Astronomy & Astrophysics manuscript no. aa62638-26
October 22, 2026
©ESO 2026

FastSpec: Fast and versatile stellar spectrum synthesis code

Oleg Kochukhov*

Department of Physics and Astronomy, Uppsala University, Box 524, SE-751 20 Uppsala, Sweden

Received 18 August 2026 / Accepted 14 October 2026

ABSTRACT

Context. Modelling data from modern spectroscopic surveys and time-resolved spectropolarimetric observations requires stellar synthetic spectra to be computed rapidly over wide wavelength ranges from line lists that increasingly reach hundreds of millions of transitions, and for stars spanning the entire Hertzsprung-Russell diagram. Spectropolarimetric analyses of magnetically active stars additionally demand a solution of the polarised radiative transfer problem.

Aims. We present FastSpec, a new local-thermodynamic-equilibrium (LTE) spectrum synthesis code designed around two goals: increased raw computational speed and broad astrophysical applicability, including the calculation of the full Stokes vector in the presence of a magnetic field.

Methods. FastSpec computes emergent intensity and flux spectra from plane-parallel and spherical model atmospheres. Chemical and ionisation equilibrium is obtained from the updated Spectroscopy Made Easy (SME) equation-of-state package, continuum opacities follow the Atlas family of routines, hydrogen bound-bound and bound-free absorption uses the occupation-probability formalism, and line absorption is evaluated with fast table-based Voigt and Faraday-Voigt profile algorithms. The transfer equation is integrated with DELO-Bézier formal solvers in both scalar and Stokes-vector form. Speed is achieved through aggressive OpenMP parallelisation, adaptive per-line wavelength windowing, and a memory-capped block-processing scheme that lets arbitrarily large line lists be handled within a user-defined RAM budget.

Results. The code synthesises spectra thousands of Å wide from line lists of 10^8 molecular and atomic transitions in under a minute on a modern multi-core laptop while reproducing benchmark calculations from established codes. We demonstrate the code's performance and versatility on the solar spectrum, on the near-infrared molecular bands of an M dwarf, on the four Stokes parameters of a magnetic star, and on the hydrogen lines of hot stars.

Conclusions. FastSpec combines the physical completeness of classical model atmosphere and spectrum synthesis codes with the speed required by large surveys and the polarised transfer needed for stellar magnetism studies into a single portable package.

Key words. radiative transfer – polarization – methods: numerical – stars: atmospheres – stars: magnetic field – line: profiles

1. Introduction

The comparison of observed stellar spectra with synthetic ones computed from model atmospheres is the backbone of quantitative stellar spectroscopy. It underlies the determination of atmospheric parameters and chemical abundances of individual stars, the interpretation of data collected by large spectroscopic surveys, and the mapping of stellar surface structures through Doppler and Zeeman-Doppler imaging. A number of mature codes address this spectrum synthesis problem in the local thermodynamic equilibrium (LTE) approximation, among them Atlas/Synthe (Kurucz 2005), MOOG (Sneden 1973), Turbospectrum (Plez 2012; Gerber et al. 2023), Spectrum (Gray & Corbally 1994), Synspec (Hubeny & Lanz 2017), and Spectroscopy Made Easy (SME, Valenti & Piskunov 1996; Piskunov & Valenti 2017). Codes that additionally solve the polarised transfer equation for magnetic stars include Zeeman (Landstreet 1988; Wade et al. 2001), MOOGStokes (Deen 2013), ZeeTurbo (Cristofari et al. 2023), Synthmag and Synmast (Piskunov 1999; Kochukhov 2007; Kochukhov et al. 2010), and the Invers family (Piskunov & Kochukhov 2002; Kochukhov et al. 2014) used for magnetic mapping.

Three developments have progressively raised the computational demands placed on such codes. First, spectroscopic surveys now deliver observations for thousands of stars so that even modest per-star costs aggregate into large computing budgets, and iterative parameter determination multiplies the number of syntheses required. Second, atomic and molecular line databases have grown enormously. The VALD database (Piskunov et al. 1995; Kupka et al. 1999; Ryabchikova et al. 2015) contains several million atomic transitions, while molecular line lists from the ExoMol project (Tennyson et al. 2024) routinely reach hundreds of millions of lines per molecule, which must be processed to model the spectra of cool stars, brown dwarfs, and exoplanets. Third, spectropolarimeters attached to large telescopes produce time series of all four Stokes parameters, whose modelling requires many polarised syntheses per rotation phase. At the same time, robust interpretation of mean polarisation profiles commonly employed in spectropolarimetry demands full Stokes parameter synthesis over broad wavelength regions (Kochukhov 2026). Recent codes such as Korg (Wheeler et al. 2023) have shown that careful algorithm design and modern hardware can accelerate synthesis by large factors; the same considerations motivate the present work.

In this paper we describe FastSpec, a spectrum synthesis code written in modern Fortran and built explicitly for speed and versatility. Its central design choices are (i) all computationally heavy loops – over wavelengths, depths, and lines – are parallelised with OpenMP; (ii) each spectral line contributes only within an adaptively determined wavelength window so that the cost scales with the number of significant line-wavelength inter-

* Corresponding author: oleg.kochukhov@physics.uu.se

actions rather than with the product of list size and grid size; (iii) the memory footprint is capped by a user-defined limit, above which the line list is processed in blocks, allowing lists of essentially unlimited size to be handled on modest hardware; and (iv) profile functions are evaluated from precomputed lookup tables. When a magnetic field and a line list with Zeeman-splitting information are supplied, FastSpec solves the polarised transfer problem and returns the full Stokes $IQUV$ spectrum.

The physical ingredients of FastSpec are, wherever possible, inherited from well-tested public codes: the equation of state from SME (Valenti & Piskunov 1996; Piskunov & Valenti 2017), the continuum opacities from the Atlas12 code (Kurucz 2005), and the hydrogen line and bound-free package Hbop/Hlinop (Barklem & Piskunov 2003). The originality of FastSpec lies in the way these ingredients are restructured, parallelised, and combined with the adaptive windowing and memory management, which make large-scale, magnetic synthesis fast.

The remainder of the paper is organised as follows. Section 2 gives an overview of the code structure and data flow. In Sect. 3 we describe the physical foundations: the equation of state, continuum and hydrogen opacities, line absorption and broadening, the Zeeman effect, and the radiative transfer and disk-integration algorithms. Section 4 presents the algorithmic ingredients responsible for the code's speed together with performance benchmarks. In Sect. 5 we summarise the input data formats and in Sect. 6 validate the code against reference calculations and illustrate its range of applications. In Sect. 7 we summarise the paper. The input keywords and output file formats are documented in the appendices.

2. Code overview

FastSpec is written in Fortran 2008[1] and parallelised with OpenMP. This choice allows the well-tested routines of established codes to be reused with minimal changes. Compiled Fortran code also remains considerably faster than interpreted languages such as Python for computationally expensive tasks. The code is organised into a small number of modules with clearly separated responsibilities: shared data structures and physical constants; general utilities; input and output; the equation of state; continuum, hydrogen, and helium opacities; the line-absorption kernel and profile functions; the radiative transfer solvers; and the disk-integration and convolution routines. A high-level driver orchestrates these components in the sequence shown in Fig. 1.

A calculation proceeds through the following stages.

1. *Initialisation.* The input parameter file is parsed, and the model atmosphere is read. The wavelength grid is constructed, either with a constant step in wavelength or a constant step in velocity (i.e. equidistant in $\ln\lambda$). The union of all chemical species appearing in the requested line lists, augmented by a fixed set of continuum absorbers, is assembled, and the equation of state is solved at every atmospheric depth. Lookup tables for the Voigt and, if needed, Faraday-Voigt functions are computed.
2. *Continuum opacity.* The continuous absorption and scattering coefficients are computed for every depth and wavelength.
3. *Line opacity.* The requested line lists are processed one at a time. For each list the lines are read, weak lines are rejected against a line-to-continuum threshold, an adaptive wavelength window is determined for every retained line, and the line opacity is accumulated onto the wavelength grid. Hydrogen lists and magnetic lists are handled by dedicated paths (Sects. 3.4 and 3.6).
4. *Radiative transfer.* The transfer equation is integrated for the requested set of limb angles μ, yielding the emergent intensity spectrum (scalar case) or the emergent Stokes vector (magnetic case) at each μ.
5. *Disk integration and convolution.* If a flux spectrum is requested, the intensity spectra are disk-integrated with rotational broadening and radial-tangential macroturbulence, and optionally convolved with a Gaussian instrumental profile.
6. *Output.* The selected products – line-strength list, line and continuum opacities, intensity spectra, and flux spectrum – are written as unformatted binary stream files.

The code has a second mode of operation intended for interactive work. Given a previously computed intensity-spectrum file and a set of broadening parameters, it performs only the disk integration and convolution described in step 5 above. This makes it inexpensive to explore the effect of projected rotational velocity, macroturbulence, and spectral resolution without recomputing the intensity spectra.

3. Physical foundations

FastSpec assumes a static atmosphere in LTE, in either plane-parallel or spherical geometry. The emergent spectrum is obtained by solving the transfer equation along rays of direction cosine $\mu = \cos\theta$ through a model atmosphere specified on a column-mass (ϱx) depth scale. In this section we describe, in the order in which they enter the calculation, the equation of state, the opacity sources, the treatment of the Zeeman effect, and the radiative transfer and disk-integration algorithms.

3.1. Model atmospheres and the wavelength grid

Model atmospheres are read in the `.krz` format.[2] The header specifies the effective temperature $T_{\rm eff}$, the surface gravity $\log g$, a model-type flag, the standard wavelength $\lambda_{\rm std}$ at which reference opacities are tabulated, and a set of switches selecting the continuum opacity sources. The header is followed by the chemical composition, given as the abundances of the 99 elements from H to Es, and by the depth-dependent structure: the column mass ϱx, temperature T, electron number density $n_{\rm e}$, total atomic number density $n_{\rm a}$, and mass density ρ at each of the N_τ depth points. The format is directly compatible with the Atlas9 (Kurucz 2005), Marcs (Gustafsson et al. 2008), and LLmodels (Shulyak et al. 2004) models routinely used in the field. A spherical model atmosphere is identified from its header, which then also specifies the stellar radius and, as an additional depth column, the geometrical height above the reference radius. These inputs define the extended geometry in which the transfer equation is solved (Sect. 3.7).

The default chemical composition can be overridden in two ways: a global metallicity [M/H] can scale the solar abundance pattern of all elements heavier than helium, or individual element abundances can be specified explicitly. The reference solar abundances are those of Asplund et al. (2021), and the atomic weights follow the compilation distributed with Turbospectrum.[3]

[1] See https://fortran-lang.org for an overview of modern Fortran.

