UiO : Institute of Theoretical Astrophysics
University of Oslo

Elements of Monte Carlo Radiative Transfer in Stellar Atmospheres

Ida Risnes Hansen

Abstract

Radiative transfer calculations are an essential part of modelling in stellar atmospheres. While the standard integral methods used today have seen great success over the last decades, they do have their limitations. In particular, they tend to be slow when applied to 3D non-LTE problems. Meanwhile, Monte Carlo radiative transfer is used many other places in astrophysics, where it is the preferred choice in 3D, scattering-dominated regimes; the areas where traditional methods fall short.

In this thesis, we test Monte Carlo radiative transfer simulations on a solar atmosphere model. One of the drawbacks of the method, is its inefficiency in high optical depth regions. To overcome this, we introduce a boundary condition that excludes thermal regions in the lower atmosphere. With some adjustment, we are able to efficiently calculate the mean radiation field for all the wavelengths covering the bound-free and bound-bound transitions of a 2-level plus continuum hydrogen model atom.

We then let this be part of an iterative radiative transfer calculation, where we use the statistical equilibrium equations to update the atom populations, and the Monte Carlo simulation to update the radiation field until the populations converge. Doing this, we achieve convergence in 9 iterations. The resulting populations and radiation field are comparable to results from an integral method, but differ from it in ways that can be traced back to the shortcomings of the method.

Acknowledgments

First, I want to thank ITA and RoCS for the opportunity to be a tiny part of what you do, and for being so welcoming. In particular my supervisor, Tiago Pereira, for all your kindness, patience, and enthusiasm, and for letting me work with you on such a fun project.

I am also grateful to the students in the cellar for all the coffee and long lunches. You are lovely. A special thanks to Christina, Jessie and Jowita for making Blindern feel like home, but just as much for dragging me off campus (everyone should have a Jessie in their lives, forcing them to bake bread and paint their nails). I also want to thank my other friends, especially Siriann, Cecilie, Håvard, Elif, HoWan, Aina Marie and Marita, for all the late night conversations and encouragement.

A big thanks to my brothers for never letting me forget that I am an idiot, my sister for reminding me not to take everything so seriously and all three of you for always being around. Finally, I cannot thank my parents and grandparents enough for all your love, relentless support and for always making me laugh (mostly with you).

You have all been stellar.

Contents

Chapter 1

Introduction

1.1 Motivation

Separated from their subject matter by astronomical units, if not light years, solar and stellar physicists are left to make theoretical models of the objects and phenomena they study. To test the rigour of these models, it is often necessary to compute the predicted spectra of the stars and compare them to observations.

Consequently, the clue to so much of what we know about the sun and other stars, is understanding how photons and matter interact in their outer layers. The two are deeply connected. This is fortunate, as it gives us access to a vast amount of information. However, it is also unfortunate, as it makes untangling this information exceedingly difficult.

At the heart of the problem is the co-dependence of the radiation field and the state of the matter. Photons are taken and added to the field at rates determined by the particle states. On the other hand, in regions outside of local thermodynamic equilibrium (LTE), these states depend on the radiation field. Adding to the problem in these non-LTE regions, is that photons can travel far.

Computationally, the consequence of the non-local nature of the problem is that the radiation field in each spatial point depends on the field in many other points. Meanwhile, the non-linearity demands an iterative solution approach.

The combination of these two aspects quickly become expensive for 3D atmosphere models at high resolution. Not to mention that the problem often needs to be solved for many wavelengths and include complicated physics.

While non-LTE radiative transfer calculations in stellar atmospheres remains a daunting problem, the field has a proud tradition going back more than half a century (see review by Hubeny & Lanz 1996). In the decades since, it has grown in sophistication. From 1D problems with limited line-blanketing (Auer & Mihalas 1969), to fully-blanketed 3D calculations with partial redistribution (Sukhorukov & Leenaarts 2016).

While the field has had great success and today's standard methods put great care into computational efficiency and accuracy, the increase in complexity and resolution

has not come for free. Sukhorukov & Leenaarts reports 50 000-200 000 CPU hours before convergence for an atmosphere model with spatial resolution around 100 km. Meanwhile, state of the art telescopes today can resolve structures as small as 22 km at certain wavelengths (Tritschler et al. 2016). While this gap in resolution has been addressed Bjørgen & Leenaarts (2017), radiative transfer calculations are still lagging behind. For this reason, it is important that the methods keep developing.

If we look to other fields in astronomy, Monte Carlo methods are widely used in radiation transport problems (see the review by Whitney 2011). Here, the idea is to follow individual energy packets in a random walk through the medium and let them interact according to probabilities found from opacity calculations. With enough packets, it is possible to get a good representation of the radiation field. Notably, Monte Carlo radiative transfer (MCRT) methods excel at 3D non-LTE problems. The statistical approach of MCRT means that the results will be subject to random noise. While this can be combated by increasing the number of packets, this does not come for free and may be computationally expensive. However, the largest drawback of MCRT is perhaps its poor performance in optically thick regions, where it tends to be both inefficient and inaccurate (Camps & Baes 2018).

As most computational methods, MCRT has both drawbacks and strengths. Whether the benefits outweighs the downsides in solar-type atmosphere problems needs to be tested. That is what we attempt to do in this thesis.

1.2 Problem statement

The aim of this thesis is to test whether a Monte Carlo radiative transfer approach can be used in model atmospheres of cool stars. The basics of the method already exists and are well-tested. What is new in this thesis, is applying it to a solar atmosphere model, and addressing the problems specific to this case. In particular, we work to overcome the highest optical depth regions in the lower atmosphere, and look for the most efficient ways to distribute packets between different wavelengths.

When we have addressed the challenges above, we use the Monte Carlo simulation to estimate the angle averaged radiation field at the wavelengths covering the transitions of a model atom. Using statistical equilibrium, we then update the level populations of the atom, re-configure the simulation set-up, and run the Monte Carlo simulation again. This iterative scheme is then continued for a maximum number of iterations, or until the populations converge. This approach follows the principles of the traditional methods. However, it will be an important test of the global success of the Monte Carlo approach. For validation, the final populations will be compared to the results from an integral method.

This is a very preliminary study. The goal is not to develop a complete and competitive radiative transfer code, but rather to explore its potential and discuss its feasibility.

1.3 Outline

Chapter 2 is dedicated to setting up the necessary background. It includes the most relevant concepts from both radiative transfer and Monte Carlo methods. Chapter 3 presents the method. It particularly focuses on the high optical depth problem and the distribution of packets. Chapter 4 presents all results. Here, the main weight is put on determining simulation configurations that works for different wavelengths, before presenting results from the full radiative transfer calculations. These are then compared to the corresponding results from integral methods. Finally, the significance, and shortcomings of the results are discussed in Chapter 5.

Chapter 2

Theory

2.1 Radiative transfer

Below follows a brief review of the some of the most central concepts in radiative transfer. Although, with a particular focus on concepts and framings relevant to solar-type atmospheres, and Monte Carlo radiative transfer. This section will roughly follow Rutten's notes (2003). For simplicity, we will not consider time evolution or polarisation, and we assume all scattering is isotropic and coherent.

While most of the description below is general and can be applied to many objects and phenomena, it is useful to keep in mind the specific question we ask in solar radiative transfer: given an atmosphere model, what will the escaping radiation look like? Comparing this synthetic intensity to observations can be a measure of how good the atmosphere model is, and improving radiative transfer methods is a way of making this answer more reliable.

2.1.1 Transport equation

The radiation field is built up of photon packets with various energies, travelling in different directions. Over time, packets are taken and added to the field, and change direction and energies due to particle interactions. With a wide range of particle species and many different types of interactions, this quickly becomes messy. For this reason, it is convenient to consider the average state of the matter and the photons.

To describe the radiation field, we use the intensity, I_λ. This can formally be defined in terms of the energy transported by photons in the wavelength band $\lambda + d\lambda$, over a time interval dt, at the position $\vec{r}$, over a solid angle $d\Omega$, in the direction $\vec{l}$, through the area dA with the normal $\vec{n}$,

$$dE_\lambda \equiv I_\lambda(\vec{r}, \vec{l}, t)\,(\vec{l} \cdot \vec{n})\,dA\,dt d\lambda\,d\Omega. \tag{2.1}$$

The intensity has units of J s^{-1} nm^{-1} m^{-2} sr^{-1}, and is the protagonist of radiative transfer. In the context of stellar atmosphere radiative transfer, it is the outgoing intensity we want to calculate, and compare to observations.

To find how the intensity changes throughout the atmosphere, we need to know the energy added and taken from the field at the different wavelengths and in different directions. This is determined by the sum of all the different interactions between photons and other particles. Macroscopically, these processes are collected in the emissivity j_λ and the extinction coefficient α_λ.

The emissivity in a region is proportional to the intensity added to a beam going through it. The intensity added to the field over a path ds in a given direction the region is,

$$dI_\lambda = j_\lambda ds, \tag{2.2}$$

where the units of the emissivity are J s^{-1} nm^{-1} m^{-3} sr^{-1}. Likewise, the intensity taken from the field is proportional to the extinction over this path. However, it is also proportional to the intensity itself. The extinction coefficient α_λ is therefore defined by

$$dI_\lambda = -\alpha_\lambda I_\lambda ds, \tag{2.3}$$

and has units of m^{-1}. In sum, the change in intensity over a path ds in a spatial region is the difference between the emission and extinction processes in the region,

$$dI_\lambda = j_\lambda ds - \alpha_\lambda I_\lambda ds. \tag{2.4}$$

This is the complete transport equation. It is easy to be deceived by its conciseness. However, a lot of complicated physics are hidden in the emissivity and the extinction. The rates of the interactions they describe depend on the state of the particles. Simultaneously, these particle states depend on the field itself. Thus, the problem is non-linear and will usually require an iterative solution approach that starts from an initial guess of either the particle states or the field. On top of this non-linearity, photons can travel far in the thinner parts of the atmosphere, making the problem non-local as well. This means that inside each iteration, the equation must be solved for many wavelengths and in many directions for each position.

