Structural Uncertainty Visualization of Morse Complexes for Time-Varying Data Prediction

Weiran Lyu*
University of Utah

Saumya Gupta†
State University of New York at Stony Brook

Chao Chen‡
State University of New York at Stony Brook

Bei Wang§
University of Utah

Abstract

Scientific simulations are essential for understanding complex physical systems, yet they are often computationally intensive and time-consuming. To address this challenge, researchers increasingly employ deep learning models to generate data efficiently and predict future system states. However, the uncertainty inherent in model outputs can undermine the reliability of these predictions, especially when analyzing structural patterns critical for scientific insight. While most existing approaches estimate uncertainty at the pixel or grid level, characterizing uncertainty in predicted topological structures provides a more intuitive and compact way to capture meaningful changes in the data. In this work, we quantify and visualize the uncertainty of Morse complexes during model prediction. Morse complexes, grounded in Morse theory, are gradient-based topological structures that offer concise abstractions of scalar fields. Given a time-varying scalar field, we use UNet-T, a U-Net-style convolutional architecture, to predict future timesteps. To assess the uncertainty of the resulting topological structures, we introduce MC-U, a joint-estimation graph neural network (GNN) that captures how uncertainty propagates into predicted Morse complexes. We demonstrate our approach on several 2D time-varying scientific datasets, showing that it effectively identifies regions of reduced structural reliability, thereby enhancing both the interpretability and the trustworthiness of the predictions.

Index Terms: Morse complex, uncertainty visualization, scientific machine learning, discrete Morse theory, deep learning, topological method.

1 Introduction

Scientific simulations are indispensable for understanding complex physical systems, but running them over long time horizons is often computationally expensive and time-consuming. To overcome this challenge, researchers have increasingly turned to deep learning models to enable efficient data generation [9, 53–55]. However, predictions from these models inherently involve uncertainty, stemming from factors such as limited training data, model approximation errors, and chaotic dynamics in the underlying system. This underscores the need for an effective uncertainty quantification framework in surrogate modeling.

Uncertainty is commonly characterized using statistical summaries—such as the mean, median, and standard deviation—which are then encoded into the data for visualization and analysis. Most existing approaches focus on point-wise or grid-wise uncertainty estimation, where uncertainty is assessed independently at each spatial location. While point-wise uncertainty quantification is commonly adopted in deep learning-based scientific data predictions, the study of structural uncertainty from a topology-based perspective remains underexplored.

Topological descriptors are fundamental tools in topological data analysis and scientific visualization, offering succinct and robust abstractions of complex datasets. Popular scalar field descriptors include persistence diagrams and barcodes [7, 15, 23], merge and contour trees [5, 8], Reeb graphs [48], and Morse/Morse-Smale complexes [13, 14, 22]. These have been applied across domains such as structural biology, climate science, combustion, neuroscience, physics, chemistry, and ecology for tasks including scalar field comparison [40, 57, 66], data compression [24, 65], feature tracking [39, 64], and ensemble analysis [3, 37]; see [67] for a survey. By providing meaningful abstractions, these descriptors reduce data complexity, enable multiscale hierarchical representations, and support progressive simplification and visual exploration.

Recent studies have integrated uncertainty visualization into topological descriptors to enable more effective analysis of underlying structures. In particular, Gupta et al. explored structural uncertainty of Morse complexes in image segmentation [28]. However, this method was designed for static image segmentation models and is not suitable for scientific visualization, particularly when estimating dynamic data. In scientific visualization, prior work has explored uncertainty in various topological descriptors—such as merge trees [68], contour trees [35, 70], Morse and Morse-Smale complexes [3]. However, few studies have addressed the visualization of uncertainty in topological structures arising from deep learning–based predictions of scientific temporal data.

Motivated by this gap, we investigate the structural uncertainty of Morse complexes in predictions generated by deep learning models for time-varying scientific data. Such data typically consist of a sequence of scalar fields, where the underlying topology evolves over time, often becoming increasingly complex. Predictions of future timesteps—conditioned on current timesteps—may exhibit structural deviations that impact scientific interpretation and decision-making. Specifically, we quantify the structural variations introduced into Morse complexes by the model's prediction.

Contributions. In this paper, we quantify and visualize the structural uncertainty of Morse complexes arising from time-varying data predictions. Our framework consists of two key components: a UNet [49] architecture for forecasting future timesteps in time-varying data, and an uncertainty model called MC-U that quantifies structural uncertainty in the Morse complexes of the predicted scalar fields. Together, these models effectively capture structural uncertainties in the predictions, enabling more interpretable and trustworthy analysis of complex scientific data. Our contributions are as follows:

- We propose UNet-T, an UNet-style convolutional architecture to perform predictions of future timesteps, demonstrating its effectiveness through accurate reconstruction.
- We introduce MC-U, a Morse complex uncertainty quantification model that employs a graph neural network (GNN) [51, 63, 69] to jointly estimate structural uncertainty in predicted scalar fields. The model is grounded in discrete Morse theory and extends a joint estimation framework [28] originally developed for image

*e-mail: wlyu@utah.edu
†e-mail: saumgupta@cs.stonybrook.edu
‡e-mail: chao.chen.1@stonybrook.edu
§e-mail: beiwang@sci.utah.edu

segmentation tasks. MC-U identifies regions of structural stability by analyzing the uncertainty values associated with predicted Morse complex skeletons.
- We demonstrate the effectiveness of MC-U both qualitatively and quantitatively through experiments on various 2D time-varying scientific datasets, in comparison with pixel-wise uncertainty models.

2 Related Work

In this section, we briefly review uncertainty visualization for topological descriptors, with a focus on structure-aware data uncertainty in scalar fields, and recent approaches to uncertainty quantification in deep learning.

Uncertainty visualization of topological descriptors. A substantial body of work has studied uncertainty visualization in topological descriptors for scalar fields. A variety of feature-wise and structure-wise uncertainty visualization methods have been proposed, with a particular focus on critical points, merge trees [5], contour trees [8], Morse and Morse–Smale complexes [42].

Mihai and Westermann [41] introduced a probabilistic framework for assessing the stability of critical points in uncertain scalar-field ensembles. They derive confidence regions by analyzing statistics of gradients and Hessians matrices, confidence intervals for the gradient vector and for the determinant and trace of the Hessian—allowing them to identify locations where critical points are likely to occur and infer their types. Günther et al. [25] introduced the concept of mandatory critical points, identifying regions in 2D uncertain scalar fields where at least one critical point is guaranteed to exist within a mandatory critical region. Favelier et al. [16] introduced persistence-based clustering to ensemble members, followed by the identification of mandatory critical regions to visualize the positional uncertainty of critical points.

To explore the structural variations of contours for noisy data, scientists study uncertainty information that requires effective visualization, enhanced by user interaction and relevant contexts. Wu et al. [62] developed interactive visualization tools for exploring data uncertainty of contour trees. Kraus [35] visualized uncertain structures in contour trees by utilizing grayscale morphology. Zhang et al. [70] proposed an efficient sampling-based Monte Carlo method to study the uncertainty of contour trees arising from terrains. Whitaker et al. [61] introduced contour boxplots for visualization and exploration of ensembles of contours of the underlying functions, providing a generalized method to quantify the uncertainty in a model or simulation process. Yan et al. [68] introduced strategies to compute structural averages of merge trees for uncertainty visualization.

Morse and Morse-Smale complexes have been used to investigate structural variations for scientific ensembles. Athawale et al. [3] proposed statistical summary maps for quantifying structural variations and visualizing the uncertainty of Morse complexes. Their statistical summary maps characterize the uncertain behaviors of gradient flows by capturing variations in destinations of gradient flows, persistence values, and directional changes during persistence simplification. Lan et al. [37] introduced a novel topological skeleton based on Morse complexes to analyze the uncertainty associated with atmospheric rivers.

Uncertainty quantification in deep learning. There has been significant work on uncertainty quantification of deep neural networks [1, 19, 21, 27]. For 2D data in particular, uncertainty estimation is often computed in the context of image-to-image and segmentation tasks, producing per-pixel uncertainty estimates. For image-to-image tasks, methods such as [2] train a neural network to predict heuristic lower and upper bounds, which are used to estimate uncertainty. For segmentation tasks, early work by Kendall and Gal [32] proposed a Bayesian approach that combines MC dropout [20] to estimate model uncertainty, with a learned loss attenuation term to handle data uncertainty. In recent years, generative models have been used to create multiple possible outputs (hypotheses), where the pixel-wise variance across these outputs is treated as uncertainty [36,50]. Following this paradigm, Probabilistic U-Net [34] combines a conditional variational autoencoder [56] with a U-Net [49] architecture to generate different segmentation hypotheses by sampling latent variables. PHiSeg [4] further extends this by introducing latent variables at each level of the U-Net, while [47] using diffusion models to generate multiple hypotheses, leading to more diverse outputs.

From a scientific visualization perspective, researchers have investigated uncertainty quantification in deep learning to support a wide range of applications. Shen et al. [53] studied uncertainty with a flow-based model for the super-resolution task. They performed uncertainty quantification by sampling from the Gaussian latent space. Shen et al. [52] investigated deep learning–based surrogate flow models for scientific applications, with a particular emphasis on uncertainty quantification during data generation.

3 Background

Our framework incorporates several technical ingredients, including Morse complexes, discrete Morse theory, and structural uncertainty estimation.