[2] https://marcs.oreme.org/documents/krzformat

[3] https://github.com/bertrandplez/Turbospectrum_NLTE

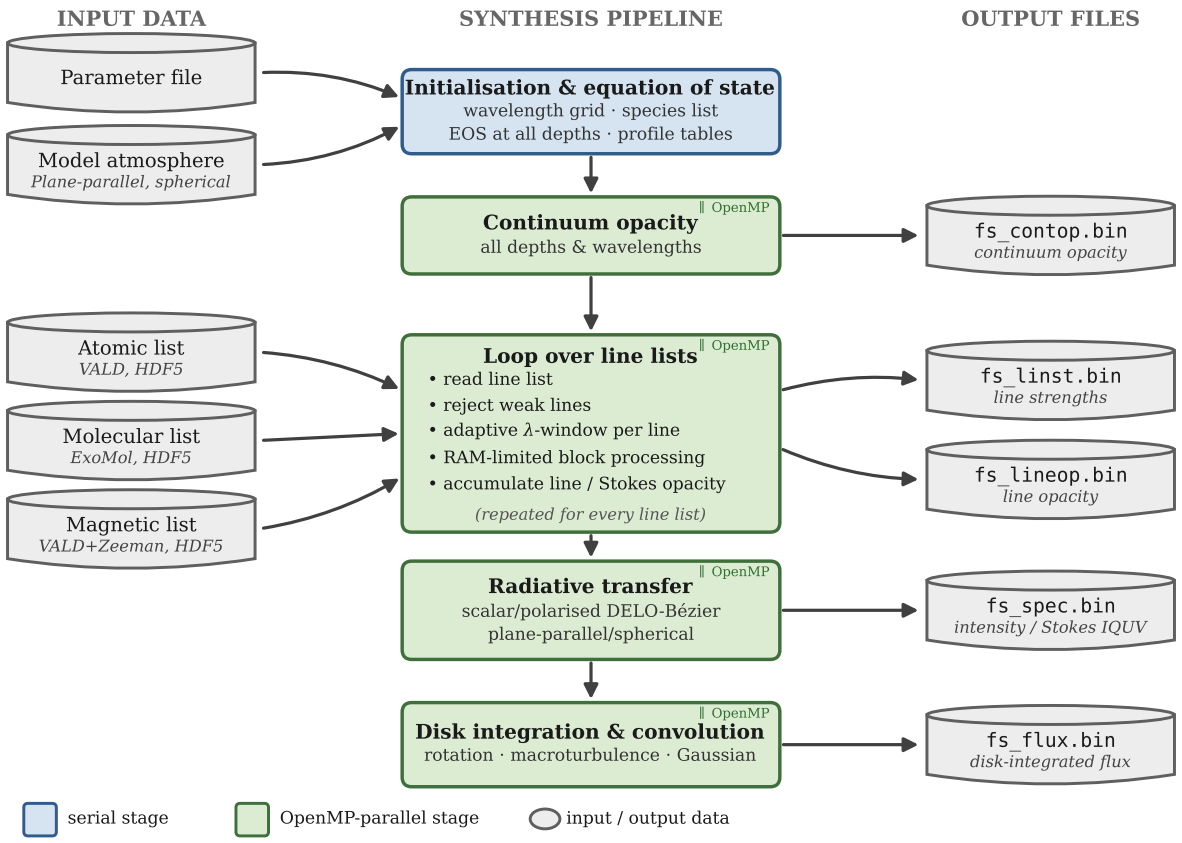


Fig. 1: Schematic data flow of FastSpec. The input data (left) are the parameter file; the model atmosphere; and the atomic, molecular, and magnetic line lists. The output products (right) are the binary line-strength, line-opacity, continuum-opacity, intensity, and flux files. The central pipeline runs from top to bottom through the initialisation and equation of state, the continuum opacity, the loop over line lists (weak-line rejection, adaptive per-line wavelength windowing, and RAM-limited block processing), the scalar or polarised radiative transfer, and the disk integration and convolution. Green boxes mark the OpenMP-parallel stages, and the blue box indicates the (largely serial) initialisation.

The wavelength grid is defined by its end points and a step that is interpreted either as a constant wavelength increment (in Å) or as a constant velocity increment (in $\mathrm{km\,s^{-1}}$), the latter giving a grid that is uniform in $\ln\lambda$ and thus natural for radial-velocity and rotational-broadening work. All calculations use air wavelengths; line lists supplied with vacuum wavelengths are converted to air when they are read.

3.2. Equation of state

Chemical and ionisation equilibrium is solved with the equation-of-state package[4] of SME (Valenti & Piskunov 1996; Piskunov & Valenti 2017), incorporated with small modifications and bug fixes. The package handles a default set of 497 species: the neutral atoms and their first two ions for the 99 elements from H to Es (only the first ion for H), about 190 molecules, and 15 negative ions, plus free electrons. The equation of state itself supports up to six ionisation stages (neutral to five times ionised), as required by the high-ionisation continuum opacity package (Sect. 3.3). Species that occur in the input line lists but not in the default set are appended to the species list automatically during the inventory pass over the lists (Sect. 5) so that their number densities are available to the line-opacity routines. The molecular partition functions and equilibrium constants are based on the compilation of Barklem & Collet (2016), supplemented by the data from Sauval & Tatum (1984) along with direct fits to ExoMol partition functions for several molecules.

At each atmospheric depth, the code solves the coupled system of particle-conservation equations for every element, the charge-neutrality equation, and the total-pressure equation by Newton-Raphson iteration, with an outer loop that enforces self-consistency of the electron pressure. For temperatures above 12 000 K, molecules are neglected and the Saha equation is used directly; for cooler layers the full molecular equilibrium is solved, using logarithmic partial pressures below 1500 K for numerical stability. The solver returns, for every species, the ratio n_i/U_i of number density to partition function and the partition function U_i itself; the number density is recovered as their product. This factorisation is convenient because the line-absorption coefficient (Sect. 3.5) depends on n_i/U_i directly. The electron density, total particle density, and mass density returned by the solver replace the values read from the model atmosphere, ensuring thermodynamic consistency between the equation of state and the opacities.

3.3. Continuum opacity

The continuous absorption and scattering coefficients follow the implementation in the Atlas12 code of Kurucz (2005), trans-

[4] https://github.com/SpectroscopyMadeEasy/SMElib

Table 1: Continuum opacity sources in FastSpec.

Process	Species / method
Bound-free, free-free	H I, H^-, H_2^+
Bound-free, free-free	He I, He II, He^-
Metal bound-free (cool)	C I, Mg I, Al I, Si I, Fe I
Metal bound-free (warm)	C II, N I, O I, Mg II, Si II, Ca II
High-ion bound-free/free-free	C, N, O, Ne, Mg, Si, S, Fe
Photodissociation	CH, OH
Collision-induced abs.	H_2-H_2, H_2-He
Rayleigh scattering	H I, He I, H_2
Thomson scattering	free electrons

Notes. Each source can be switched on or off individually through the opacity flags in the model-atmosphere header.

lated to Fortran 90 and reorganised so that all temperature-dependent quantities that do not depend on wavelength are precomputed once per model, before the parallel loop over wavelengths. The individual sources, each controlled by a switch inherited from the model-atmosphere header, are listed in Table 1. They comprise bound-free and free-free absorption by H I, H^-, H_2^+, He I, He II, and He^-; bound-free absorption by the astrophysically important neutral and singly ionised metals (C, N, O, Mg, Al, Si, Ca, Fe); photodissociation of CH and OH; collision-induced absorption by H_2-H_2 and H_2-He pairs; the high-ionisation metal opacities relevant for hot stars (up to five times ionised); Rayleigh scattering by H I, He I, and H_2; and Thomson scattering by free electrons. All number densities are taken from the equation of state, and all bound-free and free-free terms include the stimulated-emission factor $1 - e^{-h\nu/kT}$.

Where the calculation also treats hydrogen lines with the occupation-probability package (Sect. 3.4), the bound-free contribution of the low-lying hydrogen levels is removed from the continuum to avoid double counting. The continuum is otherwise identical whether or not a hydrogen line list is present. Figure 2 illustrates the relative contributions of the principal sources across the ultraviolet, optical, and near-infrared for a hot ($T_{\rm eff}$ = 20 000 K) and for a solar model, obtained directly from the code by activating one source at a time.

3.4. Hydrogen line and bound-free opacity

Hydrogen lines and the associated bound-free absorption are treated jointly in the occupation-probability formalism, using the Hbop/Hlinop package[5] of Barklem & Piskunov (2003), restructured for FastSpec so that the level populations and all line-, depth-, and (line, depth)-dependent quantities are precomputed once for fast, thread-safe evaluation over the wavelength grid. The formalism follows Däppen et al. (1987) and Hummer & Mihalas (1988), with the charged-perturber term treated as in Hubeny et al. (1994). An occupation probability w_n is assigned to each bound level n. It describes the gradual dissolution of the upper levels by the plasma microfield and provides a physically consistent merging of the bound-bound and bound-free opacity across the series limits so that the pseudo-continuum between the highest resolved lines and the true continuum is captured without discontinuities.

The Stark broadening of the hydrogen lines uses quasi-static Holtsmark profiles following Griem (1960) with corrections that approximate the unified profiles of Vidal et al. (1973), supplemented by impact electron broadening (Griem 1967), a long-range normalisation correction (Stehlé & Hutcheon 1999), and a

[5] https://github.com/barklem/hlinop

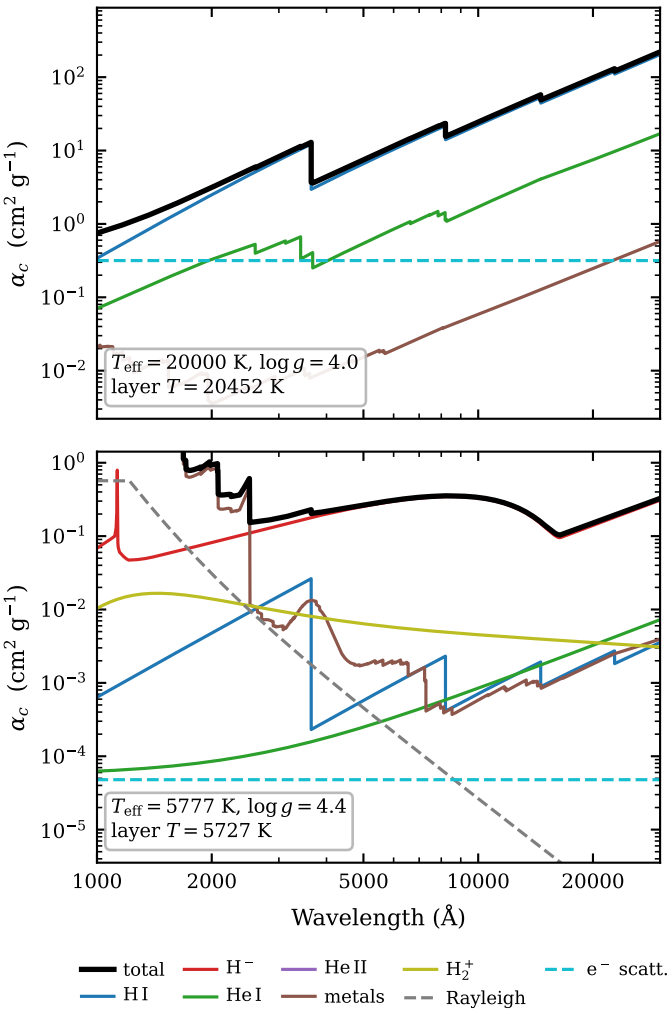


Fig. 2: Continuum absorption coefficient per gram, α_c, at a representative photospheric layer (temperature closest to $T_{\rm eff}$) decomposed into the individual sources of Table 1 for a hot model (*top*; $T_{\rm eff}$ = 20 000 K, log g = 4.0) and the Sun (*bottom*; $T_{\rm eff}$ = 5777 K, log g = 4.44). The black curve is the total absorption coefficient; the coloured curves are the isolated sources (solid: true absorption; dashed: Rayleigh and Thomson scattering), each computed by FastSpec with only the corresponding opacity switch active.

