

# LEARNING CONTINUOUS MORPHOLOGICAL TRAJECTORIES VIA LATENT PRINCIPAL CURVES

**Sebastian Magana**

University of British Columbia  
Department of Mathematics  
Vancouver, BC, Canada  
maganaaa@student.ubc.ca

**Wenjun Zhao**

Wake Forest University  
Department of Mathematics  
Winston-Salem, NC, USA  
zhaow@wfu.edu

**Khanh Dao Duc**

University of British Columbia  
Department of Mathematics  
Vancouver, BC, Canada  
khanh.daoduc@ubc.ca

## ABSTRACT

Inferring continuous morphological transformations from collections of static biological snapshots is an important, yet challenging problem. In the context of cellular biology, prevailing approaches reduce 3D shape collections to static reconstructions or hand-crafted descriptors, which fail to capture smooth, multidimensional transitions. We present **MorphCurveVAE**, a two-stage pipeline for constructing continuous morphological trajectories from sets of static, segmented 3D microscopy images. Stage 1 learns a smooth, compact latent manifold of volumetric morphologies using a multi-branch convolutional variational auto-encoder (VAE) that can encode multiple correlated substructures into disentangled subspaces. Stage 2 extracts a constrained, topologically-aware principal curve through the augmented latent space to produce directional and correlated trajectories of structural dynamics. To demonstrate our framework, we apply MORPHCURVEVAE to a large public dataset (Allen Institute WTC-11) of segmented volumetric cell and nucleus images spanning the mitosis cycle. Our results indicate high-quality reconstructions, low projection errors to the fitted principal curve, and biologically and visually plausible continuous animations. These results suggest MORPHCURVEVAE as a practical tool for modeling biological morphological trajectories, while remaining broadly applicable to other biological imaging domains where time-resolved observations are unavailable.

## 1 INTRODUCTION

Modeling continuous morphological changes is a cornerstone of many scientific domains, including developmental biology, cellular simulations, and morphometrics (Donovan-Maiye et al., 2022; Winn-Nuñez et al., 2024; Mitteroecker and Schaefer, 2022). Large repositories of segmented 3D images provide rich snapshot datasets, but converting such collections into continuous trajectories is not straightforward (Viana et al., 2023; Ulman et al., 2017). Many prior pipelines focus on per-sample reconstruction or rely on hand-crafted shape descriptors (e.g., volume, sphericity, elongation) that do not naturally capture smooth, multi-dimensional transitions (Ruan and Murphy, 2019; Kazhdan et al., 2003). Trajectory-inference methods from genomics highlight the importance of topology-aware ordering (Saelens et al., 2019), yet existing 3D generative shape models do not provide analogous capabilities for morphological datasets, leaving continuous transitions underexplored. Moreover, recent cell culture modeling work continues to often poorly handle or completely omit transitional events in shape dynamics (Dixon et al., 2025).

In this paper we propose a general pipeline with two modular components: (i) a multi-branch VAE for learning smooth latent embeddings of volumetric morphologies, and (ii) a supervised principal-curve extractor that enforces ordered traversal through phase centroids while regularizing geometric arc properties. The pipeline, illustrated in Figure 1, is intentionally general, however we evaluate our method on a large public dataset of live human iPSC cells collected and released by the Allen Institute (WTC-11 dataset) which provides volumetric masks for nucleus and cell membranes along with hand annotations for mitotic phase on a subset of cells (Viana et al., 2023). These same components apply to any domain where multiple correlated structures morphologically change together (e.g., multi-modal medical imaging, evolving organoids, time-lapse morphological studies).

## 2 RELATED WORK

Learning representation approaches of 3D morphology include voxel- and mesh-based convolutional models, implicit-function methods, and diffusion-based generative families. Voxel-based VAEs remain attractive as they directly support segmented masks, allow probabilistic decoding, and integrate smoothly with geometric constraining (Wu et al., 2015; Ruan and Murphy, 2019). Implicit approaches like DeepSDF offer compact representations but require specialized decoders (Park et al., 2019). Diffusion, and other likelihood-free models, can achieve high geometric fidelity, including for microscopy morphologies, but demand significant data, compute, and complex conditioning to enforce constraints (Waibel et al., 2022; Chu et al., 2023; Winn-Nuñez et al., 2024). Our choice of a voxel-based VAE reflects a pragmatic trade-off: interpretable, extensible, constrainable for trajectory extraction.

