SELDON

An AI Foundation Model for Transients

10⁷ public alerts every night Rubin LSST, from 2026

→ **milliseconds** TO CHARACTERIZE EACH ONE

read the paper
arXiv:2603.04392

SELDON can process 10 million nightly alerts in seconds — forecasting future light curves and spectra that flow straight into follow-up prioritization.

G. Narayan[1,2,3] · J. O'Brien[1,2] · J. Wu[1,4] · N. Samia[1,4] · A. Wasserman[1,2] · J. Li[1,3] · R. Xu[1,4,7] · Z. Zhao[1,4] · H. Abunemeh[1,2] ·. M. Krafczyk[1,3] · S. Koranda[1,3] · T. A. Manning[3] · J. Caraballo[5] · D. Muthukrishna[6] · A. Gagliano[8] · D. Apley[4]

[1]NSF–Simons SkAI Institute · [2]University of Illinois Urbana–Champaign · [3]NCSA · [4]Northwestern University · [5]MIT · [6]Harvard–Smithsonian CfA · [7]Rutgers University · [8]Physical Super Intelligence

§01 / why — A survey-scale alert stream we cannot characterize fast enough.

10 million public alerts per night from Rubin's LSST — physics-based MCMC inference takes hours per object.

SELDON is a multi-modal, multi-survey foundation model for transients, developed at the NSF–Simons SkAI Institute. It is trained entirely on simulated LSST light curves and spectra.

SELDON delivers millisecond inference at survey scale, with physically interpretable rise times, decay rates, and peak fluxes that flow into follow-up prioritization.

Rubin Observatory Cerro Pachón, Chile. Survey starts 2026.

§02 / the model — SELDON architecture.

