ML for statistically rigorous observing strategy optimization

Alex I. Malz
Carnegie Mellon University

Motivation: Observing strategy optimization

How can we choose the best possible observing strategy? (What does “best” even mean?)

- Every science case has its own metrics, with scaling properties in their own units.
- The choice is thus a convolution of nonuniform units and subjectivity of science goal importance.

Can we put them all on common footing to isolate the subjective portion of the decision?

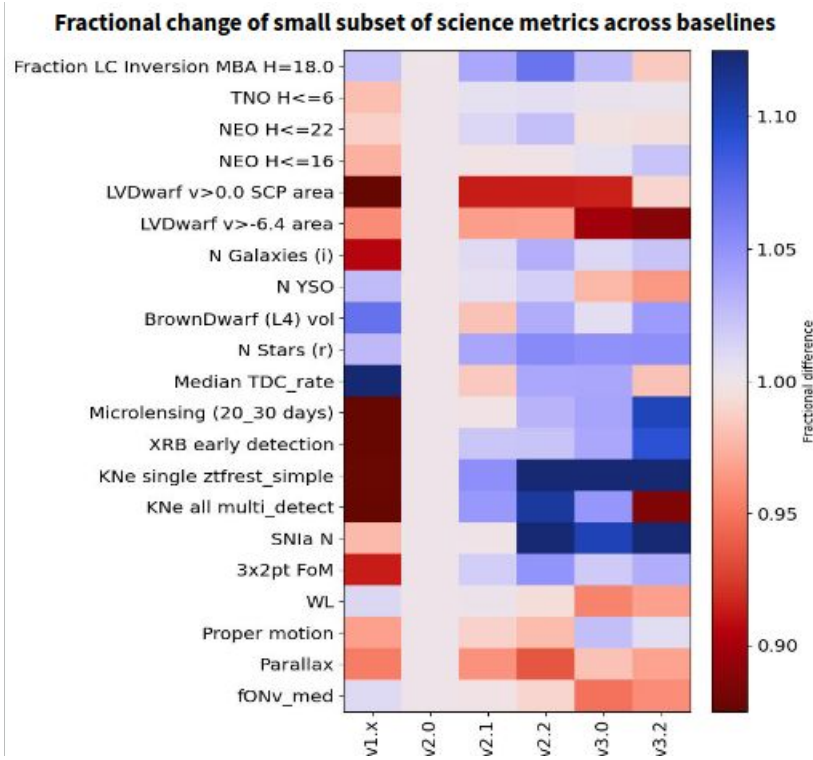

Credit: Lynne Jones

Motivation: Photometric redshifts

Why can't we just add another metric for photo-z?

- There are many estimation algorithms to choose from, with no obvious forerunner.
- LSST's estimator (and prior) has not yet been chosen.