3.1 Morse Complexes

We focus on the construction of 2D Morse complexes. Let $f : \mathbb{M} \to \mathbb{R}$ be a smooth function defined on a 2D manifold, where ∇f denotes its gradient. A point $x \in \mathbb{M}$ is called a *critical point* if $\nabla f(x) = 0$; otherwise, it is a *regular point*. A critical point is *non-degenerate* if its associated Hessian matrix is non-singular. A function f is a *Morse function* if all of its critical points are *non-degenerate* and the critical points have distinct function values [12, page 128]. At any regular point x, an *integral line* is a maximal path whose tangent vectors align with ∇f [14]. f increases along the integral line and each integral line originates and terminates at critical points, namely local minima, local maxima, or saddles.

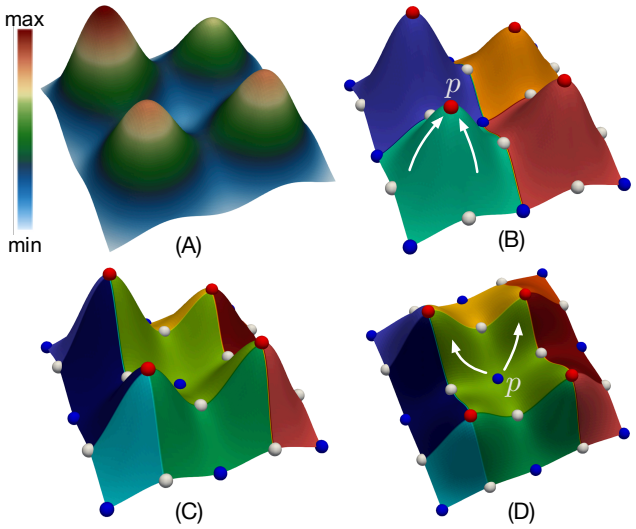

Figure 1: (A) shows a 2D scalar field f. (B) depicts the Morse complex of f, and (C) shows the Morse complex of $-f$. In (B), the stable manifold of a local maximum p is highlighted, along with white integral lines that terminate at p. In contrast, (D) shows the unstable manifold of a local minimum p, with white integral lines originating from p.

The *stable manifold* associated with a local maximum p consists of the point itself and all regular points whose integral lines terminate at p [12]. The *unstable manifold* associated with a local minimum p is the point itself together with all regular points whose integral lines originate at p [12]. A Morse function f is a

Morse-Smale function if the stable and unstable manifolds intersect transversally [12]. Assuming transversality, the stable manifolds induce a decomposition of the domain into 2-cells, where local minima and saddles correspond to 0-cells, and the integral lines connecting them define 1-cells. Together, these cells form the *Morse complex* of f. In other words, the Morse complex segments the domain into regions where gradient flow behaves coherently. Analogously, the unstable manifolds induce the *Morse complex* of $-f$. See Fig. 1 for illustrations.

We focus on the 1D skeleton of the Morse complex—often referred to as the *Morse skeleton* in this paper—which consists of the 0-cells and 1-cells and captures the connectivity between local minima and saddles.

3.2 Discrete Morse Theory

Discrete Morse theory (DMT), introduced by Forman [17], is a combinatorial analogue of classical Morse theory [42] that enables topological analysis of discrete data. In our setting, let K be a 2D cell complex. A function $f : K \to \mathbb{R}$ is a *discrete Morse function* if for each d-cell $\alpha^{(d)} \in K$, there is at most one co-face $\beta^{(d+1)}$ such that $f(\beta) \leq f(\alpha)$, or at most one face $\gamma^{(d-1)}$ such that $f(\gamma) \geq f(\alpha)$, but not both. Cells that do not participate in any such pair are called *critical cells* and they correspond to local minima, local maxima, or saddles.

A *discrete vector field* V is a collection of such pairs $\{\alpha^{(d)}, \beta^{(d+1)}\}$ with $\alpha^{(d)} < \beta^{(d+1)}$, such that each cell appears in at most one pair [18, Definition 3.3]. A *V-path* is an alternating sequence of cells

$$\alpha_0^{(d)}, \beta_0^{(d+1)}, \alpha_1^{(d)}, \beta_1^{(d+1)}, \ldots, \beta_r^{(d+1)}, \alpha_{r+1}^{(d)}$$

where each pair $\{\alpha_i < \beta_i\} \in V$ and $\beta_i > \alpha_{i+1} \neq \alpha_i$, for $0 \leq i \leq r$ [18]. Such a path is a *non-trivial closed V-path* if $r \geq 0$ and $\alpha_0 = \alpha_{r+1}$. When V contains no non-trivial closed V-paths, it is called a *discrete gradient field* [18, Theorem 3.5].

DMT provides a robust framework for extracting Morse (and Morse-Smale) complexes. Notably, it has been used for image segmentation to ensure topological accuracy [28, 31]. In our work, we focus on V-paths that connect saddles to local maxima (or local minima) through sequences of regular cells, which correspond to the Morse skeleton. These paths may be considered as topological analogs of mountain ridges in the scalar field.

3.3 Probabilistic Discrete Morse Theory

DMT, as described above, provides a deterministic framework for computing the Morse skeleton of a discrete Morse function f. However, this deterministic treatment neglects the uncertainty that often underlies f, and consequently overlooks the potential variability in the stable (or unstable) manifolds and the resulting Morse skeleton. In many real-world applications, f is not an exact, noise-free function. For instance, f may represent a likelihood map produced by a segmentation model or a scalar field predicted by a temporal forecasting model. In such cases, f can be ambiguous or noisy due to model limitations or inherent data variability.

To address this uncertainty, Gupta et al. [28] introduced the *Probabilistic DMT* for modeling the variability inherent in segmentation likelihoods from image analysis tasks. Unlike DMT, which treats each 1-cell as a single fixed path, Probabilistic DMT models each 1-cell as a sample drawn from an underlying probability distribution. This probabilistic formulation enables the capture and quantification of uncertainty in segmentation likelihoods, yielding a more robust and realistic representation of the data.

Probabilistic DMT extends DMT by incorporating a *perturb-and-walk* sampling algorithm. Directly sampling from the exponentially large space of all possible Morse skeletons (also referred to as *structures*) is computationally intractable. Hence, the perturb-and-walk strategy is used to efficiently generate diverse, plausible structures of the likelihood without exhaustively enumerating all possible structures [29, 38, 45]. By introducing controlled perturbations to the likelihood and repeatedly extracting pieces of the Morse skeleton, Probabilistic DMT simulates how slight fluctuations in the likelihood can lead to different structures. This stochastic sampling process avoids explicitly evaluating all subsets of structures, while still capturing meaningful variability, making it well-suited for modeling structure-wise uncertainty.

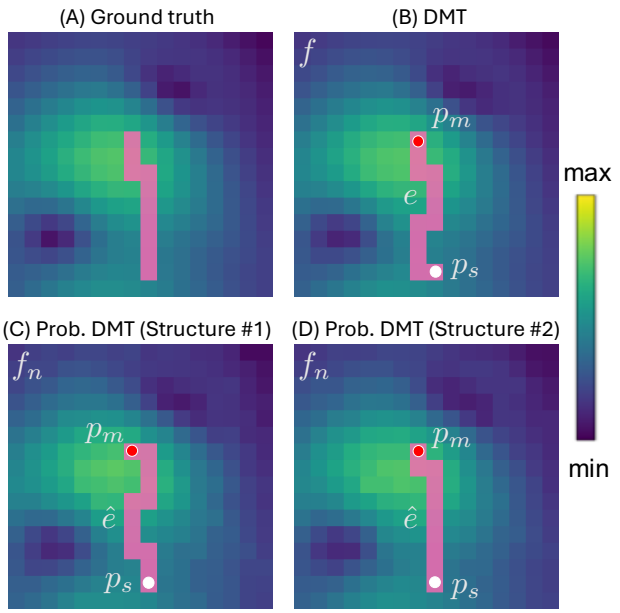

Figure 2: (A) A ground truth scalar field and a piece of a Morse skeleton (i.e., a 1-cell in pink). (B) The likelihood f predicted by a deep learning model, and a 1-cell e (in pink) obtained from DMT. Red and white points denote local maximum p_m and saddle p_s respectively. (C) Probabilistic DMT uses perturb-and-walk strategy to first sample $f_n \sim f + r$, followed by a guided walk to generate structure variation $\hat{e}$ between the same pair of critical points (p_s, p_m). (D) Another structure generated by Probabilistic DMT.

Formally, consider a 1-cell (i.e., a *structure*) e defined by the path connecting a pair of critical points (p_s, p_m) where p_s represents a saddle point and p_m represents a local maximum. In order to obtain a structural variation of e, a perturbed version f_n is first sampled from f by adding random Gaussian noise:

$$f_n \sim f + r, \qquad r \sim \mathcal{N}(0, \sigma^2)$$

This process is independent of the perturbation model r (which is Gaussian). As the variance of the Gaussian model σ is unknown, Bayesian probability theory is used to sample the variance from the Inverse Gamma distribution (its conjugate prior [43]).

Next, to generate a variation $\hat{e}$ of e, a guided walk from p_s to p_m is simulated on the perturbed field f_n. At each step, given the current location p, the next location p'' is chosen as $p'' = \text{argmax}(Q(p'))$, where, $p' \in \text{neighborhood}(p)$[1] and,

$$Q(p') = \gamma Q_d(p') + (1 - \gamma) f_n(p'),$$
$$\text{where } Q_d(p') = \frac{1}{\|p_m - p'\|_2}.$$

And $\gamma \in [0, 1]$ is a balancing parameter. The distance regularizer Q_d ensures path completion towards the destination p_m, while

[1] An 8-connectivity neighborhood for a 2D domain.

