

MAGRITTE, a modern software library for 3D radiative transfer: I. Non-LTE atomic and molecular line modelling

Frederik De Ceuster^{1,2★}, Ward Homan,² Jeremy Yates,¹ Leen Decin,^{2,3} Peter Boyle⁴
and James Hetherington^{5,6}

¹Department of Physics and Astronomy, University College London, Gower Place, London WC1E 6BT, UK

²Department of Physics and Astronomy, Institute of Astronomy, KU Leuven, Celestijnenlaan 200D, B-3001 Leuven, Belgium

³School of Chemistry, University of Leeds, Leeds LS2 9JT, UK

⁴School of Physics and Astronomy, The University of Edinburgh, Edinburgh EH9 3FD, UK

⁵Department of Computer Science, University College London, Bloomsbury, London WC1E 6EA, UK

⁶The Alan Turing Institute, 96 Euston Road, Kings Cross, London NW1 2DB, UK

Accepted 2019 December 14. Received 2019 December 13; in original form 2019 October 31

ABSTRACT

Radiative transfer is a key component in almost all astrophysical and cosmological simulations. We present MAGRITTE: a modern open-source software library for 3D radiative transfer. It uses a deterministic ray-tracer and formal solver, i.e. it computes the radiation field by tracing rays through the model and solving the radiative transfer equation in its second-order form along a fixed set of rays originating from each point. MAGRITTE can handle structured and unstructured input meshes, as well as smoothed-particle hydrodynamics (SPH) particle data. In this first paper, we describe the numerical implementation, semi-analytic tests and cross-code benchmarks for the non-LTE line radiative transfer module of MAGRITTE. This module uses the radiative transfer solver to self-consistently determine the populations of the quantized energy levels of atoms and molecules using an accelerated Lambda iteration (ALI) scheme. We compare MAGRITTE with the established radiative transfer solvers RATRAN (1D) and LIME (3D) on the van Zadelhoff benchmark and present a first application to a simple Keplerian disc model. Comparing with LIME, we conclude that MAGRITTE produces more accurate and more precise results, especially at high optical depth, and that it is faster.

Key words: radiative transfer – methods: numerical – software: development.

1 INTRODUCTION

Radiative processes play an essential role in the dynamics, chemistry, and energy balance of various astrophysical objects, from planetary and stellar atmospheres to galaxies and the Universe as a whole. Radiation can provide a radiative pressure that can drive dynamics (see e.g. Höfner et al. 2003), it can affect chemistry through various photoionization and photodissociation reactions (see e.g. Huggins & Glassgold 1982), and it can efficiently heat or cool very localized regions (see e.g. Woitke et al. 1996). Furthermore, the radiative properties determine what can and cannot be seen in observations, i.e. which regions are visible in what part of the electromagnetic spectrum.

Ever since the first detection and identification of atoms and molecules in space (Douglas & Herzberg 1941; Weinreb et al. 1963), their line emission and absorption features have been an indispensable diagnostic tool to infer the physical and chemical conditions

throughout the Universe. In order to interpret the observational data, we require high precision atomic and molecular data (Schöier et al. 2005) as well as sound theoretical models. Historically, these models quickly evolved from highly idealized equilibrium systems to more self-consistent non-equilibrium models (Mihalas & Athay 1973). Moreover, with the advent of high (spatial) resolution imaging, for instance using the Atacama Large Millimetre Array (ALMA), full 3D models are imperative to properly model the intricate structures observed in the data (see e.g. Maercker et al. 2012; Decin et al. 2015; Smith, Ginsburg & Bally 2018; Alves et al. 2019).

Given the tight coupling between radiation field and medium, it is crucial both in astrophysical and cosmological modelling to properly account for all radiative processes and their interdependence. This, however, can be highly complicated due to: (i) an intricate 3D geometrical structure shielding or exposing specific regions to radiation, (ii) the scattering of radiation by dust or free electrons yielding additional non-trivial coupling between the geometry and the radiation field, and (iii) the mixing in frequency space due to Doppler shifts caused by velocity gradients in the

* E-mail: frederik.deceuster@kuleuven.be

medium. Furthermore, the coupling between the radiative processes and the often very specialized and diverse dynamical and chemical models requires a modular radiative transfer solver that can easily be integrated with the various existing hydrodynamics and chemistry models. Finally, the ever growing size and complexity of these models requires fast and scalable algorithms that can efficiently leverage the wealth of modern computational resources.

There are two main computational strategies to solve radiative transfer problems. On the one hand there are probabilistic (Monte Carlo) solvers such as e.g. RADMC-3D (Dullemond et al. 2012), SKIRT (Verstocken et al. 2017), CMACIONIZE (Vandenbroucke & Wood 2018), and some components of TORUS (Harries et al. 2019). On the other hand there are deterministic or formal solvers such as e.g. SPHRAY (Altay, Croft & Pelupessy 2008), 3D-PDR (Bisbas et al. 2012), and LAMPFRAY (Frostholm, Haugbølle & Grassi 2018). Furthermore, there are also codes that combine ideas from both techniques such as RATRAN (Hogerheijde & van der Tak 2000) and its 3D successor LIME (Brinch & Hogerheijde 2010). The latter has been widely used to model atomic and molecular lines in 3D models of various astrophysical objects (see Andrews et al. 2012; Maercker et al. 2012; Bergin et al. 2013; Walsh et al. 2016; Homan et al. 2018; Booth, Walsh & Ilee 2019; Evans et al. 2019; Montargès et al. 2019).

Currently, most radiative transfer solvers use probabilistic methods. These methods mimic the physical photon transport by propagating a number of photon packets through the medium (see e.g. Noebauer & Sim 2019, for an extensive review). The main issue with this approach is that the trajectories of these photon packets are randomly determined by the properties of the medium. This implies that they can get trapped in opaque regions, impeding them from contributing much to the overall radiation field. Hence, a large number of packets need to be propagated which can significantly increase the computation time. Although many techniques have been devised to avoid the trapping of photon packets (see e.g. Yusef-Zadeh, Morris & White 1984), it remains challenging for probabilistic radiative transfer solvers to efficiently obtain accurate results, especially at medium to high optical depths (Camps & Baes 2018).

Deterministic or formal solvers compute the radiation field by solving the radiative transfer equation along rays through the medium. Since the optical properties of the medium often depend on the radiation field this has to be done in an iterative way. Although there are no photon packets in this approach, a problem physically very similar to photon trapping can manifest itself in the form of slow convergence of the iteration process. In this context, the problem was first identified in the 1970s by various authors (see e.g. Scharmer & Carlsson 1985, and references therein) and is more commonly known as the Lambda-iteration problem. This problem arose when attempts were made to model the radiative hydrodynamics of hot stars without assuming LTE (i.e. non-LTE). The extremely slow or false convergence produced by this effect resulted in erroneous fits to the observed data. Subsequent work, for instance, by Olson, Auer & Buchler (1986) and Rybicki & Hummer (1991), elegantly addressed these issues using a technique called accelerated Lambda iteration (ALI). For a complete overview of these methods see, for example, Hubeny & Mihalas (2014).

Although deterministic solvers could better cope with the optical depth related issues, probabilistic solvers became more popular due to their relative ease of implementation, especially in two and three spatial dimensions. However, with the development of fast solution methods it is now possible to implement a deterministic solver with comparative ease. When combined with the ability to sample rays

finer, multidimensional ray-tracing codes can now become powerful probes of objects with complex geometries, velocity fields, and optical depth ranges. Moreover, their deterministic computational scheme leads to various opportunities for optimization and facilitates utilizing the various layers of parallelism in the calculation, further reducing the computational cost.

MAGRITTE is a modern open-source software library for 3D radiative transfer. It is written in C++11, but almost all classes and functions are wrapped using PYBIND11 (Wenzel, Rhinelanders & Moldovan 2017) such that they can also be used in PYTHON. Our motivation to develop MAGRITTE is twofold. On the one hand, the ever increasing amount of high-quality observational data puts increasingly higher demands on the modelling software, while, on the other hand, advances in computer technology provide us with the means to meet these demands. Common examples are the extended use of different layers of parallelism (e.g. vector instructions, multithreading and message passing) and the growing availability of hardware accelerators such as graphics processing units (GPUs) or field programmable gate arrays (FPGAs). Using these technologies in an existing code base, however, often requires a complete rewrite of the internal data structures. Therefore, we opted to build a new code base that is flexible enough to cope with the requirements for multiple astrophysical and cosmological applications and has a modular data structure that can readily be adapted to leverage the different forms of parallelism and hardware accelerators available in modern (super)computer architectures.

Since advances in modelling are increasingly made by improved software implementations rather than new mathematical techniques, it is imperative that both the software and its source code are publicly available for the community to review and adapt. Therefore, we commit ourselves to make future releases of MAGRITTE and its source code publicly available¹ at github.com/Magritte-code.

This is the first paper in a series in which we will analyse the physical, mathematical, and computational aspects of the various components of the software library. In this first paper, we present MAGRITTE's module for atomic and molecular line radiative transfer. The radiation field is computed self-consistently with the populations of the quantized energy levels. In contrast to many early treatments of line radiative transfer, this approach does not make the assumption of local thermodynamic equilibrium (LTE), which hence classifies it as a non-LTE solver. We present MAGRITTE's ray-tracing scheme which only uses cell locations and nearest neighbour information. Therefore, it can easily cope with smoothed particle hydrodynamics (SPH) particles as well as structured and unstructured model meshes. We introduce our solution method to solve the radiative transfer equation along a ray pair and present our implementation of the ALI scheme based on Rybicki & Hummer (1991). To validate our methods, we run a set of test models for which we can obtain semi-analytical results. This way, we can get an absolute measure of the errors resulting from our methods. We further demonstrate MAGRITTE's validity, by performing a cross-code comparison with RATRAN (Hogerheijde & van der Tak 2000) and LIME (Brinch & Hogerheijde 2010) using the van Zadelhoff et al. (2002) benchmark, and some additional variations on that. Finally, we present a first application of MAGRITTE modelling the CO emission of a simple Keplerian disc.

The structure of this paper is as follows. In Section 2, we introduce the radiative transfer problem and elaborate on the tight coupling between the radiation field and the medium. Section 3

¹Under GNU General Public License v2.0.

presents our solution methods to the problem and the numerical implementations. In Section 4, we describe a set of semi-analytic tests and cross-code benchmarks to validate our methods and Section 5 describes a first application of MAGRITTE to a simple Keplerian disc model. Finally, our results are discussed in Section 6 and we conclude with Section 7.

