Atmospheric Retrieval Methods for Clear and Cloudy Brown Dwarfs

Kavli Report
$\sim$ 4,000 words

Ryan Garland

Supervisor: Prof. Mike Line

Contents

1

Introduction

> *"Reports that say that something hasn't happened are always interesting to me, because as we know, there are known knowns; there are things we know we know. We also know there are known unknowns; that is to say we know there are some things we do not know. But there are also unknown unknowns - the ones we don't know we don't know. And if one looks throughout the history of our country and other free countries, it is the latter category that tend to be the difficult ones."*
>
> – Donald Rumsfeld

When performing a retrieval on the atmospheres of exoplanets and brown dwarfs, whether it be by Markoc Chain Monte Carlo (MCMC), Nested Sampling (NS) or optimal estimation (OE), there is an implicit assumption that the forward model is 'correct'. There are ways of artificially introducing unaccounted for errors, such as adding in estimates of systematic errors into the measurement covariance matrix (Irwin et al., 2008), or allowing the *a priori* errors to inflate in a Bayesian manner during the retrieval process (Line et al., 2015). This is one way of addressing the 'known unknowns'. However they do not directly address the effect of propagating errors due to different kinds of assumptions and parameterisations. There are many assumptions and missing pieces of physics that could affect the retrieved values in atmospheric models: one-dimensional radiative transfer (Feng et al., 2016; Agúndez et al., 2014), local thermodynamic equilibrium (López-Puertas & Taylor, 2001), constant-in-altitude volume mixing ratios, homogenous cloud structure (Artigau et al., 2009), etc.

In this report we tackle the problem of a major 'known unknown' in atmospheric retrievals - the parameterisation of the temperature-pressure (TP) profile. The TP-profile is highly degenerate with the volume mixing ratios and cloud structure, and so the parameterisation used must be carefully considered. Choosing which parameterisation to use is a struggle between having the flexibility to fit the data, leading to possible ill-conditioning and

unphysical vertical oscillations in temperature, and the enforcement of physical constraints, which could smooth out real features. Here we explore the effects of these parameterisations on the retrieval of other parameters.

2

CHIMERA Model and Retrieval Method

We use the CHIMERA forward model Line et al. (2015), which computes the upwelling one-dimensional disk-integrated thermal emission spectrum given the molecular abundances, TP profile, and gravity g. We include constant-with-altitude volume (molar) mixing ratios for H_2O, CH_4, CO, CO_2, NH_3, H_2S, and alkali opacities. These are the species known to be found in cool dwarf atmospheres that have spectral signatures in the near-infrared. The alkali opacities include only sodium and potassium and are treated as only one free parameter with their ratio assumed to be solar. Hydrogen/ helium in solar ratio is assumed to make up the remainder of the gas.

Most of the opacity database is drawn from Freedman et al. (2014), but it also includes the most up-to-date methane line list (Yurchenko & Tennyson, 2014) with the line broadening coefficients from Margolis (1996). The opacities are sampled from high-resolution cross sections at 1cm^{-1} resolution (Sharp & Burrows, 2007).

Finally, the high-resolution spectra are convolved with a wavelength-dependent Gaussian instrumental profile that reflects the wavelength-dependent resolving power, and then binned to the data wavelength grid for direct data-model comparison.

The Bayesian retrieval model we use is the Markov chain Monte Carlo (MCMC) approach implemented with affine-invariant ensemble sampler EMCEE Foreman-Mackey et al. (2013). We describe our log-likelihood and priors to calculate the log-posterior separately in each temperature-pressure

profile parameterisation, as each case is unique.

3

Temperature-Pressure Profile Parameterisations

3.1 The Unconstrained Layer-By-Layer Approach

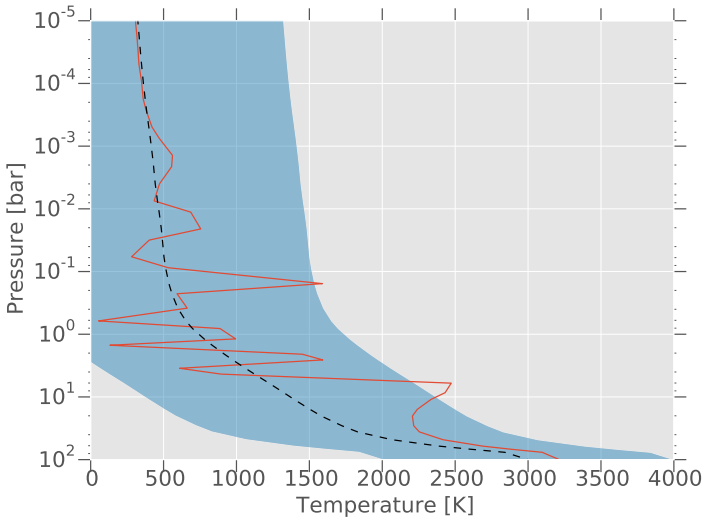

Figure 3.1.1: The retrieved temperature profile from optimal estimation of a typical T6 dwarf is shown in orange, prior temperature profile in dotted black, and the temperature errors in shaded light blue.

The usual approach to calculate the emergent spectra from planets, exoplanets and brown dwarfs is to approximate an inhomogenous atmosphere as many (∼40-100) stacked homogeneous layers. One then retrieves the relevant parameters at each of these layers (temperature, clouds), or sometimes as a globally averaged value over all layers (VMRs). Figure 3.1.1 illustates the optimal estimation retrieval of a typical brown dwarf spectrum using

extremely large errors on the temperature points, so that the retrieval is 'unconstrained' and hence ill-conditioning occurs. Here we see very unphysical oscillations due to the overfitting of noisy features in the spectrum that we are retrieving from. This is therefore not an appropriate way to parameterise the TP-profile.

However, we will now introduce a unique case where this unconstrained layer-by-layer approach can be useful in providing insight. If we produce a spectrum from our CHIMERA radiative transfer code, and then perform an MCMC to retrieve (for example) temperature on that spectrum using the exact input values used for our forward model as our starting positions, we will begin at the global maximum likelihood (i.e., the best fit to the data), and any further exploration of parameter space will therefore return to this global maximum. By doing this, we can examine the 'true information content' of our spectrum in TP- parameter space. This experiment will give us an intuitive sense for the 'true information content' of a spectrum in TP-space, from which we can measure against our other parameterisations. We will now demonstate this point.

Here we use the CHIMERA code to create a typical cloud-free, late-T ($T_{eff} \sim 800$K) dwarf forward model spectrum so that we may perform an MCMC retrieval on it. We will retrieve 15 temperature points, using a hermitian spline to interpolate between those points for the full radiative transfer calculation with 70 layers. All other values (VMRs, mass, radius, etc) are fixed to their true values. For this unconstrained layer-by-layer approach, our unnormalised log-posterior is given only by the uniform temperature prior (0-4000K, which gives a probability of either 0 or 1), and the log-likelihood function:

$$\ln\mathcal{L}(\boldsymbol{y}|\boldsymbol{x}) = -\frac{1}{2}\sum_{i=1}^{n}\frac{(y_i - F_i(\boldsymbol{x}))^2}{{\sigma_i}^2} \tag{3.1}$$

where the index i denotes the ith data point, y is the measured flux, $F(\boldsymbol{x})$ is the modelled flux that comes from the forward model with walker position $\boldsymbol{x}$, and σ is the data error.

Our starting positions for the temperature-pressure profile are shown in Figure 3.1.2. We then proceed with our MCMC for 10000 iterations. Our marginalised posterior is shown in Figure 3.1.3 where we have summarised 5000 randomly selected samples of the TP-profile after 10000 iterations. We see in Figure 3.1.3 that deep in the atmosphere there is a lack of information on the temperature structure, indicated by the large spread in the uncertainty of the temperature. This is because it is below the photosphere, where the optical depth is so large ($\sim$ unity) that the object is opaque. Similarly there is a lack of information in the upper region of the atmosphere, where there is very little absorption by the atmospheric gases.