A large part of any solution routine is to calculate the extinction and emission given some particle states. They depend on the same particle configurations, and are often tightly connected. For this reason, it is often convenient to write the transport equation on a slightly different form. We define the dimensionless quantity optical depth τ_λ,

$$d\tau_\lambda = -\alpha_\lambda ds, \tag{2.5}$$

and the source function,

$$S_\lambda = \frac{j_\lambda}{\alpha_\lambda}, \tag{2.6}$$

with the same units as the intensity. Equation (2.4) can then be written,

$$\frac{dI}{d\tau} = I_\lambda - S_\lambda. \tag{2.7}$$

While S_λ and τ_λ are still clearly dependent variables, they are less so than α_λ and j_λ (Rutten 2003, pg. 14).

2.1.2 LTE vs. non-LTE

The lower regions of the solar atmosphere are dense and the optical depth is high for most wavelengths. We say the medium becomes *optically thick*, and the matter is in local thermodynamic equilibrium (LTE). This means that most energy transition processes are caused by collisions, and the state of the matter and radiation field are determined by the local temperature. In particular, the source function is given by the Planck function,

$$B_\lambda \equiv \frac{2hc^2}{\lambda^5} \frac{1}{e^{hc/kT\lambda} - 1}. \tag{2.8}$$

Higher up in the chromosphere, the density decreases and collisions no longer dominate. Instead, energy transitions caused by the radiation field start to become important. Moreover, photons travel farther in this optically thick regime, and the problem again becomes non-local.

The nature of photon-particle interactions is subtle, and the terms *scattering* and *destruction* are not very precise. However, the distinction becomes useful in radiative transfer applications, in particular in the discussion of non-LTE. We say that a photon is scattered if it is absorbed and immediately re-emitted by a particle, not necessarily with the same energy. On the other hand, if the particle returns to its initial state without emitting a photon, we say that the photon is destroyed. Based on this we can the define the photon destruction probability,

$$\varepsilon_\lambda \equiv \frac{\alpha^a_\lambda}{\alpha^a_\lambda + \alpha^s_\lambda}, \tag{2.9}$$

where α^a_λ is the sum of all extinction from destruction processes, while α^a_λ is the extinction from scattering. On average, a photon will scatter $N_s = \varepsilon_\lambda^{-1}$ times before it is destroyed.

Another convenient quantity is the mean free path,

$$l_\lambda \equiv \alpha_\lambda^{-1}. \tag{2.10}$$

This is the average distance a photon will travel before it interacts with a particle. It is possible to show that after N_s scatterings, a photon will on average have travelled an effective path length $l^*_\lambda \approx \sqrt{N_s} l_\lambda$.

2.1.3 Source function

While the source function for collisional processes is still given by the Planck function, the one for purely scattering processes will depend on all available photons, and is given by the (direction-averaged) mean intensity,

$$J_\lambda = \frac{1}{4\pi} \int I_\lambda d\Omega, \tag{2.11}$$

with the same units as intensity. For a medium with two-level atoms and where the energy of the photons are preserved in all scatterings (coherent scattering), we can define the source function,

$$S_\lambda = \frac{\sum j_\lambda}{\sum \alpha_\lambda} = \frac{\alpha^a_\lambda B_\lambda + \alpha^s_\lambda J_\lambda}{\alpha^a + \alpha^s} = \varepsilon_\lambda B_\lambda + (1 - \varepsilon_\lambda) J_\lambda. \tag{2.12}$$

This clearly shows how the source functions becomes more locally determined in areas where collisions dominate. In situations with multi-level atoms or non-coherent scattering, (2.12) no longer strictly holds. In these cases, absorbed photons can be re-emitted at different wavelengths, and the source function will depend on quantities at multiple wavelengths.

2.1.4 Cross-talk between wavelengths

While we have already stressed the non-local and non-linear properties of radiative transfer, the redistribution of wavelengths in the different transitions is a further entanglement of the problem.

In multi-level atoms, electrons do not necessarily transition directly between two states. Instead, they can go via several levels and in the process emit lower energy photons. This is known as photon conversions, and couples different energy photons together.

However, even without photon conversion, wavelength redistribution will happen within the separate transition processes. Coherent scattering does not hold in most cases. A more realistic model is to assume complete redistribution (CRD) over the transitions. This assumption assures that the emission profiles are equal to the extinction profiles. This way, a version of (2.12) still holds, but integrated over the entire transition. In reality, most transitions are only partially redistributed (PRD). This is especially true for strong transition. For the sake of simplicity we will assume CRD and not PRD in this thesis.

In addition to wavelength redistribution, anisotropy and polarisation are important topics for realistic radiation transport. However, these will not be dealt with in this thesis.

2.1.5 Extinction processes

While the macroscopic description in the above sections makes for compact notation, computing the total extinction for each wavelength can be a comprehensive task. One reason for this, is the many different ways photons and particles can interact. We will not go into great detail, but for the purpose of this thesis, it is useful to mention the most important processes and emphasise the distinction between the destruction processes and scattering processes.

To illustrate the different types, we consider the simple case of an atomic transition with a single electron. However, the examples can readily be extended to more complicated cases.

Bound-bound transitions

A bound-bound transition is between a bound state and another bound state in a particle. It can move to a higher energy state due to a collision or after absorbing a photon at the right wavelength. We call this collisional and radiative excitations respectively. From a higher level, it can be de-excited spontaneously or collisionally, or the de-excitation can be induced by radiation. For each particle species, these five processes are associated with Einstein coefficients that gives the rate of transitions per particle. For the excitations,

$$B_{lu}\overline{J}_{\lambda_0} = \text{rate of radiative excitations from state l to u per particle in state l} \tag{2.13}$$

$$C_{lu} = \text{rate of collisional excitations from state l to u per particle in state l} \tag{2.14}$$

and for the de-excitations

$$A_{ul} = \text{rate of spontaneous de-excitations from state u to l per particle in state u} \tag{2.15}$$

$$B_{ul}\overline{J}_{\lambda_0} = \text{rate of radiatively induced de-excitations from state u to l per particle in state u} \tag{2.16}$$

$$C_{ul} = \text{rate of collisional de-excitations from state u to l per particle in state u} \tag{2.17}$$

$$\tag{2.18}$$

$\overline{J}_{\lambda_0}$ is defined as the direction averaged radiation field integrated over all wavelengths in the line. An extinction process comes about when a photon is absorbed by a particle in a radiative excitation. It can be shown that the line extinction

$$\alpha_\lambda^{line} = \frac{hc}{4\pi\lambda_0}\varphi(n_l B_{lu} - n_u B_{ul}) \tag{2.19}$$

where λ_0 is the line center wavelength, n_l and n_u is the number of particles in state l and u respectively, and φ is the line profile determined by broadening processes. The term $n_u B_{ul}$ is a correction for radiatively stimulated emission. Conventionally it is added here, because it has a very similar form to the the radiative excitation.

We can use these coefficients to define the line destruction probability. Any extinction process that is followed by a collisional de-excitation will not emit any photons, and is a destruction process. On the other hand, both the spontaneous de-excitation and induced de-excitation gives a photon. If we define the radiative transition rate $R_{ul} = A_{ul} + B_{ul}$, we can write the line destruction probability,

$$\varepsilon_\lambda^{line} = \frac{C_{ul}}{C_{ul} + R_{ul}}. \tag{2.20}$$

The shape of the line profile can be influenced by many processes. Some of the most important are natural damping due to the uncertainty in the energy state of a particle, collisional damping, broadening from thermal motion, and large scale motions in the atmosphere. Thermal broadening usually dominates in the core, and has a Gaussian profile. Meanwhile, collisional and radiative damping give a Lorentz shape. The convolution of these two gives the Voigt function, with a Doppler-core and Lorentzian wings.

$$H(a, v) \equiv \frac{a}{\pi}\int_{-\infty}^{\infty} \frac{e^{-y^2}}{(v-y)^2 + a^2} dy \tag{2.21}$$

$$v \equiv \frac{\lambda - \lambda_0 + \frac{v_0}{c}\lambda_0}{\Delta\lambda_D} \tag{2.22}$$

$$a \equiv \frac{\lambda^2\gamma}{4\pi\Delta\lambda_D} \tag{2.23}$$

where γ is the damping parameter, v_{los} is the line-of-sight velocity of the local medium, and $\Delta\lambda_D \equiv \lambda_0/c\sqrt{2kT/m}$ is the Doppler width.

Bound-free transitions

In atomic bound-free transitions, an electron is either removed from a particle, or recombines with it. There are five possible transition processes that can be viewed as parallels to the ones described for the bound-bound transitions. The particle can be ionised collisionally or by a photon with energy higher than the ionisation threshold. Meanwhile, a recombination can be induced by a photon, happen spontaneously, or collisionally.

We can define radiative and collisional transition rates corresponding to those in (2.13)- (2.17), but here, the upper state u is always the ionised state, while $\overline{J}_{\lambda_0}$ is integrated over the bound-free transition, instead of the line profile. Thus, a destruction probability on the same form as (2.20) is valid here as well.

For hydrogen-like species, the extinction drops with the cube of the wavelength moving away from the ionisation limit (also known as the ionisation edge),

$$\alpha_\lambda^{bf} \propto \left(\frac{\lambda}{\lambda_{edge}}\right)^3 n_i(1 - e^{-hc/\lambda kT})g_{bf}. \tag{2.24}$$

Here, n_i is the density of the particle in state i, and g_{bf} is the bound-free Gaunt factor that corrects for quantum mechanical effects. The terms in the parenthesis correct for induced photo-recombination, similarly to the $n_u B_{ul}$ correction in (2.19).

Free-free transitions

In free-free processes, electrons emit or absorb a photon due to the influence of an ion. In an emission process, the electron is deflected and slowed down by the electric field of the ion, and emits a photon to compensate for the loss in kinetic energy. Conversely, an electron can absorb a photon and gain kinetic energy while close to the ion. The free-free extinction,

$$\alpha_\lambda^{ff} \propto \frac{n_e n_{ion}\lambda^3}{\sqrt{T}c^3}(1 - e^{-hc/\lambda kT})g_{ff}, \tag{2.25}$$

where n_e and n_{ion} is the electron and ion density respectively. g_{ff} is the free-free Gaunt factor, while the parenthesis again correct for induced emission. In cases with Maxwellian velocity distributions, it is a good approximation to say that free-free processes are thermal (Rutten 2003, pg. 27).

Elastic scattering

It is also possible for an electron to scatter directly with electrons, without any change in wavelength. Extinction from these processes can be viewed as pure scattering.