2 PHYSICAL PROBLEM

2.1 Radiative transfer

The objective of radiative transfer is to determine the radiation field in a region, given the properties of the medium in that region and some boundary conditions. The radiation field is described in terms of its specific monochromatic intensity $I_\nu(\mathbf{x}, \hat{\mathbf{n}})$, i.e. the energy transported in a certain direction in a certain frequency bin. This is a function of frequency (ν), position (\mathbf{x}), and direction ($\hat{\mathbf{n}}$). Any interaction between the radiation field and the medium can be described in terms of the change to the specific monochromatic intensity. The radiative transfer equation relates this change in specific monochromatic intensity $I_\nu(\mathbf{x}, \hat{\mathbf{n}})$ along a ray in direction $\hat{\mathbf{n}}$ to the local emissivity $\eta_\nu(\mathbf{x})$ and opacity $\chi_\nu(\mathbf{x})$ of the medium. Scattering introduces an extra contribution to both the emissivity and opacity (Chandrasekhar 1960; Steinacker, Baes & Gordon 2013). The time-independent radiative transfer equation including scattering reads

$$\hat{\mathbf{n}} \cdot \nabla I_\nu(\hat{\mathbf{n}}) = \eta_\nu - (\chi_\nu + \chi_\nu^{\text{sca}}(\hat{\mathbf{n}})) I_\nu(\hat{\mathbf{n}}) + \int d\Omega' \int_0^\infty d\nu' \Phi_{\nu\nu'}(\hat{\mathbf{n}}, \hat{\mathbf{n}}') I_\nu(\hat{\mathbf{n}}'), \quad (1)$$

where $\chi_\nu^{\text{sca}}(\hat{\mathbf{n}})$ is the extra opacity due to scattering and $\Phi_{\nu\nu'}(\hat{\mathbf{n}}, \hat{\mathbf{n}}')$ is the scattering redistribution function which gives the probability for radiation of frequency ν' incoming along direction $\hat{\mathbf{n}}'$ to be scattered in direction $\hat{\mathbf{n}}$ and to be shifted to frequency ν .

For local radiative processes, we assume both the emissivity and opacity to be isotropic, i.e. independent of the direction $\hat{\mathbf{n}}$. However, in contrast to the classical general formulation of the transfer equation by Cannon (1971, 1972), we allowed for a directional dependence in the scattering opacity. This more general approach allows us to also treat, for instance, scattering from dust grains that are aligned by a magnetic field (see e.g. Andersson, Lazarian & Vaillancourt 2015, and the references there). The anisotropy of the scattering opacity slightly complicates the solution methods. However, we have included it to keep our solution methods as general as possible.

The radiative transfer equation (1) is a first-order integro-differential equation. Generally, it can only be solved in an iterative way, since both the emissivity and opacity depend on the radiation field. We discuss our solution strategy for solving the transfer equation in Section 3.3. First, we break down the coupling between the medium and the radiation field.

2.2 Coupling radiation field and medium

In general we can distinguish four types of interactions between a radiation field and a medium based on the frequency range on which they act: line, ionization, continuum, and scattering interactions. In this first paper, we will limit ourselves to atomic and molecular line interactions.

2.2.1 Atomic and molecular lines

Electronic, rotational, and vibrational transitions between the quantized energy levels in atomic and molecular species can lead to a significant emission and absorption in narrow frequency ranges. These transitions are referred to as line transitions due to the characteristically narrow features they induce in spectra.

The resulting emissivity and opacity due to a line transition from a level i to a level j (with level energies $E_i > E_j$) are given in terms of the Einstein A_{ij} , B_{ji} , and B_{ij} coefficients and the populations $n_i(\mathbf{x})$ of the quantized energy levels

$$\eta_\nu^{ij}(\mathbf{x}) = \frac{h\nu}{4\pi} n_i(\mathbf{x}) A_{ij} \phi_\nu^{ij}(\mathbf{x}),$$

$$\chi_\nu^{ij}(\mathbf{x}) = \frac{h\nu}{4\pi} (n_j(\mathbf{x}) B_{ji} - n_i(\mathbf{x}) B_{ij}) \phi_\nu^{ij}(\mathbf{x}), \quad (2)$$

where A_{ij} and B_{ij} account, respectively, for spontaneous and stimulated emission and B_{ji} accounts for absorption. Note that stimulated emission is treated as negative absorption. Both line emissivity and opacity are proportional to the line profile function $\phi_\nu^{ij}(\mathbf{x})$. In this paper, we assume Gaussian line profile functions² resulting from the Doppler shifts caused by the thermal and turbulent motions of the atoms and molecules in the medium,

$$\phi_\nu^{ij}(\mathbf{x}) = \frac{1}{\delta\nu_{ij}(\mathbf{x})\sqrt{\pi}} \exp\left[-\left(\frac{\nu - \nu_{ij}}{\delta\nu_{ij}(\mathbf{x})}\right)^2\right], \quad (3)$$

where the characteristic line width

$$\delta\nu_{ij}(\mathbf{x}) = \frac{\nu_{ij}}{c} \sqrt{v_{\text{thermal}}(\mathbf{x})^2 + v_{\text{turbulent}}(\mathbf{x})^2}, \quad (4)$$

is determined by the mean thermal and turbulent velocities of the medium in the co-moving frame.

Many early line radiative transfer models assumed the populations of the quantized energy levels to be in local thermodynamic equilibrium (LTE), i.e. particle velocities, level populations, and radiation field are completely determined by the local gas temperature. In contrast, we will only assume kinetic equilibrium, i.e. we only assume a Maxwell–Boltzmann distribution for the particle velocities. This situation is often referred to as non-LTE. As a result, the mean local velocity of the gas particles in the co-moving frame can be characterized by

$$v_{\text{thermal}}(\mathbf{x}) = \sqrt{\frac{2k_B T(\mathbf{x})}{m_{\text{spec}}}}, \quad (5)$$

where m_{spec} is the mass of the species of gas under consideration. If we make no further assumptions on the level populations, they can only be determined by directly solving the kinetic rate equations, which, in the co-moving frame, are given by,

$$\frac{\partial n_i(\mathbf{x})}{\partial t} = \sum_{j=1}^N n_j(\mathbf{x}) P_{ji}(\mathbf{x}) - n_i(\mathbf{x}) \sum_{j=1}^N P_{ij}(\mathbf{x}). \quad (6)$$

The components of the matrix $P_{ij}(\mathbf{x})$ denote the transition rates from level i to level j . Hence, for each level i , $P_{ii}(\mathbf{x}) = 0$. The transition rates are composed of a radiative part $R_{ij}(\mathbf{x})$ and a collisional part $C_{ij}(\mathbf{x})$, such that

$$P_{ij}(\mathbf{x}) = R_{ij}(\mathbf{x}) + C_{ij}(\mathbf{x}). \quad (7)$$

²However, all our methods can also readily be applied to all other types of line profile functions.

The radiative part can be expanded further in terms of the Einstein coefficients and the average radiation intensity in the line

$$R_{ij}(\mathbf{x}) = \begin{cases} A_{ij} + B_{ij} J_{ij}(\mathbf{x}) & \text{for } i > j \\ B_{ji} J_{ij}(\mathbf{x}) & \text{for } i < j \end{cases}, \quad (8)$$

where $J_{ij}(\mathbf{x})$ is the local mean intensity in the spectral range of the transition ij . It is computed by averaging the specific monochromatic intensity $I_\nu(\mathbf{x}, \hat{\mathbf{n}})$ over all directions ($\hat{\mathbf{n}}$) and integrating it over the line profile $\phi_\nu^{ij}(\mathbf{x})$, such that

$$J_{ij}(\mathbf{x}) = \oint \frac{d\Omega}{4\pi} \int_0^\infty d\nu \phi_\nu^{ij}(\mathbf{x}) I_\nu(\mathbf{x}, \hat{\mathbf{n}}). \quad (9)$$

The collisional part of the transition rates is composed of the collisional rates (K_{ij}^p) for each collision partner (p), weighted by their respective abundances

$$C_{ij}(\mathbf{x}) = \sum_{p \in \mathcal{C}} K_{ij}^p(\mathbf{x}) n^p(\mathbf{x}), \quad (10)$$

where \mathcal{C} is the set of collision partners. The position dependence in the collisional rates stems from their dependence on the local temperature of the gas species.

For the models in this paper, we used the energy levels, Einstein coefficients and collisional de-excitation rates from the Leiden Atomic and Molecular Database³ (LAMDA; Schöier et al. 2005). The collisional excitation rates are computed from the de-excitation rates, assuming the detailed balance relation

$$K_{ji}(\mathbf{x}) = K_{ij}(\mathbf{x}) \frac{g_j}{g_i} \exp\left(\frac{h\nu_{ij}}{k_B T(\mathbf{x})}\right), \quad (11)$$

where g_i denotes the statistical weight of the respective level.

We assume the radiative time-scales to be much smaller than any other time-scale in the system. Hence, to find the level populations given the radiation field, we solve equation (6) in the static limit, i.e. assuming that for all levels i , $\partial n_i(\mathbf{x})/\partial t = 0$. Dropping the position dependence on all variables, the resulting linear equation for each level i can be written as

$$\begin{aligned} & \sum_{j, j < i} \{ n_i A_{ij} - (n_j B_{ji} - n_i B_{ij}) J_{ij} \} \\ & - \sum_{j, j > i} \{ n_j A_{ji} - (n_i B_{ij} - n_j B_{ji}) J_{ij} \} \\ & + \sum_{j=1}^N \{ n_i C_{ij} - n_j C_{ji} \} = 0. \end{aligned} \quad (12)$$

It is important to remember that the radiation field, and thus J_{ij} , depends on the level populations through the line contributions to the emissivity and opacity (see equation 2). This dependence can be expressed mathematically using a Lambda operator. We define this operator such that it yields the mean intensity in the line when acting on the set of all level populations $N \equiv \{n_i(\mathbf{x})\}$, for all positions \mathbf{x} and energy levels i .

$$J_{ij} = \Lambda_{ij}[N]. \quad (13)$$

In most practical cases it is unfeasible to directly invert the Lambda operator (i.e. directly solve the radiative transfer equation (1)) and solve the kinetic rate equation (12) for the level populations. Instead, we solve equation (12) iteratively, by evaluating J_{ij} using

the values from the previous iteration. However, this method, known as Lambda iteration, converges notoriously slowly. Over the years, various methods have been devised to accelerate its convergence (for an overview see e.g. Hubeny & Mihalas 2014, and the references there). We use the operator splitting method (Cannon 1973a,b) in a very similar way to Rybicki & Hummer (1991). The idea is to split the Lambda operator into an approximated part (Λ_{ij}^*) that can easily be evaluated and inverted given the current level populations, and a residual part ($\Lambda_{ij} - \Lambda_{ij}^*$) that can easily be evaluated using the populations of the previous iteration. Hence,

$$J_{ij} = \Lambda_{ij}^*[N] + (\Lambda_{ij} - \Lambda_{ij}^*)[N^\dagger], \quad (14)$$

where the dagger (\dagger) indicates that the quantity is evaluated using the previous iteration. In this way, the contribution of the level populations of the previous iteration can be minimized. The kinetic rate equations (12) can thus be rewritten as

$$\begin{aligned} & \sum_{j, j < i} \{ n_i A_{ij} - (n_j B_{ji} - n_i B_{ij}) (\Lambda_{ij}^*[N] + J_{ij}^{\text{eff}}) \} \\ & - \sum_{j, j > i} \{ n_j A_{ji} - (n_i B_{ij} - n_j B_{ji}) (\Lambda_{ij}^*[N] + J_{ij}^{\text{eff}}) \} \\ & + \sum_{j=1}^N \{ n_i C_{ij} - n_j C_{ji} \} = 0, \end{aligned} \quad (15)$$

where we introduced the effective mean intensity in the line,

$$J_{ij}^{\text{eff}} = (\Lambda_{ij} - \Lambda_{ij}^*)[N^\dagger]. \quad (16)$$

Note that the effective mean intensity is now the only quantity that is evaluated using the level populations of the previous iteration. Clearly, the choice of the approximated Lambda operator (ALO) is essential for the success of this acceleration scheme. In some cases the diagonal part of the Lambda operator already suffices (Olson et al. 1986), however in 3D models, a non-local ALO is often preferred. We discuss our implementation of the ALO, following Rybicki & Hummer (1991), in Section 3.4.