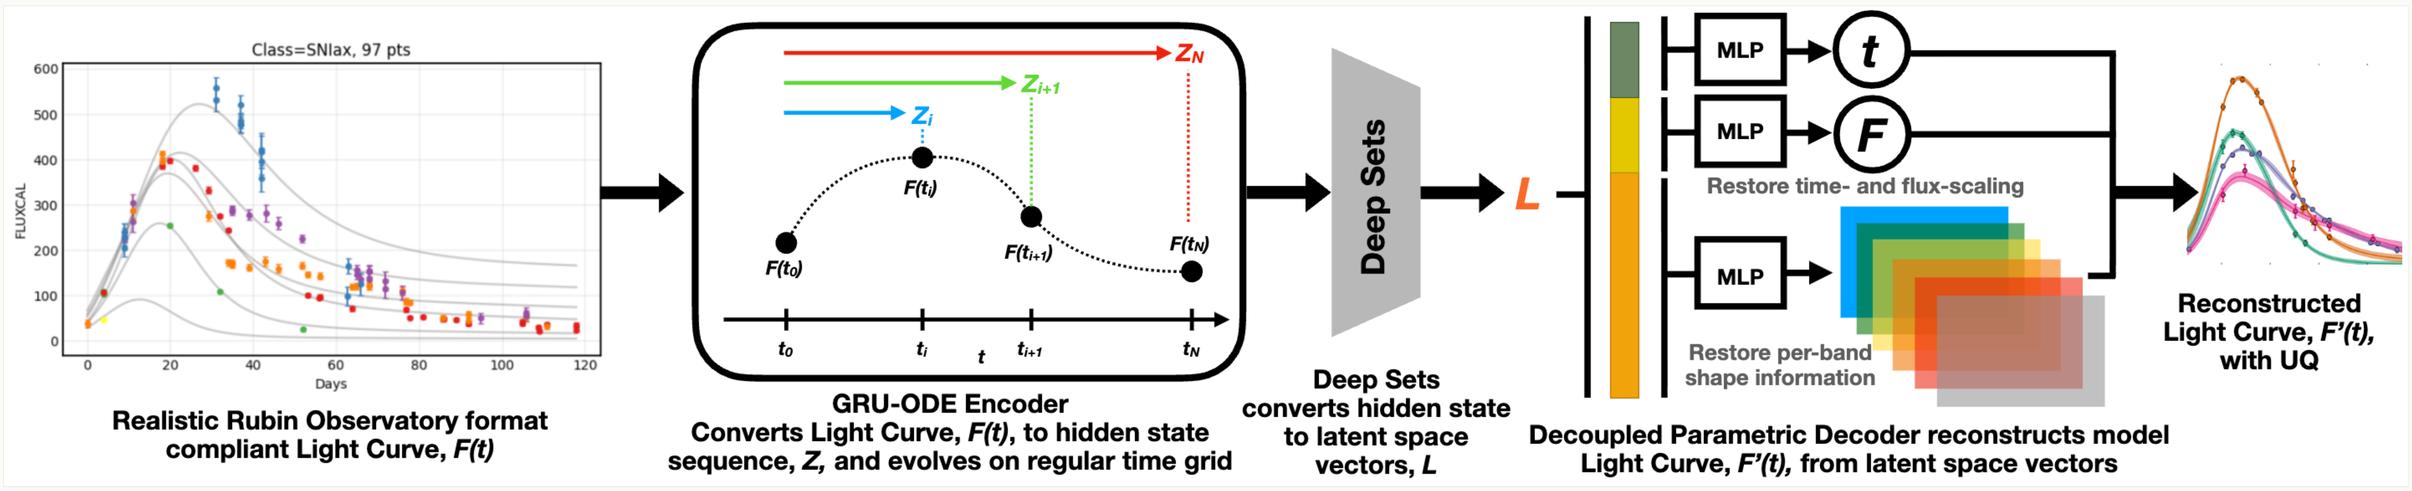

A — Masked GRU-ODE encoder

Summarizes sparse, irregularly sampled, heteroscedastic light-curve panels. Handles correlated multi-band observations even when only a handful of points are observed.

B — Latent neural ODE propagator

Integrates the hidden state forward in continuous time. Extrapolates to unseen epochs without retraining.

C — Gaussian-basis decoder

Outputs physically meaningful parameters — rise time, decay rate, peak flux. These feed downstream prioritization of spectroscopic follow-up.

Continuous-time · Sparse-aware · Physics-interpretable

§03 / what it does — Multi-survey, multi-modal inputs, one combined latent space, multiple downstream tasks.

TASK · 01/FORECASTING

Predict future evolution from partial observations.

10–35 % lower MAE Z-score error than the masked-GRU

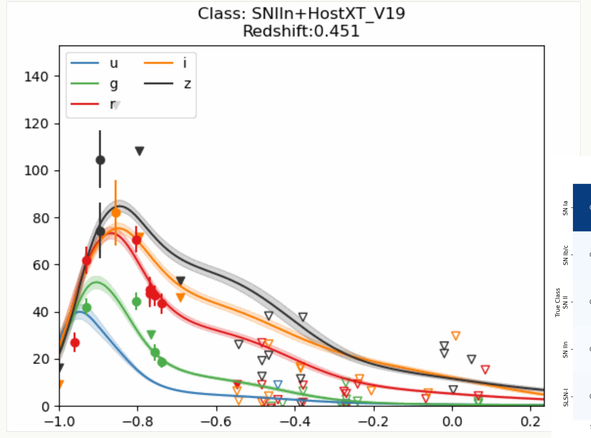

TASK · 02/CLASSIFICATION

Classify light curves or spectra

~10 % flat improvement over DASH on spectroscopic classification.

PRELIMINARY

Our light curve encoder is more mature than the spectral encoder, and we are still investigating architecture choices and optimizations. Still.

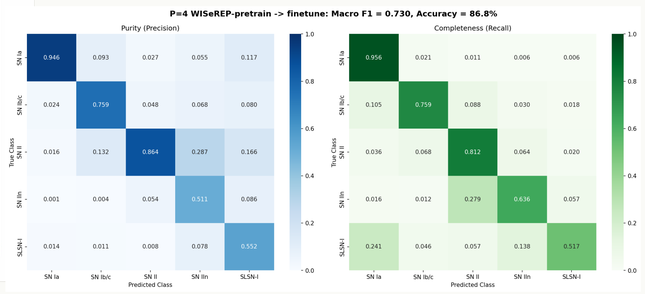

TASK · 03/RELIABLE UQ

Residual distributions on a shared task.

