Scanning for cosmological tensions across a DiRAC-enabled grid of models, datasets and samplers

Dily Duan Yi Ong
Kavli Institute for Cosmology
University of Cambridge

dlo26@cam.ac.uk

About me

- Dily Ong
- PhD Student at the University of Cambridge
- Dr. Will Handley's research group
- Research focuses on machine learning enhanced Bayesian inference in cosmological model selection and tension quantification
- Main author of the python package unimpeded

Content

- **Background and Motivation**: The Rise of Cosmological Tensions
- **Methodology**: MCMC vs nested sampling
- **The Solution**: Introducing the unimpeded package
- **Results**:
    - Tension statistics grid
    - Posterior samples from nested sampling runs
    - Machine learning emulators
- **Summary and Future Work**

Bayesian Inference in Cosmology

Parameter Estimation	Model Comparison	Tension Quantification
Using the observed data to estimate the values of cosmological parameters.	How much does the data support a particular model? Comparison between two models.	Do different datasets make consistent predictions from the same model? Study of discrepancies between different datasets.
E.g. $H_0, n_s, \sigma_8, \tau, \omega_b h^2, \omega_{cdm} h^2$	E.g. ΛCDM model vs dynamic dark energy cosmology	e.g. CMB vs Type IA supernovae data -> Hubble tension
Commonly uses the MCMC methods	Relies on the Bayesian evidence Z	Also needs evidence Z

Motivation: A Universe in Tension

- Precision cosmology => more datasets => more anomalies
- Major tensions:
  - H_0 Hubble tension (CMB vs SNe)
  - σ_8 tension (CMB vs cluster-galaxy weak lensing)
  - Ω_K tension (CMB vs lensing vs BAO)
- More cosmological models:
  - $\Lambda\mathrm{CDM}, K\Lambda\mathrm{CDM}, n\Lambda\mathrm{CDM}, m\Lambda\mathrm{CDM}, \mathrm{n_{run}}\Lambda\mathrm{CDM}, w\mathrm{CDM}, w_0 w_\mathrm{a}\Lambda\mathrm{CD}$

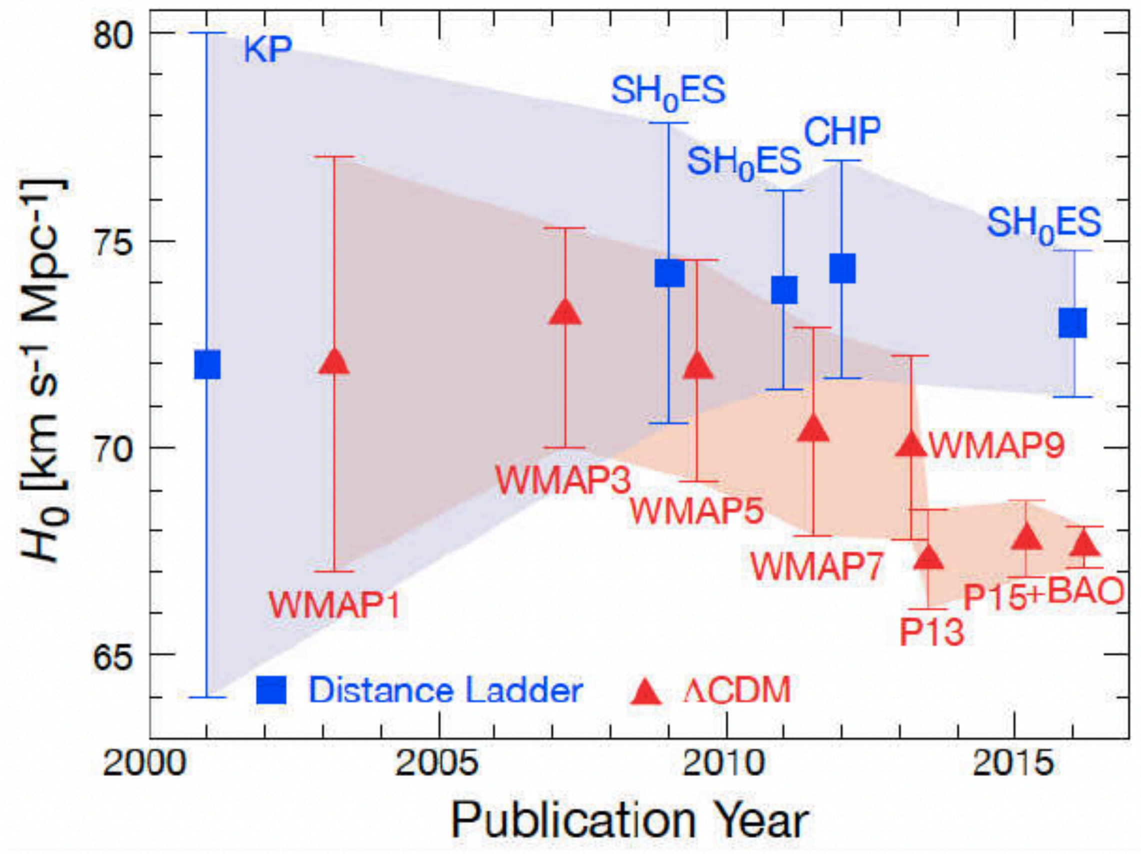

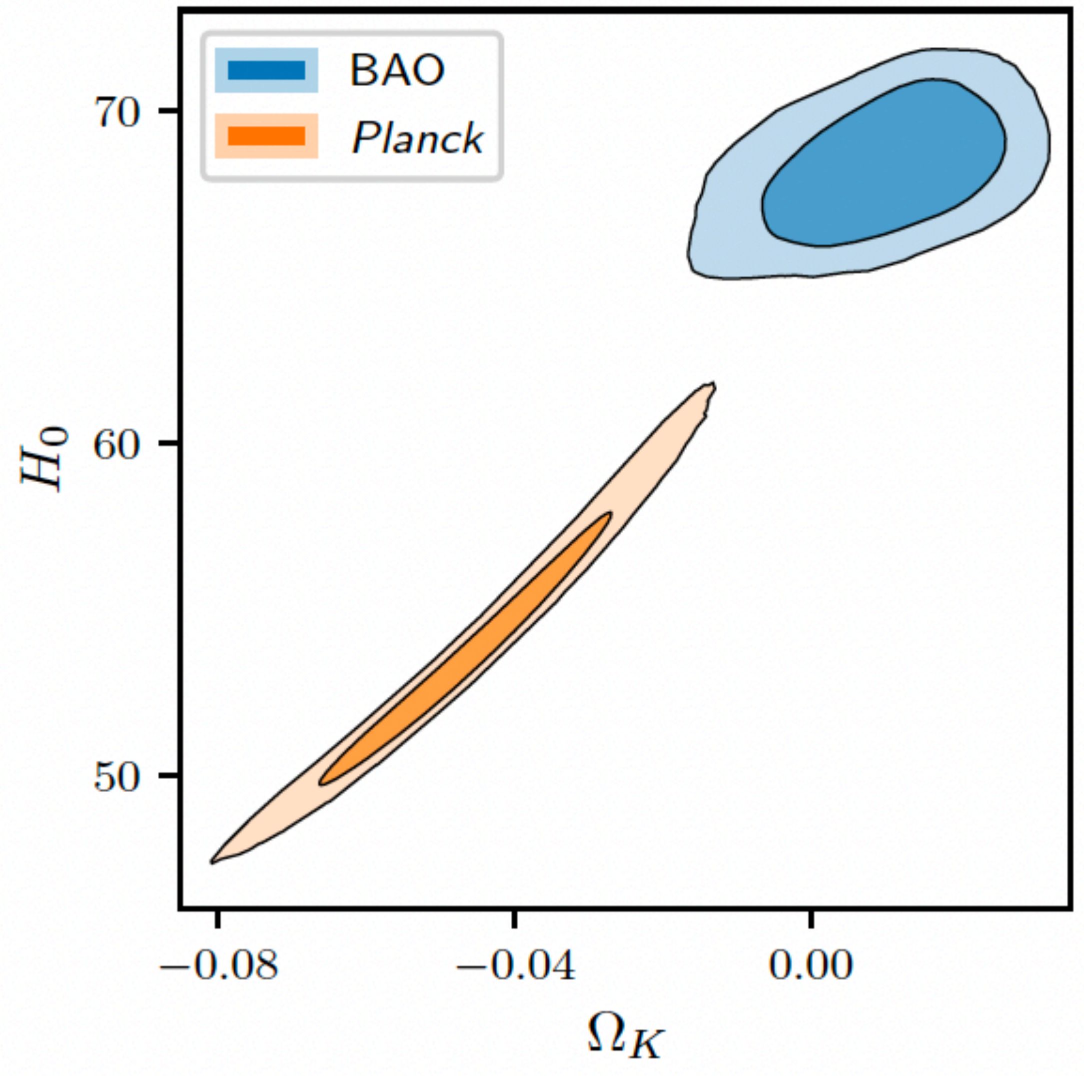

- Need: robust ways to compare, combine, and check consistency across many datasets and models

Our challenge

- The problem of cost
  - Model comparison and tension quantification are far more computationally expensive, than parameter estimation (traditionally performed by MCMC methods)
- The problem of scale
  - Dozens of nuisance parameters per dataset
  - E.g. likelihood from the Year 1 DES analysis, 20 out of 26 parameters are nuisance parameters, only 6 corresponding to cosmological parameters of interest

The Need for a New Archive

- Planck legacy archive's MCMC chains grid
- Great for parameter estimation
- BUT: cannot compute evidences in high dimensions from MCMC chains
  - Only parameter estimation
  - No model comparison and tension quantification