To see how the addition of clouds may play a role in the retrieval of the

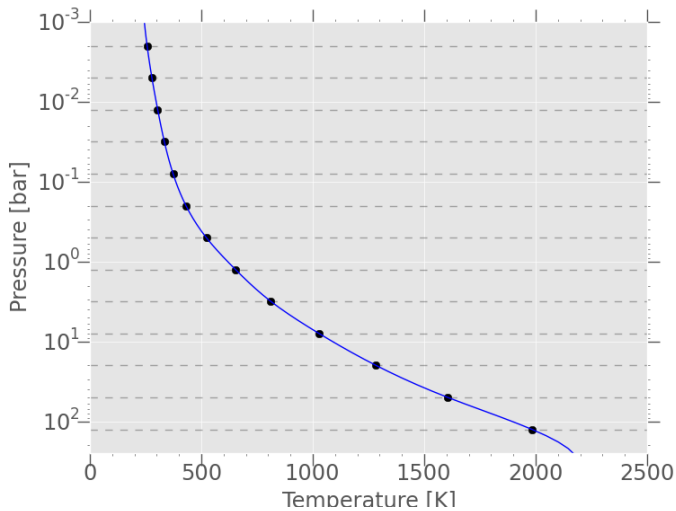

Figure 3.1.2: Unconstrained layer-by-layer approach. The starting and 'true' temperature-pressure profile to be retrieved in the MCMC of a typical late-T dwarf (cloud-free, $T_{eff} \sim 800K$), where the black dots indicate the 15 temperature-pressure starting positions/true values to be retrieved and the blue line indicates the interpolated profile. Horizontal dashed lines indicate the pressure levels at which we are retrieving.

temperature structure, we repeat this experiment with a typical mid-T dwarf (cloudy, $T_{eff} \sim 1000K$) containing a cloud placed at 10 bar with a constant opacity with respect to height to the top of the atmosphere. Figure 3.1.4 shows our starting positions and true profile, while Figure 3.1.5 shows our converged retrieval. We see that the amount of information in TP-space has lessened as the photosphere decreases in size, due to the increase in optical depth caused by the cloud. This experiment now gives us the intuitive sense for the 'true information content' of a brown dwarf spectrum in TP-space for both cloudy and clear cases, from which we can use as a measure against our other parameterisations.

3.2 The Parametric Approach

Another method of attack for retrieving the temperature-pressure profile is to parameterise the profile so that it is described by a small number of variables, while still encapsulating as much physics as possible (Parmentier & Guillot, 2014; Madhusudhan & Seager, 2009; Robinson & Catling, 2012; Line et al., 2016). Using fewer parameters to describe the TP-profile can be advantageous because convergence occurs much more quickly in the MCMC, as the number of samples required to create a statistically independent pos-

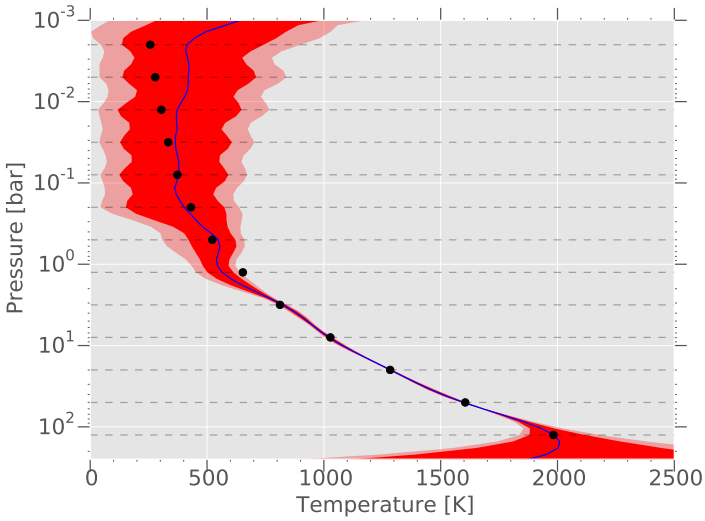

Figure 3.1.3: Unconstrained layer-by-layer approach. The retrieved temperature-pressure profile in the MCMC of a typical late-T dwarf (cloud-free, $T_{eff} \sim 800K$), where the black dots indicate the 15 true temperature-pressure points and the blue line indicates the median retrieved profile. Horizontal dashed lines indicate the pressure levels at which we are retrieving. The dark red shading indicates the area in TP-space encompassing $\pm 1\sigma$ and light red shading indicates $\pm 2\sigma$.

terior is proportional to the number of dimensions.

In this work we choose to use a modified version of the parameterisation used in Madhusudhan & Seager (2009), where Layer 3 is removed (see Figure 3.2.1). We choose this parameterisation because it has been shown to be flexible enough to work on our own Solar System planets, as well as exoplanets (Madhusudhan & Seager, 2009), and uses only four parameters to fully characterise the TP-profile, allowing for swift convergence in an MCMC.

The temperature-pressure profile parameterisation is split into two layers, described by:

$$\begin{aligned} P_0 < P < P_1 : \quad & P = P_0 \exp \alpha_1 (T - T_0)^{\beta_1} \quad \text{Layer 1} \\ P_1 < P < P_2 : \quad & P = P_2 \exp \alpha_2 (T - T_2)^{\beta_2} \quad \text{Layer 2} \end{aligned} \tag{3.2}$$

where, P is the pressure in bars, T is the temperature in K, and P_i, T_i, α_i, and β_i are free parameters. We set $\beta_1 = \beta_2 = 0.5$ as in Madhusudhan & Seager (2009). α describes the slope of the pressure-temperature curves. Rewriting this in terms of temperature we have:

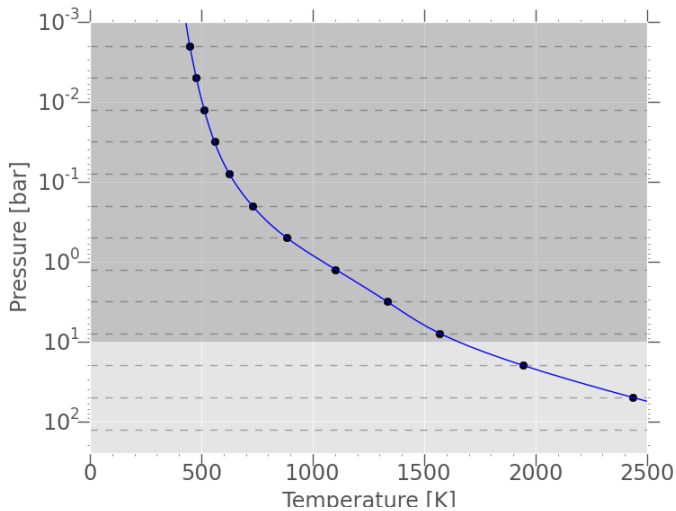

Figure 3.1.4: Unconstrained layer-by-layer approach. The starting and 'true' temperature-pressure profile to be retrieved in the MCMC of a typical mid-T dwarf (cloudy, $T_{eff} \sim 1000K$), where the black dots indicate the 15 temperature-pressure starting positions/true values to be retrieved and the blue line indicates the interpolated profile. Horizontal dashed lines indicate the pressure levels at which we are retrieving. The darker grey shading indicates the area covered by the cloud.

$$\begin{aligned} P_0 < P < P_1 : \quad T &= \left(\frac{\ln \frac{P}{P_0}}{\alpha_1}\right)^2 + T_0 \quad \text{Layer 1} \\ P_1 < P < P_2 : \quad T &= \left(\frac{\ln \frac{P}{P_2}}{\alpha_2}\right)^2 + T_2 \quad \text{Layer 2} \end{aligned} \tag{3.3}$$

and we can also determine T_2 independently through the intersection of these layers:

$$P = P_1 : \quad T_2 = \left(\frac{\ln \frac{P}{P_0}}{\alpha_1}\right)^2 + T_0 - \left(\frac{\ln \frac{P}{P_2}}{\alpha_2}\right)^2 \quad \text{Layer 1-2 Intersection} \tag{3.4}$$

If we now define our top of atmosphere, P_0, as a constant, we can fully describe our brown dwarf temperature-pressure profile using only four parameters: α_1, α_2, T_0, and P_2.

