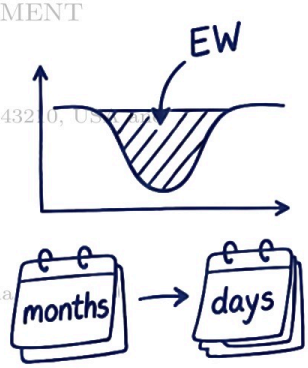
Classical Voigt fitting does the physics; an LLM inspects, repairs, or rejects difficult fits.

Version May 7, 2026
Preprint typeset using LaTeX style openjournal v. 09/06/15

Egent automates the expert judgment that makes equivalent-width measurement slow.

EGENT: AN AUTONOMOUS AGENT FOR EQUIVALENT WIDTH MEASUREMENT



Yuan-Sen Ting (丁源森)
Department of Astronomy, The Ohio State University, Columbus, OH 43210, USA
Center for Cosmology and AstroParticle Physics (CCAPP), The Ohio State University, Columbus, OH 43210, USA
Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

Serat Mahmud Saad
Department of Astronomy, The Ohio State University, Columbus, OH 43210, USA

Fan Liu (刘凡)
National Astronomical Observatories, Chinese Academy of Sciences, Beijing 100012, China

Yuting Shen
School of Computer Science, Georgia Institute of Technology, Atlanta, GA 30332, USA
Version May 7, 2026

arXiv:2512.01270v2 [astro-ph.IM] 6 May 2026

ABSTRACT

We present Egent, an autonomous agent that combines classical multi-Voigt profile fitting with large language model (LLM) visual inspection and iterative refinement. The fitting engine is built from scratch with minimal dependencies, creating an ecosystem where the LLM can reason about fits through function calls—adjusting wavelength windows, adding blend components, modifying continuum treatment, and flagging problematic cases. Egent operates directly on raw flux spectra without requiring pre-normalized continua. We validate against manual measurements from human experts using 18,615 lines from the C3PO program across 84 Magellan/MIKE spectra at SNR ∼50–250 per pixel. The raw agreement between Egent and expert measurements is MAD = 5–7 mÅ without any post hoc per-spectrum correction. Per-spectrum slopes of ∼0.85–1.19 around unity reflect differences in global continuum methodology rather than fitting failures. The LLM's primary role is quality control: it confirms good fits (∼60–65% of lines are LLM-refined and accepted), flags problematic cases (∼10–20%), and occasionally rescues edge cases where tool use improves fits. Agreement between GPT-5 and GPT-5-mini confirms reproducibility, with GPT-5-mini enabling low-cost analysis at ∼200 lines per US dollar. Every fit stores complete Voigt parameters, continuum coefficients, and LLM reasoning chains, enabling exact reconstruction without re-running. Egent compresses what traditionally requires months of expert effort into days of automated analysis, enabling survey-scale EW measurement. We provide open-source code at https://github.com/tingyuansen/Egent, including a web interface for drag-and-drop analysis and a local LLM backend for fully offline operation on consumer hardware.

Against 18,615 expert measurements, raw disagreement is typically 5–7 mÅ.

Keywords:

1. INTRODUCTION

Large spectroscopic surveys such as APOGEE (Majewski et al. 2017), GALAH (Buder et al. 2025), Gaia-ESO (Randich et al. 2022), and the ongoing SDSS-V (Kollmeier et al. 2017), 4MOST (de Jong et al. 2012), and WEAVE (Dalton et al. 2014) have propelled astronomy into the era of big spectroscopic data. The need for automated analysis has driven the adoption of full-spectrum fitting methods for stellar parameter and abundance determination (e.g., Ness et al. 2015; Holtzman et al. 2015; García Pérez et al. 2016; Ting et al. 2019; Xiang et al. 2022). Historically, however, spectra were analyzed through equivalent width (EW) measurements (e.g., Korista et al. 1997; Milone et al. 2013; Montes et al. 2018): individual spectral lines have characteristic shapes, and the integrated area under a line—the equivalent width—directly encodes the number of absorbing atoms in the stellar atmosphere, providing elemental abundances.

EW measurement persists for small-scale, high-precision studies. For high-resolution differential analysis, line-by-line EW measurement remains the standard approach, because individual spectral lines carry different systematic errors, including uncertain oscillator strengths (Scott et al. 2015a,b), unmodeled blending, and departures from local thermodynamic equilibrium (Asplund et al. 2000, 2009; Bergemann et al. 2019; Lind and Amarsi 2024). Full-spectrum fitting averages over these systematics, while differential EW analysis between similar stars cancels them. This line-by-line approach remains the consensus method for achieving precision abundance measurements (Bensby et al. 2004; Bedell et al. 2014; Nissen 2016). Solar twin studies (Ramirez et al. 2009), planet engulfment signatures (Yana Galarza

The useful role for an LLM is judgment at awkward edge cases, not replacing the physics.

et al. 2021; Liu et al. 2024; Spina 2024), and precise elemental ratios all rely on it.

Existing tools range from fully interactive programs like IRAF's `splot`, where each line requires manual continuum placement and blend decisions, to automated codes like ARES (Sousa et al. 2007), DAOSPEC (Stetson and Pancino 2008), and TAME (Kang and Lee 2012), and integrated frameworks like iSpec (Blanco-Cuaresma et al. 2014) and SMHR[1] that combine EW measurement with abundance analysis. Beyond stellar spectroscopy, automated spectral decomposition has been tackled by tools such as GaussPy (Lindner et al. 2015) and GaussPy+ (Riener et al. 2019), which use machine-learning-assisted Gaussian decomposition for ISM spectral line surveys (21 cm absorption and CO/H I emission, respectively); these share the goal of replacing manual inspection with algorithmic decisions, though they target different data regimes than stellar absorption-line work.

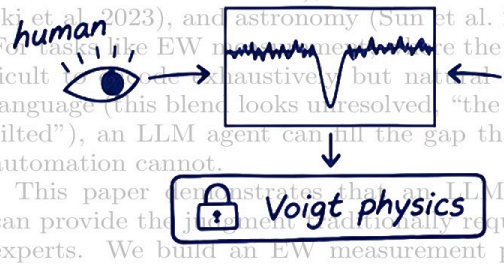
These codes can process line lists in batch mode, but they still require parameter calibration, typically expect pre-reduced spectra, and rely on visual examination of fits to identify problematic cases. The contribution of Egent is not multi-component fitting or heuristic quality control—both are well established—but the integration of an LLM visual inspector that reasons about diagnostic plots through function-calling tools, providing a quality-control loop that adapts to each fit without requiring hand-tuned rules. For a typical differential abundance study of ∼100 stars with ∼200 lines each, an expert may spend months on measurements and quality checks.

While in principle EW measurement involves fitting Voigt profiles to many lines simultaneously, full automation remains difficult—analogous to source extraction in imaging, where even with known profiles, edge cases abound. Simple algorithms fail on unrecognized blends, curved continuum from nearby strong lines, and crowded spectral regions (Ramirez et al. 2009; Scott et al. 2015b; Ting et al. 2021). Rule-based systems require extensive tuning for each spectrograph and stellar type, yet still miss unusual cases. Difficult lines have historically required human eyes and judgment.

Natural-language rules like "the continuum is tilted" are surprisingly hard to encode exhaustively.

A further complication is the instrumental response. Raw spectra do not show a flat continuum from the stellar photosphere; instead, the observed flux is modulated by the spectrograph's blaze function—a wavelength-dependent efficiency that varies smoothly across each échelle order (Škoda et al. 2008, see Figure 1a). This has led to extensive effort to "continuum normalize" spectra before fitting (Ivanov et al. 2004; Blanco-Cuaresma et al. 2014; García Pérez et al. 2016), introducing additional subjective choices that propagate into systematic errors.

Recent advances in large language models (LLMs) suggest a path forward. Modern LLMs exhibit robust reasoning capabilities (Wei et al. 2022; Achiam et al. 2023), and when equipped with function-calling tools, they become agents capable of taking actions (Yao et al. 2022), not merely generating text like a chatbot, but executing code, inspecting results, and iterating. Such agents have demonstrated autonomous scientific research capabilities in chemistry (Boiko et al. 2023; Bran et al. 2023; Caldas Ramos et al. 2024), materials science (Szymanski et al. 2023), and astronomy (Sun et al. 2024, 2025a). For tasks like EW measurement where the rules are difficult to encode exhaustively but natural to express in language (this blend looks unresolved, "the continuum is tilted"), an LLM agent can fill the gap that rule-based automation cannot.



This paper demonstrates that an LLM-based agent can provide the judgment traditionally requiring human experts. We build an EW measurement pipeline from scratch, creating an ecosystem of tools that the LLM can invoke: fitting functions, diagnostic plots, continuum adjustments, and quality flags. By fitting locally within each wavelength window, Egent avoids the need for global continuum normalization—the LLM handles continuum placement as part of its iterative refinement. All decisions are logged, providing complete provenance that manual measurements lack. Egent occupies a middle ground: classical Voigt fitting (fast, interpretable) with LLM-guided decisions for difficult cases, uniquely providing complete provenance for every measurement.

2. METHOD: THE EGENT PIPELINE

2.1. *Design Philosophy*

For an agentic pipeline, we need tools that the agent can invoke—mimicking how a human expert would interact with traditional software like IRAF: zooming in on a spectral region, choosing line centers, deciding continuum placement, adding blend components, etc. Our first approach was to wrap IRAF itself (via PyRAF) and let the LLM agent drive it, but this proved difficult due to the many interconnected components and legacy dependencies.

I first tried letting the agent drive IRAF; legacy dependencies made that far more fragile than rebuilding the fitter.

Instead, we recognized that for EW fitting specifically, the core task reduces to multi-component Voigt profile fitting—something straightforward to implement from scratch. Egent is built in pure Python using only standard scientific libraries (`numpy`, `scipy.optimize`, `scipy.special` for the Faddeeva function). We deliberately avoid legacy packages and complex dependencies; everything is self-contained and inspectable, apart from the LLM API calls—the interface through which our code sends prompts to and receives responses from external language model services.[2]

This design serves multiple purposes. First, it ensures reproducibility across computing environments with minimal installation requirements—particularly important for the offline version we present later. Second, version-controlled dependencies and a simple codebase make agentic tool use more robust; complex environ-

[1] https://github.com/andycasey/smhr

[2] For an introduction to LLM APIs and tool use, see https://tingyuansen.github.io/coding_essential_for_astronomers/.

The agent never calculates EW itself—it chooses among tightly constrained scientific tools.

Raw échelle spectra are curved, crowded, and emphatically not continuum-normalized.

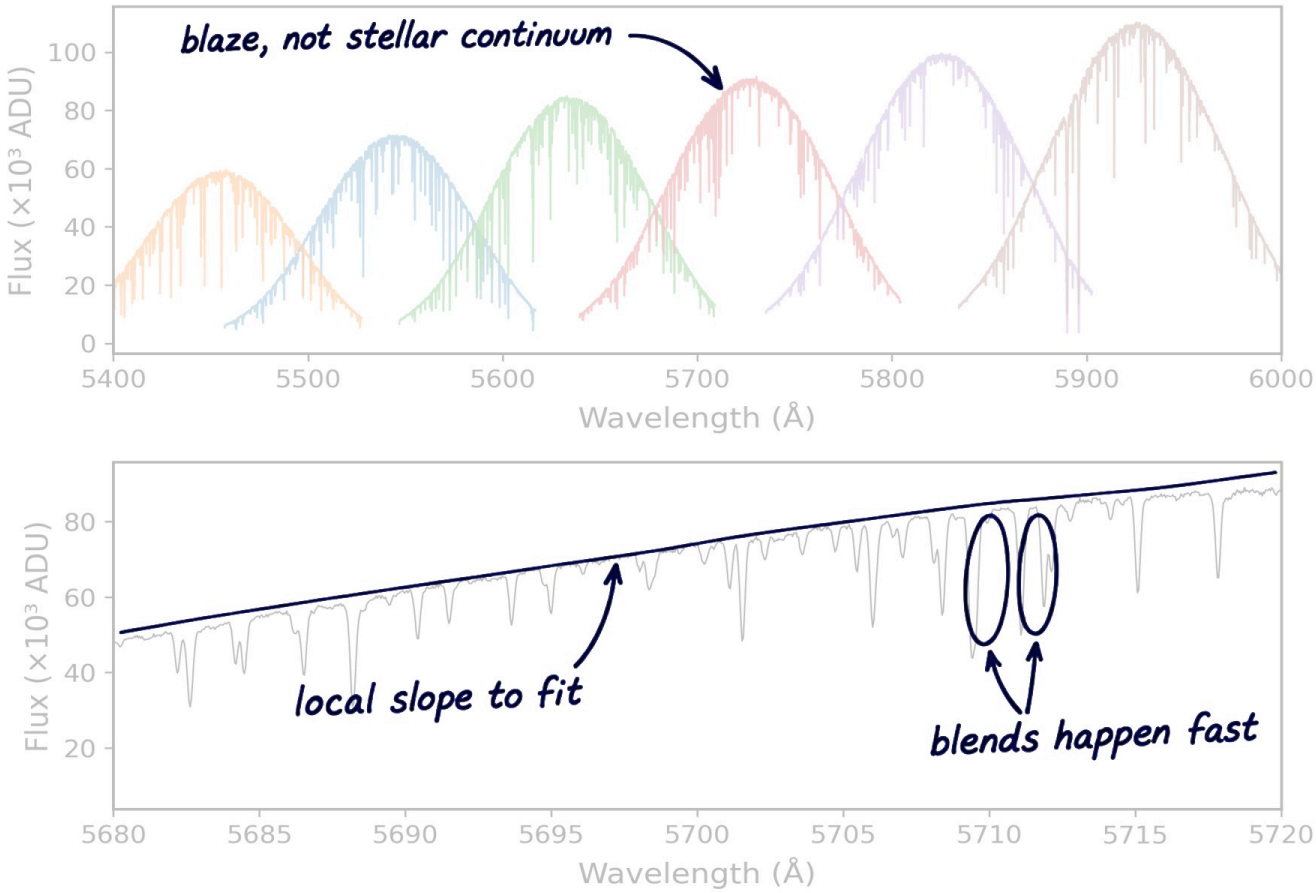


Figure 1. Raw Magellan/MIKE échelle spectrum illustrating challenges for automated EW measurement. *Top panel:* Multiple échelle orders (different colors) showing the characteristic blaze function—the instrumental response that modulates the observed flux with a curved envelope peaking near each order's center. Absorption lines appear as narrow dips superimposed on this curved background. *Bottom panel:* Single order zoom (5680–5720 Å) showing raw flux without continuum normalization. The underlying continuum slopes upward following the blaze function, demonstrating why local continuum fitting is required. Visible challenges include blended lines (multiple absorptions within <1 Å), weak lines near the noise level, and the varying continuum slope that traditional methods must remove before fitting.

The blaze shape is smooth globally, but a small local window usually needs only a simple polynomial.

ments with many package-level dependencies can fail in failure modes that degrade the agent's reliability. Third, the transparent fitting algorithm allows users to understand and modify the agent's behavior.

The division of labor is that physics defines the fitting; the LLM operates through tools. The LLM does not perform calculations or invent new methods. It inspects diagnostic plots, identifies problems that a human expert would notice, and suggests parameter changes using the same tools a human would use. This division of labor leverages the LLM's strength (visual pattern recognition and reasoning) while keeping the physics grounded in well-understood spectral line models.

We primarily use GPT-5-mini for production runs, with GPT-5 serving as a cross-check for validation. Both models support function calling and vision capabilities required for our pipeline. GPT-5-mini processes approximately 200 lines per dollar (averaging ∼5,000 input tokens and ∼1,000 output tokens per line across multiple iterations), while GPT-5 is roughly five times more expensive (∼40 lines per dollar). The consistency between models (demonstrated in Section 4) validates that our pipeline does not rely on model-specific behaviors, and the cost efficiency of GPT-5-mini makes large-scale surveys economically feasible.

Minimal dependencies mattered: agents become unreliable when every tool can fail in a different way.

2.2. *Pipeline Overview*

Figure 2 illustrates the complete Egent workflow. Egent expects input spectra already in the stellar rest frame (e.g., after applying barycentric and radial-velocity corrections). Egent requires a user-provided line list specifying target wavelengths to measure; for each target line in this list, the pipeline proceeds as follows:

1. Extract a local spectral region around the target wavelength (default ±3 Å).

2. Identify all catalog lines within this window that might contribute to the observed absorption.

3. Fit a multi-Voigt model plus polynomial continuum simultaneously to the extracted region (Section 2.5).

The spectrum must already be in the stellar rest frame; Egent does not solve wavelength calibration.

The fitter measures; the agent inspects and adjusts difficult cases.

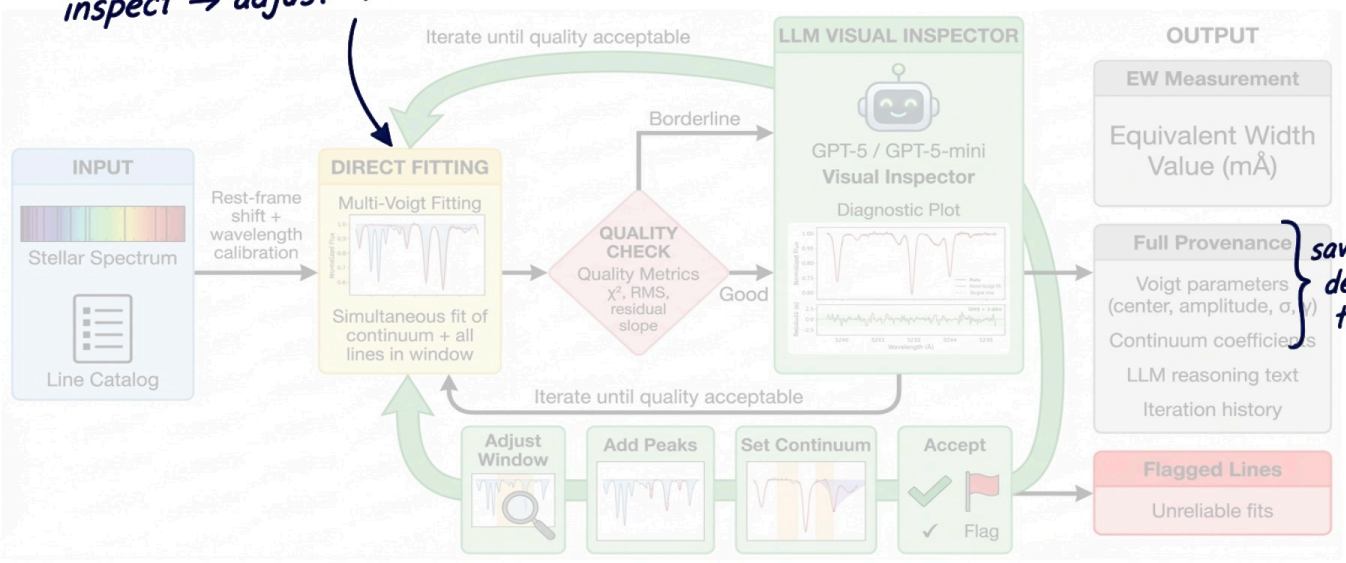


Figure 2. Schematic overview of the Egent pipeline. Input spectra (shifted to stellar rest frame with barycentric and empirical wavelength corrections) and line catalogs enter the direct multi-Voigt fitting stage. Quality metrics determine whether the fit is acceptable or requires LLM inspection. The LLM visual inspector examines diagnostic plots and can adjust the extraction window (Fig. 4), add peaks for blends (Fig. 5), propose continuum regions based on visual reasoning (Fig. 6), or flag unreliable fits. The iteration loop continues until the LLM accepts the measurement. All outputs include complete provenance: Voigt parameters, continuum coefficients, and LLM reasoning chains.

