

Modelling the Spectra of Brown Dwarfs



Ryan Garland
Trinity College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2018

Abstract

Brown dwarfs (BDs) are substellar mass objects with masses $\leq 80M_{Jup}$, the hydrogen-burning limit. Deuterium fusion occurs in objects with $\geq 13M_{Jup}$, lasting only a very short period (4-50 Myrs). This means that brown dwarfs are constantly cooling over their lifetime, passing through temperature ranges equivalent to those of exoplanets and planets in the Solar System. This aging and cooling creates a spectral sequence (hottest to coolest: M, L, T, Y) as various molecules and condensates begin to form. The atmospheres created during this sequence are extremely similar to gas giant planets. As brown dwarfs are free-floating objects, it is possible to attain much higher resolution spectra than with exoplanets because we do not have to subtract out the light from the host star, and thus they provide the perfect testbed for atmospheric modelling of objects beyond our Solar System. The aim of this thesis is to characterise the coolest brown dwarf atmospheres, spectrally classified as Y dwarfs, and the associated change in their atmospheric structure and composition with decreasing temperature. In order to achieve this ultimate goal, a number of steps are taken. The first is the updating of the molecular and atomic data to include the latest line lists and pressure broadening parameters, which enables us to more accurately model BDs with our radiative transfer calculations. Along the way, we discuss the most effective method to calculate gaseous absorption, and illustrate the effects of using sub-optimal techniques. With the radiative transfer calculations optimised, we proceed to implement a new inverse method to our NEMESIS retrieval code: Nested Sampling. This algorithm employs state-of-the-art Bayesian statistics to fit spectra and derive probability distributions for each parameter of interest. We verify our new Bayesian method and radiative transfer code against two other codes from separate groups, and investigate the model dependency of parameters using realistic atmospheric models of exoplanets that may be probed by the James Webb Space Telescope. We discuss the model-dependency of each parameter in the context of changing signal-to-noise, finding that some retrieved parameters are more model-dependent than others. Finally, having updated and verified our code through the methods previously discussed, we apply it to a range of ultracool T and Y dwarf atmospheres. The retrieved probability distributions

of each parameter representing the atmospheric structure and composition of the brown dwarfs are compared to one another to elucidate trends with decreasing temperature. For the coldest object that we analysed, we discovered a cloud that is consistent with a Na_2S cloud, which is expected to form at these temperatures. This is the first Bayesian detection of such a feature on a Y dwarf. We also discuss the plausibility that this cloud may instead be a water cloud. The analysis of these objects and the work presented here helps elucidate the atmospheric structure and composition of the Y dwarf population. We hope that this work will assist not only in future analyses of Y dwarf atmospheres, but also more generally in any atmosphere within or beyond our Solar System.

Acknowledgements

First and foremost, I must thank my primary supervisor, Prof. Patrick Irwin, for being, in my view, the best supervisor I could have hoped for. Always he asked the right questions, pushed me when I needed it, encouraged and supported my ventures into new academic areas... and above all, was good craic. This one is for the pints in the Gardener's, the bagpipes on the rooftop, the muddy summer barbecues, and the enlightening conferences. Thanks for everything, Pat. I've enjoyed every second of the last few years (except making k -tables, that was grim).

I must also thank my second supervisor, Prof. Suzanne Aigrain, for providing a keen outside eye on my niche atmospheric work, and encouraging me to work in interdisciplinary fields. I appreciate all the support that you have given me.

My time here in Oxford has been a magical one. This place imbues you with a sense of wonder and inspiration, and an immense feeling of privilege that I have tried hard not to forget. I would like to thank my college, Trinity, and my collegemates who have welcomed me into their family, and made me feel at home on my first step off the Emerald Isle.

A huge (hungover) thanks goes out to my besties throughout the years: Rohini Giles, Ellie Armstrong, Ashley Fisher, Sam Coles, Daniel Toledo Carrasco, Sophie Bauduin, Arrate Antunano Martin, Juan Alday Parejo, Ashwin Braude, and Jake Taylor. You are all the best!

I thank my girlfriend, Jessica, for trotting back and forth from Belfast to Birmingham to Oxford and back again, repeatedly, for a year and a half... I know how dreadful that is. But I am super excited to ditch those long distance woes for very short distance Netflix nights, sour beers, and fistfuls of pizza (synonyms of love). Favourites together.

I extend the biggest thanks imaginable to my family who have loved and supported

me from the moment I left, every step of the way. Thank you, Mummy and Daddy, you are the best parents a son could ask for - I love you very much!

I dedicate this thesis to my grandfather, Tim O'Regan, for his continuous loving support of me and my whole family, and to my grandmother, Margaret O'Regan, an inspirational woman to us all, whom I wish dearly could be with us today.

Contents

1	Introduction	7
1.1	The First Brown Dwarfs	7
1.2	Definition of a Brown Dwarf	8
1.3	The Life of a Brown Dwarf	10
1.3.1	Formation Theory	10
1.3.2	Spectral Classification	11
1.3.3	Summary of Spectral Types	14
1.4	The L-T Transition, Clouds and Variability	16
1.5	The T-Y Transition and Y dwarfs	22
2	Radiative Transfer & NEMESIS	28
2.1	Spectral Line Theory	28
2.1.1	Line intensity	30
2.1.2	Line broadening	31
2.1.3	Doppler broadening	31
2.1.4	Pressure broadening	32
2.1.5	Voigt Profile	33
2.2	Solving the Radiative Transfer Equation	34
2.2.1	Line-by-line	35
2.2.2	Correlated-k approximation	35
2.2.3	Cross-sections	38
2.2.4	Opacity Sampling	38

3	Inverse Methods	39
3.1	Optimal Estimation	40
3.2	Markov Chain Monte Carlo	42
3.2.1	Bayesian Inference	42
3.2.2	Markov Chain Monte Carlo	43
3.2.3	EMCEE	45
3.3	Nested Sampling	48
3.3.1	MultiNest	52
4	Gaseous Absorption in Brown Dwarfs and Exoplanets	55
4.1	Introduction	55
4.2	Line Lists	58
4.3	Pressure broadening coefficients	59
4.4	Calculation of k -Tables	63
4.5	Pressure Broadening Effect on k -Tables	65
4.5.1	CO	65
4.5.2	VO	67
4.5.3	TiO	69
4.5.4	H ₂ O	69
4.5.5	NH ₃	69
4.5.6	CO ₂	73
4.5.7	CH ₄	73
4.5.8	Summary of H ₂ -He Pressure-broadening effect on k -tables	75
4.6	Cross-section Comparison	77
4.7	Results	81
4.7.1	Primary Transit	85
4.7.2	Brown Dwarfs	85
4.7.3	Primary Transit and Secondary Eclipse	90
4.8	Conclusions	90
5	Cross-Comparison Retrievals for JWST	94
5.1	Introduction	94

5.2	Forward Model Comparison	95
5.2.1	Individual Gas Comparison	96
5.2.2	HD189733b	100
5.2.3	Super-Earth, Low Mean Molecular Weight	102
5.2.4	Super-Earth, High Mean Molecular Weight	102
5.3	Retrieval Comparison	104
5.3.1	HD189733b	105
5.3.2	Super Earth Low MMW	107
5.3.3	Super Earth High MMW	110
5.4	Conclusions	110
6	T and Y dwarf Retrievals	115
6.1	Data Sources	116
6.1.1	Data Sources	116
6.1.2	Binning Procedure	118
6.1.3	Currently Estimated Properties	118
6.2	Methods	120
6.2.1	Temperature Profile	120
6.2.2	Cloud Profile	124
6.3	Benchmark Retrieval: GL570D	129
6.3.1	Results	129
6.4	Y Dwarf Retrievals	135
6.4.1	Temperature and Spectral Typing	138
6.4.2	Gaseous Composition	142
6.4.3	Mass, Radius, Surface Gravity	145
6.4.4	Clouds	146
6.5	Conclusions	152
7	Conclusions	154
7.1	Summary	155
7.2	Future Work	159
7.2.1	Efficiently Solving the RTE	159

7.2.2	Informing WFIRST	167
A	Full T and Y dwarf Posteriors	176

Chapter 1

Introduction

1.1 The First Brown Dwarfs

In a series of abstracts published in 1962-1963, Shiv S. Kumar theoretically inferred the existence of brown dwarfs (BDs) by realising that “there exists a limiting mass below which a contracting star cannot reach the main-sequence stage because the temperature and density at the centre are too low for hydrogen-burning to start” (Kumar, 1963).

It was not until 1995, more than thirty years later, that his theoretical work was confirmed by the first definitive discovery of a brown dwarf, Teide 1 (Rebolo et al., 1995), a lithium-burning, free-floating brown dwarf straddling the limit of the hydrogen-fusion boundary. In the following edition of *Nature*, the first detection of a planet (Pegasus 51 b) orbiting a Sun-like star was published (Mayor & Queloz, 1995). In the very same edition, the discovery of the first methane-rich brown dwarf, Gliese 229B, a much cooler brown dwarf than Teide 1, was presented (Nakajima et al., 1995).

The discovery of Gliese 229B led to the first link between stellar astrophysics and planetary science. Continued work on Gliese 229B showed that it had a complex, molecular atmosphere, with absorption bands due to methane (CH_4), water (H_2O), carbon monoxide (CO) and other atomic/molecular sources (see Figure 1.1.1 taken from Oppenheimer et al. (1998)). Non-equilibrium chemistry was also identified due to the apparent imbalance between CO and CH_4 , indicating the possibility of strong winds and deep convection.

For a further five years only a handful of brown dwarfs would be discovered. Since




Figure 1.1.1: A spectrum of Gliese 229B from optical to thermal IR bands. Spectral regions with horizontal bars correspond to wavelengths where the Earth’s atmosphere is too opaque to permit collection of useful data from the ground. Labelled horizontal bars indicate absorption bands due to methane (CH₄), water (H₂O) and carbon monoxide (CO). Two atomic lines are also indicated and are due to neutral caesium (Cs I). Figure from Oppenheimer et al. (1998).

the turn of the millennium, however, we have amassed almost 1,800 confirmed objects (Gagne, 2018), leading to an immensely data-led field of research.

1.2 Definition of a Brown Dwarf

The definition of a brown dwarf has been hotly debated, but in general there are two major camps in the brown dwarf community. The International Astronomical Union (IAU) defines brown dwarfs as “substellar objects with true masses above the limiting mass for thermochemical fusion of deuterium” (substellar meaning below the hydrogen-burning limit of $\sim 0.075 M_{\odot}$ for solar composition) and stipulates that it does not matter “how they are formed nor where they are located” (Boss et al., 2007). For a solar composition object, this limiting mass is given as $13 M_{\text{Jup}}$. Any mass below this value is therefore considered a planet. The main disadvantage of this definition is that there may be some

overlap between brown dwarfs and very large planets, and it does not discriminate between any parameter other than mass. As well as this, it is very difficult to measure the mass of a free-floating brown dwarf because the only measurable signal of the mass is spectroscopic, which is a second-order effect on the spectrum. Brown dwarfs in binary systems do however have measurable dynamical masses through astrometric means (Dupuy & Liu, 2017), which are more accurate. However, this definition is advantageous to use due to its simplicity and ease of use when other defining information is not known.

The second camp resides in formation theory, stating that brown dwarfs form in the same way as stars, i.e. the fragmentation and gravitational collapse of a cloud of molecular gas and dust, while planetary formation is more consistent with core accretion (Chabrier et al., 2014). This is perhaps the more scientifically rigorous definition, but requires knowledge of the formation history of the substellar object to ascertain whether it is a brown dwarf or a planet, which may or may not be readily available.

The definition of a BD adopted by this thesis is from Chabrier et al. (2014), where they present clear arguments for their definition which is (1) consistent with observations, (2) based on a physically-motivated justification, invoking two drastically different *dominant* formation mechanisms, and (3) allows a mass overlap between these two populations, as is indeed supported by observations, but excluded by the IAU definition. It is given as follows.

Brown dwarfs are (1) all free-floating objects below the hydrogen-burning minimum mass, $\sim 0.075 M_{\odot}$ for solar composition (Chabrier et al., 2000b), no matter whether or not they are massive enough to ignite deuterium burning, and (2) objects that are companions to a parent star or other BD, but exhibit compositional and mechanical (mass-radius) properties consistent with the ones of a gaseous sphere of *global chemical composition similar to the one of the parent star/BD*. In contrast, “giant planets” are defined as companions of a central, significantly more massive object, a *necessary condition*, whose bulk properties strongly depart from the ones just mentioned.




Figure 1.3.1: Comparison of sizes and effective temperatures of the Sun, the red dwarf star Gliese 229A, the young brown dwarf Teide 1, the old brown dwarf Gliese 229B, one of the coolest known brown dwarfs WISE 1828 C 2650, and Jupiter. The effective temperature of a body such as a star or planet is the temperature of a black body of the same radius that would emit the same total amount of electromagnetic radiation. Figure from Joergens (2014).

1.3 The Life of a Brown Dwarf

1.3.1 Formation Theory

The current formation theory for brown dwarfs most consistent with observations is that of gravoturbulent fragmentation (Hennebelle & Chabrier, 2011). This theory proposes that large-scale turbulence is injected at the molecular cloud scale by various sources and cascades to smaller scales by shocks, generating a field of density fluctuations down to the dissipative scale. Overdense regions form, resulting in gravitational collapse and thus form self-gravitating cores which isolate themselves from the surrounding medium. For a more comprehensive view of formation theory, see Luhman (2012). Regardless of the finer details of formation mechanisms, it is generally agreed that brown dwarfs form in

much the same way as stars.

Objects which cannot gravitationally overcome the electron degeneracy pressure of the forming, partially degenerate electron gas core ($< 0.075 M_{\odot}$), will never reach the critical temperature necessary to ignite hydrogen fusion, and are thus classified as a “failed star” or brown dwarf. The scaling of the radius with respect to mass is also affected by electron degeneracy, giving a relationship of $R \propto M^{-1/3}$ (Burrows & Liebert, 1993). This means that for the entire range of BD masses, $1 - 75M_{\text{Jup}}$, the radius of a brown dwarf changes by no more than 50% and is always close to R_{Jup} .

Eighty years ago, Wigner & Huntington (1935) postulated that under high pressures molecular hydrogen would undergo a phase transition to the metallic state. Given suitable pressure ($P > 3$ Mbar), such a substance must form the bulk of Jupiter ($\sim 85\%$), Saturn ($\sim 50\%$), and brown dwarfs ($> 99.9\%$). Above this enormous, fully convective zone (Mohanty et al., 2007) on brown dwarfs is a thin slither of atmosphere with a swirling soup of clouds, condensates, and atomic and molecular gases.

1.3.2 Spectral Classification

Once fully formed, a brown dwarf over $13M_{\text{Jup}}$ will enjoy a brief (4-50 Myrs) period of deuterium and possibly lithium fusion (for larger masses), before cooling down for the rest of its lifetime by radiating energy from its upper atmosphere, the photosphere (Chabrier et al., 2000a). It is because of this cooling that we have four spectral classifications of BDs: M, L, T, and Y Dwarfs, which are further discussed in Subsection 1.3.3.

These spectral types are not equivalent to the standard Morgan-Keenan stellar spectral types system because their age-mass-temperature degeneracy causes the classical stellar temperature relationship to break down for such ultra-cool objects (Helling & Casewell, 2014). A BD with a specific spectral type (or effective temperature) may be higher mass and old, or lower mass and young, i.e. a young L5 dwarf ($T_{\text{eff}} = 1500\text{K}$) in the Pleiades cluster (125 Myr) has a mass of $25M_{\text{Jup}}$, whereas an old (3-5 Gyr) field L5 dwarf will have a mass of $70M_{\text{Jup}}$ (Chabrier et al., 2000b) and $T_{\text{eff}} = 1500\text{K}$. It should be made clear that, whatever the object’s atmospheric conditions, spectral types are a method of classification not implicitly related to particular atmospheric features such as the measured temperature, but instead a way of grouping similar spectra.

Spectral type	Special Characteristics	Chemistry
M8	TiO (705.3, 843.2 nm) at max. (M99, (K99))	TiO at max
M9	VO (740, 790 nm) at max (M99, (K99))	VO at max
L0	blueward TiO gone, redder TiO remains (K99,M99); Al I, Ca I disappear in J band (Mc03); Rb I (780.0, 794.8 nm); Cs I (852.1, 894.3 nm) weak (K99, M99)	Ca, Al, Ti cond
L1	VO bands (740, 790nm) weaken; Na I (589.0 & 589.6 nm) doublet weakens; Rb I, Cs I increasing in strength (K99,M99)	V cond Cs>Cs ⁺
L2	all TiO bands gone except at 843.2 nm (K99); Ca I triplet still in K band (Mc03)	
L3	Fe I at 1.189 m still present (Mc03)	
L3.5	Fe I at 1.189 m gone (Mc03); K I (1.169, 1.197, 1.25m) peak in strength (R01)	Fe cond.
L4	all VO gone (K99); CrH (861.1nm) = FeH (869.2nm); FeH at 869.2 nm & J band at max. (R01, B02, K99, M99, Mc03); K I doublets in J-band at max (B02, Mc03)	FeH $\approx$ CrH K at max.
L5	CH ₄ appears at 3.3 m (Noll et al. 2000); CaH at 685 nm still seen (M99, Mc03); FeH (869.2 nm & J band) weakens (B02,K99,M99); CrH 861.1 nm at max. (B02, K99, M99); optical Rb I, Cs I strengthening, but weaker in J band (K99, Mc03)	
L6	CrH weakens in optical (K99, M99) Li I (670.8nm) still observed in some objects (K99, M99)	Cr cond Li = LiCl
L7	TiO 843.2nm (K99, M99) essentially absent, ditto TiO 549.7, 559.7, 615.9, 615.9, 638.4nm (M99); Li I most likely absent (K99,M99)	
L8	subtle CH ₄ band in K band (G02, N04); no or subtle CH ₄ in H band (B02,G02,Mc03); Na I (589.0, 589.6 nm) very broad, Na I barely seen in optical and J band (K99,M99,Mc03); K I doublets at min. in J band (N04)	CO=CH ₄
T0	CH ₄ appears in K and H bands (G02) FeH (989.6 nm) weaker (N04) or as strong as in L8 (B02b)	CO=CH ₄
T1	weak CH ₄ in H, J, & K bands (B02,G02); strengths of CO=CH ₄ in K band; K I (1.2432, 1.2522m) present in J band (B02)	CO=CH ₄
T2	strengths of CH ₄ >CO in K band (B02); CrH (861.1 nm) & FeH (869.2nm) disappear (B03); FeH (989.6 nm) weaker (N04) or as strong as in L8 (B02b); Cs I (852.1, 894.3nm) peak in strength (B03)	
T3	CO (2.295m) very weak or absent in K band (B02, Mc03,N04)	
T4	FeH (989.6nm), CrH (996.9nm) increase again in strength (B02,B02b,B03,N04)	
T5	No CO in K band (B02); FeH (989.6nm) & CrH (996.9nm) peak in strength (B02b, B03,N04); K I at maximum in J band (B02, Mc03,N04)	Na ₂ S cond
T6	K I at in J band begins to weaken (B02)	
T7	K I in J band very weak (B02, Mc03)	
T8	no FeH (989.6nm) (B02, Mc03); no K I in J band (B02, Mc03); Cs I marginally detectable in optical (B03)	(NH ₃ =N ₂)?
Y0	excess NH ₃ absorption from 1.53 to 1.58 μ m; smooth change in width of the J-band peak; rapid fall in the flux of the blue wing of the H-band (C11)	
Y1	J-band peak narrows further; H-band to J-band peak ratio increases (K12)	
Y2	near-equal heights of the J- and H-band peaks (K12)	

B02: Burgasser et al. (2002b), Burgasser et al. (2002a); G02: Geballe et al. (2002); K99: Kirkpatrick et al. (1999) Kirkpatrick et al. (2000); M99: Martín et al. (1999); Mc03: McLean et al. (2003); N04: Nakajima et al. (2004), R01: Reid et al. (2001); C11 = Cushing et al. (2011); K12 = Kirkpatrick et al. (2012)

H band centred at $\sim 1.6 \mu\text{m}$; J band centred at $\sim 1.25 \mu\text{m}$; K band centred at $\sim 2.2 \mu\text{m}$

Table 1.3.1: Spectral Classification of L and T Dwarfs, adapted from Lodders & Fegley (2006).




Figure 1.3.2: The near-infrared spectral sequence from late-M dwarfs to early-Y brown dwarfs. All spectra are normalised to 1 at their peak and then a flux offset is applied (Helling & Casewell, 2014). Some major gaseous absorption features are labelled.

Figure 1.3.2 exemplifies visually the spectral differences between each spectral classification, and a more specific representation of the definitions is outlined in Table 1.3.1 adapted from Lodders & Fegley (2006), with additional information provided for Y dwarfs. Figure 1.3.3 translates these spectral classifications into approximate effective temperatures.

Throughout this thesis we will repeatedly refer to different photometric bandpasses, such as the commonly used infrared (IR) bandpasses J , H , and K . Hence, for clarity, the wavelengths and bandwidths are summarised in Table 1.3.2, taken from Binney & Merrifield (1998). Note that this is not a particular scheme of bandpasses, and should be largely similar to the various schemes adopted in the brown dwarf community. These bandpasses are designed to calculate photometric magnitudes by integrating directly over the spectral flux in their respective wavelength ranges, taking into account the sensitivity of the bandpass to the flux at different wavelengths.




Figure 1.3.3: Effective temperature as a function of spectral type. Color represents the $J - K$ color-magnitude. L dwarf spectral types are based on red optical spectra and T and Y dwarf spectral types are based on near-infrared spectra. A gap is left between L8 and T0 because there is no L9 subtype in the optical sequence. Temperatures for the L dwarfs and T dwarfs with spectral types earlier than T9 are from Vrba et al. (2004). However, only dwarfs with trigonometric parallax errors less than 30% of the parallax value are plotted. Temperatures for the T dwarfs with spectral types later than T8 are from Leggett et al. (2009, 2012) while the temperatures for the Y dwarfs are from Leggett et al. (2013). WISE J1828C2650 is currently typed as $\geq Y2$ and so is plotted as a solitary object beyond a spectral type of Y2. Figure from Joergens (2014).

1.3.3 Summary of Spectral Types

M dwarfs are the youngest, hottest, and most massive brown dwarfs (but not necessarily *both* young and massive, due to the age-mass-temperature degeneracy). They are spectrally similar to older, low-mass dwarf stars, with strong metal-oxide molecular bands (e.g. TiO, VO, CO and H₂O), neutral atomic line absorption (Na I, K I), and metal hydrides absorption (MgH, CaH). Late M dwarfs (i.e. $\geq M6$; early M dwarfs would be $\leq M3$, and mid-M in between) are characterised by increasingly stronger TiO, VO, H₂O, FeH, CrH, CaH, and MgH bands. Water bands grow stronger from mid-M dwarfs through

Letter	Effective Wavelength Midpoint λ_{eff} (μm)	FWHM (Bandwidth) $\Delta\lambda$ (μm)
U	0.365	0.066
B	0.445	0.094
V	0.551	0.088
R	0.658	0.138
I	0.806	0.149
Z	0.900	0.137
Y	1.020	0.120
J	1.220	0.213
H	1.630	0.307
K	2.190	0.390
L	3.450	0.472
M	4.750	0.460
N	10.50	2.500
Q	21.00	5.800

Table 1.3.2: Summary of Photometric Bandpasses, taken from Binney & Merrifield (1998). Note that this is not a particular scheme of bandpasses, and should be largely similar to the various schemes adopted in the brown dwarf community.

the L and T dwarf sequence in all spectral ranges.

The L dwarf sequence starts when TiO and VO bands begin to weaken while CrH, FeH, and water bands strengthen at optical wavelengths. In the near infrared, CO, CrH and FeH absorptions strengthen towards the mid-L subclasses, and then decline as CH₄ begins to form.

The change between the L and T dwarfs, often called the L-T transition region is characterised by the near-infrared colours of the brown dwarfs changing, while the effective temperature of the object remains the same (see Figure 1.4.2).

T dwarfs are characterised by the methane absorption seen in the near-IR that gets progressively stronger as one progresses through the subclasses, making the $J - K$ and $H - K$ colours bluer. In the optical, the spectrum is affected by collisionally-induced molecular hydrogen absorption and FeH absorption.

In 2011, the latest spectral type was discovered: the Y dwarf. According to Gagne (2018) (last accessed 5th June 2018), there are 27 classified Y dwarfs. All the Y dwarfs show deep H₂O and CH₄ absorption bands in their near-infrared spectra, similar to late-T dwarfs (Figure 1.3.2). The J band peaks of the Y dwarfs are narrower than those of the latest type T dwarfs.

The J and H spectral bandpass magnitudes become approximately equal, meaning

that the $J - H$ (and similarly $Y - J$) trend towards the blue for T dwarfs turns back towards the red for Y0, as seen in Figure 1.3.4, i.e., we are establishing color as going from blue to red as we go from left to right in the diagram. The $J - H$ change is mediated by the change in effective temperature and hence the corresponding Planck function, while the $Y - J$ shift is thought to occur as alkali atoms (which normally dominate shorter wavelength absorption) are being bound to molecules or condensates, thus reducing the alkali atom opacity (Kirkpatrick et al., 2012; Morley et al., 2014a). Lodders & Fegley (2006) predicted that as the effective temperature decreases in Y dwarfs, the H band becomes less shaped by CH_4 and H_2O , and NH_3 absorption becomes important. We measure this prediction in our own work in Chapter 6 and find consistent results.

Species including CO , CH_4 , N_2 , NH_3 , H_2O , Fe , Cr , CaTiO_3 , Na_2S , MnS , ZnS and KCl become critical, not only for modelling the opacity as a function of wavelength through the atmosphere, but also for determining the physical structure of the atmosphere (i.e. whether condensates are present).

1.4 The L-T Transition, Clouds and Variability

At the beginning of the L dwarf spectral sequence ($\sim 2200\text{K}$), refractory species including iron, silicates and metal oxides condense to form dust clouds (Burrows & Sharp, 1999; Lodders, 1999; Burrows et al., 2006). As the temperature decreases, the formation and thickening of dust clouds induces increasingly redder spectral energy distributions (SEDs), characterising the L spectral sequence. When the temperature reaches $\sim 1200\text{K}$, the dust opacity abruptly diminishes, marking the transition to cloud-free and methane-rich T spectral types. This transition from cloudy to cloud-free atmosphere is known as the L-T transition, and is characterised by its dramatic blueward shift of ~ 2 magnitudes in $J - K$ encompassing $\sim \text{L8}$ to $\sim \text{T5}$ spectral types, at a nearly constant effective temperature (Golimowski et al., 2004; Stephens et al., 2009). This is represented graphically in Figure 1.3.4.

It is thought that the disappearance of dust clouds at the L-T transition occurs as dust grains gravitationally settle below the photosphere. Analogously, Jupiter’s photosphere contains prominent ammonia clouds, but thick iron and silicate clouds are thought to re-




Figure 1.3.4: Color-magnitude ($J-K$ color against H) diagram of stars (gray) classified by the *Morgan-Keenan* classification system, M dwarfs (orange), L dwarfs (green), T dwarfs (light blue), and Y dwarfs (blue). A number of directly imaged planets or planet candidates are also shown (bold circles). Figure from Bowler (2016). Labelled ‘blue’ and ‘red’ on the x-axis for clarifying color direction.




Figure 1.4.1: Pressure-temperature profiles of brown dwarf model atmospheres by Morley et al. (2012) are plotted, with the condensation curves for a range of condensate species. The 1-6 μm photosphere (range of pressures over which contributions to the spectrum occur, i.e. optical depth $\tau < 1$) of each model is shown as a thicker line. Cloudless models are in blue, and cloudy models with sedimentation efficiency $f_{\text{sed}} = 2$ (red) and 4 (orange). A high sedimentation efficiency parameter f_{sed} results in vertically thinner clouds with larger particle sizes, whereas a lower f_{sed} results in more vertically extended clouds with smaller particle sizes. Condensation curves for each condensate species are plotted. The cloudy models include the condensates Cr, MnS, Na₂S, ZnS, and KCl. Figure from Morley et al. (2012).

side in its deep atmosphere, hidden from view (Lodders & Fegley, 2006). Unfortunately, this settling of condensates remains poorly understood, with models generally predicting a much more gradual disappearance of clouds over a wider range of temperatures (Tsuji & Nakajima, 2003; Marley et al., 2002; Allard et al., 2003). For example, the ~ 1 micron fluxes of L-T transition brown dwarfs approximately double in value, whereas models predict monotonically decreasing fluxes. Due to this discrepancy, it is suggested that, rather than uniformly settling below the photosphere, clouds are dynamically disrupted at L-T transition temperatures, opening up windows to the deep photosphere, which contribute to the abrupt decline in cloud opacity and resurgence of 1 micron flux (Ackerman & Marley, 2001; Burgasser et al., 2002a).




Figure 1.4.2: NIR spectral type versus $J - K$ colour for all targets observed in Radigan et al. (2014)’s program. Grey points show the population of known L and T dwarfs with $J < 16.5$. Purple circles show detections of large-scale photometric variability, with the symbol size drawn proportional to the peak-to-peak amplitude of variability detected in the J and K photometric bands (ranging from 0.9-9% for the smallest to largest symbols). Non-detections, where variability has been searched for but not detected, are shown in pink. A grey dashed line encircles the objects considered part of the L/T transition sample. From Radigan et al. (2014).

Cloud features should rotate in and out of view in such patchy cloud coverage, producing rotationally modulated variability in the object’s SED (rotation periods of BDs ~ 2 –10 hours, Reiners & Basri (2008)). The first of the unambiguous detections of brown dwarf variability in the NIR with amplitudes greatly surpassing the photometric noise level and/or multi-epoch evidence was reported by Artigau et al. (2009). A T2.5 dwarf was reported to be variable with a peak-to-peak amplitude of ~ 50 mmag in J , with a period of 2.4 hours (10σ confidence level). Since then a number of isolated objects within the L-T transition region have been reported with large-amplitude variations (Radigan et al., 2012; Girardin et al., 2013; Gillon et al., 2013; Biller et al., 2013), and a survey conducted by Radigan et al. (2014) inferred that $\sim 40\%$ of L9-T3.5 dwarfs are periodic variables with peak-to-peak amplitudes $> 2\%$. Figure 1.4.2 shows the findings of Radigan et al. (2014) graphically.

These results suggest that cloud dispersal at the L-T transition proceeds in a spatially


Cloud Holes : Large Color Variations**Cloud Thickness: Small Color Variations**

Figure 1.4.3: Diagram illustrating potential mechanisms for greater colour differences caused by cloud inhomogeneities. Figure from Apai et al. (2013).

non-uniform manner, leading to localised regions of clouds and clearings (Buenzli et al., 2014), or alternatively thick and thin cloud patches (Kostov & Apai, 2013), as demonstrated in Figure 1.4.3. Using Doppler imaging techniques, Crossfield et al. (2014) have made a global cloud map of Luhman 16B (an L-T transition dwarf) which shows that the spectroscopic temporal variability may be mapped by three large cloud inhomogeneities, as seen in Figure 1.4.4.

The technique of Doppler Imaging relies on the varying Doppler shifts across the face of a rotating object. As darker regions rotate across the visible face of the brown dwarf, the Doppler-broadened absorption line profiles exhibit deviations at the projected radial velocities of the darker areas. Features near the equator cause changes across the entire line profile and move across the full span of velocities; features at higher latitudes move more slowly, experience smaller Doppler shifts, and affect a narrower range of velocities.

Taking Luhman 16B’s rotation period and considering that evolutionary models predict these objects to be 1.0 ± 0.2 times the radius of Jupiter, Crossfield et al. (2014) concluded that Luhman 16B’s rotation axis must be inclined ≤ 30 deg from the plane of the sky; i.e., we are viewing this brown dwarf nearly equator-on. Note that this means




Figure 1.4.4: Surface map of brown dwarf Luhman 16B, which clearly depicts a bright near-polar region (seen in the upper-right panels) and a darker mid-latitude area (lower-left panels) consistent with large-scale cloud inhomogeneities. The lightest and darkest regions shown correspond to brightness variations of roughly $\pm 10\%$. The time index of each projection is indicated near the center of the figure. Figure from Crossfield et al. (2014).

the features have an assumed orientation, and potentially could be either northern or southern.