We now create a typical mid-T ($T_{eff} \sim 1000K$) brown dwarf spectrum whose atmosphere consists of a thick cloud at 10 bar reaching upwards a few scale heights, and a composition consisting of only 400ppm H_2O and a solar abundance mixture of H_2 and He. We are thus retrieving not only

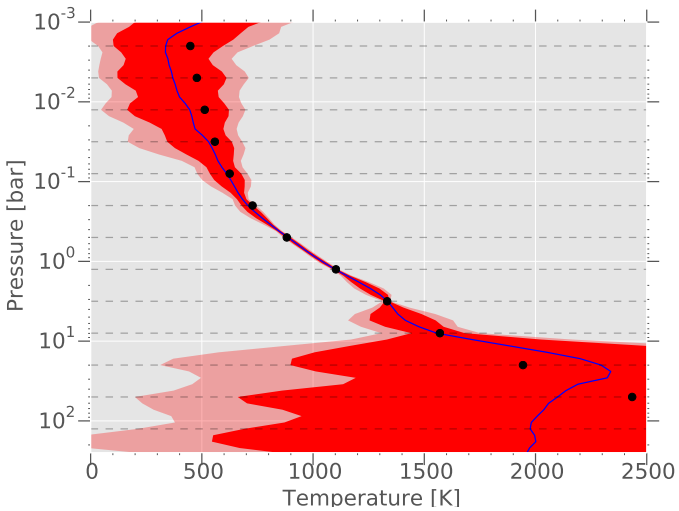

Figure 3.1.5: Unconstrained layer-by-layer approach. The retrieved temperature-pressure profile in the MCMC of a typical mid-T dwarf (cloudy, $T_{eff} \sim 1000K$), where the black dots indicate the 15 true temperature-pressure points and the blue line indicates the median retrieved profile. Horizontal dashed lines indicate the pressure levels at which we are retrieving. The dark red shading indicates the area in TP-space encompassing $\pm 1\sigma$ and light red shading indicates $\pm 2\sigma$. The darker grey shading indicates the area covered by the cloud.

the temperature profile, but also the VMR of H_2O, and the three cloud properties of opacity, base pressure, and scale height. In this case, while we do have a 'true TP-profile', we do not initialise our calculations around this global maximum. This is so that we can truly test the ability of the parameterisation to fit the data and how the uncertainties of the parameterisation propogate into other retrieved parameters. Our log-likelihod function is the same as in the unconstrained case, and our priors are all uniform. The volume mixing ratios uniform priors are between 1E-12 and 1, where all of the remaining gas is filled by H_2 and He in a solar ratio. The log opacity is between -12 and 0, the scale height is between 0 and 10, and the base pressure of the cloud is between the top and bottom of the atmosphere. α is uniform between 1e-6 - 1e6, P_2 is between the top and bottom of the atmosphere, and T_0 is between 0 and 4000K.

In Figure 3.2.3, we show the unsummarised TP-profiles of the 5000 randomly selected samples from the posterior, while in Figure 3.2.4 we show a summarised version. The darkest shades of red in Figure 3.2.3 indicate the preference for the walkers to sample that particular set of parameter space, and show that it is more likely to fit the data well. While one set

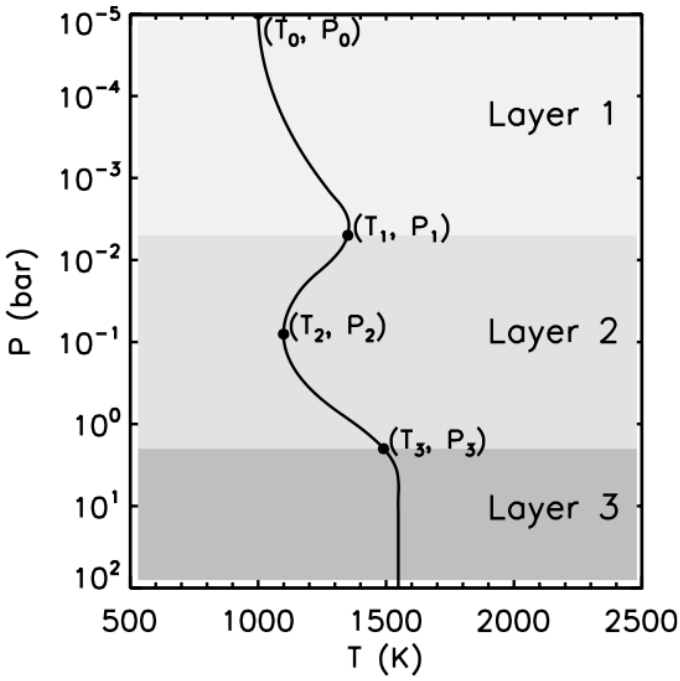

Figure 3.2.1: Parametric temperature-pressure (TP) profile. In the general form, the profile includes a thermal inversion layer (Layer 2) and has six free parameters (see Equation 3.2). For a profile with thermal inversion, $P_2 > P_1$,and for one without a thermal inversion, $P_1 \geq P_2$. An isothermal profile is assumed for pressures sabove P_3 (Layer 3). Alternatively, for cooler atmospheres with no isothermal layer, Layer 2 could extend to deeper layers and Layer 3 could be absent. This is what is done in the present work.

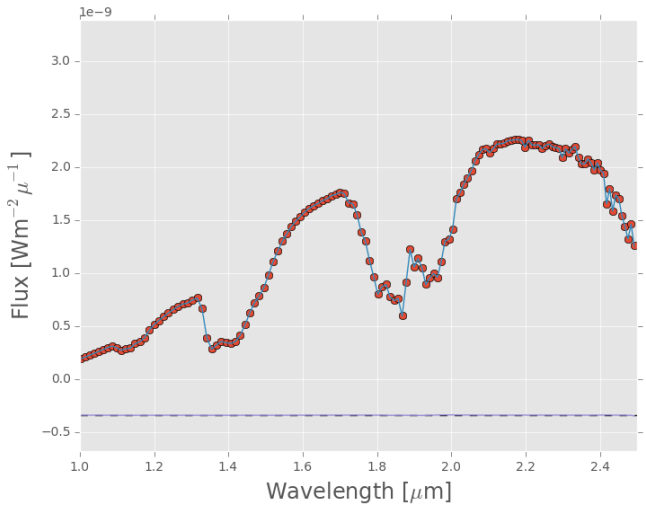

Figure 3.2.2: Parametric parameterisation. The best fit (blue) to the fake data (orange dots) produced by the parametric model. A dotted line shows the residual between the fit and the data.

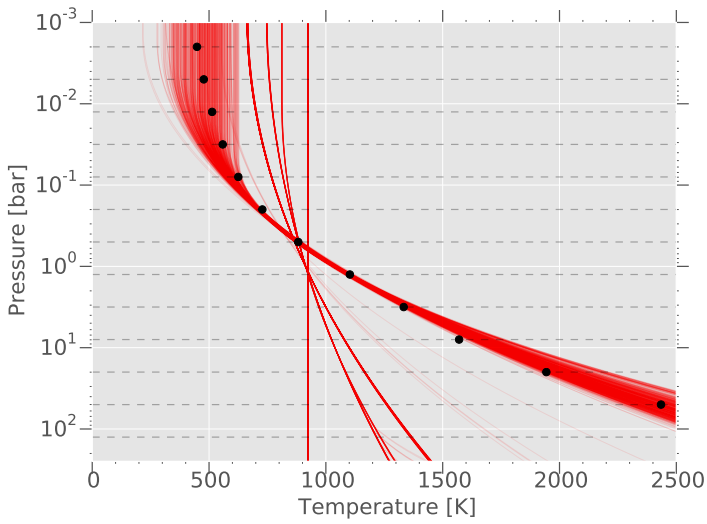

Figure 3.2.3: Parametric temperature-pressure (TP) profile retrieval showing 5000 randomly selected samples of the MCMC where the black dots indicate the 15 true temperature-pressure points. Horizontal dashed lines indicate the pressure levels at which we are retrieving. The thicker the red shading, the more samples that are stacked atop one another.

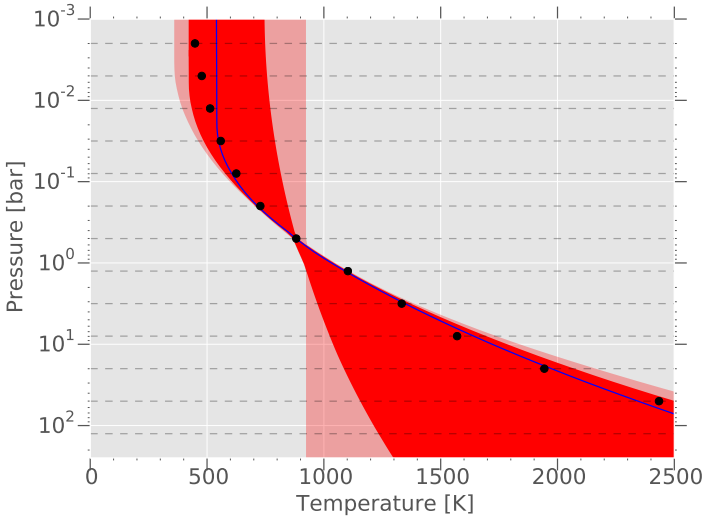

Figure 3.2.4: The retrieved parametric temperature-pressure profile where the black dots indicate the 15 true temperature-pressure points and the blue line indicates the median retrieved profile. Horizontal dashed lines indicate the pressure levels at which we are retrieving. The dark red shading indicates the area in TP-space encompassing $\pm 1\sigma$ and light red shading indicates $\pm 2\sigma$.

