Research Statement:
Phillip A. Cargile

Stellar astronomy is in the midst of a continuing data revolution. Large spectroscopic surveys such as SDSS-V, DESI, and H3, together with upcoming programs including Via, S5, WEAVE, and 4MOST, are measuring detailed spectra for hundreds of millions of stars across the Milky Way. Roman and SPHEREx will extend this view with enormous spectrophotometric data sets, while Gaia continues to provide precise positions, distances, and motions for billions of Milky Way stars. Together, these facilities provide the information needed to reconstruct the formation of the Milky Way, map the interstellar medium, measure the ages and compositions of stars, and test the physical models used throughout astrophysics.

The scientific value of these data is realized only when we can combine spectra, photometry, and astrometry in a physically consistent analysis. This is difficult for two related reasons. First, the forward models themselves remain imperfect: uncertainties in radiative transfer, stellar atmospheres, stellar evolution, and atomic and molecular data produce systematic errors in predicted observables. Second, the volume and complexity of modern data demand methods that are fast, statistically rigorous, and capable of propagating uncertainty across heterogeneous measurements. A precise observation does not guarantee an accurate physical result when either the model or the inference procedure is inadequate.

My research focuses on improving *ab initio* models of stars and stellar spectra, developing computational and machine-learning tools that make those models usable at scale, and applying the resulting systems to major astronomical surveys. I approach these activities as parts of one stellar-astrophysics program: observed stellar spectra and photometry expose limitations in stellar atmospheres, atomic and molecular data, and evolutionary models; improved stellar models enable more accurate inference of temperatures, compositions, masses, and ages; survey-scale inference creates benchmark samples spanning the Hertzsprung–Russell diagram; and those benchmarks identify the next piece of stellar physics or computation to improve. The projects below describe the principal components of this work and motivate my future research program.

Fitting All the Lines — Empirical Calibration of Atomic & Molecular Transitions:

A theoretical stellar spectrum begins with a model atmosphere, a radiative-transfer calculation, and a line list containing the atomic and molecular transitions that absorb light. In the optical spectra of cool stars, millions of transitions can contribute measurable opacity. Their wavelengths, oscillator strengths, and broadening parameters are assembled from laboratory experiments and quantum calculations, but many remain insufficiently accurate for the precision now delivered by astronomical observations. These errors are not merely cosmetic: discrepancies in individual and blended atomic and molecular lines propagate into biased measurements of stellar temperature, gravity, composition, age, and distance.

I lead *Fitting All the Lines* (FAL), a project that empirically calibrates approximately 140,000 optical atomic and molecular lines against benchmark stars such as the Sun, Arcturus, and Procyon. The global fit simultaneously adjusts line wavelengths, oscillator strengths, and van der Waals damping parameters for line groups shared across the benchmark sample. Each model evaluation uses Korg to synthesize the stellar spectra, while a robust loss function and sensitivity screening focus the optimization on informative transitions. Fitting stars that span temperature, gravity, and metallicity is especially important in blended regions: it helps distinguish an error in a line parameter from an imperfect stellar model or a calibration error in one benchmark spectrum.

The next product of this work is the CWC grid, a Korg-based library that will combine improvements in the synthesis code with the empirically calibrated FAL line list. In addition to conventional regularly spaced grids, the library will include randomly distributed synthetic spectra across the Hertzsprung–Russell diagram. This complementary sampling will improve coverage of the stellar parameter space and support accurate interpolation and neural emulation without imposing artifacts from a single grid geometry. The CWC grid is being designed both for modeling individual stars and for use within broader applications, including future releases of Fast Stellar Population Synthesis (FSPS). It will also form the basis of the FAL manuscript now in preparation, *Fitting All the Lines: Empirical Calibration of ∼140,000 Atomic and Molecular Lines for Stellar Spectral Modeling*.

This project embodies a central principle of my work: reliable inference requires treating the forward model as an object to test and improve, not as fixed truth. Expanding the calibration across stellar type will reveal transitions that are weak or absent in solar-like stars, while comparisons among Korg, SYNTHE, and ATLAS12 will distinguish limitations in the line data from differences in the radiative-transfer implementation. The long-term goal is a transparent spectral-modeling framework in which physical assumptions, empirical corrections, and benchmark performance can be examined together.

