We rebuilt ATLAS12 and SYNTHE in Python so self-consistent stellar spectra no longer require legacy Fortran.

pyKurucz: A Pure Python Reimplementation of Kurucz ATLAS12 and SYNTHE for Self-Consistent Stellar Atmospheres and Spectra

Elliot M. Kim[1] **and Yuan-Sen Ting**[2,3]

[1]Cornell University, USA
[2]The Ohio State University, USA
[3]Max Planck Institute for Astronomy, Heidelberg, Germany

11 May 2026

arXiv:2603.11693v2 [astro-ph.SR] 22 May 2026



Summary

pyKurucz is a pure Python reimplementation of Kurucz's ATLAS12 and SYNTHE — the standard tools for computing self-consistent stellar atmospheres and synthetic spectra from first principles. Stellar spectra encode information about a star's temperature, gravity, and chemical composition, but extracting it requires a forward model: starting from $(T_{\rm eff}, \log g)$ and an abundance pattern, the temperature–pressure–density stratification of the atmosphere must first be solved (radiative/convective equilibrium with opacity sampling), and then the emergent spectrum must be computed by integrating radiative transfer through that atmosphere wavelength by wavelength. The Kurucz Fortran codes have filled this role for decades. pyKurucz performs the same calculation, to the same numerical precision, entirely in Python.

Changing abundances changes opacity and atmospheric structure, so the whole spectrum must respond.

The package consists of two cooperating engines. `atlas_py` is the atmosphere engine: it iteratively solves Saha/Boltzmann populations, molecular equilibrium, continuous opacity (H^-, H^+ bound-free via Kurucz's Hminus routine, H^- free-free, Rayleigh, H_2 Rayleigh, Thomson), atomic line opacity by *direct opacity sampling* from preselected Kurucz line records (LINOP/XLINOP), the Rosseland mean, the radiative transfer equation (JOSH solver, using the same pretabulated BLOCKJ/BLOCKH coefficient matrices as ATLAS12), and mixing-length convection, then applies temperature corrections (TCORR) until the layer structure converges. Direct opacity sampling — the defining feature of ATLAS12 over the more common ATLAS9 — means the atmosphere is recomputed against the *actual* abundance pattern requested, so peculiar patterns (CEMP, Ap, individual α-element offsets, r-process enhancements) are handled identically to bulk-scaled ones, with no opacity-distribution-function (ODF) retabulation required. `synthe_py` is the synthesis engine: given a converged atmosphere it evaluates continuous opacity, line opacity from the full Kurucz GFALL atomic line list with Voigt profiles (thermal + van der Waals + Stark + radiative broadening), Stark-broadened hydrogen line profiles via the HPROF4 routine, tabulated helium line profiles (BCS, Griem, and Dimitrijević lookup tables), molecular line opacity from the Schwenke TiO list, the Partridge–Schwenke H_2O list, and the GFALL diatomic pack (CN, CO, C_2, CH, OH, MgH, FeH, and others), and the radiative transfer equation on a high-resolution wavelength grid via its own JOSH solver. The end-to-end orchestrator (`pykurucz.py`) chains an optional neural warm-start (`kurucz-al`; Li et al., 2025) into `atlas_py` into `synthe_py` to go from $(T_{\rm eff}, \log g, [{\rm M/H}], [\alpha/{\rm M}]$, per-element abundances) to a synthetic spectrum.

The neural emulator supplies an initial guess; atlas_py still solves the final atmosphere.

The implementation uses NumPy, SciPy, and Numba (JIT compilation of inner loops in both engines), requiring no Fortran compiler. End-to-end, on five representative test cases spanning 2500 K cool giants to 44000 K O stars, the Python pipeline reproduces Fortran ATLAS12 + SYNTHE to sub-0.005% median deviation in normalized flux over 300–1800 nm at resolving power $R = 300{,}000$. Full documentation — installation, user guide, architecture, physics reference, and worked examples — is available at https://pykurucz.vercel.app.

Controlled tests span five representative cases from 2500 to 44000 K.

Abundance changes must reshape the atmosphere before we synthesize the spectrum.

Statement of need

Kurucz's ATLAS12 (atmospheres) and SYNTHE (synthesis), together with the atomic and molecular line lists that accompany them, are the de facto standards for stellar-spectrum forward modeling, accumulating tens of thousands of citations over several decades (Kurucz, 2005). The original Fortran, written in a pre-Fortran-77 dialect with fixed-format source, single-character variables, computed GOTOs, and EQUIVALENCE statements, is incompatible with modern `gfortran` without significant manual patching — compiling it is a rite of passage that has defeated many research groups. Following the passing of Robert L. Kurucz in 2025, the long-term maintenance trajectory of the original code is uncertain. This work provides a faithful, ground-truth numerical reimplementation of *both* ATLAS12 and SYNTHE in pure Python — not a wrapper around Fortran, but a line-by-line translation validated against the original.

A fixed atmosphere misses the structural response, even when fitting just one element.

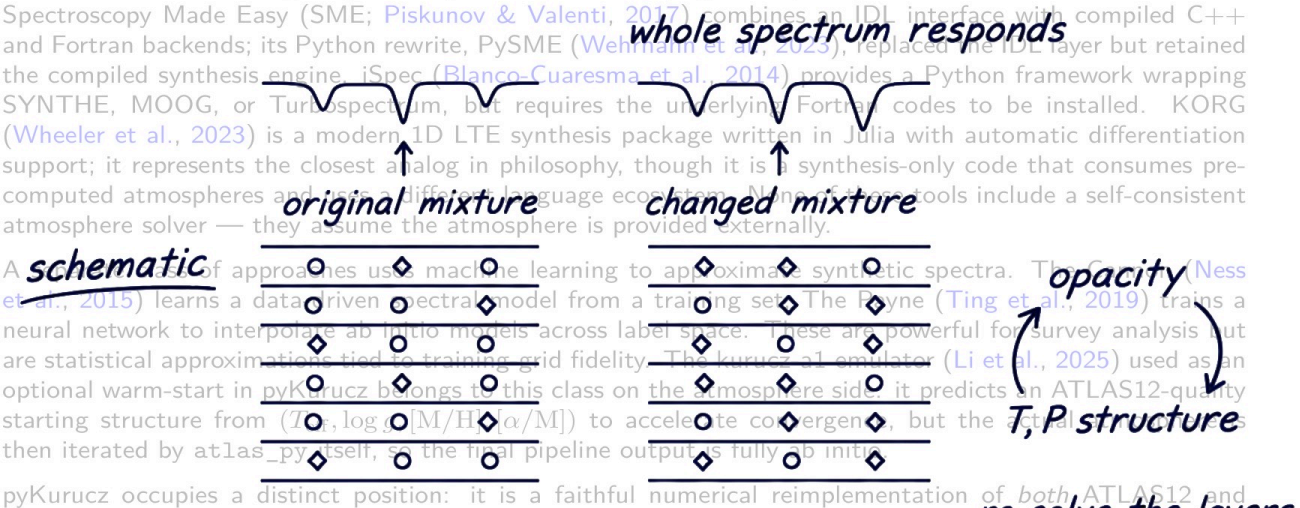
Direct opacity sampling handles peculiar mixtures without rebuilding opacity tables.