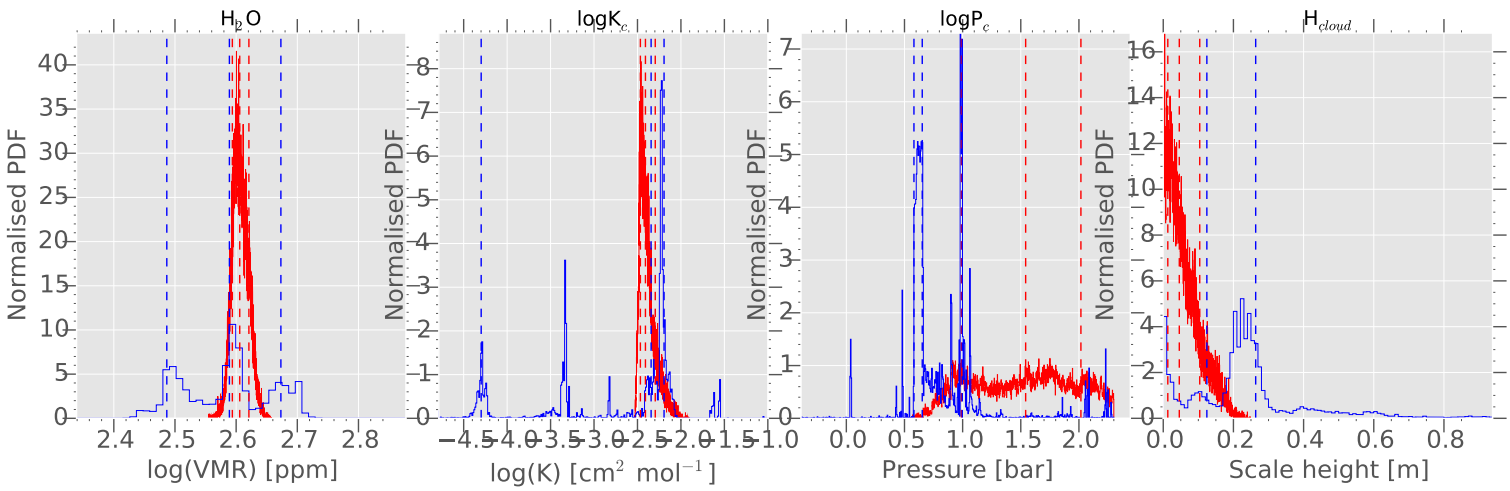

Figure 3.2.5: The marginalised posteriors for the parameters retrieved using the parametric TP-profile (blue) and the 'true' profile (red) are shown, from left-to-right: VMR of H_2O, the log of the opacity of the cloud (K_c), the log of the base pressure (P_c), and the scale height of the cloud (H_{cloud}).

of curves are clearly the most physical and overlap the black dots (true profile) closely, there is a multi-modal distribution apparent in the retrieval, which is typically due to the degeneracy between temperature and a number of other parameters. When we look at the summarised Figure 3.2.4, this multi-modality creates a very blocky shading in error space. To see how this affects the marginalised posteriors of the different retrieved parameters, we again use the uncorrelated layer-by-layer approach in Section 3.1, where we start at the global maximum likelihood (i.e., the true values), to find the 'true information content' of the marginalised posteriors. Note here that fixing the true values includes fixing the temperature profile, and so we are only retrieving the VMR of H_2O and the cloud properties. We compare the marginalised posteriors of the parametric method and the 'true information content' marginalised posteriors from the uncorrelated layer-by-layer method in Figure 3.2.5.

Figure 3.2.5 shows the overlapping marginalised posteriors for the four retrieved parameters of VMR and cloud properties. We can see here that the marginalised posteriors for the parametric approach produces extremely blocky and multi-modal posteriors, while the 'true' profiles exhibit very Gaussian behaviour. It is therefore shown that the error propagation of this parametric method is substantially altering the posteriors, and this should be noted as a major disadvantage of using this approach.

However, we note that convergence using this method is fast, and also manages to produce a perfect fit with correctly retrieved cloud properties and the H_2O abundance. The fit is shown in Figure 3.2.2. Note that the residual is zero because we are retrieving from our own foward model with very little artificial noise.

Bearing this swift convergence but lack of robust error propagation in mind, we suggest and investigate a two-step approach in the next section.

3.3 The Two-Step Approach

In this section we consider the lessons learned from Section 3.1 and Section 3.2 and attempt to combine them in a two-step process for an MCMC retrieval. The idea is simple: first we use the parametric approach to gain a swiftly converged best fit to the data, then assuming that this fit is close to the global maximum likelihood, we reinitialise the MCMC with the best-fit parameters as a starting position for an uncorrelated layer-by-layer approach with 15 temperature points. If the assumption holds that we are at or close to the global maximum likelihood, then we will retrieve the 'true information content' posteriors for the retrieved parameters.

We apply this method to our cloud-free late-T dwarf fake dataset used in Section 3.1, retrieving only 15 temperature points and fixing the other parameters to their true values. Our resulting temperature profile is shown

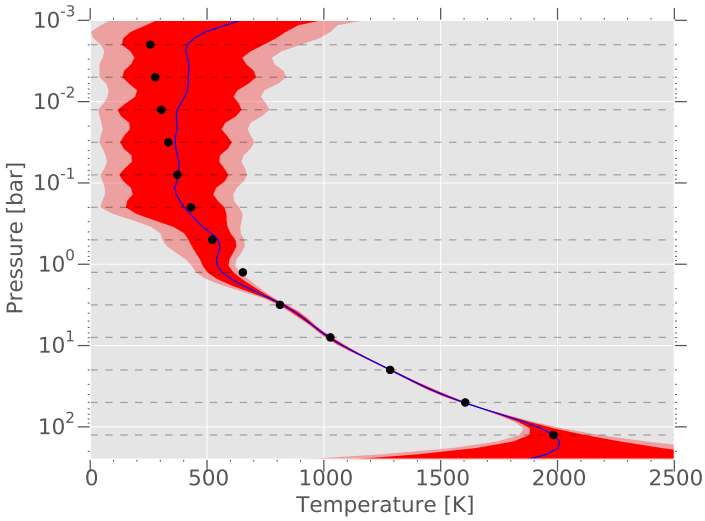

Figure 3.3.1: Two-step approach. The retrieved temperature-pressure profile in the MCMC of a typical late-T dwarf (cloud-free, $T_{eff} \sim 800K$), where the black dots indicate the 15 true temperature-pressure points and the blue line indicates the median retrieved profile. Horizontal dashed lines indicate the pressure levels at which we are retrieving. The dark red shading indicates the area in TP-space encompassing $\pm 1\sigma$ and light red shading indicates $\pm 2\sigma$.

in Figure 3.3.1. This is exactly the same retrieved result as in Section 3.1. While this result is promising, all of the retrievals discussed so far have been performed on fake datasets created by our forward model. In order to verify this method, we must test it on real data.

The two data sets we investigated include the T6 ($\sim$ 800K) dwarf SDSSp J1624+0029 and T4.5 ($\sim$ 1000K) dwarf SDSS J092615.38+584720.9AB, retrieved from the SpeX prism library (Burgasser, 2014). The reduction technique of the two data sets used in this discussion is outlined in Line et al. (2015).