A wide array of trajectory-inference (TI) methods exist, from diffusion pseudotime (DPT) (Haghverdi et al., 2016) to reversed-graph embedding (Monocle 2) (Qiu et al., 2017) and cluster-then-curve approaches like Slingshot (Street et al., 2018). Benchmark studies highlight that these methods are highly sensitive to latent topology, with cyclic trajectories posing persistent challenges (Saelens et al., 2019). Notably, most TI algorithms are not designed for image-derived morphological masks, and the cyclic nature of many biological processes exacerbates their limitations. These factors motivate our supervised principal-curve extraction, which incorporates temporal and geometric constraints to recover cyclic morphological trajectories.

Recent large curated datasets, such as WTC-11, now enable systematic evaluation of reconstruction accuracy, and trajectory extraction (Viana et al., 2023). Recent literature also highlights challenges in culture-scale modeling, such as maintaining cell identity across divisions or accounting for mitosis in lineage tracking (Dixon et al., 2025). Our pipeline complements this ecosystem by providing a constrainable, reproducible method for extracting smooth, biologically meaningful trajectories from image-derived latent spaces.

## 3 METHODS

The MORPHCURVEVAE pipeline, represented in Figure 1, consists of two core stages: (i) learning a compact latent manifold of volumetric morphologies, and (ii) extracting a supervised, principal curve through an augmented latent space, enabling the decoding of continuous animations.

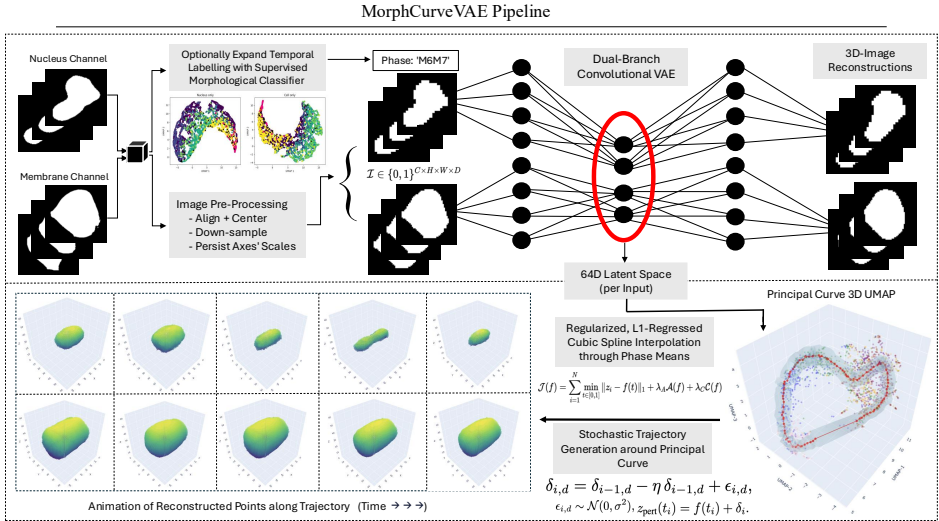


Figure 1: (top) multi-branch VAE encodes correlated channels into compact latent space; (bottom right) supervised principal-curve extractor fits an ordered latent trajectory; (bottom left) decoder projects trajectory near principal curve back to shape space, yielding reconstructed animation.

### 3.1 PREPROCESSING AND LABEL REFINEMENT

**Image preprocessing.** The pipeline accepts volumetric segmented masks (binary or probabilistic) for one or more correlated substructures. The minimal preprocessing performed by the pipeline includes (i) centering the masks and aligning the structures with ICP rigid alignment if not using pre-aligned data, and (ii) downsampling images to a fixed grid of size  $C \times D \times H \times W$  while persisting original scale axes for later augmentation (see Sec. 3.2.1). Note that for our dataset, images were previously both centered and aligned, and all inputs were downsampled to binary masks of size  $2 \times 64 \times 64 \times 64$  with externally saved scale factors  $(s_x, s_y, s_z)$ .

**Label refinement and extension.** The principal-curve stage requires at least a subset of ordered temporal labels, and benefits greatly from a large set of high quality labels. When full annotation coverage is absent or mislabeling is suspected, we perform a light-weight label-refinement procedure prior to VAE training (details and code in Appendix A). In brief, we compute and utilize a morphological image signature for supervised classification both on the unlabeled data, and for internally assessing the quality of the given labeled data.

### 3.2 MULTI-BRANCH VARIATIONAL AUTOENCODER

**Architectural overview.** We adopt a multi-branch convolutional VAE that generalizes to an arbitrary number of correlated input channels. Each branch processes one input channel through a stack of 3D convolutional encoding blocks; branch embeddings are concatenated and then mapped to Gaussian latent parameters. Branches encode channels independently, however, joint latent-space vectors are captured prior to the mirrored decoding reconstructions.

