



SCUOLA INTERNAZIONALE SUPERIORE DI STUDI AVANZATI

SISSA Digital Library

CIGaRS I: combined simulation-based inference from type Ia supernovae and host photometry

Original

CIGaRS I: combined simulation-based inference from type Ia supernovae and host photometry / Karchev, Konstantin; Trotta, Roberto; Jiménez, Raúl. - In: NATURE ASTRONOMY. - ISSN 2397-3366. - (2026).
[10.1038/s41550-026-02842-5]

Availability:

This version is available at: 20.500.11767/150692 since: 2026-05-07T13:12:01Z

Publisher:

Published

DOI:10.1038/s41550-026-02842-5

Terms of use:

Testo definito dall'ateneo relativo alle clausole di concessione d'uso

Publisher copyright

note finali coverpage

(Article begins on next page)

CIGaRS I: Combined simulation-based inference from SN α Ia and host photometry

Konstantin Karchev^{1,2*}, Roberto Trotta^{1,3,4,5} and Raul Jimenez^{2,6}

¹Theoretical and Scientific Data Science, SISSA, Trieste, Italy.

²Institute of Cosmos Sciences, University of Barcelona, Spain.

³Department of Physics, Imperial College London, UK.

⁴IFPU, Trieste, Italy.

⁵ ICSC, Casalecchio di Reno, Italy.

⁶ICREA, Barcelona, Spain.

*Corresponding author(s). E-mail(s): kkarchev@sissa.it;
Contributing authors: rtrotta@sissa.it; raul.jimenez@icc.ub.edu;

Abstract

Using type Ia supernovae (SN α Ia) as cosmological probes requires empirical corrections, which correlate with their host environment. We present a unified Bayesian hierarchical model designed to infer, from purely photometric observations, the intrinsic dependence of SN Ia brightness on progenitor properties (metallicity & age), the delay-time distribution (DTD) that governs their rate as a function of age, and cosmology, as well as the redshifts of all hosts. The model incorporates physics-based prescriptions for star formation and chemical evolution from Prospector- β , dust extinction of both galaxy and SN light, and observational selection effects. We show with simulations that intrinsic dependences on metallicity and age have distinct observational signatures, with metallicity mimicking the well-known step of SN Ia magnitudes across a host stellar mass of $\sim 10^{10} M_{\odot}$. We then demonstrate neural simulation-based inference of all model parameters from mock observations of $\approx 16\,000$ SN α Ia and their hosts up to redshift 0.9. Our joint physics-based approach delivers robust and precise photometric redshifts (< 0.01 median scatter) and improved cosmological constraints, unlocking the full power of photometric data and paving the way for an end-to-end simulation-based analysis pipeline in the LSST era.

Keywords: cosmology, dark matter and dark energy, galaxies and clusters, transient astrophysical phenomena

Introduction

It has long been known that the brightnesses, stretches, and colours of type Ia supernovæ (SNæ Ia) correlate with global properties of their host galaxies, e.g. morphology, stellar mass, star-formation rate (SFR), metallicity, and/or stellar age [1–12]. Similar dependences on the SNæ’s *local* environment have also been found [13–27]. In current cosmological analyses [e.g. 28–30], which use SNæ Ia as standardisable candles, it is common practice to (attempt to) correct for all these effects *en masse* via an *ad-hoc* “mass step”: a magnitude offset between SNæ Ia hosted in galaxies with stellar masses above and below a given threshold, usually set to the sample median or fixed around $10^{10} M_{\odot}$. The rationale is that stellar mass is a simple-to-estimate proxy for nearly all host properties. Unveiling the physical processes underlying host–SN Ia connections would thus not only enhance our understanding of SNæ Ia and their formation, but also deliver more accurate and precise constraints on dark energy by improving the standardisation procedure.

A supernova’s environment also has *extrinsic* effects on its light, which is dimmed and reddened by dust along the line of sight, both in the immediate surroundings of the progenitor system and in interstellar space within the host. It can also be scattered *into* the line of sight, significantly affecting finite-resolution photometry [31]. Accounting for dust requires careful hierarchical modelling [32] since a) extinction has a similar effect (redder–dimmer) to intrinsic colour-related standardisation (bluer–brighter) and b) there is a well-known relationship between the amount and properties of the dust in a galaxy and its SFR and stellar mass: the so-called “star-forming main sequence” [see e.g. 33]. Different studies have come to conflicting conclusions regarding this interplay: for example, Brout and Scolnic [34] attribute the mass step to differences in dust in low- and high-mass hosts rather than any intrinsic/physical/causal effect, while Thorp and Mandel [35] find a preference for a single dust law across stellar masses and a mass step of $\approx(0.06 \pm 0.03)$ mag, with more massive galaxies hosting intrinsically brighter SNæ Ia. However, Bayesian model comparison using the same data and models [36] mildly disfavours a non-zero mass step with posterior odds of 2:1.

Resolving these degeneracies and arriving at a definite conclusion regarding host–SN Ia connections requires tackling *all* interrelated effects simultaneously, including the apparent dimming due to cosmological distance and selection biases. While Bayesian hierarchical modelling (BHM) provides the principled framework for doing this, traditional Markov Chain Monte Carlo (MCMC) methods require sampling jointly the vast number of latent parameters (*tens per SN*) and explicit computation—at every step—of the selection probability, which is intractable in realistic scenarios. To date, analyses either resort to *ad-hoc* approximations [29, 37, 38] or are split in two separate steps [see e.g. 35, 39]: first, the SNæ’s distance moduli (i.e. intrinsic brightnesses) are estimated and corrected for selection effects using a fiducial, possibly mis-specified, model [40]; then, they are correlated with galaxy properties obtained in a separate analysis (rarely considering the associated uncertainties [41]) and/or used in a cosmological fit.

A two-step approach cannot account for many important statistical effects. For instance, the ages of SN Ia progenitors are not representative of their hosts’ stellar populations, and so regressing using the latter may be misleading. Moreover, standardisation implicitly gives higher weight to more massive hosts, where a larger number of SNæ occur—an instance of Eddington bias [42, see also 43]. Likewise, a preference for detecting (*apparently*) brighter objects biases the selected sample towards *intrinsically* brighter SNæ to a greater extent in dustier hosts, which are also more massive. Lastly, an *apparent* correlation—*non-existent intrinsically*—between SN brightnesses and properties of the hosts may be introduced by the use of photometric redshifts in the standardisation process due to the mass–age–redshift degeneracy. Similarly, an *apparent* redshift dependence of SN properties [44, 45] might arise due to the evolution of the host galaxies, e.g. their dust content [46, 47] or metallicity or age.

In this paper, we present the first framework for combined inference and galaxy-related standardisation (CIGaRS) of SNæ Ia and their hosts, which addresses all of the above conceptual and methodological issues and can *unequivocally* disentangle intrinsic and extrinsic effects through physics-based forward modelling. We formulate our model in the context of simulation-based inference (SBI) [for overviews, see 48, 49], a modern suite of Bayesian inference techniques that leverage the flexibility of neural networks (NNs) to obtain posterior distributions, given only training examples from a forward simulator. SBI is seeing rapid adoption across disciplines and has previously been applied in SN Ia studies to analyse dust distributions from collections of light curves [36, 50] and perform cosmological inference from future-sized data sets (up to 10^5 SNæ Ia) in the presence of photometric-like redshift uncertainties [51] and arbitrary selection effects [43].

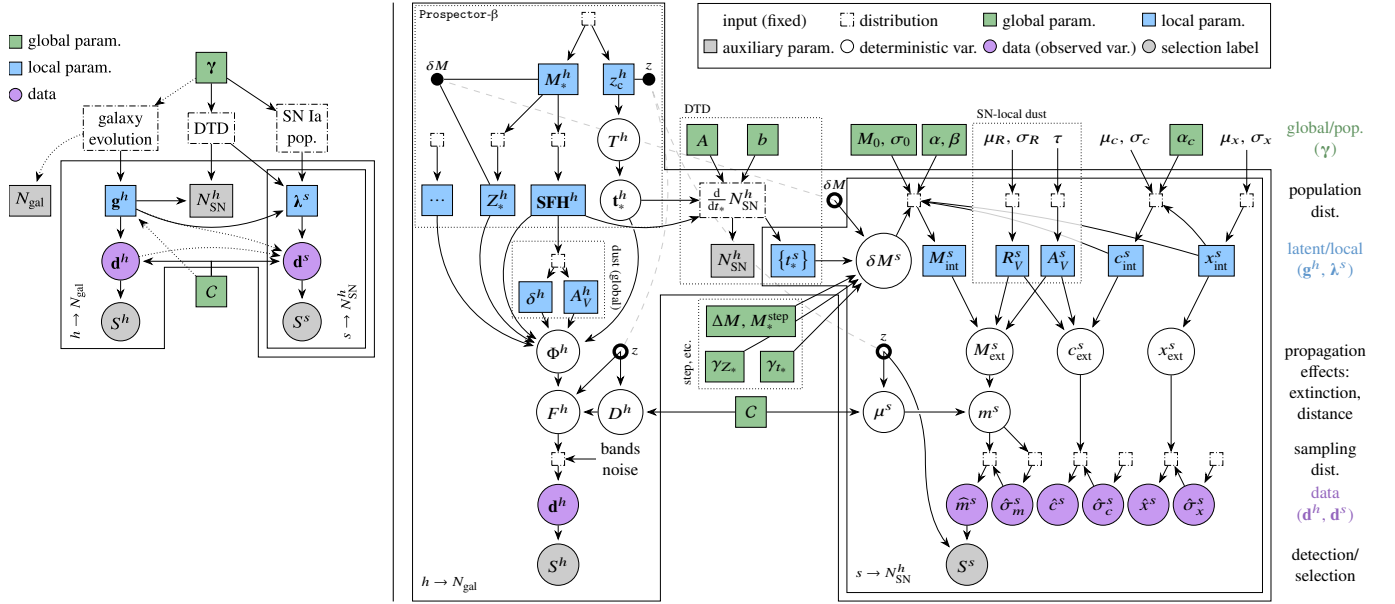


Fig. 1 *Left*: overview of unified Bayesian hierarchical modelling of SNe Ia and host galaxies. *Right*: our simulator combines Prospector- β [52, 53] and Simple-BayesSN [32] via a “bridge” formed by the DTD and intrinsic host-SN correlations to produce self-consistent SN Ia light-curve summaries and galaxy photometry (subject to sample selection). Symbol definitions and priors are listed in table 1, with full details given in [Methods: Unified forward modelling of galaxies and SNe Ia](#).

Our simulator (depicted graphically in figure 1 and fully elaborated in [Methods: Unified forward modelling of galaxies and SNe Ia](#)) adopts the prescription for galaxy evolution from Prospector- β [52, 53] and the Simple-BayesSN hierarchical model for SN Ia light-curve summaries [32]; it then applies dust extinction, cosmological distance, realistic measurement noise, and sample selection to both. We model and infer the *intrinsic* correlations between (true) host metallicity, (true) progenitor age, and (true intrinsic) SN absolute magnitude, which are latent variables in the hierarchical model; in addition, we allow for a residual mass step and infer its size (in magnitudes) and location (host mass). Lastly, we include a self-consistent formulation of the occurrence of SNe Ia within galaxies through the delay-time distribution (DTD): the rate at which SNe Ia arise from the host’s stellar populations based on their ages, which can be calculated from first principles for plausible SN Ia progenitor scenarios [54]. Our framework seamlessly integrates inference of the DTD under a suitable parametrisation, thus offering a tool to shed light on the much-debated formation mechanism of SNe Ia [55, 56].

While insights into the DTD and the host–SN Ia connections are an important output of our analysis, we envisage its primary future application to be cosmological inference from the Legacy Survey of Space and Time (LSST), soon to commence at the Vera Rubin Observatory. Due to the large spatial coverage and depth (18 000 deg² reaching redshift ≈ 1) of its primary wide-fast-deep (WFD) survey component, only $\approx 1\%$ of the $\sim 10^6$ SN candidates expected over 10 years will have spectroscopic typing and redshift, and even with dedicated campaigns like 4MOST/TiDES [57], only up to 10% of the host galaxies will be followed up.

In view of this predominantly photometric scenario, we do not consider any spectroscopic observables or explicit redshift estimates, providing the inference network only with broadband photometry of the hosts and light-curve summaries for the SNe Ia (we do assume perfect transient classification, noting that non-Ia contamination can be straightforwardly included in our framework). Owing to the physical modelling of galaxy evolution and the pooling of strength introduced by the joint analysis, our framework delivers state-of-the-art photometric redshift estimates *as a by-product*, as well as considerably more stringent cosmological constraints than using SN data alone (even with spectroscopic redshifts). We explore more thoroughly the implications of these findings for SN cosmology with future data sets in a dedicated companion paper, focusing here on host–SN Ia connections and the DTD.

Results

To examine the phenomenology and demonstrate inference with our unified model, we generate a mock data set, \mathbf{D}_0 , with global parameters as listed in table 1 and object-specific (latent) quantities sampled from their hierarchical priors. We scale the counts to be representative of the current flagship Dark Energy Survey (DES), which contains 1635 photometrically classified SNe [58]. In our particular \mathbf{D}_0 , the number of selected objects is $N_{\text{sel},0} = 1583$.

