



An Interpretable Machine-learning Framework for Modeling High-resolution Spectroscopic Data*

Michael Gully-Santiago  and Caroline V. Morley 

The University of Texas at Austin, Department of Astronomy, 2515 Speedway, Stop C1400, Austin, Texas 78712-1205, USA

Received 2022 August 12; revised 2022 October 3; accepted 2022 October 4; published 2022 December 23

Abstract

Comparison of échelle spectra to synthetic models has become a computational statistics challenge, with over 10,000 individual spectral lines affecting a typical cool star échelle spectrum. Telluric artifacts, imperfect line lists, inexact continuum placement, and inflexible models frustrate the scientific promise of these information-rich data sets. Here we debut an interpretable machine-learning framework *blasé* that addresses these and other challenges. The semiempirical approach can be viewed as “transfer learning”—first pretraining models on noise-free precomputed synthetic spectral models, then learning the corrections to line depths and widths from whole-spectrum fitting to an observed spectrum. The auto-differentiable model employs back-propagation, the fundamental algorithm empowering modern deep learning and neural networks. Here, however, the 40,000+ parameters symbolize physically interpretable line profile properties such as amplitude, width, location, and shape, plus radial velocity and rotational broadening. This hybrid data-/model-driven framework allows joint modeling of stellar and telluric lines simultaneously, a potentially transformative step forward for mitigating the deleterious telluric contamination in the near-infrared. The *blasé* approach acts as both a deconvolution tool and semiempirical model. The general-purpose scaffolding may be extensible to many scientific applications, including precision radial velocities, Doppler imaging, chemical abundances for Galactic archeology, line veiling, magnetic fields, and remote sensing. Its sparse-matrix architecture and GPU acceleration make *blasé* fast. The open-source PyTorch-based code *blasé* includes tutorials, Application Programming Interface documentation, and more. We show how the tool fits into the existing Python spectroscopy ecosystem, demonstrate a range of astrophysical applications, and discuss limitations and future extensions.

Unified Astronomy Thesaurus concepts: High resolution spectroscopy (2096); Stellar spectral lines (1630); Astronomy data modeling (1859); GPU computing (1969); Calibration (2179); Radial velocity (1332); Maximum likelihood estimation (1901); Deconvolution (1910); Stellar photospheres (1237); Atomic spectroscopy (2099); Brown dwarfs (185); Earth atmosphere (437)

1. Introduction

1.1. Spectral Fitting Past and Present

Tens of thousands of individual spectral lines or more give rise to a sea of undulations that imbue each stellar spectrum with its characteristic appearance. The identification and understanding of these lines have defined a large category of astrophysics over the past century. The field grew from by-eye catalogs of stellar templates (Cannon & Pickering 1901), to quantifying the role of atomic ionization balance (Payne 1925), to modern synthetic forward models including millions or billions of lines (e.g., Husser et al. 2013; Marley et al. 2021; van den Bekerom 2021). As technology has improved, our data and models have become more voluminous, more precise, and more complicated. The mere act of comparing models to observed spectra can now resemble a computational statistics challenge as much as a scientific one. Here we introduce a new machine-learning-based framework, *blasé*, aimed at solving computational, statistical, and scientific challenges associated with data–model comparisons for modern astronomical spectroscopy.

* Open-source code at <https://github.com/gully/blase>.

 Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

The metaphorical holy grail of astronomical spectroscopy would be a function that takes in an observed stellar spectrum and reports back the position, amplitude, width, and shape of all of its spectral lines, automatically, accurately, and precisely. The function would go further. It would report back the systemic radial velocity (RV) and rotational broadening ($v \sin i$) and fundamental stellar properties, including T_{eff} , $\log g$, and $[\text{Fe}/\text{H}]$. Finally, the function would—in a feat of artificial intelligence—provide what it believes to be the interpretable generating function that produced these data in the first place, so that we may gain insights on future examples of this or other stars. Solving this problem is hard, for at least four reasons. First, the spectral lines may overlap, and so the assignment of one line may be partially degenerate with the assignment of some other adjacent line. Second, extremely wide line wings blend into the continuum, such that the placement of the continuum level may become ill-defined. Third, the extent of line blending and realized line shape depend strongly on the spectral resolution of the spectrograph, the rotational broadening of the star, and possibly the instrumental configuration at the time of observation. Finally, telluric absorption lines commingle with the astronomical spectral lines of interest, censoring some spectral regions entirely, or partially confounding other lines with chance alignments.

Addressing these and other challenges forms the backbone of *spectral calibration*, an increasingly valuable specialty as the deficits in our models become intolerable with greater data

quantity and quality. Luckily, many scientific applications in astrophysics do not need the technically demanding noise-free template, nor a catalog of all spectral lines. A few lines suffice. For those applications, human inspection of isolated lines and semiautomated equivalent width (EW) determination have been—and will remain—adequate.

But many new and important questions in the fields of stars and exoplanets aspire to reach the margins of what the entire data set can inform. In particular, data from high-*grasp* échelle spectrographs possess simultaneously high spectral resolving power and high bandwidth, yielding tens of thousands or possibly millions of independent spectral resolution elements for each star, substar, or exoplanet. Scientific applications that seek to gain signal by “stacking” spectral lines or cross-correlating with templates can hypothetically gain huge boosts in the accessible signal-to-noise ratio (S/N) compared to a single line or a few lines. Most manual and semiautomated methods cannot take advantage of the entire spectral bandwidth, or rely on exact knowledge of the underlying templates, and may fail to achieve the hypothetical promise of these high-bandwidth spectrographs (Hood et al. 2020).

For example, exoplanet cross-correlation spectroscopy (Snellen et al. 2010; Brogi et al. 2012; Birkby et al. 2013) hinges on accurate molecular spectral templates to detect and characterize the atmospheres of exoplanets. Imperfections in these templates can mute the perceived signal strength of these atmospheric features (Hoeijmakers et al. 2015).

In extreme-precision RV (EPRV) applications, cross-correlation methods work (Dumusque 2018) but have many limitations (Zhao et al. 2022). Among the many such limitations, one pernicious noise floor stands out as enigmatic: telluric mitigation. Many practitioners today simply mask these telluric regions, yet microtellurics still inject variance into the spectrum that cannot be easily accounted for with existing methods. Instead, a robust accounting of telluric absorption at the centimeter-per-second level may require joint modeling of the star and Earth’s atmospheric absorption *before* convolution with an instrumental kernel. This telluric joint modeling capability does not yet exist at a precision that can meet these strenuous demands.

In the case of Doppler imaging, an accurate underlying spectral template is needed to detect longitudinally symmetric structures (Vogt & Penrod 1983; Luger et al. 2021a) such as polar spots (Roettenbacher et al. 2016) or zonal bands (Crossfield et al. 2014; Apai et al. 2021). There exists a nearly circular reasoning: we need to know the extent of line profile perturbations to reveal the underlying spectral template, but we need the underlying spectral template to estimate the line profile perturbations. The approaches introduced here offer a path forward on these long-standing friction points.

1.2. Automatic Differentiation Technology

Existing open-source frameworks have overcome some of these challenges, or have been purpose-built for specialized applications. These frameworks include ROBOSPECT (Waters & Hollek 2013), specmatch (Petigura 2015), specmatch-emp (Yee et al. 2017), wobble (Bedell et al. 2019), starfish (Czekala et al. 2015), sick (Casey 2016), psoap (Czekala et al. 2017), FAL (P. Cargile et al. 2022, in preparation), CHIMERA (Line et al. 2015), the Cannon (Ho et al. 2017), MOOG (Sneden et al. 2012), MOOGStokes (Deen 2013), MINEsweeper (Cargile et al. 2020), and

recently ExoJAX (Kawahara et al. 2022). The designs of these frameworks necessarily have to make a choice in the bias–variance trade-off: is the tool more *data* driven or more *model* driven? The statistical trade-off can be viewed as a concession in physical self-consistency for model flexibility: more or fewer parameters; more *accurate* or more *precise*.

A key new enabling technology breaks these classical trade-offs in data–model comparisons for astronomical spectroscopy. Automatic differentiation (“autodiff” or “auto-grad;” Gunes Baydin et al. 2015; Maclaurin 2016) and its affiliated back-propagation algorithm (Kelley 1960; Linnainmaa 1976; Rumelhart et al. 1986; Dreyfus 1990) have revolutionized machine learning and neural network architecture design and are increasingly applied in astrophysical data analysis contexts, e.g., kernel phase coronography with morphine (Pope et al. 2021; Wong et al. 2021) and ∂ lux (Desdoigts et al. 2022), stellar surface modeling with starry (Luger et al. 2021b), and exoplanet orbit fitting with exoplanet (Foreman-Mackey et al. 2021). Of the spectroscopy frameworks mentioned above, the TensorFlow-based (Abadi et al. 2016) wobble and the JAX-based (Bradbury et al. 2018) ExoJAX employ autodiff technology. wobble treats each pixel as a tunable control point, producing $\sim 10^5$ parameters for a modern stellar spectrum. The ExoJAX framework has only ~ 100 tunable parameters that describe the fundamental physical properties controlling a stratified brown dwarf atmosphere. These two autodiff-aware frameworks span the extreme ends of nonparametric and parametric modeling for spectroscopy.

In this paper, we show that autodiff-aware semiempirical models offer an appealing middle ground: informed by self-consistent models but refined with data. This sweet spot in the bias–variance trade-off can be thought of as a hybrid data-and-model-driven approach. The algorithm presented here focuses on modeling the spectra of stars and brown dwarfs. Existing stellar models (e.g., Husser et al. 2013) and substellar models (e.g., Marley et al. 2021) laboriously solve for a self-consistent thermal structure in the atmosphere given the copious opacity sources that themselves depend on temperature and pressure. Here we build on that hard work by cloning preexisting synthetic stellar or substellar models (Section 2) and optionally by cloning models of Earth’s atmospheric “telluric” absorption (Section 3). We introduce the interpretable forward-model design and its PyTorch-based (Paszke et al. 2019) implementation, blase. In Section 5 we describe how to adapt both stellar and telluric cloned models simultaneously, using a transfer-learning step. We obtain semiempirical models by comparing to real-world échelle data in Section 6. Finally, we discuss perspectives on how to think of *blasé* (Section 7) and chronicle many conceivable extensions for unlocking new science (Section 8).

2. Methodology I: Cloning Stellar Spectra

Our first stage is to *clone* a precomputed synthetic spectrum. Cloning means to mimic, emulate, or otherwise approximate the appearance of a discretely sampled spectrum with some function or combination of functions. The subject of the cloning shall be a noise-free spectrum generated from physics-based atmosphere models. The purpose of cloning may appear superficial at first glance, but it serves as a necessary gateway for our ultimate data–model comparison goals.

2.1. Overall Architecture and Design Choices

We start with a high-resolution precomputed synthetic stellar or substellar model spectrum, $S_{\text{abs}}(\lambda)$, at its native-resolution sampling and with its original absolute physical flux units. The procedure that follows is largely agnostic to the exact details of how this spectrum was made, or what physics or chemistry it may represent. For the purposes of this paper, we will showcase examples from two well-known families of precomputed synthetic astronomical spectra: PHOENIX (Husser et al. 2013) for stellar spectra ($T_{\text{eff}} \in [2300, 10,000]$ K) and Sonora (Marley et al. 2021) for substellar spectra ($T_{\text{eff}} \in [200, 2300]$ K). The algorithms in the framework may also work for precomputed synthetic spectra of reflected-light exoplanets, supernovae, galaxies, or even further afield, such as laboratory physical chemistry, plasma physics, materials science, or remote sensing.

We place mild demands on the precomputed spectra. They should have sporadic regions of discernible continuum (or pseudo-continuum) devoid of lines, and the continuum shape should vary smoothly in wavelength. The spectral lines or pseudo-lines should be resolved (and not subsampled). Many precomputed synthetic stellar spectra will meet these criteria. For those spectra that do not meet these criteria, the method should still work with some additional fine-tuning of the preprocessing steps that follow. M dwarfs in the red-optical, solar-like stars in the blue visible, and brown dwarfs at virtually any wavelength may fall into this category.

We truncate the red and blue limits of the precomputed synthetic spectrum to match a high-bandwidth echelle spectrograph, extended with a buffer at the edges of size $\pm \Delta\lambda_{\text{buffer}}$. The buffers are needed for two reasons: (1) to include the influence of spectral lines whose centers lie just outside the region of interest, but whose wings are broad enough to affect the neighboring pixels; and (2) to account for astrophysical RV shifts that may send line centers into and out of the region of interest. We set the value of the buffers to account for plausible RV and rotational broadening of real stars. A generous buffer of $v \sin i < 500 \text{ km s}^{-1}$ and $|\text{RV}| < 500 \text{ km s}^{-1}$ yields a typical buffer of about 30 Å.

The choice of limiting the bandwidth to a region of interest around a single echelle spectrograph bandwidth stems from computational constraints. In principle, there is no fundamental limit to the bandwidth one could clone with the method presented here, up to and including the entire precomputed synthetic spectral model bandwidth. We adopt the exact native wavelength sampling with no smoothing or interpolation, yielding a wavelength vector λ_S with length N_S equal to the number of pixels within the extents of our region of interest including the buffers.

At this stage, we have the choice of whether to work in linear or log scale flux units. Adopting the log of the flux would ensure that the cloned model possesses only positive flux values, a desirable trait of any physical spectral model. We have implemented both modes in `blase`, allowing users to choose their preference. We only narrate the linear flux unit description in the main text of this document for the sake of clarity, and since most practitioners may tend to think of flux in terms of linear flux units. The data–model comparison step will always take place in linear flux units, so the only operational difference is the behavior for deep and saturated lines. Appendix A lists the equations adjusted to log flux units.

2.2. Initialization

We initialize the cloned model with a series of preprocessing steps. We divide the entire spectrum by a blackbody $B(\lambda_S)$ of the same effective temperature T_{eff} as the model template. The resulting signal usually still has smooth wiggles around the continuum. An optional continuum flattening step ensures that subsequent spectral line-finding steps get applied uniformly. This high-pass filtering step should be set to capture the genuine spectral shape, without overfitting broad line wings such as those in deep hydrogen and sodium lines. Any high-pass filter will work; a Gaussian process (GP) approach would be ideal (Czekala et al. 2015). Instead, we apply a simple and familiar heuristic: fit a polynomial $P(\lambda_S)$ to a few continuum peaks and divide out the trend.

The result should be a flattened “continuum-normalized” spectrum familiar to practitioners in high-resolution spectroscopy, with the continuum level close to unity. It is this spectrum that will serve as the centerpiece of subsequent training steps. Therefore, we drop any subscript and simply refer to this flattened spectrum as S :

$$S = S_{\text{abs}} / B / P, \quad (1)$$

where the division indicates element-wise division of these arrays or “vectors” of flux values. For element-wise multiplication, we will adopt the \odot operator symbol to distinguish that the multiplication occurs element-wise between two “vectors” of spectral fluxes, as opposed to a dot product (scalar output) or matrix multiplication (matrix output). This style of vector multiplication and division is standard in numerical array manipulations such as those in `numpy`, `IDL`, `matlab`, etc.