Other mechanisms for surface brightness heterogeneities are also possible. Notably, M and early-L dwarfs are known to have elevated magnetospheric activity (Hallinan et al., 2008; Berger et al., 2010). The coupling of magnetic fields with the atmosphere could result in either hot or cold spots, which may be difficult to distinguish from cloud structures with similar temperature differentials. There is evidence from radio observations that flaring activity extends even into the T dwarfs (Route & Wolszczan, 2012). However, in general the neutral atmospheres of L and T dwarfs are considered too electrically resistive to support magnetic starspots (Mohanty et al., 2002; Gelino et al., 2002; Chabrier

& Küker, 2006). More recently, temperature fluctuations driven by deep atmospheric instabilities (Robinson & Marley, 2014) or by jet or eddy circulation in the stratified layer above the convective zone (Zhang & Showman, 2014) have also been invoked as sources of brightness heterogeneities on brown dwarfs. The former process may be important in late T dwarfs, where the radiative timescale is long enough that the temperature fluctuations dissipate on timescales longer than a minute.

1.5 The T-Y Transition and Y dwarfs

With decreasing effective temperature, it becomes more and more difficult to measure a model-independent bolometric luminosity because wavelengths that are easily accessible from the ground contain a smaller and smaller fraction of the total emergent flux. For example, at T_{eff} of 300 K, only 10% of the emergent luminosity emerges at $<2.5 \mu\text{m}$. Therefore, the effective temperature estimates of the brown dwarfs with spectral types later than T8 are based on atmospheric model fits to their near-infrared spectra Leggett et al. (2013), as is used in Figure 1.3.3. This is because at these temperatures, the spectral energy distribution has shifted considerably into the mid-infrared. We note here that, due to this difficulty, the scientific community have discovered only ~ 25 Y dwarfs, and correspondingly there are only a handful of studies (which we will mention here) that have compared measured spectral / photometric data to atmospheric models or investigated variability at the T-Y transition.

As seen in Figures 1.4.1 and 1.5.1, many different condensates are expected to form within the full range of pressure-temperature profiles expected for brown dwarfs, including Y dwarfs. Morley et al. (2012) investigated the effects of the emergence of Cr, MnS, Na_2S , ZnS, and KCl clouds in cool (late T or T-Y transition) atmospheres, and found that, particularly Na_2S may be a more natural explanation of the near- and mid-infrared colours of observed T dwarfs, rather than the re-emergence of silicate clouds that wane at the L-T transition.

Rajan et al. (2015) conducted a pilot study monitoring an initial sample of three late-T dwarfs (T6.5, T8 and T8.5) and one Y dwarf (Y0) for infrared photometric variability at multiple epochs. With J -band imaging, each target was observed for a period of 1.0-




Figure 1.5.1: Pressure-temperature profiles for three representative temperatures and two gravities are shown from Morley et al. (2014b) models. The thicker black line indicates the location of the 1-6 μm photosphere. The shaded salmon region shows where the atmosphere is convective. The dashed lines show condensation curves for each substance expected to condense in thermochemical equilibrium. The curve represents the pressure-temperature points at which the partial pressure of the gas is equal to the saturation vapour pressure; to the left of the curve, the partial pressure of each gas is higher than the saturation vapour pressure, and the excess vapour will form a cloud. The kinks in the profile in the upper atmosphere are numerical and do not represent real features in the atmospheres of Y dwarfs. Figure from Morley et al. (2014b).




Figure 1.5.2: Top: colour-magnitude diagram (CMD) for the two Spitzer photometric bands (3.6 and $4.5\mu\text{m}$) showing a variety of models, including BT-Settl (Allard et al. models (orange, dot-dashed)) and a variety of Morley models. Three varieties of Morley models are shown, one cloud-free, one with sulphide clouds ($f_{\text{sed}} = 5$), and one incorporating water clouds. Bottom: a near-IR CMD for the same set of models. The BT-Settl models provide a good fit to the $J-H$ colors and absolute magnitudes for the warmer objects, whereas the Morley objects do a better job on the colder objects at these wavelengths. Figure from Beichman et al. (2014).

4.5 hr per epoch, which covers a significant fraction of the expected rotational period. During the first epoch, the T8.5 target WISE0458 exhibited variations with a remarkable peak-to-peak amplitude of 13%, while the second epoch light curve taken ~ 2 years later did not note any variability to a 3% upper limit. The three remaining targets showed no significant variations, with a photometric precision of between 0.8-20.0%, depending on the target brightness.

Trucks et al. (2014) presented preliminary results for fourteen Y dwarfs for photometric variability at 3.6 and 4.5 microns with the Spitzer Space Telescope, which included significant Y dwarf variability in the near-infrared.

Morley et al. (2014a) have also investigated the effect of inhomogeneities in salt, sulphide and water cloud structures, similar to the L-T transition arguments made previously, as well as hot spots. They found that for objects over $T_{\text{eff}} \sim 375$ K, patchy Na_2S cloud opacity would generate the largest amplitude variability within near-infrared spectral windows. For objects under 375 K, water clouds also become important and generate larger amplitude variability in the mid-infrared.

Morley et al. (2014b) found that while water condenses high in the atmospheres of all objects below $T_{\text{eff}} \sim 400$ K, these clouds do not become optically thick until the object has cooled to 350-375 K. This result means that for the current set of Y dwarfs warmer than 400 K, water clouds will not strongly affect their spectra, and may be more affected by Na_2S . They also noted that water clouds, unlike other clouds in brown dwarf atmospheres, are very much nongray absorbers; Na_2S , while also nongray, produces much smoother absorption across its absorption spectrum. Water clouds cause a blueward shift in the near-infrared ($J - H$) because water clouds do not absorb as strongly in J as they do in H or K .

Beichman et al. (2014) investigated 15 WISE (Wide-field Infrared Survey Explorer)-discovered late-type brown dwarfs (six T8-9 and nine Y dwarfs) using the Keck-II telescope, the Spitzer Space Telescope, and the Hubble Space Telescope. In Figure 1.5.2, they show clearly the need for cloudy models to accurately model the photometry in the near-infrared and mid-infrared of Y dwarfs.

Skemer et al. (2016) obtained a 4.5-5.3 μm spectrum of WISE0855, the coldest known brown dwarf. This is the same bandpass used to study Jupiter's deep thermal emission.

Their spectrum is strikingly similar to Jupiter’s and reveals the presence of atmospheric water vapour and clouds.

Leggett et al. (2016) monitored photometrically the Y0 brown dwarf WISE1738 at both near- and mid-infrared wavelengths. They interpret their results as the Y dwarf having a 6.0 ± 0.1 hr rotation period, with one or more large-scale surface features being a source of variability (of a 3 hour period, $\leq 2\sigma$ detection level). Intriguingly, this variability and the wavelength dependence can be reproduced by atmospheric models (Morley et al., 2014a) that include patchy KCl and Na₂S clouds and associated small changes in surface temperature. The small number of large features, as well as the timescale for evolution of the features, is very similar to what is seen in the atmospheres of the solar system gas giants. Gelino & Marley (2000) used observations of Jupiter to show that, were it to be unresolved, 20% variability would be observed at $4.8\mu\text{m}$ with a period equal to the planets rotation period of 9.9 hr. The variability in this case is predominantly due to Jupiters Great Red Spot. Simon et al. (2016) present a 49-day light curve for Neptune using Kepler. The data are compared to contemporaneous images taken with the Keck telescope at $1.65\mu\text{m}$ and to Hubble Space Telescope (HST) visible imaging taken several months later. The authors find that a single large, long-lived storm that is seen in their Keck images dominates the Kepler and HST light curves. The periodicity of the long-term variability is consistent with the planets rotation and surface wind speed at the latitude of the storm. The short-term variability is interpreted as being due to smaller or fainter clouds.

Cushing et al. (2016) presented the first detection of the photometric variability in a spectroscopically confirmed Y dwarf using the 3.6 and $4.5\mu\text{m}$ Spitzer bands. They found that the light curves are nearly sinusoidal in form, and have periods of roughly 8.5 hr, with semi-amplitudes of 3.5%. They find that a simple geometric spot model with a single bright spot reproduces these observations well. After using the static 1D models of Morley and collaborators, they find that they cannot reproduce the 3.6 and $4.5\mu\text{m}$ semi-amplitudes simultaneously. This hints at the requirement of more advanced two- or three- dimensional models that include time-dependent phenomena like vertical mixing, cloud formation, and thermal relaxation.

Leggett et al. (2017) estimated properties of a sample of 24 Y dwarfs based on near-

infrared spectra and photometry. They confirm that nonequilibrium chemistry is important for these objects. Nonequilibrium cloud-free models reproduce well the near-infrared spectra and mid-infrared photometry of the warmer Y dwarfs with $425 \leq T_{\text{eff}} \leq 450\text{K}$. A small amount of cloud cover may improve the model fits in the near-infrared for the Y dwarfs with $325\text{K} \leq T_{\text{eff}} \leq 375\text{K}$. Neither cloudy nor cloud-free models reproduce the near-infrared photometry for the $T_{\text{eff}} = 250\text{K}$ Y dwarf WISE 0855. They estimate that the metallicities and ages for the sample are generally solar-like, and that the known Y dwarfs are 3-20 Jupiter-mass objects with ages 0.6-8.5 Gyr.

Chapter 2

Radiative Transfer & NEMESIS

The **N**on-linear optimal **E**stimator for **M**ultivariate **S**pectral analysis code (NEMESIS) is a radiative transfer retrieval code created by Irwin et al. (2008) as an all-purpose retrieval code that can be used to model the atmosphere of any planet, exoplanet, moon or brown dwarf, using data from any instrument. Until now it has only been used for the Solar System planets (e.g. Giles et al. (2015); Fletcher et al. (2012); Irwin et al. (2004)), moons (de Kok et al., 2008) and exoplanets (Lee et al., 2014; Barstow et al., 2014). This thesis is the first application to brown dwarfs. In order to discuss how the radiative transfer equation is solved by NEMESIS, and how the retrieval code works, a large array of physical and statistical processes must be understood. In this chapter we will discuss the physical processes. We discuss the statistical processes in Chapter 3.

2.1 Spectral Line Theory

Spectral lines originate from the interaction of a photon with an atomic nucleus or molecular nuclei. The specific structure of the molecule/atom, i.e. the number of electrons, protons, neutrons, decides the set of configurations in which these particles can coexist, and hence a set of energy levels describing their rotational, vibrational, and translational degrees of freedom. Atoms and molecules prefer to be in their ‘ground state’, where the configuration is such that the amount of potential energy is minimised. A change in energy level, δE occurs when a photon is absorbed or emitted that has an energy equal to δE , creating an absorption (decrease in intensity) or emission (increase in intensity)

spectral line.

Most transitions in the visible spectrum are electronic transitions within atoms, where electrons transition between different atomic energy levels. Molecular absorption becomes much more important in the infrared. As molecules possess both rotational and vibrational degrees of freedom, infrared absorption allows transitions between these rotational-vibrational states.

Rotational energy levels occur due to the rotation of the molecule around a fixed orientation in space, centred on the centre of mass of the molecule. Linear molecules can be described by a single moment of inertia and rotational quantum number J , while nonlinear molecules require an extra moment of inertia and rotational quantum number K (which satisfies $|K| \leq J$). (Paula, 2006).

A non-zero electric dipole moment is required for rotational transitions to occur. Some molecules, such as homonuclear diatomic molecules like H_2 , have a net zero electric dipole moment, thus pure rotational transitions cannot occur. However, collisions may produce a transient electric dipole in H_2 , leading to collision-induced absorption. For H_2 -dominated atmospheres, such as the gas giants and brown dwarfs, the infrared, contains significant H_2 - H_2 and H_2 -He collision-induced absorption. (Irwin et al., 2004)

Vibrational energy levels occur due to the stretching and bending of bonds between atoms. These systems, to a first approximation, oscillate with simple harmonic motion, and their energy levels depend on the masses of the atoms, the bond ‘stiffness’, and the vibrational quantum number ν .

Rotational-vibrational (ro-vibrational) bands occur where vibrational transitions also come with an associated rotational transition, giving rise to a combined ro-vibrational spectrum. Ro-vibrational lines are the most typical lines that will be relevant to this thesis.

Spectral lines can be used to uniquely identify the chemical composition of any medium with light passing through it. Because of their dependence on physical conditions, such as pressure and temperature, it is possible to identify the properties of the medium, i.e. the atmosphere of a brown dwarf or an exoplanet.

2.1.1 Line intensity

The transmittance $\mathcal{T}$ of a path of gas is calculated as

$$\mathcal{T} = \exp(-ku) \quad (2.1)$$

where $u = \int_0^z N dz$ is the path amount (molecule cm^{-2}) calculated along a column of length z (cm) with gaseous number density N (molecule cm^{-2}), and k is the absorption coefficient of a single line ($\text{cm}^2 \text{ molecule}^{-1}$). The absorption coefficient for a single line can be expressed as

$$k(\tilde{\nu}) = \mathcal{S}_i \Phi_i(\tilde{\nu}) \quad (2.2)$$

where we have introduced the line intensity $\mathcal{S}_i$ ($\text{cm}^{-1} (\text{molecule cm}^{-2})^{-1}$) and line profile defined as $\int_0^\infty \Phi_i(\tilde{\nu}) d\nu = 1$ for each transition i . Note that in reality we are looking at ro-vibrational molecular line absorptions, where we are summing over thousands of lines. The line intensity of a given molecular absorption line is given by

$$\mathcal{S}_i(T) = \frac{I_a}{8\pi c \tilde{\nu}_0^2} \frac{g_u e^{-E_l/k_B T}}{Q(T)} \left(1 - e^{-hc\tilde{\nu}_0/k_B T}\right) A_i \quad (2.3)$$

where c is the speed of light, I_a is the fractional isotopic abundance (technically the line strength is not dependent on I_a , however the apparent line strength in the atmosphere is), E_l is the energy of the lower level, g_u is the degeneracy of the upper level, k_B is Boltzmann's constant, T is the temperature, $Q(T)$ is the partition function at T for the species, $\tilde{\nu}_0$ is the wavenumber of the transition and A_i is the Einstein A-coefficient at a transition i . This formula assumes local thermodynamic equilibrium (LTE). This means that the medium has a well-defined temperature on a scale much greater than the mean free path of the photon, so that the Maxwellian temperature is equivalent to the Planckian temperature.

The line intensity is normally tabulated at a reference temperature $T_0 = 296\text{K}$. Thus to calculate the line intensity at any other temperature, we use the relation:

$$\mathcal{S}_i(T) = \mathcal{S}_i(T_0) \frac{Q(T_0)}{Q(T)} \frac{e^{-E_l/k_B T}}{e^{-E_l/k_B T_0}} \frac{(1 - e^{-hc\tilde{\nu}_0/k_B T})}{(1 - e^{-hc\tilde{\nu}_0/k_B T_0})} \quad (2.4)$$

2.1.2 Line broadening

For each transition i the line shape function $\Phi_i(\tilde{\nu})$ spreads the absorption coefficient over a wavenumber region. Various physical processes cause a broadening of absorption lines. Here we discuss pressure broadening and Doppler broadening, the dominant broadening processes in brown dwarf atmospheres, and the corresponding functional forms they take.

Note that for the temperature and pressure ranges involved in brown dwarf atmospheres, the broadening produced by natural broadening is negligible so we will consequently ignore natural broadening processes here. This broadening is due to the uncertainty in position when considering Heisenberg's uncertainty principle and the finite time over which a photon is absorbed/emitted. The uncertainty is of the order of $1 \times 10^{-7} \text{cm}^{-1}$. (Irwin et al., 2004).

2.1.3 Doppler broadening

While particles are absorbing and emitting photons they are in thermal motion. Particles will emit and absorb photons with wavenumber $\tilde{\nu}$ in their own frame of reference, causing a shift in the observed wavenumber of the photon. The thermal motion has no directional preference, and the accumulated effect is therefore a broadening of the line. This is known as Doppler broadening, because of the Doppler effect.

Again assuming LTE, the velocity distribution of molecules is given by the Maxwell-Boltzmann distribution for any direction, including the line of sight (Muller-Kirsten, 2013):

$$f_z(v_x) = \sqrt{\frac{m_z}{2\pi k_B T}} \exp(-v_x^2/v_0^2), \quad v_0 = \sqrt{\frac{2k_B T}{m_z}} \quad (2.5)$$

Defining the Doppler width as $\alpha_D = \tilde{\nu}_0 v_0/c$, it is possible to derive the Doppler profile:

$$\Phi_D(\tilde{\nu}) = \frac{1}{\sqrt{\pi}\alpha_D} e^{-(\tilde{\nu}-\tilde{\nu}_0)^2/\alpha_D^2} \quad (2.6)$$

and as before the associated absorption coefficient of a line assuming only Doppler broadening:

$$k(\tilde{\nu}) = \mathcal{S}_i \Phi_D(\tilde{\nu}). \quad (2.7)$$

The line shape produced is Gaussian with full width at half maximum (FWHM) is $2\ln(2)\alpha_D$. Note that $\alpha_D \propto \sqrt{T/m_z}$, so Doppler broadening increases with increasing temperature.

2.1.4 Pressure broadening

Lines will be broadened by interactions with other molecules, this is called pressure broadening. Collisions between particles effectively reduce the lifetime of the upper state of transitions, thereby causing a broadening of the line. Pressure broadening results in a Lorentzian line profile,

$$\Phi_L(\tilde{\nu}) = \frac{1}{\pi} \frac{\alpha_L}{(\tilde{\nu} - \tilde{\nu}_0)^2 + \alpha_L^2} \quad (2.8)$$

where we are considering the pressure broadening of α_L is the Lorentz width, with $2\alpha_L$ the full width at half maximum (FWHM). The pressure-broadened width α_L depends on both the pressure and temperature in a complex way, but the relationship is often approximated as (Sharp & Burrows, 2007)

$$\gamma(T) = \gamma_0 \frac{p}{p_0} \left(\frac{T_0}{T} \right)^n \quad (2.9)$$

where $\gamma(T) \equiv \alpha_L$, the pressure-broadened line half-width ($\text{cm}^{-1} \text{ atm}^{-1}$) at a standard pressure p_0 (normally 1 atm), γ_0 is the line half-width at a reference temperature T_0 (normally 296K), and n is the temperature exponent for the linewidth.

Both the reference Lorentz width γ_0 and the temperature exponent n are transition dependent, and are generally approximated via polynomial fits using the lower state of the transition's quantum total angular momentum J_{low} . The total foreign pressure broadening experienced by a gas will be given by their Lorentz width γ for each perturbing gas (for a given transition) weighted by their volume mixing ratios (see Chapter 4). The pressure

broadened width will thus be dominated by collisions with the most abundant molecules.

In HITRAN (discussed in more detail in Chapter 4) air- and self-broadening parameters are tabulated for each transition (Rothman et al., 2013). The atmospheres of brown dwarfs consist mainly of H_2 and He (approximately 85% and 15% respectively), and the broadening information provided in HITRAN is therefore not directly applicable. Information on pressure broadening by H_2 and He is extremely sparse (Freedman et al., 2008; Sharp & Burrows, 2007), and there is a large uncertainty in the values of γ_0 and n used in the literature. We discuss this issue in more detail in Chapter 4.

2.1.5 Voigt Profile

In situations where both Doppler and pressure broadening are important, as in brown dwarf atmospheres, we end up with a line profile which is a convolution of the two. This is the Voigt profile, given by

$$\Phi_V(\tilde{\nu}) = \frac{a}{\pi^{3/2}\alpha_D} \int_{-\infty}^{\infty} \frac{dy e^{-y^2}}{(v-y)^2 + a^2} \quad (2.10)$$

where

$$v = (\tilde{\nu} - \tilde{\nu}_0)/\alpha_D, \quad a = \frac{\alpha_L}{\alpha_D}, \quad \alpha_D = \frac{\tilde{\nu}_0}{c} \sqrt{\frac{2k_B T}{m_z}}. \quad (2.11)$$

For $a \rightarrow 0$, pure Doppler broadening is retrieved, while for $a \gg 1$, the shape becomes a Lorentzian. The integral must be solved numerically, and many computational algorithms exist. We use the algorithm of Humlicek (1982).

Unfortunately, real line profiles are not perfectly Voigtian (Thomas & Stamnes, 2002). The Voigt profile is fairly accurate provided weak interactions between molecules, but for stronger interactions, effects such as collisional narrowing, sub-Lorentzian effects and line shifts, where the line centre is shifted away from $\tilde{\nu}_0$, can occur. Line wings are particularly affected by this. We ignore these additional effects in our treatment of line profiles, which is the usual approach (Sharp & Burrows, 2007; Freedman et al., 2008), because the literature lacks sufficient data to represent these effects, and we expect these effects to contain errors that are secondary to other modelling errors (such as the correlated- k approximation). To avoid overestimating absorption in the line wings, we apply a

cut-off in the line profile at 25cm^{-1} from the line centre. Estimates typically chosen are in the range $20\text{-}100\text{cm}^{-1}$ range (Sharp & Burrows, 2007). 25cm^{-1} was chosen as it was relatively computationally inexpensive while providing enough width to cover the Lorentzian broadening, but not too much to overestimate the wings.

2.2 Solving the Radiative Transfer Equation

Radiative transfer models are calculated by solving the radiative transfer equation (RTE) for a given set of initial values for each variable parameter. For thermal emission the RTE is given by:

$$L_\nu = B_\nu(T_0)e^{-\tau_\nu(0)} + \int_0^{\tau_\nu(0)} B_\nu(T(\tau_\nu))e^{-\tau_\nu}d\tau_\nu(z) \quad (2.12)$$

where we have assumed the atmosphere is non-scattering. This is a reasonable assumption for brown dwarfs as the contribution to the radiance from scattering due to clouds is much smaller compared to the thermal emission of the atmosphere itself, and most helpfully this greatly simplifies the RTE. In equation 2.12, L_ν is the spectral radiance of the object (brown dwarf) at wavenumber ν due to radiation that has passed through its atmosphere from the surface to altitude z . $B_\nu(T_0)$ is the Planck function at surface temperature T_0 and wavenumber ν , $B_\nu(T)$ is the Planck function at temperature T and wavenumber ν for an altitude z , $\tau_\nu(z)$ is the optical depth (integrated opacity from space to altitude z), and $\tau_\nu(0)$ is the total optical depth at wavenumber ν of the atmosphere. The first term in the equation represents radiation emitted by the atmospheric gas, and the second term represents the attenuation of radiation between the surface and altitude z by the intervening atmosphere.

In order to solve this equation, knowledge of the strengths and positions of absorption lines for gaseous constituents in the atmosphere is required. One of the most computationally demanding parts of solving the radiative transfer equation is calculating the absorption coefficient, k . Because of this, a number of different methods are employed by various groups to calculate the gaseous absorption in radiative transfer calculations which have varying degrees of accuracy and computational cost. We outline the theory of the most prevalent techniques employed in the community in this section. The effects

of using these different techniques in solving the RTE are further explored in Chapter 4.

2.2.1 Line-by-line

The mean transmission $\bar{\mathcal{T}}$ of an atmospheric path over a particular wavenumber interval $\nu \rightarrow \nu + \Delta\nu$ can be calculated by

$$\bar{\mathcal{T}}(m) = \frac{1}{\Delta\nu} \int_{\nu_0}^{\nu_0+\Delta\nu} \exp\left(-m \sum_j k_j(\nu)\right) d\nu \quad (2.13)$$

where m is the total absorber amount (molecule cm^{-2}) and the sum is over all the lines of all the gases contributing to the absorption, where the absorption spectrum of the j th line in the sum (taking into account the gas volume mixing ratios) is $k_j(\nu)$ ($\text{cm}^2 \text{ molecule}^{-1}$). The line-by-line approach calculates the absorption coefficient, using Equation 2.2, individually for every line included over all relevant wavelengths and gases. Some hot gases, such as CH_4 (Yurchenko & Tennyson, 2014), can have over ten billion relevant lines. While this method is the most accurate to use, it is immediately apparent that the sheer number of lines makes this approach computationally expensive. For Bayesian retrieval algorithms (discussed further in Chapter 3) such as Nested Sampling (e.g. MultiNest (Feroz et al., 2013)) and Markov Chain Monte Carlo (MCMC, e.g. EMCEE (Foreman-Mackey et al., 2013)), the radiative transfer forward model will be executed millions of times while it explores a large parameter space. The line-by-line approach is therefore incompatible with these inverse methods. To increase the speed of calculating gaseous absorption, a number of different approaches may be taken.

2.2.2 Correlated-k approximation

A key observation to increase the speed of calculating the transmittance is to note that in the frequency range $d\nu$, it is not important where each absorption coefficient occurs in the range. What is important, is to realise what fraction of the frequency domain, $f(k)dk$, is occupied by absorption coefficients between $k \rightarrow k+dk$. Intuitively, if one takes several absorption lines, and bin them to create an average, it does not matter what the original orientation was - any permutation would create the same average. This intuition is fine for one layer, but, as we will discuss in more detail in Chapter 4, combining many layers,




Figure 2.2.1: The k -distribution technique applied to the absorption of methane in Saturn’s atmosphere. A high-resolution absorption coefficient spectrum of methane for the 1250-1260 cm^{-1} region is shown for typical conditions in Saturn’s atmosphere of 0.1 bar, 85 K and a volume mixing ratio of 4.5×10^{-3} . The left-hand plots show the highly detailed absorption spectra, together with the portions of the spectrum with absorption coefficients less than 0.009 and 0.2, respectively. Parts of the spectrum with absorption coefficients less than this are marked in black. The right-hand plots show the k -distributions, derived simply by ranking the absorption coefficients in increasing value, together with the sampled quadrature points (asterisks) for a 20-point numerical Gaussian integration. In this example we can see that just over 20% of the interval has absorption less than 0.009, while 84% has absorption less than 0.2. Figure from Irwin et al. (2008).

as is required for inhomogeneous atmospheres, requires a more sophisticated approach as this requires that calculations are done monochromatically (line-by-line). One of these more sophisticated approaches is called the correlated- k approximation, which we discuss here. We follow the same approach as Irwin et al. (2008).

We write equation 2.13 as:

$$\bar{\mathcal{T}}(m) = \int_0^\infty f(k) \exp(-km) dk. \quad (2.14)$$

Now defining a cumulative frequency function $g(k)$ as

$$g(k) = \int_0^k f(k) dk \quad (2.15)$$

which is a smooth, single-valued, monotonically increasing function. The inverse of $g(k)$ is also a unique smoothly varying function, $k(g)$. Using this definition, the equation 2.14 may be simplified to:

$$\bar{\mathcal{T}}(m) = \int_0^1 \exp(-k(g)m) dg \quad (2.16)$$

By discretising the equation 2.16, we can perform numerical integration using

$$\bar{\mathcal{T}}(m) = \sum_{i=1}^N \exp(-k_i m) \Delta g_i \quad (2.17)$$

where i is the i th quadrature point, k_i is the opacity of the k -distribution, and Δg_i is the quadrature weight. NEMESIS uses a Gauss-Legendre quadrature scheme to choose the quadrature points, typically with $N = 10 - 20$ ordinates. Larger N provides more accurate sampling, but is also more computationally expensive, so one must try to balance this trade-off.

Notice now that while $k(\nu)$ is a rapidly varying function of frequency, the k -distribution function $k(g)$ is a smoothly varying function of g , and hence requires far fewer integration steps. This is the biggest advantage of this scheme, and can be represented graphically, as seen in Figure 2.2.1 taken from Irwin et al. (2008). Within a single atmospheric layer, we may simply combine the k -distributions of different gases by assuming that the lines are randomly overlapping.

As was discussed briefly earlier, this is valid for a homogeneous path through an atmosphere. In real atmospheres, paths are inhomogeneous, and this inhomogeneity can be approximated by stacking homogeneous layers together. From layer-to-layer, assuming a sufficient number of homogeneous layers in the atmosphere, the transmission does not vary rapidly in adjacent layers, and so the transmission between these layers is generally well-correlated, making this approximation valid. This is where the name ‘correlated- k ’ comes from.

After calculating the k -distributions for all M layers, the mean transmission through all the layers (i.e. the entire atmosphere), can be approximated as

$$\bar{\mathcal{T}}(m) = \sum_{i=1}^N \exp \left(- \sum_{j=1}^M k_{ij} m_j \right) \Delta g_i. \quad (2.18)$$

This method is computationally inexpensive while producing errors that are less than 5% when compared to line-by-line methods (Irwin et al., 2008).

2.2.3 Cross-sections

Cross-sections are line-by-line calculations performed over a range of pressures and temperatures for a range of gases, and then subsequently binned to a specified spectral resolution (usually $\sim \Delta\nu = 1 \text{ cm}^{-1}$, Line et al. (2014)). Briefly, this method can be either very computationally inexpensive and hence inaccurate, or computationally expensive and hence accurate. This is due to the fact that transmission calculations must be performed monochromatically, and the binning process eliminates this. This is discussed in more detail in Chapter 4.

2.2.4 Opacity Sampling

The opacity sampling method (Hubeny & Mihalas, 2014) relies on having an ultra-high resolution cross-section that resolves the individual lines pre-calculated for each gas. For each wavelength bin, the high-resolution cross-section is then sampled a predefined number of times (typically at $\Delta\nu = 1 \text{ cm}^{-1}$ intervals) for each gas, and these samples are used in a pseudo-monochromatic calculation. While slower than using regular cross-sections, it can be more accurate with relatively less computational expense. It also has the advantage of being able to switch to ultra-high resolution spectra easily, as opposed to the correlated- k method which is usually binned to typical spectral resolutions of observations; and it does not need to make any assumptions about spectral correlation within the atmosphere like correlated- k . However, if the number of samples within each bin is insufficient, inaccuracies may be large. This is further discussed in Chapter 4.

Chapter 3

Inverse Methods

When attempting to establish the composition and thermal structure of an atmosphere, there are two principal approaches one can take. One is to solve the forward problem, which starts with the causes (atmospheric conditions) and then calculates the results (the spectrum). In brown dwarf and exoplanetary atmospheres, the forward problem is solved by using all available physical and chemical knowledge to calculate the state of the atmosphere, typically using either one-dimensional radiative transfer models ensuring radiative-convective equilibrium (where radiative and moist convective heating processes interact to result in a neutral thermodynamic state, e.g. Baudino et al. (2015); Saumon & Marley (2008); Morley et al. (2012)), or three-dimensional general circulation models (e.g. Amundsen et al., 2016; Sing et al., 2016; Showman & Kaspi, 2013), which apply discrete equations for fluid motion and integrate these forward in time. When the GCM is coupled to a radiative transfer scheme, it can begin to solve the forward problem by comparing its calculated spectra to observed spectra. Solving the forward problem in brown dwarf atmospheres is therefore finding the best fit to the data using precomputed spectra from the forward model.

The second approach is to solve the inverse problem, which starts with the causes (the spectrum) and then calculates the results (atmospheric conditions). Inverse methods solve the inverse problem: that is, from observational data, what is the set of parameters that enable the best fit to the data using the forward model. In our case, the forward model is the radiative transfer model, the set of parameters that describes the atmosphere of the brown dwarf/exoplanet of interest, and the observational data is the electromagnetic

spectrum collected by ground- or space-based telescopes. The forward model is typically a one-dimensional radiative transfer model (Barstow et al., 2014; Line et al., 2015; Waldmann et al., 2015). Solving the inverse problem in brown dwarf atmospheres is therefore finding the best fit to the data using spectra computed on-the-fly from the forward model.

In this chapter, we describe three prominent inverse methods used in atmospheric retrievals of brown dwarfs and exoplanets.

3.1 Optimal Estimation

NEMESIS, until now, has used non-linear optimal estimator formalism to minimise the difference between the measured and modelled spectra, subject to minimum departure from the *a priori* state vector, by minimising the cost function:

$$\phi = (\mathbf{y}_m - \mathbf{y}_n)^\top \mathbf{S}_\epsilon^{-1} (\mathbf{y}_m - \mathbf{y}_n) + (\mathbf{x}_n - \mathbf{x}_0)^\top \mathbf{S}_x^{-1} (\mathbf{x}_n - \mathbf{x}_0) \quad (3.1)$$

In equation 3.1, $\mathbf{y}_m$ is the measured spectrum, $\mathbf{y}_n$ is the spectrum calculated for the trial atmosphere, represented by a model state vector $\mathbf{x}_n$, $\mathbf{S}_\epsilon$ is the measurement covariance matrix (containing both measurement errors and estimated forward-modelling errors), $\mathbf{x}_0$ is the *a priori* state vector and $\mathbf{S}_x$ is the *a priori* covariance matrix. NEMESIS differs to usual optimal estimation formalism by adding systematic errors from the forward model to $\mathbf{S}_\epsilon$. This approach allows us to take account of the forward modelling errors in the retrieval results.