**Configuration.** After an extensive grid-search of hyper-parameter candidate values using validation-based model selection (see Additional Materials), our instantiation of the VAE was then trained with the following specifications:

- **Input format:**  $\mathcal{I} \in \{0, 1\}^{2 \times 64 \times 64 \times 64}$  (channel 0: nucleus mask; channel 1: cell mask).
- **Encoders:** three encoding blocks, each consisting of Conv3D  $\rightarrow$  GroupNorm  $\rightarrow$  LeakyReLU  $\rightarrow$  downsample (stride-2) with progressive channel growth (16, 32, 64 filters), were used in each branch.
- **Latent:** concatenated branch embeddings were mapped to Gaussian parameters  $(\mu_\phi(x), \sigma_\phi^2(x)) \in \mathbb{R}^D$ ,  $D = 64$ ; using an explicit 32/32 allocation for the two branches.
- **Decoder:** three transpose-convolutional blocks mirroring the encoder, outputting per-voxel logits for each channel, then converted to probabilities by a sigmoid before thresholding.

**Probabilistic model and loss.** The encoder defines an approximate posterior

$$q_\phi(z | x) = \mathcal{N}(z; \mu_\phi(x), \text{diag}(\sigma_\phi^2(x))), \quad (1)$$

and the decoder defines  $p_\theta(x|z)$  producing per-voxel probabilities  $\hat{x}_v$ . We minimize the  $\beta$ -VAE evidence lower bound:

$$\mathcal{L}_\beta(\theta, \phi; x) = \mathbb{E}_{q_\phi(z|x)} \left[ - \sum_v \left( x_v \log \hat{x}_v + (1 - x_v) \log(1 - \hat{x}_v) \right) \right] + \beta \text{KL}(q_\phi(z | x) \| p(z)), \quad (2)$$

with  $p(z) = \mathcal{N}(0, I)$ . The binary cross-entropy reconstruction term is suitable for binary masks. At inference we optionally threshold  $\hat{x}_v$  at  $\tau$  to obtain binary masks (default  $\tau = 0.5$ , however adjusting  $\tau$  can improve visual quality in heavily imbalanced volumes where the total voxel set is majority false, or majority true).

#### 3.2.1 LATENT SUBSPACE DESIGN AND SCALE AUGMENTATION

To preserve interpretable substructures for the individual channels’ vectors, we allocate contiguous sub-blocks of the latent vector to each branch (demo: 32 dims per branch). To preserve absolute size

information lost by downsampling, we compute axis-wise scale factors  $(s_x, s_y, s_z)$  for each sample and append them to the core latent  $z$  to form an augmented latent

$$\tilde{z} = [z; s_x; s_y; s_z] \in \mathbb{R}^{D+3}. \quad (3)$$

During decoding we ignore the appended scale coordinates when passing to  $p_\theta(\cdot | \cdot)$ , and after decoding we rescale the reconstructed volume by  $(s_x, s_y, s_z)$  to approximate original physical proportions. The appended scale coordinates are explicitly included in the principal-curve extraction and are scaled in our implementation to hold twice the variance of the latent space dimensions in order to be weighted more heavily, ensuring absolute size is considered in the trajectory geometry.

### 3.3 SUPERVISED, TOPOLOGY-AWARE PRINCIPAL-CURVE EXTRACTION

After training the VAE, we compute augmented latent vectors  $\{\tilde{z}_i\}_{i=1}^N$  and extract a smooth one-dimensional trajectory  $f : [0, 1] \rightarrow \mathbb{R}^{D'}$  summarizing morphological progression. We begin by computing centroids of the ordered temporal phases,

$$\bar{z}_k = \frac{1}{|I_k|} \sum_{i \in I_k} \tilde{z}_i, \quad (4)$$

where  $I_k$  indexes samples in stage  $k$ . These centroids are enforced as hard equality constraints on the curve, requiring the principal curve to pass through each phase centroid.

We represent the trajectory as a sequence of  $M$  points in latent space, of which the phase centroids are fixed and the remaining are optimized. A cubic spline interpolant  $f(t)$  is then fit through these points, yielding a continuous, closed curve. The optimization minimizes a composite objective,

$$\mathcal{J}(f) = L_{\text{reg}}(f) + \lambda_A \mathcal{A}(f) + \lambda_C \mathcal{C}(f), \quad (5)$$

where the regression term

$$L_{\text{reg}}(f) = \frac{1}{N} \sum_{i=1}^N \min_{t \in [0,1]} \|\tilde{z}_i - f(t)\|_1 \quad (6)$$

penalizes L1-distances to the nearest curve point, the arc-length penalty

$$\mathcal{A}(f) = \int_0^1 \|f'(t)\| dt \quad (7)$$

encourages compact parameterizations, and the curvature penalty

$$\mathcal{C}(f) = \int_0^1 \|f''(t)\|_2^2 dt \quad (8)$$

suppresses oscillations.