If the electron is free, it is called Thomson scattering,

$$\alpha_\lambda^T = \sigma^T n_e, \qquad \sigma^T = \frac{8\pi}{3}r_e^2 \tag{2.26}$$

where r_e is the electron radius. This type of scattering is independent of wavelength. Meanwhile, if the electron is bound to a particle and has an ionisation energy $E = hc/\lambda_0$, where the wavelength λ_0 is much larger than the particle, it is called Rayleigh scattering.

$$\alpha_\lambda^R = f_{lu}\sigma_\lambda^T \left(\frac{\lambda_0}{\lambda}\right)^4 n_P, \tag{2.27}$$

where f_{lu} is the oscillator strength, and n_P the density of the particle. This type of scattering is stronger at shorter wavelengths.

2.1.6 Level populations

While the atmosphere model we use may supply the overall density of a particle species, the bound-bound and bound-free processes depend on the specific level population densities of the species. In LTE, these are given by the Saha-Boltzmann distribution,

$$\left[\frac{n_c}{n_i}\right]_{\text{LTE}} = \frac{2g_c}{n_e g_i}\left(\frac{2\pi m_e kT}{h^2}\right)^{2/3} e^{-\chi_{ci}/kT} \tag{2.28}$$

where n_i and n_c are the population densities of level i and the next ionised state c. χ_{ci} is the energy required to go from state i to c, while g_i and g_c are the statistical weights of the two states.

However, outside of LTE, (2.28) does not hold. Instead, we assume statistical equilibrium (SE): that the populations stay constant over time. For a N-level plus continuum atom model, this gives N+1 equations on the form,

$$n_i \sum_{i\neq j} P_{ij} = \sum_{j\neq i} n_j P_{ij}, \tag{2.29}$$

where n_i is the population of state i, and P_{ij} the total transition rate from state i to state j, per particle in state i. For energy states l and u, where $l < u$, the total upward transition rate is given

$$P_{lu} = B_{lu}\overline{J}_{\lambda_0} + C_{lu} = R_{lu} + C_{lu}, \tag{2.30}$$

while the downward rate,

$$P_{ul} = A_{ul} + B_{ul}\overline{J}_{\lambda_0} + C_{ul} = R_{ul} + C_{ul}. \tag{2.31}$$

We have already seen how the radiation field depends on the particle states via the extinction and emission processes. Here, the statistical equilibrium equations connect the two in yet another way. This becomes important in the solution schemes we will discuss next.

2.1.7 Traditional integral methods

Considering the many complexities of radiative transfer, it is perhaps not surprising that analytical solution approaches quickly hit a wall. Since the 1960's, numerical methods have been developed and refined into complex integral schemes.

We will not go into much detail or discuss specific methods, but instead describe some general characteristics common for most solution schemes today.

Boundary problem

To solve for the intensity along some path in a medium, we need to know the boundary conditions. That is, we need to know the incoming and outgoing intensity, in both the beginning and in the end of the path. In solar radiative transfer, the two boundaries are in the top and at the bottom of the atmosphere. In the top, we can put the incoming intensity to zero, but the outgoing intensity is unknown. Meanwhile, in the dense bottom boundary, it is possible to guess the outgoing intensity with diffusion approximations, but we cannot do the same for the incoming intensity. In other words, we lack half of the boundary conditions.

To get around this problem, Feautrier's method introduces a clever way of coupling together the outgoing and incoming intensities (Rutten 2003, p. 117). Given a source function, it is then possible to solve the transport equation for the mean intensity J_λ.

Classical lambda iteration

The lambda operator represents a double integral over direction and depth, and gives the mean radiation field for a given source function,

$$J_\lambda = \int_{-1}^{1} I_\lambda d\mu \equiv \Lambda_\lambda[S_\lambda]. \tag{2.32}$$

Feautrier's method is one way of constructing such an operator. Using the source function for a two-level atom in (2.11), we can write,

$$S_\lambda = (1 - \varepsilon_\lambda)\Lambda_\lambda[S_\lambda] + \varepsilon_\lambda B_\lambda. \tag{2.33}$$

This can then be inverted, so that

$$S_\lambda = (1 - (1 - \varepsilon_\lambda)\Lambda_\lambda)^{-1}[\varepsilon_\lambda B_\lambda]. \tag{2.34}$$

However, an inversion like this can be very computationally expensive. It turns out that it is often faster to instead iterate over the source function,

$$S_\lambda^{(n+1)} = (1 - \varepsilon_\lambda)\Lambda_\lambda[S_\lambda^{(n)}] + \varepsilon_\lambda B_\lambda. \tag{2.35}$$

This is called classical lambda iteration. For a multi-level case where (2.11) does not hold, it is still possible to calculate the radiation field using (2.33), but it requires the extra step of updating the level populations, and using the emissivity and extinction relations to get an updated source function.

While lambda iteration gives a stable way of finding the source function, it tends to converge very slowly, particularly at high optical depths.

Accelerated lambda iteration

The idea behind accelerated lambda iteration (ALI) is to define an approximate operator Λ^*, so that

$$\Lambda_\lambda = \Lambda^* + (\Lambda_\lambda - \Lambda^*). \tag{2.36}$$

It then possible to write the iteration scheme as

$$S_\lambda^{(n+1)} = (1 - (1 - \varepsilon_\lambda)\Lambda^*)^{-1}[(1 - \varepsilon_\lambda)\Lambda_\lambda[S_\lambda^{(n)}] + \varepsilon_\lambda B_\lambda - (1 - \varepsilon_\lambda)\Lambda^*[S_\lambda^{(n)}]] \tag{2.37}$$

and achieve a much faster convergence at high depth. Several such operators Λ^* have been successfully developed (see Rutten 2003, p. 126). The same principles are used in the multi-level case, by linearising all the equations, and decreasing the error over successive iterations. In addition to these concepts, convergence acceleration techniques are widely used to improve the performance of the different methods.

Drawbacks and strengths

The numerical methods characterised in the previous section were first developed for 1D LTE problems. Still today, this is where they really excel. However, as atmosphere models and observations alike become more sophisticated, the idea of stellar radiative transfer as a 1D LTE problem no longer holds. While the methods have been successfully extended to 3D non-LTE problems, they do not yet scale well with dimension or resolution, and can be less reliable here than for LTE problems.

2.2 Monte Carlo radiative transfer

Radiation is our main source of information about the universe, and naturally, radiation transport is a central part of most fields in astronomy. The conditions throughout the universe vary a lot, and many fields deal with low density, scattering-dominated environments with complicated geometries. Here, the numerical methods developed for stellar atmospheres become impractical and unreliable. Conditions like these have motivated a different approach to radiation transport; Monte Carlo radiative transfer. Here, the macroscopic descriptions defined in Section 2.1 are used to simulate the average behaviour of the radiation field.

We give a brief review of the most important concepts used in radiation transport problems. For a more thorough description, see for examples the excellent reviews by Noebauer & Sim (2019) and Whitney (2011). For a more general overview of Monte Carlo methods, see Kalos & Whitlock (1986).

2.2.1 Basic idea

In Monte Carlo methods, statistical sampling is used to solve numerical problems. The problems themselves can be deterministic, but need to have a probabilistic interpretation.

The high-level idea is to use experiments to sample the state of a system, where the probability of different states are known. With enough tests, you should be able to estimate the average state of the system, with an accuracy that increases with the number of samples.

If the probability of different outcomes x of a certain event is given by a known probability distribution function $P(x)$, it is this $P(x)$ we sample from, until we get an estimate of the average outcome of the event. If the possible outcomes $x \in [a, b]$, the probability that the outcome is between a and some value x_0 is given by the cumulative probability distribution,

$$\Psi(x_0) = \frac{\int_a^{x_0} P(x)dx}{\int_a^b P(x)dx}. \tag{2.38}$$

where $\Psi(x_0) \equiv \xi \in [0, 1]$, and is uniform. If we can solve (2.39) for x_0,

$$x_0 = G(\xi), \tag{2.39}$$

where G is some function of ξ, we can effectively sample values x_0 from the PDF by drawing random numbers $\xi \in [0, 1]$ from a uniform distribution.

2.2.2 Radiation transport

Photons and matter interact according to chance, and radiation transport lends itself well to Monte Carlo methods.

Instead of using integral methods to solve for the intensity, the idea is to simulate the propagation of photon packets throughout the medium, and let exchange information with the local field.

Packet creation

While it is natural to think of the simulation as following the paths of individual photons, this is often not practical. The number of photons is simply too large. Instead, we can define larger packets of energy, where the number of packets created in each

region is proportional to the energy emission of that region. The energy of each packet is then

$$E_{\text{packet}} = \frac{E_{\text{total}}}{N_{\text{total}}}, \tag{2.40}$$

where E_{total} is the total energy emission of the medium, and N_{total} is the total number of energy packets created in the simulation. While it is desirable that N_{total} is as large as possible, this approach gives a lot more flexibility.

Propagating the packets

While the distribution of energy packets is determined by the energy emission distribution of the medium, it is in the propagation of packets that we introduce the random sampling. In a random walk, the probability that a photon travels an optical depth τ_λ before it interacts with the medium is

$$P(\tau_\lambda)d\tau_\lambda = e^{-\tau_\lambda}d\tau_\lambda. \tag{2.41}$$

Using (2.39), the cumulative probability distribution is

$$\xi = \frac{\int_0^{\tau_{\lambda,0}} e^{-\tau_\lambda}d\tau_\lambda}{\int_0^{\infty} e^{-\tau_\lambda}d\tau_\lambda} = 1 - e^{-\tau_{\lambda,0}}, \tag{2.42}$$

and we can sample optical depths

$$\tau_{\lambda,0} = -\log(1-\xi). \tag{2.43}$$

where ξ is drawn from a uniform distribution between 0 and 1. If we assume all scattering is isotropic, the direction is drawn from homogeneous distributions,

$$\phi_0 = 2\pi\xi \qquad \theta_0 = \pi\xi. \tag{2.44}$$

We then move the packet an optical depth $\tau_{\lambda,0}$ in the direction $\mathbf{r} = [\cos\phi\,\sin\theta,\ \sin\phi\,\sin\theta,\ \cos\theta]$. At the new position, the packet interacts with a particle. Here, we draw a new random number ξ. If ξ is smaller than the local destruction probability, the packet is destroyed. If, on the other hand, it is larger, the packet is scattered. We draw a new optical depth and direction, and repeat the process. This is continued for a defined maximum number of scattering, until the packet is destroyed, or until it leaves the domain. The specifics can vary, depending on the medium and the goal of the simulation. Either way, this process is repeated for all packets.