When retrieving properties that have strong measurement constraints, such as the temperature profile of the Earth from climatological records, the mean expected temperature profile and covariance, $\mathbf{x}_0$ and $\mathbf{S}_x$ respectively, are well known. However, many properties of planets beyond Earth lack strong measurement constraints (e.g. Solar System planets), or even any constraint at all (brown dwarfs and exoplanets). Here, the diagonal components of the *a priori* covariance matrix are set to the square of the estimated *a priori* errors and, for the case of retrieving continuous vertical profiles, off-diagonal elements are set to (after Rodgers (2000)):

$$S_{ij} = (S_{ii}S_{jj})^{\frac{1}{2}} \exp\left(-\frac{|\ln(P_i/P_j)|}{c}\right) \quad (3.2)$$

where P_i is the pressure at level i , and c is the correlation length scale, which decides how correlated in temperature (or any other continuous vertical quantity) pressure points i and j are. A value of $c = 1.5$ is commonly used.

The difference between the spectrum computed from the trial measurement vector $\mathbf{x}_n$ and that measured is used to calculate a new estimate of the trial vector $\mathbf{x}_{n+1}$ through the equation

$$\mathbf{x}_{n+1} = \mathbf{x}_0 + \mathbf{S}_x \mathbf{K}_n^\top (\mathbf{K}_n \mathbf{S}_x \mathbf{K}_n^\top + \mathbf{S}_\epsilon)^{-1} (\mathbf{y}_m - \mathbf{y}_n - \mathbf{K}_n (\mathbf{x}_0 - \mathbf{x}_n)) \quad (3.3)$$

where $\mathbf{K}_n$ is the matrix of functional derivatives of the n th iteration. This is also known as the Jacobian, i.e. the rate of change of calculated radiance with respect to state vector elements for all the wavelengths in the spectrum. In practice, $\mathbf{K}_n$ can vary greatly between iterations and the simple iteration scheme of Equation 3.3 can become unstable. Instead, NEMESIS uses a modified iteration scheme, based on the Marquardt-Levenberg principle where the actual modified state vector used in the next iteration, $\mathbf{x}'_{n+1}$ is calculated from $\mathbf{x}_{n+1}$ and $\mathbf{x}_n$ as

$$\mathbf{x}'_{n+1} = \mathbf{x}_n + \frac{\mathbf{x}_{n+1} - \mathbf{x}_n}{1 + \lambda} \quad (3.4)$$

λ is initially set to 1.0. If the spectrum calculated from $\mathbf{x}'_{n+1}$ reduces the cost function ϕ (Equation 3.1), then $\mathbf{x}_n$ is set to $\mathbf{x}'_{n+1}$, λ is multiplied by 0.3, and the iteration continues. If the cost function increases, then $\mathbf{x}_n$ is kept the same, λ is multiplied by 10.0, and a new $\mathbf{x}'_{n+1}$ is calculated. The choice of the multiplication parameters (0.3, 10) is somewhat arbitrary, although it is important to ensure that they are not reciprocals of one another, as this can potentially lead to endless loops.

The lower number (0.3) was chosen to be greater than the reciprocal of the larger number so that λ did not decrease too quickly. As the retrieval approaches convergence, $\lambda \rightarrow 0$, the model tends to the optimal estimate, whose error is estimated as

$$\hat{S} = (S_x^{-1} + K_n^\top S_\varepsilon^{-1} K_n)^{-1} \quad (3.5)$$

Optimal estimation has a number of advantages over other inverse methods. If a large amount of information about the system is known *a priori*, then the prior can limit the parameter space to a small region. Because of the gradient descent the algorithm performs in a localised prior-space, the best fit may be found within approximately 20 iterations. This is extremely efficient compared to the other two inverse methods discussed in this thesis. In principle, if the prior is near the global minimum in cost-function space (or global maximum in probability space), then this algorithm would retrieve, with just a few iterations, the same parameter values as other algorithms which use millions of iterations.

However, this is a big ‘if’. In practice, a grid of priors must be used to map out the posterior space. Because of the potentially large number of parameters involved, it can be difficult to probe the entire prior space efficiently. It may also be difficult or even practically impossible to prove that your best fit is a global maximum rather than a local maximum in probability space.

One other issue with this framework is that one must assume Gaussian errors for each parameter. This means that a choice for the standard deviation of the Gaussian must be made alongside a mean value, and often it is not clear what this standard deviation should be. Also, the assumption of Gaussian distributions itself is not easily justified, and could be incorrect for any number of parameters.

3.2 Markov Chain Monte Carlo

3.2.1 Bayesian Inference

Bayesian inference is about the quantification and propagation of uncertainty, defined via a probability, in light of observations of the system. They arrive at a posterior probability distribution from an *a priori* state vector.

Fundamentally, Bayesian inference is an application of Bayes’ theorem:

$$p(\mathbf{x}|\mathbf{y}) = \frac{\mathcal{L}(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{y})} \quad (3.6)$$

where $\mathbf{x}$ is the parameter vector, $\mathbf{y}$ is the data vector (i.e. the spectrum), $p(\mathbf{x}|\mathbf{y})$ is the joint posterior probability distribution where $|$ means ‘event conditional on’ so that $(\mathbf{x}|\mathbf{y})$ means $\mathbf{x}$ given $\mathbf{y}$, $\mathcal{L}(\mathbf{y}|\mathbf{x})$ is the likelihood distribution which penalises poor/improbable fits to the data, $p(\mathbf{x})$ is the prior information which can restrict the parameter space, and $p(\mathbf{y})$ is the evidence, or marginal likelihood, which is a normalisation factor required for Bayesian model comparison. The prior information is usually uniform unless stated otherwise. For example, a uniform prior on mass would be between $1\text{--}80M_{Jup}$, meaning that if the mass is within this range then the prior has probability of 1. If it outside this range, then the prior, and hence posterior, has a probability of 0.

3.2.2 Markov Chain Monte Carlo

The Markov Chain Monte Carlo method is a portmanteau of two key terms: Markov Chain, and Monte Carlo. Markov Chain Monte Carlo uses both of these concepts to conduct its Bayesian inference.

The Monte Carlo approach is a method for obtaining numerical results by repeated random sampling of the parameter space. In our case, the Monte Carlo method means randomly generating spectra from state vectors containing the atmospheric parameters of the brown dwarf or exoplanet to be modelled so that they may be compared to a measured spectrum. These atmospheric parameters include temperature structure, volume mixing ratios (VMRs), mass, radius, and so forth.

A Markov Chain is a random process that undergoes transitions from one state to another in state space. It must possess a property that is usually characterised as ‘memorylessness’: the probability distribution of the next state depends only on the current state and not on the sequence of events that preceded it. The Markov Chain can be thought of as a transition matrix or distribution, i.e. how likely one state is to transition into another. A simple example would be if the Markov Chain in a weather model stipulated that tomorrow will have a 90% chance of having the same weather as today. If today it is raining, there is a 90% chance it will rain tomorrow, or a 10% chance it will not rain. If tomorrow it rains, then the day after that will have a 90% chance of raining, and a 10% chance of not. Thus the transition from one state to the next only depends on the current state.

The Markov Chain Monte Carlo method seeks to combine the random sampling of parameter space (Monte Carlo) and the properties of guiding given by a transition matrix (Markov Chain). Therefore, MCMC techniques aim to construct clever Markov Chains which draw samples which are progressively more likely realisations of the distribution of interest - the posterior probability distribution function (PDF). How exactly this transition is constructed depends on the MCMC technique employed.

The joint posterior probability distribution function (PDF), $p(\mathbf{x}|\mathbf{y})$, is the PDF for the joint fitting of the state vectors along with the prior information, i.e., the ‘best fit’ set of parameters would come from such a distribution. However, not only do we want the ‘best’ set of parameters, we also would like to know the marginalised posterior PDF of each parameter - that is, the probability distribution for a given parameter.

Marginalisation is represented by the marginalised probability function $p(x|\mathbf{y})$, where x is a parameter from the vector $\mathbf{x}$, given the set of observations $\mathbf{y}$:

$$p(x|\mathbf{y}) = \int p(x, \boldsymbol{\alpha}|\mathbf{y}) d\boldsymbol{\alpha} \quad (3.7)$$

where $\boldsymbol{\alpha}$ is the set of nuisance parameters, i.e. all of the parameters excluding the parameter of interest.

Usually, there are many parameters in the set $\boldsymbol{\alpha}$, and is therefore expensive to integrate. However, by sampling the joint posterior distribution $p(x, \boldsymbol{\alpha}|\mathbf{y})$ with the values $(x_t, \boldsymbol{\alpha}_t)$ (where t represents the timestep/iteration, and hence number of samples), we automatically produce a sampling of values x_t from the marginalised PDF $p(x|\mathbf{y})$. In practice, the marginalised posterior is represented by the number of samples. In plain English, for atmospheric retrievals, we produce guesses for each parameter in the atmosphere $(x_t, \boldsymbol{\alpha}_t)$, and repeatedly calculate the joint posterior PDF $p(x, \boldsymbol{\alpha}|\mathbf{y})$ from the likelihood function $\mathcal{L}$ (which compares the model spectrum to the data) and the prior (usually uniform, i.e. 0 or 1). The MCMC algorithm guides the parameter guesses towards those that produce the highest likelihood (or lowest log-likelihood).

MCMC is a perfect match for planetary science, where the models often contain many free parameters and have low signal-to-noise observations. The main advantage of using MCMC over optimal estimation is its ability to quantify all aspects of uncertainty via

probability and propagate these errors robustly in the posterior probability distribution function. This is especially desirable for low resolution or wavelength-limited spectra, which is usually the case for Y dwarfs and exoplanets.

MCMC has already been used to great effect by Line et al. (2015) to calculate the probability distribution functions of many parameters relevant to brown dwarf atmospheres, concluding that only low-resolution spectra are required to unambiguously detect ammonia in T dwarfs. These methods were also performed on a survey of 11 T dwarfs, in a similar manner that we intend to use in this thesis, to detect atmospheric trends. Line et al. (2016a). They found little correlation with temperature for all gases except decreasing Na and K, which is indicative of condensation. No clouds were found in this survey. A similar analysis on L dwarfs was performed by Burningham et al. (2017), finding optically thick clouds. We note here that, in comparison to optimal estimation, MCMC can be orders of magnitude more computationally expensive and so models must be selected carefully before executing.

One thing to note about MCMC is the fact that the likelihood (or prior) function of the parameters may come from expensive calculations. Here, efficient MCMC sampling, where few function evaluations are required to create a statistically independent sample from the posterior PDF, is very valuable. The methods presented here (Section 3.2) are designed for efficiency.

One final note to make, is that the posterior distributions are usually unnormalised distributions in an MCMC. This is because MCMC techniques generally do not compute the Bayesian evidence as it is computationally expensive and unnecessary for parameter estimation. Hence Equation 3.6 in our implementation is often reduced to:

$$p(\mathbf{x}|\mathbf{y}) = \mathcal{L}(\mathbf{y}|\mathbf{x})p(\mathbf{x}) \quad (3.8)$$

where $p(\mathbf{x}|\mathbf{y})$ is now the unnormalised posterior PDF.

3.2.3 EMCEE

We use the Markov Chain Monte Carlo (MCMC) approach implemented with the affine-invariant ensemble sampler EMCEE (Foreman-Mackey et al., 2013), following Line et al.




Figure 3.2.1: Cartoon illustrating the different kinds of affine transformations on a square. Figure from: http://resources.esri.com/help/9.3/arcgisengine/java/gp_toolref/coverage_tools/how_transform_coverage_works.htm, accessed 13th March 2018.

(2015). Affine-invariant transformations preserve collinearity, so they transform parallel lines into parallel lines and preserve ratios of distances along parallel lines. Affine transformations allow the conversion of highly anisotropic probability distribution functions into isotropic distribution functions, which greatly reduces the difficulty of sampling the PDF. In plain English, the affine transformation is applied to recast probability distribution shapes that are difficult to sample into easier ones. The four kinds of affine transformation are shown in Figure 3.2.1.

This is a significant advancement over the optimal estimation used previously in NEMESIS, as we are now able to make fewer *a priori* assumptions, and the shape of the parameter uncertainties no longer need to be Gaussian.

Algorithm

The EMCEE algorithm involves simultaneously evolving an ensemble of K *walkers* in the set $S = \{X_k\}$ with positions X where the proposal distribution for one walker k is based on the current positions of the $K - 1$ walkers in the *complementary ensemble* in the set

$S_{[k]} = \{X_j, \forall j \neq k\}$, i.e. for all ($\forall$) walkers in the set ($\forall j$) that are not the walker of interest (k). Here, position refers to a vector in the N-dimensional, real-valued parameter space, i.e. the state vector containing the atmospheric parameters.

The recommended method of calculating the initial set of walkers is a tight N-dimensional Gaussian ball around a set of ‘best guess’ parameters. These initial positions then sample the parameter space, and calculate the associated posterior distribution. In EMCEE, the posterior calculations are calculated in log-space, so that the unnormalised posterior distribution is now the unnormalised log-posterior. Hence equation 3.8 becomes:

$$\ln p(\mathbf{x}|\mathbf{y}) = \ln \mathcal{L}(\mathbf{y}|\mathbf{x}) + \ln p(\mathbf{x}) \quad (3.9)$$

To update the position of a walker at position X_k , a walker X_j is drawn randomly from the remaining walkers $S_{[k]}$ and a new position Y is proposed:

$$X_{k(t)} \rightarrow Y = X_j + Z[X_{k(t)} - X_j] \quad (3.10)$$

where Z is a random variable drawn from a distribution $g(Z = z)$:

$$g(z^{-1}) \propto \begin{cases} \frac{1}{\sqrt{z}} & \text{if } z \in [\frac{1}{a}, a] \\ 0 & \text{otherwise} \end{cases}$$

and a is an adjustable scale parameter, normally set to 2. This distribution is chosen because g then satisfies

$$g(z^{-1}) = zg(z) \quad (3.11)$$

which allows the proposal of Equation 3.10 to be symmetric in the sense that

$$p(X_k(t) \rightarrow Y) = p(Y \rightarrow X_k(t)). \quad (3.12)$$

In this case, the chain will satisfy detailed balance (around any closed cycle of states, there is no net flow of probability, i.e. Equation 3.12) if the proposal is accepted with probability

$$q = \min \left(1, Z^{N-1} \frac{p(Y)}{p(X_k(t))} \right) \quad (3.13)$$

where $p(Y)$ is the probability of the proposed position Y , calculated from the likelihood function and the prior information, and N is the dimension of the parameter space. This procedure is then repeated in series for each walker. Once this procedure is completed, this is classified as one iteration. The MCMC then iterates until convergence is achieved, sampling the parameter space and hence calculating the posterior distributions along the way.

A huge benefit of using EMCEE, is that it may perform these iterations (with some modification to the algorithm) in parallel, and therefore the sampling may occur on multiple cores simultaneously, drastically reducing the amount of (real, not computational) time taken for each analysis. It is also configured for multicore-multinodal calculations using an MPI (message passing interface) setup, so that the calculations can be further spread between many computers and many cores at once.

3.3 Nested Sampling

The use of Bayesian inference is usually divided into two main categories: parameter estimation and model selection. MCMC sampling is typically used for parameter estimation, however it can be very computationally intensive and often MCMC methods experience problems in sampling efficiently from a multimodal posterior distribution with large curving degeneracies between parameters, especially in highly dimensional problems. For example, for parameters that are mostly degenerate, such as the radius and distance of a brown dwarf, it could be that the MCMC sampler gets stuck in a locally optimal mode (high probability) that explains the data using a small distance, while in fact it did not efficiently explore the parameter space to discover the globally optimal mode (highest probability) of a larger radius.

Testing for convergence can also be a difficult task. For Bayesian model selection, the Bayesian evidence (also called the marginalised likelihood or marginal density of the data, i.e. $p(\mathbf{y}) \equiv \mathcal{Z}$) must be calculated to sufficient precision, which adds an even greater computational expense to the MCMC method. As the average likelihood of a model over

its prior probability space, the evidence can be used to assign relative probabilities to different models. This is because model selection is usually done by calculating the Bayes factor. The Bayes factor is the ratio of the posterior probabilities given the observed data set. Assuming that the prior distributions are the same, then this is given by the ratio of the Bayesian evidences, $\mathcal{Z}_0/\mathcal{Z}_1$. The existing preferred evidence evaluation method, again based on MCMC techniques, is thermodynamic integration (see e.g. Ruanaidh & Fitzgerald (1996)), which is extremely computationally intensive. Briefly, thermodynamic integration involves running several separate MCMCs and uses their average log-likelihood in a quadrature scheme to estimate the evidence.

Nested Sampling (NS) (Skilling, 2004) is a Monte Carlo method which is primarily constructed to efficiently calculate the Bayesian evidence, but also produces posterior distribution inferences as a by-product. Mukherjee et al. (2006) showed that their implementation of the method requires a factor of ~ 100 fewer posterior evaluations for their cosmological modelling efforts than thermodynamic integration.

Nested Sampling exploits the relation between the likelihood and prior volume to transform the multidimensional evidence integral (where N is the number of dimensions):

$$\mathcal{Z} = \int \mathcal{L}(\mathbf{y}|\mathbf{x})p(\mathbf{x})d^N\mathbf{x} \quad (3.14)$$

into a one-dimensional integral. The prior volume X is defined by $dX = p(\mathbf{x})d^N\mathbf{x}$ so that

$$X(\lambda) = \int_{\mathcal{L}(\mathbf{y}|\mathbf{x}) > \lambda} p(\mathbf{x})d^N\mathbf{x} \quad (3.15)$$

where the integral extends over the region(s) of parameter space contained within the iso-likelihood contour $\mathcal{L}(\mathbf{y}|\mathbf{x}) = \lambda$. Equation 3.14 then becomes

$$\mathcal{Z} = \int_0^1 \mathcal{L}(X)dX \quad (3.16)$$

where $\mathcal{L}(X)$, the inverse of Equation 3.15, is a monotonically decreasing function of X . Thus, if one can evaluate the likelihoods $\mathcal{L}_i = \mathcal{L}(X_i)$ where X_i is a sequence of decreasing values,




Figure 3.3.1: Cartoon illustrating (a) the posterior of a two dimensional problem; and (b) the transformed $\mathcal{L}(X)$ function where the prior volumes X_i are associated with each likelihood $\mathcal{L}_i$. Figure from Feroz et al. (2009).

$$0 < X_M < \dots < X_2 < X_1 < X_0 = 1 \quad (3.17)$$

as shown in the illustration in Figure 3.3.1, the evidence can be approximated numerically using standard quadrature methods as a weighted sum

$$\mathcal{Z} = \sum_{i=1}^M \mathcal{L}_i w_i. \quad (3.18)$$

The trapezium rule is used for quadrature, $w_i = \frac{1}{2}(X_{i-1} - X_{i+1})$.

Equation 3.18 is performed as follows. The iteration count i is set to zero, and N ‘active’ (or ‘live’) points samples are drawn from the prior $p(\mathbf{x})$ (usually a uniform distribution over the prior range), so the initial prior volume is $X_0 = 1$. The samples are then sorted in order of their likelihood and the smallest (likelihood $\mathcal{L}_0$) is removed from the active set (hence being ‘inactive’) and replaced by a point drawn from the prior subject to the constraint that the point has a likelihood $\mathcal{L} > \mathcal{L}_0$. The corresponding prior volume contained within this iso-likelihood contour will be a random variable given by $X_1 = t_1 X_0$ where t_1 follows the distribution $\Pr(t) = Nt^{N-1}$ (i.e. the probability distri-

bution for the largest of N samples drawn uniformly from the interval $[0, 1]$). At each subsequent iteration i , the removal of the lowest likelihood point $\mathcal{L}_i$ in the active set, the drawing of a replacement with $\mathcal{L} > \mathcal{L}_i$ and the reduction of the corresponding prior volume $X_i = t_i X_{i-1}$ are repeated, until the entire prior volume has been traversed. The algorithm thus travels through nested shells of likelihood as the prior volume is reduced. The mean and standard deviation of $\log t$, which dominates the geometrical exploration, are $E[\log t_i] = -i/N$ and $\sigma[\log t_i] = i/N$. Since each value of $\log t$ is independent, after i iterations the prior volume will shrink down such that $\log X_i \approx -(i \pm \sqrt{i})/N$. Hence, taking the exponential and assuming large i , this gives $X_i = \exp(-i/N)$.

The algorithm is terminated on determining the evidence to some specified precision (usually 0.5 in log-evidence): at iteration i , the largest evidence contribution that can be made by the remaining portion of the posterior is $\Delta \mathcal{Z}_i = \mathcal{L}_{max} X_i$ where $\mathcal{L}_{max}$ is the maximum likelihood in the current set of active points. The evidence estimate (Equation 3.18), may then be refined by adding a final increment from the set of N active points, which is given by

$$\Delta \mathcal{Z} = \sum_{i=j}^N \mathcal{L}_j w_{M+j}. \quad (3.19)$$

where $w_{M+j} = X_M/N$ for all j . The final uncertainty on the calculated evidence may be estimated from a single run of the nested sampling algorithm by calculating the relative entropy of the full sequence of samples (see Feroz & Hobson (2008)). The relative entropy (also called the Kullback-Leibler divergence) is a measure of how one probability distribution (in this case, the posterior) diverges from a second, expected probability distribution (the prior volume), and is simply the log of the ratio of the two distributions.

Once the evidence $\mathcal{Z}$ is found, posterior inferences can be easily generated using the full sequence of (inactive and active) points generated in the nested sampling process. Each such point is assigned the weight

$$p_j = \frac{\mathcal{L}_j w_j}{\mathcal{Z}} \quad (3.20)$$

where the sample index j runs from 1 to $\mathcal{N} = M + N$, the total number of sampled




Figure 3.3.2: Cartoon of ellipsoidal nested sampling from a simple bimodal distribution. In (a) we see that the ellipsoid represents a good bound to the active region. In (b)-(d), as we nest inward we can see that the acceptance rate will rapidly decrease as the bound steadily worsens. Figure (e) illustrates the increased efficiency obtained by sampling from each clustered region separately. Figure from Feroz et al. (2009).

points.

Now that we have created our samples as a by-product of calculating the evidence, we can now use these samples to calculate inferences of posterior parameters such as means, standard deviations, covariances, and so on, or to construct marginalised posterior distributions, in the same fashion as we have done for the EMCEE case.

3.3.1 MultiNest

The MultiNest algorithm is a simultaneous ellipsoidal nested sampling method. Here we outline briefly the algorithm.

Firstly the prior volume is taken as a D -dimensional hypercube (each parameter value varies from 0 to 1), in which samples are drawn uniformly. A hypercube is cube that has been extended to dimensions D which are larger than three. This means that physical parameter space must be converted to unit hypercube space for the calculation. The minimum volume for the bounding ellipsoid (E , i.e. the ‘skin’ of the ellipsoid, given in orange in Figure 3.3.3, which encompasses the minimum volume required to fit all of the live points) and the volume of this ellipsoid $V(E)$ over the entire unit hypercube prior is calculated. The bounding ellipsoid is then enlarged by a factor f_k such that its volume $V(E)$ is greater than the prior volume, $X_i \equiv V(S) = \exp(-i/N)$, of the live points in the unit hypercube, S . The live points S are then partitioned into subsets S_1 and S_2 , containing n_1 and n_2 live points respectively, by applying a k -means clustering algorithm




Figure 3.3.3: Examples of the bounding ellipsoids for 1000 live points. Top: a bi-modal distribution from two different samples. Bottom: a torus from the same sample. Figure from Feroz et al. (2009).

with $K = 2$. The k -means clustering algorithm is an iterative procedure, where, for this case, $K = 2$ centroids are chosen randomly in the prior space, and all points are assigned to the centroid they are closest to. The mean of these points belonging to each centroid is then calculated, and used as the new centroid positions. This iterates until a specified minimum amount of change occurs. The bounding ellipsoid (enlarged again by f_k) and its associated volume is then calculated for each subset. If the volume of the bounding ellipsoid $V(E)$ is greater than the sum of the volumes of the bounding ellipsoids for each subset (i.e. $V(E) > V(E_1) + V(E_2)$), or $V(E)$ is greater than double the volume of the prior volume $V(S)$ (i.e. $V(E) > 2V(S)$), the algorithm repeats on each subset S_1 and S_2 for further partitioning into four subsets, and eight, and so on; otherwise $V(E)$ is returned as the optimal ellipsoid to use to sample the prior space. This iterative ellipsoidal decomposition scheme and prior sampling is illustrated in Figure 3.3.2. Essentially, this shows how partitioning into multiple ellipsoids occurs when the sampling efficiency becomes below a critical value. This describes how the sampling occurs over time; the posteriors are calculated in the same way as described as in Section 3.3, i.e. evaluating the Bayesian evidence.

In summary, this algorithm can be multiple orders of magnitude faster at calculating the Bayesian evidence than thermodynamic integration methods (MCMC). This is highly useful for model comparison and selection. For posterior distribution calculations, the number of likelihood evaluations (i.e. how many times the forward model is called) can be an order of magnitude smaller than MCMC, can easily sample multimodal distributions, and shows clear signs of convergence. One advantage of EMCEE over MultiNest, is that it can be more agnostic about the shape of the posteriors (any shape of posterior is equally well sampled), while MultiNest prefers Gaussian posteriors due to the fact that we are using ellipsoids as the primary sampling method.

Chapter 4

Gaseous Absorption in Brown Dwarfs and Exoplanets

4.1 Introduction

The Solar System contains a wonderful variety of planetary atmospheres. Each planet has its own unique formation and evolutionary history resulting in complex atmospheric compositions. In order to better understand the physical processes taking place within these celestial bodies, we extract data across almost the entire electromagnetic spectrum, primarily in the visible and infrared as these are the wavelengths where light is best reflected and thermally emitted. Radiative-transfer models allow us to simulate the radiative processes on these bodies, producing spectra that can be compared to our measured data sets. Usually these radiative-transfer models are coupled to an inverse method, which iteratively explores parameter space to retrieve the most likely set of parameters describing the atmosphere of the (exo)planet or brown dwarf (e.g., Irwin et al. (2008); Waldmann et al. (2015); Line et al. (2015)).

This method of determining the atmospheric structure and composition is reliant on the accuracy of our radiative transfer models. For Bayesian retrieval/posterior sampling algorithms, such as Nested Sampling (e.g. MultiNest, (Feroz et al., 2013)) and Markov Chain Monte Carlo (MCMC, e.g. EMCEE (Foreman-Mackey et al., 2013)), the radiative-transfer forward model will be executed millions of times while it explores a large parameter space. Because of the nature of these highly dimensional retrieval calcu-

lations, small differences in forward models could potentially produce non-trivial changes in the retrieved probability distributions. For example, if one model produces an extra, erroneous absorption band for, say, NH_3 , the retrieval would try to compensate for this error during the fitting procedure. While the ‘true’ value of NH_3 should fit the spectrum (assuming the rest of the model is perfect), the extra absorption band would mean that the total amount of retrieved NH_3 will be smaller than the ‘true’ value, as it finds a medium-ground answer which best fits all (true + erroneous) absorption bands. If we do not assume the rest of the model is perfect, then the problem is even graver - the retrieval will try to fit the spectrum by adjusting the model in a non-trivial way, such as adjusting the surface gravity or the temperature profile. In this way, one forward model error propagates through to all other retrieved posterior probability distribution functions (PDFs).

A key difference between a number of radiative transfer models is how the gaseous absorption is calculated - this is because it is a computationally expensive step. The detail of these calculations is covered much more extensively in Chapter 2, but we summarise the methods here for clarity.

A line-by-line calculation, where the absorption coefficient is calculated for the local temperature and pressure for each individual line, is the most accurate and correct as it resolves each individual line exactly. However, this method is far too slow to be used in sophisticated retrieval procedures because of the enormous number of lines involved, especially at high temperatures ($>1000\text{K}$). Instead, the absorption coefficient is usually pre-calculated at a range of temperatures and pressures for a variety of gases, and interpolated to the necessary values.

There are multiple methods of calculating the gaseous absorption, but in this Chapter we limit ourselves to discussion of just two methods: individual gas cross-sections and correlated- k (e.g., Lacis & Oinas, 1991). We also briefly consider the opacity sampling method (see e.g. Hubeny & Mihalas, 2014), though not as extensively as the previous two as implementing this would require a major overhaul of our code. We note that, as long as the pressure and temperature resolution in the cross-section look-up table is high enough, there is no distinction between a high-resolution ($\Delta\nu = 0.01 \text{ cm}^{-1}$) cross-section and a line-by-line calculation. Individual gas cross-sections have been used to solve the

radiative-transfer equation (MacDonald & Madhusudhan, 2017). The cross-sections are absorption coefficients calculated from a line-by-line calculation on a grid of pressures and temperatures for a given gas, then integrated to a (typical) resolution of $\Delta\nu = 1\text{cm}^{-1}$. Other authors calculate a high-resolution cross-section and sample it at a resolution of $\Delta\nu = 1\text{cm}^{-1}$ (Line et al., 2015; Sharp & Burrows, 2007). Hedges & Madhusudhan (2016) found that these individual gas cross-sections could have median opacity differences of $<1\%$ for low-resolution ($R \sim 100$), up to 40% for medium-resolution ($R \lesssim 5000$), and over 100% to 1000% for high-resolution cross-sections ($R \sim 1 \times 10^5$), introduced by various aspects of pressure broadening. These errors found in the individual cross-sections are before the non-trivial task of quantifying these differences over a whole atmosphere.

Another approach to calculating the spectra of brown dwarfs and exoplanets is using premixed k -coefficients (Saumon & Marley, 2008) or on-the-fly mixing for k -coefficients (Irwin et al., 2008; Barstow et al., 2014; Lee et al., 2014). Note that for the rest of the chapter, ‘cross-sections’ refers to cross-sections which are not premixed, i.e. individual cross-sections for individual gases. k -tables are produced by performing line-by-line calculations of the absorption coefficient then rewriting the absorption coefficient strength distribution in terms of a cumulative frequency distribution over bins of specified wavenumber/wavelength width by ranking and sampling the distribution according to absorption coefficient strength. The inverse of this distribution is known as the k -distribution (Lacis & Oinas, 1991). The k -distribution is a smooth, monotonically increasing function and so can be sampled with only 10-20 points, compared with $\sim 10^3$ to 10^6 for the cross-section or line-by-line methods. Within a single atmospheric layer, we may simply combine the k -distributions of different gases by assuming that the lines are randomly overlapping. Note that there are many ways of combining k -coefficients for different gases. Among brown dwarf and gas giant atmosphere models, ATMO (Tremblin et al., 2015; Drummond et al., 2016) and PETIT (Mollière et al., 2015) also use random overlap, HELIOS (Malik et al., 2017) assumes perfect correlation, and Amundsen et al. (2016) use equivalent extinction. Equivalent extinction computes k -coefficients for each gas and combine them using an equivalent grey absorption for all minor absorbers and k -coefficients for the major absorber in each band. Perfect correlation assumes that the shape of the k -distributions are the same, but can differ in magnitude. This means only one quadrature weight needs

to be used when combining different gases. Then, by assuming that the wavenumbers at which the (total) cross section takes a certain value are vertically correlated these k -distributions may be used to calculate the transmission, thermal emission or scattering of an atmosphere using the correlated- k method. From previous studies (Irwin et al., 2008), we have found the correlated- k approximation to be accurate to better than 5%. This is why we primarily use this as our benchmark absorption method for the chapter. For more information on our correlated- k method, refer to Chapter 2 and references therein.

In this Chapter, we argue that it is inaccurate to calculate the gaseous absorption in both single- and mixed-gas atmospheres by combining individual gas cross-sections with large- and moderately-sized bins ($\Delta\nu = 25, 1 \text{ cm}^{-1}$), or inefficient with high-resolution bins ($\Delta\nu = 0.01 \text{ cm}^{-1}$). Throughout this chapter we compare our calculations to our (random overlap) correlated- k method, and verify our methods with a line-by-line calculation. We also consider the effects of including H₂-He pressure-broadening in the spectral calculations.

4.2 Line Lists

A wide range of line list databases exists to provide the relevant molecular information for calculating spectra in the atmospheres of exoplanets and brown dwarfs. The key factors in deciding which line lists are most appropriate are the wavelength and temperature ranges for which they are valid. For example, a widely used line list database is the High Resolution Transmission (HITRAN) database (Rothman et al., 2013), which is collated from multiple experimental and theoretical sources. This database is, however, used mainly for representing the spectrum of the Earth, and is therefore only reliable for temperatures up to $\sim 400\text{K}$, as it removes all insignificant line intensities at this temperature. It also only contains gases that are relevant to the Earth’s atmosphere. Unfortunately, these insignificant line intensities become more significant with increasing temperature, meaning the relevant number of lines goes from thousands to billions as the temperature increases above 1500K .