$f_n(p')$ enforces preference for high intensity values (as in DMT). Thus, the process begins with $p := p_s$ and continues in this manner $p := p''$ until reaching p_m, thereby generating $\hat{e}$. Fig. 2 illustrates different structures generated by the perturb-and-walk strategy between the same pair of critical endpoints (p_s, p_m) identified by DMT.

As a result, the sampling process generates structural variations that reflect the uncertainty inherent in the likelihood. Structures with high uncertainty give rise to greater variation across samples, whereas structures with low uncertainty (high confidence) yield more consistent outcomes. This variability provides direct estimates of *intra-structural* uncertainty. A complete Probabilistic DMT procedure produces a single sample Morse skeleton from the space of possible skeletons, and multiple realizations enable explicit modeling of structure-wise uncertainty.

3.4 Structural Uncertainty Using Joint Estimation

Probabilistic DMT quantifies intra-structural uncertainty by analyzing variations within individual structures. While intra-structural uncertainty captures ambiguity within a single structure (arising from intrinsic properties such as geometry), the uncertainties of different structures are not independent. Estimating each structure's uncertainty in isolation neglects spatial relationships and may result in inconsistent or suboptimal uncertainty quantification across the entire Morse skeleton.

Therefore, Gupta et al. [28] emphasized that *inter-structural* uncertainty must be considered, as the uncertainty in one structure is often influenced by its neighboring structures. To model these dependencies, a joint estimation framework is used to reason over all structures simultaneously. Rather than predicting uncertainty independently for each structure, a graph is constructed where nodes represent individual structures and edges encode spatial adjacency. A graph neural network (GNN) [51, 63, 69] is then used to propagate information across this graph, enabling the model to capture context and refine uncertainty estimates based on inter-structure relationships.

This joint estimation framework avoids the computational burden of explicitly enumerating all possible Morse skeletons—which grows exponentially with the number of structures—while still accounting for inter-structural dependencies. Through message-passing mechanisms, the uncertainty of each structure is refined using information from its neighbors, producing more coherent and contextually aware estimates that reflect both local structural characteristics and their mutual interactions.

4 Method

Our framework for quantifying structural uncertainty in predicted Morse complexes consists of two components. The first is UNet-T, a U-Net architecture designed to predict the future state of time-varying scientific data. The second is MC-U, which takes the predictions from UNet-T as input and quantifies the uncertainty of the Morse complexes at future timesteps.

4.1 UNet-T: Time-Varying Data Prediction

We are interested in predicting a scalar field $\mathbf{x}_{t+\Delta t}$ at a future timestep $t + \Delta t$ given an input scalar field $\mathbf{x}_t$ at timestep t. To that end, we implement a U-Net [49] style convolutional architecture called UNet-T, which consists of an encoder–decoder architecture with symmetric skip connections that transfer high-resolution features from the encoder to corresponding decoder layers. The encoder compresses the input to a latent representation $\mathbf{z} = \phi(\mathbf{x})$, and the decoder attempts to reconstruct the future timestep from this representation, yielding $\hat{\mathbf{x}} = \psi(\mathbf{z})$. Specifically, the encoder consists of stacked convolutional blocks interleaved with average pooling layers that progressively reduce spatial resolution and increase feature complexity. Each convolutional block contains two convolutional layers followed by ReLU activations, which allow the network to learn rich feature representations.

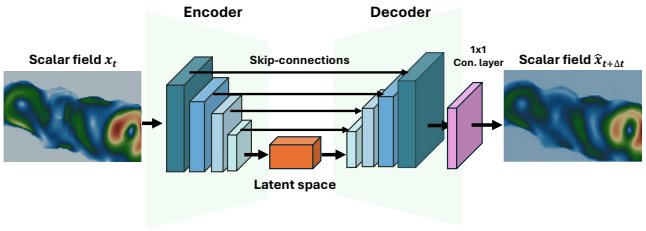

Figure 3: An illustration of the UNet-T model architecture. The model takes as input a scalar field at timestep t and reconstructs the scalar field at timestep $t + \Delta t$.

In the latent space, the network operates at the lowest spatial resolution, capturing global context. The decoder then reconstructs the high-resolution output by upsampling via bilinear interpolation and combining features through skip connections from the corresponding encoder layers. This skip-connection design helps preserve spatial details.

The final output is produced by a 1×1 convolution that maps the decoded features back to the scalar field space and resizes them to match the input dimensions. This architecture effectively predicts future timesteps by learning a compact latent representation of spatiotemporal dynamics while maintaining spatial coherence in the reconstructed data. As shown in Fig. 3, given an input $\mathbf{x}_t$, the model is trained to predict a future state $\hat{\mathbf{x}}_{t+\Delta t}$, learning a compact latent representation that reflects spatial evolution. The input $\mathbf{x}_t$ passes through four convolutional blocks during encoding and decoding. We optimize our model by minimizing a pixel-wise mean squared reconstruction loss, where the ground truth is taken to be the scalar field at future timestep $\mathbf{x}_{t+\Delta t}$. This encourages the decoded output to closely resemble the target future timestep:

$$L_{\text{rec}} = \frac{1}{N} \sum_{i=1}^{N} (\mathbf{x}_{t+\Delta t,i} - \hat{\mathbf{x}}_{t+\Delta t,i})^2$$

where $\mathbf{x}_{t+\Delta t,i}$ and $\hat{\mathbf{x}}_{t+\Delta t,i}$ denote the original and reconstructed scalar field values at spatial location i, respectively, and N is the total number of pixels or grid points in the domain. By training the model to map a field at time t to a future state at time $t + \Delta t$, the model effectively learns to predict spatial field evolution through a compressed representation.

The prediction serves as the input for our MC-U model. The predicted future scalar field $\hat{\mathbf{x}}_{t+\Delta t}$ is subsequently analyzed using a structure-aware GNN designed to estimate the uncertainty of the Morse complex of $\hat{\mathbf{x}}_{t+\Delta t}$.

4.2 MC-U: Uncertainty of Predicted Morse Complexes

The UNet-T predicts future scalar field values by encoding spatial information into compact latent representations, potentially introducing structural variations that pixel-wise uncertainty measures cannot capture. To address this limitation, we introduce a Morse complex uncertainty model, MC-U, designed to detect structural variations in the Morse complexes that arise in the predictions.

Our approach extends the Probabilistic DMT and joint estimation framework of Gupta et al. [28] (see Secs. 3.3 and 3.4) with key modifications to address the time-varying scalar field prediction task. In particular, we design the input feature vector and loss function of MC-U specifically for this task.

An overview of our MC-U model architecture is shown in Fig. 4. Given the prediction at timestep $t + \Delta t$, we first extract the Morse skeleton by computing the 1-cells of the Morse complex from the prediction. The original joint estimation framework [28] was developed for the image segmentation task, where the segmentation

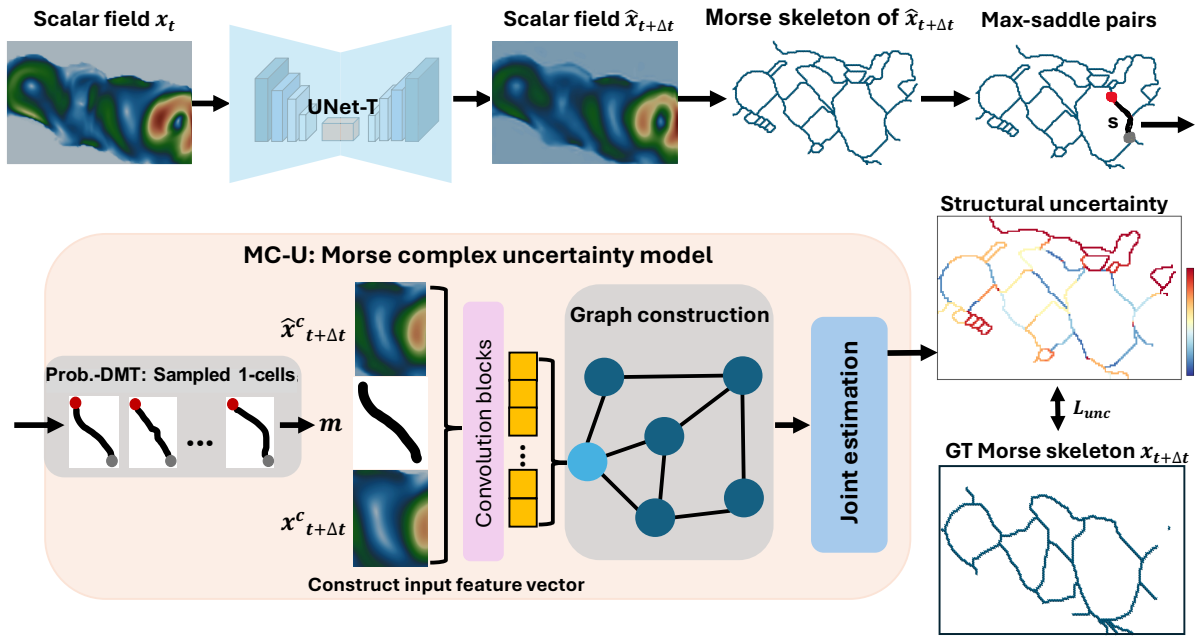

Figure 4: MC-U: Morse complex uncertainty model designed to jointly estimate the uncertainty of the Morse skeleton of the prediction at timestep $\hat{\mathbf{x}}_{t+\Delta t}$. 'GT' stands for the ground truth.

model outputs a continuous likelihood map indicating the per-pixel probability of belonging to the target structure. DMT is then applied to this likelihood map to identify pieces of Morse skeletons that highlight structural inconsistencies—such as false positives and false negatives. In our setting, the UNet-T model predicts a continuous scalar field instead of a probability map. Therefore, we directly compute the Morse skeleton from the predicted scalar field to capture the underlying topological structure.