2.2.3 Getting the radiation quantities from the simulation

The propagation of the packets is determined by PDFs that depend on the local properties of the medium (τ_λ and ε_λ). The missing step is then how we can use this to determine the radiation field. This is done by letting the packet add to the local energy in every region it goes through. This can then easily be converted to intensity and other related quantities. In the end, with enough samplings, the hope is that we can recover the macroscopic radiation field.

Strengths and drawbacks

Monte Carlo radiative transfer takes you a litter closer to the actual behaviour of photons. It is an intuitive method that is straightforward to implement. Adding more physics to the problem will often just require changing the probability distribution functions. It is not as tied to geometry in the same way integral methods are, and allows much more flexibility. In terms of performance, it excels in 3D scattering problems, and is complementary to integral methods.

Conversely, an important drawback of the method, is its poor performance at high optical depth. Here, the number of interactions increase, and packets can effectively be trapped, and take up a lot of computational time. This in turn, leads to a worse sampling of the remaining field. Camps & Baes (2018) found that Monte Carlo methods can fail to reproduce the radiation field for optical depths higher than $\tau_\lambda = 20$. While there has been several successful attempts to overcome this high optical depth problem (see for example Min et al. 2009 and Robitaille 2015), these methods are not necessarily developed with stellar atmosphere problems in mind. The lower parts of stellar atmospheres reach very high optical depths. A crucial part of applying Monte Carlo methods to stellar atmospheres will therefore be to overcome this problem.

A final point is that because the method is statistical in nature, it will always be subject to random noise. While this can be reduced by increasing the number of samplings or by other noise reduction techniques, this can be computationally expensive. Regardless, in any MCRT application, it is important to take noise into account.

Chapter 3

Method

In this section, we will describe our implementation of Monte Carlo radiative transfer in the solar atmosphere. An overview of the complete approach is shown in Figure 3.1.

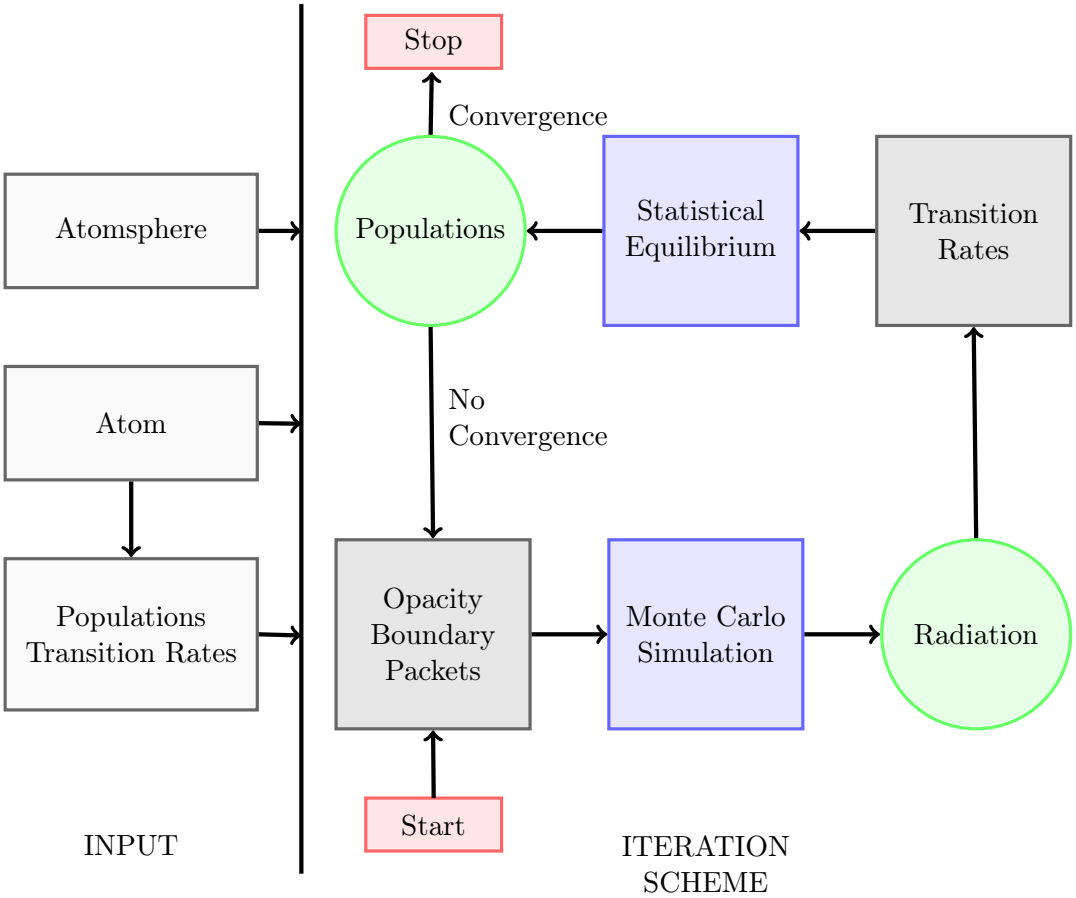

Figure 3.1: Overview of the Monte Carlo radiative transfer method.

An atmosphere model and initial populations of a model atom are taken as input. A wavelength grid is created to cover all the possible atom transitions. At each of these wavelengths, the total extinction and destruction probability are calculated from the initial atom populations. Based on these, the high optical depth regions in the lower atmosphere are excluded from the simulation domain, and energy packets distributed in the remaining atmosphere. The mean radiation field J_λ is then estimated for each λ in a Monte Carlo simulation. New populations are then estimated by solving the statistical equilibrium equations. These are then used to set up a new Monte Carlo simulation that gives a new estimate for J_λ, that again is fed to the statistical equilibrium equations. This is repeated for a set maximum number of iterations, or stopped if the populations converge. Below follows a more detailed description of each step. In Sections 3.1 - 3.4 we describe our implementation of the already existing methods, while we address the issues specific to the solar atmosphere in Section 3.6.

3.1 Input

3.1.1 Atmosphere

The simulation takes a 1D, 2D or 3D atmosphere as input. At a minimum, it must contain temperature, hydrogen density and electron density. Optionally, it can include velocities. The model grid must be rectilinear, and for simplicity, we consider each voxel homogeneous.

To test the code, we use a snapshot from a magnetohydrodynamic quiet Sun simulation, run with the Bifrost code (Gudiksen et al. 2011). It includes velocities, and has (nz,ny,nx)=(72,64,64), the height $z \in$ [-0.8, 8.0]Mm, while the sides are 5.0 Mm.

3.1.2 Model atom and wavelength sampling

In addition to the atmosphere, the code takes as input a model atom with N levels plus continuum. Apart from the atomic weight and charge, we need to know the energy and statistical weight of each level. If we are sampling lines, we also need the oscillator strength of each line.

The possible atom transitions determine the wavelengths to sample the radiation field. For the bound-free transitions, we sample linearly from the ionisation edge to the minimum wavelength defined $\lambda_{min,n} = \lambda_{max,1}(n/2)^2$ where $\lambda_{max,1}$ is the ionisation edge of the ground level and n is the level number. For the bound-bound transitions, we do a logarithmic sampling inspired by the RH code (Uitenbroek 2001), so that we have more sampling points closer to the line-core. Because we have an atmosphere with turbulence, it is necessary to sample the full line. In the end, for a N-level atom, we are left with $N(N+1)/2$ sampling regions that may or may not be overlapping.

In our tests, we use a minimal (2+1)-level hydrogen model, with atomic level data from the NIST Atomic Spectra Database (Kramida et al. 2020). Figure 3.2 shows the extinction coefficient α_λ plotted in a voxel at the bottom of the atmosphere for the

different wavelengths we will sample the field at. The yellow points show wavelengths we will use to set up and test the simulations. We later refer to them as *test-wavelengths*.

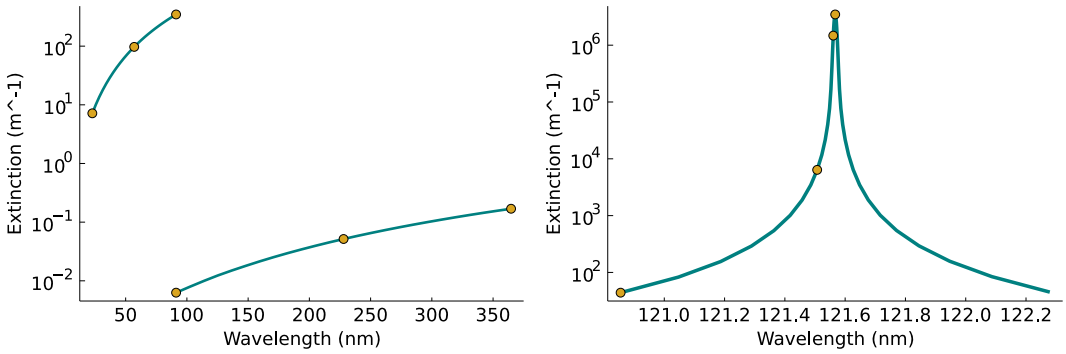

(a) Bound-free transitions.

(b) Bound-bound transition.

Figure 3.2: The extinction coefficient plotted in the lower atmosphere for each of the 150 wavelengths sampled. The points show the 10 *test-wavelengths* we use to configure the simulations.

3.1.3 Initial populations

The final input we need is the initial populations for the model atom. These can either be provided in their entirety, or they can be calculated based on a the total species density. In the latter case, they can be computed based on LTE or an assumption of zero-radiation.

3.1.4 Initial transition rates

The transition rates of the model atom become important in Section 3.4, where they go into the statistical equilibrium equation. However, they also play a part in the destruction probabilities that we need to set up the Monte Carlo simulation. For this reason, we need an initial estimation of the rates.

The collisional rates only depend on atmospheric properties, and stay the same throughout the iterations. We calculate them using the Transparency code written by Pereira (2021), based on the recipe from Johnson (1972).

