldmann, 2016). The trained network reads in a spectrum and outputs the probabilities that one of several different molecules is the dominant absorbing species in the spectrum. This could provide a useful pre-selection tool, informing which molecules to include in the forward model. However, the network struggles to deal with multiple opacity sources in a single spectrum. The algorithm is sometimes able to identify mixtures of chemical species, but only if the spectral signatures of those molecules are very different and the abundances of the different species are within about an order of magnitude of each other. If the algorithm could be generalised to higher numbers of molecules successfully, this could help to reduce the computational cost of a traditional retrieval by informing which free parameters to include (and which not to include), keeping the dimension of the parameter space to a minimum.

2.2.3.2 Random Forest

Random Forest (Breiman, 2001) is an extension of the simpler decision tree algorithm (Breiman et al., 1984). Decision trees, and hence Random Forest, can be used for either classification or regression tasks and has been widely applied in astronomy as well as other areas of science (e.g. Richards et al., 2011, Xiao et al., 2017). Random Forest was the first machine learning algorithm to be used for a complete atmospheric retrieval (Márquez-Neila et al., 2018), and has subsequently been applied in a number of studies (Fisher et al., 2020, Guzmán-Mesa et al., 2020), however the approach does have some limitations. A detailed description of the algorithm, its applications to atmospheric retrieval, and an assessment of its efficacy in this domain is presented in Chapter 3.

2.2.3.3 Generative Adversarial Networks

Generative Adversarial Networks (GANs) are a type of unsupervised generative neural network invented by Goodfellow et al. (2014). Similarly to DBNs, GANs are able to learn probability distributions over a given data set. The training process involves the use of two networks, called a *generator* and a *discriminator*. The generator aims to produce new data such that the discriminator cannot determine whether this data belongs to the original training set. Once the discriminator can no longer distinguish real data from that produced by the generator, the network has learned a close approximation to the probability distribution of the training data set. Now the generator can be used to produce new data that is similar to that found in the training set, or fill in gaps in incomplete data samples, in a technique known as *inpainting*.

Zingales & Waldmann (2018) use inpainting to predict the atmospheric properties of a spectrum. In order to do this they arrange their training data samples in a two-dimensional array containing the normalised spectral data, information about the normalisation constants used, and the values of the model parameters used to generate the spectrum. They use seven free parameters in their forward model: the mass, temperature and radius of the planet as well as mixing ratios of H₂O, CO, CO₂ and CH₄. Given only the spectral data and normalisation information, the generator attempts to reproduce the rest of the array, thus estimating the free parameters in the model.

Zingales & Waldmann (2018) perform retrievals using both the GAN and Nested Sampling, and find that the GAN predicts parameter values within the 1σ error bars of the Nested Sampling predictions. Their Nested Sampling analysis takes approximately 10 hours on 24 CPU cores. By contrast, the trained GAN takes only 2 minutes to predict the model parameters. However, the training phase of the GAN takes approximately three days per epoch on 20 CPU cores or about 9 hours per epoch on a GPU. The authors do not report how many epochs of training were required to fully train the network, or how long it took to generate the training data set of 10^7 forward models required for the training. It may be possible to produce a training data set that is applicable to a range of planetary parameters, so that a single trained GAN can estimate parameters for several different spectra, however the training set would need to cover a very large parameter space in order to include all possible chemical species, cloud models and so on that might be necessary to explain the spectra of a range of different planets.

Another issue with the GAN method of retrieval is that it is not able to accurately capture the uncertainties of its estimated parameter values. The prediction step involves repeating the inpainting process 1000 times and collecting a distribution of the results. This is not equivalent to producing a posterior distribution through MCMC or Nested Sampling, and

while the predicted median parameter values compare well between the Nested Sampling and GAN retrievals, the shapes of the distributions often differ considerably (e.g. Zingales & Waldmann, 2018, Figure 11).

2.2.3.4 Deep Neural Networks and Monte Carlo Dropout

The class of neural networks employed in Soboczenski et al. (2018) and Cobb et al. (2019) are similar in structure to the DBNs discussed previously. However, in this case the network is used to carry out parameter estimation, rather than suggesting which chemical species to include in a traditional retrieval. The network is provided with a training set of spectra labelled with their corresponding model parameters, and the training involves learning the relationship between the spectra and the parameters. This is a multivariate regression task, and here the neural network can be considered as a large combination of linear regressors combined with some non-linear functions that enable the network to learn much more complex relationships than an individual linear regressor.

Additionally, the two studies employ Monte Carlo dropout (Gal & Ghahramani, 2016) in order to approximate the uncertainties that would be found using traditional retrieval methods. This involves setting a certain proportion of the weights to zero during the prediction stage and making multiple predictions so that a distribution is generated for each parameter value. This method is an improvement over previous attempts to capture the shape of the true posterior distribution, but is found to underestimate the uncertainty in many cases. This is a known limitation of neural networks that attempt to mimic Bayesian inference methods (Blei et al., 2017).

2.3 Retrieval of a Hot Jupiter Transmission Spectrum

To close this chapter, we demonstrate a retrieval of the HST transmission spectrum of the canonical hot Jupiter HD 209458b using the AURA retrieval code (Pinhas et al., 2018). HD 209458b has been observed at optical wavelengths using STIS (Sing et al., 2008, 2016) as well as at IR wavelengths using WFC3 (Deming et al., 2013), enabling extensive study of its atmosphere (e.g. Madhusudhan et al., 2014b, Barstow et al., 2017, MacDonald & Madhusudhan, 2017, Welbanks & Madhusudhan, 2019) and making it an ideal test case to demonstrate the capabilities of atmospheric retrieval.

We retrieve the combined HST STIS+WFC3 transmission spectrum of the planet, covering a wavelength range of 0.3–1.7 μm . Our model assumes an isothermal atmosphere with inhomogeneous cloud cover following the prescription of MacDonald & Madhusudhan (2017). We choose an isothermal atmosphere since no evidence has been found for a strong

Table 2.1 Description of priors for retrieval of the transmission spectrum of HD 209458b. The priors have the same form for all chemical abundances X_i .

Parameter	Lower Bound	Upper Bound	Prior
$T_{\text{iso}} / \text{K}$	800	2710	uniform
$P_{\text{ref}} / \text{bar}$	10^{-6}	10^2	log-uniform
X_i	10^{-12}	10^{-2}	log-uniform
a	10^{-4}	10^8	log-uniform
γ	-20	2	uniform
$P_{\text{cld}} / \text{bar}$	10^{-6}	10^2	log-uniform
$\bar{\phi}$	0	1	uniform

temperature gradient in the terminator in past retrieval studies. The mixing ratios of Na, K, H₂O and HCN are included as free parameters in the model. This choice of chemical species is motivated by Welbanks & Madhusudhan (2019), who find that absorption due to Na, K, H₂O contributes to the transmission spectrum. The retrieval configuration is similar to cases 8-12 of that paper. We include HCN in order to demonstrate how the retrieval deals with a free parameter that cannot be constrained by the data. The model therefore has ten free parameters: T_{iso} , P_{ref} , X_{Na} , X_{K} , $X_{\text{H}_2\text{O}}$, X_{HCN} , a , γ , P_{cld} and $\bar{\phi}$. The priors for each free parameter are shown in Table 2.1.

We sample the space of model parameters using MultiNest, specifically its Python implementation PyMultiNest (Buchner et al., 2014). Each forward model is initially evaluated at a higher resolution than that of the data, covering 2000 wavelength points between 0.2–1.8 μm . This spectrum is then convolved with the point spread functions of the relevant instruments and integrated over the instrument functions, so that the model is binned to the resolution of the observed data so that the likelihood can be computed. For this retrieval we use 2000 live points, which is found to give a good trade-off between parameter estimation accuracy and computational cost.

Figure 2.6 shows retrieved marginalised posterior distributions for each parameter. The retrieved isothermal temperature is consistent with the equilibrium temperature of the planet. Three chemical species are clearly detected: Na, K and H₂O. The detections of Na and K are driven by the optical (STIS) data, whereas H₂O is mostly driven by the IR (WFC3) data. However, it has been shown that the inclusion of data at optical wavelengths significantly improves constraints on the H₂O abundance (Pinhas et al., 2018). The retrieved H₂O abundance is consistent with previous studies (MacDonald & Madhusudhan, 2017, Welbanks

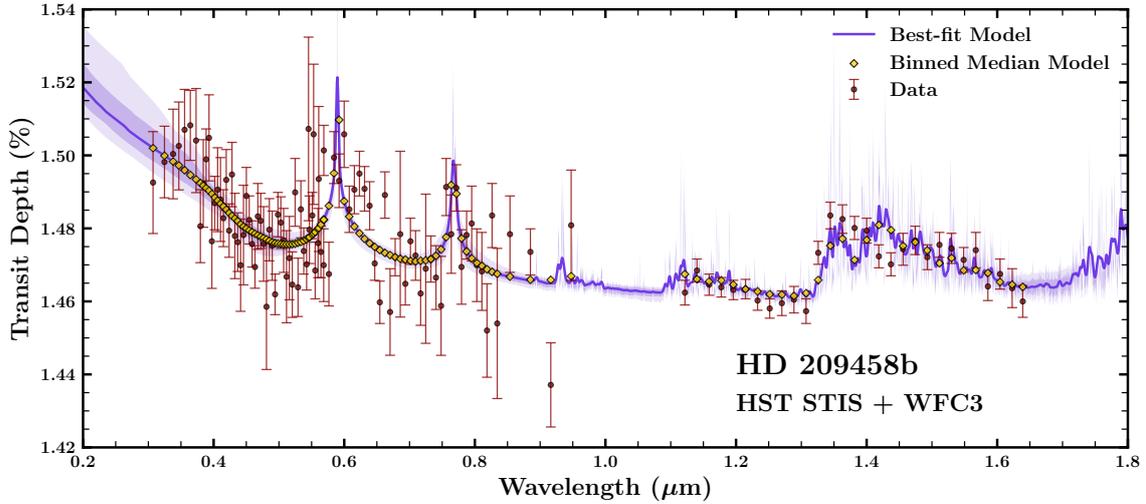


Fig. 2.5 Retrieval of the HST STIS+WFC3 transmission spectrum of HD 209458b. Observed data is indicated by red circles with error bars. The median retrieved spectrum is shown in dark purple, with lighter shaded regions indicating 1σ and 2σ confidence intervals. Yellow diamonds show the median retrieved model binned to the resolution of the data.

& Madhusudhan, 2019) and is lower than would be expected for an atmosphere with solar elemental abundances. This is consistent with a general trend that higher-mass planets tend to have lower atmospheric metallicity (Welbanks et al., 2019).

We infer a cloud coverage fraction of approximately 50%, with a high-altitude (~ 0.01 mbar) cloud deck on the cloudy fraction of the terminator. The retrieval results point to strongly enhanced Rayleigh scattering and a negative scattering slope to explain the steep increase in transit depth at shorter wavelengths. Overall, the results from this retrieval are an indication of what can be expected from state-of-the-art HST observations of a hot Jupiter that is very well-suited to atmospheric characterisation via transmission spectroscopy.

2.4 Summary and Discussion

In this chapter we have introduced the atmospheric retrieval framework, describing the theory of radiative transfer and transmission geometry required to compute model transmission spectra, as well as introducing a variety of parameter estimation methods that have been applied in past retrieval studies. We have presented an application of the AURA retrieval framework to state-of-the-art observations of the transmission spectrum of the hot Jupiter HD 209458b, demonstrating the capability of retrievals to constrain the abundances of numerous chemical species.

While the analysis shown here highlights the strengths of atmospheric retrieval, several limitations remain that may inhibit our ability to maximise the scientific output of upcoming observatories. Atmospheric retrieval is computationally intensive, and requires simplified atmospheric models in order to be feasible. As the quality of observations continues to increase, particularly with the advent of JWST (Greene et al., 2016), it is possible that more complex atmospheric models will be required, creating a computational bottleneck which could prevent us from learning about atmospheric processes that are left out of models, as well as creating the potential for biased inferences due to missing physics.

The next two chapters of this thesis represent efforts to address some of these issues. Chapter 3 is an assessment of the feasibility of using machine learning for the purpose of atmospheric retrieval. Machine learning has been proposed as an alternative to typical Bayesian retrievals in order to improve computational efficiency. While a number of machine learning-based retrieval algorithms have been developed with some success (e.g. Márquez-Neila et al., 2018, Zingales & Waldmann, 2018), they are still at times unable to accurately reproduce results from contemporary Bayesian retrievals. The goal of Chapter 3 is to determine whether it is possible to robustly reproduce the results of Bayesian retrieval algorithms with a machine learning approach.

Chapter 4 explores the need to develop more complex forward models for retrievals of next-generation observations. Atmospheric models of exoplanets used to explain observed data typically assume a one-dimensional thermally averaged structure. However, three-dimensional GCMs reveal much more complex temperature profiles that can vary with longitude and latitude across the terminator region of the atmosphere. When observing a planet's atmosphere in transmission, we are probing a region whose temperature profile may change significantly in the transition from day- to night-side across the terminator (e.g. Fortney et al., 2010). This is often not considered when creating model transmission spectra to interpret data, which can potentially lead to biases in atmospheric retrieval, as has been shown in previous studies (Caldas et al., 2019, Pluriel et al., 2020). The aim of Chapter 4 is to develop a retrieval framework capable of incorporating three-dimensional variations in a planet's temperature structure, and exploring whether such a framework will be sufficient to overcome some of the biases that have been found to be introduced when a one-dimensional atmosphere is assumed to interpret JWST-quality data.

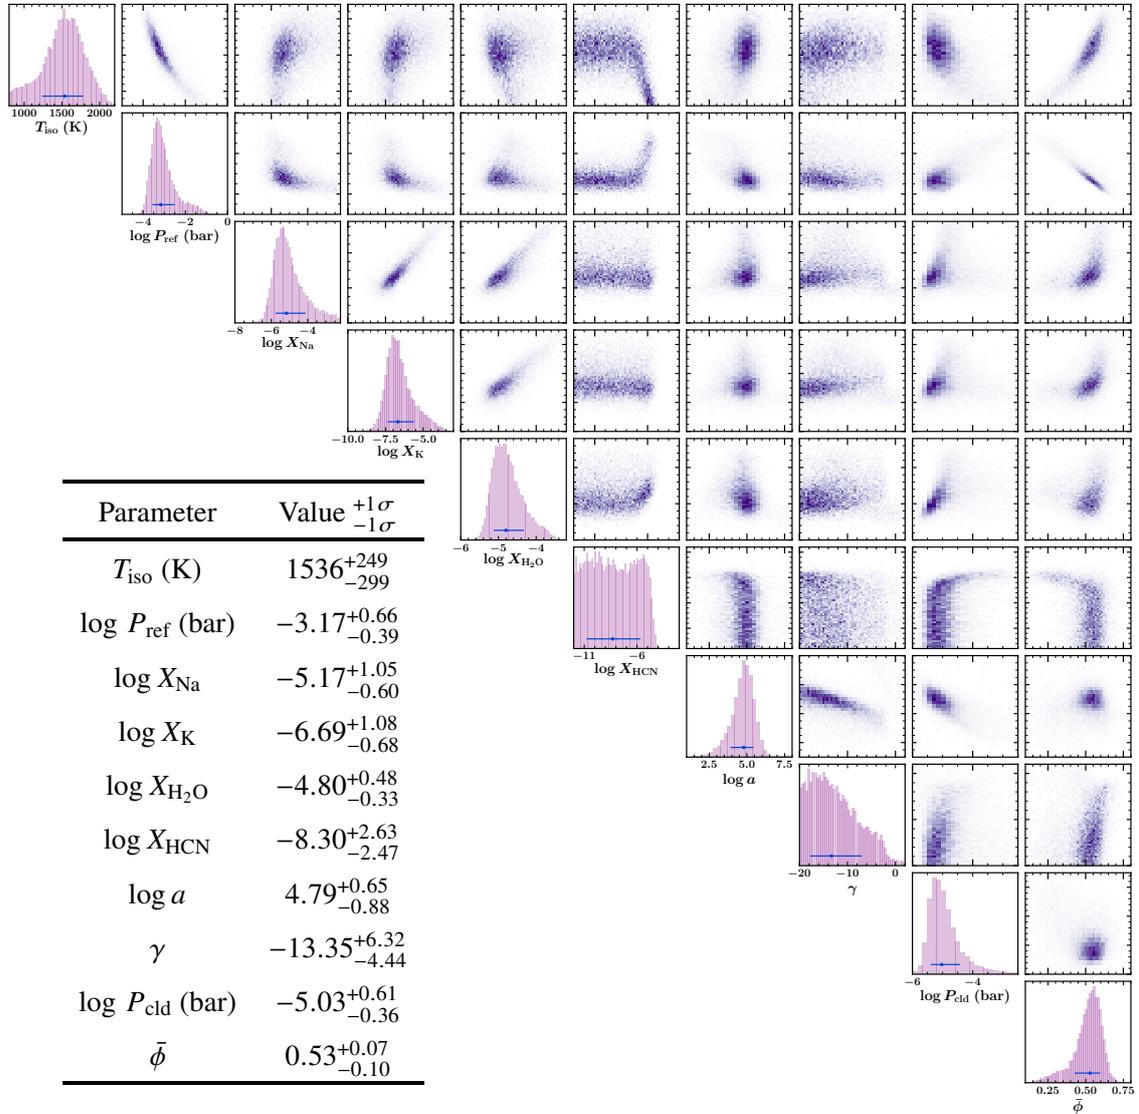


Fig. 2.6 Full posterior probability distribution from retrieval of the transmission spectrum of HD 209458b. The diagonal panels show histograms of the marginalised posteriors for each free parameter in the model, while the off-diagonal panels depict correlations between pairs of parameters. The blue points with lines indicate the median values and 1σ uncertainties. These results are consistent with Welbanks & Madhusudhan (2019), finding tight constraints on the abundances of Na, K and H_2O . The inset table shows median retrieved values and 1σ uncertainties for each parameter.

Chapter 3

Supervised Machine Learning for Atmospheric Retrieval of Exoplanets

“It’s fun to obey the machine!”

– Ralph Wiggum

3.1 Introduction

Machine learning and artificial intelligence are becoming increasingly prevalent in many areas of astrophysics. Many popular machine learning techniques have been applied to astrophysical problems including galaxy classification (Banerji et al., 2010), characterisation of supernovae (Lochner et al., 2016), and exoplanet detection (Shallue & Vanderburg, 2018). Recently a number of attempts have been made to use machine learning to retrieve properties of exoplanet atmospheres from spectroscopic data. Waldmann (2016) trained a Deep Belief Neural Network to make qualitative predictions about which molecular and atomic opacity sources to include in a traditional retrieval framework. Márquez-Neila et al. (2018) employed a supervised learning algorithm called Random Forest to retrieve atmospheric properties of the hot giant planet WASP-12b. Zingales & Waldmann (2018) developed a Generative Adversarial Network which uses unsupervised learning to predict planetary parameters as well as atomic and molecular abundances. Soboczenski et al. (2018) explored the use of Deep Neural Networks to make inference from synthetic spectra of terrestrial planets and incorporated Monte Carlo dropout in order to approximate model uncertainty. This method was further developed in Cobb et al. (2019), who used an ensemble of Neural Networks and incorporated domain-specific knowledge to improve performance.

A limitation of applying machine learning for retrievals has been the statistical interpretation of parameter predictions given the observed data. Traditionally, atmospheric retrieval has used Bayesian inference techniques to estimate the central values and uncertainties of the model parameters which fit an observed spectrum (Madhusudhan, 2018). Such techniques used for both transmission and emission spectra include MCMC (e.g. Madhusudhan & Seager, 2010, Line et al., 2013, Cubillos et al., 2013) and Nested Sampling (e.g. Benneke & Seager, 2013, Waldmann et al., 2015a, Oreshenko et al., 2017, Gandhi & Madhusudhan, 2018). When applied to atmospheric spectra, retrievals have often highlighted strong degeneracies between model parameters (e.g. Benneke & Seager, 2012, Griffith, 2014, Line & Parmentier, 2016, Welbanks & Madhusudhan, 2019). It is therefore important when carrying out a retrieval to use a method that is able to find these model degeneracies and accurately capture the inherent uncertainties in the observed spectra. Previous studies employing machine learning have produced either a set of predictions similar to running an ensemble of optimal estimation procedures (Zingales & Waldmann, 2018) or an approximation of the posterior distribution that is not shown to match the result of a Bayesian inference procedure (Cobb et al., 2019). In cases where attempts were made to compare a machine learning retrieval with a Bayesian retrieval (Zingales & Waldmann, 2018, Márquez-Neila et al., 2018), the posterior distributions between the retrievals reveal some discrepancies, as discussed later in this work.

In this paper we focus on supervised ensemble learning, similar to that employed by Márquez-Neila et al. (2018), referred to as MN18 hereafter. MN18 use the Random Forest algorithm, to train multiple estimators (or trees) to predict parameter values which best describe the transmission spectrum of WASP-12b. The distribution of predictions made by the estimators is used to find the uncertainties on the estimated parameter values. The results of their Random Forest retrieval are compared to a Nested Sampling retrieval, and whilst the two retrievals yield comparable parameter estimates, the uncertainties are not consistent between the two methods. Our goal in this study is to determine if it is possible to develop a more statistically sound retrieval framework using the Random Forest algorithm.

In Section 3.2 we first reproduce the retrieval results of MN18 (using both Nested Sampling and Random Forest), using the WASP-12b transmission spectrum. To this end, we use the same semi-analytic model used in MN18. We then modify and extend the Random Forest method to perform a retrieval of the same spectrum that produces results whose uncertainties are closer to those found in a Nested Sampling retrieval. In Section 3.3 we combine this extended Random Forest method with the fully numerical forward model described in Pinhas et al. (2018). We validate our algorithm using synthetic spectra before conducting a case study of the HST Wide-Field Camera 3 (WFC3) transmission spectrum of HD 209458b, once again comparing the results of Random Forest and Nested Sampling retrievals. In Section

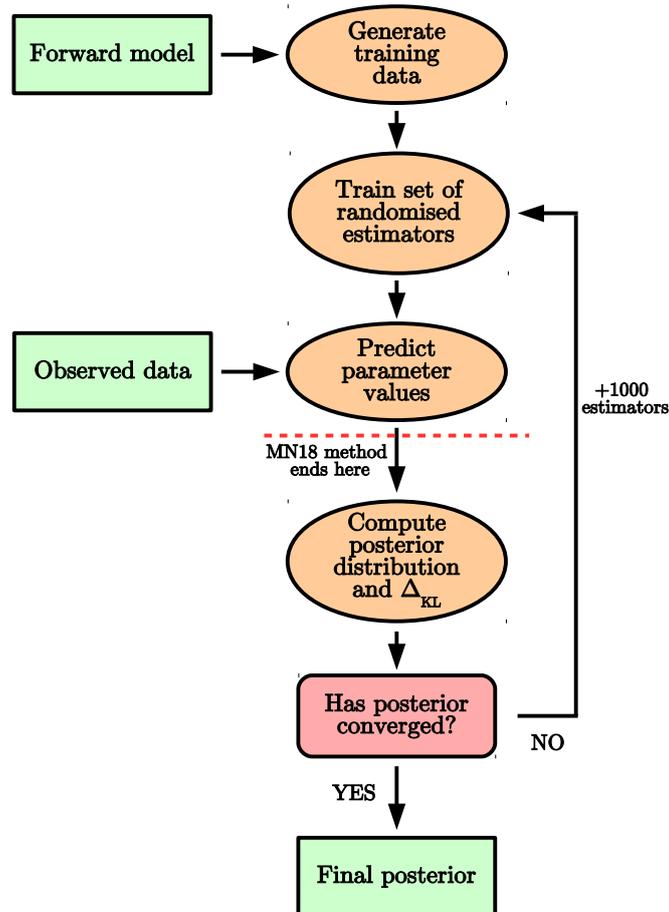


Fig. 3.1 Flowchart describing the extended Random Forest retrieval framework. As indicated, the key differences between this and the MN18 method are the calculation of the posterior distribution using the predicted parameter values and the iterative process of adding more trees until the posterior converges.

3.4 we discuss the difficulties of applying these methods to more complex cases which would require a larger parameter space to be explored than previous machine learning retrievals. We also examine more generally the possible benefits and drawbacks of incorporating machine learning, particularly ensemble learning as explored in this paper, into the retrieval process.

3.2 Methods

3.2.1 Reproduction of Previous Results

We begin by reproducing the results of MN18. We consider the same observed data as that paper, namely the WFC3 transmission spectrum of WASP-12b (Kreidberg et al., 2015).

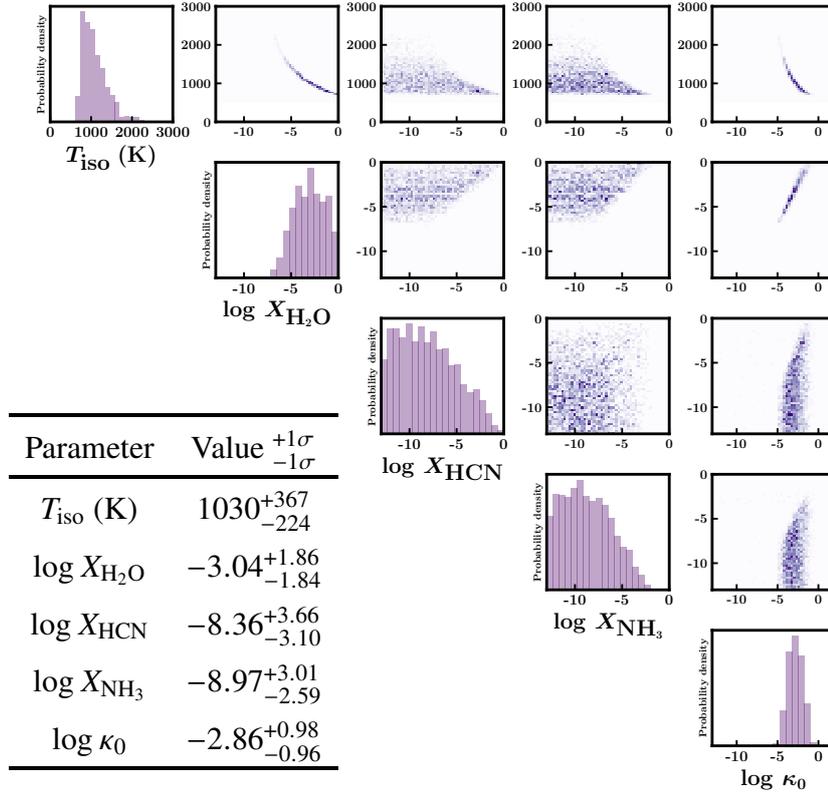


Fig. 3.2 Posterior distributions of Nested Sampling retrieval of the WFC3 transmission spectrum of WASP-12b following the methods of MN18. Inset: retrieved parameter values and associated 1σ uncertainties.

This spectrum consists of 13 binned data points in the infrared, at wavelengths ranging from $0.84 - 1.67 \mu\text{m}$. In order to retrieve atmospheric properties from the spectrum, two components are required: a forward model to calculate a transmission spectrum from a given set of parameters describing the atmospheric structure and composition, and a parameter estimation algorithm which finds the values of the model parameters that best fit the observed data. In this section we adopt the forward model of Heng & Kitzmann (2017) for consistency with the previous study. This semi-analytic model is used to produce a binned spectrum at the wavelengths of the WFC3 data given the values of five parameters: isotherm temperature T_{iso} , the abundances of H_2O , HCN and NH_3 , and a parameter to describe cloud opacity, κ_0 . For the parameter estimation we follow the two approaches considered in MN18, first using the Nested Sampling algorithm MultiNest (Feroz et al., 2009), specifically its Python implementation PyMultiNest (Buchner et al., 2014), and then using the implementation of the Random Forest algorithm (Breiman, 2001) from scikit-learn.

3.2.1.1 Nested Sampling

In a traditional retrieval, Bayesian inference is used to estimate the values of model parameters given some observed data. Suppose we want to find the probability distribution of a set of parameters, denoted θ , given some observed data, d . We can express this using Bayes' theorem:

$$p(\theta|d) = \frac{p(d|\theta)p(\theta)}{p(d)}. \quad (3.1)$$

Typically $p(\theta|d)$ is called the posterior, $p(d|\theta)$ is called the likelihood and is denoted \mathcal{L} , $p(\theta)$ is called the prior, and $p(d)$ is called the Bayesian evidence and is denoted \mathcal{Z} . Since \mathcal{Z} does not depend on θ , it simply acts as a normalisation factor and therefore is not needed for parameter estimation, however it can be used to compare different models.

Nested Sampling (Skilling, 2004) is a Monte Carlo algorithm designed to efficiently compute the Bayesian evidence of a model. It is also highly effective at sampling complex multimodal posterior distributions and is commonly used in many retrieval frameworks (e.g. Benneke & Seager, 2013, Gandhi & Madhusudhan, 2018). The algorithm initially selects a number of live points drawn from the defined prior volume, and evaluates the likelihoods of these points. Assuming Gaussian uncertainty on the measurements of the spectral data points, the likelihood is defined as

$$\mathcal{L} = \mathcal{L}_0 \exp\left(-\frac{\chi^2}{2}\right), \quad (3.2)$$

with

$$\chi^2 = \sum_i \frac{(\hat{y}_i - \bar{y}_i)^2}{\sigma_i^2}, \quad (3.3)$$

where \bar{y}_i and σ_i are the mean and standard deviation of the observed data point i , and \hat{y}_i is the model prediction for data point i .

Having calculated \mathcal{L} for each live point, the point with the lowest likelihood is discarded and replaced by a new one with a higher likelihood. This means that the volume contained within the set of live points continually shrinks, with the minimum likelihood bound by the volume progressively increasing. This process continues and \mathcal{Z} is calculated until converging to within some pre-defined tolerance. Since the evidence calculation requires a thorough sampling of the parameter space, the Nested Sampling algorithm can therefore be used to estimate posterior distributions.

Using MultiNest in conjunction with the forward model described in Heng & Kitzmann (2017), we reproduce the results from the Nested Sampling retrieval shown in MN18. The

Table 3.1 Description of priors for retrievals of the transmission spectrum of WASP-12b. The priors have the same form for all chemical abundances X_i .

Parameter	Lower Bound	Upper Bound	Prior
T_{iso} (K)	500	2900	uniform
X_i	10^{-13}	1	log-uniform
κ_0 ($\text{cm}^2 \text{g}^{-1}$)	10^{-13}	10^2	log-uniform

retrieved values and posterior distributions from this retrieval are shown in Figure 3.2. We obtain some constraints on the H_2O abundance while the HCN and NH_3 abundances remain unconstrained. The value of κ_0 is constrained to within 2 dex. The retrieved parameter values and associated uncertainties are consistent with the MN18 Nested Sampling retrieval (see table 1 of that paper).

3.2.1.2 Random Forest

Random Forest is a supervised machine learning algorithm. Supervised algorithms are trained on a data set that has been labelled in some way, and then try to predict the labels of some new, unlabelled data. The Random Forest method stems from the older decision tree algorithm (Breiman et al., 1984). Decision trees, and hence Random Forest, can be used for either classification or regression tasks; here we outline its application to a general regression problem, since this is how the algorithm is applied to atmospheric retrieval.

We define the feature space \mathcal{X} to be the vector space containing all possible input samples (binned spectra). The dimension of \mathcal{X} is equal to the number of features in a sample, i.e. the number of data points in a single spectrum \mathbf{x} . Similarly, we can define the space of all possible output labels \mathbf{y} (free parameters in the forward model) as \mathcal{Y} . In this context, the supervised machine learning problem becomes equivalent to finding the best possible partition of \mathcal{X} , where each partition corresponds to a different set of parameter estimates. The decision tree algorithm works by partitioning \mathcal{X} into subspaces and assigning different values from \mathcal{Y} to each subspace. In order to describe this further, we introduce some definitions from graph theory:

- A graph is a collection of nodes and edges, where a node can be connected to another node by an edge.

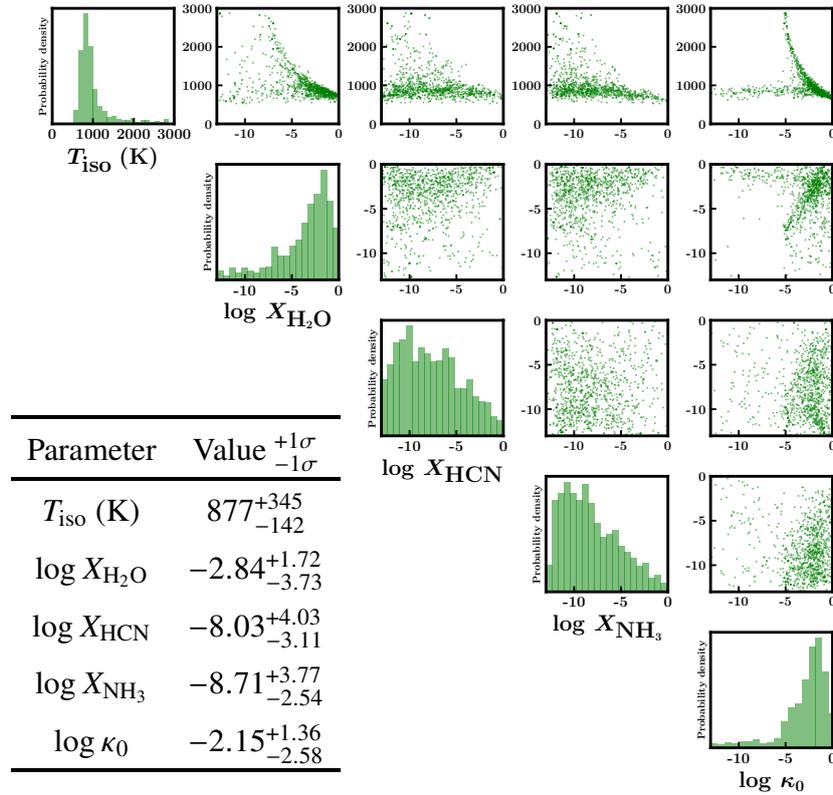


Fig. 3.3 Results of Random Forest retrieval of the WFC3 transmission spectrum of WASP-12b following the methods of MN18. While the parameter estimates are consistent to within 1σ with those in Figure 3.2, the posterior distributions have important differences in shape.

- A graph can be either undirected, meaning that if there is an edge from node a to node b then there is automatically an edge from node b to node a , or directed if this is not the case.
- If there is an edge from node a to node b but not from b to a , then a is said to be the parent of b and b the child of a .
- A tree is a graph in which there is exactly one path between any two nodes.
- If there exists a node in a tree where all edges are directed away from that node, then the node is called the root.
- A node in a tree which has no child nodes is called a leaf.

A decision tree can be defined as a directed tree in which any node n corresponds to some subspace \mathcal{X}_n of the feature space, with a root node that represents the entire space \mathcal{X} . Each leaf in the tree is assigned a value from the output space \mathcal{Y} . The aim of the learning process

is therefore to determine the tree structure which best captures the relationship between the spaces \mathcal{X} and \mathcal{Y} . We can quantify the concept of how well the model captures the relationship by defining the impurity I_n of a node n :

$$I_n = \frac{1}{N_n} \sum_{\mathbf{x} \in \mathcal{X}_n} |\mathbf{y}(\mathbf{x}) - \hat{\mathbf{y}}_n(\mathbf{x})|^2, \quad (3.4)$$

where N_n is the number of samples in the training data set which are found in the subspace \mathcal{X}_n , $\mathbf{y}(\mathbf{x})$ is the true value of the label corresponding to the sample \mathbf{x} , and $\hat{\mathbf{y}}_n(\mathbf{x})$ is the value of the label for \mathbf{x} currently predicted by the model. The impurity is similar to the χ^2 metric of equation 3.3. The algorithm proceeds by considering existing leaf nodes in the tree and splitting them into two or more child nodes (thus further partitioning the data set) such that the decrease in impurity from the parent node to the child nodes is maximised. This continues until some pre-determined tolerance in the impurity decrease is reached.

The Random Forest algorithm is an ensemble method which uses a large set of decision trees. Ensemble methods aim to improve the robustness of predictions by training multiple models that have been randomly perturbed in some way. The ensemble prediction is then a combination of the individual model predictions. Randomness can be introduced in two ways: by training each tree on a random subset of the full training data set, which is sampled with replacement, and by limiting each tree to train using a random subset of features. It can be shown that an ensemble of randomised decision trees produces a more robust prediction than using a single tree (see for example Chapter 4 of Louppe, 2014).

In order to reproduce the Random Forest results of MN18, we use the forward model of Heng & Kitzmann (2017) to generate a training set of 100000 noisy synthetic WFC3 spectra in the wavelength range $0.8 - 1.7 \mu\text{m}$. For each spectrum, the value of each free parameter is chosen at random from a uniform or log-uniform distribution from within the prior ranges specified in MN18 (see Table 3.1). The planetary and stellar radii are fixed at $R_p = 1.79R_J$ and $R_* = 1.577R_\odot$. We produce the training set at a higher wavelength resolution and larger wavelength range than the WFC3 spectrum, opting for $R = 2000$ between 0.2 and $2.0 \mu\text{m}$. This approach allows us to use the same training data set for multiple observation instances of the planet and would reduce the overall computation time of our method if other spectra of the same planet were to be analysed.

We train 1000 estimators on the training set with a minimum impurity decrease tolerance of 0.01. To begin the training phase, we bin each of the spectra in the training set to the resolution of the WFC3 spectrum, and add random Gaussian noise with a mean of 50 parts per million to each spectral data point. In order to improve the robustness of the predictions, each estimator is shown only 4 of the 13 spectral data points in each training sample. Figure

3.3 shows the distributions of the estimators' predicted parameter values for the WASP-12b spectrum, displaying a close match to figure 1 of MN18.

It is important to note that there are some discrepancies between the distributions shown in Figures 3.2 and 3.3. Most notably, the posterior distributions of H₂O abundance and of κ_0 have broad tails in the Random Forest retrieval which are not found in the Nested Sampling retrieval. These differences arise because the the distributions shown in Figure 3.3 are not true posterior distributions in the Bayesian sense; they are instead the relative densities of the predictions made by 1000 different estimators, some of which perform better than others by design. This means that this method does not necessarily capture the true shape of the posterior distributions and therefore cannot provide a robust estimate of the uncertainties of the predicted parameter values.

3.2.2 Extension of Random Forest Method

The differences between the shapes of the posterior distributions produced by the two different retrieval methods motivate the development of a new method, still employing machine learning in the form of the Random Forest algorithm, but yielding results that capture the uncertainty in parameter estimates more accurately. A diagram depicting this new approach is shown in Figure 3.1. We begin by producing a training data set in the same way as before, but we do not add noise to the model spectra. Before the training phase, we normalise the parameter values in the training data set so that they all lie between 0 and 1. This ensures that the loss function does not favour any one parameter over another.

Once the estimators have been trained on this noise-free data set and used to predict parameter values, the likelihoods of those predictions are calculated by comparing the observed spectrum to a forward model produced with the predicted parameter values (see equation 3.2). This set of predictions and associated likelihoods serves as the likelihood function for the retrieval, allowing the marginalised posterior for each parameter and pair of parameters to be computed. By calculating the likelihood of each prediction, the algorithm should no longer produce long tails that are not found in a Nested Sampling retrieval, since these predictions will have lower likelihoods and will be penalised accordingly. Since we impose a Gaussian likelihood, this method differs from other machine learning-based approaches to retrievals, which are typically likelihood-free.

The procedure initially trains a set of 1000 estimators to compute an initial posterior estimate. However, in order to ensure that enough estimators have been trained to sample the parameter space thoroughly, more estimators are added in batches of 1000 until the symmetric Kullback-Leibler divergence (Δ_{KL}) between successive posterior distributions falls below a certain tolerance. Δ_{KL} is defined as

$$\Delta_{\text{KL}} = \sum_x \left[p_i(\theta = x|d) \log \left(\frac{p_i(\theta = x|d)}{p_{i+1}(\theta = x|d)} \right) + p_{i+1}(\theta = x|d) \log \left(\frac{p_{i+1}(\theta = x|d)}{p_i(\theta = x|d)} \right) \right] \quad (3.5)$$

where $p_i(\theta = x|d)$ is the posterior probability that $\theta = x$ given by a forest of $(1000 \times i)$ trees. Figure 3.1 shows a flowchart describing our extended Random Forest retrieval method.

Figure 3.4 shows the posterior distribution from a retrieval of the same WASP-12b spectrum using the extended Random Forest method as described above. The retrieved parameter values are consistent within 1σ with those obtained in the previous two retrievals. However, the main difference between the results of this retrieval and those of the previous Random Forest retrieval is that the shape of the posterior distribution matches the Nested Sampling posterior more closely. The extended tails found in Figure 3.3 are no longer present. This is reflected in the reported uncertainties in the parameter estimates. Whereas the Random Forest retrieval following the approach from MN18 gives lower bounds far below those given by the Nested Sampling retrieval, the extended approach gives error bounds that are in line with the Nested Sampling result.

In order to obtain a good sampling of the parameter space for this problem using the extended Random Forest approach, a much larger number of estimators is required than the ensemble of 1000 used in the method outlined in MN18. Convergence is reached after 17000 estimators to produce the posterior distributions shown in Figure 3.4, with a mean tree depth of 39. A reasonable result can be produced using a higher tolerance which converges after approximately 10000 estimators have been generated. Simply increasing the number of trees without the likelihood evaluation step is not sufficient to obtain a more accurate retrieval; this is discussed in more detail in Section 3.3. While this increases the computational cost of the approach somewhat, the longest step in the retrieval is still the generation of the training data, since the forward model must be run 100000 times to create the full training set.

We investigated the effect of lowering the amount of training data to reduce computation time. This would reduce both the time taken for the training set to be produced and the training time itself, since the Random Forest algorithm trains more quickly on a smaller data set. Decreasing the amount of training data from 100000 spectra to 50000 yields resulting posterior distributions that are not well-sampled. We conclude that in this case a significant reduction in the amount of data used to train the estimators is not feasible.

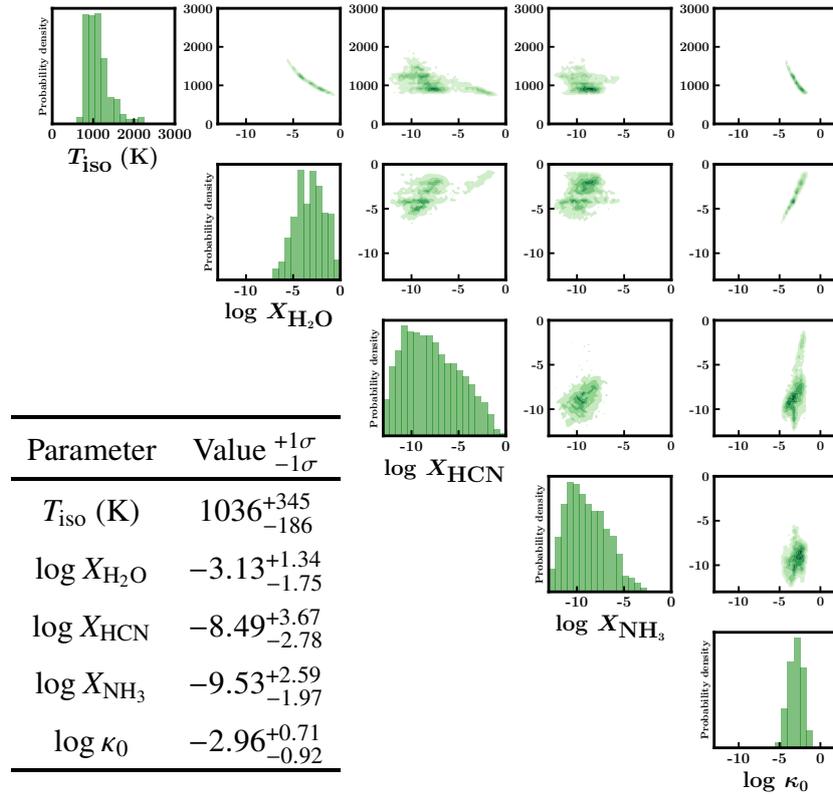


Fig. 3.4 Results of extended Random Forest retrieval of the WFC3 transmission spectrum of WASP-12b, using the same forward model as in MN18. The parameter estimates and posterior distributions provide a better match to the Nested Sampling retrieval shown in Figure 3.2, compared to Figure 3.3.

3.3 Applications

Having demonstrated that we can reproduce the results of MN18, and having extended their method to produce a result closer to that of a Nested Sampling retrieval, we now compare our new approach to a current state-of-the-art retrieval framework that uses a fully numerical forward model. We no longer use the forward model from Heng & Kitzmann (2017), instead adopting the modelling paradigm from AURA (Pinhas et al., 2018), a retrieval framework that has been validated against synthetic spectra and used to retrieve atmospheric properties, including H_2O abundances, for a range of planets (e.g. Pinhas et al., 2019, Welbanks et al., 2019).

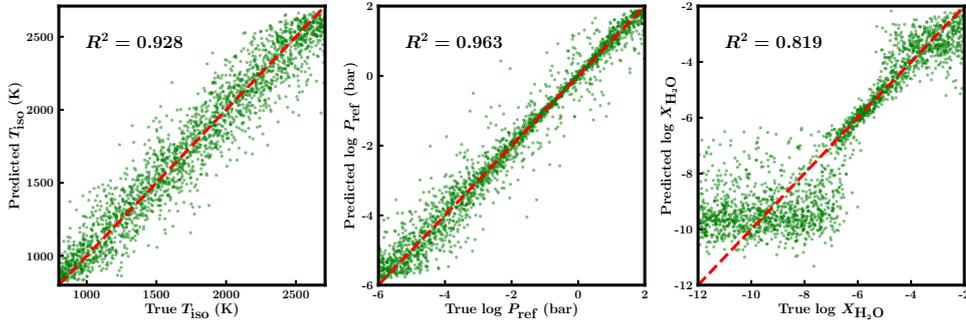


Fig. 3.5 True values versus Random Forest predictions for each parameter in the forward model, for a test set of 2000 synthetic models. The coefficient of determination (R^2) indicates the correlation between the true and predicted values, with R^2 close to 1 implying a strong correlation.

3.3.1 Validation

We begin by demonstrating our algorithm’s ability to accurately estimate parameter values from synthetic spectra. The AURA forward model is used to generate synthetic spectra in the wavelength range of WFC3 for training and testing our algorithm, assuming a cloud-free atmosphere and an isothermal temperature profile. Since WFC3 transmission spectra only provide nominal constraints on the temperature structure of the atmosphere (Barstow et al., 2013), assuming an isothermal temperature profile has little effect on retrievals with present data and is sufficient for the purposes of this study. Cloud properties are not considered since they are difficult to constrain without including data from optical wavelengths.

The model atmosphere is divided into 100 pressure layers, which are evenly log-spaced from $10^{-6} - 10^2$ bar. The main opacity sources in the model are H_2O and collision-induced absorption (CIA) due to $\text{H}_2\text{-H}_2$ and $\text{H}_2\text{-He}$. The cross-sections for these opacity sources are computed by Gandhi & Madhusudhan (2017) using line lists from the HITEMP database for H_2O (Rothman et al., 2010) and the HITRAN database for CIA (Richard et al., 2012). We fix R_p to $1.41R_J$ and leave the reference pressure P_{ref} as a free parameter to be retrieved. It was demonstrated by Welbanks & Madhusudhan (2019) that fixing one of the planetary radius and reference pressure and retrieving the other does not affect the retrieved values of the other parameters, and so we arbitrarily choose the retrieved radius of HD 209458b from Case 3 of that paper. The model therefore has 3 free parameters: T_{iso} , P_{ref} and the water abundance $X_{\text{H}_2\text{O}}$. This choice of parameterisation is appropriate for WFC3 spectra of hot Jupiters, which to date have been found to be most sensitive to planetary radius, temperature and H_2O abundance (Tsiaras et al., 2018, Welbanks & Madhusudhan, 2019). The prior ranges for each parameter are shown in Table 3.2.

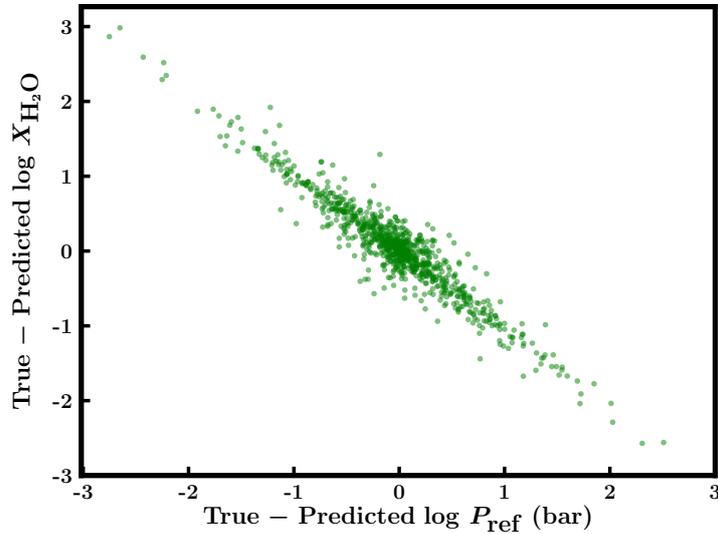


Fig. 3.6 Difference between true and predicted values of P_{ref} and $X_{\text{H}_2\text{O}}$ for synthetic models with $\log X_{\text{H}_2\text{O}} \geq -6$. We find that predictions of P_{ref} that are above the true value correspond to predictions of $X_{\text{H}_2\text{O}}$ below the true value and vice-versa.

We use this forward model to produce a training data set of 8000 model spectra and a validation data set of 2000 spectra. Each data set has randomly generated parameters, and no spectra from the validation set appear in the training set. The spectra are produced at a high resolution ($R = 1000$) in the wavelength range of WFC3 ($1.1 - 1.7 \mu\text{m}$). Experimenting with larger and smaller training data sets suggests that this is the minimum size for the algorithm to be able to accurately learn the relationship between the input spectrum and the output parameters. The Random Forest is set up in the same manner as before, using the same hyperparameters to train on normalised data. The trained Random Forest is then used to predict the parameter values for the 2000 synthetic spectra in the test set. As described in Section 3.2, we evaluate the likelihoods of the predictions made by each tree and take the median of the corresponding posterior distribution to be the predicted parameter value.

Figure 3.5 shows the outcome of our method when applied to these 2000 synthetic spectra. The R^2 coefficient of determination is close to unity for each parameter, suggesting that the retrieval is able to recover the input parameters well. For models with $\log X_{\text{H}_2\text{O}} \lesssim -6$ the correlation is much lower; this is to be expected since these cases correspond to non-detections and is consistent with the findings of MN18. For models with $\log X_{\text{H}_2\text{O}} \lesssim -6$, the spread in the results is caused by a degeneracy between P_{ref} and $X_{\text{H}_2\text{O}}$, as shown in Figure 3.6; it is possible to fit the same spectrum by increasing P_{ref} and decreasing $X_{\text{H}_2\text{O}}$. This degeneracy

Table 3.2 Description of priors for retrievals of the WFC3 transmission spectrum of HD 209458b.