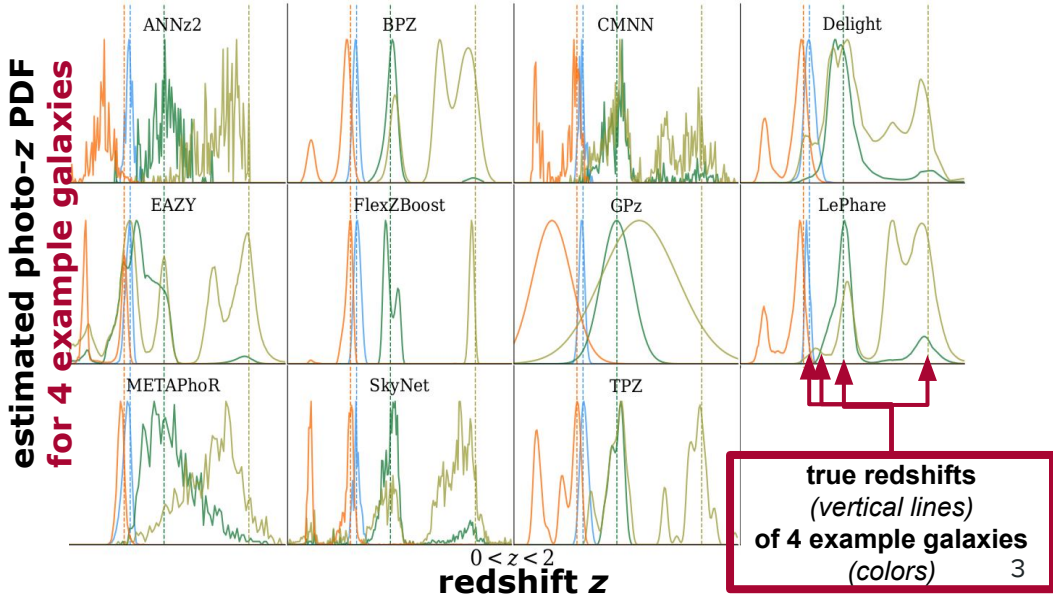

Schmidt&Malz+2001.03621

Can we derive a metric independent of the estimation algorithm?

Oh, and the implicit prior matters! See Malz&Hogg2007.12178

Alex I. Malz
aimalz@cmu.edu

What is TheLastMetric?

A science-agnostic information-theoretic metric, **ת**, is the variational lower bound on the mutual information between redshift and photometry

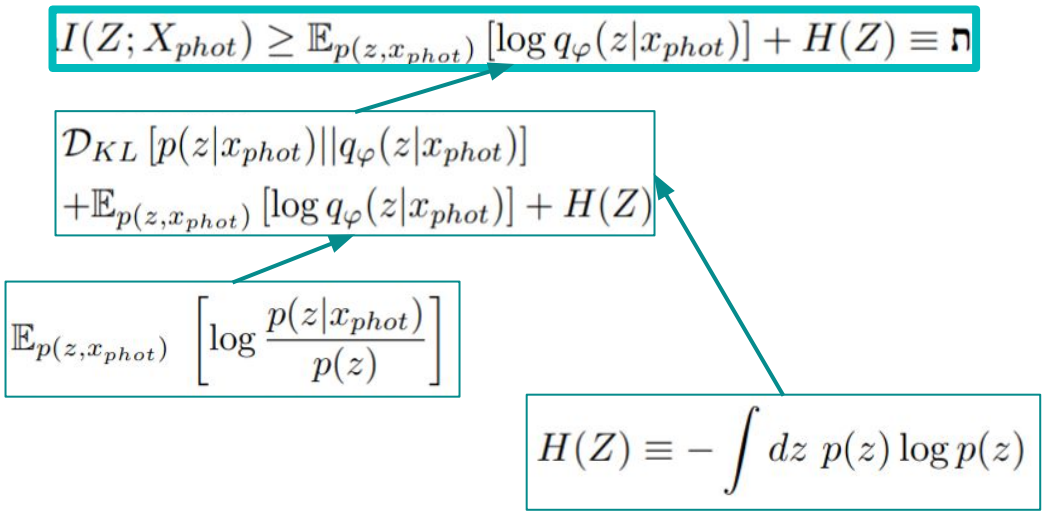

A variational approximation to **ת** uses the ***pzflow*** conditional normalizing flow (*github.com/jfcrenshaw/pzflow*)

with loss

$$\mathcal{L} = -\mathbb{E}_{p(x,y)}\left[\log q_\varphi(x|y)\right]$$

for variational approximation

$$q_\varphi(x|y) = p(z = f_\varphi^{-1}(x;y))\left|\det \frac{\partial f_\varphi}{\partial x}(x;y)\right|^{-1}$$

