The Astrophysical Journal, 982:110 (21pp), 2025 April 1 https://doi.org/10.3847/1538-4357/ad9f32

An Agnostic Approach to Building Empirical Type Ia Supernova Light Curves: Evidence for Intrinsic Chromatic Flux Variation Using Nearby Supernova Factory Data

Jared Hand[1], A. G. Kim[2], G. Aldering[2], P. Antilogus[3], C. Aragon[2,4], S. Bailey[2], C. Baltay[5], S. Bongard[3], K. Boone[2,6,7], C. Buton[8], Y. Copin[8], S. Dixon[2,6], D. Fouchez[9], E. Gangler[8,10], R. Gupta[2], B. Hayden[2,11], W. Hillebrandt[12], Mitchell Karmen[2], M. Kowalski[13,14], D. Küsters[6,14], P.-F. Léget[3], F. Mondon[10], J. Nordin[2,13], R. Pain[3], E. Pecontal[15], R. Pereira[8], S. Perlmutter[2,6], K. A. Ponder[6], D. Rabinowitz[5], M. Rigault[8], D. Rubin[2,16], K. Runge[2], C. Saunders[2,6,17,18], N. Suzuki[2,19], C. Tao[9,20], S. Taubenberger[12], R. C. Thomas[2,21], M. Vincenzi[2,22]
(The Nearby Supernova Factory)

[1] Pittsburgh Particle Physics, Astrophysics, and Cosmology Center (PITT PACC), Department of Physics and Astronomy, University of Pittsburgh, Pittsburgh, PA 15260, USA
[2] Physics Division, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA
[3] Laboratoire de Physique Nucléaire et des Hautes Energies, CNRS/IN2P3, Sorbonne Université, Université de Paris, 4 place Jussieu, 75005 Paris, France
[4] College of Engineering, University of Washington, 371 Loew Hall, Seattle, WA 98195, USA
[5] Department of Physics, Yale University, New Haven, CT 06250-8121, USA
[6] Department of Physics, University of California Berkeley, 366 LeConte Hall MC 7300, Berkeley, CA 94720-7300, USA
[7] DIRAC Institute, Department of Astronomy, University of Washington, 3910 15th Ave NE, Seattle, WA 98195, USA
[8] Univ Lyon, Université Claude Bernard Lyon 1, CNRS/IN2P3, IP2I Lyon, F-69622, Villeurbanne, France
[9] Aix Marseille Univ, CNRS/IN2P3, CPPM, Marseille, France
[10] Université Clermont Auvergne, CNRS/IN2P3, Laboratoire de Physique de Clermont, F-63000 Clermont-Ferrand, France
[11] Space Telescope Science Institute, 3700 San Martin Drive Baltimore, MD 21218, USA
[12] Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, D-85748 Garching, Germany
[13] Institut für Physik, Humboldt-Universitat zu Berlin, Newtonstr. 15, 12489 Berlin, Germany
[14] DESY, D-15735 Zeuthen, Germany
[15] Centre de Recherche Astronomique de Lyon, Université Lyon 1, 9 Avenue Charles André, 69561 Saint Genis Laval Cedex, France
[16] Department of Physics and Astronomy, University of Hawai'i, 2505 Correa Rd, Honolulu, HI 96822, USA
[17] Princeton University, Department of Astrophysics, 4 Ivy Lane, Princeton, NJ 08544, USA
[18] Sorbonne Université, Institut Lagrange de Paris (ILP), 98 bis Boulevard Arago, 75014 Paris, France
[19] Kavli Institute for the Physics and Mathematics of the Universe, The University of Tokyo Institutes for Advanced Study, The University of Tokyo, 5-1-5 Kashiwanoha, Kashiwa, Chiba 277-8583, Japan
[20] Tsinghua Center for Astrophysics, Tsinghua University, Beijing 100084, People's Republic of China
[21] Computational Cosmology Center, Computational Research Division, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA
[22] Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth PO1 3FX, UK
Received 2023 July 25; revised 2024 November 26; accepted 2024 December 13; published 2025 March 21

Abstract

We present a new empirical Type Ia supernova (SN Ia) model with three chromatic flux variation templates: one phase dependent and two phase independent. No underlying dust extinction model or patterns of intrinsic variability are assumed. Implemented with `Stan` and trained using spectrally binned Nearby Supernova Factory spectrophotometry, we examine this model's 2D, phase-independent flux variation space using two motivated basis representations. In both, the first phase-independent template captures variation that appears dust-like, while the second captures a combination of effectively intrinsic variability and second-order dust-like effects. We find that ≈13% of the modeled phase-independent flux variance is not dust-like. Previous empirical SN Ia models either assume an effective dust extinction recipe in their architecture, or only allow for a single mode of phase-independent variation. The presented results demonstrate such an approach may be insufficient, because it could "leak" noticeable intrinsic variation into phase-independent templates.

Unified Astronomy Thesaurus concepts: Observational cosmology (1146); Type Ia supernovae (1728)

1. Introduction

The discovery of dark energy with standardized Type Ia supernovae (SNe Ia) solidified these transient objects' importance to cosmology (A. G. Riess et al. 1998b; S. Perlmutter et al. 1999). SNe Ia exhibit a markedly similar peak *B*-band brightness dispersion of only ~1 mag, and this dispersion can be reduced further with multifilter photometric time series (or light curves) by exploiting correlations between light-curve duration (also called shape, width, or stretch) and color with *B*-band maximum brightness (M. M. Phillips 1993; A. G. Riess et al. 1996; S. Perlmutter et al. 1997; R. Tripp 1998). These standardization results reduce the brightness dispersion by an order of magnitude to ≈0.12 mag (D. M. Scolnic et al. 2018). Standardization using spectra can reduce the dispersion even further, to ≈0.08 mag (H. K. Fakhouri et al. 2015; K. Boone et al. 2021a).

There are two errors that limit the capacity of SNe Ia to constrain cosmological parameters: the number of observed SNe Ia (statistical uncertainty), and errors resulting from observation bias, modeling errors, and calibration (systematics). Recent SN Ia cosmology analyses have significantly reduced the statistical uncertainty by utilizing over 10^3 spectroscopically confirmed SNe Ia (M. Betoule et al. 2014;

D. M. Scolnic et al. 2018). LSST, via the Vera Rubin Observatory, will further increase our usable SN Ia sample by at least an order of magnitude over 10 yr after its commissioning (Ž. Ivezić et al 2019). Although internal photometric calibration will remain an important systematic to account for, LSST will alleviate tedious intersurvey photometric calibration systematics performed in many past analyses while still providing impressive statistics. As a result, LSST will increase the relative importance of systematics arising from SN Ia light-curve modeling and standardization for constraining cosmological parameters.

SNe Ia result from the catastrophic disruption of carbon–oxygen white dwarfs (F. Hoyle & W. A. Fowler 1960). Potential progenitor scenarios include accretion from a white dwarf's companion star or merging/collision of white dwarf binary constituents (J. Whelan & I. J. Iben 1973; I. J. Iben & A. V. Tutukov 1984). Unfortunately, accounting for observed SN Ia spectral variation while simultaneously recovering established empirical relations within the context of a detonating (or deflagrating) white dwarf framework remains a daunting and incomplete task. For example, it is suggested that progenitor mass can extend below an otherwise expected white dwarf mass limit, or Chandrasekhar mass, of 1.4 $M_{\odot}$ (R. A. Scalzo et al. 2014). See N. Soker (2019) for a recent review of and summary of SN Ia theoretical modeling progress and unanswered questions. Usable parametric theoretical SN Ia models remain elusive, leaving us reliant on empirical models trained from SN Ia observations.

1.1. Photometric Variation and Empirical Models

As mentioned, light-curve width correlates with B-band maximum brightness so that longer duration SNe Ia are systematically brighter (M. M. Phillips 1993). We refer to this as the width–luminosity relation (WLR). Similarly, bluer SN Ia light curves are systematically brighter at B-band maximum, which we similarly refer to as the color–luminosity relation (A. G. Riess et al. 1996).

One can interpret SN Ia empirical models as transforming high-dimensional sets of observations to a lower-dimensional set of parametric latent templates that inscribe dominant modes of SN Ia variability. The first generation of models attempted to describe light curves in standard rest-frame filters using templates specific to those filters (A. G. Riess et al. 1996; S. Jha et al. 2007; C. R. Burns et al. 2011), which when applied to measurements made with different rest-frame filters required a K-correction (A. Kim et al. 1996; P. Nugent et al. 2002) to convert the data to the model photometric system. The limitations of models based on broadband templates were quickly recognized: K-corrections carry their own set of uncertainties and biases that are correlated with model parameters (K. S. Mandel et al. 2022). More importantly, by restricting the templates to broad band, these models were insensitive to substantially more complex SN Ia variability revealed through spectroscopy.

SN Ia spectral variability is either intrinsic to SN Ia populations or is extrinsic, instead arising from processes external to the explosion, such as dust extinction or the SN Ia's interaction with its circumstellar environments (P. Nugent et al. 2002; S. Jha et al. 2007). Furthermore, photometric relationships emerge from spectral variation, with the temperature dependence of Fe line blanketing at least partly driving the WLR and explaining its wavelength dependence being such an example (D. Kasen & S. E. Woosley 2007). Variation in progenitor mass also contributes to the WLR, as lower-mass SN Ia progenitors are systematically dimmer and have faster-declining light curves compared to their more massive counterparts (R. A. Scalzo et al. 2014). Certain spectral features directly correlate with photometric SN Ia properties, such as the F(6420 Å)/F(4430 Å) line ratio correlating with maximum B-band brightness (S. Bailey et al. 2009), and the ratio of Si II 5972 and 6355 Å (or Si II 5972 and 3858 Å) correlating with light-curve width (P. Nugent et al. 1995).

Many SN Ia subtypes have been categorized by their spectral variation (D. Branch et al. 2009; S. Blondin et al. 2012). Grouping SNe Ia based on Si II velocity at maximum brightness during a Tripp standardization procedure (R. Tripp 1998) has been shown to reduce SN Ia dispersion poststandardization more than use of only color and stretch alone (X. Wang et al. 2009; R. J. Foley 2013). Furthermore, spectral information can improve the effective total-to-selective extinction R_V estimation, with N. Chotard et al. (2011) using spectral features to recover an effective R_V value consistent the Milky Way average of $R_V = 3.1$. Given this plethora of spectral variety within SNe Ia, and this variety's potential to further improve standardization, commonly used and recent SN Ia models make heavy use of spectroscopic observations in training.

Most recent SN Ia models reduce the dimensionality of SN Ia observations by constructing combinations of underlying spectral or color variation templates, with one template capturing the average, or fiducial, spectral evolution of SNe Ia. This approach removes any need for K-corrections and related uncertainty propagation (J. Guy et al. 2007; K. S. Mandel et al. 2022). The ubiquitous SALT model family (and its cousin SiFTO) is the canonical example of the spectral template technique (J. Guy et al. 2007; A. Conley et al. 2008; M. Betoule et al. 2014; J. D. R. Pierel et al. 2022). This family of linear SN Ia models captures variation beyond a mean spectral surface with a first-order flux variation template and a phase-independent color template (with per-SN contribution parameters x_1 and c, respectively). SALT2's success over prior models saw its widespread adoption and continuous improvement, with the most recent version SALT3 extending its wavelength coverage to the near-infrared (NIR; W. D. Kenworthy et al. 2021). More statistically rigorous linear spectral template models have also been developed, such as BayeSN with its potent hierarchical Bayesian framework (K. S. Mandel et al. 2017, 2022; S. Thorp et al. 2021).

A plethora of sophisticated models have recently been developed using The Nearby Supernova Factory (SNfactory) spectrophotometric time series (G. Aldering et al. 2002). C. Saunders et al. (2018) and their SNEMO model extracts up to 15 linear principal functional components from a set of SN Ia spectral surfaces trained using Gaussian processes to maximally explain SN Ia variation. Alternatively, P. F. Léget et al. (2020) in their SUGAR model treats SN Ia variation as a linear combination of spectral index templates, extending the initial work of N. Chotard et al. (2011) into a fully generative model. SNEMO and SUGAR, along with all the models mentioned so far, utilize linear dimensionality reduction techniques.

K. Boone et al. (2021b) apply the Isomap method with their Twins Embedding model to train a nonlinear parameterization of intrinsic SN Ia variation at maximum brightness, while

G. Stein et al. (2022) introduce a nonlinear probabilistic autoencoder that captures intrinsic variation across both wavelength and phase. Both find that a nonlinear approach requires only three intrinsic model components to describe SN Ia spectral variation where more traditional linear principal component analysis model would require seven components or more (C. Saunders et al. 2018). These two models also demonstrate noticeable improvements over SALT2 in standardized SN Ia dispersion from $\approx$0.12 to $\lesssim$0.09 mag. Improvements through nonlinear technique application are not limited to light-curve models: D. Rubin et al. (2015) introduce the hierarchical Bayesian framework UNITY that allows for nonlinear standardization, leading again to improved SN Ia dispersion poststandardization relative to the linear Tripp approach.

1.2. Shortcomings in Phase-independent Modeling

All past models either assume a dust extinction model for explicit phase-independent templates (MCLS2k2, SNooPY, SNEMO, SUGAR, BayeSN, and current nonlinear models) or include a single phase-independent template which does not differentiate between intrinsic and extrinsic variation (the SALT family). Excluding the maximum brightness model Twins Embedding, each of these models characterize phase-independent variability with only a single model component. Physical considerations alone demonstrate this to be an insufficient treatment. As summarized in J. C. Weingartner & B. T. Draine (2001), one would expect at least two extrinsic variation parameters per SN Ia: one gauging dust column density or optical depth (i.e., A_V) and the other probing second-order characteristics such as dust grain properties (i.e., R_V). Furthermore, it is plausible that empirical SN Ia models could extract intrinsic variation into a phase-independent template set. In this era of precision cosmology, modeling and standardization systematics remain stubborn obstacles to maximizing current and future SN Ia survey utility. Better understanding underlying extracted modes of phase-independent variation could answer outstanding questions about the SN Ia population (i.e., the low SN Ia R_V debate or the source of the bias associated with host properties), and improve both SN Ia modeling and standardization.

We present a new SN Ia empirical model to more deeply explore the phase-independent variability of SNe Ia. This model features three chromatic flux variation templates: one phase dependent and two phase independent. These two phase-dependent components provide the flexibility to account for multiparameter dust models while also absorbing an intrinsic time-averaged flux variation beyond that accounted for by the phase-dependent component. All templates are physics agnostic, as no assumptions are made about expected spectral features or dust treatment. This new model is trained on SNfactory's rest-frame spectrophotometric time series.

Our model bears some similarities with BayeSN. Both models are linear models implemented with `Stan`. How phase-independent variability is accounted for varies in approach, though. BayeSN implements a single-component dust extinction recipe within a hierarchical model framework, from which they recover an effective R_V that is largely consistent with the Milky Way average. In contrast, our model has two phase-independent templates, providing it two model degrees of freedom for which no physical assumptions are made. Unlike BayeSN, our model does not model spectral surface residuals.

In Section 2 we introduce the training set and its quality cuts. We describe our model and fitting technique in Section 3, with global model template and refit per-SN parameter results being presented in Section 4. Finally, we provide concluding remarks in Section 5.

2. Data

Between 2004 and 2014 SNfactory observed spectrophotometric time series of nearly 300 SNe Ia (G. Aldering et al. 2002) with the SuperNova Integral Field Spectrograph (SNIFS; B. Lantz et al. 2004). SNIFS is continuously mounted at the University of Hawai'i 2.2 m telescope, using dual-channel, moderate resolution (R ~ 600–1300) spectrographs to simultaneously observe transient events from 3200 to 5200 Å and 5100–10000 Å, respectively. This unique and homogeneous SN Ia data set is calibrated with CALSPEC and Hamuy standard stars (M. Hamuy et al. 1992, 1994; R. C. Bohlin 2014). The photometric calibration method is largely summarized in C. Buton et al. (2013) with R. Pereira et al. (2013) further describing non-photometric-night calibration. Host-galaxy-subtraction methodology is presented in S. Bongard et al. (2011). Each SNe Ia has also been fit using SALT2.4 (M. Betoule et al. 2014).

For our SN Ia training sample we generate synthetic SN-frame photometry using n_λ log-distributed top-hat filters from published rest-frame SNfactory spectra (G. Aldering et al. 2020). These spectra were first corrected for Milky Way extinction, then their wavelengths were de-redshifted (from an initial range of approximately $0.01 < z < 0.08$), and then they are placed on a relative luminosity scale (i.e., that equivalent to having been observed at $z = 0.05$). Cosmological time dilation is accounted for during de-redshifting. Observed spectra are in units of 10^{10} erg s^{-1} cm^{-2}. This work does not attempt to fit absolute magnitudes or fit for cosmological parameters, so per-SN redshifts are not used. Due to high-flux variance at wavelength boundaries, and because most objects have a higher redshift than $z = 0.05$, the per-spectra reference frame wavelength range is truncated to between 3350 and 8030 Å.

The spectral resolution of this top-hat filter synthetic photometry is flexible—for this work, we use a modest $n_\lambda = 10$ filter count. These $n_\lambda = 10$ bins are spaced at constant spectral resolution R, providing a wavelength bin size of $\approx$400 Å resolution for bluer bands and $\approx$500 Å for redder bands. SNfactory data consist of flux densities along a uniform grid of wavelengths. Because of this uniform spacing, we simply sum flux densities along the wavelength range defining our top-hat filters and then multiply said sum by the filter's wavelength range to calculate top-hat synthetic photometry:

$$F_{\rm bin} = [\lambda_{\rm max} - \lambda_{\rm min}] \sum_{i=i_{\rm min}}^{\rm max} f_{\lambda_i}. \quad (1)$$

Similarly, each corresponding variance spectrum is summed in quadrature to calculate synthetic photometry uncertainties.

