TU Delft

Delft University of Technology

Document Version
Final published version

Licence
Dutch Copyright Act (Article 25fa)

Citation (APA)
Money, R., Beferull-Lozano, B., & Isufi, E. (2025). Change Point Detection, Localization, and Imputation of Topological Signals. In M. B. Matthews (Ed.), *Proceedings of the 2025 59th Asilomar Conference on Signals, Systems, and Computers* (pp. 171-175). IEEE. https://doi.org/10.1109/IEEECONF67917.2025.11443741

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Change Point Detection, Localization, and Imputation of Topological Signals

Rohan Money, *Baltasar Beferull-Lozano*, *Elvin Isufi*

2025 59th Asilomar Conference on Signals, Systems, and Computers | 979-8-3315-8745-1/25/$31.00 ©2025 IEEE | DOI: 10.1109/IEEECONF67917.2025.11443741

***Abstract*—We propose a topology-aware online framework for detecting and localizing change points in partially observed temporal signals defined over cellular complexes. We model the process as a linear state space model. The latent dynamics follow a topology-consistent stochastic partial differential equation (SPDE), and the observations a topologically-filtered version of the state. Hidden states are inferred via a Kalman filter, while model parameters are adapted online through likelihood-based updates. Change points are detected by tracking deviations in the uncertainty parameters using a score-based criterion, and the responsible components are localized by mapping these deviations to specific cells in the complex. Under missing or noisy observations, backward Kalman smoothing—reinitialized at each detected change—provides consistent state reconstruction and reliable imputation. Experiments on synthetic and EPANET-simulated water networks demonstrate accurate detection, precise localization, and robust reconstruction under partial observability.**

***Index Terms*—Cellular Complex, Time-series, Change-point**

I. Introduction

Many real-world systems generate time-series data whose evolution is governed by structural dependencies among interconnected entities [1]–[3]. These time-series are often incomplete or comprise noisy measurements of latent processes evolving over an underlying topology [4], [5]. For example, in water distribution networks, only a limited set of flow and pressure measurements is available, while the full hydraulic state—including nodal heads, pipe flows, and demand patterns—remains latent [6]. Structural events such as leaks or valve operations modify the network connectivity and induce localized distributional shifts in these hidden dynamics. This work develops an efficient framework to model such latent processes from incomplete data, with emphasis on detecting and localizing topological regime shifts.

We model the observed time-series as structured signals supported on a cellular complex—a higher-order generalization of a graph with nodes, edges, and faces [1], [3]. In this setting, physical quantities such as nodal pressures, edge flows, and face circulations naturally arise as signals supported on different cell dimensions, whose interactions are compactly represented through Hodge Laplacians constructed from boundary maps that capture cell incidences [1]. To capture the temporal evolution of these topological signals, we model the latent dynamics using a stochastic partial differential equation (SPDE) defined on the cellular complex, where the Hodge Laplacian governs how information propagates consistently with the underlying topology [4]–[6]. The observed time-series are then obtained by passing this latent state through cellular convolution filters [7]. Although this formulation captures stationary topological dynamics, it does not accommodate localized distributional shifts that commonly arise in real-world systems. Events such as a blocked road or a leaking pipe can induce abrupt, localized changes in flow behavior, even when the underlying cellular complex remains fixed. Existing online methods for topological time-series can handle distributional drift [8]–[10], but, to the best of our knowledge, none can jointly detect structural and temporal change points in the latent dynamics.

We address this gap by developing an efficient framework that jointly estimates the latent process from incomplete observations and detects, as well as localizes, topological regime shifts that alter the underlying dynamics. We formulate the SPDE so that each state is paired with an uncertainty parameter capturing the variability of its local dynamics. By monitoring fluctuations in these parameters through a score-based criterion, the method detects and localizes changes as they occur. Detecting a structural change point enables localization of the root cause of anomalous behavior, whereas detecting a temporal change point allows the model to adapt promptly to the new dynamics—particularly important under partial observability. To achieve this, the latent states and model parameters are estimated alternatively using a Kalman filter [11] and maximum likelihood estimation (MLE), respectively. The main contributions of the paper are listed below:

- **Topology-aware online detection:** We introduce the first framework for joint change-point detection and localization in time-series defined over cellular complexes.
- **Change-point–aware time-series modeling:** We develop a principled framework for cellular time-series modeling that explicitly accounts for localized distributional shifts. By adapting online to evolving dynamics, it enables reliable inference and reconstruction even under noise, missing data, and abrupt structural changes.