A Python reimplementation offers concrete benefits over Fortran: direct integration with machine learning workflows (training emulators, differentiable spectral fitting); interoperability with the broader Python ecosystem (NumPy, SciPy, Astropy); flexibility for customization without recompiling Fortran; cross-platform installation with no compiler dependency; and improved readability for teaching and debugging.

State of the field

Stellar atmosphere modeling and spectrum synthesis have long been dominated by Fortran-based codes. On the atmosphere side, the most widely used grids and engines are Kurucz's ATLAS9/ATLAS12 (Kurucz, 2005), MARCS, and PHOENIX. On the synthesis side, the standards are SYNTHE (Kurucz, 2005), MOOG (Sneden, 1973), Turbospectrum (Plez, 2012), and SPECTRUM (Gray, 1999). While highly accurate, these codes require Fortran toolchains, are difficult to install across operating systems, often lack documentation, and do not integrate naturally with modern Python-based analysis workflows.

Several efforts have brought synthesis (but not the underlying atmosphere) into more modern environments. Spectroscopy Made Easy (SME; Piskunov & Valenti, 2017) combines an IDL interface with compiled C++ and Fortran backends; its Python rewrite, PySME (Wehmeier et al., 2023), replaced the IDL layer but retained the compiled synthesis engine. iSpec (Blanco-Cuaresma et al., 2014) provides a Python framework wrapping SYNTHE, MOOG, or Turbospectrum, but requires the underlying Fortran codes to be installed. KORG (Wheeler et al., 2023) is a modern 1D LTE synthesis package written in Julia with automatic differentiation support; it represents the closest analog in philosophy, though it is a synthesis-only code that consumes pre-computed atmospheres and lives in a different language ecosystem. None of these tools include a self-consistent atmosphere solver — they assume the atmosphere is provided externally.



A growing set of approaches uses machine learning to approximate synthetic spectra. The Cannon (Ness et al., 2015) learns a data-driven spectral model from a training set. The Payne (Ting et al., 2019) trains a neural network to interpolate ab initio models across label space. These are powerful for survey analysis but are statistical approximations tied to training-grid fidelity. The kurucz_a1_emulator (Li et al., 2025) used as an optional warm-start in pyKurucz belongs to this class on the atmosphere side: it predicts an ATLAS12-quality starting structure from $(T_{\rm eff}, \log g, [\mathrm{M/H}], [\alpha/\mathrm{M}])$ to accelerate convergence, but the actual solve is then iterated by `atlas_py` itself, so the final pipeline output is fully ab initio.

pyKurucz occupies a distinct position: it is a faithful numerical reimplementation of *both* ATLAS12 and SYNTHE — the specific codes whose line lists and physics have been the de facto standard for decades — validated against the Fortran ground truth, requiring no Fortran at any stage. It complements the other tools above by exposing self-consistent atmosphere modelling and synthesis within a pure Python environment.

The compared modern tools use externally provided atmospheres; pyKurucz solves atmosphere and synthesis together.

pyKurucz combines an ATLAS12 atmosphere loop, layered spectrum synthesis, and an optional neural warm-start.

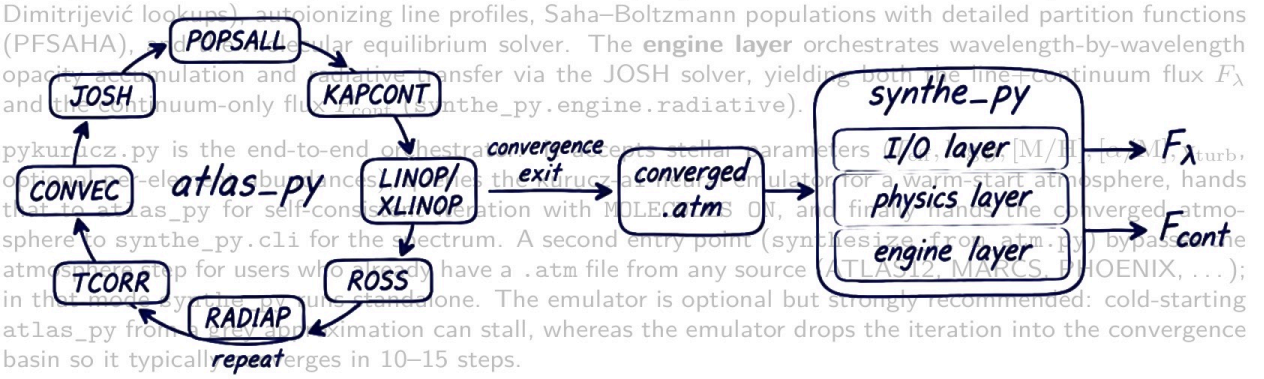
Software design

pyKurucz is organized as a modular pipeline with two cooperating engines and an optional neural warm-start.

`atlas_py` (the atmosphere engine) implements the full ATLAS12 iteration loop. Each outer iteration follows the same eight-step sequence as the Fortran code: POPSALL (Saha–Boltzmann populations + partition functions, with NMOLEC for the molecular equilibrium when MOLECULES ON), KAPCONT (continuous opacity), LINOP/XLINOP (atomic line opacity from preselected Kurucz `fort.12` records and NLTE `fort.19` lines), ROSS (Rosseland mean opacity and optical depth scale), RADIAP (radiative flux and acceleration), TCORR (temperature corrections from flux errors), CONVEC (mixing-length convection where unstable), and JOSH (frequency-by-frequency radiative transfer). HLINOP supplies hydrogen-line wing opacity. Convergence is tested on the maximum normalized change across the physical atmosphere columns (RHOX, T, P, XNE, ABROSS, VTURB) and the loop exits early once a user-supplied tolerance is met for a configurable number of consecutive iterations. The output is a converged Kurucz-format `.atm` file.