For every observation, for the binned synthetic photometry a signal-to-noise ratio (SNR) of at least SNR > 5 is required. It is also required that each SN Ia have at least eight separate days of observations. Where a single SN Ia has multiple spectra observed within a few hours, the weighted average of these flux values is used as a single effective observation. This decision was made to avoid introducing two timescales in the sampling of our light curves. Given the difficulty in constraining the date

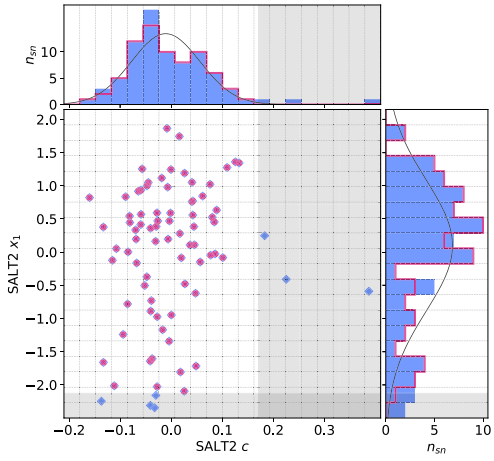

Figure 1. SALT2 c and x_1 cuts to better capture a Gaussian "core" for training. The shaded regions correspond to 2σ clippings along the longer tail of each respective c and x_1 distribution; blue points correspond to σ-clipped SNe.

of maximum in SN Ia empirical models, we demand that there exist at least one observation two days before the SALT2 maximum phase. Furthermore, no SNe Ia with an observation gap greater than 4 days within a 4 day range before and after the SALT2 maximum are used. SN Ia light curves do not have significant structure less than 4 days, so gaps of this size or smaller have no discernible impact on our results. For consistency, the chosen maximum gap size of 4 days is the same as our fixed Gaussian process mean predictor (GPMP) length scale hyperparameter later described in Section 3. These cuts leave 80 SNe Ia in the training sample.

The distribution of SN Ia color parameters for any model is asymmetric due to the positive-definite nature of dust extinction (D. M. Scolnic et al. 2014; K. S. Mandel et al. 2017; D. Brout & D. Scolnic 2021). SN Ia stretch parameters such as SALT2's stretch proxy x_1 are also best modeled with asymmetric distributions (D. M. Scolnic et al. 2014). We are interested in the Gaussian core of these distributions and partly "symmetrize" the data set by clipping the extended tails of our SALT2 c and x_1 samples. Specifically, a 2σ clipping is done on each SALT2 c and x_1 parameter sample in the direction each parameter's longer tail (Figure 1). The c clipping prevents heavily reddened SNe Ia from dominating the recovered dust-like behavior and obscuring the dust properties of the average SN Ia in training set—this c cut removes particularly reddened SNe Ia with peak apparent $B - V > 0.18$. The clipping is motivated by our interest in the core of the SN distribution that is used for cosmology, but comes at the expense of removing rarer objects that potentially provide more information on modeling SN Ia colors. A total of 73 SNe Ia remain after the SALT2 parameter σ clippings, consisting of 1155 individual spectra.

Relative to G. Aldering et al. (2020), and similar in spirit to K. Boone et al. (2021b), we remove spectra having poor extractions caused by SNR < 3; adjusting this threshold higher up to 10 did not affect our results. This step removes seven spectra, leaving 1148 to train the model.

3. Model

Global template parameters and per-SN parameters are differentiated by upper-case and lower-case characters, respectively. This model discretizes phase and wavelength space, using $n_p = 16$ phase nodes ranging from $-16 \leqslant t_{p,i} \leqslant 44$ in 4 day intervals; as mentioned in Section 2, $n_\lambda = 10$ with bins of constant R. Each phase-dependent template is an $n_p \times n_\lambda$ matrix of parameter nodes, while each phase-independent template is a length-n_λ vector.

The model prediction of the time-dependent spectral energy distribution evolution of an individual SN Ia is based on a temporal interpolation over a set of wavelength-dependent light curves at fixed phases $F_{\lambda,\mathrm{eff}}$ that are specific to that SN. The interpolation is controlled by a kernel K, which operates on $F_{\lambda,\mathrm{eff}}$ as shown in Equation (2). This equation is the same used to predict the mean in a Gaussian process, so we refer to this interpolation scheme as the GPMP:[23]

$$f_\lambda(t) = \boldsymbol{K}(t - t_0, \boldsymbol{t}_p)\boldsymbol{K}^{-1}(\boldsymbol{t}_p, \boldsymbol{t}_p)\boldsymbol{F}_{\lambda,\mathrm{eff}}. \quad (2)$$

These GPMP kernel matrices $\boldsymbol{K}$ specifically are calculated with a stationary $p = 2$ Matérn covariance function $C_{5/2}$ to ensure the interpolated curves are twice differentiable:

$$K_{ij} = C_{5/2}(|t_{p,i} - t_{p,j}|; \rho, \sigma^2). \quad (3)$$

$\boldsymbol{F}_{\lambda,\mathrm{eff}}$ is the light curve at wavelength λ on a grid of phase nodes; all $n_\lambda = 10$ light curves form the flux node matrix $\boldsymbol{F}_{\mathrm{eff}}$. $t_{p,i} \in \boldsymbol{t}_p$ is a vector indexing the model's n_p phase nodes and the per-SN parameter t_0 aligns said SN Ia's observations with the model's phase grid. Intuitively, the first kernel matrix $\boldsymbol{K}$ in Equation (2) maps each observation to our phase node space after said observation phase t_i is translated by t_0, while the second accounts for $\boldsymbol{F}_{\lambda,\mathrm{eff}}$ flux node covariance at grid phases $t_{p,i}$ and $t_{p,j}$. GPMPs provide a natural framework to translate observed phase t_i by the per-SN t_0 parameter to the model grid's phase zero-point.

Note that t_0 is not the fit date of maximum brightness—instead, $t - t_0$ aligns the observation phase with the model's phase grid $\boldsymbol{t}_p$. As we train the model using rest-frame transformed spectrophotometry, each t_0 is fit in its SN Ia's reference $z = 0.05$ frame. The kernel length scale hyperparameter ρ is fixed to match the phase node interval resolution of 4 days, although the model is insensitive to any reasonable choice in ρ (for example, $\rho \approx 1$ week). Furthermore, by fixing $\rho = 4$, the matrix $\boldsymbol{K}$ is calculated and inverted only once during sampling. The uncertainty hyperparameter σ^2 is fixed to unity since it is divided out in Equation (2).

The $\boldsymbol{F}_{\mathrm{eff}}$ of each SN is decomposed via element-wise multiplication (also called the Hadamard operation $\circ$) from a fiducial flux template matrix $\boldsymbol{F}_0$ and a warping matrix $\boldsymbol{\Omega}$:

$$\boldsymbol{F}_{\mathrm{eff}} = \boldsymbol{F}_0 \circ \boldsymbol{\Omega}. \quad (4)$$

$\boldsymbol{F}_0$ encodes the training sample's mean flux evolution via a set of fiducial light curves, while $\boldsymbol{\Omega}$ encodes deviations from these fiducial light curves for a given SN Ia. Specifically, each column $\boldsymbol{\Omega}_\lambda$ of matrix $\boldsymbol{\Omega}$, for a given λ node, is defined as:

$$\log(\boldsymbol{\Omega}_\lambda) = -0.4(\chi_0 + s_1\boldsymbol{M}_{\lambda,1} + c_1 L_{\lambda,1} + c_2 L_{\lambda,2}). \quad (5)$$

$\boldsymbol{M}_{\lambda,1}$ is the λ-node column of the phase-dependent chromatic flux variation template matrix $\boldsymbol{M}_1$. $\boldsymbol{M}_1$ therefore encodes training sample light-curve variation. $L_{\lambda,1}$ and $L_{\lambda,2}$ are the λ-

[23] The use of GPMPs over other techniques (e.g., spline interpolation) is partly historical. Our original intention was to treat model functions probabilistically, a use for which Gaussian processes would have been suited. We later opted to use the more conventional approach of using a deterministic function as is done in SALT2.

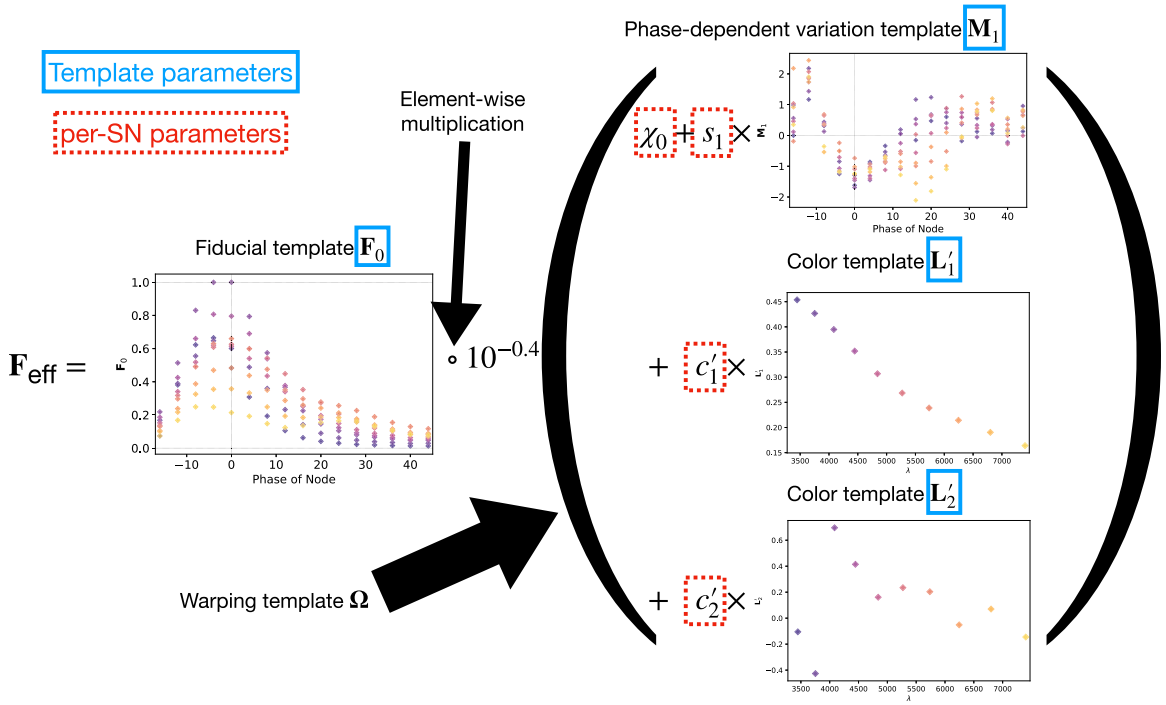

Figure 2. A schematic of our model's flux nodes. Each SN Ia has an effective flux node matrix $\boldsymbol{F}_{\text{eff}}$ that is an element-wise product of the sample's fiducial flux template $\boldsymbol{F}_0$ and a warping matrix $\boldsymbol{\Omega}$. This warping matrix includes the phase-dependent chromatic flux variation template $\boldsymbol{M}_1$, two phase-independent chromatic flux variation templates $\boldsymbol{L}_1'$ and $\boldsymbol{L}_2'$ (which make up the 2D phase-independent chromatic variation model), and per-SN parameters χ_0 (achromatic offset), s_1 (phase-dependent chromatic flux variation contribution), and c_1 and c_2 (phase-independent chromatic flux template contributions). Presented here are the $\boldsymbol{L}_1'$ and $\boldsymbol{L}_2'$ basis representation of the phase-independent templates (see Section 3.3 for more information). Per-band nodes are presented in this plot having the same figure color.

node elements of the two phase-independent chromatic flux variation template vectors $\boldsymbol{L}_1$ and $\boldsymbol{L}_2$, respectively (these appear as scalars in Equation (5) because of their phase independence). Each explicit per-SN parameter is contained within the warping template: the achromatic offset parameter χ_0, the phase-dependent chromatic flux variation parameter s_1, and the two phase-independent chromatic flux variation parameters c_1 and c_2. Both differences in intrinsic brightness and peculiar velocity effects are accounted for by the χ_0 parameters.

Figure 2 illustrates this model's architecture, specifically displaying the transformed phase-independent template basis $\boldsymbol{L}_1'$ and $\boldsymbol{L}_2'$ vectors later discussed in Section 3.3. Figure 3 is a directed acyclic graph of our model. Its conditional probability structure only connects observations to the model flux—deterministic connections (transformations and definitions) are presented with dashed arrows. Global template parameters are located in the top blue box and per-SN parameters in the bottom red box.

3.1. Template Constraints and Per-supernova Parameter Models

The third λ band, $t_{p,i} = 0$ phase node of $\boldsymbol{F}_0$ is fixed to 1—it is referred to here as fixed band 3 and corresponds to the wavelength node at 4084 Å, the top-hat filter band closest to a standard *B*-band. This constraint prevents any scaling degeneracy between $\boldsymbol{F}_0$ and the model's χ_0 parameter while also setting the specific phase node that the t_0 parameter aligns observations to. Physical consideration further requires all $\boldsymbol{F}_0$ flux node parameters be bound to greater than or equal 0, so we enforce nonnegative values for all flux values.

All chromatic flux variation templates $\boldsymbol{L}_1$, $\boldsymbol{L}_2$, and $\boldsymbol{M}_1$ have scaling degeneracies with their respective per-SN parameters c_1, c_2, and s_1. For example, the transformations $s_1 \rightarrow s_1/\alpha$ and $\boldsymbol{M}_1 \rightarrow \alpha \boldsymbol{M}_1$ leave the model unchanged; identical degeneracies exist for c_1–$\boldsymbol{L}_1$ and c_2–$\boldsymbol{L}_2$. Each scaling degeneracy is removed by requiring these three templates be normalized. This is a straightforward procedure for $\boldsymbol{L}_1$ and $\boldsymbol{L}_2$, where each is instantiated in `Stan` as unit vectors, but a more involved process is used to normalize the template matrix $\boldsymbol{M}_1$. We first define a unit vector of length $n_p \times n_\lambda$ that is then transformed into $\boldsymbol{M}_1$ by "chopping" said unit vector into n_λ column vectors (each of length n_p) that form the column space of a now normalized $\boldsymbol{M}_1$.[24] No further constraints or bounds are placed on the template parameters.

Zero-mean constraints are placed on the per-SN parameter sets c_1, c_2, and s_1. A reference SN Ia could be selected to serve as the c_1, c_2, and s_1 zero-points, but we opt instead to require these per-SN parameter sets to always have a mean of 0. These constraints are enforced structurally by instantiating these parameter sets as centered vectors (Appendix A).

3.2. Fitting the Model

As described in, e.g., C. Saunders et al. (2018) and D. Rubin et al. (2022), the SNfactory data are extracted from the 15 × 15 spaxels of the SNIFS (G. Aldering et al. 2002; B. Lantz et al. 2004) that are projected as spectra onto 2000 × 4000 CCD

[24] The unit vector constraints for $\boldsymbol{L}_1$, $\boldsymbol{L}_2$, and $\boldsymbol{M}_1$ are what prevent these templates from being described as color variation templates. If instead we constrained these templates at a reference wavelength node to be fixed to zero, then these templates would have units color, per the usual astronomical definition.

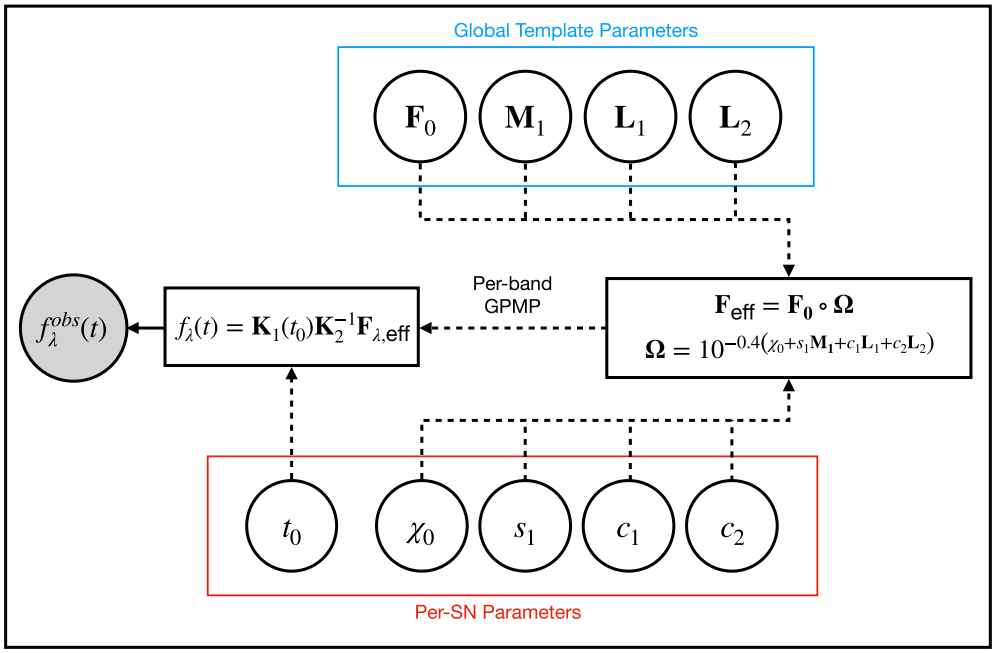

Figure 3. A directed acyclic graph representation of our model. Per-SN model parameters are located in the bottom red box; global template parameters are in the top blue box. The dashed arrows are deterministic relations (transformations and definitions). The only explicit conditional probability in the model's architecture relates observations $f_\lambda^{\text{obs}}(t)$ to modeled flux $f_\lambda(t)$. We use GPMP interpolation per band in mapping effective template nodes $\boldsymbol{F}_{\text{eff}}$ nodes to a predicted flux $f_\lambda(t)$.

detectors. In this extraction process the Poisson noise and the readout noise for each pixel are included with appropriate weights. The host galaxy is subtracted from this data cube using a reference data cube taken more than 1 yr later, as described in S. Bongard et al. (2011). Gray dimming by clouds is corrected as described in R. Pereira et al. (2013). The remaining SN-only light is fit with a point-spread function (PSF) model, as described in C. Buton et al. (2013) and D. Rubin et al. (2022), again propagating the Poisson and detector noise uncertainties. Due to small unmodeled PSF shape variations, there is additional uncertainty, which is empirically determined to be approximately 3% (P. F. Léget et al. 2020). The final calculated uncertainties are dominated either by this PSF shape noise or the Poisson plus detector noise.

Initial trial runs confirmed that the nominal uncertainties were underestimated, not having included galaxy-subtraction errors, and that the overall distribution had broader tails than a normal distribution. We thus added a further uncertainty equal to 2% of each SN Ia's maximum observed flux and treat measurement uncertainties as having a Cauchy distribution, which led to stable convergence of the fit. The credibility of this ansatz was checked though inspection of the pull distribution of photometric residuals around the final best fit. This approach has previously been used to represent model uncertainty in B. M. Rose et al. (2021).

The first source of error preferentially increases the high-flux observational errors, while the second largely affects low-flux observations, particularly those 20 or more days after peak brightness. All added uncertainties are diagonal: no covariance is injected into our data before training.

The nominal SNfactory measurement and uncertainty are used as the Cauchy distribution location and scale. For each flux observation $f_\lambda^{\text{obs}}(t)$ with its corresponding measurement uncertainty $\sigma_{f_\lambda(t)}$, our likelihood function takes the form:

$$f_\lambda^{\text{obs}}(t) \sim \text{Cauchy}[f_\lambda(t), \sigma_{f_\lambda(t)}]. \quad (6)$$

No correlations are added between measurement errors.