4. Compute quality metrics (Section 2.7) and either accept the fit directly or trigger LLM visual inspection (Section 2.8).

The default window of ±3 Å is chosen to be large relative to the line broadening kernel. At $R = 50{,}000$ (Magellan/MIKE), the instrumental FWHM is $\lambda/R \approx 0.11$ Å near 5500 Å, so the window spans ∼55 resolution elements—sufficient to capture full Voigt wings and neighboring blends. For slowly rotating solar-type stars ($v \sin i \lesssim 5$ km s^{-1}), rotational broadening adds $\lesssim 0.1$ Å, which remains well within the window. The LLM can adapt the window during iterative refinement: narrowing it (e.g., to ±2 Å) for isolated lines where edge features contaminate the continuum fit, or widening it (to ±5 Å) when strong lines are truncated.

A key insight is that simple scalar metrics like χ^2 are insufficient for quality assessment. A fit may have excellent global χ^2 while completely missing a blend at the target wavelength, or conversely, have elevated χ^2 due to edge effects that do not affect the target measurement. What matters is whether the model correctly captures the absorption *at and around the target line*—a judgment that requires inspecting the spatial pattern of residuals, not just their aggregate statistics. This is precisely the kind of visual pattern recognition that LLMs with vision capabilities now excel at.

The recent development of multimodal LLMs with strong performance on non-natural images—including scientific plots—makes this approach timely. Earlier vision models struggled with plots, charts, and diagrams; current models have improved ability to interpret axis labels, identify systematic residual patterns, and reason about what they imply for fit quality.

The following subsections detail each component: the equivalent width definition (Section 2.3), the Voigt profile model (Section 2.4), multi-component fitting (Section 2.5), the fitting procedure (Section 2.6), quality metrics (Section 2.7), LLM visual inspection (Section 2.8), and stored output (Section 2.9).

2.3. *Equivalent Width Definition*

The equivalent width is a classical concept in stellar spectroscopy, measuring integrated line absorption relative to the continuum:

$$W_\lambda = \int \frac{F_c(\lambda) - F(\lambda)}{F_c(\lambda)} \, d\lambda \tag{1}$$

where $F_c(\lambda)$ is the continuum flux and $F(\lambda)$ is the observed flux. For a fitted line profile $\phi(\lambda)$ normalized such that the continuum equals unity:

$$W_\lambda = \int [1 - \phi(\lambda)] \, d\lambda \tag{2}$$

Note that this definition does not inherently require pre-normalized spectra. Continuum normalization—with its myriad tools, subjective choices, and operator-dependent parameters—is often what makes EW measurements difficult to reproduce between researchers. Egent sidesteps this entirely. At high spectral resolution, the local continuum spanning a few angstroms can be well approximated by a low-order polynomial, often just linear. By fitting the continuum simultaneously with the line profile within each local window, we avoid the systematic errors and reproducibility issues introduced by global continuum normalization.

2.4. *The Voigt Profile*

Spectral lines are broadened by multiple physical mechanisms, making the Voigt profile—the convolution

Local continuum fitting avoids a separate global-normalization step.

of a Gaussian (thermal broadening) and a Lorentzian (pressure broadening)—the appropriate model for stellar absorption lines:

$$V(\lambda; \lambda_0, A, \sigma, \gamma) = A \cdot \frac{\mathrm{Re}[w(z)]}{\sigma\sqrt{2\pi}} \tag{3}$$

where $w(z)$ is the Faddeeva function and $z = \frac{(\lambda-\lambda_0)+i\gamma}{\sigma\sqrt{2}}$. The Gaussian width σ captures thermal broadening (from the Maxwell-Boltzmann velocity distribution of absorbing atoms) and instrumental broadening (the spectrograph's line spread function). The Lorentzian width γ captures natural broadening (from the finite lifetime of the excited state) and pressure broadening (collisions with other atoms). Rotational broadening can also be approximated as Gaussian for the moderate resolutions and slowly-rotating stars considered here; for rapid rotators at very high resolution, explicit rotational kernels would be a straightforward addition. We evaluate the Voigt profile using `scipy.special.voigt_profile`.

2.5. *Multi-Component Fitting*

A limitation of traditional manual EW measurement is fitting one line at a time. An experienced spectroscopist may identify obvious blends and fit them together, but this is labor-intensive—and even careful human analysis rarely accounts for all lines within the fitting window, each of which affects the derived continuum level. Ignoring these neighbors leads to systematic errors, particularly in crowded spectral regions.

Egent addresses this by fitting all lines in the extraction window simultaneously. The number of components N is initially determined from the input line catalog: all catalog entries within the extraction window are included. During LLM refinement, the agent can propose additional peaks if it identifies unmodeled absorption features in the residuals (Fig. 5). The model is:

$$F_{\mathrm{model}}(\lambda) = C(\lambda) \cdot \left[1 - \sum_{i=1}^{N} V_i(\lambda; \lambda_{0,i}, A_i, \sigma_i, \gamma_i)\right] \tag{4}$$

where $C(\lambda)$ is the continuum and the sum runs over all N lines in the window. Each Voigt profile V_i has four parameters: center wavelength $\lambda_{0,i}$, amplitude A_i, Gaussian width σ_i, and Lorentzian width γ_i.

The continuum $C(\lambda)$ is parameterized as a polynomial centered on the mean wavelength of the extraction window. By default, we use a linear continuum:

$$C(\lambda) = c_0 + c_1 \cdot (\lambda - \bar{\lambda}) \tag{5}$$

For regions with curved continuum (e.g., near the wing of a strong line), the LLM can request a higher-order polynomial:

$$C(\lambda) = \sum_{k=0}^{n} c_k \cdot (\lambda - \bar{\lambda})^k \tag{6}$$

with $n = 2$ (quadratic) or $n = 3$ (cubic). Higher orders are available but rarely needed and can lead to overfitting. The LLM makes this determination by inspecting residual patterns: a systematic curve in the residuals suggests the need for higher-order continuum.

The continuum level is determined from pixels identified as line-free. Initially, this is done automatically by iterative sigma-clipping: pixels more than 2.5σ below the running median are masked as absorption, and the continuum is fit to the remaining points. If the LLM identifies problems—for example, in a crowded region where no pixels are truly line-free—it can visually identify wavelength intervals that appear relatively absorption-free and propose these as continuum anchors (Fig. 6).

The total number of free parameters is $(n + 1)$ for the continuum plus $4N$ for the Voigt profiles. For a typical extraction window containing 3–5 lines with linear continuum, this corresponds to 14–22 parameters fitted to several hundred data points.

2.6. *Fitting Procedure*

The fitting procedure begins with initial parameter guesses: line centers from the input catalog, amplitudes estimated from the local flux minimum near each catalog position, and Gaussian and Lorentzian widths initialized to typical values appropriate for the spectrograph resolution ($\sigma = 0.04$ Å, $\gamma = 0.015$ Å).

Optimization is performed using the Levenberg-Marquardt algorithm via `scipy.optimize.curve_fit`. This provides both the best-fit parameters and their covariance matrix, from which we estimate parameter uncertainties. After fitting, components whose centers converge to within 0.2 Å (≈1.8 resolution elements at $R = 50{,}000$ near 5500 Å) are merged, as they cannot be meaningfully resolved and their parameters become degenerate. The covariance matrix from `curve_fit` is stored in the output JSON for every accepted fit, enabling users to assess parameter correlations and fit stability post hoc. If the fit fails to converge, we attempt restarts with perturbed initial conditions before flagging the line.

After fitting, we validate that no two Voigt profile centers have converged to within 0.2 Å of each other—a situation indicating that two components are modeling the same absorption feature (overfitting). If such pairs are found, we remove the component farther from the target wavelength and refit with fewer peaks. This ensures each absorption feature is modeled by a single Voigt profile, preventing double-counting of equivalent width.

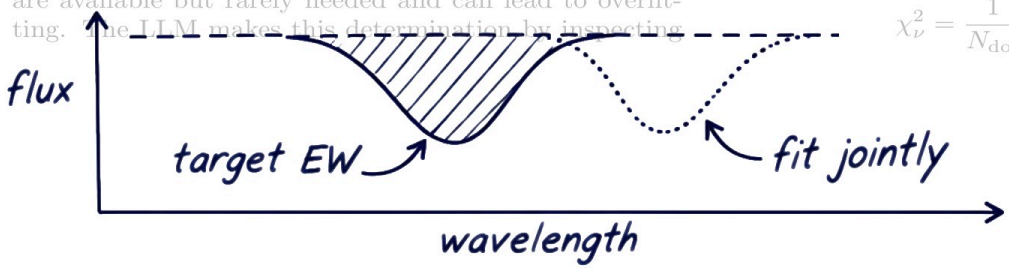
The equivalent width is then computed by numerical integration of the target line's Voigt profile. When multiple lines are fitted simultaneously, we report the EW for the component whose center is closest to the catalog wavelength.

2.7. *Quality Metrics*

We compute several diagnostics to assess fit quality and decide whether LLM review is needed.

The reduced chi-squared quantifies overall goodness of fit:

$$\chi^2_\nu = \frac{1}{N_{\mathrm{dof}}} \sum_i \frac{(F_i - F_{\mathrm{model,i}})^2}{\sigma_i^2} \tag{7}$$



Fit the neighbours together, but report the target component's area.

all neighbours in this window

Extra continuum flexibility can also absorb real line wings.

Check where the residuals sit, not just how large they are.

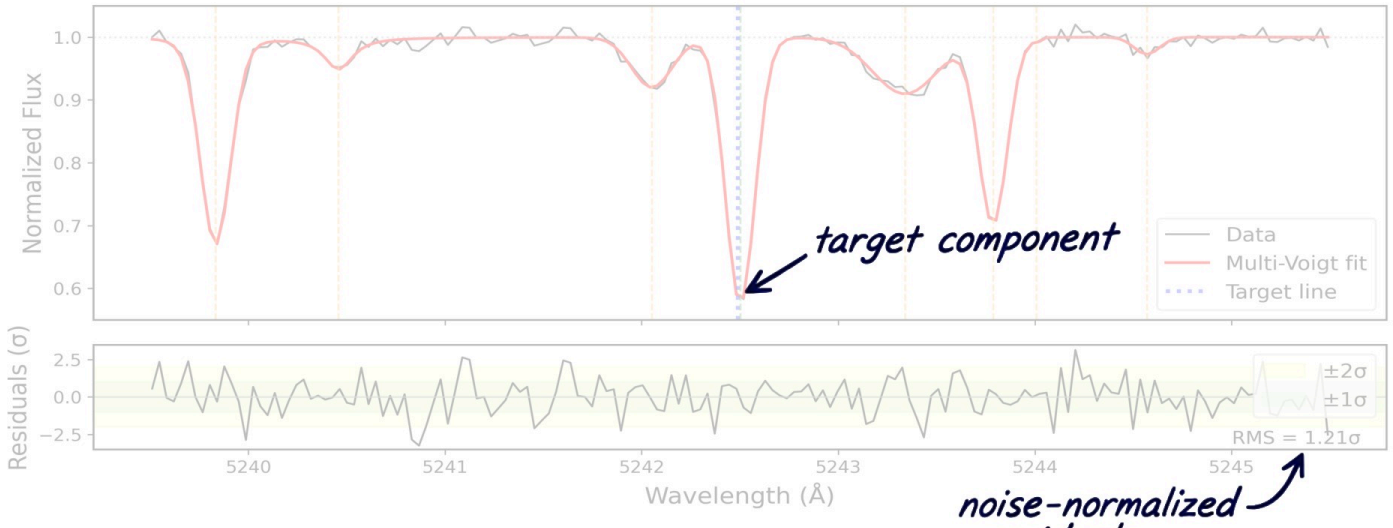


Figure 3. Multi-Voigt fitting process demonstrated on Fe I 5242.49 Å (Gaia ID 5578084051306737[illegible]). *Top panel:* The normalized spectrum (black) is fit simultaneously with 8 Voigt components (orange and green dashed lines indicate individual component centers; the target component is shown in green). The combined multi-Voigt model (red) reproduces the complex blend structure. Continuum normalization is performed iteratively alongside the Voigt fitting, allowing the model to adapt to local continuum curvature. *Bottom panel:* Normalized residuals (data minus model, divided by flux uncertainties) demonstrate fit quality. The shaded regions indicate $\pm1\sigma$ (green) and $\pm2\sigma$ (yellow); the RMS of 1.21σ confirms a statistically acceptable fit. When the linear fit to residuals shows a slope >0.3 σ/Å, a red dashed trend line is overlaid to guide the LLM's attention to continuum issues. This diagnostic panel format is also used for LLM visual inspection, where the agent examines residual patterns to detect missed blends (W-shaped residuals), continuum issues (systematic slopes indicated by the trend line), or unreliable fits (excessive scatter near the target line). Residuals are shown in σ units throughout; to convert to absolute normalized-flux units, multiply by the per-pixel uncertainty $\sigma_{\rm pixel} \approx 1/\mathrm{SNR}$ (e.g., for this spectrum at SNR $\sim$ 160, 1σ corresponds to $\sim$0.006 in normalized flux).

where $N_{\rm dof}$ is the number of data points minus the number of free parameters. Values near unity indicate a good fit; larger values suggest the model is inadequate.

The residual RMS, normalized by flux uncertainties, provides a more intuitive measure:

$$\mathrm{RMS} = \sqrt{\frac{1}{N}\sum_i \left(\frac{F_i - F_{\rm model,i}}{\sigma_i}\right)^2} \tag{8}$$

where σ_i is the flux uncertainty at each pixel. For a statistically good fit, this should be approximately unity; values above 1.5–2 indicate systematic deviations beyond what noise alone would produce.

We also compute a linear fit to the normalized residuals, yielding a residual slope in units of σ/Å. Slopes exceeding $\sim$0.5 σ/Å indicate a systematic continuum problem that requires attention. When the slope exceeds 0.3 σ/Å, we overlay the linear trend on the diagnostic plot (Figure 3b), giving the LLM a direct visual cue for continuum issues. We additionally check for correlated residuals (which might indicate a missed blend or continuum curvature) and the offset between the fitted and catalog wavelengths (which might indicate a misidentified line).

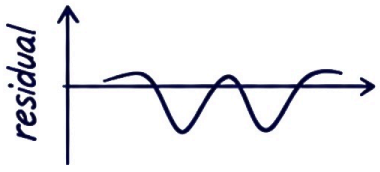
The pipeline classifies fits into three tiers. Fits with RMS $\lesssim 2\sigma$ and no quality issues are classified as *good* and accepted directly. Fits with moderate diagnostics (e.g., RMS $\sim$ 2–3σ, or 5–9 lines in the window) are classified as *acceptable* and also accepted without LLM intervention. LLM visual inspection is triggered when any of the following conditions are met: overall RMS $> 3\sigma$, $\chi^2_\nu > 5$, |residual slope| > 0.5 σ/Å, $\geq$5 consecutive residual points exceeding 2σ, fitted center offset $>$ 0.3 Å from catalog, central RMS $> 2.5\sigma$ (within ±1.5 Å of the target), or $\geq$10 lines in the extraction window. These thresholds can be relaxed (or tightened) to send more (fewer) lines to the LLM at higher (lower) cost. A complete listing of all configuration parameters is provided in Table 1.

2.8. *LLM Visual Inspection*

When quality metrics trigger inspection, the LLM receives a diagnostic plot showing the data, fitted model, individual Voigt components, and residuals (Figure 3b). The plot is annotated with the target wavelength and quality metrics. The LLM examines the plot as an expert spectroscopist would, looking for specific problems:

- W-shaped residuals centered on a wavelength offset from the target indicate a missed blend. The LLM can propose adding a Voigt component at the approximate wavelength.
- Tilted or systematically biased residuals across the window indicate a continuum problem. The LLM can adjust the extraction window or switch to a polynomial continuum.
- Large scatter in the residuals near the target wavelength, or an offset fitted center, may indicate an unreliable measurement. The LLM can flag the line as unreliable.

RMS near one assumes trustworthy pixel uncertainties.



Inspect, adjust, refit—and save the decision trail.

The LLM communicates these actions through a set of function-calling tools (Appendix B) that mirror the operations a human spectroscopist would perform: adjusting the wavelength window, adding blend components, changing continuum models, or flagging problematic lines.

After proposing an action, the pipeline executes it, refits, and returns a new diagnostic plot. The LLM visually assesses whether the fit has improved and either accepts the measurement or proposes further refinement. This iteration continues until the LLM accepts the fit or a maximum iteration count (default 10) is reached.

The reference EW stays hidden from the agent.

The LLM operates as a blind analysis: it has no access to the catalog EW values. All decisions are based solely on the fit diagnostics and visual appearance of the plot. This prevents the LLM from adjusting parameters to match expected values.

Below we show several representative examples of LLM reasoning on typical EW measurement cases (see Box 1 for verbatim outputs from GPT-5-mini):

Case 1: Blend Detection (Ca I 6696.02 Å)

"The residuals show a pronounced W-shaped pattern centered at the target (a positive excursion just blueward of 6696.02 Å and negative residuals at/just redward), indicating a missed blend very close to the target. The red fit looks symmetric while the data shows asymmetry (suggesting two close components). I will add an extra Voigt component to deblend the pair at the blue-side shoulder near 6695.9 Å."

After adding peak: "The residuals no longer show a W-shape at the target and the central RMS is 1.64σ, which is within the acceptable range."

Case 2: Window + Polynomial Continuum (Fe I 6261.10 Å)

"The fit shows very high normalized RMS (7.07) and many correlated residuals; the residuals show a strong systematic tilt across the target region indicating the continuum fit is wrong. I will try a polynomial continuum to better model curvature."

After polynomial + narrower window: "After adding nearby blends, switching to a 2nd-order polynomial continuum, and narrowing the window to 2.0 Å, the residuals near the target are flat with no W-shaped patterns and the central RMS is 1.23σ. I will accept this measurement."

Case 3: Manual Continuum Regions — Flagged (Si I 5665.55 Å)

"The region is very crowded with multiple absorptions. Automatic continuum fitting showed a large slope and correlated residuals. I will force the continuum using manually selected line-free regions that appear flat in the plot: ∼5663.25–5663.60 Å (left side), ∼5664.70–5665.05 Å (just blue of target), and ∼5668.20–5668.50 Å (right side)."

Final assessment: "Despite adding peaks and manual continuum regions, the fitter still reports very high RMS. The surrounding blends and systematic residuals make the EW unreliable. I will flag as severe_blend."