The radiative rates, on the other hand, depend on the mean radiation field and need to be updated after every Monte Carlo run. We use the combined radiative rates given by Mihalas, D (1978, p. 131). They are also given in Appendix A.

3.2 Setting up the simulation

With the atmosphere and atom input in order, we are ready to enter the iteration scheme. The first step in every iteration is preparing the quantities needed in the Monte Carlo simulation.

3.2.1 Opacity calculations

We need to calculate the opacity for all the sampled wavelengths in every voxel. Here too, we use the Transparency code.

For the hydrogen bound-bound, bound-free and free-free processes, it uses the extinction expressions given in (2.19), (2.24) and (2.25). For background processes, we include Thomson scattering and Rayleigh scattering from hydrogen. Here, the code uses the expressions in (2.26) and (2.27). Finally, we include bound-free and free-free extinction from the H^- ion and H_2^+ molecule. For H^-, the bound-free expression is taken from Geltman (1962), while the free-free is from Stilley & Callaway (1970). Finally, the expressions for H_2^+ are from Bates (1952).

For the atom bound-bound and bound-free destruction probabilities, we use the ratio of the transition rates according to (2.20). We let the Rayleigh and Thomson scattering add to the total scattering extinction, while the H^- and H_2^+ extinction, as well as the hydrogen free-free, is taken as completely thermal for simplicity.

Line broadening

We use the line broadening calculations from Transparency. It includes the natural broadening given by (2.21), van der Waals broadning (Mihalas, D 1978, pp. 282, 286-287), linear stark broadening (Sutton 1978) and quadratic stark broadening (Traving, G 1960, p. 93) .

Additionally, we need to account for the Doppler broadening from the different voxel velocities. In practice, this means shifting line-profile according to the line-of-sight velocity,

$$\Delta\lambda = \lambda_0 \frac{v_{los}}{c} = \lambda_0 \frac{\mathbf{v} \cdot \mathbf{u}}{c}, \tag{3.1}$$

where $\mathbf{v}$ is the velocity of the voxel, $\mathbf{u}$ is the unit vector of the direction of the packet and c is the speed of light. However, because we only know the direction of the packet mid-simulation, we have to make this correction scattering by scattering.

3.2.2 Creating energy packets

The next step is determining how many packets to create in each voxel. For every wavelength, we let the packets be distributed according to the thermal energy emission. For each voxel, the total energy emitted per time, per wavelength and per solid angle, is then

$$E_{\lambda,voxel} \equiv B_{\lambda,voxel}\, \alpha^a_{\lambda,voxel}\, V_{voxel}, \tag{3.2}$$

where V_{voxel} is the volume of the voxel, and $\alpha^a_{\lambda,voxel}$ is the extinction from absorption processes. Each voxel is then given a fraction of packets corresponding to its fraction of the total energy emission,

$$\frac{N_{\lambda,\,voxel}}{N_\lambda} = \frac{E_{\lambda,voxel}}{\sum_v E_{\lambda,v}}, \tag{3.3}$$

where N_λ is the total number of packets created for the wavelength λ, $N_{\lambda,\,voxel}$ is the number in a given voxel, and the sum in the denominator is over all voxels v. Each packet then carries an energy

$$E_{\lambda,\,packet} = \frac{\sum_v E_{\lambda,v}}{N_\lambda}, \tag{3.4}$$

with units J s^{-1} nm^{-1} sr^{-1}. When a packets enters a voxel, it adds to the mean radiation field by

$$\Delta J_\lambda = \frac{E_{\lambda,\,packet}}{A_{voxel}}, \tag{3.5}$$

where A_{voxel} is some area associated with each voxel. In a idealised case where each packet travels straight up, it would be natural to set this area equal to the bottom surface of each voxel, $\Delta x \Delta y$. However, because the packets will be inclined to some degree, A_{voxel} should have some dependence on the height of the voxel Δz. This dependence will likely change over height. However, for simplicity, we assume that it stays the same, an use an expression on the form,

$$A_{voxel} = \Delta z \sqrt{\Delta x^2 + \Delta y^2}. \tag{3.6}$$

3.3 Monte Carlo simulation: solving for J

Once all the set-up is taken care of, the simulation itself is straightforward. The steps below are repeated for every wavelength sampled in Section 3.1.2.

1. For each voxel, we generate the number of packets defined by (3.20).
2. For each packet, we then draw a random starting position in the voxel,

   $$\boldsymbol{r_0} = [(x_{i+1} - x_i)\xi_1,\, (y_{j+1} - y_j)\xi_2,\, (z_{k+1} - z_k)\xi_3],$$

   where x_i, y_j z_k are the voxel edges, and the ξ's are random numbers between 0 and 1, drawn from a uniform distribution.
3. From the starting position, we draw a random direction of traversal and a random optical depth, like explained in Section 2.2.2,

   $$\phi = 2\pi\xi_4 \qquad \theta = \pi\xi_5 \qquad \tau = -\log(1 - \xi_6),$$

   where the direction of traversal $\boldsymbol{u} = [\cos\phi \sin\theta,\, \sin\phi \sin\theta,\, \cos\theta]$.

4. We propagate the packet along the chosen direction until it intersects the edge of the voxel. If at this point it has traversed an optical depth greater than the τ drawn, it is placed at the correct position

$$\boldsymbol{r_{new}} = \boldsymbol{r_0} + \frac{\tau}{\alpha_{voxel}}\mathbf{u}, \tag{3.7}$$

within the original voxel. We can then move on to step 5.

If it, on the other hand, has not reached the target depth, it is propagated across voxel by voxel in its path, until one of three things happen:

a) The depth target is surpassed. The packet is then propagated to the correct position within the final voxel, and we move on to step 5.

b) The packet escapes through the top. We then move on to the next packet in step 1.

c) The packet is lost in the bottom, and we move on to the next packet.

The sides of the atmosphere are handled with periodic boundaries. Crucially, for each voxel the packet traverses, it adds to the mean radiation field J_λ, according to (3.5).

5. If the packet is moved to a new valid position in the atmosphere, we draw another random number $\xi_7 \in [0, 1]$. If ξ is smaller than the local destruction probability ε_λ in the voxel, the packet is destroyed, and we move on the the next one. If ξ_7 is larger than ε, we let the packet scatter, and repeat the process from step 3. This can be repeated for a defined maximum number scatterings.

In the end, the hope is to have a good statistical sampling of the mean radiation field J_λ at every relevant wavelength.

3.4 Statistical equilibrium: solving for the populations

With the new estimate for J_λ, we can update the radiative transitions rates, and solve for new populations with the statistical equilibrium equations. For a N-level atom plus continuum, we have N+1 equations on the form,

$$n_i \sum_{j\neq i}^{N+1} P_{ij} - \sum_{j\neq i}^{N+1} n_j P_{ji} = 0, \tag{3.8}$$

where P_{ij} is the total transition probability from level i to j, while n_i is the level population for level i. While we have $N+1$ equations and $N+1$ unknown populations, this system of equations is dependant. To get a solution, we need to use the total atom density $n_{tot} = \sum_i n_i$. We can then substitute the ground level population,

$$n_1 = n_{tot} - \sum_{j\neq 1}^{N+1} n_j, \tag{3.9}$$

into the remaining N equations from (3.8), and get

$$n_i \sum_{j\neq i}^{N+1} P_{ij} - (n_{tot} - \sum_{j\neq 1}^{k} n_j)P_{1i} - \sum_{j\neq i,1}^{N+1} n_j P_{ji} = 0 \tag{3.10}$$

$$n_i \left(P_{1i} + \sum_{j\neq i}^{N+1} P_{ij} \right) + \sum_{j\neq i,1}^{N+1} n_j \left(P_{1i} - P_{ji} \right) = n_{tot} P_{1i} \tag{3.11}$$

Thus, the solution is on the form $\vec{n} = A^{-1}\vec{b}$, where A is a $N \times N$ matrix with diagonal elements $P_{1i} + \sum_{j\neq i}^{N+1} P_{ij}$, and off-diagonal elements $P_{1i} - P_{ji}$, while the N-vector $\vec{b} = n_{tot}\left[P_{12}, ..., P_{1(N+1)}\right]$.

3.5 Convergence check

Once we have the new populations, we compare them to the old ones and demand a relative error smaller than 0.01% in each voxel.

$$\epsilon_{i+1} = \frac{|n_{i+1} - n_i|}{n_i} < 1e-4, \tag{3.12}$$

where ϵ_{i+1} is the relative error in the (i+1)-th iteration. If this criterion is not met in every voxel, we go back to the beginning of Section 3.2 and use the updated populations to set up a new simulation.

3.6 Re-distributing packets

The sections above describe a basic implementation based on the standard methods of MCRT. However, for this the Monte Carlo approach to be feasible, we need to carefully consider the conditions in the solar atmosphere, and make sure we sample it in an efficient way.

Section 3.6.1 presents a method for re-distributing packets away from the high optical depth regions in the lower atmosphere, while Section 3.6.2 gives a way of re-distributing packets between wavelengths.

3.6.1 Excluding high optical depth regions

The three central parameters in the MCRT simulation are the extinction α_λ, the photon destruction probability ε_λ, and the energy emission E_λ. They alone explain why it is necessary to address the high optical depth problem, and will also be the key to get around it. In Figure 3.3, their average column values are plotted for each of the test-wavelengths.

Imagine first that we naively distribute the packets according to the energy emission over the entire atmosphere. Figure 3.3(e) and 3.3(f) show that essentially all packets will be created below $z = 0$. At the same time, we see from the extinction plots in 3.3(a) and 4.1(b), that the mean free path ($l_\lambda = \alpha_\lambda^{-1}$) is small here, and the packets will on average go through many interactions before they leave the region. However, the high destruction probabilities we see in Figure 4.1(a) and 3.3(d) means that most of them will be destroyed before they get that far. In other words, we have to create a huge number of packets before a significant number of packets reach the higher regions. While there is a lot of variation between the different wavelengths, the same reasoning holds for them all, and it is clear that a naive distribution cannot work.

We know that the lower regions are in LTE, and the mean radiation field here follows the Planck function. Our approach is therefore to simply exclude this region from the simulation. To do this, we define a boundary for which the region below is ignored in the Monte Carlo simulation. In practice, this means that no packets are started here, and packets going below the boundary are considered lost. Instead, we set the mean radiation field equal to the Planck function in this region.