The model is implemented and trained using the statistics programming language `Stan` (B. Carpenter et al. 2017). Built into `Stan` is a No U-Turns (NUTS) Hamiltonian Monte Carlo (HMC) sampler well suited for sampling our model's high-dimensional posterior. No explicit priors are placed on the templates or per-SN parameter sets, instead leaving them with default implicit flat priors along any aforementioned bounds (Section 3.1).

`Stan` is informed with initial conditions estimated by first running simpler versions of the model. We do this only to improve sampling efficiency—it is not necessary for our model's convergence. This process is done iteratively, starting with the simplest model that only obtains the t_0, $\boldsymbol{F}_0$, and χ_0 parameters (a mean light-curve model). The results of this simplest model then become the initial conditions for a more complex model that includes the $\boldsymbol{L}_1$ and c_1 parameters. Other components (specifically, $\boldsymbol{L}_2$ and c_2, and then $\boldsymbol{M}_1$ and s_1) are then added and trained using the prior model iteration's fit as initial conditions until all the described model's components are incorporated. Note that we use SALT2's t_{max} as initial conditions for the t_0 parameters when training the simplest of these models.

With `Stan`'s default NUTS we pull 4000 samples for each of 16 types of instantiated samplers: 2000 warm-up followed by 2000 samples iterations per chain. `Stan` is run on the University of Pittsburgh's Computational Resource Center.[25] The convergence metrics "split-$\hat{R}$" and "effective sample size" were calculated after postprocessing using techniques provided in A. Vehtari et al. (2021). After training and postprocessing,

[25] https://crc.pitt.edu/

each SN Ia is refit with the template parameters fixed ($\boldsymbol{F}_0$, $\boldsymbol{M}_1$, $\boldsymbol{L}_1$, and $\boldsymbol{L}_2$) to determine the final values for the per-SN Ia χ_0, s_1, c_1, and c_2, permitting a direct comparison of these per-SN parameters against other empirical SN Ia models.

There is no selected standard $\Delta M_B = 0$ SN Ia identified in the training sample, leaving a nontrivial linear degeneracy between the achromatic offset parameter χ_0 and phase-independent parameters c_1 and c_2. For physical reasons, c_1 and c_2 should not correlate with the intrinsic magnitude, which ideally should only be captured by χ_0.

Linear transformations from P. F. Léget et al. (2020) are used to remove correlations between both the c_1 and χ_0 parameters, and the c_2 and χ_0 parameters, as summarized in Appendix B. Implementing this directly in the `Stan` model leaves the results unchanged but does reduce the sampling efficiency, so this step is performed after sampling.

3.3. Interpreting the Phase-independent Templates

Each sampler from `Stan` explores a plane spanned by the template vectors $\boldsymbol{L}_1$ and $\boldsymbol{L}_2$.[26] Even after decorrelating χ_0 from c_1 and c_2, the output basis $\{\boldsymbol{L}_1, \boldsymbol{L}_2\}$ is not unique, a consequence of this model's physics-agnostic architecture. This is because for any nonsingular linear transformation the output basis vectors (i.e., $a\boldsymbol{L}_1 + b\boldsymbol{L}_2$ and $c\boldsymbol{L}_1 + d\boldsymbol{L}_2$ for $a, b, c, d \in \mathcal{R}$ and $ab - cd \neq 0$) necessarily span the aforementioned plane. To quantify this plane's convergence (as opposed to only its basis vectors), for each posterior sample we calculate a bivector $\mathcal{L} = \boldsymbol{L}_1 \wedge \boldsymbol{L}_2$ that, by definition, spans the plane of interest. Importantly, the bivector representation $\mathcal{L}$ is no longer ambiguous.

A bivector is a geometric object representing an oriented plane element constructed from the wedge product (see Figure 4 for an illustration). Intuitively, a bivector corresponds to a plane like a vector corresponds to a line, and the wedge product is the dual to a cross product in 3D. Unlike the cross product, the wedge product generalizes to any finite-dimensional vector space greater than 2, meaning bivectors are well-defined in this model's 10D wavelength node space. Now any SN Ia'a phase-independent chromatic flux variation curve $\boldsymbol{c} = \boldsymbol{c}_1\boldsymbol{L}_1 + \boldsymbol{c}_2\boldsymbol{L}_2$ can be interpreted as residing in the plane spanned by $\mathcal{L}$, regardless of the selected $\boldsymbol{L}_1$, $\boldsymbol{L}_2$ basis.

Each component of the bivector $\mathcal{L}$ is calculated as follows:

$$\mathcal{L}_{ij} = \frac{L_{1,i}L_{2,j} - L_{1,j}L_{2,i}}{\sqrt{\sum_{k=1}^{n_\lambda}\sum_{m>k}^{n_\lambda}(L_{1,k}L_{2,m} - L_{1,m}L_{2,k})}}, \quad (7)$$

where $\hat{L}_{1,i}$ is the ith wavelength component of template $\boldsymbol{L}_1$ and $\hat{L}_{2,i}$ is the ith wavelength component of template $\boldsymbol{L}_2$. Importantly, this representation is independent of the $\boldsymbol{L}_1$, $\boldsymbol{L}_2$ choice. Note that $\mathcal{L}$ is normalized so as to represent a unit plane element. With these transformed parameters $\mathcal{L}_{ij}$, the model can unambiguously be tested for convergence and the best-fit templates be determined.

[26] In general, a plane is an affine space, not a vector space. All planes described in this paper do pass through the embedding vector space's origin, which ensures they are proper 2D vector subspaces. Because of this, all planes discussed either intersect or are parallel. For brevity, in this paper we refer to any 2D subspaces as planes.

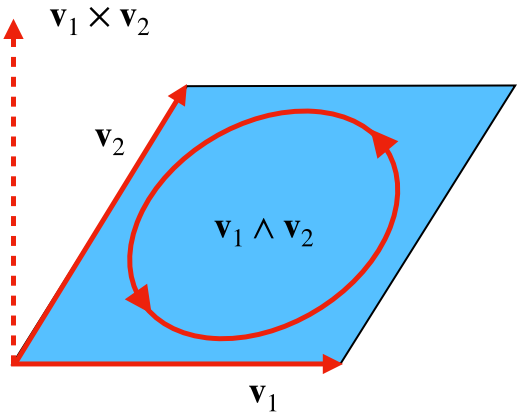

Figure 4. This is an illustration of a bivector (the blue parallelogram) $\boldsymbol{v}_1 \wedge \boldsymbol{v}_2$ constructed by the vectors $\boldsymbol{v}_1$ and $\boldsymbol{v}_2$ (the red vectors) in 3D. A 3D space allows for the corresponding cross product to be included for reference (the red dashed vector). Bivectors, like vectors, are oriented objects, with the bivector $\boldsymbol{v}_1 \wedge \boldsymbol{v}_2$ having a counterclockwise orientation determined by component vector ordering (here represented with an oriented red loop). Reversing the product order reverses a bivector's "circulation" or orientation: $\boldsymbol{v}_2 \wedge \boldsymbol{v}_1 = -\boldsymbol{v}_1 \wedge \boldsymbol{v}_2$. Note that the cross product does not generalize to all finite-dimensional vector spaces, while the wedge product $\wedge$ does.

3.4. Bases for the Phase-independent Chromatic Variation Model

We now seek a pair of vectors, $\boldsymbol{L}_1$ and $\boldsymbol{L}_2$, that span $\mathcal{L}$ and readily provide insight into the physical origin of the model's 2D phase-independent chromatic flux model. Two such bases are considered.

3.4.1. Maximum Variance Ratio Basis

The first basis, called the maximum variance ratio (MVR)[27] basis, is derived directly from the corresponding c_1–c_2 distribution.

New and uncorrelated parameter sets c_1^{mvr} and c_2^{mvr} with their relative variances maximized are found via a linear transformation. The result is a basis for $\mathcal{L}$ where $\boldsymbol{L}_1^{\mathrm{mvr}}$ accounts for the most chromatic flux diversity by a single template in $\mathcal{L}$, while $\boldsymbol{L}_2^{\mathrm{mvr}}$ captures any remaining variation. This basis amounts to the assumption that two independent physical effects affect chromatic flux variation (e.g., the amount of dust or intrinsic SN Ia diversity) while designating as much of the variance as possible to one source (e.g., dust). This assumption provides some useful insights even if it does not totally satisfy our physical expectations, as it does not consider additional possible effects (e.g., the kind of dust), nor that SNe and their progenitor environments likely correlate with said effects. This basis's solution is found by numeric minimization; it is not orthogonal.

The MVR basis is determined as follows. A basis centered at the origin can be described by the angle between the unit vectors and their orientation. Starting with an orthogonalized output basis from `Stan` $\boldsymbol{L} = [\boldsymbol{L}_1^{\mathrm{orth}}, \boldsymbol{L}_2^{\mathrm{orth}}]$ for $\mathcal{L}$ and a per-SN coefficient matrix $\boldsymbol{c}$ (here arranged as an $n_{\mathrm{sn}} \times 2$ matrix), adding an extra angle θ between the basis is achieved with the

[27] We developed this technique and are not aware of any previous equivalent use.

transformation:

$$\tilde{\boldsymbol{L}} = \boldsymbol{M}\boldsymbol{L} = \begin{pmatrix} 1 & 0 \\ \sin\theta & \cos\theta \end{pmatrix}\boldsymbol{L}. \quad (8)$$

In this $\tilde{\boldsymbol{L}}$ basis, the per-SN coefficients $\boldsymbol{c}' = \boldsymbol{c}\boldsymbol{M}^{-1}$ have an orientation given by V^T from the singular value decomposition (SVD) of $\boldsymbol{c}' = U\Sigma V^T$. Taking V^{TM} and normalizing its rows to be unit vectors gives the properly oriented basis given θ. Note that these primed components here are unrelated to those introduced below.

For this basis, the ellipticity of its corresponding c distribution is again found using SVD, given by $\log\Sigma_{11} - \log\Sigma_{22}$. An optimizer is used to determine the θ that maximizes the ellipticity.

3.4.2. Cardelli–Clayton–Mathis-derived Basis

We also desire a basis that readily separates rapidly changing chromatic flux variation (i.e., that akin to SN absorption/emission features) from continuum-like chromatic flux variation (i.e., dust-like behavior). This continuum-like variation is assumed to be dust-like, given dust extinction's ubiquitous contribution to SN Ia color/chromatic flux variation. The two phase-independent plots found in Figure 2 provide an example of what we want from this basis—specifically, one template being more smooth with respect to wavelength than the other.

For the second basis, the vector $\boldsymbol{L}_1^{'}$ simultaneously resides within the planes spanned by both $\mathcal{L}$ and the Cardelli–Clayton–Mathis (CCM89) dust model (J. A. Cardelli et al. 1989), defining the first basis component. This intersection ensures it captures smooth, CCM89-like chromatic variability. The other basis vector $\boldsymbol{L}_2^{'}$ is chosen to be perpendicular to $\boldsymbol{L}_1^{'}$, while still residing in the $\mathcal{L}$ plane; this basis is orthogonal by construction. Note that $\boldsymbol{L}_2^{'}$ will not necessarily be perpendicular to the plane spanned by CCM89, but is still guaranteed to provide the least continuum-like variability with respect to (w.r.t.) wavelength (and therefore, the most spectral-feature-like behavior) allowed by $\mathcal{L}$. This basis provides a useful representation not because it recovers a mathematically valid dust extinction curve, but instead because it clearly separates rapidly changing chromatic flux variation from SN continuum-like variability. It is important to remember that intrinsic variability could still affect the direction of the first basis vector $\boldsymbol{L}_1^{'}$, which means this basis does not guarantee a physical decomposition into exclusive dust and intrinsic components. As such, any dust-like properties inferred from $\boldsymbol{L}_1$ in isolation are physically ambiguous.

The CCM89-derived basis is calculated as follows. Since CCM89 has two basis curves $a(\lambda)$ and $b(\lambda)$ (one for each parameter A_V and A_V/R_V, respectively), one can construct a CCM89 unit bivector $\mathcal{L}_{\rm ccm}$ from discretized curves $a(\lambda) \rightarrow \boldsymbol{a}$ and $b(\lambda) \rightarrow \boldsymbol{b}$ using an appropriately modified version of Equation (7): $\hat{L}_{1,i} \rightarrow a_i$ and $\hat{L}_{2,i} \rightarrow b_i$. The two planes spanned by $\mathcal{L}$ and $\mathcal{L}_{\rm ccm}$ then intersect within the n_λ-dimensional wavelength vector space along a line. It is the vector which spans this intersecting line that defines the new $\boldsymbol{L}_1^{'}$ template:

$$\boldsymbol{L}_1^{'} = \text{Intersection}[\mathcal{L}, \mathcal{L}_{\rm ccm}]. \quad (9)$$

We are free to choose a new $\boldsymbol{L}_2$ as long as it resides within the subspace represented by $\mathcal{L}$. To minimize this $\boldsymbol{L}_2$ template's dust-like properties, $\boldsymbol{L}_2^{'}$ is defined as a $\theta = \pi/2$ radian rotation of $\boldsymbol{L}_1^{'}$ within the plane spanned by $\mathcal{L}$ via a rotation operator $\boldsymbol{R}_{\mathcal{L}}(\theta/2)$:

$$\boldsymbol{L}_2^{'} = \boldsymbol{R}_{\mathcal{L}}(\pi/4)\boldsymbol{L}_1^{'}\boldsymbol{R}_{\mathcal{L}}^{-1}(\pi/4). \quad (10)$$

This rotation maximizes the component of $\boldsymbol{L}_2^{'}$ that is perpendicular to the plane spanned by $\mathcal{L}_{\rm ccm}$. The new $\{\boldsymbol{L}_1^{'}, \boldsymbol{L}_2^{'}\}$ also transforms the c_1 and c_2 parameters sets, here labeled $c_1 \rightarrow c_1'$ and $c_2 \rightarrow c_2'$.

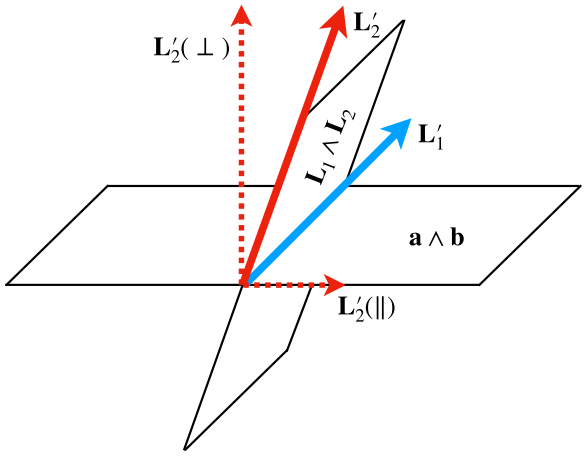

Figure 5. A 3D representation the CCM89 basis's geometric intuition. The blue solid vector is the transformed $\boldsymbol{L}_1^{'}$ template spanning the intersection of two planes, one spanned by the model bivector $\mathcal{L} = \boldsymbol{L}_1 \wedge \boldsymbol{L}_2$ and the other spanned by the CCM89 bivector $\mathcal{L}_{\rm ccm} = \boldsymbol{a} \wedge \boldsymbol{b}$. The red solid vector is a $\pi/2$ rotation in the plane spanned by $\mathcal{L}$ that defines the transformed $\boldsymbol{L}_2^{'}$ template. The decomposition $\boldsymbol{L}_2^{'} = \boldsymbol{L}_2^{'}(\|) + \boldsymbol{L}_2^{'}(\perp)$ with respect to the CCM89 plane spanned by $\mathcal{L}_{\rm ccm}$ is given with the red dashed arrows.

Figure 5 provides a 3D view of the geometric intuition involved in finding the CCM89-derived basis. All calculations discussed here are implemented using geometric algebra, which provides a novel approach to study oriented subspaces. Geometric algebraic implementations for calculating intersections, projections, and rotation operations are summarized in Appendix C. Although we exploit geometric algebra's elegance and interpretability to perform all said operations, each operation could be done using more classical linear algebraic techniques if desired.

3.5. Handling Similar Solutions

In training, the samplers converge to one of three sets of very similar solutions, as is apparent in the per-chain values of $\boldsymbol{L}_1$ shown in Figure 17. We find that simultaneously obtaining global template and per-SN parameters leaves the resulting samplers sensitive to the three least representative SNe Ia in the training sample (based on residuals). These groups are associated with slightly different solutions for these three SNe Ia, as seen by eye and quantitatively through their χ_0 solutions. Consider the following analogy: consider the inertia tensor of a space station as being the model's global templates and the coordinates of its crew members corresponding to the per-SN parameters. The moment of inertia changes only slightly if crew members move to another location on the station (assuming the space station is much more massive than the crew's total weight). Similarly, it is this "change in position" of per-SN solutions that is causing a very subtle change in template parameters, preventing complete

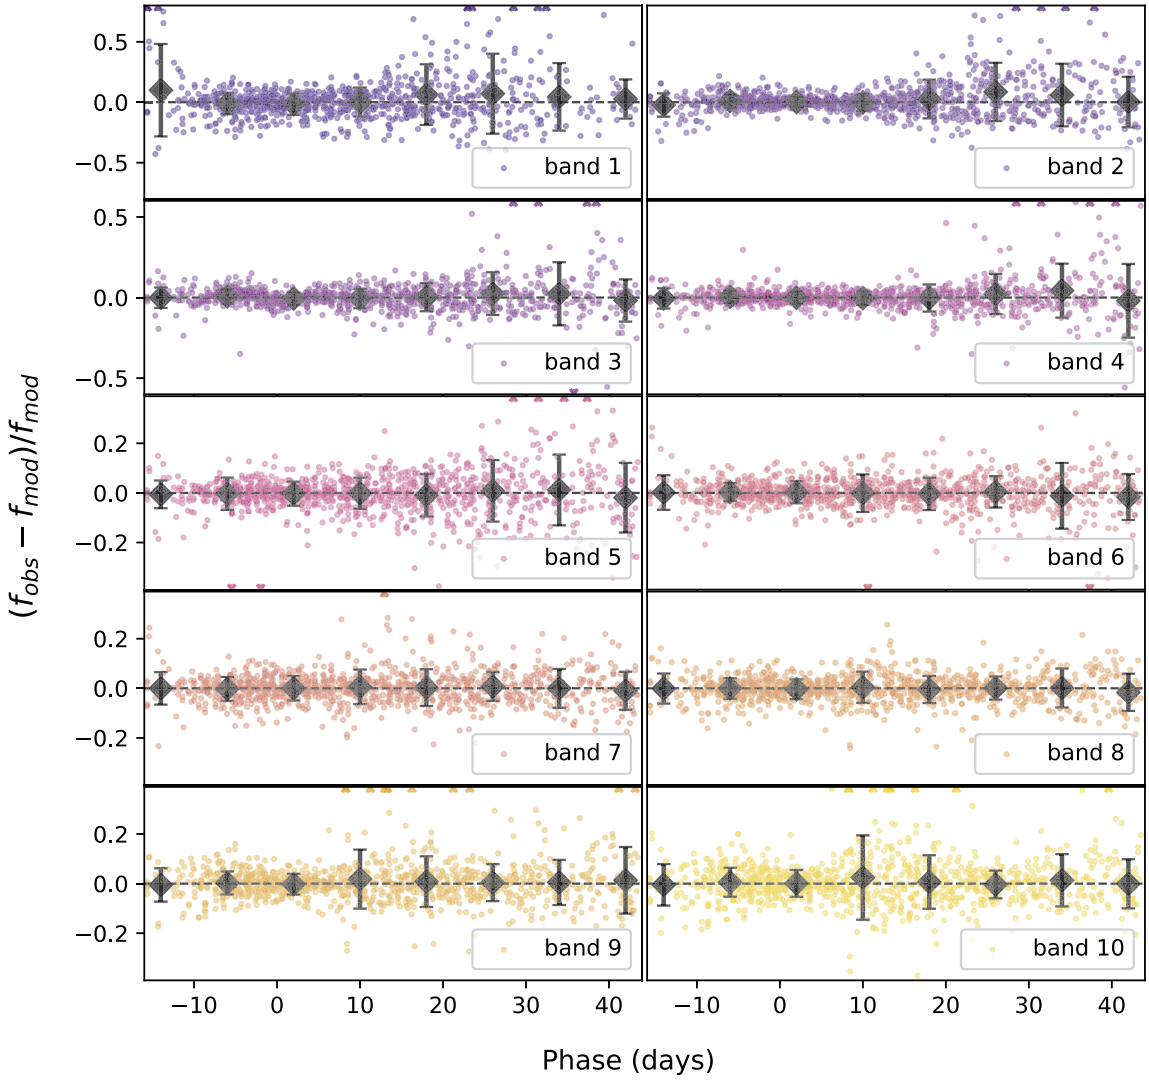

Figure 6. Best-fit model residuals with respect to observations presented for each of our 10 bands. Eight day binned averages for each band are presented as black diamonds, with error bars being binned standard deviations. The carets at the top and bottom edges represent points that lie outside the range of the plots.

convergence. Obviously this is not a fair comparison, but the resulting effects are very similar, hence this comparison.