The hyperparameters  $\lambda_A$ , and  $\lambda_C$  can be adjusted to balance curve fidelity and regularization. Higher values generally favor more compact and direct curves, resembling optimal transport transformations, whereas lower values allow the curve to explore indirect trajectories between phase centroids, which may be useful if the temporal resolution of discrete phase labels is coarse. In this work, values were selected based on visual validation of the resulting trajectories, though cell IDs across imaged phases would allow for more rigorous evaluation of how well resulting curves recapitulate complete single-cell progression across phases.

We optimize the free curve points directly with L-BFGS-B, then construct  $f$  by cubic spline interpolation through all optimized and fixed points. Initialization is given by piecewise linear interpolation through the centroids, and cyclic closure is enforced by wrapping the final segment back to the first centroid. We found this procedure numerically stable and robust across seeds.

### 3.4 STOCHASTIC TRAJECTORY SAMPLING AND DECODING

**Smooth perturbation process.** To generate realistic variant trajectories, we augment the principal curve  $f(t)$  with smooth stochastic deviations. Rather than sampling from independent Gaussians per stage, we simulate a mean-reverting perturbation process along each latent dimension. For a base curve discretized into  $N$  points  $\{f(t_i)\}_{i=1}^N$ , we construct deviations recursively as

$$\delta_{i,d} = \delta_{i-1,d} - \eta \delta_{i-1,d} + \epsilon_{i,d}, \quad \epsilon_{i,d} \sim \mathcal{N}(0, \sigma^2), \quad (9)$$

for dimension  $d$  and step  $i$ , where  $\eta > 0$  controls the reversion strength toward the unperturbed curve and  $\sigma$  controls the noise scale. The deviations  $\delta_{i,d}$  therefore follow an Ornstein–Uhlenbeck-like process, yielding smooth trajectories that remain close to the original curve. The final sampled trajectory is then obtained by adding the deviations to the base curve:

$$z_{\text{pert}}(t_i) = f(t_i) + \delta_i. \quad (10)$$

This construction produces stochastic trajectories that capture variability while avoiding divergence from the learned manifold. The parameters  $(\sigma, \eta)$  control the magnitude and smoothness of deviations, and were selected to yield trajectories visually consistent with the latent morphology progression.

**Decoding and rescaling.** To produce volumetric frames, we (i) strip appended scale coordinates from  $z_{\text{pert}}(t)$ , (ii) decode the core latent through the trained decoder  $p_{\theta}(\cdot | z)$  to obtain per-voxel probabilities, (iii) optionally threshold at  $\tau$  to obtain binary masks, and (iv) resample/rescale the decoded volume by the corresponding  $(s_x, s_y, s_z)$  to approximate physical dimensions. Resampling is implemented via trilinear interpolation to preserve smooth geometry. By sweeping  $t$  across  $[0, 1]$  (and sampling multiple perturbations per  $t$ ) we synthesize continuous animations that reflect both the mean trajectory and empirically observed phase-wise variability.

## 4 DATA AND APPLICABILITY

To demonstrate the utility of the MORPHCURVEVAE pipeline, we applied it to a public collection (Allen Institute WTC-11; Viana et al. 2023) of segmented volumetric images containing two channels (nucleus and cell membrane). The image dataset also includes comprehensive hand annotations of mitotic phases (5,764 cells labeled across a multi-stage mitotic taxonomy). These labels reflect the sequential stages of mitosis, beginning with  $M0$  (interphase) and progressing through  $M1M2$  (prophase–prometaphase),  $M3$  (metaphase),  $M4M5$  (anaphase–telophase), and  $M6M7$  (cytokinesis, further split into early/half). More details are available from the Allen Institute (Allen Institute for Cell Science, 2023). Categories for non-mitotic or corrupted examples are also included: (blob, wrong, dead). Label counts are:  $M0=2516$ ,  $M1M2=441$ ,  $M3=823$ ,  $M4M5=927$ ,  $M6M7_{\text{early}}=338$ ,  $M6M7_{\text{half}}=561$ , with 20 blob, 35 wrong, and 103 dead. To reduce potential learned imbalance in the VAE, phase label counts can optionally be balanced prior to training.