3 NUMERICAL IMPLEMENTATION

3.1 Discretization of the model

The first step in simulating an astrophysical object is finding a way to represent the object on a computer. This comes down to finding an appropriate discretization scheme for all physical parameters of the model. For radiative transfer simulations the spatial, spectral, and directional discretization schemes are most crucial.

3.1.1 Spatial discretization

There are many different types of spatial discretization schemes, each tailored to their specific use cases. Over the years, there has been a clear evolution from structured schemes, like e.g. regular Cartesian grids, to unstructured schemes, like e.g. Voronoi grids or SPH discretizations.

Since we aim to build a general-purpose library that can easily be integrated with other codes, we do not want to tie MAGRITTE to a certain discretization scheme. Instead, we designed our algorithms such that they only require data that can easily be deduced from any discretization scheme, and yet allow us to efficiently trace rays and solve the transfer equation. MAGRITTE's ray-tracing algorithm, presented in Section 3.2, only requires the positions of the cell

³Data base can be found at home.strw.leidenuniv.nl/moldata.

centres⁴ (or equivalently the positions of the particles in an SPH scheme) and the nearest neighbours lists for each cell (or particle). Hence, the input is effectively a point cloud complemented with nearest neighbour information. The boundary of the model can then be defined as the set of points in the convex hull of the point cloud.

In principle, one could use a separate spatial discretization to sample the density, velocity, and temperature distributions of the model. In practice, however, one usually samples all three on the same discretization. This is the case, especially for hydrodynamics computations, where all these parameters should be sampled equally well. However, for radiative transfer computations, especially when considering lines, it is essential to properly sample any changes in the velocity field along a certain line of sight. Since this effect depends on the velocity field in a certain direction it is difficult to fully take this into account in the spatial discretization. In MAGRITTE, when detecting large changes in the velocity field along a ray, we make an appropriate interpolation on-the-fly when ray-tracing, without adjusting the mesh (see also Section 3.2).

3.1.2 Spectral discretization

The requirements for the spectral discretization vary for different stages in the computation. For instance, when determining the level populations, we are only interested in the radiation in the lines, whereas when computing spectra we require a proper frequency sampling over the full spectrum. To accommodate this, MAGRITTE can change its spectral discretization throughout a simulation.

At the stage where the level populations are obtained, the frequency bins are distributed to suit the integration of the radiation field over the line. In MAGRITTE these integrals are evaluated using quadrature formulae. Assuming a Gaussian line profile, the corresponding Gauss–Hermite quadrature for any frequency-dependent function $y(\nu)$ is given by

$$\int_0^\infty d\nu \phi_{\nu}^{ij}(\mathbf{x}) y_{\nu} \rightarrow \sum_{n=1}^{N_q} w_n y(v_{ij} + r_n \delta v_{ij}(\mathbf{x})), \quad (17)$$

where N_q is the number of quadrature points and the quadrature weights are given by

$$w_n = \frac{2^{N_q-1} N_q!}{(N_q H_{N_q-1}(x_n))^2}, \quad (18)$$

H_{N_q-1} is the physicists' version of the Hermite polynomial and the r_n are the roots of the physicists' version of the Hermite polynomial $H_{N_q}(x)$ (see e.g. Abramowitz & Stegun 1972). To be able to easily evaluate these quadratures in MAGRITTE, we define a separate set of frequency bins for each cell, given by

$$\{v_{ij} + r_n \delta v_{ij}(\mathbf{x}), \text{ for each transition } ij \text{ and root } r_n\}, \quad (19)$$

possibly appended with additional frequency bins. Note that this set has to be different for each cell, since it depends on the local line profile width $\delta v_{ij}(\mathbf{x})$.

At the stage where spectra or images are created, extra frequency bins can be appended to the list above to improve the sampling of the spectrum. The current version of MAGRITTE allows one to append a set of user-defined frequency bins, or to add extra bins with a user-defined spacing around each line.

⁴We do not require a strict definition of the cell centre. If we define a cell as a unit in the discretization of the spatial volume, then the cell centre may be any point in that volume. (We only use the cell centre to locate the cell.)

3.1.3 Directional discretization

MAGRITTE is a ray-tracing code, i.e. the radiation field is determined by solving the radiative transfer equation along a set of rays (straight lines) originating from each cell centre. A ray can be defined by a point, in our case the cell centre, and a direction. The direction of the rays will play a key role in scattering and will determine the viewing angles for the images we can take.

In general, there are no preferred directions. Therefore, we discretize the directions uniformly. In 1D and 2D models this is trivial. In 3D, we determine the direction of a ray using the HEALPIX⁵ discretization of the sphere (Gorski et al. 2005). Given a level of refinement, ℓ , it discretizes a unit sphere in $N_{\text{rays}} = 12 \times 4^\ell$ uniformly distributed pixels of equal area. For each pixel, there is an associated unit vector pointing from the origin of the sphere to the pixel centre. These unit vectors determine the directions of MAGRITTE's rays for 3D simulations. Hence, a directional average for a quantity $y(\hat{\mathbf{n}})$ can be translated into an average over rays,

$$\frac{1}{4\pi} \oint d\Omega y(\hat{\mathbf{n}}) \rightarrow \frac{1}{N_{\text{rays}}} \sum_{r=1}^{N_{\text{rays}}} y_r. \quad (20)$$

The uniform directional sampling scheme bears the danger of missing the contributions of very localized sources of emissivity or opacity. Furthermore, there might be situations in which there is one or more preferred directions, and one might better consider a non-uniform distribution of the ray directions. Therefore, in future versions, we will investigate more advanced directional weighting schemes. In any case, the internal structure of MAGRITTE allows for any distribution of rays, allowing us to easily explore various directional distribution schemes in the future.

3.2 Ray tracing

In order to solve the radiative transfer equation along a certain ray, the emissivity and opacity of the cells that are encountered along that ray must be known. Furthermore, the path-length that the ray traces through each cell must be computed. All this must be done assuming only a point cloud with nearest neighbour information.

The idea of MAGRITTE's ray-tracing algorithm is to walk along the ray from one cell to the next and determine the path-length through each cell by projecting the cell centres on to the ray. To determine which cell is next, the set of all nearest neighbours of the current cell is considered. From this set the neighbour is chosen which is closest to the ray and whose projection on the ray lies farther than that of the current cell. This procedure is then repeated until the boundary of the mesh is reached. Fig. 1 shows a visual example of how this algorithm works. Once the rays have been traced, the transfer equation can be solved.

In each step from one cell to the next, the change in velocity along the ray is computed and checked for large variations. If the velocity field, and thus the resulting Doppler shift, changes too much the emissivity and opacity are interpolated between the cells such that the velocity steps are only a certain (user defined) fraction of the local line width. In this way, we avoid losing or improperly accounting for line contributions due to an inadequate sampling of the velocity field.

⁵Source code for HEALPIX can be found at healpix.sourceforge.net.




Figure 1. A visual representation of MAGRITTE’s ray-tracing algorithm for the ray R originating from cell O . The goal is to find which cells are encountered along the ray and hence which cell centres should be projected on the ray. Clearly, the cell O itself lies on the ray. The next cell encountered is the neighbour of O that lies closest to the ray. We call this cell P_1 . Now the next cell to be projected is the neighbour of P_1 that lies closest to the ray and that is further away from O than P_1 . The last condition is there to ensure that one proceeds along the ray towards the boundary. This process is repeated until the boundary of the mesh is reached.

3.3 Solving the transfer equation along a ray pair

In MAGRITTE, the radiation field is obtained by solving the radiative transfer equation along each pair of a ray and its antipode through the model. Although in this paper we are only concerned with line radiative transfer, we present our general solution method for the full radiative transfer equation (1) including scattering. In this way, the treatment in this paper more closely resembles the actual implementation in the code and paves the way for our future work.

For numerical stability, we solve the transfer equation in its second-order form as suggested by Feautrier (1964). We define the mean intensity-like (u) and flux-like quantity (v) along a ray as

$$\begin{aligned} u_v(\hat{n}) &\equiv \frac{1}{2}(I_v(\hat{n}) + I_v(-\hat{n})), \\ v_v(\hat{n}) &\equiv \frac{1}{2}(I_v(\hat{n}) - I_v(-\hat{n})), \end{aligned} \quad (21)$$

to describe the radiation field. To simplify notation further on, we also define new quantities to represent the scattering redistribution function and scattering opacity up and down the ray

$$\begin{aligned} \Phi_{v\nu'}^\pm(\hat{n}, \hat{n}') &\equiv \frac{1}{2}(\Phi_{v\nu'}(\hat{n}, \hat{n}') \pm \Phi_{v\nu'}(-\hat{n}, \hat{n}')), \\ \chi_v^\pm(\hat{n}) &\equiv \frac{1}{2}(\chi_v^{\text{sca}}(\hat{n}) \pm \chi_v^{\text{sca}}(-\hat{n})). \end{aligned} \quad (22)$$

Finally, to avoid lengthy integral equations we define

$$\Psi_v^\pm(\hat{n}) \equiv \oint d\Omega' \int_0^\infty dv' \Phi_{v\nu'}^\pm(\hat{n}, \hat{n}') I_{v'}(\hat{n}'). \quad (23)$$

From here onwards, we drop all ν and \hat{n} dependencies for notational simplicity. We proceed by adding and subtracting the transfer equation (1), once for \hat{n} and once for $-\hat{n}$. This yields a coupled set of first-order differential equations in u and v ,

$$\begin{aligned} \hat{n} \cdot \nabla v &= -(\chi + \chi^+)u - \chi^-v + \Psi^+ + \eta, \\ \hat{n} \cdot \nabla u &= -(\chi + \chi^+)v - \chi^-u + \Psi^-. \end{aligned} \quad (24)$$

Solving the equations, once for u and once for v , yields a set of second-order differential equations, which are our generalized

versions of the Feautrier equations (Feautrier 1964)

$$\begin{aligned} (1 - \mathcal{D}^2) u &= \frac{\eta + \Psi^+}{\chi + \chi^+} - \mathcal{D} \left(\frac{\Psi^-}{\chi + \chi^+} \right), \\ (1 - \mathcal{D}^2) v &= \frac{\Psi^-}{\chi + \chi^+} - \mathcal{D} \left(\frac{\eta + \Psi^+}{\chi + \chi^+} \right), \end{aligned} \quad (25)$$

where we defined a new differential operator \mathcal{D} as

$$\mathcal{D} \equiv \frac{1}{\chi + \chi^+} (\chi^- + \hat{n} \cdot \nabla). \quad (26)$$

Note that the order of the factors in the definition of \mathcal{D} is important, since they do not commute. Both equations in (25) are still coupled through their Ψ^\pm terms. However, the contributions between the scattering opacity, $\chi_v^{\text{sca}}(\hat{n})$, and the scattering redistribution function, $\Phi_{v\nu'}(\hat{n}, \hat{n}')$, can be arranged such that $\Phi_{v\nu'}(\hat{n}, \hat{n}) = 0$. Hence, the coupling between the equations in (25) can be weakened.