II. Problem Setup and Formulation

A. Signals Defined Over a Cell Complex

A cell complex $\mathcal{K}$ consists of k-dimensional cells— 0-cells (nodes), 1-cells (edges), 2-cells (faces), and so on. Let N_k

R. Money and B. Beferull-Lozano are with the SIGIPRO Dept., Simula Metropolitan Center for Digital Engineering, 0167 Oslo, Norway (emails: rohan@simula.no, baltasar@simula.no). E. Isufi is with Intelligent Systems Department, Delft University of Technology, Delft, Netherlands (e-mail: e.isufi-1@tudelft.nl).

The study was supported by the SURE-AI Center project number 357482 RCN, IKTPLUSS DISCO grant 338740, the TU Delft AI Labs programme, the NWO OTP GraSPA proposal #19497, and NWO project VENI 19052.

Authorized licensed use limited to: TU Delft Library. Downloaded on March 30,2026 at 07:22:16 UTC from IEEE Xplore. Restrictions apply.

denote the number of k-cells, and let $\mathbf{B}_k \in \mathbb{R}^{N_{k-1} \times N_k}$ be the boundary incidence matrix encoding how k-cells attach to $(k-1)$-cells [1]. From these boundary maps, and since $\mathbf{B}_k \mathbf{B}_{k+1} = \mathbf{0}$, the k-th Hodge Laplacian has the form

$$\mathbf{L}_k = \underbrace{\mathbf{B}_k^\top \mathbf{B}_k}_{\mathbf{L}_{k,d}} + \underbrace{\mathbf{B}_{k+1} \mathbf{B}_{k+1}^\top}_{\mathbf{L}_{k,u}} \in \mathbb{R}^{N_k \times N_k}.$$

which captures interactions among k-cells mediated through shared $(k-1)$- and $(k+1)$-cells [1], [2].

A *cell signal* assigns a scalar value to each cell, e.g., $\mathbf{x}_t^0 \in \mathbb{R}^{N_0}$ on nodes, $\mathbf{x}_t^1 \in \mathbb{R}^{N_1}$ on edges, and $\mathbf{x}_t^2 \in \mathbb{R}^{N_2}$ on faces. Throughout the paper we will consider 2-dimensional cell complexes to ease exposition and due to their practical benefits. We stack cell signals into the combined signal $\mathbf{x}_t = [\mathbf{x}_t^{0\top}, \mathbf{x}_t^{1\top}, \mathbf{x}_t^{2\top}]^\top \in \mathbb{R}^N$, and define the combined block-diagonal Laplacian $\mathbf{L} = \mathrm{blkdiag}(\mathbf{L}_0, \mathbf{L}_1, \mathbf{L}_2) \in \mathbb{R}^{N \times N}$, where $N = N_0 + N_1 + N_2$. To model cross-dimensional interactions, we also use the Dirac operator

$$\mathbf{D} = \begin{bmatrix} 0 & \mathbf{B}_1 & 0 \\ \mathbf{B}_1^\top & 0 & \mathbf{B}_2 \\ 0 & \mathbf{B}_2^\top & 0 \end{bmatrix} \in \mathbb{R}^{N \times N},$$

which couples signals between adjacent orders.

B. Modeling Objective

Let $\mathbf{y}_i = [\mathbf{y}_i^{0\top}, \mathbf{y}_i^{1\top}, \mathbf{y}_i^{2\top}]^\top \in \mathbb{R}^N$ denote the multivariate signal defined over the cellular complex $\mathcal{K}$ at discrete time index i. At each time step, we observe a noisy or incomplete version of this underlying signal, given by

$$\tilde{\mathbf{y}}_i = \mathbf{M}\mathbf{y}_i + \mathbf{v}_i \in \mathbb{R}^{N_o}, \tag{1}$$

where $\mathbf{M} \in \mathbb{R}^{N_o \times N}$ is a binary masking matrix that selects the observed components of $\mathbf{y}_i$, and $\mathbf{v}_i \in \mathbb{R}^{N_o}$ represents measurement noise. The statistical properties of the time series may experience distributional shifts resulting from unobserved changes in the underlying topological structure of $\mathcal{K}$.

We aim to develop an efficient modeling framework for this setting, assuming that the observed signal $\mathbf{y}_i$ arises from an unobserved latent process $\mathbf{x}_i$ that evolves according to the cellular topology. The proposed framework leverages this latent representation to achieve the following objectives:

1) **Detect**: Identify temporal change points, i.e., time instants at which the latent structure or its governing parameters change abruptly.
2) **Localize**: Determine the spatial or topological regions within $\mathcal{K}$ where these structural changes originate.
3) **Impute**: Reconstruct missing or corrupted observations by exploiting information from the detected change points and the learned latent dynamics.

We assume the latent process remains locally stationary within segments but may undergo abrupt changes due to structural perturbations. These transitions correspond to the change points that our framework seeks to detect and localize.