In the first stage of the retrieval, we will be retrieving 15 different parameters. The first four are the temperature parameters describing the parameteric form. We will also be retrieving 11 addditional parameters. These include 4 gases (H_2O, CO, CH_4, Na+K), 3 cloud parameters, (opacity, base pressure, scale height), surface gravity, the squared ratio of the radius to the distance, the shift in wavelength of the model spectrum, and the error inflation parameter which can adjust underestimated errors. The particulars of these various parameters are discussed in Line et al. (2015). For the second stage of the retrieval, we exchange the 4 temperature parameters, which describe the parametric function, for 15 temperature-pressure points

in an uncorrelated layer-by-layer profile. In the second stage, we are thus retrieving 26 parameters. First we will examine the T6 dwarf. The parameters already mentioned have the same priors as before. The log surface gravity is between 0 - 6.0 (cgs), ratio of radius to distance is 0-1 R_{Jup}/pc, the error inflation is uniform such that 0.01 $\min(\sigma_i^2) \leq 10^b \leq 100 \max(\sigma_i^2)$, and the wavelength shift is between -0.01 and 0.01 microns.

Because we are retrieving changing errors due to our error inflation parameter b, our log-posterior must also account for this. The log-likelihood is now given as:

$$\ln\mathcal{L}(\boldsymbol{y}|\boldsymbol{x}) = -\frac{1}{2}\Big(\sum_{i=1}^{n} \frac{(y_i - F_i(\boldsymbol{x}))^2}{{\sigma_i}^2} + \ln(2\pi s_i^2)\Big) \tag{3.5}$$

$$s_i^2 = \sigma_i^2 + 10^b \tag{3.6}$$

where σ is our data error, and b is the error inflation parameter.

Figure 3.3.2 shows the best fit acquired by the parametric approach. We note that the overall fit is good, and that there are a small number of features that are not fit perfectly. Beginning with our best fit parameters as our intialisation parameters, we then perform a layer-by-layer retrieval, producing the posteriors for temperature shown in Figure 3.3.3 and the posteriors for the other retrieved parameters. The best fit spectrum is identical to Figure 3.3.2, so we do not show it here.

The posteriors shown in Figure 3.3.4 are promisingly consistent with the expected values of a T6 dwarf, including reasonable volume mixing ratios, surface gravity, and the non-dectection of a cloud.

We can see in Figure 3.3.3 that the overall shape of the temperature profile is consistent with the original model-based experiments. There are some discrepancies, however, with a notable ringing in the photosphere of the TP-profile. We suspect that the cause of this is due to ill-conditioning, where the temperature profile is adjusting itself to fit all of the parts of the spectrum that were not fit by the parametric model. This likely happens because of the missing physics and/or other assumptions going into our model that we are not directly investigating here. It could also be that this ringing is a real feature of the brown dwarf; however, after we analyse the next object, we will revisit this point and show that it is unlikely to be a real feature.

The best fit for the T6 dwarf by the parametric approach is shown in Figure 3.3.5. We note again that the fit is good, but does not fit every feature.

We show the posteriors in Figure 3.3.7 for our T4 dwarf. The volume mixing ratios and surface gravity are typical for a T4 dwarf with a a thick cloud at 40 bar. We find similar results as for the T6, also with a consistent ringing in the photosphere.

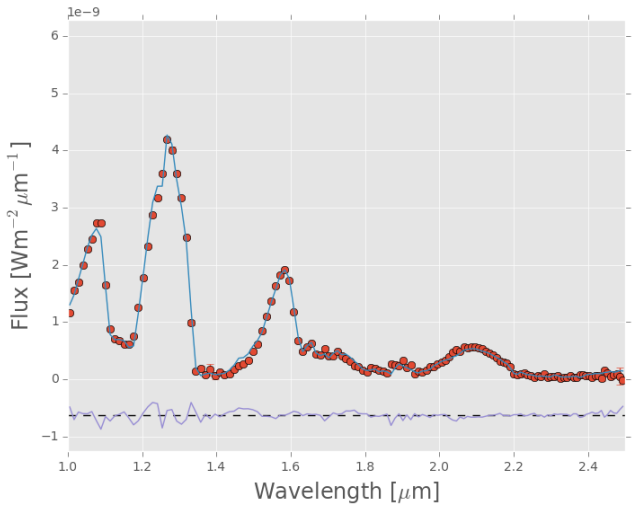

Figure 3.3.2: The best fit (blue) to the T6 data (orange dots) produced by the parametric model. A dotted line shows the residual between the fit and the data.

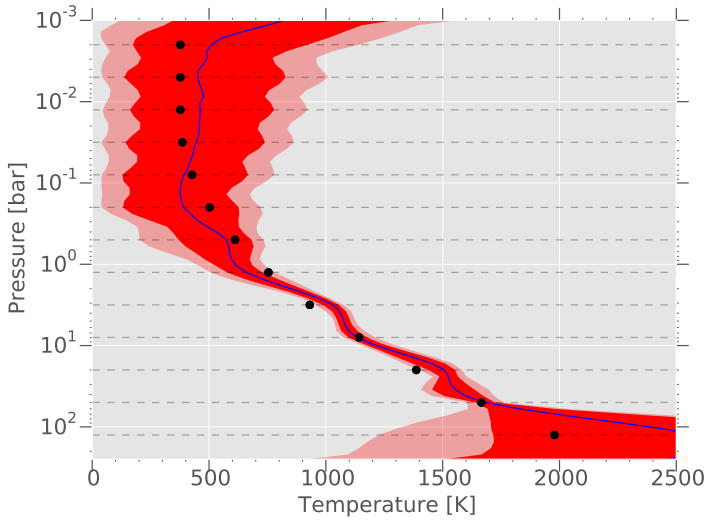

Figure 3.3.3: T6 dwarf. The retrieved two-step temperature-pressure profile where the black dots indicate the 15 parametric temperature-pressure starting positions and the blue line indicates the median retrieved profile. Horizontal dashed lines indicate the pressure levels at which we are retrieving. The dark red shading indicates the area in TP-space encompassing $\pm 1\sigma$ and light red shading indicates $\pm 2\sigma$.