The generalized Feautrier equations (25) can thus be solved in an iterative way by evaluating their right-hand sides using the solution of the previous iteration. In each iteration, two separate ordinary second-order differential equations then have to be solved. The boundary conditions can be determined from the incoming radiation on both sides of the ray pair. By making a Taylor expansion of the intensity, the incoming radiation can be related to u and v , in the same way as in the standard Feautrier procedure with improved boundary conditions for plane-parallel geometries by Auer (1967).

In MAGRITTE, the radiation field is computed by solving the equations (25) for each ray pair. Since the right-hand sides are treated in an iterative way, these act effectively as sources. The second-order differential operators on the left-hand sides will result in tridiagonal matrices on the discretized ray pairs. The form of the equations in (25) still resembles the original Feautrier equations enough that the standard solution method (Feautrier 1964) with the numerical improvements by Rybicki & Hummer (1991) can readily be adapted to this generalized case.

3.4 Accelerated lambda iteration

From the first equation in (25) we can identify the Lambda operator for our solution scheme, as defined in equation (13)

$$\Lambda_{ij}[N] = \mathcal{L}_{ij} \left[\frac{\eta + \Psi^+}{\chi + \chi^+} - \mathcal{D} \left(\frac{\Psi^-}{\chi + \chi^+} \right) \right], \quad (27)$$

where we used the auxiliary linear operator \mathcal{L}_{ij} , defined as

$$\mathcal{L}_{ij}[\cdot] = \frac{1}{2} \oint \frac{d\Omega}{4\pi} \int_0^\infty dv \phi_v^{ij} (1 - \mathcal{D}_v^2)^{-1} [\cdot]. \quad (28)$$

Following Rybicki & Hummer (1991), we can construct an approximation to the Lambda operator by considering only the diagonal band of the matrix representation of the auxiliary operator \mathcal{L}_{ij} . We call this operator \mathcal{L}_{ij}^* . The operator \mathcal{L}_{ij}^* is easy enough to invert, due to its band diagonal structure. However, using it as the ALO would render (15) into a system of non-linear equations for the level populations, which would still be hard to solve. In order to retain the linearity of (15), we instead define our ALO as

$$\Lambda_{ij}^*[N] = \frac{n_j^\dagger B_{ji} - n_i^\dagger B_{ij}}{n_j B_{ji} - n_i B_{ij}} \mathcal{L}_{ij}^* \left[\frac{\eta^{ij}(N)}{\chi(N^\dagger) + \chi^+(N^\dagger)} \right], \quad (29)$$

where we only evaluate the line emissivity in the argument of \mathcal{L}_{ij}^* with the new level populations and add an extra factor which goes to unity when the level populations converge. Since the line emissivity is linear in the level populations (see equation 2) the statistical

equilibrium equation will remain linear in the new level populations and can be written as

$$\begin{aligned} & \sum_{j, j < i} \{ n_i A_{ij} - \tilde{\Lambda}_{ij}^*[n_i] - (n_j B_{ji} - n_i B_{ij}) J_{ij}^{\text{eff}} \} \\ & - \sum_{j, j > i} \{ n_j A_{ji} - \tilde{\Lambda}_{ji}^*[n_j] - (n_i B_{ij} - n_j B_{ji}) J_{ij}^{\text{eff}} \} \\ & + \sum_{j=1}^N \{ n_i C_{ij} - n_j C_{ji} \} = 0, \end{aligned} \quad (30)$$

where the effective mean intensity, defined in (14), is given by

$$J_{ij}^{\text{eff}} = J_{ij} - \mathcal{L}_{ij}^* \left[\frac{\eta^{ij}(N^\dagger)}{\chi(N^\dagger) + \chi^+(N^\dagger)} \right] \quad (31)$$

and we introduced another auxiliary approximated operator

$$\tilde{\Lambda}_{ij}^*[n_i] = \frac{h}{4\pi} (n_j^\dagger B_{ji} - n_i^\dagger B_{ij}) \mathcal{L}_{ij}^* \left[\frac{A_{ij} v \phi_v^{ij} n_i}{\chi(N^\dagger) + \chi^+(N^\dagger)} \right]. \quad (32)$$

Note that this operator is linear in the new level populations and not symmetric in the level indices i and j . However, since our ALO (29) is symmetric in these indices, it can be implemented on the level of the transition matrix.

Apart from the ALO, we also use Ng-acceleration (Ng 1974) in MAGRITTE to speed-up convergence even more. This acceleration method introduces a special iteration step every M number of (regular) iterations. In this special iteration step, the level populations of the next iteration are predicted by a linear combination of the populations of the previous M iterations. This is done by minimizing the change in the level populations for the prediction based on the last $M - 1$ and the one to last $M - 1$ iterations. Since the Ng-iteration step does not require the computation of the radiation field, it is much faster than a regular iteration and thus accelerates the iteration process. The Ng-method allows us to specify a weight for the contribution of the different levels to the prediction (see e.g. Olson et al. 1986). For this paper, a uniform weighting scheme was applied, but MAGRITTE can readily be adapted to handle any other scheme.

3.5 Imaging the model

When modelling astrophysical objects in 3D, one often requires images of the model from several viewpoints in several frequency ranges, in order to compare the model with observations.

In MAGRITTE, these images can be obtained using the solution of the outward directed radiation field on the endpoints of each ray pair. One can construct the image by considering the outward directed radiation along a certain ray and projecting the locations of the originating points on the plane orthogonal to the ray. The result is a set of points on a plane with a corresponding intensity, which can be easily rendered into an image. Note that since every point in the model contributes to one point in the image, the spatial resolution of the resulting image is exactly equal to the highest achievable resolution for that model.

4 TESTS AND BENCHMARKS

To demonstrate the validity of our methods and to better understand their limitations, we have conducted a series of comparisons with analytical models and benchmarked against established radiative transfer codes. The analysis for these tests and benchmarks was

performed in a collection of JUPYTER notebooks (Kluyver et al. 2016), which are publicly available on GitHub.⁶

4.1 Semi-analytical tests

To assess the accuracy of MAGRITTE's ray tracer and radiative transfer solver, we first reproduce some semi-analytically solvable line radiative transfer models. This will help us later to better assess the uncertainties associated to the results of our simulations. We cannot overemphasize the importance of these analytical tests as they are the only way to obtain absolute measures of the accuracy.

4.1.1 Homogeneous Hubble-Lemaître models

As a first test, we consider the radiative transfer of a single line on a uniformly spaced grid with a constant molecular abundance and temperature distribution, and a constant velocity gradient. The velocity distribution is thus given by the Hubble–Lemaître law

$$v(r) = c \Delta\beta r, \quad (33)$$

where we parametrized the velocity gradient $\Delta\beta$ as a fraction of the speed of light c . The boundary condition is given by incoming cosmic microwave background (CMB) radiation, i.e a blackbody spectrum B_ν of temperature $T_{\text{CMB}} = 2.725$ K. If we assume LTE level populations, the line source function S_{vij} is spatially constant. In that case, one can find the mean intensity by directly integrating the transfer equation, yielding

$$J_\nu(\mathbf{x}) = S_{vij} + (B_\nu - S_{vij}) \oint \frac{d\Omega}{4\pi} e^{-\tau_\nu(\mathbf{x}, \hat{\mathbf{n}})}, \quad (34)$$

where the optical depth, assuming Gaussian line profiles centred around ν_{ij} and with a line profile width $\delta\nu_{ij}$, is given by

$$\tau_\nu(\ell) = \frac{\chi_{ij}}{2\nu\Delta\beta} \left\{ \text{Erf} \left[\frac{\nu - \nu_{ij}}{\delta\nu_{ij}} \right] - \text{Erf} \left[\frac{\nu(1 - \Delta\beta\ell) - \nu_{ij}}{\delta\nu_{ij}} \right] \right\}, \quad (35)$$

where Erf is the error function, and $\ell(\mathbf{x}, \hat{\mathbf{n}})$ is the distance from point \mathbf{x} to the boundary, as measured along the ray in direction $\hat{\mathbf{n}}$. Since the Hubble–Lemaître velocity law is both translation and rotation invariant, only the total distance to the boundary appears in the expression for the optical depth.

Considering only a single ray in the interval $[-R, R]$, the mean intensity in $r \in [0, R]$, as expressed in equation (34), reads

$$J_\nu(r) = S_{vij} + \frac{1}{2} (B_\nu - S_{vij}) [e^{-\tau_\nu(r)} + e^{-\tau_\nu(R-r)}], \quad (36)$$

where the average over all directions reduces to the average intensity flowing up and down the ray.

In three dimensions, assuming a spherical boundary with radius R , the mean intensity expressed in equation (34) reduces to

$$J_\nu(r) = S_{vij} + \frac{1}{2} (B_\nu - S_{vij}) \int_0^\pi d\theta \sin\theta e^{-\tau_\nu(\ell(r, \theta))}, \quad (37)$$

where the distance to the boundary $\ell(r, \theta)$ is given by

$$\ell(r, \theta) = r \cos\theta + \sqrt{R^2 - r^2 \sin^2\theta}. \quad (38)$$

The θ -integral in the expression for the mean intensity can easily be computed numerically. Note that introducing the spherical boundary breaks the translation invariance of the problem.

⁶Benchmarks can be found at github.com/Magritte-code/Benchmarks.

Table 1. Line data of the fictitious 2-level species. This is the same fictitious 2-level species as used in Problem 1 in van Zadelhoff et al. (2002).

$E_2 - E_1$ (cm $^{-1}$)	g_2/g_1	A_{21} (s $^{-1}$)	$K_{21}^{\text{H}_2}$ (cm 3 s $^{-1}$)
6.0	3.0	1.0×10^{-4}	2.0×10^{-10}

Table 2. MAGRITTE parameters for the semi-analytic test models.