A re-creation of the unvarnished input spectrum—if desired—can be obtained by multiplying the continuum-flattened signal by the “perturbed blackbody,” $B(\lambda_S) \odot P(\lambda_S)$, that symbolizes the blackbody modulated by continuum opacity or broadband radiative transfer effects. This smooth spectrum may be useful for applications that need to keep track of broadband flux, such as low-resolution spectra, or regions with molecular band heads. The “perturbed blackbody” continuum model contains $n_{\text{poly}} + 1$ fixed but possibly tunable lookup parameters, plus the fixed input T_{eff} . For most practitioners, these terms serve as nuisance parameters and are perfunctorily discarded.

Next, we identify the spectral lines. We apply the `find_peaks` local-maximum-finding algorithm implemented in `SciPy` (Virtanen et al. 2020) on the negative of the spectrum, in order to find the local minima. We select a “topographic prominence threshold” $P_{\text{rom}} \in (0.005, 0.02)$, defined as the vertical distance between the peak and the local baseline. The baseline is defined with topographic heuristics that act as effective pseudo-continuum detectors. This threshold dictates the number of lines that will be modeled: a lower prominence finds more, weaker lines, and a larger prominence finds fewer, deeper lines. The prominence algorithm successfully finds lines that reside on top of broad line wings, or unresolved band heads provided that the individual lines exceed the prominence threshold in their local region. The number of lines N_{lines} depends on the bandwidth, the prominence, and the intrinsic properties of the input spectrum, principally effective temperature and metallicity.

For this paper, we illustrate examples for two échelle spectrographs with particularly large spectral grasp: the Habitable-zone Planet Finder (HPF; Mahadevan et al. 2014) on the Hobby–Eberly Telescope at McDonald Observatory in Fort Davis, Texas, and the Immersion Grating Infrared Spectrograph (IGRINS; Park et al. 2014) currently on the Gemini South Telescope on Cerro Pachón in Chile. The $R = 55,000$ HPF has a native bandwidth of 8079–12785 Å, which we expand to 8049–12815 Å including the edge buffers. IGRINS has two cameras for H and K band, with the combined spectrum spanning 14267–25217 Å including the edge buffers and the region in between the two cameras, all at a resolving power of $R = 45,000$. The spectrograph acquisition, reduction, and post-processing steps yield data $D(\lambda_D)$, where λ_D is the wavelength vector at the instrumental resolution and sampling of each instrument, generally much coarser than the resolution and sampling of the precomputed synthetic spectra. The data wavelength vector may also contain gaps between échelle orders, whereas the precomputed wavelength coordinates are usually contiguous. HPF may have up to $2048 \times 28 = 57,344$ pixels, and IGRINS has typically about 75,000 pixels, after common trimming of noisy edge pixels and unusable telluric regions. Meanwhile, the HPF-truncated model spectra have $N_s = 335,849$ native-resolution samples, comparable to the IGRINS-truncated model spectra, $N_s = 330,052$.

Figure 1 shows how the number of detected lines N_{lines} scales with effective temperature and prominence threshold P_{rom} for the PHOENIX grid, truncated to the bandwidths plus buffers for HPF and IGRINS. We see between about 2000 and 20,000 lines depending on the T_{eff} and P_{rom} . HPF and IGRINS have a comparable number of lines, and halving the prominence increases the number of lines by about 20%–30% in these ranges. The number of lines monotonically increases toward cooler effective temperatures.

So far we have only one piece of information about the spectral lines: their location. Next, we derive coarse properties about each detected peak: its amplitude and width, again using the prominence algorithms implemented in `scipy` (Virtanen et al. 2020).

There does not exist a general-purpose, single-shot algorithm for obtaining the line shape in the presence of overlapping spectral lines: where do the wings of one line begin and the wings of another adjacent line end? Therefore, we do not attempt to determine anything about the line shape at this stage and instead assume that the lines resemble a Voigt profile, with a guess width about equally split between Lorentzian and Gaussian.

2.3. The Blasé Stellar Clone Model

We have now arrived at the *blasé* clone model $S_{\text{clone}}(\lambda_S)$ for a flattened synthetic spectrum S : it is the cumulative product of transmission through the sea of all overlapping spectral lines:

$$S_{\text{clone}} = \prod_{j=1}^{N_{\text{lines}}} 1 - a_j V_j, \quad (2)$$

where V_j is the Voigt profile $V(\lambda_S - \lambda_{c,j}, \sigma_j, \gamma_j)$ with Gaussian standard deviation σ , Lorentzian half-width γ , at line center position λ_c , for the j th spectral line. The amplitude a is always expected to be positive for absorption lines.

The Voigt profile $V(\lambda, \sigma_j, \gamma_j)$ can be computed in exact closed form using the Voigt–Hjerting function (Hjerting 1938)

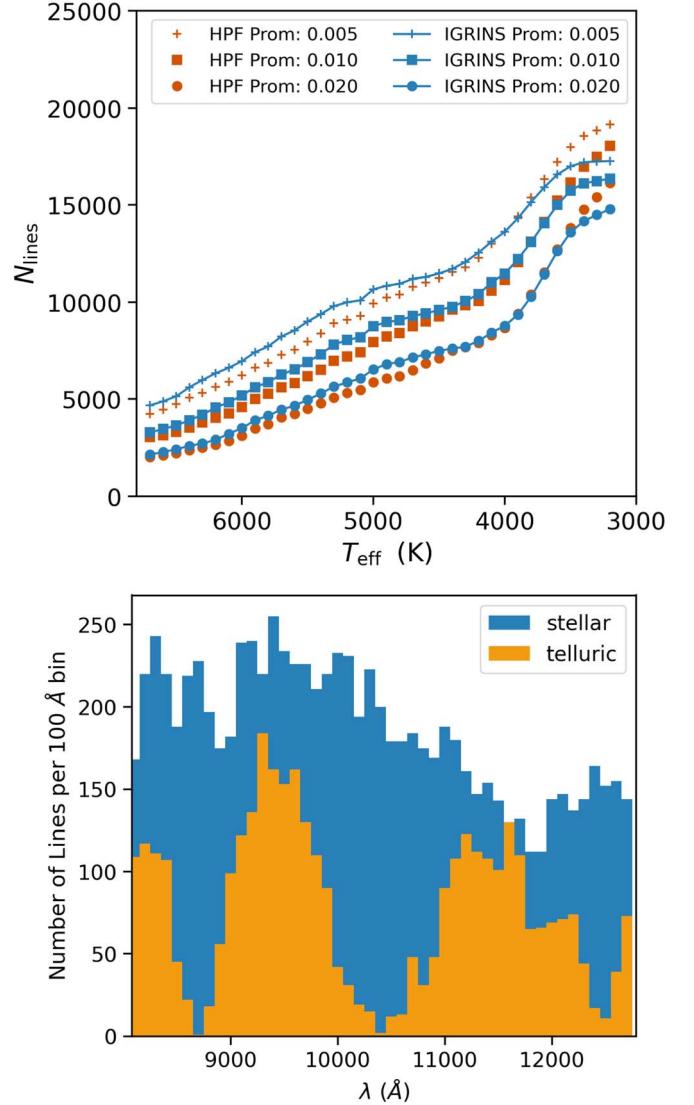


Figure 1. Scaling of prominence and density of spectral lines. Top: number of lines vs. effective temperature for PHOENIX models truncated to IGRINS (blue connected points) and HPF (red free-standing points) bandwidths, for different prominence thresholds of 0.02, 0.01, and 0.005. Bottom: number of lines per 100 Å wavelength bin for stellar (blue, upper envelope of steps) and telluric (orange, lower envelope of steps), illustrated for a $T_{\text{eff}} = 4700$ K, $\log g = 4.5$ PHOENIX model and a $T = 290$ K, relative humidity of 40% TelFit model.

as the real part of the complex Faddeeva function (e.g., Zaghloul & Ali 2011). Evaluation of the Faddeeva function can be computationally costly, so approximate forms may be desirable. Here we adopt the pseudo-Voigt approximation (Ida et al. 2000). Cloning is the stellar spectrum is step 1 out of 4 in Figure 2.

2.4. Goodness-of-fit Metric and Gradients

The model evaluated with its coarse initial values would have terrible performance: it would only vaguely resemble the synthetic spectral model, with up to $\pm 50\%$ undulations from the inexact assignment of widths, line shapes, and amplitudes. Instead, we tune the parameters of the model, starting from these coarse initial values. This model has between $N_{\text{lines}} \times 1$ and $N_{\text{lines}} \times 4$ free parameters, depending on how many of the four physical line properties you wish to tune. The threshold-

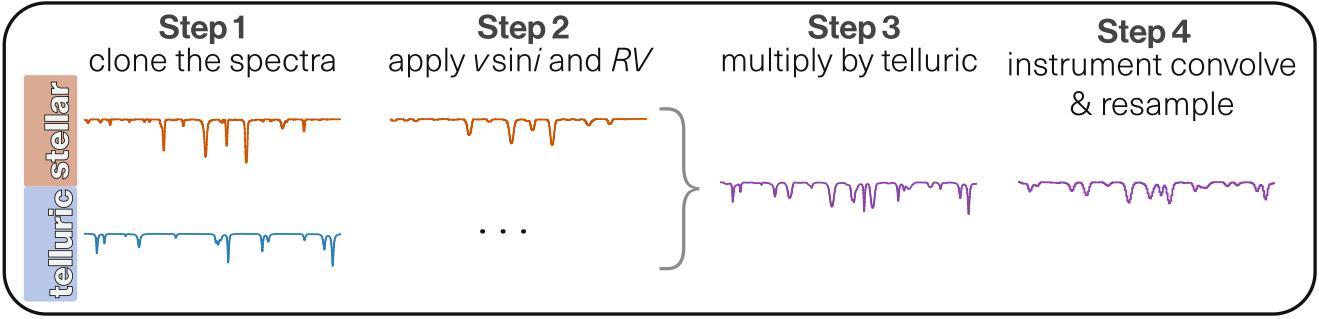


Figure 2. Visual flowchart of the *blasé* forward model. Step 0 (not shown) is to choose a precomputed synthetic stellar spectrum—and, optionally, a precomputed synthetic telluric spectrum—with physical properties close to the target and observing conditions. Both the stellar and telluric spectra get cloned (Step 1). The stellar model is warped to its extrinsic properties (Step 2), and then the stellar and telluric models get multiplied together (Step 3). Next, this joint model is convolved with an instrumental kernel and resampled to the wavelength coordinates of the data spectrum (Step 4). It is this forward model that gets directly compared to the observed spectrum (not shown).

based line-finding algorithm performs very well at finding the line center positions, so in practice we can fix the line center wavelengths. Therefore, we default to fitting three parameters per line, where the center wavelength is held fixed and the amplitude, width, and line shape are allowed to vary. We minimize a scalar “goodness-of-fit” metric, aka loss scalar \mathcal{L} , chosen as the mean squared error (MSE), which is proportional to χ^2 , the sum of the squares of the residual vector $\mathbf{R} \equiv \mathbf{S} - \mathbf{S}_{\text{clone}}$, but has no notion of per-pixel noise since the precomputed synthetic spectrum has no uncertainty:

$$\mathcal{L} = \sum_i^{N_s} (S_i - S_{\text{clone},i})^2 = \mathbf{R}^\top \cdot \mathbf{R}. \quad (3)$$

As seen in Figure 1, the number of lines can exceed 7000, meaning that the clone model has over 21,000 free parameters. Fitting more than about 300 parameters is difficult with conventional optimizers, and so historically the practice of finding the best settings for all 21,000 parameters would have been considered hopeless, a fool’s errand. A new category of “gradient-based” optimizers offers a breakthrough. Gradient-based optimizers have virtually no restriction on the number of parameters, scaling to functions of possibly millions or billions of tuning parameters. Gradient-based optimizers get vastly more information per likelihood evaluation: $N_{\text{params}} + 1$ pieces of information—the scalar value of the loss evaluated for some setting, and the derivative of the loss scalar with respect to each of the parameters, the so-called *Jacobian*, $\vec{\nabla} \mathcal{L} = \left(\frac{\partial \mathcal{L}}{\partial a_j}, \frac{\partial \mathcal{L}}{\partial \sigma_j}, \dots, \frac{\partial \mathcal{L}}{\partial \gamma_j} \right)$. The Jacobian indicates how the MSE would decrease with a change in the parameter of interest, or put simply, “which way and by how much” you have to change each individual line property to get a better fit.

Geometrically, the $\vec{\nabla} \mathcal{L}$ Jacobian resembles a compass needle guiding the direction to the best fit. This compass needle resides in, say, 21,000-dimensional space, instead of our mere three dimensions of physical space. Conventional gradient-free optimizers evaluate the Jacobian through approximate finite differences, which result in numerical instabilities. Those instabilities accrue and break down at around the 300-dimension mark. Hence, non-gradient-based optimization can be pictured as having only a coarse needle—assembled from approximate finite differences—that begins to spin aimlessly after too many dimensions.

We adopt “full-batch” gradient descent optimization, as opposed to the more popular stochastic gradient descent (SGD; Ruder 2016). In full-batch gradient descent, the entire data set is evaluated in Equation (3), as opposed to only a portion of the data set in the stochastic counterpart. Appendix C.3 discusses the trade-offs and rationale for “full-batch.”

The optimizer updates the a_j , σ_j , γ_j parameters by a small fraction of the Jacobian—called the learning rate (LR)—toward the direction that would improve the fit, for all parameters *simultaneously*. The Jacobian is calculated behind the scenes with automatic differentiation implemented as the so-called back-propagation algorithm, or simply “backprop” (Gunes Baydin et al. 2015). We choose the PyTorch framework that computes these Jacobians efficiently for all of the mathematical primitives in our *blasé* implementation (Paszke et al. 2019).

It is this ability to automatically compute Jacobians that sets PyTorch (and JAX and TensorFlow) apart. These frameworks give exact Jacobians, instantaneously, for free (or cheap). Without exact Jacobians, a conventional optimization step only obtains *one* piece of information: how much the overall loss changed. With exact Jacobians, we obtain $N_{\text{params}} + 1$ pieces of information for each evaluation of the forward model. The power of gradient descent becomes transformative as the number of parameters grows into the tens of thousands, as is the case for *blasé*.

2.5. GPU and Autodiff Specific Considerations

Forward modeling with tens of thousands of physics-informed parameters may seem like such a significant paradigm shift that it can feel too good to be true. In this section, we introduce the nonnegligible architectural design trade-offs that arise when adopting an autodiff framework—such as PyTorch—for physics-based forward modeling.

First, we make a few tweaks to the implementation for numerical purposes. We enforce that all Gaussian and Lorentzian widths are positive by tuning the natural log of the widths and then exponentiating them before inclusion in Equation (2). Based on the minus sign in Equation (2), an amplitude a_j could hypothetically take on either positive values (flux loss, absorption) or negative values (flux gain, emission). For now, we focus on photospheric absorption lines—and therefore enforce all the amplitudes to be positive by tuning $\ln a_j$ and then exponentiating in the same strategy as above. Emission lines—if desired—could be included with a mere sign flip to

Equation (2). Lines that can manifest in either absorption or emission could hypothetically relax the natural log preprocessing step for isolated lines. We narrate only the absorption scenario moving forward.

