Attentive Neural Processes model LSST-like transient light curves more accurately—and 100,000x faster—than Gaussian Processes.

Siddharth Chaini,[1] Federica B. Bianco,[1, 2] Ashish Mahabal[3]
[1]University of Delaware, [2]Rubin Observatory, [3]Caltech

LSST's cadence is a challenge for transient science

- LSST will yield ~10M alerts/night, but individual light curves will be highly sparse, irregular, and multi-band
- Most downstream tasks for transient science (classification, anomaly detection) fail on raw, gappy data.

Gaussian Processes are the standard, but...

- Require an *a priori* kernel choice (i.e., if & when we already know the physics)
- Must be fit to each object individually (slow!)

Attentive Neural Process: Amortized Inference

(1) Meta-training phase on 15,000 simulated transients

(2) Inference requires just a single forward pass

Meta-learning with Attentive Neural Processes

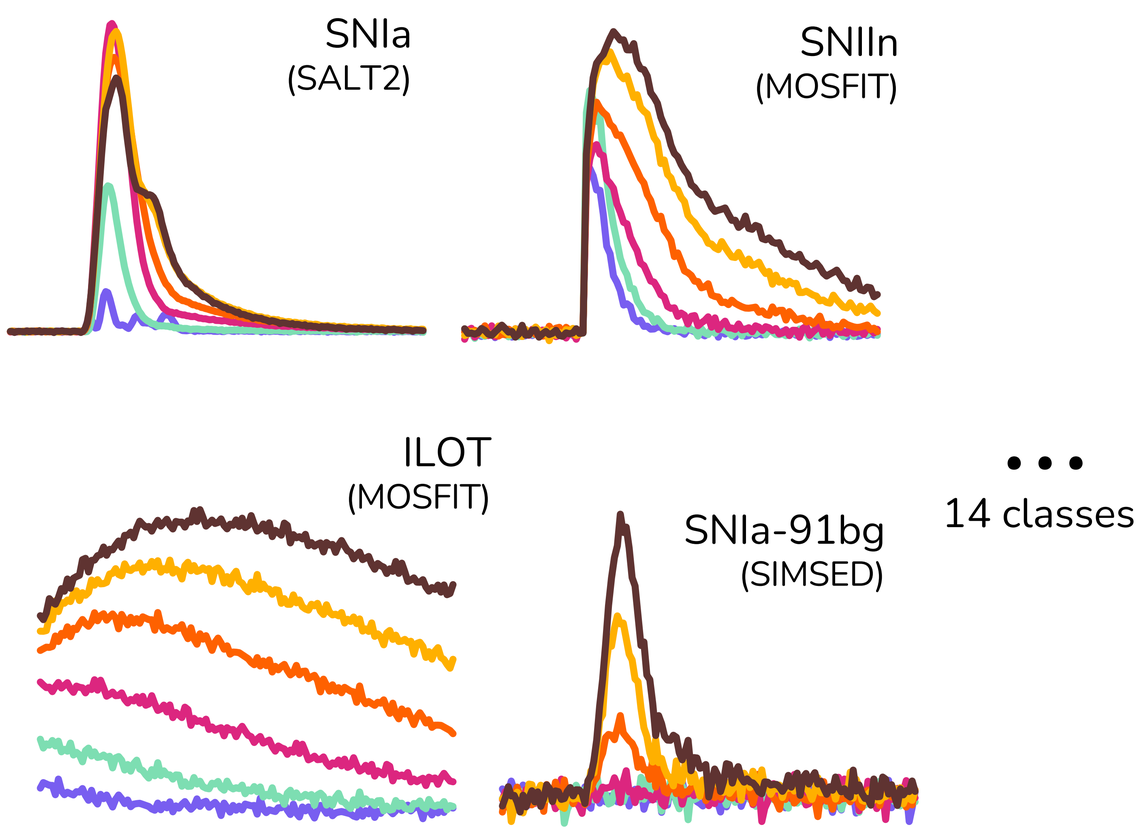

Unlike other methods, we do not assume a physical model. A single ANP learns a shared prior across 14 transient classes (SNe, ILOTs, CARTs, etc.)

Please see our paper for more details!