Model		(N_{shells})	N_{cells}	N_{rays}	N_{q}
Hubble–Lemaître	1D	50	100	2	100
	3D	50	12 528	192	100
Density distribution	1D	50	100	2	100
	3D	50	12 528	192	100

Although these are simple models, they can demonstrate some key issues in numerical radiative transfer modelling. In particular, both models can be used to directly assess the accuracy of the radiative transfer solver and to test the sampling in velocity space. Especially in line radiative transfer, it is crucial to properly sample the velocity field, since too large a step in velocity from one cell to the next can Doppler-shift a line directly from one wing to the other without capturing the effect of the core of the line. This can be tested by adjusting the velocity gradient. By considering both the single ray and full 3D model we can also assess the quality of the spatial interpolations on to the rays.

For this test we used a fictitious 2-level species in a (radially) uniformly spaced grid $[-R, R]$ with $R = 495$ km, and with a velocity gradient $c\Delta\beta = 0.01$ s $^{-1}$. The line data for the fictitious 2-level species are summarized in Table 1. We assume a constant H $_2$ number density of $n^{\text{H}_2} = 1.0 \times 10^{12}$ m $^{-3}$ and a constant fractional abundance of the fictitious 2-level species $X \equiv n^{\text{fict}}/n^{\text{H}_2} = 10^{-4}$. To obtain the level populations, we assume LTE with a constant temperature distribution $T = 45$ K. Furthermore, we assume the gas has no turbulent velocity component. The 3D model is obtained from the 1D model by mapping each 1D grid point to a shell of 3D grid points uniformly distributed over a sphere. The model parameters for MAGRITTE can be found in Table 2.

Fig. 2 shows a comparison between the solution of MAGRITTE and the semi-analytical solutions (36, 37) of the Hubble–Lemaître models. MAGRITTE’s numerical result clearly agrees with the analytic solution with a relative error well below 10^{-4} almost everywhere, where the relative error of two values is measured as twice the absolute difference over their sum.

4.1.2 Simple power-law density distribution

As a second semi-analytic test, we consider the radiative transfer of a single line on a logarithmically spaced grid, with a constant temperature distribution, with no velocity field, and a (spherically) symmetric density distribution given by a power-law

$$n^{\text{H}_2}(r) = \begin{cases} 0 & \text{for } r < R_{\text{in}} \\ n^{\text{H}_2}(R_{\text{in}}) \left(\frac{R_{\text{in}}}{r}\right)^2 & \text{for } r \geq R_{\text{in}} \end{cases}, \quad (39)$$

where r_{in} is the inner radius of the model. The boundary condition is again given by incoming CMB radiation, i.e. a blackbody spectrum B_ν of temperature $T_{\text{CMB}} = 2.725$ K. If we again assume LTE level populations, the line source function $S_{\nu_{ij}}$ is spatially constant. As a result the mean intensity is again given by equation (34). To compute




Figure 2. Comparison between MAGRITTE and the semi-analytic solution of the mean intensity as a function of frequency in the Hubble–Lemaître model, evaluated at different radii. The dots indicate MAGRITTE’s results and the line represent the analytic results. Frequencies are expressed with respect to the line centre $\nu_{21} \approx 179.88$ GHz as a fraction of the line profile width $\delta\nu_{21} \approx 519.03$ kHz. The relative error of two values is measured as twice the absolute difference over their sum.

the optical depth, one needs to integrate the density distribution along every ray. Assuming a spherical boundary with radius R , the optical depth is given by

$$\tau_\nu(r, \theta) = \frac{\chi_{ij} \phi_v^{ij} r}{\sin \theta} \left(\frac{\pi}{2} - \theta + \arccos \left(\frac{r \sin \theta}{R} \right) - f(r, \theta) \right), \quad (40)$$

where the function $f(r, \theta)$ accounts for the rays that go through the empty core ($r < R_{\text{in}}$) of the model and is given by

$$f(r, \theta) = \begin{cases} 2 \arccos \left(\frac{r \sin \theta}{R_{\text{in}}} \right) & \text{for } \theta < \theta_{\text{core}} \\ 0 & \text{for } \theta \geq \theta_{\text{core}} \end{cases} \quad (41)$$

and whether or not a ray passes through the empty core is determined by the direction of the ray at each radius, $\theta_{\text{core}} = \arcsin(R_{\text{in}}/r)$.

Considering only a single ray in the interval $[-R, R]$, the mean intensity is given by

$$J_\nu(r) = S_{\nu_{ij}} + \frac{1}{2} (B_\nu - S_{\nu_{ij}}) [e^{-\tau_\nu(r,0)} + e^{-\tau_\nu(r,\pi)}], \quad (42)$$

where the average over all directions reduces to the average intensity flowing up and down the ray. Note that one should be careful in taking the limits $\theta \rightarrow 0$ and $\theta \rightarrow \pi$, but that both are well-defined.

In three dimensions, one can simply integrate over the entire solid angle to obtain the mean intensity

$$J_\nu(r) = S_{\nu_{ij}} + \frac{1}{2} (B_\nu - S_{\nu_{ij}}) \int_0^\pi d\theta \sin \theta e^{-\tau_\nu(r,\theta)}. \quad (43)$$

However, one should be careful of distinguishing between rays that do and do not pass through the empty core of the model.

For this test we used the same fictitious 2-level species as before (Table 1) in a radially logarithmically spaced grid $[-R, R]$ with $R_{\text{in}} = 1.0 \times 10^{13}$ m and $R = 7.8 \times 10^{16}$ m, and without a velocity field. The H $_2$ number density just outside the empty core is $n^{\text{H}_2}(R_{\text{in}}) = 2.0 \times 10^{13}$ m $^{-3}$ and a constant fractional abundance of

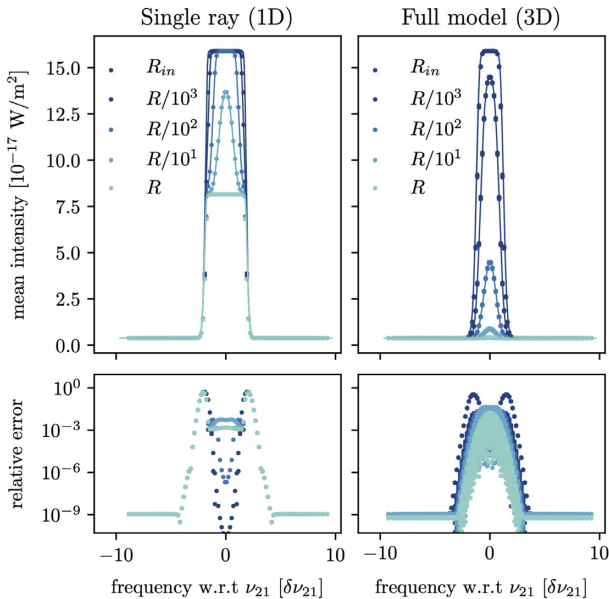


Figure 3. Comparison between MAGRITTE and the semi-analytical solution of the mean intensity as a function of frequency in a model with a simple power-law density distribution, evaluated at different radii. The dots indicate MAGRITTE’s results and the line represent the analytic results. Frequencies are expressed with respect to the line centre $\nu_{21} \approx 179.88$ GHz as a fraction of the line profile width $\delta\nu_{21} \approx 357.53$ kHz. The relative error of two values is measured as twice the absolute difference over their sum.

the fictitious 2-level species $X \equiv n^{\text{fict}}/n^{\text{H}_2} = 10^{-6}$ is used. To obtain the level populations, we assume LTE with a constant temperature distribution $T = 20$ K. Furthermore, the gas has a turbulent velocity component of $v_{\text{turb}} = 150 \text{ m s}^{-1}$. The 3D model is obtained from the 1D model by mapping each 1D grid point to a shell of 3D grid points uniformly distributed over a sphere. The model parameters for MAGRITTE can be found in Table 2. This model setup is identical to Problem 1b in van Zadelhoff et al. (2002). However, here we are only interested in the resulting radiation field when the levels are in LTE (see also Section 4.2).

Fig. 3 shows a comparison between the solution of MAGRITTE and the semi-analytical solutions (42, 43) of the simple power-law density distribution models. MAGRITTE’s numerical result clearly agrees with the analytic solution. Only at the steep edges of the line is there a larger relative error (~ 0.4), which can be attributed to the steepness of the solution and the discrete mesh.

4.2 Cross-code benchmarks

There are no analytic solutions for the full non-LTE line radiative transfer problem, so the only way to fully test MAGRITTE’s line radiative transfer module is by benchmarking it against established codes. Although this does not prove the validity of the code, it is reassuring to find the same results in different ways.

For the benchmarks we used the (1D) problems presented in van Zadelhoff et al. (2002) (from here on referred to as the benchmark paper) and compared our results with the publicly available version of the 1D Monte Carlo radiative transfer code RATRAN⁷ (Hogerheijde & van der Tak 2000). Since MAGRITTE is intrinsically multidimensional, the 1D benchmarking models were

⁷Source code can be found at personal.sron.nl/~vdtak/ratran/frames.html.

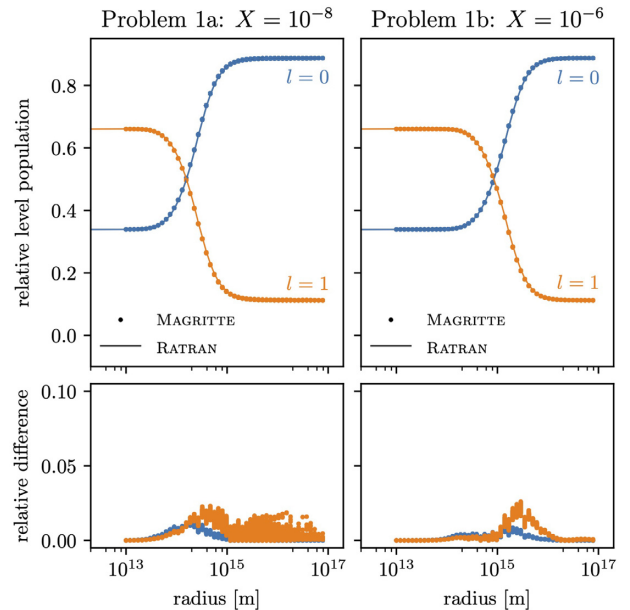


Figure 4. Comparison of the results for Problem 1 a/b of the van Zadelhoff et al. (2002) benchmark obtained with MAGRITTE (dots) and RATRAN (lines). The relative difference of two values is measured as twice the absolute difference over their sum.

mapped to their 3D equivalents by mapping each 1D grid point to a shell of 3D grid points uniformly distributed over a sphere.

