Mixture-of-Expert Variational Autoencoders for Cross-Modality Embedding of Type Ia Supernova Data

Yunyi Shen*¹² Alexander T. Gagliano*¹³⁴

Abstract

Time-domain astrophysics relies on heterogeneous and multi-modal data. Specialized models are often constructed to extract information from a single modality, but this approach ignores the wealth of cross-modality information that may be relevant for the tasks to which the model is applied. In this work, we propose a multi-modal, mixture-of-expert variational autoencoder to learn a joint embedding for supernova light curves and spectra. Our method, which is inspired by the perceiver architecture, natively accommodates variable-length inputs and the irregular temporal sampling inherent to supernova light curves. We train our model on radiative transfer simulations and validate its performance on crossmodality reconstruction of supernova spectra and physical parameters from the simulation. Our model achieves superior performance in crossmodality generation to nearest-neighbor searches in a contrastively-trained latent space, showing its promise for constructing informative latent representations of multi-modal astronomical datasets.

1. Introduction

Diversity of Type Ia Supernovae: Type Ia supernovae (SNe Ia) are the thermonuclear explosions of white dwarfs in binary systems. These explosions occur as the primary star accretes mass from its companion and approaches the Chandrasekhar limit (Hoyle & Fowler, 1960; Liu et al., 2023). Because of the relatively robust correlations among their observational properties (e.g., the "Phillips relation" between an explosion's maximum luminosity and its dimming rate; Pskovskii, 1977; Phillips, 1993; Phillips et al., 1999), SNe Ia have increasingly been used as standardizable candles to measure cosmological distances (e.g., Riess et al.

1998; Perlmutter et al. 1999; DES Collaboration et al. 2024).

As more SNe Ia are observed from wide-field imaging surveys (such as the Zwicky Transient Facility; Bellm et al. 2018), however, researchers have discovered SNe Ia with atypical light curve features (e.g., early flux excesses; Wang et al. 2024). The spectra of these SNe, which measure their wavelength-specific emission, have revealed important differences between the temperatures, velocities, and compositions of these more peculiar explosions and their normal counterparts.

Legacy Survey of Space and Time: With the upcoming Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST; Ivezić et al., 2019) imaging the full Southern Sky every 3-4 days for a decade, we expect to obtain light curves for ~ 1 M SNe each year. Our spectroscopic resources will only be able to observe the smallest subset of these discovered explosions (at most 0.1-1%). As a result, representation learning techniques applied to these datasets must be able to accommodate missing modalities at inference time. Methods able to extract multi-modal correlations from these highly imbalanced datasets will play a powerful role in helping astronomers prioritize scientifically valuable events for follow-up study while they remain bright.

Multi-Modal Feature Extraction and Reconstruction: We consider two tasks in this work: 1) inferring the timeevolving spectroscopic behavior of a supernova from measurements of its light curve (photometry) alone, and 2) learning a joint data representation for photometry and spectra that is useful for downstream tasks. Though feature learning methods using contrastive learning (Zhang et al., 2024) and pure generation methods using diffusion (Shen & Gagliano, 2025) have been applied to these modalities, a single approach has not yet been presented that can simultaneously achieve goals 1 and 2. As an obstacle toward these goals, the correlations between light curves and spectra are many-tomany: a spectrum obtained at a single time can be associated with multiple feasible light curves, and each light curve is associated with multiple spectra describing the wavelengthdependent emission of an explosion at different times. Furthermore, the dimensionality of each modality varies with the total number of observations obtained and the resolution of the instrumentation used. For these reasons, our

^{*}Equal contribution ¹IAIFI ²EECS, MIT ³MKI, MIT ⁴CfA | Harvard & Smithsonian. Correspondence to: Yunyi Shen <yshen99@mit.edu>, Alexander Gagliano <gaglian2@mit.edu>.

ML4Astro 2025, Vancouver, CA. Copyright 2025 by the author(s).

proposed method should be able to 1) learn compact private and shared representations, 2) operate on a variable number of modalities and input lengths at inference time, and 3) keep inference costs low as input dimensionality increases.