Repeat the eight physical steps until the atmosphere stays within the convergence tolerance.

`synthe_py` (the synthesis engine) is organized into three layers. The **I/O layer** (`synthe_py.io`) handles atmosphere loading (Kurucz `.atm` and preprocessed `.npz` formats), atomic line catalog parsing and compiled-binary export. The **physics layer** (`synthe_py.physics`) implements continuous opacity interpolation (KAPP), line opacity with Voigt profiles (`voigt_jit.py` for Numba-accelerated evaluation), hydrogen Stark-broadened profiles (HPROF4/HLINOP port), helium tabulated line profiles (BCS, Griem, and Dimitrijević lookups), autoionizing line profiles, Saha–Boltzmann populations with detailed partition functions (PFSAHA), and molecular equilibrium solver. The **engine layer** orchestrates wavelength-by-wavelength opacity accumulation and radiative transfer via the JOSH solver, yielding both the line+continuum flux F_λ and the continuum-only flux $F_{\rm cont}$ (`synthe_py.engine.radiative`).



`pykurucz.py` is the end-to-end orchestrator. It accepts stellar parameters ($T_{\rm eff}$, $\log g$, $[{\rm M/H}]$, $[\alpha/{\rm M}]$, $v_{\rm turb}$, optional per-element abundances), invokes the Kurucz-a1 neural emulator for a warm-start atmosphere, hands that to `atlas_py` for self-consistent iteration with MOLECULES ON, and finally hands the converged atmosphere to `synthe_py.cli` for the spectrum. A second entry point (`synthesize_from_atm.py`) bypasses the atmosphere step for users who already have a `.atm` file from any source (ATLAS12, MARCS, PHOENIX, ...); in that mode Kurucz is the only code path. The emulator is optional but strongly recommended: cold-starting `atlas_py` from the LTE approximation can stall, whereas the emulator drops the iteration into the convergence basin so it typically converges in 10–15 steps.

pyKurucz reuses the same input data as the original Fortran. The Kurucz atomic line catalogues (GFALL, predicted lines, low- and high-excitation lists, NLTE lists), molecular line lists (Schwenke TiO, Partridge–Schwenke H_2O, and the GFALL diatomic pack), partition functions, ionization potentials, continuum opacity tables, and molecular equilibrium data are stored as compact binary `.npz` archives pre-extracted from the Fortran data files, ensuring the Python pipeline operates on exactly the same physical data as the Fortran code. Numerical validation confirms that opacity, line selection, JOSH integration — use Numba JIT compilation, and the wavelength-parallel radiative transfer worker pool keeps wall-time competitive with Fortran.

Identical Kurucz input data isolate implementation differences from differences in the line lists.

Research impact statement

pyKurucz has been validated end-to-end against the Fortran ATLAS12 + SYNTHE pipeline. To isolate engine-implementation differences from initial-condition differences, both pipelines are launched from the *same* Kurucz grid-warmstart `.atm`, with the same microturbulent velocity, and `atlas_py` is run for the same 30 outer iterations as the Fortran ATLAS12 deck. Five representative comparisons spanning $T_{\rm eff} = 2500$–44000 K are included with the package (Figures 1–5). Across all five, the median normalized-flux deviation stays below 0.005%, with 99th-percentile deviations below 0.3% and sub-1% maxima — essentially the same level as the synthe-only validation against Fortran SYNTHE on a fixed atmosphere, which establishes that `atlas_py` and `synthe_py` reproduce Fortran ATLAS12 + SYNTHE end-to-end at the level of numerical noise on these representative test cases.

The user-facing default flow in `pykurucz.py` uses the kurucz-a1 neural emulator (Li et al., 2025) for the warm-start, which is fast and supports arbitrary off-grid abundance patterns; the resulting spectra agree with Fortran to within a few tenths of a percent — well below NLTE/3D/line-list systematics. Users who need bit-faithful Fortran reproducibility can instead point `synthe.py` at any existing `.atm` file, which agrees with Fortran SYNTHE at the level of numerical noise.

Numba accelerates performance-critical routines; a wavelength-parallel worker pool keeps wall-time competitive with Fortran.

Even the coolest, most molecularly crowded validation case remains nearly indistinguishable from the Fortran pipeline.

This opens several near-term research applications. First, pyKurucz enables generation of large *self-consistent* training grids for data-driven spectral models (e.g., The Payne, The Cannon) without a Fortran toolchain, lowering the barrier for groups building emulators for surveys such as SDSS-V, DESI, 4MOST, and WEAVE — and, crucially, allowing those grids to span peculiar abundance patterns (CEMP, Ap, r-process) that ATLAS9/ODF approaches cannot easily reach. Second, the pure Python pipeline can be embedded into gradient-based optimization frameworks (PyTorch, JAX) for differentiable spectral fitting where the atmosphere itself participates in the gradient. Third, the readable codebase facilitates classroom instruction and rapid prototyping of new physics. Finally, as a faithful preservation of ATLAS12 and SYNTHE in a modern language, pyKurucz serves as an archival reference implementation ensuring long-term reproducibility.

Using identical warm-starts and iteration counts makes this a controlled code comparison, not a convergence contest.

This fidelity is what makes unusual-composition training grids scientifically useful rather than merely convenient.