Predictions from the unified forward model

Physical host–SN Ia correlations and induced stellar-mass dependence

Our forward model includes explicit dependence of SN Ia magnitude offsets (δM^s) and the true progenitor metallicity (Z_*^s) and age (t_*^s). We plot those (separately) in figure 2 (with correlation parameters γ_{Z_*} , γ_{t_*} chosen consistently with observations [8, 12] and indicated in the respective panels) against stellar mass of the hosts. Due to well-known trends in galaxy formation incorporated in Prospector- β (more massive galaxies are, in general, older and more metal-rich), the SN Ia brightness *appears* to depend on stellar mass. Crucially, whereas a metallicity correlation induces a (smoothed) step across $M_* \approx 10^{10}$, the age correlation manifests as an approximately linear trend with $\log M_*$. Moreover, the distributions of magnitude offsets at a given stellar mass, $p(\delta M | M_*)$, also differ: while metallicity shows a Gaussian spread [59], progenitor ages are highly skewed towards younger populations due to the fast decaying DTD (we adopt a power law with slope $b = -1.34$ [60]).

Due to these differences, in principle, the two scenarios can be distinguished using stellar-mass estimates alone, which are easy to obtain: see our results in [Object-specific parameters](#). However, as we show next, this task is complicated by the presence of colour/stretch-related variations and residual scatter (≈ 0.3 mag in total), as well as uncertainties both in the dependent (δM^s) and independent variables (M_*^s and/or Z_*^s , t_*^s). Finally, note that splitting at the median mass — as customary — results in a diminished step (dashed lines) since both subsets are “pulled” towards the peak of the M_* distribution ($\approx 10^{11} M_\odot$): an instance of Eddington bias [42, see also 43].

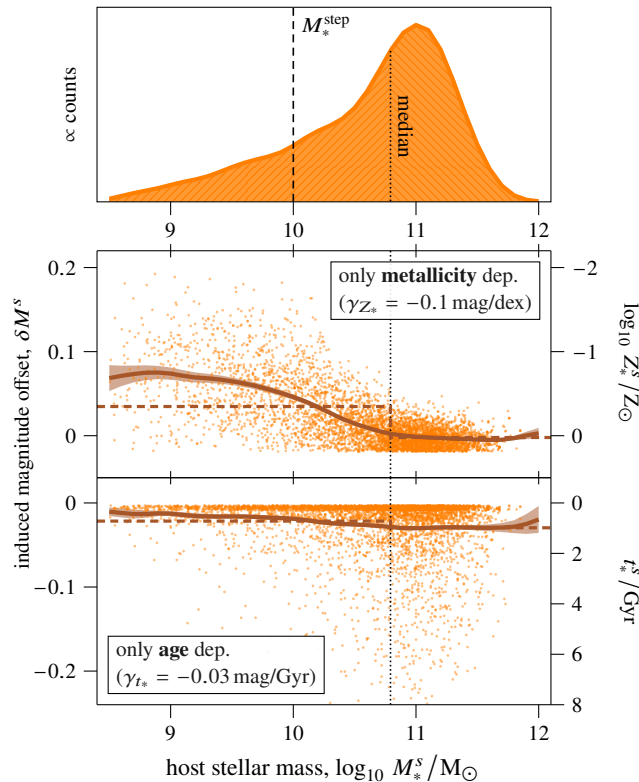


Fig. 2 *Top*: Distribution of the stellar masses of SN Ia hosts in our mock, i.e. the galaxy stellar mass function weighted by the SN Ia rate within each galaxy (using DTD slope $b = -1.34$). *Middle and bottom*: Respectively, true metallicities and ages of the progenitor stellar populations (labels on the right) and the induced SNe Ia magnitude offsets δM (labels on the left) under the default correlation parameters (γ_{Z_*} & γ_{t_*}), as indicated. Solid lines are a smooth moving average (with uncertainty due to sampling noise), whereas dashed lines depict the apparent step across the median host mass. *Note*: The plots show the full population before detection/selection and across all redshifts. We also indicate the location (M_*^{step}) of the *additional* step ($\Delta M = 0.05$) simulated in our mock data.

Redshift distribution and Hubble diagram

A key outcome of our unified approach is the predicted redshift distribution of SN $\bar{\text{a}}$ Ia, which results from the combination of galaxy evolution (according to `Prospector- β`) and a DTD-based occurrence model. At low redshifts (up to $z_c \approx 0.3$), it is qualitatively consistent with a power-law volumetric rate with exponent 1 to 1.5 as inferred by Dilday et al. [61] from SDSS and widely used in the literature (e.g. in PLAsTiCC [62]). Above $z_c \approx 0.4$, the rate starts decreasing, driven by the decline in the galaxy population. Finally, at $z_c \gtrsim 0.6$, corresponding to a universe of age $\lesssim 8$ Gyr, the frequency of SN Ia occurrence within hosts decreases, as potential hosts are smaller (since there is less time for star formation) and younger (thus reducing the range of the DTD). This effect is compounded with the drop in detection efficiency and results—for our selection function, which mimics LSST’s WFD survey (see [Detection/selection of SN \$\bar{\text{a}}\$ Ia](#))—in no SN $\bar{\text{a}}$ Ia being detected above $z_c \approx 0.8$.

In the bottom panel of figure 3, we show the different SN Ia absolute magnitude offsets included in our forward model: these would be the residuals from a perfectly specified cosmological model and standard magnitude M_0 in a traditional analysis. The dominant source of scatter are standardisable correlations with stretch and colour (≈ 0.26 mag). At high redshift, the sample exhibits typical Malmquist bias [63, 64] towards brighter SN $\bar{\text{a}}$ (downward trend) with a reduction of *observed* scatter due to the removal of dim objects. The effect of dust is also systematic: it dims (i.e. shifts upwards) the extrinsic magnitudes (M_{ext}^s , red shaded band) with respect to M_{int}^s (green shaded band) by ≈ 0.13 mag. Finally, we show the intrinsic host-related magnitude offsets (δM^s , purple shaded band), which include metallicity *and* age correlations as visualised in figure 2 *and* a residual mass step of 0.05 mag across $10^{10} M_\odot$. These effects lead to an offset due to the dominance of high-mass (and hence, higher-metallicity and older) galaxies, but their sizes are small in comparison with the above (on the order of the irreducible random scatter: $\sigma_0 = 0.1$ mag), making the host—SN Ia connections difficult to infer.

Traditional analyses need to incorporate all of these effects in an intricate likelihood that attempts to disentangle intrinsic and extrinsic influences, scatter from measurement uncertainty, trends from selection biases to derive accurate constraints on the cosmological parameters, host—SN Ia connections, and the DTD. In our SBI framework, on the other hand, it is sufficient to include them in the forward model, and the neural network learns from training data how to solve the full inverse problem.




Fig. 3 *Top*: Redshift distributions of galaxies (according to `Prospector- β`), SN $\bar{\text{a}}$ Ia within them (as predicted by our DTD-based model), and detected SN $\bar{\text{a}}$ Ia (using our LSST-inspired selection function). While the latter two are on the same scale, the galaxy counts have been re-normalised for the presentation. *Bottom*: Offsets of the extrinsic (dust-affected) and intrinsic absolute magnitudes (M_{ext}^s and M_{int}^s) from the standard M_0 and the host/progenitor-related δM^s as functions of redshift. Points represent M_{ext}^s from our mock data (non-detected SN $\bar{\text{a}}$ are also shown in a lighter shade). The shaded areas indicate the mean \pm standard deviation—for detected SN $\bar{\text{a}}$ —of the respective quantities binned by z_c , while the dashed lines delineate the irreducible/residual scatter with st. dev. σ_0 intrinsic in the SN Ia population.

Simultaneous simulation-based inference

In this section, we demonstrate the ability of CIGaRS to obtain — simultaneously — marginal posteriors for all parameters in the hierarchical model, given only SN Ia light-curve summaries (stretch and apparent brightness and colour at peak, together with the associated fit uncertainties) and LSST-like *ugrizy* photometry of the hosts. We use truncated marginal neural ratio estimation (TMNRE) [65, 66] to train a conditioned deep set++ network [67, 68, see also 43] (depicted in figure 4) and apply it to the simulated data set presented above. With minimal fine-tuning, we then scale to a tenfold larger mock catalogue of approximately 16 000 SNæ Ia and their hosts (expected to be discovered in a month’s worth of WFD observations). Details of the network architecture and the iterative training procedure are presented in [Methods: Set-based TMNRE for hierarchical models](#).

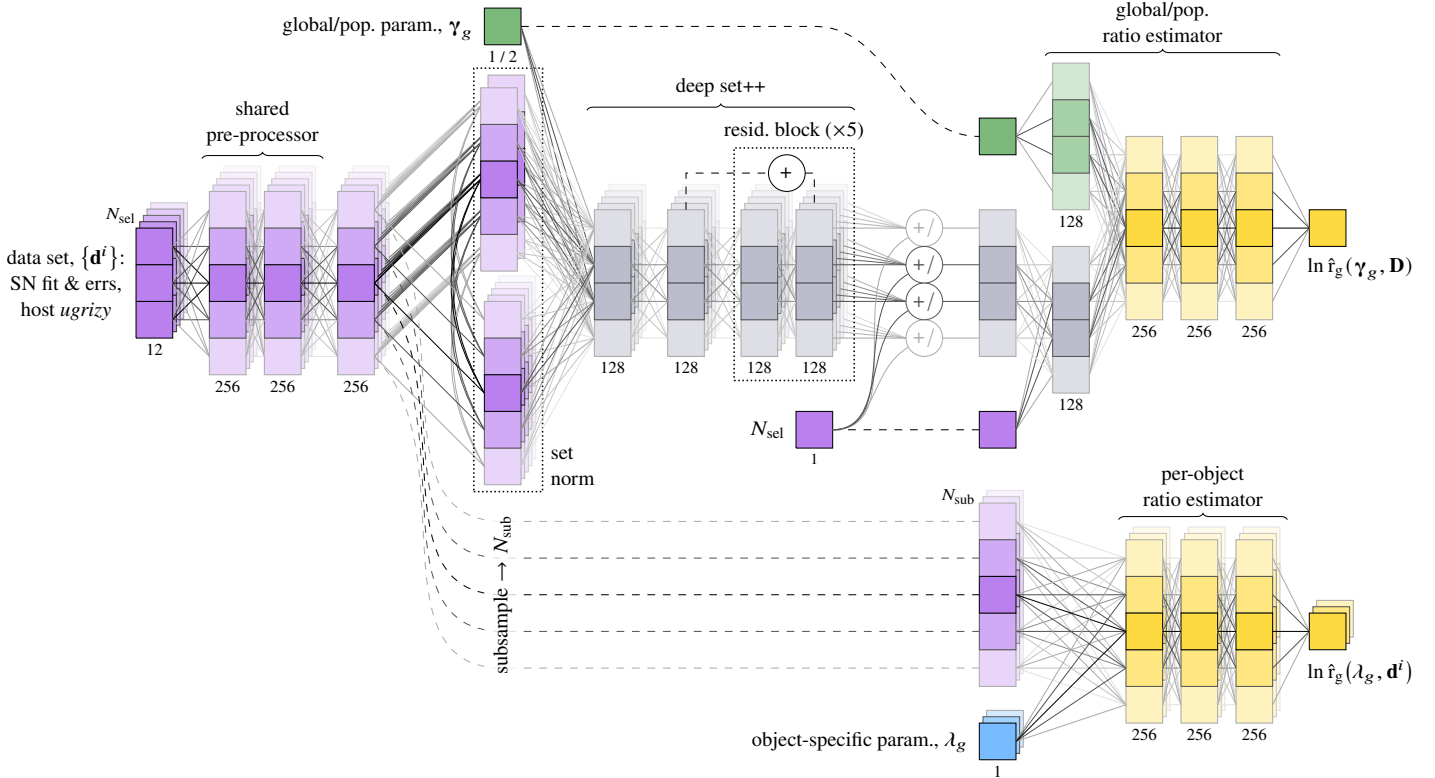


Fig. 4 Our inference network combines a conditioned deep set++ (utilising residual connections and set normalisation) [67, 68, see also 43] with standard feed-forward ratio estimators to infer simultaneously the global/population parameters in our forward model and all object-specific properties from a catalogue containing SN Ia summary statistics (with uncertainties) and photometry of their hosts.

Global and population parameters

Figure 5 shows the final posteriors for all global/population parameters. All groups are significantly constrained with respect to their initial priors, and the results are fully consistent with the values used to simulate the mock data.

We constrain the power-law slope of the DTD, b , to within ± 0.05 , comparable with the uncertainty in the DES analysis [69], which considers 809 objects but has access to spectroscopic redshifts. We also recover the parameters of all three types of intrinsic host–SN Ia dependences we consider: correlations with the progenitor metallicity and/or age and the size and location of an additional mass step. Importantly, we can disentangle a metallicity from an age dependence even when both are present (as evident from the uncorrelated posterior in the $\gamma_{Z_s} - \gamma_{t_s}$ plane). In contrast, a metallicity dependence is harder to distinguish from a mass step (hence the strong *a-posteriori* correlation in the $\Delta M - \gamma_{Z_s}$ plane), as could be anticipated from our exploration of the model’s predictions (figure 2).

In the same figure, we also show posteriors obtained from a 10 times larger mock data set, corresponding to a modest fraction of LSST observations. We verify that the precision of our fully principled and self-consistent analysis scales with the square root of the survey size (constraints are ≈ 3 times narrower, given ten times more objects). This will allow robust conclusions to be drawn about the progenitors of SNæ Ia and their environmental dependencies, even from purely photometric data, which has never been achieved before.

