Time-skipping stellar evolution:
Autoregressive deep learning emulation of the red giant branch

João Santos[1*], Diogo Capelo[1], Rui Henriques[2] and Ilídio Lopes[1]

[1] CENTRA — Centro de Astrofísica e Gravitação, Instituto Superior Técnico, Lisboa, Portugal
[2] INESC-ID, Instituto Superior Técnico, Lisboa, Portugal
*joao.nicolau.santos@tecnico.ulisboa.pt

Motivation

Stellar evolution codes offer unrivalled accuracy, yet:
- computationally heavy
- manual, expert-driven tuning
- limited parameter coverage.

Grid **interpolation**, a traditional alternative, yet:
- restricted applicability
- resolution-sensitive
- curse of dimensionality.

The **red giant branch** is a particularly complex phase to model.

Can machine learning provide a fast, accurate alternative?

Methods

$$f : \mathbf{y}(t) \mapsto \mathbf{y}(t + \delta t)$$

Data

Grid of 755 stellar tracks using MESA

$1.05 \leq M_i / M_\odot \leq 1.60$
$0.0100 \leq Z_i \leq 0.0300$
$1.00 \leq \alpha_{MLT} \leq 2.50$

Helium core mass as an age proxy $\delta t \to \delta M_{He}$

Restricted to RGB: $0.12 < M_{He} < 0.46$

Alignment to a fixed grid of 5000 timesteps

34 features: global parameters, luminosities, 24 mass fractions

Models
- eXtreme Gradient Boosting
- Feed-Forward Neural Networks
- Gated Recurrent Units
- Neural Ordinary Differential Equations

Metrics
- Baseline comparison
- Clustering
- Rollout prediction
- Uncertainty estimation
- Physical consistency

Results

FNN achieves the lowest error among all tested architectures.

He-burning luminosity is the hardest feature, yet below 0.02 dex.

Epistemic uncertainty correlates strongly with prediction error (r > 0.98), enabling estimates for unseen stars.

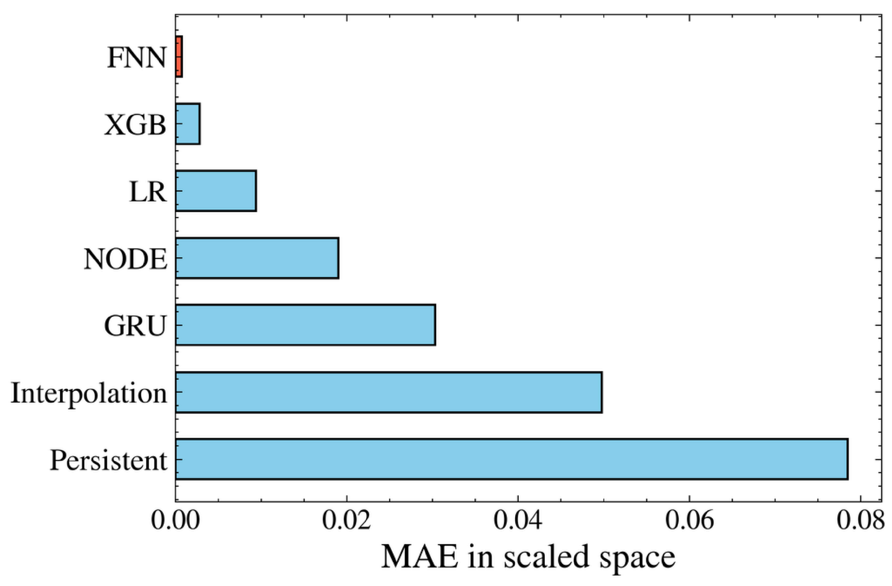

Global error across tested architectures

Rollout stability for appropriately chosen steps, for $\delta M_{He} \approx 0.013\, M_\odot$:

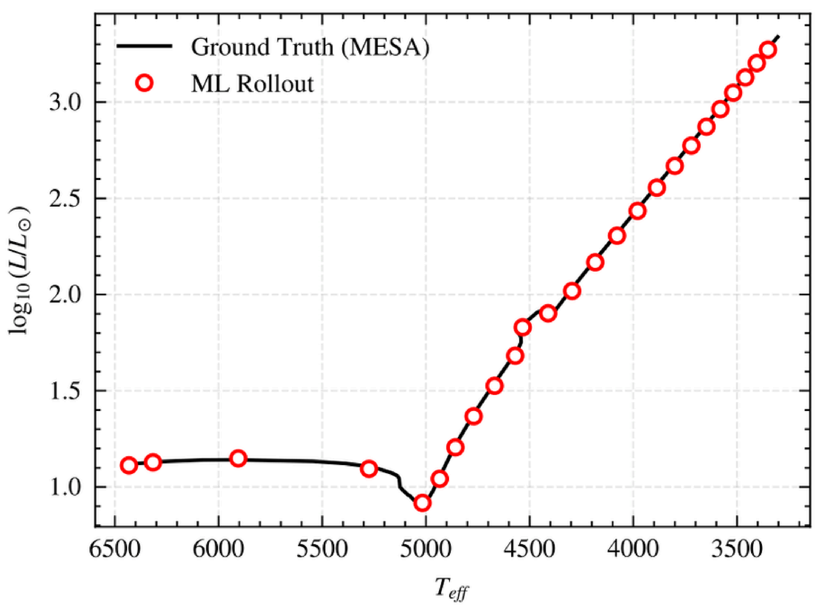

Predicted vs true evolutionary track

17ms per track

<0.20% error for global parameters

<0.05 dex error for luminosities and mass fractions

Physics check ✓

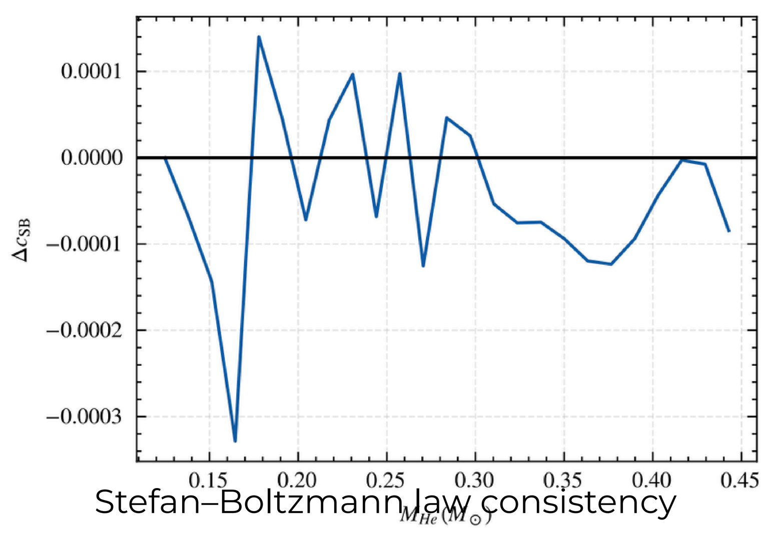

Stefan–Boltzmann law consistency

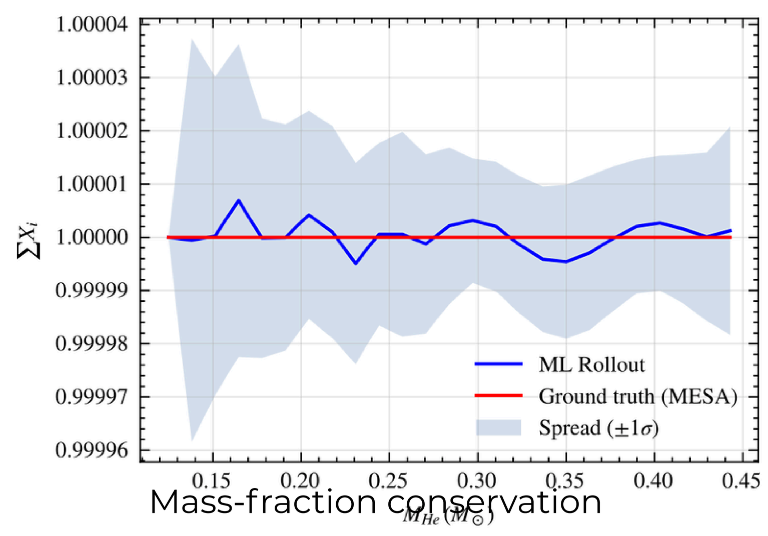

Mass-fraction conservation

Conclusion and Future directions

A fast, flexible complement to traditional codes

- Robust preprocessing pipeline for MESA outputs, resolving temporal alignment across evolutionary tracks.
- Demonstrates the potential of autoregressive prediction with machine learning for stellar evolution.
- Complements evolution codes and grid-based interpolation, enabling exploration of where modelling is impractical.
- Directly relevant to the high-throughput demands of missions like PLATO and large-scale population synthesis.

NEXT
Ongoing work extends this framework to full radial structure profiles, toward complete internal-structure emulation.

Acknowledgements
The authors thank the FCT for financial support to CENTRA and INESC-ID plurianual.

References
MESA: Paxton et al. 2011, 2013, 2015, 2018, 2019; Jermyn et al. 2023, ApJS
PLATO: Rauer et al. 2014, Experimental Astronomy, 38, 249