The ExoMol project (Tennyson & Yurchenko, 2012) contains a much more valid temperature (up to $1500\text{-}3000\text{K}$) and wavelength range for the relevant gases important to

brown dwarf and exoplanet atmospheres. However, ExoMol line data are predicted solely from *ab initio* calculations rather than measured in a laboratory and as such line positions and intensities may contain larger errors than those found experimentally in, e.g. HITRAN, for overlapping validity ranges.

For exoplanets, NEMESIS previously used a range of k -tables that assumed Earth-like pressure broadening. The line lists themselves were either unsuitable for high temperatures, or are now outdated (Lee, 2012) due to newer datasets becoming available. Therefore we created new k -tables for our calculations. We describe the new k -tables here.

In Table 4.2.1, we present a summary of the line lists that we selected according to the preceding criteria, where T_{max} is the maximum temperature to which the line list is accurate to, N is the number of lines, J_{low} is the lower state of quantum total angular momentum of the transition and $Q(T)$ is the partition function. We also use H_2 - H_2 and H_2 -He collision-induced absorption from HITRAN 2012 (Richard et al., 2012) for any calculation containing H_2 -He.

In order to use the ExoMol data, which can take up to 1TB of storage for CH_4 , a FORTRAN 77 code was written to reformat the data into a metaformat in the style of HITRAN with a two-thirds reduction in file size compared to a full HITRAN format. In this process, we calculated the line intensity at 296K, given by equation 2.3, and also inserted the relevant pressure broadening parameters, which are not given by ExoMol in their line list, which are discussed in more detail in section 4.3.

4.3 Pressure broadening coefficients

Another measure of appropriateness for each line list is the validity of its pressure-broadening coefficients. Unfortunately, line list databases which are suitable for the pressure broadening found in H_2 -He-dominated atmospheres are scarce. HITRAN’s Earth-centric lists exhibit broadening parameters suited to air-broadening (N_2 and O_2). ExoMol provides line lists where only a few molecules contain H_2 and He pressure broadening parameters (H_2O , HCN , CO). For all molecules ExoMol provides cross-sections with only Doppler (i.e. thermal) broadening, but no pressure broadening.

Table 4.2.1: Summary of Line Lists Selected

Molecule	Reference	Method ^a	Spectral Coverage (cm ⁻¹)	T_{max} (K)	N	Max. J_{low}	$Q(T)$	Available ^b
CO	Rothman et al. (2010)	expt	0-8,465	4,000	115,218	150	Rothman et al. (2010)	HITEMP
CO ₂	Tashkun & Perevalov (2011)	ai	226-8,310	4,000	628,324,454	300	Rothman et al. (2013)	CDSD
H ₂ O	Barber et al. (2006)	ai	0-30,000	3,000	505,806,202	50	Barber et al. (2006)	ExoMol
NH ₃	Yurchenko et al. (2011)	ai	0-12,000	1,500	1,138,323,351	36	Yurchenko et al. (2011)	ExoMol
CH ₄	Yurchenko & Tennyson (2014)	ai	0-12,000	1,500	9,819,604,957	39	Yurchenko & Tennyson (2014)	ExoMol
TiO	R.S. Freedman (priv. com.)	ai	0-20,000	?	125,676,312	150	Sauval & Tatum (1984)	N/A
VO	R.S. Freedman (priv. com.)	ai	3,850-25,950	?	3,153,073	175	Sauval & Tatum (1984)	N/A
Na	Kupka et al. (2000)	?	5,048-32,952	?	158	?	Sauval & Tatum (1984)	VALD
K	Kupka et al. (2000)	?	5,048-30,281	?	275	?	Sauval & Tatum (1984)	VALD

^aMethod:ai = *ab initio*

expt = experimental

^bData sources:HITEMP: <http://www.cfa.harvard.edu/hitran/HITEMP.html>CDSD: <ftp://ftp.iao.ru/pub/CDSD-4000>ExoMol: www.exomol.comVALD: <http://vald.astro.uu.se>

Molecule	Broadener	Line width reference	Temp. exponent, n^a	Temp. exponent reference
H ₂ O	H ₂	Gamache et al. (1996)	0.44	Gamache et al. (1996)
	He	Solodov & Starikov (2009); Steyvert et al. (2004)	0.44	Gamache et al. (1996)
CH ₄	H ₂	Pine (1992); Margolis (1993)	0.44	Margolis (1993)
	He	Pine (1992)	0.28	Varanasi & Chudamani (1990)
CO ₂	H ₂	Padmanabhan et al. (2014)	0.60	Sharp & Burrows (2007)
	He	Thibault et al. (1992)	0.60	Thibault et al. (2000)
CO	H ₂	Regalia-Jarlot et al. (2005)	0.60	Moal & Severin (1986)
	He	BelBruno et al. (1982); Mantz et al. (2005)	0.55	Mantz et al. (2005)
NH ₃	H ₂	Hadded et al. (2001); Pine et al. (1993)	0.64	Nouri et al. (2004)
	He	Hadded et al. (2001); Pine et al. (1993)	0.40	Sharp & Burrows (2007)
TiO	H ₂	Sharp & Burrows (2007)	0.60	Sharp & Burrows (2007)
	He	Sharp & Burrows (2007)	0.40	Sharp & Burrows (2007)
VO	H ₂	Sharp & Burrows (2007)	0.60	Sharp & Burrows (2007)
	He	Sharp & Burrows (2007)	0.40	Sharp & Burrows (2007)

Table 4.3.1: Summary of Pressure Broadening Coefficient Sources

^a Average value for all transitions due to limited data.

In order to correctly estimate the pressure broadening induced on spectral lines in brown dwarf and exoplanet atmospheres, we performed a literature search and found that Amundsen et al. (2014) had openly discussed their sources for H₂-He pressure broadening. Most of the information found in Table 4.3.1 is from Amundsen et al. (2014), with a few additional sources added. Note that we have not chosen to implement the same procedure for Na and K, as they are not experimentally well sampled for H₂ and He broadening; instead we arbitrarily set the air-broadened widths to 0.075 cm¹ atm⁻¹ based on experience with terrestrial radiative transfer studies. We also note that using Voigt lineshapes, as we have done for all gases, is especially dubious for the Na and K resonance lines (Burrows et al., 2001), because they both produce massive non-Lorentzian broadened wings in the optical. We intend on updating the Na and K *k*-tables in the future for more realistic lineshapes; however this will not drastically change the outcomes of this chapter. Methods taken by other groups tend to use or emulate *ab initio* Na and K lineshape calculations, e.g. Tremblin et al. (2015); Baudino et al. (2015).

In many cases these sources only provided broadening parameters for the lower rotational quantum number, J_{low} , for up to 8-20 for the maximum J_{low} value depending on the gas, while our line lists contain data up to $J_{low} = 300$. We implemented the broadening parameters into our line lists shown in Table 4.3.1 by first converting them into a single ‘foreign broadening parameter’ γ_0 assuming an atmospheric ratio of 85:15 for H₂:He, using the weighted sum $\gamma_0 = \gamma_{H_2} \text{VMR}_{H_2} + \gamma_{He} \text{VMR}_{He}$, where VMR represents the volume mixing ratio. We then fitted this foreign broadening coefficient with a fourth-order polynomial given by $\gamma_0(J_{low}) = \sum_{i=0}^{i=4} \alpha_i J_{low}^i$, where α_i represents each order’s constant, up until the highest J_{low} value covered by the data, then using the last available broadening coefficient for any J_{low} higher than the maximum available. Note that Amundsen et al. (2014) used a linear approach up to the maximum J_{low} , and then a constant value as we have. This is of course not ideal, but does not introduce any complex error propagation which might be the case with a more sophisticated modelling that extrapolates to higher J_{low} . An empirical approach was considered, but while many gases show a gradual flattening of broadening coefficient with increasing J_{low} (e.g., Buldyreva et al., 2011), there is no clear or simple relationship that allows us to model all of the gases after this maximum. In general, the constant value appears to be a good first-order




Figure 4.3.1: Pressure broadening coefficients for CO. Pink is air-broadening from HITRAN, blue circles are data points from relevant papers in Table 4.3.1, green lines are fourth-order polynomial fits to the data. Top panel: H_2 broadening coefficients. Middle panel: He broadening coefficients. Bottom: H_2 -He broadening coefficients with 85:15 ratio.

approximation. An example for the molecule CO is presented in Figure 4.3.1, where we compare our new H_2 -He foreign broadening to that provided for air by HITRAN.

4.4 Calculation of k -Tables

Parameter	Values for Wavelength (Wavenumber) Calculation
DELV	1 nm
WING	1.5cm^{-1}
FWHM	2 nm
CUTOFF	25cm^{-1} Na, K: 6000cm^{-1}
V_{REL}	25cm^{-1} Na, K: 6000cm^{-1}

Table 4.4.1: Summary of Spectral Parameters.




Figure 4.4.1: Diagram illustrating the parameters involved for calculating absorption, where V_i represents a wavelength bin i out of N bins with width **WING**. The considered central and adjacent bins are shaded in green. Figure from Lee (2012).

We now describe the exact details of the k -tables we use in our radiative transfer calculations.

For each gas in Table 4.2.1, we must calculate k -tables to be used in the correlated- k method of solving the radiative transfer equations. We developed a method to remove insignificant lines in the line databases at specified temperatures, as this will shorten the length of computation time of the k -tables. The method first calculates the line intensities at a particular temperature, and orders the line intensities from smallest to largest. It then creates a cumulative sum from the smallest to largest line intensities, and removes the smallest $n\%$ of contributions to the total line intensities. We use $n = 1 \times 10^{-10}$, a very small percentage, so that we do not underestimate any continuum effects. We do this for our entire temperature range, i.e. 100 - 2950K, over 20 equally spaced (150K) temperature points as different lines become important at different temperatures. At the lower temperatures, a majority of the lines are stripped away (leaving $\sim 1 \times 10^4$ lines), while at higher temperatures the line lists are essentially identical (up to $\sim 1 \times 10^{11}$).

For the k -tables, we use a long-established NEMESIS approach, and the only differences to any other k -table calculation are the numbers chosen for each parameter. We first calculate the underlying absorption spectrum monochromatically (at least $1/6$ Voigt line width) over the entire spectral region ($0.3\text{-}30\mu\text{m}$). We calculate the absorption and hence cumulative k -distributions at output wavelengths having separation $\Delta\lambda=0.001\mu\text{m}$, with square bins of width double that of the separation (i.e. a resolution of $\lambda=0.002\mu\text{m}$). We

use Gauss-Legendre quadrature to sample the distributions. To include the contribution of wings of lines centred outside the spectral region of interest, the total wavelength range considered is defined as a range between $\nu_{min} - \nu_{cut}$ to $\nu_{max} + \nu_{cut}$. The total spectral interval is then subdivided into 1.5cm^{-1} bins where line data is stored. Absorption at a particular output wavelength is then calculated by considering lines stored in the adjacent bins and the bin in the middle. Wing contributions from the lines centred outside these bins are calculated at the middle and end bins using a quadratic polynomial (because we assume the wing shape follows a Lorentzian), and added on to the absorption at the output wavelength. For each line calculation, a line wing outside the cutoff (25cm^{-1} from centre) is ignored.

We calculate a grid of spectral opacities with the 20 aforementioned temperature points, and 20 pressure levels equally spaced in logspace from $\sim 1 \times 10^{-7}$ to 100 atm, using 20 g-ordinates for Gauss-Legendre quadrature. The lineshapes are given by a Voigt profile, where the line-wing cut-off (ν_{cut}) is at 25cm^{-1} for all gases except the alkali metals Na and K, where the cut-off is at 6000cm^{-1} .


The pressure-temperature grid is linearly interpolated in temperature-log-pressure space to find the desired spectral opacity. The k -tables may also be resampled into lower resolutions to increase calculation speed, if necessary. These k -tables calculated here will be adopted in future works.

4.5 Pressure Broadening Effect on k -Tables


In this section we will examine the difference between air-broadening and H_2 -He-broadening for each gas, and comment on its effect on the opacity of each gas. The wavelengths selected were chosen to best illustrate the effect of each type of broadening for the given gas. Some comparisons involve comparing two different line lists as well as broadening types (i.e. NH_3 and CH_4), to investigate what differences this can make to each k -table.

4.5.1 CO


The CO H_2 -He broadening parameters were extracted from the references shown in Table 4.3.1, and then fit with a fourth-order polynomial as shown in Figure 4.3.1. The



(a) Temperature exponents for CO as a function of J_{low} for air broadening (HITRAN) and H_2 -He (see Table 4.3.1).



(b) CO absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1mbar. Note that the blue line is underneath the green line, but the differences are negligible.



(c) CO absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1bar.

foreign broadening used for the calculation of opacity is shown in the bottom plot of Figure 4.3.1. Note that the same J_{low} can have multiple values - this is because it depends also on K_{low} (lower vibrational quantum number). This is true for all results presented here. We can see that the pressure broadening coefficient has a higher value for most values of J_{low} for the H₂-He mixture than for air. Thus we would expect more pressure broadening to occur for H₂-He, especially for higher temperatures where a larger fraction of active lines will contain these higher values of J_{low} .


Another related element of pressure broadening is the temperature dependence, expressed in the form of a temperature exponent n as seen in equation 2.9. For CO, the overall foreign broadening temperature exponent for H₂-He is a constant 0.59 as we have taken an average of the few values over all J_{low} due to a lack of information, whereas for air it is variable with J_{low} and K_{low} (lower vibrational quantum number) as shown by Figure 4.5.1a. Overall we see that the temperature exponent is constantly lower for H₂-He than for air, and thus expect the amount of broadening for H₂-He to be larger for temperature ratios ($\frac{T_0}{T}$) < 1 in Equation 2.9, and vice versa for ratios > 1.

From approximately 0.2 bar, the differences in the pressure broadening parameters between air and H₂-He begin to become apparent (~5% change at 2950K). The differences are accentuated by increasing the pressure, as shown at 1bar in Figure 4.5.1c, where changes can be as large as 50%. In Figure 4.5.1b, we show that the two k -tables are identical at 1mbar.


Figure 4.5.1c shows two things. The first is that air broadening leads to a consistently higher computed opacity than H₂-He broadening. The second result is that the temperature dependence acts to increase this effect with increasing temperature (provided $T_0 < T$), as we can see when comparing the top and bottom plots in Figure 4.5.1c.

4.5.2 VO


Following the same method as CO, we present the pressure broadening coefficient as a function of J_{low} in Figure 4.5.2a. We use a temperature exponent for H₂-He of 0.57 and 0.75 for air. We find that VO exhibits relatively negligible differences between calculations with air or H₂-He, as shown in Figures 4.5.2b and 4.5.2c. This is unsurprising given the largely similar values of the pressure broadening coefficients.



(a) Pressure broadening coefficients for VO as a function of J_{low} for air broadening (HITRAN) and H_2 -He (see Table 4.3.1).



(b) VO absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1mbar.



(c) VO absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1bar.

Figure 4.5.2: Properties of VO k -table.

4.5.3 TiO

The pressure broadening coefficients as a function of J_{low} are presented in Figure 4.5.3a. The temperature exponent used for TiO was 0.75 for air and 0.570 for H₂-He. We find that TiO, like VO, exhibits relatively negligible differences between air and H₂-He broadening, as shown in Figure 4.5.3c and Figure 4.5.3b.


4.5.4 H₂O

The air and H₂-He pressure broadening coefficients are presented in Figure 4.5.4a. The temperature exponent used for H₂O was 0.5 for air and 0.45 for H₂-He. We find that at 1 bar, H₂O begins to exhibit significant (up to 40%) differences in its absorption spectra when comparing air and H₂-He broadening. This is due to the fact that at higher J_{low} , the pressure broadening coefficient of H₂-He is more than double that of air. One should note that this doubling is artificially instated due to a lack of data, and could be unrealistic.


4.5.5 NH₃

The pressure broadening coefficients presented in Figure 4.5.5a are very different for H₂-He and air. Note that the equation used to describe the pressure broadening coefficient also included another dimension, to represent the lower vibrational quantum number, K_{low} , as well, which results in spikes in the 2D diagram. The air-broadening data is not in the form of a polynomial equation, and here we use the air-broadening data directly for each transition. We also repeat this for CH₄. The temperature exponents are presented in Figure 4.5.5b. Here we are also comparing two separate linelists, HITRAN and ExoMol, with orders of magnitude difference in the number of lines they contain, as we are also trying to establish the changes from the last set of k -tables to the current. This was also done for practical reasons, as calculating the NH₃ (and also CH₄) k -tables take months of calculations, and would be too expensive computationally for this comparison.


We can see a complicated relationship between the pressure broadening and J_{low} in Figure 4.5.5a, with different lines experiencing different amounts of pressure broadening, likely due to its multidimensional format for NH₃. We also notice that, as expected, the ExoMol data contains much higher amounts of opacity at higher temperatures relative to



(a) Pressure broadening coefficients for TiO as a function of J_{low} for air broadening (HITRAN) and H_2 -He (see Table 4.3.1).




(b) TiO absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1mbar.




(c) TiO absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1bar.


Figure 4.5.3: Properties of TiO k -table.



(a) Pressure broadening coefficients for H₂O as a function of J_{low} for air broadening (HITRAN) and H₂-He (see Table 4.3.1).




(b) H₂O absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1mbar.




(c) H₂O absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1bar.


Figure 4.5.4: Properties of H₂O k -table.



(a) Pressure broadening coefficients for NH_3 as a function of J_{low} for air broadening (HITRAN) and $\text{H}_2\text{-He}$ (see Table 4.3.1).




(b) Temperature exponents for NH_3 as a function of J_{low} for air broadening (HITRAN) and $\text{H}_2\text{-He}$ (see Table 4.3.1).



(c) NH_3 absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1mbar. Air broadening is calculated using HITRAN and $\text{H}_2\text{-He}$ using ExoMol.

Figure 4.5.5: Properties of NH_3 k -table.



(d) NH_3 absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1bar. Air broadening is calculated using HITRAN and H_2 -He using ExoMol.

Figure 4.5.5: Properties of NH_3 k -table. (cont)

HITRAN.


4.5.6 CO_2

The difference between the estimated air and H_2 -He pressure broadening parameters of CO_2 is substantial, as shown in Figure 4.5.6a. The parameters for air are very complete as they have been extensively derived via high-temperature experiments. The temperature exponents, presented in Figure 4.5.6b, are a little closer in absolute terms, but still show significant differences, especially towards lower J_{low} .


At 1mbar, as shown in Figure 4.5.6c, the spectral differences are negligible. However, at 1bar, the spectral differences are very apparent, with a mild temperature dependence shown in Figures 4.5.6d. At high pressures the air-broadening acts to increase the overall opacity relative to H_2 -He broadening.

4.5.7 CH_4


Due to the overwhelming amount of data for air broadening shown in Figures 4.5.7a and 4.5.7b, we expect there to be a complicated lineshape departure highly dependent on any given line. Again we note that the comparison uses two line lists, where we are comparing



(a) Pressure broadening coefficients for CO₂ as a function of J_{low} for air broadening (HITRAN) and H₂-He (see Table 4.3.1).




(b) Temperature exponents for CO₂ as a function of J_{low} for air broadening (HITRAN) and H₂-He (see Table 4.3.1).



(c) CO₂ absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1mbar.

Figure 4.5.6: Properties of CO₂ k -table.



(d) CO₂ absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1bar


Figure 4.5.6: Properties of CO₂ k -table. (cont)

the ExoMol line list with the STDS line list (Wenger & Champion, 1998) to examine the overall effects of the new database. Note that this was also a practical decision, due to computational shortfallings, as stated for NH₃.


We find that, as shown in Figures 4.5.7c and 4.5.7b, each individual line exhibits an individual relationship to pressure broadening. These effects are apparent at both pressures. It appears that some lines are also slightly shifted from each line list, and we are sure this is due to the line lists themselves and not due to any physical effects caused in our calculations. Finally, we see in all cases there is a higher opacity for the ExoMol line list compared to the STDS line list for higher temperatures.

4.5.8 Summary of H₂-He Pressure-broadening effect on k -tables


All gases with the exceptions of TiO and VO show non-negligible opacity departures from air- to H₂-He broadening. It is uncertain how much of this is attributable to the lack of information available for H₂-He- broadening and the assumptions made here for higher J_{low} , but what is certain from this limited data set (i.e. low J_{low} ranges) is that using air-broadening is not an acceptable substitute for representing the lineshapes and opacity in brown dwarf atmospheres.



(a) Pressure broadening coefficients for CH_4 as a function of J_{low} for air broadening (HITRAN) and H_2 -He (see Table 4.3.1).




(b) Temperature exponents for CH_4 as a function of J_{low} for air broadening (HITRAN) and H_2 -He (see Table 4.3.1).



(c) CH_4 absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1mbar. Air broadening is calculated using STDS and H_2 -He using ExoMol.

Figure 4.5.7: Properties of CH_4 k -table.



(d) CH_4 absorption in the 10th g-ordinate in k -table at 2800, 1600, and 400K at 1bar. Air broadening is calculated using STDS and H_2 -He using ExoMol.

Figure 4.5.7: Properties of CH_4 k -table. (cont)

We have also illustrated the effects of using different line lists. When modelling high temperatures, those with lower-intensity cutoffs at room temperature (and hence having a larger number of lines) will reproduce the opacity much more accurately.

Another issue found is that the line centres shift very slightly for different line lists. This is only likely to be a problem with high resolution observations.

4.6 Cross-section Comparison

In this section we investigate two methods of solving the radiative transfer equation regarding optical depth, and show how they affect the spectra produced by forward models of brown dwarfs and exoplanets. As our retrievals are entirely dependent on our forward models being ‘correct’, this is an important point to address.

For the cross-section calculations in the literature, ExoMol (Hill et al., 2013) calculate a high-resolution ($<0.01\text{cm}^{-1}$) cross-section and the integral of the cross-section is preserved for all requested resolutions. MacDonald & Madhusudhan (2017) calculate their cross-sections at 0.01cm^{-1} (from Hedges & Madhusudhan (2016)), and bin them down to 1cm^{-1}

resolution before using them. Line et al. (2016b) use a 1cm^{-1} resolution opacity sampling method on pre-computed cross sections which have a variable resolution wavenumber grid that samples the lines at 1/4 of their Voigt half widths from Freedman et al. (2008, 2014).

One of the major issues with low- and medium-resolution cross-sections is that they cannot combine gases effectively due to their insufficient resolution. The usual approach is to sum the individual contributions of the various gases and weight them by their volume mixing ratio (Sharp & Burrows, 2007; Waldmann et al., 2015). The multiplication property of transmission is only valid when these calculations are done monochromatically. To illustrate this issue for low- and medium-resolution cross-sections, consider the simple two-gas problem in a single layer where each gas has an identical transmission spectrum over a given wavenumber interval which is larger than the bin size for the cross-section, $\Delta\nu$, shown in Figure 4.6.1. In this extreme example, the multiplication using these two methods produces mean transmissions that are a factor of two different. Correlated- k preserves this multiplication property of transmissions (Goody & Yung, 1989), i.e. the calculations are the same as for monochromaticity. Using cross-sections invokes using a mean transmission over the interval and hence the multiplication can produce fallacious results.

The opacity-sampling method requires a sufficient number of samples within a specified bin to correctly estimate the integral of the cross-section, though this number is ill-defined and depends on spectral resolution. Line et al. (2016b) show that the 1cm^{-1} resolution opacity sampling is sufficient for their purposes ($1.0\text{-}2.5\mu\text{m}$); however, this may not be a good resolution for either higher spectral resolutions or longer wavelengths. We briefly explore the effects of using opacity sampling as a means of computing the gaseous absorption in the next section, but do not include it in our main results section as we did not have the time to implement it in our code.

There are two key issues with using the Doppler-broadened cross-sections provided by the ExoMol project (Hill et al., 2013). One is that they fundamentally should not be used to solve the radiative transfer equation in mixed-gas atmospheres by combining single gas cross-sections as they are either ineffective (low- and medium-resolution) or computationally expensive (high-resolution). The second is that pressure-broadening is a key property in calculating atmospheric spectra that should not be ignored. These points




Figure 4.6.1: An illustrative example of the effects on transmission when multiplying monochromatic and non-monochromatic wavelength regions. Here we imagine combining the pink transmission with an identical distribution. Multiplying the transmissions monochromatically (as for correlated- k) gives $1 \times 1 = 1$ for 0 to $0.5\Delta\nu$, and $0 \times 0 = 0$ for 0.5 to $1.0\Delta\nu$. When averaging this pink transmission over the bin (as for cross-sections), we get an average value of 0.5 for the new green distribution. When it is multiplied by an identical green distribution, it produces a transmission of 0.25 (blue).




Figure 4.6.2: The (new, ExoMol) CH_4 k -table mean absorption coefficient at low-pressure (blue) is compared to the Doppler-broadened cross-sections at a resolution of $\Delta\nu = 1\text{cm}^{-1}$ from the ExoMol project (pink) for 1600K. The spectra are binned to $\Delta\lambda = 0.005\mu\text{m}$ resolution.

will be addressed in the next section.

At low pressures, where all of the pressure-broadening has ceased and the lineshape is dominated by Doppler broadening, the mean opacity (i.e. $\bar{k} = \sum_{i=1}^{NG} k_i \Delta g_i$) of our k -tables should be equivalent to the ExoMol cross-sections. We show this to be true for CH_4 at 1600K in Figure 4.6.2, where the cross-sections were taken from www.exomol.com at a resolution $\Delta\nu = 1\text{cm}^{-1}$. The spectra are binned to $\Delta\lambda = 0.005\mu\text{m}$ resolution. This is also the case for the other relevant gases.

By flattening our k -distribution, i.e. replacing k_i for each g -ordinate with $\bar{k}$ so that the k -distribution $k(g)$ is now a flat (constant) distribution, our k -tables at these low pressures are mathematically identical to using cross-sections, and so will fall foul to the aforementioned gas-mixing issues.

Instead of using the ExoMol cross-sections directly where we are limited to the temper-

atures ranges supplied online, we carry out our comparison using only the lowest pressure points of our flattened k -tables as a proxy. In the next section we show first that this is a good approximation to using cross-sections of resolution $\Delta\nu = 1\text{cm}^{-1}$.

4.7 Results

First we would like to show the inability to effectively mix gases with cross-sections. We create a simple one-layer model at low pressure (and therefore only Doppler-broadened), with a temperature of 1000K, and containing 50% H_2O and 50% CH_4 . We calculate the mean transmission for the correlated- k using:

$$\bar{T} = \sum_{i=1}^{NG} \sum_{j=1}^{NG} e^{-(k_i m_a + k_j m_b)} \Delta g_i \Delta g_j \quad (4.1)$$


and the cross-sections:

$$\bar{T} = e^{-(\bar{k}_a m_a + \bar{k}_b m_b)} \quad (4.2)$$


where i and j represent the index of the weights and a and b are labels for the first and second gas respectively, k is the absorption coefficient ($\text{cm}^2 \text{ molecule}^{-1}$, where $\bar{k}$ represented an average over the bin), and m is the absorber amount (molecule cm^{-2}). Absorption ($1 - \bar{T}$) is calculated with four different methods: 1) using our k -tables described earlier, the online ExoMol cross-sections directly at two different resolutions - 2) 1cm^{-1} and 3) 25cm^{-1} ; 4) our cross-section proxy k -tables described in the previous section; and 5) the opacity-sampling method at 1cm^{-1} resolution using the 0.01cm^{-1} ExoMol cross-sections. Methods 1 and 4 are combined using equation 4.1, while methods 2,3 and 5 are combined using equation 4.2. The spectra are then binned to $\Delta\lambda = 0.005\mu\text{m}$ resolution, and shown in Figure 4.7.1a. We choose these two resolutions for the cross-sections because the prior is the usual resolution used by other modellers, and the latter because it is closer to the typical resolution that we are presenting in Figure 4.7.3b and subsequent figures in the visible ($\sim 0.6\mu\text{m}$). Beyond $\sim 3\mu\text{m}$, our k -tables are in fact higher resolution than $\Delta\nu = 1\text{cm}^{-1}$. This means that we expect the cross-sections (and opacity sampling method) to become relatively less accurate at longer wavelengths. This is

also generally true because the Doppler width is proportional to the wavenumber, i.e. at long wavelengths lines become increasingly narrow and consequently require a higher resolution to be resolved.


In Figures 4.7.1a, 4.7.1b, and 4.7.1c we see that the opacity-sampling method produces absorption that is approximately of the same order as the correlated- k and line-by-line (only in Figure 4.7.1c for clarity) methods. However it also exhibits a persistent underestimation of the absorption and a large increase in noise as we go to longer, less sampled wavelengths. The effect is especially pronounced in certain bands such as the 2.5-3 μm band in Figure 4.7.1c. The opacity-sampling method therefore seems to have an accuracy and computation speed somewhere between the cross-section and correlated- k method. Unfortunately, to investigate this method further it would require a major overhaul of our code, for which we do not have the time to perform, and hence do not include it in further analysis. Figure 4.7.1a also shows that the overlapping bands can cause up to ~ 2 orders of magnitude change in absorption when comparing the correlated- k and cross-section methods. We believe these ‘ghost features’ form because the average transmissions computed from equation 4.2 do not capture the non-linear relationship between transmission and opacity. As expected, beyond 10 μm the absorption varies enormously due to an increase in the number of binned lines per resolution element in the cross-sections. Also it shows that our k -table proxy cross-sections exhibit approximately the same (if not smaller) effects than the two different ExoMol cross-sections. For the rest of this chapter when we refer to ‘cross-sections’, we will use these equivalent cross-sections calculated from the k -tables, rather than those directly from ExoMol. As shown in Figure 4.7.1c, the line-by-line calculation between 2-3 μm agrees very well with our correlated- k method and shows that there is a real concern for combining cross-sections at a resolution of $\Delta\nu = 1\text{cm}^{-1}$, especially at longer wavelengths. By comparing our correlated- k method to a high-resolution cross-section (0.01cm^{-1}) in Figure 4.7.2, it is evident that they now share a similar spectral morphology; however, the increase in resolution of the cross-section now causes a large increase in computation time. In a vast majority of real-life data resolutions, using the correlated- k method is therefore faster or more accurate.



(a) Absorption calculated for a layer consisting of 50% CH_4 , 50% H_2O at 1000K and path amount of 1×10^{20} molecule cm^{-2} . Pink is opacity-sampling at 1cm^{-1} resolution using the 0.01cm^{-1} resolution ExoMol cross-sections. Gold and blue are the cross-section absorption spectra at 1cm^{-1} and 25cm^{-1} resolution respectively, black line is the correlated- k method with a resolution of $0.002\mu\text{m}$ the green line is the cross-section derived from our correlated- k method. All absorption spectra are binned to a final resolution $0.005\mu\text{m}$ for comparison.



(b) Same as Figure 4.7.1a but for a final binned resolution of $0.1\mu\text{m}$ for easier comparison.



(c) Same as Figure 4.7.1a but zoomed in to $2\text{--}3\mu\text{m}$. The purple dots are a line-by-line calculation calculated at $\Delta\nu = 0.01\text{cm}^{-1}$.

Figure 4.7.1: Illustration of accuracy of different gaseous absorption methods.




Figure 4.7.2: Similar to Figure 4.7.1a, except pink is the cross-section absorption spectra at 0.01cm^{-1} resolution, and black is the same correlated- k calculation.

4.7.1 Primary Transit

In this example we take a simplified but more realistic atmosphere found in exoplanetary science, and compare its resulting spectra using the two different gaseous absorption methods. We do this using an isothermal (2000K) Hot Jupiter-esque primary transit as our example ($1.06M_{Sat}$, $1.4R_{Jup}$ around a Solar-size star), which we refer to as PT1.


Here we limit the k -tables to have only Doppler broadening by using the lowest pressure point in our grid (1×10^{-7} atm). This means we are directly comparing k -tables with cross-sections, i.e. the effect of flattening over the g -distribution (by assuming $k_i = \bar{k}$). From this section onwards, we also include collision-induced absorption from H_2 - H_2 and H_2 -He in our realistic atmosphere calculations.

In Figure 4.7.3a, we have two different atmospheres, each with a set of spectra calculated by the correlated- k method and the cross-section method. We have a 100% CH_4 atmosphere, and a 100% H_2O atmosphere. Here, no mixing of the gases is required, and therefore the two different methods produce almost identical results, with the differences being more pronounced for H_2O . These results are relatively similar because the propagating error in transmission multiplication throughout the pressure-varying path due to bin-averaging is small compared to those when averaging with overlapping gaseous bands. Note that the changes are not due to pressure-broadening effects, as the k -tables only contain Doppler broadening.