The question then becomes how to define a boundary like this. On the one hand, we want to exclude enough for many packets to reach the higher regions. On the other hand, if we cut off too much of the lower regions $J_\lambda \approx B_\lambda$ no longer holds, and we risk loosing information from too many packets that would otherwise have escaped. Defining a boundary will therefore be a trade-off between computational efficiency and accuracy. We cannot compromise too much on the accuracy, so the success of this approach will essentially be determined by the computational cost of the accuracy.

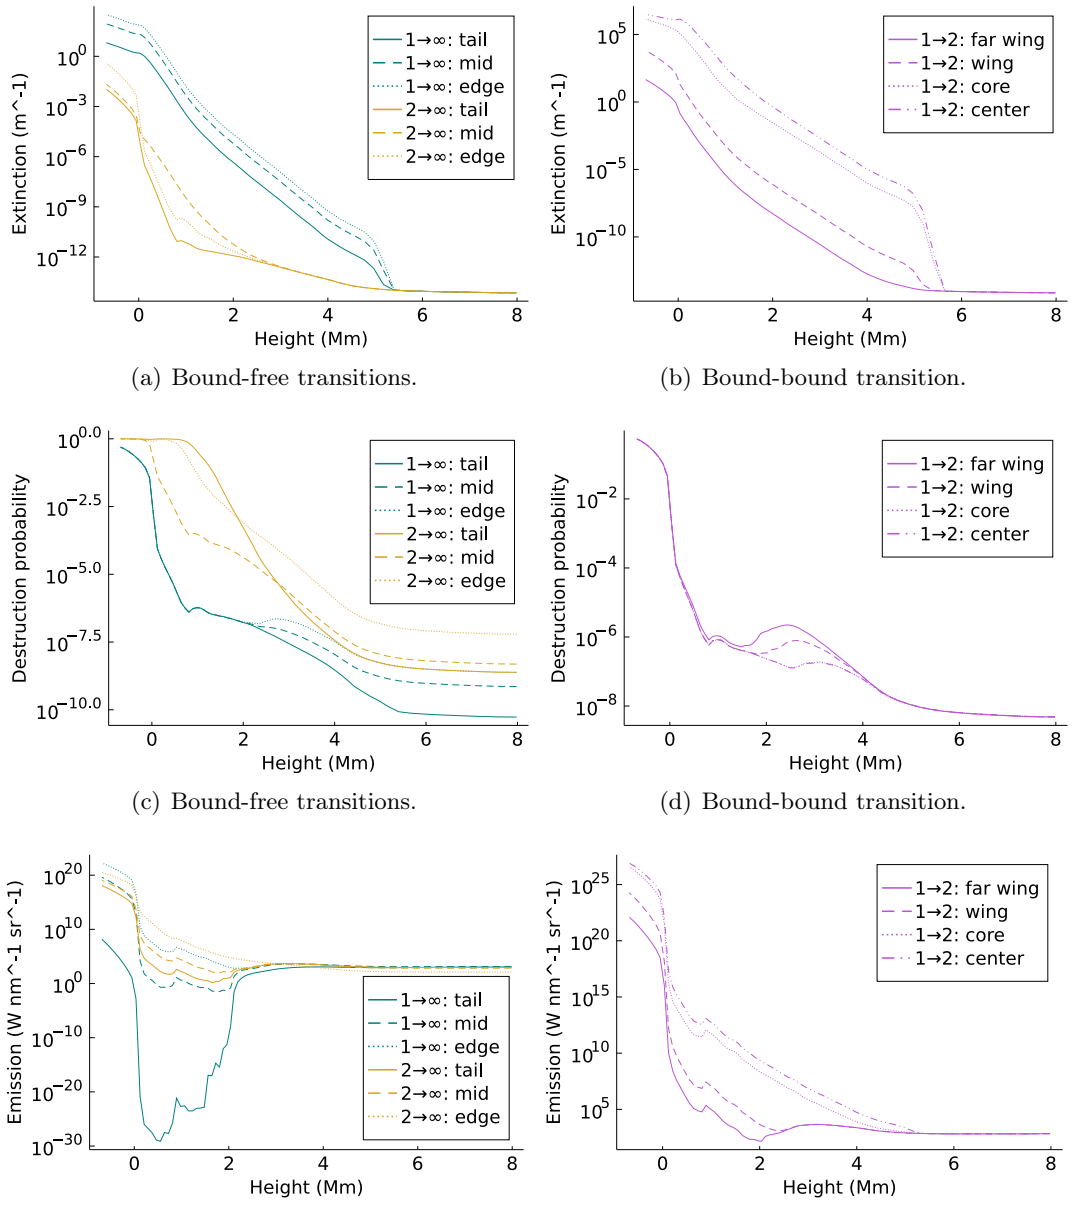

(a) Bound-free transitions.

(b) Bound-bound transition.

(c) Bound-free transitions.

(d) Bound-bound transition.

(e) Bound-free transitions.

(f) Bound-bound transition.

Figure 3.3: The column-averaged extinction coefficient, destruction probability and energy emission plotted over height for the test wavelengths.

Since this is a trial and error approach, we test a few different boundaries. First, we define

$$B_1(\tau_\lambda, \varepsilon_\lambda) \equiv \tau_\lambda \varepsilon^n, \tag{3.13}$$

and let the boundary z_b be defined by,

$$z_b \equiv \max(z \mid B_1(z) > b_1), \tag{3.14}$$

where n and b_1 are parameters we can tweak. The most obvious choice would perhaps be to let $n = 0$, and set the boundary at some optical depth b_1. However, it is also possible that a boundary with some sensitivity to the destruction probability will work better.

In addition to (3.13), we test a boundary that is more similar in spirit to MCRT. Here, the idea is to use the likelihood that a packet leaves a voxel. If we assume an average packet has to traverse half the height of the voxel it is created in, we can use the effective path $l^*_\lambda \equiv \sqrt{N}/\alpha_\lambda = \frac{1}{2}\Delta z$ to estimate how many interaction the packet will go through before it leaves the voxel. The probability that it can go through those interaction without being destroyed is then

$$B_2(\tau_\lambda, \varepsilon_\lambda) = (1 - \varepsilon_\lambda)^N, \tag{3.15}$$

where the number of scatterings $N \approx \left(\frac{1}{2}\Delta z \alpha_\lambda\right)^2$. In this case, we let the boundary be defined by

$$z_b = \min(z \mid B_2(z) > b_2). \tag{3.16}$$

We can think of the value $b_2 \in [0, 1]$, as an estimate of the fraction of photons that will escape from the boundary. While this second boundary definition is easier to interpret, it is harder to tweak. In that regard, they complement each other.

For these two boundary definitions, we look for values b_1 and b_2 that give a stable radiation field for the different wavelengths. By stable, we first mean that the field transitions continuously over the boundary, and that every voxel is visited by at least 50 packets for every wavelength. Additionally, we check that small changes in the value of b does not significantly change the radiation field. Specifically, we demand

$$\Delta J_{\lambda,i} \equiv \frac{|J_\lambda(b_{i+1}) - J_\lambda(b_i)|}{J_\lambda(b_i)} < 0.05 \tag{3.17}$$

in every voxel. There is no guarantee that stable boundary ranges like this exist. Or if they do, it is not certain that it is an indication of a realistic radiation field. However, if the radiation field is very sensitive to small changes in the boundary, it may be difficult to configure the Monte Carlo simulation in a way that works over all wavelengths.

A final note is that the change in the radiation field defined in (3.17), will depend on how many packets we use in the simulation. This means that what we call the *stable boundary range* will be limited by how long we run the simulation for. For practical reasons, we choose a number of packets that correspond to about 15 minutes for each wavelength. This is an arbitrary choice that we need to keep in mind in the analysis that follows.

3.6.2 Re-distributing packets between wavelengths

One final choice we need to make, is how many packets to create for each wavelength. The most obvious approach is to let N_λ be the same for all λ. However, this is not necessarily the most efficient solution. For instance, for wavelengths that scatter a lot, we may need fewer packets to get a good representation of the field. The same way we define a boundary to get more out of our computational resources, we want to re-distribute the packets in an efficient way.

To do this, we use the same idea as for the boundary. We look for the minimum number of packets required for the radiation field to stabilise,

$$\Delta J_{\lambda,i+1} \equiv \frac{|J_\lambda(N_i) - J_\lambda(N_{i+1})|}{J_\lambda(N_i)} < 0.05. \tag{3.18}$$

To avoid having to do this for all wavelengths, we find the minimum $N_{\lambda,min}$ for the test-wavelengths λ, and then do a least square fit to the model,

$$N_\lambda = N\,\mathrm{scale}_\lambda = N\,\frac{\sum_v N_{\lambda,v}\,\varepsilon^s_{\lambda v}}{N}, \tag{3.19}$$

where N is a packet number given to a wavelength that does not scatter, and scale_λ is the fraction of N that should be given to λ. When this model is determined, the approach is to first use (3.3) with $P_\lambda = P$ to distribute the packets for each wavelength, and then re-scale each distribution using (3.19), so that $N_\lambda = N\mathrm{scale}_\lambda$, and

$$N_{\lambda,\,voxel} = \frac{E_{\lambda,voxel}}{\sum_b E_{\lambda,b}} N\mathrm{scale}_\lambda. \tag{3.20}$$

This way, the individual packet distribution is conserved for each wavelength, but more packets are given to wavelengths with higher destruction probability over the packet creation domain. As long as enough packets are given to each wavelength, this distribution does not have to be perfect.

3.7 Code implementation

The full Monte Carlo radiative transfer routine was written in the Julia Language (Bezanson et al. 2017). In all simulations, we take advantage of its built-in multi-threading. The full code implementation is available in the repository `github.com/f0rmIdabel/SolarMCRT`.

Chapter 4

Results

The results are divided into two main parts. In Section 4.1, we present basic tests of the re-distribution methods presented in Section 3.6, and choose a simulation set-up based on this. In Section 4.2, we give the results from a full Monte Carlo radiative transfer calculation, and compare the results to the corresponding output from the MULTI3D code (Leenaarts & Carlsson 2009).