The approach balances thoroughness against computational cost: direct acceptance for clearly good fits (saving LLM costs), and iterative refinement with visual inspection for borderline cases where human-like judgment is needed.

2.9. *Stored Output*

Every measurement is stored as a JSON record with complete provenance, enabling exact reconstruction of the fit without re-running the LLM. Each record includes:

- **Measurements:** Final EW and uncertainty, direct (pre-LLM) EW for comparison, quality classification, and difference from catalog if available
- **Voigt parameters:** For every fitted line in the window—center wavelength, amplitude, σ, γ, FWHM, and individual EW—enabling exact model reconstruction
- **Continuum:** Method used (iterative linear, polynomial, or manual regions), polynomial order, and fitted coefficients
- **Region:** Échelle order, wavelength range, extraction window size
- **Diagnostics:** Reduced χ^2, normalized RMS, quality issues and warnings, number of fitted lines
- **LLM provenance:** Boolean indicating whether LLM was triggered, complete conversation transcript with reasoning text before each action, iteration history showing how EW evolved across refinements
- **Metadata:** Processing time, path to diagnostic plot, method used (direct/LLM), flag status and reason if applicable

Ten iterations cap an unresolved case; persistent trouble becomes a flag.

This level of detail serves multiple purposes: it allows verification that specific measurements are reliable, enables post-hoc replotting without re-fitting, provides training data for future fine-tuned models, and establishes a standard for reproducibility that manual measurements cannot match. The explicit reasoning text (Box 1) documents *why* the LLM made each decision, not just what it did.



A narrower window and curved continuum improve this fit together.

narrower + quadratic

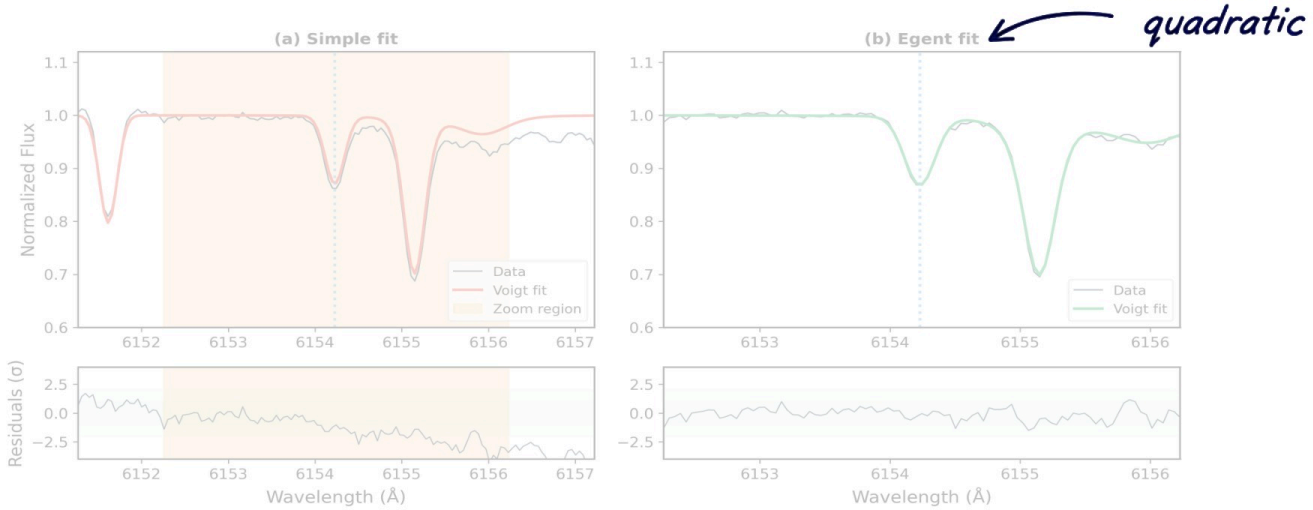


Figure 4. Example of LLM-driven window and continuum adjustment for Na I 6154.23 Å. **(a) Simple fit:** The initial 6 Å window includes edge features (visible at 6156–6157 Å) that introduce systematic bias in the residuals. The yellow shaded region indicates where the LLM determined a narrower window would improve the fit. **(b) Egent fit:** After the agent narrowed the window to 3.9 Å and switched to a quadratic continuum, the fit focuses on the region containing the target lines. The residuals improve: systematic deviations outside the $\pm 1\sigma$ band (green shading) largely disappear. This case illustrates the agent's ability to reason about fitting strategy: by excluding edge features irrelevant to the target line and adjusting the continuum model, the measurement becomes more robust. The EW error decreased from 8.9 mÅ to 5.8 mÅ (35% improvement relative to the C3PO catalog value of 36.8 mÅ).

The catalog deviation falls from 8.9 to 5.8 mÅ; this is not a fitted uncertainty.

Table 1

Default configuration parameters for the Egent pipeline. Left columns list extraction and fitting parameters that control the Voigt profile model and continuum estimation. Right columns list quality-control thresholds that determine whether a fit is accepted directly or sent for LLM visual inspection, along with pipeline settings. Fits not meeting any LLM trigger condition are accepted directly. All values can be modified by users; see the code repository for details.

half-width; this is a 6 Å total default window.

Extraction & Fitting		Quality & LLM Triggers	
Window half-width	3.0 Å	LLM trigger: high RMS	Overall RMS $> 3.0\sigma$
Continuum method	Iterative linear	LLM trigger: high χ^2	$\chi^2_\nu > 5$
Polynomial order	1 (LLM can raise)	LLM trigger: central RMS	Central RMS $> 2.5\sigma$
Sigma-clip threshold	2.5σ	LLM trigger: residual slope	$\lvert\text{slope}\rvert > 0.5$ σ/Å
Top percentile (init)	85%	LLM trigger: correlated resid.	$\geq$5 consecutive $> 2\sigma$
Min. peak height	0.02	LLM trigger: centre offset	> 0.3 Å from catalog
Min. peak prominence	0.015	LLM trigger: crowded region	$\geq$10 lines in window
σ init / bounds	0.04 / [0.003, 0.25] Å	Trend line overlay	slope > 0.3 σ/Å
γ init / bounds	0.015 / [0.001, 0.08] Å	Max LLM iterations	10 per line
Centre bounds	± 0.5 Å from catalog	Max LLM time	300 s per line
Centre merge threshold	0.2 Å	Flag retry attempts	3
Max fit evaluations	20,000	**Pipeline**	
		Default model (API)	GPT-5-mini
		Default model (local)	Qwen3-VL-8B-4bit
		Parallel workers	10 (API) / 1 (local)

3. VALIDATION DATA

Validating autonomous EW measurement requires a homogeneous dataset with expert measurements performed by a single analyst using consistent methodology. Synthetic spectra, while offering known ground truth, lack realistic complications such as instrumental blaze functions, telluric contamination, and the subtle blending patterns present in real stellar atmospheres. Conversely, measurements compiled from multiple sources suffer from analyst-to-analyst and methodology-to-methodology variations that confound meaningful comparison. The ideal benchmark combines real observational data with careful, single-analyst measurements—a combination that is rare given the labor-intensive nature of manual EW work.

We validate Egent against the C3PO (Complete Census of Co-moving Pairs Of stars) program (Yong et al. 2023; Liu et al. 2024; Sun et al. 2025b,c; Yu et al. 2025), a systematic study of 125 co-moving stellar pairs designed to probe chemical homogeneity in comoving systems and detect signatures of planet engulfment. The program obtained high-resolution spectra using three instruments: Magellan/MIKE (78 pairs, $R \sim 50{,}000$), Keck/HIRES (23 pairs, $R \sim 72{,}000$), and VLT/UVES (25 pairs, $R \sim 110{,}000$). D. Yong, with about 20 years of experience in precision abundance work, performed all EW measurements using the line-by-line differential technique that achieves relative abundance precision of

Real spectra provide a benchmark, not an unquestionable truth.

~0.01 dex. Each spectrum includes measurements for approximately 200 carefully selected lines spanning Fe I, Fe II, and 20+ other species. The C3PO measurements have enabled detection of chemical abundance differences at the 0.01 dex level between co-moving pairs, demonstrating their quality for precision differential analysis.

We analyze Magellan/MIKE spectra from the C3PO sample. For GPT-5-mini, we processed 63 spectra (13,913 lines); for GPT-5, we processed 21 spectra (4,702 lines). These numbers provide sufficient statistics to characterize both random scatter and systematic offsets. The spectra span signal-to-noise ratios of 50–250 per pixel near 5500 Å (median 174) and cover wavelengths from 3300–9400 Å, providing a representative range of data quality encountered in precision abundance studies.

The expert catalog also carries continuum-placement uncertainty.

We emphasize that expert measurements are not absolute ground truth. The "true" EW depends on subjective choices about continuum placement and blend handling that different experts would make differently. However, expert measurements from a single experienced analyst provide the fairest benchmark for real stellar spectra, as they incorporate the same challenges—blending, continuum ambiguity, noise—that any automated system must handle. An expert's years of experience inform subtle judgments about which blends are problematic and which continuum regions are trustworthy, judgments that idealized synthetic tests cannot capture.[3]

Egent assumes input spectra are in the stellar rest frame; it does not perform its own wavelength calibration. Spectra are shifted to the stellar rest frame using barycentric corrections computed from observation timestamps and observatory coordinates. We apply an additional empirical wavelength correction using ~10 strong, isolated Fe I lines per spectrum, fitting observed positions against laboratory wavelengths to correct residual offsets (~0.01–0.05 Å) arising from spectrograph flexure and reduction pipeline limitations. For the benchmark sample, we selected the subset of spectra for which this two-step correction (barycentric + Fe I empirical shift) places absorption lines at their laboratory wavelengths to within the fitting tolerance ($\lesssim$0.3 Å, the centre-bound half-width in Table 1). This selection does not bias the EW comparison, as EW is an integrated quantity insensitive to small wavelength offsets provided the line falls within the extraction window.

4. RESULTS

We present results in three parts. First, we examine the impact of LLM intervention: how often it triggers, what tools it uses, building intuition for how it improves fits. Second, we compare Egent measurements against the C3PO expert catalog, examining both aggregate statistics and individual case studies (Section 4.2). Third, we assess robustness across SNR using the natural variation in our sample (SNR~50–250), finding that performance remains stable with only modest degradation at the lowest SNR (Section 4.3).

[3] D. Yong has since transitioned out of active research.

Table 2
Pipeline statistics for GPT-5 and GPT-5-mini runs.

Metric	GPT-5	GPT-5-mini
Spectra processed	21	63
Total lines	4,702	13,913
Direct fit accepted	21.4%	26.3%
LLM triggered & accepted	59.4%	62.5%
Flagged	19.1%	11.0%
Final accepted	80.9%	88.8%

different rejection rates

4.1. *Impact of LLM Corrections*

Table 2 summarizes the pipeline behavior across GPT-5 and GPT-5-mini runs. Of the 4,702 lines processed with GPT-5 (21 spectra), 21.4% were accepted on the direct fit without LLM intervention, while 59.4% triggered LLM visual inspection and were subsequently accepted. For GPT-5-mini (63 spectra passing the $R^2 \geq 0.8$ correlation filter described below, totalling 13,913 lines), the pattern is similar: 26.3% direct acceptance, 62.5% LLM-triggered. The flagging rates show GPT-5 flagged 19.1% of lines while GPT-5-mini flagged 11.0%. The final acceptance rates are 80.9% for GPT-5 and 88.8% for GPT-5-mini.

Both models follow a similar workflow: inspect the diagnostic plot, assess quality, and either accept or intervene. Table 3 shows the usage frequency of each intervention tool. The key tools are `set_continuum_method` (switching to polynomial continuum), `extract_region` (adjusting the wavelength window), `set_continuum_regions` (manually specifying line-free pixels), and `fit_ew` with `additional_peaks` (adding blend components). GPT-5-mini uses `set_continuum_regions` more frequently (7% of LLM lines vs 1% for GPT-5), while GPT-5 uses `set_continuum_method` and `extract_region` more aggressively. These differences may reflect model-specific reasoning strategies, though both achieve comparable final accuracy.

Table 3
Intervention tool usage (percentage of LLM-triggered lines). Percentages sum to >100% because multiple tools are often used per line. The higher rates for GPT-5 partly reflect an earlier code version where the LLM explicitly called setup tools.

Tool	GPT-5	GPT-5-mini
`set_continuum_method`	49%	28%
`extract_region`	84%	24%
`fit_ew` (re-fit)	71%	32%
`set_continuum_regions`	1%	7%

Different samples and code versions complicate tool-use comparisons.

Figures 4–6 illustrate the three main types of LLM intervention.

Window and continuum adjustment (Figure 4): For Na I 6154.23 Å, the initial 6 Å window included strong features at the edges (6156–6157 Å) that pulled the linear continuum fit away from the true local continuum. The LLM recognized this pattern in the residuals—

Two concrete repairs: add missed absorbers, or choose continuum anchors.

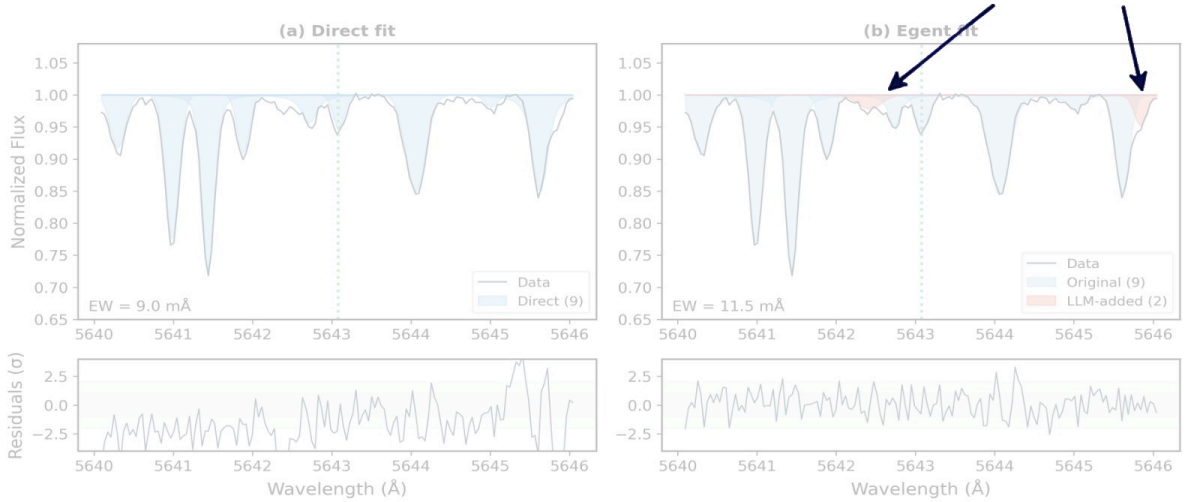


Figure 5. Example of LLM-driven peak identification for Ni I 5643.08 Å. **(a) Direct fit:** Nine Voigt components (blue shaded regions) were identified automatically from the line catalog. While these capture most absorption features, residuals at ∼5642.5 Å and ∼5645.8 Å show systematic deviations exceeding 2σ, indicating missed features. **(b) Egent fit:** By visually inspecting the diagnostic plot, the LLM identified two small absorption features not in the catalog and explicitly specified their wavelengths (5642.45 and 5645.85 Å) for inclusion in the multi-Voigt fit. The resulting 11-component model (9 original in blue, 2 LLM-identified in red) achieves flatter residuals. This demonstrates the LLM's ability to function as a visual inspector: it examines the residual panel for patterns indicating unmodeled absorption and proposes specific wavelengths for additional Voigt components. The EW improved from 9.0 mÅ to 11.5 mÅ, reducing the error relative to the catalog value (13.8 mÅ) by 52%.

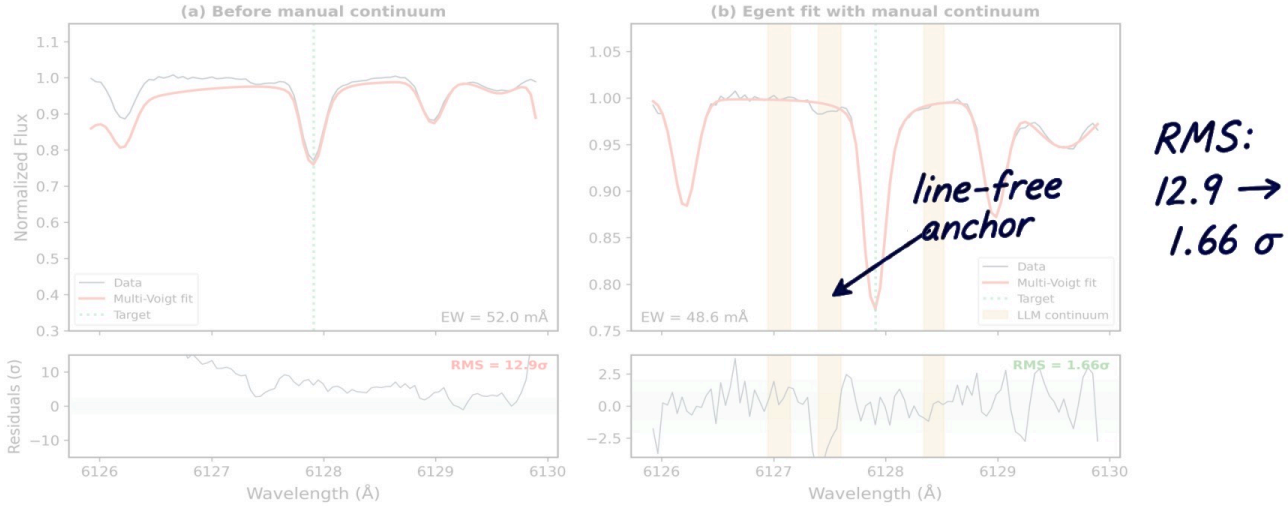


Figure 6. Before and after comparison of LLM-proposed continuum regions for Fe I 6127.91 Å. **(a) Before:** Automatic iterative continuum fitting struggles in this crowded region, resulting in a systematic offset between data and model (RMS = 12.9σ). The normalized flux appears above unity on the blue side while absorption features are incompletely modeled. **(b) After:** Based on visual inspection, the LLM identified three wavelength intervals (orange shaded regions: 6126.95–6127.15 Å, 6127.40–6127.60 Å, and 6128.34–6128.52 Å) as relatively line-free and proposed these as continuum anchors. With these constraints, the fitting procedure achieves improved residuals (RMS = 1.66σ). This approach mimics how a human expert manually identifies clean continuum pixels when automatic methods fail in crowded spectral regions.

systematic deviations at the window edges—and responded by (1) narrowing the extraction window to 3.9 Å to exclude the edge contamination, and (2) switching to a quadratic continuum to better model the local curvature. These two adjustments reduced the EW discrepancy from 8.9 mÅ to 5.8 mÅ (35% improvement).

Blend identification (Figure 5): For Ni I 5643.08 Å, the direct fit included nine Voigt components from the catalog, but residuals at ∼5642.5 Å and ∼5645.8 Å showed systematic W-shaped patterns—the telltale signature of unmodeled absorption. The LLM identified these patterns visually and proposed two additional Voigt components at specific wavelengths (5642.45 and 5645.85 Å). The 11-component fit achieved flatter residuals. This demonstrates the LLM's ability to detect blends.