We introduce a third atmosphere, composed of 50% CH_4 and 50% H_2O . In Figure 4.7.3b, we compare the spectra produced by the Doppler-broadened correlated- k and the cross-section methods. We find that not only do the cross-sections overestimate the transit depth, but indeed they change the morphology of the spectrum itself, similar to the results in Figure 4.7.1a.

4.7.2 Brown Dwarfs

In the previous two subsections we verified our correlated- k method via line-by-line calculations, and showed the impact of various resolutions of cross-sections on the morphology of the spectra and absorption profiles of simple exoplanet atmospheres. In this subsection and the next we intend to show two things for multiple types of realistic H_2 -He-dominated



(a) Spectra for two PT1 (defined in manuscript) exoplanets consisting of a single gas (100% CH_4 in pink and black, 100% H_2O in gold and green) calculated using two methods. Legend labels X-Sec and Corr- k represent cross-sections ($k_i = \bar{k}$) and correlated- k (Doppler-broadened only, $k_i(g)$). The sudden drop in CH_4 is because the linelist only contains information from $0.83\mu\text{m}$ onwards, and hence gaseous absorption due to CH_4 is zero.



(b) Spectra for one PT1 exoplanet consisting of 50% CH_4 , 50% H_2O . Pink line is the correlated- k method (Doppler-broadened only), blue is the cross-section method.

Figure 4.7.3: Illustration of overlapping gaseous absorption errors.

atmospheres: 1) the effects of using cross-sections to incorrectly mix gases, and 2) the effects of including pressure-broadening on the resulting spectra.

To illustrate the effects of using cross-sections, correlated- k with no pressure-broadening (smallest pressure level available at our k -tables, $\sim 1\text{e-}7$ atm), and correlated- k with pressure-broadening to calculate the spectra of brown dwarf atmospheres, we use the best fit parameters of a previous unpublished retrieval performed on the object GL570D, a typical late-T dwarf as an example. This T dwarf contains significant amounts of H_2O , CH_4 , and NH_3 (3.5×10^{-4} , 4.0×10^{-4} , 2.3×10^{-4} respectively for their volume mixing ratios, which are representative values that are constant in height). The spectrum also contains Na and K with VMRs of 3.5×10^{-5} and 3.5×10^{-6} respectively. The mass and radius are $41.5M_{Jup}$ and $1.38R_{Jup}$ respectively, and the temperature profile is a 673.5K grid model from Saumon & Marley (2008) with the appropriate surface gravity. The temperature profile, along with those from the next two sections, is shown in Figure 4.7.5. The exact numbers for these parameters do not significantly change the outcome of the results.

Figure 4.7.4 contains two plots. In the upper plot, we have spectra calculated using the cross-section method, correlated- k method with no pressure broadening, and the correlated- k method with pressure broadening. In the bottom plot, we have calculated the transmission in a single layer at 700K and 1 atm, weighted by the volume mixing ratios with a given path amount (1×10^{23} molecules/cm $^{-2}$) for NH_3 , CH_4 , H_2O , Na and K, so that it can be easily seen which spectral features belong to which species.

In all of the large bands we find a discrepancy between the cross-section method and the correlated- k with pressure-broadening method, usually of multiple orders of magnitude. For example, the $2.8\mu\text{m}$ water band produces luminosities that vary by up to $\sim 3\text{-}4$ orders of magnitude. The same can be said for the $3.2\mu\text{m}$ CH_4 band. In molecular bands, the addition of pressure broadening can produce an additional order of magnitude change. These are huge effects that would certainly be reflected in parameter estimation during retrievals.




Figure 4.7.4: Top plot: Disc-averaged spectra calculated from best fit parameters for GL570D using various methods of calculation. The spectra are calculated for cross-sections with no pressure-broadening (blue), correlated- k with no pressure broadening (green), and correlated- k with pressure broadening (pink). Bottom plot: the transmission calculated for the three major absorbing species in the atmosphere in a single layer of atmosphere at 700K, where each gas has been weighted by its volume mixing ratio for a given path amount.




Figure 4.7.5: Temperature profiles for the brown dwarf ('T dwarf' in pink), the primary transit of HD189733b ('Primary' in blue) and secondary eclipse of HD189733b ('Eclipse' in green).

4.7.3 Primary Transit and Secondary Eclipse


In Figure 4.7.6a we present a fiducial model from Barstow et al. (2014) of the primary transit spectrum of HD189733b, calculated using the three methods as described before for the brown dwarfs. The atmosphere primarily consists of the usual H_2 (0.87), He (0.13), H_2O (1×10^{-4}), Na (5×10^{-6}) and K (5×10^{-7}) gases, along with a haze to cover the observed Rayleigh slope.

Similar to the brown dwarf cases, we find that the total opacity and morphology of the spectra differ greatly between the methods. The transit depth can vary by up to 1%, as seen at the $2.8\mu\text{m}$ water band. Contrary to the brown dwarfs, we note that the pressure broadening effects are subdued for primary transit. This is because we are probing much higher in the atmosphere than in the brown dwarf case, where Doppler broadening is the main agent of broadening. The Na and K lines are changed vastly, although we must note that the visible region is especially subject to large changes in $\Delta\nu$ (a constant $\Delta\lambda$, as we have, produces larger $\Delta\nu$ in the visible), and thus the effects are more pronounced than for the infrared.


For secondary eclipse, the VMRs used are slightly different than for primary transit, due to the different retrieved profiles. H_2O , CO_2 , CO and CH_4 all have VMRs of 1×10^{-4} . H_2 , He, Na, and K have VMRs of 0.9, 0.1, 5×10^{-6} , and 1×10^{-7} respectively. In Figure 4.7.6b we see a familiar increase in opacity and change in morphology for the cross-section method. The pressure-broadening is also slightly more significant compared to primary transit, as we are probing lower in the atmosphere. The effects here are smaller than for the brown dwarfs, where a change in the method can increase the flux by half an order of magnitude in the mid- and far- infrared, or change it by an order of magnitude in the near-IR. This is because the contribution functions peak at a lower pressure in the atmosphere but are still appreciable enough to become apparent in the spectra.

4.8 Conclusions

A major discrepancy between different radiative transfer models is how gaseous absorption is calculated. We investigated the effects of using two different gaseous absorption



(a) A fiducial HD189733b atmosphere for primary transit. The spectra are calculated for cross-sections (pink), correlated- k with no pressure broadening (blue), and correlated- k with pressure broadening (green).



(b) A fiducial HD189733b atmosphere for secondary eclipse. The spectra are calculated for cross-sections (pink), correlated- k with no pressure broadening (blue), and correlated- k with pressure broadening (green).

Figure 4.7.6: HD189733b simulated planets for different absorption methods.

methods (correlated- k and cross-sections) to calculate spectra in a variety of atmospheres of varying complexities. We also investigated the effects of including H₂-He pressure broadening in the more complex atmospheres. These investigations are important because radiative transfer models are often coupled to inverse methods, which can iterate over millions of forward models in a large parameter space; a single forward model error, such as a ‘ghost feature’, would propagate through to all retrieved PDFs influencing them in a non-trivial way, i.e. an error in the absorption spectrum of NH₃ could influence the retrieved temperature profile.

First we investigated the effects of using H₂-He pressure broadening instead of air, and the effect they had on our k -tables. We found non-negligible effects for all gases except TiO and VO. We also showed the differences of using high-temperature line lists, which usually include more absorption at higher temperatures, and have a slight shifting of the line centres.

Following this, we first showed that for test cases with resolutions of $\Delta\nu = 1\text{cm}^{-1}$ the cross-section method overestimates the amount of absorption present in the atmosphere and should therefore be used with caution. The morphology of the spectra also changes and produces ‘ghost’ features for mixed-gas atmospheres. When considering our flattened k -table cross-sections, the effect can produce multiple orders of magnitude change in the flux received from brown dwarfs in certain wavelength regions. The effect is similar but smaller for primary transit, and is closer in order to a $\sim 1\%$ change in transit depth. The flux ratio of secondary eclipse exoplanets can find an order of magnitude change in the near-IR, with the effect becoming smaller for longer wavelengths. Correlated- k can produce similar results to line-by-line and very high-resolution cross-sections, but is much less computationally expensive. The inclusion of H₂-He pressure broadening similarly changes the total flux found in the spectra of brown dwarfs and secondary transit exoplanets by up to an order of magnitude, while only making slight changes to the spectra of transiting exoplanets.

If we take into account the sizable discrepancies found by Hedges & Madhusudhan (2016) between the different aspects of pressure broadening from medium- and higher resolutions, the issues discussed in this chapter might be even more serious than suggested. However, at higher resolutions the differences we present here will become less significant,

and so the main error source will switch from those discussed here to those discussed by Hedges & Madhusudhan (2016).

We conclude that inaccurate use of cross-sections and omission of pressure broadening can be key sources of error in the modelling of brown dwarf and exoplanet atmospheres. These sources of error in the forward model may produce strong biases in the probability distribution functions of retrieved parameters.

Chapter 5

Cross-Comparison Retrievals for JWST

5.1 Introduction

As discussed in Chapter 4, the retrieved probability distribution functions for parameters relating to exoplanetary and brown dwarf atmospheres, gained from sophisticated inverse methods, are highly dependent on the various parameterisations used, and also the accuracy of the forward model. While we have already considered the impact of calculating gaseous absorption, there may be other effects accounted for in different models that may alter the resulting spectrum when given identical inputs. These could be as technical as the sampling points in the k -tables, use of different line lists, or as simple as the definition of the pressure level where the radius is defined (1 bar or 10 bar is common). In short, every forward model is different, and the extent to which these models are different could potentially seriously affect the resulting calculated spectra and retrieved PDFs. In this chapter, we investigate the differences in forward model and retrieved probability distributions for five simulated planets with JWST instrumental specifications using three different 1D radiative transfer codes: NEMESIS, TauRex, and CHIMERA (Irwin et al., 2008; Waldmann et al., 2015; Line et al., 2014). This work not only verifies our model by comparing it to others, but also allows us to demonstrate that a number of parameters can be well constrained based upon realistic data sets. The simulated spectra is for

primary transit only, but the work could be easily adapted to other geometries, i.e. for brown dwarfs. Primary transit exoplanets were chosen instead of brown dwarfs because the other groups had a larger interest in that particular geometry (which is typical of the field in general), although we would expect to find qualitatively similar results had we investigated brown dwarfs in thermal emission instead. All NEMESIS calculations are my work, and the other calculations are from their respective groups. We also compare calculations for single-gas atmospheres so that it is easier to detect which gases are contradicting each other.

For all tests we calculate spectra in the range to be $0.5\text{-}10\mu\text{m}$, and have a spectral resolution $R = 100$ for a primary transit geometry. Each planet is calculated on exactly the same pressure grid, each with 100 layers.

5.2 Forward Model Comparison

Different approaches and input data (such as line lists) between the different codes are known *a priori*, but the extent to which these different approaches affect the resulting spectrum is not. There are two key categories of differences between the codes discussed here. CHIMERA and TauREX calculate their k -tables from ultra-high resolution cross-sections, whereas NEMESIS calculates it from the line-by-line spectrum directly. The second main difference is the line data used to calculate the k -tables for each gas. This includes the line lists themselves, as well as pressure-broadening parameters.

In Table 5.2.1, we present the line lists used for each gas by each code.

Table 5.2.1: Comparison of line lists used by each code for each gas.

Gas	NEMESIS	CHIMERA	TauREX
CH ₄	Yurchenko & Tennyson (2014)	Yurchenko & Tennyson (2014)	Yurchenko & Tennyson (2014) (updated)
H ₂ O	Barber et al. (2006)	Barber et al. (2006)	Barber et al. (2006) (updated)
NH ₃	Yurchenko et al. (2011)	Yurchenko et al. (2011)	Yurchenko et al. (2011)
CO	Rothman et al. (2010)	Rothman et al. (2010); Tipping (1976)	Rothman et al. (2010)
CO ₂	Tashkun & Perevalov (2011)	Huang et al. (2013, 2014)	Rothman et al. (2010)

To investigate the effects of these differences between our forward models, we first compare single-gas primary transit calculations for five different gases (H₂O, CH₄, NH₃, CO, CO₂). After this comparison, we concocted five semi-realistic planetary atmospheres using each code, producing fifteen (5x3) different spectra. In the following subsections,

we will outline each planet and compare the spectra produced by the different codes for individual gases and semi-realistic atmospheres.

5.2.1 Individual Gas Comparison


As there may be many problems with complicated mixed-gas atmospheres, we chose first to model an array of single-gas atmospheres using a fake planet with an isothermal atmosphere (500-2000K in 500K steps), planetary radius $1R_{Jup}$ at 10 bar, mass $1M_{Jup}$, and solar radius $1R_{Sun}$, for a range of the five different aforementioned gases, of VMR between 1×10^{-6} to 1.0 (steps being an order of magnitude). Here we will present the individual differences of each gas for selected temperatures and volume mixing ratios. In the work presented here, CHIMERA and TauREX calculations were completed by the other groups, which we compare to our NEMESIS calculations.

CH₄


The low-temperature (500K) comparison of the single-gas atmosphere of CH₄, i.e. Figure 5.2.1a, shows a remarkable similarity among all codes, with very little deviance. Each code is using a similar, however not exactly, the same line list. All use the ExoMol line list for CH₄ (Yurchenko & Tennyson, 2014). However, TauREX is uniquely positioned in UCL to use unpublished, updated linelists. Therefore TauREX uses a line list which contains approximately 40 billion lines, rather than 10 billion as for NEMESIS and CHIMERA. At low temperatures these additional lines produce negligible absorption, and so at 500K all of the spectra are almost identical. The small differences visible are due to the fact that both CHIMERA and TauREX use a method of calculating their k -tables from very high resolution cross-sections, as opposed to directly from a line-by-line calculation as NEMESIS does. The slightly different binning procedure then gives these small differences - this is true for all results presented here.

The high-temperature (2000K) comparison, i.e. Figure 5.2.1b, is also similarly good but there are some small deviations between the codes. TauRex, on average, produces slightly more absorption than the others. This is because the extra 30 billion lines in their k -tables produce non-negligible absorption at high temperatures.


Note that the sudden drop in absorption at $0.83\mu\text{m}$ is because the line list does not




(a) CH₄, 500K.




(b) CH₄, 2000K.




(c) H₂O, 500K.



(d) H₂O, 2000K.




(e) NH₃, 500K.




(f) NH₃, 2000K.


Figure 5.2.1: Primary transit spectra for modelled planet consisting of a single gas (labelled) with a VMR of 1×10^{-6} and 1.0 for top and bottom respectively, plus He and H₂ in a ratio of 85:15 (H₂:He), with an isothermal temperature of 500K (left) or 2000K (right). Spectra have been binned to have $0.01\mu\text{m}$ resolution. The smaller VMR has a constant added to the Y-axis for clarity, and hence is the top spectrum; the bottom spectrum contains the larger VMR.




(g) CO, 500K.



(h) CO, 2000K.



(i) CO₂, 500K.



(j) CO₂, 2000K.

Figure 5.2.1: Primary transit spectra for modelled planet consisting of a single gas (labelled) with a VMR of 1×10^{-6} and 1.0 for top and bottom respectively, plus He and H₂ in a ratio of 85:15 (H₂:He), with an isothermal temperature of 500K (left) or 2000K (right). Spectra have been binned to have $0.01\mu\text{m}$ resolution. The smaller VMR has a constant added to the Y-axis for clarity, and hence is the top spectrum; the bottom spectrum contains the larger VMR. (cont.)

extend below this wavelength. Similar sudden drops are seen in the rest of the spectra presented here, and are caused by the same reason.

H2O

The low-temperature (500K) comparison of the single-gas atmosphere of H₂O, i.e. Figure 5.2.1c, shows a remarkable similarity among all codes, with very little deviance. This is because each code is using the same linelist, with the very small differences due to the method of creating the *k*-tables.

The high-temperature (2000K) comparison, i.e. Figure 5.2.1d, shows that CHIMERA and NEMESIS agree very well; however, TauRex, on average, produces significantly less absorption than the others. This probably is because of the slightly different ExoMol line list available to TauREX, where line intensities have been adjusted.

NH3

The low-temperature (500K) comparison of the single-gas atmosphere of NH₃, i.e. Figure 5.2.1e, shows different results for each VMR. For the larger VMR, the codes overlap almost perfectly, except for some difference in the near-IR for CHIMERA. We believe this to be from an error in CHIMERA's *k*-table. However, in the smaller VMR case, CHIMERA and NEMESIS agree very well, while TauRex shows a lack of absorption in most bands when compared to the other two. Again, this is because of the slightly different ExoMol line list available to TauRex.

The high-temperature (2000K) comparison, i.e. Figure 5.2.1f, has the same issues as the lower temperature for the larger VMR, i.e. CHIMERA is underestimating the absorption below 2 μ m. For the smaller VMR, we find that below 2 μ m, all of the codes have different absorption rates. This is because of the error in CHIMERA's *k*-table and the slightly different ExoMol line list available to TauREX. Above 2 μ m, they are almost identical.

CO

The low-temperature (500K) comparison of the single-gas atmosphere of CO, i.e. Figure 5.2.1g, is largely the same for all codes. This is expected because all codes use mostly

the same linelist.

The high-temperature (2000K) comparison, i.e. Figure 5.2.1h, is largely similar for all codes for the small VMR case; however, the high VMR case is quite different. CHIMERA contains a lot more absorption in the first band, but is similar to NEMESIS for the last three bands. This may be due to the additional line list used by CHIMERA. TauRex exhibits an overall lack of absorption in the longer wavelength parts of the bands. At the time of writing, we are not sure of this cause; however, a line list different to the one reported may be the source.

CO₂


The low-temperature (500K) comparison of the single-gas atmosphere of CO₂, i.e. Figure 5.2.1i, is largely the same for all codes for the smaller VMR; however, the larger VMR shows that CHIMERA has multiple bands of absorption that are not apparent in TauRex and NEMESIS. Differences are easily explained by the fact that all codes are using different line lists.

The high-temperature (2000K) comparison, i.e. Figure 5.2.1j, is largely similar for all codes for the high VMR case (except from the extra absorption shortward of $1.2\mu\text{m}$ for CHIMERA). The low VMR case shows that CHIMERA and NEMESIS are quite similar, and that TauRex repeatedly shows areas of a lack of absorption. Again, differences are easily explained by the fact that all codes are using different line lists.


5.2.2 HD189733b

The first semi-realistic atmosphere that we are considering is that of HD189733b. This has the properties of planetary mass $M_p = 1.162 M_{Jup}$, planetary radius $R_p = 1.138 R_{Jup}$, solar radius $R_s = 0.781 R_{Sun}$, isothermal temperature $T_p = 1500\text{K}$, H₂O (3×10^{-4}), CO (3.5×10^{-4}), He/H₂ ratio = 0.17647, and two cases of no cloud and cloud (completely opaque cloud top at 10mbar).

The comparison without clouds, i.e. Figure 5.2.2a, produces quite similar results for CHIMERA and NEMESIS, however TauRex overestimates the amount of absorption, particularly in the Rayleigh-scattering visible part of the spectrum, and within the spectral bands.



(a) Simulated primary transit spectrum of HD189733b as described in the text for the three different codes with no clouds.



(b) Simulated primary transit spectrum of HD189733b as described in the text for the three different codes with a completely opaque cloud top at 10mbar.

Figure 5.2.2: HD189733b simulated planets.

With clouds, i.e. Figure 5.2.2b, we see exactly the same trend as before: TauRex overestimates the absorption across the board. A simple explanation of this might be that the quoted radius is not the same as the computed radius for TauRex, causing the constant shift.

5.2.3 Super-Earth, Low Mean Molecular Weight

The second semi-realistic atmosphere is that of a low mean molecular weight (MMW) Super Earth. This has the properties of planetary mass $M_p = 0.02 M_{Jup}$ ($6.36 M_{Earth}$), planetary radius $R_p = 0.238 R_{Jup}$ ($2.67 R_{Earth}$), solar radius $R_s = 0.216 R_{Sun}$, isothermal temperature $T_p = 400K$, H_2O (8×10^{-4}), NH_3 (1×10^{-4}), CH_4 (4×10^{-4}), He/H_2 ratio = 0.17647, and two cases of no cloud and cloud (completely opaque cloud top at 10mbar).


The comparison without clouds, i.e. Figure 5.2.3a, produces almost identical results across all of the codes, with TauRex producing slightly more absorption due to Rayleigh-scattering.

With clouds, i.e. Figure 5.2.3b, we see similarly that the codes produce almost identical results, although now NEMESIS seems to be underestimating the Rayleigh-scattering absorption.


5.2.4 Super-Earth, High Mean Molecular Weight

The final semi-realistic atmosphere is that of a high mean molecular weight Super Earth. This has the properties that are the same as the low MMW Super Earth, except for the mole fractions of H_2O (0.43), N_2 (0.13), CH_4 (0.14), CO_2 (0.25), and only a no cloud case is investigated.

The comparison in Figure 5.2.4 has a few vital differences between all of the codes. CHIMERA has a large increase in absorption for most of the spectrum compared to the other codes. Beyond $1\mu m$, TauRex and NEMESIS agree very well, however shortward of $1\mu m$ there is quite a large difference between all codes. This is due to different approaches to calculating the Rayleigh scattering, i.e. all codes primarily use values and expressions from Cox (2000) with some differences. NEMESIS uses additional parameters for CH_4 based on guesses from NH_3 and CO_2 data for refractive index and polarisability respectively. TauREX also uses Lorentz-Lorenz approximation for low density gases (Sneep &



(a) Simulated primary transit spectrum of Super-earth low MMW as described in the text for the three different codes with no clouds.



(b) Simulated primary transit spectrum of Super-earth low MMW as described in the text for the three different codes with a completely opaque cloud top at 10mbar.

Figure 5.2.3: Super-Earth low MMW simulated planets.




Figure 5.2.4: Simulated primary transit spectrum of Super-earth high MMW as described in the text for the three different codes with no clouds.

Ubachs, 2005), where the others do not.

5.3 Retrieval Comparison

Each of the five atmospheres described previously were then binned to a spectral resolution $R = \frac{\lambda}{\Delta\lambda} = 100$, and three sets of errors were attributed to each spectrum: 30PPM (parts per million), 60PPM, and 100PPM. These are supposed to represent a high, medium, and low signal-to-noise ratio expected by JWST. For reference, Neptune-sized planets around K-Stars have a transit depth of around 2500PPM, and so this would correspond to SNRs of approximately 80, 40, and 25 for each attributed error respectively. This will be similar to our Super-Earth case, discussed later. In the following sections, we perform retrievals on each other codes' forward model spectra, i.e. NEMESIS will perform a retrieval on a spectrum from the forward model of TauRex, and another retrieval the spectrum of CHIMERA; CHIMERA will retrieve on TauRex and NEMESIS, and TauRex on NEMESIS and CHIMERA. This is repeated for each of the three noise realisations, so that we have six retrievals per noise realisation, and eighteen retrievals in total for each object. We present ninety retrievals in total in this analysis (eighteen retrievals for each

object).

All three models used the Multinest algorithm (Feroz et al., 2013) as their inverse method, using 5000 live points for each retrieval. This eliminates any complication from different inverse methods.

Here we will describe the retrieval setup and results for each atmosphere.

5.3.1 HD189733b

The priors that all groups used for HD189733b are uniform between: $R_p = 1.046 - 1.278R_{Jup}$, $T_p = 1200 - 1800\text{K}$, gases $1 \times 10^{-12} - 1.0$, and cloud top pressure $1 \times 10^{-8} - 10$ bar.

Before performing retrievals on each model spectrum, we had a sanity check retrieval to ensure that our own model correctly retrieves on itself. We do this for the cloudy HD189733b case, which can be seen in Figure 5.3.1. As expected, the median values are almost identical to all ‘true’ values for every parameter. This verifies our inverse method is working correctly.

Figure 5.3.2 and the figures to come are all set up in a similar format. However, instead of presenting one posterior distribution for each model and noise realisation, we combine the posterior distributions of each object, giving the posterior samples equal weightings, for a given noise realisation. This means we have six posterior samples that are combined into a single posterior distribution for every noise realisation, which we will refer to as a ‘combined posterior’. Hence, each Figure contains the details of 18 retrievals, i.e. 6 per noise realisation (NEMESIS-on-TauRex, NEMESIS-on-CHIMERA, TauRex-on-NEMESIS, TauRex-on-CHIMERA, CHIMERA-on-TauRex, CHIMERA-on-NEMESIS). By combining these posteriors, we can give an account of the error caused by model dependencies. This of course causes our distributions to increase in error, and, on occasion, become multimodal. In every Figure, red represents 100PPM, black 60PPM, and blue 30PPM.

If we examine Figure 5.3.2, we find a number of interesting features. A lot of these features are present in the work to come, as we will find out.

Firstly, there is a general trend in the combined posterior from lower to higher signal-to-noise, whereby the retrieved PDFs for all parameters are almost model-independent




Figure 5.3.1: 1D marginalised posterior distributions and 2D correlations for a NEMESIS retrieval on the HD189733b spectrum with clouds, and errors of 60PPM generated by NEMESIS. Median and $\pm\sigma$ values plotted on top of 1D distributions as dotted lines. Blue dots and lines represent the true values. From top-left to bottom-right: $\log \text{H}_2\text{O}$ VMR, $\log \text{CO}$ VMR, Radius (R_{Jup}), temperature (K), $\log P_{Base}$ (bar).

at lower SNR, and become increasingly multimodal (model-dependent) at higher SNR. In this context, model-independent means that the combined posteriors produce a broad unimodal distribution; model-dependent means that the combined posteriors produce a multi-modal distribution. Typically, 60PPM seems to be relatively model-independent, and a crossover occurs to model-dependent somewhere between 60PPM and 30PPM (though we have not investigated this further).

The second feature to realise is that different parameters have different sensitivities to model-dependent / model-independent combined posterior distributions with changing SNR. That is to say, in Figure 5.3.2, retrieved H_2O , temperature, and cloud structure are model-independent (1D combined posteriors are unimodal) in the regime that we have explored. However radius, and to a lesser extent CO, in Figure 5.3.2, are both increasingly model-dependent with increasing SNR. While radius has a tri-modal PDF (and is thus the most model-dependent parameter), CO, interestingly, only has a bimodal distribution.

The radius differences are very small - retrieved values vary by approximately $\pm 0.001 R_{\text{Jup}}$ either side. Such small differences are difficult to track down, and could be explained by subtle differences in constant definitions (such as the decimal places in the definition of one Jupiter radii). The CO bimodality is due to the differences that we found for modelled CO in the previous section.

Now we examine the cloudy case, presented in Figure 5.3.3. Surprisingly, this object presents much less model-dependent combined posteriors than the cloud-free version. PDFs for all parameters and error realisations overlap very well, with only some multi-modality occurring for the highest SNR of CO. The radius multi-modality has likely disappeared because the combined posterior error increases due to the large degeneracy with the cloud top pressure. We also note a strange 2D correlation pattern seen in Figure 5.3.3 for cloud base pressure, which comes from the TauRex retrievals. We are not sure of the origin of this.

5.3.2 Super Earth Low MMW

The priors that all groups used for the low mean-molecular-weight Super Earth are: $R_p = 0.214 - 0.262 R_{\text{Jup}}$, $T_p = 200 - 7000\text{K}$, gases $1 \times 10^{-12} - 1.0$, and cloud top pressure $1 \times 10^{-8} - 10$ bar.




Figure 5.3.2: 1D marginalised posterior distributions and 2D correlations for a combined six retrievals per noise realisation (18 in total) on the HD189733b spectrum without clouds. Blue dots and lines represent the true values. From top-left to bottom-right: $\log \text{H}_2\text{O}$ VMR, $\log \text{CO}$ VMR, Radius (R_{Jup}), temperature (K), $\log P_{Base}$ (bar). Black data represents 60PPM, red represents 100PPM, and blue represents 30PPM.




Figure 5.3.3: 1D marginalised posterior distributions and 2D correlations for a combined six retrievals per noise realisation (18 in total) on the HD189733b spectrum with clouds. Blue dots and lines represent the true values. From top-left to bottom-right: $\log \text{H}_2\text{O}$ VMR, $\log \text{CO}$ VMR, Radius (R_{Jup}), temperature (K), $\log P_{Base}$ (bar). Black data represents 60PPM, red represents 100PPM, and blue represents 30PPM.

The combined posteriors for the low MMW Super Earth without clouds are presented in Figure 5.3.4. We find that the VMRs of NH_3 and CH_4 , the temperature, and the cloud top pressure are model-independent for all of the noise realisations tested. The radius is partly multi-modal for 60PPM, and fully tri-modal for 30PPM. H_2O is quite bi-modal for 60PPM, and fully bi-modal by 30PPM. The radius multimodality requires the same explanation as the HD189733b case. The bi-modality of H_2O is explained by the forward modelling differences of H_2O , as described in the previous section.

Now adding in clouds to the atmosphere, we retrieve the posteriors presented in Figure 5.3.5. In fact, all combined posteriors are unimodal up until 30PPM. At 30PPM, all parameters become either bi- or tri-modal. This can be explained by the forward modelling differences of all the gases, however it is intriguing to note that temperature’s bimodal distribution is less easy to explain. Because of the highly multimodal solutions throughout all of the other (non-temperature) parameters, the Monte Carlo propagation of these multimodal solutions has influenced the temperature retrieval in such a way that it requires a second mode.

5.3.3 Super Earth High MMW

The priors that all groups used for the high MMW Super Earth are the same as for the low MMW version; however, cloud top pressure is not retrieved at all as there are no clouds present.

The posteriors for the high MMW Super earth are presented in Figure 5.3.6. The posteriors are mostly unimodal, except for CO_2 and radius at 60PPM and 30PPM. The radius multimodality is explained in the same sense as the previous examples. The CO_2 multimodality is easily explained by the fact multiple different line lists are used by each code.

5.4 Conclusions

In this chapter we compared three different models to investigate the differences between forward models, and their influence on the retrieved posterior probability distributions using uniform Bayesian inference methods for typical resolutions of JWST at various




Figure 5.3.4: 1D marginalised posterior distributions and 2D correlations for a combined six retrievals per noise realisation (18 in total) on the low MMW Super Earth spectrum without clouds. Blue dots and lines represent the true values. From top-left to bottom-right: $\log \text{NH}_3$ VMR, $\log \text{CH}_4$ VMR, $\log \text{H}_2\text{O}$ VMR, Radius (R_{Jup}), temperature (K), $\log P_{Base}$ (bar). Black data represents 60PPM, red represents 100PPM, and blue represents 30PPM.




Figure 5.3.5: 1D marginalised posterior distributions and 2D correlations for a combined six retrievals per noise realisation (18 in total) on the low MMW Super Earth spectrum with clouds. Blue dots and lines represent the true values. From top-left to bottom-right: $\log \text{NH}_3$ VMR, $\log \text{CH}_4$ VMR, $\log \text{H}_2\text{O}$ VMR, Radius (R_{Jup}), temperature (K), $\log P_{\text{Base}}$ (bar). Black data represents 60PPM, red represents 100PPM, and blue represents 30PPM.




Figure 5.3.6: 1D marginalised posterior distributions and 2D correlations for a combined six retrievals per noise realisation (18 in total) on the high MMW Super Earth spectrum. Blue dots and lines represent the true values. From top-left to bottom-right: $\log \text{CH}_4$ VMR, $\log \text{H}_2\text{O}$ VMR, $\log \text{CO}_2$ VMR, Radius (R_{Jup}), temperature (K). Black data represents 60PPM, red represents 100PPM, and blue represents 30PPM.

levels of signal-to-noise for representative planets. We found that while different forward models produced varyingly different spectra, the vast majority of retrieved parameters are model-independent (i.e. unimodal) for errors of 60PPM and more. In general, 30PPM led to much more multimodal results for, in particular, certain gases (due to differences in line lists), and the radius. Temperature and cloud-top pressure were almost always model-independent. By combining the posterior samples into single combined posterior distributions, we, for the first time, have estimated the uncertainty in Bayesian inference due to differences in forward models. While using multiple models is the optimal method for retrievals, realistically for spectra with errors of 60PPM and more, any one model should retrieve similar posteriors to any other. This is of course limited to the models, resolution, and particular atmospheres that we tested in this study.