Boltzmann factor on the red wing to approximate quantum effects (Stehlé 1994). Self-broadening of the lower Balmer lines follows Barklem et al. (2000b), resonance broadening follows Ali & Griem (1966), and the Lyman α wing includes the quasi-molecular H_2 and H_2^+ satellites of Allard et al. (1994). Two options are provided for the hydrogen bound-free cross-section implementation: the tabulation used in Marcs and the approximation from Atlas, with both relying on the same physical foundation (Karzas & Latter 1961). The two implementations agree to better than 1%.

The package returns, at every depth and wavelength, the total (line plus continuum) and the pure continuum hydrogen opacity; their difference is the bound-bound contribution, which is added to the line opacity, while the continuum part is added to the continuum. Because the bound-free part is computed here, it is switched off in the general continuum routine (Sect. 3.3) whenever a hydrogen line list is present.

3.5. Line absorption

For a spectral line of a species with number density n_i and partition function U_i, lower-level energy $E_{\rm lo}$, and oscillator strength gf, the monochromatic line-absorption coefficient per gram at the line centre λ_0 is

$$\alpha_0 = \frac{\sqrt{\pi}\,e^2}{m_e c}\,\frac{gf}{U_i}\,n_i\,\mathrm{e}^{-E_{\rm lo}/kT}\,(1-\mathrm{e}^{-hc/\lambda_0 kT})\,\frac{1}{\rho\,\Delta\nu_{\rm D}}, \tag{1}$$

where the Doppler width combines thermal and microturbulent motions,

$$\frac{\Delta\nu_{\rm D}}{\nu_0} = \frac{\Delta\lambda_{\rm D}}{\lambda_0} = \frac{1}{c}\sqrt{\frac{2kT}{m_i} + v_{\rm mic}^2}. \tag{2}$$

The line-absorption coefficient at a wavelength offset $\Delta\lambda$ from the line centre is $\alpha_\ell = \alpha_0\,H(a,v)$, where H is the Voigt function, $v = \Delta\lambda/\Delta\lambda_{\rm D}$, and the damping parameter is

$$a = \frac{\Gamma_{\rm tot}\,\lambda_0^2}{4\pi c\,\Delta\lambda_{\rm D}}, \qquad \Gamma_{\rm tot} = \Gamma_{\rm rad} + \Gamma_{\rm Stark} + \Gamma_{\rm vdW}. \tag{3}$$

The quantities in Eq. (1) that depend only on the line, and those that depend only on depth, are precomputed separately and combined in the innermost loop so that adding a line at a new depth costs only a few floating-point operations plus one Voigt evaluation.

The three damping constants are obtained as follows. The radiative width $\Gamma_{\rm rad}$ is taken from the line list when available, otherwise from the classical value (e.g. Gray 2005), $\Gamma_{\rm rad} = 2.2233\times10^{15}/\lambda_0^2$, with λ_0 in Å. The quadratic Stark width $\Gamma_{\rm Stark}$ is taken from the line list or computed from the classical approximation scaled with electron density, distinguishing neutral and ionic species through the effective principal quantum number of the upper level (Cowley 1971); for the tabulated He I lines it is replaced by the values described in Sect. 3.5.2. The van der Waals width $\Gamma_{\rm vdW}$ is taken from the line list, or from the Unsöld approximation (Unsöld 1955) with an optional global scaling factor, or – when the line list provides the packed cross-section and velocity-exponent parameters – from the Anstee, Barklem, and O'Mara (ABO) theory of Anstee & O'Mara (1995) and Barklem et al. (2000a). Perturbations by neutral hydrogen and helium are included by default, and by molecular hydrogen optionally. For molecular lines the Stark and van der Waals broadening is switched off by default and can be enabled with a dedicated flag.

3.5.1. Fast Voigt and Faraday-Voigt functions

Because a single synthesis may require 10^{10} or more profile evaluations, the Voigt function $H(a,v)$ and its associated anomalous dispersion (Faraday-Voigt) function $L(a,v)$, needed for polarised radiative transfer, are evaluated from precomputed two-dimensional lookup tables with bilinear interpolation. The tables are node-spaced uniformly in v and in $\sqrt{a}$, which concentrates resolution where the profile varies fastest. Outside the tabulated domain, the functions are computed from the asymptotic continued-fraction expansion of the Faddeeva function. The tables themselves are built at start-up from a reference implementation based on the trapezoidal expansion of Chiarella & Reichel (1968), accurate to $\sim10^{-12}$. A per-line caching scheme stores the a-dependent interpolation state so that, when the Zeeman components of a line are summed at a given depth, only the v-dependent part of the lookup is repeated.

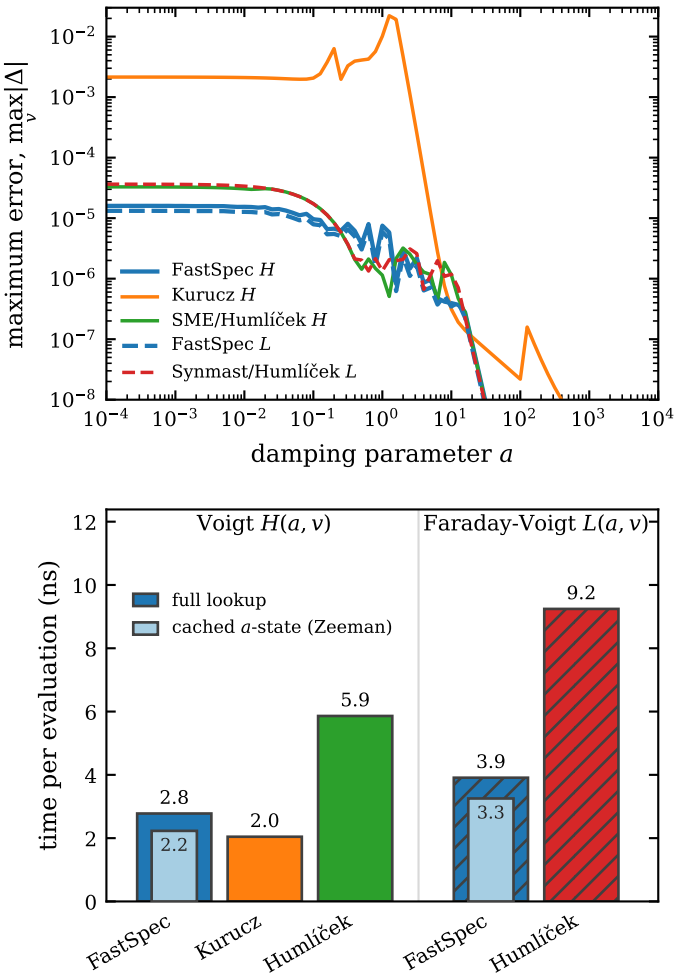


Fig. 3: Accuracy and speed of the FastSpec table-based profile functions compared with other widely used algorithms. *Top:* Maximum over v of the absolute error of the Voigt profile $H(a,v)$ (solid lines) and the Faraday-Voigt profile $L(a,v)$ (dashed lines) as a function of the damping parameter a. *Bottom:* Corresponding time per evaluation on a single core. The light-blue bars over the FastSpec entries are the cached-a-state variants used when summing over the Zeeman components of a line at fixed depth. The bar heights are the mean time per call, printed above each bar; the left group refers to $H(a,v)$ and the right hatched group to $L(a,v)$. The Atlas/Synthe tabulation has no Faraday-Voigt counterpart and therefore appears only in the left group.

Figure 3 compares the accuracy and speed of these table-based functions with several algorithms in common use: the tabulated Voigt function of Atlas/Synthe (Kurucz 2005), also used by Spectrum (Gray & Corbally 1994) and Synspec (Hubeny & Lanz 2017), the SME implementation (Piskunov & Valenti 2017) of the Humlíček (1982) w_4 rational approximation, and the Humlíček-based Faraday-Voigt routine of Synmast (Kochukhov et al. 2010). Over the full range of damping parameters, the FastSpec tables reproduce the exact profile functions to better than 2×10^{-5} – comparable to or better than the Humlíček approximations, and considerably more accurate than the Atlas/Synthe tabulation – while evaluating more than twice as fast as the former. Only the less accurate Atlas/Synthe tabulation is marginally faster. The cached-a-state variant used in the Zeeman summation is faster still. We also examined the analytic Voigt approximation used in the Korg code

(Wheeler et al. 2023); its maximum error closely follows that of the Atlas/Synthe tabulation – both peak at a few $\times 10^{-2}$ near $a \simeq 1.3$ – while its evaluation is slightly slower, so it is not shown separately in Fig. 3.

3.5.2. Helium Stark broadening

The strong He I lines, which dominate the spectra of B-type stars, are broadened predominantly by the linear and quadratic Stark effect and are poorly described by the generic quadratic-Stark prescription above. FastSpec therefore uses tabulated Stark widths and shifts for the individual He I lines, taken from the calculations of Dimitrijević & Sahal-Brechot (1984, 1990). This treatment closely follows that in SME and related codes (Valenti & Piskunov 1996; Kochukhov et al. 2010) and is reimplemented here for a larger set of He I lines. A line is matched to a table entry by its wavelength and lower-level energy, and the depth-dependent width and line-centre shift are obtained by interpolating the tabulated temperature dependence and scaling by the electron, proton, and He II number densities from the equation of state. The shift displaces the profile centre asymmetrically, which is included both in the scalar and in the polarised opacity. These per-depth widths and shifts are precomputed once during initialisation. The same table-driven framework can be extended to any other species for which detailed Stark calculations are available, for example Si I, Cr I, and Cr II (Dimitrijević et al. 2003, 2005, 2007).

3.6. The Zeeman effect and polarised line opacity

In the presence of a magnetic field of strength B, each line splits into π and σ Zeeman components. FastSpec reads the full splitting pattern – the wavelength shifts and relative strengths of the π, red-σ, and blue-σ components – from the magnetic line list (Sect. 5) so that arbitrary Zeeman patterns, and partial Paschen-Back patterns where provided, are handled exactly. The wavelength shift of a component scales as

$$\Delta\lambda_B = 4.6686 \times 10^{-13}\, g\, \lambda_0^2 B, \tag{4}$$

with λ_0 in Å, B in gauss, and the unscaled shift g given in the input line list.

From the split components, the code builds the seven independent elements of the 4×4 absorption-dispersion matrix. The three absorption profiles η_p, η_b, η_r (for the π, blue-σ, and red-σ groups) are summed over the Zeeman splitting patterns, with each Voigt function scaled by the respective component strength and the line-centre opacity coefficient α_0 of Eq. (1) and then accumulated over all contributing lines. The three anomalous dispersion profiles ρ_p, ρ_b, ρ_r are the corresponding sums of Faraday-Voigt functions. In terms of the angle γ between the field and the line of sight and the field azimuth χ in the plane of the sky, the elements of the propagation matrix are given by the expressions (e.g. Landi Degl'Innocenti & Landolfi 2004)

$$\begin{aligned}
\eta_I &= \tfrac{1}{2}\Big[\eta_p \sin^2\gamma + \tfrac{1}{2}(\eta_b + \eta_r)(1 + \cos^2\gamma)\Big],\\
\eta_Q &= \tfrac{1}{2}\Big[\eta_p - \tfrac{1}{2}(\eta_b + \eta_r)\Big] \sin^2\gamma \cos 2\chi,\\
\eta_U &= \tfrac{1}{2}\Big[\eta_p - \tfrac{1}{2}(\eta_b + \eta_r)\Big] \sin^2\gamma \sin 2\chi,\\
\eta_V &= \tfrac{1}{2}(\eta_r - \eta_b)\cos\gamma
\end{aligned} \tag{5}$$

and analogously for the dispersion terms $\rho_{Q,U,V}$, with η replaced by ρ. These profiles are accumulated on the wavelength grid in the same adaptive, memory-limited manner as the scalar line opacity, with the window width enlarged to cover the maximum Zeeman shift. The non-magnetic lines and the continuum contribute only to the diagonal absorption η_I.