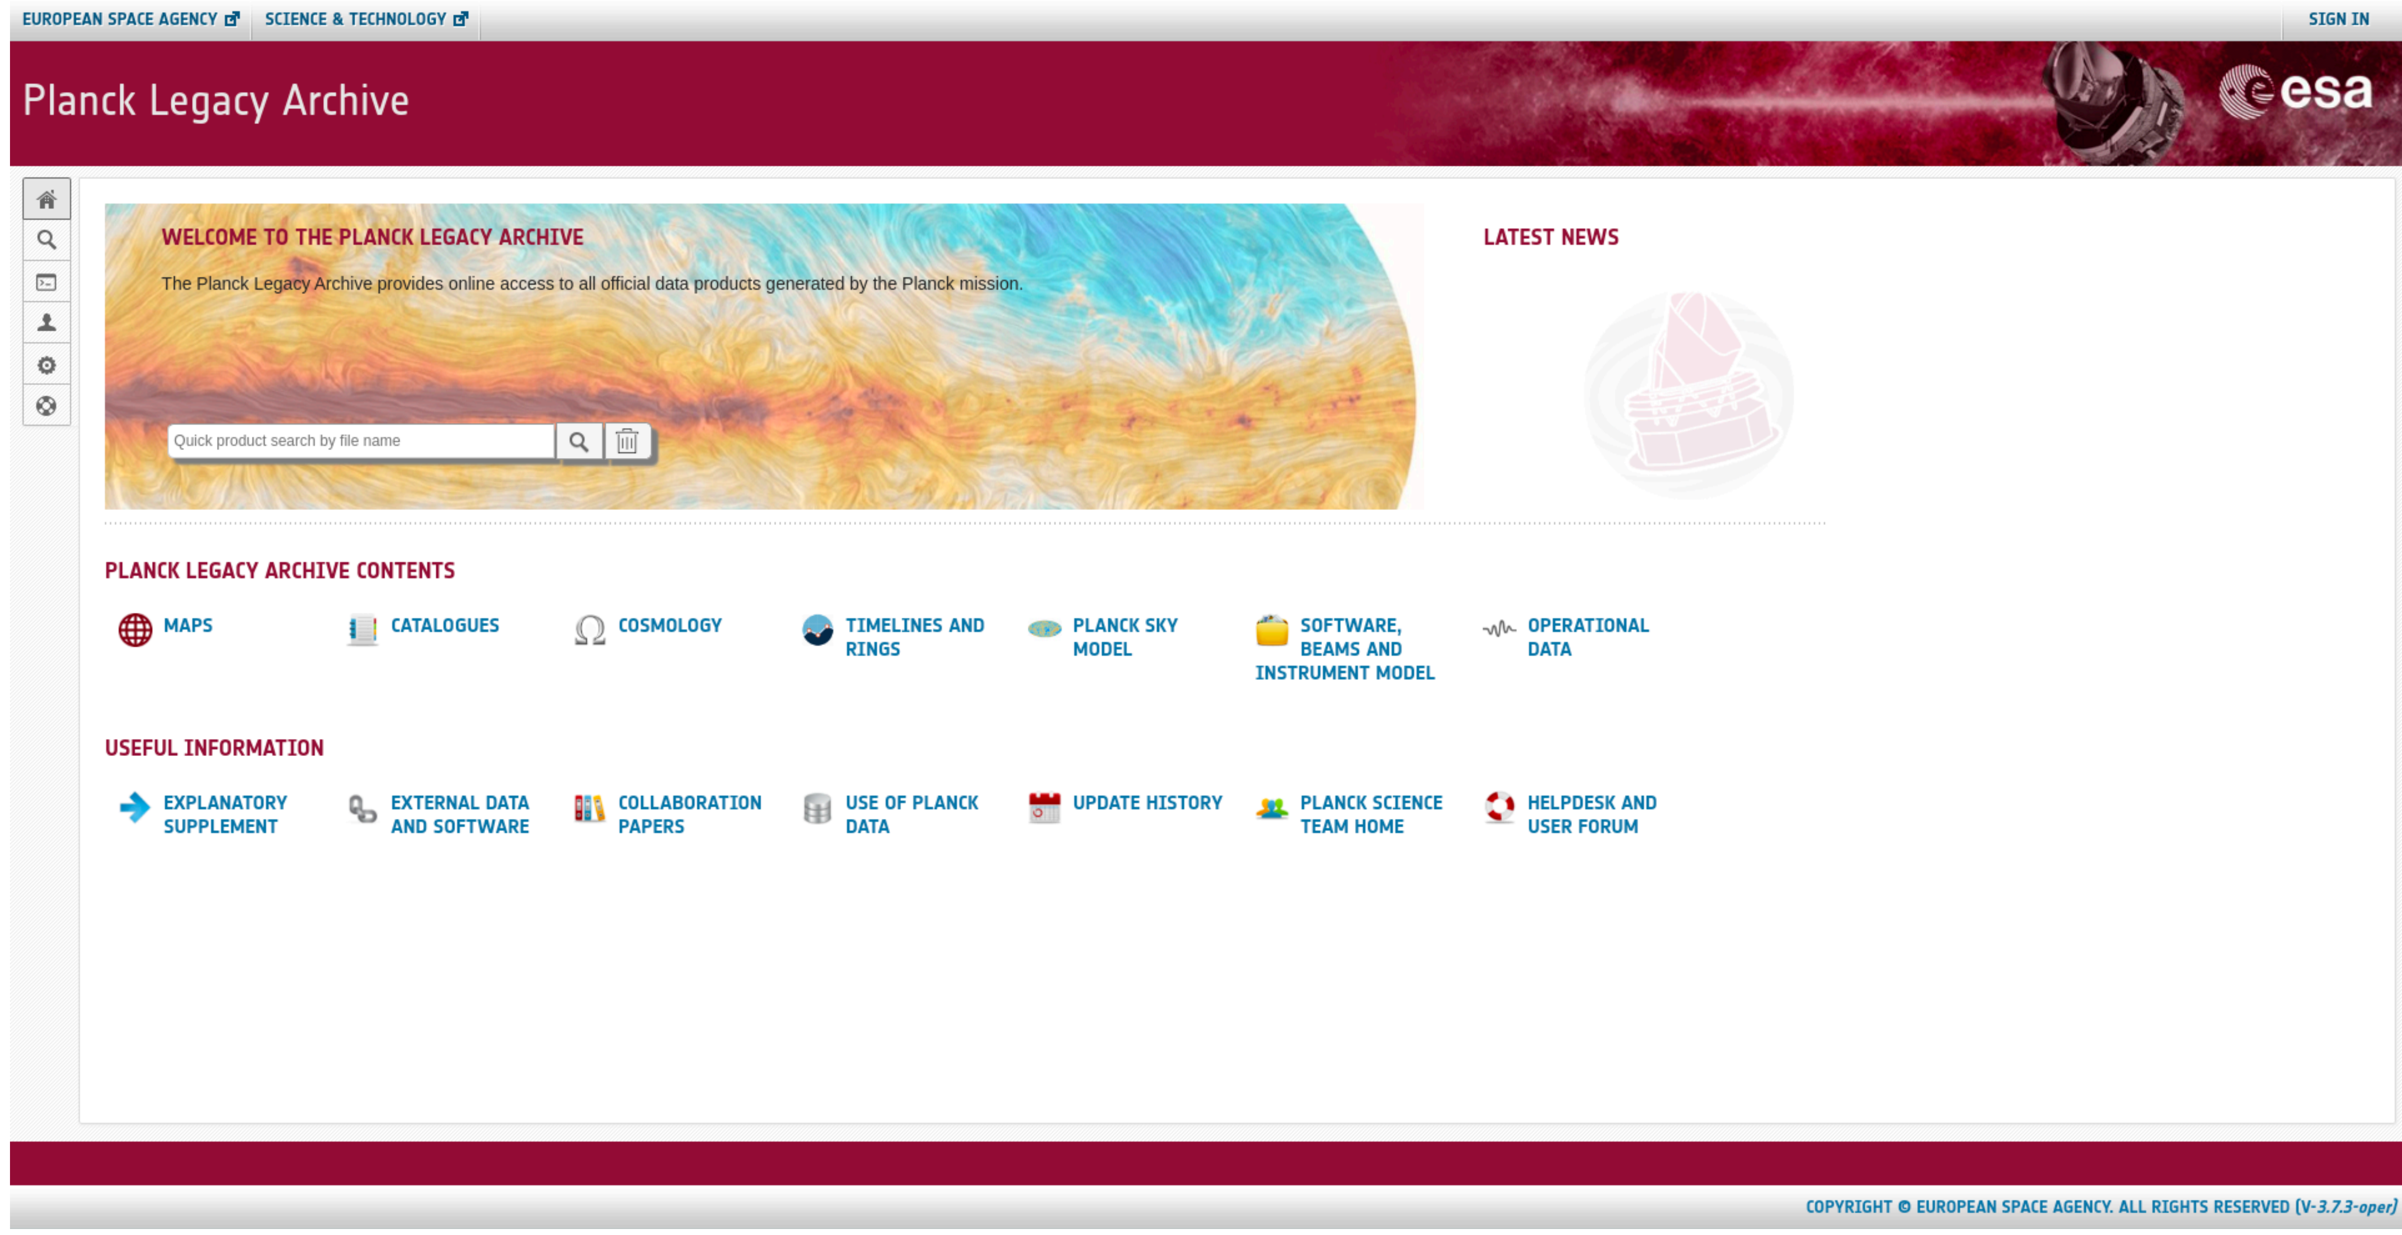

Planck legacy archive

What is Nested Sampling?

- Proposed by [Skilling 2006]

