Understanding Stellar Luminosity with Generative Deep Learning

Maria Nareklishvili[1] Nick Polson[2] Vadim Sokolov[3]

[1]Stanford GSB [2]Chicago Booth [3]George Mason University

May 2025

Outline

Stefan-Boltzmann Law

- Astrophysicists study stars primarily by observing their light with telescopes and analyzing their spectra.
- The spectrum and color of a star's emitted light are direct indicators of its surface temperature.
- Blue stars have the highest surface temperatures, while red and orange stars are cooler.

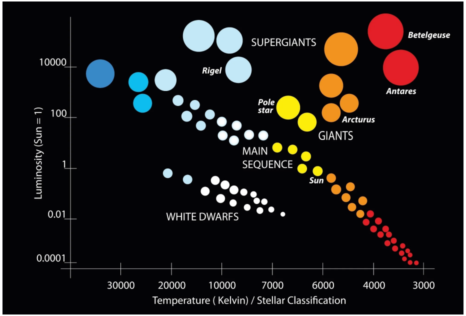

Figure: Hertzprung-Russel Diagram: Hotter stars are brighter.

The Stefan-Boltzmann Law:

- Describes the total energy radiated per unit surface area of a star per unit time.
- The total luminosity is proportional to the fourth power of the surface temperature.

$$L = 4\pi R^2 \sigma T^4$$

where:

- R = radius of a star
- T = temperature (K)
- σ = Stefan-Boltzmann constant

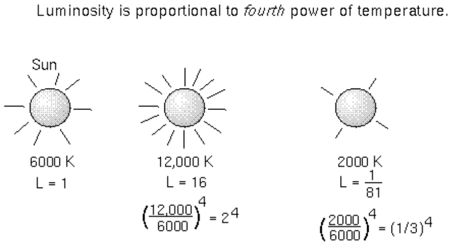

Figure: Luminosity is proportional to fourth power of temperature

Absolute Magnitude vs Stellar Brightness

The Stefan-Boltzmann Law:

- Stars are at many different distances from Earth. A very bright star that is far away might look dimmer than a faint star that is close by. Absolute magnitude removes the effect of distance, so you can compare the true brightness of stars.
- An object's absolute magnitude is defined to be equal to the apparent magnitude that the object would have if it were viewed from a distance of exactly 10 parsecs (32.6 light-years = is approximately 9.46 trillion km).

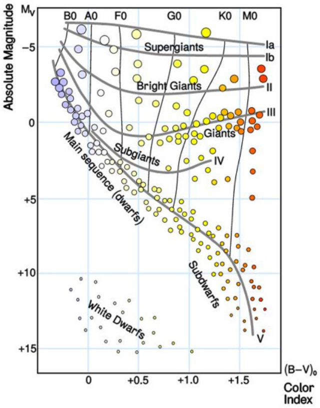

Figure: Stars with lower absolute magnitude are brighter and hotter.

Why Revisit the Stefan–Boltzmann Law? Measurement Challenges

- Cornerstone of stellar physics: $L = 4\pi R^2 \sigma T^4$.
- It is often assumed that stars are black bodies. A black body is an idealized physical object that absorbs all electromagnetic radiation (light) that falls on it, regardless of wavelength or angle.
- Real stars $\neq$ perfect black bodies: emissivity, spots, atmospheres.
- Observational noise (calibration, distance) blurs the T^4 signal.
- Goal: Quantify the entire *distribution* of temperature effects on stellar brightness, not just the mean.

What We Contribute

1. Research Questions:
   - Can a machine discover the nonlinear relationship of temperature and luminosity?
   - Can we teach machines the law and accelerate the discovery of such a relationship?
2. Physics-informed generative deep learner linking T to L.
3. Counterfactual simulation of stellar luminosities (borrowed from causal inference).
4. Full distribution of heterogeneous temperature effects.
5. Empirical validation on 240 Gaia DR3 main-sequence stars.