III. Proposed Framework

This section introduces the proposed topology-aware framework for joint state estimation, change-point detection, and localization. We begin by modeling the latent dynamics with an SPDE whose diffusion operator is defined on the cellular complex, ensuring topology-consistent evolution. We then obtain its discrete-time state-space form and detail the resulting estimation and detection steps.

A. System-Model Respecting the Topological Structure

We model the latent state $\mathbf{x}_t$ using a stochastic heat equation defined on the cellular complex $\mathcal{K}$ [4], [6]. To incorporate the topology of $\mathcal{K}$, we replace the continuous Laplacian in the classical heat equation with its discrete counterpart—the Hodge Laplacian $\mathbf{L}$. This yields the topology-aware SPDE

$$d\mathbf{x}_t = -c\,\mathbf{L}\,\mathbf{x}_t\,dt + \mathbf{D}\,\mathrm{diag}(\boldsymbol{\alpha})\,d\boldsymbol{\beta}_t, \tag{2}$$

where $c > 0$ is the diffusion coefficient, $\boldsymbol{\beta}_t$ denotes the Brownian motion, and $\mathbf{D}$ is the Dirac operator that couples adjacent topological orders. The drift term induces heat-like smoothing consistent with the cellular topology through the Hodge Laplacian $\mathbf{L}$, while the Dirac-driven noise term injects perturbations that couple adjacent topological dimensions, allowing fluctuations to propagate across nodes, edges, and faces via the boundary and coboundary relations. A key component of this model is the vector $\boldsymbol{\alpha} \in \mathbb{R}^N$ of topology-aware uncertainty intensities. Each entry $\boldsymbol{\alpha}[\ell]$ controls the local strength of the stochastic excitation injected at cell ℓ. In effect, $\boldsymbol{\alpha}$ governs the local variability of the latent diffusion dynamics. Any localized structural or dynamical modification produces a corresponding, topology-consistent change in the affected components of $\boldsymbol{\alpha}$. This makes $\boldsymbol{\alpha}$ the key quantity for detecting and localizing regime shifts in our framework.

Next, to relate the latent states to the measured signals, we employ a cellular convolution filter $\mathbf{H}(\mathbf{L}) \in \mathbb{R}^{N \times N}$ that jointly processes node-, edge-, and face-level components [7], [9]. This operator captures inter-order coupling through the boundary maps $\mathbf{B}_1$ and $\mathbf{B}_2$, yielding the block structure

$$\mathbf{H}(\mathbf{L}) = \begin{pmatrix} \mathbf{H}(\mathbf{L}_0) & \mathbf{B}_1\,\mathbf{H}(\mathbf{L}_1) & \mathbf{0} \\ \mathbf{B}_1^\top\,\mathbf{H}(\mathbf{L}_0) & \mathbf{H}(\mathbf{L}_1) & \mathbf{B}_2\,\mathbf{H}(\mathbf{L}_2) \\ \mathbf{0} & \mathbf{B}_2^\top\,\mathbf{H}(\mathbf{L}_1) & \mathbf{H}(\mathbf{L}_2) \end{pmatrix}. \tag{3}$$

Each diagonal block $\mathbf{H}(\mathbf{L}_k)$ is implemented using a polynomial graph filter [7]:

$$\mathbf{H}(\mathbf{L}_k) = \sum_{n=0}^{L_d-1} h_{n,d}\,\mathbf{L}_{k,d}^n + \sum_{n=0}^{L_u-1} h_{n,u}\,\mathbf{L}_{k,u}^n, \tag{4}$$

where $\{h_{n,d}\}$ and $\{h_{n,u}\}$ are learnable filter coefficients. This polynomial structure enables localized, multi-hop interactions among k-cells, governed by the topology encoded in $\mathbf{L}_k$. The observation model is obtained by linearly mapping the filtered state into the measurement space:

$$\tilde{\mathbf{y}}_i = \mathbf{M}\,\mathbf{H}(\mathbf{L})\,\mathbf{x}_i + \mathbf{v}_i, \tag{5}$$

where $\mathbf{M}$ selects the observed cells and $\mathbf{v}_i$ denotes measurement noise.

Authorized licensed use limited to: TU Delft Library. Downloaded on March 30,2026 at 07:22:16 UTC from IEEE Xplore. Restrictions apply.