NightLANP

PAPER

CODE

#80NSSC25K0312

Light curve interpolations: ANP v/s the rest

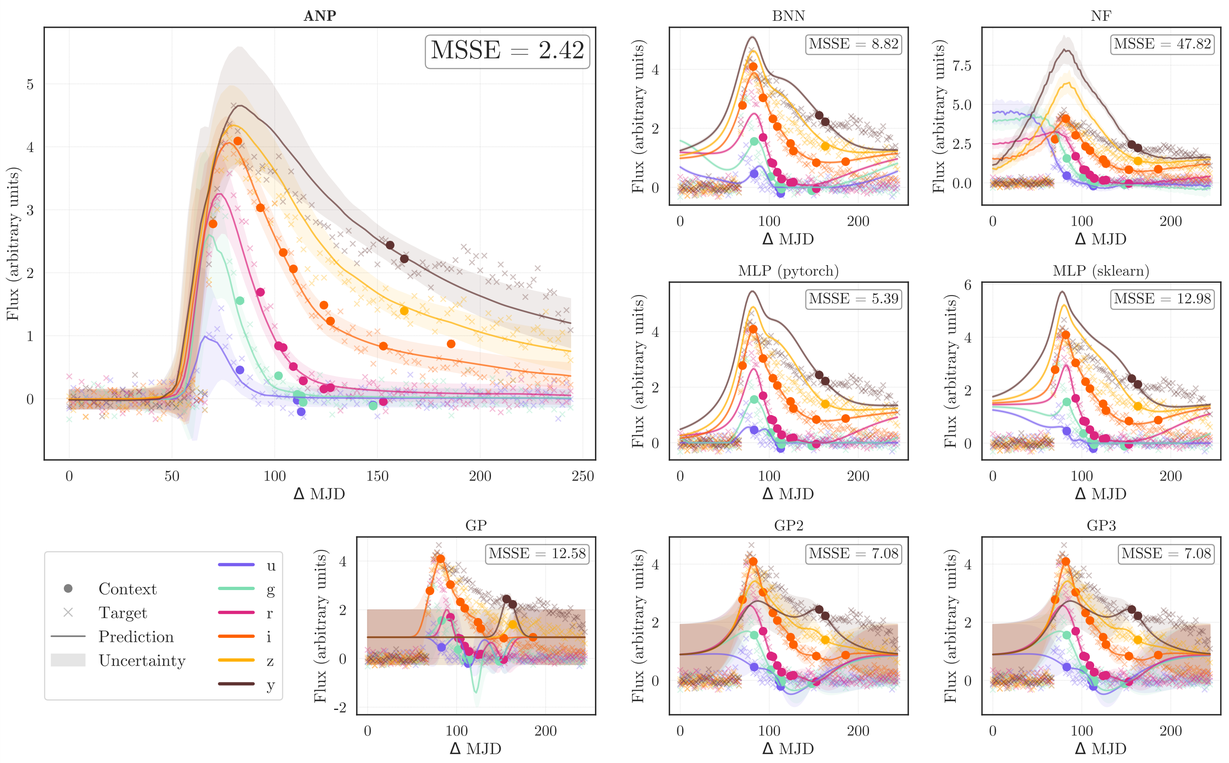

The ANP avoids the catastrophic overconfidence of Neural Networks and the under-confidence of GPs