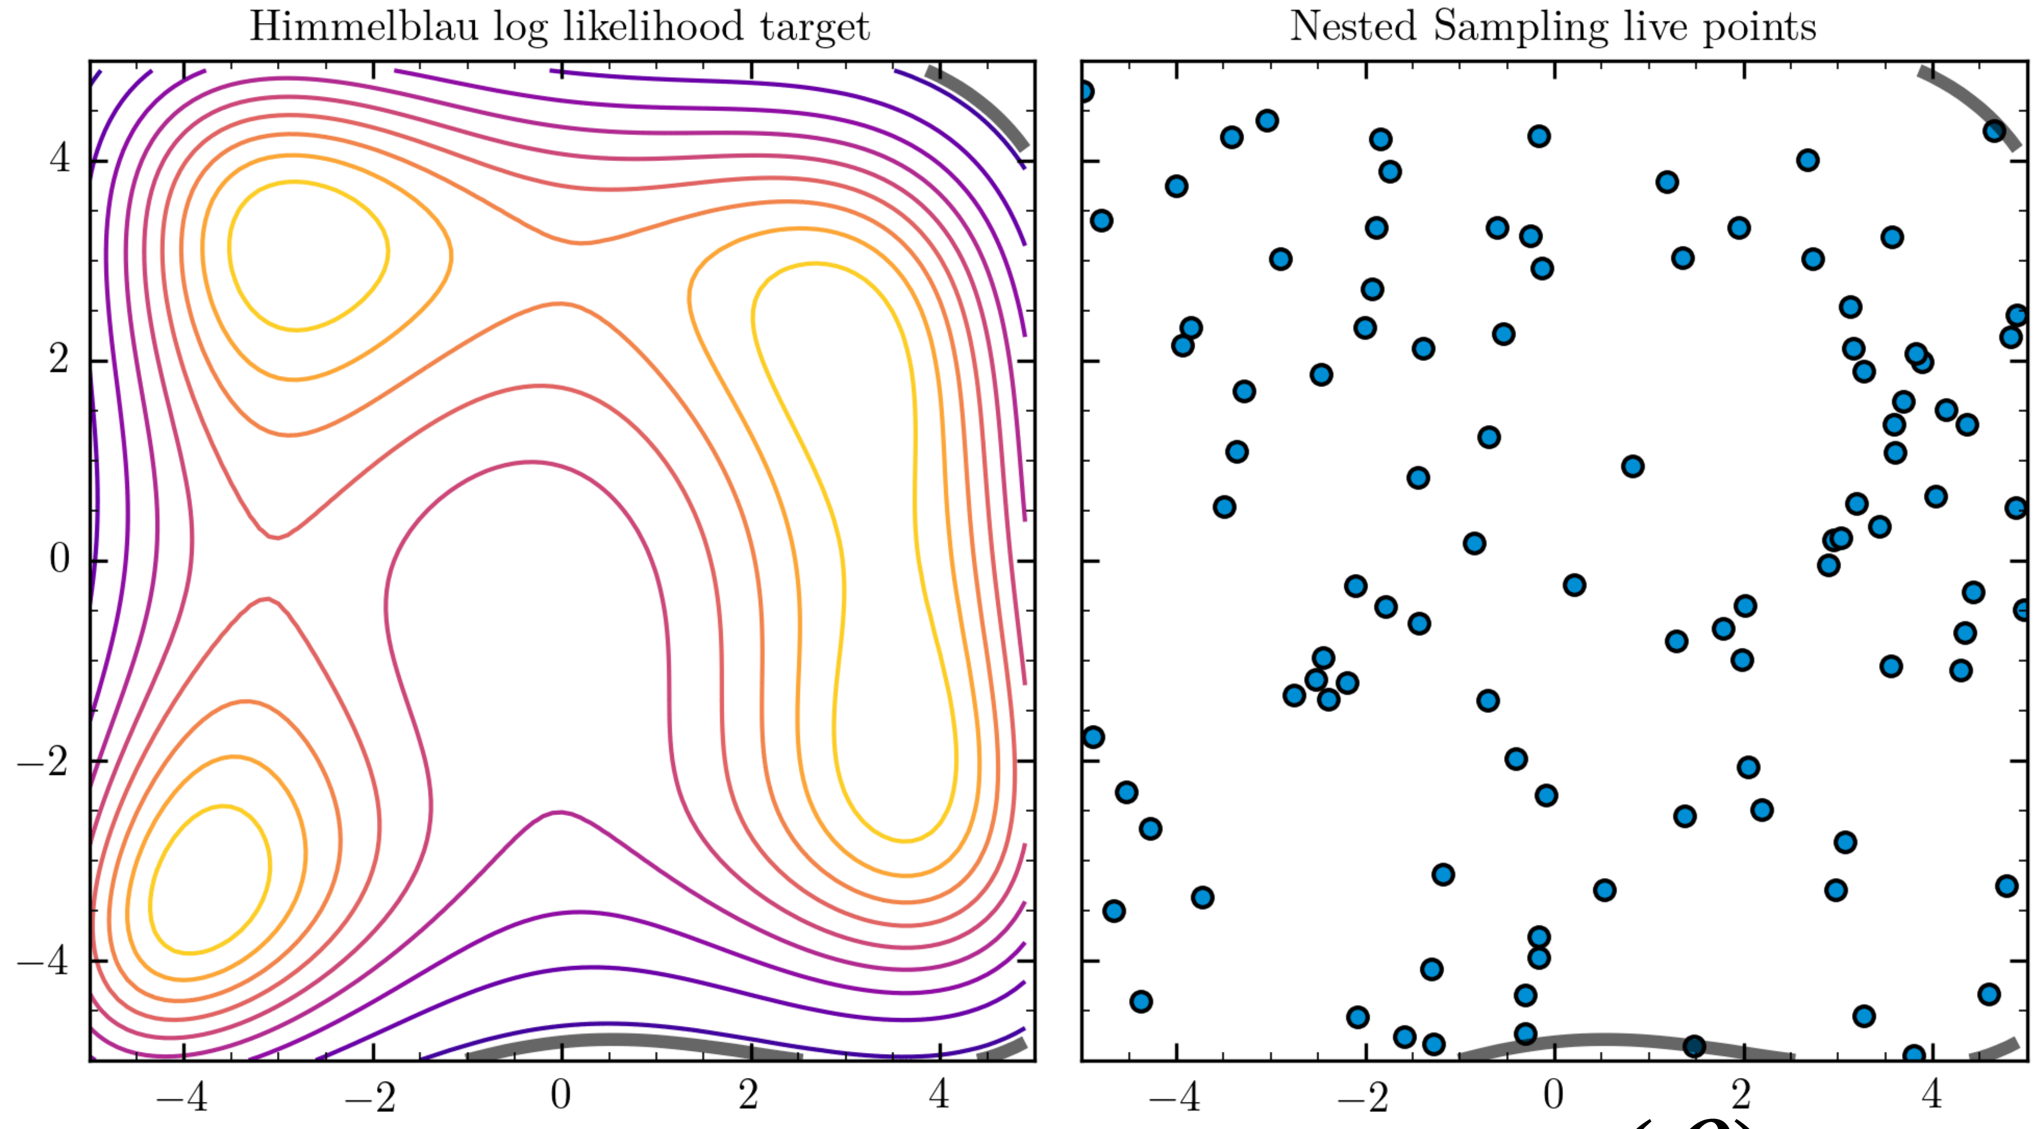

David Yallup

- The Goal: Solve the evidence integral:

$$Z = \int L(\theta)\pi(\theta)d\theta$$

- Algorithm:

1. Initialise: Start with a set of N "live points" drawn from the prior $\pi(\theta)$.
2. Iterate: At each step i:
    - Find the live point with the lowest likelihood, L_i.
    - Remove (kill) this point and store it. The remaining live points define a "nested" r higher likelihood.
    - Draw a new point from the prior, but constrained to have a likelihood
3. Terminate: Repeat until the entire prior volume has been traversed.

What is Nested Sampling?

- Transform this complex multi-dimensional integral into a simple 1-dimensional one by integrating over prior volume X instead of parameters θ.