3.7. Radiative transfer

In the scalar (non-magnetic) case, the emergent specific intensity at limb angle μ and wavelength λ is obtained by integrating

$$\mu \frac{\mathrm{d}I_\lambda}{\mathrm{d}\tau_\lambda} = I_\lambda - S_\lambda \tag{6}$$

along the ray, where the monochromatic optical depth τ_λ is built from the total (line plus continuum) absorption on the column-mass scale and the source function S_λ is the Planck function. The integration uses the DELO-Bézier formal solver developed by de la Cruz Rodríguez & Piskunov (2013). In their implementation the optical-depth increments are estimated with a monotonic Bézier quadrature, and the source function is interpolated with Bézier polynomials whose control points are set by monotonicity-preserving derivatives, which suppresses the spurious oscillations of higher-order schemes. Short-characteristic coefficients are evaluated from stable analytic expansions at small optical-depth steps.

In the magnetic case, the scalar equation is replaced by the vector transfer equation for the Stokes vector $\mathbf{I} = (I, Q, U, V)^{\mathsf{T}}$,

$$\mu \frac{\mathrm{d}\mathbf{I}}{\mathrm{d}\tau} = \mathbf{K}\,(\mathbf{I} - \mathbf{S}), \tag{7}$$

where $\mathbf{S} = (S, 0, 0, 0)^{\mathsf{T}}$ and $\mathbf{K}$ is the propagation matrix built from the η and ρ elements described in Sect. 3.6,

$$\mathbf{K} = \frac{1}{\eta_I}\begin{pmatrix} \eta_I & \eta_Q & \eta_U & \eta_V \\ \eta_Q & \eta_I & \rho_V & -\rho_U \\ \eta_U & -\rho_V & \eta_I & \rho_Q \\ \eta_V & \rho_U & -\rho_Q & \eta_I \end{pmatrix}. \tag{8}$$

Here η_I is the total of the magnetic-line diagonal absorption, the non-magnetic line opacity, and the continuum, and $\eta_{Q,U,V}$ and $\rho_{Q,U,V}$ are the polarised-absorption and anomalous dispersion coefficients. The common factor $1/\eta_I$ multiplies every element of the propagation matrix so that its diagonal reduces to unity. Equation (7) is then integrated with the vector version of the DELO-Bézier solver (de la Cruz Rodríguez & Piskunov 2013), which requires the inversion of a 4×4 matrix per depth step, performed by Gaussian elimination with partial pivoting. The two solvers share the same optical-depth and derivative routines. Both are validated in Sect. 6.

For the magnetic geometry, the field can be specified either directly in the observer's frame, as modulus, inclination to the line of sight, and azimuth – appropriate for a uniform field seen at a fixed aspect – or in the stellar surface frame as the radial, meridional, and azimuthal components, in which case FastSpec projects the field onto the line of sight separately at each limb angle before accumulating line opacities. The observer-frame specification also accepts an individual inclination and azimuth for each limb angle so that an arbitrary set of field orientations computed externally can be supplied directly.

The solvers described above assume a plane-parallel geometry. For spherically extended atmospheres, FastSpec additionally implements the spherical formal solution of SME (Piskunov & Valenti 2017), applying the same DELO-Bézier short-characteristic integration along the parallel rays that sample the curved atmosphere, in both the scalar and the polarised case.

3.8. Disk integration and convolution

Flux spectra are obtained by integrating the emergent intensity over the stellar disk. The disk is divided into concentric annuli centred on the μ points; by default these are spaced so as to sample equal projected areas. Each annulus is weighted by its fractional projected area. Within each annulus, rigid rotation is treated by convolving the local intensity with the analytic distribution of line-of-sight velocities appropriate to an annulus of finite width, which avoids the artefacts of a pixel-sampled rotation kernel (Valenti et al. 1995). Radial-tangential macroturbulence (Gray 2005) is applied as a convolution with a μ-dependent kernel that combines the radial and tangential Gaussian components in equal proportion. A final convolution with a Gaussian approximates the effect of an instrumental profile; its width may be given directly or as a resolving power. In the magnetic case, the same broadening is applied to all four Stokes parameters. The disk integration, like the opacity and radiative-transfer stages, is OpenMP-parallel.

4. Performance and algorithmic design

The speed of FastSpec is derived from four main ingredients that we describe in this section. This description is followed by benchmark measurements.

4.1. Adaptive line windowing

The cost of a naive synthesis scales as the product of the number of lines and the number of wavelength points, because every line is in principle evaluated at every wavelength. In reality a line contributes measurably only over a narrow interval around its centre. FastSpec determines this interval for every line before accumulating its opacity. The central line-to-continuum opacity ratio is first computed; if it lies below a user-controlled threshold ε (typically 10^{-4}), the line is discarded altogether. For a retained line, the window half-width is obtained in closed form by inverting the Voigt profile at the threshold, evaluated at each atmospheric depth and maximised over depth, giving a pair of grid indices that bracket the line without any profile evaluation or iterative search. In the hydrogen and magnetic paths, where the profile is not a single Voigt function, the window is instead found numerically, by growing a trial offset geometrically until the ratio drops below ε and refining the boundaries by bisection. Because the window is set individually for each line, its width varies enormously – from a few tens of mÅ for the weakest retained lines to tens of Å for the strongest lines – as shown in the top panel of Fig. 4. Each line is shown as a single bar, spanning its accumulation interval and placed at its width, so the vertical spread at any wavelength shows how differently the windows of neighbouring lines are sized. The remaining panels illustrate the effect of the threshold on a solar synthesis with a full ('extract all') VALD atomic line list. At the default $\varepsilon = 10^{-4}$, only about a third of the lines are retained (Fig. 5), yet the emergent flux spectrum is indistinguishable from one computed with a threshold two orders of magnitude smaller, whereas a coarse $\varepsilon = 10^{-1}$ visibly suppresses the weaker blends. Weak-line rejection thus removes the majority of a complete line list at essentially no cost in accuracy.

4.2. Memory-limited block processing

Storing the per-line, per-depth central opacities and profile parameters for a very large list would exceed available memory.

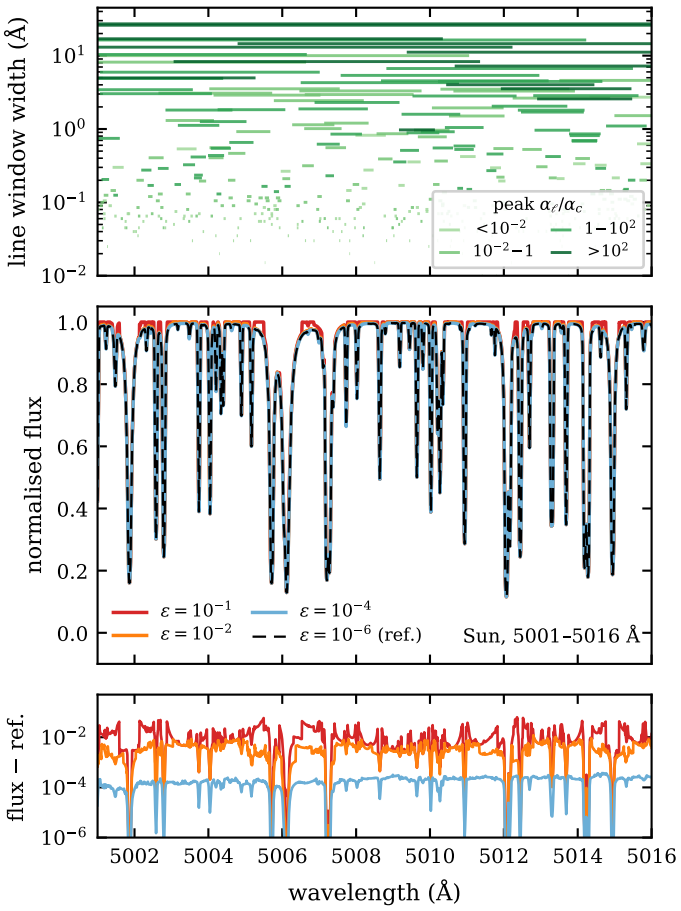


Fig. 4: Adaptive line windowing and the line-to-continuum opacity threshold ε illustrated on a segment of a solar synthetic spectrum (VALD 'extract all' atomic list, Marcs solar model, no macroscopic broadening applied). *Top:* Individual accumulation window of every line retained in the 5001–5016 Å segment, each drawn as a horizontal bar spanning $[\lambda_{\rm left}, \lambda_{\rm right}]$, positioned vertically at its width and coloured by the line's peak line-to-continuum opacity ratio $\max(\alpha_\ell/\alpha_c)$. *Middle:* Corresponding normalised flux spectrum computed with $\varepsilon = 10^{-1}$, 10^{-2} and the default 10^{-4}, with a converged reference ($\varepsilon = 10^{-6}$, dashed), which nearly coincides with the $\varepsilon = 10^{-4}$ curve. *Bottom:* Difference from the reference, with the same ε colours as in the middle panel.

FastSpec therefore caps the memory devoted to these arrays at a user-controlled limit and, when a list is larger than this cap allows, processes it in contiguous blocks of lines. The wavelength range is partitioned into segments such that all lines contributing to a segment fit within one block; the code then loads one block of per-line data at a time, accumulates its contribution over the corresponding wavelength segment, and moves on. Lines whose opacity windows are exceptionally wide are excluded from this segmentation and accumulated individually over their windows, preventing a single broad line from inflating a segment beyond the block size and from slowing the per-wavelength accumulation. In this way lists of essentially unlimited length are handled within a fixed memory footprint, at the cost of recomputing the per-line data for lines that straddle segment boundaries. The maximum block size is derived from the requested RAM limit and the dimensions of the problem. Figure 6 shows the result of computations with the ExoMol water list on a 128 GB machine.

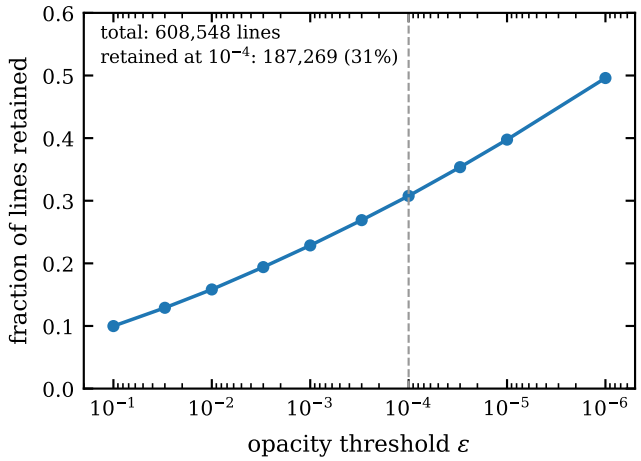


Fig. 5: Fraction of the full ('extract all') VALD atomic line list retained by the weak-line rejection as a function of the opacity threshold ε for the solar synthesis of Fig. 4. The dashed line marks the default $\varepsilon = 10^{-4}$.