Parameter	Lower Bound	Upper Bound	Prior
T_{iso} (K)	700	2810	uniform
P_{ref} (bar)	10^{-6}	10^2	log-uniform
$X_{\text{H}_2\text{O}}$	10^{-12}	10^{-2}	log-uniform

has been found previously when analysing WFC3 spectra (e.g. Pinhas et al., 2019, Welbanks & Madhusudhan, 2019).

3.3.2 Retrieval of WFC3 Transmission Spectrum

Having validated our retrieval method using synthetic data, we now apply the algorithm to a real data set for direct comparison against an AURA retrieval. We consider the observed WFC3 transmission spectrum of the hot Jupiter HD 209458b (Deming et al., 2013), which consists of 29 data points in the spectral range $1.1 - 1.7 \mu\text{m}$. We use this planet as a representative case to validate our method since it is one of the most well-studied planets in the literature, with high quality spectral data available. Additionally, the transmission spectrum of this planet has recently been analysed in numerous retrieval studies, (e.g. Barstow et al., 2017, MacDonald & Madhusudhan, 2017, Welbanks & Madhusudhan, 2019).

We first carry out a Nested Sampling retrieval using the same parameterisation as in Section 3.3.1. For this retrieval, the model is initially evaluated at a higher resolution, covering 1000 wavelength points from $1.1 - 1.7 \mu\text{m}$. The model spectrum is subsequently convolved with the WFC3 point spread function, integrated over the instrument function and binned to the resolution of the observed data. The binned spectrum is used to evaluate the likelihood function for each model (see equation 3.2). Further detail on this retrieval approach and validation using synthetic data can be found in Pinhas et al. (2018). We find that a Nested Sampling retrieval setup using 1000 live points is sufficient to obtain robust parameter estimates.

The retrieved posterior distributions, abundance estimates and uncertainties are displayed in Figure 3.7. These compare very closely to the results from Case 3 of Welbanks & Madhusudhan (2019). The only notable difference between the two retrievals is that our results show smaller error bars for the estimated reference pressure P_{ref} . This can be attributed to fixing R_p rather than retrieving it. We again find a degeneracy between P_{ref} and $X_{\text{H}_2\text{O}}$, in agreement with our findings from the model validation.

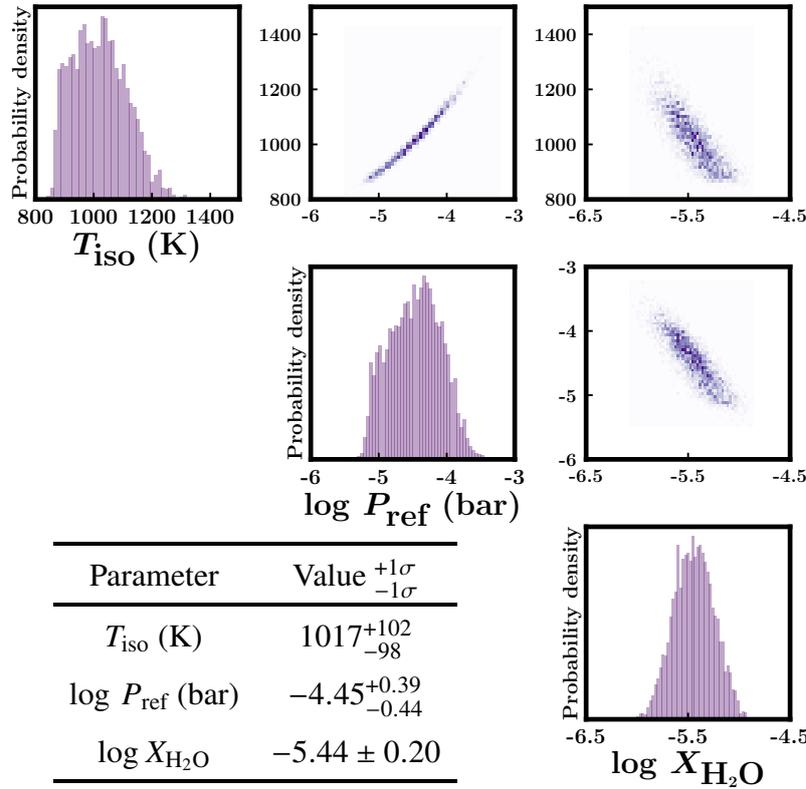


Fig. 3.7 Results of Nested Sampling retrieval of the WFC3 transmission spectrum of HD 209458b.

We attempt to reproduce the results of this Nested Sampling retrieval using the extended Random Forest method as described in Section 3.2.2. For this case study we use the same training set of 8000 spectra from Section 3.3.1. The Random Forest is set up in the same manner as before, using the same hyperparameters to train on normalised data. The likelihood function for each estimator is evaluated every time 1000 more estimators have been trained, and we find that the posterior distribution converges once 12000 estimators have been produced. Marginalised posterior distributions are then created from this likelihood function to obtain the results, which are displayed in Figure 3.8.

In this case, the extended Random Forest retrieval produces extremely similar results to the Nested Sampling retrieval. The best-fitting model spectra, along with 1σ and 2σ uncertainties, from the two retrievals are shown in Figure 3.9. The parameter estimates and uncertainties are directly compared in Table 3.3; both the retrieved median values and the 1σ uncertainties are almost identical. It is also clear from the joint distributions shown in Figure 3.8 that the extended Random Forest retrieval has found the same degeneracies between parameters as the Nested Sampling retrieval. This result demonstrates for the first time that,

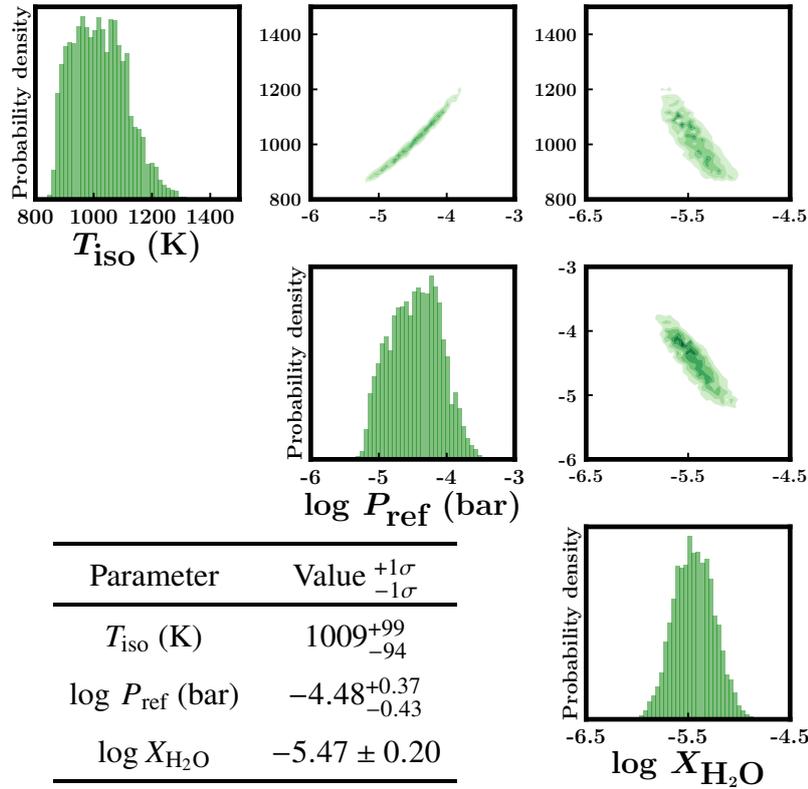


Fig. 3.8 Results of extended Random Forest retrieval of the WFC3 transmission spectrum of HD 209458b. The results are in agreement with the Nested Sampling retrieval in Figure 3.7.

for a given observation instance, a machine learning-based approach to atmospheric retrieval can not only obtain similar parameter estimates to a traditional retrieval, but that it can also deal with uncertainties and degeneracies in a robust and accurate way.

3.3.3 Addition of Unconstrained Free Parameter

The retrieval analysis of the WASP-12b transmission spectrum from MN18 has also been reproduced in Cobb et al. (2019). In that paper they find that in certain cases, the Random Forest retrieval can sometimes return a narrow posterior for a free parameter that should not be constrained; they demonstrate this by finding a synthetic spectrum following the model from Heng & Kitzmann (2017) for which the Random Forest confidently predicts H_2O , HCN and NH_3 abundances that are not the true values used to generate the model. In the present paper we investigate this issue further in order to determine whether the MN18 Random Forest approach might incorrectly infer certain parameter values in cases where a traditional

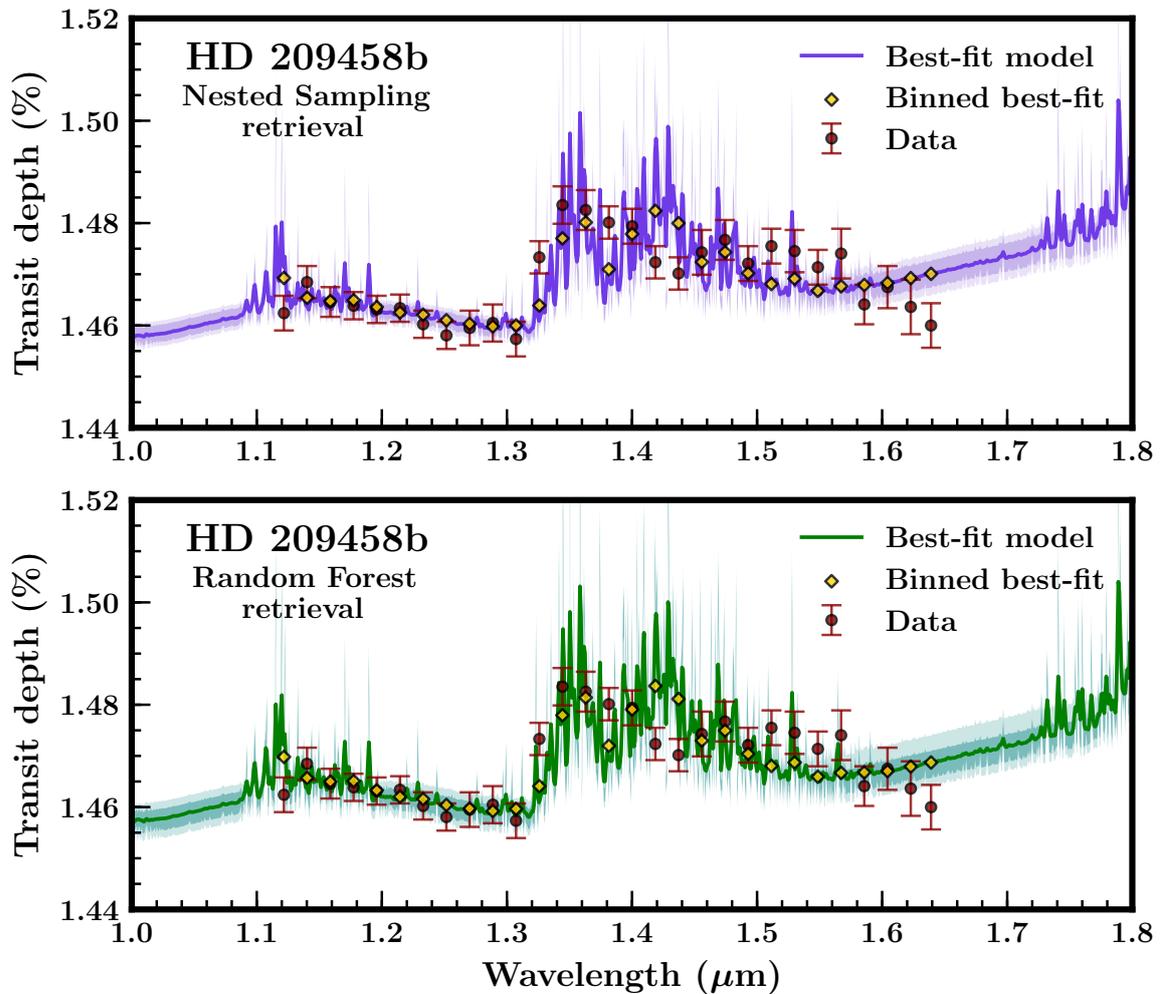


Fig. 3.9 *Top*: Best-fitting model spectrum from the Nested Sampling retrieval of the HD 209458b WFC3 transmission spectrum. *Bottom*: Best-fitting model spectrum from the extended Random Forest retrieval of the same spectrum. The shaded regions represent the 1σ and 2σ contours, produced by drawing 1000 spectra from the posterior distributions from each retrieval. The best-fitting spectra have been smoothed with a Gaussian filter for clarity.

retrieval would (correctly) not be able to constrain that value. We also consider whether a similar problem would occur if our extended Random Forest method was used instead.

In order to highlight this issue, we carry out a second set of retrievals of the WFC3 spectrum of HD 209458b, but this time including CO abundance as a free parameter in the forward model, with a log-uniform prior ranging from 10^{-12} – 10^{-2} . We choose to add CO since this molecule does not have strong features in the spectral range of the data. Previous studies such as Welbanks & Madhusudhan (2019) have therefore been unable to constrain the CO abundance from this spectrum. We verify this by first carrying out a Nested Sampling retrieval of the spectrum, whose results can be seen in Figure 3.10. As expected, the estimated values of T_{iso} , P_{ref} and $\log X_{\text{H}_2\text{O}}$ remain very close to those from the retrieval that did not include CO (see Figure 3.7), but the CO abundance itself is not unconstrained. This setup should therefore provide a test of the capabilities of both the MN18 and our Random Forest retrieval methods to deal with an unconstrained free parameter in the model.

We begin the machine learning approach by generating a training set consisting of 160000 model spectra, which we use for both the MN18 and for the extended Random Forest retrievals. We use the same training set in both cases to ensure that the only difference between the two retrievals is the implementation of the algorithm. First we employ the methods of MN18 to perform the retrieval using this data set. We add Gaussian noise to the training set and train 1000 estimators on the noisy spectra. Histograms of the results along with parameter estimates are shown in Figure 3.11. The shapes of the temperature and water abundance distributions differ somewhat from the Nested Sampling posteriors, but what is most notable is the apparent peak around -8.5 in \log CO abundance which is not present at all in the Nested Sampling case. As in Cobb et al. (2019), the algorithm is overconfident in its prediction of a parameter value which it should not be able to constrain (but cf. Fisher et al., 2020).

Next we take the same training set and apply our extended method as described in Section 3.2.2. More estimators are required to reach convergence in this case than when CO was not

Table 3.3 Comparison of Nested Sampling (NS) and extended Random Forest (RF+) retrieved parameter values from the WFC3 transmission spectrum of HD 209458b.

Parameter	NS Value $^{+1\sigma}_{-1\sigma}$	RF+ Value $^{+1\sigma}_{-1\sigma}$
T_{iso} (K)	1017^{+102}_{-98}	1009^{+99}_{-94}
$\log P_{\text{ref}}$ (bar)	$-4.45^{+0.39}_{-0.44}$	$-4.48^{+0.37}_{-0.43}$
$\log X_{\text{H}_2\text{O}}$	-5.44 ± 0.20	-5.47 ± 0.20

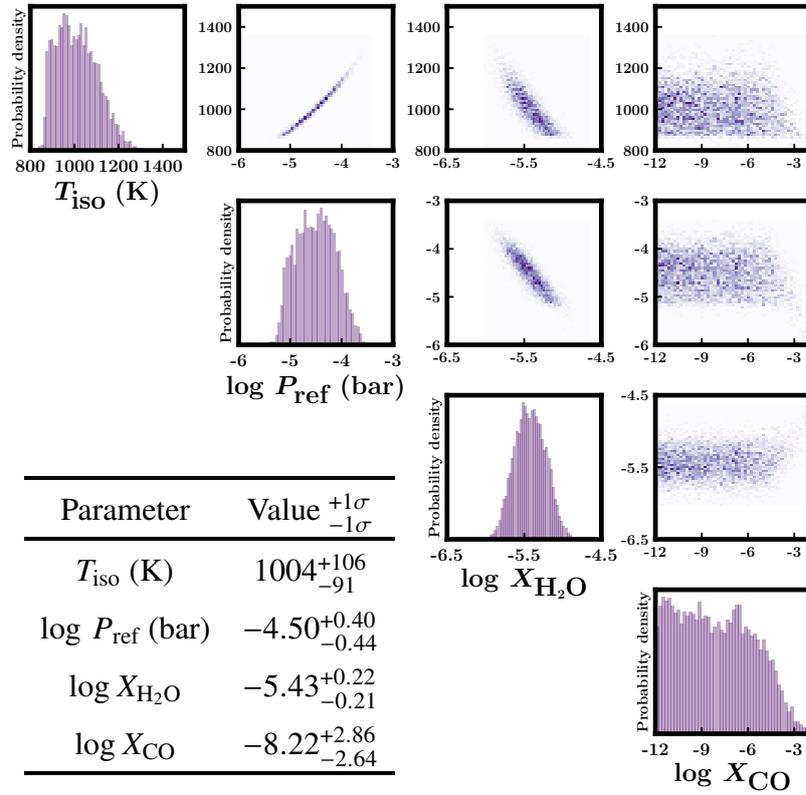


Fig. 3.10 Results of Nested Sampling retrieval of the WFC3 transmission spectrum of HD 209458b, including CO abundance as a free parameter. The non-detection of CO is consistent with expectation, given the weak CO features in the spectral range of the data.

included; in this case convergence is reached after 24000 estimators had been trained. The results from this analysis are shown in Figure 3.12. The marginalised posterior distributions of T_{iso} , P_{ref} and $\log X_{\text{H}_2\text{O}}$ are once again very similar in shape to their Nested Sampling counterparts. The extended Random Forest method produces a broad distribution, leaving the CO abundance unconstrained as in the Nested Sampling case. Since the likelihood function for each estimator is evaluated directly, this method is able to infer that the value of CO abundance does not affect how well the model fits the data. This means that our algorithm does not suffer from the flaws described in Cobb et al. (2019) and is able to deal with unconstrained free parameters in such a way that a false constraint is avoided. Figure 3.13 shows a direct comparison of the marginal distributions for CO abundance obtained in each of the retrieval studies. We also show the CO posterior for a retrieval following the methods of MN18, but using the same number of trees as in our extended method. The spurious peak around -8.5 is still present, indicating that increasing the number of trees alone is not enough to solve this problem.

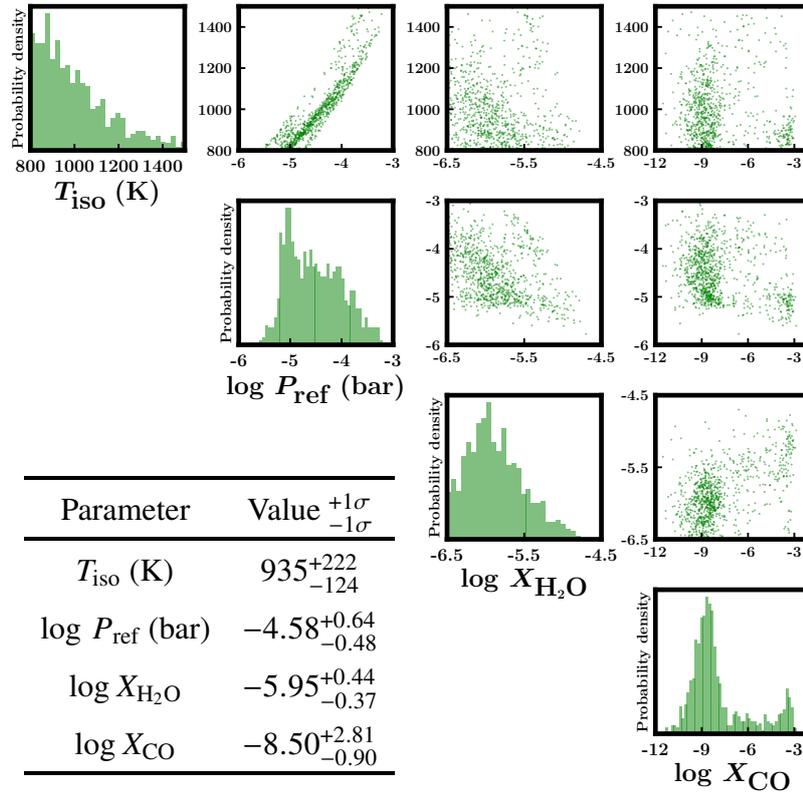


Fig. 3.11 Results of Random Forest retrieval of the WFC3 transmission spectrum of HD 209458b, including CO abundance as a free parameter, following the methods of MN18. The peak in the CO posterior is unphysical given that there is no strong CO feature in the spectral range of the data.

3.4 Discussion and Conclusions

It has previously been suggested that using machine learning to perform retrievals could significantly reduce computation time, since a trained machine learning algorithm can make predictions extremely rapidly. The Generative Adversarial Network presented by Zingales & Waldmann (2018) can predict model parameters from a spectrum in approximately 2 minutes, and Cobb et al. (2019) state that their approach can provide predictions in 1.5 seconds. The prediction time is not reported in MN18, but we find that 1000 estimators take a few seconds to make predictions for each retrieval considered in this study.

These numbers ostensibly suggest that machine learning retrievals are much faster than traditional methods, which can often take up to several hours depending on the size of the parameter space. However, these figures do not include the time taken to produce a training data set, nor do they incorporate the time taken to train the machine learning algorithm. According to Zingales & Waldmann (2018) the training phase of their GAN using a forward

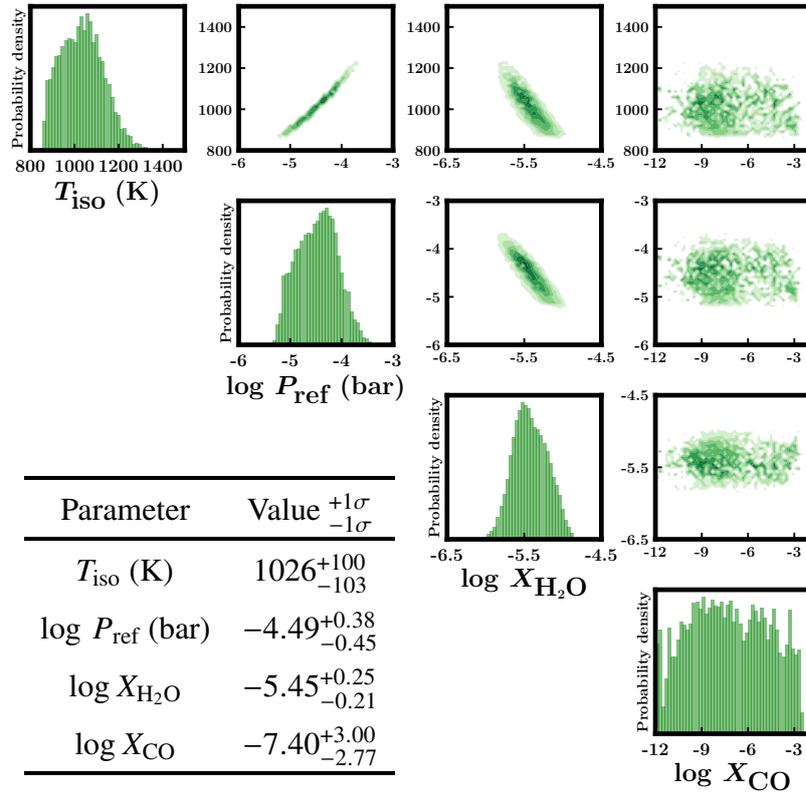


Fig. 3.12 Results of extended Random Forest retrieval of the WFC3 transmission spectrum of HD 209458b, including CO abundance as a free parameter. The results are consistent with expectations and the Nested Sampling retrieval shown in Figure 3.10.

model with 7 free parameters takes approximately three days per epoch on 20 CPU cores or about 9 hours per epoch on a GPU. The authors do not report how many epochs of training were required to fully train the network, nor do they say how long it took to generate the grid of 10^7 models that were used for the training. Cobb et al. (2019) do not produce a unique data set for their retrievals, but instead use the same training set that was used in MN18. Each of their models takes approximately 20 minutes to train.

For the present study we compare the full retrieval duration of Nested Sampling and extended Random Forest retrievals for both the three- and four-parameter cases presented in Section 3.3. We conduct retrievals on a synthetic data set binned to different resolutions from $R = 10$ to 1000, using the same computational resources for each (parallelisation over 4 CPU cores). For the Random Forest retrieval we include the time taken to produce the training data set, the training itself and the prediction step, however we note that in general only one training set would be needed to retrieve multiple observation instances of the same target.

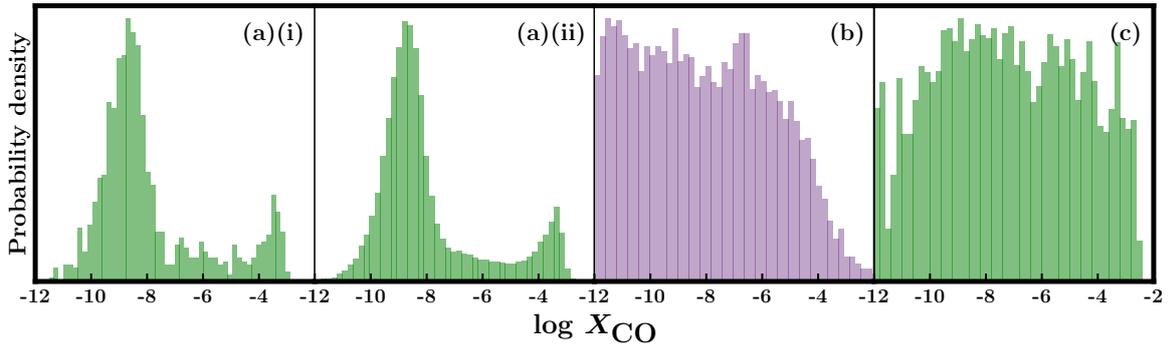


Fig. 3.13 Marginal posterior distributions for CO abundance retrieved from the transmission spectrum of HD 209458b using different methods: (a)(i) Random Forest retrieval following the methods of MN18; (a)(ii) the same approach as (a)(i) but using 24000 trees instead of 1000; (b) Nested Sampling retrieval; (c) extended Random Forest retrieval. In order to aid comparison, the histograms have been scaled so that the bin with maximum probability in each plot has the same height. Given the lack of strong CO features in the spectral range of the data, a detection of CO is unexpected. Therefore the posterior in (a) is unphysical while those in (b) and (c) match with expectation.

The results of this investigation are shown in figure 3.14, where we show the relative speedup of the extended Random Forest retrieval compared to the Nested Sampling retrieval. In the three parameter case, the Random Forest retrieval always outperforms Nested Sampling by a factor of ~ 4 to 8. At the resolution of the HD 209458b data used in Section 3.2.2, training time is approximately 4 seconds per 1000 estimators using the extended Random Forest approach. The biggest improvement over Nested Sampling is found at the lowest and highest wavelength resolutions, with a minimum at $R \sim 250$. While the duration of the Random Forest retrieval increases steadily with wavelength, following an approximate power law $\tau_{\text{RF}} \sim R^{0.4}$, the duration of the Nested Sampling retrieval increases more slowly with R up to $R \sim 250$ at which point τ_{NS} increases rapidly.

In the four parameter case, similar patterns are found in both types of retrieval, with both retrievals being slower overall. However, the addition of another parameter increases τ_{RF} much more than τ_{NS} , resulting in retrievals of comparable duration. Training a Random Forest on high-dimensional data is much slower since a larger training set is required; in this case it takes about 80s to train 1000 estimators on 4 CPU cores. This indicates that increasing the number of free parameters and the size of the training set slows down the training significantly.

The applicability of a retrieval algorithm to higher-dimensional parameter spaces is an important factor to consider when comparing machine learning and traditional retrievals. As higher quality observed data and new line lists (e.g. Tennyson & Yurchenko, 2012)

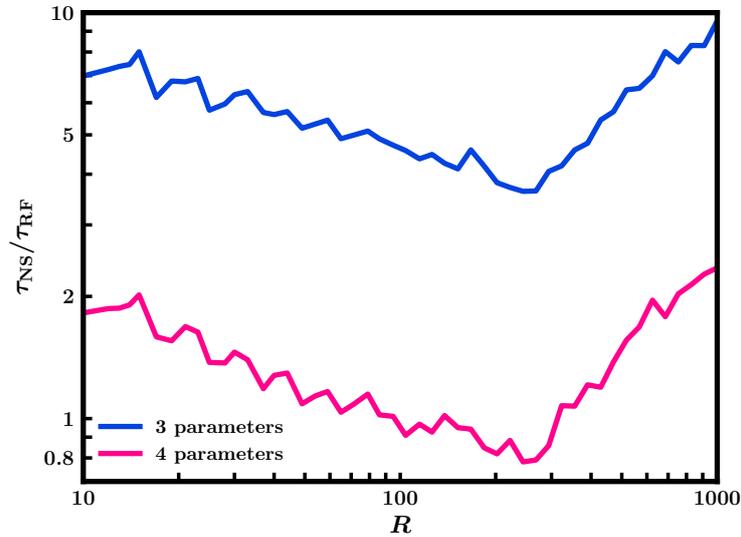


Fig. 3.14 Comparison of the duration of Nested Sampling retrievals (τ_{NS}) versus extended Random Forest retrievals (τ_{RF}) as a function of spectral resolution R . Three-parameter Random Forest retrievals are faster than their Nested Sampling counterparts, whereas four-parameter Random Forest retrievals are only faster at low and high spectral resolutions, albeit only with a factor of two. For higher dimensions, Nested Sampling retrievals tend to be more efficient.

become available it will be possible to search for an increasingly large number of atomic and molecular species, which will expand the number of possible free parameters in the forward model considerably. Additionally, extra parameters must be included to deal with other phenomena such as clouds (Wakeford & Sing, 2015, Pinhas & Madhusudhan, 2017) and stellar heterogeneity (Pinhas et al., 2018). We consider the feasibility of using the extended Random Forest method to perform a retrieval including additional chemical species and cloud/haze properties, following the prescription of MacDonald & Madhusudhan (2017). The model in this case requires 10 free parameters, and so we produce a large training data set consisting of $> 10^6$ models. When a Random Forest is trained using this data, it predicts very few points with high likelihoods, suggesting that the parameter space is not sampled finely enough in the training set. This approach already requires far more model evaluations than a Nested Sampling retrieval using the same model, which converges after approximately 500000 model evaluations. A Random Forest retrieval with n free parameters appears to require $\gtrsim 10^n$ models for an adequate training set. Full retrievals with optical and infrared data typically include up to ~ 20 free parameters (MacDonald & Madhusudhan, 2019, Madhusudhan et al., 2020), so a sufficient training set to carry out these retrievals using

this method would be prohibitively large. We therefore conclude that the Random Forest approach struggles to deal with higher dimensional parameter spaces efficiently. If a different algorithm can be used that performs more efficiently in high-dimensional parameter spaces and while still capturing uncertainties and degeneracies accurately, then it is possible that a machine learning approach could eventually compare to traditional methods for all cases, regardless of complexity.

At present, each planet being analysed requires its own training data set for our extended Random Forest retrieval. An alternative approach to address the large computation time required for Random Forest retrievals in high-dimensional spaces might be to generate a large training set with many free parameters (including bulk parameters such as surface gravity) that would be applicable to a range of planets. This training set could then be used to train a model which could retrieve properties of spectra from different planets. Each retrieval would therefore only require the prediction step to be carried out after the training has been completed once. This kind of method has been applied to retrievals using deep learning: for example, Soboczenski et al. (2018) used a large training set of 10^7 models of terrestrial planet spectra which could be applied to numerous planets. Zingales & Waldmann (2018) took a similar approach, using a training set of 10^7 hot Jupiter spectra. In the next decade, with the advent of JWST (Greene et al., 2016) and the *Atmospheric Remote-sensing Infrared Exoplanet Large-survey* (ARIEL; Tinetti et al., 2018), we expect the number of planets with high-quality spectral data suitable for retrieval to increase significantly, so a machine learning approach that could apply to a range of planets may prove to be the most efficient option for conducting population studies of exoplanet atmospheres in the future.

Other than attempting to fully reproduce traditional retrievals, there may be other scenarios in which a machine learning-based approach could prove useful in this field: for example, a small number of predictions may be able to inform the starting points for a MCMC retrieval (e.g. Hayes et al., 2020), or could inform which molecules should be included in the full retrieval process, as discussed in Waldmann (2016). Additionally, while not explored in this study, the Random Forest algorithm provides information about the information content of each data point in the spectrum, and it was mentioned in MN18 that this could be used to inform which wavelengths are most useful for future observations. We believe that combining machine learning algorithms with traditional methods can provide additional insight even when they are unable to replace existing methods entirely.

In this paper we have investigated the viability of using machine learning for atmospheric retrievals of exoplanets. We reproduced both the Nested Sampling and Random Forest results of MN18 and we extended the methods from that paper so that the resulting posterior distribution from the Random Forest retrieval more closely matches that of the corresponding

Nested Sampling retrieval. We applied this extended approach to a different planetary spectrum using a fully numerical forward model and found that once again we could accurately match the Nested Sampling and Random Forest retrievals. In addition, we found that our approach does not lead to spurious detections of parameters in cases where the parameter values should not be well-constrained, a problem found with the previous method. We have therefore developed a machine learning technique that can accurately and robustly reproduce the results of Bayesian retrievals. We investigated the potential for using this method to perform higher-dimensional retrievals and found that the algorithm requires a finely-sampled grid of training data in order to work well, making it prohibitively expensive to use this method in more complex cases. We conclude that while it is certainly possible to use machine learning techniques to reproduce traditional Bayesian retrieval results at least in low dimensions, the increased computational cost suggests that this approach does not yet provide a significant improvement on traditional methods. Future improvements in machine learning methodologies, as well as new strategies for applying these techniques to the present problem, will be required to surmount this challenge.

Chapter 4

A Three-Dimensional Retrieval Framework for Exoplanet Transmission Spectra

“It’s going to be a swell demonstration — and at no time will we be stooping to any cheap 3D tricks.”

– Kermit the Frog

“Did you say, ‘cheap 3D tricks’?!”

– Fozzie Bear

4.1 Introduction

The study of exoplanet atmospheres has seen immense progress in recent years, with spectroscopic observations enabling constraints on a range of physical and chemical processes that occur in exoplanets (Madhusudhan, 2019). One of the most successful observational techniques used to characterise exoplanet atmospheres has been transmission spectroscopy, which measures the wavelength-dependent decrease in flux from the host star as the planet transits in front of it. This approach has led to detections and abundance constraints of chemical species in numerous exoplanets (e.g. Charbonneau et al., 2002, Snellen et al., 2010, Deming et al., 2013, Mandell et al., 2013, Kreidberg et al., 2014b, Madhusudhan et al., 2014b, Wyttenbach et al., 2015, Wakeford et al., 2018, Pinhas et al., 2019) as well as providing constraints on the properties of clouds and hazes (e.g. Pont et al., 2008, Bean et al., 2010, Pont et al., 2013, Kreidberg et al., 2014a, Sing et al., 2016, Barstow et al., 2017, Pinhas et al., 2019, Benneke et al., 2019).

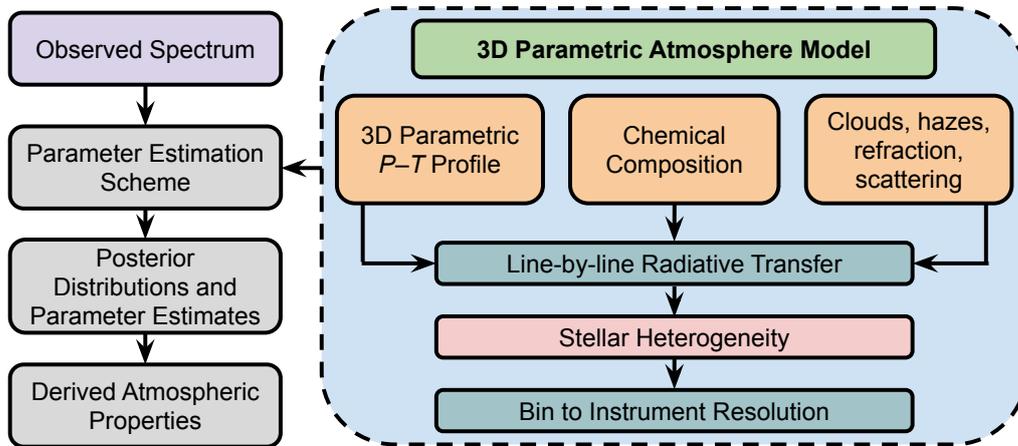


Fig. 4.1 Schematic of the AURA-3D retrieval framework. The algorithm combines an atmospheric forward model in 3D geometry with a Bayesian sampling algorithm to conduct parameter estimation. The forward model is modular in structure, with numerous features that can be incorporated as needed, such as clouds/hazes, refraction, scattering and stellar heterogeneity.

One of the most common methods for deriving atmospheric properties from spectroscopic data of exoplanets is atmospheric retrieval (Madhusudhan & Seager, 2009). An atmospheric retrieval algorithm involves combining an atmospheric forward model with a statistical sampling algorithm in order to determine estimates of the atmospheric properties that best explain the observed data (see e.g. Madhusudhan, 2018, for a recent review). The earliest retrieval algorithms employed a grid-based search to find best-fitting models (Madhusudhan & Seager, 2009), with later studies incorporating Bayesian statistical inference schemes such as Markov Chain Monte Carlo (e.g. Madhusudhan & Seager, 2010, Line et al., 2013, Cubillos et al., 2013, Zhang et al., 2019, Lacy & Burrows, 2020), and Nested Sampling (e.g. Benneke & Seager, 2013, Waldmann et al., 2015a, Oreshenko et al., 2017, MacDonald & Madhusudhan, 2017, Gandhi & Madhusudhan, 2018, Mollière et al., 2019, Zhang et al., 2020). Other approaches to parameter estimation, such as Optimal Estimation (e.g. Irwin et al., 2008, Lee et al., 2012, Barstow et al., 2017) and machine learning (e.g. Márquez-Neila et al., 2018, Zingales & Waldmann, 2018, Cobb et al., 2019, Fisher et al., 2020, Nixon & Madhusudhan, 2020), have also been implemented.

In order to create an accurate and flexible atmospheric model for use in retrievals, a large number of parameters are required to capture all possible compositions, thermal structures, and other properties such as the presence of clouds or hazes. Since statistical sampling algorithms often require a large number of model evaluations, increasing as the number of free parameters increases (Fortney et al., 2021), it is necessary to make some simplifying

assumptions to reduce computation time. Common assumptions include the use of a 1D temperature profile and uniform chemical abundances in the region of the atmosphere probed by the observations (Madhusudhan, 2018). It is important to re-examine the assumptions made in these algorithms as we enter a new era in the study of exoplanet atmospheres, in which the quality and volume of observational data will improve immensely thanks to next-generation facilities such as JWST (Greene et al., 2016), ARIEL (Tinetti et al., 2018) and the Extremely Large Telescopes (ELTs; Snellen et al., 2015).

Despite the successes of the retrieval approach for analysing data from existing observatories such as the *Hubble* Space Telescope (HST), several studies have highlighted the potential for 1D retrieval codes to lead to biased results due to variations in temperature structure and chemical abundances with latitude and longitude, particularly when considering high-quality observations of very hot Jupiters from upcoming facilities such as JWST. For example, assumptions of a 1D atmosphere have been shown to bias retrievals of emission spectra (e.g. Blečić et al., 2017, Taylor et al., 2020) and can lead to unrealistic model fits to phase curves (e.g. Irwin et al., 2020, Feng et al., 2020).

This work focuses on retrievals of transmission spectra. Transmission spectroscopy probes the day-night terminator region of the atmosphere, across which inhomogeneities may be expected (e.g. Fortney et al., 2010). These effects should be particularly prevalent in hot and ultra-hot Jupiters, since day-night temperature contrasts are expected to increase as planetary equilibrium temperature increases (Cowan & Agol, 2011, Komacek & Showman, 2016). Previous studies have already indicated that 1D retrievals can lead to biases when interpreting transmission spectra of hot giant planets. Caldas et al. (2019) developed a transmission spectrum model that could incorporate a 3D atmospheric structure, which they used to investigate the effect of day-night temperature gradients on model spectra. They used the 1D retrieval code TAUReX (Waldmann et al., 2015a) to analyse synthetic JWST spectra generated with their 3D model for which the temperature could vary significantly between the day- and nightsides. The retrievals returned terminator temperatures that were biased towards the dayside temperature, and chemical abundances which were not consistent with input values. Pluriel et al. (2021) also found day-night induced biases for a range of synthetic JWST-like spectra of hot and ultra-hot Jupiters, and Pluriel et al. (2020) used the same 3D model to explore the potential for biases caused by day–night chemical heterogeneities. They focused on the ultra-hot Jupiter WASP-121b, for which the H₂O abundance is expected to be much higher on the nightside than the dayside due to thermal dissociation (Parmentier et al., 2018). A 1D retrieval analysis showed that inferred chemical abundances could be significantly different to the input values. MacDonald et al. (2020) also discussed the potential for biased temperatures and abundances from 1D retrievals (but cf. Welbanks &

Madhusudhan, 2022). A generalised retrieval framework should be able to address the above issues in order to accurately infer atmospheric properties regardless of the presence of thermal inhomogeneities. This is the goal of the present study.

Unlike the simplified atmospheric models used in retrieval algorithms, General Circulation Models (GCMs) can be used to capture some of the more complex processes at work in planetary atmospheres (see Section 1.3.3.1). GCMs simulate the dynamics of 3D atmospheres, often incorporating chemical processes as well as radiative transport (Showman et al., 2020). GCMs have predicted a number of important features of hot Jupiter atmospheres. This includes the finding that hot Jupiters can show large day-night temperature contrasts of up to several hundred K (e.g. Showman & Guillot, 2002, Cooper & Showman, 2005, Amundsen et al., 2016). GCMs of hot Jupiters have also predicted equatorial superrotation that can lead to an eastward shift of a planet’s dayside hotspot away from the substellar point (e.g. Showman & Guillot, 2002, Cooper & Showman, 2005, Rauscher & Menou, 2010, Kataria et al., 2013). These predictions have subsequently been confirmed through comparison with observations of infrared phase curves (e.g. Knutson et al., 2007, Komacek et al., 2017, Stevenson et al., 2017).

Model transmission spectra can be generated using 3D temperature structures. Fortney et al. (2010) used pre-computed temperature profiles from SPARC/MITgcm simulations of hot Jupiters (Showman et al., 2009) to examine the differences between 1D and 3D model atmospheres. A number of more recent works have also developed multidimensional forward models for transmission spectra. Caldas et al. (2019) used a 3D radiative transfer model to investigate biases caused by day-night temperature gradients using a simple parametric temperature profile. Falco et al. (2021) presented transmission spectrum models in 1D, 2D and 3D, and MacDonald & Lewis (2022) described parametric prescriptions for 3D atmospheric models.

Although GCMs are too computationally expensive to incorporate into retrieval algorithms directly, a number of studies have aimed to bridge the gap between 1D and 3D models in the context of retrievals of transmission spectra. Lacy & Burrows (2020) conducted chemical equilibrium retrievals assuming separate dayside and nightside temperature profiles, showing that in some cases it is possible to constrain day- and nightside temperatures of hot Jupiters from their transmission spectra. Espinoza & Jones (2021) demonstrated how JWST could be used to acquire separate transmission spectra for each of a planet’s limbs, and also expanded the retrieval framework CHIMERA (Line et al., 2013) to enable retrievals of a single transmission spectrum with separate temperature profiles for each limb, under the assumption of chemical equilibrium. Welbanks & Madhusudhan (2022) also presented a 1+1D retrieval framework in which the final transmission spectrum is calculated by a linear

combination of two separate spectra, one representing the morning terminator and another representing the evening terminator.

In this work we present AURA-3D, a 3D atmospheric retrieval framework for transmission spectra of exoplanets. In Section 4.2, we describe our algorithm in detail. We present a forward model that incorporates a 3D temperature structure and which can be used to efficiently generate transmission spectra from the output of a GCM, similarly to several past studies (Fortney et al., 2010, Caldas et al., 2019). We develop a parametric pressure–temperature (P – T) profile that can match the azimuthally-averaged structure of a GCM while being very fast to compute, making it suitable for atmospheric retrieval. In Section 4.3.1 we compare our parametric temperature profile to a number of established GCMs. We explore the effects of a day-night temperature gradient on resulting transmission spectra in Section 4.3.2 by quantifying the difference in transit depths between 3D models and their 1D-averaged counterparts. In Section 4.3.3 we discuss the possible effects of chemical inhomogeneity on transmission spectra. In Section 4.3.4 we demonstrate the capability of AURA-3D to carry out retrievals with a multidimensional P – T profile by conducting a retrieval on a synthetic hot Jupiter spectrum generated using our new parametric profile. In Section 4.3.5 we retrieve a spectrum generated using a GCM temperature structure, comparing results from retrievals both with and without a day-night temperature gradient. We summarise our findings and discuss avenues for the future in Section 4.4.

4.2 Methods

The retrieval framework developed in this study is shown in Figure 4.1. Similarly to previous retrieval algorithms, this framework combines an atmospheric forward model with a Bayesian parameter estimation scheme. The key differences between our methods and those of previous retrieval algorithms concern the forward model. Our forward model enables the computation of a transmission spectrum whose temperature structure can vary in three dimensions throughout the atmosphere, and can therefore take the temperature structure calculated by a GCM as its input. However, since the calculation of a temperature structure using a GCM is very computationally expensive, we also describe a number of parameterisations that allow for rapid computation of a multidimensional temperature profile, and which can be incorporated directly into a retrieval. Our retrieval framework also incorporates several components from the existing retrieval code AURA and its subsequent implementations (Welbanks & Madhusudhan, 2019, Nixon & Madhusudhan, 2020) as well as *Aurora* (Welbanks & Madhusudhan, 2021). These features are discussed later in this section.

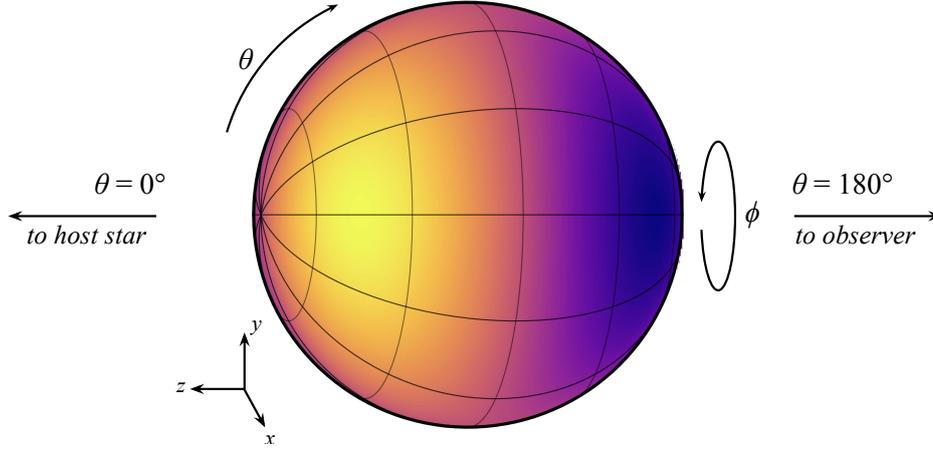


Fig. 4.2 Three-dimensional co-ordinate system adopted for our forward model. The zenith angle θ is defined to be 0 at the substellar point and π (180°) at the antistellar point. The azimuthal angle ϕ varies between 0 and 2π (360°), increasing in the direction of the trailing limb.

We note that *Aurora* can also be applied to H-poor atmospheres, which are not considered in the present work.

4.2.1 Transit Geometry

We define a spherical polar coordinate system (r, θ, ϕ) with an origin at the centre of the planet (see Figure 4.2). The coordinates $(R_p, 0, 0)$ correspond to the substellar point. The zenith angle θ varies between 0 and π , reaching $\pi/2$ at the day-night terminator and π at the antistellar point. The azimuthal angle ϕ varies between 0 and 2π , increasing in the direction of the trailing hemisphere of the planet. This geometry, based on the star-planet-observer axis, differs from a traditional latitude-longitude grid based on the rotation axis which is often used when describing the geometry of a planet. We adopt this approach since it allows for a straightforward description of day-night variations in atmospheric properties, since the zenith angle θ moves from the substellar to antistellar point. This geometry has been used in previous studies involving transmission spectroscopy (e.g. Fortney et al., 2010, Caldas et al., 2019).

Rays of light traveling from a host star to an observer may pass through the atmosphere of the transiting planet (see Figure 2.2). A single ray is assumed to have a fixed impact parameter b and a fixed azimuthal angle ϕ as it moves through the atmosphere. The ray travels along the path s , defined so that $s = 0$ when the ray is directly above the planet's terminator ($\theta = 90^\circ$). As it travels along this path, the zenith angle θ increases (moving from

the dayside to the nightside of the planet) and the distance between the ray and the centre of the planet (denoted r) also varies. The values of θ and r are given by

$$r^2 = s^2 + b^2 \quad (4.1)$$

and

$$\cos \theta = \frac{s}{r}. \quad (4.2)$$

4.2.2 Radiative Transfer

The following expression describes the transit depth Δ_λ for the generalised atmosphere of a planet with radius R_p transiting a star with radius R_* :

$$\Delta_\lambda = \frac{1}{2\pi} \int_0^{2\pi} \delta_\lambda(\phi) d\phi, \quad (4.3)$$

where

$$\delta_\lambda(\phi) = \frac{1}{R_*^2} \left[R_p^2 + 2 \int_{R_p}^{R_p+H} b \left(b(1 - e^{-\tau_\lambda(b,\phi)}) \right) db - 2 \int_0^{R_p} b e^{-\tau_\lambda(b,\phi)} db \right]. \quad (4.4)$$

In this equation, the atmospheric height is denoted by H , and $\tau(b, \phi)$ represents the optical depth at impact parameter b and azimuthal angle ϕ . The total transit depth Δ_λ is found by integrating the transit depth at each value of ϕ over all azimuthal angles.

A crucial difference between the 3D transmission spectrum model and a 1D approximation is that the attenuation coefficient at a given wavelength, μ_λ , depends on r , θ and ϕ , as opposed to only depending on r . The expression for the optical depth along a ray path s at a given (b, ϕ) is therefore

$$\begin{aligned} \tau_\lambda(b, \phi) &= \int_{-s_0}^{s_0} \mu_\lambda(r, \theta, \phi) ds \\ &= \int_{-s_0}^{s_0} \mu_\lambda(s, b, \phi) ds. \end{aligned} \quad (4.5)$$

Previous retrieval codes have included forward models with ϕ -dependence to capture properties such as inhomogeneous cloud cover (Line & Parmentier, 2016, MacDonald & Madhusudhan, 2017, Welbanks & Madhusudhan, 2021). AURA-3D also incorporates ϕ -dependence to model inhomogeneous clouds (see Section 4.2.5.1). By including a

temperature profile that depends on θ (see 4.2.3.3), our retrieval algorithm is therefore capable of constraining atmospheric properties in three dimensions.