Chapter 6

T and Y dwarf Retrievals

The main objective of this thesis is to characterise the atmospheres of brown dwarfs as they transition from high to low temperatures. This transition occurs over a timescale of billions of years, and so, in practice, we characterise a selection of objects that span a range of temperatures from hot to cold. In this chapter, we will analyse eight T and Y dwarf spectra using the Bayesian methods we have developed and verified in the previous chapters. We will compare these to the literature values where possible, but as this is the first attempt at a retrieval for many of these targets, a full comparison is not always possible. We begin by discussing the data set, then the methods, results, and finish with our conclusions.

6.1 Data Sources

The Y dwarf spectra described here are those described by Cushing et al. (2011) and kindly provided by the author. These spectra are presented in Figure 6.1.1. Here we will describe the datasets, and currently estimated properties from Cushing et al. (2011).

6.1.1 Data Sources

FIRE/Magellan

Low-resolution ($R = 250\text{--}350$), $1\text{--}2.4\mu\text{m}$ spectra for three targets (WISE J0148-7202, WISE J0410+1502, and WISE J1541-2250) were obtained between July-August 2010 with the Folded-port InfraRed Echellette (FIRE; Simcoe et al. (2008, 2013)) mounted at the auxiliary Nasmyth focus of the Magellan 6.5m Baade Telescope. The calibration and reduction of these datasets is described in Kirkpatrick et al. (2011).

NIRSPEC/Keck II

WISE J2056+1459 was observed using the Near-Infrared Spectrometer (NIRSPEC; McLean et al. (1998, 2000)) located on one of the Nasmyth platforms of the 10m Keck II telescope on Mauna Kea, Hawaii. The $0.38''$ -wide slit in the low-resolution mode provides a resolving power of $R = 2500$. WISE J2056+1459 was observed with the N3 order sorting filter ($1.143\text{--}1.375\mu\text{m}$) on the night of 2010 October 21 (UT) and with the N5 order sorting filter ($1.431\text{--}1.808\mu\text{m}$) on the night of 2010 November 22 (UT). Calibration and reduction is described in Cushing et al. (2011). The final spectrum has a peak S/N of 8 and 6 in the J and H bands, respectively.

WFC3/Hubble Space Telescope

WISE J1405+5534, WISE J1738+2732, and WISE J1828+2650 were observed with the infrared channel of the Wide Field Camera 3 (WFC3; Kimble et al. (2008)) on-board the Hubble Space Telescope (HST) as a part of a Cycle 18 program (GO-12330, PI: J. D. Kirkpatrick). The G141 grism was used to perform slitless spectroscopy of each brown dwarf covering the $1.07\text{--}1.70\mu\text{m}$ wavelength range at a resolving power of $R \approx 130$. The




Figure 6.1.1: Spectra for the Y dwarf dataset. The flux is normalised at 1.26 microns, with a constant added to each spectrum to aid visual comparison. Error bars are included.

data reduction technique is described in Cushing et al. (2011).

6.1.2 Binning Procedure

As the data come from a variety of sources, we attempt to homogenize the spectra so that they can be more easily compared and matched to a uniform modelling procedure. We integrate our spectra to contain bins with full-width at half maximum of $0.01\mu\text{m}$ over the full possible range of a given spectrum, with instrumental functions given by a Gaussian. For the modelling, we again use a FWHM of $0.01\mu\text{m}$, where our modelled spectrum is at first a tophat function for each bin, which is then convolved with another tophat to create a triangular function, to approximate our instrumental Gaussian function. This procedure is the same for all modelled objects discussed here, and is what we plot in all following spectral plots.

6.1.3 Currently Estimated Properties

Cushing et al. (2011) compared their near- infrared spectra to a new preliminary grid of model spectra generated with the model atmospheres of Marley & Saumon. A detailed description of the basic models can be found in Marley et al. (2002); Saumon & Marley (2008); Cushing et al. (2008); Stephens et al. (2009). This preliminary grid includes NH_3 absorption modelled with the NH_3 line list (Yurchenko et al., 2011) and a new prescription for the collision induced opacity for H_2 (Saumon et al., 2012).

The grid consists of solar metallicity, cloudless models with the following parameters: $T_{\text{eff}} = 200\text{-}1000\text{ K}$ in steps of 50 K, $\log g = 3.75\text{-}5$ (surface gravity) in steps of 0.25 (cm s^{-2}), and the eddy diffusion coefficient (which affects the vertical extent of the clouds) $K_{zz} = 0\text{-}10^4\text{ cm}^2\text{s}^{-1}$. Although the opacities of the condensate clouds are not included in the atmospheric models, i.e., they are cloudless, the effects on the atmospheric chemistry due to the rainout of the condensates is accounted for in the models. This assumption is reasonable for the silicate and liquid iron clouds since they form well below the observable photosphere (see however, Burgasser et al. (2010)) but may not be valid if, as expected, H_2O clouds form high in the atmosphere of the coldest models. The eddy diffusion coefficient, K_{zz} , parameterizes the vigor of mixing in the radiative layers of the atmosphere. The best-fit parameters are shown in Table 6.1.1

After most of the analysis of these objects had been performed, Leggett et al. (2017) obtained updated spectral measurements of the aforementioned Y dwarfs, and performed another grid model analysis on a range of 24 Y dwarf objects. These results are summarised in Table 6.1.2. When referencing spectral type of each object in the rest of the chapter, we took the difference between the types presented in Table 6.1.2 and Table 6.1.1, because we did not want to give extra weight to one set of grid models over the other. Ultimately, these previous spectral typings only act as a general guide for trends in the atmospheres - our Bayesian analysis will provide a much more detailed look at these trends than these spectral typings ever could.

Table 6.1.1: The parameters for the best-fitting Marley & Saumon models are listed and the range of parameters consistent given in parentheses.

Object	Type	T_{eff} (K)	$\log g$ (cm s $^{-2}$)	K_{zz} (cm 2 s $^{-1}$)	R (R_{Jup})	M (M_{Mup})
WISE J0148-7202	T9.5	500 (500-500)	4.50 (4.50-4.75)	4	1.04 (0.96-1.04)	13 (13-21)
WISE J0410+1502	Y0	450 (400-500)	3.75 (3.75-4.25)	0	1.22 (1.09-1.22)	3 (3-9)
WISE J1405+5534	Y0	350	5.00	4	0.86	30
WISE J1541-2250	Y0	350	4.50 (4.25-4.5)	4	1.01 (1.01-1.07)	12 (8-12)
WISE J1738+2732	Y0	350 (350-400)	4.75 (4.75-5.00)	4	0.93 (0.86-0.94)	20 (20-30)
WISE J1828+2650	Y0	≤ 300	...	...	...	...
WISE J2056+1459	Y0	350 (350-400)	4.75 (4.50-5.00)	4	0.93 (0.86-1.01)	20 (12-30)

Table 6.1.2: The parameters for the Leggett et al. (2017) analysis. Metallicity is the fraction of mass that is not hydrogen or helium relative to the sun.

Object	Type	T_{eff} (K)	$\log g$ (cm s $^{-2}$)	M (M_{Jup})	Metallicity [m/H]
WISE J0148-7202	N/A	N/A	N/A	N/A	N/A
WISE J0410+1502	Y0	410-440	4.25-4.75	8-20	0 - +0.3
WISE J1405+5534	Y0.5	370-400	4.3-4.8	9-21	-0.15 - + 0.15
WISE J1541-2250	Y1	360-390	4.25-4.75	8-20	-0.05 - + 0.25
WISE J1738+2732	Y0	410-440	4.0-4.5	5-14	-0.05 - +0.25
WISE1828+2650	$\geq Y2$	310-340	3.75-4.25	3-8	-0.6 - -0.3
WISE2056+1459	Y0	410-440	4.25-4.75	8-20	0 - +0.3

6.2 Methods

The objects described in Section 6.1.1 will be analysed in a similar manner to those in Chapter 5, except here we are modelling disc-averaged spectra, not primary transit. By thoroughly examining the small sample of atmospheres in a Bayesian framework, we can begin to build up a more conclusive picture of the evolution of substellar objects down to Solar System temperatures. Our work will be an improvement on these previous analyses as it is not as restricted to the assumptions found in 1D radiative-convective equilibrium grid models. The top-down approach to analysing the spectra means that our atmospheric parameters will be led by the data, rather than a rigid model. On the other hand, our solutions may not be physically self-consistent.

As our inverse method, we use the Bayesian Nested Sampling algorithm MultiNest, as described in Chapter 5 and Chapter 3. We couple this to the forward model of NEMESIS, again as in Chapter 5. We retrieve uniform-with-pressure volume mixing ratios of H_2O , CH_4 , NH_3 , Na, and K (and CO_2 and CO for our benchmark case); the radius; the mass; the temperature profile; and the cloud structure.

6.2.1 Temperature Profile

In previous work, we examined a range of temperature profile parameterisations and how they affected retrieved probability distributions. This work was completed using CHIMERA at a six-week fellowship in University of California, Santa Cruz, and only have preliminary results due to this short timeframe. From these initial results we conclude that a temperature-pressure profile parameterisation must balance complexity and freedom with some degree of restriction implemented by the natural laws of physics. As such, we decide that the five-parameter parameterisation of the temperature profile outlined in Robinson & Catling (2012) is sufficiently balanced to allow for complexity while still maintaining a reasonable amount of physics.

The temperature profile parameterisation from Robinson & Catling (2012) is an analytic solution for radiative-convective balance in an irradiated grey atmosphere. Here we derive a temperature profile parameterisation for a non-irradiated grey atmosphere by cancelling out the irradiated terms. We can break down this parameterisation into two

physically meaningful sections: radiative and convective.

Gray Approximation

First we make the gray approximation, i.e. that the gray thermal optical depth τ is related to the atmospheric pressure p by a power law in the form

$$\tau = \tau_0 \left(\frac{p}{p_0} \right)^n \quad (6.1)$$

where τ_0 is the gray infrared optical depth integrated from the top of the atmosphere to the reference pressure p_0 , and n is a parameter that controls the strength of the scaling. In the simplest case, when an absorbing gas is well mixed and the opacity does not depend strongly on pressure, we will have $n = 1$. This can physically correspond to Doppler broadening. A common scenario is to have a well-mixed gas providing collision-induced opacity (e.g. $\text{H}_2\text{-H}_2$) or pressure-broadened opacity so that $n = 2$. Typically, n takes a value between 1 and 2, as has been parameterised into some radiative transfer models (Heng et al., 2012). In some cases, where the mixing ratio of the absorbing gas depends strongly on pressure, larger values of n have been proposed, such as that for water vapor in Earth’s lower troposphere (Frierson et al., 2006). In our retrievals, we allow n to vary between 1 - 5, p_0 is given as 1000 bar, and τ_0 is retrieved to be between 1 and 1×10^5 .

Convective Zone

In the convective zone of a planetary atmosphere, we take the temperature-pressure profile to be similar to that of a dry adiabat, although somewhat less steep because of, for example, the latent heat released by condensation of volatiles during convective uplift.

The dry adiabatic temperature variation is given by Poisson’s adiabatic state equation:

$$T = T_0 \left(\frac{p}{p_0} \right)^{(\gamma-1)/\gamma} \quad (6.2)$$

Here T_0 is a reference temperature at the reference pressure p_0 and γ is the ratio of the specific heats at constant pressure (c_p) and volume (c_v):

$$\gamma = \frac{c_p}{c_v} \quad (6.3)$$

Kinetic theory allows us to relate the ratio of specific heats to the degrees of freedom, N , which dictates the number of dimensions of movement that can occur (including translation, rotation, and vibration), for the primary atmospheric constituent(s), where

$$\gamma = 1 + \frac{2}{N} \quad (6.4)$$

Most appreciable atmospheres in the solar system are dominated by linear diatomic gases, such as H_2 (Jupiter, Saturn, Uranus, Neptune), N_2 (Titan), or an $\text{N}_2\text{-O}_2$ mixture (Earth). These molecules have three translational and two rotational degrees of freedom, so that $N = 5$. Thus, $\gamma = 7/5 = 1.4$ is the same for these worlds. In the case of Venus and Mars, whose atmospheres are dominated by CO_2 , there are 3 translational, 2 rotational, and 2 vibrational degrees of freedom, so $N = 7$ and $\gamma = 1.3$. For H_2 -dominated brown dwarfs, we therefore take $\gamma = 1.4$.

Because of latent heat release due to condensing volatiles during the convection process, planetary atmospheres do not follow a dry adiabat. Thus, we modify the temperature structure in the convection-dominated region of a planetary atmosphere as

$$T = T_0 \left(\frac{p}{p_0} \right)^{\alpha(\gamma-1)/\gamma} \quad (6.5)$$

Here, α is a factor, typically around 0.6-0.9, which accounts for the deviations from the dry adiabatic lapse rate, primarily because of latent heat release. Physically, α represents the average ratio of the true lapse rate in the planet's convective region to the dry adiabatic lapse rate. While this parameterised representation of the temperature-pressure profile works well for all of the worlds of the Solar System with thick atmospheres, it breaks down for very moist atmospheres (Pierrehumbert, 2010), and for one-component condensable atmospheres, which will have temperature varying as $1/\ln(p/p_0)$ according to the Clausius-Clapeyron relation. We set our α prior bounds to be uniform between 0.5 and 2.0. While this would mean that α would physically produce a superadiabatic lapse rate for values larger than one, which is unstable over long time periods, we can

also think of this as allowing more flexibility in what is really a parameterisation (albeit a semi-physically motivated one).

Radiative Zone

Above the convection-dominated zone of the planet’s atmosphere, the temperature profile tends towards radiative equilibrium. The temperature at which the profile converges to is related to the internal energy source from the planet’s interior. In this scheme, the profile is then given as

$$T^4(\tau) = \frac{T_{\text{eff}}^4}{2}(1 + D\tau) \quad (6.6)$$

where D is the diffusivity factor, which accounts for the integration of the radiances over a hemisphere, and is often taken as 1.5, 1.66, or 2 (Rodgers & Walshaw, 1966; Weaver & Ramanathan, 1995). We use 1.5, though in actual fact this will be absorbed into τ during the retrieval anyway.

One can see quite clearly that as τ tends to zero towards the top of the atmosphere, the isothermal temperature is given by:

$$T = \frac{T_{\text{eff}}}{2^{1/4}} \quad (6.7)$$

We can see this is true if we consider a very thin layer of opacity η high in the atmosphere with negligible atmospheric opacity above it. The atmosphere below this layer effectively emits as a black body of temperature equal to the bolometric temperature T_B . Hence the heating rate of the thin layer per unit area is ηT_B . However, heat from this layer may be emitted upwards and downwards and thus in equilibrium, assuming Kirchoff’s law, we find $\eta\sigma T_B^4 = 2\eta\sigma T^4$. Rearranging, we arrive at Equation 6.7.

The Full Temperature Profile

To combine these two regimes into one single temperature profile, we calculate the radiative profile and the convective profile over a specified pressure range. The point at which the radiative and convective profiles cross, i.e. have the same temperature, is the radiative-convective boundary. At this boundary, we take 5 atmospheric layers above

(radiative) and below (convective), and smooth them using a piecewise cubic Hermite interpolating polynomial (PCHIP). This is illustrated in Figure 6.2.1.

6.2.2 Cloud Profile

A wide range of cloud condensates in Y dwarfs were investigated by Morley et al. (2012). These findings can be summarised by Figure 6.2.2. Essentially, by far the most important (optically thick) cloud that would be visible in Y dwarf atmospheres would be Na₂S. The maximum predicted optical depths of Na₂S clouds Morley et al. (2012) can be seen more clearly in Figure 6.2.3, where the effect of surface gravity and temperature is more explicit. The predicted maximum optical depth should be large enough to produce an observable change in the spectrum.

Sodium sulphide condenses via the net thermochemical reaction $\text{H}_2\text{S} + 2\text{Na} = \text{Na}_2\text{S(s)} + \text{H}_2$. The mass of the Na₂S cloud is limited by the elemental abundance of sodium, which is 13% of the abundance of sulphur in a solar-composition gas. From Visscher et al. (2006), the predicted saturation vapour pressure of Na p'_{Na} above the cloud base is given by

$$\log p'_{\text{Na}} \approx 8.550 - 13889/T - 0.50[\text{Fe}/\text{H}] \quad (6.8)$$

where the metallicity dependence ($[\text{Fe}/\text{H}]$) results from H₂S remaining in the gas phase above the Na₂S cloud and from the stoichiometry (relative quantities of reactants and products in a chemical reaction) of Na and H₂S in the condensation reaction. The condensation temperature T_{cond} is given by Visscher et al. (2006):

$$10^4/T_{\text{cond}}(\text{Na}_2\text{S}) \approx 10.045 - 0.72 \log p_t - 1.08[\text{Fe}/\text{H}] \quad (6.9)$$

with total atmospheric pressure p_t , indicating condensation near ~ 1000 K at 1 bar in a solar-metallicity gas. As there is only a weak dependence on metallicity for these equations, and it is not possible, or inaccurate, to calculate the metallicity during the retrieval, we consider the metallicity to be solar ($[\text{Fe}/\text{H}] = 0$) for Equations 6.8 and 6.9 (though not elsewhere in our calculations). We use Equation 6.9 to locate the cloud base and Equation 6.8 to modify the Na profile above this cloud base.

Now that we know the composition, condensation point, and saturated vapour pressure




Figure 6.2.1: Illustration of the radiative (blue), convective (red), and combined (black dotted) temperature profiles. The combined profile is convective until approximately one bar, where the radiative-convective boundary occurs. The rest of the profile is then radiative above this point.




Figure 6.2.2: Predicted optical depths for a range of clouds, for a variety of temperatures and surface gravities in brown dwarf atmospheres. Shaded grey area is the 1-6 μm photosphere. Figure from Morley et al. (2012).

of our cloud, we must decide its physical structure. We choose a Mie-scattering cloud with fixed imaginary (and hence real) refractive index according to Morley et al. (2012), a retrieved particle size with retrieved particle variance for a log-normal distribution, fixed cloud base pressure according to the saturation vapour pressure, integrated optical depth, and fractional scale height.




Figure 6.2.3: Predicted maximum optical depths for a range of surface gravities and effective temperatures for the Na_2S cloud in brown dwarf atmospheres. Figure from Morley et al. (2012).

6.3 Benchmark Retrieval: GL570D

GL570D is a typical T7.5 dwarf which has been well studied (Line et al., 2015) and has an accurately measured parallax (distance = 5.84 ± 0.03 pc, van Leeuwen (2007)), meaning it is possible to break the degeneracy between radius and distance. We therefore use this object as our benchmark brown dwarf to verify our methods. Our spectrum is from the SpeX database (Burgasser, 2014). It is a low-resolution ($R \sim 250$) spectrum between 1.0-2.4 microns, obtained using the low-resolution prism-dispersed mode of SpeX on the 3m NASA Infrared Telescope Facility at Mauna Kea, Hawaii, which has been previously analysed by Line et al. (2015). In fact, we are using exactly the same dataset, provided to us by the author.

In this benchmark, we retrieve seven gases: $\text{NH}_3, \text{CH}_4, \text{H}_2\text{O}, \text{CO}, \text{CO}_2, \text{Na}, \text{K}$ with uniform prior log VMR between -8.0 and -1.0 (where the rest of the remaining gas is filled with a H_2 :He mixture of ratio 85:15), uniform mass prior ($1-80 M_{Jup}$), uniform radius prior ($0.6-2.0 R_{Jup}$), uniform reference temperature prior T_0 (1000-7000K), uniform α prior (0.5-2.0), uniform T_{eff} prior (200-1400K), uniform $\log \tau_0$ prior (0-7), and uniform n prior (1-5). We do not retrieve clouds in this benchmark, so that we can directly compare to the findings of Line et al. (2015).

6.3.1 Results

Using the methods described in Section 6.2, we present the retrieved posterior distributions for the spectral fit (Figure 6.3.1), the temperature profile (Figure 6.3.2), and the posterior distributions for each parameter (Figure 6.3.3). We see that strong constraints are provided on most parameters due to a good fit to the high SNR data.

In order to put our results into context, we directly compare them to the results of Line et al. (2015). Our results are shown as the top panel of Figure 6.3.4, while Line et al. (2015) is shown in the bottom panel. The middle panel will be explained briefly. From the top and bottom panels, we can see a number of similarities as well as differences.

One main difference is that, except for Na, our error bars are significantly smaller than that of Line et al. (2015). This is due to two main factors: the different temperature profile parameterisation, and to a lesser extent, the retrieval method used. In Line




Figure 6.3.1: GL570D. Median fit in black, observed spectrum of GL570D in orange with error bars included. Technically there are also ± 1 and 2σ error bars on the median fit, however it is so small it can't be detected.

et al. (2015) they use a layer-by-layer method, with 15 free parameters for the temperature profile, and the retrieval algorithm is the MCMC code, EMCEE. Their temperature profile parameterisation retrieves the temperature at 15 specified pressures directly, and interpolates between them to create the 70 layers used in their model, while using the rate of change of temperature with respect to log-pressure to prefer straight (linear in log-pressure-temperature space) lines as a prior. We confirm that the change in errors is due to these two main differences in the middle panel in Figure 6.3.4, where we have followed the method in Line et al. (2015) exactly (i.e. same temperature profile parameterisation, and same retrieval algorithm). Note that we only perform this comparison for the benchmark case, as it is extremely computationally expensive. In the middle panel of Figure 6.3.4, we retrieve almost exactly the same posterior distributions as Line et al. (2015), with the exception of Na and K as they are retrieved separately in our case, rather than an assumed solar ratio as in Line et al. (2015). We also compare the temperature profiles directly, and show that we have very overlapping temperature profiles for the layer-by-layer case. The region between 0.2 to 2 bar is quite well matched between the two different parameterisations, however the rest of the profile is quite different. This is likely due to the fact that most of the information in the spectrum peaks in this region.




Figure 6.3.2: GL570D. Summarised temperature profiles calculated from randomly drawing 1000 temperature profiles from the posterior of the three models described in Figure 6.3.4. Blue is our Nested Sampling retrieval, green is our MCMC retrieval, and red is the retrieval from Line et al. (2015). Shaded areas are $\pm 1\sigma$. A lot is overlapping in the two MCMC retrievals, so some transparency is used to mix the colours.

The ‘loosening’ of posteriors is due to the large increase in the number of free parameters that describe the temperature profile (we use 5 parameters while Line et al. (2015) use 15, and both use their parameters to create the profile in a different way). As the number of degrees of freedom increases, the accuracy to which we can know any particular degree of freedom decreases. As a simple analogy, consider Heisenberg’s uncertainty principle: the more we know about the position (temperature profile), the less we can know about the momentum (volume mixing ratios) of the particle. While the Line et al. (2015) approach allows flexibility, we (from private communication with Mike Line and our own experience) find that it also comes at a large computational cost, and areas of low temperature information tend towards a straight line prior. This can be seen in Figure 6.3.5, where a completely straight line is visible from 1 to 300 bar for both objects, which is unphysical. Hence, we will stick to the Robinson & Catling (2012) parameterisation, which imposes at least some physical motivation.

Now that we have verified our methods with the literature, we may now analyse our Y dwarf data set.




Figure 6.3.3: GL570D. Summary of the marginalised posteriors for the retrieved parameters. This shows the marginalised posteriors (1D histograms) along the diagonals and the parameter correlations (2D histograms; contours are shown at 0.5, 1, 1.5, and 2 σ) at off-diagonal locations. Median and $\pm\sigma$ values plotted on top of 1D distributions as dotted lines.




Figure 6.3.4: G1570D. Top: Summarised VMR profiles calculated from the marginalised posteriors using the methods described in this thesis. Dark shading is $\pm 1\sigma$, and light shading is $\pm 2\sigma$. Middle: Same as Top but using the same method as Line et al. (2015). Bottom: Retrieved results from Line et al. (2015). Solid or dotted regions are comparisons to thermal equilibrium chemical models.




Figure 6.3.5: Retrieved temperature profiles from Line et al. (2015). This summarises thousands of temperature profiles drawn from the posteriors for GL570D (left) and HD 3651B (right). Posteriors are summarised with median in blue, 1σ spread in red, 2σ spread in pink. The black temperature profile shown for each object is a representative self-consistent grid model Saumon & Marley (2008) which most closely represents the retrieved temperature profiles. A completely straight line is visible from 1 to 300 bar for both objects, which is unphysical.

6.4 Y Dwarf Retrievals

Having tested and verified our retrieval method with GL570D, we now apply the same technique as described before, where we retrieve five additional parameters: the distance to the object, and the cloud parameters (particle radius, particle variance, cloud fractional scale height, and integrated optical depth at 1 micron). The cloud base is set by the condensation temperature in Equation 6.9.

The distance prior is the only non-uniform prior that we use in the retrievals. We use multiple parallax measurements as summarised in Smart et al. (2017) to form multi-modal Gaussian priors. We show one of these multi-modal Gaussian distance priors in Figure 6.4.1 for WISE1738. This enables us to break the degeneracy between distance and radius, and consequentially, to some extent, the spectroscopic mass of the object (decreased uncertainty in radius leads to decreased uncertainty in the surface gravity, and hence the mass).




Figure 6.4.1: Each histogram represents the probability distribution of each parallax measurement (converted to distance in parsecs here) for WISE1738 from Smart et al. (2017). It is clear from the diagram that while most overlap, there is some discrepancy between them. The combination of all these histograms form the multi-modal Gaussian distance prior.

For clarity, we present each retrieved parameter and its associated prior in Table 6.4.1. The

Table 6.4.1: Summary of priors.

Parameter	Prior Type	Prior Bounds
Radius	Uniform	0.5, 2.0 R_{Jup}
Mass	Uniform	1, 80 M_{Jup}
logVMR	Uniform	-8, -1
T_0	Uniform	1000, 7000K
α	Uniform	0.5, 2.0
T_{eff}	Uniform	200, 1400K
log τ_0	Uniform	0, 5
n	Uniform	1, 5
Distance	Multimodal Gaussian	Smart et al. (2017)
log FSH	Uniform	-3, 0
log $Opt.$	Uniform	-2, 4
log $Part_{rad}$	Uniform	-1, 2
log $Part_{var}$	Uniform	-2, 2

gases chosen for these retrievals were decided upon by a mixture of retrievals in the literature (i.e. Line et al. (2015)), from our own exploratory retrievals, and from theoretical works (i.e. Morley et al. (2012)).

In Figure 6.4.2, we show normalised (for clarity) fits to the spectra of the seven Y dwarfs that we analysed, as well as our benchmark object GL570D for reference. It is immediately obvious from this that our fits to the data are quite good, with only some discrepancy in WISE0410 and WISE2056. This is due to problems with the data themselves. WISE0410 and WISE2056 have issues in multiple areas of the spectrum and these wavelengths are discarded in Cushing et al. (2011). Thus, we would not expect to fit these problem areas. We keep these problem areas as a test to see if our model is physically constraining enough to prevent fitting unphysical datapoints. Looking at our fits we find quite a positive result - it does not contain the freedom (in the temperature profile) to fit unphysical bumps that more flexible models might erroneously allow.

First we will briefly show a set of typical individual results of WISE1738, so that the reader may understand the statistics gained from an individual retrieval. We will then discuss all objects simultaneously, both for brevity, and for ease of comparison. Note that almost all objects showed no preference for clouds and produce very similar retrieved posteriors. The exception is for WISE1828, our coldest object, which prefers a cloudy atmosphere that is considerably different to the others examined.

Figure 6.4.3 shows a typical summary of the parameter posterior distributions, from




Figure 6.4.2: Summary of spectral fits drawn from 100 samples of the posterior of each T/Y dwarf object and the benchmark object GL570D for comparison. Red is the median, shaded dark blue is $\pm 1\sigma$, and shaded light blue is $\pm 2\sigma$. Data are in black dots. The flux has been normalised at 1.26 microns, and shifted for each object for clarity. Data are annotated with their respective names. Order is dictated by the ranking described later in the chapter.

WISE1738. The posteriors of the rest of the objects may be found in the Appendix. Figure 6.4.4 shows the summarised temperature profile posterior. As we will see, these posteriors are not very dissimilar to those retrieved for our benchmark case GL570D - the main difference is that the larger errors and reduced wavelength range have broadened the distributions. This seems reasonable, as there should only be a temperature change between these objects of approximately 200-300K.

6.4.1 Temperature and Spectral Typing

As one of the primary concerns of this thesis is to examine how the atmospheres of brown dwarfs evolve with decreasing temperature, it would be useful to know the rank of our objects from hottest to coldest. This task is not as straightforward as it seems, as ‘temperature’, or rather the effective (blackbody) temperature, is dependent on the composition of the atmosphere and the temperature profile itself. Our calculated effective temperature is therefore model-dependent. The calculated effective temperature T_{eff} is defined by:

$$L = 4\pi R^2 \sigma T_{\text{eff}}^4 \quad (6.10)$$

where L is the luminosity of the brown dwarf (W), R is the radius (m), σ is the Stefan-Boltzmann constant ($\text{Wm}^{-2} \text{K}^{-4}$).

Fortunately, due to the physical basis of our temperature profile parameterisation, we have effectively two methods of measuring the effective temperature. That is, the retrieved effective temperature of the temperature profile parameterisation (i.e. from the isothermal temperature reached due to radiative equilibrium, Equation 6.7), and the effective temperature calculated from the total luminosity of the brown dwarf (Equation 6.10), which must assume an atmospheric solution and therefore can only be calculated after the retrieval by extrapolating the synthetic radiance to all wavelengths covered by the k -tables. Because we have not enforced (true/non-parameterised) radiative-convective equilibrium nor included all chemical species in our retrieval (as most would not be detected), we do not expect these two measures to be exactly the same, though we do expect them to be somewhat correlated as the calculated effective temperature depends on the




Figure 6.4.3: WISE1738. Summary of the marginalised posteriors for the retrieved parameters. This shows the marginalised posteriors (1D histograms) along the diagonals and the parameter correlations (2D histograms; contours are shown at 0.5, 1, 1.5, and 2 σ). at off-diagonal locations. Median and $\pm\sigma$ values plotted on top of 1D distributions as dotted lines.




Figure 6.4.4: WISE1738. Summarised temperature profile calculated from randomly drawing 1000 temperature profiles from the posterior. Blue is the median, shaded dark red is $\pm 1\sigma$, and shaded light red is $\pm 2\sigma$.

retrieved temperature profile.

The difficulty of ordering by temperature is illustrated in Figure 6.4.5. When we compare objects ordered by the median of the calculated (blackbody) T_{eff} (left) to those ordered by the median of the retrieved (radiative equilibrium) T_{eff} (right), we see that objects change positions relative to each other. To combine the information provided into one sensible ordering, we must weight each of these orderings against their merits and pitfalls, but also against other contextual information (such as previous analyses).

The calculated effective temperature suffers from very large extrema, especially in the upper ranges. One major disadvantage of using this scheme to determine temperature, is that only a very small wavelength range (in most cases $1.1\text{-}1.7\mu\text{m}$) has been used to determine the structure of the entire atmosphere. This means a number of gases are missing from the calculation, which involves calculating spectra between $0.3\text{-}30\mu\text{m}$ (the limits of our k -tables). Hence, as a number of absorbers are not present, the amount of luminosity, and hence effective temperature, is over-represented. Note that other groups extend below $0.3\mu\text{m}$ and above $30\mu\text{m}$ for their effective temperature calculations, however from preliminary tests, we found that calculations between $0.3\text{-}30\mu\text{m}$ are, at most,

5K different to the ‘true’ effective temperature using all wavelength space. Considering other uncertainties, this difference is negligible. It is also a possibility that our model is inaccurate in other ways, which would again affect the calculated effective temperature non-trivially.

The accuracy of the retrieved effective temperature depends on how much temperature-dependent information is available in the data, as well as how valid our parameterisation is, and how well we retrieve other parameters.

While both of these measures have their own pit falls, each of them do provide unique information that guides us towards a unified measure of temperature. Looking at Figure 6.4.5, it seems as though that ordering by retrieved temperature also gives a reasonable ordering of the calculated effective temperature. As the rest of this thesis gives weight to a Bayesian sense of inference, we largely agree with using this ordering. However, there are a few issues to this ordering. One is the ranking of GL570D and WISE0148. GL570D is ranked as the third-coldest object, when quite clearly it should be the hottest (as in the calculated T_{eff} ordering) due to its spectral typing, T7.5. We therefore amend this ordering such that GL570D is the hottest. Through similar reasoning, WISE0148 is a T9.5 dwarf and should theoretically be hotter than the rest of the Y dwarfs. It is not immediately obvious from these graphs, but trends in later figures suggest a later spectral typing of WISE1548 than all Y dwarfs except WISE1828. WISE1548 was suggested by Leggett et al. (2017) as a Y1, so this makes sense. Cushing et al. (2011) also suggests that these two objects are the coldest. The remaining objects are ordered by retrieved T_{eff} . This is the ordering that we will use for the rest of the chapter, which is shown in Figure 6.4.6, and represents loosely a spectral typing ordering. Note that we couldn’t do a ranking like this using only spectral types prior to this analysis, because a number of the objects have exactly the same spectral type. This discussion illustrates clearly that spectral typing is not implicitly related to the measured temperature, and instead should only be considered a classification scheme based on the spectra themselves. A stronger correlation to the spectral types may be related to the relative abundance of the different measured gases, as is clear in Figure 6.4.8.