Although not required for our priorly centered and aligned images, the pipeline does support this vital processing step using ICP-alignment. All images were then downsampled to a fixed resolution of  $2 \times 64 \times 64 \times 64$  (channel, depth, height, width). As downsampling removes absolute size information, we computed and stored an additional three scale factors (one for each physical axis).

Additionally, the pipeline requires at least a subset of labeled examples when input. To improve label quality, we trained a small probabilistic classifier using a light-weight, rotationally-invariant morphological representation (see Appendix A). Pre-annotated cells that were predicted with more than fivefold confidence in favor of a different phase were pruned from the training set, thereby reducing likely mislabels (discarded  $\sim 6\%$  of samples). After this pruning step, the same classifier can be used to extend phase annotations to unlabeled cells if necessary, prior to balancing training splits across annotation bins.

These steps ensure that both the morphological inputs and phase labels are standardized, interpretable, and robust for subsequent representation learning and trajectory extraction. Datasets suitable for this pipeline should provide: (i) per-sample segmented masks for the correlated structures of interest (binary or probabilistic), (ii) a subset of ordinal annotations to supervise trajectory extraction.

## 5 EXPERIMENTS AND RESULTS

Our pipeline evaluation focuses on (i) latent-space organization, (ii) reconstruction fidelity, (iii) principal curve trajectory extraction, and (iv) stochastic trajectory generation.

### 5.1 LATENT-SPACE ORGANIZATION

To evaluate whether the latent embeddings captured biologically meaningful variation, we projected the 64D latent space to 2D using UMAP (McInnes et al., 2018). Figure 2 shows these embeddings colored by (i) mitotic phase, (ii) nuclear volume, (iii) height, and (iv) sphericity. The latent space observably reflects the substantial progression between the continuous temporal stages, while also demonstrating the correlation of different individual structural properties to each phase. These hand-crafted descriptors are insufficient to capture the extensive variability among morphologies, however they can offer valuable and interpretable insights.

Additionally, the separation of the mitotic phases is evidently more pronounced when examining the nuclear latent space rather than that of the cell-body. This observation aligns with biological expectations: the nucleus undergoes relatively discrete structural changes during mitosis, including flattening, condensation, and division, whereas the cell membrane deforms more gradually and continuously and in response to environmental cues.

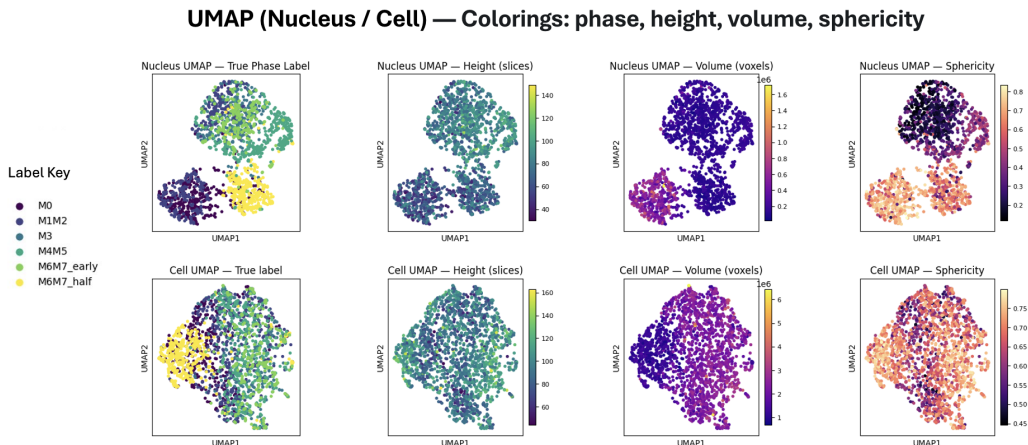


Figure 2: 2D UMAP projection of latent embeddings colored by (A) mitotic phase, (B) height, (C) nuclear volume, and (D) sphericity.

### 5.2 RECONSTRUCTION FIDELITY

We evaluated VAE reconstruction performance using both voxel-wise accuracy and Dice scores for foreground segmentation to mitigate background dominance. Table 1 summarizes these metrics per mitotic phase, demonstrating the VAE’s robustness to diverse morphologies. Figure 3 shows 2D-slices along the central depth axis for several example reconstructions, demonstrating visually accurate recovery of morphological detail across all stages.

Table 1: Reconstruction performance (test set) by mitotic phase.