Manual continuum regions (Figure 6): For Fe I 6127.91 Å in a crowded spectral region, automatic iter-

Inspection often confirms a fit; a smaller subset needs repair or rejection.

Most inspected mini-model lines change by less than 1 mÅ.

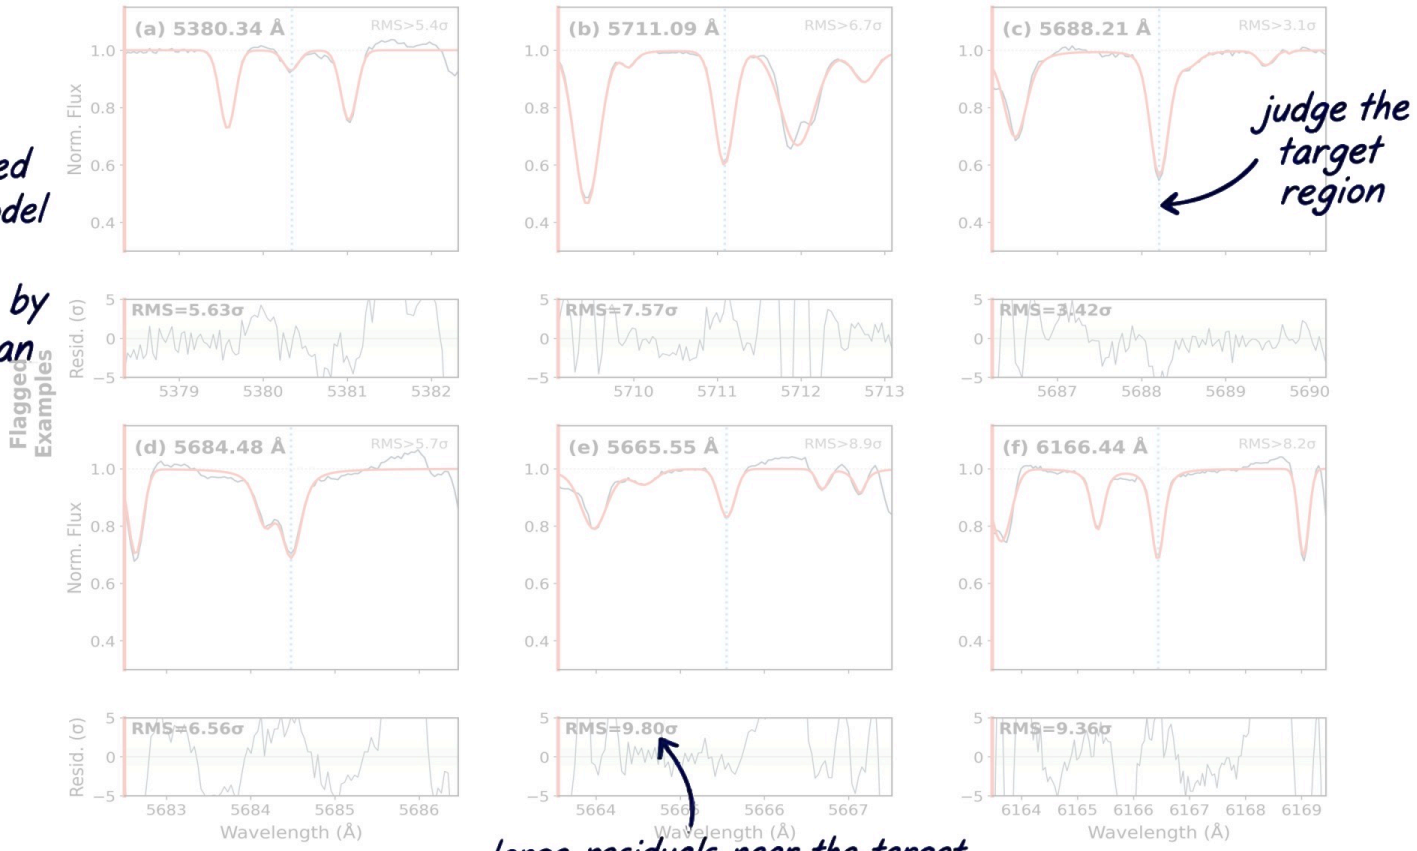


Figure 7. Representative examples of flagged lines. Lines are flagged by the LLM during iterative fitting when it determines the fit cannot be salvaged—elevated RMS values indicate continuum failures, severe blends, or model-data mismatch. The flagging mechanism identifies cases where automated fitting produces unreliable measurements, preventing these outliers from contaminating abundance analyses. All flagging uses AI visual inspection with no human intervention.

ative continuum fitting failed entirely (RMS = 12.9σ). No amount of sigma-clipping could identify true continuum pixels because absorption features dominated the entire window. The LLM recognized this failure mode and switched to manual continuum specification, visually identifying three narrow wavelength intervals (6126.95–6127.15 Å, 6127.40–6127.60 Å, and 6128.34–6128.52 Å) that appeared relatively line-free. Using only these regions as continuum anchors, the fit improved to RMS = 1.66σ. This tool is invoked in 1–7% of LLM-triggered lines (Table 3) but is essential for the most challenging spectral regions.

While Figures 4–6 demonstrate improvements in specific edge cases, statistical analysis reveals the LLM's primary role is quality control. For GPT-5-mini, 80% of LLM-triggered lines show final EW differing from the direct fit by <1 mÅ—the LLM examines the fit and confirms it is acceptable without modification. GPT-5 is more interventionist: only 59% of its LLM-triggered lines have <1 mÅ changes, with 32% showing >2 mÅ shifts (vs. 16% for GPT-5-mini). This aligns with GPT-5's higher flagging rate and slightly better agreement with catalog values (Figure 8). The LLM's value lies in (1) confirming that the majority of direct fits are reliable without human inspection, (2) flagging the 11–19% of fits that cannot be salvaged, and (3) rescuing edge cases like those shown in Figures 4–6 that would otherwise produce gross errors.

The LLM's value is most apparent in difficult cases. For the 1–7% of lines requiring manual continuum specification (`set_continuum_regions`), automatic fitting fails entirely in crowded regions. For catastrophic fits with RMS >5σ, missed blends, or wrong continuum, LLM intervention can reduce errors from tens of mÅ to a few mÅ (see Figures 4–6 for examples). These rescues represent only a few percent of all lines, but without them such cases would contribute catastrophic outliers to abundance analyses.

To quantify what the LLM adds beyond deterministic diagnostics, we compare the full pipeline against a direct-fit-only ablation. Without the LLM, applying the same RMS > 3σ flagging threshold to the direct fit accepts 10,134 lines at MAD = 5.4 mÅ; with the LLM, 12,359 lines are accepted at MAD = 5.8 mÅ. The 22% gain in completeness comes from 2,424 lines that the LLM rescues: their initial direct-fit RMS exceeded 3σ, but the LLM improved the fit (via continuum adjustment, window narrowing, or blend addition) to bring residuals below threshold. These rescued lines have MAD = 7.0 mÅ—worse than average but scientif-

Inspection accepts about 22% more lines, with MAD rising from 5.4 to 5.8 mÅ.

A bad residual pattern and a large catalog difference are different tests.

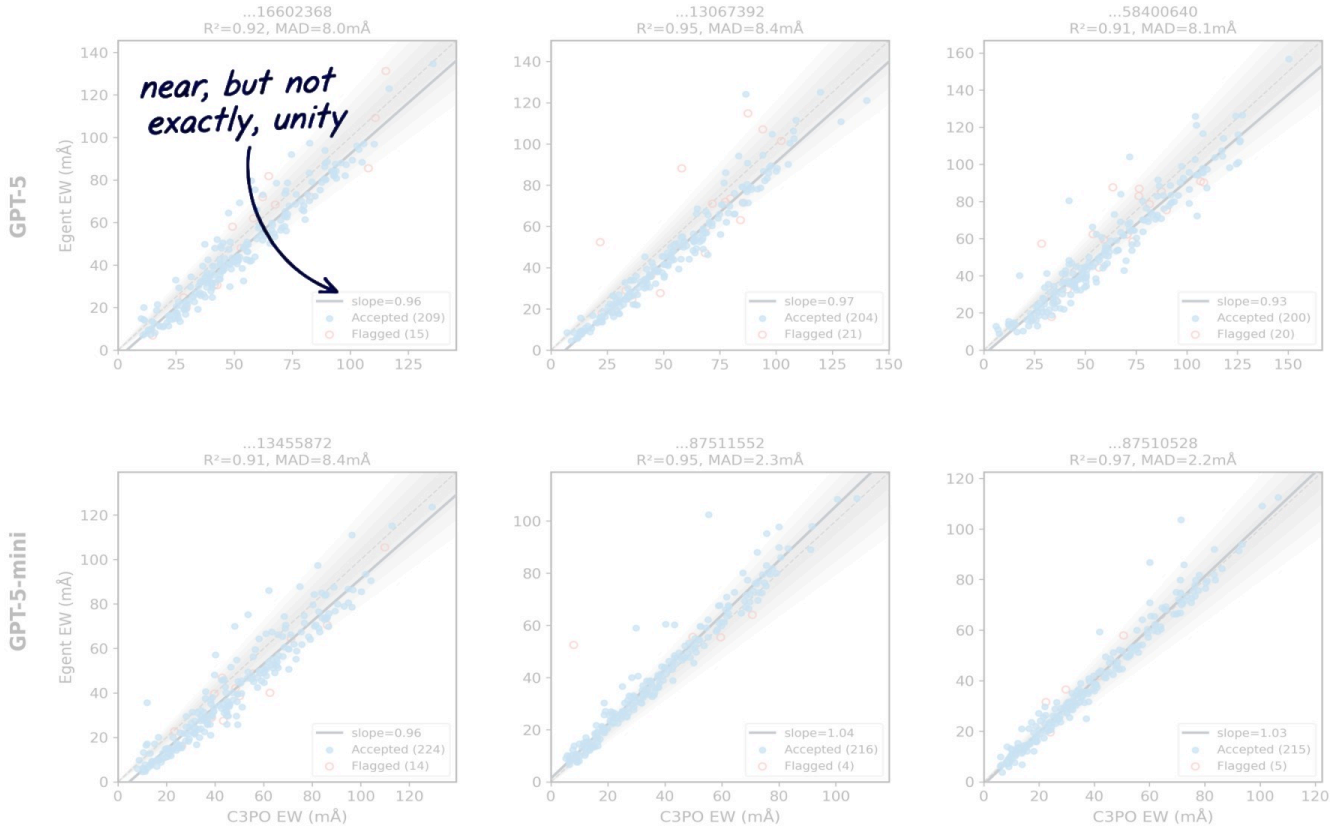


Figure 8. 1-to-1 comparison of Egent vs C3PO EW measurements. Top row: GPT-5 results for three representative spectra. Bottom row: GPT-5-mini results. Blue filled circles: accepted measurements. Hollow red circles: lines flagged by the LLM. Dark line: best-fit linear regression on unflagged points only. Shaded bands: 5%, 10%, 15%, 20% deviation from 1:1. Flagged points cluster at larger deviations, validating the AI flagging mechanism. The per-spectrum slope variations (±5–10%) arise from C3PO's global pre-normalization step—a subjective choice that introduces systematic offsets. Egent avoids this by fitting directly from raw flux with local continuum estimation, eliminating global continuum as a source of systematic error. Both models achieve $R^2 > 0.9$ and MAD $\sim$ 4–6 mÅ.

ically usable. Of the 9,602 LLM-triggered inspections, 71% result in no EW change (<0.1 mÅ), confirming the direct fit without modification—the LLM's role in these cases is purely confirmatory. Routing only the top 5% of hardest cases (ranked by direct-fit RMS) to the LLM already captures most of the MAD improvement, suggesting a tiered approach can achieve similar quality at ∼80% cost savings. The LLM's contribution is thus primarily one of *completeness*—rescuing marginal fits that deterministic rules would discard—and *nuanced flagging* that accounts for spatial residual patterns beyond what scalar diagnostics capture. Since abundance determinations average over many lines per element, this 22% gain in accepted lines translates directly into improved abundance precision.

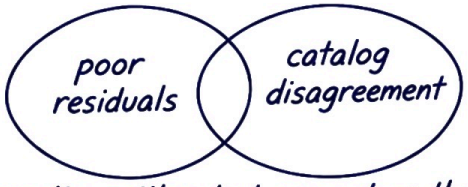
AI flagging (Figure 7): The pipeline employs fully automated quality control with no human intervention. During the iterative fitting process, the LLM can invoke the `flag_line` tool when it determines that a fit cannot be salvaged after multiple improvement attempts. Figure 7 shows representative cases where elevated RMS values indicate catastrophic failures (wrong continuum, severe blends, or persistent model-data mismatch). The LLM is explicitly instructed to focus on the target region (±0.5 Å of the target wavelength) and ignore edge effects that do not affect the measurement. This focused approach ensures thorough quality control without over-flagging lines with acceptable fits at the target.

Evaluated as a binary classifier using $|\Delta\text{EW}| > 10$ mÅ as the operational definition of a "poor" measurement, the flagging system achieves precision ≈ 46% and recall ≈ 11% for GPT-5-mini (23% for GPT-5). The false-negative rate among accepted lines is ∼28% at this threshold. These modest classifier metrics reflect the distinction between *fit quality* (what flagging targets, based on residual patterns) and *agreement with a reference catalog* (what |ΔEW| measures). A large |ΔEW| can arise from legitimate continuum methodology differences rather than a bad fit, and the ∼28% false-negative rate is consistent with the tail of a distribution whose median is 5–7 mÅ. The flagging system is designed to remove catastrophic failures, not to minimize disagreement with any particular reference.

4.2. *Comparison to Expert Measurements*

Figure 8 compares Egent measurements against C3PO values for representative Magellan/MIKE spectra, showing both GPT-5 (top row) and GPT-5-mini (bottom row)

Some accepted fits still disagree with the catalog; some flags land close by chance.



schematic: neither test guarantees the other

Strong correlation can coexist with a systematic scale offset.

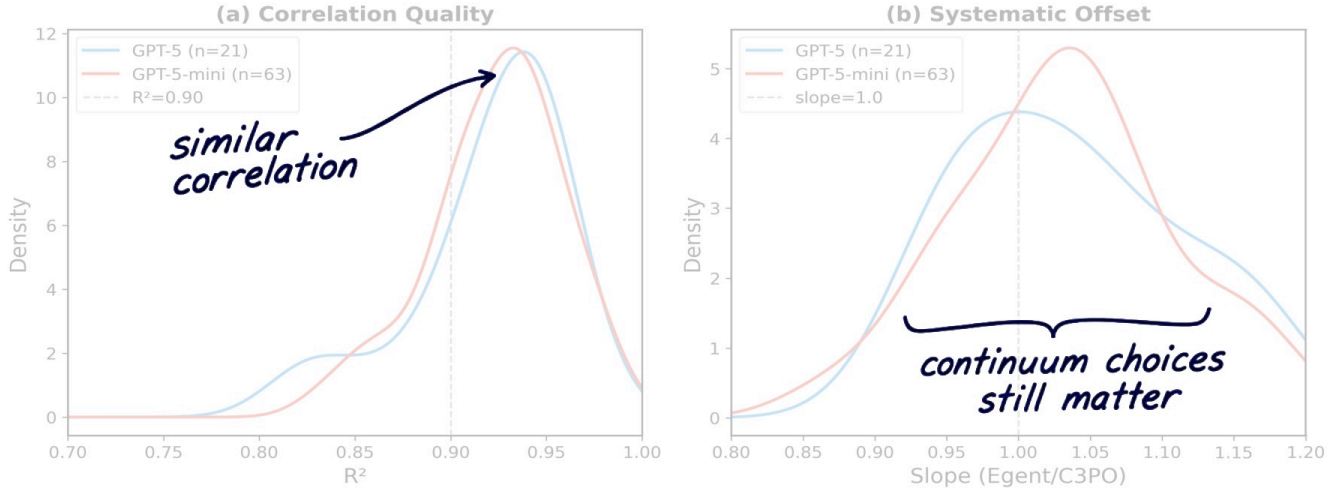


Figure 9. Distribution of correlation quality (R^2) and systematic offset (slope) across all spectra with C3PO catalog matches, comparing GPT-5 (21 spectra) and GPT-5-mini (63 spectra). A quality gate of $R^2 \geq 0.8$ is applied to the per-spectrum linear regression of Egent EW vs. C3PO EW (computed on unflagged lines with $\geq$20 matched lines); all spectra in the pre-selected sample pass this threshold (see text). (a) Both models achieve median $R^2 \approx 0.93$, with all spectra above 0.82, indicating strong linear correlation between Egent and expert measurements. (b) Slopes are centered near unity (median ≈ 1.03) with $\pm$15% spread, reflecting per-spectrum offsets from continuum methodology differences rather than fitting failures. The overlapping distributions demonstrate that GPT-5-mini achieves measurement quality comparable to GPT-5 at 5$\times$ lower cost.

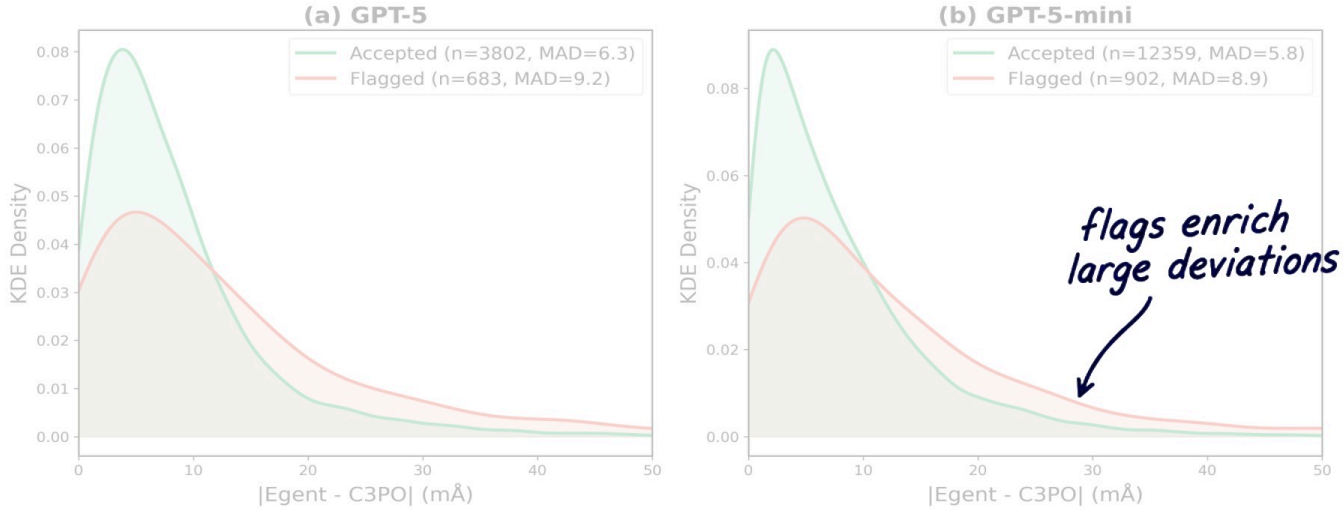


Figure 10. Distribution of absolute EW deviations |Egent − C3PO| separated by flagging status for (a) GPT-5 (21 spectra, 4,702 lines) and (b) GPT-5-mini (63 spectra, 13,913 lines). Green: accepted lines (MAD = 5–7 mÅ). Red: flagged lines (MAD = 10–20 mÅ). Both models show the same pattern: flagged lines have systematically larger deviations from catalog values, confirming that the AI flagging effectively identifies problematic fits.

results. For each spectrum, we compute a per-spectrum R^2 from the ordinary least-squares linear regression of Egent EW versus C3PO EW, using only unflagged (accepted) lines and requiring $\geq$20 matched lines per spectrum. Spectra with $R^2 < 0.8$ are excluded as a quality gate, since low R^2 indicates a systematic mismatch (e.g., from wavelength calibration errors or data reduction issues) rather than line-by-line scatter. In practice, all 21 GPT-5 spectra and all 63 GPT-5-mini spectra pass this threshold, as the input sample was pre-selected for reliable wavelength calibration before processing. Figure 9 summarizes the distribution across all spectra with C3PO matches: both models achieve $R^2 > 0.82$ (median ≈ 0.93), with slopes scattered around unity (median ≈ 1.03). Egent produces measurements with a raw median absolute deviation (MAD) of 5–7 mÅ ($\sim$11–12% of mean EW) against expert values, with no post-hoc per-spectrum correction applied—quality that traditionally requires substantial human effort.