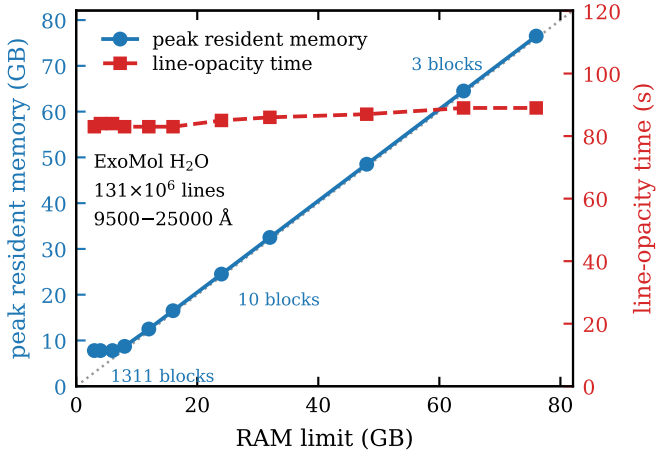


Fig. 6: Memory-limited block processing illustrated on the ExoMol H_2O line list (131×10^6 lines in the 9500–25 000 Å range, Marcs $T_{\rm eff} = 3500$ K, $\log g = 5.0$ model) on a 128 GB machine. The figure shows peak resident memory (solid, left axis) and line-opacity computation time (dashed, right axis) as a function of the user-defined RAM cap.

The peak memory closely follows the requested cap over more than a decade in memory, down to an irreducible floor (~8 GB for this large H_2O line list) set by the permanently resident line list and the spectral and index arrays. Over this whole range, the run time is essentially constant: splitting the list into as few as three or as many as ~1300 blocks changes the line-opacity time by only a few percent, because the lines recomputed at segment boundaries are a negligible fraction of the total. Block processing therefore renders the memory footprint an essentially free parameter; indeed for this list the single-block working set (~170 GB) exceeds even a 128 GB machine, so the scheme is not merely convenient but necessary.

4.3. Parallelisation and profile lookups

All heavy loops – the continuum opacity over wavelength, the per-line setup and windowing over lines, the opacity accumulation over wavelength, the radiative transfer over wavelength, and the convolutions – are parallelised with OpenMP. The precomputation of the depth-only and line-only parts of the opacity kernel, together with the shared read-only Voigt tables, ensures that the parallel regions carry no synchronisation overhead beyond the loop scheduling. Combined with the fast table-based profiles (Sect. 3.5.1), this yields near-linear scaling over the first ~ 10 cores. At higher thread counts, the opacity accumulation becomes limited by memory bandwidth rather than by synchronisation, as quantified in Sect. 4.4.

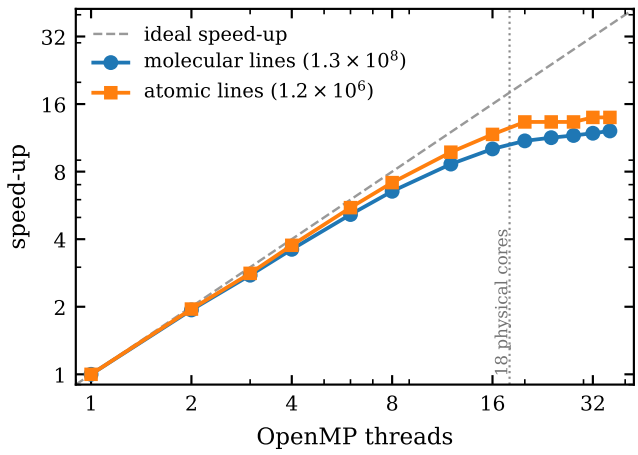


Fig. 7: Strong-scaling speed-up $S(N) = T(1)/T(N)$ of the line-opacity computation versus the number of OpenMP threads for two very different line lists: the ExoMol molecular water list (1.3×10^8 lines, cool-dwarf Marcs model) and the solar atomic list (1.2×10^6 lines). The dashed line is the ideal linear speed-up, and the dotted vertical line marks the 18 physical cores of the Intel Xeon W iMac Pro workstation used for this test (threads beyond it are logical).

Figure 7 characterises the parallel scaling of FastSpec for the line-opacity computation, which dominates a large synthesis, on a workstation with 18 physical cores (36 logical threads with simultaneous multithreading). For the ExoMol water list (131 million lines, computed for a cool-dwarf Marcs model) over the 9500–25 000 Å range, the scaling is near-linear through the first ten threads (90% parallel efficiency on four threads, 82% on eight) and then rolls off, reaching a speed-up of 12× on 36 threads. A very different workload – the solar atomic list, with full pressure broadening and the hydrogen lines treated by the occupation-probability package but some two orders of magnitude fewer lines – follows a nearly identical curve, scaling marginally better and reaching 14×. The near-coincidence of the two curves places the origin of the roll-off in the hardware rather than in the algorithm or the line list: the opacity accumulation has a low arithmetic intensity, streaming the precomputed per-line, per-depth profile coefficients from memory for only a cached Voigt-table lookup and a multiply-add per wavelength point, so once more than about ten cores are active, they saturate the shared memory bandwidth, and further threads increasingly stall waiting for data rather than computing. The final doubling from 18 to 36 threads adds only logical (hyperthreaded) cores, which share physical execution units and memory ports and barely improve a bandwidth-bound kernel. The atomic list carries more arithmetic per line and is correspondingly less bandwidth-bound, which is why it attains the slightly higher speed-up.

4.4. Benchmarks

Table 2 lists representative wall-clock times for several characteristic spectrum synthesis problems, measured on a 14-core MacBook M4 Pro laptop using all cores. For each run we separate the total wall-clock time $T_{\rm all}$ into the atomic and molecular line-opacity accumulation $T_{\rm opac}$ (the dominant cost for large lists, excluding the hydrogen-line package) and the radiative transfer $T_{\rm RT}$; the remainder covers the equation of state, continuum opacity, hydrogen lines, and I/O. In all cases, spectra were calculated at seven limb angles using a $0.3\,{\rm km\,s^{-1}}$ velocity step, corresponding to a wavelength sampling of $\lambda/\Delta\lambda \approx 10^6$. Atomic syntheses over the 3000–25 000 Å range using the full ('extract all') VALD data, whether for the Sun or for hot stars, complete in about ten seconds. Their corresponding radiative transfer – unpolarised, seven μ angles, over two million wavelength points – takes less than two seconds. Adding the full set of solar molecular lists raises the opacity time to 7 s and the total to 13 s but leaves the transfer unchanged. The two largest cases are dominated by the ExoMol water list: even with $\sim 10^8$ contributing lines, the cool-dwarf synthesis completes in about forty seconds for $\approx 10^6$ wavelengths, and the full polarised $IQUV$ synthesis in under fifty seconds for the same wavelength grid. The polarised transfer step itself stays cheap (3.0 s), scaling with the number of wavelength points rather than the line count, so most of the additional cost is the Zeeman treatment of the magnetic-line opacity.

These tests demonstrate that combined atomic and molecular line lists of hundreds of millions of lines can be synthesised over the full optical-infrared range in under a minute on a single multi-core machine. A direct comparison of the speed and memory usage of FastSpec with the modern, multi-threaded Korg code is presented in Appendix A.

5. Input data and line lists

FastSpec reads three kinds of line list. A single synthesis can combine atomic, molecular, and magnetic line data.

Atomic lists are supplied as HDF5 files containing, for each transition, the wavelength, element and ionisation stage, lower-level energy, oscillator strength, and the radiative, Stark, and van der Waals damping constants. These files can be produced by converting VALD extractions. For the calculations in this paper, we used a single 'extract all' file packing the entire VALD atomic data in a 36 MB HDF5 file. A list containing only hydrogen lines is recognised automatically and routed to the occupation-probability package (Sect. 3.4); it must be supplied first. The reader uses a two-pass strategy, first taking an inventory of the species present to build the equation-of-state species list, then reading the transitions that fall within the requested wavelength range.

Molecular lists are single-species HDF5 files, such as conversions of ExoMol (Tennyson et al. 2024), containing the wavelength, lower-level energy, oscillator strength, upper-level statistical weight, and radiative width. The isotopologue is identified from a numerical species code, and the oscillator strengths are scaled by the appropriate isotopic fraction so that a molecule can be synthesised from a set of per-isotopologue lists.

Magnetic lists are ASCII files in the format used by the Synthmag and Synmast codes (Piskunov 1999; Kochukhov et al. 2010). In addition to the usual atomic or molecular line data in VALD 'extract stellar' format, they contain, for each line, the explicit Zeeman splitting pattern: the number, shifts, and strengths of the π and σ components. This allows anomalous Zeeman and, where tabulated, partial Paschen-Back patterns to be modelled exactly. As an alternative to this precomputed form, the magnetic line data can be supplied in an HDF5 line list that stores, in addition to the usual atomic quantities, the lower- and upper-level angular momenta J and Landé factors g. FastSpec then computes the anomalous Zeeman splitting pattern of each line internally as the list is read. Lines lacking Landé factor information are treated as non-magnetic. A magnetic list triggers the polarised transfer only when a magnetic field is also specified; otherwise it is synthesised as an ordinary line list.

FastSpec is driven by an input file that specifies the wavelength range and grid type, the line lists and model atmosphere files to use, the desired outputs, disk integration and convolution parameters, elemental abundances, and optional computational settings. The full set of input keywords is documented in Appendix B, and the binary output formats in Appendix C.

6. Tests and applications

We illustrate the correctness and versatility of FastSpec with four applications. Together, these tests span most of the parameter space for which the code is intended.

The solar spectrum. Figure 8 compares synthetic solar flux spectra with the observed, telluric-corrected solar flux atlas of Wallace et al. (2011) in two diagnostic windows: the Mg I b triplet and Hα. The synthesis uses a Marcs solar model ($T_{\rm eff}$ = 5777 K, $\log g$ = 4.44), a microturbulence of $0.8\,{\rm km\,s^{-1}}$, the built-in solar abundances of Asplund et al. (2021), and the Gaia-ESO survey line list (Heiter et al. 2021) – its atomic lines together with the C_2 and MgH bands that blend into the b-triplet region. The emergent flux is disk-integrated over seven limb angles and broadened by a radial-tangential macroturbulence of $3.2\,{\rm km\,s^{-1}}$ and a projected rotational velocity of $1.9\,{\rm km\,s^{-1}}$. To provide an independent check, the same model atmosphere, abundances, line list, and broadening are also passed to PySME (Piskunov & Valenti 2017; Jian et al. 2026).[6]

The two codes agree closely across both windows, matching each other line by line and in the cores and wings of the strong features, including Hα. Both reproduce the observed metal lines and the Mg I b triplet well. Two differences stand out. The synthetic Hα core is shallower than observed because it forms in the uppermost photosphere and chromosphere, which a static one-dimensional LTE model does not describe; the mismatch therefore reflects the model atmosphere rather than the spectral synthesis. Near the C_2 Swan bandhead at 5165 Å, the two syntheses diverge slightly, PySME producing marginally stronger C_2 absorption. The difference amounts to about 2% of the continuum flux across the band, reaching 7% at the head near 5165 Å, while redward of it the two codes agree to better than 0.6%. As both read identical line data, this traces back to a slightly different C_2 dissociation energy adopted for the molecular equilibrium calculations in the two codes. The test simultaneously exercises the continuum opacity, the equation of state, and the atomic and molecular line opacity and broadening, and cross-validates them against an independent implementation.