4.2.1 Van Zadelhoff Problem 1 a/b

The first test, referred to as Problem 1 a/b in the benchmark paper, considers a fictitious two-level species in a spherically symmetric cloud, without velocity field, with a constant temperature distribution, and a density distribution given by a power law. The entire model can thus be defined analytically. The model setup is essentially the same as in the simple power-law density distribution test in Section 4.1.2. The only difference is that in Problem 1 a the relative molecular abundance $X = 10^{-8}$ results in a low optical depth, whereas in Problem 1 b the relative molecular abundance is $X = 10^{-6}$, yielding a relatively high optical depth. The model parameters for MAGRITTE can be found in Table 3.

Fig. 4 shows a comparison between the resulting level populations for Problem 1 a/b obtained with MAGRITTE and RATRAN. Both are clearly in good agreement.

4.2.2 Van Zadelhoff Problem 1 c/d

Since line radiative transfer models critically depend on a proper sampling of the velocity field along the line of sight of each ray, it is worthwhile to test if this is properly accounted for. Therefore, we consider again benchmark Problem 1 a/b from the previous paragraph, but this time with a non-zero velocity field. Although this test was not part of the van Zadelhoff et al. (2002) benchmark, we can still compare our results with RATRAN. We consider a velocity field that is pointing radially outward, given by

$$v(r) = v_{\infty} \left(\frac{r - R_{\text{in}}}{R - R_{\text{in}}} \right)^{\gamma} \hat{r}. \quad (44)$$

¹³In the benchmarks below we used $\gamma = 0.5$ and since it is the same model setup as in Sections 4.1.2 and 4.2.1 the inner radius

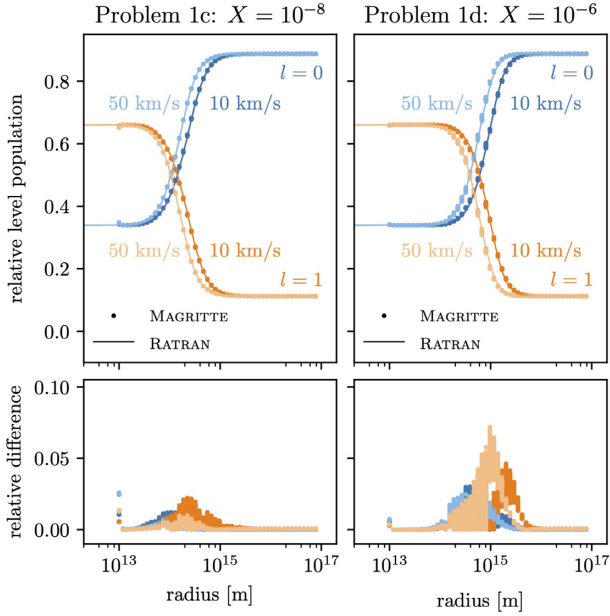


Figure 5. Comparison of the results for Problem 1 c/d obtained with MAGRITTE (dots) and RATRAN (lines). The indicated velocities are the v_∞ for each model. The relative difference of two values is measured as twice the absolute difference over their sum.

is $R_{\text{in}} = 1.0 \times 10$ m. Furthermore, we consider two different terminal velocities $v_\infty = 10$ and 50 km s^{-1} . The model parameters for MAGRITTE can be found in Table 3.

Fig. 5 shows a comparison between the resulting level populations for Problem 1 c/d obtained with MAGRITTE and RATRAN. Both are clearly in good agreement. However, in order to obtain this result, we needed to increase the number of grid points in the input for RATRAN by a factor of 10 (resulting in 500 logarithmically spaced grid points). For any lower number of grid points, RATRAN had difficulty properly sampling the velocity field and produced significantly different results from MAGRITTE.

4.2.3 Van Zadelhoff Problem 2 a/b

The third test has a more realistic setup and considers the lines of HCO^+ in a snapshot of an inside-out collapse model by Shu (1977). This is referred to as Problem 2 a/b in the benchmark paper. The parameters describing the input model were taken from the website of the benchmark⁸. The model consists of 50 logarithmically spaced grid points. In each grid point the radial velocity, gas temperature, micro-turbulence, and HCO^+ and H_2 abundances are given. Again there are two cases, one with a relatively low optical depth where the fractional HCO^+ abundance is $X = 10^{-9}$ and one with a relatively high optical depth where the relative molecular abundance is $X = 10^{-8}$. The model parameters for MAGRITTE can be found in Table 3.

Fig. 6 shows a comparison between the results for Problem 2 a/b obtained with MAGRITTE and RATRAN. Overall, both codes agree well with relative differences below 0.3 for the first five levels. This is comparable to what Brinch & Hogerheijde (2010) find in their

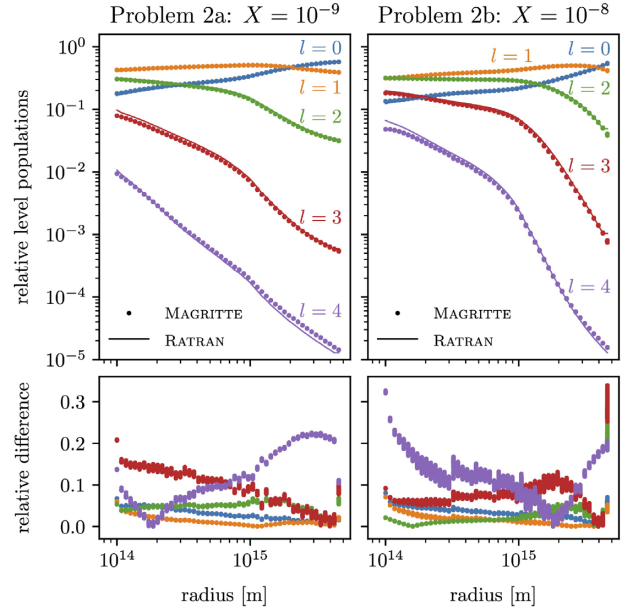


Figure 6. Comparison of the results of the first five levels (of 41) for Problem 2 a/b of the van Zadelhoff et al. (2002) benchmark obtained with MAGRITTE (dots) and RATRAN (lines). The relative difference of two values is measured as twice the absolute difference over their sum.

Fig. 10 for LIME and what Rundle et al. (2010) find in their Figs 2 and 3 for TORUS. Furthermore, Rundle et al. (2010) report that for $l = 0$ their relative deviation from the benchmark paper is less than 5 per cent, which is also comparable to what we find.

5 APPLICATION

To demonstrate the applicability of MAGRITTE in a more realistic setup, we consider the CO line radiative transfer in a simple Keplerian disc model. This is a typical use case of 3D radiative transfer modelling (see e.g. Homan et al. 2018; Booth et al. 2019). The density distribution in cylindrical (r, ϕ, z) co-ordinates is described by

$$\rho(r, \phi, z) = \begin{cases} 0 & \text{for } r < r_{\text{in}} \\ \rho_{\text{in}} \left(\frac{r}{r_{\text{in}}}\right)^p \exp\left[-\frac{1}{2}\left(\frac{z}{H(r)}\right)^2\right] & \text{for } r \geq r_{\text{in}}, \end{cases} \quad (45)$$

with $p = -2.125$, and a vertical Gaussian scale height given by

$$H(r) = r_{\text{in}} \sqrt{\frac{k_{\text{B}} T_{\text{in}}}{m_{\text{H}_2} G M_\star}} \left(\frac{r}{r_{\text{in}}}\right)^h, \quad (46)$$

where $h = 1.125$ and m_{H_2} is the mass of H_2 . The fractional CO abundance is a constant $n_{\text{CO}} = 5.0 \times 10^{-4}$. Furthermore, we assume a gas temperature distribution given by a power law

$$T(r, \phi, z) = T_\star \left(\frac{r}{r_\star}\right)^q, \quad (47)$$

in which we take $q = -0.5$, and a Keplerian velocity field

$$v(r, \phi, z) = \sqrt{\frac{GM_\star}{r}} \hat{\phi}. \quad (48)$$

The remaining physical parameters of the star and the disc are summarized in Tables 4 and 5, respectively.

⁸Benchmark website: strw.leidenuniv.nl/astrochem/radtrans/.

Table 3. MAGRITTE parameters for the benchmark models.

Model	(N_{shells})	N_{cells}	N_{rays}	N_q
Problem 1 a/b/c/d	50	23 280	192	24
Problem 2 a/b	50	23 280	192	24

Table 4. Parameters of the star in the Keplerian disc.

$M_{\star} (M_{\odot})$	$T_{\star} (K)$	$r_{\star} (AU)$
2.0	2500	2.0

Table 5. Parameters of the Keplerian disc.

$\rho_{\text{in}} (\text{kg m}^{-3})$	$T_{\text{in}} (K)$	$r_{\text{in}} (AU)$
5.0×10^{-12}	1500	10.0

Since we do not yet have a fully implemented algorithm to generate model meshes (see Section 6.4.1), we currently use the sampling algorithm and Voronoi mesher implemented in LIME⁹ (Brinch & Hogerheijde 2010).

Fig. 7 shows 16 channel maps of the CO $J = 6 - 5$ transition in an edge-on view of the Keplerian disc model produced by MAGRITTE. From left to right and top to bottom, one can clearly see the left half of the disc moving away from the observer and being red-shifted, whereas the right half of the disc is moving towards the observer and being blue-shifted.

Fig. 8 shows a composite image stacking 16 channel maps depicted in Fig. 7, as well as the relative integrated intensity for each of the channel maps as a function of the velocity along the line of sight with respect to the rest-frame of the observer. The relative integrated intensities are normalized with respect to the maximum integrated intensity of the channel-maps.

6 DISCUSSION

6.1 ALI and Ng acceleration schemes

Both ALIs and the Ng acceleration scheme are used in MAGRITTE to ensure correct results and reduce the computation time. This is done to avoid false convergence and to reduce the number of required iterations in computing the non-LTE level populations.

In general, a wider (i.e. more non-local) band diagonal matrix ALO yields a better approximation to the Lambda operator and thus will yield better convergence. However, when a non-local ALO is used, the resulting level populations from solving equation (30) cannot be guaranteed to be positive, and thus can become unphysical (Rybicki & Hummer 1991). This becomes apparent, in particular, when the solution is far from converged or when a larger bandwidth is used. It is hard to determine in advance whether a certain bandwidth for the ALO will result in unphysical level populations. Even the simplest two-cell model with a two-level species can easily be made to fail. Hence, when using a non-local ALO, one should always check the validity of the level populations after solving the statistical rate equations. If for a certain ALO, the

computation yields unphysical level populations, one can always set up and solve the system of rate equations again using only a part of the ALO. Setting up and solving the statistical rate equations only takes a fraction of the time required to compute the radiation field and the corresponding ALO. Therefore, the trade-off that should be considered in deciding the bandwidth of the ALO is the gained reduction in iterations versus the time it costs to compute the extra off-diagonal elements. Unfortunately, there is no generally applicable (problem independent) way to make this trade-off since the number of required iterations strongly depends on the model under consideration. Hauschildt, Störzer & Baron (1994) recommended optimal bandwidths for some typical model set-ups, however, these only apply to their particular implementation. By default, MAGRITTE will use a diagonal (i.e. local) ALO. Larger bandwidth ALOs can be used when specifically requested. However, MAGRITTE will always check the validity of the resulting level populations and reduce the bandwidth if required. Moreover, given a model, MAGRITTE can predict how the computation time of one iteration would change when changing the bandwidth of the ALO. This should help users to decide on an appropriate bandwidth for their particular model.