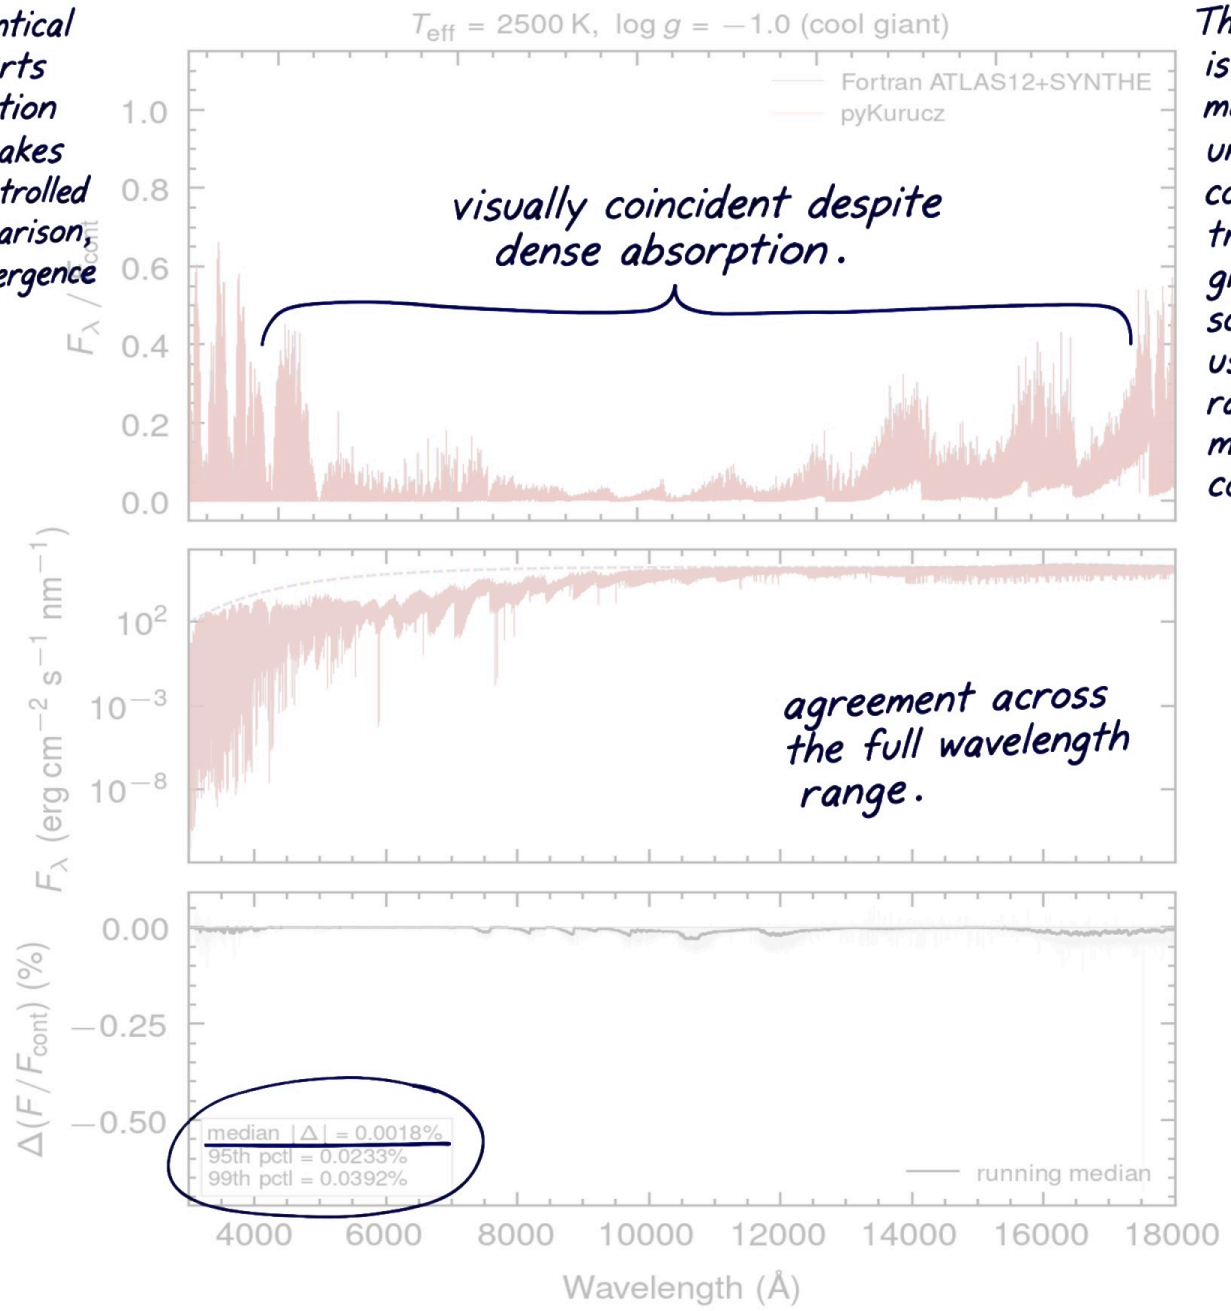
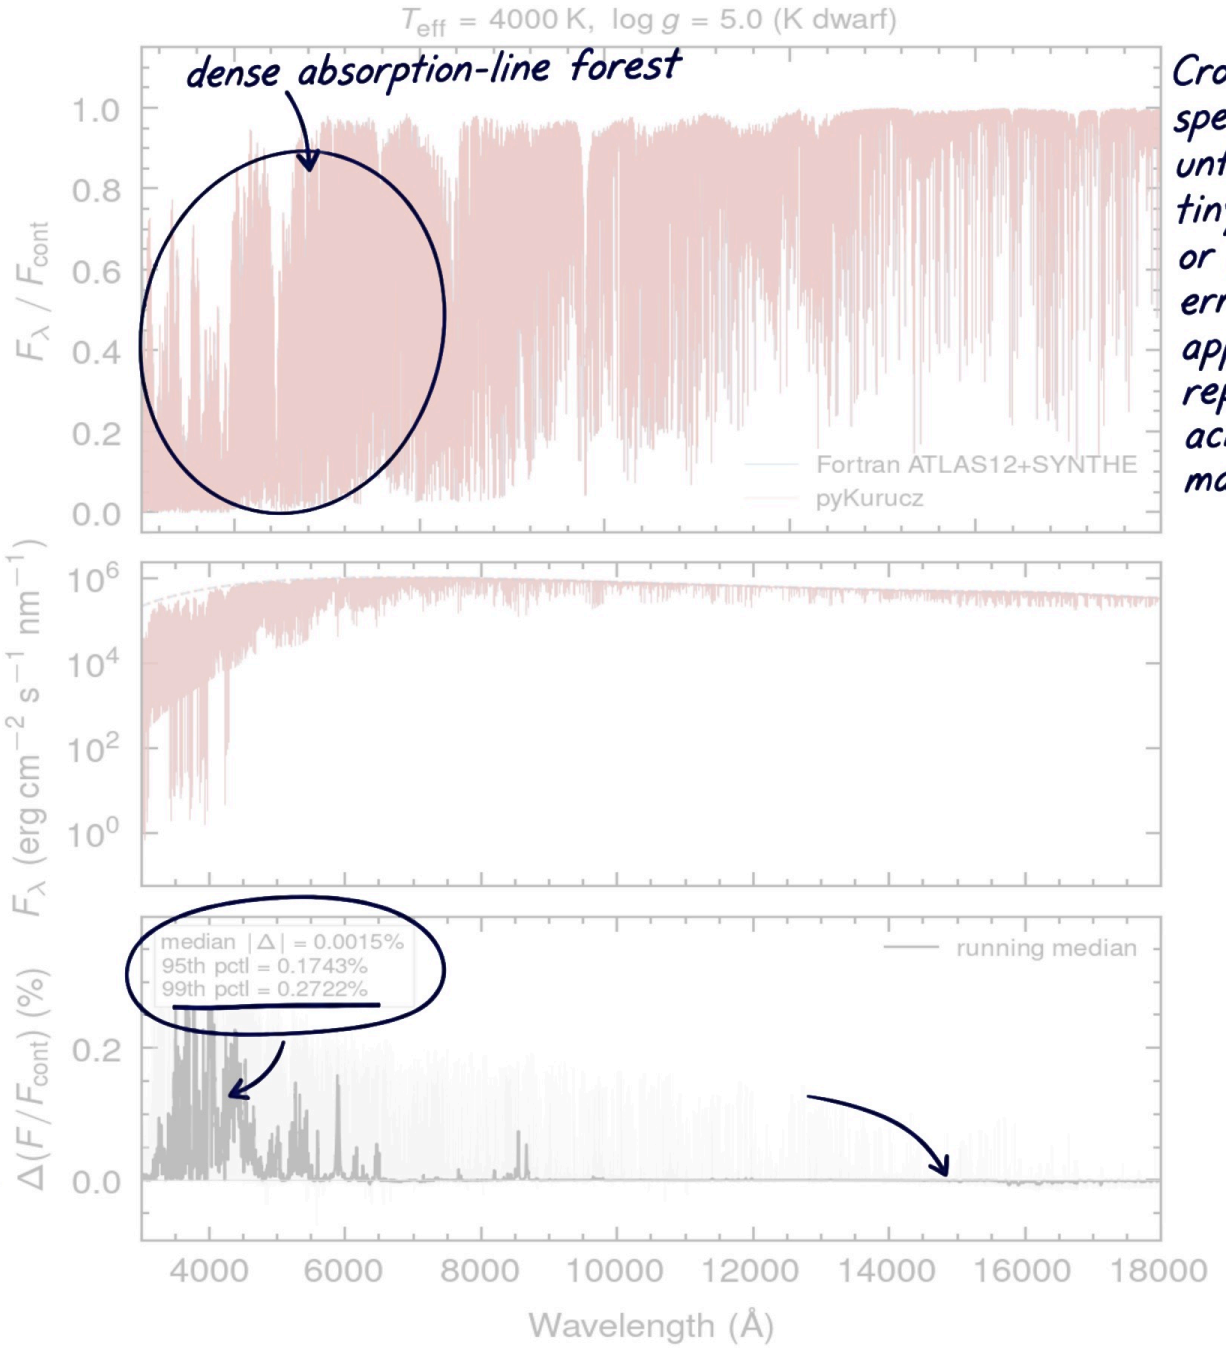