A fuller description of temperature is given by the entire retrieved temperature profile. In Figure 6.4.7 we present the summarised posteriors of the temperature profiles. We can




Figure 6.4.5: Comparison of ordering analysed brown dwarfs by retrieved effective temperature and calculated effective temperature using violin plots (kernel density estimators on 1D histograms). Red are calculated T_{eff} posteriors, blue/green are retrieved T_{eff} posteriors. Medians are represented by white dots, thick bars represent $\pm 1\sigma$, end points are the extreme values of the posteriors. Left: Objects are ordered by median values of calculated T_{eff} ; right: ordered by median values of retrieved T_{eff} .

see that the temperature profiles follow roughly three groupings. GL570D and WISE0148 represent hot objects with roughly similar temperature profiles. WISE2056, WISE1405 and WISE0410 represent cooler Y dwarf objects, with very similar temperature structures. These first two groupings make sense in terms of temperature and spectral typing. Lastly, however, WISE1541, WISE1738 and WISE1828 are all found to have hotter profiles than the other Y dwarfs, which is unexpected as they are later/colder spectral types. We suspect that this is due to the very limited datasets for these objects, a wider wavelength range might solve this problem.

6.4.2 Gaseous Composition

Now that we have established the temperature/spectral typing of our objects, we look to the second most important factor in our retrievals: gaseous composition. First we look




Figure 6.4.6: Final order of retrieved effective temperature and calculated effective temperature using violin plots (kernel density estimators on 1D histograms). Red are calculated T_{eff} posteriors, blue/green are retrieved T_{eff} posteriors. Medians are represented by white dots, thick bars represent $\pm 1\sigma$, end points are the extreme values of the posteriors.

at the three most important gases (NH_3 , CH_4 , H_2O) in Figure 6.4.8.

Here we find some very interesting and clear trends as we transition through the spectral types. NH_3 content increases in an almost log-linear fashion, over two orders of magnitude. A similar trend is seen for CH_4 , with an increase of about one order of magnitude. H_2O is the most complicated trend. It seems to increase at first, then stay roughly constant, before declining sharply at the lowest temperatures. This result would be consistent with H_2O forming clouds, condensing out the observable H_2O . While this may seem like a conflict with regards to our choice of Na_2S clouds, it is not necessarily. Figure 6.4.9 shows that Na_2S is the dominant cloud absorber at 400K. So while water may be condensing out, we may still be observing Na_2S as the main cloud deck. On the flipside, as we don't have a strong temperature constraint on our coldest object, it certainly could be in the 300K region and the detected cloud could really be made of H_2O . A water cloud would also explain well the lack of H_2O retrieved in the spectra. Leggett




Figure 6.4.7: Summarised temperature profiles calculated from randomly drawing 1000 temperature profiles from the posteriors. Thick opaque lines are the medians, shaded regions are $\pm 1\sigma$. Y dwarfs are labelled in order of ranking defined in text.

et al. (2017) suggest this object is 310-340K, so this is certainly possible.

Figure 6.4.10 presents the 2D correlations between the various gases. As expected from Figure 6.4.8, the log-linear positive correlation between NH_3 and CH_4 is clear (top plot of Figure 6.4.10). The 2D correlations involving water (middle and bottom of Figure 6.4.10) are more complex, but the plots illustrate well a very clear divide between WISE1828 and the rest of the objects. WISE1541 appears to possibly be in a state between WISE1828 and the other Y dwarfs.

Figure 6.4.11 shows that we do not detect Na in any of our objects. In addition, we only strongly detect K in our two hottest objects. These results support our narrative of condensing Na_2S and KCl clouds as we move to colder temperatures. Although, the apparent lack of these molecules could be due to a lack of data, rather than the physical absence of the molecules themselves.




Figure 6.4.8: Ordered VMRs of NH_3 , CH_4 , and H_2O using violin plots (kernel density estimators on 1D histograms). Medians are represented by white dots, thick bars represent $\pm 1\sigma$, end points are the extreme values of the posteriors.

6.4.3 Mass, Radius, Surface Gravity

The mass, radius, and hence the surface gravity, of the brown dwarf is mostly reliant on measurements from sources that are not spectroscopic. The radius is almost completely degenerate with the distance of the object. The distances of these objects are calculated from parallax measurements which can be inaccurate (large error bars) or the same object may even have multiple non-overlapping measurements. We have not restricted ourselves to a single parallax measurement, and instead provide multi-modal Gaussian priors in the retrieval conditioned to the multiple astrometric measurements quoted in Smart et al. (2017), as discussed earlier.

These parallax measurements give us some information about the radius of the object. Thanks to the effect that the gravity has on the atmospheric scale height, we can also measure the spectroscopic mass if we have some information on the radius. In Figure 6.4.12, we present the retrieved posteriors for the mass and radius, and the calculated posterior of the surface gravity when combining these.

Figure 6.4.12 shows that we have reasonable estimates of the surface gravity for GL570D, WISE0148, WISE2056, WISE1405, and a slight estimate for WISE0410. We can group these into low surface gravity objects (GL570D, WISE0148) and high surface gravity objects (WISE2056, WISE1405, WISE0410). Coincidentally, this is also the ordering of our objects, and therefore it is difficult to disentangle the effects of surface gravity on the other retrieved parameters, as gravity is much more of a secondary effect on resulting spectrum than the first-order temperature effect. A larger sample would be required to elucidate any surface gravity trends.

6.4.4 Clouds

While clouds are predicted to exist in Y dwarfs (Morley et al., 2012, 2014b), their optical depths are not expected to be very large. The largest optical depth is predicted for Na₂S for hotter Y dwarfs ($\sim 400\text{K}$), and water clouds for the colder ($\sim 275\text{K}$) Y dwarfs. In Figure 6.4.13, it is easy to see that, if a Na₂S cloud did exist at the correct condensation level, for all objects the integrated optical depth must be less than about 3 (excluding WISE1828). In WISE2056 the upper limit for the integrated optical depth is as low as 0.1.

Note that we do not attempt to retrieve clouds for GL570D and WISE0410, hence why these rows are blank in Figure 6.4.13.

Figure 6.4.14 shows the posteriors of WISE1828 in more detail. We can see that there is a strong detection of an integrated optical depth of ~ 5 , which is close to that predicted by Morley et al. (2012). It also detects a reasonably large particle variance, with particle radius centred around approximately $0.3\mu\text{m}$. This is perhaps in disagreement with Morley et al. (2012), who predict the particles to be about 100 times larger than this. It could be that, as our model is set up to retrieve Na₂S clouds, we are in fact detecting a H₂O cloud and pushing unphysical parameters (i.e. small particle sizes) on to our Na₂S prescription to fit a H₂O cloud which would need a different prescription (i.e. condensation level and refractive index spectrum). Or the particles truly are this small for a Na₂S cloud, and it is theory that is at fault. Though we should consider the limitations of the dataset before making conclusions that are so drastic.

Another point to consider is that the cloud structure might be affecting our retrieved

temperature profile. WISE1828 retrieved an unusually hot temperature profile (in contradiction to our previous notions of this object being cold), and as the temperature profile and the cloud base pressure are fundamentally linked in our model prescription, it could be that if the cloud base pressure is not where we expect it to form, the temperature profile could be made either hotter to accommodate this change in cloud base pressure by the retrieval.

The sudden depletion in the retrieved water abundance would suggest that a water cloud has indeed formed and that WISE1828 is at around 300K. However we cannot conclusively say that the cloud we are detecting is water as it is possible to fit the Na_2S with our dataset.




Figure 6.4.9: Predicted optical depths for a range of clouds for a variety of temperatures in brown dwarf atmospheres with solar metallicity. Shaded grey area is the 1-6 μm photosphere. The effective (area-weighted) mode radius of the cloud particles at each pressure is shown as dashed line. The geometric column optical depth is shown as solid lines. The quoted temperatures are effective temperatures. Figure from Morley et al. (2014b).




Figure 6.4.10: Scatter plots of 400 posterior samples to represent 2D correlations between NH_3 , H_2O , and CH_4 . Median values for each object are plotted on top in solid colored circles with black bands.




Figure 6.4.11: Ordered retrieved VMRs of Na and K using violin plots (kernel density estimators on 1D histograms). Medians are represented by white dots, thick bars represent $\pm 1\sigma$, end points are the extreme values of the posteriors.




Figure 6.4.12: Ordered retrieved mass, radius, and (calculated) surface gravity using violin plots (kernel density estimators on 1D histograms). Medians are represented by white dots, thick bars represent $\pm 1\sigma$, end points are the extreme values of the posteriors.




Figure 6.4.13: Ordered retrieved (log of all) fractional scale height, integrated optical depth normalised at $1\mu\text{m}$, particle radius, and particle distribution variance using violin plots (kernel density estimators on 1D histograms). Medians are represented by white dots, thick bars represent $\pm 1\sigma$, end points are the extreme values of the posteriors.




Figure 6.4.14: WISE1828. Posteriors of fractional scale height, integrated optical depth normalised at $1\mu\text{m}$, particle radius, and particle distribution variance. Medians and $\pm 1\sigma$ are written in black. Histograms are in red/orange, and a blue kernel density estimator is plotted on top of the histograms to aid the eye.

6.5 Conclusions

Here we have presented eight consistent retrievals on two T dwarfs and six Y dwarfs. We created a forward model with a semi-physical temperature profile and realistic Na_2S clouds, and used Nested Sampling to retrieve the posteriors on a number of atmospheric properties.

We showed two different ways of measuring the effective temperature: the retrieved stratospheric temperature, and the integrated luminosity. These two different measurements often produced different results because the first relies on the information content of the spectrum, and the second relies on knowing a perfect state of the atmosphere (which is not possible with a limited wavelength range, and our model only integrates between $0.3\text{-}30\mu\text{m}$). We argued for an ordering of these objects which included information based on the stratospheric temperature and the spectral typing, to create an ordering loosely resembling a new spectral typing of all the objects.

It is immediately apparent how useful this ordering is when comparing the posteriors of the different objects, especially when considering the volume mixing ratios of NH_3 , CH_4 , and H_2O . When examining the posteriors of these gases, one can see a clear log-linear relationship for NH_3 and CH_4 content, correlated with the ordered spectral typing/temperature. A clear trend is also visible for H_2O , where the abundance slightly increases before dramatically disappearing in the coldest object. The 2D correlations between NH_3 and CH_4 show a very strong positive correlation both in terms of the retrieval and in terms of the object ordering/temperature. 2D correlations of H_2O and the other gases are a little more complex, but show that WISE1828 is clearly an object that is vastly different to the others. WISE1541 seems to lie in-between WISE1828 and the rest of the objects.

Examining Na and K showed that only the hottest objects contained K, which fits our narrative of Na and K condensing into Na_2S and KCl (which condenses at a lower temperature than Na_2S , approximately 600K).

While we do have measurements for the mass, radius, and hence the surface gravity, we only have strong measurements of gravity for 4-5 objects. An unfortunate coincidence lines these gravities up in our regular temperature ordering, so in this small sample it is

almost impossible to disentangle the secondary effect of gravity to the primary effect of temperature.

Finally, we examined the posteriors of the cloud parameters. Our retrieved (or lack of) clouds gave an upper limit of the integrated optical depth of about 1 for all objects except WISE1828. This object is best fit with a cloud with an integrated optical depth of about 5, particle radius of $0.3\mu\text{m}$, and a particle size variance of 0.5, assuming a Na_2S -like complex refractive index spectrum from Morley et al. (2012). While the integrated optical depth is about the same as predicted for theory, the particle radius is considerably smaller. We also discuss the possibility of this object being as cold as 275K, where water clouds would be forming, and this would be the main cloud deck instead of Na_2S . The apparent sudden depletion in water abundance in this object certainly lends support to this idea. We also note that the hotter temperature profile of WISE1828 could be due to the fundamental link between cloud base pressure and temperature - a hotter profile could allow for a change in cloud base pressure which is overall advantageous for the model fit to the spectrum.

Chapter 7

Conclusions

In this Chapter we will summarise the findings presented throughout this thesis, and discuss the possibility of future work which would enhance and build upon the current state of our findings. We also discuss some unfinished work which is complementary to this thesis, but did not warrant chapters of their own.

7.1 Summary

Brown dwarfs span an enormous range of atmospheric conditions: from the super-hot M dwarfs raining iron clouds, to the ultracool Y dwarfs forming water clouds. They prove to be an invaluable resource for atmospheric characterisation beyond our Solar System with their relatively high quality spectroscopic data and simple (i.e. non-irradiated) atmospheres. In this thesis, we have explored the atmospheric conditions of a range of ultracool T and Y dwarfs using state-of-the-art physical models and Bayesian methods developed for this purpose, uncovering a number of new results that will aid our future understanding of these extremely cold extrasolar objects.

This thesis first tells the narrative of the the life of a brown dwarf, and the distinguishing spectral features which go hand-in-hand with the evolving atmospheric conditions of the catalogue of objects that are known to exist (Chapter 1). In particular, we discuss the possible atmospheric and spectral features of a newly discovered class of objects, the Y dwarfs, which are the main focus of this thesis. Alongside this history, we present a foundation of the required knowledge to help understand these objects and the methods used to analyse them, including the physical processes involved in radiative transfer (Chapter 2) and statistical methods for solving the inverse problem (Chapter 3).

We proceeded by updating our model with the latest molecular and atomic modelling data (Chapter 4. We discussed in depth how these molecular/atomic data affects the calculated spectrum of a brown dwarf. As well as this, we investigated the effects of using two different gaseous absorption methods (correlated- k and cross-sections) to calculate spectra in a variety of atmospheres of varying complexity. We also investigated the effects of including H₂-He pressure broadening in the more complex atmospheres. These investigations are important because radiative transfer models are often coupled to inverse methods that can iterate over millions of forward models in a large parameter space; a single forward model error, such as a ‘ghost feature’, would propagate through to all retrieved PDFs influencing them in a non-trivial way, i.e. an error in the absorption spectrum of NH₃ could influence the retrieved temperature profile. We first showed that for test cases with resolutions of $\Delta\nu = 1\text{cm}^{-1}$ the cross-section method overestimates the amount of absorption present in the atmosphere and should therefore be used with

caution. The morphology of the spectra also changes and produces ‘ghost’ features for mixed-gas atmospheres. When we considered cross-sections, the effect produced multiple-orders-of-magnitude change in the flux received from brown dwarfs in certain wavelength regions. The effect was similar but smaller for primary transit, and is closer in order to a $\sim 1\%$ change in transit depth. The flux ratio of secondary eclipse exoplanets produced an order of magnitude change in the near-IR, with the effect becoming lesser for longer wavelengths. Correlated- k produced similar results to line-by-line and very high-resolution cross-sections, but is much less computationally expensive. The inclusion of H_2 -He pressure broadening similarly changed the total flux calculated for brown dwarfs and secondary transit exoplanets by up to an order of magnitude, while only making slight changes to the spectra of transiting exoplanets. We concluded that inaccurate use of cross-sections and omission of pressure broadening can be key sources of error in the modelling of brown dwarf and exoplanet atmospheres. These sources of error in the forward model may produce strong biases in the probability distribution functions of retrieved parameters.

Having updated the model to include the latest molecular and atomic data, we then set out to verify our forward model (Chapter 5). We compared three different radiative transfer models to investigate the differences between their forward models, and their influence on the retrieved posterior probability distributions using uniform Bayesian inference methods for typical resolutions of JWST at various levels of signal-to-noise for a range of planets. We found that while different forward models produced varying different spectra, the vast majority of retrieved parameters are model-independent (i.e. unimodal) for errors of 60PPM and more. In general, 30PPM led to much more multi-modal results for, in particular, certain gases (due to differences in line lists), and the radius. Temperature and cloud-top pressure were almost always model-independent. By combining the posterior samples into single combined posterior distributions, we, for the first time, estimated the uncertainty in Bayesian inference due to differences in forward models. While using multiple models is the optimal method for retrievals, realistically for spectra with errors of 60PPM and more, any one model should retrieve similar posteriors to any other. This is of course limited to the models, resolution, and particular atmospheres that we tested in this study.

After having verified our model and discussed the effects of different models on the posterior distributions of various parameters, we were then in the position to confidently model the range of T and Y dwarf objects that we had received data for, to elucidate trends in atmospheric composition and structure over decreasing temperature (Chapter 6). We presented eight consistent retrievals on two T dwarfs and six Y dwarfs. We created a forward model with a semi-physical temperature profile and realistic Na_2S clouds, and used Nested Sampling to retrieve the posteriors on a number of atmospheric properties. We showed two different ways of measuring the effective temperature: the retrieved stratospheric temperature, and the integrated luminosity. These two different measurements often produced different results because the first relies on the information content of the spectrum, and the second relies on knowing a perfect state of the atmosphere (which is not possible with a limited wavelength range, and our model only integrates between $0.3\text{-}30\mu\text{m}$). We argued for an ordering of these objects which included information based on the stratospheric temperature and the spectral typing, to create an ordering loosely resembling a new spectral typing of all the objects. It became immediately apparent how useful this ordering is when comparing the posteriors of the different objects, especially when considering the volume mixing ratios of NH_3 , CH_4 , and H_2O . When examining the posteriors of these gases, one can see a clear log-linear relationship for NH_3 and CH_4 content, correlated with the ordered spectral typing/temperature. A clear trend is also visible for H_2O , where the content slightly increases before dramatically disappearing in the coldest object. Examining Na and K showed that only the hottest objects contained K, which fitted our narrative of Na and K condensing into Na_2S and KCl . Due to having strong constraints on the surface gravities of only five objects, we could not construct secondary effects / trends due to gravity. Finally, we examined the posteriors of the cloud parameters. Our retrieved (or lack of) clouds gave an upper limit of the integrated optical depth of about 1 for all objects except WISE1828. This object is best fit with a cloud with an integrated optical depth (at $1\mu\text{m}$) of about 5, particle radius of $0.3\mu\text{m}$, and a particle size variance of 0.5. While the integrated optical depth is about the same as predicted by theory (Morley et al., 2012), the particle radius is considerably smaller. We also discuss the possibility of this object being as cold as 275K , where water clouds would condense, likely forming the main cloud deck instead of Na_2S . The sudden depletion in

water abundance in this object certainly lends support to this idea. This is also the first Bayesian detection of a cloud on a Y dwarf atmosphere.

We hope that the analysis of these objects and the work presented here will help elucidate the atmospheric structure and composition of the Y dwarf population, as this thesis has always intended, and that this work will assist in future analyses of Y dwarf atmospheres in particular, and more generally any other atmosphere within or beyond our Solar System.

7.2 Future Work

The methods developed throughout this thesis have been constructed such that they lend themselves to analysing a range of objects and geometries - they are not brown dwarf-specific. Here we discuss other possible applications, with preliminary work, as well as ways to improve efficiency.

7.2.1 Efficiently Solving the RTE

Monte Carlo sampling methods such as MCMC and NS require millions to billions of iterations of the forward model to correctly evaluate the likelihood function over a wide range of parameter space, taking months of calculations even on powerful clusters. When examining a range of objects, as in this thesis, the amount of computation time becomes almost unfeasible within a reasonable timeframe. In light of this, we investigated a number of methods to increase the efficiency of the forward model.

The single biggest bottleneck in solving the radiative transfer equation, as alluded to in Chapter 4, is combining the opacity due to different gases. This is because combining transmissions of different gases is only possible if they are calculated monochromatically. In order to speed up this calculation, usually a pre-calculated opacity table of some sort is used; however, combining these pre-calculated tables is still a bottleneck, meaning a typical calculation for 5-7 gases takes 30s. A secondary bottleneck, specific to the NEMESIS code, is the read-write processes that unnecessarily take place during each iteration, which could also be improved upon in future work.

As individual gas pre-calculated opacity tables for the correlated- k method in NEMESIS is acting as a bottleneck, we investigated a number of alternative methods of speeding up this calculation. First we will discuss some of the methods considered, before presenting preliminary results on our model of choice.

Regular Grid Approach

The first method is a brute-force grid approach. That is, we calculate a regularly spaced grid of k -coefficients over pressure, temperature, and the volume mixing ratios of all concerned gases, and perform interpolations between the grid-points. First let's assume

we have 5 gases, then in total we will have 6 dimensions (the volume mixing ratios can be represented as $N - 1$ if we assume all mixing ratios are relative to the mixing ratio of water). Second, let's assume that each of these dimensions (pressure, temperature, etc) requires 20 sampling points to sufficiently resolve the non-linear behavior (in all likelihood 20 is not enough, but we will use this for the sake of argument). That means we need $20^6 = 64$ million data samples. Each of those samples has 20 g-ordinates and, at least, a few hundred, if not tens of thousands, of spectral data points. One of these grids takes up a few hundred GB of data, and adding a single gas would take it into the region of TBs. These data must all be held in memory by the cluster, and so 5 gases will be our limit for the RAM of most easily accessible clusters. It is clear that this method does not scale well to a larger number of gases, but perhaps that is not necessary, as in this thesis. Therefore we will assume 5 gases is sufficient for our calculations.

There are a few drawbacks of this method. One is that the maximum number of feasible gases is approximately 5. This is likely to be true for all further methods discussed, so we will not be too critical of this issue. Another is that it takes up a large amount of memory, and therefore locks up an entire, even quite large, cluster for this single process. Lastly, the number of sampling points for each dimension (20 here), might not sufficiently resolve the non-linearities present. As one major positive, interpolating regular grid structures is extremely quick.

Briefly we present preliminary findings on this model compared to a regular correlated- k calculation. Instead of using 5 gases, we use 3 gases (NH_3 , CH_4 , H_2O) to reduce storage size. We use 391 wavelength points in the near-IR with 20 g-ordinates for each wavelength, 20 regularly-spaced temperature points (300,2950K), 10 log-pressure points (-2.0,2.0), and 20 VMR points for two gas ratios (the log-ratio of $\text{H}_2\text{O}:\text{CH}_4$ [-1.5,0.5] and $\text{H}_2\text{O}:\text{NH}_3$ [-2.5,0.5]) as our regular sized grid. This corresponds to 80,000 samples, or 0.6256 billion data points, creating a binary file of size 17GB. For a sample atmosphere resembling WISE0140 (the atmosphere itself is not important for the speed test), the correlated- k method for calculating opacity takes, on average, 4.15 seconds; the interpolated grid method takes 0.03s. That's an improvement in speed of a factor of 140. Of course, this is just one part of the code; however the rest of the code takes approximately 0.09s. Therefore in total the correlated- k forward model takes 4.24s, while the interpolated

method takes 0.12s. That's 35 times faster than the correlated- k calculation - or in real terms, taking calculations down from a month to one day. In Figure 7.2.1, we can see that the error produced by taking the ratio of the fluxes calculated by the two methods is of the order of 1.5%. While the speed increase and high accuracy are both promising, the reader is reminded of the enormous size of these grids, which make it impractical to use on even large clusters.




Figure 7.2.1: Flux ratio spectrum (WISE0410) calculated by the regular grid method and the exact method (on-the-fly mixing).

Irregular Grid Approach

To combat the issue of the large amount of storage required for the regular grid approach, we can instead be more selective with what we choose in our grid. Based on previous retrievals, we know which probability distributions are most likely to occur for each given gas. Using this, we could construct an irregular grid which sample areas of higher likelihood by sampling a specified multivariate Gaussian with potentially non-diagonal

covariance. This would lead to higher accuracy, assuming that we are covering the same areas of probability distributions, and reduced storage (because we sample unlikely areas less). See Figure 7.2.2 as an example of the parameter samples used. Once the samples are constructed, we use a partitioning function, which splits unstructured data into partitions which are more easily searchable, such as a kd-Tree or Ball Tree, to perform a nearest neighbours search for the relevant sampling point. While some marginal improvement from storage can be gained from this method, it does add much more complexity and hyperparameters to consider (Gaussian means and errors for all parameters, plus covariances), and is slightly slower than the regular grid interpolation (about two times slower). As well as complexity, there is the possibility that retrieved results will be biased by the Gaussians selected for the samples. While an interesting idea, we believe this is, at best, a marginal improvement on the regular grid approach.

Artificial Neural Network Approach

One way to reduce storage is to derive a function that accurately represents the data, and call this function directly when necessary. For example, to represent any linear function we only require two coefficients: the gradient and the intercept. More complicated non-linear functions can be represented by a series of coefficients to describe a polynomial function. Similarly, an artificial neural network may perform this kind of regression on highly non-linear and highly dimensional problems given a sufficient training set.

Artificial neural networks can be viewed as weighted directed graphs in which artificial neurons are nodes and directed edges with weights are connections between neuron outputs and neuron inputs.

The Artificial Neural Network receives some inputs, in our case the temperature, pressure, and volume mixing ratios of the different gases. These inputs are mathematically designated by the notation x_i for n number of inputs.

Each input (temperature, volume mixing ratios, pressure) is multiplied by its corresponding weight. Weights are the information used by the neural network to solve a problem. Typically weight represents the strength of the interconnection between neurons inside the neural network. These weights are typically initially randomly distributed between $[0,1]$. The weighted inputs are all summed up inside the artificial neuron, and a




Figure 7.2.2: 1D and 2D distributions of the multivariate Gaussian that has been sampled for the irregular grid method.

threshold value or ‘bias’ is added which scales the system response.

The sum is then passed through an activation function (threshold unit), which transforms the input into the desired output (opacity). There are a range of both linear and non-linear activation functions. One of the most widely used is the sigmoidal hyperbolic function, usually approximated using the tanh function. This is because of its non-linear properties, and the fact that a small change in a weight (or bias) causes a small change in the output. This is opposed to, say, a binary activation function where a small change could change the output from 0 to 1. Some common activation functions are represented in Figure 7.2.4. This procedure and the concept of the neuron is summarised in Figure 7.2.3.










Figure 7.2.3: Diagram of an artificial neuron. Shows how the inputs (x_i) are transformed by the weights (w_i) and bias (b), before being fed into the activation function ('Threshold unit'), to produce an output. Figure from: <https://www.xenonstack.com/blog/data-science/overview-of-artificial-neural-networks-and-its-applications>, accessed 13th March 2018.

A typical neural network contains a large number of artificial neurons, called units, arranged in a series of layers (see Figure 7.2.5). The input layer contains neurons (units) which receive the input parameters (temperature, pressure, VMRs) on which the network will learn. The output layer contains units which output the desired output (opacity). The hidden layer transforms the input into something that the output layer can use.

Put mathematically, the input vector corresponds to x , the hidden layer activation corresponds to h , and the output activation y . We have a function f that maps from x to h and another function g that maps from h to y . Therefore, the hidden layer's activation is $f(x)$ and the output of the network is $g(f(x))$. By combining these functions, the level of complexity that we can compute is much greater than using f or g individually.

The neural network learns by adjusting its weights and biases (threshold) iteratively to yield desired output. For learning to take place, the neural network is trained using input-output examples, coupled to a learning algorithm. The most popular algorithms follow gradient descent. In gradient descent, the weights of the network are changed in order to minimise the difference between the actual output and the target output.

The major advantage of using a neural network, is that after training the network

Activation function	Equation	Example	1D Graph
Unit step (Heaviside)	$\phi(z) = \begin{cases} 0, & z < 0, \\ 0.5, & z = 0, \\ 1, & z > 0, \end{cases}$	Perceptron variant	
Sign (Signum)	$\phi(z) = \begin{cases} -1, & z < 0, \\ 0, & z = 0, \\ 1, & z > 0, \end{cases}$	Perceptron variant	
Linear	$\phi(z) = z$	Adaline, linear regression	
Piece-wise linear	$\phi(z) = \begin{cases} 1, & z \geq \frac{1}{2}, \\ z + \frac{1}{2}, & -\frac{1}{2} < z < \frac{1}{2}, \\ 0, & z \leq -\frac{1}{2}, \end{cases}$	Support vector machine	
Logistic (sigmoid)	$\phi(z) = \frac{1}{1 + e^{-z}}$	Logistic regression, Multi-layer NN	
Hyperbolic tangent	$\phi(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}$	Multi-layer Neural Networks	
Rectifier, ReLU (Rectified Linear Unit)	$\phi(z) = \max(0, z)$	Multi-layer Neural Networks	
Rectifier, softplus	$\phi(z) = \ln(1 + e^z)$	Multi-layer Neural Networks	

Copyright © Sebastian Raschka 2016
(<http://sebastianraschka.com>)

Figure 7.2.4: Illustrations of popular activation functions. Figure from: https://rasbt.github.io/mlxtend/user_guide/general_concepts/activation-functions/, accessed 13th March 2018.

on the data, we no longer need to load the training set to perform calculations. This means we could easily convert terrabytes of information into megabytes of information - a reduction factor of a million. This would allow us to perform many calculations on even very small clusters. Of course, the main question is how accurately the neural network can represent the data (opacity), and if this accuracy is sufficient for our purposes.

We investigated a number of settings for the architecture and activation functions of a wide range of neural networks, using the Python package Keras (Chollet et al., 2015). We trained each network on either the regular grid or irregular grid datasets as described




Figure 7.2.5: Illustration of neural network architecture. Figure from: <https://www.xenonstack.com/blog/data-science/overview-of-artificial-neural-networks-and-its-applications>, accessed 13th March 2018.

previously (the neural networks are agnostic to how the data is sampled), i.e. for a range of temperatures, pressures, and volume mixing ratios (representing an atmospheric layer), we calculated a range of corresponding k -distributions to act as the training set. We found that training a network on flattened-g ordinate spectra (i.e. the 20 g-ordinates were repeated for each wavelength bin, to produce a 1D input) produced results that were not accurate or predictable enough for reproducing spectra (errors $>5\%$).

Instead, our preferred model uses an ensemble of networks, one for each wavelength bin which contains 20 g-ordinates, to create the spectral opacity to a higher accuracy. We refer to this ensemble as a ‘superorganism’, as they are a group of synergetically interacting organisms of the same species. Our current preferred network architecture (which could easily change in the future) contains 6 input units (temperature, pressure, 4 volume mixing ratios relative to H_2O - i.e. the volume mixing ratio is now the ratio of

the water to other gas volume mixing ratios, or the water-to-gas mixing ratio), 5 hidden layers with 1024, 512, 256, 128, 64 units, and 20 output units (one for each g-ordinate per wavelength bin), which are all fully connected. This kind of architecture is referred to as a multilayer perceptron (MLP) or a deep learning model. We use the regular grid training set described before, and reserve our irregular grid for use as a validation set. We use ReLu activation functions for our hidden layers, a linear activation function for the output, the mean-squared error as our loss/cost function, and the ADAM optimiser (Kingma & Ba, 2014) as our learning algorithm. This process is repeated for each of the 391 wavelength bins, so that we have 391 trained neural networks. The average time to perform a calculation using the superorganism method is approximately 0.4 seconds. However, as this work is still in development, and is intended as future work, these results may change. At present, the network is not accurate enough to replace the on-the-fly calculations, but we have seen promising results and believe it is possible to achieve the accuracy required.

In the future, we would continue to test different network architectures to optimise the accuracy of our model, to ensure our model would then be able to be used accurately in both forward modelling and in retrievals. We would validate our model by showing multiple comparisons from the true physics calculations to our neural network’s in, firstly a forward modelling sense, and then in a retrieval sense.

7.2.2 Informing WFIRST

The Wide-Field Infrared Survey Telescope (WFIRST) mission is a space observatory that will address the most compelling scientific problems in dark energy, exoplanets and general astrophysics using a 2.4-m telescope with a wide-field infrared instrument and an optical coronagraph. The Astro2010 Decadal Survey recommended a Wide Field Infrared Survey Telescope as its top priority for a new large space mission. As conceived by the decadal survey, WFIRST would carry out a dark energy science program, a microlensing program to determine the demographics of exoplanets, and a general observing program utilizing its ultra wide field (Spergel et al., 2015). It is to be launched in the mid-2020s.

The most successful planet-finding mission yet is the Kepler space observatory, which at the time of writing has confirmed over 2000 exoplanets (Morton et al., 2016). Kepler’s

intrinsic ability to detect planets declines rapidly for small planets with large separations. As a result, Kepler has been unable to find small planets with periods longer than about one year. Simply put: Kepler is sensitive to hot and warm planets, but not to the cold planets in the outer regions of planetary systems, including analogs of all of the planets in our Solar System from Mars outward. In contrast, the exoplanet survey on WFIRST is sensitive to planets from roughly the outer habitable zone outwards, including rocky planets with the mass of Earth up to the largest gas giant planets, and analogs of the ice giant planets in our Solar System.