Molecular bands in an M dwarf. Figure 9 shows the complete near-infrared spectrum (0.8–2.5 μm) of a cool ($T_{\rm eff}$ = 3500 K, $\log g$ = 5.0) dwarf, synthesised from the full input list of atomic and molecular lines: the atomic list together with the CH, CN, CO, CrH, FeH, HF, NaH, OH, TiO, and H_2O molecular lists.

[6] https://github.com/SpectroscopyMadeEasy/PySME

Table 2: Representative wall-clock times.

Model / line list	Range (Å)	N_λ	$N_{\rm lines}$	$N_{\rm mag}$	$T_{\rm all}$ (s)	$T_{\rm opac}$ (s)	$T_{\rm RT}$ (s)
Sun / atomic	3000–25 000	2.12×10^6	3.53×10^5	…	6.9	1.0	1.3
Sun / atomic + molecular	3000–25 000	2.12×10^6	1.12×10^6	…	13.0	6.9	1.3
t20000g4.0 / atomic	3000–25 000	2.12×10^6	3.86×10^4	…	10.7	0.5	1.5
t9000g4.0 / atomic	3000–25 000	2.12×10^6	2.04×10^5	…	11.0	0.8	1.6
t9000g4.0 / Stokes $IQUV$	3000–25 000	2.12×10^6	2.04×10^5	2.04×10^5	24.7	8.0	8.2
t3500g5.0 / atomic + molecular	9500–25 000	9.67×10^5	9.67×10^7	…	43.0	40.1	0.6
t3500g5.0 / Stokes $IQUV$	9500–25 000	9.67×10^5	9.67×10^7	4.74×10^4	48.5	43.2	3.0

Notes. These test calculations were executed on a MacBook M4 Pro (14 cores: ten performance and four efficiency, 48 GB RAM) using all cores. N_λ is the number of wavelength points, $N_{\rm lines}$ is the number of lines retained above the opacity threshold, and $N_{\rm mag}$ is the number treated with the full polarised transfer. $T_{\rm all}$ is the total wall-clock time of the run, $T_{\rm opac}$ is the atomic and molecular line-opacity accumulation (the total line-opacity time excluding the hydrogen-line package), and $T_{\rm RT}$ is the (polarised, for the Stokes $IQUV$ cases) radiative transfer.

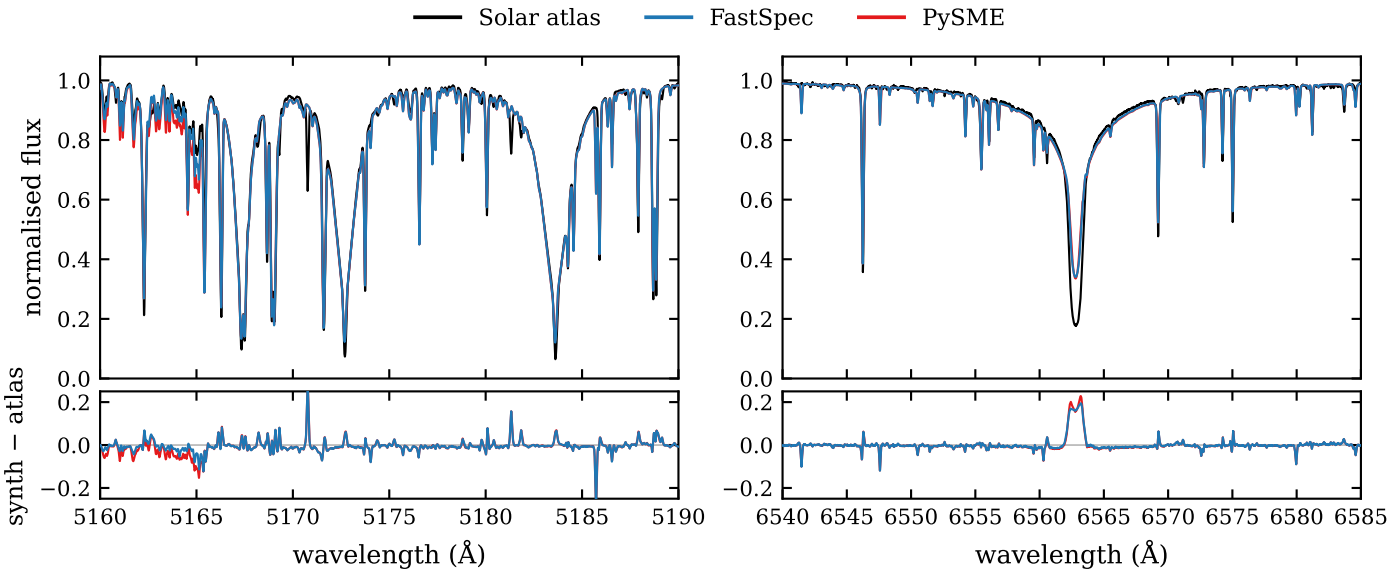


Fig. 8: Solar flux spectrum in the Mg I *b* triplet region (*left*) and around Hα (*right*). *Upper panels:* FastSpec (blue) and an independent PySME synthesis (red) computed from the same 1D Marcs solar model, solar abundances (Asplund et al. 2021), and Gaia-ESO line list (Heiter et al. 2021) are shown together with the solar flux atlas (Wallace et al. 2011) (black). *Lower panels:* Corresponding synthetic-minus-observed residuals.

About 10^8 lines are retained above the opacity threshold, overwhelmingly dominated by the ExoMol H_2O list. To identify the individual contributors, the spectrum is also recomputed with only the atomic lines (middle panel) and with each of the main molecular species in turn. The far red is shaped by the TiO bands (here from five titanium isotopologues, ^{46}TiO to ^{50}TiO), while H_2O dominates the near-infrared, together with the CO first-overtone bandhead at 2.3 μm, the OH bands longward of 1.5 μm, and FeH contributions around 1.0 and 1.6 μm. This test demonstrates the handling of very large molecular line lists and the isotopic scaling of oscillator strengths, and stresses the memory-limited block processing of Sect. 4.2.

Polarised spectra of a magnetic M dwarf. Figure 10 illustrates the polarised mode of FastSpec on the same cool dwarf as Fig. 9 ($T_{\rm eff}$ = 3500 K, $\log g$ = 5.0), now with a uniform 3 kG field inclined 45° to the line of sight with zero azimuth. The whole 0.95–2.5 μm interval is synthesised in a single run with atomic and FeH lines (Shulyak et al. 2010) treated as polarised and other molecular species as non-magnetic. Six narrow windows are extracted to show the disk-centre Stokes I, Q, U, and V profiles of representative features: the atomic lines of Ti I, K I, Ca I, and Na I, and the FeH lines within the FeH Wing-Ford band near 0.99 μm.

The comparison of the magnetic and field-free Stokes I profiles isolates the Zeeman effect: the atomic lines broaden and, for the larger effective Landé factors, split into resolved components, with strong antisymmetric V and symmetric Q signatures of up to tens of percent of the continuum. Because the field azimuth is zero, there is no direct contribution to Stokes U; the small U signal arises entirely from Faraday rotation of the linear polarisation during transfer through the atmosphere. The FeH bandhead is a contrasting molecular case: its lines differ widely in magnetic sensitivity, so within a single segment some are strongly broadened by the field while others are essentially unaffected, and the net circular polarisation stays weak and complex as the many overlapping components of opposite polarity cancel.

Hydrogen lines and the Balmer jump. Figure 11 examines the Balmer-jump region of an A-type model ($T_{\rm eff}$ = 9000 K, $\log g$ = 4.0). We compare FastSpec with an independent synthesis produced by PySME for the same model atmosphere, abundances, microturbulence, and atomic line list. The spectra are shown for

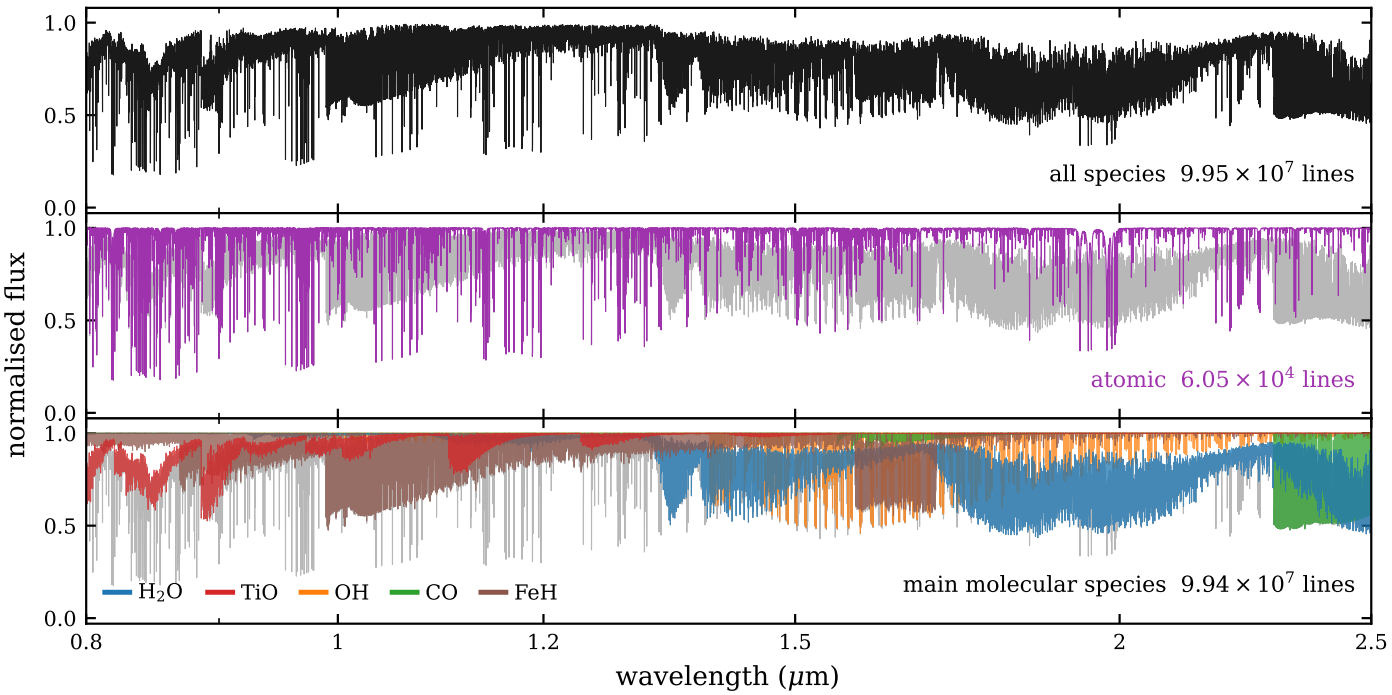


Fig. 9: Synthetic near-infrared spectrum of a cool M-dwarf model ($T_{\rm eff}$ = 3500 K, log g = 5.0). *Top:* Complete spectrum including every opacity source of the input list. *Middle:* Atomic-line contribution alone (magenta) overplotted on the full spectrum (grey). *Bottom:* Main individual molecular contributions – H_2O, TiO, OH, CO, and FeH – each computed with only that species active and overplotted in colour.