The worst performing SN Ia PTF11mkx consistently sees a 1σ difference between χ_0 values between different chain groups—considering that this χ_0 parameter fit uncertainty is only $\approx$1%, this tail wagging is very subtle. Indeed, when refitting per-SN parameters with one of the group solutions as a fixed global template (Section 4.4), all chains converge to the same solutions for the three SNe Ia. Similarly, if two different group solutions are separately held constant and refit individually, the two resulting per-SN parameter sets are indistinguishable. Because each group's solutions are consistently very similar, we opt to use the weighted average of all solutions for the best-fit template. This decision to use an average of all groups, as opposed to using a specific solution, has no impact on the remaining analysis.

Note that refitting with a fixed t_0 parameter does not prevent this subtle grouping, nor does cutting these three SNe Ia from the training set. If these three aforementioned troublesome SNe Ia are removed, the next few SNe Ia that are the least representative of this newly trimmed sample start "wagging." This instability may be a feature of our model. In particular, the flat priors we use for the c_1 and c_2 distributions may bias the results for the least representative SNe Ia of a given training sample. Considering that global template parameters again were effectively unchanged with their removal, we retain those three SNe Ia for the remainder of the analysis.

4. Best-fit Model Results

Our model, as implemented with `Stan`, consists of global template parameters and per-SN parameters. We define the best-fit solution to be the mode of a 365D template parameter space,[28] with each per-SN parameter marginalized before estimating from HMC sampling the posterior's maximum. This mode is estimated using a mean shift clustering algorithm implemented in the scikit-learn package using the default flat kernel (F. Pedregosa et al. 2011). To estimate a consistent best-fit solution, bivector components $\boldsymbol{L} = \boldsymbol{L}_1 \wedge \boldsymbol{L}_2$ (Equation 7) are used instead of the $\boldsymbol{L}_1$ and $\boldsymbol{L}_2$ components directly. This process is analogous to maximum a posteriori estimation of the HMC-sampled posterior, but allows first for the aforementioned

[28] 159 from $\boldsymbol{F}_0$, 159 from $\boldsymbol{M}_1$, nine from $\boldsymbol{L}_1$, and nine from $\boldsymbol{L}_2$, specifically.

postprocessing. The marginal posterior dispersion for each parameter is presented as 68th percentile error bars. The residuals of the best-fit model, alongside binned averages, are presented in Figure 6.

Two of the 16 HMC samplers were rejected after fitting because (1) their resulting mean t_0 parameter values were systematically inconsistent by 4 to 5 days with SALT2 date of maximum, 2) the corresponding light-curve shapes were nonphysical, and (3) these two samplers have notably inconsistent $\ln p$ values relative to the 14 retained samplers. The 14 remaining samplers converge to three very similar group solutions of which the weighted average is taken; see Section 3.5 for more details.

Table 1 provides the median values and 68th percentile upper and lower values for the per-SN light-curve model parameters.

4.1. Phase-independent Chromatic Flux Variation Templates

Ostensibly, if $\mathcal{L}$ were to only capture CCM89 dust-like behavior, then the planes spanned by $\mathcal{L}$ and $\mathcal{L}_{\rm ccm}$ would be effectively parallel and any intersection poorly constrained. This turns out to not be the case, with the best-fit solution recovering a planar separation angle of about 80°.

4.1.1. Maximum Variance Ratio Basis

Figure 7 presents the MVR basis as described in Section 3.4. Qualitatively, $\boldsymbol{L}_1^{\rm mvr}$ appears nominally more dust-like than its counterpart $\boldsymbol{L}_2^{\rm mvr}$. Figure 7 includes the best-fit CCM89 curve $R_V^{\rm mvr} = 2.18$ for reference, showcasing that its most extreme divergence from CCM89 is blueward of 5000 Å. $\boldsymbol{L}_2^{\rm mvr}$, on the other hand, captures variation that is not readily describable as dust-like: normal dust extinction should be absorptive across the optical wavelength range whereas the sign flip in $\boldsymbol{L}_2^{\rm mvr}$ produces simultaneous brightening/dimming on either side of 4000 Å. Although the degree of variation increases as wavelength decreases, there is a distinct flip in behavior around 4000 Å. Such behavior is inconsistent with dust extinction.

The conditions for the target parameter set distributions are to assign maximum variance to one component while keeping the second component uncorrelated. These conditions yield a basis consistent with the expectation that dust-like variation is the primary contributor to SN Ia phase-independent chromatic flux variation while intrinsic SN Ia diversity uncorrelated with dust accounts for additional variability.

4.1.2. Cardelli–Clayton–Mathis-derived Basis

The template $\boldsymbol{L}_1'$ is presented in the top plot of Figure 8. As summarized in Section 3.4, the intersection of the plane spanned by $\mathcal{L}$ with the plane spanned by $\mathcal{L}_{\rm ccm}$ defines the first phase-independent chromatic flux template $\boldsymbol{L}_1'$. This $\boldsymbol{L}_1'$ template, as expected, captures continuum-like variation akin to dust extinction.

From $\boldsymbol{L}_1'$ we find an intersection total-to-selective extinction ratio of $R_V^{\rm int} = 2.4$. As mentioned in 3.4, this R_V is not immediately physically interpretable without an A_V model.

Also presented in the top plot of Figure 8 is template $\boldsymbol{L}_2'$. Unlike $\boldsymbol{L}_1'$, $\boldsymbol{L}_2'$ captures the more rapidly changing flux variability allowable by $\mathcal{L}$. Specifically, $\boldsymbol{L}_2$ is capturing wavelength variation at scales smaller than that expected by continuum dust variability, at least within the optical regime. Indeed, its features nominally align with spectral features. Also note the similarity between $\boldsymbol{L}_2'$ and $\boldsymbol{L}_2^{\rm mvr}$ despite their vastly different constructions.

$\boldsymbol{L}_2'$ is not perpendicular to the CCM89 plane because the plane spanned by $\mathcal{L}$ is itself not perpendicular. As such, $\boldsymbol{L}_2'$ also captures a dust-like component alongside its dominating spectral-like component. The bottom panel presents this template's decomposition into parallel and perpendicular components defined with respect to CCM89's plane for reference.

4.1.3. Intrinsic Variation

Both of the best-fit template representations $\boldsymbol{L}_2'$ and $\boldsymbol{L}_2^{\rm mvr}$ capture pronounced phase-independent chromatic flux variation blueward of 4500 Å that is inconsistent with dust (see Figure 9). Indeed, no extrinsic phenomena readily describe this behavior. Variation blueward of 4500 Å includes the prominent Ca II H and K feature and its Si III counterpart, Si II λ4130, C II λ4267, Fe II λ4404, and Mg II λ4481. With our choice of splitting the spectral range into $n_\lambda = 10$ synthetic filters, the model cannot completely distinguish between said features, although at least one node for both $\boldsymbol{L}_2'$ and $\boldsymbol{L}_2^{\rm mvr}$ and their corresponding variation seemingly align with the wavelength bin where Ca II H and K is the dominant contributor to a strong spectral feature; Si II and Fe II features also contribute to this variation. Hints of chromatic flux variability ostensibly aligns with the SNe Ia signature Si II λ6347 feature and O I λ7774, which are visible in Figure 8, but in practice do not affect any resulting model flux predictions (again, see Figure 9).

We use the same data set used by C. Saunders et al. (2018) for their SNEMO analysis. Comparing both $\boldsymbol{L}_2'$ and $\boldsymbol{L}_2^{\rm mvr}$ to SNEMO2 and SNEMO7 eigenvectors (with two and seven components, respectively) yields indecisive insight, though. Figure 6 from C. Saunders et al. (2018) shows SNEMO eigenvectors describing similar $\boldsymbol{L}_2'$ or $\boldsymbol{L}_2^{\rm mvr}$ behavior at maximum, but all of these eigenvectors are clearly phase dependent—no eigenvector's evolution seems approximately phase independent. Figures 9 through 12 from C. Saunders et al. (2018) present some unexplained variation for SNEMO2 blueward of 4500 Å that is nominally phase independent up to 6 days postmaximum; this all but disappears with SNEMO7. That SNEMO2, itself similar to SALT2, sees unexplained phase-independent variability around maximum that aligns with the $\boldsymbol{L}_2'$ and $\boldsymbol{L}_2^{\rm mvr}$ features seems indicative of our model's performance, but such a claim taken alone is likely an excessive interpretation.

SNfactory data are also used by K. Boone et al. (2021b) with their Twins Embedding nonlinear model. More interesting insight is gained in comparing their findings with our template representations $\boldsymbol{L}_2'$ and $\boldsymbol{L}_2^{\rm mvr}$, since Twins Embedding is currently a phase-independent, maximum-phase model. Blueward of 4500 Å, spectral variation recovered by Twins Embedding loosely aligns with the both $\boldsymbol{L}_2'$ and $\boldsymbol{L}_2^{\rm mvr}$ templates (see Figures 4, 6, and 10 from K. Boone et al. 2021a for reference). Recovering this consistent variation in our phase-independent template, albeit at lower wavelength bin resolution, lends credibility that $\mathcal{L}$ as a whole is capturing intrinsic variation.

Past analyses by D. Branch et al. (1993) and A. G. Riess et al. (1998b) find Ca II H and K features are relatively stable in

Table 1
Model Per-supernova Parameter Median Values with 68th Percentile Upper and Lower Values

SN	$\chi_0 - \langle\chi_0\rangle$	s_1	c_1	c_2	t_0 (MJD)
SNF 20060512-001	-0.04 ± 0.03	-0.02 ± 0.01	-0.33 ± 0.09	0.12 ± 0.03	53882.5 ± 0.12
SNF 20060526-003	-0.06 ± 0.03	-0.07 ± 0.01	-0.31 ± 0.09	-0.04 ± 0.02	53893.4 ± 0.21
SNF 20061020-000	0.30 ± 0.03	0.21 ± 0.01	0.44 ± 0.11	0.08 ± 0.03	54037.9 ± 0.19
SNF 20061021-003	-0.08 ± 0.03	-0.05 ± 0.01	0.38 ± 0.08	-0.04 ± 0.02	54039.7 ± 0.13
SNF 20061022-005	-0.11 ± 0.03	-0.08 ± 0.02	-0.34 ± 0.10	0.15 ± 0.03	54040.8 ± 0.17
SNF 20061111-002	0.08 ± 0.03	-0.01 ± 0.01	-0.30 ± 0.12	-0.05 ± 0.03	54060.8 ± 0.18
SNF 20070424-003	0.08 ± 0.03	-0.05 ± 0.01	-0.26 ± 0.09	-0.07 ± 0.02	54225.8 ± 0.16
SNF 20070506-006	-0.19 ± 0.03	-0.07 ± 0.01	-0.94 ± 0.08	0.09 ± 0.02	54245.9 ± 0.10
SNF 20070630-006	-0.09 ± 0.02	-0.04 ± 0.01	-0.16 ± 0.08	-0.06 ± 0.02	54294.9 ± 0.11
SNF 20070717-003	-0.12 ± 0.02	0.03 ± 0.01	0.90 ± 0.08	-0.12 ± 0.02	54310.3 ± 0.11
SNF 20070802-000	-0.19 ± 0.03	-0.03 ± 0.01	1.00 ± 0.09	-0.15 ± 0.02	54326.7 ± 0.13
SNF 20070818-001	-0.24 ± 0.03	-0.04 ± 0.01	0.80 ± 0.10	-0.16 ± 0.01	54339.0 ± 0.18
SNF 20070831-015	-0.28 ± 0.03	-0.15 ± 0.01	-0.23 ± 0.11	-0.03 ± 0.03	54352.4 ± 0.26
SNF 20070902-018	0.11 ± 0.03	0.08 ± 0.01	0.66 ± 0.10	0.04 ± 0.03	54353.2 ± 0.17
SNF 20080507-000	-0.30 ± 0.02	-0.06 ± 0.01	0.44 ± 0.08	0.04 ± 0.02	54601.7 ± 0.16
SNF 20080510-001	-0.01 ± 0.03	-0.05 ± 0.01	-0.55 ± 0.08	-0.05 ± 0.02	54607.5 ± 0.12
SNF 20080510-005	-0.00 ± 0.03	-0.11 ± 0.01	-0.75 ± 0.11	-0.02 ± 0.03	54608.1 ± 0.18
SNF 20080512-010	0.25 ± 0.03	0.16 ± 0.02	-0.24 ± 0.09	0.12 ± 0.02	54603.9 ± 0.12
SNF 20080514-002	0.50 ± 0.02	0.20 ± 0.01	-0.73 ± 0.05	0.14 ± 0.01	54613.6 ± 0.06
SNF 20080522-000	-0.27 ± 0.02	-0.08 ± 0.01	-0.87 ± 0.08	0.20 ± 0.02	54624.3 ± 0.10
SNF 20080522-011	-0.11 ± 0.02	-0.05 ± 0.01	-1.02 ± 0.07	-0.05 ± 0.02	54619.1 ± 0.12
SNF 20080531-000	0.08 ± 0.02	0.05 ± 0.01	0.03 ± 0.07	-0.08 ± 0.02	54625.9 ± 0.11
SNF 20080614-010	0.24 ± 0.03	0.22 ± 0.01	0.55 ± 0.10	0.08 ± 0.03	54637.7 ± 0.10
SNF 20080620-000	0.13 ± 0.03	0.09 ± 0.01	0.64 ± 0.10	-0.03 ± 0.02	54642.7 ± 0.16
SNF 20080626-002	-0.16 ± 0.02	-0.08 ± 0.01	-0.47 ± 0.08	-0.08 ± 0.02	54654.0 ± 0.15
SNF 20080714-008	-0.17 ± 0.02	0.00 ± 0.01	1.68 ± 0.07	-0.20 ± 0.02	54670.0 ± 0.15
SNF 20080725-004	-0.07 ± 0.03	-0.10 ± 0.01	0.05 ± 0.08	-0.10 ± 0.02	54680.8 ± 0.20
SNF 20080803-000	-0.09 ± 0.03	-0.05 ± 0.01	0.88 ± 0.10	-0.04 ± 0.03	54693.1 ± 0.13
SNF 20080810-001	0.17 ± 0.02	0.14 ± 0.01	0.10 ± 0.07	0.11 ± 0.02	54700.9 ± 0.08
SNF 20080821-000	-0.13 ± 0.02	-0.06 ± 0.01	-0.23 ± 0.08	0.01 ± 0.02	54710.1 ± 0.15
SNF 20080825-010	-0.04 ± 0.03	0.11 ± 0.01	-0.20 ± 0.08	0.08 ± 0.02	54715.1 ± 0.10
SNF 20080909-030	-0.15 ± 0.05	-0.04 ± 0.01	-0.29 ± 0.17	0.05 ± 0.05	54733.8 ± 0.16
SNF 20080913-031	-0.19 ± 0.07	0.00 ± 0.02	0.55 ± 0.25	-0.09 ± 0.07	54734.6 ± 0.40
SNF 20080918-000	-0.15 ± 0.03	-0.11 ± 0.01	1.10 ± 0.11	-0.09 ± 0.03	54736.8 ± 0.17
CSS130502_01	0.29 ± 0.03	-0.01 ± 0.01	-1.23 ± 0.08	0.07 ± 0.02	56425.0 ± 0.14
LSQ12fhe	-0.45 ± 0.03	-0.07 ± 0.01	-0.41 ± 0.12	0.22 ± 0.03	56214.9 ± 0.19
LSQ14cnm	-0.00 ± 0.03	-0.00 ± 0.01	-0.88 ± 0.08	-0.03 ± 0.02	56833.9 ± 0.14
PTF09dlc	-0.00 ± 0.02	-0.03 ± 0.01	-0.56 ± 0.07	-0.08 ± 0.02	55077.4 ± 0.13
PTF09dnl	-0.23 ± 0.02	-0.07 ± 0.01	0.26 ± 0.06	-0.05 ± 0.02	55077.7 ± 0.16
PTF09fox	0.01 ± 0.02	-0.08 ± 0.01	-0.83 ± 0.07	-0.05 ± 0.01	55135.8 ± 0.14
PTF09foz	-0.06 ± 0.02	0.11 ± 0.01	0.48 ± 0.07	-0.00 ± 0.02	55133.8 ± 0.11
PTF10hmv	0.11 ± 0.02	-0.09 ± 0.01	1.04 ± 0.06	-0.01 ± 0.02	55355.3 ± 0.11
PTF10icb	0.49 ± 0.02	0.01 ± 0.01	0.32 ± 0.07	0.03 ± 0.02	55363.7 ± 0.09
PTF10mwb	0.21 ± 0.02	0.10 ± 0.01	0.19 ± 0.06	-0.02 ± 0.02	55392.6 ± 0.08
PTF10nlg	0.03 ± 0.02	-0.03 ± 0.01	1.67 ± 0.07	-0.20 ± 0.01	55394.8 ± 0.16
PTF10qyz	0.36 ± 0.03	0.10 ± 0.01	-0.43 ± 0.10	-0.08 ± 0.02	55429.0 ± 0.12
PTF10wnm	-0.00 ± 0.02	-0.09 ± 0.01	-0.66 ± 0.08	0.03 ± 0.02	55480.5 ± 0.17
PTF10wof	0.02 ± 0.03	-0.06 ± 0.01	0.12 ± 0.10	-0.12 ± 0.02	55476.7 ± 0.28
PTF11bgv	0.22 ± 0.03	0.03 ± 0.01	0.30 ± 0.08	0.14 ± 0.02	55645.6 ± 0.12
PTF11bju	-0.23 ± 0.03	-0.05 ± 0.01	-0.14 ± 0.09	0.26 ± 0.02	55652.7 ± 0.14
PTF11bnx	-0.00 ± 0.03	0.03 ± 0.01	1.44 ± 0.09	-0.16 ± 0.02	55655.6 ± 0.16
PTF11mkx	-0.50 ± 0.04	-0.00 ± 0.01	-0.33 ± 0.13	0.27 ± 0.02	55837.8 ± 0.10
PTF11mty	-0.04 ± 0.02	-0.07 ± 0.01	-0.82 ± 0.08	0.02 ± 0.02	55840.2 ± 0.12
PTF12dxm	0.34 ± 0.03	0.22 ± 0.01	0.89 ± 0.09	0.02 ± 0.02	56054.8 ± 0.09
PTF12fuu	0.18 ± 0.02	-0.05 ± 0.01	-451.28 ± 0.07	0.01 ± 0.02	56115.2 ± 0.09
PTF12grk	0.10 ± 0.03	0.12 ± 0.01	0.77 ± 0.10	-0.16 ± 0.01	56137.0 ± 0.10
PTF12iiq	0.07 ± 0.04	0.12 ± 0.01	1.17 ± 0.13	-0.17 ± 0.02	56183.6 ± 0.16
PTF13anh	-0.09 ± 0.03	0.05 ± 0.01	0.51 ± 0.12	-0.15 ± 0.01	56416.1 ± 0.16
SN 2004ef	0.07 ± 0.02	0.14 ± 0.01	1.28 ± 0.06	-0.08 ± 0.02	53265.0 ± 0.08
SN 2005hc	0.00 ± 0.03	-0.07 ± 0.01	-0.61 ± 0.09	-0.03 ± 0.02	53670.5 ± 0.18
SN 2005hj	-0.23 ± 0.03	-0.13 ± 0.01	-0.57 ± 0.09	0.11 ± 0.02	53677.9 ± 0.23
SN 2006cj	0.02 ± 0.02	-0.06 ± 0.01	-0.84 ± 0.06	0.04 ± 0.02	53882.4 ± 0.11
SN 2007bd	0.16 ± 0.03	0.03 ± 0.02	-0.58 ± 0.10	-0.02 ± 0.02	54207.6 ± 0.16
SN 2010dt	0.03 ± 0.03	0.00 ± 0.01	-0.44 ± 0.09	-0.10 ± 0.01	55363.4 ± 0.11