The autodiff machinery has a convenient way to set which parameters are held fixed and which are iteratively fine-tuned. One simply disables the autodiff flag for the fixed parameters: we set the `requires_grad=True` property for any PyTorch tensor that we want to vary. This design allows us to easily explore whether, say, allowing the λ_c parameter to vary significantly improves the fit.

The computational bottleneck occurs at the evaluation of Equation (2), which can be viewed as having an $N_{\text{lines}} \times N_S$ matrix \bar{F} assembled by stacking each Voigt absorption profile $V_j(\lambda_s)$ on top of each other:

$$\begin{pmatrix} 1 - a_1 V_1(\lambda_s) \\ 1 - a_2 V_2(\lambda_s) \\ \vdots \\ 1 - a_{N_{\text{lines}}} V_{N_{\text{lines}}}(\lambda_s) \end{pmatrix}. \quad (4)$$

An element of this matrix, F_{ji} , will have the flux value for a given j th line at a given i th wavelength coordinate. Equation (2) performs a type of matrix contraction, turning an $N_{\text{lines}} \times N_S$ matrix into a length N_S row vector. The number of Floating Point Operations (FLOPS) scales with the number of entries in this matrix. Thus, we face a trade-off of wanting to make the matrix large for accuracy and small for computational expedience.

We can rewrite Equation (2) as a sum by taking the log of both sides and dropping in this \bar{F} matrix:

$$\ln S_{\text{clone}} = \sum_{j=1}^{N_{\text{lines}}} \ln F_{ji} = \mathbf{1} \cdot \ln \bar{F}, \quad (5)$$

where $\mathbf{1}$ is a $1 \times N_{\text{lines}}$ row vector of all ones. We reemphasize that—in its current form—each spectral line has to be painstakingly evaluated across the entire spectral bandwidth. Efficient GPU algorithms exist for voluminous matrix manipulations such as this one, so this voluminous computation will proceed as quickly as possible on modern machines. In particular, the proprietary CUDA architecture for NVIDIA® GPUs contains Tensor cores with specialized matrix math. The chief bottleneck occurs when the storage of the \bar{F} matrix exceeds the available RAM of a GPU or CPU: the computation will fail with an “Out of Memory” exception. Modern NVIDIA GPUs have 8–40 GB of RAM, which translates roughly to a few thousand spectral lines across $\sim 300,000$ pixels. It is generally not possible to construct Equation (5) in its entirety in one fell swoop, even on a GPU. A remedy is needed.

2.6. Sparsity

The $\ln \bar{F}$ matrix is sparsely populated: most of the entries far from the line center are vanishingly close to zero. Here we take advantage of that mostly empty matrix using the mathematics of sparse matrices (Saad 2003).

We retain a relatively small number of pixels N_{cut} adjacent to the line center. Setting this wing cut produces a speedup by a factor of $\frac{N_S}{N_{\text{cut}}}$, which can exceed 100× for wide-bandwidth spectra. The choice of N_{cut} is nuanced. It should be set large enough that truncation effects are not seen for the broadest

lines. But even more, N_{cut} has to be future-proofed for Doppler shifting. Extreme Doppler shifts could hypothetically send line centers entirely outside the extents of N_{cut} if set too low. We therefore typically set wing cuts comparable to the buffer size $2\Delta\lambda_{\text{buffer}}$, even though most weak lines only perceptibly affect $< 1 \text{ \AA}$. We coerce all wing cuts to be the same number of pixels, typically 6000 pixels, $\sim 30\text{--}60 \text{ \AA}$ for PHOENIX, with the middle pixel being at the line center position, and about 3000 pixels to the red and blue side of the line. We populate a new approximate sparse matrix $\ln \hat{F}$ with only these 6000 pixels per line and assume zeros everywhere else.

The remapping of the sparse matrix can be pictured as having shifted all lines to the center of this new matrix \hat{F} , visualized pictorially in Equation (6). The algorithmic machinery keeps track of each (i, j, F_{ji}) trio of coordinates and flux values.

$$\bar{F} = \begin{pmatrix} \dots & \text{---} & \text{---} & \dots \\ \dots & \text{---} & \text{---} & \dots \\ \vdots & & & \vdots \\ \dots & \text{---} & \text{---} & \dots \\ \dots & \text{---} & \text{---} & \dots \end{pmatrix} \quad (6)$$

$$\bar{F} \mapsto \hat{F} = \begin{pmatrix} \dots & \text{---} & \text{---} & \dots \\ \dots & \text{---} & \text{---} & \dots \\ \vdots & & & \vdots \\ \dots & \text{---} & \text{---} & \dots \\ \dots & \text{---} & \text{---} & \dots \end{pmatrix}$$

Sparse-matrix methods generally support an operation known as *coalescing*, which sums values with repeated indices. Each pixel may get computed about ~ 100 times in this sparse implementation, which is about 70× faster than each pixel getting computed $N_{\text{lines}} \sim 7000$ times in the dense approach. Efficient algorithms for assembling and coalescing sparse matrices exist in PyTorch. Some GPUs now support additional hardware acceleration of sparse matrices, providing even greater speedups.

2.7. Optimization and Training

We use the Adam optimizer (Kingma & Ba 2014) with a typical learning rate $\text{LR} \in (0.005, 0.1)$ and all the defaults for PyTorch v1.11. We defined the number of training epochs $N_{\text{epoch}} = 100\text{--}10,000$ depending on the application. The user can optionally monitor a live view of the training progress with Tensorboard (Abadi et al. 2016) to gain an intuition for the training efficiency.

Figure 3 shows a portion of a PHOENIX spectrum cloned with `blase`. The 1000 epochs of training took 56 s on an NVIDIA® RTX2070 GPU with PyTorch v1.11, CUDA v11.1, and Intel® Core™ i7-9750H CPUs at 2.60 GHz, with all tensors as FP64. The same computation on a 2020 M1

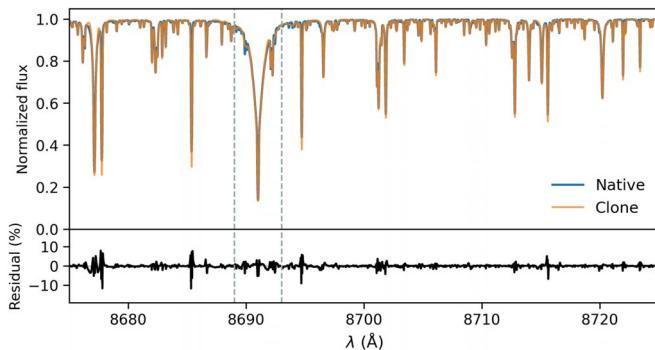


Figure 3. PHOENIX spectrum cloned with *blasé*. This $T_{\text{eff}} = 4700$ K, $\log g = 4.5$ solar-metallicity model has 9028 individual cloned spectral lines, each with three tuned parameters. The pictured 50 Å chunk contains 121 spectral lines and represents about 1% of the entire spectral bandwidth that was cloned. Some flaws can be seen near the cores of deep lines or wings of broad lines.

MacBook Air took 1^h25^m with PyTorch v1.9, 90× slower than the GPU counterpart.

We store the model parameters to disk and refer to the entire collection of parameters as a *pretrained model*. More specifically, this fine-tuned model represents an evaluable and interpretable clone of the original static pixel-by-pixel flux values.

3. Methodology II: Cloning Telluric Spectra

Ground-based near-IR échelle spectra possess thousands of depressions attributable to molecular line absorption in Earth’s atmosphere. These telluric lines hamper the unbiased interpretation of échelle spectra, so some treatment plan is needed. Often the regions of known, deep tellurics are simply discarded. In other cases, the lines are modeled with first-principles line-by-line radiative transfer (e.g., TelFit; Clough et al. 2005; Gullikson et al. 2014) or through data-driven means (e.g., wobble; Bedell et al. 2019). The most demanding EPRV applications require a precision characterization of telluric lines that the astronomical community has not yet been able to achieve and that may rival even the abilities of Earth Science practitioners. A hybrid data/model-driven approach was among the chief recommendations of the *Telluric Hack Week Workshop*¹ aimed at improving mitigation of the atmosphere’s deleterious effects. The *blasé* framework achieves a key milestone by introducing a *hybrid* approach to tellurics.

3.1. The Blasé Telluric Clone Model

We start with a precomputed synthetic telluric model, T , with associated wavelength coordinates λ_T . We employ a TelFit model, though any precomputed synthetic telluric model will work, such as MOLECFIT (Smette et al. 2015). The TelFit model does not contain any continuum sources of opacity, so we can skip the flattening procedure described in Equation (1). We orchestrate the same initialization and line finding as in the stellar models and obtain a coarse clone.

The number of pixels in the telluric model can be chosen at the time of running TelFit. Here we choose a spectral resolution $R \sim 10^6$, adequate for resolving narrow telluric lines, and yielding about 2 million pixels across the entire HPF

bandwidth. This pixel sampling is about 6× finer than the native PHOENIX pixel sampling. The number of telluric lines depends on the atmospheric properties, in particular, the local surface temperature T_{\oplus} and relative humidity RH. For a surface temperature of $T = 290$ K (62°F), relative humidity of 40%, and typical conditions for McDonald Observatory, we anticipate 3615 telluric lines across the entire HPF bandwidth, distributed as shown in the bottom panel of Figure 1.

One guiding principle departs from the stellar case: telluric lines do not require future-proofing for large RV shifts, hypothetically allowing us to reduce the number of pixels needed for a wing cut. Small RV shifts are possible owing to bulk motions in Earth’s atmosphere, but those bulk motions should be much smaller than the speeds of stars toward and away from Earth. Hence, we can hypothetically tolerate a much smaller wing cut for telluric lines. In practice, telluric lines can be saturated, and accurately cloning the resulting broad telluric wings still benefits from $N_{\text{cut}} \sim 6000$ pixels, comparable to the stellar scenario.

We optimize the sparse telluric clone, achieving comparable computational speed as the PHOENIX cloning task. We are left with $T_{\text{clone}}(\lambda_T)$, the tunable telluric clone model evaluated at its original native coordinates. Figure 4 shows a before-and-after view of the *blasé* cloning procedure for a TelFit model zoomed in on a *J*-band region with dense telluric lines. We see that the *blasé* fine-tuning appears nearly indistinguishable from the native-resolution TelFit precomputed model.

4. Cloning Performance

We compute the residual $R(\lambda_S)$ of the native-minus-cloned PHOENIX model, illustrated in the bottom panel of Figure 3. We see an rms residual of 1.2% pixel⁻¹ at native resolution. The telluric clone shows a comparable level of performance. These residuals tend to pile up in local symmetrically balanced clusters that get canceled out once convolved with coarser instrumental line profiles, so their overall effect at instrumental resolution is typically negligible: the clones are almost perfect. We identify three main categories of cloning flaws that may not be negligible depending on the science application.

The first—and expected—source of large residuals is simply missing line opacity owing to our finite prominence threshold. Lines with prominence less than P_{rom} yield residual notches with strengths comparable to P_{rom} . Including smaller prominence lines by lowering P_{rom} produces smaller residuals, at the trade-off of computing more lines and yielding higher computational cost. But at some point, turning down P_{rom} yields diminishing returns, as other imperfections provide a noise floor. We have experimentally determined this noise floor to occur near $P_{\text{rom}} = 0.01$.

Second, a conspicuous flaw occurs in the line cores of relatively narrow lines. The cloned model tends to overestimate the flux at the core and underestimate the flux along the slopes of the lines. The residuals, therefore, exhibit a W-shape ringing artifact. The scientific impact of this flaw will depend on the precision demands of the application, but it may not be as bad as it seems for many applications. Recall that the cloning occurs at native spectral resolution, much greater than typical instrumental resolution. The W-shape ringing residuals will tend to cancel each other out once the clone model undergoes convolution with an instrumental profile (Section 5.2). We ascribe this flaw to the adoption of approximate pseudo-Voigt profiles rather than exact Voigt profiles. High-precision

¹ <https://speakerdeck.com/dwhgg/telluric-line-hack-week-wrap-up>

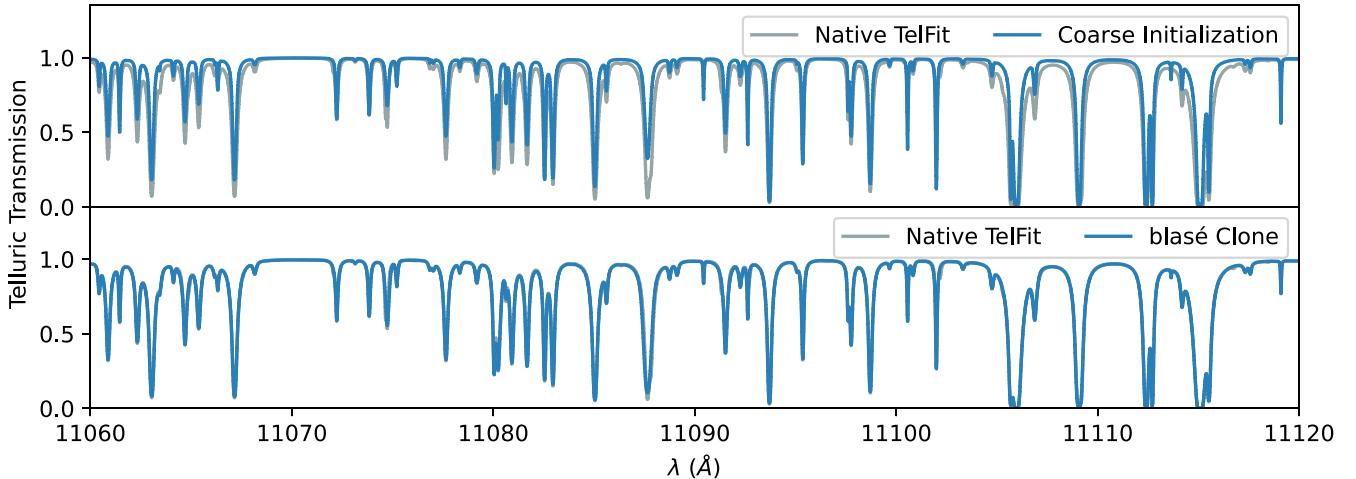


Figure 4. Precomputed TelFit model cloned with `blasé`. Before: the top panel shows a forward model evaluated with coarsely initialized line-by-line properties, based on simple threshold-based peak finding and crude estimation of the shape of the line. After: the bottom panel shows the `blasé` fine-tuned clone, with nearly identical spectral structure.

applications may wish to employ the exact Voigt profile, which we estimate to cost about $27\times$ more than this approximate form.

The overprediction of line cores in narrow lines is only partly attributable to the pseudo-Voigt approximation. Even exact Voigt profiles inadequately represent true atmospheric opacities. Typical stellar and substellar atmosphere models compute opacities for a stratified atmosphere consisting of tens or hundreds of locally isothermal, isobaric atmospheric layers. Each of those layers may exhibit a local per-line opacity that matches a Voigt profile. But the cumulative product of many slightly different Voigt profiles is not exactly equal to any single Voigt profile. Hence, the `blasé` approach of asserting lines as Voigt profiles can be best understood as a phenomenological model—that lines tend to look Voigt-like, even if they are non-Voigt in detail. In principle, `blasé` could be configured in a way to learn the requisite perturbations to the Voigt profile in a systematic manner, pooling information among lines. Such an approach is not yet implemented and may be an enticing avenue for future research.