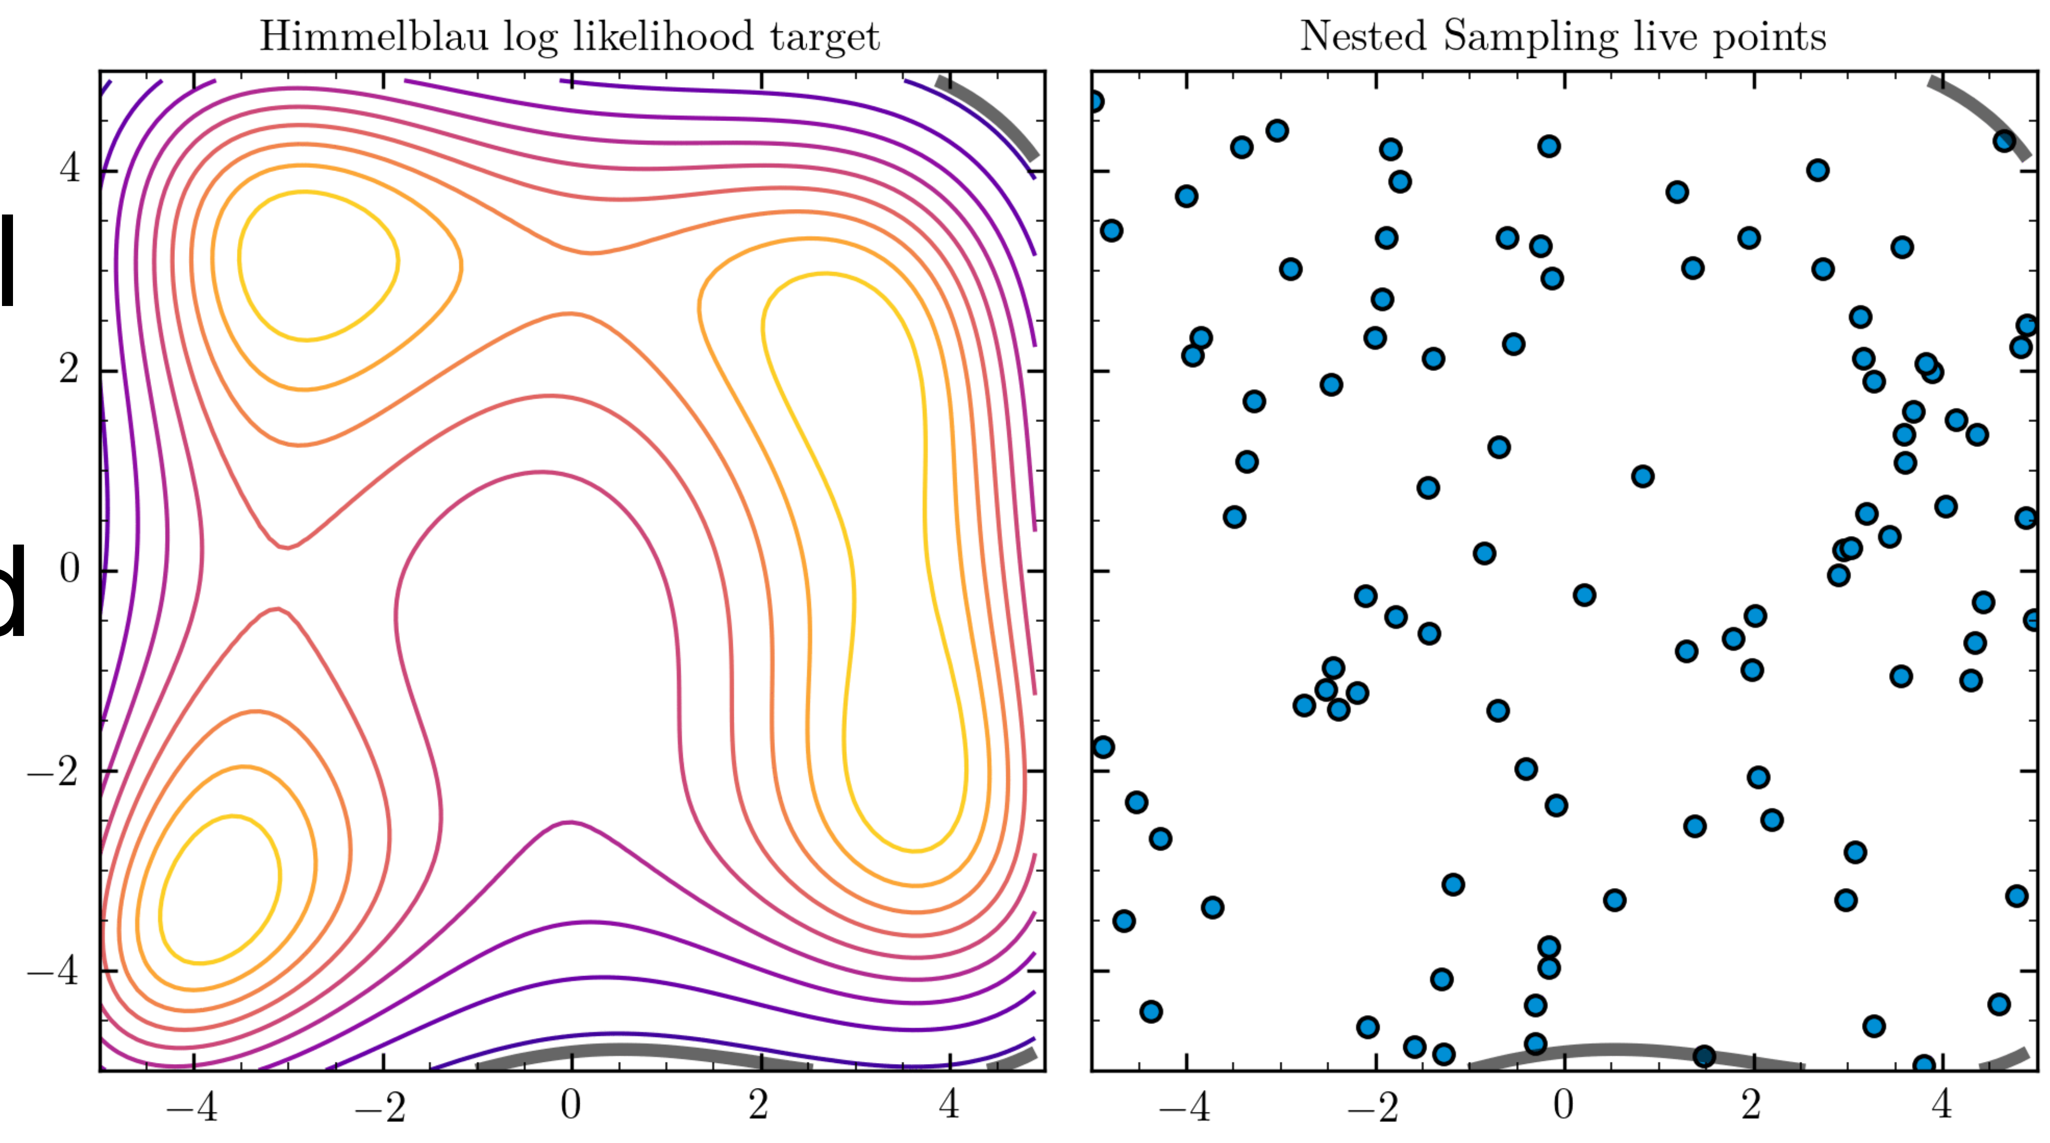

David Yallup

$$Z = \int L(\theta)\pi(\theta)d\theta \quad \Rightarrow \quad Z = \int_0^1 L(X)dX$$

Multi-dimensional   1-dimensional

What do we get?

- Numerically compute the evidence Z from the sequence of killed points
- Weighted set of posterior samples

Generating Samples: MCMC vs Nested Sampling

MCMC	Nested Sampling
• Single "walker" • Explores posterior • Cannot calculate the Bayesian Evidence Z • Use Case: Parameter Estimation, suspiciousness calculation • Fast, if proposal matrix is tuned	• Ensemble of "live points" • Scans from prior to peak of likelihood • Bayesian Evidence Z as its primary output • Use Case: Parameter Estimation, **Model Comparison, and Tension Quantification** • Slower, no tuning required

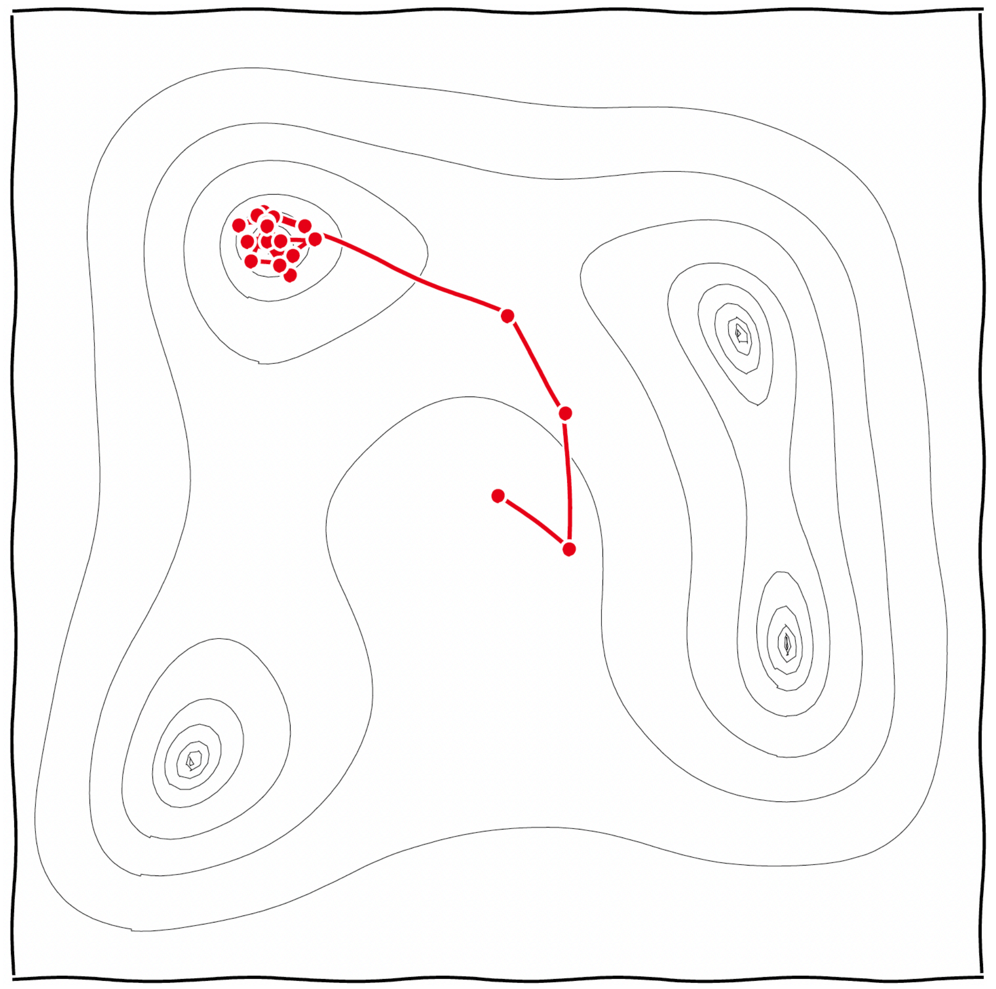

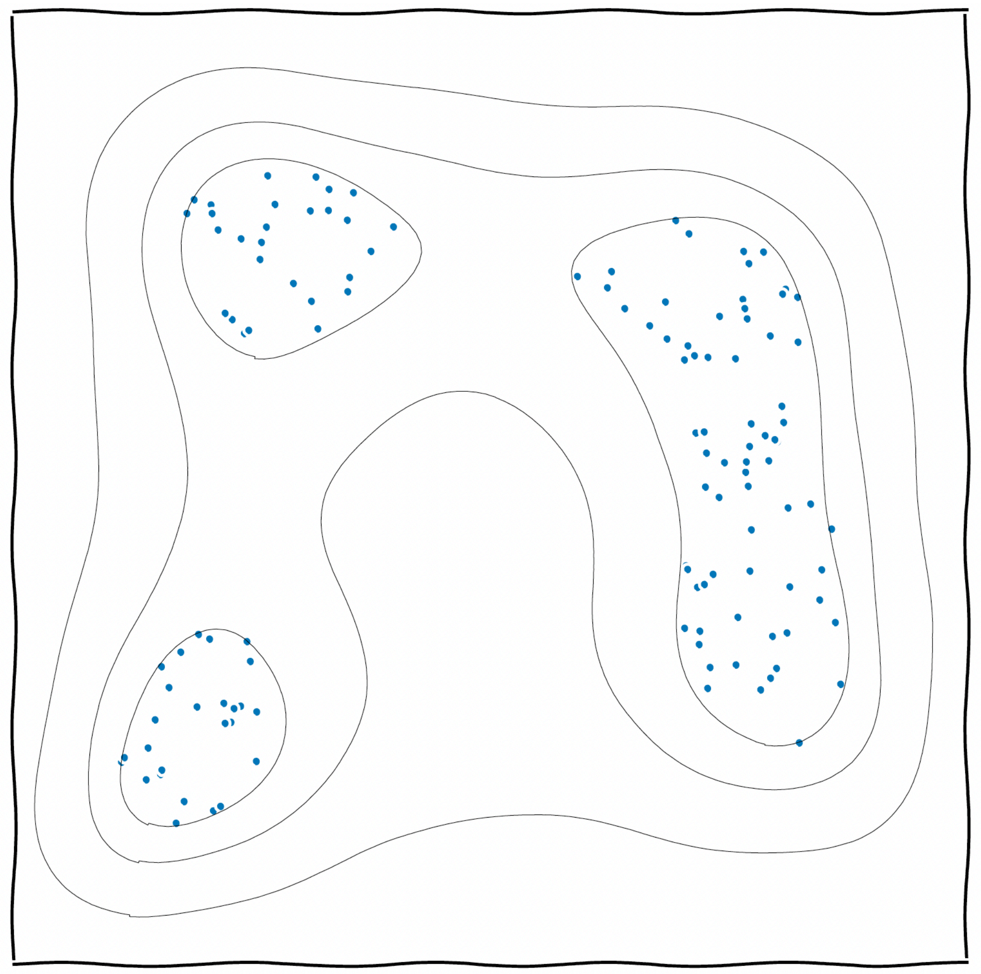

Introducing unimpeded

- A re-usable library of posterior samples, Nested Sampling chains, and machine learning emulators
- Systematic coverage of cosmological models/datasets, for tension/model comparison
- Downloadable posterior samples (MCMC or Nested Sampling)
- Downloadable machine Learning-powered “emulators” for fast reuse (e.g. Planck-like priors)
- Powered by DiRAC, hosted on Zenodo, pip-installable, available on GitHub
- Turns **weeks or months** of supercomputer time into **seconds** on your laptop