Table 1
(Continued)

SN	$\chi_0 - \langle\chi_0\rangle$	s_1	c_1	c_2	t_0 (MJD)
SN 2011bc	−0.05 ± 0.02	0.04 ± 0.01	1.18 ± 0.08	−0.05 ± 0.02	55669.6 ± 0.15
SN 2011be	0.22 ± 0.03	−0−0.03 ± 0.01	−1.20 ± 0.08	0.08 ± 0.02	55656.9 ± 0.12
SN 2011ga	0.05 ± 0.02	0.02 ± 0.01	−0.43 ± 0.08	−0.03 ± 0.02	55832.7 ± 0.13
SN 2011gf	0.07 ± 0.03	−0.07 ± 0.01	−1.54 ± 0.07	−0.04 ± 0.01	55833.5 ± 0.12
SN 2011hr	−0.25 ± 0.03	−0.03 ± 0.01	0.12 ± 0.08	0.25 ± 0.02	55891.0 ± 0.10
SN 2012cg	0.63 ± 0.02	0.00 ± 0.01	0.67 ± 0.07	0.05 ± 0.02	56084.5 ± 0.10
SN 2012dn	−0.08 ± 0.03	−0.03 ± 0.01	0.23 ± 0.09	0.20 ± 0.02	56135.1 ± 0.06
SN 2012fr	−0.56 ± 0.02	−0.09 ± 0.01	−0.60 ± 0.05	−0.09 ± 0.01	56247.1 ± 0.05
SN 2013bt	0.40 ± 0.04	0.22 ± 0.01	−0.28 ± 0.11	0.12 ± 0.03	56411.4 ± 0.11

Note. Per-SN parameters $\chi_0 - \langle\chi_0\rangle$, s_1, c_1, and c_2 have been transformed to magnitudes. The parameter t_0 is the date corresponding to maximum-phase node of the light-curve model for the third band centered at 4084 Å.

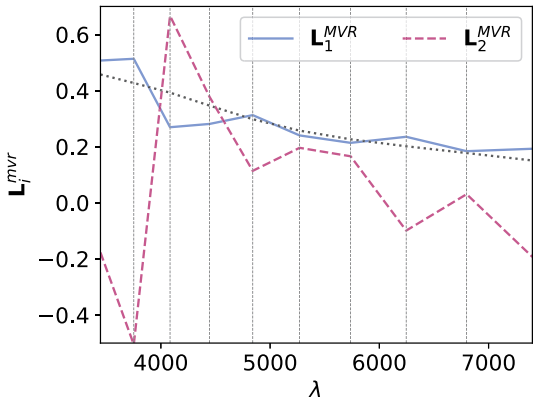

Figure 7. The blue solid line corresponds to the first MVR component $\boldsymbol{L}_1^{\rm mvr}$, which appears nominally more consistent with dust-like variation than its counterpart $\boldsymbol{L}_2^{\rm mvr}$, given as the magenta dashed. $\boldsymbol{L}_1^{\rm mvr}$ has a best-fit $R_V^{\rm mvr} = 2.18$ given as the gray dotted line, with most of its divergence from a CCM89 curve occurring blueward of 5000 Å. $\boldsymbol{L}_2^{\rm mvr}$ captures variation not readily describable as dust-like.

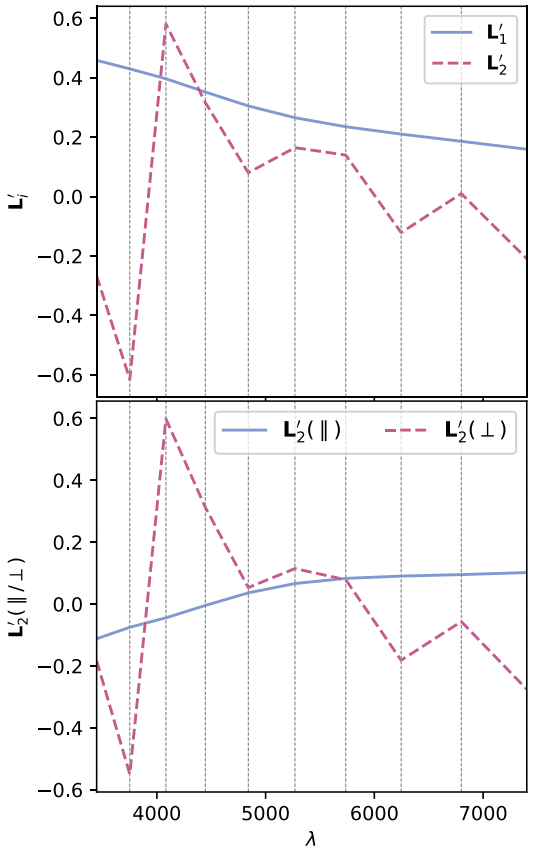

Figure 8. The top plot presents phase-independent chromatic flux variation templates $\boldsymbol{L}_1'$ and $\boldsymbol{L}_2'$. $\boldsymbol{L}_1'$ has a recovered total-to-selective extinction of $R_V^{\rm int} = 2.4$. The bottom plot presents a decomposition of $\boldsymbol{L}_2'$ into its parallel and perpendicular components with respect to the CCM89 plane. $\boldsymbol{L}_2'$ clearly captures some dust-like variability, despite being dominated by intrinsic modes. Although the low-resolution wavelength binning prevents quantification of spectral features, the most impressive $\boldsymbol{L}_2'$ variability appears in the Ca II H and K regime.

the week before and weeks after peak *B*-band brightness, with this effective phase independence being sufficiently stable to exploit for the latter's "snapshot" methodology from which they constrain luminosity distances. Both $\mathcal{L}$ basis representations do capture intrinsic, phase-independent variation around Ca II H and K, but this alone is not profound evidence of fundamental Ca II H and K time independence in the SN Ia population. SN Ia Ca II H and K features, as with all spectral features, demonstrably evolve with time. In its current form, this model cannot distinguish between phase-averaged spectral variation or truly phase-independent intrinsic variation. Indeed, a goal of this project is to demonstrate that intrinsic chromatic flux variation can "leak" into phase-independent components, something that is occurring in these results.

4.2. Fiducial Template

Figure 10 presents GPMP interpolations of each band's best-fit $\boldsymbol{F}_{\lambda,0}$ nodes as solid curves. The fixed band 3 template is the third solid curve presented in the top plot of Figure 10. The ubiquitous NIR bump is recovered for redder bands (bottom plot of Figure 10), with this second maximum occurring ≈25 days after our fixed-band peak brightness as expected (S. W. Jha et al. 2019). Apart from the reddest template curve centered at 7401 Å, the peak brightness phase per band occurs earlier for bluer bands and later for redder bands, again consistent with established trends (S. W. Jha et al. 2019). As would be expected by its *I*-band overlap, the reddest template curve exhibits somewhat more complex behavior than the other curves, such as the inflection point between its two local maxima (bottom plot, Figure 10).

Note that each fiducial template light curve's peak brightness phase does not align with our $t_{p,i} = 0$ flux node, meaning t_0

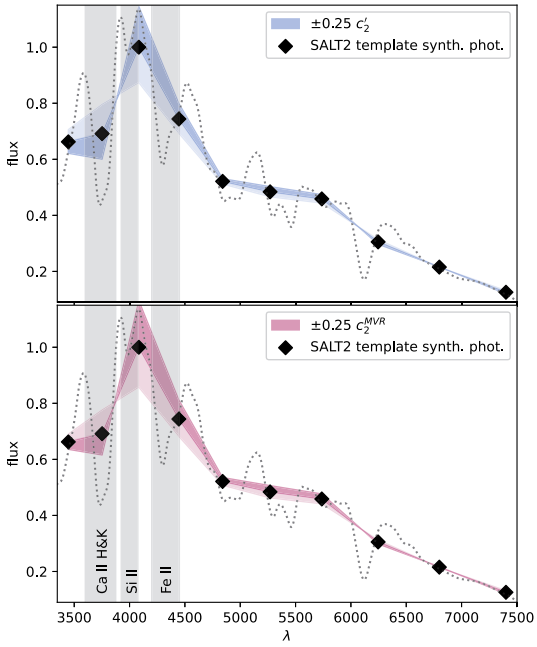

Figure 9. This figure demonstrates $a \pm 0.2$ mag c_2 variation (blue for the CCM89 basis $c_2{}'$, maroon for c_2^{MVR}) of $\boldsymbol{L}_2^{'}$ overlaid on SALT2's mean template $t = 0$ phase spectrum (dashed black line). The spectrum is binned via synthetic photometry with top-hat filters, presented as black diamonds. Flux units are normalized by synthetic photometry wavelength 4048 Å value to unity and example spectral features blueward of 4500 Å are presented for reference.

should not be interpreted as the fixed band 3's peak brightness phase. This is ultimately inconsequential, requiring only that each light curve's peak brightness phase be calculated deterministically after fitting, and has no effect on this analysis or its conclusions.

4.3. Phase-dependent Chromatic Flux Variation Template

As shown in Figure 10, the best-fit phase-dependent variation template $\boldsymbol{M}_1$ exhibits stretch-like behavior across all bands. The shaded regions in these plots show phase-dependent light-curve variation from $-0.09 < s_1 < 0.09$ mag, which approximately captures the dispersion of the fit s_1 parameter set (the set's standard deviation is 0.09 mag). The lighter shaded regions correspond to positive s_1 values, while the darker correspond to negative values. As s_1 increases (decreases), the $\boldsymbol{F}_{\mathrm{eff}}$ node values increases (decreases) with respect to $\boldsymbol{F}_0$'s node values, resulting in each GPMP interpolation curve's global maximum decreasing (increasing). This change in $\boldsymbol{F}_{\mathrm{eff}}$ node scaling is offset for by a change in χ_0, correlating the χ_0 and s_1 parameter sets (see Section 4.4).

The sign of $\boldsymbol{M}_1$ template's contribution is a function of phase: for each curve there are two phases where the $\boldsymbol{M}_1$ template's contribution reverses in sign. For positive (negative) s_1, the result is a narrowing (broadening) of the effective flux curve. The phase and degree of this broadening varies between bands in a manner consistent with D. Kasen & S. E. Woosley (2007), being more extreme for bluer wavelengths. Furthermore, Figure 11 plots our model's $\Delta m_3(15)$ as a function of s_1, demonstrating our $\boldsymbol{M}_1$ template indeed recovers stretch-like behavior for this model's B-band analog fixed band 3.

For redder bands, stretch-like behavior is convolved with NIR bump variation (bottom plot of Figure 10). A 3D mesh plot (Figure 12) best illustrates these two modes of NIR

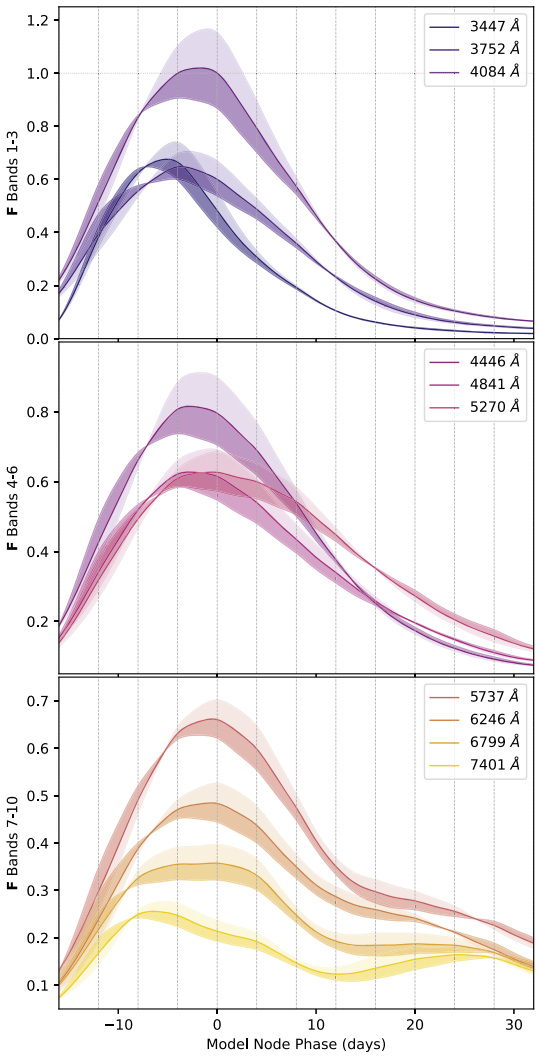

Figure 10. A visualization of ±0.09 mag variation in s_1 on the model's fiducial flux template $\boldsymbol{F}_0$, as warped by the phase-dependent chromatic flux template $\boldsymbol{M}_1$. Positive s_1 contribution is given by light shaded regions, while negative s_1 contribution is given by the dark shaded regions. Solid lines are the GPMP-interpolated light curve for that band's fiducial template nodes, which is deterministic (not stochastic) in our model. The top two plots illustrate recovered stretch-like behavior by the template $\boldsymbol{M}_1$, with broadening to narrowing of the effective light curve as s_1 increases in value. The bottom plot captures stretch-like behavior further convolved with NIR bump variational modes (bump location and size). Note that these figures are not portraying model uncertainty, only model response to s_1 parameter variation.

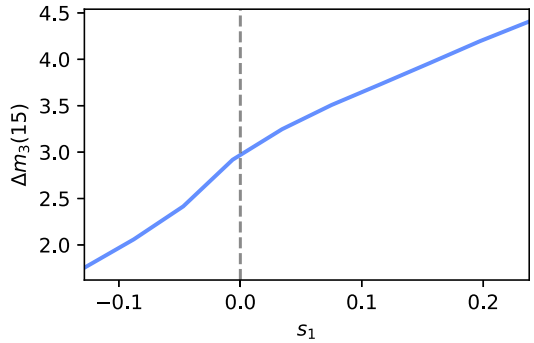

Figure 11. The model's $\Delta m_3(15)$ (the $\Delta m_B(15)$ analog for the fixed band 3) as a function of s_1 calculated for the fixed band 3 along the training sample's obtained s_1 value range.

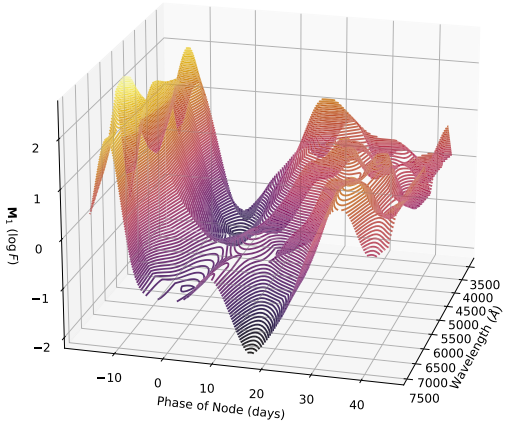

Figure 12. A contoured 3D view of our phase-dependent variation template—this is our model's equivalent to SALT2's M_1 stretch template. The valley-like structure corresponds to stretch-like behavior extracted by our $\boldsymbol{M}_1$ template.

variation: bump depth and bump location. As expected, these variation modes also correlate with stretch (C. R. Burns et al. 2011; S. Dhawan et al. 2015), with stretch appearing as the valley-like feature in Figure 12.