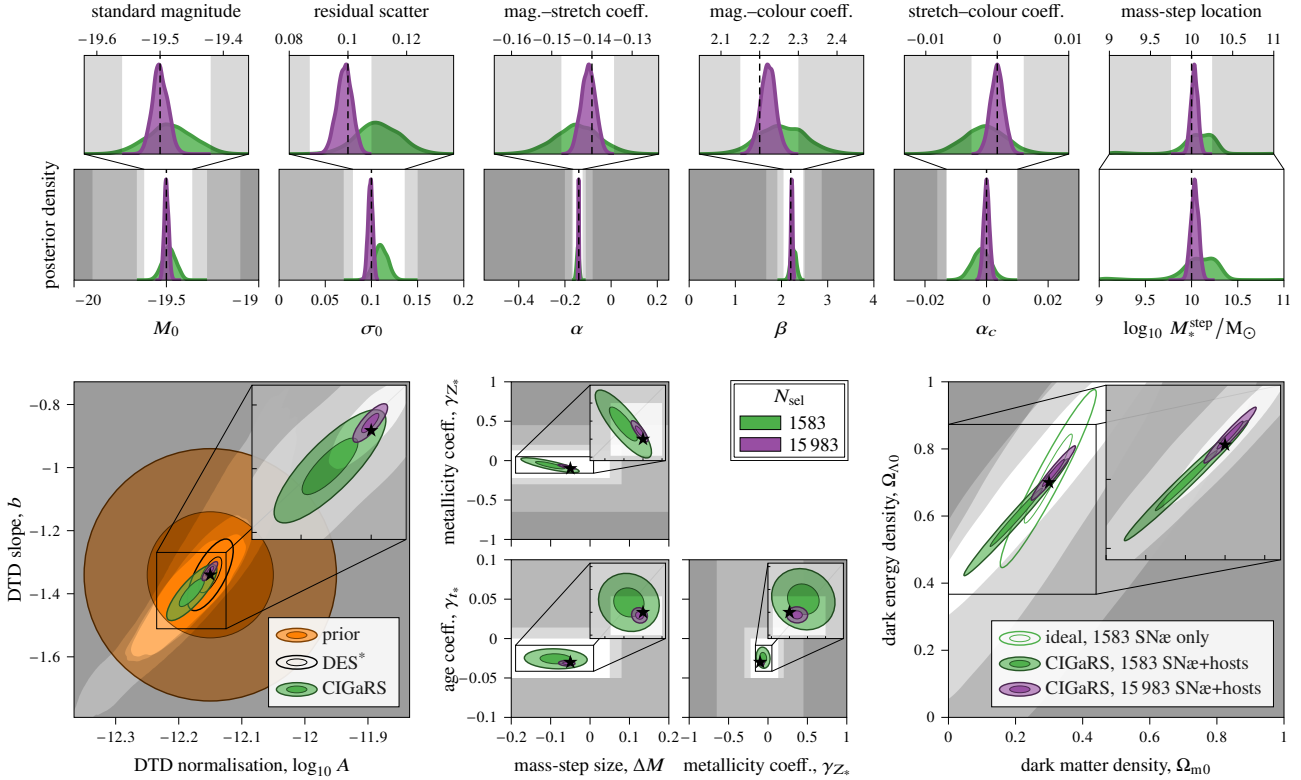


Fig. 5 Inference results—*top*: marginal posterior densities; *bottom*: 1- and 2-sigma contours (with credibility 39% and 86%, respectively)—obtained from the two mock data sets we consider, with true values indicated by dashed lines and stars. For all parameters, the prior is uniform over the plotted range, except for the DTD, where it is an uncorrelated normal distribution (orange). Gray-shaded areas indicate successive stages of prior truncation. For comparison, we show the DTD constraints (shifted to our true value) from the Dark Energy Survey (DES) [69], using 809 objects with spectroscopic redshifts, and an *idealised* cosmological analysis (given the ground-truth values for *all* object-specific latent parameters, which are not available in reality) of *only* the SN Ia observables from our mock data (excluding hosts).

Finally, we present marginal posteriors for the cosmological parameters (dark matter and dark energy densities, Ω_{m0} and $\Omega_{\Lambda 0}$) derived directly from light-curve summaries and host photometry. We compare our contours to an *idealised* analysis of *only* the SN \ddot{a} (hosts excluded, as in all current analyses) in our mock data set, in the best-case—and completely unrealistic—scenario that all latent parameters (redshifts, stretches, intrinsic colours, dust, and host-related magnitude offsets) are exactly known. Even though this leaves only the irreducible magnitude scatter of $\sigma_0 = 0.1$ mag and measurement noise as sources of uncertainty, the resulting cosmological constraints are significantly *weaker* than those delivered by CIGaRS. In a cosmology-centric companion work, we test the conjecture that this is due to the network having access to additional information in (i.e. being able to standardise also) the *host brightnesses*.

Object-specific parameters

Marginal posteriors for parameters of all 1583 objects—derived simultaneously with the global results—are presented in figure 6. True values are well recovered across the full ranges. For host stellar mass and metallicity, our uncertainties match those from previous applications of Prospector [52, 53]. When the “signal” is intrinsically low, e.g. for low metallicities, our results revert to the prior, ensuring proper Bayesian uncertainty propagation. Moreover, photometric redshifts are inferred with exceptionally high precision: a median posterior standard deviation ≈ 0.01 , *no* outliers with $|\Delta z| > 0.1$, and no signs of a systematic bias, owing to a) the Prospector- β priors informed by galaxy evolution, which break the age–mass–redshift degeneracy [see 53], and b) the extraction of redshift information from SN (and host) brightnesses, given the globally inferred cosmological model. This is on par with analyses of UV–optical–IR data with much wider coverage [70] and represents a major improvement over the expected scatter ≈ 0.04 from LSST photometry alone [71]. Despite using only summarised SN observables, our results already surpass the state of the art with full light curves and host photometry: see Chen et al. [72], who achieve photo- z scatter of ≈ 0.02 for DES.

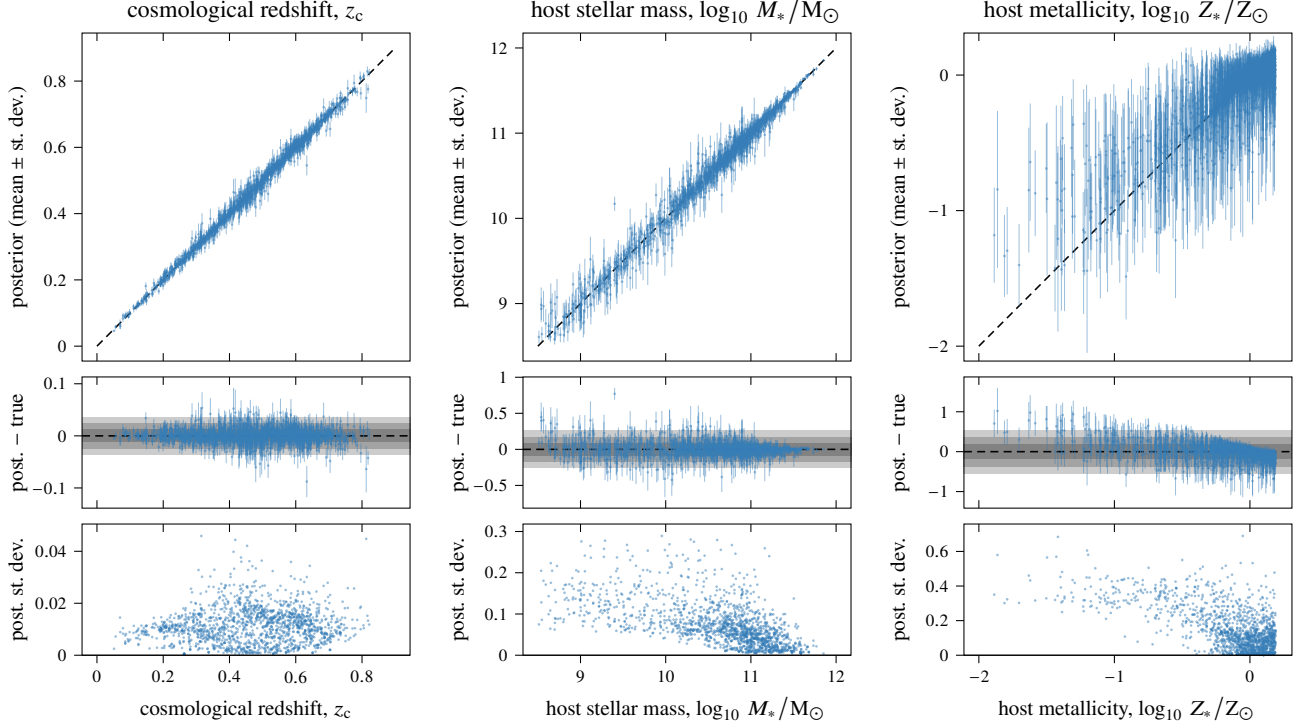


Fig. 6 Inference results for host-related parameters of the 1583 objects in our mock data, obtained simultaneously with the global constraints. Everywhere, the horizontal axis is the true parameter value in the simulation. *Top*: All 1583 marginal posteriors (mean \pm standard deviation). *Middle*: Residuals from true values; the shaded bars represent the distribution (1, 2, & 3 standard deviations) of the posterior means. *Bottom*: The posterior uncertainties (standard deviations). *Note*: We only show z_c , M_* , Z_* since they are relevant to the current study, but infer δ , τ_2 , x_{int} , c_{int} as well since this helps train the shared pre-processor: see [Methods: Set-based TMNRE for hierarchical models](#).

Conclusion and outlook

We have presented CIGaRS: the first unified and fully self-consistent framework for rigorous Bayesian analysis of SN Ia summary statistics and photometry of their hosts. It is based on a forward simulator that combines two hierarchical models: *Prospector- β* and *Simple-BayesSN*, via a prescription for the occurrence of SN α Ia (according to a delay-time distribution (DTD) convolved with non-trivial star-formation histories) and physical connections between their brightnesses and progenitor metallicity and/or age. CIGaRS qualitatively reproduces the observed SN Ia rate with redshift and the distribution of host stellar masses and makes detailed predictions about the apparent stellar-mass dependence induced by the underlying physical relations: a metallicity correlation resembles a smoothed step function across $10^{10} M_\odot$ with symmetric Gaussian scatter, whereas an age correlation leads to a linear trend with stellar mass and highly skewed residuals.

We have demonstrated, for the first time, simultaneous neural simulation-based inference of the DTD parameters, brightness–metallicity/age correlations, a residual mass step, and cosmology purely from LSST-like host photometry and SN Ia summaries (brightness, stretch, and colour) of up to $\approx 16\,000$ objects. Our approach implicitly handles redshift estimation, host-related standardisation, and selection effects, while propagating all uncertainties, and thus remains accurate, precise, and scalable. Moreover, due to the unified hierarchical modelling and combination of SN and host data, we achieve robust photometric redshift estimates with state-of-the-art scatter and cosmological constraints that improve upon spectroscopic SN-only analyses.

We envisage extensions of the model to include spatial localisation of the stellar populations (and hence, of the SN α) and of interstellar dust, peculiar velocities of the hosts and of the SN α within them (which will allow inference of cosmological parameters controlling large-scale structure), the presence of non-Ia transients, their detection and characterisation directly from a time-series of telescope images, and the incorporation of spectroscopic data, where available. Remarkably, our inference framework can be applied to an arbitrarily complicated simulator and collections of observables without significant modifications, while retaining its scalability, and we expect it to be instrumental in maximising the scientific insight extracted from upcoming vast and detailed LSST observations.

Methods: Unified forward modelling of galaxies and SNaE Ia

CiGARS is the first unified forward model for SN Ia cosmology. It is a diptych (see figure 1) of two state-of-the-art Bayesian hierarchical simulators for photometric galaxy observations (\mathbf{d}^h) and SN Ia summary statistics (\mathbf{d}^s), joined on three levels:

1. by the probability of occurrence (equivalently, number N_{SN}^h) of SNaE Ia in a given host h , based on the latter’s intrinsic properties \mathbf{g}^h ;
2. by the distribution of intrinsic properties λ^s of the SNaE, which also depend on the $\mathbf{g}^{h(s)}$ of their respective hosts;
3. by extrinsic effects of the host environment (e.g. dust) and other properties (e.g. cosmological redshift and peculiar velocity) on the SN light as it travels towards the observer.

The former two, which affect λ , can be considered *causal* connections, whereas the latter, which only affects the observables and data \mathbf{d} , are *coincidental*. One last connection can arise during data reduction (e.g. subtraction of the host light); we plan to implement it in a future extension that implements the image analysis procedure of transient surveys. In what follows, we describe in detail all components of CiGARS (see also the full hierarchy in figure 1 and the list of parameters in table 1), discussing also the planned improvements, mainly related to intra-host localisation of the SNaE.

Host model

Comprehensive treatment of the photo–spectral evolution of galaxies — see the pioneering studies using the MOPED method [73–75] and e.g. Mo et al. [76] for a review — is beyond the scope of this work. At a high level, the intrinsic properties of a galaxy h , such as its star-formation history (SFH), its metallicity (Z_*^h), and interstellar dust content, are largely determined by the galaxy’s total stellar mass (M_*^h) and its cosmological redshift (z_C^h) — a proxy (given a cosmological model) for the age of the Universe (T^h) at the time the galaxy’s light is emitted. To represent the intrinsic correlations inherent in galaxy formation, we adopt an off-the-shelf Bayesian framework for galaxy photometry: *Prospector* [77], as updated in *Prospector- β* [53].

Star-formation history and chemical enrichment: *Prospector- β*

Prospector encodes a non-trivial SFH by discretising it in seven time bins (found to be the optimal number for describing moderate signal-to-noise observations [78]) spanning $[0; T^h]$: the first two fixed to $[0; 30 \text{ Myr}]$ and $[30 \text{ Myr}; 100 \text{ Myr}]$, the next four logarithmically spaced within $[100 \text{ Myr}; 0.85 T^h]$, and the last covering $[0.85 T^h; T^h]$. The galaxy’s total mass is then distributed into seven stellar populations that represent the stars born during each history bin:

$$\mathbf{SFH}^h \equiv [\text{SFH}^{h,j}]_{j=1}^7, \quad \text{with} \quad \sum_{i=1}^7 \text{SFH}^{h,j} = M_*^h. \quad (1)$$

Assuming that all stars in a given population were born at approximately the same time ($t_*^{h,j}$, taken at the (linear) centre of the respective bin), the process of stellar population synthesis (SPS) produces the integrated light (before extinction) from all stars in a galaxy by convolving the emission spectra of single stellar populations with the SFH.

