Emma Torres Chickles

Cambridge, MA | echickle@mit.edu | GitHub: github.com/emmachickles | emmachickles.github.io | LinkedIn: linkedin.com/in/emma-chickles

Summary

ML researcher working on generative priors, inverse problems, and representation learning at survey scale. My research searches hundreds of millions of stellar light curves for binary stars that orbit in minutes, and much of the work is establishing what those measurements can and cannot support — building the diagnostics that separate a reconstruction from an assumption, and the visualizations that make the difference legible.

Experience

PhD Researcher — MIT Department of Physics Aug 2021–Present

Generative priors and inverse problems, representation learning, and large-scale time-series ML

- Built a **score-based diffusion prior** over 2-D physical fields and coupled it to a known linear forward operator for posterior sampling; benchmarked four samplers (**DPS**, **DAPS**, **PnP-DM**, Tikhonov) and showed reconstruction quality is **sampler-agnostic but prior-dominated** (fidelity 0.74–0.78 across all four).
- Quantified **prior-induced error** that posterior uncertainty cannot express: two equally defensible learned priors agree to 21% on the data-constrained subspace while differing 61% on the operator's null space, and single-run error bars understate that disagreement by **1.6–2.4**× over a third to half of the reconstruction.
- Developed self-supervised continuous-time transformer and contrastive models for irregularly sampled time series, with diagnostics separating genuine temporal structure from class-template matching — showing strong downstream performance can be driven by template matching rather than genuine temporal understanding (ICML workshop, 2026).
- Benchmarked pretrained and from-scratch time-series foundation models (**Chronos-T5**, **MOMENT**, **PatchTST**) under matched-scale controls to isolate the effect of cross-domain pretraining on downstream performance.
- Built a streaming pipeline over **11 TB** of HDF5, indexing **1.4B+ sources** for self-supervised training on irregular, heteroscedastic real-world time series.
- Deployed a GPU-accelerated search across **16 NVIDIA A100s** via Slurm job arrays, processing **1.4M+ time series** and cutting end-to-end runtime from months to **∼6 hours**.
- Led discovery of two compact binary systems (ApJ 2025, 2026) — designed the search strategy, coordinated follow-up observations, and integrated heterogeneous datasets across varying cadences and noise regimes.

Research Assistant — Wellesley College Jun 2018–May 2021

Representation learning and computer vision

Trained convolutional autoencoders on ∼200k light curves with hyperparameter-sweep infrastructure for unsupervised representation learning; applied transfer learning with pretrained CNNs to classify terrain in orbital imagery (**92%** on ∼80k labeled maps).

Selected Projects

- **Starspot mapping as an inverse problem** — extended **InverseBench** (ICLR 2025) with a new scientific inverse problem: reconstructing a 4,608-parameter surface map from a single brightness time series on a real survey cadence. Showed the data constrain only **19 of 4,608** parameters (99.6% null space), so such reconstructions are overwhelmingly prior-driven, and that a mismatched prior yields confident structure *anti-correlated* with truth. **Best Visualization**, IAIFI Summer School hackathon (2026).
- **Interactive embedding explorer** — browser tool for inspecting learned representations, letting users compare how embedding spaces organize real physical structure versus survey systematics.

Selected Publications

- **Chickles, E.**, et al. *ICML AI4Physics Workshop* (2026, accepted) — Template matching, not time learning: a diagnostic for self-supervised light curve encoders
- **Chickles, E.**, et al. *ApJ* (2026) — An eclipsing 8.56-minute orbital period mass-transferring binary
- **Chickles, E.**, et al. *ApJ* (2025) — A gravitational-wave–detectable Type Ia supernova progenitor

Skills

Languages & tools: Python (PyTorch, NumPy, SciPy, Pandas, scikit-learn), Git, Slurm, LaTeX

ML: score-based/diffusion generative priors; Bayesian inverse problems and posterior sampling; uncertainty quantification and calibration; self-supervised & contrastive representation learning; transformers for irregularly sampled sequences; time-series foundation models

Scale: distributed multi-GPU training (PyTorch DDP), multi-node job arrays, TB-scale data pipelines (HDF5)

Awards

Best Visualization, IAIFI Summer School hackathon (2026) | MIT School of Science Community Service Fellowship (2026) | Principal Investigator, Magellan Telescope observing time, multiple competitive allocations (2022–2025) | Sarah Frances Whiting Medal for Achievement in Astronomy (Wellesley, 2020)

Leadership

Mentored 3 researchers (high school to graduate level), 2 projects yielding co-authored manuscripts | President, MIT Physics Graduate Student Council (2023–24) | President, MIT Physics Graduate LGBTQ+ Organization (2024–present)

Education

MIT — Ph.D. Physics, expected May 2027 | **Wellesley College** — B.A. Physics, *magna cum laude*, 2021