We achieve these goals by constructing a Perceiver-IO (Jaegle et al., 2021a) style model for time-domain inputs, which we call the transceiver (for "transient perceiver"). We combine this model with a mixture-of-expert VAE (MM-VAE, Shi et al., 2019) for private and shared feature learning, and validate its use for cross-modality generation. Our model and associated experiments can be found at the github repository for this work ¹.

2. Method

The heterogeneous nature of supernova observations demands custom encoding methods. We begin by describing our data and the chosen encoder-decoder architecture. Then, we introduce the architecture and training objective associated with our MMVAE.

Data Encoding: Our approach follows the encoding scheme of Zhang et al. (2024), but augments it with learnable positional encodings and phase information to enable multiple spectra to be linked to a single light curve.

Photometry: For observation *i* obtained with photometric filter *n* during an imaging survey containing *B* unique photometric filters, $(t_{n,i}, m_{n,i}, b_{n,i}) \in \mathbb{R} \times \mathbb{R} \times [B]$ represents the observation time, filter, and measured magnitude of the explosion. Consequently, the full light curve is a sequence of $(t_{n,i}, m_{n,i}, b_{n_i})$ with length $L_{\text{photo},n}$.

Our light curves are encoded to a model dimension d by summing a sinusoidal encoding for time, a class embedding of filter, and linear projection of magnitude, resulting in an embedded sequence of dimensionality $\mathbb{R}^{L_{\text{photo},n} \times d}$ that was passed into perceiver. We encode observation times with a sinusoidal positional encoding followed by a learnable MLP (Peebles & Xie, 2023). Magnitude is encoded using a learnable linear embedding and the photometric filter is encoded using a simple categorical embedding. For parallelization, we also pad our input light curves to a consistent length; we retain a mask during training to indicate padded positions.

We visualize the full architecture of our model in Figure 1.

Spectroscopy: Spectra encode the flux from the supernova as a function of wavelength. Each spectrum is taken at a fixed point in time, which we represent as phase with respect to the time of supernova peak light. The *n*th spectrum associated with each supernova is represented as a sequence of flux values $f_{n,i}$ measured at wavelengths $\lambda_{n,i}$. As a result, the spectra for each supernova have dimensionality $L_{\text{spec},n}$.

We encode each spectrum by separately embedding wavelengths and fluxes to the model dimension d and summing them. Wavelengths are encoded using sinusoidal embeddings followed by a learnable MLP, while flux values are embedded linearly. This results in an embedding of shape $\mathbb{R}^{L_{\text{spec},n} \times d}$. Unlike photometry, we separately embed the phase using a sinusoidal embedding followed by an MLP projection to dimension d, and then append it to the sequence as a special token. This special token approach allows for more flexible information passing through crossattention. The full spectrum is thus embedded into dimensionality $\mathbb{R}^{(L_{\text{spec},n}+1)\times d}$. As with photometry, we pad input spectra and retain a mask indicating which positions are padded.

Transceiver Encoder-Decoder: We use an architecture akin to Perceiver and Perceiver-IO (Jaegle et al., 2021b;a) for our encoder-decoder setup. This architecture handles **any-length sequences** via cross-attention. Specifically, in the encoder, a fixed-size latent sequence (as query) attends to the input spectrum/light curve embedding (as key and value) to produce the latent posterior mean and variance. We refer to this architecture for transients as the transceiver. This differs from the vanilla transformer encoder used in models like Maven (Zhang et al., 2024), in which the input data itself serve as query and therefore must be fixed-size.