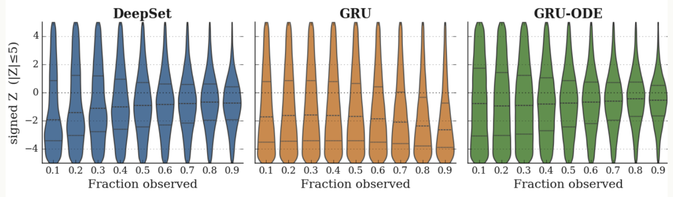

tighter SELDON's residuals are narrower and better-centered than Deep Sets and masked-GRUs (state-of-the-art for time-series)

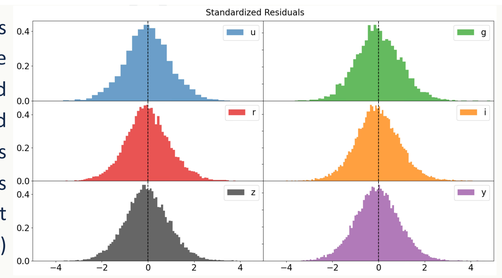

TASK · 04/SIMILARITY SEARCH

Find the nearest archival twin or model in milliseconds.

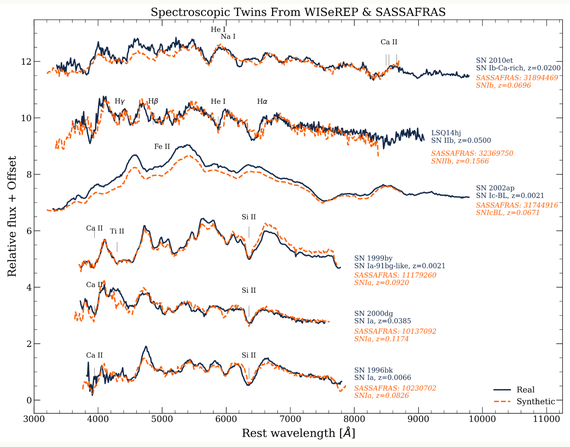

THROUGHPUT — **Process all of ELAsTiCC — 4 million objects — in a few seconds on a single H100.**

§04 / beyond light curves — Energy-based extensions and cross-survey transfer.

An **energy-based extension improves latent-space separation** between subtypes.

SELDON **reconstructs coarse spectral energy distributions** from light curves alone, a step toward inferring spectra where we have no spectroscopic coverage.

Domain adaptation carries simulated training to ZTF and YSE survey data, **without per-survey retraining**.

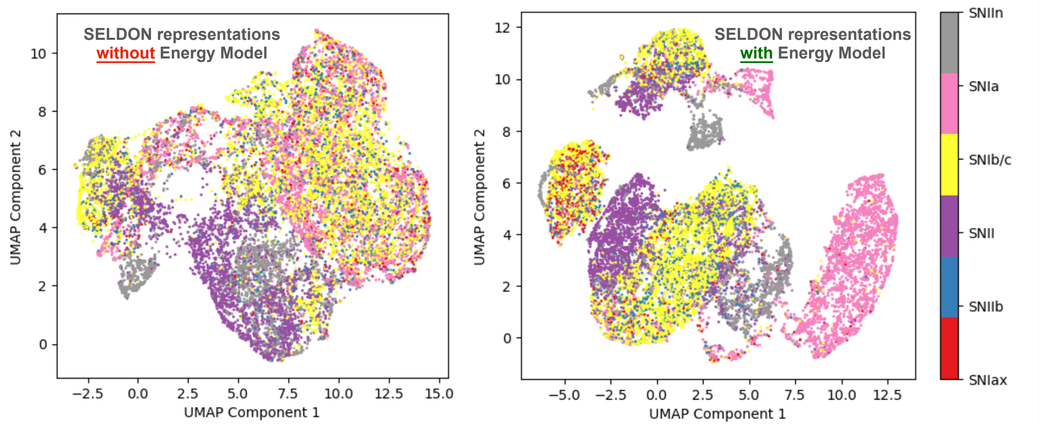

§05 / aion-1 — One source, multiple surveys, one model.

We are in the process of combining SELDON for time-series with a foundation model for the static sky – AION-1. The combined model will span Rubin, Roman, Euclid, ZTF, and YSE; combines photometry, spectra, and host-galaxy images; and shares a single representation across surveys and modalities.