The ANP succeeds across various metrics

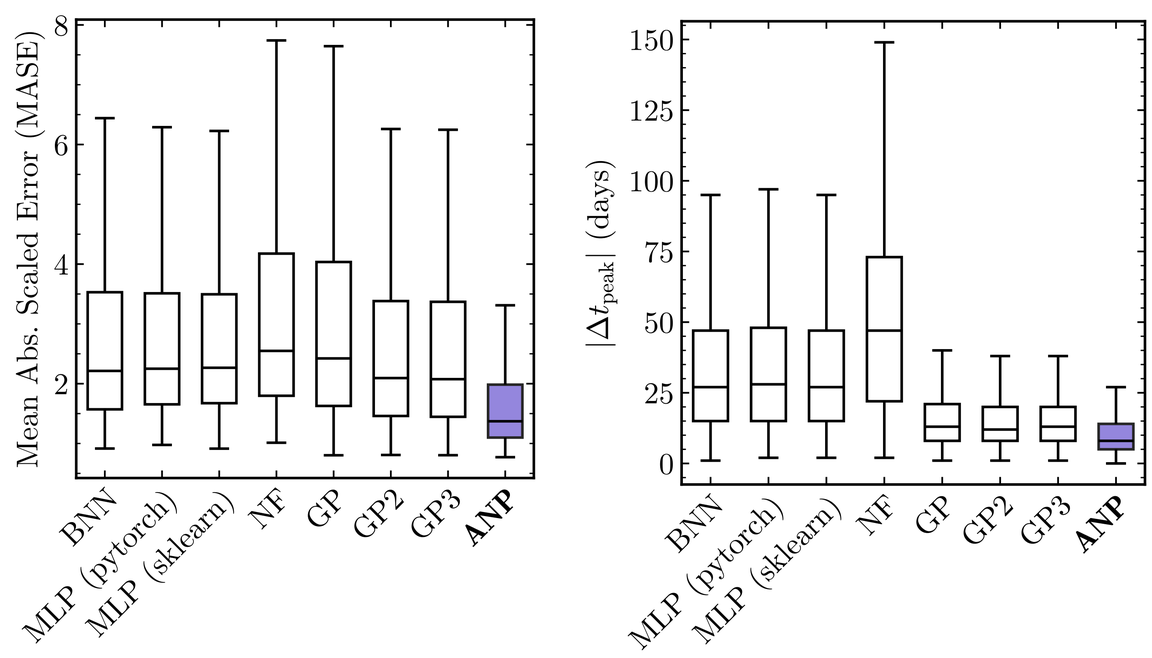

(a) Reconstruction Error

(b) Peak-time prediction

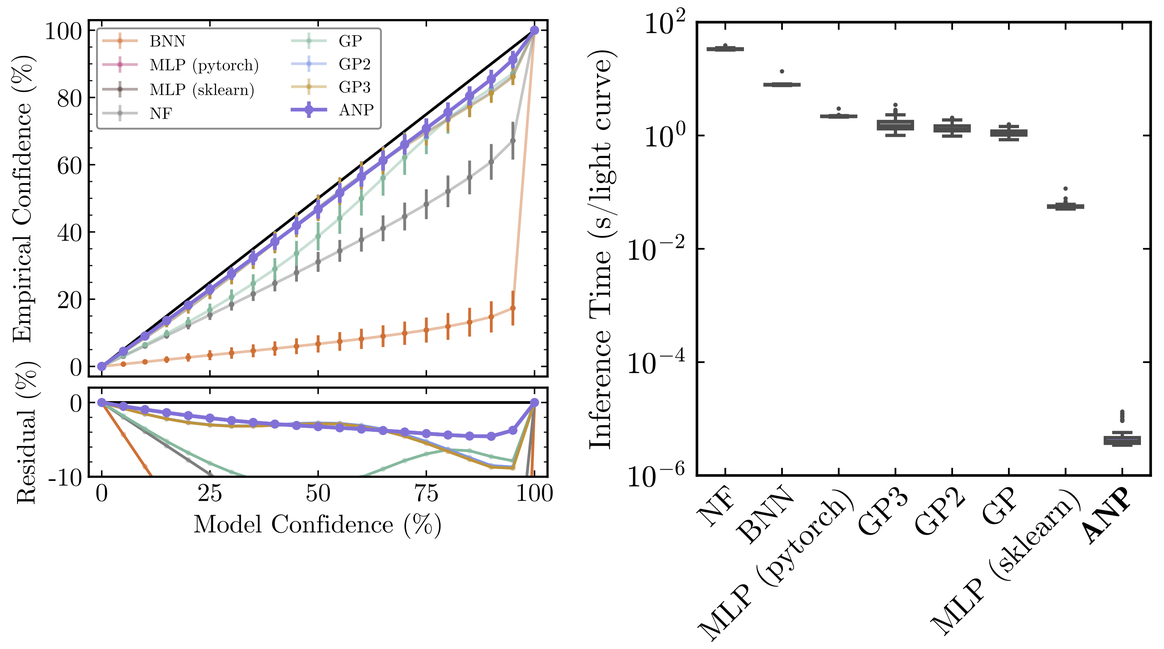

(c) Probabilistic Calibration

(d) Micro-second Inference speed
(~1 min for a night of LSST alerts!)