Multi-Modal Mixture-of-Expert VAE: Shi et al. (2019) proposed using a mixture of-expert VAE (MMVAE) to combine different modalities, enabling learning of both shared and private information. The key idea in MMVAE is to model the latent variable z as a mixture of distributions associated with the different modalities. Formally, let z_n denote the latent variable for the *n*th training pair and $x_{n,m}$ denote the data from modality $m = 1, \ldots, M$. Our MMVAE approximates the conditional latent distribution, denoted q_{ϕ_m} , using transceiver encoders parameterized by ϕ_m . Each encoder outputs the mean and variance of a Laplace distribution:

$$p(z_n|x_{n,1},\ldots,x_{n,M}) \approx q_\phi := \frac{1}{M} \sum_{m=1}^M q_{\phi_m}(z_n|x_{n,m})$$
(1)

Though learned jointly, we can use the encoder for each modality independently or combine them via the mixtureof-experts model at inference time. Each modality also has its own decoder. That is, with a decoder parameterized by ψ_m ,

$$p(x_{n,1,...,M}|z_n) \approx \prod_{m=1}^M q_{\psi_m}(x_{n,m}|z_n)$$
 (2)

We use the transceiver decoder to parameterize the mean of a Laplace distribution for each modality. We model the prior p(z) as a standard Laplace distribution with latent space size $z \in \mathbb{R}^{4 \times 4}$. Training uses the IWAE objective

¹https://github.com/YunyiShen/VAESNe-dev



Figure 1. Embedding method used for supernova spectra and light curves (see text for details). The approach is similar to that used in Zhang et al. (2024), but with learnable positional encodings for observation times of photometry and spectroscopy.



Figure 2. Encoder-decoder architecture with perceiver-IOstyle architecture for individual modalities. Modality encodings are cross attended by a latent bottleneck sequence.

proposed by Shi et al. (2019):

$$\mathcal{L}(x_{1,\dots,M}) = \mathbb{E}_{z^{1:K} \sim q\phi} \left[\log \sum_{k=1}^{K} \frac{1}{K} \frac{p(z) \prod_{m} q_{\psi_{m}}}{q_{\phi}} \right]$$
(3)

We train using AdamW in PyTorch using a learning rate of 0.001 for 500 iterations. For our transceiver encoder and decoder, we have 4 layers of multi-headed attention with 4 heads each. The model operates in the dimension of 32, with an MLP dimension of 32 in the attention layers. We choose a latent dimension of 4 by 4.

Uni-Modal VAE for Spectra Generation: To evaluate our model's capacity for cross-modality conditional generation, we build a reference baseline: a VAE trained directly on spectral data using the same transceiver encoderdecoder architecture but trained with the standard ELBO loss. This baseline ignores light curve information completely.

Multi-Modal Transceiver Trained via Contrastive Learning: As a comparison, we also jointly train transceiver encoders using a constrastive objective. We project light curves and spectra into a shared $\mathbb{R}^{4\times4}$ latent space, further project them to $\mathbb{R}^{4\times8}$ using a learned MLP, compute the Frobenius distance between pairs, and train the model to minimize the InfoNCE loss (Khosla et al., 2020).

3. Experiments

Data: We use the simulation grid from Goldstein & Kasen (2018) to train and validate our MMVAE. Goldstein & Kasen (2018) generates 4,500 Type Ia supernovae (SNe Ia) using the Sedona radiation transport code (Kasen, 2006). Each simulation yields a full spectral energy distribution (SED) surface, sampled at 1-day intervals in time and \sim 30 Å intervals in wavelength. We present the physical parameters associated with the simulation grid in Table 1 in Appendix A.

We simulate LSST-like photometry by integrating the SED surfaces across the transmission curves of the LSST filters, obtained from the SVO Filter Profile Service². The light curve cadence is based on the baseline v3.3 simulations of LSST's Wide-Fast-Deep survey. We assume a random distribution of events in both time since survey start and position across the survey footprint.