The Morse skeleton of the prediction consists of a set of 1-cells, each defined by a path connecting a saddle–maximum pair (p_s, p_m). To analyze structural variability within the skeleton, we perturb the predicted scalar field by adding Gaussian noise. For each 1-cell, we fix its endpoints (the saddle–maximum pair) and perform a guided walk starting from the saddle and ascending toward the local maximum. This walk is directed by both the scalar field values and the distance to the destination, following the procedure described in Probabilistic DMT (Sec. 3.3). Through this process, we obtain a collection of sampled 1-cells (paths) that capture possible structural variations of the scalar field. Specifically, if the sampled 1-cells for a given saddle–maximum pair exhibit little variation under perturbations, the 1-cell is associated with low uncertainty; conversely, high variability in the sampled 1-cells indicates greater structural uncertainty.

Joint estimation and training procedure. Probabilistic DMT estimates *intra-structural* uncertainty, that is, uncertainty arising from structural variations internal to each 1-cell. To account for the spatial relationships between the 1-cells, we also model *inter-structural* uncertainty using a GNN. To train the GNN, we modify the input feature vectors to include information from the predicted scalar field. Specifically, for each node in the graph, we construct the input feature vector as

$$[\hat{\mathbf{x}}^c_{t+\Delta t},\ m,\ \mathbf{x}^c_{t+\Delta t}]$$

and then process it through convolutional and pooling layers to obtain a fixed-length vector representation. Here, m denotes the sampled 1-cells generated by the guided walk technique. The term $\hat{\mathbf{x}}^c_{t+\Delta t}$ refers to a small block c extracted from the predicted scalar field $\mathbf{x}_{t+\Delta t}$, centered around the sampled 1-cell m. $\mathbf{x}^c_{t+\Delta t}$ denotes the small block centered around the corresponding locations of the original scalar field $\mathbf{x}_{t+\Delta t}$ at timestep $t+\Delta t$. Note that we take the block from $\mathbf{x}_{t+\Delta t}$, rather than from the original input scalar field $\mathbf{x}_t$ used by the UNet-T, because the uncertainty model MC-U treats timestep $t+\Delta t$ as its input. Similar to Probabilistic DMT, edges between two nodes encode spatial adjacency.

After constructing the graph for GNN, we utilize the joint estimation approach to estimate the uncertainty associated with each 1-cell s of the predicted Morse skeleton. The network head is divided into two components: a mean μ_s representing the predicted probability that s is a true positive, and a variance σ_s^2 representing the uncertainty of s in the prediction. For supervision, we extract the Morse skeleton from the ground truth scalar field at timestep $t+\Delta t$, which allows us to compute the proportion of overlap between a sampled 1-cell m (that shares the same endpoints as s) and the ground truth 1-cell w. This overlap serves as a target measure for training. The training loss function is formulated as

$$L_{unc} = \frac{1}{|S|}\sum_{s\in S}\left(\frac{1}{2\exp(q_s)}\|\mu_s - r_s\|^2 + \frac{1}{2}q_s\right)$$

where S contains a set of 1-cells forming the predicted skeleton, $r_s = \frac{\sum(w\odot m)}{\sum m}$ is the average overlap between the ground truth 1-cell w at timestep $t+\Delta t$ and a set of sampled 1-cell m, $\odot$ denotes the Hadamard product, and $q_s = \log\sigma_s^2$. The variance σ_s^2 is modeled to capture the structural uncertainty inherent in scalar field $\mathbf{x}_{t+\Delta t}$ and prediction uncertainty introduced in UNet-T.

During inference, we perform k runs of the uncertainty model to compute the average uncertainty value for each 1-cell. For each run, we then take the union of the predicted Morse skeletons—represented by $\bigcup\mu_s$—and compute their average across runs. Based on this averaged skeleton, we generate the *structural uncertainty heatmap* by overlaying it with the corresponding average uncertainty values. The final structural uncertainty heatmap is shown in Fig. 4. Each 1-cell in the Morse skeleton is associated with an uncertainty value that reflects the structural variation in the predicted scalar field.

Implementation details. We take the pre-trained UNet-T model and use its predicted future timesteps as input for the MC-U model.

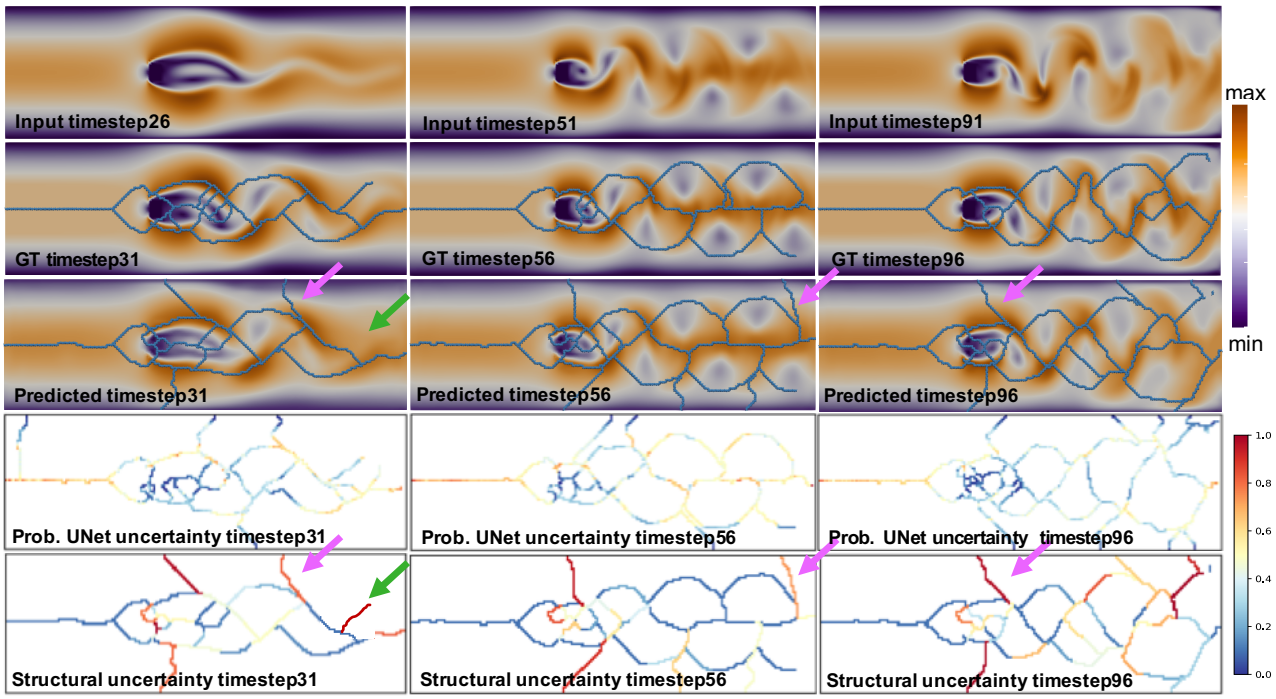

Figure 5: Navier Stokes dataset: Structural uncertainty of the Morse skeleton of predicted scalar field at timesteps 31, 56, and 96. The Morse skeleton is captured in blue. Pink and green arrows highlight false-positives and false-negatives, respectively.

We employ the Topology Toolkit (TTK) [58] to generate the Morse complexes of these predictions. We then extract the connectivity information of the Morse skeletons in advance to perform guided walks within each 1-cell (a saddle-maximum pair). We take 10 runs to generate an average structural uncertainty value for each 1-cell.

5 Experimental Results

We qualitatively and quantitatively evaluate our Morse complex uncertainty quantification framework MC-U for the time-varying scalar field prediction task. Here are our **highlighted results**:

- We demonstrate the effectiveness of MC-U in identifying highly variant structures, which are highlighted with high uncertainty values in the heatmap, while assigning low uncertainty values to more stable regions in the prediction. Compared to the baseline, MC-U is able to clearly capture and distinguish the structural variations present in the predicted Morse complex.
- We quantitatively evaluate our uncertainty heatmaps using a range of evaluation metrics, where MC-U consistently outperforms the baseline. In addition, we report the reconstruction quality of the predicted scalar fields generated by our UNet-T model to further support the validity of our structural uncertainty assessment.

Datasets. We conduct experiments on four 2D time-varying scientific datasets defined on structured grids, as presented in Tab. 1: Heated Cylinder [26] from [46] (available at [10]), Red Sea [30], Navier Stokes [6], and Ionization Front [59]. We apply persistence simplification [15] with a threshold $\epsilon = 0.01$ across all datasets during the preprocessing of extracting Morse complexes, in order to separate features from noise.

Training configurations. The training, validation, and inference samples are randomly chosen and split among all timesteps with a ratio of $0.7/0.1/0.2$ to ensure full coverage of the data. We perform our experiments on a standard laptop with i7 processor with 20 threads running at 3.5 GHz, with 32 GB memory. The UNet-T and MC-U are implemented in PyTorch and trained on a single NVIDIA RTX 2070 Super GPU. We use the Adam optimizer [33] for all datasets. The learning rate for both UNet-T and MC-U is 0.001. In Tab. 2, we provide the total training time for both UNet-T and MC-U, and the inference time for MC-U across all datasets.

Table 1: Detailed descriptions of all datasets.