The Coronagraph Instrument (CGI) on WFIRST will be the first instrument to characterize the atmospheres of gaseous planets around nearby Sun-like stars using reflected visible-light images (430-970nm) and low resolution ($R \sim 70$) spectra in three bands from 600 to 970 nm. This characterization technique has been refined and calibrated over the past 100 years on our Solar System planets Jupiter, Saturn, Uranus, and Neptune. These data will give us unique information on planetary compositions, temperatures, and cloud layers that are not available via transmission and emission spectroscopy of transiting planets that have been observed by Hubble Space Telescope and Spitzer.

Reflected light spectra can probe deep into the atmospheres of cloud-free giant planets, up to ~ 10 bar depths (e.g., Marley et al. (2014)). Near-IR spectra of transiting planets taken with HST WFC3 (and soon JWST) probe only much higher altitude regions, only to ~ 1 mbar (e.g., Kreidberg et al. (2014)). Therefore WFIRST data will probe much greater volumes of exoplanet atmospheres than do transit spectra, and this is the only technique that can be used to do so for cool planets. Most planets observable by the WFIRST CGI have equilibrium effective temperatures of ~ 150 - 350 K, and they are expected to show numerous bands of CH_4 and a moderately strong band of H_2O in their spectra over the $\lambda = 600$ - 970 nm region, the sensitivity range of the CGI integral field spectrometer. The strengths of these bands and the shape of the continuum at shorter wavelengths ($\lambda = 430$ - 600 nm) are diagnostic of the atmospheric temperature, the mixing ratios of C and O, and the amount and location of clouds in the planet's atmosphere (Hu, 2014; Marley et al., 2014; Burrows, 2014).

As the goal of this thesis is to probe the evolution of atmospheric structure and composition as we transition from extrasolar to Solar System temperatures, the reflected-

light spectra retrieved by WFIRST would be a perfect opportunity to compare not only Y-dwarf atmospheres to those of the Solar System, but also cold exoplanet atmospheres.

WFIRST Data Challenge

We participated in a WFIRST data challenge set by the WFIRST science team. The challenge involved retrieving atmospheric parameters on simulated WFIRST exoplanet data (performed by an unknown modeller other than us), where we also participated in telecons to provide feedback for the science team. The reflected light spectra for the highest resolution cases ($\text{SNR} = 20$ at 600nm, $R \sim 70$ for HD150706 b, Ups And e, HD192310 c and $\text{SNR}=10$, $R \sim 5$ for tau Ceti a, b) are shown in Figure 7.2.6 along with their stellar spectra.

HD150706 b, Ups And e represent two Jupiter-like planets at 28.22pc and 13.47pc respectively. These types of targets would be our primary focus for comparisons between Solar System gas giants and Y dwarfs. HD192310 c is a Neptune-like planet at 8.91pc. tau Ceti is a yet-to-be-discovered simulated planet representing a super Earth at 3.65pc, where a and b signify two different masses (rather than two different planets).

At the time of writing, we have analysed only one of these objects extensively, by performing ‘blind’ retrievals on HD150706b for two different resolutions ($R = 20$, $R = 70$) at the same SNR ($=20$). The retrieved parameters vary according to the cloud structure, but all contain the volume mixing ratios of NH_3 , H_2O , and CH_4 , and the radius of the object.

We found that a two-cloud parameterisation (1 Rayleigh-scattering haze, 1 deep water-like cloud) produced the best fit to the higher resolution case ($R = 70$), and a one-cloud parameterisation best fit the lower resolution ($R = 20$) case. This is because spectra with lower information content require simpler parameterisations, otherwise the retrieval becomes wildly inaccurate. This is illustrated in Figures 7.2.7, which shows the posteriors retrieved for the same model on two different resolution spectra. The higher resolution spectrum affords the retrieval enough information to accurately identify two clouds and give sensible posteriors for all parameters, whereas the lower resolution retrieval contains posteriors that are nonsensical. Figures 7.2.8 and 7.2.9 show two more retrievals performed on the lower resolution case, demonstrating that the most sensible posteriors are retrieved




Figure 7.2.6: Left: Normalised reflected light spectra of exoplanets simulated for the WFIRST data challenge. Right: Accompanying simulated stellar spectra. Each object is also labelled with its distance and mass.

from the one-cloud case (Figure 7.2.9).

We believe these results to be helpful to the WFIRST science team, and would like to continue with retrievals on the other suggested objects to help influence the continued design and science return of the instrument.

Other Future Work

The two topics discussed previously (neural networks and the WFIRST instrument), are future issues that we have performed some analysis on already. However, there are many important topics that we would like to explore that we have not begun the analysis on. These include:

- **A Bayesian Survey of All Known Y dwarfs.** Expanding on Chapter 6, we would apply the same techniques to all known Y dwarf spectra, and compare them in a similar manner. As well as presuming a Na_2S cloud, we would explore the possibility of forming a H_2O cloud more rigorously. This would further our understanding of Y dwarfs and begin to build up a picture of atmospheric structure as we transition from extra-Solar to Solar System temperatures.
- **Hyperparameter tuning over Parameterisations.** If we were to complete, successfully, our implementation of the neural network RTE solver, we would be able to perform an enormous number of retrievals very efficiently. A highly discussed issue in this thesis is the effect of parameterisations on the posterior distributions, i.e. a certain temperature profile parameterisation will give different posteriors to another. One way to combat this issue is to perform multiple retrievals using each parameterisation, and compare the Bayesian evidence of each calculation to decide which parameterisation is, in a Bayesian sense, more applicable. By doing this over a range of objects, we can show which temperature profile parameterisations are most applicable, thereby pinning down the true nature of the temperature structure of brown dwarfs further.
- **Marginalising over Data Reduction techniques for Exoplanets.** Transiting and eclipsing exoplanetary data reduction includes modelling the stellar signal and

then removing it from the spectrum, to then be left with the atmospheric signal of the exoplanet. However, the errors in the model that performs this reduction are not propagated through to the atmospheric retrieval. A way to properly propagate these errors would be through marginalisation acquired by simultaneously modelling the atmosphere of the planet and the stellar signal. We would couple our atmospheric code to a data reduction technique, such as Gaussian processes performed by Prof. Suzanne Aigrain, and perform the entire NS/MCMC calculation at once on a range of objects, instead of the usual two steps. By doing this and comparing our retrieved posteriors to the usual two-step approach, we can show how combining techniques in a more accurate way can influence the posterior distributions.

- **Applications to Directly Imaged Exoplanets.** The methods developed here would lend themselves perfectly to directly imaged exoplanets, especially young hot objects with H_2 -dominated atmospheres. We would begin by applying our methods to the HR8799 system, which contains four well studied directly imaged exoplanets. This would place our methods in the context of directly imaged exoplanets, and allow us to easily compare our results to other models. Building on this, we would analyse a larger survey of directly-imaged planets, hoping to build a database of statistics, as we have done for Y dwarfs.
- **Exploring Variability due to Clouds.** In this thesis, we have only analysed Y dwarf objects at one epoch. However, as discussed in Chapter 1, a number of brown dwarf objects are shown to have variability. It is expected that Y dwarfs may also feature spectral variability, especially when developing water clouds. Investigating the variability of the Y dwarfs in our survey would help us better understand the cloud structure and atmospheric composition of these objects.

Any one of these topics could span an entire DPhil worth of work, and we would recommend that this work be done by a group of individuals working synergetically, to produce the most effective science outcomes. We believe these topics are vital to the further understanding of brown dwarfs atmospheres in particular, and atmospheric science in general.




Figure 7.2.7: Posterior distributions for the high-resolution (blue) and low-resolution case (red) for HD150706b. From left to right: log volume mixing ratios of NH_3 , CH_4 , H_2O , radius (Jupiter radii), pressure of cloud base (bar), fractional scale height of cloud, κ the integrated optical depth of the cloud at 0.6 microns, the cloud particle radius (μm), the real refractive index at 0.6 microns, the log of the imaginary refractive index of the cloud, the integrated optical depth of the haze, the pressure of haze base (bar).




Figure 7.2.8: Posterior distributions for the low-resolution case for HD150706b. Cloud structure represented by a haze. From left to right: log volume mixing ratios of NH_3 , CH_4 , H_2O , radius (Jupiter radii), pressure of haze base (bar), fractional scale height of haze, κ the integrated optical depth of the haze at 0.6 microns.




Figure 7.2.9: Posterior distributions for the low-resolution case for HD150706b. Cloud structure represented by a water-like cloud. From left to right: log volume mixing ratios of NH_3 , CH_4 , H_2O , radius (Jupiter radii), pressure of cloud base (bar), fractional scale height of cloud, κ the integrated optical depth of the cloud at 0.6 microns, the cloud particle radius (μm), the real refractive index at 0.6 microns, the log of the imaginary refractive index of the cloud.

Appendix A

Full T and Y dwarf Posteriors

Here we present the full T and Y dwarf posteriors for completeness. GL570D and WISE1738 are omitted here as they are presented in Chapter 6.




Figure A.0.1: WISE0148. Summary of the marginalised posteriors for the retrieved parameters. This shows the marginalised posteriors (1D histograms) along the diagonals and the parameter correlations (2D histograms; contours are shown at 0.5, 1, 1.5, and 2 σ). at off-diagonal locations. Median and $\pm\sigma$ values plotted on top of 1D distributions as dotted lines.




Figure A.0.2: WISE0410. Summary of the marginalised posteriors for the retrieved parameters. This shows the marginalised posteriors (1D histograms) along the diagonals and the parameter correlations (2D histograms; contours are shown at 0.5, 1, 1.5, and 2 σ). at off-diagonal locations. Median and $\pm\sigma$ values plotted on top of 1D distributions as dotted lines.




Figure A.0.3: WISE1405. Summary of the marginalised posteriors for the retrieved parameters. This shows the marginalised posteriors (1D histograms) along the diagonals and the parameter correlations (2D histograms; contours are shown at 0.5, 1, 1.5, and 2 σ). at off-diagonal locations. Median and $\pm\sigma$ values plotted on top of 1D distributions as dotted lines.




Figure A.0.4: WISE1541. Summary of the marginalised posteriors for the retrieved parameters. This shows the marginalised posteriors (1D histograms) along the diagonals and the parameter correlations (2D histograms; contours are shown at 0.5, 1, 1.5, and 2 σ). at off-diagonal locations. Median and $\pm\sigma$ values plotted on top of 1D distributions as dotted lines.




Figure A.0.5: WISE1828. Summary of the marginalised posteriors for the retrieved parameters. This shows the marginalised posteriors (1D histograms) along the diagonals and the parameter correlations (2D histograms; contours are shown at 0.5, 1, 1.5, and 2 σ). at off-diagonal locations. Median and $\pm\sigma$ values plotted on top of 1D distributions as dotted lines.




Figure A.0.6: WISE2056. Summary of the marginalised posteriors for the retrieved parameters. This shows the marginalised posteriors (1D histograms) along the diagonals and the parameter correlations (2D histograms; contours are shown at 0.5, 1, 1.5, and 2 σ). at off-diagonal locations. Median and $\pm\sigma$ values plotted on top of 1D distributions as dotted lines.

Bibliography

- Ackerman, A. S., & Marley, M. S. 2001, *ApJ*, 556, 872
- Allard, F., Guillot, T., Ludwig, H.-G., et al. 2003, in *IAU Symposium*, Vol. 211, Brown Dwarfs, ed. E. Martín, 325
- Amundsen, D. S., Baraffe, I., Tremblin, P., et al. 2014, *A&A*, 564, A59
- Amundsen, D. S., Mayne, N. J., Baraffe, I., et al. 2016, *A&A*, 595, A36
- Apai, D., Radigan, J., Buenzli, E., et al. 2013, *ApJ*, 768, 121
- Artigau, É., Bouchard, S., Doyon, R., & Lafrenière, D. 2009, *ApJ*, 701, 1534
- Barber, R. J., Tennyson, J., Harris, G. J., & Tolchenov, R. N. 2006, *MNRAS*, 368, 1087
- Barstow, J. K., Aigrain, S., Irwin, P. G. J., et al. 2014, *ApJ*, 786, 154
- Baudino, J.-L., Bézard, B., Boccaletti, A., et al. 2015, *A&A*, 582, A83
- Beichman, C., Gelino, C. R., Kirkpatrick, J. D., et al. 2014, *ApJ*, 783, 68
- BelBruno, J. J., Gelfand, J., Radigan, W., & Verges, K. 1982, *Journal of Molecular Spectroscopy*, 94, 336
- Berger, E., Basri, G., Fleming, T. A., et al. 2010, *ApJ*, 709, 332
- Biller, B. A., Crossfield, I. J. M., Mancini, L., et al. 2013, *ApJ*, 778, L10
- Binney, J., & Merrifield, M. 1998, *Galactic Astronomy*, Princeton paperbacks (Princeton University Press)

- Boss, A. P., Butler, R. P., Hubbard, W. B., et al. 2007, Transactions of the International Astronomical Union, Series A, 26, 183
- Bowler, B. P. 2016, PASP, 128, 102001
- Buenzli, E., Apai, D., Radigan, J., Reid, I. N., & Flateau, D. 2014, ApJ, 782, 77
- Buldyreva, J., Lavrentieva, N., & Starikov, V. 2011, Collisional Line Broadening and Shifting of Atmospheric Gases: A Practical Guide for Line Shape Modeling by Current Semi-classical Approaches (Imperial College Press)
- Burgasser, A. J. 2014, in Astronomical Society of India Conference Series, Vol. 11, Astronomical Society of India Conference Series
- Burgasser, A. J., Marley, M. S., Ackerman, A. S., et al. 2002a, ApJ, 571, L151
- Burgasser, A. J., Kirkpatrick, J. D., Brown, M. E., et al. 2002b, ApJ, 564, 421
- Burgasser, A. J., Simcoe, R. A., Bochanski, J. J., et al. 2010, ApJ, 725, 1405
- Burningham, B., Marley, M. S., Line, M. R., et al. 2017, MNRAS, 470, 1177
- Burrows, A. 2014, ArXiv e-prints, arXiv:1412.6097
- Burrows, A., Hubbard, W. B., Lunine, J. I., & Liebert, J. 2001, Reviews of Modern Physics, 73, 719
- Burrows, A., & Liebert, J. 1993, Reviews of Modern Physics, 65, 301
- Burrows, A., & Sharp, C. M. 1999, ApJ, 512, 843
- Burrows, A., Sudarsky, D., & Hubeny, I. 2006, ApJ, 640, 1063
- Chabrier, G., Baraffe, I., Allard, F., & Hauschildt, P. 2000a, ApJ, 542, L119
- . 2000b, ApJ, 542, 464
- Chabrier, G., Johansen, A., Janson, M., & Rafikov, R. 2014, Protostars and Planets VI, 619
- Chabrier, G., & Küker, M. 2006, A&A, 446, 1027

- Chollet, F., et al. 2015, Keras, <https://github.com/keras-team/keras>
- Cox, A. N. 2000, Allen’s astrophysical quantities
- Crossfield, I. J. M., Biller, B., Schlieder, J. E., et al. 2014, *Nature*, 505, 654
- Cushing, M. C., Marley, M. S., Saumon, D., et al. 2008, *ApJ*, 678, 1372
- Cushing, M. C., Kirkpatrick, J. D., Gelino, C. R., et al. 2011, *ApJ*, 743, 50
- Cushing, M. C., Hardegree-Ullman, K. K., Trucks, J. L., et al. 2016, *ApJ*, 823, 152
- de Kok, R., Irwin, P. G. J., & Teanby, N. A. 2008, *Icarus*, 197, 572
- Drummond, B., Tremblin, P., Baraffe, I., et al. 2016, *A&A*, 594, A69
- Dupuy, T. J., & Liu, M. C. 2017, *ApJS*, 231, 15
- Feroz, F., & Hobson, M. P. 2008, *MNRAS*, 384, 449
- Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, 398, 1601
- Feroz, F., Hobson, M. P., Cameron, E., & Pettitt, A. N. 2013, *ArXiv e-prints*, arXiv:1306.2144
- Fletcher, L. N., Hesman, B. E., Achterberg, R. K., et al. 2012, *Icarus*, 221, 560
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
- Freedman, R. S., Lustig-Yaeger, J., Fortney, J. J., et al. 2014, *ApJS*, 214, 25
- Freedman, R. S., Marley, M. S., & Lodders, K. 2008, *ApJS*, 174, 504
- Frierson, D. M. W., Held, I. M., & Zurita-Gotor, P. 2006, *Journal of Atmospheric Sciences*, 63, 2548
- Gagne, J. 2018, List of All Ultracool Dwarfs
- Gamache, R. R., Lynch, R., & Brown, L. R. 1996, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 56, 471
- Geballe, T. R., Knapp, G. R., Leggett, S. K., et al. 2002, *ApJ*, 564, 466

- Gelino, C., & Marley, M. 2000, in *Astronomical Society of the Pacific Conference Series*, Vol. 212, *From Giant Planets to Cool Stars*, ed. C. A. Griffith & M. S. Marley, 322
- Gelino, C. R., Marley, M. S., Holtzman, J. A., Ackerman, A. S., & Lodders, K. 2002, *ApJ*, 577, 433
- Giles, R. S., Fletcher, L. N., & Irwin, P. G. J. 2015, *Icarus*, 257, 457
- Gillon, M., Triaud, A. H. M. J., Jehin, E., et al. 2013, *A&A*, 555, L5
- Girardin, F., Artigau, É., & Doyon, R. 2013, *ApJ*, 767, 61
- Golimowski, D. A., Leggett, S. K., Marley, M. S., et al. 2004, *AJ*, 127, 3516
- Goody, R., & Yung, Y. 1989, *Atmospheric Radiation: Theoretical Basis* (Oxford University Press)
- Haddad, S., Aroui, H., Orphal, J., Bouanich, J.-P., & Hartmann, J.-M. 2001, *Journal of Molecular Spectroscopy*, 210, 275
- Hallinan, G., Antonova, A., Doyle, J. G., et al. 2008, *ApJ*, 684, 644
- Hedges, C., & Madhusudhan, N. 2016, *MNRAS*, 458, 1427
- Helling, C., & Casewell, S. 2014, *A&A Rev.*, 22, 80
- Heng, K., Hayek, W., Pont, F., & Sing, D. K. 2012, *MNRAS*, 420, 20
- Hennnebelle, P., & Chabrier, G. 2011, in *IAU Symposium*, Vol. 270, *Computational Star Formation*, ed. J. Alves, B. G. Elmegreen, J. M. Girart, & V. Trimble, 159–168
- Hill, C., Yurchenko, S. N., & Tennyson, J. 2013, *Icarus*, 226, 1673
- Hu, R. 2014, *ArXiv e-prints*, arXiv:1412.7582
- Huang, X., Freedman, R. S., Tashkun, S. A., Schwenke, D. W., & Lee, T. J. 2013, *J. Quant. Spec. Radiat. Transf.*, 130, 134
- Huang, X., Gamache, R. R., Freedman, R. S., Schwenke, D. W., & Lee, T. J. 2014, *J. Quant. Spec. Radiat. Transf.*, 147, 134

- Hubeny, I., & Mihalas, D. 2014, *Theory of Stellar Atmospheres*
- Humlicek, J. 1982, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 27, 437
- Irwin, P. G. J., Parrish, P., Fouchet, T., et al. 2004, *Icarus*, 172, 37
- Irwin, P. G. J., Teanby, N. A., de Kok, R., et al. 2008, *J. Quant. Spec. Radiat. Transf.*, 109, 1136
- Joergens, V., ed. 2014, *Astrophysics and Space Science Library*, Vol. 401, 50 Years of Brown Dwarfs: From Prediction to Discovery to Forefront of Research
- Kimble, R. A., MacKenty, J. W., O’Connell, R. W., & Townsend, J. A. 2008, in *Proc. SPIE*, Vol. 7010, *Space Telescopes and Instrumentation 2008: Optical, Infrared, and Millimeter*, 70101E
- Kingma, D. P., & Ba, J. 2014, *ArXiv e-prints*, arXiv:1412.6980
- Kirkpatrick, J. D., Reid, I. N., Liebert, J., et al. 1999, *ApJ*, 519, 802
- . 2000, *AJ*, 120, 447
- Kirkpatrick, J. D., Cushing, M. C., Gelino, C. R., et al. 2011, *ApJS*, 197, 19
- Kirkpatrick, J. D., Gelino, C. R., Cushing, M. C., et al. 2012, *ApJ*, 753, 156
- Kostov, V., & Apai, D. 2013, *ApJ*, 762, 47
- Kreidberg, L., Bean, J. L., Désert, J.-M., et al. 2014, *ApJ*, 793, L27
- Kumar, S. S. 1963, *ApJ*, 137, 1121
- Kupka, F. G., Ryabchikova, T. A., Piskunov, N. E., Stempels, H. C., & Weiss, W. W. 2000, *Baltic Astronomy*, 9, 590
- Lacis, A. A., & Oinas, V. 1991, *J. Geophys. Res.*, 96, 9027
- Lee, J.-M. 2012, PhD thesis, University of Oxford
- Lee, J.-M., Irwin, P. G. J., Fletcher, L. N., Heng, K., & Barstow, J. K. 2014, *ApJ*, 789,

- Leggett, S. K., Morley, C. V., Marley, M. S., et al. 2013, *ApJ*, 763, 130
- Leggett, S. K., Tremblin, P., Esplin, T. L., Luhman, K. L., & Morley, C. V. 2017, *ApJ*, 842, 118
- Leggett, S. K., Cushing, M. C., Saumon, D., et al. 2009, *ApJ*, 695, 1517
- Leggett, S. K., Saumon, D., Marley, M. S., et al. 2012, *ApJ*, 748, 74
- Leggett, S. K., Cushing, M. C., Hardegree-Ullman, K. K., et al. 2016, *ApJ*, 830, 141
- 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., Stevenson, K. B., Bean, J., et al. 2016a, *ArXiv e-prints*, arXiv:1605.08810
- Line, M. R., Marley, M. S., Liu, M. C., et al. 2016b, *ArXiv e-prints*, arXiv:1612.02809
- Lodders, K. 1999, *ApJ*, 519, 793
- Lodders, K., & Fegley, Jr., B. 2006, *Chemistry of Low Mass Substellar Objects*, ed. J. W. Mason, 1
- Luhman, K. L. 2012, *ARA&A*, 50, 65
- MacDonald, R. J., & Madhusudhan, N. 2017, *ArXiv e-prints*, arXiv:1701.01113
- Malik, M., Grosheintz, L., Mendonça, J. M., et al. 2017, *AJ*, 153, 56
- Mantz, A., Devi, V. M., Benner, D. C., et al. 2005, *Journal of Molecular Structure*, 742, 99
- Margolis, J. S. 1993, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 49, 71
- Marley, M., Lupu, R., Lewis, N., et al. 2014, *ArXiv e-prints*, arXiv:1412.8440
- Marley, M. S., Seager, S., Saumon, D., et al. 2002, *ApJ*, 568, 335
- Martín, E. L., Delfosse, X., Basri, G., et al. 1999, *AJ*, 118, 2466

- Mayor, M., & Queloz, D. 1995, *Nature*, 378, 355
- McLean, I. S., Graham, J. R., Becklin, E. E., et al. 2000, in *Proc. SPIE*, Vol. 4008, *Optical and IR Telescope Instrumentation and Detectors*, ed. M. Iye & A. F. Moorwood, 1048–1055
- McLean, I. S., McGovern, M. R., Burgasser, A. J., et al. 2003, *ApJ*, 596, 561
- McLean, I. S., Becklin, E. E., Bendiksen, O., et al. 1998, in *Proc. SPIE*, Vol. 3354, *Infrared Astronomical Instrumentation*, ed. A. M. Fowler, 566–578
- Moal, M., & Severin, F. 1986, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 35, 145
- Mohanty, S., Baraffe, I., & Chabrier, G. 2007, in *IAU Symposium*, Vol. 239, *IAU Symposium*, ed. F. Kupka, I. Roxburgh, & K. L. Chan, 197–204
- Mohanty, S., Basri, G., Shu, F., Allard, F., & Chabrier, G. 2002, *ApJ*, 571, 469
- Mollière, P., van Boekel, R., Dullemond, C., Henning, T., & Mordasini, C. 2015, *ApJ*, 813, 47
- Morley, C. V., Fortney, J. J., Marley, M. S., et al. 2012, *ApJ*, 756, 172
- Morley, C. V., Marley, M. S., Fortney, J. J., & Lupu, R. 2014a, *ApJ*, 789, L14
- Morley, C. V., Marley, M. S., Fortney, J. J., et al. 2014b, *ApJ*, 787, 78
- Morton, T. D., Bryson, S. T., Coughlin, J. L., et al. 2016, *ApJ*, 822, 86
- Mukherjee, P., Parkinson, D., & Liddle, A. R. 2006, *ApJ*, 638, L51
- Muller-Kirsten, H. J. W. 2013, *Basics of Statistical Physics (Second Edition)*, 2nd edn. (World Scientific Publishing Company)
- Nakajima, T., Oppenheimer, B. R., Kulkarni, S. R., et al. 1995, *Nature*, 378, 463
- Nakajima, T., Tsuji, T., & Yanagisawa, K. 2004, *ApJ*, 607, 499
- Nouri, S., Orphal, J., Aroui, H., & Hartmann, J.-M. 2004, *Journal of Molecular Spectroscopy*, 227, 60

- Oppenheimer, B. R., Kulkarni, S. R., Matthews, K., & van Kerkwijk, M. H. 1998, *ApJ*, 502, 932
- Padmanabhan, A., Tzanetakis, T., Chanda, A., & Thomson, M. 2014, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 133, 81
- Paula, J. A. P. D. 2006, *Physical Chemistry*, 8th edn. (Oxford University Press)
- Pierrehumbert, R. T. 2010, *Principles of Planetary Climate*
- Pine, A., Markov, V., Buffa, G., & Tarrini, O. 1993, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 50, 337
- Pine, A. S. 1992, *The Journal of Chemical Physics*, 97
- Radigan, J., Jayawardhana, R., Lafrenière, D., et al. 2012, *ApJ*, 750, 105
- Radigan, J., Lafrenière, D., Jayawardhana, R., & Artigau, E. 2014, *ApJ*, 793, 75
- Rajan, A., Patience, J., Wilson, P. A., et al. 2015, *MNRAS*, 448, 3775
- Rebolo, R., Osorio, M. R. Z., & Martin, E. L. 1995, *Nature*, 377, 129
- Regalia-Jarlot, L., Thomas, X., der Heyden, P. V., & Barbe, A. 2005, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 91, 121
- Reid, I. N., Burgasser, A. J., Cruz, K. L., Kirkpatrick, J. D., & Gizis, J. E. 2001, *AJ*, 121, 1710
- Reiners, A., & Basri, G. 2008, *ApJ*, 684, 1390
- Richard, C., Gordon, I. E., Rothman, L. S., et al. 2012, *J. Quant. Spec. Radiat. Transf.*, 113, 1276
- Robinson, T. D., & Catling, D. C. 2012, *ApJ*, 757, 104
- Robinson, T. D., & Marley, M. S. 2014, *ApJ*, 785, 158
- Rodgers, C. 2000, *Inverse Methods for Atmospheric Sounding: Theory and Practice*, Series on Atmospheric, Oceanic and Planetary Physics (World Scientific)

- Rodgers, C. D., & Walshaw, C. D. 1966, Quarterly Journal of the Royal Meteorological Society, 92, 67
- Rothman, L., Gordon, I., Barber, R., et al. 2010, Journal of Quantitative Spectroscopy and Radiative Transfer, 111, 2139 , {XVIth} Symposium on High Resolution Molecular Spectroscopy (HighRus-2009)XVIth Symposium on High Resolution Molecular Spectroscopy
- Rothman, L., Gordon, I., Babikov, Y., et al. 2013, Journal of Quantitative Spectroscopy and Radiative Transfer, 130, 4 , {HITRAN2012} special issue
- Route, M., & Wolszczan, A. 2012, ApJ, 747, L22
- Ruanaidh, J. J. O., & Fitzgerald, W. J. 1996, Numerical Bayesian Methods Applied to Signal Processing (Statistics and Computing), 1996th edn. (Springer)
- Saumon, D., & Marley, M. S. 2008, ApJ, 689, 1327
- Saumon, D., Marley, M. S., Abel, M., Frommhold, L., & Freedman, R. S. 2012, ApJ, 750, 74
- Sauval, A. J., & Tatum, J. B. 1984, ApJS, 56, 193
- Sharp, C. M., & Burrows, A. 2007, ApJS, 168, 140
- Showman, A. P., & Kaspi, Y. 2013, ApJ, 776, 85
- Simcoe, R. A., Burgasser, A. J., Bernstein, R. A., et al. 2008, in Proc. SPIE, Vol. 7014, Ground-based and Airborne Instrumentation for Astronomy II, 70140U
- Simcoe, R. A., Burgasser, A. J., Schechter, P. L., et al. 2013, PASP, 125, 270
- Simon, A. A., Rowe, J. F., Gaulme, P., et al. 2016, ApJ, 817, 162
- Sing, D. K., Fortney, J. J., Nikolov, N., et al. 2016, Nature, 529, 59
- Skemer, A. J., Morley, C. V., Allers, K. N., et al. 2016, ApJ, 826, L17
- Skilling, J. 2004, in American Institute of Physics Conference Series, Vol. 735, American Institute of Physics Conference Series, ed. R. Fischer, R. Preuss, & U. V. Toussaint, 395–405

- Smart, R. L., Apai, D., Kirkpatrick, J. D., et al. 2017, *MNRAS*, 468, 3764
- Sneep, M., & Ubachs, W. 2005, *J. Quant. Spec. Radiat. Transf.*, 92, 293
- Solodov, A., & Starikov, V. 2009, *Molecular Physics*, 107, 43
- Spergel, D., Gehrels, N., Baltay, C., et al. 2015, *ArXiv e-prints*, arXiv:1503.03757
- Stephens, D. C., Leggett, S. K., Cushing, M. C., et al. 2009, *ApJ*, 702, 154
- Steyert, D. W., Wang, W., Sirota, J., Donahue, N. M., & Reuter, D. C. 2004, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 83, 183
- Tashkun, S. A., & Perevalov, V. I. 2011, *J. Quant. Spec. Radiat. Transf.*, 112, 1403
- Tennyson, J., & Yurchenko, S. N. 2012, *MNRAS*, 425, 21
- Thibault, F., Boisssoles, J., Le Doucen, R., et al. 1992, *J. Chem. Phys.*, 96, 4945
- Thibault, F., Calil, B., Boisssoles, J., & Launay, J. M. 2000, *Phys. Chem. Chem. Phys.*, 2, 5404
- Thomas, G. E., & Stamnes, K. 2002, *Radiative Transfer in the Atmosphere and Ocean*
- Tipping, R. H. 1976, *Journal of Molecular Spectroscopy*, 61, 272
- Tremblin, P., Amundsen, D. S., Mourier, P., et al. 2015, *ApJ*, 804, L17
- Trucks, J., Cushing, M., Hardegree-Ullman, K., et al. 2014, in *American Astronomical Society Meeting Abstracts*, Vol. 223, *American Astronomical Society Meeting Abstracts* 223, 441.22
- Tsuji, T., & Nakajima, T. 2003, *ApJ*, 585, L151
- van Leeuwen, F. 2007, *A&A*, 474, 653
- Varanasi, P., & Chudamani, S. 1990, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 43, 1
- Visscher, C., Lodders, K., & Fegley, Jr., B. 2006, *ApJ*, 648, 1181
- Vrba, F. J., Henden, A. A., Luginbuhl, C. B., et al. 2004, *AJ*, 127, 2948

- Waldmann, I. P., Tinetti, G., Rocchetto, M., et al. 2015, ApJ, 802, 107
- Weaver, C. P., & Ramanathan, V. 1995, J. Geophys. Res., 100, 11585
- Wenger, C., & Champion, J. P. 1998, J. Quant. Spec. Radiat. Transf., 59, 471
- Wigner, E., & Huntington, H. B. 1935, J. Chem. Phys., 3, 764
- Yurchenko, S. N., Barber, R. J., & Tennyson, J. 2011, MNRAS, 413, 1828
- Yurchenko, S. N., & Tennyson, J. 2014, MNRAS, 440, 1649
- Zhang, X., & Showman, A. P. 2014, ApJ, 788, L6