Finally, and most perniciously, a large category of residuals appear near the wings of the deepest and broadest lines—such as hydrogen lines and neutral alkali metal lines. The true lines exhibit advanced line shapes, such as non-Lorentzian line wings that are not captured with the overly simplistic Voigt line profile, even an exact one. Figure 5 highlights super-Lorentzian line wings around a line at 8691 \AA . Nearby narrow lines devolve into missing line wing opacity, the favored trade-off when the continuum estimate’s poor performance outweighs the pain of a narrow but tolerably small spike. This flaw can be seen where a line initialized at 8692.5 \AA and another pair of lines at 8690.0 all melt into line wings.

Cloning telluric lines suffers from one additional problem. Telluric lines can be extremely deep, exhibiting almost vanishing transmission with saturated line cores common in between the atmospheric windows that define the *I*, *J*, *H*, and *K* bands. The `blasé` method can cope with these saturated lines, but often it treats nearby and blended saturated lines as one single line. This glomming together of lines has little practical effect since few astronomical practitioners can make use of such profoundly saturated data.

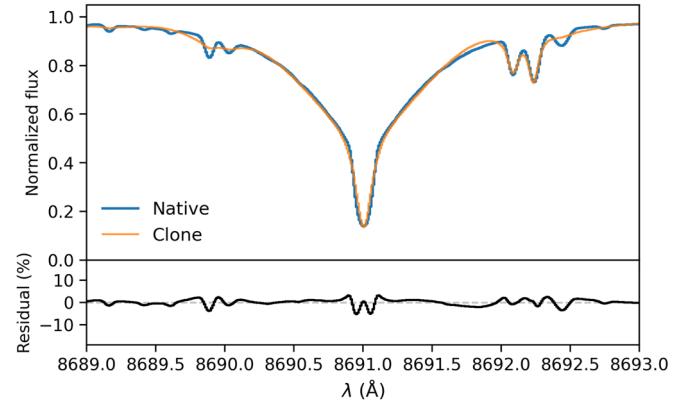


Figure 5. Zoom-in of the region between the vertical gray bars in Figure 3. The cloned model has seven spectral lines describing the 400 pixels in this 4 \AA chunk. The native PHOENIX pixel sampling can be seen as the boxy steps in both the native and cloned models.

5. Semiempirical Models with Transfer-learning Techniques

The cloned model already represents a useful intermediate product: the distillation of $N_s = 335,849 \times 2$ pixel flux values and their wavelength coordinates into a more compact quartet of properties for a list of $<10,000$ spectral lines, a dimensionality reduction of $18\times$ for the cost of 1.2% pixel $^{-1}$ in accuracy. Hence, the process so far can be myopically viewed as a physics-informed compression algorithm. But the cloned model serves as a mere stepping stone in our principal quest: the comparison of models to real data.

5.1. Augmenting the Stellar Clone with Radial Velocity and Rotational Broadening

Real stars possess two key extrinsic properties. Rotational broadening $v \sin i$ and RV depend on the observer’s viewing location. We follow Czekala et al. (2015) by emphasizing the qualifier “extrinsic,” to distinguish from stellar *intrinsic* properties, such as T_{eff} , $\log g$, and $[\text{Fe}/\text{H}]$. Intrinsic properties appear the same from any viewing location—at least for stars with isotropic surfaces—while extrinsic properties do not. The distinction is important because the extrinsic terms act as

simple convolutions and translations to the cloned spectrum and can be treated after the cloning procedure. Therefore, we define an augmented model, which we designate the “extrinsic” model, S_{ext} :

$$S_{\text{ext}}(\lambda_Z) = S_{\text{clone}}\left(\lambda_Z - \frac{RV}{c}\lambda_c\right) * \zeta\left(\frac{v}{v \sin i}\right), \quad (7)$$

where ζ is the convolution kernel for rigid-body rotation (e.g., Kawahara et al. 2022), v is the spectral axis represented as relative velocity coordinates, and $*$ denotes the convolution operator.

Most autodiff-aware convolve operators act in pixel space, approximating kernels as numerically sampled functions. There exists a special exponential spectral sample spacing that allows the convolution operators to work out of the box for rotational broadening. The design of *blasé* permits the stellar and telluric models to be reevaluated at any wavelength coordinate vector, provided that it adequately samples the underlying lines. We therefore change the sampling from the native stellar and telluric wavelength coordinate grids, λ_S and λ_T , to this special exponentially sampled wavelength grid, denoted with the subscript Z :

$$\lambda_Z = \lambda_0 \exp\left\{\frac{v - v_0}{c}\right\}, \quad (8)$$

where $v - v_0$ is the velocity vector going from zero to the velocity associated with the largest wavelength, with *linear*-spaced velocity samples. We choose a sampling in velocity space of 0.5 km s^{-1} , which corresponds to about $10 \times$ finer than the instrumental resolving power of HPF and delivers a minimum and maximum wavelength spacing of 0.013 and $0.024 \text{ \AA pixel}^{-1}$, respectively, for the HPF bandwidth.

Operationally, the RV shift gets applied to the line center positions rather than scaling the entire wavelength grid point coordinates, λ_Z . This choice yields a convenience: it cleanly makes the RV autodiff-aware, meaning that an infinitesimal change to the RV value can be sensed through back-propagation by affecting only the line center positions.

The S_{ext} spectrum is shown in Step 2 of Figure 2. There is currently no equivalent post-processing of the telluric spectrum T . As mentioned previously, we assume that the motions of Earth’s atmosphere are much less than the desired stellar RV precisions. However, demanding EPRV exoplanet searches may need to consider minuscule systematic RV shifts and broadening of the telluric templates, arising from the turbulent and bulk motions of Earth’s atmosphere. A T_{ext} could hypothetically be included in *blasé* to achieve these strenuous precision demands.

5.2. Joint Stellar and Telluric Model

Figure 2 shows a visual guide to all the steps in *blasé*. We have arrived at what may be the most intriguing and yet simple of these steps: we simply multiply the rotationally broadened and RV-shifted stellar model by the telluric transmission:

$$M_{\text{joint}} = S_{\text{ext}}(\lambda_Z) \odot T_{\text{clone}}(\lambda_Z). \quad (9)$$

It is only at this stage that we may apply the instrumental broadening kernel. The instrumental resolving power, R , acts as a convolution with a Gaussian line profile of width $\sigma = \frac{c}{2.355k}$. Real astronomical instruments usually have wavelength-dependent resolving power, which complicates the implementation for

high-grasp spectra. The extent to which this effect matters will depend on the science application. For now, *blasé* simply assumes a fixed resolving power.

Notice that the order of steps 3 and 4 in Figure 2 cannot be swapped. Mathematically speaking, element-wise multiplication and convolution do not commute. While the distinction may seem negligible, it matters at the level of precisions sought in EPRV applications (S. Mahadevan et al. 2022, private communication). Water vapor lines in our own atmosphere can “beat” with water vapor in the spectrum of, say, an M-dwarf atmosphere. The systematic telluric miscancellation would imbue a Moiré pattern of residuals that is most acute for sources with sharp lines, namely low projected rotational broadening ($v \sin i \sim \frac{c}{2R}$). The approach in *blasé* may therefore unlock a level of telluric calibration that has evaded previous efforts.

As noted, typical data-pixel sampling λ_D is much coarser than the model pixel sampling λ_Z . We therefore resample the model to the data spectrum in the following way. We evaluate the joint model M_{joint} at all of the super-resolution wavelength coordinates and then compute the mean value of those pixels within the bounds of each coarse data pixel. The resampling procedure is autodiff-aware: the same clusters of high-resolution coordinates map to the same data-pixel coordinates, no matter what the RV is. The RV only dictates what flux values are realized within those pixel bounds.

The final forward model for *blasé* is designated simply as M without subscripts to emphasize that we have achieved the desired goal of a plausible end-to-end physics-informed yet highly flexible forward model for each datum in the 1D observed spectrum:

$$M(\lambda_D) = \text{resample}[M_{\text{joint}}(\lambda_Z) * g(R)], \quad (10)$$

where g is the Gaussian instrumental convolution kernel and the `resample[]` operation indicates the average of model pixels that fall within each data pixel’s red and blue boundaries.

5.3. Regularization

Equation (10) has $\sim 21,000$ tunable parameters from the star, ~ 9000 tunable parameters from Earth’s atmosphere, plus $v \sin i$ and RV. That adds up to about 30,002 model parameters. The resolving power may also be treated as tunable if it is not known or varies slightly with, e.g., seeing, slit-or-fiber illumination, or instrumental configuration: 30,003. The pseudo-continuum polynomial term P can also be tuned to refine the continuum placement initialized at the preprocessing stage, yielding an additional 5–15 parameters, depending on the choice of polynomial order.

It may appear desirable to simply optimize all of these parameters in a laissez-faire manner, allowing them all to take on whatever value the data dictate. Such a stratagem would overfit the data, resulting in unphysically perverse line shapes that do not reflect the air of reality we aspire to impose on our synthetic spectral models. Lines would haphazardly fit noise spikes and conspire together to warp spectral shapes in unexpected ways. This overfit model may suit some rare purposes. But most of the time, we prefer to strike a better balance in the bias–variance trade-off.

We apply some amount of *regularization*, a restriction on the allowed values the model parameters can take on. Fortunately, we have a firm theoretical basis to justify this regularization.

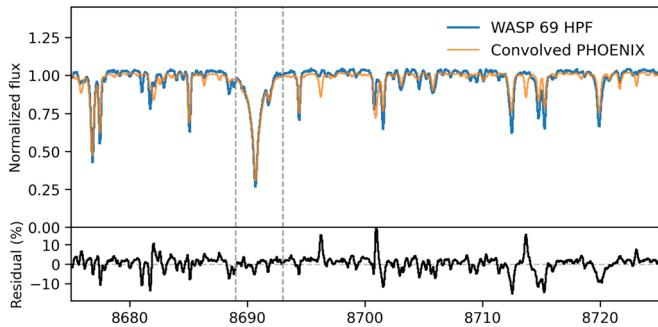


Figure 6. WASP 69 observed with HPF, compared to a $T_{\text{eff}} = 4700$ K, $\log g = 4.5$ solar-metallicity PHOENIX model warped to $v \sin i = 2.2$ km s $^{-1}$, $\text{RV} = -9.6$ km s $^{-1}$, and HPF resolving power.

We believe that our precomputed synthetic spectral models are *quasi-statically correct*: the predicted spectra resemble the unobserved “True spectrum” with lines in the correct place, but just with the wrong area under the curve. This statement may stem from the fact that it is easier to predict the mere existence of some energy transition of atoms and molecules than it is to predict their transition rates, abundances, temperature and pressure effects, and all the other line strength effects that flow down to how much light a line ultimately absorbs in a stellar atmosphere.

The degree of regularization will control the extent of overfitting or underfitting. The most extreme regularization—the antithesis of the *laissez-faire* scenario—would yield a model too rigid to respond to the data at all, yielding a model entirely unchanged from the cloned PHOENIX and TelFit models, the extreme end of underfitting. Hence, regularization constitutes the only hyperparameters worthy of tuning in *blasé*. The choice of how to set the regularization is problem specific. We default to the following choice. We fix all line parameters except for amplitude, which receives an L1 loss—namely, we penalize the absolute value of departures from an amplitude’s starting place:

$$\mathcal{L}_{\text{reg}} \equiv \sum_{j=1}^{N_{\text{lines}}} \left| \frac{\ln a_j - \ln \hat{a}_j}{\Lambda} \right|, \quad (11)$$

where the hat notation demarcates the amplitudes obtained from cloning, i.e., the “initial, theory-inspired amplitude.” We assign $\Lambda \sim 5$. The total, overall loss then becomes

$$\mathcal{L}_{\text{tot}} = \mathcal{L}_{\text{MSE}} + \mathcal{L}_{\text{reg}}. \quad (12)$$

This weak degree of regularization has the effect of permitting refinement of only the most conspicuous data-model mismatches; the weakest lines do not bother to move from their initial state because doing so would penalize the regularization without enough reduction of the overall loss. The extrinsic $v \sin i$ and RV have no regularization, but in practice they need to be initialized close to their plausible values.

6. Results: Comparison to Data

6.1. WASP 69 with HPF

The planet-host star WASP 69 makes a great benchmark because it has a low $v \sin i = 2.2 \pm 0.4$ km s $^{-1}$ (Casasayas-Barris et al. 2017), making its lines sharp and easy to perceive. The K5 dwarf has an effective temperature of about

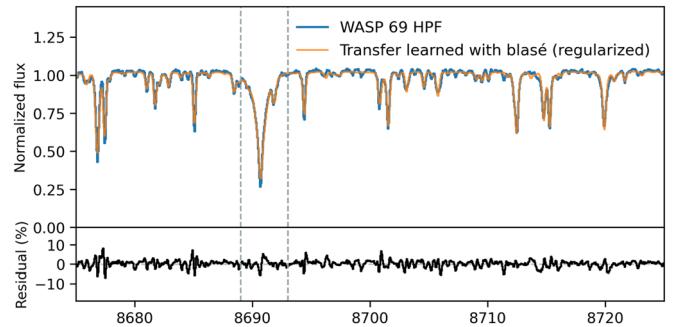


Figure 7. Semiempirical model of WASP 69 transfer-learned with *blasé*, employing a regularization prior on the learned amplitudes.

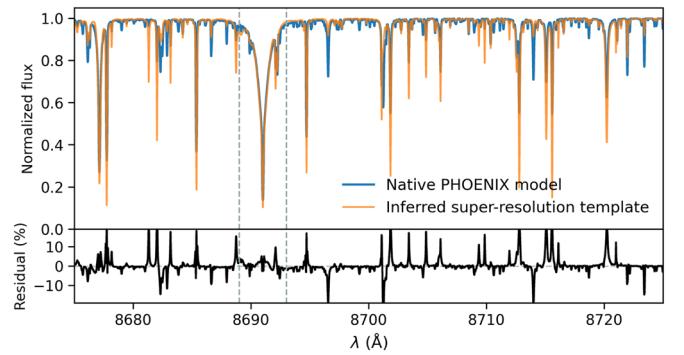


Figure 8. The native-resolution semiempirical model transfer-learned from WASP 69 HPF data. The revised model can be viewed as a super-resolution deconvolution of the HPF spectrum.

4700 K, $\log g = 4.535 \pm 0.023$, and slightly supersolar metallicity (Anderson et al. 2014).

Figure 6 shows a portion of an HPF spectrum of WASP 69 centered on a region devoid of telluric lines. This figure highlights a baseline case with no line-by-line fitting, simply conventional template matching: comparing HPF data of WASP 69 to the closest PHOENIX template, convolved and resampled to the HPF resolution and sampling. Conspicuous residuals of $\pm 10\%$ appear throughout the spectrum, with lines in the correct place, but with the amplitudes systematically biased. These line residuals arise from bona fide imperfections in the PHOENIX spectrum, with a minor contribution from the coarse sampling of the PHOENIX grid.