For fixed band 3, the phase of maximum brightness relative to our zero-phase node is a function of s_1. This movement in maximum brightness location, made clear in Figure 10, ranges from $+1$ day for our most negative $s_1 = -0.15$ mag SN Ia to -3 days for our most positive $s_1 = 0.22$ mag. Again, this has no effect on our analysis, requiring only an a posteriori calculation at the phase of maximum brightness if desired.

4.4. Per-supernova Results

Each SN Ia is refit with template parameters fixed to the previously discussed best-fit solution. Scatterplots comparing parameter sets include Spearman rank correlation coefficient (SCC) calculations alongside corresponding p-values. For color, only the CCM89-derived basis parameter sets c_1' and c_2' are presented. As the MVR and CCM89-derived bases yield parallel $\mathcal{L}$ planes, we choose to present results from the latter since the CCM89-derived basis is more readily interpreted—its first component is, by definition, a mathematically valid CCM89 curve. Note that we do not compare per-SN parameter sets from different basis decompositions.

In Figure 13, per-SN parameter sets for c_1', c_2', and s_1 are compared with their corresponding χ_0 values. The measured SCC value of 0.59 between the training sample's s_1 and χ_0 parameter sets results from a varying s_1 changing the resulting $\boldsymbol{F}_{\rm eff}$ scaling, requiring a compensating change in χ_0 to offset (see Section 4.3). By construction, the c_1' and c_2' sets are decorrelated with χ_0 (see Section 3.3). Higher rank correlations are recovered for c_1' versus s_1 and c_1' versus c_2' compared to c_2' versus s_1, as seen in the scatterplots of Figure 14.

We also quantify the fractional variance of the c_1' and c_2' bivariate distributions not explained by CCM89-like behavior. Each SN Ia phase-independent chromatic flux variation vector $\boldsymbol{c} = c_1'\boldsymbol{L}_1' + c_2'\boldsymbol{L}_2'$ is first normalized. The perpendicular component with respect to the CCM89 plane spanned by $\mathcal{L}^{\rm ccm}$ of each normalized $\boldsymbol{c}$ is then calculated via a projection operation (see Appendix C.2). This resulting distribution has a median value of 0.13 with 68th percentiles [0.05, 0.4] and provides a measure of our sample's fractional variance attributable to captured phase-independent chromatic

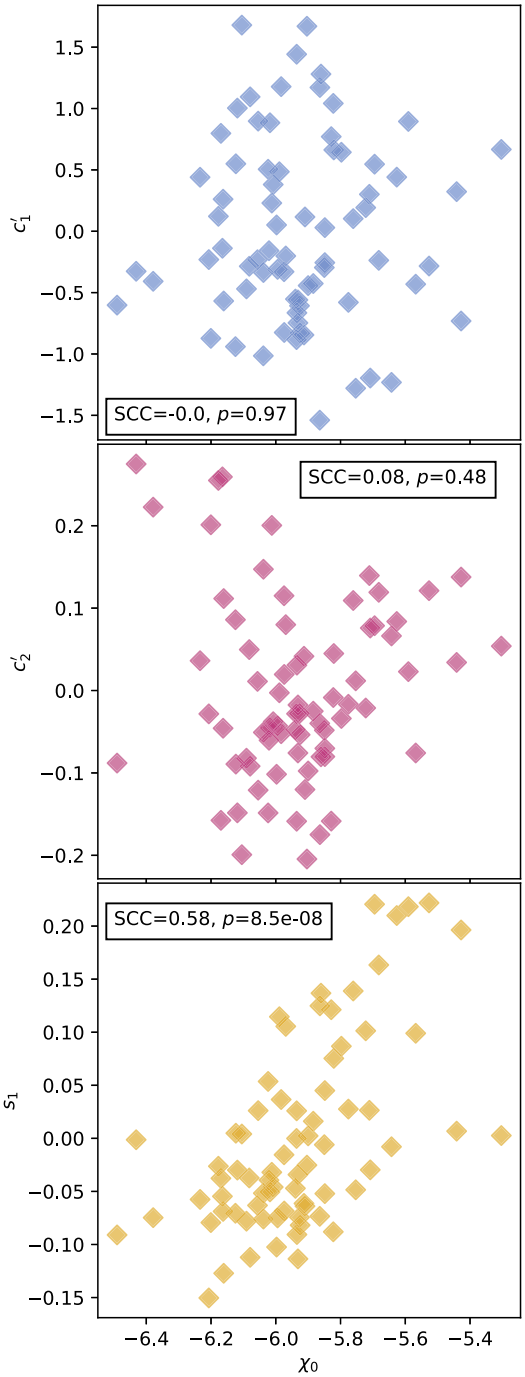

Figure 13. Comparison of fit c_1' (top, blue), c_2' (middle, magenta), and s_1 (bottom, yellow) samples against our χ_0 samples. The correlation between χ_0 and s_1 arises from s_1's changing of $\boldsymbol{F}_{\rm eff}$'s scale, which is then compensated for by a change in χ_0. There are no correlations between either c_1' and χ_0 or c_2' and χ_0 by construction.

variability which is not dust-like. Unsurprisingly, dust-like variation, which explains the remaining ≈87% variance, dominates the captured phase-independent variability. Even if this dust-like variation was the exclusive result of actual dust extinction (no "leaking" of intrinsic variability into dust-like behavior), the remaining ≈13% variance, which instead arises from intrinsic variability in the sample, is not negligible. $\mathcal{L}$, with $\boldsymbol{L}_2'$ in particular, is capturing a discernible addition of SN Ia variation over past two-component models (i.e., SALT2).

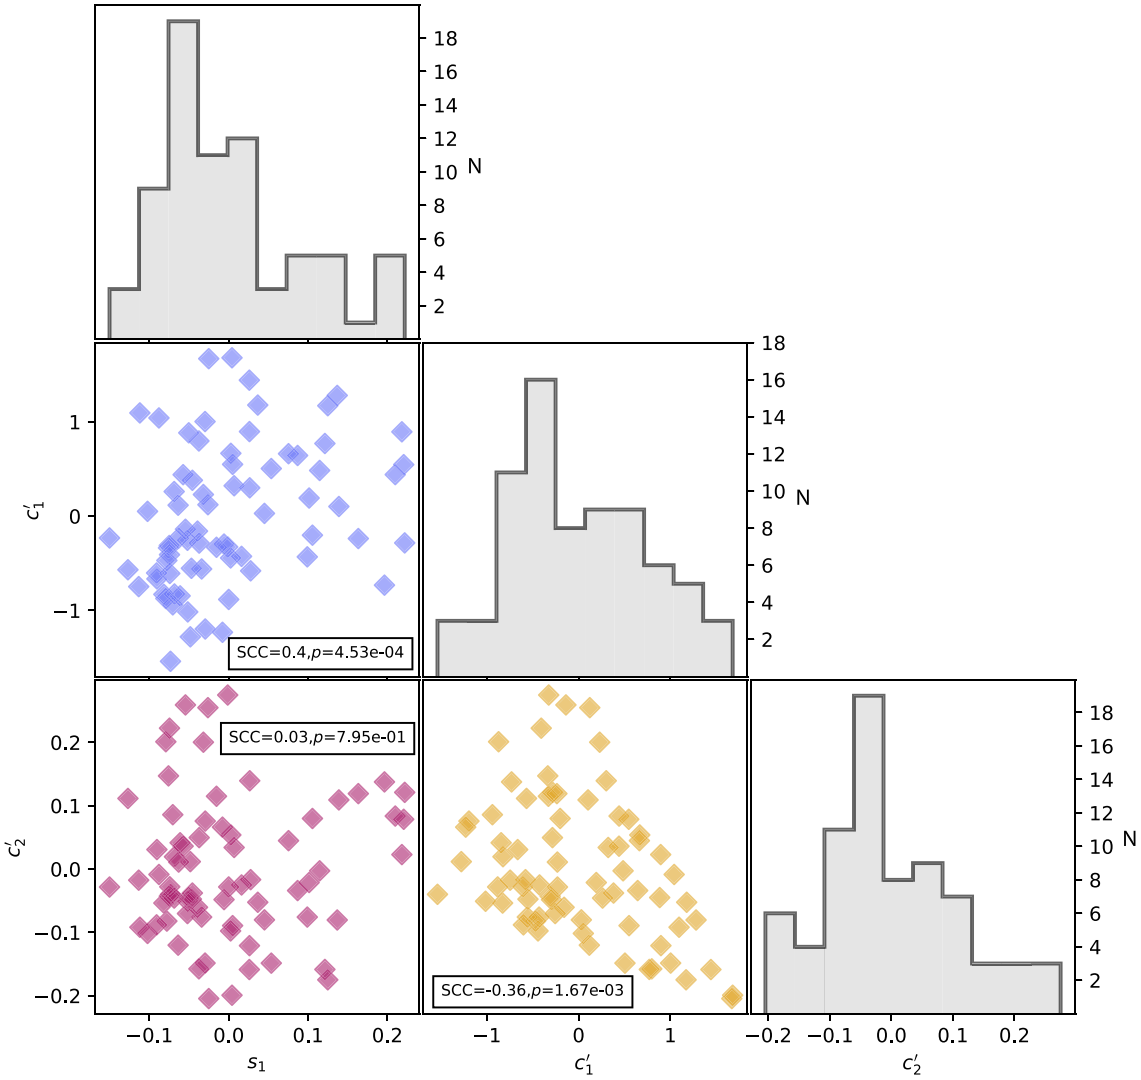

Figure 14. Corner plot for the per-SN parameters s_1, $c_1{}'$, and $c_2{}'$. Magenta points compare s_1 and $c_1{}'$, yellow points compare $c_2{}'$ and $c_1{}'$, and blue points compare $c_2{}'$ and $c_1{}'$ parameter sets. We measure only marginal rank correlations between both $c_1{}'$ vs. s_1 and $c_1{}'$ vs. $c_2{}'$. N is the number of SNe Ia.

4.4.1. Comparison to SALT2

This new SN Ia model and SALT2 are trained using optical wavelength observations, with neither making assumptions about dust extinction, making SALT2 an obvious comparator. One technical difference is our model's accounting for phase-dependent variability with a multiplicative variation template as opposed to SALT2's flux variation linear component, which is obviously additive in flux space. Nonetheless, s_1 and x_1 should correlate. This is the case as seen in Figure 15, with a rank correlation of -0.89 between x_1 and s_1. As presented in Section 4.3, this model's $\boldsymbol{M}_1$ templates obtains stretch-like behavior, just as SALT2's first-order variation template $M_1(t,\ \lambda)$ does.

The phase-independent chromatic flux variation template basis $\{\boldsymbol{L}_1^{'}, \boldsymbol{L}_1^{'}\}$ is selected without consideration of SALT2's $CL(\lambda)$ phase-independent chromatic variation model. As such, any correlations between $c_1{}'$ or $c_2{}'$ and SALT2 c are nontrivial—as seen in the top plot of Figure 16, only SALT2 c and $c_1{}'$ are correlated with a rank correlation of 0.78. Considering $\boldsymbol{L}_1^{'}$ is the maximal CCM89 dust-like vector allowed, that SALT2 c and $c_1{}'$ are strongly correlated is because SALT2's $CL(\lambda)$ template predominately captures dust-like variation. Indeed, the latest SALT3 recovers a $CL(\lambda)$ curve that is consistent with SALT2 and similarly aligns with CCM89 between 4000 and 7000 Å (W. D. Kenworthy et al. 2021). As seen in Figure 13 of W. D. Kenworthy et al. (2021), both of the SALT2 and SALT3 $CL(\lambda)$ templates begin to diverge from CCM89 near where $\boldsymbol{L}_2^{'}$ starts to exhibit most of its variability.

The bottom plot of Figure 16 demonstrates that the $c_2{}'$ and SALT2 c parameter sets are uncorrelated. As one would hope, the presented model's two-component phase-independent chromatic variation model captures SN Ia variation beyond that of SALT2. For reference, we provide best-fit linear relationships between SALT2 x_1 and s_1, and between SALT2 c and $c_1{}'$:

$$c_1{}'(c) = 8.74(\pm 0.88)c + 0.094(\pm 0.056), \tag{11}$$

$$s_1(x_1) = -0.089(\pm 0.004)x_1 + 0.0089(\pm 0.0035). \tag{12}$$

5. Conclusion

In this article, we introduce a new empirical SN Ia linear model that expands beyond similar analyses by introducing a second phase-independent chromatic flux variation template into its architecture. Our model converges to a solution with three variation components: one phase-dependent chromatic

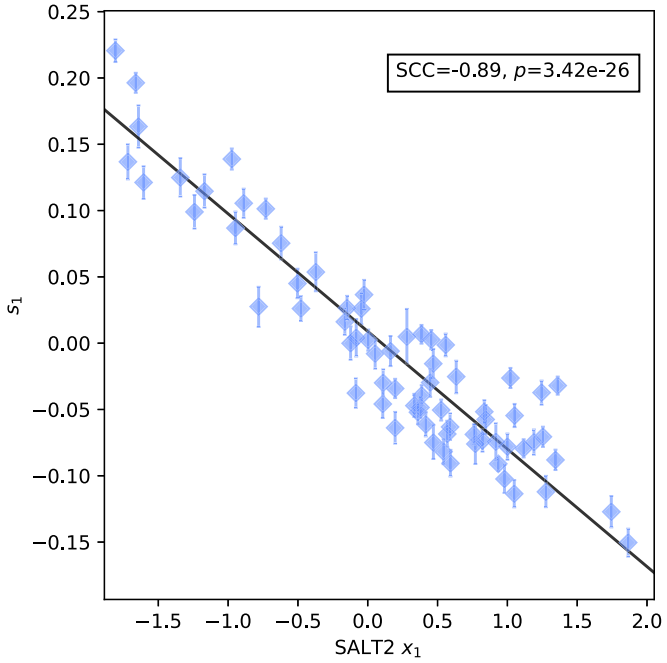

Figure 15. Per-SN comparison of our stretch parameter s_1 vs. SALT2's stretch proxy x_1. An ordinary least squares linear best fit is provided with a solid black line. Error bars correspond to 68th percentiles.

flux variation template $\boldsymbol{M}_1$ and two phase-independent chromatic flux variation templates $\boldsymbol{L}_1$ and $\boldsymbol{L}_2$, which are represented together by the bivector $\mathcal{L}$. The phase-dependent model component recovers stretch and NIR bump variation; it along with the phase-independent templates are consistent with SALT2's two model components.

The model's extended phase-independent architecture captures a nuanced combination of variability. Analysis of the 2D phase-independent chromatic flux variation plane spanned by $\mathcal{L}$ is done two bases: one that maximizes variation ratios (MVR) between the c_1 and c_2 parameter sets, and another that is defined with respect to the CCM89 dust plane. The MVR approach recovers nominally dust-like variation for its dominant component $\boldsymbol{L}_1^{\rm mvr}$, with a fit $R_V^{\rm MVR} = 2.18$, while $\boldsymbol{L}_2^{\rm mvr}$ trace potentially low-resolution intrinsic features. The CCM89 basis yields similar, albeit somewhat more physically interpretable, results. Specifically, this basis naturally decomposes $\mathcal{L}$ into one template capturing continuum-like wavelength variation ($\boldsymbol{L}_1^{'}$) and another capturing rapidly changing chromatic flux variation w.r.t. wavelength ($\boldsymbol{L}_2^{'}$). Component $\boldsymbol{L}_1^{'}$ is defined as the intersection of our model's phase-independent chromatic flux variation plane with CCM89 dust-like behavior. $\boldsymbol{L}_2^{'}$ is orthogonal to $\boldsymbol{L}_1^{'}$ and can contain phase-averaged intrinsic spectral variation not related to the $\boldsymbol{M}_1$ template; by construction, it can also contain residual CCM89-like variation.

Despite very different construction methods, $\boldsymbol{L}_2^{'}$ and $\boldsymbol{L}_2^{\rm mvr}$ are remarkably similar. As derived from the CCM89 basis, the fractional variance of intrinsic spectral variation for the 2D phase-independent chromatic flux variation plane (with 68th percentiles) is $0.13^{+0.4}_{-0.05}$—dust-like variation dominates $\boldsymbol{L}_1^{'}$ and $\boldsymbol{L}_2^{'}$ in combination, but residual spectroscopic-feature-like variation is not negligible. Despite this model's coarse $n_\lambda = 10$ wavelength bin count, we find that the features in $\boldsymbol{L}_2^{'}$ and $\boldsymbol{L}_2^{\rm mvr}$ both align with known SN Ia spectral features. Because of intrinsic variation "leaking" into these two phase-independent components, we make no strong conclusion about the recovered dust extinction curve from either $\boldsymbol{L}_1^{'}$ or $\boldsymbol{L}_1^{\rm mvr}$.

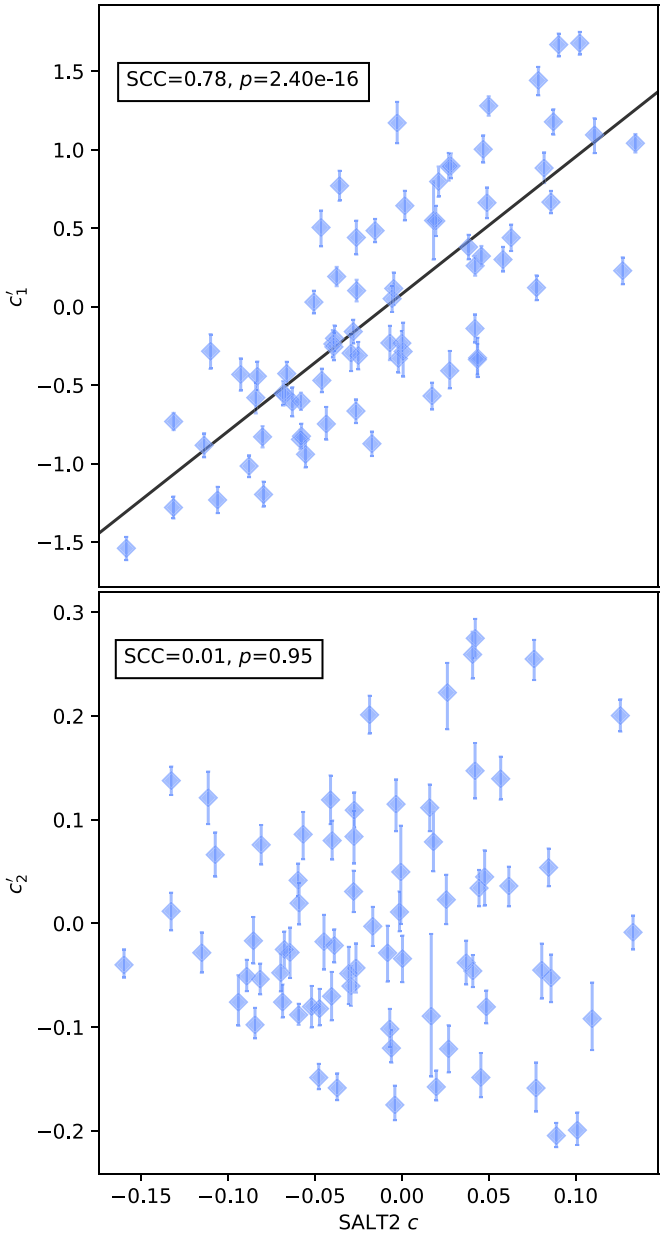

Figure 16. A per-SN comparison of our per-SN chromatic flux variation parameters c_1' (top) and c_2' (bottom) against SALT2's c parameter. We measure a clear rank anticorrelation between c_1' and SALT2 c, but measure no correlation between c_2' and SALT2 c. We interpret this as template $\boldsymbol{L}_2'$ capturing chromatic flux variation not modeled by SALT2. An ordinary least squares linear best fit between c_1' and SALT2 c is provided by a solid black line. Error bars correspond to 68th percentiles.