Despite the strong correlation, Figure 8 reveals clear spectrum-to-spectrum slope offsets: GPT-5 panels show slopes of 0.93–0.97, while GPT-5-mini panels range from 0.96–1.04. These offsets arise from continuum methodology differences, not fitting failures—visual inspection of

These examples retain useful EWs even when the data are noisy.

different stars; not injected noise

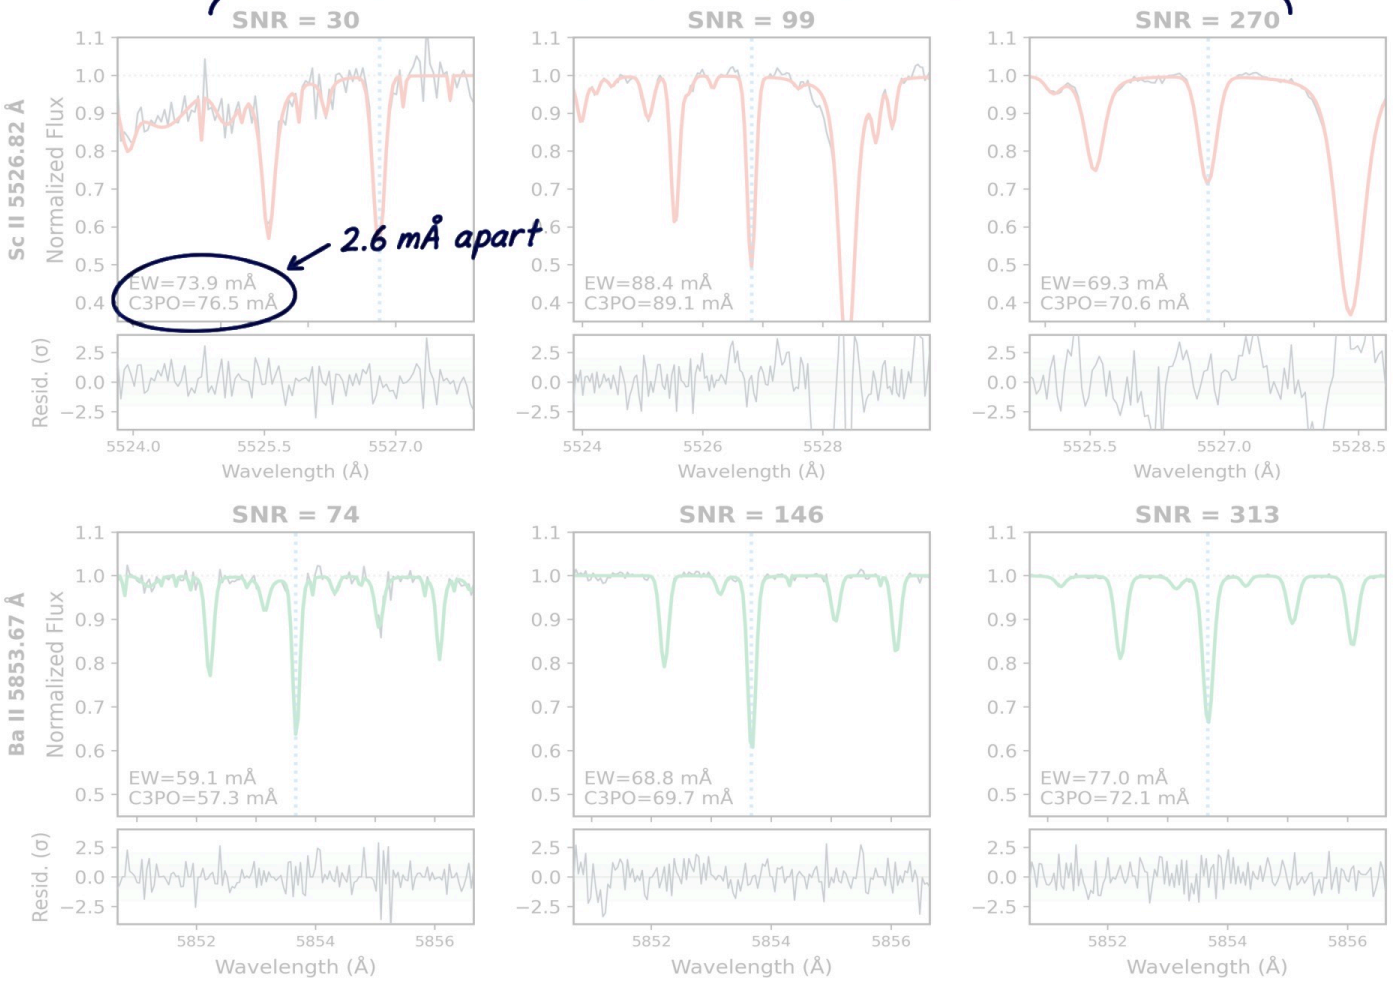


Figure 11. Fitting quality across SNR levels for two representative lines: Sc II 5526.82 Å (red, top) and Ba II 5853.67 Å (green, bottom). Each column shows the same line observed in different spectra at SNR = 30–313. Top panels: normalized flux (dark) with Voigt model overlay. Bottom panels: normalized residuals with $\pm1\sigma$ (green) and $\pm2\sigma$ (blue) bands. Even at SNR = 30, the measured EW agrees with C3PO to within 3 mÅ. Residuals scale appropriately with noise level without systematic structure.

individual fits (Appendix D) confirms that Egent's Voigt profiles trace the data precisely. In traditional manual analysis, the common practice is to first globally normalize the spectrum before local continuum adjustment—a step that facilitates human inspection but introduces systematic choices that propagate to the final EW. C3PO uses such pre-normalized spectra, while Egent fits directly from raw flux with no pre-normalization step. The slope distribution width (±15%) therefore reflects the inherent subjectivity of global continuum placement. In the linear regime of the curve of growth where EW $\propto N$, a slope deviation translates directly to abundance error via $\Delta[\mathrm{X/H}] = \log_{10}(\mathrm{slope})$—e.g., slopes of 0.93 and 1.04 correspond to −0.03 and +0.02 dex systematic offsets, comparable to or larger than random measurement precision. Our continuum-less approach explicitly avoids this source of systematic error.

Table 2 shows that GPT-5-mini achieves a higher acceptance rate (88.9% versus 80.9% for GPT-5), suggesting GPT-5 is more conservative in its flagging. Both models produce high-quality measurements on accepted lines, with comparable R^2 and MAD values. The choice between models depends on the desired balance between throughput and stringency: GPT-5-mini offers 5× cost savings with slightly higher acceptance rates, while GPT-5 may provide more conservative quality control. That modern AI, even a lightweight vision model, can achieve expert-level quality at ∼200 lines per dollar represents a fundamental shift in the cost-benefit calculus for line-by-line spectroscopic analysis.

Figure 10 demonstrates that AI flagging effectively identifies problematic measurements: accepted lines (green) cluster tightly at small deviations from C3PO, while flagged lines (red) show broader distributions extending to larger errors. In Figure 8, this manifests as hollow red markers clustering at larger deviations from the 1:1 line. However, some unflagged points also show large deviations, and some flagged lines have small deviations—visual inspection (Appendix D) reveals that occasional errors escape flagging, just as human mea-

Continuum offsets can still shift inferred abundances by hundredths of a dex.

A fitting error bar is only part of the uncertainty budget.

surements are not immune to error, while some valid flagging may catch fits that happen to land near catalog values by chance. One could impose more stringent flagging criteria, but this trades completeness for purity; the optimal threshold depends on the scientific use case. A sensitivity analysis confirms this trade-off is benign: varying the RMS flagging threshold from a strict 2.0σ to a permissive 5.0σ changes the acceptance rate from 72% to 91%, while the MAD shifts by only 0.2 mÅ (5.5 to 5.7 mÅ). The default 3.0σ threshold provides a reasonable balance, and users can adjust it with minimal impact on measurement accuracy. The key advantage is that AI flagging provides reproducible and homogeneous quality control—a form of systematic reasoning that can be version-controlled and applied uniformly across large datasets.

The reference catalog contributes to the observed scatter too.

An important practical question is whether the reported $\sigma_{\rm EW}$—propagated from the `curve_fit` covariance matrix—is well-calibrated. We assess this by comparing $\sigma_{\rm EW}$ against $|\Delta{\rm EW}|$ residuals with C3PO, noting at the outset that C3PO measurements are not ground truth: they carry their own uncertainties from manual continuum placement and line-profile judgments. The observed scatter therefore reflects contributions from both pipelines, and any miscalibration factors derived below are upper bounds on Egent's uncertainty alone.

With this caveat, we find the reported uncertainties underestimate the joint scatter by a factor of ∼2–3 overall: for GPT-5-mini, only 32% of deviations fall within $1\sigma_{\rm EW}$ (versus 68% expected for a Gaussian). The miscalibration is strongly regime-dependent. Lines with small reported $\sigma_{\rm EW} < 2$ mÅ show the largest discrepancy (ratio ≈ 3.6), because these are high-SNR clean lines where photon noise is well-captured but the dominant error source—continuum methodology differences between the two pipelines—is not. In contrast, lines with large reported $\sigma_{\rm EW} \geq 5$ mÅ are well-calibrated (ratio ≈ 1.0), because the fitting uncertainty itself is large enough to encompass the systematic floor. The reported $\sigma_{\rm EW}$ should therefore be interpreted as a lower bound capturing statistical fitting uncertainty only. For applications requiring conservative total uncertainties, one may add a systematic floor in quadrature: $\sigma_{\rm total} = \sqrt{\sigma_{\rm EW}^2 + \sigma_{\rm sys}^2}$, with $\sigma_{\rm sys} \leq 5$ mÅ as a safe upper bound—the true Egent systematic floor is likely smaller, since part of the observed scatter originates in the reference measurements.

To place these results in the context of established automated EW codes: ARES (Sousa et al. 2007, 2015) and DAOSPEC (Stetson and Pancino 2008) are the most widely used alternatives for batch EW measurement. A direct head-to-head comparison on the same spectra is not straightforward because both codes require pre-normalized input spectra, whereas Egent operates on raw flux; applying a global normalization to our data would introduce continuum choices that confound the comparison. We therefore compare against published scatter figures. On pre-normalized synthetic solar spectra ($R \sim 65{,}000$, SNR ~ 100), ARES achieves rms scatter < 2 mÅ against manual `splot` measurements, with a systematic offset below 1.5 mÅ (Sousa et al. 2007). On real FEROS spectra (SNR > 200, $R \sim 50{,}000$), the mean ARES–manual difference remains ∼1.5 mÅ. When comparing ARES v2's fully automatic continuum mode against manually tuned settings on the same code, Sousa et al. (2015) find dispersion ∼4.4 mÅ across ∼380 lines per spectrum on seven spectra spanning SNR 70–2000. For DAOSPEC, Stetson and Pancino (2008) report scatter of 3.7 mÅ against ARES on a high-SNR Solar HARPS spectrum ($R \sim 45{,}000$, SNR ~ 350, 98 lines), rising to 10.3 mÅ on moderate-SNR spectra of ω Centauri red giants ($R \sim 45{,}000$, SNR $\sim$ 100–150, 1150 lines across six stars). Egent's MAD of 5–7 mÅ against an independent reference pipeline is competitive with these benchmarks, particularly given that (1) it is measured against a separate pipeline rather than the same code with different settings, (2) it includes the full range of line strengths and spectral conditions without pre-selection, and (3) requires no pre-normalized spectra. Egent's direct-fit pathway—without LLM intervention—already achieves MAD = 5.4 mÅ on accepted lines, demonstrating that the core fitting engine performs at the level of established tools; the LLM adds completeness and provenance rather than raw precision. We emphasize that the comparison above relies on published scatter figures rather than a head-to-head run on identical inputs; a controlled benchmark—running ARES and DAOSPEC on the same C3PO spectra after a normalization step matched to each code's input requirements, with results stratified by line strength and line density—would give a more definitive answer on the relative value of Egent versus established automated EW pipelines, and is a natural future extension.

ARES and DAOSPEC were not run head-to-head on these same inputs.

4.3. *Performance Across SNR*

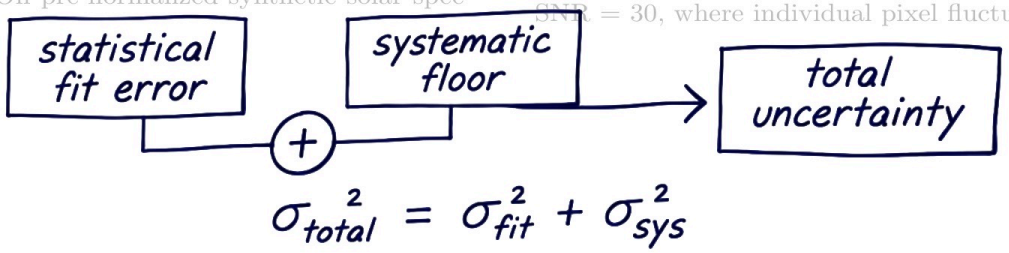
Rather than artificially degrading spectra, we assess SNR robustness using the natural variation in our GPT-5-mini sample. The 63 spectra span median SNR from 50 to 250 per pixel, providing a direct test of performance across data quality regimes. Table 4 shows that correlation quality ($R^2 \geq 0.90$, median 0.93) and precision (MAD ∼ 4–8 mÅ) remain stable across all SNR bins.

Table 4
Per-Spectrum Statistics by SNR (GPT-5-mini)

SNR Range	N Spectra	R^2 (median)	MAD (median)
<100	7	0.90	5.7 mÅ
100–150	12	0.95	3.5 mÅ
150–200	32	0.93	8.0 mÅ
>200	12	0.93	4.2 mÅ
All	63	0.93	5.7 mÅ

not a simple SNR trend

Figure 11 illustrates this robustness with two representative lines—Sc II 5526.82 Å and Ba II 5853.67 Å—observed at SNR ranging from 30 to 313. Even at SNR = 30, where individual pixel fluctuations reach 3%



In this sample, stronger lines show larger absolute catalog differences.

of the continuum, the Voigt fit successfully captures the line profile and yields EW within 3 mÅ of the expert value. The residuals naturally show larger scatter at low SNR, but remain within the expected noise envelope without systematic structure.

Weak lines still need enough signal-to-noise to be detected.

To disentangle the effects of SNR and line strength, Tables 5 and 6 present flagging rate and MAD as two-dimensional functions of both variables. The practical lower limit for reliable measurement depends on line strength: typical lines (EW $\sim$ 50–100 mÅ) remain robust down to SNR $\sim$ 30–50, while weaker lines (EW $<$ 20 mÅ) require higher SNR to distinguish from noise.

more strong-line flags

Table 5
Flagging Rate (%) by SNR and Catalog EW

	SNR					
EW (mÅ)	<100	100–150	150–200	200–300	>300	All
<30	5	3	4	7	5	5
30–60	7	5	7	7	6	7
60–100	5	9	9	9	8	8
100–150	12	9	10	12	11	11

Table 6
MAD (mÅ) by SNR and Catalog EW (unflagged lines)

	SNR					
EW (mÅ)	<100	100–150	150–200	200–300	>300	All
<30	5.2	3.7	3.7	3.7	4.0	3.9
30–60	7.0	5.0	5.3	5.7	6.6	5.9
60–100	8.8	7.3	8.2	7.5	9.1	8.1
100–150	11.6	11.1	10.6	9.7	10.1	10.4

increasing line strength

The flagging rate (Table 5) is remarkably uniform across SNR (5–12% in all bins), indicating that quality control responds to fit quality rather than raw noise level. The modest increase with line strength (5% for weak lines to 11% for strong lines) likely reflects higher blend probability in crowded spectral regions.

The MAD values (Table 6) tell a similar story: measurement precision is largely independent of SNR. At low SNR (<100), the MAD ranges from 5.2 mÅ for weak lines to 11.6 mÅ for strong lines—nearly identical to the 4.0–10.1 mÅ range at high SNR (>300). This flatness is expected given the uncertainty calibration results above: the systematic floor from continuum methodology differences ($\leq$5 mÅ) dominates over photon noise across the full SNR range of our sample. At SNR = 100, the photon-noise contribution to a typical line (EW $\sim$ 60 mÅ spanning $\sim$30 pixels) is $\sim$2–3 mÅ, below the systematic floor. A noise-injection experiment would therefore primarily confirm what the flat MAD already demonstrates—that our sample is in the systematics-dominated regime. Reaching photon-noise-dominated conditions for typical lines would require SNR $\lesssim$ 20–30, below which continuum fitting itself becomes unreliable. The MAD increases with line strength in absolute terms (4–10 mÅ) but decreases in relative terms (22% for weak lines to 9% for strong lines)—the expected behavior where absolute uncertainty is roughly constant while relative precision improves for stronger features.

5. DISCUSSION

This work demonstrates that an autonomous LLM agent, serving as both validator and adjuster, can perform EW measurement at human expert level. The key is providing the agent with appropriate tools and letting it reason about fits visually—the same process a human expert would follow, but automated and reproducible. Several implications and caveats warrant discussion.

5.1. *Implications for Surveys*

The ability to measure EWs automatically, reproducibly, and with full provenance has implications for spectroscopic surveys. Currently, even high-resolution surveys rely on full-spectrum fitting partly because manual EW measurement cannot scale to hundreds of thousands of stars. With Egent-like pipelines, line-by-line analysis becomes feasible at survey scale—at least for resolved lines at $R > 20,000$.

A key contribution of this work is demonstrating a complete end-to-end pipeline: starting from reduced flux spectra, everything proceeds autonomously to produce an EW table where every fit is visualized and flagged. No intermediate steps require human intervention.