Temperature $\Rightarrow$ Luminosity: a Causal Blueprint

Physics-anchored causality

$$T \xrightarrow{\sigma T^4} L$$

- Unidirectional law — temperature dictates energy flux; luminosity never feeds back.
- Natural experiments — stellar temperature shifts (nuclear burning, ISM accretion) are exogenous to L.
- Why it matters — fixes the causal arrow needed to learn effects, not just correlations.

Questions to a Machine:

- How do changes in temperature influence the brightness of each star in the presence of uncertainty?
- Is the relationship positive strictly for all stars? Do we observe outliers?

Fundamental Problem of Causal Inference

Only one outcome is observed per star, but we recover θ_s by simulating these alternative (counterfactual) states and iteratively update in a deep learner.

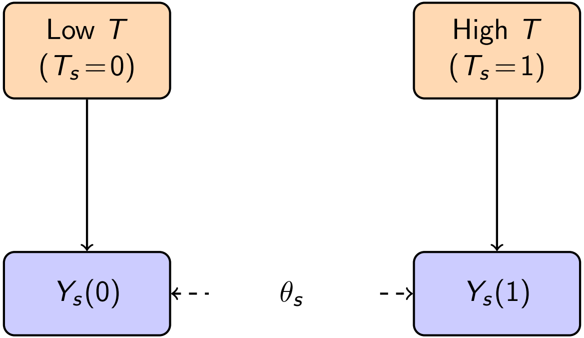

Graphical takeaway: only one outcome observed per star, but we recover θ_s via a physics-informed generative learner.

Causal Lens on Astrophysics: Potential Outcomes Framework

Potential-outcomes lens

$Y_s(1)$: Luminosity if star s at high T
$Y_s(0)$: Luminosity if star s at low T
Only a single brightness state is realised for each star:

$$Y_s = T_s Y_s(1) + (1 - T_s)\, Y_s(0).$$

Decompose potential outcomes (hypothetical brightness):

$$\begin{aligned} Y_s(0) &= \mu_s + \varepsilon_s(0), \\ Y_s(1) &= \mu_s + \theta_s + \varepsilon_s(1), \\ \implies \theta_s &= Y_s(1) - Y_s(0) \end{aligned}$$

Star-specific effect Simulate θ_s with no functional-form constraints.
Heterogeneity & noise explicitly admitted through μ_s and $\varepsilon_s(d)$.

Physics-Informed Generative Model

Stochastic radiation equation

$$Y_s = \mu_s + \theta_s T_s + \varepsilon_s, \qquad \varepsilon_s \mid T_s \sim \mathcal{N}\left(\sigma T_s^4,\ [0.1\,\sigma T_s^4]^2\right)$$

- $\theta_s \sim \mathrm{Uniform}[-\Theta, \Theta]$: flat prior reflects agnosticism about effect size.
- Noise mean $\propto T^4$ (Stefan–Boltzmann); variance $\propto T^8$ (thermodynamic dispersion).
- Physics-based constraints focus learning on *admissible* θ_s.
- In this paper: we approximate parameters (μ_s, θ_s) with a deep learner.

Take-away: Uncertainty is set by physical law, not tuning—yielding a generative model that respects stellar energy budgets.

Target Parameter

Target. For a star with $X_s = x$, estimate the conditional average temperature effect:

$$\theta_s = \mathbb{E}[Y_s(1) - Y_s(0) \mid X_s = x] = \int_0^1 \left(F^{-1}_{Y(1)|x}(q) - F^{-1}_{Y(0)|x}(q)\right) dq.$$

Strategy. Approximate each inverse CDF by a truncated Fourier series:

$$F^{-1}_{Y(d)|x}(q) \approx \sum_{k=0}^{K} \beta_k(x; \vartheta) \cos(k\pi q),$$

with β_k output by a deep network (weights ϑ).

Implementation.