Our model incorporates absorption from a wide range of chemical species following the methods described in Gandhi & Madhusudhan (2017, 2018), using cross-sections derived from line list data from a number of different sources. A Voigt function is used to apply temperature and pressure broadening. We assume a H/He-dominated atmosphere with solar abundances of H and He (Asplund et al., 2009). Additional chemical species present in the model which are used in this work include H₂O (Rothman et al., 2010), CH₄ (Yurchenko & Tennyson, 2014), NH₃ (Yurchenko et al., 2011), HCN (Barber et al., 2014), CO and CO₂ (Rothman et al., 2010). The attenuation coefficient of a given chemical species i can be expressed as

$$\mu_i(\lambda, P, T) = \rho_i \kappa_i(\lambda, P, T) = n_i \sigma_i(\lambda, P, T), \quad (4.6)$$

where ρ_i is the mass density of species i , κ_i is the opacity of species i , n_i is the number density of species i and σ_i is the absorption cross-section of species i . The number density n_i of a given species can be related to its volume mixing ratio, $X_i = n_i/n_{\text{tot}}$, where n_{tot} is the total number density. The volume mixing ratio of each chemical species, apart from H₂ and He, is a free parameter in the model. The volume mixing ratios of H₂ and He are given by

$$X_{\text{H}_2} = \frac{1 - \sum_{i, i \neq \text{H}_2, \text{He}} X_i}{1 + (X_{\text{He}}/X_{\text{H}_2})}, \quad (4.7)$$

$$X_{\text{He}} = 0.17 X_{\text{H}_2}, \quad (4.8)$$

where the value of 0.17 is derived from a solar composition (Asplund et al., 2009).

In past implementations of AURA, cross-sections of each species are stored on a three-dimensional grid of λ , P and T . In order to determine the total attenuation coefficient at a given height r in the atmosphere, the cross-sections of each chemical species are interpolated using the values of P and T found at r using the equation of hydrostatic equilibrium and the prescribed P - T profile. This yields a two-dimensional array $\mu_i(\lambda, r)$ for each species. In the present work, since each combination of values of ϕ and θ may have a different P - T profile, the cross-sections must be interpolated for each (ϕ, θ) pair, yielding a four-dimensional array $\mu_i(\lambda, r, \theta, \phi)$. The attenuation coefficient for all chemical species acting as absorbers in the model is obtained by summing the attenuation coefficients for each individual species.

We also include collision-induced absorption (CIA) due to H₂-H₂ and H₂-He (Richard et al., 2012). The attenuation coefficient due to CIA is given by the expression

$$\mu_{\text{CIA}} = X_{\text{H}_2} n_{\text{tot}}^2 [X_{\text{H}_2} \sigma_{\text{H}_2-\text{H}_2}(\lambda, T) + X_{\text{He}} \sigma_{\text{H}_2-\text{He}}(\lambda, T)] \quad (4.9)$$

where $\sigma_{\text{H}_2-\text{H}_2}$ and $\sigma_{\text{H}_2-\text{He}}$ are the H_2 - H_2 and H_2 - He cross-sections respectively. We note that unlike the cross-sections for chemical species described above, which have units of m^2 , the CIA cross-sections have units of m^5 .

4.2.3 A Multidimensional Parametric Temperature Profile

Our radiative transfer model is capable of taking any 3D temperature structure as its input. The model can therefore be used for post-processing GCMs to create transmission spectra. However, given that a single GCM typically takes at least several days to compute, such a model is not appropriate for atmospheric retrieval, which can require $\gtrsim 10^6$ model computations in order to analyse a single spectrum (Fortney et al., 2021). The temperature profile must therefore be simplified to enable rapid computation. The simplest approach is to consider an entirely isothermal temperature profile when carrying out retrievals (e.g. Waldmann et al., 2015a, Zhang et al., 2019). The analytic temperature profile described in Guillot (2010) has been incorporated into retrieval algorithms (e.g. Benneke & Seager, 2012, Line et al., 2013), and the parametric temperature profile described in Madhusudhan & Seager (2009) has also been employed in numerous retrieval frameworks (e.g. Pinhas et al., 2018, Blečić et al., 2021). While the latter two approaches result in a temperature profile that varies with height in the atmosphere, these profiles do not vary with longitude or latitude.

AURA-3D incorporates several different parametric prescriptions for P - T profiles that include temperature variations in multiple dimensions. This enables fast, flexible model computation while also allowing for varying levels of complexity in the temperature structure. We present three possible prescriptions: one in which temperature varies with height in the atmosphere but not with θ , one in which temperature varies with θ but not with height, and one in which temperature varies with height as well as with θ and ϕ .

4.2.3.1 Temperature varying with height only

In the case where the temperature profile does not vary with θ or ϕ , we adopt the 1D parametric P - T profile from Madhusudhan & Seager (2009). The atmosphere is divided into three layers, defined by $P_{1,2,3}$. The deepest layer is isothermal while the upper layers have thermal gradients controlled by $\alpha_{1,2}$. The full profile is defined as:

$$\begin{aligned}
 P &= P_0 e^{\alpha_1 \sqrt{T-T_0}}, & P_0 &< P < P_1, \\
 P &= P_2 e^{\alpha_2 \sqrt{T-T_2}}, & P_1 &< P < P_3, \\
 T &= T_3, & P &> P_3,
 \end{aligned} \tag{4.10}$$

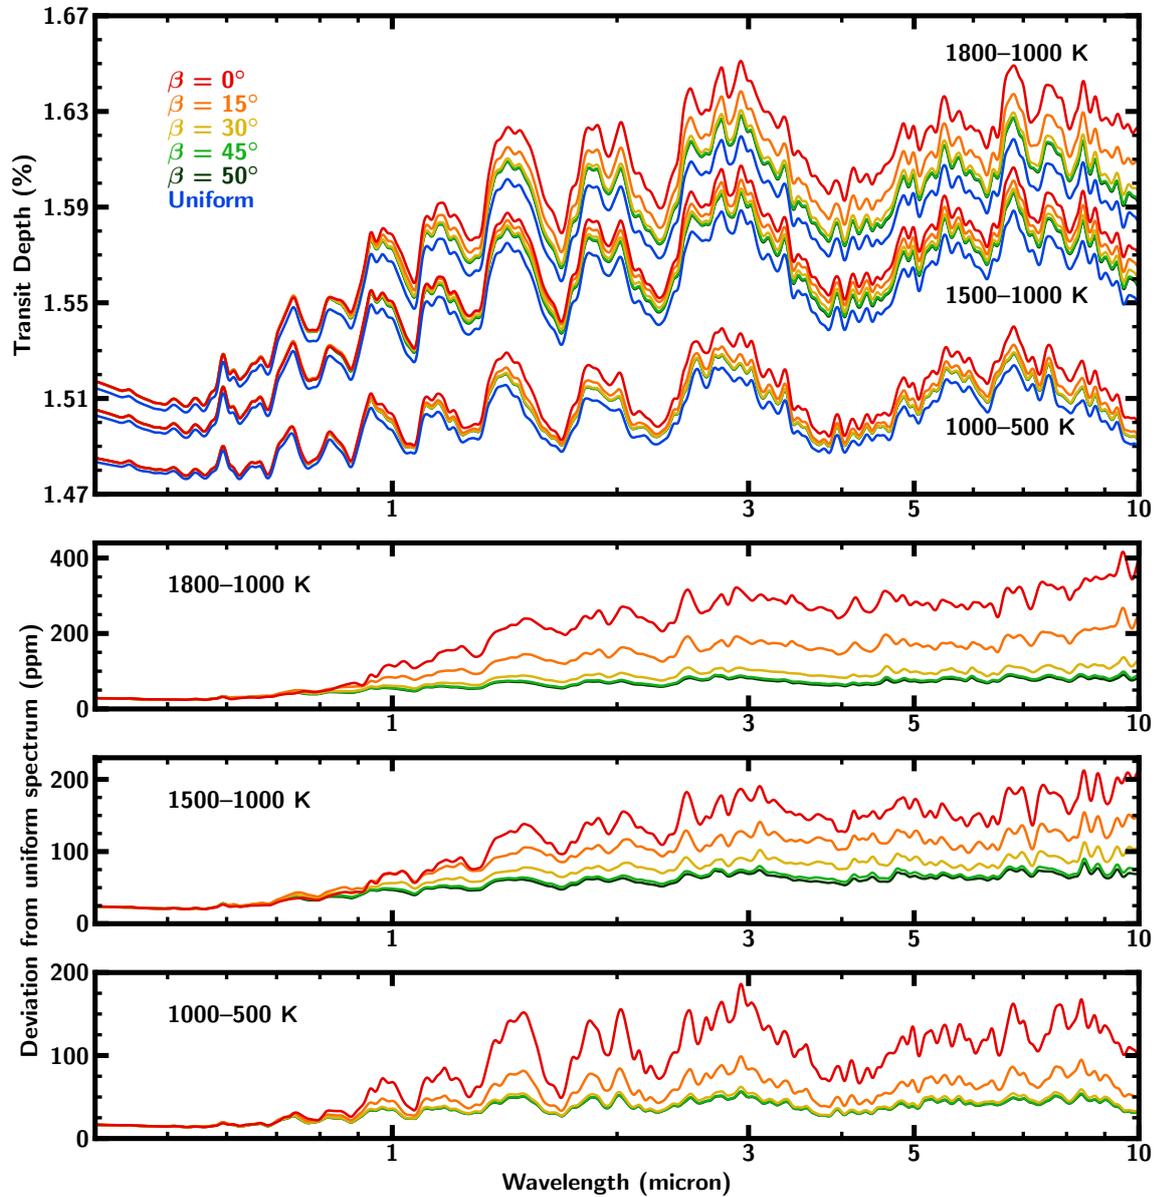


Fig. 4.3 *Upper panel*: Model transmission spectra of the canonical hot Jupiter HD 209458b with varying temperatures and transition angles (β). For these models, we adopt the parametric temperature profile described in Caldas et al. (2019). We obtain the same effects as described in their paper: namely, that the transit depth decreases as β increases, and the divergence between models increases at larger wavelengths. *Lower panels*: Difference in ppm between spectra with the non-uniform temperature profiles described above and uniform spectra with averaged temperature profiles.

where (P_0, T_0) defines the conditions at the top of the atmosphere. This can be recast to yield expressions for $T(P)$:

$$T(P) = \begin{cases} T_0 + \left(\frac{\ln(P/P_0)}{\alpha_1} \right)^2, & P_0 < P < P_1, \\ T_2 + \left(\frac{\ln(P/P_2)}{\alpha_2} \right)^2, & P_1 < P < P_3, \\ T_2 + \left(\frac{\ln(P_3/P_2)}{\alpha_2} \right)^2, & P > P_3. \end{cases} \quad (4.11)$$

This P – T profile was developed to be capable of emulating observed temperature profiles of solar system planets as well as profiles from self-consistent exoplanet atmosphere models. This flexibility, along with a rapid computation time, makes it an ideal prescription for retrievals. Our new profile presented in section 4.2.3.3 applies the same philosophy to 3D models in order to extend our retrieval framework to incorporate multidimensional effects.

4.2.3.2 Temperature varying with θ only

This model provides a simple means of incorporating day-night temperature variations across the terminator region. It is described by four main parameters. The extent of the terminator region over which the temperature varies is fixed by an angle β , similarly to Caldas et al. (2019). The other three parameters are T_{term} , the temperature in the middle of the terminator region ($\theta = \pi/2$), T_{day} , the temperature at the dayside end of the terminator ($\theta = \pi/2 - \beta/2$), and T_{night} , the temperature at the nightside end of the terminator ($\theta = \pi/2 + \beta/2$). The temperature at intermediate values of θ is found by linearly interpolating between the two appropriate temperatures. Outside of the fixed terminator region, the temperature remains uniform at either T_{day} or T_{night} . This approach differs from the formalism described in Lacy & Burrows (2020) as it uses three different temperatures rather than just two (T_{day} and T_{night}). This allows for the temperature gradient between the terminator and the dayside to be different to the gradient between the terminator and the nightside, which has been shown to be possible from GCMs, such as the model of HD 209458b presented in Fortney et al. (2010).

4.2.3.3 Temperature varying in three dimensions

Our new temperature parameterisation generalises the 1D parametric P – T profile described in Section 4.2.3.1 to allow for temperature variations in 3D. In order to achieve this, we compute three separate P – T profiles located at three different values of θ . We calculate one

profile at $\theta = \pi/2$, the exact centre of the terminator, as well as profiles at the edges of the terminator which are closest to the day- and nightsides of the planet. We label the dayside profile $T_d(P)$, the nightside profile $T_n(P)$, and the terminator profile $T_t(P)$. The terminator edges are located at $\theta = \pi/2 \pm \beta/2$, where β is defined as in the previous section. At a given (P, θ) where θ is intermediate between the two terminator edges, the temperature is obtained by linear interpolation in θ between the two appropriate profiles:

$$T(P, \theta, \phi) = G(\phi) \begin{cases} \frac{2\theta - (\pi - \beta)}{\beta} T_t(P) \\ + \frac{\pi - 2\theta}{\beta} T_d(P), & \frac{\pi - \beta}{2} < \theta < \pi/2, \\ \frac{(\pi + \beta) - 2\theta}{\beta} T_t(P) \\ + \frac{2\theta - \pi}{\beta} T_n(P), & \pi/2 < \theta < \frac{\pi + \beta}{2}. \end{cases} \quad (4.12)$$

where $G(\phi)$ is a generic function that can be used to describe the ϕ -dependence of the temperature profile. For the purposes of this study we are mostly interested in exploring day-night temperature contrasts, and so we fix $G(\phi) = 1$, however this can easily be modified to include ϕ -dependence using a formalism similar to the patchy cloud approach described in Section 4.2.5.1.

Assuming that the P - T profile is continuous throughout, it can be described using 6 free parameters. This suggests that the above parameterisation should require a total of 18 parameters. However, temperature variations with longitude/latitude are typically negligible in the deeper atmosphere (e.g. Showman et al., 2008), and so we can assign common values of P_3 and T_3 for each of the profiles, reducing the total number of free parameters to 14. A demonstration of the efficacy of this parameterisation is presented in Section 4.3.1.

4.2.4 Inhomogeneous Chemistry

Chemical inhomogeneities are likely to persist in a wide range of planets, and will therefore be important to consider when analysing upcoming data. Of particular importance will be day-night variations in chemical abundances. These variations can occur in ultra-hot Jupiters, planets with equilibrium temperatures in excess of 2000 K on which molecules such as H_2 and H_2O can be thermally dissociated on the dayside (Lothringer et al., 2018, Parmentier et al., 2018, Venot et al., 2020). They may also be present on warm to hot Neptunes, since their atmospheric composition varies strongly as a function of temperature and bulk atmospheric properties such as metallicity and C/O ratio (Moses et al., 2013). High-metallicity

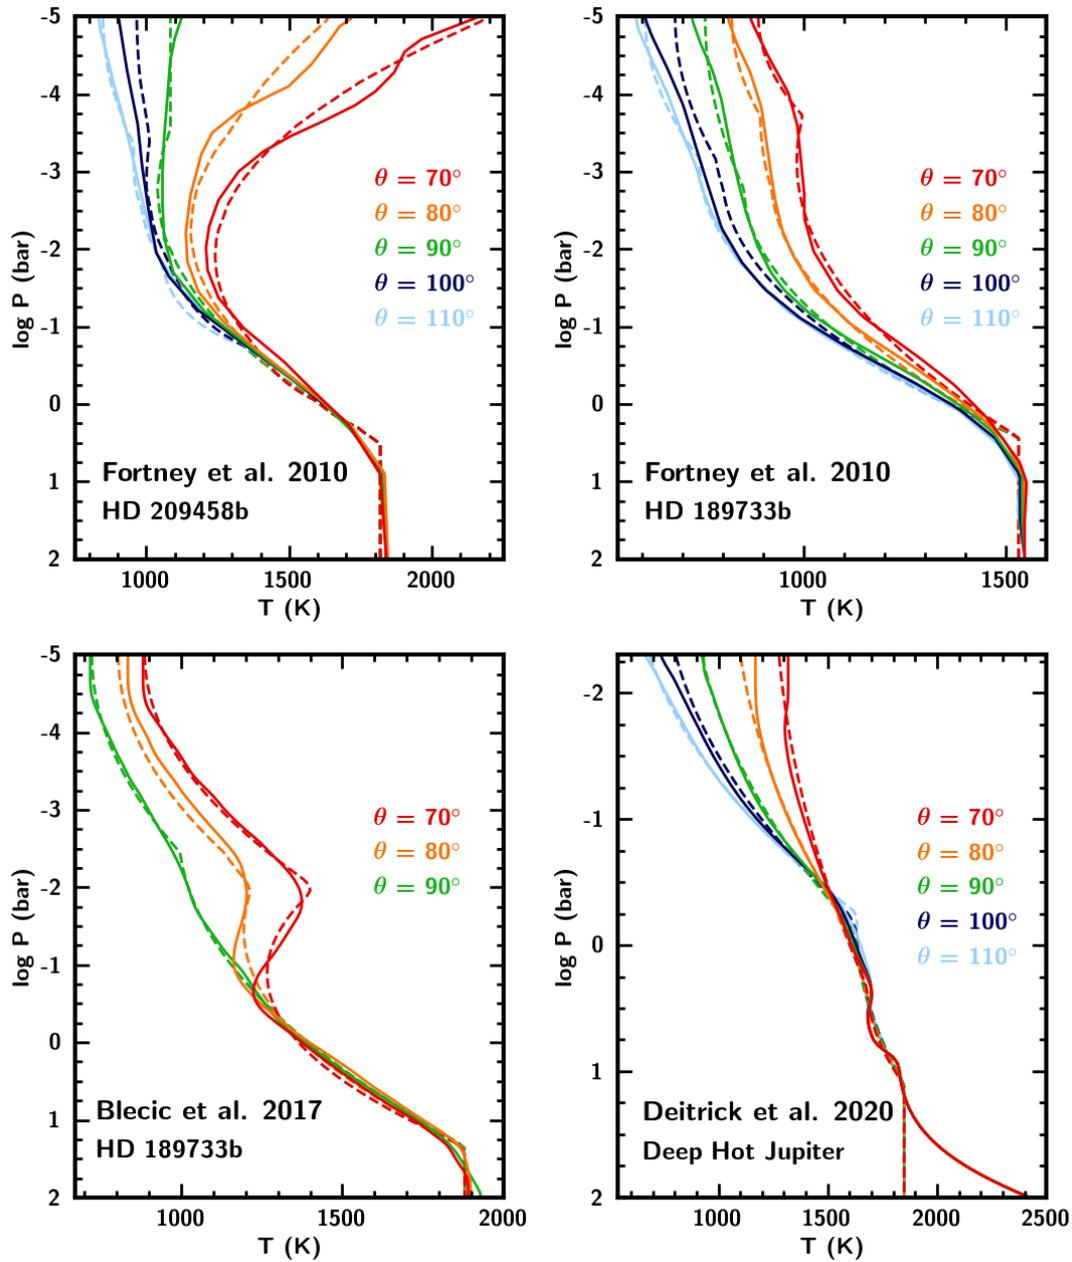


Fig. 4.4 Comparison between azimuthally-averaged P - T profiles from the terminator regions of a range of GCMs (Fortney et al., 2010, Blecic et al., 2017, Deitrick et al., 2020) and parametric profiles as described in Section 4.2.3. The GCM profiles are shown as solid lines, and parametric fits are shown as dashed lines. In each case the profiles at $\theta = 70^\circ$, 90° and 110° are fit using a modified form of the parameterisation from Madhusudhan & Seager (2009), with the parameters P_3 and T_3 shared between each profile. The profiles at $\theta = 80^\circ$ and 100° are calculated by linear interpolation between the nearest two profiles. For the Blecic et al. (2017) model, we only show profiles for $\theta \leq 90^\circ$ since nightside profiles were not presented in that paper.

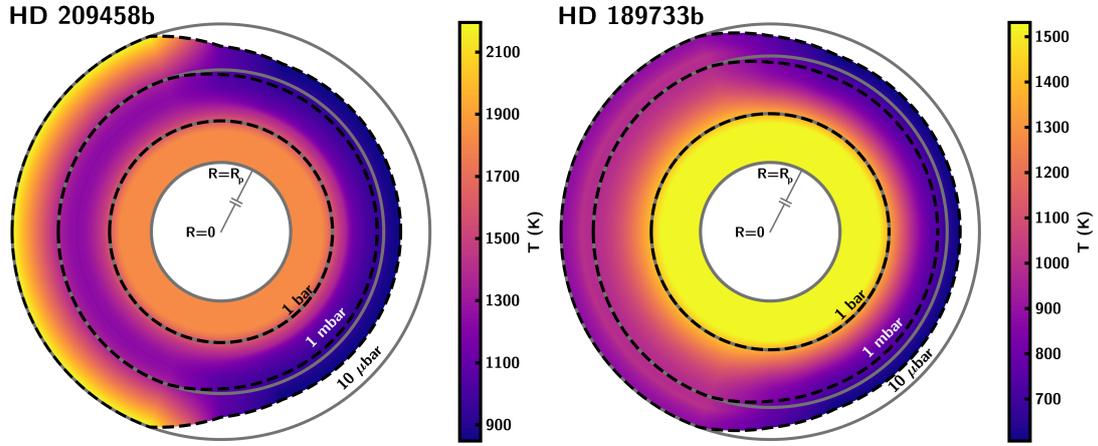


Fig. 4.5 Parametric fits to the 3D P - T profiles presented in Fortney et al. (2010) showing a slice at $\phi = 0$ (through the north-south polar plane). As in Figure 4.4, the profiles at $\theta = 70^\circ$, 90° and 110° follow our parametric prescription while the profiles all other values of θ are calculated by interpolation. For $\theta < 70^\circ$ and $\theta > 110^\circ$ the temperature remains constant, set by the temperatures at $\theta = 70^\circ$ and $\theta = 110^\circ$ respectively. Note that the inner portion of the planet (below R_p) is not to scale with the atmosphere.

atmospheres for these planets are known to exhibit strong day-night temperature contrasts (Lewis et al., 2010), which could lead to a scenario in which the atmospheric composition could vary strongly between the day- and nightsides of the planet.

Our model can incorporate chemical compositions which vary in three dimensions. This extension requires the expansion of the number density of each species n_i , as defined in equation 4.6 from a single number to a three-dimensional array $n_i(r, \theta, \phi)$. For the purpose of this study, we consider a simplified case in which two abundances are specified for each chemical species: a 'dayside' abundance ($\theta < \pi/2$) and a 'nightside' abundance ($\theta > \pi/2$). The effect of day-night chemical inhomogeneities on transmission spectra are explored in Section 4.3.3. Considerations for retrievals with inhomogeneous chemical abundances are reserved for a future study (see section 4.4.3 for further discussion).

4.2.5 Additional Physical Effects

Our model incorporates a wide range of other atmospheric properties. These are discussed in more detail in previous works (e.g. Pinhas et al., 2018, Welbanks & Madhusudhan, 2021) and are briefly described here.

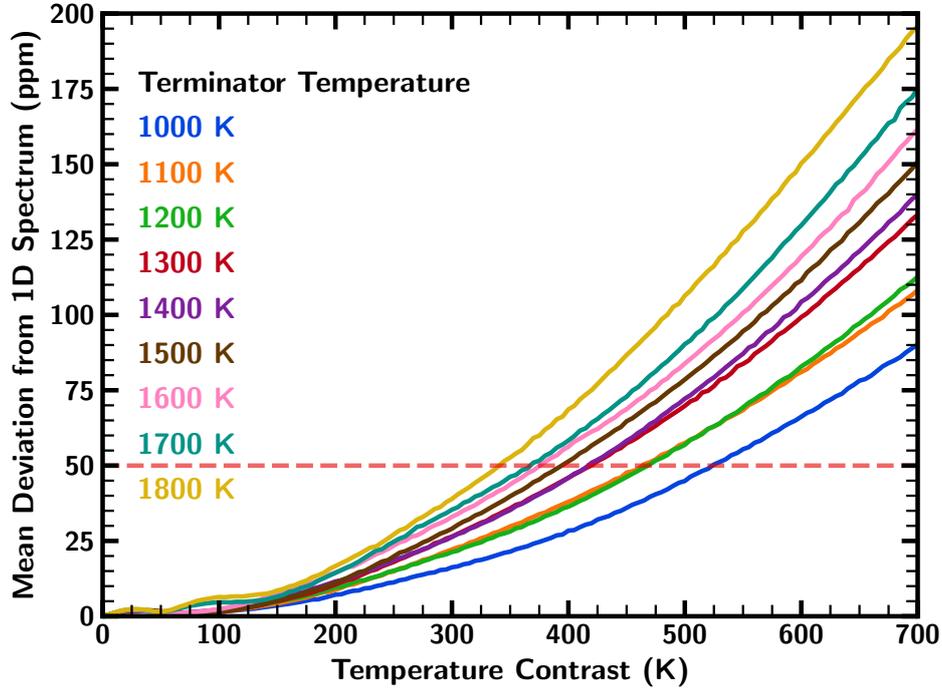


Fig. 4.6 Differences between simulated hot Jupiter spectra with a 1D thermally averaged temperature profile and spectra with different day-night temperature gradients. For a range of terminator temperatures T_t (the temperature at the top of the atmosphere at $\theta = 90^\circ$) we compute spectra with temperature contrasts $\Delta T = 0\text{--}700$ K ($\Delta T = T_d - T_t = T_t - T_n$). In each case, the spectra are calculated at $R = 5000$ and a wavelength range of $1\text{--}15\ \mu\text{m}$ and the mean difference between the two transit depths is calculated. The difference in transit depth is given in parts per million. The dashed red line at 50 ppm is an approximate indication of the expected precision that can readily be achieved by JWST observations of hot Jupiters.

4.2.5.1 Clouds and hazes

Our model includes a prescription for inhomogeneous cloud cover adapted from Welbanks & Madhusudhan (2021), which divides the atmosphere into (up to) four azimuthal slices. These slices can have one of four properties: (1) cloud/haze free, (2) covered by a grey cloud deck, (3) covered by hazes, or (4) covered by a grey cloud deck with hazes above the deck. The grey cloud deck is essentially opaque (Fortney, 2005) and is therefore implemented by setting the optical depth to infinity for all pressures larger than the cloud-top pressure P_{cloud} . Similarly to previous studies (Line & Parmentier, 2016, MacDonald & Madhusudhan, 2017, Welbanks & Madhusudhan, 2021) we adopt the haze model described in Lecavelier Des Etangs et al.

(2008), which gives a haze cross-section

$$\sigma_{\lambda,\text{haze}} = a\sigma_0 \left(\frac{\lambda}{\lambda_0} \right)^\gamma, \quad (4.13)$$

where a is the Rayleigh enhancement factor, γ is the scattering slope, and $\sigma_0 = 5.31 \times 10^{-31} \text{ m}^2$ is the cross-section due to H_2 Rayleigh scattering at a reference wavelength $\lambda_0 = 3.5 \times 10^{-7} \text{ m}$ (Dalgarno & Williams, 1962). The attenuation coefficient due to hazes, $\mu_{\lambda,\text{haze}}$ is therefore given by the following expression:

$$\mu_{\lambda,\text{haze}}(\lambda, P, T) = X_{\text{H}_2} n_{\text{tot}}(P, T) \sigma_{\lambda,\text{haze}}, \quad (4.14)$$

where X_{H_2} is the H_2 abundance and $n_{\text{tot}}(P, T)$ is the total number density.

The above prescription can easily be simplified to yield the cloud model from MacDonald & Madhusudhan (2017) as presented in Chapter 2 by considering only slices with properties (1) and (4) or the model from Line & Parmentier (2016) by considering only properties (1) and (2).

4.2.5.2 Stellar heterogeneity

In order to account for the effect of star spots and faculae on transmission spectra, we follow the approach of Pinhas et al. (2018) and incorporate the treatment of stellar heterogeneity from Rackham et al. (2017) into our atmospheric forward model. In cases where stellar heterogeneity is considered, the expression for the observed transit depth, $\Delta_{\lambda,\text{obs}}$, is

$$\Delta_{\lambda,\text{obs}} = \Delta_\lambda \mathcal{E}_{\text{het}}, \quad (4.15)$$

where Δ_λ is defined in Equation 4.3 and \mathcal{E}_{het} is a perturbative term describing stellar contamination. \mathcal{E}_{het} is defined as

$$\mathcal{E}_{\text{het}} = \left(\left(1 - \frac{\mathcal{S}_u}{\mathcal{S}_o} \right) f_{\text{het}} \right)^{-1}, \quad (4.16)$$

where \mathcal{S}_u and \mathcal{S}_o are the average spectral energy distributions of the unocculted and occulted regions of the stellar surface, and f_{het} is the areal fraction of the projected stellar disk that is covered with cool spots and/or hot faculae. The spectral components of the star are computed by interpolating from the PHOENIX grid of stellar atmospheric models (Husser et al., 2013) for stellar effective temperatures exceeding 2300 K. For stars cooler than 2300 K we instead

interpolate between the DRIFT-PHOENIX model grid (Witte et al., 2011), which enables consideration of stars with temperatures down to 1000 K.

4.2.5.3 Forward scattering and refraction

Our forward model incorporates the analytic prescriptions for forward scattering and refraction presented by Robinson et al. (2017). The optical depth is typically calculated by integrating along $d\tau_\lambda = \mu_\lambda ds$ (see equation 4.5). When forward scattering is included, the optical depth is modified to become

$$d\tau_{\text{eff}} = d\tau_\lambda(1 - f_{\text{scat}}\tilde{\omega}), \quad (4.17)$$

where f_{scat} is the correction factor for forward scattering and $\tilde{\omega}$ is the forward scattering albedo. Refraction is implemented by calculating P_{max} , the pressure at which the effect of refraction is sufficient to cause a ray of light coming from one end of the planet to bend off the far limit of the star (Robinson et al., 2017). The optical depth at pressures greater than P_{max} is set to infinity.

The effects of forward scattering and refraction on transmission spectra were investigated by Robinson et al. (2017) for Jupiter-like atmospheres and by Welbanks & Madhusudhan (2021) for mini-Neptunes. In both cases the effects were shown to be small compared to effects such as Rayleigh scattering and collision-induced absorption.

4.2.6 Synthetic Data and Statistical Inference

For this study we generate synthetic JWST observations and transmission spectra using PANDEXO (Batalha et al., 2017). For all simulated observations, we assume a noise floor of 5 ppm, and a saturation limit of 80% full well. For atmospheric retrieval, we follow the binning strategy of Pinhas et al. (2018) in order to compare the high-resolution forward model with the synthetic data.

Parameter estimation is carried out using PyMultinest (Buchner et al., 2014), a Python implementation of the Nested Sampling algorithm (Skilling, 2004, Feroz & Hobson, 2008, Feroz et al., 2009). For each parameter θ we aim to find the posterior probability distribution $p(\theta|d)$ given the data d being analysed:

$$p(\theta|d) = \frac{\mathcal{L}p(\theta)}{\mathcal{Z}}, \quad (4.18)$$

where $\mathcal{L} = p(d|\theta)$ is the likelihood, $p(\theta)$ is the prior and $\mathcal{Z} = p(d)$ is the Bayesian evidence, which is not required for parameter estimation but can be used for model comparison. We assume independently distributed Gaussian errors for each of the spectral data points, meaning

the likelihood is defined as

$$\mathcal{L} = \mathcal{L}_0 \exp\left(-\frac{\chi^2}{2}\right), \quad (4.19)$$

$$\chi^2 = \sum_i \frac{(\hat{y}_i - \bar{y}_i)^2}{\sigma_i^2}, \quad (4.20)$$

where \bar{y}_i and σ_i are the mean and standard deviation of the observed data point i , and \hat{y}_i is the value of the model data point i . The Nested Sampling algorithm samples the prior space in order to calculate the evidence \mathcal{Z} , and in the process of doing so computes the likelihood \mathcal{L} which can be used to estimate the posterior distribution.

4.2.7 Model Validation

In order to verify that our 3D forward model is accurate, we reproduce effects reported in previous works that have incorporated a 3D transmission spectrum model. We consider the set of model transmission spectra of the hot Jupiter HD 209458b presented in Figure 15 of Caldas et al. (2019). We adopt the parametric P - T profile used in that model, which consists of a two-layer vertical profile, with a constant temperature in the deep atmosphere and a variable temperature from day- to nightside in the upper atmosphere. We produce a set of models with the same temperature structures and transition angles (β) as in Caldas et al. (2019), as shown in Figure 4.3. Our models show the same dependencies on temperature and transition angle as the previous work, with higher transit depths found for lower values of β .

While we have adopted the temperature profile described in Caldas et al. (2019) for the purposes of validating our model, we note that this prescription is simplified compared to our parametric P - T profile described in Section 4.2.3.3. Unlike our parametric profile, their parameterisation is not intended to be capable of matching temperature structures from GCMs (see Section 4.3.1).

4.3 Results

In this section we demonstrate that our 3D parametric P - T profile can match azimuthally-averaged temperature structures from established GCMs. We subsequently explore the extent to which day-night temperature gradients and chemical inhomogeneity can affect resulting transmission spectra. We also perform 3D retrievals on a number of synthetic datasets, demonstrating AURA-3D's capabilities to constrain temperature profiles which vary across the terminator and to overcome biases in retrieved abundances that can arise from the assumption of a 1D temperature profile.

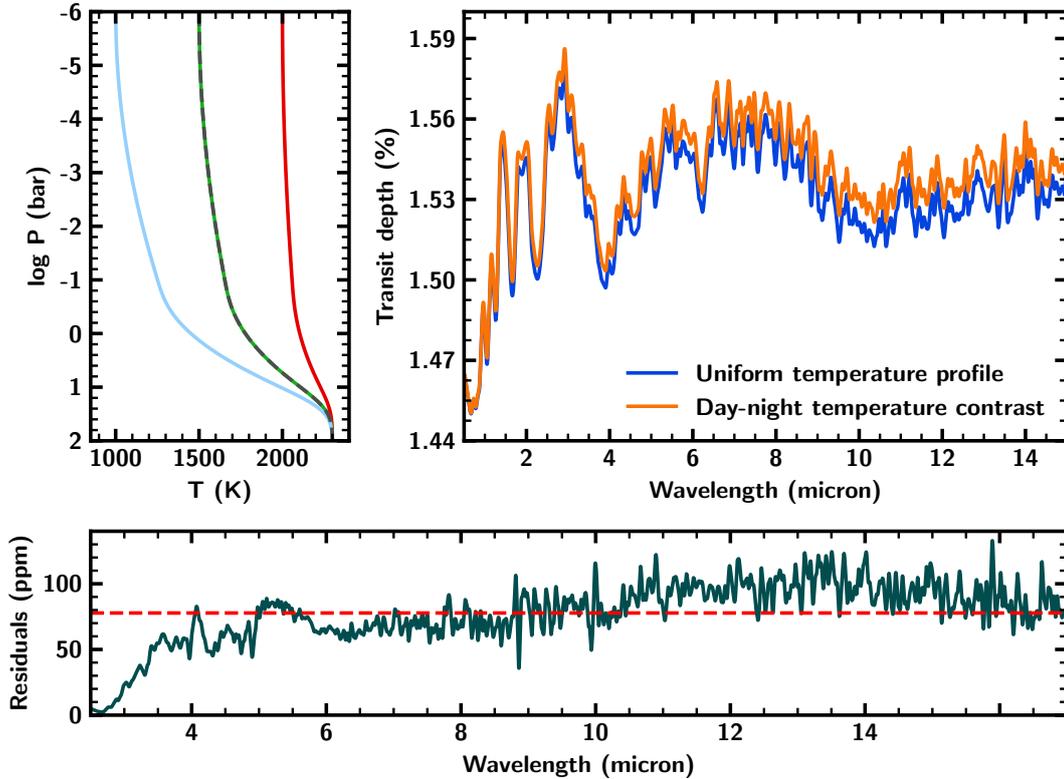


Fig. 4.7 Comparison of two theoretical transmission spectra for a HD 209458b-like planet (top right). The orange spectrum uses a parametric temperature structure with a day-night temperature contrast, shown by the blue (night), green (terminator) and red (day) temperature profiles in the top left panel, whereas the blue spectrum uses a 1D averaged structure shown by the dashed grey temperature profile. The bottom panel shows the difference between the two spectra as a function of wavelength. The spectra differ by 78 ppm on average, enough that high-quality JWST observations of this planet should require day-night temperature variations to be considered in retrievals.

4.3.1 Comparison of Parametric Profile with GCM Temperature Profiles

We demonstrate that our 14-parameter prescription for the temperature profile is able to approximate a range of azimuthally-averaged temperature structures of several GCMs. Figure 4.4 shows temperature profiles from the terminator regions of several published 3D simulations of hot Jupiter atmospheres. We consider the models of HD 209458b and HD 189733b presented in Fortney et al. (2010), the model of HD 189733b from Blečić et al. (2017), and the Deep Hot Jupiter benchmark model from Deitrick et al. (2020). For each model we show the azimuthally-averaged temperature profiles across the terminator region of the atmosphere. Since only dayside profiles are presented in Blečić et al. (2017) we only show the temperature

profiles at $\theta = 70^\circ$, 80° and 90° for this case. To fit these models we assume a value of $\beta = 40^\circ$. For each model, we therefore fit the temperature profiles at $\theta = 70^\circ$, 90° and 110° with our 14-parameter prescription, with T_3 and P_3 shared between the profiles, using a Trust Region Reflective least-squares algorithm (Sorenson, 1982). The full temperature structure is subsequently calculated by linear interpolation between the nearest of those three profiles. Figure 4.4 includes profiles at $\theta = 80^\circ$ and 100° in order to demonstrate the accuracy of the interpolated profiles. The full distribution of day-night temperature profiles for the models presented in Fortney et al. (2010) are shown in Figure 4.5.

Our parametric P - T prescription is able to fit all of the profiles shown in Figure 4.4 to a good degree of accuracy. The model of HD 209458b from (Fortney et al., 2010) is a good demonstration of why interpolation between three temperature profiles is preferred to two. In this case, the difference between the temperature profile at $\theta = 70^\circ$ and the profile at $\theta = 90^\circ$ is much greater than the difference between the profiles at $\theta = 90^\circ$ and $\theta = 110^\circ$. If the parametric model only interpolated between the profiles at $\theta = 70^\circ$ and $\theta = 110^\circ$, then the interpolated $\theta = 90^\circ$ would be much hotter than the profile from the GCM and would yield a worse fit.

For the Deep Hot Jupiter case, we did not fit our profile to the deep atmosphere (>10 bar) since this is below the region of the atmosphere probed in transmission, hence the deviation between the GCM and parametric profiles at high pressure. For atmospheric retrieval, the goal is to quickly generate a large number of model spectra which can be compared to observed data. Current retrievals require the generation of $\gtrsim 10^6$ forward models (Fortney et al., 2021). Generating a temperature structure from a GCM typically takes several days, rendering them unfeasible for retrievals. In contrast, a forward model using our parametric profile is computed in ~ 0.6 seconds on a single core. Our parameterisation therefore enables the efficient exploration of a wide range of multidimensional temperature structures.

In order to confirm that our parametric P - T profile can be used to generate transmission spectra closely matching those generated with a full 3D temperature structure, we produce a model spectrum from 0.5 – $5.5\mu\text{m}$ using the Deep Hot Jupiter temperature structure from Deitrick et al. (2020), assuming a H/He-rich atmosphere with H_2O as the sole absorbing species. We compare this to a spectrum generated using our parametric fit shown in Figure 4.4. The mean difference between the transit depths in the two spectra is 12.1 ppm, which is likely below the noise floor of most JWST instruments (Greene et al., 2016). In contrast, comparison between the full 3D model and a 1D best-fit to the globally-averaged temperature profile yields a difference of 65.5 ppm, indicating that our parameterisation provides a better match to the 3D spectrum than a 1D model. The differences between the spectra are more pronounced at longer wavelengths (see Figure 4.6).

4.3.2 Effects of a Day-night Temperature Gradient on Atmospheric Spectra

We use our parametric temperature structure to explore how a day-night temperature gradient can lead to differences in resulting transmission spectra from 1D averaged models. The benefit of using the parametric model to investigate this is that a large number of spectra can be generated quickly for comparison, without the need to run a separate GCM for each temperature structure. In order to explore this effect we generate a grid of model hot Jupiter transmission spectra. We take the planetary and stellar parameters of the hot Jupiter HD 209458b (Stassun et al., 2017), and create models with a variety of temperature structures assuming solar H₂O abundance. In the middle of the terminator region ($\theta = 90^\circ$), we vary the temperature at the top of the atmosphere in steps of 100 K from 1000 K to 1800 K, covering the range of typical hot Jupiters. We also vary the temperature contrast, ΔT , in steps of 5 K, from 0 K to 700 K. This temperature contrast is applied to both the day- and nightsides, so that $T_d = T_t + \Delta T$ and $T_n = T_t - \Delta T$. We assume $\beta = 40^\circ$, meaning the temperature varies between $\theta = 70^\circ$ and $\theta = 110^\circ$. The other parameters describing the temperature structure are held constant: the common temperature in the deepest layer of the atmosphere, set by T_3 , remains at $T_t + 800$ K in all cases. The values of $\log P_1$, $\log P_2$ and $\log P_3$ are fixed to -0.9, -1.0 and 1.4 respectively. This means that none of the temperature profiles have a thermal inversion, since $P_1 > P_2$.

We generate transmission spectra over a wavelength range of 0.5–15 μm . This covers the wavelength range of several JWST instruments, including NIRSPEC FSS & BOTS, NIRISS SOSS, NIRCAM GRISM and MIRI LRS. Each spectrum is computed at a moderately high resolution ($R = 5000$). For each combination of T_t and ΔT , we calculate the mean difference with the 1D spectrum and present our full results in Figure 4.6. An example with $T_t = 1500$ K, $\Delta T = 500$ K is shown in Figure 4.7. These tables show that for a wide range of temperature contrasts, the difference between a spectrum with a temperature gradient and one without is non-trivial. In general, the difference is more pronounced at higher temperatures and at higher temperature contrasts. For all models except $T_t = 1000$ K, the difference between the spectra is >50 ppm for a temperature contrast of 500 K or more. This level of precision would be expected for many JWST transmission spectra of hot Jupiters, depending on factors such as observing time and the brightness of the host star. However, the difference between the spectra is not uniform with wavelength, with less of a noticeable difference at short wavelengths (see Figure 4.7). This suggests that for some instruments, such as NIRISS SOSS, the temperature gradient will have less of an impact on the observed spectrum, while the effect will be more important to consider when observing with instruments such as MIRI MRS and LRS. It is also important to note that while some HST Wide Field Camera 3 (WFC3) transmission

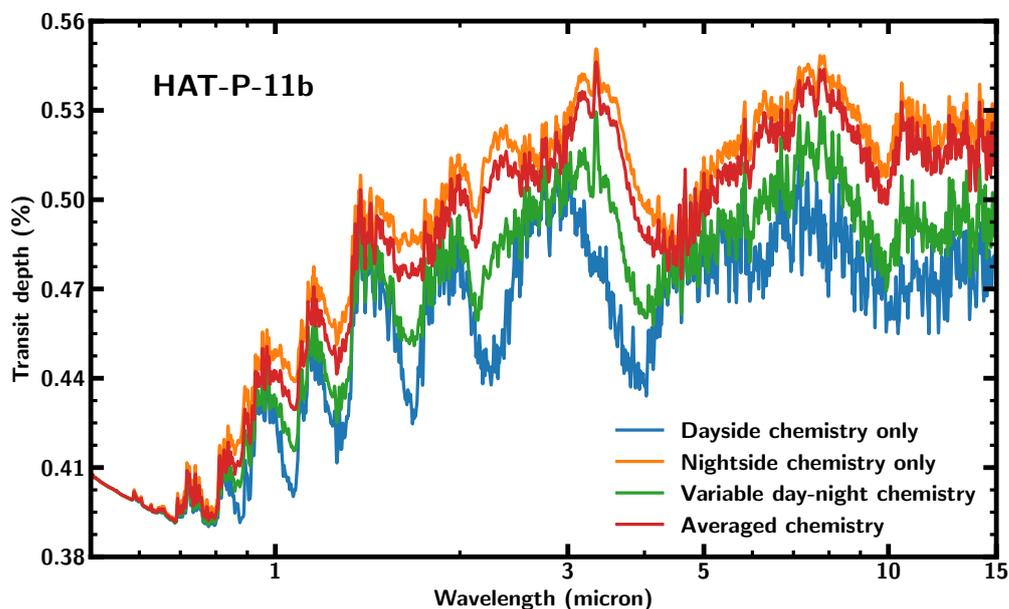


Fig. 4.8 Synthetic spectra of the hot Neptune HAT-P-11b. Four different chemical compositions are considered: (1) ‘dayside’ chemistry everywhere, (2) ‘nightside’ chemistry everywhere, (3) variable chemistry between the day- and nightsides, and (4) averaged chemistry everywhere (see Table 4.1). Model 2 generally has a higher transit depth than model 1 due to higher abundances of H_2O , CH_4 and NH_3 which are all prominent at these wavelengths. Model 3 generally lies between models 1 and 2, whereas model 4 lies closer to model 2, since the ‘nightside’ molecules are also found on the dayside of the planet, where the scale height is larger due to the higher temperatures.

spectra have also attained a level of precision of 50 ppm or better, these models indicate that WFC3 spectra will not be substantially affected by day-night temperature gradients since the difference between the spectra is small at the wavelengths probed by WFC3 (1.1–1.7 μm).

4.3.3 Effect of Chemical Inhomogeneity on Transmission Spectra

Since AURA-3D is capable of modelling atmospheres with inhomogeneous chemistry, we explore the possibility for chemical variability between the day- and nightsides of a planet to affect resulting transmission spectra. Past studies have demonstrated that thermal dissociation of molecules such as H_2O on the daysides of ultra-hot Jupiters can substantially affect transmission spectra (Lothringer et al., 2018, Parmentier et al., 2018, Pluriel et al., 2020, 2021). In this section we consider a different chemical transition, namely the transition from a CH_4 -dominated to CO -dominated atmosphere. In chemical equilibrium, the abundances of

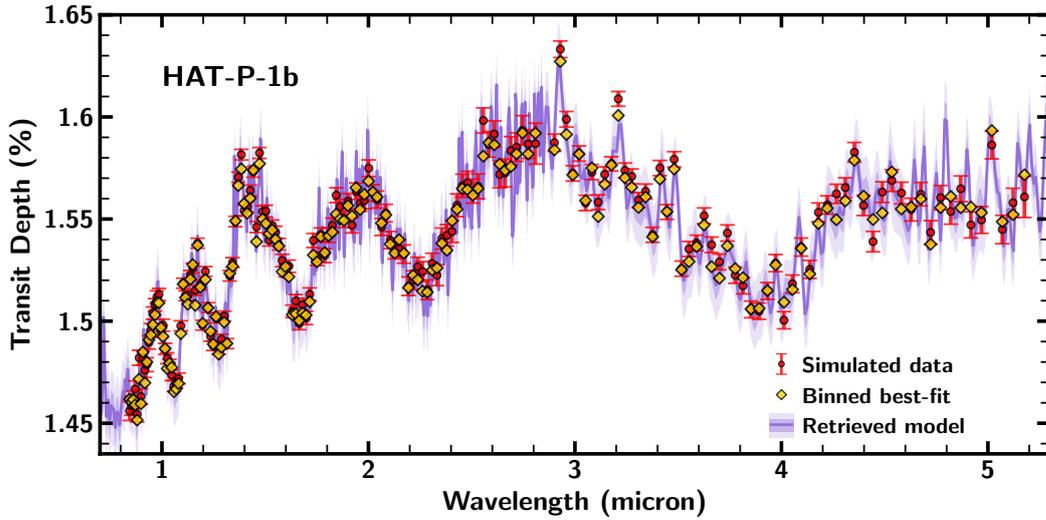


Fig. 4.9 Retrieval of a synthetic JWST spectrum of the hot Jupiter HAT-P-1b. The red points represent the simulated data, and the purple line shows the best-fit model spectrum, with 1σ and 2σ confidence intervals shown by the dark and light shaded regions. The input and retrieved models both incorporate a multidimensional parametric P - T profile. We find good agreement between the input data and the retrieved spectrum.

key atmospheric species in a H_2 -rich atmosphere can vary strongly with temperature. At temperatures below ~ 1200 K, CH_4 is expected to be a highly abundant molecule, whereas at higher temperatures CO is expected to dominate (Lodders & Fegley, 2002, Moses et al., 2011, Madhusudhan et al., 2012). If a planet's day- and nightsides are above and below this temperature respectively, this could lead to substantially different chemical abundances on either side of the planet.

In order to investigate this effect, we generate models of the hot Neptune HAT-P-11b (Bakos et al., 2010). Given this planet's equilibrium temperature of 878 K, it is possible that the day- and nightsides of the planet are above and below the transition temperature between CH_4 -dominated and CO-dominated chemistry, making the planet a useful test case to examine the effects of chemical inhomogeneity on transmission spectra. We consider two different sets of equilibrium chemical abundances: a CO-dominated regime, which we may expect to see on the dayside of the planet, and a CH_4 -dominated regime which may be expected on the nightside. The abundances for each regime are shown in Table 4.1. We choose a model P - T profile with a similar structure to the profiles used in Section 4.3.2. We take $T_t = 878$ K, the equilibrium temperature of the planet, with $\Delta T = 500$ K. The values of $\log P_1$, $\log P_2$ and $\log P_3$ are again fixed to -0.9, -1.0 and 1.4 respectively. For the chemistry we consider four different scenarios: (1) dayside abundances across the whole terminator region, (2) nightside abundances across the whole terminator, (3) inhomogeneous chemistry

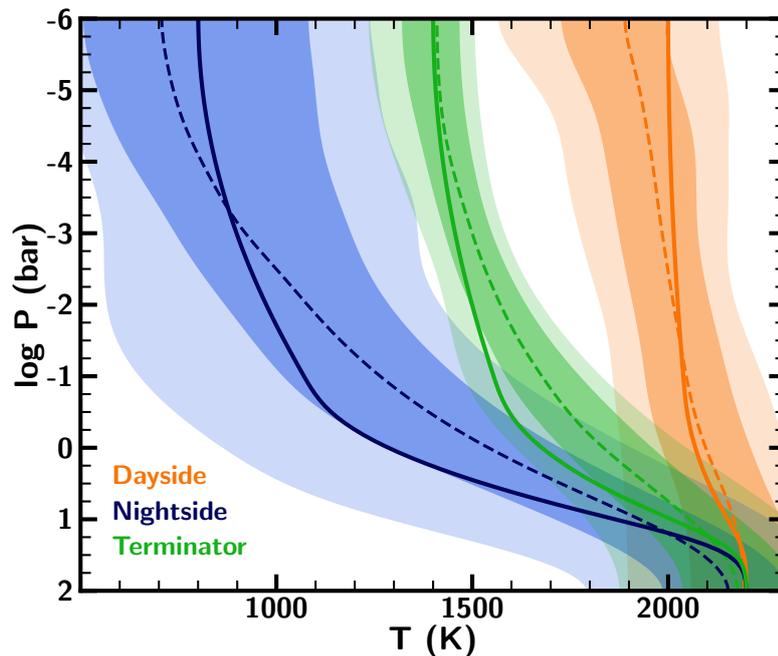


Fig. 4.10 Input and retrieved multidimensional parametric temperature profiles of a hot Jupiter. The solid lines show the input dayside, nightside, and terminator temperature profiles used to generate the synthetic data. The dashed lines correspond to the median retrieved dayside, nightside and terminator temperature profiles respectively, with shaded regions corresponding to 1σ and 2σ confidence intervals. The retrieved temperature profiles are consistent with the input profiles, with almost all of the input temperature profiles contained within the retrieved 1σ confidence intervals.

with dayside abundances for $\theta < 90^\circ$ and nightside abundances for $\theta > 90^\circ$, and (4) averaged abundances across the whole terminator (the mean of the dayside and nightside abundances). We generate models from $0.5\text{-}15\ \mu\text{m}$, which covers the wavelength ranges of several JWST instruments, including NIRSpec, NIRISS NIRCам, and MIRI LRS, and also covers many of the prominent spectral features of the molecules considered (Madhusudhan, 2019).