README   Code of conduct   MIT license

unimpeded: Universal model comparison & parameter estimation distributed over every dataset

unimpeded:	Universal model comparison & parameter estimation distributed over every dataset
Author:	Dily Ong & Will Handley
Version:	0.2.4
Homepage:	https://github.com/handley-lab/unimpeded
Documentation:	http://unimpeded.readthedocs.io/

Build Status codecov 100% docs passing pypi package 0.2.4 DOI 10.5281/zenodo.15686776 license MIT

GitHub

DiRAC
High Performance
Computing Facility

zenodo

unimpeded: Models and Datasets Coverage

- 10 cosmological models + 60 datasets & pairwise combinations (to be expanded)
- Mix-and-match any model/dataset, plus pairwise combinations

Cosmological models

- ΛCDM : H_0, τ_{reio}, $\Omega_b h^2$, $\Omega_c h^2$, A_s, n_s
- $K\Lambda$CDM : ΛCDM + Ω_K (varying curvature)
- $N\Lambda$CDM : Varying N_{eff} and total mass of 3 degenerate ν's
- $n\Lambda$CDM : Varying total mass of 3 degenerate ν's with N_{eff}=3.044
- $m\Lambda$CDM : Varying N_{eff} with two massless ν and one with m=0.06
- $\mathrm{n_{run}}\Lambda$CDM : ΛCDM + $\mathrm{n_{run}}$ (running of spectral index $dn_s/d\ln k$)
- wCDM : ΛCDM + w (constant cosmological equation of state)
- $w_0 w_a\Lambda$CDM : ΛCDM + w_0 + w_a (varying dark energy equation of state, CLP)
- $r\Lambda$CDM : ΛCDM + r (varying scalar-to-tensor ratio)

Cosmological datasets

- CMB:(Plik, Camspec, NPIPE, BICEP) $\pm$ CMB lensing
- BAO:SDSS, BOSS, eBOSS, Lyα, DESI
- SNe: Pantheon, SH0ES
- WL: DESY1

uimpeded: Easy Access to Complex Results

- Pip-installable

pip install unimpeded

- Open-source
  - unimpeded on Github
  - Chains and emulators on Zenodo

- Simple API in python for seamless download

- Easily access pre-computed results with a few lines of Python

README · Code of conduct · MIT license

unimpeded: Universal model comparison & parameter estimation distributed over every dataset

unimpeded:	Universal model comparison & parameter estimation distributed over every dataset
Author:	Dily Ong & Will Handley
Version:	0.2.4
Homepage:	https://github.com/handley-lab/unimpeded
Documentation:	http://unimpeded.readthedocs.io/

Build Status · codecov 100% · docs passing · pypi package 0.2.4 · DOI 10.5281/zenodo.15686776 · license MIT

```
from unimpeded.database import DatabaseExplorer
dbe = DatabaseExplorer()
planck = dbe.download_samples(
    method='ns', model='walcdm',
    dataset='planck_2018_CamSpec')
sdss = dbe.download_samples(
    method='ns', model='walcdm',
    dataset='bao.sdss_dr16')
planck_sdss = dbe.download_samples(
    method='ns', model='walcdm',
    dataset='bao.sdss_dr16+planck_2018_CamSpec')
```