Figure 7 shows the same data spectrum compared to a pixel-level model trained using the *blasé* technique described in Section 5. The model fit appears much better, with typical residuals approaching the photon noise of the data themselves. The model is not perfect however, especially around line cores. These residuals stem from a combination of causes. First, our finite regularization restricts line amplitudes and widths from straying too far from their values. This computational tug-of-war makes line cores land just short of the values they otherwise would have obtained in the absence of regularization. Second, the spectral resolution kernel may get biased from the domineering model mis-specification of broad lines, setting up a slightly subpar performance for all the other lines.

The inferred, semiempirical high-resolution model for WASP 69 is therefore the transfer-learned model \hat{S}_{clone} , unadorned with the extrinsic and instrumental properties. This semiempirical model is shown in Figure 8. The

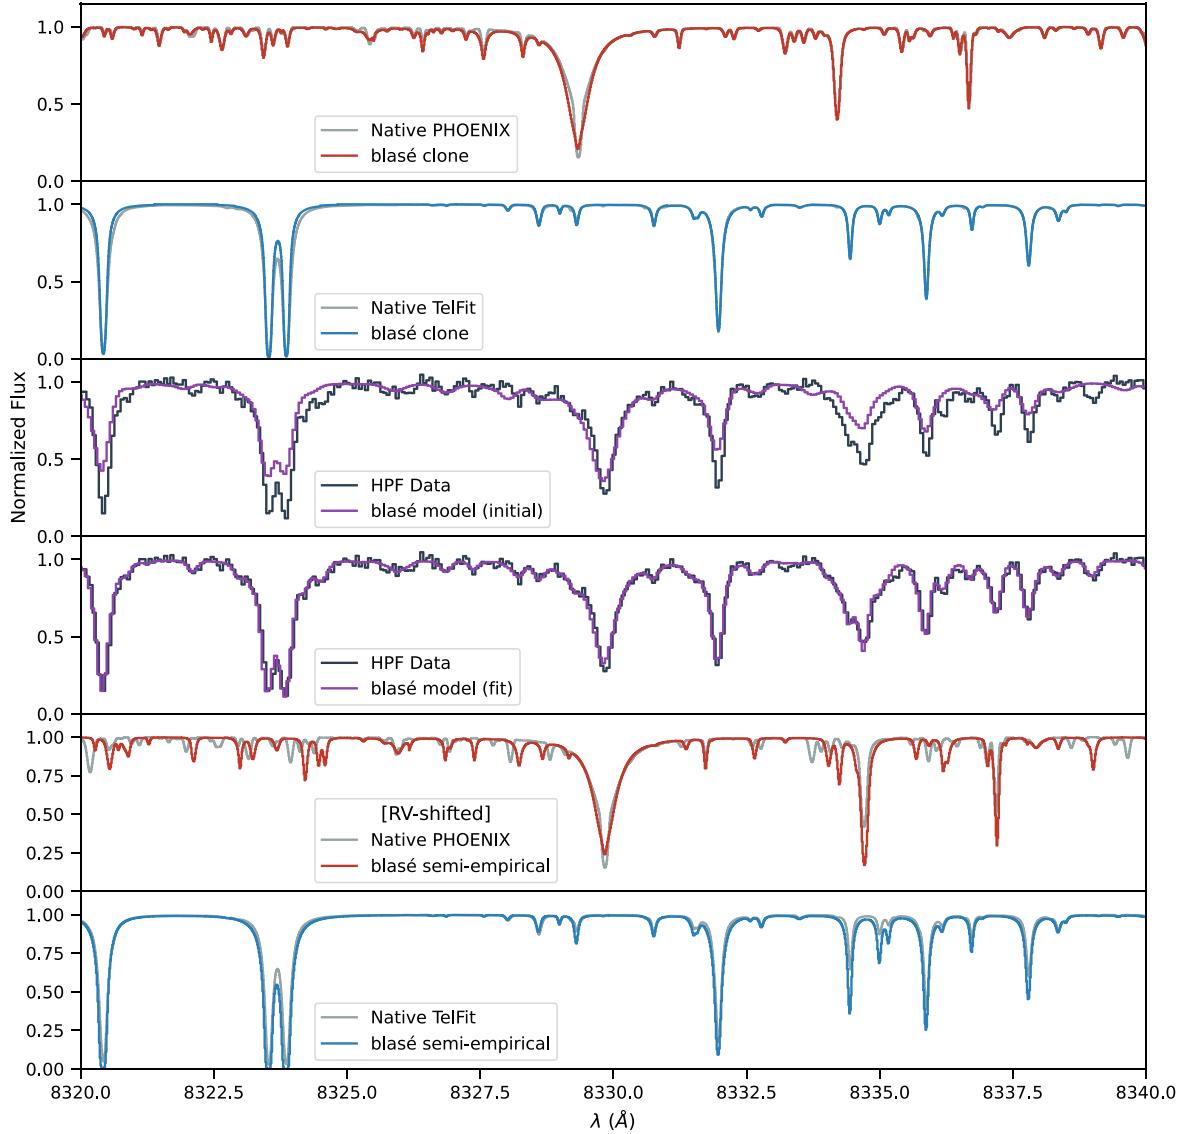


Figure 9. End-to-end training on a portion of the HPF spectrum of WASP 69. The top two panels show the cloning from Step 1 of the visual guide. The middle two panels depict the before-and-after of the instrument-convolved and resampled joint model illustrated in Step 4 of the visual guide. The final two panels show the underlying semiempirical models learned in the process.

departures from the PHOENIX model appear dramatic at this native resolution. The transfer-learned semiempirical template exhibits both deeper and shallower lines than the native PHOENIX model.

The low $v \sin i$ of WASP 69 means that the line profile broadening arises principally from the finite instrumental resolution. The interplay of telluric lines and stellar lines within an instrumental resolution element is exactly one of the challenges *blasé* was designed to solve, as discussed in Section 3. In Figure 9 we show a multipanel dissection of the HPF spectrum in a wavelength region in which conspicuous stellar and telluric lines coexist. The *blasé* end-to-end model exhibits excellent agreement with the data. The trustworthiness of the line properties inferred in locations where stellar and telluric lines exactly overlap remains an open research question: a tiebreaker is needed to distinguish these overlapping line inferences. We discuss conceivable tiebreakers later in the paper.

6.2. IGRINS Spectrum of a T6 Ultracool Dwarf

The T6 ultracool dwarf 2MASS J08173001–6155158 was recently observed with IGRINS, revealing a rich spectroscopic atlas of molecules in its $T_{\text{eff}} = 1060 \pm 50$ cloud-free atmosphere (Tannock et al. 2022). The sea of molecular lines is so rich as to blur entirely the notion of isolated lines. Many lines should instead be considered pseudo-lines. The boundaries of lines for 2MASS J08173001–6155158 become even more amorphous in the presence of its moderately high $v \sin i = 22.5 \pm 0.5 \text{ km s}^{-1}$ rotational broadening (Tannock et al. 2022). This T6 IGRINS spectrum therefore represents an extreme test case, in which the underlying notions of *blasé* go outside the comparatively safe assumptions in stellar spectra. The application of *blasé* to 2MASS J08173001–6155158 can be viewed as conducting a deconvolution step simultaneously paired with line-by-line inference.

We initialize *blasé* with the nearest Sonora template with $T_{\text{eff}} = 1100 \text{ K}$, $\log g = 5.0$, and solar metallicity

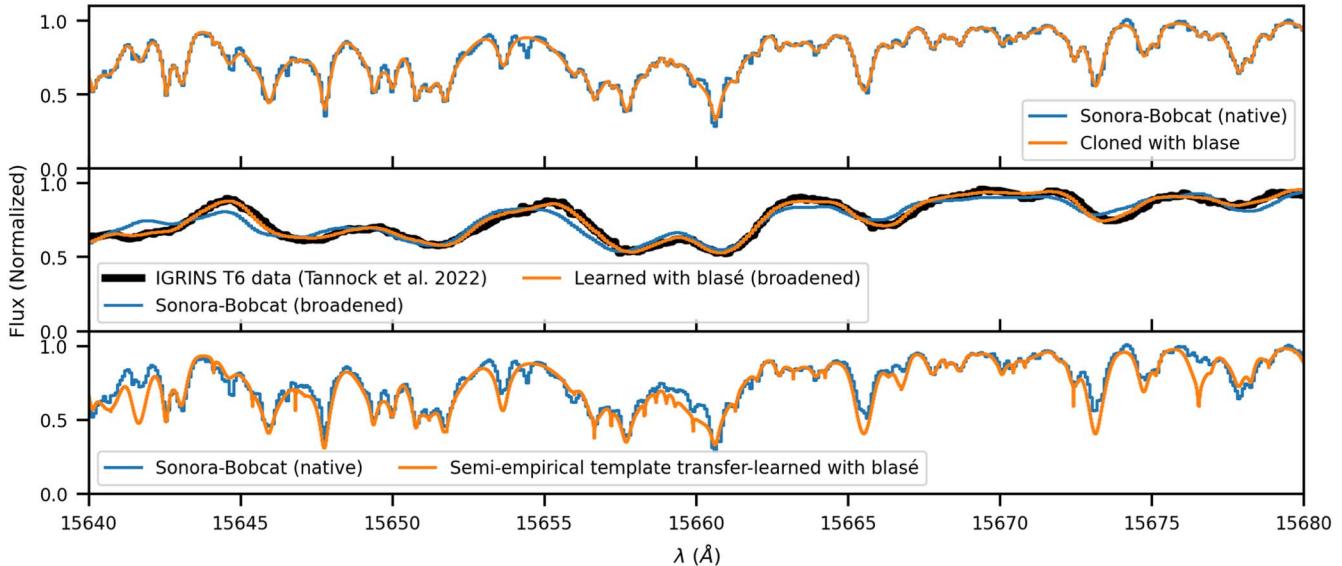


Figure 10. Line-by-line dissection of the 2MASS J08173001–6155158 IGRINS spectrum. The top panel compares the native Sonora–Bobcat spectrum to the *blasé* clone. We match the RV shift of the observed spectrum (thick black line) and convolve the initial and learned Sonora template to the large $v \sin i$. The bottom panel shows the locations and extents of line mismatches at the native resolution of Sonora.

(Tannock et al. 2022). Figure 10 shows the clone of this native spectrum near the peak of the H band. The cloning performance appears adequate to capture most of the molecular pseudo-lines, with a few notable shortcomings. The high S/N of this IGRINS spectrum lays bare the departures from the broadened Sonora–Bobcat model, albeit with surprisingly close agreement given the challenge of substellar atmosphere modeling. The line-by-line departures from Sonora exceed any minor flaws from the cloning process, indicating genuine opacity differences between the data and Sonora template. Our technique automatically detects and reports the underlying structure of data–model mismatches, a key milestone for the assembly of refined opacity tables in this cool dwarf temperature regime.

7. Discussion

7.1. How to Interpret Line-by-line Properties

Up to this point, we have presented the mechanics of *blasé* and have shown that it applies to a wide range of applications including stars, substars, and Earth’s atmosphere. The interpretation of the approach and its line-by-line outputs may carry different meaning to different practitioners. Here we caricature some of these perspectives:

“*Reverse-engineering the line lists.*”—Equation (2) (and its exponential twin Equation (A2)) could have been populated with initial line-by-line properties from atomic and molecular line lists, instead of the coarse detection-threshold approach we presented in Section 2.2. Hence, our cloning procedure can be viewed as reverse-engineering the quantum mechanical properties housed in line lists such as HITRAN (Gordon et al. 2022), HITEMP (Rothman et al. 2010), EXOMOL (Tennyson & Yurchenko 2012), VALD (Piskunov et al. 1995), or countless other primary source documents. This interpretation is only partly accurate. The line lists generally store the temperature and pressure scaling terms that get included in 1D thermal structure calculations with distinct Voigt coefficient terms computed for each layer in the stellar or substellar or Earth atmosphere. The cloned properties, therefore, represent the

flux-weighted average line property, having undergone an integral through the radiative transfer in the atmosphere.

Our deliberate choice not to initialize from the atomic and molecular lines means that we do not have to compute an expensive multilayer atmosphere at each inference step. The precomputed models have already taken care of that step, and so our cloning procedure encodes all the quantum mechanics and stellar atmosphere knowledge written down at the time of the computing of PHOENIX, Sonora, or LBLRTM.

“*Equivalent width machine.*”—The area under the curve of each j th *blasé* line represents its (neighbor-weighted) EW. That means that *blasé* can automatically and quickly distill a massive high-bandwidth spectrum into a catalog of line positions and EWs. Even more, the catalog would have some resilience to rotational broadening, line blending, and other friction points that hamper such an industrial-scale approach with existing tooling. Current applications of EWs rely on either isolated lines or—for multiphase applications—the assumption of nonvariable neighboring lines. *Blasé* does not have these limitations. The interpretations of these EWs may have to be treated with care if intercompared with EWs obtained conventionally. This approach may suit stellar abundance applications, such as Galactic archeology and chemically peculiar stars (e.g., HgMn stars; Chojnowski et al. 2020).

“*Fancy deconvolution.*”—The *blasé* approach can be viewed as a fancy deconvolution procedure, in which the convolution kernels get inferred alongside the imperfections of the model template. This perspective may suit practitioners who wish to obtain a robust rotational broadening estimate in the presence of imperfect templates, and vice versa: robust templates in the presence of imperfect rotational broadening.

“*Fancy interpolator.*”—A linear interpolator has to store the slope, offsets, and start and end points of all the line segments from pixel to pixel, merely to re-create a spectrum losslessly at some arbitrary coordinate or coordinates. *Blasé* losslessly recreates a spectrum, while having to store less information. In this way, it is fair to think of *blasé* as an interpolator that happens to have an intuitive interpretation.

“Surrogate modeling/simulation intelligence.”—Our slightly off-label use of machine learning straddles a boundary of machine learning and conventional physics-based models. The rapidly evolving field of machine learning may therefore classify the approach here using a few related terms. The phrases “surrogate modeling,” “emulation,” “discrepancy modeling,” “transfer learning,” “simulation intelligence,” and “physics-informed machine learning” could describe some aspects of *blasé*, depending on the ultimate application by the end-user. The surrogate modeling approach has been applied with mixed success for M dwarfs (Passegger et al. 2020) and brown dwarfs (Johnsen et al. 2020). The realized speedups in quantum chemistry can approach $300,000\times$ (Gilmer et al. 2017). The key idea is that information from expensive physics-based simulations can be distilled by machine learning and then applied cheaply to new examples. The performance of these surrogates should improve when more scientific domain knowledge is included at training time (Ansdel et al. 2018).

“Narrowing the synthetic gap.”—Our approach to pretraining on synthetic spectra and comparing to data resembles the technique of Passegger et al. (2020). There, a convolutional neural network is pretrained on PHOENIX spectra and then applied to CARMENES spectra of M dwarfs. Passegger et al. (2020) designate the difference in appearance between the synthetic spectra and observed spectra as the “synthetic gap,” showing that this gap leads to inaccurate metallicity inferences. *Blasé* shows a path toward shrinking the synthetic gap, with the prospect of retraining convolutional neural networks with semiempirical spectra produced via the method described here, assuming that high-fidelity metallicity labels can be obtained through co-natal binary studies.