Industrial-Scale Inference of Stellar Parameters:

Predictions from stellar-atmosphere and stellar-evolution calculations are generally distributed as finite grids because each model can be computationally expensive. Applying those predictions to arbitrary stellar parameters therefore requires interpolation, and interpolation error can become a leading systematic when the observations are precise. The problem is especially challenging for spectra, where a model may contain tens of thousands of wavelength pixels and depend on many atmospheric labels.

With Yuan-Sen Ting, I co-developed *The Payne* (Ting et al. 2019), one of the earliest examples in stellar astrophysics of using neural networks as surrogate models for *ab initio* calculations. The emulator learns the smooth, high-dimensional mapping between stellar atmospheric parameters and the flux at each wavelength. Once trained and validated, it is fast and memory efficient, preserves the structure of the underlying physical grid, and provides derivatives that can be used in optimization and inference. This approach made it practical to evaluate detailed synthetic spectra millions of times without separating the analysis from its physical foundation.

I extended this strategy to stellar evolution in *MINESweeper* (Cargile et al. 2020), a framework I created and continue to lead. MIST stellar-evolution models connect fundamental quantities such as initial mass, age, and composition to present-day atmospheric properties. MINESweeper joins those predictions to synthetic spectra and broadband photometry, allowing spectra, photometry, and astrometric information to be analyzed simultaneously. Its Bayesian framework returns a multivariate posterior rather than a single best-fitting label vector, making degeneracies explicit and allowing information such as Gaia parallaxes to enter as a prior within one coherent model.

My current development effort, *uberMS*, moves this physical model into a fast, differentiable inference framework. Automatic differentiation enables gradient-based posterior methods such as Hamiltonian Monte Carlo, NUTS, and variational inference, reducing the cost of a full spectrophotometric analysis from several hours per star to on the order of minutes. Recent applications to high-resolution TRES spectra demonstrate the performance of uberMS metallicities relative to established literature measurements (Pass et al. 2025). The manuscript now in preparation will describe fast, differentiable Bayesian inference of stellar parameters from spectrophotometric data and establish the validation needed for broader survey deployment.

This program increasingly draws on generative distribution modeling, including normalizing flows, flow matching, and score-based diffusion models. These methods offer flexible ways to represent complicated stellar populations, priors, and posterior distributions without forcing them into restrictive analytic forms. Related work on the Dartmouth Stellar Evolution Emulator (DSEE; Ying et al. 2026) uses generative modeling to broaden access to stellar-evolution predictions. My goal is not to replace physical models with generic machine-learning predictions, but to use modern ML where it expands, accelerates, or stress-tests a physically interpretable analysis.

From Methods to Astronomical Surveys:

The completed H3 Spectroscopic Survey was the principal proving ground for this end-to-end approach. H3 collected and analyzed more than 350,000 MMT/Hectochelle spectra to map Milky Way halo stars and reconstruct the Galaxy's assembly history. As a core team member, I contributed to survey planning, target selection, observations, data archiving, processing, and scientific analysis. I built and operated the MINESweeper component of the H3 pipeline on Harvard's FASRC computing cluster, jointly fitting each spectrum with public photometry and Gaia measurements. My H3 analysis products have supported more than twenty peer-reviewed papers and five Harvard PhD theses, including studies of the Milky Way's mass, merger history, dark-matter distribution, and early disk.

The same infrastructure now supports a wider range of Galactic-archaeology programs. In SDSS-V, I develop and apply stellar-inference pipelines to map distant, metal-poor populations using SDSS spectroscopy, photometry, and Gaia astrometry (Chandra et al. 2026). These distant stars probe regions of the halo that are difficult to reach with local samples and provide a record of accretion and chemical evolution across the Galaxy. Survey deployment is also an important form of model validation: selection effects, low signal-to-noise, missing data, and sparsely tested regions of parameter space reveal failures that are invisible in carefully selected benchmark samples.

For the future Via Survey, I serve as the data processing and analysis lead and am now building the calibration, spectral-extraction, and stellar-inference tools that will be applied to its observations. The Via digital twin supplies an explicit response matrix that maps a one-dimensional spectrum onto the two-dimensional detector image, turning extraction into a well-defined but ill-conditioned linear inverse problem. Conventional regularization can stabilize this inversion while subtly suppressing line cores, altering equivalent widths, or miscalibrating uncertainties. I am testing methods that combine established extraction methods with learned priors from computational imaging, including plug-and-play denoisers and score-based diffusion models. Injection–recovery experiments measure not only flux error but biases in line depth, equivalent width, radial velocity, and uncertainty coverage, while out-of-distribution tests identify when a learned prior begins to hallucinate rather than recover spectral information. This work is intended to deliver a validated extraction pipeline before Via begins survey operations.