4.1 Effect of simulation parameters on radiation field

4.1.1 Boundary

We look at the two boundary functions defined in (3.13) and (3.16). For $B_1 = \tau\varepsilon^n$, we try $n = 0$, 0.5, 1.0, 1.5, 2.0, and find that $n = 1$ gives the most stable results over different wavelengths. In Figure 4.1, this is plotted together with B_2 for the test-wavelengths.

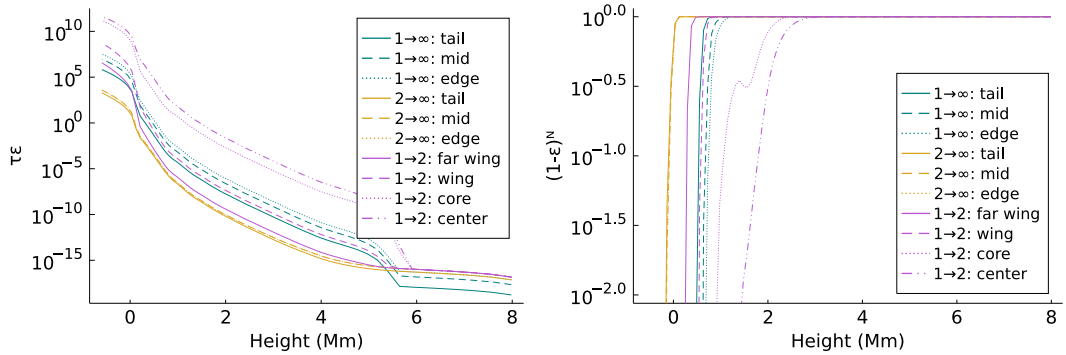

(a) B_1: Optical depht from destruction processes.

(b) B_2: Escape probability.

Figure 4.1: The column averaged boundary functions plotted over height.

For each boundary function, we looked for ranges $b \in [b_{min}, b_{max}]$, for which the mean field J_λ remains stable after the criteria given in (3.17). The results are given in Table 4.1.

B	**λ** (nm)	b_{max}	b_{min}	τ_λ
$\tau_\lambda \varepsilon_\lambda$	22.79	1e2	1e1	4e5
	56.99	1e3	1e1	4e5
	91.18	1e4	1e1	1e7
	91.19	–	–	–
	227.94	–	–	–
	364.70	–	–	–
	120.86	1e2	1e1	3e4
	121.51	5e1	1e1	2e6
	121.56	–	–	–
	121.57	–	–	–

B	**λ** (nm)	b_{max}	b_{min}	τ_λ
$(1-\varepsilon_\lambda)^{N_\lambda}$	22.79	1e-2	1e-10	2e4
	56.99	1e-3	1e-10	9e4
	91.18	1e-3	5e-9	1e5
	91.19	–	–	–
	227.94	–	–	–
	364.70	–	–	–
	120.86	5e-2	1e-7	7e4
	121.51	5e-2	1e-7	5e4
	121.56	1e-5	1e-20	3e7
	121.57	–	–	–

Table 4.1: The stable boundary range at the test-wavelengths for the the two different boundary functions. The column-averaged optical depth in the middle of b_{max} and b_{min} is given for reference.

The results are similar for the two boundary functions. Neither find stable regions for all the wavelengths. We see that the wavelengths from the second bound-free transition and the line core are particularly challenging. This behaviour is highlighted in Figure 4.2. Here, the the column-averaged mean radiation field J_λ is plotted for different boundary values b, for the second boundary function B_2. While the radiation field remains stable for the wavelength from the first bound-free transition, we see that it changes by several orders of magnitude at the second bound-free transition. For the line core wavelength, the field is unstable in the region close to the boundary, but stabilises in the uppermost layers.

We note that the optical depth at the stable regions we found is high for all the wavelengths. This means that while we are excluding the highest depth regions, the atmosphere is still thick in the regions just above the defined boundaries. Although, it is the thickest for the B_1 boundary.

In the following analysis, we use the boundary definitions

$$B_1(\tau_\lambda, \varepsilon_\lambda) = 10 \qquad B_2(\tau_\lambda, \varepsilon_\lambda) = 10^{-3}. \tag{4.1}$$

While it at this stage is difficult to evaluate their performance, they both give continuous radiation fields for all test wavelengths. Nevertheless, in the following results, it will be important to keep their shortcomings in mind.

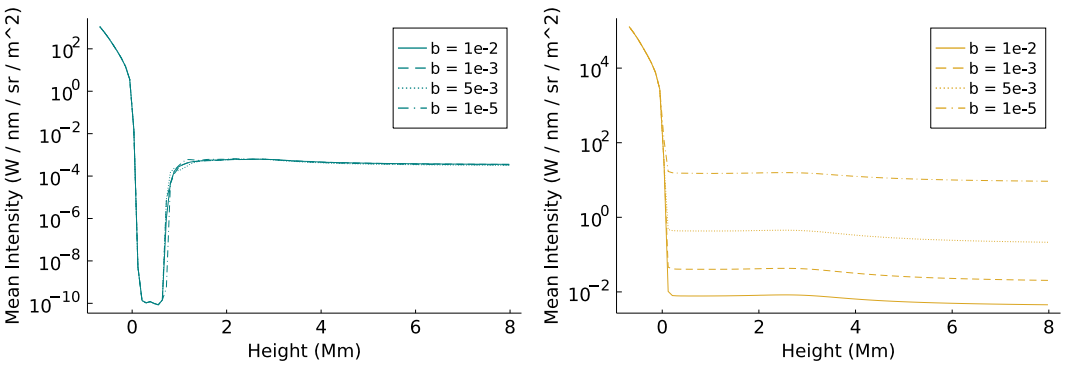

(a) $1 \rightarrow \infty$: $\lambda = 56.99$ nm

(b) $2 \rightarrow \infty$: $\lambda = 91.19$ nm

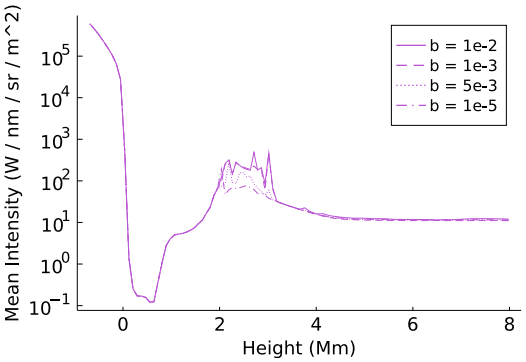

(c) $2 \rightarrow 1$: $\lambda = 121.57$ nm

Figure 4.2: The column averaged radiation field J_λ plotted for different boundary definitions $B_2 > b$ at one wavelength from each transition.

4.1.2 Packets

The least-square fit of (3.19), gives the packets scale

$$scale_\lambda = \frac{\sum_v N_{\lambda,v}\, \varepsilon_\lambda^{0.23}}{N}. \tag{4.2}$$

Table 4.2 gives the minimum packet values N_min it is calculated from. We see that the lowest opacity wavelengths require more than 200 times more packets than the highest opacity wavelengths before the field stabilises. Although, even with this re-distribution, it is the high opacity wavelengths that take up the longest time in the simulation.

4.2 Monte Carlo radiative transfer

With the newly defined boundaries in (4.1) and the packet scaling in (4.2), we are ready to run the full Monte Carlo radiative transfer simulation. In addition to the two boundaries, we try two different initial populations: LTE and zero-radiation (ZR).

Boundary	**λ** (nm)	N_{min}	N_{model}	**time (s)**
$\tau_{1.0}$	22.79	2e6	9e5	110
	56.99	1e6	1e6	90
	91.18	1e6	1e6	122
	91.18	9e7	1e8	42
	227.94	1e8	1e8	31
	364.70	7e7	1e8	29
	120.86	1e8	6e7	255
	121.51	3e8	4e7	321
	121.56	7e5	2e6	91
	121.57	4e5	2e6	340

Table 4.2: The distribution of packets between the test-wavelengths.

Finally, we run the simulation at half resolution in every dimension. Due to time limitation, we only use the minimal packet distribution, P_{min}, and set the maximum number of iterations to 25. We use 50 wavelengths for each of the bound-free transitions, and 51 wavelengths for the line transition.

Boundary	**nλ**	**Resolution**	**Initial**	**Iterations**	**Time (h)**
$(1-\varepsilon)^N$	151	Full	LTE	11	17
			ZR	9	16
		Half	ZR	–	–
$\tau\varepsilon$	151	Full	LTE	–	–
			ZR	–	–
		Half	ZR	–	–

Table 4.3: The number of iterations before convergence and the total time for every configuration.

The boundary B_1 did not give convergence within 25 iterations. Instead it stabilised around a mean relative population error 0.1. Meanwhile, B_2 gave convergence at full resolution for both initial populations. The zero-radiation populations performed slightly better, with convergence in 9 iterations, and after 16 hours of run time.

In Figure 4.3, we compare the converged mean radiation field J_λ to the output from MULTI3D, for one wavelength at each transition. The results mirror what we saw in the boundary tests. The wavelength from the first bound-free transition follows the integral method the closest. The low opacity wavelength from the second bound-free transition keeps a constant high level of error throughout the atmosphere. Finally, the

error of the line core wavelength decreases towards the higher regions.