Assuming uniform sampling with step δt, the SPDE in (2) is discretized [5] and a discrete state-space model is formed as

$$\mathbf{x}_i = \tilde{\mathbf{L}}\,\mathbf{x}_{i-1} + \mathbf{q}_i, \qquad \tilde{\mathbf{y}}_i = \mathbf{M}\,\mathbf{H}(\mathbf{L})\,\mathbf{x}_i + \mathbf{v}_i, \tag{6}$$

where $\tilde{\mathbf{L}} \approx \mathbf{I} - c\,\delta t\,\mathbf{L}$ and $\mathbf{q}_i \sim \mathcal{N}(0, \mathbf{Q})$ with $\mathbf{Q} = \delta t\,\mathbf{D}\,\mathrm{diag}^2(\boldsymbol{\alpha})\,\mathbf{D}^\top$. This structure-aware model forms the foundation for the filtering and inference procedures developed next. For brevity, we denote $\mathbf{H}(\mathbf{L})$ by $\mathbf{H}$ and its parameter vector by $\mathbf{h}$. At each discrete time index i, the goal is to jointly estimate the state $\mathbf{x}_i$, the observation filter $\mathbf{H}_i$, and the uncertainty vector $\boldsymbol{\alpha}_i$ from the available measurement $\tilde{\mathbf{y}}_i$.

B. Joint State and Parameter Estimation

The proposed inference scheme jointly updates the latent state $\mathbf{x}_i$, the uncertainty parameters $\boldsymbol{\alpha}_i$, and the observation operator $\mathbf{H}_i$ in an online fashion by alternating between state estimation and parameter adaptation. At each time step, the Kalman filter produces the state estimate under the current model, while $\boldsymbol{\alpha}_i$ and $\mathbf{H}_i$ (through $\mathbf{h}_i$) are updated using the predictive likelihood.

(a) State Estimation.: Given the parameters $\boldsymbol{\alpha}_{i-1}$ and $\mathbf{H}_{i-1}$, the latent state $\mathbf{x}_i$ and its covariance $\mathbf{P}_i$ are estimated using the standard Kalman filter [11]. The *prediction step* propagates the previous posterior mean $\mathbf{x}_{i-1|i-1}$ and covariance $\mathbf{P}_{i-1|i-1}$ through the topology-aware dynamics:

$$\mathbf{x}_{i|i-1} = \tilde{\mathbf{L}}\,\mathbf{x}_{i-1|i-1}, \tag{7}$$

$$\mathbf{P}_{i|i-1} = \tilde{\mathbf{L}}\,\mathbf{P}_{i-1|i-1}\,\tilde{\mathbf{L}}^\top + \mathbf{Q}_i, \quad \mathbf{Q}_i = \delta t\,\mathbf{D}\,\mathrm{diag}^2(\boldsymbol{\alpha}_{i-1})\,\mathbf{D}^\top. \tag{8}$$

Here, $\tilde{\mathbf{L}} = \mathbf{I} - c\,\delta t\,\mathbf{L}$ is the discrete diffusion operator, and $\mathbf{Q}_i$ captures process noise whose covariance depends on $\boldsymbol{\alpha}_{i-1}$ and the topological coupling through $\mathbf{D}$.

The *correction step* refines the predicted state using the new observation $\tilde{\mathbf{y}}_i$:

$$\mathbf{K}_i = \mathbf{P}_{i|i-1}(\mathbf{M}\mathbf{H}_{i-1})^\top(\mathbf{M}\mathbf{H}_{i-1}\mathbf{P}_{i|i-1}(\mathbf{M}\mathbf{H}_{i-1})^\top + \sigma^2\mathbf{I})^{-1} \tag{9}$$

$$\mathbf{x}_{i|i} = \mathbf{x}_{i|i-1} + \mathbf{K}_i(\tilde{\mathbf{y}}_i - \mathbf{M}\mathbf{H}_{i-1}\mathbf{x}_{i|i-1}) \tag{10}$$

$$\mathbf{P}_{i|i} = (\mathbf{I} - \mathbf{K}_i\mathbf{M}\mathbf{H}_{i-1})\mathbf{P}_{i|i-1}. \tag{11}$$

The Kalman gain $\mathbf{K}_i$ adaptively balances model uncertainty with observation reliability, yielding the posterior mean $\mathbf{x}_{i|i}$ and covariance $\mathbf{P}_{i|i}$ that summarize all available information up to time i.