Alex I. Malz aimalz@cmu.edu

Demonstration: goal

Explore the behavior of π relative to established metrics

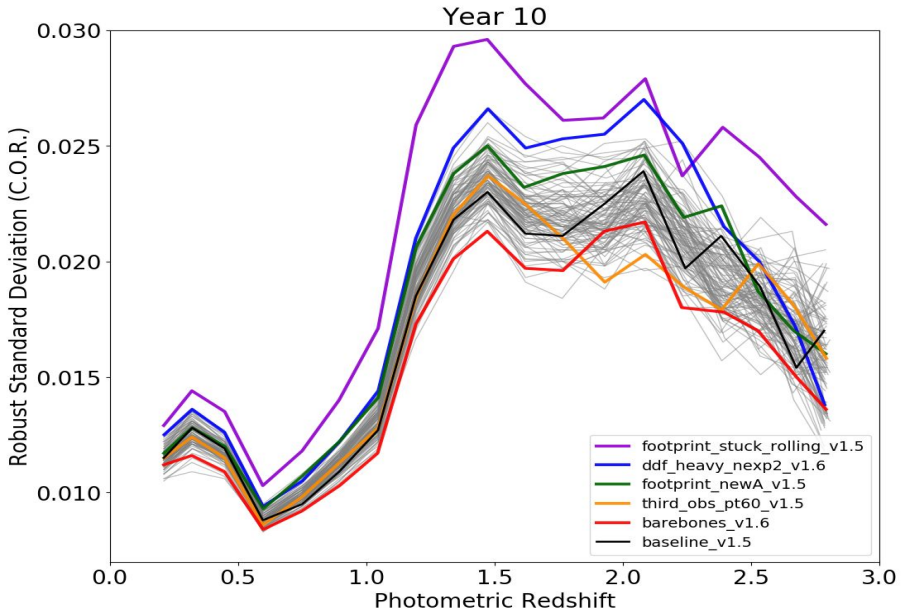

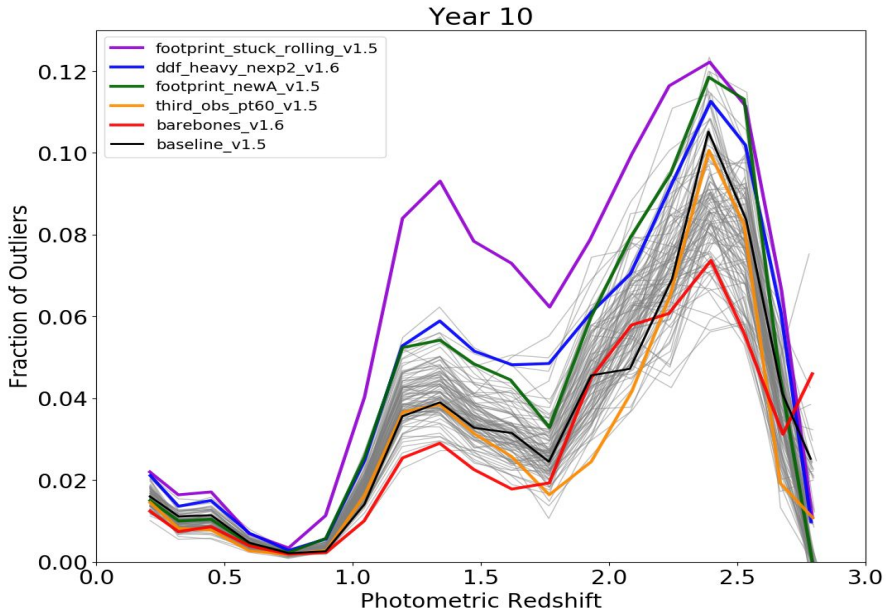

Demonstration: experimental conditions

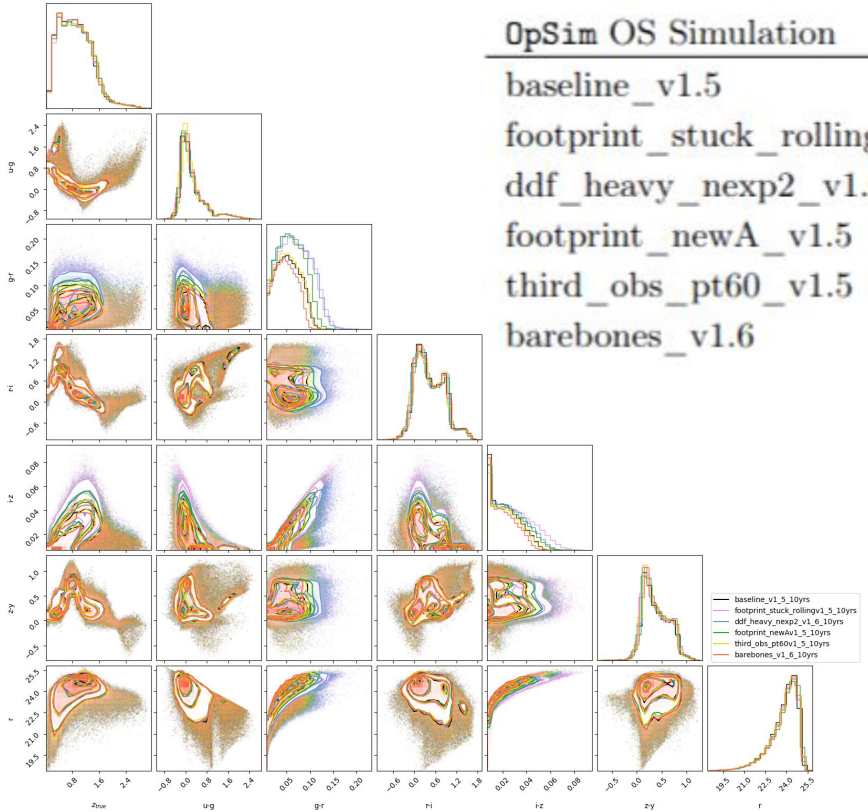

OpSim OS Simulation	m5_u	m5_g	m5_r	m5_i	m5_z	m5_y
baseline_v1.5	25.86	27.02	26.99	26.42	25.70	24.94
footprint_stuck_rolling_v1.5	25.56	26.68	26.62	26.06	25.33	24.61
ddf_heavy_nexp2_v1.6	25.57	26.82	26.84	26.26	25.57	24.82
footprint_newA_v1.5	25.75	26.87	26.85	26.29	25.55	24.78
third_obs_pt60_v1.5	25.87	27.03	26.99	26.43	25.70	24.93
barebones_v1.6	26.00	27.13	27.07	26.57	25.78	25.05