Figure 1: End-to-end pyKurucz vs Fortran ATLAS12 + SYNTHE for a cool giant ($T_{\rm eff} = 2500\,$K, $\log g = -1.0$). Both pipelines start from the same Kurucz grid-warmstart `.atm`, run 30 ATLAS12 outer iterations, and then synthesize 300–1800 nm at $R = 300{,}000$. Top: normalized flux overlay (3500–9000 Å). Middle: full flux spectrum (3000–18000 Å, log scale). Bottom: fractional deviation of normalized flux with running median.

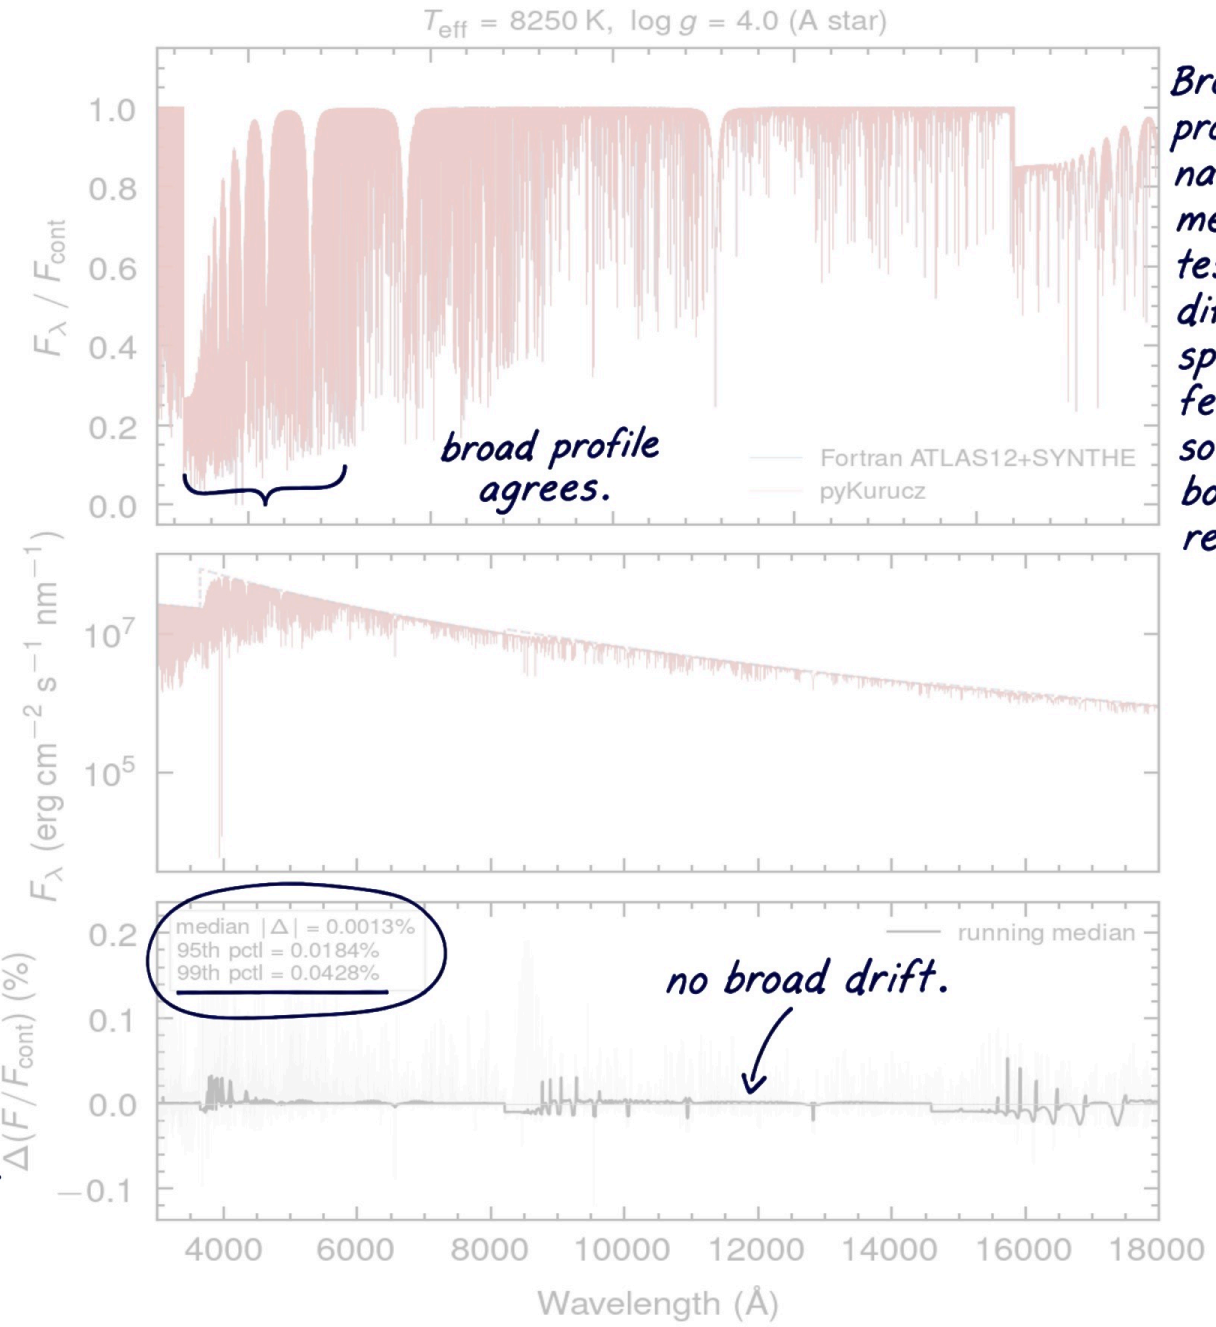
The K-dwarf test shows that agreement survives a dense forest of narrow absorption lines.