The computational requirements are modest. Direct Voigt fitting requires $\sim$0.1 seconds per line. LLM inspection adds latency—GPT-5 requires $\sim$1 second per iteration (up to 10 iterations for difficult cases), while GPT-5-mini operates at sub-second speeds. The process is massively parallelizable, limited primarily by CPU cores (for generating diagnostic PNG plots, which proves more expensive than the fitting itself) and API rate limits. In our GPT-5-mini runs, we sustained 80 concurrent workers without hitting rate limits; the bottleneck was local plot generation, easily addressed with a small computing cluster.

To estimate survey-scale feasibility: for the $\sim$200-line lists typical of differential abundance studies, at $\sim$1 second per line, a single spectrum completes in $\sim$3 minutes. With 1000 parallel workers, 500,000 spectra (comparable to APOGEE) would complete in $\sim$2 days of wall-clock time. For analyses requiring more lines per spectrum, computation time scales proportionally.

The scientific potential is considerable. As discussed in the introduction, EW-based differential analysis persists because it can cancel systematic errors that plague full-spectrum fitting. Formulating proper differential measurements that harness millions of spectra with billions of individual EW measurements represents a scale not yet explored—but one that autonomous pipelines now make accessible.

Current differential abundance studies use carefully curated line lists containing $\sim$200 lines chosen for minimal blending and reliable atomic data. We use these lines because they are easy to measure reliably, and we measure them reliably because experts have accumulated expe-

Large-survey runtimes are scaling estimates, not completed demonstrations.

rience with them. If all lines could be measured automatically, new approaches to abundance determination become possible. Self-calibration methods could use the full line list to constrain both abundances and atomic data simultaneously.

5.2. *Reproducibility and Human-in-the-Loop Development*

Every Egent measurement includes complete provenance: Voigt parameters, continuum coefficients, quality metrics, and (when triggered) the full LLM conversation. This differs qualitatively from manual measurement, where a published EW cannot be traced to specific decisions about continuum placement or blend handling.

A strength of the agentic approach is that human expertise enters through prompts and tool definitions, which can be iteratively refined. Much of Egent's development involved examining cases where initial prompts failed, understanding why, and improving the instructions. We employed an outer loop where an LLM evaluated failure cases on small test sets, iterating over different prompt formulations—effectively using AI to help develop AI prompts. This process led to the final prompts presented in Appendix A.

For example, early versions of the prompt did not explicitly mention W-shaped residuals as indicators of blends; adding this guidance improved blend detection. Users can also incorporate domain knowledge—for instance, flagging specific wavelength regions with known blends that warrant extra attention. This refinement process is more systematic than the implicit expertise accumulated in a human measurer's head—and it is version-controlled. Once prompt development is complete, everything proceeds autonomously. The agent mimics the same iterative inspection process that a human expert would perform, but without the tedium that makes manual measurement impractical at scale.

It is useful to distinguish two forms of reproducibility that the pipeline provides. The first is *archival reconstruction*: every completed run stores the full Voigt parameters, continuum coefficients, and LLM reasoning chain, so any measurement can be exactly reconstructed from the saved JSON without re-running the model. This form of reproducibility is deterministic and permanent. The second is *prospective reprocessing*: re-running the pipeline on the same data at a later date. Because LLM outputs depend on model version, sampling temperature, and stochastic decoding, re-processing will not produce bit-identical results.

We quantify this variation using the cross-model comparison as an upper bound, since GPT-5 and GPT-5-mini are architecturally different models. Among 3,640 lines measured by both models on the same spectra: 55% have exactly identical EW (both accepted the deterministic direct fit), 73% agree within 1 mÅ, and 96% agree within 10 mÅ. The 90th-percentile difference is 5.0 mÅ with no systematic bias (mean $\Delta\mathrm{EW} = -0.0$ mÅ). Even for lines where both models invoked the LLM (n = 2,632), the median difference is only 0.3 mÅ. This cross-model scatter is comparable to the systematic floor from continuum methodology, indicating that model choice is not the dominant source of uncertainty. Run-to-run variability with the same model would be smaller still. For long-term scientific use, we recommend archiving the stored outputs rather than relying on re-execution, and recording the model version used for each run. The cross-model comparison reported here is an upper bound rather than a direct measurement of stochastic variability; a dedicated same-model repeatability experiment—re-running identical prompts at fixed sampling temperature on the same lines and the same model version—would provide a more definitive lower bound on the run-to-run noise floor under model drift, and is a natural future test.

5.3. *Offline and Web Interface Versions*

While the core pipeline is compact—a few hundred lines of Python with API calls—we recognize that deploying commercial LLM APIs may present barriers for some users. We therefore provide two additional deployment options.

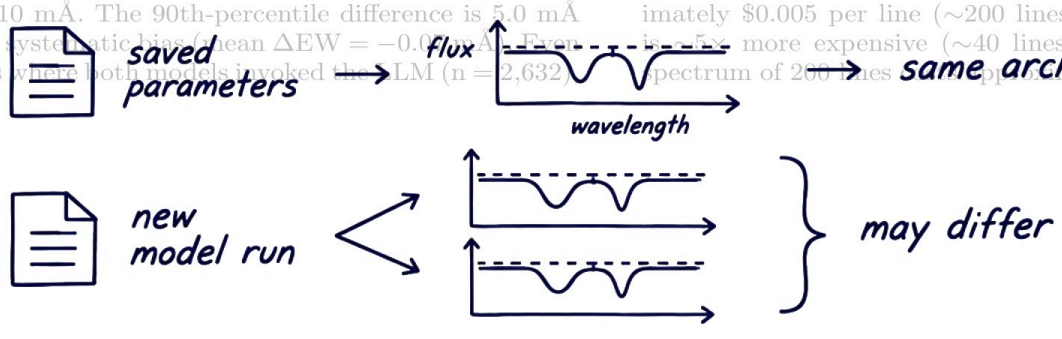
First, we offer a fully offline version using Qwen3-VL-8B (Yang et al. 2025), an open-source vision-language model deployable through the Ollama ecosystem. The entire installation is self-contained and runs on consumer hardware (e.g., a Mac with default RAM). This is particularly valuable for observatory settings where external internet access through APIs may raise security concerns. While processing a few spectra is feasible on such hardware, speed is limited without GPU acceleration—generating reasoning tokens benefits from dedicated GPU hardware. We do not provide formal benchmarks as we lack a cluster of Macbooks for systematic testing. However, given that GPT-5-mini performs well, and lightweight open-source vision-language models have been reported to achieve comparable (or superior) performance at the small-model end, we expect the offline version to produce reliable results, albeit more slowly.

Second, we provide a web interface for drag-and-drop analysis. Users can upload spectra without requiring their own API key—the web service covers the LLM costs. An initial preprocessing step converts arbitrary formats into Egent's expected input. While this interface does not support parallel processing, it enables quick testing and demonstration without any local installation or API configuration.

Installation instructions for both the Ollama backend and the web interface are available at https://github.com/tingyuansen/Egent.

5.4. *Limitations*

The monetary cost of LLM APIs represents a practical limitation. For GPT-5-mini, processing costs approximately \$0.005 per line (∼200 lines per dollar); GPT-5 is ∼5× more expensive (∼40 lines per dollar). A full spectrum of 200 lines [illegible] for GPT-5-



Reconstructing a saved fit is different from rerunning a model.

The offline version is offered, but its accuracy is not formally benchmarked here.

Validation has boundaries—and visual inspection can still miss a bad fit.

mini. For an APOGEE-scale survey (500,000 spectra × 200 lines = 100 million lines), total API cost would be ∼$500,000 for GPT-5-mini—not trivially cheap, but not impossible for a major survey either.

These estimates represent an upper limit. One straightforward way to reduce costs is relaxing the quality thresholds that trigger LLM inspection, relying more heavily on algorithmic checks and reserving visual inspection for the most uncertain cases. In our current configuration, ∼60–65% of lines trigger LLM inspection; tightening acceptance criteria for direct fits could reduce this fraction substantially. Fine-tuned lightweight open-source models (1–2B parameters) represent another promising direction: if domain-specific fine-tuning can achieve comparable performance to general-purpose frontier models, costs could decrease by orders of magnitude while enabling fully offline operation. As LLM capabilities continue to improve and costs continue to decline, these estimates will become increasingly conservative.

The validated regime is high-resolution, mostly high-SNR solar-type spectra.

Another limitation is uncertainty about how Egent scales to lower resolution or SNR. Our validation uses C3PO spectra, which by survey design represent the high-SNR regime (SNR > 50). Egent is designed for high-resolution ($R > 20,000$), high-SNR spectra of solar-type stars. At lower resolution, blending becomes severe and the "local continuum" approximation breaks down—though this is not a fundamental limitation of the agentic approach, but rather would require more sophisticated priors (e.g., from spectral synthesis). Testing has focused on solar-type stars; extension to metal-poor stars (where lines are weaker and sparser), rapid rotators (where lines are broadened and blended), and emission-line objects would require validation. The framework should generalize, but prompts and quality thresholds may need adjustment.

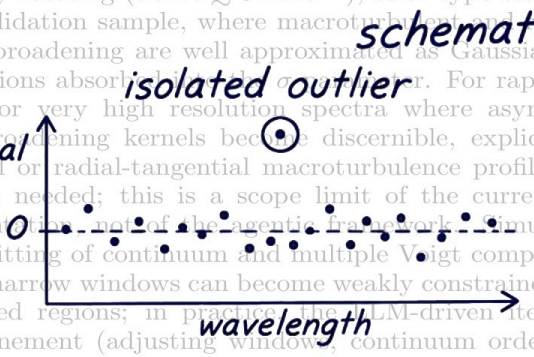
Regarding line profiles: the Voigt model is adequate for the slowly rotating ($v \sin i \lesssim 5$ km s^{-1}), solar-type stars in our validation sample, where macroturbulent and rotational broadening are well approximated as Gaussian contributions absorbed into the Gaussian width parameter. For rapid rotators or very high resolution spectra where asymmetric broadening kernels become discernible, explicit rotational or radial-tangential macroturbulence profiles would be needed; this is a scope limit of the current implementation, not of the agentic framework. Simultaneous fitting of continuum and multiple Voigt components in narrow windows can become weakly constrained in crowded regions; in practice, the LLM-driven iterative refinement (adjusting window, continuum order, and component count) serves as a practical stability diagnostic, and fits that remain unstable are flagged. Because the fitting engine is implemented in pure Python with minimal dependencies (`scipy`, `numpy`), adding new profile functions (e.g., a rotational kernel or radial-tangential macroturbulence) requires only implementing the profile and registering it as a tool—no modification of legacy compiled code is needed.

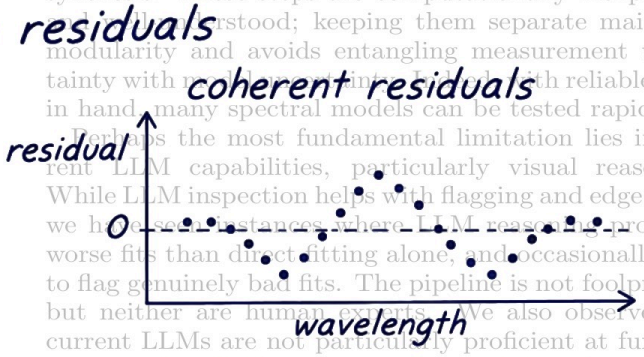
Rapid rotators and crowded low-resolution spectra need further validation.

The fitting engine uses standard Levenberg–Marquardt least squares rather than robust alternatives such as IRLS with a Huber or Tukey loss. Analysis of the quality issues that trigger LLM intervention reveals that the dominant failure modes are correlated residuals from missed blends (present in >80% of triggered fits) and wrong continuum shape—both structural model-mismatch problems that robust regression would not resolve, since the correct response is adding a Voigt component or changing the continuum model, not downweighting individual pixels. The continuum fitting already incorporates iterative 2.5σ clipping, which provides robustness against absorption features during initial continuum estimation. Robust regression and LLM inspection thus target orthogonal failure modes: the former addresses point-outlier sensitivity, while the latter handles model selection (which components, which continuum, which window). The two are complementary, and adding a robust loss function to the fitting engine is a straightforward future extension. The argument above is based on inspection of the failure-mode statistics rather than a side-by-side ablation; a controlled deterministic test—e.g., swapping the standard Levenberg–Marquardt loss for an IRLS scheme with a Huber or Tukey biweight loss within the same engine, and then re-measuring how often the LLM is still triggered and whether accepted-line scatter changes—would more definitively quantify what fraction of the inspection benefit is attributable to robust statistics versus model-guided adjustments.

Egent addresses the labor-intensive bottleneck in EW measurement, but it is not a complete end-to-end abundance pipeline. Other steps—such as rest-frame velocity correction via cross-correlation—remain outside its scope, though these are largely programmable tasks that do not require the same human judgment. The natural extension is connecting EW measurement to abundance determination via curve-of-growth analysis or spectral synthesis. These steps are computationally inexpensive and well understood; keeping them separate maintains modularity and avoids entangling measurement uncertainty with model assumptions. With reliable EWs in hand, many spectral models can be tested rapidly.

Perhaps the most fundamental limitation lies in current LLM capabilities, particularly visual reasoning. While LLM inspection helps with flagging and edge cases, we have seen instances where LLM reasoning produces worse fits than direct fitting alone, and occasionally fails to flag genuinely bad fits. The pipeline is not foolproof—but neither are human experts. We also observe that current LLMs are not particularly proficient at function tool use, contributing to the ∼10% flagging rate even at high resolution. A natural path forward would be fine-tuning models with reinforcement learning (e.g., GRPO), which does not require labeled data, but this is beyond the scope of this paper.





5.5. *Broader Context*

LLM agents have generated considerable excitement, with a rush to tackle ambitious problems like automated

hypothesis generation. Yet even in this relatively simple case—eight tools, a well-defined task, data the models were unlikely trained on—out-of-the-box models without human-in-the-loop refinement prove challenging. While visual reasoning is strong and enables effective flagging, reliable autonomous operation required careful prompt engineering and constraint building.

This is not to say that agents cannot drastically expedite research. This work demonstrates multiple layers of agentic improvement: using agents to refine prompts under human supervision, building in domain constraints, and iterating on failure cases. For tasks that require serious repetition yet demand judgment—exactly the characteristics of EW measurement—the agentic framework provides a successful model.

The approach generalizes to other astronomical tasks with similar characteristics: damped Lyman-α identification, spectral classification, and other problems where programmatic rules fail on edge cases but human judgment succeeds. As LLMs strengthen, particularly in visual reasoning and tool use, this represents a growing avenue for automating traditionally manual analysis.

6. CONCLUSION

We present Egent, an autonomous agent for equivalent width measurement that combines classical multi-Voigt profile fitting with LLM-based visual inspection. The system operates directly on raw flux spectra without requiring pre-normalized continua, and provides complete provenance for every measurement.

Validation against 18,615 lines from 84 Magellan/MIKE spectra measured by human experts in the C3PO program yields a raw MAD = 5–7 mÅ agreement without any post-hoc per-spectrum correction. Per-spectrum slopes scatter within $\sim$15% of unity (median $\approx$ 1.03), reflecting differences in global continuum methodology between Egent's local fitting and traditional approaches rather than fitting failures; the overall correlation is strong ($R^2 \approx 0.93$). The pipeline is robust across SNR$\sim$50–250, with MAD increasing only modestly from $\sim$5 mÅ at high SNR to $\sim$8 mÅ at SNR < 75.

The LLM's primary role is quality control rather than systematic improvement. Approximately 60–65% of lines trigger LLM inspection and are subsequently refined; the LLM confirms good fits, flags problematic cases ($\sim$10–20%), and occasionally rescues edge cases through tool use. Agreement between GPT-5 and GPT-5-mini demonstrates reproducibility and validates that the pipeline does not depend on model-specific behaviors.

The computational requirements are practical. Direct fitting requires $\sim$0.1 seconds per line; LLM inspection adds $\sim$1 second per iteration. A typical 200-line spectrum costs $\sim$\$1 for GPT-5-mini; survey-scale processing is feasible but not trivially cheap, motivating future work on fine-tuned lightweight models. We provide three deployment options: API-based processing for production runs, a fully offline version using Qwen-3-VL-8B through Ollama for secure environments, and a web interface for drag-and-drop analysis without local installation.

Every measurement stores complete Voigt parameters, continuum coefficients, quality metrics, and (when triggered) the full LLM reasoning chain. This provenance enables exact reconstruction of any fit without re-running the analysis—a qualitative improvement over traditional manual measurement where decisions are rarely recorded.

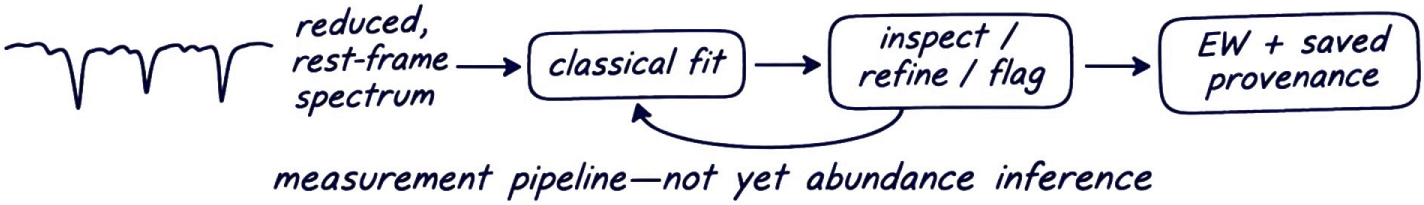
Egent demonstrates that LLM agents can perform expert-level judgment tasks in astronomy when provided with appropriate tools and domain constraints. The approach generalizes to other problems where programmatic rules fail on edge cases but human judgment succeeds. By removing the human bottleneck while maintaining reliability and transparency, Egent opens the path to line-by-line abundance analysis at survey scale—enabling differential analysis methods that harness millions of spectra with billions of individual EW measurements.

This project was originally inspired by the first author's teaching of Astron 1221 (Data Analysis in Astronomy) at OSU, a first-year Python course for astronomers covering both spectroscopic data analysis and basic agentic AI frameworks—which the author deems important even for freshman students to thrive in this era. The pipeline was initially conceived as a class project demonstrating how LLM function tools could be applied to a real astronomical problem; since freshman students tend to gravitate toward either amateur astronomy problems (lunar phases, asteroid impacts) or Nobel Prize-winning topics (S-stars near Sgr A*, 51 Pegasi b)—while all using similar tools—the author decided to take on this project himself.

The recent spectroscopy workshop held by Catherine Manea and Madeline McKenzie inspired the author to finish this project. We appreciate fruitful discussions with David Weinberg, Jennifer Johnson, Lucy Lu, Tawny Sit, Sharon Wang, Qinghui Sun, and Megan Bedell throughout the year. Part of the routine is inspired by the ongoing PASTA project (similar to C3PO), and the author's lack of funding to hire an expert for yet another painstaking EW measurement campaign—combined with OSU's recently grown fortuitous LBT time fraction, which will lead to more PEPSI spectra that will only accentuate this problem.