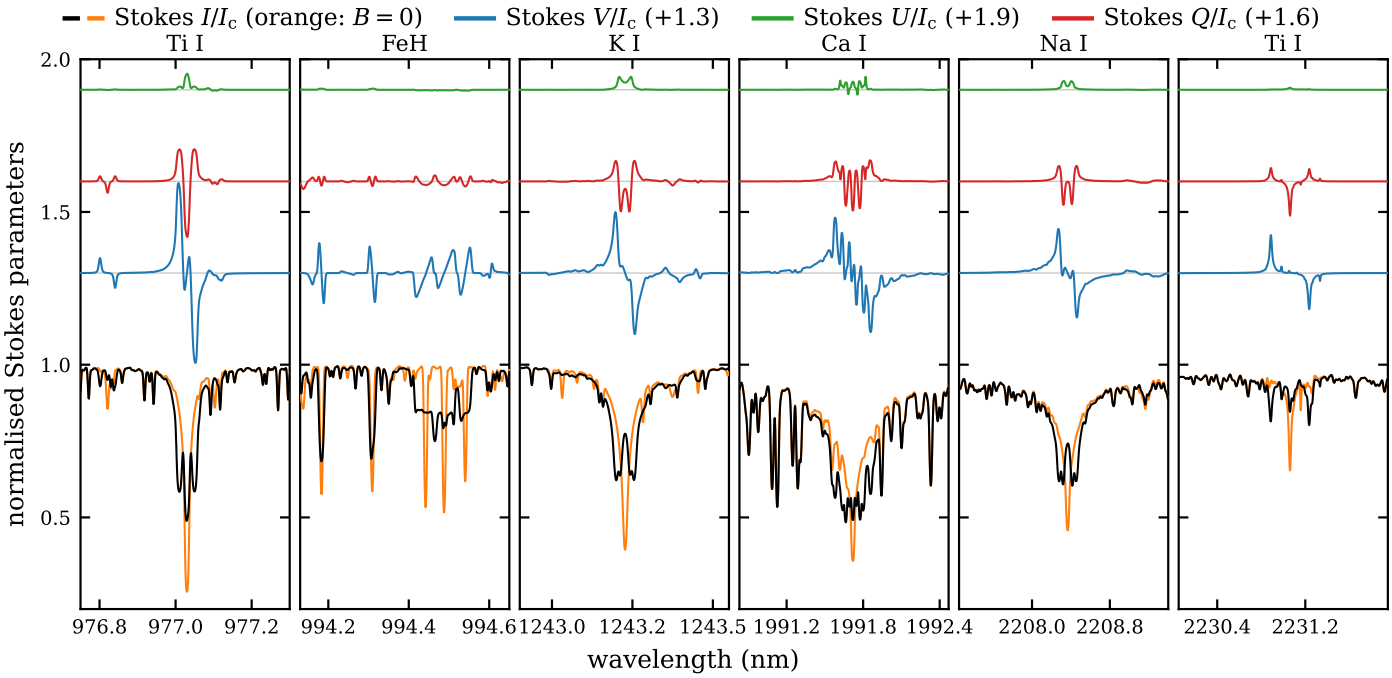


Fig. 10: Polarised disk-centre near-infrared spectrum of a magnetic M-dwarf model ($T_{\rm eff}$ = 3500 K, log g = 5.0) with a uniform 3 kG field inclined by 45° to the line of sight (zero azimuth). The six windows are each centred on a selected feature and labelled by the dominant species. In every panel the magnetic Stokes $I/I_{\rm c}$ (black) is drawn over the non-magnetic (B = 0) profile (orange); the polarised profiles $Q/I_{\rm c}$, $U/I_{\rm c}$, and $V/I_{\rm c}$ have been shifted upward by the constant given in the legend.

the 3600–4000 Å interval that brackets the jump. The two calculations differ only in their treatment of the hydrogen opacity: FastSpec synthesises the Balmer lines with its occupation-probability treatment of Barklem & Piskunov (2003) and merges them smoothly into the bound-free continuum, whereas PySME superposes the individual high Balmer lines on a sharp hydrogenic bound-free edge.

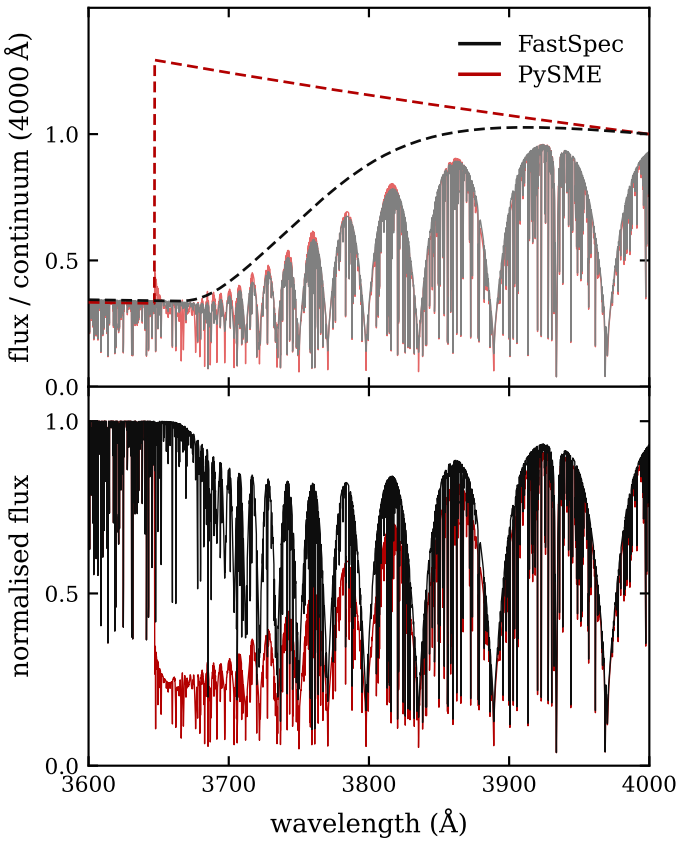


Fig. 11: Balmer-jump region of an A-type model ($T_{\rm eff}$ = 9000 K, $\log g$ = 4.0) computed with FastSpec (grey) and PySME (red) for identical inputs. *Top:* Flux (solid) and continuum (dashed) spectra of each code, divided by the respective continua at 4000 Å. *Bottom:* Continuum-normalised spectra.

Shortward of the 3646 Å series limit and beyond $\simeq$ 4000 Å, the two codes agree to better than 0.1% in continuum-normalised flux, confirming that the metal-line and lower-Balmer-line opacities are mutually consistent. They differ only over 3646–4000 Å, where PySME produces a discontinuous Balmer jump at the series limit whereas FastSpec yields a gradual one as the highest Balmer lines dissolve into the continuum.

7. Summary

We have presented FastSpec, a fast LTE stellar spectrum synthesis code that combines the physical completeness of the classical model-atmosphere codes with the computational speed demanded by large line lists and surveys, and with the polarised radiative transfer capability required to model magnetic stars. Its physical ingredients – the improved SME equation of state, Atlas-family continuum opacities, the occupation-probability hydrogen package, tabulated He I Stark broadening, and standard prescriptions for atomic and molecular line broadening – are drawn from well-tested public codes and restructured for speed. The code owes its performance to four design choices: comprehensive OpenMP parallelisation, adaptive per-line wavelength windowing with weak-line rejection, memory-capped block processing of arbitrarily large line lists, and table-based Voigt and Faraday-Voigt profiles. FastSpec synthesises intensity and flux spectra, as well as the full Stokes $IQUV$ vector when a magnetic field is present, and handles atomic, molecular, and magnetic line lists in a single run. The code supports both plane-parallel and spherical model atmospheres, with either geometry compatible with full-Stokes polarised radiative transfer. We have demonstrated the performance and range of applicability of FastSpec on the solar spectrum, on the molecular bands of an M dwarf, on the polarised spectrum of a magnetic star, and on the hydrogen lines of hot stars. The code is portable, depends only on an HDF5-enabled Fortran compiler, and is intended to serve both the modelling of individual stars and the large-scale synthesis needs of surveys and spectropolarimetric magnetic field inversions.

FastSpec is accompanied by a suite of Python utilities that convert VALD extractions and molecular line lists into its input formats and read its binary output files into arrays for analysis and plotting. The code and these utilities are made publicly available on GitHub.[7]

Acknowledgements. This work made use of the VALD database, operated at Uppsala University and the University of Montpellier, and of the equation-of-state and hydrogen-opacity packages developed by P. Barklem, J. Valenti, and N. Piskunov. We thank B. Edvardsson for sharing his collection of molecular line lists and partition functions. We are grateful to the referee, C. Sneden, for a careful reading of the manuscript and constructive comments. Support by the Swedish Research Council (grant agreement no. 2023-03667) and the Swedish National Space Agency is acknowledged.

References

Ali, A. W. & Griem, H. R. 1966, Physical Review, 144, 366
Allard, N. F., Koester, D., Feautrier, N., & Spielfiedel, A. 1994, A&AS, 108, 417
Anstee, S. D. & O'Mara, B. J. 1995, MNRAS, 276, 859
Asplund, M., Amarsi, A. M., & Grevesse, N. 2021, A&A, 653, A141
Barklem, P. S. & Collet, R. 2016, A&A, 588, A96
Barklem, P. S. & Piskunov, N. 2003, in IAU Symposium, Vol. 210, Modelling of Stellar Atmospheres, ed. N. Piskunov, W. W. Weiss, & D. F. Gray, E28
Barklem, P. S., Piskunov, N., & O'Mara, B. J. 2000a, A&AS, 142, 467
Barklem, P. S., Piskunov, N., & O'Mara, B. J. 2000b, A&A, 363, 1091
Chiarella, C. & Reichel, A. 1968, Mathematics of Computation, 22, 137
Cowley, C. R. 1971, The Observatory, 91, 139
Cristofari, P. I., Donati, J.-F., Folsom, C. P., et al. 2023, MNRAS, 522, 1342
Däppen, W., Anderson, L., & Mihalas, D. 1987, ApJ, 319, 195
de la Cruz Rodríguez, J. & Piskunov, N. 2013, ApJ, 764, 33
Deen, C. P. 2013, AJ, 146, 51
Dimitrijević, M. S. & Sahal-Brechot, S. 1984, J. Quant. Spectr. Rad. Transf., 31, 301
Dimitrijević, M. S. & Sahal-Brechot, S. 1990, A&AS, 82, 519
Dimitrijević, M. S., Ryabchikova, T., Popović, L. Č., Shulyak, D., & Tsymbal, V. 2003, A&A, 404, 1099
Dimitrijević, M. S., Ryabchikova, T., Popović, L. Č., Shulyak, D., & Khan, S. 2005, A&A, 435, 1191
Dimitrijević, M. S., Ryabchikova, T., Simić, Z., Popović, L. Č., & Dačić, M. 2007, A&A, 469, 681
Gerber, J. M., Magg, E., Plez, B., et al. 2023, A&A, 669, A43
Gray, D. F. 2005, The Observation and Analysis of Stellar Photospheres (Cambridge, UK: Cambridge University Press)
Gray, R. O. & Corbally, C. J. 1994, AJ, 107, 742
Griem, H. R. 1960, ApJ, 132, 883
Griem, H. R. 1967, ApJ, 147, 1092
Gustafsson, B., Edvardsson, B., Eriksson, K., et al. 2008, A&A, 486, 951
Heiter, U., Lind, K., Bergemann, M., et al. 2021, A&A, 645, A106
Hubeny, I., Hummer, D. G., & Lanz, T. 1994, A&A, 282, 151
Hubeny, I. & Lanz, T. 2017, arXiv e-prints, arXiv:1706.01859
Humlíček, J. 1982, J. Quant. Spectr. Rad. Transf., 27, 437
Hummer, D. G. & Mihalas, D. 1988, ApJ, 331, 794
Jian, M., Piskunov, N., Valenti, J., et al. 2026, A&A, 711, A168
Karzas, W. J. & Latter, R. 1961, ApJS, 6, 167
Kochukhov, O. 2007, in Physics of Magnetic Stars, ed. I. I. Romanyuk & D. O. Kudryavtsev, 109
Kochukhov, O. 2026, A&A, 707, A16
Kochukhov, O., Lüftinger, T., Neiner, C., Alecian, E., & MiMeS Collaboration 2014, A&A, 565, A83
Kochukhov, O., Makaganiuk, V., & Piskunov, N. 2010, A&A, 524, A5