unimpeded: Posterior samples

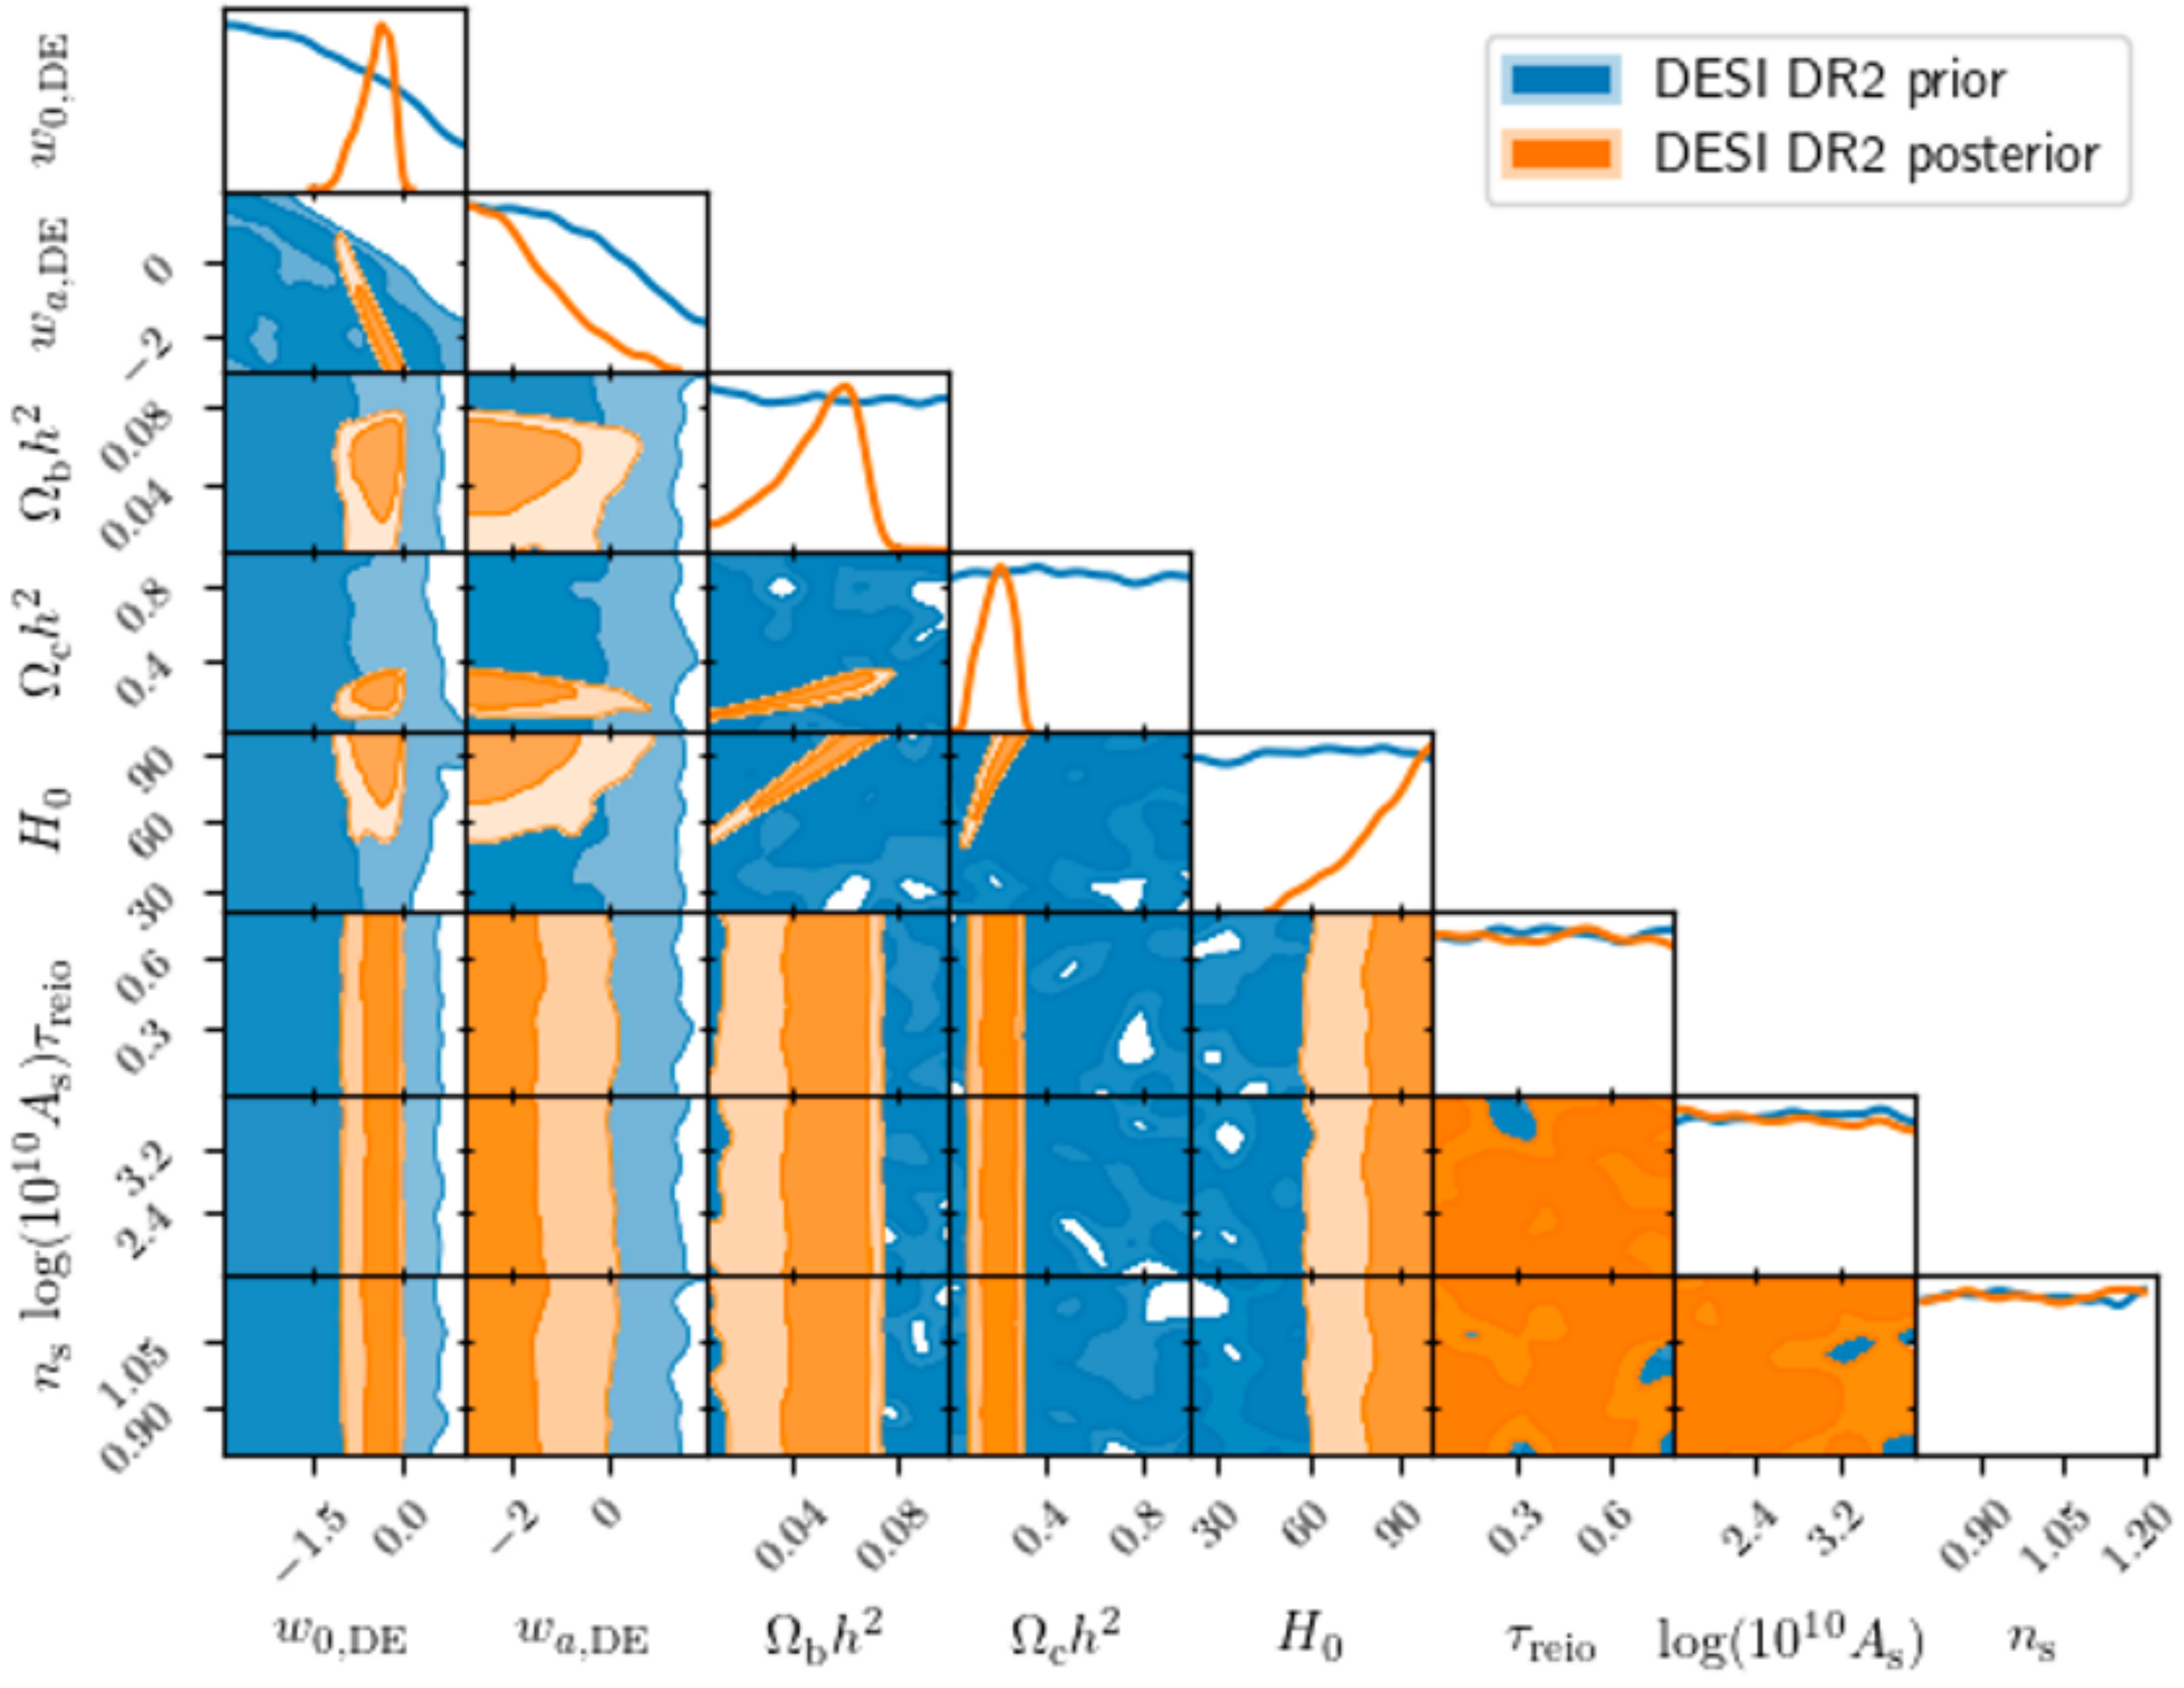

- Posteriors from nested sampling runs and MCMC runs
- Pre-computed results by HPC, accessible in your laptop in seconds, not weeks/months!
- Example:
  - DESI DR2 dataset using the w_0 w_aΛCDM cosmology model
  - Posteriors are clearly more constrained than the priors for several parameters.

unimpeded: Five Key Tension Statistics

- **R statistics** : $R = \frac{Z_{AB}}{Z_A Z_B} = \frac{P(A|B)}{P(A)} = \frac{P(B|A)}{P(B)}$ (measure confidence in combining datasets)
- **Information Ratio**: $\log I = D_A + D_B - D_{AB}$ (KL-divergence-based)
- **Suspiciousness**: $\log S = \log R - \log I$ (prior-independent tension)
- p**-value sigma**: Approximate "number of sigma" separation
- **Bayesian Model Dimensionality**: $\frac{d}{2} = \int P(\theta) \left(\log \frac{P(\theta)}{\pi(\theta)} - D \right)^2 d\theta$ (effective constrained parameter count)

unimpeded: A Global View of Tensions

- Tension statistics across 10 models (y-axis) and 29 pairwise datasets (x-axis)