Dataset	Variable	Dimensions	Timesteps
Heated Cylinder	velocity	150×450	500
Red Sea	scalars	100×150	60
Navier Stokes	speed	64×192	101
Ionization Front	density	248×600	120

An overview of experiments. As the output of the UNet-T model serves as the input to MC-U model, our primary focus is on analyzing the structural variations introduced into the Morse complexes of the predicted scalar fields. In our experimental setup, we use a UNet-T trained to predict the future timestep $t + 5$ given input at timestep t as a single-step prediction. It serves as one of the modeling configurations for generating predictions. While other prediction horizons could be adopted, our emphasis is on evaluating how structures in the Morse complex respond to the prediction process. Additionally, during inference, we perform an evaluation on multi-step predictions by recursively feeding the model's output back as input for further steps without retraining the model, allowing us to study the structural uncertainty of Morse complexes over extended prediction horizons.

We compare with the pixel-wise uncertainty estimation model, Probabilistic U-Net (Prob.UNet in short) [34], which serves as our baseline. Prob.UNet extends the traditional U-Net architecture by incorporating a conditional variational autoencoder framework, allowing it to model the distribution over plausible segmentation outputs given an input image. We adopt Prob.UNet as a baseline because it can be naturally extended to the task of time-varying scalar field prediction, where it captures pixel-wise uncertainty in the predicted fields. In contrast, our proposed external MC-U model focuses on capturing structural uncertainty in the Morse complexes of the predictions. To enable a consistent comparison of uncertainty in the context of Morse complexes, we construct an uncertainty heatmap for Prob.UNet by overlaying its pixel-wise uncertainty estimates onto the Morse skeleton extracted from its prediction.

Table 2: UNet-T training and inference time, together with the MC-U training and inference time. The training time for MC-U is reported in hours (h), whereas the rest is reported in seconds (s). We present the number of epochs trained for MC-U across all datasets.

Dataset	UNet-T training time (s)	Epochs	MC-U training time (h)	MC-U inference time (s)
Heated Cylinder	549.05	100	3.30	0.8169
Red Sea	53.59	300	0.29	0.4062
Navier Stokes	60.19	200	0.31	0.3509
Ionization Front	207.35	200	2.83	1.2435

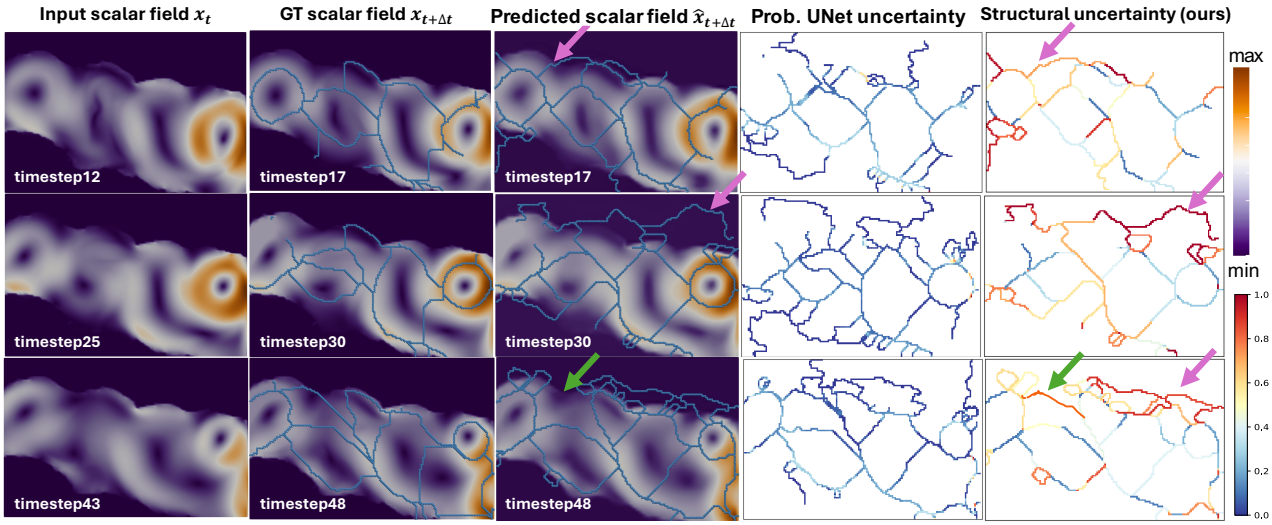

Figure 6: Red Sea dataset: Structural uncertainty visualization of the Morse skeleton of predicted scalar field at timesteps 17, 30, and 48. The Morse skeleton is captured in blue. Pink and green arrows indicate false positives and false negatives, respectively.

5.1 Structural Uncertainty of Morse Complexes

In this section, we present the Morse complex uncertainty of predicted scalar fields using several time-varying scientific datasets.

Navier Stokes dataset. The Navier Stokes dataset comes from a direct numerical Navier Stokes simulation of a viscous 2D flow around a cylinder [6]. We consider the speed as the scalar field. We present the Morse complex uncertainty of the predicted scalar fields at timesteps 31, 56, and 96, given input timesteps 26, 51, and 91, respectively. The Morse skeleton (colored in blue) is overlaid on the ground truth timesteps and the predicted timesteps. As shown in Fig. 5, MC-U captures the 1-cells with large variations in the prediction and assigns low uncertainty values to the stable ones. Red 1-cells in the skeleton represent higher uncertainty, and blue 1-cells indicate lower uncertainty. In the prediction, noise can be introduced during the reconstruction process, particularly in the regions with high gradient flows. Therefore, false-positive 1-cells will be present in the Morse skeleton. On the other hand, existing 1-cells in the original Morse skeleton may be eliminated during the reconstruction, resulting in false-negative ones in the predicted Morse skeleton. MC-U captures both false negatives and false positives, marking them as high-uncertainty 1-cells, as indicated by the pink and green arrows in Fig. 5, respectively.

In contrast, Prob.UNet assigns high uncertainty to relatively flat regions with steady gradient flows, while assigning near-zero uncertainty to both false positives and false negatives—failing to capture the structural variations present in the predicted Morse skeleton.

Red Sea dataset. The Red Sea dataset is an ensemble, which contains 50 members, of time-dependent 2D flow and scalar fields on a regular grid. We employ one ensemble member for the experiment. We present the Morse complex uncertainty of the predicted scalar fields at timesteps 17, 30, and 48, given input timesteps 12, 25, and 43, respectively. The Prob.UNet generates near-zero uncertainty everywhere in the Morse skeleton, as shown in Fig. 6. The uncertainty heatmap of our MC-U model detects the false-positive and false-negative 1-cells with high uncertainty values as highlighted by the arrows.

Heated Cylinder dataset. The Heated Cylinder dataset is a simulation of a 2D flow generated by a heated cylinder using the Boussinesq approximation [46] where the time-varying turbulent plume contains many small vortices that rotate around each other. We take the magnitude of the velocity as the scalar field. We consider 500 timesteps from 600-1100. In Fig. 7, we present the Morse complex uncertainty of the predicted scalar fields at timesteps 845, 1006, and 1084, given input timesteps 840, 1001, and 1079, respectively. The initial predictions from the UNet-T model capture the majority of key features, as evidenced by the preservation of small rotated vortices in the Morse skeleton of the predicted scalar field. MC-U accurately diagnoses false-positive 1-cells by assigning high uncertainty to them, whereas Prob.UNet treats all 1-cells uniformly, failing to differentiate between reliable and spurious features.

Ionization Front dataset. The Ionization Front dataset simulates the propagation of an ionization front instability. The simulation is done with 3D radiation hydrodynamical calculations of ionization front instabilities in which multi-frequency radiative transfer is coupled to the primordial chemistry of eight species [60]. We consider the density from the 2D slices near the center of the simulation volume for timesteps from 11-133. These timesteps show the density over time as the instability progresses toward the right. We present the Morse complex uncertainty of the predicted scalar fields at timesteps 101 and 116, given input timesteps 96 and 111, respectively in Fig. 8. While the UNet-T preserves the main Morse skeletons of the ground truth in the predicted scalar field, there are a few noisy 1-cells present on the boundary. Our structural uncertainty heatmap highlights both the false positives and false negatives with non-zero uncertainty values.

Structural uncertainty in multi-step data predictions. In the Navier Stokes dataset, by extending beyond the structural uncertainty of a single-step prediction at $t+5$, we study the structural evolution of Morse skeletons under multi-step predictions. As shown in Fig. 9, during inference, starting from an initial input at timestep $t = 26$, we recursively feed the model's output as input to predict timesteps $t + 5, t + 10, t + 15$, and $t + 20$, corresponding to