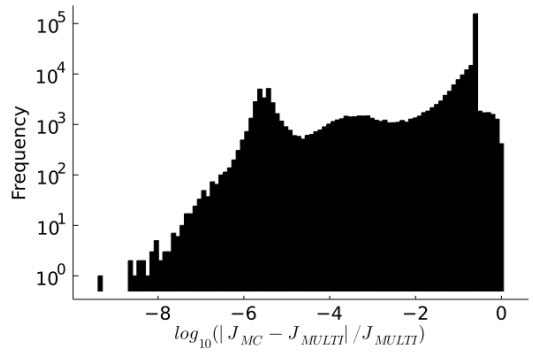

(a) 1→ ∞: $\lambda = 56.99$ nm

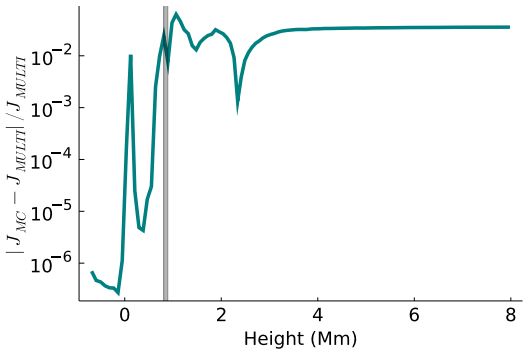

(b) 1→ ∞: $\lambda = 56.99$ nm

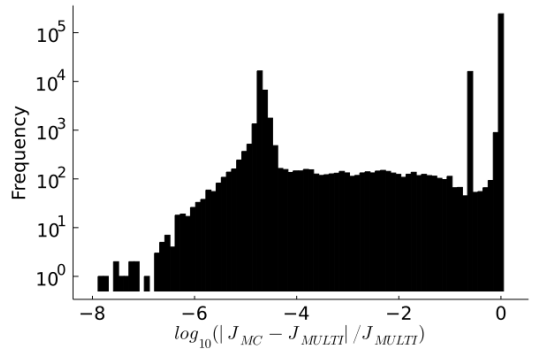

(c) 2→ ∞: $\lambda = 91.19$ nm

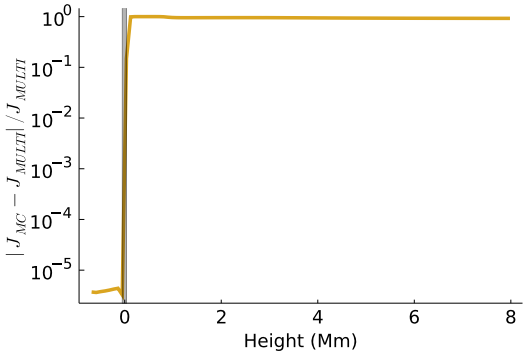

(d) 2→ ∞: $\lambda = 91.19$ nm

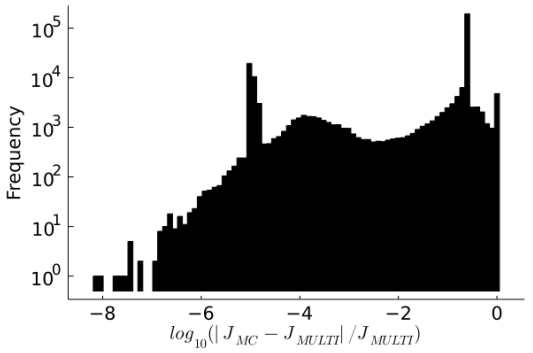

(e) 2→ 1: $\lambda = 121.57$ nm

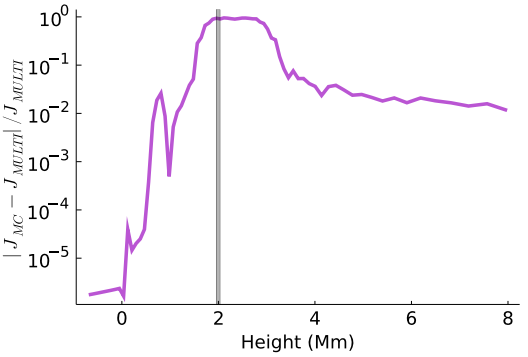

(f) 2→ 1: $\lambda = 121.57$ nm

Figure 4.3: The relative error between radiation fields from the MCRT transfer calculation and MULTI3D. First on a voxel by voxel basis, then the average error over height. The average boundary heights are shown with a grey line.

Chapter 5

Discussion and conclusions

This project has been motivated by the same things that brought other fields in astronomy to Monte Carlo methods; the need for efficient methods for 3D non-LTE radiative transfer problems. Introducing these methods to cold stellar atmospheres is challenging, because not all parts of them are in non-LTE. Monte Carlo methods tend to perform poorly at high optical depth. So in addressing one problem, Monte Carlo radiative transfer introduces the opposite one.

To get around this, we have introduced a boundary with the goal of excluding the high depth regions. Here, the challenge has first been to find configurations that are stable for a single wavelength. Then to extend this in a way that works for all wavelengths. This has not been trivial.

First, we saw that both boundary functions were unstable. $B_1 = \tau_\lambda \varepsilon_\lambda$ failed to converge within 25 iterations. While the second definition $B_2 = (1 - \varepsilon_\lambda)^N$ performed slightly better, the converged radiation field show that the instabilities of the boundary follows the system throughout the iterations.

The failing of the boundary method can be partly explained by looking back at the escape probability plot in Figure 4.1(b). Looking at the two line core wavelengths, we see that they do not have the same clear drop-off as the rest. This may be an indication that the concept of a hard boundary that cuts off information from the lower layers will be unstable. For these wavelengths, the transition from LTE to non-LTE happens more gradually. However, because they scatter a lot, they are able to recover the field towards the top of the atmosphere. This explains the high error close to the boundary.

On the other side of the problem is the low opacity wavelengths from the second bound-free transition. For these, the idea of a boundary makes more sense. The escape rate increases rapidly in a narrow height range. The problem here, is the low rate of interactions. This makes the wavelengths very sensitive to the initial boundary. It becomes a little like a shooting method. It is for the intermediate depth wavelengths the method shines. They both fulfill the condition of a clear-cut boundary and they interact more with the atmospheres. An inexact boundary will therefore effectively be corrected for.

In future work, the boundary method should either be modified to work better

with the high and low opacity wavelengths, or other methods need to be developed to overcome the high optical depths regions. Another improvement will be to get a more refined packet distribution. This could for example be done by performing a more sophisticated least-square-fit. There is likely a lot to gain here in terms of performance. In Monte Carlo methods, performance is not just about time, but also about error reduction. In future work, it will be important to get a better idea of this trade-off between noise and time.

This thesis has only scratched the surface of Monte Carlo radiative transfer. While we have seen promising results, we are still a far way from a competitive radiative transfer code. In the mean time, it can already be a great supplement to the traditional methods.

Appendices

Appendix A

Combined radiative rates

Expression for combined radiative rates are given by Mihalas, D (1978, pg. 131). Here, they are re-written in terms of wavelength instead of frequency. For upward transitions,

$$R_{ij} = \frac{4\pi}{hc} \int_0^\infty \lambda \sigma_{ij} I_\lambda d\lambda, \tag{A.1}$$

while for downward transitions,

$$R_{ji} = \frac{4\pi}{hc} \int_0^\infty \lambda \sigma_{ij} G_{ij} \left(\frac{2hc^2}{\lambda^5} + J_\lambda \right) d\lambda. \tag{A.2}$$

Where,

$$G_{ij} = \left[\frac{n_i}{n_j} \right]_{\mathrm{LTE}} \exp^{-hc/\lambda kT}. \tag{A.3}$$

For bound-bound transitions,

$$\sigma_{ij} = \frac{hc}{4\pi\lambda} B_{ij} \varphi_\lambda \tag{A.4}$$

While for bound-free transitions,

$$\sigma_{ij} = \frac{4e^2}{3\pi sqrt3 \varepsilon_0 m_e c_0^2 R_\infty} \left(\frac{\lambda}{\lambda_{edge}} \right)^3 Z^4 g_{bf} \tag{A.5}$$

Bibliography

Auer, L. H., & Mihalas, D. 1969, ApJ, 158, 641, doi: 10.1086/150226

Bates, D. R. 1952, MNRAS, 112, 40, doi: 10.1093/mnras/112.1.40

Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. 2017, SIAM review, 59, 65. https://doi.org/10.1137/141000671

Bjørgen, J. P., & Leenaarts, J. 2017, A&A, 599, A118, doi: 10.1051/0004-6361/201630237

Camps, P., & Baes, M. 2018, ApJ, 861, 80, doi: 10.3847/1538-4357/aac824

Geltman, S. 1962, ApJ, 136, 935, doi: 10.1086/147447

Gudiksen, B. V., Carlsson, M., Hansteen, V. H., et al. 2011, A&A, 531, A154, doi: 10.1051/0004-6361/201116520

Hubeny, I., & Lanz, T. 1996, in PASP, Vol. 96, Hydrogen Deficient Stars, ed. C. S. Jeffery & U. Heber, 249

Johnson, L. C. 1972, ApJ, 174, 227, doi: 10.1086/151486

Kalos, M. H., & Whitlock, P. A. 1986, Monte Carlo methods (New York, NY: Wiley)

Kramida, A., Yu. Ralchenko, Reader, J., & NIST ASD Team. 2020, NIST Atomic Spectra Database (ver. 5.8), [Online]. Available: https://physics.nist.gov/asd [2021, April 18]. National Institute of Standards and Technology, Gaithersburg, MD.

Leenaarts, J., & Carlsson, M. 2009, in Astronomical Society of the Pacific Conference Series, Vol. 415, The Second Hinode Science Meeting: Beyond Discovery-Toward Understanding, ed. B. Lites, M. Cheung, T. Magara, J. Mariska, & K. Reeves, 87

Mihalas, D. 1978, Stellar atmospheres (W.H. Freeman)

Min, M., Dullemond, C. P., Dominik, C., de Koter, A., & Hovenier, J. W. 2009, A&A, 497, 155, doi: 10.1051/0004-6361/200811470

Noebauer, U. M., & Sim, S. A. 2019, L. Rev. Comp. Astrophys., 5, doi: 10.1007/s41115-019-0004-9

Pereira, T. M. D. 2021, tiagopereira/Transparency.jl, v0.1.5, Zenodo, doi: 10.5281/zenodo.4746629

Robitaille, T. 2015, A&A, 576, doi: 10.1051/0004-6361/201015025e

Rutten, R. J. 2003, Radiative Transfer in Stellar Atmospheres

Stilley, J. L., & Callaway, J. 1970, ApJ, 160, 245, doi: 10.1086/150423

Sukhorukov, A. V., & Leenaarts, J. 2016, A&A, 597, A46, doi: 10.1051/0004-6361/201629086

Sutton, K. 1978, J. Quant. Spec. Radiat. Transf., 20, 333, doi: https://doi.org/10.1016/0022-4073(78)90102-4

Traving, G. 1960, Über die Theorie der Druckverbreitung von Spektrallinien (Karlsruhe, G. Braun)

Tritschler, A., Rimmele, T., Berukoff, S., et al. 2016, Astron. Nachr., 337, 1064, doi: 10.1002/asna.201612434

Uitenbroek, H. 2001, ApJ, 557, 389, doi: 10.1086/321659

Whitney, B. A. 2011, Bull. Astron. Soc. India., 39, 101. https://arxiv.org/abs/1104.4990