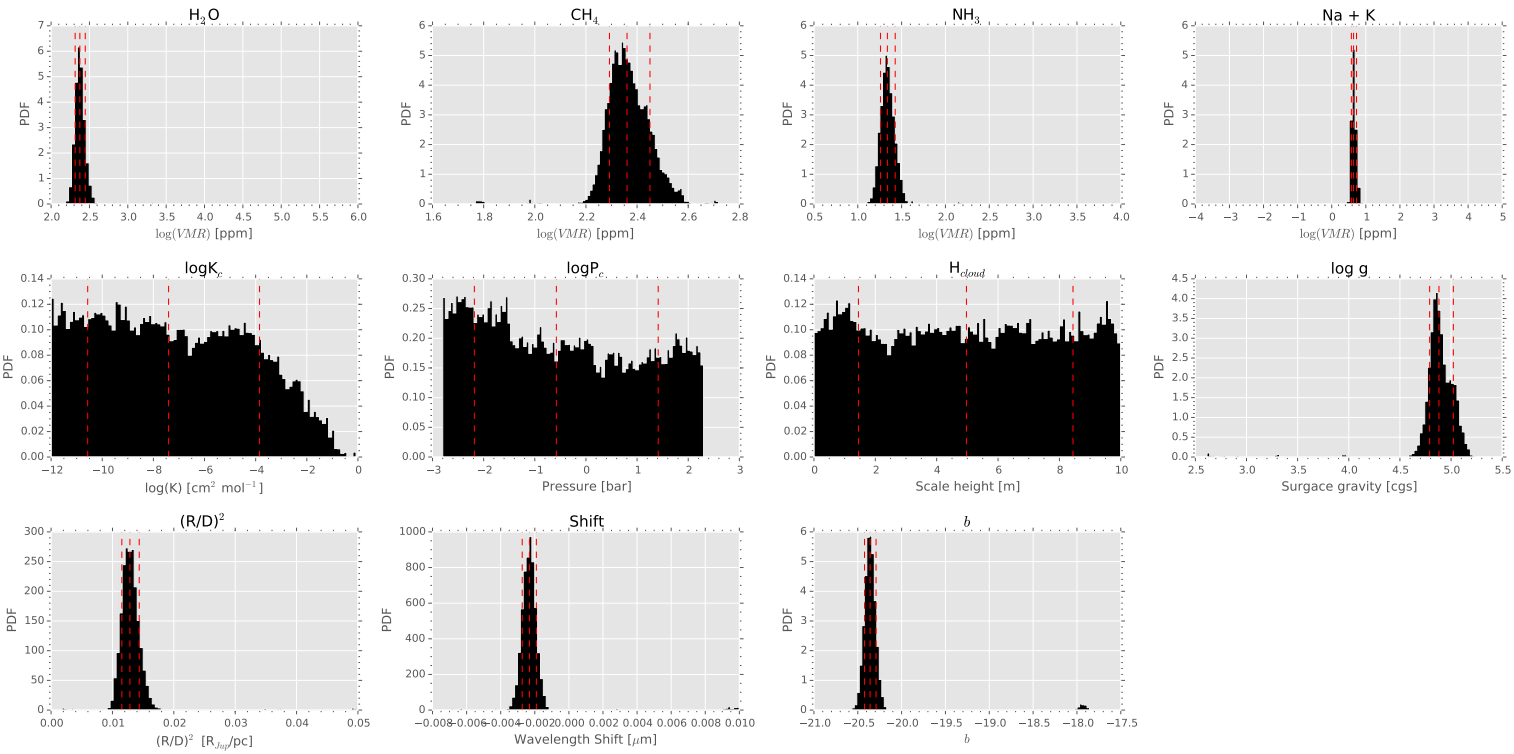

Figure 3.3.4: T6 dwarf. The marginalised posteriors for the parameters retrieved using the two-step TP-profile. Median $\pm 1\sigma$ is indicated by the red dotted lines.

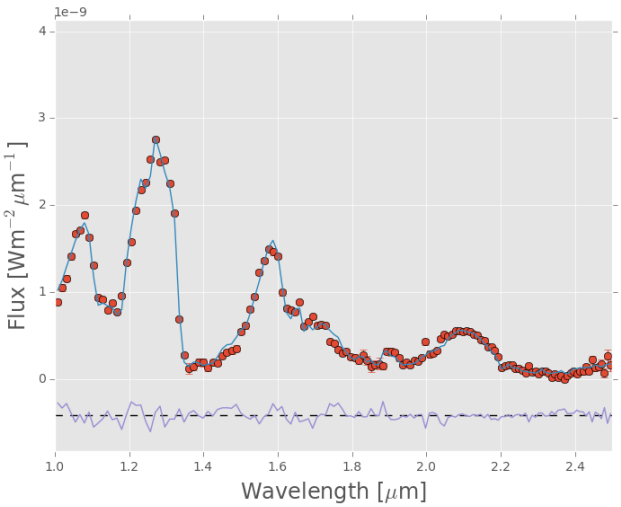

Figure 3.3.5: The best fit (blue) to the T4 data (orange dots) produced by the parametric model. A dotted line shows the residual between the fit and the data.

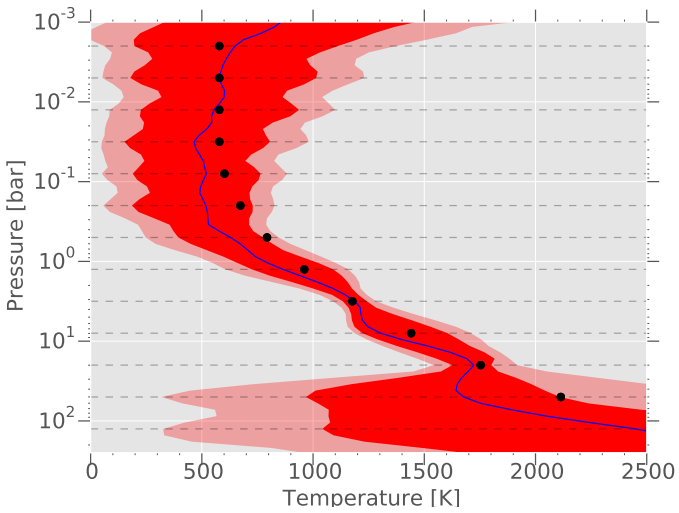

Figure 3.3.6: T4 dwarf. The retrieved two-step temperature-pressure profile where the black dots indicate the 15 parametric temperature-pressure starting positions and the blue line indicates the median retrieved profile. Horizontal dashed lines indicate the pressure levels at which we are retrieving. The dark red shading indicates the area in TP-space encompassing $\pm 1\sigma$ and light red shading indicates $\pm 2\sigma$.

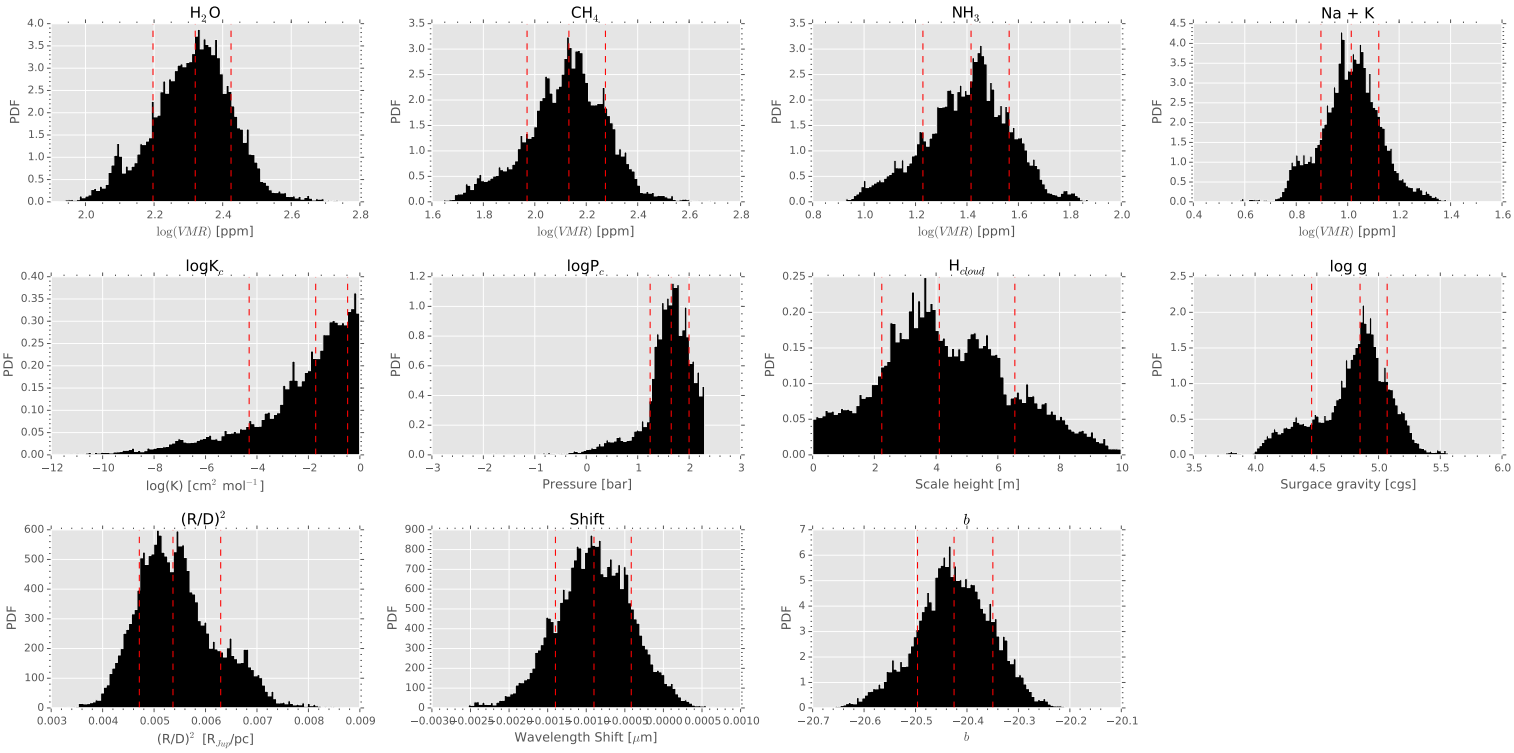

Figure 3.3.7: T4 dwarf. The marginalised posteriors for the parameters retrieved using the two-step TP-profile. Median $\pm 1\sigma$ is indicated by the red dotted lines.

Now that we have noticed the ringing in both of our modelling efforts of real data, we will assess if this is a real feature of brown dwarf atmospheres or a residual of our modelling approach. We do this by redoing our calculation for the T6 dwarf but with double the number of retrieved temperature points.