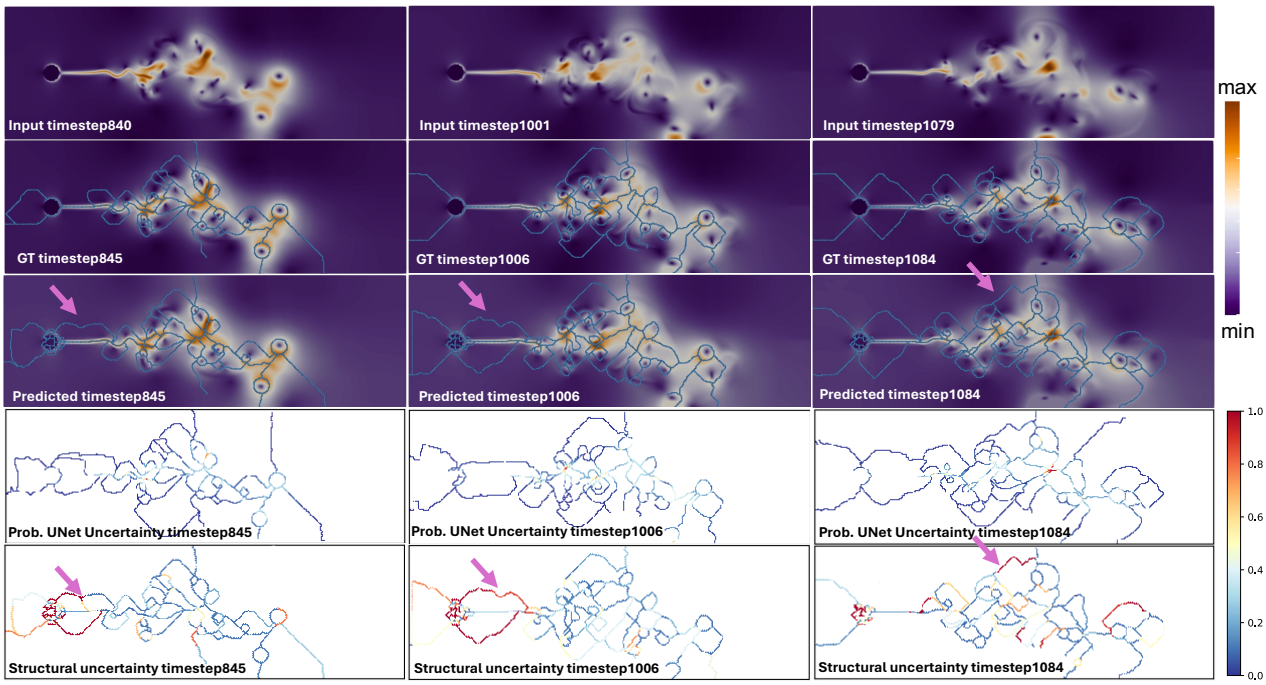

Figure 7: Heated Cylinder dataset: Structural uncertainty visualization of the Morse skeleton of predicted scalar field at timesteps 845, 1006, and 1084. The Morse skeleton is captured in blue. Pink arrows highlight false positives.

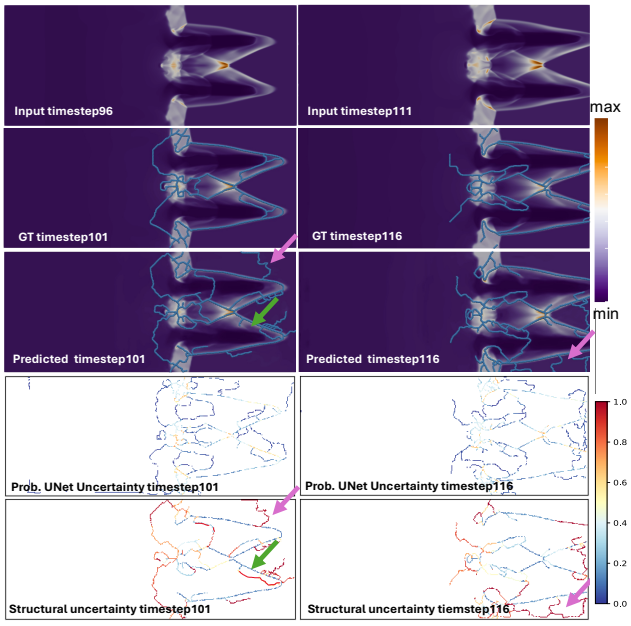

Figure 8: Ionization Front dataset: Structural uncertainty visualization of the Morse skeleton of predicted scalar field at timesteps 101 and 116. The Morse skeleton is captured in blue. Pink arrows highlight false positives and green arrows highlight false negatives.

timesteps 31, 36, 41, and 46. This procedure enables uncertainty quantification at each predicted timestep, allowing us to track how structural uncertainties evolve and accumulate throughout the prediction horizon. While the predictions generated by the UNet-T model largely preserve the major features of the ground truth scalar fields throughout this recursive process, the structural uncertainty heatmaps precisely capture both false positives and false negatives. These uncertainties reflect the increasing complexity of the scalar fields, particularly as vortex structures give rise to more intricate topological structures.

Table 3: Quantitative evaluation of the UNet-T model. We report the average Peak Signal-to-Noise Ratio (PSNR) and mean squared error (MSE) over all test samples for predicted future timesteps across various datasets.

Dataset	MSE	PSNR
Heated Cylinder	0.0002	40.5223
Red Sea	0.0006	35.0206
Navier Stokes	0.0004	36.8684
Ionization Front	0.0003	37.7844

5.2 Quantitative Evaluation

We quantitatively evaluate the uncertainty heatmap of MC-U compared to the uncertainty heatmap of Prob.UNet by utilizing various evaluation metrics. Additionally, we assess the prediction quality of the UNet-T model to further validate the reliability of the structural uncertainty heatmap produced by MC-U.

Evaluation metrics. To evaluate the quality of the predictions, we analyze how well the scalar field is reconstructed pixel-wise. We compute the mean square error (MSE) between the prediction $\hat{\mathbf{x}}_{t+\Delta t}$ and the ground truth $\mathbf{x}_{t+\Delta t}$ at timestep $t + \Delta t$. We also provide the peak signal-to-noise ratio (PSNR) between $\hat{\mathbf{x}}_{t+\Delta t}$ and $\mathbf{x}_{t+\Delta t}$. Tab. 3 shows the average MSE and PSNR evaluated on all test samples across all datasets. The high average PSNR and nearly zero MSE across all datasets indicates that our UNet-T model successfully captures the majority of features in future timesteps.

To evaluate the quality of the structural uncertainty estimation, we use Expected Calibration Error (ECE) [44] and reliability diagrams [11]. The ECE metric quantifies how well a model's predicted confidence aligns with actual accuracy (where confidence is defined as $1-$ uncertainty). It bins predictions by confidence and computes the weighted average of the absolute difference between accuracy and confidence in each bin. The lower the ECE the better, with the ideal ECE value being 0. A reliability diagram visually assesses calibration by plotting predicted confidence against estimations across confidence bins.

As shown in Tab. 4, the ECE of MC-U outperforms the

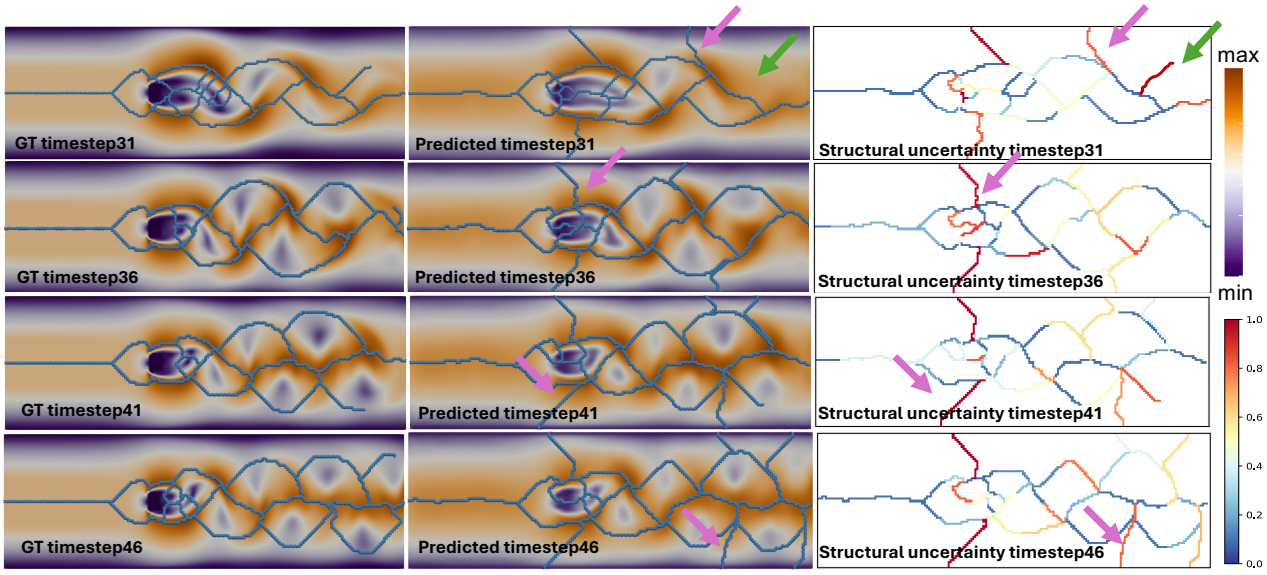

Figure 9: Structural uncertainty visualization of the Morse skeleton for multi-step data predictions for Navier Stokes dataset. The input scalar field at timestep 26 is used to recursively predict future timesteps, each prediction serving as input for the next. The Morse skeleton is captured in blue. Pink and green arrows indicate false positives and false negatives, respectively.

Table 4: Quantitative evaluation of MC-U. We report the average expected calibration error (ECE) of the baseline Prob.UNet and MC-U over all test samples for predicted future timesteps across all datasets. MC-U achieves lower ECE error is highlighted in bold.

Datasets	Heated Cylinder	Red Sea	Navier Stokes	Ionization Front
ECE of Prob.UNet	0.6571	0.8166	0.6061	0.7642
ECE of MC-U	**0.0929**	**0.1587**	**0.1106**	**0.1950**

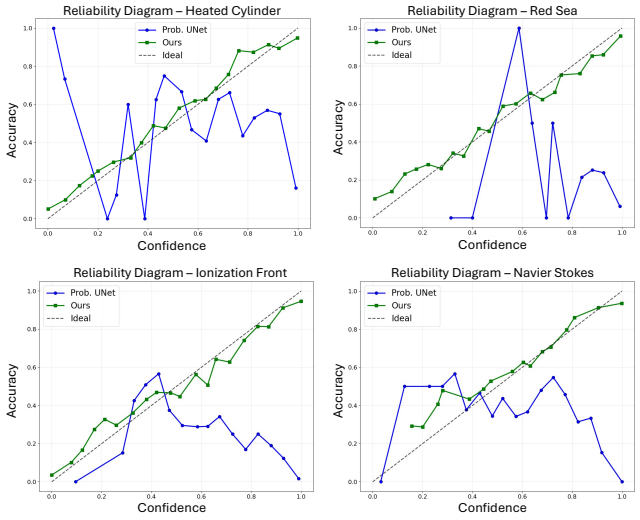

Figure 10: Reliability diagrams of test samples across various datasets, comparing MC-U with the baseline, Prob.UNet. The dotted line $x = y$ represents the ideal reliability curve.