After sampling photometry, we require at least 10 total photometric observations across all bands to include an event in our train/validation/test sets. For simulated spectra, we extract fluxes directly from the SED surface in 10-day windows from 10 days before peak brightness to 30 days after it. We apply a median filter with bandwidth 3 to approximate the finite resolution of spectrographs. We then train the model on the log-transformed flux values in units of erg s⁻¹ cm⁻² Hz⁻¹. All flux, wavelength, phase, and magnitude values are standardized to z-scores during training and converted back to physical units at inference time. Light curves are zero-padded to 10 measurements per band during training solely for parallelization, yielding uniform arrays of length 60. We split the dataset into fractions of 80/10/10 for training, validation, and testing, respectively.

Prior Sampling: To evaluate the spectroscopic diversity learned by the model, we sample from both the latent (standard Laplace) prior and the conditionally-generated spectra from events in our test set. We present the results in Appendix B.1. The samples resemble the observed diversity in the test sample with some smoothing of high-frequency information, as is commonly observed in VAEs.

²http://svo2.cab.inta-csic.es/

Cross-Modality Generation: In this setup, we encode photometry using the trained encoder, sample from the posterior distribution in latent space, and decode into spectra using the trained decoder. We evaluate our conditional generation using three metrics: the residual error in the posterior mean, the credible interval (CI) coverage, and the CI width. As a baseline, we compare to (i) a uni-modal VAE trained to encode spectra alone, (ii) a nearest neighbor search in a learned contrastive latent space³, and (iii) the average spectrum across the full training set at each phase. The first benchmark allows us to confirm that our model architecture has sufficient capacity to reconstruct each spectrum, while the third ensures that our reconstructions are successfully conditioned by the event's photometry.

Our results are shown in Figure 3. Our MMVAE achieves reconstruction performance comparable to the spectra-only VAE and exceeding that of the contrastive method, which essentially retrieves the average SN Ia spectrum at each phase. Further, the model's performance remains robust under increased masking of light curve inputs (see Appendix B.2 for details).The posterior samples of the model, however, are not well-calibrated, as the CIs exhibit undercoverage. This may reflect the limited flexibility of the encoder in approximating complex posterior distributions.

We show a representative example of a spectrum decoded by our MMVAE and the baseline models in Figure 4.



Figure 3. Metrics for the generation of supernova spectra conditioned on LSST light curves. Our MMVAE ('LC2spec') achieves similar reconstruction accuracy to a model trained directly on spectra, and outperforms a nearest neighbor search in the contrastive latent space. However, the low CI coverage (middle row) indicates that the posterior samples are not well-calibrated.

Feature Extraction: In this task, we evaluate whether sparsely-sampled light curves can be used to recover the underlying physical parameters from the Goldstein simulation grid. We encode each light curve and fit a series of small MLPs to regress the simulation parameters in Table 1. As a baseline, we train an end-to-end transceiver encoder to regress the model parameters directly from both photometry and spectra. Our results can be found in Appendix B.3.



Figure 4. Sample reconstruction of a spectrum 10 days after supernova peak light by our MMVAE (right, 'LC2spec'), and conditioned on an input light curve (left). Baseline spectra are shown at right for the training average at +10 days in green and the nearest spectrum recovered by the contrastive model in orange (see text for details).

4. Discussion

We have introduced an MMVAE trained to learn joint embeddings between supernova modalities and perform probabilistic cross-modality generation. Our proposed transceiver encoder and decoder has independent value beyond the MMVAE, and can be used to train alternative multi-modal architectures like flamingo (Alayrac et al., 2022) for representation learning of time-domain datasets.

The model presented can be used to inform real-time followup of supernovae. By producing posteriors over spectra conditioned on photometry, it allows astronomers to rapidly infer an explosion's general physical properties and to triage follow-up targets – either to confirm predicted spectral features or collect data to improve future predictions.

Beyond reconstruction, the model's latent features can also serve as inputs to lightweight, task-specific models. Empirically, our MMVAE can reconstruct data modalities more faithfully than a similar model trained with a constrastive objective. As a result, our approach provides a promising alternative for training multi-modal "foundation models" in astrophysics (e.g., Parker et al., 2024; Zhang et al., 2024; Rizhko & Bloom, 2025). The intermediate cross-modal reconstructions can also be used for mechanistic interpretability studies, by linking predicted spectral features to the output of downstream models optimized for classification and redshift prediction.