The effectiveness of the Ng acceleration scheme depends on the quality of the solutions in previous iterations. Therefore, it is advisable to start acceleration only after a certain level of convergence is already reached. Also, the optimal balance between regular and Ng accelerated iteration steps is highly problem dependent. It is, however, less critical than the choice of ALO bandwidth since the computational cost of taking into account more iterations in the Ng acceleration scheme is negligible (compared to the cost of computing the radiation field and the corresponding ALO). By default, MAGRITTE will perform an Ng acceleration step after every four regular iterations, and use the level populations of all four previous iterations.

6.2 Accuracy, precision, and re-sampling invariance

The results of the semi-analytic tests and cross-code benchmarks in Section 4 clearly demonstrate the accuracy and precision of the radiative transfer and level population solver of MAGRITTE. All models were run for various numbers of cells, rays, and frequency bins. Also the uniform distributions of the mesh points over the spherical shell was varied, as well as the relative distributions of the mesh points over the shells. These variations did not induce any significant differences in the results of MAGRITTE, demonstrating its re-sampling invariance. The parameters of the models presented in this paper are all about five times larger than the coarsest model that produces reasonable results.

The van Zadelhoff et al. (2002) Problem 1 benchmarks with the additional velocity field, presented in Section 4.2.2, emphasize the importance of a proper sampling of the velocity field. In MAGRITTE this is automatically taken care of by the ray-tracer which will interpolate the source and optical depth if the velocity changes too rapidly, whereas in RATRAN the user has to provide a 10 times finer model mesh in order to obtain accurate results (comparing Figs 5 and A1). This is acceptable for a 1D solver such as RATRAN, because the sampling of a 1D velocity field can still be assessed with comparative ease by the user. In 3D, however, velocity structures can become extremely complex along the various lines of sight. Therefore, the on-the-fly assessment and interpolation of the velocity field, as implemented in MAGRITTE, is essential to ensure accurate results.

⁹Source code can be found at github.com/lime-rt/lime.




Figure 7. Channel maps with contours of the CO $J = 6-5$ transition in an edge-on view of the Keplerian disc model.

Fig. 9 shows a comparison between the results of MAGRITTE and LIME¹⁰ (Brinch & Hogerheijde 2010) for the van Zadelhoff et al. (2002) benchmark Problem 1 a for different numbers of grid points. The bottom plots show the errors on the solutions, i.e. the relative difference with respect to the RATRAN solution, which is assumed to be the most accurate. Both solvers used the exact same model mesh. In order to do this, the model was first run with LIME which created the mesh that could then also be used by MAGRITTE. The results for both solvers are plotted after the same number of iterations (which was around 30). The exact number was determined by the number of iterations that MAGRITTE required to reach a relative change in level populations below 10^{-7} . In order to make a fair comparison the Ng acceleration in MAGRITTE was disabled and only a local (diagonal)

ALO was used. Nevertheless, the results of MAGRITTE are clearly more accurate than the results of LIME. This is apparent especially for the coarser meshes with fewer mesh points. Furthermore, the results of MAGRITTE are more precise, i.e. less spread at a given radius, than the results of LIME. This can be attributed to the Monte Carlo noise present in the solution of LIME.

Fig. 10 shows a comparison between the results of MAGRITTE and LIME for the van Zadelhoff et al. (2002) benchmark Problem 1 b, which is similar to Problem 1 a, but has a higher optical depth. Both solvers again use the same model mesh consisting of 30 000 grid points. The plot shows the results after 147 iterations (when MAGRITTE reached a relative change in level populations below 10^{-7}). Clearly, the result of MAGRITTE is much more accurate than the result of LIME. This can partly be attributed to the mesh, since also the results of MAGRITTE are slightly worse than the results obtained on the mesh of shells in the comparison with RATRAN (see Section 4.2.1). Nevertheless, there is clear discrepancy in the

¹⁰Throughout this paper any reference to LIME refers to (currently latest) release version 1.9.5 (see github.com/lime-rt/lime/releases).

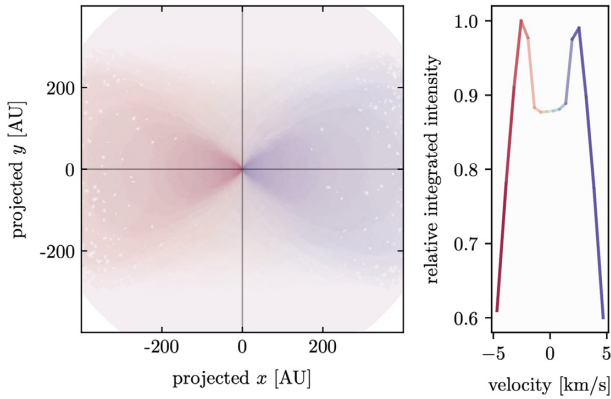


Figure 8. Composite image of the channel maps of the CO $J = 6-5$ transition in the edge-on view of the Keplerian disc model (*left-hand panel*) and the relative integrated intensity for each of the channel maps as a function of velocity (*right-hand panel*).

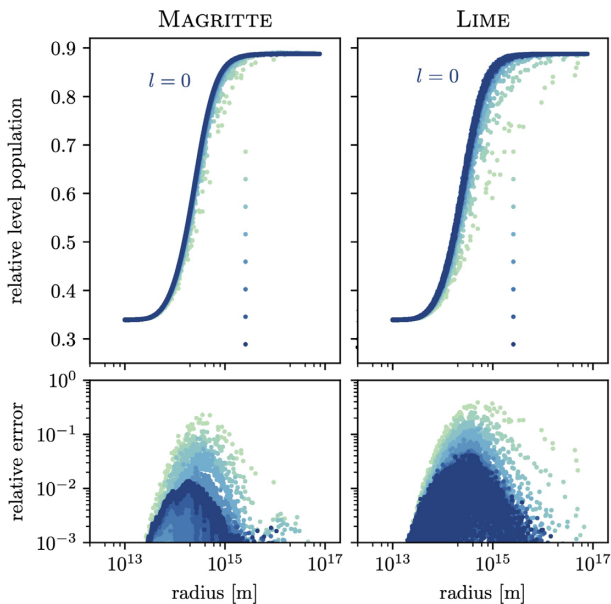


Figure 9. Comparison of the results for Problem 1 a of the van Zadelhoff et al. (2002) benchmark obtained with MAGRITTE and LIME for different numbers of grid points. The relative error of two values is measured as twice the absolute difference with respect to the solution of RATRAN over their sum.

accuracy that can be achieved with MAGRITTE and LIME in this high optical depth problem. We tried increasing the number of mesh points with a factor of 10 but did not see any improvement in the performance of LIME.

6.3 Computational performance

MAGRITTE was especially designed to achieve good scalability of performance on modern distributed computer architectures and to leverage hardware acceleration. However, since a significant part of radiative transfer research is mainly performed on commercial workstations (laptops and desktops), MAGRITTE should also perform well on these (shared memory) systems. Fig. 11 shows the




Figure 10. Comparison of the results for Problem 1 b of the van Zadelhoff et al. (2002) benchmark obtained with MAGRITTE and LIME. The relative error of two values is measured as twice the absolute difference with respect to the solution of RATRAN (dotted line) over their sum.




Figure 11. Plot of the (strong) scaling of MAGRITTE's parallelization for shared-memory systems. The dots indicate the relative timings and the grey line indicates the ideal scaling behaviour. Timings are averages over seven runs performed on one 32-core Intel Skylake node of the CSD3 cluster. (Hyper-threading was disabled for these runs such that the number of threads effectively corresponds to the number of cores used.)

preliminary (strong) scaling of MAGRITTE for the Keplerian disc model of Section 5 on a shared memory system (32-core Intel Skylake, with hyper-threading disabled). The fact that the run time is almost perfectly inversely proportional to the number of threads shows that MAGRITTE can both effectively and efficiently use the available computational resources. The (strong) scaling in Fig. 11 is only preliminary in the sense that no effort was made to ensure load balancing over the cores, which could improve the scaling. Future versions of MAGRITTE will include an active load balancing algorithm to ensure good (strong) scaling results independent of the model geometry (De Ceuster et al. in preparation).

To gauge the computational speed of MAGRITTE, we performed and timed the van Zadelhoff et al. (2002) benchmark Problem 1 (see also Section 4.2.1) with the established 3D radiative transfer

code LIME and compared the results with MAGRITTE. We used the Voronoi model mesh produced by LIME as input for MAGRITTE to ensure that both solvers got the exact same input. To further ensure a fair comparison we disabled the Ng acceleration in MAGRITTE and only used a local (diagonal) ALO. We performed 35 iterations (which in MAGRITTE corresponded to a relative change in level populations below 10^{-7}). We found that MAGRITTE was about 1.6 times faster than LIME¹¹ on the same mesh for the same number of iterations. This is mainly due to the implementation with the explicit vectorization and despite the fact the formal solver used in MAGRITTE is more precise and intrinsically slower than the one in LIME. Note that we only measured the time spent in the computation of the radiation field and the level population solver and not the time spent in creating, reading, or writing the model mesh. Considering that MAGRITTE can already obtain accurate and precise results for much coarser grids (see Section 6.2), we could conclude that MAGRITTE is more than 1.6 times faster. However, how much more largely depends on the required accuracy and is hard to compare between MAGRITTE and LIME because of the intrinsic difference in precision.

When large velocity fields are included, MAGRITTE will be slightly slower than LIME because of its careful treatment of the Doppler shifts along each ray. However, this careful treatment is required to obtain accurate results (see Section 6.2). The current implementation heavily prioritizes accuracy over speed. In future versions, the new meshing algorithm will allow us much better control over the accuracy of the radiative transfer solver, which will allow us to better balance the trade-off between accuracy and speed.

6.4 Future development of MAGRITTE

This paper only reports on the first step in the development of MAGRITTE. The code base is still under active development and will be extended and improved over the next few years. The design strategy will be twofold, on the one hand focusing on developing a complete radiative transfer library with a complete modelling pipeline to confront simulations with observations, and on the other hand achieving higher performance by leveraging modern computer architectures.

6.4.1 Meshing algorithm

The next step towards a complete radiative transfer modelling library will be to develop a mesher to generate 3D meshes for a given model (distributions for density, velocity, temperature, etc.). Since MAGRITTE is a formal solver we can build the mesh based on the numerical error it will induce in the radiative transfer computations. The algorithms used by the mesher can then also be used to assess and improve model meshes coming from hydrodynamical simulations. A strong handle on the model mesh and thus on the induced numerical errors will greatly improve the accuracy and reliability of radiative transfer computations.

6.4.2 Including more physics

The next piece of physics to be included in MAGRITTE will be to account for scattering within the existing radiative transfer solver.