The resulting spectra are shown in Figure 4.8. We see that the model with nightside chemistry across the entire terminator has a higher transit depth overall than the model with dayside chemistry across the entire terminator. This is a result of the high abundances of H_2O , CH_4 and NH_3 expected in chemical equilibrium at low temperatures, each of which has prominent absorption features at the wavelengths considered here. The model with averaged chemistry also differs substantially from the model with chemical abundances varying between the day- and nightsides of the planet across the wavelength range considered. The model with averaged chemistry has larger transit depths, since this model has large

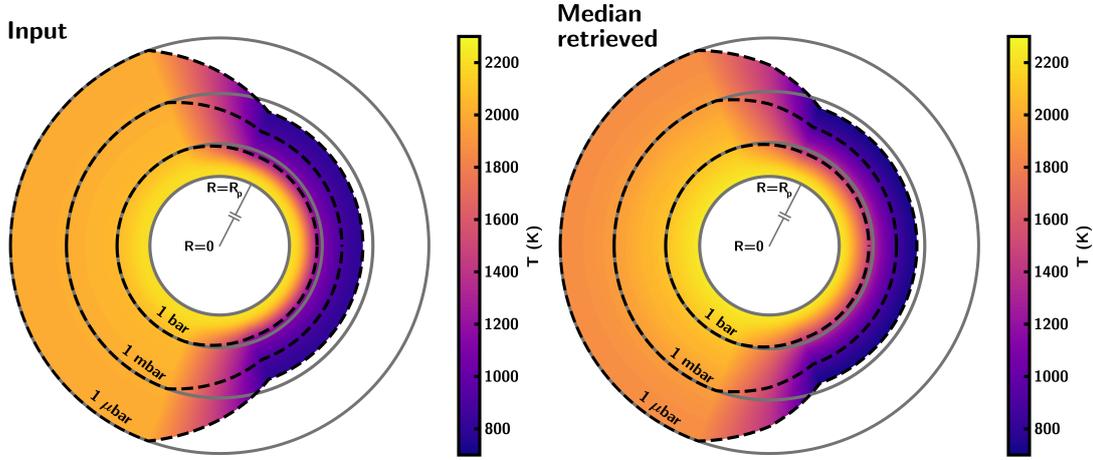


Fig. 4.11 Two-dimensional representations of the input and median retrieved parametric P - T profiles for HAT-P-1b. The temperature structure is computed by interpolating between the temperature profiles shown in Figure 4.10. Note that the inner portion of the planet (below R_p) is not to scale with the atmosphere.

amounts of H_2O , CH_4 and NH_3 on the dayside as well as the nightside, and the dayside has a larger scale height since it is at a higher temperature. In the model with variable chemistry, these molecules (particularly CH_4 and NH_3) are mostly present on the cooler side of the planet, leading to smaller features overall.

Whether such chemical inhomogeneities can be detected with JWST will need to be determined in future studies. For the remainder of this work we will focus on the possibility of constraining temperature inhomogeneities using JWST-quality observations.

4.3.4 Retrieval of a Multidimensional Temperature Structure

Here we demonstrate the capability of our new retrieval framework to recover a multidimensional temperature structure. We create a parametric temperature profile using the prescription described in Section 4.2.3.3. We use this temperature structure to produce a synthetic transmission spectrum, and carry out a retrieval of this spectrum using the same parameterisation. Our goal is to determine the extent to which the retrieval can constrain the input multi-dimensional temperature structure using JWST-quality data.

The input parameters used to generate the temperature structure are shown in the table embedded in Figure 4.13. The model interpolates between temperature profiles defined for the dayside ($\theta \leq 70^\circ$), nightside ($\theta \geq 110^\circ$) and the middle of the terminator ($\theta = 90^\circ$). We

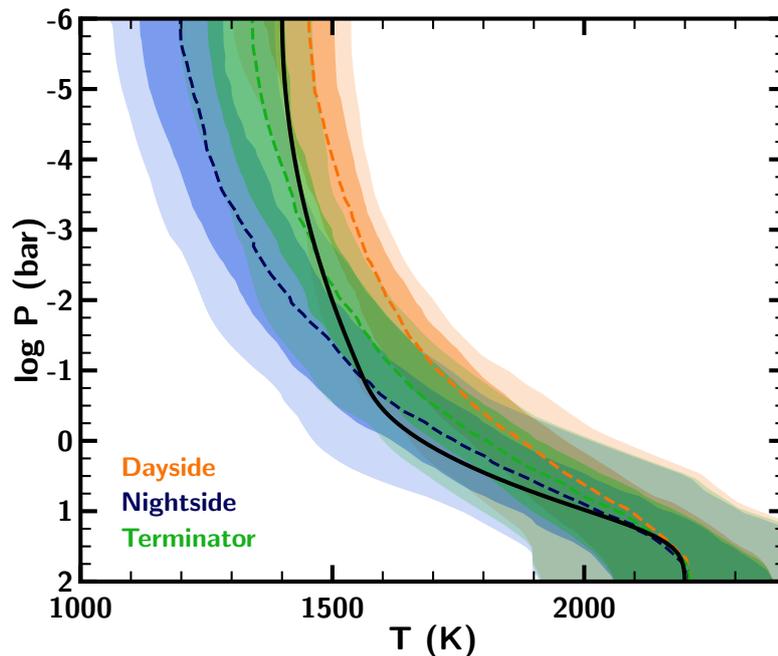


Fig. 4.12 Input and retrieved multidimensional parametric temperature profiles of a hot Jupiter with no day-night temperature contrast. The solid line shows the input temperature profile used to generate the synthetic data. The dashed lines correspond to the median retrieved dayside, nightside and terminator temperature profiles respectively, with shaded regions corresponding to 1σ and 2σ confidence intervals. The retrieved temperature profiles overlap significantly with each other and are broadly consistent with the input profile, suggesting that a multidimensional retrieval approach is not required in this case.

consider a cloud-free, solar composition atmosphere with uniform abundances of H_2O , CH_4 , NH_3 , HCN , CO and CO_2 . Input chemical abundances are also shown in the table embedded in Figure 4.13.

To generate our synthetic data set we assume planetary and stellar properties that are similar to the hot Jupiter HAT-P-1b (Bakos et al., 2010). We choose this planet as an intermediate case that is representative of the majority of hot Jupiters that may be observed with JWST. HAT-P-1b has similar bulk properties to HD 209458b, and we find day-night temperature gradients to produce similar differences in the transmission spectrum to those shown in Figure 4.7. The forward model is first computed at a moderately high resolution ($R = 3000, 0.5 - 5.5\mu\text{m}$). Synthetic observations are then generated using PANDEXO (Batalha et al., 2017), assuming a single transit observed with NIRISS SOSS and another single transit observed with NIRSpec G395H. This provides a broad wavelength coverage

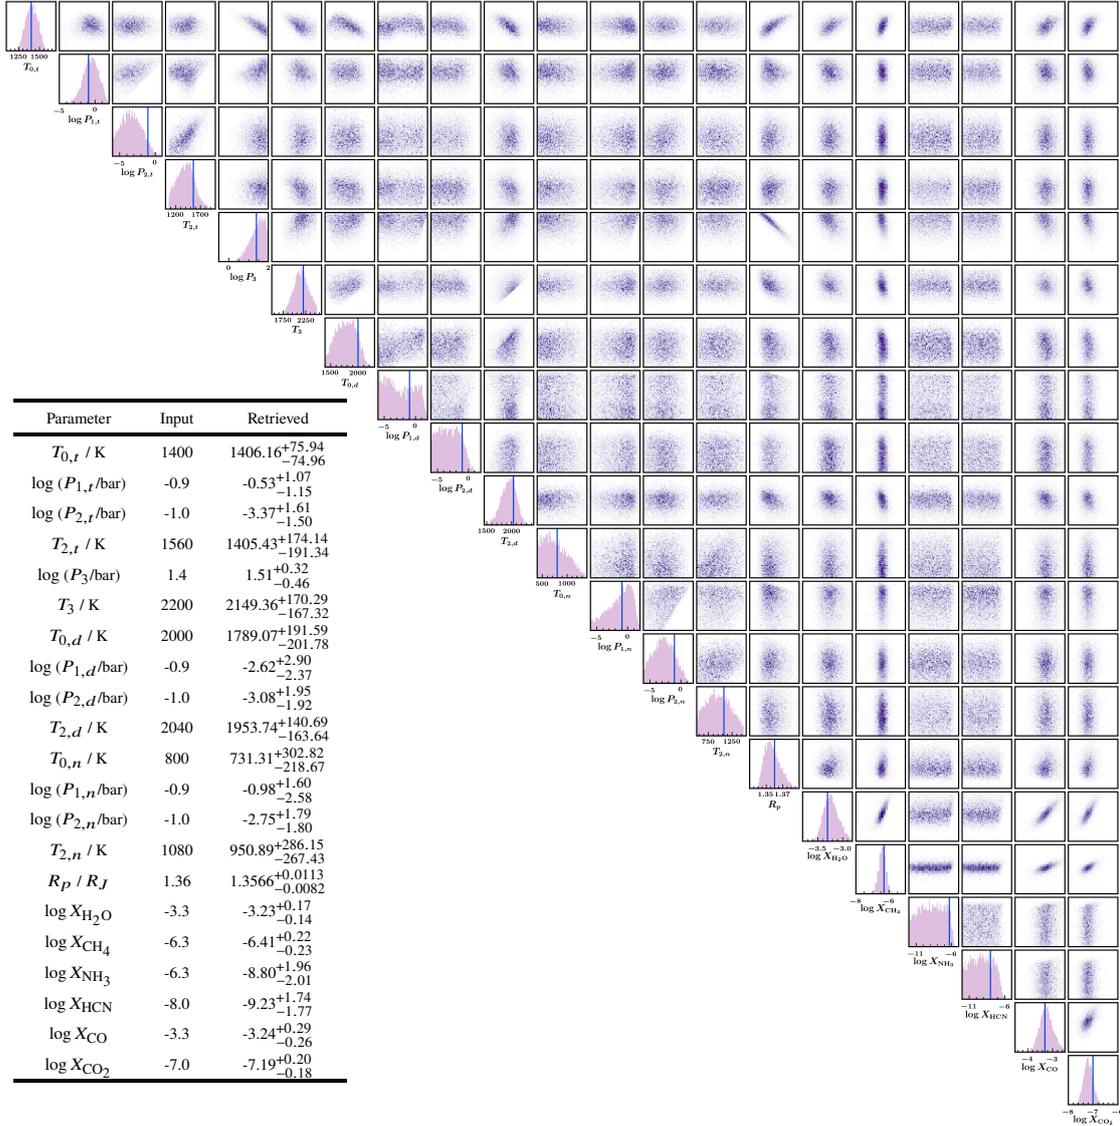


Fig. 4.13 Marginalised posterior probability distributions for the retrieval of a synthetic JWST transmission spectrum of the hot Jupiter HAT-P-1b. Blue vertical lines indicate true parameter values. *Inset*: Input and retrieved parameter values with 1σ uncertainties.

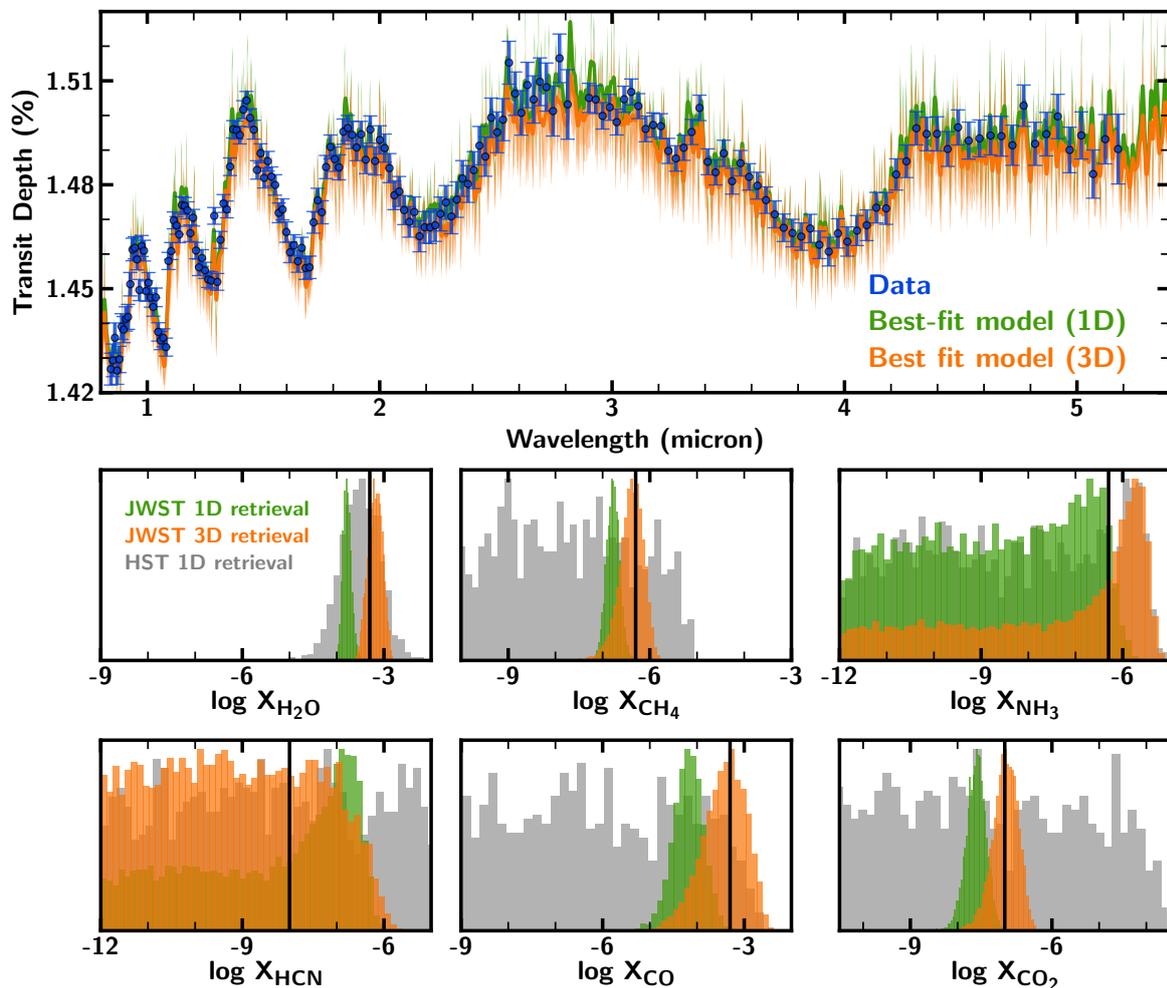


Fig. 4.14 Atmospheric retrieval of a synthetic hot Jupiter transmission spectrum generated using a 3D temperature structure. *Top*: Observations (blue) and retrieved model spectra (green and orange) for the two different model considerations. Shaded regions represent 1σ and 2σ confidence intervals. *Bottom*: Posterior distributions for the retrieved volume mixing ratios of each molecular species in the model. Input (equilibrium solar) values are shown by solid black lines. The model which does not include a day-night temperature contrast finds abundances that are not consistent with the input values. The model that does include a day-night temperature contrast is capable of inferring accurate abundances. The grey posterior distributions represent a retrieval of the HST WFC3 spectrum of the same model planet for comparison, using a 1D forward model.

Table 4.1 Chemical abundances used for the day- and nightsides of our HAT-P-11b models. The values shown are the logarithm of the volume mixing ratios.

Species	Dayside abundance	Nightside abundance
H ₂ O	-3.3	-3.0
CH ₄	-6.3	-3.3
NH ₃	-6.3	-3.6
HCN	-8.0	-20.0
CO	-3.3	-16.0
CO ₂	-7.0	-17.0

with JWST that contains multiple spectral features of the chemical species (Madhusudhan, 2019) included in the forward model. The Nested Sampling retrieval uses 2000 live points.

The input data and retrieved spectrum are shown in Figure 4.9. We find excellent agreement between the input and retrieved spectra across the full wavelength range considered. Marginalised posterior distributions for each model parameter, as well as the retrieved median values with associated 1σ uncertainties, are shown in Figure 4.13. The estimates of most free parameters are consistent with input values to within 1σ uncertainty, and all retrieved parameters are consistent with input values to within 2σ .

Our input and retrieved temperature profiles are shown in Figure 4.10. Each of the dayside, nightside and terminator input temperature profiles almost entirely remain within the retrieved 1σ confidence intervals of the retrieved profiles, and lie completely within the 2σ confidence intervals. The uncertainty estimates for the temperature profiles in the model range from ~ 100 – 300 K. We also compare 2D representations of the input and median retrieved temperature structures in Figure 4.11. We see that the temperature structures are very similar, highlighting the retrieval’s ability to recover a detailed multidimensional temperature profile using high-quality data.

The retrieved abundances of all chemical species are shown in Figure 4.13. The abundances of H₂O, CH₄, CO and CO₂ are all constrained to within 0.5 dex and are consistent with input values to 1σ . We obtain upper limits on the abundances of NH₃ and HCN which are also consistent with input values, noting that these species have low abundances in the input model ($\log X_{\text{NH}_3} = -6.3$, $\log X_{\text{HCN}} = -8.0$).

This result demonstrates that using JWST observations of a nominal hot Jupiter, it is possible to obtain accurate constraints on a multidimensional temperature profile across the

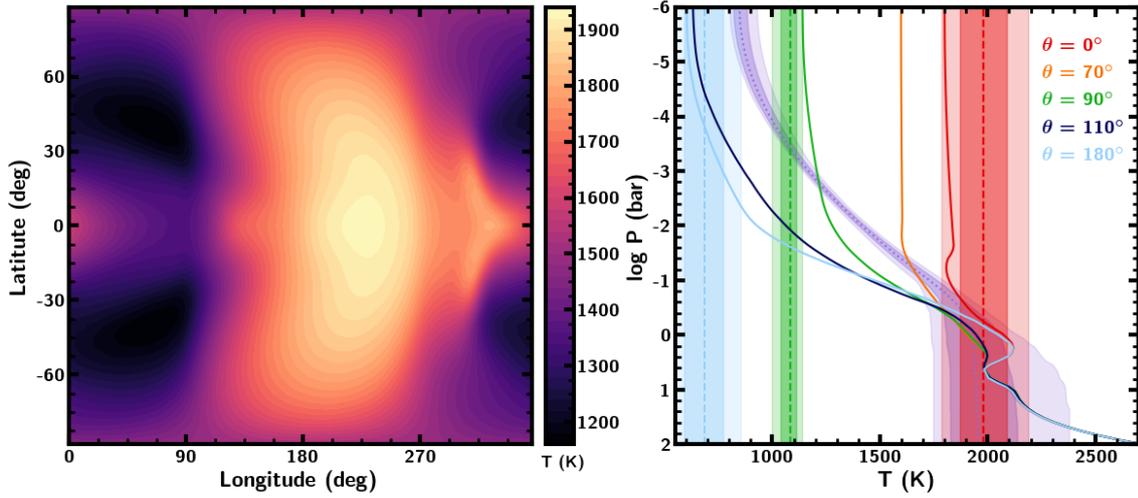


Fig. 4.15 *Left*: Temperature structure at a pressure of 0.1 bar of the forward model used to generate the synthetic spectrum. *Right*: Azimuthally-averaged temperature profiles from the forward model (solid lines), along with retrieved temperature profiles from the two retrievals. The retrieved 1D P - T profile is shown as a purple dotted line, with shaded 1σ and 2σ confidence intervals. The retrieved terminator (green), dayside (red) and nightside (light blue) temperatures from the retrieval incorporating day-night temperature variations are shown as dashed lines, also with shaded 1σ and 2σ confidence intervals.

terminator region with reasonable precision. Furthermore, these observations can be used to obtain very precise abundance estimates for a number of chemical species.

As an additional validation test for our retrieval pipeline, we also conduct a retrieval of a synthetic spectrum for a planet with no day-night temperature gradient (i.e. a 1D P - T profile) using our 3D framework. The retrieved dayside, nightside and terminator temperature profiles are shown in Figure 4.12. We find that there is considerable overlap between each of the retrieved profiles, as would be expected for a planet with no temperature contrast. In this case, the lack of strong distinction between the three profiles suggests that this spectrum is better suited to a 1D retrieval (see Section 4.4.1).

4.3.5 Retrieval of a Synthetic Spectrum from a GCM

In this section we present retrievals of a simulated JWST transmission spectrum of a hot Jupiter generated with a 3D temperature structure. We conduct retrievals using two different modelling paradigms: one using a 1D parametric P - T profile, and one using a temperature profile that is allowed to vary between the day- and nightsides of the planet.

The temperature structure of our input forward model is adapted from the open-source GCM THOR (Mendonça et al., 2016, Deitrick et al., 2020). We use THOR to reproduce

the "Deep Hot Jupiter" scenario described in Deitrick et al. (2020). In order to create a temperature structure suitable for our forward model, we subsequently extrapolate the output temperature profiles from the GCM to lower pressures. In order to do this, we fit a parabolic P - T curve to each temperature profile in the longitude/latitude grid of outputs. The resulting 3D P - T structure is shown in Figure 4.15.

Having adapted the GCM temperature structure to be appropriate for our retrieval case study, we now use it to create a model JWST transmission spectrum. Our forward model in this instance is cloud-free and assumes chemical abundances of a $1\times$ solar composition H_2 -dominated atmosphere (see Gandhi & Madhusudhan, 2017), including H_2O , CH_4 , NH_3 , HCN , CO and CO_2 . We assume the same planetary and stellar properties for the planet HAT-P-1b as in Section 4.3.4.

We carry out two retrievals of our simulated JWST spectrum. The first such retrieval uses the 1D parametric P - T prescription as described in Madhusudhan & Seager (2009), while the second uses a temperature profile that varies across the terminator, but not with height (see Section 4.2.3.2). We fix β to 40° , which is representative of the extent of the terminator region which affects a transmission spectrum (Fortney et al., 2010). We do not allow cases in which the dayside temperature is colder than the terminator temperature, or where the nightside temperature is hotter than the terminator.

A comparison of the two retrieval results is shown in Figure 4.14. We find that the retrieval assuming a 1D parametric P - T profile leads to some inaccuracies in the retrieved abundances using JWST-quality data. None of the retrieved abundances found using the 1D model are consistent with the input values to within 1σ . The only abundance retrieved accurately to $\leq 2\sigma$ is that of HCN , while the input abundance of NH_3 is within 3σ of the retrieved posterior. Note that these abundances are also the least well-constrained of the species considered. All other abundances fall more than 3σ away from the retrieved posteriors.

This finding is similar to that of Caldas et al. (2019), who found that when applying a 1D retrieval framework to a high-quality spectrum generated with a 3D input model, the true abundance would often be outside the error bars of the retrieval. They also find that the magnitude and direction of the error depends on the specific case; in this particular example, the retrieved abundances of most species are lower than the true values.

In contrast, for the retrieval in which the temperature is allowed to vary across the terminator, each of the retrieved chemical abundances is consistent within 1σ of the true values. Furthermore, the uncertainties on the parameter estimates in this case are generally larger than in the 1D retrieval. This suggests that the additional flexibility introduced by incorporation of a day-night temperature gradient allows the model to explore a wider range of models that can explain the data at hand. This demonstrates that our prescription is able

to mitigate against biases that can arise when only a 1D temperature profile is allowed in a retrieval of JWST-quality data.

We also carry out a retrieval of a simulated HST WFC3 spectrum of the same planet for comparison, using a 1D atmospheric model. In this case, the only well-constrained abundance is that of H_2O , with a precision of 0.8 dex. The retrieved abundance of H_2O is consistent with the input value, albeit with a larger uncertainty. This suggests that the constraints obtained from retrievals of HST data are not sufficiently precise to necessitate a multidimensional retrieval algorithm.

The retrieved temperature profiles acquired from each retrieval are shown on the right-hand panel of Figure 4.15. The retrieved 1D temperature profile exhibits a strong temperature gradient with pressure that is not found at the terminator of the input model. For the model including a day-night temperature gradient, the retrieved temperatures are broadly consistent with the dayside, nightside and terminator temperatures in the upper atmosphere. We note that both of these prescriptions are heavily simplified in comparison to a full 3D temperature structure. Therefore, it is difficult to gain complete information about the temperature structure of the planet's atmosphere from these retrievals alone, beyond the fact that a strong day-night temperature gradient appears to exist. However, this example demonstrates importance of having the capability to include a day-night temperature gradient in retrievals in order to ensure that we can obtain accurate abundance estimates when working with high-quality spectra of hot Jupiters.

4.4 Summary and Discussion

We introduce AURA-3D, a 3D modeling and retrieval framework for transmission spectra of exoplanetary atmospheres. The high data quality expected from upcoming JWST observations of exoplanet transmission spectra motivated the development of this new framework. Our new forward model enables calculation of transmission spectra with any 3D temperature structure, including GCM outputs, as well as computation of model spectra using parametric temperature profiles that include day-night inhomogeneities. Here we summarize the key functionalities of AURA-3D retrieval framework.

- The framework includes a 3D parametric temperature profile for use in atmospheric retrieval that can fit a wide range of GCM temperature structures. This parameterisation enables the computation of transmission spectrum models with highly flexible, physically realistic temperature structures incorporating thermal inhomogeneities.

Table 4.2 Description of priors for retrievals of simulated JWST transmission spectra of HAT-P-1b generated using a GCM. The temperature profile for the 1D retrieval uses the first six parameters ($T_0, \alpha_1, \alpha_2, P_1, P_2, P_3$) while the retrieval with a day-night temperature gradient uses the parameters $T_0, T_{0,d}, T_{0,n}$.

Parameter	Lower Bound	Upper Bound	Prior
T_0 (K)	600	2200	uniform
α_1	0.02	1.0	uniform
α_2	0.02	1.0	uniform
P_1 (bar)	10^{-6}	10^2	log-uniform
P_2 (bar)	10^{-6}	10^2	log-uniform
P_3 (bar)	10^{-2}	10^2	log-uniform
$T_{0,d}$ (K)	1200	2600	uniform
$T_{0,n}$ (K)	400	1800	uniform
X_i	10^{-12}	10^{-2}	log-uniform

Table 4.3 Input and retrieved abundances for the simulated JWST transmission spectrum of HAT-P-1b. The abundances are shown as the logarithm of the volume mixing ratios.

Species	True value	No T gradient	T gradient
H ₂ O	-3.3	$-3.78^{+0.09}_{-0.08}$	$-3.18^{+0.15}_{-0.16}$
CH ₄	-6.3	$-6.78^{+0.13}_{-0.14}$	$-6.38^{+0.23}_{-0.28}$
NH ₃	-6.3	$-8.72^{+1.89}_{-2.16}$	$-6.95^{+1.24}_{-3.33}$
HCN	-8.0	$-8.10^{+1.34}_{-2.57}$	$-9.21^{+1.93}_{-1.86}$
CO	-3.3	$-4.18^{+0.30}_{-0.32}$	$-3.39^{+0.43}_{-0.50}$
CO ₂	-7.0	$-7.60^{+0.20}_{-0.22}$	$-6.98^{+0.25}_{-0.28}$

- The radiative transfer in 3D geometry is computationally efficient for retrievals, with computation time for a single model spectrum in 3D geometry $\lesssim 1$ s. We have demonstrated retrievals using a 3D geometry requiring over 10^6 model evaluations. This is easily scalable to higher order depending on the requirement.
- When applied to synthetic JWST data of hot Jupiters, AURA-3D is able to accurately constrain separate dayside, nightside and terminator temperature profiles probed in transmission geometry, yielding meaningful information on day-night temperature contrasts.
- AURA-3D can be used to accurately retrieve chemical abundances from transmission spectra whose temperature structures are calculated with 3D GCMs, even with JWST-quality data of hot Jupiters where 1D retrievals may lead to biased estimates.

We use AURA-3D to investigate constraints that can be placed on hot giant exoplanets using JWST-quality spectra, finding the following results:

- For hot Jupiters with photospheric terminator temperatures $\gtrsim 1100$ K and temperature contrasts $\gtrsim 500$ K, the day-night temperature gradient can cause differences in transmission spectra that could be detectable with JWST.
- We demonstrate that it is possible to constrain a multidimensional temperature profile across the day-night terminator of a nominal hot Jupiter to a precision of ~ 100 - 300 K using nominal JWST-quality data.
- We demonstrate that for JWST spectra of some hot Jupiters, a 1D retrieval can lead to biased abundance estimates, in agreement with previous studies. However, we find that in the case presented in this work, the introduction of a multidimensional temperature profile can overcome these biases and retrieve accurate abundances.
- For atmospheres where the terminator may transition between CO/CH₄ dominated composition, models with inhomogeneous composition may cause an observable difference to transmission spectra compared to those with globally-averaged compositions.

4.4.1 Applicability of 1D vs 3D Retrievals

The developments presented in this work allow for retrievals with a much more complex forward model than has previously been possible. Our new framework will be extremely useful for analysing JWST-quality spectra of hot Jupiters with inhomogeneous terminators. While it is tempting to apply the most complex forward model available to all retrievals,

care needs to be taken in ascertaining whether a 3D framework is necessary to explain a given observed spectrum. Furthermore, simply starting from the most complex possible model may not be computationally feasible. Although we have shown that retrievals with a 3D geometry can be carried out with high computational efficiency, the combination of 3D modelling alongside inhomogeneous clouds/hazes, stellar heterogeneity, large numbers of chemical species and other possible features such as inhomogeneous chemistry would lead to incredibly computationally expensive retrievals with an extremely high number of free parameters. It is therefore prudent to assess a priori what degree of complexity is required for a given spectrum, finding an optimal parameter set for analysing the data at hand.

While we have demonstrated that thermal inhomogeneities will lead to substantial deviations between 1D and 3D models for hot giant planets with large day-night temperature contrasts, we also find that for cooler planets the effect of thermal inhomogeneities is smaller and, therefore, may not be detectable even with JWST-quality data. When analysing future observations, several properties of the planet should be taken into account when considering the best P - T profile prescription, including the planetary radius, gravity and equilibrium temperature, all of which affect the amplitude of spectral features and therefore the degree to which effects such as day-night temperature gradients will be detectable. The extent of the terminator region being probed in transmission is also important to consider. Hotter planets tend to have larger opening angles, allowing for a greater variation in atmospheric conditions across the region probed in transmission (Wardenier et al., 2022). When considering retrievals of real observed data it will be important to determine whether the opening angle is sufficiently large to justify the inclusion of a 3D model.

As well as the planet being targeted, it is important to consider the quality of the observed data when determining the appropriate retrieval paradigm. This study has largely focused on high-quality JWST observations with a broad infrared wavelength coverage. We have found that deviations between 1D and 3D models are less significant when considering HST-quality data. Furthermore, other studies have demonstrated that the assumption of a 1D temperature profile will generally not lead to biases when analysing HST transmission spectra (Welbanks & Madhusudhan, 2022) of most hot Jupiters. For future observations, the optimal retrieval paradigm will need to be chosen taking into account the wavelength range being probed as well as the resolution and precision of the spectrum.

4.4.2 Complementary Observations using Emission Spectra and Phase Curves

In an ideal scenario, in order to obtain complete information about the structure and properties of an exoplanet's atmosphere, we would acquire its transmission and emission spectra as well as its phase curve. This combination of observations would allow for detailed characterisation of a planet's terminator region through transmission spectroscopy and of its dayside through emission spectroscopy, as well as providing insight into spatially-resolved properties through phase curve analysis. However, undertaking such an extensive observing campaign for a single planet would require very large amounts of telescope time, and if the planet does not exhibit strong spatial inhomogeneities, then time-consuming phase curve observations are unlikely to yield additional insights into the nature of its atmosphere.

The results from this work indicate that it should be possible to detect thermal inhomogeneities in certain exoplanet atmospheres using only transmission spectroscopy. We therefore suggest that an optimal observing strategy would be to first observe a planet's transmission spectrum to look for hints of day-night variability, possibly complementing this observation with an emission spectrum to constrain dayside properties and break certain degeneracies (see Section 4.4.3). Using these observations it will be possible to ascertain the degree to which spatial variations in temperature and chemistry permeate the planet's atmospheric structure, thus determining whether observations of the full phase-curve will lead to further meaningful results.

4.4.3 Future Developments

The retrieval framework presented here enables 3D atmospheric characterisation using transmission spectra, a crucial step towards robust analysis of JWST data. However, there are additional atmospheric properties that will be important to consider when interpreting upcoming observations. For example, while inhomogeneous cloud/haze coverage has been considered in past works (e.g. Line & Parmentier, 2016, MacDonald & Madhusudhan, 2017, Welbanks & Madhusudhan, 2021), the combined effect of patchy clouds/hazes and inhomogeneous temperature profiles is yet to be fully explored in a retrieval context.

The parametric P - T profile incorporated into AURA-3D has been sufficiently complex to analyse the synthetic JWST spectra considered in this study. AURA-3D has the capability to model thermal variations with both θ and ϕ , however only variations with θ have been explored to carry out the retrievals presented in this work. However, for certain targets it is possible that an even higher signal-to-noise will be achieved, either due to observations of atmospheres with a very large scale height or via observing multiple transits of a given

target. In these cases it is possible that more complexity, including a ϕ -dependent temperature profile, will be required to accurately retrieve certain parameters. Our current framework has a built-in ϕ -dependence which has already been used to model patchy clouds (Pinhas et al., 2019), and has the functionality to include a parametric temperature profile with variations in both ϕ and θ .

It is also important to assess the impact of inhomogeneous chemical abundances on retrievals of transmission spectra. This topic has been explored for ultra-hot Jupiters (Pluriel et al., 2021) but may also be relevant for cooler planets, as discussed in this work. While AURA-3D is capable of modelling planets with different day- and nightside chemical compositions, further work is needed to ascertain the feasibility of retrieving inhomogeneous chemical abundances directly from a transmission spectrum. A retrieval allowing for independent dayside and nightside chemical abundances for each species would be subject to very strong degeneracies, since low dayside abundances could be compensated for by high nightside abundances and vice-versa. It would therefore be important to limit the allowed parameter space to physically realistic dayside and nightside abundances motivated by detailed self-consistent models. As discussed above, combining transmission spectra with observations of other regions of a planet's atmosphere, such as emission spectra, could also be useful in breaking this degeneracy.

Ultimately, we should strive to incorporate the most sophisticated forward models possible into retrieval algorithms, so that they can be invoked if necessary to explain observations. Any additional increases in model complexity must be accompanied by concurrent efforts to make these models efficient for retrievals, as we have done in the present work. Efforts are underway to improve the computational efficiency of forward models and retrievals through means such as GPU implementations (e.g. Malik et al., 2017) and machine learning (e.g. Márquez-Neila et al., 2018, Nixon & Madhusudhan, 2020). These avenues for improvement may enable further steps towards reducing other simplifying assumptions that still remain in place in current retrieval frameworks.

The developments presented in this work represent an important step forwards in our ability to identify and constrain multidimensional effects in exoplanet atmospheres. This opens the field to comprehensive atmospheric characterisation of exoplanets with large spectral features that will be accessible with transmission spectroscopy in the JWST era. Our retrieval framework will be an important tool as we work to uncover exciting new insights into the nature of exoplanet atmospheres using upcoming observations of unprecedented detail.

Chapter 5

Exploring the Phase Structure of Water-Rich Sub-Neptunes

“How much do I love you? I’ll tell you no lie. How deep is the ocean? How high is the sky?”
– Irving Berlin

5.1 Introduction

One of the most intriguing findings of recent planetary detection missions is the ubiquity of sub-Neptune planets, whose radii lie between $1\text{--}4R_{\oplus}$, i.e., larger than Earth but smaller than Neptune (Batalha et al., 2013, Fulton et al., 2017). These planets are often grouped into two categories, super-Earths and mini-Neptunes, distinguished by whether their radii are greater than or less than $\sim 1.75R_{\oplus}$ (Fulton & Petigura, 2018). With no solar system analogues, the characterisation of such planets is an important challenge in exoplanet science. In particular, understanding whether these planets are typically large rocky planets, scaled-down ice giants, or some intermediate between the two is crucial in order to place these planets in the context of the broader exoplanet population. Planets in this regime are likely to contain a substantial amount of H/He and/or H₂O (Rogers, 2015, Zeng et al., 2019), and so detailed forward modelling of planetary internal structures including these components is important for their characterisation.

Recently, Madhusudhan et al. (2020) conducted a joint analysis of the interior and atmosphere of K2-18b, a habitable-zone mini-Neptune (Foreman-Mackey et al., 2015, Cloutier et al., 2017, Benneke et al., 2019, Tsiraras et al., 2019). One of the key findings of that study was that for certain solutions, K2-18b could host a liquid water ocean beneath its hydrogen-rich atmosphere. This result provides motivation for a detailed investigation of

the thermodynamic conditions and phase structures of the H₂O layers of super-Earths and mini-Neptunes in general. In this paper, we aim to characterise the interiors of planets with a significant H₂O mass fraction, both with and without H/He envelopes. We use detailed internal structure models to explore the range of phases of H₂O that are accessible within planetary interiors. This allows us to address a number of topics including the possible depths of liquid water oceans on such planets, and the phases of water that may be found at the interface between the H₂O and H/He layers of an exoplanet.

As discussed in Section 1.3.1, internal structure modelling has long been used to link a planet's composition to its observable bulk properties such as mass, radius and equilibrium temperature. In this paper we focus on planets which contain a substantial H₂O component, and which may also possess an extended H/He envelope. Since we are interested in the phase structure of the water components of these planets, it is important to accurately model the effect of temperature variations within the H₂O layer. A number of previous studies have considered the importance of thermal effects when modelling such planets. Early studies suggested that temperature variations did not significantly alter the M – R relation for water worlds (Sotin et al., 2007, Grasset et al., 2009), however these works were generally limited to low-temperature planets with either liquid or icy surfaces. More recent work has suggested that thermal effects can have a sizeable impact on the M – R relation. Madhusudhan & Redfield (2015) showed that for highly irradiated H₂O-rich planets, the atmospheres can contribute significantly to the observed radii. Thomas & Madhusudhan (2016) explored thermal effects on the M – R relation for water worlds in more detail, focusing on planets with surfaces in liquid or supercritical phases, and finding that higher surface temperatures could lead to a large increase in radius. For example, changing the surface temperature from 300–1000 K could increase the radius of a 1–10 M_{\oplus} planet by up to 25%. Otegi et al. (2020) also demonstrated that changes in a planet's temperature profile can substantially alter the M – R relation for sub-Neptunes. Mousis et al. (2020) modelled planets with steam atmospheres and supercritical H₂O layers, which can inflate the radii of water worlds without invoking a H/He envelope, while Turbet et al. (2020) also found that irradiated planets could possess inflated H₂O layers, applying their results to the TRAPPIST-1 system.

Another important aspect is the possible extent of oceans on H₂O-rich planets. A number of past studies have explored this to some extent. For example, Léger et al. (2004) calculated the depths of oceans on planets with fixed mass (6 M_{\oplus}) and an adiabatic H₂O layer with fixed mass fraction (50%) across several different surface temperatures, reporting depths of 60 km for $T_0 = 273$ K, 72 km for $T_0 = 280$ K and 133 km for $T_0 = 303$ K. Similarly, planets with 50% H₂O by mass were considered by Sotin et al. (2007), however in this case they used an isothermal temperature profile and the surface temperature was fixed to 300 K while the planet

mass was varied. They found that a $1M_{\oplus}$ planet should have an ocean that is 150 km deep, decreasing to 50 km for a $10M_{\oplus}$ planet. The decrease in ocean depth with increasing mass is a result of the higher surface gravity of more massive planets with the same composition. Alibert (2014) considered the limiting case where a planet has the maximum amount of H_2O possible while avoiding a high-pressure ice layer. In this scenario, the total H_2O mass remains approximately constant (at $\sim 0.03M_{\oplus}$) as the planet mass changes. A study by Noack et al. (2016) noted that the maximum possible ocean depth for a given planet varies with its mass, composition and surface temperature. At 300 K their results agree with Sotin et al. (2007), but they found that between 290 and 370 K, a 10 K increase in surface temperature leads to a 14–16% increase in ocean depth.

More broadly, understanding a planet’s phase structure can provide insight into its bulk geophysical properties, as has been demonstrated for icy moons in the Solar System (e.g. Hsu et al., 2015, Soderlund et al., 2020). The presence or absence of a liquid water layer has important consequences for planetary habitability (Lammer et al., 2009). Other works have explored the general phase structures of H_2O -rich planets. For example, Zeng & Sasselov (2014) studied the temperature evolution of the interiors of water-rich planets. They found that the phase structures of these planets may change as the planet cools, and that planets older than ~ 3 Gyr should have mostly solid H_2O layers, assuming they are not highly irradiated.

As well as H_2O -rich planets without an extended atmosphere, we are also interested in the interiors of planets smaller than Neptune which retain an extended H/He envelope. Only a small amount of H/He is required to have a large impact on planetary radius (Lopez & Fortney, 2014). Rogers et al. (2011) showed that increasing the equilibrium temperature of such planets also inflates their radii significantly, with the effect being more pronounced for lower-mass planets. The extent of such an envelope as well as its temperature structure determines the thermodynamic conditions at the boundary between the envelope and the remaining interior. This is a crucial factor in determining the internal structure and surface conditions of a water layer. Various approaches have been taken when incorporating hydrogen-rich envelopes into internal structure models, ranging from isothermal H/He layers (Seager et al., 2007) to envelopes following an analytic temperature profile (Rogers et al., 2011, Valencia et al., 2013) or a temperature profile derived from self-consistent atmospheric modelling (Madhusudhan et al., 2020).

Internal structure models have been used to characterise observed super-Earths and mini-Neptunes, both through population studies and application to specific planets. At the population level, Rogers (2015) used M – R relations along with a hierarchical Bayesian analysis of the *Kepler* sample of planets to show that planets with radii $\gtrsim 1.6R_{\oplus}$ are most likely volatile-rich. Dorn et al. (2015, 2017) developed a Bayesian framework to infer

super-Earth compositions and place constraints on the extent of a volatile envelope, finding that solutions are often highly degenerate, with a range of compositions able to explain a given mass and radius. Lopez et al. (2012) predicted that the H₂O component of planets in the Kepler-11 system would be in vapour, molecular fluid or ionic fluid phases. As one of the first exoplanets smaller than Neptune to have both a measured mass and radius, the planet GJ 1214b has been the subject of numerous studies aiming to characterise its interior. The bulk properties of the planet suggest the presence of a H/He envelope, with the possibility of a H₂O layer in the interior. Rogers & Seager (2010b) showed that the planet most likely has a substantial H/He layer above a layer of H₂O, which would be in supercritical and sublimated ice phases. The presence of a gaseous layer was also suggested by Nettelmann et al. (2011), whose models favoured a metal-enriched H/He atmosphere with a significant H₂O mass fraction. An upper limit on the H/He mass fraction of 7% was proposed by Valencia et al. (2013). As mentioned previously, Madhusudhan et al. (2020) analysed the interior and atmosphere of K2-18b, a habitable-zone temperature mini-Neptune, constraining the planet's H/He mass fraction to $\lesssim 6\%$. Constraints on the planet's interior structure showed that conditions at the surface of the H₂O layer could range from supercritical to liquid phases.

In this paper we present internal structure models for super-Earths and mini-Neptunes, with the aim of thoroughly exploring the phase structures of H₂O layers on such planets. We describe our model in Section 5.2, including our compilation of the H₂O equation of state (EOS) that is valid across a wide pressure and temperature range, and approaches to incorporating the temperature structure of a H/He envelope. In Section 5.3 we validate our model against previously published results. We then use our models to explore the phase structures of H₂O-rich planets in Section 5.4. We investigate in detail how the ocean depth depends on observable properties such as surface gravity and temperature. We also constrain the range of masses and radii of planets that might host liquid water beneath H/He envelopes. We explore the wide range of internal structures that may be present on H₂O-rich planets with different surface conditions. We also show M – R relations for planets with mixed H/He–H₂O envelopes. Finally in Section 5.5 we summarise our results and discuss possible caveats, implications of our findings and avenues for future study.

5.2 Methods

The canonical planetary internal structure model consists of a differentiated two-component nucleus made up of an iron core and silicate mantle, a layer of H₂O and/or a H/He envelope (e.g. Léger et al., 2004, Fortney et al., 2007, Seager et al., 2007, Sotin et al., 2007, Valencia et al., 2007). We adopt this approach and follow the standard assumption that the core and

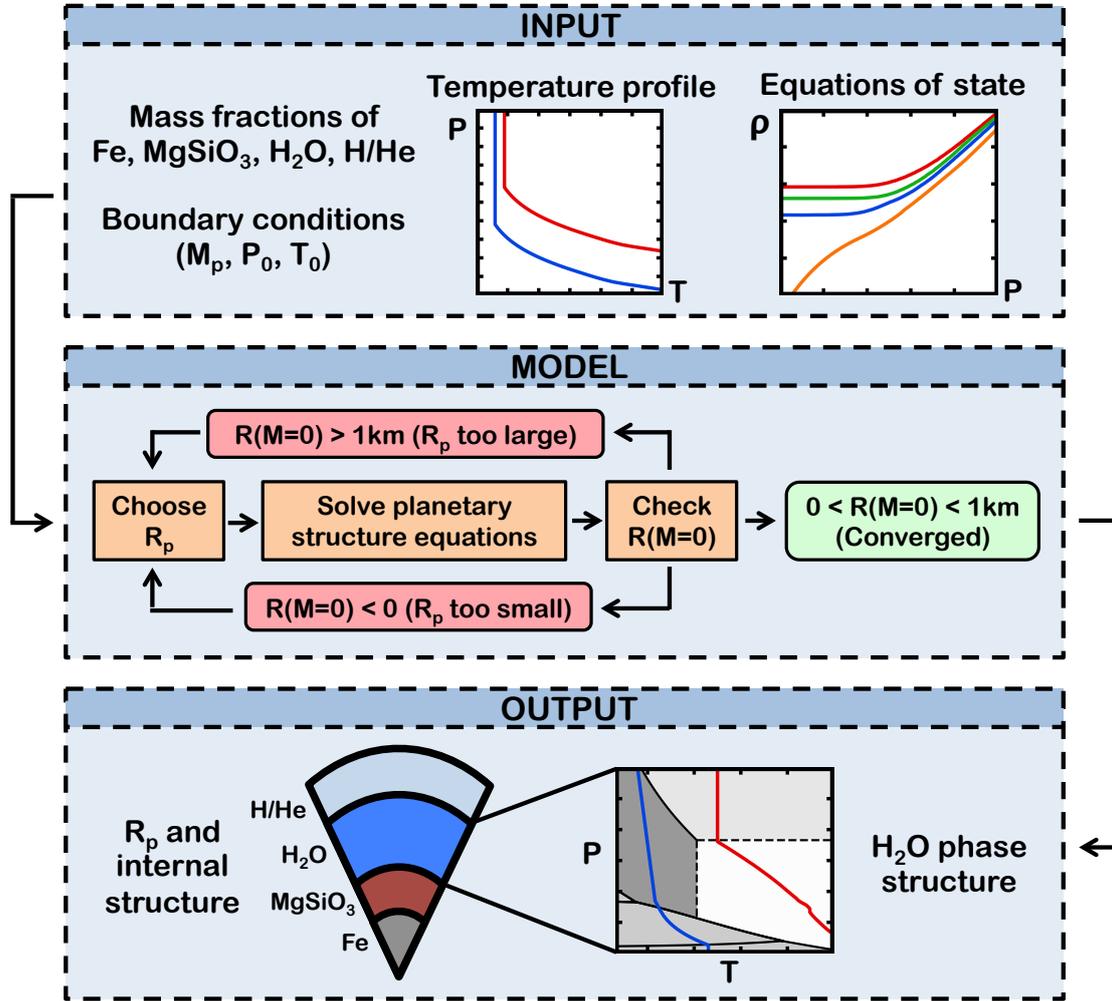


Fig. 5.1 Architecture of the internal structure model used in this study.

mantle are fully differentiated. However, since it has been suggested that water and hydrogen are miscible in the interiors of giant planets (Soubiran & Militzer, 2015), we therefore allow for a mixed H/He and H₂O envelope as well as a differentiated structure with an outer H/He envelope on top of the H₂O layer. The components of our model are depicted in Figure 5.1 and described in detail in this section.

5.2.1 Planetary Structure Equations

Our model solves the equations which determine the interior structure of a planet: the mass continuity equation,

$$\frac{dR}{dM} = \frac{1}{4\pi R^2 \rho}, \quad (5.1)$$

Table 5.1 Parameters for the EOS of Fe (ϵ) from Anderson et al. (2001) and MgSiO₃ (perovskite) from Karki et al. (2000).

Component	B_0 (GPa)	B'_0	B''_0 (GPa ⁻¹)	ρ_0 (kg m ⁻³)
Fe	156.2	6.08	n/a	8300
MgSiO ₃	247	3.97	-0.016	4100

where M is the mass of a spherical shell of material internal to a radius R and density ρ , and the equation of hydrostatic equilibrium,

$$\frac{dP}{dM} = -\frac{GM}{4\pi R^4}, \quad (5.2)$$

where P is the pressure at the shell. Linking these equations requires an EOS $\rho = \rho(P, T)$ as well as a pressure–temperature (P – T) profile $T = T(P)$, or simply $\rho = \rho(P)$ for a temperature-independent EOS. The EOS for each component is described in Section 5.2.2 and the temperature profiles used in this study are discussed in Section 5.2.3.