“A good enough substrate.”—The imperfections in precomputed stellar and substellar models have hampered their use in high-fidelity applications, where model flaws overwhelm the S/N in the data. The adaptability of *blasé* can be viewed as increasing the number of applications that could conceivably benefit from the rich albeit imperfect information encoded in these model grid spectra. In the next section, we illustrate how these adaptable models can serve as a substrate for applications in which the imperfections in precomputed models would have otherwise precluded their use.

7.2. PyTorch and the Autodiff Ecosystem

One key limitation stems from the need for autodiff to “know about” or sense each step of the end-to-end computation. Practically speaking, this requirement means that *blasé* had to be written in a single differentiable (i.e., autodiff-aware) programming framework. This requirement is so strong as to forbid the intermingling of frameworks and languages. The choice of which machine-learning framework to adopt must therefore be made up front, with some risk of *vendor lock-in* once that choice has been set. We considered three such differentiable frameworks: TensorFlow, JAX, and PyTorch. We chose PyTorch because of its ease of use and its ability to structure differentiable models in an object-oriented style familiar within the Python community, with robust optimization support. PyTorch had already been applied to several astronomy projects by this time, including Exonet (Ansdel et al. 2018), PySR (Cranmer 2020), and MPoL (Czekala et al. 2021). Ultimately, the choice of PyTorch made it easy to employ sparse matrices.

This vendor lock-in gives rise to a frustrating incompatibility among exoplanet-relevant autodiff applications. For example, no differentiable piece of code in *blase* can call an exoplanet orbital RV, a starry surface map, an astropy spectral coordinate transformation, a wobble spectral residual model, or an ExoJAX Voigt profile. All of those tools use either Theano, TensorFlow, Numpy, or JAX, which are incompatible with PyTorch. We use SciPy, but only for preprocessing, before the differentiable model has been entered. Some heuristics exist for interchanging code among the frameworks, but they pose some implementation complexity and/or some performance hit.

It may appear attractive to coordinate the relatively small autodiff-for-astrophysics community to adopt a single autodifferentiable machine-learning framework to facilitate compatibility. In practice, such coordination may be ill-fated, since the rapid pace of machine-learning development has not yet settled on a single best way forward. Each framework offers some strengths and demerits, and exploring those trade-offs will continue to yield new experimental approaches. For example, JAX may offer speedups attributable to Just In Time (JIT) compilation, at the expense of a more rigid functional programming style. Access to scalable GP implementations may be another deciding factor (Aigrain & Foreman-Mackey 2022).

8. Unlocking New Science

The scaffolding of *blase* is designed in a way to promote extensibility, so many scientific questions can be written down in the flexible language presented in this paper. In some cases, the flexibility, precision, speed, and ease of use may unlock new approaches to long-standing astrophysical questions and practical challenges. Here we enumerate some planned or conceivable extensions. We break up the themes into two categories: scientific extensions (this section) and technical innovations (Appendix C).

8.1. Extreme-precision Radial Velocity

Already *blase* is equipped to fit every single line with its own systemic RV, by tuning the line center position λ_c at each training epoch. There exist both empirical evidence and some theoretical motivation that RV jitter varies from line to line (Dumusque 2018). Importantly, the extent of this line-by-line RV jitter could be predicted in part by the depth of line formation (Cretignier et al. 2021; Al Moulla et al. 2022). A future EPRV version of *blase* could leverage this information. As an example, the depth of formation for all (or a subset of) lines could be obtained and associated with each spectral line. A regularization scaling term could be introduced to allow the line positions to vary, but only in proportion to their depth of formation.

8.2. Line-by-line Fundamental Parameter Estimates

The conversion of spectral line strengths to spectral type (Cannon & Pickering 1901)—and by extension line strengths to T_{eff} (Payne 1925)—has occupied a large chunk of stellar spectroscopy in the past century. The measurement of fundamental stellar properties remains one of the principal applications of stellar spectroscopy. *Blasé* as it currently stands does not output fundamental properties, and in fact it takes them *for granted*: you must specify, e.g., the T_{eff} , $\log g$, $[\text{Fe}/\text{H}]$, and $[\alpha/\text{Fe}]$ of your PHOENIX template as Step 0 before that

template gets cloned and warped to match data. Hence, obtaining a T_{eff} estimate from *blasé* is impossible, at least as discussed so far.

The clearest way forward would be to calibrate the *blasé*-derived line properties. Simply put, we seek a set of functions, f , that relate the line properties to fundamental properties:

$$f_j(a_j, \sigma_j, \gamma_j) \rightarrow (T_{\text{eff}}, \log g, [\text{Fe}/\text{H}]).$$

There are N_{lines} such functions—one for each spectral line—because each line has its own temperature, gravity, and metallicity dependence. For visualization simplicity purposes, we have assumed that the PHOENIX grid is merely three-dimensional ($T_{\text{eff}}, \log g, [\text{Fe}/\text{H}]$), but the same logic applies to 4D and higher, and it works for sparse and irregularly sampled grids too.

Fascinatingly, there are two related ways forward toward this function. They both involve ensembles of spectra—such as a spectral atlas, spectral sequence, or library—and they both involve first obtaining the inverse function:

$$g_j \equiv f_j^{-1} \quad (13)$$

$$g_j(T_{\text{eff}}, \log g, [\text{Fe}/\text{H}]) \rightarrow (a_j, \sigma_j, \gamma_j). \quad (14)$$

A purely model-based approach would start by *blasé*-cloning every single PHOENIX spectrum in the grid dimensions of interest, over the wavelength range of interest. The resulting product would be a line-by-line catalog of cloned spectral properties for each 3D coordinate on this grid. We could then assemble a heat map of how each j th line property changes across this heat map, g_j . In this way, we are reverse-engineering, e.g., the temperature dependence of each spectral line, as encoded by the PHOENIX atmosphere models. The final step could involve finding the nearest neighbor of each j th observed spectral line to each j th grid heat map point, that is, finding a way to invert the function g_j to get f_j . The information-weighted mean, median, or mode of these nearest neighbors could then be reported as a revised “best-fit T_{eff} ,” for instance.

A semiempirical approach could improve on this purely model-based approach. We know that the PHOENIX line depths, widths, and shapes are imperfect, and so this ($T_{\text{eff}}, \log g, [\text{Fe}/\text{H}]$) heat map will have large flaws in the line-by-line properties: the contours are systematically too bright or dark. There are many conceivable ways to quantify these heat map flaws. *Blasé* provides an expedient route. One can pull the heat map toward the locus of points established by running *blasé* on benchmark stars. The known T_{eff} of such systems would anchor the trend. We suspect that the directionality of the purely model-based heat map must be correct, and simply the slope, offset, and concavity of the trends may be wrong. With new information from each benchmark, the heat map would get lifted like a central tent pole propping up an undersupported tent.

The construction and calibration of these ensembles of spectra represent a tremendous amount of work beyond the scope of this paper. But its creation could yield an extremely precise, fast, interpretable, and reasonably accurate way to measure the properties of stars based solely on their high-bandwidth échelle spectra. The overall accuracy hinges on the accuracy of the stellar benchmark labels. The mechanics of this approach could be adapted to fit within existing pipelines such as SAPP (Gent et al. 2022).

There exists a corollary from the method described here to the approach of Czekala et al. (2015). There, a 3D heat map was created for relating the eigenweights of a Principal Component Analysis (PCA) basis to the stellar fundamental properties, for spectrum emulation purposes. Here, the 3D heat map is created on a line-by-line basis and is therefore interpretable. The PCA eigenspectra were generally uninterpretable, at least not easily. Hence, in the limiting case of obtaining a densely calibrated semiempirical heat map for all spectral lines across all grid dimensions, we will have achieved a spectral emulator that would obviate the need to clone spectra in the first place. We would simply start with this powerful line-by-line emulator as a forward model and go directly to the extrinsic warping and data–model comparison.

8.3. Abundances

Stellar abundance work involves precisely measuring line strengths (often reported as EWs) of different chemical constituents evinced in a star or substar’s photosphere. Relating those EWs to physics can be done in a few ways. Most easily, trends and patterns in the EWs can be assembled, and metal-rich and metal-poor clusters can be identified. There exists a precision/accuracy trade-off in the measurement process, usually stemming from the placement of the continuum, or assumptions about line blending. *Blasé* offers an immediate solution to these challenges since it starts from our best guess for how nearby lines may be shaping the continuum. *Blasé* appears to be a gateway to extremely fast and intervention-free industrial-scale abundances, potentially useful for large surveys like APOGEE (Majewski et al. 2017), Gaia-ESO (Gilmore et al. 2012), RAVE (Steinmetz et al. 2006), and more.

8.4. Identifying Missing Lines

So far we have not addressed the inevitable prospect that some spectral line or lines reside in the observed data spectrum but are absent entirely from the precomputed model. This “missing line” scenario arises from our incomplete and evolving knowledge of quantum mechanics and molecular chemistry. ROBOSPECT (Waters & Hollek 2013) and the Cannon (Ho et al. 2017) have data-driven ways to identify these lines. The Cannon can go further to indicate how these previously unknown lines correlate with the stellar labels, providing some physical understanding of their cause. *Blasé* as it currently stands has no mechanism for identifying these lines. Here we consider that there are N_{missing} lines with two scenarios: (I) the identification, line center location, and possibly other ancillary information about the lines have become newly available; or (II) the lines are truly anomalous and unknown.

For the first scenario, Voigt profiles could be instantiated at the newly established line center locations $\lambda_{c,m}$ of the previously unknown line. Hence, Equation (2) would become a product of the new and old physics:

$$S_{\text{clone}} = \left(\prod_{m=1}^{N_{\text{missing}}} 1 - a_m V_m \right) \left(\prod_{j=1}^{N_{\text{lines}}} 1 - a_j V_j \right). \quad (15)$$

Scenario (II)—truly unknown lines—would not have the luxury of knowing where to initialize the line center location, at least not based on precise theory. Hence, instead, an

educated guess for $\lambda_{c,m}$ could be made based on the data centroid.

In either case, the initialization of A_m , σ_m , and γ_m would be a matter of taste, and they would therefore be difficult to regularize, at least in the manner that we have described so far. For these reasons, an even better strategy may be to treat the data residuals not as lines but in some other basis. For example, it could be possible to model the *blasé* residual spectrum with, say, wobble, or with a GP. These questions remain an open research area.

8.5. Doppler Imaging

The fixed $v \sin i$ approximation breaks down for stars with large-scale surface features. Doppler imaging attempts to infer the surface map from the extent to which observed line profiles depart from a pristine rotational broadening kernel. This inference procedure suffers from a vast number of geometrical degeneracies but still provides useful constraints on stellar surfaces (Luger et al. 2021a). We emphasize a distinction between (a) longitudinally symmetric surface features and (b) longitudinally asymmetric surface features. Most RV practitioners think about the latter, since longitudinally asymmetric surface features imbue changing-in-time skewness to the line profiles, causing RV perturbations easily detectable in RV time series. These confound exoplanet searches.

Longitudinally symmetric surface features, on the other hand, do not change as the star rotates on its axis. The existence of these features manifests as static-in-time *kurtosis* of the spectral line. For example, a hypothetical non-emitting (black) polar starspot exhibits a deficit in flux at the line core, resulting in less zero-velocity flux than its homogeneous counterpart. A dark zonal band results in equal-sized bites out of the red and blue sides of the line.

It is easier to infer longitudinally asymmetric features than longitudinally symmetric ones, since we assume that we occasionally catch a glimpse of the spot-free limbs and their pristine line profiles. The latter requires exact knowledge of the underlying spectral template. Isolated, deep, well-calibrated spectral lines constitute the only practical scenario where exact knowledge can plausibly be claimed. Isolated spectral lines may be scarce or absent for M dwarfs and brown dwarfs where lines blend ostensibly in an inseparable way, confounding Doppler imaging.

Blasé offers a new approach to Doppler imaging that may overcome these historical limitations by simultaneously fitting both the imperfections in the underlying spectrum and its line profile perturbations. This approach is analogous to the linearized model in Luger et al. (2021a), but with the benefit of also handling nonlinear properties of the spectrum such as line widths, shapes, and locations, while also handling telluric contamination. Hypothetically the Luger et al. (2021a) approach could be partially absorbed into *blasé*, or vice versa, though such a merger may be complicated to implement.

8.6. Starspots and Magnetic Fields

One current assumption of starspot spectral decomposition is that the starspot spectrum itself resembles the stellar photosphere of a cooler star. This assumption appears adequate for detecting starspot spectra and measuring their physical properties (Gully-Santiago et al. 2017). But to second order, starspots

should exhibit some spectral peculiarities that make them depart from a “normal” stellar photosphere. We may be probing deeper into the photosphere, and so the lines may experience higher pressure, with slightly different line widths (O’Neal et al. 1996). Or maybe the finite convective velocity shift can be directly seen as systematic shifts of the spectral lines (as stated in Section 8.1).

Rather than applying a mixture model of fixed PHOENIX templates, one could adopt a mixture model of precloned *blasé* models. Then, the imperfections in the starspot spectrum can be learned alongside the filling factor of the starspots.

Relatedly, magnetic field information could be incorporated into *blasé* as line-by-line Landé g -factor “labels,” when available. The use of these labels could take a few different forms. The labels could simply serve as flags, to indicate that regularization should be weakened to allow these magnetic sensitive lines to wander farther from their naive nonmagnetic expectations. Or the labels could be used to infer the extent of Zeeman broadening or Zeeman splitting by parameterizing these operations as convolutions. The best choice for how to incorporate magnetic sensitivity into the flexibility of *blasé* remains an open research question.

8.7. Circumstellar Disk and Accretion Veiling

Circumstellar disk veiling suppresses the strengths of all spectral lines en masse, as the stellar photosphere gets outshined by a hot disk and/or envelope.

The bulk properties of that disk/envelope could be incorporated with a simple physical model: a blackbody of temperature T_{disk} and solid angle Ω_{disk} (Greene et al. 2018). Revealing the spectral shape of the veiling would require sufficiently high bandwidth spectra, such as Xshooter (Vernet et al. 2011) or possibly IGRINS. The extent of veiling would be fit alongside all the other stellar spectral lines. However, an MCMC approach may outperform *blasé* in accuracy, since the choice of picking an underlying spectral template is partially degenerate with the derived veiling.

Accretion veiling can be treated analogously, but with some more complications. Emission lines can be easily incorporated as described in Section 2.5. These emission lines could be initialized with a line list of hydrogen lines, forbidden lines, and other conspicuous features. The shapes of those lines can be affected by winds and other physical phenomena, yielding a variety of physically interpretable line shapes (Erkal et al. 2022). *Blasé* could incorporate those line shapes as templates shared among several lines originating from similar physical environments, but scaled and shifted based on the details of each line’s radiative transfer properties.