- Sample $q \sim \mathrm{Unif}(0, 1)$ at each forward pass.
- Shared network maps covariates; two heads ($d = 0, 1$) yield $\{\beta_k\}$ per each temperature state.
- Averaging over q gives $\hat{\theta}_s$ nonparametrically, capturing heterogeneous effects of stellar temperature on brightness.

Generative Deep Learning

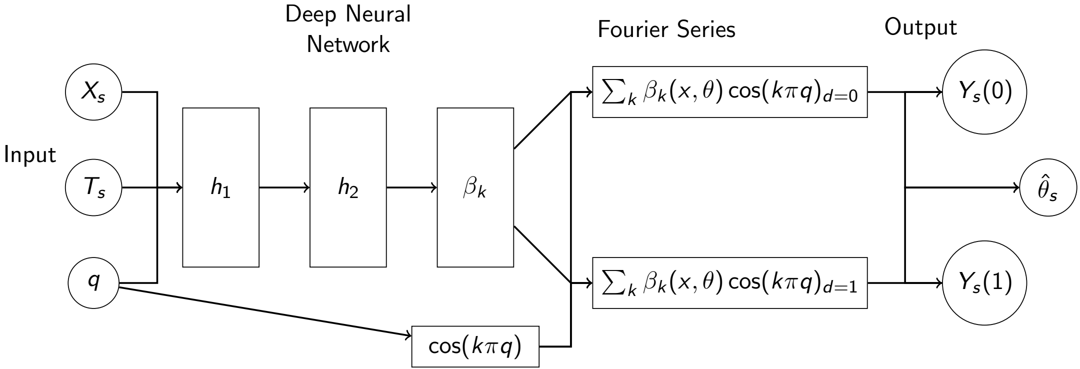

Architecture of the generative deep learning. ReLU activation functions are used everywhere.

Multi-Task Loss

$$\mathcal{L} = w_1 \, \mathrm{BCE}(T_s, \hat{\pi}_s) + w_2 \, \mathrm{QuantileLoss} + w_3 \, \mathrm{MSE}$$

- Binary cross entropy (BCE): adversarially predicts T to deal with selection.
- Quantile loss matches conditional CDFs.
- MSE stabilizes point predictions.
- Stochastic gradient descent (Adam)

Gaia DR3 Sample

- 1.8B stars $\Rightarrow$ focus on **main sequence**.
- Filter quality flags, get $S = 240$ with
  - Temperature (K)
  - Radius $R/R_{\odot}$
  - Absolute magnitude M_V
  - Luminosity $L/L_{\odot}$
- Train/test split: 50/50.

Main sequence stars fuse hydrogen atoms to form helium atoms in their cores (https://science.nasa.gov/universe/stars/types/)
.

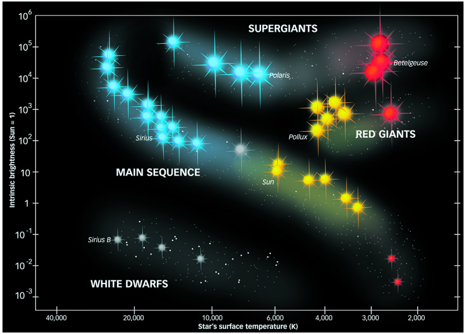

Figure: Main sequence stars.

Heterogeneous Temperature Effects

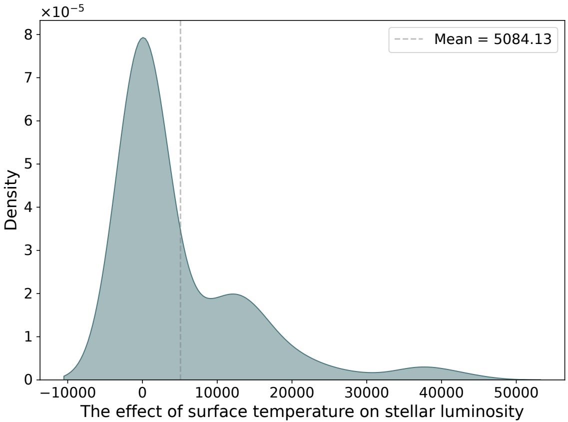

- Right-skewed: most stars modest effects, long high-tail.
- Mean $> 0 \Rightarrow$ hotter stars brighter (confirms T^4).