These equations are solved using a fourth-order Runge-Kutta scheme, with the mass interior to a shell taken as the independent variable. Previous internal structure models have performed the integration either by starting at the surface of the planet and proceeding inward (e.g. Rogers & Seager, 2010a, Madhusudhan et al., 2012, Thomas & Madhusudhan, 2016, Madhusudhan et al., 2020), or by integrating outward from the centre (e.g. Seager et al., 2007, Sotin et al., 2007, Noack et al., 2016). We choose to integrate inward from the surface since in this case the boundary conditions to be specified are surface conditions of the planet, which are more closely linked to observable parameters than the conditions at the centre of the planet. For example, Madhusudhan et al. (2020) used the retrieved reference pressure from the transmission spectrum of K2-18b as a boundary condition for the interior models. Our boundary conditions are therefore the temperature and pressure at the photosphere of the planet (i.e. the conditions at $R = R_p$.)

We solve for the planetary radius R_p at a given mass M_p and mass fractions $x_i = M_i/M_p$ of iron, silicates, H₂O and H/He. R_p is found using a bisection root-finding scheme. The value of R_p is updated iteratively for solving the structure equations until the conditions $0 < R(M = 0) < 1$ km are satisfied.

5.2.2 Equations of State

Previous works have explored a number of EOS prescriptions for different model components. Some early studies used an isothermal EOS for each component (Zapolsky & Salpeter, 1969, Seager et al., 2007), while others considered temperature dependence throughout the interior

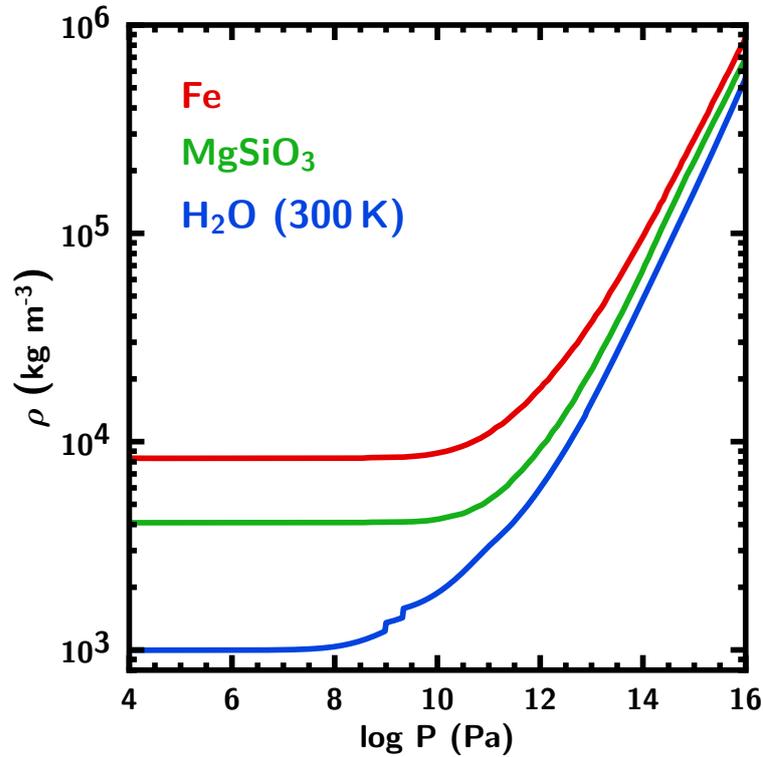


Fig. 5.2 Equations of State for Fe and MgSiO₃ used in our model. We use a temperature-independent EOS adopted from Seager et al. (2007) for each of these components. The EOS for H₂O at 300 K is also shown for comparison.

(Valencia et al., 2006, Sotin et al., 2007). It has been demonstrated that thermal effects in the iron and silicate layers do not substantially alter the $M-R$ relation (Grasset et al., 2009, Howe et al., 2014), and therefore several models treat the inner layers as isothermal with a temperature-dependent prescription for other components (e.g. Rogers & Seager, 2010a, Zeng & Sasselov, 2013, Thomas & Madhusudhan, 2016). Although a number of past models have used isothermal H₂O layers (e.g. Hubbard & Marley, 1989, Seager et al., 2007), several works have shown that thermal effects in this layer can significantly affect the $M-R$ relation (Thomas & Madhusudhan, 2016, Mousis et al., 2020, Turbet et al., 2020). Constructing an EOS for H₂O can be challenging due to the need to accurately describe the behaviour of the many different phases of water. For example, the International Association for the Properties of Water and Steam (IAPWS) provide a detailed functional EOS (Wagner & Pruß, 2002) covering liquid, vapour and some of the supercritical phase, and French et al. (2009) present an EOS that is valid for high-pressure ice. One way of avoiding this problem is to compile a patchwork EOS that uses different prescriptions for different phases. This approach was taken

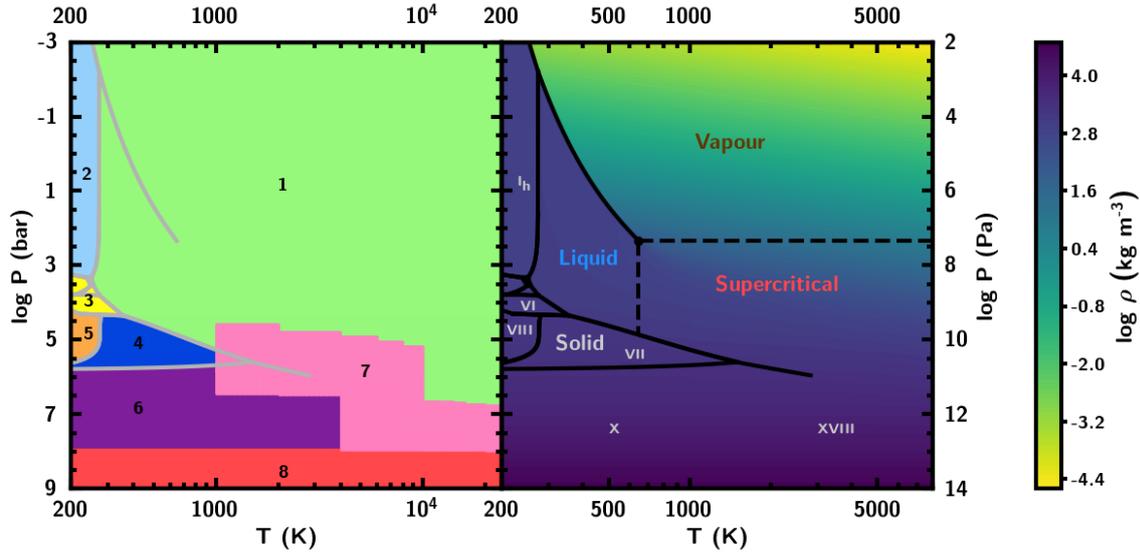


Fig. 5.3 *Left*: Data sources for the H_2O EOS used in our model: 1. Wagner & Pruß (2002). 2. Feistel & Wagner (2006). 3. Journaux et al. (2020a). 4. Fei et al. (1993). 5. Klotz et al. (2017), Fei et al. (1993). 6. Seager et al. (2007). 7. French et al. (2009). 8. Salpeter & Zapolsky (1967). The EOS of Salpeter & Zapolsky (1967) is also used for $P > 10^{14}$ Pa. *Right*: Phase diagram of H_2O (Wagner & Pruß, 2002, Dunaeva et al., 2010). The contour plot shows the EOS $\rho = \rho(P, T)$ used in our model. Regions of P - T space are labelled with their corresponding phase. The transition from ice X to ice XVIII (also called superionic ice) occurs at approximately 2000 K (Milot et al., 2019).

by Thomas & Madhusudhan (2016), who compiled a temperature-dependent H_2O EOS valid at pressures from $10^5 - 10^{22}$ Pa and temperatures from 300 – 24 000 K). A similar method was adopted by Mazevet et al. (2019), who constructed an EOS covering the liquid, gas and superionic regimes.

For the iron and silicate layers we adopt the isothermal (room-temperature) EOS described in Seager et al. (2007). We present a revised and updated version of the temperature-dependent H_2O EOS from Thomas & Madhusudhan (2016) and we use the recently published H/He EOS from Chabrier et al. (2019). We now describe each of these prescriptions in detail.

5.2.2.1 Iron

Figure 5.2 shows the temperature-independent EOS that we use for each component of the nucleus, alongside our H_2O EOS at 300 K for comparison. We adopt the EOS of the ϵ phase of Fe from Seager et al. (2007), which consists of a Vinet fit (Vinet et al., 1989) to experimental data from Anderson et al. (2001) at lower pressures ($P < 2.09 \times 10^{13}$ Pa), and

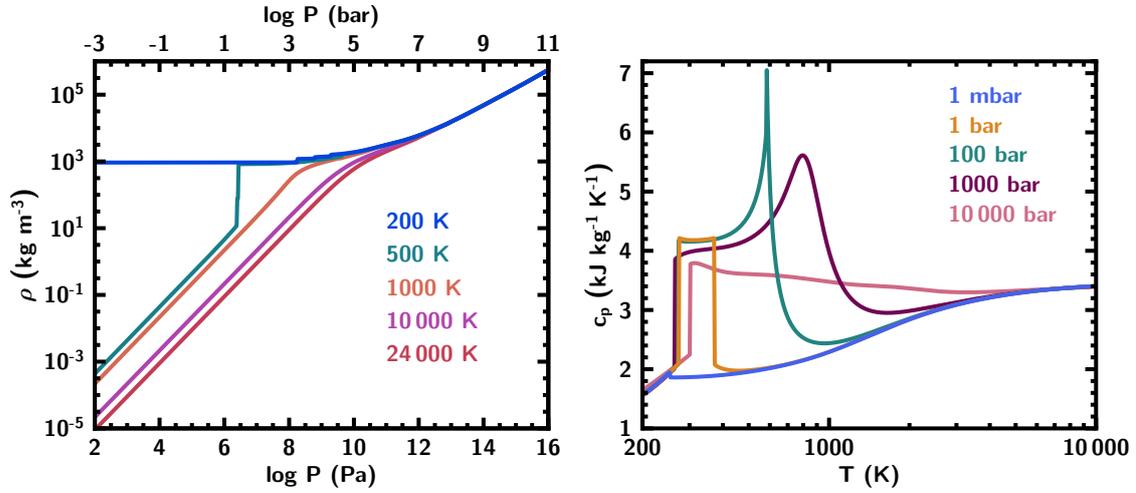


Fig. 5.4 *Left*: Density of H_2O as a function of pressure for several different isotherms. At 500 K there is a phase transition from vapour to liquid at 26.4 bar. *Right*: Specific heat capacity c_p of H_2O as a function of temperature for a range of isobars. The value of c_p increases sharply across the ice-liquid phase transition and decreases across the liquid-vapour boundary. At high temperatures away from the phase boundaries c_p does not vary significantly.

transitions to the Thomas-Fermi-Dirac (TFD) EOS as described in Salpeter & Zapolsky (1967) at higher pressures. The Vinet EOS takes the functional form

$$P = 3B_0\eta^{2/3} \left(1 - \eta^{-1/3}\right) \exp \left[\frac{3}{2} (B'_0 - 1) \left(1 - \eta^{-1/3}\right) \right] \quad (5.3)$$

where $B_0 = \rho(\partial P/\partial \rho)|_T$ is the isothermal bulk modulus, B'_0 is the pressure derivative of B_0 and $\eta = \rho/\rho_0$, where ρ_0 is the ambient density. The values used in this study for those quantities are shown in Table 5.1. For the TFD EOS, electrons are treated as a gas of non-interacting particles in a slowly-varying potential. At lower pressures chemical bonds and the crystal structure of a material play an important role in determining the EOS, and the TFD approximation breaks down since these effects are not considered. However at very high pressures, when kinetic energy dominates over Coulomb energy, TFD theory can yield an accurate EOS. The low-pressure Fe EOS smoothly transitions into the TFD EOS allowing the pressure range to be extended beyond what is obtainable solely from experimental data.

5.2.2.2 Silicates

The silicate EOS used in our model is also adopted from Seager et al. (2007) and comes from the perovskite phase of MgSiO_3 . For $P < 1.35 \times 10^{13} \text{ Pa}$ this takes the form of a fourth-order

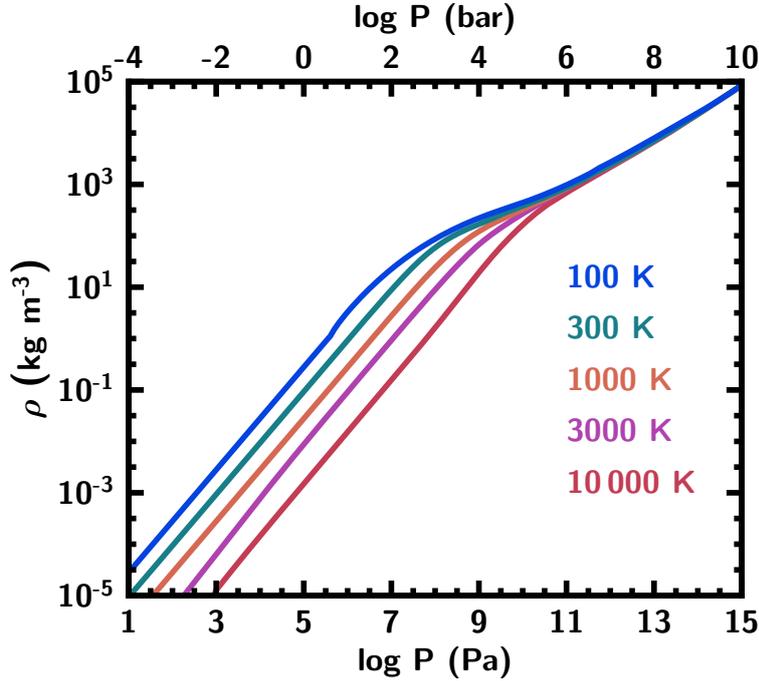


Fig. 5.5 EOS for H/He used in our model for different isotherms. At low pressures ($\lesssim 10^7$ Pa) the EOS behaves like an ideal gas, and at high pressures ($\gtrsim 10^{12}$ Pa) it becomes temperature-independent. The EOS data is taken from Chabrier et al. (2019).

finite strain Birch-Murnaghan EOS (Birch, 1952) fit to a density functional calculation from Karki et al. (2000). The fourth-order Birch-Murnaghan equation is

$$\begin{aligned}
 P = \frac{3}{2}B_0 \left(\eta^{7/3} - \eta^{5/3} \right) & \left\{ 1 + \frac{3}{4} (B'_0 - 4) \left(\eta^{2/3} - 1 \right) \right. \\
 & \left. + \frac{3}{8} B_0 \left(\eta^{2/3} - 1 \right)^2 \left[B_0 B''_0 + B'_0 (B'_0 - 7) + \frac{143}{9} \right] \right\}
 \end{aligned} \tag{5.4}$$

where B''_0 is the second pressure derivative of B_0 . The values used in this study for the relevant quantities are shown in Table 5.1. As with the Fe EOS, this smoothly transitions into a TFD EOS at high pressures.

5.2.2.3 Water

For the liquid, vapour and supercritical phases of H_2O we use the prescription presented in Thomas & Madhusudhan (2016). However, we have used different sources for several ice phases, both to extend the validity of the EOS to lower pressures and temperatures and to

incorporate newer data. The resulting EOS is valid for pressures from $10^2 - 10^{22}$ Pa and temperatures from 200 – 24 000 K. Figure 5.3 shows the various sources used to compile the EOS and their regions of validity in P – T space. We now describe our choice of EOS for each of the relevant phases of H_2O .

Liquid and vapour. We use the functional EOS from the International Association for the Properties of Water and Steam (IAPWS, Wagner & Pruß, 2002). The IAPWS EOS has been validated by numerous experiments and covers the region of P – T space above the melting curve of H_2O (lowest temperature 251.2 K).

Supercritical. Wagner & Pruß (2002) showed that the IAPWS formulation can be extrapolated to pressures and temperatures beyond the critical point of H_2O . We therefore adopt their formulation for some of the supercritical phase. However, for $T > 1000$ K and $P > 1.86 \times 10^9$ Pa we smoothly transition to the EOS presented by French et al. (2009). This EOS is derived from quantum molecular dynamics simulations of high-pressure ice, supercritical and superionic H_2O , and has been validated by experiments (Knudson et al., 2012).

Low-pressure ice. We incorporate the EOS for ice Ih from Feistel & Wagner (2006) which is widely considered to be the best available formulation for this phase (see e.g. Journaux et al., 2020b). For ices II, III, V and VI we use the latest available EOS from Journaux et al. (2020a) which is derived from experiments conducted at a range of pressures and temperatures.

High-pressure ice. We use the French et al. (2009) EOS where applicable, covering parts of the ice VII, X, and XVIII phases. Ice XVIII exists at $P > 10^{11}$ Pa, $T > 2000$ K and is also called superionic ice (Millot et al., 2019). For the remainder of the ice VII phase we follow the approach of Fei et al. (1993), who used a Vinet EOS with a thermal correction. Klotz et al. (2017) determined a functional form of the coefficient of volumetric thermal expansion, α , that allows for the extrapolation of the ice VII EOS down to the ice VIII phase at lower temperatures, finding good agreement with experimental data. We therefore use their prescription for α to calculate an EOS for ice VIII. Thermal effects become negligible for very high pressures, and so we switch to the temperature-independent TFD theoretical EOS as described in Salpeter & Zapolsky (1967) for pressures above 7.686×10^{12} Pa. For intermediate regions not covered by another data source we use the EOS from Seager et al. (2007) in order to smoothly transition to the TFD regime.

5.2.2.4 Hydrogen/Helium

We use the temperature-dependent H/He EOS from Chabrier et al. (2019) for a solar helium mass fraction ($Y = 0.275$), which covers pressures from $1-10^{22}$ Pa and temperatures from

100–10⁸ K. At temperatures relevant to our model, the hydrogen EOS is a combination of the semi-analytical model from Saumon et al. (1995) at low densities ($\rho \leq 50 \text{ kg m}^{-3}$), the model based on ab initio electronic structure calculations from Caillabet et al. (2011) at intermediate densities ($300 < \rho \leq 5000 \text{ kg m}^{-3}$), and the model for fully ionised hydrogen from Chabrier & Potekhin (1998) at high densities ($\rho > 10^4 \text{ kg m}^{-3}$). Similarly, the helium EOS is derived using a combination of models from Saumon et al. (1995) for $\rho \leq 100 \text{ kg m}^{-3}$ and Chabrier & Potekhin (1998) for $\rho > 10^5 \text{ kg m}^{-3}$, and ab initio calculations based on quantum molecular dynamics for $1000 < \rho \leq 10^5 \text{ kg m}^{-3}$. In both cases a bicubic spline procedure is used to interpolate the thermodynamic quantities between the given regimes.

The combined H/He EOS is produced using an additive volume law, which takes the form

$$\frac{1}{\rho_{\text{mix}}(P, T)} = \sum_i \frac{x_i}{\rho_i(P, T)}, \quad (5.5)$$

where x_i is the mass fraction of the i^{th} component. This prescription does not consider interactions between the two species, but Chabrier et al. (2019) claim that the correction to the EOS from taking this into account should only be of the order of a few per cent. The resulting EOS at several different temperatures is shown in Figure 5.5. The same additive volume law is used to compute the density of a mixed envelope consisting of H/He and H₂O in this study.

5.2.3 Temperature Profiles

The inclusion of EOS data across a wide range of pressures and temperatures enables us to consider any reasonable temperature profile within the H₂O and H/He layers. Deep within the interior of the planet we expect convection to dominate energy transport, leading to an adiabatic temperature profile with constant specific entropy S . However, at lower pressures near the planetary surface an adiabatic temperature profile may no longer be appropriate. Previous studies have used different methods to model the temperature structure of the outer envelope: for example, Fortney et al. (2007) took P – T profiles from a grid of atmospheric models, while Rogers et al. (2011) and Valencia et al. (2013) incorporated the analytic atmospheric model from Guillot (2010). Recent works have also coupled interior models of water worlds to steam atmosphere models (Mousis et al., 2020, Turbet et al., 2020). Here we describe some of the most common approaches used to model the temperature structures of planetary interiors: analytic models and self-consistent models. We subsequently describe the approach that we use throughout most of this paper, in which the temperature profile consists of an isothermal layer above an adiabatic layer. While our model has the capability to incorporate both analytic P – T profiles as well as those produced using a self-consistent

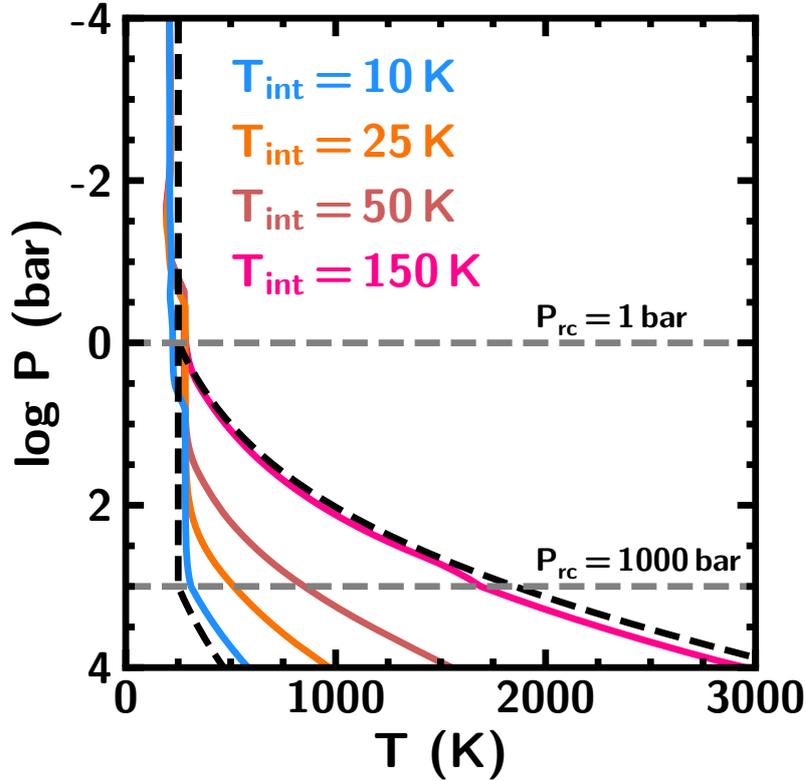


Fig. 5.6 Comparison of analytic and self-consistent P - T profiles. The solid curves show profiles from self-consistent models (Piette & Madhusudhan, 2020b), color-coded by the intrinsic temperature T_{int} used for the model. The self-consistent models span a broad range of possible internal energies for sub-Neptunes, assuming nominal planet properties based on the mini-Neptune K2-18b. Black dashed lines show isothermal/adiabatic profiles with $P_{\text{rc}} = 1 \text{ bar}$ and $P_{\text{rc}} = 1000 \text{ bar}$. We consider these as end-member scenarios when modelling planets with H/He envelopes.

atmospheric model, we find that isothermal/adiabatic temperature profiles are the most appropriate for this study, as we explain in Section 5.2.4.

5.2.3.1 Analytic profiles

The temperature profile of a planet's outer envelope can be calculated using an analytic model, such as the one described in Guillot (2010), which takes the form

$$\begin{aligned}
T^4 = & \frac{3T_{\text{int}}^4}{4} \left[\frac{2}{3} + \tau \right] \\
& + \frac{3T_{\text{irr}}^4}{4} f \left[\frac{2}{3} + \frac{1}{\gamma\sqrt{3}} + \left(\frac{\gamma}{\sqrt{3}} - \frac{1}{\gamma\sqrt{3}} \right) \exp(-\gamma\tau\sqrt{3}) \right]
\end{aligned} \tag{5.6}$$

which is a solution to the equations of radiative transfer assuming a grey atmosphere and the two-stream approximation. The irradiation temperature $T_{\text{irr}} = f^{-1/4}T_{\text{eq}}$ characterises the irradiation intensity from the host star and is related to the planetary equilibrium temperature via the redistribution factor f . The intrinsic temperature T_{int} characterises the planetary intrinsic heat flux. The ratio of visible to thermal opacities is represented by γ . This approach requires the inclusion of another differential equation in the model to solve for the optical depth τ :

$$\frac{d\tau}{dM} = -\frac{\kappa}{4\pi R^2}, \tag{5.7}$$

where the opacity κ can be specified as a function of P and T using, for example, the tabulated values of Freedman et al. (2008) for H/He. The analytic model described here has some limitations: for example, there is no treatment of clouds, which can have a significant impact on the form of the temperature profile (e.g. Kitzmann et al., 2010) and may be prevalent in super-Earth atmospheres (Kreidberg et al., 2014b). In this study, we only use analytic temperature profiles in order to validate our model against previous work that used this prescription (see Section 5.3).

5.2.3.2 Self-consistent profiles

Another approach is to use a temperature profile that has been calculated using self-consistent atmospheric modelling (e.g. Gandhi & Madhusudhan, 2017, Malik et al., 2019, Piette & Madhusudhan, 2020b). Self-consistent atmospheric models solve the equations of radiative transfer numerically under the assumptions of hydrostatic, radiative-convective and thermal equilibrium. These models are able to account for many more phenomena than the analytic prescription, such as atmospheric dynamics, clouds and particle scattering. Detailed self-consistent atmospheric modelling is not incorporated directly into our internal structure model, but temperature profiles calculated in this way can be used to obtain the density profile for the outer layers of a planet (Madhusudhan et al., 2020). Self-consistent atmospheric modelling requires many planet-specific parameters and is more time-consuming than the other approaches discussed here, so while this method is useful for exploring the structure

of a particular planet, it is less well-suited to theoretical calculations across a wide range of masses and radii. We show in Section 5.2.3.3 that an isothermal/adiabatic profile can be used in place of a self-consistently modelled profile with little change to the M – R relation.

5.2.3.3 Isothermal/adiabatic profiles

This P – T profile consists of an isotherm at the photospheric temperature T_0 down to the radiative-convective boundary, at which point the temperature profile becomes adiabatic. This approach to calculating the temperature profile allows for a high degree of flexibility while remaining simple to compute. The pressure at the radiative-convective boundary (P_{rc}) is a free parameter.

The adiabatic temperature gradient is

$$\left. \frac{\partial T}{\partial P} \right|_S = \frac{\alpha T}{\rho c_p}, \quad (5.8)$$

where c_p is the specific heat capacity at constant pressure and α is the coefficient of volumetric thermal expansion. Chabrier et al. (2019) present the adiabatic gradient for H/He along with their EOS and so we incorporate this directly into our model. For H_2O we require prescriptions for c_p and α . Thomas (2016) incorporated c_p data for the liquid and vapour phases from (Wagner & Pruß, 2002) and extrapolated this to cover all other phases of H_2O . We also use this data for liquid and vapour, but we do not extrapolate beyond these regions. Instead, we add data from Feistel & Wagner (2006) for ice Ih, Journaux et al. (2020a) for ices II, III, V and VI, Fei et al. (1993) for ices VII and VIII and French et al. (2009) for the ice VII–X transition. The behaviour of c_p is summarised in the right-hand panel of Figure 5.4. For higher pressures where c_p data is unavailable (i.e. sources 6 and 8 in Figure 5.3), we assume that c_p is equal to its value at the nearest point in P – T space with available data. The true value of c_p is not required here since the EOS used for these pressures is not temperature-dependent.

We calculate α directly from our EOS:

$$\alpha = \left. \frac{1}{V} \frac{\partial V}{\partial T} \right|_P = - \left. \frac{\partial \ln \rho}{\partial T} \right|_P. \quad (5.9)$$

Transitions between different phases of H_2O can lead to significant discontinuities in the EOS, causing α to become undefined at phase boundaries. In order to avoid this, the derivative is calculated separately for each phase and smoothly interpolated across the boundary, yielding adiabats that remain continuous.

For mixed envelopes consisting of both H/He and H₂O, the adiabatic gradient is calculated by linear interpolation using the following formula:

$$\left(\frac{\partial \log T}{\partial \log P}\right)_{S, \text{mix}} = -\frac{\sum_i x_i S_i \frac{\partial \log S_i}{\partial \log P} \Big|_T}{\sum_i x_i S_i \frac{\partial \log S_i}{\partial \log T} \Big|_P}, \quad (5.10)$$

with values of the specific entropy of H/He and H₂O taken from the same sources as the values of c_p .

5.2.4 Choice of Temperature Profiles in This Study

In the rest of this paper we use isothermal/adiabatic P – T profiles as described in the previous subsection. Here we show that these P – T profiles are a reasonable approximation to those generated by self-consistent models. We take a number of P – T profiles generated using the self-consistent model `GENESIS` (Gandhi & Madhusudhan, 2017), which was recently updated to model atmospheres of sub-Neptunes (Piette & Madhusudhan, 2020b). We consider a number of models with a wide range of internal energies, which are determined by T_{int} . T_{int} can be calculated using evolutionary models. For example, Lopez & Fortney (2014) find that, for a low-mass 5 Gyr-old planet, T_{int} can be as low as ~ 10 K. Conversely, Valencia et al. (2013) find that for the mini-Neptune GJ 1214b, T_{int} may be up to 80 K at an age of 0.1 Gyr, and Morley et al. (2017a) consider even higher values for the Neptune-mass GJ 436b, whose interior may be warmed by tidal heating. We therefore explore a range of 10–150 K for T_{int} .

The resulting temperature profiles are shown in Figure 5.6, and can be closely matched by isothermal/adiabatic profiles with P_{rc} lying between 1 bar and 1000 bar. We therefore take 1–1000 bar to be a reasonable range of values for P_{rc} when considering a general sub-Neptune atmosphere in the remainder of this paper.

To further illustrate that a model incorporating an isothermal/adiabatic temperature profile can yield very similar results to a model with a temperature profile generated by a self-consistent model, we consider one of the interior models for the planet K2-18b from Madhusudhan et al. (2020), which used P – T profiles produced by `GENESIS`. We take case 2 from that paper, which has a composition of 45% Earth-like nucleus, 54.97% H₂O, and 0.03% H/He. We fit the temperature profile used for that case with an isothermal/adiabatic profile, finding best-fit parameters $T_0 = 300$ K, $P_0 = 0.05$ bar, and $P_{\text{rc}} = 3$ bar. The model using an isothermal/adiabatic profile gives a radius of $2.613 R_{\oplus}$ at the mean observed planet mass ($8.63 M_{\oplus}$, Cloutier et al., 2019) which, like the model from Madhusudhan et al. (2020), agrees with the observed planetary radius to well within the observational uncertainty ($2.610 \pm 0.087 R_{\oplus}$; Benneke et al., 2019).

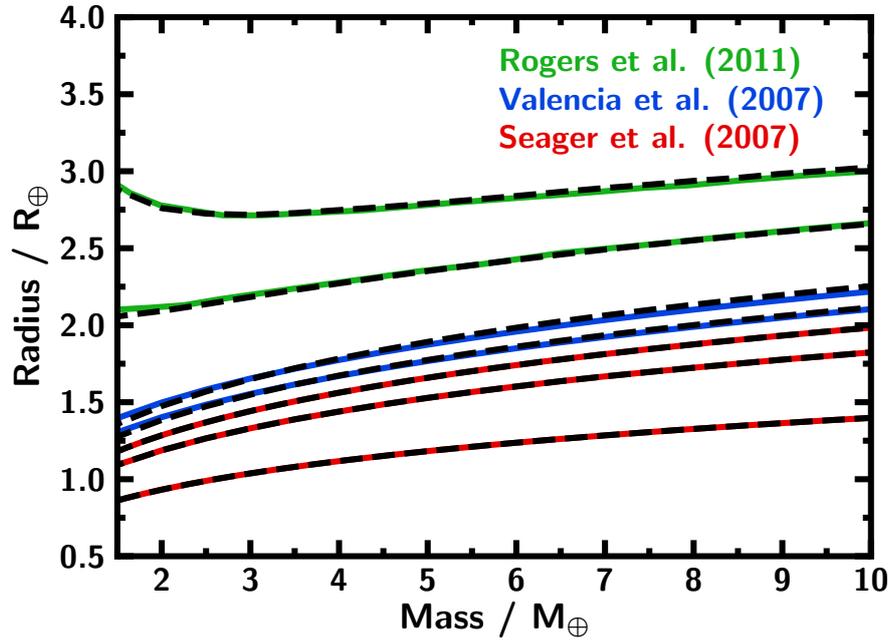


Fig. 5.7 Comparison between M – R relations produced by our model and previously published results. The red curves are taken from Seager et al. (2007) for planets of pure Fe, Earth-like and pure silicate composition. The blue curves are taken from Valencia et al. (2007) for planets with an Earth-like core and adiabatic H_2O layers of 30% and 50%. The green curves are taken from Rogers et al. (2011) for planets with a 33% Earth-like nucleus, a $\sim 67\%$ isothermal H_2O layer and H/He envelopes of 0.1% and 1%. The black dashed lines show our model reproductions of each of these cases. All of our results agree with the previously published M – R relations to within $0.05R_{\oplus}$.

5.3 Model Validation

In this section we validate our model by reproducing results from a number of previous works concerning the internal structures of sub-Neptunes. We also examine the effect of a temperature-dependent interior on the radii of H_2O -rich planets and compare our results to other studies that have followed similar approaches.

5.3.1 Comparison with Previous Studies

Figure 5.7 shows several M – R relations from previous works alongside the results of our model given the same input parameters. We begin by reproducing the results of Seager et al. (2007) for isothermal planets comprised of iron and silicates. Since we use the same EOS for Fe and MgSiO_3 as that paper, we expect to find total agreement between their results and

our model when considering planets consisting solely of iron and rock. We reproduce $M-R$ relations for three compositions shown in figure 4 of Seager et al. (2007): a pure iron planet, a pure silicate planet, and a planet consisting of 32.5% iron and 67.5% silicates. In each of these cases our results agree completely with the published $M-R$ relations.

Next we consider $M-R$ relations shown in figure 3 of Valencia et al. (2007) for planets at $T_0 = 550\text{K}$ with H_2O mass fractions of 10% and 30%. The H_2O layer consists of ice VII and X and follows an adiabatic temperature profile. The inner layers of the planet are made up of silicates and iron in a 2:1 ratio. Valencia et al. (2007) do not provide a surface pressure for their models, so in order to reproduce their results we take $P_0 = 10^{10}\text{Pa}$, since this forces the phase at the surface to be ice VII rather than liquid. The ice VII EOS used in our models differs slightly from Valencia et al. (2007), however we find good agreement between the two sets of $M-R$ curves. Across all masses and compositions considered, the largest discrepancies in radius are below $0.04R_\oplus$, smaller than the observational uncertainties of even the best super-Earth radius measurements ($\sim 0.1R_\oplus$). The small differences that do appear are likely due to the different formulation of the ice VII EOS and the fact that Valencia et al. (2007) incorporate conduction in the mantle.

Finally we compare our model to equilibrium models of planets with H/He envelopes from Rogers et al. (2011). We consider $M-R$ relations shown in figure 4 of that paper for planets with $T_{\text{eq}} = 500\text{K}$ and H/He mass fractions of 0.1% and 1%. The underlying composition in each case is 10% Fe, 23% silicates and 67% H_2O . Rogers et al. (2011) did not use a temperature-dependent H_2O EOS, instead adopting the isothermal EOS from Seager et al. (2007), and so for our reproduction we take an isothermal temperature profile at 300 K in the H_2O layer. Rogers et al. (2011) also used the H/He EOS from Saumon et al. (1995), which differs from the Chabrier et al. (2019) EOS at high densities ($\rho > 50\text{kg m}^{-3}$). The temperature profile used in the H/He envelope is analytic, taking the form of Equation 5.6 with $\gamma = 0.6\sqrt{T_{\text{irr}}/2000\text{K}}$, $f = 1/4$ and $T_{\text{int}} = (L_p/4\pi R_p^2\sigma)^{1/4}$, where L_p is the intrinsic luminosity of the planet and σ is the Stefan-Boltzmann constant. For the models reproduced here, L_p is determined by fixing $L_p/M_p = 10^{-10.5}\text{W kg}^{-1}$. For the purposes of this reproduction we do not correct for the transit radius effect. Again we find good agreement between our models and those of Rogers et al. (2011), with maximum discrepancies less than $0.05R_\oplus$. These may be a result of the differing H/He EOS. Regardless, our $M-R$ relations and all those from previous studies shown here agree to well within typical observational uncertainties for super-Earths.

5.3.2 Mass–radius Relations

Here we use our model to produce mass–radius (M – R) relations for sub-Neptune exoplanets with varying temperature structures and compositions, and compare these to previous work in the field.

5.3.2.1 Water worlds with no H/He envelope

Figure 5.8 shows M – R curves for H_2O -rich planets with no H/He envelope. We assume that the underlying nucleus is Earth-like in composition (1/3 Fe, 2/3 MgSiO_3 by mass). For the liquid, ice and supercritical phases of H_2O we assume an adiabatic temperature profile. We consider two end-member temperature profiles in the vapour phase: an isothermal profile and an adiabatic profile. We consider the effect of several different parameters on the planetary radius, including the temperature T_0 and pressure P_0 at the photosphere as well as the water mass fraction $x_{\text{H}_2\text{O}}$. The photospheric pressure and temperature of a planet can be constrained from observations: T_0 is closely related to the planetary effective temperature, which can be estimated from the orbital separation of the planet and the luminosity of the host star. P_0 is defined as the pressure at the photosphere (where $R = R_p$). It can be retrieved from an atmospheric spectrum (Welbanks & Madhusudhan, 2019, Nixon & Madhusudhan, 2020) and subsequently used as a boundary condition when analysing a planet’s internal structure (Madhusudhan et al., 2020).

In agreement with the findings of Thomas & Madhusudhan (2016), it is clear from Figure 5.8 that variations in T_0 can have a significant effect on planetary radius in both the isothermal and adiabatic cases. For model scenarios with a large H_2O mass fraction of 0.9 and an isothermal vapour layer, increasing T_0 from 300 K to 1000 K inflates the planetary radius by $0.25R_\oplus$ – $0.4R_\oplus$ depending on the planet mass. If the vapour layer is adiabatic the radius increases more dramatically, by up to $2.2R_\oplus$ in the most extreme case. The inflation is enhanced at lower masses due to the low surface gravity of the planet allowing for a more extended envelope (Rogers et al., 2011). This effect is still significant for planets with a lower H_2O mass fraction: even in the isothermal vapour case, at $x_{\text{H}_2\text{O}} = 0.3$ the radius is inflated by $0.1 - 0.2R_\oplus$ for the 1000 K model planets compared to those at 300 K. We note that even in the case of a minimal H_2O layer, the effect of changing T_0 may be non-negligible: for $x_{\text{H}_2\text{O}} = 0.01$, the radius of a $1M_\oplus$ planet with $T_0 = 300$ K is $0.981R_\oplus$, increasing to $1.076R_\oplus$ at $T_0 = 1000$ K with an isothermal vapour layer or $1.171R_\oplus$ with an adiabatic vapour layer. Even for mostly rocky planets that may host thin H_2O envelopes, temperature dependence within the H_2O layer is still important to consider.

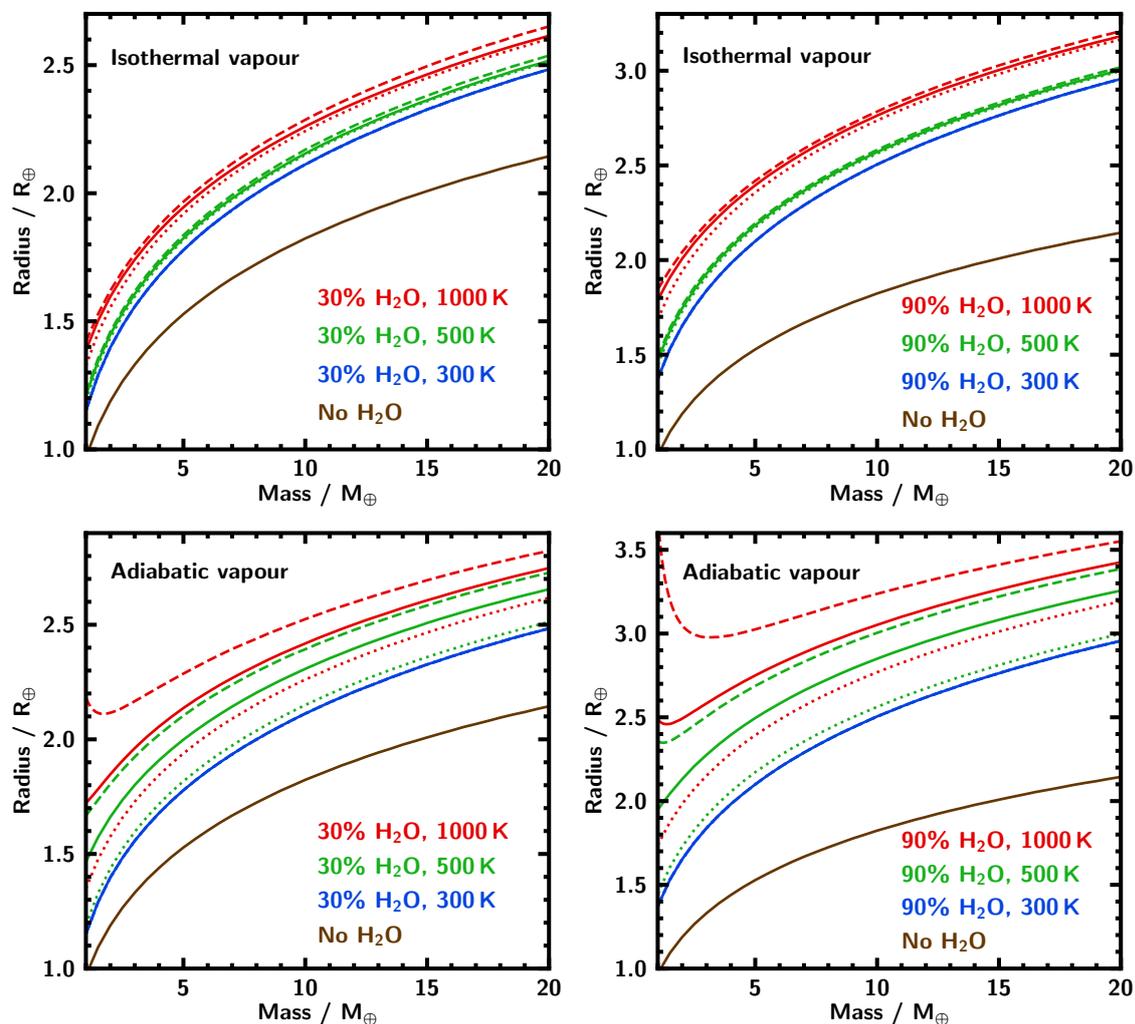


Fig. 5.8 Mass–radius (M – R) relations for water worlds without H/He envelopes. The model planets consist of a H_2O layer of 30% (left) or 90% (right) above an Earth-like nucleus (1/3 iron, 2/3 silicates by mass). All planets follow an adiabatic temperature profile in the liquid, supercritical and ice phases of the H_2O layer. In the top panel, planets with water vapour follow an isothermal temperature profile in the vapour layer, whereas in the bottom panel the temperature profile in this layer is also adiabatic. The line colour denotes the photospheric temperature T_0 and the line style indicates the photospheric pressure: solid lines show models with $P_0 = 1$ bar, dotted lines show models with $P_0 = 100$ bar and dashed lines show models with $P_0 = 0.1$ bar. Increasing T_0 from 300–1000 K increases the planet radius considerably, with a more pronounced effect if the vapour layer is adiabatic. For planets with isothermal vapour layers, changing P_0 does not strongly affect the radius, but decreasing P_0 for a planet with an adiabatic vapour layer can substantially increase its radius.

Figure 5.8 also shows how P_0 affects the planetary radius. In general, a higher P_0 leads to a smaller radius if all other parameters are left unchanged. If an isothermal temperature profile for vapour is assumed, this change is quite small: $\lesssim 0.001R_\oplus$ for planets with $T_0 = 300$ K, and approaching $0.1R_\oplus$ for the lowest masses ($M \sim 1M_\oplus$) at $T_0 = 1000$ K. However, the choice of P_0 affects the M – R relation more significantly if the temperature profile in the vapour layer is adiabatic. At $10M_\oplus$, the radius of a 90% water planet with $T_0 = 1000$ K increases by $0.47R_\oplus$ if P_0 changes from 0.1–100 bar. The effect is more pronounced at higher temperatures since the density of vapour and supercritical H_2O varies more strongly with pressure than the density of liquid (see Figure 5.3). If the vapour layer is adiabatic, then the temperature profile increasing from a lower P_0 leads to a hotter H_2O layer throughout which further magnifies this effect.

We find that a $20M_\oplus$, 90% H_2O planet with a solid or liquid photosphere can have a radius of up to $3R_\oplus$, in agreement with the findings of Zeng et al. (2019) that planets $> 3R_\oplus$ generally require a gaseous envelope. Planets with $x_{\text{H}_2\text{O}} \gtrsim 0.9$ are likely unrealistic from a planet formation perspective, since some amount of refractory material is required to initiate ice and gas accretion (Zeng & Sasselov, 2014, Lee & Chiang, 2016).

The M – R relations reported in Mousis et al. (2020) yield larger radii than our model for a similar composition. This is a result of different methods being used to compute the temperature profile in the water layer. In order to compare the models more directly we consider a $15M_\oplus$ planet with a 20% H_2O layer above a silicate mantle and an adiabatic steam atmosphere with $P_0 = 0.1$ bar, $T_0 = 430$ K, which gives a temperature profile close to that of Mousis et al. (2020) with $T_{\text{eq}} = 300$ K (see figure 1 of that paper). The radius of this model planet is $2.6R_\oplus$, close to the $2.7R_\oplus$ shown in figure 2 of Mousis et al. (2020).

5.3.2.2 Water worlds with H/He envelopes

Next we model planets with an Earth-like nucleus (1/3 iron, 2/3 silicates) and an adiabatic H_2O layer of equal mass to the nucleus underneath a H/He envelope. We adopt isothermal/adiabatic temperature profiles for the H/He layer as described in Section 5.2.3.3, meaning the temperature profile is determined by three parameters: T_0 , P_0 and P_{rc} .

Figure 5.9 shows M – R relations for planets with different photospheric temperatures, radiative-convective boundaries and H/He mass fractions. The main factors governing the M – R relation for water worlds with H/He envelopes are the temperature profile in the envelope and the mass fraction of H/He. Many features of the M – R relations presented in this figure are well-documented in the literature (e.g. Rogers et al., 2011, Lopez & Fortney, 2014), such as the inflation of radii at low masses, which is a result of low surface gravity increasing the atmospheric scale height, and the significant effect of T_0 on the planetary radius.

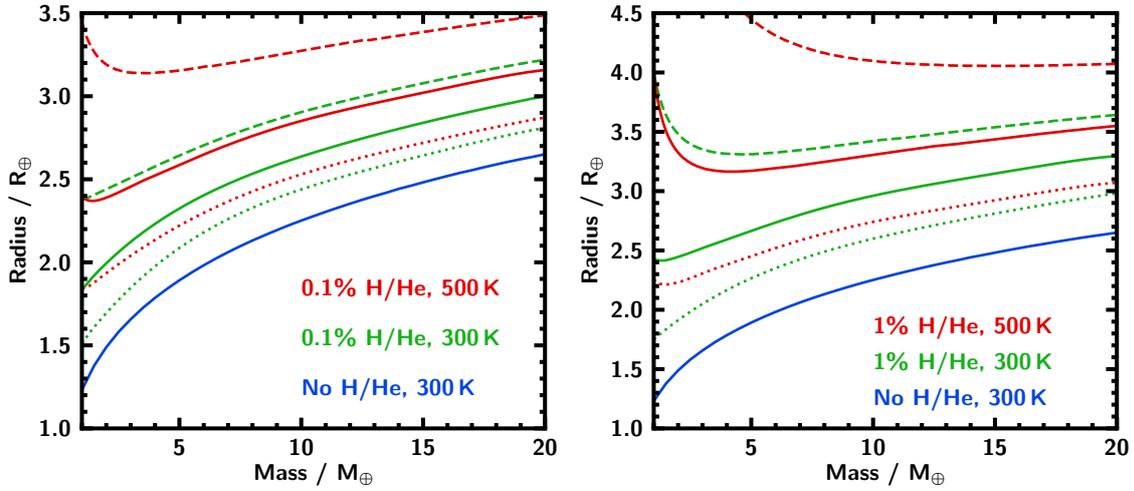


Fig. 5.9 M – R relations for planets with H/He envelopes. The line style indicates the location of the radiative-convective boundary: solid lines show models with $P_{\text{rc}} = 10$ bar, dotted lines show models with $P_{\text{rc}} = 1000$ bar and dashed lines show models with $P_{\text{rc}} = 1$ bar. Even a small amount of H/He can inflate the radius of a planet substantially, and this effect is more pronounced at higher temperatures. Increasing P_{rc} decreases the planet radius.

We consider two values of T_0 , 300 K and 500 K. We also include several different radiative-convective boundaries, with P_{rc} ranging from 1–1000 bar (see Section 5.2.3.3). The results indicate that only a small amount of H/He by mass is required to substantially increase the radius of a planet. Planets with $T_0 > 500$ K may have larger radii than Neptune ($3.88R_{\oplus}$), even with a H/He mass fraction as low as 1%. Changing P_{rc} also has a very significant effect on the M – R relation. This highlights the importance of using an accurate atmospheric temperature profile when modelling planets with a gaseous envelope. In cases where a planet’s temperature profile is not well-constrained, a wide variety of P – T profiles should be considered in order to fully explore its possible internal structures. These thermal effects become more significant as the envelope mass increases.

Throughout this section we assume a nominal photospheric pressure of 0.1 bar. Analysis of the transmission spectra of hydrogen-rich planets yield values of P_0 ranging from ~ 1 mbar– ~ 1 bar (e.g. Welbanks et al., 2019). It is straightforward to compute the change in radius resulting from a change in P_0 , since at pressures below ~ 1000 bar the EOS for the H/He envelope is that of an ideal gas. Combining this with equation 5.2 shows that the change in adjusting the pressure at the photosphere from P_1 to P_2 yields a change in radius of NH_{sc} , where $N = \ln(P_1/P_2)$ and $H_{\text{sc}} = k_B T / \mu g$ is the atmospheric scale height. Adjusting P_0 from 1 bar to 1 mbar increases the planetary radius by $\sim 7 H_{\text{sc}}$.