SPS also requires a prescription for the galaxy’s chemical enrichment history: the metallicities of all its stellar populations. *Prospector* assumes a single metallicity within a galaxy (Z_*^h), determining it (with appropriate scatter) from the total stellar mass as in Gallazzi et al. [59]. This has also been shown to be appropriate for modelling the progenitors of SNaE Ia [79], which is our ultimate goal.

Host dust extinction

The light emitted by stars is then extinguished (and reddened) by dust within the host. *Prospector* considers separate birth-cloud and diffuse dust components. The former only affects the youngest stellar population (i.e. acting only on the emission from the first time bin) and has a fixed wavelength dependence $\propto \lambda^{-1}$ and optical depth τ_1^h . Extinction due to diffuse (interstellar) dust is described by the **KC13** law [80, following 81, 82] with shape parameter δ^h and optical depth τ_2^h related to the column density of dust along the line of sight (since it represents the total extinction in the (rest-frame) V band, we will label it A_V^h instead).

In *Prospector- β* , the dust parameters are *a priori* independent from other galaxy properties; however, many studies have found *a posteriori* relationships from large galaxy surveys. We enforce the empirical relations of

Alsing et al. [83, (14) and (15)]:

$$\tau_2^h \equiv A_V^h \sim \mathcal{N}\left(0.2 + 0.5 \text{ReLU}\left(\log_{10} \text{SFR}^h / \text{M}_\odot/\text{yr}\right), 0.2^2\right), \quad (2)$$

$$\delta^h \sim \mathcal{N}\left(-0.095 + 0.111 A_V^h - 0.0066 (A_V^h)^2, 0.4^2\right), \quad (3)$$

where the star-formation rate (SFR) is derived from the most recent (current) component of \mathbf{SFH}^h : $\text{SFR}^h \equiv \text{SFH}^{h,1} / 30 \text{ Myr}$. Prospector then sets the optical depth of birth-cloud dust according to $\tau_1^h / \tau_2^h \sim \mathcal{N}(1, 0.3^2)$.

It is important to note that the above description of dust extinction is global, i.e. averaged over the spatial light distribution of the galaxy, since host photometry is usually integrated. It may thus not represent faithfully the *local* extinction law or the optical density of dust at *any* particular location within the galaxy, e.g. in the proximity of a given SN Ia. The spatial distribution of dust properties is a topic of active research [e.g. 84–87], with mounting evidence for its effect on SN Ia standardisation [88, 89]. We will address this in a future refinement of the model, in coordination with the dust effects applied to SN light.

Photometry and detection of hosts

Given the above quantities (and parameters that describe the presence and emission of interstellar gas and an active galactic nucleus (AGN), which we randomly sample from the Prospector- β priors since they are not directly relevant to SN Ia studies), we compute the rest-frame spectral energy distribution (SED)—more accurately, spectral flux— $\Phi^h(\lambda_r)$ using the NN-based SPS emulator Speculator- α [90]. The speed-up this brings over traditional (i.e., from first principles) SPS is crucial for simulating the significant amount of training data we need.

We then redshift Φ^h (in this study, we assume no peculiar velocities ($z_c^h = z^h = z^s$) since the fraction of objects in our simulations and mock data with $z_c \lesssim 0.02$, where the distinction would be important, is negligible) and apply cosmological (transverse comoving) distance $D^h \equiv D(z_c^h, C)$, where C are the cosmological parameters, to convert it to the observer-frame spectral flux density (SFD):

$$F^h(\lambda_o) \equiv \frac{\Phi^h(\lambda_o / (1 + z^h))}{(1 + z^h)^3 4\pi (D^h)^2}. \quad (4)$$

In principle, at this stage we need to account for extinction due to dust in the Milky Way (MW), but given that it is a well-understood and modelled phenomenon (both in terms of wavelength dependence and optical depth), we ignore it, assuming the final observations have been corrected for it as is routine. Instead, we directly integrate (numerically) F^h within the Rubin *ugrizy* passbands, convert to AB magnitudes, and apply 0.01 mag Gaussian noise to arrive at the simulated galaxy data \mathbf{d}^h —a vector of length 6.

Finally, to simulate detection, i.e. the indicator variable S^h , we apply a simple magnitude cut, stipulating that *all* 6 measured magnitudes must be brighter than the respective 5σ detection thresholds expected for LSST 10 yr co-added imaging [91]. Including arbitrarily complex selection criteria to obtain a fully realistic selection function is straightforward: it is sufficient to apply them to the simulations as if they were real data.

SN Ia occurrence

The first connection between galaxies and SN $\ddot{\text{a}}$ Ia (or any other transient) is the latter’s emergence from the former’s stellar populations, which we describe through the delay-time distribution (DTD): the rate (per unit rest-frame time) of occurrence of SN $\ddot{\text{a}}$ Ia from a given stellar population (per unit stellar mass within it), as a function of the population’s age t_* . Therefore, within a survey of (observer-frame) duration \mathcal{T} , a number

$$\langle N_{\text{SN}}^{h,j} \rangle = \frac{\mathcal{T}}{1 + z^h} \times \text{SFH}^{h,j} \times \text{DTD}(t_*^{h,j}) \quad (5)$$

of SN $\ddot{\text{a}}$ Ia arise (*on expectation*) from galaxy h ’s stellar population j (recall that we assumed all stars within it have the same age $t_*^{h,j}$). We adopt a power-law DTD (set to zero in the first 100 Myr, i.e. for the two youngest stellar populations, to allow for the creation of a white dwarf):

$$\text{DTD}(t) = A \times (t/\text{Gyr})^b \times \text{M}_\odot^{-1} \text{yr}^{-1} \quad \text{for } t > 0.1 \text{ Gyr}, \quad (6)$$

and independent priors representing observational constraints [60, 92]:

$$\log_{10} A \sim \mathcal{N}(-12.15, 0.1^2), \quad (7)$$

$$b \sim \mathcal{N}(-1.34, 0.2^2). \quad (8)$$

To realise this prescription, in each forward pass of the model, we iterate (in parallel on a GPU) over all stellar populations of all galaxies, i.e. over all h, j , and generate

$$N_{\text{SN}}^{h,j} \sim \text{Poisson}\left(\left\langle N_{\text{SN}}^{h,j} \right\rangle\right) \quad (9)$$

supernovæ, sampling their ages, $\{t_*^{h,j,k}\}_{k=1}^{N_{\text{SN}}^{h,j}}$, uniformly within the time bin associated with population h, j . To simplify notation, we will label uniquely each SN with $s \in \{1, \dots, \sum_{h,j} N_{\text{SN}}^{h,j}\}$ and keep track of the host $h(s)$ of each, so that we can associate the relevant intrinsic galaxy properties: $z_c^s \equiv z_c^{h(s)}$, $Z_*^s \equiv Z_*^{h(s)}$, $M_*^s \equiv M_*^{h(s)}$, and host observations $\mathbf{d}^{h(s)}$. We note that this assumes perfect host association, which is often not the case in real observations, and that we ignore the additional information contained in instances of multiple SNæ arising in *the same* host [see e.g. 93]; we plan to improve in these respects by extending the simulator towards detection and characterisation of the transients directly from telescope images.

Finally, it is also possible to keep track—in the simulator—of the stellar population $j(s)$ from which each SN has arisen and associate corresponding properties if these differ from one stellar population to another within the same host (e.g. to represent a fully realistic enrichment history through different $Z_*^{h,j}$). Similarly, we can simulate the spatial distribution of stellar populations, giving rise to metallicity/age gradients [e.g. 94] and differences in dust extinction [e.g. 84–87]. We defer an expansion of the simulator *within* the hosts for future work.

SN Ia model

Once we have determined how many SNæ Ia have occurred during a survey, we proceed to simulate their observables. This follows a similar path from intrinsic properties derived from a population model *informed by the host* of each object, through extinction, redshift, and distance, to noisy measurements and detection / sample selection.

Causal host–SN Ia connection(s)

The cornerstone of CIGaRS is an explicit parametrised dependence of the intrinsic SN Ia properties on the characteristics of their host and/or of their progenitor stellar population: in general, $p(\lambda^s | \mathbf{g}^{h(s),j(s)}, \gamma)$. Our choice of which links to include is motivated by observational evidence [4, 6–12] and takes the form of a host-dependent magnitude offset δM^s , but we could similarly introduce e.g. a parametrised host-dependent distribution of x_{int}^s [26].

Since the physical source of δM^s is not yet established, we allow for multiple possibilities (in isolation or combination)—namely, metallicity Z_*^s and progenitor age t_*^s —with correlation coefficients, respectively γ_{Z_*} and γ_{t_*} , whose priors include 0 and wide ranges of positive and negative values. In keeping with current practice and as a “catch-all” parameter, we also allow for an *additional/residual* “mass step” ΔM with a similar wide prior and a free (inferred) stellar-mass location M_*^{step} . The full intrinsic host connection is, therefore:

$$\delta M^s = \underbrace{\gamma_{Z_*} \times \log_{10} Z_*^s / Z_{\odot}}_{\text{metallicity}} + \underbrace{\gamma_{t_*} \times t_*^s}_{\text{age}} + \underbrace{\Delta M \times \mathbb{I}(M_*^s > M_*^{\text{step}})}_{\text{ad-hoc mass step}}, \quad (10)$$

where \mathbb{I} returns 1 or 0 if its argument is true or false, respectively, and $\gamma_{Z_*}, \gamma_{t_*}, \Delta M, M_*^{\text{step}} \in \gamma$ are global parameters. Inference of the presence and strength of the respective effects is performed by examining their posterior(s) and/or performing Bayesian model selection, which is most conveniently and rigorously achieved through simulation-based methods [36]. Moreover, with SBI, all underlying correlations are accounted for and marginalised over when inferring cosmology, ensuring robustness of dark-energy inference against systematics arising from the host–SN Ia dependences.

Intrinsic SN Ia properties and dust extinction: Simple-BayeSN

Bayesian hierarchical modelling replaces empirical standardisation with *a-priori* correlated *latent parameters* sampled from hierarchical priors [32, 95]:

$$\text{stretch: } x_{\text{int}}^s \sim \mathcal{N}(\mu_x, \sigma_x^2), \quad (11)$$

$$\text{colour: } c_{\text{int}}^s \sim \mathcal{N}(\mu_c + \alpha_c x_{\text{int}}^s, \sigma_c^2), \quad (12)$$

$$\text{absolute } B\text{-band magnitude: } M_{\text{int}}^s \sim \mathcal{N}(M_0 + \alpha x_{\text{int}}^s + \beta c_{\text{int}}^s + \delta M^s, \sigma_0^2), \quad (13)$$

where $\mu_x, \sigma_x, \mu_c, \sigma_c, M_0, \sigma_0, \alpha, \alpha_c, \beta \in \gamma$ are population parameters assigned fixed hyper-priors as listed in table 1.

Like starlight, SNæ are affected by the dust surrounding them and further along the line of sight, i.e. in intergalactic space and in the Milky Way. We will ignore the latter two, assuming they have been perfectly corrected for, and adopt the Bayesian formulation of host-dust extinction from Simple-BayeSN [32], which acts on the intrinsic parameters introduced in (11) to (13) to obtain their “extrinsic” versions:

$$x_{\text{ext}}^s = x_{\text{int}}^s, \quad (14)$$

$$c_{\text{ext}}^s = c_{\text{int}}^s + E_{B-V}^s = c_{\text{int}}^s + A_V^s / R_V^s, \quad (15)$$

$$M_{\text{ext}}^s = M_{\text{int}}^s + A_B^s = M_{\text{int}}^s + A_V^s (R_V^s + 1) / R_V^s, \quad (16)$$

where $E_{B-V} \equiv A_B - A_V$ is the selective extinction, with A_B the total extinction in the rest-frame B band (in which the peak SN Ia magnitude is typically standardised), and $R_V \equiv A_V / E_{B-V}$. Note that the colour–magnitude standardisation coefficient β in (13) refers only to the c_{int} rather than the dust-affected c_{ext} .

Ideally, in a unified model, one would coordinate the extinction applied to SNæ Ia with the dust properties of the host. However, due to the complicated distribution of dust *within* galaxies [e.g. 84–87] and its non-trivial attenuation effects, which include not only extinction but also scattering of galactic *and* SN light into the line of sight [31], the simple approach of adopting the host-global A_V^h and $R_V(\delta^h)$ predicts an order of magnitude larger optical depths than empirically observed in SN data. We will explore the effect of dust localisation in a further study and here adopt instead the host-*independent* dust populations from Simple-BayeSN:

$$R_V^s \sim \mathcal{N}(\mu_R, \sigma_R^2), \quad (17)$$

$$A_V^s \sim \text{Exponential}(1/\tau), \quad (18)$$

with $\mu_R, \sigma_R, \tau \in \gamma$ global parameters.

Cosmological distance

The extrinsic absolute (rest-frame B -band) magnitude M_{ext}^s is then transformed into an apparent (still rest-frame B -band) magnitude m^s through the usual distance-modulus relation:

$$m^s = M_{\text{ext}}^s + \mu^s \quad \text{with} \quad \mu^s = \mu(z_c^s, C), \quad (19)$$

where z_c^s is the SN’s cosmological redshift, and C are the cosmological parameters. In this study, we use Λ CDM, which is described by the present-day dimensionless densities of cold dark matter (CDM) and dark energy (DE) in the form of a cosmological constant: $C \equiv [\Omega_{\text{m}0}, \Omega_{\Lambda 0}]$. As when modelling galaxies, we disregard peculiar velocities, including those of the SN within its host: these are negligible at high redshifts and best accounted for at the level of light curves. We note the recent application of SN Ia observations to inference of large scale structure (LSS) parameters [96, 97], which can be seamlessly and effortlessly incorporated into CIGaRS.