In future work, we will incorporate host-galaxy images of the observed supernova as a third modality and extend our architecture to explicitly enforce public and private information using, e.g., MMVAE+ (Palumbo et al., 2023; Märtens & Yau, 2024) or diffusion decoders (Palumbo et al., 2024). We also aim to improve the calibration of our output samples. Finally, we plan to fine-tune the model for common downstream tasks in transient science and benchmark performance against existing models (e.g., ParSNIP for classification, Boone 2021; SALT3 for joint fitting of photometry and spectra, Kenworthy et al. 2021).

³The nearest neighbor method does not provide uncertainty estimates, so we exclude it from our CI-based comparisons.

Acknowledgements

We are grateful to the anonymous reviewers of the ICMLcolocated ML4Astro 2025 workshop for their constructive feedback. This work is supported by the National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, http://iaifi.org/).

Impact Statement

The proposed method can be used to learn joint representations between multiple modalities while keeping a probabilistic formalization. The method can likely be generalized to other datasets with covariate modalities.

References

- Alayrac, J.-B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al. Flamingo: a visual language model for fewshot learning. *Advances in neural information processing* systems, 35:23716–23736, 2022.
- Bellm, E. C., Kulkarni, S. R., Graham, M. J., Dekany, R., Smith, R. M., Riddle, R., Masci, F. J., Helou, G., Prince, T. A., Adams, S. M., et al. The zwicky transient facility: system overview, performance, and first results. *Publications of the Astronomical Society of the Pacific*, 131(995): 018002, 2018.
- Boone, K. ParSNIP: Generative Models of Transient Light Curves with Physics-enabled Deep Learning. , 162(6): 275, December 2021. doi: 10.3847/1538-3881/ac2a2d.
- DES Collaboration, Abbott, T. M. C., Acevedo, M., Aguena, M., Alarcon, A., Allam, S., Alves, O., Amon, A., Andrade-Oliveira, F., Annis, J., Armstrong, P., Asorey, J., Avila, S., Bacon, D., Bassett, B. A., Bechtol, K., Bernardinelli, P. H., Bernstein, G. M., Bertin, E., Blazek, J., Bocquet, S., Brooks, D., Brout, D., Buckley-Geer, E., Burke, D. L., Camacho, H., Camilleri, R., Campos, A., Carnero Rosell, A., Carollo, D., Carr, A., Carretero, J., Castander, F. J., Cawthon, R., Chang, C., Chen, R., Choi, A., Conselice, C., Costanzi, M., da Costa, L. N., Crocce, M., Davis, T. M., DePoy, D. L., Desai, S., Diehl, H. T., Dixon, M., Dodelson, S., Doel, P., Doux, C., Drlica-Wagner, A., Elvin-Poole, J., Everett, S., Ferrero, I., Ferté, A., Flaugher, B., Foley, R. J., Fosalba, P., Friedel, D., Frieman, J., Frohmaier, C., Galbany, L., García-Bellido, J., Gatti, M., Gaztanaga, E., Giannini, G., Glazebrook, K., Graur, O., Gruen, D., Gruendl, R. A., Gutierrez, G., Hartley, W. G., Herner, K., Hinton, S. R., Hollowood, D. L., Honscheid, K., Huterer, D., Jain, B., James, D. J., Jeffrey, N., Kasai, E., Kelsey, L., Kent, S., Kessler, R., Kim, A. G., Kirshner, R. P., Kovacs, E., Kuehn, K., Lahav, O., Lee, J.,