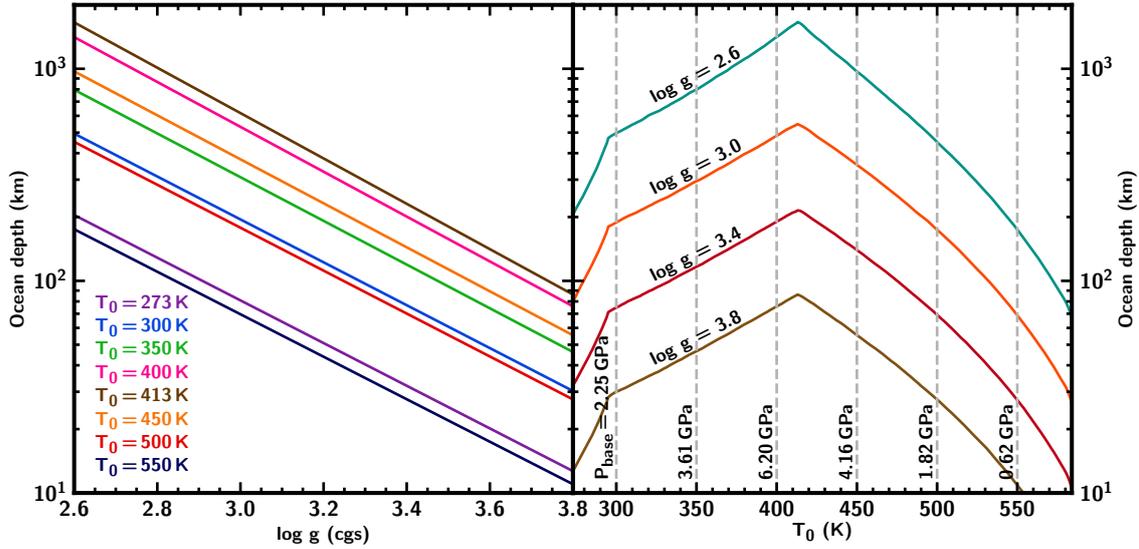


Fig. 5.10 Extent of oceans on H_2O -rich planets with different surface gravities g and temperatures T_0 . The left-hand panel shows the ocean depth against $\log g$ at different values of T_0 from 273–550 K. Ocean depth is inversely proportional to surface gravity for planets without gaseous envelopes. The right-hand panel shows ocean depth against T_0 for different values of $\log g$ (cgs) from 2.6–3.8. The grey dashed lines indicate the pressure at the bottom of the ocean, P_{base} , which depends on T_0 . A higher P_{base} leads to a deeper ocean.

5.4 Results

Here we present our results exploring in detail the internal phase structures of H_2O -rich super-Earths and mini-Neptunes. We investigate four key aspects. We begin by calculating the range of possible ocean depths on planets with a large H_2O component across a wide range of possible bulk properties and temperature structures. We allow for the full extent of the liquid phase of H_2O , reaching temperatures as high as 647 K at pressures from 218– 7×10^4 bar. Next we determine the range of masses, radii and surface conditions for which a mini-Neptune with a H/He envelope may possess a liquid water ocean underneath. We also consider how different temperature profiles can affect the phase structures of water worlds, investigating planets with ice and vapour surfaces as well as those with surface oceans. Finally, we consider how miscibility of H/He and H_2O within a planet’s envelope might affect the M – R relation.

5.4.1 Depth of Oceans on Water Worlds

Our primary goal is to estimate the possible depths of liquid H_2O layers in water-dominated planets over a range of surface conditions. For reference, the Earth’s oceans have an average depth of 3.7 km, extending to 11 km at their deepest point (Charette & Smith, 2010, Gardner

et al., 2014). Given that the Earth’s H₂O mass fraction is ~0.02%, one might expect that planets with a significant portion of their mass in H₂O should possess deeper oceans. However, it is also the case that much of the H₂O layer of a water world can be in supercritical or high-pressure ice phases due to high pressures in the interior. Therefore, the possible extent of an ocean is not simply limited by the amount of H₂O available.

Here we investigate the range of ocean depths that may be achieved in the case where a planet has enough H₂O that the size of the ocean is not limited by water mass fraction. While other works have explored this to some extent as discussed previously, for this study we consider planets across a wider range of masses, compositions and surface conditions. We aim to determine which parameters are important in controlling the extent of an ocean, and how the ocean depth varies across the full parameter space encompassed by super-Earths and mini-Neptunes that may host H₂O layers. We parametrise the temperature profile of the H₂O layer by assuming an adiabatic profile with a H₂O surface of $P_0 = 100$ bar, where $T_0 = T(P_0)$ is a free parameter that determines the adiabat. While the thermodynamic conditions at the surface of the H₂O layer are determined by numerous factors, including the level of irradiation received from the host star, the planet’s intrinsic temperature T_{int} , and atmospheric properties such as opacity from molecular and atomic chemical species, by varying T_0 across the full set of possibilities for a liquid surface we can encapsulate all feasible cases.

Before exploring the extent of oceans found by our internal structure models, we can examine the approximate behaviour of ocean depths by returning to the equation of hydrostatic equilibrium (Equation 5.2), but with radius rather than mass as the independent variable:

$$\frac{dP}{dR} = -\rho g, \quad (5.11)$$

where g is the gravitational acceleration. Consider a planet with a liquid H₂O surface. From Section 5.4.3 we can see that the ocean is unlikely to constitute a large portion of the planet interior, and so we can assume that g is constant throughout the ocean. If we also consider the density of liquid water to remain constant (see Figure 5.3), then we find that the depth of the ocean is proportional to change in pressure from the surface to the base. If the pressure at the base of the ocean is much larger than the surface pressure, then we expect the base pressure P_{base} to strongly affect the ocean depth. For a planet with a large H₂O mass fraction, the value of P_{base} is determined by the location where the adiabatic P - T profile crosses from liquid to either supercritical water or high-pressure ice. This in turn depends on the chosen surface temperature. Therefore we would expect the key parameters in determining the ocean depth to be the gravity and temperature at the ocean’s surface.

We now turn to the full models to explore this in more detail. Our analysis above suggests that planets with the same surface temperature and gravity should have approximately the

same ocean depth, regardless of other factors like the planet mass, radius and composition. We find this to be the case when using the full models. For example, consider two ocean planets corresponding to points from the M – R curves from Figure 5.8: a $5M_{\oplus}$ planet with 30% H_2O by mass, which has a radius of $1.78R_{\oplus}$, and a $10M_{\oplus}$ planet with 90% H_2O by mass, whose radius is $2.51R_{\oplus}$. Both planets have a surface temperature of 300 K and the same surface gravities, $\log g$ (cgs) = 3.19, despite their differing masses, radii and compositions. We find that the oceans on both planets are 125 km deep. Similarly, we find that ocean depth is not strongly affected by the surface pressure or the presence of an atmosphere. However, we note for a planet with a gaseous envelope, the ocean depth is determined by the gravity at the surface of the ocean rather than the gravity with the envelope included.

Having confirmed that surface temperature and gravity are the main parameters which determine the extent of an ocean, we explore the (g, T_0) parameter space in order to find the range of possible ocean depths for H_2O -rich planets. We consider values of $\log g$ (cgs) ranging from 2.6–3.8. A 90% H_2O planet at $1M_{\oplus}$ with $T_0 = 500$ K, has $\log g = 2.67$, which gives a reasonable lower bound for the surface gravity of a water world. The upper limit of $\log g = 3.8$ corresponds to an extreme case of a $20M_{\oplus}$ iron planet with a thin H_2O layer. For reference, the water-rich planets considered later in Figure 5.14 all have $\log g \approx 3.2$. We also consider surface temperatures from 273–584 K, spanning the liquid phase assuming a surface pressure of 100 bar.

Figure 5.10 shows the depths of oceans for planets across this parameter space. As expected from the arguments at the beginning of this section, we find that ocean depth is inversely proportional to gravity and approximately directly proportional to P_{base} , which is determined by the surface temperature. From the adiabatic temperature profiles shown in Figure 5.13 we would expect that the optimum surface temperature in order to maximise P_{base} , and hence ocean depth, lies between 350 and 450 K. From the grid of models we see that the surface temperature which maximises ocean depth assuming an adiabatic temperature profile starting from 100 bar is 413 K. At this temperature, and a minimal $\log g = 2.6$, the ocean depth is 1654 km, about 450 times the average depth of the Earth’s ocean (3.7 km, Charette & Smith, 2010). This can be considered an extreme upper limit to the depth of an ocean on a water world.

In scenarios that have been explored previously, our results are in agreement with past work. For example, for a $6M_{\oplus}$ planet with 50% H_2O we find ocean depths of 66 km for $T_0 = 280$ K and 125 km for $T_0 = 303$ K, which are similar to the values of 72 km and 133 km from Léger et al. (2004).

Looking at the broader parameter space, the models show that a wide variety of ocean depths are possible. For example, a water world with a 300 K surface can have an ocean

depth from 30–500 km, or 8–135 times deeper than the Earth’s ocean, depending on its mass and composition. For a given planet mass, higher H₂O mass fractions lead to deeper oceans, since planets with more H₂O relative to iron and silicates will have lower surface gravities allowing for a more extended liquid water layer. A $1M_{\oplus}$ planet with a 30% water layer and $T_0 = 300$ K has $\log g = 2.83$ and an ocean depth $R_{\text{ocean}} = 283$ km, about 76 times deeper than the average depth on Earth, whereas a $1M_{\oplus}$ planet with the same T_0 but a water mass fraction of 90% has $\log g = 2.70$ and $R_{\text{ocean}} = 388$ km. We can also see that for a fixed composition, more massive planets have higher surface gravities and so R_{ocean} decreases as the mass of the planet increases. For example, a $20M_{\oplus}$ planet with a 30% water layer and $T_0 = 300$ K has $\log g = 3.50$ and $R_{\text{ocean}} = 59$ km.

5.4.2 Potential for Liquid Water on Mini-Neptunes

We now consider the extent to which planets possessing H/He envelopes may host liquid H₂O at their surface. Assuming that the H/He and H₂O layers of the planet do not mix, the pressure and temperature at the H₂O-H/He boundary (HHB) depends on the H/He mass fraction and the atmospheric temperature profile (Madhusudhan et al., 2020). This in turn determines the phase of H₂O at the HHB. The liquid phase of H₂O is accessible at temperatures up to 647 K and pressures up to 7.3×10^4 bar (see Figure 5.3).

Madhusudhan et al. (2020) demonstrated that the mini-Neptune K2-18b, despite having a mass and radius indicative of a planet with a substantial gaseous envelope, could have liquid water at the HHB. This scenario arises in the case where K2-18b consists mostly of H₂O, with a small Earth-like nucleus and a low but non-negligible H/He mass fraction. Figure 5.11 shows the internal structure of one such model planet, the equivalent of Case 3 from Madhusudhan et al. (2020), which has a nucleus mass fraction of 10% and a H/He mass fraction of 6×10^{-5} , with the rest of the planet consisting of H₂O. The phase structure of the H₂O layer is also shown.

These results hint at a region in M – R space where planets must host non-negligible gaseous envelopes to explain their inflated radii, but could still possess a liquid H₂O layer. Here we aim to determine the extent of this region of the parameter space and investigate which, if any, other observed planets fall into this category. To begin, we construct the M – R curve for a theoretical planet with an Earth-like nucleus of 10% and a H₂O layer of 90% by mass, with a surface temperature of 300 K. We consider this to be an upper limit to the radius of a planet with surface liquid water and no H/He envelope. We choose this composition since a 100% H₂O planet is unrealistic from a planet formation perspective (see Section 5.3.2.1). Next we give the planet a H/He envelope, which we assume is fully differentiated from the H₂O layer (see Section 5.4.4 for a discussion of mixed envelopes). The pressure

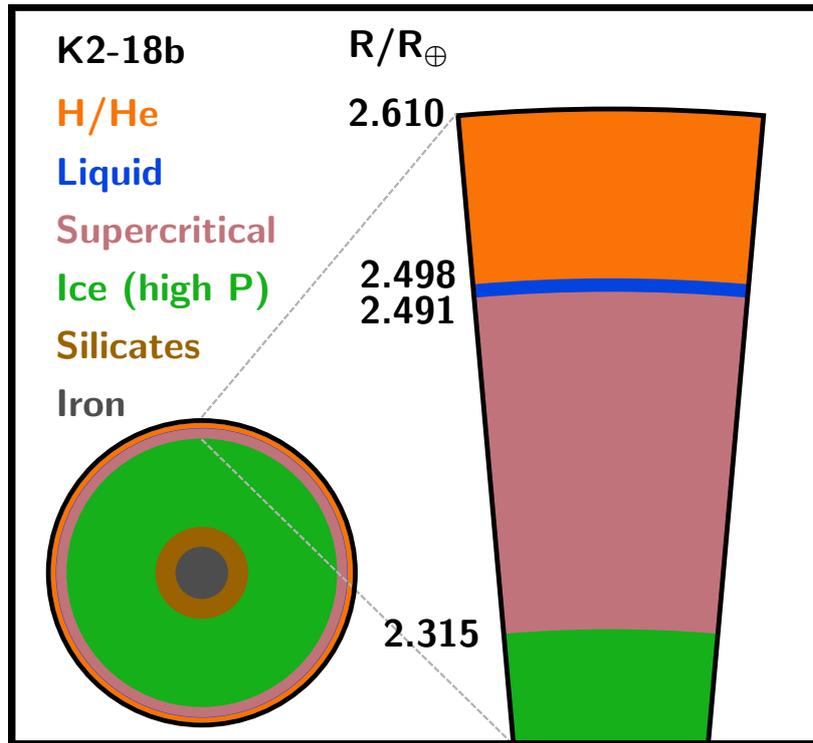


Fig. 5.11 Internal structure for one possible interior composition of K2-18b, corresponding to Case 3 from Madhusudhan et al. (2020) with a H/He mass fraction of 6×10^{-5} , an Earth-like nucleus of 10% by mass and a $\sim 90\%$ H₂O layer. A layer of liquid water is present directly beneath the H/He envelope.

and temperature at the HHB increases with the mass of the envelope, since we assume a temperature profile consisting of an isotherm and an adiabat; i.e., the temperature either remains constant or increases monotonically with depth (see Section 5.3.2.2). Therefore, for a habitable-zone temperature planet with a fixed core mass fraction, there must be a maximum amount of H/He that it can possess before the HHB becomes too hot for liquid water. For a given mass and temperature profile we can solve for this H/He mass fraction and therefore obtain the maximum radius of a planet that could host liquid water.

We find solutions for a finely-spaced grid of masses and a range of temperature profiles. For illustration, we consider planets from $1\text{--}20M_{\oplus}$ and nominally assume a photospheric pressure $P_0 = 0.1$ bar, with photospheric temperatures T_0 ranging from 100–647 K and P_{rc} from 1–1000 bar; the parameter ranges are motivated by the $P\text{--}T$ profiles shown for a wide range of mini-Neptune atmospheres in Piette & Madhusudhan (2020b), as discussed in Section 5.2.4 and Figure 5.6. Although we consider photospheric temperatures up to the critical temperature of H₂O, i.e. the maximum temperature at which liquid H₂O can exist,

we note that for planets with $T_0 \geq 319$ K at our chosen P_0 an inflated radius could also be attributed to a steam atmosphere (see Section 5.4.3).

Our results are shown in Figure 5.12, which shows the region of M – R space where a planet hosting a H/He envelope could possess a liquid water ocean. We also show measured masses and radii of planets near this region. We find that for a given T_0 and P_{rc} , there is a maximum total mass of H/He envelope, M_{env} , that can allow for liquid water at the HHB. This maximum M_{env} increases as the HHB becomes deeper, meaning the shaded region in Figure 5.12 represents planets with a maximal $P_{\text{rc}} = 1000$ bar. For planets whose atmospheres have a radiative-convective boundary at a lower pressure, the temperature starts to increase from higher in the atmosphere. Therefore, the envelope must be smaller to maintain liquid water at the HHB. For example, consider a planet with $T_0 = 300$ K. If $P_{\text{rc}} = 1000$ bar, then a $10M_{\oplus}$ planet can host up to $7.0 \times 10^{-2}M_{\oplus}$ of H/He while retaining liquid water at the HHB. However, if $P_{\text{rc}} = 10$ bar, then the maximal $M_{\text{env}} = 4.1 \times 10^{-4}M_{\oplus}$. The mass of envelope permitted also decreases as the photospheric temperature increases, following an approximate power law behaviour: for example, for $M_p = 10M_{\oplus}$ and $P_{\text{rc}} = 10$ bar we have $(M_{\text{env}}/M_{\oplus}) \approx 3.3 \times 10^6 (T_0/K)^{-4}$. M_{env} only weakly depends on the total planet mass M_p : for example, at $T_0 = 300$ K, $P_{\text{rc}} = 10$ bar, a $20M_{\oplus}$ planet has at most $M_{\text{env}} = 3.4 \times 10^{-4}M_{\oplus}$ while a $1M_{\oplus}$ planet has up to $M_{\text{env}} = 5.8 \times 10^{-4}M_{\oplus}$ while retaining a surface ocean. This behaviour also follows an approximate power law, with $M_{\text{env}} \approx 6 \times 10^{-4} M_p^{-0.18}$ in this case. Despite the maximum permissible mass of the H/He envelope decreasing with increasing T_0 , the increase in radius is approximately the same across all values of T_0 considered. Although a hotter H/He envelope is less dense and therefore should be more inflated, the maximum P_{HHB} that can permit a water ocean also decreases as T_0 increases, cancelling out the increase in radius from the warmer, less dense envelope. The maximum permitted H/He envelope across all temperature profiles considered increases the radius of the planet by 0.23–1.19 R_{\oplus} , with larger envelopes seen for lower mass planets due to the smaller surface gravity.

Our results hold for a H_2O mass fraction of 0.9 and can be treated as an upper limit for the radius of a planet with a liquid water ocean. For planets with less H_2O relative to the nucleus, the increase in radius permitted from a H/He envelope while retaining a liquid HHB is smaller, since the planet has a higher surface gravity. For example, in the 90% H_2O case, a $10M_{\oplus}$ planet has $R_p = 2.51R_{\oplus}$ and therefore $\log g = 3.192$ (cgs). For $P_{\text{rc}} = 10$ bar, the maximum increase in radius allowed from a H/He envelope while retaining an ocean is $0.238R_{\oplus}$. A planet with the same mass but with only 50% H_2O (and $\sim 50\%$ nucleus) has $\log g = 3.287$, and the maximum increase in radius for this planet is $0.156R_{\oplus}$.

From Figure 5.12 we see that there are several planets other than K2-18b whose observed masses and radii could be explained with a small H/He envelope, possibly with liquid H_2O at

the HHB. However, most of these planets have equilibrium temperatures high enough that their radii could be also explained with a steam atmosphere. One planet which falls into this category is Kepler-20d (Gautier et al., 2012), which has a mass of $10.07^{+3.97}_{-3.70} M_{\oplus}$ and a radius of $2.744^{+0.073}_{-0.055} R_{\oplus}$, with an equilibrium temperature of 401 ± 13 K (Buchhave et al., 2016). A primarily H_2O planet with a steam atmosphere would be consistent with these measurements, as would a planet with a H/He envelope of up to a few per cent. Due to its surface temperature, whether such a planet could host a liquid H_2O layer is very sensitive to the choice of temperature profile: while there are solutions allowing for liquid water when using the temperature profiles considered here, the atmosphere of Kepler-20d has not been observed, and so it is not known whether the hydrogen-rich temperature profiles used in this section would be appropriate to model this particular planet. This problem will be alleviated as more mini-Neptune atmospheres are observed in the future.

5.4.3 Diversity of Water World Phase Structures

Our results thus far have focused on planets with a liquid water component, however in reality these planets represent just a fraction of the possible phase structures of water worlds. Here we discuss the wide variety of H_2O phase structures that may be present on water-rich planets. To achieve this we examine each phase of H_2O that can be found at the planet's surface and determine the various phase structures that the H_2O layer may possess depending on its internal temperature profile.

The phase structure of the H_2O layer of a planet depends strongly on the choice of internal temperature profile. Most of the planet's interior will be convective and so we use an adiabatic temperature profile for the liquid, ice and supercritical phases (Valencia et al., 2007, Thomas, 2016). We consider two end-member cases for vapour layers: an isothermal or an adiabatic profile. Since the true temperature profile of a steam atmosphere is likely to lie between these extremes, this should allow us to explore the set of possible phase structures for water-rich planets with steam atmospheres.

For illustration, we consider planets of $8M_{\oplus}$ consisting of a 70% H_2O layer by mass over an Earth-like nucleus (1/3 Fe, 2/3 $MgSiO_3$). We exhaustively search the parameter space of surface conditions to identify the various phase structures that can arise in water world interiors, and choose a representative case for each unique structure that we find. Each of these cases has a surface in one of three phases: ice, liquid or vapour. In the remainder of this section we discuss the structure for each case in turn. Table 5.2 and Figure 5.13 summarise the temperature structures of our representative model planets, and Figure 5.14 shows the resulting phase structures for each of these cases. In this figure we refer to ice Ih as

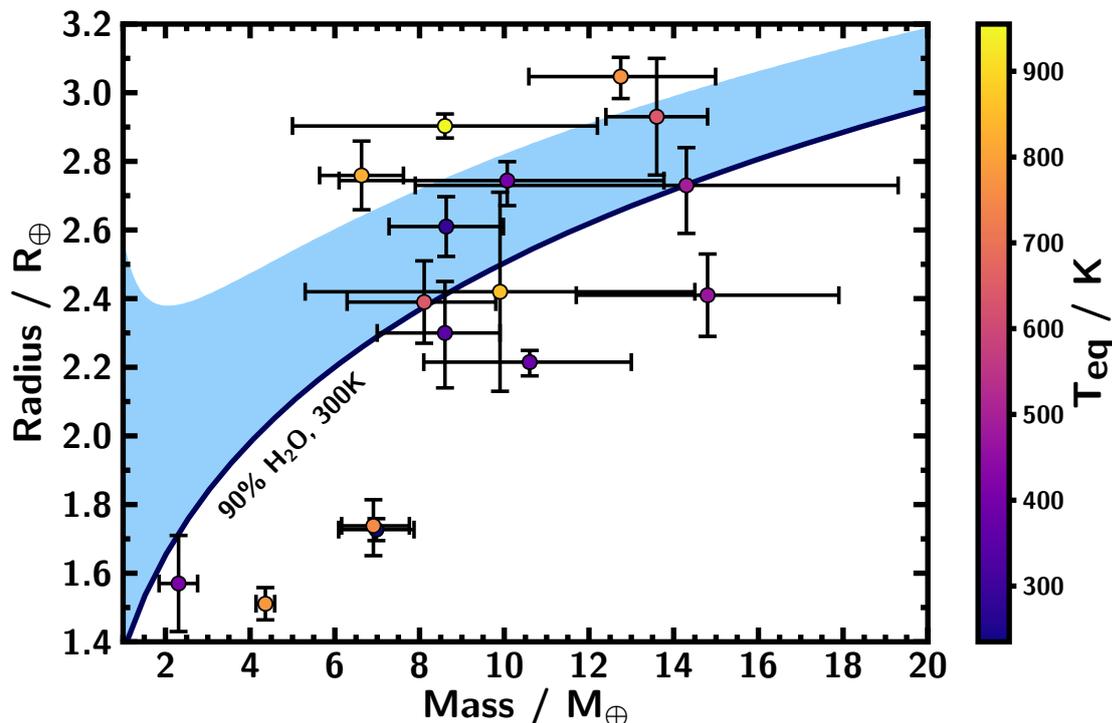


Fig. 5.12 M – R diagram showing the range of planet masses and radii consistent with a H/He envelope and liquid H₂O at the HHB (light blue shaded region). Temperatures at the HHB can lie anywhere within the liquid phase of water, which extends to 647 K at high pressures. The dark blue line shows a theoretical M – R relation for a planet consisting of a 10% Earth-like nucleus and a 90% H₂O layer with a 300 K surface temperature, which we take as the upper limit for a planet with no H/He envelope. We also show planets whose masses and radii have been reported with at least 2σ confidence, with $T_{\text{eq}} \leq 1000$ K. T_{eq} is indicated by the colour of each planet. Data on planetary masses, radii and equilibrium temperatures is taken from the NASA Exoplanet Archive.

"low-pressure ice" and other ice phases as "high-pressure ice". The various cases are now discussed in more detail.

Case 1: Ice surface

Planets with ice Ih surfaces may have one of two different phase structures: they can remain in ice throughout the interior (Case 1a) or host a liquid layer in between low- and high-pressure ices (Case 1b). A completely icy structure prevails in model planets with $T_0 < 251.2$ K. This is represented by Case 1a, which has a surface temperature of 220 K and surface pressure of 1 bar. The interior passes through ices Ih, III, V, VI, VII and X. However, a hotter surface temperature that is closer to the ice-liquid phase boundary yields a sub-surface

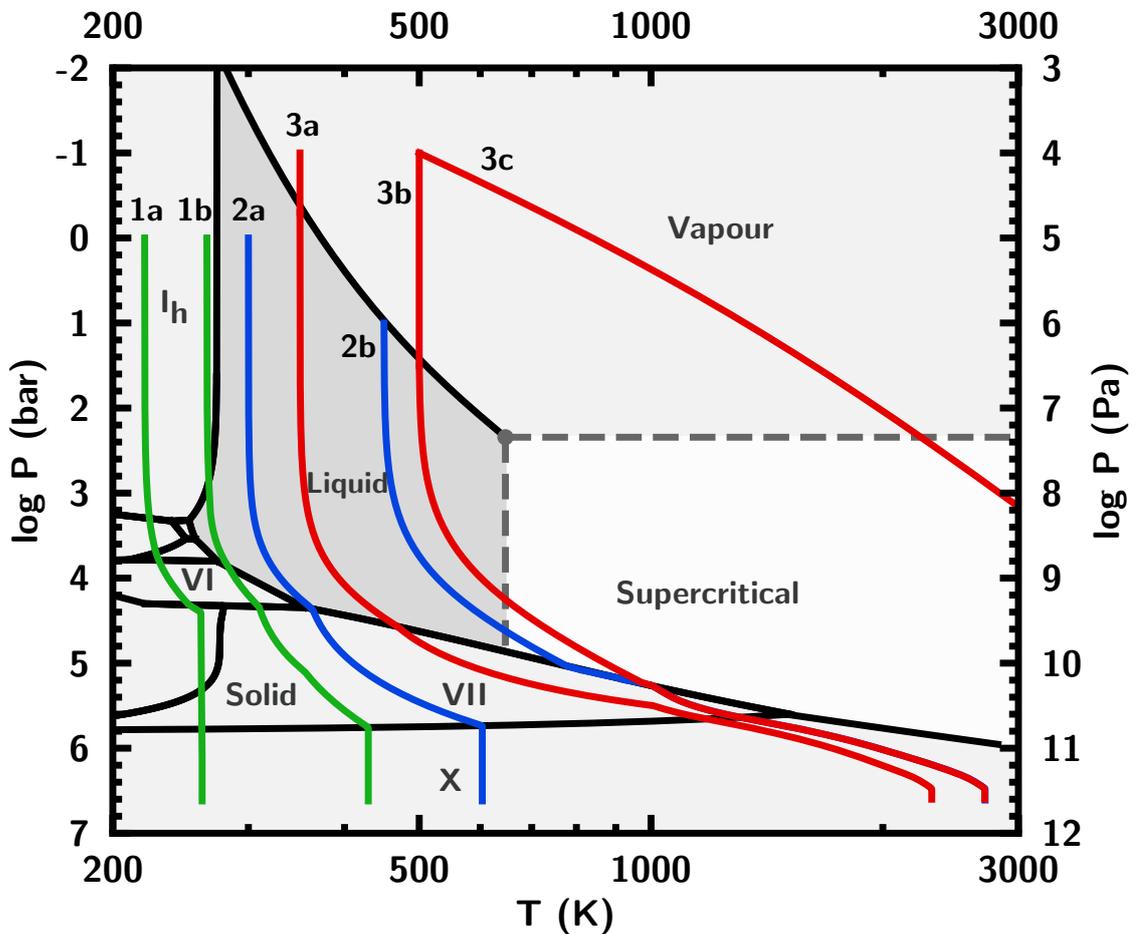


Fig. 5.13 P - T profiles of the planets shown in Figure 5.14. An adiabatic profile is used for all phases other than vapour. In the vapour phase, an isothermal profile is used in cases 3a and 3b, and an adiabatic profile is used in case 3c.

ocean in between ices I_h and the high-pressure ices (in this case ices VI, VII and X). This is demonstrated with Case 1b, for which the model inputs are identical to 1a except for a higher T_0 of 270 K. A structure similar to Case 1b has been proposed previously for water-rich exoplanets at large orbital separations (Ehrenreich et al., 2006). Allowing for deviations from an adiabatic temperature profile by considering processes such as conduction in the ice layers could heat the interior further, increasing the probability of a subsurface ocean.

Case 2: Liquid surface

Water-rich planets with a liquid surface host predominantly icy H_2O layers, either moving directly from liquid to high-pressure ice (Case 2a) or hosting a supercritical layer in between

Table 5.2 Surface conditions and radiative convective boundaries for the model planets considered in Section 5.4.3. Each model planet represents a unique phase structure that a water-rich sub-Neptune may possess.

Case	Phase at R_p	T_0 (K)	$\log P_0$ (bar)	$\log P_{rc}$ (bar)
1a	Ice	220	0	0
1b	Ice	270	0	0
2a	Liquid	300	0	0
2b	Liquid	450	1	1
3a	Vapour	350	-1	-0.3
3b	Vapour	500	-1	1.5
3c	Vapour	500	-1	-1

(Case 2b). For most surface conditions, the H₂O layer transitions directly from liquid to ices VII and X, as demonstrated by Case 2a which has $T_0 = 300$ K and $P_0 = 1$ bar. In a small number of cases where T_0 and P_0 are large it is possible to obtain a layer of supercritical H₂O between the liquid and ice phases. This is shown in Case 2b which has $T_0 = 450$ K, $P_0 = 10$ bar.

Case 3: Steam atmosphere

For planets whose photospheric pressure and temperature correspond to the vapour phase, leading to a steam atmosphere, the underlying phase structure depends strongly on the atmospheric temperature profile, with interiors consisting of liquid water (Case 3a), supercritical water (Case 3c), or both (Case 3b) before transitioning to high-pressure ice. In order to demonstrate this we consider end-member cases of a purely isothermal and purely adiabatic profile, each with the same photospheric pressure, $P_0 = 0.1$ bar. For a planet to host a surface ocean, the temperature throughout the envelope must remain below the critical temperature of water, 647 K. Cases 3a and 3b demonstrate this for two different values of T_0 : in Case 3a, which has a lower T_0 of 350 K, the interior moves directly from liquid to high-pressure ice, giving a similar structure to Case 2a underneath the atmosphere. By contrast, the hotter $T_0 = 500$ K used for Case 3b yields an additional layer of supercritical H₂O, which is similar to Case 2b. Finally, if an adiabatic temperature profile in the atmosphere is assumed, the vapour-liquid transition is forbidden and so the interior changes directly from vapour to supercritical H₂O, as shown in Case 3c. This is also the case for any planet with $T_0 \geq 647$ K. Note that this also inflates the radius of the planet significantly.

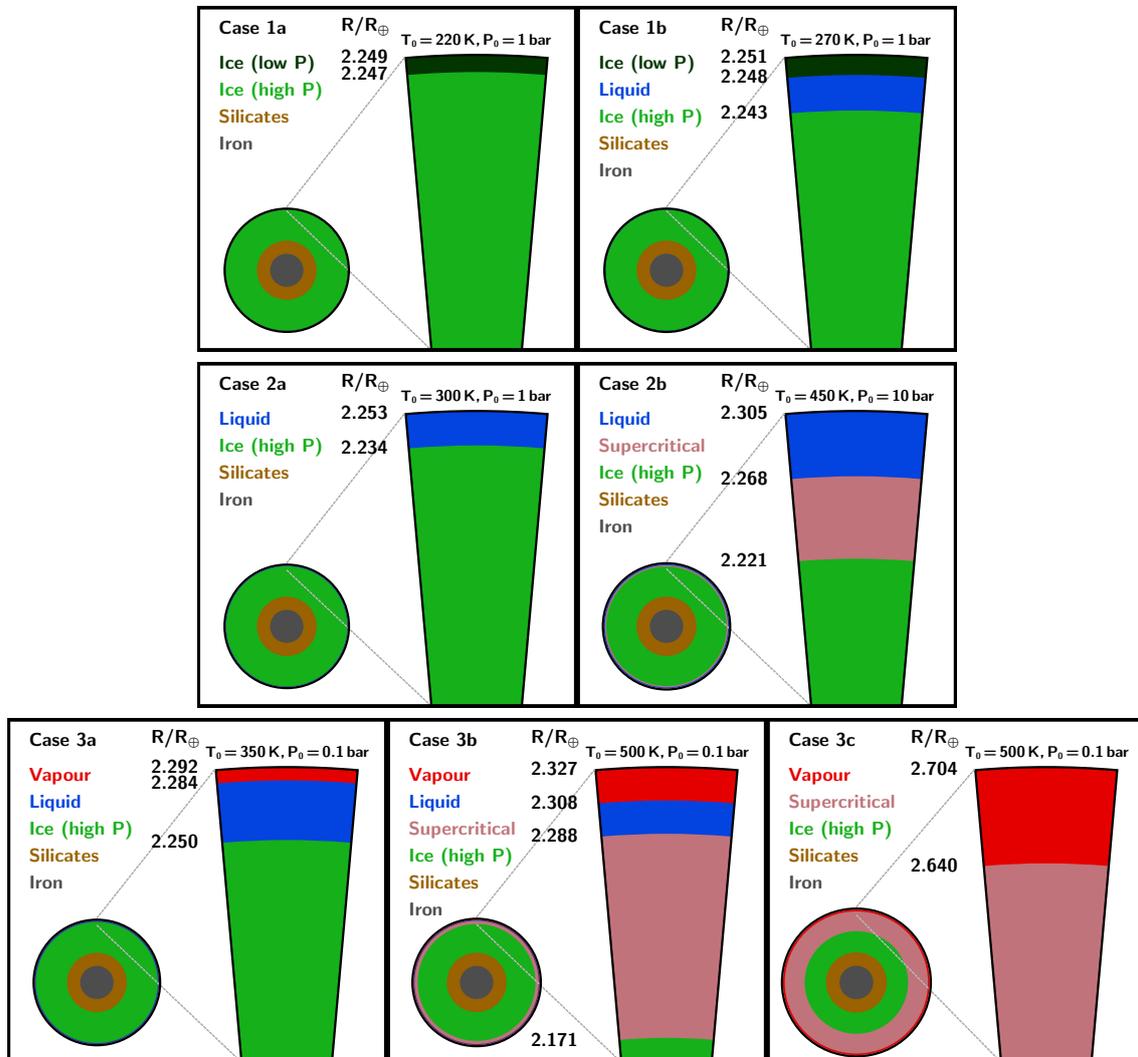


Fig. 5.14 Phase structures of H_2O -rich planets with different surface conditions and temperature structures. Each of these planets has the same mass ($8M_{\oplus}$) and H_2O mass fraction (0.7), but different surface conditions. The P - T profiles of the H_2O layers of these planets are shown in Figure 5.13. We find that a diverse set of phase structures are possible depending on the surface conditions and temperature profile, from a sub-surface ocean between low- and high-pressure ice (Case 1b) to a steam atmosphere above a layer of supercritical water (Case 3c).

5.4.4 Planets with Mixed Envelopes

So far we have considered models where each layer in the planetary interior is fully differentiated. However, it is possible that sub-Neptunes could host envelopes consisting of mixed H_2O and H/He , as has been suggested for the interiors of giant planets (Soubiran & Militzer, 2015). H_2O is regularly detected in exoplanet atmospheres, and mass–metallicity trends derived from atmospheric observations suggest that lower-mass planets should have higher H_2O abundances (Welbanks et al., 2019). In this section we aim to determine the difference between M – R relations for mixed and unmixed H_2O – H/He envelopes, where the envelope is defined here to mean the outer H_2O and H/He components of the planet. This will indicate whether interior structure models need to take into account the presence of atmospheric species as well as highlighting possible degeneracies between mixed and unmixed envelopes of different compositions. We produce M – R relations for mixed H_2O – H/He envelopes with different quantities of each component, and compare these to the equivalent unmixed relations. The methods for modelling mixed envelopes are described in Section 5.2. The helium mass fraction within the H/He component is held constant throughout ($Y = 0.275$).

For all models considered here we take $T_0 = 500$ K, $P_0 = 0.1$ bar and $P_{\text{rc}} = 10$ bar. This ensures that the temperature profile is sufficiently hot for the H_2O to be in vapour or supercritical phase throughout the envelope, where miscibility with H/He is more likely. The model planets consist of an Earth-like nucleus with a mass fraction of 0.95. The remaining 5% of the planetary mass is divided between H_2O and H/He in different proportions, ranging from a pure H_2O to a pure H/He envelope. For each composition we produce M – R relations for both a mixed and an unmixed envelope.

The resulting M – R relations are shown in Figure 5.15. We find that small amounts of H_2O in a H/He -rich envelope do not significantly alter a planet’s radius. A planet with a mixed envelope containing 1% H_2O by mass has a radius that is on average $0.028R_{\oplus}$ lower than a planet of the same total mass with a pure H/He atmosphere, which is small compared to the measurement uncertainties of sub-Neptune radii. An unmixed envelope with 1% H_2O lies in between these cases. This mass fraction corresponds to a volume mixing ratio of $\log X_{\text{H}_2\text{O}} \approx -2.9$, close to solar H_2O abundance at this temperature, $\log X_{\text{H}_2\text{O}} = -3.0$ (Asplund et al., 2009, Madhusudhan, 2012), and suggests that when modelling the interior structures of planets with hydrogen-dominated atmospheres that may contain small amounts of other chemical species, the assumption of a pure H/He outer envelope provides a sufficient density profile. Conversely, while the difference between a pure H_2O envelope and a mixed envelope with 1% H/He is also minimal, the radius increases substantially when considering an unmixed rather than a mixed envelope. In this case, switching from a mixed to unmixed

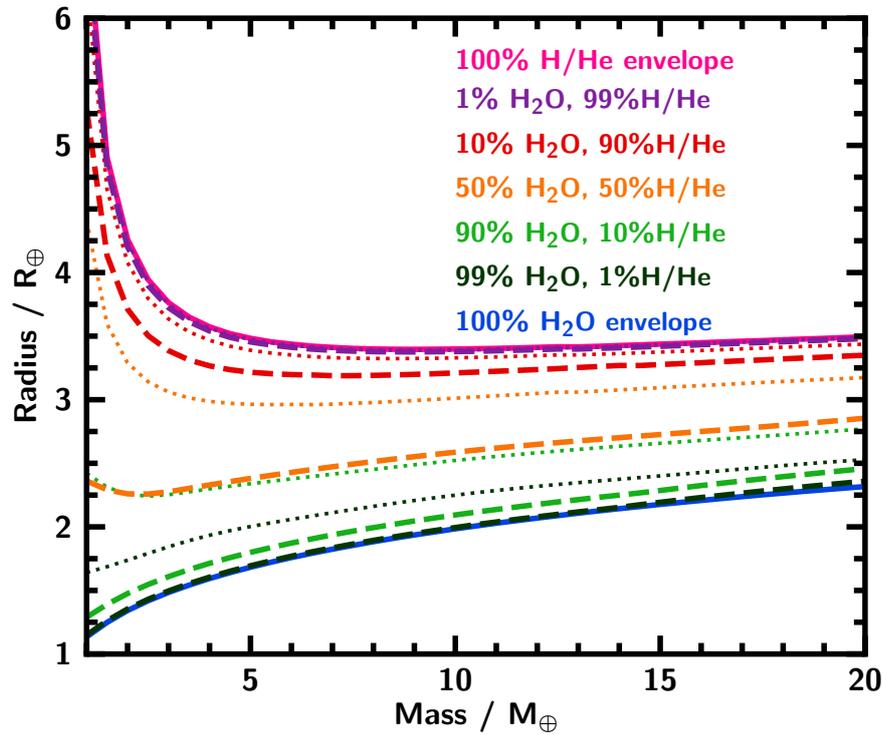


Fig. 5.15 M – R relations for planets with mixed (dashed) and unmixed (dotted) envelopes of fixed mass fraction ($x_{\text{env}} = 0.05$) above an Earth-like nucleus. All models have $T_0 = 500$ K, $P_0 = 0.1$ bar and $P_{\text{rc}} = 10$ bar. The solid blue and pink lines show end-member cases of a pure H_2O and H/He envelope respectively. For H/He-rich envelopes with small amounts of H_2O , there is little difference between the mixed and unmixed cases. However, for planets with mostly H_2O envelopes, mixing in the envelope leads to a much smaller radius than the unmixed case.

envelope drastically changes the planet’s structure: a mixed envelope has a mostly steam atmosphere throughout, whereas the unmixed case with $T_0 = 500$ K has a H/He atmosphere over a layer of supercritical water. It is worth noting that the M – R relation for a planet with a mixed envelope consisting of equal amounts of H/He and H_2O is very close to the relation for an unmixed envelope with 90% H_2O , 10% H/He. This provides another source of uncertainty when constraining a planet’s composition and interior structure. However, this can potentially be resolved using atmospheric observations.

5.5 Summary and Discussion

We have investigated the internal structures of water-rich exoplanets, motivated by the recent suggestion of the possibility of a habitable water ocean under a H/He atmosphere in the mini-Neptune K2-18b (Madhusudhan et al., 2020). We have presented a modelling framework for sub-Neptune planets, spanning super-Earths and mini-Neptunes, and used this to explore the H₂O phase structures of such planets in detail. We have found the following key results:

- For planets with a liquid H₂O layer, the depth of the ocean is determined by the planet’s surface gravity and temperature, and can extend up to ~ 1600 km. Planets with gravities comparable to Earth’s can have oceans that are over one hundred times deeper than the Earth’s ocean, which has an average depth of ~ 3.7 km.
- Planets with H/He envelopes can allow for significant liquid H₂O layers underneath the envelope. This phenomenon can occur over a wide region in M – R space. For example, liquid H₂O may exist on such planets with $T_0 = 300$ K, $P_{rc} = 10$ bar provided the total mass of the H/He envelope is $\lesssim 4 \times 10^{-4} M_{\oplus}$. The mass of H/He permitted increases as the photospheric temperature decreases and as the radiative-convective boundary moves deeper in the atmosphere.
- More broadly, a diverse range of phase structures are possible in H₂O-rich super-Earths and mini-Neptunes. Besides surface liquid H₂O, other phase structures include liquid water sandwiched between two ice layers or a steam atmosphere above supercritical water and high-pressure ice, depending on surface conditions.
- Planets with mixed H/He-H₂O envelopes have significantly lower radii than planets with the same composition but differentiated H/He and H₂O if the mass fractions of the two components are comparable, but this effect is minimal for a small mass fraction of H₂O ($\sim 1\%$) in a mostly H/He atmosphere.

5.5.1 Potential Habitability of Water Worlds

Traditionally, a planet is classed as habitable if it has the right conditions for liquid water at its surface (Hart, 1979, Kasting et al., 1993). Living organisms have been found to survive in liquid water at $T \lesssim 400$ K and $P \lesssim 1000$ bar (Merino et al., 2019). Some of the possible interior structures found for K2-18b in Madhusudhan et al. (2020) had thermodynamic conditions in this range at the surface of the water layer. Many of the model planets considered in this work also fit this definition, highlighting the possibility of a family of habitable planets in the mini-Neptune regime.

It is important to consider whether other factors could preclude habitability for water worlds. All of the planets with H₂O oceans shown in Section 5.4.3 have very deep layers of high-pressure ice which separate the ocean from the silicate mantle. Kitzmann et al. (2015) found that this separation, which prevents a carbonate-silicate cycle that can regulate the CO₂ inventory of a planet, can have a destabilising effect on the climate with negative consequences for habitability. However, other CO₂ exchange mechanisms have been proposed that may have a regulatory effect (Levi et al., 2017, Ramirez & Levi, 2018). The potential for habitable conditions on H₂O-rich planets should therefore not be ruled out despite the lack of an Earth-like carbonate-silicate cycle.

In Section 5.4.1 we consider ocean depths for surface temperatures ranging from 273–584 K. Past studies have been restricted to lower temperatures; for example Noack et al. (2016) explores temperatures from 290–370 K, stating that for higher temperatures the liquid water in the ocean would evaporate due to the runaway greenhouse effect. However, this rests on the assumption of a surface pressure of 1 bar similar to Earth. By assuming a surface pressure of 100 bar we can explore higher temperatures, and we find that the increase in ocean depth with surface temperature reported by Noack et al. (2016) no longer holds beyond $T_0 = 413$ K. The pressure at the surface of the ocean on a water world is not known and could be greater than 1 bar, especially if the planet possesses a H/He envelope. Relaxing this assumption allows for a widening of the parameter space for water worlds with global oceans and suggests that liquid water layers could be a possibility even for planets considerably warmer than Earth at the bottom of their atmospheres.

In Figure 5.12 we show a region of M – R space where planets require H/He envelopes but could still host liquid water oceans. This region is found assuming a maximal H₂O mass fraction of 90%. While it is understood that a minimal amount of rocky material is necessary to form a planet by core accretion, the exact range of possible ice/rock ratios for exoplanets is not known, and estimates vary widely. It is possible that a planet consisting of as much as 90% H₂O and only ~10% iron and silicates could be unfeasible from a planet formation perspective. If this is the case, a region of M – R space similar to the one described here will still exist, but at smaller radii due to the lower maximum radius of a planet without any H/He. Whether planets that fall into this category are indeed habitable depends on what effect a hydrogen-rich atmosphere has on a planet's climate, which is still poorly understood. However, a number of studies have argued in favour of rocky planets with H/He atmospheres being potentially habitable (Pierrehumbert & Gaidos, 2011, Koll & Cronin, 2019) and other works (e.g. Seager et al., 2020) have shown that single-celled organisms can survive in a pure H₂ atmosphere. The potential for habitability of sub-Neptunes with H₂-rich atmospheres

is considered in several recent and upcoming works (Madhusudhan et al., 2020, Piette & Madhusudhan, 2020b, Madhusudhan et al., in press).

5.5.2 Future Directions and Applications

There are many challenges to overcome in order to characterise water-rich exoplanets in detail, both from a theoretical and observational perspective. Our model makes a number of standard assumptions that may be challenged in the future. For example, while the vast majority of internal structure models assume an adiabatic temperature profile throughout most of the interior, thermal boundary layers that inhibit convection have been proposed for the interiors of Uranus and Neptune (Nettelmann et al., 2016, Podolak et al., 2019). Additionally, while we consider mixed H/He-H₂O envelopes in this study, other chemical species will be present in planetary interiors that might affect their internal structure. Levi et al. (2014) modelled water-rich planets with a methane component, and while they found that the inclusion of methane did not significantly change the $M-R$ relation, they did see a noticeable effect on the atmosphere. Shah et al. (2021) modelled planets with hydrated interiors, and while the impact on $M-R$ relationships was small compared to current measurement uncertainties, future missions such as PLATO (Rauer et al., 2016) may lead to a higher precision for masses and radii that could make these effects important to consider. Many aspects of the behaviour of high-pressure ices, such as possible interactions between rock and ice in this regime, are still unknown (Journaux et al., 2020b, Huang et al., 2021), and further understanding in this area could have important consequences for the potential habitability of ocean worlds (Noack et al., 2016).

Characterising observed sub-Neptunes using internal structure models is an inherently degenerate problem, with many different compositions consistent with a given mass and radius. For many planets in this regime, it is impossible to distinguish between a water world scenario and a rocky planet with a thick H/He envelope (e.g. Luque et al., 2021). Determining the atmospheric compositions of these planets may allow us to break this degeneracy: while a small number of mini-Neptune atmospheres have been characterised using spectroscopic observations (e.g. Benneke et al., 2019), it is hoped that upcoming facilities such as JWST will provide the opportunity to observe the atmospheres of low-mass planets in much greater detail than has previously been possible (e.g. Morley et al., 2017b, Welbanks & Madhusudhan, 2021), which should in turn allow better constraints on the interior composition and structure.

The results from this study, as well as those of Madhusudhan et al. (2020) and Piette & Madhusudhan (2020b), highlight the diversity of exoplanets that have the potential to be habitable. These results point to the exciting possibility of the right conditions for life being present on planets much larger than Earth. As next-generation instruments make potential

ocean planets more amenable to characterisation and facilitate the search for biosignatures, we hope that our findings can further motivate the quest to detect signs of life on other worlds, even those which bear little resemblance to our own.

Chapter 6

Internal Structures of JWST Targets in the Sub-Neptune Regime

“You know, I don’t think there’s such a thing as ‘space earth’. If it’s from space, then it’s not really earth.”

– Katara

6.1 Introduction

The previous chapter highlighted the wide diversity of possible structures of sub-Neptune exoplanets. In order to further understand which of the scenarios presented in that work are actually present in nature, we need to acquire high-quality observations of the atmospheres of these planets. To date, informative observations of sub-Neptune atmospheres have proved challenging, often returning featureless spectra which may be a result of high mean molecular weight atmospheres leading to a small scale height, or high-altitude clouds and hazes (e.g. Kreidberg et al., 2014a).

However, our understanding of the atmospheric and interior properties of sub-Neptunes is set to be revolutionised with JWST. Simulated JWST observations indicate that information-rich spectra of sub-Neptune atmospheres can be obtained with several different JWST instruments, even in the presence of clouds (Constantinou & Madhusudhan, 2022). In order to maximise the signal acquired from spectroscopic observations, it is optimal to target planets orbiting M type stars, since the low masses, sizes and temperatures of these stars mean that the planet–star contrast is favourable for planetary detection and characterization. Recent detection surveys indicate that occurrence rates of low-mass planets orbiting M dwarfs are high (Dressing & Charbonneau, 2015, Mulders et al., 2015), providing numerous

opportunities to find small planets that are ideal for characterisation. This opportunity has been recognised with a large number of approved JWST Guaranteed Time Observations (GTO) and Cycle 1 General Observers (GO) programs dedicated to observations of spectra of sub-Neptunes orbiting M-dwarfs.

In this chapter we present efforts to characterise the internal structures of a number of sub-Neptunes that will be targeted by upcoming JWST observations¹. In Section 6.2 we present a study of the habitable-zone mini-Neptune K2-18b, a planet with a strong detection of atmospheric water vapour (Benneke et al., 2019, Tsiraras et al., 2019, Madhusudhan et al., 2020). We conduct a systematic study of the possible internal structures of the planet, including identifying the possible locations and thermodynamic conditions of the boundary between the planet’s H/He envelope and H₂O layer. We discuss implications for the possible origins of the planet as well as its potential habitability. In Section 6.3 we consider a pair of mini-Neptunes in orbit around the M-dwarf TOI-776. We constrain the possible interior compositions of both planets and discuss the limitations of our analysis based on extant data. Finally, in Section 6.4 we briefly summarise the planned JWST observations of each of these planets, and explain how these observations will enable us to better understand the nature of each of these worlds.

6.2 The Habitable-zone Exoplanet K2-18b

The habitable-zone transiting exoplanet K2-18b is a particularly good example of a planet orbiting an M-dwarf that is well-suited to detailed characterisation (Foreman-Mackey et al., 2015, Montet et al., 2015). The brightness and small size of its host star make precise measurements of the planetary mass, radius, and atmospheric spectra viable (Benneke et al., 2017, Cloutier et al., 2019). The habitable-zone temperature of K2-18b provides further impetus to uncover the nature of its interior and atmosphere.