Prob.UNet across all datasets that are highlighted in bold. The structural uncertainty estimated by MC-U accurately reflects the actual correspondence between the Morse skeleton of the predictions and the ground truth structures, demonstrating the effectiveness of our approach in capturing well-calibrated structural uncertainty relative to the true topological features. In Fig. 10, the structural uncertainty estimations of MC-U highly align with the ideal reliability curve. In contrast, the reliability curves of the baseline Prob.UNet show deviations from the diagonal, indicating that its uncertainty estimates are less consistent with the true structure. The curves of Prob.UNet mainly stay in the bottom-right quadrant, implying over-confident predictions (high confidence, low accuracy). This is in line with Prob.UNet's qualitative results of near-zero uncertainty in most locations.

6 Conclusion and Discussion

Our framework introduces a novel approach to reasoning about uncertainty in scientific data predictions. It quantifies and visualizes structural uncertainty in predicted Morse complexes. It integrates the UNet-T model, which forecasts future states from the current state, with the MC-U model, which incorporates Morse complex structures derived from predicted data. We extend probabilistic discrete Morse theory by performing guided walks on the scalar field and jointly estimating structural uncertainty using a GNN. By operating on Morse complexes extracted from predicted scalar fields, the MC-U model effectively captures and interprets structural variations introduced by data predictions.

While our current work primarily focuses on Morse complexes, the MC-U model demonstrates the potential to estimate uncertainty in other topological descriptors, such as merge trees and Reeb graphs. These examples highlight the flexibility of using GNNs to model structural variability across different topological representations. Furthermore, although we adopt UNet-T for single- and multi-step predictions to study structural uncertainty, our framework is adaptable and can be integrated with a broader class of deep learning-based surrogate models, including those used for reconstruction, super-resolution, or generative tasks.

Another important direction is to investigate how structural uncertainty affects downstream scientific tasks, such as feature tracking and event detection, where reliable topological interpretations are essential. Additionally, enhancing the MC-U framework to incorporate temporal consistency—ensuring smooth and coherent uncertainty estimates across timesteps—may further improve the robustness and interpretability of structural uncertainty analysis. We leave these extensions for future work.

Acknowledgement

This research was partially supported by grants from the National Science Foundation (NSF) CCF-2144901, DMS-2301361, IIS-2205418, and DMS-2134223; the Department of Energy (DOE) DE-SC0021015 and DE-SC0023157; and the National Institute of Health (NIH) R01NS143143 and R01CA297843.

References

[1] M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya, et al. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information fusion*, 76:243–297, 2021. doi: 10.1016/j.inffus.2021.05.008 2

[2] A. N. Angelopoulos, A. P. Kohli, S. Bates, M. Jordan, J. Malik, T. Alshaabi, S. Upadhyayula, and Y. Romano. Image-to-image regression with distribution-free uncertainty quantification and applications in imaging. In *International Conference on Machine Learning*, pp. 717–730. PMLR, 2022. doi: 10.48550/arXiv.2202.05265 2

[3] T. M. Athawale, D. Maljovec, L. Yan, C. R. Johnson, V. Pascucci, and B. Wang. Uncertainty visualization of 2D Morse complex ensembles using statistical summary maps. *IEEE Transactions on Visualization and Computer Graphics*, 28(4):1955–1966, 2022. doi: 10.1109/TVCG.2020.3022359 1, 2

[4] C. F. Baumgartner, K. C. Tezcan, K. Chaitanya, A. M. Hötker, U. J. Muehlematter, K. Schawkat, A. S. Becker, O. Donati, and E. Konukoglu. Phiseg: Capturing uncertainty in medical image segmentation. In *Medical Image Computing and Computer Assisted Intervention*, pp. 119–127, 2019. doi: 10.1007/978-3-030-32245-8_14 2

[5] K. Beketayev, D. Yeliussizov, D. Morozov, G. Weber, and B. Hamann. Measuring the distance between merge trees. In *Topological Methods in Data Analysis and Visualization III: Theory, Algorithms, and Applications*, Mathematics and Visualization, pp. 151–165. 2014. 1, 2

[6] S. Camarri, M. Buffoni, A. Iollo, and M. V. Salvetti. Simulation of the three-dimensional flow around a square cylinder between parallel walls at moderate reynolds numbers. In *XVII Congresso di Meccanica Teorica ed Applicata*, 2005. 6, 7

[7] G. Carlsson, A. J. Zomorodian, A. Collins, and L. J. Guibas. Persistence barcodes for shapes. In *Proceedings of the Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, pp. 124–135, 2004. doi: 10.1145/1057432.1057449 1

[8] H. Carr, J. Snoeyink, and U. Axen. Computing contour trees in all dimensions. *Computational Geometry*, 24(2):75–94, 2003. doi: 10.1016/S0925-7721(02)00093-7 1, 2

[9] Y.-T. Chen, H. Li, N. Shi, X. Luo, W. Xu, and H.-W. Shen. Explorable INR: An Implicit Neural Representation for Ensemble Simulation Enabling Efficient Spatial and Parameter Exploration. *IEEE Transactions on Visualization and Computer Graphics*, 2025. doi: 10.1109/TVCG.2025.3567052 1

[10] Computer Graphics Laboratory. Visualization data repository. https://cgl.ethz.ch/research/visualization/data.php, 2024. 6

[11] M. H. DeGroot and S. E. Fienberg. The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 1983. 8

[12] H. Edelsbrunner and J. Harer. *Computational Topology: An Introduction*. American Mathematical Society, 2010. 2, 3

[13] H. Edelsbrunner, J. Harer, V. Natarajan, and V. Pascucci. Morse-Smale complexes for piece-wise linear 3-manifolds. In *Annual Symposium on Computational Geometry*, pp. 361–370, 2003. doi: 10.1145/777792.777846 1

[14] H. Edelsbrunner, J. Harer, and A. J. Zomorodian. Hierarchical Morse-Smale complexes for piecewise linear 2-manifolds. *Discrete & Computational Geometry*, pp. 87–107, 2003. doi: 10.1145/378583.378626 1, 2

[15] H. Edelsbrunner, D. Letscher, and A. J. Zomorodian. Topological persistence and simplification. *Proceedings 41st Annual Symposium on Foundations of Computer Science*, 28:454–463, 2000. doi: 10.1109/SFCS.2000.892133 1, 6

[16] G. Favelier, N. Faraj, B. Summa, and J. Tierny. Persistence atlas for critical point variability in ensembles. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):1152–1162, 2019. 2

[17] R. Forman. Morse theory for cell complexes. *Advances in mathematics*, 134(1):90–145, 1998. 3

[18] R. Forman. A user's guide to discrete Morse theory. *Séminaire Lotharingien de Combinatoire*, 48:B48c, 2002. 35 pages. 3

[19] Y. Gal. Uncertainty in deep learning. 2016. 2

[20] Y. Gal and Z. Ghahramani. Bayesian convolutional neural networks with bernoulli approximate variational inference. *arXiv preprint*, 2015. doi: 10.48550/arXiv.1506.02158 2

[21] J. Gawlikowski, C. R. N. Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruspe, R. Triebel, P. Jung, R. Roscher, et al. A survey of uncertainty in deep neural networks. *Artificial Intelligence Review*, 56(Suppl 1):1513–1589, 2023. doi: 10.1007/s10462-023-10562-9 2

[22] S. Gerber and K. Potter. Data analysis with the Morse-Smale complex: The `msr` package for R. *Journal of Statistical Software*, 50(2):1–25, 2012. doi: 10.18637/jss.v050.i02 1

[23] R. Ghrist. Barcodes: The persistent topology of data. *Bulletin of the American Mathematical Society*, 45:61–75, 2008. doi: 10.1090/S0273-0979-07-01191-3 1

[24] N. Gorski, X. Liang, H. Guo, L. Yan, and B. Wang. A general framework for augmenting lossy compressors with topological guarantees. *IEEE Transactions on Visualization and Computer Graphics*, 2025. doi: 10.1109/TVCG.2025.3567054 1

[25] D. Günther, J. Salmon, and J. Tierny. Mandatory critical points of 2d uncertain scalar fields. *Computer Graphics Forum*, 33(3):31–40, 2014. doi: 10.1111/cgf.12359 2

[26] T. Günther, M. Gross, and H. Theisel. Generic objective vortices for flow visualization. *ACM Transactions on Graphics*, 36(4):141:1–141:11, 2017. doi: 10.1145/3072959.307368 6

[27] Z. Guo, Z. Wan, Q. Zhang, X. Zhao, Q. Zhang, L. M. Kaplan, A. Jøsang, D. H. Jeong, F. Chen, and J.-H. Cho. A survey on uncertainty reasoning and quantification in belief theory and its application to deep learning. *Information Fusion*, 101:101–987, 2024. doi: 10.1016/j.inffus.2023.101987 2

[28] S. Gupta, Y. Zhang, X. Hu, P. Prasanna, and C. Chen. Topology-aware uncertainty for image segmentation. *Advances in Neural Information Processing Systems*, 36, 2024. doi: 10.48550/arXiv.2306.05671 1, 3, 4