SN Ia observables

SN observations take the form of multi-band light curves (LCs): collections of flux measurements in different filters at irregularly spaced times, which vary from SN to SN. However, cosmological analyses are typically performed after the LCs have been *summarised* (independently of one another) by subtracting the host’s constant contribution and fitting a LC model, which produces parameter *estimates* $\hat{x}(\mathbf{LC})$, $\hat{c}(\mathbf{LC})$, $\hat{m}(\mathbf{LC})$ and fit uncertainties $\hat{\sigma}_x(\mathbf{LC})$, $\hat{\sigma}_c(\mathbf{LC})$, $\hat{\sigma}_m(\mathbf{LC})$. It is then assumed that these *summary statistics* are related to the extrinsic SN *parameters* x_{ext}^s , c_{ext}^s , m^s by Gaussian sampling distributions:

$$\hat{x}^s \sim \mathcal{N}\left(x_{\text{ext}}^s, (\hat{\sigma}_x^s)^2\right), \quad (20)$$

$$\hat{c}^s \sim \mathcal{N}\left(c_{\text{ext}}^s, (\hat{\sigma}_c^s)^2\right), \quad (21)$$

$$\hat{m}^s \sim \mathcal{N}\left(m^s, (\hat{\sigma}_m^s)^2\right). \quad (22)$$

Since the uncertainties only depend on instrumental properties like noise and cadence, their distributions can be robustly determined and treated as fixed simulator inputs. We will use the model of Boyd et al. [98], which is based on the results of current surveys and LSST simulations:

$$\ln \hat{\sigma}_x^s \sim \mathcal{N}\left(-1.5, 0.5^2\right), \quad (23)$$

$$\ln \hat{\sigma}_c^s \sim \mathcal{N}\left(-3.5, 0.3^2\right), \quad (24)$$

$$\ln \hat{\sigma}_m^s \sim \mathcal{N}\left(0.1(m^s - 56), 0.6^2\right). \quad (25)$$

It is possible to extend this description to a dense covariance matrix that represents the systematic correlations between SN $\bar{\alpha}$, as demonstrated in [SICRET](#). However, with hierarchical modelling and SBI, these can instead be accounted for explicitly in the simulator, replacing the likelihood-centric covariance formulation. Issues related to summary statistics can also be circumvented altogether by extending the simulation and analysis to full light curves, as demonstrated in [SIDE-real](#).

Detection/selection of SN $\bar{\alpha}$ Ia

The detection and inclusion of SN $\bar{\alpha}$ in modern cosmological analyses is a complex process that considers the raw data quality, the goodness of light-curve fits, and the estimated summaries. Using host observations (e.g. for deriving redshifts and/or studying host–SN Ia connections) relies on a further procedure for identifying and associating them correctly [see e.g. 99]. A major advantage of our framework is that arbitrary detection/selection/association (and classification) criteria can be straightforwardly integrated into the forward model and accounted for with simulation-based inference, as recently shown in [STAR NRE](#).

Since the present study does not specifically focus on selection effects, we adopt the same simple SN detection/selection procedure as in [STAR NRE](#) [43, section 3.2], deriving from expected LSST observing conditions a probability $p(S^s | \hat{m}^s, z^s)$ that depends on the SN’s observed brightness and its redshift (due to the different detection limits in the different (observer-frame) bands in which the peak occurs). We treat as “detected/selected” all host–SN Ia pairs (“objects”) where both S^h and S^s are sampled true. We label their count N_{sel} and treat it as an observable as in [STAR NRE](#). Finally, the output of the simulator is a *set* of length-12 vectors that combine the host- and SN-related observables for each of the detected/selected objects:

$$\mathbf{D} \equiv \left\{ \left(\mathbf{d}^{h(i)}, \hat{m}^i, (\hat{\sigma}_m^i)^2, \hat{x}^i, (\hat{\sigma}_x^i)^2, \hat{c}^i, (\hat{\sigma}_c^i)^2 \right) \right\}_{i=1}^{N_{\text{sel}}}. \quad (26)$$

Implementation details and simulated counts

Here, we discuss two technical details of our forward simulator.

Firstly, we note that `Prospector` relies on the cosmological model to map redshift to age: $T(z_c, C)$. In principle, this calculation can be repeated for each sampled cosmology in our simulator; however, we do not expect this to have a noticeable effect on our results since SNæ Ia are mainly informative of cosmological distances rather than times; i.e., for values of C consistent with a given SN Ia sample, $T(z_c, C)$ does not vary significantly with C . Therefore, we treat $T(z_c)$ as fixed to that consistent with the WMAP cosmology [100], as when training/creating `Prospector-β`.

The second point concerns the number of objects that we simulate, which, in principle, depends on (and hence is informative of) the astrophysical and cosmological models (e.g. through a parameter S_8), as well as the surveyed sky area (Ω). Currently, our high-level SN-focused simulator does not include these details since SN observations do not carry information about this connection, beyond the distribution of their redshifts, which we assume is extracted mainly from the hosts.

Hence, the purely host-related part of CIGaRS has no free global parameters up to the point where the *redshifted* (i.e. observer-frame) spectral flux

$$\Phi_o^h(\lambda_o) \equiv \frac{\Phi^h(\lambda_o/(1+z^h))}{(1+z^h)^3} \quad (27)$$

is converted to SFD through cosmological distance as in (4). This means that we can generate a *fixed* “bank” of galaxies and associated observables—noiseless *absolute ugrizy* magnitudes obtained by integrating (27) through the respective bandpass filters—to serve as potential hosts for simulating SNæ. We choose a bank size of 1 000 000: note that this represents the *total* galaxy population (N_{gal}) rather than a detected/selected sample.

When compiling a mock survey, given C , we calculate the apparent brightnesses of the galaxies, add noise, and evaluate detection as in [Photometry and detection of hosts](#). We then seed *only* the detected/selected subset with SNæ Ia as in [SN Ia occurrence](#). This is an implicit selection criterion *on the SNæ*: we only consider SNæ for which we can observe and uniquely identify the host (with certainty); for our (current) method to be applicable, we need to apply the same cut to the real data, but this is usually not a stronger requirement than the cuts typically applied on the SN data. For the default cosmological model and DTD (see table 1), only about 39% of the galaxies pass selection, and within them the SN Ia rate is approximately 2500/yr (before selection).

Lastly, we need to set a survey duration (\mathcal{T} , see (5)), but this is only well defined in combination with Ω or a physically meaningful N_{gal} , which we replaced with the fixed size of the galaxy “bank”. Therefore, we set an arbitrary scaling in (5) so as to achieve (on expectation) a given number of detected/selected objects: e.g. 1600 or 16 000. Importantly, since Ω and \mathcal{T} are well known for real surveys, we can easily extend the simulator to properly calculate the galaxy and SN Ia counts when analysing real data; and in the present setup, we are allowed to set the same scaling when generating training data as for the test mocks.

Table 1 Parameters in the model, their (hierarchical) priors, and the values used to simulate the analysed mock data. The given ranges indicate the support of the respective priors, within which object-specific parameters are drawn.

	parameter	symbol	value / range	unit	prior	
C	DM density ($z_c = 0$)	Ω_{m0}	0.3	—	$\mathcal{U}(0; 1)$	
	DE density ($z_c = 0$)	$\Omega_{\Lambda 0}$	0.7	—	$\mathcal{U}(0; 1)$	
galaxy model	cosmological redshift	z_c^h	[0.02; 1.52]	—	Prospector- β	
	stellar mass (total)	$\log_{10} M_*^h / M_{\odot}$	[8.5; 12.5]	dex		
	stellar metallicity	$\log_{10} Z_*^h / Z_{\odot}$	[-1.98; 0.19]	dex		
	dust optical depth	A_V^h	[0; 4]	mag		(2)
	KC13 dust law	δ^h	[-1; 0.4]	—		(3)
DTD	normalisation	$\log_{10} A$	-12.15	dex	$\mathcal{N}(-12.15, 0.1^2)$	
	slope	b	-1.34	—	$\mathcal{N}(-1.34, 0.2^2)$	
“causal” host–SN Ia connection	“mass step” size	ΔM	-0.05	mag	$\mathcal{U}(-0.2; 0.2)$	
	“mass step” location	$\log_{10} M_*^{\text{step}} / M_{\odot}$	10	dex	$\mathcal{U}(9; 11)$	
	magnitude–metallicity	γZ_c	-0.1	mag dex $^{-1}$	$\mathcal{U}(-2; 2)$	
	magnitude–age	γt_c	-0.03	mag Gyr $^{-1}$	$\mathcal{U}(-0.1; 0.1)$	
SN Ia pop.	“causal” mag. offset	δM^s	$(-\infty; \infty)$	mag	(10)	
	abs. magnitude (B -band)	M_{int}^s	$(-\infty; \infty)$	mag	$\mathcal{N}(M_0 + \alpha x_{\text{int}}^s + \beta c_{\text{int}}^s + \delta M^s, \sigma_0^2)$	
	standard abs. mag.	M_0	-19.5	mag	$\mathcal{U}(-20; -19)$	
	residual mag. scatter	σ_0	0.1	mag	$\mathcal{U}(0.01; 0.2)$	
	“stretch”	x_{int}^s	$(-\infty; \infty)$	—	$\mathcal{N}(\mu_x, \sigma_x^2)$	
	pop. mean	μ_x	0	—	fixed	
	pop. st. dev.	σ_x	1	—	fixed	
	“colour”	c_{int}^s	$(-\infty; \infty)$	mag	$\mathcal{N}(\mu_c + \alpha_c x_{\text{int}}^s, \sigma_c^2)$	
	pop. mean	μ_c	0	mag	fixed	
	pop. st. dev.	σ_c	0.1	mag	fixed	
intrinsic correlation coefficients	magnitude–stretch	α	-0.14	mag	$\mathcal{U}(-0.55; 0.25)$	
	magnitude–colour	β	2.2	—	$\mathcal{U}(0; 4)$	
	colour–stretch	α_c	0	mag	$\mathcal{U}(-0.03; 0.03)$	
extrinsic effects	extinction (V -band)	A_V^s	[0; $+\infty$)	mag	Exponential($1/\tau$)	
	pop. mean	τ	0.1	mag	fixed	
	total-to-selective extinction	R_V^s	[1.2; $+\infty$)	—	$\mathcal{N}(\mu_R, \sigma_R^2)$	
	pop. mean	μ_R	3	—	fixed	
	pop. st. dev.	σ_R	0.5	—	fixed	

Methods: Set-based TMNRE for hierarchical models

Bayesian hierarchical models feature a large number of free parameters, which scales with the number of observed objects. Accurate, precise, and fast inference from future data is thus a major computational challenge that likelihood-based methods are ill-suited to address. Moreover, the likelihood requires explicit calculations of intractable probabilities, often leading to *ad-hoc* approximations. Simulation-based inference (SBI) addresses all these issues—scalability, realism, and rigour—by delivering marginal Bayesian posteriors, given data \mathbf{D}_0 , for any parameters of interest θ , implicitly integrated over *all* nuisance parameters (\mathbf{v}) and relevant stochastic processes (e.g. selection and other systematic effects):

$$p(\theta | \mathbf{D}_0) \propto p(\theta) p(\mathbf{D}_0 | \theta) = p(\theta) \times \int p(\mathbf{D}_0 | \mathbf{v}, \theta) p(\mathbf{v} | \theta) d\mathbf{v}. \quad (28)$$

SBI requires only *samples* from the prior ($p(\theta)$) and the marginal likelihood ($p(\mathbf{D} | \theta)$), which are provided by a stochastic forward simulator that represents the Bayesian joint model $p(\theta, \mathbf{D})$. There are different techniques (flavours of neural SBI) for using simulated pairs (θ, \mathbf{D}) to train a neural network (NN) and later perform inference from observed data \mathbf{D}_0 . We adopt the approach called neural ratio estimation (NRE) [66] since it allows the greatest freedom in the NN architecture and in choosing priors for training and evaluation. It trains a network $\hat{r}(\theta, \mathbf{D})$ to approximate the single real number

$$r(\theta, \mathbf{D}) \equiv \frac{p(\theta, \mathbf{D})}{p(\theta) p(\mathbf{D})} = \frac{p(\theta | \mathbf{D})}{p(\theta)} \quad (29)$$

by minimisation of the binary cross-entropy (BCE) loss commonly used for classification tasks. Once trained, $\hat{r}(\theta, \mathbf{D}_0)$ evaluated at the observed data can be simply multiplied by the prior or used to re-weight prior samples to represent the target posterior, according to the second equality in (29).

In principle, θ can represent any group of parameters we wish to derive a posterior for. However, scientific interpretations are usually based on one- or two-dimensional marginal distributions, and consequently, we will define the following groups of (global) parameters of interest γ_g :

- one-dimensional: $M_0, \sigma_0, \alpha, \beta, \alpha_c, M_*^{\text{step}}$;
- two-dimensional: $[A, b]$ for the DTD and $[\Omega_{m0}, \Omega_{\Lambda 0}]$ for cosmology;
- two-dimensional: $[\Delta M, \gamma_{Z_*}]$, $[\Delta M, \gamma_{t_*}]$, and $[\gamma_{Z_*}, \gamma_{t_*}]$, forming all combinations of the host–SN Ia-connection parameters: a so-called “corner plot”.

For each of them, we train—simultaneously—separate ratio estimators \hat{r}_g , following shared data pre-processing.