(b) Parameter Adaptation.: Given the predictive distribution $p(\tilde{\mathbf{y}}_i \mid \tilde{\mathbf{y}}_{1:i-1}) \sim \mathcal{N}(\mathbf{M}\mathbf{H}\mathbf{x}_{i|i-1},\, \mathbf{S}_i)$, the instantaneous log-likelihood of the new observation is

$$\mathcal{L}_i(\boldsymbol{\alpha}, \mathbf{H}) = -\frac{1}{2}\log\det(\mathbf{S}_i) - \frac{1}{2}\big(\tilde{\mathbf{y}}_i - \mathbf{M}\mathbf{H}\mathbf{x}_{i|i-1}\big)^\top\mathbf{S}_i^{-1}\big(\tilde{\mathbf{y}}_i - \mathbf{M}\mathbf{H}\mathbf{x}_{i|i-1}\big), \tag{12}$$

where $\mathbf{S}_i = \mathbf{M}\mathbf{H}\mathbf{P}_{i|i-1}\mathbf{H}^\top\mathbf{M}^\top + \sigma^2\mathbf{I}$. The parameters $\mathbf{h}$ and $\boldsymbol{\alpha}$ are updated via stochastic gradient ascent applied to the instantaneous log-likelihood:

$$\mathbf{h}_i = \mathbf{h}_{i-1} + a_1\,\nabla_{\mathbf{h}}\,\mathcal{L}_i(\boldsymbol{\alpha}_{i-1}, \mathbf{H}_{i-1}), \tag{13}$$

$$\boldsymbol{\alpha}_i = \boldsymbol{\alpha}_{i-1} + a_2\,\nabla_{\boldsymbol{\alpha}}\big[\mathcal{L}_i(\boldsymbol{\alpha}_{i-1}, \mathbf{H}_{i-1}) + \log p(\boldsymbol{\alpha})\big], \tag{14}$$

Here, $\nabla_{\mathbf{h}}$ and $\nabla_{\boldsymbol{\alpha}}$ denote gradients with respect to $\mathbf{h}$ and $\boldsymbol{\alpha}$, respectively, and a_1, a_2 are the corresponding learning rates. A Gaussian prior $p(\boldsymbol{\alpha}) = \mathcal{N}(\boldsymbol{\alpha} \mid \boldsymbol{\alpha}_{i-1}, \boldsymbol{\Sigma}_\alpha)$ regularizes the uncertainty parameters, keeping $\boldsymbol{\alpha}_i$ close to its previous value. This prior is applied only to $\boldsymbol{\alpha}$, since enforcing smooth temporal evolution in these parameters is essential for distinguishing genuine structural changes from random fluctuations. In contrast, no prior is placed on $\mathbf{h}$ because $\mathbf{H}$ serves only as the observation map and does not carry structural information relevant for detecting regime shifts. Regularizing $\mathbf{h}$ would constrain this mapping without providing statistical benefit, and would unnecessarily limit the model's ability to adapt to measurement variations.

C. Change-Point Detection and Topological Localization

The adaptive updates of $\boldsymbol{\alpha}_i$ capture how uncertainty propagates through the cellular topology. When the system undergoes a structural or parametric transition—such as a local reconfiguration of the complex or an abrupt change in diffusion behavior—specific components of $\boldsymbol{\alpha}_i$ deviate sharply from their recent steady-state values. This property enables online detection and localization of topological change points.

For each component l (corresponding to a node, edge, or face), we compute a local moving-average baseline:

$$\bar{\boldsymbol{\alpha}}_i[l] = \frac{1}{\tau}\sum_{k=i-\tau+1}^{i}\boldsymbol{\alpha}_k[l], \tag{15}$$

where τ is the averaging window length representing typical stationary behavior. The instantaneous deviation at cell l is then measured by

$$\mathbf{score}_i[l] = \big|\bar{\boldsymbol{\alpha}}_i[l] - \boldsymbol{\alpha}_i[l]\big|, \tag{16}$$

quantifying how much the current uncertainty differs from its recent history.

A change point is declared whenever the deviation of any cell exceeds its threshold γ:

$$\mathrm{CP}_i = \bigvee_l \big\{\,\mathbf{score}_i[l] > \gamma\,\big\}, \tag{17}$$

where CP_i is a binary indicator that evaluates to 1 when a change point is detected at time i and 0 otherwise, and $\bigvee$ denotes a logical OR over cells. The active set

$$\mathcal{S}_i = \{\, l : \mathbf{score}_i[l] > \gamma \,\} \tag{18}$$

identifies the subset of simplices responsible for the detected change. This procedure not only detects the occurrence of a topological transition but also *localizes* the structural region where it originates—providing an interpretable mapping between statistical deviations and the physical topology.

D. Post-Detection Smoothing and Reconstruction

Once a change point is detected, the online filter enters a transient adaptation phase during which the parameters $(\mathbf{h}_i, \boldsymbol{\alpha}_i)$ gradually converge to their new steady-state values. After the deviation scores $\mathbf{score}_i[l]$ fall below their thresholds, marking the onset of a new stable regime, we apply a fixed-interval Rauch–Tung–Striebel (RTS) smoother retrospectively over the preceding transition window. The smoother reuses the newly stabilized parameters and incorporates information

Authorized licensed use limited to: TU Delft Library. Downloaded on March 30,2026 at 07:22:16 UTC from IEEE Xplore. Restrictions apply.