This research was supported by NSF Grant AST-2406729 and a Humboldt Research Award from the Alexander von Humboldt Foundation. The GPT-5 and GPT-5-mini API allocation in this development was substantial and is supported by the NSF NAIRR program. The authors acknowledge the National Artificial Intelligence Research Resource (NAIRR) Pilot and OpenAI for contributing to this research result.

While the ideas, draft, and layout are all from the authors, this project has benefited—in line with its own ideals—from collaboration with Claude-Opus-4.5 and



Cursor for both coding development and writing refinement. Google's Nanobanana-Pro is used to generate the schema in Figure 2.

We thank D. Yong's right hand for the C3PO measurements that enabled validation.

APPENDIX

A. SYSTEM PROMPTS

We provide the complete system prompt used for LLM visual inspection, enabling full reproducibility.

A.1. *Visual Inspection Prompt*

The following prompt instructs the LLM on its role as an expert visual reviewer for borderline EW measurements. When the direct multi-Voigt fit produces quality metrics outside acceptable thresholds (Section 2.7), the LLM is invoked to inspect the diagnostic plot and decide whether to accept, improve, or flag the measurement. A key feature of this prompt is the explicit instruction to focus on the target region (±0.5 Å) and ignore edge effects that do not affect the measurement.

You are an expert reviewer for stellar EW measurements. You're called for BORDERLINE cases where automated fitting is uncertain.

YOUR ROLE: Use visual inspection to make judgment calls that algorithms can't.

WORKFLOW: (1) The fit is already done. Call get_fit_plot to VISUALLY INSPECT. (2) Make a judgment: ACCEPT, IMPROVE, or FLAG.

⇒ CRITICAL: FOCUS ONLY ON TARGET REGION ⇐
The TARGET LINE is marked with a BLUE DASHED vertical line. ONLY care about fit quality within ±0.5 Å of this target.
EDGE vs TARGET distinction:

- IGNORE: Spikes/noise at window edges (>1.5 Å from target)—these are normal
- CHECK: Residuals within ±0.5 Å of the blue target line
- If target region is clean but edges are messy → ACCEPT
- If target region has W-pattern or blend → FIX or FLAG

RMS INTERPRETATION:

- **RMS < 1.5σ:** LIKELY GOOD—still check for W-shaped residuals at target
- **RMS 1.5–2.0σ:** GOOD—check target region only, accept if clean
- **RMS 2.0–2.5σ:** MARGINAL—try to improve, accept if target is clean after
- **RMS > 2.5σ:** POOR—needs work or FLAG

BLEND DETECTION (CHECK EVEN WITH LOW RMS):
Look for these patterns in residuals NEAR THE TARGET (±1 Å):

- **W-SHAPED residuals** at target = MISSED BLEND. Fix: fit_ew(additional_peaks=[wavelength])
- **ASYMMETRIC wing/shoulder** on target line. Fix: add the wing as additional_peak
- **Systematic offset** at target center. May need continuum adjustment or nearby blend

ACCEPT (call record_measurement) if: RMS < 1.5σ (excellent fit), OR RMS < 2.0σ AND target region residuals within ±2σ.

IMPROVE by trying: (1) Add missing peaks near target. (2) Polynomial continuum for curved residuals. (3) Window adjustment (narrower for edge issues, wider for truncated lines). (4) Manual continuum regions for crowded areas.

FLAG (use sparingly!) only when: Problem is AT THE TARGET (±0.5 Å), not at edges; RMS > 2.5σ that cannot be improved; target line is completely blended or absent.

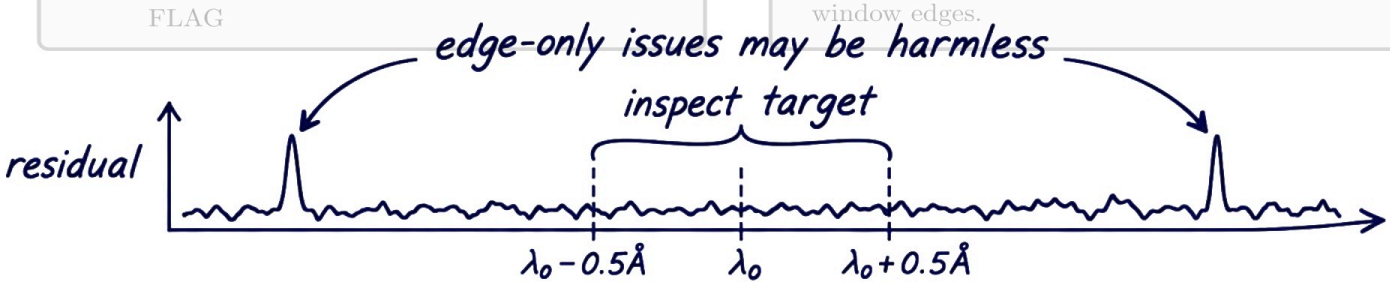
DO NOT FLAG for: Edge spikes/noise far from target (>1.5 Å away); Low RMS (<2.0σ) with clean target region.

B. TOOL SCHEMAS

Eight tools are available to the LLM agent, each implemented as an OpenAI function-calling schema. These tools mirror the operations a human spectroscopist would perform in IRAF or similar software: loading data, adjusting the wavelength window, configuring continuum fitting, adding blend components, inspecting results visually, and recording the final measurement.

`load_spectrum(gaia_id: int)` — Loads a stellar spectrum from the data repository with barycentric correction applied. Returns metadata including SNR, wavelength range, and number of échelle orders available.

`extract_region(line_wavelength: float, window: float = 3.0)` — Extracts the spectral region centered on the target wavelength. The `window` parameter specifies the half-width in Å (default ±3 Å). The LLM uses this to "zoom in" (`window=2.0`) for isolated lines or "zoom out" (`window=5.0`) when strong lines are truncated at window edges.



`set_continuum_method(method: str, order: int = 1)` — Configures the continuum fitting approach. Options: `iterative_linear` (default), `iterative_poly` (polynomial of specified order for curved continuum), `top_percentile` (uses brightest pixels as anchors in crowded regions).

`set_continuum_regions(regions: list)` — Last-resort tool for crowded regions where automatic continuum detection fails. The LLM visually identifies wavelength intervals that appear line-free and specifies them as `[[start1, end1], [start2, end2], ...]`. The continuum is then fit using only pixels within these regions.

`fit_ew(additional_peaks: list = [])` — Performs the multi-Voigt fit using current extraction region and continuum settings. All catalog lines in the window are included automatically. The `additional_peaks` parameter allows the LLM to specify wavelengths for Voigt components not in the catalog (for missed blends identified via W-shaped residuals). Returns EW, uncertainty, quality metrics, and full Voigt parameters for all fitted lines.

`get_fit_plot()` — Generates the diagnostic plot for visual inspection. Returns a base64-encoded PNG showing: (1) normalized spectrum with model overlay, (2) individual Voigt component centers, (3) target wavelength indicator, and (4) normalized residuals with $\pm 1\sigma$ and $\pm 2\sigma$ bands.

`flag_line(line_wavelength: float, reason: str)` — Flags a line as unreliable when no satisfactory fit can be achieved. Valid reasons: `no_data`, `severe_blend`, `fit_failed`, `bad_continuum`, `wrong_line`. The flag and reason are stored in the output JSON.

`record_measurement(line_wavelength: float, ew_mA: float, ew_err_mA: float, quality: str)` — Records the final EW measurement with uncertainty and quality assessment (`good`, `acceptable`, or `poor`). This concludes the LLM inspection for accepted measurements.

C. COMPLETE LLM CONVERSATION EXAMPLE

Below we present a complete, verbatim conversation from GPT-5-mini processing Ca I 6696.02 Å (Gaia ID 5578084051306739 2). This case demonstrates blend detection via W-shaped residuals, peak injection to deblend, and the LLM's reasoning at each step. The initial direct fit triggered LLM inspection due to poor quality (central RMS = 3.63σ).

Iteration 1: Initial Visual Inspection

LLM reasoning: "I will visually inspect the existing fit plot around 6696.02 Å to check: (a) whether the red fit follows the black data at the target, (b) the residuals for any W-shaped pattern indicating missed blends (especially near 6696.0 Å), and (c) whether the continuum shows a slope or offset. I'll call get_fit_plot to view the diagnostic."

Action: `get_fit_plot({})`

Iteration 2: Blend Detection and Peak Injection

LLM reasoning: "I observe a poor central fit: the central RMS is 3.63σ ($\gg 2.5\sigma$) and the residuals around 6696 Å show a W-like pattern / asymmetric structure suggesting the target is blended (the fit treats the feature as a single symmetric component). In the main panel the data shows a slight asymmetry/shoulder near 6695.9–6696.1 Å that the red fit doesn't fully reproduce. I will attempt to deblend by adding an additional Voigt component near 6695.9 Å (keeping the main component at the target) to model the close blend."

Action: `fit_ew({"additional_peaks": [6695.9]})`

neighbour added, target retained

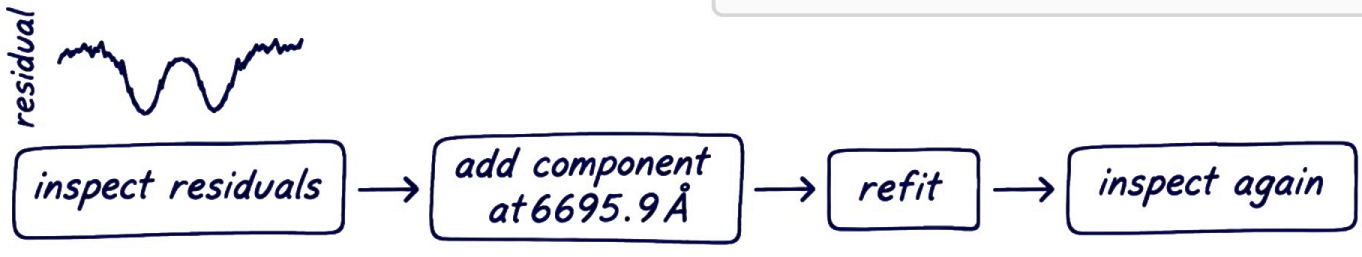
Iteration 3: Re-inspection After Deblending

LLM reasoning: "The updated fit deblended into multiple components but quality flags remain: correlated residuals and central RMS $> 1.5\sigma$. I will inspect the updated diagnostic plot to locate remaining mismatches (any residual W-shapes or unmodeled dips around 6696.0–6696.3 Å and at the right-side feature near 6698.65 Å). If residuals show continuum tilt, I'll try a polynomial continuum next."

Action: `get_fit_plot({})`

Iteration 4: Accept Measurement

LLM reasoning: "After adding a deblending component at 6695.9 Å the fit follows the data closely at the target. The residuals no longer show a W-shape at the target and the central RMS is 1.64σ (overall RMS 1.49σ), which is within the acceptable range (1.5–2.5). There



Deblending changes which absorption area belongs to the target.

remains a neighboring fitted line at ∼6696.283 Å that was modeled. I will record the measured EW from the improved fit."

Action:

```
record_measurement({"line_wavelength":
6696.02, "ew_mA": 24.6, "ew_err_mA":
2.1, "quality": "acceptable"})
```

In this example, target EW changes from 38.8 to 24.6 mÅ.

This sequence—measure, diagnose, adjust, re-inspect—mimics the workflow of a human expert but executes in seconds and produces a reproducible record. The explicit reasoning text documents *why* each decision was made, not just what action was taken. The final EW (24.6 mÅ) differs from the initial direct fit (38.8 mÅ) by 37%, demonstrating the importance of proper deblending.

D. INDIVIDUAL LINE FITS

Accepted mosaic tiles can still contain imperfect fits.

Figures D1–D3 show all individual line fits for three representative spectra from the GPT-5-mini sample. Each small panel displays the normalized flux (black) and multi-Voigt model (red), with the target wavelength indicated by the blue vertical line. Background colors distinguish three fit categories: gray indicates direct fits accepted without LLM intervention, green indicates LLM-improved fits, and red indicates flagged lines.

These figures demonstrate that Egent produces high-quality fits across the full wavelength range and line strength distribution. The majority of panels show excellent agreement between data and model, with the fitted Voigt profiles accurately capturing both strong and weak absorption features. Direct fits (gray) dominate in clean spectral regions, while LLM intervention (green) was required for more complex cases involving blends, continuum curvature, or edge effects. The flagged lines (red) represent cases where no satisfactory fit could be achieved during the iterative process—typically due to severe blending, contamination, or low SNR.

Rarely, problematic fits may evade flagging—visual inspection of these figures reveals occasional cases where the model does not perfectly match the data but was not flagged. The flagging threshold can be made more stringent via prompt adjustment, but this involves a trade-off: overly strict criteria flag many acceptable fits, reducing the final sample size without meaningful improvement in quality. The current calibration (flagging ∼10–20% of lines) represents a reasonable balance, but users with more conservative requirements can adjust the flagging thresholds or perform additional manual review of the retained lines.

A key feature visible in these plots: the normalized flux (raw flux divided by locally-estimated continuum) consistently lies near unity across all panels. This demonstrates the quality of Egent's approach—fitting directly from raw flux spectra with simultaneous local continuum estimation, completely bypassing the need for a separate global continuum normalization step. In contrast, C3PO relies on pre-normalized spectra where continuum placement is performed as a separate, subjective step. The systematic offsets in Figure 8 therefore reflect C3PO's global continuum choices, not Egent's fitting quality. Egent's single-step approach mitigates this source of systematic error.

REFERENCES

S. R. Majewski, R. P. Schiavon, P. M. Frinchaboy, C. Allende Prieto, R. Barkhouser, D. Bizyaev, B. Blank, S. Brunner, A. Burton, R. Carrera, et al., AJ **154**, 94 (2017), 1509.05420.

S. Buder, J. Kos, X. E. Wang, M. McKenzie, M. Howell, S. Martell, M. R. Hayden, D. B. Zucker, T. Nordlander, B. Montet, et al., PASA **42**, e051 (2025), 2409.19858.

S. Randich, G. Gilmore, L. Magrini, G. G. Sacco, R. J. Jackson, R. D. Jeffries, C. C. Worley, A. Hourihane, A. Gonneau, C. Viscasillas Vazquez, et al., A&A **666**, A121 (2022), 2206.02901.

J. A. Kollmeier, G. Zasowski, H.-W. Rix, M. Johns, S. F. Anderson, N. Drory, J. A. Johnson, R. W. Pogge, J. C. Bird, G. A. Blanc, et al., arXiv e-prints arXiv:1711.03234 (2017), 1711.03234.

R. S. de Jong, O. Bellido-Tirado, C. Chiappini, É. Depagne, R. Haynes, D. Johl, O. Schnurr, A. Schwope, J. Walcher, F. Dionies, et al., in *Ground-based and Airborne Instrumentation for Astronomy IV*, edited by I. S. McLean, S. K. Ramsay, and H. Takami (2012), vol. 8446 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 84460T, 1206.6885.

G. Dalton, S. Trager, D. C. Abrams, P. Bonifacio, J. A. López Aguerri, K. Middleton, C. Benn, K. Dee, F. Sayède, I. Lewis, et al., in *Ground-based and Airborne Instrumentation for Astronomy V*, edited by S. K. Ramsay, I. S. McLean, and H. Takami (2014), vol. 9147 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 91470L, 1412.0843.

M. Ness, D. W. Hogg, H.-W. Rix, A. Y. Q. Ho, and G. Zasowski, ApJ **808**, 16 (2015), 1501.07604.

J. A. Holtzman, M. Shetrone, J. A. Johnson, C. Allende Prieto, F. Anders, B. Andrews, T. C. Beers, D. Bizyaev, M. R. Blanton, J. Bovy, et al., AJ **150**, 148 (2015), 1501.04110.

A. E. García Pérez, C. Allende Prieto, J. A. Holtzman, M. Shetrone, S. Mészáros, D. Bizyaev, R. Carrera, K. Cunha, D. A. García-Hernández, J. A. Johnson, et al., AJ **151**, 144 (2016), 1510.07635.

Y.-S. Ting, C. Conroy, H.-W. Rix, and P. Cargile, ApJ **879**, 69 (2019), 1804.01530.

M. Xiang, H.-W. Rix, Y.-S. Ting, R.-P. Kudritzki, C. Conroy, E. Zari, J.-R. Shi, N. Przybilla, M. Ramirez-Tannus, A. Tkachenko, et al., A&A **662**, A66 (2022), 2108.02878.

K. Korista, J. Baldwin, G. Ferland, and D. Verner, ApJS **108**, 401 (1997), astro-ph/9611220.

A. Milone, R. da Silva, A. Sansom, and P. Sanchez-Blazquez, arXiv e-prints arXiv:1309.0548 (2013), 1309.0548.

D. Montes, R. González-Peinado, H. M. Tabernero, J. A. Caballero, E. Marfil, F. J. Alonso-Floriano, M. Cortés-Contreras, J. I. González Hernández, A. Klutsch, and C. Moreno-Jódar, MNRAS **479**, 1332 (2018), 1805.05394.

P. Scott, M. Asplund, N. Grevesse, M. Bergemann, and A. J. Sauval, A&A **573**, A26 (2015a), 1405.0287.

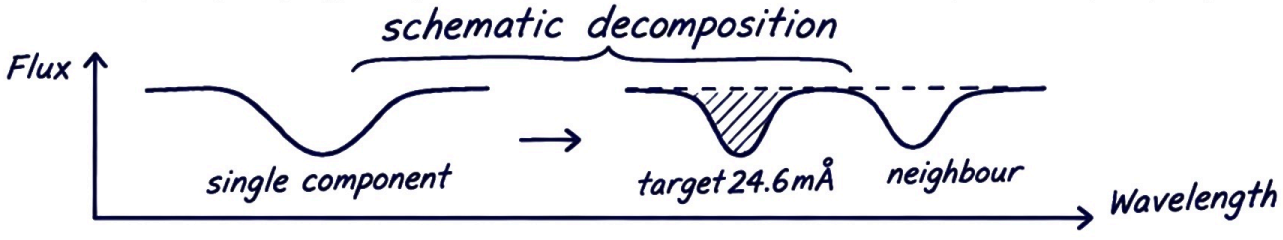
P. Scott, N. Grevesse, M. Asplund, A. J. Sauval, K. Lind, Y. Takeda, R. Collet, R. Trampedach, and W. Hayek, A&A **573**, A25 (2015b), 1405.0279.

M. Asplund, Å. Nordlund, R. Trampedach, and R. F. Stein, A&A **359**, 743 (2000), astro-ph/0005321.

M. Asplund, N. Grevesse, A. J. Sauval, and P. Scott, ARA&A **47**, 481 (2009), 0909.0948.

M. Bergemann, A. J. Gallagher, P. Eitner, M. Bautista, R. Collet, S. A. Yakovleva, A. Mayriedl, B. Plez, M. Carlsson, J. Leenaarts, et al., A&A **631**, A80 (2019), 1905.05200.

K. Lind and A. M. Amarsi, ARA&A **62**, 475 (2024), 2401.00697.



The mosaic exposes ordinary fits as well as rejected cases.