Object-specific parameters

In addition to the global parameters, we will—simultaneously—train ratio estimators for all $O(N_{\text{sel}})$ object-specific (local) parameters in the hierarchical model (here we will use the unified label λ^i , which represents \mathbf{g}^h and λ^s in CIGaRS collectively). While we can, as before, form local-parameter groups λ_g from parameters of *the same object*, we will infer each object-specific parameter $\lambda_g \in \{z_c, M_*, Z_*, \delta, \tau_2, x_{\text{int}}, c_{\text{int}}\}$ marginally. While in figure 6 we only show a subset of the results that represent scientific interest in the present study, estimating *all* local parameters during training helps extract informative features and ultimately improves inference of the global parameters.

In **SICRET** and **SIDE-real**, we demonstrated simultaneous marginal inference of all $\{\lambda^i\}$ for models in which the data-set size N_{sel} was known *a priori*, i.e. it could be fixed after a survey is performed because selection effects were not considered. Here, we extend the approach to simulators that produce data sets with varying N_{sel} , which introduces two complications.

The first concerns identifying the objects. Within a BHM, the $\{\lambda^i\}$ are *a-priori* independent and identically distributed (i.i.d.), conditional on the set of global parameters γ , and each influences the sampling distribution of only one observed “object”; i.e. we have the general model structure

$$p(\gamma, \{\lambda^i\}, \{\mathbf{d}^i\}) = \left[\prod_{i=1}^{N_{\text{sel}}} p(\mathbf{d}^i | \lambda^i, \gamma) p(\lambda^i | \gamma) \right] p(\gamma). \quad (30)$$

Previously, we used a fixed ordered collection of *auxiliary* variables $[\mathbf{a}^i]_{i=1}^{N_{\text{sel}}}$ to disambiguate the assignment of labels i , which effectively individualised the sampling distributions:

$$p(\mathbf{d}^i | \lambda^i, \gamma) \rightarrow p(\mathbf{d}^i | \mathbf{a}^i, \lambda^i, \gamma) \rightarrow p_i(\mathbf{d}^i | \lambda^i, \gamma). \quad (31)$$

However, the present simulator produces unordered (*exchangeable*) sets of *a-priori* undetermined sizes, and so the auxiliary variables need to become an *output* of the model, i.e. be incorporated into the observable \mathbf{d}^i . Indeed, CIGaRS explicitly models the host photometry (from which redshift is ultimately derived) and the observational (fit) uncertainties, which comprised \mathbf{a}^i in **SICRET** and **SIDE-real**. Thus, we can treat λ^i and \mathbf{d}^i as *realisations* of singular random variables λ and \mathbf{d} :

$$\left[\prod_i p(\mathbf{d} = \mathbf{d}^i | \lambda = \lambda^i, \gamma) p(\lambda = \lambda^i | \gamma) \right] p(\gamma). \quad (32)$$

rather than collections of separate (i.i.d.) random variables, allowing us to train a *single* network to represent *all* posteriors

$$p(\lambda^i | \{\mathbf{d}^i\}) \rightarrow p(\lambda | \mathbf{d} = \mathbf{d}^i, \{\mathbf{d}^i\}), \quad (33)$$

where we have ignored the dependence on the full $\{\mathbf{d}^i\}$ since it is approximately redundant with the truncation of global parameters, as argued in **SICRET**.

Secondly, a technical complication arises since we need training pairs

$$\lambda, \mathbf{d} \sim p(\lambda, \mathbf{d}) = \int p(\lambda, \mathbf{d} | \gamma) p(\gamma) d\gamma, \quad (34)$$

but the simulator samples in proportion to $\langle N_{\text{sel}} \rangle (\gamma) \times p(\gamma)$ rather than simply $p(\gamma)$. This is not a problem when N_{sel} is a fixed input and can otherwise be rectified by randomly selecting a fixed number N_{sub} of objects from the simulator output. In fact, $N_{\text{sub}} := 1$ would be the natural choice in line with the above treatment of $\{\lambda^i\}$ and $\{\mathbf{d}^i\}$ as realisations of λ and \mathbf{d} , and we choose this for generating a validation set, which also represents the prior $p(\lambda)$. Still, by setting a larger N_{sub} , we can cheaply (i.e. without extra simulation) enlarge the effective batch size for training local-parameter inference networks, and so we choose $N_{\text{sub}} := 100$ in this case. Finally, for evaluating the results for the $N_{\text{sel},0}$ objects in \mathbf{D}_0 , we simply skip the sub-selection and use the full set ($N_{\text{sub}} := N_{\text{sel},0}$).

Network architecture: conditioned deep set++

The neural network we use (depicted, together with details of the sizes of its various layers, in figure 4) is based on the conditioned deep-set architecture from **STAR NRE** (following Zaheer et al. [68]) but augmented in depth as in Zhang et al. [67] due to the sophistication of the simulator and the inference tasks and with the addition of local-parameter estimators as in **SICRET** and **SIDE-real**.

Given a data set $\mathbf{D} \equiv \{\mathbf{d}^i\}$ with cardinality N_{sel} , we first pre-process each element \mathbf{d}^i with a small feed-forward network to (automatically) derive (non-linear) “features”:

$$\mathbf{d}^i = \text{DataPre}(\mathbf{d}^i), \quad (35)$$

which are generally useful across all inference tasks. These may include galaxy colours, (fiducially) standardised SN magnitudes, observational uncertainty (from the provided variances), as well as estimates of object-specific parameters directly usable in the respective downstream ratio estimators. While they are not forced to be directly “meaningful” to a human scientist, we plan to explore and interpret these features in future work. The $\{\mathbf{d}^i\}$ are then input for all ratio estimators.

Following Zaheer et al.’s representation theorem [68] and the considerations laid out in **STAR NRE**, the estimator for a group of global parameters θ_g makes use of two feed-forward networks $\hat{\phi}_g, \hat{\rho}_g$ and takes the form

$$\ln \hat{r}_g(\theta_g, \mathbf{D}) \equiv \hat{\rho}_g(\theta_g, \mathbf{S}_g, N_{\text{sel}}), \quad (36)$$

with

$$\mathbf{S}_g \equiv \frac{1}{N_{\text{sel}}} \sum_i \hat{\phi}_g\left(\theta_g, \frac{\mathbf{d}^i - \mathbf{m}}{\mathbf{s}}, \mathbf{m}, \mathbf{s}\right), \quad (37)$$

where \mathbf{m} and \mathbf{s} are, respectively, the mean and standard deviation of the $\{\mathbf{d}^i\}$. This implements the `SetNorm` operation advocated by Zhang et al. [67] but explicitly preserves the information contained in the set moments \mathbf{m} and \mathbf{s} . For expressivity (which allows the set aggregation to be a simple averaging operation), $\hat{\phi}_g$ is a deep residual network that implements “skip connections” as proposed by Zhang et al. [67] after combining its inputs $(\theta_g, \frac{\mathbf{d}^i - \mathbf{m}}{\mathbf{s}}, \mathbf{m}, \mathbf{s})$ into one vector through a single hidden network layer: see figure 4. Finally, $\hat{\rho}_g$ similarly combines its inputs $(\theta_g, \mathbf{S}_g, N_{\text{sel}})$ and feeds them through a few fully connected layers to output the single number $\ln \hat{r}_g(\theta_g, \mathbf{D})$.

The ratio estimator for a local/object-specific parameter λ_g is simpler since it does not need to aggregate information across the set, and so it consists of a single feed-forward network:

$$\ln \hat{r}_g(\lambda_g, \mathbf{d}^i) = \text{LocalNRE}_g(\lambda_g, \mathbf{d}^i), \quad (38)$$

which can be applied in parallel over N_{sub} input data at the same time.

Truncation and fine-tuning

Neural SBI is *amortised*: once the network is trained, results can be quickly derived from numerous simulated (or real, if available) data realisations. While this can be exploited to verify and calibrate inference [see e.g. 50, 51], it requires training over a wide range of data, which can be inefficient if one is only interested in the results for a single — *the real/target* — data set \mathbf{D}_0 .

In such cases, the simulation budget, training time, and network capacity can be optimised through various sequential SBI strategies, which modify (either continuously, or in a succession of *stages*) the prior $p(\theta) \rightarrow \tilde{p}(\theta)$ from which parameters are drawn, based on intermediate results during training. A simple yet effective prescription is *prior truncation* [65], in which the shape/form of $\tilde{p}(\theta)$ is unchanged, but its support is restricted to a region $T(\mathbf{D}_0)$ in which the posterior density is significantly different from zero:

$$\theta \notin T(\mathbf{D}_0) \implies p(\theta | \mathbf{D}_0) \approx 0, \quad (39)$$

as approximated by a previously trained network evaluated at \mathbf{D}_0 . Thus, the only needed modification is a trivial re-normalisation to account for the excluded prior probability mass:

$$p(\theta) \rightarrow \tilde{p}_{T(\mathbf{D}_0)}(\theta) = \begin{cases} \frac{p(\theta)}{\int_{T(\mathbf{D}_0)} p(\theta) d\theta} & \text{if } \theta \in T(\mathbf{D}_0), \\ 0 & \text{otherwise,} \end{cases} \quad (40)$$

which leaves the posterior (given \mathbf{D}_0) unchanged while restricting simulated examples to be more similar to \mathbf{D}_0 .

We apply truncation separately for each *global* inference group: in one dimension we use a simple contiguous interval that contains 99.99% approximate credibility and then sample training examples by analytically modifying the (simple: uniform or normal) prior; for two-dimensional groups, we either define an iso-likelihood contour (that again contains 99.99% approximate posterior mass) and use rejection sampling within it, or apply one-dimensional truncation separately for the two parameters (i.e. we truncate to an axis-aligned rectangular region) when they are not strongly correlated. The truncation regions are illustrated with gray shading in figure 5.

Once new training data is generated, we resume — instead of re-initialising — training, which means that the network is simply fine-tuned to give more accurate results in the “zoomed-in” parameter space, instead of forced to re-learn the analysis from scratch. This is especially relevant for the pre-processor and deep-set featurisers (described below), whose computations should not significantly depend on the parameter ranges.

Training details

Analysing the mock data with ≈ 1600 objects requires three rounds of truncation (starting from the priors in table 1 and with the network in figure 4) for results to converge (judged by the difference between posteriors from successive rounds). In each, we generate 64 000 full example data sets (with 6400 more for validation and to represent the priors when evaluating results) and train using the Adam optimizer [101] and learning rate 10^{-3} (reduced to 10^{-4} when fine-tuning) for at most 200 epochs, although the loss typically plateaus earlier. On two NVIDIA A100 GPUs with a combined mini-batch size of 64, one stage takes ≈ 24 h, with the majority of the computation taken up by *back-propagating* through all 11 deep sets (for the separate global-parameter groups).

For the larger ($N_{\text{sel}} \approx 16\,000$) data set, we take the final network from above and fine-tune it on simulations from the same final truncated prior region but with a tenfold increase in output size (on expectation). Since the deep-set architecture processes the full data at once, training it now requires roughly ten times more compute and memory, so we use 8 GPUs and a batch size of 32, which again requires 24 h per truncation stage (we only need one in this case).

Finally, for each global-parameter group γ_g , we pick the checkpoint with the lowest validation loss for the specific group, whereas for object-specific parameters λ_g , we pick the best checkpoint overall (lowest loss averaged over all groups). To represent the posteriors, we evaluate $\hat{f}(\theta_g, \mathbf{D}_0)$ over the parameter samples in the validation set, which represent $p(\theta_g)$, and use the results as importance weights for plotting contours and calculating posterior moments.

Software used. CIGaRS is based on the CLIPPY convenience layer for inference and probabilistic programming in Python (<https://github.com/kosiokarchev/clippy>). The forward simulator is implemented in Pyro [102] and makes use of the cosmographic utilities of *pytorch* (<https://github.com/kosiokarchev/pytorch>), the *prospector* code (<https://github.com/bd-j/prospector>), and the pre-trained *Speculator- α* emulator (<https://github.com/justinalsing/speculator>), embedded in SLICsim (<https://github.com/kosiokarchev/slicsim>). Defining and training the neural network uses PyTorch [103] and PyTorch Lightning [104].

Acknowledgements. RT acknowledges funding from Next Generation EU, in the context of the National Recovery and Resilience Plan, Investment PE1 – Project FAIR “Future Artificial Intelligence Research”, co-financed by the Next Generation EU [DM 1555/11.10.22]. RT is partially supported by the Fondazione ICSC, Spoke 3 “Astrophysics and Cosmos Observations”, Piano Nazionale di Ripresa e Resilienza Project ID CN00000013 “Italian Research Center on High-Performance Computing, Big Data and Quantum Computing” funded by MUR Missione 4 Componente 2 Investimento 1.4: Potenziamento strutture di ricerca e creazione di “campioni nazionali di R&S (M4C2-19)” - Next Generation EU (NGEU). Funding for the work of RJ was partially provided by project PID2022-141125NB-I00, grant CEX2024-001451-M funded by MICIU/AEI/10.13039/501100011033. This work was supported by a grant from the Simons Foundation (Grant Award ID BD-Targeted-00017375, UB).