Algorithm 1: CCP-SPDE: Change-point–aware cellular SPDE.

Given : $\mathbf{L}, \mathbf{M}$
Initialize: $a_1, a_2, c, \delta t,\ \mathbf{x}_{0|0}, \mathbf{P}_{0|0},\ \mathbf{h}_0,\ \boldsymbol{\alpha}_0,\ \gamma$
for $i = 1, 2, \ldots$ **do**
  Obtain observation $\tilde{\mathbf{y}}_i$
  State Estimation:
    KF prediction (7)–(8),
    KF correction (9)–(11)
  Parameter Adaptation:
    Update $\mathbf{h}_i$ via (13),
    Update $\boldsymbol{\alpha}_i$ via (14)
  Change-Point Detection:
    Compute baseline $\tilde{\boldsymbol{\alpha}}_i$ and scores using (15)–(16)
    Detect CP_i via (17), update active set $\mathcal{S}_i$ via (18)
  Post-Stabilization Smoothing:
  if *parameters stabilized after a CP* **then**
    Run RTS smoother (19)–(21);
    Reconstruct $\hat{\mathbf{y}}_j = \mathbf{H}_T\, \mathbf{x}_{j|T}$
  end
end

from future observations to correct the state estimates obtained during adaptation. For each time index i in this transition window, processed backward from its last time step $T-1$ to its first, we apply a Rauch–Tung–Striebel (RTS) smoother. The corresponding smoother gain $\mathbf{J}_i$, the smoothed state $\mathbf{x}_{i|T}$ and covariance $\mathbf{P}_{i|T}$ updates are

$$\mathbf{J}_i = \mathbf{P}_{i|i}\,\tilde{\mathbf{L}}^\top \mathbf{P}_{i+1|i}^{-1}, \tag{19}$$

$$\mathbf{x}_{i|T} = \mathbf{x}_{i|i} + \mathbf{J}_i\big(\mathbf{x}_{i+1|T} - \mathbf{x}_{i+1|i}\big), \tag{20}$$

$$\mathbf{P}_{i|T} = \mathbf{P}_{i|i} + \mathbf{J}_i\big(\mathbf{P}_{i+1|T} - \mathbf{P}_{i+1|i}\big)\mathbf{J}_i^\top. \tag{21}$$

Then, the reconstructed observations follow as

$$\hat{\mathbf{y}}_i = \mathbf{H}_T\, \mathbf{x}_{i|T}. \tag{22}$$

where $\mathbf{H}_T$ is estimated cellular filter at time T.

This post-stabilization smoothing serves two purposes: (i) it improves the latent trajectory over the transition period by incorporating future observations; and (ii) it provides coherent estimates of missing or noisy observations using the stabilized model parameters together with the smoothed states. We refer to the entire procedure as Change-point-aware cellular SPDE (CCP-SPDE) and is described in **Algorithm** 1.

IV. Experiments

We evaluate the proposed framework on both synthetic data and EPANET-generated water–network data, assessing its ability to detect and localize change points as well as to reconstruct partially observed signals.

Performance Metrics: Reconstruction quality is evaluated using the normalized mean squared error $\mathrm{NMSE} = \frac{\sum_{t=1}^{T} \|\mathbf{y}_t - \hat{\mathbf{y}}_t\|_2^2}{\sum_{t=1}^{T} \|\mathbf{y}_t\|_2^2}$. Change-point detection performance is quantified using the area under the ROC curve (AUC), computed from the true-positive rate $\mathrm{TPR}(\gamma) = \frac{\mathrm{TP}}{\mathrm{TP+FN}}$ and the false-positive rate $\mathrm{FPR}(\gamma) = \frac{\mathrm{FP}}{\mathrm{FP+TN}}$ as the detection threshold varies. Here: (i) TP denotes correctly detected change points, (ii) FP denotes false alarms, (iii) FN denotes missed detections, and (iv) TN denotes correctly identified non-change points.

Dataset	p_m	CCP-SPDE	CC-SPDE	G-SPDE
Synthetic	0.05	0.02	0.09	0.52
	0.10	0.08	0.21	0.76
	0.15	0.18	0.37	1.30
	0.20	0.35	0.65	2.60
EPANET	0.05	0.12	0.30	2.41
	0.10	0.32	0.45	2.82
	0.15	0.41	0.52	2.95
	0.20	0.46	0.64	2.99

TABLE I: Reconstruction accuracy under missing data. CCP-SPDE is the propose approach; CC-SPDE is its ablated version without smoothing; G-SPDE is a graph-based baseline.

Baselines: We compare the proposed framework with methods targeting detection/localization and reconstruction.