Now plotting the temperature-pressure profile in Figure 3.3.8 we can see that the profile suffers from even more extreme vertical oscillations, which are a classic indication of ill-conditioning, where each unfitted feature in the parametric spectrum is now being fitted by altering the temperature profile in an unphysical way. Hence, we do not recommend this approach in fitting real data. In the next section we will describe an in-between parameterisation which allows using the layer-by-layer approach in a more physically constrained manner.

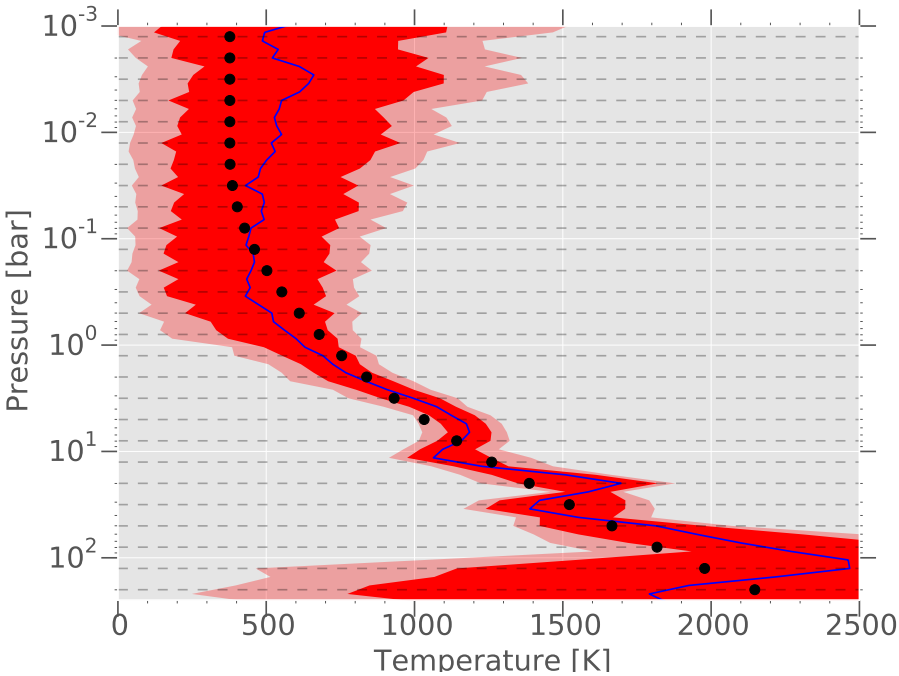

Figure 3.3.8: T6 dwarf. The retrieved two-step temperature-pressure profile where the black dots indicate the 30 parametric temperature-pressure starting positions and the blue line indicates the median retrieved profile. Horizontal dashed lines indicate the pressure levels at which we are retrieving. The dark red shading indicates the area in TP-space encompassing $\pm 1\sigma$ and light red shading indicates $\pm 2\sigma$.

3.4 The Correlated or Gaussian Process Approach

In this section we explore the optimal estimation method employed by Rodgers (2000); Irwin et al. (2008) for temperature-pressure profile parameterisation in an MCMC framework. The cost function (Irwin et al., 2008) is given as:

$$\phi = (\boldsymbol{y_m} - \boldsymbol{y_n})^{\mathsf{T}} \boldsymbol{S_\varepsilon^{-1}} (\boldsymbol{y_m} - \boldsymbol{y_n}) + (\boldsymbol{x_n} - \boldsymbol{x_0})^{\mathsf{T}} \boldsymbol{S_x^{-1}} (\boldsymbol{x_n} - \boldsymbol{x_0}) \tag{3.7}$$

In equation 3.7, $\boldsymbol{y_m}$ is the measured spectrum, $\boldsymbol{y_n}$ is the spectrum calculated for the trial atmosphere, represented by a model state vector $\boldsymbol{x_n}$, $\boldsymbol{S_\varepsilon}$ is the measurement covariance matrix (containing both measurement errors and estimated forward modelling errors), $\boldsymbol{x_0}$ is the *a priori* state vector and $\boldsymbol{S_x}$ is the *a priori* covariance matrix.

We now rewrite this in an unnormalised log-posterior form. Starting from the unnormalised log-form Bayes' theorem:

$$\ln p(\boldsymbol{x}|\boldsymbol{y}) = \ln\mathcal{L}(\boldsymbol{y}|\boldsymbol{x}) + \ln p(\boldsymbol{x}) \tag{3.8}$$

where $\boldsymbol{x}$ is the parameter vector, $\boldsymbol{y}$ is the data vector (i.e. the spectrum), $p(\boldsymbol{x}|\boldsymbol{y})$ is the posterior probability distribution, $\mathcal{L}(\boldsymbol{y}|\boldsymbol{x})$ is the likelihood distribution which penalises poor/improbable fits to the data, $p(\boldsymbol{x})$ is the prior information which can restrict the parameter space.

The likelihood function is given by:

$$\ln\mathcal{L}(\boldsymbol{y}|\boldsymbol{x}) = -\frac{1}{2}\Big(\sum_{i=1}^{n}\frac{(y_i - F_i(\boldsymbol{x}))^2}{{s_i}^2} + \ln(2\pi {s_i}^2)\Big) \tag{3.9}$$

where the index i denotes the ith data point, y is the measured flux, $F(x)$ is the modelled flux that comes from the forward model. s is the data error given by:

$$s_i^2 = \sigma_i^2 + 10^b \tag{3.10}$$

where σ is the measurement error, and b is a free parameter which accounts for missing forward model physics and underestimated uncertainties.

One can briefly note at this point that the first term in equation 3.9 corresponds to a penalty term for a model's failure to fit observed values and the second term to a penalty term that increases proportionally to a model's complexity.

The prior, $p(\boldsymbol{x})$, can be broken up into pieces as

$$\ln p(\boldsymbol{x}) = \ln p(\boldsymbol{T}) + \ln p(\boldsymbol{x'}) \tag{3.11}$$

where $p(\boldsymbol{T})$ is the prior on the temperature profile and $p(\boldsymbol{x'})$ is the prior on all other retrieved parameters. $p(\boldsymbol{x'})$, in our case, is a set of uniform priors on all parameters except temperature. For example, the uniform prior on the mass of a brown dwarf is between 1 and 80 Jupiter masses. Uniform means that any mass between that range is given a probability of one, and anything outside the range is given a probability of zero. $p(\boldsymbol{T})$ is given by:

$$\ln p(\boldsymbol{T}) = -\frac{1}{2}\Big((\boldsymbol{T} - \boldsymbol{T_0})^{\mathsf{T}}\boldsymbol{K^{-1}}(\boldsymbol{T} - \boldsymbol{T_0}) + \ln(\det(\boldsymbol{K})\Big) \tag{3.12}$$

where $\boldsymbol{T_0}$ is the temperature prior and $\boldsymbol{K}$ is the temperature covariance matrix (previously $\boldsymbol{S_x}$), which penalises distancing from the prior, and the second term penalises more complex models. The temperature covariance matrix is given by:

$$\boldsymbol{K} = A\exp\Big(-\frac{\ln(P_i/P_j)}{l}\Big) \tag{3.13}$$

A is a constant error amplitude, P_i is the pressure at level i, and l is the correlation length scale, which decides how correlated in temperature pressure points i and j are.

Combining all of these equations, our joint log-posterior is given by:

$$\ln p(\boldsymbol{x}|\boldsymbol{y}) = -\frac{1}{2}\left(\sum_{i=1}^{n}\frac{(y_i - F_i(\boldsymbol{x}))^2}{{s_i}^2} + (\boldsymbol{T}-\boldsymbol{T_0})^{\mathsf{T}}\boldsymbol{K^{-1}}(\boldsymbol{T}-\boldsymbol{T_0}) + \ln(2\pi {s_i}^2) + \ln(\det(\boldsymbol{K})) + \ln p(\boldsymbol{x'})\right) \tag{3.14}$$

Now that we have our probability equations, we can implement them and assess how well they can fit the models and the robustness of their error propagation. We go back to our model-generated, typical cloudy mid-T dwarf and perform another layer-by-layer calculation using this correlated approach, where now we have added an additional three parameters to the retrieval: the correlation length scale l, the error amplitude A, and the prior temperature T_0, where we assume T_0 is an isotherm. We again initialise at the true answer, so that we are investigating only the influence of our prior rather than its ability to converge.