In our approach of SN standardization, we avoid setting conditions on the properties of unextinguished SNe or of dust, and hence are unable to directly make conclusions about foreground dust, apart from recovering a model component whose behavior is like that of dust (specifically, CCM89). Either forward modeling via additional postprocessing, or directly integrating a dust model into an improved version of our SN Ia model are options to consider—such augmentations will be the focus of work in the future. Such future work will also increase synthetic photometry spectral resolution to allow for better identification of spectral features.

Acknowledgments

The authors thank Prof. Michael Wood-Vasey for his continuous feedback about model implementation and interpretation of results. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists, Office of Science Graduate Student Research (SCGSR) 2018 S2 program. The SCGSR program is administered by the Oak Ridge Institute for Science and Education for the DOE under contract number DE-SC0014664. This work is also supported in part by the US Department of Energy Office of Science under DE-SC0007914. This research is supported in part by the University of Pittsburgh Center for Research Computing through the resources provided.

This work was supported in part by the Director, Office of Science, Office of High Energy Physics of the US Department of Energy under contract No. DE-AC02-05CH11231. Support in France was provided by CNRS/IN2P3, CNRS/INSU, and PNC and French state funds managed by the National Research Agency within the Investissements d'Avenir program under grant reference numbers ANR-10-LABX-0066, ANR-11-IDEX-0004-02, and ANR-11-IDEX-0007. Additional support comes from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 759194-USNAC). Support in Germany was provided by DFG through TRR33 "The Dark Universe" and by DLR through grants FKZ 50OR1503 and FKZ 50OR1602. In China support was provided by Tsinghua University 985 grant and NSFC grant No. 11173017. Some results were obtained using resources and support from the National Energy Research Scientific Computing Center, supported by the Director, Office of Science, Office of Advanced Scientific Computing Research of the US Department of Energy under contract No. DE-AC02-05CH11231. We also thank the Gordon & Betty Moore Foundation for their support.

M.R. has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 759194-USNAC).

The first author extends their gratitude to Prof. Wood-Vasey and Dr. Kim for their constant, unconditional, and unwavering support during a period of health struggles. This project would not have come to fruition without their mentorship, constant advocating, and kindness.

Appendix A
Centered Vectors for Type Ia Supernova Parameters

A centered vector is a vector $\boldsymbol{v}$ with the mean of its components being zero. There are numerous ways to structurally enforce centered vectors in `Stan`.[29] For performance reasons we opt for a shifted simplex method.

A simplex $\boldsymbol{\Delta}$ of dimension n is a vector with the constraint that its components sum to one: $\sum_n^i \Delta_i = 1$. Dividing the sum of the simplex by its component count n gives that simplex's mean: $\langle \boldsymbol{\Delta} \rangle = 1/n$. Therefore, any simplex $\boldsymbol{\Delta}$ instantiated in `Stan` can be transformed via translation by a factor $-1/n$ into a centered simplex $\boldsymbol{\Delta}^{\langle \mathbf{0} \rangle}$ with a mean of zero:

$$\Delta_i^{\langle 0 \rangle} = \Delta_i - 1/n. \tag{A1}$$

Finally, any centered vector $\boldsymbol{r}$ with component r_i can be defined as the product of a centered simplex $\boldsymbol{\Delta}^{\langle \mathbf{0} \rangle}$ and a scaling parameter r:

$$r_i = r\Delta_i^{\langle 0 \rangle} = r(\Delta_i - 1/n). \tag{A2}$$

[29] https://mc-stan.org/docs/2_18/stan-users-guide/parameterizing-centered-vectors.html

Appendix B
Achromatic Offset and Chromatic Parameter Degeneracy

As described in P. F. Léget et al. (2020), for all SNe Ia there is a degeneracy between an achromatic offset parameter (i.e., χ_0) and a phase-independent chromatic variation vector (i.e., $c_1\boldsymbol{L}_1 + c_2\boldsymbol{L}_2$). With an explicit length-n_λ vector of ones $\mathbf{1}$ to represent achromatic offset behavior in wavelength space, one can define a vector for each SN Ia:

$$\boldsymbol{v} \equiv \chi_0\mathbf{1} + c_1\boldsymbol{L}_1 + c_2\boldsymbol{L}_2. \tag{B1}$$

For two arbitrary constants $\{\alpha_1, \alpha_2\}$, the vector $\boldsymbol{v}$ is invariant under the following transformations:

$$\begin{aligned}
\boldsymbol{v} &= \chi_0{}'\mathbf{1} + c_1\boldsymbol{L}_1{}' + c_1\boldsymbol{L}_2{}',\\
\boldsymbol{L}_1{}' &= \boldsymbol{L}_1 - \alpha_1\mathbf{1},\\
\boldsymbol{L}_2{}' &= \boldsymbol{L}_2 - \alpha_2\mathbf{1},\\
\chi_0{}' &= \chi_0 + \alpha_1 c_1 + \alpha_2 c_2.
\end{aligned}$$

At each iteration during sampling or in postprocessing, this degeneracy can be removed by choosing a fixed $\{\alpha_1, \alpha_2\}$ and then recalculating the components of $\boldsymbol{v}$. Both the c_1 and c_2 parameter sets should be uncorrelated with peak brightness dispersion, so α_1 or α_2 are calculated so that both c_1 and χ_0, and c_2 and χ_0, are uncorrelated. Representing parameter sets as length-n_{sn} vectors $\boldsymbol{c}_1$, $\boldsymbol{c}_2$, and $\boldsymbol{\chi}_0{}'$, and starting from the transformation definition for $\chi_0{}'$ above, we require:

$$\begin{aligned}
\begin{bmatrix} \sigma_{\mathbf{c}_1\chi_0{}'} \\ \sigma_{\mathbf{c}_2\chi_0{}'} \end{bmatrix} &= \begin{bmatrix} \sigma_{\mathbf{c}_1\chi_0} \\ \sigma_{\mathbf{c}_2\chi_0} \end{bmatrix} + \begin{bmatrix} \alpha_1\sigma_{\mathbf{c}_1\mathbf{c}_1} + \alpha_2\sigma_{\mathbf{c}_1\mathbf{c}_2} \\ \alpha_1\sigma_{\mathbf{c}_2\mathbf{c}_1} + \alpha_2\sigma_{\mathbf{c}_2\mathbf{c}_2} \end{bmatrix}\\
&= \begin{bmatrix} \sigma_{\mathbf{c}_1\chi_0} \\ \sigma_{\mathbf{c}_2\chi_0} \end{bmatrix} + \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} \begin{bmatrix} \sigma_{\mathbf{c}_1\mathbf{c}_1} & \sigma_{\mathbf{c}_1\mathbf{c}_2} \\ \sigma_{\mathbf{c}_2\mathbf{c}_1} & \sigma_{\mathbf{c}_2\mathbf{c}_2} \end{bmatrix} = 0,
\end{aligned} \tag{B2}$$

or after solving for the vector $[\alpha_1, \alpha_2]^T$:

$$\begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} = -\begin{bmatrix} \sigma_{\mathbf{c}_1\chi_0} \\ \sigma_{\mathbf{c}_2\chi_0} \end{bmatrix} \begin{bmatrix} \sigma_{\mathbf{c}_1\mathbf{c}_1} & \sigma_{\mathbf{c}_1\mathbf{c}_2} \\ \sigma_{\mathbf{c}_2\mathbf{c}_1} & \sigma_{\mathbf{c}_2\mathbf{c}_2} \end{bmatrix}^{-1}, \tag{B3}$$

where the matrix on the right-hand side of this equation is the precision matrix for our c_1 and c_2 samples.

Both $\boldsymbol{L}_1^{'}$ and $\boldsymbol{L}_2^{'}$ are maintained as unit vectors. This is enforced after each transformation by normalizing $\boldsymbol{L}_1^{'}$ and $\boldsymbol{L}_2^{'}$ by factors $|\boldsymbol{L}_1 - \alpha_1\mathbf{1}|$ and $|\boldsymbol{L}_2 - \alpha_2\mathbf{1}|$, respectively. Each parameter set for c_1 and c_2 is multiplied by their respective factors to preserve the products $c_1\boldsymbol{L}_1^{'}$ and $c_2\boldsymbol{L}_2^{'}$, respectively.

Appendix C
Brief Introduction to Geometric Algebra

Geometric algebra provides a robust, elegant, and easy to interpret framework to perform geometric operations. It extends linear algebra by introducing the geometric product, a combination of the inner product and outer product. Although

we directly calculate intersections, rotations, and projections/rejections using geometric algebra, these operations can be performed using linear algebra instead.

We will assume through that our vector space is finite dimensional and that each vector's elements is real-valued (specifically, we assume our vector space's field are the real numbers $\mathbb{R}$). We will also utilize a canonical basis representation of our vector, such a Cartesian basis for $\mathbb{R}^n$, $\{\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \ldots, \hat{\mathbf{e}}_n\}$. All geometric algebra calculations are implemented with the Clifford package (Robert Kern and the Clifford Team 2021). For a more detailed introduction, see E. Hitzer (2013).

Geometric algebra extends elementary vector algebra using the geometric product. The geometric product of two vectors $\{\boldsymbol{x}, \boldsymbol{y}\} \in \mathbb{R}^n$ combines a symmetric inner (or dot) product with an antisymmetric outer (or wedge) product:

$$\begin{aligned}\boldsymbol{x}\boldsymbol{y} &= \boldsymbol{x} \cdot \boldsymbol{y} + \boldsymbol{x} \wedge \boldsymbol{y} = \sum_{i=1}^{n} x_i y_i \\ &\quad + \sum_{i=1}^{n}\sum_{j>i}^{n} (x_i y_j - y_i x_j)\, \hat{\boldsymbol{e}}_i \wedge \hat{\boldsymbol{e}}_j, \qquad \text{(C1)}\end{aligned}$$

$$= |\boldsymbol{x}||\boldsymbol{y}| \cos\theta_{xy} + |\boldsymbol{x}||\boldsymbol{y}| \sin\theta_{xy} \hat{\boldsymbol{x}} \wedge \hat{\boldsymbol{y}}. \qquad \text{(C2)}$$

In the last step we define the angle θ_{xy} as the angle between the two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$. The inner product maps our vectors to our underlying field $\mathbb{R}$, yielding a scalar and is a symmetric operation: $\boldsymbol{x} \cdot \boldsymbol{y} = \boldsymbol{y} \cdot \boldsymbol{x}$. The wedge product of two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ is an antisymmetric operation: $\hat{\boldsymbol{e}}_i \wedge \hat{\boldsymbol{e}}_j = -\hat{\boldsymbol{e}}_j \wedge \hat{\boldsymbol{e}}_i = 0$ (if $i = j$). The object output by the wedge product of two vectors is called a bivector and is an oriented plane element.

It is common in geometric algebra to refer to a scalar as a 0-blade and a bivector as a 2-blade. Similarly, vectors such as $\boldsymbol{x}$ and $\boldsymbol{y}$ are referred to as 1-blades. In general, a k-blade is a k-dimensional object with can be generated by k-wedge (or k-geometric) products of k independent vectors. These k-vectors are themselves oriented k-dimensional subspace elements. The produced object $\boldsymbol{x}\boldsymbol{y}$ is a linear combination of two grades of k-blades (a 0-vector and a 2-vector, specifically) that we call a multivector. In general, a multivector is a linear combination of k-blades. Also, each multivector $\boldsymbol{M}$ has an inverse such that via the geometric product $\boldsymbol{M}\boldsymbol{M}^{-1} = 1$. Furthermore, any multivector defined as a product of vectors $\boldsymbol{M} = \boldsymbol{a}_1\boldsymbol{a}_2 \ldots \mathbf{a}_m$ has a corresponding multivector called its reverse: $\boldsymbol{M}^{\dagger} = \boldsymbol{a}_m\boldsymbol{a}_{m-1} \ldots \mathbf{a}_1$.

We can derive a canonical basis of our geometric algebra using the geometric product on our starting vector space's canonical basis as a generating set. Because canonical basis is orthogonal, the geometric product of canonical basis components reduces to a wedge product for $i \neq j$: $\hat{\boldsymbol{e}}_i\, \hat{\boldsymbol{e}}_j = 0 + \hat{\boldsymbol{e}}_i \wedge \hat{\boldsymbol{e}}_j$, or to an inner product for $i = j$: $\hat{\boldsymbol{e}}_i\, \hat{\boldsymbol{e}}_i = \hat{\boldsymbol{e}}_i \cdot \hat{\boldsymbol{e}}_i + 0$. As an example, for a vector space of dimension $n = 3$, the generated geometric algebraic canonical basis has eight basis elements: $\{1;\ \hat{\boldsymbol{e}}_1,\ \hat{\boldsymbol{e}}_2,\ \hat{\boldsymbol{e}}_3;\ \hat{\boldsymbol{e}}_1\hat{\boldsymbol{e}}_2,\ \hat{\boldsymbol{e}}_1\hat{\boldsymbol{e}}_3,\ \hat{\boldsymbol{e}}_2\hat{\boldsymbol{e}}_3;\ \hat{\boldsymbol{e}}_1\hat{\boldsymbol{e}}_2\hat{\boldsymbol{e}}_3\}$. This set of basis elements naturally fall into grades of equal subspace dimension: one scalar element, three 1D basis vectors (the starting vector space's canonical vector basis), three 2D basis bivectors, and one 3D basis trivector. The geometric algebra over $\mathbb{R}^3$ basis elements therefore prescribes corresponding sets of orthogonal bases for all possible subspaces that exist within $\mathbb{R}^3$. The generation of basis elements generalizes naturally for any finite dimension vector space $\mathbb{R}^n$, with the resulting geometric algebra $\mathbb{G}^n$ having 2^n generated basis components. The respective subspaces, or grades, of the geometric algebra (i.e., the bivector or grade-2 components) have dimensionality $\binom{n}{k}$, where k is the grade of interest. Note that the last element in this geometric algebra basis is called the pseudoscalar i, since $ii = -1$ for any dimension n in a way analogous with an imaginary number. Finally, we can decompose any multivector $\boldsymbol{M}$ using a grade projection operation:

$$\boldsymbol{M} = \sum_{k=1}^{n} \langle \boldsymbol{M} \rangle_k, \qquad \text{(C3)}$$

where the k-grade projection operator $\langle \boldsymbol{M} \rangle_k$ returns the k-blade component of $\boldsymbol{M}$.

C.1. Reflections and Rotations

Consider two vectors $\boldsymbol{r}, \boldsymbol{x} \in \mathbb{R}^n$. Say we reflect $\boldsymbol{r}$ along $\boldsymbol{x}$. This operation can be interpreted as negating the component of $\boldsymbol{r}$ perpendicular to $\boldsymbol{x}$. The geometric product can be used to decompose $\boldsymbol{r}$ into its parallel and perpendicular components relative to the unit vector $\hat{\boldsymbol{x}}$:

$$\boldsymbol{r}\hat{\boldsymbol{x}} = \boldsymbol{r} \cdot \hat{\boldsymbol{x}} + \boldsymbol{r} \wedge \hat{\boldsymbol{x}}, \qquad \text{(C4)}$$

$$\boldsymbol{r} = (\boldsymbol{r} \cdot \hat{\boldsymbol{x}})\hat{\boldsymbol{x}}^{-1} + (\boldsymbol{r} \wedge \hat{\boldsymbol{x}})\hat{\boldsymbol{x}}^{-1}. \qquad \text{(C5)}$$

The component $(\boldsymbol{r} \cdot \hat{\boldsymbol{x}})\hat{\boldsymbol{x}}^{-1}$ is proportional to the projection of $\boldsymbol{r}$ onto $\boldsymbol{x}$, while the component $(\boldsymbol{r} \wedge \hat{\boldsymbol{x}})\hat{\boldsymbol{x}}^{-1}$ is the perpendicular (or rejected) remainder. From this decomposition we can define a reflection $\boldsymbol{r}'$ of $\boldsymbol{r}$ along $\boldsymbol{x}$:

$$\begin{aligned}\boldsymbol{r}' &= -(\boldsymbol{r} \cdot \hat{\boldsymbol{x}})\hat{\boldsymbol{x}}^{-1} + (\boldsymbol{r} \wedge \hat{\boldsymbol{x}})\hat{\boldsymbol{x}}^{-1} \\ &= -(\hat{\boldsymbol{x}} \cdot \boldsymbol{r})\hat{\boldsymbol{x}}^{-1} - (\hat{\boldsymbol{x}} \wedge \boldsymbol{r})\hat{\boldsymbol{x}}^{-1} \\ &= -(\hat{\boldsymbol{x}} \cdot \boldsymbol{r} + \hat{\boldsymbol{x}} \wedge \boldsymbol{r})\hat{\boldsymbol{x}}^{-1} \\ &= -\hat{\boldsymbol{x}}\boldsymbol{r}\hat{\boldsymbol{x}}^{-1} \\ &= -\hat{\boldsymbol{x}}\boldsymbol{r}\hat{\boldsymbol{x}}.\end{aligned}$$

Here we exploit the unit vector's geometric algebra inverse being itself ($\hat{\boldsymbol{x}}\hat{\boldsymbol{x}} = \hat{\boldsymbol{x}} \cdot \hat{\boldsymbol{x}} = 1$).