[29] T. Hazan and T. Jaakkola. On the partition function and random maximum a-posteriori perturbations. *International Conference on Machine Learning*, pp. 1667 – 1674, 2012. doi: 10.48550/arXiv.1206.6410 3

[30] I. Hoteit. The ieee scivis contest. https://kaust-vislab.github.io/SciVis2020/, 2020. 6

[31] X. Hu, Y. Wang, F. Li, D. Samaras, and C. Chen. Topology-aware segmentation using discrete morse theory. In *International Conference on Learning Representations*, 2021. doi: 10.48550/arXiv.2103.09992 3

[32] A. Kendall and Y. Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *International Conference on Neural Information Processing Systems*, pp. 5580 – 5590, 2017. doi: 10.48550/arXiv.1703.04977 2

[33] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. doi: 10.48550/arXiv.1412.6980 6

[34] S. Kohl, B. Romera-Paredes, C. Meyer, J. De Fauw, J. R. Ledsam, K. Maier-Hein, S. M. A. Eslami, D. Jimenez Rezende, and O. Ronneberger. A probabilistic U-Net for segmentation of ambiguous images. In *Proceedings of International Conference on Neural Information Processing Systems*, pp. 6965–6975. Curran Associates Inc., 2018. doi: 10.48550/arXiv.1806.05034 2, 6

[35] M. Kraus. Visualization of uncertain contour trees. In *Proceedings of the International Conference on Information Visualization Theory and Applications*, pp. 132–139, 2010. doi: 10.5220/0002817201320139 1, 2

[36] B. Lakshminarayanan, A. Pritzel, and C. Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *International Conference on Neural Information Processing Systems*, pp. 6405 – 6416, 2017. doi: 10.48550/arXiv.1612.01474 2

[37] F. Lan, B. Gamelin, L. Yan, J. Wang, B. Wang, and H. Guo. Topological characterization and uncertainty visualization of atmospheric rivers. *Computer Graphics Forum*, 2024. doi: 10.1111/cgf.15084 1, 2

[38] M. Lazaro-Gredilla, A. Dedieu, and D. George. Perturb-and-max-product: Sampling and learning in discrete energy-based models. *Advances in Neural Information Processing Systems*, pp. 928–940, 2021. 3

[39] M. Li, X. Yan, L. Yan, T. Needham, and B. Wang. Flexible and probabilistic topology tracking with partial optimal transport. *IEEE Transactions on Visualization and Computer Graphics*, 2025. doi: 10.1109/TVCG.2025.3561300 1

[40] W. Lyu, R. Sridharamurthy, J. M. Phillips, and B. Wang. Fast comparative analysis of merge trees using locality-sensitive hashing. *IEEE Transactions on Visualization and Computer Graphics*, 31(1):141–151, 2025. doi: 10.1109/TVCG.2024.3456383 1

[41] M. Mihai and R. Westermann. Visualizing the stability of critical points in uncertain scalar fields. *Computers & Graphics*, 41:13–25, 2014. 2

[42] J. Milnor. *Morse Theory*. Princeton University Press, 1963. 2, 3

[43] K. P. Murphy. Conjugate bayesian analysis of the gaussian distribution. 2007. 3

[44] M. P. Naeini, G. Cooper, and M. Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2015. doi: 10.1609/aaai.v29i1.9602 8

[45] G. Papandreou and A. L. Yuille. Perturb-and-map random fields: Using discrete optimization to learn and sample from energy models. In *International Conference on Computer Vision*, pp. 193–200, 2011. doi: 10.1109/ICCV.2011.6126242 3

[46] S. Popinet. Free computational fluid dynamics. *ClusterWorld*, 2(6), 2004. 6, 7

[47] A. Rahman, J. M. J. Valanarasu, I. Hacihaliloglu, and V. M. Patel. Ambiguous medical image segmentation using diffusion models. In *IEEE/CVF conference on computer vision and pattern recognition*, pp. 11536–11546, 2023. doi: 10.1109/CVPR52729.2023.01110 2

[48] G. Reeb. Sur les points singuliers d'une forme de pfaff complètement intégrable ou d'une fonction numérique (on the singular points of a complete integral pfaff form or of a numerical function). *Comptes Rendus Acad. Science Paris*, 222:847–849, 1946. 1

[49] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer Assisted Intervention*, 2015. doi: 10.1007/978-3-319-24574-4_28 1, 2, 4

[50] C. Rupprecht, I. Laina, R. DiPietro, M. Baust, F. Tombari, N. Navab, and G. D. Hager. Learning in an uncertain world: Representing ambiguity through multiple hypotheses. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3611–3620. Venice, Italy, 2017. doi: 10.1109/ICCV.2017.388 2

[51] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008. doi: 10.1109/TNN.2008.2005605 1, 4

[52] J. Shen, Y. Duan, and H.-W. Shen. Surroflow: A flow-based surrogate model for parameter space exploration and uncertainty quantification. *IEEE Transactions on Visualization and Computer Graphics*, 2024. doi: 10.1109/TVCG.2024.3456372 2

[53] J. Shen and H.-W. Shen. Psrflow: Probabilistic super resolution with flow-based models for scientific data. *IEEE Transactions on Visualization and Computer Graphics*, pp. 986–996, 2023. doi: 10.1109/TVCG.2023.3327171 1, 2

[54] N. Shi, J. Xu, H. Li, H. Guo, J. Woodring, and H.-W. Shen. VDL-Surrogate: A View-Dependent Latent-based Model for Parameter Space Exploration of Ensemble Simulations. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):820–830, 2023. Proc. IEEE VIS 2022. doi: 10.1109/TVCG.2022.3209384 1

[55] N. Shi, J. Xu, S. W. Wurster, H. Guo, J. Woodring, L. V. Roekel, and H.-W. Shen. GNN-Surrogate: A Hierarchical and Adaptive Graph Neural Network for Parameter Space Exploration of Unstructured-Mesh Ocean Simulations. *IEEE Transactions on Visualization and Computer Graphics*, 28(6):2301–2313, 2022. Proc. IEEE PacificVis 2022. doi: 10.1109/TVCG.2022.3150932 1

[56] K. Sohn, X. Yan, and H. Lee. Learning structured output representation using deep conditional generative models. pp. 3483–3491, 2015. 2

[57] R. Sridharamurthy, T. B. Masood, A. Kamakshidasan, and V. Natarajan. Edit distance between merge trees. *IEEE Transactions on Visualization and Computer Graphics*, 26(3):1518–1531, 2020. doi: 10.1109/TVCG.2018.2873612 1

[58] J. Tierny, G. Favelier, J. A. Levine, C. Gueunet, and M. Michaux. The topology toolkit. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):832–842, 2018. doi: 10.1109/TVCG.2017.2743938 6

[59] D. Whalen and M. L. Norman. The ieee scivis contest, 2008. 6

[60] D. J. Whalen and M. L. Norman. Ionization front instabilities in primordial h ii regions. *The Astrophysical Journal*, 673:664 – 675, 2007. doi: 10.1086/524400 7

[61] R. Whitaker, M. Mirzargar, and R. Kirby. Contour boxplots: A method for characterizing uncertainty in feature sets from simulation ensembles. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2713–2722, 2013. doi: 10.1109/TVCG.2013.143 2

[62] K. Wu and S. Zhang. A contour tree based visualization for exploring data with uncertainty. *International Journal for Uncertainty Quantification*, 3(3):203–223, 2012. doi: 10.1615/Int.J.UncertaintyQuantification.2012003956 2

[63] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020. doi: 10.1109/TNNLS.2020.2978386 1, 4

[64] L. Yan, H. Guo, T. Peterka, B. Wang, and J. Wang. TROPHY: A topologically robust physics-informed tracking framework for tropical cyclones. *IEEE Transactions on Visualization and Computer Graphics*, 30:1302–1312, 2024. doi: 10.1109/TVCG.2023.3326905 1

[65] L. Yan, X. Liang, H. Guo, and B. Wang. Toposz: Preserving topology in error-bounded lossy compression. *IEEE Transactions on Visualization and Computer Graphics*, 30:1302–1312, 2024. Presented at IEEE VIS 2023. doi: 10.1109/TVCG.2023.3326920 1

[66] L. Yan, T. B. Masood, F. Rasheed, I. Hotz, and B. Wang. Geometry-aware merge tree comparisons for time-varying data with interleaving distances. *IEEE Transactions on Visualization and Computer Graphics*, 29(8):3489–3506, 2023. doi: 10.1109/TVCG.2022.3163349 1

[67] L. Yan, T. B. Masood, R. Sridharamurthy, F. Rasheed, V. Natarajan, I. Hotz, and B. Wang. Scalar field comparison with topological descriptors: Properties and applications for scientific visualization. *Computer Graphics Forum*, 40(3):599–633, 2021. doi: 10.1111/cgf.14331 1

[68] L. Yan, Y. Wang, E. Munch, E. Gasparovic, and B. Wang. A structural average of labeled merge trees for uncertainty visualization. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):832–842, 2020. doi: 10.1109/TVCG.2019.2934242 1, 2

[69] S. Zhang, H. Tong, J. Xu, and R. Maciejewski. Graph convolutional networks: a comprehensive review. *Computational Social Networks*, 6(1):1–23, 2019. doi: 10.1186/s40649-019-0069-y 1, 4

[70] W. Zhang, P. K. Agarwal, and S. Mukherjee. Contour trees of uncertain terrains. In *Proceedings of the 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems*, vol. 43, 2015. doi: 10.1145/2820783.2820823 1, 2