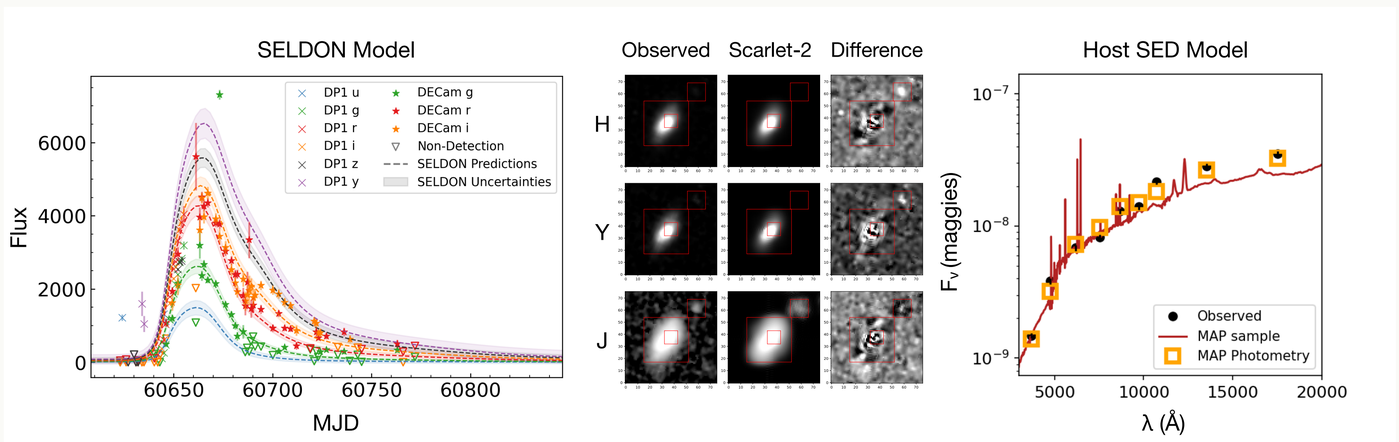

And this already works: SN 2024ahzc , jointly characterized with Scarlet2 + SELDON on Rubin DP1 + Euclid Q1 + DECam imaging. One source, three surveys, one model.

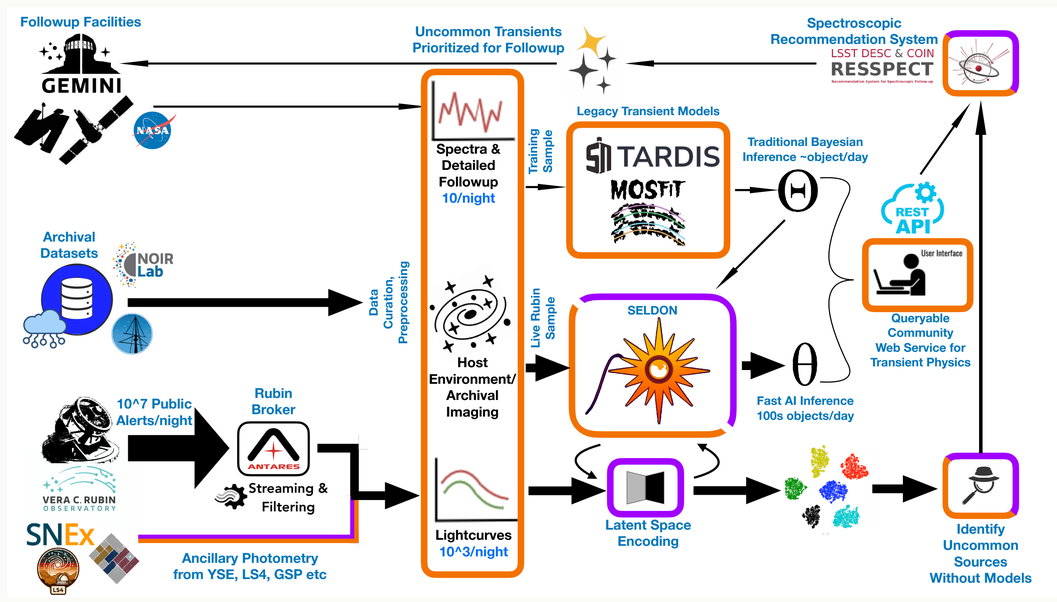

Funded by the U.S. National Science Foundation and the Simons Foundation

→ gsn@illinois.edu · https://gnarayan.github.io/