Lee, S., Lewis, G. F., Li, T. S., Lidman, C., Lin, H., Malik, U., Marshall, J. L., Martini, P., Mena-Fernández, J., Menanteau, F., Miquel, R., Mohr, J. J., Mould, J., Muir, J., Möller, A., Neilsen, E., Nichol, R. C., Nugent, P., Ogando, R. L. C., Palmese, A., Pan, Y. C., Paterno, M., Percival, W. J., Pereira, M. E. S., Pieres, A., Malagón, A. A. P., Popovic, B., Porredon, A., Prat, J., Qu, H., Raveri, M., Rodríguez-Monroy, M., Romer, A. K., Roodman, A., Rose, B., Sako, M., Sanchez, E., Sanchez Cid, D., Schubnell, M., Scolnic, D., Sevilla-Noarbe, I., Shah, P., Smith, J. A., Smith, M., Soares-Santos, M., Suchyta, E., Sullivan, M., Suntzeff, N., Swanson, M. E. C., Sánchez, B. O., Tarle, G., Taylor, G., Thomas, D., To, C., Toy, M., Troxel, M. A., Tucker, B. E., Tucker, D. L., Uddin, S. A., Vincenzi, M., Walker, A. R., Weaverdyck, N., Wechsler, R. H., Weller, J., Wester, W., Wiseman, P., Yamamoto, M., Yuan, F., Zhang, B., and Zhang, Y. The Dark Energy Survey: Cosmology Results with ~1500 New High-redshift Type Ia Supernovae Using the Full 5 yr Data Set., 973(1): L14, September 2024. doi: 10.3847/2041-8213/ad6f9f.

- Goldstein, D. A. and Kasen, D. Evidence for subchandrasekhar mass type ia supernovae from an extensive survey of radiative transfer models. *The Astrophysical Journal Letters*, 852(2):L33, 2018.
- Hoyle, F. and Fowler, W. A. Nucleosynthesis in supernovae. Astrophysical Journal, vol. 132, p. 565, 132:565, 1960.
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., Abel, B., Acosta, E., Allsman, R., Alonso, D., AlSayyad, Y., Anderson, S. F., Andrew, J., et al. Lsst: from science drivers to reference design and anticipated data products. *The Astrophysical Journal*, 873(2):111, 2019.
- Jaegle, A., Borgeaud, S., Alayrac, J.-B., Doersch, C., Ionescu, C., Ding, D., Koppula, S., Zoran, D., Brock, A., Shelhamer, E., et al. Perceiver io: A general architecture for structured inputs & outputs. arXiv preprint arXiv:2107.14795, 2021a.
- Jaegle, A., Gimeno, F., Brock, A., Vinyals, O., Zisserman, A., and Carreira, J. Perceiver: General perception with iterative attention. In *International conference on machine learning*, pp. 4651–4664. PMLR, 2021b.
- Kasen, D. Secondary maximum in the near-infrared light curves of type ia supernovae. *The Astrophysical Journal*, 649(2):939, 2006.
- Kenworthy, W. D., Jones, D. O., Dai, M., Kessler, R., Scolnic, D., Brout, D., Siebert, M. R., Pierel, J. D. R., Dettman, K. G., Dimitriadis, G., Foley, R. J., Jha, S. W., Pan, Y. C., Riess, A., Rodney, S., and Rojas-Bravo, C. SALT3: An Improved Type Ia Supernova Model for Measuring Cosmic Distances. , 923(2):265, December 2021. doi: 10.3847/1538-4357/ac30d8.

- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., and Krishnan, D. Supervised contrastive learning. *Advances in neural information* processing systems, 33:18661–18673, 2020.
- Liu, Z.-W., Röpke, F. K., and Han, Z. Type ia supernova explosions in binary systems: A review. *Research in Astronomy and Astrophysics*, 23(8):082001, 2023.
- Märtens, K. and Yau, C. Disentangling shared and private latent factors in multimodal variational autoencoders. In *Machine Learning in Computational Biology*, pp. 60–75. PMLR, 2024.
- Palumbo, E., Daunhawer, I., and Vogt, J. E. Mmvae+: Enhancing the generative quality of multimodal vaes without compromises. In *The Eleventh International Conference* on Learning Representations. OpenReview, 2023.
- Palumbo, E., Manduchi, L., Laguna, S., Chopard, D., and Vogt, J. E. Deep generative clustering with multimodal diffusion variational autoencoders. In *The Twelfth International Conference on Learning Representations*, 2024.
- Parker, L., Lanusse, F., Golkar, S., Sarra, L., Cranmer, M., Bietti, A., Eickenberg, M., Krawezik, G., McCabe, M., Morel, R., Ohana, R., Pettee, M., Régaldo-Saint Blancard, B., Cho, K., Ho, S., and Polymathic AI Collaboration. AstroCLIP: a cross-modal foundation model for galaxies. , 531(4):4990–5011, July 2024. doi: 10.1093/mnras/ stae1450.
- Peebles, W. and Xie, S. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.
- Perlmutter, S., Aldering, G., Goldhaber, G., Knop, R. A., Nugent, P., Castro, P. G., Deustua, S., Fabbro, S., Goobar, A., Groom, D. E., et al. Measurements of ω and λ from 42 high-redshift supernovae. *The Astrophysical Journal*, 517(2):565, 1999.
- Phillips, M., Lira, P., Suntzeff, N. B., Schommer, R., Hamuy, M., and Maza, J. The reddening-free decline rate versus luminosity relationship for type ia supernovae. *The Astronomical Journal*, 118(4):1766, 1999.
- Phillips, M. M. The absolute magnitudes of type ia supernovae. Astrophysical Journal, Part 2-Letters (ISSN 0004-637X), vol. 413, no. 2, p. L105-L108., 413:L105–L108, 1993.
- Pskovskii, I. P. Light curves, color curves, and expansion velocity of type i supernovae as functions of the rate of brightness decline. Soviet Astronomy, vol. 21, Nov.-Dec. 1977, p. 675-682. Translation. Astronomicheskii Zhurnal, vol. 54, Nov.-Dec. 1977, p. 1188-1201., 21:675– 682, 1977.

- Riess, A. G., Filippenko, A. V., Challis, P., Clocchiatti, A., Diercks, A., Garnavich, P. M., Gilliland, R. L., Hogan, C. J., Jha, S., Kirshner, R. P., et al. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *The astronomical journal*, 116 (3):1009, 1998.
- Rizhko, M. and Bloom, J. S. AstroM³: A Self-supervised Multimodal Model for Astronomy., 170(1):28, July 2025. doi: 10.3847/1538-3881/adcbad.
- Shen, Y. and Gagliano, A. T. Variational diffusion transformers for conditional sampling of supernovae spectra. *arXiv preprint arXiv:2505.03063*, 2025.
- Shi, Y., Paige, B., Torr, P., et al. Variational mixture-ofexperts autoencoders for multi-modal deep generative models. Advances in neural information processing systems, 32, 2019.
- Wang, Q., Rest, A., Dimitriadis, G., Ridden-Harper, R., Siebert, M. R., Magee, M., Angus, C. R., Auchettl, K., Davis, K. W., Foley, R. J., Fox, O. D., Gomez, S., Jencson, J. E., Jones, D. O., Kilpatrick, C. D., Pierel, J. D. R., Piro, A. L., Polin, A., Politsch, C. A., Rojas-Bravo, C., Shahbandeh, M., Villar, V. A., Zenati, Y., Ashall, C., Chambers, K. C., Coulter, D. A., de Boer, T., DiLullo, N., Gall, C., Gao, H., Hsiao, E. Y., Huber, M. E., Izzo, L., Khetan, N., LeBaron, N., Magnier, E. A., Mandel, K. S., McGill, P., Miao, H.-Y., Pan, Y.-C., Stevens, C. P., Swift, J. J., Taggart, K., and Yang, G. Flight of the Bumblebee: the Early Excess Flux of Type Ia Supernova 2023bee Revealed by TESS, Swift, and Young Supernova Experiment Observations. *The Astrophysical Journal*, 962(1): 17, February 2024. doi: 10.3847/1538-4357/ad0edb.
- Zhang, G., Helfer, T., Gagliano, A. T., Mishra-Sharma, S., and Villar, V. A. Maven: a multimodal foundation model for supernova science. *Machine Learning: Science and Technology*, 5(4):045069, 2024.