- σ>2 in red

$$p = \int_{d-2\log S}^{\infty} \chi_d^2(x)\,\mathrm{d}x$$

$$\sigma = \sqrt{2}\mathrm{Erfc}^{-1}(p)$$

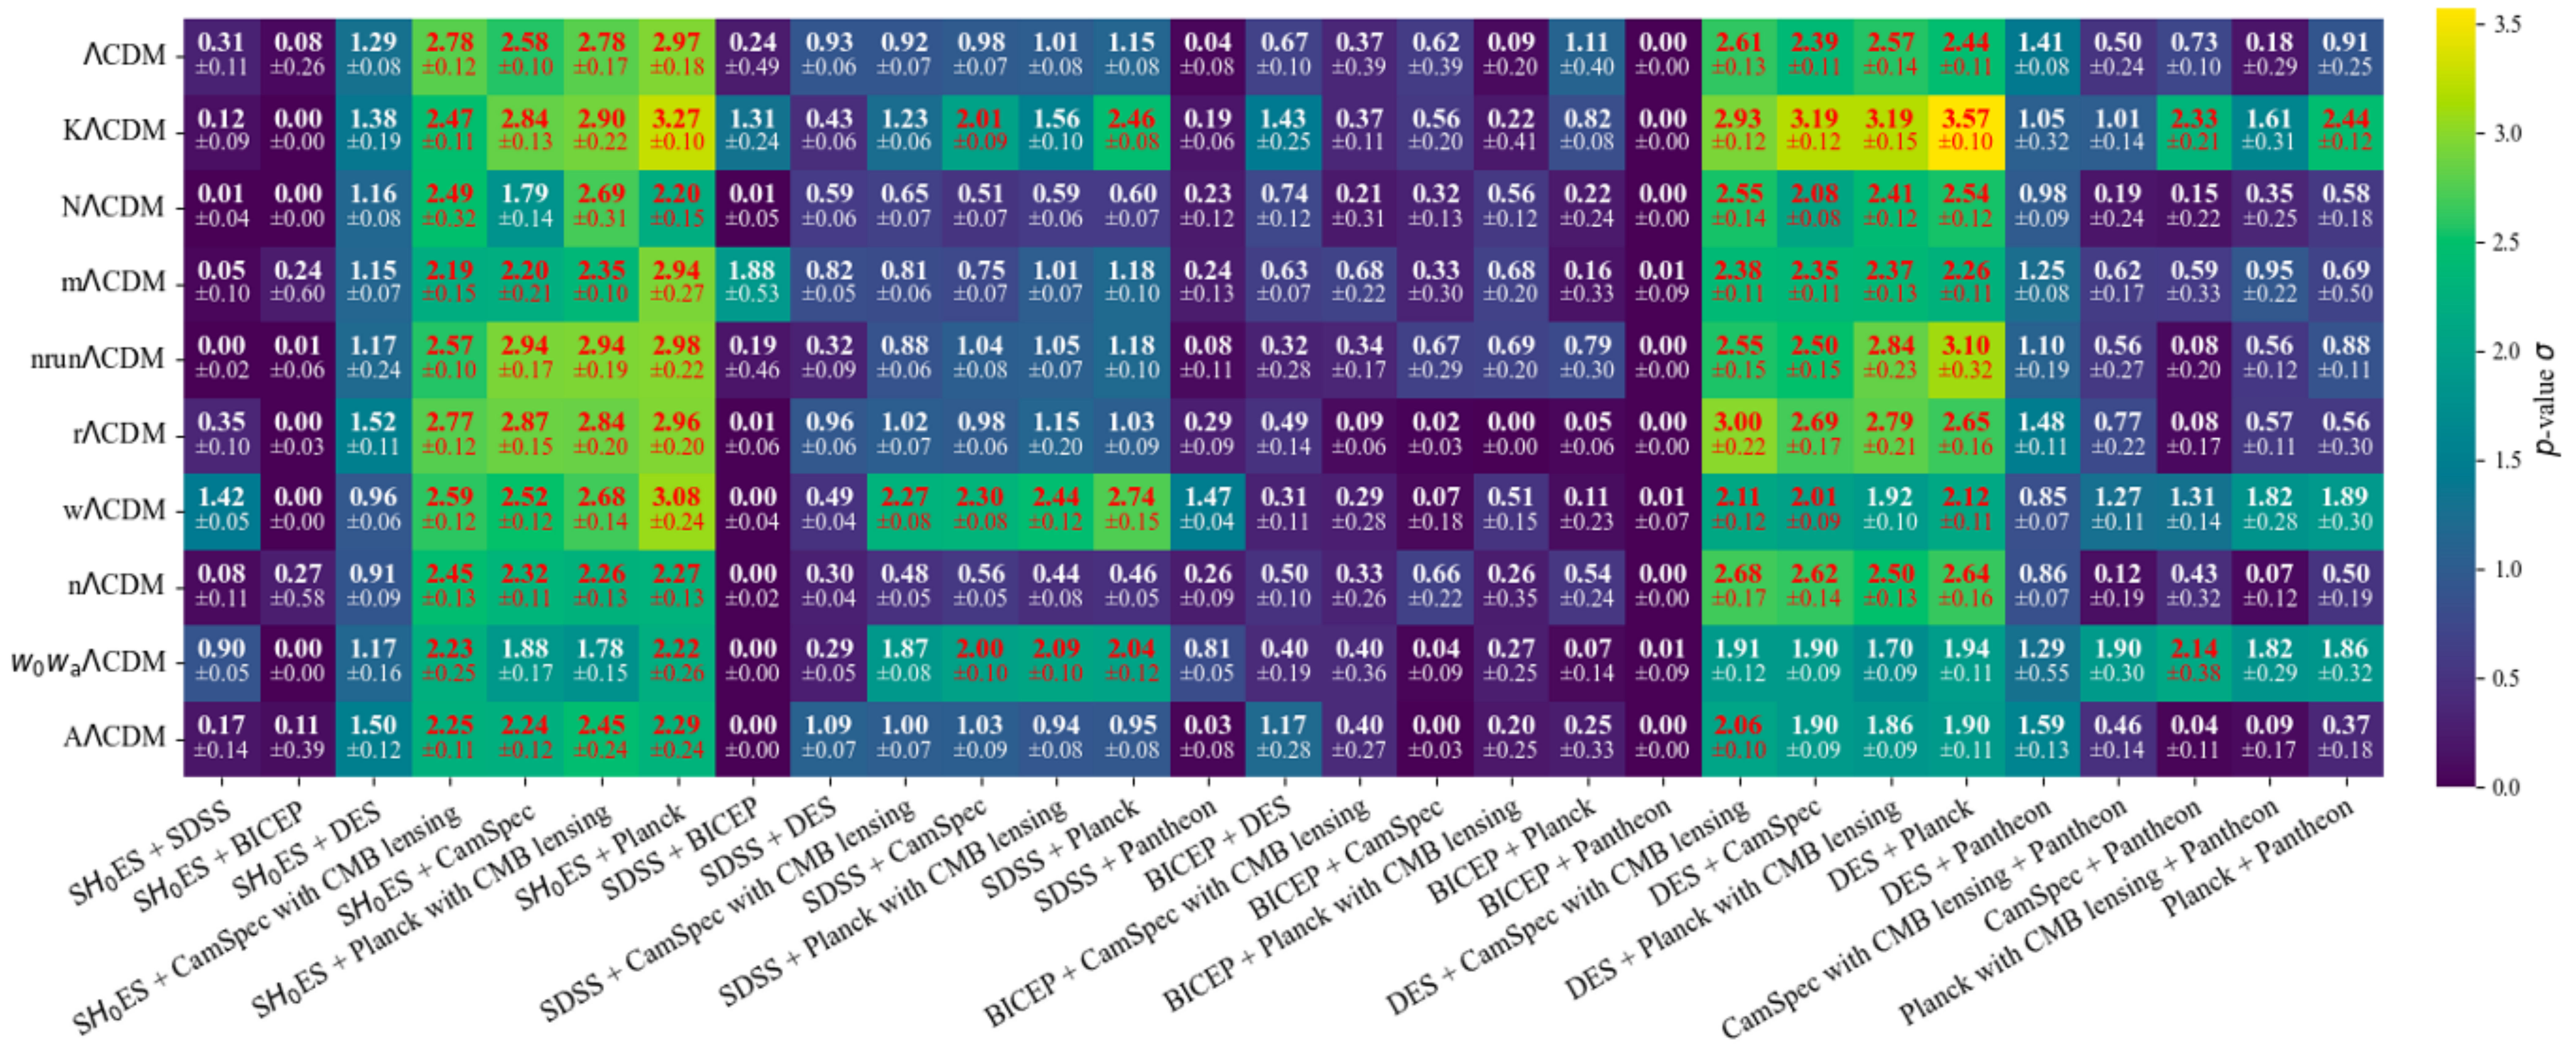

A systematic, quantitative search for tensions across a grid of models and data

unimpeded: Machine learning emulators

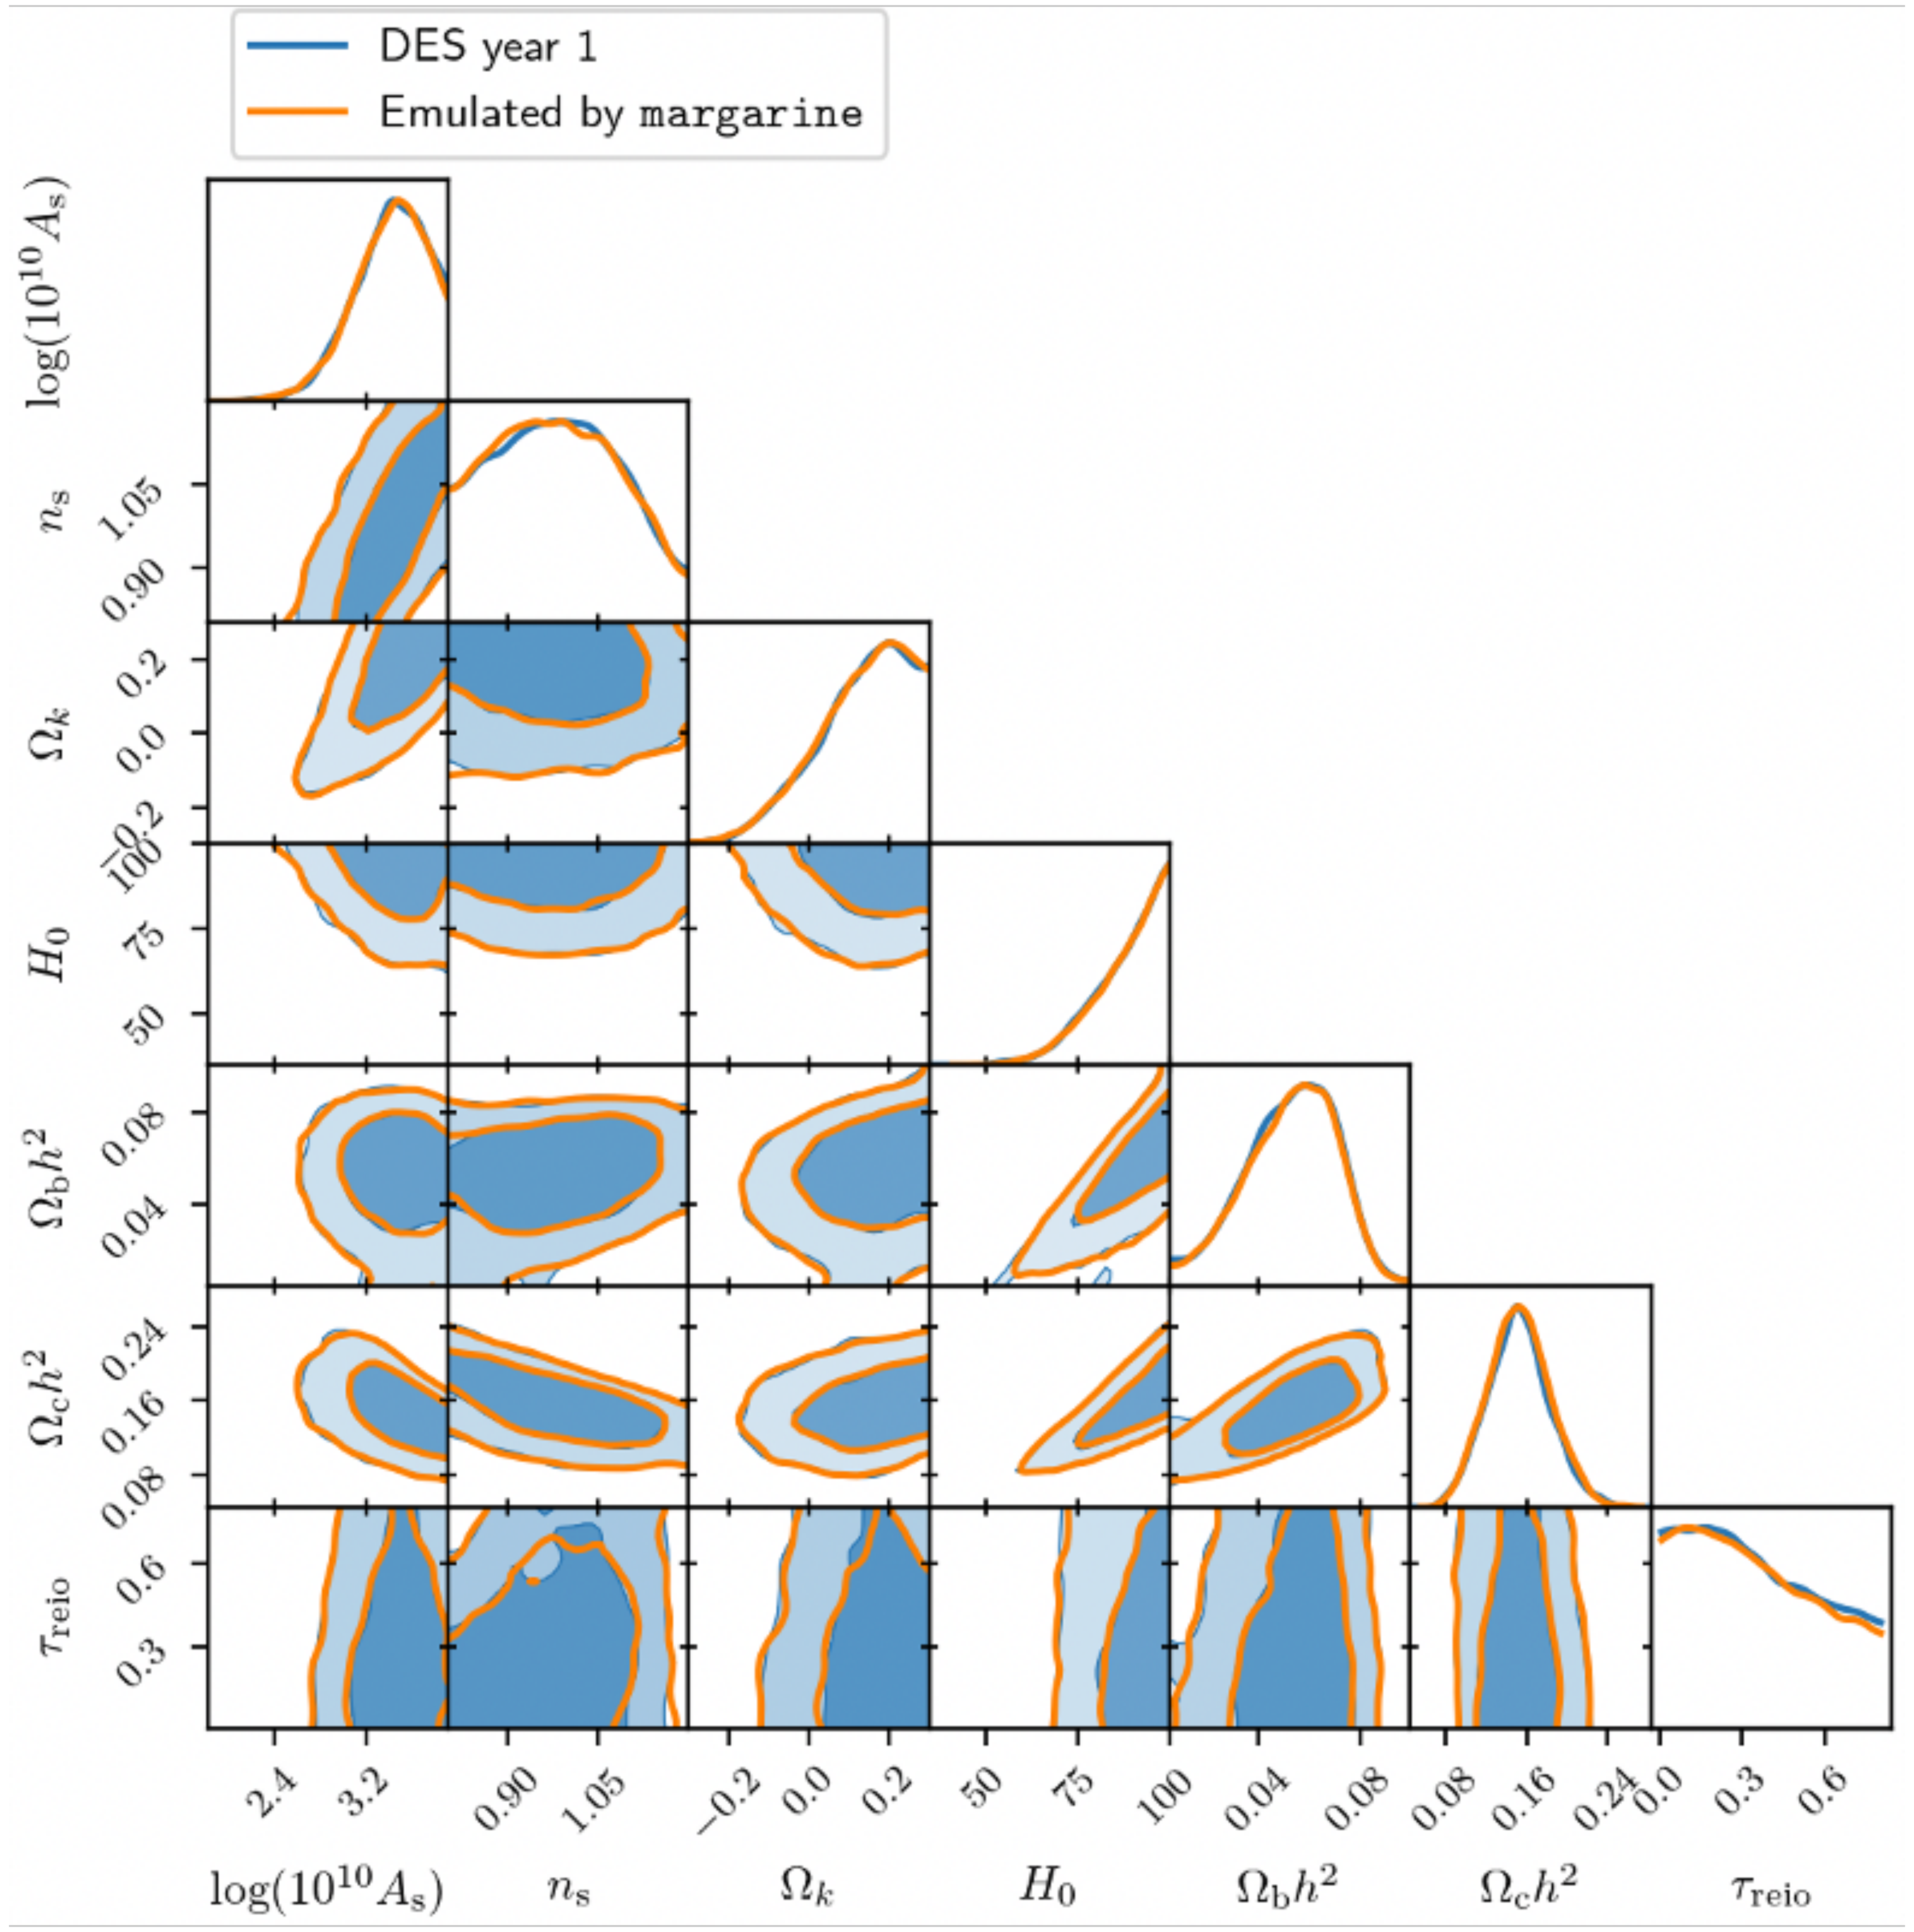

- Machine learning emulators
- Emulate marginalised likelihoods or posteriors.
- Benefits:
  - Dramatically speeds up inference for re-use in new analyses
  - Provides a fast and flexible alternative to full MCMC/NS runs
  - Provides a real ’planck prior’ rather than a Gaussian approximation

What’s next?

- Future of unimpeded:
  - Expanding the grid with DiRAC 17 to include next-generation datasets (DESI, Euclid, etc.).
  - Integrating more advanced ML and Simulation-Based Inference techniques.
- Join Us!
  - We are actively seeking α-testers and collaborators. If this tool could be useful for your research, please get in touch!
  - The code is open-source on GitHub.
- About Me:
  - I am a final-year PhD student and am currently exploring postdoctoral opportunities for next year. I'd be delighted to discuss my work further.

What's next?

- Future of unimpeded:
  - Expanding the grid with DiRAC 17 to include next-generation datasets (DESI, Euclid, etc.).
  - Integrating more advanced ML and Simulation-Based Inference techniques.
- Join Us!
  - We are actively seeking α-testers and collaborators. If this tool could be useful for your research, please get in touch!
  - The code is open-source on GitHub.
- About Me:
  - I am a final-year PhD student and am currently exploring postdoctoral opportunities for next year. I'd be delighted to discuss my work further.

Thank you!

Bonus Slide: Prior range

Parameter	Prior range	Baseline	Definition
$\omega_b \equiv \Omega_b h^2$	$[0.005, 0.1]$	...	Baryon density today
$\omega_c \equiv \Omega_c h^2$	$[0.001, 0.99]$	...	Cold dark matter density today
$100\theta_{MC}$	$[0.5, 10.0]$	...	$100 \times$ approximation to r_*/D_A (CosmoMC)
τ	$[0.01, 0.8]$	...	Thomson scattering optical depth due to reionization
Ω_K	$[-0.3, 0.3]$	0	Curvature parameter today with $\Omega_{tot} = 1 - \Omega_K$
$\sum m_\nu$	$[0, 5]$	0.06	The sum of neutrino masses in eV
$m^{\text{eff}}_{\nu,\,\text{sterile}}$	$[0, 3]$	0	Effective mass of sterile neutrino in eV