My survey work also extends from stars to the Galactic interstellar medium. In collaborations with the Roman Galactic Plane Survey and SPHEREx, I provide stellar modeling of spectrophotometric data to infer stellar properties and construct detailed maps of interstellar dust. Stars at different distances act as probes through the dust distribution; analyzing them consistently makes it possible to recover both stellar parameters and the intervening extinction. Roman's depth and angular resolution and SPHEREx's all-sky coverage offer complementary views of the dusty Galaxy, creating an opportunity to connect stellar physics, Galactic structure, and the cycling of material through the interstellar medium.

I have also contributed to the JWST Advanced Deep Extragalactic Survey (JADES) through collaboration with the NIRCam imaging team. I helped process, register, and stack JADES and complementary imaging of GOODS-North and GOODS-South. This work resulted in the JADES Origins Field, the deepest JWST imaging yet achieved, and supported searches that identified the highest-redshift galaxies known at the time. Although the scientific target differs from my stellar-spectroscopy work, the methodological theme is the same: careful processing and detailed calibration of heterogeneous large-scale observational data sets are prerequisites for robust and trustworthy inference when probing the limits of the data.

<u>Future Directions:</u>

My future research will build on four connected directions:

- ***A community-ready calibrated spectral grid:*** I will complete the optical FAL calibration and develop the Korg-based CWC grid across the stellar parameter space required by current surveys. The central validation will be differential performance across benchmark stars, not simply agreement with the spectra used during calibration. Comparisons with SYNTHE and ATLAS12, tests of updated atomic and molecular data, and explicit tracking of empirical corrections will make it possible to identify which improvements are robust across synthesis codes and stellar types.
- ***Spectral extraction as a physics-informed inverse problem:*** Building on Via's digital twin, I will develop extraction methods that combine an explicit detector-response model with learned priors from computational imaging. The goal is to recover the spectrum and calibrated uncertainty directly from the detector image while measuring the bias introduced by regularization. Injection–recovery tests, prior–data tension diagnostics, and out-of-distribution benchmarks will determine when learned information improves a low-signal spectrum and when the data should instead override the prior.
- ***Differentiable inference with calibrated uncertainty:*** I will extend uberMS to larger and more heterogeneous stellar samples while testing when faster posterior approximations remain

scientifically reliable. Exact or high-fidelity samplers can audit representative cases; posterior predictive checks and simulation-based coverage tests can identify regimes where an approximation fails. Generative distribution models can then capture non-Gaussian population structure and accelerate inference without hiding the connection to the underlying stellar models.

– ***Mapping stars and dust across the Milky Way:*** Roman, SPHEREx, Gaia, SDSS-V, H3, and Via provide complementary measurements of stars over different distances, wavelengths, and Galactic environments. I will combine these data through a consistent stellar forward model to improve three-dimensional maps of extinction, characterize distant stellar populations, and test whether residuals arise from the dust model, the stellar model, or the calibration. This joint analysis will turn survey overlap into a scientific constraint rather than a cross-matching exercise.

These directions require close collaboration among experts in stellar atmospheres and evolution, instrumentation, survey science, statistics, and machine learning. I helped found AstroAI at the Center for Astrophysics and serve on its executive committee in order to build precisely this kind of exchange. Through AstroAI, invited tutorials, and direct work with graduate and early-career researchers, I aim to make advanced computational methods understandable and useful to scientists who also bring essential domain knowledge to the problem.

The unifying objective of my research is to make the path from stellar photons to physical conclusions more accurate, scalable, and transparent. That path begins with the forward model for stars—stellar atmospheres and evolution, radiative transfer, and atomic and molecular data—and continues through detector calibration and spectral extraction to the joint interpretation of spectra, photometry, and astrometry. By improving each connected stage, I aim to deliver accurate stellar temperatures, compositions, masses, ages, distances, and extinctions for samples ranging from individual benchmarks to entire surveys. Those measurements will support concrete astrophysical goals: reconstructing the Milky Way's assembly, mapping its stars and dust, and testing where current models of stars and their observations break down.