References

- [1] Sullivan M, Ellis RS, Aldering G, et al (2003) The Hubble diagram of type Ia supernovae as a function of host galaxy morphology. *Monthly Notices of the Royal Astronomical Society* 340:1057–1075. <https://doi.org/10.1046/j.1365-8711.2003.06312.x>
- [2] Sullivan M, Le Borgne D, Pritchett CJ, et al (2006) Rates and Properties of Type Ia Supernovae as a Function of Mass and Star Formation in Their Host Galaxies. *The Astrophysical Journal* 648:868–883. <https://doi.org/10.1086/506137>
- [3] Aubourg É, Tojeiro R, Jimenez R, et al (2008) Evidence of short-lived SN Ia progenitors. *Astronomy and Astrophysics* 492:631–636. <https://doi.org/10.1051/0004-6361:200809796>
- [4] Neill JD, Sullivan M, Howell DA, et al (2009) The Local Hosts of Type Ia Supernovae. *The Astrophysical Journal* 707:1449–1465. <https://doi.org/10.1088/0004-637X/707/2/1449>
- [5] Brandt TD, Tojeiro R, Aubourg É, et al (2010) The Ages of Type Ia Supernova Progenitors. *The Astronomical Journal* 140:804–816. <https://doi.org/10.1088/0004-6256/140/3/804>
- [6] Lampeitl H, Smith M, Nichol RC, et al (2010) The Effect of Host Galaxies on Type Ia Supernovae in the SDSS-II Supernova Survey. *The Astrophysical Journal* 722:566–576. <https://doi.org/10.1088/0004-637X/722/1/566>
- [7] Sullivan M, Conley A, Howell DA, et al (2010) The dependence of Type Ia Supernovae luminosities on their host galaxies. *Monthly Notices of the Royal Astronomical Society* 406:782–802. <https://doi.org/10.1111/j.1365-2966.2010.16731.x>
- [8] Childress M, Aldering G, Antilogus P, et al (2013) Host Galaxy Properties and Hubble Residuals of Type Ia Supernovae from the Nearby Supernova Factory. *The Astrophysical Journal* 770:108. <https://doi.org/10.1088/0004-637X/770/2/108>
- [9] Johansson J, Thomas D, Pforr J, et al (2013) SN Ia host galaxy properties from Sloan Digital Sky Survey-II spectroscopy. *Monthly Notices of the Royal Astronomical Society* 435:1680–1700. <https://doi.org/10.1093/mnras/stt1408>
- [10] Lee YW, Chung C, Demarque P, et al (2022) Evidence for strong progenitor age dependence of type Ia supernova luminosity standardization process. *Monthly Notices of the Royal Astronomical Society* 517:2697–2708. <https://doi.org/10.1093/mnras/stac2840>
- [11] Chung C, Yoon SJ, Park S, et al (2023) On the Root Cause of the Host “Mass Step” in the Hubble Residuals of Type Ia Supernovae. *The Astrophysical Journal* 959:94. <https://doi.org/10.3847/1538-4357/ad0121>

- [12] Chung C, Park S, Son J, et al (2024) Strong progenitor age bias in supernova cosmology. I. Robust and ubiquitous evidence from a larger sample of host galaxies in a broader redshift range. <https://doi.org/10.48550/arXiv.2411.05299>, [arXiv:2411.05299](https://arxiv.org/abs/2411.05299)
- [13] Rigault M, Copin Y, Aldering G, et al (2013) Evidence of environmental dependencies of Type Ia supernovae from the Nearby Supernova Factory indicated by local $H\alpha$. *Astronomy and Astrophysics* 560:A66. <https://doi.org/10.1051/0004-6361/201322104>
- [14] Rigault M, Aldering G, Kowalski M, et al (2015) Confirmation of a Star Formation Bias in Type Ia Supernova Distances and its Effect on the Measurement of the Hubble Constant. *The Astrophysical Journal* 802:20. <https://doi.org/10.1088/0004-637X/802/1/20>
- [15] Rigault M, Brinnel V, Aldering G, et al (2020) Strong dependence of Type Ia supernova standardization on the local specific star formation rate. *Astronomy and Astrophysics* 644:A176. <https://doi.org/10.1051/0004-6361/201730404>
- [16] Jones DO, Riess AG, Scolnic DM (2015) Reconsidering the Effects of Local Star Formation on Type Ia Supernova Cosmology. *The Astrophysical Journal* 812:31. <https://doi.org/10.1088/0004-637X/812/1/31>
- [17] Jones DO, Riess AG, Scolnic DM, et al (2018) Should Type Ia Supernova Distances Be Corrected for Their Local Environments? *The Astrophysical Journal* 867:108. <https://doi.org/10.3847/1538-4357/aae2b9>
- [18] Moreno-Raya ME, López-Sánchez ÁR, Mollá M, et al (2016) Using the local gas-phase oxygen abundances to explore a metallicity dependence in SNe Ia luminosities. *Monthly Notices of the Royal Astronomical Society* 462:1281–1306. <https://doi.org/10.1093/mnras/stw1706>
- [19] Moreno-Raya ME, Mollá M, López-Sánchez ÁR, et al (2016) On the Dependence of Type Ia SNe Luminosities on the Metallicity of Their Host Galaxies. *The Astrophysical Journal* 818:L19. <https://doi.org/10.3847/2041-8205/818/1/L19>
- [20] Kim YL, Smith M, Sullivan M, et al (2018) Environmental Dependence of Type Ia Supernova Luminosities from a Sample without a Local-Global Difference in Host Star Formation. *The Astrophysical Journal* 854:24. <https://doi.org/10.3847/1538-4357/aaa127>
- [21] Kim YL, Kang Y, Lee YW (2019) Environmental Dependence of Type Ia Supernova Luminosities from the YONSEI Supernova Catalog. *Journal of Korean Astronomical Society* 52:181–205. <https://doi.org/10.5303/JKAS.2019.52.5.181>
- [22] Roman M, Hardin D, Betoule M, et al (2018) Dependence of Type Ia supernova luminosities on their local environment. *Astronomy and Astrophysics* 615:A68. <https://doi.org/10.1051/0004-6361/201731425>
- [23] Rose BM, Garnavich PM, Berg MA (2019) Think Global, Act Local: The Influence of Environment Age and Host Mass on Type Ia Supernova Light Curves. *The Astrophysical Journal* 874:32. <https://doi.org/10.3847/1538-4357/ab0704>
- [24] Kelsey L, Sullivan M, Smith M, et al (2021) The effect of environment on Type Ia supernovae in the Dark Energy Survey three-year cosmological sample. *Monthly Notices of the Royal Astronomical Society* 501:4861–4876. <https://doi.org/10.1093/mnras/staa3924>
- [25] Kelsey L, Sullivan M, Wiseman P, et al (2023) Concerning colour: The effect of environment on type Ia supernova colour in the dark energy survey. *Monthly Notices of the Royal Astronomical Society* 519:3046–3063. <https://doi.org/10.1093/mnras/stac3711>
- [26] Ginolin M, Rigault M, Copin Y, et al (2024) ZTF SN Ia DR2: Colour standardisation of Type Ia Supernovae and its dependence on environment. <https://doi.org/10.48550/arXiv.2406.02072>, [arXiv:2406.02072](https://arxiv.org/abs/2406.02072)
- [27] Ginolin M, Rigault M, Smith M, et al (2024) ZTF SN Ia DR2: Environmental dependencies of stretch and luminosity of a volume limited sample of 1,000 Type Ia Supernovae. URL <http://arxiv.org/abs/2405.20965>,

arXiv:2405.20965

- [28] Brout D, Scolnic D, Popovic B, et al (2022) The Pantheon+ Analysis: Cosmological Constraints. *The Astrophysical Journal* 938:110. <https://doi.org/10.3847/1538-4357/ac8e04>
- [29] Rubin D, Aldering G, Betoule M, et al (2025) Union through UNITY: Cosmology with 2000 SNe Using a Unified Bayesian Framework. *The Astrophysical Journal* 986:231. <https://doi.org/10.3847/1538-4357/adc0a5>
- [30] Vincenzi M, Brout D, Armstrong P, et al (2024) The Dark Energy Survey Supernova Program: Cosmological Analysis and Systematic Uncertainties. *The Astrophysical Journal* 975:86. <https://doi.org/10.3847/1538-4357/ad5e6c>
- [31] Duarte J, González-Gaitán S, Mourão A, et al (2025) Assessing differences between local galaxy dust attenuation and point source extinction within the same environments. <https://doi.org/10.48550/arXiv.2503.04906>, arXiv:2503.04906
- [32] Mandel KS, Scolnic DM, Shariff H, et al (2017) The Type Ia Supernova Color-Magnitude Relation and Host Galaxy Dust: A Simple Hierarchical Bayesian Model. *The Astrophysical Journal* 842:93. <https://doi.org/10.3847/1538-4357/aa6038>
- [33] Speagle JS, Steinhardt CL, Capak PL, et al (2014) A Highly Consistent Framework for the Evolution of the Star-Forming “Main Sequence” from $z \sim 0-6$. *The Astrophysical Journal Supplement Series* 214:15. <https://doi.org/10.1088/0067-0049/214/2/15>
- [34] Brout D, Scolnic D (2021) It’s Dust: Solving the Mysteries of the Intrinsic Scatter and Host-galaxy Dependence of Standardized Type Ia Supernova Brightnesses. *The Astrophysical Journal* 909:26. <https://doi.org/10.3847/1538-4357/abd69b>
- [35] Thorp S, Mandel KS (2022) Constraining the SN Ia host galaxy dust law distribution and mass step: Hierarchical BAYESN analysis of optical and near-infrared light curves. *Monthly Notices of the Royal Astronomical Society* 517:2360–2382. <https://doi.org/10.1093/mnras/stac2714>
- [36] Karchev K, Trotta R, Weniger C (2023) SimSIMS: Simulation-based Supernova Ia Model Selection with thousands of latent variables. <https://doi.org/10.48550/arXiv.2311.15650>, arXiv:2311.15650
- [37] Rubin D, Aldering G, Barbary K, et al (2015) UNITY: Confronting Supernova Cosmology’s Statistical and Systematic Uncertainties in a Unified Bayesian Framework. *The Astrophysical Journal* 813:137. <https://doi.org/10.1088/0004-637X/813/2/137>
- [38] Hinton SR, Davis TM, Kim AG, et al (2019) Steve: A Hierarchical Bayesian Model for Supernova Cosmology. *The Astrophysical Journal* 876:15. <https://doi.org/10.3847/1538-4357/ab13a3>
- [39] Thorp S, Mandel KS, Jones DO, et al (2021) Testing the consistency of dust laws in SN Ia host galaxies: A BAYESN examination of Foundation DR1. *Monthly Notices of the Royal Astronomical Society* 508:4310–4331. <https://doi.org/10/gm9wv5>
- [40] Kessler R, Scolnic D (2017) Correcting Type Ia Supernova Distances for Selection Biases and Contamination in Photometrically Identified Samples. *The Astrophysical Journal* 836:56. <https://doi.org/10.3847/1538-4357/836/1/56>
- [41] Shariff H, Jiao X, Trotta R, et al (2016) BAHAMAS: New Analysis of Type Ia Supernovae Reveals Inconsistencies with Standard Cosmology. *The Astrophysical Journal* 827:1. <https://doi.org/10.3847/0004-637X/827/1/1>
- [42] Eddington AS (1913) On a formula for correcting statistics for the effects of a known error of observation. *Monthly Notices of the Royal Astronomical Society* 73:359–360. <https://doi.org/10.1093/mnras/73.5.359>

- [43] Karchev K, Trotta R (2024) STAR NRE: Solving supernova selection effects with set-based truncated auto-regressive neural ratio estimation. <https://doi.org/10.48550/arXiv.2409.03837>, arXiv:2409.03837
- [44] Lee YW, Chung C, Kang Y, et al (2020) Further Evidence for Significant Luminosity Evolution in Supernova Cosmology. *The Astrophysical Journal* 903:22. <https://doi.org/10.3847/1538-4357/abb3c6>
- [45] Nicolas N, Rigault M, Copin Y, et al (2021) Redshift evolution of the underlying type Ia supernova stretch distribution. *Astronomy and Astrophysics* 649:A74. <https://doi.org/10.1051/0004-6361/202038447>
- [46] Popovic B, Rigault M, Smith M, et al (2024) ZTF SN Ia DR2: Evidence of Changing Dust Distributions With Redshift Using Type Ia Supernovae. <https://doi.org/10.48550/arXiv.2406.06215>, arXiv:2406.06215
- [47] Thorp S, Mandel KS, Jones DO, et al (2024) Using rest-frame optical and NIR data from the RAISIN survey to explore the redshift evolution of dust laws in SN Ia host galaxies. *Monthly Notices of the Royal Astronomical Society* 530:4016–4031. <https://doi.org/10.1093/mnras/stae1111>
- [48] Cranmer K, Brehmer J, Louppe G (2020) The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences* 117(48):30055–30062. <https://doi.org/10.1073/pnas.1912789117>
- [49] Lueckmann JM, Boelts J, Greenberg D, et al (2021) Benchmarking Simulation-Based Inference. In: Banerjee A, Fukumizu K (eds) *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*. PMLR, pp 343–351, URL <https://proceedings.mlr.press/v130/lueckmann21a.html>
- [50] Karchev K, Grayling M, Boyd BM, et al (2024) SIDE-real: Supernova Ia Dust Extinction with truncated marginal neural ratio estimation applied to real data. *Monthly Notices of the Royal Astronomical Society* 530:3881–3896. <https://doi.org/10.1093/mnras/stae995>
- [51] Karchev K, Trotta R, Weniger C (2023) SICRET: Supernova Ia Cosmology with truncated marginal neural Ratio Estimation. *Monthly Notices of the Royal Astronomical Society* 520:1056–1072. <https://doi.org/10.1093/mnras/stac3785>, arXiv:2209.06733 [astro-ph]
- [52] Leja J, Johnson BD, Conroy C, et al (2017) Deriving Physical Properties from Broadband Photometry with Prospector: Description of the Model and a Demonstration of its Accuracy Using 129 Galaxies in the Local Universe. *The Astrophysical Journal* 837:170. <https://doi.org/10.3847/1538-4357/aa5ffe>
- [53] Wang B, Leja J, Bezanson R, et al (2023) Inferring More from Less: Prospector as a Photometric Redshift Engine in the Era of JWST. *The Astrophysical Journal* 944:L58. <https://doi.org/10.3847/2041-8213/acba99>
- [54] Greggio L (2005) The rates of type Ia supernovae. I. Analytical formulations. *Astronomy and Astrophysics* 441:1055–1078. <https://doi.org/10.1051/0004-6361:20052926>
- [55] Palicio PA, Matteucci F, Della Valle M, et al (2024) Cosmic Type Ia supernova rate and constraints on supernova Ia progenitors. *Astronomy and Astrophysics* 689:A203. <https://doi.org/10.1051/0004-6361/202449740>
- [56] Ruiter AJ, Seitzzahl IR (2025) Type Ia supernova progenitors: A contemporary view of a long-standing puzzle. *Astronomy and Astrophysics Review* 33:1. <https://doi.org/10.1007/s00159-024-00158-9>
- [57] Frohmaier C, Vincenzi M, Sullivan M, et al (2025) TiDES: The 4MOST Time Domain Extragalactic Survey. <https://doi.org/10.48550/arXiv.2501.16311>, arXiv:2501.16311
- [58] Sánchez BO, Brout D, Vincenzi M, et al (2024) The Dark Energy Survey Supernova Program: Light Curves and 5 Yr Data Release. *The Astrophysical Journal* 975:5. <https://doi.org/10.3847/1538-4357/ad739a>
- [59] Gallazzi A, Charlot S, Brinchmann J, et al (2005) The ages and metallicities of galaxies in the local universe. *Monthly Notices of the Royal Astronomical Society* 362:41–58. <https://doi.org/10.1111/j.1365-2966.2005.09321.x>