(A) Detection and Localization. (i) *CP–Mean*, which flags a change when the running mean of the observed signal departs from its historical average; and (ii) *CP–TIRSO* [12], a node-separable graph-learning method that treats each cell independently and uses its learned parameters as proxies for $\boldsymbol{\alpha}$. Both baselines ignore higher-order interactions between nodes.

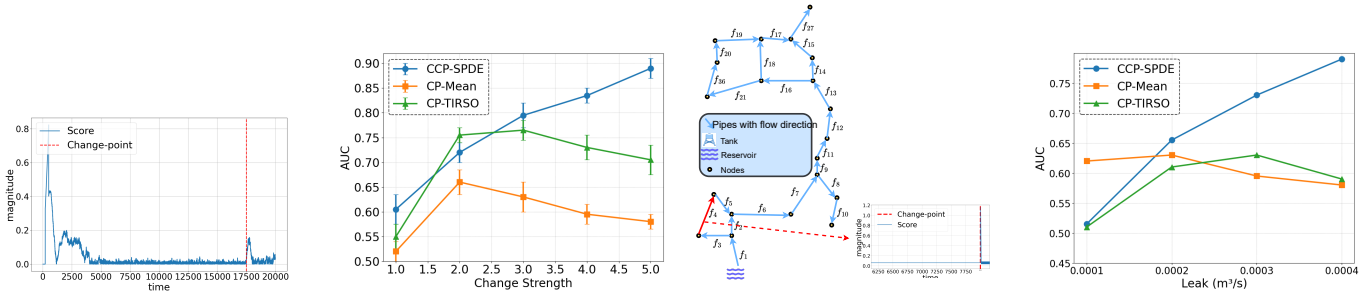
(B) Reconstruction. (i) *G-SPDE*, a graph-based SPDE that models only pairwise dependencies and omits multi-order couplings; and (ii) *CC-SPDE*, an ablated variant of our method that uses the same cell-complex dynamics but performs only forward Kalman filtering without RTS smoothing, making it vulnerable to parameter drift during regime transitions.

A. Synthetic Data Experiments

Synthetic data are generated using the topology-aware SPDE model in (6) on a cellular complex with 10 nodes, 17 edges, and 9 faces. The underlying graph is drawn from an Erdős–Rényi model, and all admissible 2-cells are subsequently filled to form a complete 2-dimensional complex. The latent dynamics use parameters $c = 0.5$ and $\delta t = 0.1$, and observations follow $\mathbf{y}_i = \mathbf{H}(\mathbf{L})\,\mathbf{x}_i + \mathbf{v}_i$, with Gaussian noise of variance $\sigma_o = 1$. To evaluate change-point detection, each trial introduces a localized structural perturbation at time $t^\star$ by modifying a single uncertainty entry $\boldsymbol{\alpha}[l]$, thereby altering the diffusion behavior of cell l while all other cells remain unchanged. Detection and localization performance are reported using AUC scores over 25 Monte Carlo runs. Hyperparameters $a_1 = 10^{-4}$ and $a_2 = 10^{-4}$ are selected via grid search on the first 5,000 samples to optimize NMSE.

Fig. 1a shows the score evolution over time, where a clear and isolated peak appears at the true change point. Fig. 1b summarizes the detection accuracy: all methods operate near chance level for small perturbations, but as the perturbation grows, the proposed method displays a sharp rise in AUC and consistently outperforms both baselines. This gain stems from modeling how uncertainty propagates through the Dirac operator and the Hodge Laplacian, enabling discrimination between genuine structural changes and noise that diffuses across adjacent cells—an effect the baselines cannot capture. Recon-

Authorized licensed use limited to: TU Delft Library. Downloaded on March 30,2026 at 07:22:16 UTC from IEEE Xplore. Restrictions apply.



(a) Score variation (synthetic) (b) AUC (synthetic) (c) Score variation (Cherry Hills) (d) AUC (Cherry Hills)

Fig. 1: (a) Temporal evolution of the score function on synthetic data; the red line marks the true change point. The score exhibits an immediate sharp rise after the change, then stabilizes as the system settles into the new regime. (b) AUC variation with perturbation magnitude for the synthetic experiments. (c) Cherry Hills water-network topology (portion shown for clarity); the red pipe indicates the leak location. Its score trajectory shows a pronounced peak precisely after the change point. (d) AUC trend as leakage strength increases.

struction is evaluated under partial observability using random masks with missing-data ratios $p_m \in \{0.05, 0.10, 0.15, 0.20\}$. NMSE values in Table I are computed from the detected change point to the stabilized post-change interval. The proposed *CCP-SPDE* achieves the lowest reconstruction error across all settings by smoothing the states after the parameters restabilize, thus avoiding the accumulation of transitional errors. Note that because the synthetic data are generated under our proposed SPDE model, these results should be interpreted primarily as model-consistency validation.