8.8. Sky Emission Lines and Wavelength Calibration

Night-sky emission lines—arising from OH, for example—add emission-line spikes on top of the astronomical target spectrum of interest. These lines get subtracted in most pipeline packages, but in some cases a more careful treatment may be desirable. In HPF, a sky reference fiber points toward blank sky, needing some careful calibration before subtraction from the target fiber (Gully-Santiago et al. 2022). Remote-sensing applications may wish to measure such emission lines as the primary science of interest. It would be straightforward for *blasé* to handle emission lines, following the augmentation for emission lines in Section 2.5. The main benefit over

conventional sky subtraction methods may be the ability to disentangle the underlying spectrum at super-resolution. A predictive model for sky emission could be built up over time, building data-driven corrections to theoretical models.

The well-known center wavelengths of these lines could be used for additional wavelength calibration, as is already done in IGRINS (Lee & Gullikson 2016). An imperfect instrumental wavelength calibration could be diagnosed based on the departure of the predicted versus realized wavelength center positions, enabling a relatively fast and intervention-free correction strategy. Hypothetically, a prescription for the wavelength calibration could be built directly into *blasé*, with the coefficients for a polynomial rerouting the realized wavelength center locations from their known wavelength centers.

9. Conclusions

We have introduced an interpretable machine-learning approach to forward-modeling stellar, substellar, and telluric spectroscopic data. The line-by-line approach relies on a key enabling technology, automatic differentiation, that allows a nearly unlimited number of spectral lines to be forward-modeled simultaneously. We initialize these lines to match precomputed synthetic stellar spectra, achieving excellent performance, and lending some confidence that the approach has a capacity to capture a tremendous amount of information at once.

We demo the framework on two sources: the K5 exoplanet host star WASP 69 using a precomputed PHOENIX model, and the T6 ultracool dwarf 2MASS J08173001–6155158 using a precomputed Sonora–Bobcat model. We discuss how *blasé* can be used to measure EWs for thousands of lines automatically, understand line lists, measure rotation rates, generate surrogate models, and construct semiempirical models. This tool could readily have applications across stellar and substellar astronomy, including for PRV work, stellar compositions, Doppler imaging, and stellar activity.

We thank the anonymous referee for comments that improved the paper. We thank Phill Cargile, Hajime Kawahara, and Dan Jaffe for preprint comments.

This material is based on work supported by the National Aeronautics and Space Administration under grant Nos. 80NSSC21K0650 for the NNH20ZDA001N-ADAP:D.2 program and 80NSSC20K0257 for the XRP program issued through the Science Mission Directorate.

We acknowledge the National Science Foundation, which supported the work presented here under grant No. 1910969.

These results are based on observations obtained with the Habitable-zone Planet Finder Spectrograph on the HET. The HPF team was supported by NSF grants AST-1006676, AST-1126413, AST-1310885, AST-1517592, AST-1310875, AST-1910954, AST-1907622, AST-1909506, ATI 2009889, and ATI-2009982 and the NASA Astrobiology Institute (NNA09-DA76A) in the pursuit of precision radial velocities in the NIR. The HPF team was also supported by the Heising-Simons Foundation via grant 2017-0494.

The Hobby–Eberly Telescope (HET) is a joint project of the University of Texas at Austin, the Pennsylvania State University, Ludwig-Maximilians-Universität München, and Georg-August-Universität Göttingen. The HET is named in

honor of its principal benefactors, William P. Hobby and Robert E. Eberly.

M.G.S. thanks Ian Czekala, Phill Cargile, Rodrigo Luger, Greg Herczeg, Will Best, Dan Foreman-Mackey, Megan Bedell, Mark Marley, Erwan Pannier, Dirk C. M. van den Bekerom, Dan Clemens, Dan Jaffe, Tom Greene, Adam Kraus, David Hogg, Greg Zeimann, Ben Montet, Christina Hedges, Brittany Miles, Arpita Roy, Ben Pope, and Kevin Gullikson for conversations and resources that shaped his thinking on spectral calibration and telluric mitigation.

Facility: HET (HPF).

Software: pandas (McKinney 2010), matplotlib (Hunter 2007), astropy (Astropy Collaboration et al. 2013, 2018), exoplanet (Foreman-Mackey et al. 2021), numpy (Harris et al. 2020), scipy (Virtanen et al. 2020), ipython (Pérez & Granger 2007), starfish (Czekala et al. 2015), seaborn (Waskom 2021), pytorch (Paszke et al. 2019), muler (Gully-Santiago et al. 2022).

Appendix A Log Flux Scaling Mode

Here we illustrate how *blasé* gets altered when applying the logarithmic flux preprocessing step. First, we compute the natural log of the flux directly on the precomputed synthetic spectrum in its absolute flux scaling and native pixel sampling:

$$\ln S = \ln S_{\text{abs}} - \ln B - P. \quad (\text{A1})$$

We simply “rebrand” P as residing in logarithmic flux units, and we disregard it since it is largely a nuisance parameter anyway. We then treat the *blasé* clone model as a sum of opacities, retaining the Voigt profile:

$$\ln S_{\text{clone}} = - \sum_{j=1}^{N_{\text{lines}}} a_j V_j. \quad (\text{A2})$$

Here the a_j values have also been slightly rebranded from their meaning in Equation (2). We still want to enforce only absorption lines—and not spurious emission lines—so we use the sample trick of sampling the a_j values in log and then exponentiating them to get guaranteed positive values. Note that Equations (2) and (A2) carry modified meanings for the Voigt profile. Specifically, Equation (2) can be viewed as the Taylor series expansion for Equation (A2) in the limit of small opacities:

$$e^{-a_j V_j} \approx (1 - a_j V_j). \quad (\text{A3})$$

Both equations are approximate. A real stellar atmosphere’s line shape arises from a sum of disparate Voigt profiles weighted along a nonuniform column of gas, whereas here we have assumed that the column of gas is approximated as a single uniform isothermal backlit layer. A sum of unlike-Voigt profiles is not exactly equal to any single Voigt profile. Theoreticians may resonate with this more “first principles” representation, while data practitioners may find Equation (2) more natural, so to some extent, the choice is a matter of taste. Table A1 summarizes the notation used throughout the paper for reference.

The sparse matrix gets rebranded as filled with opacity values, instead of log-fluxes, but operationally remains the same. All subsequent steps operate on the summed and exponentiated opacities, behaving identically to their linear counterparts. For example, we exponentiate before computing the residuals and data–model comparison, $R = e^{\ln S} - e^{\ln S_{\text{clone}}}$.

Table A1
Notation Used in This Paper

Symbol	Meaning
Spectra	
λ_s	Native wavelength coordinates of the precomputed stellar spectrum
λ_t	Native wavelength coordinates of the telluric spectrum
λ_d	Native wavelength coordinates of the data spectrum
S_{abs}	Flux values of the precomputed synthetic stellar spectral model λ_s
B	Blackbody of temperature T_{eff} to coarsely normalize S_{native}
P	Smooth polynomial to refine continuum normalization
S	Continuum-normalized augmentation of S_{abs}
T	Transmission values of the precomputed synthetic telluric model
D	The observed data spectrum flux values
ϵ	The estimated uncertainties in the data spectrum
S_{clone}	Evaluable and tunable cloned flux model of S
T_{clone}	Evaluable and tunable cloned transmission model of T
S_{ext}	An augmentation of S_{clone} with $v \sin i$ convolution and RV translation
M_{joint}	The joint stellar and telluric model: $S_{\text{ext}} \odot T_{\text{clone}}(\lambda_s)$
M	Joint model convolved with instrumental kernel and resampled to λ_d
R	The residual spectrum between a pair of inputs, e.g., $D - M$
v	The spectral coordinate axis λ expressed as a velocity difference
Line Properties	
$\lambda_{c,j}$	Line center position of the j th spectral line
a_j	Gaussian line profile amplitude of the j th spectral line
σ_j	Gaussian line profile scale of the j th spectral line
γ_j	Lorentzian line profile half-width of the j th spectral line
V_j	The Voigt profile of the j th spectral line
\bar{F}	The dense ($N_{\text{lines}} \times N_x$) matrix of all line fluxes stacked vertically
\hat{F}	The sparse ($N_{\text{lines}} \times N_{\text{sparse}}$) matrix of all line fluxes stacked vertically
ζ	The rotational broadening convolution kernel
g	The instrumental broadening convolution kernel, typically a Gaussian
Scalars	
N_{lines}	Number of spectral lines
N_x	Number of pixel coordinates in the precomputed spectrum λ_x
N_{sparse}	Number of nonzero pixels computed in the sparse implementation
$\pm \Delta \lambda_{\text{buffer}}$	Buffer exceeding the red and blue limits of the data spectrum
P_{rom}	The prominence threshold of spectral lines to include in cloning
$v \sin i$	Rotational broadening for stellar inclination i and equatorial velocity v
RV	Radial velocity of the star
R	Spectrograph resolving power $\lambda/\delta\lambda$
L	The loss scalar, usually the sum of the squares of the residuals
Operators	
$\text{resample}[F(\lambda_x)]$	The resample operator, takes in a flux spectrum F evaluated at λ_x coordinates and returns the mean flux within the pixel boundaries of coordinate λ_x
\circledast	The convolution operator
\odot	<i>Hadamar product</i> , an element-wise product of two same-length vectors

Appendix B

Comparison to Existing Spectroscopy Frameworks

Several astronomical spectral frameworks share similar aims to *blasé*. These existing frameworks will have enduring value for the wide range of problems in the field of stellar spectroscopy. Here we scrutinize the differences among some of these approaches to clarify how this work fits in.

The *specmatch* synthetic template matching tool produces noise-free nearest neighbor templates given an input spectrum (Petigura 2015). Several practical barriers limit the accuracy of using precomputed synthetic spectral models alone. First and

foremost, real stars are usually more complicated than our simplified models of them. Real spectra often vary over more dimensions than our models do. Conspicuous examples of these hidden variables can be found in protostars: starspots, accretion veiling, dust extinction, and magnetic Zeeman splitting. Jointly modeling all of these phenomena alongside the intrinsic stellar photosphere is challenging.

The empirical version, *specmatch-emp* (Yee et al. 2017), matched spectra better than the synthetic templates, but it is still too rigid for some applications and requires the assembly of hundreds of standardized high-S/N templates, ideally with low intrinsic rotational broadening. Such a large number of high-quality templates with high resolving power and low $v \sin i$ has not yet been established in the near-infrared.

The *wobble* framework (Bedell et al. 2019) modernized the construction of high-S/N templates to account for temporally variable telluric lines. The tool requires dozens of high-S/N spectra acquired at a range of barycentric Earth RVs. The final telluric-free combined spectrum would still have to be compared to models for absolute calibration or can be used out of the box for precision relative RVs. The *wobble* framework also pioneered the off-label application of automatic differentiation frameworks—in this case TensorFlow—toward their physically motivated use in stellar spectra. *Blase* can be viewed as an evaluable and interpretable super-resolution version of *wobble*, which accepts more bias in the bias-variance trade-off.

The *starfish* framework (Czekala et al. 2015) provides a robust likelihood function for data–model comparisons and retires many of the problems in this domain. *Starfish* pioneered the use of whole-spectrum fitting with resilience to model imperfections by addressing the problem of what to do when the underlying atomic and molecular data were wrong or approximate or missing. It has been extended to inferring starspot physical properties (Gully-Santiago et al. 2017), measuring veiling in Class 0 protostars (Greene et al. 2018), and quantifying imperfections in brown dwarf models (Zhang et al. 2021). The Spectral Inference Crank (*sick*; Casey 2016) shares similar aims to *starfish* and provides additional useful grid search capabilities.

For very large bandwidths and very many spectral lines, the problem of identifying and cataloging line imperfections essentially becomes a bookkeeping and continuum assignment problem. *Blase* and *starfish* provide different strategies for orchestrating the line-mismatch identification procedure, with each route having trade-offs depending on the application.

Appendix C

Conceivable Technical Improvements

Blasé already performs very well under a wide range of cloning and transfer-learning tasks. However, some precision applications may demand even more strenuous performance than what the current implementation can accommodate. Here we describe some of these technical improvements, discuss their design, and mention some science cases they may unlock.

C.1. Exact Instead of Pseudo-Voigt Profile

We currently employ the pseudo-Voigt profile for its low computational cost. We have a prototype exact Voigt–Hjerting implementation following Kawahara et al. (2022). We coarsely estimate that moving to this exact Voigt implementation could

decrease some residual regions by $\sim 30\%$, while increasing the computational cost by more than $10\times$ over the existing pseudo-Voigt approximation. The exact Voigt–Hjerting implementation still outbids the higher cost of a direct numerical convolution of a Gaussian and Lorentzian profile.

C.2. Addressing the Pseudo-continuum with Gaussian Process Regression

We currently assume that the input spectra are adequately normalized to the continuum. We have a few options to relax this assumption. We could simply tune the P term that represents the wavelength-dependent pre-factor to Equation (2). Tuning P would correct for large-scale imperfections in the otherwise-fixed continuum flattening procedure. This change would be easy and effective, but it has some challenges with model selection and flexibility: how to set the polynomial order to avoid over- and underfitting. GPs offer many advantages for continuum fitting (Czekala et al. 2015). In short, a GP likelihood relaxes the assumption that the continuum has been perfectly normalized, in favor of the more realistic statement “the continuum has been coarsely normalized, with some characteristic but as-yet-unknown correlation and scale length and amplitude of the imperfections.” That statement translates to the following modification to Equation (3):

$$\mathcal{L} = \frac{1}{2} \mathbf{R}^\top \mathbf{C}^{-1} \mathbf{R} + \frac{1}{2} \ln \det\{\mathbf{C}\}, \quad (\text{C1})$$

where we introduce the covariance matrix \mathbf{C} , with its associated kernel and collection of typically two to three parameters. We anticipate that this GP likelihood would have the greatest impact on stars with significant band heads and line blanketing: spectra with a so-called “pseudo-continuum.” M dwarfs and brown dwarfs fall into this challenging category.

The main demerit of moving to a GP likelihood is computational cost. Fortunately, a few efficient autodiff-aware implementations of GPs exist. The *celerite* algorithm *celerité* (Foreman-Mackey et al. 2017) has an exact back-propagation implementation (Foreman-Mackey 2018) that scales linearly with the number of data points. The *celerité* algorithm does not currently have a PyTorch implementation. The GPyTorch framework (Gardner et al. 2018) has a large category of approximate and exact GPs that could be straightforwardly dropped into *blase*. Even still, these GPs could increase the computation cost by of order $10\times$.

C.3. Minibatches and Stochastic Gradient Descent

Currently, each training epoch sees the entire data set, a setup dubbed *full-batch* gradient descent. An alternative scheme allows training with only a portion of the entire data set at a time in *minibatches*. The massive data volumes in modern neural network applications cannot fit into the GPU memory, so minibatches are a necessity. Our meager 1 MB data set can easily fit into the GPU memory, but our model can be large if we have a large number of pixels or lines or both. Hence, while minibatches may not be required owing to data size limitations, they may be useful for particularly large models. Minibatches also act as a form of regularization, the principal source of stochasticity in the SGD algorithm, which tends to have better convergence than full-batch gradient descent (Ruder 2016).