Given its mass ($M_p = 8.63 \pm 1.35 M_\oplus$, Cloutier et al. 2019) and radius ($R_p = 2.610 \pm 0.087 R_\oplus$, Benneke et al. 2019), K2-18b has a bulk density ($2.67^{+0.52}_{-0.47} \text{ g/cm}^3$, Benneke et al. 2019). This density, between that of Earth and Neptune, may be thought to preclude a purely rocky or icy interior and require a hydrogen-rich outer envelope. However, the extent of such an envelope and the conditions at the interface between the envelope and the underlying interior

¹The contents of this chapter are based on published work (Madhusudhan et al., 2020, Luque et al., 2021). The text in this chapter is adapted from my contributions to the text of these two papers. I contributed the H₂O equation of state for the interior models of K2-18b, and prepared figures 6.1–6.5. I used my internal structure model (presented in Chapter 5) to characterise TOI-776b and c, and prepared figures 6.6 and 6.7. The co-authors of these publications also contributed to the text and figure preparation for the manuscripts from which the contents of this chapter are derived.

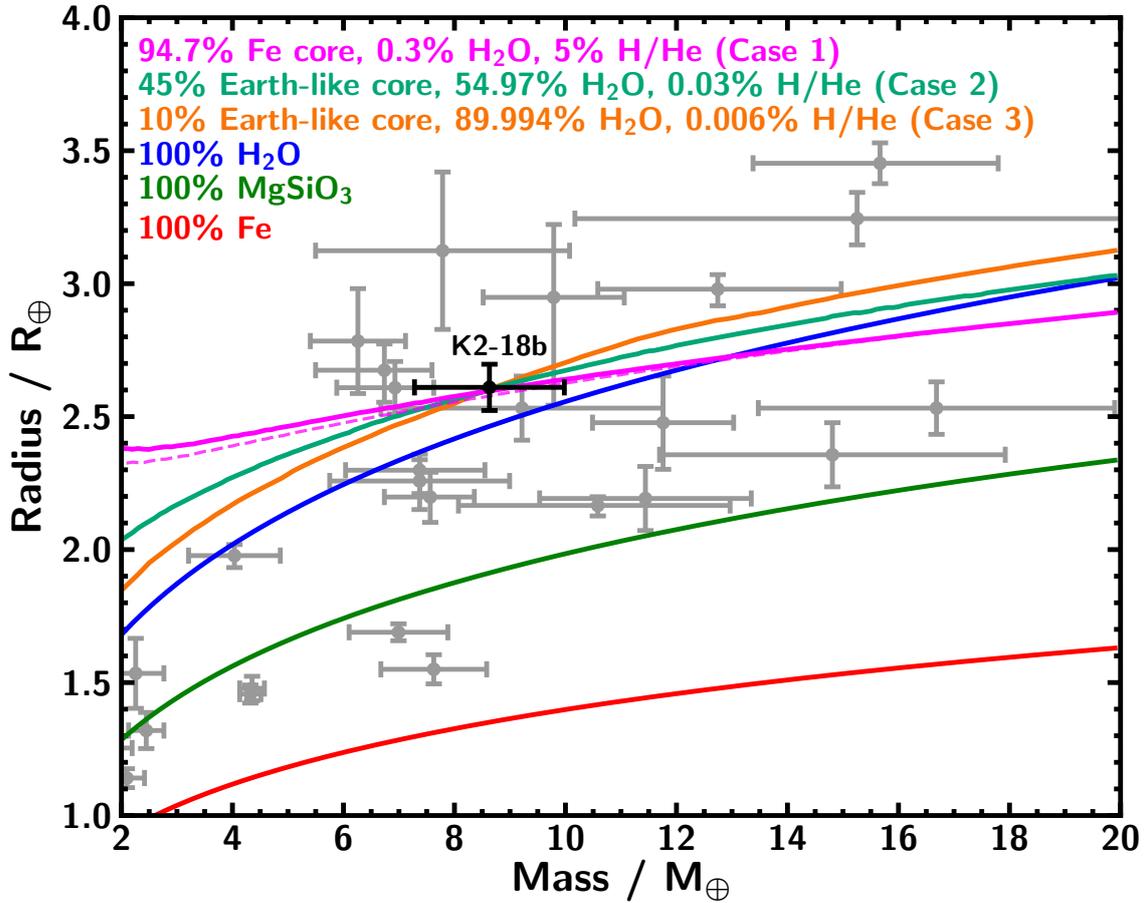


Fig. 6.1 Mass–radius diagram. The solid magenta, teal, and orange curves show interior models with three representative compositions that fit the mass and radius of K2-18b. The dashed magenta line represents the same composition as the solid magenta line, but with a mixed H₂O-H/He envelope. Also shown are exoplanets whose masses and radii are known to $\geq 3\sigma$ with $T_{\text{eq}} < 1000$ K, from TEPCat (Southworth, 2011).

have not been explored. We note that the mass and radius of the planet have recently been revised (Benneke et al., 2019), which may have impacted inferences made using previous values (Cloutier et al., 2017, Tsiaras et al., 2019).

Previous studies of planets with similar masses and radii, such as GJ 1214b, suggested envelope mass fractions $\lesssim 7\%$ (Rogers & Seager, 2010b, Nettelmann et al., 2011, Valencia et al., 2013). GJ 1214b is expected to host super-critical H₂O below the envelope at pressures and temperatures too high to be conducive for life (Rogers & Seager, 2010b). However, while GJ 1214b has an equilibrium temperature (T_{eq}) of ~ 500 K, K2-18b may be more favourable given its lower $T_{\text{eq}} \sim 250 - 300$ K.

The broadband spectrum of K2-18b was observed using HST WFC3 (Benneke et al., 2019), covering a wavelength range of 1.1–1.7 μm . Analyses of this spectrum along with data from *Kepler* and *Spitzer* has led to multiple high-confidence detections of H_2O in the atmosphere (Benneke et al., 2019, Tsiaras et al., 2019, Madhusudhan et al., 2020). For the purposes of this work, we use the abundance estimates presented in Madhusudhan et al. (2020), which are consistent to within 1σ with Benneke et al. (2019). The derived H_2O volume mixing ratio ranges between 0.02–14.80%, with median values of 0.7–1.6% depending on model assumptions. The derived H_2O abundance range corresponds to an O/H ratio of 0.2–176.8 \times solar, assuming all the oxygen is in H_2O as expected in H_2 -rich atmospheres at such low temperatures (Burrows & Sharp, 1999). These results cannot be directly compared to Tsiaras et al. (2019) as their retrievals used older measurements of the planetary mass and radius which could have biased their inferences.

Madhusudhan et al. (2020) find a depletion of CH_4 and NH_3 in the atmosphere, suggesting that the atmosphere is not in chemical equilibrium, and do not find strong evidence for clouds/hazes. The retrieved value of P_{ref} , the pressure at the observed planetary radius, is 12 – 174 mbar. The median value of 0.05 bar is used as the surface boundary condition, pressure P_0 , for the internal structure models in Section 6.2.1.

6.2.1 Constraints on Internal Structure

We model the interior of the planet with a canonical four-layer structure. The model comprises a two-component Fe+rock core consisting of an inner Fe layer and an outer silicate layer, a layer of H_2O , and an outer H/He envelope. Such a model spans the possible internal structures and compositions of super-Earths and mini-Neptunes (e.g. Valencia et al., 2010, 2013, Rogers et al., 2011, Lopez & Fortney, 2014), as well as terrestrial planets and ice giants in the solar system (Guillot & Gautier, 2014). The mass fractions of the different components (x_{Fe} , x_{rock} , $x_{\text{H}_2\text{O}}$, x_{env}) are free parameters in the model and sum to unity. Our present model is adapted from a three-layer model for super-Earths from Madhusudhan et al. (2012) comprising of iron, rock, and H_2O , with the H/He envelope added in the present work, and is very similar to the model described in Chapter 5.

The model solves the standard internal structure equations of hydrostatic equilibrium and mass continuity assuming spherical symmetry. The equation of state (EOS) for each of the two inner layers is adopted from Seager et al. (2007) who use the Birch-Murnaghan EOS (Birch, 1952) for Fe (Ahrens, 2000) and MgSiO_3 perovskite (Karki et al., 2000). For the H_2O layer we use the temperature-dependent H_2O EOS compiled by Thomas & Madhusudhan (2016) from French et al. (2009), Sugimura et al. (2010), Fei et al. (1993), Seager et al. (2007)

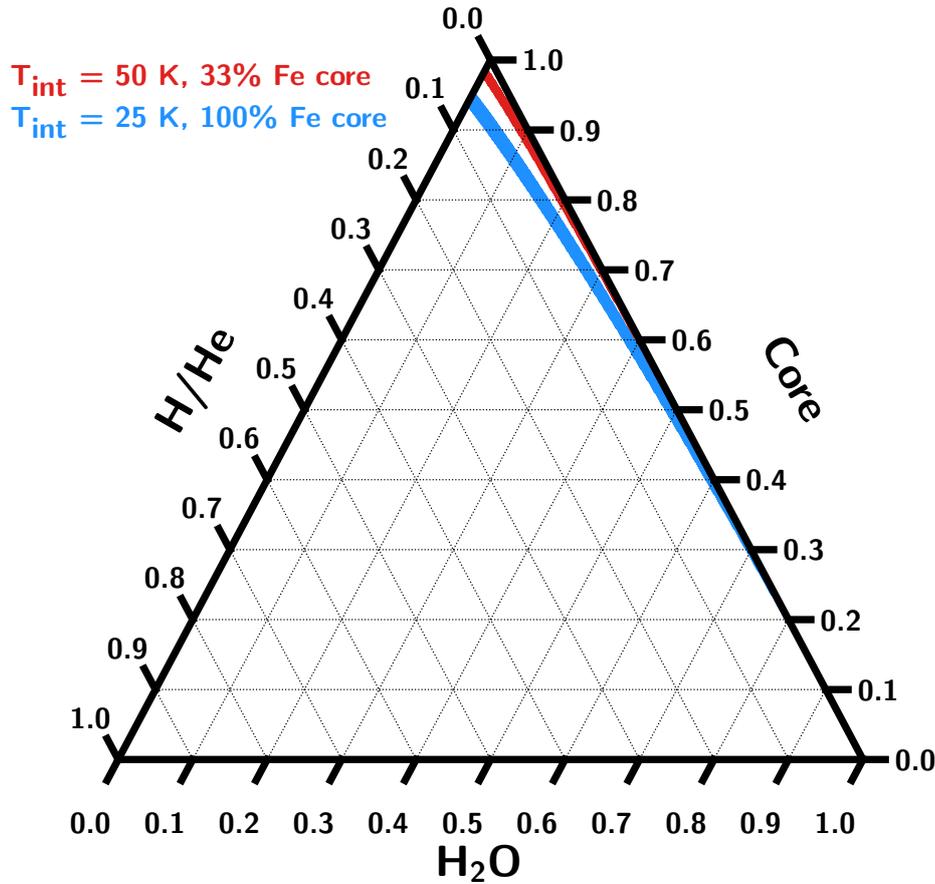


Fig. 6.2 Ternary diagram showing best-fitting ($\leq 1\sigma$) interior compositions allowed by the mass and radius of K2-18b for two end-member core compositions and interior temperatures.

and Wagner & Pruß (2002). For the gaseous envelope we use the latest H/He EOS from Chabrier et al. (2019) for a solar helium mass fraction ($Y = 0.275$).

The EOS in the H/He and H₂O layers can have a significant temperature dependence which we consider in our model. Past studies (Rogers et al., 2011, Valencia et al., 2013) considered analytic pressure–temperature (P – T) profiles for irradiated atmospheres derived using double grey approximations (Hansen, 2008, Guillot, 2010) with the internal and external fluxes and opacities as free parameters. Similarly to Chapter 5, we calculate self-consistent dayside P – T profiles for K2-18b in the H/He envelope using the GENESIS code (Gandhi & Madhusudhan, 2017). We include opacity due to H₂O (Rothman et al., 2010), as detected in the transmission spectrum, H₂ Rayleigh scattering, clouds and H₂–H₂ and H₂–He collision-induced absorption. Our assumed H₂O abundance is 10×solar (see Section 6.2) and we also use 10×solar abundances for the cloud species. We include KCl, ZnS and Na₂S clouds (Morley et al., 2013), for which we obtain opacities from Pinhas & Madhusudhan (2017). We further include water ice clouds using opacities from Budaj et al. (2015).

The P - T profile also depends on the planetary internal flux, which is characterised by the internal temperature T_{int} . We consider values of T_{int} which span the range expected for a planet with the mass and radius of K2-18b and an age of 1 – 10 Gyr, with envelope compositions from solar to water-rich. We choose end-member cases of $T_{\text{int}} = 25\text{K}$ and 50K , consistent with previous studies on planets of similar mass and radius, such as GJ 1214b (e.g. Valencia et al., 2013). The GENESIS models are calculated between pressures of $10^{-5} - 10^3$ bar, and assume full redistribution of the incident stellar irradiation. We explore a range of P - T profiles and choose two representative cases with different values T_{int} . Where required by the internal structure model, the bottom of the P - T profile of the H/He envelope is continued to deeper pressures using the adiabatic gradient from Chabrier et al. (2019). We also employ an adiabatic temperature profile in the H_2O layer.

Figure 6.1 shows mass–radius relations for models with different interior compositions. We explore the full range of plausible interior compositions in three components: $x_{\text{core}} = x_{\text{Fe}} + x_{\text{rock}}$, $x_{\text{H}_2\text{O}}$, and x_{env} , where $x_i = M_i/M_p$ is the mass fraction of each component i . For each atmospheric P - T profile considered, we explore two different core compositions: (1) an Earth-like core made of 33% Fe, 67% rock by mass, and (2) a pure Fe core, the densest possible composition. Here, we discuss results from two end-member cases: (1) a pure Fe core with $T_{\text{int}} = 25\text{K}$, and (2) an Earth-like (33% Fe) core with $T_{\text{int}} = 50\text{K}$. Solutions for all other cases lie between these two cases.

As shown in Figure 6.2, while a wide range of core and H_2O mass fractions are permitted, we place a stringent upper limit on the mass fraction of the H/He envelope: $x_{\text{env}} = 6.2\%$. This maximal x_{env} corresponds to the case of a pure Fe core, with $x_{\text{core}} \sim 94\%$, underlying the H/He envelope with no $x_{\text{H}_2\text{O}}$; here it is assumed that the atmospheric H_2O is not mixed in the envelope. However, if the retrieved atmospheric H_2O abundance is assumed to be well mixed in the envelope then the maximal $x_{\text{env}} = 6\%$ with $x_{\text{H}_2\text{O}} = 0.4\%$ by mass; low, but still significantly higher than that of the Earth’s oceans ($\sim 0.02\%$).

We find that a substantial gaseous H/He envelope is not necessary to explain the density of K2-18b. Figure 6.3 shows the x_{env} required for different x_{core} . At one extreme, a $\sim 100\%$ H_2O interior with no rocky core can explain the data with an x_{env} of just $\sim 10^{-6}$, comparable to the mass fraction of the Earth’s atmosphere. The presence of a rocky core would necessitate at least a thin H/He envelope. However, even considering a reasonable $x_{\text{core}} = 10 - 50\%$ still requires x_{env} of only $\sim 10^{-5} - 10^{-2}$, as shown in Figure 6.3. Model solutions with the hotter P - T profile and/or lower Fe content in the core require smaller x_{env} for a given x_{core} .

We have also considered models with miscible H_2O and H/He envelopes. We follow the approach of Soubiran & Militzer (2015), using an additive volume law for mixtures. Assuming that the median H_2O mixing ratio in the atmosphere is representative of the mixed

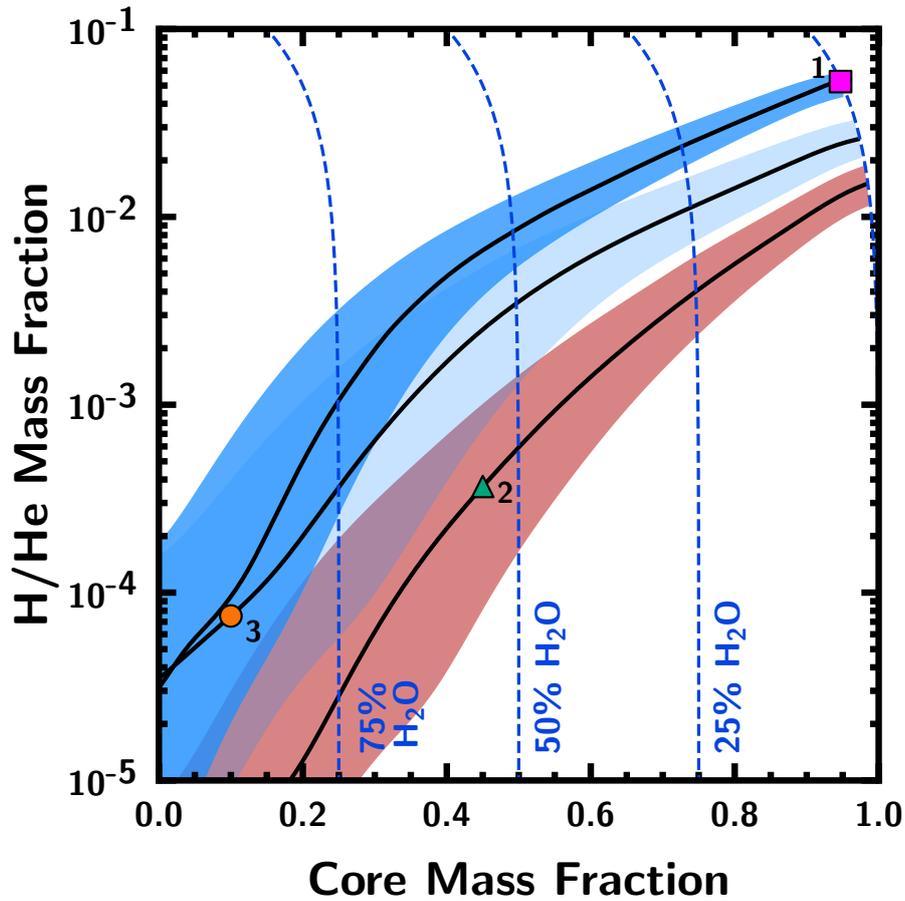


Fig. 6.3 Envelope vs core mass fraction for model solutions. The dark red and blue shaded regions show the same cases as in the ternary diagram. The pale blue region shows an additional case with $T_{\text{int}} = 25\text{K}$ and an Earth-like core for comparison. Black lines in each case show best-fit solutions. The magenta square, teal triangle and orange circle represent the rocky world, intermediate and water world scenarios discussed in section 6.2.3.

(H_2O - H/He) envelope, we find that the difference in radius between the mixed and non-mixed models is less than half of the measured uncertainty (see Figure 6.1). The constraint on the envelope mass fraction from this mixed case is $x_{\text{env}} = 2.5 - 6.4\%$, consistent with, and a subset of, the constraints discussed above. Note that in this case x_{env} includes both the H/He and H_2O mass fraction.

6.2.2 Atmosphere-ocean Boundary

Our constraints on the interior compositions of K2-18b result in a wide range of thermodynamic conditions at the H_2O - H/He boundary (HHB). The pressure (P_{HHB}) and Temperature (T_{HHB}) at the HHB for the model solutions are shown in Figure 6.5. Each point on the HHB loci

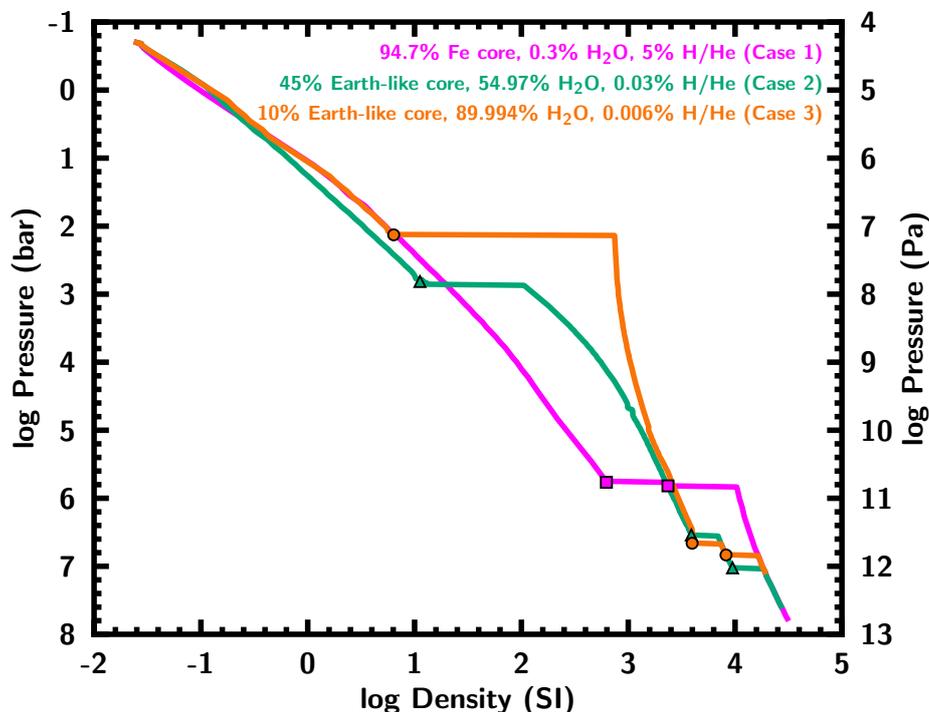


Fig. 6.4 Pressure–density profiles for three possible compositions of K2-18b discussed in Section 6.2.3. The transitions between components are marked.

denotes the transition from the P – T profile in the H/He envelope to the corresponding H_2O adiabat. The P_{HHB} and T_{HHB} depend on the H/He envelope mass fraction. For a given P – T profile, larger envelopes result in higher P_{HHB} and T_{HHB} . For example, solutions with $x_{\text{env}} \gtrsim 1\%$ lead to P_{HHB} and T_{HHB} corresponding to the super-critical phase of H_2O . As shown in Figure 6.2, solutions with higher x_{env} correspond to higher x_{core} and lower $x_{\text{H}_2\text{O}}$.

Conversely, solutions with lower x_{core} and, hence, lower x_{env} and higher $x_{\text{H}_2\text{O}}$, lead to lower P_{HHB} and T_{HHB} ; with H_2O in vapour or liquid phases at the HHB. For example, an $x_{\text{core}} \lesssim 30\%$ leads to a P_{HHB} and T_{HHB} corresponding to the liquid phase of H_2O , for the cooler P – T profile (with $T_{\text{int}} = 25\text{K}$). For $x_{\text{core}} \sim 10\%$ or less, the P_{HHB} and T_{HHB} approach STP conditions for liquid H_2O . Below the HHB, H_2O is found in increasingly dense phases spanning liquid, vapour, super-critical, and ice states depending on the location of the HHB and the extent of the H_2O layer, as shown in Figure 6.5. In the case of a mixed H_2O –H/He envelope, the HHB is undefined as it corresponds to an extreme case with no pure H_2O layer.

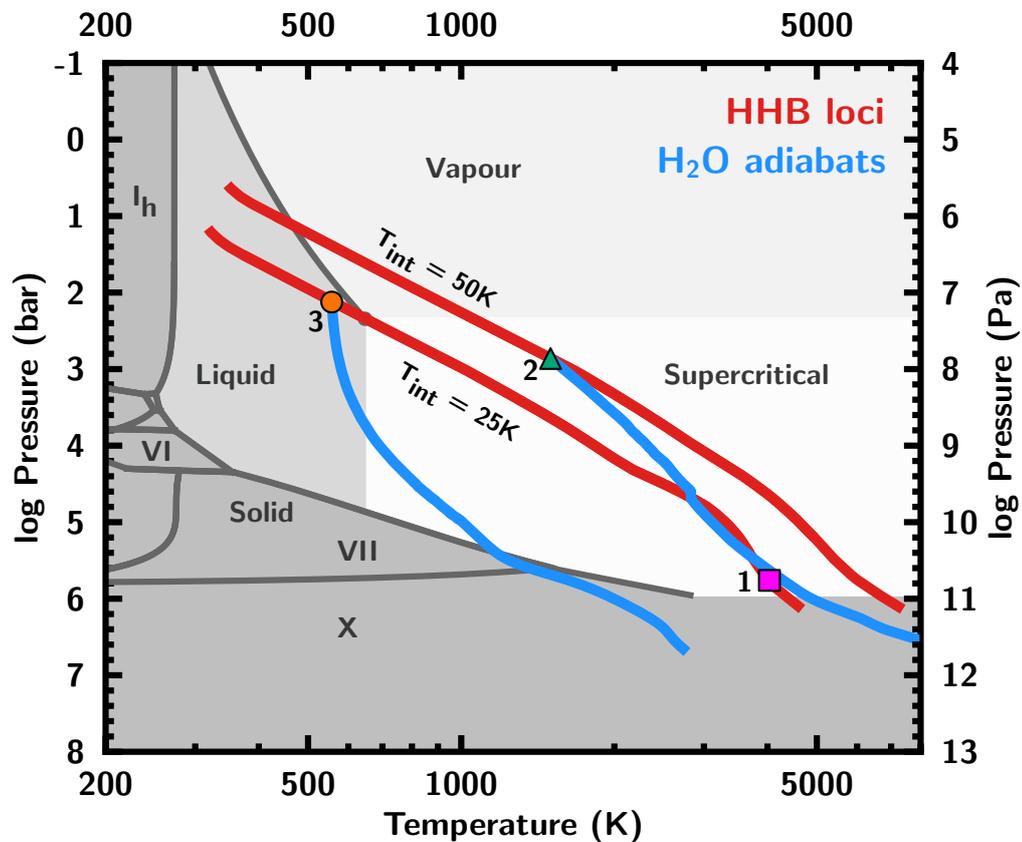


Fig. 6.5 Thermodynamic conditions at the H₂O-H/He boundary (HHB) for K2-18b. The red lines indicate possible pressures and temperatures at the HHB for two values of T_{int} . The phase diagram of H₂O is shown in the background. The square, circle and triangle correspond to the representative cases from Figure 6.4 with the same colour. We only show solutions with core mass fractions $\geq 10\%$; less massive cores lead to lower P and T at the HHB. Blue lines show adiabatic temperature profiles in the H₂O layer for the three examples.

6.2.3 Possible Composition and Origins

Our constraints on the interior and atmospheric properties of K2-18b provide insights into its physical conditions, origins, and potential habitability. Here we discuss three representative classes that span the range of possible compositions, as indicated in Figures 6.1, 6.2, 6.4 and 6.5. The specific cases chosen here fit the M_p and R_p exactly, as shown in figure 6.1. A wider range of solutions exist in each of these classes within the 1σ uncertainties.

Case 1: Rocky World. One possible scenario is a massive rocky interior overlaid by a H/He envelope. For example, a pure Fe core of 94.7% by mass with an almost maximal H/He envelope of 5% explains the data with minimal $x_{\text{H}_2\text{O}} = 0.3\%$, consistent with our retrieved H₂O abundance in the atmosphere. The HHB in this case is at $\sim 10^6$ bar, yielding

supercritical H₂O close to the ice X phase. It is also possible in this case that the H₂O and H/He are mixed, meaning the HHB is not well-defined. Such a scenario is consistent with either H₂ outgassing from the interior (Elkins-Tanton & Seager, 2008, Rogers & Seager, 2010b) or accretion of an H₂-rich envelope during formation (Lee & Chiang, 2016).

Case 2: Mini-Neptune. There are a range of plausible compositions consisting of a non-negligible H/He envelope in addition to significant H₂O and core mass fractions, akin to canonical models for Neptune and Uranus (Guillot & Gautier, 2014). One such example is a 45% Earth-like core with $x_{\text{env}} = 0.03\%$ and $x_{\text{H}_2\text{O}} = 54.97\%$. In this case the HHB is at $P_{\text{HHB}} = 700\text{bar}$ and $T_{\text{HHB}} = 1500\text{K}$, with H₂O in the supercritical phase.

Case 3: Water World. A ~100% water world with a minimal H₂-rich atmosphere ($x_{\text{env}} \sim 10^{-6}$) is permissible by the data. However, such an extreme case is implausible from a planet formation perspective; some amount of rocky core is required to initiate further ice and gas accretion (Rogers et al., 2011, Lee & Chiang, 2016). For example, a planet with $x_{\text{core}} = 10\%$, $x_{\text{H}_2\text{O}} = 89.994\%$ and a thin H/He envelope ($x_{\text{env}} = 0.006\%$) can explain the data. For this case, $P_{\text{HHB}} = 130\text{bar}$ and $T_{\text{HHB}} = 560\text{K}$, corresponding to liquid H₂O. For the same core fraction, solutions with even smaller H/He envelopes are admissible within the 1σ uncertainties on M_p and R_p , leading to P_{HHB} and T_{HHB} approaching habitable STP conditions.

6.2.4 Potential Habitability

A notional definition of habitability argues for a planetary surface with temperatures and pressures conducive to liquid H₂O (e.g. Kasting et al., 1993, Meadows & Barnes, 2018). Living organisms are known to thrive in Earth's extreme environments (extremophiles). Their living conditions span the phase space of liquid H₂O up to ~1000 bar pressures at the bottom of the Marianas Trench and ~400 K temperatures near hydrothermal vents (e.g. Merino et al., 2019).

Whether or not habitable conditions prevail on K2-18b depends on the extent of the H/He envelope. The thermodynamic conditions at the surface of the H₂O layer span a wide range in the H₂O phase diagram. While most of these solutions lie in the super-critical phase, many others lie in the liquid and vapour phases. Model solutions with core mass fractions <15% and H/He envelopes $\lesssim 10^{-3}$ allow for liquid H₂O at Earth-like habitable conditions discussed above. One plausible scenario is an ocean world, as discussed in section 6.2.3, with liquid water approaching STP conditions (300 K, ~1–10 bar) underneath a thin H/He atmosphere ($x_{\text{env}} \lesssim 10^{-5}$).

A number of studies in the past have argued for potential habitability on planets with H/He-rich atmospheres orbiting M-dwarfs (e.g. Pierrehumbert & Gaidos, 2011, Seager et al.,

2013, Koll & Cronin, 2019). Given our constraints above, we find that K2-18b has a realistic chance of being habitable. Furthermore, our constraints on CH₄ and NH₃ suggest chemical disequilibrium. Among other possibilities for chemical disequilibrium, e.g. photochemistry, the potential influence of biochemical processes may not be entirely ruled out (e.g. Seager et al., 2013). Future observations, e.g. with JWST, will have the potential to refine our findings (see Section 6.4). We argue that planets such as K2-18b can indeed have the potential to approach habitable conditions and searches for biosignatures should not necessarily be restricted to smaller rocky planets.

6.3 A Pair of Mini-Neptunes Orbiting the Bright M-dwarf TOI-776

6.3.1 Planetary Composition and Internal Structure

For a better understanding of the nature of the two exoplanets, we produce detailed models of their interior compositions, using their masses, radii, and surface temperatures. Our internal structure model is the same as the model used in Chapter 5. We assume that the core is Earth-like in composition (1/3 iron, 2/3 silicates by mass). The temperature profile in the H/He envelope is isothermal from the surface down to some radiative-convective boundary, where it becomes adiabatic. The pressure at the radiative-convective boundary P_{rc} is a free parameter in the model. For this study, we considered values of P_{rc} ranging from 1–100 bar.

We explore the parameter space of possible compositions in $(x_{core}, x_{H_2O}, x_{H/He})$ space. For each composition, we consider a range of masses that agree with the observed mass of the planet to within 1σ . For a given mass \hat{M} , the model radius \hat{R} is computed, and the χ^2 statistic is calculated:

$$\chi^2 = \frac{(M_p - \hat{M})^2}{\sigma_M^2} + \frac{(R_p - \hat{R})^2}{\sigma_R^2}, \quad (6.1)$$

where (σ_M, σ_R) are the observed uncertainties on the mass and radius of each planet.

The bulk densities of TOI-776b and c ($3.4_{-0.9}^{+1.1}$ g cm⁻³ and $3.5_{-1.3}^{+1.5}$ g cm⁻³, respectively) are too low for either planet to have a purely terrestrial (iron plus rock) composition. Therefore, the planets must possess an envelope with some amount of H₂O and/or H/He, in order to explain their masses and radii. The right panel of Fig. 6.6 shows limiting cases for each planet in which the envelope composition is either purely H₂O or purely H/He. The mass and radius of TOI-776b can be explained to within 1σ ($\chi^2 \leq 1$) with a pure H₂O envelope of 12–73% by mass or a pure H/He envelope with a mass fraction of 1.1×10^{-4} – 5.2×10^{-3} . Best-fit solutions (those which minimize χ^2) for pure envelopes are found at $x_{H_2O} = 0.3$ and $x_{H/He} = 1.1 \times 10^{-3}$.

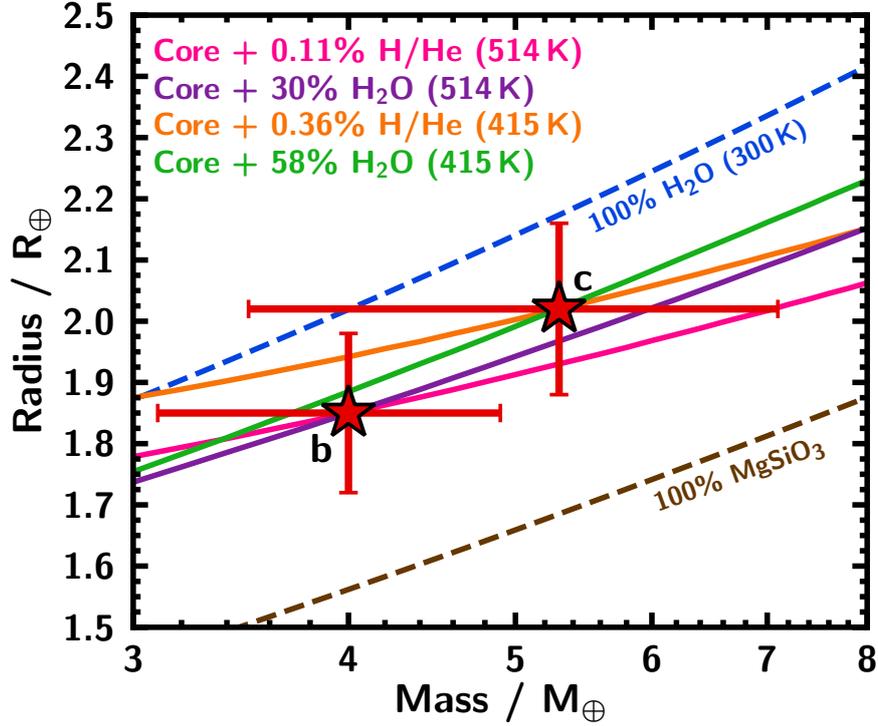


Fig. 6.6 Locations of TOI-776b and c on a mass–radius diagram. The solid pink and purple lines show end-member models that are consistent with the mass and radius of TOI-776b, and the orange and green lines show end-member compositions consistent with the mass and radius of TOI-776c, assuming an Earth-like core (1/3 iron, 2/3 silicates).

TOI-776c might have larger envelopes; within 1σ , it is consistent with a pure H₂O layer of $\geq 18\%$ or a pure H/He envelope with a mass fraction of 5.4×10^{-4} – 1.2×10^{-2} . The best-fit pure-envelope solutions for TOI-776c are $x_{\text{H}_2\text{O}} = 0.58$ and $x_{\text{H/He}} = 3.6 \times 10^{-3}$. Each of the best-fit models, shown in the right panel of Fig. 6.6, have a radiative-convective boundary at $P_{\text{rc}} = 10$ bar.

It is also possible that the planets in this system have both H₂O and H/He components, as well as an iron/rock core. For the three components, we explored the full range of plausible values (x_{core} , $x_{\text{H}_2\text{O}}$, and $x_{\text{H/He}}$) that could explain the interior compositions of each planet. We consider two different temperature profiles for each planet, with $P_{\text{rc}} = 1$ and 100 bar. Figure 6.7 shows the mass fractions of water and H/He compatible to within 1σ ($\chi^2 \leq 1$) with the masses and radii of TOI-776b and c. We obtain upper limits on the total H₂O and H/He mass fractions for TOI-776b: $x_{\text{H}_2\text{O}} \leq 73\%$ and $x_{\text{H/He}} \leq 0.52\%$. These correspond to cases with pure H₂O or H/He envelopes, as previously discussed. For TOI-776c, we find that $x_{\text{H/He}} \leq 1.2\%$. A 100% H₂O planet would theoretically be consistent with the mass and

radius of TOI-776c, but this would be unrealistic from a planet formation perspective, as some rocky material is needed for further accretion of ice and gas (Lee & Chiang, 2016). Figure 6.7 also shows a significant overlap between the best-fit shaded regions for the two planets, meaning that the planets could also share the same composition.

The masses and radii of TOI-776b and c allow for a wide range of possible solutions, from water worlds with steam atmospheres to mostly rocky planets with hydrogen-rich envelopes; however, they are inconsistent with bare rocks without atmospheres. Our models assume a surface pressure of 0.1 bar, meaning a water-world solution for either planet yields a steam atmosphere. On the other hand, a higher surface pressure could result in liquid H₂O at the surface. A rocky planet with an outgassed secondary atmosphere which includes carbon compounds is unlikely: Elkins-Tanton & Seager (2008) placed an upper limit on the mass fraction for this type of atmosphere at 5%. The lower mass limits in the case of pure H₂O envelopes are 8% and 18% for TOI-776b and c, respectively. On the other hand, in a carbon-rich atmosphere, the dominant species, CO₂, has a higher mean molecular weight than H₂O, leading to a lower atmospheric scale height. All things considered, we can infer that a 5% carbon-rich atmosphere is less than what would be needed to explain the planet radii. However, determining whether the two planets have H₂O- or H/He-rich atmospheres is impossible with the present data. Atmospheric observations of the planets will be required in order to break this degeneracy.

6.4 Future Prospects for Characterisation

Each of the planets considered in this section require more detailed atmospheric observations in order to answer key questions about their characteristics. Although a transmission spectrum of K2-18b has been observed using HST, many unsolved puzzles regarding its atmosphere remain. Given its low temperature, CH₄ and NH₃ would be expected to be the prominent carriers of carbon and nitrogen in the atmosphere of K2-18b, assuming chemical equilibrium. However, HST observations led to non-detections of CH₄ and NH₃, suggesting strong disequilibrium in the atmosphere (Madhusudhan et al., 2020). Without more precise observations, it is not possible to measure the causes extent of this disequilibrium, or assess whether it affects the abundances of other chemical species such as CO and CO₂. Meanwhile, there are no extant observations of the atmospheres of the planets in the TOI-776 system. Based on mass and radius measurements it is highly likely that these planets possess atmospheres, but it is not possible to tell if they are primary or secondary atmospheres, meaning little is known about the possible formation and evolutionary histories of these planets.

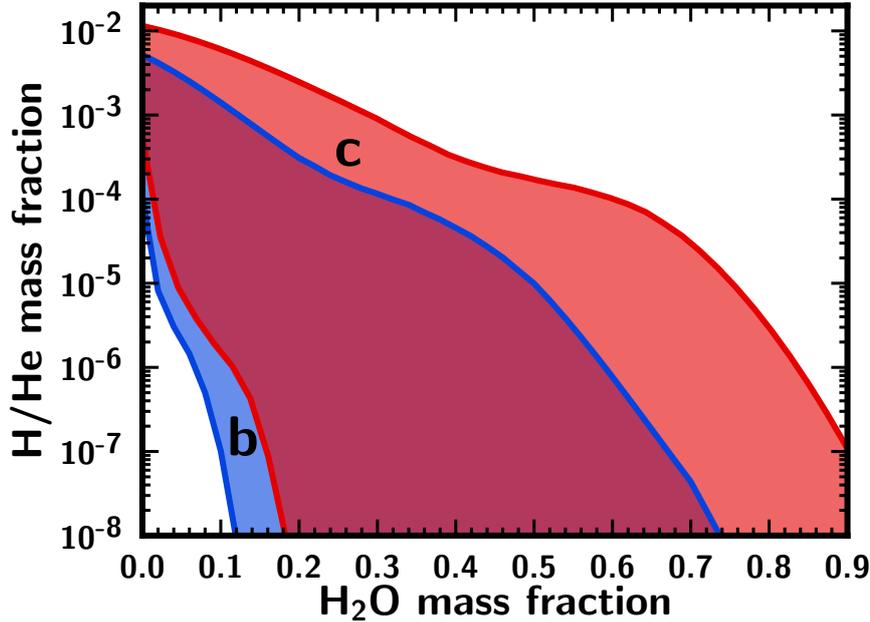


Fig. 6.7 H/He vs. H_2O mass fractions for the best-fit interior compositions ($\leq 1\sigma$) permitted by the masses and radii of TOI-776b and c, assuming an Earth-like core, for two different pressure–temperature profiles with radiative–convective boundaries at 1 and 100 bar. The blue shaded region indicates possible compositions for TOI-776b, and the red shaded region shows compositions for TOI-776c. The darker red shaded area between the two corresponds to the range of possible compositions that could explain both planets. For TOI-776b, the H_2O mass fraction is constrained to be $\leq 73\%$ and the H/He mass fraction is $\leq 0.52\%$. For TOI-776c the upper limit for H/He is 1.2%. A purely H_2O planet would be consistent with this mass and radius, but we only show H_2O mass fractions up to 90%.

K2-18b will be the target of two JWST Cycle 1 GO programs (2372, PI: Hu and 2722, PI: Madhusudhan), both of which will observe the planet’s transmission spectrum. Program 2372 will be conducted using the Near Infrared Spectrograph (NIRSpec), using the G235H and G395H gratings to cover the wavelength range 1.7–5.2 μm . The observations will consist of two transits with G235H and four transits with G395H to achieve a precision of ~ 20 ppm per spectral element at a resolution of $R = 100$. Program 2722 will use three different instruments, incorporating the Near Infrared Imager and Slitless Spectrograph (NIRISS) and the mid-Infrared Instrument (MIRI) as well as NIRSpec. NIRISS will be used in Single Object Slitless Spectroscopy (SOSS) mode, NIRSpec will use the G395H grating and the SUB2048 subarray, and MIRI’s Low Resolution Spectrometer will be used in slitless mode. Each instrument will be used to observe one transit, leading to a combined transmission spectrum covering 1–8 μm .

The transmission spectrum acquired from program 2372 will be sufficiently precise to detect and measure the abundances of numerous chemical species, including H₂O, CH₄, NH₃, CO₂ and CO. The observations should also be able to constrain the cloud-top pressure and detect byproducts of photochemistry, such as HCN and N₂O. The spectrum acquired from program 2722 will be complementary to the spectrum from program 2372, and should be able to make very accurate measurements of the abundances of CH₄ and NH₃ in the atmosphere of K2-18b. This will enable unprecedented constraints on the degree of chemical disequilibrium in the planet's atmosphere.

TOI-776b and c will be observed as part of a JWST Cycle 1 GO program targeting 11 sub-Neptunes (2512, PI: Batalha). Two transits of each planet will be observed using the G395H grating, yielding transmission spectra with a wavelength coverage of 2.9–5.3 μm . These spectra will allow for a unique comparative study of two planets in the same system. It is currently unclear whether these planets have hydrogen-rich primary atmospheres or secondary atmospheres with high mean molecular weights. The planned observations will answer this question for both planets. If both atmospheres are H-rich, then by measuring the mixing ratios of chemical species, it will be possible to determine elemental abundance ratios between the two planets, providing a direct comparison of their formation pathways. If however one of the two planets does not have a H-rich atmosphere, this would also be an extremely valuable result as it would suggest a different formation and evolutionary history for the two planets in the system. If both planets have high mean molecular weight atmospheres, these observations will indicate whether it will be feasible to measure the properties of secondary atmospheres by observing a small number of transits using JWST.

Chapter 7

Conclusions and Future Developments

“There are special times in your scientific career where Truly Landmark Things Are Happening. There will be other times where it feels like not much is happening. It’s all part of the journey. Don’t forget to appreciate the Landmark times.”

– Jonathan Fortney, 14th July 2022

The aim of this thesis has been to explore the characteristics of exoplanetary atmospheres and interiors. In order to achieve this, we have developed and expanded on a range of methods for characterising planetary atmospheric and internal properties. For exoplanet atmospheres, we have made advances in both the forward modelling and parameter estimation components of the atmospheric retrieval algorithm, incorporating three-dimensional inhomogeneities into the forward model and performing parameter estimation in a novel way using machine learning. For interiors, we have developed and applied a detailed internal structure model with the goal of characterising the possible phase structures of sub-Neptunes with a high H₂O component. In the remainder of this chapter, we summarise the main conclusions of my work before describing future developments in these areas.

7.1 Machine Learning for Atmospheric Retrieval

In Chapter 3, we presented a new approach to atmospheric retrieval using supervised machine learning, specifically the Random Forest algorithm (Nixon & Madhusudhan, 2020). This method expanded on previous efforts to incorporate Machine Learning into the retrieval framework with the goal of more closely approximating existing Bayesian methods in order to ensure that the results were reliable. Using my new framework we reproduced the results of a traditional Bayesian retrieval of the HST spectrum of the canonical hot Jupiter HD 209458b, demonstrating that machine learning could be used to accurately and robustly

reproduce traditional retrievals. My method is highly computationally efficient in low-dimensional parameter spaces, with computation times a factor of a few lower than retrievals using MultiNest. However, we found that the Random Forest algorithm is computationally prohibitive in higher-dimensional spaces that are commonly explored with Bayesian retrievals, due to the need to produce an increasingly large training data set. Further developments will be needed before it will be possible to construct a widely applicable atmospheric retrieval framework using machine learning.

A possible avenue for incorporating machine learning into retrievals would be to create a hybrid machine learning/Bayesian algorithm which uses machine learning techniques to emulate numerical models of transmission spectra while retaining a Bayesian parameter estimation scheme. It is possible that the trained model could produce spectra much more quickly than a radiative transfer code, which would lead to faster evaluations of the likelihood function and enable more computationally efficient retrievals. Ultimately, this could allow more complex forward models with fewer simplifying assumptions to be used to carry out retrievals.

7.2 Three-dimensional Atmospheric Retrieval

In Chapter 4 we introduced AURA-3D, a retrieval framework incorporating three-dimensional transmission geometry as well as a multidimensional temperature profile that is computationally efficient for retrievals of JWST data (Nixon & Madhusudhan, 2022). Using AURA-3D we showed that for simulated JWST data of hot Jupiters, it is possible to constrain thermal inhomogeneities across the terminator region to a precision of a few hundred K. While a 3D treatment does not appear to be necessary for retrievals of HST spectra, we did find that in certain cases, a 1D retrieval of a JWST hot Jupiter transmission spectrum could lead to biased abundance estimates, in agreement with previous studies. However, by incorporating my new prescription for the temperature structure, we were able to overcome these biases and obtain accurate retrieval results for a simulated planet with a fully 3D temperature structure. This will be an important tool as we come to analyse the first JWST spectra of hot giant planets.

While AURA-3D represents a significant advance in retrieval capabilities, there are many further developments that may be necessary in order to capture the complex nature of exoplanet atmospheres. We have already described how inhomogeneous chemistry in a planet's terminator has a strong impact on transmission spectra, meaning this effect may need to be incorporated into retrievals. However, it is important that concurrent advances in self-consistent modelling provide information about what kinds of atmospheric phenomena we can realistically expect to find on different planets, in order to ensure that models used in

retrievals remain physically motivated. Another natural development of AURA-3D would be to add capabilities for analysing other kinds of observations, such as emission spectra and phase curves. It could then be possible to place joint constraints on atmospheric properties across the whole planet, rather than solely focusing on the terminator region. Such an approach would be complementary to analysis of these kinds of observations using GCMs.

The goal of any atmospheric retrieval algorithm should be to allow the data to speak for itself. In order to achieve this, we should continue to explore the wide array of physical and chemical processes at work in atmospheres, so that we are ready to learn as much as possible from the glimpses of planetary atmospheres that we are able to observe.

7.3 Phase Structures of Water-rich Super-Earths

In Chapter 5 we presented a new internal structure modelling code for sub-Neptunes, incorporating a newly-compiled equation of state for H_2O that covers the full range of pressures, temperatures and phases that we might expect to find in the interiors of planets with radii between $\sim 1\text{--}4R_\oplus$ (Nixon & Madhusudhan, 2021). Using this model, we explored the range of possible ocean depths on water-rich worlds, identifying the key parameters that determine the extent of a planet's ocean to be surface temperature and gravity. We found that, for certain combinations of these parameters, these oceans could be hundreds of times deeper than Earth's. We also considered a hypothetical class of sub-Neptune consisting of a H/He envelope over a liquid water ocean, exploring the possible masses, radii and temperatures at which such a planet could exist. We found that certain planets could have H/He envelopes that contribute substantially to the overall planet radius while retaining an underlying liquid ocean. This finding expands the parameter space of planets that could host habitable conditions beyond the terrestrial bodies which are normally considered.

Although my study did not focus primarily on habitability, further considerations of the habitability of sub-Neptunes could be a possible future application of my internal structure model. Similar models have already been used to explore the concept of a habitable zone for sub-Neptunes (Madhusudhan et al., 2021). Furthermore, recent laboratory experiments have suggested that microbial life could thrive in H_2 -dominated environments (Seager et al., 2020). However, there are many more factors that could impede the development of life on such planets; for example, many of the geological processes that take place on Earth and are crucial for the development of life could not take place on a planet with an extremely large internal H_2O component (Kitzmann et al., 2015). Therefore, continuing to develop a detailed understanding of the internal phase structures of water worlds will be paramount as we seek

to determine the extent to which the search for extraterrestrial life could be expanded to more exotic worlds.

7.4 Internal Structures of JWST Targets

In Chapter 6 we applied internal structure models to a number of exoplanets whose atmospheric spectra are set to be observed with JWST (Madhusudhan et al., 2020, Luque et al., 2021). We used the planetary bulk parameters and retrieved atmospheric properties of the mini-Neptune K2-18b to constrain its interior structure and thermodynamic conditions. These constraints span multiple scenarios, from rocky worlds with massive H/He to water worlds with thin envelopes. The thermodynamic conditions at the surface of the H₂O layer range from the supercritical to liquid phases, with a range of solutions allowing for habitable conditions on K2-18b. We also constrained the possible internal structures of two sub-Neptunes orbiting the M dwarf TOI-776. Like K2-18b, bulk densities of TOI-776b and c allow for a wide range of possible interior compositions, from water worlds to rocky planets with H/He-rich atmospheres, but they are too low for either planet to have a purely terrestrial (iron plus rock) composition. Thus, an atmosphere is expected for both planets.

Our understanding of the atmospheres of sub-Neptunes is set to be revolutionised with JWST, an opportunity which has been recognised with a large number of approved JWST programs dedicated to observations of their spectra, including each of the three planets described above. In the future, it will be important to consider the extent to which interior-atmosphere interactions can affect the atmospheric structure, since the atmospheres of sub-Neptune planets are thinner than their hot Jupiter counterparts. Theoretical work is being undertaken to understand these effects (Kite et al., 2020, Schlichting & Young, 2022), and combining my atmospheric and interior models could be a very useful endeavour. Possible areas of study include assessing the detectability of atmospheric species whose atmospheric abundances will be affected by the interior, as well as the effect of a planet's surface on its temperature structure.