B. EPANET Water Network Evaluation

We evaluate the framework on the Cherry Hills water network simulated in EPANET, consisting of 36 nodes, 40 pipes, and 5 faces [9]. At each time step, flow rates (m^3/s) are observed on the pipes. A controlled leakage is introduced at a known pipe, with its magnitude gradually increased to assess detection sensitivity. For change-point detection, the flow data are corrupted with 10% random missing entries and additive Gaussian noise of variance 0.001. All signals are standardized using statistics from the first 2,000 samples. To ensure comparability across pipes, each deviation-score trajectory is normalized by its maximum value within this baseline window, eliminating scale differences. The hyperparameters $a_1 = a_2 = 10^{-3}$, $\delta t = .1$, and $c = 0.5$ are selected via grid search on the same initial window to optimize NMSE.

Fig. 1c shows the network topology, with the leaking pipe highlighted. The score trajectories show a sharp, peak after change point at the leaking edge, indicating accurate spatial localization. This is reflected in the AUC scores, which increase with leakage strength and consistently surpass the baselines (Fig. 1d). Across the 25 Monte Carlo runs, the AUC varies by roughly ± 0.2, primarily due to random masking: when the leaking pipe or its neighbors are heavily masked, the evidence of leakage diminishes, reducing detection accuracy. For clarity, this variability is omitted from Fig. 1d. Reconstruction performance is evaluated following the same protocol as in the synthetic experiments, using progressively larger missing-data ratios $p_m \in \{0.05, 0.10, 0.15, 0.20\}$. The NMSE results in Table I show that the proposed method consistently outperforms baselines across all levels of missing data.

V. Conclusion

We introduced an online, change-point–aware framework for modeling time-series data over cellular complexes. The proposed topology-aware state-space formulation captures latent dynamics that evolve over structured domains, while jointly estimating hidden states and model parameters in real time. By tracking temporal variations in process-uncertainty parameters, the method reliably detects and localizes topological regime shifts. Experiments on both synthetic and EPANET-generated datasets demonstrate strong performance in detection, localization, and reconstruction under partial and noisy observations. Future work includes extending the framework to nonlinear state dynamics and more expressive observation models, further broadening its range of potential applications.

References

[1] E. Isufi, G. Leus, B. Beferull-Lozano, S. Barbarossa, and P. Di Lorenzo, "Topological signal processing and learning: Recent advances and future challenges," *Signal Processing*, vol. 233, p. 109930, 2025.

[2] S. Barbarossa and S. Sardellitti, "Topological signal processing over simplicial complexes," *IEEE Transactions on Signal Processing*, 2020.

[3] J. Hoppe, V. P. Grande, and M. T. Schaub, "Don't be afraid of cell complexes! an introduction from an applied perspective," 2025. [Online]. Available: https://arxiv.org/abs/2506.09726

[4] R. Money, M. Sabbaqi, J. Krishnan, B. Beferull-Lozano, and E. Isufi, "Kalman filtering for simplicial processes," in *2024 58th Asilomar Conference on Signals, Systems, and Computers*, 2024.

[5] M. Sabbaqi and E. Isufi, "Inferring time varying signals over uncertain graphs," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024.

[6] M. Sabbaqi, R. Taormina, and E. Isufi, "Gknet: Graph kalman filtering and model inference via model-based deep learning," 2025. [Online]. Available: https://arxiv.org/abs/2506.22004

[7] M. Yang, E. Isufi, M. T. Schaub, and G. Leus, "Simplicial convolutional filters," *IEEE Transactions on Signal Processing*, 2023.

[8] R. Money, J. Krishnan, B. Beferull-Lozano, and E. Isufi, "Online edge flow imputation on networks," *IEEE Signal Processing Letters*, 2023.

[9] J. Krishnan, R. Money, B. Beferull-Lozano, and E. Isufi, "Simplicial vector autoregressive models," *IEEE Transactions on Signal Processing*, 2024.

[10] L. Marinucci, C. Battiloro, and P. D. Lorenzo, "Topological adaptive least mean squares algorithms over simplicial complexes," 2025. [Online]. Available: https://arxiv.org/abs/2505.23160

[11] R. E. Kalman, "A new approach to linear filtering and prediction problems," 1960.

[12] B. Zaman, L. Ramos, D. Romero, and B. Beferull-Lozano, "Online topology identification from vector autoregressive time series," *IEEE Transactions on Signal Processing*, vol. 69, pp. 210–225, 2021.

Authorized licensed use limited to: TU Delft Library. Downloaded on March 30,2026 at 07:22:16 UTC from IEEE Xplore. Restrictions apply.