We experimented with minibatches by assembling and evaluating only a portion of the dense $\bar{\mathbf{F}}$ matrix at a time, in minibatches. The choice to evaluate only a portion of lines at a time would mean that the model is inaccurately evaluated at all wavelength pixels. Instead, we choose to evaluate all lines, but only on a random subset N_{batch} of the total pixels N_s , so that the model can eventually converge to exact at those points. All lines are allowed to update at each glimpse of a minibatch, but many lines with cores far from minibatch pixels will provide only weak information about how the loss scalar changes for their parameters.

Overall minibatches as implemented above performed worse than the sparse implementation, with both lower accuracy and slower computation time.

C.4. Broad Lines and Advanced Line Shapes

Some lines—such as those arising from hydrogen, sodium, potassium, and others—have extremely broad line wings, approaching larger than the ~ 6000 pixels we allocate for the sparse implementation. These special lines should be handled separately from the weak lines, both from a computational performance perspective and from an accuracy perspective.

Extremely broad lines will exhibit truncation effects if the sparse window is small compared to the line wing size. The truncation effects will look like top-hat functions severing the asymptotic wings, imbuing artificial step function kinks in the emulated spectrum. We can afford to increase the sparse window on a few (say, $N_{\text{broad}} \sim 20$) of the broadest spectral lines. We then construct and evaluate the entire dense matrix for those lines: $\sim 330,000 \times 20$. The number of FLOPS in each category scales as about 6,000,000 for the 20 broad and dense lines versus about 36,000,000 for the sea of about 7000 narrow and sparse lines, depending on the exact choices for wing cuts and the number of lines.

One could introduce advanced line shapes for these ~ 20 broad lines, perturbing the Voigt line wings with a smooth wavelength-dependent correction term \mathbf{G} :

$$\tilde{\mathbf{V}}(\lambda) = \mathbf{V} \odot \mathbf{G} \quad (\text{C2})$$

$$\mathbf{G} = 1 + (e^{\alpha_j} - 1) \cdot \mathcal{S}\left(\frac{|\lambda - \lambda_{c,j}| - \lambda_{t,j}}{b_j}\right), \quad (\text{C3})$$

where \odot is again the element-wise product (aka Hadamard product), \mathcal{S} is the sigmoid function, and we have introduced three new tunable parameters for each of the j broad lines; λ_t is the truncation wavelength, b is a scale parameter for how slowly or how rapidly in wavelength space the transition from non-Lorentzian proceeds, and α is a possibly negative stretch parameter that controls whether the line wing is sub- or super-Lorentzian.

This functional form has a few advantages. It is smooth. The smoothness of the transition is controlled by a tunable parameter, b . It can handle either sub- or super-Lorentzian shapes. In the limit $\lim_{\alpha \rightarrow 0} \mathbf{G}$, the line shape becomes exactly Lorentzian. The sigmoid is efficiently implemented in PyTorch. Finally, it enforces that the perturbation only produces absorption and not emission profiles.

C.5. Using Native Line Lists Rather Than Clones

For many practitioners, the choice to clone precomputed synthetic models in the first place may seem roundabout: “Why

not just use the line lists?” Adopting the line lists would have many advantages: it would provide chemical and molecular provenance tags. Metadata associated with the quality of the atomic and molecular data could be used to assign physics-informed regularization. Many other benefits would effortlessly accrue from adopting the native line lists. The FAL project (P. Cargile et al. 2022, in preparation) follows such a principled prescription.

As already emphasized, there exist at least a few demerits of adopting the line lists and therefore supporting the *blasé* strategy. First, these line lists need to undergo expensive multilevel radiative transfer calculations in order to obtain their amplitudes, so adopting the line lists would mean a laborious and computationally expensive pursuit simply to get close to what has already been computed. Second, as the effective temperature scales to ultracool dwarfs (Figure 1), the number of lines skyrockets, tending toward the billions for T dwarfs. The methane line list alone (Hargreaves et al. 2020) represents a prohibitive data volume. The ExoJAX and Radis (Pannier & Laux 2019; van den Bekerom 2021) libraries offer a breakthrough solution to the voluminous line list problem. Even still, *blasé* deals with the less pure but more practical “pseudo-line” that gets closer to the astronomical observables anyway and offers a middle ground between the extremes of interpretability and performance.

C.6. Wavelength-dependent Limb Darkening

Currently, the extrinsic model step possesses up to four parameters: the $v \sin i$ and RV, and two optional parameters for limb darkening. These four parameters may adequately parameterize a star with a uniform stellar disk. Extremely high S/N spectra of rapidly rotating stars may require additional flexibility. The limb darkening is generally wavelength dependent, and so a panchromatic spectrum may require a different limb darkening from the blue end to the red end. The limb darkening may instead depend on physical properties of the spectral line formation, such as physical depth of formation, and so the extent of limb darkening may jump haphazardly from line to line to line, rather than as a predictably smooth function across wavelength. *Blase* could be built to handle such a seemingly pathological scenario by adding a vector of limb-darkening parameters, one for each line. One would have to regularize the fits with some typical limb darkening and a heuristic penalty for departures from this mean.

ORCID iDs

Michael Gully-Santiago  <https://orcid.org/0000-0002-4020-3457>
 Caroline V. Morley  <https://orcid.org/0000-0002-4404-0456>

References

Abadi, M., Agarwal, A., Barham, P., et al. 2016, arXiv:1603.04467

Aigrain, S., & Foreman-Mackey, D. 2022, arXiv:2209.08940

Al Mouda, K., Dumusque, X., Cretignier, M., Zhao, Y., & Valenti, J. A. 2022, *A&A*, 664, A34

Anderson, D. R., Collier Cameron, A., Delrez, L., et al. 2014, *MNRAS*, 445, 1114

Ansdell, M., Ioannou, Y., Osborn, H. P., et al. 2018, *ApJL*, 869, L7

Apai, D., Nardiello, D., & Bedin, L. R. 2021, *ApJ*, 906, 64

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33

Bedell, M., Hogg, D. W., Foreman-Mackey, D., Montet, B. T., & Luger, R. 2019, *AJ*, 158, 164

Birkby, J. L., de Kok, R. J., Brogi, M., et al. 2013, *MNRAS*, 436, L35

Bradbury, J., Frostig, R., Hawkins, P., et al. 2018, JAX: Composable Transformations of Python+NumPy programs, v.0.2.5, <http://github.com/google/jax>

Brogi, M., Snellen, I. A. G., de Kok, R. J., et al. 2012, *Natur*, 486, 502

Cannon, A. J., & Pickering, E. C. 1901, *AnHar*, 28, 129

Cargile, P. A., Conroy, C., Johnson, B. D., et al. 2020, *ApJ*, 900, 28

Casasayas-Barris, N., Palle, E., Nowak, G., et al. 2017, *A&A*, 608, A135

Casey, A. R. 2016, *ApJS*, 223, 8

Chojnowski, S. D., Hubrig, S., Hasselquist, S., et al. 2020, *MNRAS*, 496, 832

Clough, S. A., Shephard, M. W., Mlawer, E. J., et al. 2005, *JQSRT*, 91, 233

Cranmer, M. 2020, PySR: v0.11.11 Fast & Parallelized Symbolic Regression in Python/Julia, Zenodo, doi:10.5281/zenodo.4041459

Cretignier, M., Dumusque, X., Hara, N. C., & Pepe, F. 2021, *A&A*, 653, A43

Crossfield, I. J. M., Biller, B., Schlieder, J. E., et al. 2014, *Natur*, 505, 654

Czekala, I., Andrews, S. M., Mandel, K. S., Hogg, D. W., & Green, G. M. 2015, *ApJ*, 812, 128

Czekala, I., Mandel, K. S., Andrews, S. M., et al. 2017, *ApJ*, 840, 49

Czekala, I., Zawadzki, B., Loomis, R., et al. 2021, MPoL-dev/MPoL: v0.1.1 Release, v0.1.1, Zenodo, doi:10.5281/zenodo.4939048

Deen, C. P. 2013, *AJ*, 146, 51

Desdoits, L., Pope, B., & Tuthill, P. 2022, *Proc. SPIE*, 12180, 1218032

Dreyfus, S. E. 1990, *JGCD*, 13, 926

Dumusque, X. 2018, *A&A*, 620, A47

Erkal, J., Manara, C. F., Schneider, P. C., et al. 2022, *A&A*, 666, A188

Foreman-Mackey, D. 2018, *RNAAS*, 2, 31

Foreman-Mackey, D., Agol, E., Ambikasaran, S., & Angus, R. 2017, *AJ*, 154, 220

Foreman-Mackey, D., Luger, R., Agol, E., et al. 2021, *JOSS*, 6, 3285

Gardner, J. R., Pleiss, G., Bindel, D., Weinberger, K. Q., & Wilson, A. G. 2018, arXiv:1809.11165

Gent, M. R., Bergemann, M., Serenelli, A., et al. 2022, *A&A*, 658, A147

Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., & Dahl, G. E. 2017, arXiv:1704.01212

Gilmore, G., Randich, S., Asplund, M., et al. 2012, *Msngr*, 147, 25

Gordon, I. E., Rothman, L. S., Hargreaves, R. J., et al. 2022, *JQSRT*, 277, 107949

Greene, T. P., Gully-Santiago, M. A., & Barsony, M. 2018, *ApJ*, 862, 85

Gullikson, K., Dodson-Robinson, S., & Kraus, A. 2014, *AJ*, 148, 53

Gully-Santiago, M. A., Herczeg, G. J., Czekala, I., et al. 2017, *ApJ*, 836, 200

Gully-Santiago, M. A., Luna, J. L., Morley, C. V., et al. 2022, *JOSS*, 7, 4302

Gunes Baydin, A., Pearlmuter, B. A., Andreyevich Radul, A., & Siskind, J. M. 2015, arXiv:1502.05767

Hargreaves, R. J., Gordon, I. E., Rey, M., et al. 2020, *ApJS*, 247, 55

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, 585, 357

Hjerting, F. 1938, *ApJ*, 88, 508

Ho, A. Y. Q., Ness, M. K., Hogg, D. W., et al. 2017, *ApJ*, 836, 5

Hoeijmakers, H. J., de Kok, R. J., Snellen, I. A. G., et al. 2015, *A&A*, 575, A20

Hood, C. E., Fortney, J. J., Line, M. R., et al. 2020, *AJ*, 160, 198

Hunter, J. D. 2007, *CSE*, 9, 90

Husser, T.-O., Wende-von Berg, S., Dreizler, S., et al. 2013, *A&A*, 553, A6

Ida, T., Ando, M., & Toraya, H. 2000, *JApCr*, 33, 1311

Johnsen, T. K., Marley, M. S., & Gulick, V. C. 2020, *PASP*, 132, 044502

Kawahara, H., Kawashima, Y., Masuda, K., et al. 2022, *ApJS*, 258, 31

Kelley, H. J. 1960, *Am. Rocket Soc. J.*, 30, 947

Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980

Lee, J.-J., & Gullikson, K. 2016, plp: v2.1 alpha 3, Zenodo, doi:10.5281/zenodo.56067

Line, M. R., Teske, J., Birmingham, B., Fortney, J. J., & Marley, M. S. 2015, *ApJ*, 807, 183

Linnainmaa, S. 1976, *BIT Numer. Math.*, 16, 146

Luger, R., Bedell, M., Foreman-Mackey, D., et al. 2021a, arXiv:2110.06271

Luger, R., Foreman-Mackey, D., Hedges, C., & Hogg, D. W. 2021b, *AJ*, 162, 123

Maclaurin, D. 2016, PhD thesis, Harvard Univ., Massachusetts

Mahadevan, S., Ramsey, L. W., Terrien, R., et al. 2014, *Proc. SPIE*, 9147, 91471G

Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, *AJ*, 154, 94

Marley, M. S., Saumon, D., Visscher, C., et al. 2021, *ApJ*, 920, 85

McKinney, W. 2010, in Proc. 9th Python in Science Conf., ed. S. van der Walt & J. Millman, 56

O’Neal, D., Saar, S. H., & Neff, J. E. 1996, *ApJ*, **463**, 766

Pannier, E., & Laux, C. O. 2019, *JQSRT*, **222**, 12

Park, C., Jaffe, D. T., Yuk, I.-S., et al. 2014, *Proc. SPIE*, **9147**, 91471D

Passegger, V. M., Bello-García, A., Ordieres-Meré, J., et al. 2020, *A&A*, **642**, A22

Paszke, A., Gross, S., Massa, F., et al. 2019, arXiv:1912.01703

Payne, C. H. 1925, PhD thesis, Radcliffe College

Pérez, F., & Granger, B. E. 2007, *CSE*, **9**, 21

Petigura, E. A. 2015, PhD thesis, Univ. California, Berkeley

Piskunov, N. E., Kupka, F., Ryabchikova, T. A., Weiss, W. W., & Jeffery, C. S. 1995, *A&AS*, **112**, 525

Pope, B. J. S., Pueyo, L., Xin, Y., & Tuthill, P. G. 2021, *ApJ*, **907**, 40

Roettenbacher, R. M., Monnier, J. D., Korhonen, H., et al. 2016, *Natur*, **533**, 217

Rothman, L. S., Gordon, I. E., Barber, R. J., et al. 2010, *JQSRT*, **111**, 2139

Ruder, S. 2016, arXiv:1609.04747

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986, *Natur*, **323**, 533

Saad, Y. 2003, *Iterative Methods for Sparse Linear Systems*, Other Titles in Applied Mathematics (2nd ed.; Philadelphia, PA: SIAM)

Smette, A., Sana, H., Noll, S., et al. 2015, *A&A*, **576**, A77

Sneden, C., Bean, J., Ivans, I., Lucatello, S., & Sobeck, J. 2012, MOOG: LTE Line Analysis and Spectrum Synthesis, Astrophysics Source Code Library, ascl:1202.009

Snellen, I. A. G., de Kok, R. J., de Mooij, E. J. W., & Albrecht, S. 2010, *Natur*, **465**, 1049

Steinmetz, M., Zwitter, T., Siebert, A., et al. 2006, *AJ*, **132**, 1645

Tannock, M. E., Metchev, S., Hood, C. E., et al. 2022, *MNRAS*, **514**, 3160

Tennyson, J., & Yurchenko, S. N. 2012, *MNRAS*, **425**, 21

van den Bekerom, D. C. M., & Pannier, E. 2021, *JQSRT*, **261**, 107476

Vernet, J., Dekker, H., D’Odorico, S., et al. 2011, *A&A*, **536**, A105

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261

Vogt, S. S., & Penrod, G. D. 1983, *PASP*, **95**, 565

Waskom, M. 2021, *JOSS*, **6**, 3021

Waters, C. Z., & Hollek, J. K. 2013, *PASP*, **125**, 1164

Wong, A., Pope, B., Desdoigts, L., et al. 2021, *JOSAB*, **38**, 2465

Yee, S. W., Petigura, E. A., & von Braun, K. 2017, *ApJ*, **836**, 77

Zaghoul, M. R., & Ali, A. N. 2011, arXiv:1106.0151

Zhang, Z., Liu, M. C., Marley, M. S., Line, M. R., & Best, W. M. J. 2021, *ApJ*, **921**, 95

Zhao, L. L., Fischer, D. A., Ford, E. B., et al. 2022, *AJ*, **163**, 171