Figure D1. All line fits for a representative spectrum. Each panel shows the target wavelength (Å) above the plot, with normalized flux (black), multi-Voigt model (red), and target line position (blue dashed). Background colors: gray = direct fit accepted, green = LLM-improved, red = flagged. Lines are sorted by fit type.

Colour records the processing category, not guaranteed correctness.

Review can leave an already useful fit nearly unchanged.

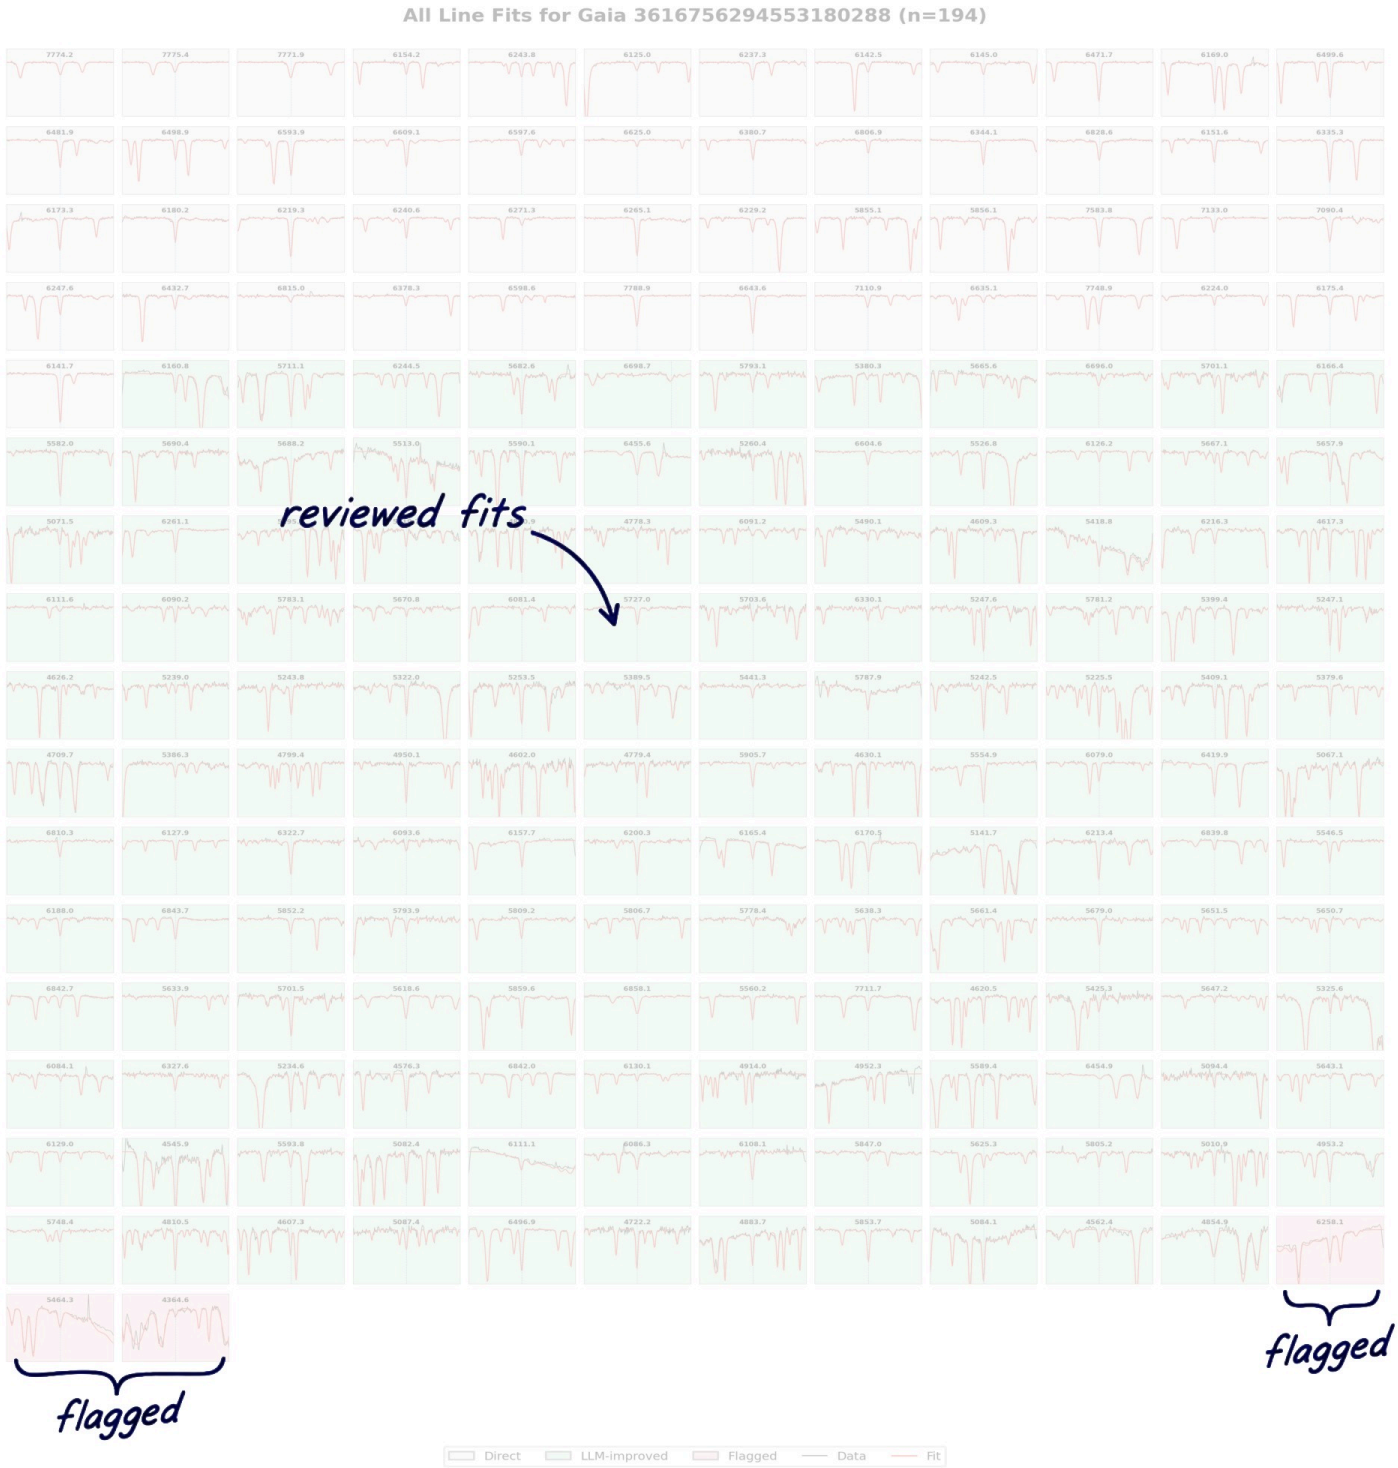


Figure D2. All line fits for a second representative spectrum. Background colors: gray = direct fit accepted, green = LLM-improved, red = flagged. Lines are sorted by fit type.

Green records LLM involvement; it does not measure the size of the EW correction.

A retained fit is still worth checking when the target looks awkward.

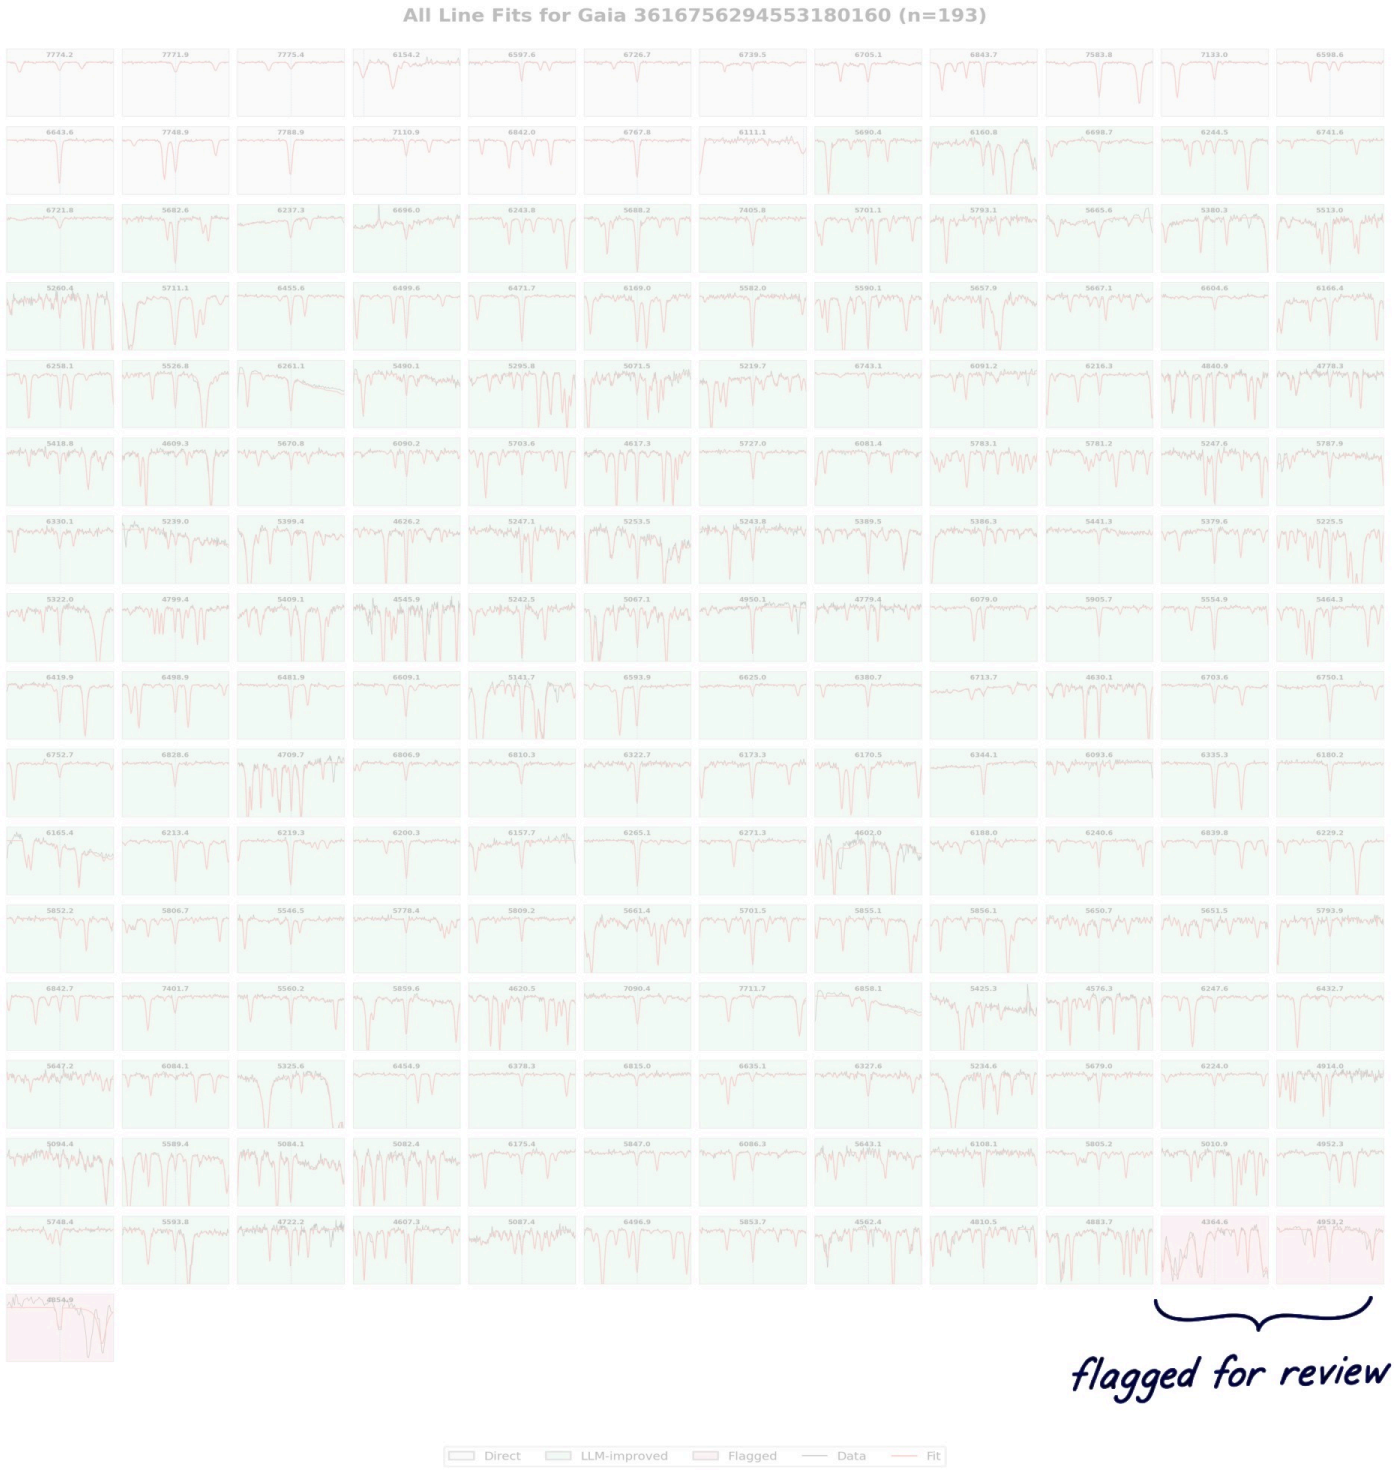


Figure D3. All line fits for a third representative spectrum. Background colors: gray = direct fit accepted, green = LLM-improved, red = flagged. Lines are sorted by fit type.

The complete mosaics make spot checks possible; they do not certify every fit.

T. Bensby, S. Feltzing, and I. Lundström, A&A **415**, 155 (2004), astro-ph/0310741.
M. Bedell, J. Meléndez, J. L. Bean, I. Ramírez, P. Leite, and M. Asplund, ApJ **795**, 23 (2014), 1409.1230.
P. E. Nissen, A&A **593**, A65 (2016), 1606.08399.
I. Ramirez, J. Melendez, and M. Asplund, arXiv e-prints arXiv:0911.1893 (2009), 0911.1893.
J. Yana Galarza, R. López-Valdivia, J. Meléndez, and D. Lorenzo-Oliveira, ApJ **922**, 129 (2021), 2109.00679.
F. Liu, Y.-S. Ting, D. Yong, B. Bitsch, A. Karakas, M. T. Murphy, M. Joyce, A. Dotter, and F. Dai, Nature **627**, 501 (2024), 2403.13209.
L. Spina, arXiv e-prints arXiv:2401.12296 (2024), 2401.12296.
S. G. Sousa, N. C. Santos, G. Israelian, M. Mayor, and M. J. P. F. G. Monteiro, A&A **469**, 783 (2007), astro-ph/0703696.
P. B. Stetson and E. Pancino, PASP **120**, 1332 (2008), 0811.2932.
W. Kang and S.-G. Lee, MNRAS **425**, 3162 (2012), 1207.2845.
S. Blanco-Cuaresma, C. Soubiran, U. Heiter, and P. Jofré, A&A **569**, A111 (2014), 1407.2608.
R. R. Lindner, C. Vera-Ciro, C. E. Murray, S. Stanimirović, B. Babler, C. Heiles, P. Hennebelle, W. M. Goss, and J. Dickey, AJ **149**, 138 (2015), 1409.2840.
M. Riener, J. Kainulainen, J. D. Henshaw, J. H. Orkisz, C. E. Murray, and H. Beuther, A&A **628**, A78 (2019), 1906.10506.
F. Liu, B. Bitsch, M. Asplund, B.-B. Liu, M. T. Murphy, D. Yong, Y.-S. Ting, and S. Feltzing, MNRAS **508**, 1227 (2021), 2108.11001.
P. Škoda, B. Šurlan, and S. Tomić, in *Ground-based and Airborne Instrumentation for Astronomy II*, edited by I. S. McLean and M. M. Casali (2008), vol. 7014 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, p. 70145X, 1112.4012.
V. D. Ivanov, M. J. Rieke, C. W. Engelbracht, A. Alonso-Herrero, G. H. Rieke, and K. L. Luhman, ApJS **151**, 387 (2004), astro-ph/0311596.
J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. Chi, Q. Le, and D. Zhou, arXiv e-prints arXiv:2201.11903 (2022), 2201.11903.
J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al., arXiv e-prints arXiv:2303.08774 (2023), 2303.08774.
S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, and Y. Cao, arXiv e-prints arXiv:2210.03629 (2022), 2210.03629.
D. A. Boiko, R. MacKnight, and G. Gomes, arXiv e-prints arXiv:2304.05332 (2023), 2304.05332.
A. M. Bran, S. Cox, O. Schilter, C. Baldassari, A. D. White, and P. Schwaller, arXiv e-prints arXiv:2304.05376 (2023), 2304.05376.
M. Caldas Ramos, C. J. Collison, and A. D. White, arXiv e-prints arXiv:2407.01603 (2024), 2407.01603.
N. J. Szymanski, B. Rendy, Y. Fei, R. E. Kumar, T. He, D. Milsted, M. J. McDermott, M. Gallant, E. D. Cubuk, A. Merchant, et al., Nature **624**, 86 (2023).
Z. Sun, Y.-S. Ting, Y. Liang, N. Duan, S. Huang, and Z. Cai, arXiv e-prints arXiv:2406.01391 (2024), 2406.01391.
Z. Sun, Y.-S. Ting, Y. Liang, N. Duan, S. Huang, and Z. Cai, arXiv e-prints arXiv:2510.08354 (2025a), 2510.08354.
D. Yong, F. Liu, Y.-S. Ting, M. Joyce, B. Bitsch, F. Dai, A. Dotter, A. I. Karakas, and M. T. Murphy, MNRAS **526**, 2181 (2023), 2309.01546.
Q. Sun, Y.-S. Ting, F. Liu, S. X. Wang, B. J. Anthony-Twarog, B. A. Twarog, J.-Y. Yang, D.-C. Chen, A. I. Karakas, J.-W. Xie, et al., ApJ **978**, 107 (2025b), 2410.20632.
Q. Sun, Y.-S. Ting, B. J. Anthony-Twarog, B. A. Twarog, F. Liu, and Y. Lu, ApJ **991**, 185 (2025c), 2508.08671.
J. Yu, Y.-S. Ting, L. Casagrande, F. Liu, S. X. Wang, Q. Sun, D. Huber, B. Chen, G. Cordoni, G. Da Costa, et al., MNRAS **538**, 2408 (2025), 2503.10339.
S. G. Sousa, N. C. Santos, V. Adibekyan, E. Delgado-Mena, and G. Israelian, A&A **577**, A67 (2015), 1504.02725.
A. Yang, A. Li, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Gao, C. Huang, C. Lv, et al., arXiv e-prints arXiv:2505.09388 (2025), 2505.09388.

This paper was built using the Open Journal of Astrophysics LaTeX template. The OJA is a journal which provides fast and easy peer review for new papers in the `astro-ph` section of the arXiv, making the reviewing process simpler for authors and referees alike. Learn more at http://astro.theoj.org.