[7] https://github.com/okochukhov/FastSpec

Kupka, F., Piskunov, N., Ryabchikova, T. A., Stempels, H. C., & Weiss, W. W. 1999, A&AS, 138, 119
Kurucz, R. L. 2005, Mem. Soc. Astron. Italiana, 8, 14
Landi Degl'Innocenti, E. & Landolfi, M. 2004, Astrophysics and Space Science Library, Vol. 307, Polarization in Spectral Lines (Dordrecht: Kluwer Academic Publishers)
Landstreet, J. D. 1988, ApJ, 326, 967
Piskunov, N. 1999, in Astrophysics and Space Science Library, Vol. 243, Polarization, ed. K. N. Nagendra & J. O. Stenflo, 515–525
Piskunov, N. & Kochukhov, O. 2002, A&A, 381, 736
Piskunov, N. & Valenti, J. A. 2017, A&A, 597, A16
Piskunov, N. E., Kupka, F., Ryabchikova, T. A., Weiss, W. W., & Jeffery, C. S. 1995, A&AS, 112, 525
Plez, B. 2012, Turbospectrum: Code for spectral synthesis, Astrophysics Source Code Library, record ascl:1205.004
Ryabchikova, T., Piskunov, N., Kurucz, R. L., et al. 2015, Phys. Scr, 90, 054005
Sauval, A. J. & Tatum, J. B. 1984, ApJS, 56, 193
Shulyak, D., Reiners, A., Wende, S., et al. 2010, A&A, 523, A37
Shulyak, D., Tsymbal, V., Ryabchikova, T., Stütz, C., & Weiss, W. W. 2004, A&A, 428, 993
Sneden, C. 1973, ApJ, 184, 839
Stehlé, C. 1994, A&AS, 104, 509
Stehlé, C. & Hutcheon, R. 1999, A&AS, 140, 93
Tennyson, J., Yurchenko, S. N., Zhang, J., et al. 2024, J. Quant. Spectr. Rad. Transf., 326, 109083
Unsöld, A. 1955, Physik der Sternatmosphären, mit besonderer Berücksichtigung der Sonne (Berlin: Springer)
Valenti, J. A., Marcy, G. W., & Basri, G. 1995, ApJ, 439, 939
Valenti, J. A. & Piskunov, N. 1996, A&AS, 118, 595
Vidal, C. R., Cooper, J., & Smith, E. W. 1973, ApJS, 25, 37
Wade, G. A., Bagnulo, S., Kochukhov, O., et al. 2001, A&A, 374, 265
Wallace, L., Hinkle, K. H., Livingston, W. C., & Davis, S. P. 2011, ApJS, 195, 6
Wheeler, A. J., Abruzzo, M. W., Casey, A. R., & Ness, M. K. 2023, AJ, 165, 68

Table A.1: Timing and peak memory of FastSpec and Korg for three benchmark syntheses.

				$T_{\rm all}$ (s)		RAM (GB)	
Test	Range (Å)	N_λ	$N_{\rm lines}$	FastSpec	Korg	FastSpec	Korg
Sun / atomic	3000–25 000	2.20×10^6	1.18×10^6	12.9	45.4	3.65	14.05
M dwarf / TiO	5000–16 000	1.10×10^6	1.20×10^7	13.1	49.1	4.67	11.40
M dwarf / H_2O	9500–25 000	1.55×10^6	1.31×10^8	66.4	649.2	16.56	16.71

Notes. Both codes were run on the same hardware using all logical cores. N_λ is the number of wavelength points, $N_{\rm lines}$ the number of lines in the input list over the range, $T_{\rm all}$ the total wall-clock time, and RAM the peak memory footprint.

Appendix A: Performance comparison with Korg

Korg[8] (Wheeler et al. 2023) is a modern open-source spectrum synthesis code that, like FastSpec, is written for speed and supports native multi-threaded execution, which makes it a natural reference point for assessing the performance of FastSpec. The two codes do not share the same physics genealogy – they differ in the equation of state, the continuum opacities, and the details of the line broadening – so we make no attempt at a point-by-point comparison of the emergent spectra. Instead we configure both codes as closely as possible for three large benchmark syntheses, matching the line lists, wavelength ranges, and sampling, using the same Marcs model atmospheres, and setting equivalent line-inclusion thresholds, and we compare the total wall-clock time and the peak resident memory. Both codes were run on the same machine (different from the one used for the timings in Table 2). Korg (version 1.2.1) was launched with `julia -t auto`, which activates all logical cores, and FastSpec likewise used all OpenMP threads. To reproduce the internal line-selection cutoff of Korg, we set the FastSpec line-to-continuum opacity threshold to $\varepsilon = 3 \times 10^{-4}$ (Sect. 4.1). Both codes read their line lists from native HDF5 files – FastSpec from its converted VALD and molecular lists and Korg from lists saved in its own HDF5 format – so that the reported times reflect the synthesis rather than the parsing of large ASCII line lists.

The first test is a solar synthesis over 3000–25 000 Å at a constant step $\Delta\lambda = 0.01$ Å (2.2×10^6 wavelength points), with the full VALD atomic line list (1.18×10^6 lines over the range). As summarised in Table A.1, FastSpec completes the synthesis in 12.9 s using 3.65 GB, against 45.4 s and 14.05 GB for Korg, a factor of 3.5 in speed and 3.8 in memory. The second test targets the molecule-dominated spectrum of a cool M dwarf ($T_{\rm eff} = 3500$ K, $\log g = 5.0$) over 5000–16 000 Å at $\Delta\lambda = 0.01$ Å (1.1×10^6 points), using the five TiO isotopologue line lists (^{46}TiO to ^{50}TiO) also employed in Sect. 6, 1.2×10^7 lines in total. Here FastSpec runs in 13.1 s using 4.67 GB, whereas Korg takes 49.1 s and 11.40 GB, corresponding to factors of 3.7 in speed and 2.4 in memory. The third test is the most computationally demanding and isolates the molecular opacity that shapes the near-infrared spectra of cool dwarfs. Using the same M-dwarf model, we synthesise the 9500–25 000 Å region at $\Delta\lambda = 0.01$ Å (1.55×10^6 points) with the ExoMol water line list, which contributes 1.31×10^8 lines over this interval. FastSpec completes the synthesis in 66.4 s using 16.56 GB (close to the user-defined 16 GB RAM limit applied in this test), against 649.2 s and 16.71 GB for Korg, a factor of 9.8 in speed at essentially the same peak memory.

Across the three tests, FastSpec is between 3.5 and 9.8 times faster than Korg, the advantage growing with the size of the line list and reaching almost an order of magnitude for the water synthesis, where line handling dominates the total cost. Its memory footprint is 2–4 times smaller in the atomic and TiO tests and comparable to that of Korg in the H_2O test, where the shared cost of storing more than 10^8 lines dominates the working set of both codes. These gains are consistent with the design choices described in Sect. 4: adaptive line windowing makes the cost scale with the number of significant line-wavelength interactions rather than with the product of list size and grid size, while memory-limited block processing keeps the working set bounded. These results are not a criticism of Korg, whose design emphasises automatic differentiability and interoperability within the Julia ecosystem and therefore adopts different implementation trade-offs; both codes are orders of magnitude faster than the legacy synthesis programs they build upon.

Appendix B: Input parameter file

The input parameter file uses a `keyword=value` syntax; comment lines beginning with `#` or `!` and inline comments are ignored. The recognised keywords are summarised in Table B.1. The mandatory keywords are `WAVE`, `SAVE`, `LINES`, and `MODEL`; the remaining keywords are optional and take the default values given in the last column.

Appendix C: Output files

The outputs are written as unformatted Fortran stream files, each beginning with the array dimensions. The line-strength list (`fs_linst.bin`) contains one block per line list, each holding the species names, ionisation stages, central wavelengths, and central line depths of the retained lines. The line- and continuum-opacity files (`fs_lineop.bin`, `fs_contop.bin`) store the wavelength grid, the opacity as a function of depth and wavelength, and the model depth structure. The intensity file (`fs_spec.bin`) stores the μ grid, the wavelength grid, and the line and continuum intensities at each angle – or, in the magnetic case, the four Stokes parameters and the continuum. The flux file (`fs_flux.bin`) stores the wavelength grid and the disk-integrated line and continuum fluxes. Python routines that read these files into arrays are included in the distribution.

[8] https://github.com/ajwheeler/Korg.jl

Table B.1: Input keywords for the FastSpec control file.

Keyword	Meaning	Default
`WAVE`	Start and end wavelength, step, and mode (0: constant in wavelength, 1: constant in velocity)	–
`SAVE`	Five flags selecting the outputs: line-strength list, line opacity, continuum opacity, intensity, flux	–
`LINES`	Block of line-list entries, one per line: species code and file name	–
`MODEL`	Model-atmosphere file name (`.krz`)	–
`MUS`	Number of μ angles, or an explicit list of μ values	7
`VMIC`	Microturbulent velocity (km s^{-1})	0.0
`CONVOL`	$v \sin i$, $v_{\rm mac}$ (km s^{-1}) and Gaussian width (km s^{-1}) or resolving power	none
`BFIELD`	Magnetic field vector (G): three numbers for `BFRAME` 0 or 1, or $2N_\mu + 1$ numbers for `BFRAME` 2	none
`BFRAME`	Field vector frame: stellar (0), observer (1), or a list of inclination/azimuth pairs (2)	0
`[M/H]`	Metallicity scaling of the default solar abundance pattern	–
`ABUNDANCE`	Individual element abundances overriding those in the model header/default solar pattern	–
`OPACITY`	Precomputed line-opacity file to add	none
`LIMITRAM`	RAM limit (GB) for the large arrays	16
`EPSOPAC`	Line-to-continuum threshold ε for line inclusion	10^{-4}
`VWFUDGE`	Scaling factor of the Unsöld van der Waals width	1.0
`H2BRD`	Include H_2 in the van der Waals broadening	no
`MOLBRD`	Apply Stark/van der Waals broadening to molecular lines	no
`BFKURUCZ`	Choice of hydrogen bound-free cross-sections (Atlas or Marcs)	Atlas