The blue end is the hardest region here, where blanketing is strongest and the residual median rises.

Figure 2: End-to-end comparison for a K dwarf ($T_{\rm eff} = 4000\,\mathrm{K}$, $\log g = 5.00$).

For the A star, the Python calculation matches both broad absorption structure and narrow lines across the spectrum.



Broad profiles and narrow metal lines test different spectral features, so matching both is reassuring.

The remaining residual structure is localized rather than a wavelength-wide continuum offset.

Figure 3: End-to-end comparison for an A-type star ($T_{\rm eff} = 8250\,\mathrm{K}$, $\log g = 4.00$).

The late-B dwarf confirms that the same implementation remains stable as ionization and continuum balance shift with temperature.

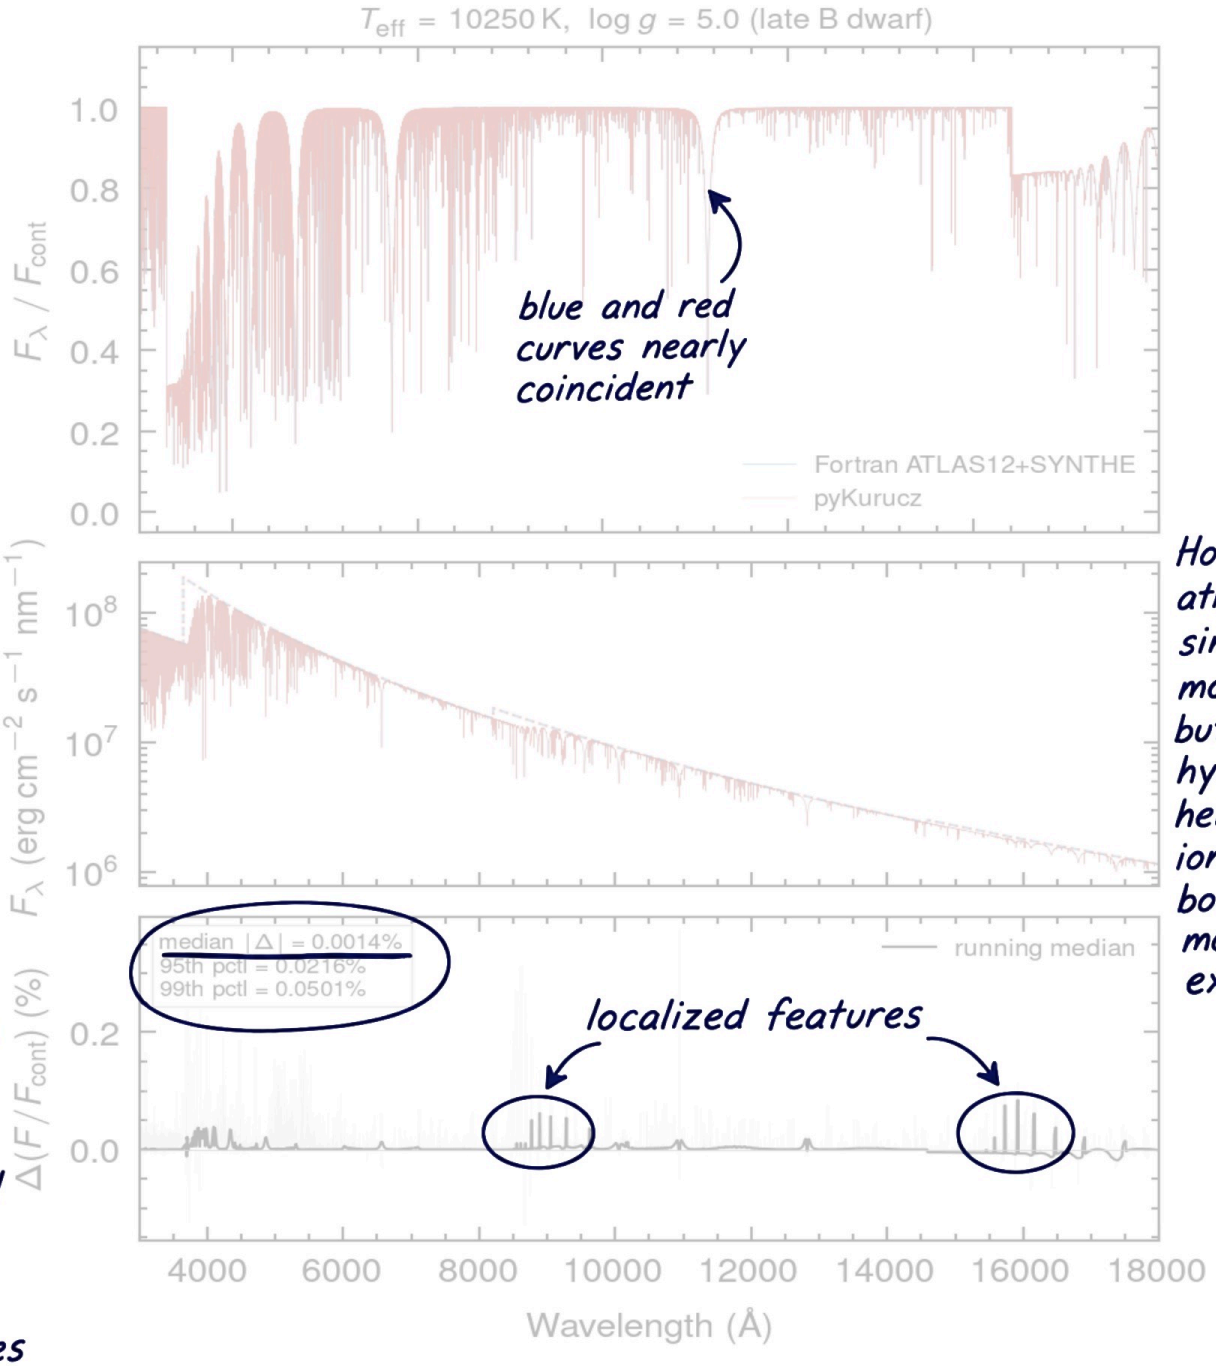


Figure 4: End-to-end comparison for a late B dwarf ($T_{\rm eff} = 10250\,\mathrm{K}$, $\log g = 5.00$).

The hot O-star test still agrees closely, but both engines' shared LTE assumption warrants caution for specific lines.

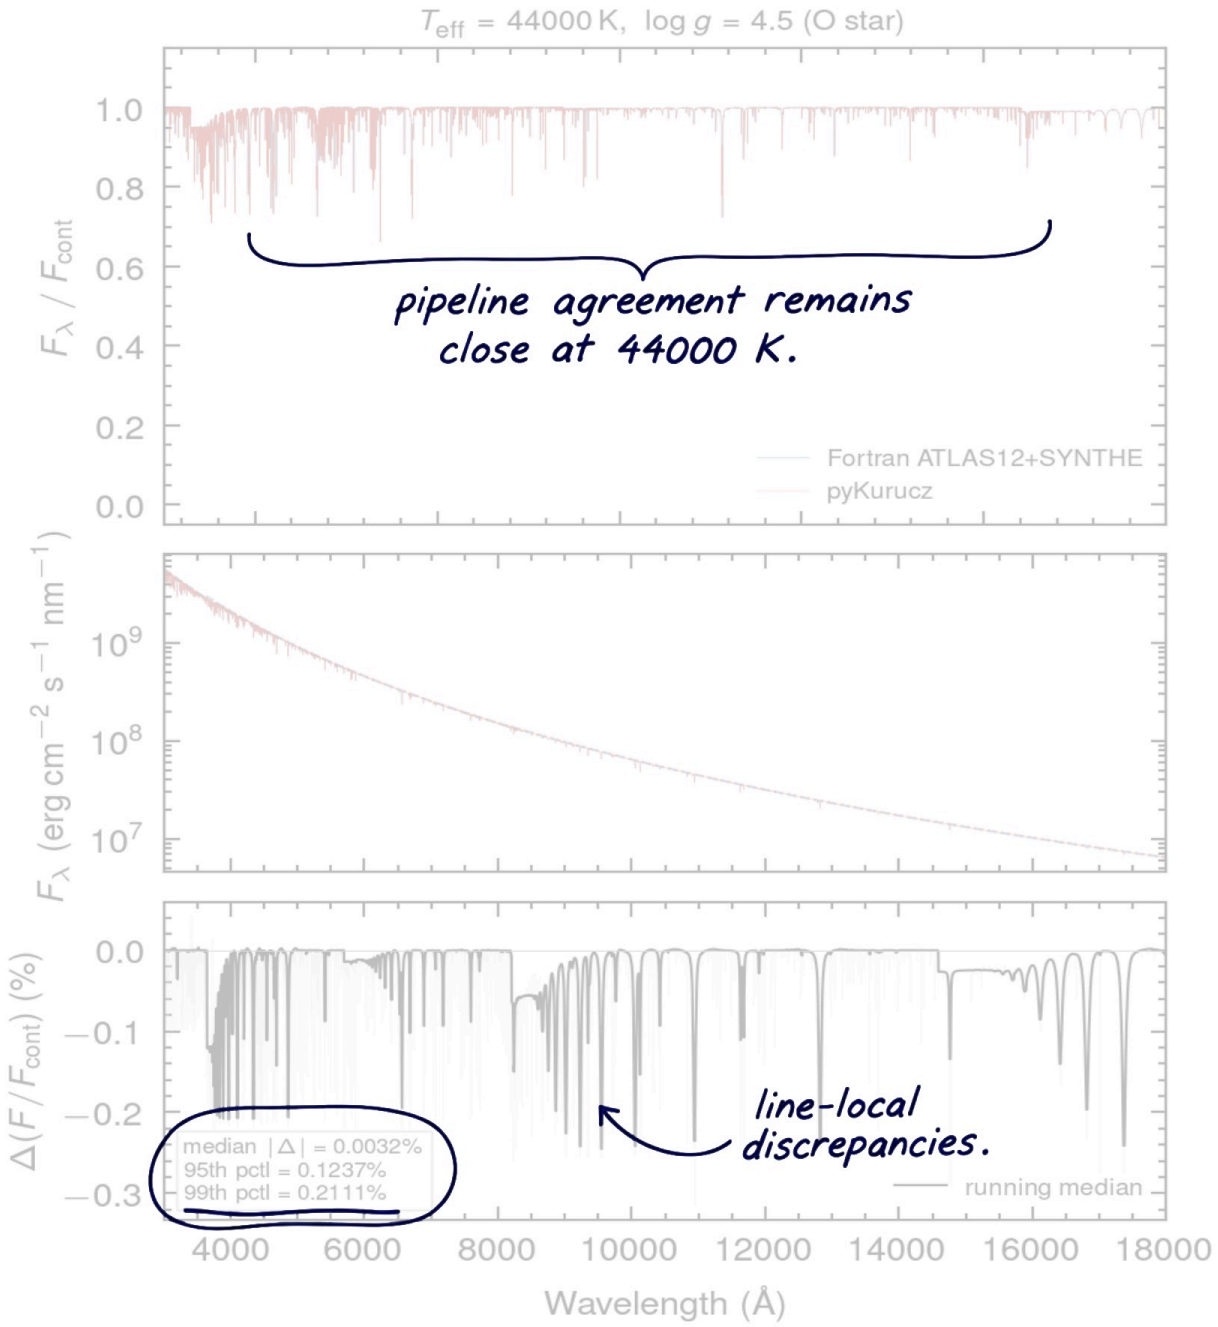


Figure 5: End-to-end comparison for a hot O-type star ($T_{\rm eff} = 44000\,\mathrm{K}$, $\log g = 4.50$).