Demonstration: results

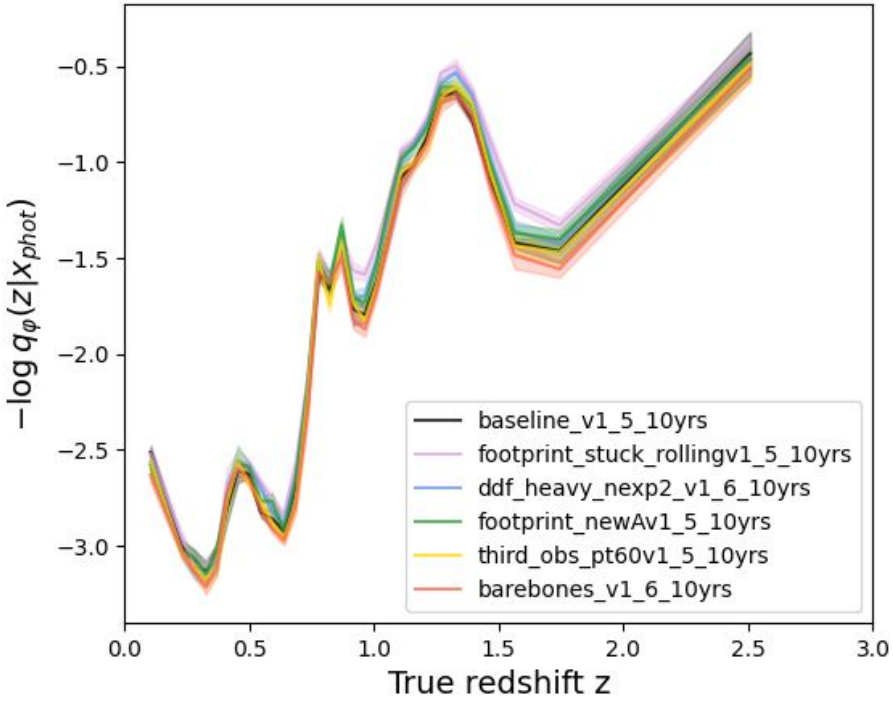

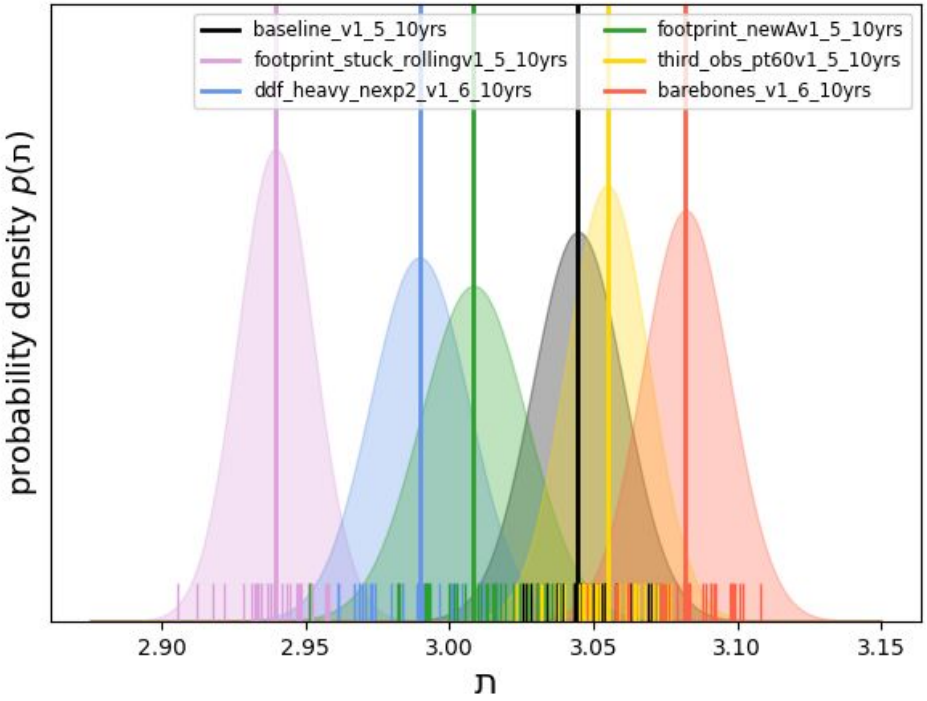

Demonstration: results

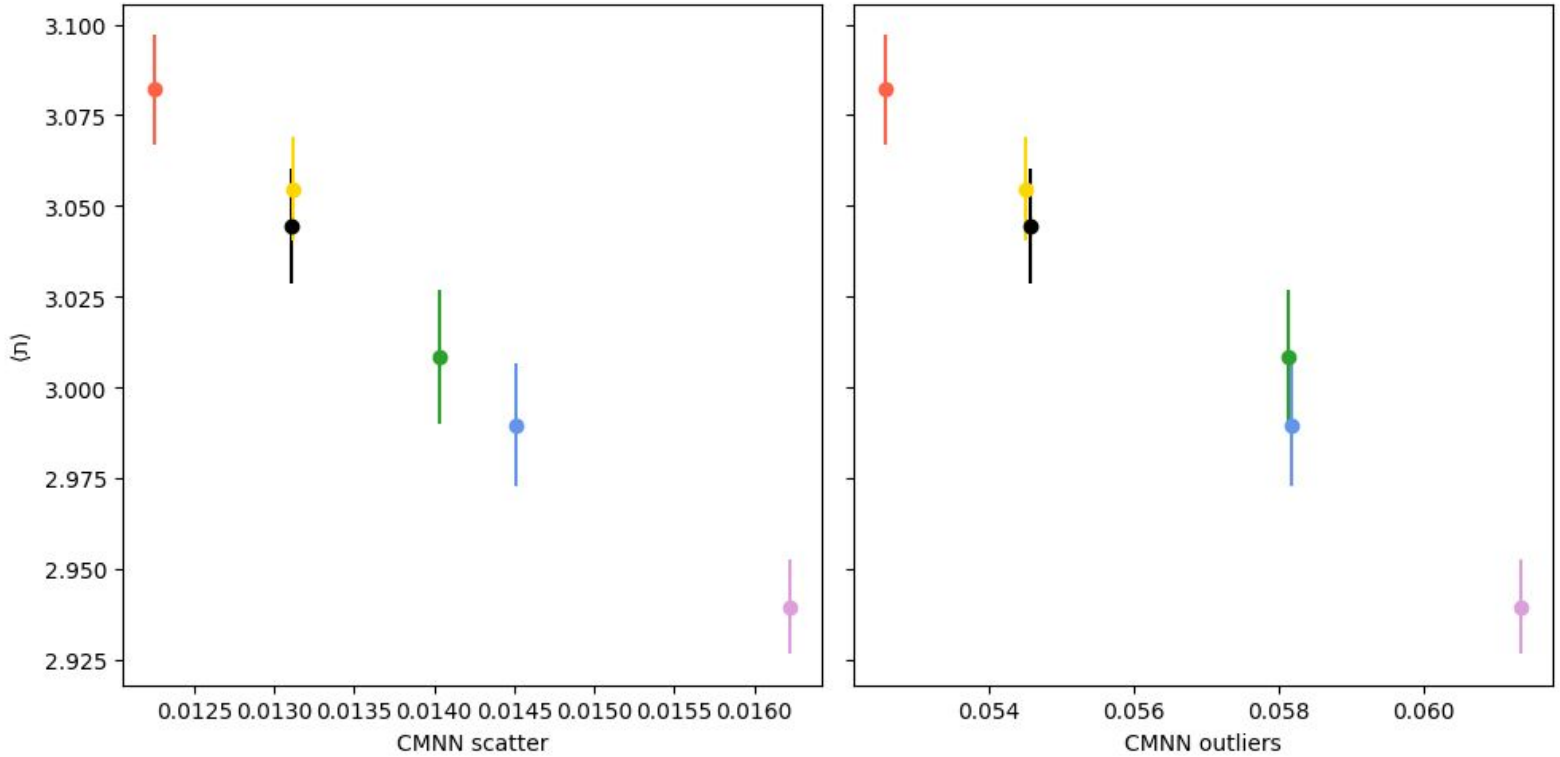

Summary

Malz, Lanusse, Crenshaw, Scott & Graham (sub. to ApJS) arXiv:2104.08229

Could ת really be "the last metric"?

- TheLastMetric quantifies the redshift information content of a photometric galaxy catalog.
- TheLastMetric distinguishes observing strategies without committing to any estimator (but it does have a prior).
- TheLastMetric has moderate computational cost, above the metrics of magnitude limits but below most photo-z estimators.

Status of ongoing work

- Currently using TheLastMetric to perform comparisons of joint photometry between Rubin and other surveys (w/Bryan Scott)
- Beginning to implement TheLastMetric in MAF (w/Xiaolong Li)

Next steps

- Adapt TheLastMetric to other scalar parameters derived from photometry.