Figure 3.4.1 shows the summarised temperature posterior for 5000 random samples. The model has the flexibility to fit the TP-space well, in a similar manner to our previous level-by-level approach, but struggles beyond the photosphere to produce a spread which covers the true TP-profile. This is an artifact of using a varying temperature prior which is isothermal. Because of the lack of information beneath the photosphere, the profile relaxes to prior. Because of the strong information in the photosphere (due to peaking weighting function) the TP-profile is driven to encompasses a certain range of temperatures ($\sim 500 - 1500$K) with a strong certainty. This certainty then prefers samples of priors that are within this region, hence the ± 2 sigma is within this range at the bottom of the atmosphere.

To assess if this type of unphysical bending towards an isothermal prior affects the retrieved values of the other parameters, we examine the marginalised posteriors in Figure 3.4.2, comparing them to the uncorrelated layer-by-layer approach. We see that overall, most of the posteriors are the same, except for the retrieved cloud height, which has moved a number of bar higher into the atmosphere. This therefore means that this isothermal prior is indeed affecting the retrieval of other parameters, because it is robustly propogating a certainty in temperature which is incorrect in the deeper portion of the atmosphere. Unfortunately, we did not have time to attempt to fix this issue by investigating different types of priors. In the next section, we will discuss the future work that will be to investigate this and other

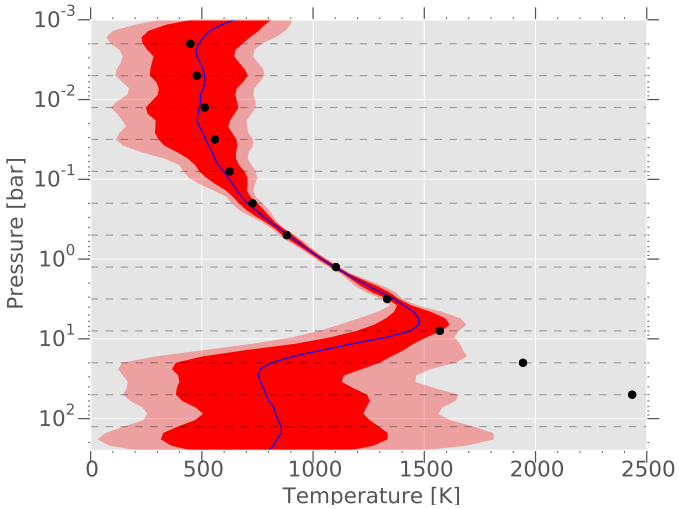

Figure 3.4.1: Correlated layer-by-layer approach. The retrieved correlated temperature-pressure profile where the black dots indicate the 15 parametric temperature-pressure starting/true profile points and the blue line indicates the median retrieved profile. Horizontal dashed lines indicate the pressure levels at which we are retrieving. The dark red shading indicates the area in TP-space encompassing $\pm 1\sigma$ and light red shading indicates $\pm 2\sigma$.

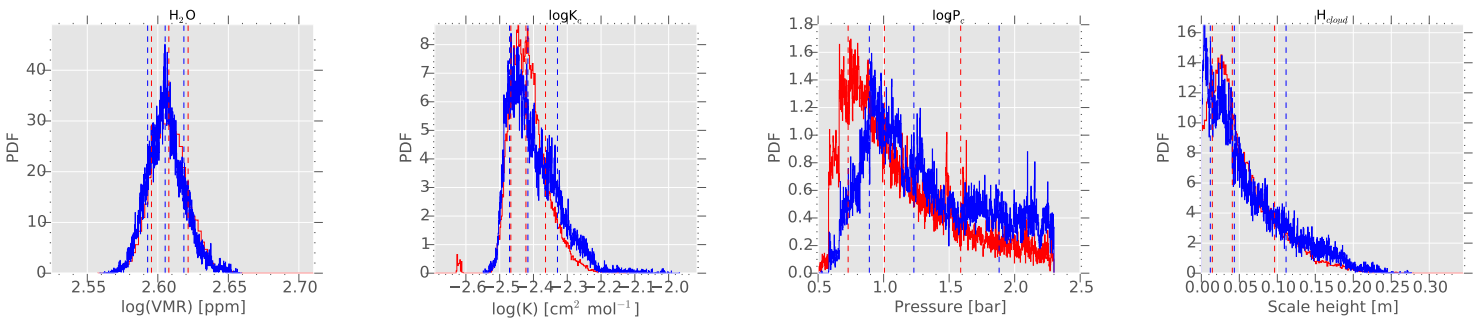

Figure 3.4.2: The marginalised posteriors for the parameters retrieved using the correlated TP-profile are shown in red, with the 'true information content' posteriors in blue. Median $\pm 1\sigma$ is indicated by the red/blue dotted lines.

temperature-pressure profile parameterisations.

4

Future Work

In the future we would like to futher explore this type of comparison with a variety of different parameterisations, ultimately in search for the perfect method that will work on all types of planetary atmospheres. The first among this will be to use a self-consistent grid model, such as the 1D radiative-convective Saumon & Marley (2008) models, to physically constrain the atmosphere from first principles. This method only contains two retrieved parameters, T_{eff} and $\log g$, and should therefore converge very quickly. The forseeable downside is that we are making the assumption that all of the first-principle physics is correct in this model.

Another very simple parameterisation to consider is the Eddington approximation for a non-irradiated, grey atmosphere (Parmentier & Guillot, 2014) which again is only described by two parameters:

$$T^4 = \frac{3}{4}T_{int}^4\Big(\frac{2}{3} + \tau\Big) = \frac{3}{4}T_{int}^4\Big(\frac{2}{3} + \frac{\kappa P}{g}\Big) \tag{4.1}$$

where T is the temperature at pressure level P, T_{int} in the internal effective temperature, τ is the optical depth, κ is the opacity, and g is gravitational acceleration. T_{int} and κ are our retrieved parameters. Again this parameterisation will converge quickly but is maybe too simple for complex atmospheres with sufficient resolution data.

The final goal would be to utilise the two-step approach starting with the best parametric approach, followed by the layer-by-layer method demonstrated in Section 3.4, where the temperature prior T_0, is instead given by,

say, $T_0(T_{eff}, \log g)$, as in the Saumon & Marley (2008) models.

5

Acknowledgements

We would like to thank first and foremost the Kavli Summer Program in Astrophysics for a world-class, once-in-a-lifetime opportunity. We thank The Scientific Organizing Committee for the wonderful set of inspirational lectures given throughout the program. Thank you to the Local Organizing Committee for overseeing our very pleasant stay at the University of California, Santa Cruz. We thank the many insightful and engrossing lectures from the invited speakers. We thank the supervisors for their stellar supervision. Finally, we thank the Kavli Fellows themselves, for the great discussions, friendships, and wonderful memories.

Bibliography

Agúndez, M., Parmentier, V., Venot, O., Hersant, F., & Selsis, F. 2014, A&A, 564, A73

Artigau, É., Bouchard, S., Doyon, R., & Lafrenière, D. 2009, ApJ, 701, 1534

Burgasser, A. J. 2014, in Astronomical Society of India Conference Series, Vol. 11, Astronomical Society of India Conference Series

Feng, Y. K., Line, M. R., Fortney, J. J., et al. 2016, ArXiv e-prints, arXiv:1607.03230

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

Irwin, P. G. J., Teanby, N. A., de Kok, R., et al. 2008, J. Quant. Spec. Radiat. Transf., 109, 1136

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. 2016, ArXiv e-prints, arXiv:1605.08810

López-Puertas, M., & Taylor, F. 2001, Non-LTE Radiative Transfer in the Atmosphere, Series on Atmospheric, Oceanic and Planetary Physics

Madhusudhan, N., & Seager, S. 2009, ApJ, 707, 24

Margolis, J. S. 1996, J. Quant. Spec. Radiat. Transf., 55, 823

Parmentier, V., & Guillot, T. 2014, A&A, 562, A133

Robinson, T. D., & Catling, D. C. 2012, ApJ, 757, 104

Rodgers, C. 2000, Inverse Methods for Atmospheric Sounding: Theory and Practice, Series on Atmospheric, Oceanic and Planetary Physics (World Scientific)

Saumon, D., & Marley, M. S. 2008, ApJ, 689, 1327

Sharp, C. M., & Burrows, A. 2007, ApJS, 168, 140

Yurchenko, S. N., & Tennyson, J. 2014, MNRAS, 440, 1649