Any rotation can be decomposed into two reflections, with the rotational plane being spanned by two vectors we reflect against. Picking another vector $\boldsymbol{y} \in \mathbb{R}^n$, we write the rotation $\boldsymbol{r}'$ of $\boldsymbol{r}$ in the plane corresponding to grade-2 (or bivector component) of the multivector $\boldsymbol{R} = \hat{\boldsymbol{x}}\hat{\boldsymbol{y}}$ as:

$$\begin{aligned}\boldsymbol{r}' &= \hat{\boldsymbol{y}}\hat{\boldsymbol{x}}\boldsymbol{r}\hat{\boldsymbol{x}}\hat{\boldsymbol{y}} \\ &= \boldsymbol{R}^{\dagger}\boldsymbol{r}\boldsymbol{R}.\end{aligned}$$

This bilinear operation on $\boldsymbol{r}$ by the unit 2-vector $\boldsymbol{R}$ is indeed a rotation. To see this, let us insert a more illuminating form of $\boldsymbol{R}$:

$$\begin{aligned}\boldsymbol{r}' = &\left(\cos\frac{\theta_{xy}}{2} - \sin\frac{\theta_{xy}}{2}\hat{\boldsymbol{x}} \wedge \hat{\boldsymbol{y}}\right)\boldsymbol{r} \\ &\times \left(\cos\frac{\theta_{xy}}{2} + \sin\frac{\theta_{xy}}{2}\hat{\boldsymbol{x}} \wedge \hat{\boldsymbol{y}}\right). \qquad \text{(C6)}\end{aligned}$$

To interpret this result, consider first the geometric algebra of vector space $\mathbb{R}^2$ with its four generated basis elements $\{1;\ \hat{\boldsymbol{e}}_1,\ \hat{\boldsymbol{e}}_2;\ \hat{\boldsymbol{e}}_1\hat{\boldsymbol{e}}_2\}$. In 2D any bivector $\hat{\boldsymbol{x}} \wedge \hat{\boldsymbol{y}}$ is proportional our unit pseudoscalar $\hat{\boldsymbol{e}}_1\hat{\boldsymbol{e}}_2$, itself which behaves identically to the imaginary number (similarly denoted as i). Substituting i for $\hat{\boldsymbol{x}} \wedge \hat{\boldsymbol{y}} = \hat{\boldsymbol{e}}_1\hat{\boldsymbol{e}}_2$ into Equation (C6), we recover the canonical form of a rotation by an angle θ in 2D:

$$\boldsymbol{r}' = e^{-i\theta/2}\boldsymbol{r}e^{i\theta/2}. \qquad \text{(C7)}$$

For higher dimensions we can similarly treat the bivector $\hat{\boldsymbol{x}} \wedge \hat{\boldsymbol{y}}$ as being the imaginary number of the 2D subspace spanned by said bivector. Therefore, any rotation of a vector $\boldsymbol{r}$ by an angle θ within a plane spanned by the unit bivector $\hat{\mathcal{A}} = \hat{\mathbf{x}} \wedge \hat{\mathbf{y}}$ can be written:

$$\boldsymbol{r}' = e^{-\hat{\mathcal{A}}\theta/2}\boldsymbol{r}e^{\hat{\mathcal{A}}\theta/2} = \boldsymbol{R}^{\dagger}\boldsymbol{r}\boldsymbol{R}. \tag{C8}$$

This unit bivector is interpreted as a generator of rotation within the plane spanned by $\hat{\mathcal{A}}$; the multivector $\boldsymbol{R}$ (which has scalar and bivector components) is called a rotor.

C.2. Projection and Rejection of Vectors onto a Bivector

Consider a vector $\boldsymbol{v}$ and a bivector $\mathcal{A}$. Similar to Appendix C.1, we can decompose a $\boldsymbol{v}$ into components parallel and perpendicular to the plane spanned by $\mathcal{A}$:

$$\boldsymbol{v}\mathcal{A} = \boldsymbol{v}\cdot \mathcal{A} + \boldsymbol{v} \wedge \mathcal{A}, \tag{C9}$$

$$\boldsymbol{v} = (\boldsymbol{v} \cdot \mathcal{A})\mathcal{A}^{-1} + (\boldsymbol{v} \wedge \mathcal{A})\mathcal{A}^{-1}. \tag{C10}$$

If the wedge product of any vector with a $\mathcal{A}$ is 0, then said vector must exist within the plane spanned by that $\mathcal{A}$. Alternatively, any vector with no projection onto $\mathcal{A}$ requires said vector be perpendicular to $\mathcal{A}$. Just as in Appendix C.1, we therefore identify this decomposition into parallel (projected) and perpendicular (rejected) components relative to a plane spanned by $\mathcal{A}$:

$$\boldsymbol{v}^{\parallel} = (\boldsymbol{v} \cdot \mathcal{A})\mathcal{A}^{-1}, \tag{C11}$$

$$\boldsymbol{v}^{\perp} = (\boldsymbol{v} \wedge \mathcal{A})\mathcal{A}^{-1}. \tag{C12}$$

C.3. Intersection of Planes

This section describes the formalism readily displayed in Figure 5. It is the most complicated of the operations we implement using geometric algebra and is included for completeness, despite its complexity. As with other operations, it is possible to calculate this intersection using conventional linear algebra techniques. Indeed, we confirmed this geometric algebraic implementation with such a conventional approach.

Intuitively, a bivector is to a plane what a vector is to a line. Just as one can find the intersection of two lines using their representing vectors, one also can find the intersection of two planes using their representing bivectors.

Consider two bivectors $\mathcal{A}$ and $\mathcal{B}$, each of which represent two different planes. If these two planes intersect (which we will assume they do), then there are at most two linearly independent vectors $\boldsymbol{a}$ and $\boldsymbol{b}$ that exist each of the respective planes spanned by $\mathcal{A}$ and $\mathcal{B}$, but that do not reside along these two planes' intersection. Furthermore, the intersection of $\mathcal{A}$ and $\mathcal{B}$ is a well-defined, 1D subspace necessarily spanned by some vector, a vector that we will call $\boldsymbol{c}$. This means that $\mathcal{A} = \boldsymbol{a} \wedge \boldsymbol{a}$ and $\mathcal{B} = \boldsymbol{c} \wedge \boldsymbol{b}$, and taken together, these two bivectors span a 3D subspace (because they share the intersection subspace spanned by $\boldsymbol{c}$). This volume and its orientation can be represented using a trivector $\boldsymbol{a} \wedge \boldsymbol{b} \wedge \boldsymbol{c}$. Normalizing this then trivector gives us the intersection's unit volume element:

$$i = \frac{\boldsymbol{a} \wedge \mathcal{B}}{|\boldsymbol{a} \wedge \mathcal{B}|} = \frac{\boldsymbol{b} \wedge \mathcal{A}}{|\boldsymbol{b} \wedge \mathcal{A}|}. \tag{C13}$$

We want a formal expression for $\boldsymbol{c}$ from the starting bivectors $\mathcal{A}$ and $\mathcal{B}$, though. To get this, we first find a bivector whose plane is simultaneously perpendicular to both $\mathcal{A}$ and $\boldsymbol{B}$, which can be done by taking the grade projection of the geometric product $\mathcal{A}\mathcal{B}$ to its bivector component: $\langle \mathcal{A}\boldsymbol{B}\rangle_2$. The vector perpendicular to this bivector $\langle \mathcal{A}\boldsymbol{B}\rangle_2$, but which still exists within the volume spanned by the intersection unit volume i, is the intersection vector $\boldsymbol{c}$ we are looking for. It turns out that taking the geometric product of i with $\langle \mathcal{A}\boldsymbol{B}\rangle_2$ performs exactly this operation:

$$\mathbf{c} = i\langle \mathcal{A}\mathcal{B}\rangle_2. \tag{C14}$$

Once way to think about this is that the unit volume i rotates the bivector $\langle \mathcal{A}\boldsymbol{B}\rangle_2$ by $\pi/2$ and projects out the vector $\boldsymbol{c}$ perpendicular to the bivector $\langle \mathcal{A}\boldsymbol{B}\rangle_2$.

Appendix D
Hamiltonian Monte Carlo Sample Chains

As mentioned and interpreted in Sections 3 and 4, we identify three HMC sampler groupings during postprocessing. Here we present the HMC posterior. Sampling chains for the raw $\boldsymbol{L}_1$ parameters before achromatic offset de-correlation (Appendix B) or basis selection (Section 4.1). As such, the sampler chains seen here are not yet physically interpretable, nor are they used to measure convergence. Although also present in postprocessed basis representations, the sampler groupings are far more visible before postprocessing, providing easier "by eye" identification. See Figure 17 for a presentation of the HMC sample chains.

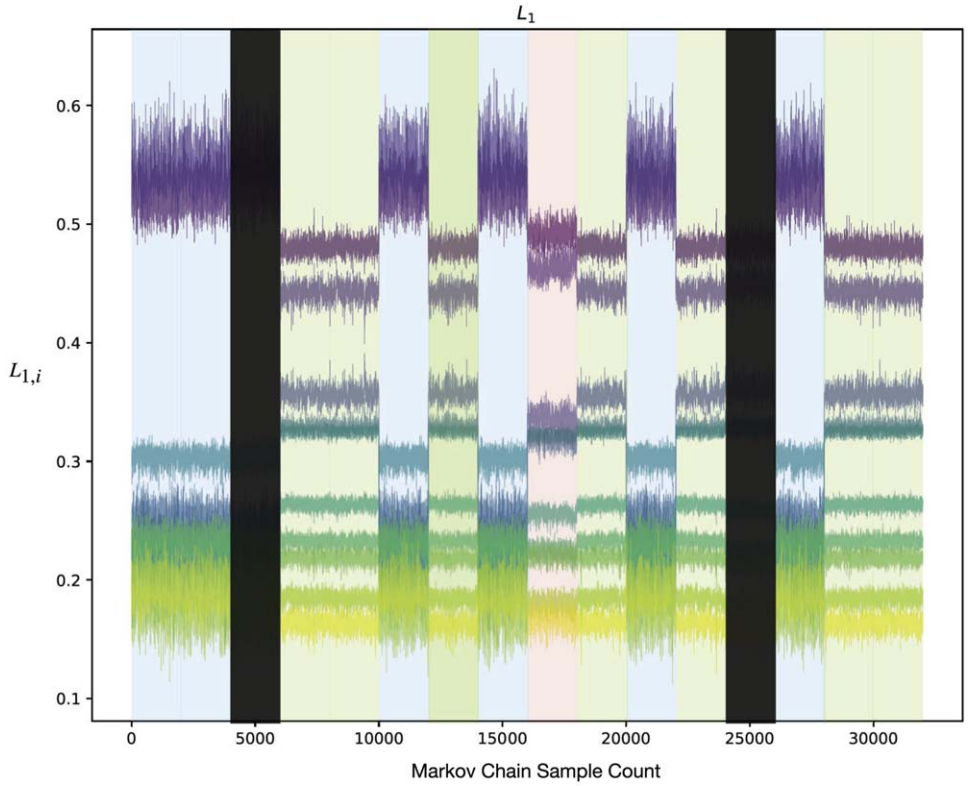

Figure 17. Posterior samplings for each of the $\boldsymbol{L}_1$ template's 10 components. The x-axis labels the 2000 samples for each of the 16 samplers, concatenated in sequence for 32,000 total draws. The aforementioned groupings are shaded blue, green, and pink. The dark gray rectangles mask the samples drawn by two samplers which became lost during sampling and were subsequently excluded from our analysis.

ORCID iDs

Jared Hand https://orcid.org/0000-0001-7260-4274
A. G. Kim https://orcid.org/0000-0001-6315-8743
P. Antilogus https://orcid.org/0000-0002-0389-5706
C. Aragon https://orcid.org/0000-0002-9502-0965
S. Bailey https://orcid.org/0000-0003-4162-6619
C. Baltay https://orcid.org/0000-0003-0424-8719
K. Boone https://orcid.org/0000-0002-5828-6211
C. Buton https://orcid.org/0000-0002-3780-7516
Y. Copin https://orcid.org/0000-0002-5317-7518
S. Dixon https://orcid.org/0000-0003-1861-0870
D. Fouchez https://orcid.org/0000-0002-7496-3796
E. Gangler https://orcid.org/0000-0001-6728-1423
R. Gupta https://orcid.org/0000-0003-1820-4696
B. Hayden https://orcid.org/0000-0001-9200-8699
Mitchell Karmen https://orcid.org/0000-0003-2495-8670
M. Kowalski https://orcid.org/0000-0001-8594-8666
D. Küsters https://orcid.org/0000-0002-9207-4749
P.-F. Léget https://orcid.org/0000-0002-8357-3984
J. Nordin https://orcid.org/0000-0001-8342-6274
R. Pain https://orcid.org/0000-0003-4016-6067
S. Perlmutter https://orcid.org/0000-0002-4436-4661
K. A. Ponder https://orcid.org/0000-0002-8207-3304
D. Rabinowitz https://orcid.org/0000-0003-4961-7653
M. Rigault https://orcid.org/0000-0002-8121-2560
D. Rubin https://orcid.org/0000-0001-5402-4647
C. Saunders https://orcid.org/0000-0002-4094-2102
N. Suzuki https://orcid.org/0000-0001-7266-930X
S. Taubenberger https://orcid.org/0000-0002-4265-1958

References

Aldering, G., Adam, G., Antilogus, P., et al. 2002, Proc. SPIE, 4836, 61
Aldering, G., Antilogus, P., Aragon, C., et al. 2020, RNAAS, 4, 63
Bailey, S., Aldering, G., Antilogus, P., et al. 2009, A&A, 500, L17
Betoule, M., Kessler, R., Guy, J., et al. 2014, A&A, 568, A22
Blondin, S., Matheson, T., Kirshner, R. P., et al. 2012, AJ, 143, 126
Bohlin, R. C. 2014, AJ, 147, 127
Bongard, S., Soulez, F., Thiébaut, É., & Pecontal, É. 2011, MNRAS, 418, 258
Boone, K., Aldering, G., Antilogus, P., et al. 2021a, ApJ, 912, 71
Boone, K., Aldering, G., Antilogus, P., et al. 2021b, ApJ, 912, 70
Branch, D., Chau Dang, L., & Baron, E. 2009, PASP, 121, 238
Branch, D., Fisher, A., & Nugent, P. 1993, AJ, 106, 2383
Brout, D., & Scolnic, D. 2021, ApJ, 909, 26
Burns, C. R., Stritzinger, M., Phillips, M. M., et al. 2011, AJ, 141, 19
Buton, C., Copin, Y., Aldering, G., et al. 2013, A&A, 549, A8
Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, ApJ, 345, 245
Carpenter, B., Gelman, A., Hoffman, M. D., et al. 2017, Journal of Statistical Software, 76, 1
Chotard, N., Gangler, E., Aldering, G., et al. 2011, A&A, 529, L4
Conley, A., Sullivan, M., Hsiao, E. Y., et al. 2008, ApJ, 681, 482
Dhawan, S., Leibundgut, B., Spyromilio, J., & Maguire, K. 2015, MNRAS, 448, 1345
Fakhouri, H. K., Boone, K., Aldering, G., et al. 2015, ApJ, 815, 58
Foley, R. J. 2013, MNRAS, 435, 273
Guy, J., Astier, P., Baumont, S., et al. 2007, A&A, 466, 11
Hamuy, M., Suntzeff, N. B., Heathcote, S. R., et al. 1994, PASP, 106, 566
Hamuy, M., Walker, A. R., Suntzeff, N. B., et al. 1992, PASP, 104, 533
Hitzer, E. 2013, arXiv:1306.1660
Hoyle, F., & Fowler, W. A. 1960, ApJ, 132, 565
Iben, I. J., & Tutukov, A. V. 1984, ApJS, 54, 335
Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, ApJ, 873, 111
Jha, S., Riess, A. G., & Kirshner, R. P. 2007, ApJ, 659, 122
Jha, S. W., Maguire, K., & Sullivan, M. 2019, NatAs, 3, 706
Kasen, D., & Woosley, S. E. 2007, ApJ, 656, 661
Kenworthy, W. D., Jones, D. O., Dai, M., et al. 2021, ApJ, 923, 265

Robert Kern and the Clifford Team 2021, clifford: Geometric Algebra for Python, v1.4.0, GitHub, https://github.com/pygae/clifford
Kim, A., Goobar, A., & Perlmutter, S. 1996, PASP, 108, 190
Lantz, B., Aldering, G., Antilogus, P., et al. 2004, Proc. SPIE, 5249, 146
Léget, P. F., Gangler, E., Mondon, F., et al. 2020, A&A, 636, A46
Mandel, K. S., Scolnic, D. M., Shariff, H., Foley, R. J., & Kirshner, R. P. 2017, ApJ, 842, 93
Mandel, K. S., Thorp, S., Narayan, G., Friedman, A. S., & Avelino, A. 2022, MNRAS, 510, 3939
Nugent, P., Kim, A., & Perlmutter, S. 2002, PASP, 114, 803
Nugent, P., Phillips, M., Baron, E., Branch, D., & Hauschildt, P. 1995, ApJL, 455, L147
Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, JMLR, 12, 2825
Pereira, R., Thomas, R. C., Aldering, G., et al. 2013, A&A, 554, A27
Perlmutter, S., Aldering, G., Goldhaber, G., et al. 1999, ApJ, 517, 565
Perlmutter, S., Gabi, S., Goldhaber, G., et al. 1997, ApJ, 483, 565
Phillips, M. M. 1993, ApJL, 413, L105
Pierel, J. D. R., Jones, D. O., Kenworthy, W. D., et al. 2022, ApJ, 939, 11
Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998a, AJ, 116, 1009
Riess, A. G., Nugent, P., Filippenko, A. V., Kirshner, R. P., & Perlmutter, S. 1998b, ApJ, 504, 935
Riess, A. G., Press, W. H., & Kirshner, R. P. 1996, ApJ, 473, 88
Rose, B. M., Rubin, D., Strolger, L., & Garnavich, P. M. 2021, ApJ, 909, 28
Rubin, D., Aldering, G., Antilogus, P., et al. 2022, ApJS, 263, 1
Rubin, D., Aldering, G., Barbary, K., et al. 2015, ApJ, 813, 137
Saunders, C., Aldering, G., Antilogus, P., et al. 2018, ApJ, 869, 167
Scalzo, R. A., Ruiter, A. J., & Sim, S. A. 2014, MNRAS, 445, 2535
Scolnic, D. M., Jones, D. O., Rest, A., et al. 2018, ApJ, 859, 101
Scolnic, D. M., Riess, A. G., Foley, R. J., et al. 2014, ApJ, 780, 37
Soker, N. 2019, NewAR, 87, 101535
Stein, G., Seljak, U., Bohm, V., et al. 2022, ApJ, 935, 5
Thorp, S., Mandel, K. S., Jones, D. O., Ward, S. M., & Narayan, G. 2021, MNRAS, 508, 4310
Tripp, R. 1998, A&A, 331, 815
Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., & Bürkner, P.-C. 2021, BayAn, 16, 667
Wang, X., Filippenko, A. V., Ganeshalingam, M., et al. 2009, ApJL, 699, L139
Weingartner, J. C., & Draine, B. T. 2001, ApJ, 548, 296
Whelan, J., & Iben, I. J. 1973, ApJ, 186, 1007