Phase	Voxel Accuracy (%)	Voxel STD (%)	Dice Score (%)	Dice STD (%)
M0	98.38	0.48	88.99	3.08
M1M2	97.99	0.48	88.41	3.26
M3	98.54	0.32	91.99	1.87
M4M5	98.70	0.32	92.99	1.57
M6M7-early	98.69	0.25	93.35	1.41
M6M7-half	98.81	0.26	92.16	1.81
Overall	98.52	0.46	91.29	2.98




Figure 3: Central slices of several example reconstructions.

### 5.3 PRINCIPAL CURVE TRAJECTORY EXTRACTION

Within the 64-dimensional latent space, we fit a constrained principal curve through the cell embeddings, explicitly enforcing that the curve passes through the phase-wise centroids. This curve captures the dominant morphological progression across mitosis, though a one-dimensional trajectory cannot fully account for all latent variability.

To assess how well the curve represents the embedding, we computed projection distances from each latent point to the curve. Key quantiles are reported in Table 2, with distances falling between 0.27 and 0.47 times the mean latent norm (10th–90th percentiles), substantially tighter than would be expected for a random one-dimensional curve in high dimensionality, which typically lies at roughly the same scale as the full latent norm itself. This contrast indicates that the constrained principal curve successfully captures meaningful low-dimensional structure.

Table 2: Quantiles of latent projection distances to the constrained principal curve, measured in median vector norms.

Quantile	10%	25%	50%	75%	90%
Distance	0.2713	0.32126	0.3793	0.4413	0.4747

Figure 4 shows a 3D UMAP projection of the latent space with cell embeddings colored by mitotic phase, overlaid with the principal curve. Individual points along the curve illustrate reconstructed morphologies, with animations available in the Additional Materials. A tubular envelope around the curve encodes local variance of projections, highlighting regions of higher morphological heterogeneity. Notably, a pronounced gap occurs in the latent space corresponding to cytokinesis, where segmentation captures only one daughter cell, producing a discrete jump within the otherwise smooth, cyclic trajectory. This visualization clearly conveys the principal morphological progression throughout mitosis and the overall cyclic nature captured by the model.

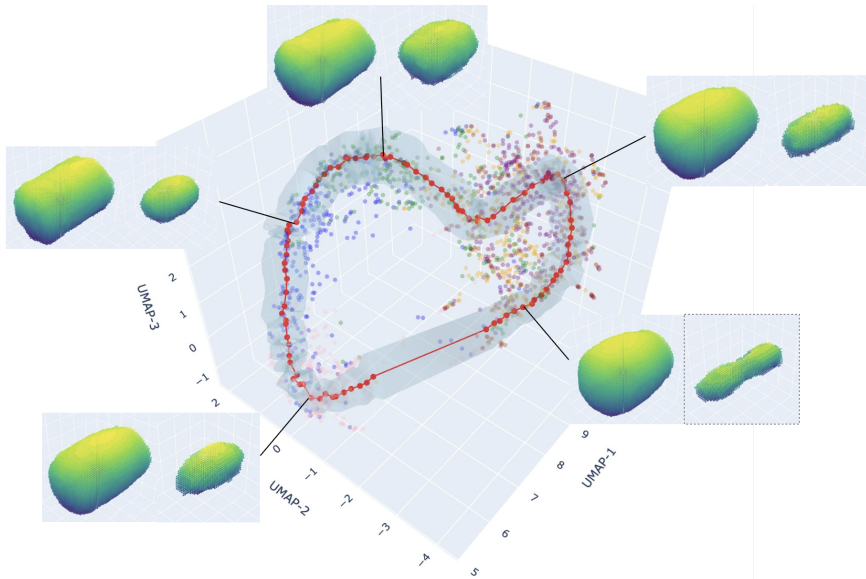


Figure 4: 3D UMAP projection of latent embeddings with the principal curve overlaid as a tube, where thickness represents local variance of projections.

## 6 CONCLUSION

We have presented MORPHCURVEVAE, a reproducible pipeline for transforming collections of static, segmented 3D images into continuous structural trajectories, enabling stochastic generation of plausible morphological populations. Our contribution addresses the lack of general-purpose tools that convert repositories of volumetric shapes into interpretable morphological dynamics with identifiable cycle phases and data-point trajectory positions, thereby extending the ecosystem of generative and trajectory-based modeling methods. Compared with state-of-the-art 3D generative models such as diffusion-based or implicit-function approaches (Waibel et al., 2022; Chu et al., 2023; Winn-Nuñez et al., 2024), MORPHCURVEVAE prioritizes interpretability, constraint enforcement, and computational efficiency.