w_0	$[-3.0, -0.3]$	-1	Dark energy equation of state[a], $w(a) = w_0 + (1-a)w_a$
w_a	$[-2, 2]$	0	As above (perturbations modelled using PPF)
N_{eff}	$[0.05, 10.0]$	3.046	Effective number of neutrino-like relativistic degrees of freedom (see text)
Y_P	$[0.1, 0.5]$	BBN	Fraction of baryonic mass in helium
A_L	$[0, 10]$	1	Amplitude of the lensing power relative to the physical value
n_s	$[0.9, 1.1]$	...	Scalar spectrum power-law index ($k_0 = 0.05\text{Mpc}^{-1}$)
n_t	$n_t = -r_{0.05}/8$	Inflation	Tensor spectrum power-law index ($k_0 = 0.05\text{Mpc}^{-1}$)
$dn_s/d\ln k$	$[-1, 1]$	0	Running of the spectral index
$\ln(10^{10}A_s)$	$[2.7, 4.0]$	...	Log power of the primordial curvature perturbations ($k_0 = 0.05\,\text{Mpc}^{-1}$)
$r_{0.05}$	$[0, 2]$	0	Ratio of tensor primordial power to curvature power at $k_0 = 0.05\,\text{Mpc}^{-1}$

Planck 2013 results. XVI. Cosmological parameters

Bonus Slide: Zenodo

zenodo | unimpeded: | Communities | My dashboard

324 result(s) found

Versions

View all versions

Access status

- Open 321
- Restricted 3

Resource types

- Dataset 279
- Publication 40
- Image 2
- Poster 1
- Presentation 1
- Software 1

February 9, 2025 (v1) | Dataset | Open

unimpeded: mlcdm planck_2018_CamSpec

Ong, Dily

cosmological model:mlcdm, dataset:planck_2018_CamSpec

Uploaded on February 10, 2025

February 9, 2025 (v1) | Dataset | Open

unimpeded: wlcdm des_y1.joint+planck_2018_plik

Ong, Dily

cosmological model:wlcdm, dataset:des_y1.joint+planck_2018_plik

Uploaded on February 11, 2025

February 9, 2025 (v1) | Dataset | Open

unimpeded: rlcdm bao.sdss_dr16+sn.pantheon

Ong, Dily

cosmological model:rlcdm, dataset:bao.sdss_dr16+sn.pantheon

Uploaded on February 10, 2025

Files (316.5 MB)

Name	Size	Download all
mcmc_wlcdm_des_y1.joint+planck_2018_plik.csv md5:63243fd154170721340d6b126d3cb51d	134.2 MB	Preview Download
mcmc_wlcdm_des_y1.joint+planck_2018_plik.yaml md5:3a75b2be0d65788b98bc7ac674170508	17.4 kB	Download
ns_wlcdm_des_y1.joint+planck_2018_plik.csv md5:b33ed2091153ce302f4970905fa29218	182.2 MB	Preview Download
ns_wlcdm_des_y1.joint+planck_2018_plik.prior_info md5:85010fb0ba899aad91d5650e1abf1a7e	48 Bytes	Download
ns_wlcdm_des_y1.joint+planck_2018_plik.yaml md5:9d2e88284645351fed2338593c290040	17.8 kB	Download