Once we can account for scattering, we can do dust continuum radiative transfer and include a thermal balance module to iteratively determine the dust temperature. Later, we will focus on coupling MAGRITTE with (photo)chemistry and hydrodynamics codes to provide fully self-consistent radiation-hydro-chemical models.

6.4.3 Computational aspects

The initial motivation to develop MAGRITTE was to create a general-purpose software library for 3D radiative transfer, that could leverage modern computer architectures, such as highly distributed systems with accelerators (e.g. GPUs and FPGAs), to improve the performance of astrophysical and cosmological simulations. Therefore, MAGRITTE was vectorized and parallelized for both shared and distributed memory systems, and can off-load certain computations to accelerators. The full optimization and parallelization strategy will be presented in a forthcoming paper. All future releases of MAGRITTE and its source code, including the optimized and accelerated versions, will be made publicly available¹ at github.com/Magritte-code.

7 CONCLUSIONS

In this first paper in a series on MAGRITTE: a modern open-source software library for 3D radiative transfer modelling, we presented and tested its non-LTE line radiative transfer module. MAGRITTE uses a deterministic ray-tracer and formal solver that computes the radiation field by (iteratively) solving the radiative transfer equation along a fixed set of rays originating from each point. The ray-tracing algorithm only requires the locations of the cell centres and the nearest neighbour lists. Hence, it can readily be applied to SPH particles, as well as structured and unstructured model meshes. We formulated an elegant solution method for the second-order form of the radiative transfer equation along a ray pair based on Feautrier (1964) and Cannon (1971, 1972), treating the scattering contributions from other rays in an iterative way. Furthermore, we presented our implementation of the ALI scheme by Rybicki & Hummer (1991) in this context. We demonstrated the validity of MAGRITTE by comparing its results against both semi-analytical model solutions and the established (1D) radiative transfer solver of RATRAN (Hogerheijde & van der Tak 2000) on the van Zadelhoff et al. (2002) benchmark for line radiative transfer. As an example application, we used MAGRITTE to generate channel maps of CO lines in a simple Keplerian disc model. Comparing our results with the established 3D radiative transfer solver LIME (Brinch & Hogerheijde 2010), we conclude that MAGRITTE produces more accurate and more precise results, especially at high optical depth, and that it is faster.

ACKNOWLEDGEMENTS

We are very grateful to Michiel Hogerheijde and Daniel Harsono for fruitful discussions on the line radiative transfer benchmarks and the comparison with RATRAN and LIME. We also would like to thank the anonymous referee for their helpful suggestions. FDC is supported by the EPSRC iCASE studentship programme, Intel Corporation and Cray Inc. FDC, WH, and LD acknowledge support from the ERC consolidator grant 646758 AEROSOL. This work was performed using the Cambridge Service for Data Driven Discovery (CSD3), part of which is operated by

¹¹As measured with eight threads on the eight cores of a decent but standard laptop with an Intel Core i7-7700HQ CPU clocked at 2.8 GHz.

the University of Cambridge Research Computing on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The DiRAC component of CSD3 was funded by BEIS capital funding via STFC capital grants ST/P002307/1 and ST/R002452/1 and STFC operations grant ST/R00689X/1. DiRAC is part of the National e-Infrastructure.

REFERENCES

- Abramowitz M., Stegun I. A., 1972, *Handbook of Mathematical Functions*. Dover Publications, New York.
- Altay G., Croft R. A., Pelulessy I., 2008, *MNRAS*, 386, 1931
- Alves F. O., Caselli P., Girart J. M., Segura-Cox D., Franco G. A. P., Schmiedeke A., Zhao B., 2019, *Science*, 366, 90
- Andersson B.-G., Lazarian A., Vaillancourt J. E., 2015, *ARA&A*, 53, 501
- Andrews S. M. et al., 2012, *ApJ*, 744, 162
- Auer L., 1967, *ApJ*, 150, L53
- Bergin E. A. et al., 2013, *Nature*, 493, 644
- Bisbas T. G., Bell T. A., Viti S., Yates J., Barlow M. J., 2012, *MNRAS*, 427, 2100
- Booth A. S., Walsh C., Ilee J. D., 2019, *A&A*, 629, A75
- Brinch C., Hogerheijde M. R., 2010, *A&A*, 523, A25
- Camps P., Baes M., 2018, *ApJ*, 861, 80
- Cannon C., 1973a, *J. Quant. Spectrosc. Radiat. Transfer*, 13, 627
- Cannon C. J., 1971, *Aust. J. Phys.*, 24, 341
- Cannon C. J., 1972, *Aust. J. Phys.*, 25, 177
- Cannon C. J., 1973b, *ApJ*, 185, 621
- Chandrasekhar S., 1960, *Radiative Transfer*. Dover Publications Inc, New York
- Decin L., Richards A. M. S., Neufeld D., Steffen W., Melnick G., Lombaert R., 2015, *A&A*, 574, A5
- Douglas A. E., Herzberg G., 1941, *ApJ*, 94, 381
- Dullemond C. P., Juhasz A., Pohl A., Sereshti F., Shetty R., Peters T., Commercon B., Flock M., 2012, RADMC-3D: A multi-purpose radiative transfer tool, Astrophysics Source Code Library, record ascl:1108.016
- Evans M. G., Hartquist T. W., Caselli P., Boley A. C., Ilee J. D., Rawlings J. M. C., 2019, *MNRAS*, 483, 1266
- Feautrier P., 1964, *C. R. Acad. Sci.*, 258, 3189
- Frostholm T., Haugbølle T., Grassi T., 2018, preprint ([arXiv:1809.05541](https://arxiv.org/abs/1809.05541))
- Gorski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
- Harries T. J., Haworth T. J., Acreman D., Ali A., Douglas T., 2019, *Astron. Comput.*, 27, 63
- Hauschildt P. H., Störzer H., Baron E., 1994, *J. Quant. Spectrosc. Radiat. Transfer*, 51, 875
- Höfner S., Gautschi-Loidl R., Aringer B., Jørgensen U. G., 2003, *A&A*, 399, 589
- Hogerheijde M. R., van der Tak F. F. S., 2000, *A&A*, 64, 137
- Homan W., Danilovich T., Decin L., De Koter A., Nuth J., Van De Sande M., 2018, *Astron. Astrophys.*, 614, A113
- Hubeny I., Mihalas D., 2014, *Theory of Stellar Atmospheres*. Princeton University Press, Princeton.
- Huggins P. J., Glassgold A. E., 1982, *AJ*, 87, 1828
- Kluyver T. et al., 2016, in Loizides F., Schmidt B., eds, *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, Vol. 16. IOS Press, Amsterdam, p. 87
- Maercker M. et al., 2012, *Nature*, 490, 232
- Mihalas D., Athay R. G., 1973, *ARA&A*, 11, 187
- Montargès M. et al., 2019, *MNRAS*, 485, 2417
- Ng K. C., 1974, *J. Chem. Phys.*, 61, 2680
- Noebauer U. M., Sim S. A., 2019, *Living Rev. Comput. Astrophys.*, 5, 1
- Olson G. L., Auer L. H., Buchler J. R., 1986, *J. Quant. Spectrosc. Radiat. Transfer*, 35, 431
- Rundle D., Harries T. J., Acreman D. M., Bate M. R., 2010, *MNRAS*, 407, 986
- Rybicki G. B., Hummer D. G., 1991, *A&A*, 245, 171
- Scharmer G. B., Carlsson M., 1985, *J. Comput. Phys.*, 59, 56
- Schöier F. L., van der Tak F. F. S., van Dishoeck E. F., Black J. H., 2005, *A&A*, 432, 369
- Shu F. H., 1977, *ApJ*, 214, 488
- Smith N., Ginsburg A., Bally J., 2018, *MNRAS*, 474, 4988
- Steinacker J., Baes M., Gordon K. D., 2013, *ARA&A*, 51, 63
- van Zadelhoff G.-J. et al., 2002, *A&A*, 395, 373
- Vandenbroucke B., Wood K., 2018, *Astron. Comput.*, 23, 40
- Verstocken S., Van De Putte D., Camps P., Baes M., 2017, *Astron. Comput.*, 20, 16
- Walsh C. et al., 2016, *ApJ*, 823, L10
- Weinreb S., Barrett A. H., Meeks M. L., Henry J. C., 1963, *Nature*, 200, 829
- Wenzel J., Rhineland J., Moldovan D., 2017, pybind11 – Seamless operability between C++ and Python. Available at: <https://github.com/pybind/pybind11>
- Woitke P., Krueger D., Sedlmayr E., Woitke P., Krueger D., Sedlmayr E., 1996, *A&A*, 311, 927
- Yusef-Zadeh F., Morris M., White R. L., 1984, *ApJ*, 278, 186

APPENDIX A: ADDITIONAL FIGURE

In this appendix we present an additional figure supporting our claims regarding the accuracy, precision, and re-sample-invariance of MAGRITTE with respect to RATRAN.

Fig. A1 shows a comparison between the results obtained with MAGRITTE (in 3D) and RATRAN (in 1D) for a mesh with 50 shells, in contrast to Fig. 5 in the main body of the paper where the results for a 50 shell mesh for MAGRITTE was compared to a 500 shell mesh for RATRAN. The relative differences for the coarser RATRAN model are about four times larger than for the finer model. This is due to the insufficient sampling of the velocity field which MAGRITTE can and RATRAN cannot account for.

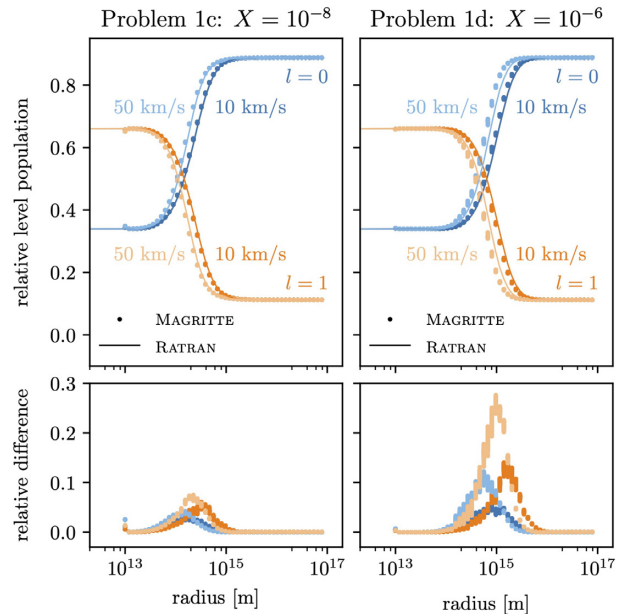


Figure A1. Comparison of the results for Problem 1 c/d, obtained with MAGRITTE (dots) and RATRAN (lines) both on a model mesh with 50 shells. The indicated velocities are the v_∞ for each model. The relative difference of two values is measured as twice the absolute difference over their sum.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.