- [60] Heringer E, Pritchett C, van Kerkwijk MH (2019) The Delay Times of Type Ia Supernova. *The Astrophysical Journal* 882:52. <https://doi.org/10.3847/1538-4357/ab32dd>
- [61] Dilday B, Kessler R, Frieman JA, et al (2008) A Measurement of the Rate of Type Ia Supernovae at Redshift $z \approx 0.1$ from the First Season of the SDSS-II Supernova Survey. *The Astrophysical Journal* 682:262–282. <https://doi.org/10.1086/587733>
- [62] Kessler R, Narayan G, Avelino A, et al (2019) Models and Simulations for the Photometric LSST Astronomical Time Series Classification Challenge (PLAsTiCC). *Publications of the Astronomical Society of the Pacific* 131:094501. <https://doi.org/10.1088/1538-3873/ab26f1>
- [63] Malmquist KG (1922) On some relations in stellar statistics. *Meddelanden fran Lunds Astronomiska Observatorium Serie I* 100:1–52. URL <https://ui.adsabs.harvard.edu/abs/1922MeLuF.100....1M>
- [64] Malmquist KG (1925) A contribution to the problem of determining the distribution in space of the stars. *Meddelanden fran Lunds Astronomiska Observatorium Serie I* 106:1–12. URL <https://ui.adsabs.harvard.edu/abs/1925MeLuF.106....1M>
- [65] Miller BK, Cole A, Forré P, et al (2021) Truncated marginal neural ratio estimation. In: Ranzato M, Beygelzimer A, Dauphin Y, et al (eds) *Advances in Neural Information Processing Systems, Neural Information Processing Systems*, vol 34. Curran Associates, Inc., pp 129–143, URL <https://proceedings.neurips.cc/paper/2021/hash/01632f7b7a127233fa1188bd6c2e42e1-Abstract.html>
- [66] Hermans J, Begy V, Louppe G (2020-07-13/2020-07-18) Likelihood-free MCMC with amortized approximate ratio estimators. In: III HD, Singh A (eds) *Proceedings of the 37th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol 119. PMLR, pp 4239–4248, URL <https://proceedings.mlr.press/v119/hermans20a.html>
- [67] Zhang L, Tozzo V, Higgins J, et al (2022-07-17/2022-07-23) Set norm and equivariant skip connections: Putting the deep in deep sets. In: Chaudhuri K, Jegelka S, Song L, et al (eds) *Proceedings of the 39th International Conference on Machine Learning, Proceedings of Machine Learning Research*, vol 162. PMLR, pp 26559–26574, URL <https://proceedings.mlr.press/v162/zhang22ac.html>
- [68] Zaheer M, Kottur S, Ravanbakhsh S, et al (2017) Deep Sets. In: Guyon I, Von Luxburg U, Bengio S, et al (eds) *Advances in Neural Information Processing Systems, Neural Information Processing Systems*, vol 30. Curran Associates, Inc., URL <https://proceedings.neurips.cc/paper/2017/hash/f22e4747da1aa27e363d86d40ff442fe-Abstract.html>
- [69] Wiseman P, Sullivan M, Smith M, et al (2021) Rates and delay times of Type Ia supernovae in the Dark Energy Survey. *Monthly Notices of the Royal Astronomical Society* 506:3330–3348. <https://doi.org/10.1093/mnras/stab1943>
- [70] Leistedt B, Alsing J, Peiris H, et al (2023) Hierarchical Bayesian Inference of Photometric Redshifts with Stellar Population Synthesis Models. *The Astrophysical Journal Supplement Series* 264:23. <https://doi.org/10.3847/1538-4365/ac9d99>
- [71] Merz G, Liu X, Schmidt S, et al (2024) DeepDISC-photoz: Deep Learning-Based Photometric Redshift Estimation for Rubin LSST. <https://doi.org/10.48550/arXiv.2411.18769>, [arXiv:2411.18769](https://arxiv.org/abs/2411.18769)
- [72] Chen R, Scolnic D, Vincenzi M, et al (2024) Evaluating Cosmological Biases using Photometric Redshifts for Type Ia Supernova Cosmology with the Dark Energy Survey Supernova Program. URL <http://arxiv.org/abs/2407.16744>, [arXiv:2407.16744](https://arxiv.org/abs/2407.16744)
- [73] Panter B, Heavens AF, Jimenez R (2003) Star formation and metallicity history of the SDSS galaxy survey: Unlocking the fossil record. *Monthly Notices of the Royal Astronomical Society* 343:1145–1154. <https://doi.org/10.1046/j.1365-8711.2003.06722.x>

- [74] Heavens A, Panter B, Jimenez R, et al (2004) The star-formation history of the Universe from the stellar populations of nearby galaxies. *Nature* 428:625–627. <https://doi.org/10.1038/nature02474>
- [75] Panter B, Jimenez R, Heavens AF, et al (2007) The star formation histories of galaxies in the Sloan Digital Sky Survey. *Monthly Notices of the Royal Astronomical Society* 378:1550–1564. <https://doi.org/10.1111/j.1365-2966.2007.11909.x>
- [76] Mo H, van den Bosch F, White S (2010) *Galaxy Formation and Evolution*. Cambridge University Press
- [77] Johnson BD, Leja J, Conroy C, et al (2021) Stellar Population Inference with Prospector. *The Astrophysical Journal Supplement Series* 254:22. <https://doi.org/10.3847/1538-4365/abef67>
- [78] Tojeiro R, Heavens AF, Jimenez R, et al (2007) Recovering galaxy star formation and metallicity histories from spectra using VESPA. *Monthly Notices of the Royal Astronomical Society* 381:1252–1266. <https://doi.org/10.1111/j.1365-2966.2007.12323.x>
- [79] Bravo E, Badenes C (2011) Is the metallicity of their host galaxies a good measure of the metallicity of Type Ia supernovae? *Monthly Notices of the Royal Astronomical Society* 414:1592–1606. <https://doi.org/10.1111/j.1365-2966.2011.18498.x>
- [80] Kriek M, Conroy C (2013) The Dust Attenuation Law in Distant Galaxies: Evidence for Variation with Spectral Type. *The Astrophysical Journal* 775:L16. <https://doi.org/10.1088/2041-8205/775/1/L16>
- [81] Calzetti D, Armus L, Bohlin RC, et al (2000) The Dust Content and Opacity of Actively Star-forming Galaxies. *The Astrophysical Journal* 533:682–695. <https://doi.org/10.1086/308692>
- [82] Noll S, Burgarella D, Giovannoli E, et al (2009) Analysis of galaxy spectral energy distributions from far-UV to far-IR with CIGALE: Studying a SINGS test sample. *Astronomy and Astrophysics* 507:1793–1813. <https://doi.org/10.1051/0004-6361/200912497>
- [83] Alsing J, Peiris H, Mortlock D, et al (2023) Forward Modeling of Galaxy Populations for Cosmological Redshift Distribution Inference. *The Astrophysical Journal Supplement Series* 264:29. <https://doi.org/10.3847/1538-4365/ac9583>
- [84] Mosenkov AV, Baes M, Bianchi S, et al (2019) Dust emission profiles of DustPedia galaxies. *Astronomy and Astrophysics* 622:A132. <https://doi.org/10.1051/0004-6361/201833932>
- [85] Hahn C, Starkenburg TK, Anglés-Alcázar D, et al (2022) IQ Collaboratory. III. The Empirical Dust Attenuation Framework-Taking Hydrodynamical Simulations with a Grain of Dust. *The Astrophysical Journal* 926:122. <https://doi.org/10.3847/1538-4357/ac4253>
- [86] Zhang J, Wuyts S, Cutler SE, et al (2023) Dust attenuation, dust content, and geometry of star-forming galaxies. *Monthly Notices of the Royal Astronomical Society* 524:4128–4147. <https://doi.org/10.1093/mnras/stad2066>
- [87] Baes M, Gebek A, Trčka A, et al (2024) The TNG50-SKIRT Atlas: Post-processing methodology and first data release. *Astronomy and Astrophysics* 683:A181. <https://doi.org/10.1051/0004-6361/202348418>
- [88] Hill R, Shariff H, Trotta R, et al (2018) Projected distances to host galaxy reduce SNIa dispersion. *Monthly Notices of the Royal Astronomical Society* 481:2766–2777. <https://doi.org/10.1093/mnras/sty2510>
- [89] Toy M, Wiseman P, Sullivan M, et al (2025) Reduction of the type Ia supernova host galaxy step in the outer regions of galaxies. *Monthly Notices of the Royal Astronomical Society* p staf248. <https://doi.org/10.1093/mnras/staf248>
- [90] Alsing J, Peiris H, Leja J, et al (2020) SPECULATOR: Emulating Stellar Population Synthesis for Fast and Accurate Galaxy Spectra and Photometry. *The Astrophysical Journal Supplement Series* 249:5. <https://doi.org/10.3847/1538-4365/ab917f>

- [91] Bianco FB, Jones L, Ivezić Ž, et al (2022) Updated estimates of the Rubin system throughput and expected LSST image depth. Tech. Rep. PSTN-054, Vera C. Rubin Observatory
- [92] Heringer E, Pritchett C, Kezwer J, et al (2017) Type Ia Supernovae: Colors, Rates, and Progenitors. *The Astrophysical Journal* 834:15. <https://doi.org/10.3847/1538-4357/834/1/15>
- [93] Dhawan S, Mortsell E, Johansson J, et al (2024) ZTF SN Ia DR2: Cosmology-independent constraints on Type Ia supernova standardisation from supernova siblings. <https://doi.org/10.48550/arXiv.2406.01434>, [arXiv:2406.01434](https://arxiv.org/abs/2406.01434)
- [94] Poetrodjojo H, Groves B, Kewley LJ, et al (2021) The SAMI Galaxy Survey: Reconciling strong emission line metallicity diagnostics using metallicity gradients. *Monthly Notices of the Royal Astronomical Society* 502:3357–3373. <https://doi.org/10.1093/mnras/stab205>
- [95] March MC, Trotta R, Berkes P, et al (2011) Improved constraints on cosmological parameters from Type Ia supernova data. *Monthly Notices of the Royal Astronomical Society* 418:2308–2329. <https://doi.org/10.1111/j.1365-2966.2011.19584.x>
- [96] Carreres B, Chen RC, Peterson ER, et al (2025) Type Ia supernova growth-rate measurement with LSST simulations: Intrinsic scatter systematics. <https://doi.org/10.48550/arXiv.2505.13290>, [arXiv:2505.13290](https://arxiv.org/abs/2505.13290)
- [97] Rosselli D, Carreres B, Ravoux C, et al (2025) Forecast for growth-rate measurement using peculiar velocities from LSST supernovae. <https://doi.org/10.48550/arXiv.2507.00157>, [arXiv:2507.00157](https://arxiv.org/abs/2507.00157)
- [98] Boyd BM, Grayling M, Thorp S, et al (2024) Accounting for Selection Effects in Supernova Cosmology with Simulation-Based Inference and Hierarchical Bayesian Modelling. <https://doi.org/10.48550/arXiv.2407.15923>, [arXiv:2407.15923](https://arxiv.org/abs/2407.15923)
- [99] Gupta RR, Kuhlmann S, Kovacs E, et al (2016) Host Galaxy Identification for Supernova Surveys. *The Astronomical Journal* 152:154. <https://doi.org/10.3847/0004-6256/152/6/154>
- [100] Hinshaw G, Larson D, Komatsu E, et al (2013) Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. *The Astrophysical Journal Supplement Series* 208:19. <https://doi.org/10.1088/0067-0049/208/2/19>
- [101] Kingma DP, Ba J (2017) Adam: A Method for Stochastic Optimization. <https://doi.org/10.48550/arXiv.1412.6980>, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
- [102] Bingham E, Chen JP, Jankowiak M, et al (2019) Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research* 20(1):973–978
- [103] Paszke A, Gross S, Massa F, et al (2019) PyTorch: An Imperative Style, High-Performance Deep Learning Library. In: Wallach H, Larochelle H, Beygelzimer A, et al (eds) *Advances in Neural Information Processing Systems* 32. Curran Associates, Inc., p 8024–8035
- [104] Falcon W, The PyTorch Lightning team (2023) PyTorch Lightning. Zenodo, <https://doi.org/10.5281/zenodo.3530844>