When evaluated on a relevant dataset, MORPHCURVEVAE produced high-fidelity reconstructions, pronounced latent progression between phases, and visually coherent animations, demonstrating both robustness and practical utility. In summary, MORPHCURVEVAE provides a principled framework for bridging static morphological datasets with temporally structured generative modeling, offering a foundation for future advances in dynamic shape analysis.

Several limitations of the pipeline, as well as its application to the dataset, should be noted. Firstly, our design choice to use voxel-based VAEs incurs a trade-off of weaker geometric fidelity compared to signed-distance or diffusion-based generators, for the sake of interpretability and constrainable flexibility. Pertaining to the dataset (Allen Institute WTC-11), no persisting image IDs exist across temporal phases, obstructing attempts to benchmark accuracies for complete reconstructed trajectories of individual structures. Lastly, any discrete labeling for intrinsically continuous trajectories is noisy and approximate, thus constraining the curve to pass through discrete centroids is imprecise.

## 7 REPRODUCIBILITY STATEMENT

Please see the code repository referenced below:

- **Code Repository:** <https://github.com/anonymous67552/MorphCurveVAE>
- **Additional Materials:** Hyperparameter grid-search logs, extended figures and animations: <https://github.com/anonymous67552/AdditionalMaterials>

The associated README file contains all necessary instructions to reproduce these results.

## 8 MEANINGFULNESS STATEMENT

Many biological processes are inherently continuous, yet are observed only as collections of static snapshots due to experimental constraints. This work contributes a framework for learning smooth, ordered latent representations that recover biologically meaningful progression from such data. By modeling continuous morphological trajectories from segmented microscopy images, our approach enables interpretable analysis of cellular dynamics, including mitotic shape changes, without requiring time-resolved recordings. The framework applies to diverse biological imaging modalities, but a specific motivating problem is modeling mitotic cells, which are often ignored or oversimplified in standard segmentation pipelines and cellular simulations.

### ACKNOWLEDGMENTS

SM was supported by an NSERC Canada USRA fellowship. WZ was supported by a Pacific Institute of Mathematical Sciences (PIMS) Postdoctoral fellowship and PIMS Kantorovich Initiative. KDD was supported by a NSERC Discovery Grant RGPIN-2020-05348. This research was supported in part through the computational resources and services provided by Advanced Research Computing at the University of British Columbia. The authors also thank Felix Zhou and Matheus Viana for early feedback.

### REFERENCES

- M. P. Viana, J. Chen, T. A. Knijnenburg et al. Integrated intracellular organization and its variations in human iPSCs. *Nature*, 613:345–354, 2023. DOI:10.1038/s41586-022-05563-7.
- X. Ruan and R. F. Murphy. Evaluation of methods for generative modeling of cell and nuclear shape. *Bioinformatics*, 35(14):2475–2483, 2019. DOI: 10.1093/bioinformatics/bty983.
- J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 165–174, 2019.
- M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3D shape descriptors. In *Eurographics Symposium on Geometry Processing*, pp. 156–165, 2003. DOI: 10.2312/SGP/SGP03/156-165.
- L. McInnes, J. Healy, and J. Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. arXiv:1802.03426, 2018. DOI: 10.48550/arXiv.1802.03426.
- Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3D ShapeNets: A deep representation for volumetric shapes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1912–1920, 2015.
- J. C. Dixon et al. Colony context and size-dependent compensation mechanisms give rise to variations in nuclear growth trajectories. *Cell Systems*, 16(5):101265, 2025. DOI: 10.1016/j.cels.2025.101265.
- L. Haghverdi, M. Büttner, F. A. Wolf, F. Büttner, and F. J. Theis. Diffusion pseudotime robustly reconstructs lineage branching. *Nature Methods*, 13:845–848, 2016. DOI: 10.1038/nmeth.3971.
- X. Qiu, Q. Mao, Y. Tang, L. Wang, R. Chawla, H. A. Pliner, and C. Trapnell. Reversed graph embedding resolves complex single-cell trajectories. *Nature Methods*, 14:979–982, 2017. DOI: 10.1038/nmeth.4402.
- K. Street, D. Risso, R. B. Fletcher, D. Das, J. Ngai, N. Yosef, E. Purdom, and S. Dudoit. Slingshot: cell lineage and pseudotime inference for single-cell transcriptomics. *BMC Genomics*, 19:477, 2018. DOI: 10.1186/s12864-018-4772-0.
- W. Saelens, R. Cannoodt, H. Todorov, and Y. Saeyns. A comparison of single-cell trajectory inference methods. *Nature Biotechnology*, 37(5):547–554, 2019. DOI: 10.1038/s41587-019-0071-9.