7.5 Concluding Remarks

On 12th July 2022, the first transmission spectrum observed using JWST was presented to the public. The spectrum of the hot Jupiter WASP-96b was taken using NIRISS SOSS, covering a wavelength range of 0.6–2.9 μm . A reproduction of this spectrum is shown in Figure 7.1, along with a fiducial model fit. The spectrum was presented primarily as a demonstration of

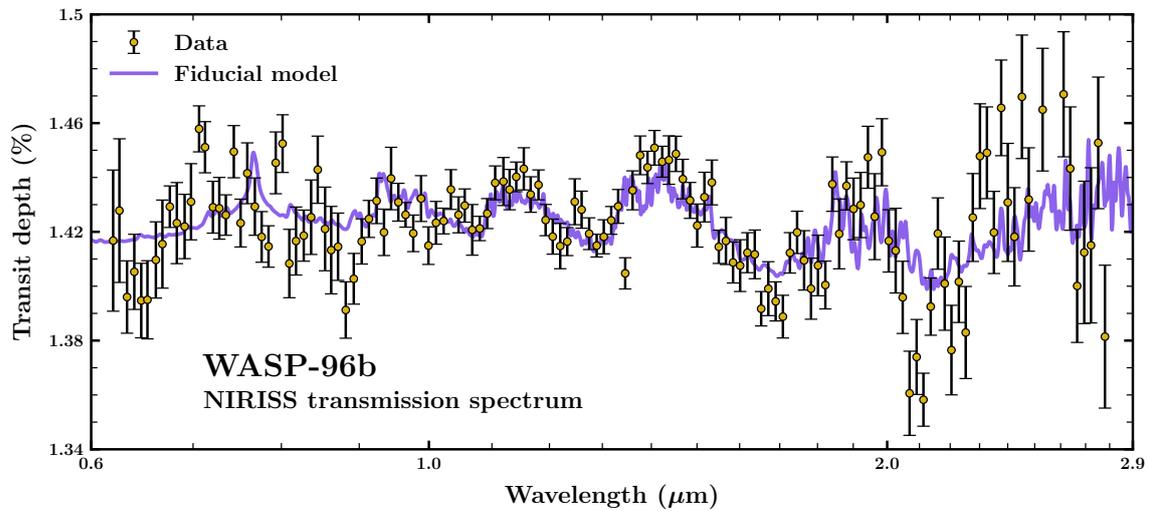


Fig. 7.1 JWST NIRISS transmission spectrum of the hot Jupiter WASP-96b, adapted from its initial presentation by NASA (nasa.gov/webbfirstimages). The gold points with error bars represent the observed data, and the purple line represents a fiducial model that achieves a moderately good fit to the observations. H_2O absorption is evident from the features at 1–2 μm .

the capabilities of JWST, with a number of caveats relating to the data reduction process, and so we do not conduct a full retrieval of the spectrum here.

However, the spectrum that was displayed indicates the vast improvements in data quality that we are about to realise as the JWST program commences. The wavelength coverage and resolution of this spectrum far exceeds any spectrum taken using a single HST instrument, and even this early reduction of the spectrum clearly exhibits a number of H_2O absorption features from 1–2 μm , underlining the ability of JWST to make highly confident detections of molecular species. Furthermore, this spectrum is far from the highest-quality data set that JWST will be able to provide. Transit observations are very time-sensitive, and this target was selected to fit a very narrow time window soon after JWST became operational in order to be ready for the initial press release, meaning many planets which are expected to be better suited to transmission spectroscopy could not be chosen. Also, this observation was only of a single transit event using one instrument; by observing multiple transits with a range of instruments, it will be possible to achieve more precise spectra with a much broader wavelength coverage.

With this first spectrum, we have entered a new era in exoplanet science, in which the quality and volume of observational data will improve immensely. The coming months and years of JWST observations will lead to a flurry of new discoveries, never-before-seen features, and a host of further questions and problems to consider. It is even possible that this

era will see such landmark events as the first indications of signs of life on worlds beyond our solar system. Theoretical models of planetary atmospheres and interiors will play a crucial role in the interpretation of this new stream of data. It is my hope that the work presented in this thesis, alongside future developments in atmospheric retrieval and internal structure models, will advance our ability to characterise exoplanet atmospheres and interiors using current instruments such as HST, contributing to its lasting legacy, while setting the stage for a revolution in our understanding of other worlds using JWST.

References

- Agol E., Steffen J., Sari R., Clarkson W., 2005, MNRAS, 359, 567
- Ahrens T. J., 2000, Mineral Physics & Crystallography: A Handbook of Physical Constants. American Geophysical Union
- Alessi M., Pudritz R. E., Cridland A. J., 2017, MNRAS, 464, 428
- Alibert Y., 2014, A&A, 561, A41
- Alibert Y., et al., 2018, Nature Astronomy, 2, 873
- Allard N. F., Spiegelman F., Kielkopf J. F., 2016, A&A, 589, A21
- Allard N. F., Spiegelman F., Leininger T., Molliere P., 2019, A&A, 628, A120
- Allen M., Yung Y. L., Waters J. W., 1981, J. Geophys. Res., 86, 3617
- Amundsen D. S., et al., 2016, A&A, 595, A36
- Anderson O. L., Dubrovinsky L., Saxena S. K., LeBihan T., 2001, Geophysical Research Letters, 28, 399
- Anglada-Escudé G., et al., 2016, Nature, 536, 437
- Arcangeli J., et al., 2018, ApJ, 855, L30
- Ashton G., et al., 2022, Nature Reviews Methods Primers, 2
- Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, ARA&A, 47, 481
- Atreya S. K., Crida A., Guillot T., Lunine J. I., Madhusudhan N., Mousis O., 2018, in Baines K. H., Flasar F. M., Krupp N., Stallard T., eds, Saturn in the 21st Century. Cambridge University Press, pp 5–43, doi:10.1017/9781316227220.002
- Bakos G. Á., et al., 2010, ApJ, 710, 1724
- Banerji M., et al., 2010, MNRAS, 406, 342
- Baraffe I., Chabrier G., Fortney J., Sotin C., 2014, in Beuther H., Klessen R. S., Dullemond C. P., Henning T., eds, Protostars and Planets VI. p. 763, doi:10.2458/azu_uapress_9780816531240-ch033

- Barber R. J., Strange J. K., Hill C., Polyansky O. L., Mellau G. C., Yurchenko S. N., Tennyson J., 2014, *MNRAS*, 437, 1828
- Barman T. S., Macintosh B., Konopacky Q. M., Marois C., 2011, *ApJ*, 733, 65
- Barstow J. K., Aigrain S., Irwin P. G. J., Fletcher L. N., Lee J. M., 2013, *MNRAS*, 434, 2616
- Barstow J. K., Aigrain S., Irwin P. G. J., Sing D. K., 2017, *ApJ*, 834, 50
- Batalha N. M., et al., 2013, *ApJS*, 204, 24
- Batalha N. E., et al., 2017, *PASP*, 129, 064501
- Batygin K., Bodenheimer P. H., Laughlin G. P., 2016, *ApJ*, 829, 114
- Bayes T., 1763, *Phil. Trans. R. Soc.*, 53, 370
- Bean J. L., Miller-Ricci Kempton E., Homeier D., 2010, *Nature*, 468, 669
- Beaugé C., Nesvorný D., 2012, *ApJ*, 751, 119
- Bell T. J., et al., 2021, *MNRAS*, 504, 3316
- Benneke B., Seager S., 2012, *ApJ*, 753, 100
- Benneke B., Seager S., 2013, *ApJ*, 778, 153
- Benneke B., et al., 2017, *ApJ*, 834, 187
- Benneke B., et al., 2019, *ApJ*, 887, L14
- Bennett D. P., Anderson J., Gaudi B. S., 2007, *ApJ*, 660, 781
- Biller B. A., Bonnefoy M., 2018, in Deeg H. J., Belmonte J. A., eds, *Handbook of Exoplanets*. Springer, p. 101, doi:10.1007/978-3-319-55333-7_101
- Birch F., 1952, *J. Geophys. Res.*, 57, 227
- Blecic J., Dobbs-Dixon I., Greene T., 2017, *ApJ*, 848, 127
- Blecic J., et al., 2021, arXiv e-prints, p. arXiv:2104.12525
- Blei D. M., Kucukelbir A., McAuliffe J. D., 2017, *Journal of the American Statistical Association*, 112, 859
- Bodenheimer P., Lin D. N. C., Mardling R. A., 2001, *ApJ*, 548, 466
- Bond I. A., et al., 2004, *ApJ*, 606, L155
- Bonnefoy M., et al., 2014, *A&A*, 567, L9
- Borucki W. J., et al., 2010, *Science*, 327, 977
- Boss A. P., 1997, *Science*, 276, 1836

- Bowler B. P., Nielsen E. L., 2018, in Deeg H. J., Belmonte J. A., eds, *Handbook of Exoplanets*. Springer, p. 155, doi:10.1007/978-3-319-55333-7_155
- Breiman L., 2001, *Machine Learning*, 45, 5
- Breiman L., Friedman J., Stone C., Olshen R., 1984, *Classification and Regression Trees*. Taylor & Francis
- Buchhave L. A., et al., 2016, *AJ*, 152, 160
- Buchner J., et al., 2014, *A&A*, 564, A125
- Budaj J., Kocifaj M., Salmeron R., Hubeny I., 2015, *MNRAS*, 454, 2
- del Burgo C., Allende Prieto C., 2016, *MNRAS*, 463, 1400
- Burningham B., Marley M. S., Line M. R., Lupu R., Visscher C., Morley C. V., Saumon D., Freedman R., 2017, *MNRAS*, 470, 1177
- Burrows A., Sharp C. M., 1999, *ApJ*, 512, 843
- Cabot S. H. C., Madhusudhan N., Welbanks L., Piette A., Gandhi S., 2020, *MNRAS*, 494, 363
- Caillabet L., Mazevet S., Loubeyre P., 2011, *Phys. Rev. B*, 83, 094101
- Caldas A., Leconte J., Selsis F., Waldmann I. P., Bordé P., Rocchetto M., Charnay B., 2019, *A&A*, 623, A161
- Cameron A. G. W., 1978, *Moon and Planets*, 18, 5
- Carleo G., Cirac I., Cranmer K., Daudet L., Schuld M., Tishby N., Vogt-Maranto L., Zdeborová L., 2019, *Reviews of Modern Physics*, 91, 045002
- Cassan A., et al., 2012, *Nature*, 481, 167
- Chabrier G., Potekhin A. Y., 1998, *Phys. Rev. E*, 58, 4941
- Chabrier G., Johansen A., Janson M., Rafikov R., 2014, in Beuther H., Klessen R. S., Dullemond C. P., Henning T., eds, *Protostars and Planets VI*. p. 619, doi:10.2458/azu_uapress_9780816531240-ch027
- Chabrier G., Mazevet S., Soubiran F., 2019, *ApJ*, 872, 51
- Changeat Q., Al-Refaie A., 2020, *ApJ*, 898, 155
- Chapman S., 1930, *Memoirs of the Royal Meteorological Society*, 3, 103
- Charbonneau D., Brown T. M., Latham D. W., Mayor M., 2000, *ApJ*, 529, L45
- Charbonneau D., Brown T. M., Noyes R. W., Gilliland R. L., 2002, *ApJ*, 568, 377
- Charette M. A., Smith W. F., 2010, *Oceanography*, 23

- Chauvin G., Lagrange A. M., Dumas C., Zuckerman B., Mouillet D., Song I., Beuzit J. L., Lowrance P., 2004, *A&A*, 425, L29
- Chen G., Pallé E., Nortmann L., Murgas F., Parviainen H., Nowak G., 2017, *A&A*, 600, L11
- Chen G., et al., 2018, *A&A*, 616, A145
- Chilcote J., et al., 2017, *AJ*, 153, 182
- Chubb K. L., Min M., 2022, arXiv e-prints, p. arXiv:2206.09738
- Cloutier R., et al., 2017, *A&A*, 608, A35
- Cloutier R., et al., 2019, *A&A*, 621, A49
- Cobb A. D., et al., 2019, *AJ*, 158, 33
- Constantinou S., Madhusudhan N., 2022, *MNRAS*, 514, 2073
- Cooper C. S., Showman A. P., 2005, *ApJ*, 629, L45
- Cowan N. B., Agol E., 2008, *ApJ*, 678, L129
- Cowan N. B., Agol E., 2011, *ApJ*, 729, 54
- Cubillos P., et al., 2013, *ApJ*, 768, 42
- Cubillos P. E., et al., 2022, *PSJ*, 3, 81
- D'Angelo G., Durisen R. H., Lissauer J. J., 2010, in Seager S., ed., *Exoplanets*. University of Arizona Press, pp 319–346
- Dalgarno A., Williams D. A., 1962, *ApJ*, 136, 690
- Deibert E. K., et al., 2021, *ApJ*, 919, L15
- Deitrick R., Mendonça J. M., Schroffenegger U., Grimm S. L., Tsai S.-M., Heng K., 2020, *ApJS*, 248, 30
- Deleuil M., et al., 2018, *A&A*, 619, A97
- Deming D., et al., 2013, *ApJ*, 774, 95
- Demory B.-O., et al., 2013, *ApJ*, 776, L25
- Des Marais D. J., et al., 2002, *Astrobiology*, 2, 153
- Des Marais D. J., et al., 2008, *Astrobiology*, 8, 715
- Désert J. M., et al., 2011, *A&A*, 526, A12
- Dittmann J. A., et al., 2017, *Nature*, 544, 333
- Dobbs-Dixon I., Agol E., 2013, *MNRAS*, 435, 3159

- Dorn C., Khan A., Heng K., Connolly J. A. D., Alibert Y., Benz W., Tackley P., 2015, *A&A*, 577, A83
- Dorn C., Venturini J., Khan A., Heng K., Alibert Y., Helled R., Rivoldini A., Benz W., 2017, *A&A*, 597, A37
- Dressing C. D., Charbonneau D., 2015, *ApJ*, 807, 45
- Dunaeva A. N., Antsyshkin D. V., Kuskov O. L., 2010, *Solar System Research*, 44, 202
- Ehrenreich D., Lecavelier des Etangs A., Beaulieu J.-P., Grasset O., 2006, *ApJ*, 651, 535
- Einstein A., 1936, *Science*, 84, 506
- Elkins-Tanton L. T., Seager S., 2008, *ApJ*, 685, 1237
- Espinoza N., Jones K., 2021, *AJ*, 162, 165
- Evans J., 1998, *The History and Practice of Ancient Astronomy*. Oxford University Press, <https://ezp.lib.cam.ac.uk/login?url=https://search.ebscohost.com/login.aspx?direct=true&db=nlebk&AN=368809&site=ehost-live&scope=site>
- Evans T. M., et al., 2017, *Nature*, 548, 58
- Everall A., Evans N. W., Belokurov V., Schönrich R., 2019, *MNRAS*, 489, 910
- Falco A., Zingales T., Pluriel W., Leconte J., 2021, arXiv e-prints, p. arXiv:2110.11799
- Fei Y., Mao H., Hemley R. J., 1993, *The Journal of Chemical Physics*, 99, 5369
- Feistel R., Wagner W., 2006, *Journal of Physical and Chemical Reference Data*, 35, 1021
- Feng Y. K., Line M. R., Fortney J. J., 2020, *AJ*, 160, 137
- Feroz F., Hobson M. P., 2008, *MNRAS*, 384, 449
- Feroz F., Skilling J., 2013, in von Toussaint U., ed., *American Institute of Physics Conference Series Vol. 1553, Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 32nd International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*. pp 106–113, doi:10.1063/1.4819989
- Feroz F., Hobson M. P., Bridges M., 2009, *MNRAS*, 398, 1601
- Fischer D. A., et al., 2016, *PASP*, 128, 066001
- Fisher C., Hoeijmakers H. J., Kitzmann D., Márquez-Neila P., Grimm S. L., Sznitman R., Heng K., 2020, *AJ*, 159, 192
- Ford E. B., et al., 2012, *ApJ*, 750, 113
- Foreman-Mackey D., Montet B. T., Hogg D. W., Morton T. D., Wang D., Schölkopf B., 2015, *ApJ*, 806, 215
- Fortney J. J., 2005, *MNRAS*, 364, 649

- Fortney J. J., Marley M. S., Barnes J. W., 2007, *ApJ*, 659, 1661
- Fortney J. J., Lodders K., Marley M. S., Freedman R. S., 2008, *ApJ*, 678, 1419
- Fortney J. J., Shabram M., Showman A. P., Lian Y., Freedman R. S., Marley M. S., Lewis N. K., 2010, *ApJ*, 709, 1396
- Fortney J. J., Mordasini C., Nettelmann N., Kempton E. M. R., Greene T. P., Zahnle K., 2013, *ApJ*, 775, 80
- Fortney J. J., Barstow J. K., Madhusudhan N., 2021, arXiv e-prints, p. arXiv:2109.01976
- Freedman R. S., Marley M. S., Lodders K., 2008, *ApJS*, 174, 504
- French M., Mattsson T. R., Nettelmann N., Redmer R., 2009, *Phys. Rev. B*, 79, 054107
- Fulton B. J., Petigura E. A., 2018, *AJ*, 156, 264
- Fulton B. J., et al., 2017, *AJ*, 154, 109
- Gal Y., Ghahramani Z., 2016, in Balcan M. F., Weinberger K. Q., eds, *Proceedings of Machine Learning Research Vol. 48, Proceedings of The 33rd International Conference on Machine Learning*. PMLR, New York, New York, USA, pp 1050–1059, <http://proceedings.mlr.press/v48/gal16.html>
- Galle J. G., 1846, *MNRAS*, 7, 153
- Gandhi S., Madhusudhan N., 2017, *MNRAS*, 472, 2334
- Gandhi S., Madhusudhan N., 2018, *MNRAS*, 474, 271
- Gardner J. V., Armstrong A. A., Calder B. R., Beaudoin J., 2014, *Marine Geodesy*, 37, 1
- Gaudi B. S., 2012, *ARA&A*, 50, 411
- Gautier T. N. I., et al., 2012, *ApJ*, 749, 15
- Gibson N. P., Nugroho S. K., Lothringer J., Maguire C., Sing D. K., 2022, *MNRAS*, 512, 4618
- Gillon M., et al., 2017, *Nature*, 542, 456
- Goldreich P., Tremaine S., 1980, *ApJ*, 241, 425
- Goodfellow I. J., Pouget-Abadie J., Mirza M., et al., 2014, in *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2. NIPS'14*. MIT Press, Cambridge, MA, USA, pp 2672–2680, <http://dl.acm.org/citation.cfm?id=2969033.2969125>
- Grasset O., Schneider J., Sotin C., 2009, *ApJ*, 693, 722
- Greene T. P., Line M. R., Montero C., Fortney J. J., Lustig-Yaeger J., Luther K., 2016, *ApJ*, 817, 17

- Griffith C. A., 2014, *Philosophical Transactions of the Royal Society of London Series A*, 372, 20130086
- Grillmair C. J., et al., 2008, *Nature*, 456, 767
- Guillot T., 2010, *A&A*, 520, A27
- Guillot T., Gautier D., 2014, arXiv e-prints, p. arXiv:1405.3752
- Guillot T., Showman A. P., 2002, *A&A*, 385, 156
- Guillot T., Burrows A., Hubbard W. B., Lunine J. I., Saumon D., 1996, *ApJ*, 459, L35
- Guillot T., Fletcher L. N., Helled R., Ikoma M., Line M. R., Parmentier V., 2022, arXiv e-prints, p. arXiv:2205.04100
- Guo X., et al., 2020, *AJ*, 159, 239
- Gupta A., Schlichting H. E., 2019, *MNRAS*, 487, 24
- Guzmán-Mesa A., et al., 2020, *AJ*, 160, 15
- Habets G. M. H. J., Heintze J. R. W., 1981, *A&AS*, 46, 193
- Handley W. J., Hobson M. P., Lasenby A. N., 2015, *MNRAS*, 453, 4384
- Hansen B. M. S., 2008, *ApJS*, 179, 484
- Harper M., et al., 2015, Zenodo 10.5281/zenodo.594435
- Harrington J., et al., 2022, *PSJ*, 3, 80
- Harris C. R., et al., 2020, *Nature*, 585, 357
- Hart M. H., 1979, *Icarus*, 37, 351
- Hayes J. J. C., et al., 2020, *MNRAS*, 494, 4492
- Haynes K., Mandell A. M., Madhusudhan N., Deming D., Knutson H., 2015, *ApJ*, 806, 146
- Hedges C., Madhusudhan N., 2016, *MNRAS*, 458, 1427
- Helled R., Morbidelli A., 2021, in Madhusudhan N., ed., *ExoFrontiers; Big Questions in Exoplanetary Science*. Institute of Physics Publishing, pp 12–1, doi:10.1088/2514-3433/abfa8fch12
- Helled R., et al., 2014, in Beuther H., Klessen R. S., Dullemond C. P., Henning T., eds, *Protostars and Planets VI*. p. 643, doi:10.2458/azu_uapress_9780816531240-ch028
- Heng K., Kitzmann D., 2017, *MNRAS*, 470, 2972
- Heng K., Lyons J. R., Tsai S.-M., 2016, *ApJ*, 816, 96
- Herschel W., Watson D., 1781, *Philosophical Transactions of the Royal Society of London Series I*, 71, 492

- Hinton G. E., Osindero S., Teh Y.-W., 2006, *Neural Computation*, 18, 1527
- Hogg D. W., Foreman-Mackey D., 2018, *ApJS*, 236, 11
- Howe A. R., Burrows A., Verne W., 2014, *ApJ*, 787, 173
- Hsu H.-W., et al., 2015, *Nature*, 519, 207
- Huang C., et al., 2021, *MNRAS*, 503, 2825
- Hubbard W. B., Marley M. S., 1989, *Icarus*, 78, 102
- Hubeny I., Burrows A., Sudarsky D., 2003, *ApJ*, 594, 1011
- Hunter J. D., 2007, *Computing in Science & Engineering*, 9, 90
- Husser T. O., Wende-von Berg S., Dreizler S., Homeier D., Reiners A., Barman T., Hauschildt P. H., 2013, *A&A*, 553, A6
- Irwin P. G. J., et al., 2008, *J. Quant. Spectrosc. Radiative Transfer*, 109, 1136
- Irwin P. G. J., Parmentier V., Taylor J., Barstow J., Aigrain S., Lee E. K. H., Garland R., 2020, *MNRAS*, 493, 106
- Jacob W. S., 1855, *MNRAS*, 15, 228
- Journaux B., et al., 2020a, *Journal of Geophysical Research (Planets)*, 125, e06176
- Journaux B., et al., 2020b, *Space Sci. Rev.*, 216, 7
- Karki B. B., Wentzcovitch R. M., de Gironcoli S., Baroni S., 2000, *Phys. Rev. B*, 62, 14750
- Kasim M. F., Bott A. F. A., Tzeferacos P., Lamb D. Q., Gregori G., Vinko S. M., 2019, *Phys. Rev. E*, 100, 033208
- Kasting J. F., Whitmire D. P., Reynolds R. T., 1993, *Icarus*, 101, 108
- Kataria T., Showman A. P., Lewis N. K., Fortney J. J., Marley M. S., Freedman R. S., 2013, *ApJ*, 767, 76
- Kite E. S., Fegley Bruce J., Schaefer L., Ford E. B., 2020, *ApJ*, 891, 111
- Kitzmann D., Patzer A. B. C., von Paris P., Godolt M., Stracke B., Gebauer S., Grenfell J. L., Rauer H., 2010, *A&A*, 511, A66
- Kitzmann D., et al., 2015, *MNRAS*, 452, 3752
- Kitzmann D., Heng K., Oreshenko M., Grimm S. L., Apai D., Bowler B. P., Burgasser A. J., Marley M. S., 2020, *ApJ*, 890, 174
- Klotz S., Komatsu K., Kagi H., Kunc K., Sano-Furukawa A., Machida S., Hattori T., 2017, *Phys. Rev. B*, 95, 174111

- Knudson M. D., Desjarlais M. P., Lemke R. W., Mattsson T. R., French M., Nettelmann N., Redmer R., 2012, *Phys. Rev. Lett.*, 108, 091102
- Knutson H. A., et al., 2007, *Nature*, 447, 183
- Knutson H. A., et al., 2009, *ApJ*, 690, 822
- Koll D. D. B., Cronin T. W., 2019, *ApJ*, 881, 120
- Komacek T. D., Showman A. P., 2016, *ApJ*, 821, 16
- Komacek T. D., Showman A. P., Tan X., 2017, *ApJ*, 835, 198
- Konopacky Q. M., Barman T. S., Macintosh B. A., Marois C., 2013, *Science*, 339, 1398
- Kratter K., Lodato G., 2016, *ARA&A*, 54, 271
- Kreidberg L., et al., 2014a, *Nature*, 505, 69
- Kreidberg L., et al., 2014b, *ApJ*, 793, L27
- Kreidberg L., et al., 2015, *ApJ*, 814, 66
- Kreidberg L., et al., 2018, *AJ*, 156, 17
- Lacy B. I., Burrows A., 2020, *ApJ*, 905, 131
- Lagrange A. M., et al., 2010, *Science*, 329, 57
- Lambrechts M., Johansen A., 2012, *A&A*, 544, A32
- Lammer H., et al., 2009, *A&ARv*, 17, 181
- Laughlin G., 2018, in Deeg H. J., Belmonte J. A., eds, *Handbook of Exoplanets*. Springer, p. 1, doi:10.1007/978-3-319-55333-7_1
- Lavie B., et al., 2017, *AJ*, 154, 91
- Lecar M., Podolak M., Sasselov D., Chiang E., 2006, *ApJ*, 640, 1115
- Lecavelier Des Etangs A., Pont F., Vidal-Madjar A., Sing D., 2008, *A&A*, 481, L83
- Lederberg J., 1965, *Nature*, 207, 9
- Lee E. J., Chiang E., 2016, *ApJ*, 817, 90
- Lee J. M., Fletcher L. N., Irwin P. G. J., 2012, *MNRAS*, 420, 170
- Lee J.-M., Heng K., Irwin P. G. J., 2013, *ApJ*, 778, 97
- Léger A., Mariotti J. M., Mennesson B., Ollivier M., Puget J. L., Rouan D., Schneider J., 1996, *Icarus*, 123, 249
- Léger A., et al., 2004, *Icarus*, 169, 499

- Levi A., Sasselov D., Podolak M., 2014, *ApJ*, 792, 125
- Levi A., Sasselov D., Podolak M., 2017, *ApJ*, 838, 24
- Lewis N. K., Showman A. P., Fortney J. J., Marley M. S., Freedman R. S., Lodders K., 2010, *ApJ*, 720, 344
- Lin D. N. C., Bodenheimer P., Richardson D. C., 1996, *Nature*, 380, 606
- Line M. R., Parmentier V., 2016, *ApJ*, 820, 78
- Line M. R., et al., 2013, *ApJ*, 775, 137
- Line M. R., Fortney J. J., Marley M. S., Sorahana S., 2014, *ApJ*, 793, 33
- Line M. R., Teske J., Burningham B., Fortney J. J., Marley M. S., 2015, *ApJ*, 807, 183
- Line M. R., et al., 2016, *AJ*, 152, 203
- Lochner M., McEwen J. D., Peiris H. V., Lahav O., Winter M. K., 2016, *ApJS*, 225, 31
- Lodders K., Fegley B., 2002, *Icarus*, 155, 393
- Lopez E. D., Fortney J. J., 2014, *ApJ*, 792, 1
- Lopez E. D., Fortney J. J., Miller N., 2012, *ApJ*, 761, 59
- Lothringer J. D., Barman T., Koskinen T., 2018, *ApJ*, 866, 27
- Loupe G., 2014, PhD thesis, University of Liège
- Luger R., Agol E., Foreman-Mackey D., Fleming D. P., Lustig-Yaeger J., Deitrick R., 2019, *AJ*, 157, 64
- Lupu R. E., Marley M. S., Lewis N., Line M., Traub W. A., Zahnle K., 2016, *AJ*, 152, 217
- Luque R., et al., 2021, *A&A*, 645, A41
- Lustig-Yaeger J., Meadows V. S., Lincowski A. P., 2019, *AJ*, 158, 27
- MacDonald R. J., Lewis N. K., 2022, *ApJ*, 929, 20
- MacDonald R. J., Madhusudhan N., 2017, *MNRAS*, 469, 1979
- MacDonald R. J., Madhusudhan N., 2019, *MNRAS*, 486, 1292
- MacDonald R. J., Goyal J. M., Lewis N. K., 2020, *ApJ*, 893, L43
- Madhusudhan N., 2012, *ApJ*, 758, 36
- Madhusudhan N., 2018, *Atmospheric Retrieval of Exoplanets*. Springer, p. 104, doi:10.1007/978-3-319-55333-7_104
- Madhusudhan N., 2019, *ARA&A*, 57, 617

- Madhusudhan N., Redfield S., 2015, *International Journal of Astrobiology*, 14, 177
- Madhusudhan N., Seager S., 2009, *ApJ*, 707, 24
- Madhusudhan N., Seager S., 2010, *ApJ*, 725, 261
- Madhusudhan N., et al., 2011, *Nature*, 469, 64
- Madhusudhan N., Lee K. K. M., Mousis O., 2012, *ApJ*, 759, L40
- Madhusudhan N., Knutson H., Fortney J. J., Barman T., 2014a, in Beuther H., Klessen R. S., Dullemond C. P., Henning T., eds, *Protostars and Planets VI*. p. 739, doi:10.2458/azu_uapress_9780816531240-ch032
- Madhusudhan N., Crouzet N., McCullough P. R., Deming D., Hedges C., 2014b, *ApJ*, 791, L9
- Madhusudhan N., Amin M. A., Kennedy G. M., 2014c, *ApJ*, 794, L12
- Madhusudhan N., Nixon M. C., Welbanks L., Piette A. A. A., Booth R. A., 2020, *ApJ*, 891, L7
- Madhusudhan N., Piette A. A. A., Constantinou S., 2021, *ApJ*, 918, 1
- Malik M., et al., 2017, *AJ*, 153, 56
- Malik M., Kempton E. M. R., Koll D. D. B., Mansfield M., Bean J. L., Kite E., 2019, *ApJ*, 886, 142
- Mandell A. M., Haynes K., Sinukoff E., Madhusudhan N., Burrows A., Deming D., 2013, *ApJ*, 779, 128
- Mansfield M., et al., 2022, *AJ*, 163, 261
- Mao S., Paczynski B., 1991, *ApJ*, 374, L37
- Marley M. S., McKay C. P., 1999, *Icarus*, 138, 268
- Marley M. S., Ackerman A. S., Cuzzi J. N., Kitzmann D., 2013, in Mackwell S. J., Simon-Miller A. A., Harder J. W., Bullock M. A., eds, *Comparative Climatology of Terrestrial Planets*. University of Arizona Press, pp 367–392, doi:10.2458/azu_uapress_9780816530595-ch015
- Marois C., Macintosh B., Barman T., Zuckerman B., Song I., Patience J., Lafrenière D., Doyon R., 2008, *Science*, 322, 1348
- Marois C., Zuckerman B., Konopacky Q. M., Macintosh B., Barman T., 2010, *Nature*, 468, 1080
- Márquez-Neila P., Fisher C., Sznitman R., Heng K., 2018, *Nature Astronomy*, 2, 719
- Mayne N. J., Baraffe I., Acreman D. M., Smith C., Wood N., Amundsen D. S., Thuburn J., Jackson D. R., 2014, *Geoscientific Model Development*, 7, 3059
- Mayor M., Queloz D., 1995, *Nature*, 378, 355

- Mazevet S., Licari A., Chabrier G., Potekhin A. Y., 2019, *A&A*, 621, A128
- Meadows V. S., Barnes R. K., 2018, in Deeg H. J., Belmonte J. A., eds, *Handbook of Exoplanets*. Springer, p. 57, doi:10.1007/978-3-319-55333-7_57
- Mendonça J. M., Grimm S. L., Grosheintz L., Heng K., 2016, *ApJ*, 829, 115
- Merino N., Aronson H. S., Bojanova D. P., Feyhl-Buska J., Wong M. L., Zhang S., Giovannelli D., 2019, *Frontiers in Microbiology*, 10, 780
- Millot M., Coppari F., Rygg J. R., Correa Barrios A., Hamel S., Swift D. C., Eggert J. H., 2019, *Nature*, 569, 251
- Mollière P., Wardenier J. P., van Boekel R., Henning T., Molaverdikhani K., Snellen I. A. G., 2019, *A&A*, 627, A67
- Montet B. T., et al., 2015, *ApJ*, 809, 25
- Morgan W. W., Keenan P. C., Kellman E., 1943, *An atlas of stellar spectra, with an outline of spectral classification*. University of Chicago Press
- Morley C. V., Fortney J. J., Kempton E. M. R., Marley M. S., Visscher C., Zahnle K., 2013, *ApJ*, 775, 33
- Morley C. V., Knutson H., Line M., Fortney J. J., Thorngren D., Marley M. S., Teal D., Lupu R., 2017a, *AJ*, 153, 86
- Morley C. V., Kreidberg L., Rustamkulov Z., Robinson T., Fortney J. J., 2017b, *ApJ*, 850, 121
- Moses J. I., Fouchet T., Bézard B., Gladstone G. R., Lellouch E., Feuchtgruber H., 2005, *Journal of Geophysical Research (Planets)*, 110, E08001
- Moses J. I., et al., 2011, *ApJ*, 737, 15
- Moses J. I., et al., 2013, *ApJ*, 777, 34
- Mousis O., Deleuil M., Aguichine A., Marcq E., Naar J., Aguirre L. A., Brugger B., Gonçalves T., 2020, *ApJ*, 896, L22
- Mróz P., et al., 2018, *AJ*, 155, 121
- Mulders G. D., 2018, in Deeg H. J., Belmonte J. A., eds, *Handbook of Exoplanets*. Springer, p. 153, doi:10.1007/978-3-319-55333-7_153
- Mulders G. D., Pascucci I., Apai D., 2015, *ApJ*, 814, 130
- Murgas F., Chen G., Nortmann L., Palle E., Nowak G., 2020, *A&A*, 641, A158
- Murray C. D., Correia A. C. M., 2010, in Seager S., ed., *Exoplanets*. University of Arizona Press, pp 15–23
- Nayak M., Lupu R., Marley M. S., Fortney J. J., Robinson T., Lewis N., 2017, *PASP*, 129, 034401

- Nesvorný D., Morbidelli A., 2008, *ApJ*, 688, 636
- Nettelmann N., Fortney J. J., Kramm U., Redmer R., 2011, *ApJ*, 733, 2
- Nettelmann N., Wang K., Fortney J. J., Hamel S., Yellamilli S., Bethkenhagen M., Redmer R., 2016, *Icarus*, 275, 107
- Nixon M. C., Madhusudhan N., 2020, *MNRAS*, 496, 269
- Nixon M. C., Madhusudhan N., 2021, *MNRAS*, 505, 3414
- Nixon M. C., Madhusudhan N., 2022, arXiv e-prints, p. arXiv:2201.03532
- Noack L., et al., 2016, *Icarus*, 277, 215
- Öberg K. I., Murray-Clay R., Bergin E. A., 2011, *ApJ*, 743, L16
- Oreshenko M., et al., 2017, *ApJ*, 847, L3
- Otegi J. F., Dorn C., Helled R., Bouchy F., Haldemann J., Alibert Y., 2020, *A&A*, 640, A135
- Owen J. E., 2019, *Annual Review of Earth and Planetary Sciences*, 47, 67
- Owen J. E., Wu Y., 2013, *ApJ*, 775, 105
- Parmentier V., Crossfield I. J. M., 2018, in Deeg H. J., Belmonte J. A., eds, *Handbook of Exoplanets*. Springer, p. 116, doi:10.1007/978-3-319-55333-7_116
- Parmentier V., et al., 2018, *A&A*, 617, A110
- Pedregosa F., et al., 2011, *Journal of Machine Learning Research*, 12, 2825
- Petigura E. A., Howard A. W., Marcy G. W., 2013, *Proceedings of the National Academy of Science*, 110, 19273
- Pierrehumbert R., Gaidos E., 2011, *ApJ*, 734, L13
- Piette A. A. A., Madhusudhan N., 2020a, *MNRAS*, 497, 5136
- Piette A. A. A., Madhusudhan N., 2020b, *ApJ*, 904, 154
- Piette A. A. A., Madhusudhan N., Mandell A. M., 2022, *MNRAS*, 511, 2565
- Pinhas A., Madhusudhan N., 2017, *MNRAS*, 471, 4355
- Pinhas A., Rackham B. V., Madhusudhan N., Apai D., 2018, *MNRAS*, 480, 5314
- Pinhas A., Madhusudhan N., Gandhi S., MacDonald R., 2019, *MNRAS*, 482, 1485
- Pluriel W., Zingales T., Leconte J., Parmentier V., 2020, *A&A*, 636, A66
- Pluriel W., Leconte J., Parmentier V., Zingales T., Falco A., Selsis F., Borde P., 2021, arXiv e-prints, p. arXiv:2110.09080
- Podolak M., Helled R., Schubert G., 2019, *MNRAS*, 487, 2653

- Polichtchouk I., Cho J. Y. K., 2012, *MNRAS*, 424, 1307
- Pollack J. B., Hubickyj O., Bodenheimer P., Lissauer J. J., Podolak M., Greenzweig Y., 1996, *Icarus*, 124, 62
- Pont F., Knutson H., Gilliland R. L., Moutou C., Charbonneau D., 2008, *MNRAS*, 385, 109
- Pont F., Sing D. K., Gibson N. P., Aigrain S., Henry G., Husnoo N., 2013, *MNRAS*, 432, 2917
- Quirrenbach A., Reffert S., Bergmann C., 2011, in Schuh S., Drechsel H., Heber U., eds, *American Institute of Physics Conference Series Vol. 1331, Planetary Systems Beyond the Main Sequence*. pp 102–109, doi:10.1063/1.3556189
- Rackham B., et al., 2017, *ApJ*, 834, 151
- Ramirez R. M., Levi A., 2018, *MNRAS*, 477, 4627
- Rauer H., et al., 2014, *Experimental Astronomy*, 38, 249
- Rauer H., Aerts C., Cabrera J., PLATO Team 2016, *Astronomische Nachrichten*, 337, 961
- Rauscher E., Menou K., 2010, *ApJ*, 714, 1334
- Richard C., et al., 2012, *J. Quant. Spectrosc. Radiative Transfer*, 113, 1276
- Richards J. W., et al., 2011, *ApJ*, 733, 10
- Ricker G. R., et al., 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 014003
- Robert C., Casella G., 2004, *Monte Carlo statistical methods*. Springer Verlag
- Robinson T. D., Fortney J. J., Hubbard W. B., 2017, *ApJ*, 850, 128
- Rodgers C. D., 2000, *Inverse Methods for Atmospheric Sounding - Theory and Practice*, *Inverse Methods for Atmospheric Sounding - Theory and Practice. Series: Series on Atmospheric Oceanic and Planetary Physics*, ISBN: 9789812813718. World Scientific Publishing Co. Pte. Ltd., Edited by Clive D. Rodgers, vol. 2, doi:10.1142/9789812813718
- Rogers L. A., 2015, *ApJ*, 801, 41
- Rogers L. A., Seager S., 2010a, *ApJ*, 712, 974
- Rogers L. A., Seager S., 2010b, *ApJ*, 716, 1208
- Rogers L. A., Bodenheimer P., Lissauer J. J., Seager S., 2011, *ApJ*, 738, 59
- Rothman L. S., et al., 1998, *J. Quant. Spectrosc. Radiative Transfer*, 60, 665
- Rothman L. S., et al., 2010, *J. Quant. Spectrosc. Radiative Transfer*, 111, 2139
- Rothman L., et al., 2013, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 130, 4

- Roux N. L., Bengio Y., 2008, *Neural Computation*, 20, 1631
- Safronov V. S., 1969, *Evoliutsiia doplanetnogo oblaka*. Nauka
- Salpeter E. E., Zapolsky H. S., 1967, *Phys. Rev.*, 158, 876
- Saumon D., Chabrier G., van Horn H. M., 1995, *ApJS*, 99, 713
- Schad D. J., Nicenboim B., Burkner P.-C., Betancourt M., Vasishth S., 2022, *Psychological Methods*, 26
- Schlichting H. E., Young E. D., 2022, *PSJ*, 3, 127
- Seager S., Deming D., 2010, *ARA&A*, 48, 631
- Seager S., Sasselov D. D., 1998, *ApJ*, 502, L157
- Seager S., Kuchner M., Hier-Majumder C. A., Militzer B., 2007, *ApJ*, 669, 1279
- Seager S., Bains W., Hu R., 2013, *ApJ*, 777, 95
- Seager S., Huang J., Petkowski J. J., Pajusalu M., 2020, *Nature Astronomy*, 4, 802
- Sedaghati E., Boffin H. M. J., Csizmadia S., Gibson N., Kabath P., Mallonn M., Van den Ancker M. E., 2015, *A&A*, 576, L11
- Selsis F., Despois D., Parisot J. P., 2002, *A&A*, 388, 985
- Shah O., Alibert Y., Helled R., Mezger K., 2021, *A&A*, 646, A162
- Shallue C. J., Vanderburg A., 2018, *AJ*, 155, 94
- Sheppard K. B., Mandell A. M., Tamburo P., Gandhi S., Pinhas A., Madhusudhan N., Deming D., 2017, *ApJ*, 850, L32
- Showman A. P., Guillot T., 2002, *A&A*, 385, 166
- Showman A. P., Cooper C. S., Fortney J. J., Marley M. S., 2008, *ApJ*, 682, 559
- Showman A. P., Fortney J. J., Lian Y., Marley M. S., Freedman R. S., Knutson H. A., Charbonneau D., 2009, *ApJ*, 699, 564
- Showman A. P., Tan X., Parmentier V., 2020, *Space Sci. Rev.*, 216
- Shporer A., Hu R., 2015, *AJ*, 150, 112
- Sing D. K., Vidal-Madjar A., Désert J. M., Lecavelier des Etangs A., Ballester G., 2008, *ApJ*, 686, 658
- Sing D. K., et al., 2012, *MNRAS*, 426, 1663
- Sing D. K., et al., 2016, *Nature*, 529, 59

- Skilling J., 2004, in Fischer R., Preuss R., Toussaint U. V., eds, American Institute of Physics Conference Series Vol. 735, Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering. pp 395–405, doi:10.1063/1.1835238
- Skilling J., 2006, *Bayesian Analysis*, 1, 833
- Smolensky P., 1986, in Rumelhart D. E., McClelland J. L., PDP Research Group C., eds, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 1. MIT Press, Cambridge, MA, USA, pp 194–281, <http://dl.acm.org/citation.cfm?id=104279.104290>
- Snellen I. A. G., de Kok R. J., de Mooij E. J. W., Albrecht S., 2010, *Nature*, 465, 1049
- Snellen I., et al., 2015, *A&A*, 576, A59
- Soboczenski F., et al., 2018, *Bayesian Deep Learning for Exoplanet Atmospheric Retrieval*. Neural Information Processing Systems Foundation, Inc.
- Soderlund K. M., et al., 2020, *Space Sci. Rev.*, 216, 80
- Sorenson D. C., 1982, *SIAM Journal on Numerical Analysis*, 19, 409
- Sotin C., Grasset O., Mocquet A., 2007, *Icarus*, 191, 337
- Soubiran F., Militzer B., 2015, *ApJ*, 806, 228
- Southworth J., 2011, *MNRAS*, 417, 2166
- Spergel D., et al., 2015, arXiv e-prints, p. arXiv:1503.03757
- Spyratos P., et al., 2021, *MNRAS*, 506, 2853
- Stassun K. G., Collins K. A., Gaudi B. S., 2017, *AJ*, 153, 136
- Stevenson K. B., et al., 2017, *AJ*, 153, 68
- Sugimura E., Komabayashi T., Hirose K., Sata N., Ohishi Y., Dubrovinsky L. S., 2010, *Phys. Rev. B*, 82, 134103
- Swain M. R., Vasisht G., Tinetti G., 2008, *Nature*, 452, 329
- Taylor J., Parmentier V., Irwin P. G. J., Aigrain S., Lee E. K. H., Krissansen-Totton J., 2020, *MNRAS*, 493, 4342
- Tennyson J., Yurchenko S. N., 2012, *MNRAS*, 425, 21
- Tennyson J., et al., 2016, *Journal of Molecular Spectroscopy*, 327, 73
- Thomas S. W., 2016, PhD thesis, University of Cambridge
- Thomas S. W., Madhusudhan N., 2016, *MNRAS*, 458, 1330
- Thrastarson H. T., Cho J. Y.-K., 2010, *ApJ*, 716, 144

- Tinetti G., et al., 2007, *Nature*, 448, 169
- Tinetti G., et al., 2018, *Experimental Astronomy*, 46, 135
- Trotta R., 2017, arXiv e-prints, p. arXiv:1701.01467
- Tsiaras A., et al., 2018, *AJ*, 155, 156
- Tsiaras A., Waldmann I. P., Tinetti G., Tennyson J., Yurchenko S. N., 2019, *Nature Astronomy*, 3, 1086
- Tsiganis K., Gomes R., Morbidelli A., Levison H. F., 2005, *Nature*, 435, 459
- Turbet M., Bolmont E., Ehrenreich D., Gratier P., Leconte J., Selsis F., Hara N., Lovis C., 2020, *A&A*, 638, A41
- Udalski A., Szymanski M., Kaluzny J., Kubiak M., Mateo M., 1992, *Acta Astron.*, 42, 253
- Valderrama-Bahamóndez G. I., Fröhlich H., 2019, *Frontiers in Applied Mathematics and Statistics*, 5
- Valencia D., O'Connell R. J., Sasselov D., 2006, *Icarus*, 181, 545
- Valencia D., Sasselov D. D., O'Connell R. J., 2007, *ApJ*, 665, 1413
- Valencia D., Ikoma M., Guillot T., Nettelmann N., 2010, *A&A*, 516, A20
- Valencia D., Guillot T., Parmentier V., Freedman R. S., 2013, *ApJ*, 775, 10
- Venot O., Hébrard E., Agúndez M., Dobrijevic M., Selsis F., Hersant F., Iro N., Bounaceur R., 2012, *A&A*, 546, A43
- Venot O., et al., 2020, *ApJ*, 890, 176
- Venturini J., Helled R., 2017, *ApJ*, 848, 95
- Vinet P., Rose J. H., Ferrante J., Smith J. R., 1989, *Journal of Physics Condensed Matter*, 1, 1941
- Virtanen P., et al., 2020, *Nature Methods*, 17, 261
- Visscher C., Moses J. I., 2011, *ApJ*, 738, 72
- Wagner W., Pruß A., 2002, *Journal of Physical and Chemical Reference Data*, 31, 387
- Wakeford H. R., Sing D. K., 2015, *A&A*, 573, A122
- Wakeford H. R., et al., 2017, *Science*, 356, 628
- Wakeford H. R., et al., 2018, *AJ*, 155, 29
- Waldmann I. P., 2016, *ApJ*, 820, 107
- Waldmann I. P., Tinetti G., Rocchetto M., Barton E. J., Yurchenko S. N., Tennyson J., 2015a, *ApJ*, 802, 107

- Waldmann I. P., Rocchetto M., Tinetti G., Barton E. J., Yurchenko S. N., Tennyson J., 2015b, *ApJ*, 813, 13
- Wardenier J. P., Parmentier V., Lee E. K. H., 2022, *MNRAS*, 510, 620
- Way M. J., et al., 2017, *ApJS*, 231, 12
- Weaver I. C., et al., 2021, *AJ*, 161, 278
- Welbanks L., Madhusudhan N., 2019, *AJ*, 157, 206
- Welbanks L., Madhusudhan N., 2021, *ApJ*, 913, 114
- Welbanks L., Madhusudhan N., 2022, *ApJ*, 933, 79
- Welbanks L., Madhusudhan N., Allard N. F., Hubeny I., Spiegelman F., Leininger T., 2019, *ApJ*, 887, L20
- Wetherill G. W., 1980, *ARA&A*, 18, 77
- Williams J. P., Cieza L. A., 2011, *ARA&A*, 49, 67
- Wilson J., et al., 2020, *MNRAS*, 497, 5155
- Winn J. N., 2010, in Seager S., ed., *Exoplanets*. University of Arizona Press, pp 55–77
- de Wit J., Seager S., 2013, *Science*, 342, 1473
- Witte S., Helling C., Barman T., Heidrich N., Hauschildt P. H., 2011, *A&A*, 529, A44
- Wolszczan A., Frail D. A., 1992, *Nature*, 355, 145
- Wong M. H., Mahaffy P. R., Atreya S. K., Niemann H. B., Owen T. C., 2004, *Icarus*, 171, 153
- Wordsworth R., Pierrehumbert R., 2014, *ApJ*, 785, L20
- Wordsworth R. D., Forget F., Selsis F., Millour E., Charnay B., Madeleine J.-B., 2011, *ApJ*, 733, L48
- Wright J. T., Gaudi B. S., 2013, in Oswalt T. D., French L. M., Kalas P., eds, *Planets, Stars and Stellar Systems. Volume 3: Solar and Stellar Planetary Systems*. Springer, p. 489, doi:10.1007/978-94-007-5606-9_10
- Wytttenbach A., Ehrenreich D., Lovis C., Udry S., Pepe F., 2015, *A&A*, 577, A62
- Xiao L.-H., Chen P.-R., Gou Z.-P., Li Y.-Z., Li M., Xiang L.-C., Feng P., 2017, *Asian Journal of Andrology*, 19, 586
- Youdin A. N., 2011, *ApJ*, 742, 38
- Yu X., et al., 2021, *Nature Astronomy*, 5, 822
- Yurchenko S. N., Tennyson J., 2014, *MNRAS*, 440, 1649

-
- Yurchenko S. N., Barber R. J., Tennyson J., 2011, MNRAS, 413, 1828
- Zalesky J. A., Line M. R., Schneider A. C., Patience J., 2019, ApJ, 877, 24
- Zapolsky H. S., Salpeter E. E., 1969, ApJ, 158, 809
- Zeng L., Sasselov D., 2013, PASP, 125, 227
- Zeng L., Sasselov D., 2014, ApJ, 784, 96
- Zeng L., et al., 2019, Proceedings of the National Academy of Science, 116, 9723
- Zhang M., et al., 2018, AJ, 155, 83
- Zhang M., Chachan Y., Kempton E. M. R., Knutson H. A., 2019, PASP, 131, 034501
- Zhang M., Chachan Y., Kempton E. M. R., Knutson H. A., Chang W. H., 2020, ApJ, 899, 27
- Zhu W., Dong S., 2021, ARA&A, 59, 291
- Zingales T., Waldmann I. P., 2018, ApJ, 156, 268