On Average Effect Increases with Radius

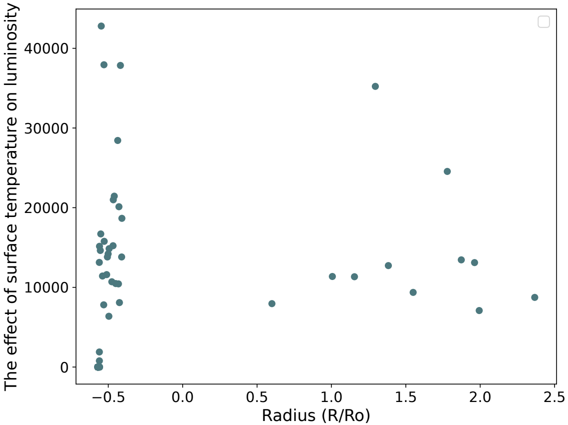

Larger radius $\Rightarrow$ surface area multiplier reinforces T sensitivity.

Effect Decreases with Absolute Magnitude

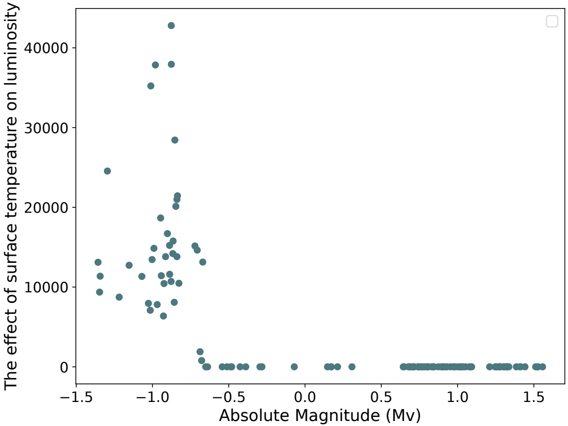

Lower M_V (intrinsic brightness) $\rightarrow$ stronger T effects.

Comparison with Generalized Random Forest

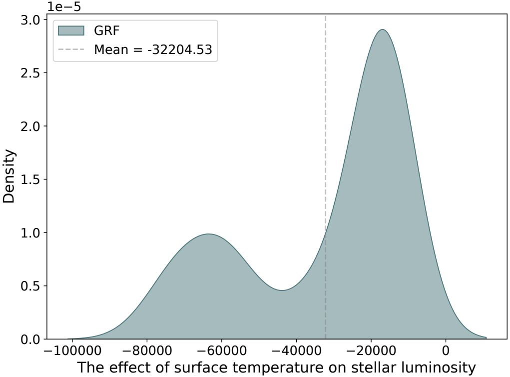

- GRF produces implausible negative effects in small samples.
- Our physics-informed learner rediscovers Stefan-Boltzmann law.

Physical Insights Recovered

1. θ_s scales with R^2: surface area matters.
2. Absolute magnitude inversely related to T sensitivity.
3. Distributional view reveals minority of stars with extreme responses.

Takeaways

Generative Deep Learning validates Stefan–Boltzmann Law

- Recovers physics law under measurement noise.
- Provides star-specific effect estimates.
- Bridges causal thinking and astrophysics.
- Generative deep learning avoids unknown prior distributions in variational inference.

Broader Impact

- Data-driven stellar models may help understand exoplanet habitability.
- Framework generalizes to any noisy physical law.
- Opens door to policy applications of astronomy (e.g., satellite calibration).

Questions?

Thank you!

Contact: marnar@stanford.edu