- D. J. E. Waibel, E. Röell, B. Rieck, R. Giryes, and C. Marr. A diffusion model predicts 3D shapes from 2D microscopy images. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 1–5, 2023. DOI: 10.1109/ISBI53787.2023.10230752.
- R. Chu, E. Xie, S. Mo, Z. Li, M. Nießner, C.-W. Fu, and J. Jia. DiffComplete: Diffusion-based generative 3D shape completion. In *Advances in Neural Information Processing Systems*, vol. 36, 2023.
- E. T. Winn-Nuñez, H. Witt, D. Bhaskar, R. Y. Huang, J. S. Reichner, I. Y. Wong, and L. Crawford. Generative modeling of biological shapes and images using a probabilistic  $\alpha$ -shape sampler. *bioRxiv*, 2024.01.09.574919, 2024. DOI: 10.1101/2024.01.09.574919.
- R. M. Donovan-Maiye, J. M. Brown, C. K. Chan, L. Ding, C. Yan, N. Gaudreault, J. A. Theriot, M. M. Maleckar, T. A. Knijnenburg, and G. R. Johnson. A deep generative model of 3D single-cell organization. *PLOS Computational Biology*, 18(1):e1009155, 2022. DOI: 10.1371/journal.pcbi.1009155.
- P. Mitteroecker and K. Schaefer. Thirty years of geometric morphometrics: Achievements, challenges, and the ongoing quest for biological meaningfulness. *American Journal of Biological Anthropology*, 178(Suppl. 74):181–210, 2022. DOI: 10.1002/ajpa.24531.
- V. Ulman, M. Mařka, K. Magnusson et al. An objective comparison of cell-tracking algorithms. *Nature Methods*, 14:1141–1152, 2017. DOI: 10.1038/nmeth.4473.
- W. Zhao, D.J. Sutherland, and K. Dao Duc. Fast and interpretable quantification of biological shape heterogeneity via stratified Wasserstein kernel. *biorxiv*, 2025. DOI: <https://doi.org/10.1101/2025.11.05.686865>
- Allen Institute for Cell Science. Mitotic annotation dataset [Dataset]. Associated publication: M. P. Viana et al., *Nature*, 613:345–354, 2023. DOI: 10.1038/s41586-022-05563-7. Available at [https://open.quiltdata.com/b/allencell/packages/aics/mitotic\\_annotation](https://open.quiltdata.com/b/allencell/packages/aics/mitotic_annotation).

## A APPENDIX: ROTATIONALLY-INVARIANT MORPHOLOGICAL EMBEDDINGS FOR LABEL REFINEMENT

To ensure balanced training splits across mitotic phases, the pipeline requires reliable phase annotations prior to autoencoder training. Because the latent space is not yet available at this stage, we construct a rotationally invariant morphological embedding directly from the original volumetric masks following Zhao et al. (2025). Specifically, surface points are sampled from each binary mask, and all pairwise Wasserstein-like distances between points are computed. From the resulting distance distribution, 100 quantiles are extracted to summarize shape geometry in a way that is robust to outliers. This procedure is performed independently for the nucleus and cell body, after which the two 100-dimensional summaries are concatenated to yield a 200-dimensional embedding vector for each cell. While this embedding has no decoder and cannot reconstruct shapes, it provides a compact, rotationally invariant signature of the morphology. Empirically, these embeddings exhibit meaningful separability across phases in low-dimensional UMAP projections (available in Additional Materials).

We then train a supervised classifier on these embeddings using available phase annotations. The classifier is implemented as a stacked ensemble of diverse base learners, including random forests,  $k$ -nearest neighbors, support vector machines, and gradient-boosted trees, with a random forest meta-learner.

Training is performed in a two-fold procedure (50/50 split), where the ensemble is trained on one half of the labeled data and used to predict the other. Any cell for which the classifier predicted a phase different from the annotation with more than fivefold higher confidence was flagged as a probable mislabel. These samples (approximately 6% of annotated cells) were excluded from training, reducing label noise while retaining the majority of the dataset. After pruning, the same classifier can be applied to extend phase annotations across the entire dataset when desired.

Despite the intrinsic ambiguity of phase boundaries, the ensemble achieved over 90% classification accuracy even without excluding mislabeled or low-quality images (e.g., cells annotated as “dead,” “blob,” or “wrong”). This indicates that the rotationally invariant embedding provides a reliable morphological signature for phase classification and that the pruning criterion effectively filters a small fraction of likely erroneous labels, yielding cleaner supervision for downstream representation learning.