Limitations and future work

Agreement with Fortran validates our translation; it does not validate LTE for real O-star atmospheres.

pyKurucz reimplements the full ATLAS12 + SYNTHE path through atomic and molecular line opacity, but several limitations remain, each defining a clear direction for future development.

LTE only. Both engines assume local thermodynamic equilibrium. Non-LTE effects matter for specific lines (Li I 6707 Å, Na D, O I triplet) in metal-poor and hot stars. Departure coefficients from external NLTE codes can be applied as corrections.

External departure coefficients can correct selected lines, but they do not turn atlas_py into an NLTE atmosphere solver.

The remaining boundaries are 1D geometry, a fragile cool-star corner case, and performance optimization.

1D plane-parallel geometry. Like the original Fortran, all models assume plane-parallel stratification. 3D hydrodynamic atmospheres (Stagger, CO^{5}BOLD) can be ingested by `synthe_py` for post-processing, but `atlas_py` itself is 1D.

atlas_py solves a 1D atmosphere; synthe_py can use supplied 3D atmospheres for post-processing.

Cool RSG α-perturbed corner cases. In a narrow region of parameter space — cool ($T_{\rm eff} \lesssim 4500\,$K), low-gravity ($\log g \lesssim 0.0$) models with [$\alpha$/Fe] $\gtrsim +0.3$ — the Newton solver's element partial emulator's prior is far enough from the true converged solution that `atlas_py` can iterate into all-NaN before recovering. Defensive guards trap the failure with a clean error rather than silently writing a degenerate atmosphere; the recommended escape hatch is a *neighbour warm-start*, in which a converged neighbour cell's atmosphere is used as the initial guess. This [illegible] emission-line stars.

For the cool, low-gravity corner case, start from a converged neighbour with the target chemistry.

Performance. The current pure-Python implementation is deliberately written as a straightforward, original for equivalent calculations. Since the codebase is readable, well-structured Python with Numba-JIT inner loops, we expect that AI coding agents — the same tools used to create this reimplementation — can systematically profile and optimize the bottlenecks, making this a tractable engineering task rather than a fundamental limitation.

AI usage disclosure

A pure Python reimplementation of ATLAS12 and SYNTHE has been a long-standing aspiration in the stellar spectroscopy community — discussed informally for decades — but the software engineering effort has been prohibitive. The original Fortran is written in a pre-Fortran-77 style with fixed-format source, single-character variable names, pervasive use of computed GOTOs and EQUIVALENCE statements, and implicit typing throughout — making it incompatible with modern `gfortran` without significant manual patching, and [illegible] Kurucz's [illegible] This project attempted this reimplementation repeatedly since the early days of large language models, but earlier models lacked the sustained context and code reasoning needed to navigate Kurucz's deeply interlinked subroutines. It was only with Anthropic's Claude Opus 4.6 (via Cursor) that the project became tractable — though "tractable" still required extensive human guidance: several thousand US dollars in API tokens and several months of the authors' time directing, debugging, and validating the output. Every AI-generated routine was tested against the Fortran reference, first at the subroutine level and then end-to-end across a grid of atmosphere models. We view this as a demonstration of AI-assisted scientific software preservation: not fully autonomous, but a qualitative leap in what is feasible for a small team tackling a large legacy codebase.

AI helped translate the code; subroutine tests and end-to-end Fortran comparisons established fidelity.

Acknowledgements

This project is dedicated to the memory of Robert L. Kurucz (1944–2025), whose ATLAS and SYNTHE codes, atomic and molecular line lists, and freely shared data have served the astronomical community for over half a century.

References

Blanco-Cuaresma, S., Soubiran, C., Heiter, U., & Jofré, P. 2014, Astronomy & Astrophysics, 569, A111, doi: 10.1051/0004-6361/201423945

Gray, R. O. 1999, SPECTRUM: A stellar spectral synthesis program. https://ui.adsabs.harvard.edu/abs/1999ascl.soft10002G

Karzas, W. J., & Latter, R. 1961, The Astrophysical Journal Supplement Series, 6, 167, doi: 10.1086/190063

Kurucz, R. L. 2005, in Memorie della Società Astronomica Italiana Supplementi, Vol. 8, 14. https://ui.adsabs.harvard.edu/abs/2005MSAIS...8...14K

Li, J., Jian, M., Ting, Y.-S., & Green, G. M. 2025, arXiv e-prints

Ness, M., Hogg, D. W., Rix, H.-W., Ho, A. Y. Q., & Zasowski, G. 2015, The Astrophysical Journal, 808, 16, doi: 10.1088/0004-637X/808/1/16

Piskunov, N., & Valenti, J. A. 2017, Astronomy & Astrophysics, 597, A16, doi: 10.1051/0004-6361/201629124

Plez, B. 2012, Turbospectrum: Code for spectral synthesis. https://ui.adsabs.harvard.edu/abs/2012ascl.soft05004P

Sneden, C. A. 1973, The Astrophysical Journal, 184, 839, doi: 10.1086/152374

Ting, Y.-S., Conroy, C., Rix, H.-W., & Cargile, P. 2019, The Astrophysical Journal, 879, 69, doi: 10.3847/1538-4357/ab2331

Wehrhahn, A., Piskunov, N., & Ryabchikova, T. 2023, Astronomy & Astrophysics, 671, A171, doi: 10.1051/0004-6361/202244482

Wheeler, A. J., Abruzzo, M. W., Casey, A. R., & Ness, M. K. 2023, The Astronomical Journal, 165, 68, doi: 10.3847/1538-3881/acaaad