A. Simulation Parameters

We present the physical parameters associated with the Goldstein & Kasen (2018) simulation grid in Table 1.

Parameter	Symbol	Range
Kinetic energy	E_k	$[4.514, 5.144] \times 10^{51} \mathrm{erg}$
Ejecta Mass	$m_{\rm ej}$	$[0.7,2.5]M_{\odot}$
⁵⁶ Ni Mass	$m_{ m Ni}$	$[0.75, 2.16] M_{\odot}$
IME Mass	$m_{\rm IME}$	$[0.0036, 2.07] M_{\odot}$
Unburnt C/O Mass	$m_{\rm CO}$	$[0,0.17]M_\odot$
Progenitor Mass	m	$[1.0, 2.5] M_{\odot}$

Table 1. Physical parameters associated with the SN Ia simulation grid used in this work.

B. Additional Results

B.1. Prior sampling

We can sample the latent prior and decode directly into spectra to inspect the model's diversity in generation. We show 50 prior samples of spectra at phase +10 days relative to peak light in Figure 5. The model's prior qualitatively matches the test sample but retains less high-frequency structure.



Figure 5. Spectral samples from the model prior (left) and after conditioning on photometry for events in the test set (right). Our model captures key spectroscopic diversity in SNe Ia e.g., the blended absorption near 4200 Å, and this diversity is enhanced by the inclusion of cross-modality information.

B.2. Masking

In this experiment, we randomly mask 0%, 10%, 30%, 50%, 70%, and 90% of the conditioned light curves and attempt to reconstruct a spectrum. An example is shown in Figure 6. Our method shows some robustness to masking until nearly 70% of photometric observations are masked, beyond which model predictions appear to collapse to the training-set mean spectrum.

Next, we more directly assess whether our model reverts to predicting an average SN Ia spectrum when conditioned on less photometric information. In Figure 7, we show the absolute difference between 1) the model prediction and an average SN Ia spectrum from our grid, 2) the model prediction and the ground-truth spectra, and 3) the average and ground-truth spectra. We observe that as the majority of observations are masked, the difference between predicted and average spectra systematically decreases. This confirms that the model regresses toward the prior when conditioned on less information.





MMVAE-generated spectra (blue), the average SN Ia spectrum in the training set at each phase (green), and the nearest-neighbor spectra from the contrastive model (orange). Spectra are shown in units of erg s⁻¹ cm⁻² Hz⁻¹ and photometry is shown in units of absolute magnitudes. Figure 6. Sample spectra reconstructions (right columns) obtained by conditioning on photometry at different levels of masking (left column). Rows correspond to 0, 10, 30, 50, 70, and 90% masking fractions and right columns correspond to spectra at phases of -10, 0, +10, +20, and +30 days relative to peak luminosity. Right panels show ground truth spectra (red),





B.3. Feature extraction

The baseline end2end model is a simple encoder following the texttttransceiver architecture, similar to the encoder used in MMVAE and the contrastively-trained model.

We train identical MLPs to predict each physical parameter in Table 1 using two-thirds of the test set. We evaluate on the remaining third. Our results are shown in Figure 8.

Recovering most physical parameters from photometry alone proves challenging, and most methods perform only slightly better than a training-set average. This may be the result of the sparse sampling of the conditioned LSST light curves. Parameters closely tied to brightness —such as 56 Ni mass — are better predicted. Both the MMVAE and the contrastively-trained model outperform the end-to-end baseline on most parameters, suggesting that a structured latent representation is better able to preserve physical information.



Figure 8. Mean and 1σ standard deviations of absolute residuals for recovering six parameters from Goldstein & Kasen (2018)'s simulation grid. Both MMVAE and contrastive models outperform end-to-end regression on most parameters, but with large scatter.